École des Hautes Études en Sciences Sociales

EHESS Doctoral School

Doctoral School no. 465
Économie Panthéon Sorbonne

Analysis and Policy in Economics

Amory Gethin

Essays on the Political Economy of Global Inequality

Supervised by: Thomas Piketty

Defense date: December 6, 2023

Referees : Daron Acemoglu, Massachusetts Institute of Technology
Emmanuel Saez, University of California, Berkeley

Jury : Julia Cagé, Sciences Po Paris
Thomas Piketty, École des hautes études en sciences sociales
Ekaterina Zhuravskaya, Paris School of Economics
Gabriel Zucman, Paris School of Economics

École des Hautes Études en Sciences Sociales

École doctorale de l'EHESS

École doctorale n°465
Économie Panthéon Sorbonne

Discipline : Analyse et Politique Économique

Amory Gethin

Essais sur l'économie politique des inégalités mondiales

Thèse dirigée par : Thomas Piketty

Date de soutenance: 6 décembre 2023

Rapporteurs : Daron Acemoglu, Massachusetts Institute of Technology
Emmanuel Saez, University of California, Berkeley

Jury : Julia Cagé, Sciences Po Paris
Thomas Piketty, École des hautes études en sciences sociales
Ekaterina Zhuravskaya, Paris School of Economics
Gabriel Zucman, Paris School of Economics

Remerciements

Cette thèse s'est nourrie du soutien des nombreuses personnes qui m'ont accompagné ces dernières années. Je tiens avant tout à remercier Thomas Piketty, qui a été un directeur de thèse exceptionnel. Sa bienveillance, sa passion et son approche unique de la recherche ont toujours été un exemple. De mes débuts au World Inequality Lab à nos travaux communs sur les clivages politiques jusqu'à mes propres travaux sur les inégalités mondiales, il n'a cessé de m'accompagner, de remodeler ma pensée et de m'aiguiller vers de nouvelles pistes toujours plus captivantes. Ma dette envers lui est immense.

Je veux ensuite remercier Daron Acemoglu et Emmanuel Saez, qui m'ont fait l'honneur d'être les rapporteurs de cette thèse. Leur vaste perspective sur les questions économiques a toujours suscité mon admiration et continue de m'inspirer au quotidien. Je tiens également à remercier les autres membres du jury, Julia Cagé, Ekaterina Zhuravskaya et Gabriel Zucman, dont les travaux me fascinent toujours plus chaque jour ; je me réjouis d'avoir pu bénéficier de leurs précieux conseils.

Cette thèse est aussi le fruit des deux années que j'ai eu la chance de passer outre-Atlantique. Je suis éternellement reconnaissant à Vincent Pons de m'avoir accueilli à la Harvard Business School. Sa vision d'ensemble, sa rigueur et sa justesse n'ont eu de cesse de m'inspirer, notamment lors de notre travail commun sur les mouvements sociaux aux États-Unis. Je veux également remercier Abhijit Banerjee, qui m'a généreusement accueilli au MIT et dont l'immense savoir m'impressionnera toujours. Notre article commun sur les clivages politiques en Inde a été l'un des propulseurs de mon engouement pour la recherche comparative. Il me faut aussi à nouveau remercier Emmanuel Saez et Gabriel Zucman, qui m'ont pris sous leur aile à Berkeley et ont su patiemment me guider avec une si grande bienveillance.

C'est au World Inequality Lab, à l'école d'économie de Paris, qu'ont émergé les idées qui constituent l'essentiel de cette thèse. Je dois beaucoup à Lucas Chancel, pour m'y avoir accueilli et soutenu, d'abord comme assistant chercheur, puis comme économiste

et enfin comme coordinateur. Son travail de coordination et sa participation à deux des articles de cette thèse ont été des apports précieux. Le World Inequality Lab est un magnifique cocon où se mêlent idées foisonnantes, passion intellectuelle et une inévitable jovialité. Merci à tous ceux avec qui j'ai pu y partager cette expérience humaine au quotidien. Beaucoup sont devenus mes co-auteurs – Thomas Blanchet, Denis Cogneau, Léo Czajka, Carmen Durrer de la Sota, Matthew Fisher-Post, Yajna Govind, Thanasak Jenmana, Marc Morgan, Alix Myczkowski et Anne-Sophie Robilliard. Tous ont participé à faire de ces années une aventure enrichissante : Facundo Alvaredo, Marie Andreescu, Oscar Barrera, Nitin Bharti, Philipp Bothe, Luis Estévez Bauluz, Richard Clarke, Mauricio de Rosa, Alice Fauvel, Ignacio Flores, Bertrand Garbinti, Diane Gabeloteau, Jonathan Goupille-Lebret, Sehyun Hong, Zhexun Mo, Rowaida Moshrif, Theresa Neef, Gastón Nievas Offidani, Olivia Ronsain, Alice Sodano, Li Yang et Álvaro Zúñiga-Cordero. Je ne saurais oublier Aroop Chatterjee, grâce à qui mon intérêt pour l'Afrique du Sud a pu s'épanouir si pleinement ; trois chapitres de cette thèse résultent de cet inoubliable séjour à Johannesburg.

Cette thèse s'est aussi nourrie d'un travail collectif sur les clivages politiques mené aux côtés de multiples compagnons de route. Je dois remercier tout particulièrement Clara Martínez-Toledano, avec qui j'ai partagé l'inattendue mais géniale épopée de coordonner notre livre collectif dont le neuvième chapitre de cette thèse est issu. Outre de nombreux membres du World Inequality Lab qui ont joué un rôle fondamental dans la constitution de cet ouvrage, mes pensées vont aux autres participants à ce vaste projet : Lydia Assouad, Jules Baleyte, Yonatan Berman, Fabian Kosse, Ana Leiva, Attila Lindner, Filip Novokmet, Sultan Mehmood, Juliet-Nil Uraz et Tom Zawisza.

Ce voyage intellectuel n'aurait pas été le même sans les multiples rencontres que j'ai pu faire. À Paris d'abord, au cours de presque huit années passées à l'école d'économie de Paris : merci Kentaro Asai, Pierre Bachas, Cécile Bonneau, Alice Chiocchetti, Sebastián García, Louise de Gaudemaris, Juan Ferrer, Julius Goedde, Clément Gras, Gustave Kenedi, Yagan Hazard, Claire Leroy, Antoine Levy, Rosanne Logeart, Mathilde Muñoz, Nicolas Marx, Julia Paul-Venturine, Madeleine Péron, Tom Raster, Yannic Rehm, Éléonore Richard, Manpreet Singh, Louis Sirugue, Morten Støstad, Eric Teschke, Guillaume Thébaudin, Basile Vidalenc et Duncan Webb.

À Cambridge ensuite, où j'ai eu la chance de rencontrer tant d'autres joyeux aspirants chercheurs issus de tous horizons : Thomas Bézy, Andrea Cornejo, Edgard Dewitte, Grace Finley, Lars Harhoff Andersen, Frederik Kjøller Larsen, Romaine Loubes,

Simon Löwe, Oda Nedregård, Justine Knebelmann, Ignas Maželis, Clément Mazet-Sonilhac, Stephen Osadetz, Leonardo Peñaloza-Pacheco, Vincent Rollet, Julian Schärer, Sarah Schneider-Strawczynski, Garima Shastra, Clara Sievert, Emily Silcock et Marie Truffier Blanc.

À Berkeley, *last but not least*, où le grand soleil de la matrice m'a convaincu que la vie californienne n'a pas son pareil : merci Léa Bignon, Eva Davoine, Jeremy Do Nascimento Miguel, Sebastian Hager, Ségal Le Guern Herry, Wouter Leenders, Lukas Lehner, Jakob Miethe, Petra Orešković, Cristóbal Otero, Jon Piqueras, Camille Semelet, Jeffrey Yusof. Tant d'autres belles rencontres que je ne pourrais toutes lister ici ont ponctué cette aventure.

Merci, enfin, à mes amis et à ma famille. Antonio, Florent, Gabriel, Jack, Khenpo La, pour la joie et la folle énergie que vous continuez de sans cesse m'apporter. Mes parents, pour votre amour inconditionnel, votre soutien sans faille et surtout vos extraordinaires qualités humaines auxquelles je ne cesserai d'aspirer.

Résumé

Cette thèse consiste en un ensemble de neuf essais portant sur divers thématiques liées à la pauvreté, aux inégalités et aux clivages politiques.

Le premier chapitre étudie le rôle de l'éducation dans la réduction de la pauvreté mondiale et des inégalités de genre depuis 1980. À partir d'enquêtes couvrant 95 % de la population mondiale, de séries historiques sur les inégalités de revenus et d'un modèle stylisé reliant éducation et structure des salaires, nous y estimons à quoi ressemblerait la distribution des revenus mondiaux si l'éducation revenait à son niveau observé en 1980. Sur la base d'hypothèses conservatrices, l'expansion éducative explique environ la moitié de la croissance économique mondiale, deux-tiers de la croissance des revenus des 20 % les plus pauvres du monde et plus de la moitié de l'amélioration des inégalités hommes-femmes. Les politiques éducatives apparaissent ainsi au cœur du remarquable déclin de la pauvreté et des inégalités de genre observé au cours des dernières décennies.

Le deuxième chapitre porte sur la construction de mesures de la pauvreté et des inégalités mondiales incorporant la consommation de biens publics. L'approche traditionnelle de mesure des inégalités se focalise sur le revenu disponible ou la consommation, excluant par là même l'ensemble des transferts perçus sous forme d'éducation, de santé et d'autres services publics. À partir de diverses sources de données, nous construisons une nouvelle base de données historique sur la valeur et la progressivité des transferts publics en nature reçus dans le monde depuis 1980. La consommation croissante de biens publics explique environ 20 % de la réduction de la pauvreté mondiale. La redistribution gouvernementale dans son ensemble, incluant à la fois transferts monétaires et transferts en nature, en explique près de 30 %. La prise en compte des biens publics dans la mesure du bien-être économique a ainsi d'importantes conséquences sur l'évaluation de la pauvreté mondiale, des inégalités et des différences de niveaux de vie entre pays.

Le troisième chapitre examine l'évolution de la redistribution gouvernementale dans le

monde depuis 1980. Combinant de multiples sources de données, nous y construisons de nouvelles mesures de la progressivité des impôts payés et transferts perçus par groupe de revenu dans 151 pays. Trois faits stylisés ressortent de notre analyse. Tout d'abord, les transferts publics expliquent l'essentiel des différences de redistribution entre pays. En effet, les impôts n'ont généralement qu'un effet mineur sur les inégalités : dans la plupart des pays du monde, le système fiscal n'est que rarement progressif et apparaît même souvent régressif. Deuxièmement, la redistribution a augmenté dans la plupart des régions du monde, à l'exception de l'Europe de l'Est et de l'Afrique, où elle n'a que peu évolué depuis 1980. Malgré de considérables différences entre systèmes socio-fiscaux à travers le monde, enfin, la redistribution apparaît n'expliquer qu'environ 20 % des différences de niveaux d'inégalités entre pays, tandis que la "prédistribution" (les inégalités avant impôts et transferts) rend compte de 80 % de ces variations. Les pays les plus égalitaires tendent à redistribuer davantage, cependant, ce qui suggère que les politiques redistributives pourraient jouer un rôle indirect important dans la formation des inégalités primaires.

Dans le quatrième chapitre, nous associons données d'enquêtes, données fiscales et comptes nationaux pour constituer une nouvelle base de données sur les inégalités avant et après impôts et transferts dans vingt-six pays européens depuis 1980. Notre approche se base sur la méthode des comptes nationaux distribués, récemment développée et appliquée au cas des États-Unis, ce qui nous permet de comparer de manière systématique l'évolution des inégalités dans ces deux régions du monde. Les inégalités se sont accrues dans la plupart des pays européens au cours de cette période, en particulier au sommet de la pyramide des revenus, mais bien moins qu'aux États-Unis. Nous montrons également que les inégalités sont aujourd'hui plus élevées aux États-Unis non pas du fait de différences de progressivité des impôts et transferts (ou "redistribution"), mais avant tout du fait d'écarts importants en termes d'inégalités avant impôts et transferts (ou "prédistribution") entre les deux régions.

Le cinquième chapitre étudie comment la prise en compte de la consommation des biens publics affecte la mesure de la pauvreté et des inégalités dans le cas spécifique de l'Afrique du Sud post-apartheid. À partir de données budgétaires, d'enquêtes et de recensements, nous estimons la répartition de tous les transferts publics reçus par groupe de revenu entre 1993 et 2019, dont l'éducation, la santé, les services de police, les infrastructures de transport, les subventions au logement ou encore les services publics locaux. Notre analyse révèle une progression considérable des transferts en nature perçus par les ménages les plus pauvres : après comptabilisation de la consommation des biens publics, le taux de croissance du revenu des 50 %

les plus modestes passe d'environ 65 % à 90 %. Les services publics apparaissent ainsi comme un des leviers fondamentaux de la réduction de la pauvreté et d'une croissance plus inclusive depuis la fin de l'apartheid.

Le sixième chapitre analyse l'évolution des inégalités en Afrique du Sud depuis 1993 et le rôle joué par la redistribution gouvernementale dans la réduction de celles-ci. Combinant données d'enquêtes, données fiscales et comptes nationaux, nous documentons une forte croissance des écarts de revenus avant impôts et transferts. Cette croissance a cependant été plus que compensée par une expansion majeure de l'état social sud-africain, conduisant à une réduction des inégalités après impôts et transferts au cours de cette période. L'Afrique du Sud continue néanmoins de figurer parmi les pays les plus inégalitaires au monde, l'appartenance raciale jouant un rôle persistant dans la structuration de ces extrêmes disparités.

Le septième chapitre se tourne vers l'évolution des inégalités de patrimoine en Afrique du Sud. En associant des micro-données couvrant l'ensemble des déclarations d'impôts sur le revenu à des données macroéconomiques sur l'évolution du patrimoine agrégé, nous construisons de nouvelles séries historiques sur la distribution des actifs et passifs des ménages depuis la fin de l'apartheid. Ces séries révèlent des niveaux d'inégalités extrêmes : les 10 % de sud-africains les plus aisés détiennent 86 % du patrimoine total, tandis que le patrimoine moyen des 50 % les moins aisés est négatif (leurs dettes excèdent la valeur de leurs actifs). Rien n'indique que la concentration du patrimoine ait baissé de manière significative depuis 1993.

Le huitième chapitre examine l'évolution des inégalités de revenus en Afrique depuis 1990. En l'absence de données satisfaisantes dans la plupart des pays du continent, nous corrigeons les indicateurs d'inégalités de consommation disponibles à partir d'études traitant de la relation entre revenu et consommation, ainsi que de la sous-estimation des hauts revenus dans les enquêtes réalisées auprès des ménages. Nos résultats suggèrent que les inégalités en Afrique sont élevées et sont restées relativement stables depuis 1990. Colonisation de peuplement et diffusion de l'islam ressortent parmi les deux corrélats les plus significatifs rendant compte des différences de niveaux d'inégalités entre pays.

Le neuvième chapitre porte sur l'évolution de long terme des clivages politiques dans vingt-et-une démocraties occidentales. À partir d'une nouvelle base de données couvrant les déterminants du vote au cours de plus de 300 élections organisées entre 1948 et 2020, nous documentons une divergence complète des effets du revenu et du diplôme sur les comportements électoraux. Dans les années 1950, les partis sociaux-démocrates et affiliés obtenaient de meilleurs scores parmi les électeurs les moins aisés et les moins diplômés. Si l'association entre revenu et vote est restée remarquablement stable depuis lors, le clivage éducatif s'est quant à lui complètement renversé : en 2020, les électeurs les plus diplômés étaient devenus beaucoup plus enclins à voter pour les partis de "gauche" dans la plupart des démocraties occidentales. L'analyse de données portant sur les programmes des partis politiques suggère que cette transformation a été étroitement associée à la prise d'importance croissante d'une nouvelle dimension "socioculturelle" du conflit politique.

Abstract

This thesis is a collection of nine essays covering topics related to poverty, inequality, and political divides.

The first chapter studies the role played by education in the historical reduction of global poverty and gender inequality since 1980. Combining survey microdata covering 95% of the world's population, historical inequality statistics, and a simple model of education and the wage structure, I estimate how the world distribution of income would look like if educational attainment was to come back to its 1980 level in each country. Under conservative assumptions, education accounts for about half of global economic growth, two-thirds of income gains for the world's poorest 20% individuals, and over half of improvements in the share of labor income accruing to women. This puts education policies at the center of the remarkable reduction of poverty and gender inequality observed in the past decades.

In the second chapter, I construct measures of global poverty and inequality that incorporate the consumption of public goods. Traditional income distribution statistics focus on household disposable income or consumption, entirely excluding government transfers received by individuals in the form of education, healthcare, and other public services. Combining various data sources, I build a novel historical database on the value and progressivity of all government in-kind transfers received worldwide since 1980. I find that the consumption of public goods accounts for about 20% of global poverty reduction. Total government redistribution, including cash and in-kind transfers, accounts for 30%. Incorporating public goods in measures of economic welfare has important implications for the measurement of global poverty, inequality, and cross-country differences in living standards.

The third chapter sheds new light on the evolution of government redistribution worldwide since 1980. Combining various data sources, we build new measures of the progressivity of taxes and transfers in 151 countries. We establish three main facts. First, transfers explain the bulk of cross-country differences in redistribution. Taxes

have little effect on inequality in most countries in the world: most tax systems are either flat or regressive. Second, redistribution has increased in most world regions, but not in Eastern Europe and Africa, where it has virtually stagnated since 1980. Third, despite large differences in tax-and-transfer systems, redistribution accounts for only 20% of cross-country differences in inequality; "predistribution" (pretax income inequality) explains 80%. Countries with higher redistribution display lower levels of pretax inequality, however, pointing to a potentially large role of redistributive policies in indirectly shaping the distribution of market incomes.

In the fourth chapter, we combine survey, tax, and national accounts data to build new estimates of pretax and posttax income inequality covering twenty-six European countries since 1980. Our series are based upon the Distributional National Accounts framework, recently developed to track inequality in the United States, which allows us to systematically compare the evolution of income disparities in the two regions. We find that inequality increased in most European countries during this period, especially at the top of the income distribution, but much less than in the United States. Most importantly, we show that inequality is higher in the United States not because of differences in the progressivity of taxes and transfers ("redistribution"), but primarily because of differences in pretax inequality ("predistribution").

The fifth chapter investigates the implications of incorporating measures of the consumption of public goods into poverty and inequality statistics, focusing on the case of post-apartheid South Africa. Combining budget data with census and survey microdata, I estimate the distribution of all government transfers received by income group from 1993 to 2019, including education, healthcare, police services, transport infrastructure, housing subsidies, and local government services. I find that there have been considerable increases in in-kind transfers received by low-income households: accounting for the consumption of public goods raises the real income growth rate of the poorest 50% from about 65% to 90%. This puts public services as a fundamental driver of inclusive growth and poverty reduction since the end of apartheid.

The sixth chapter focuses on the evolution of inequalities in South Africa since 1993 and the role played by government redistribution in mitigating them. Combining survey, tax, and national accounts data, we find that there has been a large increase in pretax inequality, but that this rise has been overcompensated by major expansions in government redistribution. However, South Africa still stands out as one of the most unequal countries in the world, with race playing a persistent role in structuring these extreme disparities.

The seventh chapter centers on the evolution of wealth inequality in South Africa. With microdata covering the universe of income tax returns, surveys, and macroeconomic balance sheets, we construct new estimates of the distribution of household assets and liabilities since the end of apartheid. We document extreme levels of wealth concentration: the 10% richest South Africans own 86% of aggregate wealth, while the average wealth of the poorest 50% is negative (their debts exceed their assets). We find no evidence of a decrease in wealth inequality since 1993.

The eighth chapter studies the evolution of income inequality in Africa since 1990. In the absence of high-quality data in most countries, we correct available consumption distributions using information from a selection of studies covering the relationship between income and consumption and the typical underestimation of top-income earners in household surveys. We find that inequality in Africa is high and has remained relatively stable since 1990. Historical settler colonialism and the spread of Islam stand out as the strongest correlates of cross-country differences in inequality.

The ninth chapter provides new evidence on the long-run evolution of political cleavages in twenty-one Western democracies. We assemble a new survey database covering the determinants of the vote in over 300 elections from 1948 to 2020. We document a striking long-run divergence of the effects of income and education. In the 1950s, social democratic and affiliated parties were more popular among low-income and lower-educated voters. While the income gradient has remained remarkably stable, there has been a complete reversal of the educational divide: by 2020, higher-educated voters have become much more likely to vote for the "left" in most Western democracies. Drawing on manifesto data covering political platforms, we find that this transformation is tightly linked to the growing salience of a new "sociocultural" dimension of political conflict.

Table of Contents

General Introduction

The past decades witnessed deep economic and political transformations in the world economy. Access to essential public services among the world's poorest individuals improved dramatically, from basic education to healthcare, drinkable water, and electricity. The rise of China, India, and other developing economies reshuffled the global organization of production. Sustained economic growth was accompanied by rising inequality in many parts of the world, putting into question the benefits of globalization. The imminent threat of climate change brought to light the unsustainability of prevailing economic systems, generating new ideological movements but also new social conflicts. At the same time, the spread of democracy was met by the revival of authoritarian movements expanding in countries as diverse as France, the United States, India, and Brazil.

This PhD thesis is motivated by two principles. First, understanding these profound mutations requires analyses that extend beyond the present day and beyond the borders of narrow territories. To dissect the mechanisms underlying the fall of extreme poverty, the rise of inequality, or the transformation of political conflict, comparative and historical perspectives are key. In this spirit, this thesis gathers a collection of cross-country analyses and more narrow case studies that are to be explicitly connected and combined. Careful comparisons can teach us so much about the making of contemporary societies. With attention to commonalities and differences, they can illuminate new prospects in unexpected directions.

Second, answering complex questions requires interpretations that are anchored in well-identified facts. As such, this thesis presents itself as an advocate of the power of description. It contains few new theoretical insights. Instead, it hopes to modestly provide empirical clarity on questions such as: what do we mean by "the rise of inequality" or "the reduction of global poverty?" By how much do taxes and transfers reduce poverty and inequality today compared to several decades ago? What are the long-run transformations hiding behind the "decline of class divides" and development of "illiberal movements?" Addressing these challenging questions inevitably requires turning to history. The contribution of this collection of essays is to put together data and statistical analyses covering specific aspects of this history.

Unpacking the Dynamics of Global Poverty

The first part of this thesis, consisting of two closely related chapters, centers on the factors driving the historical evolution of the world distribution of income. In chapter 1, I investigate the role played by education in the reduction of global poverty and gender inequality since 1980. Despite considerable improvements in access to schooling in the past decades, how useful these investments have been at reducing poverty remains a topic of considerable debate. Combining a new microdatabase covering education and earnings for 95% of the world's population with measures of the returns to schooling, historical income distribution statistics, and a simple model of education and the wage structure, I estimate the contribution of education to economic growth and its distribution worldwide. Under conservative assumptions, education explains about half of average economic growth, two-thirds of income gains for the world's poorest 20% individuals, and over half of improvements in the share of labor income accruing to women since 1980. Given the dominant role that governments have had at providing education, this puts education policy at the center of the remarkable reduction of poverty and gender inequality observed in the past decades.

Public services do not only contribute to pretax income growth, however. They also have a direct effect on poverty by allowing households to save money on services they would otherwise have to pay for. Drawing on this intuition, I construct in chapter 2 measures of global poverty that incorporate the consumption of public goods. Combining historical budget data, household surveys, and fiscal incidence studies, I estimate the distributional incidence of spending on education, healthcare, and other public services received worldwide since 1980. I also incorporate cash transfers received and taxes paid, providing a unique perspective on the role played by government redistribution in shaping the world distribution of income. I find that the consumption of public services accounts for about 20% of global poverty reduction since 1980. Total government redistribution, including cash and in-kind transfers, accounts for 30%.

Combining these two perspectives, I provide in chapter 1 a minimalist estimate of the total contribution of public policies to worldwide poverty reduction. At any given point in time, poor households benefit from both greater pretax incomes through education received during their childhood, and net government transfers increasing their posttax incomes today. Combining direct redistribution and indirect investment benefits from education brings the total contribution of public policies to global poverty reduction to at least 50%. This should probably be seen as a lower bound, given that other public services, such as healthcare or transport infrastructure, arguably contributed to pretax income growth during this period too.

The Power and Limits of Government Redistribution

Beyond this global perspective, government redistribution has had varying success at reducing poverty and inequality in different countries. Chapters 3 to 8 investigate these successes and failures in more detail, adopting various angles to zoom on specific contexts.

Chapter 3, written with Matthew Fisher-Post, complements chapter 2 by providing a more granular examination of cross-country differences in redistribution and their evolution in the past decades. Combining surveys, tax simulators, budget data, and estimates from fiscal incidence studies, we construct a new database covering the distributional incidence of taxes and transfers in nearly all countries in the world since 1980. We document large variations in tax-and-transfer systems across countries and significant progress in the inequality-reducing effects of taxes and transfers in most world regions. Transfers are by far the most important component of redistribution across countries and over time; progressive tax systems are rare. Overall, however, taxes and transfers do not appear to be the main driver of cross-country differences in inequality: about 80% of variations in inequality after taxes and transfers can already be accounted for by differences in inequality before taxes and transfers. "Predistribution," not "redistribution," stands out as the dominant driver of national differences in the distribution of incomes.

Chapter 4, written with Thomas Blanchet and Lucas Chancel, turns to a more precise analysis of predistribution and redistribution comparing the sources of inequality in Europe and the United States. Following the Distributional National Accounts methodology recently developed to track inequality in the United States, we combine surveys, tax, and national accounts data to produce pretax and posttax income inequality series covering twenty-six European countries from 1980 to 2017. We find that pretax inequality has risen in Europe, but this increase has been much less pronounced than in the United States and has been restricted to the very top of the income distribution. Contrary to a widespread view, we also demonstrate that Europe's lower inequality levels cannot be explained by more equalizing tax-and-transfer systems. Across several measures, the U.S. even turns out to redistribute more income to low-income households than any European country. This puts predistribution policies, such as minimum wages, industrial regulation, corporate governance, or education at the center of the large inequality gap that increasingly separates the two regions.

Chapters 5 through 7 present themselves as an even more detailed inquiry into the sources and nature of economic inequality in post-apartheid South Africa. This case study turns out to be particularly interesting for understanding the power and limits of redistribution. On the one hand, South Africa went through a remarkable transition in the 1990s, ending centuries of extreme racial oppression and paving the way to a new social state. On the other hand, this new era has not been without challenges, from low economic growth to the persistence of chronic unemployment. I investigate in chapter 5 trends in inequality in access to public services since 1993. In chapters 6 and 7, written with Aroop Chatterjee and Léo Czajka, we develop a more comprehensive analysis of the evolution of income and wealth inequalities. These three chapters rely on a unique combination of surveys, tax data, national accounts data, and newly digitized budget reports, which allow us to adopt a comprehensive view of the different facets of South African inequality. The main finding is that substantial increases in government redistribution have over-compensated the significant rise of pretax income inequalities observed since the end of apartheid. Public services have played a particularly important role in that transformation. In that sense, South Africa somewhat represents an interesting exception to the superiority of predistribution, with cash and in-kind transfers standing out as the dominant drivers of poverty and inequality reduction. That being said, the country remains one of the most unequal in the world, even after redistribution, with the racial dimension persistently structuring these disparities. There is also no evidence that wealth inequality has decreased since the end of apartheid. In 2017, the richest 10% owned over 85% of household wealth.

Chapter 8, written with Lucas Chancel, Denis Cogneau, Alix Myczkowski, and Anne-Sophie Robilliard, closes the second part of this thesis with an analysis of income inequality in Africa. It is sometimes thought that inequality is relatively low in Africa in comparison to other world regions. We show that this finding largely results from inequality being measured in terms of consumption, which is mechanically lower than in terms of income, as well as from substantial underreporting of top earners in household surveys. With an approximate correction for these two sources of bias, we find Africa to be one of the most unequal regions in the world, with little change since 1990. In line with chapters 1 and 2, our findings highlight the importance of within-country inequality in explaining levels and trends in extreme deprivation, even in the poorest countries.

Representing Inequalities: The Complexities of Modern Multidimensional Political Conflict

The third part of this thesis expands the scope of the analysis to the sphere of politics, with a focus on the democratic representation of inequalities. Chapter 9, written with Clara Martínez-Toledano and Thomas Piketty, studies the long-run evolution of the structure of political cleavages in twenty-one Western democracies. Combining electoral surveys covering over 300 elections held from 1948 to 2020, we construct a new database covering the socioeconomic determinants of the vote. We document a striking divergence between the effects of education and income on voting behaviors. In the early postwar decades, both low-income and lower-educated voters were much more likely to support social democratic and affiliated parties than high-income and higher-educated voters. While income-related divides have remained remarkably constant, there has been a complete reversal of the education cleavage, generating what might be called "multi-elite party systems:" economic elites still vote for the "right," while educated elites now decisively support the "left." The emergence of new green and anti-immigration parties has accelerated this transition, although it cannot account for it alone. Linking our microdatabase to manifesto data covering political parties' programs, we provide evidence that the reversal of educational divides is tightly linked to the growing salience of a new "sociocultural" dimension of political conflict.

These striking facts open many avenues for analyses of inequalities and cleavage structures that go beyond advanced Western democracies. Under which conditions do social inequalities become politicized? Who are the core supporters of growing environmental and authoritarian movements in old and new democracies? How does the relative salience of economic and sociocultural conflicts vary across contemporary democracies and why? Some of these questions are explored in the nineteen chapters of our collective book, *Political Cleavages and Social Inequalities. A Study of Fifty Democracies, 1948–2020*, co-edited with Thomas Piketty and Clara Martínez-Toledano (Harvard University Press, 2021). The author of the present thesis had the fortune of writing or co-writing thirteen of these chapters covering trends in political divides in various regions of the world. In a comparative chapter, "Political Cleavages and Social Inequalities in 50 Democracies, 1948-2020," co-written with Thomas Piketty and Clara Martínez-Toledano, we start by providing a global map of electoral divides and their relationship to socioeconomic inequalities.

Three chapters are then dedicated to covering the specific trajectories of Western democracies: “Political Cleavages, Class Structures, and the Politics of Old and New Minorities in Australia, Canada, and New Zealand,” “Historical Political Cleavages and Post-Crisis Transformations in Italy, Spain, Portugal and Ireland, 1958-2020” (with Luis Bauluz, Clara Martínez-Toledano, and Marc Morgan), and “Party System Transformation and the Structure of Political Cleavages in Austria, Belgium, the Netherlands, and Switzerland” (with Carmen Durrer de la Sota and Clara Martínez-Toledano).

Five additional chapters turn to case studies of Asian democracies: “Caste, Class, and the Changing Political Representation of Social Inequalities in India, 1962-2019” (with Abhijit Banerjee and Thomas Piketty), “Social Inequality and the Dynamics of Political and Ethnolinguistic Divides in Pakistan, 1970-2018” (with Sultan Mehmood and Thomas Piketty), “Political Cleavages and the Representation of Social Inequalities in Japan, 1953-2017,” “Democratization and the Construction of Class Cleavages in Thailand, the Philippines, Malaysia, and Indonesia, 1992-2019” (with Thanasak Jenmana), and “Inequality, Identity, and the Structure of Political Cleavages in South Korea, Taiwan, and Hong Kong, 1996-2016” (with Carmen Durrer de la Sota).

Finally, four chapters analyze voting behaviors in Brazil (“Democracy and the Politicization of Inequality in Brazil, 1989-2018”, with Marc Morgan) and eight African and Middle Eastern countries: “Extreme Inequality, Elite Transformation, and the Changing Structure of Political Cleavages in South Africa, 1994-2019,” “Social Inequalities and the Politicization of Ethnic Cleavages in Botswana, Ghana, Nigeria, and Senegal, 1999-2019” (with Jules Baleyte, Yajna Govind, and Thomas Piketty), and “Political Cleavages and Social Inequalities in Algeria, Iraq, and Turkey, 1990-2019” (with Lydia Assouad, Thomas Piketty, and Juliet-Nil Uraz).

Although not included in the present document, this material is closely related to this thesis and represents a direct extension of the analysis developed in chapter 9.

Chapter 1

Distributional Growth Accounting: Education and the Reduction of Global Poverty, 1980-2019

The past decades witnessed dramatic improvements in access to basic public services among the global poor. These improvements were reflected in major progress made on indicators as diverse as school enrollment, literacy, vaccination rates, and access to drinkable water (United Nations, 2023). Valued as a transfer received by households, the direct consumption of public goods can account for about 20% of global poverty reduction since 1980. Total government redistribution, including cash and in-kind transfers, can account for 30% (Gethin, 2023b).

How useful these policies have been at generating pretax income growth remains, however, an open question. Education, in particular, has expanded massively in the past decades, yet its contribution to global poverty reduction remains unclear. The quality of education is low in many developing countries, raising doubts on the ability of schooling to generate productivity gains. Many competing factors could also have played a more important role than schooling, from globalization and technical progress to other government policies. The economic effects of education will ultimately depend on who benefits from educational expansion, the associated returns to schooling, and general equilibrium effects. Because of difficulties in quantifying these different channels, we lack estimates of how large benefits from schooling have been for the global poor. This question is of fundamental importance for policy, in a world where the vast majority of children from low-income households are enrolled in public schools.

This article makes a first attempt at estimating the aggregate and distributional effects of worldwide educational expansion since 1980. The starting point is a new microdatabase representative of 95% of the world's population, which I assemble from various repositories of household surveys and country-specific sources. The data cover individual labor income, education, age, gender, and other socioeconomic variables for a sample of 10 million individuals in 150 countries in 2019, providing a unique snapshot on the contemporary structure of global poverty and inequality.

Starting from this new dataset, I estimate what the world distribution of income would look like if there had been no improvement in schooling since 1980. I then compare the resulting counterfactual to the actual evolution of pretax incomes, yielding an estimate of the contribution of education to growth for different groups within each country. To construct this counterfactual, I combine standard growth accounting tools with a simple model of education and the wage structure *à la* Goldin and Katz (2007). In this "distributional growth accounting" framework, expanding education increases aggregate labor income by the private return to schooling. Educational expansion also has distributional effects by pushing down the relative wage of skilled workers as their relative supply increases. The magnitude of these effects is governed by a long-run elasticity of substitution between skill groups, which I calibrate from the recent macroeconomics literature.

I bring this framework to the data by constructing a counterfactual world distribution of income in three steps. First, I downgrade education levels in each survey until matching the distribution of educational attainment that prevailed in 1980. Second, I reduce individual incomes accordingly, using new country-specific measures of returns to primary, secondary, and tertiary education estimated from the microdata. I also exploit causal estimates of the returns to schooling from a collection of fifteen papers as a validation exercise. Finally, I estimate general equilibrium effects: the supply of skilled workers would be lower, and hence their relative wage higher, if education had not improved. This approach is analogous to the canonical growth accounting exercise, which typically combines cross-country data on average years of schooling with a uniform 10% return to derive the same counterfactual (e.g., Barro and Lee, 2015). The main contribution of this paper lies in the use of rich survey microdata representative of nearly all of the world's population with a model embedding imperfect substitution between skilled and unskilled workers. Together, these two ingredients allow for a much more granular estimation of how the economic benefits of education vary within and across countries.

I validate this methodology with new quasi-experimental evidence from three large-scale schooling initiatives in India, Indonesia, and the United States. Combining data on the distribution of income with differential exposure to each program across subnational regions, I document two main facts. First, educational expansion induced by these policies had large causal effects on aggregate regional incomes, comparable to individual returns found in the same contexts. Second, the three policies disproportionately benefited low-income earners, generating large reductions in inequality. The distributional growth accounting framework reproduces these two findings with a remarkable degree of accuracy—if anything, it slightly underestimates economic benefits of schooling—, suggesting that it provides a good methodological foundation to study the role of education in shaping pretax income growth.

In my benchmark specification, I find that private returns to schooling can account for about 50% of global economic growth and 70% of income gains for the world's poorest 20% individuals since 1980 (see Figure 1.1). They also explain about 40% of the reduction in the share of the world's population living in extreme poverty. Given the predominant role that governments have had at providing education and other basic services to low-income households, this puts public policies at the center of the historical fall of global poverty. Combining measures of direct government redistribution from a companion paper (Gethin, 2023b) with indirect investment benefits from education estimated in this paper brings the total contribution of public policies to global extreme poverty reduction to at least 50%.

These estimates should be considered as conservative. The estimation relies on standard Mincerian returns to schooling, which are typically lower than causal estimates derived from natural experiments. It is based on a relatively high elasticity of substitution between skill groups, limiting redistributive effects generated by the growing relative supply of skilled workers. It assumes that education only affects labor income, ignoring potential effects on capital income and productivity through savings, innovation, or other channels (Gennaioli et al., 2013; Queiró, 2022). It also ignores human capital externalities, on which there is now significant empirical evidence.[1] All in all, my findings are governed by two main sets of parameters: the private returns to schooling and the degree of imperfect substitutability between skill groups. With plausible values for these parameters, I bound the contribution of education to the world's poorest 20% growth between 60% and 90%. Moving below 60% would require assuming either that workers are perfect substitutes, or that

[1]Existing studies have generally found strong indications of externalities from higher education. Evidence on other levels of schooling is more debated. See in particular Acemoglu and Angrist (2000), Chauvin et al. (2017), Ciccone and Peri (2006), Glaeser and Lu (2018), Guo, Roys, and Seshadri (2018), Iranzo and Peri (2009), Moretti (2004), and Wantchekon, Klašnja, and Novta (2015).

returns to schooling are significantly below those found in the data, in contradiction with much of the labor economics literature and my own analysis of the three natural experiments mentioned above.

Methodologically, accounting for distributional effects within countries appears to be crucial to adequately measure the role of schooling in global poverty reduction. A standard growth decomposition relying on cross-country data, as in Barro and Lee (2015), would underestimate the contribution of education to poverty reduction by a factor of three.[2] One reason is that cross-country data cannot accurately measure poverty: the poorest individuals in the world do not all live in the poorest countries. More importantly, the classic approach fails to account for important dimensions of educational expansion, such as labor income shares being greater at the bottom of the distribution and general equilibrium effects redistributing schooling gains from high-skilled to low-skilled workers. That being said, relying on microdata instead of aggregate data also affects some findings in the opposite direction. Moving from a constant return of 10%, as often assumed in the literature, to heterogeneous returns by level reduces the contribution of education to global poverty reduction by about 20%. The main reason is that the return to basic education is particularly low in developing countries, ranging from just 3% per year in India to 6% in Sub-Saharan Africa, potentially reflecting low education quality.

One should also stress that the large contribution of schooling to global poverty reduction does not preclude that other factors, such as physical capital or technology, may have played a significant role too. In the model, for instance, the return to schooling increases with the skill bias of technology. Had technology not improved, the return to schooling would likely be substantially lower than the one observed today. Skill-biased technical change has thus potentially played a key role in *amplifying* the economic benefits of education. Quantifying this interaction effect between schooling and technology would necessitate survey data going back to the 1980s for all countries, which are unfortunately not available. For a subset of countries with historical survey data, I provide suggestive evidence that skill-biased technical change has had such positive effects, typically enhancing the returns to schooling by 20-30%.

Finally, I extend distributional growth accounting to the study of another major historical transformation: the decline of global gender inequality. To do so, I quantify

[2] The standard approach also underestimates the contribution of education to aggregate growth by about 40%, for two main reasons. First, it relies on labor income shares that exclude mixed income entirely, while I provide evidence that mixed income is affected by schooling just as much as wages. Second, I account for imperfect substitution, which increases the contribution of schooling by magnifying losses from not expanding education.

how large gender labor income gaps would be today if education had not improved. The estimation accounts for differential educational expansion by gender, but also for heterogeneous effects of schooling on earnings and labor force participation. This counterfactual is then compared to the actual evolution of female labor income shares, on which data is available since 1991. The main conclusion is that education can explain a large share of reductions in gender inequality observed in the past decades, typically 50% to 80% depending on the specification and world region considered. Education has thus played a key role in the historical empowerment of women observed in most parts of the world.

A large literature in labor economics uses the canonical labor supply-and-demand framework to relate changes in the wage distribution to educational expansion.[3] Concurrently, a considerable literature in macroeconomics investigates the contribution of human capital to development and economic growth.[4] These two methodological perspectives, one focused on within-country inequality and the other on cross-country dynamics, have remained relatively independent from one another. The main contribution of this article is to bring them together into a unified "distributional growth accounting" framework, which I use to quantify the role of education in the reduction of global poverty and gender inequality.

This article also contributes to our understanding of the forces shaping the long-run evolution of the world distribution of income. Global inequalities have undergone profound transformations in recent decades, including rapidly declining poverty and cross-country income convergence (Chen and Ravallion, 2010; Hammar and Waldenström, 2020; Pinkovskiy and Sala-i-Martin, 2016; Sala-i-Martin, 2006), the emergence of a new "global median class" (Lakner and Milanovic, 2016), skyrocketing top income inequality (Chancel and Piketty, 2021), and moderately decreasing gender inequality (Neef and Robilliard, 2021). Amongst the numerous factors shaping these

[3] This framework has been used extensively to account for trends in wage inequality in the United States (Acemoglu and Autor, 2011; Autor, Goldin, and Katz, 2020; Deming, 2023; Goldin and Katz, 2007, 2008; Hershbein, Kearney, and Pardue, 2020; Katz and Murphy, 1992). A growing literature successfully extends this analysis to low- and middle-income countries: see for instance Fernández and Messina (2018), Khanna (2023), and Vu and Vu-Thanh (2022). A few studies also investigate the role of education in explaining cross-country differences in gender inequality (e.g., Kleven and Landais, 2017).

[4] The recent literature has focused more heavily on explaining cross-country differences in development: see for instance Caselli and Coleman (2006), Gennaioli et al. (2013), Hall and Jones (1999), Hanushek and Kimko (2000), Hendricks and Schoellman (2018), Hsieh and Klenow (2010), Jones (2014), and Rossi (2020). Growth accounting decompositions have also been used extensively since Solow (1957), although worldwide perspectives are more recent (e.g., Barro and Lee, 2015; Mankiw, Romer, and Weil, 1992). Most closely related to this paper is recent work by Collin and Weil (2020), who estimate that accelerating human capital accumulation could have significant effects on global poverty reduction in coming decades.

dynamics, I isolate the contribution of one of them: education. I find that it has been a powerful source of convergence and can account for a large share of real income gains for both women and the world's poorest individuals.

Finally, this paper relates to the large empirical evidence on the economic effects of education. A vast literature documents positive impacts of education on individual earnings (Card, 2001; Deming, 2022). A more limited number of studies examine the general equilibrium effects of education policies (e.g., Che and Zhang, 2018; Duflo, 2004; Khanna, 2023). I contribute to this literature by extending previous work on India (Khanna, 2023), Indonesia (Duflo, 2001), and the United States (Acemoglu and Angrist, 2000), moving beyond individual returns to quantify the causal effect of education on macroeconomic growth and inequality. I also draw heavily on existing studies to calibrate the parameters guiding my results, such as elasticities of substitution between skill groups and differential economic effects of education by gender. In doing so, this article is an attempt at taking the best of microeconomic evidence to draw conclusions on the aggregate effects of education on global poverty and gender inequality. This approach speaks to the growing literature highlighting the need to bridge the micro-macro gap in the study of economic development (Buera, Kaboski, and Townsend, 2023).

The rest of the paper is organized as follows. Section 1.1 outlines the conceptual framework used for distributional growth accounting. Section 1.2 presents the data and methodology. Section 1.3 describes the main results on the role of education in shaping the distribution of global economic growth. Section 1.4 turns to the study of global gender inequality. Section 1.5 provides a general discussion and additional results. Section 1.6 concludes.

1.1 Distributional Growth Accounting

This section presents the framework used to estimate the aggregate and distributional effects of human capital accumulation. Section 1.1.1 formulates the problem of interest. Sections 1.1.2 and 1.1.3 expose simple formulas relating the distribution of educational attainment to aggregate earnings and inequality. Section 1.1.4 outlines the methodology used to estimate the contribution of educational expansion to the distribution of economic growth.

1.1.1 Setup

1.1.1.1 Research Question

Output is produced by combining physical capital K and workers with different levels of educational attainment:

$$Y = F(K, L) = F(K, L_0, L_1, ..., L_m) \tag{1.1}$$

Throughout the paper, I define workers as including both wage earners and the self-employed. Labor income includes both compensation of employees and mixed income, referred to as "wages" for simplicity. Capital income includes all remaining national income components.

Workers of skill s are paid their marginal product w_s. Skill groups may be imperfectly substitutable in production, implying that wages depend on relative supplies:

$$w_s = w_s(L) \tag{1.2}$$

Consider a group p, receiving income from both labor and capital, and composed of individuals with different levels of skills. This can correspond to a group of the income distribution, such as the poorest 20%, or to a social group such as women. The average income of group p is the sum of their capital income and labor income:

$$y^p = y_K^p + \sum_s w_s(L)\ L_s^p \tag{1.3}$$

With L_s^p the share of workers of type s in group p. In the benchmark specification, I make the conservative assumption that capital income is not affected by schooling, as in standard growth accounting (Barro and Lee, 2015). We are interested in estimating a counterfactual income $\tilde{y}^p$, given a counterfactual distribution of educational attainment $\tilde{L} = (\tilde{L}_1, ..., \tilde{L}_m)$:

$$\tilde{y}^p = y_K^p + \sum_s w_s(\tilde{L})\ \tilde{L}_s^p \tag{1.4}$$

To estimate $\tilde{y}^p$, we therefore need to characterize four sets of parameters: the initial joint distribution of labor and capital incomes; the initial joint distribution of wages $w_s(L)$ and schooling L; counterfactual education levels $\tilde{L}_s^p$; and counterfactual wages $w_s(\tilde{L})$.

1.1.1.2 Model Specification

Throughout the paper, I consider variants of the CES production function with two skill groups:

$$Y = \left(A_H L_H^{\frac{\sigma-1}{\sigma}} + A_L L_L^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \tag{1.5}$$

With A_H and A_L labor-augmenting technology terms, L_H and L_L the supplies of high-skill and low-skill labor, and σ the elasticity of substitution between H and L.

Assuming that workers are paid their marginal product, the return to schooling is:

$$r(L) = \log\left(\frac{w_H}{w_L}\right) = \frac{\sigma-1}{\sigma}\log\left(\frac{A_H}{A_L}\right) - \frac{1}{\sigma}\log\left(\frac{L_H}{L_L}\right) \tag{1.6}$$

The relative wage of skilled workers depends on their relative supply, as well as on the skill bias of technology. In particular, a 1% increase in the relative share of skilled workers is associated with a $\frac{1}{\sigma}$% decline in the skill premium. A higher elasticity of substitution implies greater substitutability between skill groups and a lower sensitivity of wages to relative supplies.

1.1.2 Aggregate Returns to Schooling

1.1.2.1 Individual Returns to Schooling and Supply Effects

I now characterize the aggregate effects of educational expansion in this model. We are interested in the effect of increasing L_H from an initial level L_H to a new level $\tilde{L}_H$. Denote the resulting change in the supply of skilled workers as $\Delta L_H = \tilde{L}_H - L_H$ and corresponding changes in wages as Δw_L and Δw_H. The effect of skill upgrading on output is:

$$\Delta Y = \tilde{Y} - Y = \Delta L_H \left(\underbrace{r(L)}_{\substack{\text{Initial Returns} \\ \text{to Schooling} \\ >0}} + \underbrace{\Delta r(L)}_{\substack{\text{Supply Effects} \\ \text{on Newly Skilled} \\ <0}} \right) + \underbrace{L_H \Delta w_H}_{\substack{\text{Supply Effects} \\ \text{on Always Skilled} \\ <0}} + \underbrace{L_L \Delta w_L}_{\substack{\text{Supply Effects} \\ \text{on Unskilled} \\ >0}} \tag{1.7}$$

With $r(L) = w_H - w_L$ the return to schooling at baseline and $\Delta r(L) = r(\tilde{L}) - r(L)$ the change in return to schooling induced by skill upgrading. The effect of human capital accumulation on earnings can be separated into four parts. First, educational upgrading increases the earnings of newly skilled workers through returns to schooling observed at baseline (initial returns to schooling). However, the resulting increase in the supply of skilled workers exerts downward pressure on these returns, mitigating the final benefits for this group (supply effects on newly skilled). For the same reasons, it also reduces the earnings of workers who were already skilled (supply effects on always skilled). Finally, low-skilled workers see their wage increase: the decline in their relative supply drives up their marginal productivity (supply effects on unskilled).

Supply effects thus mitigate the returns to schooling for those benefiting from educational expansion, at the same time as they redistribute income from high-skilled to low-skilled workers. In the case in which low-skilled and high-skilled workers are perfect substitutes, supply effects boil down to zero, and the change in output is $\Delta Y = \Delta L_H \, r(L) = \Delta L_H \, r$.

1.1.2.2 Initial, Final, and True Returns to Schooling

How do these four effects play out in practice, and where does the true aggregate effect of schooling lie? There are two natural options.

Initial Return One option would be to use the initial individual return to schooling observed *before* educational expansion, $r(L)$. This amounts to assuming that supply effects redistribute income between skill groups, but do not affect aggregate output. For instance, if the individual return declined from 10% to 8%, then 10% is the return that would be used. The change in output is then:

$$\Delta \bar{Y} = \Delta L_H r(L)$$

Final Return An alternative option would be to use the final return to schooling observed *after* educational expansion, $r(\tilde{L}) = r(L) - \frac{1}{\sigma}\Delta L_H$. This amounts to assuming that all supply effects are a net loss for the economy. If the individual return declined from 10% to 8%, then 8% is the return that should be used:

$$\Delta \underline{\mathrm{Y}} = \Delta L_H r(\tilde{L})$$

True Return The true effect turns out to lie in-between. The aggregate effect of education on output is lower than the initial individual return to schooling. Indeed, with imperfect substitution, there are decreasing returns to human capital accumulation: positive supply effects on unskilled workers are more than offset by negative supply effects on skilled workers (Caselli and Ciccone, 2013). Yet, because of decreasing returns, the true effect is also higher than the final return. Newly skilled workers end up with lower benefits than they might have hoped, but part of this loss benefits, on net, the unskilled:

$$\underbrace{r(\tilde{L})}_{\text{Final Returns}} \quad \leq \quad \underbrace{\frac{\Delta Y}{\Delta L_H}}_{\substack{\text{Aggregate Gains}\\ \text{Per Newly Skilled}}} \quad \leq \quad \underbrace{r(L)}_{\text{Initial Returns}} \tag{1.8}$$

A graphical illustration providing the main intuition can be found in appendix figure A.23. In the context of this paper, one is interested in estimating the effect of reducing education back to its 1980 level. One could use the 2019 (final) return to schooling, corresponding to the derivative of the production function before reducing education. Alternatively, one could use marginal gains from schooling observed after reducing education, corresponding to counterfactual (initial) returns. With imperfect substitution, the log of output is concave in schooling, implying that the true effect of education lies in-between these two estimates. As the elasticity of substitution increases, the production function becomes less concave: the 2019 returns to schooling become a better approximation of the output loss that would result from reducing education.

Assuming that the parameters of the production function are known, it is possible to re-express the true change in output as a function of a "true individual return" r^* that should be used. This return satisfies:

$$\tilde{Y} \quad = \quad w_H L_H + w_L \tilde{L}_L + \exp\left(\log(w_L) + r^*\right) \Delta L_H \tag{1.9}$$

Put simply, new output is the sum of wages received by the always skilled L_H at baseline (first term), wages received by the always unskilled $\tilde{L}_L$ at baseline (second term), and wages received by the newly skilled ΔL_H (third term), whose educational upgrading increases output by r^* log points. Rearranging yields a closed-form solution for the true aggregate return to schooling:

$$r^* \quad = \quad \log\left(\frac{\tilde{Y} - w_H L_H - w_L L_L}{\Delta L_H}\right) - \log(w_L) \tag{1.10}$$

Appendix A.2.2 provides a theoretical discussion, as well as results from a simple simulation using a CES production function, illustrating how the optimal return to schooling differs from initial and final returns depending on the elasticity of substitution and the skill bias of technology. For parametrizations similar to those found in the data, the optimal return to schooling is a weighted average of initial and final returns, with a typical weight on initial returns of 50-70%.

1.1.3 Distributional Effects of Schooling

I now turn to the effect of educational expansion on the income distribution. Consider two groups, rich R and poor P, who differ in their relative proportions of skilled and unskilled workers: $L_H^R > L_H^P$. Using equation 1.7, the effect of educational expansion on the rich-poor income gap can be expressed as:

$$\Delta Y^R - \Delta Y^P = \underbrace{\left(L_H^R - L_H^P\right)\Delta w_H}_{\substack{\text{Differential Supply} \\ \text{Effects on Always Skilled} \\ \leq 0}} + \underbrace{\left(L_L^R - L_L^P\right)\Delta w_L}_{\substack{\text{Differential Supply} \\ \text{Effects on Unskilled} \\ \leq 0}} + \underbrace{\left(\Delta L_H^R - \Delta L_H^P\right) r(\tilde{L})}_{\substack{\text{Differential Selection} \\ \text{Into Education}}} \tag{1.11}$$

The first two terms reveal that supply effects tend to reduce inequality. The intuition is simple. Increasing the share of skilled workers puts downward pressure on their wage. Because the high-income group has a greater fraction of skilled workers than the low-income group, this negative effect will be greater for them (differential supply effects on always skilled). Conversely, supply effects benefit unskilled workers. Because unskilled workers are concentrated at the bottom of the distribution, low-income groups will see a greater rise in their earnings as a result (differential supply effects on unskilled).

The third term of the equation highlights another important fact: the distributional effects of education also depend on which *type* of unskilled workers benefits most. If all low-skilled workers benefiting from better education originally come from R, in particular, this will increase inequality. In other words, *who* exactly benefits from educational expansion matters significantly for estimating the distributional effects of education. This mechanism will be particularly important for studying the relationship between education and gender inequality.

1.1.4 Estimation

I now introduce the methodology used to bring this framework to the data. The objective is to estimate the contribution of education to the real earnings growth of different groups p, which differ in their relative supply of workers belonging to skill groups s. A useful way to conceptualize this problem empirically is to formulate it as a counterfactual question: what would have been the distribution of income in 2019, had there been no progress in education since 1980? I propose to estimate this counterfactual in five steps, starting from microdata reporting information on the joint distribution of labor income and education (see appendix A.2.1 for more details).

1.1.4.1 Downgrade Education Levels

The first step is to reduce educational attainment to match its distribution in 1980: absent human capital accumulation, the distribution of skill groups would be $L^{1980} = (L_1^{1980}, ..., L_m^{1980})$ instead of $L^{2019} = (L_1^{2019}, ..., L_m^{2019})$. In practice, I implement this in the microdata by randomly sampling individuals and downgrading their education levels until reaching the counterfactual. I always give priority to workers whose education is closest to the targeted level. For instance, if the share of workers with no schooling rose from 20% in 1980 (counterfactual) to 40% in 2019 (observed), I construct the counterfactual by first reducing the education of primary-educated workers from primary education to no schooling. I then reduce the education levels of secondary- and tertiary-educated workers only if necessary, until reaching the targeted share of 20%. The resulting database thus contains "untreated" workers, whose education is unchanged, and "treated" workers whose educational attainment is downgraded by one or several levels. This method is similar to the one recently used by Hershbein, Kearney, and Pardue (2020) to simulate the economic effects of expanding access to higher education in the United States.

1.1.4.2 Reduce Wages Using Returns to Schooling

The second step is to reduce the earnings of treated workers by an estimate of the return to schooling. This implies calculating the return to schooling that should be used, which, as discussed in section 1.1.2, lies in-between initial and final returns.

Estimation of Initial and Final Returns In the present context, the final return is the return prevailing in 2019, that is, after educational expansion. This return can be estimated in the data using, for instance, a standard Mincerian wage regression.

In contrast, the initial return corresponds to the return that would prevail in 2019, had there been no educational expansion since 1980. This return is unobserved and has to be recovered from the model. With a CES production function:

$$r(L^{1980}) = r(L^{2019}) + \frac{1}{\sigma}\Delta \log\left(\frac{L_H}{L_L}\right) \tag{1.12}$$

Changes in the relative supply of skilled and unskilled workers are observed, so initial returns can be calculated for any chosen elasticity using this formula. The initial return is higher than the return observed in 2019: if education was to come back to its 1980 levels, the gains from human capital accumulation would appear substantially higher than those observed today.[5]

Estimation of the True Return The true return that should be used lies in-between initial and final returns, and can be calculated using equation 1.10. Wages and relative supplies are observed, so the only missing parameter to do so is A_H/A_L. Rearranging equation 1.6:

$$\frac{A_H}{A_L} = \exp\left[\frac{\sigma}{\sigma-1}\log\left(\frac{w_H}{w_L}\right) + \frac{1}{\sigma-1}\log\left(\frac{L_H}{L_L}\right)\right] \tag{1.13}$$

For a given value of the elasticity, A_H/A_L can thus be recovered from 2019 returns to schooling $\log(w_H/w_L)$ and 2019 relative supplies $\log(L_H/L_L)$ (as in, e.g., Rossi, 2022). This completely closes the model, and allows for a direct calculation of the true return r^* that should be used.

Application to the Data Once the true return to schooling is estimated, I directly reduce the earnings of treated workers by this return. For instance, workers downgraded from primary education to no schooling see their earnings reduced by the return to primary education.

1.1.4.3 Adjust Relative Wages to Account for Supply Effects

The third step is to account for the distributional incidence of supply effects. Changes in relative wages are given by:

$$\Delta\log\left(\frac{w_H}{w_L}\right) = -\frac{1}{\sigma}\Delta\log\left(\frac{L_H}{L_L}\right) \tag{1.14}$$

Assuming that σ is known, one can directly adjust relative wages in the microdata. For instance, for an increase in the relative supply of skilled workers of 1 log point and an elasticity of substitution of 6, the average wage gap between skilled and

[5] It is important to stress that this exercise does not amount to calculating the returns that actually prevailed in 1980. As evident from equation 1.6, returns to schooling are a function of both relative supplies L_H/L_L and the skill bias of technology A_H/A_L. The return observed in 1980 is thus the product of both 1980 relative supplies and 1980 technology. In contrast, the return we are interested in here is the return that would prevail in 2019 if relative supplies were to come back to their 1980 levels, but technology was to remain unchanged at its 2019 level. This return is by construction never observed and has to estimated. I come back to this in section 1.5.

unskilled workers is reduced by about 0.17 log points. Notice that the aggregate effect of educational expansion is entirely captured in step 2, so average income is left unchanged in this step of the estimation.

1.1.4.4 Derivation of Total Income

Steps 1 to 3 yield a counterfactual distribution of labor income absent educational expansion from 1980 to 2019. The fourth step is to move from this counterfactual distribution of labor income to a counterfactual distribution of total income. Assuming that we know the joint distribution of labor and capital income, this simply amounts to calculating $\tilde{y}^p = y_K^p + \tilde{y}_L^p$ for any given social group or quantile of the income distribution.

1.1.4.5 Growth Accounting

The final step is to calculate the share of growth in the real income of group p that can be attributed to human capital. This is equal to the gap between the counterfactual and actual growth rate, expressed as a fraction of actual growth from 1980 to 2019:

$$\text{Success}^p = \frac{g^p - \tilde{g}^p}{g^p} \tag{1.15}$$

With $g_t^p = \frac{y_{2019}^p - y_{1980}^p}{y_{1980}^p}$ and $\tilde{g}_{it} = \frac{\tilde{y}_{2019}^p - y_{1980}^p}{y_{1980}^p}$. If real income growth had been exactly the same absent human capital accumulation, then $\tilde{g}_{it} = g_{it}$ and hence $\text{Success}^p = 0\%$. On the contrary, if there would have not been any growth at all in the absence of educational progress, then $\tilde{g}_{it} = 0$ and $\text{Success}^p = 100\%$: human capital can explain all of economic growth. Notice that as in standard growth accounting, success can be higher than 100%, if actual growth rates fall below those predicted by changes in educational attainment.

1.2 Data and Methodology

This section presents the data sources and methodology used to estimate the contribution of education to global income and gender inequality reduction since 1980. First, I combine a new set of surveys covering education and wages in 150 countries (section 1.2.1) with data on the evolution of educational attainment since 1980 (section 1.2.2). Using estimates of returns to schooling (section 1.2.3) and of supply effects induced by educational expansion (section 1.2.4), I estimate by how much lower would wages be, within and between countries, had there been no educational progress since 1980.

Finally, I exploit data on the actual evolution of the global income distribution to construct the distributional growth accounting decomposition (section 1.2.5). Section 1.2.6 validates the overall methodology using new evidence from three large-scale education policies in India, Indonesia, and the United States.

1.2.1 Survey Microdata

The starting point is a unique set of household surveys covering the joint distribution of personal income and education by age and gender in 150 countries, which I have assembled for this paper. These surveys come from two main sources.

The first data source is the International Labor Organization's database of labor force surveys. Based on a considerable data collection effort and with the collaboration of national statistical institutes, ILOSTAT have harmonized over 1,300 household surveys, covering 130 countries over the 1990-2022 period. The database records individual-level information on wages, self-employment income, education, and other sociodemographic variables such as age, gender, occupation, industry, and hours worked. All surveys are nationally representative. Most surveys are labor force surveys, which aim specifically at capturing the dynamics of employment and earned income. In some countries where labor force surveys do not exist, the ILO has instead collected consumption or living standards surveys, which generally aim to record household expenditure but also contain information on individual wages and self-employment. In the main analysis, I use the latest survey available in each country.

The coverage of ILO microdata is remarkable, but the information collected on education is limited in some countries. Furthermore, a number of countries are missing, including big countries such as China and Russia. To expand the coverage and quality of the database, I turn to the websites of national statistical institutes and other sources, from which I collect additional household surveys for 55 countries. These include the European Union statistics on income and living conditions (EU-SILC), providing individual microdata for 32 European countries, as well as the Life in Transition Survey (LITS), which covers 10 additional countries in Eastern Europe and Central Asia. I complement these two cross-national datasets with surveys available from country-specific data portals. These sources allow me to cover 13 additional countries: China, Iraq, India, Japan, Mozambique, Morocco, Russia, Somalia, South Africa, South Korea, South Sudan, Tunisia, and the United States. In each case, I collect detailed information on personal income, education, and other sociodemographic variables, which I harmonize in the same way as the ILO. More details can be found in appendix A.5.

Figure 1.2 maps the geographic coverage of the resulting database. Table 1.1 provides descriptive statistics. The data cover about 9.6 million individuals surveyed in 150 countries. These surveys are representative of over 95% of the population of each world region. The exception is the Middle East and North Africa, where microdata is well known to be either non-existent or inaccessible to researchers (Ekhator-Mobayode and Hoogeveen, 2022). Overall, the microdata cover about 95% of the world's population, and about 93% of the world's GDP.[6] To the best of my knowledge, this represents the first micro-database on the world distribution of income ever built in economics research.

1.2.2 Educational Attainment Data

The first step of the estimation consists in downgrading the education of individuals to match its distribution observed in 1980. This requires data on the evolution of educational attainment in each country. Here, the primary source is the database compiled by Barro and Lee (2013) and updates.[7] It records estimates of the share of individuals with no schooling, primary education, secondary education, and higher education by 10-year age groups and gender, in 146 countries, from 1950 to 2015. A number of countries are covered by the survey microdata but absent from the Barro-Lee database. For these, I construct my own estimates, using either census data from IPUMS International, cohort-level trends observed in the labor force surveys, or other sources (see appendix A.6).

For the counterfactual to be valid, it is important to ensure that education categories recorded in surveys match those observed in the Barro-Lee database. This is particularly challenging for a number of reasons. For instance, the ILO sometimes includes incomplete degrees in each education level and sometimes does not, depending on the information available in each survey. Meanwhile, the Barro-Lee database sometimes includes lower secondary education with primary education. Comparing the distribution of educational attainment in the two sources, I manually map categories in each country, one by one, until estimates from the two datasets coincide as close as possible. Finally, to ensure that the construction of the counterfactual is perfectly consistent, I re-calibrate the sample weights in each survey to make them match the distribution of educational attainment by age and gender recorded in the Barro-Lee

[6] In nearly all countries, the survey was fielded after 2015, ensuring that the database is broadly representative of the global distribution of income over the 2015-2019 period. Appendix figure A.44 provides a map of the corresponding survey years in each country.

[7] See `https://barrolee.github.io/BarroLeeDataSet/BLv3.html`.

database. This last step only marginally affects weights, given that the distribution of educational attainment is very close in the two sources after reclassification.[8]

Figure 1.3 plots the distribution of global educational attainment from 1980 to 2019. There has been a dramatic expansion of secondary education, from about 25% to 60%. This rise was mirrored by a significant decline in the share of adults with primary education or no schooling. In 2019, less than 15% of working-age adults had not attended at least some years of basic education. Tertiary education also expanded significantly, from less than 5% in 1980 to over 10% in 2019.[9]

1.2.3 Returns to Schooling

The second step is to reduce individual incomes by the returns to schooling. To do so, one needs estimates of returns to schooling observed in 2019 in each country (corresponding, as discussed in section 1.1.4, to final returns). There are two options, each of which with advantages and disadvantages: returns to schooling estimated by OLS in my data, or causally identified returns available from the existing literature.

1.2.3.1 OLS Returns

In the main analysis, I rely on OLS estimates of returns to schooling by education level, derived from a modified Mincerian equation of the form:

$$\ln y_{ic} = \alpha_c + \beta_c^{pri} D_{ic}^{pri} + \beta_c^{sec} D_{ic}^{sec} + \beta_c^{ter} D_{ic}^{ter} + X_{ic}\beta_c + \varepsilon_{ic} \tag{1.16}$$

With y_{ic} earned income of individual i in country c, D_{ic}^{pri}, D_{ic}^{sec}, and D_{ic}^{ter} dummies for having reached primary, secondary, and tertiary education, and X_{ic} a vector of controls including gender, an experience quartic, and interactions between gender and the experience quartic (as in Lemieux, 2006; Autor, Goldin, and Katz, 2020). I restrict the sample to all individuals with positive personal income, including both wage earners and self-employed individuals. The dependent variable is total annual earned income from all jobs.

[8]Appendix figures A.45, A.46, A.47, and A.48 compare estimates of the share of the working-age population with no schooling, primary education, secondary education, and tertiary education in the Barro-Lee database and the survey microdata, after manually mapping educational categories in the two sources.

[9]Improvements in schooling have coincided with significant declines in overall educational attainment inequalities, both between and within countries: see appendix figure A.21, which provides a Theil decomposition of global educational attainment inequality following the methodology first proposed by Morrisson and Murtin (2013). Appendix figure A.22 charts average years of schooling of the working-age population by world region since 1980.

Appendix A.7 provides results for other empirical specifications, variable definitions, and sample restrictions; the resulting returns are largely insensitive to alternative methodological choices. In particular, returns to schooling on annual per capita household consumption, available for a subset of eleven low- and middle-income countries, are almost identical to returns estimated on personal income. This provides reassuring evidence that estimates are not biased by selection into formal employment or hours worked, both of which vary significantly with schooling and across countries (e.g., Bick, Fuchs-Schündeln, and Lagakos, 2018).

The main advantage of OLS returns lies in their coverage and comparability: they can be estimated for all 150 countries using a unified methodology. The potential disadvantages are twofold.

A first source of concern is that returns estimated by OLS may suffer from omitted variable bias. A classical argument is that ability bias may lead to overestimating the return to schooling: more educated workers have higher earnings because of greater innate abilities, not because of schooling itself. Reassuringly, several decades of labor economics literature teaches us that omitted variable bias is typically small and, if anything, tends to go in the opposite direction (for reviews, see Card, 1999; Deming, 2022; Gunderson and Oreopolous, 2020). Returns estimated by OLS are thus likely to provide a good approximation of average individual returns to schooling, if not a lower bound.

A second source of concern is that our parameter of interest is not the *average* return to schooling: it is the return to schooling for those who benefited from educational expansion since 1980.[10] Given potentially large heterogeneity in returns across individuals (e.g., Heckman, Humphries, and Veramendi, 2018), our parameter of interest may differ significantly from the one estimated over the entire population. In particular, improved access to schooling has predominantly benefited children from lower socioeconomic backgrounds (Gethin, 2023b), a population that is often found to have higher returns. The next section provides evidence that this is indeed likely to be the case.

Figure 1.4 plots the resulting estimates of returns to schooling by world region. Returns are strongly convex, in particular in low-income countries. In Sub-Saharan

[10]Consider for instance a country in which the share of workers with primary education increased from 10% to 20% between 1980 and 2019. Intuitively, this implies that by 2019, 10% of workers correspond to workers who would have obtained primary education anyway in a world with no educational expansion (the "always skilled"), while another 10% benefited from increased access to schooling (the "newly skilled"). One would like to estimate the return for these 10% of newly skilled, yet an OLS estimate over the entire sample will capture both groups.

Africa, the average return to a year of primary education is 6%, while the average return to a year of tertiary education is 22%. Returns to primary education are extremely low in India, barely reaching 3%. A direct implication is that using the average return would likely lead to overestimating the contribution of education to global poverty reduction, given that the global poor have mostly benefited from expansions in access to basic education, which displays the lowest returns. There are also significant variations in average returns across regions. Europe and the United States have the highest average return to schooling, at over 12%, while it falls below 6% in the Middle East and North Africa.[11]

1.2.3.2 IV Returns

An alternative option is to use instrumental variable estimates of returns to schooling, in particular those derived from differential exposure to large-scale education programs. The estimation of these returns generally relies on comparing cohorts or regions that were more or less exposed to specific policies, such as compulsory schooling laws or school construction programs. They are causally identified, so they do not suffer from omitted variable bias. Another major advantage is that they focus by construction on the newly skilled, since they are based on comparing the earnings of those who marginally gained access to education to those who did not. The main drawback is that they can only cover specific programs expanding access to specific levels of education (although the estimates compiled below do end up covering about 60% of the world's population).

The labor economics literature has made considerable efforts at expanding such estimates to multiple contexts and policies in recent years, in particular in developing countries. For the purpose of this paper, I have assembled a collection of IV estimates of the returns to schooling, drawing on a number of recent studies.

I select estimates based on four criteria. First, I give priority to articles studying episodes of large-scale expansions in access to schooling. Second, I select studies for which a comparable OLS estimate is available for comparison. Third, I restrict the sample to relatively recent studies, covering policies that expanded access to education during my period of interest (see Card (1999) for similar findings based on older studies mostly conducted in rich countries). Fourth, I select one estimate per

[11] Appendix figure A.50 plots the cross-country distribution of returns to schooling, while appendix figures A.51, A.52, A.53, and A.54 map average returns to schooling and returns to primary, secondary, and tertiary education in all countries with available data. The median returns to primary, secondary, and tertiary education are 5%, 9%, and 13%, respectively. See also figure A.49, which presents results of cross-country regressions by education level and further decomposes secondary education into upper secondary and lower secondary.

country and education level.

Figure 1.5 plots the resulting OLS and IV estimates of the return to schooling. Two results stand out. First, OLS and IV returns are highly correlated. This provides reassuring evidence that OLS estimates capture true variations in returns to schooling across countries and education levels relatively well. Second, IV returns are almost always higher than OLS returns or not significantly different (in line with previous estimates compiled by Card (1999) for developed countries). In some countries, such as China, IV returns are two to three times larger. Appendix figure A.55 plots the ratio of IV to OLS estimates across studies. The gap ranges from close to 0% in Nigeria to almost 250% in China (tertiary). The average gap is in the order of 80%.

As mentioned above, I use OLS returns in the main analysis. As an alternative specification, I exploit IV estimates to correct returns to schooling upwards in three steps. First, I multiply OLS estimates by the ratios plotted in appendix figure A.55 for each country-level covered by the data (for instance, I increase the return to tertiary education in Vietnam by 50%). Second, I multiply the return to other levels of education by the same ratio (for instance, I increase the returns to primary and secondary education in Vietnam by 50% too). This amounts to assuming that returns to other levels of schooling are underestimated by the same factor in each country. Third, I increase returns to schooling in missing countries by the average correction factor observed, that is, 80%.

1.2.4 Supply Effects

The third step of the methodology is to account for supply effects, which are necessary to estimate the true return to schooling (section 1.1.2) and the distributional effects of schooling expansion (section 1.1.3). This implies extending the CES production function to more than two levels of schooling, and calibrating elasticities of substitution between skill groups.

1.2.4.1 Production Function Specification

Until now, I have worked with a CES production function with two skill groups, yet the data provide the distribution of educational attainment for four: workers with no schooling, primary education, secondary education, and tertiary education. To incorporate supply effects on these four skill groups, I introduce nests in the CES production function, in line with previous work in labor economics (see in particular Goldin and Katz (2007) on the United States and Fernández and Messina (2018) on Latin America).

At the top level, output is produced by combining workers with tertiary education and workers with below tertiary education:

$$L = \left(A_{\bar{ter}} L_{\bar{ter}}^{\frac{\sigma_1 - 1}{\sigma_1}} + A_{ter} L_{ter}^{\frac{\sigma_1 - 1}{\sigma_1}} \right)^{\frac{\sigma_1}{\sigma_1 - 1}} \tag{1.17}$$

With L_{ter} the share of college-educated workers and $L_{\bar{ter}} = 1 - L_{ter}$. The subgroup of workers with less than tertiary education is split into workers with secondary education and workers with below secondary education:

$$L_{\bar{ter}} = \left(A_{\bar{sec}} L_{\bar{sec}}^{\frac{\sigma_2 - 1}{\sigma_2}} + A_{sec} L_{sec}^{\frac{\sigma_2 - 1}{\sigma_2}} \right)^{\frac{\sigma_2}{\sigma_2 - 1}} \tag{1.18}$$

With L_{sec} the share of secondary-educated workers and $L_{\bar{sec}} = 1 - L_{sec} - L_{ter}$ the share of workers with less than secondary education. This corresponds exactly to the specification adopted by Goldin and Katz (2007) and Fernández and Messina (2018). Finally, because my data also include countries with a significant fraction of workers with below primary education, I introduce a third nest separating workers with primary education from those with no schooling:

$$L_{\bar{sec}} = \left(A_{non} L_{non}^{\frac{\sigma_3 - 1}{\sigma_3}} + A_{pri} L_{pri}^{\frac{\sigma_3 - 1}{\sigma_3}} \right)^{\frac{\sigma_3}{\sigma_3 - 1}} \tag{1.19}$$

With L_{pri} the share of workers with primary education and L_{non} the share of workers with no schooling. This cutoff has been adopted to study episodes of expansions in access to primary education (see in particular Khanna (2023) on India).

These three equations yield three formulas for the returns to primary education, secondary education, and tertiary education, which map directly onto the categories reported in the Barro-Lee database and the returns to schooling plotted in figure 1.4. True returns to schooling (equation 1.10) and relative wage adjustments (equation 1.14) can then be calculated separately for each of the three nests.[12]

[12]Appendix figures A.28, A.29, and A.30 plot relative skill efficiency for the three levels of the production function versus GDP per capita. In line with recent evidence by Rossi (2022), there is a strong correlation between the relative efficiency of skilled labor and economic development: skilled workers are substantially more efficient in rich countries. The United States stand out as displaying a particularly large relative efficiency of tertiary-educated workers. This directly results from the return to tertiary education being very large in spite of the U.S. having a high share of tertiary-educated workers.

1.2.4.2 Elasticities of Substitution

To close the model, the only parameters that need to be calibrated are the elasticities of substitution between skill groups. A rich empirical literature in labor economics has attempted to estimate these elasticities in various contexts, with resulting values ranging from 1.5 to 5 and typically close to 2, depending on the methodology used, the chosen skill cutoff, and the country considered. Appendix table A.18 reports selected estimates from existing empirical studies.

Most studies rely on short-run variations in relative skill supply to estimate elasticities of substitution. As a result, they identify a relatively short-run elasticity, corresponding to the case in which the skill bias of technology (A_H/A_L) is held approximately constant. Yet in the long run, one should expect firms to adjust their technological mix as a response to the greater supply of skilled workers (e.g., Acemoglu, 1998). This would imply less sensitivity of relative wages to relative supplies, and hence greater values of the elasticity of substitution after accounting for endogenous technical change.

Motivated by this fact, a recent literature in macroeconomics attempts to estimate long-run elasticities of substitution between skill groups. Two recent articles, in particular, have made significant progress in this direction. Exploiting data on wage gains at migration to the United States for different skill groups, Hendricks and Schoellman (2023) find a long-run elasticity ranging from 4.5 to 8, depending on the cutoff chosen to define high-skilled workers. Bils, Kaymak, and Wu (2022) instead exploit data on worldwide trends in education and returns to schooling to pin down values of the elasticity of substitution consistent with both no worldwide technological regress and a greater skill bias of technology in rich countries. These two conditions yield lower and upper bounds of 4 and 6, very similar to the numbers obtained by Hendricks and Schoellman (2023).

To the extent that this article analyses large improvements in educational attainment over several decades, long-run elasticities of substitution are the relevant parameters. In my main analysis, I thus calibrate $\sigma_1 = \sigma_2 = \sigma_3 = 6$, in line with estimates reported in Hendricks and Schoellman (2023) and Bils, Kaymak, and Wu (2022). I discuss the robustness of my results to alternative specifications in section 1.3.4.

1.2.5 Global Labor and Capital Income Inequality Data

The final step of the estimation consists in moving from labor income to total income, and comparing counterfactual to actual real income growth rates. This requires data on the distribution of income in each country, aggregate labor and capital income shares, and the share of income received from labor and capital by income group within each country.

1.2.5.1 Global Income Inequality Data

Data on the world distribution of income come from the World Inequality Database (WID). The database covers average per-capita income by percentile in all countries in the world, every year from 1980 to 2019. The income concept is pretax national income, that is, total income received by individuals before accounting for taxes and transfers, but after accounting for the operation of pension and unemployment systems. Importantly, all components of net national income (GDP, minus consumption of fixed capital, plus net foreign income) are allocated to individuals, following the Distributional National Accounts (DINA) framework (see Chancel et al., 2022b; Piketty, Saez, and Zucman, 2018). This ensures that all income distributions are consistent with macroeconomic growth rates and aggregate capital and labor income shares recorded in the national accounts. The database is constructed by compiling estimates from detailed national or regional studies, which combine surveys, tax data, and national accounts to construct distributions that are conceptually comparable across countries (see for instance Piketty, Saez, and Zucman (2018) on the United States, Blanchet, Chancel, and Gethin (2022) on Europe, and De Rosa, Flores, and Morgan (2022a) on Latin America).

1.2.5.2 Aggregate Labor and Capital Income Shares

Aggregate factor income shares come from Bachas et al. (2022), who combine a number of sources to build a new database on the components of net national income worldwide since 1965. Their database provides a decomposition of net domestic product into compensation of employees, mixed income, the operating surplus of households (actual and imputed rental income), and the operating surplus of corporations (profits net of depreciation).

I define the labor income share as the share of income attributable to compensation of employees and mixed income. This is the definition of the labor share that is the most conceptually meaningful in my context, given that my microdata cover individual income and returns to schooling for both wage earners and the self-employed. In the main analysis, I thus make the conservative assumption that human capital only affects wages and mixed income, while leaving capital income unchanged.

1.2.5.3 Capital Income Concentration

The last step is to estimate how the capital income share varies alongside the income distribution in each country. High-quality data on this decomposition are scarce, given well-known issues with the underestimation of capital income in household surveys. Drawing on the few studies that were able to mobilize tax and national accounts data to estimate such decomposition with a relatively good level of precision, I was able to derive profiles of the capital income share by percentile for the United States (Piketty, Saez, and Zucman, 2018), South Africa (Chatterjee, Czajka, and Gethin, 2022), and 10 Latin American countries (De Rosa, Flores, and Morgan, 2022a). The corresponding series are plotted in appendix figure A.20. The profiles look very similar across these three cases. The capital share is always below 20% for the bottom 90% of earners, corresponding mostly to imputed rental income. It rises exponentially at the very top of the distribution, where the main source of income is from bonds and stock. Given these similarities, I use the average profile observed across countries, which I rescale in each country-year to match the aggregate capital income share.

1.2.6 Validation from Three Natural Experiments

Despite the relative popularity that the canonical labor demand and supply framework has encountered in the literature, the ability of my estimates to capture the true aggregate and distributional effects of educational expansion may naturally be questioned. There are two main sources of concern. First, one might be worried that individual returns to schooling reflect signaling rather than true increases in productivity. This would imply that the returns to schooling used in this paper overestimate the effect of schooling on aggregate growth. Second, there might be heterogeneity in who benefits from educational expansion, the associated returns to schooling, and general equilibrium effects. This would imply that the relatively simple methodology adopted here might underestimate or overestimate benefits of schooling for low-income groups.

To address these two potential limitations and assess the overall validity of the distributional growth accounting approach, I turn to causal evidence from three natural experiments. I focus on outlining the main results. The interested reader will find more details in appendix A.3.

1.2.6.1 Contexts, Data, and Methodology

I investigate the aggregate and distributional effects of three large-scale education policies: India's District Primary Education Program (1990s-2000s), Indonesia's INPRES school construction program (1970s), and U.S. state compulsory schooling laws (1870s-1960s). These three sets of policies have been extensively studied to estimate individual returns to schooling, human capital externalities, and general equilibrium effects affecting different skill groups (e.g., Acemoglu and Angrist, 2000; Duflo, 2001, 2004; Khanna, 2023). Less is known of their exact effects on real income growth by income group.

Combining data from existing studies and additional sources, I exploit these natural experiments to estimate the causal effect of educational expansion on aggregate economic growth and its distribution across subnational regions. I run variants of the following specification:

$$\ln y_{rt}^i = \gamma_0^i + \gamma_1^i S_{rt} + X_{rt}^i \beta + \delta_r + \delta_t + \varepsilon_{rt} \quad (1.20)$$

$$S_{rt} = \alpha_0 + \alpha_1 Z_{rt} + \eta_{rt} \quad (1.21)$$

Where y_{rt}^i denotes the average income of income group i (such as the bottom 20% of earners) in subnational region r at time t. The objective is to estimate the impact of an exogenous increase in S_{rt}, average years of schooling of the working-age population in region r. X_{rt}^i is a vector of controls, such as the demographic composition of the region, δ_r are subnational region fixed effects, and δ_t are time fixed effects. The parameters of interest are γ_1^i for different groups i, which provide reduced-form estimates of the effect of increasing average years of schooling on real average income.

S_{rt} is instrumented by Z_{rt}, a variable capturing quasi-experimental variation in exposure to the education program. In India, I rely on Khanna (2023), who estimates the impact of the DPP using a regression discontinuity design around the cutoff district literacy rate used to allocate the program. In Indonesia, I instrument district average years of schooling by the number of schools built under the INPRES program, following Duflo (2001). In the United States, the instrument is average required years of schooling across cohorts born in different states, as in the existing literature (Acemoglu and Angrist, 2000; Guo, Roys, and Seshadri, 2018).

After estimating the effects of each program, I then compare these results to aggregate and distributional effects predicted by the model. In India, for instance, I simulate the effect of increasing average years of schooling by one year through primary education, following each step of the methodology outlined in section 1.1.4. This yields simulated estimates of γ_1^i, which can be compared to those obtained empirically from the natural experiment.

1.2.6.2 Main Results

Figures 1.6, 1.7, and 1.8 plot the main results, comparing the estimated and simulated effects of educational expansion on the average income of each income quintile. All three policies led to large reductions in income inequality. In India, for instance, increasing average years of schooling by one year in a treated district is associated with a 20% increase in the average income of the bottom quintile in this district, compared to a null effect on the average income of the top 20%. Aggregate effects of education on earnings are found to range from 8% to 15% per average year of schooling (as shown by the dashed line in each figure), which is relatively close to individual Mincerian returns. This provides direct causal evidence against the idea that individual returns to schooling only reflect signaling.

The model performs remarkably well at reproducing results from these natural experiments. In all three cases, simulations predict significantly higher returns to schooling expansion at the bottom of the distribution. This is because (i) these policies targeted basic education, which disproportionately benefits low-income earners in the simulation, and (ii) supply effects magnify this redistributive effect. If anything, the model slightly underestimates benefits at the bottom of the distribution, in particular in the United States. Together, these results provide reassuring evidence that the methodology developed in this paper performs well at capturing the distributional effects of educational expansion, and may even provide a lower bound on true benefits for low-income earners.

1.3 Education and the World Distribution of Income

This section presents the main results on the role of education in reducing global poverty and inequality. Section 1.3.1 focuses on the overall distribution of global economic growth since 1980. Section 1.3.2 decomposes the effects of education by world region, time period, and between and within countries. Section 1.3.3 compares the results to those of a standard growth accounting decomposition, isolating the contribution of each step of the distributional growth accounting methodology. Section 1.3.4 provides various robustness checks and extensions.

1.3.1 Education and the Distribution of Global Economic Growth

I start by presenting results on the role of education in shaping the distribution of global economic growth since 1980. Table 1.2 presents a distributional growth accounting decomposition of the world distribution of income for the 1980-2019 period.

1) Education Explains 50% of Average Economic Growth Global average income per capita approximately doubled over this period (+98%). Absent educational expansion, growth would have been significantly lower, at 45%. Education thus contributed 53 percentage points of growth. Taking the ratio between the contribution of education and actual economic growth, private returns to schooling explain about 54% of average per capita income growth since 1980.

2) Education Explains 70% of Growth for the Global Poorest 20% This average figure hides significant heterogeneity by global income group. For the poorest 50% of individuals in the world, growth has been markedly higher, exceeding 150%. Yet, benefits from educational expansion have also been higher for this group, so that the share of growth explained by education reaches almost 60%. Overall, private returns to schooling can account for 55% to 75% of growth for the global bottom 90%. This share is highest for the bottom 20% (71%) and middle 40% of the income distribution (74%), two groups that have witnessed lower growth and large gains from increased access to schooling. It is lowest at the very top of the distribution, mainly because the bulk of income at the top is received from capital—which by assumption is not affected by schooling.

Figure 1.1 provides a more granular picture of the distribution of global economic growth since 1980. All individuals in the world are ranked from the poorest 1% to the richest 0.01%. Total pretax income growth is then calculated for each percentile, together with growth explained by private returns to schooling (lower shaded area) and residual growth coming from other factors (upper shaded area). Real income gains have been greatest at the middle and the very top of the global income distribution, generating what has often been referred to as the "elephant curve" of global inequality and growth (Lakner and Milanovic, 2016). This pattern reflects the conjunction of trends in inequality between and within countries, including the rise of China and India (middle of the distribution), sluggish economic growth in low-income countries (bottom of the distribution), weak income gains for most households living in high-income countries (upper middle of the distribution), and skyrocketing top income inequality in many parts of the world (top end of the distribution).[13] The main contribution of this paper is to isolate gains from education, represented by the lower shaded area. These gains have been particularly large for most income groups, ranging from 80 points to 120 points for most percentiles within the bottom 90%.

Taking the ratio of the contribution of education to actual growth rates yields figure 1.9, which represents the share of growth explained by education by global income percentile. This share ranges from 55% to 90% for all income groups within the bottom 90%. It is highest at the bottom and the upper middle of the income distribution, exceeding two-thirds for the poorest 20% and for groups ranging from the 70$^{\text{th}}$ to the 90$^{\text{th}}$ percentiles.

3) Education Explains 40% of Extreme Poverty Reduction Beyond growth for specific groups, another indicator that has received considerable attention is the share of the world's individuals living in extreme poverty. A difficulty in the context of this paper is that poverty headcount ratios are based on counting the number of individuals whose income falls below a certain threshold rather than on actual income gains. This makes the calculation of the share of poverty reduction explained by education less conceptually meaningful (since it implies counting people rather than comparing growth rates) and more sensitive to the choice of a specific threshold (given that poverty rates are not necessarily linear in growth rates).

Another limitation is that estimated poverty rates can differ significantly across sources. Most commonly used estimates are those provided by the World Bank, but these are not ideal in the context of this paper for two main reasons. First, they rely exclusively on data from household surveys, which generally miss capital income entirely (Blanchet, Flores, and Morgan, 2022) and can display growth rates that differ substantially from those reported in the national accounts (Pinkovskiy and Sala-i-Martin, 2016). Second, they are based on household expenditure rather than pretax income, which may bias the results depending on the size and distributional incidence of taxes, government transfers, and saving rates. An alternative solution is to rely on pretax income distributions from the World Inequality database, which have the advantage of covering the correct income concept and being consistent with

[13] The interested reader will find additional figures describing these trends in more detail in the appendix. In particular, appendix figure A.14 provides a Theil decomposition of global income inequality since 1980, while figure A.15 charts the average share of pretax income received by the richest 10% by world region. Figures A.16, A.17, A.18, and A.19 provide further graphical evidence on the geographical breakdown of global income groups in 1980 and 2019, as well as the geographical composition of the world's poorest 20% and richest 20% since 1980.

macroeconomic growth rates. The difficulty is that poverty thresholds were designed by the World Bank to match deprivation levels reported in surveys, not GDP per capita levels.

With these limitations in mind, table 1.3 extends the growth accounting decomposition to global poverty headcount ratios at $2.15, $3.65, and $6.85 per day (the three thresholds typically used by the World Bank), calculated using pretax income distributions from the World Inequality Database. Poverty at $2.15 per day fell from 20% in 1980 to 9% in 2019. Absent educational expansion, it would be about 4 percentage points higher today, implying a decline in the global poverty headcount ratio of 32% instead of the 55% observed. By this measure, private returns to schooling can account for 42% of global extreme poverty reduction. The corresponding figures are 32% at $3.65 per day and 44% at $6.85 per day.

As an alternative, I reproduce this analysis using World Bank poverty rates.[14] The results are presented in appendix table A.3. Education explains 40% of global poverty reduction at $2.15 per day, 56% at $3.65 per day, and 73% at $6.85 per day, similar to or higher than the figures obtained with the WID data. The takeaway is again that education has been a major driver of improved living standards for the world's poorest individuals.

1.3.2 Decomposing Global Schooling Gains

Faced with these results, one may naturally wonder what are the different factors driving them. Why does education explain a higher share of growth at the bottom and upper-middle of the world distribution of income? Are the results primarily driven by the distribution of growth within countries or by differences in aggregate gains from schooling across countries? This section attempts to answer these questions.

[14] The World Bank does not publish data on the world distribution of income. I thus reconstruct it myself by collecting income and consumption distributions from the World Bank's website and extrapolating the average income of each country-percentile to missing years using real GDP per capita growth rates. This yields trends in global poverty almost identical to those officially reported by the World Bank. Finally, I construct counterfactual income distributions using the same methodology as in the rest of the paper. The main difference is that capital income is absent from World Bank surveys, implying a 100% passthrough of schooling on income instead of a passthrough equal to the labor share. In addition to the poverty analysis, appendix table A.2 reproduces the main distributional growth accounting decomposition using World Bank data. With this dataset, education accounts for about 60% of global bottom 20% growth since 1980.

1.3.2.1 Distributional Growth Accounting by World Region

A first way of better understanding the results is to decompose them by world region. Table 1.4 displays growth decompositions for selected geographical regions and country income groups, distinguishing between average economic growth and real income growth of the poorest 20%. Two main results stand out.

1) Education Explains More of Growth in Low- and High-Income Countries Looking at aggregate growth figures, education explains the totality of growth for low-income countries and about 60% of growth for high-income countries. Sub-Saharan Africa and Latin America are the two world regions where education explains the highest share of growth, despite the fact that gains from education have been the lowest. The reason is simply that these are the two regions that have witnessed the lowest average growth rates over the period considered. Indeed, one should remind the reader that in countries where growth has been close to zero or negative, education explains over 100% of growth by construction, as in any growth accounting exercise (Barro and Lee, 2015). The interpretation is simply that other factors than education—such as civil wars and economic crises—have affected growth negatively, leading to a negative counterfactual growth rate.

In contrast, private returns to schooling alone cannot account for the exceptional growth rates witnessed by China and India, despite substantial improvements in schooling in these two countries. Interestingly, although economic growth has been about two times slower in India than in China, education can explain about the same share of income gains in the two countries. The main reason is that a significant fraction of educational expansion in India has occurred at the level of basic education, which displays exceptionally low returns (figure 1.4), while China has primarily benefited from expansions in secondary and tertiary education.

2) Education Explains Over 50% of Growth For Low-Income Groups in All Regions In each region or country, education almost always explains more growth at the bottom of the distribution than for the average individual, for two main reasons. First, because of rising inequality in many countries, actual growth has been significantly lower at the bottom of the distribution (in particular in Europe, Northern America, China, and India). Education thus explains more of growth at the bottom of the distribution, simply because there is less growth to be explained. Secondly, gains from schooling have been greater for low-income earners than for high-income earners in most world regions, mainly because supply effects tend to reduce inequality and because capital income is concentrated at the top of the distribution

(I come back to this in section 1.3.3). As a result, education can explain more than 100% of growth for the bottom 20% of earners in Western economies, Latin America, MENA, and Sub-Saharan Africa.[15] Even in India, which displays extraordinarily low returns to basic education, private returns to schooling can explain about half of real income gains for the poorest 20%.

Combining these two facts enables a better understanding of the patterns presented in figure 1.1. Economic growth has been the highest for Chinese and Indian middle classes, corresponding to income groups near the median of the global income distribution. Although these groups have benefited from some of the highest gains from schooling, these gains cannot fully account for such exceptional growth rates. At the upper-middle of the global income distribution, stagnating real incomes for European and US low-income earners have coincided with relatively high gains from schooling for these groups, which is why education explains the bulk of growth for percentiles 70 to 90. Finally, growth has been weak for the world's poorest individuals, mainly due to low or even negative economic growth in Sub-Saharan Africa, but also because of rapidly rising inequality in middle-income countries (in particular India). Schooling gains have not been particularly impressive in low-income countries either, but they have been high for low-income earners of both low- and middle-income countries. This explains why education can account for such a large share of growth at the bottom of the global income distribution.

1.3.2.2 Distributional Growth Accounting Within and Between Countries

Given these complex patterns of educational expansion affecting inequality within and between countries, has education been a driver of higher or lower global income inequality overall? One way of answering this question is to perform a Theil decomposition of global inequality into a between-country and a within-country component. This decomposition is reported in table 1.5.

1) Education Has Prevented the Rise of Global Inequality Since 1980
The first two rows of table 1.5 compare the evolution of the Theil index of worldwide inequality to its counterfactual evolution absent schooling expansion. From 1980 to 2019, global inequality more or less stagnated. Absent educational progress, it would have instead risen dramatically, from 1.06 to 1.33. Educational progress has thus contributed to strongly reducing global inequality in the past decades.

[15] Again, it is important to mention that figures exceeding 100% should not be seen as extraordinary. In the United States, for instance, the real average income of the bottom 20% *declined* over the 1980-2019 period, which implies that education necessarily explains over 100% of growth for that group.

2) Education Has Reduced Inequality Within Countries, But Not Between Countries The next two rows of table 1.5 compare actual and counterfactual trends in inequality between countries, measured by the between-country component of the Theil index. Education has not had much effect on cross-country income convergence. Inequalities between countries declined strongly, from 0.6 to 0.34, mainly because of the rise of China and India. This decline would have been the same absent educational progress. This result mirrors the complex patterns highlighted above, with education explaining more growth in low-income and high-income countries than in middle-income countries. The effect of education on the overall dispersion of cross-country average incomes appears to have been close to zero.

While the effect of education on between-country inequality is unclear, schooling has been an unambiguous driver of convergence within countries. The Theil index of within-country inequalities grew by 0.28 points over the 1980-2019 period, from 0.46 to 0.74. This increase would have been twice as large absent educational expansion. In other words, education has sufficiently mitigated the rise of within-country inequality to keep overall global inequality constant.[16]

1.3.2.3 Distributional Growth Accounting by Time Period

The structure of the data also allows me to estimate the contribution of education by time period. This analysis delivers two main results.

1) Schooling Gains Have Increasingly Benefited the Global Poor First, the benefits of worldwide improvements in schooling have increasingly accrued to the global poor since 1980. The best way to see this is to compare the distribution of schooling gains over the 1980-2019 and 2000-2019 periods. Figure 1.10 plots the contribution of education to the annual income growth of each global percentile for these two periods. The figure is obtained by comparing the distribution of income in 2019 to its counterfactual distribution absent educational progress, and then annualizing the resulting ratio.[17]

[16]Appendix figure A.3 compares average gains from schooling to gains for the poorest 20% individuals in each country. In nearly all countries, education appears to have increased income more for low-income than high-income earners, mainly because of supply effects redistributing income from high-skilled to low-skilled workers.

[17]Formally, let y be income in 2019 and $\tilde{y}$ counterfactual income. The contribution of education is then equal to $(\frac{y}{\tilde{y}})^{(1/T)} - 1$, with T the number of years corresponding to the period considered (40 years for 1980-2019, 20 years for 2000-2019). Notice that my estimates always rely on comparing a given year to 2019, given that I only have microdata for the latter. This implies that I unfortunately cannot estimate the distributional incidence of educational expansion over the 1980-2000 period, for instance.

Consistently with previous results, schooling gains have been greatest at the middle and bottom of the global income distribution since 1980, generating a 1.2-1.4% annual increase in earnings for these groups, compared to less than 0.8% for all groups within the top decile. In the recent period, growth effects of education have declined overall, but not at the bottom of the world distribution of income. The past decades have thus seen persistently large human capital accumulation for the poor and a relative slowdown of educational progress in the rest of the world. Since 2000, education has been an unambiguous driver of global inequality reduction.

2) Education Explains Over 50% of Global Bottom 20% Growth in All Periods Actual economic growth has also accelerated since 2000, a period characterized by the surge of the Chinese and Indian economies and higher growth rates in most other parts of the world. Appendix table A.1 presents a distributional growth accounting decomposition of the world distribution of income for the 2000-2019 period. Education explains 27% of global average economic growth since 2000, about two times lower than the corresponding figure since 1980. For the global bottom 20%, however, education can still account for about 50% of income gains, because actual growth for this group has not been particularly high, while gains from schooling have remained about the same. The same pattern holds over the 2010-2019 period.[18] All in all, private returns to schooling can always explain at least half of growth for the world's poorest individuals, regardless of the period considered.

1.3.2.4 Distributional Growth Accounting Across Cohorts

The results presented until now focus on changes in educational attainment of the working-age population from 1980 to 2019. This involves comparing education for cohorts born between the 1910s and 1960s with those born between the 1960s and 1990s. The effect of education on economic growth thus results from two separate mechanisms: the replacement of new cohorts by old ones (e.g., 1910s cohorts are observed in 1980 but not in 2019), and the fact that new cohorts are more educated than their elders (e.g., 1990s cohorts are more educated than 1980s cohorts). In this section, I isolate the specific contribution of the second channel, namely, educational progress among new generations that arrived on the labor market from 1980 to 2019. This approach is interesting from a historical and policy point of view, because it allows answering the following question: how much lower would incomes be if cohorts

[18]Appendix figure A.1 plots the share of average growth and global bottom 20% growth that can be explained by education over the 1980-2019, 1990-2019, 2000-2019, and 2010-2019 periods. Appendix figure A.2 represents the distribution of global economic growth together with the contribution of education in 2000-2019, as in figure 1.1.

arriving on the labor market after 1980 had not been more educated than the 1980 cohort? Put simply, it allows capturing the specific contribution of improvements in education among new generations since 1980.

To derive this counterfactual, I use the same methodology as in the rest of the paper, with the only difference that counterfactual educational attainment is that of the 1980 cohort instead of that of the 1980 working-age population. I start by calculating education of the 1980 cohort, estimated as that of individuals aged 60 to 65 in 2019. I then downgrade education levels of each post-1980 cohort until reaching the 1980 counterfactual. Finally, I reduce earnings using returns to schooling and estimate general equilibrium effects as in the rest of the paper.

The main results of this exercise are presented in appendix table A.7. Appendix figures A.4 and A.5 present the corresponding distribution of gains from schooling and share of growth explained by global income percentile. Worldwide educational progress appears to have been particularly progressive when narrowing the focus to post-1980 generations. Generational progress explains about 12% of global economic growth since 1980, but 61% of income gains for the world's poorest 20% individuals, 37% for the global bottom 50%, and 4% for the global top 10%. These results are in line with those of the previous section, which highlighted the increasingly progressive nature of educational progress. Given substantial improvements in school enrollment observed in low-income countries in the 1990s and 2000s (Barro and Lee, 2015), this pattern can only be expected to intensify in the future. In the coming decades of the twenty-first century, education could well become an even stronger force of decreasing global inequality than it already was at the turn of the twentieth century.

1.3.3 Standard Versus Distributional Growth Accounting

Another way of better understanding the results is to compare my estimates to those obtained before and after applying each of the estimation steps outlined in section 1.1. This is useful to isolate the different mechanisms driving the results, from differential returns to schooling to changes in within-country inequality and general equilibrium effects. It also enables comparing my results to those one would obtain from a canonical growth accounting decomposition.

Table 1.6 displays the share of global average economic growth and global bottom 20% income growth that can be explained by education with different data sources and assumptions.

1.3.3.1 The Standard Growth Accounting Decomposition

I start by presenting results from a standard growth accounting decomposition in its simplest form. To do so, I follow the same methodology as Barro and Lee (2015), who exploit cross-country GDP and educational attainment data to estimate the fraction of global economic growth explained by human capital accumulation from 1960 to 2010.[19] This decomposition only requires three ingredients: cross-country per-capita GDP data (taken from the World Inequality Database), capital income shares (taken from the Penn World Tables as in Barro and Lee, 2015), and an estimate of the Mincerian return to schooling (assumed to be 10% per year). Counterfactual income absent educational progress is then calculated as:

$$\tilde{y}^c = v_L \frac{y^c}{(1+r)^{\Delta S}} + v_K y^c \tag{1.22}$$

With y^c GDP per capita in 2019 in country c, v_L and v_K the labor and capital income shares, $\Delta S = S^{2019} - S^{1980}$ the change in average years of schooling of the working-age population, and $r = 0.1$ the return to schooling. Put simply, if average years of schooling increased by one year, then labor income would have been $\frac{1}{1.1} = 0.91$ times lower absent educational progress, while capital income would have remained unchanged.

The first line of table 1.6 presents the results. Mincerian returns to schooling can explain about a third of global average economic growth. The second column shows corresponding results for the global bottom 20%. Because this growth accounting decomposition relies only on cross-country data, the poorest 20% has to be defined as the poorest 20% of countries (population-weighted). Schooling gains have been relatively small for these countries. As a result, with this methodology, education can explain less than a quarter of growth for the global bottom 20%.

1.3.3.2 Adjusting the Labor Income Share

A first problem with this approach is that capital income shares reported in the Penn World Tables treat all mixed income as capital income. The implicit assumption is that wages are the only source of income affected by schooling; there is no return to schooling on mixed income. This does not appear to be true. OLS returns estimated

[19] The data sources and exact steps of the methodology used here differ slightly from those in Barro and Lee (2015). Barro and Lee (2015) do not report results over the 1980-2019 period, but I can compare my results to theirs for the 2001-2010 period. The two estimates are very close: education explains 15.7% of global economic growth according to their estimates, versus 19.7% according to mine. See Barro and Lee (2015), table 4.5.

in section 1.2.3 include mixed income and are, on average, indistinguishable from those estimated on wages only.[20] Aggregate effects of schooling estimated using the natural experiments studied in section 1.2.6 also include mixed income. The labor income share used for the estimation should thus include mixed income, because this is the income concept Mincerian returns to schooling apply to.

The second line of table 1.6 presents the results when mixed income is included in the labor income share. The Penn World Tables do not provide this decomposition, so I turn to the factor income shares estimated by Bachas et al. (2022). This adjustment alone increases the contribution of schooling to average economic growth to 43%, and its contribution to bottom 20% growth to 35%.

1.3.3.3 Incorporating Within-Country Inequality

In a third step, I incorporate within-country inequality in the estimation: the global poorest 20% correspond no more to the poorest 20% of countries. All other methodological ingredients stick to the standard growth accounting exercise. In particular, the average income of each income group is reduced by the same proportion within each country.

By construction, accounting for within-country inequality leaves the share of aggregate economic growth explained unchanged. However, it raises the contribution of education to bottom 20% growth from 35% to 41%, for two main reasons. First, the bottom 20% is now a composite of individuals from low-income and middle-income countries, some of which witnessed significant schooling gains. Second, inequality has risen rapidly in many countries: growth for low-income earners has been lower than average economic growth. This second factor increases the contribution of education because there is less growth among the global poor to be explained than what cross-country data suggest.

1.3.3.4 Incorporating Within-Country Capital Income Concentration

Fourth, I account for the fact that capital income is concentrated at the top of the distribution in each country, as discussed in section 1.2.5.3. For the majority of individuals belonging to the bottom 90%, almost all of income consists in wages or mixed income. The passthrough from schooling to income is thus close to 100% for low-income earners, rather than equal to the labor income share as assumed until now.

[20]See appendix table A.29, which compares returns to schooling by level estimated with and without including mixed income in total personal income.

As in step 3, this does not affect the share of aggregate growth explained by education. However, it raises the contribution of education to bottom 20% growth significantly, from 41% to 51%.

1.3.3.5 Bringing in the Microdata

Fifth, I bring in the micro-database covering education and earnings in 150 countries collected for the purpose of this paper. I then estimate the impact of educational expansion on earnings in each country, using Mincerian returns to schooling estimated by OLS. This introduces three main differences with the previous exercise. First, returns to schooling are allowed to vary by country. Second, returns to schooling are allowed to vary by level in each country. Third, educational expansion is allowed to benefit income groups differentially in each country. For instance, expanding primary education only benefits workers who would otherwise have had no schooling, so it tends to generate more growth at the bottom than at the top of the income distribution (see section 1.1.4.1).

Moving from aggregate data to microdata slightly decreases the contribution of education to average growth, mainly because the average Mincerian return in my data is closer to 9% than 10%. It reduces much more strongly the contribution of education to the growth of the global bottom 20%, which falls from 51% to 41%. The main reason is that the world's poorest individuals have mostly benefited from expansions in primary and secondary education since 1980, whose average yearly returns fall well below 10%. In particular, the expansion of primary education in India has been one of the major transformations occurring during this period, with yearly returns estimated to reach less than 3% in my data (see figure 1.4). Heterogeneity in returns by country and level thus appears to have important effects on growth accounting estimates, implying much lower gains from schooling for the global poor than with a homogeneous 10% return.

1.3.3.6 Accounting for General Equilibrium Effects: Distributional Effects

Sixth, I account for general equilibrium effects redistributing income between skill groups: the expansion of education has increased the supply of skilled workers, thus depressing their wages relative to those of unskilled workers (see section 1.1.4.3). This step of the methodology requires data on the joint distribution of wages and education, so it can only be estimated with the microdata.

To isolate this particular distributional effect, I only adjust relative wages, leaving the average income unchanged in each country. The contribution of education to average growth is therefore the same as in the previous step. As shown in line 6 of table 1.6, general equilibrium effects have strongly benefited the global bottom 20%. For this group, accounting for changes in relative wages raises the contribution of education from 41% to 52%. This large impact is consistent with the empirical literature documenting important effects of changes in the supply of skilled workers on inequality and the distribution of economic growth (e.g., Goldin and Katz, 2007; Khanna, 2023; Moretti, 2004).

1.3.3.7 Accounting for General Equilibrium Effects: Aggregate Effects

Finally, I account for the fact that not expanding education would have been more detrimental to growth than suggested by 2019 returns. The true returns to schooling that should be used are thus higher than the returns observed in 2019 (see section 1.1.4.2).

This final step adjusts returns to schooling differentially by country and level, generating both aggregate and distributional effects. The contribution of education to average economic growth rises significantly, from 40% to 54%. For the bottom 20%, it increases from 52% to 71%, yielding the benchmark estimate presented at the beginning of this section.

For comparison, table 1.6 also reports results in which IV estimates of returns to schooling are used instead of OLS estimates. The contribution of education to average growth remains about the same, while the share of global bottom 20% growth explained rises to 75%.

1.3.3.8 Summary

In summary, the results presented in table 1.6 tell us two key facts on the role of education in reducing global poverty.

First, education explains 3 times more growth for the global poor than a standard growth accounting exercise would suggest. Changes in inequality within countries, capital income concentration, differential returns, and general equilibrium effects together imply a much more complex picture of educational expansion than that depicted by a traditional decomposition relying on aggregate data. Almost all of these additional layers of detail imply a greater contribution of education to global poverty reduction.

Second, accounting for the distributional effects of schooling within countries is of paramount importance for understanding the effect of education on global poverty reduction. Low returns to basic education imply that the contribution of schooling to global bottom 20% income growth is about 20% lower than a constant 10% return would suggest. General equilibrium effects explain over 50% of the contribution of education to global poverty reduction, by redistributing a large share of schooling gains from high-skilled to low-skilled workers.

1.3.4 Robustness

I now discuss the sensitivity of my main finding to alternative specifications and other sources of concern.

1.3.4.1 Alternative Elasticities of Substitution

My benchmark estimates assume an elasticity of substitution between skill groups of 6, drawing on recent estimates derived in the macroeconomics literature on long-run substitutability between skill groups. Appendix table A.4 presents the share of growth explained by education by global income group under two alternative specifications. The low substitutability specification assumes an elasticity of 4, corresponding to the lower bound on long-run substitutability derived in Bils, Kaymak, and Wu (2022). In contrast, I set $\sigma_1 = 5$, $\sigma_2 = 7$, and $\sigma_3 = 9$ in the high substitutability scenario, corresponding to elasticities at the upper bound of those found in the literature (Hendricks and Schoellman, 2023).

The low substitutability scenario implies greater redistributive effects of educational expansion from high- to low-skill workers. As a result, the share of growth explained by education since 1980 rises from 71% to 85% for the global bottom 20%. Conversely, the high substitutability scenario implies weaker general equilibrium effects, but only slightly so: the share of global bottom 20% growth explained remains as high as 66%.

1.3.4.2 Alternative Nesting of the CES Production Function

Another concern is that the results might be sensitive to the way the production function is specified. The specification of CES nests outlined in 1.2.4 is fairly standard and has been successfully used in the empirical literature (Fernández and Messina, 2018; Goldin and Katz, 2007). Yet, one may still be concerned that alternative patterns of imperfect substitutability may yield different distributional effects of educational expansion.

In appendix table A.5, I consider an alternative specification in which firms first pick between workers with below and above secondary education, and then choose between subcategories of workers within these two nests.[21] With this specification, education explains 53% of global economic growth and 65% of growth for the global bottom 20%, slightly less than in the main specification but of the same order of magnitude.

1.3.4.3 Alternative Patterns of Schooling Expansion

Another step of the simulation has to do with identifying who benefits from private returns to schooling. In the main specification, I randomly sample individuals within age-gender cells and downgrade their education levels until matching those observed in 1980 (see section 1.1.4.1). While the results from the three natural experiments studied in section 1.2.6 suggest that this approach does a good job at capturing the distributional incidence of expansions in access to schooling, one might still be concerned about unobserved heterogeneity. For instance, if educational expansion mostly benefited children from disadvantaged socioeconomic backgrounds, one might underestimate benefits for low-income earners (who tend to come from more disadvantaged backgrounds) and overestimate aggregate effects (since individuals from more disadvantaged backgrounds might have lower expected incomes).

Extending the estimation beyond age-gender cells is not possible for the world as a whole, given the lack of data on access to schooling by socioeconomic characteristic since 1980. However, I can investigate the implications of using more refined categories for India (1983-2019), South Africa (2002-2019), and the United States (1980-2019). For India, I use historical waves of the National Sample Survey to estimate variations in educational expansion by state. For South Africa and the United States, I can similarly expand the analysis to cover differential educational progress by race and region (states in the U.S., provinces in South Africa). The results are presented in appendix figures A.6, A.7, and A.8.[22] Using more refined categories turns out to have almost no effect on the results. It marginally reduces gains from schooling for top income earners in India and the United States, while it raises them slightly

[21]More specifically, the production function is: $L = \left(A_L L_L^{\frac{\sigma_1 - 1}{\sigma_1}} + A_H L_H^{\frac{\sigma_1 - 1}{\sigma_1}} \right)^{\frac{\sigma_1}{\sigma_1 - 1}}$, with nests given by $L_L = \left(A_{non} L_{non}^{\frac{\sigma_2 - 1}{\sigma_2}} + A_{pri} L_{pri}^{\frac{\sigma_2 - 1}{\sigma_2}} \right)^{\frac{\sigma_2}{\sigma_2 - 1}}$ and $L_H = \left(A_{sec} L_{sec}^{\frac{\sigma_3 - 1}{\sigma_3}} + A_{ter} L_{ter}^{\frac{\sigma_3 - 1}{\sigma_3}} \right)^{\frac{\sigma_3}{\sigma_3 - 1}}$.

[22]These results correspond to those obtained after going through the first step of the methodology only. That is, I randomly sample individuals, downgrade educational attainment, and reduce their earnings using 2019 returns to schooling. I then plot gains from schooling as the percent difference between actual income and counterfactual income absent educational expansion.

for all income groups in South Africa. I take this as reassuring evidence that the methodology adopted in this paper provides a good first-order approximation of actual distributional effects of schooling.

1.3.4.4 Capital Income Affected by Schooling

In my benchmark estimates, I assume that education has no effect on capital income at all, in line with standard growth accounting decompositions. This is arguably a very conservative assumption: with a constant saving rate, one should expect a fraction of schooling gains to be saved and later received in the form of capital income. There is also evidence that education can have potentially large effects on entrepreneurial income (Gennaioli et al., 2013; Queiró, 2022) or innovation (e.g., Nimier-David, 2023), both of which should translate into capital income gains. Appendix table A.6 presents results in which returns to schooling are assumed to apply to both labor and capital income. This has two effects on distributional growth accounting results: it increases the contribution of education for all groups, and it increases it more for top-income groups within each country, who earn a greater fraction of income from capital. In this scenario, education explains about two-thirds of global economic growth and 80% of income gains for the world's poorest 20% individuals. It still explains a relatively lower fraction of growth at the top of the world distribution of income, about 40% for the top 1%, mainly because of slower educational progress in the United States and because supply effects reduce inequality within countries.

1.3.4.5 Education Quality

A last source of concern is that the quality of education might have changed during the period considered. This implies that changes in educational attainment since 1980 used to derive the counterfactual might not be comparable. If quality has decreased, then changes in years of schooling of the working-age population should not be valued in the same way as if it has remained the same. In the extreme case in which quality has declined sufficiently so as to fully cancel increases in quantity, one might be estimating large benefits of educational expansion even when there have been none.

I investigate trends in education quality and potential implications for the results of this paper at length in appendix A.4. Data on the evolution of quality are scarce, especially in the developing world. Recent test scores suggest that it has stagnated or increased in most countries (Altinok, Angrist, and Patrinos, 2018; Angrist et al., 2021), while long-run trends in literacy by cohort suggest some decline in a subset of developing economies (Le Nestour, Moscoviz, and Sandefur, 2022). All in all, there is little evidence of widespread changes in quality that would alter the main findings of this paper. Even under conservative assumptions on a potential decline in quality, I show that my main results on the world distribution of income are unaffected.

1.4 Education and Global Gender Inequality

This section studies the role of education in shaping the evolution of worldwide gender inequality since 1991. Sections 1.4.1 and 1.4.2 outline the conceptual framework and methodology. Section 1.4.3 presents the main results.

1.4.1 Conceptual Framework

Before moving on to the results, it is useful to conceptually distinguish the different channels through which human capital accumulation can affect gender inequality. Consider an individual i with a level of schooling s_i. They can be employed or inactive with a probability e_i, which depends on schooling: $e_i = e_i(s_i)$. When employed, their wage depends on schooling, which has a return r_i per year. Their expected income y_i is thus:

$$y_i = e_i(s_i) \cdot r_i^{s_i} \tag{1.23}$$

The gender gap is:

$$\begin{aligned} \Delta y &= \ln y_f - \ln y_m & (1.24) \\ &= \ln e_f(s_f) - \ln e_m(s_m) + (s_f \ln r_f - s_m \ln r_m) & (1.25) \end{aligned}$$

Where f denote women and m men. Rewriting returns to schooling for men as $r_m = \alpha r_f$:

$$\Delta y = \underbrace{\underbrace{(s_f - s_m)\ln r_f}_{\text{Differential Educational Expansion}} - \underbrace{s_m \ln \alpha}_{\text{Differential Returns to Schooling}}}_{\text{Intensive Margin}} + \underbrace{\ln e_f(s_f) - \ln e_m(s_m)}_{\text{Extensive Margin}} \tag{1.26}$$

There are three main channels through which education can reduce gender inequality. First, expanding schooling differentially in the favor of women will increase their relative income. This effect will be stronger when the returns to schooling for women are high. Second, holding the distribution of educational attainment constant, the relative income of women will be greater if their returns to schooling are greater than men's. Third, schooling may have an additional impact on the gender gap by differentially affecting the labor force participation of men and women. If schooling increases the propensity of being employed, this will magnify the effect of differential educational expansion on gender inequality. I now present the data and methodology used to estimate the contribution of each of these three channels.

1.4.2 Methodology

To estimate the role of schooling in the reduction of global gender inequality, I apply the same methodology as the one presented in section 1.2. The only difference is that I allow for differential returns by gender, as well as a potential additional effect of schooling on female labor force participation. I thus construct three separate estimates of counterfactual gender inequality.

1) Differential Educational Expansion First, I consider a case in which only differential trends in educational attainment matter. As in section 1.2, I start by reducing education levels by age $\times$ gender cell in each country. I then reduce earnings of both men and women by the same returns to schooling, estimated over the entire working-age population.

Figure 1.11 shows that with the exception of Sub-Saharan Africa, there has been a significant decline in gender schooling inequalities since 1991 in all regions of the world. This reduction has been greatest in China, where the gender gap in years of schooling declined from 1.5 to 0.5. In Europe, Northern America, and Latin America, the gender education gap has reversed, with working-age women now being slightly more educated than men. Convergence in educational attainment by gender throughout the world can thus be expected to have acted as a significant driver of the reduction in gender inequality.

2) Differential Returns to Schooling Second, I incorporate differential returns to schooling conditional on being employed. To do so, I estimate Mincerian returns by gender, decomposed by education level, as in equation 1.16. I then reduce the earnings of men and women by gender-level-specific returns, so as to construct counterfactual earnings absent educational progress since the beginning of the period considered.

Figure 1.12 reproduces a well-known fact: in all regions of the world, returns to schooling are higher for women than for men (e.g., Montenegro and Patrinos, 2021).[23] This gap can be large: women's returns are 2-4 percentage points higher than men's in Latin America, India, and the MENA region. Heterogeneity in returns is thus expected to amplify the inequality-reducing effects of improved access to schooling for women.

3) Extensive Margin Third, I incorporate differential effects on employment. In the benchmark specification, I run the following OLS regression in each country:

$$e_{ic} = \alpha_c + \vartheta_c \text{Educ}_{ic} \cdot \text{Gender}_{ic} + \xi_c \text{Educ}_{ic} + \digamma_c \text{Gender}_{ic} + X_{ic}\beta_c + \varepsilon_{ic} \qquad (1.27)$$

With e_{ic} a dummy for being economically active, Educ_{ic} years of schooling, Gender_{ic} a dummy taking one for women and zero for men, and X_{ic} a vector of controls including interactions between gender and an experience quartic. The coefficient of interest is ϑ_c, capturing the differential effect of increasing education by one year on female labor force participation. This coefficient turns out to be positive in 109 countries and negative in 47 countries. In the average country, the effect of an additional year of schooling on employment is 0.7 points greater for women: increasing the average education of women by one year relative to that of men increases female labor force participation by 0.7 percentage points.

In alternative specifications, I rely on quasi-experimental evidence from a number of recent studies; their results are presented in appendix table A.13. For instance, Elsayed and Shirshikova (2023) find that the staggered construction of public universities in Egypt increased female labor force participation (FLFP) for women but not for men, with an implied effect of 8 percentage points per additional year of schooling. Similarly, Keats (2018) finds that schooling increases FLPF by 7 percentage points in Uganda, using differential exposure to the 1997 elimination of primary school fees as an instrument. Two recent studies on Indonesia (Akresh, Halim, and Kleemans, 2023) and Pakistan (Khan, 2021), however, find no effect of increased access to schooling on female labor force participation. Overall, the different studies listed in table A.13 find widely varying effects, ranging from 0 to 10 percentage points, with an average of about 6. Expressed as a percent increase relative to baseline, this corresponds to an average increase in FLFP of 19% per year of schooling (typically ranging from 15% to 30%).

[23] Appendix figure A.10 extends this comparison to all countries in the dataset. Returns are higher for women in nearly all countries in the world. Of course, one may naturally be worried about selection bias. Unfortunately, the literature comparing OLS and IV estimates of differential returns to schooling by gender is scarce, although some of the studies outlined in table A.31 do find comparable or higher gender gaps in returns in IV specifications.

4) Distributional Growth Accounting by Gender The final step is to estimate the contribution of education to global gender inequality reduction. This requires data on the evolution of gender inequality in each country. I rely on recent work by Neef and Robilliard (2021), who combine various sources to build the first database on the evolution of gender labor income inequality in nearly all countries in the world from 1991 to 2019. The indicator of interest is the female labor income share, defined as:

$$y_L^f(s_f, s_m) = \frac{N_f \cdot e_f(s_f) \cdot r_f^{s_f}}{\sum_{i \in \{f,m\}} N_i \cdot e_i(s_i) \cdot r_i^{s_i}} \tag{1.28}$$

With N_i the share of the working-age population of gender i. This indicator corresponds to the total share of labor income accruing to women, coding labor income as zero for individuals who are out of the labor force. To estimate the contribution of education to global gender inequality reduction, I thus construct counterfactual female labor income shares $\tilde{y}_L^f(\tilde{s}_f, \tilde{s}_m)$ using the methodology outlined above. I then compare the actual evolution of gender inequality to its evolution absent educational progress since 1991.

1.4.3 Education and Global Gender Inequality, 1991-2019

1.4.3.1 Education and the Global Female Labor Income Share

I start by presenting results on the contribution of schooling to gender inequality reduction from a global perspective. Table 1.7 compares the evolution of the global female labor income share since 1991 to its counterfactual evolution absent educational expansion. The global female labor share is calculated by taking the ratio of total labor income received by women to total labor income received by both men and women in the world as a whole.

1) Education Explains 50-80% of Global Gender Inequality Reduction There has been a decline in global gender inequality, albeit small: women received about 29% of labor income in 1991, compared to 32% in 2019.[24] The second row estimates by how much lower the female labor income share would have been if the distribution of educational attainment had remained unchanged, assuming returns to schooling are the same for men and women and no differential effect of schooling on employment. This compositional factor alone explains about half of global gender

[24] Appendix figure A.9 compares the female labor income share in 1991 and 2019 in all countries. There has been a decline in gender inequality in 140 countries out of the 174 covered by the data. China stands out as the only large country where the female labor income share has declined.

inequality reduction: the female labor income share would have increased by 1.3 points instead of the 2.8 observed. The third row incorporates heterogeneous returns to schooling by gender. Because returns to schooling are higher for women than for men in nearly all countries, this raises the contribution of education to over two-thirds. The last row incorporates the effect of education on female labor force participation. By this last measure, education accounts for approximately 80% of global gender inequality reduction since 1991.

While the magnitude of these effects might seem surprising at first, one should not forget that the dynamics of the global female labor share depend on both inequality within and between countries. Even in a world with stable gender income gaps in each country, global gender inequality can still decrease if aggregate economic growth is greater in countries with higher initial female labor income shares. The results presented in table 1.7 are thus capturing two separate effects of education: the effects of differential educational expansion by gender within countries, and the fact that educational attainment rose particularly rapidly in countries with lower gender inequality to begin with (such as China).

2) Education Explains 50-60% of Gender Inequality Reduction in the Average Country For the study of gender inequality, one might be more interested in understanding the particular role that education has played in reducing gender income gaps within countries. The last column of table 1.7 isolates this channel by presenting the average share of gender inequality reduction explained by education. To construct this indicator, I calculate the evolution of actual and counterfactual female labor income shares in each country, together with the share of gender inequality reduction explained.[25] I then take the population-weighted average of this indicator over all 150 countries covered by the data. From this angle, differential educational expansion alone accounts for 50% of gender inequality reduction in the average country. Incorporating heterogeneous returns brings this share to 58%, while accounting for employment effects raises it to almost 60%.

Robustness I investigate the sensitivity of these two main results to alternative assumptions and sample restrictions in the appendix. First, one might be worried that these two sets of results are driven by countries where gender inequality increased. In these countries, by construction, education can explain over 100% of reductions

[25]In countries where gender inequality increased and education mitigated this increase, I bound the share of gender inequality reduction explained by education at 100%. In countries where education increased gender inequality, I set the share of gender inequality reduction explained by education at 0%.

in gender inequality. China stands out as a potentially important source of concern, given its large population. Appendix table A.14 reproduces table 1.7 after excluding China from the analysis: the results are similar. Appendix table A.15 considers a more narrow restriction of the sample, excluding all 28 countries in which gender inequality declined, while still keeping countries in which education explains none of gender inequality reduction (because gender schooling gaps stagnated or increased). The contribution of education in the average country declines to about 40%.

Second, it is useful to investigate how sensitive are these results to assumptions on the effect of education on female labor force participation. Appendix table A.16 compares the benchmark estimates to alternative specifications of employment effects commensurable to those found in the existing literature (see appendix table A.13). I consider six specifications, in which an additional year of schooling increases employment either in absolute terms (by 4, 6, or 8 percentage points) or in relative terms (by 15%, 20%, or 25% relative to baseline employment rates). The share of global gender inequality reduction explained exceeds 90% under all these specifications, while results for the average country exceed 65%.

Taken together, these findings suggest that education has been one of the most important drivers of improvements in gender labor income inequality worldwide since the 1990s. Given limited causal evidence on employment effects and differential returns to schooling by gender in each country, the results should be considered as somewhat more uncertain than those on global poverty. However, under reasonable assumptions, education accounts for at least 50% of the rise in the share of labor income accruing to women, and potentially as much as 80-90%.

1.4.3.2 Education and Gender Inequality by World Region and Time Period

1) Schooling Has Reduced Gender Inequality in All Regions of the World Figure 1.13 extends this analysis to the evolution of gender inequality in different regions of the world.[26] Education accounts for more than 40% of declining gender income gaps in all regions, with estimates ranging from a bit below 45% in Latin

[26]Appendix figure A.11 plots income gains for women relative to men, calculated by taking the ratio of actual to counterfactual incomes by gender in each country. Educational expansion alone has generated about 20% to 60% more growth for women than for men in all regions. Accounting for heterogeneous returns and employment effects brings this ratio to 50-150%. Appendix figure A.12 plots annualized income gains from schooling for men and women in each country. In nearly all countries, schooling has generated more growth for women than for men. Accordingly, appendix figure A.13 shows that the female labor income share would be significantly lower in nearly all countries in 2019 if there had been no educational expansion since 1991.

America to over 70% in the MENA region. Differential educational expansion alone can explain 20% to 50% of gender inequality reduction in all regions.

2) The Effect of Educational Expansion on Gender Inequality Has Increased Finally, I present results by world region and for the world as a whole for different time periods in appendix table A.17. Two results stand out. First, the effect of education on gender inequality has increased over time. Income gains from schooling have been about 2 times greater for women since 1991 in the average country, compared to 2.5 since 2010. This rising convergence is visible in most regions, although it has been most pronounced in China and India. Second, education explains over 40% of gender inequality reduction in nearly all regions and time periods. This indicator does not display any clear time trend, given that actual gender inequality reduction has accelerated at the same time as schooling gains. Since 1991, education steadily explains about 60% of reductions in the gender income gap in the average country.

1.5 Discussion

This section presents a general discussion in three directions. Section 1.5.1 investigates what might have been the various factors explaining educational progress since 1980 and implications for the results presented in this paper. Section 1.5.2 discusses complementarities between education and technology and provides an empirical analysis of the role of skill-biased technical change in magnifying the growth effects of education. Section 1.5.3 draws on results from a companion paper to estimate the total contribution of public policies to global poverty reduction, combining direct redistribution and indirect investment benefits from education.

1.5.1 Where Does Schooling Come From?

In developing distributional growth *accounting*, I have done nothing else than to estimate how much lower would incomes be had education not improved, holding constant all other characteristics of the economic environment. An objection to this approach is that education might be determined by other proximate drivers of development (e.g., Hsieh and Klenow, 2010). In particular, if skill-biased technical change is a key determinant of educational expansion, separating the contribution of education loses a lot of its interest. A policymaker interested in enhancing economic growth should focus on technological progress: schooling would naturally follow. Estimating the overall contribution of technical change to human capital accumulation (and vice versa) goes far beyond this paper. In this section, I instead provide suggestive evidence that technology is unlikely to have been the dominant driver of global human capital accumulation since 1980.

1) The Global Poor Overwhelmingly Rely on Public Schools A first fact to keep in mind is that the overwhelming majority of the world's poor children are enrolled in public schools. In the 1970s and 1980s, corresponding to the period of interest to this paper, over 90% of worldwide primary school enrollment was public (World Bank, 2023). Most of the remaining 10% corresponded to children from high-income households within each country, probably putting the contribution of public schools to improved access to schooling for the global poor at over 95%.[27] In a world where governments have and continue to be the primary providers of education for poor children, the idea that market forces are the main force behind human capital accumulation since 1980 seems difficult to sustain.

2) Technological Progress Does Not Necessarily Lead to More Schooling If schooling was primarily determined by economic incentives, then one should also expect skill-biased technical change and school enrollment to closely follow each other. Yet, there are clear examples of disconnections between the two. Perhaps the most illuminating of these is the recent history of the United States, where skill-biased technical change has been exceptionally pronounced at the same time as educational progress has been among the slowest in the world. The result has been an enormous rise in wage inequality, with little indication of any endogenous adjustment in the supply of skilled workers (Autor, Goldin, and Katz, 2020; Goldin and Katz, 2008). This is not to say that market incentives do not play any role at all: there is plenty of evidence that they do.[28] The main argument is that the counterfactual world with no educational progress studied in this paper is not an impossible world. There are concrete historical examples of disconnections between education and technology that justify the use of growth accounting as a useful tool.

[27]India, where private schools have often been credited as the key driver of schooling expansion in recent decades, represents a particularly interesting example. Survey microdata covering the 1986-2017 period, which I have collected for a companion paper (Gethin, 2023b), provide information on school attendance by type of school. There has been a large increase in the share of children enrolled in private schools over time, but this increase has been entirely driven by children from middle- and high-income households. About 85% of children coming from the poorest 20% of households were enrolled in public schools in 2017, almost the exact same share as in 1986.

[28]See for instance Atkin (2016), Blanchard and Olneyb (2017), Foster and Rosenzweig (1996), Li (2018), and Oster and Steinberg (2013).

3) An Empirical Investigation As a last piece of complementary evidence, I investigate correlates of schooling across countries and over time. I collect data on three complementary indicators of access to schooling: expected years of schooling (corresponding to the number of years a child can hope to stay in school), net primary school enrollment, and net secondary school enrollment. I then regress these indicators on selected variables capturing different dimensions of the economic environment, from the role of government (public education spending, government effectiveness) to trade (trade-to-GDP ratio) and technology (internet usage, mobile cellular subscriptions, and the skill bias of technology estimated from the surveys compiled in this paper).

Table A.8 runs this regression across countries, for the last year available. Table A.9 extends this analysis to the 1980-2019 period with country and year fixed effects. The takeaway is that public education expenditure stands out as the only variable robustly correlated with schooling in both cross-sectional and panel data. The skill bias of technology is not significantly related to schooling across countries, and if anything enters the regression with the wrong sign. There is also no clear evidence that countries with greater economic growth or faster adoptions of new technologies have had larger increases in schooling since 1980 (if anything, the opposite seems to be true). In contrast, real public education expenditure per child is strongly associated with improvements in access to schooling across all three indicators. These results should of course not be interpreted causally, but they suggest that technological progress is unlikely to have been behind the major schooling expansions of the past decades studied in this paper.

1.5.2 Education and Skill-Biased Technical Change

The main results of this paper are based on building counterfactual income distributions in 2019, bringing education levels back to their 1980 value. This amounts to answering the following question: how much lower would incomes be if education had not improved, holding all other factors to their 2019 values? An alternative way of estimating gains from schooling would instead be to use surveys from the 1980s to estimate the effect of increasing education levels to their 2019 value. This amounts to answering a slightly different question: how much higher would incomes be if education had improved, holding all other factors to their 1980 values? If we stick to the CES production function with labor-augmenting technology terms, the difference between these two estimates turns out to provide useful insights into the role of skill-biased technical change in amplifying the growth effects of schooling.

1.5.2.1 Backward Accounting: Schooling With Relative Efficiency Gains

The main results of this paper correspond to what one might call "backward accounting," bringing education levels back to their 1980 value. Formally, let $A^t = \frac{A_H^t}{A_L^t}$ be the skill bias of technology (the relative efficiency of skilled workers) and L^t the vector of skill supplies at time t. Output is a function of both: $Y^t = Y(A^t, L^t)$. The main exercise of this paper amounts to constructing:

$$\tilde{Y}^{2019}_{\text{backward}} = Y(A^{2019}, L^{1980}) \tag{1.29}$$

The share of growth explained by education is:

$$\text{Success}_{\text{backward}} = 1 - \frac{\frac{Y(A^{2019},L^{1980})}{Y(A^{1980},L^{1980})} - 1}{\frac{Y(A^{2019},L^{2019})}{Y(A^{1980},L^{1980})} - 1} \tag{1.30}$$

The estimation thus amounts to comparing growth rates with and without educational expansion, *given the actual evolution of the relative efficiency of skilled workers*, from A^{1980} to A^{2019}. In this model, not expanding education has large negative effects on growth and inequality. In particular, supply effects magnify the growth effects of education, because the loss in income from not expanding education would have been greater than what 2019 returns suggest.

1.5.2.2 Forward Accounting: Schooling Without Relative Efficiency Gains

An alternative would be to work with "forward accounting," taking incomes in 1980 and estimating the growth effects of moving to 2019 education levels. This amounts to constructing:

$$\tilde{Y}^{2019}_{\text{backward}} = Y(A^{1980}, L^{2019}) \tag{1.31}$$

The share of growth explained by education is:

$$\text{Success}_{\text{forward}} = \frac{\frac{Y(A^{1980},L^{2019})}{Y(A^{1980},L^{1980})} - 1}{\frac{Y(A^{2019},L^{2019})}{Y(A^{1980},L^{1980})} - 1} \tag{1.32}$$

The interpretation is now different: it amounts to comparing actual growth to a world in which only education would have increased, *while technology would have remained to its 1980 value.* Supply effects now decrease the contribution of education to aggregate growth: the return to schooling declines as education expands when the skill bias of technology does not.

If relative efficiency A has not changed, then the two estimates should be identical. Indeed, because the supply of skilled workers was lower in 1980, the return to schooling should be much higher in 1980 than in 2019. Aggregate gains from schooling estimated by increasing incomes in 1980 (using 1980 returns diminished by supply effects) should then be equal to those estimated by reducing incomes in 2019 (using 2019 returns augmented by supply effects).

If relative efficiency has changed, however, the two estimates will differ. In the presence of skill-biased technical change ($A^{2019} > A^{1980}$), in particular, we should expect backward accounting to deliver greater growth effects of education: $\text{Success}_{\text{backward}} > \text{Success}_{\text{forward}}$. The reason is intuitive: in a world with rising relative efficiency of skilled workers, expanding schooling turns out to be a much more profitable investment than in a world where it remains unchanged. Comparing estimates of backward and forward growth accounting across countries can then tell us something about the role of skill-biased technical change in enhancing the growth effects of educational expansion.

1.5.2.3 An Empirical Investigation

Unfortunately, 1980s survey data are available for almost none of the countries covered by this paper. However, I was able to find and harmonize comparable surveys covering personal income and education in the early 2000s for 33 countries: 14 European countries, 14 Latin American countries, the United States, Indonesia, Thailand, Ghana, and South Africa. I can thus investigate how estimates from backward and forward accounting differ across these countries for the 2000-2019 period. Surveys fielded at the beginning and end of the period are not always perfectly comparable, so the results should be interpreted with some care, but they can still provide suggestive evidence.

I estimate backward accounting results using the same methodology as in the rest of the paper. For forward accounting, I take 2000 surveys and increase education levels to their 2019 values in each country. I then estimate returns to schooling in 2000 and reduce them to reach true returns to schooling, using the exact opposite step as the one outlined in section 1.1.2. Finally, I estimate distributional effects by adjusting relative wages as in the rest of the paper.

1.5.2.4 Results

Appendix table A.10 compares estimated income gains from schooling under backward (with efficiency gains) and forward (without efficiency gains) accounting, separately for the population as a whole and the bottom 50%.

Gains from schooling are almost always lower under forward than backward growth accounting. This is consistent with significant increases in the relative efficiency of skilled workers, which have made the benefits of expanding access to schooling much larger than an analysis holding technology fixed would suggest (which is what forward accounting does). The gap between the two estimates can be particularly large for the bottom 50% of earners. Indeed, skill-biased technical change has both aggregate and distributional effects. When the skill bias of technology is rising rapidly, expanding access to schooling becomes a particularly powerful tool for ensuring that the benefits of technological progress are broadly shared. This is just an example of the "race between education and technology" (Goldin and Katz, 2008).

There are significant variations across countries. Europe stands out as the region where the gap between the two estimates is the largest: for the bottom 50%, they differ by a factor of 3. This is consistent with the fact that skill-biased technical change has been particularly pronounced in high-income countries in recent decades (this does not stand out in the case of the United States mainly because educational progress was particularly small during this period). In contrast, forward and backward accounting yield almost identical results in Brazil and Mexico, in line with recent evidence pointing to stagnating or even declining demand for skilled labor in these countries (Fernández and Messina, 2018).

Appendix table A.11 extends this analysis to the share of growth explained by education. Forward accounting typically explains 25-30% less growth than backward accounting, with significant variations across countries. For the bottom 50%, forward accounting still explains an important fraction of growth in most countries covered in this analysis, including about 60% in Europe, 50% in Indonesia, and over 100% in Brazil and Mexico. Even absent skill-biased technical change, expanding education would still have had significant effects on inequality and growth for the poorest individuals in these regions of the world.

In summary, education and labor-augmenting technology act as strong complements. Skill-biased technology without education will generate large increases in inequality with little growth for low-income earners. A perfect example is the trajectory of the United States since 1980, where pretax incomes have literally stagnated for the bottom 50% of earners despite strong macroeconomic growth (Piketty, Saez, and Zucman, 2018). Schooling expansions without technology will reduce inequality but will display lower and decreasing returns (as was perhaps visible in Latin America in the 2000s). With labor-augmenting technology and imperfect substitution, the classic separation between education and total factor productivity becomes more complex and non-additive.[29] Arguably, the historical contribution of education to economic growth should be evaluated in light of how skill-biased technical change has actually evolved. This is what I have attempted to do in this paper.

1.5.3 An Estimate of the Total Contribution of Public Policies to Global Poverty Reduction

I conclude this paper with a combined analysis of direct effects of government transfers on poverty and indirect effects of education on pretax incomes. Public policies contribute to reducing poverty through two channels at a given point in time. First, individuals benefit from cash transfers and in-kind transfers that increase their posttax incomes. In a companion paper, I show that such direct government redistribution can account for 30% of global poverty reduction since 1980 (Gethin, 2023b). Second, public policies can also contribute to increasing future pretax earnings; studying this indirect contribution was the objective of this paper, with a focus on education. Putting these two estimates together can then give us an approximate estimate of the total contribution of public policies to global poverty reduction. This estimate is arguably partial, given that public education is not the only type of policy contributing to pretax earnings growth, but it can at least be viewed as a lower bound.

Table 1.8 compares the evolution of global poverty, the average income of the world's poorest 20%, and the average income of the world's poorest 50% before and after accounting for direct redistribution and indirect benefits from education. Absent educational expansion since 1980, poverty would have declined by about 32% in terms of pretax income, compared to an actual reduction of 55%. Removing taxes from individual incomes and adding government transfers yields an estimate of global poverty in terms of posttax income, which declined by 70%. By this measure, direct redistribution and indirect benefits from education together account for about 54% (1 - $\frac{32}{70}$) of global poverty reduction since 1980. Using a similar reasoning, public policies account for about 80% of real income gains for the world's poorest 20% individuals, and about two-thirds of gains for the poorest half of the world's population.

[29] Caselli and Ciccone (2019) and Jones (2019) provide an interesting discussion along these lines.

As a robustness check, I reproduce this analysis using World Bank data instead of data from the World Inequality Database. The results are presented in appendix table A.12. With this data source, public policies account for about 46% of global poverty reduction since 1980, slightly less than with the WID data. This is not surprising, given that poverty at this threshold already declined by almost 80% in terms of pretax income, which mechanically puts an upper bound on the role played by direct redistribution. Public policies explain 70% of income gains for the global bottom 20% and bottom 50%, corresponding to the same orders of magnitude as with the WID data. The takeaway is that public policies can account for 50-80% of the reduction of global poverty since 1980, and potentially much more if we were to account for the indirect growth effects of public healthcare, transport infrastructure, housing policies, public order and safety, and increasing investments in other public goods observed in the past decades.

1.6 Conclusion

This article represented a first attempt at estimating the role played by education in the historical reduction of global poverty. Combining a stylized model of education and the wage structure with tools borrowed from growth accounting, I proposed a "distributional growth accounting" framework identifying the contribution of education to real income growth within and across countries. Under conservative assumptions, private returns to schooling can explain a large fraction of real income gains among the world's poorest individuals, in the order of 60-70% and potentially more. It can also account for over half of the rise in the share of labor income accruing to women. This puts public education policies at the center of the remarkable reduction of poverty and gender inequality observed in the past decades.

The focus of this article was on private returns to schooling, yet much remains to be done on other dimensions of human capital such as work experience, indirect effects of education on technology, and human capital externalities. How has educational expansion contributed to innovation and its diffusion worldwide? To what extent has population ageing affected economic growth through returns to experience, and how do these effects vary across skill groups in developed and developing economies? Understanding these important economic questions would allow for a fuller understanding of the role played by human capital in shaping global inequalities.

More generally, the results presented in this article call for further research on the structural drivers of changes in the world distribution of income since 1980. The approach adopted in this paper could be extended to other key transformations of the past decades, such as trade globalization, structural change, financialization, and even democratization and changing gender norms. The microeconomics literature provides ample and growing empirical evidence on the economic effects of these factors in specific contexts. Combined with the microdata collected in this paper, additional data collection efforts, and adequate theoretical frameworks, this evidence could be aggregated to shed light on the role played by these long-run processes in the reduction of global poverty and inequality.

Bibliography

Acemoglu, Daron (1998). “Why do new technologies complement skills? Directed technical change and wage inequality”. In: *The Quarterly Journal of Economics* 113.1, pp. 1055–1089.

Acemoglu, Daron and Joshua Angrist (2000). “How Large Are Human-Capital Externalities? Evidence from Compulsory Schooling Laws”. In: *NBER Macroeconomics Annual* 15.1, pp. 9–59.

Acemoglu, Daron and David Autor (2011). “Skills, Tasks and Technologies: Implications for Employment and Earnings”. In: *Handbook of Labor Economics, Volume 4B*. Ed. by Orley Ashenfelter and David Card. North Holland, pp. 1043–1171.

Akresh, Richard, Daniel Halim, and Marieke Kleemans (2023). “Long-Term and Intergenerational Effects of Education: Evidence from School Construction in Indonesia”. In: *The Economic Journal* 133.650, pp. 582–612.

Altinok, Nadir, Noam Angrist, and Harry A. Patrinos (2018). “Global Data Set on Education Quality (1965-2015)”. World Bank Policy Research Working Paper No. 8314.

Angrist, Noam et al. (2021). “Measuring human capital using global learning data”. In: *Nature* 592.1, pp. 403–408.

Atkin, David (2016). “Endogenous Skill Acquisition and Export Manufacturing in Mexico”. In: *American Economic Review* 106.8, 2046–2085.

Autor, David, Claudia Goldin, and Lawrence F. Katz (2020). “Extending the Race between Education and Technology”. In: *AEA Papers and Proceedings* 110.1, pp. 347–351.

Bachas, Pierre et al. (2022). “Globalization and Factor Income Taxation”. NBER Working Paper 29819.

Barro, Robert J. and Jong-Wha Lee (2013). “A new data set of educational attainment in the world, 1950–2010”. In: *Journal of Development Economics* 104.1, pp. 184–198.

Barro, Robert J. and Jong-Wha Lee (2015). *Education Matters: Global Schoolig Gains from the 19th Century to the 21st Century*. Oxford University Press.

Bick, Alexander, Nicola Fuchs-Schündeln, and David Lagakos (2018). "How Do Hours Worked Vary With Income? Cross-Country Evidence and Implications". In: *American Economic Review* 108.1, pp. 170–199.

Bils, Mark, Baris Kaymak, and Kai-Jie Wu (2022). "Labor Substitutability Among Schooling Groups". NBER Working Paper.

Blanchard, Emily J. and William W. Olneyb (2017). "Globalization and human capital investment: Export composition drives educational attainment". In: *Journal of International Economics* 106.1, pp. 165–183.

Blanchet, Thomas, Lucas Chancel, and Amory Gethin (2022). "Why Is Europe More Equal Than the United States?" In: *American Economic Journal: Applied Economics* 14.4, pp. 1–40.

Blanchet, Thomas, Ignacio Flores, and Marc Morgan (2022). "The weight of the rich: improving surveys using tax data". In: *The Journal of Economic Inequality* 20.1, pp. 119–150.

Buera, Francisco J., Joseph P. Kaboski, and Robert M. Townsend (2023). "From Micro to Macro Development". In: *Journal of Economic Literature* 61.2, pp. 471–503.

Card, David (1999). "The Causal Effect of Education on Earnings". In: *Handbook of Labor Economics*. Ed. by Orley C. Ashenfelter and David Card. Elsevier, pp. 1801–1863.

– (2001). "Estimating the Return to Schooling: Progress on Some Persistent Econometric Problems". In: *Econometrica* 69.5, pp. 1127–1160.

Caselli, Francesco and Antonio Ciccone (2013). "The contribution of schooling in development accounting: Results from a nonparametric upper bound". In: *Journal of Development Economics* 104.1, pp. 199–211.

– (2019). "The Human Capital Stock: A Generalized Approach: Comment". In: *American Economic Review* 109.3, pp. 1155–1174.

Caselli, Francesco and Wilbur John II Coleman (2006). "The World Technology Frontier". In: *American Economic Review* 96.3, pp. 499–522.

Chancel, Lucas and Thomas Piketty (2021). "Global Income Inequality, 1820–2020: the Persistence and Mutation of Extreme Inequality". In: *Journal of the European Economic Association* 19.6, 3025–3062.

Chancel, Lucas et al. (2022b). *World Inequality Report 2022*. Belknap Press of Harvard University Press.

Chatterjee, Aroop, Léo Czajka, and Amory Gethin (2022). "Wealth Inequality in South Africa, 1993-2017". In: *World Bank Economic Review* 36.1, pp. 19–36.

Chauvin, Juan Pablo et al. (2017). "What is different about urbanization in rich and poor countries? Cities in Brazil, China, India and the United States". In: *Journal of Urban Economics* 98.1, pp. 17–49.

Che, Yi and Lei Zhang (2018). "Human Capital, Technology Adoption and Firm Performance: Impacts of China's Higher Education Expansion in the Late 1990s". In: *The Economic Journal* 128.614, pp. 2282–2320.

Chen, Shaohua and Martin Ravallion (2010). "The Developing World Is Poorer Than We Thought, But No Less Successful In the Fight Against Poverty". In: *Quarterly Journal of Economics* 125.4, pp. 1577–1625.

Ciccone, Antonio and Giovanni Peri (2006). "Identifying Human-Capital Externalities: Theory with Applications". In: *Review of Economic Studies* 73.1, pp. 381–412.

Clay, Karen, Jeff Lingwall, and Melvin Jr. Stephens (2021). "Laws, educational outcomes, and returns to schooling evidence from the first wave of U.S. state compulsory attendance laws". In: *Labour Economics* 68.1, pp. 1019–1035.

Collin, Matthew and David Weil (2020). "The Effect of Increasing Human Capital Investment on Economic Growth and Poverty: A Simulation Exercise". In: *Journal of Human Capital* 14.1, pp. 43–83.

De Rosa, Mauricio, Ignacio Flores, and Marc Morgan (2022a). "More Unequal or Not as Rich? Revisiting the Latin American exception". Stone Center Working Paper.

Deming, David J. (2022). "Four Facts about Human Capital". In: *Journal of Economic Perspectives* 36.3, pp. 75–102.

– (2023). "Multidimensional Human Capital and the Wage Structure". NBER Working Paper.

Duflo, Esther (2001). "Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment". In: *American Economic Review* 91.4, pp. 795–813.

– (2004). "The medium run effects of educational expansion: evidence from a large school construction program in Indonesia". In: *Journal of Development Economics* 74.1, pp. 163–197.

Ekhator-Mobayode, Uche E. and Johannes Hoogeveen (2022). "Microdata collection and openness in the Middle East and North Africa". In: *Data and Policy* 4.31, pp. 1–17.

Elsayed, Ahmed and Alina Shirshikova (2023). "The women-empowering effect of higher education". In: *Journal of Development Economics* 163.1, pp. 1031–1041.

Fernández, Manuel and Julián Messina (2018). "Skill premium, labor supply, and changes in the structure of wages in Latin America". In: *Journal of Development Economics* 135.1, pp. 555–573.

Foster, Andrew D. and Mark R. Rosenzweig (1996). "Technical Change and Human-Capital Returns and Investments: Evidence from the Green Revolution". In: *American Economic Review* 86.4, pp. 931–953.

Gennaioli, Nicola et al. (2013). "Human Capital and Regional Development". In: *The Quarterly Journal of Economics* 128.1, 105–164.

Gethin, Amory (2023b). "Revisiting Global Poverty Reduction: Public Goods and the World Distribution of Income, 1980-2022". Working Paper.

Glaeser, Edward and Ming Lu (2018). "Human-Capital Externalities in China". NBER Working Paper.

Goldin, Claudia and Lawrence F. Katz (2007). "The Race between Education and Technology: The Evolution of U.S. Educational Wage Differentials, 1890 to 2005". NBER Working Paper.

– (2008). *The Race Between Education and Technology*. Harvard University Press.

Gunderson, Morley and Philip Oreopolous (2020). "Returns to education in developed countries". In: *The Economics of Education*. Ed. by Steve Bradley and Colin Green. Academic Press, pp. 39–51.

Guo, Junjie, Nicolas Roys, and Ananth Seshadri (2018). "Estimating Aggregate Human Capital Externalities". Working Paper.

Hall, Robert E. and Charles I. Jones (1999). "Why Do Some Countries Produce So Much More Output Per Worker Than Others?" In: *The Quarterly Journal of Economics* 114.1, pp. 83–116.

Hammar, Olle and Daniel Waldenström (2020). "Global Earnings Inequality, 1970-2018". In: *The Economic Journal* 130.632, pp. 2526–2545.

Hanushek, Eric A. and Dennis D. Kimko (2000). "Schooling, Labor-Force Quality, and the Growth of Nations". In: *American Economic Review* 90.5, pp. 1184–1208.

Heckman, James J., John Eric Humphries, and Gregory Veramendi (2018). "Returns to Education: The Causal Effects of Education on Earnings, Health, and Smoking". In: *Journal of Political Economy* 126.1, pp. 197–246.

Hendricks, Lutz and Todd Schoellman (2018). "Human Capital and Development Accounting: New Evidence from Wage Gains at Migration". In: *The Quarterly Journal of Economics* 133.2, 665–700.

– (2023). "Skilled Labor Productivity and Cross-Country Income Differences". In: *American Economic Journal: Macroeconomics* 15.1, pp. 240–268.

Hershbein, Brad, Melissa S. Kearney, and Luke W. Pardue (2020). “College Attainment, Income Inequality, and Economic Security: A Simulation Exercise”. In: *AEA Papers and Proceedings* 110.1, pp. 352–355.

Hsieh, Chang-Tai and Peter J. Klenow (2010). “Development Accounting”. In: *American Economic Journal: Macroeconomics* 2.1, pp. 207–223.

Iranzo, Susana and Giovanni Peri (2009). “Schooling Externalities, Technology, and Productivity: Theory and Evidence from U.S. States”. In: *The Review of Economics and Statistics* 91.2, pp. 420–431.

Jones, Benjamin F. (2014). “The Human Capital Stock: A Generalized Approach”. In: *American Economic Review* 104.11, 3752–3777.

– (2019). “The Human Capital Stock: A Generalized Approach: Reply”. In: *American Economic Review* 109.3, pp. 1175–1195.

Katz, Lawrence F. and Kevin M. Murphy (1992). “Changes in Relative Wages, 1963-1987: Supply and Demand Factors”. In: *The Quarterly Journal of Economics* 107.1, pp. 35–78.

Keats, Anthony (2018). “Women’s schooling, fertility, and child health outcomes: Evidence from Uganda’s free primary education program”. In: *Journal of Development Economics* 135.1, pp. 142–159.

Khan, Salman (2021). “Unmet Expectations: The Impacts of School Construction on Female Outcomes in Rural Punjab, Pakistan”. Working Paper.

Khanna, Gaurav (2023). “Large-Scale Education Reform in General Equilibrium: Regression Discontinuity Evidence from India”. In: *Journal of Political Economy* 131.2, pp. 549–591.

Kleven, Henrik and Camille Landais (2017). “Gender Inequality and Economic Development: Fertility, Education, and Norms”. In: *Economica* 84.334, pp. 180–209.

Lakner, Christoph and Branko Milanovic (2016). “Global Income Distribution: From the Fall of the Berlin Wall to the Great Recession”. In: *The World Bank Economic Review* 30.2, pp. 203–232.

Le Nestour, Alexis, Laura Moscoviz, and Justin Sandefur (2022). “The Long-Run Decline of Education Quality in the Developing World”. Working Paper.

Lemieux, Thomas (2006). “The ”Mincer Equation” Thirty Years After Schooling, Experience, and Earnings”. In: *Jacob Mincer. A Pioneer of Modern Labor Economics.* Ed. by Shoshana Grossbard. Springer. Chap. 11, 127–145.

Li, Bingjing (2018). “Export expansion, skill acquisition and industry specialization: evidence from China”. In: *Journal of International Economics* 114.1, pp. 346–361.

Mankiw, Gregory N., David Romer, and David N. Weil (1992). “A Contribution to the Empirics of Economic Growth”. In: *Quarterly Journal of Economics* 107.2, pp. 407–437.

Montenegro, Claudio E. and Harry Anthony Patrinos (2021). “A data set of comparable estimates of the private rate of return to schooling in the world, 1970–2014”. In: *International Journal of Manpower* 1.1, pp. 1–20.

Moretti, Enrico (2004). “Estimating the social return to higher education: evidence from longitudinal and repeated cross-sectional data”. In: *Journal of Econometrics* 121, pp. 175–212.

Morrisson, Christian and Fabrice Murtin (2013). “The Kuznets curve of human capital inequality: 1870–2010”. In: *The Journal of Economic Inequality* 11.1, 283–301.

Neef, Theresa and Anne-Sophie Robilliard (2021). “Half the Sky? The Female Labor Income Share in a Global Perspective”. In: *Journal of Comparative Economics* 47.1, pp. 669–692.

Nimier-David, Elio (2023). “Local Human Capital and Firm Creation: Evidence from the Massification of Higher Education in France”. Working Paper.

Oster, Emily and Bryce Millett Steinberg (2013). “Do IT service centers promote school enrollment? Evidence from India”. In: *Journal of Development Economics* 104.1, 123–135.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman (2018). “Distributional National Accounts: Methods and Estimates for the United States”. In: *Quarterly Journal of Economics* 133.2, pp. 553–609.

Pinkovskiy, Maxim and Xavier Sala-i-Martin (2016). “Lights, Camera … Income! Illuminating the National Accounts-Household Surveys Debate”. In: *The Quarterly Journal of Economics* 131.2, pp. 579–631.

Queiró, Francisco (2022). “Entrepreneurial Human Capital and Firm Dynamics”. In: *The Review of Economic Studies* 89.4, 2061–2100.

Rossi, Federico (2020). “Human Capital and Macroeconomic Development: A Review of the Evidence”. In: *The World Bank Research Observer* 35.2, pp. 227–262.

– (2022). “The Relative Efficiency of Skilled Labor across Countries: Measurement and Interpretation”. In: *American Economic Review* 112.1, pp. 235–266.

Sala-i-Martin, Xavier (2006). “The World Distribution of Income: Falling Poverty and... Convergence, Period”. In: *The Quarterly Journal of Economics* 121.2, pp. 351–397.

Solow, Robert M. (1957). “Technical Change and the Aggregate Production Function”. In: *Review of Economics and Statistics* 39.1, pp. 312–320.

United Nations (2023). *The Sustainable Development Goals Report*. United Nations.

Vu, Khoa and Tu-Anh Vu-Thanh (2022). “Higher Education Expansion, Labor Market, and Firm Productivity in Vietnam”. Working Paper.

Wantchekon, Leonard, Marko Klašnja, and Natalija Novta (2015). “Education and Human Capital Externalities: Evidence from Colonial Benin”. In: *The Quarterly Journal of Economics* 130.2, 703–757.

World Bank (2023). *World Development Indicators*. World Bank.

Figure 1.1: Education and the Distribution of Global Economic Growth, 1980-2019

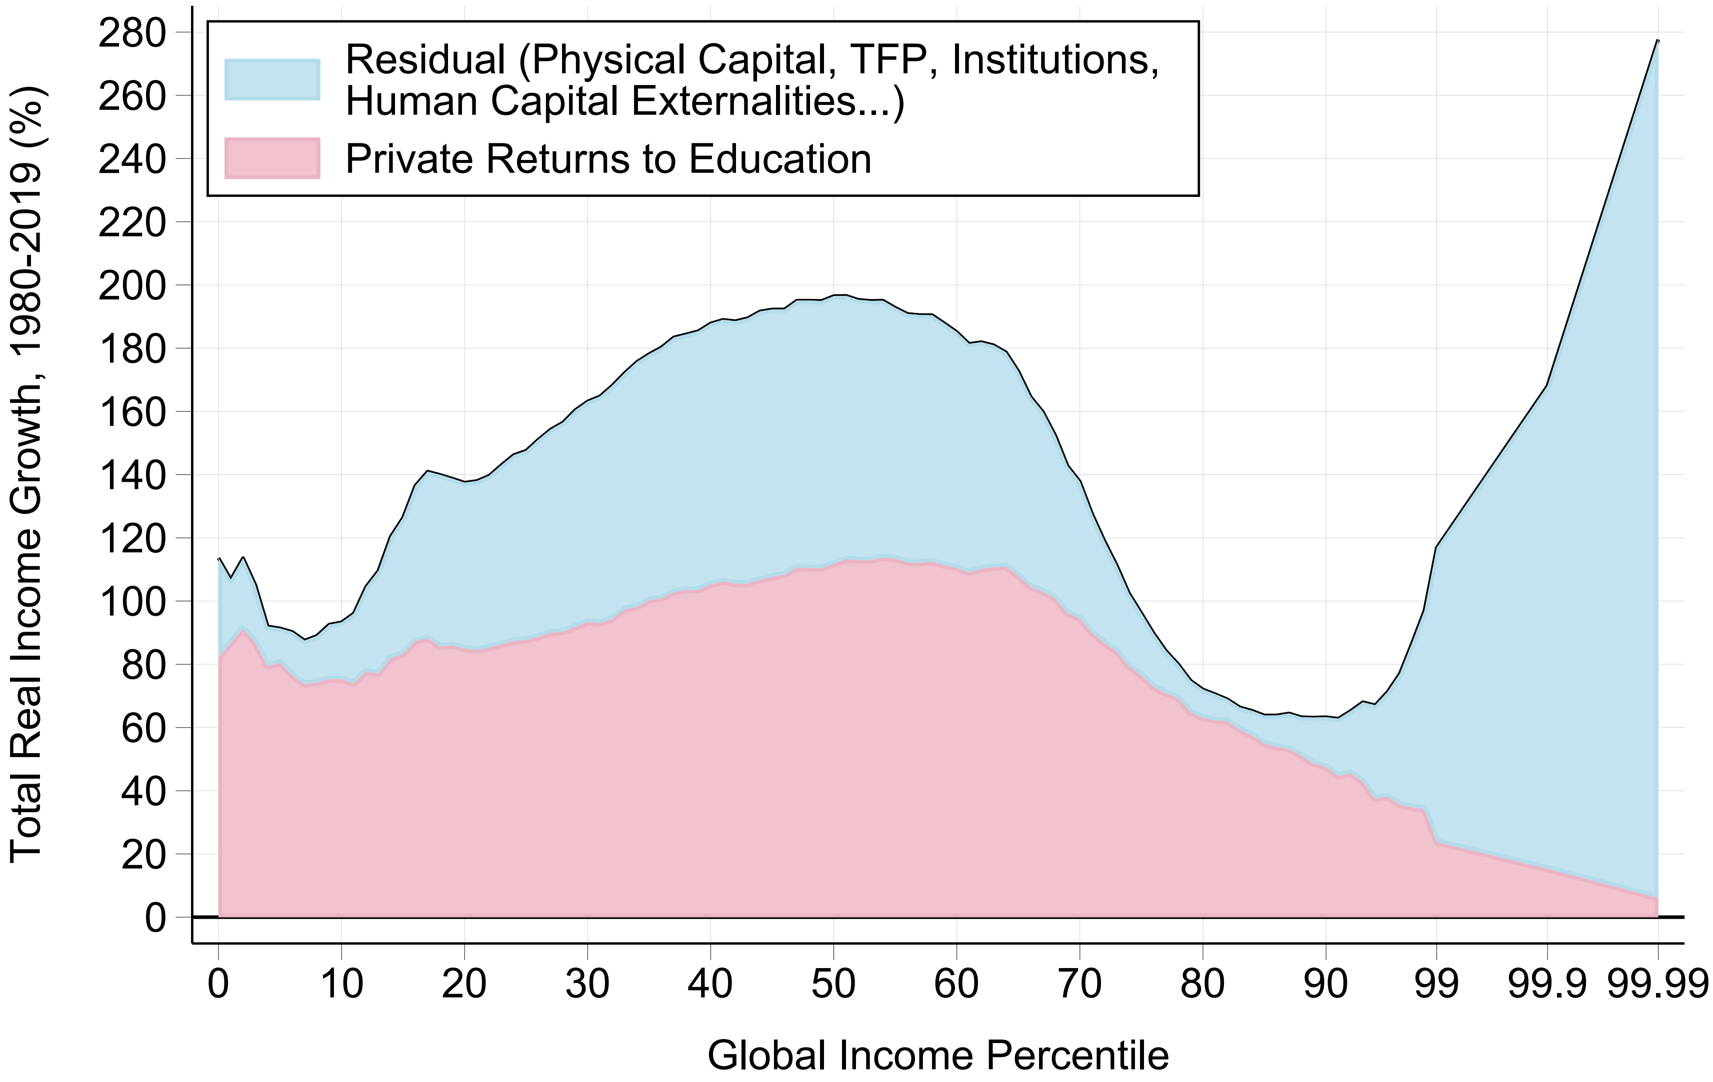

Notes. The figure plots total real income growth by global income percentile from 1980 to 2019, decomposing it into a part that can be explained by private returns to schooling and an unexplained component. The upper shaded area represents the growth rates that would have prevailed absent any improvement in the education of the world's working-age population since 1980. The lower shaded area represents the corresponding contribution of education to economic growth. Taking the ratio between this contribution and actual growth rates, education explains about 70% of growth for the world's 20% poorest individuals. The income concept is pretax income per capita.

Figure 1.2: Survey Microdata Coverage

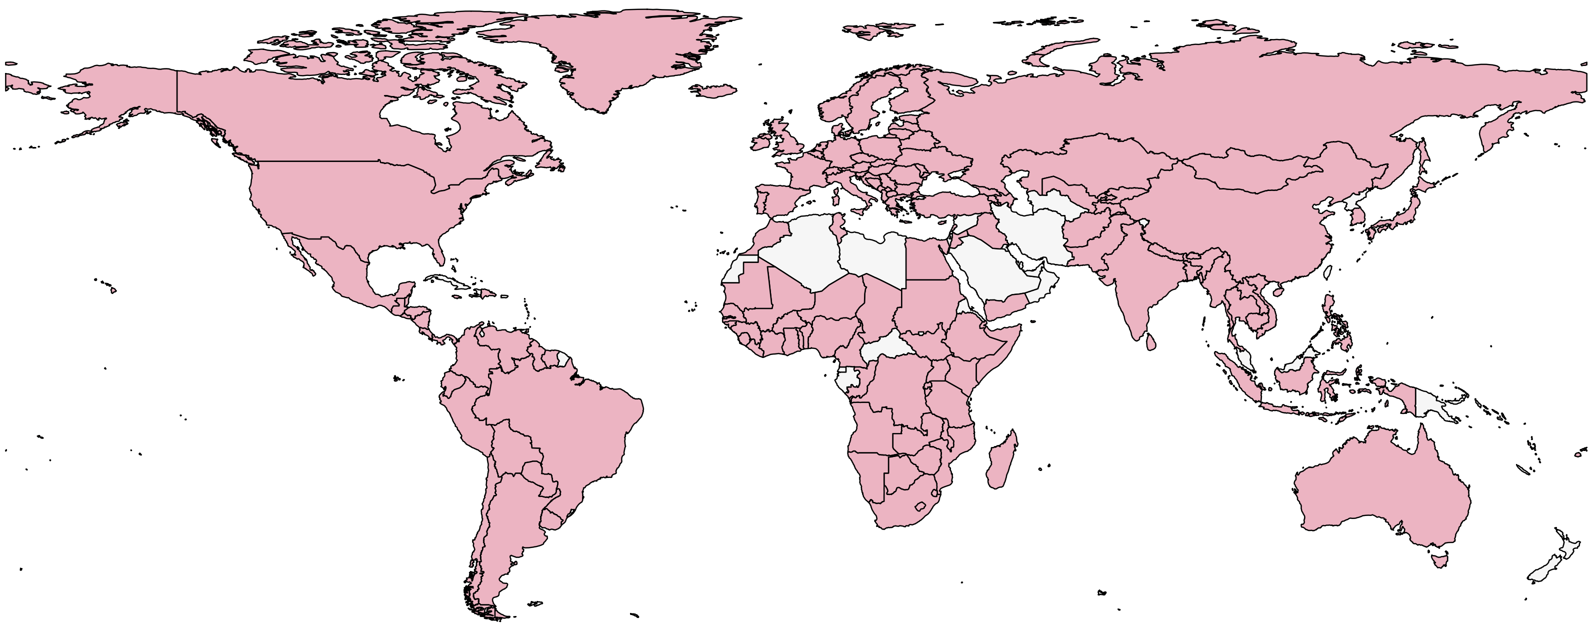

Figure 1.3: Educational Attainment of the Global Working-Age Population, 1980-2019

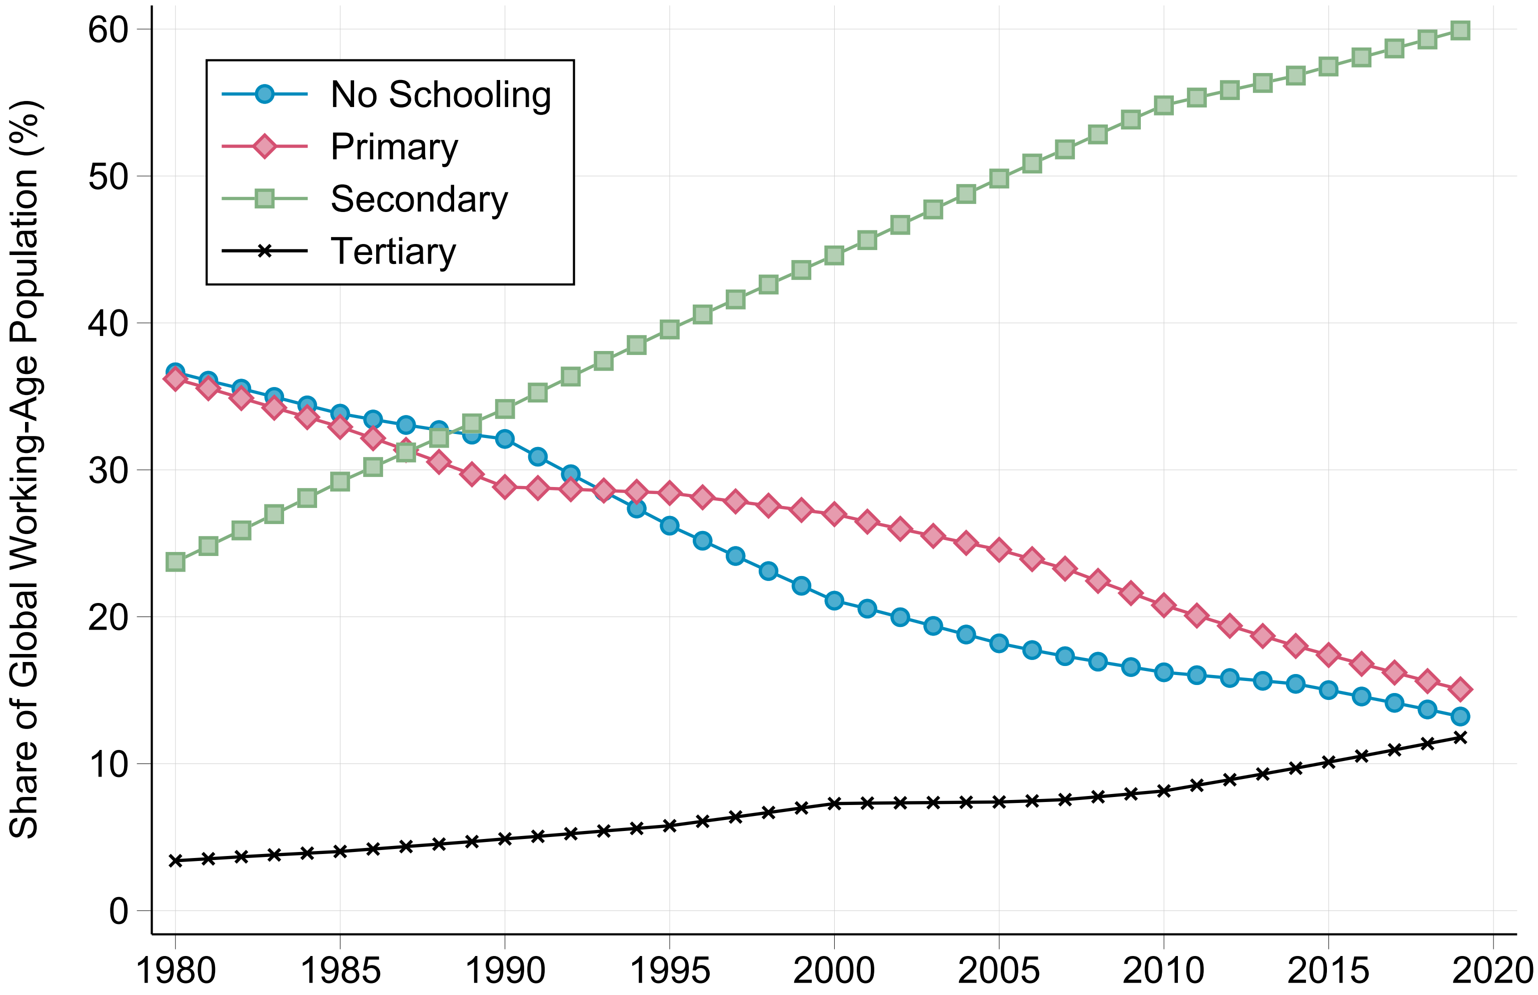

Notes. Author's computations combining data from Barro and Lee (2013) and updates, own sources, and working-age population estimates from the United Nations' World Population Prospects statistics. The figure plots the distribution of educational attainment of the working-age population in the world as a whole.

Figure 1.4: Returns to Schooling by World Region

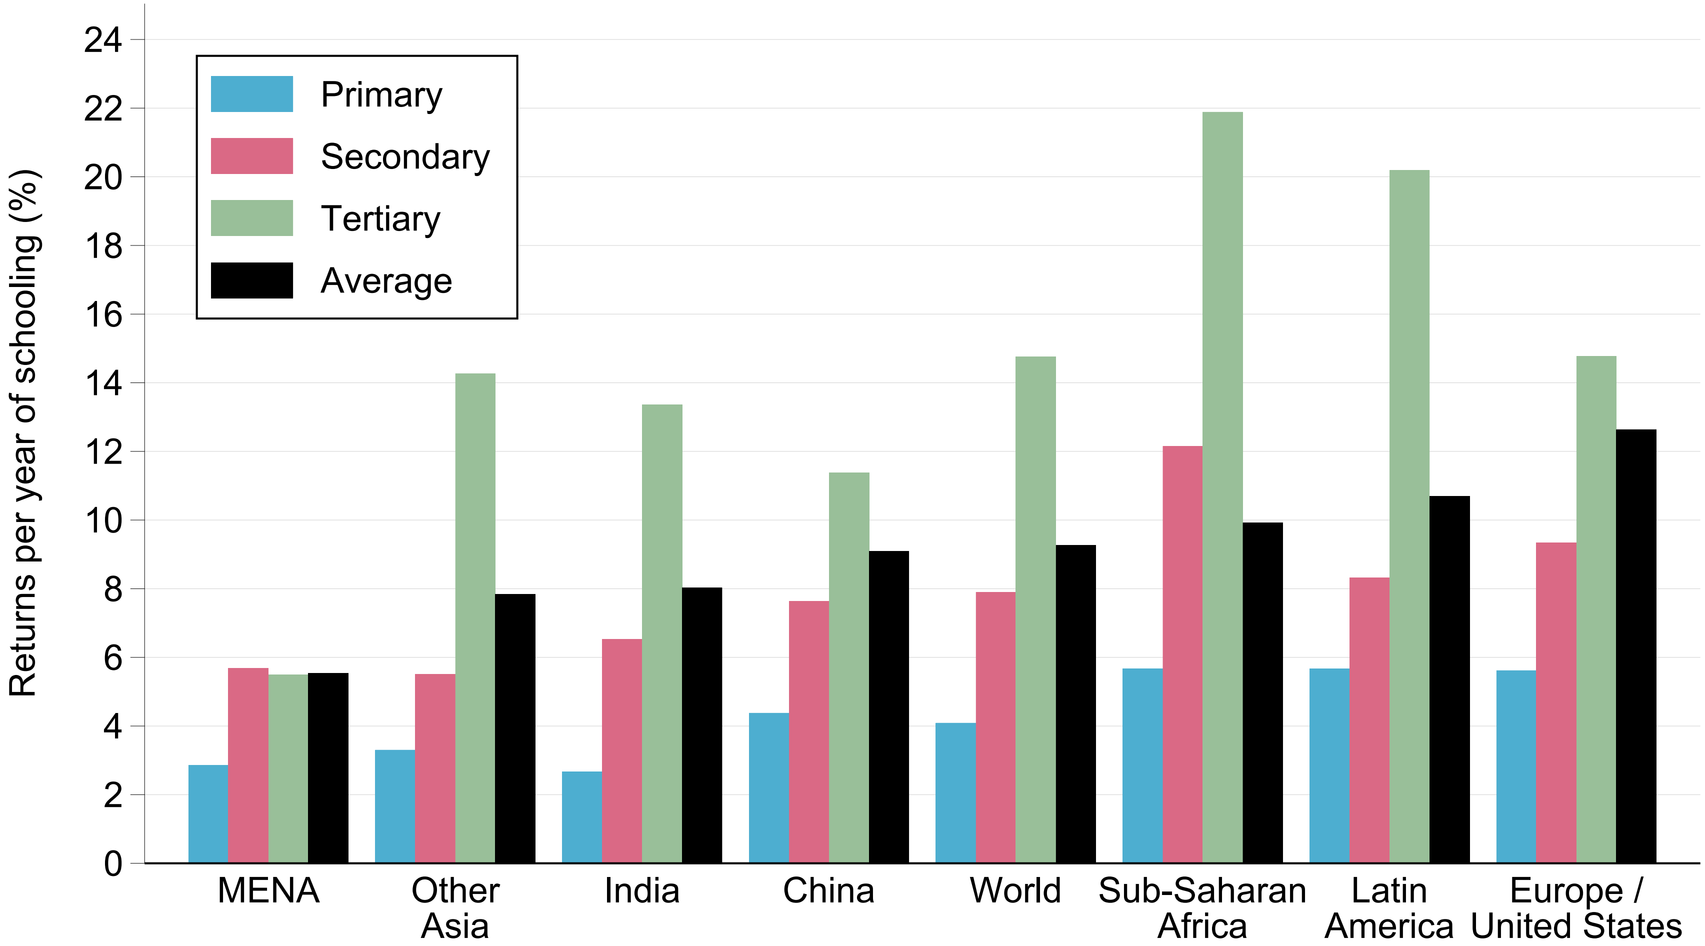

Notes. Author's computations using labor force survey microdata. Estimates correspond to the effect of one additional year of schooling on the log of personal income, estimated using modified Mincerian equations controlling for an experience quartic, gender, and interactions between the experience quartic and gender. Primary: return to an additional year of primary education. Secondary: return to an additional year of secondary education. Tertiary: return to an additional year of tertiary education. Population-weighted averages of coefficients estimated in each country.

Figure 1.5: Returns to Schooling: OLS versus IV Estimates

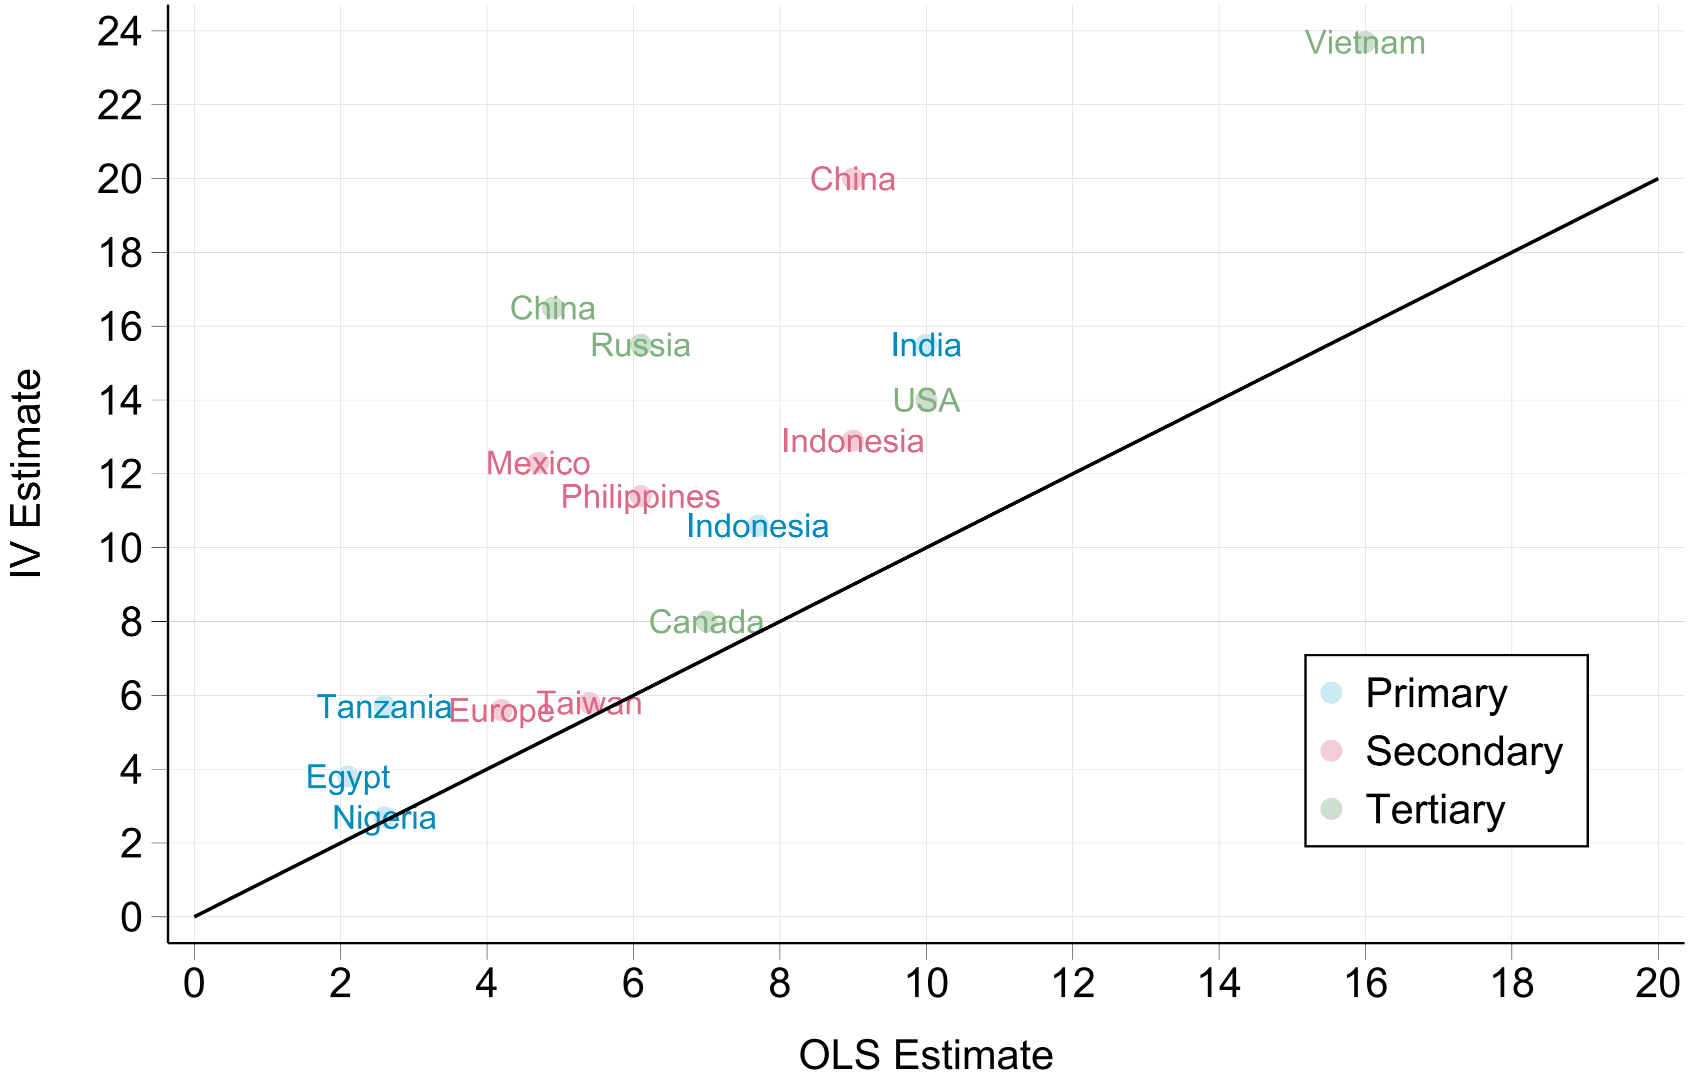

Notes. Author's elaboration compiling estimates from a number of empirical studies: see appendix table A.31. The figure compares ordinary least squares (x-axis) and instrumental variable (y-axis) estimates of the return to an additional year of schooling. OLS estimates generally correspond to results from a Mincerian equation of the log of earnings on years of schooling, estimated over the entire working-age population. In contrast, IV estimates typically rely on quasi-experimental variation in access to a specific level of education (primary, secondary, or tertiary).

Figure 1.6: Validation: Actual Versus Simulated Distributional Effects of India's District Primary Education Program

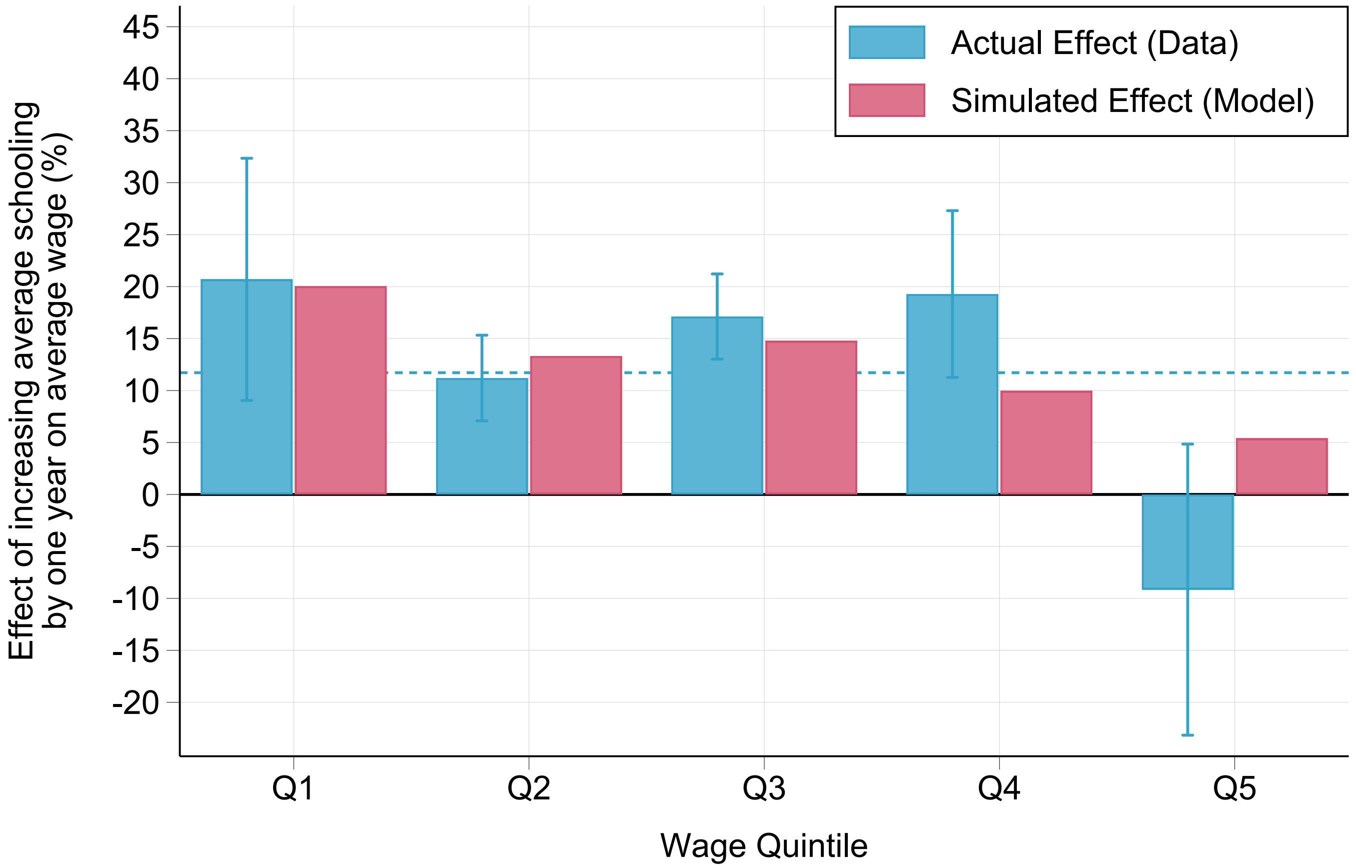

Notes. Actual effect: effect of increasing average district schooling by one year on individual wages by decile, instrumenting schooling with exposure to the District Primary Education Program. Capped spikes correspond to 95% confidence intervals. The dashed line shows the estimated effect of average years of schooling on average earnings. Estimates combine NES microdata with exposure to the policy from Khanna (2023). Simulated effect: predicted effect of increasing average schooling by one year (through primary education) on personal income by decile, calibrating the model on 2019 labor force survey microdata. Simulations assume returns to schooling of 13%.

Figure 1.7: Validation: Actual Versus Simulated Distributional Effects of Indonesia's INPRES School Construction Program

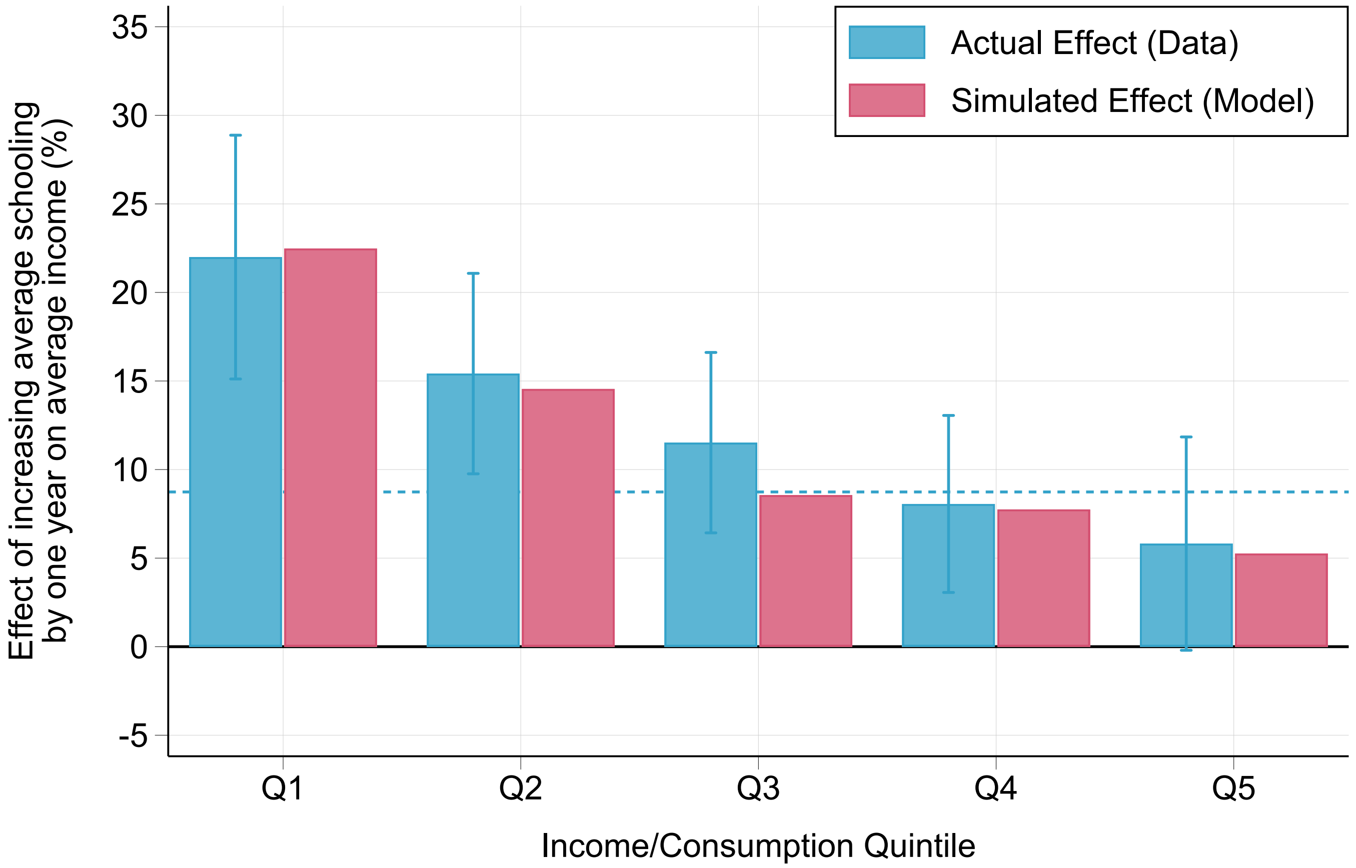

Notes. Actual effect: effect of increasing average district schooling by one year on per-adult consumption by decile, instrumenting schooling with exposure to the INPRES program. Capped spikes correspond to 95% confidence intervals. The dashed line shows the estimated effect of average years of schooling on average consumption. Estimates combine SUSENAS 1993-2019 microdata with INPRES program intensity from Duflo (2001). Simulated effect: predicted effect of increasing average schooling by one year (through primary education) on personal income by decile, calibrating the model on 1996 labor force survey microdata. Simulations assume returns to schooling of 11%.

Figure 1.8: Validation: Actual Versus Simulated Distributional Effects of U.S. Compulsory Schooling Laws

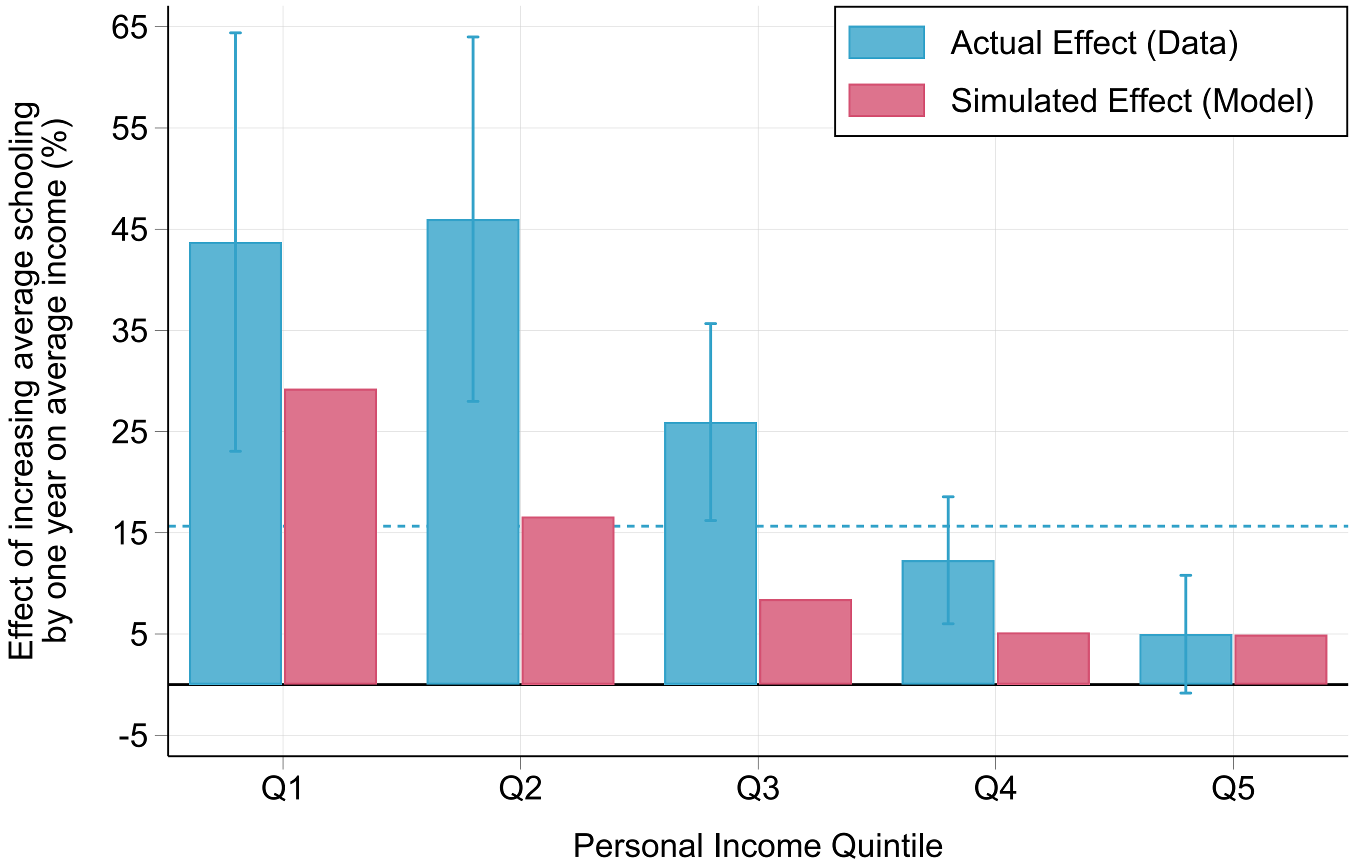

Notes. Actual effect: effect of increasing average state schooling by one year on personal income by decile, instrumenting schooling with state compulsory schooling laws. Capped spikes correspond to 95% confidence intervals. The dashed line shows the estimated effect of average years of schooling on average consumption. Estimates combine 1940-2000 census microdata with compulsory schooling laws from Acemoglu and Angrist (2000) and Clay, Lingwall, and Stephens (2021). Simulated effect: predicted effect of increasing average schooling by one year (through secondary education) on personal income by decile, calibrating the model on 1960 census microdata. Simulations assume returns to schooling of 12%.

Figure 1.9: Share of Growth Explained by Education by Global Income Percentile, 1980-2019

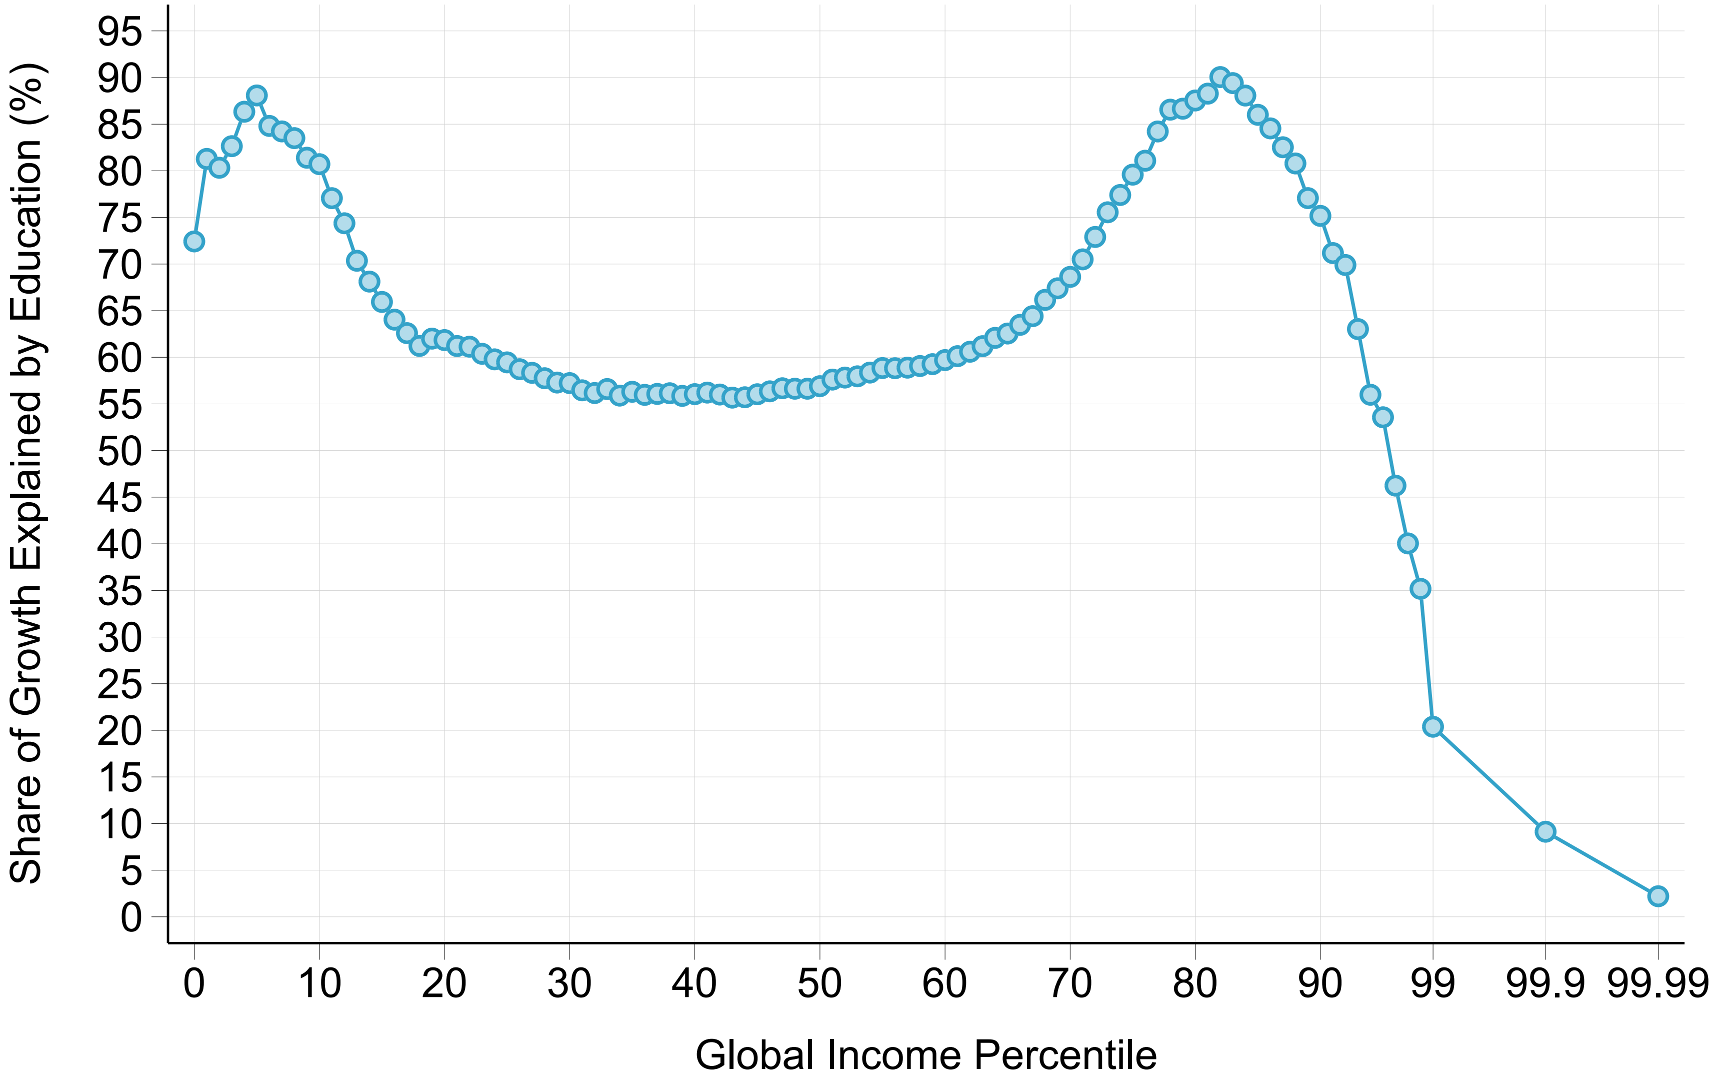

Notes. The figure plots the share of real pretax income growth that can be explained by improvements in educational attainment by global income percentile. Each point corresponds to the ratio of gains from schooling—equal to actual minus counterfactual income growth absent educational expansion—over actual income growth over the 1980-2019 period, calculated for each percentile of the world distribution of income.

Figure 1.10: Annualized Real Income Gains from Schooling by Global Income Group, 1980-2019

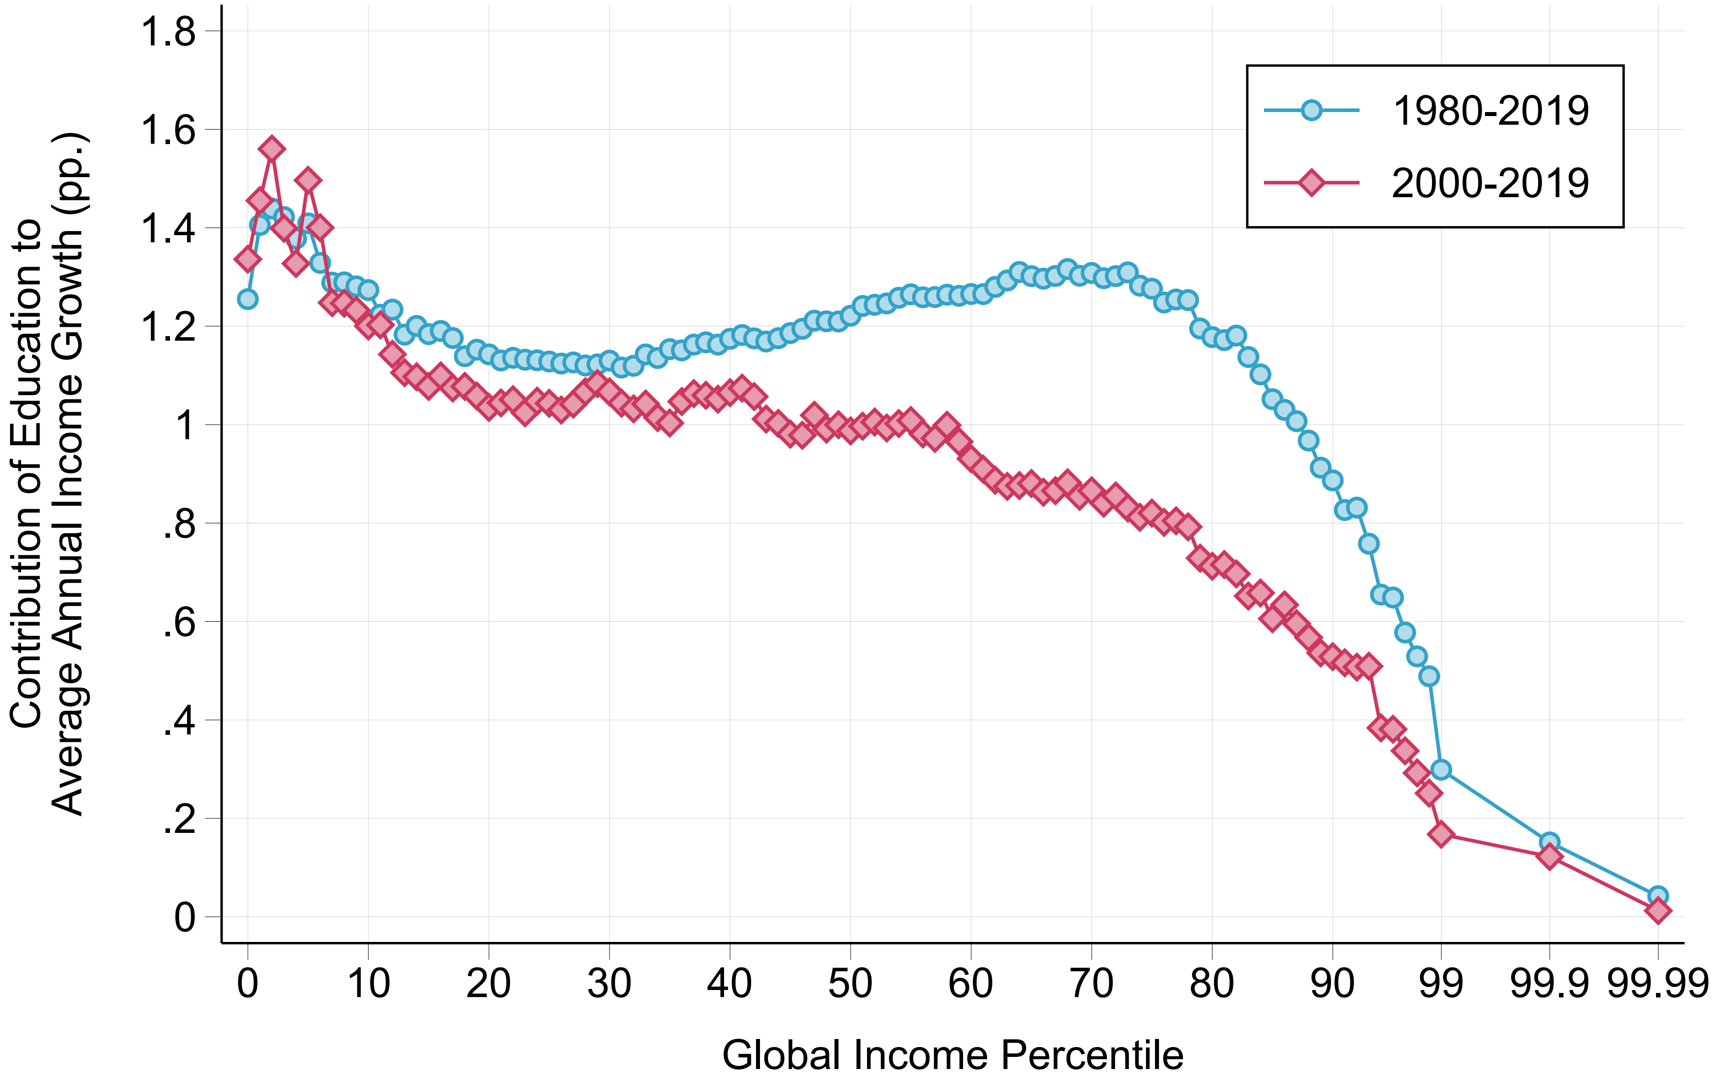

Notes. The figure plots annualized gains from schooling by global income percentile, calculated by taking the percent difference between actual and counterfactual income growth absent educational expansion, and annualizing the resulting figure over the period considered. Interpretation: from 1980 to 2019, education contributed to increasing average incomes by 1-1.4% per year for the world's poorest 80% individuals.

Figure 1.11: Global Gender Schooling Inequality, 1991-2019
Gender Gap in Years of Schooling (Men - Women)

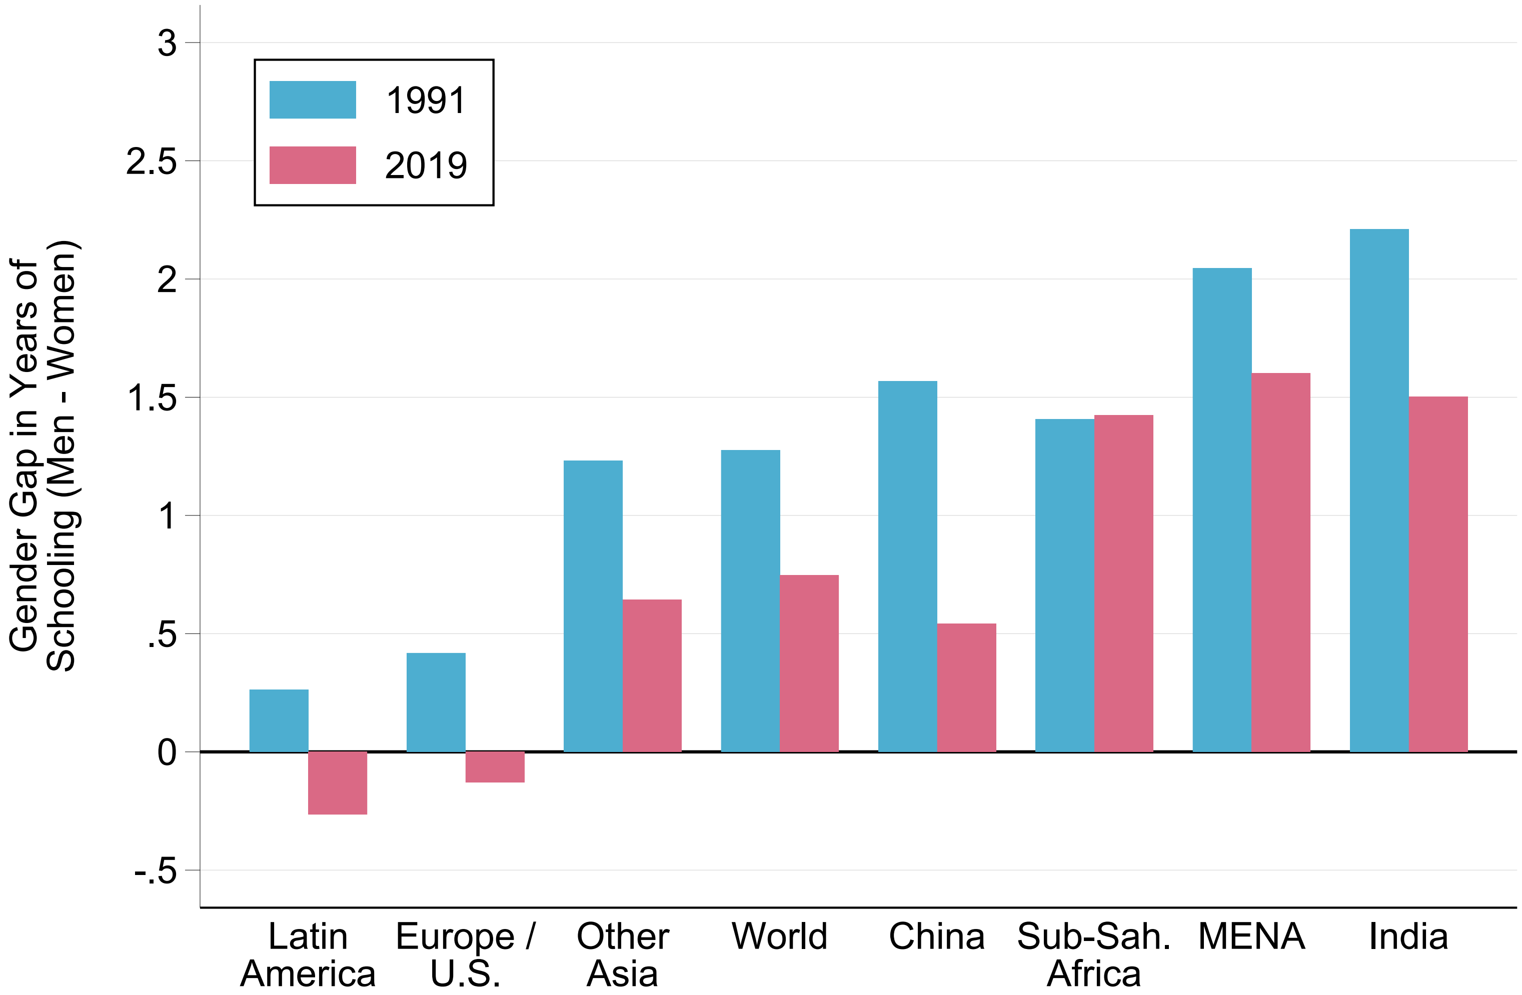

Notes. Author's computations using data from Barro and Lee (2013) and updates. The figure shows the population-weighted average gap in years of schooling between working-age men and women by world region and in the world as a whole.

Figure 1.12: Returns to Schooling by Gender and World Region

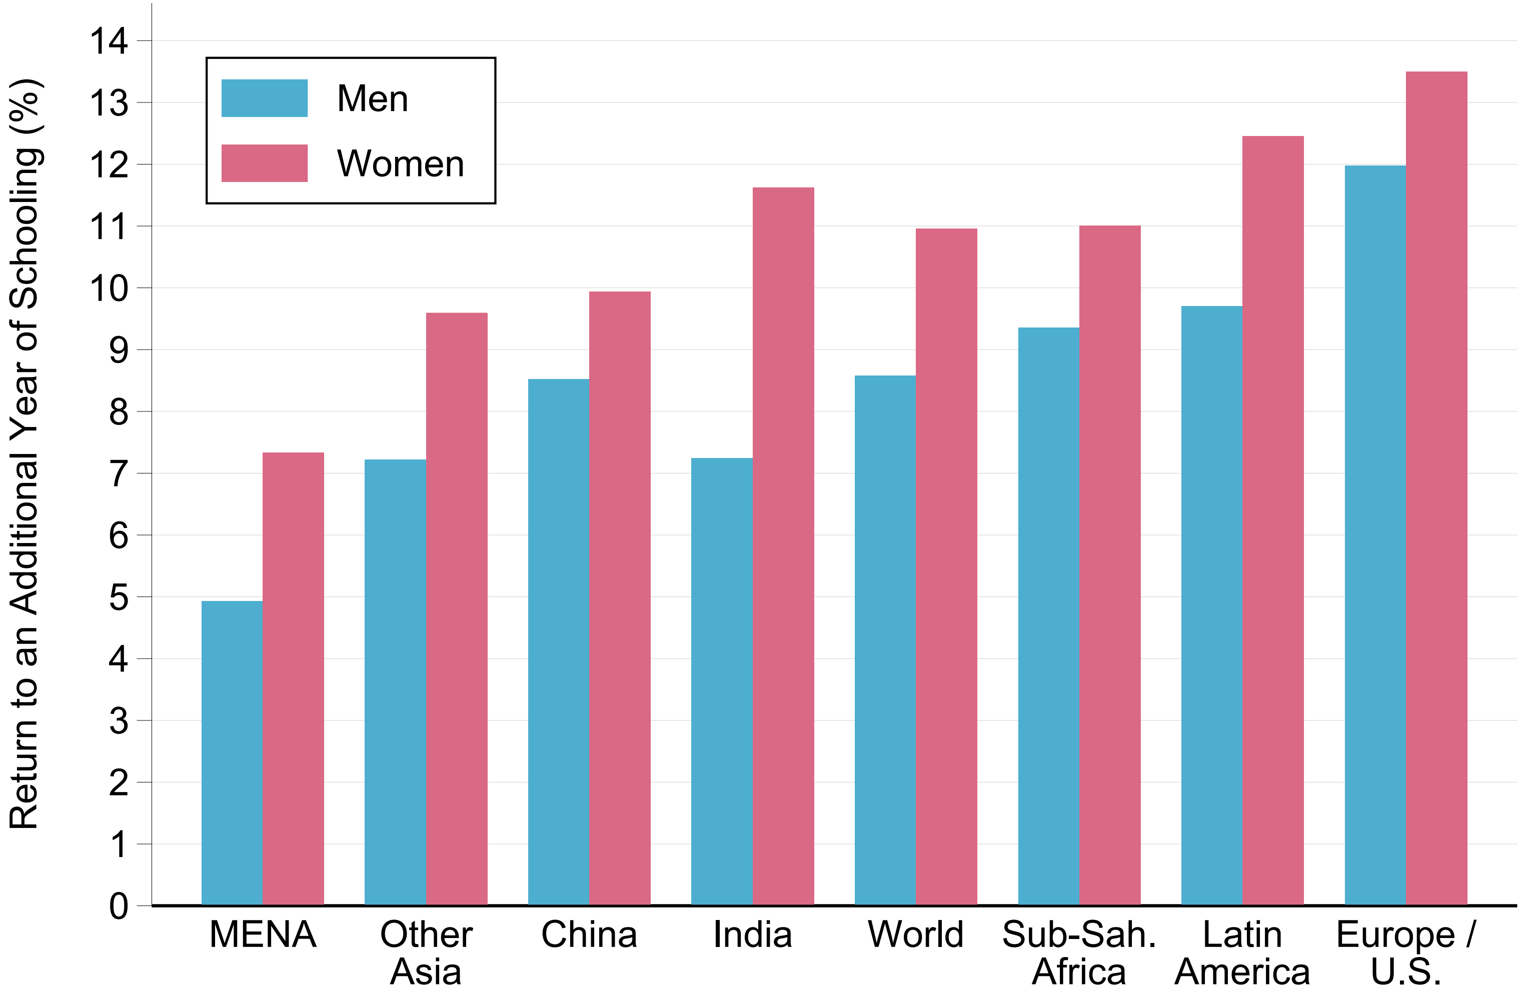

Notes. Author's computations using labor force survey microdata. Estimates correspond to the effect of one additional year of schooling on the log of personal income, estimated separately by gender using modified Mincerian equations controlling for an experience quartic. Population-weighted averages of coefficients estimated in each country.

Figure 1.13: Share of Gender Inequality Reduction Explained by Education by World Region, 1991-2019

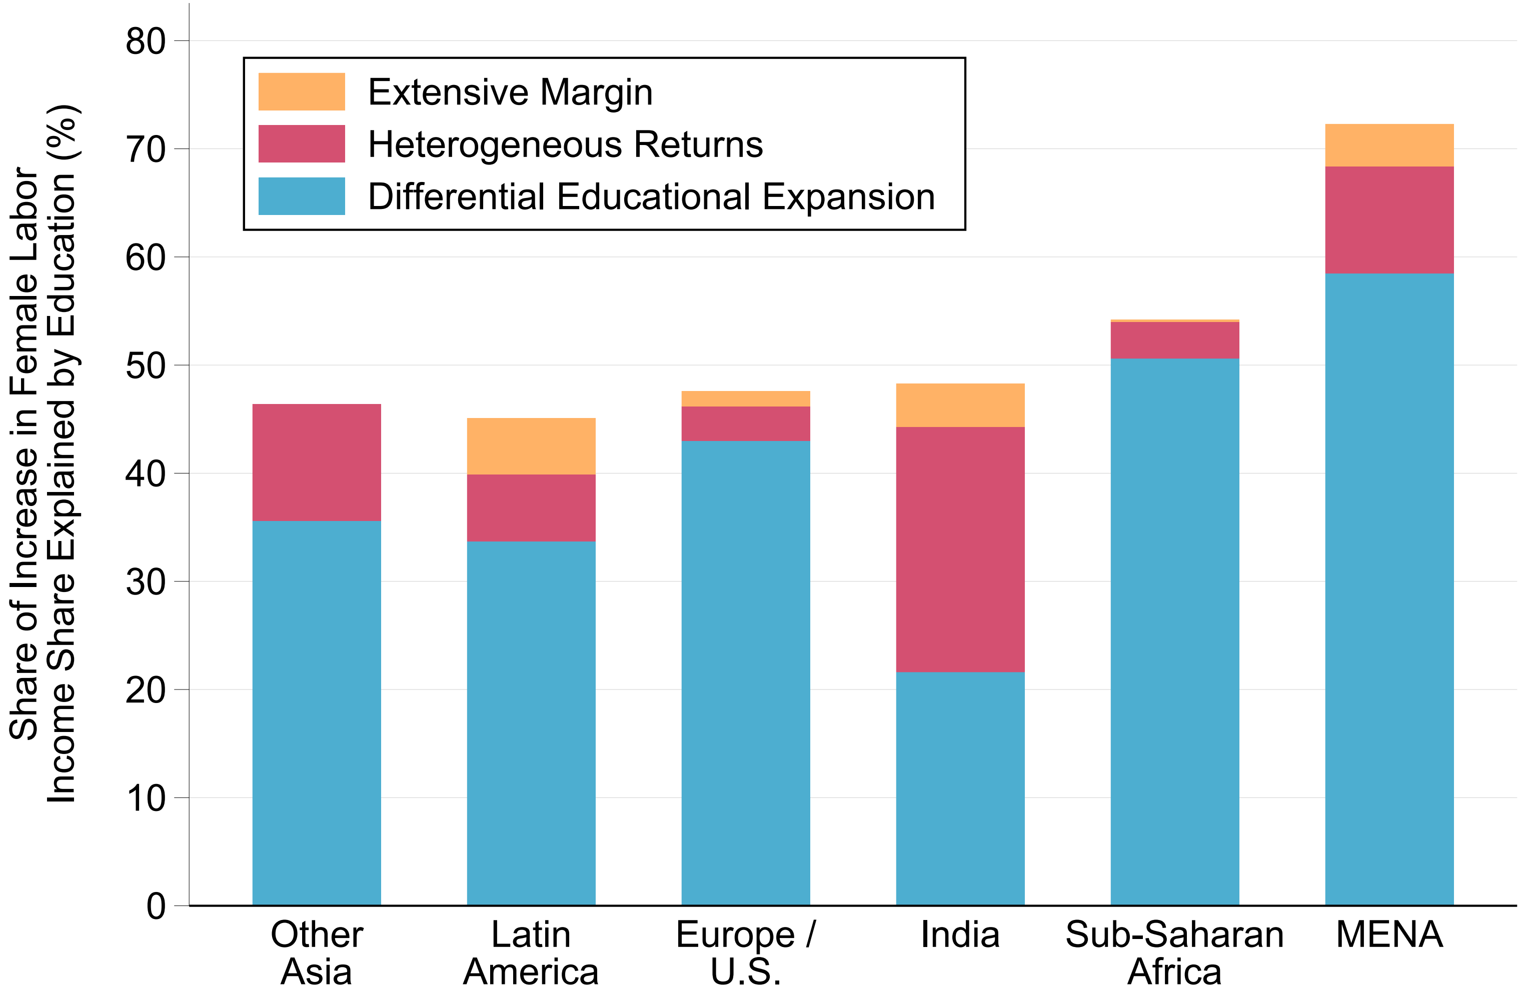

Notes. Author's computations using labor force survey microdata. Population-weighted averages of gains from schooling estimated in each country.

Table 1.1: Survey Microdata Descriptive Statistics

	Countries	Observations	Share of Population Covered
Europe	39	743,328	100.0%
Northern America	2	539,862	100.0%
Latin America	24	4,126,194	96.5%
Asia-Pacific	29	2,524,531	95.5%
Middle East and North Africa	14	817,958	74.2%
Sub-Saharan Africa	42	876,054	98.9%
World	150	9,627,927	95.2%

Notes. The table reports the number of countries covered by the survey microdata, the total number of observations, and the share of the total population covered by world region and in the world as a whole (last row).

Table 1.2: Distributional Growth Accounting, World, 1980-2019

	Total Income Growth (%) g	Growth Without Education (%) $\tilde{g}$	Contribution of Education (pp.) $g - \tilde{g}$	Share of Growth Explained (%) $\frac{g-\tilde{g}}{g}$
Full Population	+98%	+45%	53	54%
Bottom 50%	+164%	+68%	96	59%
Bottom 20%	+115%	+34%	81	71%
Next 30%	+176%	+76%	100	57%
Middle 40%	+94%	+24%	70	74%
Top 10%	+91%	+57%	34	38%
Top 1%	+131%	+109%	22	17%
Top 0.1%	+173%	+158%	15	9%
Top 0.01%	+278%	+272%	6	2%

Notes. The table reports actual real income growth rates, counterfactual growth rates absent educational expansion, and the corresponding share of growth explained by education for different groups of the world distribution of income.

Table 1.3: Education and Global Poverty Reduction

	1980	2019	Difference (%)	Share of Decline Explained (%)
Global Poverty: $2.15 / Day				
Actual	20%	9%	-55%	
Counterfactual	20%	13%	-32%	42%
Global Poverty: $3.65 / Day				
Actual	40%	14%	-65%	
Counterfactual	40%	22%	-44%	32%
Global Poverty: $6.85 / Day				
Actual	58%	27%	-53%	
Counterfactual	58%	41%	-30%	44%

Notes. The table compares the actual evolution of the global poverty headcount ratio to the evolution it would have followed absent educational expansion since 1980. All global poverty headcount ratios calculated using 2017 PPP USD. The income concept is pretax national income, as reported in the World Inequality Database. See appendix table A.3 for comparable results using per-capita consumption distributions from the World Bank.

Table 1.4: Distributional Growth Accounting by World Region and Country Income Group, 1980-2019

	Full Population			Bottom 20%		
	Actual Growth	Contrib. of Education	Share Explained	Actual Growth	Contrib. of Education	Share Explained
	g	$g - \tilde{g}$	$\frac{g-\tilde{g}}{g}$	g	$g - \tilde{g}$	$\frac{g-\tilde{g}}{g}$
Europe / Northern America	+81%	50	62%	+30%	69	¿100%
Latin America	+38%	33	86%	+39%	60	¿100%
China	+988%	289	29%	+377%	205	54%
India	+410%	124	30%	+208%	107	52%
Other Asia-Pacific	+117%	56	48%	+263%	125	48%
Middle East and North Africa	+118%	46	39%	+35%	46	¿100%
Sub-Saharan Africa	+15%	31	¿100%	+51%	65	¿100%
Low-income	+16%	35	¿100%	+50%	70	¿100%
Low-middle-income	+182%	68	37%	+198%	100	51%
High-middle-income	+210%	83	40%	+218%	143	65%
High-income	+89%	50	57%	+77%	94	¿100%

Notes. The table reports actual real income growth rates and the corresponding share of growth that can be explained by education, for the full population and the poorest 20%, by world region.

Table 1.5: Education and Inequality Between and Within Countries

	1980	2019	Difference
Theil Index of Global Inequality			
Actual	1.06	1.07	0.02
Counterfactual	1.06	1.33	0.27
Between-Country Component			
Actual	0.60	0.34	-0.26
Counterfactual	0.60	0.34	-0.26
Within-Country Component			
Actual	0.46	0.74	0.28
Counterfactual	0.46	0.99	0.53
Within-Country Share (%)			
Actual	43%	69%	25
Counterfactual	43%	74%	31

Notes. The table compares the actual evolution of global inequality since 1980 to the evolution it would have followed absent educational expansion, decomposing these transformations into a between-country component and a within-country component. Within-country share: share of global inequality explained by inequality within countries.

Table 1.6: From Standard to Distributional Growth Accounting

	Share of Growth Explained, 1980-2019	
	Global Average	Global Bottom 20%
Standard Growth Accounting Cross-Country Data, 10% Return	33%	23%
+ Adjusted Labor Share	43%	35%
+ Within-Country Inequality	43%	41%
+ Within-Country Labor Shares	43%	51%
+ Microdata, 2019 Returns	40%	41%
+ Distributional Effects, 2019 Returns	40%	52%
+ Distributional Effects, Adjusted Returns	54%	71%
+ Distributional Effects, IV Returns	54%	75%

Notes. The table reports the share of global economic growth and real income growth of the global bottom 20% that can be explained by education, depending on methodological assumptions and data sources used. Adjusted labor share: labor income includes mixed income. Within-country inequality: income distribution data from the World Inequality Database. Within-country labor shares: labor share varies by income group within each country. 2019 returns: Mincerian returns by level estimated using the microdata. Adjusted returns: true returns estimated by adding supply effects to 2019 returns. IV returns: 2019 returns adjusted using instrumental variable estimates of returns to schooling. Distributional effects: relative wage adjustments due to supply effects (imperfect substitution between skill groups).

Table 1.7: Education and Global Gender Inequality, 1991-2019

	1991	2019	Diff.	Share Explained By Education	Share Explained (Cross-Country Average)
Global Female Labor Income Share	29.3%	32.1%	2.8		
Counterfactual: No Educational Progress	29.3%	30.6%	1.3	55%	50%
Counterfactual: + Heterogeneous Returns	29.3%	30.1%	0.8	71%	58%
Counterfactual: + Extensive Margin	29.3%	29.9%	0.6	78%	59%

Notes. The table reports actual versus counterfactual global female labor income shares under different assumptions. Global female labor income: total share of labor income received by women in the world as a whole. Change in education: only account for differential trends in schooling by gender, applying the same returns to schooling for men and women to build the counterfactual. Heterogeneous returns: account for differential returns by gender. Extensive margin: account for differential effects of schooling on employment by gender. Cross-country average: population-weighted average of the share of gender inequality reduction explained by education in each country.

Table 1.8: Public Policies and Global Poverty Reduction:
Combining Direct Redistribution and Indirect Investment Benefits from Education

	1980	2019	Change (%)	Total Share of Change Explained (%)
Global Poverty Rate ($2.15/Day)				
Pretax Income Absent Educational Expansion	20%	13%	-32%	
Pretax Income	20%	8.7%	-55%	
Posttax Income	17%	5.1%	-70%	54%
Global Bottom 20% Average Income ($/Day)				
Pretax Income Absent Educational Expansion	1.3	1.7	+32%	
Pretax Income	1.3	2.8	+115%	
Posttax Income	1.5	4.0	+164%	80%
Global Bottom 50% Average Income ($/Day)				
Pretax Income Absent Educational Expansion	2.7	4.5	+67%	
Pretax Income	2.7	7.1	+163%	
Posttax Income	2.9	8.8	+200%	66%

Notes. The table compares the evolution of global poverty and the average income of the global bottom 20% and bottom 50% under three scenarios. The first one considers the evolution of each indicator if there had been no educational progress since 1980 ("pretax income absent educational expansion"). The second one corresponds to the actual evolution of each indicator in terms of pretax income ("pretax income"). The third one corresponds to the actual evolution of each indicator in terms of posttax income, that is, after removing all taxes and adding all cash and in-kind transfers (see Gethin, 2023b). The last column displays the corresponding share of global poverty reduction or real income gains that can be attributed to public policies, combining indirect investment benefits from education (moving from "pretax income absent educational expansion" to pretax income) and direct redistribution (moving from pretax to posttax income), calculated as one minus the ratio of the first row to the third row of the fourth column within each panel. Global poverty rate calculated at $2.15 per day in 2017 PPP USD. Real incomes of the bottom 20% and bottom 50% expressed in 2021 PPP USD as in the rest of the paper. Estimates of the world distribution of income from the World Inequality Database. See appendix table A.12 for comparable results using per-capita consumption distributions from the World Bank.

Chapter 2

Revisiting Global Poverty Reduction: Public Goods and the World Distribution of Income, 1980-2019

Government redistribution is rising around the world. Between 1980 and 2019, real government expenditure per world citizen doubled, from about \$2500 to \$5000 at purchasing power parity. Cash transfers cannot be held responsible: they represent less than 10% of global public expenditure and have scarcely increased since 1980. Instead, the bulk of the growth of government redistribution has been driven by investments in public education, healthcare, housing, police services, transport infrastructure, and other public goods. Together, these transfers represented some 30% of global GDP in 2019.[1]

This dramatic transformation remains largely absent from poverty and inequality statistics. The standard concept used to measure global poverty is household final consumption expenditure, defined as the market value of all goods and services purchased by households. By construction, it excludes public goods, since these goods are not bought on a market. As a result, it remains difficult to understand how macroeconomic growth reduces poverty, in a world where almost a third of global GDP is redistributed by governments in unaccounted ways. It also limits our ability to answer some of the most basic questions of human development, such as: who

[1]See figure 2.4, which plots the evolution of real worldwide government expenditure per capita since 1980. Appendix figure B.25 plots global government expenditure as a share of global GDP.

benefits from public goods? How does the provision of public services vary across time and space? And to what extent have public goods contributed to global poverty reduction in the past decades?

This paper represents a first attempt at answering these questions. I propose a simple framework for studying the distribution of public goods that combines two parameters: a cost parameter and a progressivity parameter. The cost parameter corresponds to how much governments spend on each type of transfer. The progressivity parameter governs the share of this transfer that is received by different income groups. I also investigate the robustness of my results to accounting for a productivity parameter, capturing the fact that holding cost constant, the quality of public goods provided may vary across countries, over time, and throughout the income distribution.

I apply this framework to the study of global poverty reduction since 1980. The starting point is a new database on the world distribution of public spending, which I construct by combining data from about twenty different sources. To cover the cost component, I draw on budget data to build new aggregate series on the level and composition of general government expenditure. To cover the progressivity component, I rely on estimates of the distributional incidence of public education and healthcare from various fiscal incidence studies and surveys. To account for potential variations in productivity, I construct measures of cost efficiency by benchmarking the value of in-kind transfers to government performance measures. The resulting dataset yields new estimates of the monetary value of public goods received by income group in most countries in the world from 1980 to 2019. It also covers the distribution of taxes and cash transfers, allowing me to compare the incidence of public services to that of these other traditional redistributive tools.

I find that the rising consumption of public goods has played a major role in improving the living conditions of the world's poorest individuals. Figure 2.1 plots the evolution of the global poverty headcount ratio since 1980 before and after accounting for taxes and transfers. The share of the world's population living with less than $2.15 per day in 2017 PPP US dollars declined from 23% to 13% in terms of pretax income, representing a 43% decline. After deducting taxes from individual incomes and adding cash and in-kind transfers, this figure rises to 63%. By this measure, government redistribution accounts for about 30% of global poverty reduction since 1980. Public goods alone account for 20%.

Public goods have also played a key role in making global economic growth more inclusive. Public goods tend to strongly reduce inequality within countries, because they are almost always more equally distributed than pretax incomes. As a result, increasing spending on public education, healthcare, and other public services has strongly reduced global inequalities. All income groups within the bottom 60% of the world distribution of income have benefited from greater net government transfers since 1980. The global top 10% to bottom 50% income ratio has declined by 30% before accounting for public services, compared to 36% after doing so. Public goods thus explain almost 20% of total global inequality reduction since 1980. Today, they reduce global income disparities as much as taxes and cash transfers combined.

I also find that dimensions of government redistribution are correlated across countries. In particular, low-income countries score lower on most dimensions of government redistribution. Not only do they spend less on public services, they also invest more heavily in services that are more regressive and provide each of them more unequally. There is also evidence that they provide public services less efficiently than high-income countries, even after accounting for differences in cost of provision. This "triple curse" comes with extreme inequalities in the quality of public services received worldwide. In 2019, only about 0.5% of global GDP was redistributed to the poorest 10% of world citizens, while almost 10% of global GDP accrued to the richest global income decile. As a result, accounting for public goods increases the share of global income disparities explained by inequalities between countries. The share of the global poor living in poor countries is greater than we thought, because citizens of poor countries benefit from public services of much lower quality than those of the rich world.

Together, these results highlight the critical role played by public-private complementarities in reducing poverty. Economic growth not only improves the labor market and consumption opportunities of low-income households. It also comes with greater government revenue through taxation, a significant fraction of which ends up being redistributed in the form of improved public services. In directly accounting for public goods consumption in the measurement of poverty, my results thus uncover and quantify an important channel—enhanced public spending—through which economic growth contributes to global poverty reduction. It is also important to stress that investments in education, healthcare, and other public services are likely to have also contributed to pretax income growth, in addition to their direct effects on posttax inequality. Accounting for this indirect channel would lead to putting even more weight on public goods in explaining global poverty reduction. In a companion paper, I show that private returns to schooling alone can account for over half of real pretax income gains for the world's poorest 20% individuals since 1980 (Gethin, 2023a).

Despite their relative robustness, two important limitations of these findings should be acknowledged. A first limitation is empirical: due to the lack of comprehensive data, our understanding of the incidence of many public goods remains quite limited. The approach I adopt thus consists in deriving *lower bounds* on the progressivity and productivity of government expenditure. For instance, I distribute spending on a number of public services proportionally to posttax disposable income, which amounts to assuming that high-income groups benefit from substantially higher transfers. I also make the conservative assumption that the productivity of governments is never higher than that of the private sector. My results can be easily updated and improved as better data becomes available, with the likely conclusion that public goods have contributed to the decline of global poverty and inequality to an even greater extent than estimated here.

A second limitation is more conceptual in nature. While the results presented here provide useful information on the distribution of public goods, they tell us little of their value from the perspective of economic welfare. A classic result of economic theory states that the value of an in-kind transfer should not be higher than that of cash, because cash allows consumers to choose what they consume (Atkinson and Stiglitz, 1976). Yet, a growing literature questions the validity of this claim. For instance, in-kind transfers may be preferable to cash if they insure households against commodity price risk (Gadenne et al., 2022), have larger spillover effects onto children (Hendren and Sprung-Keyser, 2020), or if recipients have a desire for self-control mechanisms (Liscow and Pershing, 2022). There is also survey evidence that individuals may prefer public goods to cash, in particular education and health, both in rich and poor countries (Khemani, Habyarimana, and Nooruddin, 2019; Thesmar and Landier, 2022). I do not attempt to disentangle these different factors here. Put simply, this paper studies the incidence of public services on the distribution of *total consumption*, including both privately and publicly provided goods, in the same way as GDP is used to compare total production across countries and over time. Moving from consumption-production to economic welfare would require estimating individuals' willingness to pay for the private and public goods that they consume. I discuss challenges in doing so and avenues for future research in this direction in section 2.4.3.[2]

This article contributes to our understanding of the evolution of global poverty in

[2]The results presented in this paper can be interpreted as mirroring economic welfare if willingness to pay is exactly equal across all types of private and public goods. It is also important to mention that standard poverty statistics do already incorporate a number of in-kind incomes that are not necessarily optimally "chosen." These include, for instance, own consumption of food produced by the household and gifts received in kind from other households, both of which may be valued significantly less than cash.

the past decades. The classic approach to measuring monetary poverty is to compute the share of individuals whose consumption falls below a given threshold (e.g., Chen and Ravallion, 2010; Deaton, 2010; Ravallion, 2012). While such measures provide invaluable information on the living standards of the poor, they fail to capture dimensions of economic well-being that are not typically bought on a market. Well aware of this limitation, international organizations and statistical institutes have started developing a number of indicators of multidimensional poverty.[3] These different measures have provided useful insights, yet they tend to suffer from limited space and time coverage and are not directly comparable with growth statistics. In this paper, I tackle some of these limitations by constructing measures of monetary poverty and inequality that incorporate public services. I provide evidence that doing so contributes to reconciling monetary and multidimensional approaches to measuring living standards, precisely because public services are major determinants of cross-country differences in deprivation in health, education, and other non-monetary dimensions of quality of life.

This article also provides new evidence on the evolution of global income inequality. A number of studies have attempted to estimate the world distribution of income, generally focusing on household consumption or pretax income (e.g., Bourguignon and Morrisson, 2002; Chancel and Piketty, 2021; Lakner and Milanovic, 2016; Sala-i-Martin, 2006). I contribute to these efforts by estimating the incidence of all types of taxes and transfers on global poverty and inequality. To the best of my knowledge, this study is the first to analyze how government redistribution in its various forms has contributed to shaping global income disparities since the 1980s.

My methodology is directly inspired by the growing literature attempting to bridge gaps between micro- and macro-approaches to the measurement of living standards. Piketty, Saez, and Zucman (2018) construct Distributional National Accounts (DINA) for the United States, allocating the entirety of national income, taxes, and government expenditure to individuals every year since 1913. A number of studies following this framework have been conducted on other countries since then.[4] The major

[3]Such measures have become increasingly available and mobilized in both developed in developing countries: by 2017, 16 countries used multidimensional poverty indices as official measures of poverty (Glassman, 2019). Since 2010, the Oxford Poverty and Human Development Initiative has published cross-country measures of multidimensional poverty that combine indicators on deprivations in health, education, and living standards (Alkire, Kanagaratnam, and Suppa, 2021). In the same spirit, the World Bank has recently released a multidimensional poverty measure that incorporates both monetary and non-monetary components (World Bank, 2018).

[4]See in particular Blanchet, Chancel, and Gethin (2022) on Europe, Bozio et al. (2022) on France, and De Rosa, Flores, and Morgan (2022b) on Latin America. See also Germain et al. (2021), Bruil et al. (2022), and Jestl and List (2022), who cover posttax income for a limited number of years in France, the Netherlands, and Austria, respectively. See Chancel et al. (2022b) for a

advantage of this methodology is that it produces estimates of income inequality that are consistent with macroeconomic growth. Its main limitation is that it does not generally account for the progressivity and productivity of public goods. Instead, studies typically assume that all public goods are valued at cost, and received either proportionally to posttax disposable income or as a lump sum.[5] In this article, I go beyond these simplifying assumptions by explicitly accounting for the progressivity and productivity of public education and healthcare in a national accounts framework.

More generally, this paper extends our knowledge of who benefits from in-kind transfers. A large body of literature has attempted to estimate the distributional incidence of specific public services in specific contexts.[6] While many of the methods used in this article are directly inspired from this work, I depart from existing studies in taking a long-run, historical perspective on the incidence of all forms of government redistribution on global poverty.

The rest of the paper is organized as follows. Section 2.1 presents motivating evidence and the general framework used to study the distribution of public goods. Section 2.2 applies this framework to build a new database on public goods provision worldwide since 1980. Section 2.3 presents the results. Section 2.4 investigates the role of potential differences in public sector productivity, discusses how public services can help solving well-known discrepancies between surveys and national accounts in the measurement of poverty, and provides a general discussion. Section 2.5 concludes.

2.1 Motivating Evidence and Conceptual Framework

This section presents motivating evidence for studying the distribution of public goods (section 2.1.1) and introduces the general framework used in the paper (section

presentation of other studies following the DINA methodology.

[5]For instance, Piketty, Saez, and Zucman (2018) allocate all non-health expenditure proportionally to posttax disposable income. Blanchet, Chancel, and Gethin (2022) consider two polar scenarios, one in which public goods are distributed proportionally to posttax disposable income, and one as a lump sum.

[6]See for instance Benhenda (2019), Lustig (2018), Paulus, Sutherland, and Tsakloglou (2010), Verbist, Förster, and Vaalavuo (2012), and Wagstaff et al. (2014) on education and health, Aaberge and Atkinson (2010) and Aaberge et al. (2019) on local government services, and Mladenka and Hill (1978) on police expenditure. To the best of my knowledge, O'Dea and Preston (2010) represents the only attempt at conceptualizing and providing guidelines on how all public services could be allocated to individuals (although they do not attempt to actually do so). My approach is largely inspired by theirs, and in many cases directly follows their recommendations.

2.1.2).

2.1.1 Motivating Evidence

I start by providing motivating evidence for incorporating estimates of public goods delivery in poverty and inequality statistics. I establish two simple stylized facts. First, public and private goods are substitutes: in countries with lower public goods provision, households tend to rely on market alternatives to a greater extent. Second, public goods have large effects on dimensions of well-being that are not captured by private consumption. As a result, standard poverty statistics underestimate poverty in countries with small welfare states relatively to those with higher public goods provision. They also tend to structurally underestimate the growth elasticity of poverty, given that economic growth allows governments to invest in public goods that are not recorded in private consumption.

2.1.1.1 Public and Private Goods Are Substitutes

The standard approach to measuring poverty and inequality focuses on household disposable income or household final consumption expenditure (disposable income minus savings). Disposable income is equal to the sum of labor and capital incomes, minus direct taxes paid, plus cash transfers received. By definition, it excludes public services, which amounts to implicitly assuming that their value to households is exactly zero.

This assumption can lead to implausible conclusions when analyzing the incidence of public policies on poverty. Consider for instance a government that decides to fully subsidize healthcare, effectively bringing down all private out-of-pocket healthcare expenditure to zero. Theoretically, individual incomes should be adjusted by adding the corresponding new in-kind transfer received by the government to their incomes. Yet, in the standard framework, poverty will remain unchanged, because the value of subsidized healthcare is recorded as being exactly zero. More generally, every policy subsidizing the provision of a good that was previously privately bought will be measured as having no incidence on poverty or inequality.

Figure 2.2a provides evidence that this channel is empirically relevant and quantitatively important. There is a strong negative correlation between the share of households pushed into extreme poverty by out-of-pocket healthcare expenditure and the size of public health spending across countries. In Bangladesh, where the government spends less than 0.5% of national income on health, 7% of the population see their daily expenditure fall below PPP $3.65 per day because of private health spending. Meanwhile, less than 0.3% of the South African population ends up poor because of out-of-pocket health spending, in a country where almost 7% of national income is spent on government-provided health services. Private and public expenditure are therefore not independent. In-kind transfers do allow poor households to save money, and not accounting for such money leads to overestimating poverty in countries with large welfare states.

2.1.1.2 Public Goods Matter for Non-Monetary Dimensions of Quality of Life

Public goods do not only matter for private consumption: they also contribute to improving non-monetary dimensions of well-being. The need to go beyond strictly monetary measures of poverty has been increasingly recognized in the past decades. Accordingly, researchers and international organizations have started developing a number of indicators of multidimensional poverty, which typically involve aggregating individual-level measures of well-being across a number of domains. For instance, Alkire, Kanagaratnam, and Suppa (2021) combine measures of deprivation in health, education, and access to a number of basic goods, each of which is assigned a weight of one-third.[7]

Figure 2.2b provides suggestive evidence that accounting for in-kind transfers contributes to bridging the gap between monetary and multidimensional poverty statistics. The x-axis plots general government expenditure on education, health, and housing and community amenities as a fraction of net national income. The y-axis represents the difference between the share of households living in multidimensional poverty and the share of households living in monetary poverty. There is a strong negative correlation between the two variables: multidimensional poverty is lower than monetary poverty in countries with large welfare states, while it is significantly higher in countries with low government expenditure. This suggests that in-kind transfers strongly improve the well-being of the global poor in dimensions of quality of life that are not captured by monetary poverty statistics.

The framework adopted in this paper can thus be viewed as one way of incorporating non-monetary dimensions of poverty in a monetary framework, through the value of the public services that largely determine them. The major advantage of this approach is its conceptual consistency with macroeconomic statistics. Unlike multidimensional

[7]More precisely, the index is constructed by attributing a weight of 1/3 to two health indicators (nutrition and child mortality), 1/3 to two education indicators (years of schooling and school attendance), and 1/3 to six "living standards" indicators (access to cooking fuel, sanitation, drinking water, electricity, housing, and basic assets.)

measures of poverty, it is based on an internationally agreed upon framework, the system of national accounts, which remains the most commonly used source for tracking incomes across countries and over time. Unlike classic monetary poverty measures, it accounts for all forms of government spending, which ensures that income estimates incorporate the large fraction of national incomes that is redistributed in the form of public goods.

2.1.2 Conceptual Framework

I propose to value public goods by combining data on their cost and their incidence throughout the income distribution. Consider individual i receiving pretax labor and capital income m_i, paying taxes $\tau(m_i)$, and receiving cash and in-kind transfers from the government $g(m_i)$. Her posttax income is:

$$y_i = m_i - \tau(m_i) + g(m_i) \tag{2.1}$$

The value of public goods received is defined as:

$$g(m_i) = \sum_j G^j \times \gamma^j(m_i) \tag{2.2}$$

G^j is a *cost* component equal to total government expenditure on function j (e.g., education).

$\gamma^j(m_i)$ is a *progressivity* component equal to the share of expenditure on function j received by individual i. By definition, $\gamma^j(m_i) \in [0, 1]$.

My benchmark estimates thus amount to valuing cash and in-kind transfers equally, in line with the approach adopted by the national accounts and the existing fiscal incidence literature (e.g., Lustig, 2018; Piketty, Saez, and Zucman, 2018). A natural concern is that governments may differ in their ability to provide public goods even after accounting for differences in cost of provision. I thus investigate the potential role played by differences in cost efficiency by introducing a third parameter into the estimation:

$$g(m_i) = \sum_j G^j \times \gamma^j(m_i) \times \theta^j(m_i) \tag{2.3}$$

With $\theta^j(m_i)$ a *productivity* component adjusting expenditure received by i for the quality of the service provided. It equals zero if the transfer is completely useless (for instance, if the value added of teachers at the school attended by i is exactly 0).

On the contrary, it may be greater than one if the government is more efficient than a benchmark production unit at providing a given service (for instance, if public schools are more cost-efficient than private schools). Hence, $\theta^j(m_i) \in [0, +\infty)$, and $\theta^j(m_i) = 1$ corresponds to the case in which public goods are valued at cost of provision.

Given difficulties at conceptualizing and measuring productivity (which explains why national accounts and GDP growth figures do not generally attempt to do so), I start by presenting results with $\theta^j(m_i) = 1$ in sections 2.2 and 2.3. I investigate the robustness of my results to departing from this assumption in section 2.4.

2.2 Methodology

I now turn to the methodology used to construct a new database on the provision of public services worldwide. I first cover the distribution of pretax income (section 2.2.1), followed by the estimation of cost (section 2.2.2) and progressivity (section 2.2.3). Table 2.1 provides summary statistics on the data sources and methodology used to distribute government expenditure.

2.2.1 Pretax Income

The starting point of the construction of the database consists in measuring the distribution of pretax income. Data on global pretax income inequality come from the World Inequality Database (Chancel and Piketty, 2021), which draws from studies combining surveys, tax, and national accounts data from various sources to build a new database on the distribution of income in all countries in the world since 1980. Average income in each country-year is scaled up to match net national income per capita: poverty and inequality statistics are consistent with macroeconomic growth rates. The concept of income observed is pretax national income, that is, income before accounting for the operation of the tax-and-transfer system, but after accounting for the operation of the pension and unemployment systems.

2.2.2 Cost G^j

The first step required to distribute public goods is to measure how much governments spend and on which types of policies. To do so, I build a new database on the level and composition of general government expenditure since 1980 by combining various data sources. My primary source for total expenditure as a share of GDP is Mauro et al. (2015), which I complement with other series from the IMF and the

IFPRI-SPEED database (Yu, Magalhaes, and Benin, 2015). For the composition of public spending, I primary rely on IMF series, which breakdown government expenditure by Classification of the Functions of Government (COFOG). I combine them with additional data on education, health, and social protection spending from the World Bank, the OECD, and the United Nations Economic Commission for Latin America and the Caribbean.

2.2.3 Progressivity $\gamma^j(m_i)$

2.2.3.1 Allocation Principles

Measuring the progressivity of public goods is conceptually and empirically challenging, given that their ultimate beneficiaries cannot always be unambiguously identified. I rely on two key allocation principles to estimate the distributional incidence of public goods, which directly follow the existing literature (e.g., Lustig, 2018; O'Dea and Preston, 2010). First, public services accrue to individuals based on who receives them at a given point in time. Second, public goods benefit households based on the price they would have to pay to benefit from this service if it was not provided as a public good. These two principles are necessary to ensure conceptual consistency with standard poverty and inequality statistics.

1) Cash Flow Principle First, I distribute public goods to individuals based on their beneficiaries at a given point in time. For instance, education spending is distributed to households who send their children to school, while health spending is distributed to individuals using more intensively the public healthcare system. This ensures that public goods are valued in a way that is conceptually consistent with standard fiscal incidence analysis, which focuses on the incidence of taxes and transfers over a given period. Put differently, public services are allocated in the same way as they would theoretically be if households were to receive a cash transfer at time t and immediately use it to buy the corresponding service on a private market.

Departing from this assumption would require moving away from the cross-sectional analysis that forms the basis of international poverty statistics. For instance, high-income earners may benefit from greater public education spending during their lifetime because of longer studies, which implies that education expenditure might be more unequally distributed than generally thought (although only modestly so: see Riedel and Holger, 2022). Yet, allocating education in this way would also conceptually require moving from the analysis of current income to that of permanent income, incorporating estimates of how much taxes individuals pay over their lifetime and how much cash transfers they receive. Unfortunately, available data does not allow for such a detailed analysis when studying the evolution of global poverty.

2) Equivalent Pricing Principle Second, public goods accrue to households based on the price that they *would have to pay* for the corresponding service, rather than the price they *would be willing to pay.* This ensures again that cash transfers and public goods are valued in a conceptually comparable way: if the household was to receive cash instead of the public good, it would have to pay the market price of the corresponding service to benefit from it, not the maximum value it would be willing to pay. Moving from income to welfare would require accounting for the unobserved value that consumers put on *both* market and public goods. Willingness to pay is higher than the observed price for all consumers located to the left of the demand curve, who would continue buying the good if its price was to marginally increase (e.g., Aaberge et al., 2019).

In line with standard poverty statistics, which focus on consumption and do not attempt to estimate the individual welfare value of each good bought by each household, I will thus distribute public goods based on who benefits more from them, rather than who might put greater or lower value on each type of service. For example, the welfare perspective would imply that high-income households might be willing to pay significantly more for education, because the real income gains that they would get from returns to schooling might be higher.[8] This would call for putting a greater value on each dollar of public education received by children from high-income parents. In contrast, assuming that the cost of providing education is the same across income groups, the income perspective implies that education should benefit households proportionally to the number of children attending school. Consistency with standard consumption aggregates thus requires allocating education proportionally to school attendance, not expected real income gains from schooling, because a household willing to send a child to school would have to pay the price of the school, not the price of its returns to schooling, if it was to buy the same service from a private provider.

2.2.3.2 Education

Public education spending represented about 4.4% of national income in the average country in 2019. Following the existing fiscal incidence literature (e.g., Lustig, 2018),

[8] If the return to schooling is proportional and constant (e.g., 10%) and children from high-income parents can expect to have greater income regardless of education, for instance, then the real expected gains from schooling will be higher from children from high-income parents than those from low-income parents.

I distribute education expenditure to individuals proportionally to school attendance of children in the household. The data source is a unique historical micro-database that I have contributed to construct in a companion paper in collaboration with the World Bank (Gethin, Kofi Tetteh Baah, and Lakner, forthcoming). The database consists of over 1,300 nationally representative surveys fielded in 155 countries from 1980 to 2021. It records detailed information on the structure of the household, school attendance, age, and total household income (or consumption). Based on this information, Gethin, Kofi Tetteh Baah, and Lakner (forthcoming) provide detailed indicators of inequality in access to education and intensity of use of the education system by per-capita household income decile and age in each country. Drawing on this database, I calculate the transfer received by decile d in country c at time t as:

$$\mathrm{G}_{dct}^{\mathrm{educ}} = n_{dct}^{pri} g_{dct}^{pri} + n_{dct}^{sec} g_{dct}^{sec} + n_{dct}^{ter} g_{dct}^{ter}$$

Where n_{dct}^{j} denotes the average number of children in school at level j, g_{dct}^{j} denotes average spending per child on function j, and $j \in \{pri, sec, ter\}$ refers to primary education, secondary education, and tertiary education, respectively. Data on the relative costs of primary, secondary, and tertiary education per child come from the World Bank's World Development Indicators. The number of children in school by level and per-capita household income is recorded in the Gethin, Kofi Tetteh Baah, and Lakner (forthcoming) database.[9] Finally, all transfers received in each country-year are proportionally rescaled to match total public education expenditure in the database constructed in section 2.2.2.

While this approach is straightforward and arguably captures first order differences in access to education, there are two potential sources of concern. The first one is that public education spending may vary not only by level, but also across subnational regions. In particular, poorer regions may benefit from lower spending, leading public education spending inequalities to be underestimated. The second concern is that these estimates do not account for children in private schools, which typically benefit from less (or no) public education spending. This will lead to overestimating public education spending inequalities, since children in private schools tend to disproportionately come from high-income households. I investigate the sensitivity of my results to these two concerns by comparing my estimates to those of the Commitment to Equity Institute (CEQ) Database. The CEQ compiles estimates of

[9]When missing, relative costs are assumed to have remained constant before or after the last year available. In the absence of detailed information on school attendance by grade in the microdata, individuals in primary school are taken as those aged 6 to 12, individuals in secondary school as those aged 12 to 18, and individuals in tertiary education as those aged above 18.

tax-and-transfer progressivity from a number of fiscal incidence studies following a comparable methodology (see Lustig, 2018). Education spending is allocated in the exact same way as above, except that these more detailed studies do exclude children in private schools when allocating transfers and generally also account for variations in spending by subnational region. The CEQ database provides this indicator for one or two years in 45 countries.

A comparison of the two datasets is displayed in appendix figure B.2, focusing on the share of public education spending received by the poorest 50%. The two estimates are strongly correlated, suggesting that the simplified methodology does succeed at capturing broad cross-country variations in education spending inequalities similar to those found by the CEQ. On average, my measures of the bottom 50% share of education spending are slightly lower, mainly because I do not exclude children in private schools from the allocation, while CEQ studies generally assume that they benefit from no public education subsidy at all.[10] This provides reassuring evidence that my estimates provide a good approximation, and if anything likely yield a lower bound on public education transfers received by the global poor.

2.2.3.3 Health

I distribute health expenditure (3.5% of NNI) proportionally to use intensity of the public healthcare system. Here, I rely directly on the CEQ database, which provides estimates of the distributional incidence of health expenditure from a number of studies. These estimates are typically constructed by using survey microdata covering indicators of frequency of use of public healthcare, such as the number of visits to a public health institution in the past month, or the total amount of user fees paid. These indicators are then aggregated at the household level to derive measures of healthcare use intensity by pretax income decile. The data cover 45 countries for one or two years in the 2010s.

2.2.3.4 Other Public Goods

Other expenditure includes spending on public order and safety (2% of NNI), transport and other economic affairs (5.8% of NNI), general public services (5.5% of NNI), and defense, housing and community, recreation and culture, and environmental protection (4.6%). In the absence of data on their distributional incidence, I make the conservative assumption that they are received by individuals proportionally to

[10]In practice, the government does contribute to the funding of private schools in many countries, although it usually provides lower funding than to public schools. The true transfer received thus likely falls in-between.

posttax disposable income, that is, in a highly unequal way. I view this as a lower bound. Indeed, there is a case for allocating some of these public services in a much more equal way: for instance, police services can be thought of benefiting households proportionally to the crimes that they experience (e.g., O'Dea and Preston, 2010), while housing policies include many public housing programs that disporportionally benefit low-income households. In Gethin (2023c), I provide evidence that under reasonable assumptions, nearly all in-kind transfers are more equally distributed than pretax income in the case of South Africa.

2.2.3.5 Other Dimensions of Redistributions: Social Assistance and Taxes

Finally, to have a complete perspective on the role of government redistribution in shaping poverty and inequality, I incorporate in the database estimates of the distributional incidence of social assistance and taxes.

Social Assistance I distribute social assistance expenditure (2.9% of NNI on average in 2019) to beneficiaries of cash transfers and in-kind social benefits. The main data sources are Piketty, Saez, and Zucman (2018) for the United States, Blanchet, Chancel, and Gethin (2022) for European countries, the CEQ database (40 countries), and the World Bank's ASPIRE database (108 countries). In each case, I only distribute social assistance expenditure and exclude pensions and unemployment benefits, given that these transfers are already included in estimates of the pretax income distribution (see section 2.2.1).

Taxes Finally, I allocate taxes in each country-year by combining data on total tax revenue with estimates of the distributional incidence of taxes. Aggregate data come from Bachas et al. (2022), who build a new database on the level and composition of tax revenue in 150 countries from 1965 to 2018. Data on the share of taxes paid by pretax income decile come from a companion paper (Fisher-Post and Gethin, 2023).

2.2.3.6 Imputation of Missing Data

I consider three scenarios for the distribution of public goods, cash transfers, and taxes in countries with missing data. In my benchmark scenario, I fill missing values with the average tax or transfer incidence profile observed in all country-years. I then consider an upper bound in which missing countries are attributed the average incidence profile of the five countries with the most progressive profiles, and a lower bound in which missing countries are attributed the profile of the five countries with the most regressive profiles.

2.2.3.7 Validation: Comparison With Detailed South African Series

Given the relative scarcity of data, especially when it comes to the time dimension, it is useful to get a sense of how accurately my estimates capture broad trends in government redistribution in countries where more detailed information exists. Appendix figure B.1 compares two estimates of the share of national income redistributed to the bottom 50% in the form of public goods in South Africa. The first one corresponds to the "simplified" series estimated in this paper, which exclusively rely on aggregate budget data from the IMF and the World Bank, estimates of the progressivity of education covering the 2002-2019 period from Gethin, Kofi Tetteh Baah, and Lakner (forthcoming), and estimates of the distribution of health spending for one year from the CEQ database (Goldman, Woolard, and Jellema, 2020). The second corresponds to "detailed" series constructed in Gethin (2023c). These series combine survey, census, and newly digitized budget data to allocate all public goods to individuals every year since 1993. Unlike simplified series, they cover each function of government in much greater detail, allowing for a precise allocation of local government spending, housing subsidies, public transport, transport infrastructure, police services, and different kinds of subsidies received by households. They cover the evolution of progressivity over time, while simplified series extrapolate the incidence of transfers from one year of data in the case of healthcare. They also account for variations in spending by province, while the simplified series do not.

Despite their limitations, simplified series appear to track remarkably well the evolution of redistribution in South Africa. In both simplified and detailed series, public services received by the bottom 50% are found to have significantly increased over time, from about 7% of national income in 2000 to 10-11% in 2019. If anything, simplified series do slightly underestimate the rise of redistribution, mainly because progressivity is assumed to have remained constant, while Gethin (2023c) finds that it has significantly increased across all functions of government. They also slightly underestimate redistribution in 2019, mainly because housing subsidies and local government expenditure are assumed to be distributed proportionally to posttax disposable income, while Gethin (2023c) finds them to be much more progressive. These results provide reassuring evidence that the simplified allocation developed in this paper provides a very good first-order approximation of levels and trends in government redistribution around the world.

2.3 Public Goods and the World Distribution of Income

This section presents the main results on the incidence of public goods on poverty and inequality across countries and in the world as a whole. Section 2.3.1 discusses cross-national variations in the size and progressivity of government redistribution around the world since 1980. Section 2.3.2 studies the incidence of public services on global poverty and inequality.

2.3.1 The Distribution of Public Goods Around the World

I start by exploiting my new database to document three stylized facts on the distribution of public goods. First, public goods are progressive: they systematically reduce inequality. Second, public goods have grown since 1980, in particular those public goods that are most progressive. Third, redistribution in the form of public goods correlates strongly with economic development: low-income countries spend less on public goods than high-income countries and in ways that are less progressive.

2.3.1.1 Public Goods are Progressive

Figure 2.3 plots the distribution of the progressivity of government redistribution across countries, measured as the share of total expenditure received by the bottom 50% (see also table 2.1). Education, healthcare, and social assistance are all *relatively* progressive (less concentrated than pretax income): they systematically reduce income inequality. However, there are significant variations both across categories and across countries within each category. In particular, cash transfers appear to be *absolutely progressive* in most countries: the bottom 50% receive on average a greater fraction of these transfers than their share in the population. Meanwhile, public education and healthcare tend to be slightly absolutely regressive: higher-income earners benefit from greater transfers than low-income groups.

Social assistance is the most progressive function of government, due to the often explicitly pro-poor design of the corresponding programs (such as conditional cash transfers or food stamps). On average, the bottom 50% receives about 64% of social assistance expenditure, However, there are large variations across countries, with the share of social assistance transfers accruing to the bottom 50% ranging from only 16% in Haiti to as much as 92% in Peru.

Education is less progressive than cash transfers but still substantially reduces inequality, as it falls close to a lump sum allocation. In the majority of countries, the bottom 50% benefit from 45-50% of public education expenditure. This figure is the product of two countervailing forces. On the one hand, inequality in access to schooling implies that children from high-income households tend to stay longer in school. The fact that spending per child is higher as higher levels of education reinforces these inequalities. On the other hand, fertility is often slightly higher among low-income households, which increases the progressivity of public education through a demographic effect. These two effects more or less compensate each other on average, yielding a quasi-egalitarian distribution of public education spending.[11]

Public healthcare is about as progressive as education in the average country, although there are significant variations. In some countries, low-income households use relatively less intensively the public healthcare system, partly because user fees may act as a barrier to access. In others, they do so to a greater extent, partly because they suffer from poorer health, and partly because high-income households tend to rely on private healthcare services to a greater extent.

Combining social assistance, education, healthcare, and other public services distributed proportionally to posttax disposable income implies that about 30% of total government expenditure ends up accruing to the bottom 50% in the average country. In nearly all countries in the world, government transfers are relatively progressive (less concentrated than income), but absolutely regressive (accruing in greater proportion to the rich than to the poor). There are large variations in the progressivity of expenditure, with the bottom 50% share of total spending varying from only 15-16% (Angola, Somalia, Republic of the Congo) to almost 50% (Denmark, Sweden, United Kingdom).

2.3.1.2 Public Goods Have Grown

The second stylized fact is that governments have dedicated growing resources to public services in the past decades. Between 1980 and 2019, average general government expenditure as a share of national income increased from about 26% to 29% (see table 2.1). This rise cannot be explained by cash transfers: social assistance spending almost stagnated at about 2.5-3% of NNI on average, which represents about 10% of total government expenditure. Much of the rise of government intervention was instead driven by significant increases in public goods, and especially education and

[11]As discussed in section 2.2.3, in the absence of data, I make the conservative assumption that all children in school benefit from public education spending. If one was to exclude children in private schools, public education would be more progressive, because private schools are used much more intensively by high-income households (Lustig, 2018).

healthcare. Meanwhile, expenditure on economic affairs and general public services slightly declined. Overall, net national incomes increased significantly, leading public services to expand considerably in real value in the world as a whole. As shown in figure 2.4, real government expenditure per world citizen approximately doubled from 1980 to 2019.

Figure 2.4 breaks down the evolution of government expenditure on social assistance and public services by country income group from 1980 to 2019. There are three main results.

First, low-income countries spend significantly less on both social assistance and public goods as a share of national income than high-income countries. In 2019, total expenditure amounted to about 24% of national income in low-income countries, 27-28% in middle-income countries, and 36% in high-income countries. Poorer countries also dedicate a lower fraction of total expenditure to education and healthcare. Less than 6% of national income is spent on public education and health in low-income countries, compared to almost 15% in high-income countries. Meanwhile, low-income countries actually dedicate a greater share of national income to other public goods than high-income countries (17% versus 15%).

Second, there has been a slight convergence in public goods provision between countries with different levels of economic development. Total expenditure on public services expanded by about 4 percentage points in low-income countries and 6 percentage points in lower-middle-income countries, compared to about 3 percentage points in high-income countries. It stagnated in upper-middle-income countries, mainly because total expenditure as a share of national income was approximately the same in China in 2019 as in 1980.

Third, there has been a general trend towards devoting greater resources in the most progressive forms of public goods. Regardless of the level of economic development, spending on education and healthcare expanded as a share of national income. In contrast, expenditure on other public goods declined in upper-middle-income and high-income countries and stagnated in lower-middle-income countries. Low-income countries stand out as having invested about as much in education and healthcare as in other public goods.

Combining these results with cross-country differences in macroeconomic growth, middle-income countries appear to have seen expenditure on public services increase most significantly, by about 180% in real terms from 1980 to 2019, mainly due to the rise of China and India.[12] In high-income countries, public goods have expanded almost two times slower, by about 100% from 1980 to 2019. In low-income countries, finally, real expenditure on public goods has almost stagnated, mainly due to exceptionally low or even negative growth in the poorest countries in the world. Differences in public expenditure remain substantial in 2019, with average spending on public goods reaching almost $15,000 at purchasing power parity in high-income countries, about three times more than in upper-middle-income countries, and over thirty times more than in low-income countries.[13]

2.3.1.3 Low-Income Countries Score Lower on All Dimensions of Redistribution

The third stylized fact is that there are large variations in redistribution in the form of public goods, which correlate strongly with economic development. Figure 2.5a maps the share of national income received by the bottom 50% in 2019 around the world. Progressive spending on public goods is generally highest in North America and Western Europe, exceeding more than 8% of national income in most countries. It is also relatively high in South America and Southern Africa, where some countries redistribute similar or even higher shares of national income to the bottom 50% than in Western countries. Public goods provision is significantly lower in Asia: less than 7% of national income is received by the poorest half of the population in most countries. Finally, in-kind redistribution is lowest in Western, Central, and Eastern Africa, where it often falls below 4% of national income.

Figure 2.5b plots the level and composition of public services received by the bottom 50% in fifteen selected countries or regions, which together represented about two-thirds of the world's population in 2019. There are huge differences in in-kind redistribution to the bottom 50% across countries. In Bangladesh, Nigeria, and Indonesia, only about 4% of national income is received by the poorest half of the population in the form of public goods. The corresponding figure exceeds 11% in Western Europe and the United States. Redistribution in the US is slightly higher than in Western Europe, in line with the findings of Blanchet, Chancel, and Gethin (2022). It is also interesting to note that most differences in in-kind redistribution across countries can be explained by spending on education and health. Less than

[12] See appendix figure B.26, which plots real expenditure on public goods by country income group. Figure B.27 plots the same figures expressed as a share of each country's national income.

[13] Appendix figures B.28 and B.29 plot the corresponding series by world region. Spending on public goods has increased in all regions of the world. This rise has been most pronounced in China and India, and lowest in Africa. See also figure B.40, which maps changes in general government expenditure as a share of national income in each country from 1980 to 2019.

1% of national income is received in public education and healthcare by the bottom half of the population in Bangladesh and Nigeria. The corresponding figures are higher than 6% in Brazil, South Africa, Western Europe, and the United States.

While these differences arise from combining data on the size and progressivity of government expenditure in each country, they generally extend to each of these parameters taken separately. Table 2.2 decomposes the distribution of public goods into these two drivers by country income group and world region. Both dimensions of redistribution increase significantly with economic development. Total expenditure on public goods is about 30% of national income in high-income countries, compared to 23% in low-income countries. 33% of spending accrues to the bottom 50% in the former group, compared to 23% in the latter. Combining these parameters, the bottom 50% ends up benefiting from only 5% of national income in the form of public goods in low-income countries, about two times lower than in high-income countries. Poor countries thus not only invest less in public goods than rich countries; they also provide them much more unequally than in the rich world.[14]

Similarly, variations in overall redistribution across geographical regions tend to be reproduced across different dimensions of redistribution. African and Asian countries display significantly lower levels of general government expenditure as a share of national income. They also tend to invest a lower fraction of that expenditure in education and health, the two most progressive functions of government.[15] China and India stand out as interesting cases. Expenditure on public goods is higher in India, but redistribution is operated in a significantly more progressive way in China. As a result, both countries end up redistributing about 6% of their national incomes to the bottom 50%.

2.3.2 Public Goods and Global Economic Growth

I now turn to analyzing the incidence of public goods on the distribution of global economic growth since 1980. I first show that public goods have played a major role in making global economic growth more progressive. I then analyze the incidence

[14] Appendix table B.1 reports pairwise correlation coefficients between dimensions of redistribution and net national income per capita across countries, including measures of productivity discussed in section 2.4. Nearly all dimensions of redistribution are significantly positively correlated: countries spending less also spend in more regressive and more inefficient ways. All four parameters are also positively correlated with economic development, in particular progressivity ($\rho = 0.7$) and aggregate productivity ($\rho = 0.65$).

[15] Online appendix figure B.39 maps general government expenditure as a share of NNI around the world in 2019. Figure B.41 plots the share of education and health spending in the government budget.

of public goods on global inequality. Finally, I decompose redistribution into its different components.

2.3.2.1 Public Goods and Global Poverty Reduction

To what extent has the rise of public goods contributed to the decline of global poverty? Figure 2.1 plots the evolution of the global poverty headcount ratio at $2.15 per day, expressed in 2017 PPP USD, before and after accounting for cash transfers and public goods. Following the distributional national accounts methodology (Piketty, Saez, and Zucman, 2018), I compare three concepts of income: pretax national income, posttax disposable income, and posttax national income. Posttax disposable income removes direct taxes from pretax income and adds cash transfers, which corresponds to the standard concept used to measure poverty. Posttax national income removes all taxes, including indirect taxes, and adds all government expenditure, which ensures that average incomes are consistent with net national income growth. Global poverty has declined by about 43% in terms of pretax income, from 23% in 1980 to 13% in 2019. Adding cash transfers lifts about 2% of the world population out of poverty. It also increases the rate of poverty reduction since 1980 to 50%. Finally, adding public goods further reduces poverty by about 4 percentage points, and yields a total rate of global poverty decline of 63%. Hence, government redistribution contributes to reducing the global poverty rate by about a third today, and it has contributed to accelerating the rate of global poverty decline since 1980 by almost 50%. About two-thirds of these effects are driven by public services. Overall, they have contributed to about 20% of the decline in global poverty.[16]

The key role played by public goods in global poverty reduction can mainly be explained by the rise of public education and healthcare services, which have increasingly accrued to the global poor in the past decades. Figure 2.6 plots the level and composition of public services received by global bottom 20% since 1980, expressed as a share of total global income. Although redistribution to the bottom 20% in the form of public goods remains extremely low, it has steadily increased in the past decades: about 0.8% of global GDP was redistributed to the global income quintile in 2019, compared to 0.35% in 1980. The bulk of these gains was driven by education and healthcare, whose value was multiplied by three, from about 0.2% to

[16]It is important to stress that these ratios depend on which transfer is allocated first. If one was to first allocate public goods and then cash transfers, then the contribution of the former would appear substantially higher in comparison to the latter. In allocating public goods after cash transfers throughout the paper, I provide a lower bound on the contribution of in-kind transfers relative to cash transfers.

0.6% of global income. In 2019, they represented over two-thirds of public goods received by the bottom 20%. The combination of increased redistribution with global GDP growth has implied large gains in the real value of public services received by the global poor. From 1980 to 2019, the per capita transfer received by the global bottom 20% was multiplied by about 4, growing from only $30 to $120 per year at purchasing power parity.[17]

An alternative way of looking at the role of public goods in shaping global poverty reduction is to compare the growth rates of specific groups before and after accounting for public goods. Figure 2.7 plots the real average income of the world's poorest 20% before and after cash and in-kind transfers. The global bottom 20% average income approximately doubled in terms of pretax income. Adding cash transfers increases this growth rate to over 130%, while incorporating education, health, and other in-kind transfers raises it further to 170%. By this view, cash transfers account for about 25% of global bottom 20% growth, public goods account for 20%, and total transfers account for as much as 40%. Appendix figure B.8 extends this analysis to the world's poorest 50% individuals, with similar conclusions.[18]

My main result is robust to polar assumptions on the distribution of public goods. On the one hand, one may argue that only education and health eventually accrue to the poor, while other forms of public goods have little value and mostly benefit richer households. On the other hand, there is a case to make for an egalitarian allocation of collective public goods. After all, poorer households do indirectly benefit from services as diverse as street lighting, post offices, environmental protection, local and national administrations, and garbage removal in many countries around the world. Appendix table B.2 shows how sensitive is my result on global poverty reduction to these two scenarios. In my benchmark estimates, accounting for public services increases the rate of poverty reduction from 50% to 63%. Restricting public goods to education and health leaves this result unchanged. Assuming that all collective public goods are received on a lump sum basis raises the rate of poverty reduction even further, to 79%. My main conclusion is thus relatively robust to different scenarios on the progressivity of other public goods: public services account for 20-30% of global poverty reduction since 1980 and potentially more.

A second concern is that my findings might be driven by a specific country. The

[17]See appendix figure B.13, which plots the per capita real value of public services received by the global bottom 20% in 2021 PPP US dollars. Appendix figures B.12 and B.14 plot the same figures for the global bottom 50%. The results are broadly similar, although education and healthcare represent a slightly smaller fraction of transfers received.

[18]More specifically, transfers account for about 20% of real bottom 50% income growth, about 15 points of which is due to public goods and 5% to cash transfers.

obvious candidates are China and India, which together represent over at third of the world's population and have both significantly invested in public services in the past decades. Appendix table B.2 reproduces my results on global poverty reduction after excluding China, after excluding India, and after excluding both countries from the sample. The results are qualitatively similar: public services account for about 15% of global poverty reduction when excluding China, 30% when excluding India, and 25% when excluding both countries.

Finally, I investigate the sensitivity of my results to using World Bank data instead of data from the World Inequality Database.[19] The World Bank data cover consumption or posttax disposable income per capita distributions that are not consistent with growth rates reported in the national accounts, so it is not the most adequate data source to study the impact of government redistribution on poverty and inequality. I attempt to reconstruct measures of pretax and posttax income nonetheless, using data available on the World Bank's website.[20] The main results are presented in appendix figures B.5, B.6, and B.7 for poverty thresholds at \$2.15, \$3.65, and \$6.85 per day. The results are qualitatively similar to the main findings presented above: redistribution is found to have accelerated global poverty reduction at all thresholds.[21]

2.3.2.2 Public Goods and Global Inequality

I now turn to analyzing the incidence of public goods on global income inequality and the distribution of global economic growth.

Figure 2.8a plots the real income growth rate experienced by each global income

[19]Both the levels and trends in global poverty in the WID data differ from those of the World Bank for at least four main reasons. First, World Bank estimates focus on consumption (posttax disposable income minus net household saving), while my focus here is on income. Second, the estimates presented here are consistent with national income growth rates, while World Bank estimates are based on surveys and do not attempt to bridge gaps between survey and national accounts aggregates. Third, some of the estimates used in this paper are based on studies relying on data sources that may differ from those of the World Bank in a number of countries, including China (Piketty, Yang, and Zucman, 2019), India (Chancel and Piketty, 2019), and Brazil (Morgan, 2017). See Chancel and Piketty (2021). Fourth, I use GDP purchasing power parity conversion factors, while the World Bank only corrects for price differences in household final consumption expenditure.

[20]The World Bank does not publish data on the world distribution of income. I thus reconstruct it myself by collecting distributions from the World Bank's website and extrapolating the average income of each country-percentile to missing years using real GDP per capita growth rates. This yields trends in global poverty almost identical to those officially reported by the World Bank. Finally, I reconstruct measures of pretax income as consumption or disposable income, minus cash transfers, plus direct taxes.

[21]Poverty at \$2.15 per day declined by already 77% in terms of pretax income, so it is unsurprisingly difficult to explain much more of poverty reduction with government redistribution. For the two other thresholds, public goods account for a substantial fraction of poverty reduction (about 20% at \$3.65 per day and 30% at \$6.85 per day)

percentile from 1980 to 2019. As is well-known (e.g., Chancel and Piketty, 2021; Lakner and Milanovic, 2016), the distribution of global economic growth has taken the shape of an "elephant curve," being highest at the middle of the global income distribution, lowest for the global upper-middle class, and relatively high among the richest 1%. Yet, little is known of how changes in government redistribution have shaped this general fact. My new database allows for the first time to make progress in answering that question. As shown in figure 2.8a, the distribution of global income growth has been relatively similar in terms of pretax and posttax disposable income. Higher cash transfers have led to negligible increases in growth rates at the bottom, financed by higher direct taxes paid by global middle- and top-income groups. By this measure, which corresponds to the standard way of studying the incidence of government policies on poverty, redistribution has done little to increase real incomes at the bottom since 1980.

In contrast to cash transfers, public goods have played an important role in making global economic growth more inclusive. The upper line of figure 2.8a adds public goods to the analysis and removes all taxes so as to reach posttax national income. Moving from posttax disposable income to posttax national income shifts the total growth rate of the 10^{th} percentile from about 100% to 160%. All percentiles within the bottom 60% see their growth rate rise substantially.[22] While in terms of posttax disposable income, most percentiles within the bottom 20% grew at a rate lower than that of the top 1%, the opposite is true in terms of posttax national income. Public goods thus appear to have been a major force of inclusive growth since 1980.

Figure 2.8b represents the evolution of global income inequality since 1980, measured as the ratio of the average income of the top 10% to that of the bottom 50% in the world as a whole, for different income concepts. There are two main results.

First, taxes and transfers significantly reduce global inequality: in 2019, the top 10% to bottom 50% income ratio was 39 in terms of pretax income, compared to 26 in terms of posttax national income. Taxes and transfers all contribute to reducing global inequality, but transfers have the strongest impact. Indeed, cash transfers reduce the indicator by about 5 percentage points; adding in-kind transfers further decreases it by 5 percentage points; finally, removing taxes pushes it down by 4 percentage points. By this measure, transfers account for about 70% of the impact

[22] Appendix figure B.9 compares the growth rates of average disposable income and of public services received by global income percentile from 1980 to 2019. In line with the results presented above, public goods have grown significantly faster than posttax disposable incomes, especially at the bottom of the global income distribution. Total disposable income growth ranges from 80% to 180% within the global bottom 50%, while total growth in public services received ranges from 220% to 360%.

of government redistribution on global inequality, while taxes account for about 30%.

Second, public goods have been the strongest driver of the rise of global government redistribution since 1980. Global pretax income inequality has fallen in the past decades: the richest decile earned 53 times more than the poorest half of the world's population in 1980 compared to 39 times today, amounting to a 26% decline. The corresponding figures are 30% after cash transfers, 37% after cash and in-kind transfers, and 37% after all taxes and transfers. In other words, accounting for government redistribution increases the total decline in global income disparities since 1980 by 40%. About two-thirds of this effect is driven by public goods.[23]

Because public goods provision varies so widely across countries, it does not only affect poverty and inequality in the world as a whole: it also shapes their distribution across space.

The upper panel of table 2.3 provides a Theil decomposition of global inequality into its between-country and within-country components for the main income concepts of interest. In line with the results presented above, taxes and transfers reduce global inequality: the Theil index is 1.13 in terms of pretax income, 0.98 in terms of posttax disposable income (or 13% lower), and 0.8 in terms of posttax national income (or 29% lower). However, because poor countries tend to have less progressive tax-and-transfer systems, and because redistribution only reduces inequality within countries, it increases the share of global income disparities explained by inequality between countries. The between-country component accounts for 30% of global inequality in terms of pretax income, but 33% in terms of posttax disposable income, and as much as 39% in terms of posttax national income. Accounting for government redistribution, in particular public goods, thus increases the weight of national differences in net national incomes per capita in explaining global inequality.[24]

The lower panel of table 2.3 focuses more specifically on the bottom of the distribution by breaking down the geographical location of the world's poorest 20% by world region. Accounting for government redistribution significantly increases the share of the global poor living in India, Pakistan, Bangladesh, Ethiopia, Nigeria, and other Sub-Saharan African countries, all of which were identified previously as having weak and regressive tax-and-transfer systems. On the contrary, it improves the relative

[23]Appendix figures B.15 and B.16 plot the evolution of the Gini and Theil indices of global inequality for different income concepts. The conclusions are qualitatively similar: for both indicators, accounting for public goods leads to a faster decline in global income disparities since 1980.

[24]Appendix figure B.17 plots the share of global inequality explained by average income differences between countries from 1980 to 2019 for different income concepts. This share has significantly declined across all income concepts.

positions of low-income individuals living in China, Latin America, and the Western world. These differences are quantitatively large. For instance, moving from pretax income to posttax national income increases the share of the global bottom quintile living in India from 18% to 24%, while this share drops from 7% to almost zero in Western Europe and North America.

In the end, lower redistribution in low-income countries translates into huge inequalities in the quality of public services received around the world. In 2019, public goods benefiting the poorest 10% of the world's citizens represented less than 0.5% of global GDP. The share of global GDP received by global bottom 50% as a whole increased significantly throughout the period, from about 1.5% to 3.5% of global GDP, mainly due to greater education and health transfers. However, this remains extremely small in comparison to the quality of services enjoyed by the richest world citizens: in 2019, public goods received by the upper decile of the global income distribution amounted to over 10% of global GDP. [25]

2.3.2.3 Decomposing Redistribution

Figure 2.9 further breaks down the incidence of government redistribution by showing how the global poverty rate behaves under a number of counterfactual scenarios on the size and progressivity of taxes and transfers.[26] The three leftmost bars show that taxes and transfers reduce global poverty from about 13% to 7%, as in figure 2.1.

The next bar considers a radical scenario in which government expenditure would be distributed on a lump sum basis, that is, in a perfectly egalitarian way ($\gamma^j(m_i) = \gamma$). This would reduce global poverty by 4 percentage points. This large effect is consistent with the significant inequalities in the distribution of public goods documented above and the fact that these inequalities are particularly high in poor countries, which spend less on the types of public goods that are most progressive. The last two bars further impose that all countries in the world move to a "Nordic welfare state," redistributing 50% of their national income, and that the global poor do not have to pay taxes to finance this expenditure. Moving to a Nordic welfare state would have a large effect on global poverty, while removing taxes would reduce it only marginally. This finding is consistent with the fact that both taxes and transfers are substantially lower in poor countries than in the rich world. Overall, applying all these scenarios jointly would reduce the global poverty rate from about 7% to below 1%.

[25] See appendix figures B.10 and B.11.

[26] Appendix figures B.22 and B.23 extend this analysis to poverty at \$3.65 and \$6.85 per day and also incorporate scenarios on public sector productivity discussed in section 2.4.1. The results are similar.

In summary, about 3-6% of the world's population falls below the poverty line because of inequalities in access to public services. Equalizing the distribution of transfers and increasing government capacity would have the biggest incidence on global poverty, followed by improving tax progressivity. Even under extreme scenarios on the size and progressivity of government transfers, however, the global poverty rate would still reach about 1%. This points to the roles of both cross-country macroeconomic convergence and reductions in pretax income inequality within countries as necessary complementary factors for improving the living conditions of the global poor.

2.4 Discussion and Extensions

This section briefly discusses some implications of the results presented in this article and avenues for future research. Section 2.4.1 explores the robustness of my results to accounting for public sector productivity. Section 2.4.2 investigates how accounting for public services can shed new light on a key debate in development economics: whether surveys or national accounts should be used to track poverty and economic development. Section 2.4.3 discusses challenges in moving from measures of consumption to measures of the welfare value of public goods. Section 2.4.4 explores the potential of my new measures of public goods redistribution for the study of the political economy of inequality.

2.4.1 Accounting for Public Sector Productivity

A natural concern is that cost of provision may not be an accurate indicator of the quality of public services received, because the productivity of governments may vary across time and space. In this section, I investigate the robustness of my results to adjusting transfers received for variations in public sector productivity. I focus on the main results and leave an extended presentation of the methodology to appendix B.2.

2.4.1.1 Methodology

Conceptual Framework I consider an extension in which the value of public goods is allowed to differ from cost of provision. The value of public goods received by individuals can theoretically be broken down into three components:

$$g(m_i) = \sum_j G^j \times \gamma^j(m_i) \times \theta^j(m_i) \tag{2.4}$$

With G^j government expenditure and $\gamma^j(m_i)$ the share of expenditure received by i. $\theta^j(m_i)$ captures the fact that for a given cost of provision, individuals may receive services of different quality. Empirically, it is useful to make a distinction between two notions of productivity:

$$\theta^j(m_i) = \Theta^j \times q^j(m_i) \tag{2.5}$$

Θ^j is the *aggregate productivity* of expenditure on function j, which does not depend on m_i. It captures the fact that the government may be more or less efficient at providing a given service than a benchmark production unit. For instance, public schools in country A may be on average less cost-efficient than public schools in country B, which implies that all public education transfers should be reduced by a constant factor in country A.

$q^j(m_i)$ is a *heterogeneous productivity* parameter. It captures the fact that the quality of services provided, holding cost constant, may differ between income groups. For instance, teachers teaching in poorer areas may be more or less qualified than those teaching in richer areas, independently from the wages they receive.

Aggregate Productivity Θ^j I propose to estimate the productivity of public education and healthcare by anchoring cost of provision to educational and health outcomes. For education, the outcome of interest is expected human capital at age 5, which I derive by combining data on school attendance and test scores from international databases. For health, the outcome is the healthcare access and quality index provided by the global burden of disease study (GBD, 2022), which ranks healthcare systems from 0 to 100 based on death rates from 32 causes of death that could be avoided by timely and effective medical care. I choose these indicators for two main reasons. First, they are among the only education and health indicators for which data is available for almost all countries in the world and with some time dimension. Second, they are relatively good measures of the output of the public sector, in contrast to other measures such as life expectancy, which are arguably more contaminated by unobserved factors.

I then compare these outcomes to spending on education and healthcare to derive measures of cost efficiency in each country-year. Appendix figures B.32 and B.33 provide a concrete illustration. Education spending per capita is strongly correlated with expected human capital, but there is also significant variation in educational outcomes for a given level of spending. Country-years that perform best for a given cost are attributed $\Theta^j = 1$: they are at the "efficient frontier". Meanwhile, country-years below the frontier are attributed lower values of Θ^j the further they are from the frontier. In this approach, no country-year has a score higher than 1, implying that the best government in the world is assumed to do just as well as the private sector and never better.

I discuss the limitations and implications of this approach in appendix 2.4.1. I view these estimates of productivity as a lower bound for three reasons. First, PPP conversion factors already make an adjustment for public sector productivity, so this approach holds the risk of "double-counting" inefficiencies (World Bank, 2013). Second, they imply necessarily reducing transfers in all countries that are not at the frontier ($\Theta \leq 1$). This is equivalent to assuming that governments are never more efficient than the private sector: absent any government, education and healthcare would be delivered at the same price or lower in any country-year. Third, omitted variable bias implies that productivity is likely to be underestimated in low-income countries, whose lower educational and health outcomes are arguably the product of other factors than government performance (such as lower income *per se*). That being said, I find that my measures of productivity correlate positively with existing indicators of government efficiency, which I view as reassuring evidence that this approach captures cross-country differences in public sector productivity relatively well.

Heterogeneous Productivity $q^j(m_i)$ Heterogeneous productivity is arguably even more challenging to estimate. In the absence of better data, I investigate using subjective perceptions of public services from international survey data to derive estimates of heterogeneous productivity by income group around the world. The data source is the Gallup World Poll, a yearly survey conducted since 2005 in 165 countries, which asks respondents whether they are satisfied with different types of public services in their area. I aggregate average responses by income quintile to measure differences in satisfaction with local public education, healthcare, police, and transport services. I then use relative responses as a scaling parameter, to increase or decrease the transfer received by each income group, for each of these four functions of government. This approach is arguably far from being satisfying. Nonetheless, existing empirical evidence on inequalities in service delivery (conditional on access) suggest that heterogeneous productivity is likely to be quantitatively modest (see appendix B.2).

2.4.1.2 Main Results

Cross-Country Differences in Redistribution The main takeaway is that accounting for productivity magnifies cross-country differences in redistribution. Appendix table B.3 extends table 2.2 to productivity-adjusted estimates. Mechanically, because $\Theta^j \leq 1$ by assumption, all countries end up redistributing a lower fraction of national income to the bottom 50%. The gap is particularly large in the case of low-income countries: in-kind transfers are reduced by 40%, compared to about 15% in high-income countries. By this view, poor countries thus suffer from a "triple curse" of redistribution in the form of public goods: not only do they spend less on public services and distribute them more unequally, they also provide them less efficiently. The gap in total spending between low-income and high-income countries is about 30% (23% versus 30%), while the gap in the value of the transfer eventually accruing to the bottom 50% exceeds 250% (3% versus 8%).

In terms of regional patterns, Asian and Sub-Saharan African countries are characterized by significantly lower aggregate and heterogeneous public sector productivity than Western countries and Latin America.[27] The China-India comparison is also striking: although India spends more on public goods than China, the productivity-adjusted transfer received by the bottom 50% ends up being a third lower in India than in China as a share of national income. This finding is consistent with the literature documenting the exceptionally low performance of the Indian public sector (Das et al., 2016; Muralidharan, 2019; Muralidharan and Sundararaman, 2015).

Global Poverty and Inequality I now turn to implications of productivity adjustments for the analysis of global poverty and inequality. The main conclusion is that accounting for productivity substantially reduces in-kind transfers received by the global poor. However, it does not significantly alter the trend; as a result, it only marginally affects my results on the role of public goods in reducing global poverty and inequality.

Appendix figure B.18 plots the share of global income received by the world's poorest 20% before and after adjusting for productivity. Adjusting for productivity reduces the total transfer received by the global bottom 20% by about a third but does not affect the trend. Productivity-adjusted estimates suggest that the share of global

[27] See online appendix figures B.43 and B.44, which map aggregate education and health productivity scores in each country. Figure B.47 maps average differences in satisfaction with public services across countries. In both dimensions, however, available data suggests that there has been a convergence over time: see figures B.45 and B.46 for aggregate productivity, and figure B.48 for income differences in satisfaction with public services.

income accruing to the poorest quintile rose from about 0.2% to over 0.5%.

Appendix figures B.19 and B.20 turn to global poverty reduction and the distribution of global economic growth. Adjusting for productivity reduces the rate of global poverty reduction from 63% to 60%. It also reduces the growth rate of percentiles within the bottom 50% of the world distribution of income by 5 to 10 percentage points out of growth rates ranging from 150 to 220. My main findings thus appear to be relatively robust to accounting for a potentially lower productivity of the public sector in low-income countries.

Finally, to get a sense of the importance of productivity in shaping the relationship between public goods and global poverty reduction today, appendix figure B.21 reproduces figure 2.9 with additional steps in which aggregate and heterogeneous productivity differences would be eliminated. With productivity adjustments, the poverty rate in 2019 is about 8% after accounting for all taxes and transfers. Removing heterogeneous productivity differences (setting $q^j(m_i) = 1$) would reduce poverty by less than half a percentage point. Remove aggregate productivity differences (setting $\Theta^j = 1$) would have a larger effect, reducing poverty by about one percentage point. These effects are significant, but still much lower than the effect of equalizing all transfers. These results suggest that improving productivity can be useful to reduce global poverty, but reducing inequalities in access to public services is likely to have quantitatively larger effects.

2.4.2 Surveys, National Accounts, and Public-Private Complementarities: Public Goods and Measurement Discrepancies in Poverty Statistics

A major debate in development economics centers around whether national accounts or surveys should be used in priority to measure economic development and poverty in the developing world. For reasons that continue to not be well understood, persistent discrepancies between GDP and survey incomes can lead to conflicting conclusions on the evolution of living standards in the past decades (Deaton, 2005).

Recent studies point to GDP as providing a better benchmark for tracking economic development than household surveys. Combining data from various sources, Pinkovskiy and Sala-i-Martin (2016) provide evidence that GDP correlates much more significantly to satellite-recorded nighttime lights than survey means. It also accounts for a much greater fraction of variations in a number of indicators of quality of life, such as life expectancy, access to safe water, and primary school enrollment.

Most importantly, the difference between GDP and survey means is positively associated with achievements on these indicators. In other words, "countries with higher and growing well-being tend to suffer from progressively greater mismeasurement of income by surveys." While the authors suggest that this finding could be due to the complexity of survey questionnaires, the exact reasons underlying this result remain unclear.

There is one natural candidate for explaining this discrepancy: public goods. As was made clear from the results presented in this article, surveys entirely miss services provided by governments in the form of education, health, transport, and other public services, which are not bought on a market and are thus absent from standard consumption measures. Arguably, these services play a key role in improving quality of life in the exact dimensions studied by Pinkovskiy and Sala-i-Martin (2016), as was already suggested in Figure 2.2b. The share of national income spent on public goods also appears to have significantly risen in the past decades, which could partly explain why surveys and GDP have become increasingly disconnected from each other.

I investigate this possibility in appendix table B.4. In the spirit of Pinkovskiy and Sala-i-Martin (2016), I regress five indicators of quality of life on the gap between GDP per capita and survey means: expected years of schooling, youth literacy, the secondary school enrollment rate, infant mortality, and life expectancy. I then compare the coefficient obtained before and after controlling for public spending on education and health, taken as a proxy for public goods provision in these two dimensions of well-being.

In line with Pinkovskiy and Sala-i-Martin (2016), I find that the gap between GDP and surveys tends to be positively correlated with greater quality of life, both before and after adding country fixed effects (panels A and B). For instance, a 1% increase in the gap between GDP per capita and average survey income is associated with a 0.16% increase in expected years of schooling. However, controlling for spending on education or health considerably reduces the size of the coefficient and renders it statistically non-significant in most specifications. Put differently, one of the main reasons why GDP estimates track indicators of quality of life better than surveys is that they incorporate consumption of public goods while surveys do not. In directly incorporating this "missing consumption" into poverty and inequality statistics, this article contributes to correcting some of the conceptual discrepancies between these two approaches to the measurement of living standards. My results also highlight the critical role played by public-private complementarities in global poverty reduction:

by enhancing public spending possibilities through greater tax revenue, GDP growth allows governments to increase public goods provision. Accounting for this channel, as was done in this article, leads to a more positive view of the role of macroeconomic growth in reducing poverty than the one pictured by household surveys alone.

2.4.3 From Consumption to Welfare: Challenges in Measuring the Value of Public Goods

A key limitation of the results presented in this article is that they do not account for how "valuable" public services actually are. While receiving free education might be useful to low-income households, it might not be as useful as receiving food or cash. In this section, I briefly discuss conceptual and empirical challenges in estimating the value of public goods.

There are at least three alternative ways of measuring the value of public services: through stated preferences, through revealed preferences, and through outcome-based estimation.

Stated Preferences Stated preferences refer to what households actually consider the value of public goods to be. To the best of my knowledge, only two studies have attempted to explicitly ask households whether they would prefer receiving cash than public services, and in what proportions (Khemani, Habyarimana, and Nooruddin, 2019; Thesmar and Landier, 2022).[28] In both cases, public services are found to be preferred to cash by a majority of households, in particular education and health.[29] By this measure, at least some public services should be attributed a greater value than cash transfers, which would reinforce my finding on the role of public goods in reducing global poverty.

Revealed Preferences Revealed preferences approaches use various methods to derive implicit measures of households' willingness to pay for public goods from behavioral patterns. The underlying principle is quite simple: if households receiving a cash transfer do not use it entirely to buy more education, then any increase in education spending should be attributed a lower value than a cash transfer of the

[28] See also Liscow and Pershing (2022), who test the preferences of US citizens for in-kind transfers compared to cash, but focusing on a basket of basic necessities, not on public goods.

[29] Thesmar and Landier (2022) ask respondents in France, Germany, and the United States to compare the actual composition of the government budget to the one they would prefer. They find clear majority support in favor of greater spending in education and health, and lower spending in cash transfers and defense. Khemani, Habyarimana, and Nooruddin (2019) perform a similar exercise in the context of Bihar, India.

same amount. This is a classic finding of economic theory (Atkinson and Stiglitz, 1976): cash transfers are superior to in-kind transfers, because they allow households to choose what they consume.[30] Based on this general result, public services should be attributed a significantly lower value than cash when being incorporated into poverty statistics, because they are not "freely chosen" by households.

Outcome-Based Measures Finally, outcome-based approaches value public services based on their actual effects. For instance, Finkelstein, Hendren, and Luttmer (2020) propose to measure the value of public policies by comparing the cost of each policy to total returns for its beneficiaries. In this context, the relative value of public services with respect to cash transfers depends on their ability to improve welfare. Focusing on 133 policy changes in the United States, Hendren and Sprung-Keyser (2020) provide evidence that investments in health and education targeted to low-income children display the highest marginal value of public funds, because they end up paying for themselves through substantial increases in earnings in later life. This would call for potentially putting a greater value on education and health expenditure than on cash transfers. Extending this approach to the study of global poverty would ideally require estimating the marginal value of an extra dollar spent in different types of public services in each country. These estimates could then be used to value public services by comparing their marginal value to that of cash transfers.

Understanding Discrepancies In a world with full information, perfect rationality, and perfectly competitive markets, these three measures of the value of public services should coincide, because households would be willing to pay a price equal to expected returns. However, this is rarely the case for at least three reasons.

First, many of the assumptions underlying the Atkinson-Stiglitz theorem do not hold in practice. Poor households may spend little on education and health not because returns are low, but because of many other factors such as limited information on their actual benefits, liquidity constraints, and market imperfections or spatial frictions that limit the supply of private education and healthcare services. All these factors are likely to lead to downward-biased estimates of willingness to pay when

[30] Another approach consists in using housing prices to derive implicit valuations of public services. For instance, Eshaghnia, Heckman, and Razavi (2021) find, drawing on granular data on housing prices and school characteristics in Denmark, that low- and high-income households are willing to pay a relatively similar fraction of their income for an increase in school quality (see Eshaghnia, Heckman, and Razavi (2021), Figure 4). By this measure, high-income households put a much greater monetary value on education than low-income households, which would imply distributing education spending in a more unequal way than done in this article.

measured from revealed preferences.

Second, individuals may value public services beyond the direct value that they get from consuming them. Support for government provision of services is not only dictated by personal benefits, but also strongly responds to beliefs about what constitutes a just society (Thesmar and Landier, 2022). Public goods may have positive externalities, such as lower inequality, of which individuals are well aware; knowledge of these externalities causally increases support for redistribution (Lobeck and Støstad, 2022). Outcome-based measures do not generally account for these externalities, which could lead to underestimating the true value of public services.

A third discrepancy comes from the fact that stated preferences may be subject to considerable measurement error, depending on the way questions are framed and other characteristics of survey design. As in the case of revealed preferences, individuals may also not be fully informed about how valuable public services are compared to one another and compared to cash. This makes it difficult to use stated preferences as a benchmark for valuing public goods.

All these inconsistencies make it difficult to evaluate the exact value that should be attributed to public services, both theoretically and empirically. This value ultimately depends on what one believes should matter, whether it is what individuals want (stated preferences), what they actually do (revealed preferences), or the benefits that they eventually get and what kinds of benefits are most important (outcome-based measures). Arguably, all of these three dimensions of welfare matter and should be studied jointly in future research.

2.4.4 The Correlates of Public Goods Provision: An Exploratory Analysis

I conclude this article with an exploratory analysis of the cross-country correlates of public goods redistribution. The objective is not to provide any new causal evidence, but merely to illustrate how the measures constructed in this article could contribute to shedding new light on the political economy of inequality. I hope that the methodology developed in this article, focusing not only on how much governments spend but also on how progressively and efficiently they do so, can inspire new studies on the different modalities through which public policies can reduce poverty. Combining subnational data on political outcomes with indicators on the size, progressivity, and productivity of public goods provision would be a particularly fruitful avenue for future research.

I investigate the correlates of redistribution in the form of public goods by combining my new measures with selected political and economic indicators available from international datasets. The outcome of interest is the share of national income received by the bottom 50% in the form of public goods in each country, computed from the database constructed in this article. I consider five explanatory variables. The first two capture political regime characteristics: the electoral democracy index available from the V-Dem database and the political competition index produced by the Polity5 project. The next two are measures of public sector corruption (V-Dem) and government effectiveness (World Bank Worldwide Governance Indicators), which relate more closely to the quality of governance. The last variable is the log of GDP per capita, expressed in 2021 PPP US dollars. All models control for the level of inequality, the total population, the demographic structure, and the trade to GDP ratio.

The results of this exercise are presented in appendix table B.5. The first three columns correspond to pooled OLS regressions on the full sample (column 1), the 2000-2019 sample only (column 2), and the 2000-2019 sample after excluding advanced Western democracies (column 3). Columns 3 to 6 repeat the same three specifications with country fixed effects.

Pooled OLS regressions point to the electoral democracy index, government effectiveness, and economic development as being significantly associated with redistribution. Electoral democracy and GDP per capita predict greater pro-poor spending on public goods, while government effectiveness has the opposite effect. The latter result might be driven by the fact that more effective governments spend less on public goods because they are able to provide them in more cost-efficient ways. Public sector corruption is associated with lower redistribution, but the effect is smaller and only statistically significant at the 10% level in the second specification.

Electoral democracy stands out as the only robust correlate of redistribution when adding country fixed effects. This effect is large, statistically significant, and relatively stable across specifications. Moving from the least democratic to the most democratic regime is associated with an increase in public goods received by the bottom 50% of 0.7 to 1.4 percentage points of national income. In contrast, political competition, public sector corruption, government effectiveness, and GDP per capita all display smaller and statistically non-significant coefficients in most specifications.

While these results are only suggestive and should be interpreted with care, they resonate well with the large literature pointing to the key role of political representation in fueling the rise of the welfare state (e.g., Cascio and Washington, 2014; Fowler, 2013; Fujiwara, 2015; Lindert, 1994; Meltzer and Richard, 1981). They are also in line with recent evidence ruling out the "luxury good hypothesis," according to which social protection would be a luxury good mechanically growing over the course of economic development (Lokshin, Ravallion, and Torre, 2022). After controlling for political variables and including country fixed effects, GDP per capita is not significantly associated with more or less redistribution.

2.5 Conclusion: Three Proposals to Improve Poverty Statistics

Public goods matter. They have been major drivers of human development in the past decades, contributing to improved access to education, healthcare, security, and other dimensions of quality of life. Yet, still little is known of who exactly benefits from these services, not only in a given country but even less so in the world as a whole. This article represented a first attempt at incorporating measures of public service delivery in global poverty statistics. I showed that doing so leads to a more positive view of global poverty reduction since 1980, because public goods are strongly progressive and governments have been increasingly investing in them. Nonetheless, the share of the world's GDP accruing to the global poor remains extremely limited, because low-income countries suffer from a curse of providing public goods in lower quantities, less progressively, and also potentially less efficiently than in the rich world. There is space for improvement in all three of these dimensions of government redistribution. Enhancing tax revenue, improving equity in access to public services, and raising government productivity should be seen as necessary complementary tools in the fight against global poverty.

This article has taken a large, global perspective on poverty reduction in the past decades, yet much remains to be done to better track public goods delivery around the world. First, there is an urgent need for more transparency on what governments actually do. The data exploited in this article cover spending on large categories, such as education or health, without much detail on the underlying policies. Unfortunately, information on these policies remains very limited; even when it exists, it often ends up buried under a multitude of documents published by different institutions. The publication of regular reports consolidating and harmonizing data on government budgets, with precise information on the corresponding policies, should be viewed as a priority not only for government accountability, but also for the measurement of global poverty. Too often, researchers and statistical institutes aiming to track living standards face no other option than to ignore public services, simply because of a critical lack of data on what these public services actually are.

Second, more attention should be given to public goods in the design of living standards surveys. Surveys routinely fielded by statistical institutes spend considerable time and effort compiling detailed data on household expenditure, yet the information that they collect on access to basic public goods remains rudimentary at best. Adding regular questions on both objective indicators and subjective perceptions of public service delivery would allow for a much more complete view of the well-being of low-income households. These questions should be designed in ways that would make them directly comparable with spending data on the different kinds of public services provided by governments.

Third, much more research should be conducted on how individuals actually value public services, not only in comparison to cash but also in comparison to one another. Under which conditions do households prefer to receive a transfer in the form of public healthcare, rather than education or cash? How do these priorities vary across countries, over time, and throughout the income distribution? Evidence on these questions remains extraordinarily scarce. Designing surveys eliciting such preferences would represent an important contribution to our understanding of the role of public services in reducing global poverty. Ideally, specific modules could be directly added to the questionnaires of living standards surveys, so as to regularly collect information on citizens' needs and priorities when it comes to public goods delivery.

Bibliography

Aaberge, R. and A. B. Atkinson (2010). “Top incomes in Norway”. In: *Top incomes: a global perspective*. Ed. by A. B. Atkinson and Thomas Piketty. Oxford University Press. Chap. 9, pp. 448–481.

Aaberge, Rolf et al. (2019). “Local governments, in-kind transfers, and economic inequality”. In: *Journal of Public Economics* 180.1, pp. 1–22.

Alkire, Sabina, Usha Kanagaratnam, and Nicolai Suppa (2021). “The Global Multidimensional Poverty Index (MPI) 2021”. Oxford Poverty and Human Development Initiative.

Atkinson, Anthony B. and Joseph E. Stiglitz (1976). “The Design of Tax Structure: Direct Versus Indirect Taxation”. In: *Journal of Public Economics* 6.1, pp. 55–75.

Bachas, Pierre et al. (2022). “Globalization and Factor Income Taxation”. NBER Working Paper 29819.

Benhenda, Asma (2019). “Teaching Staff Characteristics and Spending per Student in French Disadvantaged Schools”. Working Paper.

Blanchet, Thomas, Lucas Chancel, and Amory Gethin (2022). “Why Is Europe More Equal Than the United States?” In: *American Economic Journal: Applied Economics* 14.4, pp. 1–40.

Bourguignon, François and Christian Morrisson (2002). “Inequality among World Citizens: 1820-1992”. In: *American Economic Review* 92.4, pp. 727–744.

Bozio, Antoine et al. (2022). “Predistribution vs. Redistribution: Evidence from France and the U.S.” WID.world Working Paper.

Bruil, Arjan et al. (2022). “Inequality and Redistribution in the Netherlands”. Working Paper.

Cascio, Elizabeth U. and Ebonya Washington (2014). “Valuing the Vote: The Redistribution of Voting Rights and State Funds following the Voting Rights Act of 1965”. In: *Quarterly Journal of Economics* 129.1, pp. 379–433.

Chancel, Lucas and Thomas Piketty (2019). "Indian Income Inequality, 1922-2015: From British Raj to Billionaire Raj?" In: *Review of Income and Wealth* 65.S1, S33–S62.

– (2021). "Global Income Inequality, 1820–2020: the Persistence and Mutation of Extreme Inequality". In: *Journal of the European Economic Association* 19.6, 3025–3062.

Chancel, Lucas et al. (2022b). *World Inequality Report 2022*. Belknap Press of Harvard University Press.

Chen, Shaohua and Martin Ravallion (2010). "The Developing World Is Poorer Than We Thought, But No Less Successful In the Fight Against Poverty". In: *Quarterly Journal of Economics* 125.4, pp. 1577–1625.

Das, Jishnu et al. (2016). "Quality and Accountability in Health Care Delivery: Audit-Study Evidence from Primary Care in India". In: *American Economic Review* 106.12, pp. 3765–3799.

De Rosa, Mauricio, Ignacio Flores, and Marc Morgan (2022b). "More Unequal or Not as Rich? Revisiting the Latin American Exception". Stone Center for Inequality Studies Working Paper Series.

Deaton, Angus (2005). "Measuring Poverty in a Growing World (Or Measuring Growth in a Poor World)". In: *The Review of Economics and Statistics* 87.1, pp. 1–19.

– (2010). "Price Indexes, Inequality, and the Measurement of World Poverty". In: *American Economic Review* 100.1, pp. 5–34.

Eshaghnia, Sadegh, James J. Heckman, and Goya Razavi (2021). "The Willingness to Pay for School Quality, Neighborhood Attributes and Later Life Outcomes". Working Paper.

Finkelstein, Amy, Nathaniel Hendren, and Erzo F. P. Luttmer (2020). "Welfare Analysis Meets Causal Inference". In: *Journal of Economic Perspectives* 34.4, pp. 146–167.

Fisher-Post, Matthew and Amory Gethin (2023). "Government Redistribution and Development: Global Estimates of Tax-and-Transfer Progressivity, 1980-2019". Working Paper.

Fowler, Anthony (2013). "Electoral and Policy Consequences of Voter Turnout: Evidence from Compulsory Voting in Australia". In: *Quarterly Journal of Political Science* 8.1, pp. 159–182.

Fujiwara, Thomas (2015). "Voting Technology, Political Responsiveness, and Infant Health: Evidence from Brazil". In: *Econometrica* 83.2, pp. 423–464.

Gadenne, Lucie et al. (2022). "In-Kind Transfers As Insurance". Working Paper.

GBD (2022). “Assessing performance of the Healthcare Access and Quality Index, overall and by select age groups, for 204 countries and territories, 1990–2019: a systematic analysis from the Global Burden of Disease Study 2019”. In: *Lancet Glob Health* 10.12, pp. 1715–1743.

Germain, Jean-Marc et al. (2021). *Rapport du groupe d’experts sur la mesure des inégalités et de la redistribution*. Tech. rep. INSEE.

Gethin, Amory (2023a). “Distributional Growth Accounting: Education and the Reduction of Global Poverty, 1980-2022”. Working Paper.

– (2023c). “Who Benefits from Public Goods? Public Goods and Inequality in Post-Apartheid South Africa, 1993-2019”. Working Paper.

Gethin, Amory, Samuel Kofi Tetteh Baah, and Christoph Lakner (forthcoming). “Global Schooling Inequality, 1980-2022”. Working Paper.

Glassman, Brian (2019). “Multidimensional Deprivation in the United States: 2017”. American Community Survey Report.

Goldman, Maya, Ingrid Woolard, and Jon Jellema (2020). “The Impact of Taxes and Transfers on Poverty and Income Distribution in South Africa 2014/2015”. AFD Research Paper.

Hendren, Nathaniel and Ben Sprung-Keyser (2020). “A Unified Welfare Analysis of Government Policies”. In: *The Quarterly Journal of Economics* 135.3, pp. 1209–1318.

Jestl, Stefan and Emanuel List (2022). “Inequality, Redistribution, and the Financial Crisis: Evidence from Distributional National Accounts for Austria”. In: *The Review of Income and Wealth* forthcoming.

Khemani, Stuti, James Habyarimana, and Irfan Nooruddin (2019). “What Do Poor People Think About Direct Cash Transfers?” Brookings Future Development.

Lakner, Christoph and Branko Milanovic (2016). “Global Income Distribution: From the Fall of the Berlin Wall to the Great Recession”. In: *The World Bank Economic Review* 30.2, pp. 203–232.

Lindert, Peter H. (1994). “The Rise of Social Spending, 1880-1930”. In: *Explorations in Economic History* 31.1, pp. 1–37.

Liscow, Zachary and Abigail Pershing (2022). “Why Is So Much Redistribution In-Kind and Not in Cash? Evidence from a Survey Experiment”. In: *National Tax Journal* 75.2.

Lobeck, Max and Morten Nyborg Støstad (2022). “The Consequences of Inequality: Beliefs and Redistributive Preferences”. Working Paper.

Lokshin, Michael, Martin Ravallion, and Iván Torre (2022). “Is Social Protection a Luxury Good?” World Bank Policy Research Working Paper.

Lustig, Nora, ed. (2018). *Commitment to Equity Handbook: Estimating the Impact of Fiscal Policy on Inequality and Poverty*. Brookings Institution Press.

Mauro, Paolo et al. (2015). “A modern history of fiscal prudence and profligacy”. In: *Journal of Monetary Economics* 76.1, pp. 55–70.

Meltzer, Allan H. and Scott F. Richard (1981). “A Rational Theory of the Size of Government”. In: *Journal of Political Economy* 89.5, pp. 914–927.

Mladenka, Kenneth R. and Kim Quaile Hill (1978). “The Distribution of Urban Police Services”. In: *The Journal of Politics* 40.1, pp. 112–133.

Morgan, Marc (2017). “Extreme and Persistent Inequality: New Evidence for Brazil Combining National Accounts , Surveys and Fiscal Data, 2001-2015”.

Muralidharan, Karthik (2019). “The State and the Market in Education Provision: Evidence and the Way Ahead”. Columbia SIPA Working Paper 2019-06.

Muralidharan, Karthik and Venkatesh Sundararaman (2015). “The Aggregate Effect of School Choice: Evidence from a Two-Stage Experiment in India”. In: *The Quarterly Journal of Economics* 130.3, 1011–1066.

O’Dea, Cormac and Ian Preston (2010). “The distributional impact of public spending in the UK”. IFS Working Paper W12/06.

Paulus, Alari, Holly Sutherland, and Panos Tsakloglou (2010). “The distributional impact of in-kind public benefits in European countries”. In: *Journal of Policy Analysis and Management* 29.2, pp. 243–266.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman (2018). “Distributional National Accounts: Methods and Estimates for the United States”. In: *Quarterly Journal of Economics* 133.2, pp. 553–609.

Piketty, Thomas, Li Yang, and Gabriel Zucman (2019). “Capital Accumulation, Private Property, and Rising Inequality in China, 1978-2015”. In: *American Economic Review* 109.7, pp. 2469–2496.

Pinkovskiy, Maxim and Xavier Sala-i-Martin (2016). “Lights, Camera … Income! Illuminating the National Accounts-Household Surveys Debate”. In: *The Quarterly Journal of Economics* 131.2, pp. 579–631.

Ravallion, Martin (2012). “Why Don’t We See Poverty Convergence?” In: *American Economic Review* 102.1, pp. 504–523.

Riedel, Lukas and Stichnoth Holger (2022). “Government expenditure in the DINA framework: Allocation methods and consequences for post-tax income inequality”. ZEW Discussion Papers, No. 22-004.

Sala-i-Martin, Xavier (2006). “The World Distribution of Income: Falling Poverty and... Convergence, Period”. In: *The Quarterly Journal of Economics* 121.2, pp. 351–397.

Thesmar, David and Augustin Landier (2022). *Le prix de nos valeurs*. Flammarion.

Verbist, Gerlinde, Michael Förster, and Maria Vaalavuo (2012). "The Impact of Publicly Provided Services on the Distribution of Resources: Review of New Results and Methods". OECD Working Paper.

Wagstaff, Adam et al. (2014). "Who Benefits from Government Health Spending and Why?" World Bank Policy Research Working Paper 7044.

World Bank (2013). *Measuring the Real Size of the World Economy: The Framework, Methodology, and Results of the International Comparison Program – ICP*. The World Bank.

– (2018). *Poverty and Shared Prosperity 2018: Piecing Together the Poverty Puzzle*. Tech. rep.

Yu, Bingxin, Eduardo Magalhaes, and Samuel Benin (2015). "A Database User Manual for SPEED: Statistics on Public Expenditures for Economic Development". International Food Policy Research Institute.

Figure 2.1: Global Poverty and Public Goods: Global Poverty Headcount Ratio, 1980-2019

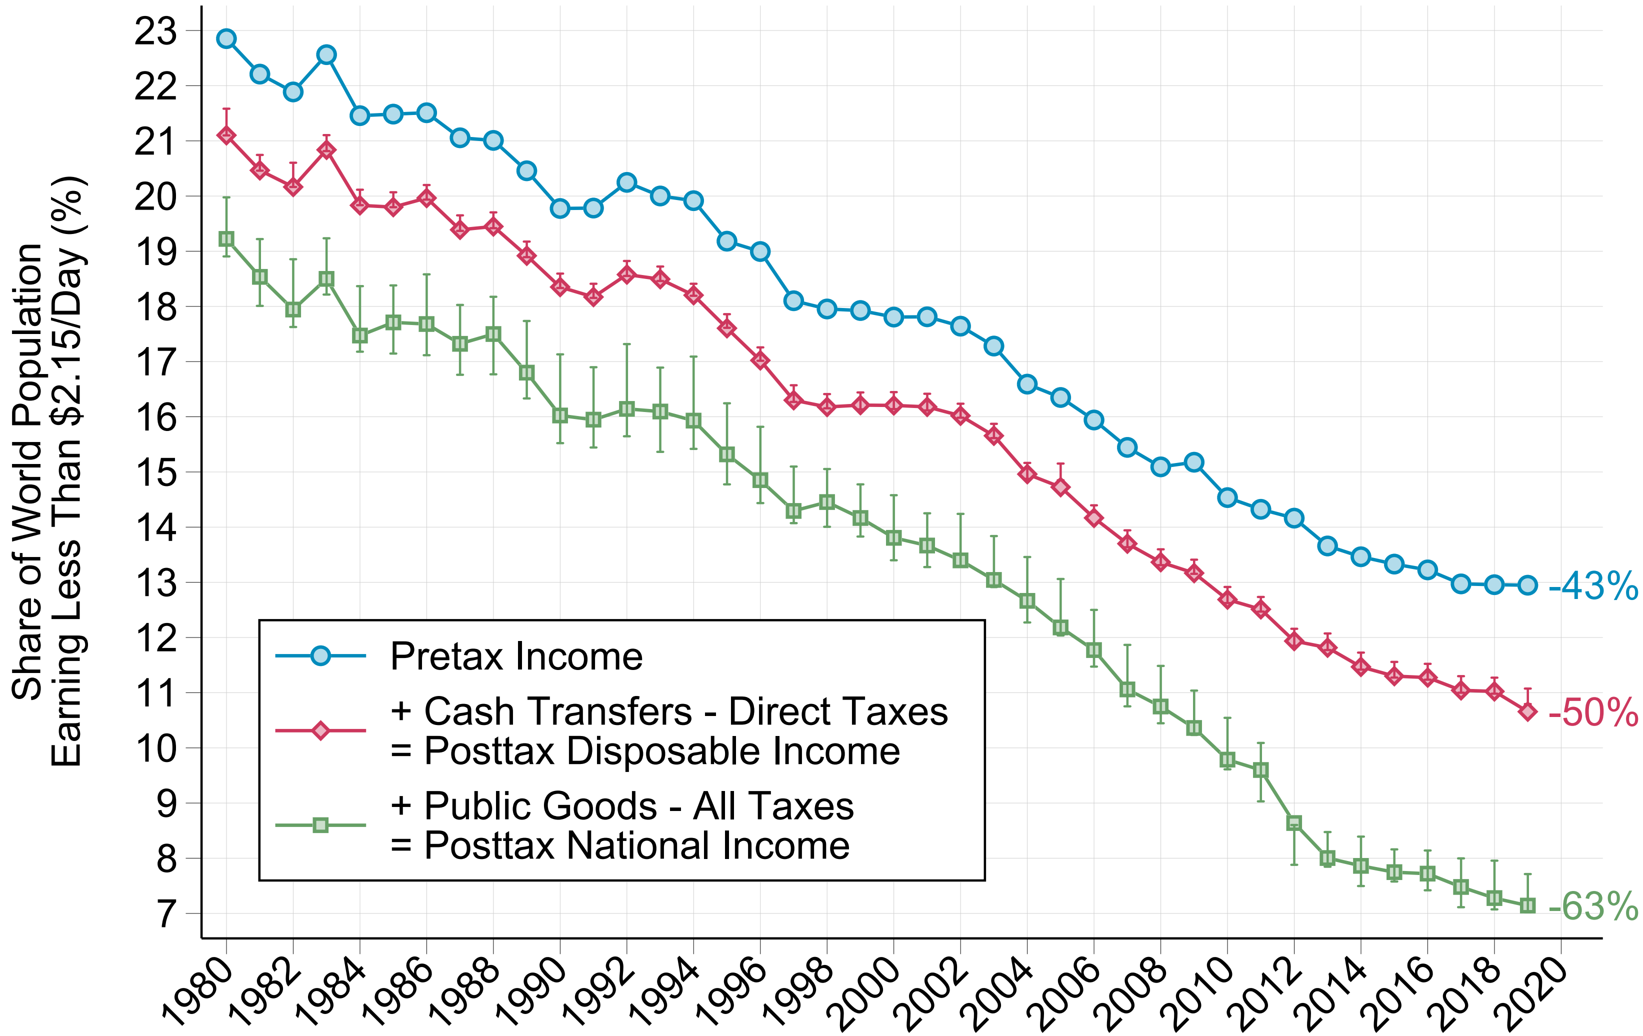

Notes. The figure plots the evolution of the poverty headcount ratio at $2.15 per day (2017 PPP USD) in the world as a whole, for different income concepts. Posttax disposable income removes direct taxes and adds cash transfers. Posttax national income removes all taxes and adds all cash and in-kind transfers. Spikes correspond to lower and upper scenarios on the distribution of transfers. The unit of observation is the individual. Income is split equally between all household members.

Figure 2.2: Public Goods and Poverty Measurement

(a) Public and Private Goods are Substitutes

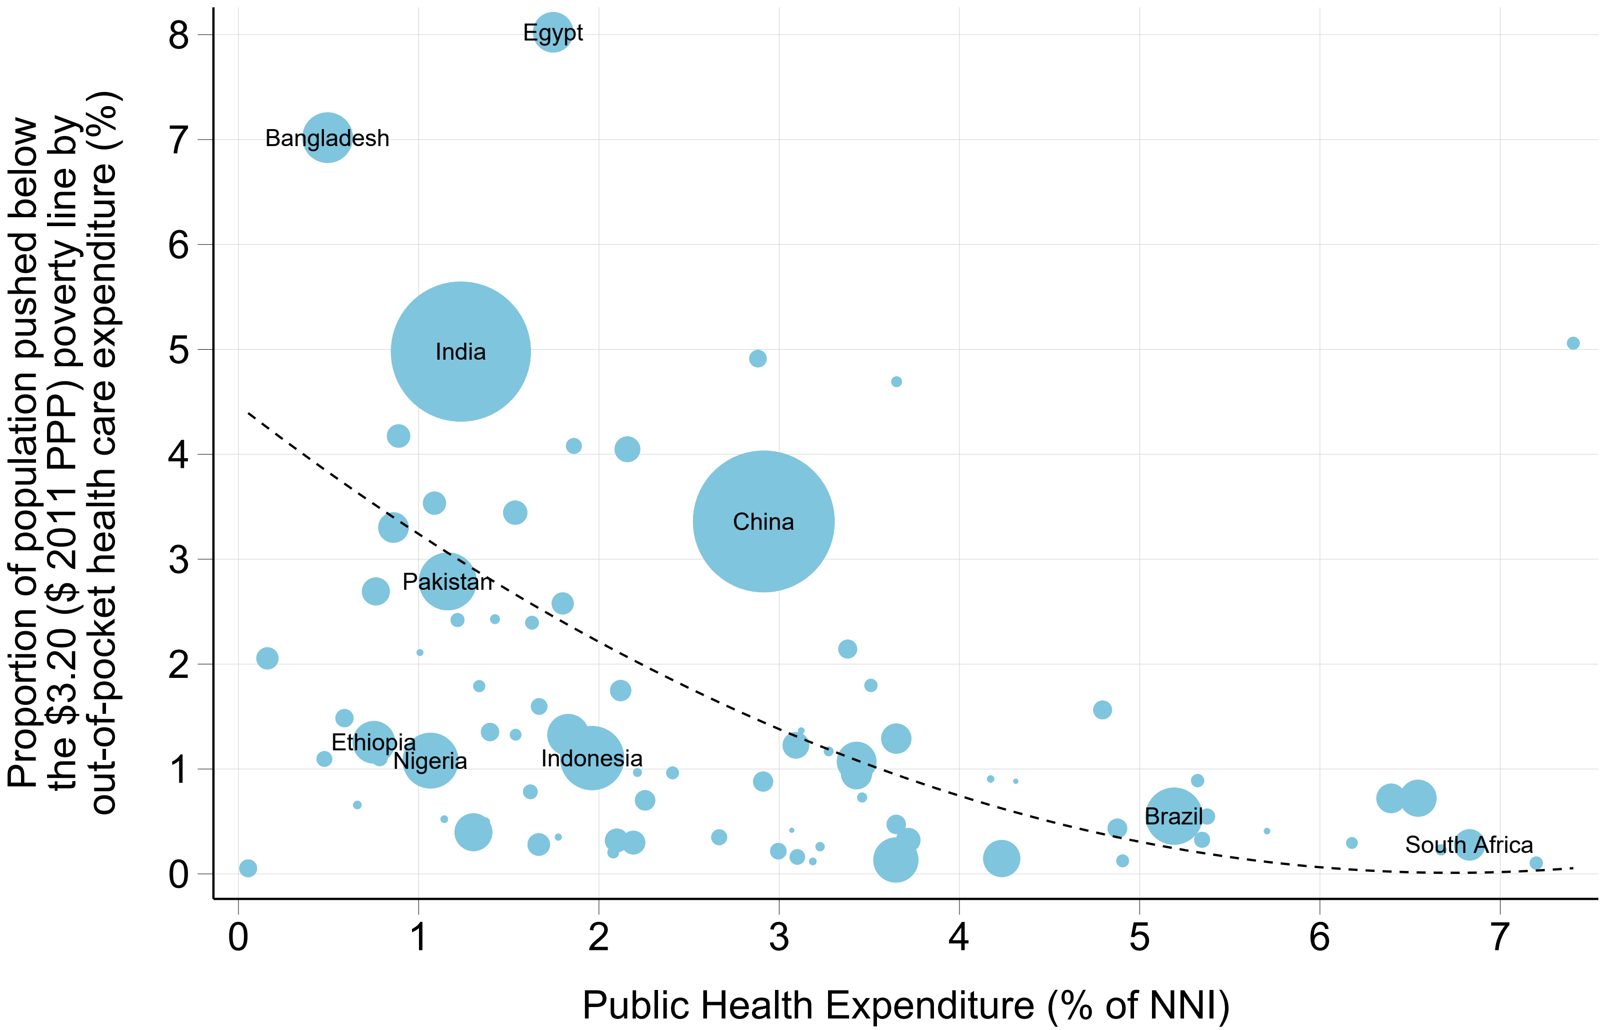

Notes. Author's computations combining national budget data (public health expenditure) and World Bank estimates (healthcare-driven poverty). The figure plots the relationship across countries between public health spending, expressed as a share of national income, and healthcare-driven poverty, measured as the share of the population falling into poverty due to out-of-pocket health expenditure. In countries spending more on public healthcare, fewer households fall into poverty due to own spending on healthcare.

Figure 2.2: Public Goods and Poverty Measurement

(b) Public Goods Matter for Non-Monetary Dimensions of Quality of Life

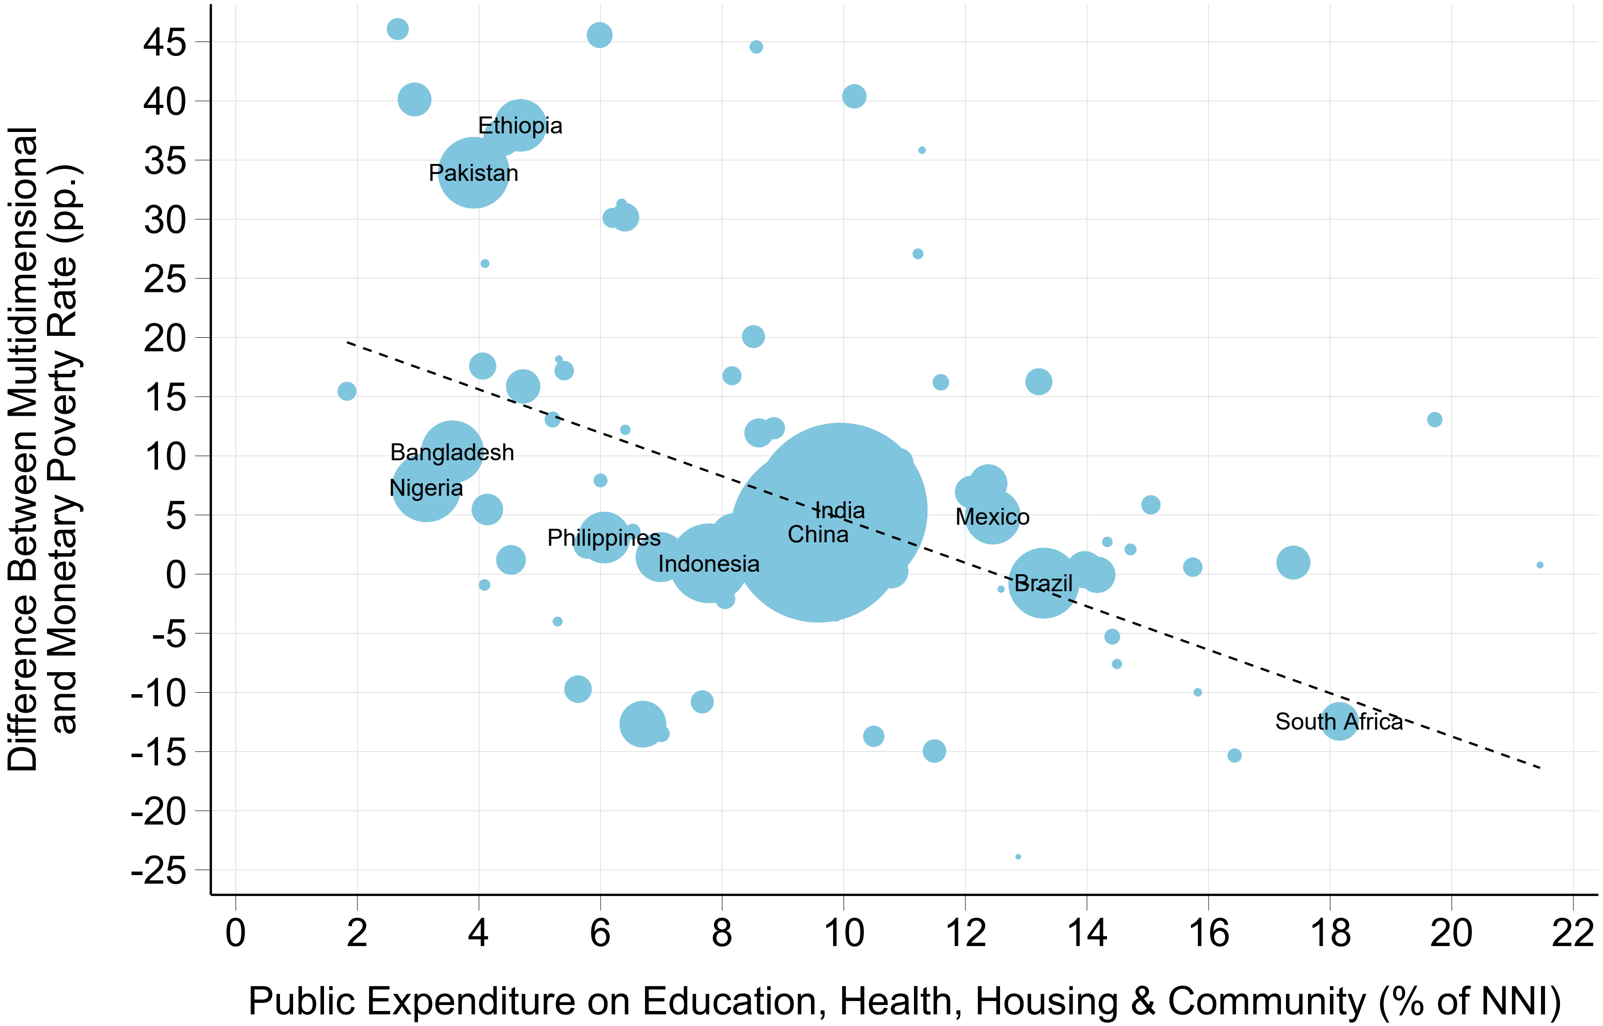

Notes. Author's computations combining national budget data (public expenditure), World Bank estimates (monetary poverty rate), and Oxford Poverty and Human Development Initiative estimates (multidimensional poverty rate). The figure plots the relationship across countries between public expenditure on education, health, housing, and community services, and the gap between monetary and multidimensional measures of poverty. Monetary poverty: share of population spending less than \$2.15 per day (2017 PPP USD). Multidimensional poverty: index combining deprivation in health, education, and living standards (see Alkire, Kanagaratnam, and Suppa, 2021). In countries with greater spending on basic public services, fewer households fall in multidimensional poverty relative to those falling in monetary poverty.

Figure 2.3: Public Goods Are Progressive
Distribution of Share of Transfers Received by the Bottom 50%

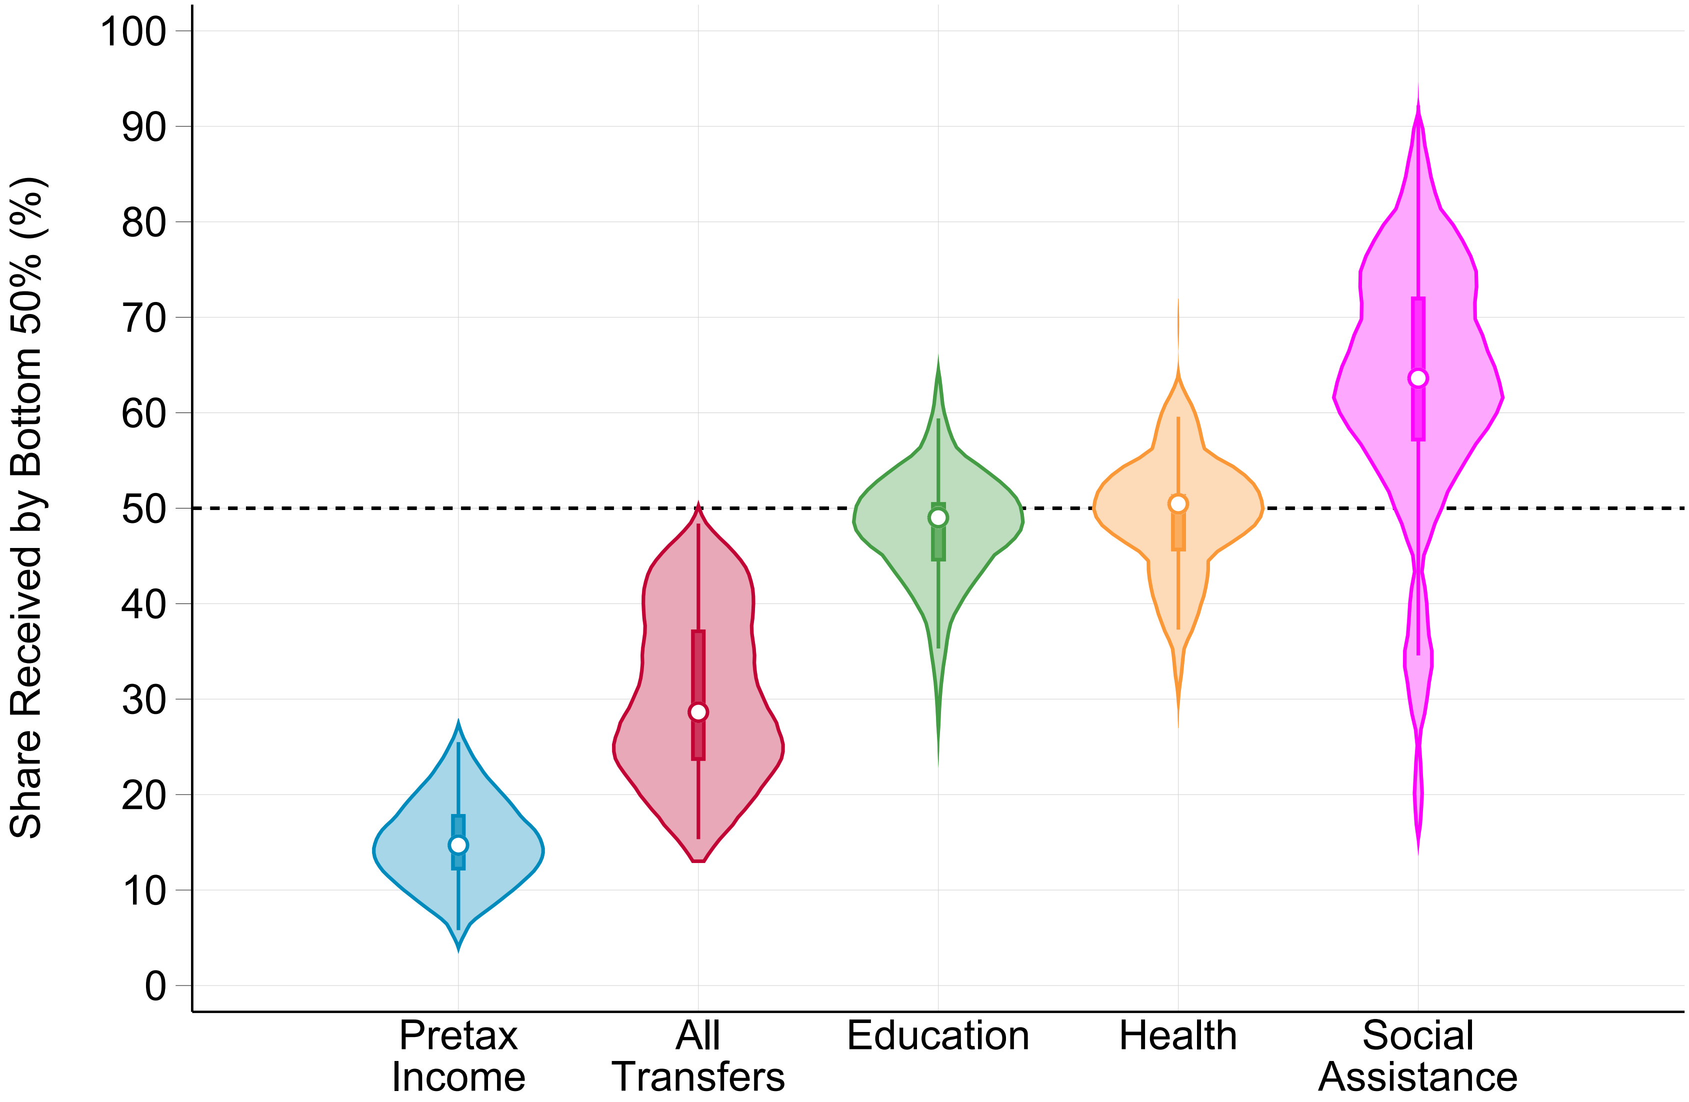

Notes. The figure represents the distribution of the share of government transfers and the share of pretax income received by the bottom 50% of the pretax income distribution in each country.

Figure 2.4: Public Goods Have Grown
Global Real Public Expenditure Per Capita, 1980-2019

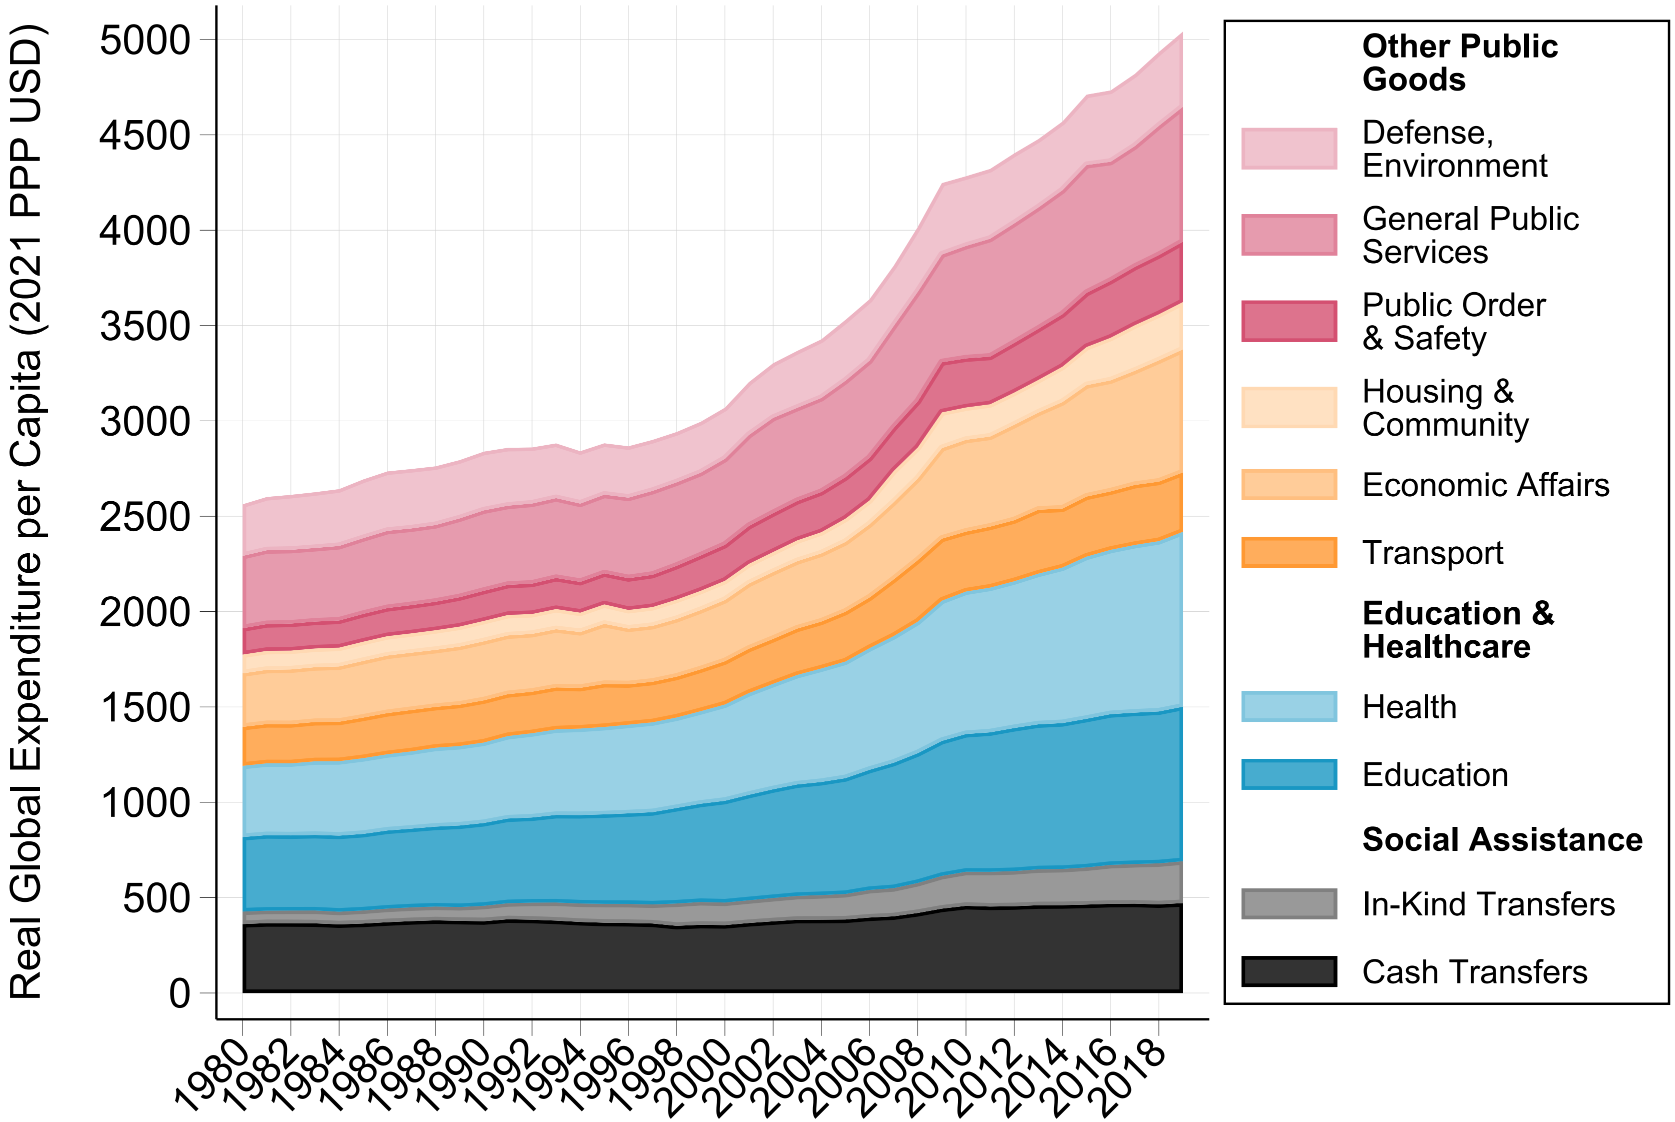

Notes. Author's computations using national budget data. The figure represents the evolution of general government expenditure per capita by function, expressed in 2021 PPP US dollars, in the world as a whole. Economic affairs include recreation and culture.

Figure 2.4: Public Goods Have Grown
Expenditure on Public Goods by Country Income Group, 1980-2019

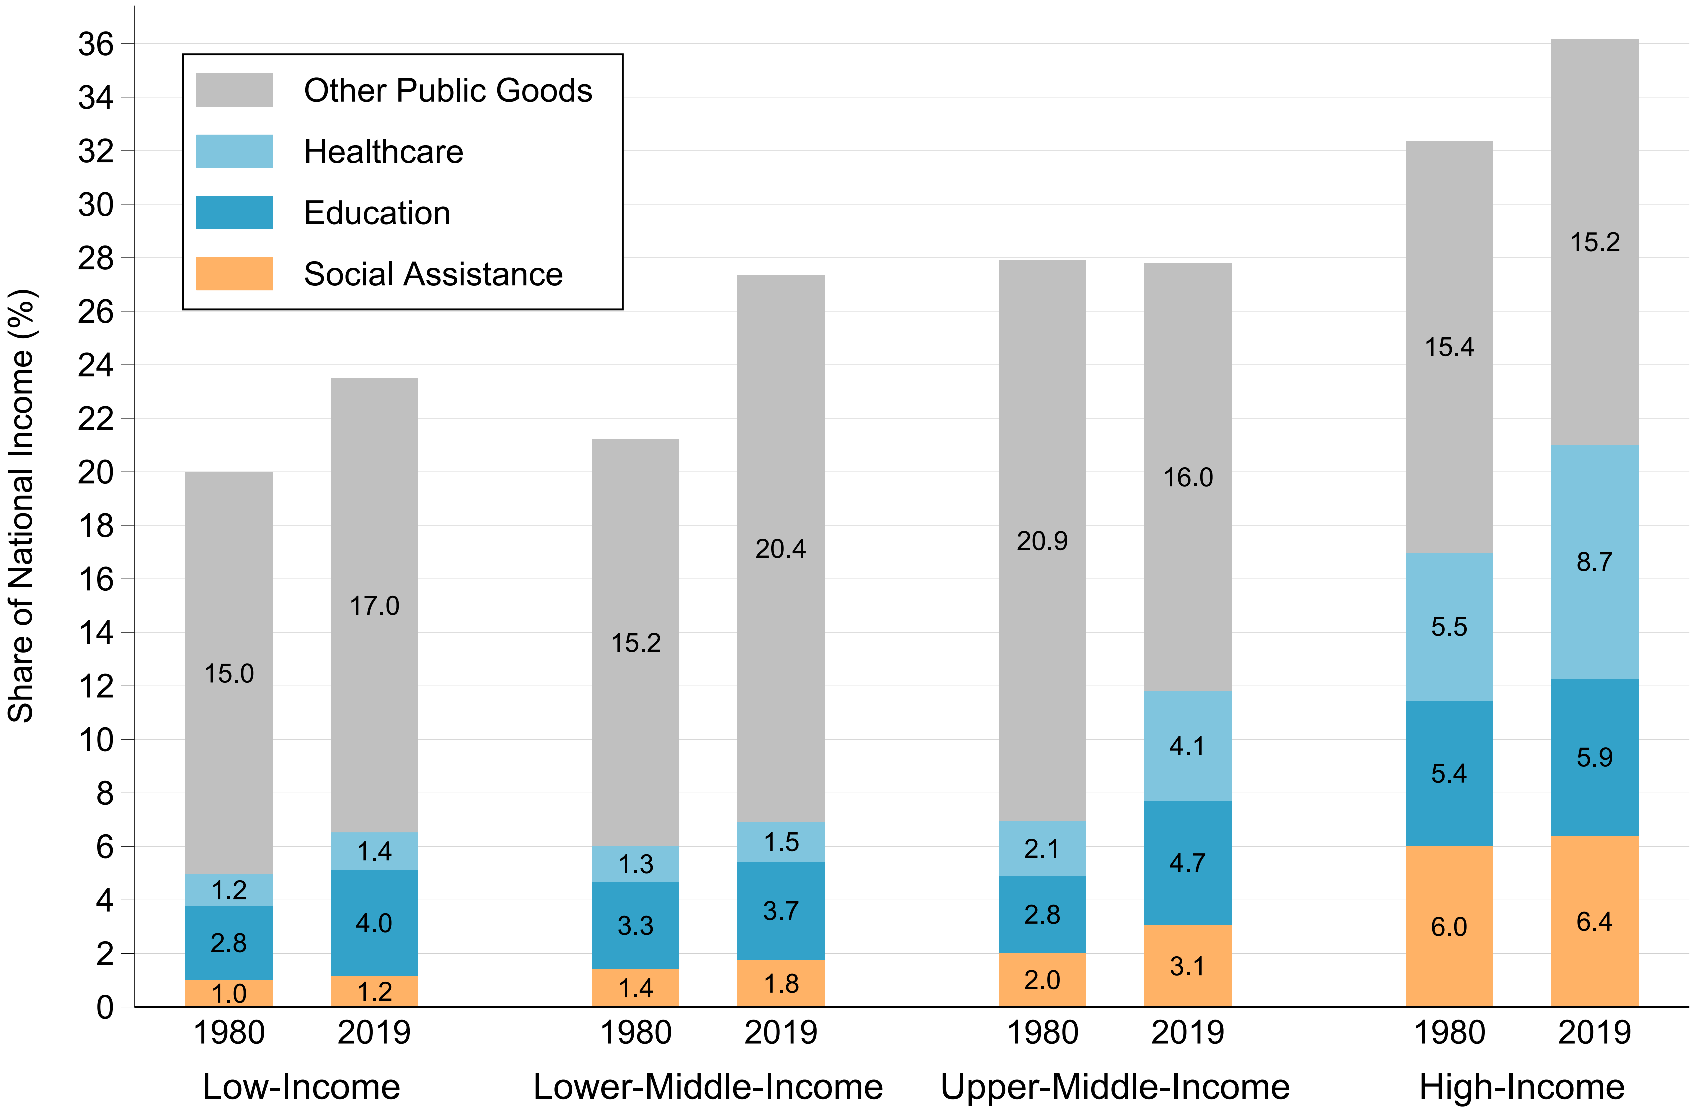

Notes. Author's computations combining national budget data. The figure represents the average share of national income spent on social assistance and public goods by country income group. Population-weighted averages across all countries in each group. See appendix figure B.38 for the composition of country income groups.

Figure 2.5: The Distribution of Public Goods in International Perspective

(a) Public Goods Received by the Bottom 50% Around the World

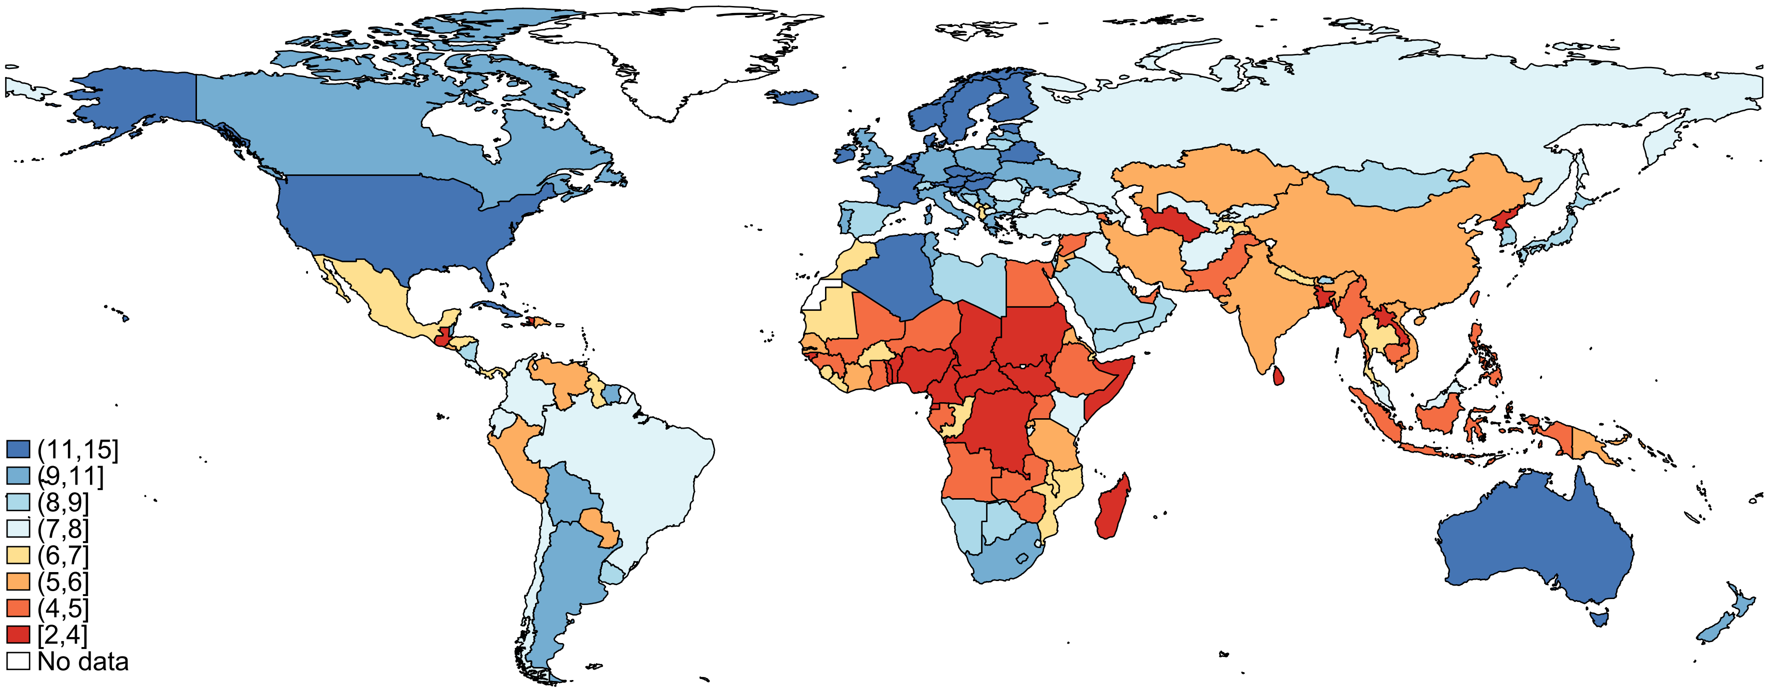

Notes. The figure maps total in-kind transfers received by the bottom 50% in each country in 2019, expressed as a share of national income. The unit of observation is the individual. Income is split equally between all household members.

Figure 2.5: The Distribution of Public Goods in International Perspective

(b) Public Goods Received by the Bottom 50% in Selected Countries

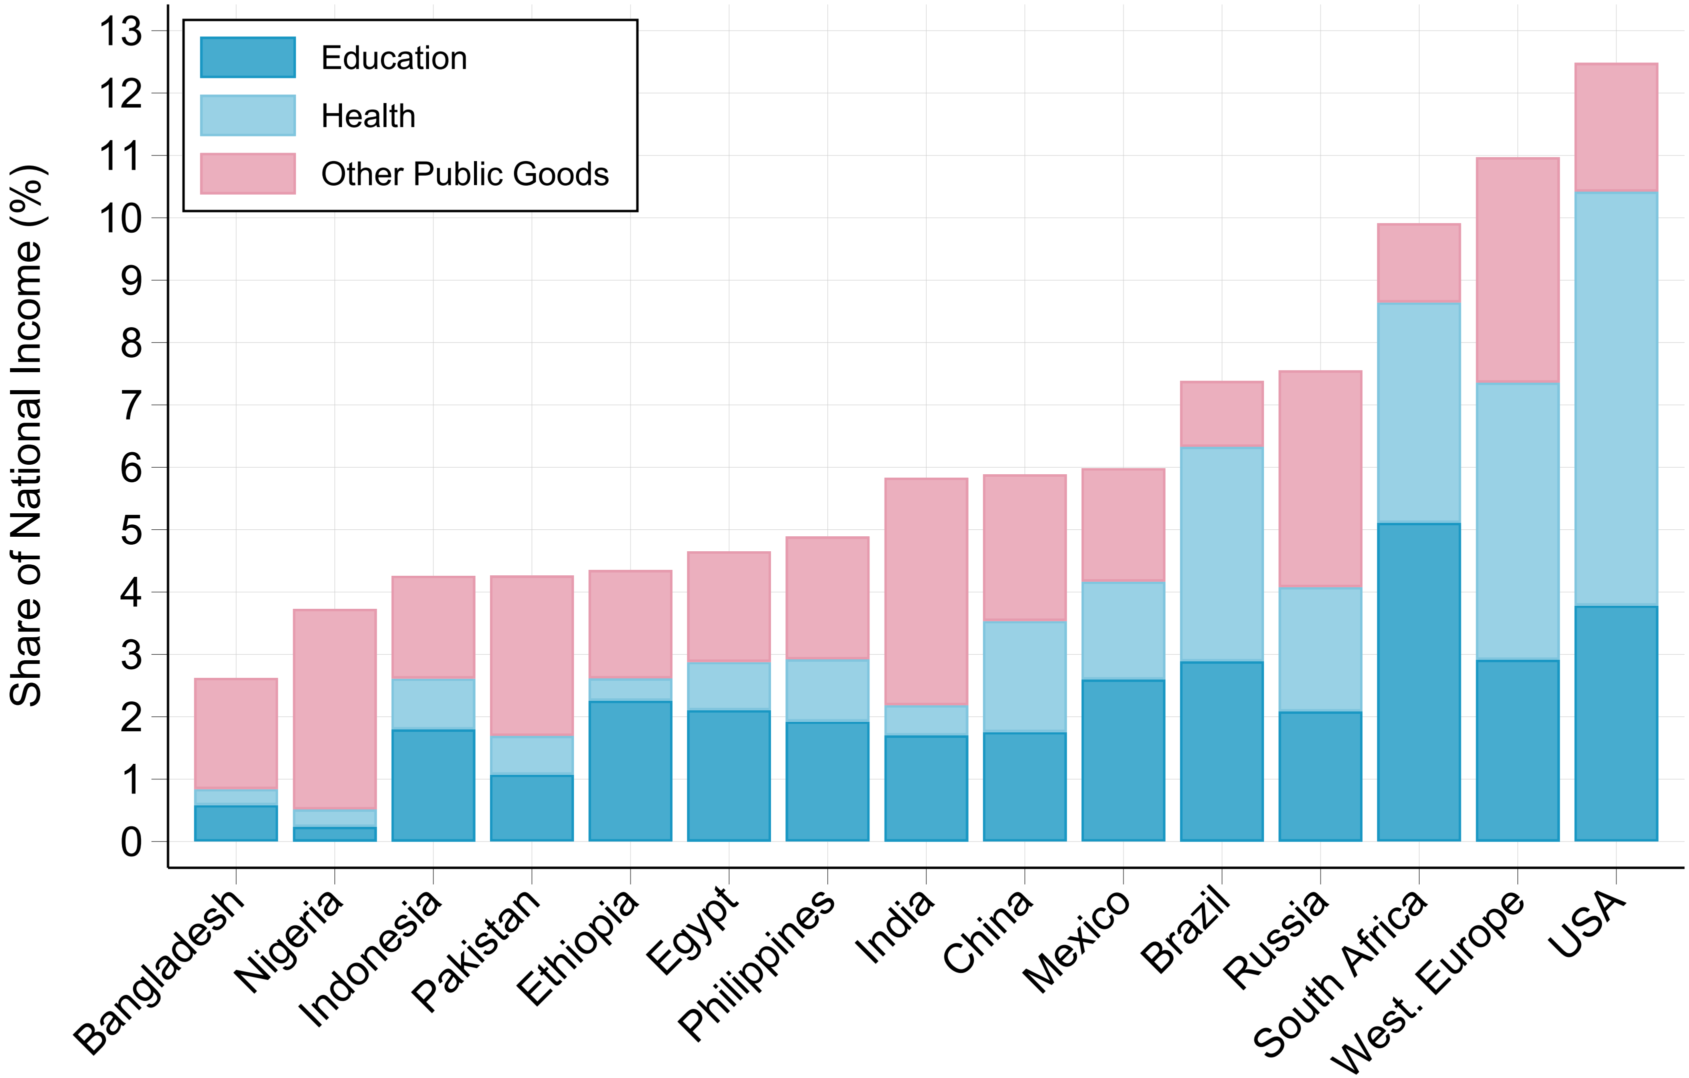

Notes. The figure shows the level and composition of in-kind transfers received by the bottom 50% in each country or region in 2019, expressed as a share of national income. The unit of observation is the individual. Income is split equally between all household members.

Figure 2.6: Public Goods Received by the Global Bottom 20%, 1980-2019 (% of Global Income)

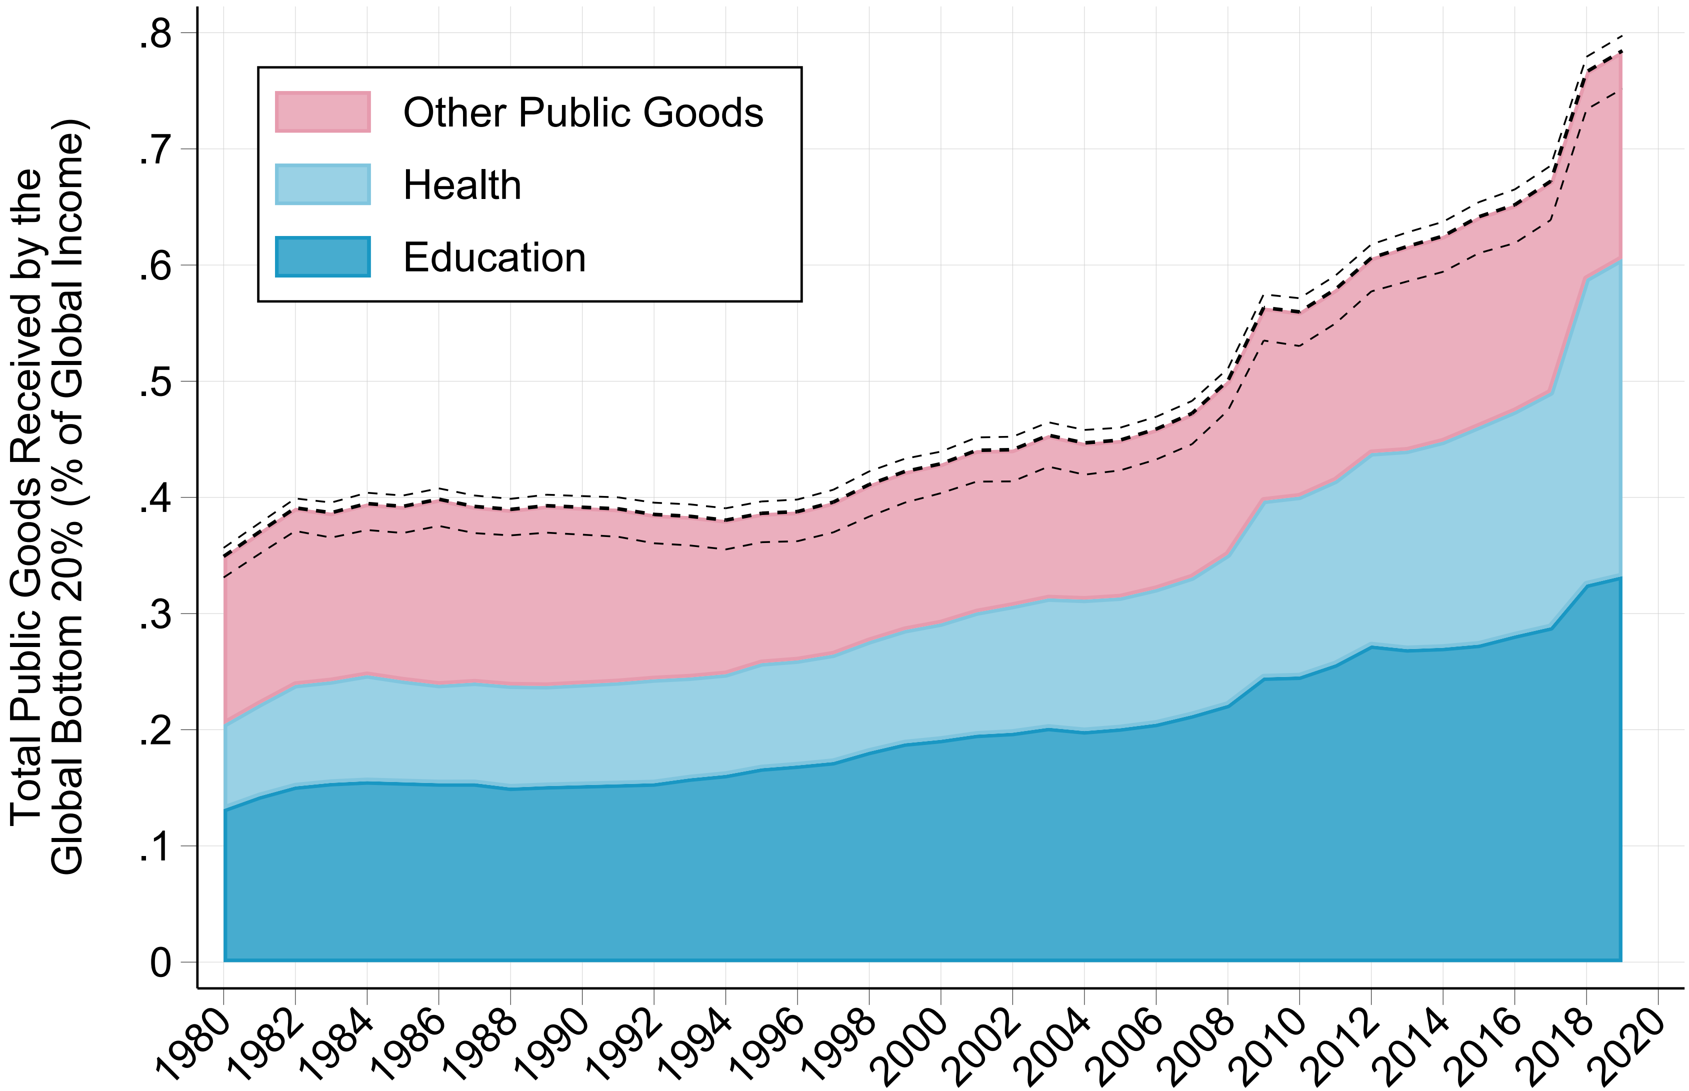

Notes. The figure plots the level and composition of public goods accruing to the global bottom 20%, expressed as a share of global income. The unit of observation is the individual. Income is split equally between all household members.

Figure 2.7: Public Goods and Global Poverty Reduction:
Real Average Income of the Global Bottom 20%, 1980-2019

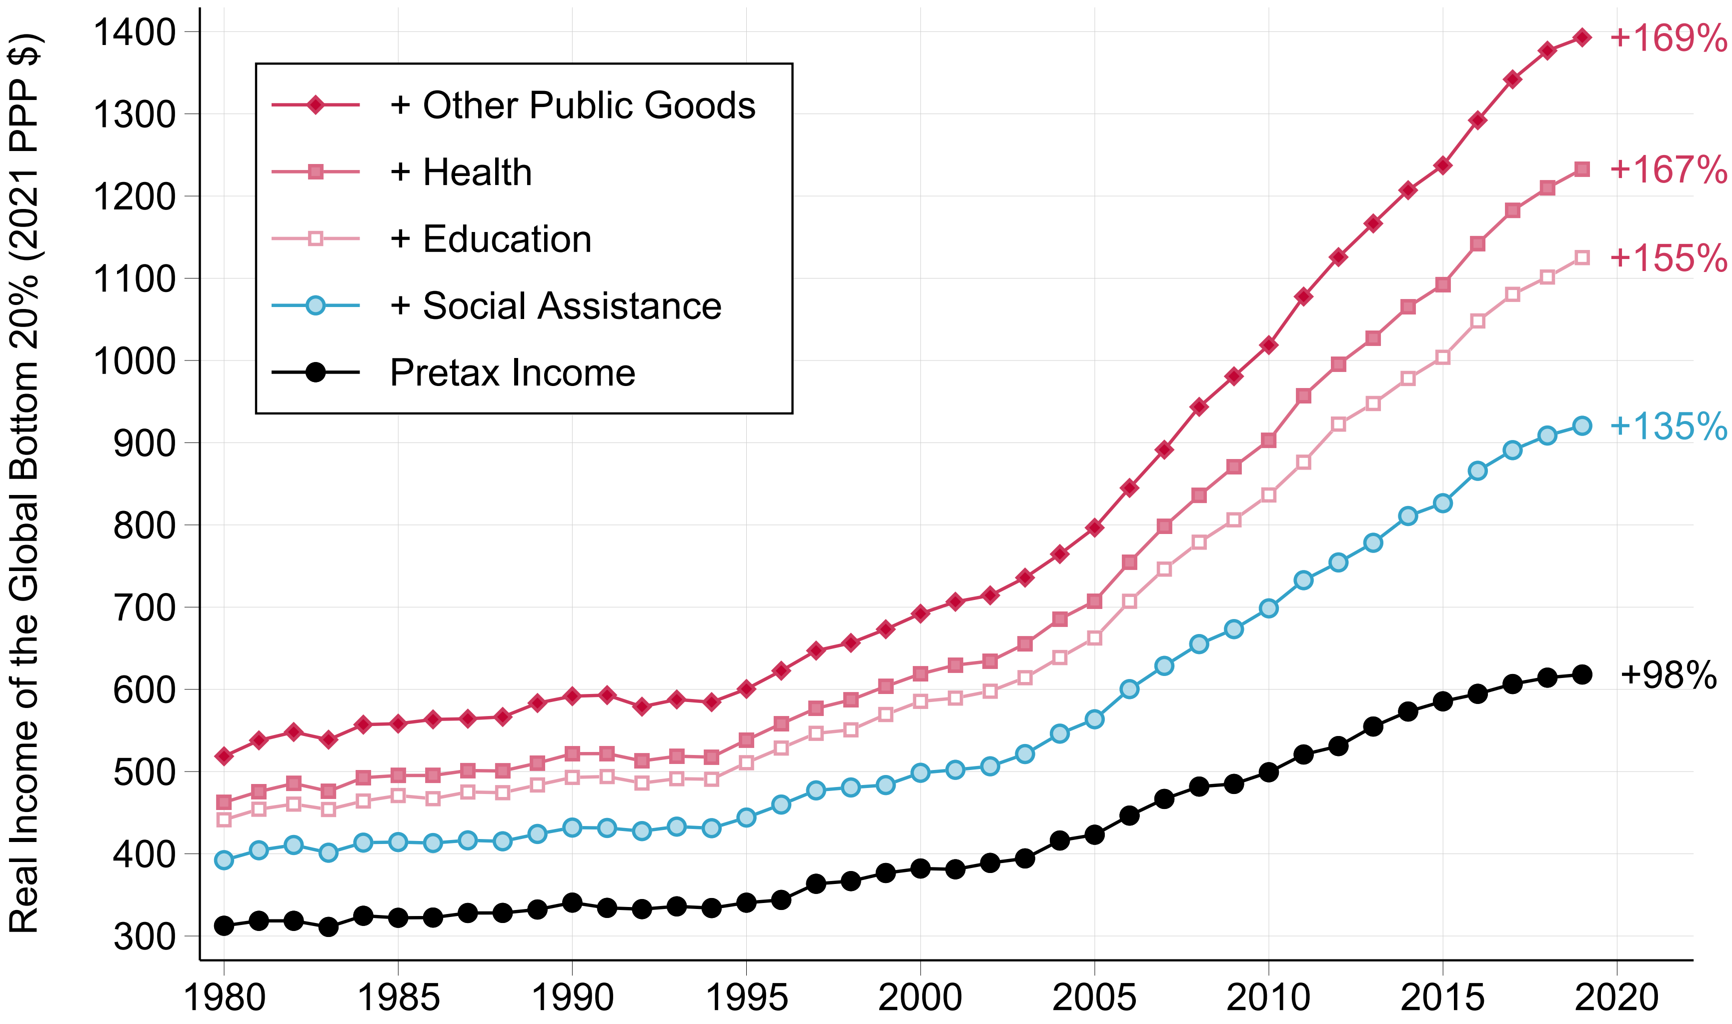

Notes. The figure plots the evolution of the global bottom 20% real average income from 1980 to 2019, before and after accounting for cash transfers and public goods. The unit of observation is the individual. Income is split equally between all household members.

Figure 2.8: Global Inequality and Public Goods

(a) Real Income Growth Rate by Global Income Percentile, 1980-2019

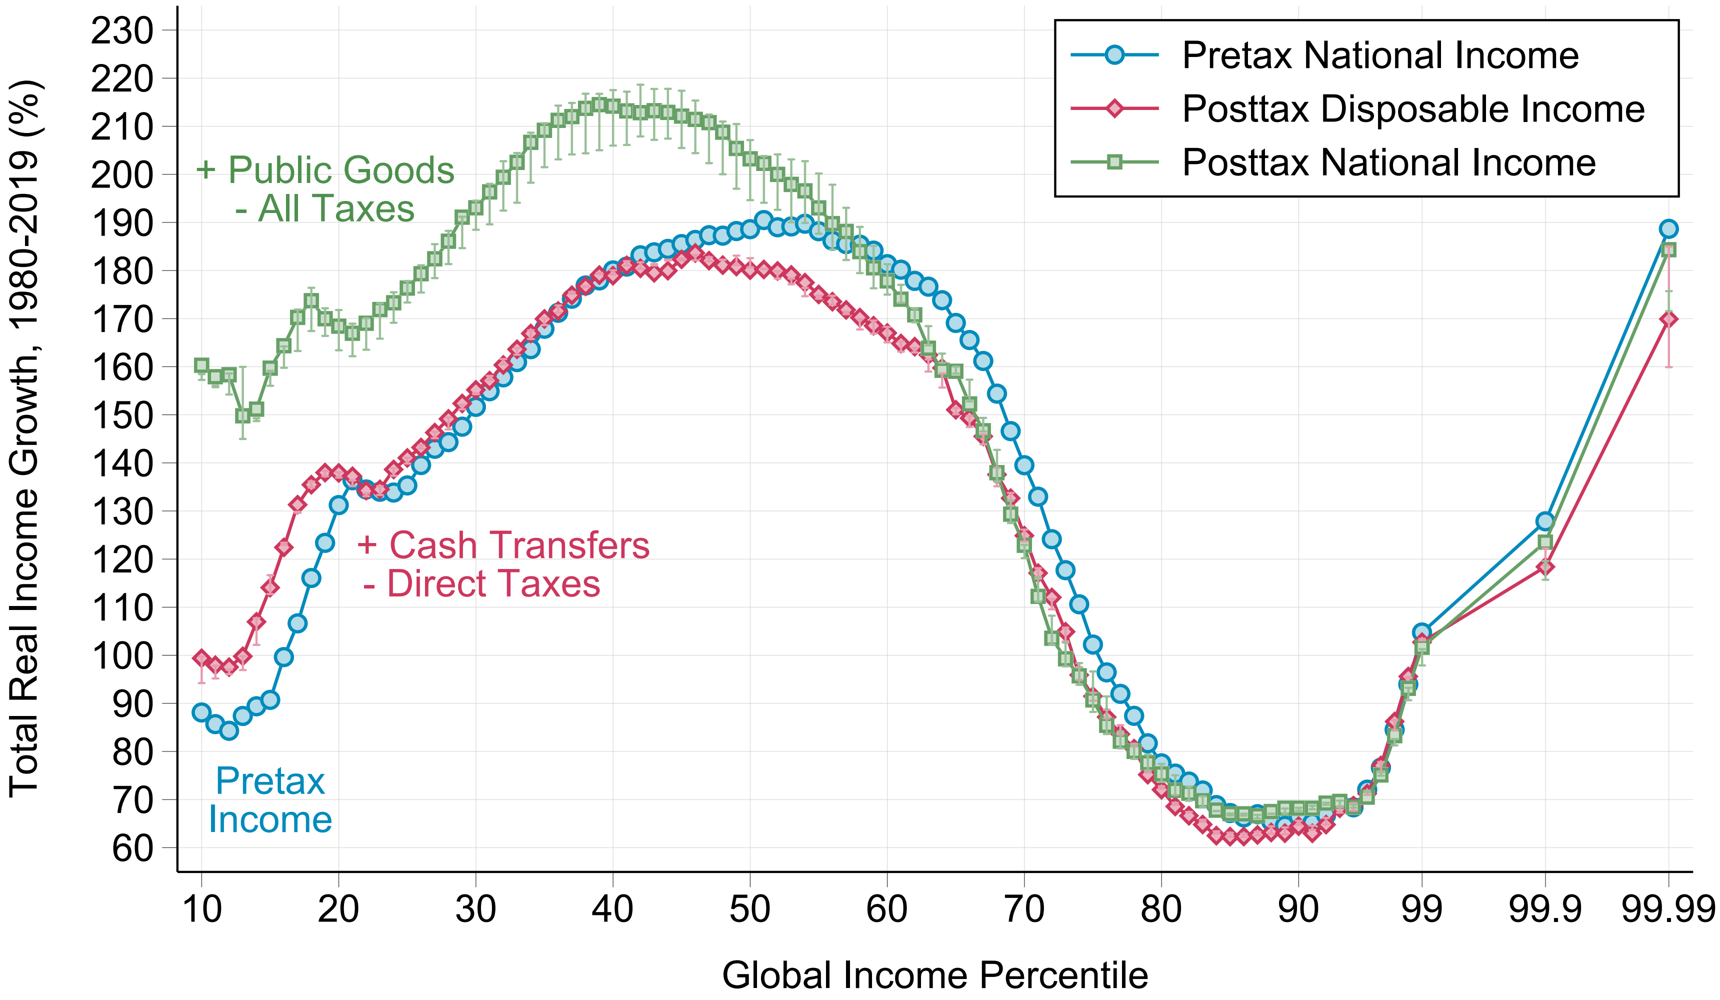

Notes. The figure plots total real income growth by global income percentile from 1980 to 2019 for different income concepts. Posttax disposable income removes direct taxes and adds cash transfers. Posttax national income removes all taxes and adds all cash and in-kind transfers. Capped spikes correspond to lower and upper scenarios on the progressivity of transfers. The unit of observation is the individual. Income is split equally between all household members.

Figure 2.8: Global Inequality and Public Goods

(b) Global Top 10% to Bottom 50% Average Income Ratio

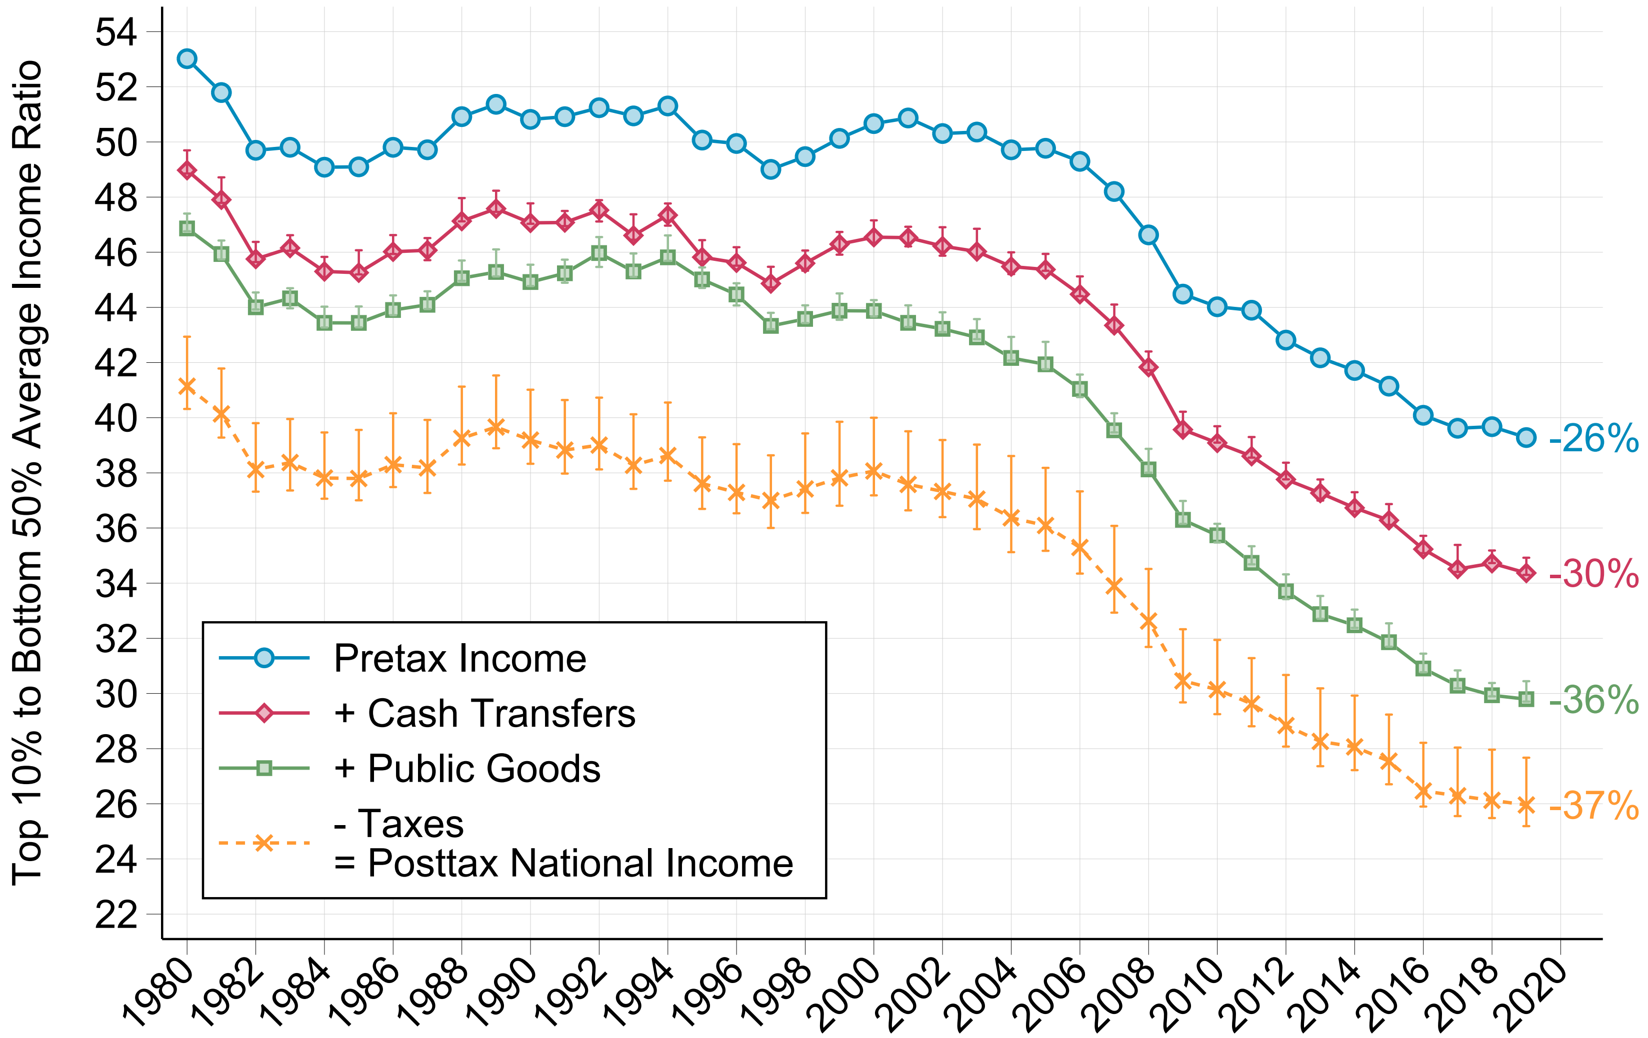

Notes. The figure plots the ratio of the average income of the top 10% to that of the bottom 50% in the world as a whole, for different income concepts. Capped spikes correspond to lower and upper scenarios on the progressivity of transfers. The unit of observation is the individual. Income is split equally between all household members.

Figure 2.9: Decomposing the Incidence of Public Goods on Global Poverty

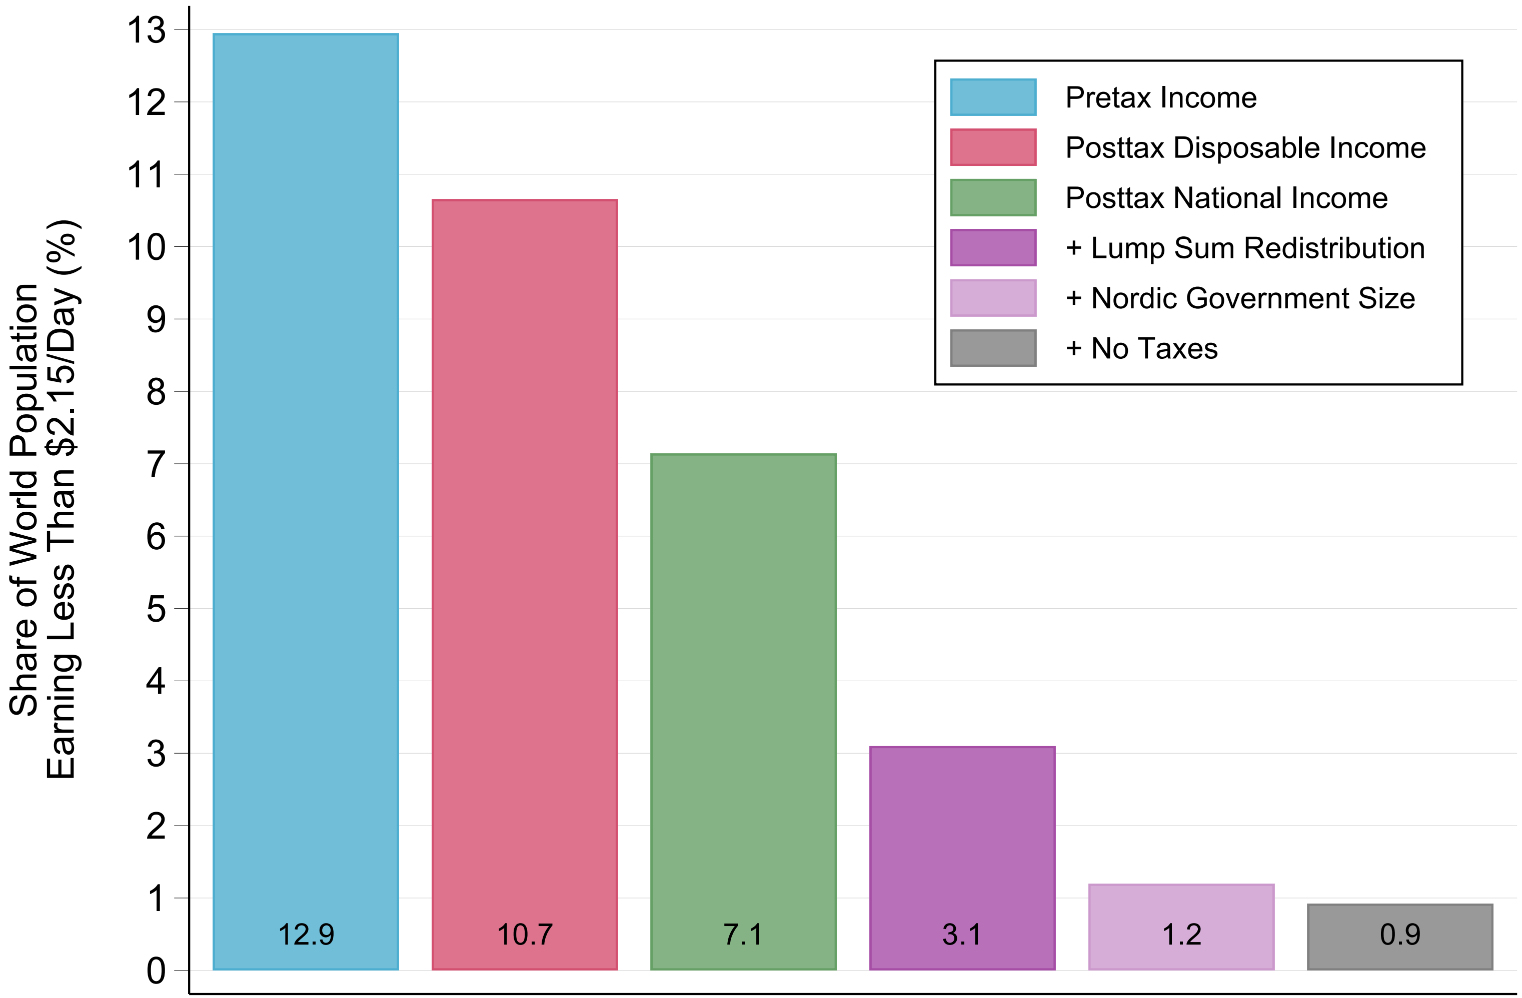

Notes. The figure plots the share of the world population living with less than $2.15 per day in 2019, measured in 2017 PPP USD, by income concept. Posttax disposable income removes direct taxes and adds cash transfers. Posttax national income removes all taxes and adds all cash and in-kind transfers. The fourth bar assumes that all transfers are received on a lump sum basis: $\gamma(m_i) = \gamma$. The next bar further considers that all countries have welfare states similar to that of Nordic countries, that is, general government expenditure is set at 50% of national income in each country. The last bar considers that no taxes are paid to finance transfers. The unit of observation is the individual. Income is split equally between all household members.

Table 2.1: Methodology Used to Distribute Global Government Expenditure

	Source / Method	Avg. Share of NNI (%) G^j		Share of Transfer Received (%) (γ^j, Bottom 50%)		
		1980	2019	Min	Mean	Max
Social Assistance	WID/CEQ/ASPIRE	2.6%	2.9%	16%	64%	92%
Education	GKL	3.5%	4.4%	25%	46%	64%
Health	WID/CEQ	2.5%	3.5%	29%	50%	69%
All Others	Prop. disposable income	17.4%	17.8%	8%	16%	30%
Economic Affairs		6.3%	5.8%			
General Public Services		5.6%	5.5%			
Public Order & Safety		1.4%	2.0%			
Other		4.1%	4.6%			
Total		26.0%	28.6%	15%	29%	48%

Notes. The table reports the sources used to distribute global government expenditure, together with summary statistics on expenditure by function as a share of national income and the share of expenditure received by the bottom 50% in each country. GKL: Gethin, Kofi Tetteh Baah, and Lakner (forthcoming). WID: Blanchet, Chancel, and Gethin (2022) for Europe and Piketty, Saez, and Zucman (2018) for the US. CEQ: Commitment to Equity Institute Database. Prop. disposable income: component distributed proportionally to posttax disposable income (pretax income, minus direct taxes, plus social assistance transfers).

Table 2.2: Dimensions of Redistribution by Country Income Group and World Region

	Expenditure (% NNI) G	Share of Transfer Received (%) (γ, Bottom 50%)	Net Transfer Received (% NNI) (g, Bottom 50%)
Country Income Group			
Low-Income	23.3%	22.8%	5.3%
Lower-Middle-Income	26.3%	24.0%	6.3%
Upper-Middle-Income	25.6%	29.2%	7.4%
High-Income	30.4%	33.2%	10.1%
World Region			
Sub-Saharan Africa	25.9%	22.3%	5.7%
Middle East and Northern Africa	28.6%	25.5%	7.2%
China	23.3%	27.0%	6.3%
India	31.4%	19.2%	6.0%
Other Asia / Oceania	23.3%	27.7%	6.5%
Latin America	25.8%	30.1%	7.7%
US / Canada / Western Europe	30.3%	34.9%	10.6%

Notes. The table reports statistics on dimensions of in-kind redistribution by country income group (defined based on the World Bank's classification) and world region. All figures focus on public goods, that is, total government expenditure excluding social protection spending.

Table 2.3: Public Goods and the Geography of Global Inequality

	Pretax National Income	Posttax Disposable Income	Posttax National Income
Theil Decomposition			
Theil Index	1.13	0.98	0.89
Between-Country Component	30%	33%	39%
Within-Country Component	70%	67%	61%
Share in Global Bottom 20%			
India	18%	21%	24%
China	11%	11%	8%
Pakistan	19%	24%	31%
Bangladesh	19%	20%	30%
Ethiopia	58%	66%	74%
Nigeria	23%	28%	34%
Other Asia / Oceania	17%	17%	17%
Other Sub-Saharan Africa	62%	65%	67%
Middle East and Northern Africa	19%	19%	17%
Latin America	17%	11%	6%
US / Canada / Western Europe	7%	2%	0%

Notes. The table reports a Theil decomposition of global inequality into a between-country and a within-country component, as well as the geographical composition of the global bottom 20% in 2019, for different income concepts. Posttax disposable income removes direct taxes and adds cash transfers. Posttax national income removes all taxes and adds all cash and in-kind transfers. The unit of observation is the individual. Income is split equally between all household members.

Chapter 3

Government Redistribution and Development: Global Estimates of Tax-and-Transfer Progressivity, 1980-2019

Despite a momentous renewal of attention to inequality, even the most recent studies often fail to account for the distributional effects of government taxes and transfers—above all in the developing world. Publicly available inequality statistics generally provide data on the distribution of household disposable income or consumption, with little information on the extent to which government intervention affects poverty and inequality. While significant recent efforts have been made in specific countries, there is a critical lack of cross-country, long-run data on how redistribution in its different forms has evolved in the past decades. As a result, it remains difficult to answer questions as simple as: which countries do the most to reduce income disparities through taxes and transfers? Is redistribution higher than it was forty years ago? Are differences in inequality primarily driven by differences in the distribution of market incomes ("predistribution"), or by differences in tax-and-transfer systems ("redistribution")?

This article makes a first step towards answering these questions. Combining new data sources and methods, we assemble a comprehensive database on the distribution of taxes and transfers in 151 countries since 1980. Our estimates of redistribution account for all forms of taxes and transfers, including personal income taxes, corporate taxes, consumption taxes, local taxes, cash transfers, and public education and health expenditure. We distribute all taxes and transfers using a common methodological framework, Distributional National Accounts (DINA; Blanchet et al., 2021), which ensures that our estimates are comparable across countries and over time, and consistent with national income and government budget aggregates.

In the absence of survey or tax microdata, which largely do not exist for our sample, several methodological innovations allow us to estimate the distributional incidence of taxes and transfers. Tax revenue aggregates, by type of tax, are drawn from Bachas et al. (2022), while pretax income distributions are available from the World Inequality Database (Blanchet et al., 2021). We model the distributional incidence of taxes from a number of parameters on *inter alia* statutory tax schedules, functional income concentrations, and the relative weights of disaggregated tax components, for which we put together data from several sources. Similarly, we complement our new series on total government expenditure, by function, with information on the distributional incidence of social assistance, education, and healthcare, drawing on related work by Gethin (2023b). We validate our estimates against those of existing studies where those exist, ensuring that our simplified methodology accurately reproduces results from preexisting work.

Our database reveals five new stylized facts on worldwide fiscal progressivity, in levels and trends. First, tax-and-transfer systems always reduce inequality. One way to measure this is to compare the top 10% to bottom 50% average income ratio in terms of pretax and posttax income. Taxes and transfers reduce this ratio in all 151 countries in our sample. This effect varies considerably, however, from 15% in the average African country to over 30% in Europe and the United States.

Second, transfers are the dominant driver of this redistributive effect. Taxes appear to have almost no effect on inequality in most regions of the world: low-income households face about the same effective tax rate as high-income households. As a result, removing taxes from individual incomes reduces inequality by about 2% in the average country. In contrast, transfers always strongly reduce inequality, typically by about 20%. Putting these two facts together, we estimate that over 90% of the effect of tax-and-transfer systems on inequality comes from transfers, while less than 10% comes from taxes.

Third, redistribution rises with development, but this is entirely due to transfers. Tax progressivity is uncorrelated with per capita income, despite noticeable regional patterns. For instance, Western European and Anglosphere countries have slightly progressive tax systems, while the distribution of taxes is strongly regressive in Eastern Europe and Latin America, mainly due to the prevalence of high indirect taxes and less progressive personal income taxes. In contrast, the impact of transfers on inequality rises sharply with development: the raw correlation between the total transfer received by the bottom 50% as a share of national income and GDP per capita exceeds 0.6. This finding mainly arises from the fact that high-income countries spend more on cash and in-kind transfers, but can also be explained by their greater reliance on more progressive forms of public spending—in particular social assistance and healthcare. In the average African country, less than 2% of national income is transferred to the poorest 50% of the population in the form of government transfers, compared to over 6% in Europe and the United States.

Fourth, there has been no cross-country convergence in redistribution. The net effect of taxes and transfers on inequality has increased significantly in the average country, from a reduction of approximately 10% in 1980, to 20% in 2019. However, this average figure masks considerable heterogeneity. Redistribution has risen significantly in Western Europe, the Anglosphere, and Latin America, while it has stagnated in Eastern Europe and Africa. The gap in redistribution between low- and high-income countries has remained about the same. Upper-middle-income countries have caught up with high-income countries, but this is mainly due to the rise of fiscal progressivity in China.

Fifth, despite large cross-country differences in tax-and-transfer systems, variations in inequality are primarily driven by differences in pretax inequality ("predistribution") rather than by variations in taxes and transfers ("redistribution"). In line with existing work focusing on Europe and the United States (Blanchet, Chancel, and Gethin, 2022; Bozio et al., 2022), we find that countries displaying the highest levels of pretax inequality also end up displaying the highest levels of posttax inequality. A simple cross-country regression of the bottom 50% posttax income share on the bottom 50% pretax income share yields an R-Squared of over 0.8. By this measure, predistribution accounts for over 80% of cross-country variations in inequality, while redistribution accounts for less than 20%. We do find a strong correlation between predistribution and redistribution, however: countries with more progressive tax-and-transfer systems display lower levels of pretax inequality. This suggests that while the *direct* effect of taxes and transfers explains little of variations in posttax inequality, redistributive policies might still play a much more important role in *indirectly* shaping the distribution of market incomes.

Our work stands at the confluence of two main strands of the literature on inequality and fiscal policy: one that has studied the incidence and impact of taxes and transfers, and another that has aimed to measure inequality in a way consistent with measures of growth and total national income.

In the former, tax incidence analysis maintains an illustrious tradition, from Musgrave (1953), Tax Foundation (1967) and Kakwani (1977) through Lambert (1992), Fullerton and Metcalf (2002) and Saez, Slemrod, and Giertz (2012). The central question of this literature has been to ask on whom the burden of taxation falls. Studies in this line have emphasized context-specific behavioral responses to taxation, and the role of taxes and transfers to equalize income distributions. Few studies have taken comprehensive account of all taxes, all transfers, and all incomes, measuring the movement from pretax to posttax income distributions in a way that is consistent with macroeconomic estimates of national income.

In the latter tradition of inequality measurement, a slew of recent DINA studies have generated worldwide evidence on pretax income inequality levels and trends (see Chancel et al., 2022b).[1] Gathered together in the World Inequality Database, these data series represent a scholarly benchmark as the preeminent long-run, worldwide, harmonized estimates of total national income distributions. However, the majority of these income distributions are estimated only pretax[2]—before the operation of government tax and transfer policies—leaving open an important empirical question, on the ability of fiscal policy to impact inequality.

The central contribution of this paper is to close that gap and estimate comprehensive posttax income distributions, worldwide since 1980. As such, our work relates perhaps most directly to the Commitment to Equity initiative (CEQ Institute; see Lustig, 2018 and World Bank, 2022a), whose pioneering efforts have made important strides to estimate the incidence of taxes and transfers in the developing world.[3] Our main contribution beyond their work is to cover all countries, all incomes, and all taxes and transfers, as well as the evolution of redistribution over time.

The remainder of this article is organized as follows. Section 3.1 establishes our methods to estimate worldwide fiscal progressivity since 1980, and demonstrates the robustness of the approach. Section 3.2 presents our analysis and the main findings that emerge. Section 3.3 concludes.

[1] Pretax income is the income that accrues to all earners directly on the marketplace, before taxes and transfers (but after social insurance), with the distribution of income adding to 100% of annual national income in the national accounts. For background and further details on the concept of pretax income and its estimation, refer to Blanchet et al. (2021) and the World Inequality Database.

[2] Several important exceptions are discussed in sections 3.1.4 and 3.1.5 below.

[3] CEQ studies generally do not precede the year 2010, and usually cover but one year per country. Equity income from ownership of corporations, as well as corporate income taxes, are usually excluded from this framework.

3.1 Data and Methodology

This section covers the methodology used to build our new database on government redistribution worldwide. Section 3.1.1 covers general methodological principles. Section 3.1.2 outlines the data sources used for the distribution of pretax income and government revenue and expenditure aggregates. It also presents our core "calibration" and "validation" database on government redistribution in 45 countries, compiled from seven studies following the DINA framework—which we use to inform and to test several distributional incidence assumptions. Section 3.1.3 describes the methodology used to allocate taxes and transfers. Finally, section 3.1.4 investigates the ability of our methodology to reproduce estimates from seven existing DINA studies.

3.1.1 Conceptual Framework

Concepts Our methodology follows the distributional national accounts (DINA) framework (Blanchet et al., 2021; Piketty, Saez, and Zucman, 2018), which offers a foundation to estimate the distribution of income, taxes, and transfers in a way that is consistent with national accounting principles (UN SNA, 2008). Unlike previous approaches to the measurement of inequality, the DINA methodology distributes all income flows to all individuals, as well as all types of taxes paid and transfers received, to arrive at both pretax and posttax income distributions that match 100% of national income.

The DINA approach generally establishes three income concepts: factor national income, pretax national income, and posttax national income, all of which add up to net national income. Factor national income refers to market income flows deriving from labor and capital, before any form of government intervention.[4] Pretax national income corresponds to income after the operation of the pension and unemployment systems, but before the operation of the tax-and-transfer system. It is equal to factor income, minus social contributions paid, plus social insurance benefits received. Finally, posttax national income corresponds to income after the operation of the tax-and-transfer system. All taxes are allocated and removed from individual pretax incomes, including personal income taxes, corporate taxes, property and wealth taxes, and indirect taxes. Similarly, moving from pretax to posttax national income implies distributing the entirety of general government expenditure, including cash

[4] It can be expressed net or gross of indirect taxes on production. It involves allocating incomes usually observed in surveys and tax data, such as compensation of employees and dividends, but also income flows only received indirectly by households, such as imputed rents or the retained earnings of corporations, which are also part of net national income.

transfers, in-kind benefits (e.g., healthcare), and collective government expenditure (e.g., public order and safety).

Objective We focus on measures of government redistribution that compare the distribution of pretax national income to that of posttax national income.[5] Starting with data on the distribution of pretax income z, we aim to measure the distribution of taxes $T(z)$ and government transfers $G(z)$, so as to reach posttax income y:

$$y = z - T(z) + G(z) \tag{3.1}$$

Our analysis therefore relies on three key ingredients: data on the distribution of pretax income, data on total taxes collected and transfers disbursed in each country, and data on the distributional incidence of each type of tax and transfer. We turn to each of these three ingredients in turn.

3.1.2 Data Sources

Data on Pretax Income Distributions Our starting point on the distribution of pretax national income is the World Inequality Database, which covers 174 countries over the 1980-2019 period. The database was constructed by compiling estimates from existing DINA studies, which have been systematically harmonized and combined to yield comparable distributional statistics (see Chancel and Piketty, 2021). For each country-year, the data cover pretax income thresholds and averages for 127 generalized percentiles (g-percentiles), corresponding to each percentile within the bottom 99% (*p0p1* through *p98p99*), followed by a more detailed decomposition of incomes within the top 1%. By construction, following the DINA framework, average income is consistent with net national income, as recorded in the World Inequality Database (see Blanchet and Chancel, 2016; UN SNA, 2008). The database also provides information on the share of pretax income coming from capital income and labor income, for each g-percentile (Blanchet, 2022b). This decomposition is consistent with aggregate factor income shares estimated in Bachas et al. (2022).

[5] As in the existing studies that apply the DINA framework, we prefer to measure the distance between pretax income and posttax income, rather than between factor income and posttax income. This comparison has the advantage of not making estimates of redistribution too sensitive to demographic factors, such as the size of the elderly population (where retired persons earn zero factor income but do receive significant social security benefits). Furthermore, even if social insurance contributions do resemble a tax—as a *compulsory* levy, *unrequited* at the time of payment—social insurance benefits resemble less of a redistributive transfer, and rather may be considered as deferred compensation, similar to any private-sector pension or annuity.

Data on Tax Revenue Aggregates To study the distribution of taxes paid by individuals, we first need to know the magnitude and composition of government revenue. We rely on aggregate tax revenue series recently constructed by Bachas et al. (2022), who combine national accounts data with government revenue statistics to estimate the evolution of macroeconomic tax rates in more than 150 countries since 1965. Their database provides information on total tax revenue as a share of national income, disaggregated into six categories: personal income taxes (code 1100 in the OECD classification of taxes; OECD, 2022), corporate income taxes (1200), social insurance contributions (2000, 3000), property and wealth taxes (4000), indirect taxes (5000), and other taxes (6000).

Data on Public Expenditure Aggregates To study the distribution of transfers, we similarly need to know the magnitude and composition of government expenditures. We use data from Gethin (2023b), who estimates harmonized series on the level and composition of general government expenditure by function of government (COFOG). The database provides information on government expenditure on social protection, education, healthcare, and other public spending in about 170 countries since 1980. Social protection is itself disaggregated into social insurance (pension and unemployment benefits) and social assistance.

Data for Validation Having compiled data on pretax income inequality and disaggregated government revenue and expenditure, we need to estimate the distributional incidence of taxes and transfers in each country-year for which the above aggregates are observed. We start by collecting data on the incidence of taxes and transfers in countries for which detailed, high-quality estimates are available from existing DINA studies. Table 3.1 provides information on the data collected from these studies: in total, the database covers 657 country-years over 45 countries, with significant time and geographical variation. From each study, we collect information on tax and transfer incidence profiles, that is, the share of taxes paid and transfers received by pretax income generalized percentile.

Taken together, the fiscal incidence data from these studies provides unique insights into variations in tax-and-transfer progressivity over time and space. We use these different estimates for validation of our estimates, as discussed further in section 3.1.4 below.

3.1.3 Distribution of Taxes and Transfers

Each tax and transfer, for each country-year, has a unique distributional profile. We now discuss the distributional estimates for each type of tax and transfer in turn. To introduce our method, consider the following equation:

$$T_i = \int_{p \geq K}^{p100} \tau_i(z)\ dz \tag{3.2}$$

For each type of tax and overall, the aggregate revenue received by the government is equivalent to the sum of taxes paid by all tax units, or the definite integral of effective tax rates applied to incomes over the distribution. The function $\tau_i(z)$ gives the taxes of type i paid by pretax income z, for each g-percentile p. The equivalent is true for transfers (negative taxes). By construction, our estimates always match revenue and expenditure totals T_i on aggregate. Our goal is to estimate the shape of $\tau_i(z)$ over the income distribution, for each type of tax and transfer i.

Personal Income Taxes For personal income taxes (PIT), only taxpayers with income above the PIT exemption threshold K pay any taxes. We estimate K for all country-years from Bachas et al. (2022) and Jensen (2022). Above the PIT exemption threshold, we simulate the structure of personal income tax incidence using statutory rate schedules from the World Tax Indicators (WTI) database (Peter, Buttrick, and Duncan, 2010). This database provides information on the average and marginal statutory income tax rates at average income (where taxable income equals per capita national income), then at two and three and four times that level, and finally the top marginal tax rate. We complement the WTI with inputs from Strecker (2021) and Vegh and Vuletin (2015) and online sources. From this basis, we can approximate a continuous schedule of statutory personal income tax rates.

Drawing on additional data sources (see Appendix C.1), we also make three critical distinctions: (1) between countries whose PIT systems tax married couples' joint income vs. those that only tax individual incomes; (2) between countries whose PIT systems tax capital income differently from labor income, noting differential rates on dividends and on capital gains; and (3) between the pretax and *taxable* income distributions (since (1) and (2) may occasion some re-ranking).

In this simplified simulation, the elements of the PIT system can be summarized as follows, to estimate the tax rate τ for any g-percentile p and its corresponding income level z:

$$\tau(z)_{PIT} = \sum_{j=1}^{3} \frac{\tau_j z_j}{z} \tag{3.3}$$

Where j refers to PIT on labor income (employee compensation and mixed income), dividend income, and capital gains (with taxable incomes z_j taxed at rate τ_j).

After modeling this statutory PIT schedule, we fit its "predicted" revenues proportionally to actual revenues observed in Bachas et al. (2022) and corresponding to T_{PIT} in equation (3.2) above.

Corporate Income Taxes Following Blanchet et al. (2021), we allocate the corporate income tax (CIT) proportionally to income from corporate equity. High-quality estimates of corporate equity ownership (and, therefore, corporate income tax burdens) by generalized percentile are available for the Netherlands (Bruil et al., 2022), the United States (Piketty, Saez, and Zucman, 2018), and South Africa (Chatterjee, Czajka, and Gethin, 2023).[6] In our benchmark estimates, in the absence of better information, we thus take the average of the three corresponding tax incidence profiles. We then proportionally scale up the CIT incidence profile in each country-year so as to match total CIT revenue.

Property and Wealth Taxes Property and wealth taxes include taxes on immovable property, wealth taxes, inheritance and gift taxes, and taxes on financial and capital transactions. They are by far the least significant revenue item, averaging 2% of national income and rarely exceeding 4%. Like Piketty, Saez, and Zucman (2018), we assume that residential property taxes are paid by households proportionally to housing wealth, while business property taxes and inheritance, wealth, and financial transaction taxes are distributed proportionally to capital income excluding mixed income and imputed rents (that is, in the same way as corporate taxes).

Unfortunately, we do not observe the concentration of housing wealth, so we assume that residential property taxes are paid proportionally to pretax income. This is consistent with evidence from South Africa and the United States suggesting that the distribution of housing property taxes is relatively flat (Chatterjee, Czajka, and Gethin, 2023; Piketty, Saez, and Zucman, 2018). For other wealth taxes, we use the same corporate tax stylized profile as above.

The data source for total property and wealth tax revenue is Bachas et al. (2022),

[6]See Appendix Figure C.1, which plots these three profiles by generalized percentile.

while we use the OECD tax database (OECD, 2022) to decompose these taxes into housing property, business property, and other taxes on wealth. For countries and years missing in the OECD database, we assume that 50% of property and wealth taxes fall on residential property, while 50% fall on business property and net wealth.

Indirect and Other Taxes As in Blanchet et al. (2021), we assume that indirect taxes are paid by consumers, but we also account for the fact that part of consumption goes untaxed because it is made in the informal sector. First, we estimate income-to-consumption ratios along the income distribution. Second, we estimate the share of informal consumption in total consumption by generalized percentile.

For the first step, our benchmark scenario assumes that the income-to-consumption ratio is logit-shaped and about two times higher for the 99^{th} percentile than for the median (see Appendix Figure C.2). This is in line with evidence from Chancel et al. (2023), who combine data on income-consumption ratios by pretax income percentile from a number of studies and show that this profile provides a good approximation of the typical empirical profile observed.

For the second step, we account for the fact that low-income households tend to purchase goods in informal markets to a greater extent than high-income households. This implies that a greater fraction of their consumption goes untaxed, especially in low-income countries where informality is high. Here, we draw on recent empirical evidence by Bachas, Gadenne, and Jensen (2022), who estimate the share of consumption made in informal markets, by income percentile, in a sample of developing countries. Informality is relatively greater among low-income earners in poor countries than in rich countries.[7] Drawing on this empirical regularity documented in Bachas, Gadenne, and Jensen (2022), we estimate the share of consumption $s_{ct}(p)$ made in the formal market for percentile p in country c at time t as a linear function, whose slope depends on the level of economic development:

$$s_{ct}(p) = p \times \theta_{ct} \tag{3.4}$$

$$\theta_{ct} = \alpha + \beta GDP_{ct} \tag{3.5}$$

Where GDP_{ct} denotes GDP per capita, expressed in constant 2021 PPP USD. Accounting for informality makes indirect taxes significantly less regressive, in particular in low-income countries, although this effect is generally not sufficiently strong to make them progressive as a share of income.[8]

[7]See Appendix Figure C.3.

[8]Appendix Figure C.4 illustrates how accounting for informality changes the progressivity of

Finally, other residual taxes include a number of miscellaneous items, such as user fees, penalties, fines, and poll taxes, which usually represent less than 0.5% of national income. These taxes are generally not conditioned on income or consumption, which implies that their burden is much higher among low-income groups than high-income groups when expressed in proportion of their income. Accordingly, we make the simplifying (and probably conservative) assumption that they are distributed similarly to indirect taxes, that is, in a regressive way.

Social Contributions We also construct estimates of the distribution of social contributions. Social insurance systems are already accounted for in pretax income, so we do not need to deduct social contributions to reach posttax income. However, we still estimate their incidence to arrive at a more comprehensive view of the magnitude and progressivity of the tax system in each country.

We assume that social contributions are paid proportionally to labor income, excluding income that is not taxed due to exemptions or evasion. To do so, we rely on a unique database provided by the International Labor Organization (ILO), which compiles labor force surveys fielded in about 150 countries since the 1990s. For approximately 110 countries, we observe whether individuals paid social contributions, and estimate the propensity to do so along the labor income distribution. Informal work and exemptions are generally more prevalent at the bottom of the distribution, while capital income is more prevalent at the top. As a result, middle-income groups often display the highest effective tax rates.[9]

Social Assistance Benefits Social assistance expenditure consists in both cash and in-kind transfers received by households, such as conditional cash transfers and food stamps, as defined in the system of national accounts (see Eurostat, 2019). Note that social assistance excludes social insurance transfers (mainly unemployment and pension benefits), which are already included in our definition of pretax income, as discussed above. Data on aggregate expenditure come from Gethin (2023b), who draws on various sources to derive harmonized series on the evolution of spending on social assistance programs around the world.

Data on the incidence of social transfers come from four sources: Piketty, Saez, and Zucman (2018) for the United States, Blanchet, Chancel, and Gethin (2022) for 30 European countries, the World Bank's ASPIRE database for 101 countries (World

indirect taxes in Niger, one of the poorest countries in our sample.

[9]Appendix Figure C.5 illustrates how accounting for informality and exemptions changes our estimates of the incidence of social contributions, in the context of Argentina in 2019.

Bank, 2018), and the database of the Commitment to Equity Institute for 3 countries (Iran, Togo, and Venezuela; Lustig, 2023). For the 45 countries not covered by any of these sources, our benchmark scenario allocates transfers using the average profile observed in all countries.

Education We consider two alternative scenarios for the distribution of education spending. One option is to allocate education proportionally to posttax disposable income (pretax income, minus direct taxes, plus cash transfers), in line with what was done for DINA studies covering the United States (Piketty, Saez, and Zucman, 2018) and Europe (Blanchet, Chancel, and Gethin, 2022). Another option is to allocate education spending to children attending school in the household. This approach has been adopted by DINA studies covering Latin America (De Rosa, Flores, and Morgan, 2022b) and South Africa (Gethin, 2022), among others, as well as by the CEQ institute in a number of studies (Lustig, 2018). Gethin (2023b) extends this approach to all countries in the world since 1980, combining data on education spending with a unique set of surveys covering school attendance and household income worldwide.

The school attendance approach has the advantage of allocating education expenditure to individuals actually benefiting from the education system at a given point in time. The main disadvantage is that it can be sensitive to various demographic and compositional factors overestimating the progressivity of education spending. For instance, education spending may appear progressive mainly because low-income households tend to have more children, or because households with children tend to have young parents with lower incomes. Students attending university and living alone may also appear in survey data as a particularly poor household, making tertiary education spending implausibly progressive. There may also be large inequalities in school spending and school quality across geographical areas, which are generally not observed. For all these reasons, while education spending is probably more equally distributed than posttax disposable income, it should also probably be allocated in a more unequal way than the school attendance approach suggests.

In the main results, we thus present series with education spending allocated proportionally to posttax disposable income. We reproduce all findings with the school attendance approach in the appendix, drawing on estimates from Gethin (2023b). We view the construction of more precise measures of the distribution of education spending, such as indicators relying on public education transfers that children can expect to receive as a function of their socioeconomic background, as an important target for future research.[10]

Health and Other Transfers Data on the distributional incidence of healthcare come from Gethin (2023b), who mostly relies on series from the CEQ database (Lustig, 2018). In line with other DINA studies, all other government expenditure is distributed proportionally to posttax disposable income, that is, in a distributionally neutral way. This includes spending on transport, public order and safety, administration, defense, and all other types of public goods.

3.1.4 Comparison With Existing DINA Studies

Our compilation of data from earlier DINA studies covering 45 countries allows us to verify to what extent our simplified methodology provides a good approximation of patterns of fiscal progressivity across countries and over time. If the validation exercise shows that our new estimates match the sample of existing estimates, we can more confidently trust these new estimates outside of that sample.

One major difficulty is that the DINA studies collected for this validation exercise are not always perfectly comparable with one another. Two main issues should be stressed in particular. First, existing DINA studies do not always use the exact same methodology to allocate each type of tax. For instance, Piketty, Saez, and Zucman (2018) distribute business property taxes proportionally to corporate equity, while other DINA studies most often distribute them either proportionally to pretax income or in ways undocumented by the authors. Similarly, the quality of data available to measure the concentration of corporate equity varies tremendously across countries, from exceptionally detailed administrative data in the Netherlands (Bruil et al., 2022) to dividends and employer income reported in surveys in the case of Latin America (De Rosa, Flores, and Morgan, 2022b).

Second, and partly because of limitations in data sources available, effective tax rates paid by percentile can be very noisy in a number of existing DINA studies. For instance, Blanchet, Chancel, and Gethin (2022) rely on surveys to measure the distribution of direct taxes, which makes estimates of their progressivity quite noisy from one year to another, especially at the top of the distribution. More importantly, all DINA studies rely on surveys reporting the joint distribution of pretax income and consumption to allocate indirect taxes. Because of the existence of many zero or very low pretax incomes in such surveys, consumption-to-income ratios can easily diverge,

[10] See for instance Piketty (2022), Figure 32, documenting large inequalities in public education spending received by French cohorts.

making estimates of the distributional incidence of consumption taxes particularly volatile. In South Africa, for instance, the bottom 50% pretax income share is less than 3%, leading effective tax rates as a share of pretax income to diverge towards infinity for most households within this group (Chatterjee, Czajka, and Gethin, 2023).

With these limitations in mind, Figure 3.1 compares our estimates of the effective tax rates faced by percentiles *p50*, *p75*, *p90* and *p99* to those of existing DINA studies. With few exceptions, our estimates are clustered along the 45-degree line, suggesting that our simplified approach does a good job at reproducing broad cross-country and time variations in taxes paid by different pretax income groups.

We provide three additional validation exercises in the appendix. First, we compare our measures of absolute progressivity by type of tax to those reported in existing DINA studies (see Appendix Figure C.6). The two estimates fall very close to each other in the case of personal income taxes and corporate taxes. However, because of the issue of low pretax incomes highlighted above, the fit of indirect taxes is much more variable. Given well-known challenges at measuring the relationship between income and consumption in surveys (Chancel et al., 2023), whether our smoothed estimates or those of existing DINA studies are more reliable is difficult to say. On average, however, it is reassuring that our measures of the progressivity of indirect taxes falls quite close to average progressivity found in existing work.

Second, we zoom in on effective tax rates paid by income group in specific countries, focusing on DINA studies with the least volatile estimates. Appendix Figures C.7, C.8, and C.9 present this comparison for the United States, the Netherlands, and South Africa, respectively. Although our estimates are not perfect, our simplified methodology reproduces the strong regressivity of taxes in the Netherlands and the relatively more progressive tax systems of the United States and South Africa remarkably well.

In a third validation of our method, we compare our estimates of overall tax progressivity against those of existing studies in each country (see Appendix Figure C.10). Progressivity is measured as the percent change in the top 10% to bottom 50% average income ratio obtained when removing taxes from pretax incomes (see section 3.2.1 for more details on this indicator). Because of issues highlighted above, our estimates of this indicator unsurprisingly do not correlate perfectly with those of existing papers, yet there does appear to be a strong and positive relationship. We view this as additional reassuring evidence that our methodology captures broad cross-country variations in tax systems relatively well.

3.1.5 Integration of Existing DINA Studies in our Database

Finally, while our estimates accurately capture broad variations in tax and transfer progressivity, existing DINA studies should be considered as of better quality, given that they rely on actual country-specific surveys and tax data to allocate taxes and transfers. We thus replace our series with those of Piketty, Saez, and Zucman (2018) for the United States, Blanchet, Chancel, and Gethin (2022) for Europe, and De Rosa, Flores, and Morgan (2022b) for Latin America. Given the lack of detailed tax and transfer incidence profiles comparable to ours, we only replace series covering total taxes paid by percentile and posttax income distributions (with the exception of the United States, for which we also replace transfers).

European and Latin American series only cover a subset of our period of interest, generally corresponding to the post-2000 period. To ensure time consistency, we thus adjust our 1980-2000 series based on the difference observed between our series and theirs in the first year available. For taxes paid, we rescale effective tax rates paid by generalized percentile based on the ratio of ETRs between the two sources. For posttax inequality series, we rescale the average income of each generalized percentile based on the ratio of average incomes between the two sources.

A last adjustment comes from the fact that Piketty, Saez, and Zucman (2018) and Blanchet, Chancel, and Gethin (2022) allocate education spending proportionally to posttax disposable income, while we allocate it based on school attendance of children in the household, as in De Rosa, Flores, and Morgan (2022b). To ensure that the final series are conceptually consistent, we thus remove education distributed proportionally from the European and U.S. series and add back education distributed based on the school attendance approach (taken from Gethin, 2023a). For Latin America, we leave the series unchanged, given that education is allocated using a method conceptually similar to ours. Appendix Figures 3.9 to 3.18, as well as Appendix Table 3.2, show that our main findings remain robust to distributing education spending proportionally to posttax disposable income.

3.2 A Global Perspective on Government Redistribution

This section presents the main results on levels and trends in government redistribution around the world. Section 3.2.1 presents facts on tax progressivity, while section 3.2.2 turns to the analysis of transfers and the overall effect of government redistribution on inequality. Finally, section 3.2.3 investigates the role played by differences in the distribution of pretax incomes ("predistribution") versus taxes and transfers ("redistribution") in explaining cross-country differences in inequality.

3.2.1 Levels and Trends in Tax Progressivity

3.2.1.1 A Global Map of Tax Progressivity

Taxes Are Weakly Progressive or Regressive in Most World Regions We start by documenting worldwide differences in the size and structure of taxes. Figure 3.2 shows the evolution of aggregate tax revenue by world region between 1980 and 2019. For simplicity and tractability, we divide the world in six groups of countries throughout the paper: the Anglosphere (United States, United Kingdom, Canada, Australia, and New Zealand), Western Europe, Eastern Europe (including Russia), Latin America, Asia, and Africa. We then calculate total tax revenue as a share of national income in each country and plot the resulting population-weighted average by world region.

Total taxation has increased in Asia, Latin America, and Western Europe, while it has remained stable in Africa, the Anglosphere, and Eastern Europe. Western Europe and Anglosphere countries stand out as having much larger tax revenue from personal income taxes, while indirect taxes are more widespread in other world regions. Overall, there have not been major changes in the composition of taxes within each region, although there are some exceptions. In Eastern Europe, in particular, corporate tax revenue has declined significantly at the same time as indirect taxation has expanded as a share of national income.

Figure 3.3 plots the 2019 average effective tax rate (ETR) faced by each percentile of the pretax income distribution in different regions of the world. Throughout this section, we include social contributions in our analysis of tax progressivity (results excluding social contributions are qualitatively similar). Two main results stand out. First, consistently with Figure 3.2, there are large differences in aggregate tax rates between regions, with macroeconomic tax rates being lowest in Sub-Saharan Africa (10-20%) and highest in Western Europe (over 40%). Second, differences between income groups are small in most regions: nowhere in the world does the average ETR of the top 10% earners exceed that of the bottom 50% by more than 10 percentage points. In Africa, Asia, and Western Europe, taxes paid are essentially flat throughout the income distribution, while they are slightly more progressive at the top in the Anglosphere. Latin America and especially Eastern Europe are the only regions where tax systems are unambiguously regressive. Indeed, Eastern

European (and ex-Soviet) countries tend to rely heavily on indirect taxes as a source of revenue (approximately 15% of national income, while closer to 9% in the rest of the world), and have moved toward flat taxation of household income in recent decades.

Taxes Have Little Effect on Inequality in Most Countries Given limited variations in effective tax rates along the income distribution in most regions of the world, one should not expect taxes to play a substantial role in reducing inequality. Figure 3.4 presents a global map of tax progressivity in 2019, providing a more granular picture on cross-country differences in the distribution of taxes worldwide. We summarize the progressivity of taxes with a simple indicator: the percent difference in inequality, measured as the top 10% to bottom 50% average income ratio, before and after removing taxes from individual incomes:

$$\gamma_T = \frac{r_{pre} - r_{net}}{r_{pre}} \tag{3.6}$$

Where *pre* refers to pretax income, *net* refers to net-of-tax income (pretax income minus taxes), and $r = \frac{\bar{y}_{p90p100}}{\bar{y}_{p0p50}}$ is the ratio of the average income (pretax or net of taxes) of the top 10% richest to that of the bottom 50% poorest individuals in each country-year.

Positive values thus indicate progressive tax systems, while negative values indicate regressive tax systems. As shown in Figure 3.4, taxes have little effect on inequality: in many countries, they reduce the inequality ratio r by less than 5%. The geographical patterns documented in Figure 3.3 clearly stand out. Latin American and Eastern Europe countries have strongly regressive tax systems. Western European and Southern African countries display the most progressive tax systems, although the magnitude of the effect is generally small, on the order of 5-15%.

Robustness to Other Indicators A concern with this analysis is that this indicator of tax progressivity may be not be perfectly comparable across countries. In countries with higher pretax inequality, in particular, taxes may appear mechanically more progressive. The overall impact of taxes may also end up being mechanically higher in countries with greater aggregate tax revenue (see discussion in Appendix C.2). As an alternative to this measure of "absolute" progressivity, we thus consider two other indicators, "relative" and "normalized" progressivity. Relative progressivity corresponds to the percent difference in the effective tax rates of the top 10% and bottom 50% in each country. Normalized progressivity corresponds to absolute progressivity computed over a single, "normalized" distribution, which ensures that it is insensitive to differences in pretax inequality across countries. Maps comparable to Figure 3.4 are presented for these indicators in Appendix Figures C.24 and C.25. The results are similar.

3.2.1.2 Trends in Tax Progressivity Since 1980

Tax Progressivity Has Stagnated in Most World Regions We now turn to documenting trends in tax progressivity worldwide. To start, consider Figures 3.5 and 3.6, which plot the level and composition of taxes paid by percentile in the average country in 1980 and 2019. This figure is constructed by dividing taxes by pretax income for each percentile in each country, and then taking the population-weighted average of this indicator over all countries in the world.

Two results stand out. First, there has been an increase in worldwide taxation, which ranged from 18-22% of income in 1980, and increased to 22-26% by 2019. Second, there has been no clear change in average worldwide tax progressivity since 1980; if anything, tax progressivity has declined. Overall, top-income groups face slightly higher effective tax rates than earners at the middle of the income distribution, because of the particularly progressive nature of personal income and corporate income taxes. Yet taxes are also slightly higher at the very bottom of the distribution, where consumption is high relative to pretax income and the burden of indirect taxes is thus particularly large. While direct taxes have grown (and PIT systems have become slightly more progressive), so have indirect taxes, leading to little change in average tax progressivity.

Figure 3.7 decomposes this general result geographically by showing the evolution of tax progressivity by world region. Eastern Europe has seen a particularly pronounced and steady decline in progressivity: taxes had more or less no effect on the income distribution in 1990, while they increased inequality by over 25% in 2019. In all other regions, tax progressivity has remained remarkably stable since 1980, mirroring the overall pattern documented in Figures 3.5 and 3.6.

There Has Been No Cross-Country Convergence in Effective Tax Rates Increases in average tax rates coupled with differences in progressivity imply that taxation has changed differentially for different income groups. We bring these dynamics into focus at the regional level, charting top 1%, top 10%, and bottom 50% effective tax rates since 1980, by region and on average, in Appendix Figures C.11, C.12, and C.13. Top 1% effective tax rates have declined substantially in the

Anglosphere and Eastern Europe. Western Europe has overtaken the Anglosphere as the region that taxes the richest the most, but the gap is even greater among low incomes, which explains why tax progressivity is still higher in the latter. Eastern Europe began the post-Soviet era on a par with Western neighbors for top-income taxation, but since then have reverted toward the global mean. No countries tax their poorest citizens as much as do the countries of Eastern Europe. Africa stands out as the only region with no significant change in taxation at all: on average, effective tax rates have remained low and stable for all income groups. All in all, there is no clear convergence between countries in effective tax rates paid.

3.2.2 Levels and Trends in Total Government Redistribution

3.2.2.1 The Distribution of Government Transfers

We now turn to the analysis of transfers, including social assistance, education, healthcare, and other public goods. Figure 3.8 plots the average share of national income received by the bottom 50%, the middle 40%, and the top 10% in the form of cash and in-kind transfers by world region in 2019. Appendix Figure C.33 reproduces this figure with education distributed using the school attendance approach.

The Size of Transfers Varies Substantially Across Regions There are large differences across regions in the amount of transfers received by low-income groups, with total expenditure received by the bottom 50% ranging from about 6% of national income in Africa to 18% in Western Europe. On average, cash transfers, healthcare, education, and other public goods each represent about a quarter of transfers received, but with substantial variations across regions. Redistribution in the form of social assistance is particularly developed in Europe, while public healthcare spending is exceptionally large in the United States (and targeted to the poor lacking private insurance). In Africa and Asia, in-kind transfers represent the bulk of redistribution.

The Progressivity of Transfers Varies Substantially Across Regions The countries of Western Europe and the Anglosphere particularly stand out for both *relative* and *absolute* progressivity. In Latin America, Asia, and Africa, on the other hand, top earners receive a greater share of government transfers than do the bottom 50% of the income distribution. This is mainly the result of our assumption that transfers other than social assistance and healthcare are received proportionally to disposable income, that is, in a very unequal way. Because Latin America, African, and Asian countries spend little on these functions of government, public expenditures appear to be the least progressive in these regions. Even under this conservative assumption on the low progressivity of public goods other than healthcare, however, government transfers are unambiguously progressive.

3.2.2.2 The Net Impact of Taxes and Transfers on Inequality

Tax-and-Transfer Systems Always Reduce Inequality, But With Large Variations Combining taxes and transfers, our database allows us to provide a global map of government redistribution, in Figure 3.9. Appendix Figure C.34 reproduces this figure with education distributed using the school attendance approach. The "extent of redistribution" is measured as the percent difference in the top 10% to bottom 50% average income ratio, as in equation (3.6) above (and in, e.g., Bozio et al., 2022).

Two results stand out. First, tax-and-transfer systems always reduce inequality: the indicator is strictly positive in all countries in the world. Second, there are large variations in the extent of redistribution, ranging in 2019 from less than 10% in several Sub-Saharan countries to over 30% in countries such as the United States, Norway, and South Africa. Overall redistribution follows clear regional patterns, being highest in Northern America and Europe, and lowest in Latin America, Sub-Saharan African (excluding Southern Africa), and Asia.

Figure 3.10 shows that, in all regions of the world, tax-and-transfer systems mostly redistribute income from the top 10% to the bottom 50%. Appendix Figure C.35 reproduces this result with education distributed using the school attendance approach. On net, the middle 40% generally neither benefit nor lose much from the tax-and-transfer system. The net transfer received by the bottom 50% is highest in the Anglosphere and Western Europe, and lowest in Asia and Africa.

Transfers Account for 90% of Redistribution Combining our previous results on the lack of strong tax progressivity and large differences in the size and distributional incidence of transfers, we can expect transfers to be the dominant drivers of redistribution. We formalize this in Table 3.2, which compares how inequality changes before and after removing taxes and adding transfers to individual incomes. Appendix Table C.2 reproduces these findings with education distributed using the school attendance approach. In 2019, the top 10% to bottom 50% income ratio was approximately $r = 18$ in the average country (calculated as the population-weighted average of the indicator across all countries). Removing taxes barely affects inequality, while adding government transfers reduces inequality by over 3 percentage points. By this measure, taxes account for less than 10% of the effect of government redistribution on inequality, while transfers account for over 90%. There are significant variations across regions: the contribution of taxes reaches about 30% in the Anglosphere and Africa, while it is negative in Eastern Europe and Latin America, where taxes increase inequality. Overall, transfers largely dominate taxes in reducing inequality in most countries in the world.

Table 3.3 provides more detailed results on the redistributive impact of different categories of taxes and transfers.[11] Appendix Table C.3 reproduces these findings with education distributed using the school attendance approach. Estimates from existing DINA studies do not allow us to derive such a detailed decomposition, so this table uses our own estimates for Europe and Latin America, which explains why the results differ slightly from those in Table 3.2. We calculate the progressivity of each type of tax or transfer as the percent reduction in inequality it occasions (as in equation 3.6 above). For instance, the statistic for personal income taxes γ_{PIT} corresponds to the percent reduction in the top 10% to bottom 50% ratio before and after removing personal income taxes from pretax income. Positive values indicate that the tax or transfer reduces inequality, while negative values indicate that it increases inequality.

The first column displays the results in the average country, taking the population-weighted average of the corresponding indicators across all countries in the world. Personal income taxes and corporate taxes each reduce inequality by about 4%, while indirect taxes increase inequality by about 8%. The effect of property and wealth taxes is negligible.

The effect of transfers on inequality is significantly higher: social assistance and healthcare expenditure each reduce inequality by about 10%. All in all, the progressivity of personal income taxes and corporate taxes thus appears to be more or less cancelled by the regressivity of indirect taxes, leading to a tax system that reduces inequality by only 3% in the average country. Meanwhile, all transfers are strongly progressive, which explains why they play a dominant role in reducing inequality.

Interesting regional variation stands out. Personal income taxes play a key role in reducing inequality in the Anglosphere and Western Europe, while indirect taxes increase inequality most in Europe and Latin America. Social assistance is the most significant driver of redistribution in Europe, while healthcare plays a more important role in Africa.

[11] See Appendix Table C.1 for similar results in 1980.

3.2.2.3 Trends in Government Redistribution Since 1980

We now present results on the evolution of overall redistribution using two complementary indicators. Figure 3.11 plots the evolution of the extent of redistribution by world region, measured as the percent reduction in the top 10% to bottom 50% income ratio operated by the tax-and-transfer system. Appendix Figure C.36 reproduces this figure with education distributed using the school attendance approach. This figure tells us whether government redistribution reduces inequality more today than in the past. Meanwhile, Figure 3.12 plots the evolution of the share of national income redistributed to the bottom 50%, which tells us to what extent redistribution increases the incomes of the poorest individuals in each region (see Appendix Figure C.37 for similar results with education distributed using the school attendance approach).

Redistribution has increased in most regions. In the average country, the extent of redistribution increased from about 10% to 20% from 1980 to 2019. This average figure hides considerable heterogeneity, with significant increases in redistribution in Western Europe, the Anglosphere, and Asia compared to complete stagnation in Eastern Europe and Africa. The same result extends to the net transfer received by the bottom 50%, which increased from about 2% to 2.5% of national income in the average country but barely changed in Eastern Europe and Africa. Overall, there is no evidence of cross-country convergence in the redistributive power of tax-and-transfer systems.

3.2.2.4 Government Redistribution Over the Course of Development

We conclude this section with a correlational analysis of the relationship between government redistribution and economic development.

Tax Progressivity Is Uncorrelated With GDP per capita There is little correlation between tax progressivity and per capita income (Figure 3.13). The raw correlation between tax progressivity and GDP per capita is approximately $\rho = -0.09$. In other words, total taxation increases as countries develop, but there is little progressivity in the increase, and little tax progressivity overall: effective taxation on the poorest rises in parallel to effective taxation on the richest, and started at a similar rate. Overall, the tax system appears to increase or reduce inequality by less than 10%, throughout the vast majority of countries in the world.

Transfer Progressivity Is Positively Correlated With GDP By contrast, low-income households benefit from much greater government transfers in rich countries than in poor countries. Figure 3.14 plots the share of national income received by the bottom 50% in the form of cash and in-kind transfers (expressed as a share of national income), against GDP per capita.[12] Appendix Figure C.38 reproduces this figure with education allocated using the school attendance approach. The raw correlation between the two variables is $\rho = 0.64$. In Anglosphere and Western European countries, the bottom 50% receive 15-20% of national income, versus 2-8% in most African countries. Transfers thus appear to reduce inequality much more in high-income countries than in low-income countries. There are interesting exceptions, however. For instance, the bottom 50% benefit from about the same transfer in South Africa as in China, despite the latter being slightly richer.

This positive relationship between transfers and development is not only driven by the fact that high-income countries have larger governments: high-income countries also provide more progressive transfers. Appendix Figure C.14 reproduces Figure 3.14, but focusing on transfers received by the bottom 50% as a fraction of total public spending. There is a large positive relationship between the two variables. In many African countries, less than 25% of government expenditure accrues to the bottom 50%, while this share exceeds 40% in nearly all Anglosphere and Western European countries. This result is driven by the fact that high-income countries spend much more on social assistance and healthcare than low-income countries. The bulk of transfers in low-income countries correspond to other forms of public goods, such as administration or public order and safety, which we distribute proportionally to disposable income, that is, in a highly unequal way.

Net Redistribution Is Positively Correlated with GDP Putting these two results together yields Figure 3.15, which plots GDP per capita versus the percent reduction in the top 10% to bottom 50% income ratio through taxes and transfers (see Appendix Figure C.39 for the same figure with education distributed using the school attendance approach). The raw correlation between total tax-and-transfer progressivity and development is $\rho = 0.53$. Outliers exist—where income is low but progressivity high, or vice versa—but the general trend looks more like that of Figure 3.14 than that of Figure 3.13. The progressivity of transfers dominates that of taxes, and high-income countries generally redistribute through transfers.

[12]This figure slightly differs conceptually from the previous one in that it shows the absolute level of spending rather than transfers expressed as a percentage of income. The result would be similar if we were to express transfers received as a share of pretax income.

High-income countries thus appear to redistribute significantly more than low-income countries, both today and in 1980. This can be seen more clearly in Appendix Figures C.15 and C.16, which plot the evolution of total fiscal progressivity and net transfers received by the bottom 50%, respectively, by country income group. High-income countries redistribute more than lower-income countries, and this gap has not changed much over time—if anything, it has widened. Upper middle-income countries have been catching up since the turn of the century, but the effect is almost entirely explained by China's fiscal transformation.[13]

3.2.3 Predistribution versus Redistribution: A Global Perspective

We conclude this paper with a brief analysis of the relationship between pretax and posttax income inequality. We start by showing that pretax inequality is the dominant driver of cross-country differences in posttax inequality. While tax-and-transfer systems do vary substantially across countries, they do not significantly alter the ranking of which countries are the most or least unequal in the world. Moving beyond this direct effect of taxes and transfers, we then provide suggestive evidence that redistribution may have significant indirect effects on pretax inequality. Accounting for this indirect effect would potentially lead to putting a much greater weight on redistributive policies in accounting for cross-country differences in inequality.

3.2.3.1 Pretax Versus Posttax Inequality

We start by comparing the bottom 50% share in terms of pretax national income and posttax national income in all 151 countries in 2019 (see Figure 3.16, and Appendix Figure C.40 for comparable results with education distributed based on school attendance).[14] This comparison provides direct suggestive evidence on the role of pretax inequality ("predistribution") versus taxes and transfers ("redistribution") in shaping the final distribution of income. If posttax inequality is entirely driven by taxes and transfers and pretax inequality played no role, then pretax and posttax inequality should be uncorrelated. On the contrary, if posttax inequality is entirely

[13]While China's macroeconomic tax rate (i.e., total public revenue from taxes) hovered near 15% of national income in the 1980s and 1990s, it has since risen to more than 25% of national income. See Bachas et al. (2022) for further discussion on the case of China. Taxes have not become more progressive in China, nor are transfers much more targeted towards the poor than they were pre-2000, but the aggregate revenue of China's government allows it to more effectively transfer a larger share of national income to the poorest.

[14]See also Appendix Figure C.17 for comparable results on the top 10% to bottom 50% average income ratio.

driven by the distribution of income before taxes and transfers, then we should expect the ranking of countries to remain exactly the same before and after accounting for taxes and transfers.

The main takeaway is that there is a very strong correlation between pretax and posttax inequality: notwithstanding a few exceptions, the ranking of countries in terms of pretax and posttax income inequality is almost exactly the same. This finding goes in line with previous evidence focusing on Europe and the United States (Blanchet, Chancel, and Gethin, 2022). A useful way of quantifying this relationship is to run a cross-country regression of the posttax bottom 50% income share on the bottom 50% pretax income share in 2019. This regression delivers an R-Squared of over 0.8. By this measure, “predistribution” accounts for over 80% of cross-country variations in income inequality, while “redistribution” accounts for less than 20%.

We extend this analysis to the bottom 50%, top 10%, and top 1% income shares by region in the appendix (see Appendix Figures C.18, C.19, and C.20). The results are similar: regions with the most equal pretax income distributions generally also have the most equal posttax income distributions.

3.2.3.2 Redistribution Versus Pretax Inequality

A natural limitation of the previous analysis is that redistribution might indirectly affect pretax inequality. For instance, greater investments in social assistance, education, and healthcare may play a key role in generating higher pretax income growth for low-income households. Answering this question rigorously would require data sources and identification strategies that go beyond those mobilized in this paper. However, it is still interesting to investigate whether countries redistributing more are also those that display the lowest levels of pretax inequality.

Figure 3.17 plots the extent of redistribution versus the bottom 50% pretax income share across countries in 2019 (see Appendix Figure C.41 for the specification with education distributed based on school attendance).[15] The correlation between the two variables is positive and significant ($\rho = 0.33$): countries with more progressive tax-and-transfer systems display lower levels of pretax inequality on average. There are important exceptions, however, including highly unequal countries with substantial government redistribution (such as the United States and South Africa), but also equal countries with weakly progressive tax-and-transfer systems (such as many Eastern European countries). This modest but positive correlation is again consistent

[15] See also Appendix Figure C.21 for comparable results on the top 10% to bottom 50% average income ratio.

with previous evidence focusing on Europe and the United States (Blanchet, Chancel, and Gethin, 2022).

One concern is that it may be easier to reduce inequality through taxes and transfers in more unequal countries, given that relative incomes at the bottom of the distribution are particularly low in these countries. We thus complement this analysis with a focus on the net transfer received by the bottom 50%, expressed as a share of national income, in Figure 3.18 (see Appendix Figure C.42 for the specification with education distributed based on school attendance). The correlation between this measure of redistribution and the bottom 50% pretax income share is now much higher, reaching $\rho = 0.54$.[16]

The takeaway is that taxes and transfers could well contribute to strongly reducing pretax inequality indirectly. This would potentially lead to putting a much greater weight on redistributive policies in explaining cross-country differences in inequality. There are still important exceptions, however: for instance, South Africa redistributes more than India, yet displays dramatically higher levels of pretax inequality. Similarly, Latin American countries are characterized by high levels of pretax inequality at the same time as quite progressive tax-and-transfer systems. Our analysis suggests that higher redistribution can lead to lower pretax inequality, but this is far from an iron law. Understanding the conditions under which redistributive policies successfully curb income disparities and their exact contribution to cross-country differences in predistribution represents a fruitful avenue for future research.

3.3 Conclusion

In this paper, we have constructed new estimates of the distributional incidence of taxes and transfers in 151 countries from 1980 to 2019. Combining data from several sources on tax-and-transfer progressivity, we derived estimates of redistribution that are consistent, comprehensive, and comparable across countries and over time. We showed that our simplified methodology is able to replicate results from existing work remarkably well.

Drawing on this database, we have uncovered a number of new stylized facts on the evolution of fiscal progressivity around the world since 1980. Most strikingly, we have documented that the global profile of taxation was and has remained essentially flat. Anglosphere countries, despite recent well-documented decreases in tax progressivity,

[16] See Appendix Figure C.22 for comparable results on the top 10% to bottom 50% average income ratio.

remain the countries whose taxes do the most to reduce inequality. Other regions' tax profiles are less progressive—and, in the case of many Latin American and Eastern European countries, even regressive overall.

Because transfers strongly benefit low-income households, however, tax-and-transfer systems always reduce inequality. They do so much more in high-income than in low-income countries, mainly because the former display larger welfare states, but also because they better target government transfers towards low-income households. There has been little cross-country convergence in redistribution. If anything, the gap has only widened: from 1980 to 2019, the share of national income transferred to low-income households increased in Western Europe and the Anglosphere while it stagnated in Africa.

As a result, taxes and transfers have done little to change the global picture of inequality. In a static sense, predistribution matters demonstrably more than redistribution, explaining about 80% of cross-country variations in posttax income inequality. And the consequences of inequality in redistribution across countries are stark: the poorest people, in the poorest countries, benefit less from redistribution and public services than do the poorest in richer countries.

There remains a need to better understand what drives differences in distribution and redistribution, across countries and over time. For any society the optimal levels, composition, and distributional incidence of taxes and transfers must surely depend on a range of factors whose investigation lies beyond the scope of this article. We hope that the new database constructed in this paper—which estimates the levels, composition, and distributional incidence of all taxes and transfers, worldwide since 1980—will contribute to further evidence-based examination of efficiency and equity in fiscal policy.

Bibliography

Bachas, Pierre, Lucie Gadenne, and Anders Jensen (2022). "Informality, Consumption Taxes and Redistribution". NBER Working Paper 27429.

Bachas, Pierre et al. (2022). "Globalization and Factor Income Taxation". NBER Working Paper 29819.

Blanchet, Thomas (2022b). "Estimates of Capital and Labor Income Distributions in the World Inequality Database". Unpublished Technical Note.

Blanchet, Thomas and Lucas Chancel (2016). "National Accounts Series Methodology".

Blanchet, Thomas, Lucas Chancel, and Amory Gethin (2022). "Why Is Europe More Equal Than the United States?" In: *American Economic Journal: Applied Economics* 14.4, pp. 1–40.

Blanchet, Thomas et al. (2021). *Distributional National Accounts Guidelines: Methods and Concepts Used in the World Inequality Database*. Tech. rep.

Bozio, Antoine et al. (2022). "Predistribution vs. Redistribution: Evidence from France and the U.S." WID.world Working Paper.

Bruil, Arjan et al. (2022). "Inequality and Redistribution in the Netherlands". Working Paper.

Chancel, Lucas and Thomas Piketty (2021). "Global Income Inequality, 1820–2020: the Persistence and Mutation of Extreme Inequality". In: *Journal of the European Economic Association* 19.6, 3025–3062.

Chancel, Lucas et al. (2022b). *World Inequality Report 2022*. Belknap Press of Harvard University Press.

Chancel, Lucas et al. (2023). "Income inequality in Africa, 1990–2019: Measurement, patterns, determinants". In: *World Development* 163.1, pp. 106–162.

Chatterjee, Aroop, Léo Czajka, and Amory Gethin (2023). "Redistribution without Inclusion? Inequality in South Africa Since the End of Apartheid". Working Paper.

De Rosa, Mauricio, Ignacio Flores, and Marc Morgan (2022b). “More Unequal or Not as Rich? Revisiting the Latin American Exception”. Stone Center for Inequality Studies Working Paper Series.

Eurostat (2019). *Manual on sources and methods for the compilation of COFOG statistics*. Eurostat.

Fisher-Post, Matthew, Nicolas Herault, and Roger Wilkins (2022). “Distributional National Accounts for Australia, 1991-2018”. Working Paper.

Fullerton, Don and Gilbert E Metcalf (2002). “Tax incidence”. In: *Handbook of public economics* 4, pp. 1787–1872.

Gethin, Amory (2022). “Who Benefits from Public Goods? Evidence from South Africa”. Working Paper.

– (2023a). “Distributional Growth Accounting: Education and the Reduction of Global Poverty, 1980-2022”. Working Paper.

– (2023b). “Revisiting Global Poverty Reduction: Public Goods and the World Distribution of Income, 1980-2022”. Working Paper.

Jensen, Anders (2022). “Employment structure and the rise of the modern tax system”. In: *American Economic Review* 112.1, pp. 213–234.

Kakwani, Nanak C (1977). “Measurement of tax progressivity: an international comparison”. In: *The Economic Journal* 87.345, pp. 71–80.

Lambert, Peter J (1992). *The distribution and redistribution of income*. Springer.

Lustig, Nora, ed. (2018). *Commitment to Equity Handbook: Estimating the Impact of Fiscal Policy on Inequality and Poverty*. Brookings Institution Press.

– (2023). *Commitment to equity handbook: Estimating the impact of fiscal policy on inequality and poverty*. Brookings Institution Press.

Musgrave, Richard A (1953). “On incidence”. In: *Journal of Political Economy* 61.4, pp. 306–323.

OECD (2022). *Revenue Statistics 2022*. OECD.

Peter, Klara Sabirianova, Steve Buttrick, and Denvil Duncan (2010). “Global reform of personal income taxation, 1981-2005: Evidence from 189 countries”. In: *National Tax Journal* 63.3, pp. 447–478.

Piketty, Thomas (2022). *A Brief History of Equality*. Harvard University Press.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman (2018). “Distributional National Accounts: Methods and Estimates for the United States”. In: *Quarterly Journal of Economics* 133.2, pp. 553–609.

Saez, Emmanuel, Joel Slemrod, and Seth H. Giertz (2012). “The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review”. In: *Journal of Economic Literature* 50.1, pp. 3–50.

Strecker, Nora (2021). Personal communication.

Tax Foundation (1967). “Tax Burdens and Benefits of Government Expenditure by Income Classes 1961 and 1965”. In: *Research Publication* 9.

UN SNA (2008). *The System of National Accounts*. United Nations.

Vegh, Carlos A and Guillermo Vuletin (2015). “How is tax policy conducted over the business cycle?” In: *American Economic Journal: Economic Policy* 7.3, pp. 327–370.

World Bank (2018). *The state of social safety nets 2018*.

World Bank (2022a). *Correcting Course: Poverty and Shared Prosperity 2022*. World Bank.

Figure 3.1: Validation: Comparison of Effective Tax Rates to Existing DINA Studies
at p50, p75, p90 and p99

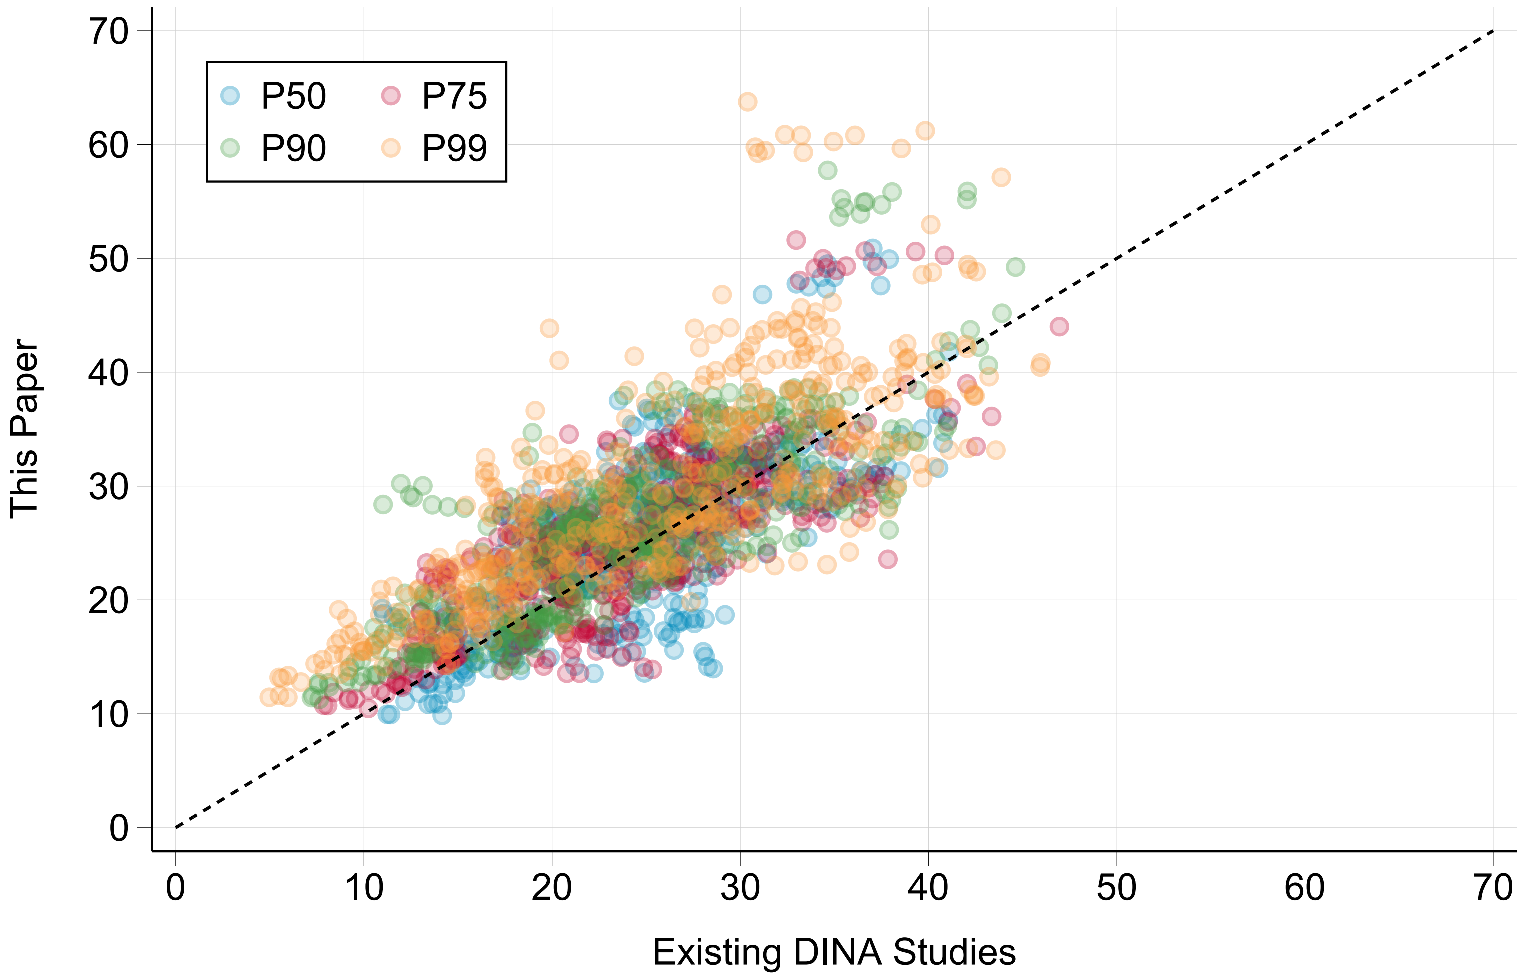

Notes. Axes represent effective tax rate at indicated points along the income distribution.

Figure 3.2: Tax Revenue by World Region, 1980-2019

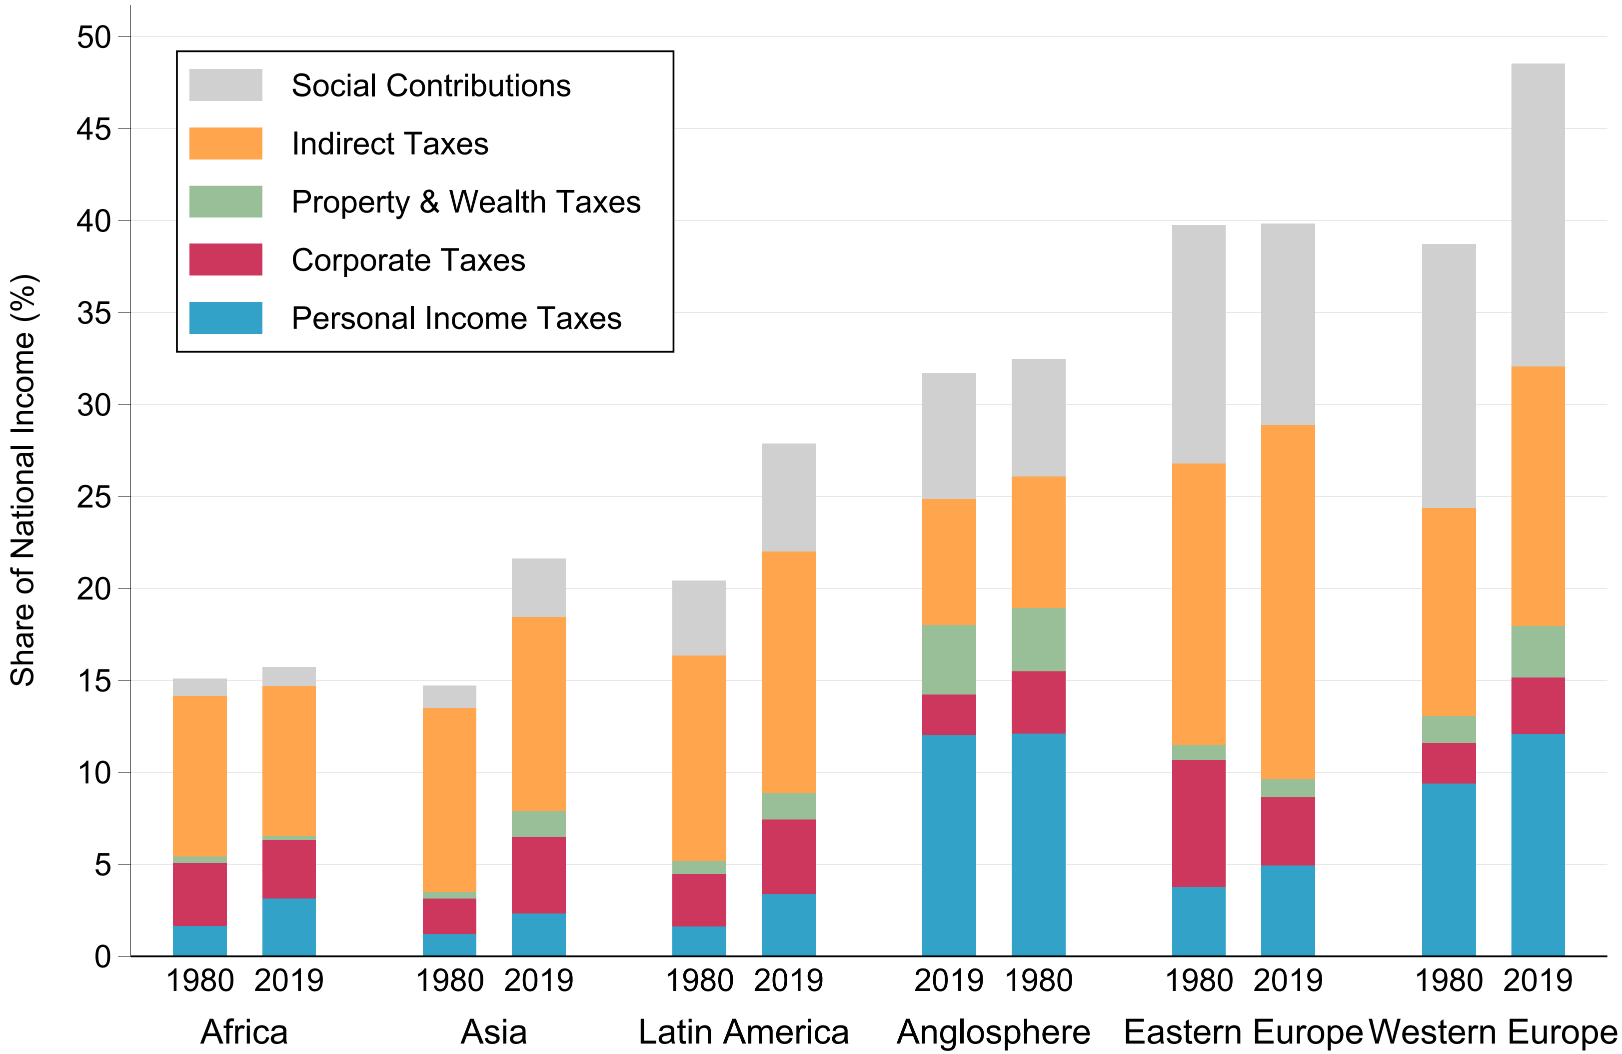

Notes. Population-weighted averages of tax revenue aggregates in each country. Data from Bachas et al. (2022).

Figure 3.3: Effective Tax Rate by Income Group and World Region, 2019

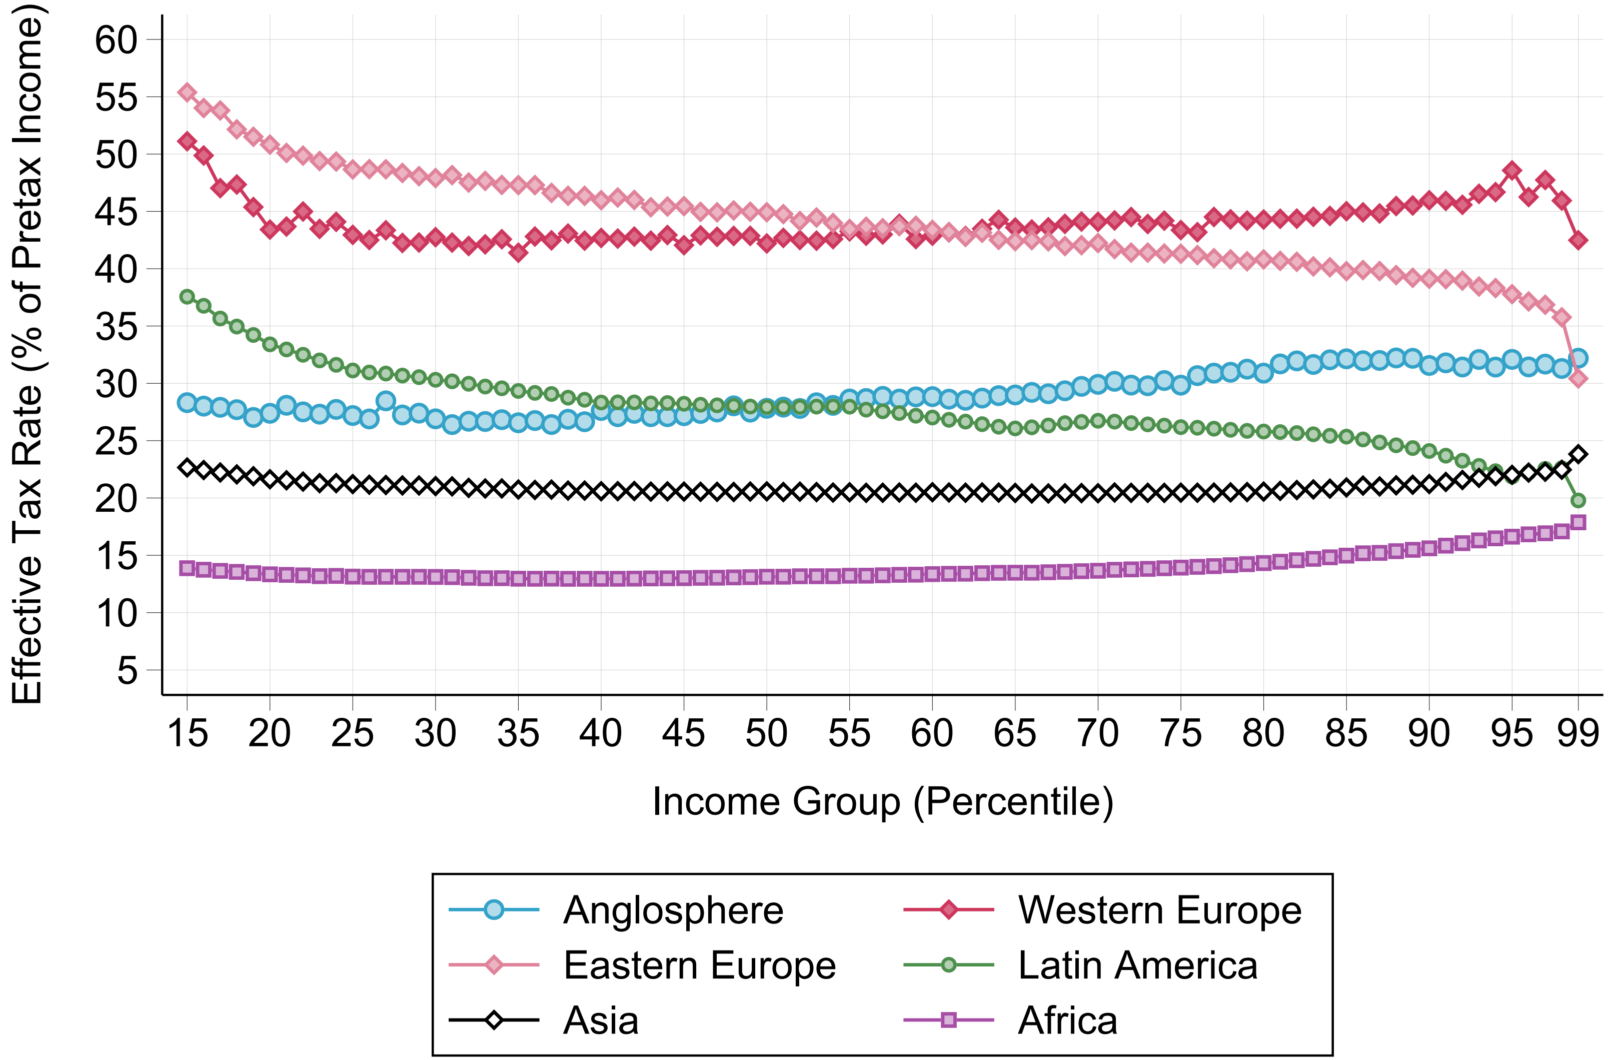

Notes. Population-weighted averages of effective tax rates by percentile in each country. Taxes include social contributions.

Figure 3.4: Tax Progressivity Around the World:
Percent Reduction in Top 10% to Bottom 50% Average Income Ratio (Pretax versus Net-of-tax Income)

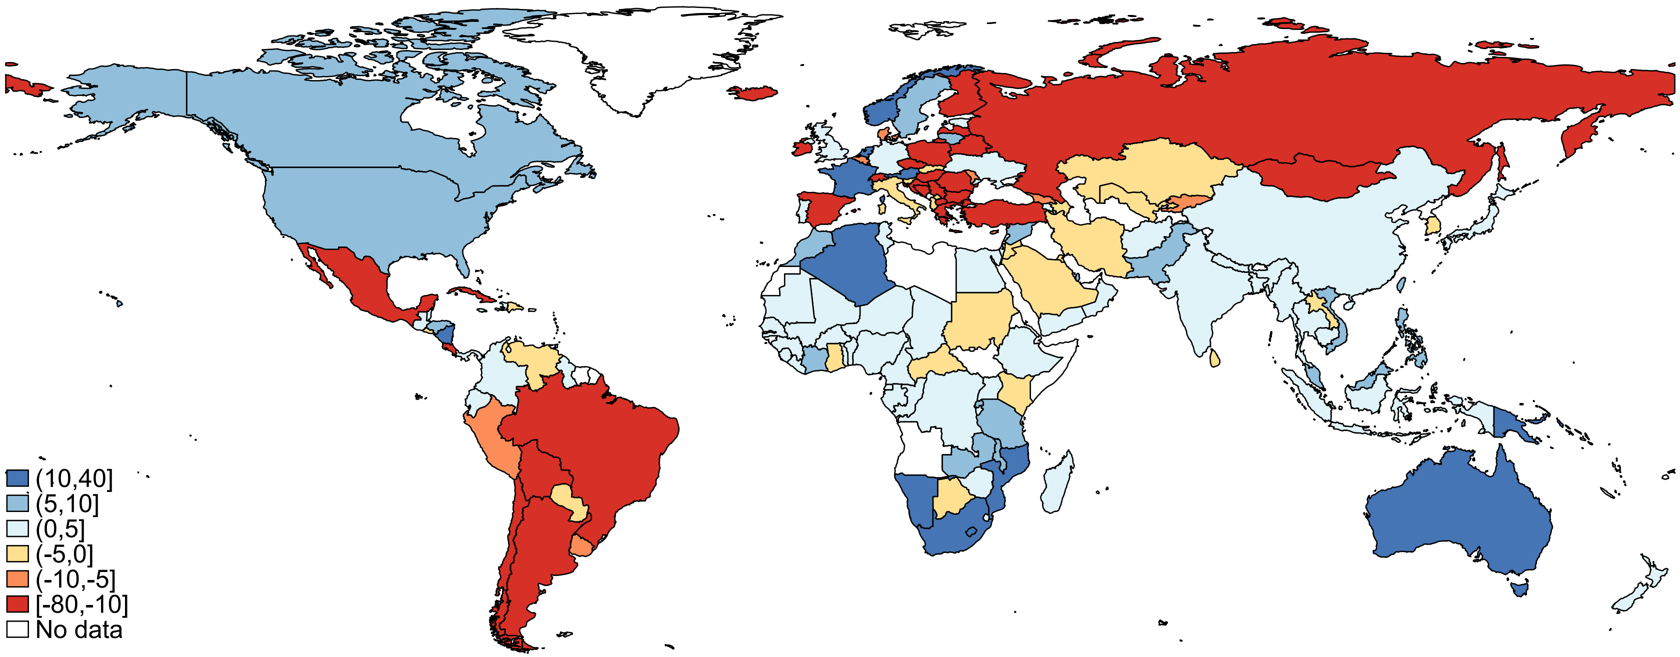

Notes. Net-of-tax income: pretax income minus taxes. Taxes include social contributions.

Figure 3.5: Composition of Taxes Paid by Percentile: Global Average, 1980

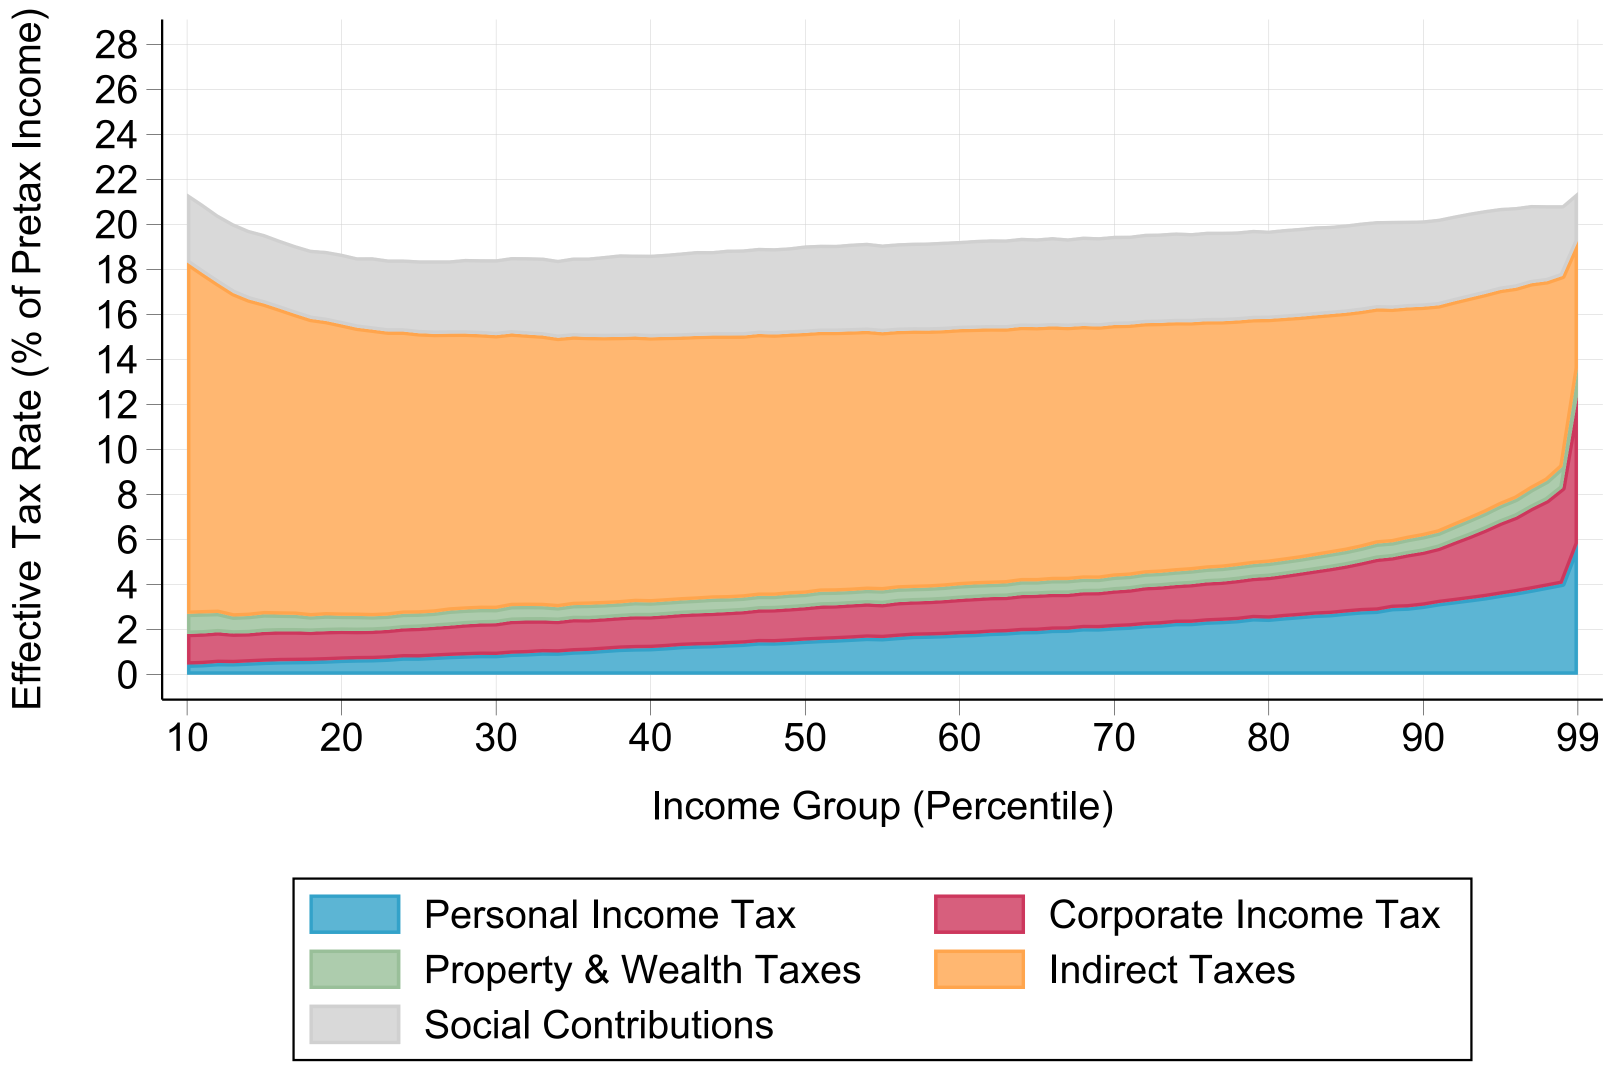

Notes. Population-weighted averages of effective tax rates by percentile in each country.

Figure 3.6: Composition of Taxes Paid by Percentile: Global Average, 2019

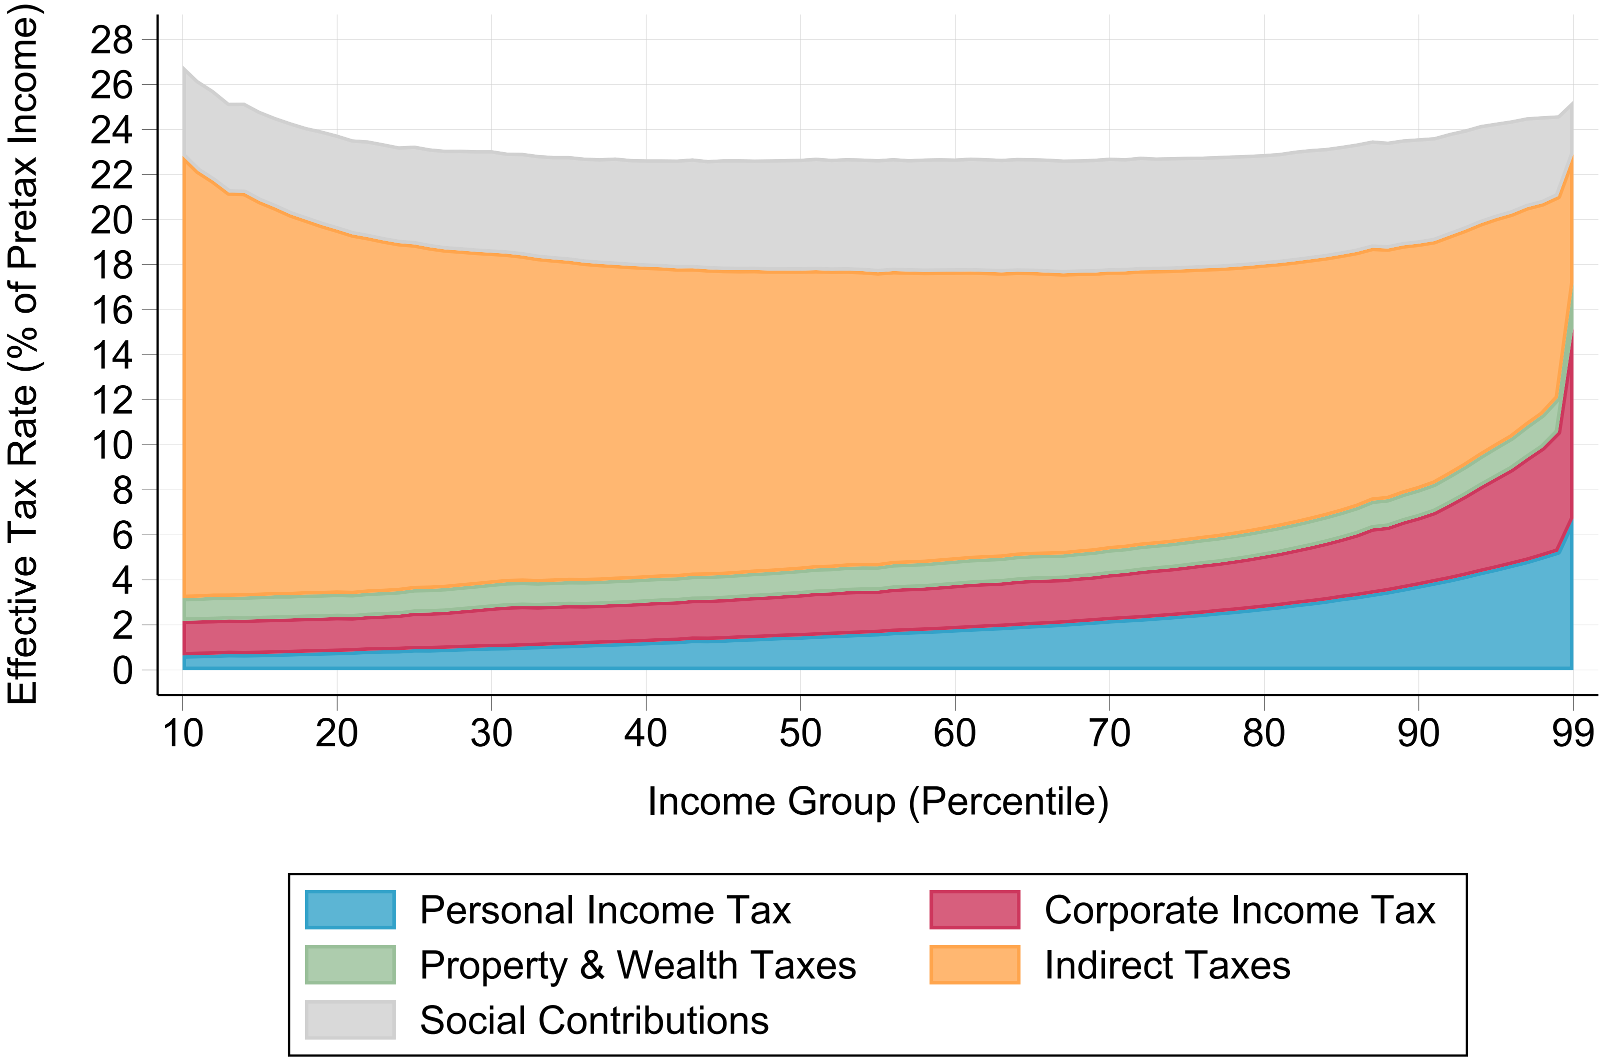

Notes. Population-weighted averages of effective tax rates by percentile in each country.

Figure 3.7: Tax Progressivity by World Region, 1980-2019:
Percent Reduction in Top 10% to Bottom 50% Average Income Ratio (Pretax versus Net-of-tax Income)

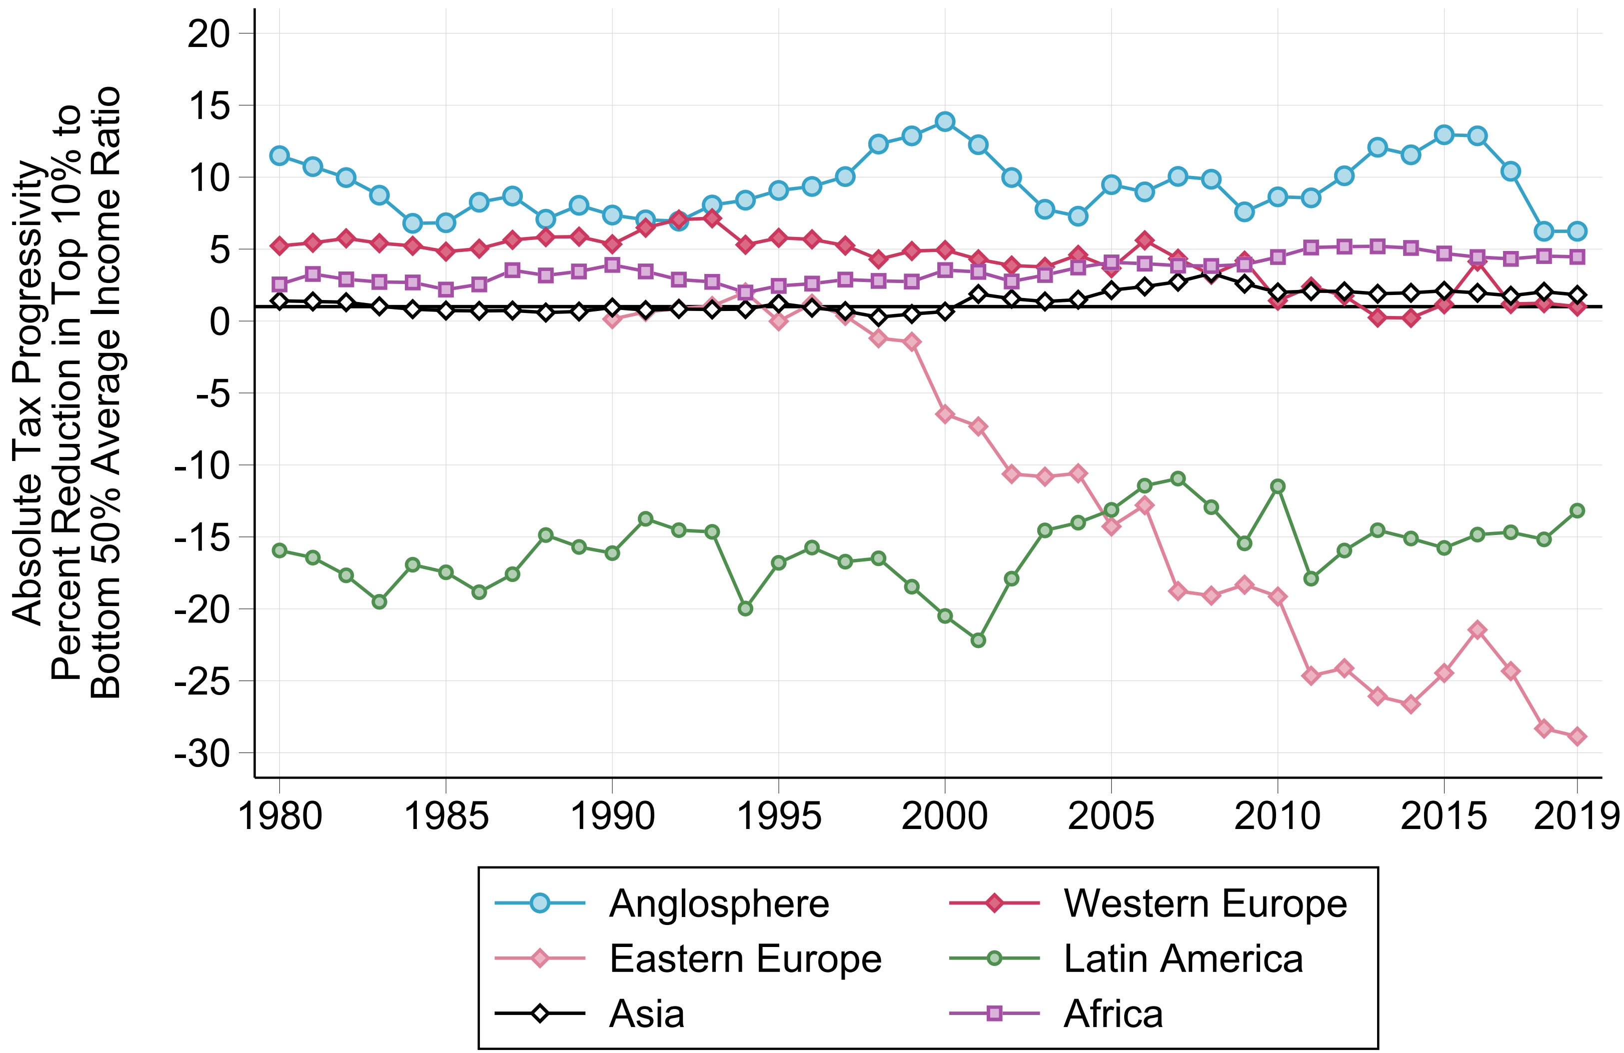

Notes. Net-of-tax income: pretax income minus taxes. Taxes include social contributions. Population-weighted averages of tax progressivity in each country.

Figure 3.8: Government Transfers Received by Income Group and World Region, 2019

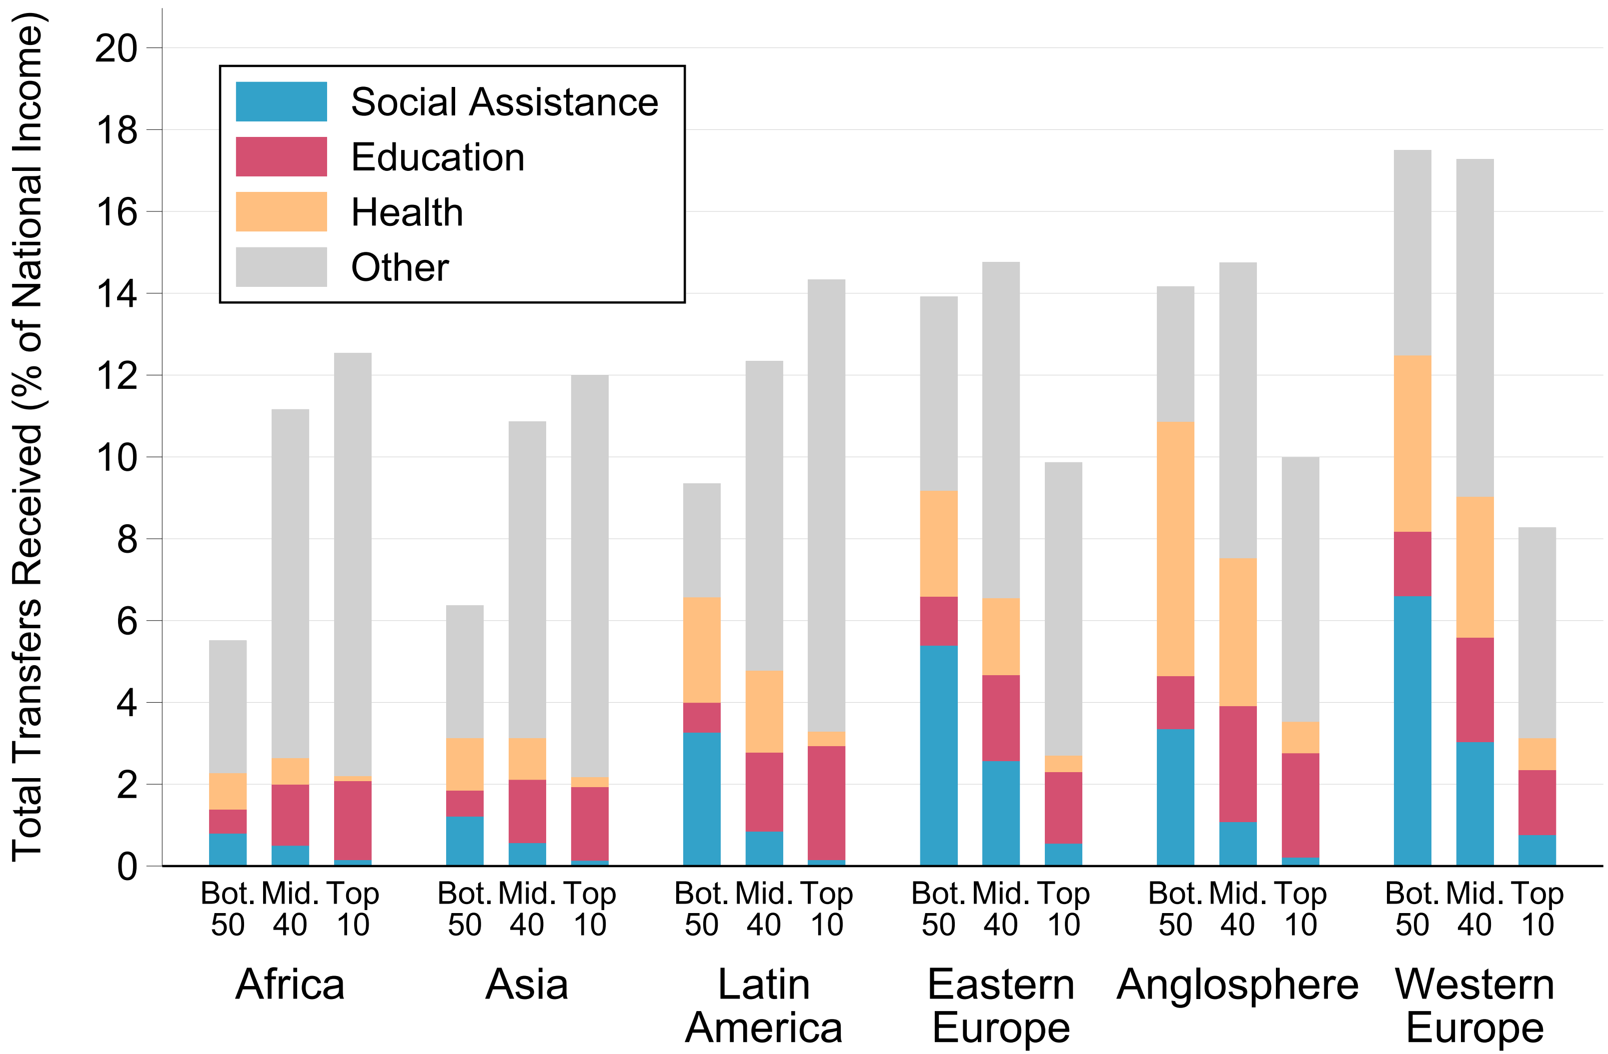

Notes. Population-weighted average of transfers received by income group in each country. Bot. 50: bottom 50% (p0p50); Mid. 40: middle 40% (p50p90); top 10: top 10% (p90p100).

Figure 3.9: A Global Map of Redistribution
Percent Reduction in Top 10% to Bottom 50% Income Ratio, Pretax - Posttax

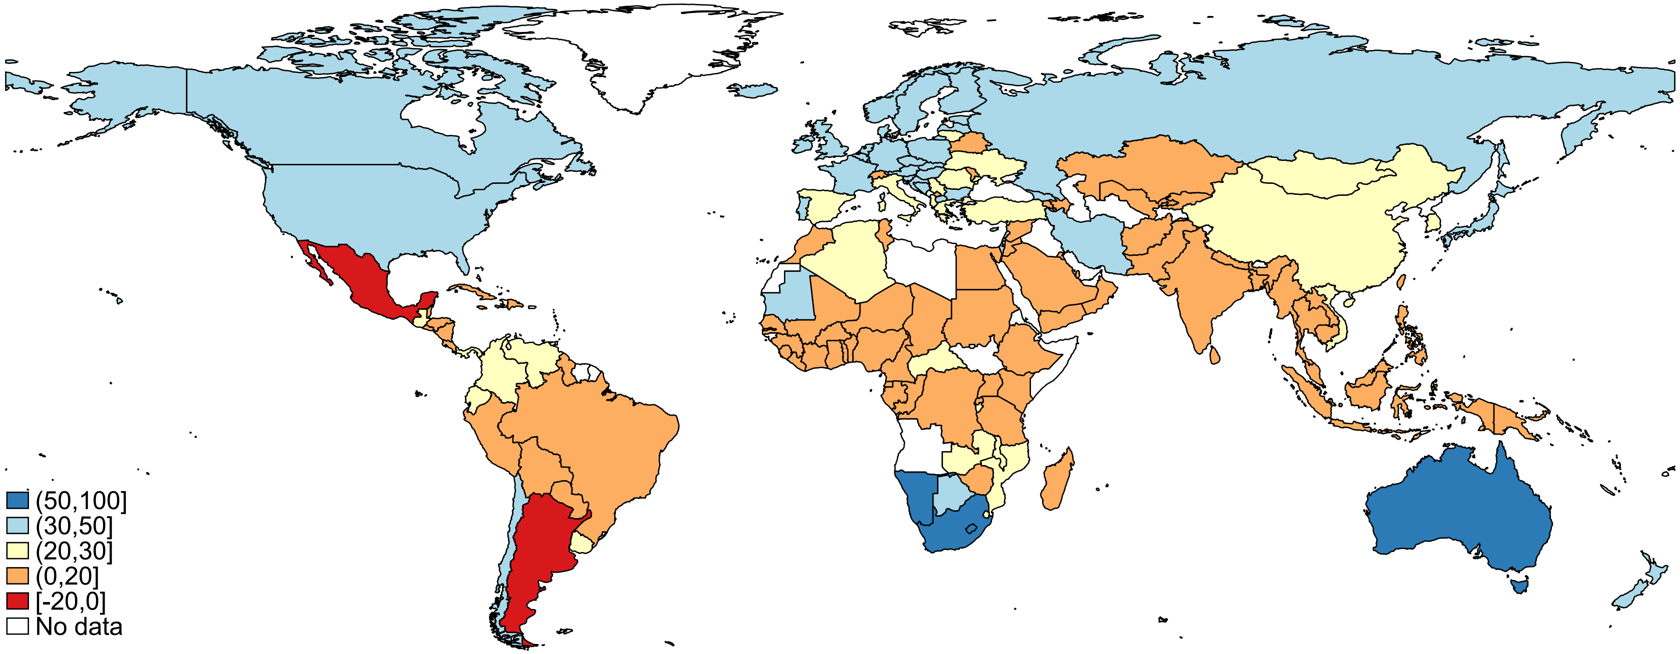

Notes. Posttax income: pretax income, minus all taxes, plus all transfers. Taxes exclude social contributions.

Figure 3.10: A Global Map of Redistribution: Net Transfers Operated by the Tax-and-Transfer System Between Pretax Income Groups, 2019

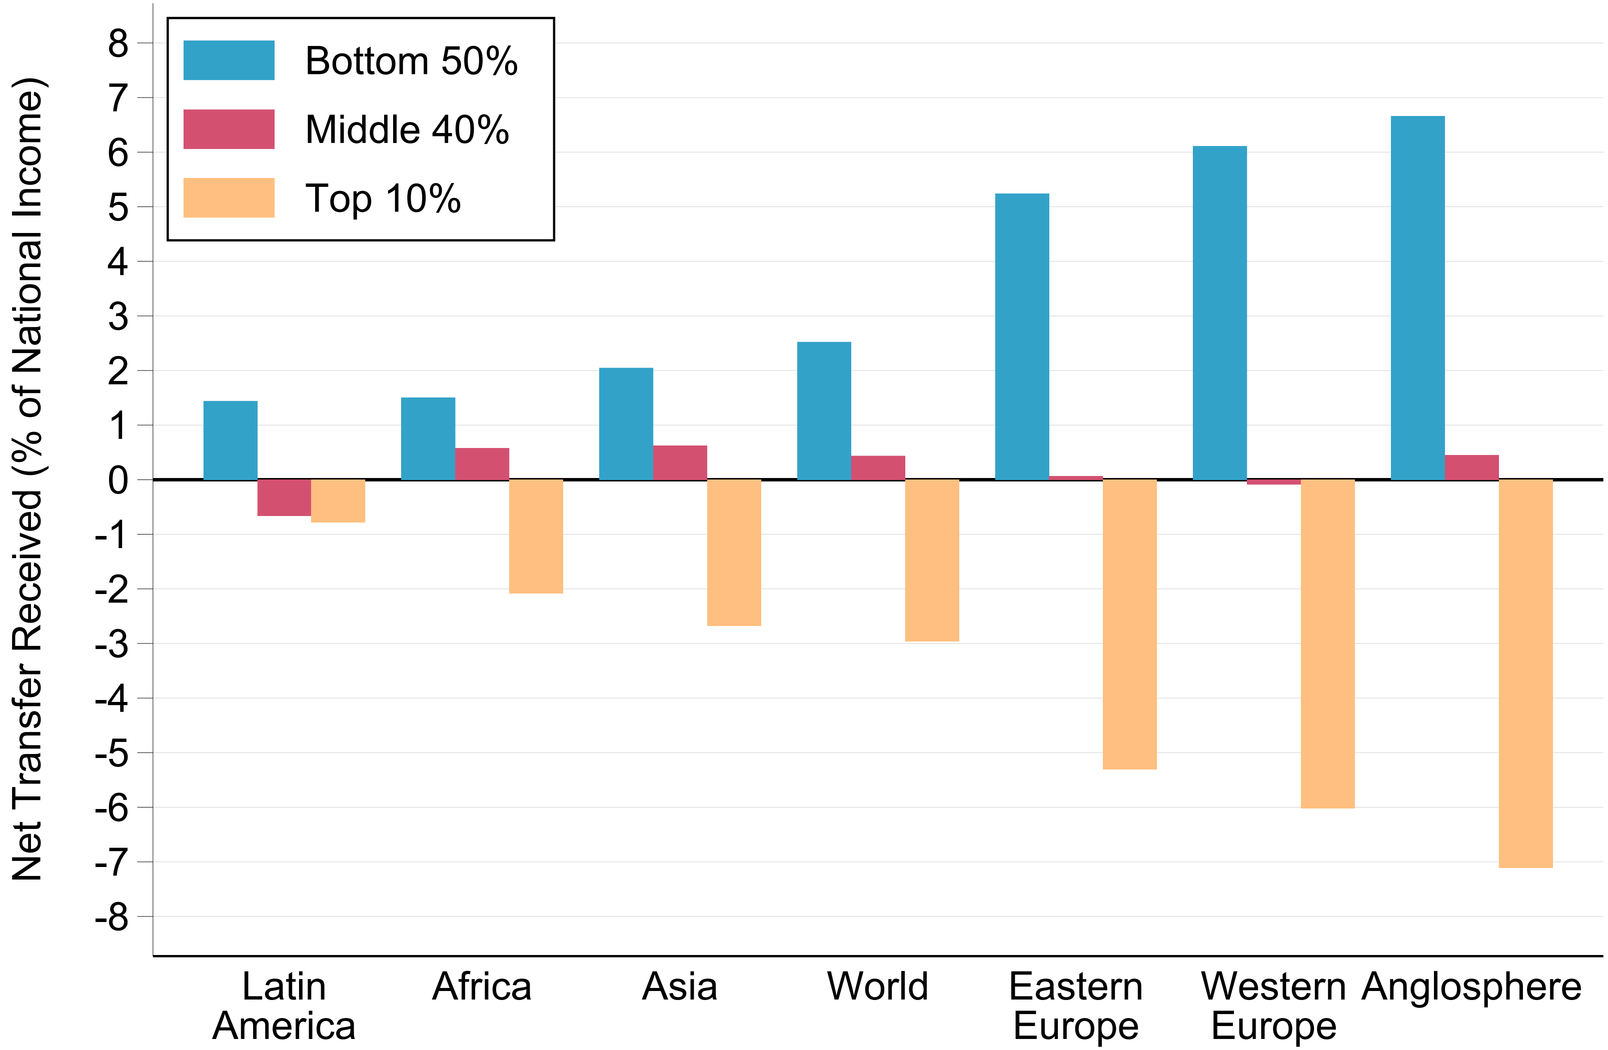

Notes. Net transfer: all transfers received minus all taxes paid, expressed as a share of national income. Taxes exclude social contributions. Population-weighted averages of net transfers received by income group in each country.

Figure 3.11: Extent of Redistribution by World Region, 1980-2019:
Percent Reduction in Top 10% to Bottom 50% Income Ratio, Pretax - Posttax

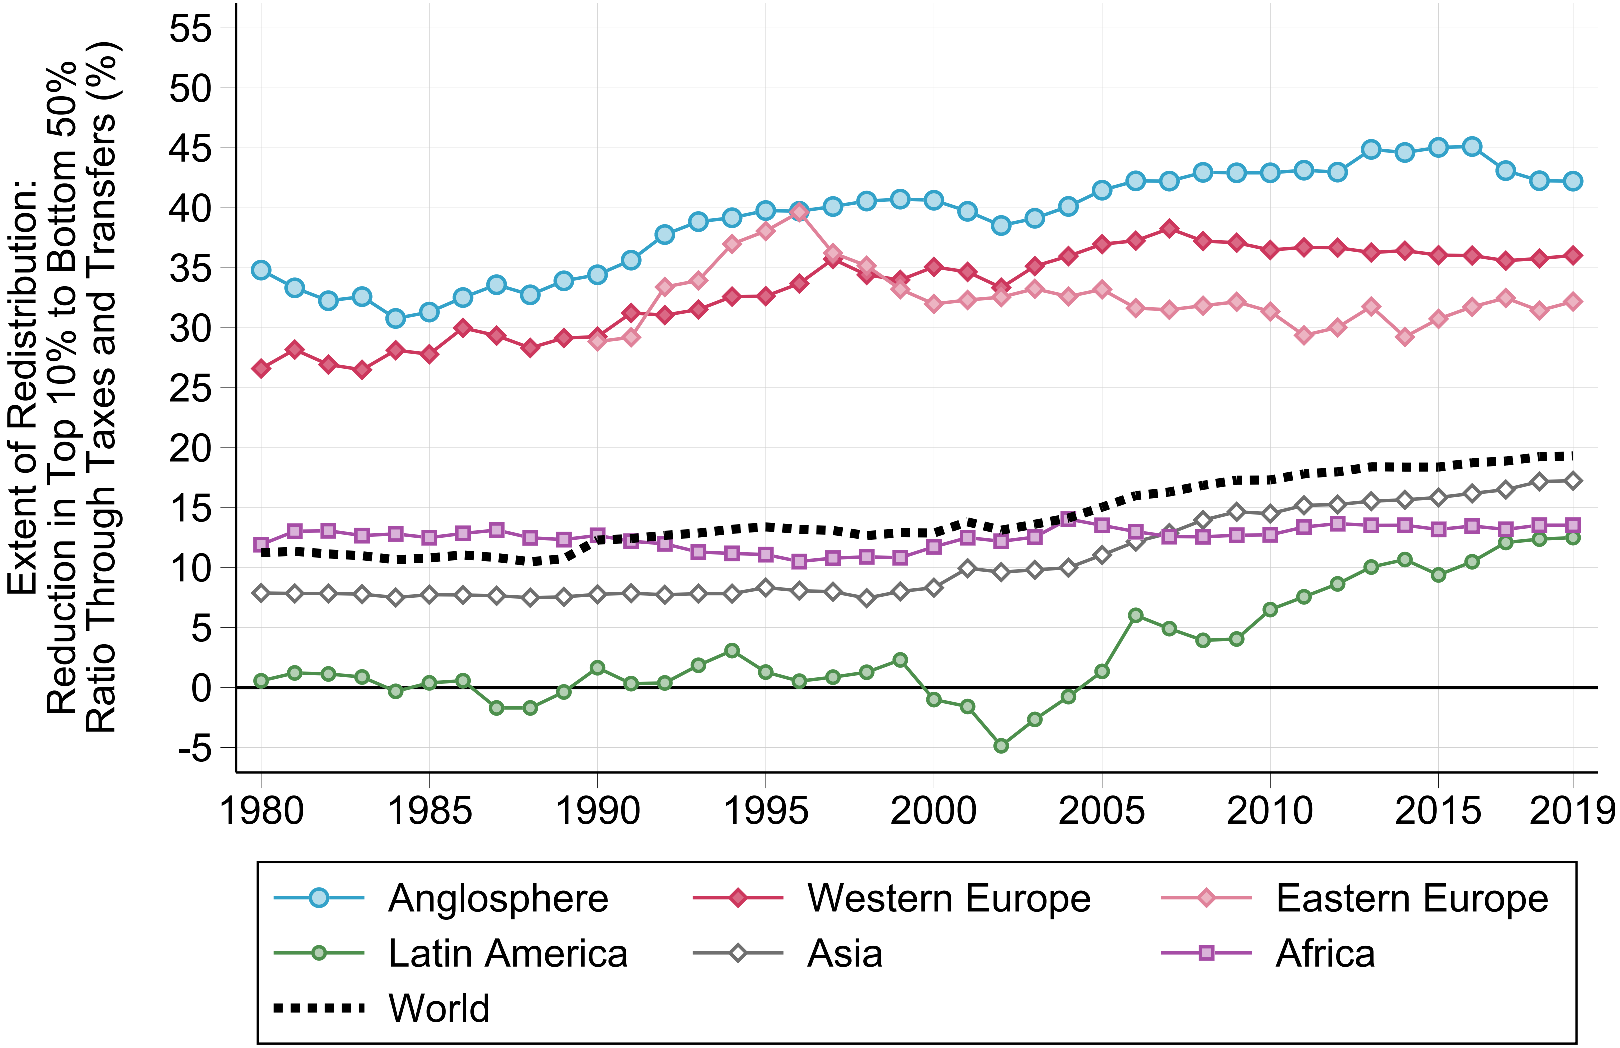

Notes. Population-weighted averages of the extent of redistribution in each country.

Figure 3.12: Extent of Redistribution by World Region, 1980-2019:
Net Transfer Received by the Bottom 50% (% of National Income)

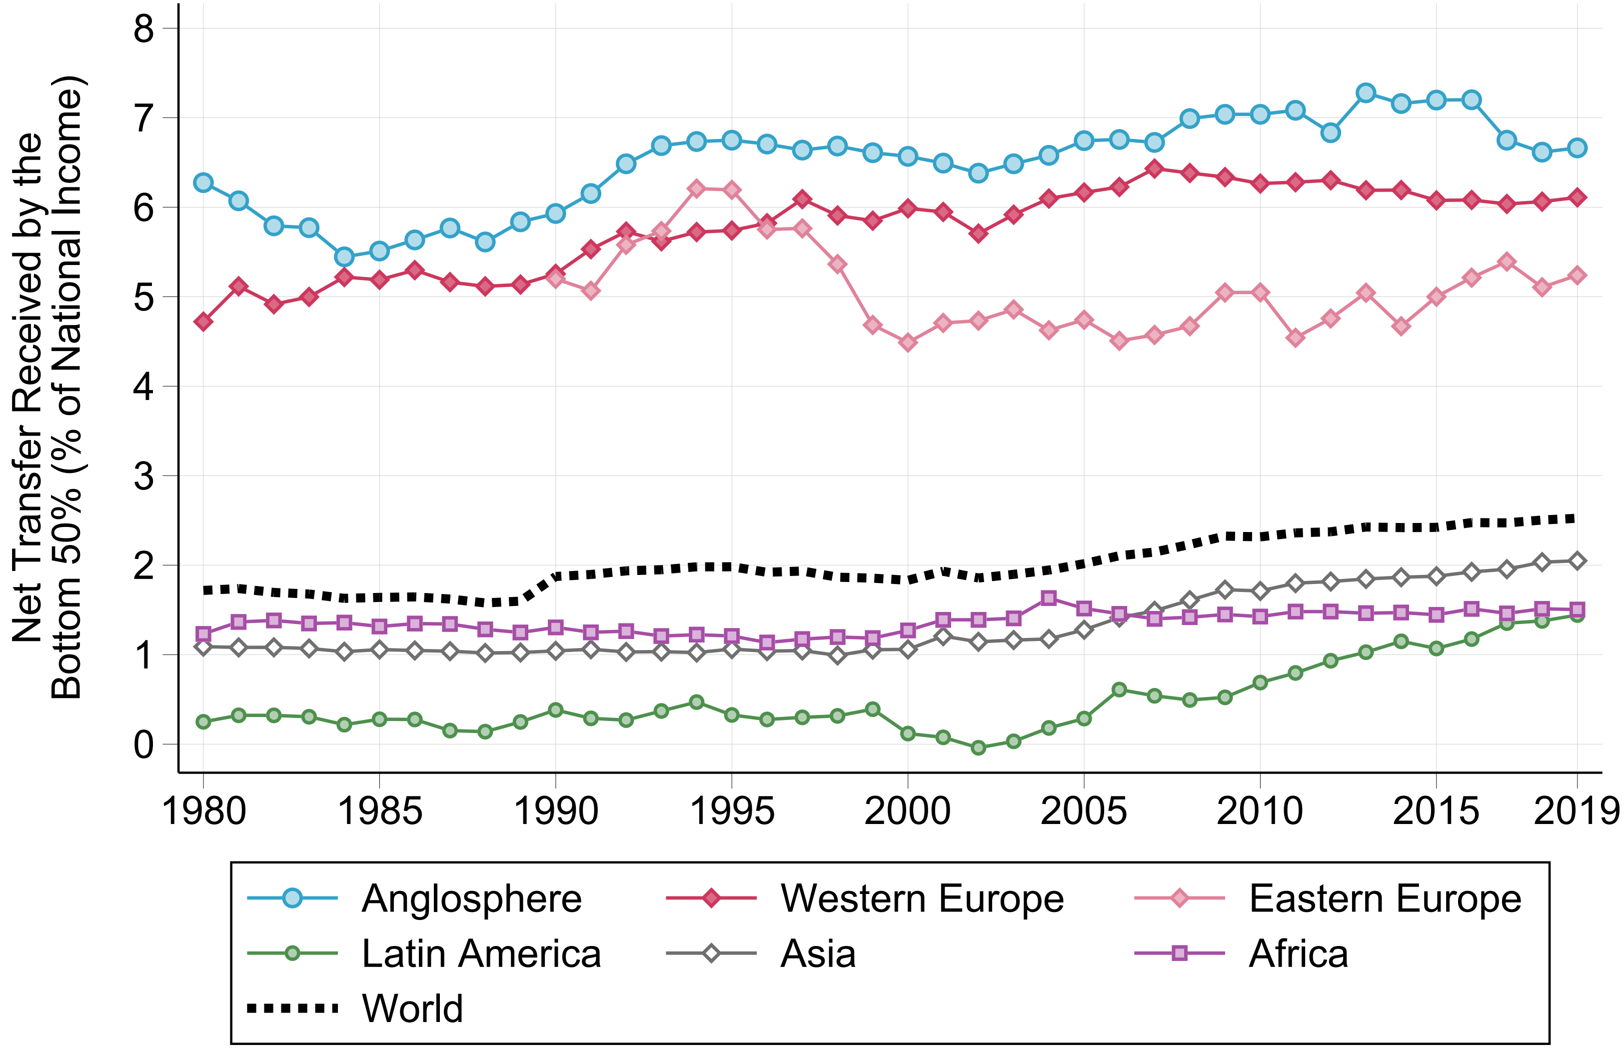

Notes. Net transfer: all transfers received minus all taxes paid, expressed as a share of national income. Population-weighted averages of net transfers received in each country.

Figure 3.13: Tax Progressivity Over the Course of Development:
Percent Reduction in Top 10% to Bottom 50% Average Income Ratio (Pretax versus Net-of-tax Income)

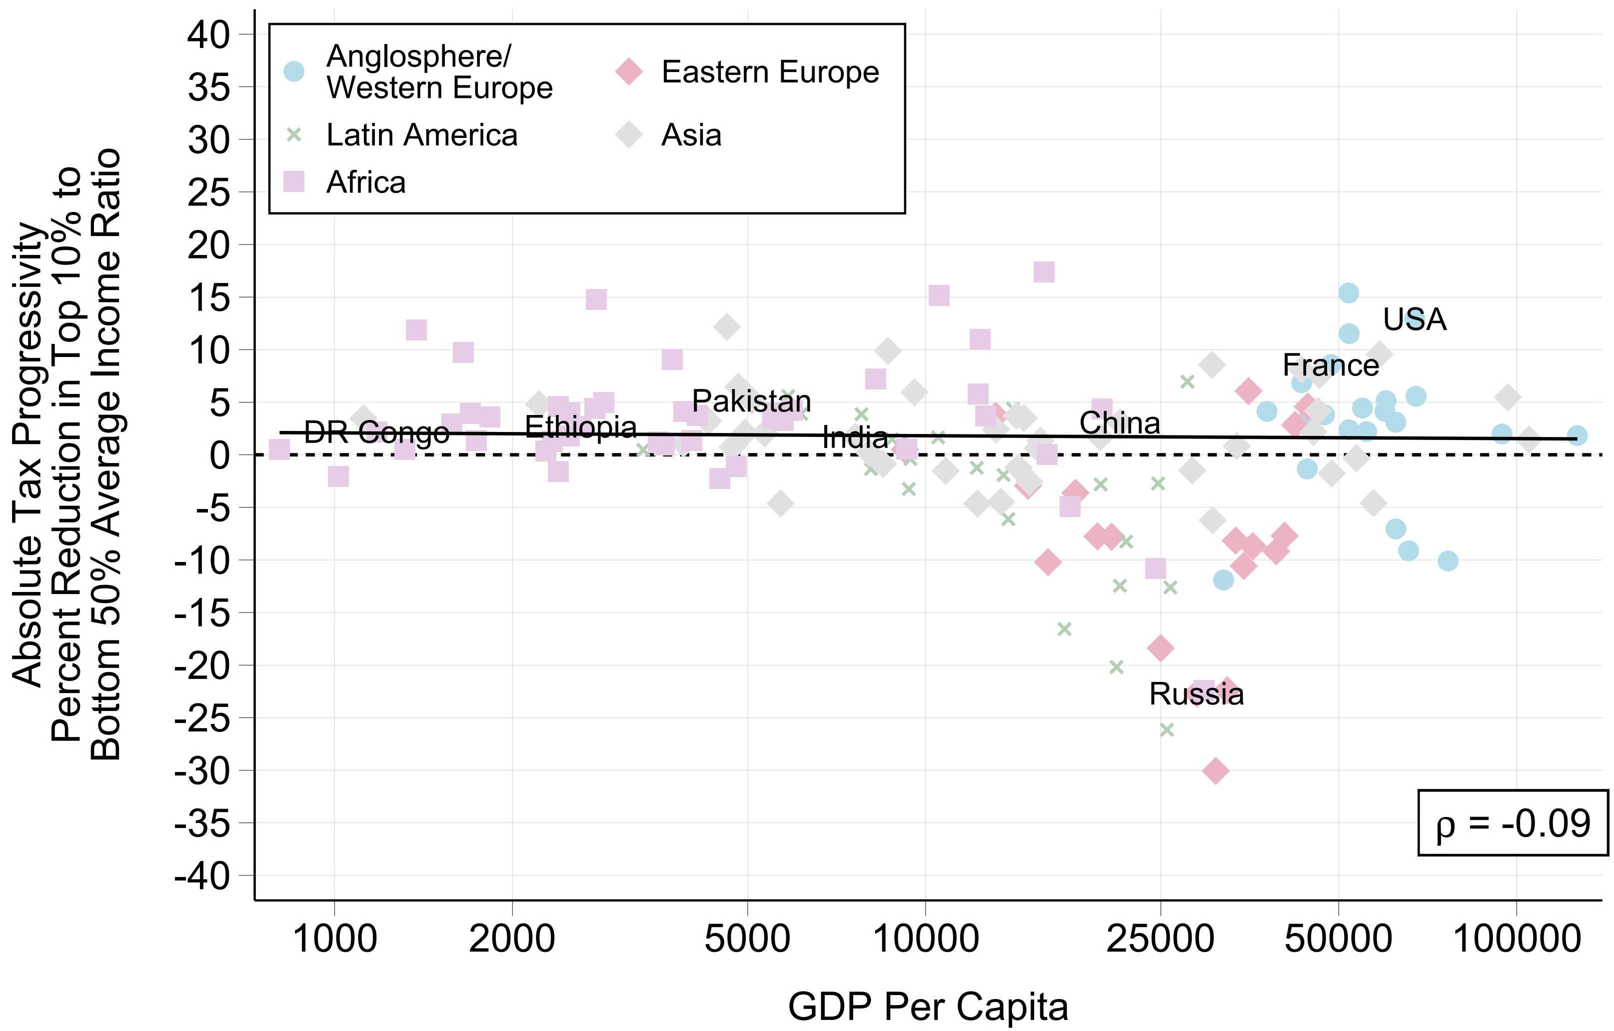

Notes. Net-of-tax income: pretax income minus taxes. Taxes exclude social contributions.

Figure 3.14: Transfer Progressivity Over the Course of Development:
Total Transfer Received by the Bottom 50% (% of National Income)

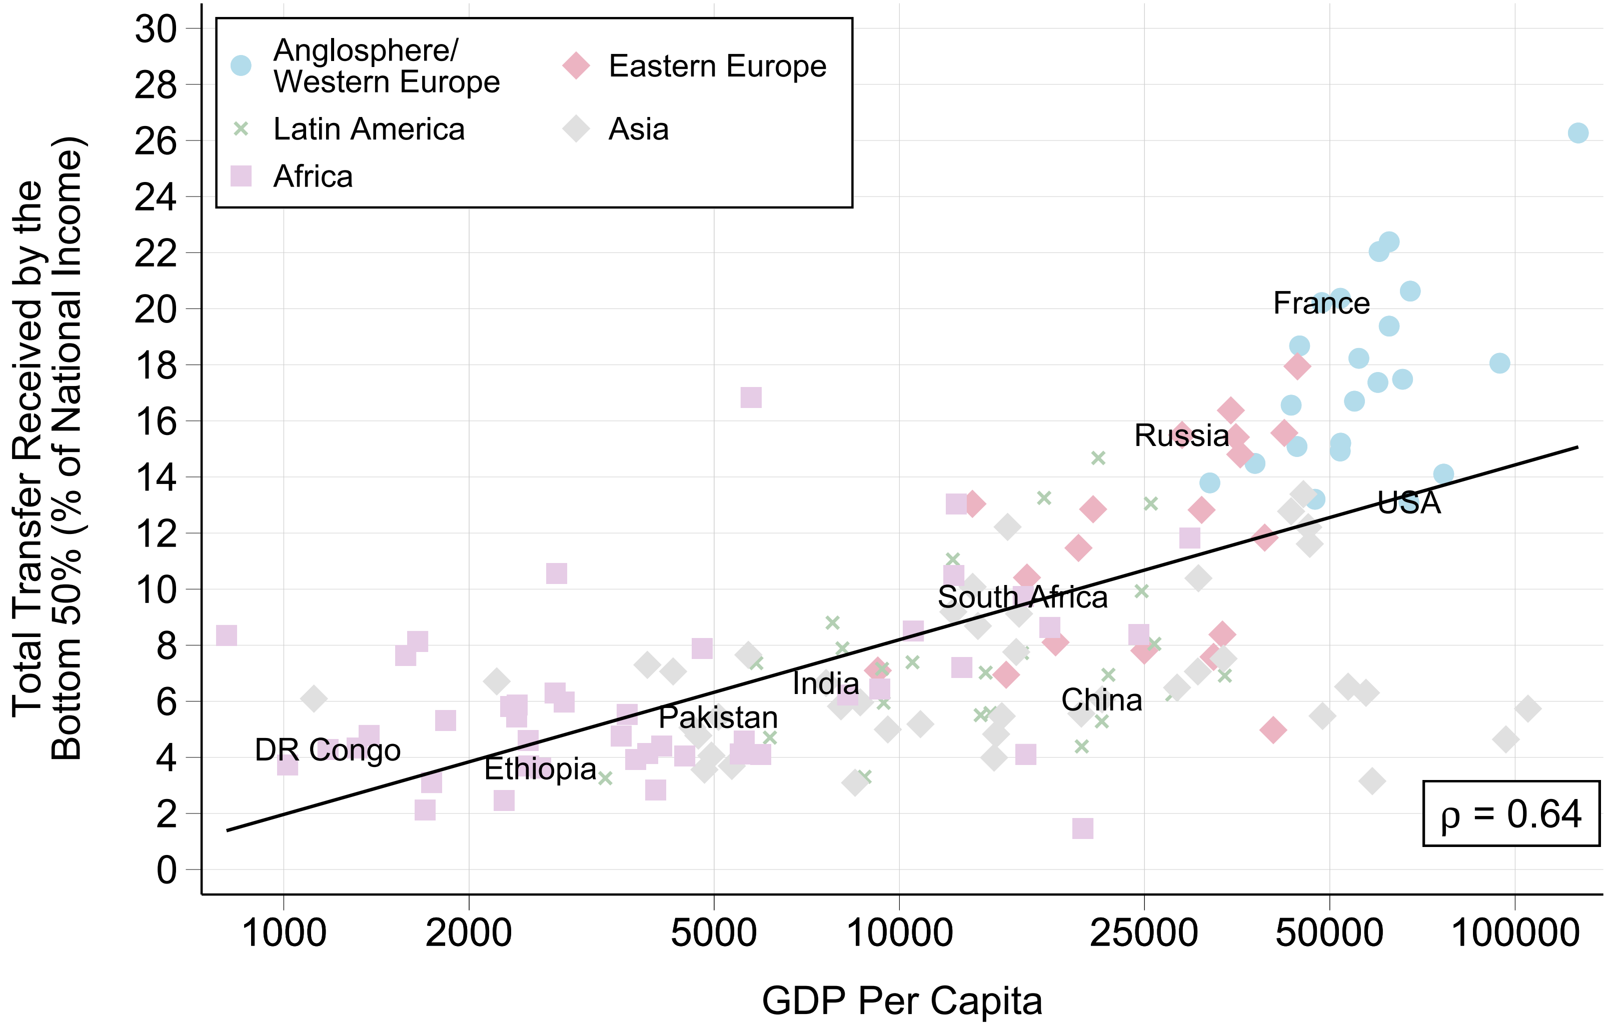

Notes. Total transfer received: sum of all transfers received (before paying taxes), expressed as a share of national income.

Figure 3.15: Net Redistribution Over the Course of Development:
Percent Reduction in Top 10% to Bottom 50% Income Ratio, Pretax - Posttax

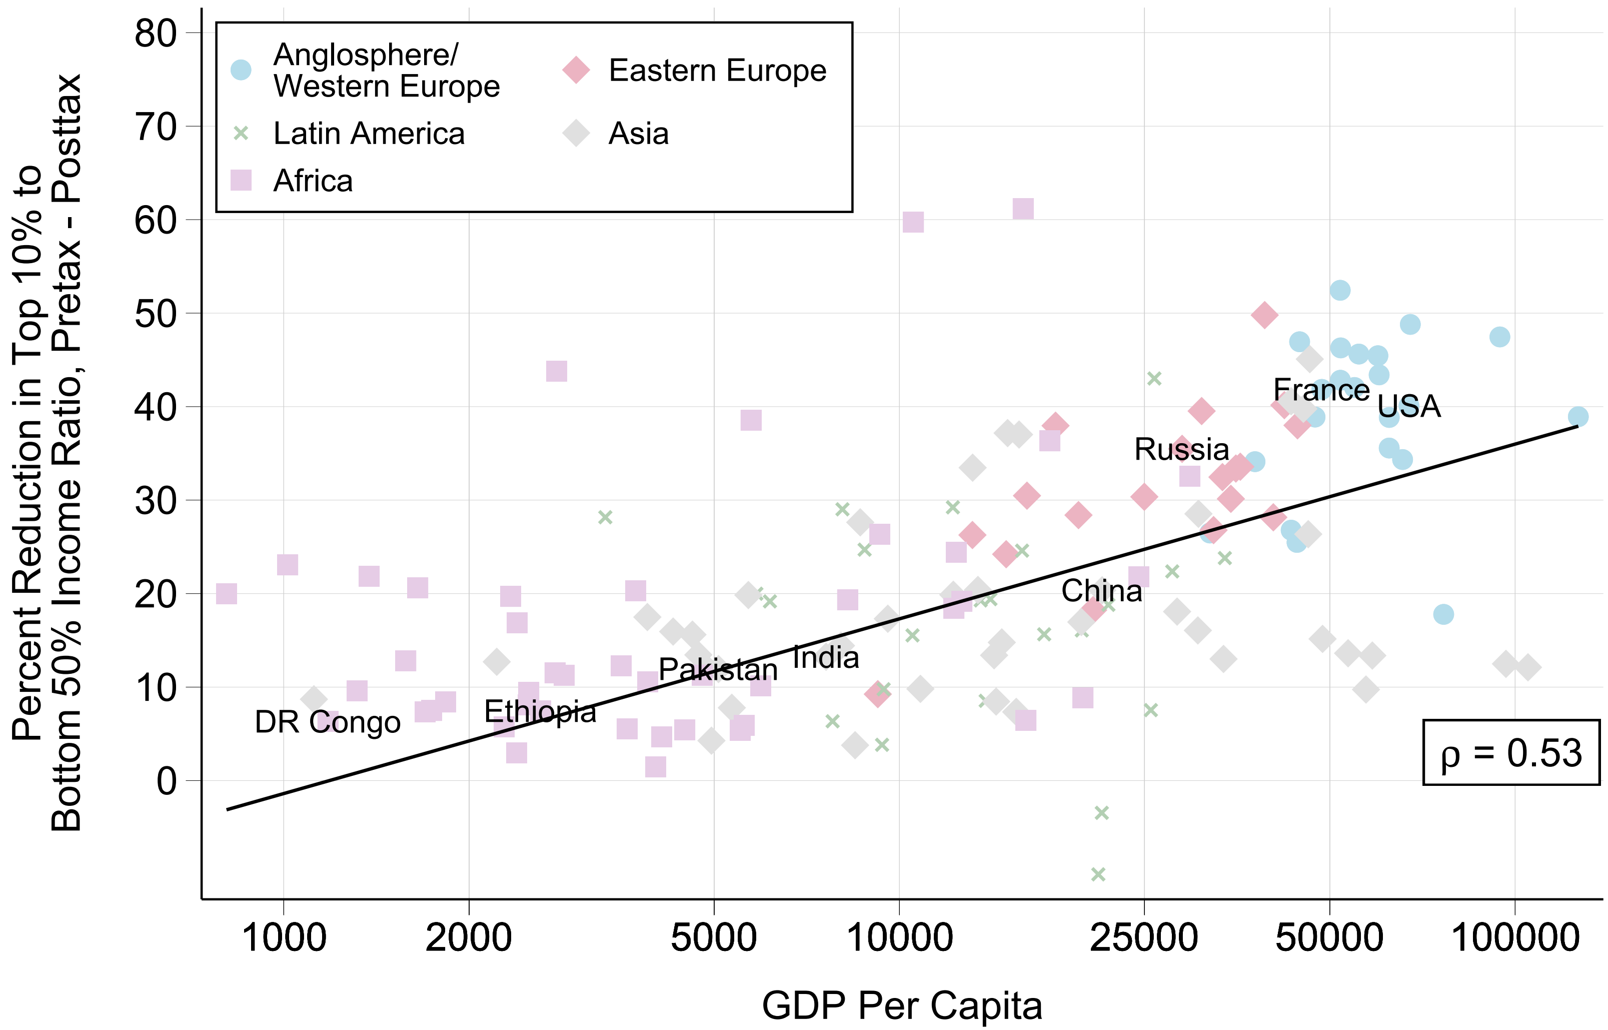

Figure 3.16: Predistribution versus Redistribution:
Bottom 50% Pretax versus Posttax National Income Shares by Country, 2019

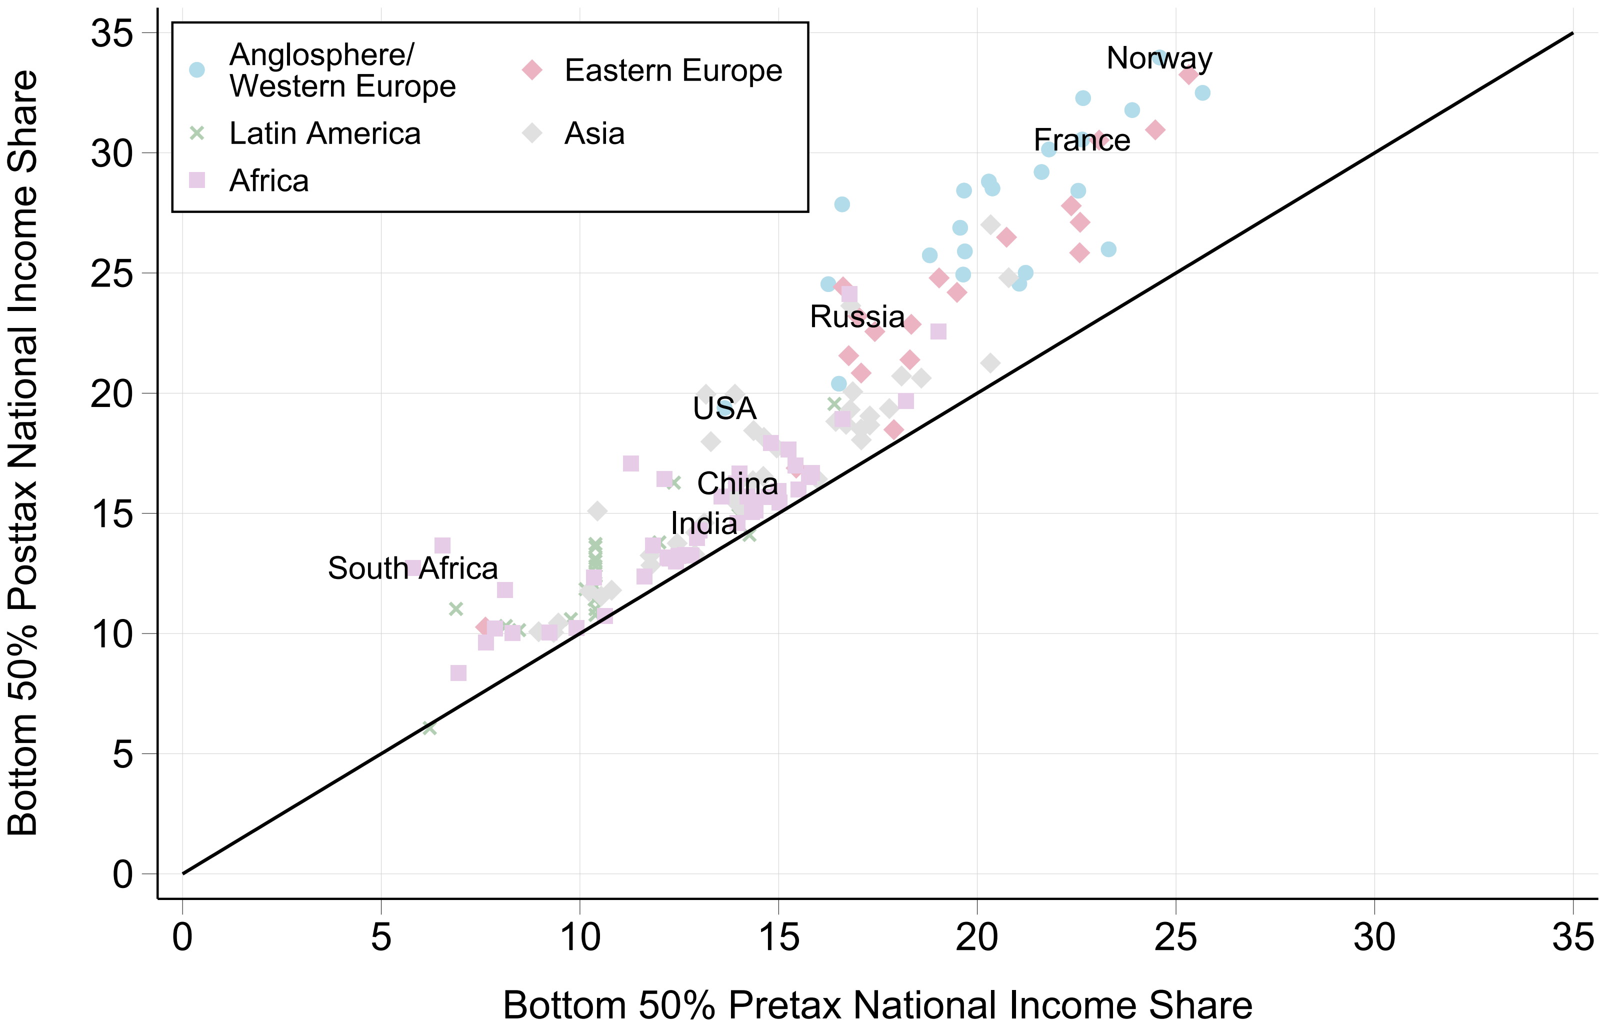

Figure 3.17: Predistribution versus Redistribution:
Bottom 50% Pretax Income Share versus Extent of Redistribution, 2019

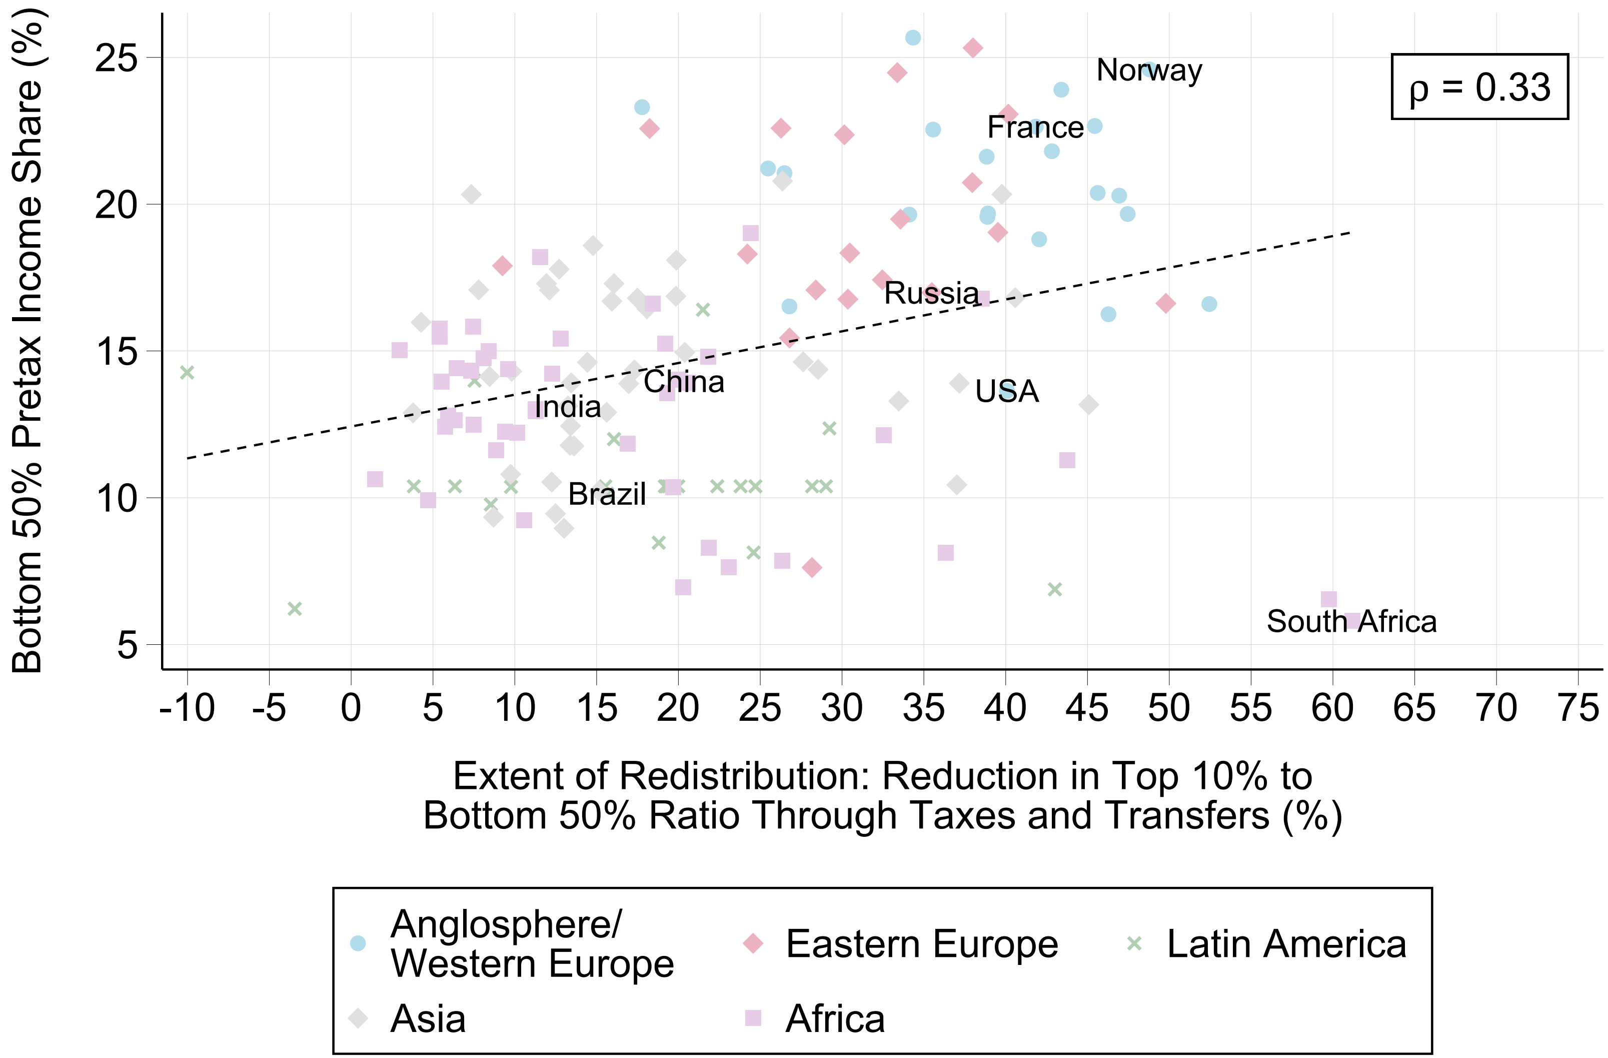

Figure 3.18: Predistribution versus Redistribution:
Bottom 50% Pretax Income Share versus Net Transfer Received by the Bottom 50%, 2019

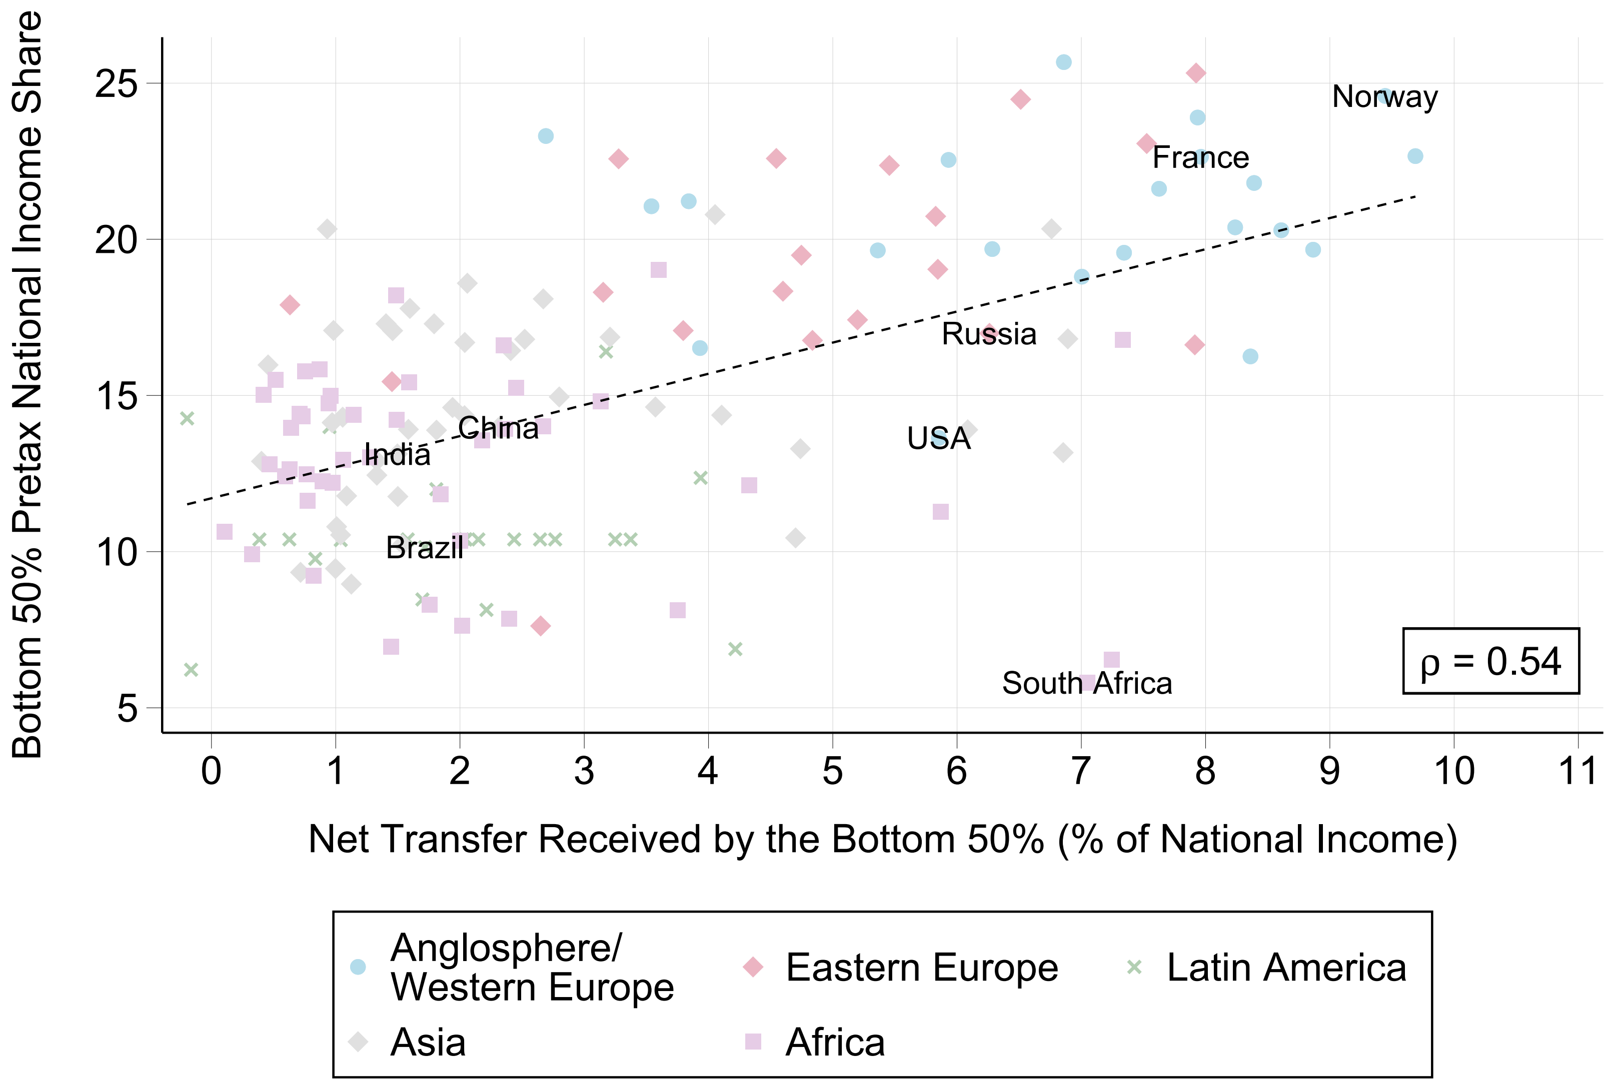

Table 3.1: Country and Time Coverage of Fiscal Incidence Estimates in Existing DINA Studies

Study	Countries	Years
Piketty, Saez, and Zucman (2018)	United States of America	1962-2019
Chatterjee, Czajka, and Gethin (2023)	South Africa	1993-2019
Bozio et al. (2022)	France	1990-2018
Fisher-Post, Herault, and Wilkins (2022)	Australia	1991-2018
Bruil et al. (2022)	Netherlands	2016
De Rosa, Flores, and Morgan (2022b)	Argentina, Brazil, Chile, Colombia, Costa Rica, Ecuador, El Salvador, Mexico, Peru, Uruguay	2000-2020*
Blanchet, Chancel, and Gethin (2022)	Austria, Belgium, Bulgaria, Croatia, Cyprus, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Iceland, Ireland, Italy, Latvia, Lithuania, Luxembourg, Malta, Netherlands, Norway, Poland, Portugal, Romania, Serbia, Slovakia, Slovenia, Spain, Sweden, Switzerland, United Kingdom	2007-2017*

Notes. * signifies unbalanced panel.

Table 3.2: Extent of Redistribution by World Region: the Dominant Role of Transfers

	Top 10% / Bottom 50% Average Income Ratio			Extent of Redistribution: Percent Reduction in Inequality		
	Pretax Income	After Taxes	After Taxes & Transfers	Through Taxes	Through Taxes & Transfers	Tax Share of Redistribution
Africa	20.0	18.9	16.3	4.2%	13.5%	30.9%
Anglosphere	14.8	13.0	8.6	11.6%	42.2%	27.4%
Asia	17.4	17.0	14.5	2.9%	17.3%	16.6%
Eastern Europe	11.2	13.0	7.6	-13.7%	32.2%	-42.6%
Latin America	31.6	35.0	28.1	-10.6%	12.5%	-84.4%
Western Europe	8.7	8.4	5.6	3.8%	36.0%	10.7%
World Average	18.2	18.0	14.9	1.8%	19.3%	9.3%

Notes. Population-weighted averages of indicators in each country. After taxes: top 10% to bottom 50% average income ratio in terms of net-of-tax income (pretax income minus all taxes). After taxes and transfers: top 10% to bottom 50% average income ratio in terms of posttax income (pretax income minus all taxes plus all transfers). Tax share of redistribution: ratio of extent of redistribution through taxes over extent of redistribution through taxes and transfers. Estimates for Eastern Europe, Western Europe, Latin America, the United Kingdom, and the United States come from existing DINA studies. All other series from this paper. Taxes exclude social contributions.

Table 3.3: Extent of Redistribution by World Region: Decomposition by Tax and Transfer, 2019

	World Average	Anglosphere	Western Europe	Eastern Europe	Latin America	Asia	Africa
Personal Income Taxes	4.4%	12.4%	14.0%	3.7%	4.6%	3.1%	3.2%
Corporate Taxes	4.2%	3.7%	3.7%	4.4%	4.0%	4.6%	3.3%
Property & Wealth Taxes	0.6%	0.8%	1.3%	0.6%	0.4%	0.6%	0.0%
Indirect Taxes	-7.7%	-7.3%	-14.7%	-23.4%	-10.2%	-6.9%	-3.3%
Social Contributions	-1.3%	-5.7%	-2.5%	-6.6%	-0.7%	-0.9%	0.2%
All Taxes	3.1%	12.1%	9.5%	-12.3%	0.9%	2.9%	4.2%
Social Assistance	10.4%	16.6%	22.9%	20.7%	23.5%	7.5%	5.5%
Healthcare	10.3%	28.4%	15.8%	11.2%	20.3%	7.5%	6.5%
All Transfers	18.3%	36.7%	33.4%	28.2%	34.7%	14.2%	10.9%

Notes. Population-weighted averages of indicators in each country. The table reports the negative of the percent change in the top 10% to bottom 50% income ratio before and after removing the corresponding tax or adding to corresponding transfer to pretax income. For instance, the top row reports the percent reduction in inequality resulting from removing personal income taxes from individual incomes. Positive values indicate that the corresponding tax or transfer reduces inequality. All series from this paper (existing DINA studies do not provide comparable, detailed decompositions by type of tax).

Chapter 4

Why Is Europe More Equal than the United States?

The evolution of inequality in Europe and the United States has attracted considerable attention in recent academic and policy debates, yet basic questions about the distribution of growth in the two regions remain unanswered. How did Europe and the US compare in terms of their distributional outcomes? What have been the respective roles of pretax income inequality and redistribution in explaining differences between the two regions? The comparative study of the distribution of growth, taxes, and transfers can provide critical insights into such debates. However, because of a lack of conceptual and empirical consistency, existing estimates of the income distribution have been hard to interpret and compare across countries. These shortcomings have led to a series of misunderstandings on the drivers of inequality in rich nations.

The standard source to compare economic growth across countries is the national accounts, while the standard source to measure inequality and redistribution is household surveys. Surveys are known to underrepresent top incomes and do not add up to macroeconomic income totals, leading to potential inconsistencies in the study of growth, inequality, and redistribution. In order to address some of these limitations, Piketty, Saez, and Zucman (2018) and Alvaredo et al. (2020) developed Distributional National Accounts (DINA) that combine various sources to distribute the entirety of a country's net national income, and established guidelines to carry out this work.

This new methodology has attracted significant interest, but unfortunately, with the exception of the United States (Piketty, Saez, and Zucman, 2018) and France

(Bozio et al., 2018; Garbinti, Goupille-Lebret, and Piketty, 2018), similar work in comparable countries remains rare.[1] In Europe, the absence of estimates of the income distribution that integrate survey, tax, and national accounts data are not the result of a lack of data *per se*. In fact, there is a fair amount of data available, at least since the 1980s. The problem is that these data are scattered across a variety of sources, taking several forms and using different concepts and methodologies. As a result, researchers and policymakers find themselves with a disparate set of indicators that are not always comparable, are hard to aggregate, provide uneven coverage, and can tell conflicting stories.

This article addresses these substantive and methodological issues by constructing distributional national accounts for twenty-six European countries from 1980 to 2017. To our knowledge, this is the first attempt at doing so. Our estimates combine virtually all existing data sources on the income distribution of European countries in a consistent way. These include household surveys, tax data, and national accounts, but also additional databases on social insurance benefits and contributions, and government spending on health that have been compiled by several institutions over the years (OECD, Eurostat, WHO). Our methodology exploits the strengths of each data source to correct for the weaknesses of others, making all assumptions explicit and as transparent as possible. It avoids the kinds of systematic errors and implausible assumptions that arise from the comparison of different income concepts, statistical units, or methodologies. Crucially, our series are fully comparable with recently produced US distributional national accounts, allowing us to compare the dynamics of inequality and redistribution in the two regions in great detail.

Two key findings emerge from the analysis of our new database.

First, we show that, over the past four decades, inequality has increased in nearly all European countries as well as in Europe as a whole, both before and after taxes, but much less than in the United States. Between 1980 and 2017, the share of pretax income that accrued to the richest 1% Europeans rose from 8% to 11% before taxes and transfers and from 7% to 9% after taxes and transfers. In the US, the top 1% pretax income share rose from 11% to 21% over the same period, and

[1]Statistical institutes, international organizations, and researchers have increasingly recognized the need to bridge the micro-macro gap in inequality studies. Since 2011, an expert group on the Distribution of National Accounts mandated by the OECD has been working on methods to allocate gross disposable household income to income quintiles (Fesseau and Mattonetti, 2013; Zwijnenburg, Bournot, and Giovannelli, 2019). In a similar fashion, experimental statistics on the distribution of personal income and wealth have been recently published by Eurostat (2018), Statistics Netherlands (2014), Statistics Canada (2019) and the Australian Bureau of Statistics (2019). These exercises have improved upon traditional survey-based estimates, but do not make systematic use of tax data and are restricted to the household sector.

the top 1% posttax income share from 9% to 16%. We also find that European inequality increased between 1980 and 1990, but less so afterwards, while the rise was sustained in the US. In Europe as a whole, inequality levels are mostly explained by within-country inequality, rather than by average income differences between Western, Northern, and Eastern European countries.[2] Between-country average income differentials are also found to explain close to none of the inequality trends observed in Europe in the past four decades. Still, regional dynamics vary: Eastern Europe has experienced the highest inequality increase, while the trend has been more muted in Western Europe. Northern Europe also experienced a significant increase in inequality but remains the most equal region, both before and after redistribution.

Second, the main reason for Europe's relative resistance to the rise of inequality has little to do with the direct impact of taxes and transfers. While Western and Northern European countries redistribute a larger fraction of output than the US (about 47% of national income is taxed and redistributed in Europe versus 35% in the US), the distribution of taxes and transfers does not explain the large gap between Europe and US posttax inequality levels. Quite the contrary: after accounting for all taxes and transfers, the US appears to redistribute a greater fraction of its national income to the poorest 50% than any European country. This finding stands in sharp contrast with the widespread view that "redistribution", not "predistribution", explains why Europe is less unequal than the US (e.g., OECD, 2008; OECD, 2011). In other words, Europe has been much more successful than the US at ensuring that its low-income groups benefit from relatively good-paying jobs. We show that the differences between our conclusions and those of the OECD are driven by several factors, including the greater underrepresentation of top incomes in US surveys, the fact that we account for indirect taxes and in-kind transfers, which are more progressive in the US than in Europe overall, and the inclusion of pensions in the definition of pretax income.

This paper contributes to the existing literature on the evolution of income inequality in Europe and the US in three ways.

First, we provide novel estimates on the distribution of growth in Europe as a whole and within European countries. While it has generally been acknowledged that income inequality has grown in Europe since the 1980s (OECD, 2008), little is known of how this rise compares across countries, across income groups in the distribution,

[2] In 2017, more than 80% of European inequality is due to within-country differences according to a Theil index decomposition. See Figure 4.4.

or across time periods. The efforts made by the Luxembourg Income Study (LIS) to harmonize existing surveys, for instance, have been extremely helpful to improve the comparability of pre-2000 inequality statistics in Europe. Yet, because of sampling issues and misreporting at the top of the income distribution, surveys can picture evolutions that are inconsistent with those suggested by tax data. In this paper, we combine for the first time all these sources in a meaningful way, using new techniques and a consistent methodology. We show that correcting for the weaknesses of existing estimates does lead to substantively different conclusions on the level and evolution of inequality in Europe, the distributive impact of taxes and transfers, and how inequality and redistribution compare across European countries.

Second, we compare how growth has been distributed before and after taxes in Europe and the United States since 1980. While most studies suggest that posttax income inequality is greater in the US than in European countries today, it remains unclear whether this gap is due to differences in pretax income inequality or to differences in government redistribution. International organizations such as the OECD (OECD, 2008; OECD, 2011), in line with other research (e.g., Jesuit and Mahler, 2010; Immervoll and Richardson, 2011), find that the lower posttax income inequality levels of European countries are mostly due to redistribution. This contrasts with Bozio et al. (2018), who use the DINA methodology, distribute all taxes and transfers and find that redistribution reduces inequality less in France than in the US.[3] Whether the US is more unequal than Europe as a whole (i.e., as a region) also remains an open question.[4] Thanks to our new dataset, we are able to provide new insights into these questions, decomposing precisely the contributions of spatial integration, pretax income inequality, and redistribution in explaining differences between Europe and the US and their evolution over time.

Third, we contribute to the distributional national accounts literature by enriching its methodology. Piketty, Saez, and Zucman (2018) and Garbinti, Goupille-Lebret, and Piketty (2018) start with tax data, to which they progressively add information from surveys and national accounts. This “top-down” approach exploits all the richness of tax microdata and yields very detailed and precise estimates. However, while this type of work should be extended to as many European countries as possible, there are many countries and time periods for which tax microdata are simply not

[3]See also Guillaud, Olckers, and Zemmour (2019), who find results similar to Bozio et al. (2018) without using the DINA framework.

[4]Works on the distribution of income in the EU-15 (Atkinson, 1996) or the Eurozone (Beblo and Knaus, 2001) suggested that income inequality was higher in the US, but recent studies extending the analysis to new, poorer Eastern European member states have found mixed results (e.g. Brandolini, 2006; Dauderstädt and Keltek, 2011; Salverda, 2017; Filauro and Parolin, 2018).

available. This justifies our "bottom-up" approach, which starts from surveys and gradually incorporates information from top incomes shares, estimated from income tax tabulations, and unreported national income components. As such, we view our methodology as well-suited to estimating the distribution of national income in countries gathering a mix of survey microdata, tabulated tax returns, and a variety of other heterogeneous data sources. This case corresponds to the majority of countries beyond Europe and the US.[5]

The rest of this paper is organized as follows. Section 4.1 presents our conceptual framework, data sources, and methodology. Section 4.2 summarizes our findings on the distribution of pretax incomes in Europe. Section 4.3 discusses the impact of taxes and transfers on inequality in Europe and the US. Section 4.4 concludes.

4.1 Data Sources and Methodology

This section introduces the data sources and methodology used to estimate the distribution of national incomes in Europe. Section 4.1.1 outlines our conceptual framework and the assumptions used to distribute the components of net national income. Section 4.1.2 presents the data sources used. Section 4.1.3 explains how we harmonize and combine these data sources to derive estimates of factor income, pretax income, and posttax income inequality.

4.1.1 Conceptual Framework

Universe We study the distribution of national income in twenty-six European countries from 1980 to 2017. The choice of countries considered in this paper has been dictated by the availability of comparable, high-quality data sources allowing us to estimate pretax and posttax inequality statistics with a sufficient degree of certainty.[6] Our geographical area of interest includes all fifteen members of the European Union before its 2004 extension (Austria, Belgium, Denmark, Finland,

[5] In a similar fashion, Piketty, Saez, and Zucman (2019) have recently proposed a simplified method for recovering estimates of top pretax national income shares based on the fiscal income shares of Piketty and Saez (2003) and very basic assumptions on the distribution of untaxed labor and capital income components. Our methodology follows the same spirit.

[6] More precisely, we exclude from our sample all European countries for which no tax data was available to correct incomes at the top end of the distribution: Albania, Bosnia and Herzegovina, Bulgaria, Cyprus, Kosovo, Lithuania, Latvia, Malta, Moldova, Montenegro, North Macedonia, and Slovakia. However, we still provide results for each of these countries in the extended online appendix, estimated with and without an imputed top income correction profile. Including or excluding these countries from the analysis barely affects our estimates of European income inequality (see appendix figures D.20 and D.21).

France, Germany, Greece, Ireland, Italy, Luxembourg, the Netherlands, Portugal, Spain, Sweden, and the United Kingdom), seven Central and Eastern European countries that joined the EU in 2004 or in the years that followed (Croatia, the Czech Republic, Estonia, Hungary, Poland, Romania, and Slovenia), and four countries that are not part of the EU but have maintained tight relationships with it (Iceland, Norway, Serbia, and Switzerland).

Methodological Framework We follow the principles of the DINA guidelines (Alvaredo et al., 2020), which provide a set of methods to distribute the totality of net national income—GDP minus capital depreciation plus net foreign income—in a way that is consistent with the concepts defined in the System of National Accounts (United Nations Statistics Commission, 2008). The DINA framework acknowledges three levels of distribution: factor national income, pretax national income, and posttax national income. We report in table 4.1 how these concepts are derived, which data sources are used to allocate their various components, the distributional assumptions made in this paper to do so, and the share of national income they represent.

Factor Income Factor national income corresponds to all income flows that accrue to individuals before any form of government redistribution.[7] It is equal to the sum of the primary incomes of the different sectors of the economy: households, corporations, and the government.

The primary income of households (79% of national income on average across European countries) can be decomposed into four main components: compensation of employees, mixed income, net property income, and the net imputed rents of owner-occupiers. The distribution of these income flows is generally observed in survey and tax data, although data on imputed rents are not systematically collected and are usually not included in inequality measures published by statistical institutes or international organizations.

The primary income of corporations (8%) corresponds to the income that companies retain after having paid suppliers, employees, shareholders, and corporate taxes, and that we also refer to as "retained earnings" or "undistributed profits". Following other DINA studies, we consider that the undistributed profits of privately owned

[7]We refer in this paper to "redistribution" as the operation of the tax-and-transfer system, measured by the difference between pretax and posttax income inequality. By contrast, "predistribution" refers to all forms of government interventions (such as labor market regulations, minimum wages, educational investments, etc.) that drive pretax inequality levels (see section 4.3 as well as Hacker and Pierson (2010) for a discussion of these concepts).

corporations belong to the owners of these corporations. We separate the share of retained earnings that accrues to shareholder households, to the government, and to pension funds, proportionally to the total amount of equity they own. We distribute the retained earnings of shareholder households proportionally to their equity ownership. We distribute the retained earnings that accrue to pension funds proportionally to wage and pension income. And we distribute the government's share like the primary income of the government.[8,9]

The primary income of the general government (12%) is the sum of taxes less subsidies on production and imports and of net property income. In our benchmark series, we distribute it proportionally to pretax income, in line with DINA recommendations (Alvaredo et al., 2020).[10]

Pretax Income Pretax income corresponds to income after the operation of social insurance systems, but before other types of redistribution. It is equal to factor income, plus pension benefits (17% of national income on average) and unemployment and disability benefits (1.7%), minus the social contributions that pay for them. Contributions and transfers are generally observed in survey data and can therefore be directly removed from or added to individual factor incomes.

Notice that for pretax income to sum up to national income, it is important to remove the same amount of social contributions as the amount of social benefits that we distribute. In most countries, social contributions exceed pension and unemployment benefits, because contributions also pay for health or family-related benefits that we classify as non-insurance-based redistribution. In these cases, we only deduct a fraction of social contributions from pretax income (their "contributory" part). On the contrary, in a few countries such as Denmark, social contributions are virtually non-existent. In these cases, we assume that social insurance is financed by the income tax by deducting a fraction of the income tax from factor income to get to

[8]This can be justified by the fact that retained earnings correspond to profits that are kept within the company rather than distributed to shareholders as dividends. This income ultimately increases the wealth of shareholders and therefore represents a source of income to them. Several papers have documented the impact of including retained earnings in the United States (Piketty, Saez, and Zucman, 2018), Canada (Wolfson, Veall, and Brooks, 2016), and Chile (Atria et al., 2018; Fairfield and Jorratt De Luis, 2016). In Norway, Alstadsæter et al. (2017) showed that the choice to keep profits within a company or to distribute them is highly dependent on tax incentives, and therefore that failing to include them in estimates of inequality makes top income shares and their composition artificially volatile.

[9]This approach assumes that the wealthiest shareholders do not own stock in companies that systematically have higher retained earnings than the rest.

[10]We provide variants in which taxes on products are distributed proportionally to consumption in the appendix (see appendix figure D.22).

pretax income.

Posttax Income Posttax income accounts for other forms of redistribution operated by the government. We consider two types of posttax income concepts. Posttax disposable income removes all taxes from pretax income but only adds back cash transfers and therefore does not sum up to national income. Posttax national income also adds back collective government expenditure and therefore adds up to national income.

To move from pretax to posttax income, we first remove all taxes and social contributions that remain to be paid by individuals. These include non-contributory social contributions (1% of national income) and direct taxes on income and wealth (11%), which are directly observed in survey and tax data. They also include indirect taxes (14%) and corporate income taxes (3%), which are not directly observed. We assume that indirect taxes are paid by consumers and distribute them proportionally to household final consumption expenditure. Corporate income taxes are paid out of corporate profits, so we distribute them similarly to undistributed profits.

We then allocate all types of government transfers to individuals. Social assistance transfers (5% of national income) are observed in survey data, so they can be added directly to individual incomes. We distribute other public spending proportionally to posttax disposable income (17% of national income), with the exception of public health expenditure (8%), which we distribute in a lump-sum way, considering that the insurance value provided by health systems is similar for everyone. While this remains a simplification, the existing literature suggests that it does represent a good first-order approximation of who benefits from the public healthcare system.[11] We use the proportionality assumption for non-health in-kind transfers as a benchmark for simplicity, transparency, and comparability with US distributional national accounts (Piketty, Saez, and Zucman, 2018), but we discuss at greater length the robustness of our findings to this assumption in section 4.3. In particular, we consider an alternative scenario in which all collective expenditure is distributed on a lump-sum basis and find that it does not alter our main conclusions.[12]

We distribute the budget balance of the government (the discrepancy between what it collects in taxes and what it pays in transfers, representing -0.7% of national

[11]See in particular Germain et al. (2020), who combine household surveys and administrative data with a simulation model of health payments to distribute health expenditure in France.

[12]This assumption affects the levels of posttax inequality but is unlikely to affect the trends, as government final expenditures have remained fairly stable in Europe with no major changes in their decomposition by functions: see figure D.8 in appendix.

income on average) proportionally to posttax disposable income.

Unit of Analysis In our benchmark series, the statistical unit is the adult individual (defined as being 20 or older) and income is split equally among spouses.[13]

4.1.2 Data sources

National Accounts For total net national income, we use series compiled by the World Inequality Database based on data from national statistical institutes, macroeconomic tables from the United Nations System of National Accounts, and other historical sources (see Blanchet and Chancel, 2016). For the various components of national income, we collect national accounts data from Eurostat, the OECD, and the UN. Additional data comes from the OECD health and social expenditures databases. We provide a detailed view of the coverage that these data provide in the appendix.

Survey Microdata We collect and harmonize household survey microdata from several international and country-specific datasets. Our most important source of survey data is the European Union Statistics on Income and Living Conditions (EU-SILC), which has been conducted on a yearly basis since 2004 in thirty-two countries. We complement EU-SILC by its predecessor, the European Community Household Panel (ECHP), which covers the 1994-2001 period for thirteen countries in Western Europe. Our second most important source of survey data is the Luxembourg Income Study (LIS), which provides access to harmonized survey microdata covering twenty-six countries since the 1970s. Most Western European countries are covered from 1985 until today, and several countries from Eastern Europe have been surveyed since the 1990s.

Survey Tabulations We complement survey microdata sources with a number of tabulations available from the World Bank's PovcalNet portal, the World Income Inequality Database (WIID), and other sources. PovcalNet provides pre-calculated survey distributions by percentile of posttax income or consumption per capita. The WIID gathers inequality estimates obtained from various studies, and gives information on the share of income received by each decile or quintile of the population. Finally, we collect historical survey data on posttax income inequality in former

[13] We also compute additional series in which income is split between all adult household members, not just members of a couple (i.e., a "broad" rather than a "narrow" equal-split)—see appendix figure D.23.

communist Eastern European countries provided by Milanovic (1998). In all cases, we use generalized Pareto interpolation (Blanchet, Fournier, and Piketty, 2021) to recover complete distributions from the tabulations.

Tax Data To better capture the evolution of incomes at the top end of the distribution, we rely on known top income shares estimated from administrative data and compiled in the World Inequality Database. In general, tax data is only reliable for the top of the distribution, and this is why these series do not cover anything below the top 10%. At the time of writing, data series were available for nineteen European countries. We complete this database by gathering and harmonizing a new collection of tabulated tax returns covering Austria (1980–2015), East Germany (1970–1988), Estonia (2002–2017), Iceland (1990–2016), Italy (2009–2016), Luxembourg (2010, 2012), Portugal (2005–2016), Romania (2013), and Serbia (2017). We use these tabulations to add new top income shares to our database (see appendix section D.1.7).

4.1.3 Methodology

We now explain how we combine these various data sources to estimate the distributions of factor income, pretax income, and posttax income in Europe. First, we derive measures of household income inequality from survey microdata. Second, we train a machine learning algorithm to correct conceptual inconsistencies in survey tabulations. Third, we combine survey data with tax data to correct incomes at the top end of the distribution. Fourth, we combine external data sources with national accounts aggregates to distribute unreported national income components. We summarize the different steps of this methodology in table 4.2. Similar tables for each of the countries covered in this paper are available in the appendix.

1) Direct Measurement of Income Concepts in Survey Microdata When we have access to survey microdata (from EU-SILC or LIS), we can in most cases estimate income concepts that are close to our concepts of interest. As a result, we have survey data on both pretax and posttax income inequality for almost all countries since 2007, and for a longer period of time for a number of countries.

A significant exception concerns employee and employer social contributions in EU-SILC, which are not always reported separately from income and wealth taxes. We use the social contribution schedules published in the OECD Tax Database to impute social contributions separately. This only has a marginal effect on estimates of pretax income inequality.

2) Harmonization of Survey Tabulations Contrary to microdata, tabulations only provide distributions covering specific income concepts and equivalence scales. For these data sources, as well as for survey microdata for which information on taxes and transfers is incomplete, we have to develop a strategy to transform the distribution of the observed "source concept" (e.g., posttax income among households) into an imputed distribution measured in a "target concept" (pretax or posttax income per adult).

To tackle this prediction problem, we choose to rely on XGBoost (Chen and Guestrin, 2016), a state-of-the-art implementation of a standard, high-performing machine learning algorithm called boosted regression trees. The key idea behind our harmonization procedure is that while the income or consumption concepts we observe are different, they are also related. Using all the cases in which the income distribution is simultaneously observed for two different concepts, we can thus map the way they tend to relate to one another, and convert any source concept to our concept of interest. We provide a detailed overview of the method and results of this imputation procedure in appendix section D.1.3. In particular, we show that this approach performs better than more naive ones, such as assuming a single correction coefficient by percentile.[14] Overall, this harmonization only has a small impact on our results, given that we observe both pretax and posttax income in the majority of cases and that corrections to equivalence scales only have limited impact on estimates of the income distribution.

3) Combination of Surveys and Tax Data Survey data are known to often miss the very rich. For our purpose it is important to distinguish two reasons for that: non-sampling error and sampling error.

Non-Sampling Error. Non-sampling error refers to the systematic biases that affect survey estimates in a way that is not directly due to sample size. These mostly include people refusing to answer surveys and misreporting their income in ways that are not observed, and therefore not corrected, by survey producers. We correct the survey data for non-sampling error by combining them with top income shares estimated from tax data using standard survey calibration methods. Statistical institutes already routinely apply these methods to ensure survey representativity in terms of age or gender. We directly extend them to enforce representativity in terms

[14]While this approach is certainly not perfect, the existing literature has often chosen to ignore these issues altogether, and directly compare and combine, say, income and consumption data (e.g. Lakner and Milanovic, 2016). We feel that our approach is preferable, because it corrects at least for what can be corrected.

of income, by adding top income shares based on tax data as a calibration margin.[15]

Sampling Error. Sampling error refers to problems that arise purely out of the limited sample size of survey data. The sample size of surveys varies a lot and can sometimes be quite low: this, in itself, can affect estimates of inequality at the top. Borrowing methods from extreme value theory, we correct sampling error by modeling the top 10% of the income distribution as a generalized Pareto distribution (see appendix section D.1.4). Note that by construction, this adjustment has no impact on the top 10% income share (which we know from the tax data), but only refines the income distribution within the top 10%.

Correcting survey-based estimates using top income shares derived from tax data has a large impact on our estimates of the income distribution, because surveys tend to significantly underestimate both the level of top income inequality and its rise since the 1980s in most European countries.

4) Distribution of Unreported National Income Components Once we have harmonized and corrected survey data with tax data, we find ourselves with more accurate and comparable inequality series. However, these series still lack some components of national income from the household sector (imputed rents), the corporate sector (undistributed profits), and the government sector (taxes on products and government spending) (see table 4.1).

Imputed Rents. Imputed rents are not always recorded in household surveys, and they are not included in the income concepts used in survey tabulations. To distribute them, we rely on EU-SILC surveys, which do record imputed rents, and perform a simple statistical matching procedure, using income as a continuous variable, to add imputed rents to the rest of our series (see appendix section D.1.5). The method preserves the rank dependency between income and imputed rents in EU-SILC, the distribution of imputed rents in EU-SILC, the distribution of income in the original data, and the imputed rents total in the national accounts.

Undistributed Profits. As we explain in section 4.1.1, undistributed profits are

[15] One advantage of calibration procedures, in particular, is that they allow to perform survey correction with a taxable income concept that may differ from the income concept of interest—either in terms of income definition or statistical unit. Accordingly, we always perform the correction by matching income concepts in the tax data and in the survey data. Importantly, this allows us to account for top incomes while retaining the wealth of information included in the surveys, notably on taxes and transfers, so that we can still calculate both pretax and posttax incomes after correction. For the historical period (typically before 2007), for which we do not have survey microdata to match precisely to the tax data concepts, we retropolate the adjustment observed in recent years to the tax-based top share series (see appendix section D.1.4.3).

distributed partially in proportion to the ownership of corporate stocks (including both private and public shares held directly or indirectly through mutual funds), partially in proportion to labor and pension income (for the fraction that accrues to pension funds), and partially like government primary income (for the fraction that accrues to the government). These respective shares correspond the fraction of corporate equity owned by households, governments and pension funds. The distribution of stock ownership comes from the Household Finance and Consumption Survey (HFCS).[16] We first calibrate that survey on the top income shares as we do for other surveys to make it representative in terms of income. We then use the same statistical matching procedure as above to allocate undistributed profits alongside the distribution of income.[17]

Corporate Income Taxes. Because the corporate income tax is paid out of corporate profits, we distribute it similarly to undistributed profits.

Indirect Taxes. Indirect taxes (including VAT and excise taxes) are eventually paid by consumers, so we allocate them proportionally to household final consumption expenditure. For that, we rely on the Household Budget Surveys (HBS) from Eurostat to get the distribution of consumption and its dependency to income. We then use the same statistical matching procedure as above to allocate indirect taxes to individuals.

4.1.4 Validation of our Methodology

Impact of the Different Methodological Steps Our estimates differ from existing survey-based estimates for two main reasons: because we use tax data at the top of the distribution, and because we incorporate forms of income that are traditionally absent from inequality statistics. How do these elements impact our results? Figure 4.1a gives the answer.[18] Based only on survey data, which do not add up to national income, we would conclude that inequality has been slightly declining in Europe after a one-time increase in the early 1990s: the top 10% income share has been stable after 1995, while the bottom 50% share has been slightly but consistently on the rise. When using tax data to correct the top of the distribution, we get a fairly different picture: the increase in the top 10% share has been much more significant, while the share of the bottom 50% has been stable. Adding missing national income

[16]In the United Kingdom we use its equivalent, the Wealth and Assets Survey (WAS).

[17]The HFCS only started around 2013, so before that year we keep the distribution of retained earnings constant and only change the amount of retained earnings to be distributed.

[18]See the extended appendix for the impact of our different methodological steps country by country.

components further modifies the distribution of income. Some components (such as undistributed profits) have a strong unequalizing impact, while others (such as imputed rents) have more equalizing effects. Overall, we distribute between one fifth and one quarter of national income in the form of additional income components. This leads to our DINA series, which show a slightly higher top 10% income share in recent years than survey and tax data alone. Most of the difference with raw survey estimates, however, comes from tax data.[19]

Comparison with Earlier Works Existing studies comparing inequality levels between the US and Europe have typically relied on surveys.[20] This implies making strong assumptions on the distribution of missing incomes in one region or the other, typically considering that these sources of income are distributionally neutral. While our method is not perfect, it has the advantage of making these assumptions explicit and ground them in the latest empirical evidence.

In particular, we wish to provide results that are conceptually similar to other works on distributional national accounts, yet in practice our methodology is quite different. In France, Garbinti, Goupille-Lebret, and Piketty (2018) and Bozio et al. (2018) estimated the distribution of pretax national income and posttax disposable income using detailed tax microdata, combined with various surveys and microsimulation models for taxes and benefits, and rescaling income component by component to the national accounts. By contrast, we only use tax tabulations to correct survey data, and rescale our results to the national accounts at a coarser level. The advantage of our method is that it is applicable much more widely and rapidly, in particular in countries in which no tax microdata is available.

To what extent can our approach yield results that are comparable to more complex and detailed works? As figure 4.1b shows, we get results that are very similar to these earlier works in the case of France. Concretely, our methodological approach starts from the raw survey series shown on the bottom line, which suggest that the top 10% share has fluctuated between 22% and 26%. In a second step, we calibrate

[19]Moving from survey to DINA estimates does not only increase estimates of income concentration: it also significantly affects the ranking of European countries in terms of pretax income inequality and in terms of the intensity of the rise of top income concentration since the 1980s. See in particular appendix figures D.16, D.17, D.18, and D.19. In 2017, for instance, accounting for misreporting of top incomes in surveys and unreported national income components increases the estimated top 1% share by 10 percentage points in Poland, compared to only half a percentage point in the Netherlands. As a result of this correction, Poland moves from being one of the least unequal countries of Europe to the most unequal in terms of pretax income. More generally, surveys tend to better capture top incomes in Northern European countries, where survey responses are often corrected *ex post* using administrative data, than in Eastern Europe.

[20]See footnote 4.

these distributions to the top income shares measured from tax data. In a third step, we impute additional sources of income, such as retained earnings and imputed rents. This yields the DINA top 10% pretax income share, which closely follows the series estimated by Garbinti, Goupille-Lebret, and Piketty (2018). Finally, we impute all taxes and cash transfers to derive the top 10% posttax disposable income share, which is also remarkably similar to that obtained by Bozio et al. (2018).[21]

Notice in particular that we obtain these results in spite of the fact that our data sources for France are not of especially high quality and are also very different from the ones used by Garbinti, Goupille-Lebret, and Piketty (2018) and Bozio et al. (2018).[22] All these results provide strong evidence that our methodology performs very well at reproducing more detailed DINA studies, despite the differences between our "bottom-up" approach combining survey microdata with tabulated tax data and "top-down" approaches that rely primarily on tax microdata.

4.2 The Distribution of Pretax National Incomes in Europe and the United States, 1980-2017

In this section, we show that pretax income inequality has risen much less in Europe than in the US since 1980. This is true for most European countries taken separately but also for Europe taken as a whole—a block that is broadly similar in terms of population size and aggregate economic output as the US. Section 4.2.1 presents results on the distribution of pretax income in Europe and the United States in 2017. Section 4.2.2 discusses the evolution of pretax income inequality in the two regions since 1980. Section 4.2.3 analyzes the role of spatial integration in accounting for the dynamics of inequality in Europe and the US.

[21]See appendix figure D.15 for similar results on the bottom 50% of the distribution.

[22]The SILC statistics for France are a transcription of a survey (called SRCV) that is used for its extensive set of questions on material poverty, but is not considered the best survey for income inequality. For that purpose, the French statistical institute relies on another survey, called ERFS. However, that survey is not part of any international scheme, such as EU-SILC, nor is it available through portals such as the Luxembourg Income Study, so we do not include it in our estimations. Before SILC is available, we rely on France's Household Budget Survey, which has been made available through LIS. While France's HBS is a key source for consumption data, it is not viewed as the best source for income data either. It is also separate from EU-SILC data, which explains the inconsistent trend. Therefore, there is no reason to think that our methodology would work better for France than other countries just because of the quality of the data in input.

4.2.1 The Distribution of Pretax Income in 2017

How do pretax incomes vary in Europe and the United States today? Table 4.3 provides a first answer to this question by displaying the average incomes and income shares of key income groups in Western Europe, Northern Europe, Eastern Europe, and the US in 2017. The average national income per adult stood at €52,700 in the US at purchasing power parity, compared to €44,900 in Northern Europe, €35,300 in Western Europe, and €21,700 in Eastern Europe. In Europe, only Norway (€55,000) and Luxembourg (€102,000) have higher average national incomes than the US.[23]

Things look very different at the bottom of the pretax income distribution. The bottom 50% earned only about €12,300 in the US in 2017, compared to €21,600 in Northern Europe and €14,600 in Western Europe. Of the twenty-seven countries considered in this paper, the US thus ranks third in terms of average national income per adult but nineteenth when it comes to the average income of the poorest 50%.[24] On average, pretax income inequality at the bottom is lowest in Northern Europe (with a bottom 50% share of 24%), followed by Western Europe (21%) and Eastern Europe (20%). With a bottom 50% pretax income share of only 11.7%, the US is by far the most unequal of all countries, followed by a distant Serbia (16%) and very far from the values observed in the Czech Republic, Iceland, Norway, and Sweden (all above 25%).[25] These differences appeared even more pronounced at the very bottom of the distribution: the average income of the poorest 20% was €11,600 in Northern Europe in 2017, more than three times larger than its counterpart in the United States (€3,800).

The same differences are visible at the top end of the distribution: the top 1% captured 21% of total pretax income in the US in 2017, compared to 12% in Eastern Europe, 11% in Western Europe, and less than 9% in Northern Europe. In 2017, the top 0.001% average pretax income exceeded €92 million in the US, nearly ten times the value observed in Northern Europe. The European countries with lowest top 1% income shares are the Netherlands, Slovenia, Iceland, Belgium, and Finland (less than 9%), while those with highest top income concentration are Germany, the United Kingdom, Greece, and Poland (13-15%).[26]

In summary, while the US stands out as being richer than most European countries today, differences in average national incomes mask substantial heterogeneity. With inequality levels surpassing by far those observed in any European country, the US

[23] See appendix table D.1.9.
[24] See appendix figure D.50.
[25] See appendix figure D.43.
[26] See appendix figure D.42.

displays bottom pretax average incomes that barely exceed those observed in poorer Eastern European countries. In contrast, the lower inequality levels and higher average incomes observed in Northern Europe imply significantly better standards of living for the majority of the population than in the United States.

4.2.2 The Distribution of Pretax Income Growth

We now turn to documenting the evolution of pretax income inequality in Europe and the US. Figure 4.2a shows the evolution of the top 10% pretax income share in the US, Eastern Europe, Western Europe, and Northern Europe from 1980 to 2017. The United States remained more unequal than most European countries throughout the entire period, but the gap between Europe and the US has widened significantly over time.[27] Indeed, the top 10% rose most rapidly and steadily in the US (from 34% to 48%), followed by Eastern Europe (from 24% to 36%), Western Europe (from 30% to 35%), and Northern Europe (from 26% to 31%). From 1980 to 2017, Eastern Europe shifted from being the least unequal to the most unequal European region. A significant part of this change occurred between 1989 and 1995, following the disintegration of the Soviet Union and the transition of Eastern European countries to market economies.[28]

The rise of top incomes has been a widespread phenomenon, yet there has been significant heterogeneity in the intensity of this rise across countries. Figure 4.2b plots the percentage point change in the top 10% pretax income share by country over the 1980-2017 period.[29] In Europe, inequality rose most strongly in Hungary, Poland, Romania, the Czech Republic, and Estonia, five Central and Eastern European countries that saw their economies shift from communist to capitalist systems during the 1990s. The US ranks third of all the countries considered here, with an increase in the top 10% share of almost 14 percentage points. In Western Europe, Germany is the country where the top 10% share grew the most (+ 9 percentage points),

[27]In 1980, the top 10% share was higher in Spain and Greece than in the US, and several Western European countries had inequality levels close to those observed in the US. This contrasts with the more recent period, when the US clearly stands out as being the most unequal of all countries studied in this paper: see appendix figure D.41.

[28]Let us stress here that we focus solely on monetary income inequality, which was unusually low in Russia and Eastern Europe under communism. Other forms of inequality prevalent at the time, in terms of access to public services or consumption of other forms of in-kind benefits, may have enabled local elites to enjoy higher standards of living than what their income levels suggest. That being said, the survey tabulations at our disposal do partially account for forms of in-kind income, so this limitation should not be exaggerated (see Milanovic, 1998). Furthermore, the top 10% income share did continue to rise in many Eastern European countries after 1995.

[29]See appendix figures D.45 and D.46 for similar results on top 1% and bottom 50% pretax income shares. Appendix figures D.32 to D.40 compare the evolution of pretax income inequality across countries by five-year intervals.

followed by Portugal and Italy (+ 8 pp). Meanwhile, several European countries saw pretax income inequality barely change in the past decades, including Spain, Greece, France, and Austria. In no European country, however, do we observe a significant long-run decline in the top 10% pretax income share.

Table 4.4 provides a more detailed picture of the rise of pretax income inequality by showing the real average annual income growth of selected income groups in our four regions of interest over the 1980–2017 and 2007–2017 periods.[30] National incomes grew at a modest yearly rate in the past four decades in Europe and the US: 1% in Western Europe, 1.2% in Eastern Europe, 1.4% in the US, and 1.8% in Northern Europe. In all regions, however, growth rates have been markedly higher the further one moves towards the top end of the distribution. The average pretax income of the top 1% thus rose at a rate of 1.9% in Western Europe, 3.2% in Northern Europe, 3.3% in the US, and 3.8% in Eastern Europe. Meanwhile, middle-income groups saw their average pretax incomes grow at a rate closer to the average of the full population in all regions. The bottom 20% benefited the least from real national income growth: their average income increased at a rate of 1.2% in Northern Europe and 0.7% in Western Europe, while it decreased at a rate of 1.3% in Eastern Europe and fell on average by 1.1% every year in the United States.

While the long run picture reveals a clear increase in inequality, the period of stagnation that followed the 2007–2008 crisis has been less detrimental to the European middle class than to other income groups. In Western Europe and Northern Europe, average earnings increased or stagnated for middle-income groups, while they decreased significantly at both tails of the distribution. Eastern European countries were less affected by the crisis but experienced a similar evolution: the bottom 20% grew at an annual rate of 1.6% between 2007 and 2017, lower than the regional average of 2.2%. Therefore, while the financial crisis has to some extent halted the rise of top income inequality in Europe, income gaps between the middle and the bottom of the distribution have continued to widen, and low incomes have consistently lagged behind the expansion of the overall economy. The rise of inequality has been much clearer and more pronounced in the United States: between 2007 and 2017, the bottom 20% saw their average pretax income decrease by 2.9% every year, while that of the top 1% expanded at an annual rate of 1%.

[30]The cumulated income growth rates of selected pretax income groups in Western Europe, Northern Europe, Eastern Europe, and the US are respectively represented in appendix figures D.26, D.27, D.28, and D.27.

4.2.3 The Distribution of Pretax Income in Europe as a Whole and the Role of Between-Country Inequalities

Our findings show that pretax income differences are lower and have risen less in most European countries than in the US in the past decades. Does this result hold, however, once considering inequality in Europe at large, that is after accounting for the important differences in average national incomes between Western, Northern, and Eastern European countries?

Figure 4.3a compares the levels and evolution of the top 1% and bottom 50% pretax income shares in the US, Europe as a whole, and Western and Northern Europe from 1980 to 2017.[31] Income inequality was unambiguously larger in the US than in Europe in 2017, even after accounting for differences in average incomes between European countries. The share of regional income received by the top percentile was almost twice as high in the United States (21%) as in Western and Northern Europe (11%) and Europe at large (11.5%). Meanwhile, the bottom 50% pretax income share reached 17% in Europe and 20% in Western and Northern Europe, compared to less than 12% in the US. This was not always the case: in 1980, the bottom 50% share was actually slightly higher in the US than in Europe as a whole (about 20% of national income) and only two percentage points lower than in Western and Northern Europe.

A more detailed picture of the distribution of growth in Europe and the US is displayed in Figure 4.3b, which plots the average annual income growth rate by percentile in the two regions from 1980 to 2017, with a further decomposition of the top percentile.[32] Average income growth has been slightly higher in the US (1.4% per year) than in Europe (1.1%) in the past four decades, yet this average gap hides substantial differences throughout the distribution. The average pretax income of the top 0.001% grew at a rate of 3.7% in Europe as a whole and as much as 5.4% per year in the US. Meanwhile, low-income groups have benefited significantly more from macroeconomic growth in Europe than in the US: the average income of the bottom 50% grew positively in Europe, while it stagnated in the US and even declined for

[31]We estimate pretax income distributions for Europe as a whole and for Northern and Western Europe by aggregating country-level distributions after converting average national incomes at market exchange rates euros rather than at purchasing power parity. This approach is justified by the fact that PPP conversion factors exist for European countries but not for US states: it would be unclear why one would correct for spatial differences in the cost of living in the former case but not in the latter. Estimating the distribution of European-wide income at purchasing power parity slightly reduces European inequality levels, as well as the share of inequality explained by between-country income disparities, so it does not affect our main conclusions.

[32]See appendix figure D.25 for similar results on each European region.

the bottom 30% of the population. The two growth incidence curves cross at the 67th percentile, that is, while average pretax income growth has been higher in the US than in Europe, it has been lower for the bottom 67% of the US population than for all corresponding European income groups.

To what extent are these differences driven by pretax income inequality between US states and between European countries, rather than within states and within countries? A Theil decomposition of within-group and between-group inequality in Europe and the US is shown in figure 4.4. The Theil index has risen much more in the US than in the Europe, and this change has been entirely due to increases in inequality within US states. In 1980, the Theil index in the US was almost perfectly equal to that of Europe at large, reaching about 0.45; by 2017, it had become higher than 1 in the US, whereas it did not exceed 0.6 in Europe. The overall Theil index and the Theil index of within-state inequality are almost indistinguishable in the US: within-state inequality explained 97% of overall US inequality in 1980 and 98% in 2017. The share of inequality explained by the between-group component has remained larger in Europe, but it has decreased from about 24% in 1980 to 17% in 2017, due mainly to the rise of pretax income inequality within European countries. In other words, macroeconomic convergence in Europe has become increasingly insufficient to reduce inequalities between European residents, and within-country inequality continues to matter the most.

4.3 The Impact of Taxes and Transfers on Inequality

We now turn to discussing the impact of taxes and transfers on inequality in Europe and the United States. Section 4.3.1 and section 4.3.2 present results on the distribution of taxes and transfers. Section 4.3.3 studies the net direct impact of the tax-and-transfer system on pretax income inequality. Section 4.3.5 investigates to what extent taxes and transfers indirectly contribute to reducing pretax income inequality in Europe and the US.

4.3.1 The Structure and Distribution of Taxes

Before investigating the distributional impact of taxes, it is useful to briefly compare the size and composition of government revenue in Europe and the United States.[33]

[33]Appendix table D.7 presents the structure of taxes and transfers in Europe and the United States, expressed as a share of national income, over the 2007-2017 period. Appendix figures D.10,

In 2007–2017, taxes and social contributions amounted to 47% of national income in Europe, compared to 28% in the United States. The United States collected less tax revenue than any European country, from Romania (32%), the country with lowest tax revenue, to Denmark (57%), which displayed the highest tax to national income ratio. The gap between the two regions was driven by two components of revenue: social contributions, which represented 19% of national income in Europe versus 8% in the US, and indirect taxes (14% versus 7%). Meanwhile, both regions collected comparable amounts of revenue from income and wealth taxes (10-11%) and from corporate income taxes (3%). The macroeconomic tax rate was larger in Northern Europe (52%) than in Western Europe (48%) and Eastern Europe (41%), due mostly to the larger share of national income collected in income and corporate taxes. If one excludes contributory social contributions from the analysis (that is, contributions financing the pension and unemployment systems), then the gap between Europe and the US decreases but remains significant: 23% of national income was collected in non-contributory taxes in the US in 2007-2017, compared to 30% in Europe.

Figure 4.5a represents the level and composition of non-contributory taxes paid by pretax income group in Eastern Europe, Western and Northern Europe, and the United States in the past decade.[34] Two results clearly stand out. First, while taxes paid are lower in the US than in Europe for most pretax income groups, the taxation profile is unambiguously more progressive in the United States. The top 1% face a tax rate higher than 30% in the US, which is relatively comparable to what we observe in Western and Northern Europe. Meanwhile, bottom income groups are taxed at an average rate that is nearly twice as small in the US as in Europe. Second, the difference in tax progressivity between the two regions is mainly driven by indirect taxes, which represent a significantly larger share of national income in Europe than in the US. These taxes tend to be regressive, because they are paid proportionally to consumption. Eastern Europe is the region with the least progressive tax system, due to the importance of indirect taxes and to the low progressivity of income and wealth taxes. This reflects the fact that many Eastern European countries have opted for flat (or almost flat) income taxes, whereas Western and Northern European countries and the US have a relatively long history of progressive income taxes and have so far maintained increasing marginal income tax rates.

D.11, and D.12 present similar results disaggregated by country.

[34] This way to look at tax incidence is useful for international comparisons focusing on the entire support of the adult distribution (including pensioners and the unemployed), as it allows us to better analyze the distribution of taxes independently from demographic (pensions) or economic (unemployment) factors that might artificially blow up or reduce tax progressivity. A complementary view, focusing on the distribution of all taxes as a share of factor income among the working-age population, is presented at the end of this section.

Figure 4.5b ranks European countries and the United States according to a simple measure of tax progressivity: the ratio of the total tax rate faced by the top 10% to that of the bottom 50%. The composition of bars correspond to the composition of taxes paid by the top 10%. The US stands out as the country with the highest level of tax progressivity: the top decile faces a tax rate that is more than 70% higher than that of the poorest half of the population. By this measure, the European country with the most progressive tax system is the United Kingdom, followed by Norway, the Czech Republic, and France. Many European countries have values close to 1 on this indicator, corresponding to relatively flat tax systems, in which top income groups face a tax rate approximately equal to that of the bottom 50%. Several countries, in particular Serbia, Croatia, and Romania, are characterized by unambiguously regressive tax systems. As shown in the figure, the US also stands out as one of the countries where the top 10% pay the largest share of their pretax income in the form of income and wealth taxes, which points to the role of the income tax in enhancing tax progressivity at the top end of the distribution.

Looking at non-contributory taxes as a share of pretax income is useful to study tax progressivity independently from the pension and unemployment systems, whose significance may depend on demographic and economic factors that are not directly related to redistribution (such as the size of the elderly population). The downside of this approach is that it misses a share of payments that can legitimately be considered as taxes by individuals. We address this issue by reporting in the online appendix the distribution of total taxes paid as a share of factor income.[35] By doing so and by narrowing down the analysis to the employed and working-age (20–64) population, the analysis remains consistent and cross-country comparisons meaningful. The main conclusions are unchanged. Because social contributions fall on labor income and are generally set at fixed rates, they tend to be flat for most groups within the bottom 90% and regressive at the top. This turns the tax systems of Western and Northern European countries into approximately flat tax systems, while those of most Eastern European countries become strongly regressive at the top end of the distribution. Because social contributions are smaller in the United States than in Europe, the US tax system remains more progressive than that of all European countries (with the exception of the UK).

[35] See in particular appendix figures D.53 and D.59, which reproduce the results of figures 4.5a and 4.5b in terms of factor income.

4.3.2 The Structure and Distribution of Transfers

As for taxes, total government expenditure is significantly lower in the US (35% of national income) than in Europe (47%).[36] The difference between the two regions is due to cash transfers, which represent 9% of national income in the US compared to 23% in Europe. Within cash transfers, pensions are the aggregate that differs the most between the two regions (16% of national income in Europe versus 5% in the US), followed by family and social assistance transfers (5% vs. 3%) and unemployment and disability benefits (1.6% vs. 1.4%). Meanwhile, in-kind transfers in health, education, and other collective government expenditure are very similar in Europe and the US (25-26%, of which about 7% goes to health). Total government expenditure is higher in Northern Europe (51% of national income) and Western Europe (48%) than in Eastern Europe (42%), due mainly to the larger size of social assistance transfers (5% in Western and Northern Europe vs. 3.5% in Eastern Europe) and in-kind transfers (29% vs. 25% vs. 23%, respectively) in Western and Northern Europe.

Figure 4.6a presents the distribution of transfers across posttax income groups in Europe and the US, expressed as a share of posttax national income. Unsurprisingly, transfers are progressive in both the US and Europe: they represent over 60% of the posttax incomes of bottom deciles, compared to less than 30% of those of the top 1%. Pensions represent a smaller share of posttax income in the US than in Europe for all posttax income groups, while the distribution of other cash transfers is relatively similar between the two regions. Health payments are the most progressive type of transfers. In Europe, this is directly due to the fact that we distribute health expenditure on a lump-sum basis, assuming as a first approximation that all individuals benefit from the same in-kind transfer (see methodology). Health expenditure is also highly progressive in the United States, where public health spending is significant and targeted towards the very poor (via Medicaid). Other in-kind transfers are neither progressive nor regressive, because we assume that they are distributed proportionally to posttax disposable income (we come back to this assumption in the next section).

Figure 4.6b provides a complementary picture of the magnitude and progressivity of government expenditure by plotting total transfers received by the bottom 50% in European countries and the United States, expressed as a share of national income. The US ranks third in terms of the smallest share of national income transferred to the bottom 50% (about 13%), due mainly to lower expenditure on pensions. In

[36] See appendix table D.7.

Europe, transfers received by the poorest half of the population are smallest in Serbia (11%), followed by Romania (12%), Estonia (14%), and Poland (14%). Meanwhile, Denmark, the Czech Republic, Sweden, the Netherlands, Belgium, and Finland stand out as the European countries allocating the greatest share of national income to the poorest half of the population (22-23%, corresponding to slightly less than half of all government revenue in these countries).

4.3.3 The Net Impact of Taxes and Transfers on Inequality

On the one hand, taxes are lower and more progressive in the United States than in Europe. On the other hand, Europe redistributes a significantly greater fraction of national income to low-income groups than the US, although transfers are about as progressive in the two regions. What is the net impact of the tax-and-transfer system on inequality, and is it more progressive in the US or in Europe overall?

Figure 4.7a directly answers this question by representing the share of national income transferred by the tax-and-transfer system between pretax income groups in Eastern Europe, Western and Northern Europe, and the United States in 2017. The bottom 50% and the middle 40% are net beneficiaries of redistribution in all three regions, but the US tax-and-transfer system appears to be unequivocally more progressive. The bottom 50% in the US received a positive net transfer of 6% of national income in 2017, compared to about 4% in Western and Northern Europe and less than 3% in Eastern Europe. Meanwhile, the top 10% saw their average income decrease by 8% of national income in the US after taxes and transfers, compared to about 4% in Western and Northern Europe and 3% in Eastern Europe. The middle 40% benefits slightly more from redistribution in the US (2%) than in Europe (less than 1%).

Figure 4.7b represents the net transfer received by the bottom 50% in all European countries and the United States in 2017. Again, the US stands out as the country that redistributes the greatest fraction of national income to the bottom 50% (6%), followed by the United Kingdom, Norway, the Netherlands, France, and Belgium (4-5%). In all countries considered in this paper, the bottom 50% end up being net beneficiaries of redistribution. Serbia, Croatia, Spain, Switzerland, Estonia, and Hungary are the European countries that redistribute the lowest share of national income to bottom income groups.

Assumptions made on the allocation of collective consumption expenditure can have a large impact on estimated posttax inequality levels across countries. As discussed in the methodology section, our benchmark series follow Piketty, Saez, and Zucman

(2018) and allocate non-health in-kind transfers proportionally to posttax disposable income. However, we also consider alternative series in which we distribute all collective expenditure in a lump-sum way. Our main conclusions are unchanged. Under that scenario, the US is still the country redistributing the largest fraction of national income to the bottom 50% (about 11%), and the ranking of European countries on this measure also remains broadly the same (ranging from 5% to 10%). Our results on the evolution of posttax income inequality in the two regions are also maintained.[37] These are not surprising results, given that collective government expenditure represents approximately the same share of national income in Europe and the US and has remained relatively constant since the 1980s. As an additional robustness check, we make the polar assumption that all government consumption is distributed in a lump-sum way in Europe and proportionally to posttax income in the US. Even under this extreme and highly implausible scenario, we find that redistribution is not dramatically and unambiguously more progressive in Europe than in the US (see appendix figure D.72).

That being said, we acknowledge that the way we allocate this large component of government spending remains unsatisfactory. Ideally, one would like to distribute one by one specific types of expenditure in education, housing, infrastructure, and other areas of government intervention by combining microdata on individual use with macrodata on total spending by program. Unfortunately, while this should be done in the context of more precise country-level studies (for promising attempts, see for instance Aaberge and Atkinson, 2010; Germain et al., 2020; O'Dea and Preston, 2010), the data at our disposal simply does not allow us to do so for all the countries considered in this paper. We leave this for future research. For our purpose, what is important is that allocating collective expenditure in two polar ways (proportionally vs. lump sum) only marginally affects our comparison of the US and European countries, both in terms of trends and levels of inequality and redistribution.

4.3.4 Predistribution vs. Redistribution: Revisiting the Europe-US inequality gap

When comparing inequality in Europe and the United States, landmark publications on inequality such as OECD (2008) and OECD (2011) reached two main conclusions:

[37]See appendix figure D.80, which reproduces figure 4.7b assuming that all collective expenditure is allocated on a lump-sum basis. Appendix figure D.70 compares the evolution of top 10% and bottom 50% posttax income shares in Europe and the US under these two polar scenarios. Allocating collective expenditure in a lump-sum way reduces inequality significantly in both regions, but does not affect the trends observed.

that income is less concentrated in most European countries than in the United States, and that this gap is substantially larger in terms of posttax income than in terms of pretax income.[38] The policy implications of these findings are relatively clear: if high income inequality countries were to increase redistribution to its level observed in less unequal countries, they would get significantly closer to the inequality levels observed in the latter. Our results challenge this claim. As documented in previous sections, pretax income inequality appears to be considerably higher in the US than in Europe, and accounting for redistribution only marginally affects the US-Europe inequality gap. If anything, taxes and transfers reduce inequality more in the US than in Europe.

Why do our conclusions contradict the standard view on redistribution in Europe and the US? We find that this is the case for three main reasons.

First, OECD estimates rely exclusively on surveys, while we systematically distribute the entire national income by combining surveys with tax data and national accounts. Because household surveys tend to underestimate top income inequality more in the US than in Europe, our estimates lead to a significant upward revision of the gap in pretax income inequality between the two regions.

Second, standard estimates of redistribution only allocate direct taxes and transfers to individuals, thereby ignoring corporate taxes, indirect taxes, and in-kind transfers. Distributing these components of government revenue and expenditure reverts the rankings of Europe and the US in terms of redistribution. This is because indirect taxes are much higher in Europe than in the US and fall disproportionately on low-income earners.

Third, our benchmark measure of redistribution compares pretax incomes to posttax incomes, while the standard view tends to compare factor incomes (sometimes referred to as market incomes) to posttax incomes. Because many European countries have a greater share of pensioners than the United States, and because public pension systems are much more developed in Europe than in the US, including pensions in the analysis leads to increasing estimates of redistribution more in the former than in the latter. However, as we now show, our conclusions are robust to using one or the other of these two income concepts.

To illustrate the role of these three factors in explaining the differences between our

[38] In OECD, 2011, redistribution as measured by the difference between market Gini and disposable income Gini is found to be 18% in the US vs. 40% in Sweden and 33% in Norway (p. 270). Similar findings are obtained in more recent OECD publications such as Causa and Norlem Hermansen (2017).

conclusions and the standard view, we compare in table 4.5 several estimates of the top 10% and bottom 50% income shares in Europe and the US in 2017. The table reports results obtained using three different methodologies (relying only on surveys, combining surveys and tax data, and following the DINA framework) and for three different income concepts (factor income, pretax income, and posttax income).

Survey-based estimates suggest that the top 10% factor income share is only slightly higher in the US (35.9%) than in Europe (33.3%). This gap is significantly larger in terms of posttax disposable income (4.7 percentage points) than in terms of factor income (2.6 pp.). By this measure, about 45% of the US-Europe inequality gap can be explained by redistribution, if we define redistribution as the gap between factor income and posttax income inequality. The differential impact of redistribution in the two regions appears even stronger when looking at the bottom 50% income share, which is lower in Europe than in the US in terms of factor income, but becomes higher when moving to pretax and posttax incomes. This corresponds relatively well to the standard view: by moving to European redistribution levels, the US would close a significant share of the US-Europe posttax income inequality gap.

If we combine surveys with tax data, we get a relatively different picture. The estimated top 10% factor income share increases by 7.6 percentage points in the US, compared to only 4.4 percentage points in Europe. As a result, the US-Europe gap in factor income inequality more than doubles, from 2.6 pp. in surveys to 5.8 pp. in estimates combining surveys with tax data. While taxes and transfers do continue to reduce inequality more in Europe than in the US, redistribution now appears to only explain about 19% of the posttax income inequality gap between the two regions (although the results continue to some extent to conform to the standard view when focusing on the bottom 50% income share).

Moving to DINA estimates further modifies the distribution of income in both regions. The allocation of unreported national income components (undistributed profits and imputed rents) increases factor income inequality more in the US than in Europe, shifting the gap in the top 10% income share from 5.8 to 8.1 percentage points. It also reverts the US-Europe gap at the bottom of the distribution: the bottom 50% factor income share now appears to be higher in Europe than in the US. By contrast, the difference between the two regions in terms of top posttax income inequality actually decreases from 7.2 to 6.7 pp (and from 5.8 to 5 pp in terms of the bottom 50% share). This is because moving from standard estimates of posttax income to DINA series implies allocating corporate taxes, indirect taxes, and in-kind transfers, which are more progressive in the US than in Europe overall. DINA estimates reveal that taxes and transfers reduce top inequality less in Europe than in the United States: the gap in the top 10% share between the two regions is 8.1 percentage points in terms of factor income, compared to 6.7 percentage points in terms of posttax income. Predistribution, not redistribution, explains why Europe is less unequal than the US.

Until now, we have compared factor income inequality to posttax income inequality for greater comparability with the existing literature. If we define redistribution as the gap between pretax income inequality and posttax income inequality, as in section 4.3, then the picture gets even clearer. Estimates from surveys, surveys and tax data, and DINA series all point to redistribution being higher in the United States than in Europe, both at the top and at the bottom of the income distribution. This is even more the case in DINA series than in surveys. According to our DINA estimates, greater redistribution in the US thus succeeds in closing 41% ($\frac{11.4-6.7}{11.4}$) of the US-Europe top pretax income inequality gap (and 42% of the gap in the bottom 50% share). This is a radically different conclusion from the one obtained by the OECD. Redistribution does not explain why Europe is less unequal than the US: it actually contributes to *reducing* the inequality gap between the two regions.

In our view, pretax income is more comparable across countries, because it avoids artificially inflating inequality and redistribution in countries with a large elderly population and public pension systems.[39] That being said, we acknowledge that pensions may contribute to reducing inequality within the elderly population, and social contributions may also have significant distributional consequences in some cases. Whether factor income or pretax income should be used as the benchmark concept remains an open question. What is important for our analysis, however, is that our results are robust to adopting one or the other of these two approaches to the measurement of redistribution.[40]

Finally, we do not find any evidence that redistribution has mitigated the rise of pretax income inequality more in Europe than in the US. Figure 4.8 represents the evolution of the top 10% and bottom 50% pretax and posttax income shares in

[39]In particular, rich pensioners may earn little factor income and therefore may appear to be lifted out of poverty by the pension system, even in a system in which pension benefits are proportional to income.

[40]We report in appendix table D.12 comparable results for the top 1% income share, the Gini index, and the Theil index. The results are in line with those discussed above. The Gini coefficient estimated from survey data, for instance, is 4.1 pp lower in the US than in Europe in terms of factor income, while it is 6.8 pp higher in terms of posttax disposable income. This conforms to the standard view. When moving to DINA estimates, by contrast, it appears to be unambiguously higher in the US across all income concepts (by 8.8 pp in terms of factor income, 14.4 pp in terms of pretax income, and 8.6 pp in terms of posttax national income).

the two regions from 1980 to 2017. In 1980, redistribution already appeared to be greater in the US than in Europe: for instance, the bottom 50% pretax income share stood at about 20% in both regions, while the bottom 50% posttax income share was significantly higher in the US (26%) than in Europe (22%). By 2017, the bottom 50% share has become significantly lower in the US, and the gap between the two regions is much larger in terms of pretax income (12% in the US versus 17% in Europe) than in terms of posttax income (18% versus 20%). Seen from this perspective, the greater inequality levels observed in the United States today appear to be a relatively recent phenomenon. When considering the European continent as a whole and properly accounting for redistribution, we find that posttax income disparities at the bottom of the distribution were in fact larger in Europe than in the United States only a few decades ago.

4.3.5 The Indirect Impact of Taxes and Transfers on Pretax Income Inequality

While the distinction between predistribution and redistribution is widespread and useful, it should be approached with care. Indeed, redistribution policies may have an impact on the distribution of pretax incomes themselves, not only on the gap between pretax and posttax income inequality. For example, high top marginal tax rates can limit top earners' incentives to bargain for higher pay, decreasing pretax inequality. Transfers at the bottom of the distribution can also change incentives to work or acquire skills. To what extent could these considerations change our conclusion regarding the role of redistribution? This section provides an exploration of this question. We investigate two channels: changes in pretax inequality due to changes in top marginal tax rates, and changes in pretax inequality due to net redistribution at the bottom.

Top Marginal Income Tax Rates The idea that high top marginal tax rates reduce top incomes has been suggested before (Piketty, Saez, and Stantcheva, 2014), and supported by cross-country evidence tying top marginal rates to reduced income concentration at the top. Using our own data, we indeed observe that higher top marginal tax rates are associated with lower top 1% pretax income shares. In appendix section D.1.8, we study different specifications for estimating the elasticity of the top 1% share with respect to (one minus) the top marginal tax rate across European countries. Across specifications, we find estimates ranging between $\sigma = 0.12$ and $\sigma = 0.45$, somewhat more muted than the findings of Piketty, Saez, and Stantcheva

(2014), but nonetheless significant overall.[41]

Let us assume that $\sigma = 0.5$, close to the benchmark of Piketty, Saez, and Stantcheva (2014) and at the high end of our own estimates. Based on this assumption of a rather strong impact of tax rates on pretax inequality, can we explain the evolution of European inequality and the difference between Europe and the United States? Figure 4.9a simulates two counterfactual evolutions of the top 1% pretax income share in Europe to answer this question: one that applies the United States' top marginal tax rate to every European country, and another that fixes top tax rates at their 1981 value in every country. The first scenario shows that the lower top marginal tax rates observed in the United States can only explain a small fraction of the Europe–US inequality differential. The second scenario shows that the decrease of top marginal tax rates generally observed in European countries can explain about 40% of the rise in within-country inequality observed since the 1980s. Therefore, the decrease of top marginal tax rates does contribute to explaining the rise of top income concentration in Europe, but it cannot account for the higher pace at which pretax inequality rose in the US. For top marginal tax rates to explain the entire difference between Europe and the United States, we would have to assume extremely high elasticities of the order $\sigma = 2$.

Net Transfers at the Bottom Now, focusing on the bottom of the distribution, it could be argued that there is a tradeoff between redistribution and predistribution, and that policymaking is about setting the equilibrium between the two. To assess this view, one can measure the correlation between the bottom 50% pretax income share and the net transfers received by the bottom 50%, as measured by the difference between their posttax and pretax income share. As shown in figure 4.9b, the cross-country correlation suggests a positive link between these two variables, with a small but positive elasticity of 0.10.[42] In other words, we find no evidence that redistribution and predistribution are substitutes, and if anything they may be complements. Since the United States redistributes a larger share of national income to the bottom 50% than European countries, a positive relationship between lower pretax inequality and higher redistribution cannot explain the differential between Europe and the United States.

The exploratory results of this section should be interpreted with care and not in a

[41]Differences with the results of Piketty, Saez, and Stantcheva (2014) arise mostly due to their inclusion of non-European countries and their longer time frame. See appendix D.1.8.

[42]We stress that this elasticity is only mildly significant and not robust to the inclusion of country fixed effects (see appendix section D.1.9).

strictly causal way. That being said, they do suggest the existence of some indirect effects of redistribution on predistribution, but not in a way that would overturn our key conclusions.

4.4 Conclusion

This article developed a new methodology to estimate the distribution of national income in 26 European countries between 1980 and 2017 by combining all available surveys, tax data, and national accounts in a systematic manner. The resulting dataset was then used to study the joint evolution of growth, inequality, and redistribution in Europe and the United States in the past decades.

Our results revealed that pretax income inequality has risen in almost all European countries since 1980. This rise has been concentrated at the top end of the distribution and has been most pronounced in Eastern Europe. However, income concentration has grown much less in Europe than in the United States. This is true of each European country taken separately but also of Europe as a whole. While inequalities between European countries remain significant, they only explain a small and decreasing fraction of European-wide income disparities.

Against a widespread view, we documented that the structure of taxes and transfers cannot explain why Europe is less unequal than the United States today. On the contrary, redistribution appears to reduce inequality more in the US than in Europe, despite the lower aggregate levels of taxes and transfers observed in the US. The novelty of this conclusion mainly arises from accounting for the underrepresentation of top incomes in surveys, which is more acute in the US than in Europe; from distributing the totality of national income, which leads to revising inequality estimates upwards more in the US than in Europe; and from allocating indirect taxes and in-kind transfers, which are more progressive in the US than in Europe. Given that the two regions have been exposed in a relatively similar way to technological change and globalization in the past decades, our results thus shed light on the importance of predistribution policies, such as access to education and healthcare or labor market regulations, in explaining international differences in the distribution of pretax income growth.

We see at least two avenues for future research. First, there is a need to better understand to what extent collective government expenditure in education, health, and other spheres of public intervention reduces inequality in the long run. While we have shown that our main conclusions are robust to polar assumptions on the distributional incidence of this form of redistribution, much remains to be done when it comes to precisely estimating it. Doing so would require combining distributional national accounts with more disaggregated data on who benefits from specific policies and programs.

Our dataset could also be used to better assess the distributional impact of taxes and transfers on inequality. Drawing on simple correlations and estimates from the existing literature, we have shown that changes in top marginal income tax rates or in net redistribution cannot entirely rationalize the diverging trajectories of Europe and the United States observed in the past decades. In the same spirit, further analyses could more systematically simulate, for instance, the effect of adopting specific tax-and-transfer systems of the distribution of pretax and posttax incomes. Such an enterprise would be particularly useful to better understand the sources of rising pretax income inequalities and to identify which policies affect them in the long run.

Bibliography

Aaberge, R. and A. B. Atkinson (2010). “Top incomes in Norway”. In: *Top incomes: a global perspective*. Ed. by A. B. Atkinson and Thomas Piketty. Oxford University Press. Chap. 9, pp. 448–481.

Alstadsæter, Annette et al. (2017). “Accounting for Business Income in Measuring Top Income Shares: Integrated Accrual Approach Using Individual and Firm Data from Norway”.

Alvaredo, Facundo et al. (2020). “Distributional National Accounts Guidelines: Concepts and Methods used in the World Inequality Database”. In: *World Inequality Lab Working Paper*.

Atkinson, Anthony B. (1996). “Income distribution in Europe and the United States”. In: *Oxford Review of Economic Policy* 12.1.

Atria, Jorge et al. (2018). “Top incomes in Chile: A Historical Perspective of Income Inequality (1964-2015)”.

Australian Bureau of Statistics (2019). “Australian National Accounts: Distribution of Household Income, Consumption and Wealth, 2003-04 to 2017-18”. In.

Beblo, Miriam and Thomas Knaus (2001). “Measuring income inequality in Euroland”. In: *Review of Income and Wealth* 47.3.

Blanchet, Thomas and Lucas Chancel (2016). “National Accounts Series Methodology”.

Blanchet, Thomas, Juliette Fournier, and Thomas Piketty (Apr. 16, 2021). “Generalized Pareto Curves: Theory and Applications”. In: *Review of Income and Wealth*, roiw.12510.

Bozio, Antoine et al. (2018). “Inequality and Redistribution in France, 1990-2018: Evidence from Post-Tax Distributional National Accounts (DINA)”. In: *WID.working paper series 2018/10*.

Brandolini, Andrea (2006). “Measurement of income distribution in supranational entities: The case of the European Union”. In: *LIS Working Paper Series, No. 452*.

Causa, Orsetta and Mikkel Norlem Hermansen (2017). *Income redistribution through taxes and transfers across OECD countries*. Tech. rep. OECD Economics Department Working Papers, No. 1453.

Chen, Tianqi and Carlos Guestrin (2016). "XGBoost: A Scalable Tree Boosting System".

Dauderstädt, Michael and Cem Keltek (2011). "Immeasurable Inequality in the European Union". In: *Intereconomics* 46.1, pp. 44–51.

European Central Bank (2010–2017). *Household Finance and Consumption Survey*.

Eurostat (1994–2001). *European Community Household Panel (ECHP)*.

– (2004–2020). *European Union Statistics on Income and Living Conditions (EU-SILC)*.

– (2010). *Household Budget Surveys (HBSs)*.

– (2018). "Income comparison: social surveys and national accounts". In.

Fairfield, Tasha and Michel Jorratt De Luis (2016). "Top Income Shares, Business Profits, and Effective Tax Rates in Contemporary Chile". In: *Review of Income and Wealth* 62.August, S120–S144.

Fesseau, Maryse and Maria Liviana Mattonetti (2013). "Distributional Measures Across Household Groups in a National Accounts Framework: Results from an Experimental Cross-country Exercise on Household Income, Consumption and Saving". In: *OECD Statistics Working Papers 2013/4*.

Filauro, Stefano and Zachary Parolin (2018). "Unequal unions? A comparative decomposition of income inequality in the European Union and United States". In: *Journal of European Social Policy*.

Garbinti, B., J. Goupille-Lebret, and T. Piketty (2018). "Income inequality in France, 1900-2014: Evidence from Distributional National Accounts (DINA)". In: *Journal of Public Economics* 162.1. Dataset, pp. 63–77.

Germain et al. (2020). "Rapport sur les comptes nationaux distributionnels". In: *INSEE Rapport*.

Guillaud, Elvire, Matthew Olckers, and Michaël Zemmour (2019). "Four Levers of Redistribution: The Impact of Tax and Transfer Systems on Inequality Reduction". In: *Review of Income and Wealth*.

Hacker, Jacob S. and Paul Pierson (2010). "Winner-Take-All Politics: Public Policy, Political Organization, and the Precipitous Rise of Top Incomes in the United States". In: *Politics & Society* 38.2, pp. 152–204.

Immervoll, Herwig and Linda Richardson (2011). "Redistribution Policy and Inequality Reduction in OECD Countries: What has changed in two decades?" In: *OECD Social, Employment and Migration Working Papers No. 122*.

Jesuit, David K. and Vincent A. Mahler (2010). “Comparing Government Redistribution Across Countries: The Problem of Second-Order Effects”. In: *Social Science Quarterly* 91.5, pp. 1390–1404.

Kleven, Henrik et al. (2020). “Taxation and Migration: Evidence and Policy Implications”. In: *Journal of Economic Perspectives* 34.2, pp. 119–142.

Lab, World Inequality (2020). *World Inequality Database.*

Lakner, Christoph and Branko Milanovic (2016). “Global Income Distribution: From the Fall of the Berlin Wall to the Great Recession”. In: *The World Bank Economic Review* 30.2, pp. 203–232.

LIS (n.d.). *Luxembourg Income Study (LIS) Database.*

Milanovic, Branko (1998). *Income, Inequality, and Poverty during the Transition from Planned to Market Economy.* Dataset. World Bank Regional and Sectoral Studies.

O’Dea, Cormac and Ian Preston (2010). “The distributional impact of public spending in the UK”. IFS Working Paper W12/06.

OECD (2008). *Growing Unequal? Income Distribution and Poverty In OECD Countries.* OECD Publishing.

OECD (2011). *Divided we stand: Why inequality keeps rising?* OECD Publishing.

Office for National Statistics, Social Survey Division (2006–2018). *Wealth and Assets Survey.*

Piketty, Thomas and Emmanuel Saez (2003). “Income inequality in the United States, 1913-1998”. In: *Quarterly Journal of Economics* 118.1, pp. 1–39.

Piketty, Thomas, Emmanuel Saez, and Stefanie Stantcheva (Feb. 2014). “Optimal Taxation of Top Labor Incomes: A Tale of Three Elasticities”. In: *American Economic Journal: Economic Policy* 6.1. Dataset, pp. 230–271.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman (2018). “Distributional National Accounts: Methods and Estimates for the United States”. In: *Quarterly Journal of Economics* 133.2, pp. 553–609.

– (2019). “Simplified Distributional National Accounts”. In: *AEA Papers and Proceedings* 109, pp. 289–295.

Saez, Emmanuel and Gabriel Zucman (2019). *The Triumph of Injustice: How the Rich Dodge Taxes and How to Make Them Pay.* W. W. Norton & Company.

Salverda, Wiemer (2017). “The European Union’s income inequalities”.

Statistics Canada (2019). “Distributions of Household Economic Accounts, estimates of asset, liability and net worth distributions, 2010 to 2018, technical methodology and quality report”. In.

Statistics Netherlands (2014). “Measuring Inequalities in the Dutch Household Sector”.

United Nations Statistics Commission (2008). *System of National Accounts 2008*. United Nations.

UNU-WIDER (n.d.). *World Income Inequality Database (WIID).*

Wolfson, Michael, Mike Veall, and Neil Brooks (2016). “Piercing the Veil – Private Corporations and the Income of the Affluent”. In: *Canadian Tax Journal* 64.1, pp. 1–25.

World Bank (n.d.[b]). *PovcalNet: the on-line tool for poverty measurement developed by the Development Research Group of the World Bank.*

Zwijnenburg, Jorrit, Sophie Bournot, and Federico Giovannelli (2019). “OECD Expert Group on Disparities in a National Accounts Framework: Results from the 2015 Exercise.” In: *OECD Statistics Working Papers 2019/76.*

Figure 4.1: Measuring Inequality: From Surveys to Distributional National Accounts

(a) Pretax Income Inequality in All 26 European Countries, 1980–2017

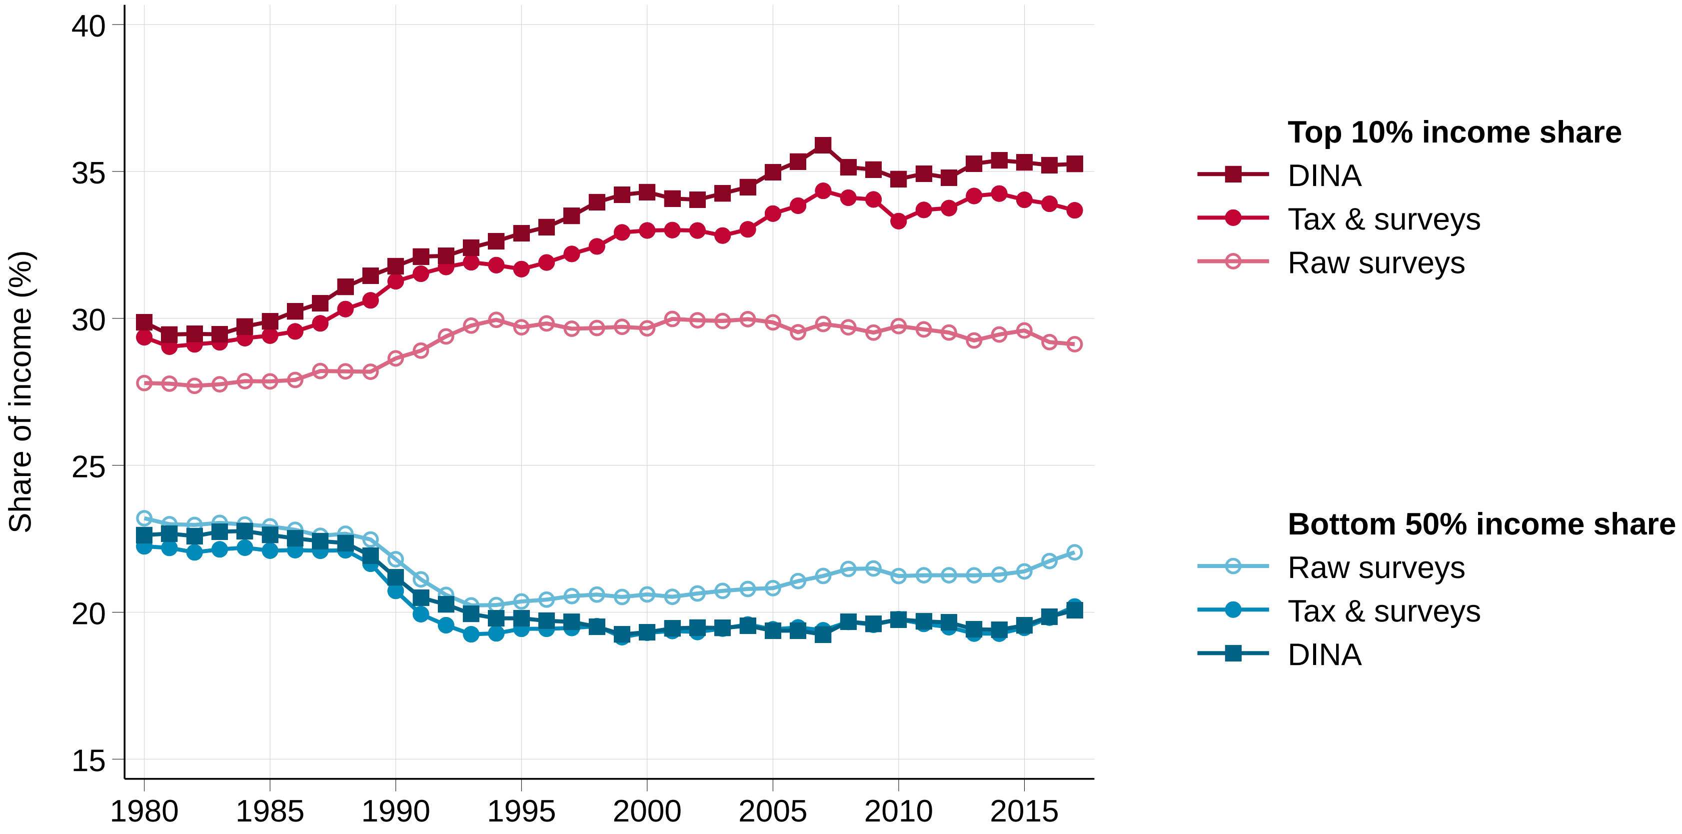

Source: Authors' computations combining surveys, tax data and national accounts. *Note*: Incomes measured at purchasing power parity. The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses, except for the "raw survey income" series in panel (b), for which income is split equally among all adult household members. Posttax DINA series distribute taxes on products proportionally to income for consistency with Bozio et al. (2018), see the online appendix for other approaches that follow the latest DINA guidelines.

Figure 4.1: Measuring Inequality: From Surveys to Distributional National Accounts

(b) Top 10% Income Share in France, 1978-2015

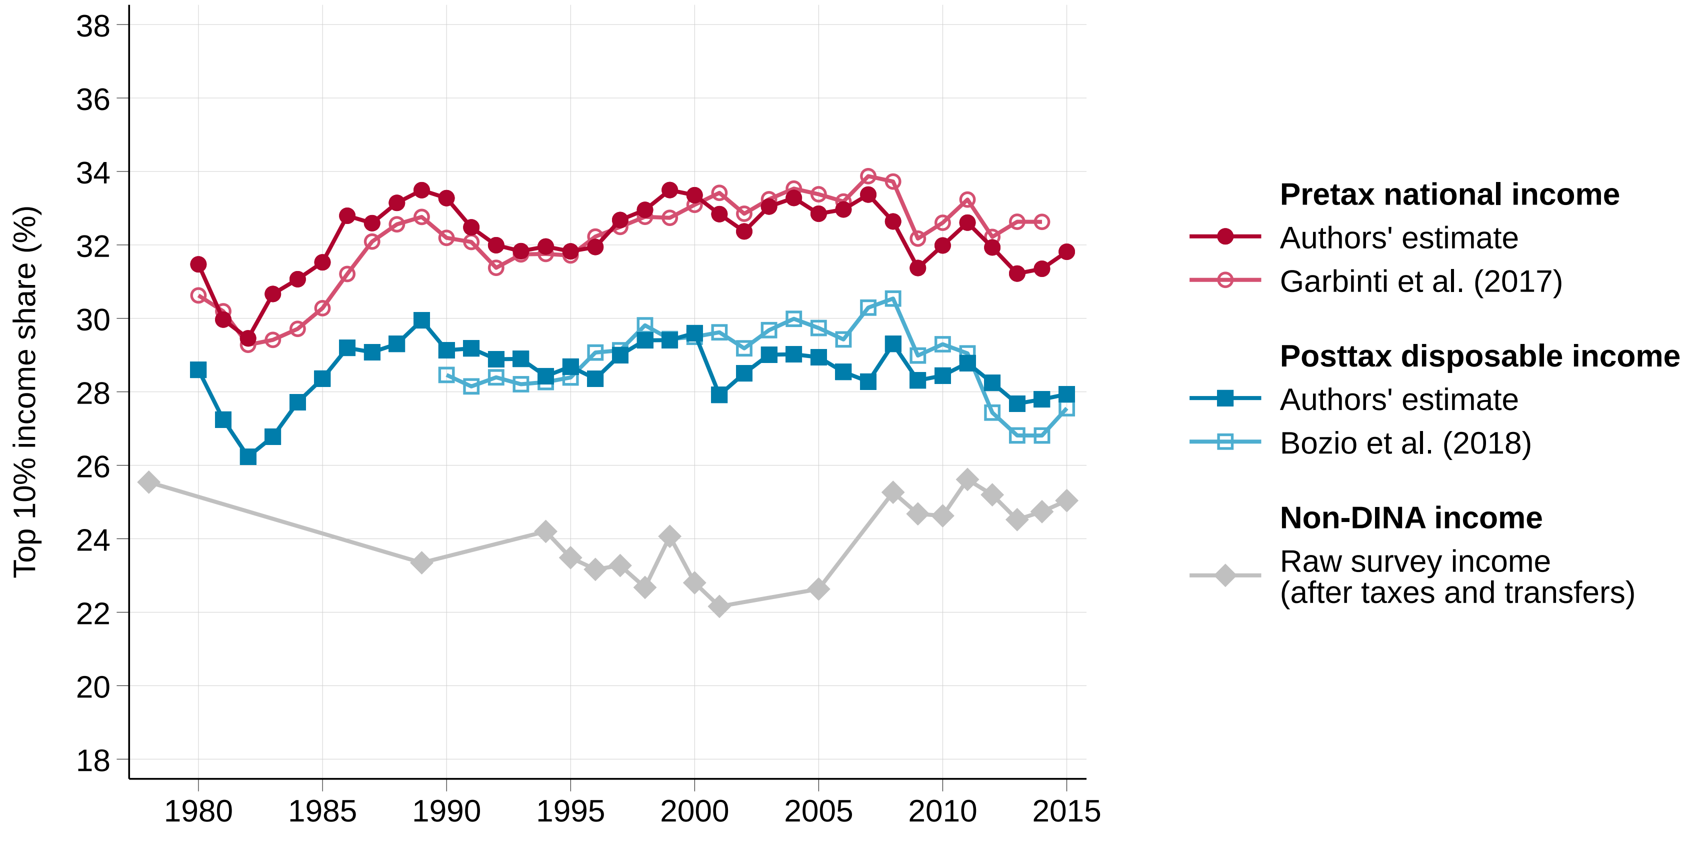

Source: Authors' computations combining surveys, tax data and national accounts. *Note*: Incomes measured at purchasing power parity. The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses, except for the "raw survey income" series in panel (b), for which income is split equally among all adult household members. Posttax series distribute taxes on products proportionally to income for consistency with Bozio et al. (2018), see the online appendix for other approaches that follow the latest DINA guidelines.

Figure 4.2: The Rise of Top Incomes in Europe and the United States, 1980-2017

(a) Top 10% Pretax Income Share, 1980-2017

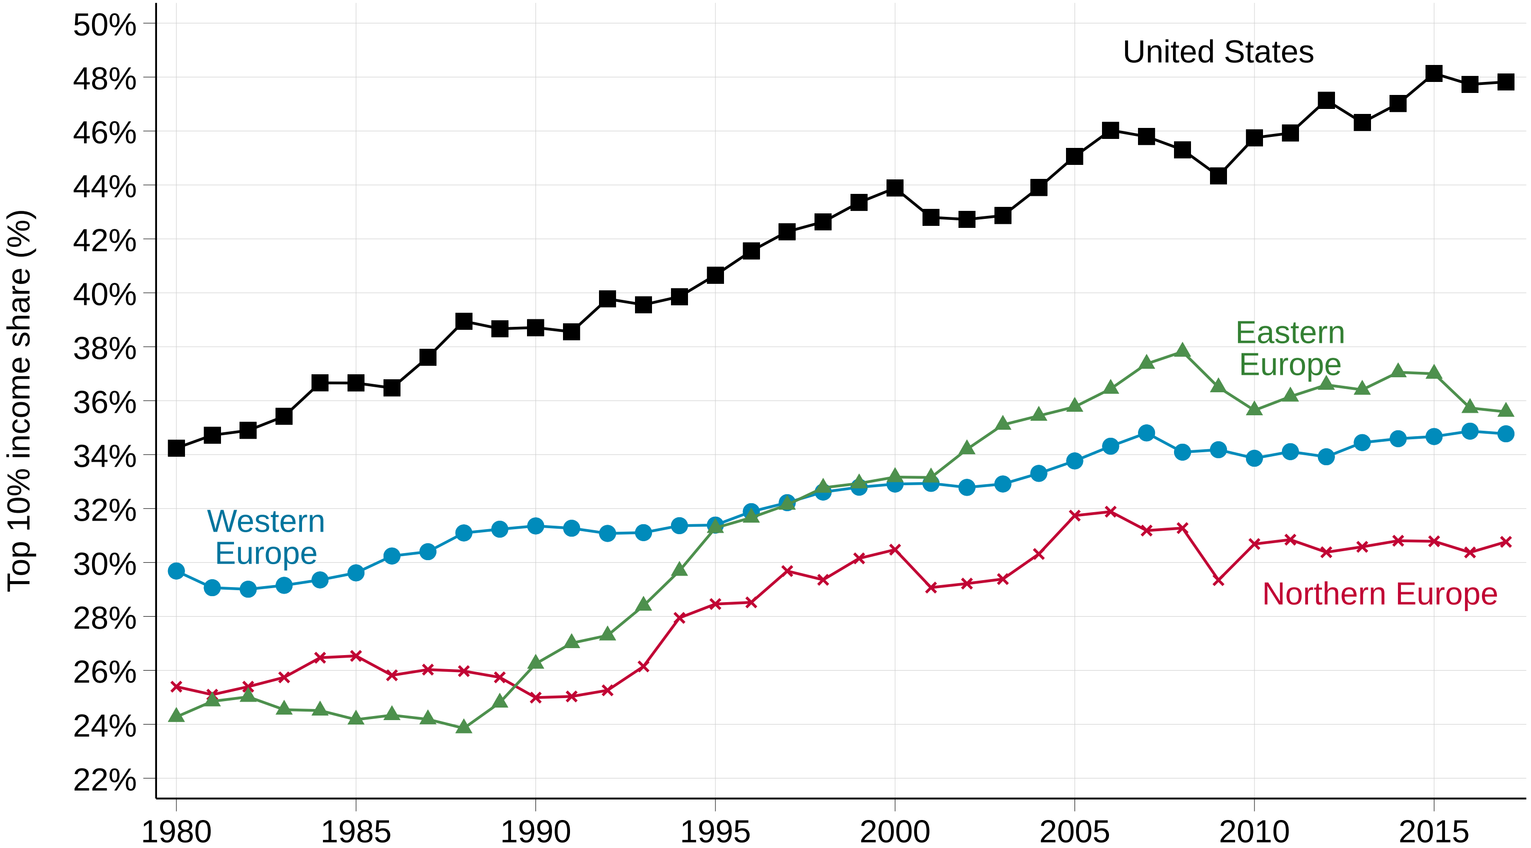

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* Panel (a) represents the evolution of the share of pretax income received by the top 10% in Western Europe, Northern Europe, Eastern Europe, and the United States. Panel (b) plots the percentage point change in the top 10% pretax income share by country between 1980 and 2017. The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. See online appendix table D.6 for the composition of European regions.

Figure 4.2: The Rise of Top Incomes in Europe and the United States, 1980-2017

(b) Percentage Point Change in Top 10% Pretax Income Share by Country, 1980-2017

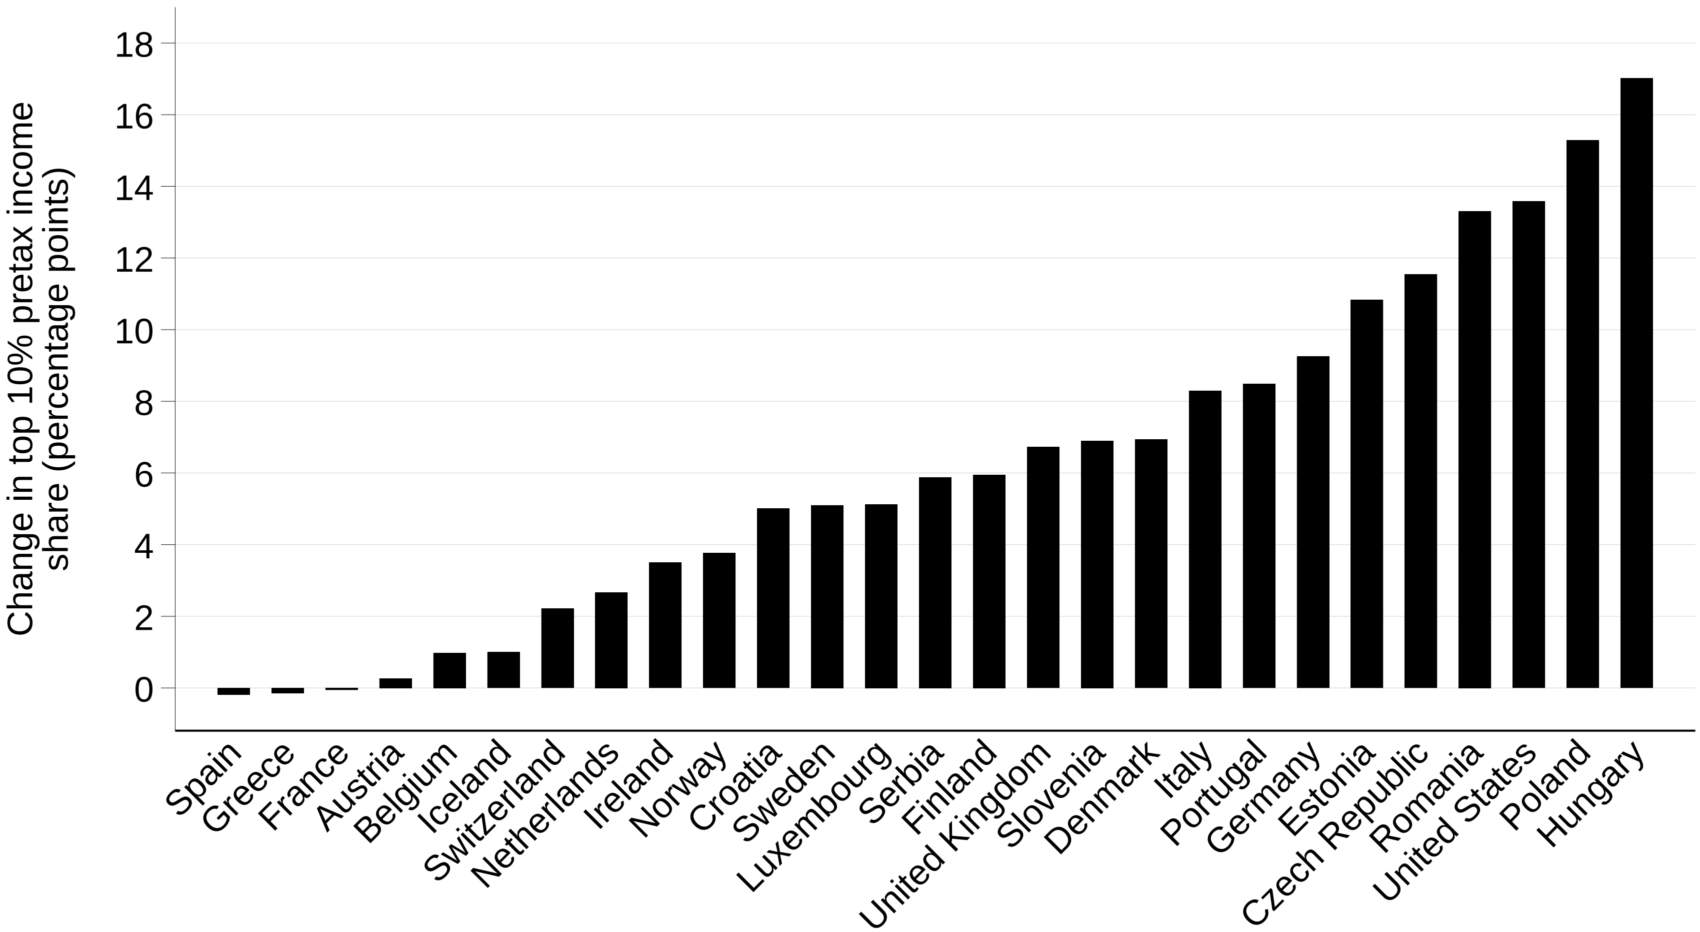

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* Panel (a) represents the evolution of the share of pretax income received by the top 10% in Western Europe, Northern Europe, Eastern Europe, and the United States. Panel (b) plots the percentage point change in the top 10% pretax income share by country between 1980 and 2017. The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. See online appendix table D.6 for the composition of European regions.

Figure 4.3: The Distribution of Pretax Income Growth in Europe and the United States, 1980-2017

(a) Top 1% versus Bottom 50% Pretax Income Shares

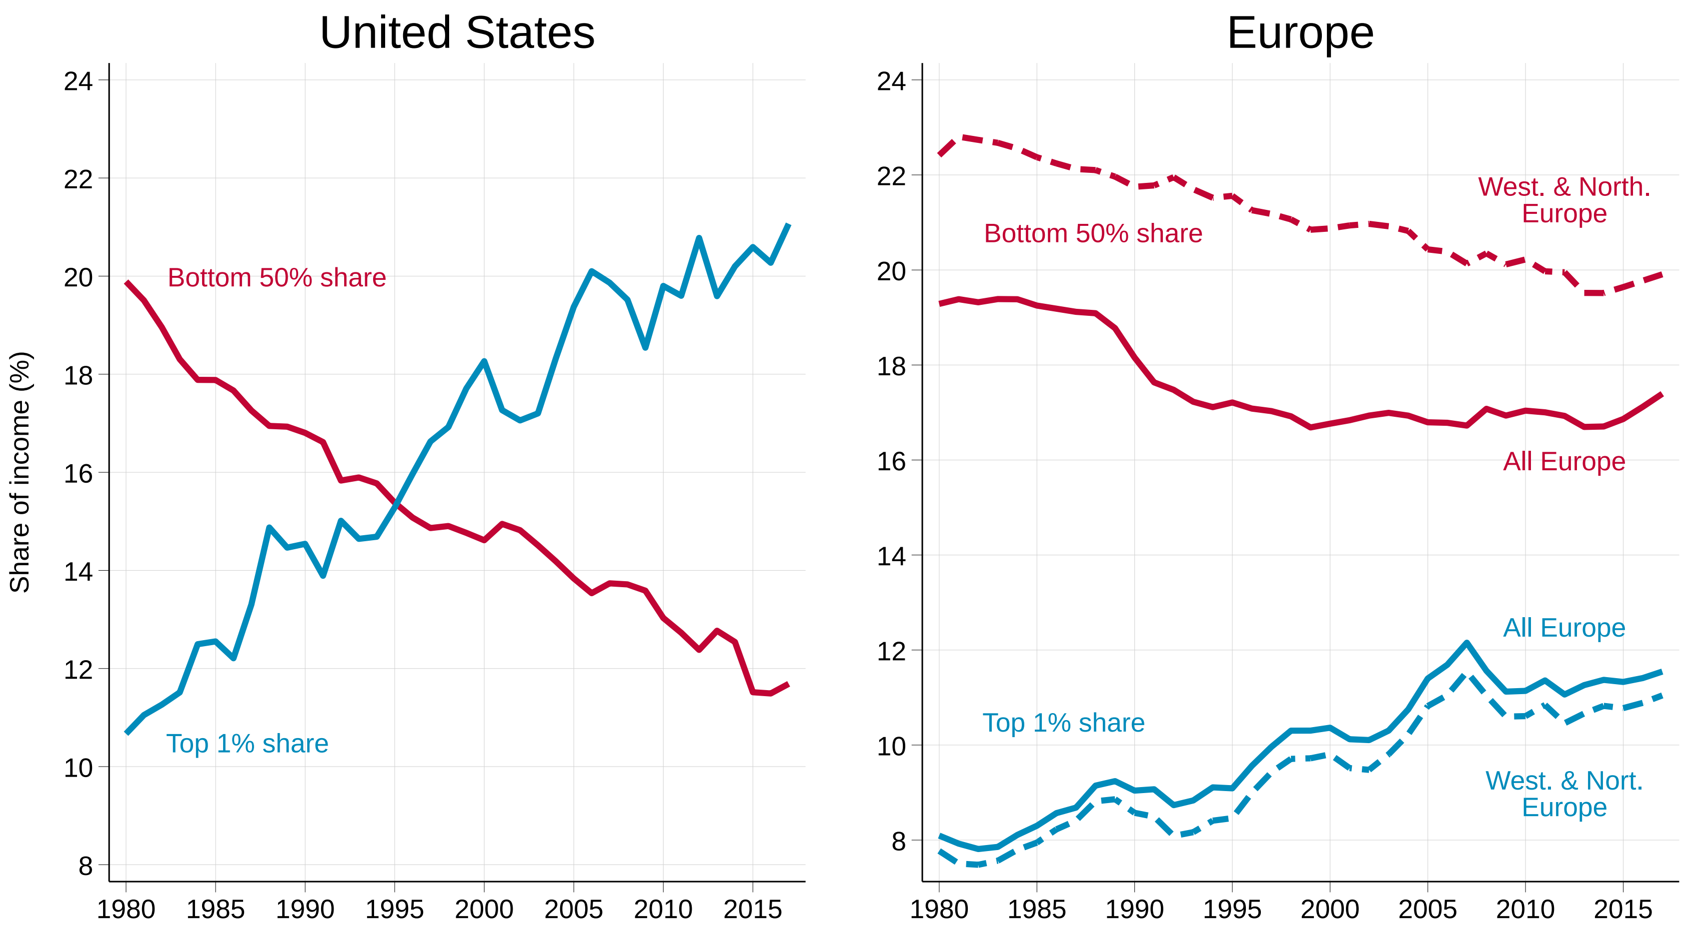

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* Panel (a) compares the share of pretax income received by the bottom 50% to that received by the top 1% of the regional population in Europe and the United States. Panel (b) plots the average annual pretax income growth rate by percentile in Europe and the US, with a further decomposition of the top percentile. Figures for the US come from Piketty, Saez, and Zucman (2018). Figures for Europe correspond to Europe at large, that is, after accounting for differences in average national incomes between European countries, measured at market exchange rates. The same holds for Western and Northern Europe. The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. See online appendix table D.6 for the composition of European regions.

Figure 4.3: The Distribution of Pretax Income Growth in Europe and the United States, 1980-2017

(b) Average Annual Pretax Income Growth by Percentile

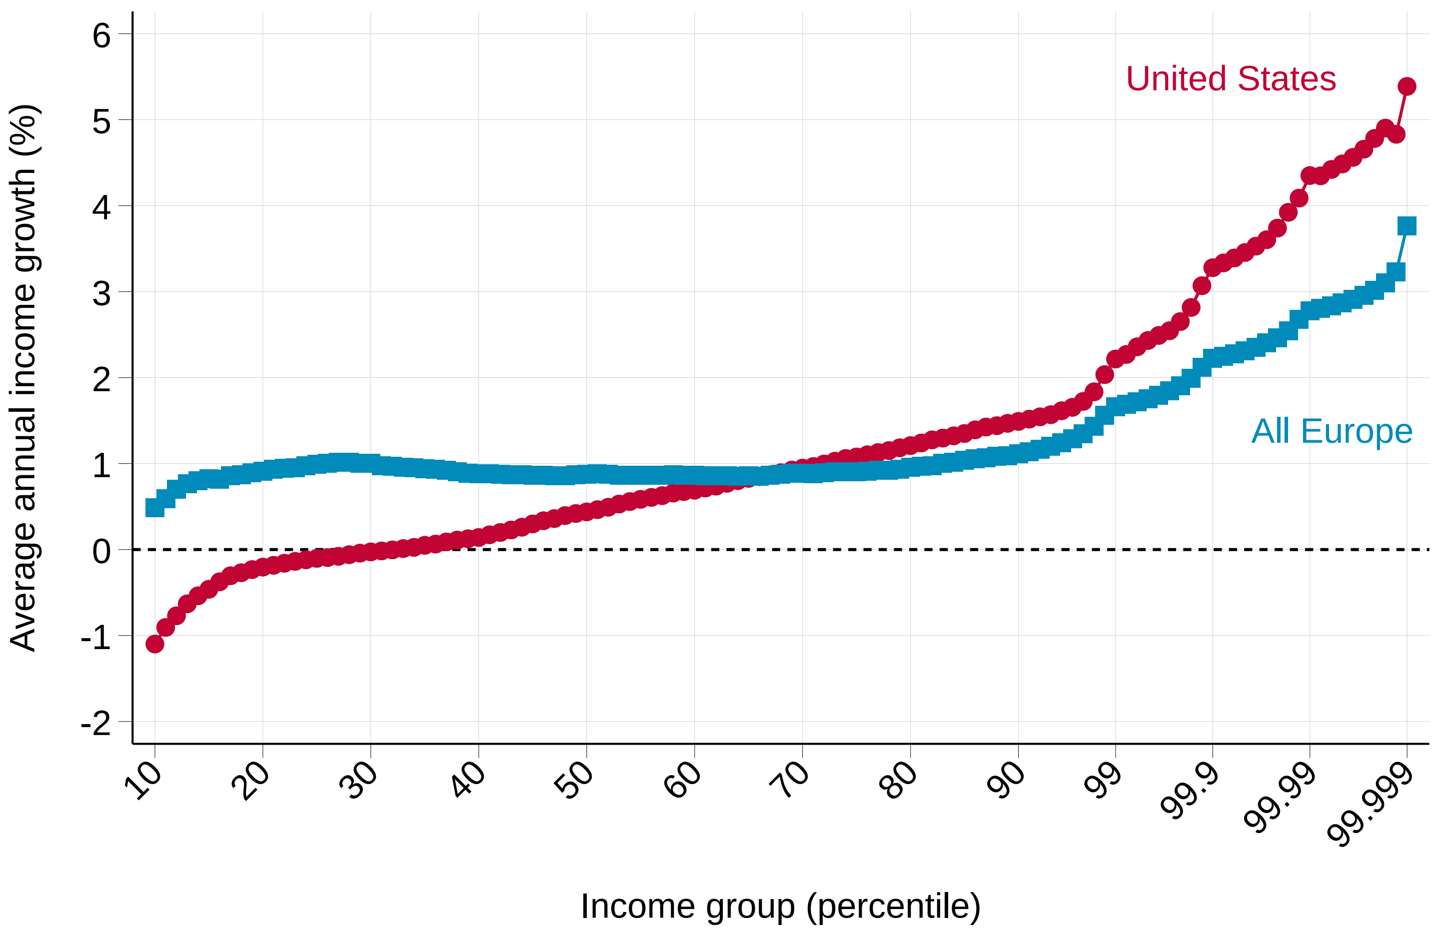

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* Panel (a) compares the share of pretax income received by the bottom 50% to that received by the top 1% of the regional population in Europe and the United States. Panel (b) plots the average annual pretax income growth rate by percentile in Europe and the US, with a further decomposition of the top percentile. Figures for the US come from Piketty, Saez, and Zucman (2018). Figures for Europe correspond to Europe at large, that is, after accounting for differences in average national incomes between European countries, measured at market exchange rates. The same holds for Western and Northern Europe. The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. See online appendix table D.6 for the composition of European regions.

Figure 4.4: Pretax Income Inequality in Europe and the United States, 1980-2017: Theil Decomposition

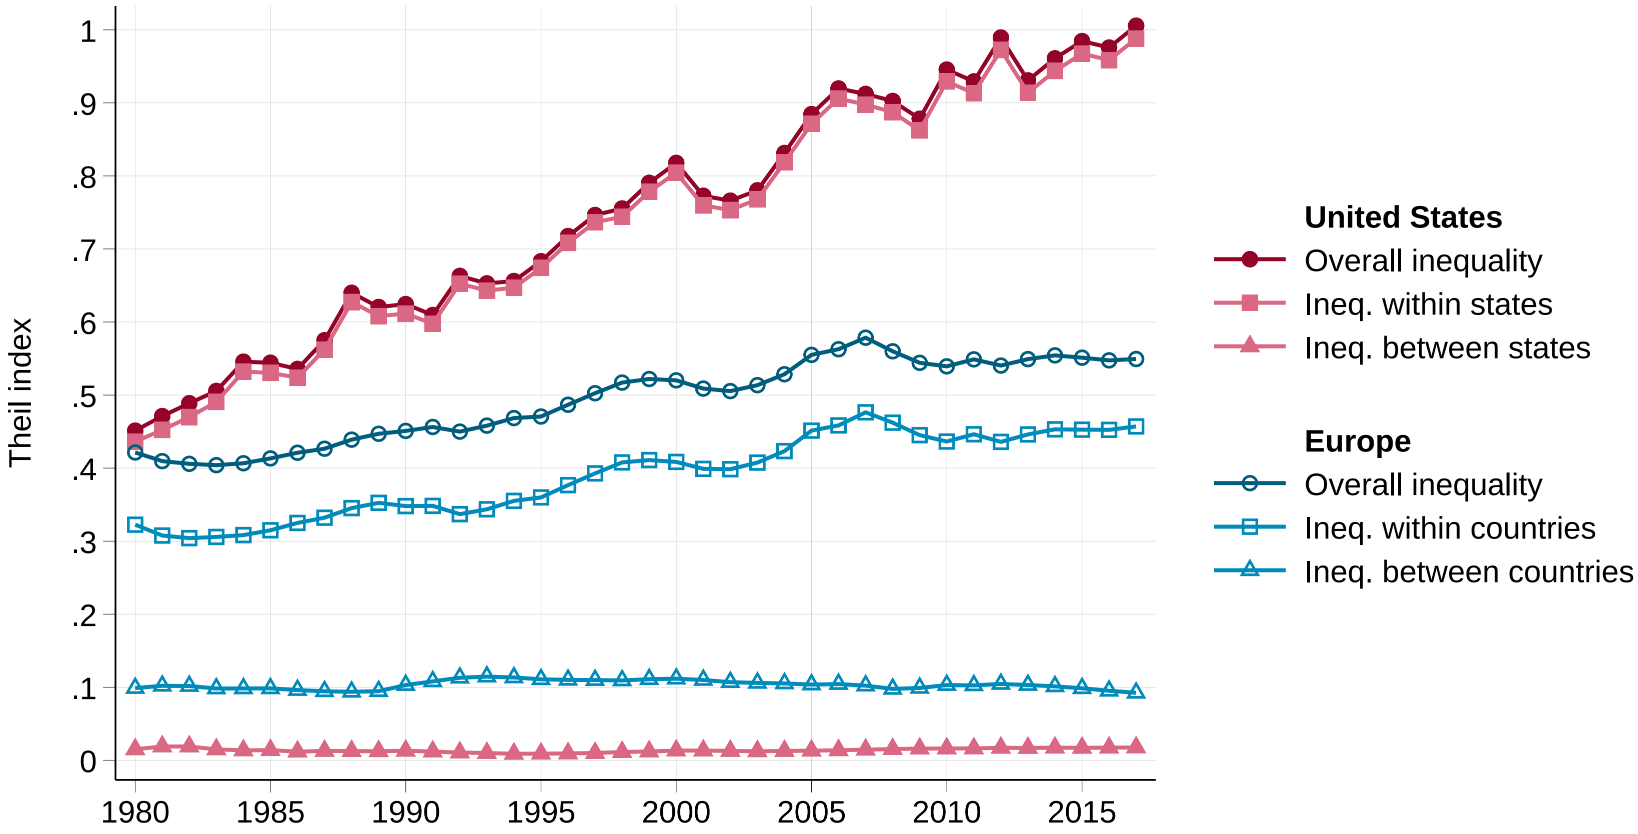

Source: Authors' computations combining surveys, tax data and national accounts for European countries. Figures for the US come from Piketty, Saez, and Zucman (2018) for the overall Theil index, and from state GDP estimates of the Bureau of Economic Analysis for the US between-group component. *Notes:* Figures for Europe correspond to Europe at large, that is, after accounting for differences in average national incomes between European countries, measured at market exchange rates. The income concept is pretax income. The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure 4.5: The Distribution of Taxes in Europe and the United States

(a) Non-contributory Taxes Paid as a Share of Pretax Income

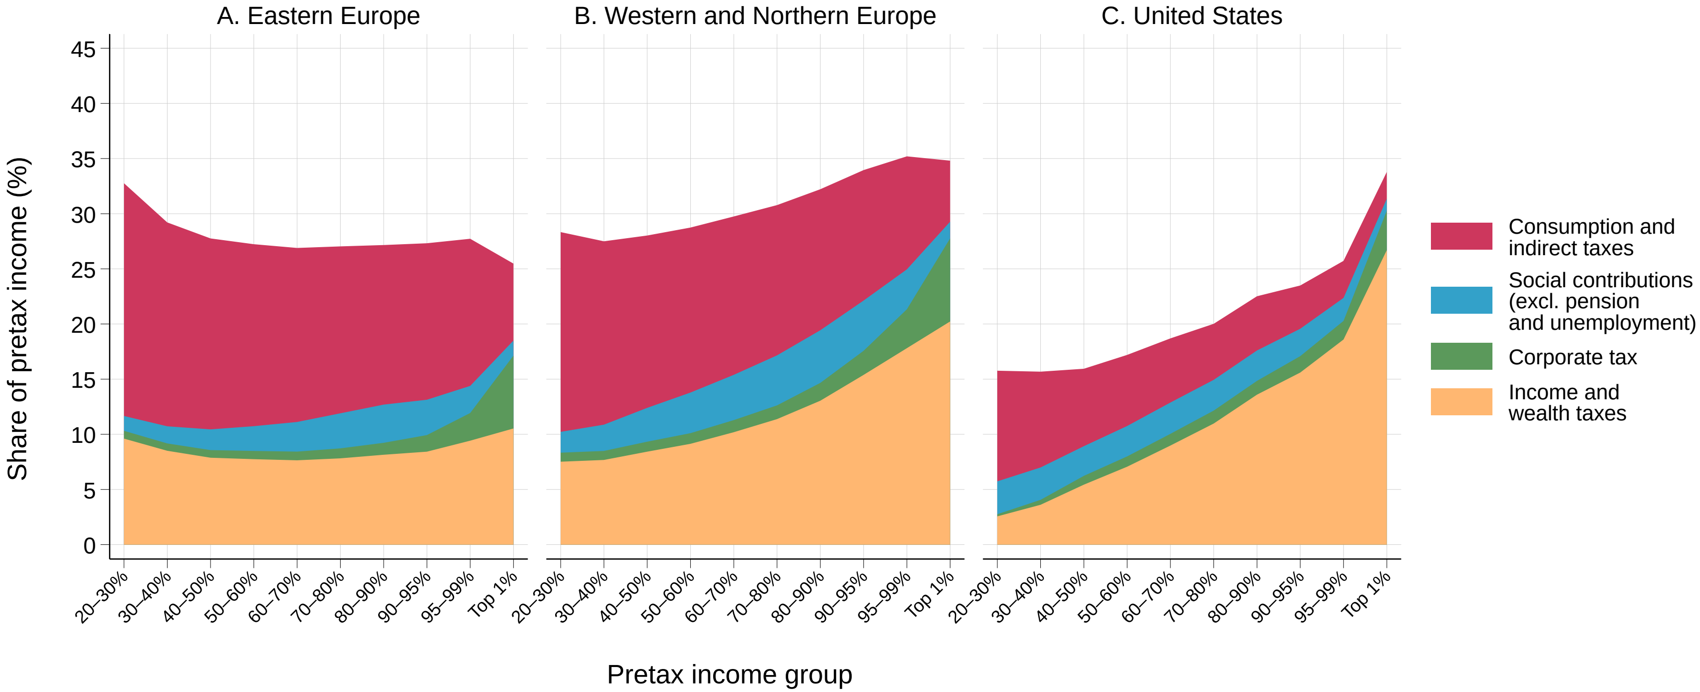

Source: Authors' computations combining surveys, tax data and national accounts for European countries; Piketty, Saez, and Zucman (2018) for the United States. *Notes:* Figures correspond to averages over the 2007–2017 period for European countries (population-weighted average of country-specific estimates in the case of European regions), and to 2017–2018 for the US. In panel (b), the composition of bars corresponds to the composition of taxes paid by the top 10%. The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. See online appendix table D.6 for the composition of European regions.

Figure 4.5: The Distribution of Taxes in Europe and the United States

(b) Level and Composition of Taxes Paid by the Top 10% Relative to the Bottom 50% by Country

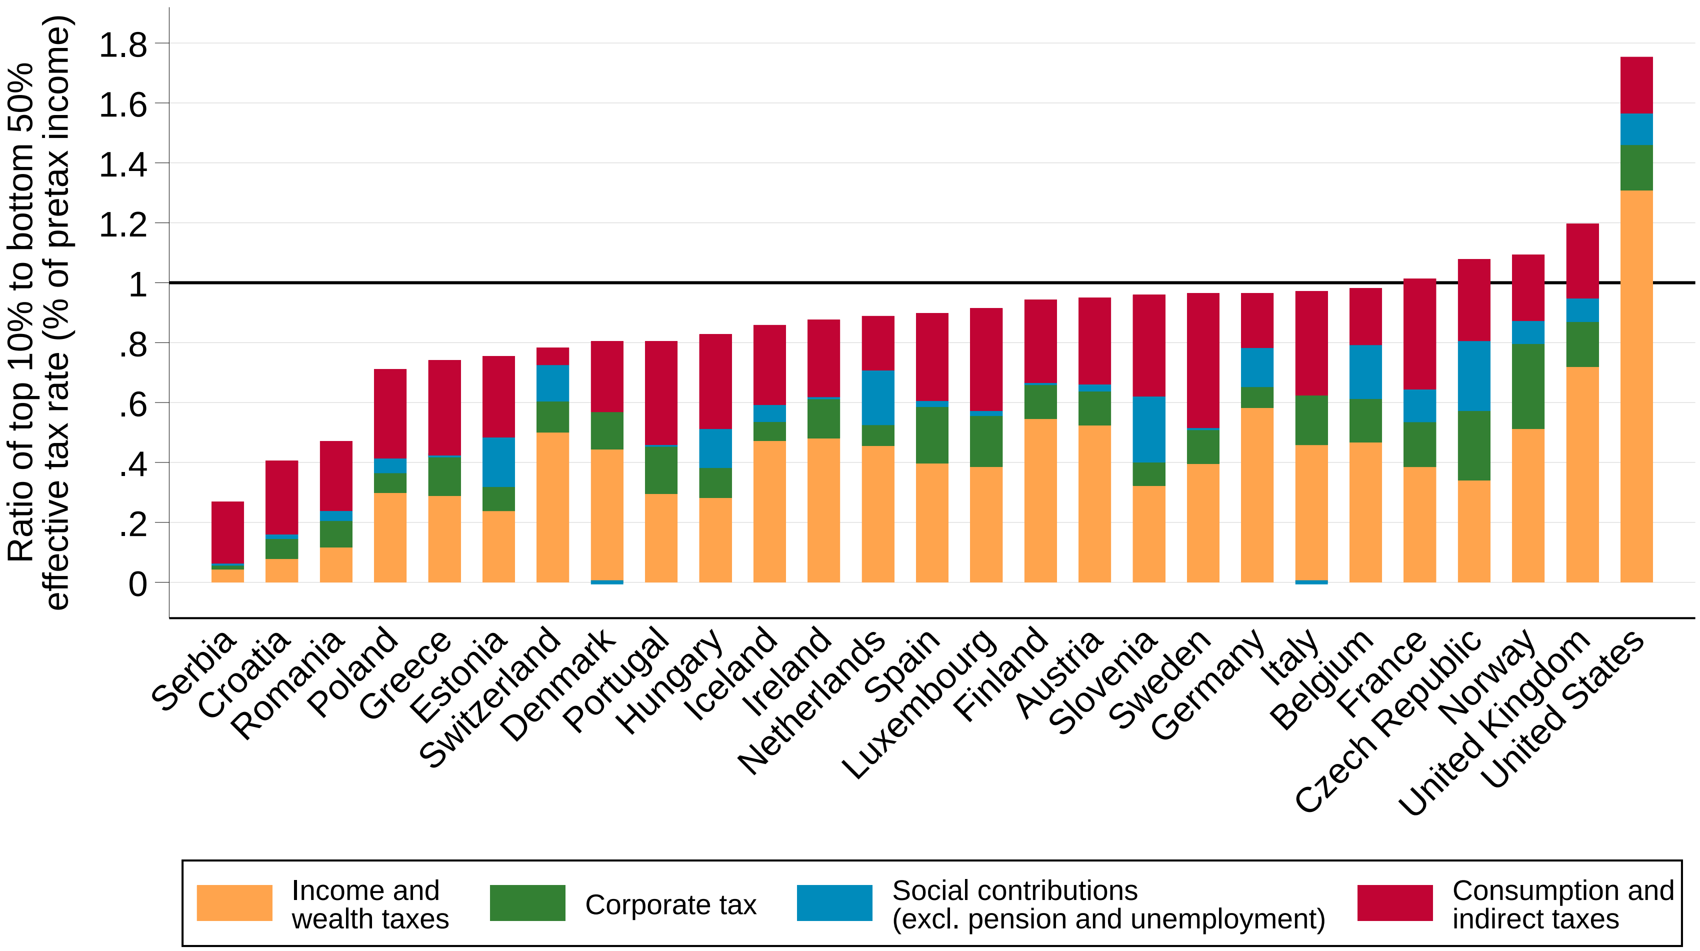

Source: Authors' computations combining surveys, tax data and national accounts for European countries; Piketty, Saez, and Zucman (2018) for the United States. *Notes:* Figures correspond to averages over the 2007–2017 period for European countries (population-weighted average of country-specific estimates in the case of European regions), and to 2017–2018 for the US. In panel (b), the composition of bars corresponds to the composition of taxes paid by the top 10%. The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. See online appendix table D.6 for the composition of European regions.

Figure 4.6: The Distribution of Transfers in Europe and the United States

(a) Total Transfers Received by Posttax Income Group (% of posttax income)

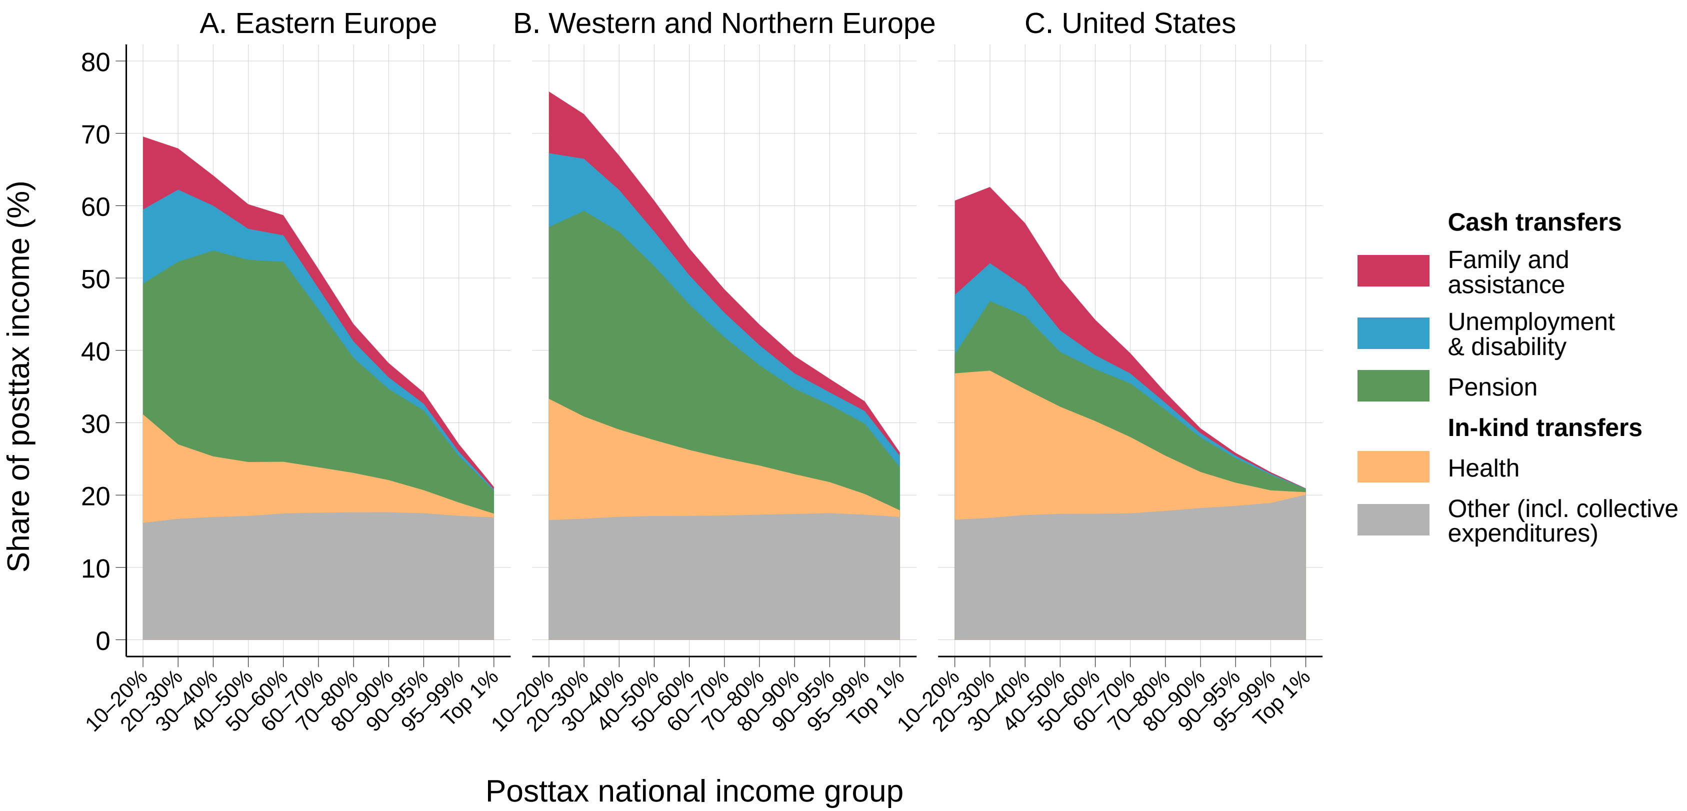

Source: Authors' computations combining surveys, tax data and national accounts for European countries; Saez and Zucman (2019) for the US. *Notes:* Figures correspond to averages over the period 2007–2017 for European countries (population-weighted average of country-specific estimates in the case of European regions), and to 2017–2018 for the US. The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. See online appendix table D.6 for the composition of European regions.

Figure 4.6: The Distribution of Transfers in Europe and the United States

(b) Total Transfers Received by the Bottom 50% by Country (% of national income)

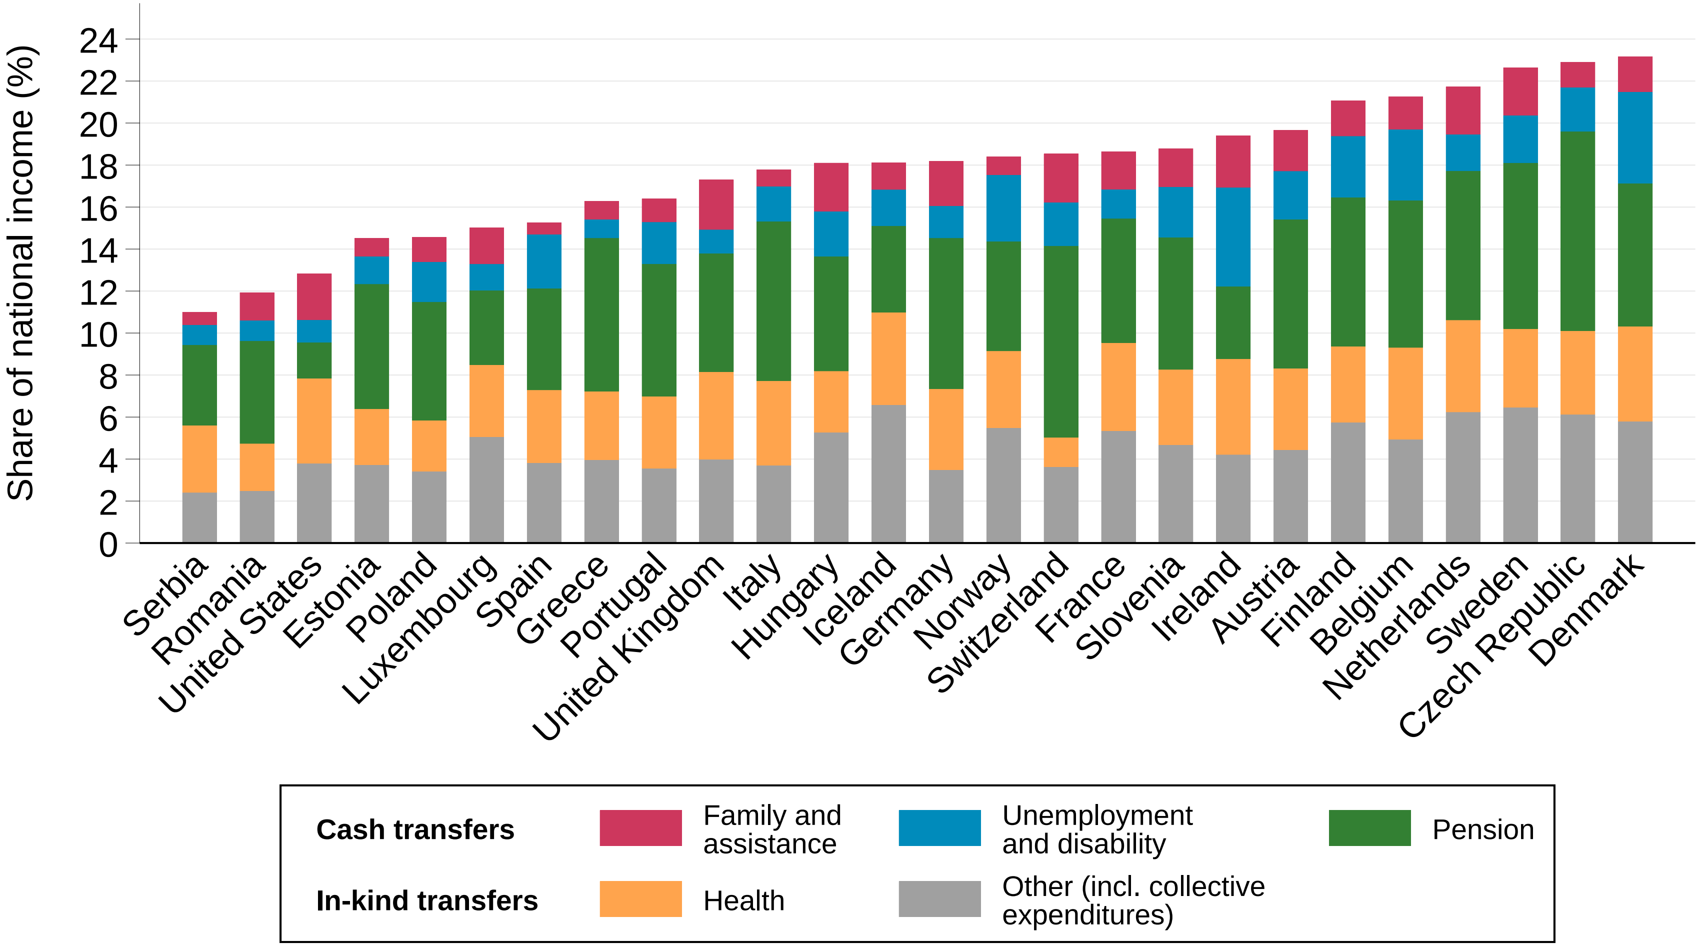

Source: Authors' computations combining surveys, tax data and national accounts for European countries; Piketty, Saez, and Zucman (2018) for the US. *Notes:* Figures correspond to averages over the period 2007–2017 for European countries (population-weighted average of country-specific estimates in the case of European regions), and to 2017–2018 for the US. The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. See online appendix table D.6 for the composition of European regions.

Figure 4.7: Net Redistribution in Europe and the United States

(a) Net Transfers Operated by the Tax-and-Transfer System Between Pretax Income Groups (% of National Income)

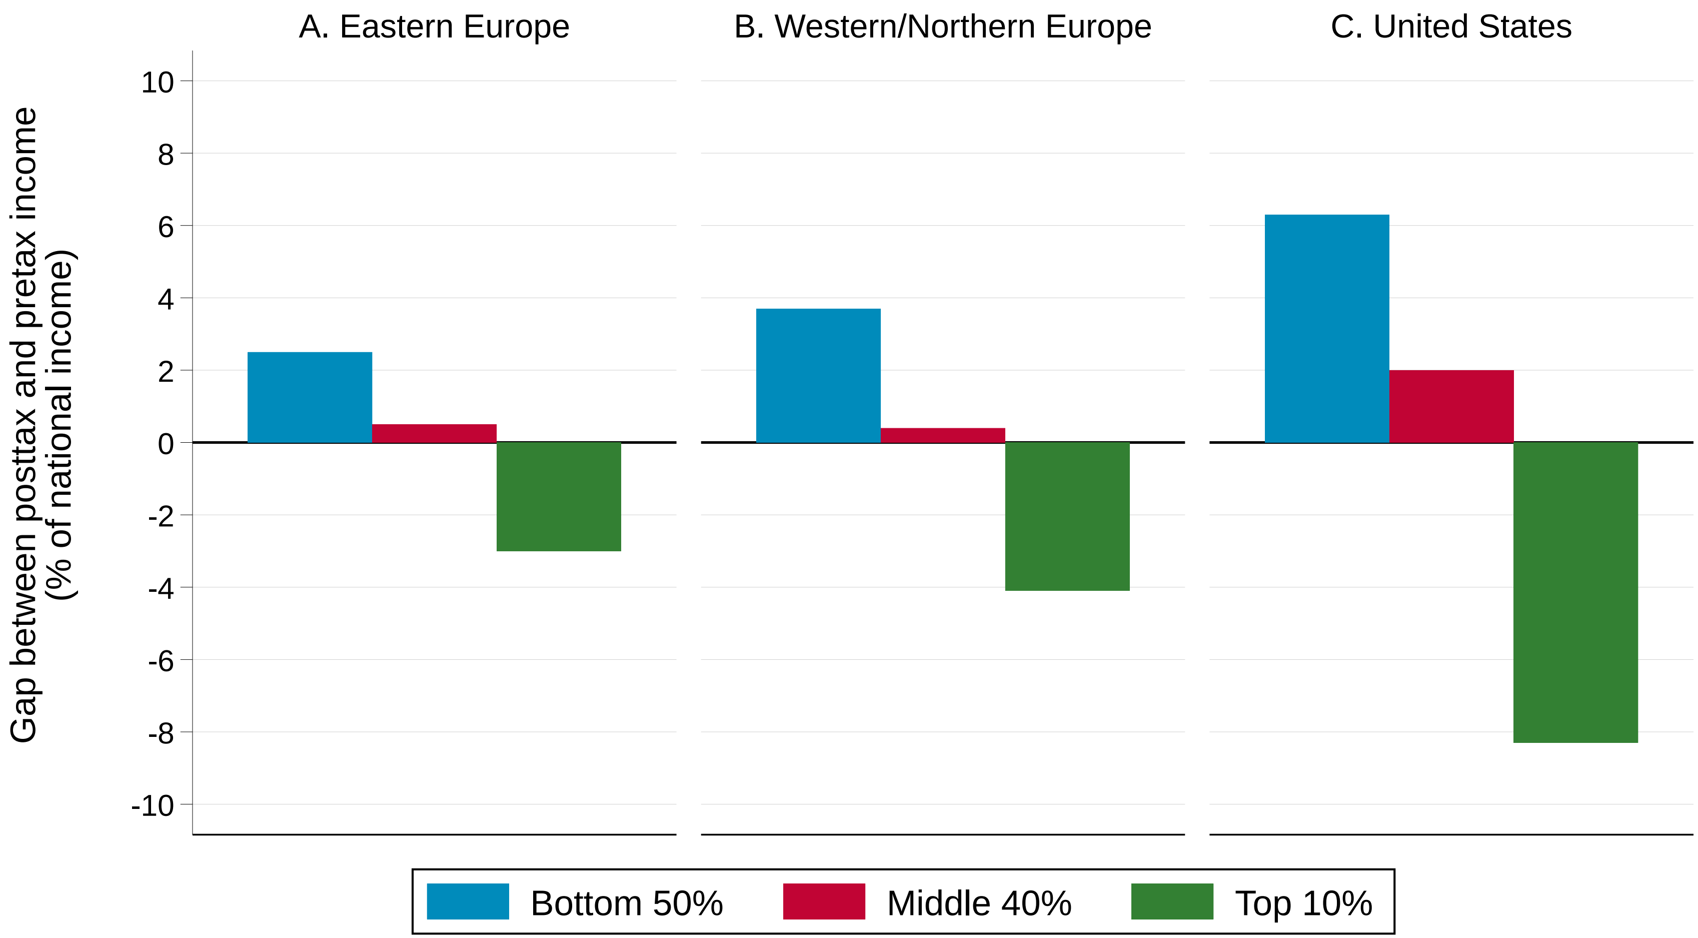

Source: Authors' computations combining surveys, tax data and national accounts for European countries; Piketty, Saez, and Zucman, 2018 for the US.
Notes: Panel (a) represents the net transfer received or paid by pretax income group in Eastern Europe, Western and Northern Europe, and the United States in 2017. Panel (b) represents the net transfer received by the bottom 50% by country, expressed as a share of national income, in 2017. The unit of observation is the adult individual aged 20. Income is split equally among spouses. See online appendix table D.6 for the composition of European regions.

Figure 4.7: Net Redistribution in Europe and the United States

(b) Net Transfer Received by the Bottom 50% by Country

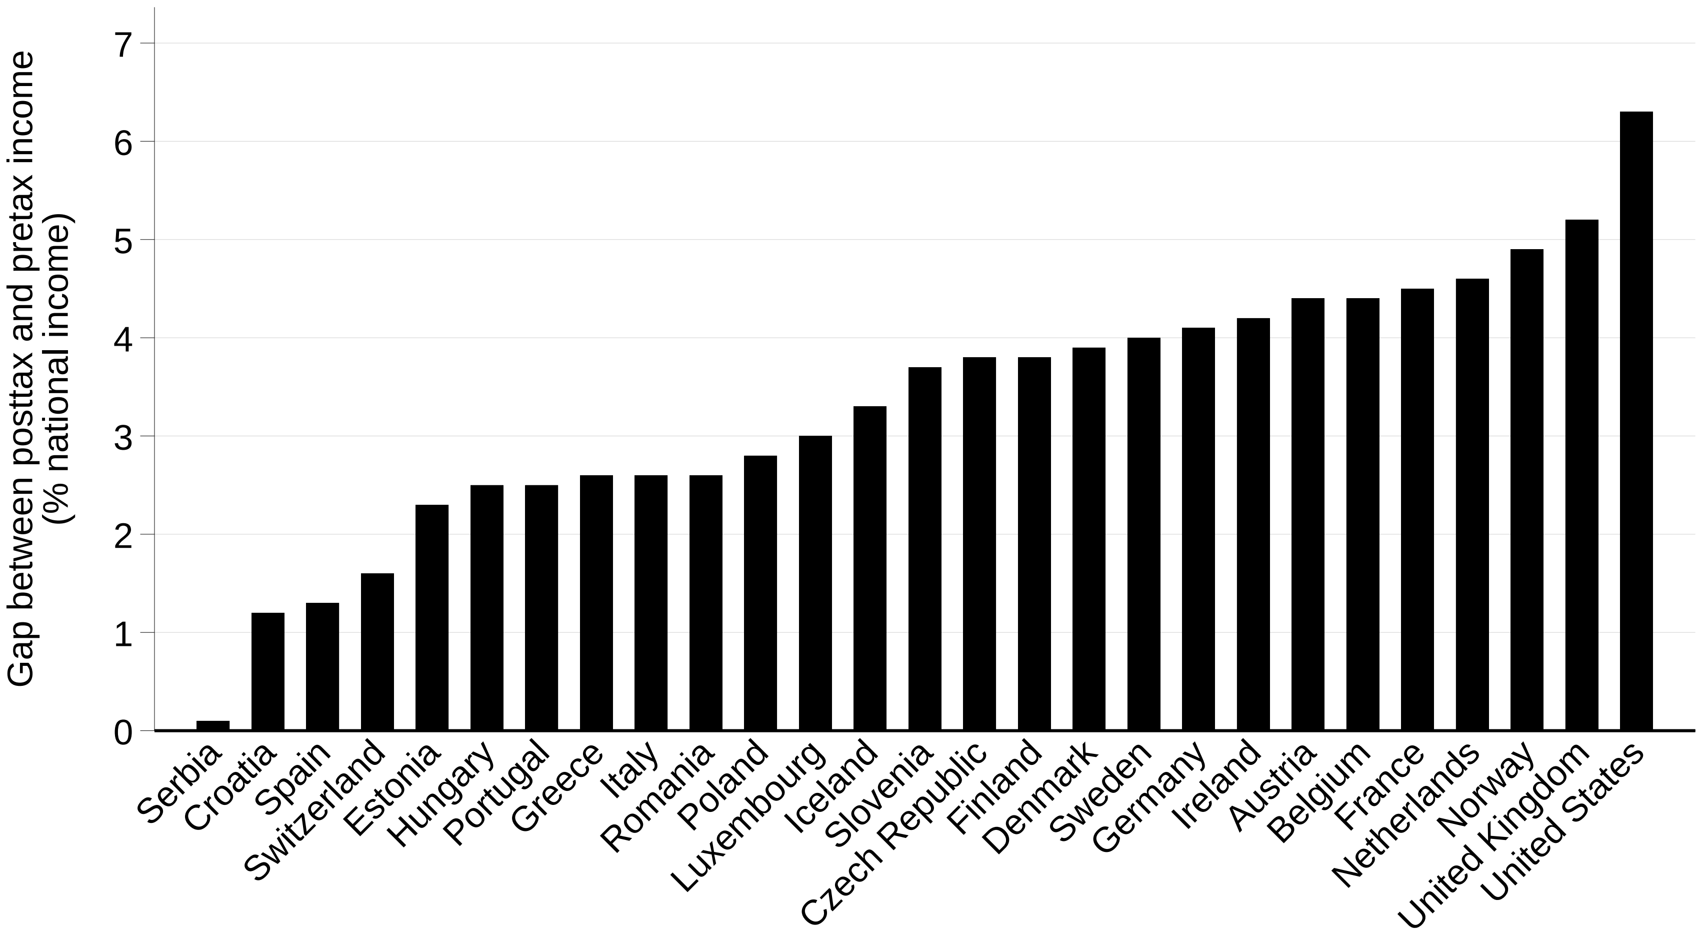

Source: Authors' computations combining surveys, tax data and national accounts for European countries; Piketty, Saez, and Zucman, 2018 for the US. *Notes:* Panel (a) represents the net transfer received or paid by pretax income group in Eastern Europe, Western and Northern Europe, and the United States in 2017. Panel (b) represents the net transfer received by the bottom 50% by country, expressed as a share of national income, in 2017. The unit of observation is the adult individual aged 20. Income is split equally among spouses. See online appendix table D.6 for the composition of European regions.

Figure 4.8: Pretax and posttax income inequality in Europe and the United States, 1980-2017

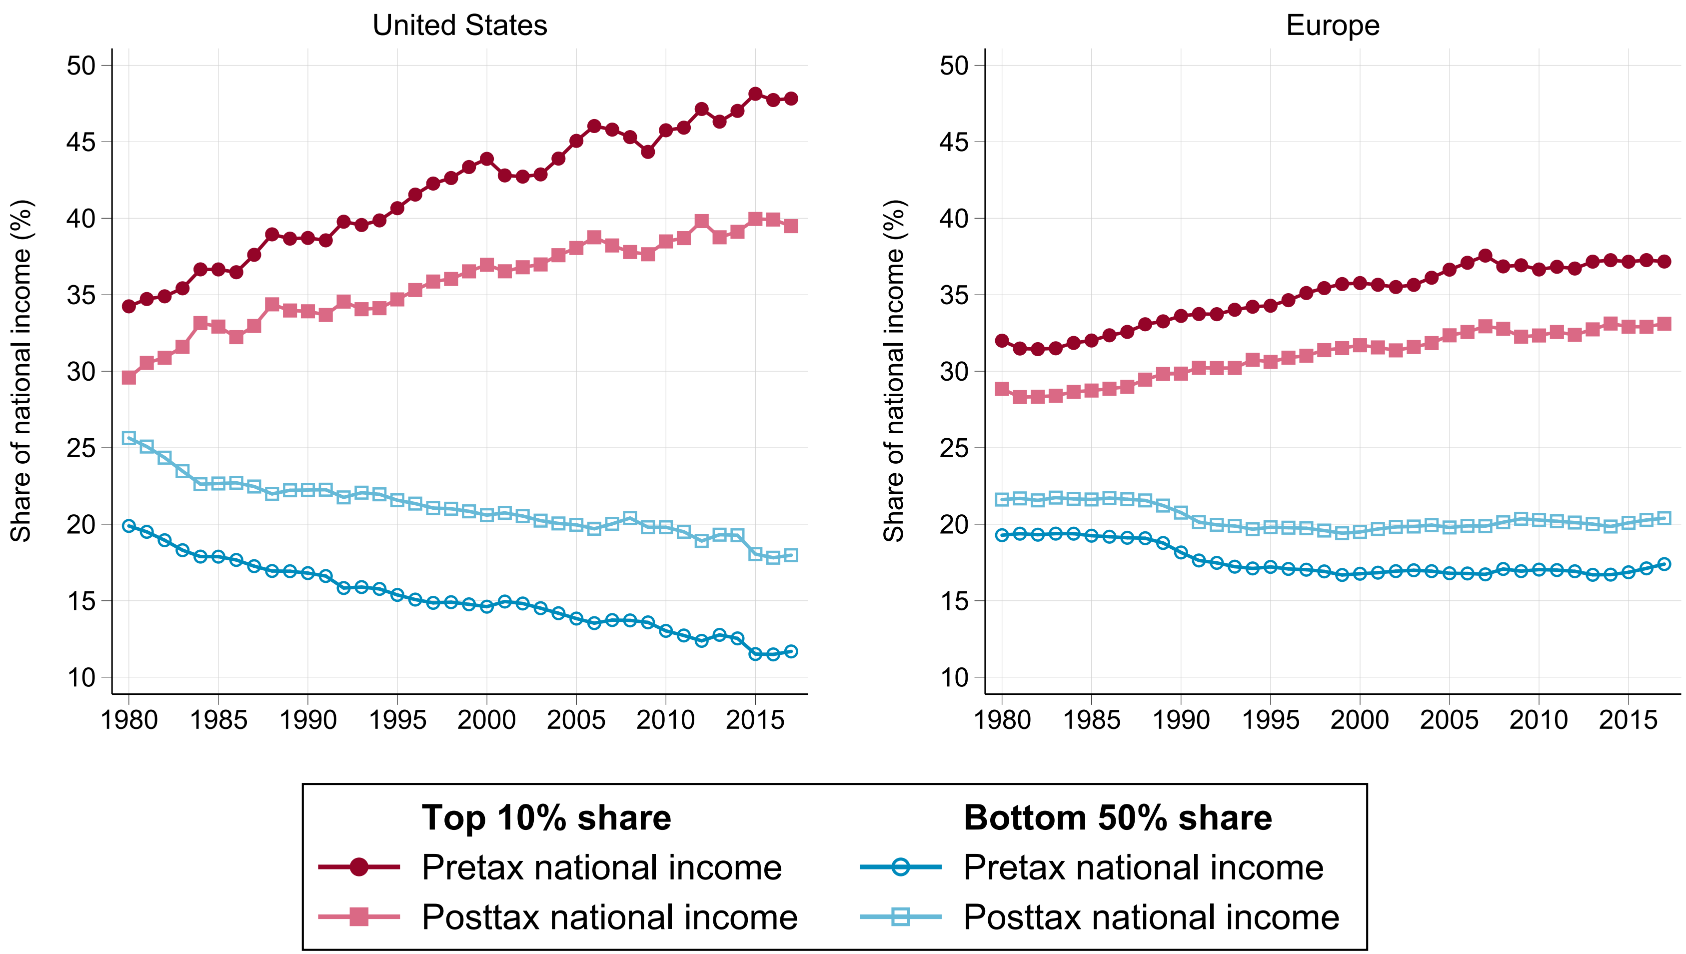

Source: Authors' computations combining surveys, tax data and national accounts for Europe and Piketty, Saez, and Zucman, 2018 for the US. *Notes:* The figure represents the evolution of the top 10% and bottom 50% shares in Europe and the United States in terms of pretax national income and posttax national income from 1980 to 2017. The unit of observation is the adult individual aged 20. Income is split equally among spouses. See online appendix table D.6 for the composition of European regions.

Figure 4.9: The Indirect Impact of Redistribution on Predistribution

(a) Evolution of Top 1% Share Under Different Top Marginal tax Rates

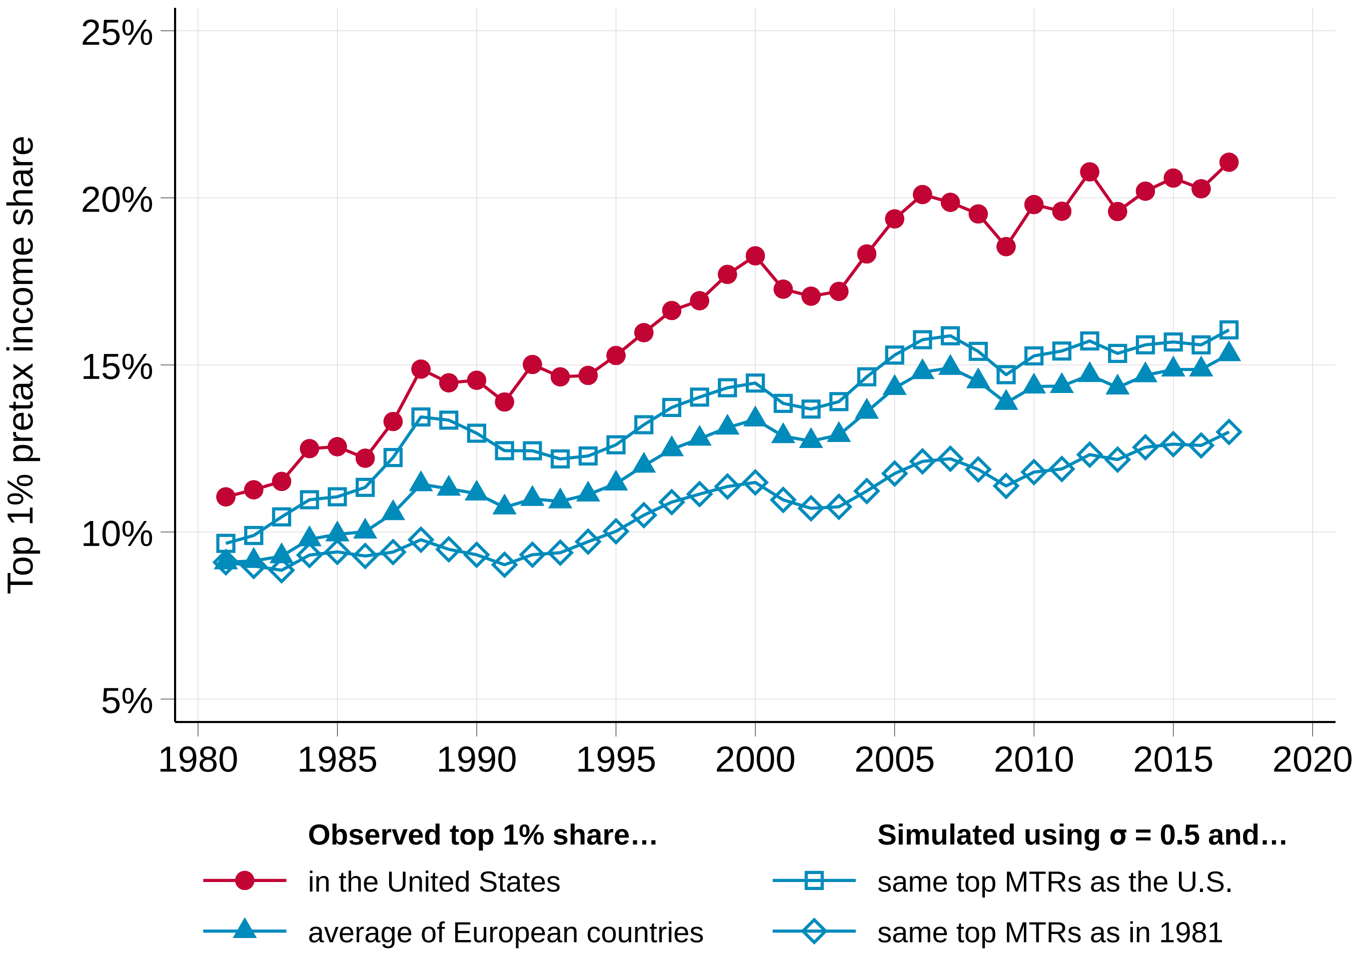

Source: Authors' computations combining surveys, tax data and national accounts for Europe; Piketty, Saez, and Zucman, 2018 for the US. Top marginal tax rates data extended from Kleven et al. (2020) using OECD data (see appendix D.1.6.2). *Notes:* European estimates refer to a population-weighted average of European countries with data available since 1981 (Austria, Belgium, Denmark, Finland, France, Germany, Greece, Ireland, Italy, Luxembourg, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, and United Kingdom). Counterfactual top 1% share estimated using the model $\Delta(\text{top 1\% share}) = (\Delta(1 - MTR))^{\sigma}$.

Figure 4.9: The Indirect Impact of Redistribution on Predistribution

(b) Redistribution to the Bottom 50% and Pretax Income Inequality

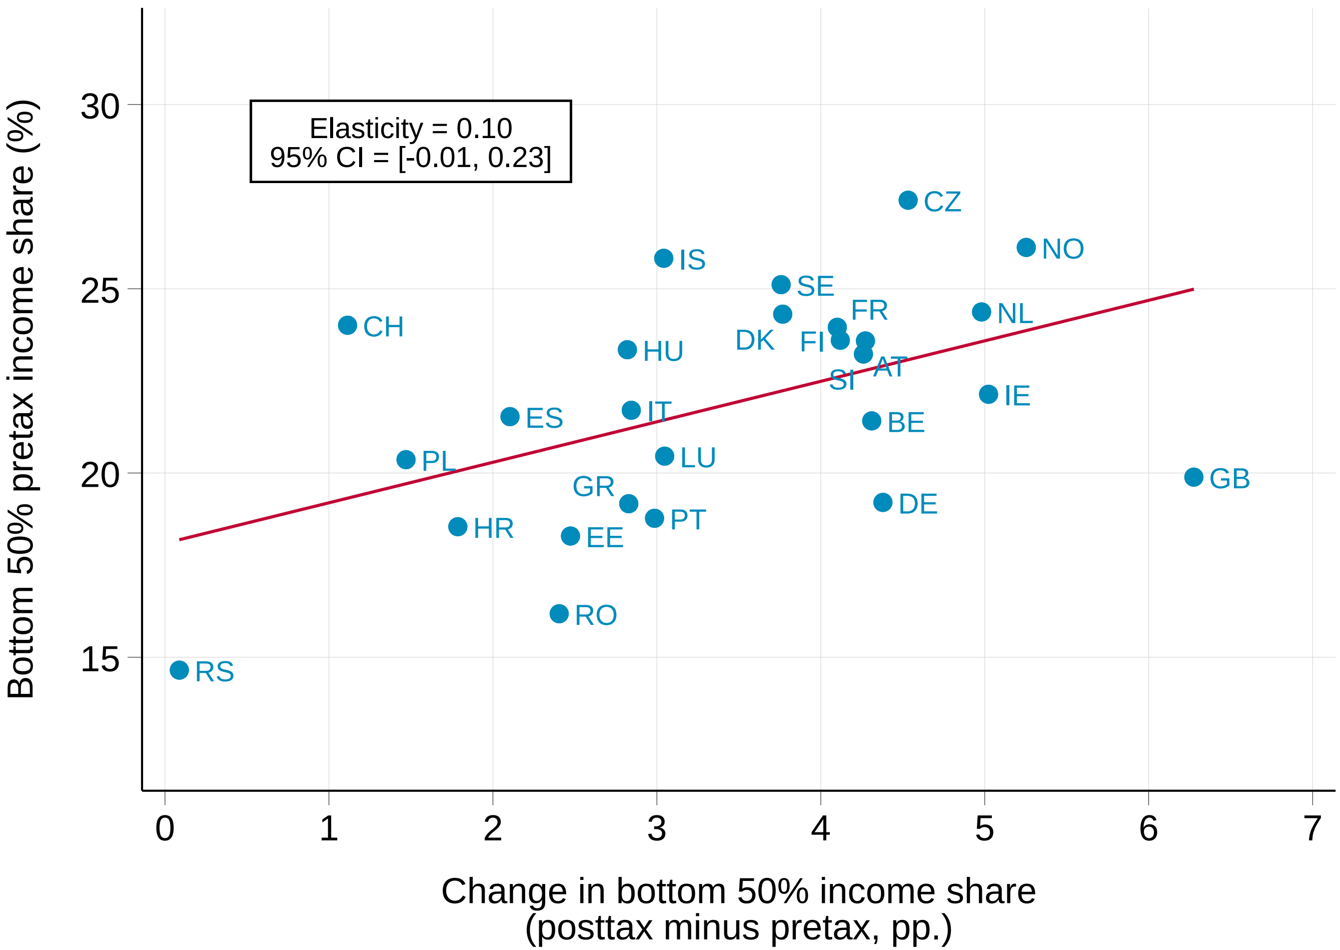

Source: Authors' computations combining surveys, tax data and national accounts. *Notes*: Averages over the 2007-2017 period. Elasticity refers to specification (2) in table D.5, appendix D.1.9.

Table 4.1: Methodology Used to Distribute Factor Income, Pretax Income, and Posttax Income in Europe

	Income concept	Source	Method	Share of income
	Factor national income			100%
(+)	Household primary income			79.2%
	Compensation of employees, mixed and property income	Survey + tax data	Observed	76.9%
	Net imputed housing rents	Survey + tax data	Observed	2.3%
(+)	Corporate primary income	National accounts	Proportional to equity ownership / wages and pension for equity held through pension funds	8.3%
(+)	Government primary income	National accounts	Proportional to pretax income	12.4%
	Pretax national income			100%
(+)	Factor national income			100%
(−)	Contributory social contributions	Survey + tax data	Observed/simulated	18.2%
(+)	Pension benefits	Survey + tax data	Observed	16.6%
(+)	Unemployment benefits	Survey + tax data	Observed	1.7%
	Posttax national income			100%
(+)	Pretax national income			100%
(−)	Taxes			29.3%
	Non-contributory social contributions	Survey + tax data	Observed/simulated	1.3%
	Direct taxes on income and wealth	Survey + tax data	Observed	11.1%
	Taxes on products	National accounts	Proportional to consumption	14%
	Corporate income tax	National accounts	Proportional to equity ownership / wages and pension for equity held through pension funds	3%
(+)	Transfers			30%
	Cash transfers	Survey + tax data	Observed	5.1%
	Public health expenditures	National accounts	Lump sum	7.7%
	Other public expenditures	National accounts	Proportional to posttax income	17.3%
(+)	Budget balance	National accounts	Proportional to posttax income	−0.7%

Notes: The table reports the methodology used to distribute the various components of factor national income, pretax national income, and posttax national income in European countries, together with the share of net national income each component typically represents (population-weighted average over all European countries over the 2010-2017 period).

Table 4.2: Methodology Used to Combine Survey, Tax, and National Accounts Data in Europe

Methodological Step	Detailed Steps	Sources and Coverage	Discussion / Impact
Step 1: Direct Measurement of Income Concepts in Survey Microdata.	Construction of pretax and posttax income variables.	EU-SILC (2004–2017); LIS (1980–2017); ECHP (1994–2001)	
	Imputation of social contributions.	Employee contributions (OECD, 2004–2017); Employer contributions (OECD, 2004–2005, EU-SILC, 2006–2017)	*Negligible impact* • Top 10% pretax income share decreases on average by 0.1 pp. after deduction of contributory social contributions.
Step 2: Harmonization of Survey Tabulations.	Collection and interpolation of survey tabulations, and harmonization using a machine learning algorithm.	World Income Inequality Database, PovcalNet, other survey data sources (1980–2017).	*Small impact* • 28% of cases: pretax income estimated from posttax income. • 1.5% of cases: income estimated from consumption.
Step 3: Combination of Surveys and Tax Data.	Calibration of survey microdata using top income shares series estimated from tax data.	World Inequality Database, various research articles, authors (1980-2017).	• Matching of income concepts and statistical units in surveys and tax data. • Calibration of surveys on tax data.
	Application of the correction to all survey distributions.		*Large impact* • Correction increases top 10% pretax income share by 2.3 pp. on average.
Step 4: Distribution of Unreported National Income Components.	Estimation and calibration of consumption, imputed rents, and stock ownership.	HFCS/WAS surveys for stock ownership; HBS for consumption; EU-SILC for imputed rents.	Top 10% pretax income earners account on average for: • 36% of stock. • 19% of consumption. • 16% of imputed rents.
	Missing incomes matched statistically to calibrated survey distributions.		*Moderate impact* • Retained earnings increase top 10% pretax income share by 1.0 pp. • Corporate tax increases top 10% pretax income share by 0.7 pp. • Imputed rents decrease top 10% pretax income share by 0.4 pp. • Taxes on products increase top 10% posttax income share by 1.5 pp. • Health spending decreases top 10% posttax income share by 1.5 pp.

Notes: The table reports the methodology used to combine survey, tax, and national accounts data to create European distributional national accounts, together with the impact of each methodological step on estimates of pretax and posttax income distributions. Numbers in the table refer to population-weighted averages across all countries and all years included in the database.

Table 4.3: The distribution of pretax income in Europe and the United States, 2017

	Eastern Europe		Northern Europe		Western Europe		United States	
	Average income	Income share	Average income	Income share	Average income	Income share		
Full population	€21,700	100%	€44,900	100%	€35,300	100%	€52,700	100%
Bottom 50%	€8,700	20.1%	€21,600	24.1%	€14,600	20.8%	€12,300	11.7%
Bottom 20%	€3,100	2.8%	€11,600	5.2%	€6,800	3.8%	€3,800	1.4%
Next 30%	€12,500	17.3%	€28,300	18.9%	€19,900	16.9%	€18,000	10.2%
Middle 40%	€24,100	44.3%	€50,600	45.1%	€39,200	44.5%	€53,300	40.5%
Top 10%	€77,300	35.6%	€138,000	30.8%	€123,000	34.8%	€252,000	47.8%
Top 1%	€261,000	12.0%	€395,000	8.8%	€384,000	10.9%	€1,110,000	21.1%
Top 0.1%	€892,000	4.1%	€1,140,000	2.5%	€1,230,000	3.5%	€5,190,000	9.8%
Top 0.01%	€3,060,000	1.4%	€3,290,000	0.7%	€3,970,000	1.1%	€23,830,000	4.5%
Top 0.001%	€10,490,000	0.5%	€9,490,000	0.2%	€12,840,000	0.4%	€92,020,000	1.7%

Source: Authors' computations combining surveys, tax data and national accounts for Europe; Piketty, Saez, and Zucman (2018) for the United States. *Notes:* The table shows the average annual real pretax income of various groups of the population in Western and Northern Europe, Eastern Europe and the United States in 2017. Incomes measured at purchasing power parity, €1 = $1.3. The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Table 4.4: Average annual pretax income growth in Europe and the United States, 1980-2017

	Eastern Europe		Northern Europe		Western Europe		United States	
	1980-2017	2007-2017	1980-2017	2007-2017	1980-2017	2007-2017	1980-2017	2007-2017
Full population	1.2%	2.2%	1.8%	0.7%	1.0%	0.0%	1.4%	0.4%
Bottom 50%	0.3%	2.8%	1.5%	0.3%	0.7%	0.0%	-0.1%	-1.2%
Bottom 20%	-1.3%	1.6%	1.2%	-0.5%	0.7%	-0.6%	-1.1%	-2.9%
Next 30%	0.6%	3.0%	1.5%	0.5%	0.7%	0.1%	0.1%	-0.9%
Middle 40%	1.1%	2.3%	1.7%	1.0%	0.8%	0.0%	1.0%	0.5%
Top 10%	2.2%	1.7%	2.4%	0.6%	1.4%	0.0%	2.3%	0.9%
Top 1%	3.8%	1.1%	3.2%	-0.6%	1.9%	-0.3%	3.3%	1.0%
Top 0.1%	5.7%	0.1%	4.3%	-1.9%	2.3%	-1.0%	4.2%	1.3%
Top 0.01%	7.7%	-1.0%	5.4%	-3.3%	2.6%	-1.7%	4.9%	1.4%
Top 0.001%	9.8%	-2.1%	6.6%	-4.6%	2.9%	-2.5%	5.4%	0.5%

Source: Authors' computations combining surveys, tax data and national accounts for Europe; Piketty, Saez, and Zucman (2018) for the United States. *Notes:* The table shows the average annual real pretax income growth of various groups of the population in Western and Northern Europe, Eastern Europe and the United States over the 1980-2017 and 2007-2018 periods. Incomes measured at purchasing power parity. The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Table 4.5: Predistribution versus redistribution in Europe and the United States: estimates of the top 10% and bottom 50% income shares using different concepts and data sources

	Top 10%			Bottom 50%		
	United States	Europe	Difference	United States	Europe	Difference
Surveys						
Factor income	35.9%	33.3%	+2.6 pp.	15.0%	12.1%	+2.9 pp.
Pretax income	33.1%	26.9%	+6.2 pp.	20.2%	25.9%	−5.7 pp.
Posttax income	28.9%	24.3%	+4.7 pp.	23.7%	29.2%	−5.5 pp.
Surveys + Tax data						
Factor income	43.5%	37.7%	+5.8 pp.	11.2%	8.5%	+2.7 pp.
Pretax income	41.7%	32.1%	+9.6 pp.	15.1%	21.8%	−6.7 pp.
Posttax income	35.9%	28.8%	+7.2 pp.	18.9%	24.7%	−5.8 pp.
DINA						
Factor income	46.0%	37.9%	+8.1 pp.	11.2%	12.5%	−1.4 pp.
Pretax income	45.7%	34.3%	+11.4 pp.	12.7%	21.4%	−8.6 pp.
Posttax income	37.1%	30.4%	+6.7 pp.	19.8%	24.9%	−5.0 pp.

Source: Authors' computations combining surveys, tax data and national accounts for Europe (population-weighted average). Survey-based estimates for the United States come from the Luxembourg Income Study. Surveys + Tax data and DINA estimates for the United States come from Piketty, Saez, and Zucman (2018). *Notes:* The table shows how estimates of top 10% and bottom 50% factor income, pretax income, and posttax income shares in Europe and the United States in 2017 vary depending on whether they are observed in household surveys, computed by combining surveys and tax data, or estimated using the distributional national accounts methodology.

Chapter 5

Who Benefits from Public Goods? Public Services and Inequality in Post-Apartheid South Africa

The standard concept used to track poverty and inequality within countries is posttax disposable income, defined as the sum of labor and capital incomes, plus cash transfers received, minus direct taxes paid. This concept has the advantage of capturing money that effectively ends up in households' bank accounts and can be used to purchase goods and services. Yet, it suffers from a key limitation: it entirely ignores in-kind transfers received by households in the form of services freely provided by the government. As a result, standard income distribution statistics still provide a very partial picture of the ways through which government redistribution reduces inequality. This is especially true in developing countries, where cash transfers tend to only represent a tiny fraction of public spending. Instead, much of redistribution involves transfers in public goods as diverse as education, healthcare, transport infrastructure, police services, and water supply.

This article makes a first attempt at incorporating detailed estimates of public goods provision in poverty and inequality statistics. The context is post-apartheid South Africa, which provides a particularly ideal case study to analyze government redistribution in kind. Since 1993, newly elected governments have massively invested in education, healthcare, and other public services, often with the explicit objective of reducing the extreme inequalities inherited from the apartheid regime of racial segregation. Drawing on various surveys, census microdata, and newly digitized budget reports, I build a comprehensive database covering the joint distribution of pretax incomes, taxes, cash transfers, and in-kind transfers every year from 1993 to 2019. Unlike existing studies, which focus on specific types of public services at a specific point in time, I allocate all public goods to individuals and account for changes in their progressivity over time. While these estimates still suffer from significant limitations, I view them as a useful first step towards more comprehensive measures of public service delivery, which can be refined as better data sources become available in the future.

I find that most government policies tend to be strongly progressive (less concentrated than income), but with large variations across functions of government. In 2019, the poorest 50% received about 77% of cash transfers, compared to 61% of education spending, 52% of public healthcare, 38% of local public goods, 38% of police services, and only 10% of transport expenditure. Overall, they benefit from about 43% of total government spending. This is less than their share in the South African population, but substantially higher than their share of pretax income, which falls below 3%. In other words, public services unambiguously reduce inequality.

Redistribution in kind is not only progressive; it is quantitatively substantial. In 2019, over 14% of South Africa's national income accrued to the bottom 50% in the form of in-kind transfers. This represented over three times total cash transfers received. As a result, incorporating public services in measures of posttax income significantly changes estimates of poverty and inequality. The share of income received by the bottom 50% is only 6.5% in terms of posttax disposable income. After accounting for public goods, it rises to almost 15%, corresponding to a threefold increase in the average income of the poorest half of the South African population.

Finally, I find that there has been a dramatic rise in government redistribution since the end of apartheid. From 1993 to 2019, total cash and in-kind transfers received by the bottom 50% expanded by over 50%, from 12% to 19% of national income. This transformation results from the combination of three factors. First, total government expenditure rose significantly, both in real terms and as a fraction of national income. Second, the share of public spending dedicated to the most progressive types of policies also increased, in particular education and healthcare. Third, there were significant improvements in the progressivity of most government policies, which increasingly accrued to low-income groups. This transformation has acted as a major driver of inclusive growth: accounting for in-kind transfers almost doubles the growth rate of the real income of the bottom 50% since 1993.

This article connects to a growing literature attempting to bridge conceptual gaps between surveys and national accounts in the measurement of inequality. Piketty, Saez, and Zucman (2018) estimate Distributional National Accounts (DINA) for the United States every year since 1913, yielding distributional statistics that are consistent with macroeconomic growth rates. A number of studies following this framework have been conducted on other countries since then, including detailed studies of government redistribution covering Europe (Blanchet, Chancel, and Gethin, 2022), France (Bozio et al., 2018), and Latin America (De Rosa, Flores, and Morgan, 2022b).[1] The main limitation of these studies is that they do not attempt to estimate who benefits from public services; instead, they typically assume that all in-kind government expenditure is distributed proportionally to posttax disposable income.[2] To the best of my knowledge, this paper is the first to build detailed estimates of the progressivity of all public services and its evolution over time. I show that public goods act as a major driver of inequality reduction, which calls for the necessity of better incorporating them in measures of poverty and inequality (see Gethin (2023b) for a preliminary attempt at expanding this analysis to the study of the world distribution of income).

This paper also contributes to extending our knowledge of who benefits from public services. Some studies have attempted to estimate the distributional incidence of specific public goods, in particular health and education (see Goldman, Woolard, and Jellema (2020) in the context of South Africa).[3] I depart from these studies in two ways. First, I follow the DINA methodology and allocate in-kind transfers in a framework that is rooted in the national accounts. This contrasts with the existing literature, which tends to scale down public services to match aggregates observed in surveys, in ways that tend to be variable and inconsistent with macroeconomic statistics. Second, I focus on *all* public goods, while existing studies typically restrict themselves to specific types of public spending. In doing so, I directly follow some of the principles outlined in O'Dea and Preston (2010), who provide a set of potential guidelines to estimate the distributional incidence of all government policies.

[1]See also Germain et al. (2021), Bruil et al. (2022), and Jestl and List (2022) on France, the Netherlands, and Austria, respectively. See Chancel et al. (2022b) for a presentation of other studies following the DINA methodology.

[2]Piketty, Saez, and Zucman (2018) allocate all non-health in-kind transfers proportionally to posttax disposable income. Blanchet, Chancel, and Gethin (2022) consider two polar scenarios, one in which in-kind transfers are distributed proportionally to posttax disposable income, and one in which they are received as a lump sum. De Rosa, Flores, and Morgan (2022b) allocate education and health spending based on fiscal incidence studies, as in this paper, and all other government spending proportionally to posttax disposable income.

[3]See for instance Benhenda (2019), Lustig (2018), Paulus, Sutherland, and Tsakloglou (2010), Verbist, Förster, and Vaalavuo (2012), and Wagstaff et al. (2014) on education and health, Aaberge and Atkinson (2010) and Aaberge et al. (2019) on local government services, and Mladenka and Hill (1978) on police expenditure.

The rest of the article is organized as follows. Section 5.2 describes the methodology used to estimate the distributional incidence of all taxes and transfers to individuals. Section 5.3 presents the results. Section 5.4 concludes.

5.1 Conceptual Framework

Measuring the progressivity of public goods is conceptually and empirically challenging, given that their ultimate beneficiaries cannot always be unambiguously identified. I rely on three simple allocation principles to estimate the distributional incidence of public goods, which directly follow the existing literature (e.g., Lustig, 2018; O'Dea and Preston, 2010; Piketty, Saez, and Zucman, 2018). First, public services accrue to individuals based on who receives them at a given point in time. Second, public goods benefit households based on the price they would have to pay to benefit from this service if it was not provided as a public good. Third, public goods are valued in a way that is consistent with the national accounts, that is, at cost of provision (potentially adjusted for government productivity). These three principles are necessary to ensure conceptual consistency with both standard poverty and inequality statistics and macroeconomic growth rates reported in the national accounts.

5.1.1 Cash Flow Principle

First, I distribute public goods to individuals benefiting from their consumption at a given point in time. For instance, education spending is distributed to households who send their children to school, health spending is distributed to individuals using more intensively the public healthcare system, and public transport expenditure is distributed to individuals relying more extensively on public transportation. This ensures that public goods are valued in a way that is conceptually consistent with standard fiscal incidence analysis, which focuses on the incidence of taxes and transfers over a given period. Put differently, public services are allocated in the same way as they would theoretically be if households were to receive a cash transfer at time t and immediately use it to buy the corresponding service on a private market.

5.1.2 Equivalent Pricing Principle

Second, public goods accrue to households based on the price that they *would have to pay* for the public service, rather than the price they *would be willing to pay.*

This ensures again that cash transfers and public goods are valued in a conceptually comparable way: if the household was to receive cash instead of the public good, it would have to pay the market price of the corresponding service to benefit from it, not the maximum value it would be willing to pay.

Standard poverty statistics focus on consumption and do not attempt to estimate the individual welfare value of each good bought by each household. Accordingly, I distribute public goods based on who benefits more from them, not based on who puts lower or greater value on each type of service. For example, the welfare perspective would imply that high-income households may be willing to pay significantly more for police services, as they may have more to lose from burglaries and other property crimes than low-income households.[4] This would call for allocating police services proportionally to wealth. In contrast, assuming that the cost of solving a crime is the same across income groups, the income perspective implies that detective services should benefit households proportionally to the number of crimes that they experienced. Consistency with standard consumption aggregates thus requires allocating police expenditure proportionally to reported crimes, not wealth, because a household suffering a crime would have to pay the price of solving the crime, not the price of its entire wealth, if it was to buy the same service from a private investigator.

5.1.3 Consistency With National Accounts Principle

Third, I follow the principles of Distributional National Accounts (Piketty, Saez, and Zucman, 2018), which aim to provide a set of guidelines to allocate the totality of net national income to individuals. The guiding principle of this methodology is to close the gap between micro and macro estimates of the income distribution: inequality statistics should be consistent with macroeconomic growth rates. In the context of this article, this implies valuing public services at cost of provision, simply because this is what national statistical institutes do when constructing estimates of GDP growth.

Departing from cost of provision would imply revising estimates of GDP growth, specifically "deflating" public services in a way that is different from the average good consumed. This represents a particularly challenging task, which probably explains why national accountants have preferred to use cost of provision as a reasonable assumption until now. That being said, I investigate at the end of this paper the

[4]Notice however that low-income households tend to suffer from significantly higher violent crime, including murders, whose cost may be valued at an equally, if not higher level than property crime.

robustness of my results to adjusting for the quality of public services received over time and throughout the income distribution (controlling for cost of provision), drawing on recent work in which I benchmark public spending to educational and health outcomes (Gethin, 2023b).

5.2 Data and Methodology

This section outlines the methodology used to estimate the distributional incidence of public services in South Africa. I start from the microfile constructed in Chatterjee, Czajka, and Gethin (2023), which provides comprehensive information on the distribution of income in South Africa since 1993. I then present in turn the methods used to allocate public spending on education, healthcare, local government services, housing, transport, other economic affairs, public order and safety, and other functions of government.

5.2.1 Distribution of Factor Income and Pretax Income

The starting point of this paper is a microfile covering the distribution of factor and pretax incomes at the individual level, every year from 1993 to 2019. Chatterjee, Czajka, and Gethin (2023) construct this file by combining surveys, tax data, and national accounts to allocate the entirety of net national income to individuals. The database records information on the composition of the household, sociodemographic characteristics of each household member, and income received from different sources. It also covers household wealth, expenditure by type of good, different types of taxes paid, and cash transfers received from the government.

The bulk of my analysis focuses on allocating in-kind transfers to individuals. In broad strokes, I first identify different functions and policies of the South African government, and collect new budget data on spending in each of them. I then combine different microdata sources to estimate who benefits from spending on these functions alongside the income distribution. Finally, I incorporate these estimates into the microfile, so as to get a comprehensive picture of their joint distribution with respect to income since 1993.

Figure 5.1 plots the level and composition of general government expenditure in South Africa since 1993, expressed as a share of national income. Table 5.1 provides an overview of the microdata and macrodata sources used to allocate government expenditure to individuals, as well as the corresponding distributional assumptions. I now turn to presenting these sources and methods for each function of government in more detail.

5.2.2 Education (9% of NNI)

I distribute public education proportionally to use intensity of the public education system, accounting for differences in public education spending by province and level of education.

The 1996, 2001, 2011, and 2016 censuses, as well as the 2007 community survey, provide information on school attendance, current grade, and the type of school attended (private/public).[5] I match them with provincial budget data to allocate expenditure to individuals following public education by five levels: early childhood development, primary education, secondary education, tertiary education, and adult basic education. I assume that each individual within a given province-function cell receives the same transfer, equal to the per-student expenditure on this function.[6] Finally, I proportionally rescale education transfers received so as to match total national education expenditure.

5.2.3 Health (5.1% of NNI)

I distribute health expenditure proportionally to use intensity of the public healthcare system, accounting for differences in public health spending by province and type of institution (clinics versus hospitals).

To measure intensity of use of public healthcare, I combine two different sets of surveys: the October Household Surveys (1995-1996) and the General Household Surveys (2004-2019). Both surveys have collected data on (1) whether household members have used the public healthcare system in the past month (2) the type of institution (private/public) usually visited by household members and (3) whether the institution usually visited is a clinic or a hospital.

First, I combine these variables to generate cells of public hospital and public clinic use intensity by pretax income ventile, race, and province. Second, I interpolate and extrapolate these cells so as to cover the entire 1993-2019 period. Third, I merge these cells with the harmonized DINA microfile, so as to a get a measure of the

[5] The 1996 census microfile unfortunately does not provide information on type of institution, so I assume that all individuals attending school benefit from public education expenditure.

[6] Administrative data on the distribution of education expenditure in South Africa shows that this is a reasonable assumption. See for instance Motala and Carel (2019), table 4.3, who show that personnel expenditure per learner is highly equalized across school quintiles (which are defined by the living standards of the community around the school and used by the South African government to allocate resources).

intensity of use of public clinics and public hospitals by group. Fourth, I combine the microdata with estimates of public health spending by province and by institution (clinics/hospitals), which I have collected from the archives of the National Treasury. Finally, I proportionally rescale health transfers received so as to match total national health expenditure.

5.2.4 Local Government Services (9.6% of NNI)

The local government sector is large in South Africa, and has been growing in the past decades thanks to increasing transfers from the central government. Municipalities are in charge of providing households with electricity, water, sanitation, waste removal, and other basic services, some of which are distributed free of charge to poor households in the form of "free basic services" since 2001. They also deliver a number of local public services related to public safety, healthcare, administration, and other public goods.[7]

To distribute local government spending from 1993 to 2019, I first collect new historical budget data from a number of sources. I then allocate them to individuals by matching budget series with survey and census microdata. Finally, I incorporate these estimates into the DINA microfile.

Harmonization of Local Government Budget Data I combine data on local government expenditure from four sources. The first one are tables A2 published by the National Treasury in the context of the Medium Term Revenue and Expenditure Framework (MTREF), which cover operating expenditure by function in each of South Africa's municipalities from 2006 to 2019. The second one are tables A1 from the same source, which specifically cover expenditure made by municipalities for the provision of Free Basic Services. The third one are tables published in the 2008 Local Government Budgets and Expenditure Review, which cover total expenditure by municipality from 2003 to 2006. Finally, I digitize data on consolidated municipal operating expenditure by district council over the 1996-1999 period from the 2000 edition of the Local Government Budgets and Expenditure Reviews. As above, I interpolate and extrapolate total expenditure as a share of national income, so as to cover every year from 1993 to 2019.

The last step of harmonization consists in incorporating local government expenditure series into general government expenditure. Indeed, part of local expenditure is

[7]See appendix figure E.27, which plots the level and composition of total local government expenditure from 2001 to 2019.

already accounted for in consolidated budgets: the part that is financed by transfers to municipalities from the central government. While most of these transfers are included in the "Community Development" function, some transfers, especially in recent years, consist in conditional grants that appear indirectly in other consolidated government functions. Unfortunately, exact estimates of which fraction of each function is spent through municipalities are not available. For simplicity, I assume that all transfers are spent through either Community Development or Water Supply. That is, I completely remove these two expenditure items from the national budget and replace them by the series of total local government spending estimated above.

Allocation of Local Government Expenditure I allocate local government expenditure to individuals by matching these newly constructed budget series with census microdata. To do so, I rely on the 1996, 2001, 2011, and 2016 censuses, as well as the 2007 community survey. I incorporate municipal expenditure into the microfile in three steps.

First, I match budget and census data at the municipal level in each census, recoding municipality names and codes when necessary. I do so for both local/metro and district municipalities, so as to distribute these two layers of local government one after the other. For 1996, I match individuals at the district level, given that I have no information on expenditure at a lower geographical level.

Second, I distribute local/metro and district municipal expenditure on a lump basis to individuals, assuming that all adults living in a given municipality benefit from the same transfer.[8]

Finally, I incorporate these estimates into the DINA microfile. First, I aggregate municipal expenditure by pretax income ventile, race, and province. Second, I interpolate and extrapolate average expenditure received by each cell so as to cover the entire 1993-2019 period. Third, I match these cells with the DINA microfile. Finally, I rescale proportionally the average transfer received by individuals in each year so as to perfectly match yearly aggregate municipal operating expenditure at the national level.

[8] One can compare this strategy to a more complex one, distributing separately free basic services, water expenditure, electricity expenditure, and other expenditure separately in each municipality, based on households' access to these different types of services. I also compare the results to those obtained by allocating municipal expenditure at the district level instead of the municipal level, using either census (2001-2011) or NIDS (2008-2016) data. I find that these three alternative strategies yield virtually identical results in terms of the distribution of municipal expenditure by income, race, and province.

5.2.5 Housing (0.9% of NNI)

Housing development expenditure in South Africa mainly corresponds to the Reconstruction and Development Programme (RDP), a large national housing programme initiated in 1994 that allows low-income households to acquire a house built and provided by the government. To distribute public housing expenditure, I rely on the General Household Survey (2002-2019), which has consistently asked survey respondents whether any household member received a government housing subsidy to obtain this dwelling or any other dwelling. First, I aggregate the share of individuals who declared having received a subsidy by cells of consumption decile and province. Second, I merge these cells with the DINA microfile. Third, I proportionally rescale each cell so as to match total government housing expenditure.

5.2.6 Transport (2.5% of NNI)

Expenditure on transport services can be separated into two components: public transport expenditure and expenditure on transport infrastructure.

Public transport expenditure corresponds to expenditure on the public transport system, including buses and commuter rail, and represents about 20-25% of total transport expenditure.[9] I distribute total public transport expenditure proportionally to household expenditure on public trains and buses, which is directly reported in household income and expenditure surveys (COICOP codes 07311110 to 07321210).

Infrastructure expenditure corresponds to expenditure on roads, railroads, and other infrastructure used by households, firms, and publicly owned vehicles to transport goods and people. Accordingly, I split the benefit received by individuals into a household part, a firm part, and a government part. First, I use input-output tables provided by the OECD and the South African statistical institute to derive an estimate of what fractions of transport infrastructure are used by the household, corporate, and government sectors. Second, I distribute each of these fractions to their ultimate beneficiaries.

For the household sector, I assume that infrastructure expenditure benefits individuals proportionally to their fuel consumption, as reported in income surveys. This amounts to assuming that households disproportionately using their car, for instance, benefit from a greater government transfer on transport infrastructure.

[9]Unfortunately, budget reports only provide a decomposition of transport expenditure into public transport and infrastructure from 2007 to 2019, so I assume that this decomposition was the same throughout the 1993-2006 period.

For the corporate sector (mainly corresponding to the transport of goods), I use input-output tables to derive measures of the "transport intensity" of household consumption by expenditure category (COICOP). I then allocate infrastructure expenditure proportionally to this intensity measure, observed at the household level. This amounts to assuming that households disproportionately consuming goods that need to be transported (for instance, goods produced in another country) indirectly benefit from public expenditure on the roads used to transport these goods.

Finally, I distribute the public sector fraction proportionally to the public transport transfer received by each individual, as estimated above. This amounts to assuming that individuals using public transport not only directly benefit from using public vehicles, but also indirectly benefit from the fact that these public vehicles use roads or railways provided by the government.

5.2.7 Other Economic Affairs (2.8% of NNI)

Expenditure on other economic affairs mainly include subsidies to specific economic sectors and other policies dedicated to supporting production. The South African budget decomposes it into six functions: General economic, commercial, and labour affairs; Agriculture, forestry, fishing and hunting; Fuel and energy; Mining, manufacturing and construction; Communication; and Recreation and Culture.

As in the case of taxes on international trade and transport infrastructure expenditure, I allocate expenditure on these different sectors proportionally to their consumption intensity.[10] First, I use input-output tables to estimate the indirect consumption intensity of these different sectors by COICOP category. I then allocate total government expenditure on these sectors proportionally to the total intensity of household consumption expenditure in this sector. This amounts to assuming, for instance, that households consuming goods that require more energy to be produced benefit proportionally more from energy subsidies provided to firms.

5.2.8 Public Order and Safety (3.7% of NNI)

Expenditure on public order and safety includes police services, law courts, and prisons. Police services are in turn broken down by the South African government into "Visible policing," which aims to "Enable police stations to institute and preserve safety

[10]The exception is general economic, commercial, and labour affairs, for which no sector can be clearly identified. I distribute this component proportionally to the total transfer received in other economic affairs.

and security," and "Detective services" and "Crime Intelligence," whose objective is to investigate and solve crimes (South African Treasury, 2022).

Accordingly, I split public order and safety expenditure into two functions: an "insurance" function equal to visible policing, and a "use" function equal to the sum of detective services, crime intelligence, law courts, and prisons. The insurance function relates to crime prevention and security provision, which primarily benefit households through police presence and responsiveness to emergencies. In contrast, the use function corresponds to the set of services that are provided to households once crimes are already committed, from police investigations to justice and incarceration.

I distribute the insurance function of public order and safety proportionally to police presence by income group.[11] To do so, I rely on Victims of Crime surveys (1998-2017), which have consistently asked individuals about the frequency at which they see a police officer in uniform or on duty in their area or neighborhood.[12] First, I aggregate police presence intensity by income decile in Victims of Crime surveys. Second, I interpolate and extrapolate between years to cover the full 1993-2019 period. Finally, I match these cells with the DINA microfile, and distribute expenditure on the insurance function proportionally to police presence intensity in each cell.

I distribute the use function of public order and safety proportionally to crimes reported to the police by income group. This corresponds to the fact that individuals directly benefit from government services, in the form of police investigations and law courts, when they are victims of a crime (O'Dea and Preston, 2010). I rely again on Victims of Crime surveys, which record all crimes suffered by survey respondents in the past year. First, I aggregate total crimes reported by the police by cells of income decile. Second, I interpolate and extrapolate between years to cover the full 1993-2019 period. Finally, I match these cells with the DINA microfile, and distribute expenditure on the use function proportionally to the number of reported crimes in each cell.

[11] This strategy can be motivated by the literature on the crime-reducing effects of police manpower and police presence on crime (Chalfin and McCrary, 2017; Di Tella and Schargrodsky, 2004; Levitt, 1997).

[12] Respondents are given a choice between "At least once a day," "At least once a week," "At least once a month," "Less than once a month," or "Never." I combine these options to derive a proxy for the number of days per year a respondent sees a police officer (coding each option as 365, 52, 12, 6, and 0, respectively).

5.2.9 Other Expenditure (3.7% of NNI)

Other government expenditure in South Africa consists in spending on general public services (2.5%) and defense (1.2% of NNI). I consider two polar scenarios: one in which they are distributed on a lump sum basis, and one in which they are distributed proportionally to posttax disposable income (that is, in an extremely unequal way).

5.3 Public Goods and Inequality in Post-Apartheid South Africa

I now turn to documenting trends in the distribution of public services in South Africa. Section 5.3.1 presents the main results on the progressivity of in-kind transfers and how it varies by government function. Section 5.3.2 studies the incidence of public goods on inequality and the distribution of growth. Section 5.3.3 investigates the robustness of the results to adjusting government transfers received for public sector productivity.

5.3.1 Who Benefits From Public Goods?

How large is government redistribution in South Africa, and how has it evolved since 1993? Table 5.2 provides a first answer to this question by documenting the share of total government expenditure, the share of national income, and the average transfer received by the poorest 50% by government function in 1993 and 2019. Three main conclusions can be drawn.

5.3.1.1 In-Kind Transfers Are Large and Strongly Progressive

Following the standard approach to the analysis of tax or transfer incidence, let us define a transfer as *relatively progressive* if it reduces inequality, that is, is less concentrated than income. Based on this definition, government redistribution in South Africa appears to be very strongly progressive. In 2019, the poorest half of the population received only 2.7% of pretax income, but over 40% of public spending. Every single category of government spending was relatively progressive, both in 1993 and 2019. In other words, government transfers systematically reduce inequality.

In-kind transfers also appear to be very large. In 2019, total transfers received by the bottom 50% amounted to about $3200 at PPP after excluding social protection, corresponding to about five times their average pretax income. Spending on education alone represented twice their average income, and was about 35% higher than total social protection expenditure. Overall, in-kind transfers accounted for almost 80% of total expenditure accruing to the bottom 50% in 2019.

5.3.1.2 Progressivity Varies Significantly by Function of Government

Beyond this general result, there are major differences in progressivity across types of government transfers. In particular, only social protection, education, health, and housing expenditure are *absolutely progressive*, that is, received in greater proportion by the poor than by the rich.

Social protection stands out as the most progressive spending category, with over three quarters of expenditure accruing to the bottom 50% in 2019. This is consistent with the fact that most cash grants are explicitly targeted towards the poor. Indeed, the bulk of social protection expenditure in South Africa consists in the old age grant, the child support grant, and the disability grant, all of which are means-tested (see appendix figures E.6).

Public education and healthcare also appear to be slightly progressive in South Africa, for two main reasons. First, both services are used more extensively by poor households, who overwhelmingly send their children to public schools and rely on public clinics for healthcare, while top earners primarily rely on private alternatives. Second, they are also used more intensively by low-income groups, who tend to have more children and visit health institutions more frequently. As a result, the bottom 50% received about 61% of public education spending and 52% of public health spending in 2019.[13]

Public housing expenditure is also absolutely progressive, with 58% of spending received by the bottom 50%. Indeed, low-income households are much more likely to live in a state-subsidized dwelling, although some middle- and high-income households do benefit from public housing too. In 2018-2019, as much as 22% of the poorest 50% individuals declared having received some assistance from the government to obtain a dwelling (see appendix figure E.34).

Local government spending is regressive in absolute terms: the poorest 50% receive less than 40% of total local government expenditure. This is a direct consequence

[13]In 2016, the average number of children attending public schools exceeded 2 among the poorest 50%, compared to less than 0.4 among the top 10% (see appendix figure E.15). Over 30% of children within the top 10% attend private schools, compared to less than 10% of children within the bottom 50% (see appendix figure E.16). The same differences are visible for public healthcare. The share of individuals having visited a public health institution in the past three months strongly declines with income (see appendix figure E.22). Over half of South Africans within the top income quintile are covered by private health insurance and rely primarily on private healthcare, compared to less than 5% of those in the bottom quintile (see appendix figures E.23 and E.24).

of richer municipalities having access to greater resources through larger local tax collections, which enables them to spend more on public services. In 2019, the top 10% thus benefited from nearly PPP $1700 per capita in local government expenditure, compared to 800$for the bottom 10% (see appendix figure E.31).

Public order and safety expenditure is absolutely regressive too. This is true of spending on both visible policing and law enforcement. Richer households are significantly more likely to suffer from crimes and report them to the police (see appendix figure E.39). They also tend to live in neighborhoods with greater police presence (see appendix figure E.40). The correlation between these indicators and income is mild, however, leading top income groups to benefit from only slightly greater transfers. In 2019, the bottom 50% received just below 40% of public order and safety expenditure.

Transport expenditure and expenditure on economic affairs are the most regressive of all functions of government (although there are still progressive in relative terms). Only about a fifth of public transport expenditure accrues to the bottom 50%, mainly because public transport is more intensively used by middle-class households in richer urban areas.[14] Infrastructure scarcely benefits the poor at all, with only 7% of expenditure accruing to the bottom 50% in 2019. This results from the fact that richer households use private vehicles to a much greater extent, and also benefit from higher consumption of transported goods.

Putting all cash and in-kind transfers together, how does total public spending received vary alongside the income distribution? As shown in figure 5.2, which plots the share of national income received by pretax income decile in 2019, all income groups benefit from large transfers. Overall, the distribution of public spending ends up being slightly absolutely regressive: the top 10% received 6% of national income in 2019, while other deciles each received approximately 4% of NNI each. Cash transfers are strongly concentrated among the poorest 30%, while public spending on transport and other economic affairs is heavily concentrated among the top 20%. Other public goods are more broadly shared.

5.3.1.3 Government Redistribution Has Increased

There has been a dramatic rise in redistribution since the end of apartheid. Between 1993 and 2019, the average transfer received by the bottom 50% grew by over 100%, from about $2000 to almost $4500 at purchasing power parity (see table 5.2).

[14]See appendix figures E.45 and E.46: the average number of bus and train trips realized per week is highest among the second, third, and fourth quintiles, and is lowest among the top 20%.

This increase was the outcome of three factors. First, average national income per capita expanded by 37% in real terms (see appendix figure E.1) Second, general government expenditure grew as a share of national income, from about 37% to 43% of NNI. This rise was concentrated in the functions of government that are the most equally distributed. Social protection spending rose from 3% to 5% of NNI, education spending from 8% to 9%, and health spending from 4% to 5% (see table 5.1). Third, the progressivity of transfers increased: from 1993 to 2019, the share of total government expenditure accruing to the bottom 50% expanded from 32% to 43%. The rise of progressivity happened for virtually all functions of government and can be accounted for by a number of factors, including improved access to education and healthcare and significantly lower spatial inequalities in the provision of local public goods.[15] The outcome of these three forces has been a large increase in the real value of transfers received by the bottom 50%, which extends to all categories of public spending.

Figure 5.3 plots the level and composition of transfers received by the bottom 50% since 1993, expressed as a share of national income, providing a more detailed perspective on the rise of government transfers received by low-income households. As evident from the figure, the expansion of redistribution has been primarily driven by cash transfers, education, healthcare, and local government spending. Public order and safety, transport, and other economic affairs only represent a minor fraction of public services received by low-income households. Furthermore, a standard analysis of government redistribution focusing only on cash grants would miss an enormous part of transfers received by the poor, both in terms of levels and trends. In 2019, cash transfers represented less than a quarter of total public expenditure accruing to the poorest 50% individuals in South Africa.

5.3.2 The Incidence of Public Goods on the Distribution of Growth

I now turn to analyzing the distributional impact of redistribution on income and growth. I derive two main conclusions: in-kind transfers strongly reduce inequality, and they have significantly contributed to income growth at the bottom since 1993.

[15] See for instance appendix figure E.15: from 1996 to 2016, the average number of children attending public schools remained the same within the bottom 50%, while it was divided by more than two within the top 10%. Figures E.29 and E.30 show that there has been a dramatic convergence of local government spending across municipalities, as the rise of overall expenditure was strongly driven by the catch-up of low-spending municipalities.

5.3.2.1 Public Goods Substantially Reduce Inequality

Given that in-kind transfers are large and progressive, it naturally follows that they strongly contribute to reducing inequality. To get a sense of their redistributive power, consider table 5.3, which provides information on the contemporary distribution of income in South Africa before and after taxes and transfers.

Pretax income is extremely unequally distributed.[16] In 2019, the top 0.1% captured over 8% of pretax income, more than three times the share of income received by the bottom 50% as a whole. The top 10% income share stood at almost 69% (compared to about 47% in the US: see Piketty, Saez, and Zucman, 2018). Meanwhile, the average pretax income of the poorest quintile was not far from an exact zero; this may look striking but should not come as a surprise, in a country where the unemployment rate has regularly exceeded 25% since the end of apartheid. Together, these figures confirm South Africa's position as one of the most unequal countries in the world (see Chatterjee, Czajka, and Gethin, 2022).

Columns 4 and 5 remove direct taxes and add cash transfers to reach posttax disposable income. Cash transfers are large and progressive in South Africa, while direct taxes are mostly borne by the top 10%. As a result, moving from pretax to posttax disposable income increases the average income of the poorest half of the population by over 50%. Both the middle 40% and the top 10% see their average incomes decrease, due to higher direct taxes paid than cash transfers received.

Columns 6 and 7 remove all remaining taxes (including indirect taxes and the corporate income tax) and add in-kind transfers to reach posttax national income. Inequality is substantially lower in terms of posttax national income than in terms of posttax disposable income. In the benchmark scenario, moving from posttax disposable to posttax national income multiplies the average income of the bottom 50% by more than 3. As a result, the bottom 50% income share more than doubles, from about 6.5% pretax to almost 15% posttax. The bottom 20% average income increases from about \$400 to \$2000; in other words, 80% of the final income of the poorest quintile consists in income received in the form of in-kind transfers. South Africa's poorest individuals thus receive very little cash income, but they benefit from much more significant indirect transfers received in the form of free education, healthcare, electricity, water supply, public housing, and police services.

[16] Pretax income is the sum of all primary incomes received by individuals, before accounting for taxes and transfers, but after accounting for the operation of the pension and unemployment systems, which are very small in South Africa: see Chatterjee, Czajka, and Gethin (2023).

5.3.2.2 Public Goods Account for a Large Share of Low-Income Households' Income Gains

Not only do in-kind transfers reduce inequality, they have contributed to significantly increasing incomes at the bottom of the distribution since the end of apartheid. Figure 5.4 represents the evolution of the bottom 50% average income from 1993 to 2019, before and after adding different layers of government transfers to the analysis. Average factor incomes grew by 14% over this period, which is only about 40% of the average national income growth rate. Adding pensions and unemployment benefits leaves this picture unchanged, since these transfers are very small and almost entirely received by top-income groups.[17]

Accounting for other cash transfers pushes the bottom 50% real income growth rate to 53%. This effect is almost entirely due to the adoption of the child support grant in 1998, which was followed by a gradual rise in take-up rates until today.[18] Accounting for in-kind social protection further increases this figure to 67%, mainly due to the development of various provincial social development programs.

Education, health, and local government spending account for the bulk of in-kind government redistribution. Adding education transfers pushes the average income of the bottom 50% from about $1,500 to over $2,500, and its growth rate from 67% to 80%. Health transfers add about another $500, and local government and housing expenditure bring the bottom 50% average income to over $4,000. Finally, accounting for spending on public order and safety, transport, and other economic affairs increases it to $4,500. The total growth rate of the bottom 50% after all transfers reaches 95%, which is nearly 7 times that of factor income and 80% higher than that of pretax income plus cash transfers. Notice that this figure mechanically underestimates the true contribution of in-kind transfers to bottom real income growth, since it adds them after market income and cash transfers in the analysis.

Figure 5.5 reproduces the same analysis for the poverty headcount ratio at $6 per

[17] See appendix figures E.3, E.4, and E.5. Private pension contributions and benefits are almost exclusively paid and received by the top 30%, with contributions being approximately equal to benefits within each income decile. The unemployment insurance fund is extremely small and has run large surpluses, with total unemployment benefits paid falling below 0.1% of national income in 2019.

[18] See appendix figure E.6, which shows that the bulk of the rise of social protection expenditure since 1993 has been driven by the child support grant. The growth of cash transfers cannot be explained by increases in the value of grants allocated per beneficiary: in fact, their real monthly value has stagnated or even decreased (see appendix figure E.8). Instead, there has been a significant increase in coverage: by 2019, about 10% of the adult population received an old age grant from the government, and almost two-thirds of all South African children benefited from a child support grant (see appendix figure E.9).

day in 2011 purchasing power parity dollars, corresponding to the poverty threshold generally used in middle-income countries. Absolute poverty declined by about 14% in terms of pretax income, 21% when adding cash transfers, and 81% when adding all cash and in-kind transfers. In 1993, government transfers lowered the poverty headcount ratio from about 75% to 55%. By 2019, it reduced it by over 50 percentage points, from 65% to 10%. Public goods have thus played a key role in the historical reduction of poverty in post-apartheid South Africa.

5.3.3 Accounting for Public Sector Productivity

The above analysis focuses on the distribution of public services evaluated at cost of provision. A natural concern is that the quality of services received, controlling for cost, may vary over time and throughout the income distribution and may also differ between the private sector and the public sector. Following Gethin (2023b), I thus investigate the sensitivity of my results to adjusting in-kind transfers using two productivity parameters: aggregate productivity, which refers to the overall efficiency of the South African government at providing public goods, and heterogeneous productivity, which captures inequality in the quality of services received by income group.

5.3.3.1 Aggregate Productivity

To account for potential inefficiencies in public goods provision in South Africa compared to other countries in the world, I rely on estimates by Gethin (2023b), who combines a number of data sources to estimate levels and trends in public education and public healthcare productivity around the world since 1980. First, data is collected on public education and health spending, as well as on educational and health outcomes. Second, an efficient frontier is estimated, corresponding to the maximum educational or health outcome observed for a given level of expenditure. In other words, the country-year performing best at a given cost of provision is attributed the highest productivity. Finally, other country-years are attributed a productivity score based on their distance to the frontier. For instance, a government performing two times worse than the government at the frontier for the corresponding cost level is attributed a productivity of 0.5. The resulting aggregate productivity indicator ranges from 0 to 1, with 0 corresponding to a completely useless government, and 1 corresponding to the most efficient government observed (which is equivalent to making the conservative assumption that this government is just as efficient as the private sector).

Drawing on Gethin (2023b), I construct measures of productivity-adjusted public goods received in South Africa by multiplying cost of provision by the corresponding indicators. With this approach, education and healthcare productivity are found to be particularly low in South Africa, about 0.4-0.5 over the period considered. This correction thus amounts to reducing transfers received by as much as 50-60%. In the absence of better information, I also multiply other in-kind transfers received by the average of the two productivity indicators in each year.

5.3.3.2 Heterogeneous Productivity

Another potential issue is that the quality of public services may vary by income group, even after accounting for differences in spending received. For instance, teachers teaching in poorer areas may be less qualified, even if they are paid the same as teachers in richer areas. Accounting for such "heterogeneous productivity" is an extremely challenging task, as it would ideally imply deriving monetary indicators of how the value added of each type of government service varies by income group.

In the absence of better information, I combine a number of data sources to get a sense of how important variations in the quality of public services alongside the distribution of income might be in South Africa. Table 5.4 reports data on how service delivery varies by income quintile, based on a battery of indicators covering three complementary dimensions: subjective perceptions of public services, objective indicators of government output, and distance to public institutions. Two main conclusions can be drawn from these figures.

First, there is evidence that poorer households benefit from public services of lower quality in most dimensions of government intervention. With the exception of public schools, local public institutions are always perceived as being of significantly lower quality by the bottom income quintile than by the rest of the population. Low-income households are also characterized by public school teachers with lower knowledge of mathematics, more frequent water and electricity interruptions, and public housing of lower quality. They tend to live further away from public institutions, in particular police stations and hospitals (but not public schools and public transport services).

Second, despite these differences, inequalities in access to public services remain relatively small. In particular, the data point to clear bounds on the maximum potential gap between top and bottom income groups. There is not a single indicator on which the bottom 20% scores less than 70% of the sample mean. The ratio exceeds 0.85 for most measures, in particular when it comes to subjective perceptions. There are some indicators, such as the success of the police at making an arrest after the household reported a crime, on which the government does not appear to perform better for the rich than for the poor.

It is also important to stress that some of these indicators do not account for the fact that higher quality might be the result of greater resources, which are already captured in estimates of progressivity. For instance, estimates of school teachers' knowledge of mathematics are based on the entire South African population, including private and fee-paying schools, which are disproportionately concentrated in the top quintile and benefit from substantial private resources (Venkat and Spaull, 2015). Similarly, quality differentials in local government services largely reflect the major differences in resources that exist between richer and poorer municipalities (see section 5.3), which are not accounted for here either. Correcting for differential resources would thus lead to revising inequalities in access to public services downwards. In this context, estimates of heterogeneous productivity derived from these indicators should be taken as upper bounds on the degree of heterogeneous productivity by income group. In the results that follow, I aggregate these different subjective and objective measures by government function, and correct the transfer received by each income group accordingly.

5.3.3.3 Results

Because adjusting for productivity implies strongly reducing the value of in-kind transfers, it naturally follows that their redistributive impact is lower than when they are valued at cost of provision. Appendix table E.1 reports estimates of the distribution of pretax and posttax income in South Africa after adjusting in-kind transfers for aggregate and heterogeneous productivity. Moving from posttax disposable income to posttax national income now increases the bottom 50% income share from 6.5% to 9%. Public goods thus end up having a lower redistributive power, but still very significant, almost as large as that of cash transfers.

Appendix figure E.2 reproduces figure 5.4 after adjusting in-kind transfers for public sector productivity. The average income of the bottom 50% now reaches about $3,000 after accounting for all transfers, compared to about $4,500 in figure 5.4. By this measure, public sector inefficiencies reduce the average income of the bottom 50% by a third. However, adjusting figures for productivity does not alter the trend: the bottom 50% average income rose by 53% before accounting for in-kind transfers, compared to 92% after doing so.

In summary, large public sector inefficiencies in South Africa could imply that public goods do not reduce poverty and inequality as much as an analysis relying on cost of provision would suggest. However, even under conservative assumptions on the productivity of the South African government, they still end up having large effects on the income distribution and have been major drivers of inclusive growth since the end of apartheid. Given the large rise of spending on public goods and progress made in terms of access to public services among low-income households, the idea that estimates of government redistribution should be restricted to cash transfers is difficult to sustain.

5.4 Conclusion

Public services remain largely absent from standard poverty and inequality statistics, despite representing the bulk of government redistribution in low- and middle-income countries. Focusing on the case of post-apartheid South Africa, this article argued that accounting for the distribution of public goods is essential to accurately track poverty and the distribution of growth. Not only do public services powerfully reduce inequality; they have become increasingly progressive, contributing to generating large real income gains for low-income households since the end of apartheid. The takeaway is that standard distributional analysis focusing exclusively on cash transfers is likely to miss crucial information on the evolution of the living standards of the poorest individuals. Developing tools to regularly and accurately track the distribution of public goods should be seen as an imperative for policy and future research.

These results call for future research in at least two directions. First, the fact that wages have remained so low at the bottom of the income distribution bears the question of how useful these public services have been at generating pretax income growth. Arguably, they have strongly contributed to improving the quality of life of South African citizens in a number of dimensions, from greater access to electricity and water to better education and health outcomes. At the same time, the fact that better access to these services does not seem to have enabled a fairer distribution of employment and pretax incomes is puzzling. One possibility is that pretax income inequality would have grown even faster in the absence of the rise of government redistribution. Another possibility is that of "redistribution without inclusion," whereby the legacy of apartheid and spatial segregation continues to weigh so heavily in access to economic opportunities that public services have failed to truly enable low-income households to escape the poverty trap.

Another natural avenue for future research is to better understand how low-income households actually value public services, not only in comparison to cash transfers, but also in comparison to one another. Evidence on this question remains extremely scarce, although some surveys suggest that individuals do strongly value public goods, in particular health and education (Khemani, Habyarimana, and Nooruddin, 2019; Thesmar and Landier, 2022). Answering this question would require new methods and data sources that go beyond those mobilized in this article.

Bibliography

Aaberge, R. and A. B. Atkinson (2010). “Top incomes in Norway”. In: *Top incomes: a global perspective*. Ed. by A. B. Atkinson and Thomas Piketty. Oxford University Press. Chap. 9, pp. 448–481.

Aaberge, Rolf et al. (2019). “Local governments, in-kind transfers, and economic inequality”. In: *Journal of Public Economics* 180.1, pp. 1–22.

Benhenda, Asma (2019). “Teaching Staff Characteristics and Spending per Student in French Disadvantaged Schools”. Working Paper.

Blanchet, Thomas, Lucas Chancel, and Amory Gethin (2022). “Why Is Europe More Equal Than the United States?” In: *American Economic Journal: Applied Economics* 14.4, pp. 1–40.

Bozio, Antoine et al. (2018). “Inequality and Redistribution in France, 1990-2018: Evidence from Post-Tax Distributional National Accounts (DINA)”. In: *WID.working paper series 2018/10*.

Bruil, Arjan et al. (2022). “Inequality and Redistribution in the Netherlands”. Working Paper.

Chalfin, Aaron and Justin McCrary (2017). “Criminal Deterrence: A Review of the Literature”. In: *Journal of Economic Literature* 55.1, pp. 5–48.

Chancel, Lucas et al. (2022b). *World Inequality Report 2022*. Belknap Press of Harvard University Press.

Chatterjee, Aroop, Léo Czajka, and Amory Gethin (2022). “Wealth Inequality in South Africa, 1993-2017”. In: *World Bank Economic Review* 36.1, pp. 19–36.

– (2023). “Redistribution without Inclusion? Inequality in South Africa Since the End of Apartheid”. Working Paper.

De Rosa, Mauricio, Ignacio Flores, and Marc Morgan (2022b). “More Unequal or Not as Rich? Revisiting the Latin American Exception”. Stone Center for Inequality Studies Working Paper Series.

Di Tella, Rafael and Ernesto Schargrodsky (2004). "Do Police Reduce Crime? Estimates Using the Allocation of Police Forces after a Terrorist Attack". In: *American Economic Review* 94.1, pp. 115–133.

Germain, Jean-Marc et al. (2021). *Rapport du groupe d'experts sur la mesure des inégalités et de la redistribution*. Tech. rep. INSEE.

Gethin, Amory (2023b). "Revisiting Global Poverty Reduction: Public Goods and the World Distribution of Income, 1980-2022". Working Paper.

Goldman, Maya, Ingrid Woolard, and Jon Jellema (2020). "The Impact of Taxes and Transfers on Poverty and Income Distribution in South Africa 2014/2015". AFD Research Paper.

Jestl, Stefan and Emanuel List (2022). "Inequality, Redistribution, and the Financial Crisis: Evidence from Distributional National Accounts for Austria". In: *The Review of Income and Wealth* forthcoming.

Khemani, Stuti, James Habyarimana, and Irfan Nooruddin (2019). "What Do Poor People Think About Direct Cash Transfers?" Brookings Future Development.

Levitt, Steven D. (1997). "Using Electoral Cycles in Police Hiring to Estimate the Effects of Police on Crime". In: *American Economic Review* 87.3, pp. 270–290.

Lustig, Nora, ed. (2018). *Commitment to Equity Handbook: Estimating the Impact of Fiscal Policy on Inequality and Poverty*. Brookings Institution Press.

Mladenka, Kenneth R. and Kim Quaile Hill (1978). "The Distribution of Urban Police Services". In: *The Journal of Politics* 40.1, pp. 112–133.

Motala, Shireen and David Carel (2019). "Educational Funding and Equity in South African Schools". In: *South African Schooling: The Enigma of Inequality*. Ed. by Nic Spaull and Jonathan D. Jansen. Springer. Chap. 4, pp. 67–85.

O'Dea, Cormac and Ian Preston (2010). "The distributional impact of public spending in the UK". IFS Working Paper W12/06.

Paulus, Alari, Holly Sutherland, and Panos Tsakloglou (2010). "The distributional impact of in-kind public benefits in European countries". In: *Journal of Policy Analysis and Management* 29.2, pp. 243–266.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman (2018). "Distributional National Accounts: Methods and Estimates for the United States". In: *Quarterly Journal of Economics* 133.2, pp. 553–609.

South African Treasury (2022). *Budget Review 2022*. Tech. rep. National Treasury of the Republic of South Africa.

Thesmar, David and Augustin Landier (2022). *Le prix de nos valeurs*. Flammarion.

Venkat, Hamsa and Nic Spaull (2015). "What do we know about primary teachers' mathematical content knowledge in South Africa? An analysis of SACMEQ 2007". In: *International Journal of Educational Development* 41.1, pp. 121–130.

Verbist, Gerlinde, Michael Förster, and Maria Vaalavuo (2012). "The Impact of Publicly Provided Services on the Distribution of Resources: Review of New Results and Methods". OECD Working Paper.

Wagstaff, Adam et al. (2014). "Who Benefits from Government Health Spending and Why?" World Bank Policy Research Working Paper 7044.

Figure 5.1: Government Expenditure in South Africa, 1993-2019

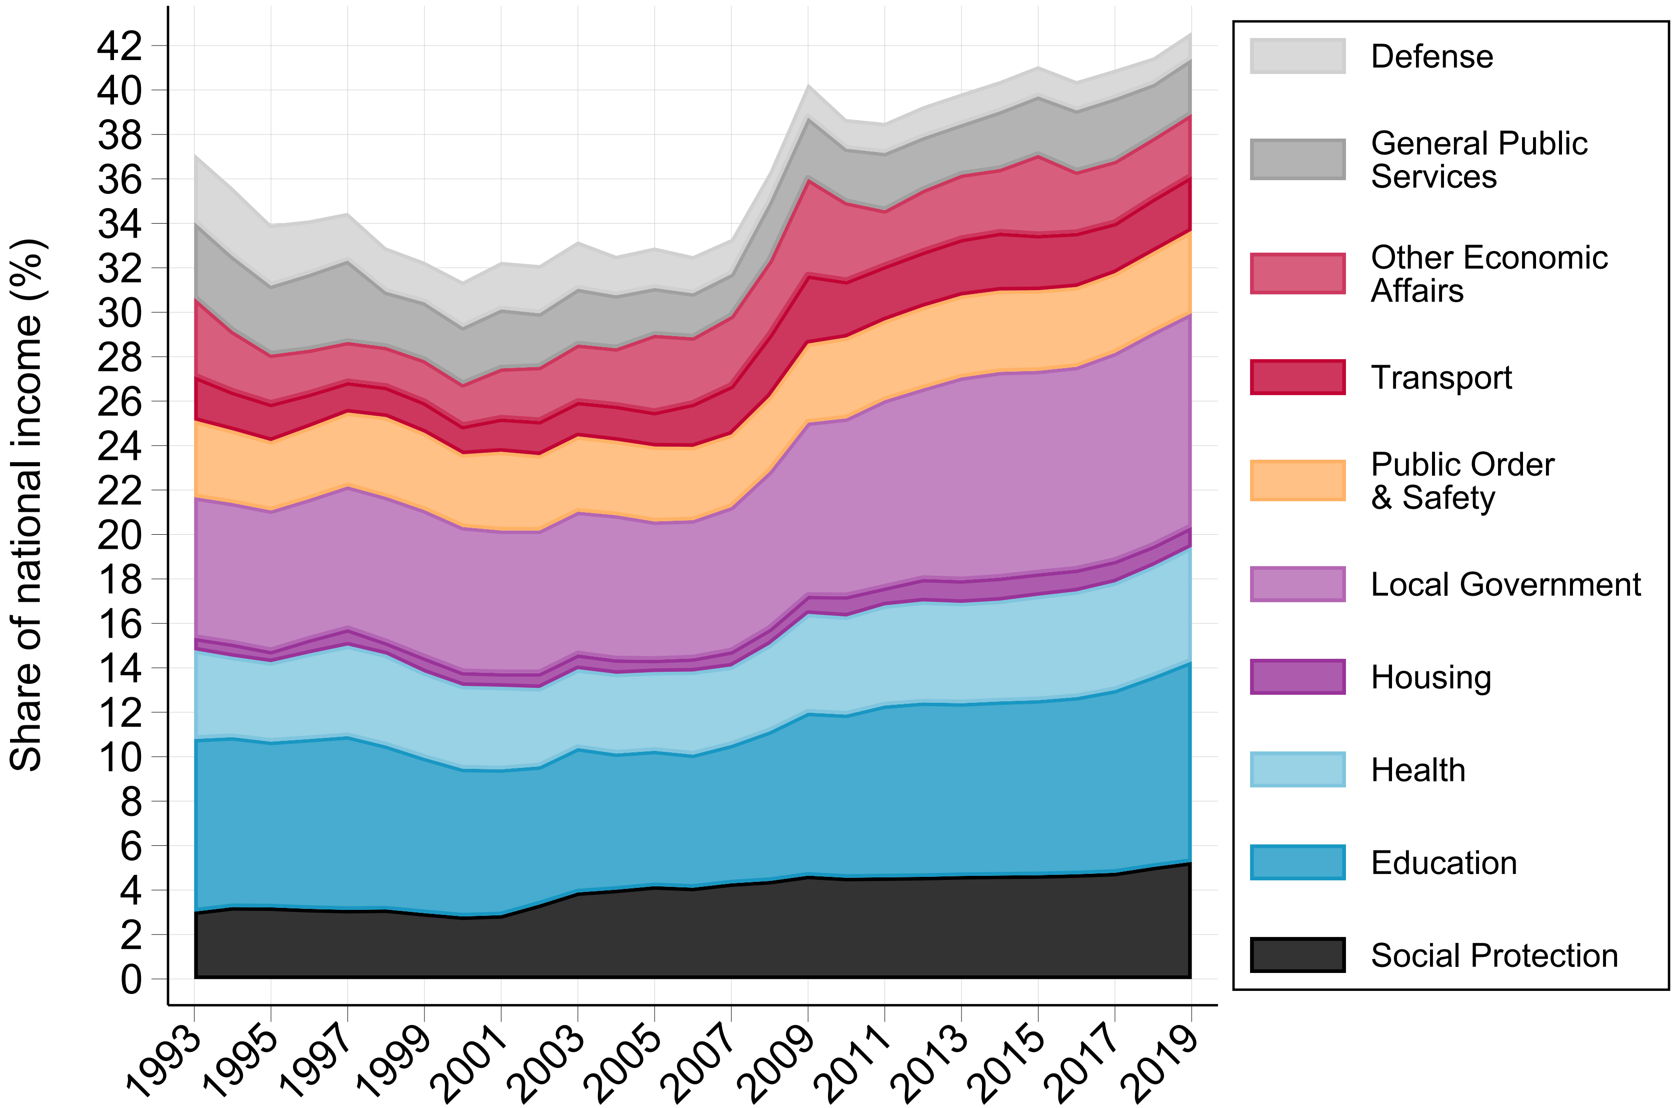

Notes. Author's computations combining data from the South African Reserve Bank, the South African National Treasury, and Local Government Budget Reports.

Figure 5.2: Government Transfers Received by Pretax Income Decile, 2019

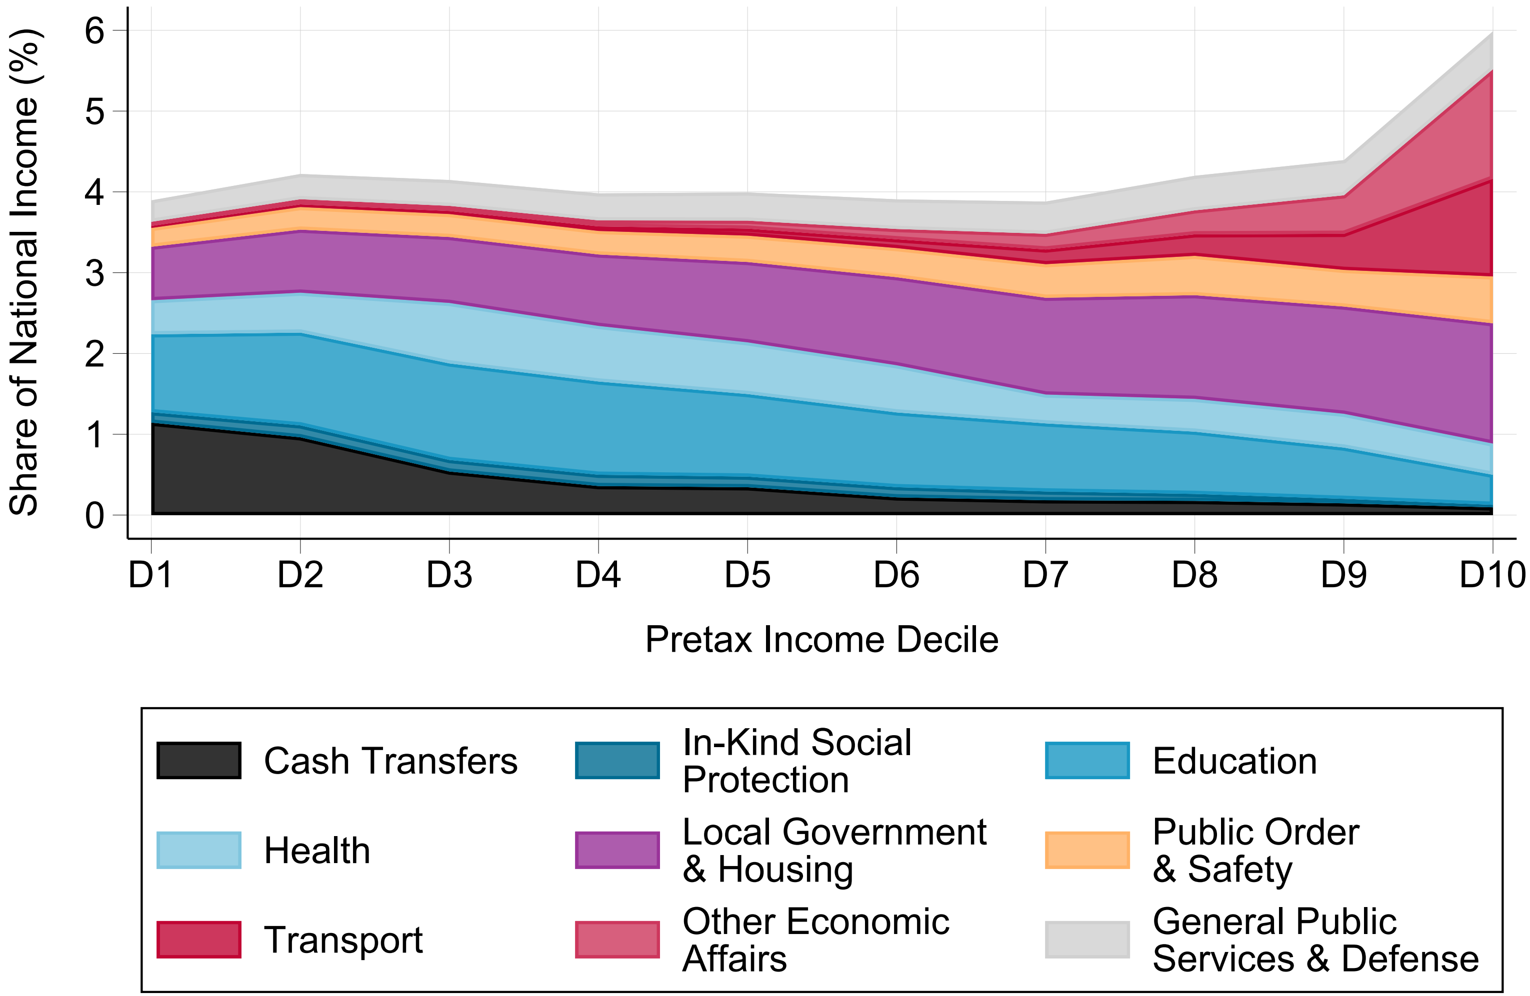

Notes. The figure represents the level and composition of total government transfers received by pretax income decile in 2019. The unit of observation is the individual. Income is split equally among all household members. General public services and defense are assumed to be distributed on a lump sum basis.

Figure 5.3: The Rise of Redistribution: Government Transfers Received by the Bottom 50%, 1993-2019

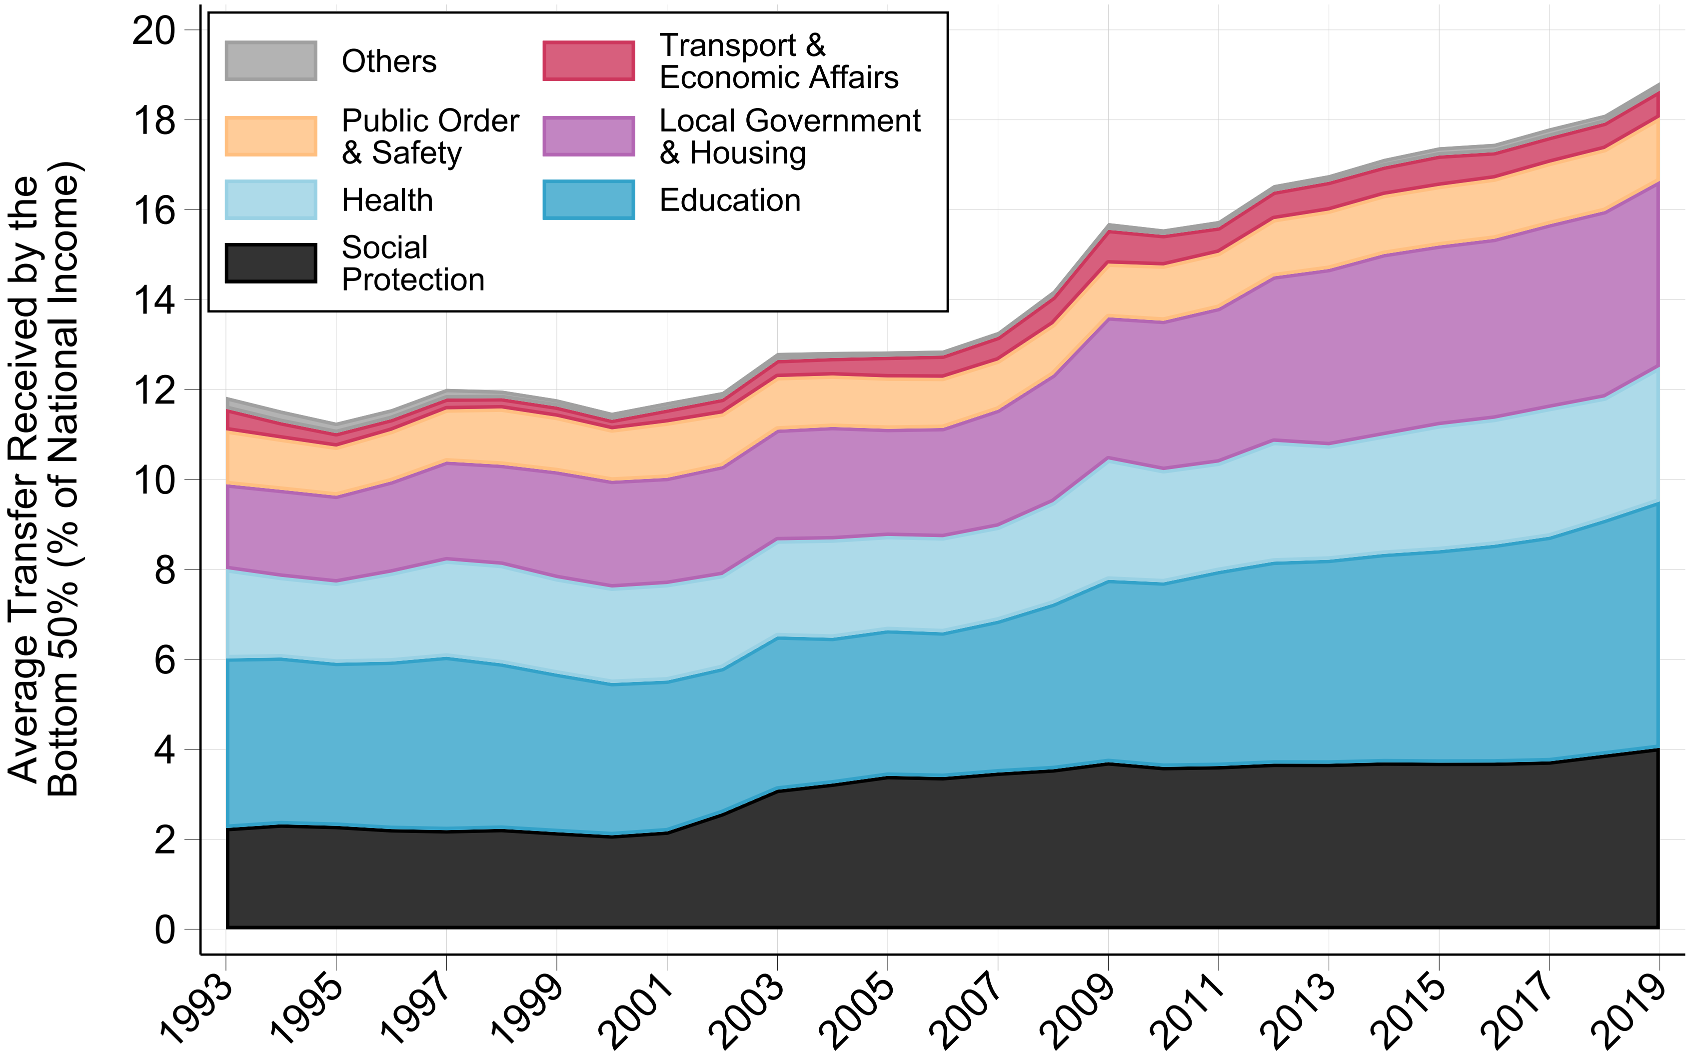

Notes. Author's computations using distributional national accounts microfile. The unit of observation is the individual. Income is split equally between all household members.

Figure 5.4: In-Kind Transfers and Poverty Reduction:
Bottom 50% Average Income Before and After Transfers, 1993-2019

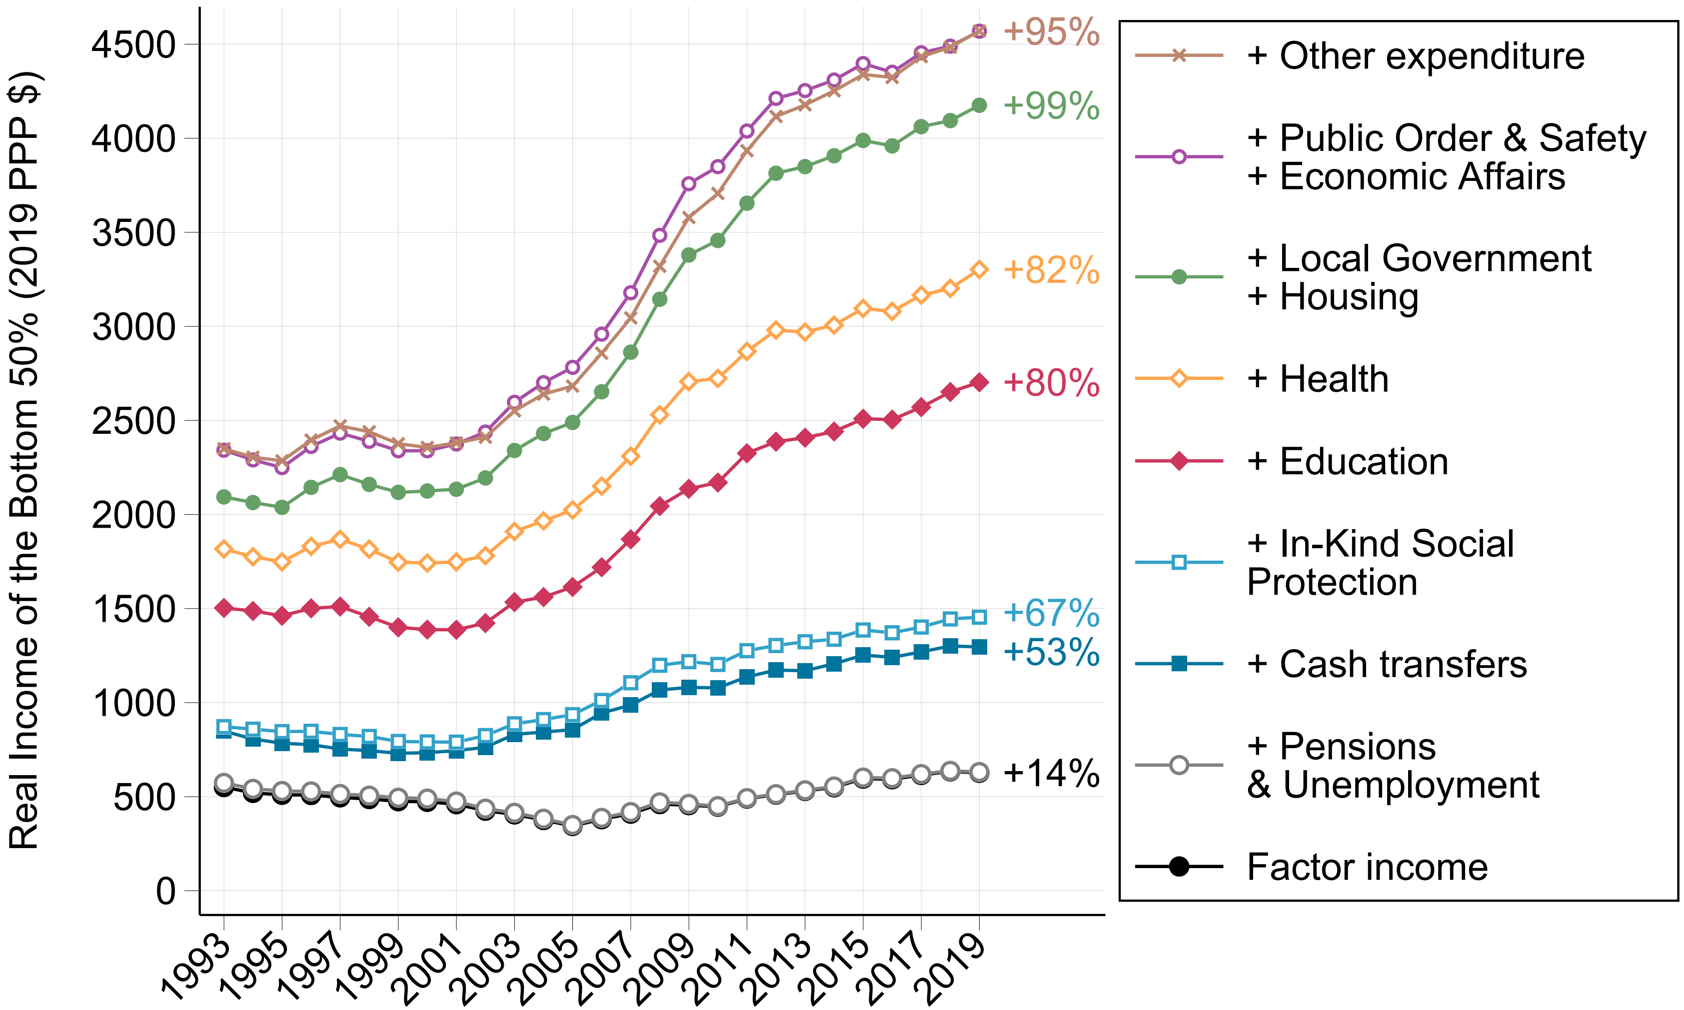

Notes. Author's computations using distributional national accounts microfile. The figure represents the evolution of the real average income of the bottom 50%, before and after adding cash and in-kind transfers one by one to the analysis. Other expenditure corresponds to general public services and defense, distributed proportionally to posttax disposable income. The unit of observation is the individual. Income is split equally between all household members.

Figure 5.5: In-Kind Transfers and Poverty Reduction:
Poverty Headcount Ratio at \$6/Day, 1993-2019

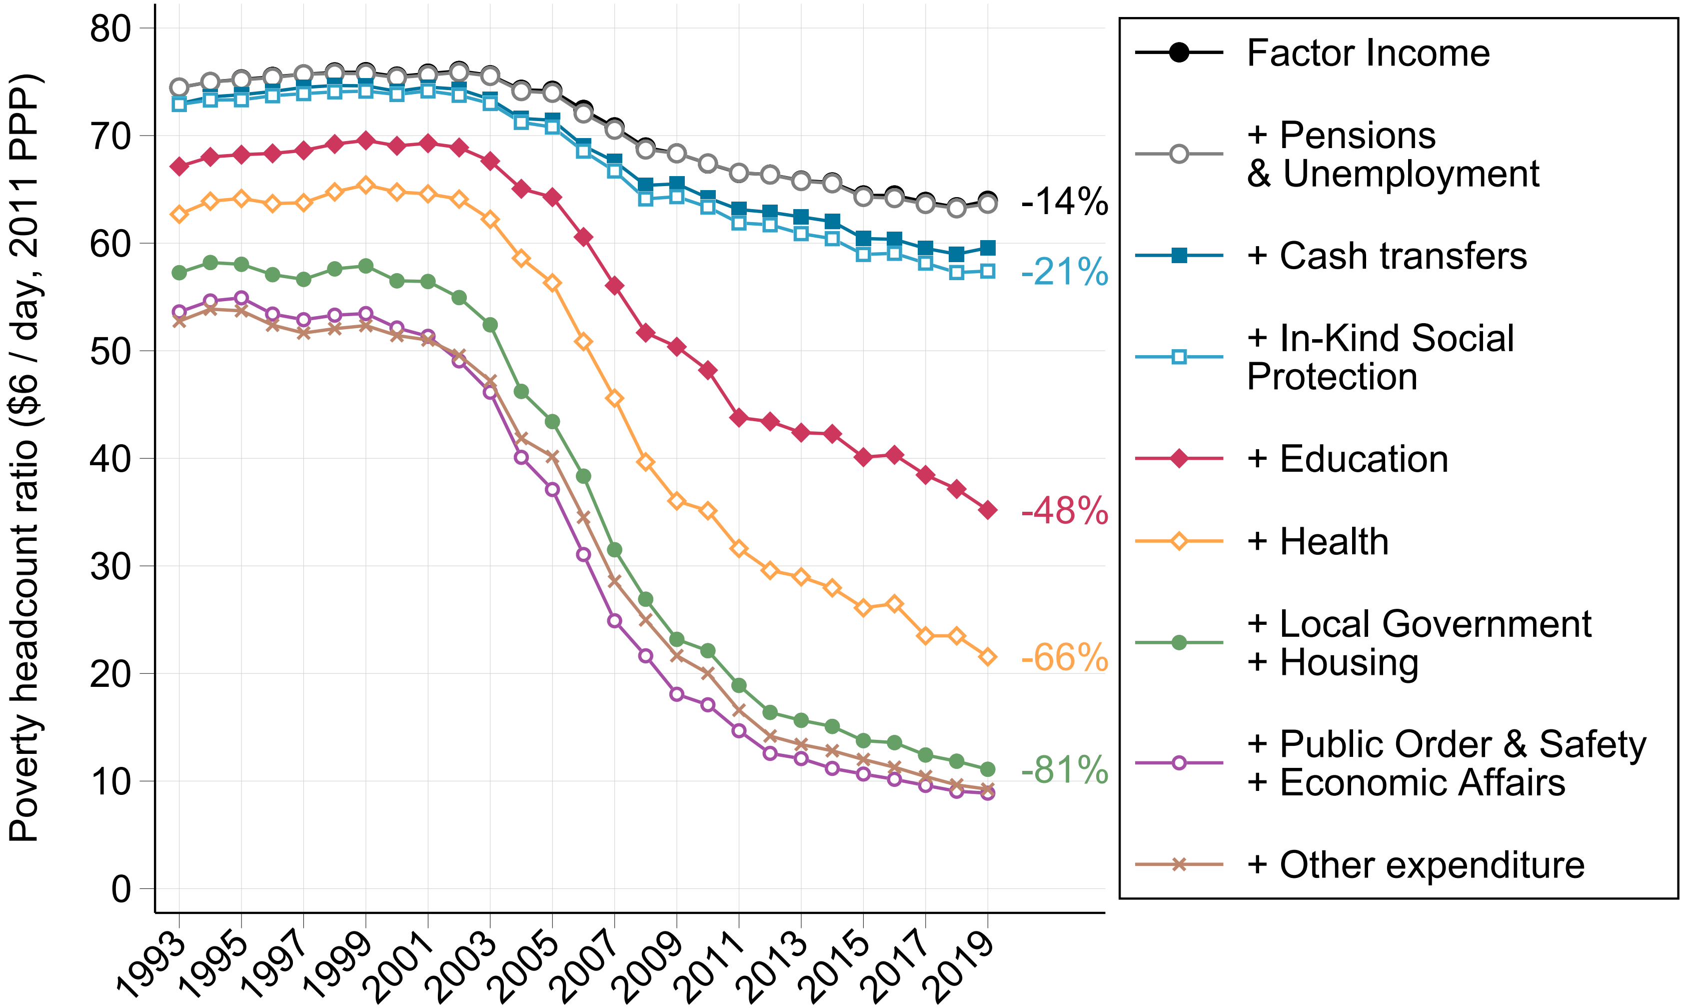

Notes. Author's computations using distributional national accounts microfile. Other expenditure corresponds to general public services and defense, distributed proportionally to posttax disposable income. The unit of observation is the individual. Income is split equally between all household members.

Table 5.1: Methodology Used to Distribute Government Expenditure in South Africa

	Method	Microdata	Macrodata	% NNI	
				1993	2019
Social Protection				3.0%	5.3%
Cash Transfers	Microsimulation	IES/LCS	National Budget	2.8%	4.2%
In-Kind Transfers	Proportional to cash transfers	IES/LCS	National Budget	0.2%	1.1%
Education	Lump sum per student, by function and province	Census	Provincial Budgets	7.8%	9.0%
Health	Proportional to healthcare use, by function and province	GHS/OHS	Provincial Budgets	4.0%	5.1%
Housing	Lump sum per beneficiary	GHS	National Budget	0.6%	0.9%
Local Government	Lump sum per municipality	Census	Local Gov. Budgets	6.3%	9.6%
Public Order and Safety				3.5%	3.7%
Visible Policing	Proportional to police presence	VCS	National Budget	1.9%	1.8%
Law Enforcement	Proportional to reported crimes	VCS	National Budget	1.5%	2.0%
Transport				2.0%	2.5%
Public Transport	Proportional to public transport expenditure	IES/LCS	National Budget	0.5%	0.5%
Infrastructure	Proportional to transport-intensive consumption	IES/LCS	National Budget Input-Output Tables	1.5%	1.9%
Other Economic Affairs	Proportional to sector-intensive consumption	IES/LCS	National Budget Input-Output Tables	3.5%	2.8%
All Others	Lump sum / proportional to income	Microfile	National Budget	6.5%	3.7%
Total				37.1%	42.6%

Notes. The table reports the methodology used to distribute the South African government budget from 1993 to 2019, together with the corresponding microdata sources, macrodata sources, and expenditure on each government function as a share of net national income in 1993 and 2019. GHS: General Household Surveys; IES: Income and Expenditure Surveys; LCS: Living Conditions Surveys; OHS: October Household Surveys; VCS: Victims of Crime Surveys.

Table 5.2: Government Redistribution in South Africa, 1993-2019:
Level, Composition, and Progressivity of Transfers Received by the Bottom 50%

	Share of Total Expenditure Received (%)			Share of National Income Received (%)			Average Transfer Received (2021 PPP USD)		
	1993	2019	1993-2019	1993	2019	1993-2019	1993	2019	1993-2019
Social Protection	74%	77%	+3%	2.3%	4.0%	+79%	390	950	+146%
Education	49%	61%	+24%	3.8%	5.5%	+44%	650	1290	+97%
Health	47%	52%	+11%	1.9%	2.7%	+44%	320	630	+97%
Housing	45%	58%	+28%	0.3%	0.5%	+105%	40	120	+181%
Local Government	26%	38%	+46%	1.6%	3.6%	+122%	280	850	+205%
Public Order and Safety	35%	38%	+10%	1.2%	1.4%	+18%	210	330	+62%
Visible Policing	38%	38%	+1%	0.7%	0.7%	-7%	120	160	+27%
Law Enforcement	31%	38%	+22%	0.5%	0.7%	+56%	80	170	+113%
Transport	7%	10%	+38%	0.1%	0.2%	+70%	20	60	+134%
Public Transport	14%	21%	+51%	0.1%	0.1%	+62%	10	30	+122%
Infrastructure	5%	7%	+38%	0.1%	0.1%	+77%	10	30	+143%
Other Economic Affairs	10%	13%	+33%	0.3%	0.4%	+5%	60	80	+44%
Total	32%	43%	+37%	11.7%	18.5%	+58%	2010	4350	+116%
Pretax Income				3.3%	2.7%	-20%	570	630	+10%

Notes. The table reports the level and composition of government transfers received by the bottom 50% of the pretax income distribution in South Africa in 1993 and 2019. Columns 2 to 4 show the share of total transfers received by the bottom 50%. Columns 5 to 7 report the corresponding share of net national income received. Columns 8 to 10 report the average annual transfer received by the bottom 50%, expressed in 2021 PPP USD. The unit of observation is the individual. Income and transfers are split equally between all household members. "Total" adds spending on defense and general public services to other rows, assuming that these two components are distributed proportionally to posttax disposable income. The last row shows the pretax income share and the average pretax income of the bottom 50%.

Table 5.3: The Distribution of Income in South Africa in 2019

	Pretax National Income		Posttax Disposable Income		Posttax National Income	
	Average Income	Income Share	Average Income	Income Share	Average Income	Income Share
Full population	$ 11,800	100%	$ 7,780	100%	$ 11,800	100%
Bottom 50%	$ 630	2.7%	$ 1,020	6.5%	$ 3,440	14.6%
Bottom 20%	$ 45	0.1%	$ 410	1.1%	$ 1,950	3.3%
Next 30%	$ 1,020	2.6%	$ 1,420	5.5%	$ 4,430	11.3%
Middle 40%	$ 8,410	28.6%	$ 6,530	33.6%	$ 10,300	35.2%
Top 10%	$ 80,700	68.7%	$ 46,600	59.9%	$ 59,000	50.2%
Top 1%	$ 329,000	28.0%	$ 170,000	21.9%	$ 219,000	18.6%
Top 0.1%	$ 970,000	8.3%	$ 519,000	6.7%	$ 633,000	5.4%

Notes. The table reports statistics on the distribution of income in South Africa in 2019 for different income concepts. Posttax disposable income is the sum of primary incomes, minus direct taxes, plus cash transfers. Posttax national income deducts all taxes and adds all transfers. General public services and defense are distributed proportionally to posttax disposable income. The unit of observation is the individual. Income is split equally between all household members.

Table 5.4: Indicators of Heterogeneous Public Service Delivery by Income Quintile in South Africa

	Q1	Q2	Q3	Q4	Q5	$q^j(Q1)$	Source
Subjective Indicators (% Positively Rating)							
Local public school	69%	69%	69%	68%	69%	1.01***	Census
Local public clinic	46%	45%	46%	46%	50%	0.98***	Census
Local public hospital	47%	47%	47%	48%	51%	0.97***	Census
Local police services	43%	43%	44%	45%	48%	0.97***	Census
Electricity supply	63%	63%	63%	64%	67%	0.99***	Census
Water supply	50%	54%	58%	62%	68%	0.85***	Census
Refuse removal services	49%	54%	57%	60%	66%	0.85***	Census
Sanitation services	52%	56%	59%	64%	74%	0.85***	Census
Government-subsidized dwelling	48%	49%	50%	51%	53%	0.96***	Census
Police response to reported crime	52%	53%	52%	53%	56%	0.98	VCS
Objective Indicators							
School teacher mathematics test success rate	38%	40%	40%	47%	67%	0.82***	SACMEQ
Share of reported crimes leading to arrest	24%	20%	21%	18%	20%	1.15	VCS
Asked to pay a bribe in past 12 months	5%	9%	8%	11%	15%	1.78***	VCS
Water interruption in past 3 months	19%	19%	17%	16%	14%	0.90***	Census
Electricity interruption in past 3 months	32%	28%	25%	21%	16%	0.76***	Census
Value of subsidized dwelling (R 1,000)	177	178	267	308	305	0.72***	GHS
Distance to Nearest Public Services (km)							
Primary school	1.5	1.5	1.6	1.8	2.0	1.12***	LCS
Secondary school	2.9	2.8	2.6	2.4	2.8	0.93***	LCS
Clinic	4.7	4.5	3.8	3.5	3.8	0.86***	LCS
Hospital	13.2	12.6	10.2	8.6	7.3	0.79***	LCS
Police station	8.6	8.1	6.1	4.9	4.6	0.75***	LCS
Public transport	1.1	1.0	1.1	1.0	1.3	1.04*	LCS

Notes. The table reports estimates of heterogeneous government productivity by income group, based on a number of subjective and objective indicators of public service delivery. Q1 to Q5 refer to income quintiles. $q^j(Q_1)$ is the corresponding measure of the relative quality of services received by the bottom quintile, equal to the ratio of the value of the indicator for Q1 to the overall sample mean (or its inverse when the scale of the variable is inverted). Statistical significance stars correspond to a regression of the indicator of interest on a dummy taking one if the individual belongs to the bottom quintile. *** p<0.01, ** p<0.05, * p<0.1. Census: 2016 national census. GHS: 2019 General Household Survey. VCS: 2017 Victims of Crime Survey. LCS: 2014-2015 Living Conditions Survey. SACMEQ: The Southern and Eastern Africa Consortium for Monitoring Educational Quality (estimates from Venkat and Spaull, 2015).

Chapter 6

Redistribution without Inclusion? Inequality in South Africa Since the End of Apartheid

Numerous studies have provided new insights into the determinants of economic deprivation in recent years, yet considerable challenges remain when it comes to accurately understanding the link between poverty, inequality, and growth. How inclusive has economic growth been in the developing world in the past decades? To what extent have cash transfers and government investments in health, education, and infrastructure development accrued to low-income groups, and what fraction of these benefits has been mitigated by an increased tax burden? Because of a critical lack of data on the joint distributions of income, consumption, taxes, and transfers, answering these questions has until today proved to be a remarkably challenging task. At the heart of this difficulty lies major differences in data sources, methods, and research communities. At the micro level, studies investigating poverty and inequality have almost exclusively relied on household surveys, often the only source at our disposal to observe the distributions of income, consumption, and wealth. At the macro level, researchers studying the determinants of growth have mostly worked with national accounts, which provide crucial information on key macroeconomic aggregates and the size of government intervention in the economy. The sometimes inconsistent and conflicting stories arising from these two sources have made it particularly difficult to understand how economic growth is shared over time and to what extent government redistribution in its various forms effectively benefits the poor.

This paper attempts to make progress in this direction by constructing a new micro dataset on the distribution of macroeconomic growth in South Africa from 1993 to 2019. Combining available data sources—surveys of various kinds, income tax data, national accounts, and historical administrative data on government taxes and expenditure from budget reports—we systematically allocate all components of the net national income, all government taxes, and all government expenditure to individuals. The resulting dataset is consistent with income, expenditure, and wealth aggregates reported in the national accounts. It is also consistent with what we know from administrative reports on various key parameters of government intervention, including the number of recipients of social grants, the total spending on each of these grants, the distribution of top taxable incomes reported in personal income tax data, and other key statistics on the size and targeting of taxes and transfers. Importantly, it covers in-kind transfers and public goods, which are particularly large and progressive in South Africa, drawing on recent related work (Gethin, 2023c). It also incorporates all these key parameters while keeping the richness of the information reported in household surveys, allowing us to decompose the evolution of inequality, redistribution, and growth according to various economic variables such as consumption, labor income, capital income, and wealth, or sociodemographic variables such as age, gender, race, and geography.

The case of South Africa is particularly revealing of the limitations we face in our understanding of the links between inequality, redistribution, and growth. On the one hand, the country is widely acknowledged as standing at the upper frontier of contemporary inequality today (Alvaredo et al., 2018). The richest 10% own a striking 85% of total household wealth, with an average net worth exceeding $440,000 at purchasing power parity (Chatterjee, Czajka, and Gethin, 2022), while 57% of the population lives with less than $5.5 per day (World Bank, 2020b). These extreme disparities, despite the end of the apartheid regime of racial segregation and exclusion at the beginning of the 1990s, have been found to have increased significantly in the past decades, driven by the boom of top incomes, chronic unemployment, and persisting household indebtedness (Bassier and Woolard, 2018; Leibbrandt et al., 2010).

On the other hand, South Africa is often regarded as displaying one of the most ambitious and efficient welfare states of the developing world. It has developed a highly progressive personal income tax, which collects substantial revenue in comparison to the majority of other emerging and developing economies. It has invested growing resources in education, health, and social protection, and its relatively well-targeted social grants system has provided critical social relief to the poor and the elderly (Bassier et al., 2020; Duflo, 2000; Maboshe and Woolard, 2018; Tondini, 2021). The reductions in inequality and poverty operated by South Africa's tax-and-transfer system have even been found to be the largest achieved among all emerging economies with comparable data (Inchauste et al., 2015).

This contrasting trajectory, mixing rising pretax income inequality, low growth, and large and increasing government redistribution leaves us with a puzzling and unclear track record of South Africa's success in improving the living conditions of the poor since the end of apartheid. We do not know, for instance, whether the decline in absolute poverty observed in the 2000s, as measured by consumption expenditure reported in household surveys, was driven by higher market incomes, improved access to credit, or social transfers (such as the Child Support Grant introduced in 1998). Commonly used consumption or income aggregates do not account for in-kind transfers, such as education and healthcare, hence leaving aside crucial elements of government redistributive policy. We know even less of the distributional incidence of taxes, in particular indirect taxes (such as VAT or excise duties) and the corporate income tax, which are generally excluded from studies tracking the evolution of inequality and poverty over time. The objective of this paper is to make advances in filling these gaps by making the best of all the available data sources at our disposal (surveys, tax data, national accounts, and historical data on the structure of taxes and transfers) to get a more complete picture of the distribution of growth and redistribution over time. While we still face considerable challenges in measuring these various components and our analysis is not devoid of limitations and uncertainties, we hope that it can contribute to improving our knowledge of how inequality, poverty, and redistribution interact in the long run.

Our analysis reveals a number of striking findings. National income per capita grew by about 35% from 1993 to 2019, yet this figure masks large heterogeneity across income groups. The average pretax income of the top 1% increased by almost 80%, while that of the poorest 20% declined. The share of pretax income accruing to the top 10% of the population thus shifted from about 64% to 69%, putting South African inequality levels much higher than anywhere else in the world, including countries such as Brazil (57%), India (57%), or the United States (45%). This dramatic rise of top incomes was driven by both capital and labor income, although labor income played a more decisive role after the 2007-2008 crisis.

Turning to the impact of taxes and transfers, we find that major increases in government redistribution more than compensated the rise of pretax inequality. This transformation was driven by both cash and in-kind transfers, which increased in size and became more progressive over time. The rise of redistribution was in part financed by higher taxes on the top 1%, who saw their effective tax rate shift from about 25% to 40%, mainly through expansions in the personal income tax and the corporate income tax. A significant part of this redistribution, however, was annulled by increases in taxes paid by the poorest 50% in the form of VAT, excise duties, and local taxes. In 2019, the profile of taxes paid by pretax income group was thus distinctly U-shaped, with higher tax rates paid by the bottom and the top of the distribution than by middle income groups, who saw their tax burden remain nearly unchanged in the past twenty-five years.

All in all, we find that growing redistribution generated substantial improvements in the living standards of South Africa's poorest individuals, but only had small effects on overall inequality. After accounting for taxes and transfers, the top 1% income share stood at nearly 20% in 2019, almost the exact same level as in 1993. The rise of South Africa's welfare state has thus succeeded at redistributing the fruits of economic growth, but it has been insufficient to curb the extreme inequalities inherited from apartheid.

Finally, we decompose inequality and redistribution into two key historical determinants of South Africa's extreme economic disparities: race and geography. We document a significant decline in the income gap between White and Black South Africans since the mid-2000s: the ratio of average White income to average Black income fell from 14 in 2005-2009 to 8 in 2015-2019. However, much of this decrease can be accounted for by the top 10% of Black earners, who witnessed exceptional income gains. When we exclude this group from the analysis, the racial income gap appears to have stood at about the same level in 2019 as since 1993. Racial inequalities are substantially larger in terms of wealth than in terms of income or consumption and are only moderately reduced by the tax-and-transfer system. Turning to geography, we find that the South African state operates significant transfers from the two richest provinces, Gauteng and Western Cape, to the rest of the country, although posttax spatial inequalities remain large. Another dimension of inequality for which redistribution seems to have succeeded at fully absorbing the rise of pretax inequalities is the rural-urban gap, which grew dramatically in terms of pretax income but remained stable after accounting for taxes and transfers. This striking rise of redistribution from urban to rural areas is mostly attributable to higher in-kind transfers and public goods, which disproportionately improved in rural areas over the period considered.

This paper contributes to the growing literature attempting to bridge the micro-macro gap in poverty and inequality studies. Piketty, Saez, and Zucman (2018) combine surveys, tax, and national accounts data to create Distributional National Accounts (DINA) allocating the entirety of national income growth to individuals in the United States since 1913. A number of studies following the DINA framework (see Alvaredo et al., 2018) have been conducted since then on other countries or regions of the world, with the objective of constructing comparable, yearly statistics on the long-run distribution of income and wealth.[1] The major innovation of DINA studies is their consistency with macroeconomic figures reported in the national accounts and their allocation of all taxes and transfers (including indirect taxes, in-kind transfers, and collective government expenditure) to individuals. One of their limitations, however, has been the degree of precision with which taxes and transfers are distributed. Piketty, Saez, and Zucman (2018), for instance, distribute education spending as a lump sum per child, leaving aside variations in expenditure across space and level of education. Blanchet, Chancel, and Gethin (2022) distribute health expenditure on a lump sum basis in the context of Europe, assuming that all adults benefit from the same amount of health investment regardless of age, location, or socioeconomic status. Similarly, Bozio et al. (2022) allocate all consumption taxes on value added, energy, or tobacco proportionally to overall consumption expenditure, regardless on the type of goods on which these taxes fall.

Several studies have made significant efforts to refine our understanding of the distribution of indirect taxes and in-kind transfers, but they have typically not done so in a way that is consistent with the national accounts. In the context of South Africa, Inchauste et al. (2015) exploit data from the Living Conditions Survey to allocate government taxes and social spending to individuals in a particularly granular way, combining for instance microdata on educational attendance by program with figures on total expenditure on each of these programs by province to allocate total education spending. This allows the authors to derive a much more precise estimate of the distributional incidence of some of the elements of the tax-and-transfer system. However, this estimate covers only one year, is not consistent with national accounts, relies exclusively on surveys (which tend to underestimate income at the top end), and excludes key components of government revenue and spending (e.g., the corporate income tax).

In this paper, we attempt to take the best from all of these contributions to derive a

[1]See for instance Blanchet, Chancel, and Gethin (2022) on Europe, Garbinti, Goupille-Lebret, and Piketty (2018) and Bozio et al. (2022) on France, or Piketty, Yang, and Zucman (2019) on China. The results of these studies have been compiled in the World Inequality Database (see http://wid.world).

comprehensive picture of the distribution of growth, taxes, and transfers in South Africa given the data at our disposal. We directly follow the Distributional National Accounts framework and distribute, component by component, the national income between 1993 and 2019. Our estimates account for incomes that are never directly received by individuals, such as imputed rents or corporate undistributed profits, yet are key components of macroeconomic growth figures. We allocate all taxes to individuals, accounting for key features of the tax system such as VAT-exempt goods, the types of expenditure facing excise taxes, the heterogeneous effects of trade duties through variations in import densities by type of good, expenditure made in the informal sector, and personal income tax exemptions. We distribute all government expenditure as precisely as possible, incorporating information on the value and the number of recipients of each social grant from historical budget reports, excluding individuals relying on private health insurance or going to private schools from public spending, and decomposing education and health transfers by province and function following recent work by Gethin (2023c). Although our estimates are far from being perfect and could be improved as better data becomes available, we hope that these methodological insights can contribute to make new steps towards the much needed reconciliation between macro and micro sources in economics research.

Section 6.1 covers data sources and methodology. Section 6.2 presents results on the distribution of income before accounting for taxes and transfers. Section 6.3 studies the impact of taxes and transfers on inequality and the distribution of growth since 1993. Section 6.4 decomposes inequality and redistribution by race and geography.

6.1 Data and Methodology

This section presents the data sources and methodology used to estimate the distribution of pretax income, posttax income, consumption, and wealth in South Africa between 1993 and 2019. Section 6.1.1 outlines our conceptual framework. Section 6.1.2 explains how we distribute factor national income combining surveys, tax, and national accounts data. Section 6.1.3 details how we move from factor income to pretax income. Sections 6.1.4 and 6.1.5 cover the allocation of taxes and transfers. Section 6.1.6 describes how we estimate the distribution of household final consumption expenditure and household wealth.

6.1.1 Conceptual Framework: Distributional National Accounts

We are interested in distributing the consumption, income, and wealth aggregates codified in the United Nations' System of National Accounts (UN SNA), which are routinely estimated by statistical institutes and used to estimate and decompose macroeconomic growth. These include net national income, household final consumption expenditure, and household net worth.

Net National Income. Our benchmark income concept is net national income. National income equals GDP minus capital depreciation plus net foreign income. It is the sum of the primary incomes of the different sectors of the economy: households, corporations, and the government (see Table 6.1). The primary income of households can itself be decomposed into four main components: compensation of employees, mixed income, net property income, and the imputed rents of owner-occupiers. The primary income of corporations corresponds to the net benefit that companies retain after having paid suppliers, employees, shareholders, and taxes, and that we refer to interchangeably as "retained earnings" or "undistributed profits". The primary income of the general government is the sum of taxes less subsidies on production and imports (i.e., indirect taxes collected during the production process) and of its net property income.

Distributional Income Concepts. Following the DINA framework (Blanchet et al., 2021), we consider three main income concepts to distribute national income at the individual level. Factor national income is the sum of all income flows accruing to individuals before any tax or transfer. Pretax national income equals factor income after the operation of unemployment and pension systems, that is, after payment of social contributions and distribution of pension and unemployment benefits. Posttax national income equals pretax income after deduction of all taxes (including indirect taxes and the corporate income tax), payment of all kinds of transfers (including collective government expenditure in health, education, defense etc.), and allocation of the general government deficit of surplus. By definition, individual factor incomes, pretax incomes, and posttax incomes all add up to the net national income.

Distributional Consumption and Wealth Concepts. In addition to income, we also distribute consumption and wealth concepts consistent with national accounts definitions. Household final consumption expenditure (HFCE) is the sum of all purchases made by resident households. The net saving of households is the difference between net disposable income (posttax income excluding collective government expenditure) and HFCE. Personal wealth is the net wealth of the households sector, that is, the sum of all financial and non-financial assets held by households, minus their financial liabilities.

6.1.2 From Reported Household Income to Factor National Income

We now outline our methodology to distribute factor national income. We first combine survey and tax data to measure the distribution of reported household income (wages, property income, and mixed income). We then allocate unreported income components (imputed rents, property income attributed to policy insurance holders, undistributed profits, and government primary income) to individuals. Table 6.1 outlines the methodology used to distribute each of these subcomponents of factor national income. We discuss in greater detail these methodological steps in appendix F.1.

Harmonization of Survey Data. Household surveys represent our main data source to distribute income at the individual level. Seven surveys collecting detailed information on all components of household income and expenditure have been conducted in South Africa since 1993. We combine these "income surveys" with labor force surveys, which provide more detailed data on wages and self-employment income on an annual basis, to build a microfile covering the distribution of "reported household income" every year since 1993.

Combination of Survey and Tax Data. Surveys can be well-suited to measure income and expenditure at the bottom of the distribution, yet they are well-known to underestimate inequality at the top end (e.g., Blanchet, Chancel, and Gethin, 2022). To better capture the levels and dynamics of top incomes, we combine our survey microfile with tabulated income tax returns available from the South African Revenue Service. The available tabulations report the number of taxpayers and total taxable income by income tax bracket every year since 2002. We correct the survey data with the tax data in four steps. First, we approximate full distributions from the tax tabulations using Generalized Pareto Interpolation (Blanchet, Fournier, and Piketty, 2021). Secondly, we define a "taxable income" concept in the survey data that is comparable to that observed in the tax data (excluding in particular dividends, which are not subject to personal income tax in South Africa). Thirdly, we calibrate the survey microdata on the tax tabulations using the algorithm developed by Blanchet, Flores, and Morgan (2022), which reweighs survey observations so as to match the distribution of top taxable incomes reported in the tax data. This method has the major advantage of preserving the survey microdata and the dependency between its different variables (such as income components and sociodemographics), while enforcing that the survey becomes fully representative of top taxable incomes, in the same way that statistical institutes routinely adjust survey weights to make them more representative in terms of age or gender. Finally, we extrapolate the correction to the 1993-2001 period, for which no tax data is available, assuming that top incomes were underrepresented during this period to the same extent as in 2002.

Rescaling of Household Income Components to National Accounts Totals. Having combined survey and tax data, we now have a microfile covering reported household income for the full South African population since 1993. However, for various reasons linked to sampling, mismeasurement of income flows, and non-response, income aggregates reported in this microfile do not necessarily match those recorded in the national accounts. Following other DINA studies, we rescale proportionally each of the five income flows reported in survey and tax data—compensation of employees, mixed income, rental income, interest, and dividends—to their corresponding national accounts totals. This step only has minor distributional implications at the bottom of the distribution, but it leads to significantly increasing the income share of the top 1%. This is because capital incomes, in particular interest and dividends, are both massively underreported in household surveys and by construction mostly absent from South African income tax data (Chatterjee, Czajka, and Gethin, 2022).

Imputed Rents. The imputed rents of owner-occupiers represent about 3% of national income. Imputed rents are not recorded consistently in South African surveys as such, but income surveys have asked households to give an approximate value of the value of their home since 1993. We use this information to distribute imputed rents proportionally to the market value of owner-occupied housing wealth.

Other property income. Other property income, also referred to as property income attributed to insurance holders and pension entitlements, corresponds to investment income indirectly received by individuals through their ownership of unmatured insurance and pension assets. Accordingly, we assume that it is distributed proportionally to pension and life insurance assets, estimated by combining data on wages, social contributions, and self-reported wealth data from the National Income Dynamics Study (see Chatterjee, Czajka, and Gethin, 2022). This component represents a significant share of national income in South Africa (6% in 2018), where private pensions, life insurance policies, and investment funds are widespread and have been growing in the past decades.

Interest Paid by Households. Household debts in the form of mortgages and other loans are significant in South Africa (53% of national income in 2018), and particularly widespread at the bottom of the wealth distribution (Chatterjee, Czajka, and Gethin, 2022). Accordingly, interest paid by households represents a sizable component of national income, reaching 5% of NNI in 2018. Data on debt balances have been recorded in income surveys since 1993, but debt repayments are only partially and inconsistently measured. To avoid artificially creating too many households with negative income, we therefore choose to distribute interest paid proportionally to factor income among individuals who declare having unpaid debts.

Corporate Undistributed Profits. Undistributed profits correspond to profits that are kept within the company rather than distributed to shareholders as dividends. These income flows ultimately increase the wealth of shareholders and therefore represent a source of income to them. Accordingly, we allocate retained earnings proportionally to stock ownership, including both directly held shares and shares held indirectly through pension funds. We only distribute the share of retained earnings attributable to the private domestic sector, hence excluding that held by the government.

Taxes less Subsidies on Production and Imports. We allocate the primary income of the government proportionally to factor income, assuming this component of national income is distributionally neutral. This assumption is meaningful to the extent that one could replicate our entire analysis by relying on a definition of net national income at factor cost (instead of market prices), excluding indirect taxes and subsidies from the final measure of output. Our inequality series is thus insensitive to adopting one or the other of these approaches to national accounting.

Remaining components of factor national income. The remaining components of national income (3 % of NNI) mainly include government and foreign shares of corporate retained earnings, as well as other small income flows such as miscellaneous government transfers. In the absence of better information on the incidence in these items, we assume for simplicity that they are distributionally neutral and allocate them proportionally to factor income.

6.1.3 From Factor National Income to Pretax National Income

To recover pretax income from factor income, we remove all pension and unemployment contributions from individual income and we add all corresponding pension and unemployment benefits. This has only minor distributional incidence in South Africa, given that private pension benefits are received by a small share of the population and that the unemployment insurance system only redistributes a tiny fraction of national income (see Table 6.1).

Pension Contributions. Contributions to private pension plans (6% of national income) are recorded in income surveys, so we directly deduct them from individual factor incomes.

Pension Benefits. Private pension benefits (3% of national income) are also recorded in income surveys. However, these surveys tend to significantly underestimate the share of adults receiving private pension income (2-3% in income surveys vs. 5-6% according to administrative data). We use predictive mean matching to impute incomes to individuals declaring no pension income but with characteristics similar to those who do, in such a way that the total number of pension income recipients becomes exactly equal to that observed in administrative data sources. This ensures that our microfile is representative of what we know about the actual number of recipients of pension benefits in South Africa, while preserving the observed relationships between pension income and the other covariates recorded in the surveys.

Unemployment Insurance Contributions. Unemployment insurance contributions are set at a fixed rate of 2% of gross wage in South Africa and capped at a maximum amount in Rand. About 25% of adults contribute to the Unemployment Insurance Fund (UIF), collecting some 0.4% of national income in 2018. UIF contributors are well identified in labor force surveys, so we directly impute contributions based on statutory rules.

Unemployment Insurance Benefits. Unemployment insurance benefits are only available to adults having previously made monthly contributions to the UIF. This explains why they only cover a small fraction of the population (1.9% in 2018) and represent only 0.4% of national income. Unemployment benefits and beneficiaries are recorded in income surveys but are typically underrepresented. As in the case of private pension income, we therefore impute UIF benefits to additional recipients and we proportionally rescale the value of these benefits, so as to perfectly match both the official number of recipients and total UIF expenditure recorded in administrative data sources.

Pension and Unemployment Deficits or Surpluses. To ensure that pretax national income equals factor national income, we have to distribute the surpluses or deficits of the pension and unemployment systems. Following other DINA studies, we distribute 50% of the gap between contributions and benefits to contributors proportionally to contributions paid, and 50% to recipients proportionally to benefits received. This corresponds to assuming that the burden of the deficit (or the benefits of the surplus) will eventually be shared 50/50 by contributors and recipients.

6.1.4 From Pretax National Income to Posttax National Income: Taxes

To move from pretax income to posttax income, we start by deducting all taxes paid (see Table 2). These include all direct taxes (including the personal income tax and the corporate income tax) and all indirect taxes (including the Value Added Tax and excise duties).

Personal Income Tax. The personal income tax (PIT) is the tax collecting the highest share of government revenue in South Africa, amounting to 11% of national income in 2018. We microsimulate the income tax at the individual level, for each year since 1993, exploiting information on statutory rules, thresholds and marginal tax rates collected from historical administrative sources. As our microfile is calibrated on tabulated income tax returns, it is perfectly representative of taxable incomes at the top. It is therefore fully consistent with administrative data, both in terms of the number of taxpayers and total income tax receipts.

Corporate Income Tax. The corporate income tax (CIT) is the second biggest direct tax on income in South Africa (6% of national income in 2018). The CIT is paid on corporations' profits, so we distribute it similarly to retained earnings, that is, proportionally to directly and indirectly held shares.

Other Direct Taxes on Income and Wealth. Other direct taxes on individual income and wealth only represent a small fraction of national income (1.7% in 2018). We distribute the dividends tax (0.8% of NNI), a flat tax of 20% paid by individuals on dividends received from South African companies, proportionally to dividends received. The Skills Development Levy (0.4% of NNI) is a flat tax of 1% paid on the wages of employees registered with the UIF, so we impute it directly based on rules. We allocate the remaining direct taxes to their corresponding tax bases: transfer duties to housing wealth (0.2% of national income), the securities transfer tax to equity ownership (0.1%), the estate duty and the donations tax to net wealth (0.1%), and the remaining taxes on income to pretax income (0.1%).

Value Added Tax. The value added tax (VAT) is the largest indirect tax in South Africa, enforced at a standard rate of 15% and collecting 8% of national income in 2018. In line with DINA studies and with standard tax incidence analyses, we assume that the VAT is paid by consumers. However, we refine our VAT tax incidence model in two ways. First, we exclude 19 "basic food goods" , which are zero-rated and therefore not subject to VAT, as well as all other VAT-exempt goods and services (including housing rents, transport services, petrol products, educational expenditure, and financial services: see South African Reserve Bank, 2019). Household expenditure on each of these items has been recorded in all income surveys, so we can directly remove them from our consumption aggregate. Secondly, following Bachas, Gadenne, and Jensen (2022), we exclude goods and services bought on the informal market, approximated by the type of store in which purchases occur. These two steps significantly mitigate the regressive impact of VAT, although not sufficiently to make it progressive, given the particularly high gap between consumption and income at the bottom of the distribution and the small size of the informal sector in South Africa.

General Fuel Levy and Excise Duties. Other indirect taxes on domestic products include the general fuel levy (1.8% of NNI), other excise duties (1.1%), and other taxes on goods and services (0.3%). The general fuel levy is a tax on fuel consumption, so we distribute it proportionally to fuel and transport expenditure. Other excise duties correspond to taxes on tobacco and alcohol, paid at production, so we distribute them proportionally to spending on these two goods. Other taxes on goods and services include a number of other minor indirect taxes, which we distribute proportionally to overall household expenditure.

Taxes on International Trade. Import duties and other taxes on international trade together represent about 1.4% of national income. A simplified way to distribute these taxes would be to assume that they are borne by consumers as VAT. However, the nature of imported goods might differ significantly from that bought by a typical consumer, leading to biased estimates. To correct for heterogeneity in consumption of domestic vs. imported goods, we use input-output tables published by Statistics South Africa to derive an estimate of import density by COICOP category of household expenditure. We then distribute taxes on international trade proportionally to import-density-corrected consumption.

Local Taxes Local government revenue in South Africa consists mainly in property rates, service charges for the provision of electricity, water, and other services such as refuse removal, and transfers received from the central government. Since the latter are financed by central government revenue, we do not allocate them to individuals (doing so would lead to double counting, as transfers to municipalities are indirectly financed by national taxes). Property rates, electricity charges, and water charges are directly reported by households in income surveys, so we allocate budget totals proportionally to these reported values. We distribute the remaining components of municipal operating revenue proportionally to the total municipal tax burden of each individual, so as to match total revenue reported in municipal budgets.

Other Tax and Non-Tax Revenue. To reach total consolidated government revenue, we distribute the remaining tax and non-tax revenue proportionally to pretax income (i.e., in a distributionally neutral way). These include all other taxes not previously mentioned (less than 0.1% of national income), payments to the Southern African Customs Union (-1.2%), non-tax revenue (0.8%), and revenue collected by provinces and other public entities (2.4%).

6.1.5 From Pretax National Income to Posttax National Income: Transfers

Having removed all taxes from pretax income, we now allocate all government expenditure—including both direct and in-kind transfers, as well as the government deficit—to individual incomes to reach posttax national income.

Direct Social Transfers. Social protection spending represents about 5% of the national income in South Africa, the majority of which consists in three social grants: the old age grant (1.8%), the child support grant (1.5%), and the disability grant (0.6%). The old age grant is a means-tested monthly benefit available to South Africans older than 60. The child support grant is granted to a child's primary caregiver whose income falls below a specific threshold. The disability grant is provided to workers suffering from a permanent disability. As in the case of pension and unemployment benefits, data on social grants is available in income surveys, but the number of self-reported recipients tends to be lower than in administrative data (although only slightly). For consistency, we attribute social grants to additional eligible beneficiaries using a linear probability model, and we impute the value of grants received based on statutory rules, each year since 1993 (see appendix F.1). This ensures that our microfile is fully consistent with both the number of grant beneficiaries and total government expenditure on grants.

Other In-Kind Transfers and Public Goods. We take the distribution of other government expenditure, including education, healthcare, transport infrastructure, police services, and other public goods from Gethin (2023c). This related paper combines various surveys and historical budget data to identify beneficiaries of public services and the corresponding distribution of public spending by function of government from 1993 to 2019.

Government Deficit. As in the case of the deficits of pension or unemployment systems, we assume that 50% of the general government deficit (6% of NNI in 2019) is borne by taxpayers proportionally to total taxes paid, and 50% proportionally to total transfers received.

6.1.6 Household Expenditure and Household Wealth

In addition to factor income, pretax income, and posttax income, we also distribute consumption and wealth concepts consistent with national accounts definitions.

Household Expenditure. Following our approach for income, we distribute household final consumption expenditure by proportionally rescaling subcomponents of consumption reported in income surveys to their corresponding totals recorded in the national accounts, for each of the 12 COICOP categories available in both micro and macro data. This allows us to document the joint dynamics of consumption and income at the individual level, as well as to derive estimates of net saving (net household disposable income minus HFCE) by income group that are consistent with macroeconomic figures.

Household Wealth. Finally, we combine survey data on income and wealth with households balance sheets statistics published by the South African Reserve Bank to add an estimate of household net worth and its composition to our microfile since 1993. We do so by applying a "mixed method" combining rescaling and income capitalization (following Saez and Zucman, 2016), whereby specific household wealth components from the balance sheets are distributed proportionally either to the corresponding income flows they generate, or to the market values of assets or liabilities reported by survey respondents. For more information on this methodology, we refer to our companion paper dedicated to the estimation of wealth inequality in South Africa (Chatterjee, Czajka, and Gethin, 2022).

6.2 The Distribution of Factor Income

In this section, we present results on the distribution of factor national income, that is, income arising from the use of production factors (capital and labor) before any form of government redistribution. This analysis serves as a basis for understanding the evolution and structure of income inequality in South Africa, which play a key role in determining the allocation of taxes and transfers. Section 6.2.1 provides background information on macroeconomic growth in South Africa. Section 6.2.2 describes on the dynamics of factor income inequality. Section 6.2.3 decomposes the distribution of factor income into its labor and capital components.

6.2.1 National Income Growth in South Africa since 1993

Figure 6.1 plots the evolution and composition of real national income per capita in South Africa since 1993. Average national income was equal to $12,800 (80,000 Rand) at purchasing power parity in 2019, up by 37% from its 1993 level. Macroeconomic growth can be decomposed into three main phases: a phase of economic stagnation between 1993 and 2000, during which real national income remained relatively stable; a phase of fast growth between 2000 and 2011; and a phase of decline since 2012, characterized by a negative average annual growth rate.

Following the income approach to national accounting, the national income can be decomposed into its different income components. These components can be grouped into four main aggregates: compensation of employees (57% of NNI in 2019), mixed income and imputed rents (14%), household property income and corporate undistributed profits (17%), and government primary income (12%). Gross wages have followed a U-shaped curve over the period, dropping from about 60% of the national income to 50% in 2006, before bouncing back since then. The share of mixed income and imputed rents in national income has fluctuated with no clear trend. Conversely to wages, property income received by household and corporate undistributed profits have followed an inverted U-shaped curve, growing from 20% of

NNI to about 25% from 1993 to the mid-2000s, before falling back to 17% in 2019.

6.2.2 The Distribution of Factor National Income

We now turn to documenting changes in the distribution of national income, before the operation of the pension, unemployment insurance, and tax-and-transfer systems. Table 6.3 reports data on the distribution of factor national income in 2019 across selected income groups, revealing extreme income disparities. About one third of total income accrued to the poorest 90% of the population in 2019, compared to two-thirds for the richest 10% and over 28% for the top 1% alone. The top 0.01% of the population (5,860 individuals) received about 2% of factor national income, almost as much as the poorest 50% as a whole (29 million individuals). The bottom 50% have an average income of $600 per year at purchasing power parity, which is 20 times lower than the national income per adult. Meanwhile, the top 10% received $80,000 (7 times the national average) and the top 0.1% almost 1 million PPP dollars (80 times the national average).

In Figure 6.2a, we represent the evolution of income inequality in South Africa over time and compare it to that observed in other countries for which comparable distributional national accounts studies have been conducted. South Africa stands at the upper frontier of global income inequality today: the share of income accruing to the top 10% exceeded 65% in 2019, compared to 55-60% in Brazil and India, 40-45% in China and the United States, and below 35% in France. Furthermore, the top 10% share has increased significantly since 1993, moving up by almost ten percentage points between 1993 and the early 2000s, before stabilizing thereafter. The 2007-2008 crisis has been associated with a slight drop in top income inequality, as observed for instance in a number of European countries (see Blanchet, Chancel, and Gethin, 2022).

Figure 6.2b plots the cumulative income growth of the top 1%, the top 10%, the middle 40%, and the bottom 50%. A striking divergence in real factor incomes has taken place between the bottom and the top of the distribution since the end of apartheid. Between 1993 and 2019, the average national income grew by 37%, yet it rose by almost 50% for the top decile and by over 70% for the richest percentile. Meanwhile, the average factor income of the middle 40% grew by 20% and that of the bottom 50% almost stagnated. Coming back to the three phases of national income growth outlined above, we see that the stagnant decade of the 1990s was associated with dramatically different trajectories across income groups, as the boom of top incomes was almost perfectly compensated by income losses among the bottom 90%.

Economic growth in the early 2000s benefited both the top and the bottom of the distribution, albeit significantly more the former than the latter. Finally, the drop in real incomes after 2011 was mostly driven by the top of the distribution, but this decline was insufficient to bring back inequality to its 1993 level.

6.2.3 Decomposing Factor Income Inequality: Labor versus Capital

To shed light on some of the factors behind the rise of income inequality, one can start by decomposing income into its labor and capital components. This decomposition also directly informs the tax incidence analysis conducted in section 6.3, given that the distribution of taxes is highly dependent on the distribution of the various income components on which taxes fall.

Figure 6.3a represents the evolution of the top 1% capital income, labor income, and total factor income shares since 1993. Three results stand out. First, in line with what we observe in the majority of countries with available data (e.g., Garbinti, Goupille-Lebret, and Piketty, 2018; Piketty, Saez, and Zucman, 2018), capital income inequality has remained substantially higher than labor income inequality throughout the entire period. Second, the boom of top 1% incomes of the 1990s and 2000s was driven by both labor and capital income: the top 1% capital income share grew from 55% to 60-65% from 1993 to 2007 and the top 1% labor income share from 15% to 25%. Third, the stabilization of top incomes after 2007 masks a divergence between a continued increase in labor income concentration and a decline in top capital income inequality. This decline partly mirrors dynamics at the macro level, in particular the fall of household property income and the growing share of wages in the national income since the 2007-2008 crisis that we documented in section 6.2.1.

As shown in 6.3b, a direct consequence of these differential dynamics has been a significant increase in the labor share of income at the top. In 1993, investment income (interest, dividends, and rental income) and undistributed profits represented 60% of factor income of the top 1%, while labor income amounted to less than 30%. The share of wages in top 1% incomes has grown to nearly 50% in 2019, due in large part to the shrinking size of investment income since the early 2010s.

6.3 The Distribution of Taxes and Transfers

To what extent have taxes and transfers curbed the rise of inequality in South Africa? To answer this question, we now analyze the distribution of taxes and transfers and its impact on inequality and the real incomes of income groups throughout the distribution since 1993. Section 6.3.1 discusses the impact of the pension and unemployment system. Sections 6.3.2 and 6.3.3 respectively analyze the distributional incidences of taxes and transfers. Section 6.3.4 documents how the tax-and-transfer system as a whole has shaped the distribution of macroeconomic growth since the end of apartheid.

6.3.1 From Factor to Pretax Income

In most advanced economies, the pension and unemployment insurance systems redistribute a substantial share of the national income the elderly and the unemployed every year, leading to very large reductions in inequality when moving from factor income to pretax income (Blanchet, Chancel, and Gethin, 2022). This is not the case in South Africa, where the private pension system and unemployment insurance through the Unemployment Insurance Fund (UIF) only benefit to a small fraction of the population.

In 2019, about 20% of the population contributed a total of about 6% of the net national income to private pension and provident funds. Meanwhile, on the income side, about 6% of adults received private pension income (about 3% of NNI). The size of the unemployment insurance system was even smaller: in 2019, about 25% of adults contributed a total of 0.5% of NNI to the fund, enabling some 2% of the population to benefit from unemployment benefits. This share stood in sharp contrast with the 29% unemployment rate, a gap that can be explained in large part by the relative stringent conditions required to benefit from unemployment insurance in South Africa (Bhorat, Goga, and Tseng, 2013).

Figure 6.4 represents the net transfers operated by the pension and UIF systems between income deciles, expressed as a share of NNI, in 2019. The private pension system appears to mostly redistribute small fractions of the national income from the top decile to the middle class. The top 1% are the main contributors to the system, losing 0.5% of NNI more in contributions than they receive in benefits, while the ninth decile (p80p90) receive a net transfer of about 0.2% of NNI. Redistribution operated through the UIF system is even smaller, with no group receiving more than 0.1% of NNI in unemployment benefits. Furthermore, while it is overall progressive, it is notably regressive at the top end of the distribution, with the top 1% contributing less than 0.02% of NNI in unemployment contributions. This can be explained both by the relatively lower share of labor income among top incomes and by the maximum cap set on UIF contributions (R178,464, or about $28,500 at PPP, in

2019), which effectively turn the contribution into a regressive tax at high wage levels.

In summary, moving from factor national income to pretax national income has almost no impact on inequality at all in South Africa (see also figure 6.10a below), given that the sums transferred are very small and mostly redistribute income from top to middle income groups.

6.3.2 The Distribution of Taxes

We now consider the distributional incidence of all taxes collected in South Africa throughout our period of interest. Figure 6.5 plots the level and composition of general government revenue, expressed as a share of national income, from 1993 to 2019. Total revenue has grown significantly in the past quarter of century, from 35% of national income in 1993 to 43% in 2019. This represents a 23% increase in the share of national income extracted from economic output every year by the general government. Of the three most important taxes in South Africa—the Personal Income Tax (PIT), the Corporate Income Tax (CIT), and the Value-Added Tax (VAT)—the CIT is the one whose revenue has grown most rapidly in relative terms, from 3.6% of NNI in 1993 to 5.8% in 2019, followed by VAT (6.3% to 7.8%) and finally by the PIT (9.7% to 11.6%). If one groups taxes in South Africa into three broad categories, direct taxes (including the PIT, the CIT, and other taxes on income and wealth), indirect taxes (including VAT, other taxes on goods and services, and taxes on international trade), and other government revenue (including other taxes, non-tax revenue, and local government revenue), direct taxes appear to represent the largest and most rapidly growing component of government revenue. Direct taxes rose from 14% to 19% of NNI between 1993 and 2019, while indirect taxes expanded from 11% to 13% and other revenue from 10% to 12%.

How have these changes in the magnitude and structure of taxation affected the distribution of taxes paid by income group? Figures 6.6a and 6.6b provides a first answer to this question by decomposing total taxes paid by pretax national income group in 1993 and 2019. In 1993, the profile of taxation was relatively flat, except for the upper-middle of the income distribution, where effective taxation was slightly higher. Nearly all deciles transferred between 20% and 35% of their pretax incomes in taxes to the government. Bottom income groups paid almost all of their taxes in indirect and local taxes. Meanwhile, the personal income tax and the corporate income tax represented the bulk of the tax burden of the top decile. It is also interesting to note that the personal income tax was regressive at the very top, which is directly due to the fact that top income groups relied heavily on non-taxable capital incomes, in particular corporate undistributed profits held through stock ownership. The corporate income tax did not compensate sufficiently for this regressive aspect of the tax system, leading the top 0.1% to pay lower taxes than the rest of top 10% earners.

Moving to 2019, we see that the increase in taxation has been almost entirely concentrated at two parts of the distribution: the very bottom and the very top. At the top, taxation is no longer regressive, mostly due to the rise in the share of corporate income tax paid by the top 0.1% (from about 9% of its pretax income in 1993 to 19% in 2019). At the bottom, the share of income paid by low-income groups in indirect and local taxes has grown substantially, with the tax burden of the third decile more than doubling. This can be explained both by the increase in total revenue collected from indirect and local taxes and by the rising gap between income and consumption among the poor in the past decades, to which we come back below. Meanwhile, the effective tax rate faced by middle income groups has barely changed, with individuals located between the median and the 90th percentile still paying less than 35% of their pretax income in taxes.

Figures 6.7a and 6.7b provide another perspective on this transformation by representing the yearly evolution of total taxes paid by top 1% and bottom 50% pretax income earners since 1993. In 1993, the top 1% faced a slightly higher effective tax rate than the bottom 50%. By 2019, the tax burden of the top 1% had increased to 45%, while that of the bottom 50% had surged to 60% of their total pretax income. The increase in top income taxation has been driven by the corporate income tax (from 9% to 13% of pretax income), but also by the personal income tax (from 9% to 16%). This latter evolution reflects both the fact that top taxable incomes have grown faster than the threshold required to enter top marginal income tax rates and the declining share of non-taxable capital income (dividends and undistributed profits) in top 1% pretax incomes. On the contrary, we see that the bottom 50% pay almost no personal income tax or corporate income tax at all, while local taxes, VAT, and excise duties have driven nearly all of the increase in their tax burden.

At this stage, let us discuss a bit further our results on the very high tax rates faced by bottom pretax income groups. It might look surprising and even unrealistic at first sight to observe such extremely high effective tax rates, given in particular that some of these rates are higher than the statutory rates of the taxes considered (for instance, the bottom 50% pay 20% of their pretax income in VAT while the statutory VAT rate is 15%). This is a mechanical result of our allocation strategy, which implies distributing indirect taxes proportionally to consumption (excluding exempted goods and the informal sector). Given that low-income groups have consumption levels that can greatly exceed their pretax incomes, the tax base on which these taxes are applied (consumption) may be substantially higher than the denominator considered for tax incidence analysis (pretax income). The presence of such a large discrepancy between the consumption and income distribution profiles, leading to extreme negative (respectively positive) savings at the bottom (respectively top) is not new (see Chancel et al., 2023; Czajka, 2017; Deaton, 1997), yet it is not fully understood.

If a large fraction of the poor are effectively consuming from their savings or from consumer debt, such tax rates may then not seem extraordinary. On the one hand, one cannot exclude that some measurement issues in household surveys (underreporting of income at the bottom of the distribution, overreporting of consumption at the bottom, or alternatively underreporting of consumption at the top) may lead to biased estimates of savings across income groups, implying an overestimation of the regressivity of indirect taxes. On the other hand, there is suggestive evidence of strongly negative and deteriorating savings rates among the poor in South Africa. According to national accounts published by the South Africa Reserve Bank, the ratio of households' saving to their disposable income has remained systematically negative since the mid-2000s, fluctuating between 0 and -2% after a sharp decline in the 1990s, so that households have, in aggregate, consumed more goods and services than their disposable income allows alone. In 2019, as much as 5.7% of the entire national income (or 8% of household disposable income) was absorbed in interest repayments by households on previously contracted loans (authors' computations using national accounts data). Chatterjee, Czajka, and Gethin (2022), combining microdata on income, assets, and debts with macrodata on households' balance sheets, estimate that the total net worth of the poorest 50% is negative, that is, the total market value of the assets they own is lower than the debts they owe. This is consistent with data from the 2008 Living Conditions Survey, in which 72% of adults, and an overwhelming share of respondents at the bottom of the income distribution, declared having "no regular savings for emergencies."

6.3.3 The Distribution of Transfers

We now analyze how government expenditure has been distributed since 1993. We focus on the main stylized facts; an extended analysis of changes in the size and progressivity of government transfers, with a particular focus on in-kind transfers, can be found in Gethin (2023c). As shown in Figure 6.8, the rise of public spending has mirrored that of revenue in the past decades: total consolidated government expenditure grew from 36% to 42% of NNI between 1993 and 2019. Even more so than in the case of taxes, this transformation has been accompanied by significant changes in the nature of government intervention. General public services and defense are the two only types of spending that have declined as a share of NNI, from 8.7% to 7.0% and from 2.6% to 1.2% respectively. Meanwhile, spending on social protection is the item that has grown the fastest, nearly doubling from 3.3% to 6.5%, followed by health, local government expenditure, public order and safety, and education.

Figure 6.9a plots the share of total transfers in grants, education, healthcare, and other public goods received by the top 10% and bottom 50% as a share of national income. Consistently with the fact that the South African government has invested a rising share of NNI in individualized transfers that primarily benefit the poor, the share of national income transferred to the bottom 50% has grown much faster than that received by the top 10%. The bottom 50% received almost 19% of national income in the form of cash and in-kind transfers in 2019, representing an increase of over 50% since 1993. Meanwhile, the share of national income redistributed to the top 10% has declined, from 14% of NNI in 1993 to 11% in 2019.

Figure 6.9b plots the cumulative growth rate of the bottom 50% before and after transfers. The rise of redistribution has generated substantial real income gains for low-income households. Bottom 50% average income growth is barely affected by the inclusion of the old age grant and the disability grant, mainly because these grants already existed in 1993 and have not increased significantly in real terms since then. In contrast, the introduction of the child support grant in 2002 and its progressive deployment over the course of the 2000s has strongly benefited the bottom 50%, whose total growth rate shifts from below 20% to over 50% when accounting for child support grants received. Finally, substantial increases in the size and progressivity of public services have further contributed to improvements in the living standards of low-income households (see Gethin, 2023c). Accounting for in-kind transfers and public goods—including in-kind social protection, education, healthcare, local government services, and other public services—bring bottom 50% real income growth to almost 100%.

6.3.4 The Overall Impact of the Tax-and-Transfer System

Our analysis of taxes and transfers has shown mixed results. On one hand, in-kind transfers have grown substantially since 1993, and this rise has primarily benefitted bottom income groups. On the other hand, the bottom 50% have faced increasing effective tax rates, driven by the rise of indirect and local taxes. Combining these two pictures, who has benefited most from the rise of South Africa's welfare state since the end of apartheid?

Figure 6.10a compares the top 1% and bottom 50% shares in terms of factor, pretax, posttax disposable, and posttax national income since 1993. South Africa's tax-and-transfer system is progressive overall and has become significantly more progressive over time. Between 1993 and 2019, the top 1% factor income share grew from 22% to 28%, while the top 1% posttax national income share first rose but then came back to its 1993 level, at about 18%. This result directly mirrors the rising tax burden of the top 1%, which has not come with greater transfers.

Turning to the bottom 50%, redistribution appears to have increasingly benefited this group, due in particular to the rising role of in-kind transfers and public goods. In 2019, moving from pretax to posttax disposable income (that is, removing all taxes but only adding back cash transfers) increases the bottom 50% share from about 3% to 5%, while moving from posttax disposable income to posttax national income (that is, adding in-kind transfers and all other government expenditure) raises it to 15%. In terms of pretax and posttax disposable income, the bottom 50% share first dropped from 1993 to the mid-2000s, before coming back to approximately the same level. In terms of posttax national income, in contrast, it declined from 11% in 1993 to 10% in 2005, but then rose steadily until reaching 15% in 2019. Rising redistribution in the form of education, healthcare, local government services, and other public goods has thus acted as a powerful equalizer since the end of apartheid, although inequality remains high even after accounting for taxes and transfers.

Figure 6.10b provides a more granular picture of redistribution in South Africa by representing the share of national income transferred by the tax-and-transfer system between income deciles in 1993 and 2019. Two results stand out. First, in 2019, all deciles within the bottom 80% were net beneficiaries, while the top 10% saw its pretax income reduced by a net total of 20% of national income. Redistribution in South Africa thus appears substantial, transferring about a fifth of the entire national income from the top decile to the rest of the population. Second, redistribution operated by the tax-and-transfer system has intensified over time. At the top, the net transfer of the top 10% grew by over 50%. Meanwhile, all deciles within the bottom 80% received significantly higher net transfers in 2019 than in 1993. The net transfer received by the bottom 50% grew from 10% to 15% of national income.

Having considered the impact of taxes and transfers on overall inequality, let us focus more specifically on the evolution of real incomes. Figure 6.11 provides a granular picture of the distribution of growth throughout the South Africa population by representing the cumulative evolution of real income by percentile between 1993 and 2019. The dramatic rise of pretax income inequality, combined with low macroeconomic growth, have implied drastically different trajectories at the top and bottom of the distribution. The top 1% has grown at the fastest pace, experiencing an almost 80% increase in average pretax income, compared to about 20% for most percentiles at the middle of the income distribution and strongly negative growth rates within the bottom 25%. The rise of redistribution, however, has more than compensated increases in pretax inequality. Removing all taxes and adding all cash and in-kind transfers from individual incomes reduces top 1% real income growth to 40%, while it raises median growth to about 45% and growth for the poorest 25% to over 100%. The rise of the South African welfare state has thus turned the distribution of economic growth since the end of apartheid from very regressive to unambiguously progressive.

In summary, our analysis of inequality and growth has revealed a striking surge in both pretax income inequality and government redistribution in South Africa since the end of the apartheid regime. This "chase between inequality and redistribution" has to some extent been won by the latter, as substantial improvements in tax progressivity at the top of the distribution and rising cash and in-kind transfers have made the final distribution of growth strongly progressive. This positive assessment should not be exaggerated, however. Even after taxes and transfers, income inequality remains exceptionally high in comparative perspective, with the bottom 50% as a whole still receiving less income than the top 1% in 2019.

6.4 The Evolution of Racial and Spatial Inequality

Our new dataset does not only cover income, taxes and transfers, it also preserves all the richness of household surveys and thus allows us to decompose inequality and redistribution by a number of sociodemographic variables. In this section, we study the evolution of income concentration along two key dimensions of South African inequality: race and geography.

6.4.1 Racial Inequality

Race has always been at the heart of economic and political conflicts since the making of the South African state. Throughout the twentieth century, inequalities between racial groups stood at unparalleled levels. These inequalities were institutionalized through the political domination of the White minority, which culminated in the apartheid regime of strict racial segregation established in 1948. Between the early twentieth century and the late 1980s, the per capita income of African South Africans thus remained stable at a level reaching less than 10% of that of the White population (Leibbrandt et al., 2010). This represents some of the most extreme inequalities between racial or sociocultural groups observed in contemporary history. By comparison, the White-Black income gap has fluctuated between 50% and 60% in the United States between the 1950s and today (Piketty, 2020).

How have the end of apartheid and the transition to democracy in the mid-1990s, rising inequality, and enhanced redistribution reshaped South Africa's historical legacy of extreme racial disparities? To answer this question, we first provide a long-run view on racial inequality in figure 6.12a by representing the evolution of the share of White and Black South Africans in top income groups since 1955. The figure combines historical tabulated tax returns collected by Alvaredo and Atkinson (2022), census data (1970, 1980, 1990), and our distributional national accounts data after 1993. Under apartheid, Whites represented over 95% of top 1% earners and over 90% of the top 10%, while the share of Black South Africans in upper income groups was nearly zero. A remarkable transformation in the composition of top incomes has taken place since the early 1990s: the share of Africans in the top 10% jumped from 2% in 1980 to 15% in 1990-1994, and then increased monotonically until reaching about 45% in 2019. A similar evolution occurred within the top 1%, although racial inequalities continue to be higher in the top 1% than in the top 10%.

Figure 6.12b turns to the evolution of the overall White-to-Black income ratio, focusing on the role played by changes at the top of the distribution in the decline of racial inequality. As shown by the bottom line of the figure, White South Africans' average factor income was about 14 times higher than Black South Africans' in the early 1990s. This ratio remained stable until the 2010s, before declining to 8 in 2015-2019. However, the picture looks very different if one excludes top Black earners from the analysis: excluding Black earners belonging to the top 1% leads to a decline in the gap from 14 to 11, while removing all those in the top 5% of the factor income distribution leads to an even smaller change, from 15 to 13. If one excludes completely all Black South Africans belonging to the top 10% from the analysis, then the White-Black income ratio appears to have remained constant, at about 17. In other words, racial inequalities have decreased in South Africa, but this decrease is mostly attributable to the emergence of a new Black elite, who has occupied a growing share of the top 10% of the income distribution.

Relatedly, figure 6.12c shows how factor income growth has been distributed within each population group from 1993 to 2019. Two results stand out. First, inequality has risen dramatically within each group. The top 10% of Asian, Black, Coloured, and White earners saw their average factor income grow by 70 to 220%, compared to growth rates of approximately 0 to 40% for the bottom 50%. Second, the average factor income of Black South Africans grew substantially faster than that of other groups: it rose by 150% over the period considered, compared to 50-80% for Asians, Coloureds, and Whites. However, much of this dynamic was driven by differential trajectories at the top of the distribution: the average factor income of top 10% Black earners increased much faster than that of the top 10% of other groups, while growth rates of the middle 40% and bottom 50% of each group are of the same order of magnitude.

Figure 6.13a provides more detail on the contemporary structure of racial inequality in South Africa by decomposing the White-Black gap by economic concept. Two results stand out. First, racial inequality remains substantially larger in terms of wealth than in terms of income or consumption: the White-Black income ratio reaches almost 14 in terms of personal net wealth versus 7-9 in terms of consumption, factor income, and pretax income. Second, the tax-and-transfer system strongly reduces racial inequalities, in particular in-kind transfers, yet posttax income gaps remain high. The ratio decreases from 9 to 8 when moving from pretax to posttax disposable income, and drops to 5 in terms of posttax national income. Taxes and transfers thus significantly reduce racial inequalities, but they do little to change the overall relationship between race and economic status. As shown in figure 6.13b, which represents the racial composition of posttax national income groups in 2019, White earners continue to be massively overrepresented at the top end of the distribution even after accounting for taxes and transfers. In 2019, they represented over 70% of the top 1% compared to less than 5% of all posttax income percentiles within the poorest half of the population. Put differently, taxes and transfers do not significantly alter the racial dimension of economic inequalities in South Africa. They primarily reduce inequality between population groups by reducing inequality between income groups, without substantially affecting their racial composition.

6.4.2 Spatial Inequality

To conclude this paper, we consider another dimension of inequality: geography. How large are spatial inequalities in South Africa and how are they affected by the tax-and-transfer system?

Figure 6.14a compares the relative average incomes of South Africa's provinces before and after accounting for government taxes and transfers. Regional inequalities are significant in South Africa, and clearly separate the country into two groups: that of the richer provinces of Western Cape and Gauteng, whose average factor incomes exceed the average national income by 60-80%, and the rest of the country, with incomes falling between 40% and 60% of the national average. These regional disparities are larger, for instance, than inequalities between European countries, and substantially wider than differences in average incomes across US States (see Blanchet, Chancel, and Gethin (2022)). In line with our finding on the overall progressivity of the tax-and-transfer system, we find that the government also operates redistribution between provinces, although only to a moderate extent. Western Cape and Gauteng are net contributors, while all other provinces are net beneficiaries. The provinces that benefit most from the tax-and-transfer system are Limpopo, KwaZulu-Natal, and Eastern Cape, whose relative income increases by 30-50% after accounting for taxes and transfers. Meanwhile, Gauteng sees its relative average income decrease by over 15% between factor and posttax income.

In addition to regional inequality, the rural-urban income gap has been found to be significant in many countries throughout the world, often determining a substantial share of overall income inequality, migration patterns, and human capital accumulation (Young, 2013). South Africa is no exception to this general pattern, yet we find that rural-urban disparities have risen significantly since 1993. The average factor income of urban earners was almost 6 times higher than that of rural areas in 2015-2019, compared to about 4 in 1993-1994 (see figure 6.14b). However, the rise of government redistribution has prevented posttax inequality from increasing: the rural-urban gap grew from 3.6 to 4.8 in terms of posttax disposable income, while it stagnated at about 2.8 in terms of posttax national income. Growing spending on in-kind transfers and public services has thus disproportionately benefited rural areas since the end of apartheid, fully compensating the rise of rural-urban pretax income inequality.

6.5 Conclusion

By most contemporary measures, South Africa continues to stand out as the most unequal country in the world, yet this paper has documented dramatic changes in the structure of these inequalities since the end of the apartheid regime in the 1990s. The surge of pretax income inequality has implied radically different growth trajectories across income groups. The top 1% experienced an 80% increase in their average pretax income, while that of the bottom 20% declined. However, increasing government redistribution in the form of progressive taxation, cash grants, and public services has overcompensated the rise of pretax inequality, generating large real income gains for low-income households mostly at the expense of the richest decile. That being said, the expansion of South Africa's welfare state has been largely insufficient to significantly alter the extreme disparities inherited from a century of racial discrimination and oppression. The share of posttax income accruing to the richest 1% was about the same in 2019 as in 1993, while the bottom 50% only received 15% of national income, even after accounting for all cash transfers, in-kind transfers, and public goods received. While racial inequalities have declined, this decrease has been entirely driven by the income gains of a few Black earners at the top end of the distribution, thereby excluding the majority of the poor. These inequalities continue to be much larger in terms of wealth than in terms of income and have not been substantially affected by the growing progressivity of the tax-and-transfer system.

We see at least two avenues for future research. First, this paper has demonstrated the crucial importance of allocating indirect taxes and in-kind transfers when estimating the impact of taxes and transfers on poverty, inequality, and the distribution of economic growth. Yet, while we believe we have made significant advances in facing this challenge, the data sources at our disposal to properly understand who pays government taxes, and who gains from spending in health, education, and other collective expenditure remain largely unsatisfactory. Who benefits from investments in infrastructure development, industrial policy, or housing programs at the macro level, and how has this changed over time? What kinds of government spending most effectively accrue to low-income groups and how? These are important questions on which our knowledge remains all too limited.

Second, while our results have shed new light on the interactions between taxes, transfers, and the distribution of growth, much remains to be understood when it comes to the behavioral and general equilibrium mechanisms underlying the persistence of extreme economic inequalities and the ability of government redistribution to reduce these inequalities in the long run. To what extent can progressive taxation contribute to limiting income and wealth concentration beyond their immediate impact on top pretax incomes? Can cash and in-kind transfers truly reduce poverty and inequality beyond the short-term relief they provide, especially in countries where the poor are highly leveraged and vulnerable to transitory income shocks as in South Africa? To what extent taxes and transfers shape future pretax incomes? Answering these questions requires going beyond the descriptive analysis conducted in this paper and modelling the joint relationships between income, wealth, savings, and household debt (for recent fruitful attempts, see for instance Blanchet, 2022a; Mian, Straub, and Sufi, 2021). We hope that our new database and the stylized facts presented in this paper will contribute to research in these multiple directions.

Bibliography

Alvaredo, Facundo and Anthony B. Atkinson (2022). “Top incomes in South Africa in the twentieth century”. In: *Cliometrica* 16, pp. 477–546.

Alvaredo, Facundo et al. (2018). *World Inequality Report 2018*. Harvard University Press.

Bachas, Pierre, Lucie Gadenne, and Anders Jensen (2022). “Informality, Consumption Taxes and Redistribution”. NBER Working Paper 27429.

Bassier, I. and I. Woolard (2018). “Exclusive Growth: Rapidly Increasing Top Incomes Amidst Low National Growth in South Africa”. REDI3x3 Working Paper 47.

Bassier, I. et al. (2020). “Locked Down and Locked Out: Repurposing Social Assistance as Emergency Relief to Informal Workers”. SALDRU Working Paper 261.

Bhorat, H., S. Goga, and D. Tseng (2013). “Unemployment Insurance in South Africa: A Descriptive Overview of Claimants and Claims”. University of Cape Town Working Paper.

Blanchet, T. (2022a). “Uncovering the Dynamics of the Wealth Distribution”. Working Paper.

Blanchet, Thomas, Lucas Chancel, and Amory Gethin (2022). “Why Is Europe More Equal Than the United States?” In: *American Economic Journal: Applied Economics* 14.4, pp. 1–40.

Blanchet, Thomas, Ignacio Flores, and Marc Morgan (2022). “The weight of the rich: improving surveys using tax data”. In: *The Journal of Economic Inequality* 20.1, pp. 119–150.

Blanchet, Thomas, Juliette Fournier, and Thomas Piketty (Apr. 16, 2021). “Generalized Pareto Curves: Theory and Applications”. In: *Review of Income and Wealth*, roiw.12510.

Blanchet, Thomas et al. (2021). *Distributional National Accounts Guidelines: Methods and Concepts Used in the World Inequality Database*. Tech. rep.

Bozio, Antoine et al. (2022). “Predistribution vs. Redistribution: Evidence from France and the U.S.” WID.world Working Paper.

Chancel, Lucas et al. (2023). “Income inequality in Africa, 1990–2019: Measurement, patterns, determinants”. In: *World Development* 163.1, pp. 106–162.

Chatterjee, Aroop, Léo Czajka, and Amory Gethin (2022). “Wealth Inequality in South Africa, 1993-2017”. In: *World Bank Economic Review* 36.1, pp. 19–36.

Czajka, Léo (2017). “Income Inequality in Côte d’Ivoire: 1985-2014”. In: *WID.world Working Paper 2017/8.*

Deaton, A. (1997). *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy.* World Bank Publications.

Duflo, E. (2000). “Child Health and Household Resources in South Africa: Evidence from the Old Age Pension Program”. In: *American Economic Review* 90.2, pp. 393–398.

Garbinti, B., J. Goupille-Lebret, and T. Piketty (2018). “Income inequality in France, 1900-2014: Evidence from Distributional National Accounts (DINA)”. In: *Journal of Public Economics* 162.1. Dataset, pp. 63–77.

Gethin, Amory (2023c). “Who Benefits from Public Goods? Public Goods and Inequality in Post-Apartheid South Africa, 1993-2019”. Working Paper.

Inchauste, G. et al. (2015). “The Distributional Impact of Fiscal Policy in South Africa”. World Bank Policy Research Working Paper.

Leibbrandt, M. et al. (2010). *Trends in South African Income Distribution and Poverty since the Fall of Apartheid.* OECD Social, Employment and Migration Working Papers 101.

Maboshe, M. and I. Woolard (2018). “Revisiting the Impact of Direct Taxes and Transfers on Poverty and Inequality in South Africa”. WIDER Working Paper 2018/79.

Mian, A., L. Straub, and A. Sufi (2021). “Indebted Demand”. In: *The Quarterly Journal of Economics.*

Piketty, Thomas (2020). *Capital and Ideology.* Harvard University Press.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman (2018). “Distributional National Accounts: Methods and Estimates for the United States”. In: *Quarterly Journal of Economics* 133.2, pp. 553–609.

Piketty, Thomas, Li Yang, and Gabriel Zucman (2019). “Capital Accumulation, Private Property, and Rising Inequality in China, 1978-2015”. In: *American Economic Review* 109.7, pp. 2469–2496.

Saez, Emmanuel and Gabriel Zucman (2016). “Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data”. In: *Quarterly Journal of Economics* 131.May, pp. 519–578.

South African Reserve Bank (2019). *Tax Chronology of South Africa: 1979–2019*. Technical Report.

Tondini, A. (2021). "The Lasting Labor Market Effects of Cash Transfers: Evidence from South Africa's Child Support Grant". Working Paper.

World Bank (2020b). *South Africa*. Poverty and Equity Brief.

Young, A. (2013). "Inequality, the Urban-Rural Gap, and Migration". In: *The Quarterly Journal of Economics* 128.4, pp. 1727–1785.

Figure 6.1: Average national income per capita, 1993-2019

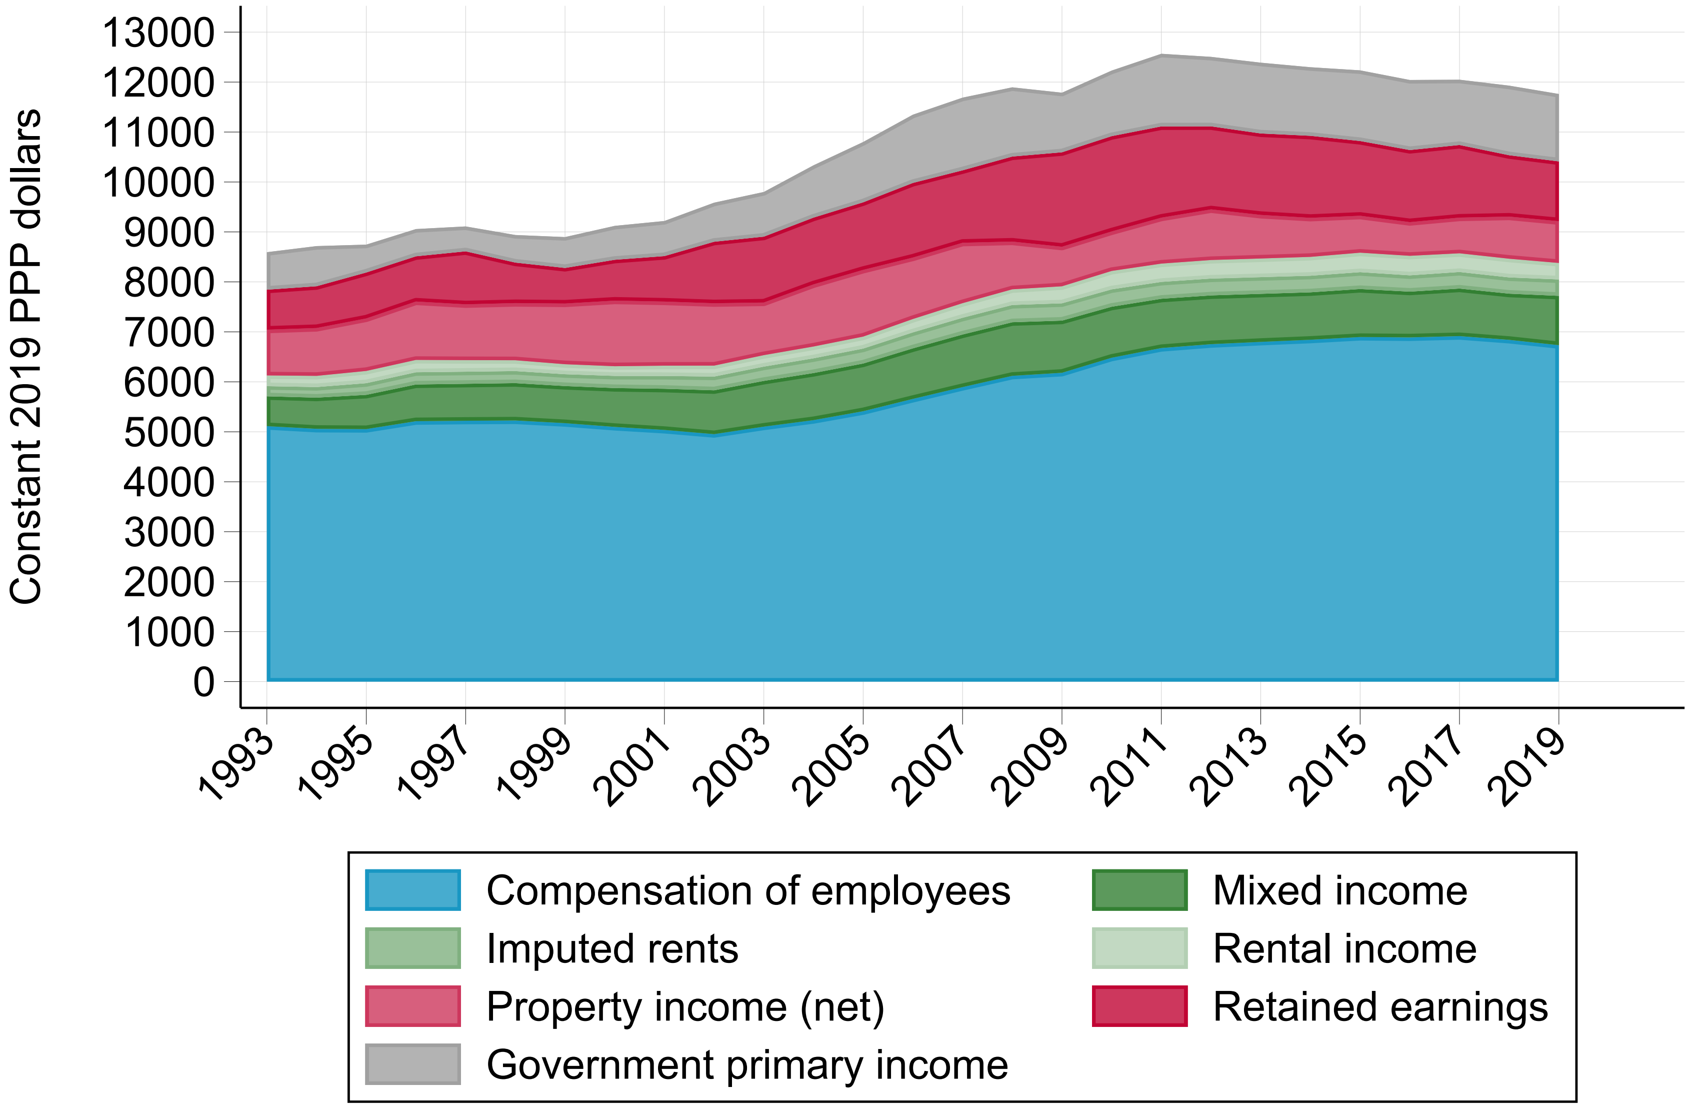

Notes. Authors' computations using national accounts series from the South African Reserve Bank Quarterly Bulletin.

Figure 6.2: The distribution of factor national income, 1993-2019

(a) South Africa in comparative perspective: top 10% pretax income share

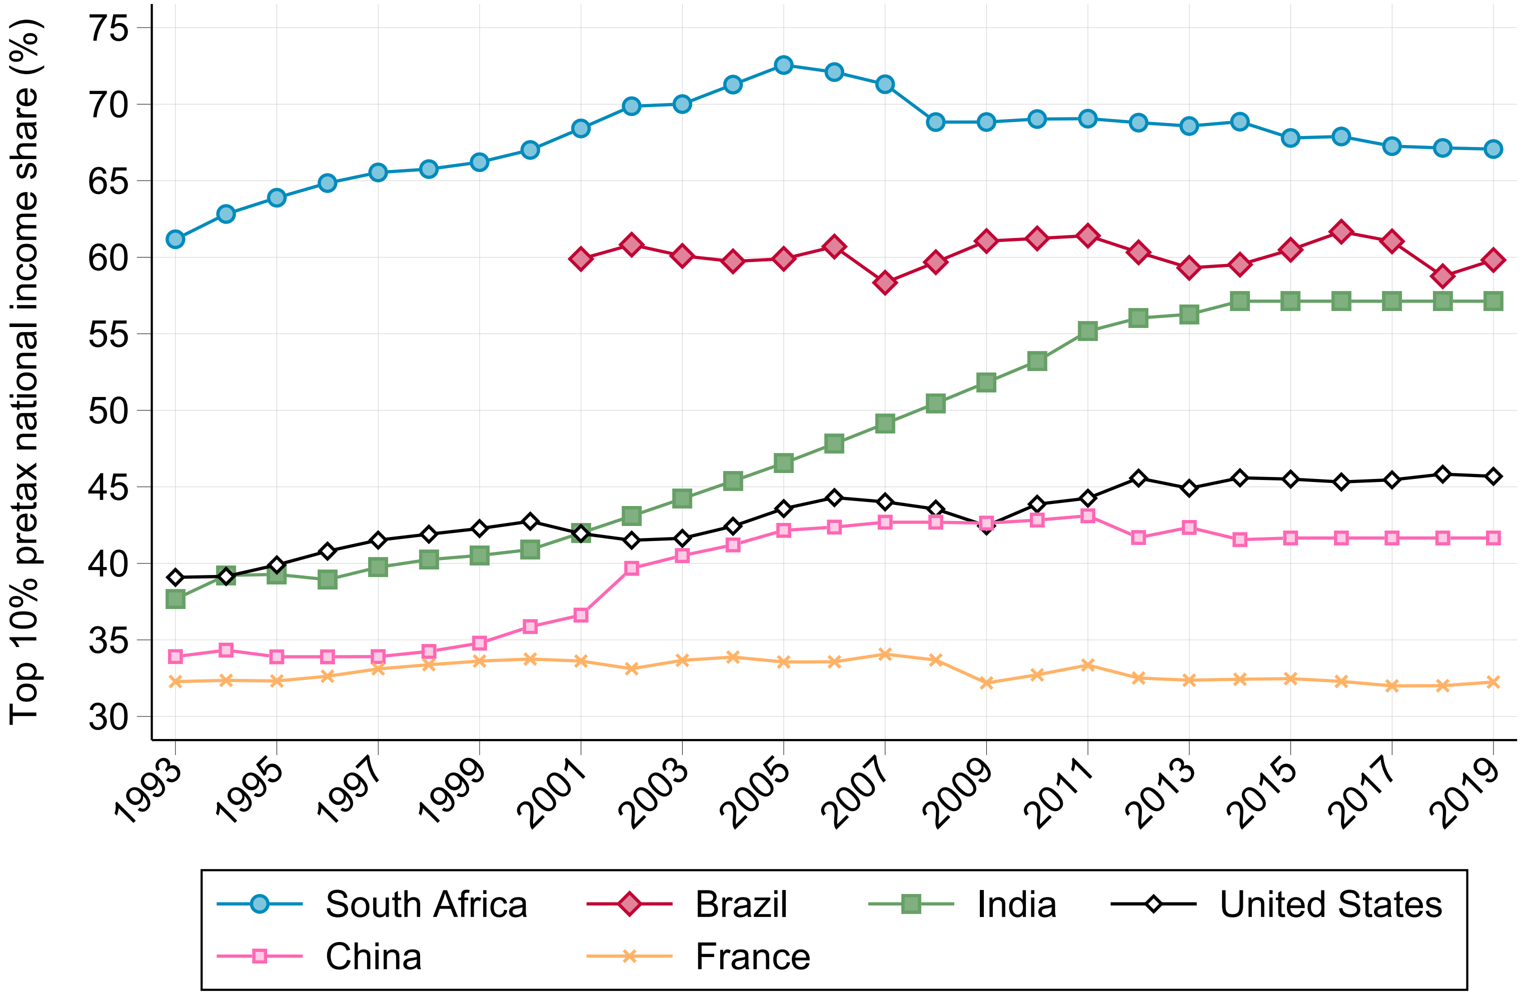

Notes. Authors' computations combining survey, tax, and national accounts data (South Africa); World Inequality Database (other countries).

Figure 6.2: The distribution of factor national income, 1993-2019

(b) Cumulated income growth by factor income group

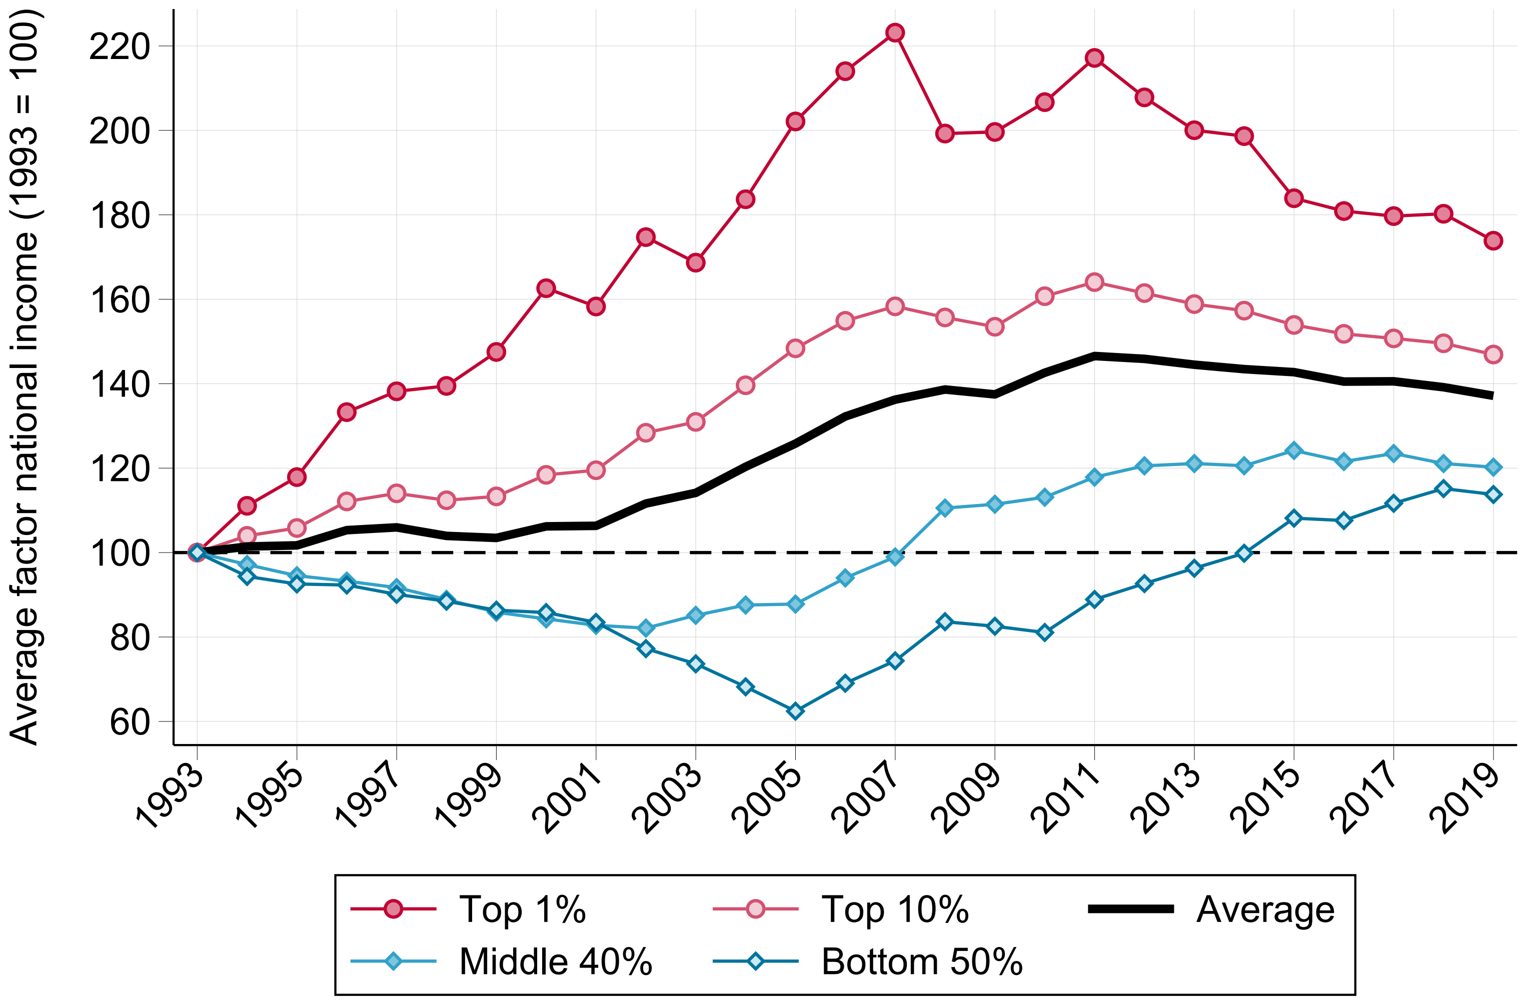

Notes. Authors' computations combining survey, tax, and national accounts data.

Figure 6.3: Decomposing top factor income inequality

(a) Top 1% income share: labor versus capital

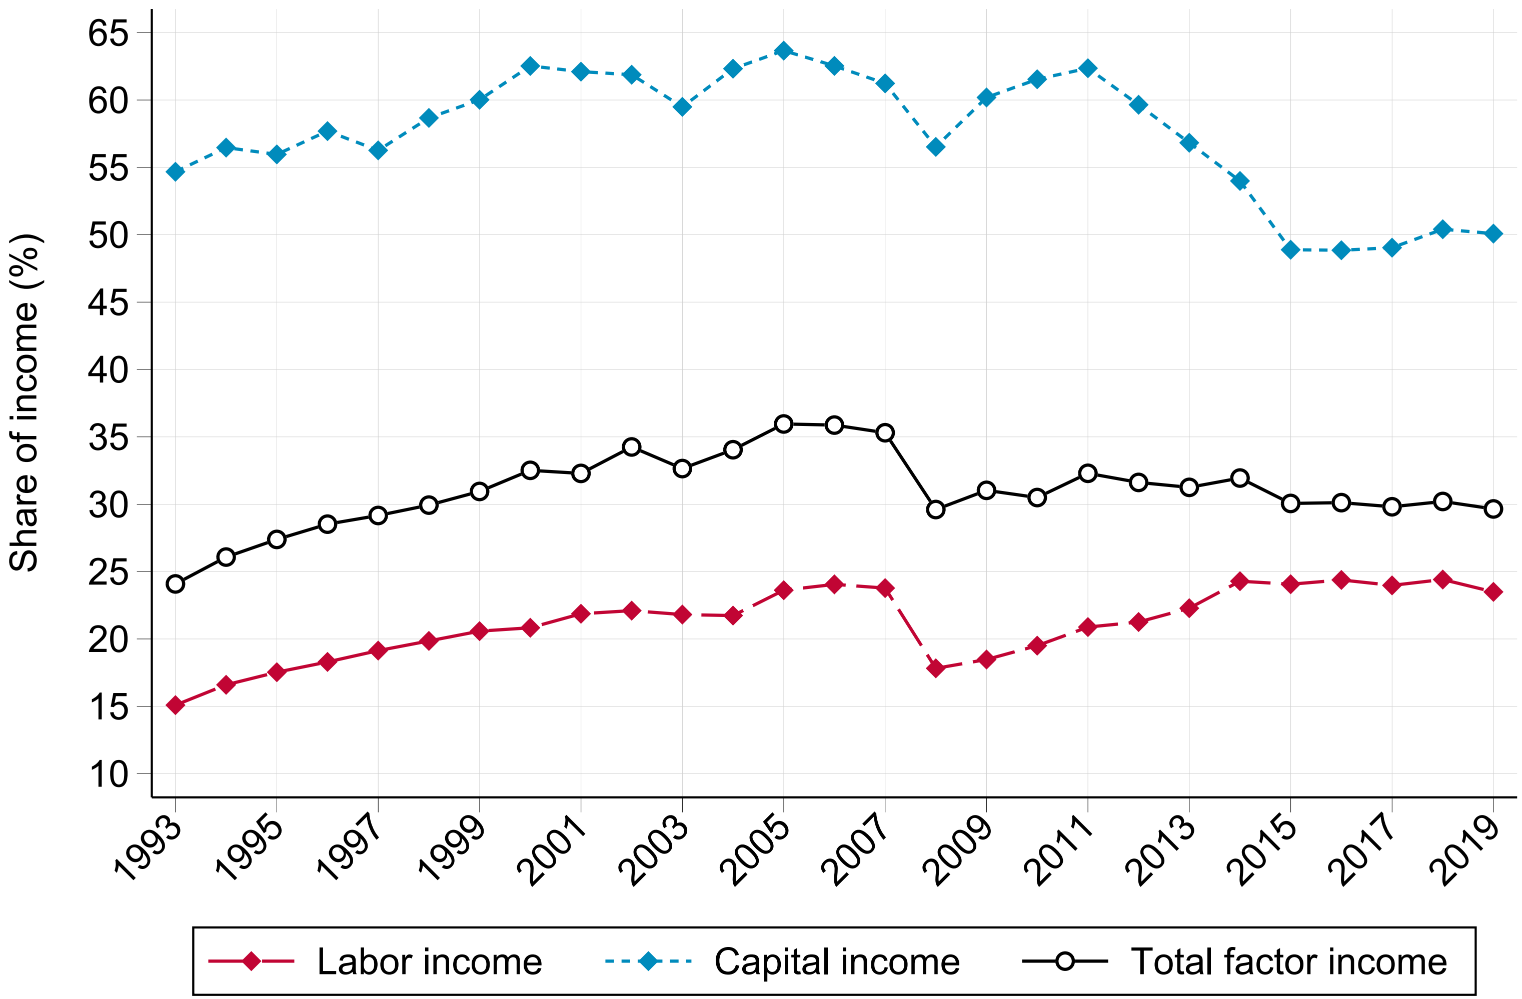

Notes. Authors' computations combining survey, tax, and national accounts data. Labor income is defined as the sum of compensation of employees and 70% of mixed income. Capital income is defined as the sum of 30% of mixed income, property income (rental income, interest, dividends, and other property income), and the private share of corporate undistributed profits.

Figure 6.3: Decomposing top factor income inequality

(b) Composition of top 1% factor income

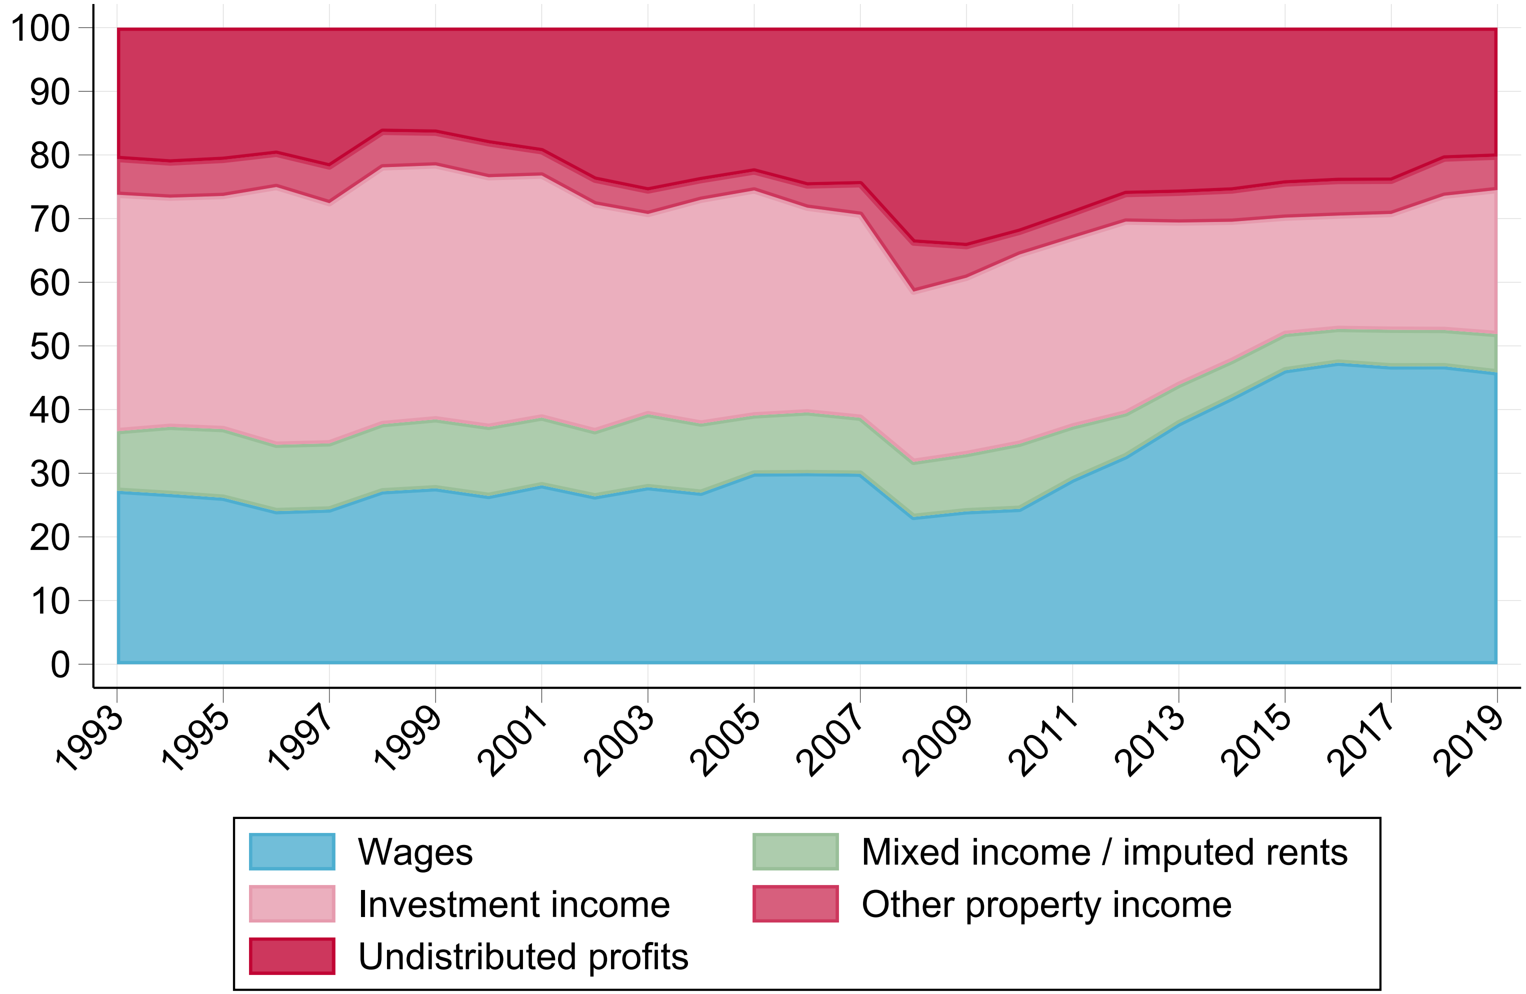

Notes. Authors' computations combining survey, tax, and national accounts data. The figure represents the composition of the factor income of top 1% earners.

Figure 6.4: From factor to pretax income: net transfers operated between factor income groups by the pension and unemployment insurance systems

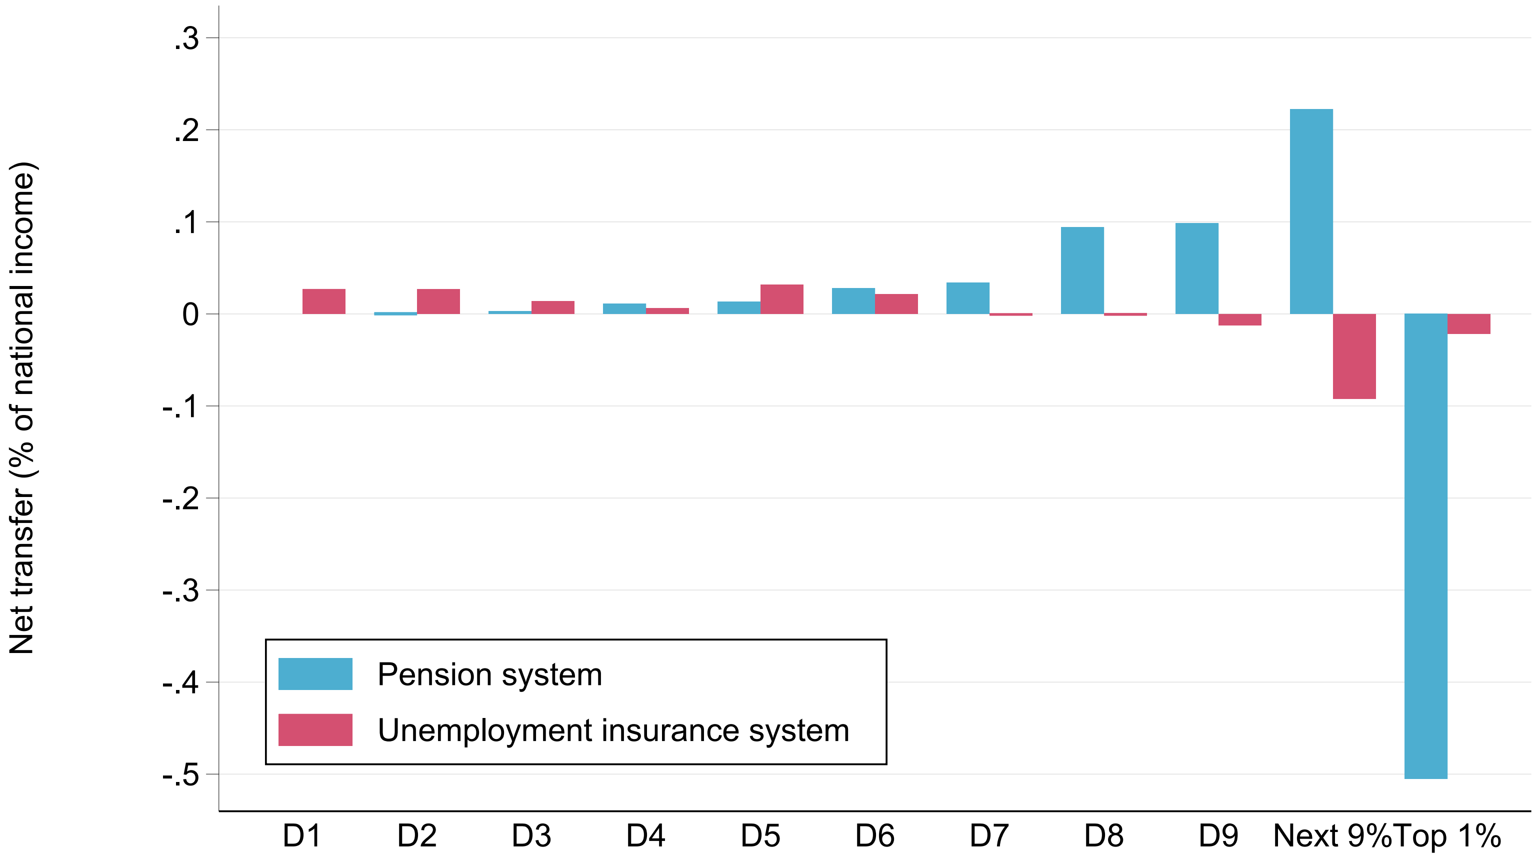

Notes. Authors' computations combining survey, tax, and national accounts data. The figure represents the net transfers received or paid by factor income group through the pension and unemployment insurance systems (that is, the difference between total benefits received and total contributions paid), expressed as a share of national income.

Figure 6.5: Government revenue in South Africa

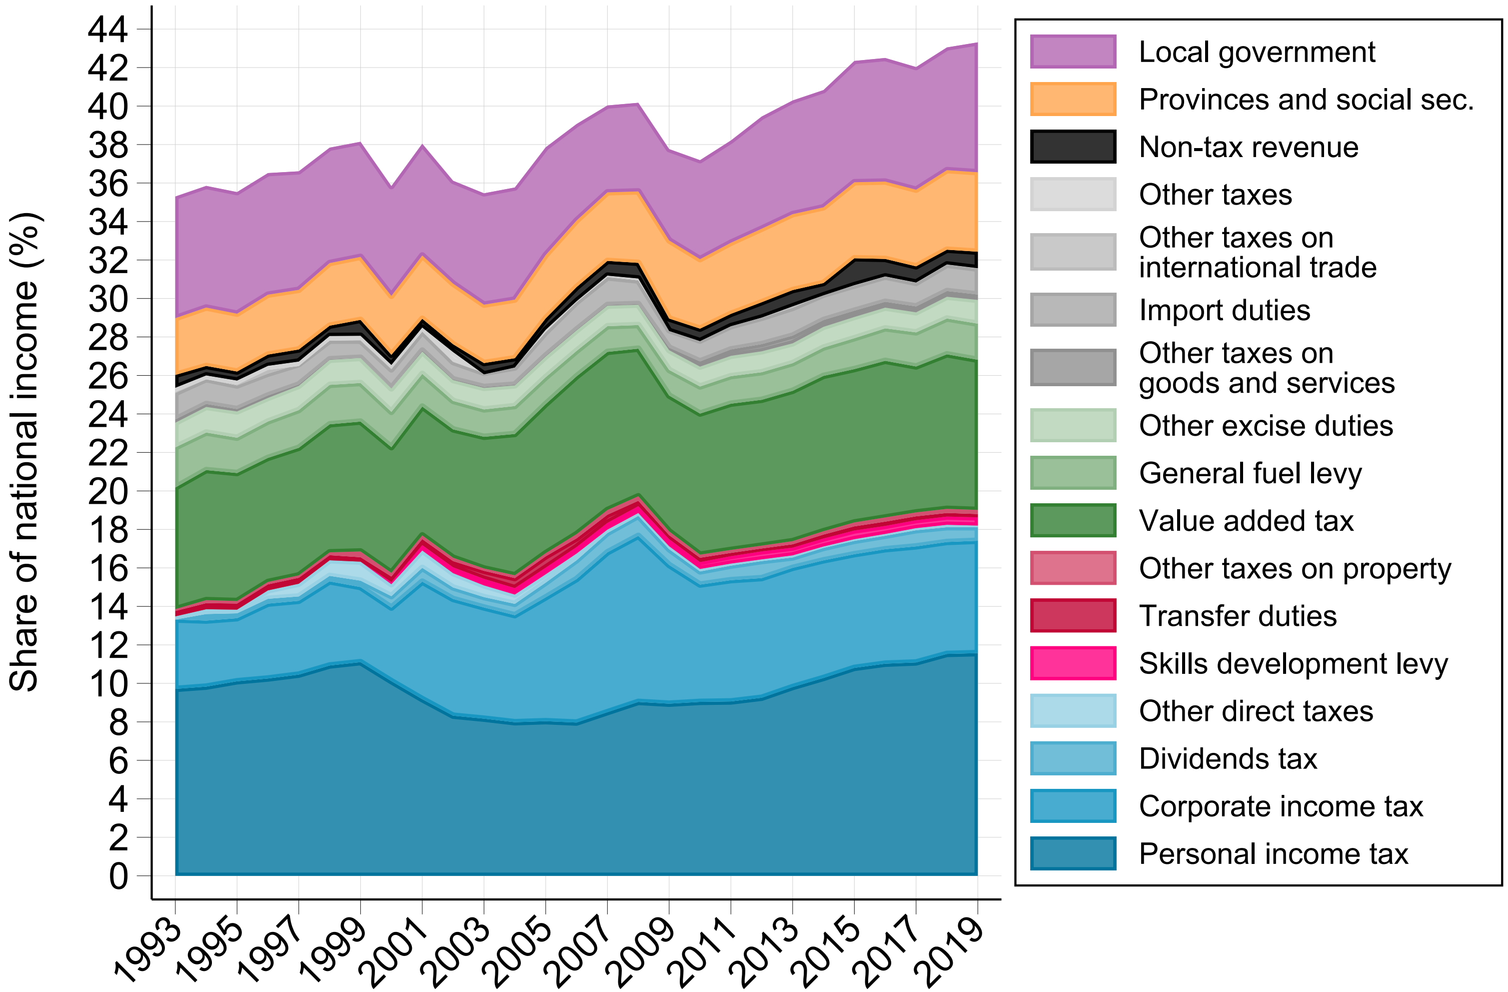

Notes. Authors' computations combining national accounts series from the South African Reserve Bank Quarterly Bulletin with government budget data collected from Treasury National Budget Reports.

Figure 6.6: Taxes paid by pretax national income group: 1993 versus 2019

(a) 1993

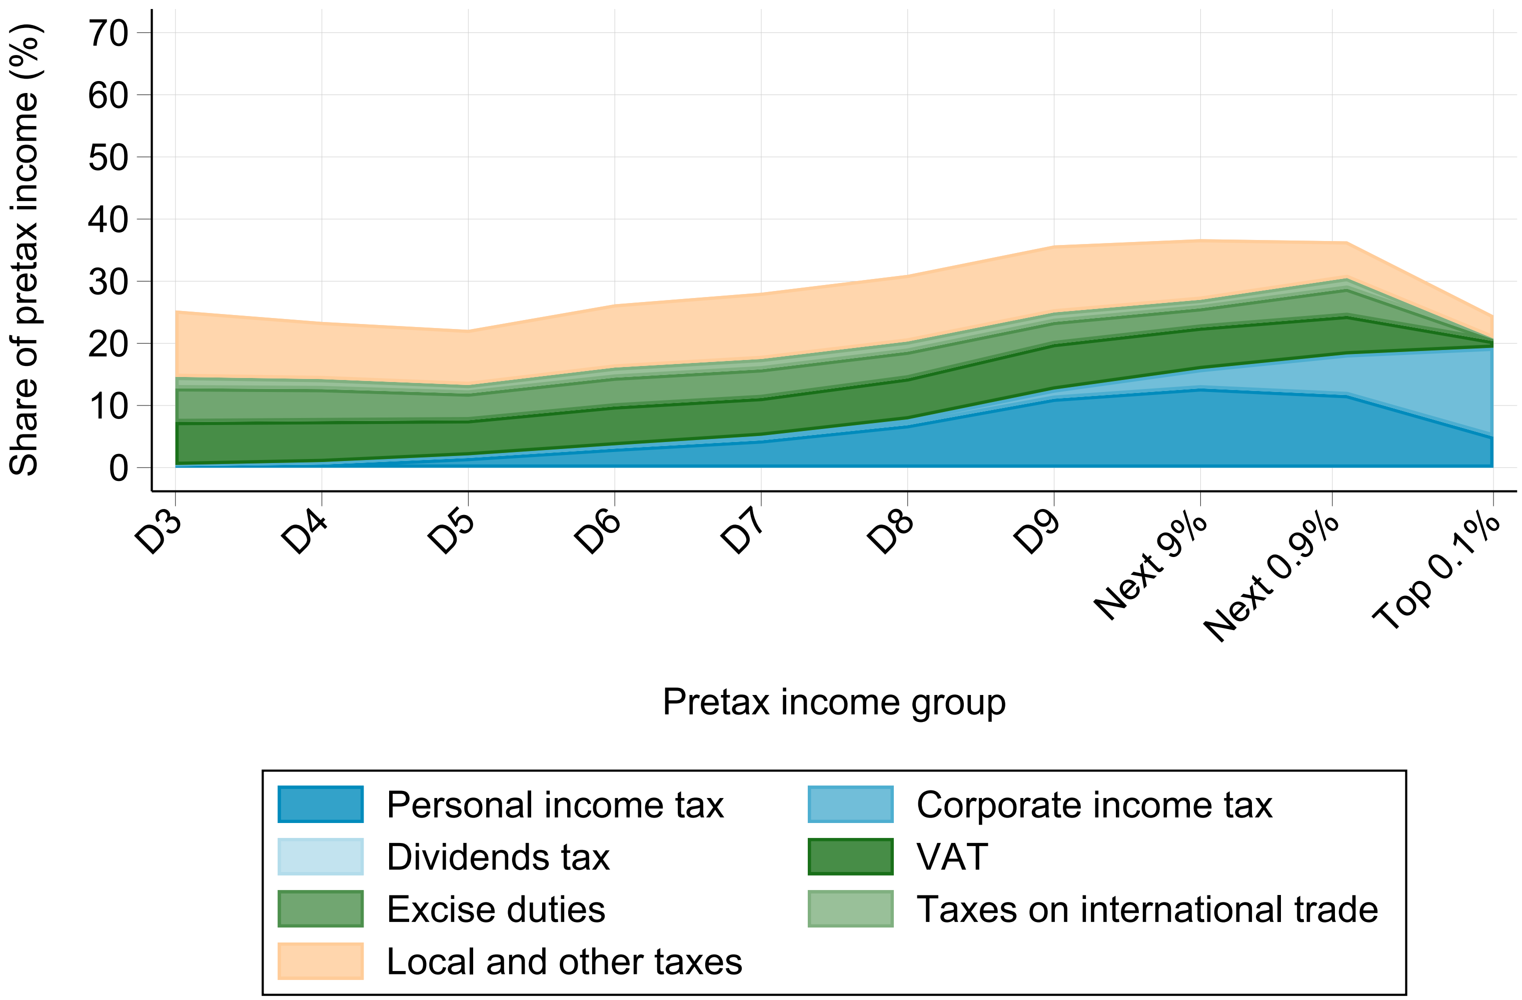

Notes. Authors' computations combining survey, tax, and national accounts data. The figure represents the effective tax rate faced by pretax income group in South Africa in 1993, expressed as a share of pretax income.

Figure 6.6: Taxes paid by pretax national income group: 1993 versus 2019

(b) 2019

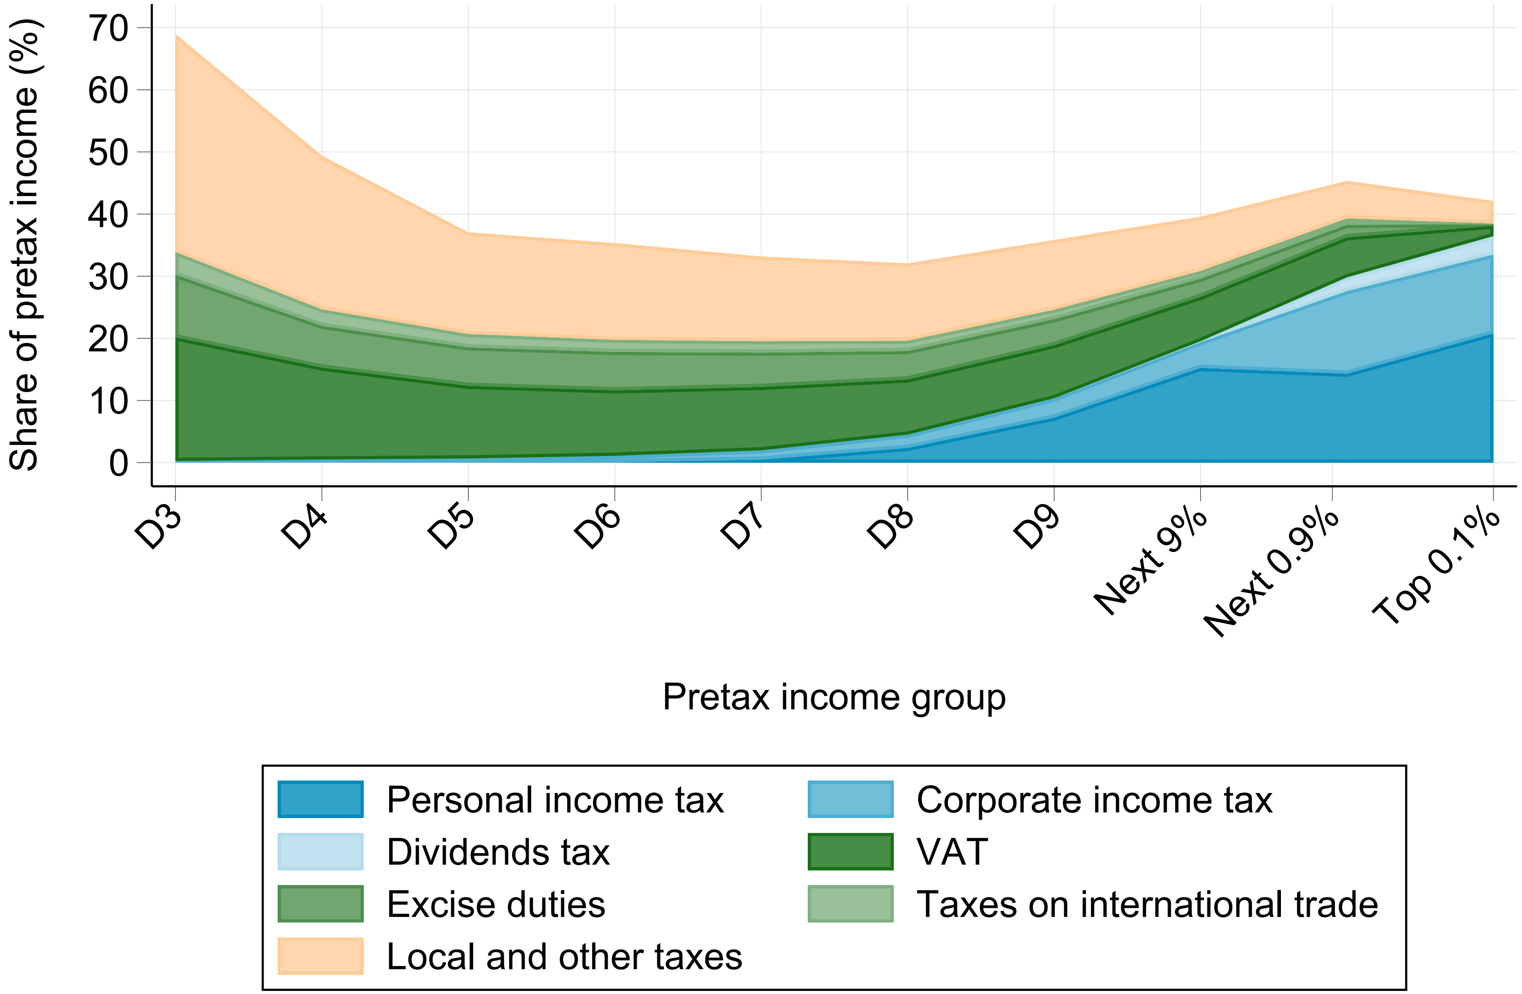

Notes. Authors' computations combining survey, tax, and national accounts data. The figure represents the effective tax rate faced by pretax income group in South Africa in 2019, expressed as a share of pretax income.

Figure 6.7: Taxes paid by the top 1% and the bottom 50%

(a) Top 1%

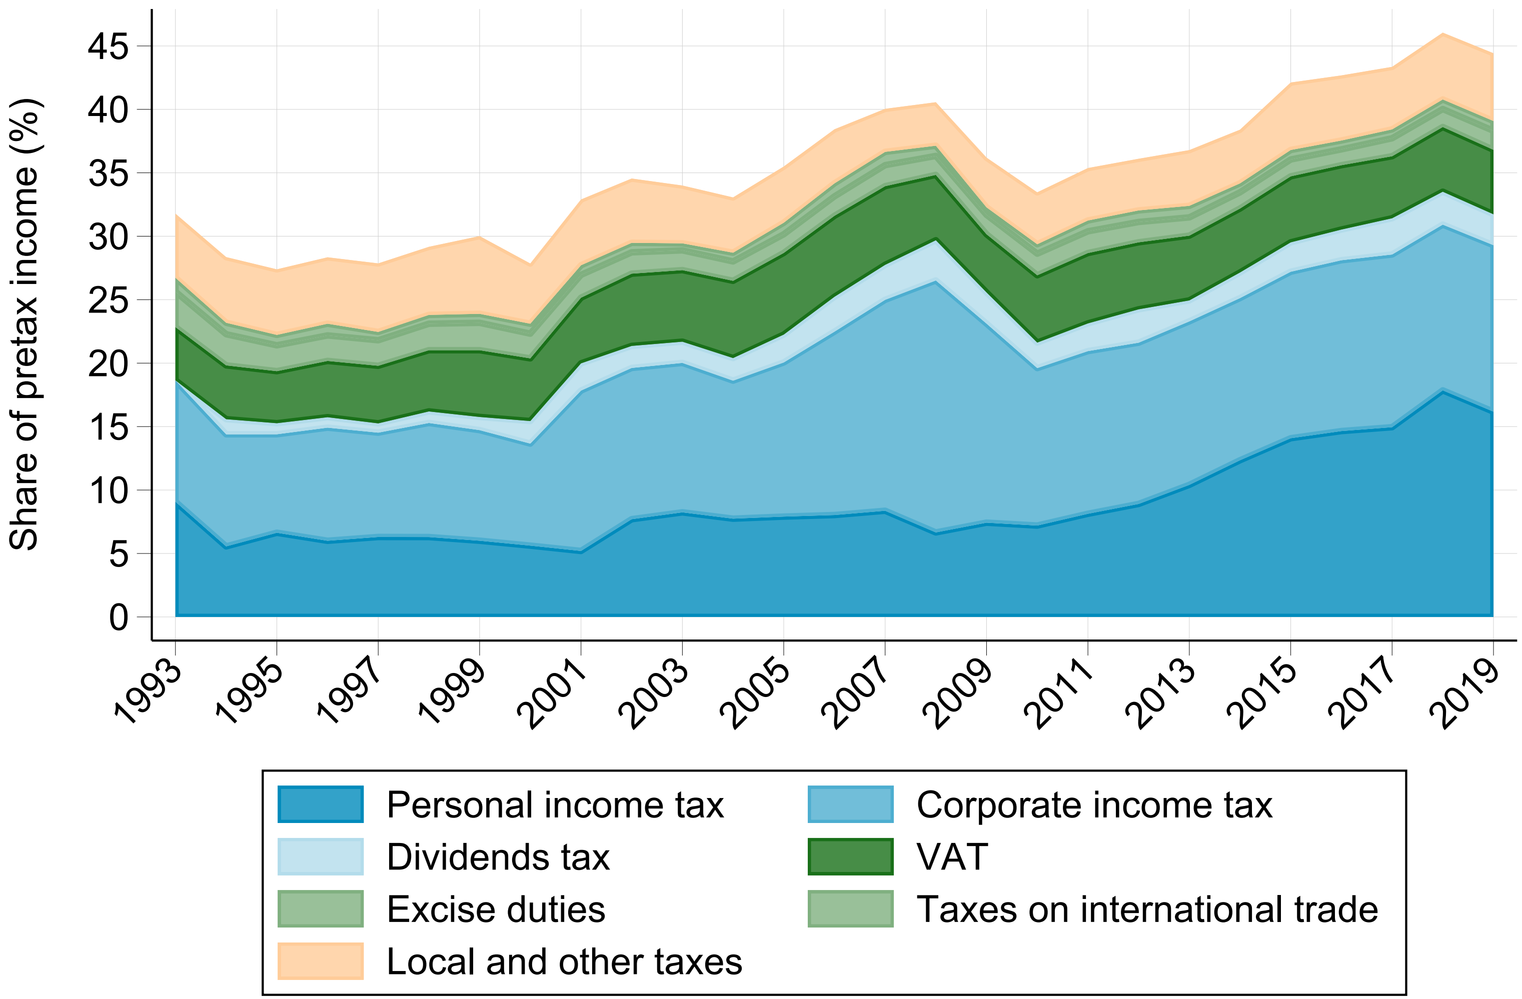

Notes. Authors' computations combining survey, tax, and national accounts data. The figure represents the effective tax rates faced by top 1% pretax income earners (panel A) and bottom 50% pretax income earners (panel B), expressed as a share of pretax income.

Figure 6.7: Taxes paid by the top 1% and the bottom 50%

(b) Bottom 50%

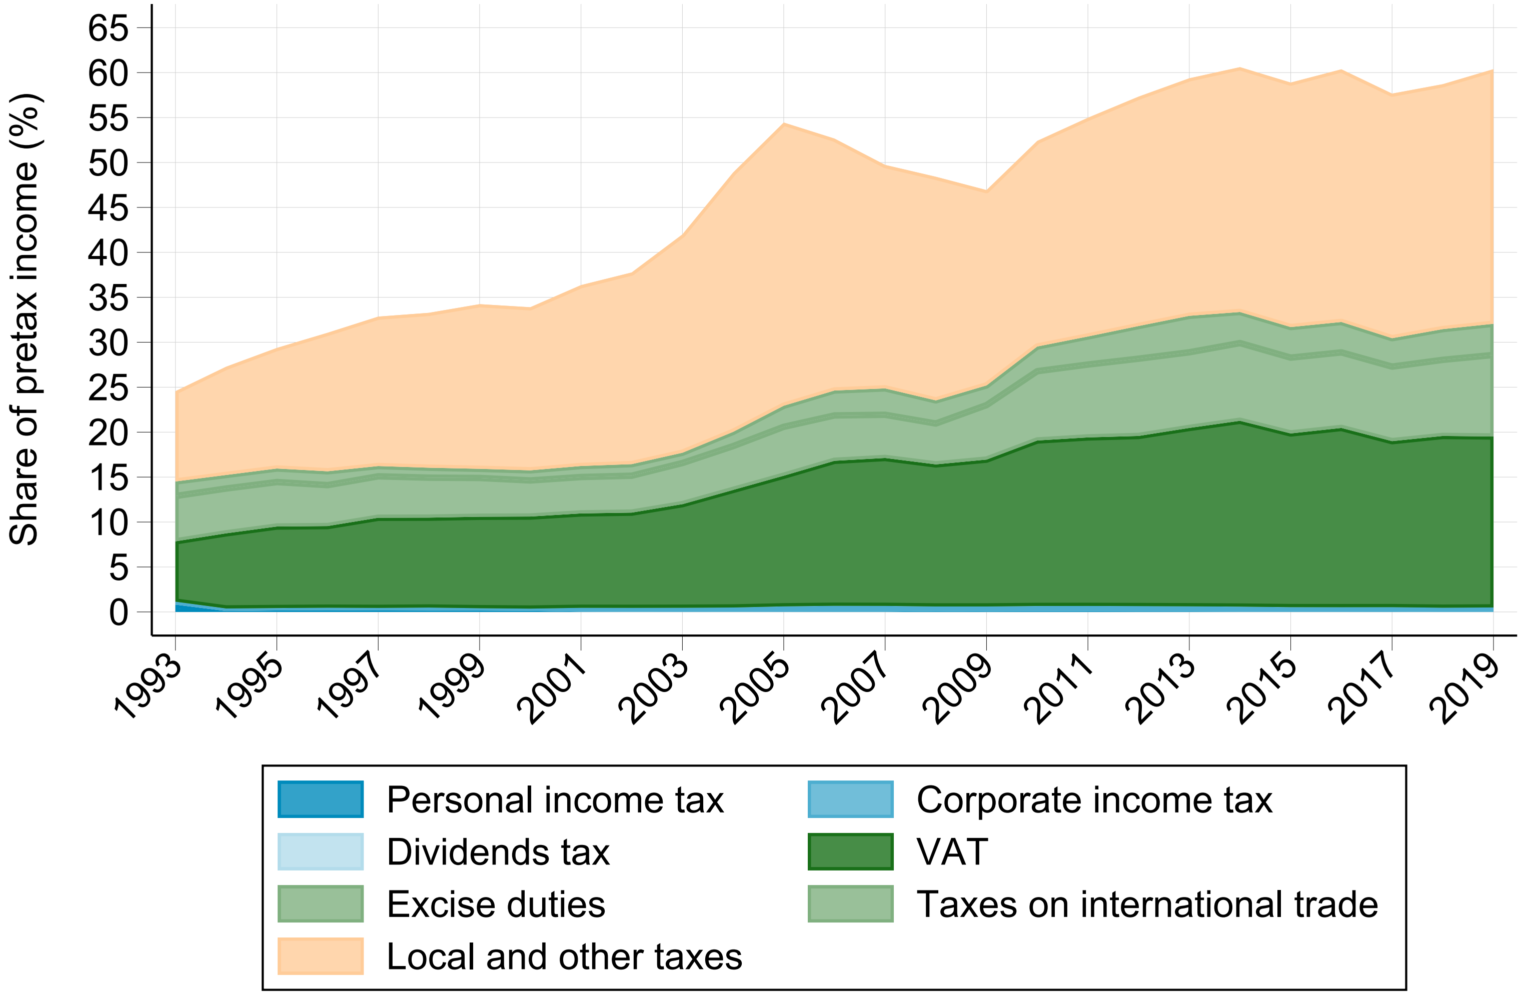

Notes. Authors' computations combining survey, tax, and national accounts data. The figure represents the effective tax rates faced by top 1% pretax income earners (panel A) and bottom 50% pretax income earners (panel B), expressed as a share of pretax income.

Figure 6.8: Government expenditure in South Africa

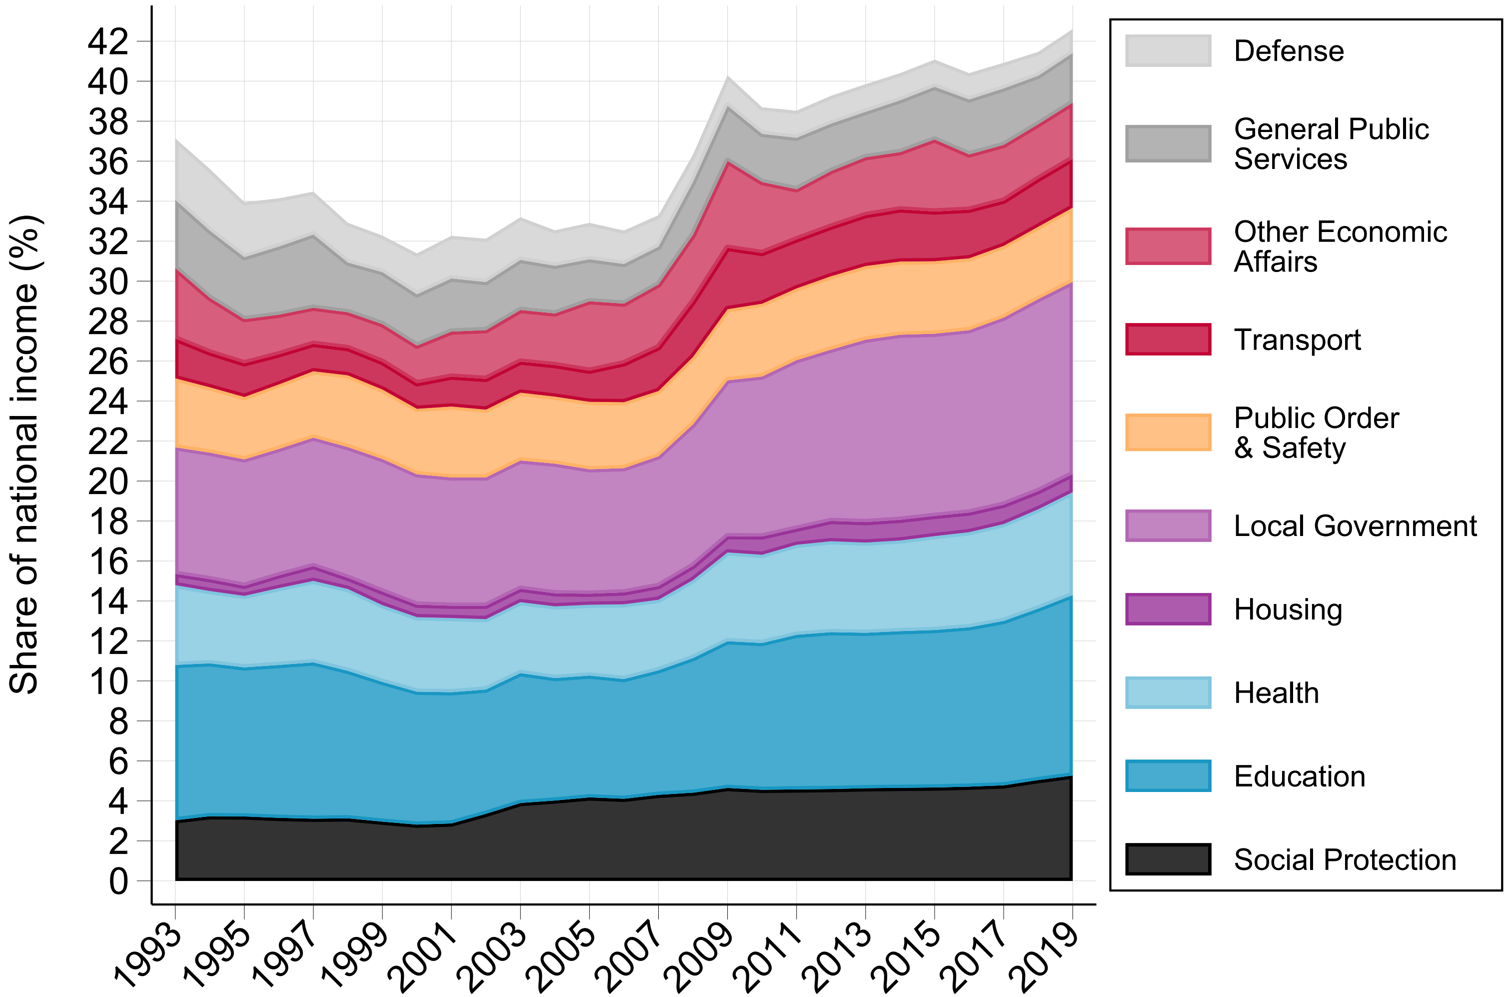

Notes. Authors' computations combining national accounts series from the South African Reserve Bank Quarterly Bulletin with government budget data collected from Treasury National Budget Reports.

Figure 6.9: The rise of social transfers

(a) Total individualized transfers received by pretax income group

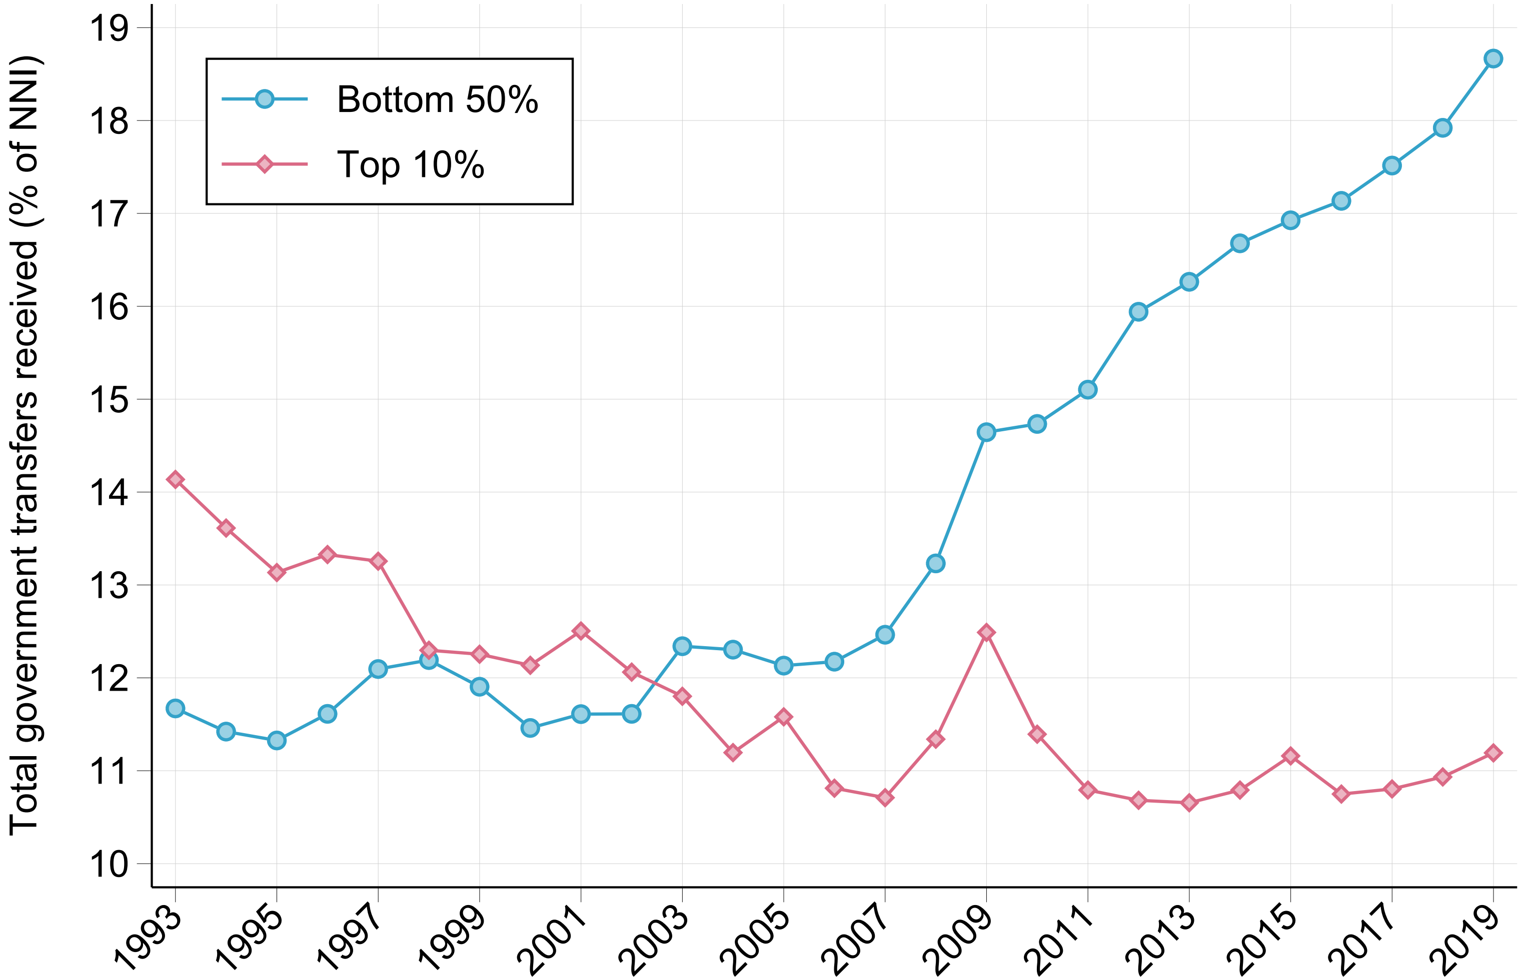

Notes. Authors' computations combining survey, tax, and national accounts data. The figure represents the total individualized transfers (social protection, education, and health transfers) received by bottom 50%, middle 40%, and top 10% pretax income earners, expressed as a share of national income.

Figure 6.9: The rise of social transfers

(b) Bottom 50% average income growth, before and after transfers

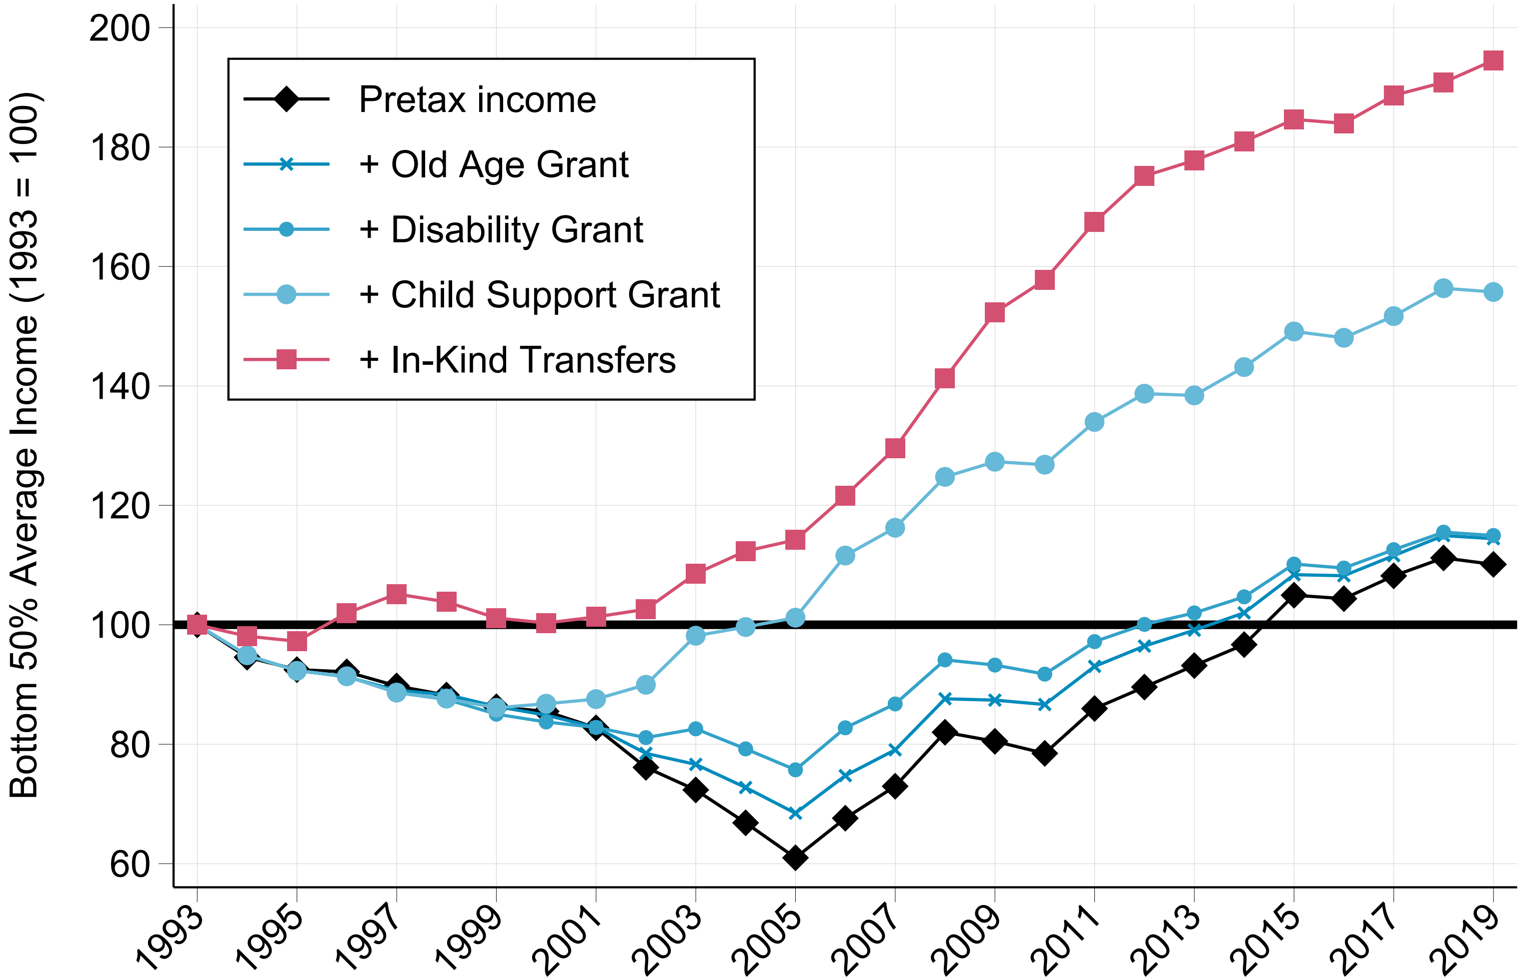

Notes. Authors' computations combining survey, tax, and national accounts data. The figure represents the cumulated income growth rate of bottom 50% earners between 1993 and 2019, before and after adding specific social transfers to individual pretax incomes.

Figure 6.10: The overall impact of taxes and transfers on inequality

(a) Top 1% versus bottom 50%: from factor to posttax national income

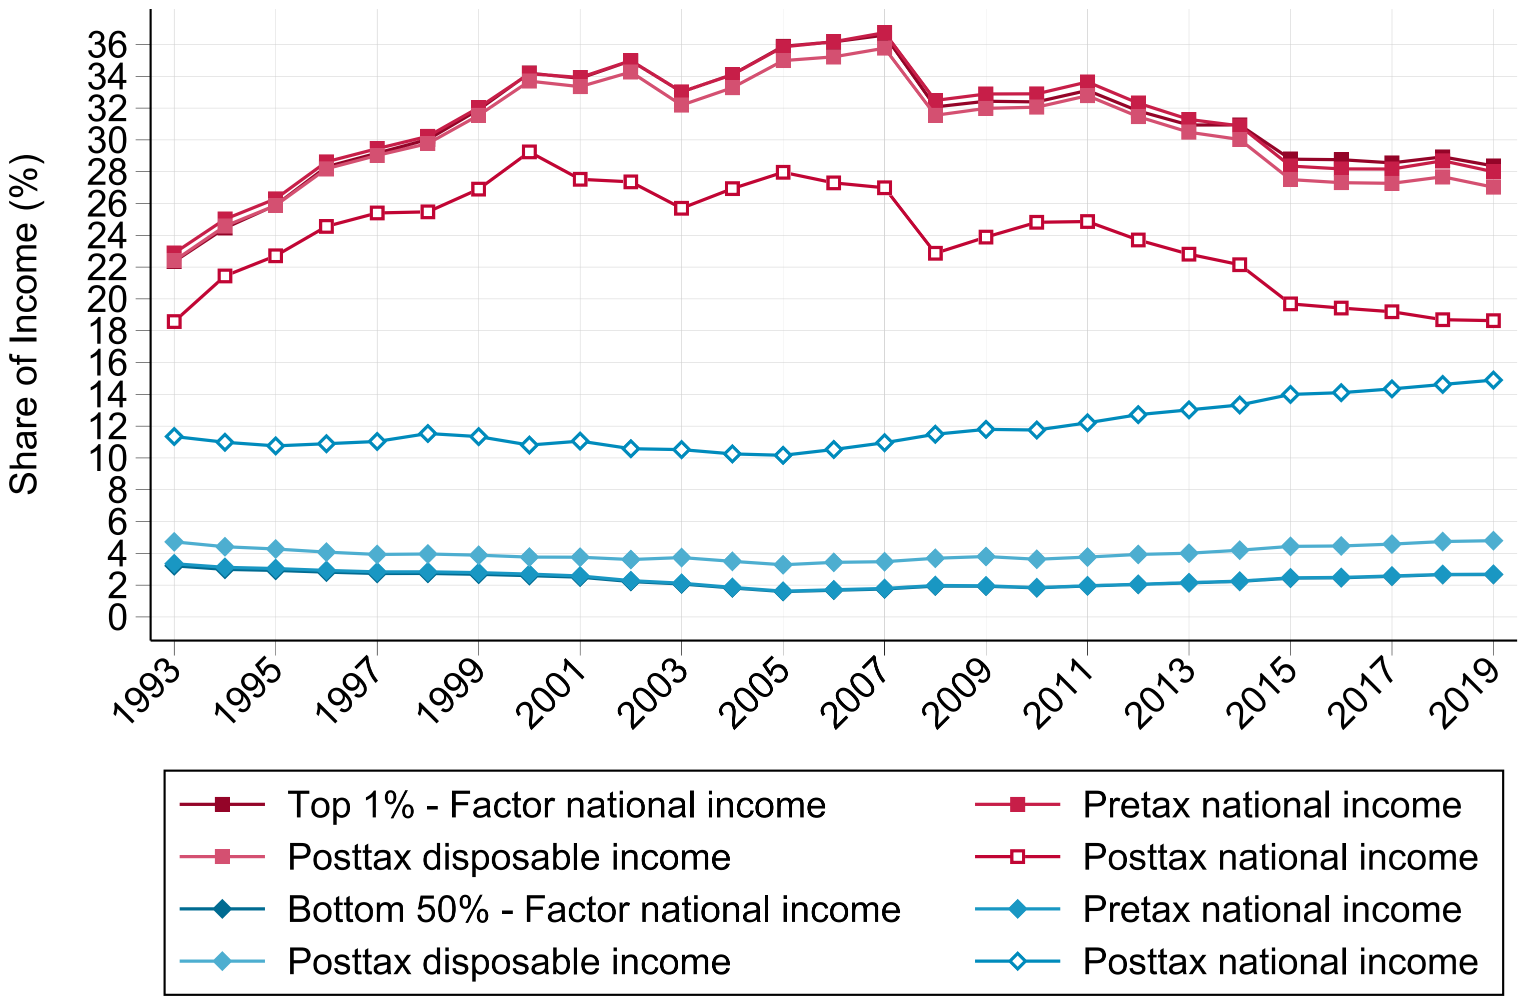

Notes. Authors' computations combining survey, tax, and national accounts data.

Figure 6.10: The overall impact of taxes and transfers on inequality

(b) Net transfers operated by the tax-and-transfer system by factor income group

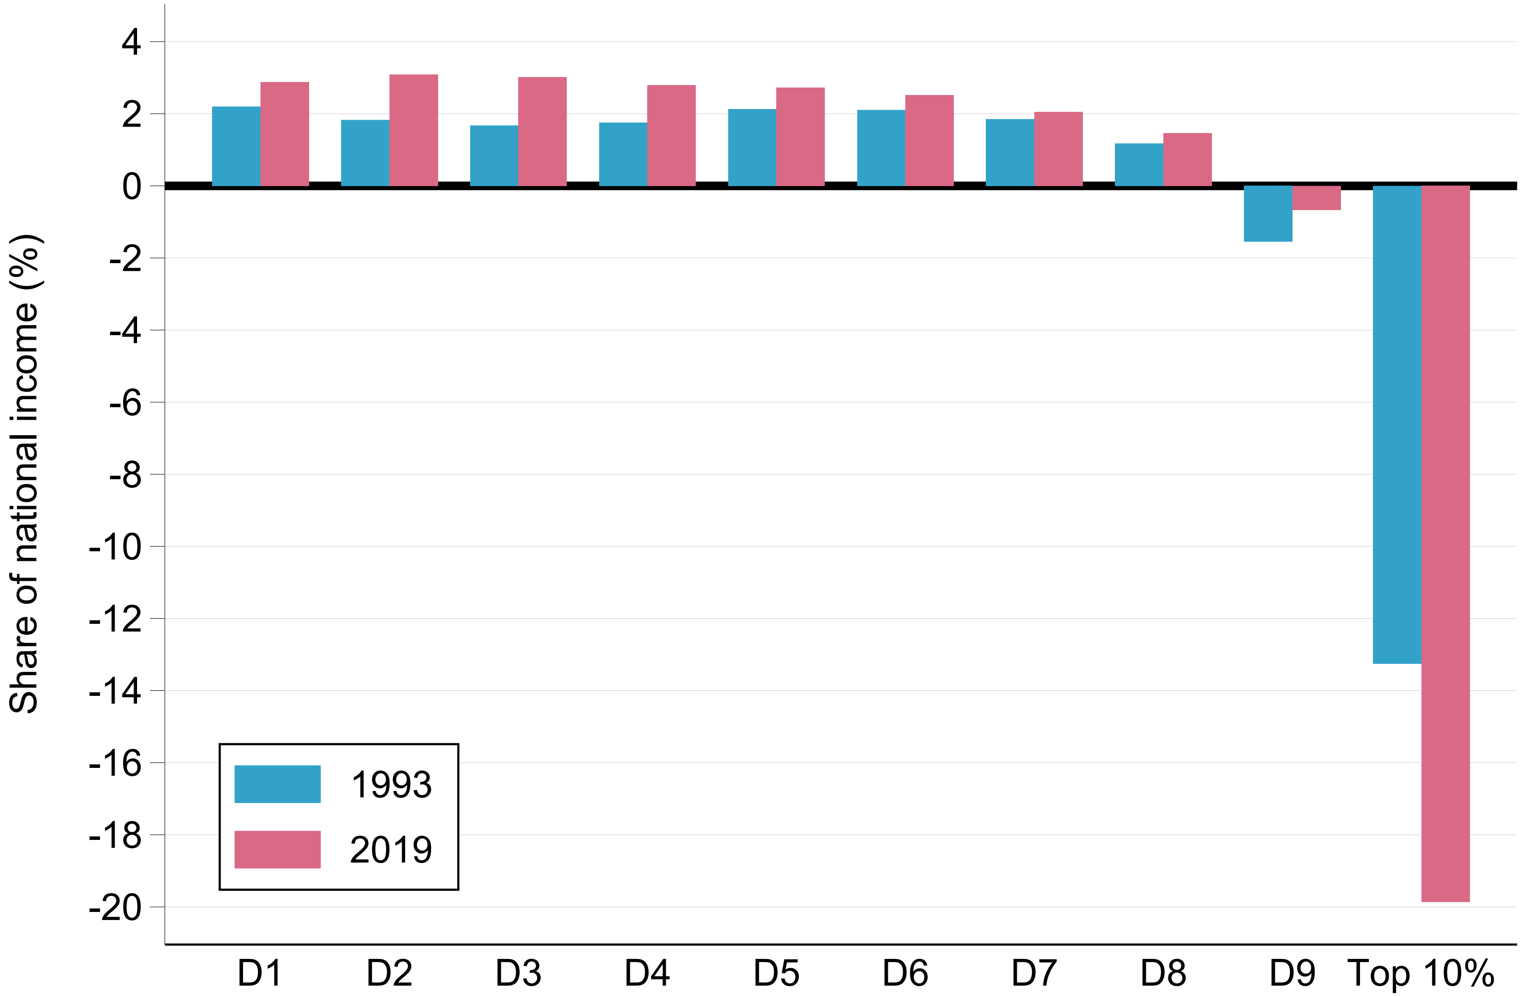

Notes. Authors' computations combining survey, tax, and national accounts data. The figure represents the net transfers operated between factor income deciles by the tax-and-transfer system, that is, the difference between total transfers received and total taxes paid, expressed as a share of national income.

Figure 6.11: Redistribution, inequality, and growth: cumulated income growth by percentile, 1993-2019

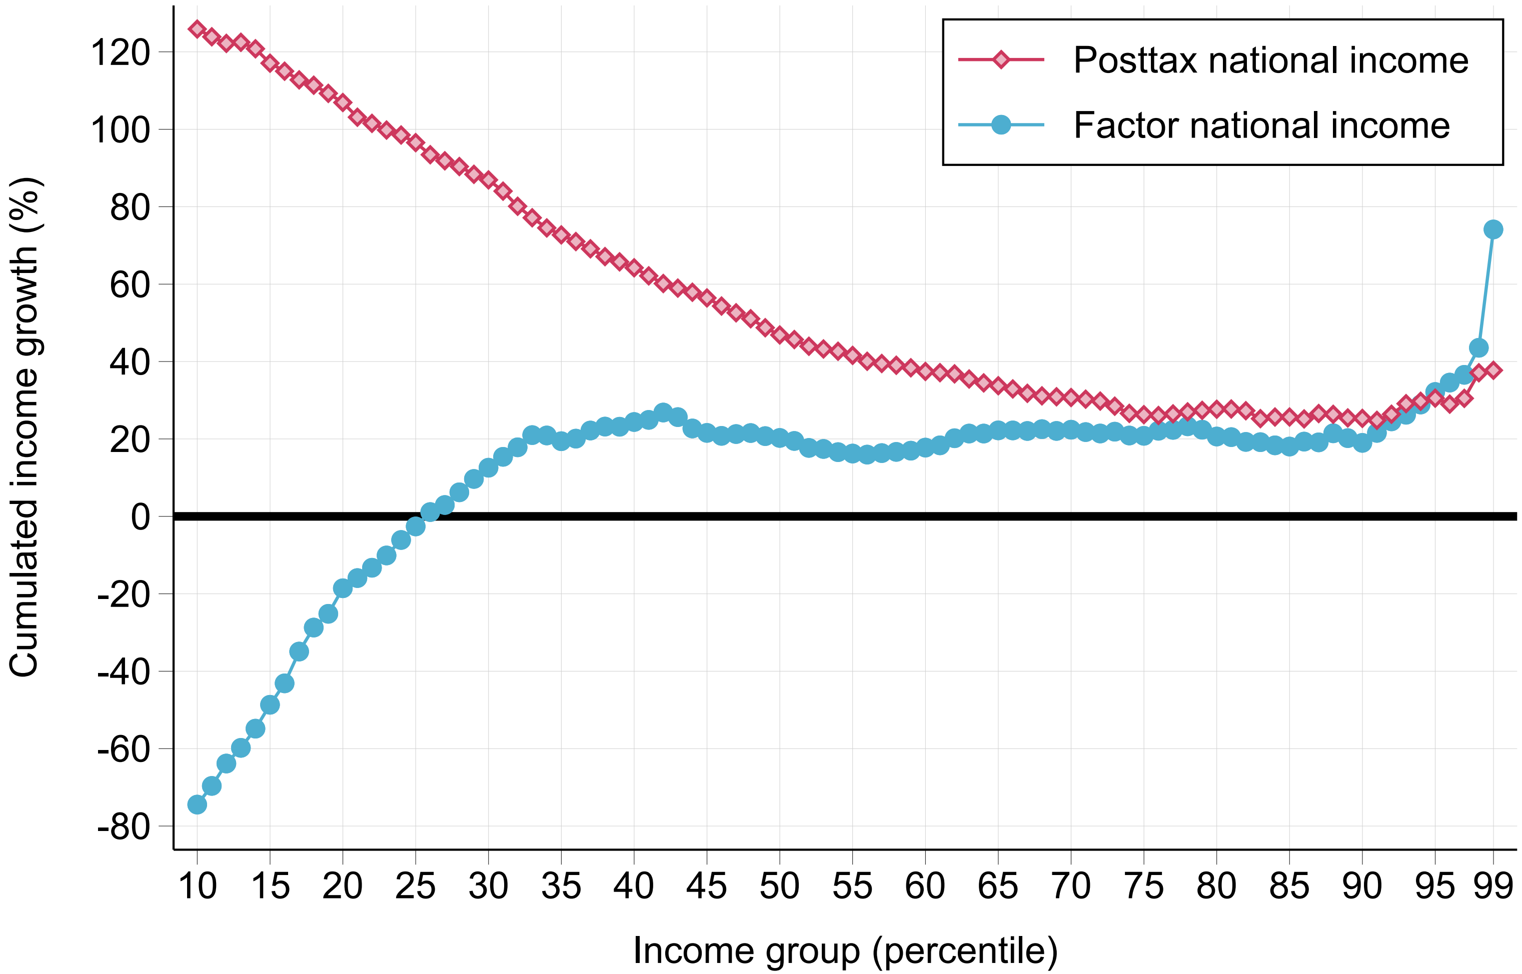

Notes. Authors' computations combining survey, tax, and national accounts data. The figure represents the cumulated income growth rate by percentile between 1993 and 2019 in terms of factor national income and posttax national income.

Figure 6.12: Racial inequality and top incomes

(a) Share of Black versus White earners in top factor income groups, 1955-2019

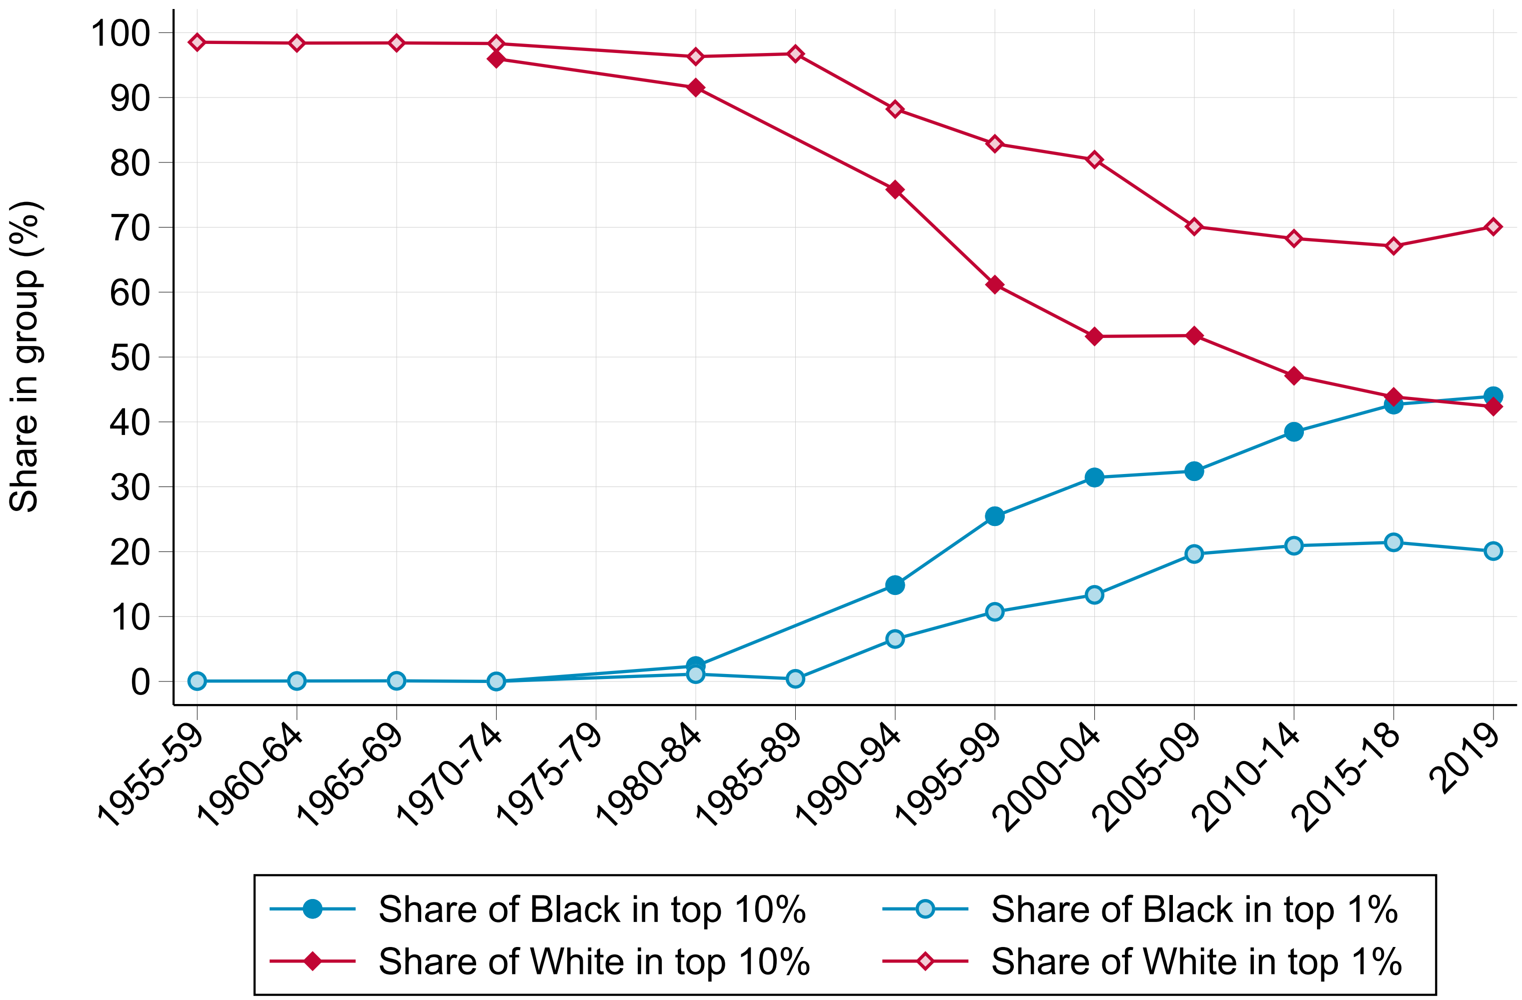

Notes. Authors' computations using from Alvaredo and Atkinson (2010) for the top 1% before 1993; census data for 1970, 1980, and 1990 (top 10%); distributional national accounts for 1993-2019. The figure represents the share of Black earners and White earners in top factor income groups between 1955 and 2019.

Figure 6.12: Racial inequality and top incomes

(b) Top Black incomes and the decline in the racial factor income gap

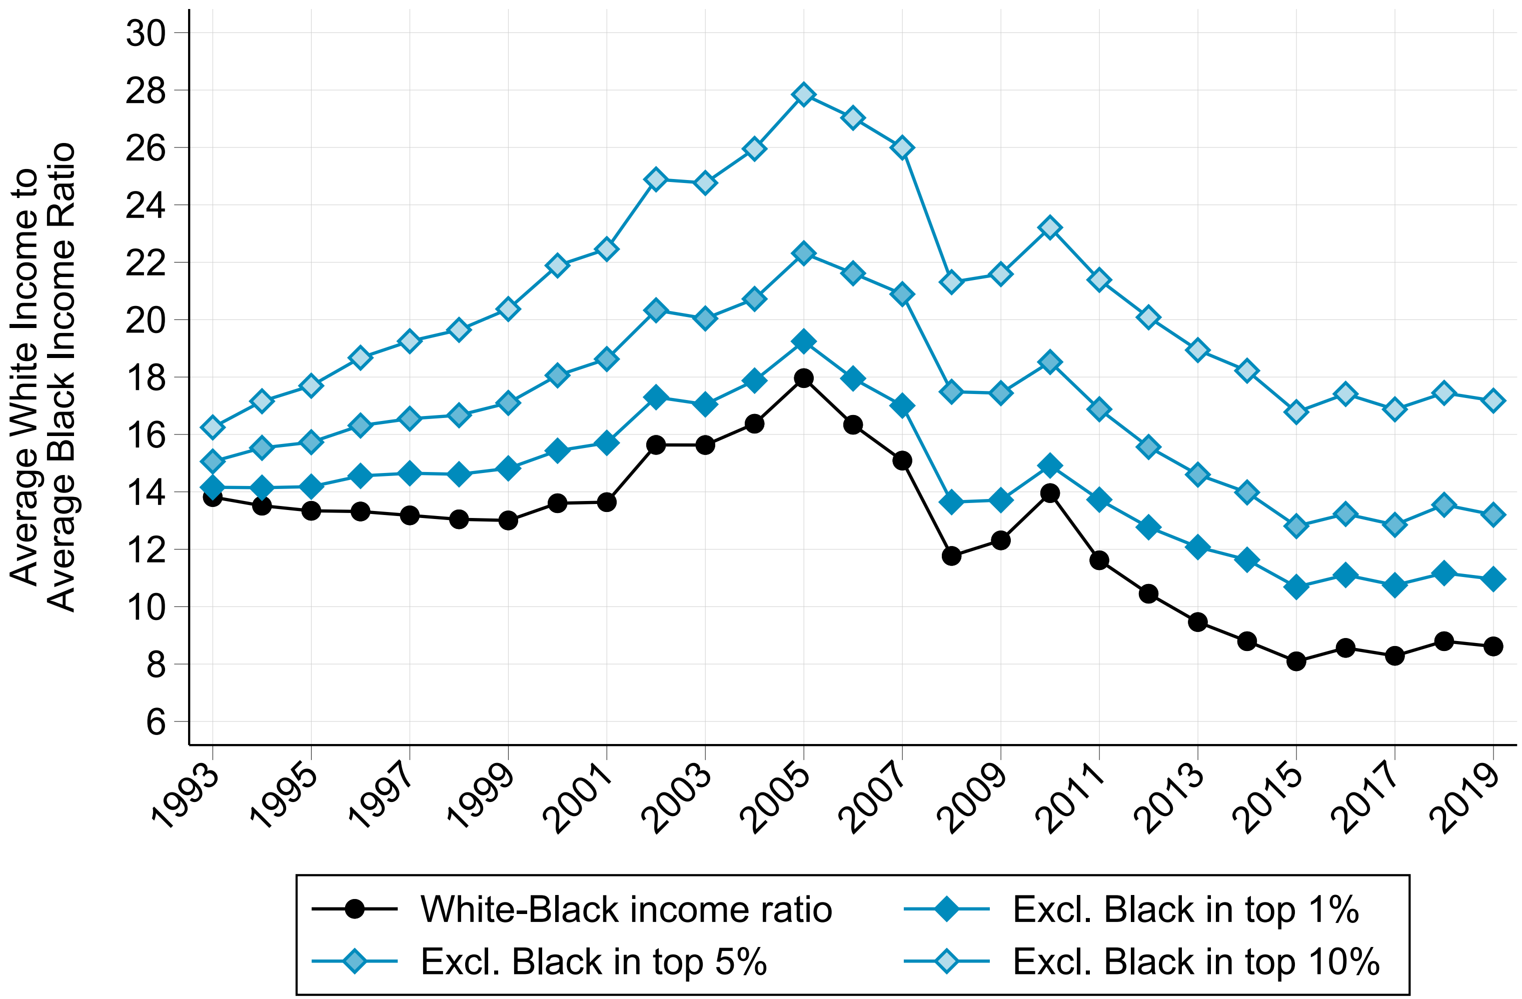

Notes. Authors' computations combining survey, tax, and national accounts data. The figure represents the ratio of the average factor income of White earners and the average factor income of Black earners between 1993 and 2019, before after excluding all Black earners located in the top 10%, the top 5%, and the top 1% of the overall factor income distribution.

Figure 6.12: Racial inequality and top incomes

(c) The distribution of growth within population groups, 1993-2019

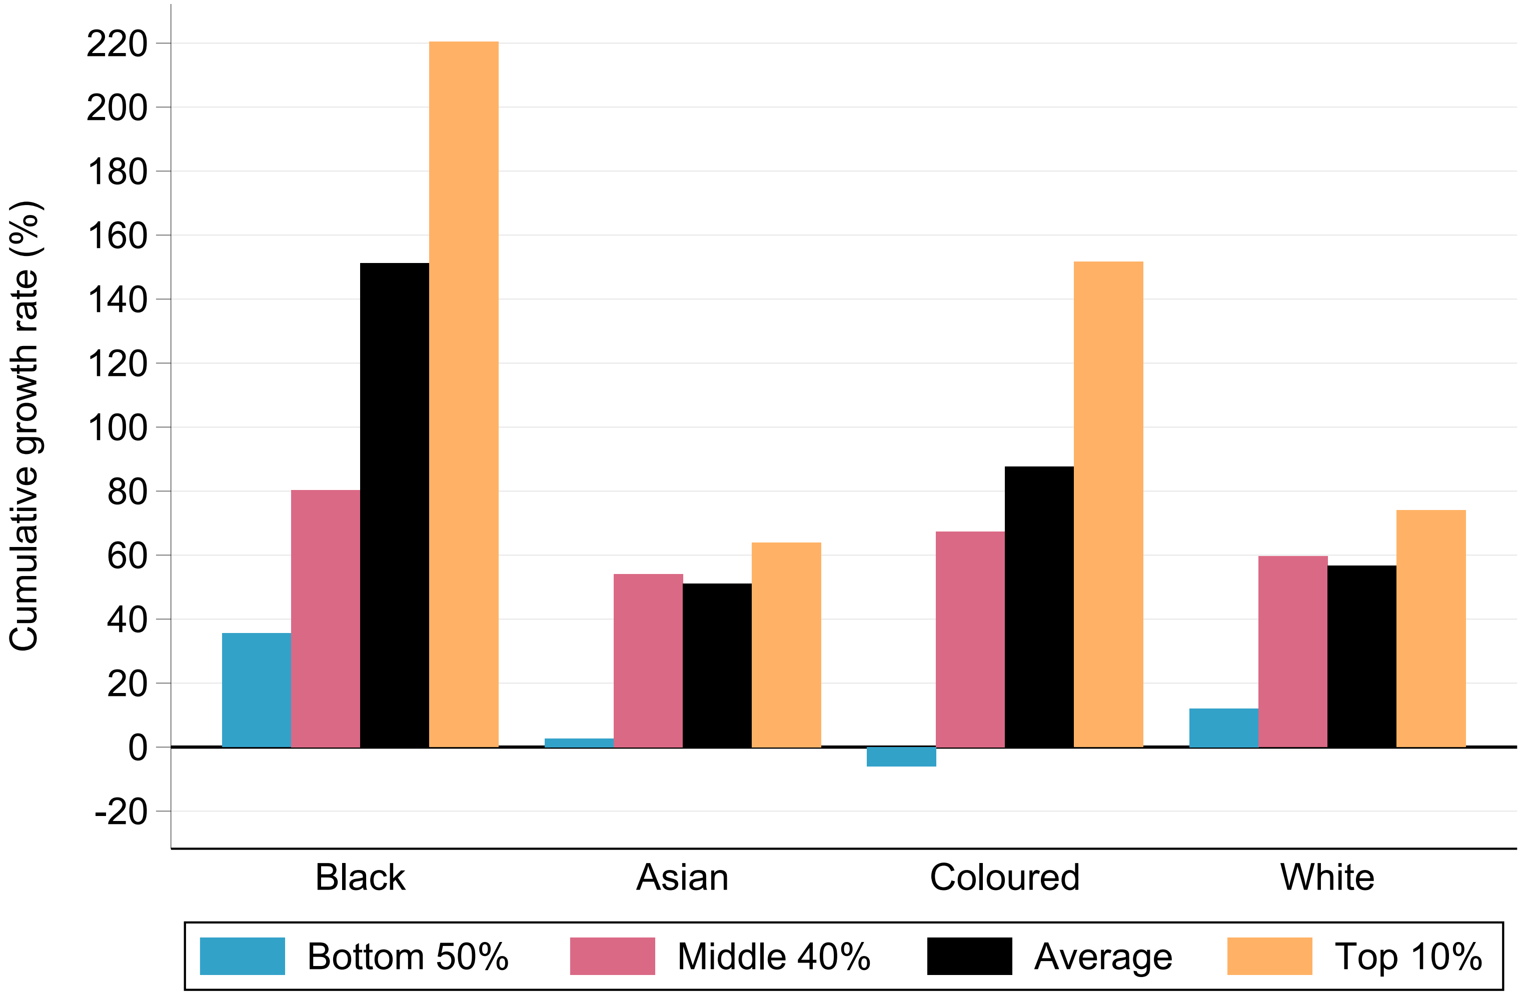

Notes. Authors' computations combining survey, tax, and national accounts data. The figure represents the total growth rate of selected factor income groups within each population group from 1993 to 2019. The top 10% of Black earners saw their average income grow by 200% during this period.

Figure 6.13: The structure of racial inequality in 2019

(a) The White-Black gap in income, consumption, and wealth

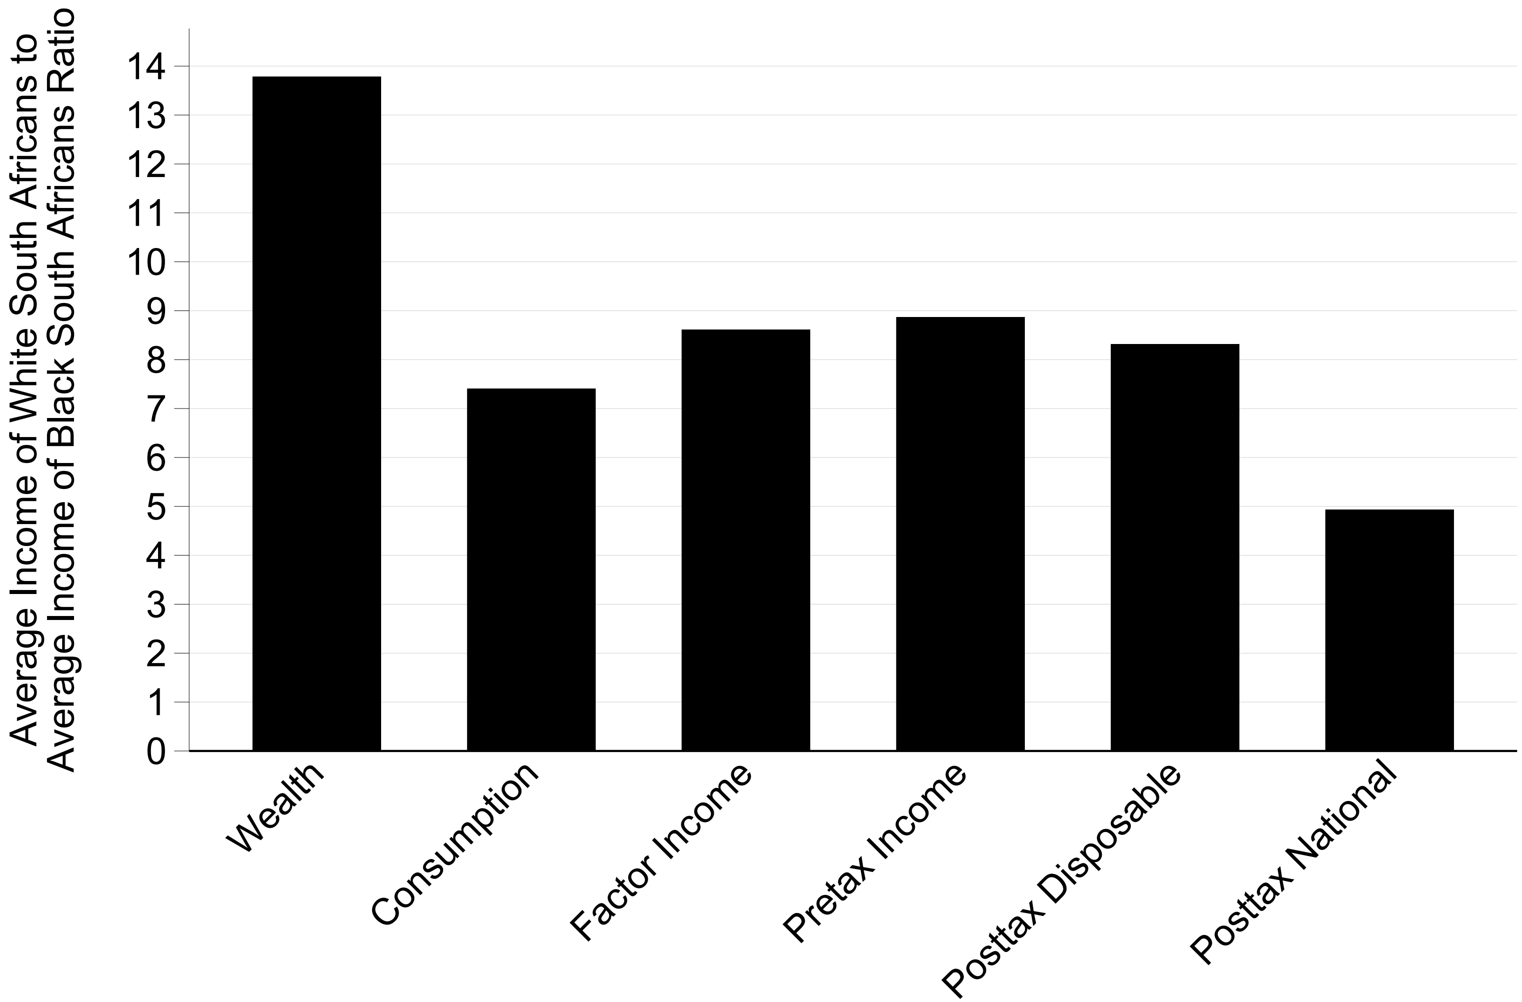

Notes. Authors' computations combining survey, tax, and national accounts data. The figure represents the ratio of White to Black average income, consumption, and wealth in 2019.

Figure 6.13: The structure of racial inequality in 2019

(b) Racial composition of posttax national income groups

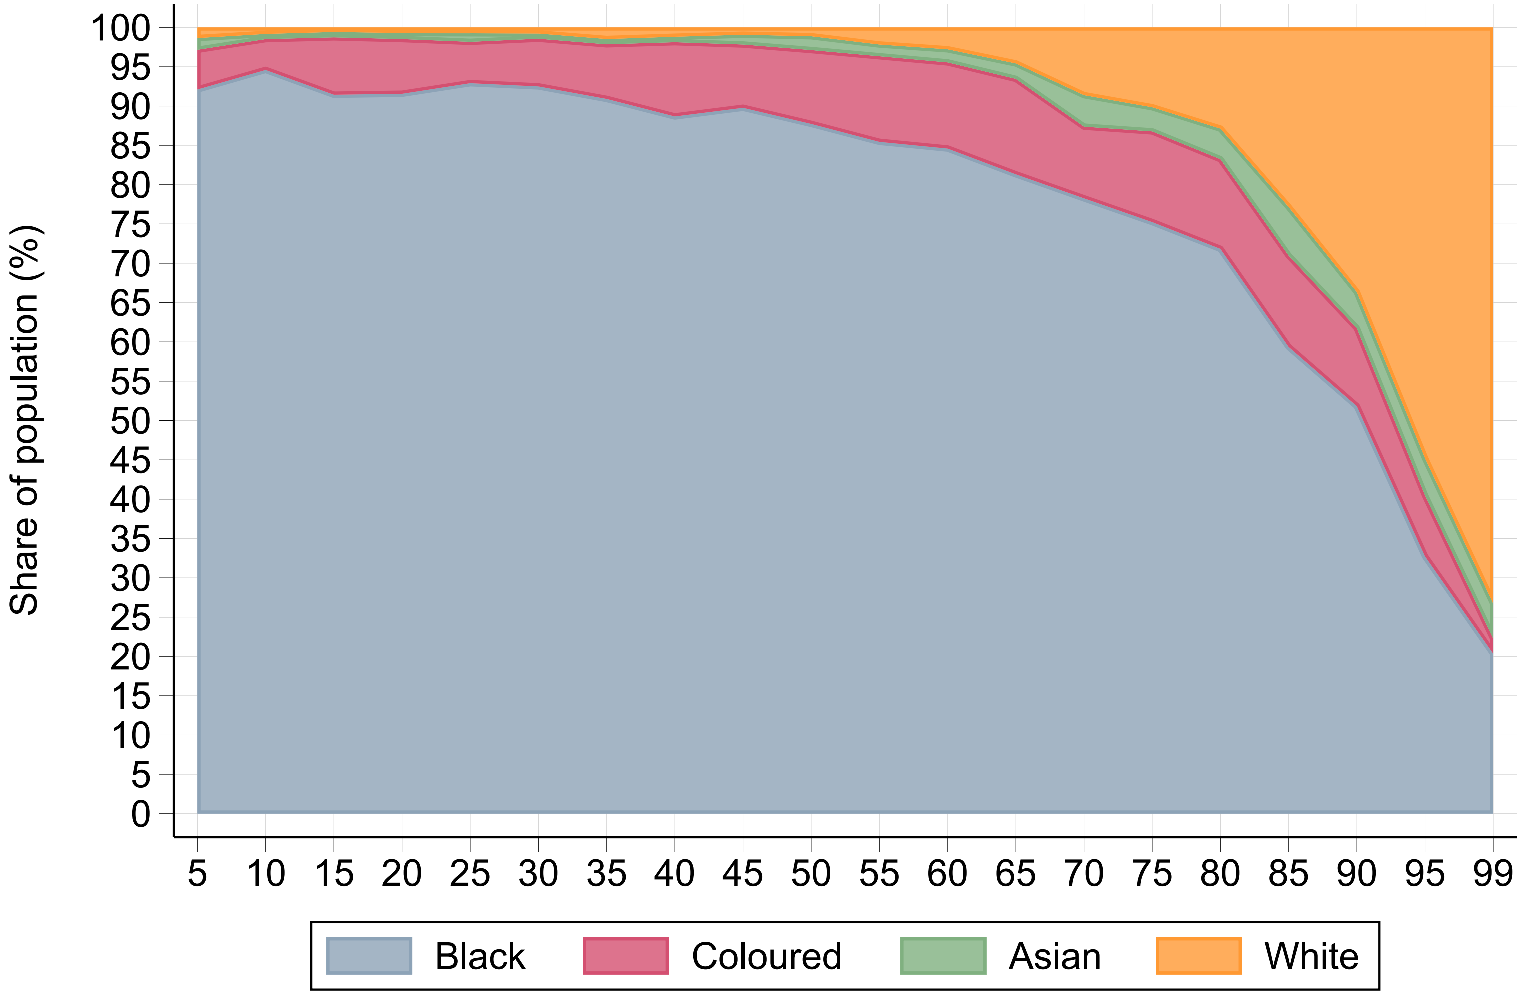

Notes. Authors' computations combining survey, tax, and national accounts data. The figure represents the composition of posttax national income groups (ventiles) by population group in 2019.

Figure 6.14: Spatial inequality and redistribution

(a) Average income by province relative to national income, 2019

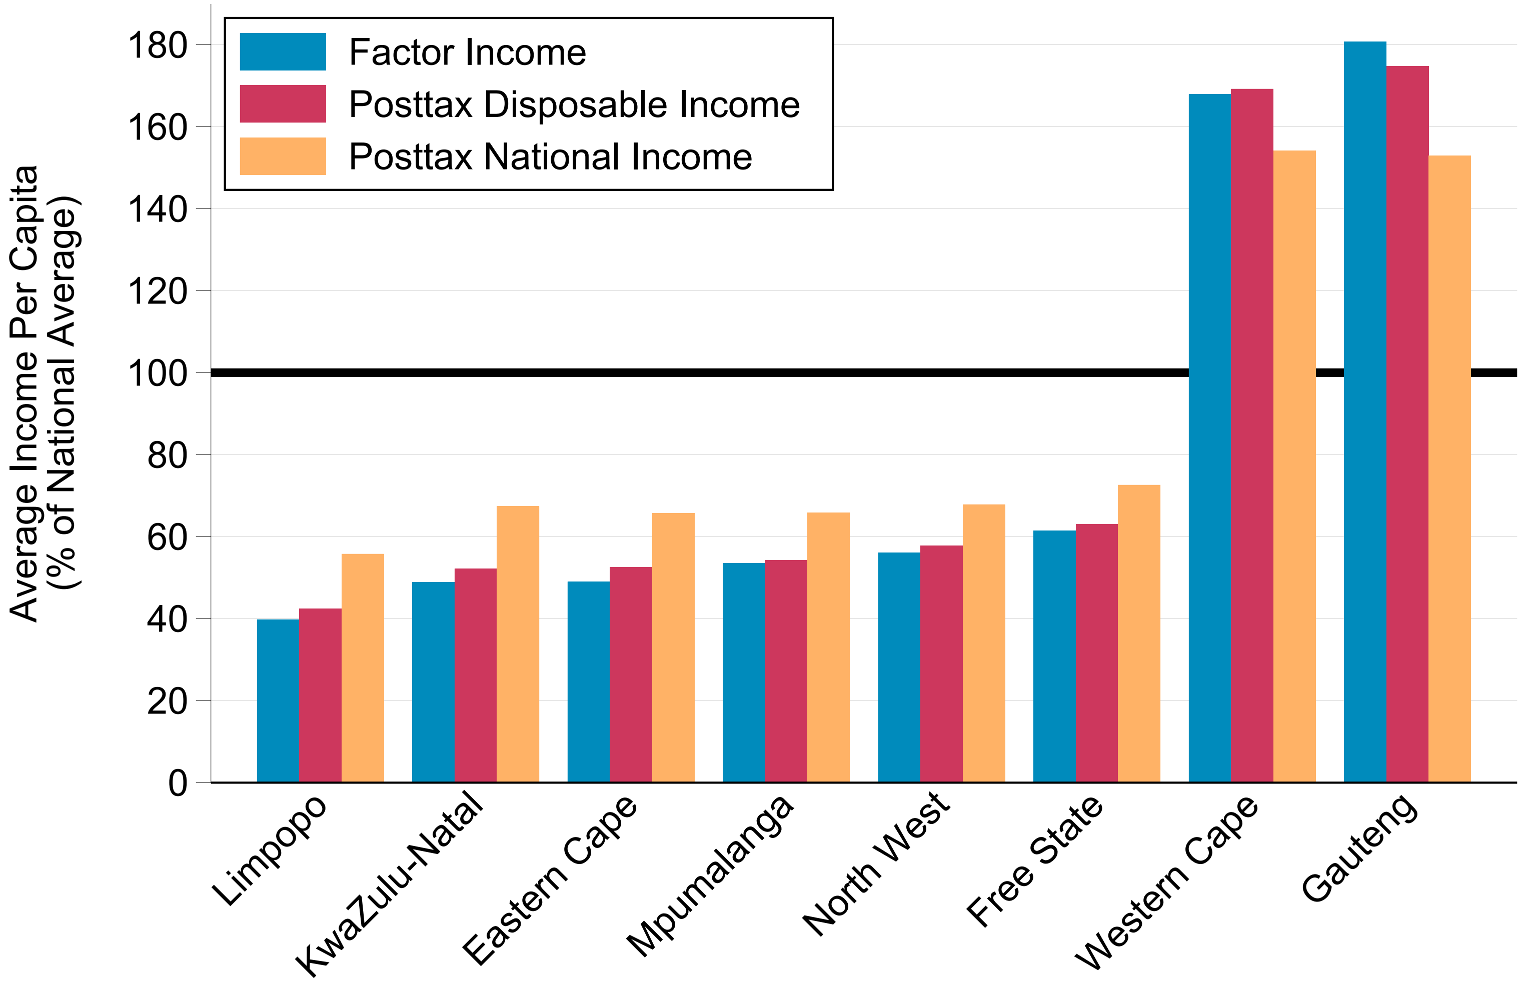

Notes. Authors' computations combining survey, tax, and national accounts data. Limpopo includes the North West province. The figure represents the average income of South African provinces, before and after taxes and transfers, relative to the national average in 2019.

Figure 6.14: Spatial inequality and redistribution

(b) Social transfers and the rural-urban income gap

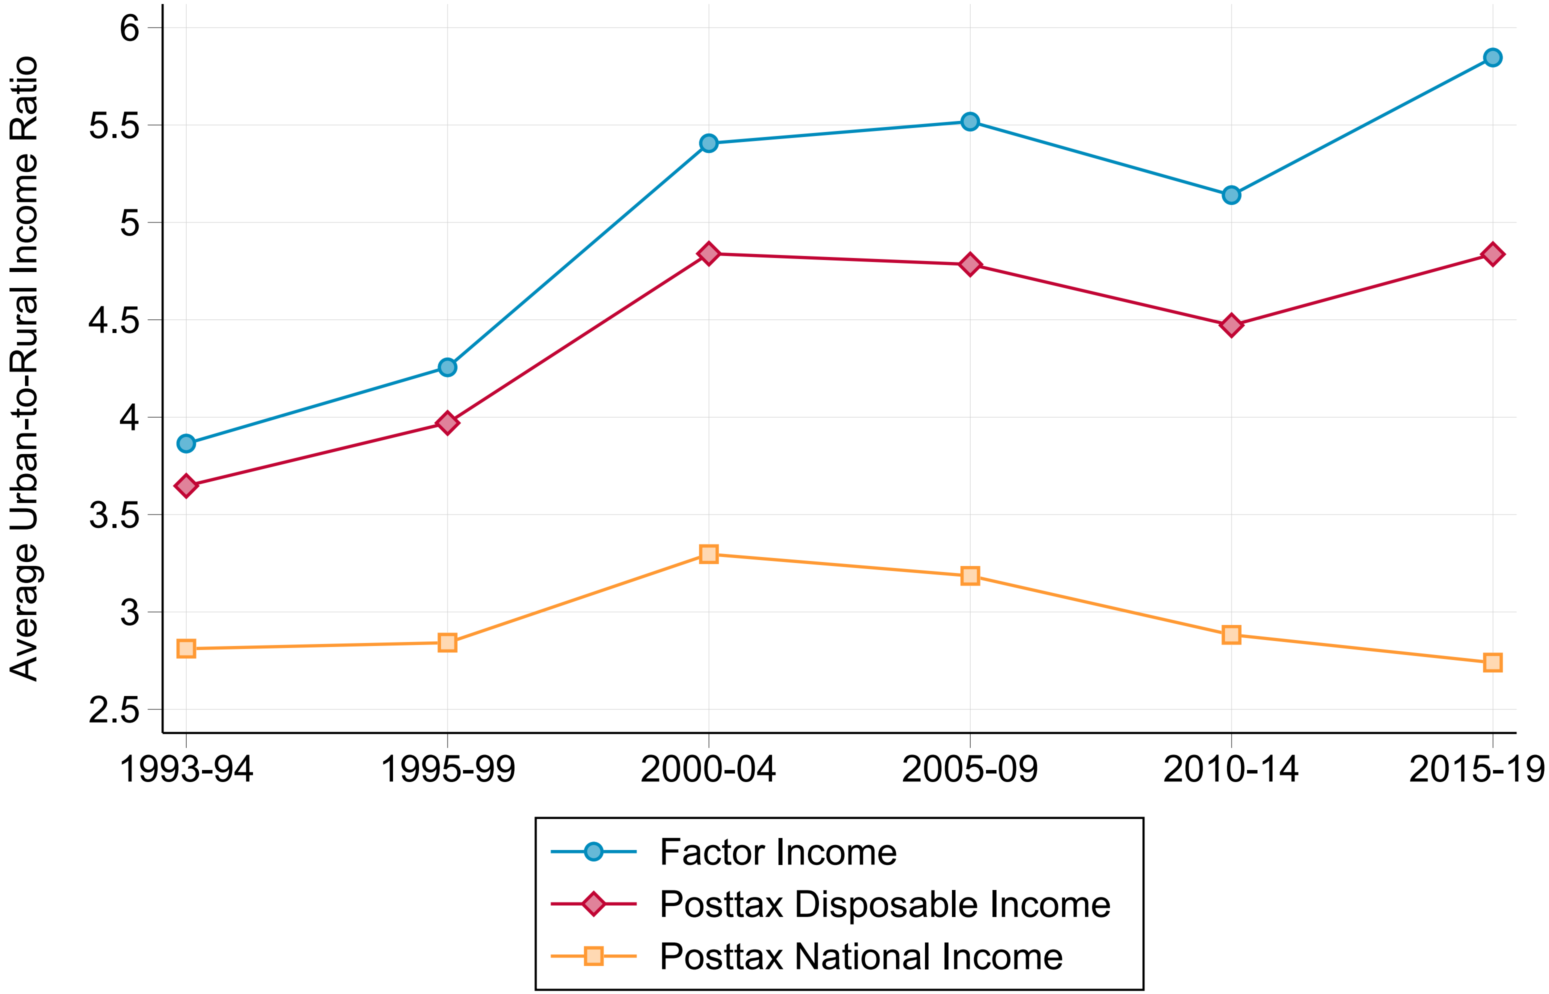

Notes. Authors' computations combining survey, tax, and national accounts data. The figure represents the ratio of the average income of urban areas to the average income of rural areas, before and after accounting for taxes and transfers.

Table 6.1: The distribution of factor national income and pretax national income

Item	Distribution method	% of NNI (2019)
Factor national income		100%
Compensation of employees	Proportional rescaling	57%
Mixed income	Proportional rescaling	9%
Property income, net		9%
Rents	Proportional rescaling	2%
Interest	Proportional rescaling	2%
Dividends	Proportional rescaling	4%
Other property income	Proportionally to pension and life insurance wealth	6%
Interest paid by households	Proportionally to factor income of debtors	-5%
Imputed rents of owner-occupiers	Proportionally to housing wealth of owner-occupiers	3%
Corporate undistributed profits	Proportionally to equity	8%
Taxes less subsidies on production and imports	Proportionally to factor income	11%
Remaining national income components	Proportionally to factor income	3%
Pretax national income		100%
Pension contributions	Observed	6%
Pension benefits	Observed	3%
Pension deficit or surplus	50% prop. to contributions, 50% prop. to benefits	3%
Unemployment insurance contributions	Rule-based imputation	0.5%
Unemployment insurance benefits	Observed	0.4%
Unemployment insurance fund deficit or surplus	50% prop. to contributions, 50% prop. to benefits	0.1%

Notes. The table reports the methodology used to distribute the various components of factor national income and pretax national income (for more details, see sections II.B and II.C), along with the size of each component expressed as a share of net national income (NNI) in 2019. Factor national income is the sum of all income flows accruing directly or indirectly to individuals, before accounting for the operation of the tax-and-transfer system, and before accounting for the operation of the pension and unemployment systems. Pretax national income is equal to factor income after the operation of the pension and unemployment systems. Both factor national income and pretax national income sum to the net national income.

Table 6.2: The distribution of taxes

Item	Distribution method	% of NNI (2019)
Direct taxes		19.0%
Personal income tax	Rule-based imputation	11.2%
Corporate income tax	Proportionally to equity	6.1%
Dividends tax	Proportionally to dividends	0.8%
Skills development levy	Rule-based imputation	0.4%
Transfer duties	Proportionally to housing wealth	0.2%
Securities transfer tax	Proportionally to equity	0.1%
Estate duty	Proportionally to net wealth	0.1%
Donations tax	Proportionally to net wealth	0.0%
Other taxes on income	Proportionally to pretax income	0.1%
Indirect taxes		12.6%
Value added tax	Proportionally to expenditure (excl. zero-rated / informal market)	8.0%
General Fuel Levy	Proportionally to fuel and transport expenditure	1.8%
Other excise duties	Proportionally to tobacco and alcohol expenditure	1.1%
Other taxes on goods and services	Proportionally to total expenditure	0.3%
Taxes on international trade	Proportionally to import-density-corrected expenditure	1.4%
Other government revenue	Proportionally to pretax income	2.0%
Total consolidated revenue		33.6%

Notes. The table reports the methodology used to distribute all taxes in South Africa at the individual level (for more details, see section II.D), along with the size of each component, expressed as a share of net national income (NNI), in 2019.

Table 6.3: The distribution of factor income in South Africa in 2019

	Number of individuals	Income threshold	Average income	Income share
Full population	58,600,000	$ 0	$ 11,700	100%
Bottom 90% (p0p90)	52,740,000	$ 0	$ 4,100	31.3%
Bottom 50% (p0p50)	29,300,000	$ 0	$ 600	2.7%
Middle 40% (p50p90)	23,440,000	$ 2,200	$ 8,400	28.7%
Top 10% (p90p100)	5,860,000	$ 26,200	$ 80,600	68.7%
Top 1% (p99p100)	586,000	$ 129,000	$ 332,000	28.3%
Top 0.1% (p99.9p100)	58,600	$ 662,000	$ 973,000	8.3%
Top 0.01% (p99.99p100)	5,860	$ 1,370,000	$ 2,400,000	2.0%

Notes. The table reports the distribution of factor national income in 2019, providing information for each income group on the number of adults belonging to this group, the minimum income required to belong to this group, the average income of this group expressed in 2019 PPP US dollars ($1 = R6.3), and the share of factor national income received. Factor national income is the sum of all income flows accruing directly or indirectly to individuals, before accounting for the operation of the tax-and-transfer system, and before acounting for the operation of the pension and unemployment systems. Income is split equally among all adults members of the household (aged 20 or above).

Chapter 7

Wealth Inequality in South Africa, 1993-2017

A growing number of studies have made significant progress in measuring the distribution of household income and consumption within countries and over time, yet still little is known on the dynamics of household wealth. This knowledge gap is particularly acute in the developing world, where available data sources are scarce, often insufficiently detailed and prone to important measurement error. Given the rise of global wealth concentration (Alvaredo et al., 2018; Zucman, 2019) and the policy challenges it poses in terms of tax evasion (Alstadsæter, Johannesen, and Zucman, 2019; Kleven et al., 2020; Londoño-Vélez and Ávila-Mahecha, 2021) and political equilibrium (Bertrand et al., 2020; Bombardini and Trebbi, 2020; Esteban and Ray, 2006) there is a pressing need to address this shortcoming and improve our knowledge of the wealth distribution.

This paper estimates the distribution of household wealth in South Africa from 1993 to 2017 by combining household survey data, tax microdata, and macroeconomic balance sheets statistics. A number of results emerge from our analysis.

First, South Africa displays unparalleled levels of wealth concentration. The top 10% of South African wealth holders own more than 85% of household wealth, while the top 1% wealth share reaches 55%. The top 0.01% (about 3,500 adults) own a higher share of wealth than the bottom 90% as a whole (about 32 million individuals). The average wealth of the bottom 50% is negative: the market value of their assets is lower than their liabilities. Such levels of wealth inequality are higher than in any other country for which comparable, high-quality estimates of the wealth distribution are available (namely France, the United Kingdom, the United States, Russia, China, and India).

Secondly, there is no evidence that wealth inequality has decreased since the end of the apartheid regime. The top 10% wealth share has fluctuated between 80% and 90% between 1993 and 2017, largely as the result of the rise and fall of household debt before and after the 2007-2008 crisis, with no sign of long-run trend. If anything, the available evidence suggests that the share of wealth captured by the top 1% and the top 0.01% may even have increased. This result is particularly striking considering South Africa's recent history of positive growth (real average income and wealth per adult respectively increased by 19% and 33% from 1993 to 2017) and greater racial inclusiveness (all discriminatory laws against oppressed racial groups had been abolished by 1991).

Thirdly, these inequalities are reproduced at the level of all asset classes. The top 10% of wealth holders own more than 55% of business assets and housing wealth, and over 99% of bonds and stock. Financial assets constitute the bulk of the assets of the top 0.1%, while owner-occupied housing and pension wealth are the main holdings of the bottom 90%. Significant wealth accumulation is visible over the life cycle, but levels of wealth concentration within each age group are almost perfectly similar to those measured for the full population. This suggests that individuals across the wealth distribution do accumulate at relatively similar paces but start from very different initial endowments, hence pointing to the importance of inheritance.

Previous studies on post-apartheid economic inequality have focused on income, but the literature on wealth remains extremely scarce. Two studies have attempted to measure the distribution of wealth in South Africa (Daniels and Augustine, 2016; Mbewe and Woolard, 2016), yet they suffer from two major limitations.[1] First, they cover only one (2015) or two years (2010, 2015) of data and therefore cannot assess any long-run trends in wealth inequality since the end of apartheid. Secondly, they rely exclusively on the National Income Dynamics Study, a wealth survey that greatly underestimates wealth concentration within the top 10% (this issue and its implications are discussed in more detail in section 5). This is in large part due to substantial under-reporting of financial assets by survey respondents, a limitation that has now been extensively documented in the inequality measurement literature (Blanchet, Flores, and Morgan, 2018; Blanchet, Fournier, and Piketty, 2017; Blanchet et al., 2021; Korinek, Mistiaen, and Ravallion, 2006), as well as by the authors of the previous studies themselves (Daniels and Augustine, 2016).

[1] See Chatterjee (2019) for a broader review. Orthofer (2016) is sometimes cited as an additional study, exploiting tax microdata. However, given the method applied, the resulting estimates correspond to the distribution of financial incomes, not to the distribution of household wealth.

By contrast, following income capitalisation approaches recently applied in the United States (Saez and Zucman, 2016) and France (Garbinti, Goupille-Lebret, and Piketty, 2018), our methodology combines survey and tax microdata with macrodata on household wealth totals. Unlike previous studies, it ensures that average wealth and the portfolio composition of assets across the distribution are fully consistent with the household balance sheets statistics published by the South African Reserve Bank. It allows us to obtain a much more reliable picture of wealth inequality within the top 10% and especially within the top 1%, which is key to understanding wealth dynamics in countries such as South Africa where wealth concentration is extreme. Importantly, it allows us to cover the entire 1993-2017 period, as well as to compare wealth inequality in South Africa to other countries where similar exercises have been performed.

Finally, this paper also contributes to the methodological literature on the measurement of wealth inequality in developing countries. By comparing estimates of the wealth distribution obtained with three different methodologies—direct measurement of net worth, rescaling of reported wealth components to balance sheets totals, and capitalisation of income flows—we show that capitalising reported income flows to match macroeconomic wealth totals can yield relatively good results, even in the absence of income tax microdata. Crucially, these estimates appear to be much more reliable than those solely relying on survey-based self-reported wealth, which omit the bulk of financial wealth. In other words, bridging the micro-macro gap in wealth measurement appears to be an essential step to accurately measure the wealth distribution. This opens new avenues for estimating the dynamics of wealth inequality in low- and middle-income countries, where wealth microdata are unavailable or unreliable, yet where macroeconomic balance sheet statistics can be usefully combined with surveys collecting data on household income. In that respect, we hope that this paper can serve as a useful guide for future studies aiming to measure wealth inequality in countries with limited data such as South Africa.

The rest of the paper is organized as follows. Section 2 defines the key concepts and presents the different data sources we use. Section 3 explains the methods we apply to combine these data sources. Section 4 presents our main results and compares our estimates with that of other countries. Section 5 contrasts our results with those obtained from alternative methodologies.

7.1 Concepts and Data Sources

Following the United Nations System of National Accounts (UN SNA) guidelines (United Nations, 2009), we define household wealth as the total market value of the assets and liabilities held by the household sector. Using this concept is central to produce comparable estimates over time and across countries. Assets can be classified into eight broad categories: owner-occupied housing, tenant-occupied housing, unincorporated business assets, pensions, life insurance, bonds, equity, and currency (deposits, notes and coins). Liabilities can be divided into mortgage debt and all other debts (including consumer credits, credit cards, and informal loans).[2] As with most countries in the world, there exists no unified administrative database in South Africa measuring wealth at the micro level for the full population.[3] In the absence of such information, the distribution of household wealth in South Africa has to be measured by combining several complementary data sources.

Macroeconomic data. In South Africa, the first comprehensive attempt to estimate the value of total household wealth in the economy goes back to Muellbauer and Aron (1999), who collect and combine a number of data sources to provide figures on the assets and liabilities of the household sector since 1975. The South African Reserve Bank (SARB) has since then updated and revised these figures on a yearly basis. The only alternative data source that would allow to approximate total household wealth are waves 4 (2015) and 5 (2017) of the National Income Dynamics Study.[4] As it covers only two years, this survey offers little scope to study the evolution of wealth inequality in the long run. Moreover, it suffers from several limitations (internal inconsistencies, measurement errors, implausibly low aggregates) documented in section 5 (see also appendix G.2). For these reasons, we prefer not to rely on this source. Throughout our series, all wealth totals thus come from macroeconomic balance sheets published by the SARB. They are then combined with diverse microdata sources to estimate how these aggregates are distributed.

Personal Income Tax data. We exploit Personal Income Tax (PIT) data compiled by the South African Revenue Service (SARS) to measure the distribution of wages, pension income, pension contributions, mixed income, and capital income (rents, interest, and dividends) for the top 30% of the population. This individual panel

[2]This classification is the most precise common decomposition that could be achieved after harmonisation of all the data sources. Notice that land directly owned by the household sector is classified in housing (owner- or tenant-occupied), not in business assets. Liabilities include all debts contracted with both formal (e.g. commercial banks) and informal creditors.

[3]The few countries still collecting direct information on wealth include Switzerland, Spain, France, Norway, and Colombia. These countries are the only ones still enforcing a tax on net wealth. For other countries in the world, most of what we know about wealth either comes from wealth surveys, estate duty data, or, as in this study, via the income capitalisation method applied on income surveys or personal income tax data.

[4]Other surveys collecting information on income and consumption sometimes include some information on some wealth components (mostly house value or debt), but never encompass total wealth.

covers two types of tax statements over the 2010-2017 period: IRP5 forms, which are submitted to SARS by employers on behalf of their employees and cover wages and pension contributions, and ITR12 forms, which are self-assessed by all taxpayers who need to disclose information on mixed, rental, interest, and dividend incomes.[5] Due to its administrative nature, this data covers the full tax paying population, including individual observations at the very top of the distribution, which greatly increases the granularity of measured income flows. This is an advantage over surveys, which often suffer from sample biases and higher non-response rates among the wealthiest.

Household surveys. Finally, we combine a number of household surveys to cover individuals and income or wealth concepts not captured by the tax data. We use surveys for three main purposes: to measure the distribution of key income variables for the bottom 70% of the population; to estimate the distribution of debts and assets that do not generate income flows and hence cannot be capitalised (owner-occupied housing, currency); and to extrapolate our 2010-2017 series back to 1993. These include two main types of surveys: seven "income surveys"[6] covering all forms of incomes received by individuals (as well as certain wealth components such as housing and debts), and fifty-four "labour force surveys"[7] conducted on a more regular basis since 2000 and mainly covering wages and mixed income.

7.2 Methodology

This section presents the methodology used to estimate the distribution of household wealth in South Africa since 1993. First, a harmonized survey microfile is built by merging existing household surveys. Surveys are then combined with tax data to better capture the top end of the distribution. Finally, measures of net worth are derived by capitalising relevant income flows and rescaling other assets and liabilities to macro totals.

Harmonization of household surveys. We begin by combining household surveys to estimate the distribution of available income and wealth components, on a yearly basis, throughout the 1993-2017 period. Starting from available income surveys (1993, 1995, 2000, 2005, 2008, 2010, 2015), we first interpolate missing years from 1993

[5] The IRP5 and ITR12 data are presented in the form of source codes corresponding to specific taxable income concepts, exemptions and deductions. See the appendix for more details about our classification and Ebrahim and Axelson (2019) for an overview and discussion of the dataset.

[6] The Project for Statistics on Living Standards and Development (PSLSD - 1993), the Income and Expenditure Surveys (IES - 1995, 2000, 2005, 2010) and the Living Conditions Surveys (LCS - 2008, 2015).

[7] The Labour Force Surveys (LFS - twice a year from 2000 to 2007) and the forty Quarterly Labour Force Surveys (QLFS - every three months since 2008).

to 2017 by creating new datasets resulting from the combination and proportional reweighting of the two adjacent surveys. Yearly distributions of gross wages and mixed incomes are then corrected to make them match those reported in the Labour Force Survey series since 2000. In broad strokes, this process allows us to obtain a harmonized survey microfile covering every year from 1993 to 2017, in which the distribution of available income and wealth components are fully consistent with information reported in both income surveys (for all income concepts excluding wages and mixed income) and labour force surveys (for wages and mixed income). More details on these methodological steps are available in the appendix.

Combination of household surveys with tax data Survey distributions are combined with PIT data to better capture the top end of the distribution in two steps. First, we derive an income concept that is comparable between the survey and tax data, which we refer to as "merging income"[8], and we merge the two data sources based on the exact rank of merging income observed at the individual level. We then identify the quantile of the South African income distribution q above which reported merging incomes become higher in the tax data than in the survey data, and we assume that the tax data is more reliable than the survey data only above q. In practice, this implies keeping all variables from the survey data below q, and replacing all comparable variables from the tax data above q (wages, mixed income, rental income, interest, dividends, private pension income, and contributions to pension funds). Between 2010 and 2017, we find q to be consistently located between the 70^{th} and the 75^{th} percentiles, so that we use the tax microdata to cover the top 25-30% of the income distribution.[9]

Income capitalisation and rescaling The income capitalisation method consists in using capital income flows (e.g. dividends) to approximate the distribution of

[8]Defined as the sum of wages, mixed income, rental income, interest income, and pension income.

[9]See appendix Figures G.8 and G.9. Our choice of a merging point based on an income concept differs slightly from the approach of Hundenborn, Woolard, and Jellema (2018), who rather derive a taxable income concept from survey data, and then keep the tax data above the filing threshold of taxable income. The main reason for merging our two datasets based on a broad income concept is twofold. First, our IRP5-ITR12 panel covers a large number of individuals who are below the filing threshold, given that all employers in South Africa are now required to file an IRP5 tax form for all their employees, regardless of their level of remuneration. However, as is emphasised in the SARS' Tax Statistics, this rule was not followed strictly by all employers, so that the tax data cannot be considered to be representative of the universe of formal wage earners. In other words, our data covers relatively well the top of the distribution up to a certain point, below which it contains a mix of low- and middle-income wage earners. It seems therefore most useful to keep as many individuals as possible from the tax data, while removing those whose location in the distribution of income cannot be identified precisely, which is what our method does in a simple way. Secondly, defining taxable income remains a complex task, and it remains unclear whether this can be done with a sufficient level of precision and consistency, in particular given that surveys tend to not properly capture the top of the distribution.

households' assets and liabilities (e.g. shares). In our case, given that the SARB balance sheet is the best available data source to capture the level and composition of total household wealth in South Africa, this implies distributing each aggregate in proportion to its income flow measured at the micro level. The core assumption is that of constant rates of return by asset class. Six types of assets can be capitalised: tenant-occupied housing from the rental income received by individual landowners; unincorporated business assets from the mixed income received by self-employed individuals; pension assets from the pension contributions and pension income of formal wage earners and pensioners; life insurance assets from factor income; bonds and interest deposits from interest income; and corporate shares and equity from dividends.[10]

The capitalisation method cannot be applied to liabilities nor to owner-occupied housing and currency, as these components of wealth do not generate any income flow. We therefore measure these components directly from available household surveys and rescale them proportionally to match SARB totals. To mitigate measurement issues and the risk of creating outliers with excessively negative net worth,[11] however, we do not directly rescale debts: we assume instead that mortgage debt is distributed proportionally to the value of the house of mortgagors, and that other forms of debts are distributed proportionally to the consumption of those declaring having contracted debts. These are conservative assumptions, as mortgages and other forms of debt are likely to be more unequally distributed than house values and consumption respectively. We refer to this combination of rescaling and income capitalisation as a "mixed approach" (see table 7.1).

Finally, to extrapolate our series backwards to 1993, we first apply our methodology to the years 2010-2017, with and without PIT data. We then compare the wealth distribution resulting from these alternative specifications to extract average correction coefficients at the quantile level, and use these coefficients to adjust the wealth

[10]In the case of pension assets, we follow the approach proposed by Saez and Zucman (2016) and allocate them to wage earners and pensioners so as to match their distribution recorded in the NIDS. In our case, this implies distributing 75% of pension assets to formal wage earners proportionally to pension contributions paid, and 25% to pensioners proportionally to pension income received. As shown in the appendix (figure G.6), this capitalisation technique applied to the NIDS data yields results which are very similar to those obtained from direct measurement. Similarly, we assume that 50% of life insurance assets belong to wage earners proportionally to factor income—the sum of wages, mixed income and pension income—and that 50% belong to all other adults proportionally to factor income. This again reproduces well the distribution of life insurance assets reported in the NIDS (see appendix figure G.7).

[11]Mortgage debt and other forms of debts have been recorded in surveys but the coverage is often partial and inconsistent. As a result, rescaling debts to balance sheets totals results in seriously overestimating the number of individuals with negative net worth and generating implausibly high debt values.

distributions estimated from survey data over the 1993-2010 period (see appendix G.2).

7.3 The distribution of wealth in South Africa: key results and comparative perspectives

This section presents our main results on wealth inequality in South Africa. We first provide an overview of aggregate household wealth and how it is distributed across broad wealth groups. We then present figures on the concentration of specific assets and on the dynamics of wealth accumulation over the life cycle. Finally, we discuss how wealth inequality in South Africa has evolved since 1993, and how it compares to other countries.

7.3.1 The level and composition of aggregate wealth in South Africa, 1993-2018

Before presenting figures on the distribution of wealth, it is useful to provide basic facts on the level and composition of household net worth in South Africa and its evolution since 1993 (see Figure 7.1). Before the early 2000s, real average wealth per adult stagnated at around 240,000 Rand. It then rapidly increased by about 30%, before stabilizing at some 320,000 Rand after the 2008 financial crisis. The net wealth to national income ratio has remained relatively stable since 1993, ranging from 2.5 (before 2003) to 2.8 (after 2008).

In 2018, financial and non-financial assets respectively amounted to two years and one year of national income. Pension assets represented the biggest component of financial assets (73% of national income), closely followed by equities and fund shares (51%), bonds and interest deposits (45%), and life insurance assets (35%). Meanwhile, the bulk of non-financial assets consisted of owner-occupied housing (75% of national income), followed by tenant-occupied housing (24%) and business assets (12%). The total liabilities of the household sector amounted to about 54% of national income, divided into mortgage debt (25%) and non-mortgage debt (28%). Household debt rose significantly between 2000 and 2008, in large part due to a boom in mortgage advances (see appendix figure G.5).

Finally, based on the estimation made by Alstadsæter, Johannesen, and Zucman (2018), we assume that 11.8% of South African GDP was held offshore in 2007, and, in the absence of data on the evolution of wealth held in offshore tax havens, that this share has remained constant throughout the period. This is a conservative assumption, given that global offshore wealth is known to have steadily risen in the past decades. Given the relative stability of wealth-income ratios, this implies that offshore wealth represented about 5% of net wealth throughout the period of interest (see appendix figure G.1).

7.3.2 The distribution of wealth in South Africa in 2017

Table 7.2 provides information on the number of adults (above 20 years old), the entry thresholds, the average wealth and the share of wealth of various groups of the wealth distribution in 2017.

Average wealth varies hugely across the distribution. The bottom 50% of the South African population have negative net worth: the levels of the debts that they owe exceeds the market value of the assets they own. The middle 40% of the distribution—individuals located between the median and the 90^{th} percentile—have a net worth more than twice lower than the national average. Together, the bottom 90% of the South African adult population own about 14% of total personal wealth in the economy, while the remaining 86% belong to the top decile. The average wealth of the bottom 90% of the population is about six times lower than the national average, compared to nine times higher among the top 10%.

Ownership is not only polarised between top and bottom wealth groups, it is also extremely concentrated within the top 10%. The top 1% of the South African adult population (350,000 individuals) own 55% of aggregate personal wealth, and the top 0.1 % alone (35,000 individuals) own almost a third of wealth. The top 0.01% of the distribution, amounting to some 3,500 individuals, own about 15% of household wealth, greater than the share of wealth owned by the bottom 90% as a whole (32 million individuals). Their average wealth is more than 1,500 times greater than the national average, and 9,000 times greater than the average of the bottom 90%.

7.3.3 The composition of personal wealth across the distribution

The extreme degree of wealth inequality observed in South Africa is in large part driven by the relative exclusion of poorer wealth groups from any form of wealth accumulation, and by the concentration of all forms of assets at the top end. Table 7.3 provides some insights into this polarisation by showing the share of different types of assets held by wealth groups across the distribution. The top 10% own more than

55% of all forms of assets, including pension assets, housing wealth, unincorporated business assets and currency, notes and coins. They own virtually all (99.8%) bonds and stock in the economy. The top 1% alone holds more than a tenth of all forms of assets and a bit more than 95% of all bonds and stocks. Currency and housing wealth are the least concentrated forms of wealth, yet low wealth groups only possess a small share of them: the bottom 50% of the wealth distribution own about 10% of currency, notes and coins, and less than 15% of housing assets.

Figure 7.2 provides another view of the link between asset types and wealth groups by representing the portfolio composition of percentiles of the wealth distribution in 2017. Currency, notes and coins are the main form of assets held by poorest South African adults, while owner-occupied housing, pensions and life insurance form the majority of assets for most of the distribution within the bottom 90%. Unincorporated business assets represent a small share of portfolios for the upper-middle class. Bonds and stocks, finally, represent a large share of wealth for the top 1% and the bulk of assets held within the top 0.1%.

7.3.4 Wealth and age

Based on available information on age from the PIT data, we can document to what extent wealth accumulation through the life cycle contributes to reducing or exacerbating inequalities.[12] Figure 7.3 shows a stable relationship between age and average wealth over the 2012-2017 period. Average net worth rises significantly and linearly between ages 20 and 55: individuals aged between 20 and 25 have an average net worth lower than 25% of the national average, while those aged between 50 and 55 are between 50% and two times wealthier than the average adult. Average wealth then stabilises between ages 50 and 65 and decreases slightly for older individuals, but still remains more than 50% higher than the national average for individuals older than 75. Interestingly, this pattern is almost perfectly similar to that found in the case of France see Garbinti, Goupille-Lebret, and Piketty, 2017, figure 5.

Although average wealth does vary significantly across age groups, age differences cannot account for observed wealth disparities. Indeed, levels of wealth concentration within each age group are almost perfectly similar to those measured among the full population. The share of wealth held by the top 10% exceeds 85%, and the top 1% share is higher than 55%, whether one restricts the analysis to those aged between

[12]There are many other important categories to investigate in the context of wealth inequality in South Africa. Unfortunately, the only relevant covariate present in PIT data is age. We leave the study of other dimensions of wealth inequality (race, gender, geography, etc.) for future research.

20 and 39, between 40 and 59, or older than 60 (figure 7.4). Altogether, this implies that individuals across the wealth distribution do accumulate at relatively similar paces but start from very different initial endowments. This suggests that inherited wealth could play a central role in explaining levels of wealth concentration observed in South Africa.[13]

7.3.5 Long-run trends and comparative perspectives

We conclude this section by highlighting the most notable facts arising from the comparison of our results over time and across countries. Figure 7.5 plots the evolution of the share of wealth accruing to the top 10% in South Africa (our estimates), together with that from all other countries where a similar method could be applied: China, Russia, India, the United Kingdom, France and the United States. In the long run, and despite a 30% growth in real average wealth per adult, wealth concentration has remained remarkably stable in South Africa, increasing between 2005 and 2010 before gradually stabilizing back to its pre-2000 level. Notwithstanding these short-term fluctuations and the fact that wealth concentration has increased in all other countries, South Africa has remained significantly more unequal than all these countries throughout the entire period. The South African top 10% wealth share has fluctuated between 80% and 90% during the 1993-2017 period, while it has remained below 75% in the US, 70% in Russia and China, 65% in India and 55% in France or the United Kingdom. The same result holds for the top end of the distribution: the top 1% wealth share was 55% in South Africa in 2017, compared to 43% in Russia, 39% in the United States, 31% in India, 30% in China and less than 25% in France and the UK (figure 7.6).

Having a closer look at our series, we can bring out two additional observations. First, the rapid increase in wealth concentration between 2005 and 2008 was in large part due to a strong fall in the bottom 90% share driven by the boom and bust in mortgage advances in the 2000s, which temporarily drove a higher share of households into negative net worth. Between 2004 and 2008, in particular, mortgage debt increased from 9% of net household wealth to almost 15%, before decreasing back to 9% in 2018 (see appendix figure G.5). This temporary fall in bottom wealth shares driven by expanding debts mirrors that observed in the US at about the same period (see appendix figure G.4).

[13] Notice that the estimates presented here correspond to individual series, rather than to "equal-split" series where wealth would be split equally among household adult members. In practice, splitting wealth among household members would imply redistributing wealth to younger individuals, thereby making the wealth-age profile less steep. This would reinforce our argument that age is not a primary determinant of wealth inequality in South Africa.

Secondly, it is worth noticing that while the top 10% share has remained broadly stable, there seems to have been an increase in wealth concentration within the top 10%. Between 1993 and 2017, the top 1% share grew from 54% to 57% and the top 0.1% share from 22% to 31% (see appendix figure G.3). This is likely due to the combination of two factors: the rise in the share of non-pension financial assets, from 19% to 24% of net household wealth between 1992 and 2018, and the increase in wage inequality in South Africa during this period, which indirectly affected the distribution of pension assets.

Overall, it is particularly striking that wealth inequality has remained at extreme and stable levels in South Africa in spite of the many progressive policies that have been pursued since the early 1990s. All discriminatory laws were abolished by 1991 and a new constitution was adopted in 1994. Since then, South Africa's successive governments endorsed several ambitious socio-economic policy frameworks whose primary objectives consistently included reducing economic inequality inherited from colonial and apartheid regimes.[14] Yet, wealth inequality has remained remarkably stable over the past three decades. In line with our observations on the role of inheritance in explaining constant wealth disparities within age groups, our long-term series suggest that asset allocations before 1993 may still contribute to shape wealth inequality in recent years, despite the many reforms to address these lasting disparities.

7.4 Robustness checks

In this section, we contrast our results with those obtained using alternative methodologies. We then discuss how sensitive our estimates are to different assumptions regarding the distribution of debts, the measurement of housing wealth, and equivalence scales.

7.4.1 Comparing methodologies: direct measurement, rescaling, and survey-based mixed approaches

In our baseline "mixed approach" to estimate wealth inequality in South Africa, we have combined surveys and exhaustive tax microdata to capitalise income flows and match wealth aggregates to macroeconomic balance sheets. To shed light on the

[14] Including the Reconstruction and Development Programme (RDP - 1994); Growth, Employment and Redistribution (GEAR - 1996); Accelerated and Shared Growth Initiative for South Africa (ASGISA - 2005); New Growth Path (NGP - 2010); and National Development Plan (NDP - 2013).

contributions of these various data sources and methodological steps, it is useful to compare our benchmark series with three alternative specifications: one in which we estimate wealth inequality from self-reported assets and liabilities in household surveys ("direct measurement"), one in which we rescale these reported assets and liabilities to macro totals ("rescaling"), and one in which we apply our mixed approach directly to surveys, without combining them with tax data.

Direct measurement. In South Africa, the only publicly available data source allowing direct measurement for the entire spectrum of household wealth components is the NIDS survey. The direct measurement approach implies that figures are not consistent with macroeconomic statistics, both in terms of levels and composition of household wealth. In the case of the NIDS, this implies overstating the total value of housing assets and understating the significance of non-pension financial assets (see appendix section G.2).

Rescaling. A second way of measuring the distribution of wealth consists in assuming that the distribution of recorded wealth components and their correlation is relatively well measured by the household survey, but that it is mainly their average amounts that are understated or overstated. In this case, one can obtain an estimate of the wealth distribution by effectively scaling up individual-level assets and liabilities in the NIDS surveys to match the totals recorded in the national balance sheets. This has the advantage of ensuring consistency with macroeconomic aggregates, as in our mixed approach. The drawback is that self-reported wealth components may be more prone to measurement error than self-reported income flows, potentially creating a number of outliers and yielding implausible levels of wealth inequality.

Survey-based mixed approach. A third way of measuring wealth inequality, in the absence of tax microdata, is to directly apply our mixed methodology to household surveys, capitalising relevant income flows and rescaling assets that do not generate income flows to macro totals. To the extent that household surveys tend to underestimate top income inequality (albeit much less than top wealth inequality), we may expect estimated wealth inequality to be lower when relying solely on surveys than when combining surveys with tax data.

Results. Table 7.4 compares estimates of the share of wealth held by the bottom 50%, the middle 40%, the top 10%, the top 1% and the top 0.1% derived from these different methodologies. Wave 4 and 5 of the NIDS are the only surveys collecting direct data on wealth and thus for which estimates from the three methodologies can be compared. Three main results stand out from these figures.

First, all approaches converge in revealing an extreme degree of wealth concentration. Regardless of the methodology, the share of wealth held by the bottom 50% is estimated to be consistently negative, while the top 10% is higher than 80%. The fact that wealth inequality in South Africa is substantially larger than in any other country for which a similar measurement method has been applied is therefore robust to alternative methodologies.

Secondly, while methodologies converge when it comes to large groups (e.g. the top 10% and the bottom 90%), they yield much more variable results when it comes to measuring wealth concentration at the top of the distribution. Direct measurement in the NIDS surveys implies a top 0.1% share below 10%, i.e. more than twice lower than most of the results obtained from rescaling or the mixed approach. This is due to the extremely poor coverage of non-pension financial assets in the NIDS: the total reported value of bonds and stock, two types of assets that are overwhelmingly concentrated at the top end of the wealth distribution, does not exceed 4% of macro totals in both waves of the survey (see appendix table G.2). Rescaling financial assets to balance sheets totals or capitalising income flows corrects for this micro-macro discrepancy, moving the estimates closer to those obtained with our benchmark methodology.[15]

Thirdly, the survey-based mixed approach yields relatively close results across years and data sources: the top 10% share lies between 85% and 90%, and the top 1% is estimated to be between 50% and 60% in most cases. Most importantly, these estimates are very close to those obtained when combining surveys with PIT data: despite their tendency to underestimate top income inequality, surveys can still be usefully exploited to estimate wealth concentration using the mixed approach. A careful look at the particular structure of capital income concentration can help solve this apparent paradox. The relative consistency between the two sources is mainly due to the fact that both in the surveys and the tax data, financial incomes (interest, dividends and rental income) are extremely concentrated, so that both sources imply attributing a substantial share of wealth—and in particular of tenant-occupied housing, bonds and shares—to the top 0.1% of the distribution.

In summary, our results point to the key significance of bridging the micro-macro gap. Because surveys tend to omit the bulk of financial assets, studies solely relying on

[15] Also notice that wealth inequality between the top 10% and the bottom 90% is significantly larger under the rescaling approach than when relying on the mixed approach. This is essentially due to the fact that scaling up debts to balance sheets totals creates a large number of households with strongly negative net worth (the bottom 50% goes down by several percentage points), especially in the NIDS where assets and liabilities suffer from important underreporting issues.

self-reported household wealth are likely to very strongly underestimate top wealth inequality. By contrast, capitalising income flows to match macro totals can prove to be a more reliable methodology, even in the absence of income tax microdata. This opens new avenues for estimating wealth inequality in other emerging countries, where tax microdata might not be available yet where surveys collecting data on income can be usefully combined with data from national accounts.

7.4.2 Debts, housing wealth and equivalence scales

We conclude this paper by briefly discussing three sources of concern related to the mismeasurement of household debt, the underestimation of total housing wealth, and the distribution of wealth within households.

Mismeasurement of household debt. One concern with our estimates is that debt is self-reported in household surveys. By rescaling reported debts to macro totals, we might overestimate the number of households with negative net worth, especially given that surveys tend to only capture a small fraction of private debt (see appendix table G.3). In order to evaluate the potential significance of this bias, we compare the evolution of household net worth inequality with that of household assets inequality (excluding debts) in appendix figure G.14.

Two key results emerge from this comparison. First, excluding debt systematically reduces wealth inequality, but only moderately: the top 10% have owned a consistent 80% of assets and the top 1% about 45% of assets since 1993. Secondly, debt dynamics appear to drive virtually all fluctuations in wealth inequality over time: wealth concentration has followed ups and downs, while the concentration of assets has remained remarkably stable. This points to the role of credit dynamics in accounting for short-run trends in wealth disparities. The rise and fall of wealth inequality visible in our series before and after the 2007-2008 financial crisis, in particular, coincides with the mortgage credit boom and bust (see appendix figure G.5).

Underestimation of housing wealth. A second concern relates to the aggregate value of housing wealth in South Africa. Indeed, housing appears to be the only asset class for which reported values in surveys are substantially *higher* than in balance sheets totals (see appendix table G.2). Whether this inconsistency arises from survey respondents overestimating the value of their home or from the SARB underestimating housing wealth remains an open question.[16] For consistency and

[16]Notice that this issue is not one specific to South Africa—in the United States too, survey values have been found to be higher than in balance sheets. Which source of information provides

comparability with existing studies, we choose to rely on SARB statistics. However, we report in the appendix series in which we assume that total housing wealth is underestimated by a factor of 2 (see figures S12 and S13). Unsurprisingly, as housing is one of the least unequally distributed asset in South Africa, increasing its average value reduces wealth inequality. Yet, because all assets are strongly concentrated at the top end, including housing (see table 7.3), it affects our main results only moderately, with the top 10% share still reaching about 80% and the top 1% about 40%.

Equivalence scales. Lastly, one might be concerned that the equivalence scale used in this paper—allocating wealth components directly to individuals, and therefore not accounting for wealth sharing within households—may lead to overestimating wealth inequality. It might also lead to overstating wealth inequality more in South Africa than in countries such as France, given that multi-generational households and intra-familial sharing agreements might be more common in the former than in the latter.

We investigate this concern in appendix figures G.10 and G.11, which compare our ”individual” series to that obtained when splitting wealth equally among all household members (”per capita” series), or among all adult household members (”broad equal-split” series). We find that changes in equivalence scales only moderately affect wealth inequality, which is highest in the individual series and lowest in the broad equal-split series. The top 10% share exceeds 80%, and the top 1% share 45%, in all three specifications.

7.5 Conclusion

This paper systematically estimated the distribution of household wealth in South Africa since 1993 by combining all relevant macro and micro data sources. Our results have revealed unparalleled levels of wealth concentration, with the top 1% owning a higher share of wealth than the bottom 99%. These extreme inequalities have remained remarkably stable since the end of the apartheid regime, despite the significant economic growth and the major social transformations that the country has undergone since then. They extend to all forms of assets, from housing to financial capital, which are consistently held by individuals located at the top end.

Methodologically, our results point to the substantial limitations of wealth surveys,

the most accurate estimate of the market value of housing wealth remains debated (Blanchet, 2016; Dettling et al., 2015; Henriques and Hsu, 2014).

which vastly underestimate financial assets and are therefore incapable of properly measuring wealth inequality within the top 10%. Instead, we have shown that bridging the micro-macro gap by capitalising relevant income flows, even in the absence of tax microdata, can yield more consistent and meaningful estimates of the wealth distribution. This comes as good news for researchers aiming at tracking the dynamics of wealth concentration in countries where tax microdata might not be accessible, yet where household income surveys and macroeconomic balance sheets exist and can be combined.

We see at least two avenues for future research. First, our estimates of wealth inequality could be refined if better information on dividends and income received through unit trusts were made available to researchers (see the discussion in appendix G.3). Information on these forms of income are collected on a regular basis by the South African Revenue Service, but are not yet accessible. We hope that access to these data sources will enable future studies to have a more granular picture of the composition of wealth and its dynamics at the very top of the distribution.

Secondly, our findings on the stability of wealth inequality since 1993 call for further research on the dynamics and weight of inherited wealth relative to that of newly created and accumulated wealth in the post-apartheid era. This would likely require combining other complementary data sources—such as estate duty data, credit data or panel data on income and savings—and modelling the joint dynamics of savings, inter-generational transmission, and household debt.

Bibliography

Alstadsæter, Annette, Niels Johannesen, and Gabriel Zucman (2019). "Tax evasion and inequality". In: *American Economic Review* 109.6, pp. 2073–2103.

Alstadsæter, Annette, Niels Johannesen, and Gabriel Zucman (2018). "Who owns the wealth in tax havens? Macro evidence and implications for global inequality". In: *Journal of Public Economics* 162.1, pp. 89–100.

Alvaredo, Facundo et al. (2018). *World Inequality Report 2018*. Harvard University Press.

Bertrand, Marianne et al. (2020). "Tax-exempt Lobbying: Corporate Philanthropy as a Tool for Political Influence". In: *American Economic Review* 110.7, pp. 2065–2102.

Blanchet, Thomas (2016). "Wealth inequality in Europe and in the United States: estimations from surveys, national accounts and wealth rankings". Paris School of Economics Master Thesis.

Blanchet, Thomas, Ignacio Flores, and Marc Morgan (2018). "The Weight of the Rich: Improving Surveys Using Tax Data".

Blanchet, Thomas, Juliette Fournier, and Thomas Piketty (2017). "Generalized Pareto Curves: Theory and Applications".

Blanchet, Thomas et al. (2021). *Distributional National Accounts Guidelines: Methods and Concepts Used in the World Inequality Database*. Tech. rep.

Bombardini, Matilde and Francesco Trebbi (2020). "Empirical models of lobbying". In: *Annual Review of Economics* 12, pp. 391–413.

Chatterjee, Aroop (2019). "Measuring wealth inequality in South Africa: An agenda". In: *Development Southern Africa* 36.6, pp. 839–859.

Daniels, R. and T. Augustine (2016). "The measurement and distribution of household wealth in South Africa using the National Income Dynamics Study (NIDS) Wave 4". NIDS Discussion Paper 2016/10.

Dettling, L. J. et al. (2015). "Comparing Micro and Macro Sources for Household Accounts in the United States: Evidence from the Survey of Consumer Finances".

Finance and Economics Discussion Series, Divisions of Research & Statistics and Monetary Affairs of the Federal Reserve Board.

Ebrahim, Amina and Chris Axelson (2019). “The creation of an individual level panel using administrative tax microdata in South Africa”. UNU-WIDER Working Paper 2019/27.

Esteban, Joan and Debraj Ray (2006). “Inequality, lobbying, and resource allocation”. In: *American Economic Review* 96.1, pp. 257–279.

Garbinti, B., J. Goupille-Lebret, and T. Piketty (2018). “Income inequality in France, 1900-2014: Evidence from Distributional National Accounts (DINA)”. In: *Journal of Public Economics* 162.1. Dataset, pp. 63–77.

Garbinti, Bertrand, Jonathan Goupille-Lebret, and Thomas Piketty (2017). “Accounting for wealth inequality dynamics: Methods, estimates and simulations for France (1800-2014)”. WID.world Working Paper 2016/5.

Henriques, A. M. and J. W. Hsu (2014). “Analysis of Wealth Using Micro- and Macrodata: A Comparison of the Survey of Consumer Finances and Flow of Funds Accounts”. In: *Measuring Economic Sustainability and Progress*. Ed. by D. W. Jorgenson, J. S. Landefeld, and P. Schreyer. University of Chicago Press. Chap. 9, pp. 245–274.

Hundenborn, Janina, Ingrid Woolard, and Jon Jellema (2018). “The effect of top incomes on inequality in South Africa”. UNU-WIDER Working Paper 2018/90.

Kleven, Henrik et al. (2020). “Taxation and Migration: Evidence and Policy Implications”. In: *Journal of Economic Perspectives* 34.2, pp. 119–142.

Korinek, Anton, Johan A. Mistiaen, and Martin Ravallion (2006). “Survey nonresponse and the distribution of income”. In: *Journal of Economic Inequality* 4.1, pp. 33–55.

Londoño-Vélez, Juliana and Javier Ávila-Mahecha (2021). “Enforcing Wealth Taxes in the Developing World: Quasi-Experimental Evidence from Colombia”. In: *American Economic Review: Insights* 3(2).

Mbewe, Samson and Ingrid Woolard (2016). “Cross-sectional features of wealth inequality in South Africa: Evidence from the National Income Dynamics Study”. NIDS Discussion Paper 2016/12.

Muellbauer, John and Janine Aron (1999). “Estimates of personal sector wealth for South Africa”. CEPR Discussion Paper No. 4646.

Orthofer, Anna (2016). “Wealth inequality in South Africa: Evidence from survey and tax data”. REDI3x3 Working Paper 15.

Saez, Emmanuel and Gabriel Zucman (2016). "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data". In: *Quarterly Journal of Economics* 131.May, pp. 519–578.

United Nations (2009). *System of National Accounts 2008*. United Nations, European Commission, International Monetary Fund, Organisation for Economic Co-operation and Development, and World Bank joint publication.

Zucman, Gabriel (2019). "Global Wealth Inequality". In: *Annual Review of Economics* 11.1, pp. 109–138.

Figure 7.1: The evolution of household wealth in South Africa, 1993-2018

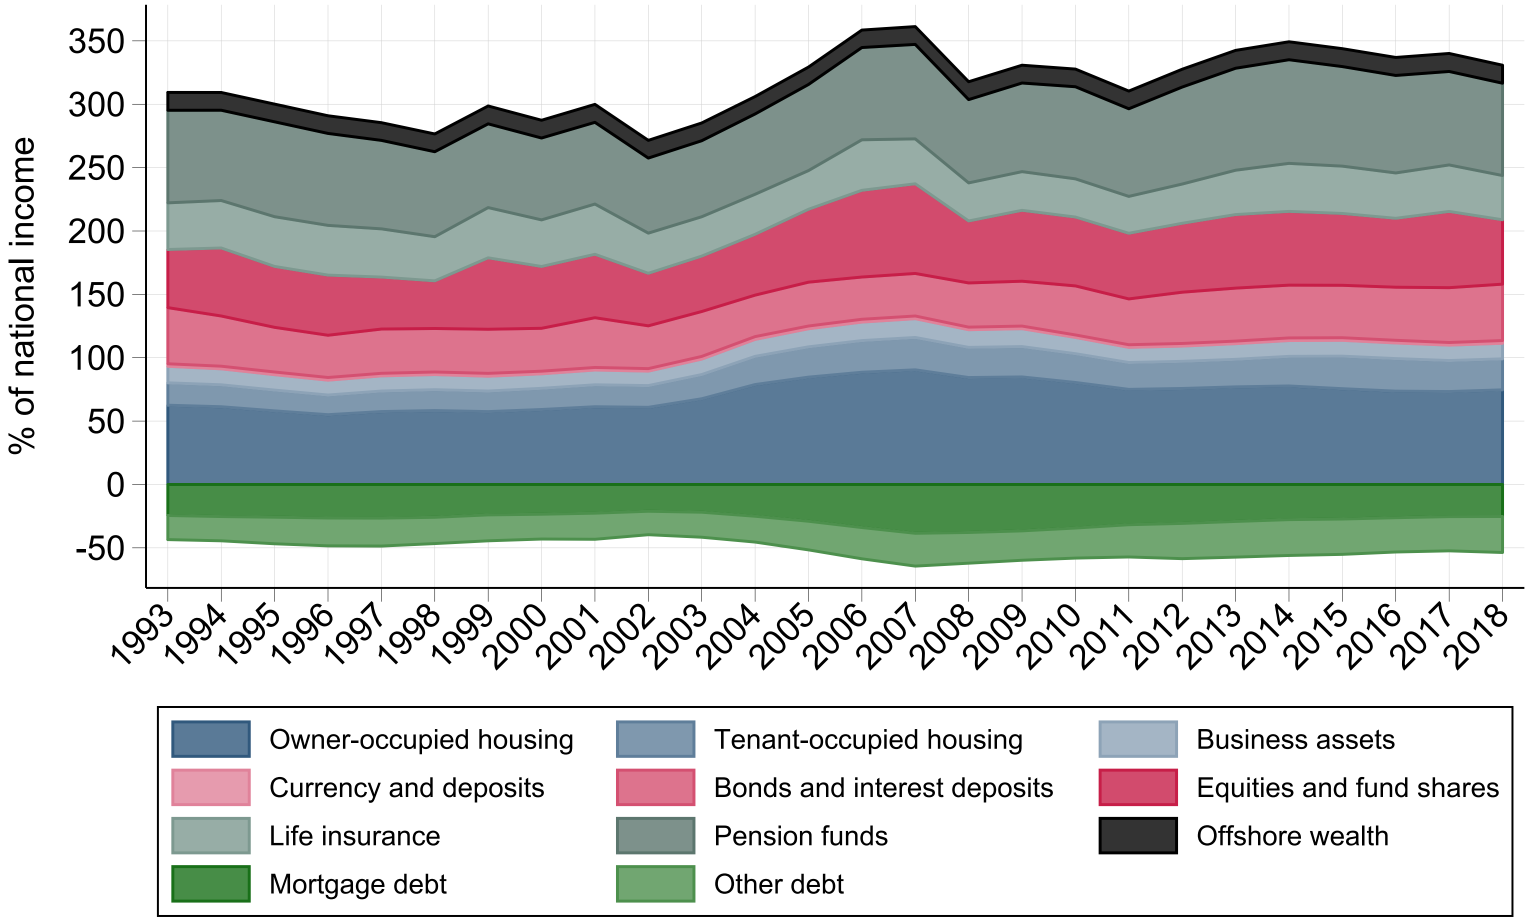

Notes. Authors' compilation based on data from the South African Reserve Bank. This figure shows the level and composition of household wealth in South Africa between 1993 and 2018, expressed as a share of the net national income.

Figure 7.2: The composition of assets by wealth group in 2017

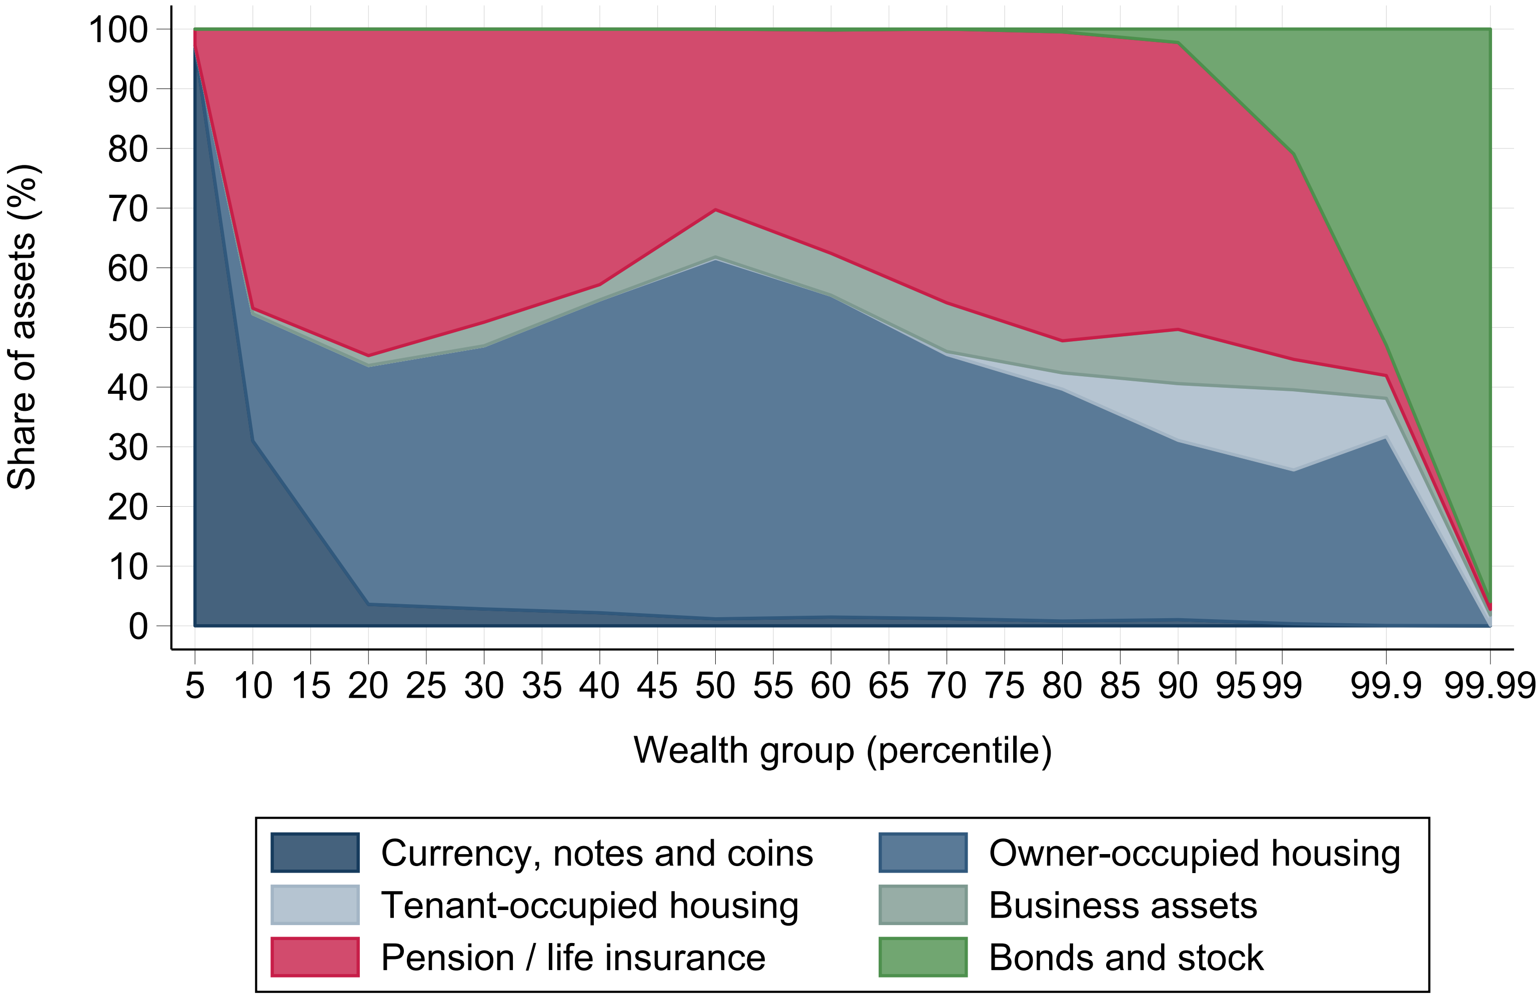

Notes. Authors' computations combining surveys, tax microdata and macroeconomic balance sheets statistics. The figure shows the composition of assets of various groups in the distribution of household assets in South Africa in 2017. The unit of observation is the adult aged 20 or above. The results come from the harmonised survey data file, and wealth is split equally among adult members of the household, except for the top 1% and above for which the individual data built from the combined survey and tax microdata are used.

Figure 7.3: Average wealth by age relative to average wealth per adult, 2012-2017

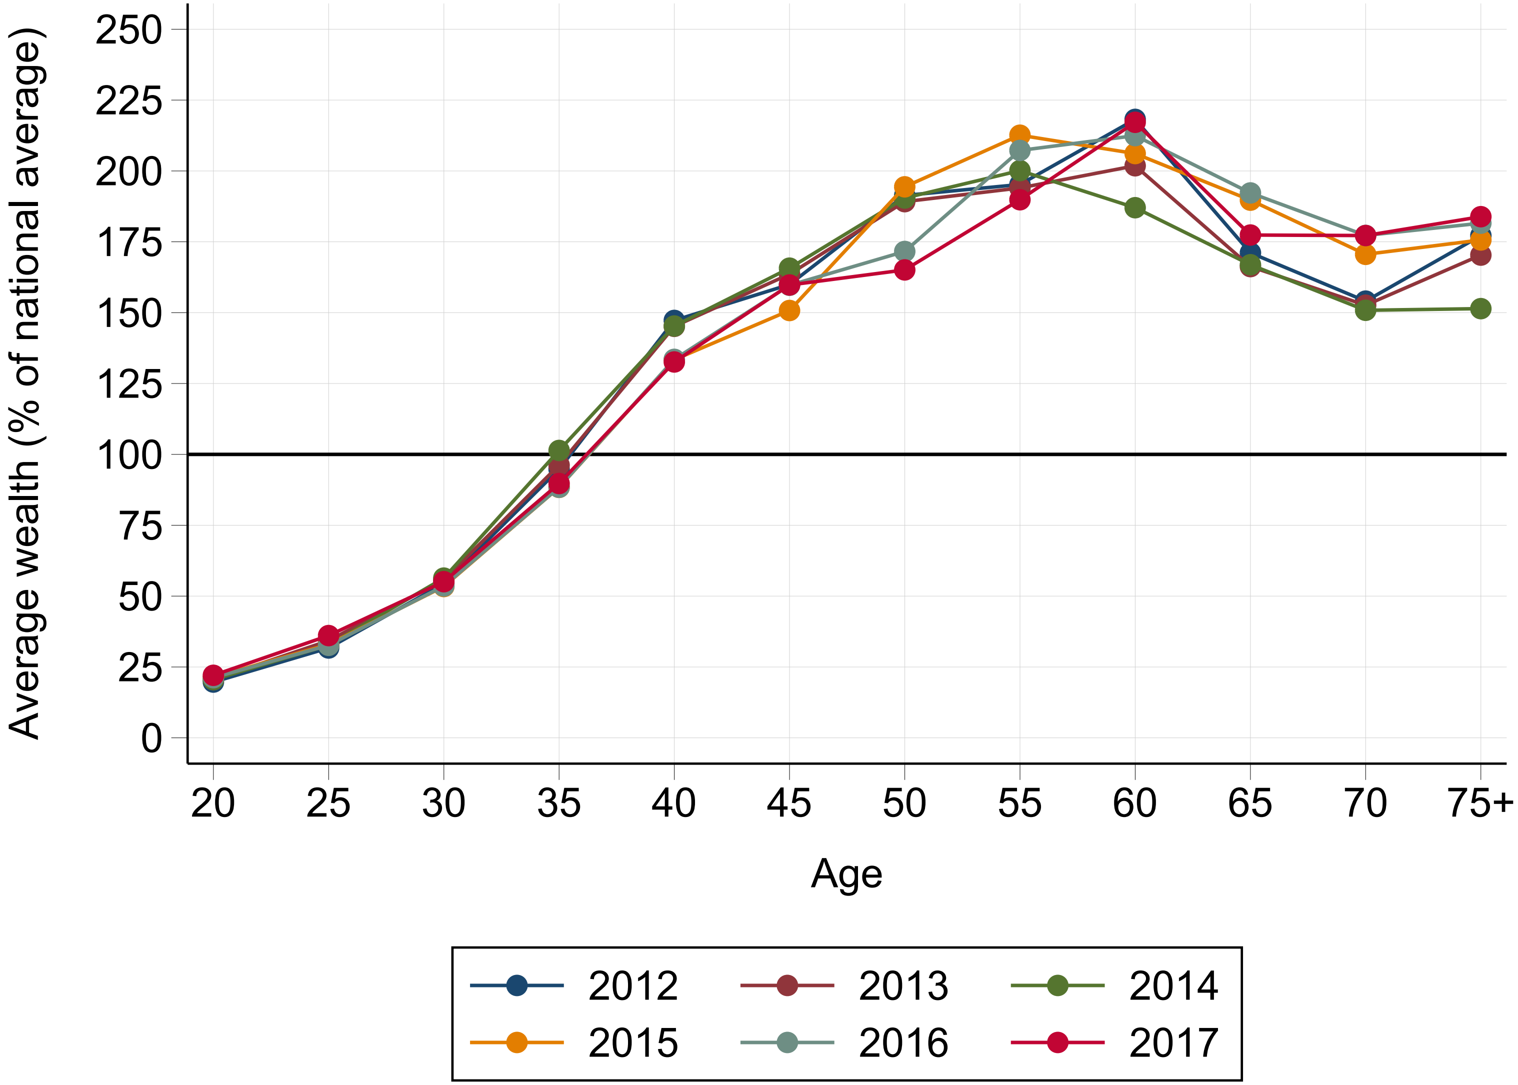

Notes. Authors' computations combining surveys, tax microdata and macroeconomic balance sheets statistics. The figure shows the mean net worth of South African adults by age group relative to the national average. The unit of observation is the individual adult aged 20 or above.

Figure 7.4: Wealth inequality within age groups, 2010-2017

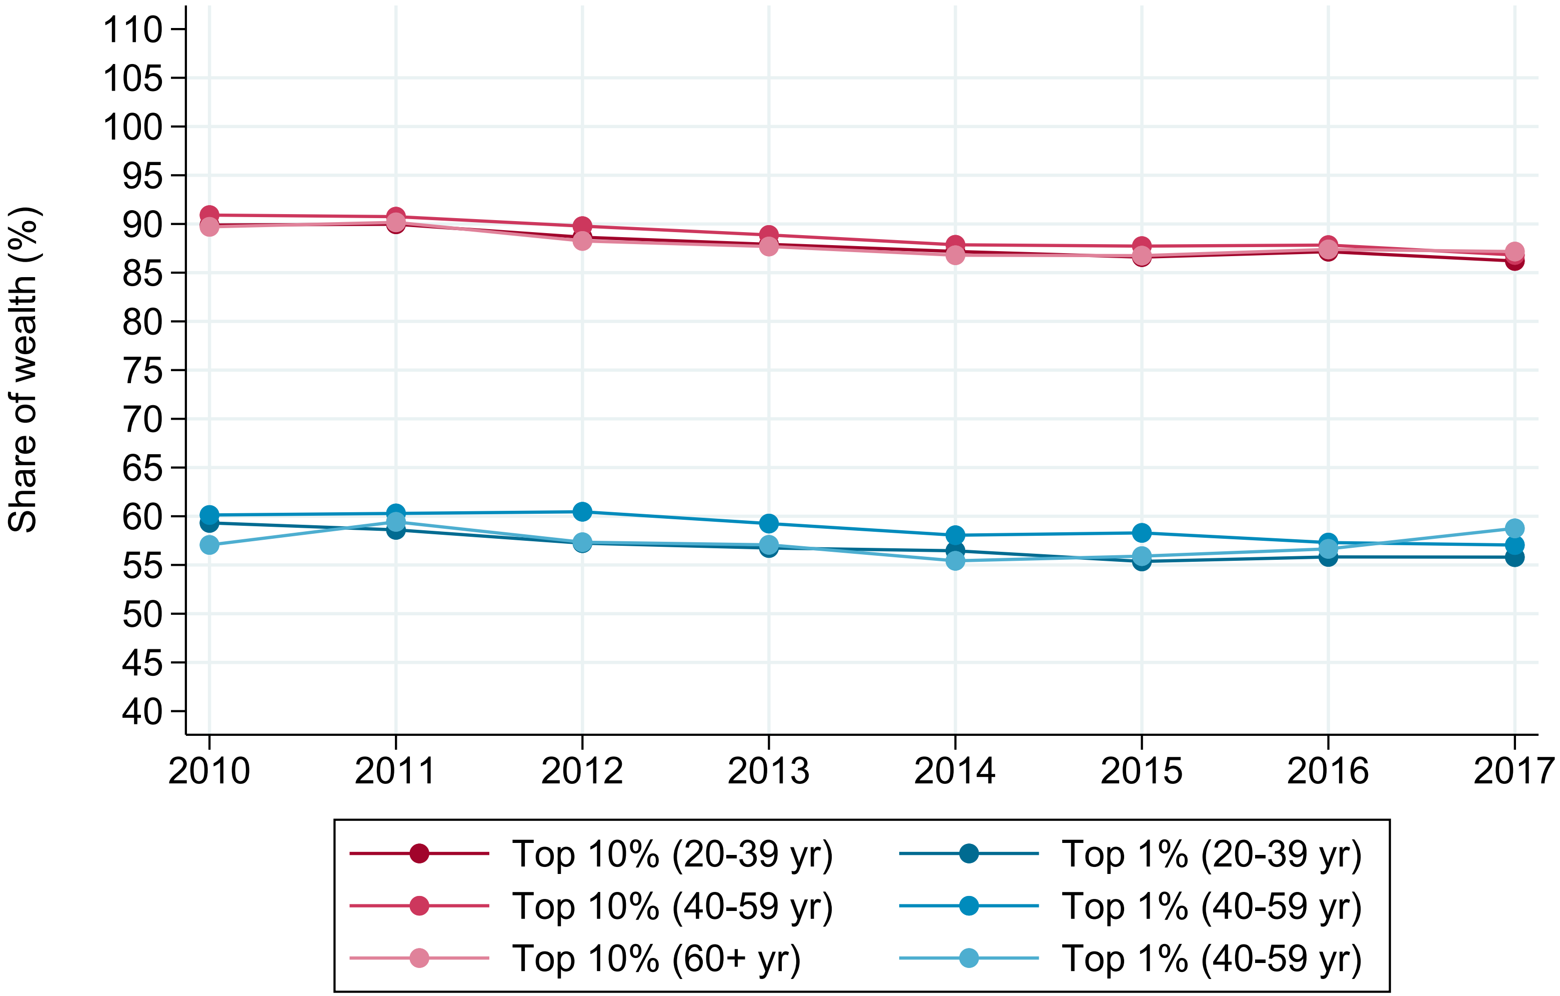

Notes. Authors' computations combining surveys, tax microdata and macroeconomic balance sheets statistics. The figure shows top 10% wealth share and the top 1% wealth share estimated when splitting the South African population into three age groups (20-39 years old, 40-59 years old, and 60+ years old). The unit of observation is the individual adult aged 20 or above.

Figure 7.5: South African wealth inequality in comparative perspective: Top 10% wealth share

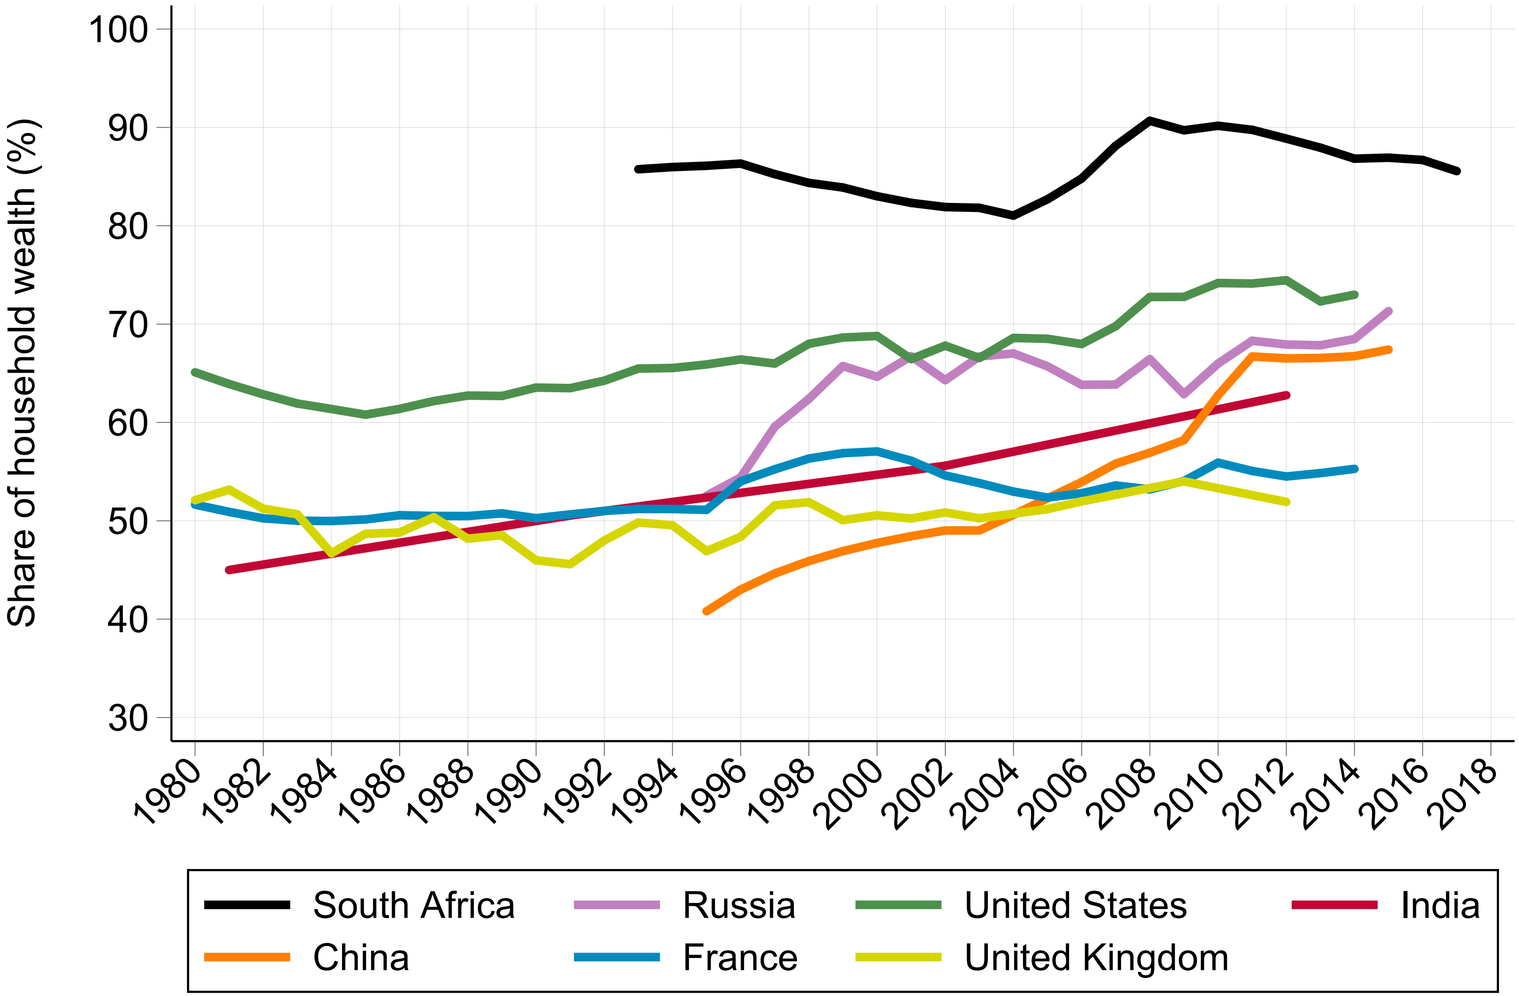

Notes. Authors' computations combining surveys, tax microdata and macroeconomic balance sheets statistics for South Africa; World Inequality Database (`http://wid.world`) for other countries. The figure compares the top 10% wealth share in South Africa to that of other countries. The unit of observation is the individual adult aged 20 or above. Wealth is individualised (South Africa) or split equally among adult household members (other countries).

Figure 7.6: South African wealth inequality in comparative perspective: Top 1% wealth share

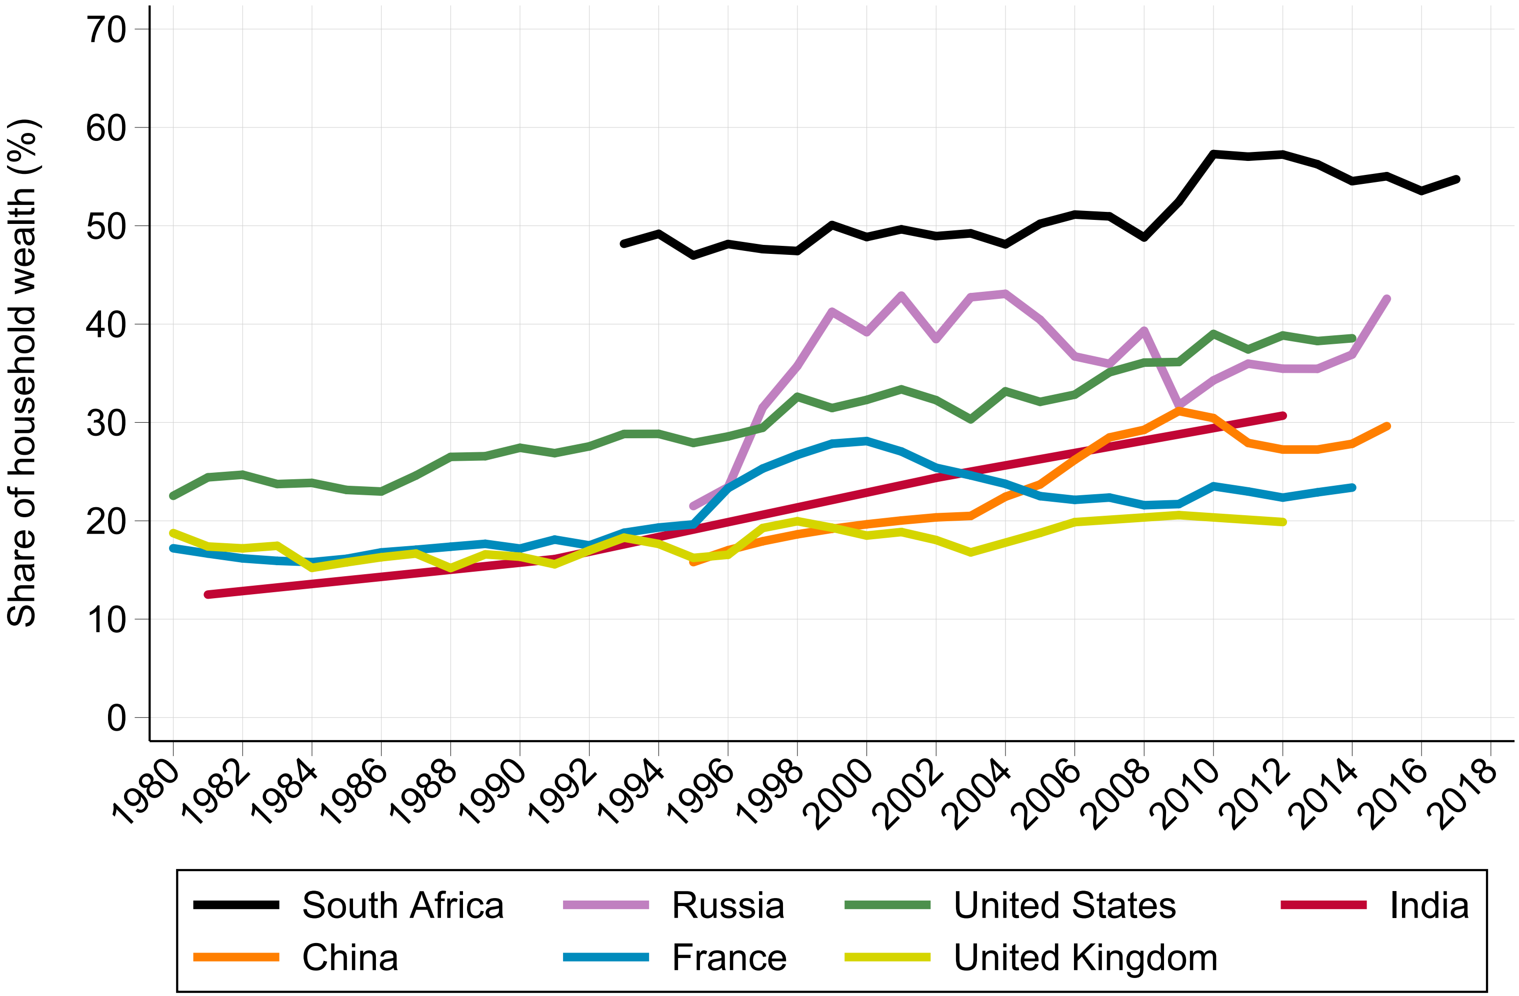

Notes. Authors' computations combining surveys, tax microdata and macroeconomic balance sheets statistics for South Africa; World Inequality Database (`http://wid.world`) for other countries. The figure compares the top 1% wealth share in South Africa to that of other countries. The unit of observation is the individual adult aged 20 or above. Wealth is individualised (South Africa) or split equally among adult household members (other countries).

Table 7.1: Estimating the distribution of personal wealth in South Africa: a mixed approach

Asset / liability	Variable	Measurement method
Non-financial assets		
Owner-occupied dwellings	Value of home	Rescaling
Tenant-occupied dwellings	Rental income	Capitalisation
Business assets	Business income	Capitalisation
Financial assets		
Pension assets	Pension contributions and pension income	Mixed method
Life insurance assets	Factor income	Mixed method
Currency, notes and coins	Bank account balance	Rescaling
Bonds and interest deposits	Interest income	Capitalisation
Corporate shares and equity	Dividends	Capitalisation
Liabilities		
Mortgage debt	Reported debt and house value	Mixed method
Other debts	Reported debts and consumption	Mixed method

Notes: The table shows the methodological approach used to estimate the distribution of the different assets and liabilities reported in the household balance sheets. Direct measurement corresponds to reported data on the market value of assets or liabilities in household surveys. Capitalisation corresponds to assuming that the distribution of an asset follows that of one or several corresponding income flows.

Table 7.2: The distribution of personal wealth in South Africa in 2017

	Number of adults	Threshold (2018 R)	Average (2018 R)	Average (2018 PPP $)	Wealth Share (%)
Full population	35,600,000		326,000	52,200	100
Bottom 90% (p0p90)	32,040,000		52,300	8,400	14.4
Bottom 50% (p0p50)	17,800,000		-16,000	-2,600	-2.5
Middle 40% (p50p90)	14,240,000	27,700	138,000	22,000	16.9
Top 10% (p90p100)	3,560,000	496,000	2,790,000	447,000	85.6
Top 1% (p99p100)	356,000	3,820,000	17,830,000	2,860,000	54.7
Top 0.1% (p99.9p100)	35,600	30,350,000	96,970,000	15,540,000	29.8
Top 0.01% (p99.99p100)	3,560	146,890,000	486,200,000	77,920,000	14.9

Source: Authors' computations combining surveys, tax microdata, and macroeconomic balance sheets statistics. *Notes*: The table shows the distribution of household wealth in South Africa in 2017. The unit of observation is the individual adult aged 20 or above. Wealth thresholds are in 2018 Rands.

Table 7.3: Share of total assets held by wealth group by asset class (%), 2017

	Currency	Business assets	Housing	Pensions / life insurance	Bonds & Stock
Bottom 90% (p0p90)	37.3	40.4	41.2	36.2	0.2
Bottom 50% (p0p50)	9.7	1.4	14.0	5.3	0.0
Middle 40% (p50p90)	27.7	39.1	27.2	30.9	0.2
Top 10% (p90p100)	62.7	59.6	58.8	63.8	99.8
Top 1% (p99p100)	10.6	41.9	27.8	14.1	95.2
Top 0.01% (p99.99p100)	1.5	13.4	8.5	2.1	62.7
% of total assets	0.6	3.6	28.8	32.5	34.6

Source: Authors' computations combining surveys, tax microdata, and macroeconomic balance sheets statistics.
Notes: The table shows the shares of different types of assets held by specific wealth groups in 2017. The unit of observation is the individual adult aged 20 or above. In 2017, the top 1% of South Africans in terms of net worth owned 95% of the bonds and corporate shares in the economy. Bonds and shares represented 34.1% of total household assets in the economy at this date. Figures may not add up due to rounding.

Table 7.4: Shares of household wealth held by groups in South Africa (%): survey-based results

	Bottom 50%	Middle 40%	Top 10%	Top 1%	Top 0.1%
Direct measurement					
NIDS, wave 4	-3.3	18.4	84.9	41.3	9.7
NIDS, wave 5	-0.5	16.9	83.6	40.2	8.6
Rescaling					
NIDS, wave 4	-8.2	10.9	97.3	58.3	24.6
NIDS, wave 5	-7.0	8.0	99.0	63.9	29.3
Mixed approach					
NIDS, wave 4	-4.5	14.5	90.0	58.5	25.1
NIDS, wave 5	-3.3	12.5	90.8	60.6	30.1
PSLSD, 1993	-1.3	11.9	89.4	51.7	20.6
IES, 1995	-5.1	15.2	89.8	50.6	23.7
IES, 2000	-1.7	14.4	87.3	52.9	26.1
IES, 2005	-0.3	13.5	86.7	54.1	28.6
LCS, 2008	-8.0	14.0	93.9	52.2	22.4
IES, 2010	-7.3	14.8	92.5	60.0	31.7
LCS, 2015	-3.3	14.2	89.0	51.1	20.0

Source: authors' computations from survey microdata. *Notes*: The table compares estimates of the share of household wealth owned by the bottom 50% (p0p50), the middle 40% (p50p90), the top 10% (p90p100), the top 1% (p99p100 and the top 0.1% (p99.9p100) obtained from household surveys using different methodological approaches. The unit of observation is the individual adult aged 20 or above. PSLSD: Project for Statistics on Living Standards and Development. IES: Income and Expenditure Survey. LCS: Living Conditions Survey. NIDS: National Income Dynamics Study.

Chapter 8

Income Inequality in Africa, 1990-2019: Measurement, Patterns, Determinants

Despite strong economic growth in many African countries during the last decades, human development and poverty indicators have not improved as expected. Indeed, reports by the World Bank on the attainment of Millennium Development Goal targets have shown that poverty has been decreasing in all regions of the world with the exception of the African continent (World Bank, 2015). This stands in contrast with statistics showing that African countries have enjoyed a significant resurgence in growth since the mid-1990s (Fosu, 2015). Solving this puzzle has fueled an interest in the study of inequality as one of the potential factors driving the weak poverty-reduction elasticity of growth in Africa (Fosu, 2009; Thorbecke, 2013).

Is Africa a high-inequality region? Given its high and persistent poverty levels, as well as its expected position in a worldwide Kuznets curve, poverty has long been the main focus of global development and research efforts in Africa (Barrett, Carter, and Little, 2006). Even if the Kuznets curve is no longer considered as a well-grounded empirical regularity, African inequality levels remain debated.[1] Analyses are typically made on the basis of household surveys, which provide a rich set of socioeconomic

[1] According to Bhorat and Naidoo (2017), the average Gini coefficient in Africa, based on household surveys, is 0.43 in 2014, whereas it is 0.39 in the rest of the developing world. However, heterogeneity is high across countries, and this high average level is driven by seven highly unequal countries with a Gini above 0.55, located mostly in Southern Africa: Angola, the Central African Republic, Botswana, Zambia, Namibia, Comoros, and South Africa. In terms of trends, the reported average African Gini has declined (it was 0.48 in the early 1990s), but this fall is largely due to trends in relatively low-inequality countries.

information but also have several important limitations when it comes to comparing income inequality levels across countries. From one country to another, studies using household surveys may inform on different types of welfare concepts (e.g., disposable income, taxable income, or consumption) and may use different ranking concepts (individuals, households, or other equivalence scales). Moving from one concept to another is likely to significantly modify the income distribution in a country and the estimated level of inequality. As a consequence, when studying inequality across countries or regions, it is key to compare distributions using consistent concepts and methodologies. In addition, household surveys tend to misreport top incomes due to both sampling and non-sampling errors, which typically leads to underestimating inequality. Average income or consumption levels reported in surveys are also often at odds with values reported in the national accounts. As a result, relying only on household surveys to compare inequality levels between Africa and other regions may lead to inaccurate estimates and conclusions (Blanchet et al., 2021).

A combination of sources is likely to provide a better approximation of Africa's true inequality levels and how it compares to other regions. Combining sources is preferable, for all sources have their own and specific drawbacks. Yet, this is a challenging task and it should be performed with care and transparency, as many issues remain imperfectly addressed (Ravallion, 2022). This paper makes a first attempt in that direction by putting together surveys, tax data, and national accounts in a systematic manner to estimate the present level of pretax income inequality in the continent, and more tentatively its evolution from 1990 to 2019. Our main finding is that Africa stands out as a region of extreme income inequality by international standards: with a top 10% income share of 55% and a bottom 50% share below 10%, Africa exhibits the highest gap between average incomes of the top 10% and bottom 50% (Figure 8.1). This overall high inequality level masks relatively large regional variations. These can in part be explained by historical determinants such as settler colonialism, postcolonial land reforms and socialist policies, and also potentially the influence of Islam. They might also reflect more proximate differences in productivity and employment in the agriculture and service sectors. We hope future research will be better able to disentangle the exact weight played by these different factors in accounting for the very high levels of African inequality.

The rest of the paper is organized as follows. Section 2 provides a brief overview of the literature on the distribution of income in Africa. Section 3 develops and implements a simple statistical method for combining (noisy) household survey data with (scarce) income tax data and (imperfect) national accounts. Section 4 exploits these new estimates to compare inequality in Africa to the rest of the world, explore historical correlates of African inequality, and discuss the link between redistribution policies and inequality.

8.1 Related Literature on Income Inequality and Growth in Africa

Research on the drivers of inequality in African countries is hindered by the lack of good-quality data, both on the distribution of living standards and on other economic or social indicators, but a few potential lines of explanation have been investigated.

A first strand of the literature has explored the link between the so-called "sub-optimal" structural transformation of the vast majority of African countries and inequality (Cornia, 2017, 2019). In theory, the growth of labor-intensive sectors, such as manufacturing or labor-intensive service activities should boost wage employment and thus reduce income inequality (Bhorat and Naidoo, 2017). Yet, unlike Europe during the Industrial Revolution or East Asia more recently, African economies did not experience a gradual shift from agriculture to manufacturing. Instead, the decline of the share of agriculture in GDP went to mining industries and to services.[2] As a result, the decrease in agricultural employment was entirely absorbed in services or in urban unemployment, as mining industries are very capital-intensive. Polarization of the service sector increased because of the development of informal activities, with very poor working conditions and low incomes. This led to a gradual "urbanization of poverty", as informal employment and urban unemployment spread (Ravallion, Chen, and Sangraula, 2007). This pattern is also consistent with the "urbanization without growth" documented by Fay and Opal (2000) in the late 20th century and by Jedwab and Vollrath (2015) in historical perspective. However, the urban-rural gap did not decrease significantly, because of the persistent urban bias in public spending, and because skilled urban residents were more able to exploit the opportunities brought about by liberalization, in particular cheaper food products. Besides, inequalities increased within the rural and within the urban sectors even when they decreased between sectors (Cornia, 2017).

The impact of African growth patterns on income inequality has been studied by looking at the joint evolution of sectoral value-added shares and Gini coefficients (Cornia, 2017). Gini coefficients fell in countries where the value-added share of modern agriculture, labor-intensive manufacturing, and modern services stagnated

[2] McMillan, Rodrik, and Inigo (2014) estimated that structural change in Africa between 1990 and 2005 made a negative contribution to overall economic growth of 1.3% per year on average.

or rose (for example in Ethiopia, Cameroon, Madagascar); it increased in countries with stagnating land yields, a decline of manufacturing sectors, a rise in the resource enclave and skill-intensive services, and urban informalization.

Focusing on agriculture, a strand of the literature argues that raising agricultural productivity could reduce inequality through increased rural incomes and the diversification of rural activities towards non-agricultural activities, thus favoring industrialization (Estudillo and Otsuka, 2010; Gollin, 2010; Pingali, 2010). Accordingly, a series of empirical papers underline the role of agricultural modernization in triggering growth, development, and reducing poverty and inequality (Bourguignon and Morrisson, 1998; Christiaensen, Demery, and Kuhl, 2011; Imai and Gaiha, 2014). This can be particularly relevant for Africa, which has not fully completed its agricultural transition yet. Yet, some authors underline that agricultural modernization has more impact on poverty than on inequality (Herault and Thurlow, 2009; Imai, Gaiha, and Cheng, 2016). Some also stress that equal land distribution is key to enabling agriculture to reduce both poverty and inequality (Christiaensen, Demery, and Kuhl, 2011; Griffin, Khan, and Ickowitz, 2002; Manji, 2006). Furthermore, although there are few economies of scale in agricultural production, these can be very important in the transport sector, especially for the international transport of agricultural products. Then, the exploitation of these rents can go hand in hand with increasing inequalities.

The influence of extractive industries on inequality has often been pointed out. In theory, extractive industries fuel income inequality, both through economic and institutional channels. According to Bhorat et al. (2017a), extractive industries are characterized by high capital intensity and limited employment creation, and mainly for skilled labor. Besides, the high cost of entry leads to monopolistic or oligopolistic market structures that favor high pricing and profits. A boom in the resource price can lead to the appreciation of the local currency, which can then disadvantage employment-intensive and often export-reliant sectors, or attract the best workers, draining them from the other sectors (the so-called "Dutch disease"). Extractive industries can also lead to the crowding out of non-resource investment (Papyrakis and Gerlagh, 2004), or hamper financial sector growth (Isham et al., 2003), and tend to fuel urbanization without industrialization, by sustaining the existence of "consumption cities" (Gollin, Jedwab, and Vollrath, 2016).

The links between institutions, public policies, and inequality have also been explored in the literature. Colonial legacy is a central issue in this regard (Walle, 2009). Under colonial rule, a minority of settlers held a very large fraction of wealth and positioned themselves at the top of the income distribution (Alvaredo, Cogneau, and Piketty, 2021). High wages were paid in a small formal sector formed by colonial administrations and a few companies specializing in the trade of natural resource exports (Cogneau, Dupraz, and Mesplé-Somps, 2021). This dualistic structure partly survived after independence and settlers' departure, giving rise to an "oligarchic bourgeoisie." By comparing five countries using carefully harmonized household survey data, Cogneau et al. (2007) find that income dualism between agriculture and other sectors explains much of cross-country differences in inequality; dualism is higher in the three former French colonies of Côte d'Ivoire, Guinea, and Madagascar than in the two former British colonies of Ghana and Uganda; using the same data, Bossuroy and Cogneau (2013) show that intergenerational mobility between agriculture and other sectors is also lower in the former French colonies, due to higher employment dualism and the concentration of non-agricultural activities in large cities. Cogneau (2007) argues that the decentralized management of colonial empires also produced large spatial inequalities, and Roessler et al. (2022) show that colonial investments in some cash crop producing areas have left a long-lasting imprint.

In terms of redistribution policies, Odusola (2017) shows that the fiscal space has been increasing over time in part due to an increase in the tax-to-GDP ratio. Institutions played a significant role in this increase: the Open Budget Index is highly correlated with fiscal space, which was also boosted by debt relief.[3] However, fiscal space in Africa remains low compared to the rest of the developing world, and, despite recent improvements in domestic taxation, in many countries tax revenue remains highly dependent on mineral extraction (Cogneau et al., 2021). Further, the distributional effectiveness of fiscal policy remains highly questionable in most countries. Indeed, Odusola (2017) shows that the difference between the gross Gini (before taxes and transfers) and the net Gini (after taxes and transfers) has declined in most countries, which implies that the efficiency of tax-and-transfer systems has also decreased.

According to Bhorat et al. (2017b), there has been a general increase in social protection expenditure, but social protection coverage, quality and level of assistance remain critical issues. The expenditure increase is more pronounced in Southern African countries, is variable across countries and does not appear to be correlated with economic growth. Current social protection expenditure is highly related to the quality of democratic governance (as captured by the Mo Ibrahim Index) and to resource dependence (non-resource dependent countries spend more on average).[4]

[3]The Open Budget Index is issued from the Open Budget Survey, which measures budget transparency, participation, and oversight.

[4]The Ibrahim Index of African Governance (IIAG) score aggregates four categories: safety

The comprehensive review of social protection in Africa by the African Development Bank et al. (2011) has shown the positive impact of many specific transfer programs on poverty and inequality reduction, suggesting that social protection can be a key driver of inequality reduction. Bhorat et al. (2017b) look at the correlation between inequality reduction (measured by the difference between pre-transfer and post-transfer Gini coefficients) and various characteristics of social protection. They find no clear impact of public social spending on inequality, but a positive impact of both pro-poor coverage of social protection and transfer average amount on inequality reduction.

Regarding educational inequalities, the quality of education is still low, despite significant progress in primary schooling enrolment (Bhorat and Naidoo, 2017; Bold et al., 2017). In addition, except in some Southern and Northern African countries, progresses of secondary education have been slow, and important enrolment differentials by income groups persist. This fosters high wage premiums for a few skilled workers in some occupations, which fuels income inequality.

8.2 Data and Methodology

In this section, we present the data sources used to estimate income inequality in Africa and our methodology to combine them. Section 3.1 presents our data sources. Section 3.2 describes the method used to convert consumption inequality estimates into income inequality estimates. Section 3.3 explains how we correct for under-representation of top incomes in surveys. Section 3.4 outlines how we reconcile our results with national accounts.

8.2.1 Data Sources

8.2.1.1 Survey Data

Our primary data source consists in survey tabulations from the World Bank, which are made publicly available on the PovcalNet website.[5] These tabulations provide information on the distribution of consumption per capita. We use Generalized Pareto Interpolation (Blanchet, Fournier, and Piketty, 2021) to harmonize these tabulations and estimate the distribution of consumption by percentile.[6] We complete

and rule of law, participation and human rights, sustainable economic opportunity, and human development (Mo Ibrahim Foundation, 2014).

[5] `http://iresearch.worldbank.org/PovcalNet/povOnDemand.aspx`.

[6] The objective of this interpolation technique is to produce a "smooth" distribution starting from either tabulated income tax data or non-exhaustive individual data, as is typically available from

our database with eight surveys from Côte d'Ivoire, which have been used by Czajka (2017) for his study on the evolution of income inequality in the country since the mid-1980s.[7] Finally, we use additional surveys conducted in Ghana (1988, 1998), Guinea (1994), Madagascar (1993), and Uganda (1992), which were compiled by Cogneau et al. (2007) and are especially useful to model the relationship between consumption inequality and income inequality. We also exploit surveys available from Jenmana (2018) for Thailand (2001-2016) and from Chancel and Piketty (2017) for India (2005, 2011), to have a broader perspective on the joint distribution of income and consumption.

Figure 8.2 shows that there are large variations in data coverage across African countries. In Morocco, Nigeria and Madagascar, surveys have been more or less conducted on a regular basis since the early 1980s. In central African countries, by contrast, only one or two surveys are available, in general after 2000. Overall, if we pool together all surveys in our dataset and interpolate between years, we are able to cover about 60% of the continental population in the early 1990s, and 80-90% from 2000 onward.[8]

8.2.1.2 Tax Data

In contrast to developed countries, where tax data can be used to correct for the under-representation of top incomes in a number of countries (Alvaredo et al., 2018), publicly available tax tabulations are close to non-existent in Africa. We use South African tax tabulations covering the 2002-2014 period provided by Alvaredo and Atkinson (2022) and updates, as well as a similar tabulation covering the formal sector in 2014 Côte d'Ivoire available from Czajka (2017), to study to what extent accounting for the "missing rich" affects income inequality estimates. We also extend our analysis to other developing countries using Thai and Indian tax tabulations provided by Jenmana (2018) and Chancel and Piketty (2017). Given the lack of income tax data in most African countries, we make strong but transparent assumptions in order to correct survey data on the basis of comparable countries

survey tabulations. Compared to other methods of interpolation, Generalized Pareto Interpolation has been shown to guarantee the smoothness of the estimated distribution, particularly for the top of the distribution (Blanchet, Fournier, and Piketty, 2021).

[7] See also Cogneau, Houngbedji, and Mesplé-Somps (2016) and Cogneau, Czajka, and Houngbedji (2018).

[8] The collection of data on household living conditions, on which the estimation of our inequality indicators is based, is not carried out every year in all countries due to its high cost. The database resulting from combining available surveys does not therefore cover all years for a given country. We thus interpolate income distributions between two years to cover every year from 1990 to 2019, by linearly interpolating the average income of each percentile.

where both tax and survey data are available. As additional tax data becomes available, our series can be updated accordingly. In the meantime, given that top-end corrections have a comparable and sizable magnitude in most countries, we feel that it is more adequate to apply a simple and transparent correction method to countries with missing tax data than to make no correction at all.

8.2.1.3 National Accounts

We account for inequalities between African countries by using macroeconomic series available from the World Inequality Database[9], which cover the 1950-2017 period. These series were constructed by Blanchet and Chancel (2016) by combining various historical data sources. In line with the Distributional National Accounts methodology (Blanchet et al., 2021), which aims to provide income inequality estimates that are consistent with macroeconomic growth rates, we use these series to scale our country-level inequality estimates to the national income per adult at purchasing power parity.

8.2.2 From Survey Consumption to Survey Income

The first issue with available inequality statistics in Africa is that they rely almost exclusively on consumption. This makes systematic comparisons between developed and developing countries difficult, since inequality is most often measured in terms of pretax or posttax income in the former. From a theoretical perspective, income inequality is expected to be higher than consumption inequality, as (i) high-income earners tend to save more than poorer individuals (ii) income has a transient component that some households are able to smooth in order to maintain a stable level of consumption and (iii) income is often less accurately measured than consumption and measurement error can inflate inequality. The consumption-income gap is likely to be large at the bottom of the distribution, where the proportion of households incurring transient negative income shocks and with mismeasured incomes is generally higher. It is also likely to be important at the top of the distribution, since the very rich tend to save a large proportion of their current earnings, benefit from large transient positive income shocks such as capital gains, and underreport their income in surveys. Yet, very little is known on how income-consumption profiles vary across countries and across time.

Our primary objective is to make estimates of the distributions of consumption and income comparable. Accordingly, if we know to what extent consumption is higher or

[9] `http://wid.world`.

lower than income at all points of a given distribution, we can use this relationship to "transform" consumption distributions into income distributions (Blanchet, Chancel, and Gethin, 2022). In other words, our aim is to model income-consumption profiles $c_1(.)$ of the form:

$$c_1(p) = \frac{Q^I(p)}{Q^C(p)}$$

Where $Q^I(.)$ is the quantile function associated with a given distribution of income, $Q^C(.)$ is the quantile function associated with a given distribution of consumption, and $p \in [0, 1]$.

We start by estimating the empirical shape of $c_1(p)$ for countries and years for which we have reliable survey data on both survey pretax income and consumption. Following our definition of $c_1(p)$, computing income-consumption ratios is straightforward: it simply amounts to dividing the bracket average of each percentile of the pretax income distribution by its consumption counterpart. In order to make profiles comparable, we systematically normalize average pretax income or consumption to 1. Notice that since our aim is to use $c_1(p)$ as a multiplicative factor, the ratio of aggregate consumption to aggregate income is irrelevant: what matters is how $c_1(p)$ varies with p.

Figure 8.3 plots income-consumption profiles in Côte d'Ivoire, Ghana, Guinea, Madagascar, Uganda, Thailand and India for various years. In nearly all surveys, the relationship between income inequality and consumption inequality is distinctively S-shaped. Average income is in general substantially lower than average consumption for the poorer half of the population. The ratio of income to consumption then increases more or less linearly up to percentiles 80 and 90, before rising exponentially at the top of the distribution. This is consistent with the intuitive mechanisms outlined above: poorer individuals tend to smooth their consumption, while the very rich tend to save a significant proportion of their current earnings. As a result, consumption inequality is generally lower than income inequality.

In order to characterize more precisely consumption-income profiles across surveys, we formulate $c_1(.)$ parametrically by using a scaled logit function of the form:

$$c_1(p) = \alpha + \beta \log(\frac{p}{1-p}) \tag{8.1}$$

For $p \in (0, 1)$. α is a constant which determines the starting point of the curve. It is irrelevant to our imputation problem, since multiplying the quantile function by α only affects the overall mean of the distribution. β is our parameter of interest: it determines how the ratio of income to consumption increases with p and is therefore a proxy of the extent to which income inequality is higher than income inequality.

Appendix table H.1 reports the results of $\hat{\alpha}$ and $\hat{\beta}$ estimated by ordinary least squares, along with the corresponding adjusted R-squared. In nearly all cases, our scaled logistic function provides an excellent fit of income-consumption profiles, explaining over 90% of variations in the data. Our coefficient of interest $\hat{\beta}$ is always positive and varies little across surveys. Consumption series underestimate income inequality most in Thailand at the beginning of the 2000s ($\hat{\beta} = 0.16$), and least in Madagascar and Uganda at the beginning of the 1990s ($\hat{\beta} = 0.05$ in Madagascar and $\hat{\beta} = 0.06$ in Uganda). Beyond these two extremes, a majority of correction profiles range between 0.10 and 0.14.

Our objective is to provide a reasonable approximation of income inequality in Africa by transforming all available consumption distributions into pretax survey income distributions. To do so, we define three theoretical profiles reflecting the variability in $\hat{\beta}$ observed in the data, allowing us to derive "confidence intervals" for our income inequality estimates. For our benchmark scenario (scenario A henceforth), we use $\hat{\beta}_A = 0.12$; in scenario B, we correct distributions more moderately by imposing $\hat{\beta}_B = 0.10$; and we correct them more strongly in scenario C by using $\hat{\beta}_C = 0.14$. Figure 8.4 plots our three correction profiles (setting $\alpha = 0.85$ to make them easily comparable with observed profiles).

8.2.3 From Survey Income to Fiscal Income

The second correction we apply to our survey distributions consists in correcting the average income of top earners. We refer to these top-corrected distributions as "fiscal income" in what follows. It is well-known that the rich are under-represented in surveys, because of both sampling and misreporting issues (e.g., Blanchet, Chancel, and Gethin, 2022). In some cases, the representativeness of survey samples can be very questionable. In Côte d'Ivoire, for instance, surveys tend to underestimate specific groups when compared to population censuses. Among the poor, these include migrants from Burkina-Faso and Mali; among the rich, some surveys completely miss French expatriates and the Lebanese minority (Czajka, 2017). When some groups had a zero probability to be surveyed, no reweighting procedure will solve the problem (Ravallion, 2022). Many studies have attempted to correct for these biases by combining surveys with tax data, either in the form of tabulations or microdata. Tax data only cover a limited part of the population but provide better coverage of the very top of the distribution. While corrections based on tax data almost systematically yield higher inequality levels, little is known on the typical shape of these corrections and how this shape varies across countries.

Following the method used for consumption, our aim is to use existing data to define "plausible" profiles correcting income levels at the top of the distribution. In the African case, correcting for the under-representation of the rich in surveys is particularly challenging. To our knowledge, one of the only research papers combining surveys and tax data in an African country at the time of writing is Czajka (2017).[10] The paper exploits recently released tax tabulations from Côte d'Ivoire, and shows that the average pretax income of the top 1% could be underestimated by about 75% in the private sector. In other developing countries, the correction profiles of top pretax incomes obtained from matching surveys with tax data vary greatly across studies. In Brazil, Morgan (2017) finds that the average taxable income of the top 1% is 1.5 to 3 times higher than in surveys, with variations across years. Corresponding figures are found to be between 1.5 and 2.5 in Thailand (Jenmana, 2018) and as high as 3.5 in Lebanon (Assouad, 2017).

We look at variations in the underestimation of top incomes in Africa by bringing together surveys and tax tabulations from Côte d'Ivoire (Czajka, 2017) and South Africa (Alvaredo and Atkinson, 2022; Chatterjee, Czajka, and Gethin, 2023). For South Africa, we match the 2008, 2010 and 2012 surveys compiled in the Luxembourg Income Study (LIS) with the fiscal income series provided by Alvaredo and Atkinson (2022) and subsequent updates available from the World Inequality Database (Chatterjee, Czajka, and Gethin, 2023). We then use the method developed by Blanchet, Flores, and Morgan (2022) to combine surveys and tax data in order to get corrected pretax survey income distributions. The method essentially compares the distributions of survey pretax income and fiscal income, and finds a merging point where they cross. It then reweights survey observations so that the information on top incomes in the survey matches that observed in the tax data.

Exactly as in the case of consumption and income, our objective is to estimate "survey-fiscal" profiles $c_2(.)$ of the form:

$$c_2(p) = \frac{Q^F(p)}{Q^I(p)}$$

Where $Q^I(p)$ is the quantile function associated with the distribution of income observed in the survey, and $Q^F(p)$ is the quantile function of the distribution obtained

[10]See also Chatterjee, Czajka, and Gethin (2023) and Bassier and Woolard (2020) for preliminary evidence in the context of South Africa.

after correcting for the under-representation of top incomes. The South African profiles can be computed by dividing the average incomes observed in the corrected distributions by their corresponding values in the surveys. In Côte d'Ivoire, the ratio of fiscal income to survey income by percentile is obtained from Chancel and Czajka (2017).

Figure 8.5 plots survey-fiscal profiles in our two countries of interest. In Cote d'Ivoire, the ratio of corrected income to survey income is close to 1 before the 90^{th} percentile, and then increases exponentially. In South Africa, the correction starts much earlier (before the 80^{th} percentile), but rises more moderately. In both countries, surveys tend to largely underestimate top incomes, especially at the very top of the distribution. Correcting for this bias amounts to increasing the average of the top 1% by between 50% and 125%.

The correction profile of top incomes can be formally conceptualised as depending on two dimensions: the size of the group which is corrected, and the magnitude of the correction applied to top earners within this group. One way to formulate these two dimensions parametrically is to model survey-fiscal profiles by the quantile function of the Lomax (or Pareto Type II) distribution:

$$c_2(p) = \mu + \sigma(p^{1/\gamma} - 1)$$

For $p \in [0, 1]$. μ is a constant which determines the starting point of the curve; as in the case of consumption-income profiles, it is irrelevant to our problem. Since it makes sense to let $c_2(p)$ take the value 1 before a certain percentile p_0, one can set $\mu = 1 + \sigma$, so that $c_2(0) = 1$ and:

$$c_2(p) = 1 + \sigma p^{1/\gamma}$$

σ is the scale parameter. It controls the slope of the curve: the higher σ, the more top incomes are underestimated by surveys. γ is the shape parameter: as it decreases, the slope becomes more convex, so that a smaller fraction of top incomes is corrected.

While it is difficult to find regularities in the correction of top incomes given the paucity of comparable data across countries and across years, we believe that some correction is better than no correction at all, given what we know of countries with better data availability. In our benchmark scenario, we set $\sigma = 0.9$ and $\gamma = 0.05$. We then let σ vary from 0.6 to 1.2. As figure 8.6 shows, this approximately corresponds to rescaling incomes exponentially above the 80^{th} percentile (γ) and multiplying the average income of the top 1% by between 1.5 and 2 (σ). These bounds are in line with the different corrections observed in Côte d'Ivoire and South Africa. They are arguably sufficiently large to represent plausible variations in the correction of top incomes in Africa across countries and across time. If anything, this correction profile is likely to be a lower bound: in other developing countries such as Brazil, Lebanon or Thailand, it was not uncommon to find that the top 1% average was underestimated by a factor of 2 to 3 (Assouad, 2017; Jenmana, 2018; Morgan, 2017).

We illustrate the effect of the different adjustments presented in sections 3.2 and 3.3 for the case of Morocco. Figure 8.7 plots the top 10% share across years in Morocco adding up the corrections for conceptual discrepancies and underestimation of inequality at the top. Using the consumption distribution provided by PovcalNet, the highest decile received about 30% of total consumption, with no clear trend over the period. Moving from consumption to pretax income (section 3.2) increases this value to 35-40%, while correcting top incomes (section 3.3) further increases it to above 45% in our benchmark scenario. These results suggest that consumption-based measures from PovcalNet tend to underestimate the share of national income accruing to top 10% earners by as much as 40%.

8.2.4 From Fiscal Income to National Income

Under the assumption that our method for improving the measurement of income inequality is valid, the distribution we obtain corresponds to the distribution of pretax household income – that is, the sum of compensation of employees, mixed income and property income received by the household sector in the national accounts. To reach national income and obtain figures on individual incomes that are consistent with macroeconomic growth, we have to make assumptions on the distribution of unreported income components. These mainly include the taxes on production received by the general government and the retained earnings of corporations, which can represent a significant fraction of the national income in both developed and developing economies (Alvaredo et al., 2018).

In developed countries, and in some emerging economies, the levels of unreported income components can generally be observed from national accounts, and various methods can be used to impute these components indirectly on the basis of household surveys. Unfortunately, this is not the case for most African countries, where national accounts are still in their infancy. As a result, we do not have access to reliable data on unreported income. We choose to distribute the gap between surveys and the net national income proportionally to individual income.[11] We stress that this step is far from optimal, given the relatively low quality of national accounts in some countries (see for instance Anand and Segal (2015) and Assouad, Chancel, and Morgan (2018) on this matter, and more specifically Jerven (2013) in the context of Africa). This choice is nevertheless motivated by the fact that national accounts remain the best comparable macroeconomic estimates available at the international level. This step therefore has the advantage of making average incomes and growth rates more comparable across countries and over time while keeping the overall distribution of pretax incomes unchanged.[12]

We also stress that this assumption is conservative: in most existing distributional national accounts studies, the imputation of unreported income leads to higher inequality levels, mainly because retained earnings are concentrated at the top the distribution (e.g., Blanchet, Chancel, and Gethin, 2022; Chatterjee, Czajka, and Gethin, 2023; Jenmana, 2018; Morgan, 2017; Piketty, Saez, and Zucman, 2018). As better national accounts data, survey microdata, and tax data become available, our estimates can be updated to account for such discrepancies.

8.3 The Distribution of Income and Growth in Africa, 1990-2019

8.3.1 How Unequal is Africa?

8.3.1.1 Inequality in African Countries

Is Africa a low or high inequality continent? Although our estimates should be interpreted with care, they suggest income inequality is very high in most African countries, especially in international perspective. The income earned by the top 10% of the distribution ranges from 37% in Algeria to 67% in Botswana (Figure 8.8), while the bottom 40% is at most 14% in Algeria, and is about 4% in South Africa (Figure 8.9).

Significant regional differences appear across the African continent. Southern Africa

[11]Net national income is equal to GDP, minus consumption of fixed capital, plus net foreign income. For more details, see the distributional national accounts guidelines (Blanchet et al., 2021).

[12]Appendix table H.4 presents the gap between survey means and net national income per capita in each country, revealing that this gap remains relatively large in most countries, with significant variations. That being said, the ranking of countries in terms of economic development remains relatively similar across measures. Our estimates of levels and trends in inequality in Africa as a whole are also barely affected by the use of survey means instead of national accounts aggregates (see appendix figure H.1).

is by far the most unequal region, with the top 10% share exceeding 65% in South Africa and Botswana. Inequality is slightly lower in Central Africa, but remains very high by international standards: for instance, in Congo in 2011, 56% of national income accrued to the 10% income earners, while the bottom 40% income share was 7%. Eastern African countries appear less unequal, especially at the bottom of the income distribution: in Kenya in 2015, for instance, the top 10% received 48% of national income and the bottom 40% about 9%.

Income inequality tends to decrease as one moves towards the North and the West of the continent. In Sierra Leone in 2011, the top 10% owned 42% of national income, and the bottom 40% owned 12%; its neighbors display comparable income shares. The lowest inequality levels can be found in Northern Africa; Algeria appears as the least unequal country in Africa, as in 2011 37% of national income was captured by the top 10% of the distribution, while the bottom 40% received 14%.[13]

8.3.1.2 Inequality in Africa as a Whole

Africa stands out as one of the continents with the highest levels of regional income inequality. According to our estimates, the top 10% of Africans captured 54% of national incomes in 2019, while the bottom 50% received only 9% (Figure 8.10). From an international perspective, the top 10% income share is 34% in Europe (550m individuals), 41% in China (1.4bn individuals), 47% in the United States (330m individuals), 55% in Brazil and the rest of Latin America (210m individuals), 56% in India (1.3bn individuals), and 61% in the Middle East (420m individuals). A particularly striking characteristic of the pan-African distribution is the extent of the gap between the top 10% and the bottom 50% income shares. Average incomes of the top 10% are about 30 times higher than those of the bottom 50%, well above the value found in other extreme inequality regions (the ratio is around 20 in the Middle East, India, or Brazil: see Figure 8.1). This finding reveals the dual and polarized nature of the pan-African income distribution, with extremely low incomes at the bottom and relatively high incomes at the top. As shown in Figure 8.10, overall income inequality in Africa seems to have remained very stable since the 1990s. The top 10% income share decreased from 55% to 54%, while the bottom 50% share increased from 8% to 9%.

[13]Regarding Algeria, whose inequality level appears very low by regional standards, the lack of transparency and the absence of recent data (the last available survey dates back to 2011) make it difficult to properly evaluate the reliability of inequality estimates. Going further back in time, inequality seems to have decreased since the 1990s. However, at this stage, we lack elements to assess this evolution.

Is inequality on the African continent mostly due to inequality within African countries or to cross-country differences in average national incomes? Figure 8.11 decomposes overall African inequality into its between-country and within-country components by plotting two counterfactual scenarios: one in which countries would have the same average national income, and one in which individuals within each country would have the same income. Inequality within countries stands out as explaining the bulk of pan-African income inequality. If there was no inequality between countries, keeping current within-country inequality levels constant, the top 10% income share in Africa would be 48%, only slightly lower than its actual value (54%). Conversely, if all individuals had the same income within each country, keeping national average income differences constant, the top 10% income share would drop to only 24%. A Theil decomposition of African inequality levels shows that 25% of African inequality can be attributed to the between-country component and as much as 75% to the within-country component.

The slight decline in overall African inequality since the 1990s has been mostly due to the dynamics of between-country inequality. This reduction was caused by several phenomena. Since the 1990s, several countries located at the middle of the African distribution in terms of national income per capita, such as Nigeria, Morocco, Ghana, Angola, Tunisia, or Namibia have seen their average income increase significantly. On the other hand, the average income of Africa's richest countries (Algeria, South Africa, or Libya for example) stagnated in the 1990s, and increased only moderately in the 2000s. Meanwhile, the poorest countries did not experience any significant increase in average income. This explains why the top 10% between-country income share decreased more than the bottom 50% increased.

The dynamics of between- and within-country inequality in Africa contrast with those observed at the global level, in Europe, or in Asia. At the global level, we observe a significant reduction of between-country inequality, which has been partially or entirely offset by a rise in within-country inequality (see Chancel et al., 2022b).[14] In Europe, contrary to Africa, most of the evolution in pan-European income inequality stems from within-country dynamics. Turning to Asia, the huge rise of inequality recorded in China and India (which amount to about 60% of the regional population) over the past four decades meant that a significant share of the rise of pan-Asian income inequality is explained by within-country changes. That being said, the African exception could also reflect the quantity of noise that plagues survey measurements and blur the evolution of within-country inequality.

[14]On the evolution of global income inequality in recent decades, see also Anand and Segal (2017) and Lakner and Milanovic (2016).

8.3.2 Accounting for Differences in Inequality Patterns across Africa

Why are inequality levels in Africa so high? This question is particularly challenging to address because of strong data limitations, as well as of the specificity and diversity of Africa's economic and political structures, shaped by both colonial heritage and its recent history. In the following two subsections, our objective is not to provide a definitive explanation for the diversity of inequality levels found in Africa, but merely explore the role of historical factors on the one hand, and of government redistribution policies on the other.

8.3.2.1 Historical Determinants: Settler Colonialism, Socialism and Islam

Contemporary African inequality levels could reflect both the situation at the moment of countries' independence, and the political economy and institutions that followed (Cornia, 2019; Heldring and Robinson, 2018). In this section, we examine to what extent regional patterns of income inequality may be explained by long-term history. The evidence we present is only suggestive, and its interpretation can only be speculative.

First, our analysis suggests that high levels of inequality are typically found in countries that experienced European settler colonization, a type of colonization that resulted in high land and capital concentrations and in many cases restricted the access of natives to education and good jobs (Alvaredo, Cogneau, and Piketty, 2021). The long-run impact of settler colonialism might account for the high inequality levels in Southern Africa. Second, we uncover a large and robust negative correlation between income inequality and the spread of Islam. This negative correlation, whose interpretation will require further research, might account for the lower levels of inequality observed in the Western and North-Eastern regions of Sub-Saharan Africa, as well as in North Africa to some extent. Last, we show that other long-term factors, such as geography, precolonial history, and colonizers' identity do not correlate with country-level income inequality.

Among the countries with the highest income share of the richest 10%, South Africa and Namibia are still today inhabited by a significant number of people of European descent. The direct descendants of British, Dutch, French, and German settlers now make up 8% of the population in South Africa, and around 6% in Namibia. In 2019, the top 10% income share was estimated at 66% in South Africa and 64% in Namibia. As is well-known, apartheid in South Africa was only terminated in 1994 (and in

1990 in Namibia). In 1987, white South Africans represented 90.5% of top 5% income earners, while Coloured, Asians, and Blacks represented 4, 3, and 2.5%, respectively (Alvaredo and Atkinson, 2022). In the same region, European descendants still represent around 2% of population in Eswatini, and 1.2% in Botswana (Putterman and Weil, 2010) and these two countries also display rather high top 10% shares (respectively, 59.5 and 58.9%).[15] At the world level, Putterman and Weil (2010) show that in countries where people of European descent are mixed with natives and with people of other origins, income inequality is higher, while descendants of Europeans tend to lie at the upper end of the income distribution. Apart from South Africa and Namibia, the most salient cases are found in Latin America and the Caribbean.

Yet, Easterly and Levine (2016) have also argued that the consequences of settler colonialism extend after the departure of Europeans. Settler colonialism had a long-term impact on institutions, human and physical capital accumulation, and finally on GDP per capita. It could also have left a persistent imprint on inequality. Outside of South Africa and Namibia, although significant numbers of European expatriates can be found in some countries, Africa-natives of European descent are now very small minorities. Nonetheless, many other countries received significant numbers of European settlers in the past. We make use of the data set built by Easterly and Levine (2016) to identify countries that experienced settlement colonialism between 1870 and 1970. Over this period of a hundred years, we categorized countries as former settlement colonies if the share of Europeans in the total population went above 2.5% at some point in time.[16] This threshold of 2.5% is not too arbitrary. Only a few countries exhibit shares between 1% and 2.5%: Egypt (1.4%), Gabon (1.3%), Senegal (1.2%), where Europeans were mostly administrators and traders, and the islands of Cabo Verde (2%) and São Tomé and Principe (1.9%), which were uninhabited before Europeans arrived; for 30 countries, the maximal European share is just below 0.25%. With the 2.5% threshold, we are left with 12 settler countries out of 54. Over 1870-1970, the maximum share of Europeans reached 21% in South

[15]In the colonial era, Eswatini (former Swaziland) and Botswana (former Bechuanaland) were largely administered by white South Africans, like Namibia (former South West Africa) between 1920 and 1990; on Botswana, see Bolt and Hillbom (2016).

[16]We complemented the Easterly and Levine (2016) data set for Libya and Mozambique. We also corrected their data for Djibouti, Kenya, and Malawi, which contained obvious overestimates. Easterly and Levine (2016) used the share of Europeans in the population fifty years before independence. Yet, their data are patchy, and for many countries the share they retain actually corresponds to a later date (1956 for Tunisia and Morocco, versus 1860-1911 for Algeria). We think their criterion fits better for the early colonialism in Latin America and the Caribbean (that lasted a longer time) than for the late colonialism of the 19th century, in Africa or Asia. In many countries again, such as Morocco and Zambia, significant inflows of settlers came in during the Interwar period, or even after 1945. We capture a kind of settlement colonialism that was more short-lived than the one they measure.

Africa and 14% in Namibia.[17] In the Northern neighborhood of South Africa, the two former Rhodesias, now Zimbabwe and Zambia, belong to our group of settler colonies. Their poor neighbor Malawi (former Nyasaland), with which they formed a Federation between 1953 and 1963, also received Scottish settlers, yet the figure of 2.7% from Easterly and Levine (2016) for 1956, taken from Curtin et al. (1995), is overestimated. Further North, the highlands of Kenya received British settlers who captured a significant fraction of arable land, yet their share in population never went above 1% (Bigsten, 1986).[18] In Zimbabwe, white power remained until 1979 and settlers started to migrate out right after, then at an accelerated pace in the 21st century. White Zimbabweans now constitute a very tiny group, estimated at less than 0.2% of population, much like white Zambians. Portuguese Angola and Mozambique were also settlement colonies, until the independence wars that ended in 1975, after which most of the settlers quickly left. At the other end of the continent, the three French colonies of North Africa were also exposed to large European settlement, first Algeria, then Tunisia, and Morocco (Cogneau, Dupraz, and Mesplé-Somps, 2021). In Algeria and Tunisia in the late colonial era (1950s), top income inequality was as high as in South Africa, as income tax tabulations reveal (Alvaredo, Cogneau, and Piketty, 2021).[19] Again, most settlers had left at the end of the 1960s, not long after the countries' independence. The neighboring Italian colony of Libya also received a large number of settlers. Italians left in two waves, first in the late 1940s after independence, and then in the 1970s after Muammar Gaddafi took power. According to our criterion, we also categorize the island of Mauritius as former settler colony, where French and British settlers owned plantations, even after the abolition of slavery.

A first direct consequence of settler colonialism is the unequal distribution of land for agriculture (Frankema, 2010). We gathered data on Gini coefficients of the land size distribution from various sources. Only 33 countries have non-missing data for some year after independence.[20] In this subsample, the nine former settler colonies come out with an average land Gini of 0.65 that is higher by almost 0.15 (p-value=0.007) than the average of non-settler countries (0.49). If we exclude South Africa, i.e., one of the two countries where descendants of European settlers still weight more than

[17] According to our criterion, Eswatini is a former settler colony, but Botswana is not. The patchy nature of the data prevents us from exploiting a continuous measure of the *intensity* of settlement.

[18] If we disregard the demographic threshold of 2.5%, and classify Kenya and Malawi as settler colonies, our results are very little changed.

[19] The top 1% share was even higher in Zambia and Zimbabwe.

[20] 2 in the 1960s, 8 in the 1970s, 5 in the 1980s, 12 in the 1990s, and 6 in the 2000s. We combine data assembled by the NGO Grain, in particular from FAO reports, Frankema (2010), and Vollrath (2007) from agricultural censuses.

2.5% of population, the difference is maintained at 0.13 (p =0.016).[21] Settler colonies of North Africa (Algeria, Libya, Morocco, Tunisia) make no exception in this regard, with an average Gini of 0.68 (all data are from 1987 to 2001). In contrast with many Asian countries, land reforms in Africa have been limited, even in socialist Algeria and Tunisia (Bessaoud, 2007).

Land inequality is not the only channel through which the legacy of settler colonialism can impact present income inequality. Inequality in other assets (capital, education), the dualistic or segmented structure of the labor market, as well as economic or political institutions are other potential channels. When contrasting the 12 former settler colonies with other countries, we find a significant difference in the top 10% share, of 5.5 percentage points (p-value=0.017); the bottom 50% share is lower by 2.1 p.p. (p=0.054). However, as pointed out before, North Africa is the least unequal region, which drags this correlation down. When restricting the analysis to Sub-Saharan Africa (49 countries out of 54), the differences between former settler colonies and other countries doubles, reaching 11 p.p. for the top 10% ($p<0.001$), and -4.9 p.p. for the bottom 50% ($p<0.001$). In North Africa, independent Algeria, Tunisia, Libya, and Egypt all embraced, at least for some time, some form of socialism that maintained a state-controlled economy and relatively high levels of public spending. Instead, the Kingdom of Morocco remained under a monarchical and conservative government, which could partly explain its relatively higher level of inequality today. The presence of strong states that adopted a socialist orientation at some point might be one explanation for North Africa's exceptionalism.

When looking at the maps of Figures 8.8 and 8.9, another historical correlate of inequality is revealed, that is the extension of Islam. Most of the countries in which the majority of population is Muslim appear in green or light yellow colors: in North Africa; on the Western coast from Mauritania to Guinea; and in the Sahel strip, from Mali to Sudan. The only exceptions are Morocco (99% Muslim) and Chad (56%), yet their estimated top 10% share is just above the 48% upper threshold of light yellow color, at 49%. Indeed, the negative correlation of the top 10% share with an estimate of the proportion of Muslim population in 2010 (Kettani, 2010) stands at -0.57 ($p<0.001$); the positive correlation with the bottom 50% is 0.62 ($p<0.001$). When setting apart North Africa, where the share of Muslims is above 94% in all five countries, the correlations are only slightly lower (respectively -0.50 and +0.53, $p<0.001$).

[21] Namibia is missing, yet a recent World Bank Report notes that "70% of Namibia's 39.7 million hectares of commercial farmland is still owned by Namibians of European descent" (World Bank, 2022, pp. 4 and 60-66).

The interpretation of these correlations is more difficult than for settler colonialism. Note first that Islam may have interacted with colonialism. European colonizers and missions tended to favor non-Muslim areas (e.g., Cogneau and Moradi, 2014). In contrast, Islamized areas experienced lower investments by missionaries in education and health, and lower penetration by the colonial state in terms of administration and social services (Bauer, Platas, and Weinstein, 2022).[22] Muslim elites were less often involved in colonial rule than evangelized elites, and, if in power after independence, they could have more strongly broken with the unequal legacy of colonialism. Moreover, Islamic thought shows a tradition of egalitarianism that may influence state policies as well as individual behavior (Marlow, 1997). Yet, one can first note that, outside of Africa, in majority Muslim Middle East, income inequality is large (Alvaredo, Assouad, and Piketty, 2019). Second, if Islamic charity translates into large private transfers to the poor, then the income-consumption profile might be steeper than what we have assumed, so that income inequality would be underestimated; or else, egalitarianism and the culture of charity may lead rich individuals to under-report their income more often, out of shame. Nonetheless, the income-consumption profile of majority Muslim Guinea is even flatter than neighboring Côte d'Ivoire (see Figure 8.3). Third, in Sub-Saharan Muslim Africa, households are larger, so that part of intra-household inequality is missed by standard surveys, leading to a significant underestimation of total inequality (De Vreyer and Lambert, 2021). While it is too early to conclude, the negative correlation between Islam and (measured) income inequality certainly deserves further research.

Finally, we ask whether settler colonialism and the spread of Islam are robust correlates of income inequality, when compared with other potential historical correlates. As mentioned above, we measure settler colonialism with a binary variable that is equal to one if the European population represented more than 2.5% of the total population, at some point in time between 1870 and 1970. This is true for 12 African countries out of 54; four are in North Africa, and seven belong to a large Southern cone (Angola, Eswatini, Mozambique, Namibia, South Africa, Zambia, and Zimbabwe), the island of Mauritius being the last one. We measure the spread of Islam with the share of Muslims in the total population circa 2010. 18 African countries have a majority Muslim population (more than 50%).

[22]Indeed, Sub-Saharan Muslim countries feature lower levels of education today. Yet, we found no correlation between income shares and mean years of schooling in 2015. Therefore, despite its positive correlation with settler colonialism and its negative correlation with the share of Muslims, average education does not explain the correlations of the two historical variables with inequality. Data for mean years of schooling in 2015 are from the Human Development Report 2021/2022 (UNDP, 2022).

We first restrict the analysis to Sub-Saharan Africa, and show simple OLS regressions of the top 10% and bottom 50% income shares on these two variables (see Tables 8.1 and 8.2). Both coefficients are very significant, both economically and statistically speaking. Having been exposed to settler colonialism is associated with a 8.9 percentage points higher top 10% share, and with a 3.8 p.p. lower bottom 50% share. Going from 0 to 100% of Muslims lowers the top 10% share by 6.6 p.p., and adds 3.4 p.p. to the bottom 50% (Table 8.1, column A). The two variables alone explain more than 40% of the variance of income shares across countries (adjusted R-squared). They are quite correlated with the regional patterns that are visible in Figures 8.8 and 8.9, as settler colonialism mainly affected countries in Southern Africa, and Islam is more widespread in Western Africa. If we break down Sub-Saharan Africa into four regions (North-Eastern, Western, Eastern and Southern), regional differences also explain 35 to 42% of the variance in income shares (Table 8.1, column B). Small islands, which are specific in that they were uninhabited before slavery and colonization, display significantly lower levels of inequality. Yet, a horse race between our two historical variables and regional dummies shows that the former are not subsumed under the latter. Both historical variables remain very significant, both coefficients are just slightly reduced by around 15% (Table 8.1, column C). Furthermore, they are almost able to explain all the contrast between the most equal (North-Eastern and Western) and the least equal regions (Southern). If we except the small islands' specificity, regional dummies turn statistically insignificant as a whole.

In Table 8.2, we then confront our two historical variables with other potential long-term correlates of present-day income inequality. We consider three groups of alternative factors (see Table 8.1 footnotes for a precise description of the variables). The first one is geography. Although rainfall, temperature and distance to the sea should not directly impact income inequality, they could for example condition agricultural productivity and the potential earnings of farmers. A second group relates to precolonial history (inside present-day borders, which were delineated by colonizers): the slave trade, precolonial polities, and ethnic fractionalization. The three dimensions are potentially intertwined, as the slave trade may have affected the political structures that were observed by anthropologists at the end of the 19th century or at the beginning of the 20th, and ethnic fractionalization as well; ethnicities can also be characterized by diverse political cultures. By enriching local traders, the slave trade might have had a long-term unequalizing impact; conversely, by increasing labor scarcity, after abolition it might have led to higher earnings for unskilled free labor, hence reducing inequality. Although the effect of centralized precolonial structures is perhaps ambiguous, hierarchical political structures, which we also distinguish, may be hypothesized to be more unequal. Ethnic fractionalization may generate vertical inequalities in some places, whereby politically dominant groups would be economically advantaged. Finally, a third set of historical factors is the national identity of the colonizer (Belgian, British, French, etc.). Colonizers' effects may go through different educational policies and local elite formation (Ricart-Hughuet, 2021). Past works have argued that these three groups of long-term factors could explain differences in GDP per capita, or the quality of institutions (e.g., Sachs and Warner, 1997; Nunn, 2008; Gennaioli and Rainer, 2007; Easterly and Levine, 1997; La Porta, Silanes, and Shleifer, 2008). Here, we ask whether they correlate with income inequality. It seems that they do not. None of the three groups of variables comes out with statistically significant coefficients, and none is able to explain a significant share of the variance among African countries. In this respect, settlement colonialism and the spread of Islam do a much better job than geography, precolonial history, or the identity of the colonizer.

In Appendix Tables H.5 and H.6, we run the same regressions on the whole sample of African countries, hence adding the five countries of North Africa. Both the European settlement variable and the Muslims share preserve their high significance, even if, as expected, the coefficient of the former is reduced. In this case, the two variables do not suffice to erase regional differences, in particular between Northern and Southern Africa. In the countries of North Africa that received a lot of French and Italian settlers (Algeria, Libya, Morocco, and Tunisia), the equalizing effect of Islam is not high enough to explain why inequality is low. To get there, we would need to allow Islam to be more inequality-reducing in former settler colonies. It is not impossible that the type of Arab socialism that was experimented in North Africa (with the exception of Morocco), as it combined with Islam as a state religion, was quite effective in mitigating inequalities and in cancelling out part of the unequal legacy of settler colonialism. More research is warranted in order to go beyond the mere speculation developed in this section.

8.3.2.2 Redistribution Policies and Inequality in Africa

Most African countries have still significant progress to make regarding government redistribution, from increasing the fiscal space to improving tax progressivity, implementing efficient social protection systems, and providing high-quality quality public services. These issues are all the more pressing as existing research suggests that improvements along these margins are key drivers of inequality reduction in Africa.

In terms of government revenue, Africa is lagging behind all developed and many developing world regions (Figure 8.12). A large group of countries in Middle, Western and Eastern Africa is characterized by low government revenue, below 20% of GDP. Only richer Northern and Southern African countries succeed in collecting more than 30% of GDP in taxes. For most African governements, low state capacity hinders their ability to reduce income inequality. In some countries, fiscal capacity has improved during the two last decades, in particular on the side of domestic taxation; yet in many countries, government revenue remains highly dependent on mineral resources and their volatile international prices (Cogneau et al., 2021).

The impact of progressive taxation on posttax income inequality is straightforward, but its role in shaping pretax income inequality is also real, through capital accumulation and wage bargaining (Piketty, Saez, and Stantcheva, 2014). In Africa, redistribution through taxation is limited. Personal top income tax rates are lower than in the developed world in most African countries (Figure 8.13b). For a quarter of the countries for which data is available, top personal income tax rates amount to 25% or less. For half of countries studied, top personal income tax rates lie between 30 and 40%. Only eight countries have top marginal tax rates higher than or equal to 40%, comparable to those observed in rich countries. According to Odusola (2017), more generally, African tax systems tend to be regressive.

Social protection and assistance coverage are still minimal. African Development Bank et al. (2011) provide a comprehensive review of social protection in Africa, demonstrating that it can have a significant impact on poverty and inequality. Nonetheless, only a fifth of countries where data is available, mostly located in the South and the North, provide social insurance, social safety nets, or unemployment benefits to more than 45% of their population (Figure 8.13c). This figure was 54% in Brazil in 2015, and 63% in China in 2013.

Public services can also strongly impact income inequality through their influence on education and health inequalities. This issue is particularly relevant in Africa, where despite a substantial rise in primary enrollment rates in the last decades, the quality of public education remains low (Bhorat and Naidoo, 2017; Bold et al., 2017). In most African countries, total government expenditure on education falls below 5% of GDP. This is particularly true in Central and Eastern Africa, but also in comparatively rich countries such as Egypt and Algeria (Figure 8.13d).

Given the relative scarcity of data, estimating the incidence of taxes and transfers on inequality in each African country would require methods and data collection efforts that go far beyond those exploited in this paper. That being said, recent fiscal incidence studies (e.g., Lustig, 2018) and historical data collection efforts (e.g., Bachas et al., 2022) have shed new light on the potential impact of taxes and transfers on inequality in developing countries. Drawing on these various data sources, Gethin (2023b) constructs a new database covering estimates of the distribution of taxes and transfers worldwide since 1980. Although results should be interpreted with care given their preliminary nature and previously mentioned data limitations, appendix figure H.2 suggests that taxes and transfers only have a minimal impact on the level and evolution of inequality in Africa. Moving from pretax income to posttax disposable income (pretax income, minus direct taxes, plus social assistance transfers) reduces the top 10% income share by only a couple of percentage points, while only marginally increasing that of the bottom 50%.

8.4 Conclusion

Existing data sources on economic inequality in Africa are scarce and raise many challenges. We have tried to respond to one of the main challenges, namely the strong underestimation of inequalities by consumption-based indicators. The resulting estimates, though far from perfect, are at least conceptually comparable with the rest of the world.

The pan-African income distribution built from these estimates appears to be particularly unequal compared to other world regions. Within-country inequality accounts for a large part of pan-African inequality, and indeed many African countries rank among the most unequal in the world. Southern African countries are the most unequal of the continent, while inequality tends to be lower towards the North and the West.

Historical and institutional determinants may account for part of the geographical patterns of African inequality. Settler colonialism seems to cast its long shadow on Southern Africa even after the demise of apartheid, even in countries where white settlers have left for long. In North Africa, postcolonial policies inspired by socialism may have contributed to mitigating this legacy. The egalitarian spirit of Islam is also a potential candidate for explaining the lower levels of inequality observed in Northern and Western Africa.

The evolution of inequality since 1990 is even harder to measure, because data reliability becomes even more questionable as we go back in time. There has been a very modest decrease in inequality in Africa as a whole, which is entirely accounted for by a slight decrease in between-country inequality. Within-country inequality shows no clear trend overall, due to a very wide variety of trajectories that cannot even be summed up in clear regional patterns. Understanding potential drivers of the evolution of inequality over time in Africa remains an open issue.

We stress that further research on the subject requires African countries to cooperate to produce more reliable, transparent, and harmonized distributional data, on pretax and posttax income inequality as well as on wealth distributions. Recent digitization and tax data sharing efforts in certain countries (Côte d'Ivoire, Senegal, Mali, or South Africa, for instance) are interesting examples that could be expanded to other parts of the continent.

Bibliography

African Development Bank et al. (2011). *Assessing progress in Africa toward the Millenium Development Goals: MDG Report 2011*.

Alesina, Alberto et al. (2003). "Fractionalization". In: *Journal of Economic Growth* 8 (2), pp. 115–194.

Alvaredo, Facundo, Lydia Assouad, and Thomas Piketty (2019). "Measuring Inequality in the Middle East 1990-2016: The World's Most Unequal Region?" In: *Review of Income and Wealth* 65 (4), pp. 685–711.

Alvaredo, Facundo and Anthony B. Atkinson (2022). "Top incomes in South Africa in the twentieth century". In: *Cliometrica* 16, pp. 477–546.

Alvaredo, Facundo, Denis Cogneau, and Thomas Piketty (2021). "Income Inequality under Colonial Rule. Evidence from French Algeria, Cameroon, Indochina and Tunisia and comparison with the British Empire 1920-1960". In: *Journal of Development Economics* 152.

Alvaredo, Facundo et al. (2018). *World Inequality Report 2018*. Harvard University Press.

Anand, Sudhir and Paul Segal (2015). "The global distribution of income". In: *Handbook of income distribution*. Vol. 2. Elsevier, pp. 937–979.

– (2017). "Who are the global top 1%?" In: *World Development* 95, pp. 111–126.

Assouad, Lydia (2017). "Rethinking the Lebanese economic miracle: The extreme concentration of income and wealth in Lebanon, 2005-2014". In: *WID.world Working Paper 2017/13*.

Assouad, Lydia, Lucas Chancel, and Marc Morgan (2018). "Extreme Inequality: Evidence from Brazil, India, the Middle East, and South Africa". In: *AEA Papers and Proceedings*. Vol. 108, pp. 119–23.

Bachas, Pierre et al. (2022). "Globalization and Factor Income Taxation". NBER Working Paper 29819.

Barrett, Christopher B, Michael R Carter, and Peter D Little (2006). “Understanding and reducing persistent poverty in Africa: Introduction to a special issue”. In: *The journal of development studies* 42.2, pp. 167–177.

Bassier, Ihsaan and Ingrid Woolard (2020). “Exclusive growth? Rapidly increasing top incomes amid low national growth in South Africa”. In: *WIDER Working Paper* 53.

Bauer, Vincent, Melina R Platas, and Jeremy M Weinstein (2022). “Legacies of Islamic Rule in Africa: Colonial Responses and Contemporary Development”. In: *World Development* 152, p. 105750.

Bessaoud, Omar (2007). “Les réformes agraires postoloniales au Maghreb : un processus inachevé”. In: *Revue d’histoire moderne et contemporaine* 63 (4-4bis), pp. 115–137.

Bhorat, Haroon and Karmen Naidoo (2017). “Drivers of Inequality in the Context of the Growth-Poverty-Inequality Nexus in Africa: an Overview of Key Issues”. In: *Income Inequality Trends in sub-Saharan Africa: Divergence, determinants and consequences*. United Nations Development Programme.

Bhorat, Haroon et al. (2017a). “Resource Dependence and Inequality in Africa: Impacts, consequences and potential solutions”. In: *Income Inequality Trends in sub-Saharan Africa: Divergence, determinants and consequences*. United Nations Development Programme.

Bhorat, Haroon et al. (2017b). “Social Protection and Inequality in Africa: Exploring the interactions”. In: *Income Inequality Trends in sub-Saharan Africa: Divergence, determinants and consequences*. United Nations Development Programme.

Bigsten, Arne (1986). “Welfare and economic growth in Kenya, 1914-76”. In: *World Development* 14 (9), pp. 1151–1160.

Blanchet, Thomas and Lucas Chancel (2016). “National Accounts Series Methodology”.

Blanchet, Thomas, Lucas Chancel, and Amory Gethin (2022). “Why Is Europe More Equal Than the United States?” In: *American Economic Journal: Applied Economics* 14.4, pp. 1–40.

Blanchet, Thomas, Ignacio Flores, and Marc Morgan (2022). “The weight of the rich: improving surveys using tax data”. In: *The Journal of Economic Inequality* 20.1, pp. 119–150.

Blanchet, Thomas, Juliette Fournier, and Thomas Piketty (Apr. 16, 2021). “Generalized Pareto Curves: Theory and Applications”. In: *Review of Income and Wealth*, roiw.12510.

Blanchet, Thomas et al. (2021). *Distributional National Accounts Guidelines: Methods and Concepts Used in the World Inequality Database.* Tech. rep.

Bold, Tessa et al. (2017). "Enrollment without Learning: Teacher Effort, Knowledge, and Skill in Primary Schools in Africa". In: *Journal of Economic Perspectives* 31 (4), pp. 185–204.

Bolt, Jutta and Ellen Hillbom (2016). "Long-term trends in economic inequality: lessons from colonial Botswana, 1921–74". In: *Economic History Review* 69 (4), pp. 1255–1284.

Bossuroy, Thomas and Denis Cogneau (2013). "Social Mobility in Five African Countries". In: *Review of Income and Wealth* 59(S1), pp. 84–110.

Bourguignon, François and Christian Morrisson (1998). "Inequality and development: the role of dualism". In: *Journal of Development Economics* 57 (2), pp. 233–257.

Chancel, Lucas and Léo Czajka (2017). "Estimating the regional distribution of income in sub-Saharan Africa". In: *WID.world Technical Note 2017/6.*

Chancel, Lucas and Thomas Piketty (2017). "Indian income inequality, 1922-2014: From British Raj to Billionaire Raj?" In: *WID.world Working Paper 2017/11.*

Chancel, Lucas et al. (2022b). *World Inequality Report 2022.* Belknap Press of Harvard University Press.

Chatterjee, Aroop, Léo Czajka, and Amory Gethin (2023). "Redistribution without Inclusion? Inequality in South Africa Since the End of Apartheid". Working Paper.

Christiaensen, Luc, Lionel Demery, and Jesper Kuhl (2011). "The (evolving) role of agriculture in poverty reduction–An empirical perspective". In: *Journal of Development Economics* 96 (2).

Cogneau, Denis (2007). *L'Afrique des inégalités. Où conduit l'histoire.* Editions Rue d'Ulm.

Cogneau, Denis, Léo Czajka, and Kenneth Houngbedji (2018). "The return of the triumphant elephant? Growth and income inequality in Cote d'Ivoire (1988-2015)". In: *Afrique Contemporaine* 263-264, pp. 221–225.

Cogneau, Denis, Yannick Dupraz, and Sandrine Mesplé-Somps (2021). "Fiscal Capacity and Dualism in Colonial States: The French Empire 1830-1962". In: *Journal of Economic History* 81 (2), pp. 441–480.

Cogneau, Denis, Kenneth Houngbedji, and Sandrine Mesplé-Somps (2016). "The Fall of the Elephant: two decades of poverty increase in Cote d'Ivoire (1988-2008)". In: *Growth and Poverty in Sub-Saharan Africa.* Ed. by C. Arndt, A. McKay, and F. Tarp. Oxford University Press.

Cogneau, Denis and Alexander Moradi (2014). “Borders that Divide: Education and Religion in Ghana and Togo since Colonial Times”. In: *Journal of Economic History* 74 (3), pp. 694–728.

Cogneau, Denis et al. (2007). *Inequalities and Equity in Africa*. Agence Française de Développement (AFD).

Cogneau, Denis et al. (2021). “Taxation in Africa from Colonial Times to Present. Evidence from former French colonies 1900-2018”. In: *Paris School of Economics WP 2021/62*.

Cornia, Giovanni Andrea (2017). “Inequality Levels, Trends and Determinants in sub-Saharan Africa: an Overview of Main Changes since the early 1990s”. In: *Income Inequality Trends in sub-Saharan Africa: Divergence, determinants and consequences*. United Nations Development Programme.

– (2019). “Sub-Saharan Africa’s Diverging Inequality Trends, 1991-2011: A Structuralist Interpretation”. In: *Revue d’économie du développement* 27(2), pp. 9–43.

Curtin, Philip et al. (1995). *African History: From Earliest Times to Independence*. Longman.

Czajka, Léo (2017). “Income Inequality in Côte d’Ivoire: 1985-2014”. In: *WID.world Working Paper 2017/8*.

De Vreyer, Philippe and Sylvie Lambert (2021). “Inequality, Poverty and the Intra-Household Allocation of Consumption in Senegal”. In: *World Bank Economic Review* 35 (2), pp. 414–435.

Easterly, Willam and Ross Levine (1997). “Africa’s Growth Tragedy: Policies and Ethnic Divisions”. In: *Quarterly Journal of Economics* 112 (4), pp. 1203–1250.

– (2016). “The European origins of economic development”. In: *Journal of Economic Growth* 21 (3), pp. 225–257.

Estudillo, Joanna P. and Keijiro Otsuka (2010). “Rural poverty and income dynamics in Southeast Asia”. In: *Handbook of Agricultural Economics*. Ed. by P. Pingali and R. Evenson. Elsevier, pp. 3435–3468.

Fay, Marianne and Charlotte Opal (2000). *Urbanization without growth: A not so uncommon phenomenon*. Vol. 2412. World Bank Publications.

Fosu, Augustin Kwasi (2009). “Inequality and the impact of growth on poverty: comparative evidence for sub-Saharan Africa”. In: *The Journal of Development Studies* 45.5, pp. 726–745.

– (2015). “Growth, inequality and poverty in Sub-Saharan Africa: recent progress in a global context”. In: *Oxford Development Studies* 43.1, pp. 44–59.

Frankema, Ewout (2010). “The colonial roots oof land inequality: geography, factor endowments, or institutions?” In: *Economic History Review* 63 (2), pp. 418–451.

Gennaioli, Nicola and Ilia Rainer (2007). “The modern impact of precolonial centralization in Africa”. In: *Journal of Economic Growth* 12 (3), pp. 185–234.

Gethin, Amory (2023b). “Revisiting Global Poverty Reduction: Public Goods and the World Distribution of Income, 1980-2022”. Working Paper.

Gollin, Douglas (2010). “Agricultural productivity and economic growth”. In: *Handbook of Agricultural Economics*. Ed. by P. Pingali and R. Evenson. Elsevier, pp. 3825–3866.

Gollin, Douglas, Rémi Jedwab, and Dietrich Vollrath (2016). “Urbanization with and without industrialization”. In: *Journal of Economic Growth* 21 (1), pp. 35–70.

Griffin, Keith, Azizur Rahman Khan, and Amy Ickowitz (2002). “Poverty and the Distribution of Land”. In: *Journal of Agrarian Change* 2 (3), pp. 279–330.

Heldring, Leander and James A. Robinson (2018). “Colonialism and Development in Africa”. In: *The Oxford Handbook of the Politics of Development*. Ed. by Carol Lancaster and Nicolas van de Walle. Oxford University Press, pp. 295–327.

Herault, Nicolas and James Thurlow (2009). “Agricultural distortions, poverty and inequality in South Africa”. In: *Agricultural Distortions Working Paper 104*.

Imai, Katsushi S. and Raghav Gaiha (2014). “Dynamic and long-term linkages among growth, inequality and poverty in developing countries”. In: *Discussion Paper Series DP2014-33. REIB Kobe University.*

Imai, Katsushi S., Raghav Gaiha, and Wenya Cheng (2016). “Does Agricultural Growth Reduce Poverty and Inequality in Developing Countries?” In: *Discussion Paper Series DP2015-23. REIB Kobe University.*

Isham, Jonathan et al. (2003). *The Varieties of Resource Experience: How Natural Resource Export Structures Affect the Political Economy of Economic Growth.* Middlebury College Working Paper Series 0308.

Jedwab, Remi and Dietrich Vollrath (2015). “Urbanization without growth in historical perspective”. In: *Explorations in Economic History* 58, pp. 1–21.

Jenmana, Thanasak (2018). “Income inequality in Thailand, 2001-2016: Democratisation and the Emergence of Class Conflicts”.

Jerven, Morten (2013). “Comparability of GDP estimates in S ub-S aharan A frica: The effect of Revisions in Sources and Methods Since Structural Adjustment”. In: *Review of Income and Wealth* 59, S16–S36.

Kettani, Houssain (2010). “Muslim Population in Africa 1950-2020”. In: *International Journal of Environmental Science and Development* 1 (2), pp. 136–142.

La Porta, Rafael, Florencio Lopez-de Silanes, and Andrei Shleifer (2008). “The Economic Consequences of Legal Origins”. In: *Journal of Economic Literature* 46 (2), pp. 285–332.

Lakner, Christoph and Branko Milanovic (2016). “Global Income Distribution: From the Fall of the Berlin Wall to the Great Recession”. In: *The World Bank Economic Review* 30.2, pp. 203–232.

Lustig, Nora, ed. (2018). *Commitment to Equity Handbook: Estimating the Impact of Fiscal Policy on Inequality and Poverty*. Brookings Institution Press.

Manji, Ambreena (2006). *The Politics of Land Reform in Africa: From Communal Tenure to Free Markets*. New York: Zed Books.

Marlow, Louise (1997). *Hierarchy and Egalitarianism in Islamic Thought*. NCambridge: Cambridge University Press.

McMillan, Margaret, Dani Rodrik, and Verduzco-Gallo Inigo (2014). “Globalization, Structural Change, and Productivity Growth, with an Update on Africa”. In: *World Development* 63, pp. 11–32.

Michalopoulos, Stellios and Elias Papaioannou (2013). “Pre-colonial Ethnic Institutions and Contemporary African Development”. In: *Econometrica* 81 (1), pp. 113–152.

Morgan, Marc (2017). “Extreme and Persistent Inequality: New Evidence for Brazil Combining National Accounts , Surveys and Fiscal Data, 2001-2015”.

Nunn, Nathan (2008). “The Long Term Effects of Africa’s Slave Trade”. In: *Quarterly Journal of Economics* 123 (1), pp. 139–176.

Odusola, Ayodele (2017). “Fiscal Policy, Redistribution, and Inequality in Africa”. In: *Income Inequality Trends in sub-Saharan Africa: Divergence, determinants and consequences*. United Nations Development Programme.

Papyrakis, Elissaios and Reyer Gerlagh (2004). “The resource curse hypothesis and its transmission channels”. In: *Journal of Comparative Economics* 32.1, pp. 181–193.

Piketty, Thomas, Emmanuel Saez, and Stefanie Stantcheva (Feb. 2014). “Optimal Taxation of Top Labor Incomes: A Tale of Three Elasticities”. In: *American Economic Journal: Economic Policy* 6.1. Dataset, pp. 230–271.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman (2018). “Distributional National Accounts: Methods and Estimates for the United States”. In: *Quarterly Journal of Economics* 133.2, pp. 553–609.

Pingali, Prabhu (2010). “Agriculture renaissance: making ”agriculture for development” work in the 21st century”. In: *Handbook of Agricultural Economics*. Ed. by P. Pingali and R. Evenson. Elsevier, pp. 3867–3894.

Putterman, Louis and David N. Weil (2010). “Post-1500 Population Flows and The Long-Run Determinants of Economic Growth and Inequality”. In: *Quarterly Journal of Economics* 125 (4), pp. 1627–1682.

Ravallion, Martin (2022). “Missing Top Income Recipients”. In: *Journal of Economic Inequality* 20, pp. 205–222.

Ravallion, Martin, Shaohua Chen, and Prem Sangraula (2007). “New evidence on the urbanization of global poverty”. In: *Population and development review* 33.4, pp. 667–701.

Ricart-Hughuet, Joan (2021). “Colonial Education, Political Elites, and Regional Political Inequality in Africa”. In: *Comparative Political Studies* 54 (14).

Roessler, Philip et al. (2022). “The Cash Crop Revolution, Colonialism and Economic Reorganization in Africa”. In: *Working Paper ETH-Zürich*.

Sachs, Jeffrey D. and Andrew Warner (1997). “Sources of Slow Growth in African Economies”. In: *Journal of African Economies* 6 (3), pp. 335–376.

Thorbecke, Erik (2013). “The interrelationship linking growth, inequality and poverty in sub-Saharan Africa”. In: *Journal of African economies* 22.suppl_1, pp. i15–i48.

UNDP (2022). *Human Development Report 2021-2022*. UNDP.

Vollrath, Dietrich (2007). “Land distribution and international agricultural productivity”. In: *American Journal of Agricultural Economics* 89 (1), pp. 202–216.

Walle, Nicolas van de (2009). “The Institutional Origins of Inequality in Sub-Saharan Africa”. In: *Annual Review of Political Science* 12 (1), pp. 307–327.

World Bank (2015). *World development indicators 2015*. The World Bank.

– (2022b). *Inequality in Southern Africa: An Assessment of the Southern African Customs Union*. World Bank.

Figure 8.1: Inequality Levels Across World Regions, 2019

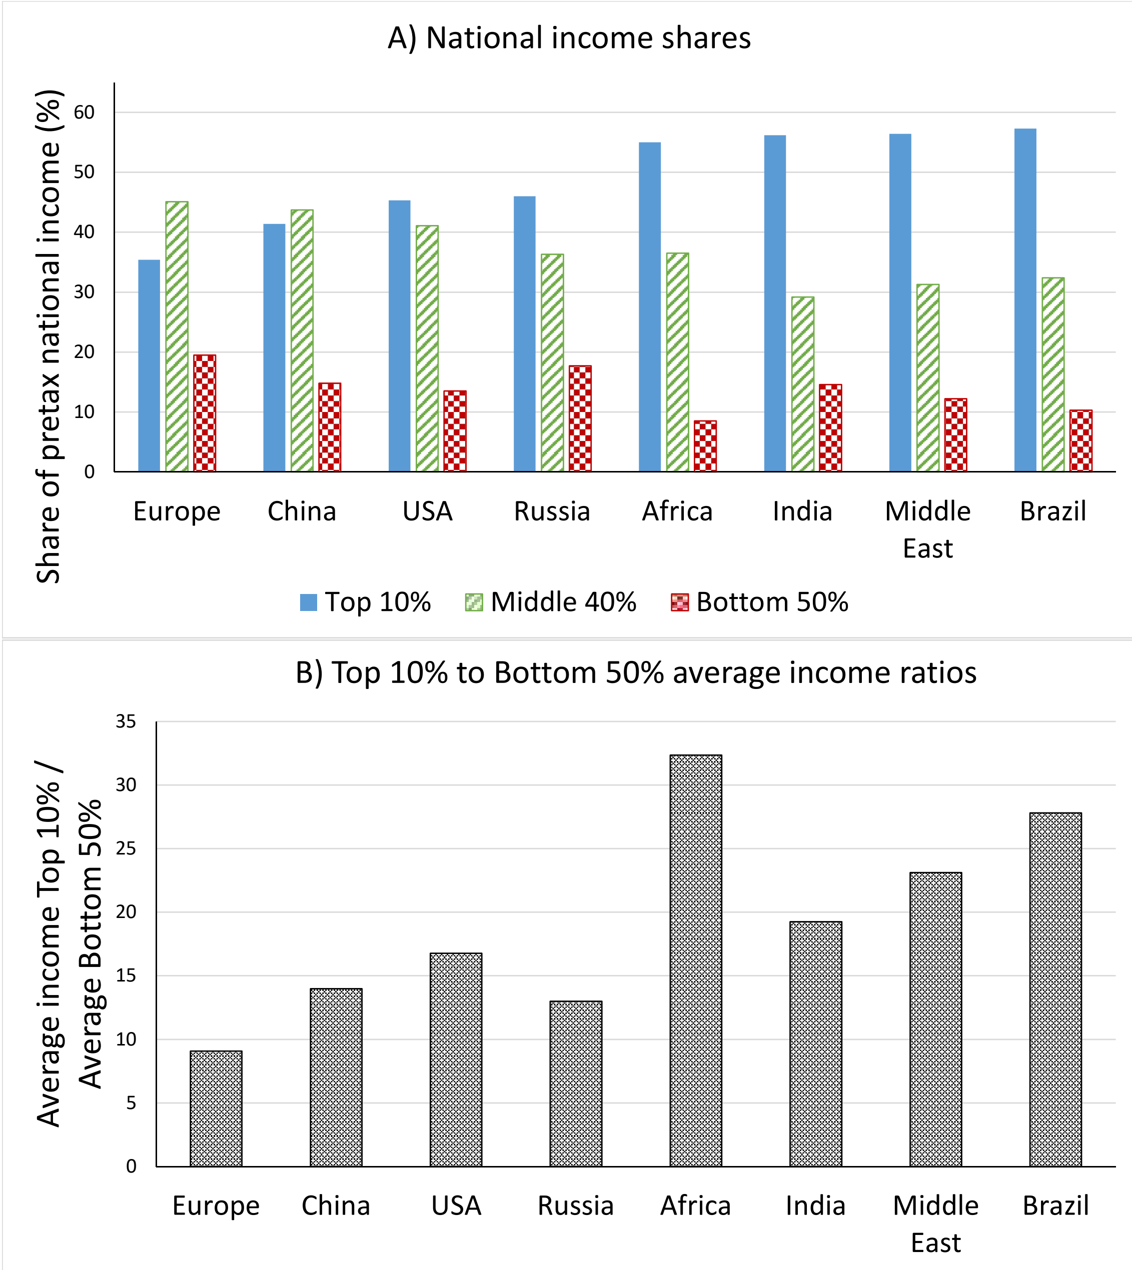

Source: Authors' computations based on WID.world (2021) and own estimates. Distribution of pretax income per adult.

Figure 8.2: Coverage of Survey Data Sources
First Year of Available Household Survey Data by Country

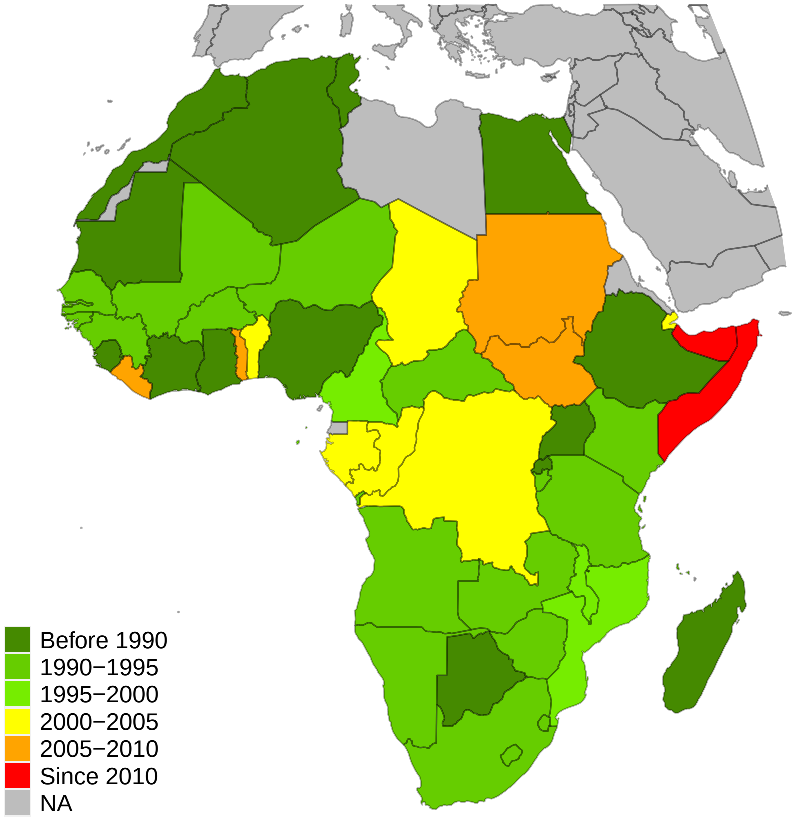

Notes. Authors' computations using available survey data from PovcalNet.

Figure 8.3: Empirical Consumption-Income Profiles in Eight Countries

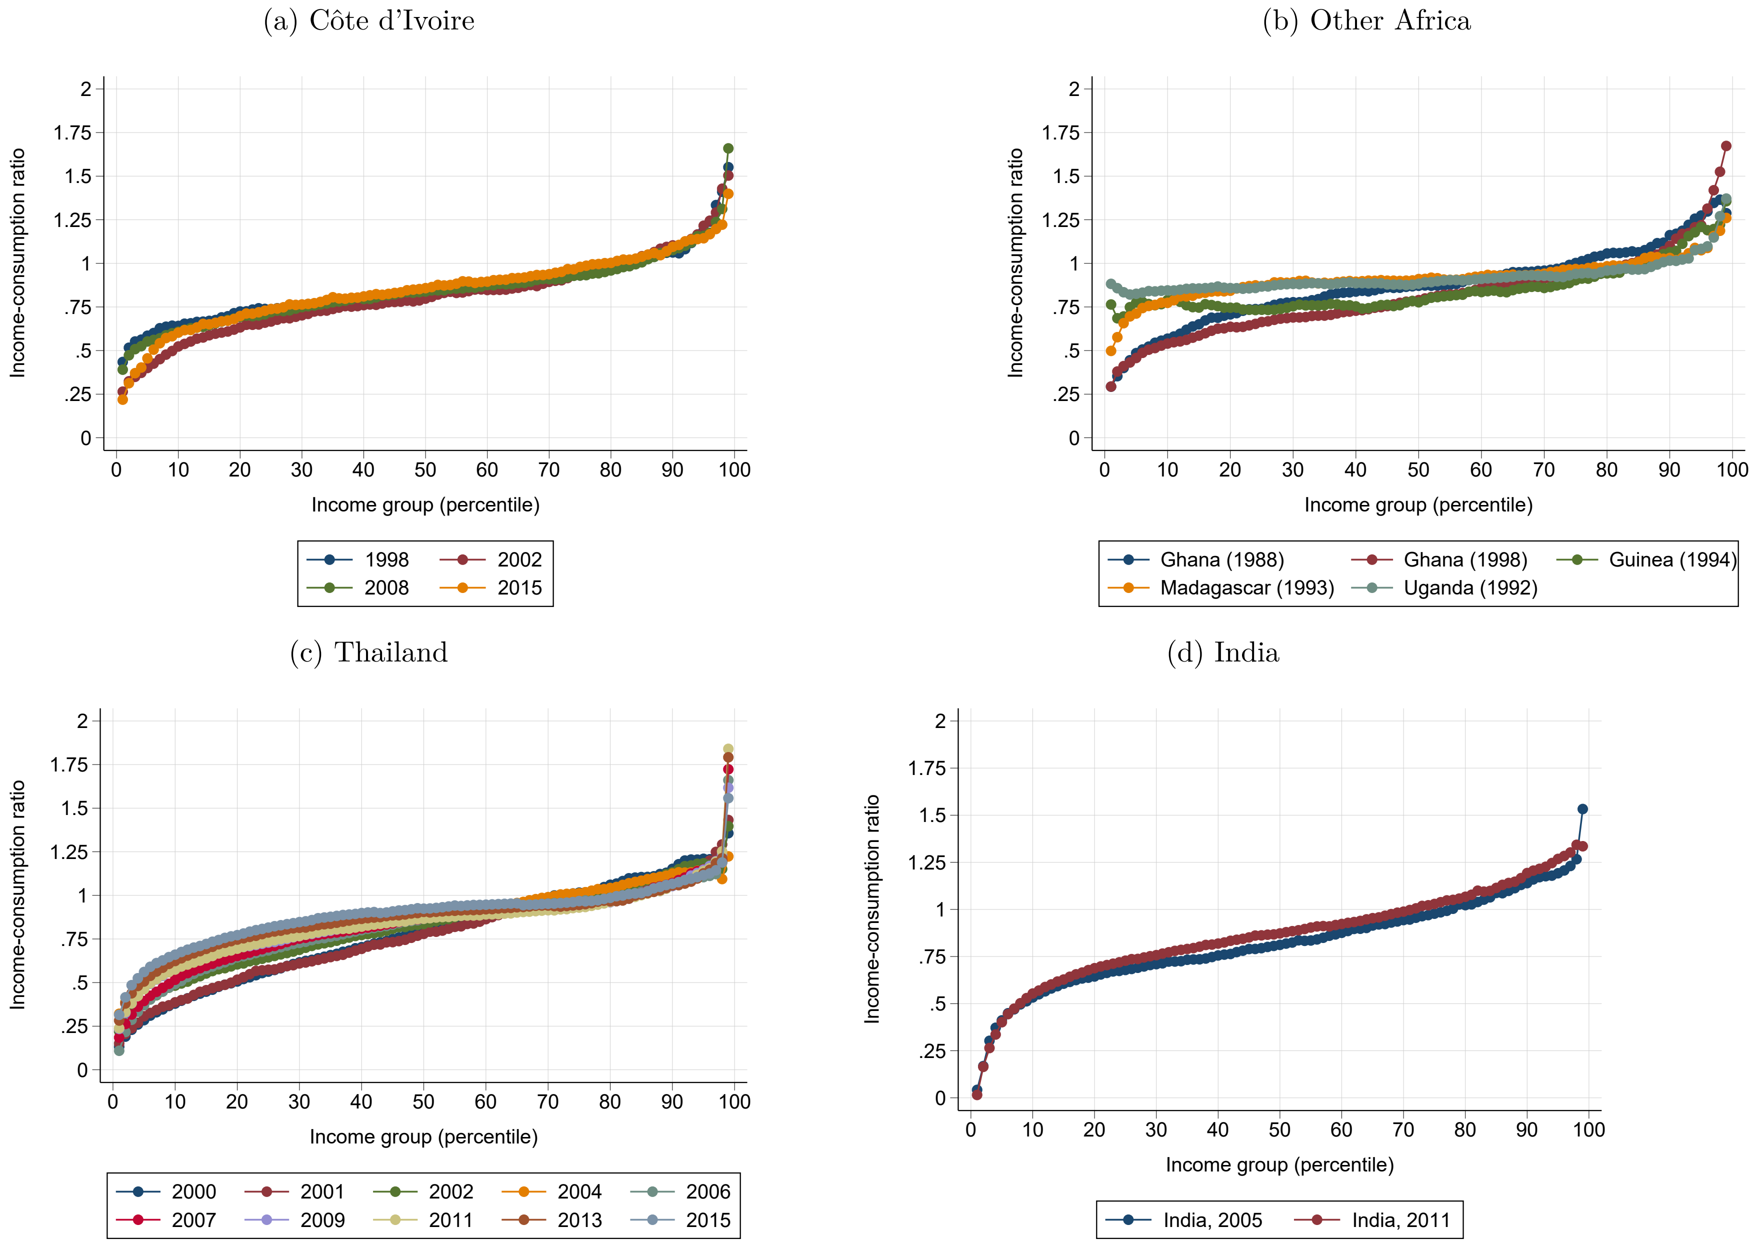

Notes. Authors' computations using survey data. The figure shows the ratio of average income to average consumption by percentile in each survey.

Figure 8.4: Theoretical Income-Consumption Profiles

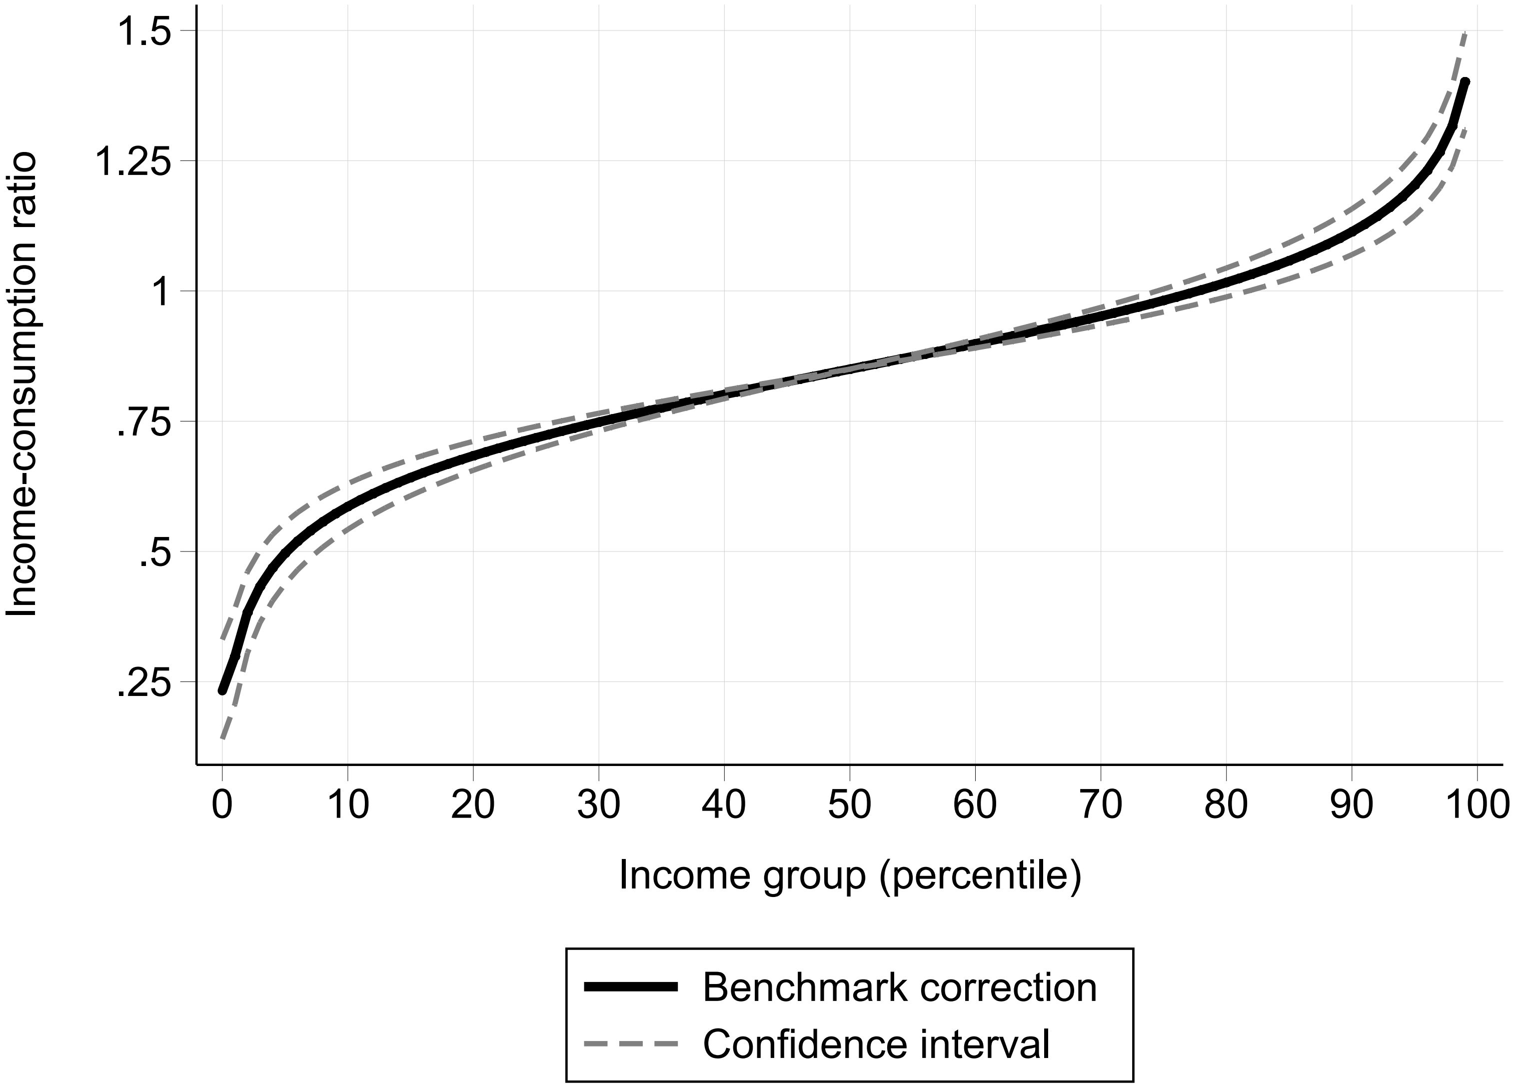

Notes. Authors' elaboration. The figure represents the three income-consumption profiles used to transform consumption distributions into income distributions. These profiles correspond to logistic functions of the form $Q_i(p) = \alpha + \beta_i \log \frac{p}{1-p}$ for $i \in A, B, C$. We set $\alpha = 0.85$ and $\beta_A = 0.12, \beta_B = 0.10, \beta_C = 0.14$.

Figure 8.5: Empirical Survey-Fiscal Profiles in Côte d'Ivoire and South Africa

(a) Côte d'Ivoire (2014) (b) South Africa (2008-2012)

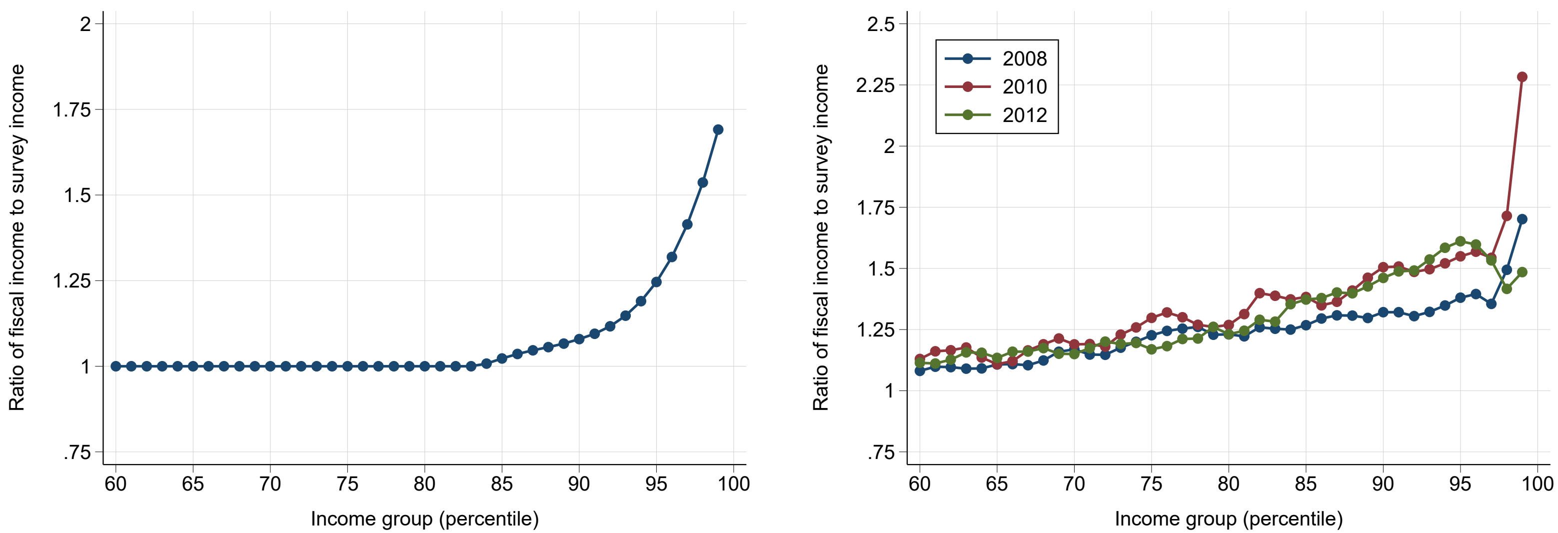

Notes. Authors' computations combining survey and tax data. The figure represents the ratio of survey income to taxable income by percentile in each country.

Figure 8.6: Theoretical Survey-Fiscal Profile

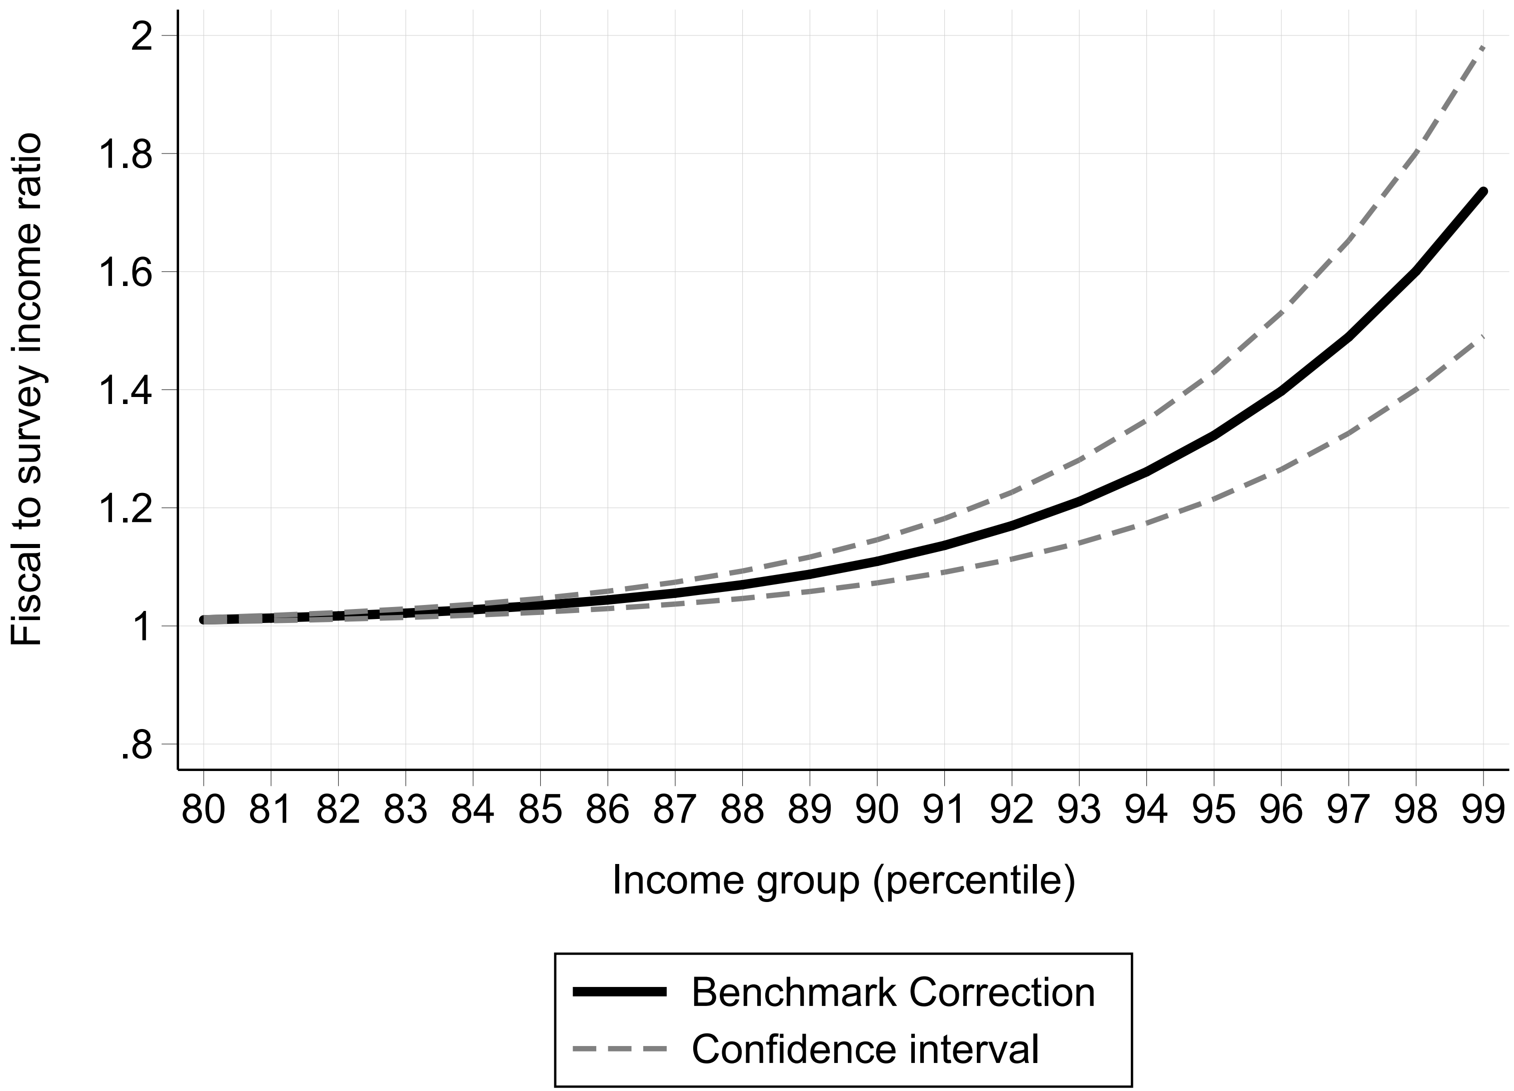

Notes. Authors' elaboration. Profiles correspond to functions of the form $c_2(p) = 1 + \sigma p^{1/\gamma}$, with $\gamma = 0.05$ and σ taking 0.05, 0.6, and 1.2.

Figure 8.7: Top 10% Income Share in Morocco, 1984-2014
From Consumption Inequality to Corrected Income Inequality

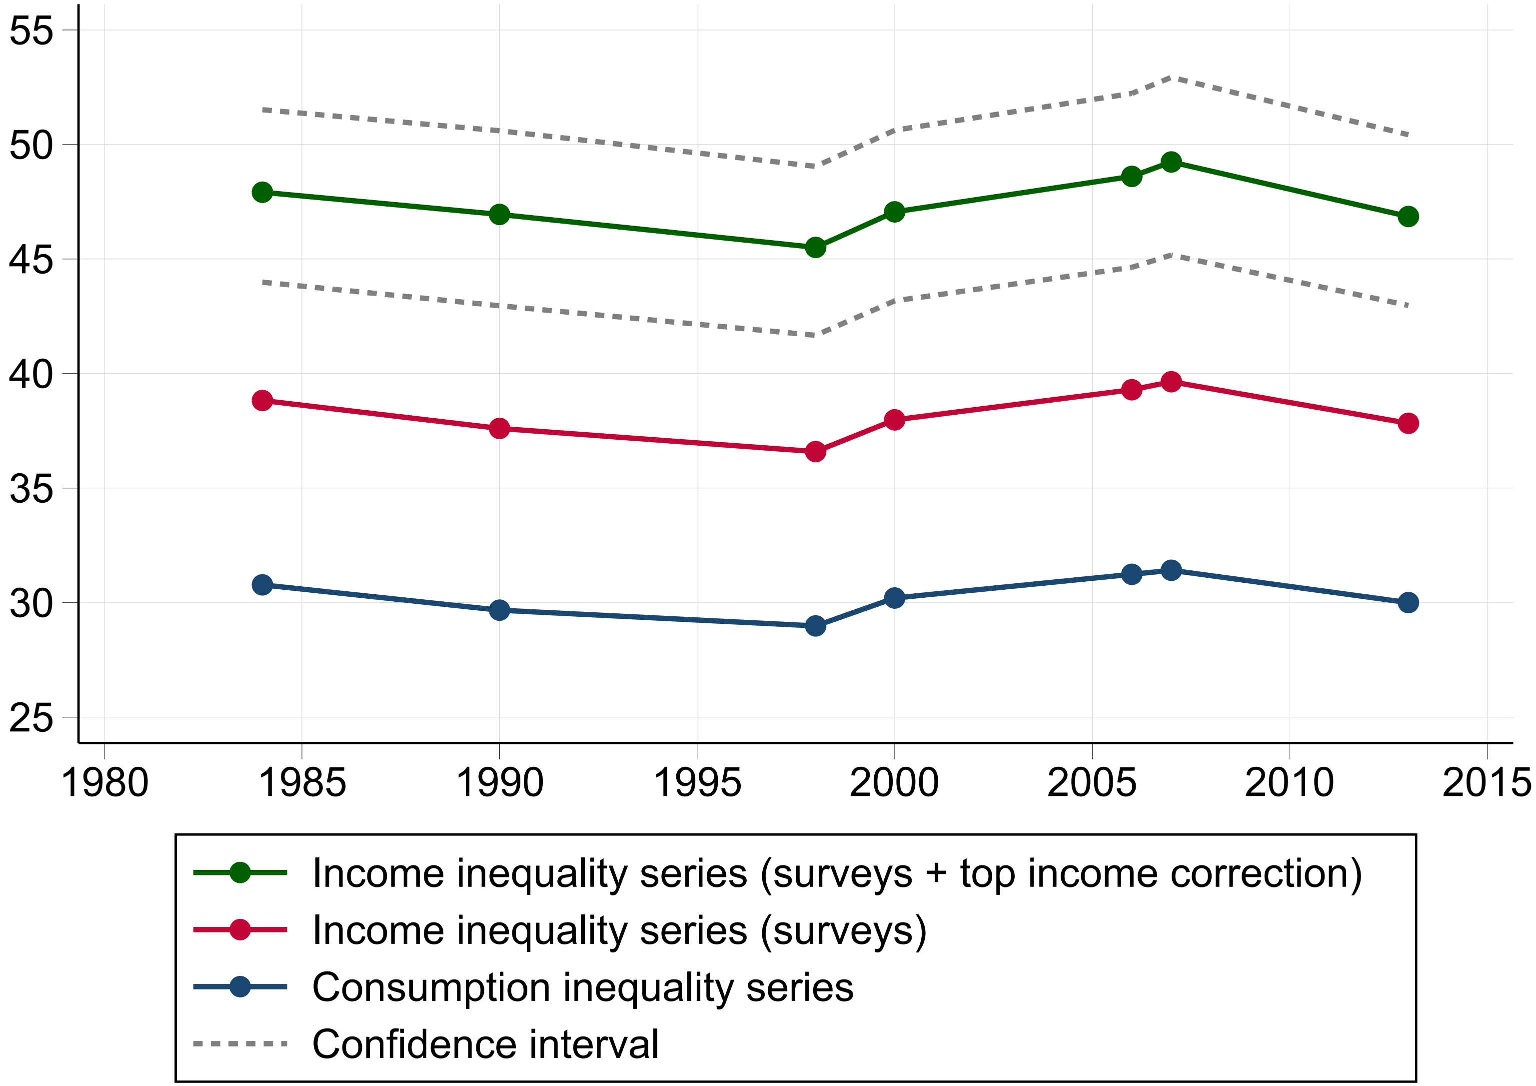

Notes. Authors' computations.

Figure 8.8: Top 10% Income Shares in Africa in 2019

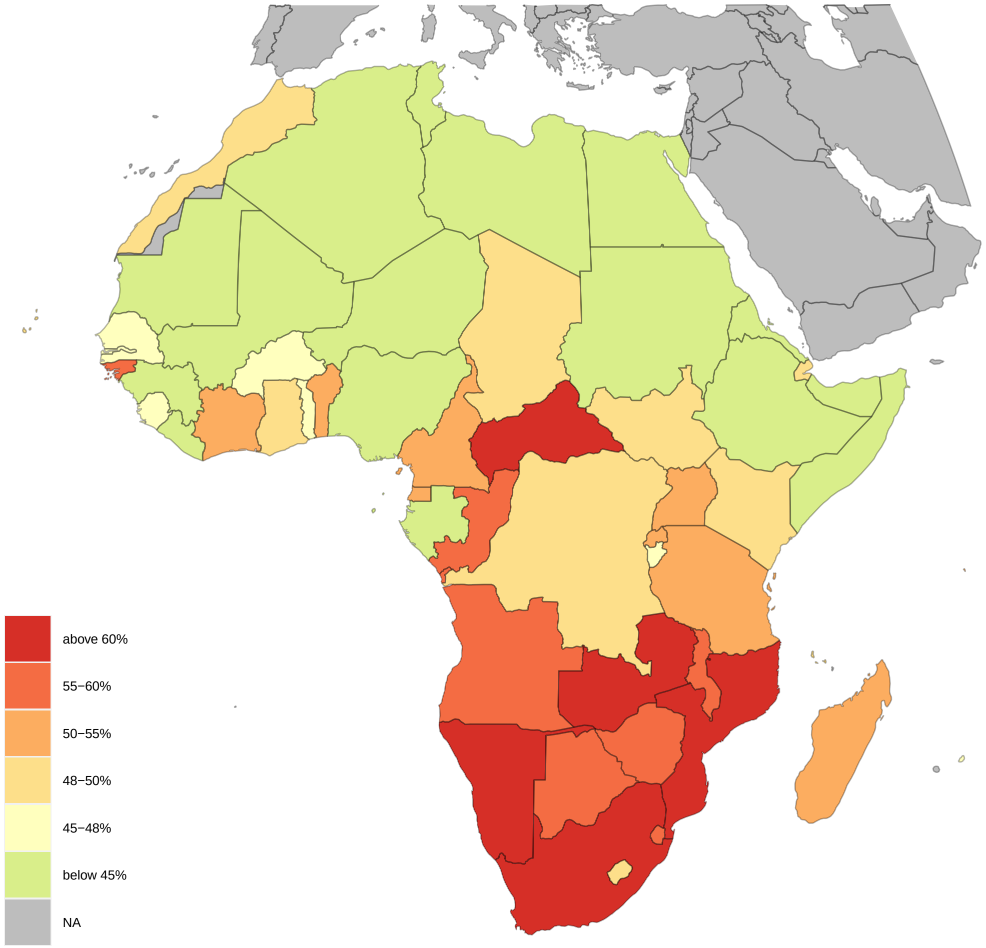

Notes. Authors' computations combining survey, tax, and national accounts data using the different methods presented. Interpolation between survey years and straightforward extrapolation are implemented to estimate current levels of inequality.

Figure 8.9: Bottom 40% Income Shares in Africa in 2019

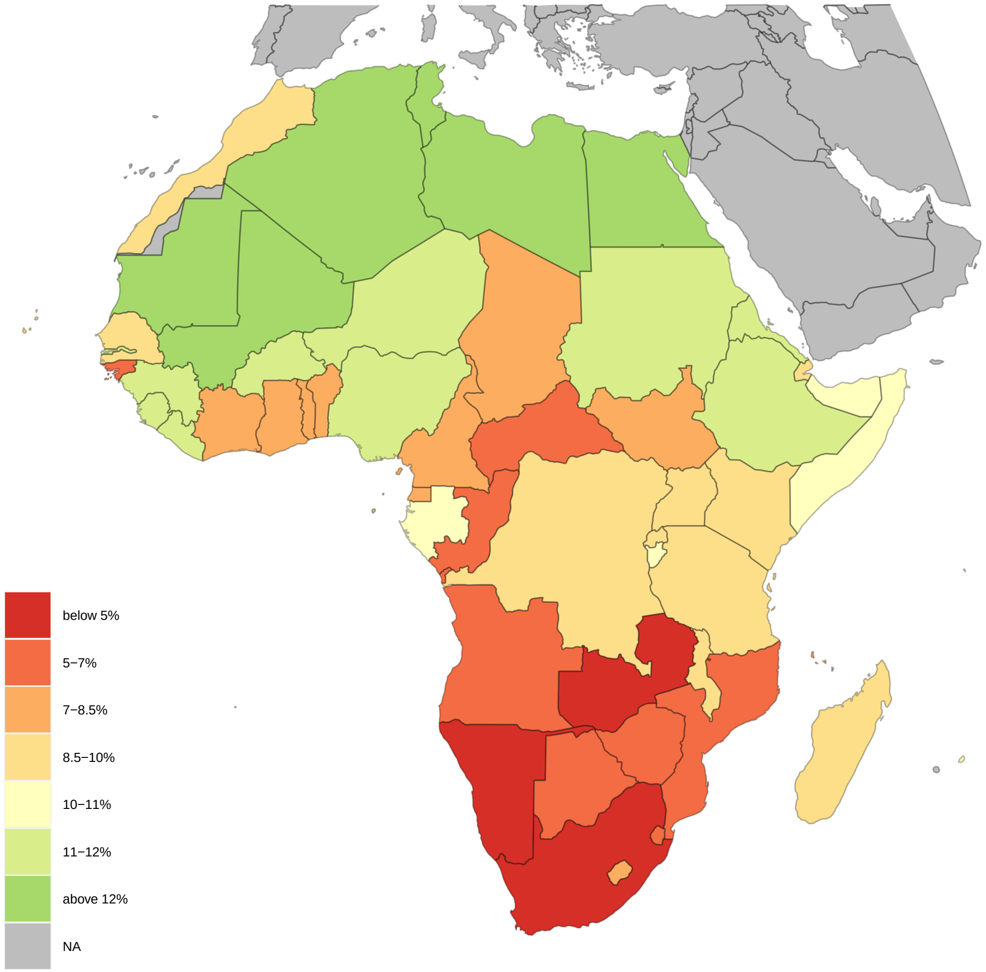

Notes. Authors' computations combining survey, tax, and national accounts data using the different methods presented. Interpolation between survey years and straightforward extrapolation are implemented to estimate current levels of inequality.

Figure 8.10: Evolution of the Pan-African Income Distribution

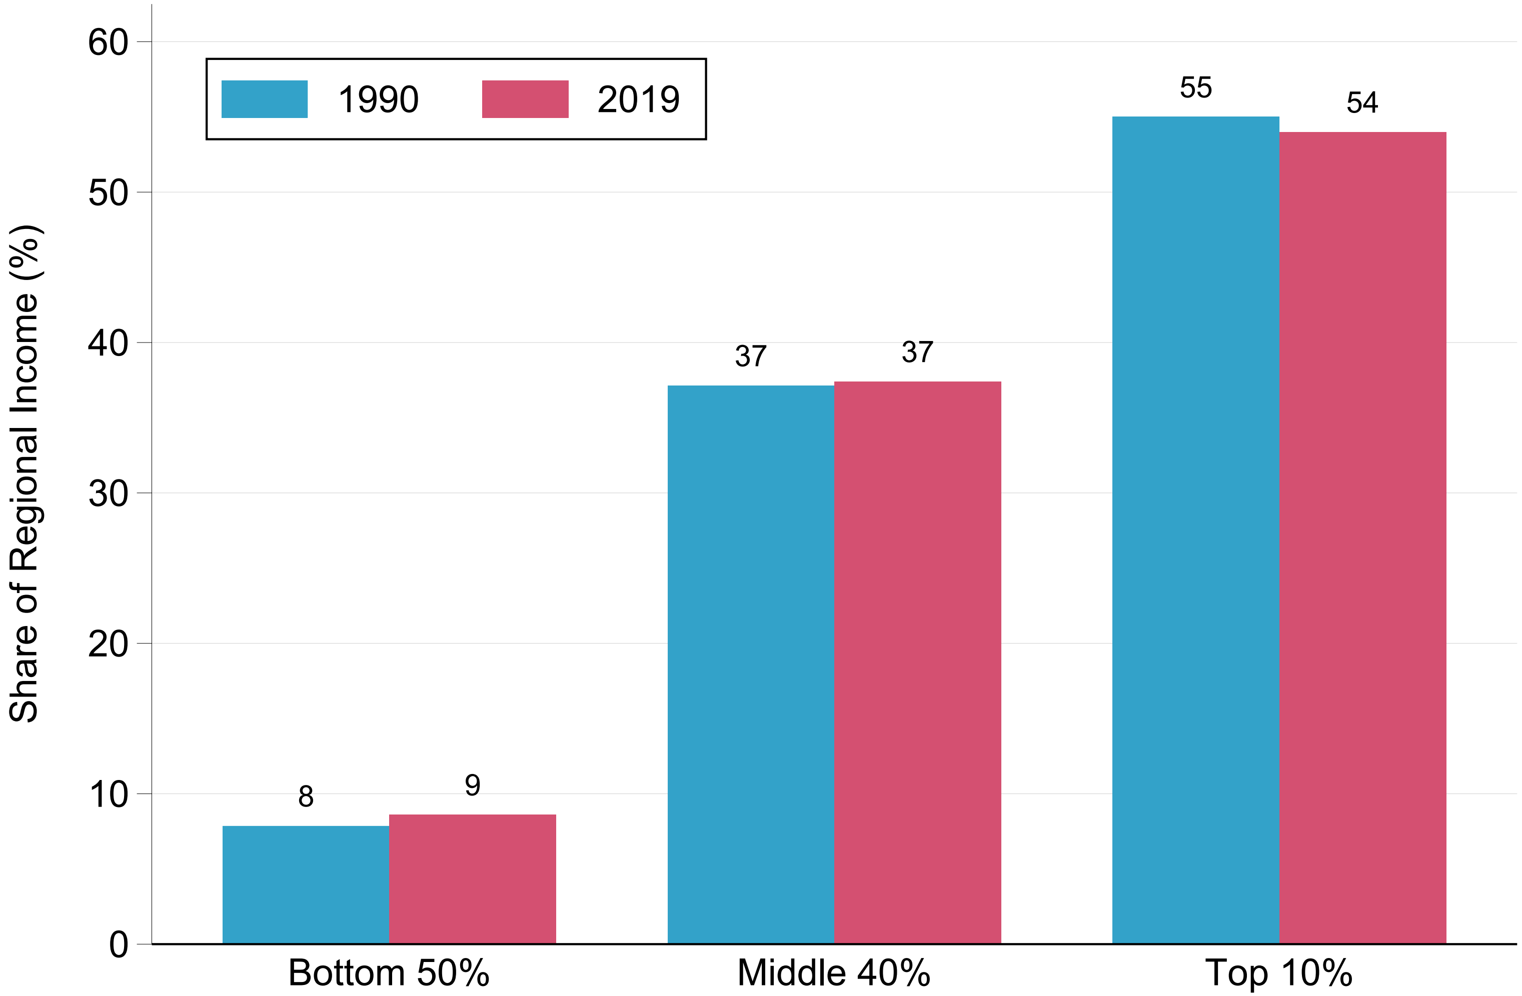

Notes. Authors' computations combining survey, tax, and national accounts data using the different methods presented. Interpolation between survey years and straightforward extrapolation are implemented to estimate current levels of inequality.

Figure 8.11: Decomposing Pan-African Inequality: Top 10% Income Share (1990-2019)

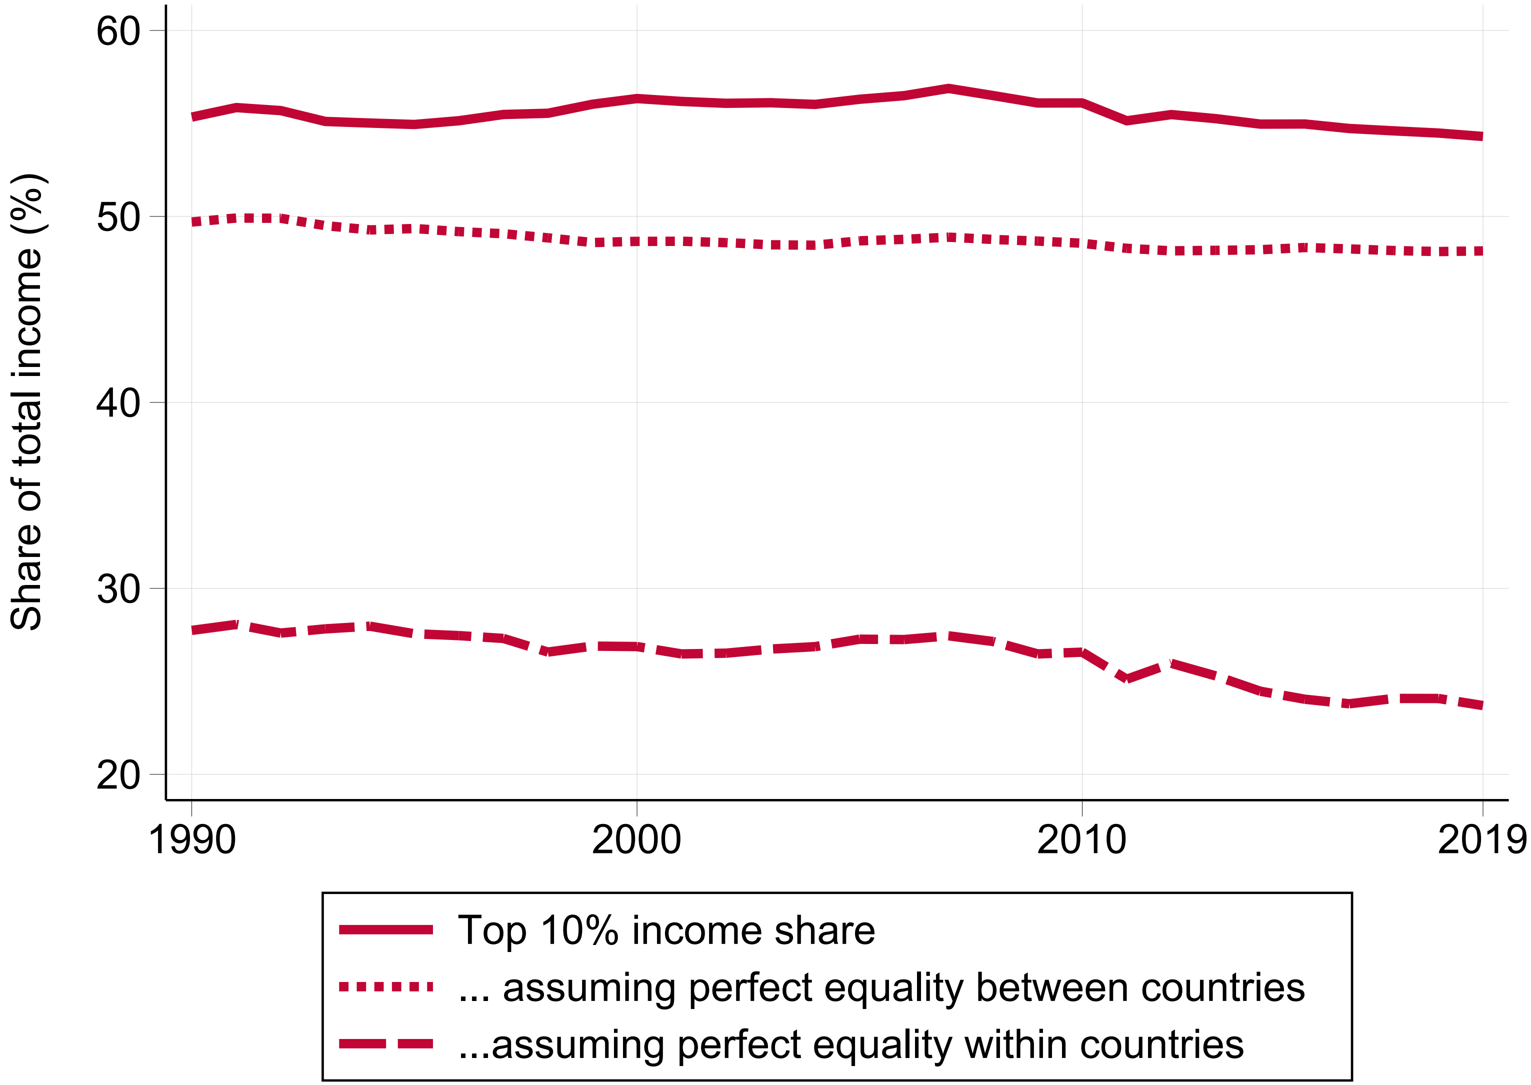

Notes. Authors' computations combining survey, tax, and national accounts data using the different methods presented. Interpolation between survey years and straightforward extrapolation are implemented to estimate current levels of inequality.

Figure 8.12: General Government Revenue in 2019 (% of GDP)

(a) In the world

(b) In Africa

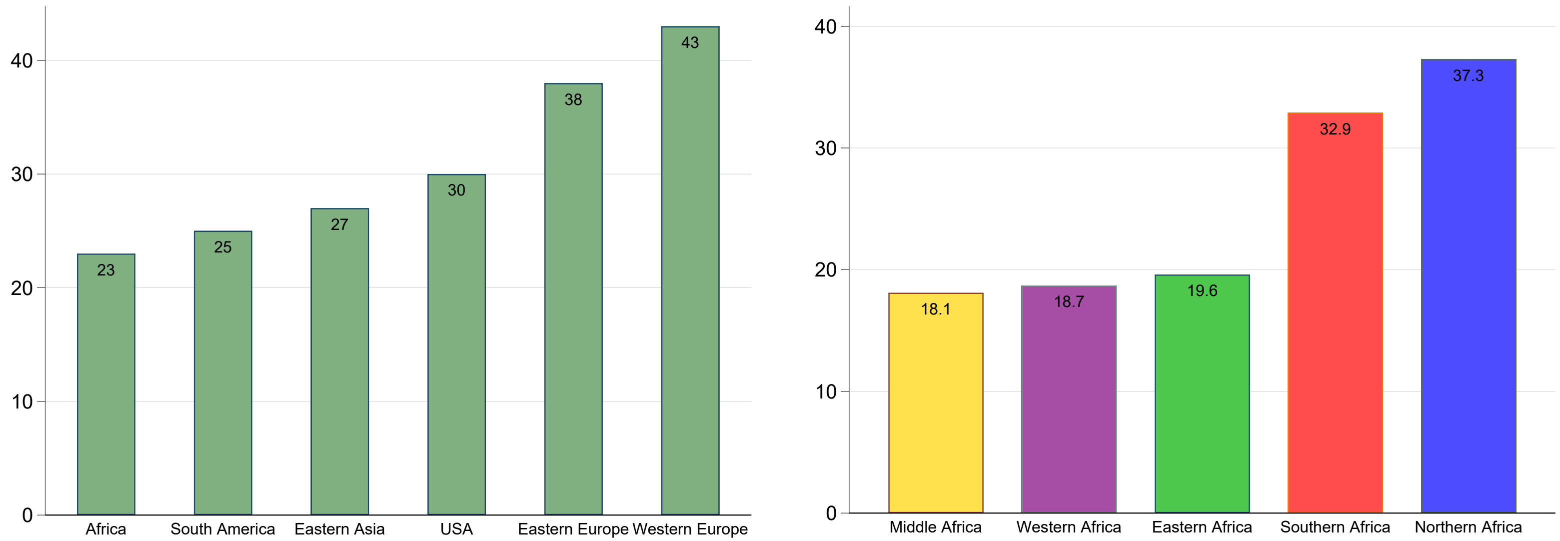

Notes. Authors' computations using data from the World Economic Outlook (International Monetary Fund). General government revenue consists of taxes, social contributions, grants receivable, and other revenue.

Figure 8.13: Characteristics of African Tax-and-Transfer Systems

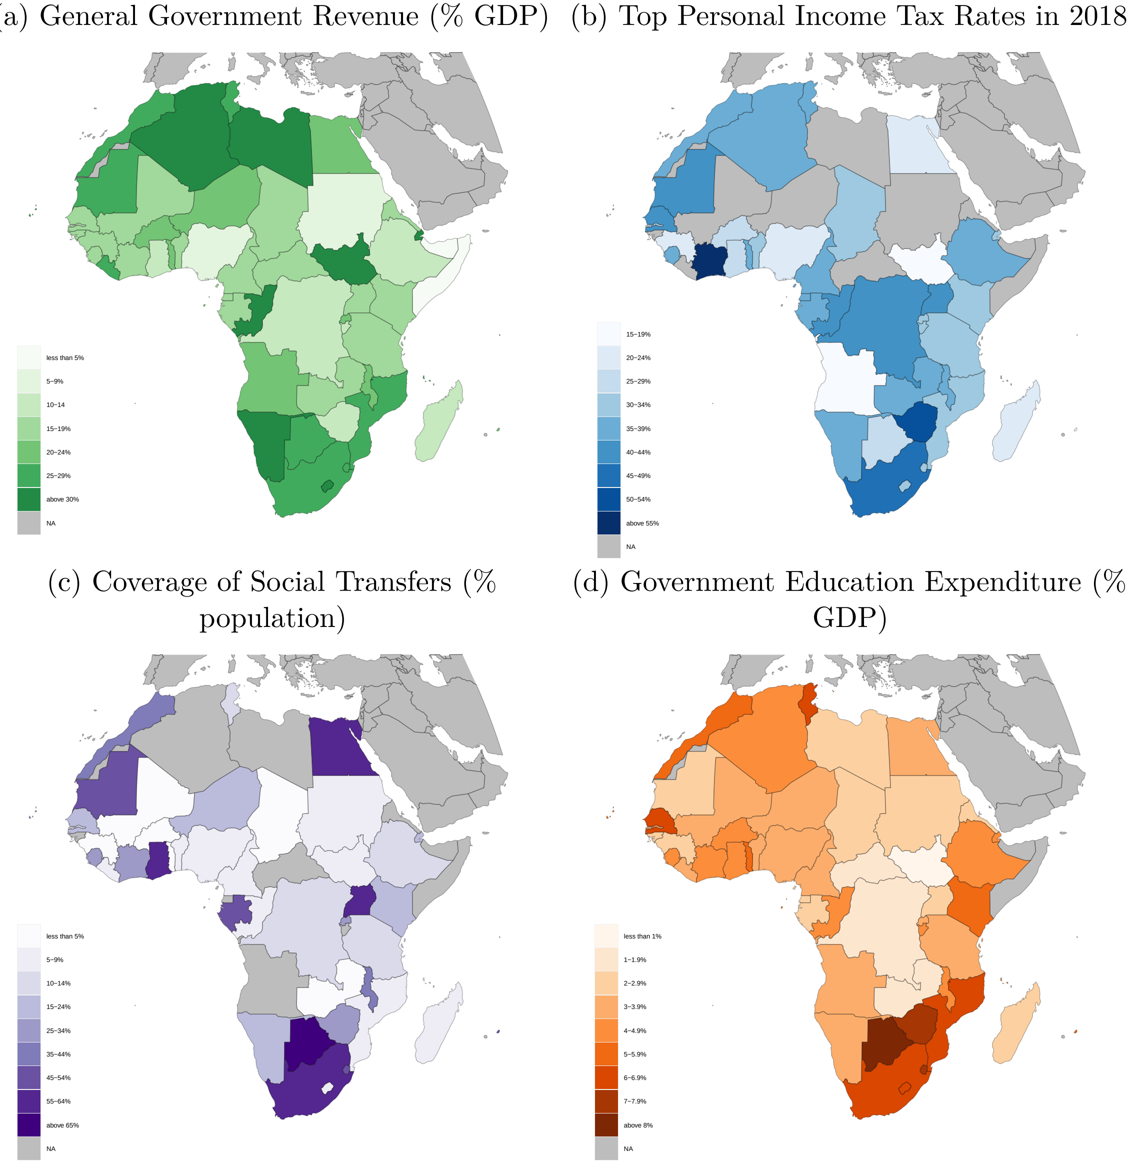

Notes. Authors' computations combining data from the World Economic Outlook (International Monetary Fund), Deloitte (Guide to fiscal information: Key economies in Africa, 2018), the Ernst & Young 2018-19 Worldwide Personal Tax and Immigration Guide, 2019, and the World Development Indicators (World Bank). General government revenue consists of taxes, social contributions, grants receivable, and other revenue. Data is from 2018 for government revenue, from most recent available years for other variables. Coverage of social protection and labor programs (SPL) shows the percentage of population participating in social insurance, social safety net, and unemployment benefits and active labor market programs. Estimates include both direct and indirect beneficiaries. Government expenditure on education (current, capital, and transfers) corresponds to all expenditure by the general government, including expenditure funded by transfers from international sources.

Table 8.1: European settlement and Islam correlates versus regional differences. Sub-Saharan Africa

	Top 10% income share			Bottom 50% income share		
	A	B	C	A	B	C
European settlement	+0.089***		+0.075***	-0.038***		-0.028**
	(0.021)		(0.023)	(0.010)		(0.010)
Muslims share	-0.066***		-0.058**	+0.034***		+0.029**
	(0.022)		(0.026)	(0.010)		(0.012)
North-Eastern		-0.066**	-0.025		+0.033***	+0.014
		(0.026)	(0.026)		(0.012)	(0.012)
Western		-0.056***	-0.020		+0.025***	+0.009
		(0.019)	(0.020)		(0.008)	(0.009)
Southern		+0.065**	+0.018		-0.035***	-0.017
		(0.026)	(0.026)		(0.012)	(0.012)
Small islands		-0.048	-0.071**		+0.024*	+0.034***
		(0.029)	(0.027)		(0.013)	(0.012)
F-test regional variables (p-value)		0.000	0.592		0.000	0.277
N	49	49	49	49	49	49
Adj. R^2	0.439	0.348	0.508	0.436	0.415	0.543

Source: authors' computations. Standard errors in parentheses; *: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$. European settlement: Dummy for whether European settlers went above 2.5% of total population between 1870 and 1970 (Easterly and Levine, 2016). Eur. settlement: Angola, Eswatini, Mozambique, Mauritius, Namibia, South Africa, Zambia, Zimbabwe. Muslim share: proportion of Muslims in total population circa 2010. Muslims > 50%: Burkina Faso, Chad, Comoros, Djibouti, Guinea, Guinea Bissau, Mali, Mauritania, Niger, Sudan, Senegal, Sierra Leone, Somalia. North-Eastern: Djibouti, Eritrea, Ethiopia, Somalia, Sudan, South Sudan. Western: Benin, Burkina Faso, Cabo Verde, Côte d'Ivoire, Gambia, Ghana, Guinea, Guinea Bissau, Liberia, Mali, Mauritania, Niger, Nigeria, Senegal, Sierra Leone, Togo. Eastern (omitted): Burundi, Comoros, Kenya, Madagascar, Mauritius, Mozambique, Malawi, Rwanda, Seychelles, Tanzania, Uganda, Zambia. Southern: Botswana, Eswatini, Lesotho, Namibia, South Africa, Zimbabwe. Small islands: Islands that were uninhabited before slave trade and colonization: C. Verde, Mauritius, São Tome & P., Seychelles. F-test for regional variables does not include the small islands dummy.

Table 8.2: European settlement and Islam correlates versus geography, precolonial history, and colonizers' identity. Sub-Saharan Africa

	Top 10% income share					Bottom 50% income share				
	A	B	C	D	E	A	B	C	D	E
European settlement	+0.089***	+0.073***	+0.089***	+0.074***	+0.058**	-0.038***	-0.029**	-0.038***	-0.032**	-0.024**
	(0.020)	(0.028)	(0.021)	(0.021)	(0.029)	(0.009)	(0.010)	(0.012)	(0.009)	(0.013)
Muslims share	-0.079***	-0.092***	-0.080***	-0.090***	-0.101***	+0.041***	+0.050***	+0.039***	+0.050***	+0.057***
	(0.020)	(0.024)	(0.027)	(0.022)	(0.032)	(0.009)	(0.010)	(0.012)	(0.009)	(0.013)
Controls:		p-value	p-value	p-value	p-value	p-value	p-value	p-value	p-value	p-value
Geography		0.587			0.775		0.173			0.475
Slave exports			0.199		0.194			0.155		0.106
Precolonial pol.			0.809		0.684			0.874		0.588
Ethnic fract.			0.863		0.743			0.972		0.976
Colonizer ident.				0.211	0.274				0.058	0.078
N	49	49	49	49	49	49	49	49	49	49
Adj. R^2	0.520	0.506	0.495	0.545	0.506	0.533	0.568	0.520	0.594	0.624

Source: authors' computations. Standard errors in parentheses; *: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$.
European settlement and Muslims share: see Table 8.1 and text. Geography: Abs. latitude, longitude, min month. avg rainfall, max month. afternoon avg humidity, min avg month. low temp, log(coastline/area). (Nunn, 2008).
Slave exports: Log total slave exports normalized by historic population (Nunn, 2008); results are similar with slave exports normalized by land area. Precolonial polities: Percentages of population from Centralized Stratified, Centr. Egalitarian, and Fragmented Strat.groups; Frag. and Egal. being omitted (Gennaioli and Rainer, 2007). The variables were constructed using the dataset from Michalopoulos and Papaioannou (2013), as some countries were missing in Gennaioli and Rainer (2007). Ethnic fractionalization: Alesina et al. (2003). Sao Tome and Principe was set at the value for Cabo Verde. Colonizer identity: Dummy variables for the last colonizer being either Belgian, British, French, or Portuguese (Somalia has 0.5 for British as it was shared with Italy), and for non-colonized (Ethiopia and Liberia).
In all regressions, a "small island" dummy is included: Cabo Verde, Mauritius, São Tome & P., Seychelles. These islands were uninhabited before slavery and colonization. For them, the precolonial dummies were set at zero (meaning 100% was fragmented and egalitarian); given the small island dummy, this has no impact on reported point estimates.

Chapter 9

Brahmin Left Versus Merchant Right: Changing Political Cleavages in 21 Western Democracies, 1948–2020

Western democracies have undergone deep transformations in recent years, embodied by political fragmentation, the increasing salience of environmental issues, and the growing success of anti-establishment authoritarian movements (Trump, Brexit, Le Pen, etc.). Yet, much remains to be understood about the nature and origins of these political upheavals. On what dimensions of political conflict (education, income, age, etc.) have such transformations aligned? Is the rise of "populism" the outcome of recent trends (such as the 2007-2008 crisis, immigration waves, or globalization), or can we trace it back to longer-run structural changes? Beyond country-specific factors, can we find evolutions that are common to all Western democracies?

This paper attempts to make progress in answering these questions by exploiting a new dataset on the long-run evolution of electoral behaviors in 21 democracies. Drawing on nearly all electoral surveys ever conducted in these countries since the end of World War II, we assemble microdata on the individual determinants of the vote for over 300 elections held between 1948 and 2020. Together, these surveys provide unique insights into the evolution of voting preferences in Western democracies. The contribution of this paper is to establish a new set of stylized facts on these preferences, as well as to explore some mechanisms underlying their transformation in the past decades.[1]

Comparing the evolution of electoral cleavages requires grouping political parties in such a way that the coalitions considered are as comparable across countries and over time as possible. To do so, we start by making a distinction between two large groups of parties: social democratic, socialist, communist, and green parties ("left-wing" or "social democratic and affiliated" parties) on one side, and conservative, Christian democratic, and anti-immigration parties ("right-wing" or "conservative and affiliated" parties) on the other side.[2]

The most relevant result that emerges from our analysis is the existence of a gradual process of disconnection between the effects of income and education on the vote. In the 1950s-1960s, the vote for social democratic and affiliated parties was "class-based," in the sense that it was strongly associated with the lower-income and lower-educated electorate. It has gradually become associated with higher-educated voters, giving rise in the 2010s to a divergence between the influences of income (economic capital) and education (human capital): high-income voters continue to vote for the "right", while high-education voters have shifted to supporting the "left." This separation between a "Merchant right" and a "Brahmin left" is visible in nearly all Western democracies, despite their major political, historical, and institutional differences.[3] We also find that the rise of both green and anti-immigration parties since the 1980s-1990s has accelerated this transition—although it can only explain about 15% of the overall shift observed—, as education, not income, most clearly distinguishes support for these two families of parties today.

As a result, many Western democracies now appear to have shifted from "class-based" to "multidimensional" or "multiconflictual" party systems, in which income and education differentially structure support for competing political movements. One might also call these systems "multi-elite" party systems, in which governing

[1]This paper is part of a broader collective project dedicated to tracking political cleavages in fifty democracies throughout the world: see Gethin, Martínez-Toledano, and Piketty (2021). Several chapters of this collective volume are dedicated to discussing at greater length the results introduced in this paper in the case of specific countries, in particular Piketty (2021); Kosse and Piketty (2021); Martínez-Toledano and Sodano (2021); Gethin (2021); Bauluz et al. (2021); and Durrer, Gethin, and Martínez-Toledano (2021). All the data series, computer codes, and microfiles of this collaborative project can be publicly accessed online as part of the World Political Cleavages and Inequality Database (`http://wpid.world`).

[2]We also include parties commonly classified as liberal or social-liberal in this latter group, such as the Liberal Democrats in Britain and the Free Democratic Party in Germany. In Section II.B, we perform several robustness checks to ensure that our classification is consistent both in terms of parties' programmatic supply and voters' own perceptions of the political space.

[3]In India's traditional caste system, upper castes were divided into Brahmins (priests, intellectuals) and Kshatryas/Vaishyas (warriors, merchants, tradesmen), a division that modern political conflicts in Western democracies therefore seem to follow to some extent.

coalitions alternating in power tend to reflect the views and interests of a different kind of elite (intellectual versus economic), assuming that elites do have a greater influence on political programs and policies than the rest of the electorate.[4]

To shed light on the factors underlying the divergence of the effects of income and education on the vote, we match our dataset with the Comparative Manifesto Project database, the most comprehensive available data source on the evolution of political parties' programs since the end of World War II. Drawing on two indicators of party ideology from the political science literature (Bakker and Hobolt, 2013), corresponding to parties' relative positions on an "economic-distributive" axis and a "sociocultural" axis, we provide evidence that the separation between these two dimensions of political conflict and the divergence of income and education are tightly related phenomena. Specifically, we document that the correlation between parties' income gradient and their position on the economic-distributive dimension has remained very stable since the 1960s: parties emphasizing "pro-free-market" issues receive disproportionately more votes from high-income voters today, just as they used to sixty years ago. Meanwhile, the correlation between the education gradient and parties' positions on the sociocultural axis has dramatically increased over time, from 0 in the 1960s to nearly 0.5 in the 2010s.

In other words, parties promoting "progressive" policies (green and traditional left-wing parties) have seen their electorate become increasingly restricted to higher-educated voters, while parties upholding more "conservative" views on sociocultural issues (anti-immigration and traditional right-wing parties) have on the contrary concentrated a growing share of the lower-educated electorate. We also find a strong and growing cross-country association between ideological polarization on sociocultural issues and the reversal of the education cleavage. In particular, the two countries in our dataset where this reversal has not yet occurred, Portugal and Ireland, are also those where partisan divides over these issues remain the weakest today. Taken together, these results suggest that changes in political supply, in particular the increasing emphasis on sociocultural factors among old and new parties, appear to be an important factor behind the progressive disconnection between educational and income divides.

We should stress, however, that the limitations of available information on party

[4] A large literature in economics and political science has documented the existence of unequal political representation and the distortion of politicians' and legislators' beliefs toward their most privileged constituencies: see, for instance, Adams and Ezrow (2009); Bartels (2017b); Bertrand et al. (2020); Bonica et al. (2013); Cagé (2020); Kuhner (2014); Gilens (2012); Gilens and Page (2014); Pereira (2021).

manifestos constrain our ability to carry a causal analysis or fully test the hypotheses behind the empirical regularities we uncover. In particular, the sociocultural axis puts together many different items that may also involve various forms of economic conflict over the consequences of environmental, migration, or education policies. The manifesto data do not provide information on the actual policies implemented by governing coalitions either. For instance, social democratic and affiliated parties may continue emphasizing redistributive policies just as they used to in the past, but their credibility in effectively pursuing these policies may have declined since then. Another complementary interpretation of our findings is that left-wing parties have gradually developed a more elitist approach to education policy, in the sense that they have increasingly been viewed by less well-off voters as parties defending primarily the winners of the higher education competition.[5] Unfortunately, the data at our disposal makes it difficult to provide a direct test for these various hypotheses. The fact that turnout has fallen sharply among both the bottom 50% least educated and poorest voters in a number of countries, but not among the top 50%, could be interpreted as a sign that socially disadvantaged voters have felt left aside by the rise of "multi-elite" party systems.[6]

We also investigate to what extent shifts in the composition of education groups in terms of gender, age, or other socioeconomic variables could account for the reversal of the education cleavage. To do this, we compare the education gradient before and after controlling for all available covariates in our database. We also carry a Kitagawa-Oaxaca-Blinder decomposition of the education gradient, which allows us to formally estimate what fraction of the reversal can be accounted for by structural changes in educational achievement. Both methods yield identical results: compositional effects can only predict 16% to 17% of the transformation of educational divides observed since the 1950s.

[5]This risk was identified as early as in 1958 by Michael Young in his famous dystopia about "the rise of the meritocracy" (Young, 1958). In this book, Young expresses doubts about the ability of the British Labour Party (of which he was a member) to keep the support of lower educated classes in case the party fails to combat what he describes as the rise of "meritocratic ideology" (a strong view held by higher education achievers about their own merit, which Young identifies as a major risk for future social cohesion). For a simple theoretical model along these lines, see Piketty (2018). It is based upon a two-dimensional extension of the Piketty (1995) model about learning the role of effort and a distinction between education-related effort and business-related effort. The model can account for the simultaneous existence of "Brahmin left" voters (i.e., dynasties believing strongly in the role of education-related effort) and "Merchant right" voters (i.e., dynasties believing strongly in the role of business-related effort).

[6]See Piketty (2018), figures A1-A2. Turnout rates among bottom 50% voters have always been relatively low in the US (at least during the post-World War II period). To some extent the British and French pattern has moved toward the US pattern since the 1970s-1980s. Unfortunately, the surveys at our disposal do not allow us to analyze in a consistent manner the evolution of turnout in our sample of 21 countries, so we do not push any further our analysis of turnout.

We do find, however, some heterogeneity in the reversal when further decomposing voters into subgroups by different socioeconomic characteristics. Generational dynamics appear to have mattered tremendously in generating the reversal of the education cleavage: while older lower-educated voters continue to vote “along class lines” and thus to support the left, social democratic and green parties have attracted a growing share of the higher-educated electorate among the youth. The reversal in the educational divide has also been highest among non-religious voters and among men, although it has happened within other subgroups too. Overall, the disconnection of income and education cleavages has been a relatively independent and widespread phenomenon, in the sense that it cannot be accounted for by other socioeconomic variables and is not linked to any particular subgroup of voters.

Finally, we also exploit the other variables in our dataset to study cleavages related to age, geography, religion, gender, and other socioeconomic variables. The main conclusion is that there has been no major realignment of voters along these other dimensions, comparable to the one observed in the case of education. Younger voters are more likely to vote for social democratic and affiliated parties, but this was already the case by a comparable magnitude in the 1950s. Similarly, rural-urban and religious cleavages have remained stable or have decreased in most countries in our dataset: rural areas and religious voters continue to be supportive of conservative parties, as they used to in the past. The major exception is gender, the only variable other than education for which we find a clear reversal of electoral divides: in nearly all countries, women used to be more conservative than men and have gradually become more likely to vote for left-wing parties.

This paper directly relates to the growing literature on the sources of political change and the rise of “populism” in Western democracies. Recent studies have emphasized the role of various economic and sociocultural factors, including globalization and trade exposure (Autor et al., 2020; Colantone and Stanig, 2018a,b; Malgouyres, 2017), economic insecurity and unemployment (Algan et al., 2017; Becker and Fetzer, 2018; Becker, Fetzer, and Novy, 2017; Dehdari, 2021; Fetzer, 2019; Funke, Schularick, and Trebesch, 2016; Guiso et al., 2020; Liberini et al., 2019), immigration (Becker and Fetzer, 2016; Dustmann, Vasiljeva, and Damm, 2019; Halla, Wagner, and Zweimüller, 2017; Tabellini, 2020), and cultural and moral conflicts (Bonomi, Gennaioli, and Tabellini, 2021; Enke, 2020; Norris and Inglehart, 2019). We contribute to this body of evidence by adopting a broader, long-run historical perspective on the evolution of political cleavages since the end of World War II. We find little evidence that the shifts in electoral divides we observe were driven by single, major events such as the end of the Cold War, the increasing salience of immigration since the 2000s, trade shocks, or the 2007-2008 crisis. What seems to have happened instead is a very progressive, continuous reversal of educational divides, which unfolded decades before any of these events took place and has carried on uninterruptedly until today.

We also contribute to the literature on multidimensional political competition and its impact on redistribution and inequality. A key result from this literature is that political support for redistribution should be inversely proportional to the strength of other political cleavages crosscutting class divides (Alesina, Baqir, and Easterly, 1999; Alesina, Glaeser, and Sacerdote, 2001; Bonomi, Gennaioli, and Tabellini, 2021; Roemer, 1998; Roemer et al., 2007). The divergence of the effects of income and education on the vote documented in this paper, two highly correlated measures of inequality, could in this context contribute to explaining why the rise of economic disparities in the past decades has not been met by greater redistribution or renewed class conflicts.

Finally, this paper relates to the large political science literature on the determinants of the vote in comparative and historical perspective. Numerous studies have highlighted that Western democracies have undergone a process of growing polarization over a new "sociocultural," "universalistic-particularistic," or "green / alternative / libertarian versus traditional / authoritarian / nationalist" dimension of political conflict in the past decades (Bornschier, 2010a; Dalton, 2018; Evans and De Graaf, 2012; Hooghe, Marks, and Wilson, 2002; Inglehart, 1977; Kitschelt, 1994; Kriesi et al., 2008; Norris and Inglehart, 2019). There is also extensive evidence that education has been playing a major role in restructuring electoral behaviors and collective beliefs along this new dimension in recent decades (Bornschier, 2010b; Bovens and Wille, 2012; Dolezal, 2010; Duch and Taylor, 1993; Ford and Jennings, 2020; Kitschelt and Rehm, 2019; Langsæther and Stubager, 2019; Rydgren, 2013, 2018; Stubager, 2008, 2010, 2013; Waal, Achterberg, and Houtman, 2007). We contribute to this literature by gathering the largest dataset ever built on the socioeconomic determinants of the vote in Western democracies[7]; by focusing explicitly on the distinction between income and education, two variables whose effects are rarely studied jointly in comparative studies; and by directly matching this dataset with historical data on party ideology to document the dynamic links between political supply and demand.[8] In doing so,

[7]Our work directly draws on previous data collection and harmonization efforts. See in particular Bosancianu (2017), Elff (2007), Evans and De Graaf (2013), Franklin, Mackie, and Valen (1992), Thomassen (2005), and Önudottir, Schmitt, and al. (2017), and the collections of post-electoral surveys compiled by the Comparative Study of Electoral Systems (`http://cses.org`) and the Comparative National Elections Project (`https://u.osu.edu/cnep/`).

[8]In matching survey and manifesto data, we follow recent political science studies seeking to understand how political supply influences class and religious divides. See in particular Elff (2009), Evans and De Graaf (2013), Evans and Tilley (2012, 2017), Jansen, De Graaf, and Need (2011,

we confirm many of the findings of the existing literature, but we also provide new insights into the transformation of political cleavages in Western democracies. In particular, we gather for the first time cross-country, long-run historical evidence of a gradual dissociation of the effects of education and income on the vote. This dissociation appears to have started as early as the 1950s and to have unfolded uninterruptedly since then, and can be related to the growing salience of a large and complex set of policy issues, including the environment, migration, gender, education, and merit, which divide voters along educational but not income lines.

Section II presents the new dataset exploited in this paper. Section III documents the divergence of the income and education effects and discusses the role of green and anti-immigration parties in explaining the reversal of the education cleavage. Section IV matches our survey dataset with manifesto data to study the link between this transformation and the emergence of a new axis of political conflict. Section V explores alternative explanations and heterogeneity in the reversal of the education cleavage and analyzes the evolution of other determinants of electoral behaviors. Section VI concludes.

9.1 Data and Methodology

9.1.1 A New Dataset on Political Cleavages in Western Democracies, 1948-2020

The dataset we exploit in this paper consists in a collection of electoral surveys conducted between 1948 and 2020 in Western democracies. These surveys have one main point in common: they contain information on the electoral behaviors of a sample of voters in the last (or forthcoming) election, together with data on their main sociodemographic characteristics such as income, education, or age. While they suffer from limitations typical to surveys (such as small sample sizes), they provide an invaluable source for studying the long-run evolution of political preferences in contemporary democracies.

Universe. Our area of study encompasses 21 countries commonly referred to as "Western democracies", for which we can cover a total of about 300 national elections (see Table I). These include 17 Western European countries, the United States, Canada, Australia, and New Zealand. For seven countries in our dataset (France, Germany, Italy, Norway, Sweden, the UK, and the US), available surveys allow us

2012), Jansen, Evans, and de Graaf (2013), and Rennwald and Evans (2014).

to go back as early as the 1950s. The majority of remaining countries have data going back to the 1960s or the early 1970s, with the exception of Spain and Portugal, which did not hold democratic elections between the 1940s and the late 1970s.

The focus of this paper is on national (general or presidential) elections, which determine the composition of government and the head of the State. In the majority of Western democracies, they have been held on a regular basis every four or five years since at least the end of World War II. Depending on their frequency and the availability of electoral surveys, we are able to cover political attitudes in 9 to 21 of these elections in each country.

Data sources. Our primary data source consists in so-called National Election Studies, most of which have been conducted by a consortium of academic organizations (see Table I). The majority of these surveys are post-electoral surveys: they are fielded shortly after the corresponding national election has been held, with sample sizes generally varying between 2,000 and 4,000 respondents, and they collect detailed information on voting behaviors and the sociodemographic characteristics of voters.

In all Western democracies except Austria, Ireland, and Luxembourg, we have been able to get access to such high-quality data sources. For these three countries, we rely instead on more general political attitudes surveys, which were not specifically conducted in the context of a given election but did ask respondents to report their previous voting behaviors: the Eurobarometers, the European Social Survey, and the European Election Studies. Furthermore, in a few countries such as Australia or Belgium, where national election studies were not conducted prior to the 1970s or 1980s, we complement them with other political attitudes surveys conducted in earlier decades. While these sources do not allow us to accurately track election-to-election changes, they are sufficient to grasp long-run changes in party affiliations, which is the objective of this paper.[9]

Harmonization. Starting from raw data files, we extract in each survey all sociodemographic characteristics that are sufficiently common and well-measured to be comparable across countries and over time. Based on these criteria, we were able to build a harmonized dataset covering the following variables: income, education, age, gender, religious affiliation, church attendance, race or ethnicity (for a restricted number of countries), rural-urban location, region of residence, employment status, marital status, union membership, sector of employment, home ownership, self-perceived social class, and (in recent years) country of birth.

[9]A complete list of all data sources used by country can be found in appendix Table A1.

Income and education, the two variables that form the core part of our analysis in section III, deserve special attention. Indeed, one reason why income and education variables are not often studied jointly in large-scale comparative studies on electoral behaviors is that they tend to be difficult to harmonize. Education systems and educational attainments vary significantly across countries and over time, and they are not always perfectly comparable across surveys. The same limitations apply to income, which is only collected in discrete brackets in the majority of the sources used in this paper.

We address this shortcoming by normalizing these two variables and focusing on specific education and income deciles. Appendix A introduces the method we use to move from discrete categories (education levels or income brackets) to deciles. In broad strokes, our approach consists in allocating individuals to the potentially multiple income or education deciles to which they belong, in such a way that average decile-level vote shares are computed assuming a constant vote share within each education- or income-year cell. This is a conservative assumption, as vote shares for specific parties are likely to also vary within education groups or income brackets. The levels and changes in education and income cleavages documented in this paper should thus be considered as lower bounds of the true effects of education and income on the vote.

Lastly, in order to make surveys more representative of election outcomes, we systematically reweight respondents' answers to match official election results. Given that post-electoral surveys capture relatively well variations in support for the different parties, this correction leaves our results unchanged in the majority of cases.

9.1.2 Party Classification

Our objective is to compare the long-run evolution of electoral cleavages in Western democracies. This requires grouping political parties in such a way that the size of the coalitions considered and their historical affiliations are as comparable and meaningful as possible. To do so, we make a distinction between two large groups of parties in our main specification (see the coalitions delineated by dashed lines in Figure IV).[10]

On one side of the political spectrum are social democratic, socialist, communist,

[10]See appendix Tables A2 and A3 for more information on the classification of the main parties in each country. Parties not classified in either of these two groups mainly correspond to independent candidates and regional parties (such as the Bloc Québécois in Canada or the Scottish National Party). These parties or candidates have received about 7% of votes since 1945, with no clear trend (see Figure 4).

and green parties, often classified as "left-wing" and that we also refer to as "social democratic and affiliated parties" in what follows. These include the Democratic Party in the US, labor parties in countries such as the UK, Australia, or Norway, as well as various parties affiliated to socialist and social democratic traditions in Western European countries. It also includes environmental parties in their various forms, together with several new left-wing parties that emerged after the 2008 crisis (such as Podemos in Spain, Die Linke in Germany, or La France Insoumise in France).

On the other side are conservative, Christian democratic, and anti-immigration parties, often classified as "right-wing" and that we also refer to as "conservative and affiliated parties." These include the Republican Party in the US and other conservative parties such as those of the UK, Norway, and Spain; Christian democratic parties, which are common in Western European multi-party systems such as those of Austria, Belgium, and Switzerland; and anti-immigration parties such as the French Rassemblement National or the Danish People's Party. We also include parties commonly classified as liberal or social-liberal in this group, such as the Liberal Democrats in Britain, the Free Democratic Party in Germany, and the Liberal Party in Norway, but our results are robust to not doing so.[11]

This binary classification has one major advantage: it allows us to directly compare electoral divides in two-party systems, such as the UK or the US, to those observed in highly fragmented party systems such as France or the Netherlands. Aggregating parties into two large groups of comparable size in each country is thus useful to get a first perspective on the long-run evolution of political cleavages that is consistent both over time and across countries. These groups also correspond in many cases to the coalitions of parties that have effectively built political majorities, whether in coalition governments or through direct parliamentary support.

To make sure that this distinction between "left" and "right" is meaningful when it comes to differentiating parties and voters, we contrast two indicators for all parties: the average self-reported left-right position of voters supporting each of these parties, and the score of each of these parties on the left-right ideological index available from the Comparative Manifesto Project database. The first of these indicators is

[11]The exceptions are Austria, Canada, Denmark, and the Netherlands, for which we classify as "left-wing" parties generally considered to be liberal (NEOS in Austria, the Liberal Party in Canada, the Social Liberal Party in Denmark, and D66 in the Netherlands). This choice is motivated by our objective to compare coalitions of significant and comparable size across countries. Liberal parties have received about 10% of the vote in Western democracies since 1945 (see Figure 4), with no clear trend, and have consistently been supported by both high-income and higher-educated voters (see appendix Figures A26 and A28). Our results are thus robust to excluding them or not from the analysis.

available in most post-electoral surveys used in this paper, which have directly asked respondents to position themselves on a 0 (left) to 10 (right) scale. The second is a measure of parties' left-right positions that theoretically ranges from -100 (right) to 100 (left). It was first computed from manifesto data and validated by factor analysis by Budge and Laver (1992), and it has been widely used in comparative political science research since then (e.g., Evans and De Graaf, 2013).

We find that our categorization of political parties into two groups is very consistent with these two indicators. Every single party that we have classified as "social democratic and affiliated" is supported by voters who declare being more left-wing than the average voter, and is more left-wing than the average party on the CMP left-right ideological index.[12] This is true for social democratic and socialist parties, but also for green parties, which are all ranked as left-wing in survey and manifesto data. The same holds in the case of conservative, Christian democratic, and anti-immigration parties, which are nearly all identified as more right-wing than the average party or voter. Moreover, the two indicators of parties' positions on a left-right scale are also consistent with one another (the correlation between the two variables is 0.82). We are thus confident that our classification is meaningful in terms of both parties' programmatic supply and voters' own perceptions of the political space.

That being said, we are not claiming that these two groups are ideologically or programmatically homogeneous in any way, neither internally nor over time. Our objective is, on the contrary, to document how such large families or parties have aggregated diverse and changing coalitions of voters in the past decades. In section III, we thus consider in greater detail how specific subfamilies of parties, in particular, green and anti-immigration movements, have contributed to reshaping electoral divides in countries with multi-party systems.

9.1.3 Empirical Strategy

In the rest of the paper, we present results from simple linear probability models of the form:

$$y_{ict} = \alpha + \beta x_{ict} + C_{ict}\gamma \ + \varepsilon_{ict} \tag{9.1}$$

[12] See appendix Figures B16 and B17. The one single exception here is Fianna Fáil in Ireland, which we still choose to classify with left-wing parties to study a coalition of sufficient size (if we were to exclude it, the total vote share of the "left" would fall below 30% throughout the period considered).

Where y_{ict} is a binary outcome variable of interest (e.g. voting for left-wing parties) for individual i in country c in election t, x_{ict} is a binary explanatory variable of interest (e.g. belonging to top 10% educated voters), and C_{ict} is a vector of controls.

In the absence of controls, the coefficient β simply equals the difference between the share of top 10% educated voters voting for left-wing parties and the share of other voters (bottom 90% educated voters) voting for left-wing parties:

$$\beta = E(y_{ict} = 1, x_{ict} = 1) - E(y_{ict} = 1, x_{ict} = 0) \tag{9.2}$$

With controls, the interpretation is also straightforward: all things being equal, belonging to the top 10% of educated voters increases one's propensity to vote for left-wing parties by β percentage points. All control variables in our dataset are specified as dummy variables, so that the model is fully saturated and can be estimated by OLS using heteroscedasticity-robust standard errors.

9.2 The Disconnection of Education and Income Cleavages in Western Democracies

This section presents our main results on the evolution of electoral divides related to income and education. Section III.A documents the reversal of the education cleavage and the stability of income divides. Section III.B studies how the fragmentation of party systems and the rise of green and anti-immigration parties has contributed to this transformation.

9.2.1 The Divergence of Income and Education

To document the evolution of the influences of income and education on the vote, we start by relying on a simple indicator: the difference between the share of the 10% most educated voters and the share of the 90% least educated voters voting for social democratic, socialist, communist, and green parties (that is, β in equation 1). We use the same indicator for income, defined as the difference between the share of richest 10% voters and the share of poorest 90% voters voting for social democratic and affiliated parties. These two indicators have the advantage of measuring the evolution of the voting behaviors of two groups of equal size, which makes the estimates more comparable.[13]

[13]As discussed in section II.A, deciles of education are computed using all educational categories available in surveys, which implies that the composition of "top 10% educated voters" changes over

Figure I depicts the average quinquennial evolution of these two indicators, after controls, in the twelve Western democracies for which data is available since the 1960s.[14] As shown in the upper line, highest-educated voters were less likely to vote for social democratic parties than lowest-educated voters by 15 percentage points in the 1960s. This gap has shifted very gradually from being negative to becoming positive, from -10 in the 1970s to -5 in the 1980s, 0 in the 1990s, +5 in the 2000s, and finally +10 in 2016-2020. Higher-educated voters have thus moved from being significantly more right-wing than lower-educated voters to significantly more left-wing, leading to a complete reversal in the educational divide.

The evolution has been dramatically different in the case of income. The bottom line shows that top-income voters have always been less likely to vote for social democratic and affiliated parties than low-income voters. This gap has decreased from -15 in the 1960s to about -10 in the past decade, but it remains negative. High-income voters have thus remained closer to conservative parties than low-income voters over the past fifty years.

Combining these two evolutions, a striking long-run transformation in the structure of political cleavages emerges. In the early postwar decades, the party systems of Western democracies were "class-based," in the sense that social democratic and affiliated parties represented both the low-education and the low-income electorate, whereas conservative and affiliated parties represented both high-education and high-income voters. These party systems have gradually evolved towards what we propose to call "multiconflictual" or "multi-elite" party systems: higher-educated voters now vote for the "left," while high-income voters still vote for the "right."

Note that the two indicators shown in the figure control for all available variables at the micro level (education/income, age, gender, religion, church attendance, rural/urban location, region, race/ethnicity, employment status, and marital status). The evolution of these two indicators without controls displays a larger decline in the influence of income on the vote, from nearly -20 in the 1960s to about -5 in 2016-2020. The main reason is that higher-educated voters have on average higher incomes, so that the reversal of the educational divide has mechanically led to a reduction in the difference between top-income and low-income voters. Nonetheless, what is important for our analysis is that the transition observed is robust to the

time. At the beginning of period, this category is mainly composed of university graduates and voters with secondary education; in the 2010s, it gives more weight to individuals with masters or doctorates. See appendix A for more details.

[14]The corresponding regression coefficients by country and decade are displayed in appendix Tables D1 and D2.

inclusion or exclusion of controls.[15]

The divergence of divides related to income and education is common to nearly all Western democracies, but it has happened at different speeds and with different intensities. Figure II shows that the support of higher-educated voters for social democratic parties was lowest in Norway, Sweden, and Finland between the 1950s and 1970s, three democracies well known for having stronger historical class-based party systems than most Western democracies. The reversal of the education cleavage has not yet been fully completed in these countries, as social democratic parties have managed to keep a non-negligible fraction of the low-income and lower-educated electorate (Martínez-Toledano and Sodano, 2021).

This delay is also common to recent democracies such as Spain or Portugal or late industrialized countries such as Ireland, where left-wing parties continue to be more class-based. Portugal and to a lesser extent Ireland represent two major exceptions in our dataset, where we do not observe a clear tendency towards a reversal of the educational divide. Among several factors, this unique trajectory can be explained by the polarization of mainstream parties and the success of new left-wing parties after the onset of the 2008 financial crisis (Bauluz et al., 2021). In contrast, the gap in left votes between higher-educated voters and lower-educated voters is today highest in countries such as the United States, Switzerland, and Netherlands, due largely to the particular salience of identity-based concerns and the strength of anti-immigration and green movements in the latter two countries (Durrer, Gethin, and Martínez-Toledano, 2021).

Figure III shows that top-income voters have also remained more likely than low-income voters to vote for conservative and affiliated parties in nearly all Western democracies, but with important variations. The influence of income on the vote was largest in Northern European countries, Britain, Australia, and New Zealand in the 1950s and 1960s, consistently with their histories of early industrialization and class polarization. It has declined in these countries since then, although income continues to be negatively associated with support for the left.

Meanwhile, low-income voters have supported less decisively left-wing parties in countries with weak historical class cleavages and crosscutting religious (Italy) or ethnolinguistic (Canada) cleavages (Bauluz et al., 2021; Gethin, 2021). Despite these variations, the tendency of high-income voters to support the right in contemporary Western democracies has proved remarkably resilient over time, pointing to the

[15] See appendix Figure A1. We come back to the influence of other covariates in generating the evolution of the education cleavage in Section V.

persistence of conflicts over economic issues and redistributive policy. The only country where a flattening of the income effect could well be underway is the United States (as well as Italy, due to the recent success of the Five Star Movement among the low-income electorate), where in 2016 and 2020 top 10% earners became for the first time since World War II not significantly less likely to vote for the Democratic Party.

Our findings on the reversal of educational divides and the stability of the income effect are extremely robust to alternative specifications. The pattern observed is virtually identical whether one considers the top 50% of education and income voters, other discrete categories such as primary-educated voters or university graduates, or continuous measures of education and income, before and after controls.[16] It also holds in absolute values, not only in relative terms: between 1948-1960 and 2016-2020, for instance, the share of least educated 50% voters voting for social democratic and affiliated parties declined from about 50% to 40%, while it rose linearly from 25% to almost 50% among the top 10% educated.[17] We also find that our results hold when considering a continuous measure of left-right voting derived from the Comparative Manifesto Project database instead of a binary dependent variable.[18] Finally, we report in the appendix full regression tables on the determinants of the vote for social democratic and affiliated parties by country, as well as simple descriptive statistics on support for these parties by education level and income group in each country.[19] With the exception of the few cases highlighted above, we find a complete reversal of the education effect and a stability of the income effect in nearly all countries, regardless of the indicator considered to measure the influence of these two variables.[20]

[16]See appendix Figures A5 to A20. Continuous measures of income and education are derived as the rank of individuals in the income and education distributions, defined from all available income brackets and education categories available in each survey. If 25% of voters are primary-educated, 50% are secondary-educated, and 25% are tertiary-educated, for instance, then voters belonging to each of these categories are attributed quantile values of 0, 0.25, and 0.75, respectively.

[17]See appendix Figure A29.

[18]See appendix Tables D5 to D8.

[19]Regression results by country are reported in appendix Tables E1 to E21, descriptive statistics by education group in appendix Figures EA1 to EA21, and descriptive statistics by income group in appendix Figures EB1 to EB21.

[20]In some cases, the effect of income is non-linear, especially at the beginning of the period: support for left-wing parties is higher among middle-income groups than at the bottom of the distribution. This is mainly due to the fact that farmers and the self-employed, many of which have low incomes, have always been substantially more likely to vote for conservative parties. However, income remains an only imperfect and partial measure of economic resources. In particular, we find in the case of France (the only country with high-quality wealth data) that the effect of wealth on support for the left is much larger and linear, and has remained more stable in the past decades (see appendix Figures EC1 and EC2).

9.2.2 The Fragmentation of Political Cleavage Structures

The emergence of multi-party systems has come together with a significant reshuffling of political forces in most Western democracies. As shown in Figure IV, traditional socialist and social democratic parties have seen their average vote share across Western democracies decline from about 40% to 34% since the end of World War II, while that received by Christian democratic and conservative parties has decreased from 38% to 30%. Communist parties, who used to gather 7% of the vote in the 1940s, have almost completely disappeared from the political scene. Although immigration issues were already present in political debates in many Western democracies, anti-immigration parties started to grow in the late 1970s and have seen their support increase uninterruptedly since then, reaching on average 11% of votes in the past decade. Green parties made their entry in the political landscape in the 1970s and 1980s and have also progressed steadily, reaching on average 8% of votes in the past decade. Support for social-liberal and liberal parties has remained more stable, even though there are important variations across countries.

Figure V displays the evolution of our previous education (Panel A) and income (Panel B) indicators, decomposed for each of these families of parties from 1948 to 2020. In the 1950s-1960s, both top 10% educated voters and top 10% income voters were significantly less likely to vote for social democratic, socialist, communist, and other left-wing parties and more likely to vote for conservative, Christian democratic, and liberal parties than other voters. By 2016-2020, income continues to clearly distinguish these two groups of parties, but their education gradient has completely reversed. Meanwhile, support for anti-immigration and green parties does not differ significantly across income groups (their income gradient is close to zero), but it does vary substantially across educational categories. This has been a constant fact since these parties started taking on a growing importance in the political space in the 1970s and 1980s. In 2016-2020, top 10% educated voters were more likely to vote for green parties by 5 percentage points and less likely to vote for anti-immigration parties by a comparable amount. In other words, the increasing support for green parties on the left and anti-immigration parties on the right has clearly contributed to the reversal of the education cleavage. This finding goes in line with the large political science literature that has shown education to be an important determinant of support for green and anti-immigration parties in recent years (Abou-Chadi and Hix, 2021; Dolezal, 2010; Rydgren, 2013, 2018).

We should stress, however, that the rise of new parties cannot explain alone the reversal of the education cleavage for at least two reasons. First, this reversal started several decades before most of these parties even existed: as Figure V shows, we can date it back to as early as the 1950s. Second, as also shown in Figure V, there have been major transformations in the structure of the vote for traditional left-wing and right-wing parties too, even in the most recent decades. One way to formally decompose the respective influences of traditional left-wing parties and green parties in generating the reversal of the education cleavage is to compare our main indicator of interest including and excluding green parties from the analysis. We find that the gradient has moved from -19.1 to +8.2 between 1948-1960 and 2016-2020 when including green parties, and from -19.1 to +4.3 when excluding them. In other words, the rise of green parties explains about 15% of the reversal observed during this period, and it explains about half of the positive link between education and support for the left in the most recent years. The same holds when it comes to the increase in support for anti-immigration parties in generating the reversal of the link between education and support for the right: it explains about 14% of the overall shift and 55% of the negative gradient in 2016-2020.[21]

Figure VI provides another perspective on this transformation by representing the income and education gradients of these different families of parties in a two-dimensional space in 1961-1965 (panel A) and 2016-2020 (panel B). In the 1960s, the effects of income and education on the vote were aligned: higher income and higher education were both associated with higher support for conservative and affiliated parties. By 2016-2020, these two variables now have opposite effects: higher income is associated with higher support for conservative parties, while higher education is associated with higher support for social democratic parties. Anti-immigration and green parties differ primarily in their tendency to attract voters belonging to different education groups (they are distant on the x-axis but not on the y-axis).

Figure VII further decomposes this two-dimensional structure of political conflict by country in the last decade, distinguishing between traditional right-wing and left-wing parties in panel A and between anti-immigration and green parties in panel B.[22] The two-dimensional split of the electorate can be seen in nearly all countries in our dataset: social democratic and other left-wing parties systematically make better relative scores among low-income voters, conservative and other right-wing parties among high-income voters, anti-immigration parties among lower-educated voters, and green parties among higher-educated voters.[23]

[21]See appendix Figures A25 (left-wing parties) and A26 (right-wing parties).

[22]The corresponding regression coefficients by country and decade, after controls, are displayed in appendix Tables D3 and D4.

[23]In two countries, Italy and New Zealand, lower-educated voters are not significantly more or less likely to vote for anti-immigration parties. In Italy, this is driven by the fact that support for

Despite these commonalities, however, there are large differences across countries in these two indicators. In particular, while nearly all green parties make better scores among higher-educated voters than among the lower educated, they differ in their tendency to attract low- or high-income voters. Similarly, anti-immigration parties have attracted a particularly high share of the lower-educated vote in several Western democracies in the past decade, but we also observe variations in the income profile of far-right voting. These variations are likely to reflect cross-country differences in political fragmentation and voting systems, which create different incentives for parties of the traditional left and the traditional right to adapt their policy proposals in the face of growing electoral competition from new political movements. To better understand these dynamics and the role of political supply in shaping education and income divides, we now turn to manifesto data.

9.3 The Origins of the Transformation of Political Cleavages: Evidence from Manifesto Data

This section investigates the relationship between the divergence of income and education cleavages and ideological polarization by matching our survey dataset with manifesto data. Section IV.A introduces the Comparative Manifesto Project data and the indicators we consider. Section IV.B presents our results on the link between political supply and demand.

9.3.1 Manifesto Project Data and Methodology

Manifesto Data. To make a first step towards understanding the mechanisms underlying the transformation documented in section III, we match our survey dataset with the Comparative Manifesto Project (CMP: Volkens et al., 2020), a hand-coded historical database on the programmatic supply of political parties. The CMP is the result of a collective effort to collect and code the manifestos published by parties just before general elections. Each manifesto is first divided into "quasi-sentences" conveying a specific claim or policy proposal. These quasi-sentences are then assigned to broad ideological or policy categories using a common coding scheme. The resulting dataset presents itself in the form of items (such as "social justice" or

Fratelli d'Italia (which we classify as an anti-immigration party alongside the Lega) was particularly concentrated among higher-educated voters in the 2018 election. In the case of New Zealand, the only significant anti-immigration party, New Zealand First, receives support mainly from the Māori minority and is often considered to be a centrist party, which may explain why its position on the income-education quadrant differs from that of other anti-immigration parties (Gethin, 2021).

"law and order"), with scores corresponding to the share of quasi-sentences dedicated to a specific issue in a party's manifesto. The CMP is the largest available database on political programs in contemporary democracies at the time of writing, and the only one covering nearly all elections held in our 21 countries of interest since the end of World War II.

Combination of Manifesto and Survey Data. We proceed by matching one by one every single party reported in both the CMP and our dataset. This was possible for a total of 459 parties, allowing us to cover over 90% of votes cast in nearly all elections contained in the survey data. The remaining correspond either to independent candidates, or to small parties for which data was not available in the CMP. To the best of our knowledge, this represents the most comprehensive mapping between political supply and demand ever built in comparative research.

Indicators of Interest. Following the political science literature, we consider two main indicators of political supply proposed by Bakker and Hobolt (2013). The indicators correspond to parties' positions on two axes of political conflict: an "economic-distributive" axis representing divides over economic policy and inequality, and a "sociocultural" axis mapping conflicts over issues such as law and order, the environment, multiculturalism, or immigration.[24]

The economic-distributive indicator is equal to the difference between the percentage of "pro-free-market" statements and "pro-redistribution" statements in a given party's manifesto. Pro-redistribution emphases include, among others, proposals to expand social services, nationalize industries, or enhance social justice. Meanwhile, pro-free-market statements encompass references to the limitation of social services, economic incentives, and free enterprise.

Conversely, the sociocultural indicator is defined as the difference between the percentage of "progressive" emphases and "conservative" emphases. Conservative emphases include categories such as political authority, positive evaluations of traditional morality, or negative attitudes towards multiculturalism. Progressive emphases cover issues related to environmentalism, the protection of minority groups, or favorable mentions of multiculturalism.

Given that manifesto items sum by definition to 100%, both indicators theoretically range from -1 to 1, with 1 representing a case of a party exclusively emphasizing pro-free-market/conservative values, and -1 that of a party exclusively emphasizing pro-redistribution/progressive values. While these measures of political ideology

[24]The manifesto items used to derive these two indicators are reported in appendix Table B1.

remain broad and are not exempt from measurement error, they represent the best data at our disposal to study the link between political supply and demand in the long run.

Let us also stress that by operating this distinction between economic and sociocultural dimensions of political conflict, we are not suggesting that sociocultural divides are purely conflicts over identity or morality that would be fully exempt from material concerns. Immigration, environmental, and cultural policies can have strong distributional implications, for instance by disproportionally affecting low-skilled workers or by mostly benefitting residents of large cities, who tend to concentrate a larger share of the higher-educated electorate. In that respect, the emergence of a secondary dimension of political conflict linked to education should also be understood as incorporating new forms of socioeconomic cleavages.

9.3.2 The Evolution of Ideological Polarization

How has the structure of economic and sociocultural conflicts changed in Western democracies since the end of World War II, and to what extent can this account for the growing disconnection between the influences of income and education on the vote? Figure VIII provides a first answer to this question by displaying the evolution of the average economic-distributive and sociocultural scores of specific families of parties between 1945 and 2020.[25] Indices are normalized by the average score by decade so as to better highlight the dynamics of polarization.

Polarization on economic issues has remained remarkably stable in the past decades. The economic-distributive score of social democratic and socialist parties has remained 9 to 14 points below average, while that of conservative parties has fluctuated between +8 and +11. Green parties, which started gaining electoral significance at the beginning of the 1980s, have held economic positions that are comparable to that of traditional left-wing parties. Anti-immigration parties have moved closer to the average position of conservative parties, after a period of particularly marked emphasis on pro-free-market policies. This is consistent with qualitative accounts on the ideological transformation of far-right movements in Western Europe, from the Freedom Party of Austria (Durrer, Gethin, and Martínez-Toledano, 2021) to the French Rassemblement National (Piketty, 2018) and the True Finns (Martínez-Toledano and Sodano, 2021), which have shifted to defending economic redistributive policies in recent years.

[25]The underlying figures are reported in appendix Table B2. See appendix Figures B2 to B8 for a complete representation of the political space by decade.

Meanwhile, polarization on the sociocultural axis of political conflict has dramatically risen since the 1970s, after a brief period of convergence in the early postwar decades. This polarization has been driven by both old and new parties. Between 1970 and 2020, social democratic and socialist parties increasingly emphasized progressive issues, as their deviation from the mean sociocultural score declined linearly from -0.6 to -5.4, while conservative parties shifted to more conservative positions. Green parties have consistently emphasized progressive issues to much greater extent than other parties since their emergence in the 1980s, with a stable score of about -25. Finally, anti-immigration parties have seen their score on the sociocultural axis surge, from +4 in the 1970s to +20 in the 2010s.

Beyond these two indicators of party ideology, we provide more detailed results on the structure of the manifestos of each of these party families in the appendix.[26] Two key results stand out from these disaggregated figures. First, the conservative turn of both anti-immigration and other right-wing parties has been mainly driven by three items coded in the database: positive emphases of "national way of life" (including appeals to nationalism and patriotism), positive emphases of "law and order" (corresponding to favorable mentions of strict law enforcement and stricter actions against crime), and negative mentions of multiculturalism.[27] Meanwhile, green and other left-wing parties have dedicated a growing share of their manifestos to environmental issues and to positive emphases of an "anti-growth economy" (including calls for a more sustainable development path). Second, we find that left-wing and right-wing parties do continue to differ on many issues on the economic-distributive dimension. In particular, both green and other left-wing parties tend to put greater emphasis on welfare, equality, and social justice, while the manifestos of both anti-immigration and other right-wing parties contain a larger share of sentences promoting a free-market economy and welfare state limitation.

9.3.3 Ideological Polarization and the Transformation of Electoral Divides

The stability of economic-distributive conflicts and the rise of sociocultural divides resonates well with our finding on the stability of the income gradient and the reversal

[26]See appendix Tables B3 to B7.

[27]Consistently with the idea that new ethnoreligious minorities perceive conservative and anti-immigration parties as particularly hostile to their integration, we find that immigrants and Muslim voters have been substantially more likely to vote for social democratic and affiliated parties than other voters in the past decade (see appendix Figures CE1 and CE2). We also find deep and persistent divides between voters belonging to different racial or ethnic groups in countries with available data (see appendix Tables E14, E20, and E21).

of the education cleavage. In particular, if the two phenomena are related, one might expect to observe that (1) parties with more progressive positions attract a relatively higher share of higher-educated voters, (2) this relation should rise over time as sociocultural issues gained prominence, and (3) countries that are more polarized on sociocultural issues should have higher education gradients, thereby accounting for the cross-country variations documented in section III.

Figure IX provides descriptive evidence that the reversal of the education cleavage and the rise of a second dimension of political conflict are tightly associated. The upper line represents the correlation between the education gradient of a given party and the sociocultural index of this party by decade, computed across all parties available in the database. This correlation was close to zero and not statistically significant in the 1960s. It has risen monotonically since then, from 0.1 in the 1970s to 0.3 in the 1990s and finally 0.46 in the past decade. Meanwhile, as represented in the bottom line, the correlation between the income gradient and the position of a given party on the economic-distributive axis has remained very stable and negative over the entire period. In other words, higher-educated voters have gradually converged in supporting parties with progressive positions, while high-income voters continue to vote for parties with pro-free-market positions just as much as they used to in the immediate postwar era.[28]

We also investigate in greater detail how these correlations vary across all items available in the Comparative Manifesto Project database.[29] We find that the transformation documented above is visible in nearly all subcategories. In the 1960s-1970s, the education gradient was not significantly correlated to any of the items composing the sociocultural index. By 2010-2020, it has become strongly negatively correlated to positive emphases of “law and order,” “national way of life,” and “traditional morality,” and to negative mentions of “multiculturalism.” At the same time, it has become strongly positively correlated to positive emphases of “culture,” “anti-growth economy,” “freedom and human rights,” “environmentalism,” and “multiculturalism.” These results suggest that the emergence of a new sociocultural axis of political conflict cannot be narrowed down to a single topic of divergence: it involves conflicting visions and priorities over a complex and diverse set of issues.

Figure X plots the cross-country relation between a simple measure of ideological

[28]This transformation is robust to controlling for the composition of parties’ electorates in terms of other variables, as well as to accounting for country, year, and election fixed effects (see appendix Table B9).

[29]See appendix Table B10, which reports correlation coefficients between all items available in the CMP dataset and our education and income indicators.

polarization, defined as the standard deviation of the sociocultural index across all parties in a given election, and the education gradient in the past decade. The relation between the two indicators is strongly positive: countries in which parties compete more on sociocultural issues also display a greater propensity of higher-educated voters to support social democratic, socialist, green, and affiliated parties. In particular, we see that Portugal and Ireland, which were identified as two exceptions showing no clear trend toward a reversal of the education cleavage, are the two countries where sociocultural polarization is today the lowest.[30] While the small number of countries makes it difficult to precisely identify the evolution of this relationship, we also find that it has grown over time, in line with our party-level analysis.[31]

Results combining data on political supply and demand therefore suggest that the emergence of a new sociocultural axis of political conflict is tightly linked to the reversal of the education cleavage in Western democracies. As parties have progressively come to compete on sociocultural issues, electoral behaviors have become increasingly clustered by education group. This relation holds at the country level, with the divergence between education and income being more pronounced in democracies where parties compete more fiercely on this new dimension of electoral divides.

9.4 Electoral Change in Western Democracies: Alternative Explanations and Other Dimensions of Political Conflict

This section studies alternative explanations and heterogeneity in the reversal of the education cleavage and analyzes other dimensions of political conflict. Section V.A investigates to what extent the reversal of educational divides can be explained by changes in the composition of education groups. Section V.B explores heterogeneity in this reversal in terms of age, gender, religion, and other variables available in our dataset. Section V.C briefly discusses the evolution of other electoral cleavages in Western democracies, independently from education and income.

[30]Notice that the indicator mechanically "overestimates" polarization in highly fragmented party systems such as that of Denmark, while it underestimates it in countries with fewer parties such as the United Kingdom, New Zealand, and the United States. This may explain why these countries have lower levels of sociocultural polarization than one might expect.

[31]See appendix Figure B15, which reproduces Figure 8 at the country level.

9.4.1 Can Compositional Changes Explain the Reversal of Educational Divides?

In previous sections, we have studied the reversal of the education cleavage across all Western democracies, with little consideration for changes in the link between education and the other variables in our dataset. While we have shown that this reversal is robust to accounting for all available controls, it remains unclear to what extent shifts in the composition of education groups could account for some of the transformation. It is well-known, for instance, that women have become both more educated (Parro, 2012; Riphahn and Schwientek, 2015; Vincent-Lancrin, 2008) and more left-wing than men in the past decades (see section V.C below). The realignment of gender divides could thus have contributed to generating the move of higher-educated voters toward social democratic and affiliated parties. Similarly, the secularization of Western societies and the associated increase in the share of non-religious voters, who tend to be more educated, could have facilitated the transformation of the education cleavage.

To investigate the role of these various factors, we conduct two complementary analyses: a comparison of the education gradient before and after controls, and a Kitagawa-Oaxaca-Blinder decomposition of the education cleavage. To derive meaningful comparisons, we restrict the analysis in this section to countries for which we have data since the 1960s and the richest comparable set of covariates (Australia, Denmark, Finland, France, the Netherlands, Norway, New Zealand, Sweden, the United Kingdom, and the United States).

We find that the inclusion of control variables only marginally affects the overall change in the link between education and the vote of the past decades.[32] More precisely, top 10% educated voters were less likely to vote for social democratic and affiliated parties by 21.6 percentage points in the 1960s, while they were more likely to do so by 5.3 points in the 2010s. This represents an overall change in the education gradient of 26.9 percentage points. Adding controls does slightly affect the level of the coefficient, but it does not significantly affect the trend: the education gap after controlling for all available covariates has moved from about -18.8 to 3.6, amounting to a shift of 22.4 percentage points. By this measure, changes in the composition of education groups can only account for about 16% of the transformation of educational divides.

Another, more formal way of evaluating what fraction of the reversal is due to changes

[32] See appendix Table D9.

in the composition of groups is to directly estimate a two-way Kitagawa-Oaxaca-Blinder decomposition of the education gradient (Blinder, 1973; Kitagawa, 1955; Oaxaca, 1973). This allows us to decompose the marginal effect of education into two components: one that can be explained by group differences in predictors (that is, differences in the composition of education groups in terms of age, gender, etc.), and one that remains unexplained. As above, we find that other variables largely fail to account for the reversal of educational divides: the actual coefficient shifts from -22.5 to +10.4 between 1961-65 and 2016-20 (corresponding to a 32.9 points change), while the unexplained component increases from -19.6 to +7.6 (corresponding to a 27.2 points change).[33] This implies that these covariates can only predict 17% of the reversal observed over the period considered.

9.4.2 Heterogeneity in the Reversal of Educational Divides

Although compositional changes only explain between 16 and 17% of the reversal of the educational divide, we find some heterogeneity in the reversal when further decomposing voters into subgroups by different socioeconomic characteristics.

In particular, generational dynamics appear to have played a major role in generating the reversal of the education cleavage. Figure XI decomposes the evolution of the education gradient by cohort of voters born at different decades of the twentieth century. Higher-educated voters have been more likely to vote for social democratic and affiliated parties than lower-educated voters within generations born after the 1940s, while the opposite is true among generations born before World War II. New generations have thus become increasingly divided along educational lines, suggesting that the education cleavage could continue rising in the future, as old generations voting along historical class lines gradually disappear from the political landscape. The reversal of the education cleavage has, however, also taken place within recent cohorts, which points to the role of other factors potentially related to political supply or ideological change, as documented in Section IV.

We also find some heterogeneity in the education gradient across other subgroups of voters.[34] In the 2010s, the educational divide is higher among men than women, among non-religious voters than religious voters, among public sector than private sector employees, and in rural areas than in urban areas. The reversal in the educational divide has also been highest among non-religious voters and among men, although it has occurred within nearly all groups. Overall, this evidence reveals that

[33]This decomposition is represented in appendix Figure A51.
[34]See appendix Table D10.

while there exist interesting heterogeneities, the reversal of the educational divide has been a widespread phenomenon that is not restricted to a particular subgroup of voters.

9.4.3 The Evolution of Other Electoral Cleavages

We conclude this paper by briefly discussing the evolution of other determinants of electoral behaviors. Our main finding is that there has been either a stability or a decline of their impact on vote choices. The major exception is gender, for which we do find a significant reversal, comparable in magnitude to that of the education cleavage.

Generational Cleavages. Young voters have always been more likely to vote for left-wing parties than older cohorts in the majority of Western democracies. However, while there are fluctuations across countries and over time, we do not find any evidence that this cleavage has deepened in recent decades.[35] We also document variations in the profile of the vote for anti-immigration parties by age across Western democracies: the share of votes received by these parties increases with age in Denmark, Italy, Norway, New Zealand, Switzerland, and Sweden, but it clearly decreases with age in Austria, Spain, Finland, and France.[36] These findings put into question the strand of the political science literature that has argued that political change in Western democracies would have a major generational dimension, and that the emergence of populist authoritarian leaders in recent years would have partly represented a "backlash" against social progress among the older generations (see Inglehart, 1977; Norris and Inglehart, 2019). As shown in the previous section, educational divides within recent cohorts, rather than conflicts between generations, seem to represent a more important source of electoral realignment in contemporary democracies.

Rural-urban Cleavages. We also find that rural-urban divides have remained relatively stable in the past seven decades: rural areas continue to be more likely to vote for conservative and affiliated parties by 5 to 15 percentage points in most Western democracies, just as they used to in the 1950s-1960s.[37] Furthermore, the fragmentation of the political space in multi-party systems has been associated with a reshuffling of rural-urban divides within rather than across left-right blocs: support for green parties tends to be concentrated in cities today, just like other left-wing

[35]See appendix Figures CA1 to CA4.
[36]See appendix Figures CA5 to CA7.
[37]See appendix Figure CB1.

parties, while anti-immigration parties generally fare better in rural areas, as is the case of other conservative parties. The stability of the rural-urban cleavage thus rules out this dimension as the primary driver of electoral change since the end of World War II.[38]

Religious Cleavages. Religious divides do not seem to have undergone any clear reversal in the past decades either. In all countries with available data, religious voters have always been much less likely to vote for social democratic and affiliated parties than non-religious voters.[39] This gap has slightly declined in most countries since the 1960s, but it remains decisively negative. Moreover, while green movements often disproportionately attract non-religious individuals, this does not make them very different from other left-wing parties, which have always found greater support among secular voters too. Support for anti-immigration parties appears to vary little across religious groups in most countries, so that their progression in recent decades has contributed to further weakening the religious cleavage.[40]

Gender Cleavages. We also corroborate across all Western democracies a well-known fact (Abendschön and Steinmetz, 2014; Edlund and Pande, 2002; Inglehart and Norris, 2000): women used to be more conservative than men and have gradually become more likely to vote for social democratic and affiliated parties.[41] This transition, as in the case of the education cleavage, has been very gradual and is visible as early as the 1950s. In line with the existing literature, we find that much of the negative gradient of the early postwar decades can be explained by the fact that women used to be more religious than men (Blondel, 1970; Goot and Reid, 1984). In particular, this explains why the gender divide was exceptionally large in Italy in the 1950s, where religious cleavages were historically more pronounced than in most Western democracies. However, the reversal does hold even after controlling for all available variables.[42] Along with education, gender is thus one of the only two variables in our dataset for which a complete reversal of electoral divides seems to have taken place.[43]

[38]The share of votes received by green and anti-immigration parties by rural-urban location is represented in appendix Figures CB2 and CB3, respectively. Notice, however, that a few Western democracies (in particular Australia, Belgium, Britain, and France) seem to have witnessed a significant transformation of center-periphery cleavages in recent years, as left-wing parties have concentrated a growing share of the vote in capital cities (see appendix Figures CB4 to CB8).

[39]See appendix Figure CC1.

[40]See appendix Figures CC5 (green parties) and CC6 (anti-immigration parties).

[41]See appendix Figure CD2.

[42]See appendix Figures CD1 and CD3.

[43]Several explanations have been given to this reversal. In the US and Western Europe, the decline of marriage, the rise of divorce, and the economic fragility of women have been shown to be important drivers behind the emergence of the modern gender gap (Abendschön and Steinmetz,

Other socioeconomic cleavages. Finally, our dataset also makes it possible to study the evolution of the vote by union membership, public-private sector of employment, and home ownership. Union members have always been more likely to vote for social democratic and affiliated parties than non-union members, although this gap has slightly declined in most Western democracies since the 1960s.[44] This is also the case of public sector workers and homeowners, which have remained more supportive of social democratic and affiliated parties than other voters in the past decades.[45]

9.5 Conclusion

The new historical database on political cleavages in 21 Western democracies introduced in this article reveals some important facts. In the early postwar decades, social democratic and affiliated parties represented both the low-education and the low-income electorates, while conservative and affiliated parties represented both high-education and high-income voters. These party systems have gradually evolved towards "multiconflictual" or "multi-elite" party systems in most Western democracies, in which higher-educated voters vote for the "left", whereas high-income voters still vote for the "right."

Results combining our database on political demand with political supply data from the Comparative Manifesto Project suggest that the emergence of a new sociocultural axis of political conflict has been tightly associated with the reversal of the education cleavage in Western democracies. As parties have progressively come to compete on sociocultural issues, electoral behaviors have become increasingly clustered by education group. This transformation has been most pronounced in democracies where parties compete most fiercely on this new dimension of electoral divides.

The divergence of political conflicts related to income and education documented in this paper, two strongly correlated measures of socioeconomic status, could also contribute to explaining why rising income and wealth disparities have not led

2014; Edlund and Pande, 2002). In Northern Europe, the expansion of women's employment in the public sector has also been an important factor behind the increase in the vote for the left among women in recent decades (Knutsen, 2001; we reproduce this result in appendix Figure CD4). Women have also been more attracted by environmental issues, which have spurred women's support for green parties, while anti-immigration parties have generally found greater support among men (Givens, 2004; see appendix Figures CD5 and CD6).

[44] See appendix Figures CF5 (before controls) and CF6 (after controls).

[45] See appendix Figures CF7 (before controls) and CF8 (after controls) for the sectoral cleavage and Figures CF9 (before controls) and CF10 (after controls) for support for left-wing parties among homeowners.

to renewed class conflicts. It might also shed light on the reasons why growing inequalities have not been met by greater redistribution in many countries, as political systems could come to increasingly oppose two coalitions embodying the interests of two kinds of elites.

While multiple lessons have emerged from this new database, we acknowledge the analysis remains insufficient and is not exempted from limitations. First, the indicators of political supply used in this paper and more generally the CMP data capture the tendency of parties to emphasize specific issues and are therefore unable to perfectly measure their position on these issues. Moreover, the policy categories coded in the CMP database unfortunately remain very broad, which precludes us from analyzing in greater detail more specific types of issues such as gender equality, immigration, trade protectionism, or education policy. Addressing these two shortcomings would require going back to the original manifestos and deriving new indicators from text analysis or alternative coding techniques.

Secondly, while our descriptive analysis has provided suggestive evidence that the reversal of the education cleavage and the rise of a new sociocultural axis of political conflict were interrelated phenomena, much remains to be understood when it comes to the mechanisms underlying this transformation. In particular, it remains unclear whether the reversal of educational divides was driven by a change in political supply independently from the structure of collective beliefs, or whether shifting supply was on the contrary driven by changing social attitudes across education groups. While some studies have suggested that social divides between groups have remained stable on a number of issues in the long run (Bertrand and Kamenica, 2020; Evans and Tilley, 2017), which would point to the role of shifts in supply, the data at our disposal does not allow us to disentangle these different channels of causality. A promising avenue for future research lies in establishing more directly the causal impact of political supply on the transformation of political cleavages. This would require identifying quasi-experimental settings in which parties exogenously change position on specific issues or suddenly move to emphasizing new concerns.

Finally, the electoral surveys exploited in this paper rely on samples of a few thousands of voters available since the end of World War II that are sufficient to reveal major trends at the national level, but prevent us from carrying more refined and long-run analyses. Other sources and methods, such as localized election results linked to census data, could be mobilized to broaden the historical perspective and perform more granular analyses.

All of these issues raise important challenges that we hope will contribute to stimulating new research in these multiple directions.

Bibliography

Abendschön, Simone and Stephanie Steinmetz (2014). "The Gender Gap in Voting Revisited: Women's Party Preferences in a European Context". In: *Social Politics* 21.2, pp. 315–344.

Abou-Chadi, Tarik and Simon Hix (2021). "Brahmin Left versus Merchant Right? Education, Class, Multiparty Competition, and Redistribution in Western Europe". In: *The British Journal of Sociology* 72.1, pp. 79–92.

Adams, James and Lawrence Ezrow (2009). "Who Do European Parties Represent? How Western European Parties Represent the Policy Preferences of Opinion Leaders". In: *The Journal of Politics* 71.1, pp. 206–223.

Alesina, Alberto, Reza Baqir, and William Easterly (1999). "Public Goods and Ethnic Divisions". In: *The Quarterly Journal of Economics* 114.4, pp. 1243–1284.

Alesina, Alberto F., Edward L. Glaeser, and Bruce Sacerdote (2001). "Why Doesn't the United States Have a European-Style Welfare State?" In: *Brookings Papers on Economic Activity* 2, pp. 187–278.

Algan, Yann et al. (2017). "The European Trust Crisis and the Rise of Populism". In: *Brookings Papers on Economic Activity*, pp. 309–382.

Autor, David et al. (2020). "Importing Political Polarization? The Electoral Consequences of Rising Trade Exposure". In: *American Economic Review* 110.10, pp. 3139–3183.

Bakker, Ryan and Sara Hobolt (2013). "Measuring Party Positions". In: *Political Choice Matters: Explaining the Strength of Class and Religious Cleavages in Cross-National Perspective*. Ed. by Geoffrey Evans and Nan Dirk De Graaf. Oxford University Press.

Bartels, Larry M. (2017b). "Political Inequality in Affluent Democracies: The Social Welfare Deficit". Working Paper.

Bauluz, Luis et al. (2021). "Historical Political Cleavages and Post-Crisis Transformations in Italy, Spain, Portugal and Ireland, 1953-2020". In: *Political Cleavages and*

Social Inequalities. A Study of 50 Democracies, 1948-2020. Ed. by Amory Gethin, Clara Martínez-Toledano, and Thomas Piketty. Harvard University Press.

Becker, Sascha O. and Thiemo Fetzer (2016). "Does Migration Cause Extreme Voting?" CAGE Working Paper.

– (2018). "Why an EU Referendum? Why in 2016?" Warwick Economics Research Paper Series.

Becker, Sascha O., Thiemo Fetzer, and Dennis Novy (2017). "Who Voted for Brexit? A Comprehensive District-Level Analysis". In: *Economic Policy* 32.92, pp. 601–650.

Bertrand, Marianne and Emir Kamenica (2020). "Coming apart? Cultural distances in the United States over time". Working Paper.

Bertrand, Marianne et al. (2020). "Tax-exempt Lobbying: Corporate Philanthropy as a Tool for Political Influence". In: *American Economic Review* 110.7, pp. 2065–2102.

Blinder, Alan S. (1973). "Wage Discrimination: Reduced Form and Structural Estimates". In: *Journal of Human Resources* 8.4, pp. 436–455.

Blondel, Jean (1970). *Votes, Parties and Leaders*. Penguin.

Bonica, Adam et al. (2013). "Why Hasn't Democracy Slowed Rising Inequality?" In: *Journal of Economic Perspectives* 27.3, pp. 103–124.

Bonomi, Giampaolo, Nicola Gennaioli, and Guido Tabellini (2021). "Identity, Beliefs, and Political Conflict". In: *The Quarterly Journal of Economics* 136.4, pp. 2371–2411.

Bornschier, Simon (2010a). *Cleavage Politics and the Populist Right*. Temple University Press.

– (2010b). "The New Cultural Divide and the Two-Dimensional Political Space in Western Europe". In: *West European Politics* 33.3, pp. 419–444.

Bosancianu, Constantin M. (2017). "True OECD Voter".

Bovens, Mark and Anchrit Wille (2012). "Education: The Contours of a New Cleavage? Comparing 23 European Countries". In: *Stabiliteit en verandering in Europa*. Ed. by Kees Aarts and Marcel Wittenberg. DANS.

Budge, Ian and Michael J. Laver (1992). *Party Policy and Government Coalitions*. St. Martins Press.

Cagé, Julia (2020). *The Price of Democracy. How Money Shapes Politics and What to Do about It*. Harvard University Press.

Colantone, Italo and Piero Stanig (2018a). "Global Competition and Brexit". In: *American Political Science Review* 112.2, pp. 201–218.

– (2018b). "The Trade Origins of Economic Nationalism: Import Competition and Voting Behavior in Western Europe". In: *American Journal of Political Science* 62.4, pp. 936–953.

Dalton, Russell J. (2018). *Political Realignment: Economics, Culture, and Electoral Change*. Oxford University Press.

Dehdari, Sirus H. (2021). "Economic Distress and Support for Radical Right Parties - Evidence from Sweden". In: *Comparative Political Studies*.

Dolezal, Martin (2010). "Exploring the Stabilization of a Political Force: The Social and Attitudinal Basis of Green Parties in the Age of Globalization". In: *West European Politics* 33.3, pp. 534–552.

Duch, Raymond M. and Michaell A. Taylor (1993). "Postmaterialism and the Economic Condition". In: *American Journal of Political Science* 37.3, pp. 747–779.

Durrer, Carmen, Amory Gethin, and Clara Martínez-Toledano (2021). "Party System Transformation and the Structure of Political Cleavages in Austria, Belgium, the Netherlands and Switzerland, 1967-2019". In: *Political Cleavages and Social Inequalities. A Study of 50 Democracies, 1948-2020*. Ed. by Amory Gethin, Clara Martínez-Toledano, and Thomas Piketty. Harvard University Press.

Dustmann, Christian, Kristine Vasiljeva, and Anna Piil Damm (2019). "Refugee Migration and Electoral Outcomes". In: *The Review of Economic Studies* 86.5, pp. 2035–2091.

Edlund, L. and R. Pande (2002). "Why have women become left-wing? The political gender gap and the decline in marriage". In: *Quarterly Journal of Economics* 117.3, pp. 917–961.

Elff, Martin (2007). "Social Structure and Electoral Behavior in Comparative Perspective: The Decline of Social Cleavages in Western Europe Revisited". In: *Perspectives on Politics* 5.2, pp. 277–294.

– (2009). "Social Divisions, Party Positions, and Electoral Behaviour". In: *Electoral Studies* 28.2, pp. 297–308.

Enke, Benjamin (2020). "Moral Values and Voting". In: *Journal of Political Economy* 128.10, pp. 3679–3729.

Evans, Geoffrey and Nan Dirk De Graaf (2012). *Political Choice Matters: Explaining the Strength of Class and Religious Cleavages in Cross-National Perspective*. Oxford University Press.

– (2013). *Political Choice Matters: Explaining the Strength of Class and Religious Cleavages in Cross-National Perspective*. Oxford University Press.

Evans, Geoffrey and James Tilley (2012). "The Depoliticization of Inequality and Redistribution: Explaining the Decline of Class Voting". In: *The Journal of Politics* 74.4, pp. 963–976.

– (2017). *The New Politics of Class: The Political Exclusion of the British Working Class*. Oxford University Press.

Fetzer, Thiemo (2019). "Did Austerity Cause Brexit?" In: *American Economic Review* 109.11, pp. 3849–3886.

Ford, Robert and Will Jennings (2020). "The Changing Cleavage Politics of Western Europe". In: *Annual Review of Political Science* 23.1, pp. 295–314.

Franklin, Mark N., Thomas T. Mackie, and Henry et al. Valen (1992). *Electoral Change: Responses to Evolving Attitudinal Structures in Western Countries*. Cambridge University Press.

Funke, Manuel, Moritz Schularick, and Christoph Trebesch (2016). "Going to Extremes: Politics after Financial Crises, 1870–2014". In: *European Economic Review* 88, pp. 227–260.

Gethin, Amory (2021). "Political Cleavages, Class Structures and the Politics of Old and New Minorities in Australia, Canada, and New Zealand 1963-2019". In: *Political Cleavages and Social Inequalities. A Study of 50 Democracies, 1948-2020*. Ed. by Amory Gethin, Clara Martínez-Toledano, and Thomas Piketty. Harvard University Press.

Gethin, Amory, Clara Martínez-Toledano, and Thomas Piketty (2021). *Political Cleavages and Social Inequalities. A Study of 50 Democracies, 1948-2020*. Harvard University Press.

Gilens, Martin (2012). *Affluence and Influence: Economic Inequality and Political Power in America*. Princeton University Press.

Gilens, Martin and Benjamin I. Page (2014). "Testing Theories of American Politics: Elites, Interest Groups, and Average Citizens". In: *Perspectives on Politics* 12.3, pp. 564–581.

Givens, Terri E. (2004). "The Radical Right Gender Gap". In: *Comparative Political Studies* 37.1, pp. 30–54.

Goot, Murray and Elizabeth Reid (1984). "Women: If Not Apolitical, Then Conservative". In: *Women and the Public Sphere*. Ed. by Janet Siltanen and Michelle Stanworth. Hutchinson.

Guiso, Luigi et al. (2020). "Economic Insecurity and the Demand of Populism in Europe". TBD.

Halla, Martin, Alexander F. Wagner, and Josef Zweimüller (2017). “Immigration and Voting for the Far Right”. In: *Journal of the European Economic Association* 15.6, pp. 1341–1385.

Hooghe, Liesbet, Gary Marks, and Carole J. Wilson (2002). “Does Left/Right Structure Party Positions on European Integration?” In: *Comparative Political Studies* 35.8, pp. 965–989.

Inglehart, Ronald (1977). *The Silent Revolution: Changing Values and Political Styles Among Western Publics.* Princeton University Press.

Inglehart, Ronald and Pippa Norris (2000). “The Developmental Theory of the Gender Gap: Women’s and Men’s Voting Behavior in Global Perspective”. In: *International Political Science Review* 21.4, pp. 441–463.

Jansen, Giedo, Nan Dirk De Graaf, and Ariana Need (2011). “Class Voting, Social Changes and Political Changes in the Netherlands”. In: *Electoral Studies* 40.3, pp. 510–524.

– (2012). “Explaining the Breakdown of the Religion–Vote Relationship in The Netherlands, 1971–2006”. In: *West European Politics* 35.4, pp. 756–783.

Jansen, Giedo, Geoffrey Evans, and Nan Dirk de Graaf (2013). “Class voting and Left-Right party positions: A comparative study of 15 Western democracies, 1960-2005”. In: *Soc Sci Res* 42.2, pp. 376–400.

Kitagawa, Evelyn M. (1955). “Components of a Difference Between Two Rates”. In: *Journal of the American Statistical Association* 50.272, pp. 1168–1194.

Kitschelt, Herbert (1994). *The Transformation of European Social Democracy.* Cambridge University Press.

Kitschelt, Herbert P. and Philipp Rehm (2019). “Secular Partisan Realignment in the United States: The Socioeconomic Reconfiguration of White Partisan Support since the New Deal Era”. In: *Politics and Society* 47.3, 425–479.

Knutsen, Oddbjørn (2001). “Social Class, Sector Employment, and Gender as Party Cleavages in the Scandinavian Countries: A Comparative Longitudinal Study, 1970–95”. In: *Scandinavian Political Studies* 24.4, pp. 311–350.

Kosse, Fabian and Thomas Piketty (2021). “Electoral Cleavages and Socioeconomic Inequality in Germany 1949-2017”. In: *Political Cleavages and Social Inequalities. A Study of 50 Democracies, 1948-2020.* Ed. by Amory Gethin, Clara Martínez-Toledano, and Thomas Piketty. Harvard University Press.

Kriesi, Hanspeter et al. (2008). *West European Politics in the Age of Globalization.* Cambridge University Press.

Kuhner, Timothy K. (2014). *Capitalism v. Democracy: Money in Politics and the Free Market Constitution.* Stanford Law Books.

Langsæther, Peter Egge and Rune Stubager (2019). "Old wine in new bottles? Reassessing the Effects of Globalisation on Political Preferences in Western Europe". In: *European Journal of Political Science Research* 58.4, pp. 1213–1233.

Liberini, Federica et al. (2019). "Was Brexit Triggered by the Old and Unhappy? Or by Financial Feelings?" In: *Journal of Economic Behavior and Organization* 161, pp. 287–302.

Malgouyres, Clément (2017). "Trade Shocks and Far-Right Voting: Evidence from French Presidential Elections". EUI Working Paper.

Martínez-Toledano, Clara and Alice Sodano (2021). "Changing Party Systems, Socioeconomic Cleavages, and Nationalism in Northern Europe, 1956-2017". In: *Political Cleavages and Social Inequalities. A Study of 50 Democracies, 1948-2020*. Ed. by Amory Gethin, Clara Martínez-Toledano, and Thomas Piketty. Harvard University Press.

Norris, Pippa and Ronald Inglehart (2019). *Cultural Backlash. Trump, Brexit, and Authoritarian Populism*. Cambridge University Press.

Oaxaca, Ronald (1973). "Male-Female Wage Differentials in Urban Labor Markets". In: *International Economic Review* 14.3, pp. 693–709.

Parro, Francisco (2012). "International Evidence on the Gender Gap in Education over the Past Six Decades: A Puzzle and an Answer to It". In: *Journal of Human Capital* 6.2, pp. 150–185.

Pereira, Miguel M. (2021). "Understanding and Reducing Biases in Elite Beliefs about the Electorate". In: *American Political Science Review*.

Piketty, Thomas (1995). "Social Mobility and Redistributive Politics". In: *Quarterly Journal of Economics* 110.3, pp. 551–584.

– (2018). "Brahmin Left vs Merchant Right: Rising Inequality and the Changing Structure of Political Conflict (Evidence from France, Britain and the US, 1948-2017)". WID.world Working Paper.

– (2021). "Brahmin Left versus Merchant Right: Rising Inequality and the Changing Structure of Political Conflict in France, Britain, and the US, 1948-2020". In: *Political Cleavages and Social Inequalities. A Study of 50 Democracies, 1948-2020*. Ed. by Amory Gethin, Clara Martínez-Toledano, and Thomas Piketty. Harvard University Press.

Rennwald, Line and Geoffrey Evans (2014). "When Supply Creates Demand: Social Democratic Party Strategies and the Evolution of Class Voting". In: *West European Politics* 37.5, pp. 1108–1135.

Riphahn, Regina T. and Caroline Schwientek (2015). “What Drives the Reversal of the Gender Education Gap? Evidence from Germany”. In: *Applied Economics* 47.53, pp. 5748–5775.

Roemer, John E. (1998). “Why the Poor do not Expropriate the Rich: An Old Argument in New Garb”. In: *Journal of Public Economics* 70.3, pp. 399–424.

Roemer, John E. et al. (2007). *Racism, Xenophobia, and Distribution. Multi-Issue Politics in Advanced Democracies*. Harvard University Press.

Rydgren, Jens, ed. (2013). *Class Politics and the Radical Right*. Routledge.

– ed. (2018). *The Oxford Handbook of the Radical Right*. Oxford University Press.

Stubager, Rune (2008). “Education Effects on Authoritarian–libertarian Values: A Question of Socialization”. In: *The British Journal of Sociology* 59.2, pp. 327–350.

– (2010). “The Development of the Education Cleavage: Denmark as a Critical Case”. In: *West European Politics* 33.3, pp. 505–533.

– (2013). “The Changing Basis of Party Competition: Education, Authoritarian–Libertarian Values and Voting”. In: *Government and Opposition* 48.3, pp. 372–397.

Tabellini, Marco (2020). “Gifts of the Immigrants, Woes of the Natives: Lessons from the Age of Mass Migration”. In: *The Review of Economic Studies* 87.1, pp. 454–486.

Thomassen, Jacques, ed. (2005). *The European Voter: A Comparative Study Of Modern Democracies*. Oxford University Press.

Vincent-Lancrin, Stéphan (2008). “The Reversal of Gender Inequalities in Higher Education: An On-going Trend”. In: *OECD, Higher Education to 2030, Volume 1, Demography*. OECD, pp. 265–298.

Volkens, Andrea et al. (2020). *The Manifesto Data Collection. Manifesto Project (MRG/CMP/MARPOR), Version 2020b*. https://manifesto-project.wzb.eu/datasets.

Waal, Jeroen Van der, Peter Achterberg, and Dick Houtman (2007). “Class Is Not Dead—It Has Been Buried Alive: Class Voting and Cultural Voting in Postwar Western Societies (1956–1990)”. In: *Politics and Society* 35.3, pp. 403–426.

Young, Michael (1958). *The Rise of the Meritocracy, 1870-2033: An Essay on Education and Inequality*. Penguin.

Önudottir, E., H. Schmitt, and F. Vegetti et al. (2017). “The New European Voter database”. GESIS Data Archive.

Table I - A New Dataset on Political Cleavages in Western Democracies, 1948-2020

	Time period	Elections	Main data source	Data quality	Avg. sample size
Australia	1963-2019	18	Australian Election Studies	High	2382
Austria	1971-2017	10	Eurobarometers, European Social Survey	Medium	3831
Belgium	1971-2014	14	Belgian National Election Study	High	4817
Canada	1963-2019	17	Canadian Election Studies	High	3302
Denmark	1960-2015	21	Danish Election Studies	High	2819
Finland	1972-2015	11	Finnish Voter Barometers	High	2452
France	1956-2017	17	French Election Studies	High	3208
Germany	1949-2017	19	German Federal Election Studies	High	2782
Iceland	1978-2017	12	Icelandic National Election Studies	High	1488
Ireland	1973-2020	13	Eurobarometers, European Social Survey	Medium	7115
Italy	1953-2018	14	Italian National Election Studies	High	2147
Luxembourg	1974-2018	9	Eurobarometers, European Election Studies	Low	3890
Netherlands	1967-2017	15	Dutch Parliamentary Election Studies	High	2068
New Zealand	1972-2017	16	New Zealand Election Studies	High	2555
Norway	1957-2017	15	Norwegian Election Studies	High	1964
Portugal	1983-2019	10	Portuguese Election Studies	High	1822
Spain	1979-2019	14	CIS Election Surveys	High	8996
Sweden	1956-2014	19	Swedish National Election Studies	High	3088
Switzerland	1967-2019	14	Swiss Election Studies	High	3328
United Kingdom	1955-2017	16	British Election Studies	High	5262
United States	1948-2020	18	American National Election Studies	High	2179

Source: authors' elaboration.

Note: the table presents, for each country, the time coverage of the dataset, the number of elections covered, the main data source used, the quality of electoral surveys, and the average sample size of these surveys.

Figure I - The Disconnection of Income and Education Cleavages in Western democracies

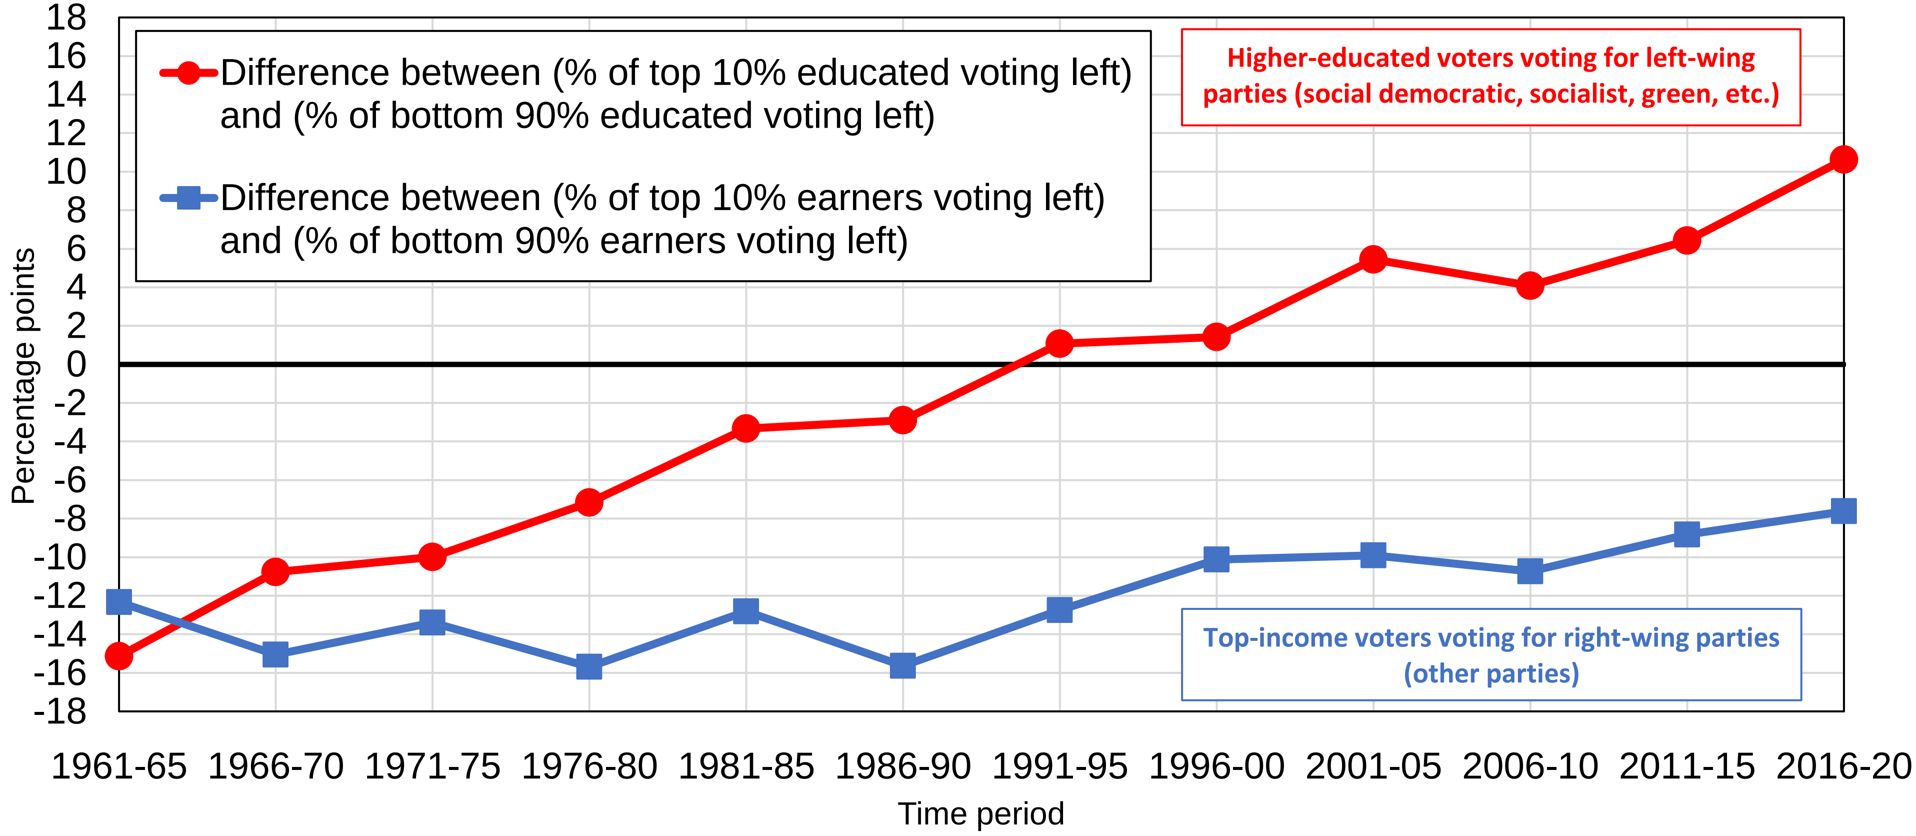

Source: authors' computations using the World Political Cleavages and Inequality Database.

Note: in the 1960s, both higher-educated and high-income voters were less likely to vote for left-wing (social democratic / socialist / communist / green / other left-wing) parties than lower-educated and low-income voters by more than 10 percentage points. The left vote has gradually become associated with higher education voters, giving rising to a complete divergence of the effects of income and education on the vote. Figures correspond to five-year averages for Australia, Britain, Canada, Denmark, France, Germany, Italy, the Netherlands, Norway, Sweden, Switzerland, and the US. Estimates control for income/education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure II - The reversal of educational divides in Western democracies
Panel A. English-speaking and Northern European countries

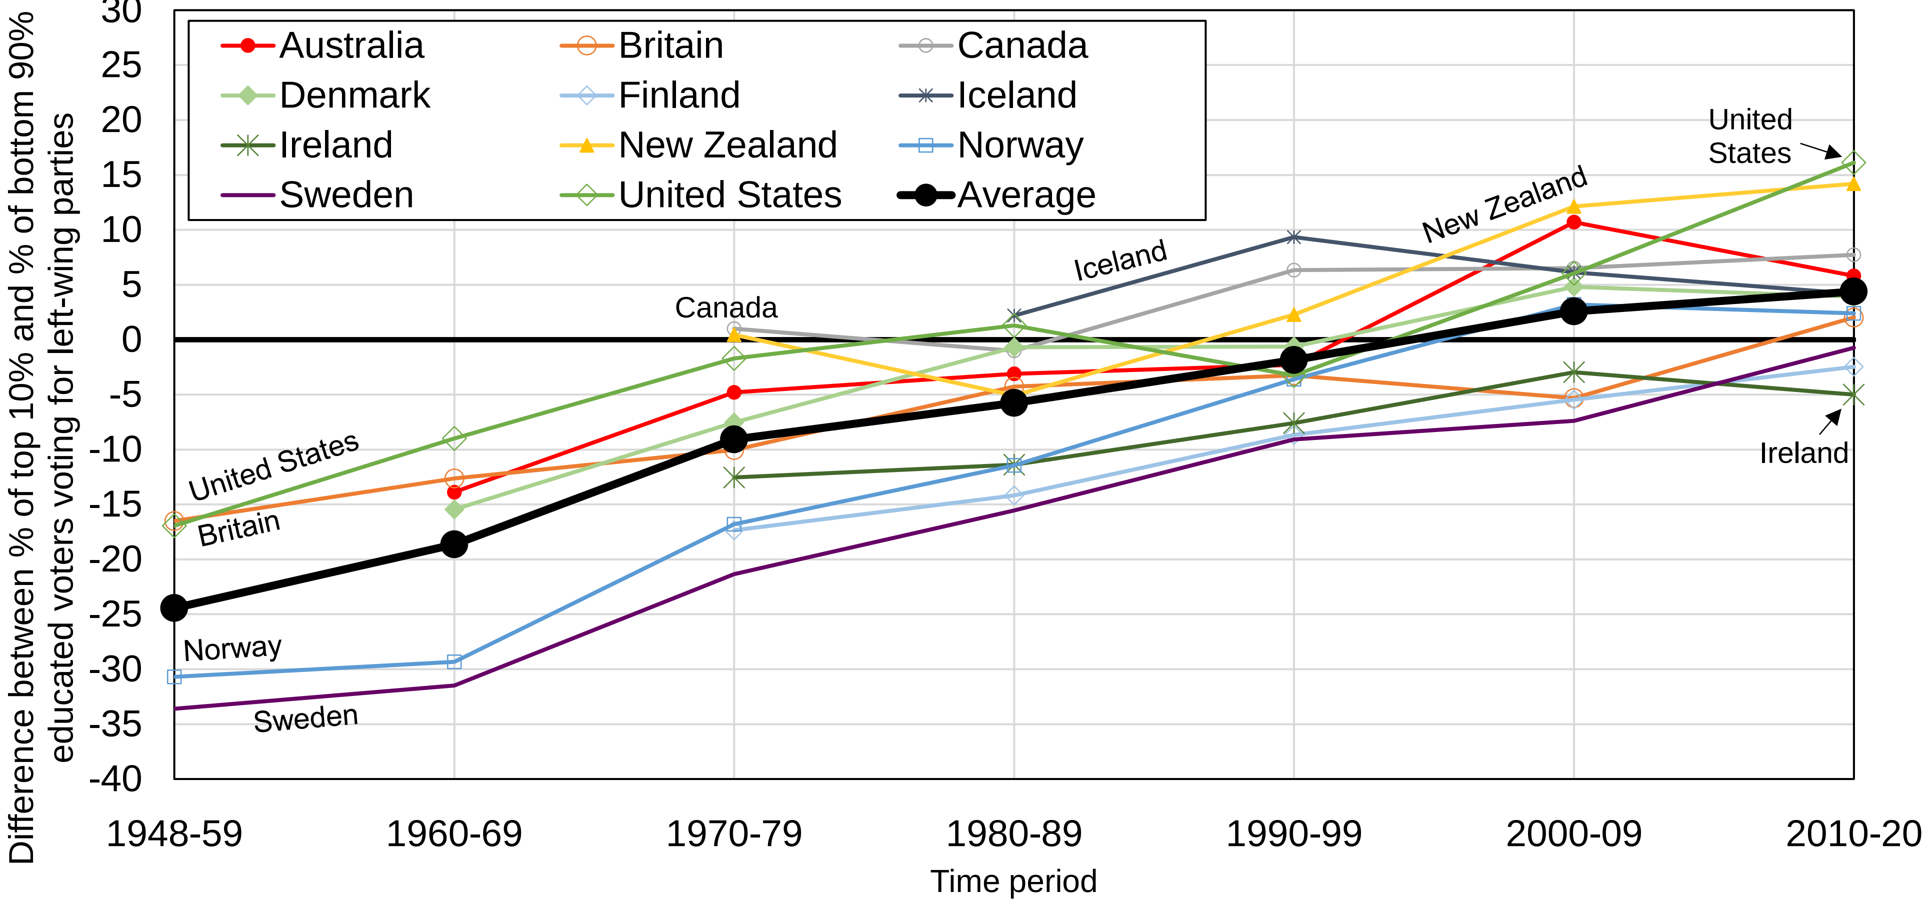

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of higher-educated (top 10%) and lower-educated (bottom 90%) voters voting for social democratic / socialist / communist / green / other left-wing parties in English-speaking and Northern European countries. In nearly all countries, higher-educated voters used to be significantly more likely to vote for conservative parties and have gradually become more likely to vote for these parties. Estimates control for income, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure II - The reversal of educational divides in Western democracies
Panel B. Continental and Southern European countries

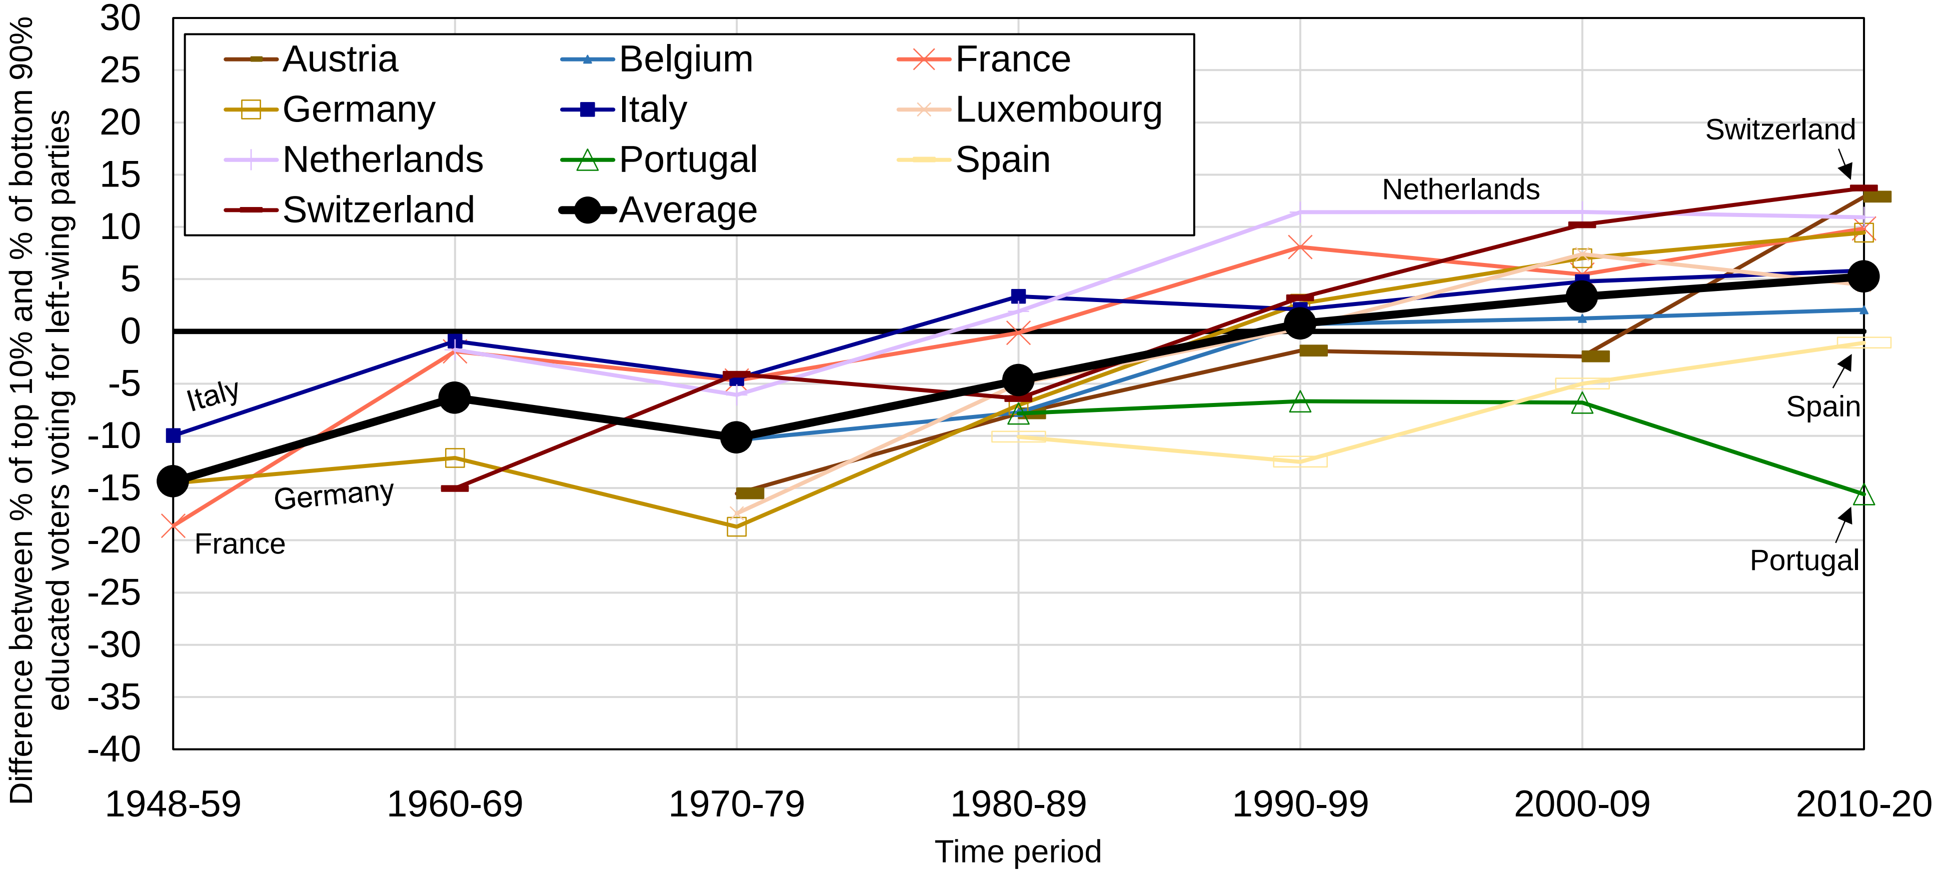

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of higher-educated (top 10%) and lower-educated (bottom 90%) voters voting for social democratic / socialist / communist / green / other left-wing parties in Continental and Southern European countries. In nearly all countries, higher-educated voters used to be significantly more likely to vote for conservative parties and have gradually become more likely to vote for these parties. Estimates control for income, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure III - The stability/decline of income divides in Western democracies
Panel A. English-speaking and Northern European countries

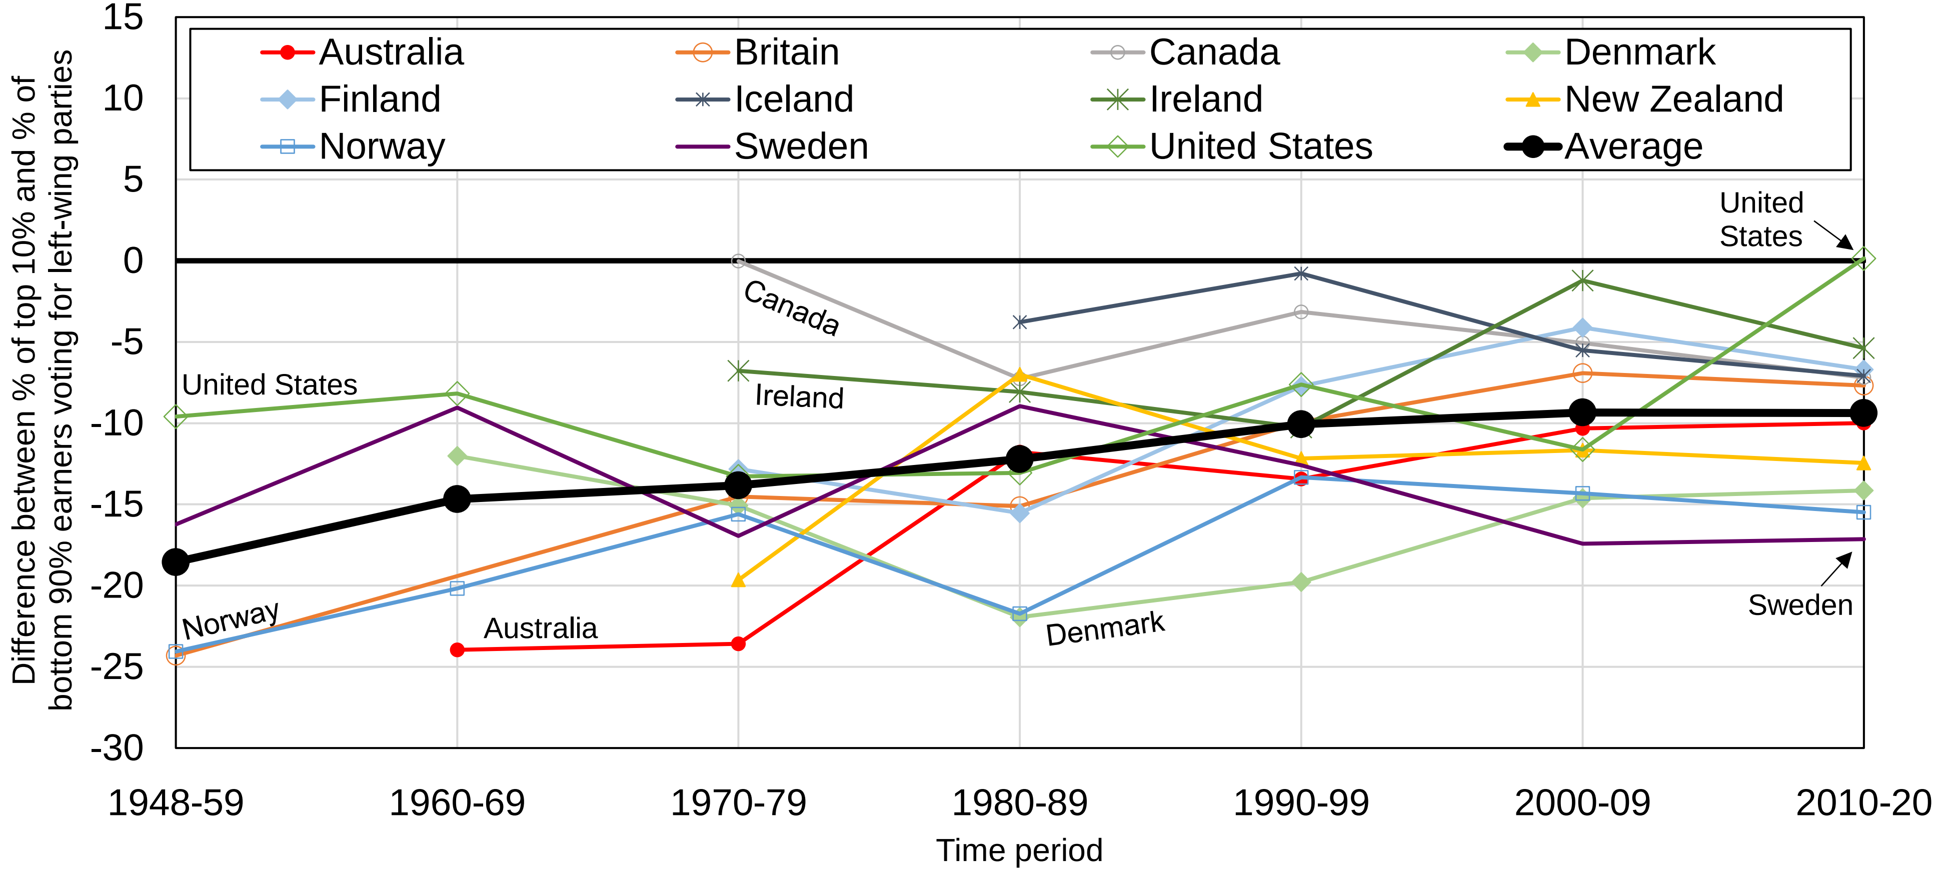

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of high-income (top 10%) and low-income (bottom 90%) voters voting for social democratic / socialist / communist / green / other left-wing parties in English-speaking and Northern European countries. In all countries, top-income voters have remained significantly less likely to vote for these parties than low-income voters. Estimates control for education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure III - The stability/decline of income divides in Western democracies
Panel B. Continental and Southern European countries

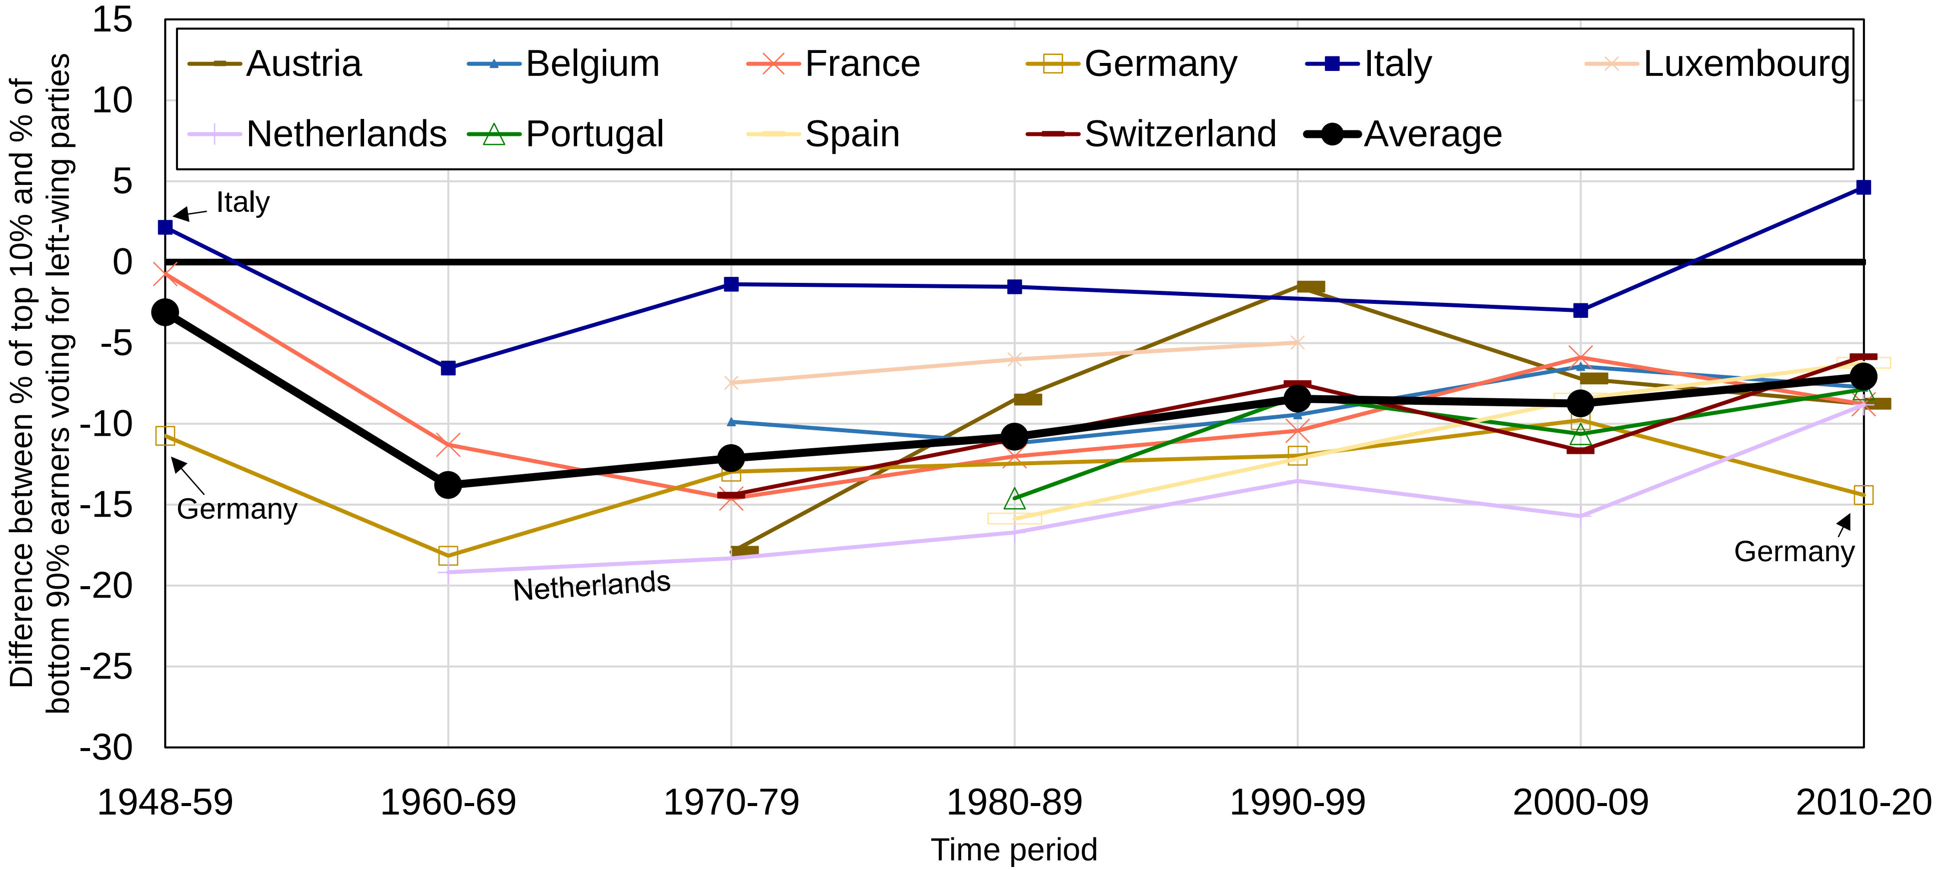

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of high-income (top 10%) and low-income (bottom 90%) voters voting for social democratic / socialist / communist / green / other left-wing parties in Continental and Southern European countries. In all countries, top-income voters have remained significantly less likely to vote for these parties than low-income voters. Estimates control for education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure IV - The transformation of Western party systems, 1945-2020

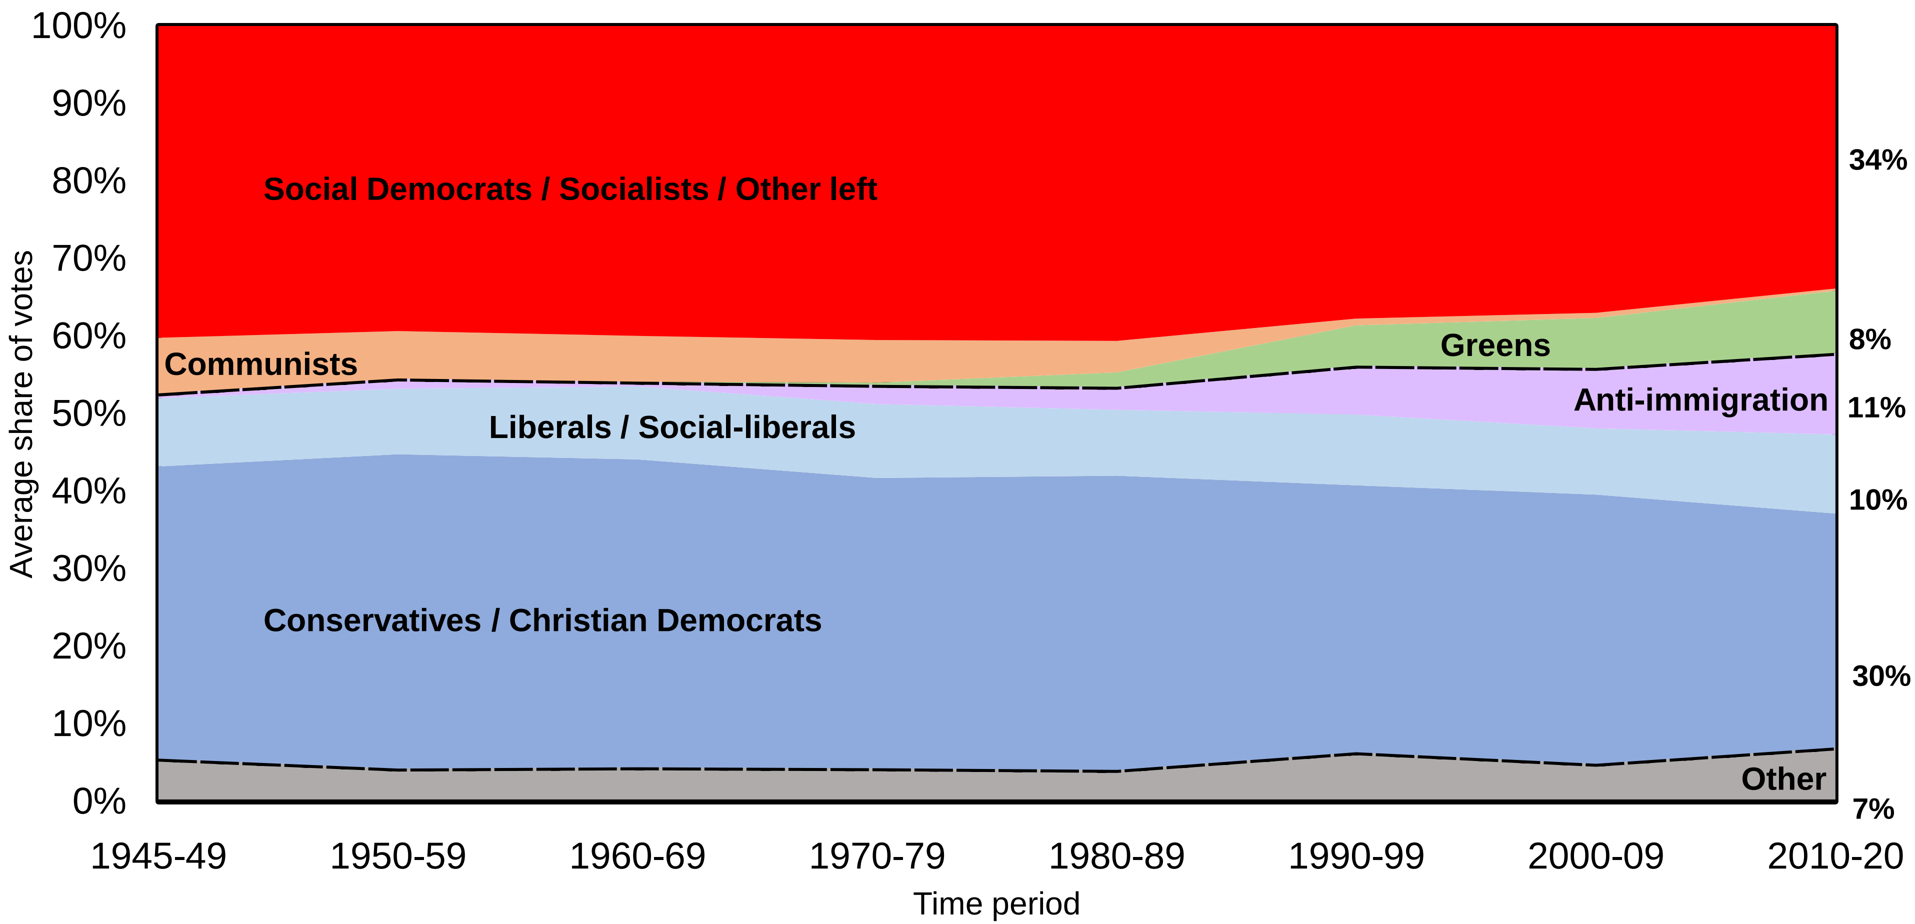

Source: authors' computations using official election results data.

Note: the figure represents the average share of votes received by selected families of political parties in Western democracies between the 1940s and the 2010s. Communist parties saw their average scores collapse from 7% to less than 0.5%, while green and anti-immigration parties rose until reaching average vote shares of 8% and 11%, respectively. Decennial averages over all Western democracies except Spain and Portugal (no democratic elections before 1970s) and the United States and the United Kingdom (two-party systems). The dashed lines delimit the categorization of parties considered in the main specification (social democrats and affiliated, conservatives and affiliated, and other parties).

Figure V - Decomposition by party family
Panel A. The reversal of educational divides in Western democracies

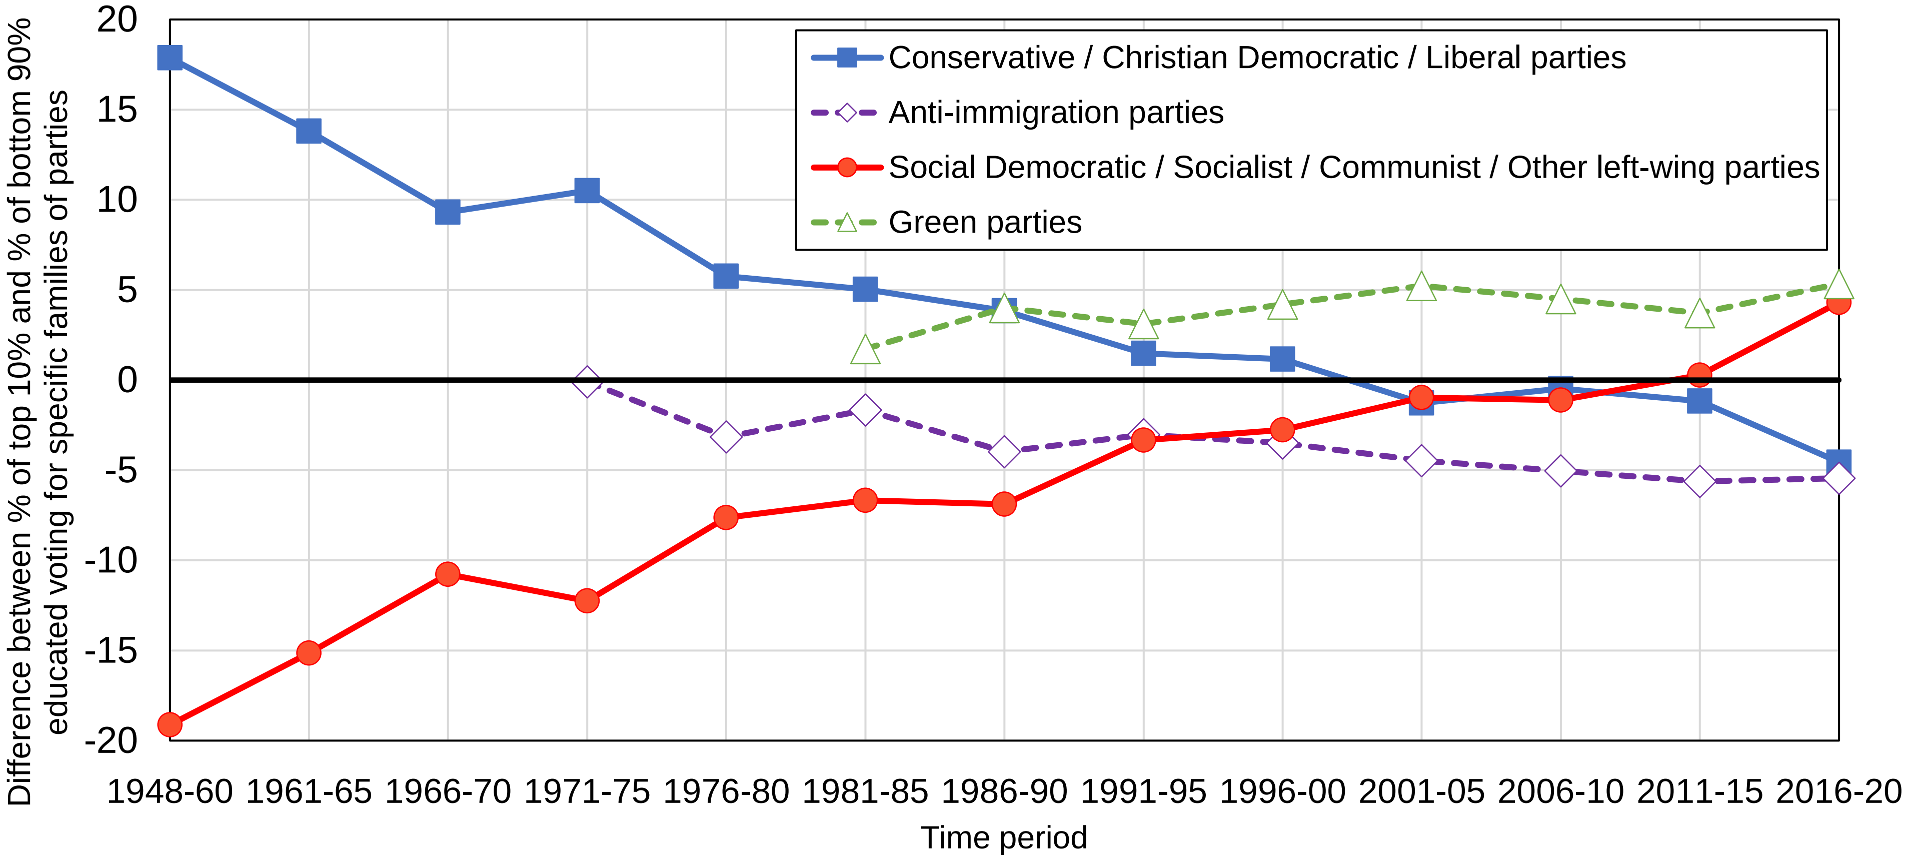

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of top 10% educated voters and the share of bottom 90% educated voters voting for specific families of parties. Figures correspond to five-year averages over all countries available for a given time period (unbalanced panel of all 21 Western democracies). The estimates are presented after controlling for income, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure V - Decomposition by party family
Panel B. The decline/stability of income divides in Western democracies

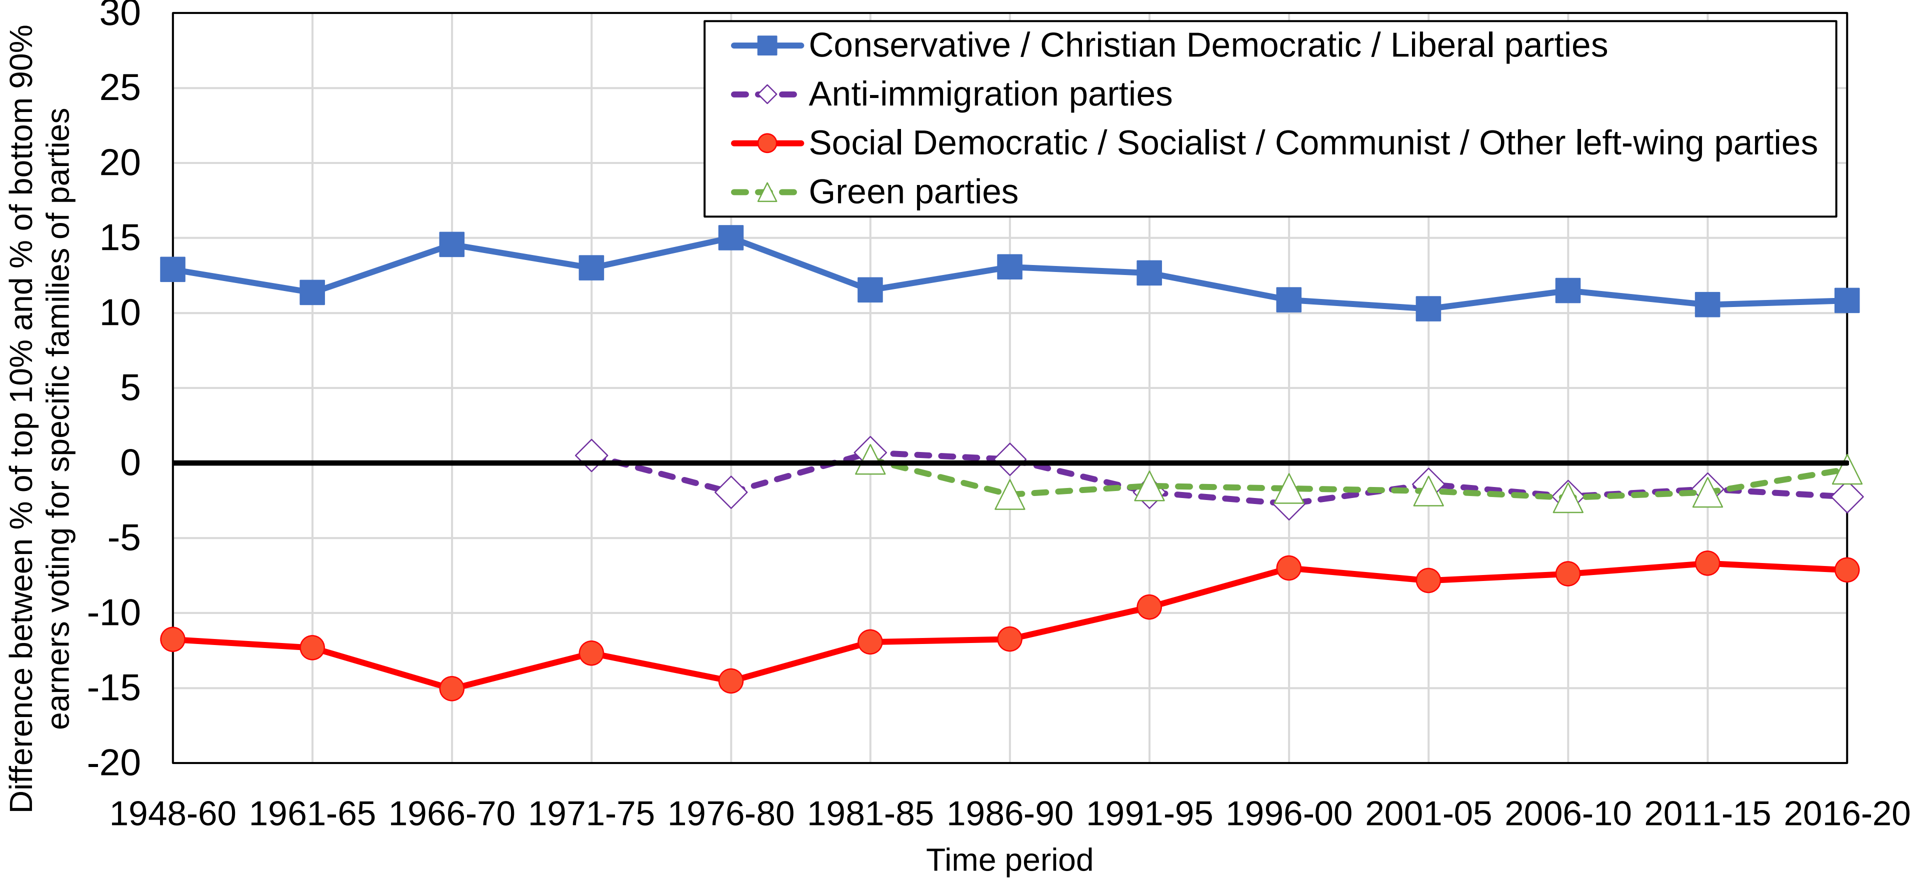

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of top 10% income voters and the share of bottom 90% income voters voting for specific families of parties. Figures correspond to five-year averages over all countries available for a given time period (unbalanced panel of all 21 Western democracies). The estimates are presented after controlling for education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure VI - The fragmentation of political cleavage structures.
Panel A. 1961-1965

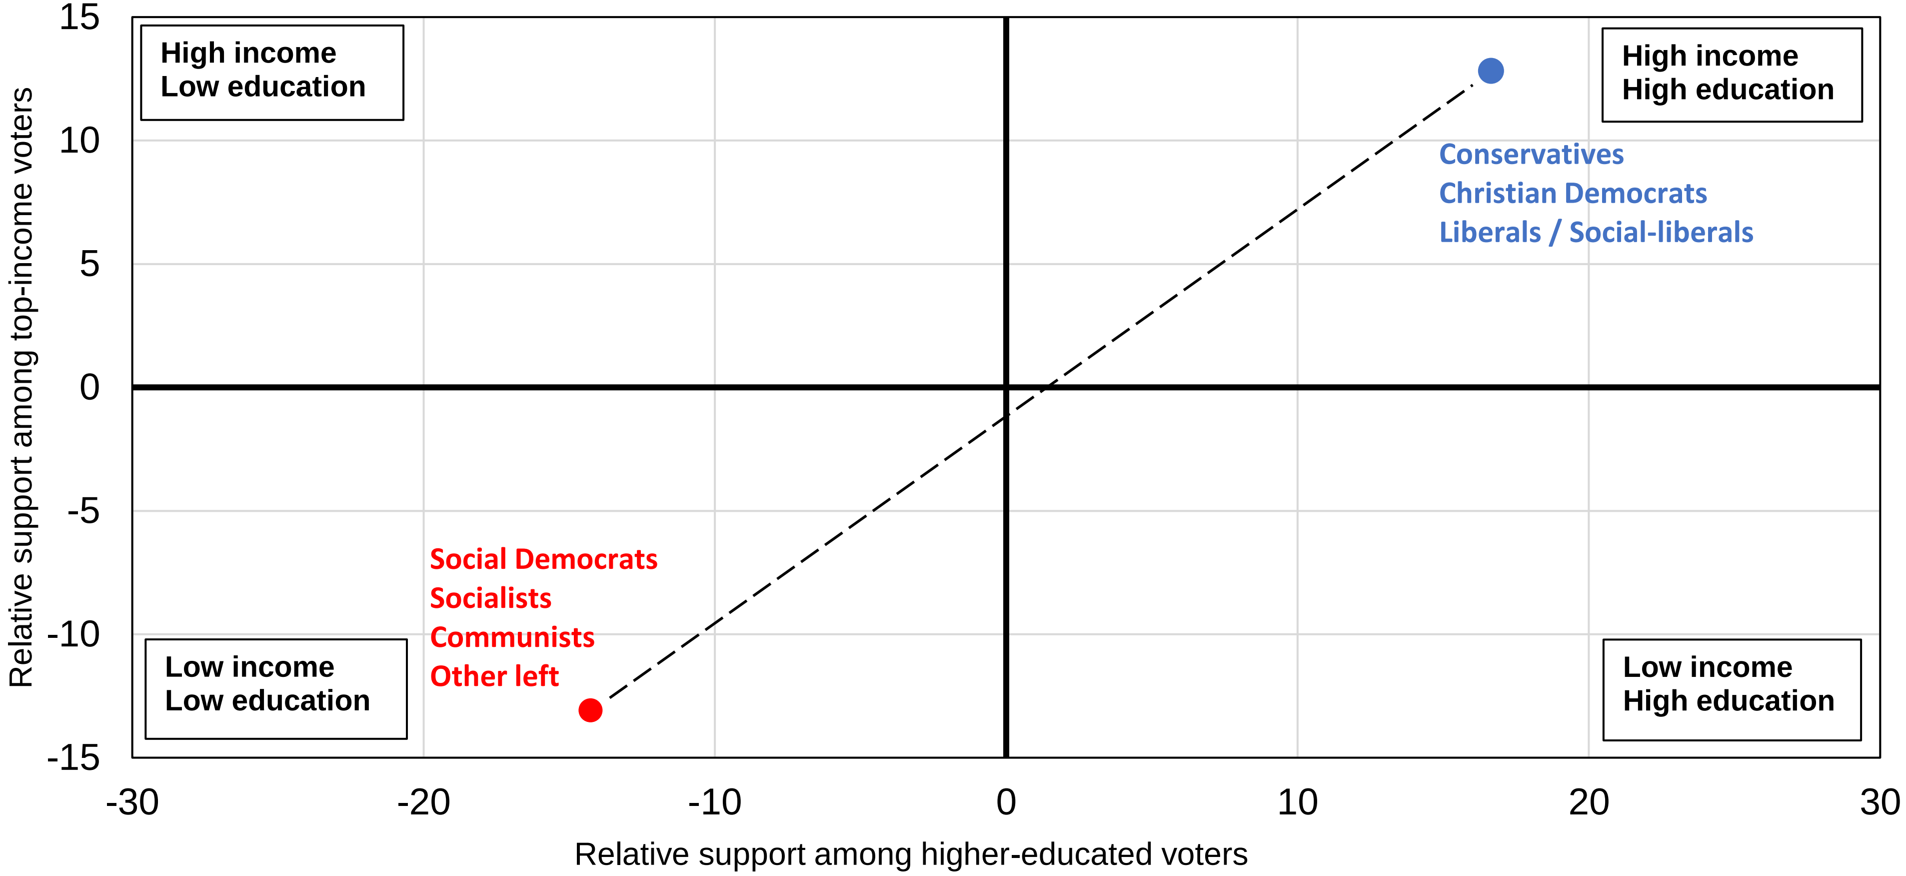

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of high-income (top 10%) and low-income (bottom 90%) voters voting for selected groups of parties on the y-axis, and the same difference between higher-educated (top 10%) and lower-educated (bottom 90%) voters on the x-axis. In the 1960s, social democratic, socialist, and communist parties were supported by both low-income and lower-educated voters, while conservative, Christian, and liberal parties were supported by both high-income and higher-educated voters. Averages over all Western democracies. Estimates control for income/education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure VI - The fragmentation of political cleavage structures.
Panel B. 2016-2020

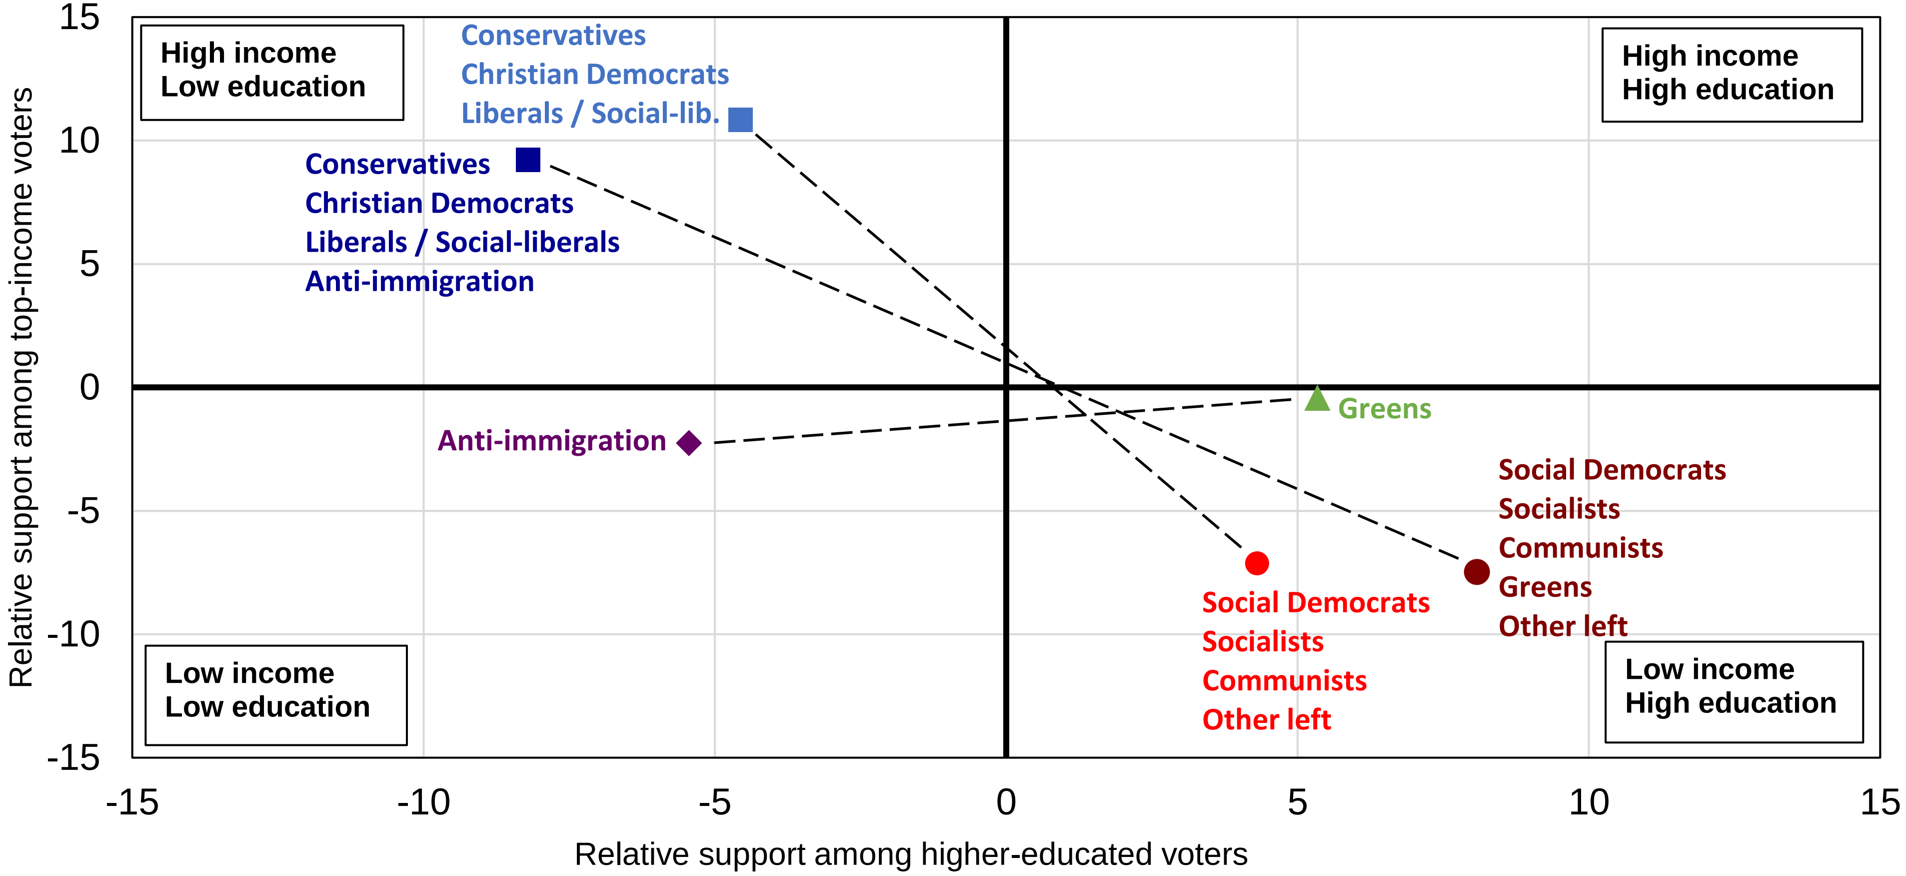

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of high-income (top 10%) and low-income (bottom 90%) voters voting for selected groups of parties on the y-axis, and the same difference between higher-educated (top 10%) and lower-educated (bottom 90%) voters on the x-axis. Education most clearly distinguishes anti-immigration from green parties, while both income and education most clearly distinguishes conservative and Christian democratic parties from socialist, social democratic, and communist parties. Averages over all Western democracies. Estimates control for income/education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure VII - Decomposing income and education cleavages
Panel A. Social Democrats / Socialists / Communists vs. Conservatives / Christian democrats / Liberals

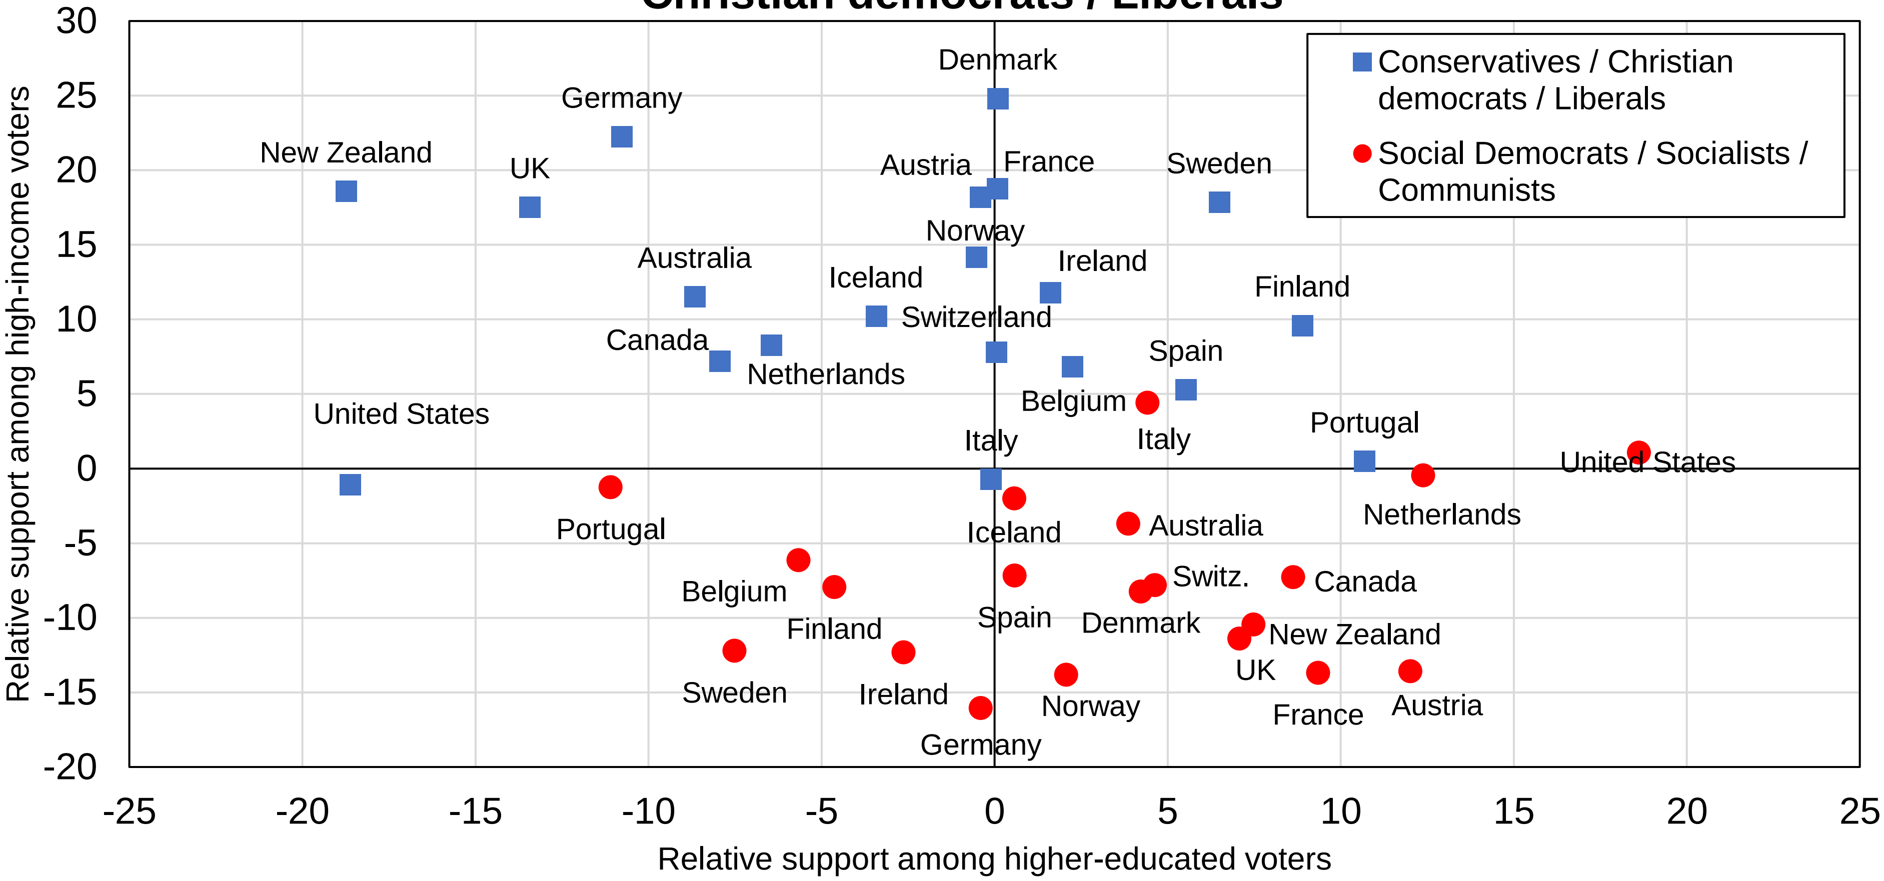

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of high-income (top 10%) and low-income (bottom 90%) voters voting for selected groups of parties on the y-axis, and the same difference between higher-educated (top 10%) and lower-educated (bottom 90%) voters on the x-axis, in the last election available (between 2014 and 2020). Estimates control for income/education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure VII - Decomposing income and education cleavages
Panel B. Green vs. Anti-immigration parties

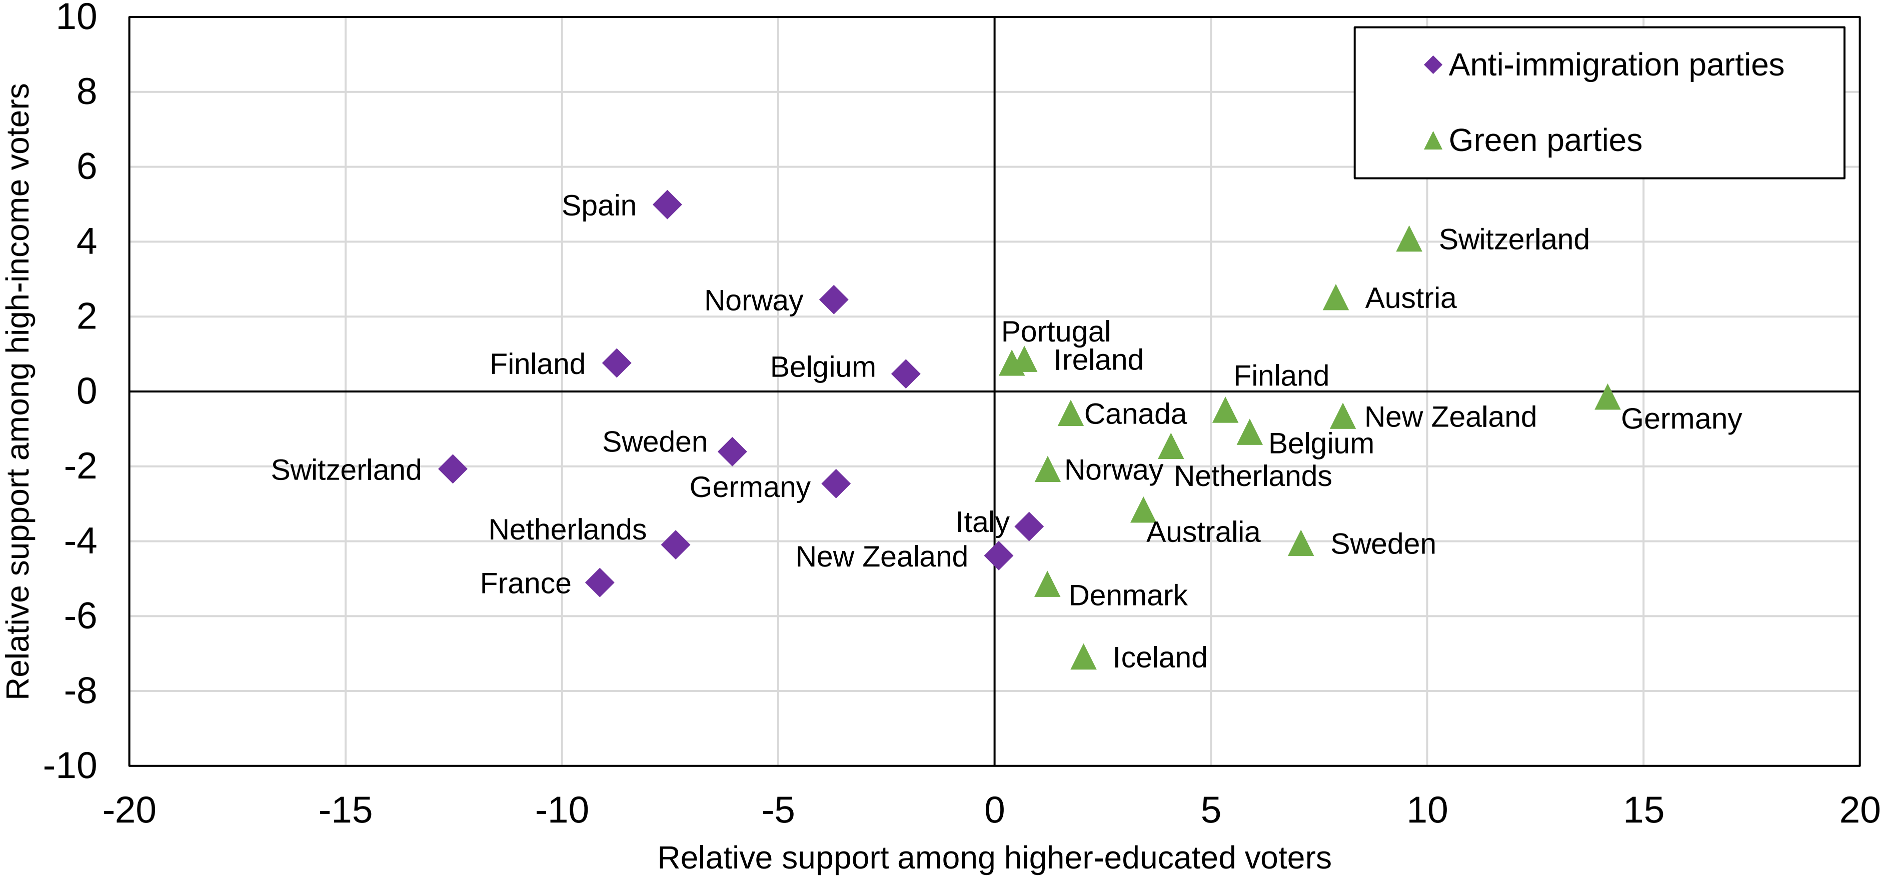

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of high-income (top 10%) and low-income (bottom 90%) voters voting for selected groups of parties on the y-axis, and the same difference between higher-educated (top 10%) and lower-educated (bottom 90%) voters on the x-axis, in the last election available (between 2014 and 2020). Estimates control for income/education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure VIII - The evolution of ideological polarization in Western democracies, 1945-2020
Panel A. Economic-distributive score by party family

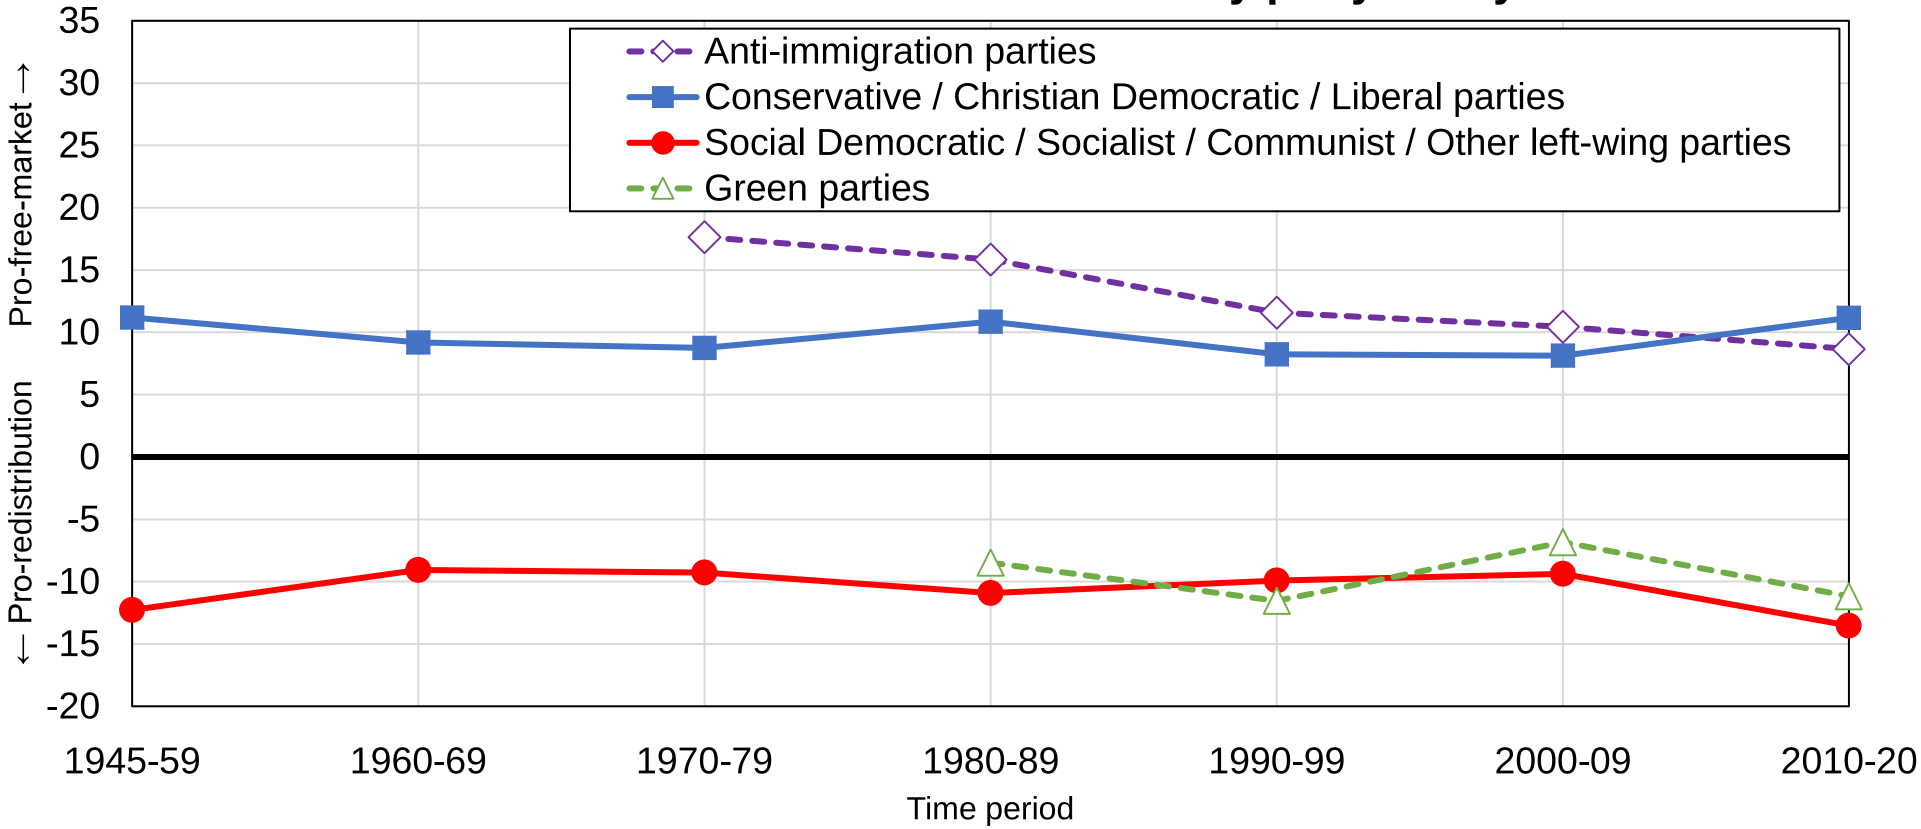

Source: authors' computations using the Comparative Manifesto Project database.
Note: the figure displays the average economic-distributive scores by decade for four families of parties across all Western democracies: social democratic, socialist, communist, and other left-wing parties; conservative, Christian democratic, and liberal parties; anti-immigration parties; and green parties. Negative values on the economic-distributive index correspond to greater proportions of pro-redistribution emphases relatively to pro-free-market emphases. Indices are normalized by the average score by decade so as to better highlight the dynamics of polarization.

Figure VIII - The evolution of ideological polarization in Western democracies, 1945-2020
Panel B. Sociocultural score by party family

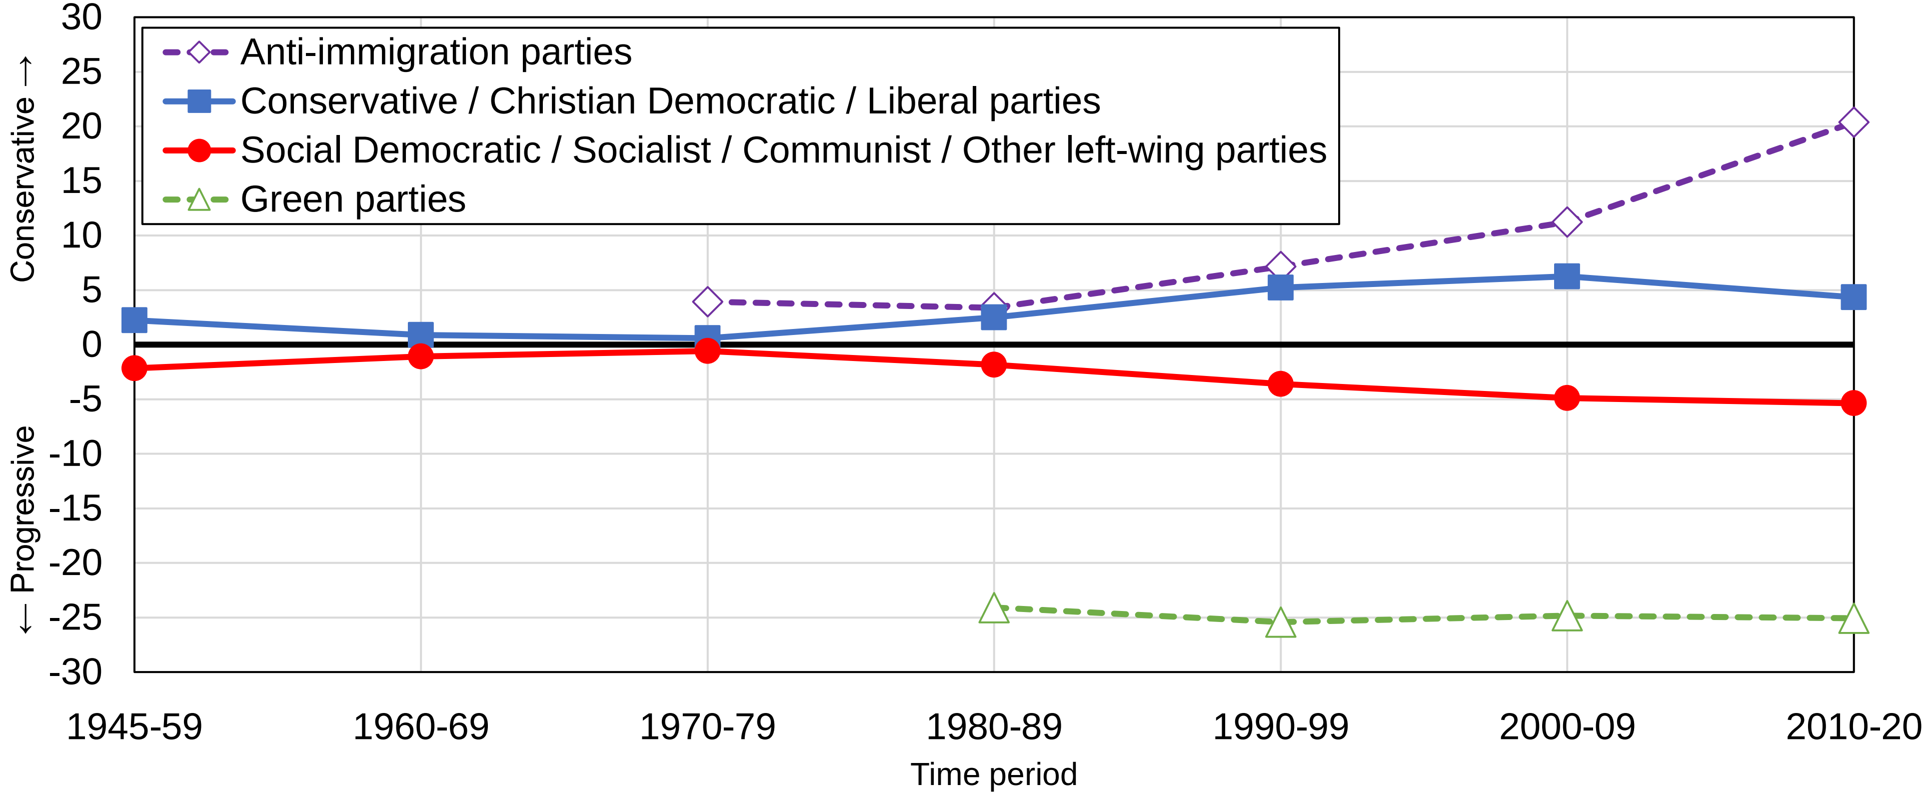

Source: authors' computations using the Comparative Manifesto Project database.
Note: the figure displays the average sociocultural scores by decade for four families of parties across all Western democracies: social democratic, socialist, communist, and other left-wing parties; conservative, Christian democratic, and liberal parties; anti-immigration parties; and green parties. Negative values on the sociocultural index correspond to greater proportions of progressive emphases relatively to conservative emphases. Indices are normalized by the average score by decade so as to better highlight the dynamics of polarization.

Figure IX - Multidimensional political conflict and the divergence of income and education cleavages

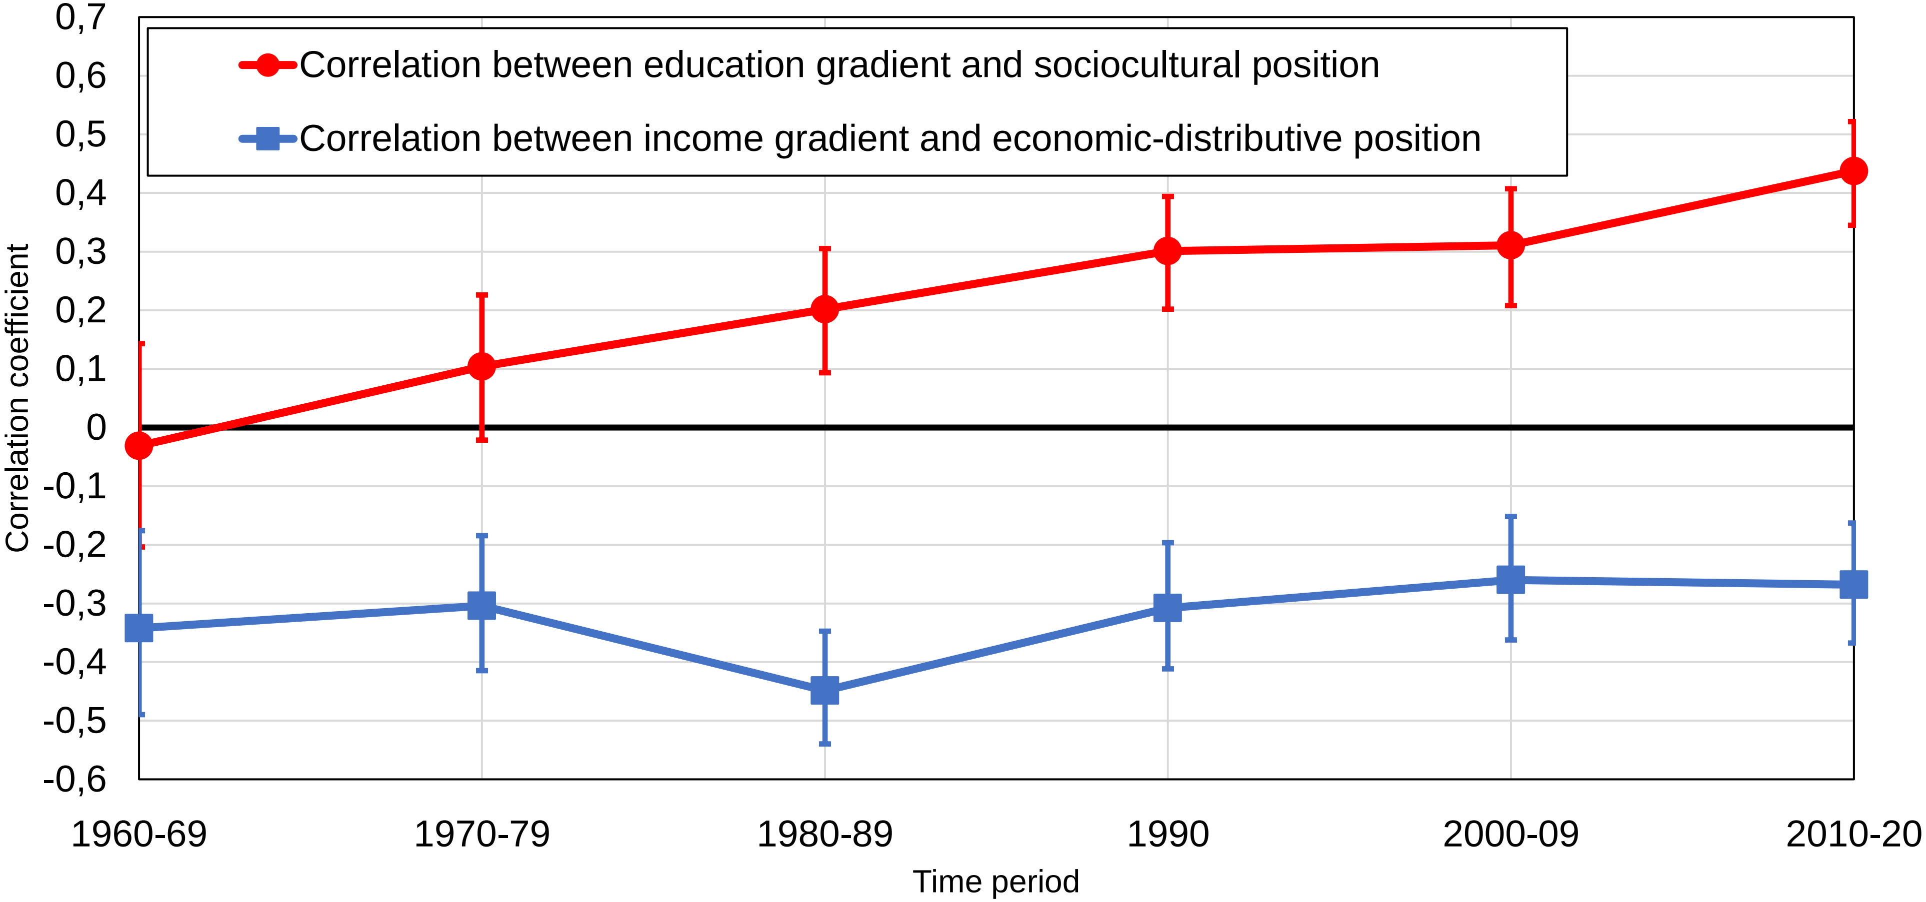

Source: authors' computations combining the World Political Cleavages and Inequality Database with Manifesto Project data.
Note: the upper lines plots the raw correlation between the education gradient (defined as the share of top 10% educated voters within the electorate of a given party) and the sociocultural index across all parties in the database. The bottom line plots the raw correlation between the income gradient (defined as the share of top 10% income voters within the electorate of a given party) and the economic-distributive index (inverted, so that higher values correspond to greater pro-redistribution emphases). The unit of observation is the political party. Error bars represent 95% confidence intervals.

Figure X - Sociocultural polarization and educational divides

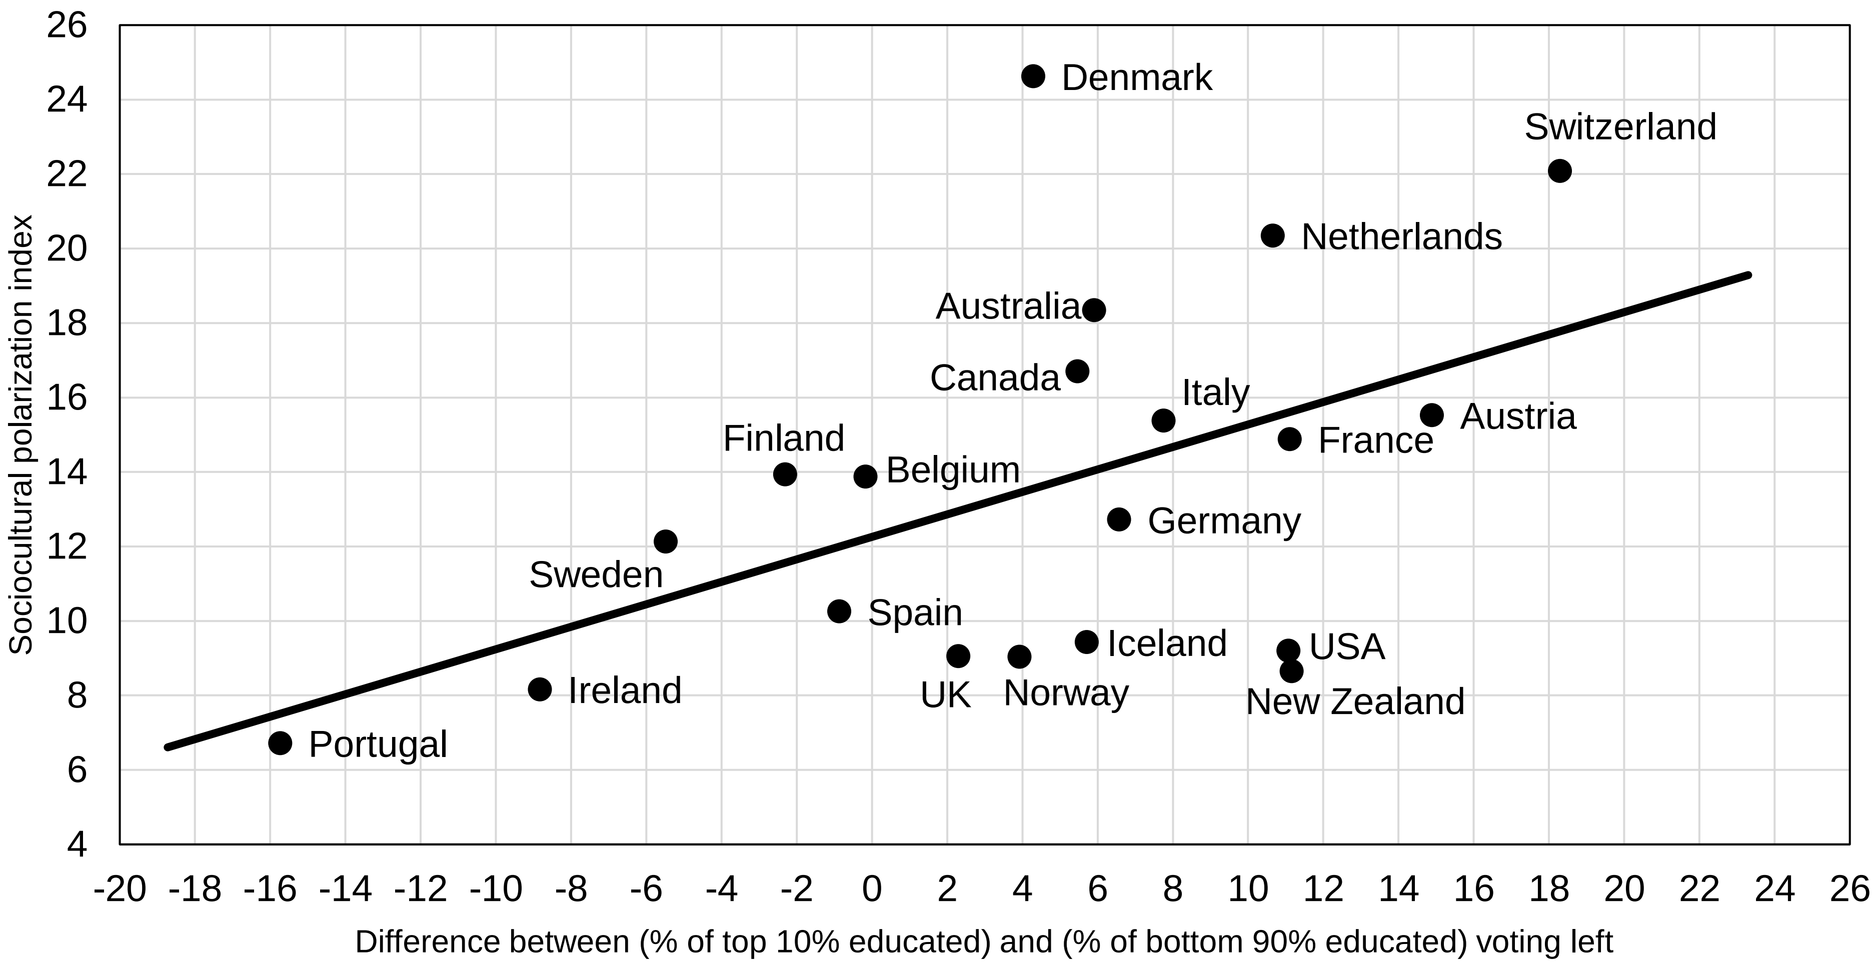

Source: authors' computations combining the World Political Cleavages and Inequality Database with Manifesto Project data.
Note: the figure represents the relationship between sociocultural polarization (defined as the standard deviation of the sociocultural index across all parties in a given country) and the education cleavage for all 21 Western democracies in the 2010s. Higher-educated voters are significantly more likely to support left-wing parties in countries where polarization on the sociocultural axis is higher.

Figure XI - Generational dynamics and educational divides
The education cleavage by birth cohort

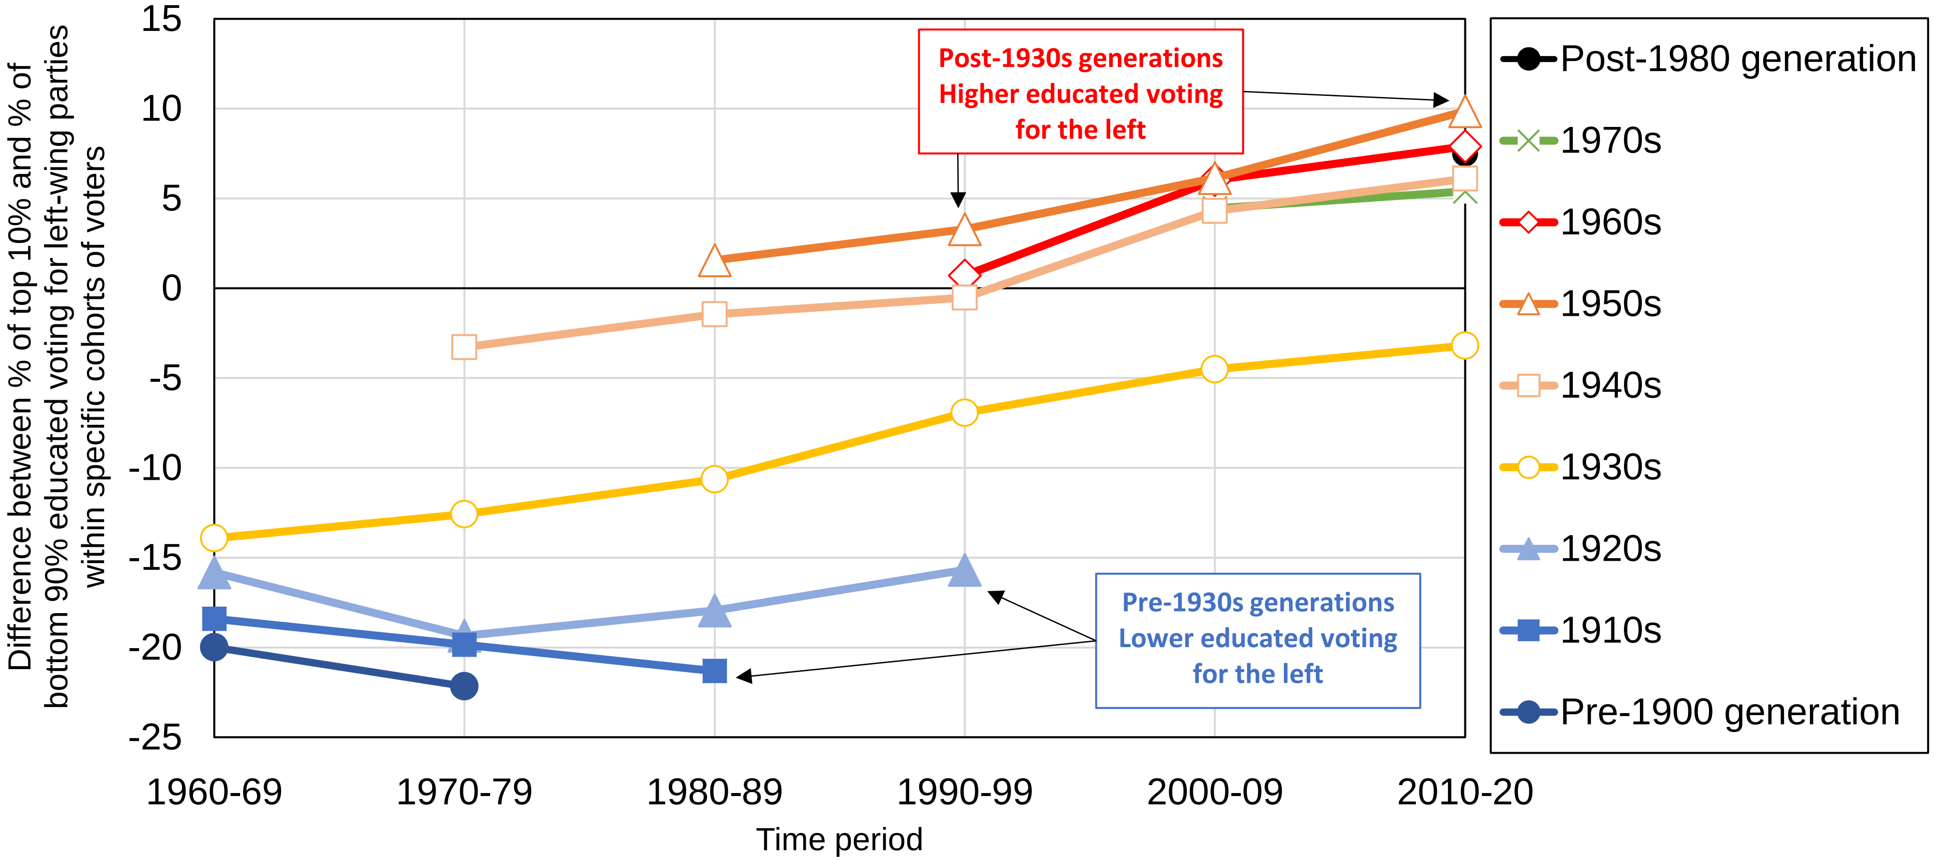

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of higher-educated (top 10%) and lower-educated (bottom 90%) voters voting for social democratic / socialist / communist / green parties within specific cohorts of voters. Between the 1960s and the 1990s, lower-educated voters born in the early decades of the twentieth century remained significantly more likely to vote for these parties than higher-educated voters born during the same period. In the last decade, on the contrary, young lower-educated voters were significantly less likely to vote for these parties than young higher-educated voters. Figures correspond to ten-year averages for Australia, Britain, Canada, Denmark, France, Germany, Italy, the Netherlands, Norway, Sweden, Switzerland, and the US.

Appendix A

Appendix to "Distributional Growth Accounting: Education and the Reduction of Global Poverty, 1980-2019"

A.1 Additional Figures and Tables

A.1.1 Additional Results: World Distribution of Income

Figure A.1: Growth Accounting, 1980-2019: Global Average vs. Poorest 20%

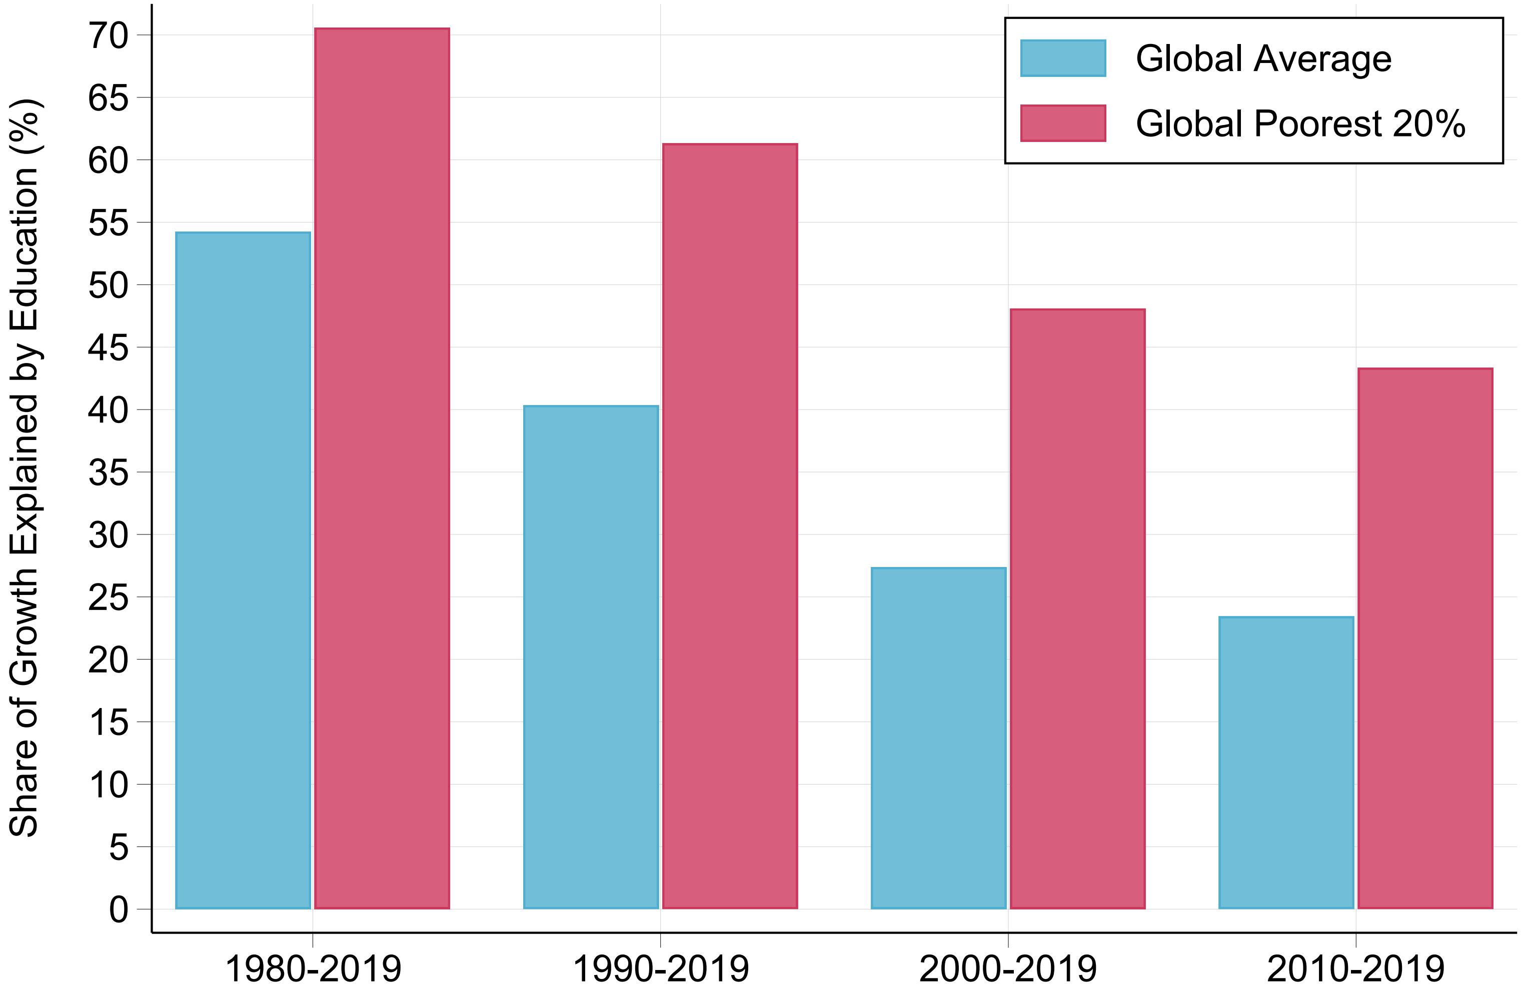

Notes. Author's calculations.

Figure A.2: The Distribution of Global Economic Growth, 2000-2019

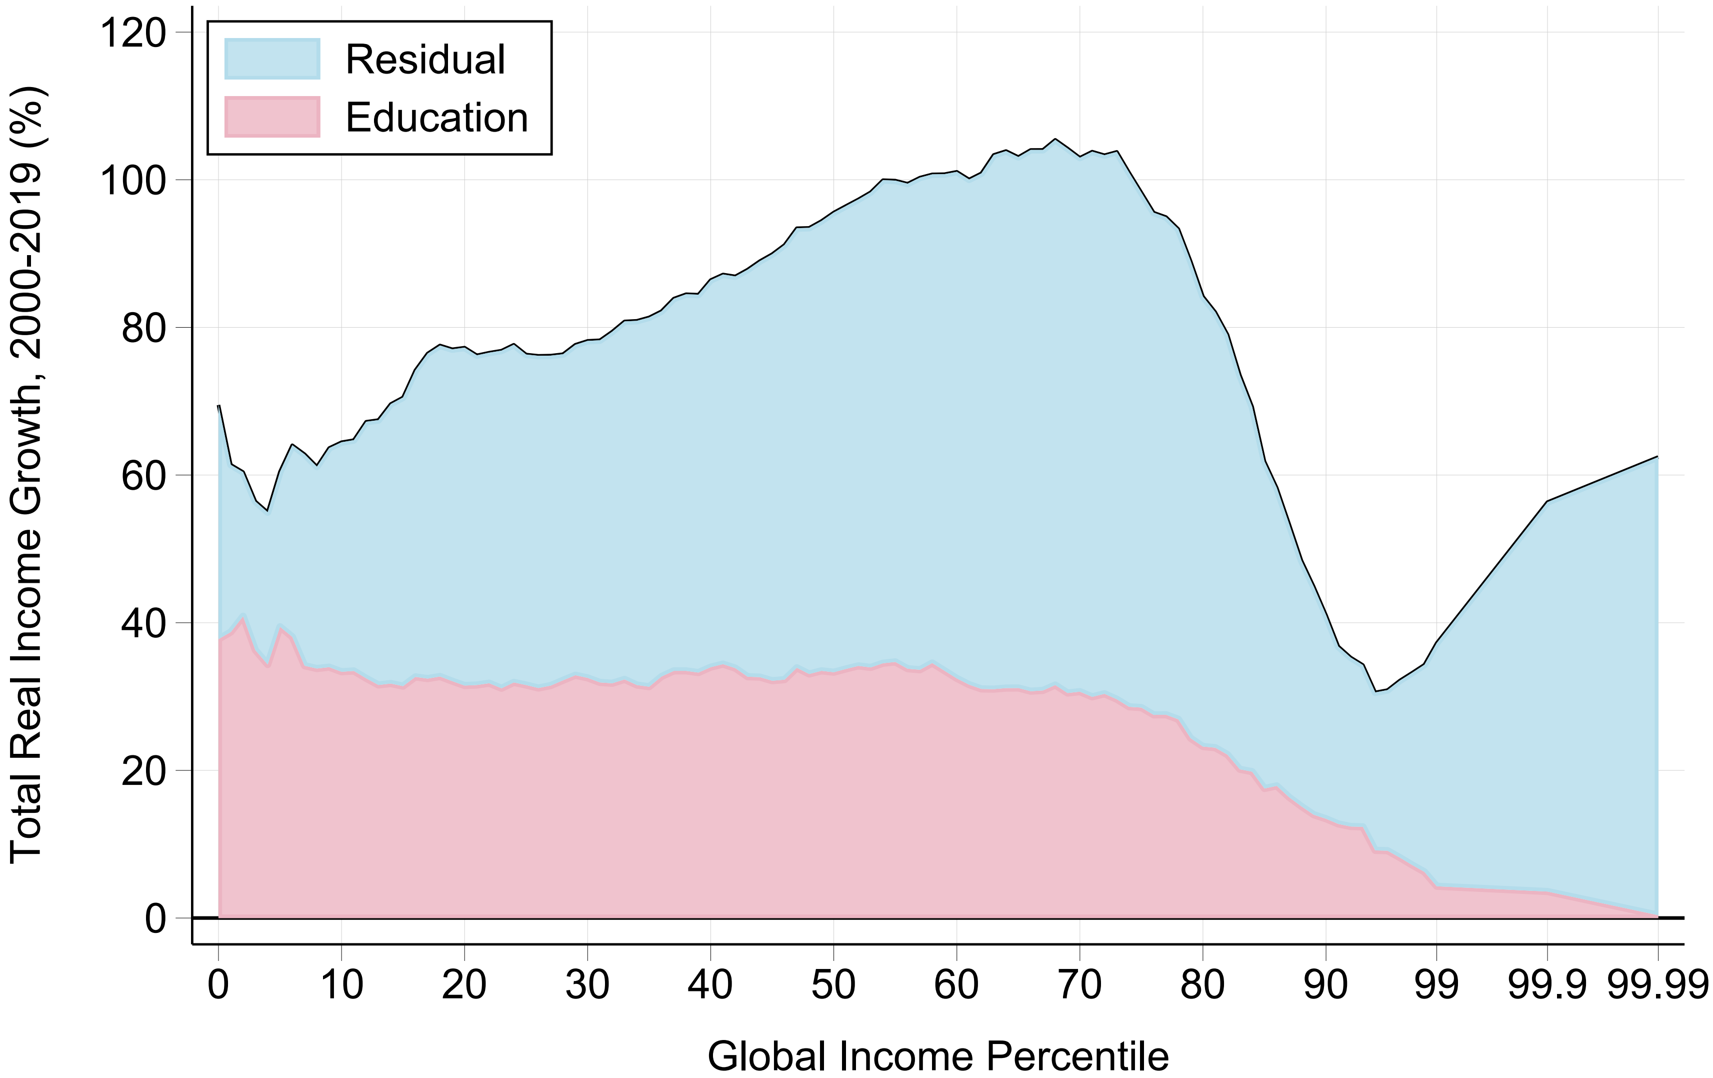

Notes. The figure shows total real income growth by global income percentile, and decomposes it into a part that can be explained by education and an unexplained component. The income concept is pretax national income.

Figure A.3: Income Gains From Schooling by Country, 1980-2019: Average Versus Bottom 50%

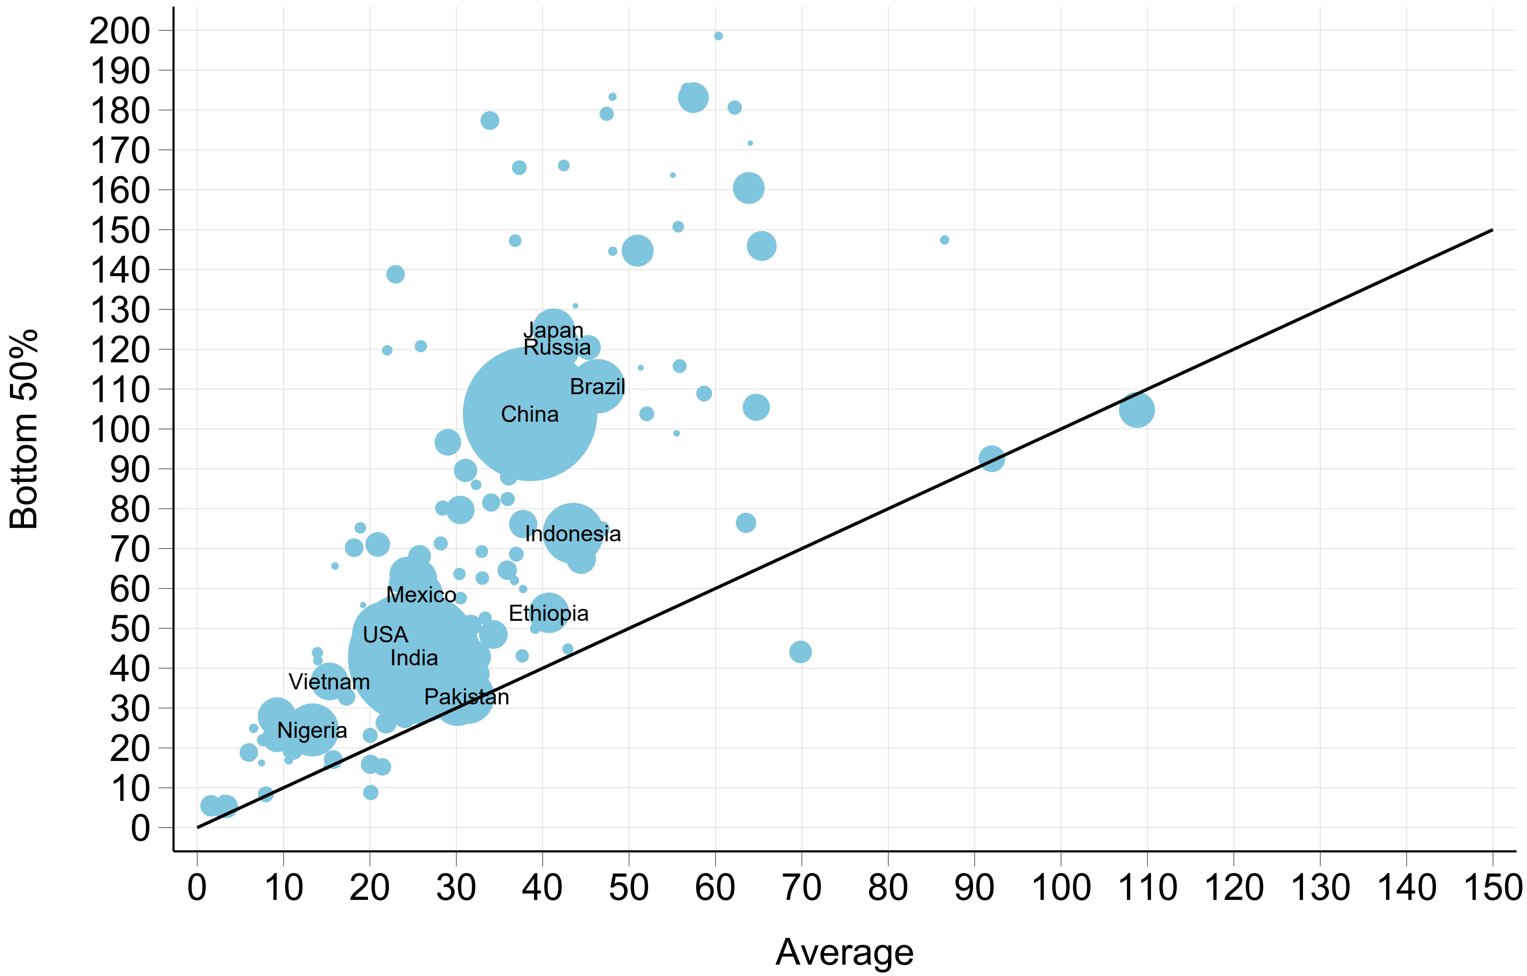

Notes. Author's calculations. The figure compares gains from schooling for the population as a whole (average) versus the bottom 50% in each country over the 1980-2019 period. Gains from schooling correspond to the percent increase in income generated by educational expansion since 1980.

Figure A.4: Contribution of Education to Global Average Annual Income Growth, 1980-2019: Educational Progress Among Post-1980 Generations Only

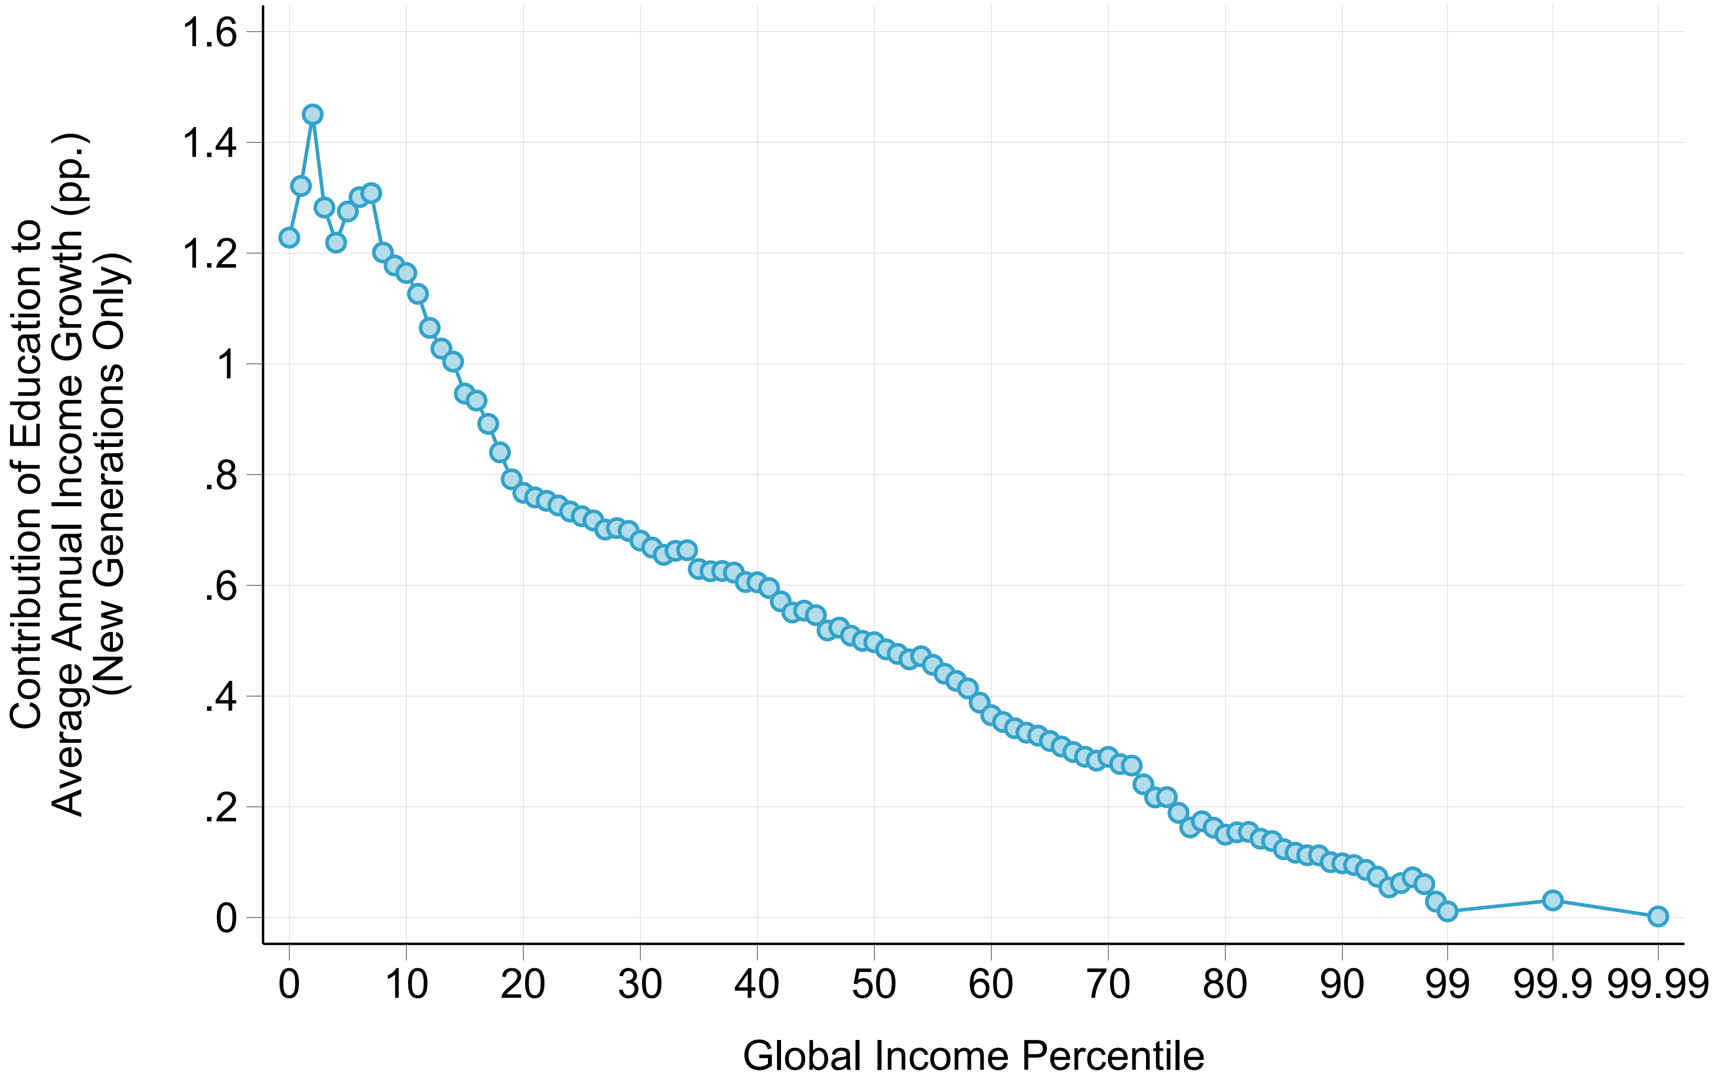

Notes. Author's calculations. The figure reports schooling gains by global income percentile, focusing on educational expansion among post-1980 generations only.

Figure A.5: Share of Growth Explained by Education by Global Income Percentile, 1980-2019:
Educational Progress Among Post-1980 Generations Only

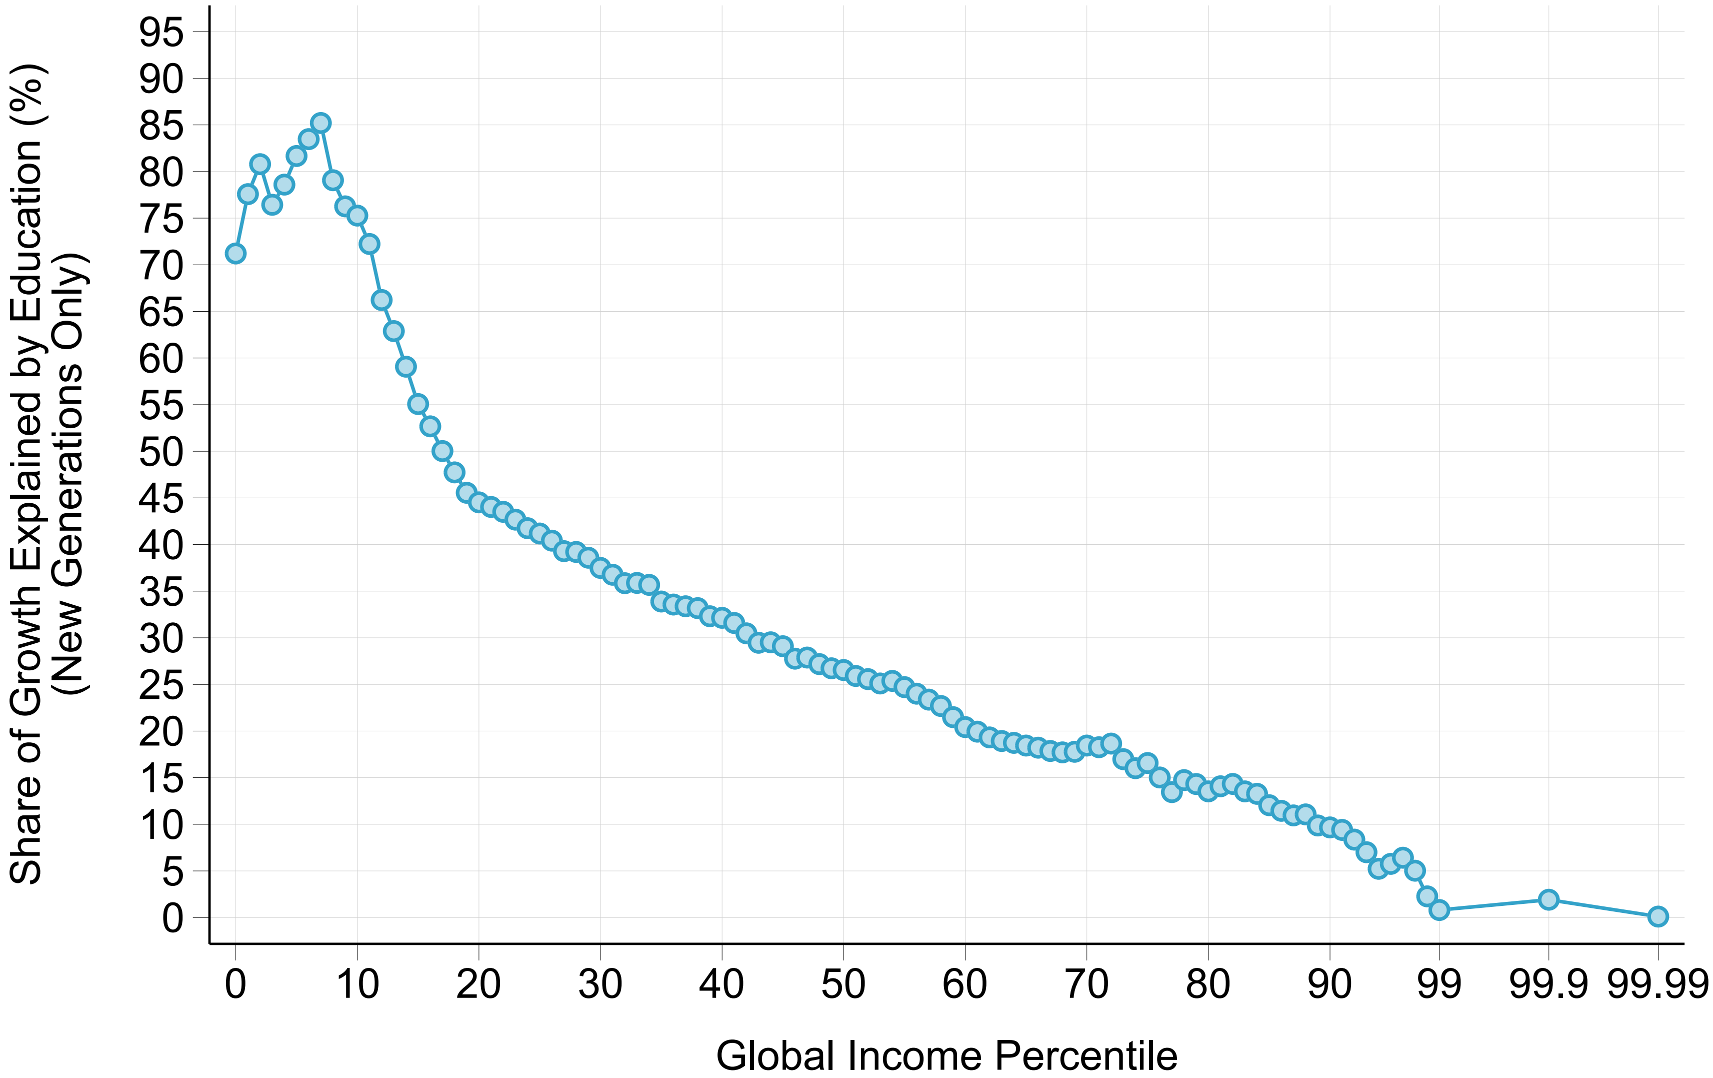

Notes. Author's calculations. The figure reports the share of growth explained by education by global income percentile, focusing on educational expansion among post-1980 generations only.

Figure A.6: Gains from Schooling With and Without Heterogeneous Educational Expansion by Socioeconomic Characteristic: India, 1983-2019

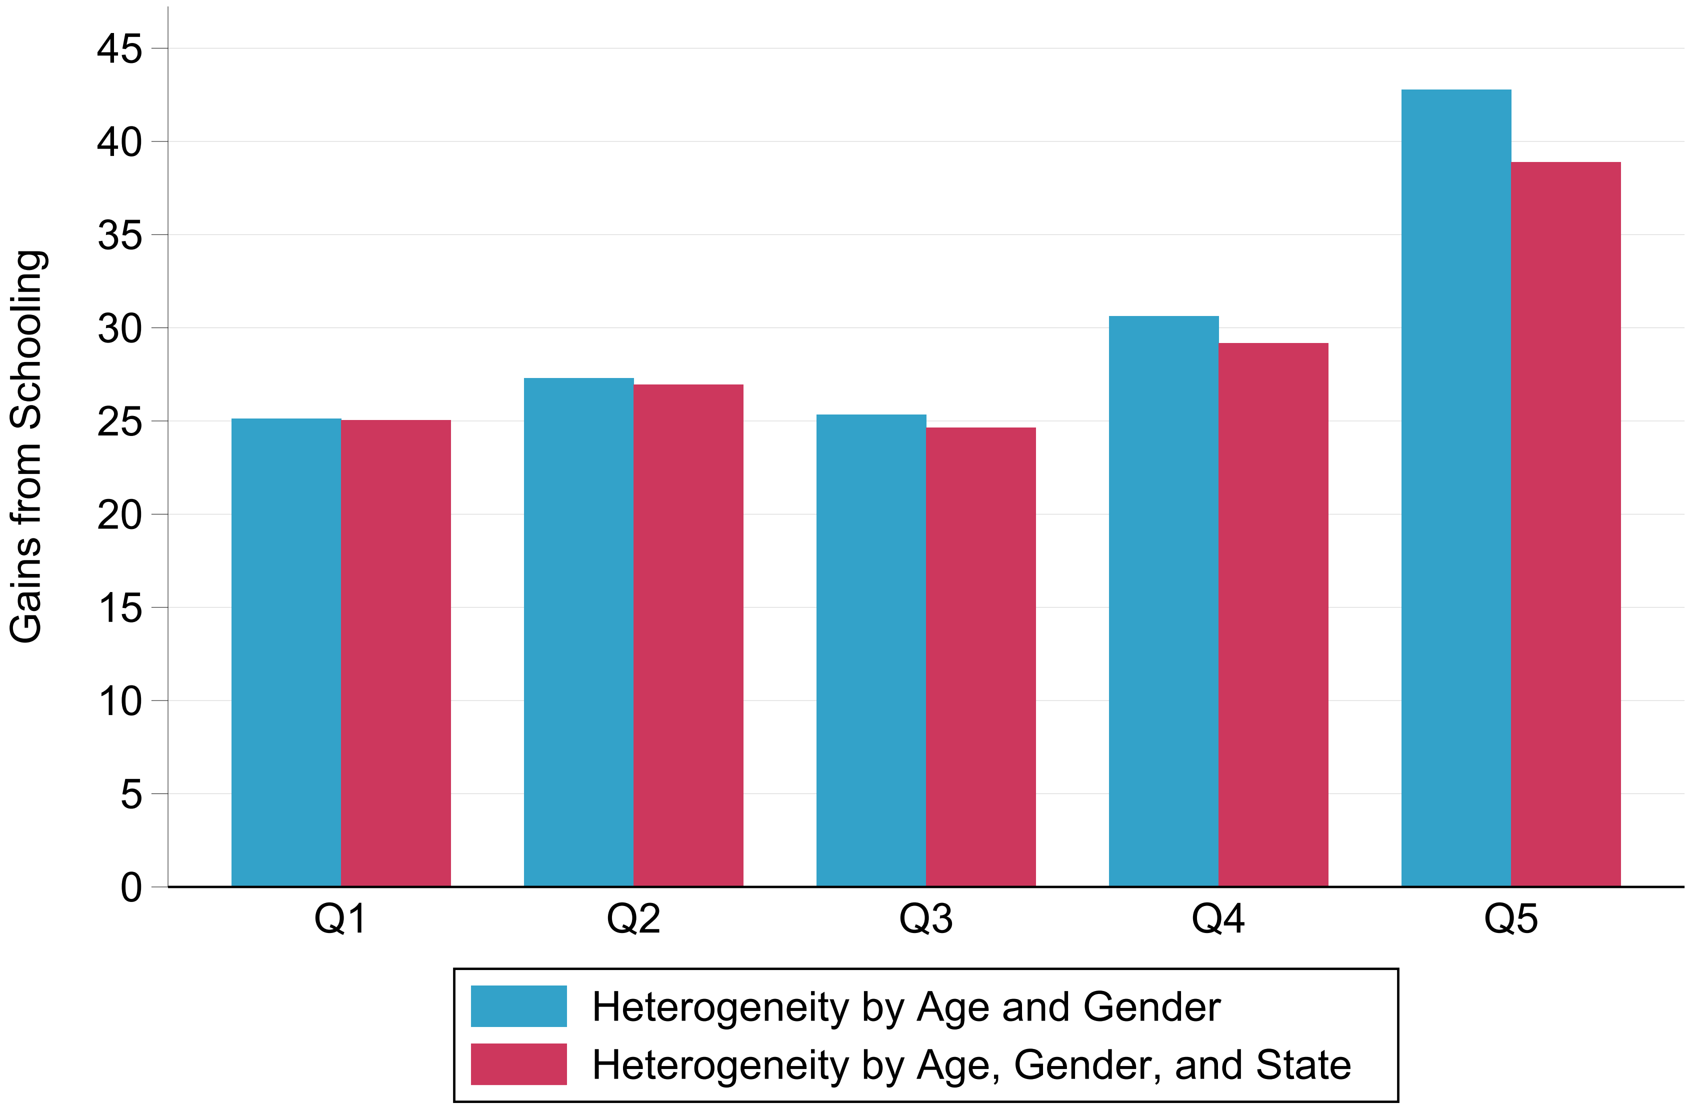

Notes. Educational attainment by age, gender, and state of residence in 1983 is estimated using the 1983 National Sample Survey. Education levels of individuals in the 2019 Periodic Labor Force Survey are then downgraded by age-gender cell (specification 1) or age-gender-state cell (specification 2) until reaching 1983 levels. Their earnings are reduced using estimates of returns to schooling by level. Finally, the figure plots schooling gains by income quintile, defined as the percent difference between actual income and counterfactual income absent educational expansion since 1983.

Figure A.7: Gains from Schooling With and Without Heterogeneous Educational Expansion by Socioeconomic Characteristic: South Africa, 2002-2019

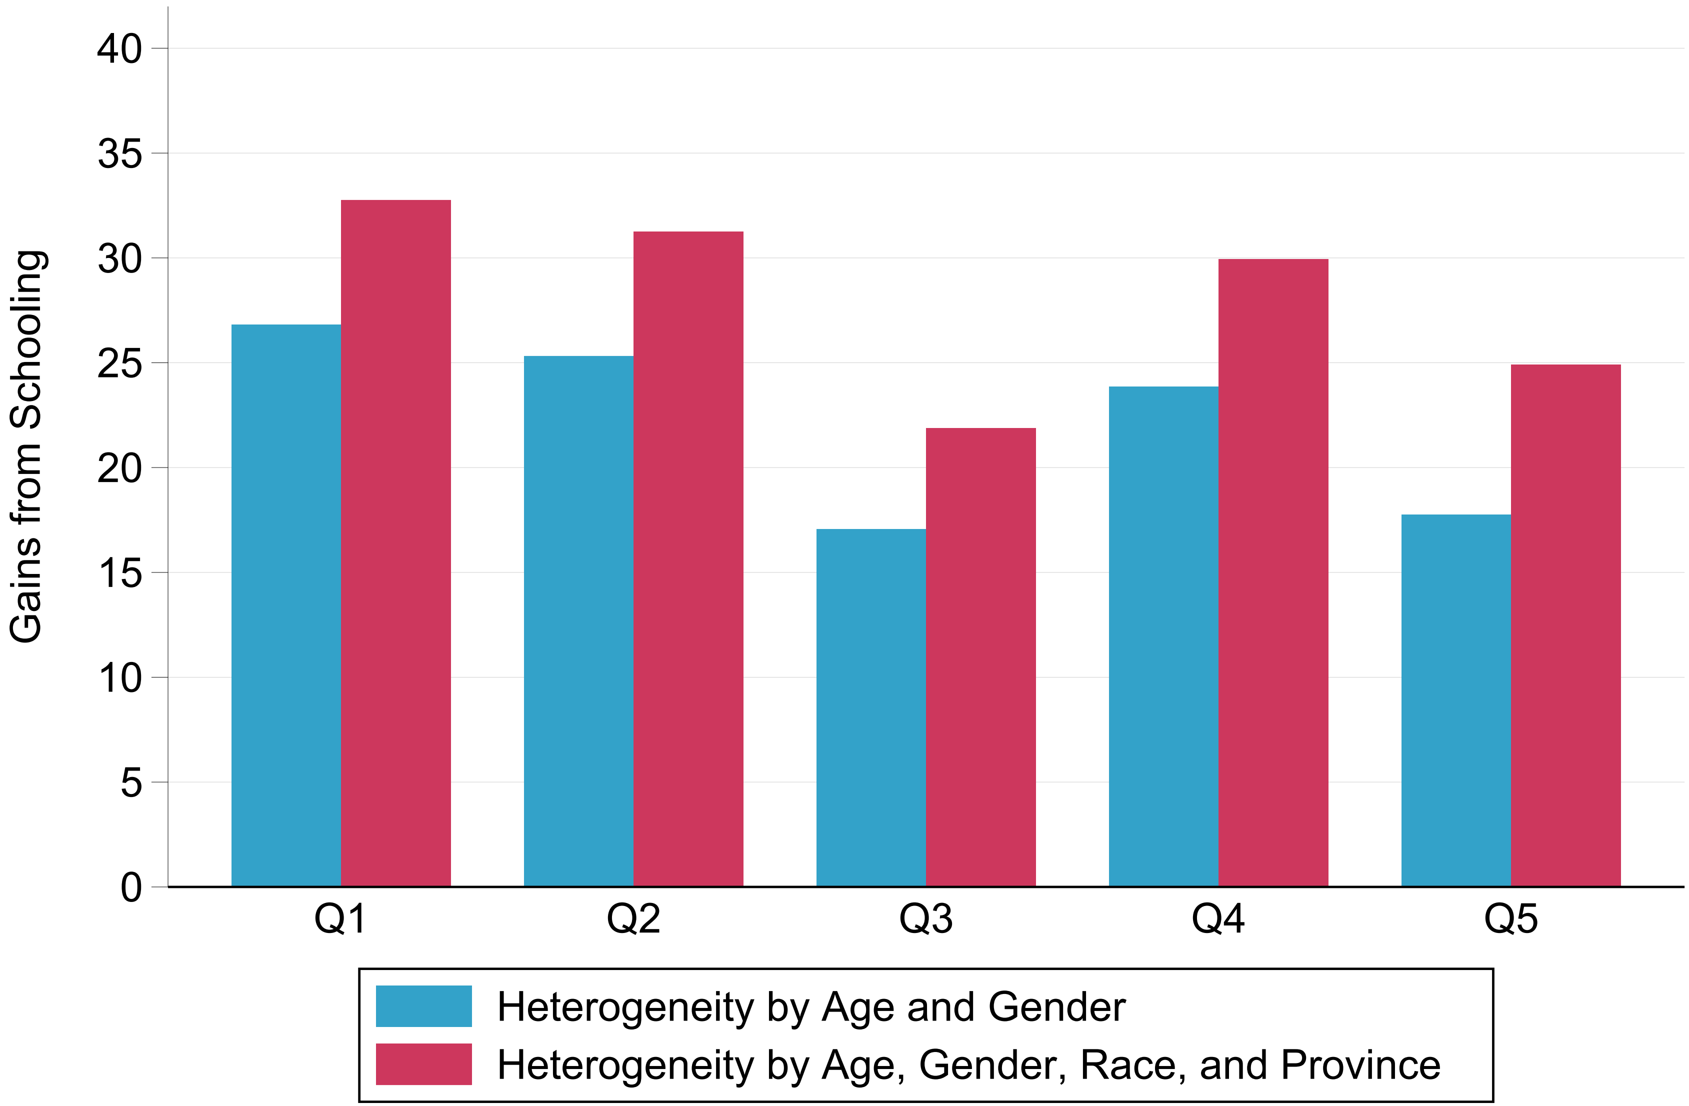

Notes. Educational attainment by age, gender, race, and province of residence in 2002 is estimated using the 2002 General Household Survey. Education levels of individuals in the 2019 General Household Survey are then downgraded by age-gender cell (specification 1) or age-gender-race-province cell (specification 2) until reaching 2002 levels. Their earnings are reduced using estimates of returns to schooling by level. Finally, the figure plots schooling gains by income quintile, defined as the percent difference between actual income and counterfactual income absent educational expansion since 2002.

Figure A.8: Gains from Schooling With and Without Heterogeneous Educational Expansion by Socioeconomic Characteristic: United States, 1980-2019

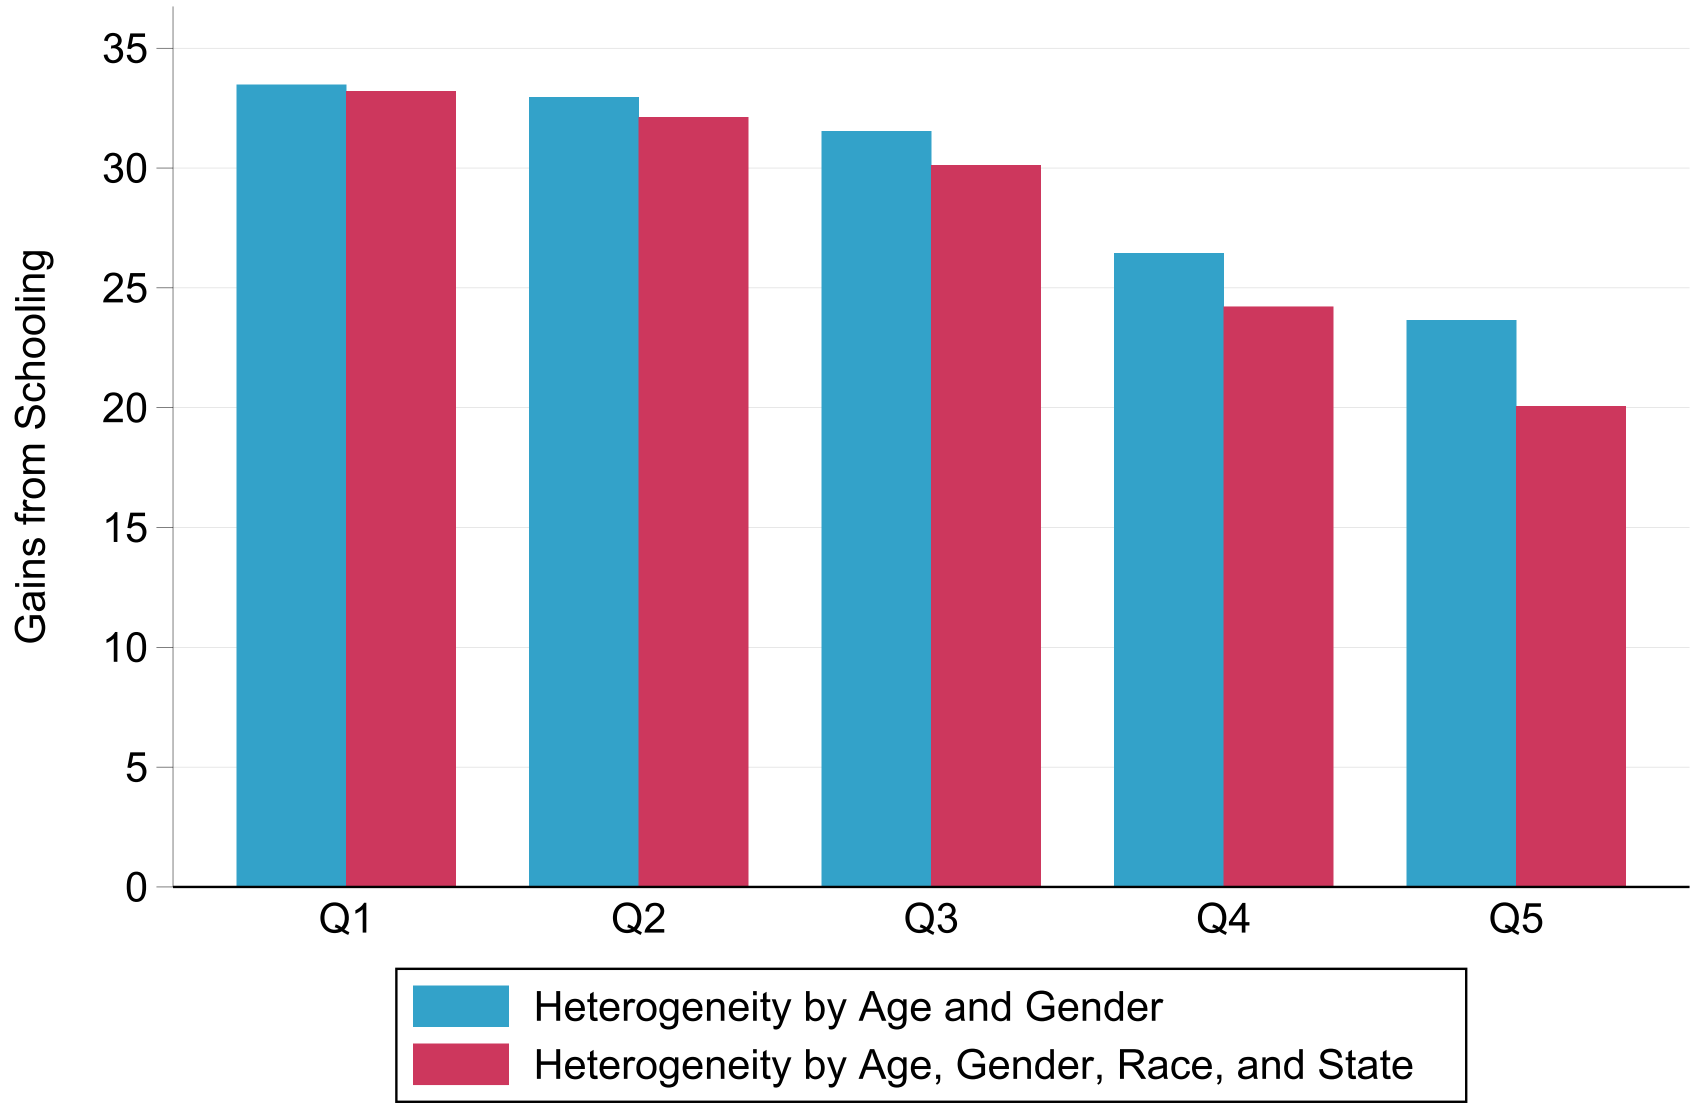

Notes. Educational attainment by age, gender, race, and state of residence in 1980 is estimated using 1980 IPUMS census sample microdata. Education levels of individuals in the 2019 Current Population Survey are then downgraded by age-gender cell (specification 1) or age-gender-race-state cell (specification 2) until reaching 1980 levels. Their earnings are reduced using estimates of returns to schooling by level. Finally, the figure plots schooling gains by income quintile, defined as the percent difference between actual income and counterfactual income absent educational expansion since 1980.

Table A.1: Distributional Growth Accounting, World: 2000-2019

	Total Income Growth (%)	Growth Without Education (%)	Contribution of Education (pp.)	Share of Growth Explained (%)
	g	$\tilde{g}$	$g-\tilde{g}$	$\frac{g-\tilde{g}}{g}$
Full Population	+55%	+40%	15	27%
Bottom 50%	+82%	+49%	33	40%
Bottom 20%	+69%	+36%	33	48%
Next 30%	+84%	+52%	33	39%
Middle 40%	+78%	+55%	23	30%
Top 10%	+37%	+29%	8	21%
Top 1%	+42%	+38%	4	10%
Top 0.1%	+57%	+53%	3	6%
Top 0.01%	+62%	+62%	0.3	0.6%

Notes. The table reports actual real income growth rates, counterfactual growth rates absent educational expansion, and the corresponding share of growth explained by education for different groups of the world distribution of income.

Table A.2: Distributional Growth Accounting, World: World Bank Data

	1980-2019			2000-2019		
	Total Income Growth (%)	Growth Without Education (%)	Share of Growth Explained (%)	Total Income Growth (%)	Growth Without Education (%)	Share of Growth Explained (%)
	g	$\tilde{g}$	$1-\frac{\tilde{g}}{g}$	g	$\tilde{g}$	$1-\frac{\tilde{g}}{g}$
Full Population	+77%	+17%	77%	+48%	+30%	39%
Bottom 50%	+206%	+79%	61%	+97%	+57%	41%
Bottom 20%	+196%	+74%	62%	+83%	+43%	48%
Next 30%	+208%	+81%	61%	+102%	+62%	39%
Middle 40%	+84%	+7%	91%	+78%	+50%	36%
Top 10%	+61%	+17%	72%	+27%	+16%	42%
Top 1%	+77%	+40%	47%	+26%	+17%	35%
Top 0.1%	+102%	+66%	36%	+20%	+11%	49%
Top 0.01%	+155%	+118%	24%	+5%	+2%	66%

Notes. The table reports actual real income growth rates, counterfactual growth rates absent educational expansion, and the corresponding share of growth explained by education for different groups of the world distribution of income. Estimates of the world distribution of income constructed from World Bank per-capita income and consumption data.

Table A.3: Education and Global Poverty Reduction: World Bank Data

	1980	2019	Difference (%)	Share of Decline Explained (%)
Global Poverty: $2.15 / Day				
Actual	46%	10%	-80%	
Counterfactual	46%	24%	-48%	40%
Global Poverty: $3.65 / Day				
Actual	60%	25%	-59%	
Counterfactual	60%	45%	-26%	56%
Global Poverty: $6.85 / Day				
Actual	70%	47%	-32%	
Counterfactual	70%	64%	-9%	73%

Notes. The table compares the actual evolution of the global poverty headcount ratio to the evolution it would have followed absent educational expansion since 1980. All global poverty headcount ratios calculated using 2017 PPP USD. Estimates of the world distribution of income constructed from World Bank per-capita income and consumption data.

Table A.4: Distributional Growth Accounting, World:
Alternative Elasticities of Substitution

	1980-2019			2000-2019		
	Low Substitutability	Benchmark	High Substitutability	Low Substitutability	Benchmark	High Substitutability
Full Population	61%	54%	53%	30%	27%	27%
Bottom 50%	72%	59%	56%	51%	40%	39%
Bottom 20%	85%	71%	66%	61%	48%	46%
Next 30%	70%	57%	54%	49%	39%	38%
Middle 40%	92%	74%	73%	35%	30%	30%
Top 10%	35%	38%	38%	17%	21%	21%
Top 1%	10%	17%	16%	5%	10%	9%
Top 0.1%	7%	9%	8%	5%	6%	6%
Top 0.01%	2%	2%	2%	0.4%	0.6%	0.5%

Notes. The table reports the share of growth explained by education for different groups of the world distribution of income, depending on assumptions made on the substitutability of skilled and unskilled workers. Low substitutability: $\sigma_1 = \sigma_2 = \sigma_3 = 4$. Benchmark: $\sigma_1 = \sigma_2 = \sigma_3 = 6$. High substitutability: $\sigma_1 = 5$, $\sigma_2 = 7$, $\sigma_3 = 9$.

Table A.5: Distributional Growth Accounting, World, 1980-2019:
Alternative Nesting of CES Production Function

	Total Income Growth (%)	Growth Without Education (%)	Contribution of Education (pp.)	Share of Growth Explained (%)
	g	$\tilde{g}$	$g - \tilde{g}$	$\frac{g-\tilde{g}}{g}$
Full Population	+98%	+45%	52	53%
Bottom 50%	+164%	+79%	85	52%
Bottom 20%	+115%	+40%	74	65%
Next 30%	+176%	+89%	87	50%
Middle 40%	+94%	+28%	66	70%
Top 10%	+91%	+54%	37	41%
Top 1%	+131%	+107%	24	18%
Top 0.1%	+173%	+157%	17	10%
Top 0.01%	+278%	+270%	8	3%

Notes. The table reports actual real income growth rates, counterfactual growth rates absent educational expansion, and the corresponding share of growth explained by education for different groups of the world distribution of income.

Table A.6: Distributional Growth Accounting, World, 1980-2019:
With Capital Income Affected by Education

	Total Income Growth (%)	Growth Without Education (%)	Contribution of Education (pp.)	Share of Growth Explained (%)
	g	$\tilde{g}$	$g - \tilde{g}$	$\frac{g-\tilde{g}}{g}$
Full Population	+98%	+32%	66%	67%
Bottom 50%	+164%	+54%	110%	67%
Bottom 20%	+115%	+23%	92%	80%
Next 30%	+176%	+61%	114%	65%
Middle 40%	+94%	+14%	80%	85%
Top 10%	+91%	+42%	49%	54%
Top 1%	+131%	+81%	50%	38%
Top 0.1%	+173%	+111%	62%	36%
Top 0.01%	+278%	+194%	84%	30%

Notes. The table reports actual real income growth rates, counterfactual growth rates absent educational expansion, and the corresponding share of growth explained by education for different groups of the world distribution of income. Returns to schooling are assumed to affect both labor income and capital income by the same amount.

Table A.7: Distributional Growth Accounting, World, 1980-2019:
Educational Progress Among Post-1980 Generations Only

	Total Income Growth (%)	Growth Without Education (%)	Contribution of Education (pp.)	Share of Growth Explained (%)
	g	$\tilde{g}$	$g-\tilde{g}$	$\frac{g-\tilde{g}}{g}$
Full Population	+98%	+86%	12	12%
Bottom 50%	+164%	+103%	61	37%
Bottom 20%	+115%	+45%	70	61%
Next 30%	+176%	+117%	59	33%
Middle 40%	+94%	+79%	15	16%
Top 10%	+91%	+87%	4	4%
Top 1%	+131%	+130%	1	1%
Top 0.1%	+173%	+170%	3	2%
Top 0.01%	+278%	+277%	0.2	0.1%

Notes. The table reports actual real income growth rates, counterfactual growth rates absent educational expansion, and the corresponding share of growth explained by education for different groups of the world distribution of income. Educational expansion is defined as improvements in educational attainment among post-1980 generations only.

Table A.8: Cross-Country Correlates of Schooling: Cross-Sectional Estimates

	Expected Years of Schooling	Primary School Enrollment	Secondary School Enrollment
Log Public Education Expenditure Per Child	1.351*	10.916***	4.115**
	(0.699)	(1.391)	(1.949)
Government Effectiveness	0.554*	1.961	-2.093
	(0.317)	(1.638)	(2.288)
Trade-to-GDP Ratio	-0.547***	-0.624	-3.220***
	(0.150)	(0.785)	(1.136)
Internet Usage	0.292	-1.899	1.969
	(0.372)	(1.961)	(2.778)
Mobile cellular subscriptions (per 100 people)	-0.045	3.584**	1.987
	(0.294)	(1.535)	(2.152)
Skill Bias of Technology	-0.124	-1.207	-0.151
	(0.196)	(1.025)	(1.425)
Log GDP Per Capita	0.600	-4.813**	7.720**
	(0.539)	(2.368)	(3.392)
Child Population (% Total)	-0.820**	-1.457	-13.819***
	(0.347)	(1.699)	(2.426)
Constant	12.062***	84.302***	54.134***
	(0.510)	(2.632)	(3.714)
Treament			
Binary			
N	139	140	135
Adj. R-squared	0.75	0.54	0.80

Notes. All variables are standardized to have a mean of 0 and standard deviation of 1. Log public education expenditure per child: data from Gethin (2023b). Skill bias of technology: average of relative efficiency terms A_H/A_L for primary, secondary, and tertiary education, weighted by the share of workers in each category, estimated using labor force survey microdata. Log GDP per capita: data from the World Inequality Database. Education expenditure and GDP expressed in 2019 PPP USD. All other variables: data from the World Bank Development Indicators. Standard errors in parenthesis. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.9: Cross-Country Correlates of Schooling: Panel Estimates

	Expected Years of Schooling	Primary School Enrollment	Secondary School Enrollment
Log Public Education Expenditure Per Child	1.490***	6.110***	8.032***
	(0.246)	(1.321)	(1.527)
Government Effectiveness	-0.170*	0.328	-0.734
	(0.097)	(0.625)	(0.669)
Trade-to-GDP Ratio	0.157**	0.398	-0.446
	(0.061)	(0.378)	(0.353)
Internet Usage	-0.542***	-5.183***	-4.089***
	(0.053)	(0.319)	(0.370)
Mobile cellular subscriptions (per 100 people)	0.146***	0.228	1.369***
	(0.052)	(0.307)	(0.327)
Log GDP Per Capita	0.137	-2.125	2.472
	(0.241)	(1.381)	(1.639)
Child Population (% Total)	-0.527***	3.996***	-8.557***
	(0.147)	(0.910)	(1.054)
Constant	12.617***	91.593***	64.970***
	(0.110)	(0.506)	(0.801)
Treament			
Binary			
N	1,888	2,060	1,561
Adj. R-squared	0.95	0.86	0.97

Notes. All variables are standardized to have a mean of 0 and standard deviation of 1. Log public education expenditure per child: data from Gethin (2023b). Log GDP per capita: data from the World Inequality Database. Education expenditure and GDP expressed in 2019 PPP USD. All other variables: data from the World Bank Development Indicators. Standard errors in parenthesis. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.10: Income Gains from Schooling With and Without Relative Efficiency Gains, 2000-2019

	Average			Bottom 50%		
	Without Efficiency Gains	With Efficiency Gains	Ratio	Without Efficiency Gains	With Efficiency Gains	Ratio
Europe	+9%	+17%	1.88	+13%	+41%	3.12
United States	+4%	+4%	1.06	+6%	+8%	1.29
Brazil	+24%	+23%	0.93	+49%	+50%	1.03
Mexico	+10%	+8%	0.87	+20%	+18%	0.92
Other Latin America	+9%	+8%	0.96	+15%	+21%	1.43
Indonesia	+15%	+27%	1.86	+31%	+45%	1.44
Thailand	+16%	+28%	1.77	+27%	+42%	1.57
Ghana	+4%	+3%	0.74	+7%	+5%	0.77
South Africa	+13%	+9%	0.70	+37%	+38%	1.04
Average	+9%	+13%	1.38	+16%	+30%	1.88

Notes. The table compares income gains from schooling with and without relative efficiency gains in selected countries and groups of countries. With efficiency gains (backward growth accounting): income gains from schooling estimated by reducing incomes in 2019 to match education levels observed in 2000 (holding the relative skill bias to its 2019 level). Without efficiency gains (forward growth accounting): income gains from schooling estimated by increasing incomes in 2000 to match education levels observed in 2019 (holding the relative skill bias to its 2000 level). Europe: Austria, Belgium, Denmark, Estonia, Finland, France, Iceland, Ireland, Luxembourg, Norway, Portugal, Sweden, Switzerland, United Kingdom. Other Latin America: Argentina, Bolivia, Chile, Colombia, Costa Rica, Dominican Republic, Ecuador, Guatemala, Panama, Paraguay, Peru, Uruguay.

Table A.11: Share of Growth Explained by Education Without and With Relative Efficiency Gains, 2000-2019

	Average			Bottom 50%		
	Without Efficiency Gains	With Efficiency Gains	Ratio	Without Efficiency Gains	With Efficiency Gains	Ratio
Europe	51%	73%	1.45	60%	89%	1.49
United States	18%	19%	1.06	44%	57%	1.28
Brazil	91%	86%	0.95	¿100%	¿100%	1.00
Mexico	¿100%	¿100%	1.00	¿100%	¿100%	1.00
Other Latin America	22%	21%	0.96	28%	39%	1.35
Indonesia	21%	35%	1.67	52%	68%	1.30
Thailand	31%	49%	1.60	36%	50%	1.40
Ghana	8%	6%	0.75	15%	12%	0.79
South Africa	58%	41%	0.72	¿100%	¿100%	1.00
Average	36%	45%	1.25	46%	63%	1.37

Notes. The table compares the share of growth explain by education with and without relative efficiency gains in selected countries and groups of countries. With efficiency gains (backward growth accounting): income gains from schooling estimated by reducing incomes in 2019 to match education levels observed in 2000 (holding the relative skill bias to its 2019 level). Without efficiency gains (forward growth accounting): income gains from schooling estimated by increasing incomes in 2000 to match education levels observed in 2019 (holding the relative skill bias to its 2000 level). Europe: Austria, Belgium, Denmark, Estonia, Finland, France, Iceland, Ireland, Luxembourg, Norway, Portugal, Sweden, Switzerland, United Kingdom. Other Latin America: Argentina, Bolivia, Chile, Colombia, Costa Rica, Dominican Republic, Ecuador, Guatemala, Panama, Paraguay, Peru, Uruguay.

Table A.12: Public Policies and Global Poverty Reduction:
Combining Direct Redistribution and Indirect Investment Benefits from Education (World Bank Data)

	1980	2019	Change (%)	Total Share of Change Explained (%)
Global Poverty Rate ($2.15/Day)				
Pretax Income Absent Educational Expansion	46%	24%	-48%	
Pretax Income	46%	9.5%	-80%	
Posttax Income	44%	4.9%	-89%	46%
Global Bottom 20% Average Income ($/Day)				
Pretax Income Absent Educational Expansion	0.7	1.2	+73%	
Pretax Income	0.7	2.1	+194%	
Posttax Income	0.8	2.8	+240%	70%
Global Bottom 50% Average Income ($/Day)				
Pretax Income Absent Educational Expansion	1.3	2.2	+79%	
Pretax Income	1.3	3.8	+205%	
Posttax Income	1.4	5.0	+264%	70%

Notes. The table compares the evolution of global poverty and the average income of the global bottom 20% and bottom 50% under three scenarios. The first one considers the evolution of each indicator if there had been no educational progress since 1980 ("pretax income absent educational expansion"). The second one corresponds to the actual evolution of each indicator in terms of pretax income ("pretax income"). The third one corresponds to the actual evolution of each indicator in terms of posttax income, that is, after removing all taxes and adding all cash and in-kind transfers (see Gethin, 2023b). The last column displays the corresponding share of global poverty reduction or real income gains that can be attributed to public policies, combining direct redistribution (moving from pretax to posttax income) and indirect investment benefits from education (moving from "pretax income absent educational expansion" to pretax income), calculated as one minus the ratio of the first row to the third row of the fourth column within each panel. Global poverty rate and real incomes expressed in 2017 PPP USD. Estimates of the world distribution of income constructed from World Bank per-capita income and consumption data. See table 1.8 for comparable results using data from the World Inequality Database.

A.1.2 Additional Results: Global Gender Inequality

Figure A.9: Female Labor Income Share, 1991 vs. 2019

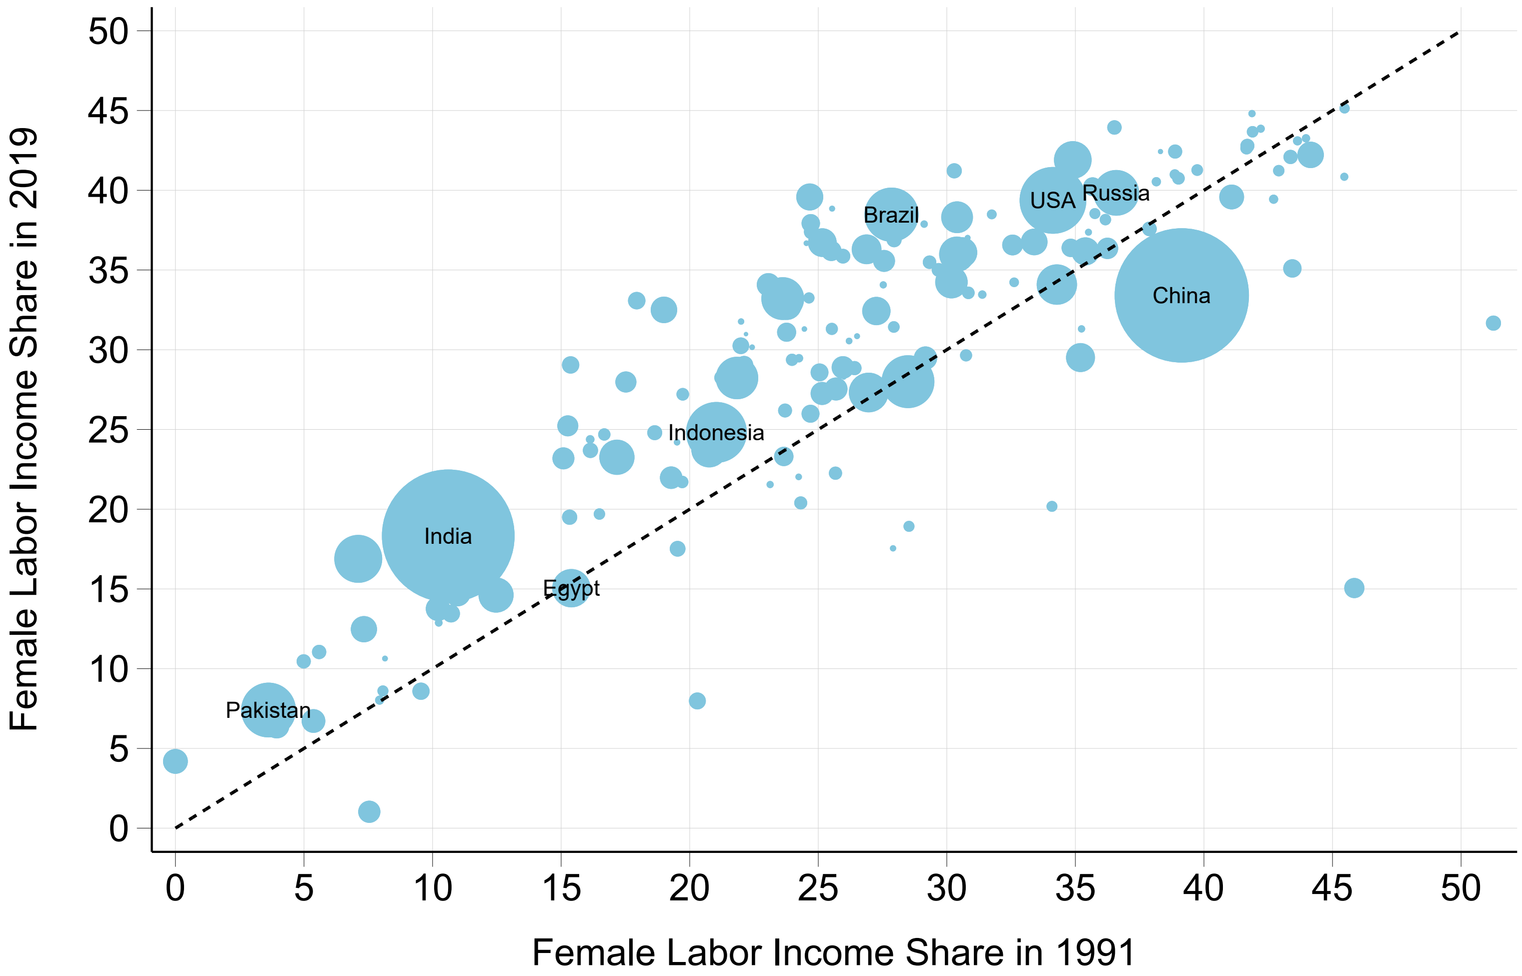

Notes. Data from Neef and Robilliard (2021).

Figure A.10: Returns to Schooling: Men vs. Women

Notes. Author's computations using labor force survey microdata.

Figure A.11: Contribution of Schooling to Gender Inequality Reduction by World Region
Gender Ratio of Income Gains from Schooling (Women / Men), 1991-2019

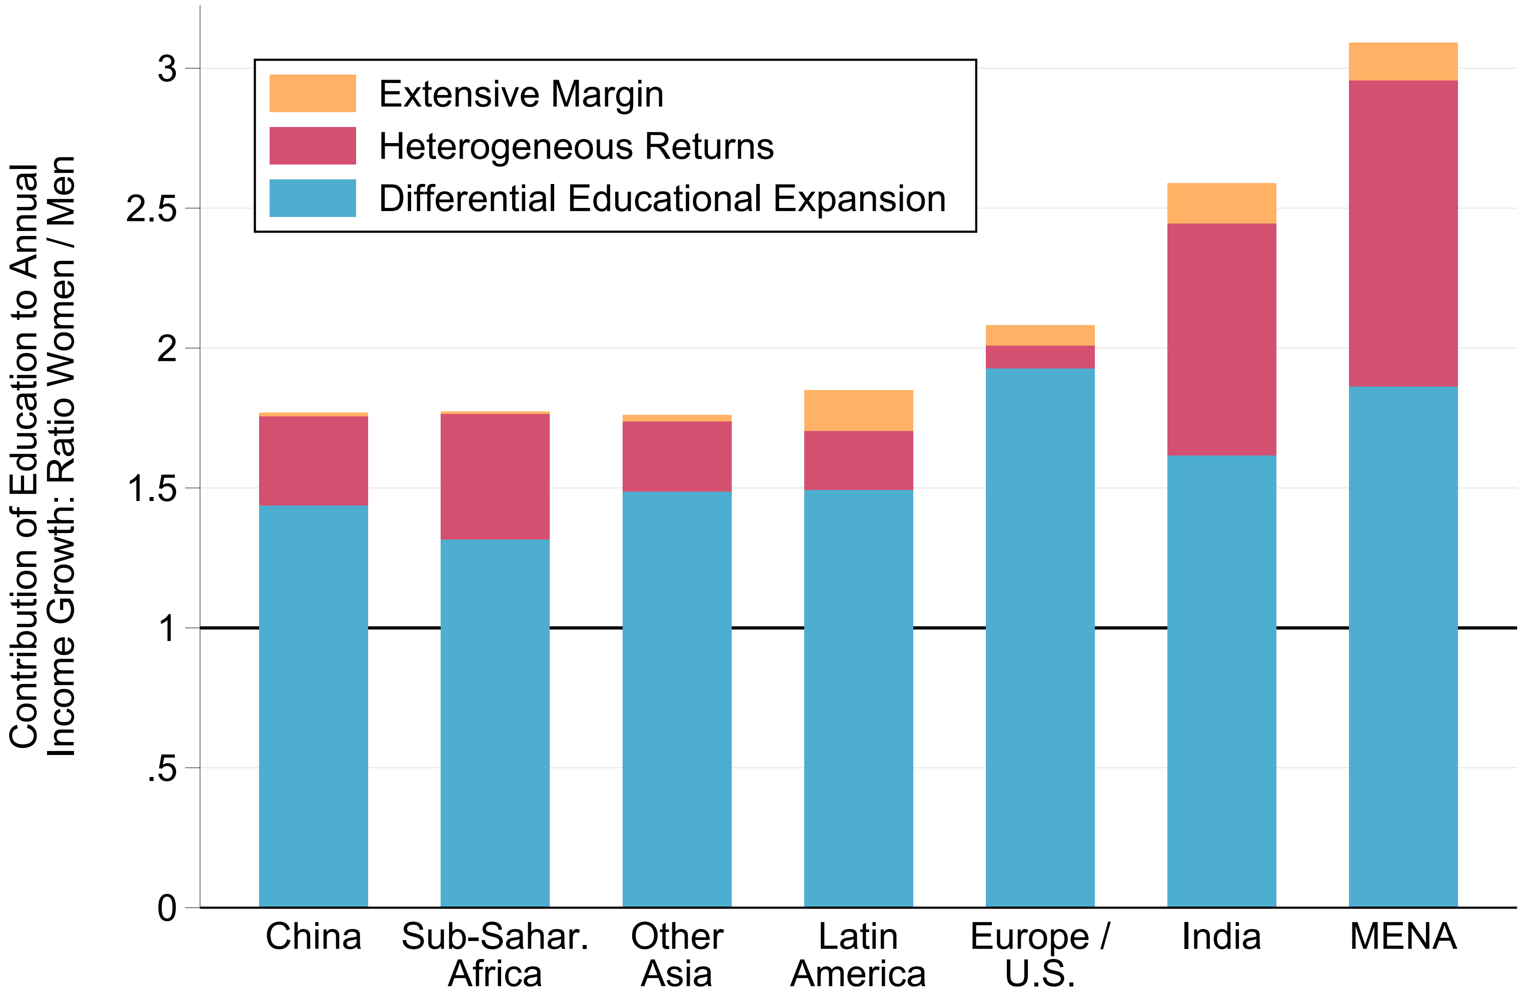

Notes. Author's computations using labor force survey microdata. Each bar corresponds to the ratio of gains from schooling among women over gains from schooling among men. Gains from schooling correspond to the percent difference between actual income and counterfactual income absent educational expansion. Population-weighted averages of gains from schooling estimated in each country.

Figure A.12: Annualized Income Gains from Schooling, 1991-2019: Men vs. Women

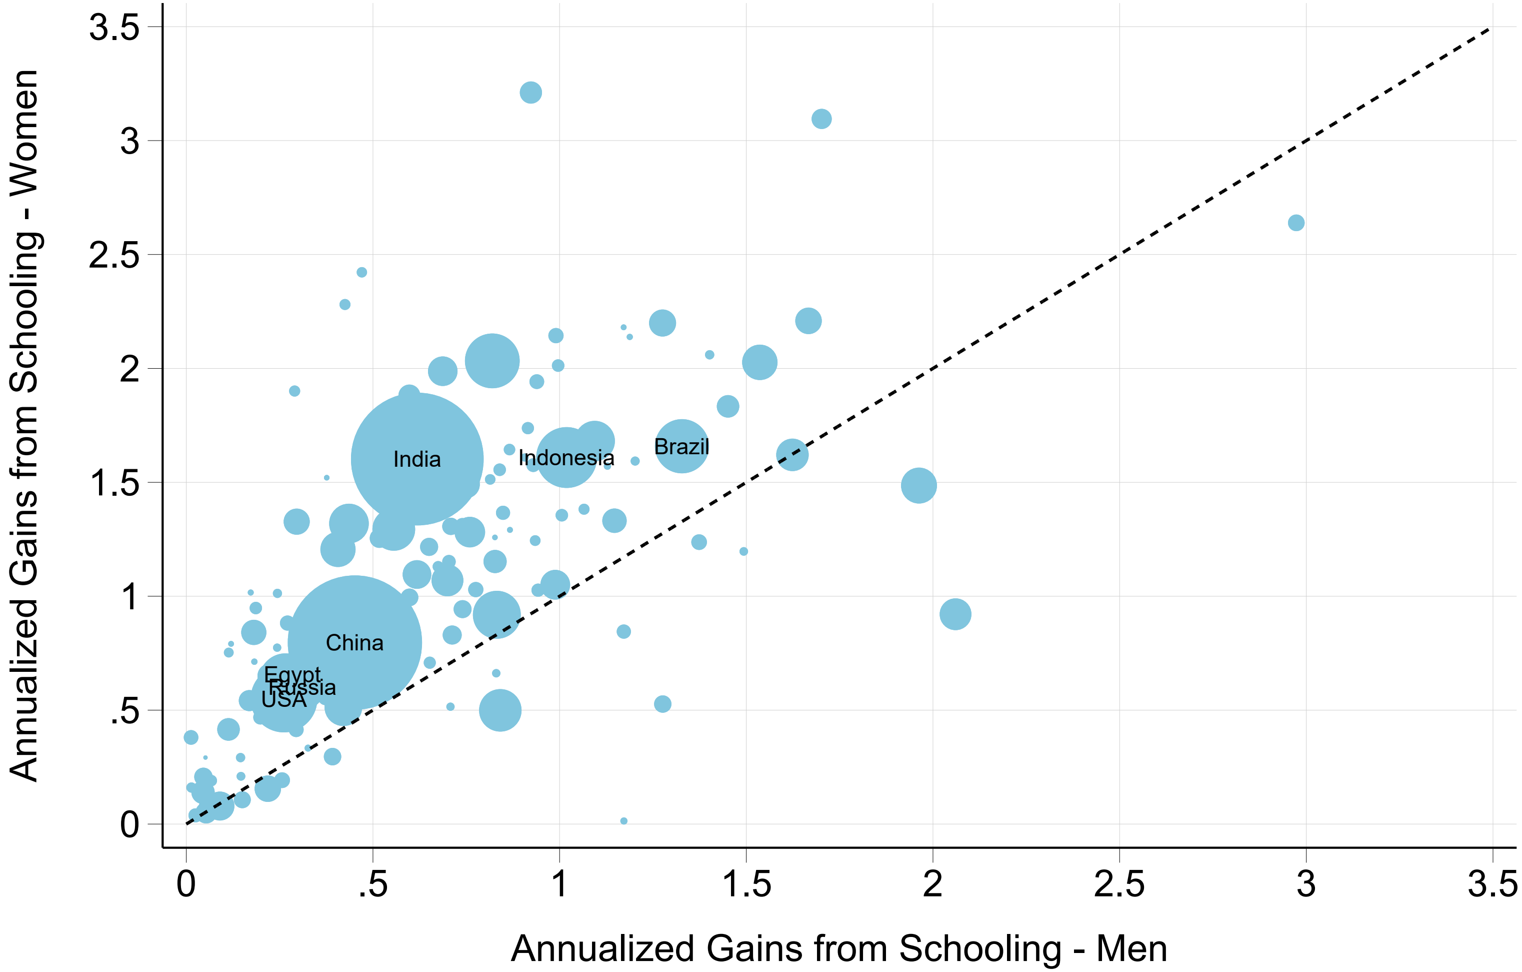

Notes. Author's computations using labor force survey microdata.

Figure A.13: Education and Gender Inequality: Actual vs. Counterfactual Female Labor Income Share in 2019

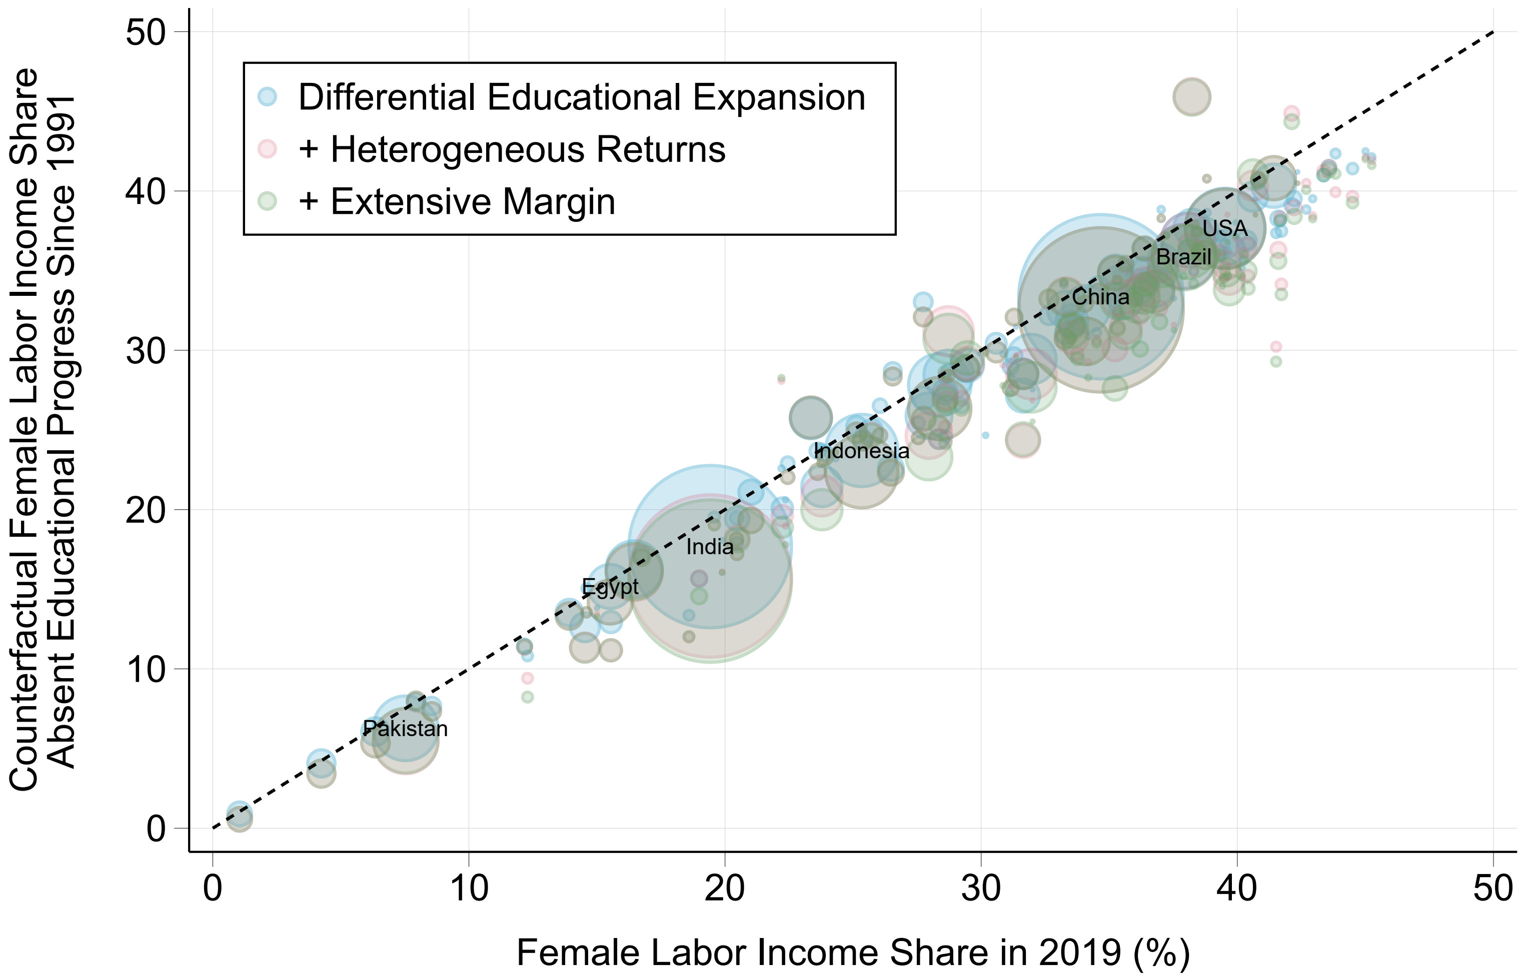

Notes. Author's computations using labor force survey microdata.

Table A.13: Effect of an Additional Year of Schooling on Female Labor Force Participation: Selected Causal Estimates

Source	Country	Level	β	SE	Baseline	Δ (%)
Akresh, Halim, and Kleemans (2023)	Indonesia	Primary	0		21%	+0%
Khan (2021)	Pakistan	Primary	0		25%	+0%
Grépin and Prashant (2015)	Zimbabwe	Secondary	3	1.7	11%	+27%
Delesalle (2021)	Tanzania	Primary	3.9	1.3		
Cui, Liu, and Zhao (2019)	China	Secondary	4.8	2.7	21%	+23%
Erten and Keskin (2018)	Turkey	Secondary	5		14%	+36%
Spohr (2003)	Taiwan	Secondary	5.8	3.2	45%	+13%
Keats (2018)	Uganda	Primary	7.3	3.1	79%	+9%
Elsayed and Shirshikova (2023)	Egypt	Tertiary	8		31%	+26%
Oliobi (2022)	Nigeria	Tertiary	8.7	2.3	65%	+13%
Chicoine (2021)	Ethiopia	Primary	9.3	5.8	33%	+28%
Hicks and Duanc (2023)	Jordan	Secondary	9.6	3.1	31%	+31%
Kim (2023)	Korea	Tertiary	9.8	2.7	41%	+24%
Overall Average			5.8		35%	+19%

Notes. The table reports estimates of the impact of increasing women's education by one year on female labor force participation (FLFP), based on studies relying on various natural experiments generating quasi-random variation in access to schooling. β: estimated effect of an additional year of schooling on labor force participation. SE: standard error. Baseline: overall labor force participation of women for the sample and definition of employment considered. Δ (%): corresponding percent increase in labor force participation per year of schooling.

Table A.14: Education and Global Gender Inequality, 1991-2019: Excluding China

	1991	2019	Diff.	Share Explained By Education	Share Explained (Cross-Country Average)
Global Female Labor Income Share	28.7%	31.7%	3.1		
Counterfactual: No Educational Progress	28.7%	30.1%	1.5	52%	38%
Counterfactual: + Heterogeneous Returns	28.7%	29.8%	1.1	63%	48%
Counterfactual: + Extensive Margin	28.7%	29.5%	0.9	71%	49%

Notes. The table reports actual versus counterfactual global female labor income shares under different assumptions. China is excluded from the analysis. Global female labor income: total share of labor income received by women in the world as a whole. Change in education: only account for differential trends in schooling by gender, applying the same returns to schooling for men and women to build the counterfactual. Heterogeneous returns: account for differential returns by gender. Extensive margin: account for differential effects of schooling on employment by gender. Cross-country average: population-weighted average of the share of gender inequality reduction explained by education in each country.

Table A.15: Education and Global Gender Inequality, 1991-2019:
Average Country, Excluding Countries With Rising Gender Inequality

	1991	2019	Diff.	Share Explained By Education
Actual Female Labor Income Share	20.9%	27.3%	6.4	
Counterfactual: No Educational Progress	20.9%	25.8%	4.9	24%
Counterfactual: + Heterogeneous Returns	20.9%	25.1%	4.2	35%
Counterfactual: + Extensive Margin	20.9%	24.8%	3.9	39%

Notes. The table reports actual versus counterfactual female labor income shares under different assumptions. Figures correspond to the population-weighted average of all countries in the world, excluding all countries where the female labor income share declined, but keeping countries in which educational expansion increased gender inequality. Change in education: only account for differential trends in schooling by gender, applying the same returns to schooling for men and women to build the counterfactual. Heterogeneous returns: account for differential returns by gender. Extensive margin: account for differential effects of schooling on employment by gender.

Table A.16: Education and Global Gender Inequality, 1991-2019:
By Specification of Employment Effects

	Share Explained (World)	Share Explained (Cross-Country Average)
No Employment Effect	71%	58%
Benchmark: OLS Employment Effects	78%	59%
Alternative: +4pp. per Year	97%	68%
Alternative: +6pp. per Year	111%	73%
Alternative: +8pp. per Year	124%	76%
Alternative: +15% per Year	117%	70%
Alternative: +20% per Year	130%	73%
Alternative: +25% per Year	143%	75%

Notes. The table reports the share of the decline in gender inequality that can be explained by education, focusing on the global female labor income share (second column) and the average country (third column; population-weighted), depending on assumptions made on the impact of education on female labor force participation. OLS employment effects: effects of schooling on employment estimated by OLS in each country. Alternative: uniform effect of schooling on female labor force participation, either in terms of percentage points or in terms of relative increases in employment, corresponding to the range of quasi-experimental estimates reported in table A.13.

Table A.17: Education and Global Gender Inequality: By Region and Time Period

	Gains from Schooling Ratio Women / Men			Share of Gender Inequality Reduction Explained by Education		
	1991-2019	2000-2019	2010-2019	1991-2019	2000-2019	2010-2019
China	1.8	2.2	2.0	¿100%	¿100%	¿100%
Europe / U.S.	1.9	1.5	2.5	48%	33%	50%
India	2.6	4.8	4.3	48%	50%	68%
Latin America	1.8	1.9	2.1	45%	44%	44%
MENA	3.1	2.9	2.6	72%	73%	66%
Other Asia-Pacific	1.8	1.7	1.8	42%	59%	48%
Sub-Saharan Africa	1.8	1.8	1.8	54%	47%	44%
World Average	2.0	2.4	2.5	59%	59%	62%

Notes. The table reports relative gains from schooling by gender together with the share of gender inequality reduction explained by education for various world regions and time periods. Gains from schooling ratio: women-to-men ratio of annualized income gains from schooling. All numbers correspond to population-weighted cross-country averages of the corresponding indicators in each region. Estimates account for differential educational expansion, heterogeneous returns to schooling, and extensive margin effects.

A.1.3 Stylized Facts on Educational Attainment and the World Distribution of Income

Figure A.14: The Growing Importance of Income Inequality Within Countries
Theil Decomposition of Global Income Inequality, 1980-2019

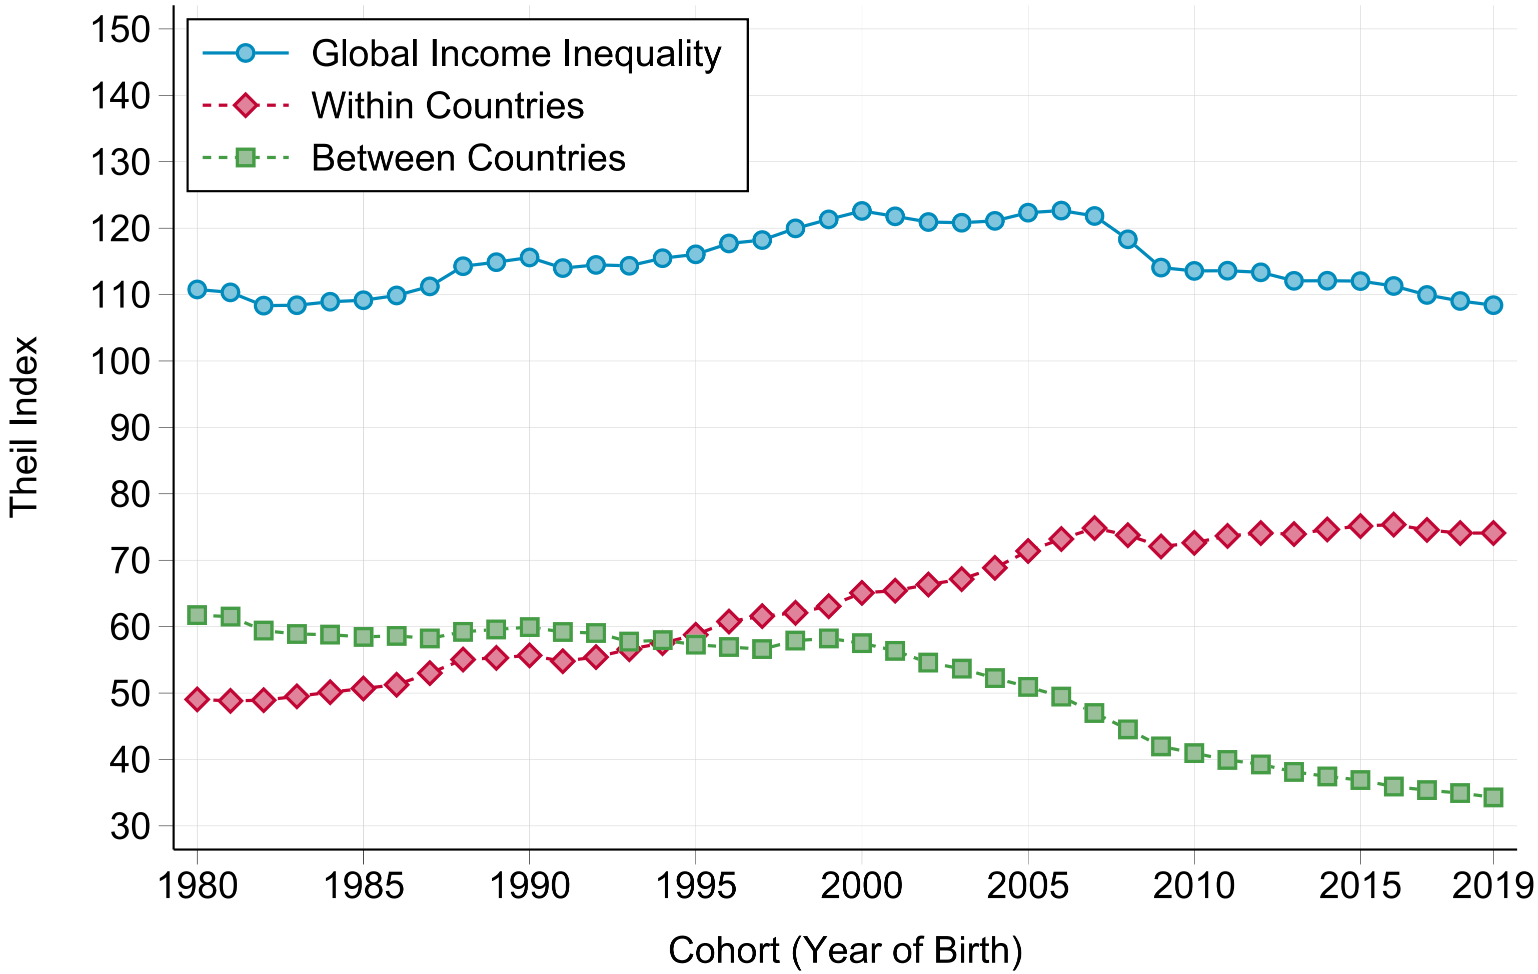

Source: Author's computations using data from the World Inequality Database. The figure plots the evolution of the Theil index of global income inequality from 1980 to 2019, as well as its decomposition into a between-country component and a within-country component.

Figure A.15: The Rise of Within-Country Inequality: Average Top 10% Pretax Income Share by World Region, 1980-2019

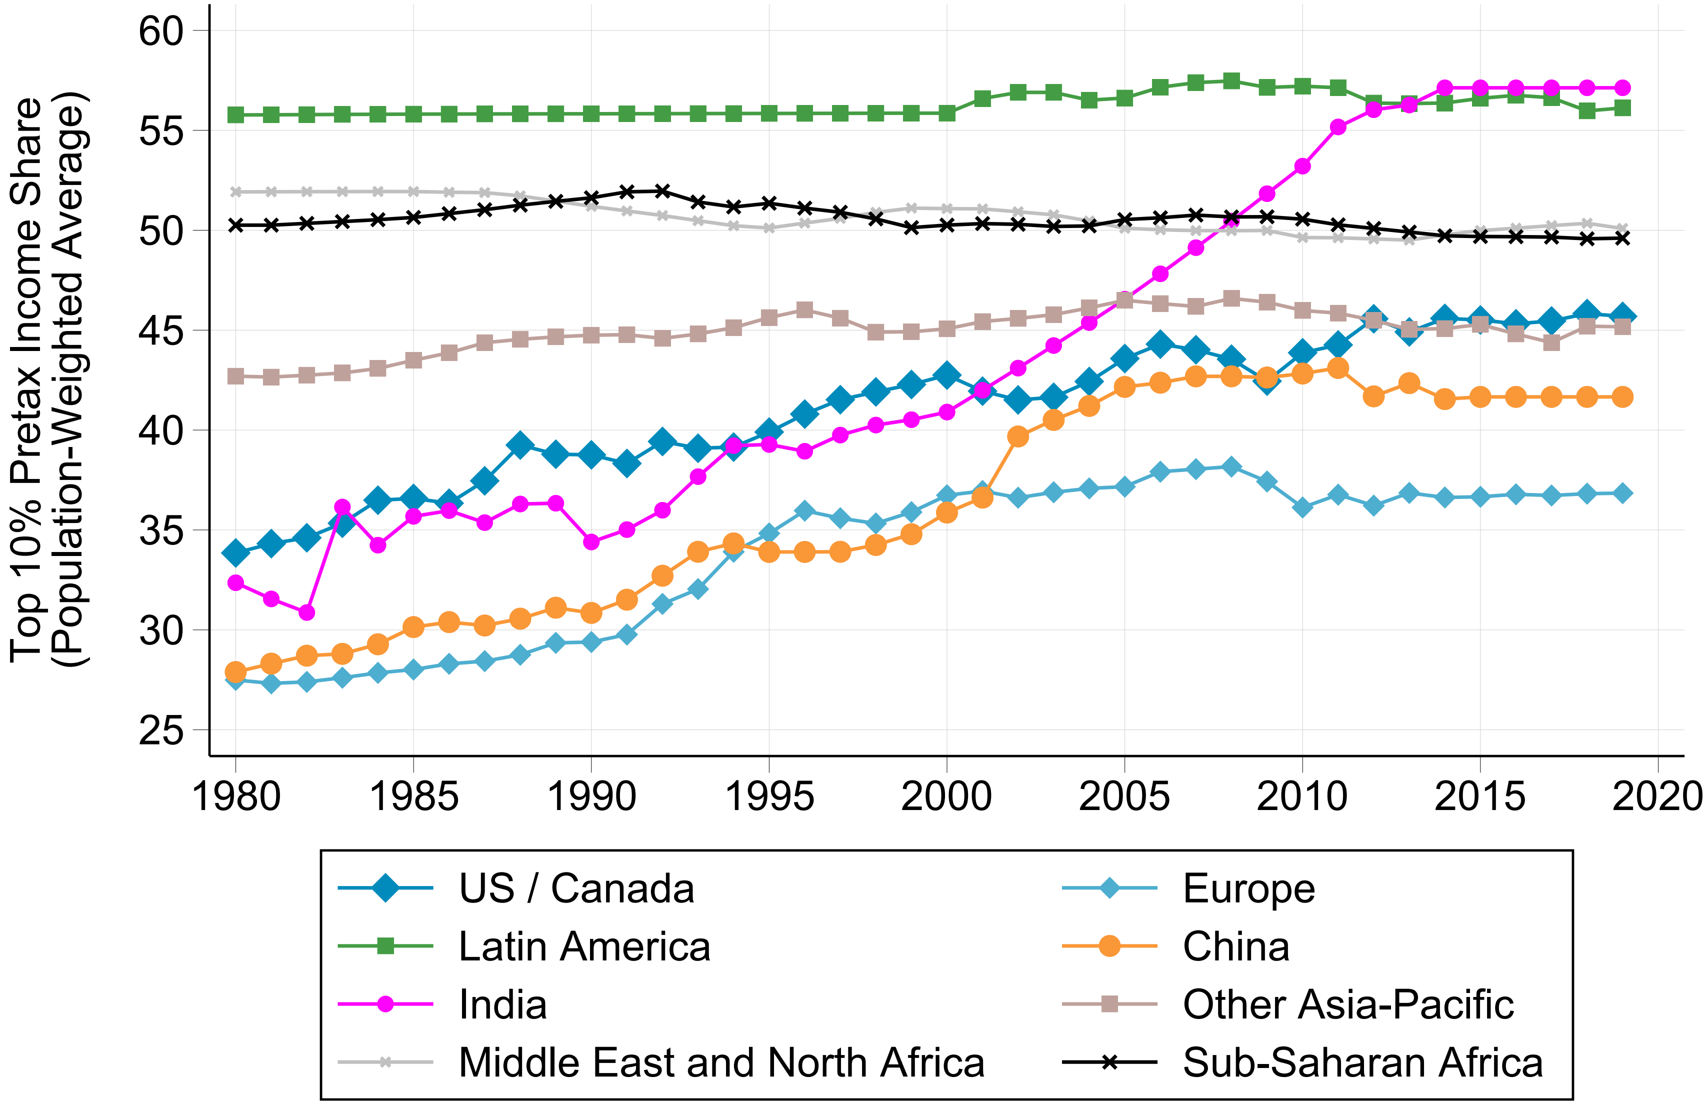

Notes. Author's computations using data from the World Inequality Database.

Figure A.16: Geographical Breakdown of Global Income Groups, 1980

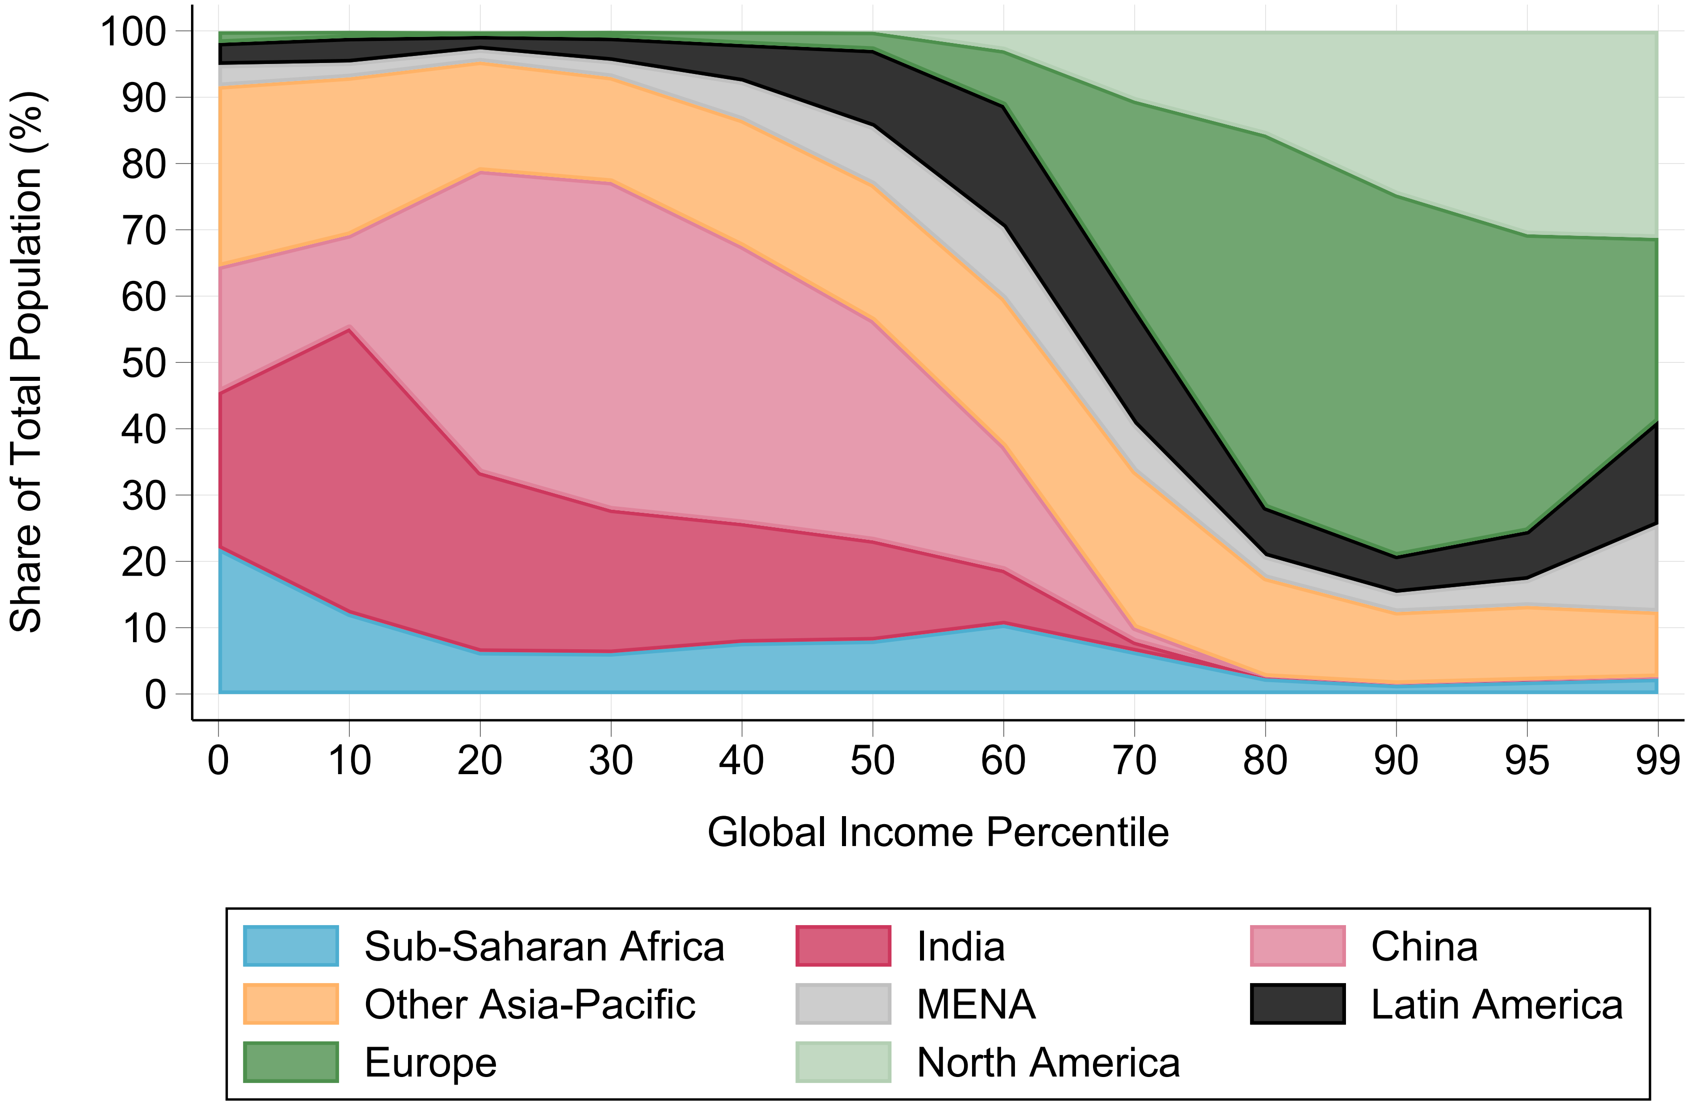

Notes. Author's computations using data from the World Inequality Database.

Figure A.17: Geographical Breakdown of Global Income Groups, 2019

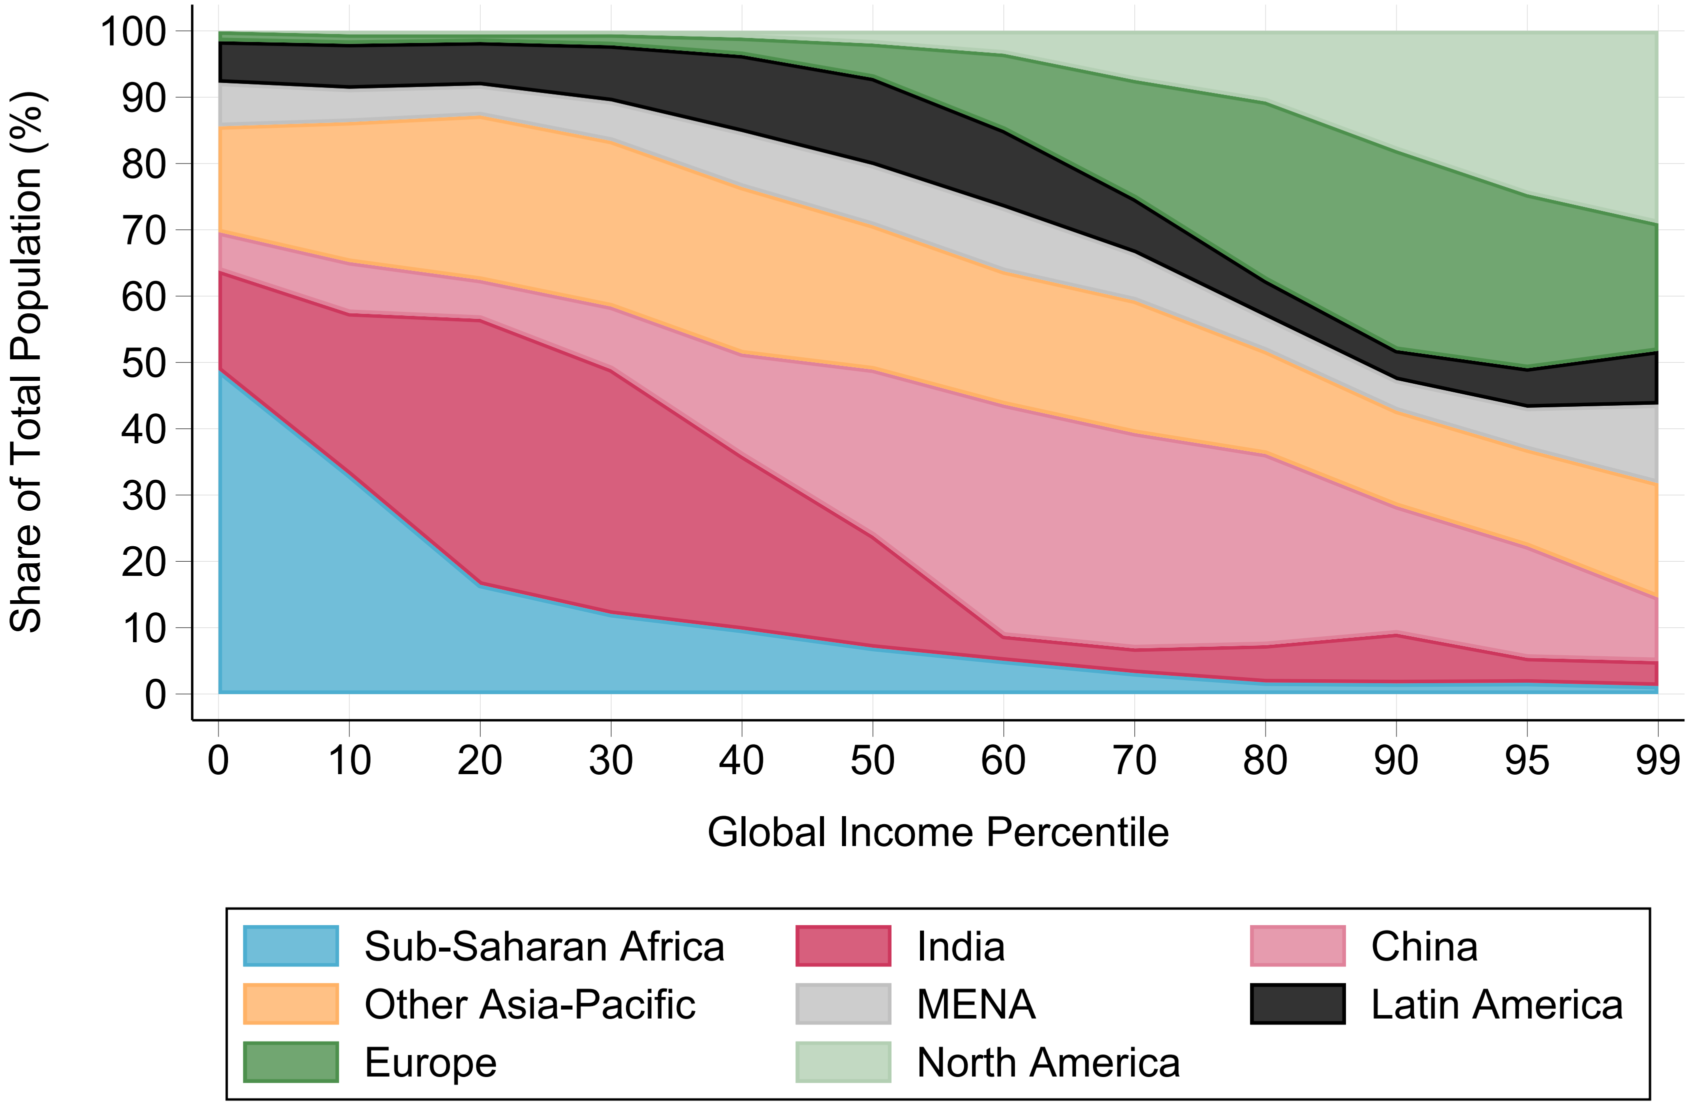

Notes. Author's computations using data from the World Inequality Database.

Figure A.18: Geographical Breakdown of the Global Bottom 20%, 1980-2019

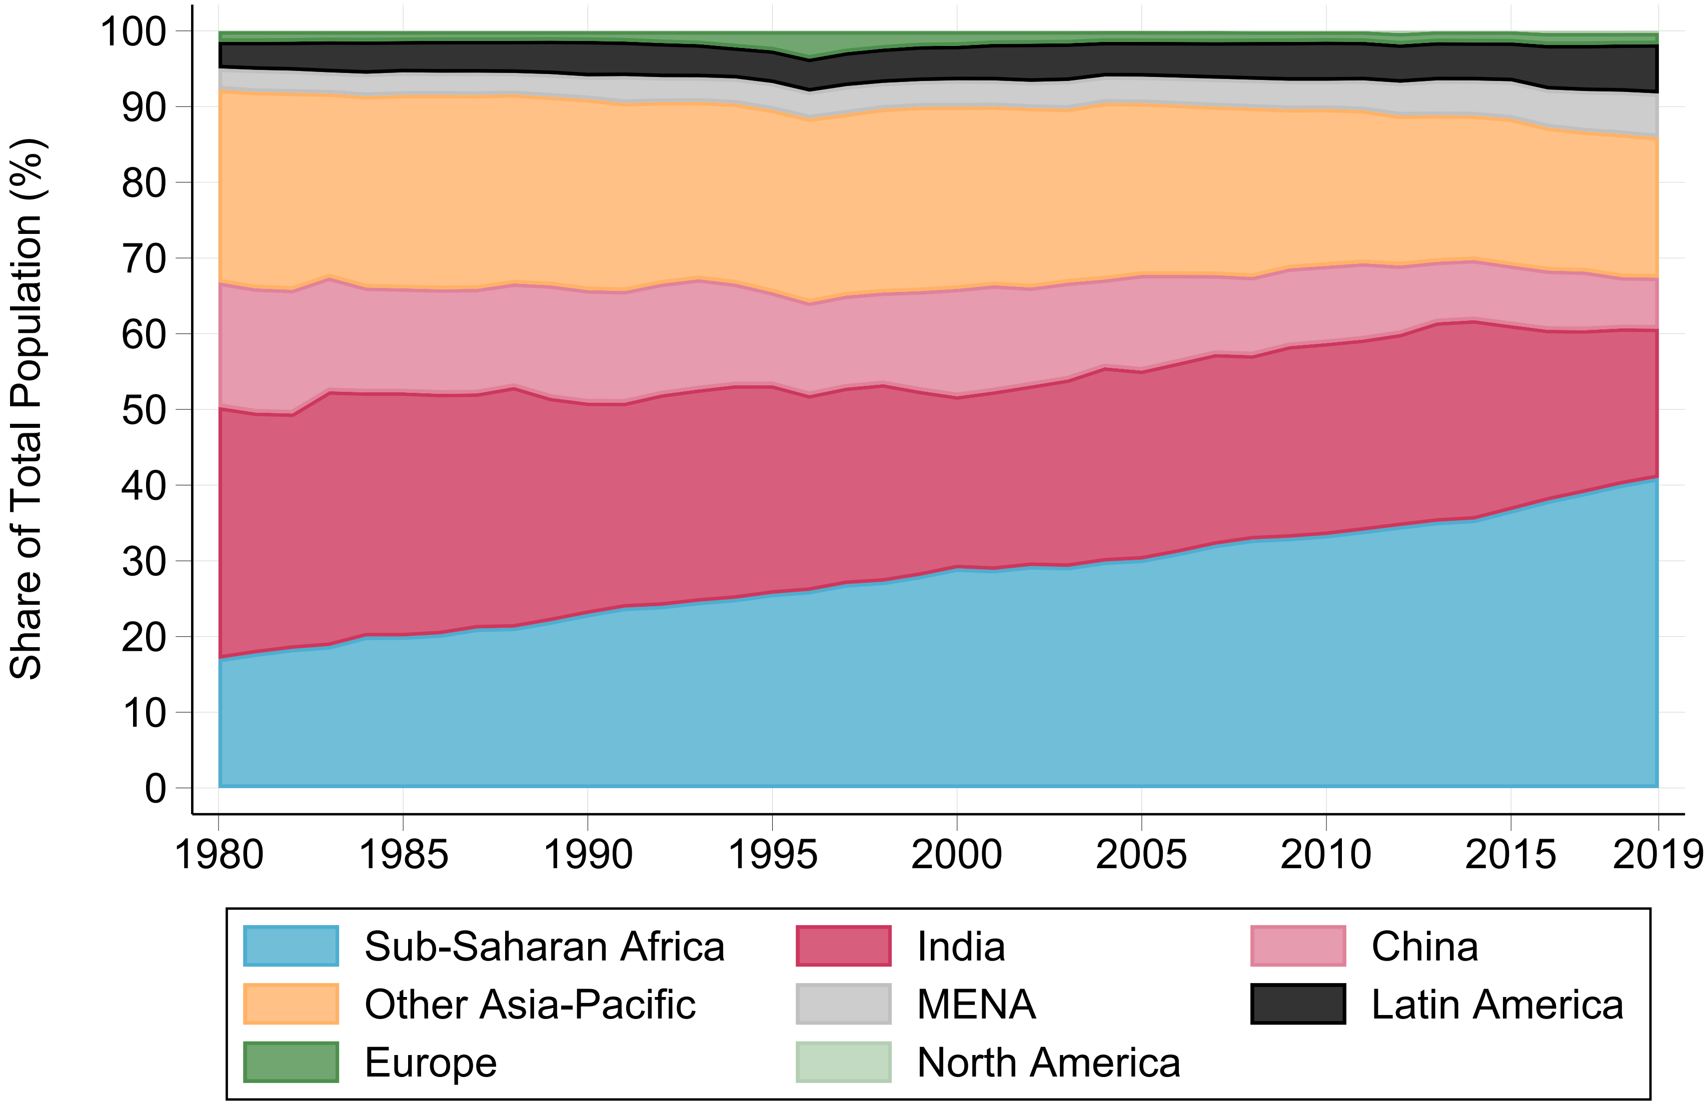

Notes. Author's computations using data from the World Inequality Database.

Figure A.19: Geographical Breakdown of the Global Top 20%, 1980-2019

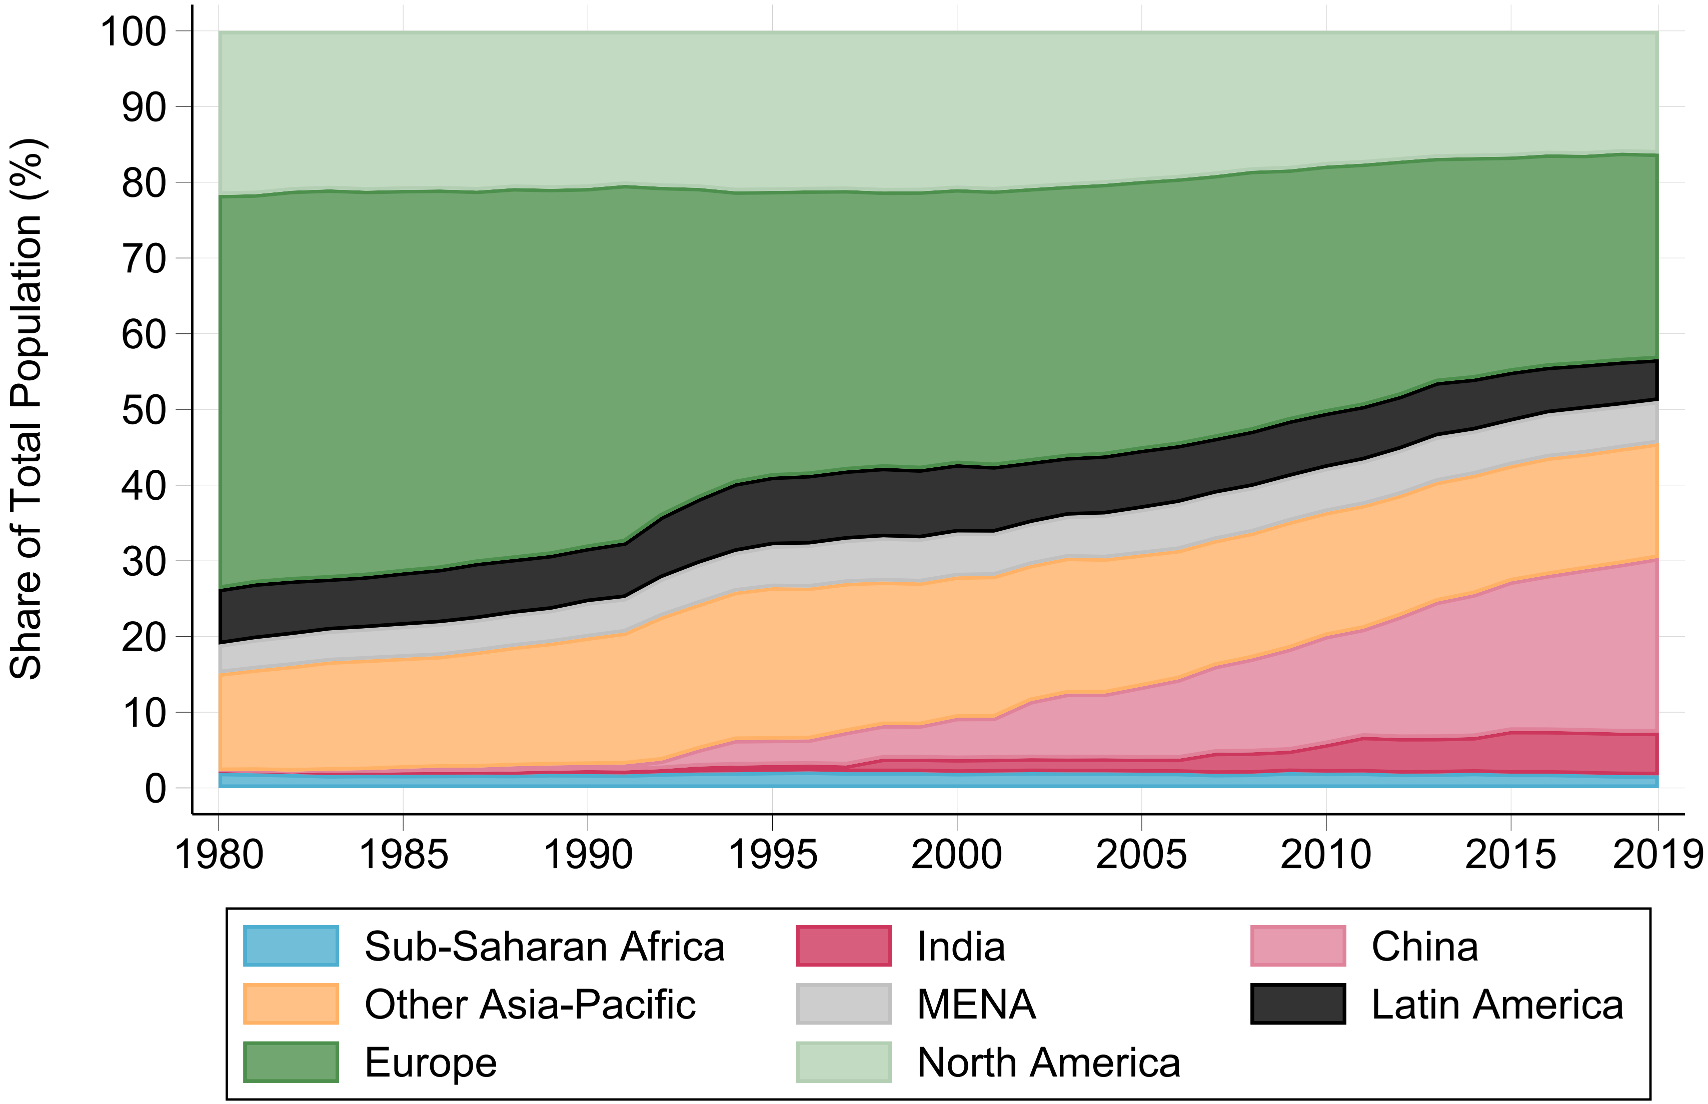

Notes. Author's computations using data from the World Inequality Database.

Figure A.20: The Concentration of Capital Income
in the United States, Latin America, and South Africa

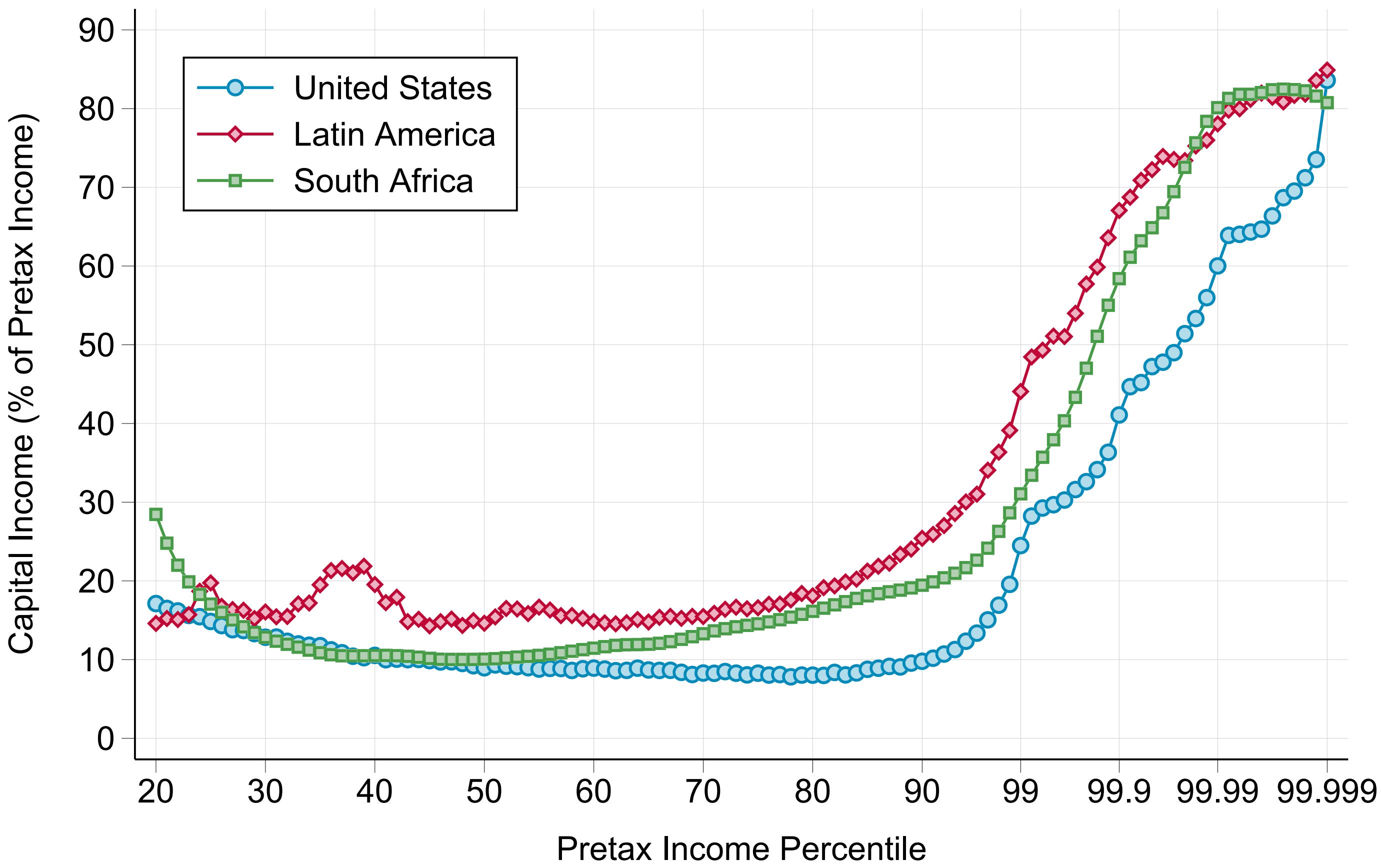

Source: author's elaboration combining data from Piketty, Saez, and Zucman (2018) for the United States, De Rosa, Flores, and Morgan (2022b) for Latin America, and Chatterjee, Czajka, and Gethin (2023) for South Africa.

Figure A.21: The Decline of Global Educational Attainment Inequality
Theil Decomposition of Global Human Capital Inequality, 1980-2019

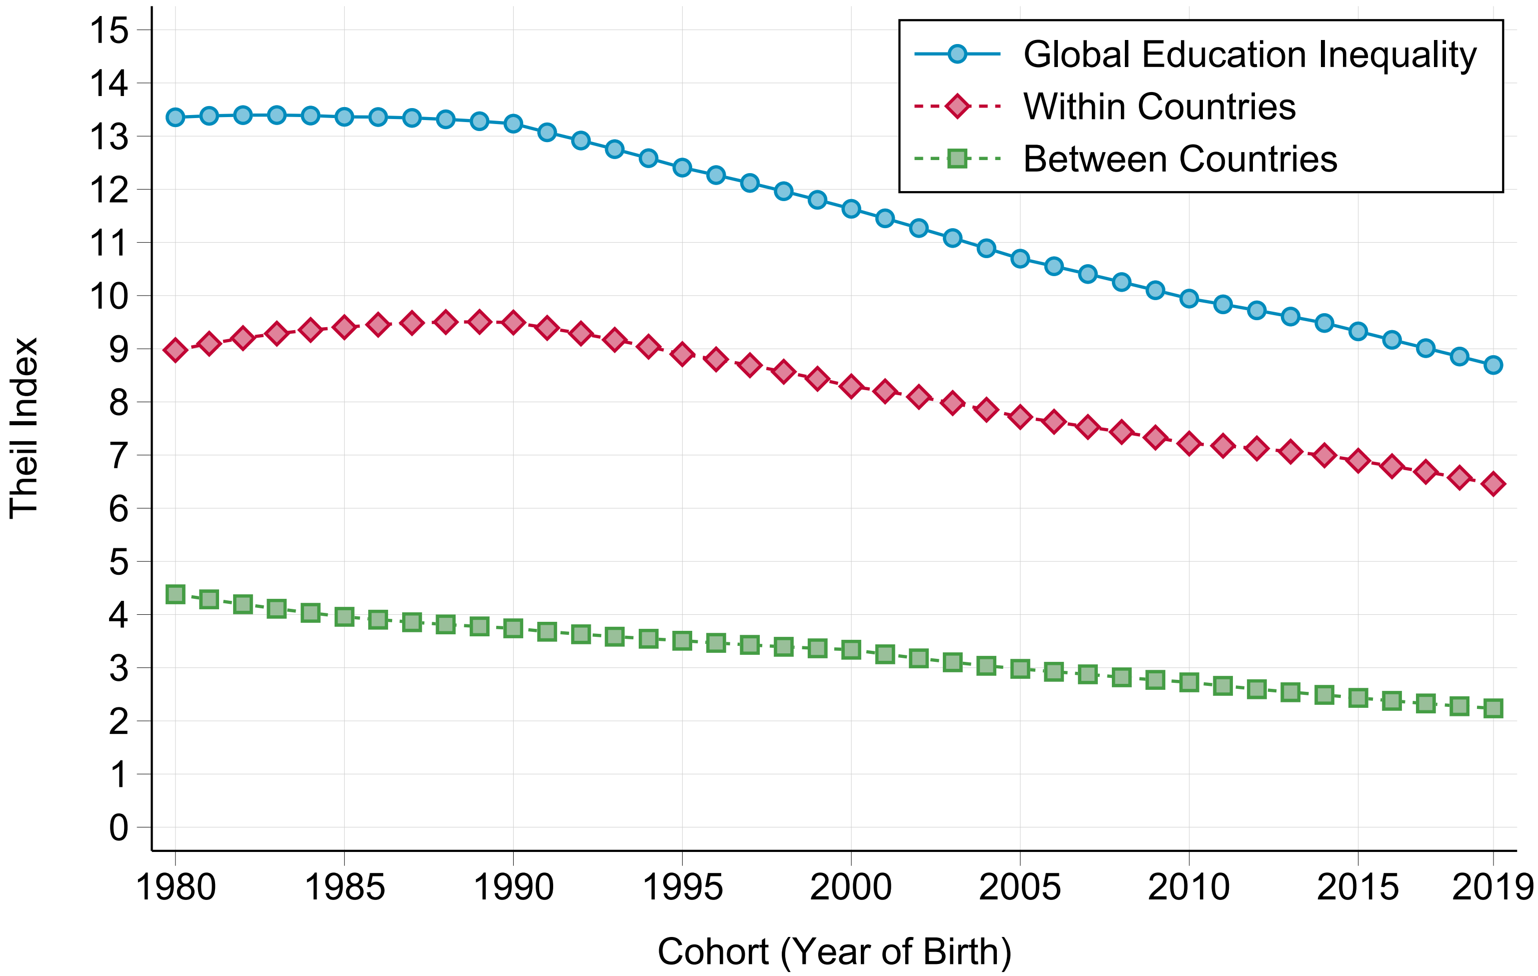

Source: Author's computations combining data from Barro and Lee (2013), updates, and other sources. The figure plots the evolution of the Theil index of global human capital inequality, as well as its decomposition into a between-country component and a within-country component. Human capital at time t in country c with average years of schooling s_{ct} is computed as $H_{ct} = e^{rs_{ct}}$, with r the returns to schooling (set to 10%): see Morrisson and Murtin (2013). This indicator is analogous to the standard deviation of years of schooling.

Figure A.22: Average Years of Schooling by World Region

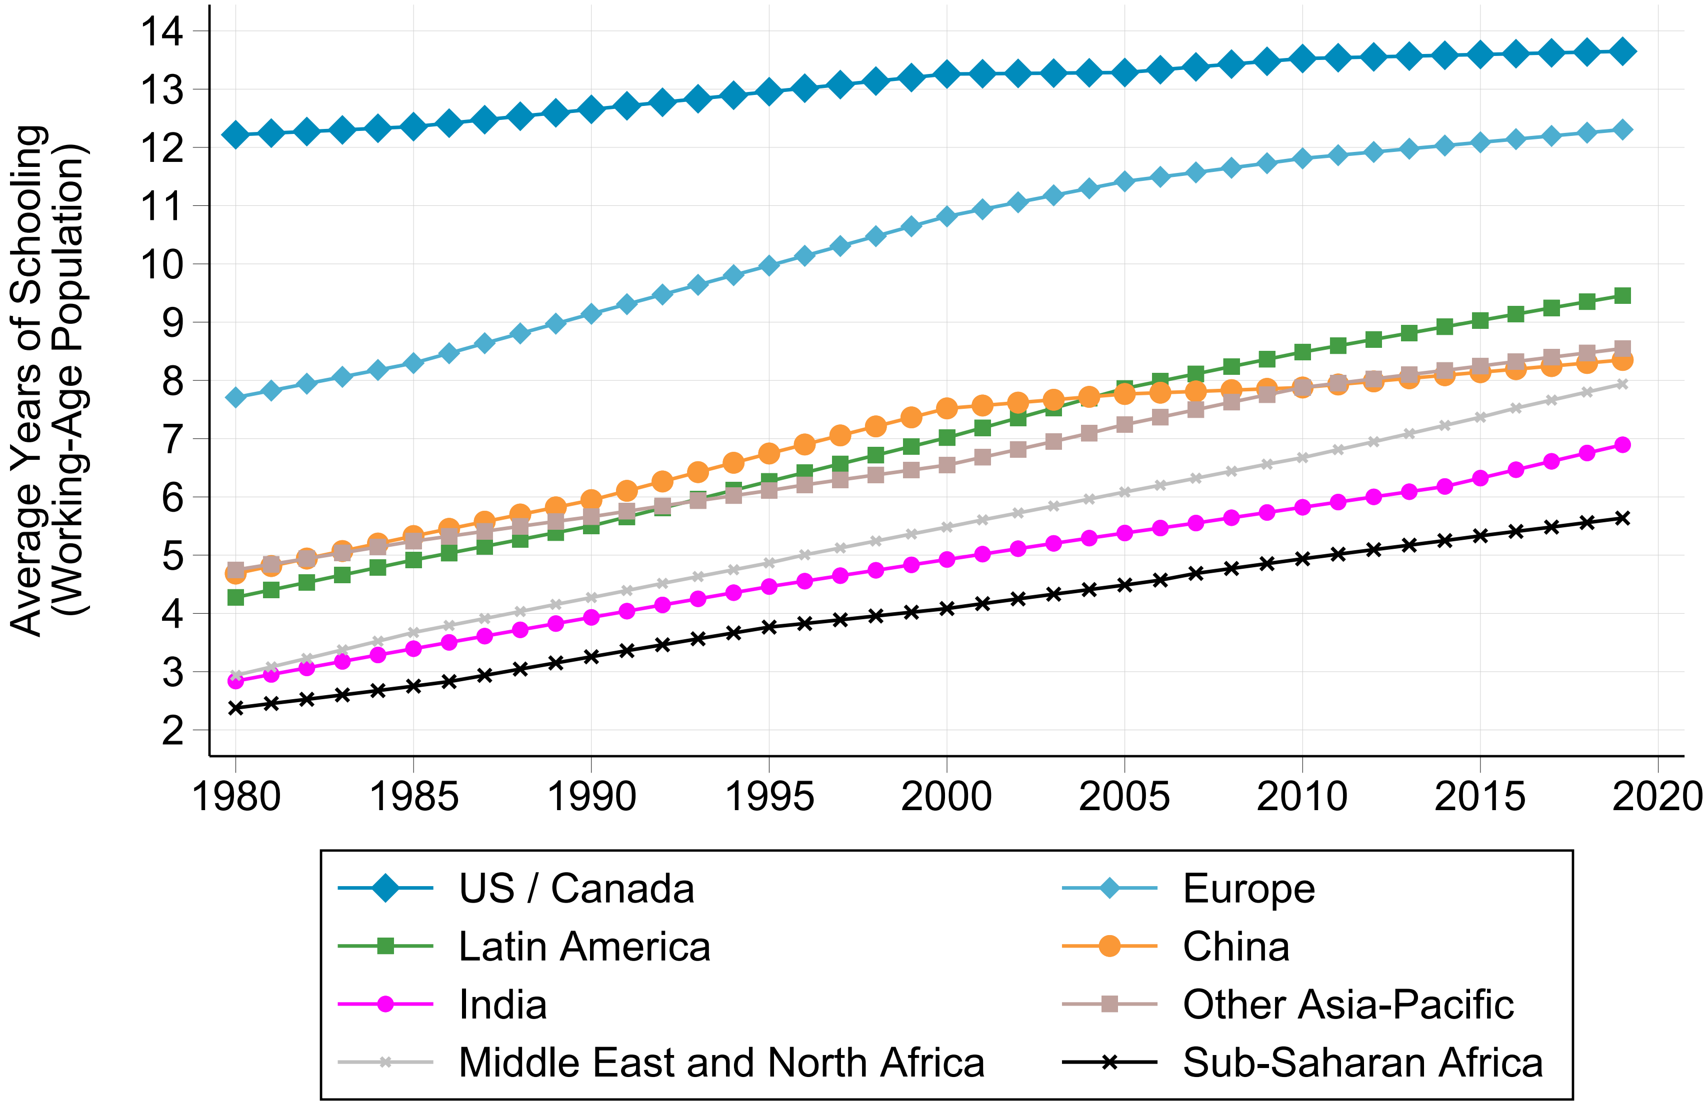

Notes. Author's computations combining data from Barro and Lee (2013), updates, and other sources.

A.2 Estimation Details

A.2.1 Estimation Steps

Sample Restriction The starting point is individual-level data on wages and education. In each survey, I keep individuals aged 25 to 65, with no missing information on age, gender, or education, and with positive reported income.

1) Downgrade Education Levels Next, I match the microdata with information on the distribution of educational attainment by age-gender cell in 1980, covering four education levels: no schooling, incomplete or complete basic education, incomplete or complete secondary education, and incomplete or complete tertiary education. To move from observed educational attainment to counterfactual educational attainment, I randomly sample individuals and downgrade their education levels, until matching 1980 totals by age-gender-level cell.

Individuals belonging to closest education categories are given priority in the simulation. For instance, if 20% of individuals in a given age-gender cell had no schooling in 1980, compared to 10% today, I randomly sample 10% of individuals among the primary education group and downgrade their education level to no schooling. When closest education levels do not contain enough individuals (e.g., only 5% of individuals in this survey have primary education), I instead sample individuals from the category above (secondary education in this example). The outcome is a modified survey, in which the distribution of education levels by age-gender cell corresponds to that observed in 1980. This survey contains "unaffected" observations, corresponding to individuals with unchanged education, as well as "treated" individuals whose education has been downgraded. This approach is very similar to the one recently adopted by Hershbein, Kearney, and Pardue (2020) to estimate the distributional effects of expanding access to college in the United States.

2) Reduce Wages Using Returns to Schooling The second step is to reduce the income of "treated" individuals using estimates of returns to schooling. More precisely, consider an individual i with income y_i, whose education level is downgraded from s_2 to s_1. Denote $r_{s_1,s_2} = \ln(w_1) - \ln(w_2)$ the returns to schooling of moving from s_1 to s_2. Then, the counterfactual income of individual i is:

$$\tilde{y}_i = \exp[\ln(y_i) - r_{s_1,s_2}] \tag{A.1}$$

I use separate estimates of returns to primary education $r_{non,pri}$, returns to secondary education $r_{pri,sec}$, and returns to tertiary education $r_{sec,ter}$. In the main analysis, these returns are estimated in the microdata using a modified Mincerian equation of the form:

$$\ln y_i = \alpha + \beta^{pri} D_i^{pri} + \beta^{sec} D_i^{sec} + \beta^{ter} D_i^{ter} + X_i \beta + \varepsilon_i \tag{A.2}$$

Which implies that:

$$r_{non,pri} = \beta^{pri} \tag{A.3}$$
$$r_{pri,sec} = \beta^{sec} - \beta^{pri} \tag{A.4}$$
$$r_{sec,ter} = \beta^{ter} - \beta^{sec} \tag{A.5}$$

For individuals whose education is downgraded by several levels, I use the corresponding cumulative returns. For instance, an individual downgraded from secondary education to no schooling will have a counterfactual income given by $\tilde{y}_i = \exp[\ln(y_i) - r_{sec,pri} - r_{pri,non}]$.

As explained in section 1.2, I consider lower and upper bounds for returns to schooling. The lower bound corresponds to returns to schooling in 2019, estimated using a modified Mincerian equation. Hence, it corresponds to returns to schooling prevailing under the current distribution of educational attainment:

$$\underline{\mathrm{r}}_{s_1,s_2} = r_{s_1,s_2}(L) \tag{A.6}$$

With $L = (L_1, ..., L_m)$ the distribution of educational attainment in 2019. In contrast, the upper bound corresponds to counterfactual returns to schooling that would prevail if, all other things equal, education levels were to come back to their 1980 levels:

$$\bar{r}_{s_1,s_2} = r_{s_1,s_2}(\tilde{L}) \tag{A.7}$$

With $\tilde{L} = (\tilde{L}_1, ..., \tilde{L}_m)$ the distribution of educational attainment in 1980. These counterfactual returns to schooling are by construction not observed and have to be estimated. Assuming that the production technology is CES and that we know the elasticity of substitution σ_{s_1,s_2} between s_1 and s_2, the upper bound can be calculated as:

$$\bar{r}_{s_1,s_2} = r_{s_1,s_2}(L) - \frac{1}{\sigma_{s_1,s_2}} \Delta \ln \left(\frac{L_2}{L_1} \right) \tag{A.8}$$

With $\Delta \ln\left(\frac{L_2}{L_1}\right) = \ln\left(\frac{L_2}{L_1}\right) - \ln\left(\frac{\tilde{L}_2}{\tilde{L}_1}\right)$ the change in the relative supply of skilled workers between 1980 and 2019. An increase in educational attainment from 1980 to 2019 will translate into a decrease in returns to schooling, which implies that the counterfactual return absent educational progress would be higher than the one observed: $\bar{r}_{s_1,s_2} > \underline{r}_{s_1,s_2}$.

In practice, I use the three nests of the production function to calculate three counterfactual returns to primary education, secondary education, and tertiary education:

$$\bar{r}_{non,pri} = r_{non,pri}(L) - \frac{1}{\sigma_3}\Delta \ln\left(\frac{L_{pri}}{L_{non}}\right) \tag{A.9}$$

$$\bar{r}_{pri,sec} = r_{pri,sec}(L) - \frac{1}{\sigma_2}\Delta \ln\left(\frac{L_{sec}}{L_{pri}}\right) \tag{A.10}$$

$$\bar{r}_{sec,ter} = r_{sec,ter}(L) - \frac{1}{\sigma_1}\Delta \ln\left(\frac{L_{sec}}{L_{ter}}\right) \tag{A.11}$$

Finally, using the CES production function, it is possible to recover the relative weight that should be put on counterfactual versus observed returns to obtain the true return to schooling (see section 1.1.2). Figure A.24 plots the empirical distribution of these weights across all countries for each of the three nests. The weight put on initial (counterfactual) returns mostly ranges from 0.5 to 0.7 for all three nests, with typical values in-between 0.55 and 0.65. The true return is thus close to the average of observed and counterfactual return, with a slightly greater weight given to the latter. Figures A.25, A.26, and A.27 display the corresponding observed versus true returns to primary, secondary, and tertiary education in each country.

3) Adjust Relative Wages The third step is to adjust relative wages of both unaffected and treated individuals, using the nested CES specification presented in section 1.2. This step of the estimation leaves the average income in the survey unchanged, since aggregate effects are captured in the previous step of the estimation. It then suffices to adjust relative wages using the three elasticities of substitution, while keeping average income constant. In practice, I do this in three steps in the microdata.

First, I reduce the log average income of primary-educated workers by the product of the primary/no schooling supply shift and the elasticity of substitution σ_3, and readjust earnings within this nest ($L_{\bar{s}ec}$) to leave the average unchanged:

$$\Delta \log\left(\frac{w_{pri}}{w_{non}}\right) = -\frac{1}{\sigma_3}\Delta \log\left(\frac{L_{pri}}{L_{non}}\right) \tag{A.12}$$

Second, I repeat the same procedure for the secondary/below-secondary nest ($L_{\bar{t}er}$):

$$\Delta \log\left(\frac{w_{sec}}{w_{s\bar{e}c}}\right) = -\frac{1}{\sigma_2}\Delta \log\left(\frac{L_{sec}}{L_{s\bar{e}c}}\right) \tag{A.13}$$

Third, I repeat the same procedure for the upper level of the CES production function:

$$\Delta \log\left(\frac{w_{ter}}{w_{t\bar{e}r}}\right) = -\frac{1}{\sigma_1}\Delta \log\left(\frac{L_{ter}}{L_{t\bar{e}r}}\right) \tag{A.14}$$

In a handful of countries, the share of the working-age population with no schooling or primary education declined to almost zero in 2019, leading some relative supply shifts to diverge to infinity. To avoid this divergence, I bound the absolute value of changes in relative supply to 4 log points. This does not affect the results significantly, given that concerned countries are those where the initial level of low-skilled workers was already very small.

4) Growth Accounting The final step is to use this counterfactual to estimate the share of growth explained by education. I first aggregate actual labor income Y_L and counterfactual labor income $\tilde{Y}_L$ in the survey microdata by decile, and calculate the corresponding ratio of counterfactual income to actual income: $\psi^d = \frac{\tilde{Y}_L^d}{Y_L^d}$. This yields a measure of how much lower labor income would be if education had not improved.

I then incorporate these estimates into global income distribution data. I start with distributions from the World Inequality Database, which provide information on the average pretax income of each generalized percentile (all percentiles from p0 to p99, followed by a further decomposition of top incomes up to p99.999p100). I merge estimates of ψ^d by country-year-decile with this database.[1] I then calculate

[1]To get smoother profiles of counterfactual income by generalized percentile, I assume that ψ^d for each decile corresponds to the ratios observed for p5, p15, p25, p35, p45, p55, p65, p75, p85, and p95. I then interpolate ψ^d between percentiles to fill in missing values. I assume that values observed for percentiles within the bottom 5% and the top 5% are those observed for p5 and p95, respectively. Finally, I drop generalized percentiles with zero average pretax income in the World Inequality Database.

counterfactual total pretax income of generalized percentile p as:

$$\tilde{Y}^p = Y_K^p + \psi^p Y_L^p \tag{A.15}$$

Finally, I construct separate actual and counterfactual world distributions of income from 1980 to 2019, by ranking all individuals in the world by each income concept and aggregating average income by global generalized percentile.

A.2.2 Aggregate and Individual Returns to Schooling: CES Simulations

A.2.2.1 Theoretical Background

The objective of this section is to shed light on the following question: which returns to schooling should be used to estimate the effect of educational expansion on total output? And how far is this return from the return observed before (initial return) versus after (final return) increasing education? Consider a CES production function with two skill types:

$$Y = \left(A_H L_H^{\frac{\sigma-1}{\sigma}} + A_L L_L^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}$$

We normalize $L_H + L_L = 1$, so that L_H corresponds to the share of skilled workers in the economy. The objective is to compare output under an initial distribution of skills $\{L_{H1}, L_{L1}\}$ and a final distribution of skills $\{L_{H2}, L_{L2}\}$, with $L_{H2} > L_{H1}$.

One possibility is to predict the change in output using initial returns to schooling r_1:

$$r_1 = \log\left(\frac{w_{H1}}{w_{L1}}\right) = \frac{\sigma-1}{\sigma}\log\left(\frac{A_H}{A_L}\right) - \frac{1}{\sigma}\log\left(\frac{L_{H1}}{L_{L1}}\right)$$

Predicted output can then be calculated as a weighted average of the wages of always skilled workers, always unskilled workers, and newly skilled workers:

$$Y = w_{H1}L_{H1} + w_{L1}L_{L2} + w_{H1}(L_{L1} - L_{L2})$$

Where $w_{H1} = \exp\left(\log(w_{L1}) + r_1\right)$ is the wage of high-skilled workers in the initial period. This amounts to considering that supply effects change relative wages, but do not reduce the effect of educational expansion on output. Initial returns to schooling then capture exactly the effect of skill upgrading. As demonstrated by Caselli and Ciccone (2013), this is an upper bound under standard assumptions on the production technology (which are satisfied in the CES case). This is because supply effects decrease the marginal product of skilled workers and increase that of unskilled workers, but the former effect dominates the latter.

An alternative possibility is to use final returns to schooling r_2:

$$r_2 \quad = \quad \log\left(\frac{w_{H2}}{w_{L2}}\right) = \frac{\sigma-1}{\sigma}\log\left(\frac{A_H}{A_L}\right) - \frac{1}{\sigma}\log\left(\frac{L_{H2}}{L_{L2}}\right)$$

Predicted output can then be calculated as:

$$Y = w_{H1}L_{H1} + w_{L1}L_{L2} + \exp\Bigg(\log(w_{L1}) + r_2\Bigg)\Bigg(L_{L1} - L_{L2}\Bigg)$$

This amounts to assuming that supply effects reduce the effect of educational expansion on output by the same amount as the decrease in returns to schooling. The benefits of skill upgrading then correspond exactly to the returns to schooling observed at the end of the period. This constitutes a lower bound on the actual effect of educational expansion, which may underestimate it significantly. In particular, consider the extreme case in which a large shock to the supply of skilled workers would bring returns to schooling to zero. This approach would then predict no change in output from educational expansion, which is impossible in this model as long as we assume that $A_H > A_L$.

A.2.2.2 Simulation

To investigate which weight should be put on final versus initial returns to schooling, I simulate predicted and actual changes in output under different parametrizations of the production function. More specifically, I run simulations jointly varying the elasticity of substitution σ from 1.5 to 8, the relative efficiency of skilled workers A_H/A_L from 1.2 to 5, and the final share of skilled workers L_{H2} from 0.15 to 0.95 (assuming an initial value $L_{H1} = 0.1$).

Given that all parameters are specified, it is also possible to calculate the actual individual return r^* that should be used to predict changes in output. This return satisfies:

$$
\begin{aligned}
Y &= w_{H1}L_{H1} + w_{L1}L_{L2} + \exp\left(\log(w_{L1}) + r^*\right)\left(L_{L1} - L_{L2}\right) \\
\Rightarrow r^* &= \log\left(\frac{Y - w_{H_1}L_{H1} - w_{L_1}L_{L1}}{L_{L1} - L_{L2}}\right) - \log(w_{L1})
\end{aligned}
$$

Finally, we also observe initial and final returns r_1 and r_2, which means that one can calculate the "optimal weight" γ that should be put on each return:

$$
\begin{aligned}
r^* &= \gamma r_1 + (1-\gamma)r_2 \\
\Rightarrow \gamma &= \frac{r^* - r_2}{r_1 - r_2}
\end{aligned}
$$

A weight of 0.5 means that the average of initial and final returns provides a good approximation for the true effect of educational expansion on output. A weight greater than 0.5 means that we should give more importance to initial returns in the estimation.

A.2.2.3 Results

Figure A.31 plots the resulting distribution of optimal weights on initial returns across all combinations of parameters. For elasticities ranging from 1.5 to 8 and relative skill efficiencies ranging from 1.2 to 5, the weight on initial returns ranges from about 0.45 to 0.8, with a mean of 0.65. Almost no estimate falls below 0.5, meaning that initial returns are almost always closer to the true effect of educational expansion on output than final returns.

Figure A.32 shows how the weight on initial returns varies across specific combinations of relative skill efficiencies and the elasticity of substitution. The weight is higher for lower elasticities of substitution and for higher levels of relative skill efficiencies. For a long-run elasticity of substitution of 4, the weight ranges from about 0.5 in case of very low differences in efficiencies to 0.7 for a scenario in which high-skill workers are 5 times more efficient than low-skill workers.

Figure A.23: Initial, Final, and True Returns to Schooling: Graphical Illustration

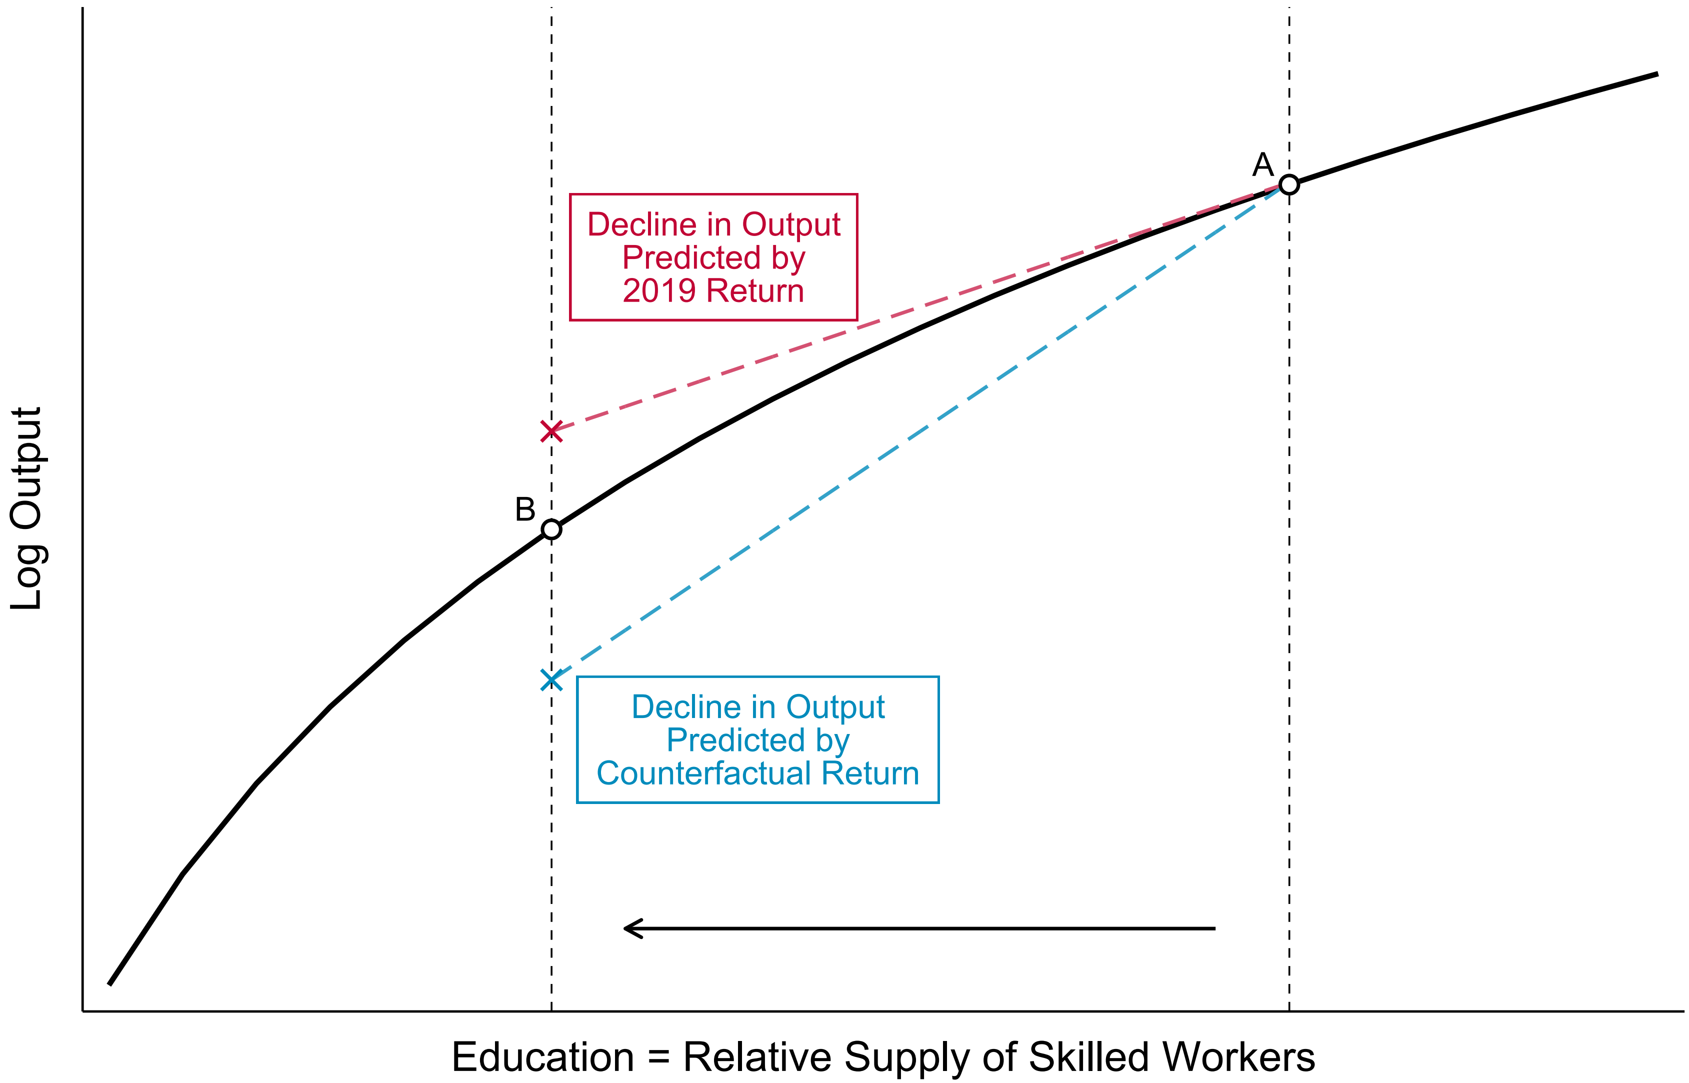

Notes. The figure provides a graphical illustration of why true returns to schooling are in-between initial and final returns. The upper dashed line shows the decline in output from reducing education as predicted by the return to schooling observed in 2019 (corresponding to final returns). The lower dashed line shows the output loss predicted by the return to schooling that prevails after reducing education (corresponding to initial returns). The true decline in output lies in-between the two estimates.

Figure A.24: Returns to Schooling: Empirical Distribution of Optimal Weights on Initial Returns

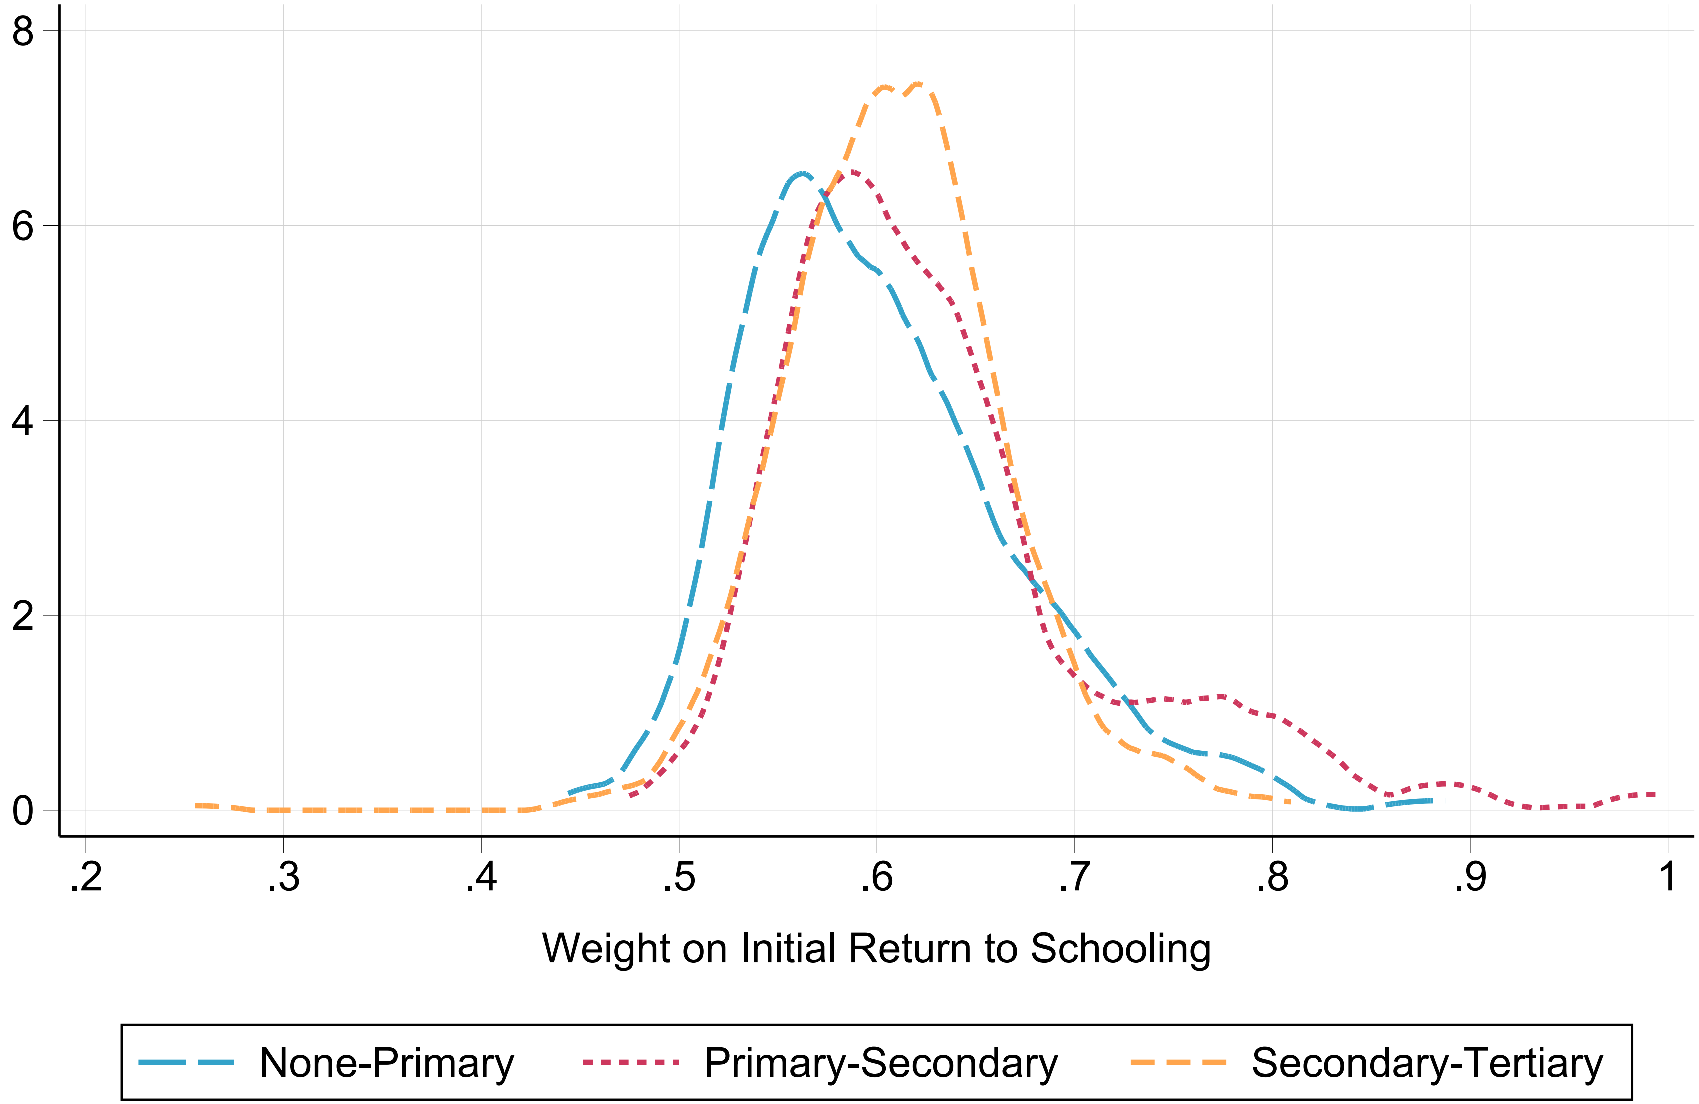

Notes. The figure plots the empirical distribution of optimal weights on initial returns to schooling required to estimate the true effect of educational expansion on output, for each of the three nests of the production function. Estimates assume an elasticity of substitution of 5.

Figure A.25: Returns to Schooling: 2019 vs. True Total Return to Primary Education

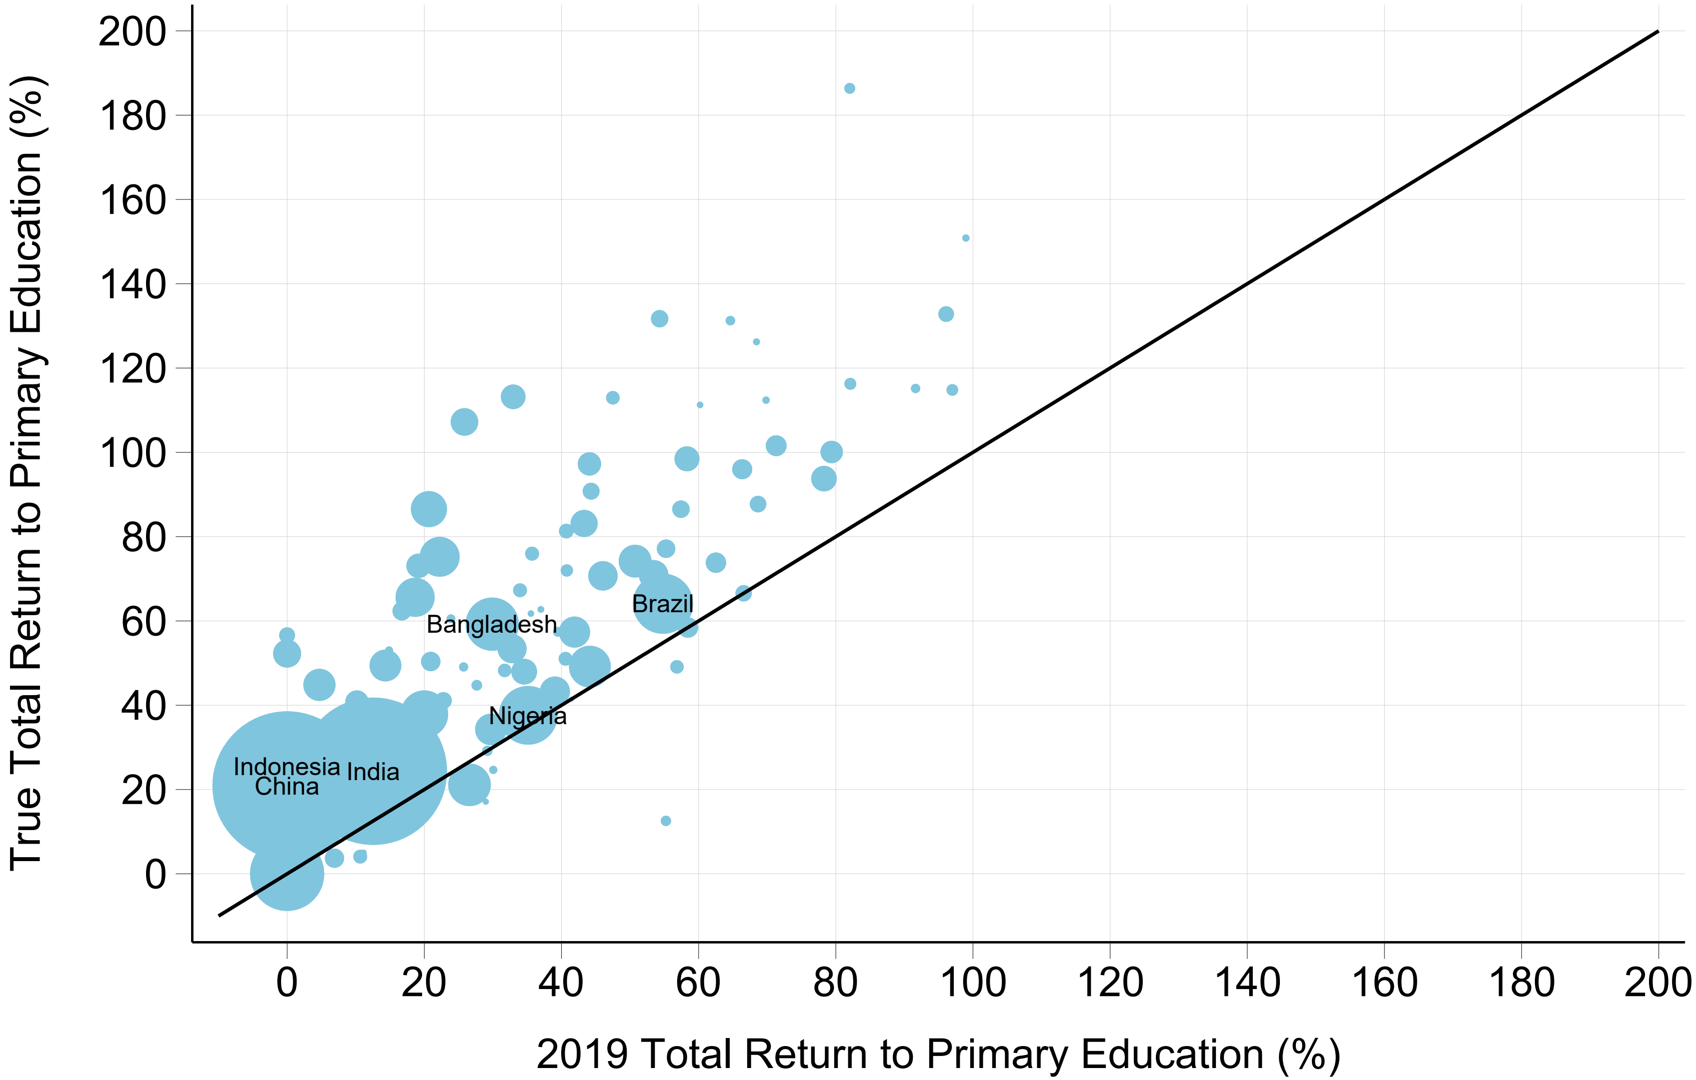

Notes. The figure compares the 2019 (final) and true total return to primary education (the percent increase in personal income of moving from no schooling to primary education) in each country. True returns are estimated assuming an elasticity of substitution of 5.

Figure A.26: Returns to Schooling: 2019 vs. True Total Return to Secondary Education

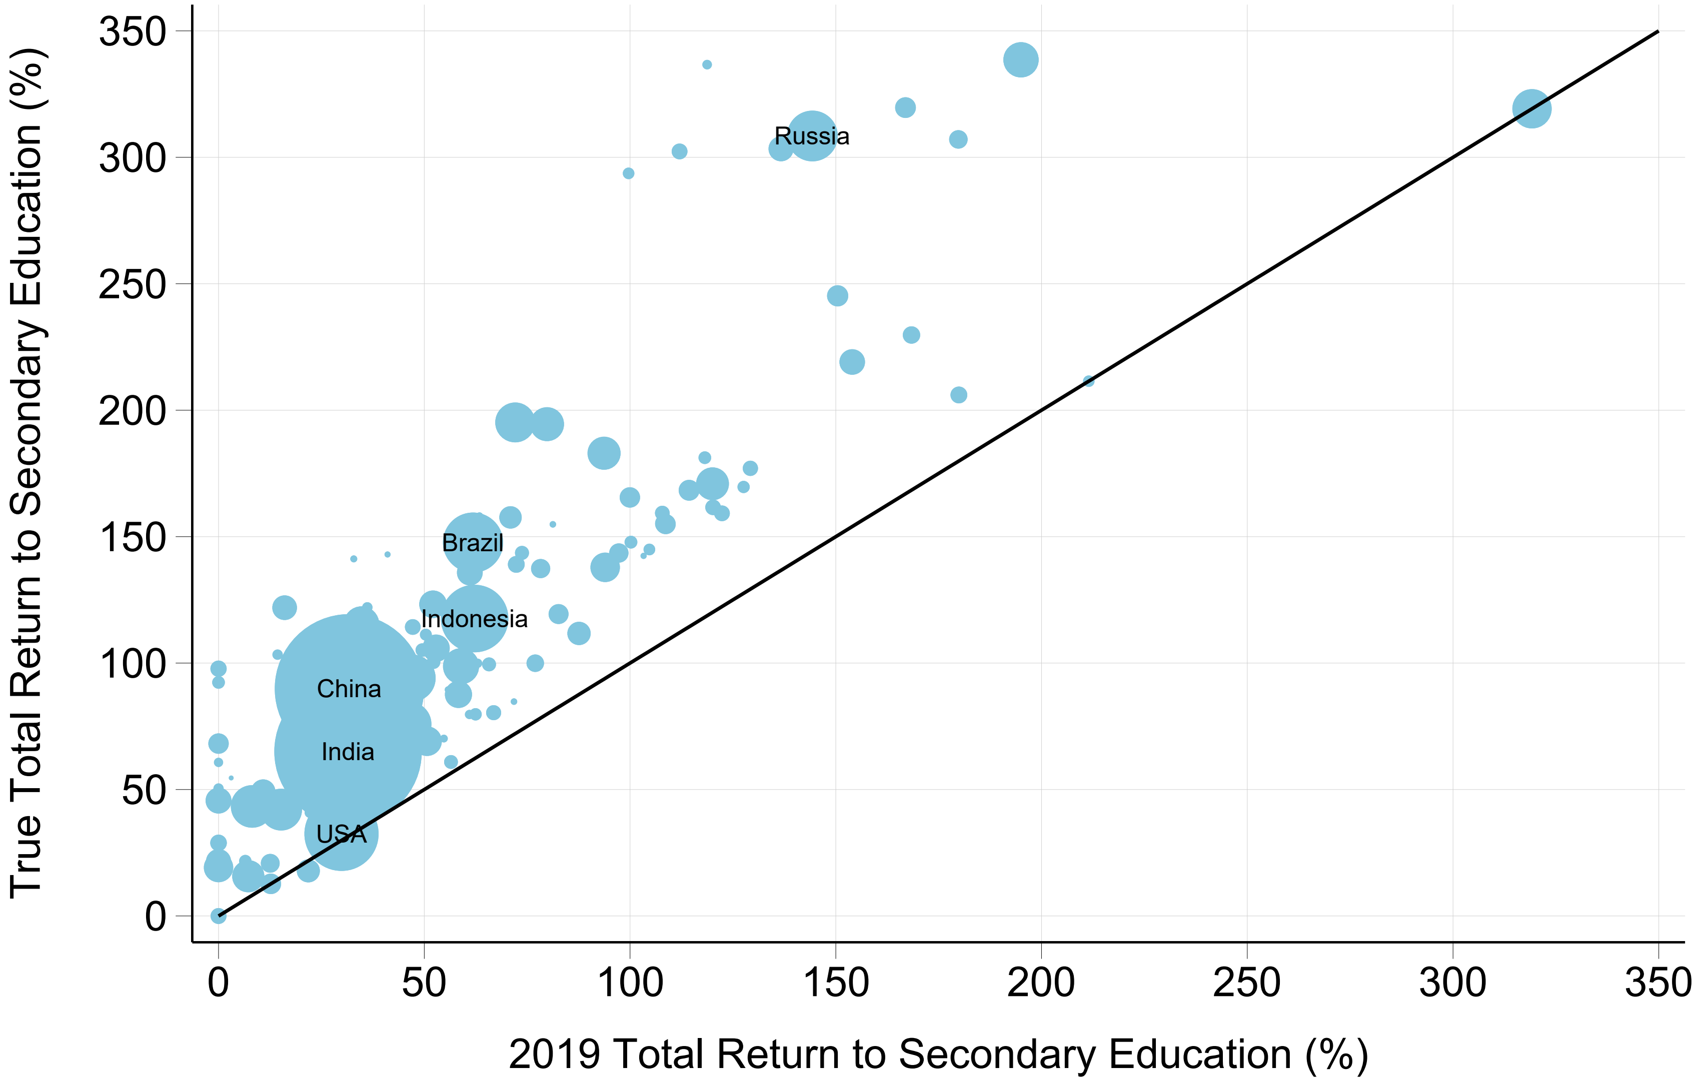

Notes. The figure compares the 2019 (final) and true total return to secondary education (the percent increase in personal income of moving from primary to secondary education) in each country. True returns are estimated assuming an elasticity of substitution of 5.

Figure A.27: Returns to Schooling: 2019 vs. True Total Return to Tertiary Education

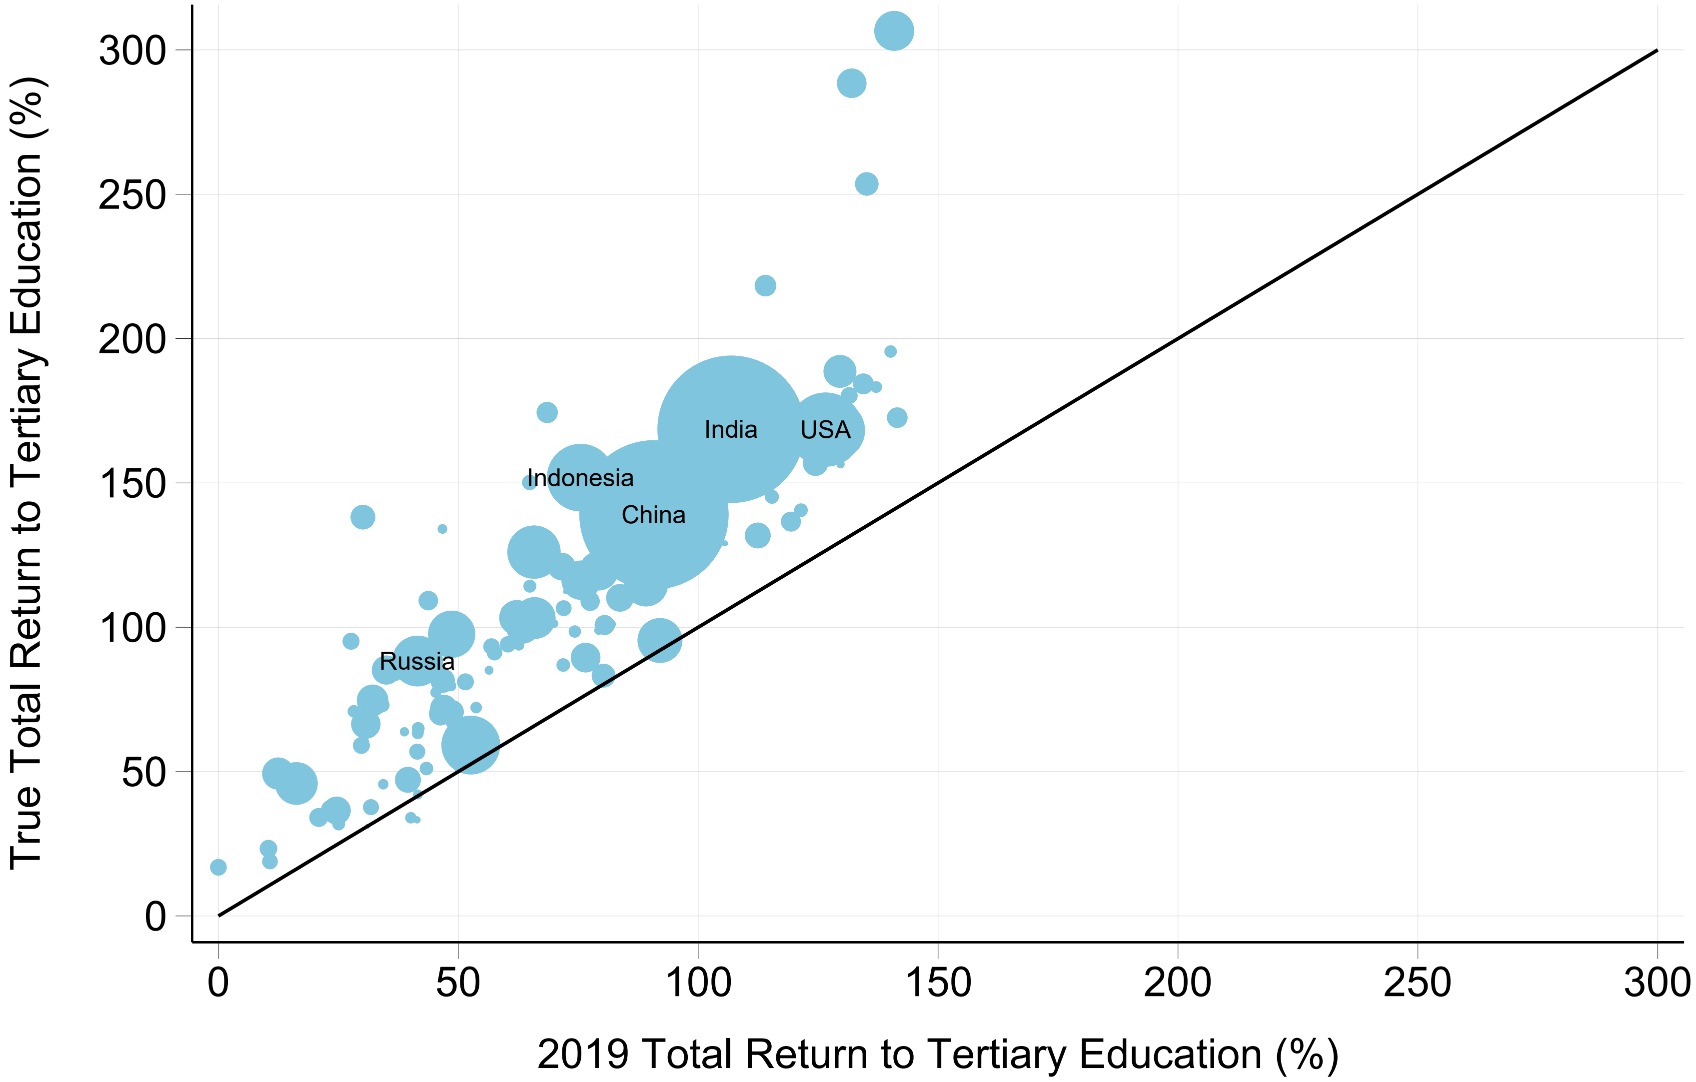

Notes. The figure compares the 2019 (final) and true total return to tertiary education (the percent increase in personal income of moving from secondary to tertiary education) in each country. True returns are estimated assuming an elasticity of substitution of 5.

Figure A.28: Relative Efficiency of Skilled Labor Versus GDP Per Capita (Tertiary Versus Below Tertiary)

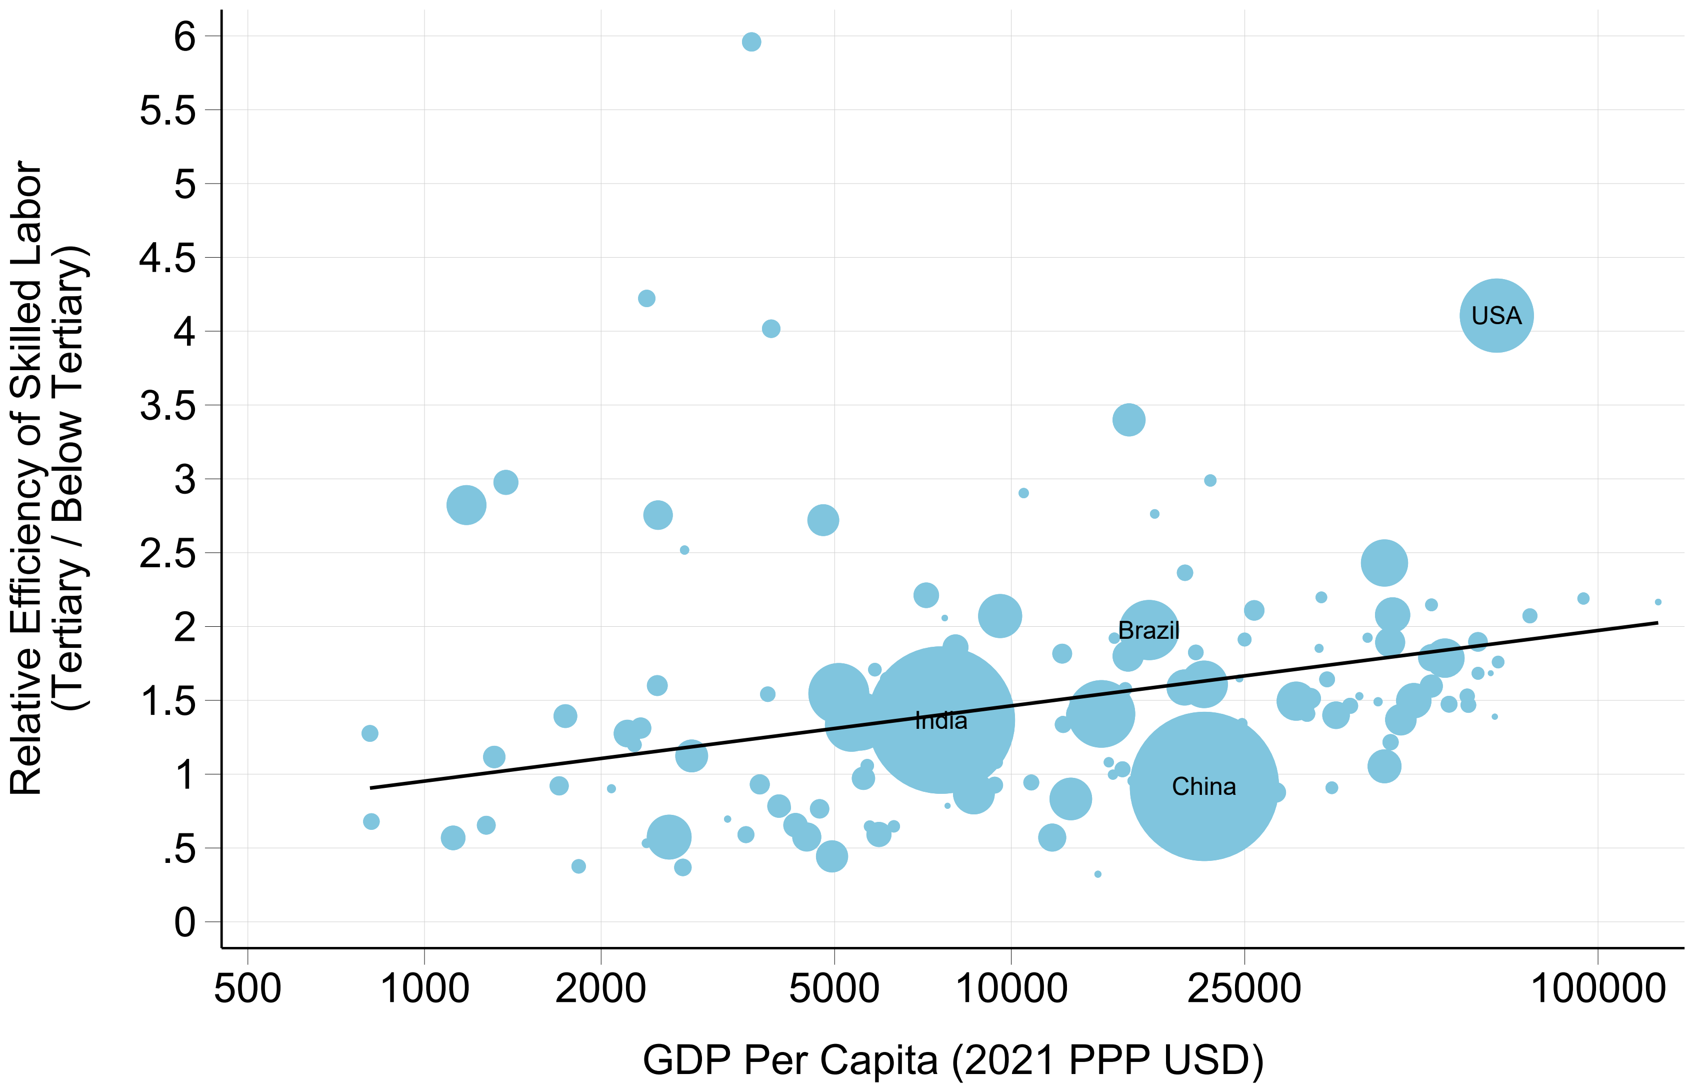

Notes. Author's calculations using survey microdata.

Figure A.29: Relative Efficiency of Skilled Labor Versus GDP Per Capita (Secondary Versus Below Secondary)

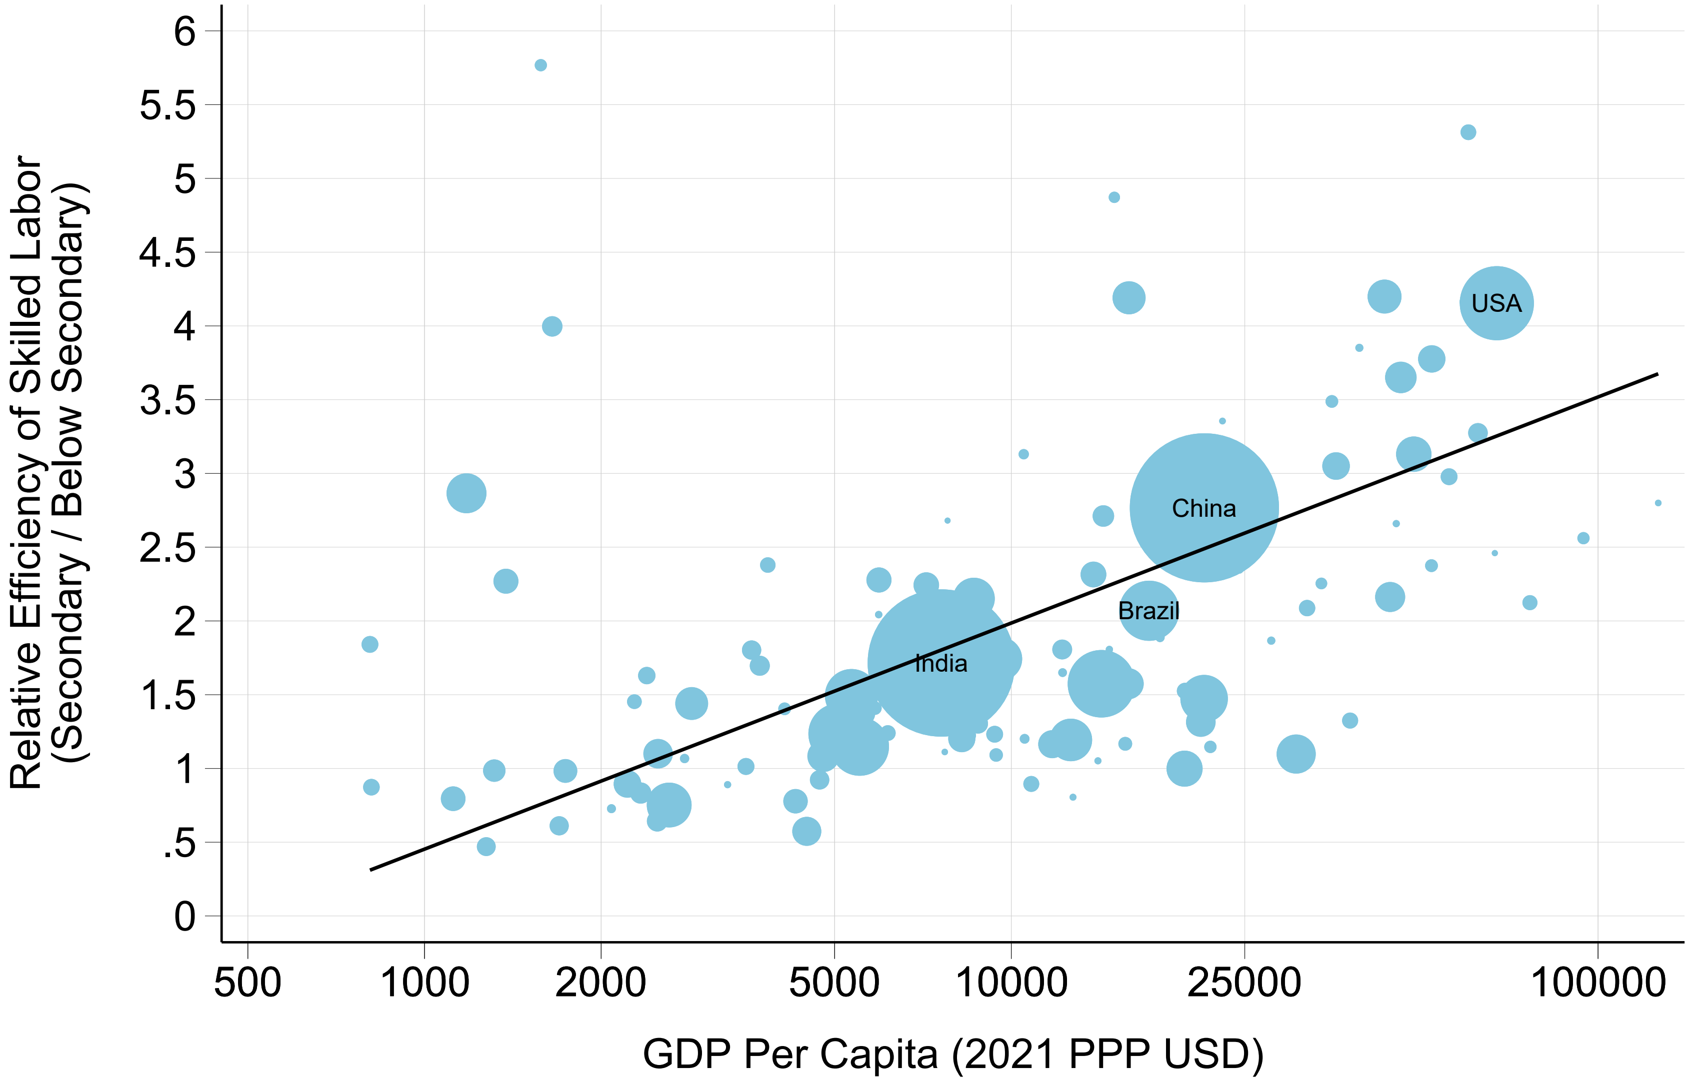

Notes. Author's calculations using survey microdata.

Figure A.30: Relative Efficiency of Skilled Labor Versus GDP Per Capita (Primary Versus No Schooling)

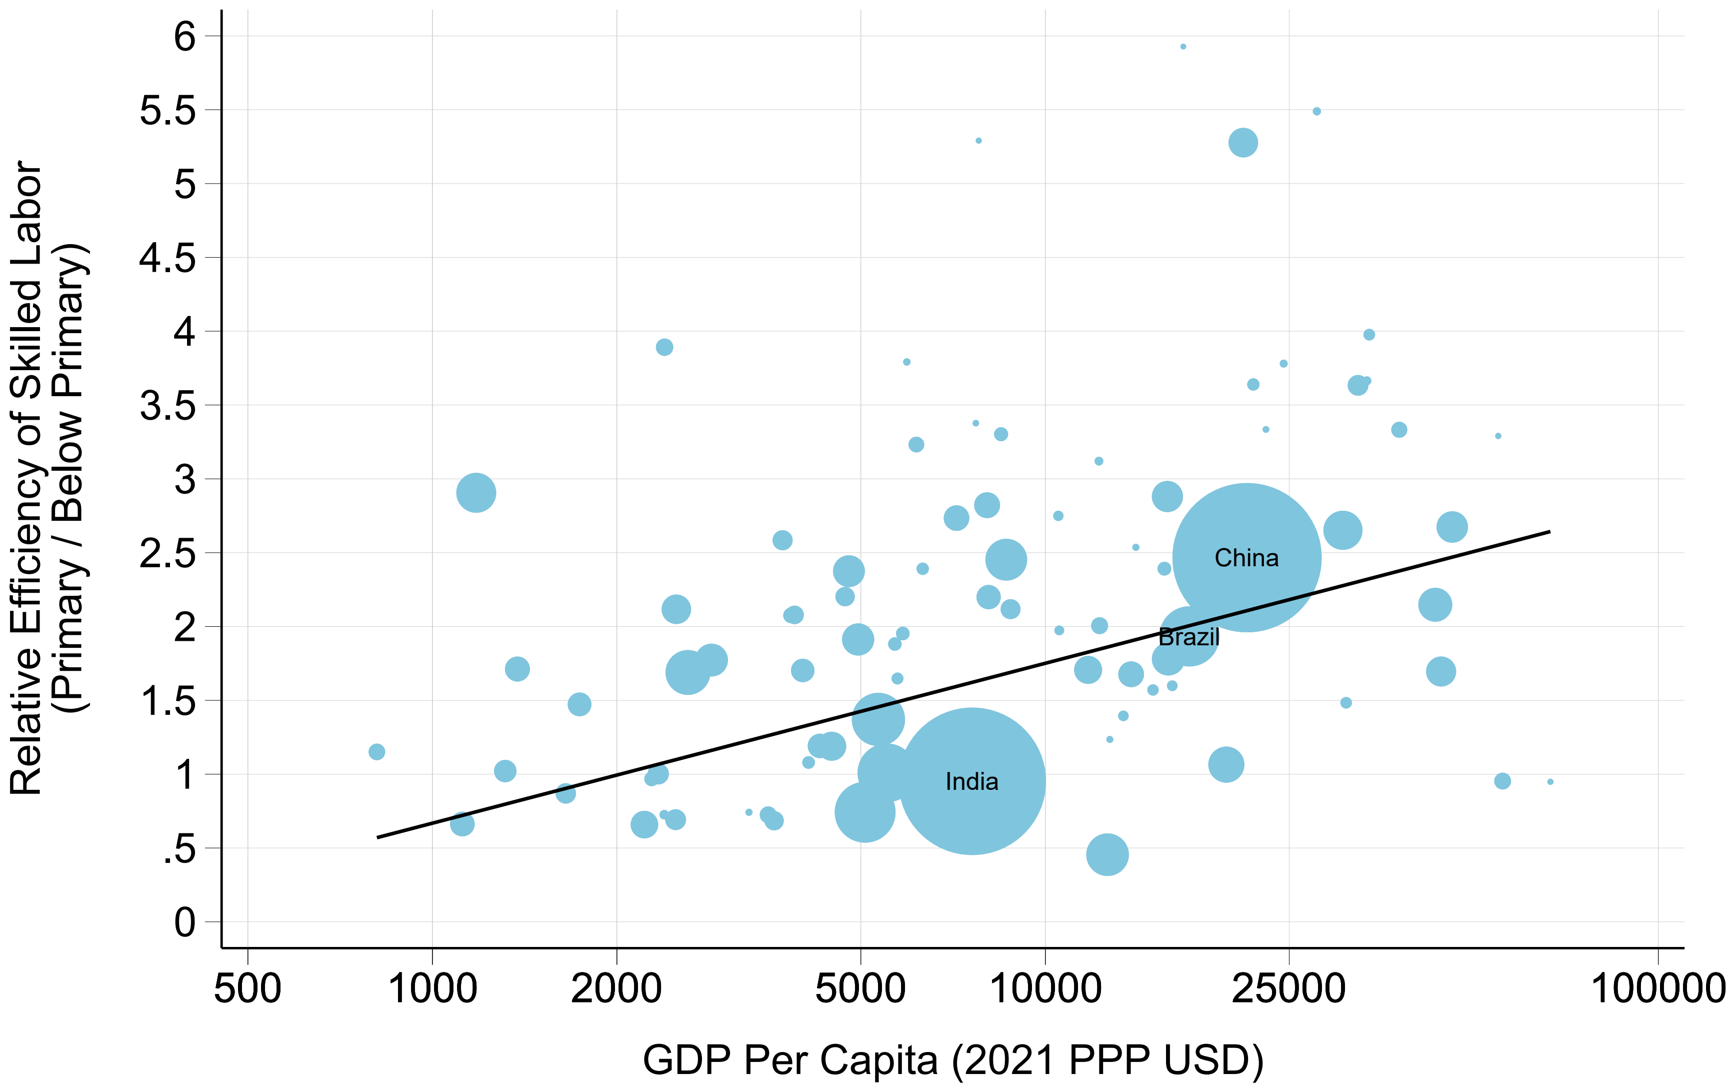

Notes. Author's calculations using survey microdata.

Figure A.31: Returns to Schooling: Simulated Distribution of Optimal Weights on Initial Returns

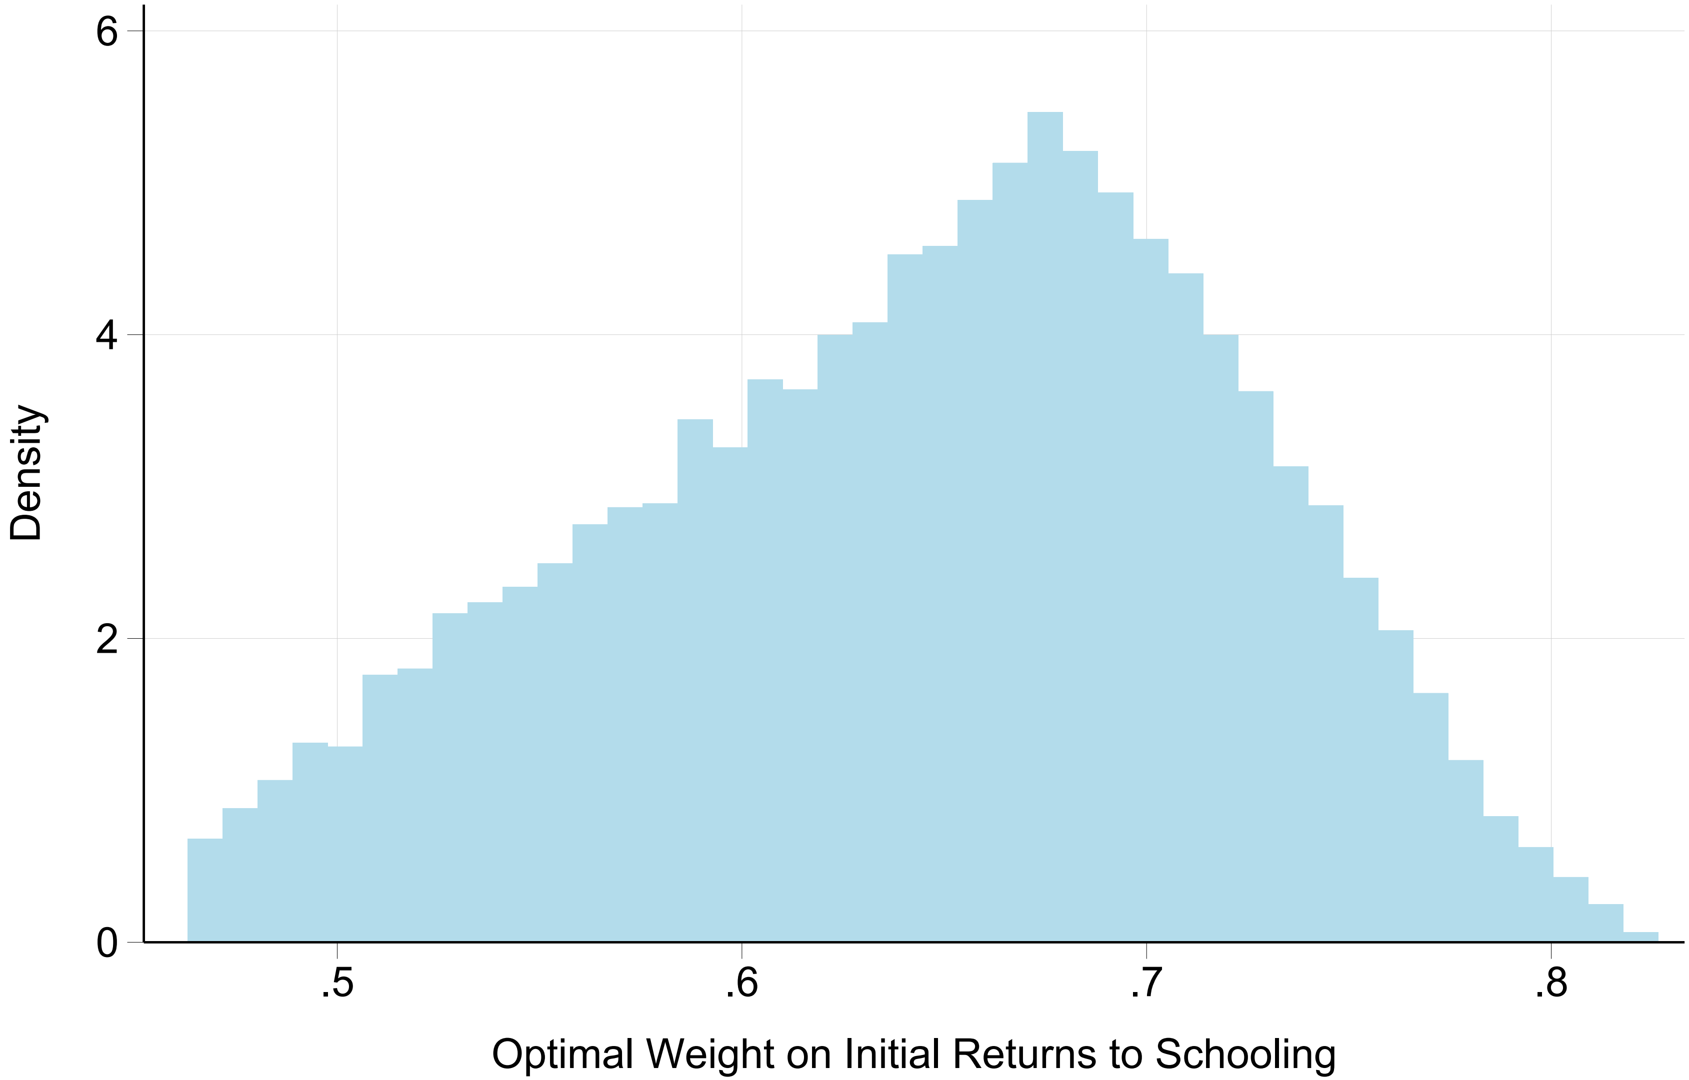

Notes. The figure plots the distribution of optimal weights on initial returns to schooling required to estimate the true effect of educational expansion on output, based on simulations varying the elasticity of substitution, the relative efficiency of skilled workers, and the magnitude of educational expansion.

Figure A.32: Returns to Schooling: Optimal Weights on Initial Returns Under Different Parametrizations of the CES Production Function

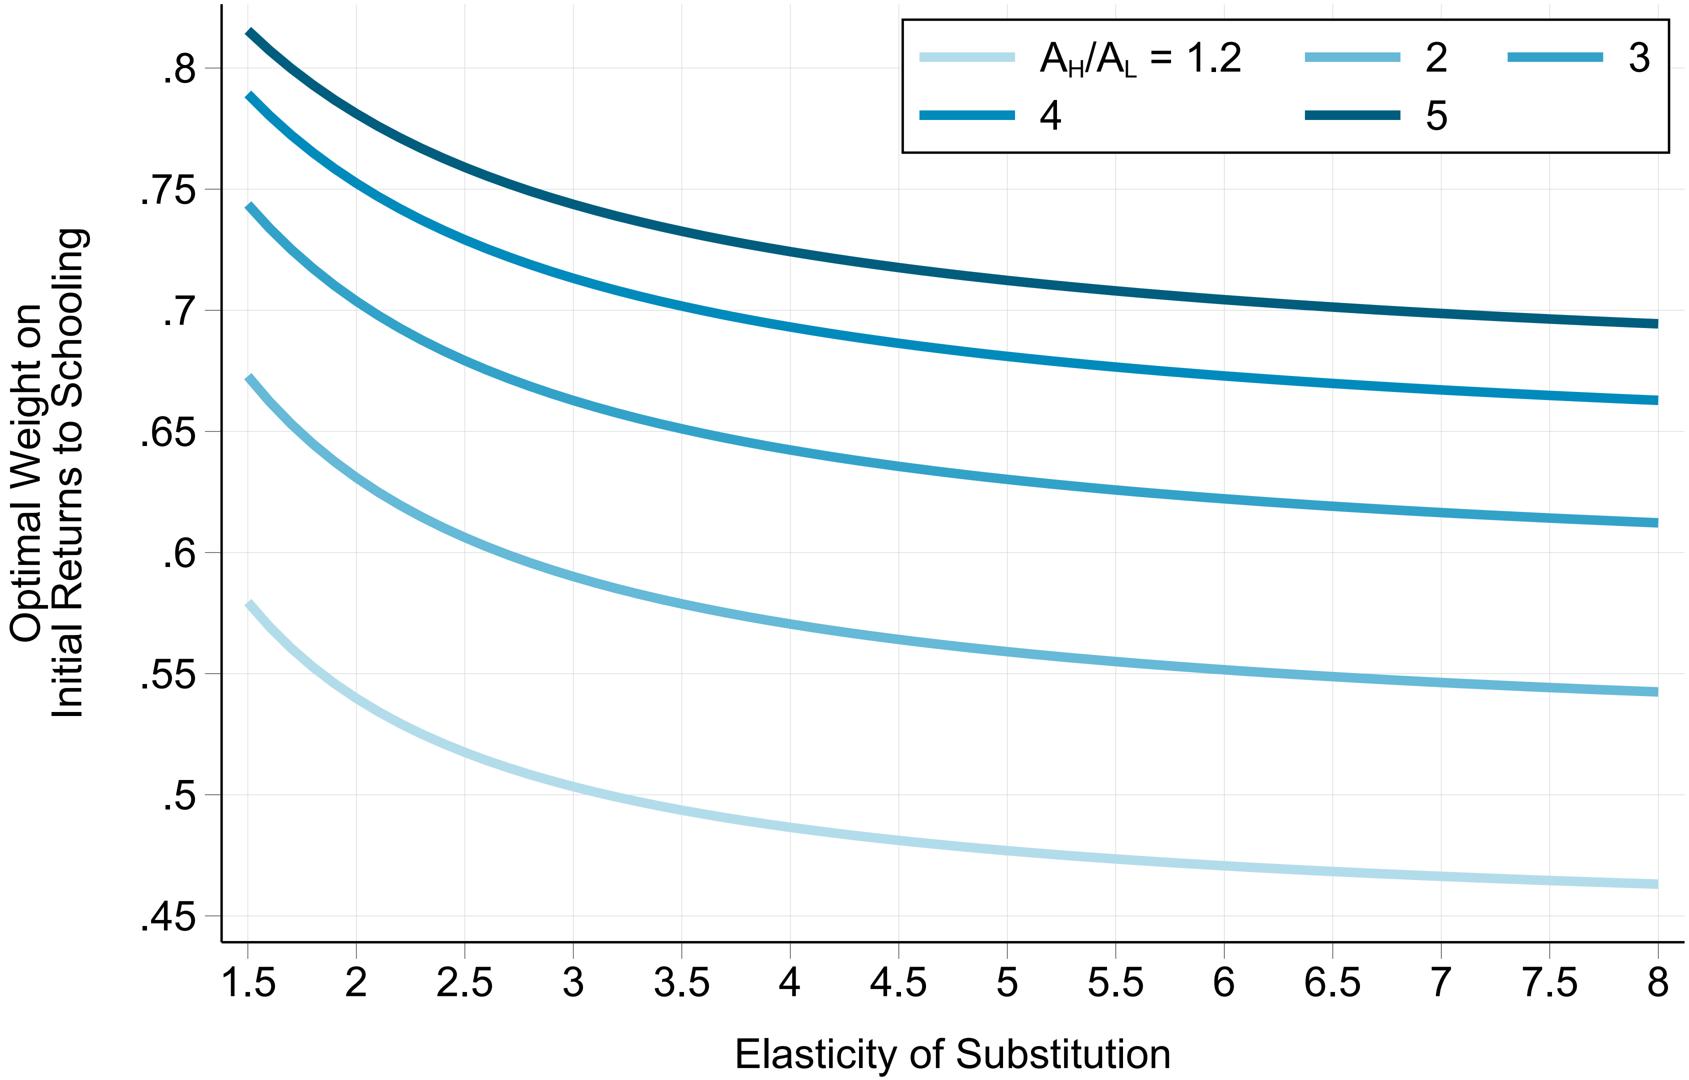

Notes. The figure plots how optimal weights on initial returns to schooling required to estimate the true effect of educational expansion on output vary, depending on parametrizations of relative skill efficiency A_H/A_L and the elasticity of substitution σ.

Table A.18: Empirical Estimates of the Elasticity of Substitution Between Skill Groups

Source	Country	Tertiary/Below	Secondary/Below
Long-run elasticity			
Bils, Kaymak, and Wu (2022)	Cross-Country	4 to 6	4 to 6
Hendricks and Schoellman (2023)	Cross-Country	4.5	7.8
Short-run elasticity			
Bowlus et al. (2021)	United States	5.3	
Autor, Goldin, and Katz (2020)	United States	1.62	
Hershbein, Kearney, and Pardue (2020)	United States	1.4	
Acemoglu and Autor (2011)	United States	1.6	
Goldin and Katz (2007)	United States	1.6	2 to 5
Ciccone and Peri (2005)	United States	1.5	
Heckman, Lochner, and Taber (1998)	United States	1.4	
Katz and Murphy (1992)	United States	1.41	
Murphy, Riddell, and Romer (1998)	Canada	1.36	
Angrist (1995)	Palestine	2	
Vu and Vu-Thanh (2022)	Vietnam	2.67	
Fernández and Messina (2018)	Latin America	1.25	2.3
Khanna (2023)	India		4.24
Caselli and Coleman (2006)	Cross-Country	1.3	

Notes. The table reports selected estimates of the elasticity of substitution between skill groups from various empirical studies. Bils, Kaymak, and Wu (2022): unique elasticity of substitution for all skill groups.

A.3 Natural Experiments

This appendix exploits evidence from three natural experiments to shed light on the ability of the model to reproduce results from real-world episodes of educational expansion. Section A.3.1 outlines the general econometric framework. Sections A.3.2, A.3.3, and A.3.4 turn to analyzing the Indian District Primary Education Program, the Indonesian INPRES school construction program, and U.S. compulsory schooling laws. Overall, the model does a remarkable job at reproducing aggregate and distributional effects of actual policies. If anything, it tends to underestimate the effect of human capital accumulation on real earnings at the bottom of the income distribution. Estimates relying on the simulation method outlined in section 1.1.4 should thus be considered a lower bound.

A.3.1 General Methodology

A large literature focuses on causally identifying individual returns to schooling. Less is known of the distributional effects of increasing human capital at the level of regions or countries. This section attempts to shed some light on these effects by studying three large-scale natural experiments in India, Indonesia, and the United States. More specifically, consider the following empirical specification:

$$\ln y_{rt}^i = \gamma_0^i + \gamma_1^i S_{rt} + X_{rt}^i \beta + \delta_r + \delta_t + \varepsilon_{rt} \tag{A.16}$$

$$S_{rt} = \alpha_0 + \alpha_1 Z_{rt} + \eta_{rt} \tag{A.17}$$

Where i denotes income groups, such as quintiles, in subnational regions r at time t. The objective is to estimate the impact of increasing average regional schooling S_{rt} on y_{rt}^i, the log average income of income group i. X_{rt}^i is a vector of controls, such as the demographic composition of the region, δ_r are subnational region fixed effects, and δ_t are time fixed effects.

The parameter of interest is γ_1^i, the semi-elasticity of average income of group i to regional average years of schooling. One option is to directly estimate equation A.16 by OLS. This is analogous to the usual cross-country or cross-region growth regression specification (e.g., Gennaioli et al., 2013). Alternatively, average schooling S_{rt} can be instrumented using an instrument Z_{rt}, such as compulsory schooling laws, which generates quasi-random differential trends in average schooling across regions. This approach has also been used, in particular in the case of U.S. compulsory schooling laws, mainly with the objective of estimating human capital externalities (Acemoglu and Angrist, 2000; Ciccone and Peri, 2006; Guo, Roys, and Seshadri, 2018). The main addition here is the focus on distributional effects, which amounts to estimating heterogeneous treatment effects by income group.

Estimating the distributional effects of educational expansion is empirically challenging, because it requires two sets of data that are rarely jointly available: data on the distribution of income within subnational regions, and an instrument that can predict quasi-random variation in regional schooling. Drawing on existing work, I study three such sources of variation: the India District Primary Education Program, the Indonesian School Construction Program, and U.S. state compulsory schooling laws.

A.3.2 India District Primary Education Program, 1994-2004

A.3.2.1 Context

Between the 1990s and the beginning of the 2000s, India engaged in a massive expansion of public schooling, the District Primary Education Program (DPEP), targeting low-literacy regions. Districts with a female literacy rate below the national average were more likely to benefit from the policy. Exploiting this allocation rule, Khanna (2023) estimates the general equilibrium effects of the program using a regression discontinuity design. He finds a return to schooling of about 13% per year (after accounting for general equilibrium effects). General equilibrium effects induced by the greater relative supply of skilled workers depress returns by one-third, while indirectly benefiting unskilled workers. To the best of my knowledge, this represents one of the only studies providing quasi-experimental evidence on the aggregate effects of schooling expansion initiatives. The design and data make it particularly well-suited for estimating the distributional incidence of human capital accumulation.

A.3.2.2 Data

I exploit data from the replication package provided by Khanna (2023). Exposure to the program is determined by district female literacy in 1991. There are 571 districts, 271 of which were treated by the program. Individual outcomes are obtained from the 2009 National Sample Survey (NSS). The microdata covers information on wages and education at the district level, allowing for a direct estimation of the impact of the program on the distribution of labor income. As in Khanna (2023), the sample is restricted to all adults aged 17 to 75 with positive wage income.

A.3.2.3 Empirical Specification

I follow Khanna (2023) and estimate the impact of the policy using the same regression discontinuity design as in the paper, comparing districts below and above the average female literacy rate. Optimal bandwidths are calculated using either the Calonico, Cattaneo, and Titiunik (2014) method or the Imbens and Kalyanaraman (2012) method (henceforth CCT and I and K, respectively).

The main addition is that I focus on the effect of the program on the average wage of each wage quintile, to directly get a reduced-form estimate of the distributional incidence of primary education expansion. Khanna (2023) centers his analysis on the estimation of individual returns to schooling, as well as spillovers to other skill groups. In contrast, I use the RD to directly instrument average district schooling and estimate its impact on the average wage of each wage quintile within each district.

Figure A.33 plots the first stage, comparing district average years of schooling among adults with positive wage income below and above the literacy cutoff. Districts below the cutoff were more likely to benefit from the program. Adults living in districts that were just below this cutoff have significantly higher levels of education than those living in districts just above.

A.3.2.4 Results

Table A.19 presents the main results. Increasing district average years of schooling by one year is associated with a 12% increase in wages in treated districts (CCT method). This effect is almost two times larger for the bottom 20% of earners, who benefit from a 21% increase in wages. In contrast, the top 20% see their average wage decline, although the coefficient is not statistically significant. Results relying on the I and K method are similar, but the aggregate effect of educational expansion appears even larger, reaching 26% for average wages and 32% for the bottom 20%. Aggregate returns to schooling estimated using this method are in the range of individual returns estimated by Khanna (2023), who finds returns of 16% (CCT) to 21% (I and K) using conventional 2SLS estimates, and 13% (CCT) after accounting for general equilibrium effects.

Table A.20 compares the CCT estimates to simulated effects of expanding primary education, under different parametrizations of the return to schooling and the elasticity of substitution between skilled and unskilled workers. Figure 1.6 plots the corresponding coefficients in the specific case where the return to schooling is set to 13% and the elasticity of substitution to 4 (corresponding approximately to the values obtained by Khanna, 2023). The simulation is done by upgrading the education of randomly sampled individuals from no schooling to primary education, increasing their earnings using the return to schooling, and finally adjusting relative wages for general equilibrium effects, using the method outlined in section 1.1.

Simulated estimates fall close to the true effects of the policy. Simulation results show that with a return to schooling of 13%, increasing average district education by one year through basic education is associated with an increase in average wages of about 9%. The aggregate effect is lower than the individual return, because those benefiting from the expansion are workers with no schooling, whose wages are significantly lower than average. To simulate an aggregate effect similar to the one estimated using the natural experiment, a higher individual return to schooling is required, in the order of 16%. Relying on individual returns to estimate the effect of basic education expansion thus provides a lower bound on the true aggregate effect of educational expansion.

The simulation also predicts distributional effects that are very similar to those estimated with the RD design. Both in the simulation and in the natural experiment, benefits appear relatively similar for the first four quintiles and significantly lower for the top 20%. This can be rationalized by the fact that in India, workers with no schooling and workers with basic education are both prevalent among the bottom 80% of the distribution, so that upgrading some workers from no schooling to basic education benefits this entire group. Simulated effects do not vary much with the elasticity of substitution, although lower values of the elasticity are associated with greater gains for the bottom quintile, where the concentration of workers with no schooling in the greatest. All in all, the model performs remarkably well at reproducing the observed economic effects of primary education expansion in India.

A.3.3 Indonesia School Construction Program, 1973-1978

A.3.3.1 Context

Between 1973 and 1978, Indonesia engaged in a massive school construction program aiming to expand access to basic education throughout the country. Exploiting differences in exposure to newly built schools across cohorts and regions, Duflo (2001) estimates individual returns to schooling ranging from 7% to 11%. A number of studies have updated and extended her analysis since then, focusing on intergenerational effects (Akresh, Halim, and Kleemans, 2023), structural transformation (Karachiwalla and Palloni, 2019), or rural-urban migration (Hsiao, 2023).

Duflo (2004) also moves beyond individual outcomes to focus on spillovers of the program to non-treated groups. Combining labor force surveys covering the 1986-1999 period, she estimates the impact of the greater supply of young skilled workers on older generations' formal employment and wages. Her analysis shows mixed findings, suggesting a decline in the wages of non-treated groups, but an increase in employment in the formal sector.

A.3.3.2 Data

Drawing on the work of Duflo (2004), I exploit differential exposure to the program by district to estimate the aggregate and distributional effects of primary education expansion. My analysis expands her work in two ways. First, I significantly expand the time coverage of the data, which increases statistical power and allows me to get closer to long-run effects. To do so, I collect and harmonize every round of the SUSENAS, a household survey covering about a million individuals every year, from 1993 to 2019. The result is a balanced panel of 230 districts, for which I have yearly data on the education level of the labor force, the distribution of consumption, and other sociodemographic variables over a twenty-six-year period.[2] Second, I study the effects of the program on total district consumption and its distribution by consumption quintile, while Duflo (2004) focuses on spillover effects on older cohorts. The sample is restricted to all adults aged 15 to 70; consumption is then split equally all members of the household.

A.3.3.3 Empirical Specification

The empirical specification corresponds to the one in equation A.16, with schooling being instrumented as in Duflo (2004). I estimate the effect of average years of schooling in district r on the log average consumption of decile i, controlling for district and year fixed effects. Average years of schooling is instrumented by the interaction between survey years and the number of schools built per 5-14 population between 1974 and 1978.[3] The school construction program is thus taken as an instrument for *differential trends* in the education of the working-age population across districts from 1993 to 2019. Districts with greater treatment intensity are expected to see a faster secular increase in average schooling, because of the greater

[2]Some districts have undergone splits and merges over the period of interest. I rely on crosswalks provided by Roodman (2022) to ensure consistent boundaries over time.

[3]Given significant noise introduced by the low sample size available for each district-year cell, I specify survey years as a continuous variable in the first stage. Indeed, as shown by Duflo (2004), we should expect the program to have introduced smooth, secular differential trends in educational expansion. Constraining the interaction of survey years and treatment intensity to follow such a secular trend makes the results less sensitive to different empirical strategies.

access to schooling enabled by the policy for cohorts educated 20-50 years ago. The identification assumption, analogous to Duflo (2004), is that there is no unobserved shock both correlated with the program and affecting household expenditure during that period.

Figure A.34 plots the first stage. The dependent variable is average years of schooling in a given district-year; each point corresponds to the coefficient of the interaction of a survey year dummy with treatment intensity, with 1993 taken as the baseline year. This figure is analogous to Duflo (2004), figure 2. In districts with greater exposure to the school construction program in 1974-1978, average years of schooling among the working-age population have risen at a significantly faster pace. The estimates are slightly noisy, because of the relatively low number of observations available in each district-year cell, but they confirm that the program had long-lasting effects on regional rates of human capital accumulation.

A.3.3.4 Results

Table A.21 presents the main results. The baseline specification controls for the demographic and gender composition of each district, the share of college graduates, and district and year fixed effects. Increasing average district schooling by one year is associated with an 8.7% rise in average consumption in the district. This effect is almost four times larger for the bottom quintile (22%) than for the top quintile (5.8%). Columns 4 to 6 add controls for 1971 primary school enrollment and water and sanitation spending interacted with survey year, as in Duflo (2004). These estimates are underpowered and the standard errors much larger, because of limited sample size in each district, but the results obtained are qualitatively similar. The effect on bottom 20% average consumption rises to 51%, while that on the top 20% boils down to zero. Columns 7 to 9 add further controls for 1971 child population and population density interacted with survey year. This model is even more underpowered, but the point estimates remain of the same order of magnitude. In particular, the coefficient on the average income of the bottom 20% remains large (45%) and statistically significant at the 5% level. While it is clear that the sample size is not sufficient to precisely estimate aggregate returns to schooling, the progressive nature of the policy stands out across all specifications.

Table A.34 compares the benchmark estimates to simulated effects of the policy using the 1996 Indonesian labor force survey (SAKERNAS). Figure 1.7 plots the coefficients by quintile when the return to schooling is set to 12% and the elasticity of substitution to 4.[4] The simulation is done exactly as in the Indian case, upgrading the education of randomly sampled individuals from no schooling to primary education, increasing their earnings using the return to schooling, and finally adjusting relative wages.

As in India, the simulation does a good job at reproducing results from the natural experiment. The expansion of primary education is estimated to be progressive in all specifications, with orders of magnitude similar to those found in the data. Lower values for the elasticity of substitution are associated with significantly higher growth for the bottom 40% relative to the top 60%. The benchmark specification, with a return of 11% (close to the estimate of Duflo, 2001) and an elasticity of 4, matches both aggregate and distributional effects particularly well. Higher elasticities of substitution generally imply inequality-reducing effects of the policy that are too low in comparison to those observed in the data.

A.3.4 U.S. Compulsory Schooling Laws, 1875-1961

A.3.4.1 Context

Between the mid-19th and the mid-20th century, U.S. states gradually implemented laws limiting child labor and enforcing compulsory school attendance for newly educated cohorts. The effect of these laws were first studied by Acemoglu and Angrist (2000), who combined data on laws implemented from 1914 to 1965 with census microdata to estimate the magnitude of human capital spillovers. Their analysis gave rise to a rich literature exploiting compulsory schooling laws to estimate individual returns to schooling (Clay, Lingwall, and Stephens, 2021; Stephens and Yang, 2014), elasticities of substitution between skill groups (Ciccone and Peri, 2006), and human capital externalities (Ciccone and Peri, 2006; Guo, Roys, and Seshadri, 2018; Iranzo and Peri, 2009).

A.3.4.2 Data

My analysis relies on similar sources than those used in the existing literature, but extends previous work in two ways. First, I study the total aggregate and distributional effects of educational expansion, while existing studies focus on estimating

[4] Duflo (2004) finds no evidence that schooling expansion led to a significant decline in the skill premium (Duflo, 2004, table 6), which would point to an infinite elasticity. However, the standard errors are very large, implying confidence intervals that include both very low and negative elasticities. Rather than evidence in favor of perfect substitution, these findings point to the fact that limitations in the sample size unfortunately make it difficult to estimate such elasticity with the available data.

different dimensions of these effects separately. Second, I exploit recently compiled data by Clay, Lingwall, and Stephens (2021), covering compulsory schooling laws over the entire 1875-1961 period. This represents an important improvement over the previous literature, which only covered laws implemented after 1915, based on the database of Acemoglu and Angrist (2000).[5] To estimate the impact of schooling on the distribution of income, I rely on the 1940 to 2000 census microdata samples available from IPUMS USA, which cover personal income, state of birth, state of residence, education, and other sociodemographic variables by ten-year interval. The sample is restricted to all adults aged 25 to 65 with positive personal income (wage income in 1940) living in the contiguous United States.

A.3.4.3 Empirical Specification

As in the Indian and Indonesian case studies, I regress the average income of each personal income decile on average state schooling, instrumented by compulsory schooling laws. More specifically, consider the following instrument for average years of schooling S_{st} in state s at time t:

$$S_{st} = \pi_0 + \pi_1 \sum_c \sum_{s'} N_{css't} RS_{cs'} + \theta_s + \theta_t + u_{st} \tag{A.18}$$

Where $RS_{cs'}$ is required years of schooling for cohort c born in state s', $N_{css't}$ is the number of individuals living in state s at time t who were born in state s', and θ_s and θ_t are state and year fixed effects. Required years of schooling correspond to the time a children born in a given year is required to stay in school, calculated by combining information on required attendance at each year of life (see Clay, Lingwall, and Stephens, 2021; Stephens and Yang, 2014). The instrument is thus equal to the average *required* years of schooling of the working-age population of state s, calculated by averaging required years of schooling across all cohort-state-of-birth cells, weighted by their relative populations at time t. This approach is analogous to the one recently adopted by Guo, Roys, and Seshadri (2018), who instrument average state education by required years of schooling in each state-age cell.

The interpretation of the instrumentation strategy is similar to that of the Indonesian case study. Differences in required years of schooling across cohorts born from 1875 to 1961 are used to predict differential trends in average schooling across states from 1940 to 2000.

[5]Consider in particular the 1940 to 1960 censuses, which cover periods of the twentieth century during which basic education mattered most for explaining cross-state variations in human capital. Post-1915 compulsory schooling laws fail to capture variations in schooling for all workers older than 25-45 during that period, so they end up missing important sources of variations.

Figure A.35 provides a concrete illustration of how required years of schooling evolved across cohorts born in Alabama, California, Indiana, and Massachusetts from 1875 to 1965. In 1875, Massachusetts was the only state imposing compulsory education, for a duration of 6 years. All states saw the implementation of increasingly restrictive laws, but with significant variations in timing and intensity. Indiana rapidly shifted from no compulsory schooling to nine years in 1899, while Alabama followed the same transition much more gradually, from no compulsory law until 1902 to four years in 1905, six years in 1909, eight years in 1919, and finally nine years by 1933.

Table A.23 shows the first stage. Column 1 controls for the demographic, gender, and racial composition of each state, as well as the share of college graduates. An additional average required year of schooling is associated with a 0.19 increase in actual average years of schooling among the working-age population. Column 2 add census × year fixed effects, which have been shown to potentially matter significantly when estimating the effect of U.S. compulsory schooling laws (Stephens and Yang, 2014). This reduces the effect to 0.14. Finally, column 3 adds further controls for initial conditions, interacting census year fixed effects with average income and average years of schooling in 1940. This is a very ambitious specification, as it implies estimating over 100 coefficients on a sample of only 343 observations. The coefficient on required years of schooling is reduced to 0.12, and remains statistically significant at the 1% level.

A.3.4.4 Results

Table A.24 presents the main results. In the baseline specification, an additional average year of schooling is associated with a 0.16 log-point increase in average income. The corresponding values are 0.44 for the bottom 20%, compared to 0.05 for the top 20%. Educational expansion thus appears to have been a powerful driver of inequality reduction, even more so in the U.S. than in India and Indonesia. Adding interacted census region and year fixed effects leaves the results almost unchanged (columns 4 to 6). Columns 7 to 9 further add controls for initial conditions. The aggregate effect is slightly lower (0.08), and the estimates are unsurprisingly underpowered. Even under this highly demanding specification, however, the coefficient on the bottom 20% remains large (0.27) and statistically significant at the 5% level.

Table A.25 compares observed and simulated effects of the policy. Figure 1.8 plots the coefficients by quintile when the return to schooling is set to 12% and the elasticity of substitution to 4, as in the previous case studies. The simulation is done by upgrading the education of randomly sampled individuals with either no schooling or primary education to secondary education, given that required years of schooling range from 0 to 9 years.

Here, in contrast to the two previous case studies, the model appears to strongly underestimate the aggregate and inequality-reducing effects of the policy. Even with returns to schooling of 16% and an elasticity of substitution of 2, it can generate an effect on the average income of the bottom quintile of "only" 0.4 log points, while overestimating growth for the top quintile.

There are at least three reasons why this might be the case. First, state compulsory schooling laws extended both primary and secondary school attendance, with significant variations in timing and intensity across states. This makes it more difficult to accurately simulate the overall effect of these policies. Indeed, who exactly benefited from them (individuals who would have had either no schooling, some primary education, or some secondary education in the absence of these policies) is less clear than in the Indian and Indonesian cases. Second, there is evidence that returns to schooling were substantially higher at the bottom of the income distribution during the first wave of compulsory schooling laws (Clay, Lingwall, and Stephens, 2021). In contrast, the simulation assumes a constant return by income group. This limits by construction its ability to capture higher returns for low-income earners. Third, recent evidence points to potentially large human capital externalities from schooling expansion in the United States, as high as 6-8% per year of schooling (Guo, Roys, and Seshadri, 2018). This might explain why even with an individual return to schooling as high as 16%, the simulation ends up underestimating the aggregate return by 4-5 percentage points.

Figure A.33: India DPEP: First Stage

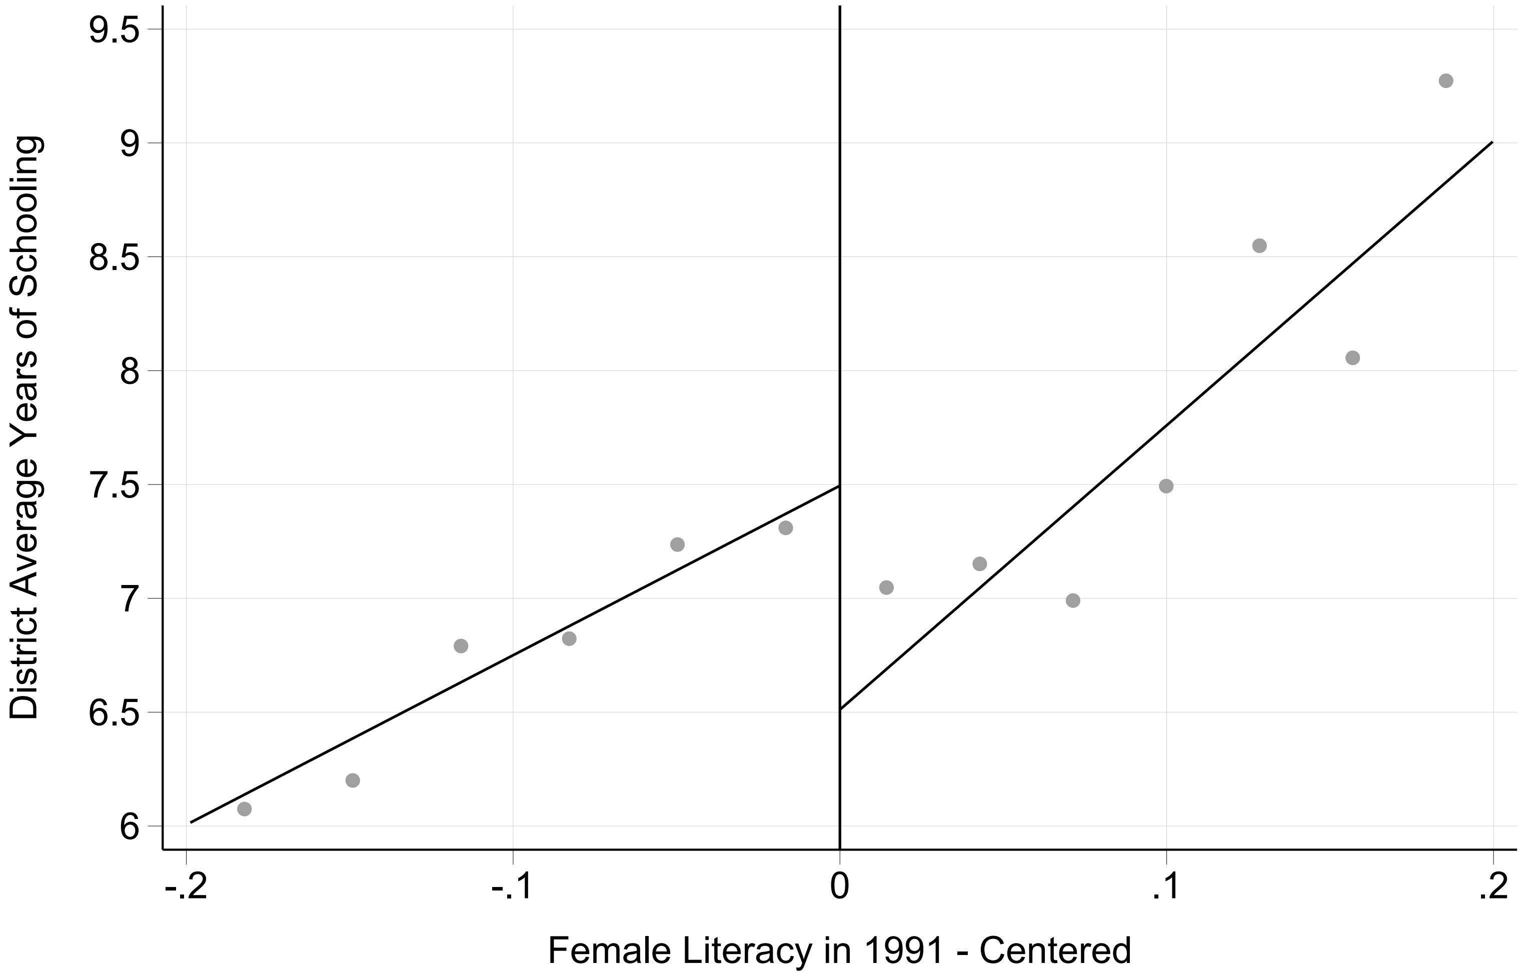

Notes. The figure compares average years of schooling among adult wage earners below and above the literacy cutoff used to allocate the program. Data from Khanna (2023).

Table A.19: India DPEP: Aggregate and Distributional Effects of Schooling

	Bandwidth Selection: CCT Method			Bandwidth Selection: I and K Method		
	Average Income	Bottom 20% Income	Top 20% Income	Average Income	Bottom 20% Income	Top 20% Income
Average Years of Schooling	0.117*	0.207***	-0.092	0.257***	0.316***	0.012
	(0.061)	(0.059)	(0.071)	(0.058)	(0.058)	(0.068)
N	46314	9007	9515	46314	9007	9515

Notes. The table reports the effect of district average years of schooling on district average income, the average of the bottom 20%, and the average income of the top 20%. Bandwidths: "CCT" indicates the Calonico, Cattaneo, and Titiunik (2014) method, "I and K" the Imbens and Kalyanaraman (2012) method. Data from Khanna (2023). Standard errors in parenthesis. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.20: India DPEP: Actual vs. Simulated Effects of Educational Expansion

	Parameters		Effect of Increasing Average District Schooling by One Year (%)					
	Return to Schooling	Elasticity of Substitution	Average Income	Q1	Q2	Q3	Q4	Q5
Actual Effect			11.7	20.7	11.2	17.1	19.3	-9.2
Simulated Effect	13%	∞	9.3	17.4	13.5	14.8	10.0	5.6
	13%	6	9.3	19.2	13.4	14.8	10.0	5.5
	13%	4	9.3	20.0	13.3	14.8	10.0	5.4
	13%	2	9.3	22.6	12.9	14.9	9.9	5.2
	16%	∞	12.4	18.9	15.0	18.5	14.3	8.3
	16%	6	12.4	20.9	15.0	18.6	14.3	8.1
	16%	4	12.4	21.9	15.0	18.6	14.3	8.0
	16%	2	12.4	24.8	14.7	18.9	14.2	7.7
	20%	∞	17.1	20.0	16.5	21.9	21.0	13.4
	20%	6	17.1	22.3	16.8	22.2	21.0	13.1
	20%	4	17.1	23.4	16.9	22.4	20.9	12.9
	20%	2	17.1	26.8	17.1	22.9	20.8	12.4

Notes. Actual effect: estimated effect of the policy on average district income and the average income of each wage quintile, using data from Khanna (2023). Simulated effect: effect of the policy predicted using 2019 LFS data, under different assumptions on the return to a year of schooling and the elasticity of substitution between skilled and unskilled workers.

Figure A.34: Indonesia INPRES: First Stage: Effect of the Program on
District Average Years of Schooling, 1993-2019

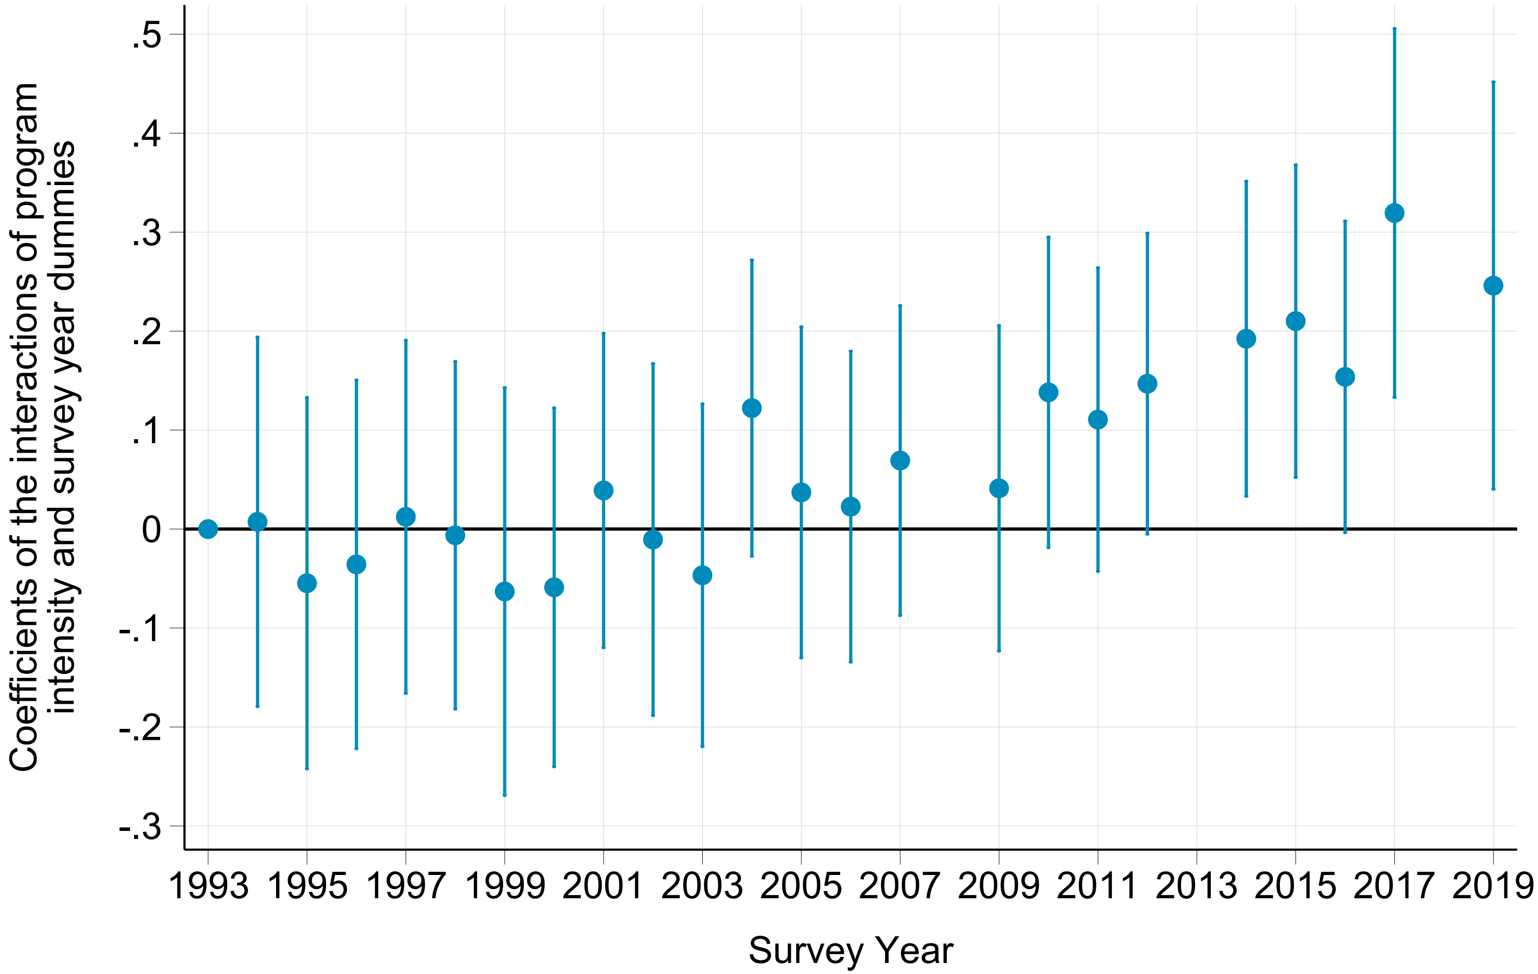

Notes. The figure compares the evolution of average years of schooling in districts with more or less exposure to the INPRES school construction program. The dependent variable is average years of schooling in each district-year. Estimates combine 1993-2019 SUSENAS microdata with treatment intensity by district from Duflo (2001).

Table A.21: Indonesia INPRES: Aggregate and Distributional Effects of Schooling

	Baseline			+ Controlling for 1971 Primary School Enrollment and Water & Sanitation Spending			+ Controlling for 1971 Child Population and Population Density		
	Average Income	Bottom 20% Income	Top 20% Income	Average Income	Bottom 20% Income	Top 20% Income	Average Income	Bottom 20% Income	Top 20% Income
Average Years of Schooling	0.087***	0.220***	0.058*	0.133	0.505***	-0.002	0.084	0.445**	-0.029
	(0.026)	(0.035)	(0.031)	(0.083)	(0.145)	(0.097)	(0.108)	(0.179)	(0.131)
N	5520	5520	5520	5352	5352	5352	5304	5304	5304

Notes. The table reports the effect of regency average years of schooling on regency average income, the average income of the bottom 20%, and the average income of the top 20%. Columns 1 to 3 control for the demographic composition of the regency, the share of women, and the share of workers with tertiary education. Columns 4 to 6 add controls for 1971 primary school enrollment rates and water and sanitation spending, interacted with survey year. Columns 7 to 9 further add controls for the share of the population aged 5 or below in 1971 and population density in 1971, interacted with survey year. Data from Duflo (2001) and Roodman (2022). Standard errors in parenthesis. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.22: Indonesia INPRES: Actual vs. Simulated Effects of Educational Expansion

	Parameters		Effect of Increasing Average District Schooling by One Year (%)					
	Return to Schooling	Elasticity of Substitution	Average Income	Q1	Q2	Q3	Q4	Q5
Actual Effect			8.7	22.0	15.4	11.5	8.1	5.8
Simulated Effect	9%	∞	5.7	15.4	10.6	5.3	4.6	4.0
	9%	6	5.7	19.1	11.6	6.4	5.2	3.7
	9%	4	5.7	20.4	12.3	7.1	5.5	3.5
	9%	2	5.7	25.4	13.9	4.4	3.2	3.0
	11%	∞	7.7	17.2	12.9	7.6	7.3	5.7
	11%	6	7.7	20.6	13.8	7.5	7.1	5.4
	11%	4	7.7	22.5	14.6	8.6	7.8	5.3
	11%	2	7.7	27.9	16.4	6.0	5.8	5.1
	13%	∞	10.5	18.8	15.1	9.6	10.0	8.9
	13%	6	10.5	22.0	15.8	9.3	9.8	8.6
	13%	4	10.5	23.9	16.2	8.9	9.7	8.6
	13%	2	10.5	28.8	18.6	7.6	8.2	8.2

Notes. Actual effect: estimated effect of the policy on average district income and the average income of each wage quintile, using data from Duflo (2001). Simulated effect: effect of the policy predicted using 1996 SAKERNAS microdata, under different assumptions on the return to a year of schooling and the elasticity of substitution between skilled and unskilled workers.

Figure A.35: U.S. Compulsory Schooling Laws: Examples

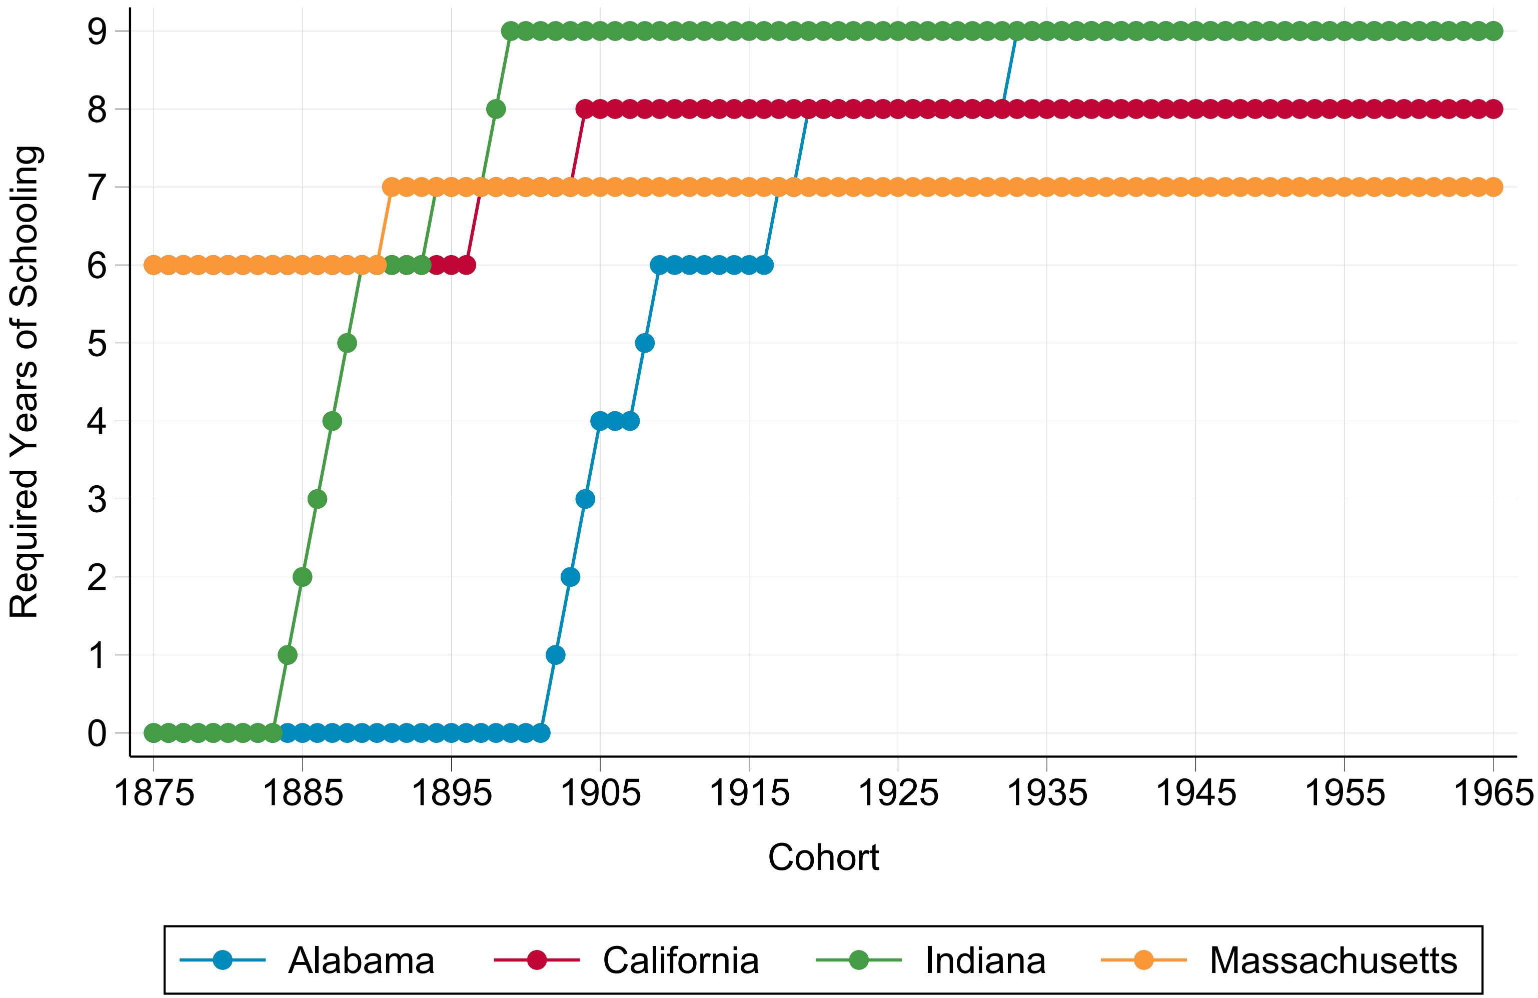

Notes. Author's elaboration based on data from Clay, Lingwall, and Stephens (2021).

Table A.23: U.S. Compulsory Schooling Laws: First Stage
Effect of Required Years of Schooling on State Average Years of Schooling

	(1)	(2)	(3)
Required Years of Schooling	0.191***	0.141***	0.115***
	(0.032)	(0.038)	(0.032)
Region × Year FE	No	Yes	Yes
Extended Controls	No	No	Yes
N	343	343	343

Notes. The unit of observation is the state-year. Required years of schooling: average required years of schooling in each state-year, instrumented using required years of schooling for each state-cohort. Region × Year FE: interacted census region and census year fixed effects. Extended controls: additional controls for 1940 average years of schooling and average personal income interacted with census year fixed effects.

Table A.24: U.S. Compulsory Schooling Laws: Aggregate and Distributional Effects of Schooling

	Baseline			+ Census Region × Year FE			+ Controls for 1940 Educational Attainment and Average Income × Year FE		
	Average Income	Bottom 20% Income	Top 20% Income	Average Income	Bottom 20% Income	Top 20% Income	Average Income	Bottom 20% Income	Top 20% Income
Average Years of Schooling	0.157***	0.437***	0.050*	0.147***	0.431***	0.079*	0.082	0.272**	0.063
	(0.032)	(0.095)	(0.027)	(0.045)	(0.110)	(0.044)	(0.051)	(0.114)	(0.057)
N	343	343	343	343	343	343	343	343	343

Notes. The table reports the effect of state average years of schooling on state average income, the average income of the bottom 20%, and the average income of the top 20%. Columns 1 to 3 control for the demographic, gender, and racial composition of each state, as well as the share of workers with tertiary education. Columns 4 to 6 add census region × year fixed effects. Columns 7 to 9 further add controls for 1940 average years of schooling and average personal income, interacted with survey year dummies. Data from IPUMS census microdata combined with information on compulsory schooling laws from Acemoglu and Angrist (2000) and Clay, Lingwall, and Stephens (2021). Standard errors in parenthesis. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.25: U.S. Compulsory Schooling Laws: Actual vs. Simulated Effects of Educational Expansion

	Parameters		Effect of Increasing Average State Schooling by One Year (%)					
	Return to Schooling	Elasticity of Substitution	Average Income	Q1	Q2	Q3	Q4	Q5
Actual Effect			15.7	43.7	46.0	25.9	12.3	5.0
Simulated Effect	8%	∞	4.5	17.4	9.6	5.2	3.6	2.8
	8%	6	4.5	20.9	11.4	5.7	3.3	2.1
	8%	4	4.5	22.6	12.3	6.0	3.2	1.8
	8%	2	4.5	27.8	14.8	6.8	2.8	0.9
	12%	∞	7.4	24.0	14.0	7.6	5.5	5.9
	12%	6	7.4	27.5	15.7	8.2	5.3	5.3
	12%	4	7.4	29.2	16.6	8.4	5.2	4.9
	12%	2	7.4	34.4	19.1	9.2	4.7	4.0
	16%	∞	11.0	29.2	18.3	9.7	7.4	10.6
	16%	6	11.0	32.7	20.1	10.3	7.2	9.9
	16%	4	11.0	34.5	20.9	10.6	7.1	9.5
	16%	2	11.0	39.7	23.4	11.4	6.7	8.6

Notes. Actual effect: estimated effect of the policy on average state income and the average income of each personal income quintile, combining IPUMS census microdata with information on compulsory schooling laws from Acemoglu and Angrist (2000) and Clay, Lingwall, and Stephens (2021). Simulated effect: effect of the policy predicted using 1960 census microdata, under different assumptions on the return to a year of schooling and the elasticity of substitution between skilled and unskilled workers.

A.4 Education Quality

A natural question is how education quality might have changed from 1980 to 2019, and potential implications for the results presented in this paper. The main source of concern is that if education quality has increased or decreased, then educational attainment becomes a biased measure of actual changes in the education of the labor force. If quality has changed, 1980 and 2019 levels of attainment are not comparable indicators anymore. To make them comparable, one would need to adjust them for changes in quality by, for instance, re-expressing 2019 years of schooling in 1980 quality-adjusted equivalents.

This section discusses existing evidence on the evolution of education quality in developed and developing economies, and attempts to quantify how sensitive are the main results to accounting for these changes. Existing sources provide conflicting stories: some indicators show signs of improvements, while others suggest quality may have declined. Overall, however, there is little evidence of widespread declines in cognitive gains from schooling around the world. Furthermore, accounting for the potential decline in quality observed in some sources leaves the main results unchanged, because this decline appears to have been minor in comparison to the large observed increases in the quantity of schooling.

A.4.1 Trends in Education Quality: Comparison of Available Estimates

A.4.1.1 International Test Scores

The first set of available estimates on changes in quality come from international test scores, which have been increasingly conducted in most countries in the world since the 1990s-2000s. Drawing from various international sources, Angrist et al. (2021) compile test score results for 163 countries over the 2000-2017 period, 122 of which have at least one data point in the 2000s and another data point in the 2010s. The data suggest that education quality has remained broadly stable in most regions, despite some noticeable increases in quality observed in Sub-Saharan Africa and Latin America.

Figure A.36 compares average test scores in the 2000s and 2010s for all countries with available data, based on the database of Angrist et al. (2021). Each data point corresponds to a test score in a given country, for a given education level (primary/secondary) and subject (mathematics/science/reading). All points are very close to the 45-degree line, suggesting that there has been little change in quality over the period. If anything, there has been a slight improvement in average quality: test scores have improved for 170 country-level-subject cells, while they have declined for 100.

For a more restricted number of countries, it is also possible to look at longer-run trends in education quality, based on the database of harmonized test scores compiled by Altinok, Angrist, and Patrinos (2018). Figure A.37 plots the evolution of this indicator since 1970 for a selected number of high- and middle-income countries. The picture that arises is again one of remarkable stability, although some countries have undergone important long-run improvements in schooling quality, including Brazil, Chile, Iran, and South Korea.

A.4.1.2 Conditional Literacy

Test scores arguably provide the best available information, yet they suffer from a critical lack of historical depth for most countries in the world. To make a first step towards closing this gap, Le Nestour, Moscoviz, and Sandefur (2022) exploit information on literacy reported in the Demographic and Health Surveys and the Multiple Indicator Cluster Surveys. These surveys have repeatedly collected information on ability to read in many developing countries since 2000. The enormous advantage of these sources is that they cover adults, which allows tracking education quality across cohorts. This considerably expands the time period, given that the first cohorts covered by the data were born as early as the 1950s. To the best of my knowledge, this represents the only available approach to track historical trends in education quality in the developing world. Based on this, Le Nestour, Moscoviz, and Sandefur (2022) exploit repeated cross-sections to identify changes in education quality, defined as expected literacy at grade 5, across cohorts.

Figure A.38 shows the main result of this exercise, comparing expected literacy at age 5 for cohorts born in 1950-1960s versus 1980-2000. The estimates of Le Nestour, Moscoviz, and Sandefur (2022) point to a clear decline in quality in a number of developing countries. In India, for instance, five years of schooling are found to be associated with about 50% of 1980-2000 cohorts being able to read, compared to 90% of 1950-1960 cohorts.

These results are insightful, but it is important to stress that they do not necessarily imply that the results presented in this paper should be revised downwards for at least four reasons.

First, ability to read is arguably a very partial and noisy measure of quality. For instance, Hermo et al. (2022) show that the decline of vocabulary knowledge in Sweden since the 1960s has been accompanied by a significant increase in logical reasoning skills, which can be rationalized by increasing labor market returns to the latter. In this context, relying solely on one dimension of quality (such as reading) could provide an inaccurate picture of changes in education quality.

Second, identifying trends in the quality of education from repeated cross sections of surveys requires explicitly modeling age, period, and cohort effects. This makes the results much more sensitive to methodological choices, measurement error, and potential sampling differences across survey waves, all of which can be particularly acute in developing countries. The results presented on South Africa in the next section suggest that this is an important concern.

Third, such estimates are not immune to standard problems associated with causal identification (which is also true of test scores). An important source of bias is that improvements in access to schooling in the developing world have been overwhelmingly concentrated among children coming from low-income and lower-educated families (Gethin, 2023b). As a result, lower performance among newly educated cohorts may primarily be the result of greater cognitive and socioeconomic barriers to learning, rather than to changes in the value added of schooling.

Finally, changes in average performance may not necessarily imply lower returns to schooling. Even if newly educated cohorts may have lower levels of cognitive skills, economic returns to schooling for them may still be equal, or even greater (as suggested by the IV returns to schooling presented in the main text), than returns for the rest of the population. What matters is not whether newly skilled workers have lower or higher levels of skills, but instead what are the returns to increased access to schooling for this specific subset of the population. Put differently, differences in average skills may be very different from differences in marginal returns to skill.

A.4.1.3 Country-Specific Sources: Insights from South Africa

Despite these limitations, the cohort-based approach developed by Le Nestour, Moscoviz, and Sandefur (2022) has another advantage: it can be extended to other countries and data sources reporting information on education quality. Indeed, the DHS/MICS are not the only surveys recording information on adult skills. Applying the same methodology to other datasets and country-specific sources provides a fruitful avenue for future research.

While engaging in such a vast data collection and harmonization effort goes beyond the objective of this paper, I draw on previous work (Gethin, 2023c) to document long-run trends in quality in one context: South Africa. Indeed, the General Household Survey has collected detailed information since 2009 on adults' ability to perform six basic operations: writing one's name, reading newspapers and other documents, filling in a form, writing a letter, calculating how much change should be received when buying something, and reading road signs. Information is collected for each household member, with four values ranging from "No difficulty" to "Unable to do."

Drawing on these cross-sections, I run simple regressions relating scores on these indicators to completed years of schooling, controlling for gender, race, province of residence, and survey year fixed effects. Regressions are run by decade of birth to capture cohort changes in education quality. I normalize each dependent variable to range from 0 to 1. Coefficients of interest can then be interpreted as expected literacy obtained from an additional year of schooling.

Figure A.39 plots the resulting evolution of coefficients by decade of birth. Despite some fluctuations and differences in expected gains across items, education quality is estimated to have remained extremely stable from the 1940s to the 1980s. On average, a year of schooling is associated with a 10 percentage point increase in literacy.

This result is puzzling, given that South Africa is one of the countries with the largest estimated decline in quality in the Le Nestour, Moscoviz, and Sandefur (2022) data. Indeed, conditional literacy at grade 5 is found by the authors to have decreased by as much as 20 percentage points, from about 70% to 50% (see figure A.38).

This conflicting evidence suggests that much more research is needed before reaching decisive conclusions on trends in education quality in the developing world. Test scores are perhaps the best data available, but they do not exist before the 2000s in many countries. Cohort trends in literacy are arguably a promising indicator, but data sources and methodologies remain to be further tested and compared.

A.4.1.4 Returns to Schooling Among U.S. Migrants

A last piece of evidence comes from returns to schooling among U.S. migrants. Schoellman (2012) argues that differences in returns to schooling among U.S. migrants originating from different countries provides a good proxy for education quality, because it captures income gains from schooling for individuals having been educated in different countries but working in the same labor market. For instance, returns to schooling are expected to be higher among Swedish migrants than among Congolese migrants, as differences in educational attainment reflect greater differences in accumulated human capital in the former group than in the latter. Schoellman (2012) provides evidence that this indicator is a good proxy for education quality, strongly correlating with GDP per capita and available test scores (see also Rossi, 2022).

The advantage of returns to schooling among migrants is that they can be estimated for an even greater number of countries than cohort trends in literacy studied in Le Nestour, Moscoviz, and Sandefur (2022). Pooling several waves of U.S. censuses, it is also possible to estimate returns to schooling for different cohorts of migrants. Although this analysis is evidently not devoid of limitations—in particular small sample sizes and potential differential selection into schooling across cohorts of a given country—, it can still hopefully shed light on broad long-run trends.

I pool 1980, 1990, and 2000 U.S. censuses, together with all American Community Surveys from 2001 to 2021. I restrict the sample to individuals aged 25 to 65 with positive earned income, who were born outside of the U.S. between 1950 and 1980, and arrived in the U.S. after age 20. I then run the following regressions:

$$y_{icyt} = \zeta_{cy} s_{icyt} + X_{icyt}\beta_{cy} + \mu_t \tag{A.19}$$

With y_{icyt} the log of total yearly earned income of individual i born in country c in decade y (1950s, 1960s, 1970s, or 1980s) and observed in year t. s_{icyt} is completed years of schooling, X_{icyt} are control variables (gender, state of residence, and year of immigration), and μ_t are census/ACS year fixed effects. The parameter of interest is ζ_{cy}, the return to a year of schooling for individuals born in country c in decade y. If education quality has declined substantially, then we should expect ζ_{cy} to have declined over time: a year of schooling should deliver greater returns for migrants born in the 1950s than for migrants born in the 1980s. I run this regression separately for each country of origin × decade of birth cell.

The results of this exercise are presented in figure A.40, which plots population-weighted averages of the estimated returns to schooling by world region of birth and decade of birth. Returns to schooling are lowest among migrants from Latin America and Sub-Saharan Africa and highest among migrants from Europe and the Anglosphere (Canada, Australia, New Zealand, United Kingdom). There are fluctuations across decades, but no clear trend in quality in most regions. Returns have fluctuated at about 4-6% per year of schooling among Latin American migrants, compared to 9-11% among European and Anglosphere natives. The world average varies from 7% to 9% with no clear long-run evolution.[6] This suggests again that changes in education quality are unlikely to play a substantial role in affecting the results presented in this paper.

A.4.2 A Quantification Exercise

While it remains unclear which data source should be preferred, it is still useful to test how sensitive are my main findings to accounting for the potential decline in quality documented in Le Nestour, Moscoviz, and Sandefur (2022). This is somewhat of a heroic task, because it requires (1) extrapolating cohort trends to cover education quality for the entire 1980-2019 working-age populations (2) putting a monetary value on literacy, to build measures of quality-adjusted years of schooling, and (3) extrapolating changes in quality to countries with no available data. This section represents an exploratory attempt at doing so.

A.4.2.1 Methodological Framework

Constructing estimates of quality-adjusted years of schooling requires mapping education quality into equivalent years of schooling. Following the existing literature, I consider the following standard extension of the Mincer-type human capital stock (e.g., Hanushek, Ruhose, and Woessman, 2017):

$$h = \exp(r_L L + r_Q Q) \tag{A.20}$$

With r_L the return to a year of schooling, L average years of schooling, r_Q the return to education quality, and Q an indicator of education quality. The objective is to convert a change in quality from Q to $\tilde{Q}$ into an equivalent change in years of schooling from L to $\tilde{L}$. This equivalence satisfies:

$$\exp(r_L L + r_Q \tilde{Q}) = \exp(r_L \tilde{L} + r_Q Q) \tag{A.21}$$

Rearranging:

$$\tilde{L} = L - \frac{r_Q}{r_L}(Q - \tilde{Q}) \tag{A.22}$$

[6] It is also interesting to investigate differences between cohort trends in returns to schooling among migrants and in literacy rates estimated by Le Nestour, Moscoviz, and Sandefur (2022). The raw cross-country correlation between changes in returns and changes in literacy from the 1960s to the 1980s cohorts is 0.33. This suggests that both sources tell a broadly similar story on which countries have seen education quality decline or improve most.

Calculating quality-adjusted changes in years of schooling thus requires data on changes in education quality ($Q - \tilde{Q}$), as well as the relative value of schooling quality (r_Q) compared to schooling quantity (r_L). I now turn to estimating each of these two components.

A.4.2.2 Estimation of Global Trends in Conditional Literacy

The first step is to estimate ($Q - \tilde{Q}$), the evolution of quality of schooling for the working-age population from 1980 to 2019. The database of Le Nestour, Moscoviz, and Sandefur (2022) provides information on literacy at grade 5 in 86 countries for two cohorts born during the 1952-1999 period (see Le Nestour, Moscoviz, and Sandefur, 2022, Table 7). Starting from these two data points by country, I estimate average conditional literacy for the working-age population.

First, I divide all figures by 5, so that the indicator corresponds to expected literacy per year of education. This ensures that the education quality indicator is comparable to years of schooling.

Second, I linearly interpolate and extrapolate this indicator backwards and forwards, to cover all cohorts born from 1915 to 1994. This is a very conservative assumption: it amounts to considering that education quality continued to decline at the same pace after the last cohort observed, and was already declining at the same pace from 1915 until the first cohort observed. This is unlikely to be true, given evidence documented above on the stability or even rise of education quality in many countries since the 2000s.

Third, I construct measures of average education quality of the working-age population. To do this, I average the indicator over all cohorts aged 25 to 65 in a given year, weighted by the population of each cohort. Data on population by age is taken from the United Nations' World Population Prospects. The result is an indicator of education quality covering the working-age population of each country from 1980 to 2019, corresponding to average expected literacy per year of schooling.

Finally, in the absence of data for the rest of the world, I impute the indicator for missing countries using three polar scenarios. The benchmark scenario assumes that education quality in missing countries has declined at the same pace as the average decline observed over the 86 countries. The upper bound assumes that it has not declined. The lower bound assumes that it has declined at the speed of India, that is, at a very fast pace (see figure A.38). I view this last case as an extreme and implausible scenario, given above-mentioned evidence on the stability or rise of test scores in many countries.

Figure A.41 compares education quality of the working-age population in 1980 and 2019 for countries with available data. The overall pattern and ranking of countries is similar to the one visible in figure A.38. However, the change in quality appears less dramatic, because this figure compares education quality for the overall population rather than across cohorts. In India, for instance, literacy per year of schooling declined by about 8 percentage points, from 22 to 14.

A.4.2.3 Estimation of Returns to Literacy

The second step is to estimate r_Q/r_L, the returns to literacy relative to a year of schooling. This requires data on personal income, years of schooling, and literacy at the individual level. I was able to find four high-quality surveys covering these three variables: the Brazilian 2015 PNAD survey, the Indonesian 1998 SUSENAS survey, the Pakistani 2018 HIES survey, and the South African 2019 GHS survey. In each of these four countries, I estimate the relative returns to literacy by running two regressions: a regression relating the log of total personal income to literacy, and a regression relating the log of total personal income to years of schooling, controlling for gender, potential experience, and potential experience squared in each case. I restrict the sample to workers with either no schooling or basic education, to make sure that the two estimates are comparable (nearly all workers with more than basic education are literate).

The results are presented in table A.26. Returns to schooling range from 3% to 8% per year of basic education, while returns to literacy range from 18 to 39 log points. The ratio between the two coefficients, corresponding to r_Q/r_L, is very similar across countries, ranging from 5 in Pakistan to about 6.5 in Indonesia. I take a value of 6 to construct measures of quality-adjusted years of schooling in what follows. This amounts to assuming that moving the entire population from being illiterate to literate is equivalent to increasing average schooling by 6 years.

A.4.2.4 Results

Having estimated changes in education quality and its price relative to education quantity, one can now construct measures of quality-adjusted years of schooling. In practice, I set 1980 as the benchmark year, and adjust estimates of average years of schooling in all other years from 1981 to 2019 so that they reflect the quality observed in 1980. For instance, quality-adjusted years of schooling in 2019 are calculated as $\tilde{L}_{2019} = L_{2019} - \frac{r_Q}{r_L}(Q_{2019} - Q_{1980})$, with L_{2019} unadjusted years of schooling observed in 2019, $\frac{r_Q}{r_L} = 6$, and $Q_{2019} - Q_{1980}$ the change in expected literacy per year of schooling from 1980 to 2019. This approach thus amounts to "deflating" years of schooling observed from 1981 to 2019 to express them in 1980 equivalents.

Figure A.42 compares the evolution of average years of schooling in the world as a whole, before and after adjusting for changes in education quality. The unadjusted indicator rose from 5 to 8.5. Years of schooling expressed in 1980 equivalents rose from 5 to 8.1-8.3. Adjusting for education quality thus reduces average years of schooling today by at most 0.5 years (or 6%), and the overall increase in education since 1980 by at most 0.4 years (or 9%).

Figure A.43 compares the share of growth explained by global income percentile before and after making the lower bound adjustment (assuming that education quality declined as fast as in India for all countries with missing data). The two lines are barely distinguishable: even under strong assumptions on the decline in education quality, the main result remains almost unchanged. Overall, the share of growth explained by education declines by about 2 to 6 percentage points depending on the percentile considered, with the greatest changes observed at the upper-middle of the income distribution. The results presented in this paper thus appear to be strongly robust to potential changes in education quality observed since 1980.

Figure A.36: Harmonized Test Scores by Country: 2000-2009 vs. 2010-2019

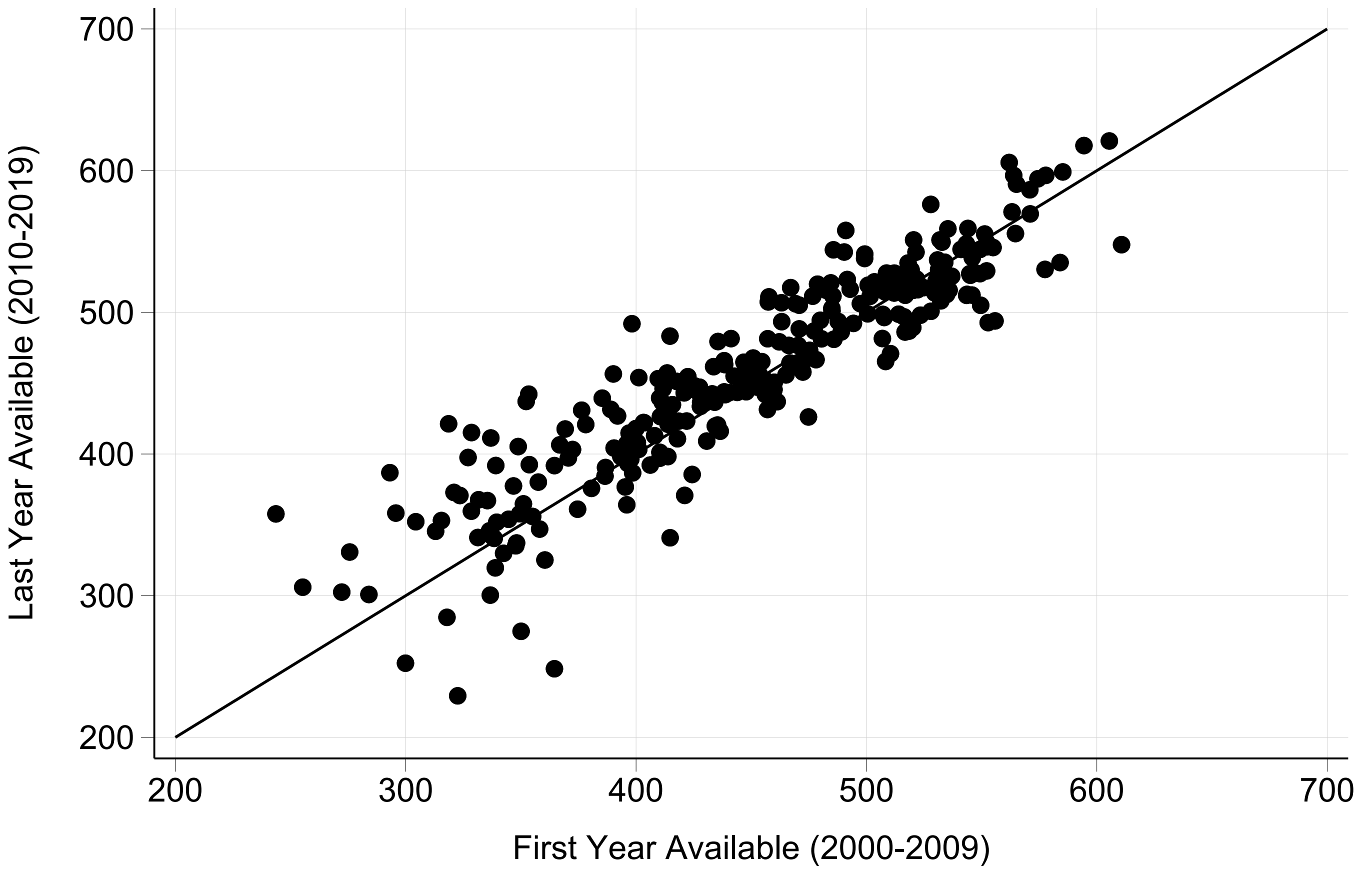

Source: Angrist et al. (2021). Each point corresponds to a test score reported for a given country × education level (primary/secondary) × subject (maths/science/reading).

Figure A.37: Long-Run Trends in Test Scores in Selected Countries, 1970-2015

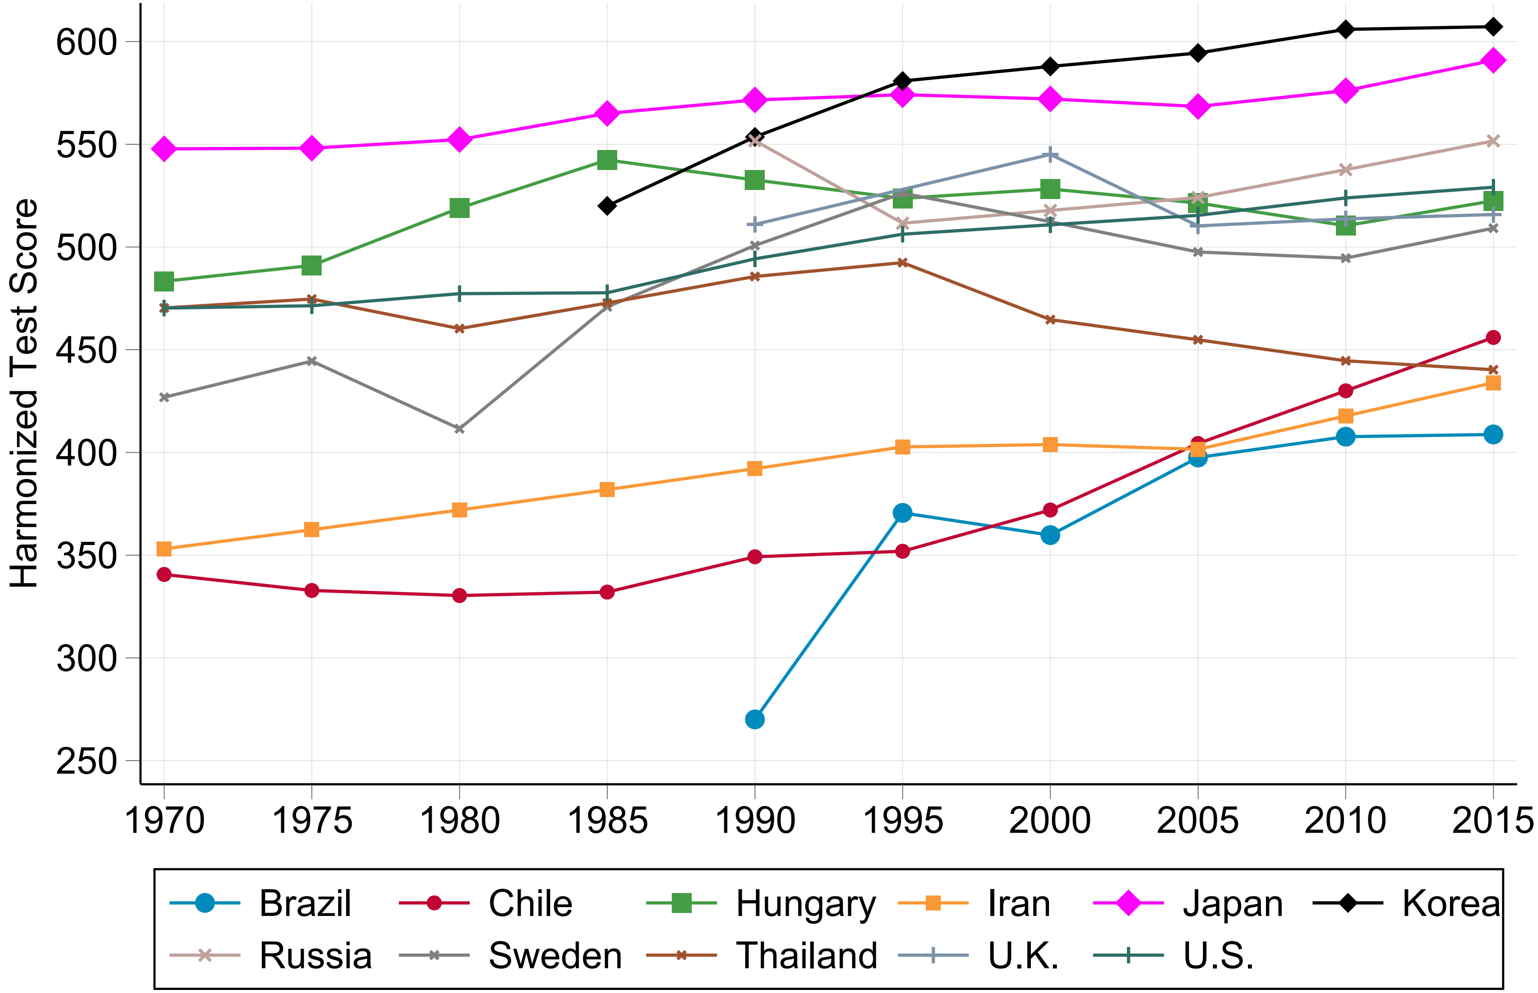

Source: Altinok, Angrist, and Patrinos (2018).

Figure A.38: Literacy at Grade 5: 1950-1960 versus 1980-2000 Cohorts

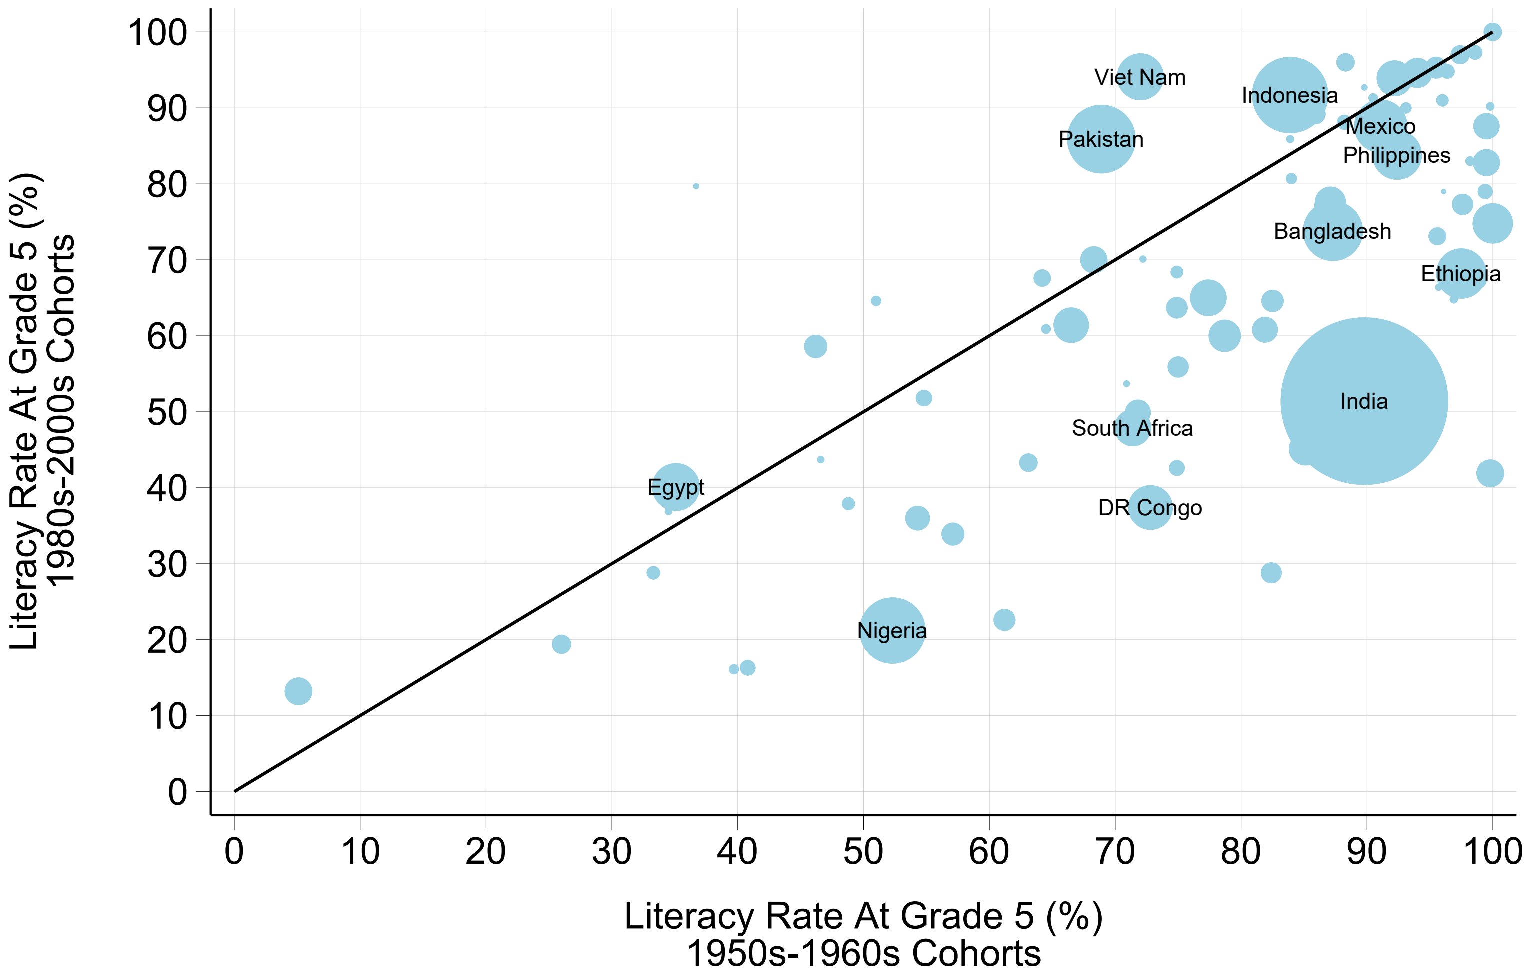

Source: Le Nestour, Moscoviz, and Sandefur (2022).

Figure A.39: Long-Run Trends in Education Quality in South Africa:
Cognitive Gains Per Year of Basic Education, 1940-1980 Cohorts

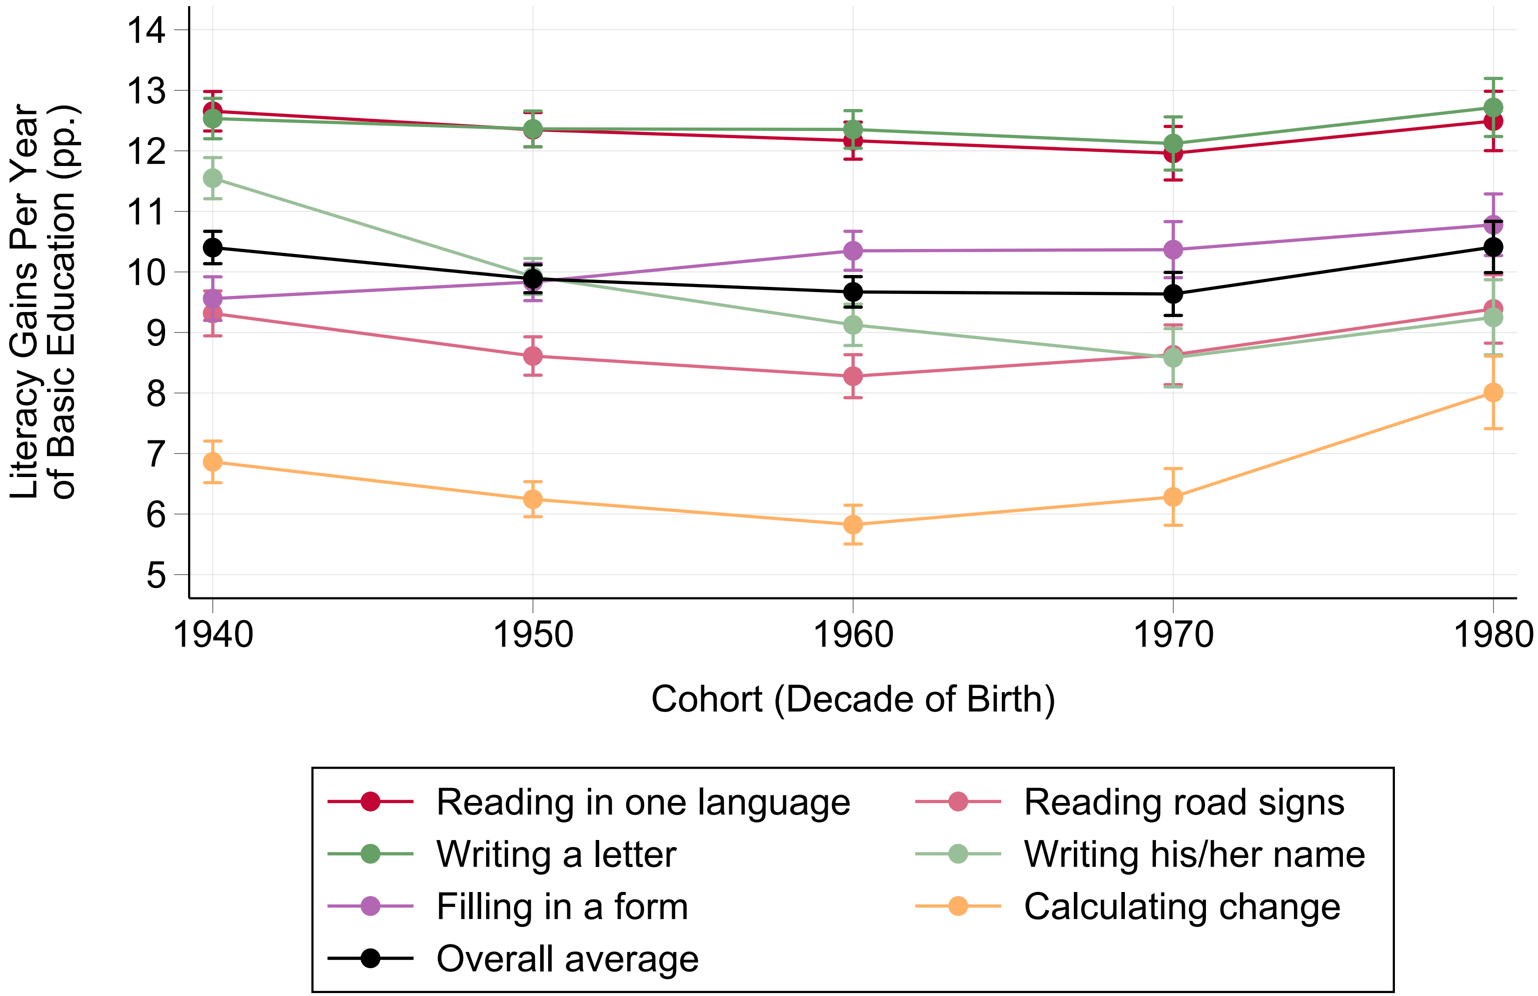

Note: author's calculations using South African General Household Surveys (2009-2013). Each data points corresponds to the coefficient of a regression of corresponding literacy scores (0-1) on years of schooling, restricting the sample to individuals with 0 to 6 years of education.

Figure A.40: Trends in Returns to Schooling Across Cohorts of U.S. Migrants

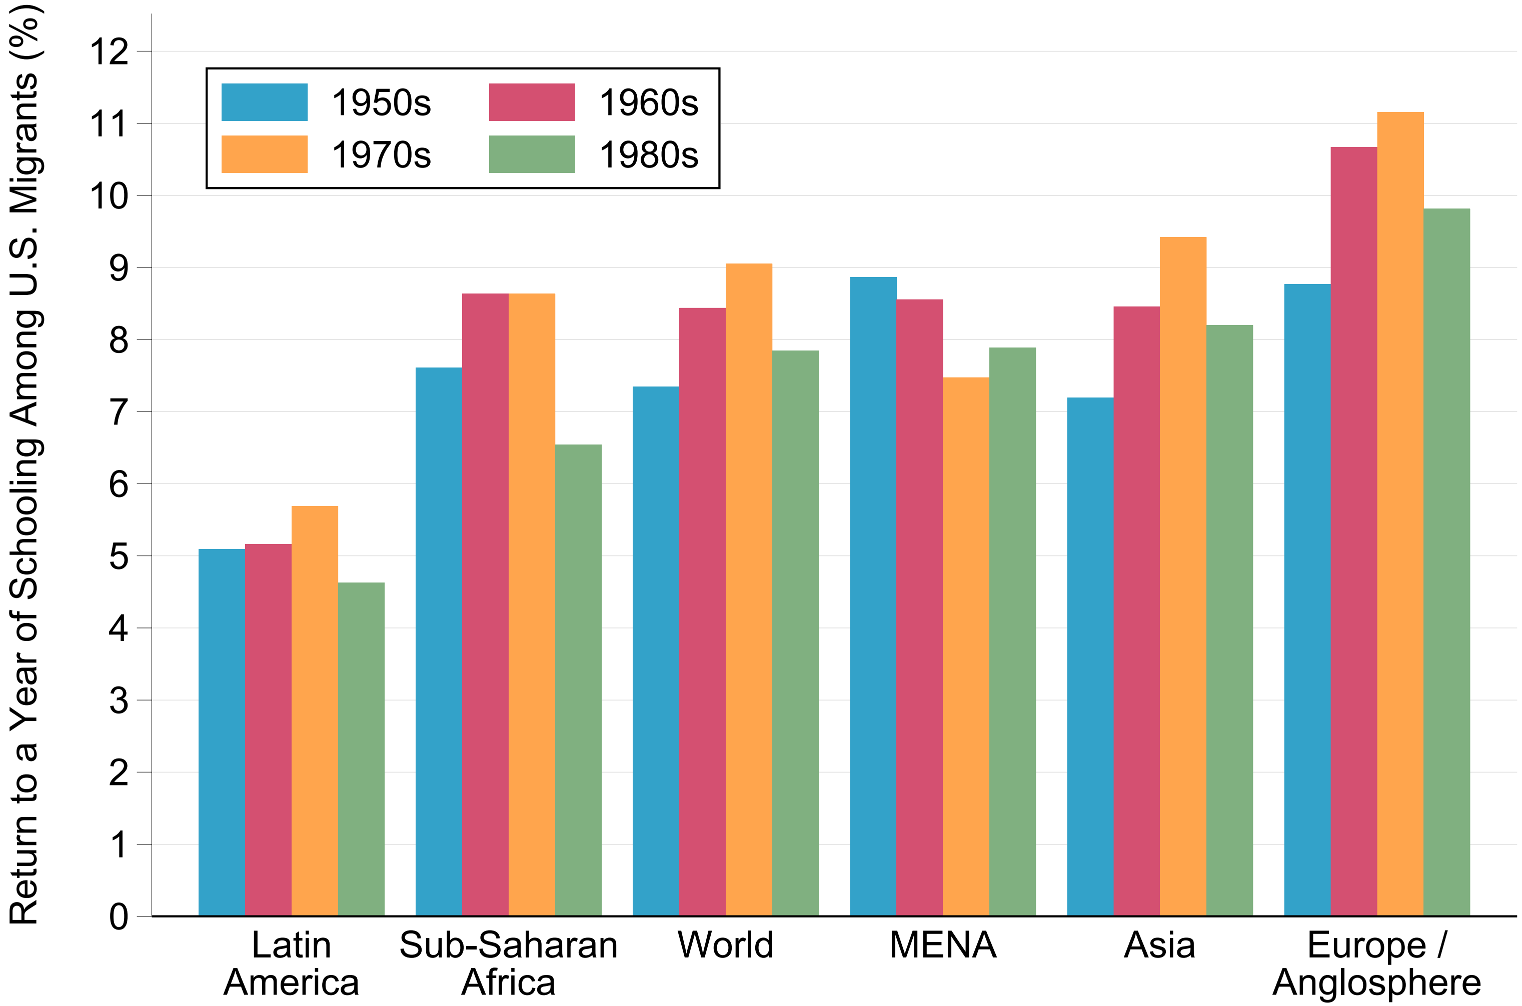

Note: author's calculations using U.S. censuses and American Community Surveys. Each bar corresponds to the population-weighted average of returns to schooling estimated for a given country of origin × decade of birth cell.

Figure A.41: Literacy Gains Per Year of Schooling: 1980 versus 2019 Working-Age Population

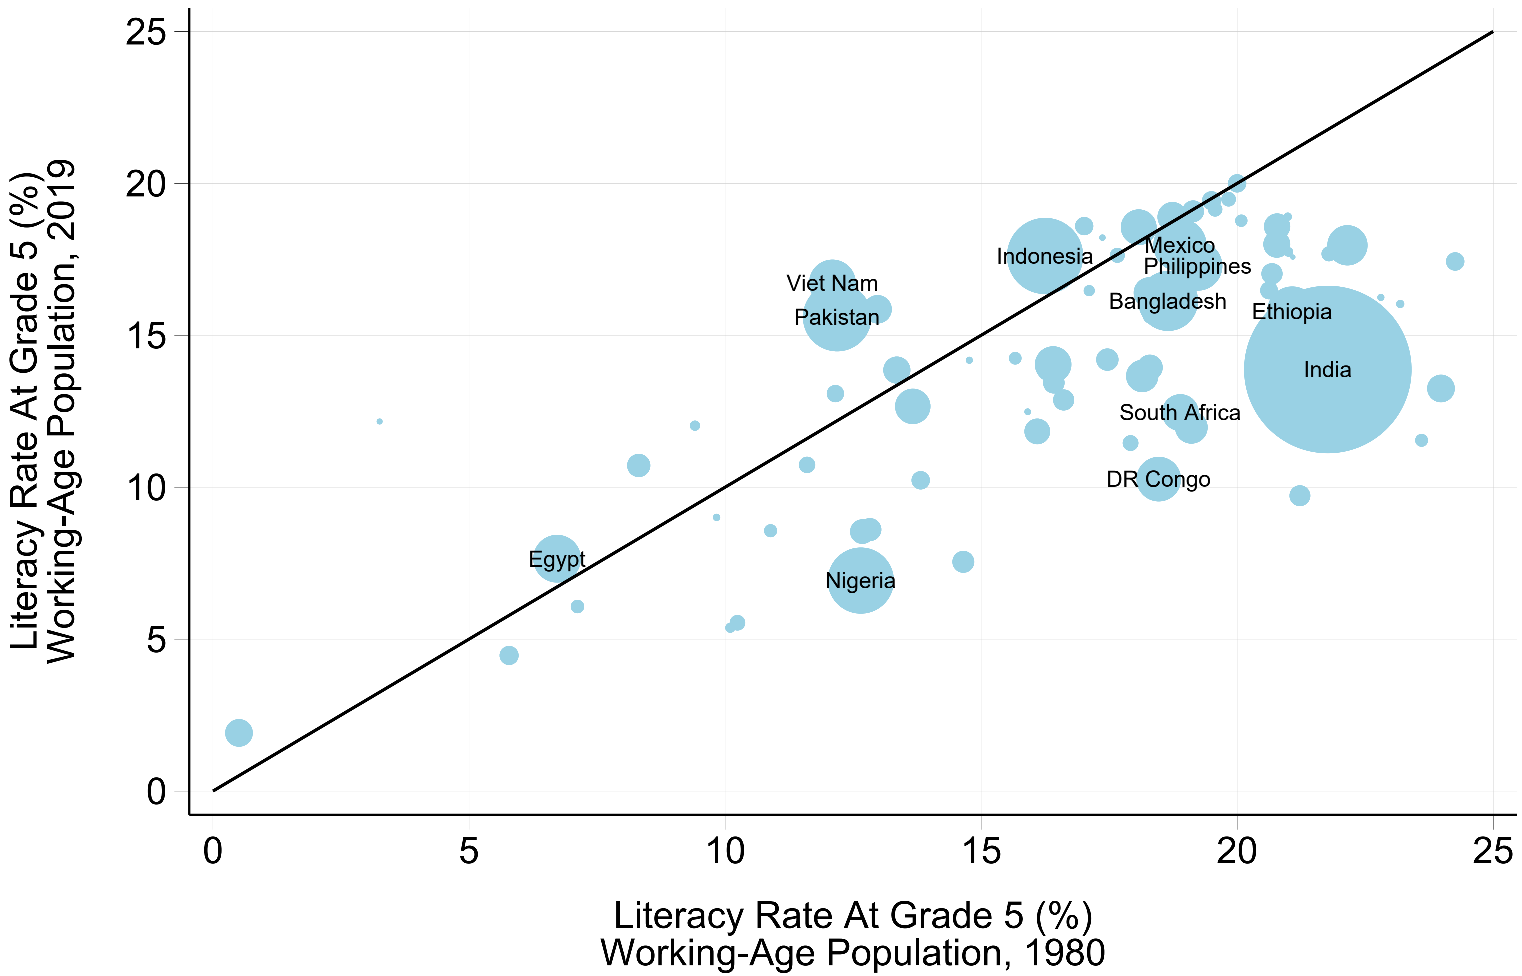

Note: author's estimation combining data from Le Nestour, Moscoviz, and Sandefur (2022) and United Nations World Population Prospects.

Table A.26: Returns to Literacy

	Brazil	Indonesia	Pakistan	South Africa
Return to Literacy	0.39***	0.31***	0.31***	0.18***
	(0.01)	(0.01)	(0.01)	(0.05)
Return to Schooling	0.08***	0.05***	0.06***	0.03***
	(0.00)	(0.00)	(0.00)	(0.01)
Literacy / Schooling	5.10	6.53	5.03	5.66

Notes. The table reports estimates of returns to literacy, returns to schooling, and the ratio between the two. The coefficient on literacy corresponds to a regression of the log of personal income on literacy; the coefficient on years of schooling corresponds to a separate regression of the log of personal income on years of schooling. Both regressions control for gender, potential experience, and potential experience squared in each country. Data sources: 2015 Brazil PNAD survey, 1998 Indonesia SUSENAS survey, 2018 Pakistan HIES survey, 2019 South Africa GHS survey. Standard errors in parenthesis. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure A.42: Global Average Years of Schooling: Unadjusted versus Quality-Adjusted Estimates Using Cohort Conditional Literacy Rates

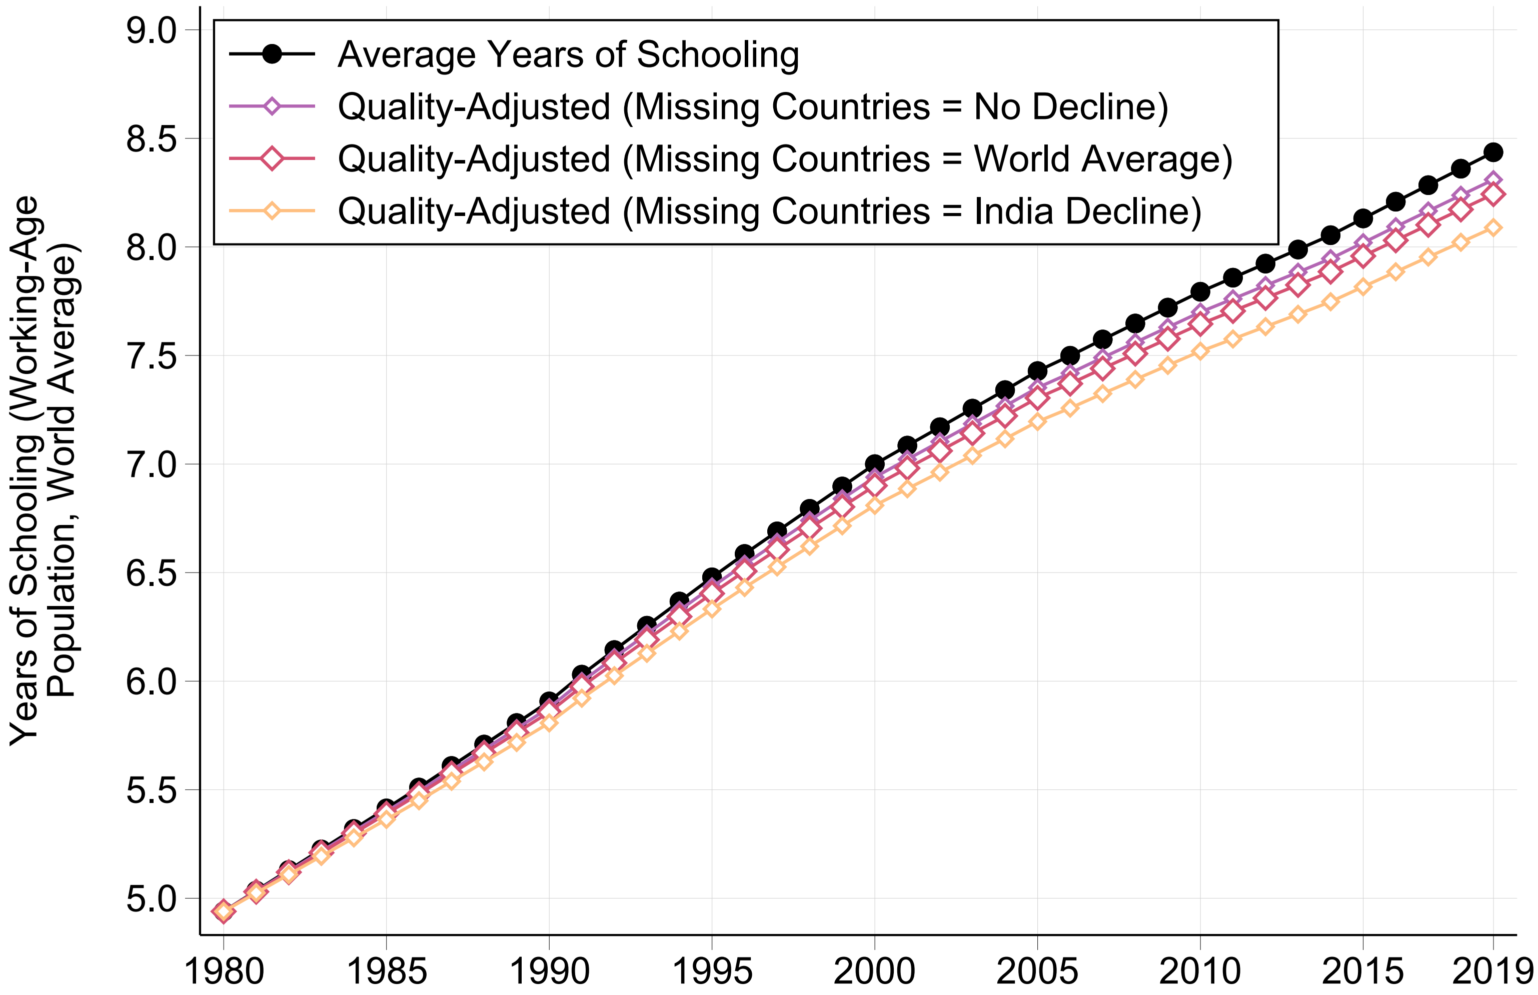

Note: unadjusted estimates correspond to average years of schooling of the world working-age population. Quality-adjusted estimates correct years of schooling for the decline in education quality estimated by Le Nestour, Moscoviz, and Sandefur (2022), so that years of schooling are expressed in 1980 equivalents throughout the period. Upper bound: education quality assumed to have remained constant for all countries with no data from Le Nestour, Moscoviz, and Sandefur (2022). Benchmark: countries with missing data are attributed the average yearly change in quality estimated across all countries covered. Lower bound: countries with missing data are attributed the decline in quality observed in India.

Figure A.43: Share of Growth Explained by Education: Benchmark Versus Lower Bound on Decline in Education Quality

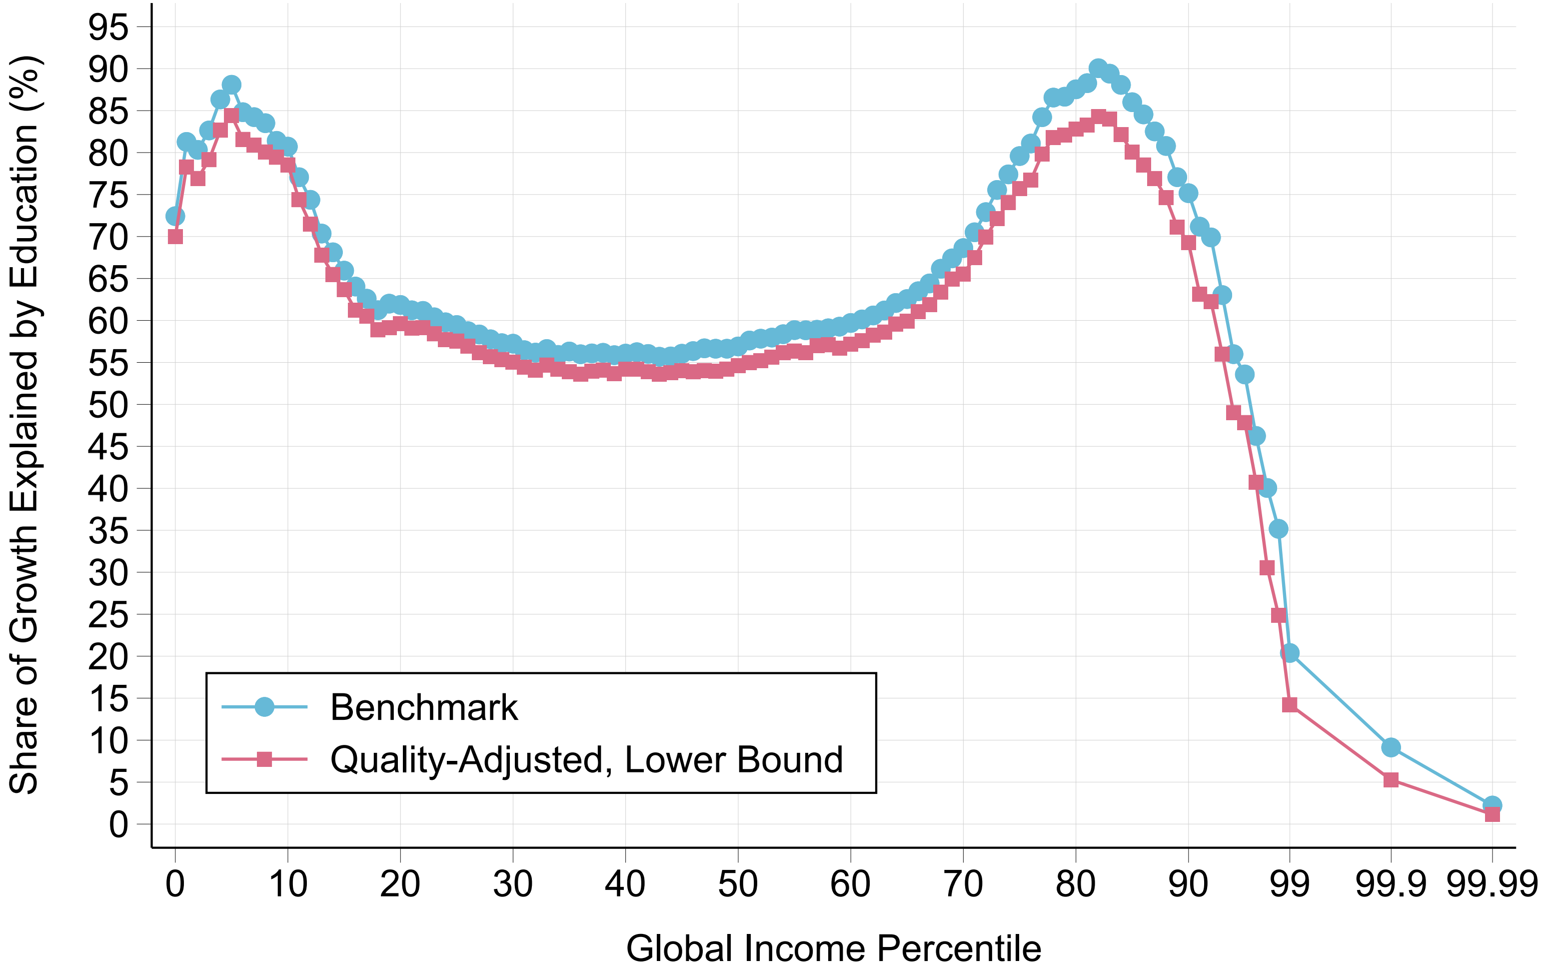

Note: Quality-adjusted estimates correct years of schooling for the decline in education quality estimated by Le Nestour, Moscoviz, and Sandefur (2022), so that years of schooling are expressed in 1980 equivalents throughout the period. Countries with missing data are attributed the decline in quality observed in India.

A.5 Data Appendix: Survey Microdata

The survey microdata covering education and earnings in 150 countries used in this paper come from four main data sources.

ILO Microdata The main data source is a set of harmonized household surveys that were collected and compiled by the International Labor Organization. The ILO database covers over 1,400 surveys fielded in 136 countries from 1990 to 2022. In the main analysis of this paper, I use the last survey available in each country. However, I also exploit historical surveys in the analysis of backward versus forward accounting presented in section 1.5.2. The database presents itself as a single harmonized microfile. The main variables are country, year, household ID, sample weight, wage income (from main job, second job, and all jobs combined), self-employment income (from main job, second job, and all jobs combined), age, gender, education, labor force participation, occupation (ISCCO-08), industry, and rural-urban location. I define personal income as the sum of all wage and self-employment income received by an individual. I drop all zeros and missing values, so that the sample is restricted to all individuals with strictly positive personal income.

European Statistics on Income and Living Conditions Although the ILO microdata do cover European countries, the coding of educational attainment is broader than in the original microfiles, so I decide to rely on my own data collection. The European Statistics on Income and Living Conditions (EU-SILC) cover detailed information on personal income and education in 32 countries every year from 2003 to 2020. I harmonize EU-SILC surveys in the same way as those of the ILO, defining personal income as the sum of individual wage and mixed income. I then replace all ILO surveys by this microfile, with the exception of France, Portugal, and Switzerland, for which the ILO provides national labor force surveys of even better quality.

Life in Transition Survey For 10 Eastern European and Central Asian countries not covered by the ILO, I rely on the Life in Transition Survey (LITS). These are Azerbaijan, Belarus, Georgia, Kyrgyzstan, Kosovo, Kazakhstan, North Macedonia, Montenegro, Ukraine, and Uzbekistan, for which labor force or household living standards surveys are unfortunately not publicly accessible at the time of writing. The LITS is far from being ideal, with sample sizes of only 3,000-5,000 in each country, yet it is to the best of my knowledge the only data source available to measure individual incomes and education. I use the last wave of the LITS, fielded in 2016, which I harmonize in the same way as the ILO.

Country-Specific Surveys Finally, I collect and harmonize surveys from country-specific data portals to cover 13 additional countries: China, Iraq, India, Japan, Mozambique, Morocco, Russia, Somalia, South Africa, South Korea, South Sudan, Tunisia, and the United States.

For seven countries, I was available to find and harmonize a high-quality survey providing detailed information on individual incomes and education. This type of survey was available for China (2018 Chinese Household Income Project), India (2019 Periodic Labor Force Survey), Russia (2019 Russia Longitudinal Monitoring Survey), South Korea (2019 Korean Labor and Income Panel Study), Tunisia (2014 Labor Force Survey), South Africa (2019 General Household Survey), and the United States (2019 Current Population Survey).

For the remaining six countries, I rely on surveys of lower quality or only providing information on household expenditure. For Japan, in the absence of better publicly available data, I use the 2017 general household survey, which does cover individual income and education but has a small sample size (about 1,000). I use household income and expenditure surveys for Iraq (Household Socio-Economic Survey), Mozambique (Inquérito aos orcamentos familiares), Morocco (Household Expenditure Survey), Somalia (High Frequency Survey), and South Sudan (High Frequency Survey), which provide information on individual employment and education, as well as total household expenditure, but not on individual incomes. In the absence of better information, I proxy personal income by splitting equally household expenditure among adults in employment, excluding unemployed or inactive individuals as well as children.

Figure A.44: Survey Data Coverage: Year Covered by Each Survey

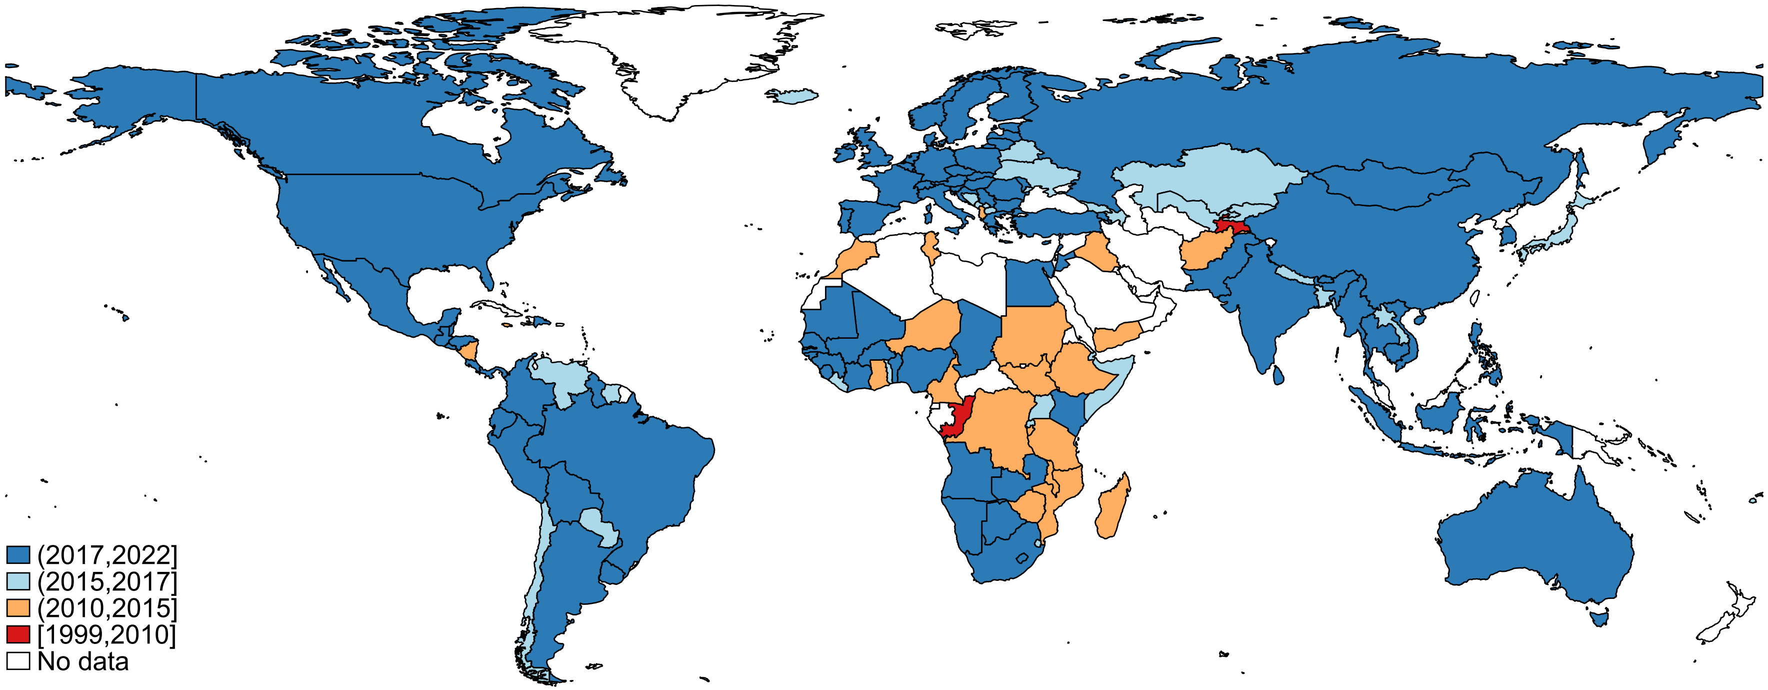

Notes. Colored countries are those covered by the survey microdata. Colors correspond to the year during which each survey was fielded.

Table A.27: Survey Data Sources

Country	Source	Survey Year
Europe		
Albania	Living Standards Survey	2012
Austria	EU Statistics on Income and Living Conditions	2019
Belarus	Life in Transition Survey	2016
Belgium	EU Statistics on Income and Living Conditions	2019
Bosnia and Herzegovina	Labour Force Survey	2016
Bulgaria	EU Statistics on Income and Living Conditions	2019
Croatia	EU Statistics on Income and Living Conditions	2019
Czechia	EU Statistics on Income and Living Conditions	2019
Denmark	EU Statistics on Income and Living Conditions	2019
Estonia	EU Statistics on Income and Living Conditions	2019
Finland	EU Statistics on Income and Living Conditions	2019
France	Employment Survey	2019
Germany	EU Statistics on Income and Living Conditions	2019
Greece	EU Statistics on Income and Living Conditions	2019
Hungary	EU Statistics on Income and Living Conditions	2019
Iceland	EU Statistics on Income and Living Conditions	2018
Ireland	EU Statistics on Income and Living Conditions	2019
Italy	Labour Force Survey	2019
Latvia	EU Statistics on Income and Living Conditions	2019

Lithuania	EU Statistics on Income and Living Conditions	2019
Luxembourg	EU Statistics on Income and Living Conditions	2019
Malta	EU Statistics on Income and Living Conditions	2019
Moldova	Labour Force Survey	2019
Montenegro	Life in Transition Survey	2016
Netherlands	EU Statistics on Income and Living Conditions	2019
North Macedonia	Life in Transition Survey	2016
Norway	EU Statistics on Income and Living Conditions	2019
Poland	EU Statistics on Income and Living Conditions	2019
Portugal	Employment Survey	2019
Romania	EU Statistics on Income and Living Conditions	2019
Russia	Russia Longitudinal Monitoring Survey	2019
Serbia	Labour Force Survey	2019
Slovakia	EU Statistics on Income and Living Conditions	2019
Slovenia	EU Statistics on Income and Living Conditions	2019
Spain	EU Statistics on Income and Living Conditions	2019
Sweden	EU Statistics on Income and Living Conditions	2019
Switzerland	Labour Force Survey	2019
Ukraine	Life in Transition Survey	2016
United Kingdom	Labour Force Survey	2018
Northern America		
Canada	Labour Force Survey	2019

USA	Current Population Survey	2019
Latin America		
Argentina	Permanent Household Survey, Urban	2019
Barbados	Survey on Living Conditions	2016
Belize	Labour Force Survey	2019
Bolivia	Continuous Employment Survey	2019
Brazil	Continuous National Household Sample Survey	2019
Chile	National Survey on Socio-Economic Conditions	2017
Colombia	Integrated Household Survey	2019
Costa Rica	National Household Survey	2019
Dominican Republic	Continuous National Labour Force Survey	2019
Ecuador	National Survey on Employment	2019
El Salvador	Multi-purpose Household Survey	2019
Guatemala	Monthly Employment and Income Survey	2019
Guyana	Labour Force Survey	2019
Honduras	Continous Multi-Purpose Household Survey	2019
Jamaica	Labour Force Survey	2014
Mexico	National Occupation and Employment Survey	2019
Nicaragua	National Household Survey on Measuring Living Conditions	2014
Panama	Labour Market Survey	2019
Paraguay	Continous Household Survey	2017
Peru	National Household Survey	2019
Suriname	Survey on Living Conditions	2016

Trinidad and Tobago	Continuous Sample Survey of the Population	2016
Uruguay	Continous Household Survey	2019
Venezuela	Household Sample Survey	2017
Asia		
Afghanistan	Households Living Conditions Survey	2014
Australia	Household, Income and Labour Dynamics Survey	2019
Bangladesh	Labour Force Survey	2017
Bhutan	Labour Force Survey	2019
Brunei Darussalam	Labour Force Survey	2014
Cambodia	Labour Force Survey	2019
China	China Household Income Project	2018
Fiji	Employment, Unemployment Survey	2016
India	Periodic Labour Force Survey	2019
Indonesia	National Labour Force Survey	2019
Japan	General Social Survey	2017
Kazakhstan	Life in Transition Survey	2016
Kosovo	Life in Transition Survey	2016
Kyrgyzstan	Life in Transition Survey	2016
Lao	Labour Force Survey	2017
Maldives	Household Income and Expenditure Survey	2019
Mongolia	Labour Force Survey	2019
Myanmar	Labour Force Survey	2019
Nepal	Labour Force Survey	2017

Pakistan	Labour Force Survey	2019
Philippines	Labour Force Survey	2018
South Korea	Korean Labor and Income Panel Study	2019
Sri Lanka	Labour Force Survey	2018
Tajikistan	Living Standards Survey	2009
Thailand	Household Socio-Economic Survey	2019
Timor-Leste	Labour Force Survey	2016
Tonga	Labour Force Survey	2018
Uzbekistan	Life in Transition Survey	2016
Vietnam	Labour Force Survey	2019
Middle East and North Africa		
Armenia	Household Labour Force Survey	2019
Azerbaijan	Life in Transition Survey	2016
Cyprus	EU Statistics on Income and Living Conditions	2019
Egypt	Labour Force Sample Survey	2018
Georgia	Life in Transition Survey	2016
Iraq	Household Socio-Economic Survey	2012
Jordan	Employment and Unemployment Survey	2019
Lebanon	Labour Force Survey	2019
Morocco	Household Expenditure Survey	2014
Palestine	Labour Force Survey	2019
Sudan	Household Survey	2011
Tunisia	Labor Force Survey	2014

Turkey	Household Labour Force Survey	2019
Yemen	Labour Force Survey	2014
Sub-Saharan Africa		
Angola	Employment Survey	2019
Benin	Integrated Survey of Household Living Conditions	2018
Botswana	Multi-Topic Household Survey	2019
Burkina Faso	Regional Integrated Survey on Employment and the Informal Sector	2018
Burundi	Living Standards Survey	2014
Cabo Verde	Continuous Multi-Objective Survey	2015
Cameroon	Household Survey	2014
Chad	Modular and Integrated Household Survey on Living Conditions	2018
Comoros	National Survey on Employment and the Informal Sector	2014
Côte d'Ivoire	National Survey on the Employment Situation	2019
Democratic Republic of the Congo	Survey on Employment and household's living conditions	2012
Djibouti	Djiboutian Household Survey	2017
Eswatini	Labour Force Survey	2016
Ethiopia	National Labor Force Survey	2013
Gambia	Labour Force Survey	2018
Ghana	Labour Force Survey	2015
Guinea	National Survey on Employment and the Informal Sector	2019
Guinea-Bissau	Harmonized Survey on Household Living Conditions	2018
Kenya	Household Budget Survey	2019
Lesotho	Labour Force Survey	2019

Liberia	Labour Force Survey	2017
Madagascar	National Survey on Employment and the Informal Sector	2015
Malawi	Labour Force Survey	2013
Mali	Continous Household Employment Survey	2018
Mauritania	Living Standards Survey	2019
Mauritius	Continuous Multi-Purpose Household Survey	2019
Mozambique	Inquérito aos orçamentos familiares	2014
Namibia	Labour Force Survey	2018
Niger	National Survey on Household Living Conditions	2014
Nigeria	Socio Economic Survey	2019
Republic of the Congo	Employment Survey	2009
Rwanda	Labour Force Survey	2017
Senegal	National Employment Survey	2019
Sierra Leone	Integrated Household Survey	2018
Somalia	High Frequency Survey	2017
South Africa	General Household Survey	2019
South Sudan	High Frequency Survey	2015
Tanzania	National Household Budget Survey	2012
Togo	Regional Integrated Survey on Employment and the Informal Sector	2017
Uganda	National Labour Force Survey	2017
Zambia	Labour Force Survey	2019
Zimbabwe	Labour Force Survey	2014

A.6 Data Appendix: Educational Attainment Data

A.6.1 Data Sources

Barro-Lee Database The primary data source used to measure the evolution of educational attainment is the database compiled by Barro and Lee (2013) and updates.[7] The database covers the distribution of educational attainment by age group and gender in 146 countries at five year intervals from 1950 to 2015. It covers 123 countries out of the 150 countries studied in this paper. The education categories are no schooling, incomplete primary, complete primary, incomplete secondary, complete secondary, incomplete tertiary, and complete tertiary. I interpolate linearly the share of individuals belonging to each category between missing years, and extrapolate linearly educational attainment by age and gender after 2015, so as to cover the entire 1980-2019 period.

IPUMS and Survey Data For the 27 countries absent from the Barro-Lee database, I rely on census and survey data. For Burkina Faso (1986-2006), Ethiopia (1984-2007), Guinea (1983-2014), and Palestine (1997-2017), the data source is the census microdata samples available from IPUMS International. For India, which is covered by the Barro-Lee database but displays somewhat erratic trends, I rely instead on the education modules of the national sample survey (1983-2017), which I collected and harmonized for the purpose of this paper. For the remaining 22 countries, in the absence of better data, I use cohort-level trends in educational attainment observed in the surveys collected in this paper.[8] I first aggregate the distribution of educational attainment by cohort and gender in each survey. I then derive estimates of educational attainment of the 1980 to 2019 working-age populations by taking the weighted average across cohorts belonging to the working-age population in the corresponding year.

A.6.2 Matching Survey and Aggregate Data

To derive accurate estimates of counterfactual income absent educational expansion, it is important to make sure that educational attainment in the survey data matches perfectly aggregate data used to derive the counterfactual. Although education levels do correlate strongly in the two sources, some inconsistencies remain. For instance,

[7] See `http://www.barrolee.com/`.

[8] The countries are Angola, Azerbaijan, Bosnia and Herzegovina, Bhutan, Belarus, Cabo Verde, Cambodia, Chad, Djibouti, Georgia, Guinea-Bissau, Kosovo, Lebanon, Montenegro, Madagascar, Macedonia, Nigeria, Somalia, Suriname, South Sudan, Timor-Leste, and Uzbekistan.

aggregate and survey data sometimes report incomplete degrees as complete and sometimes do not, or code lower secondary education as primary education. To make sure that the two sources coincide, I first manually recode some categories in survey and/or aggregate data, country by country, by visually inspecting the distribution of educational attainment in the two sources. The result of this manual recoding process is displayed in figures A.45, A.46, A.47, and A.48, which compare the share of the working-age population with no schooling, primary education, secondary education, and tertiary education in survey versus aggregate data. The two sources end up very close to each other after recoding.

Second, I perform a final small adjustment to the sample weights of each survey to make sure that education levels by age and gender match perfectly in the two sources. I combine aggregate data on the distribution of attainment with data on total population by age and gender from the UN to derive estimates of the total number of individuals belonging to each of 40 education-age-gender cells. I then use linear calibration to ensure that total weights match the total population belonging to each cell in each country. The result is a new weight variable that ensures that the distribution of educational attainment by age and gender (and for the working-age population as a whole) in the survey data matches perfectly that observed in aggregate data.

Figure A.45: Barro-Lee Versus Survey Data: Share of Working-Age Population With No Schooling

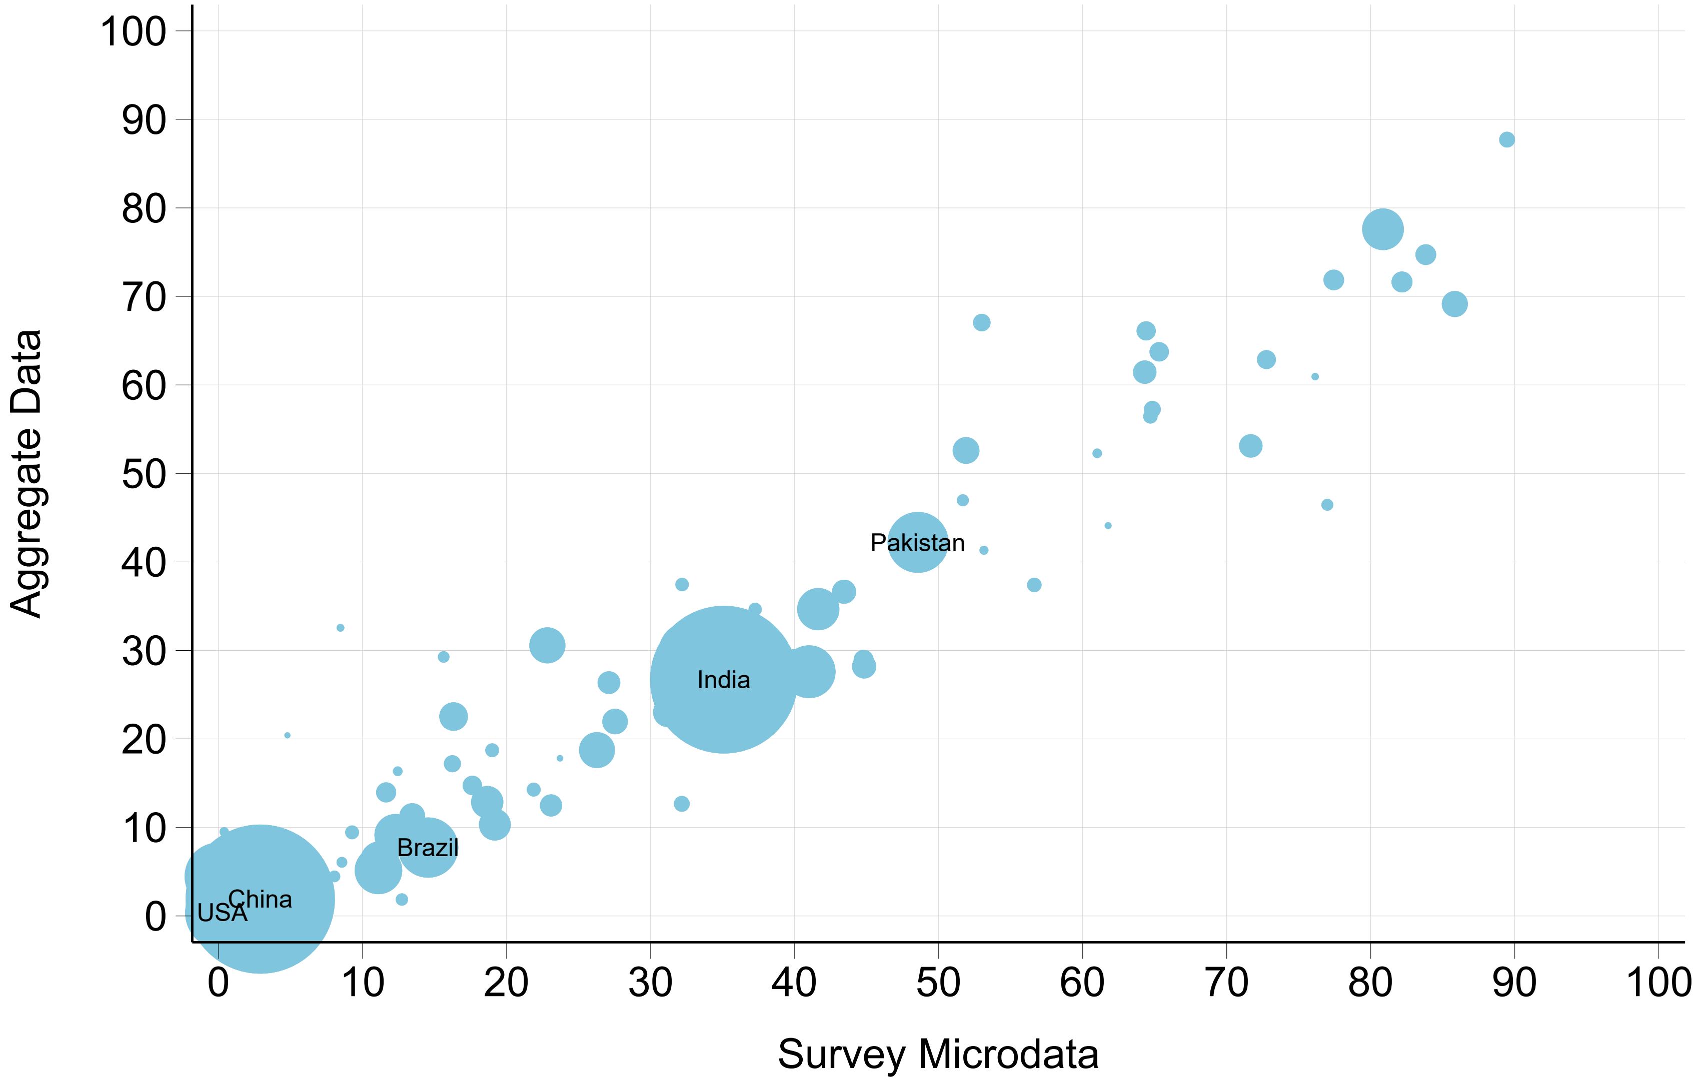

Notes. The figure compares estimates of the share of the working-age population with no schooling in the survey microdata (x-axis) and aggregate data from Barro and Lee (2013) and other sources (y-axis), after manual reclassification of educational categories in each country.

Figure A.46: Barro-Lee Versus Survey Data: Share of Working-Age Population With Primary Education

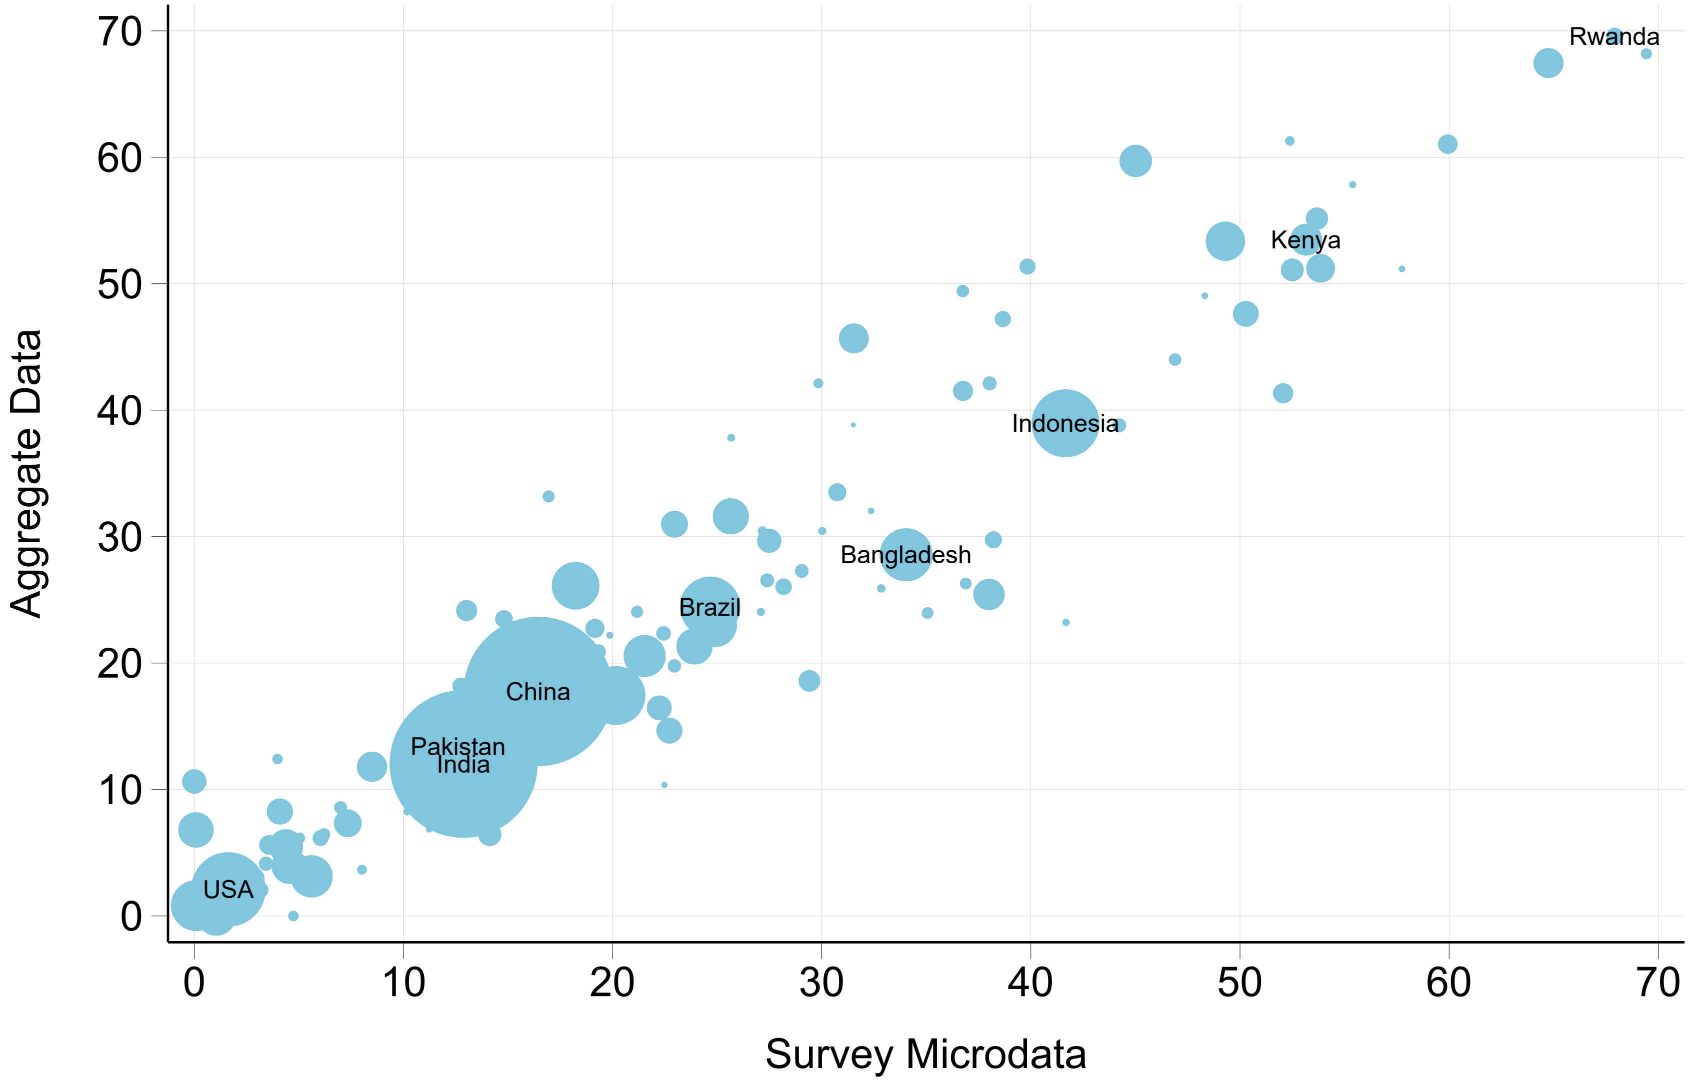

Notes. The figure compares estimates of the share of the working-age population with primary/basic education in the survey microdata (x-axis) and aggregate data from Barro and Lee (2013) and other sources (y-axis), after manual reclassification of educational categories in each country. Primary/basic education includes lower secondary education in some countries.

Figure A.47: Barro-Lee Versus Survey Data: Share of Working-Age Population With Secondary Education

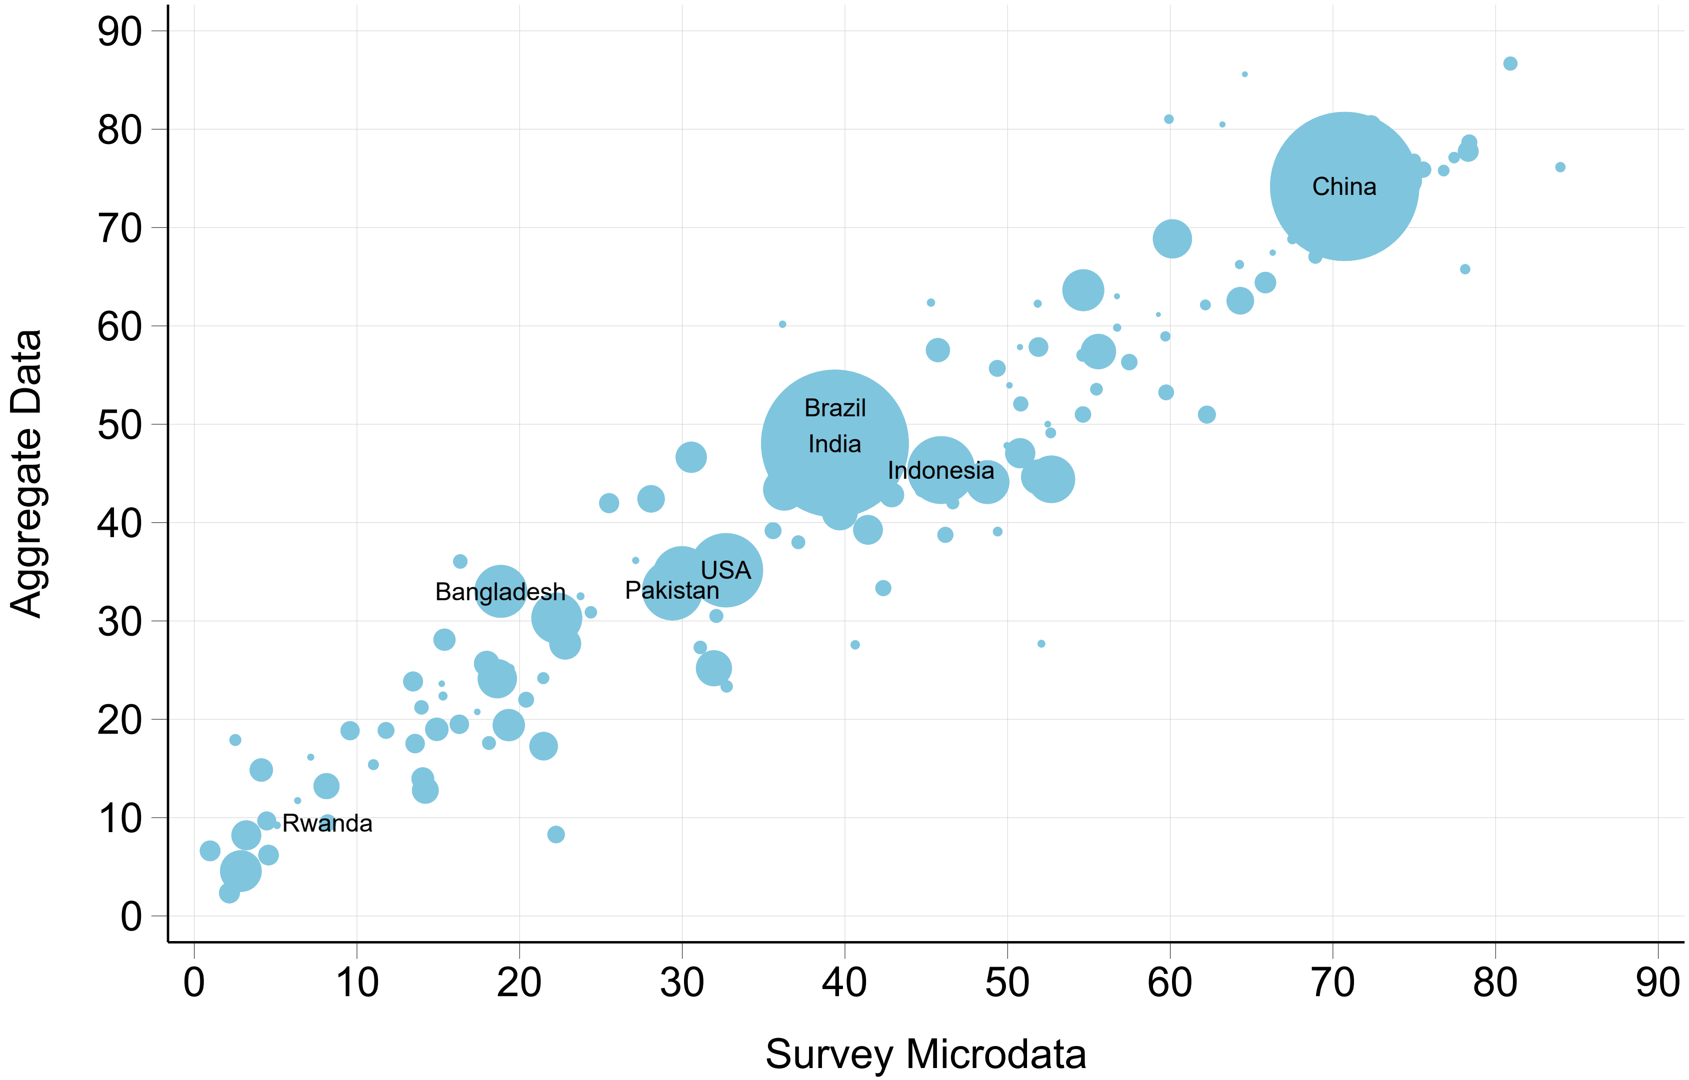

Notes. The figure compares estimates of the share of the working-age population with secondary education in the survey microdata (x-axis) and aggregate data from Barro and Lee (2013) and other sources (y-axis), after manual reclassification of educational categories in each country. Secondary education excludes lower secondary education in some countries.

Figure A.48: Barro-Lee Versus Survey Data: Share of Working-Age Population With Tertiary Education

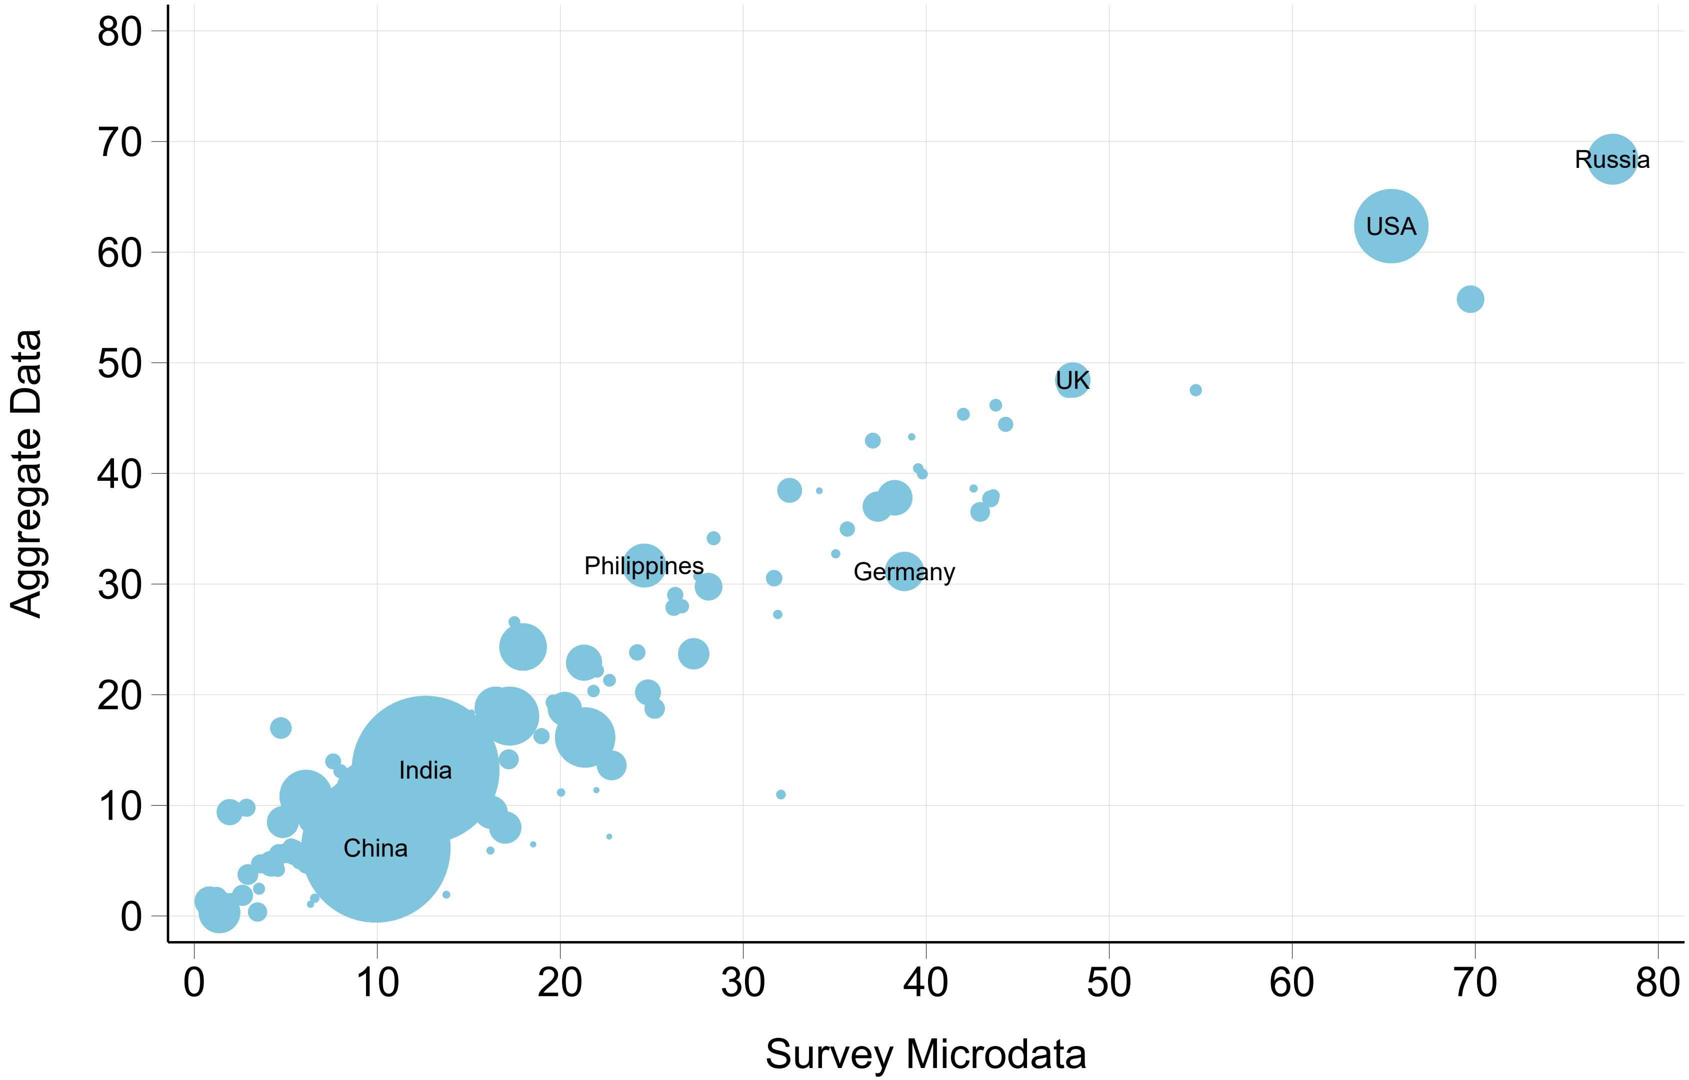

Notes. The figure compares estimates of the share of the working-age population with tertiary education in the survey microdata (x-axis) and aggregate data from Barro and Lee (2013) and other sources (y-axis), after manual reclassification of educational categories in each country.

A.7 Data Appendix: Returns to Schooling

A.7.1 OLS Estimates of Returns to Schooling

In the main analysis, I use estimates of returns to schooling by level estimated in each country. I rely on the following modified Mincerian equation:

$$\ln y_{ict} = \alpha_t + \beta_{ct}^{pri} D_{ict}^{pri} + \beta_{ct}^{sec} D_{ict}^{sec} + \beta_{ct}^{ter} D_{ict}^{ter} + X_{ict}\beta + \varepsilon_{ict} \tag{A.23}$$

With y_{ict} earned income of individual i in country c at time t, D_{ict}^{pri}, D_{ict}^{sec}, and D_{ict}^{ter} dummies for having reached primary, secondary, and tertiary education, and X_{ict} a vector of controls including gender, an experience quartic, and interactions between gender and the experience quartic. Earned income is the sum of all wage and self-employment income received by a given individual. I restrict the sample to all individuals aged above 15 with strictly positive income. I estimate this regression separately in each country and extract estimates of β_{ct}^{pri}, β_{ct}^{sec}, and β_{ct}^{ter}. In 47 countries with too few observations to estimate the return to primary education, I make the very conservative assumption that the return observed in 2019 is exactly zero. The same holds in 12 countries with too few observations to estimate the return to secondary education. Note that primary and secondary education do still end up having posiive effects on earnings in these countries in the benchmark specification, because imperfect substitution implies that the true return lies above the return observed in 2019 (see section 1.1.2).

Figure A.50 plots the distribution of annualized returns to schooling by level, while figures A.51 to A.54 map these returns in all countries with available estimates. Average returns to a year of schooling, estimated using a Mincerian equation with individual years of schooling on the right-hand side, typically range from 3% to 20%, with a median of 9%. Returns to primary education are typically lower than returns to secondary education, which are themselves below returns to tertiary education. The return to primary education is just 5% in the median country, compared to 9% for secondary education and 13% of tertiary education.

Table A.28 investigates the robustness of these results to using a standard Mincerian equation with only gender, potential experience, and potential experience squared as controls. Table A.29 compares baseline estimates pooling labor and self-employment income to a specification restricting the analysis to wage income. The results are almost identical: the average return to schooling is 8.9-9.7%, while the returns to primary, secondary, and tertiary education are 4.5-4.7%, 8.4-8.8%, and 14.4-15%, respectively.

Another concern is that workers and self-employed individuals declaring positive personal income might only represent a subset of the population. This is particularly concerning in low-income countries, where a large fraction of the population often relies on subsistence agriculture and thus ends up excluded from my estimation of the returns to schooling. I investigate this concern in appendix table A.30 by comparing three specifications. The first one corresponds to a standard Mincerian equation estimated at the individual level, restricting the sample to individuals declaring positive personal income. The second specification corresponds to a "household-level Mincerian equation," regressing per-capita expenditure on adults' average years of schooling. The third specification repeats the second specification, but after restricting the sample to households with at least one adult declaring positive personal income, which is useful to check whether the results are driven by selection into reporting positive income. I estimate these returns for eleven countries characterized by high poverty rates and large agricultural sectors: India, Pakistan, the Democratic Republic of the Congo, Burkina Faso, Côte d'Ivoire, Guinea-Bissau, Mali, Niger, Sénégal, and Togo. For each of these eleven countries, I was able to collect and manually harmonize survey microdata covering personal income, household expenditure, and educational attainment.

The three estimates end up falling very close to each other, amounting to a Mincerian return typically varying from 7% to 10%. Individual returns are slightly higher than household-level returns in some countries, such as India, Pakistan, and Côte d'Ivoire, which is to be expected given that variations in consumption are more driven by other factors, such as savings and transfers received by other households and the government. Yet there are also countries where individual returns are lower, such as Burkina Faso and Mali. Household-level returns before and after excluding households with no reported income are virtually identical in most countries. Together, these findings provide reassuring evidence that the returns estimated in this paper provide a good approximation of the true returns to schooling for the population as a whole.

A.7.2 IV Estimates of Returns to Schooling

In an alternative specification, I rely on instrumental variable estimates of the returns to schooling from a number of existing studies (see table A.31), which I use to adjust Mincerian OLS returns estimated with my data. Indeed, given that IV estimates from collected studies were generally computed at a different period and using a different sample than mine, they cannot directly be used in the estimation. Another difficulty is that these returns are annualized, while the returns used in my analysis correspond to total log-point increases in earnings from reaching specific levels of educational attainment. I thus incorporate IV estimates into the estimation in two steps. First, I use the ratio of IV to OLS estimates of yearly returns to schooling from these studies to adjust yearly returns (see figure A.55). Second, I exponentiate this ratio by level-specific average years of schooling to adjust total returns by level.

Formally, the total return of moving from level s_1 to level s_2 is:

$$r(s_1, s_2) = \ln(w_2) - \ln(w_1) \tag{A.24}$$

Which implies that the ratio of w_2 to w_1 is $w_2/w_1 = \exp(r(s_1, s_2))$. The corresponding annualized return to schooling is thus:

$$\beta(s_1, s_2) = \exp\left(r(s_1, s_2)\right)^{1/\mathrm{T}(s_1,s_2)} - 1 \tag{A.25}$$

With $\mathrm{T}(s_1, s_2)$ the difference in average years of schooling between individuals with educational attainment s_2 and s_1. We know from studies relying on quasi-experimental designs that IV estimates are higher than OLS estimates by a factor γ: $\beta^{IV}(s_1, s_2) = \gamma\beta(s_1, s_2)$. Hence, the adjusted total return to schooling is:

$$\begin{aligned} r^{IV}(s_1, s_2) &= \mathrm{T}(s_1, s_2) \times \ln\left(1 + \gamma\beta(s_1, s_2)\right) && \text{(A.26)} \\ &= \mathrm{T}(s_1, s_2) \times \ln\left[1 + \gamma\left(\exp\left(r(s_1, s_2)\right)^{1/\mathrm{T}(s_1,s_2)} - 1\right)\right] && \text{(A.27)} \end{aligned}$$

Figure A.49: Returns to Schooling: Pooled Estimates by Level

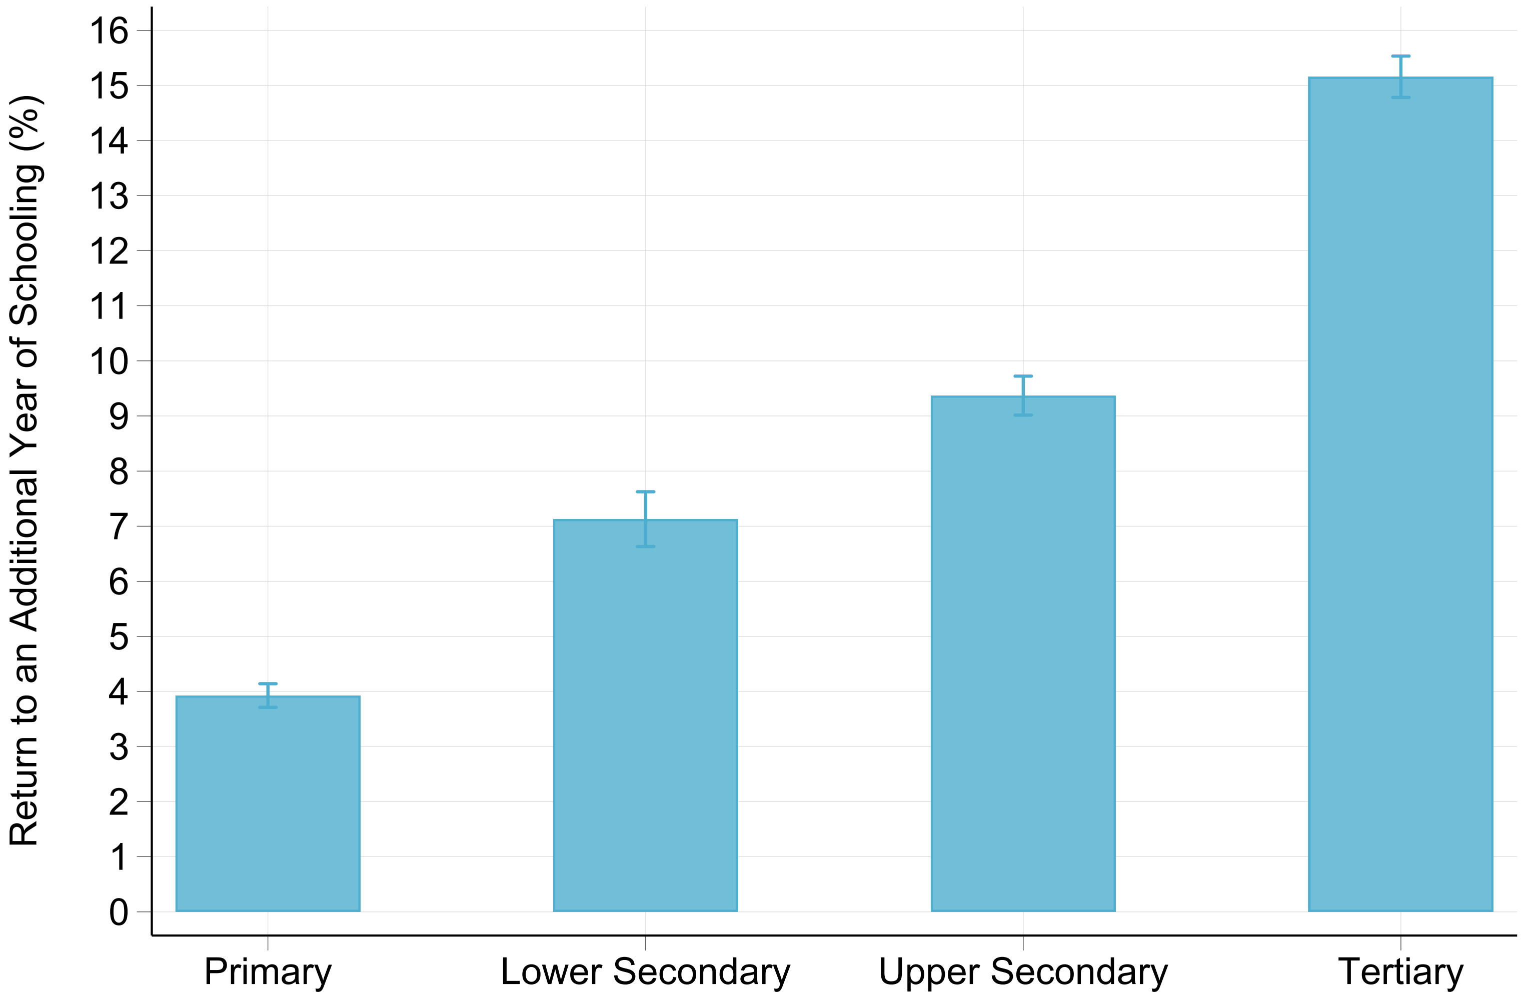

Notes. The figure reports estimates of an additional year of schooling by education level, based on a pooled regression on the full micro dataset. Primary: returns to a year of primary education. Lower secondary: return to a year of lower secondary education, restricting the sample to individuals with either primary or lower secondary education. Upper secondary: return to a year of upper secondary education, restricting the sample to individuals with either lower secondary or upper secondary education. Tertiary: return to a year of higher education, restricting the sample to individuals with either upper secondary or tertiary education. All models include controls for gender, an experience quartic, interactions between gender and the experience quartic, and country fixed effects. Observations are weighted to match each country's total population. Capped spikes represent 95% confidence intervals.

Figure A.50: Returns to Schooling: Distribution of Estimates by Level

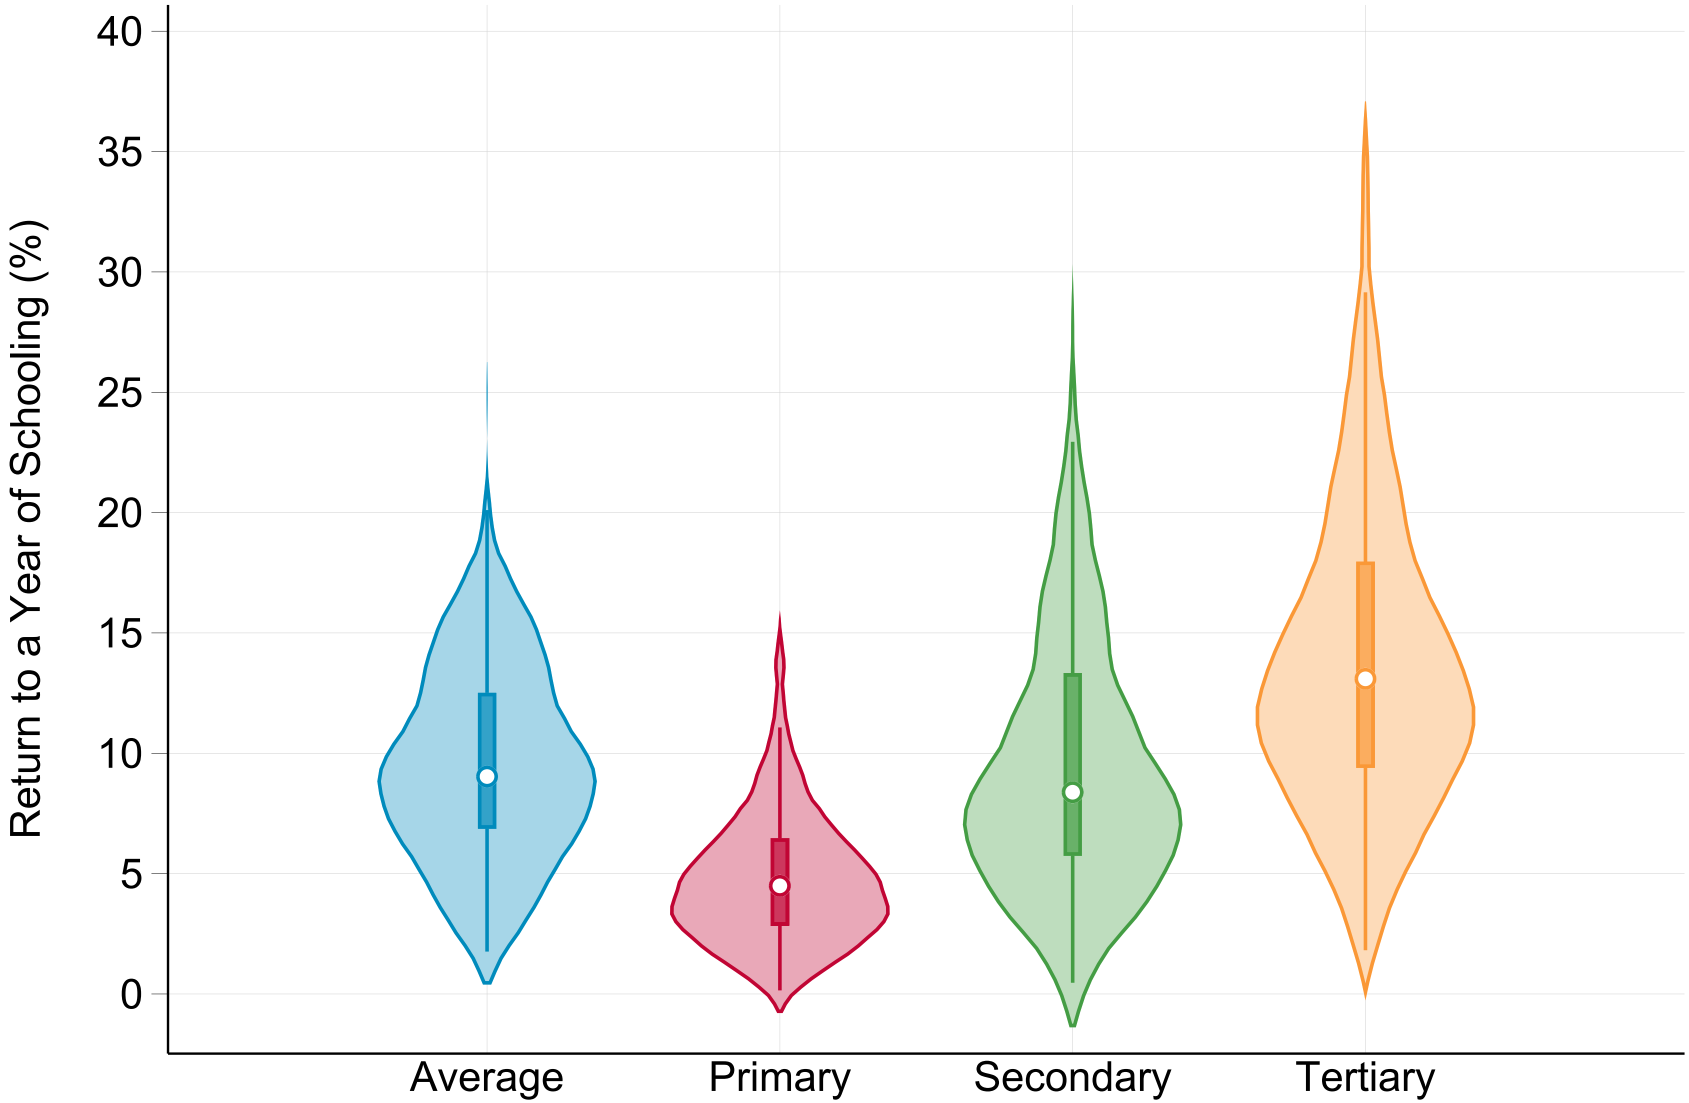

Notes. Author's computations using labor force survey microdata. The figure plots the cross-country distribution of returns to schooling by level. Estimates correspond to the effect of one additional year of schooling on the log of personal income, estimated using modified Mincerian equations controlling for an experience quartic, gender, and interactions between the experience quartic and gender. Primary: return to a year of schooling among individuals with either no schooling, some primary education, or completed primary education. Secondary: return to a year of schooling among individuals with either some primary education, completed primary education, some lower or upper secondary education, or completed upper secondary education. Tertiary: return to a year of schooling among individuals with some upper secondary education, completed upper secondary education, some tertiary education, or completed tertiary education.

Figure A.51: Return to an Additional Year of Schooling

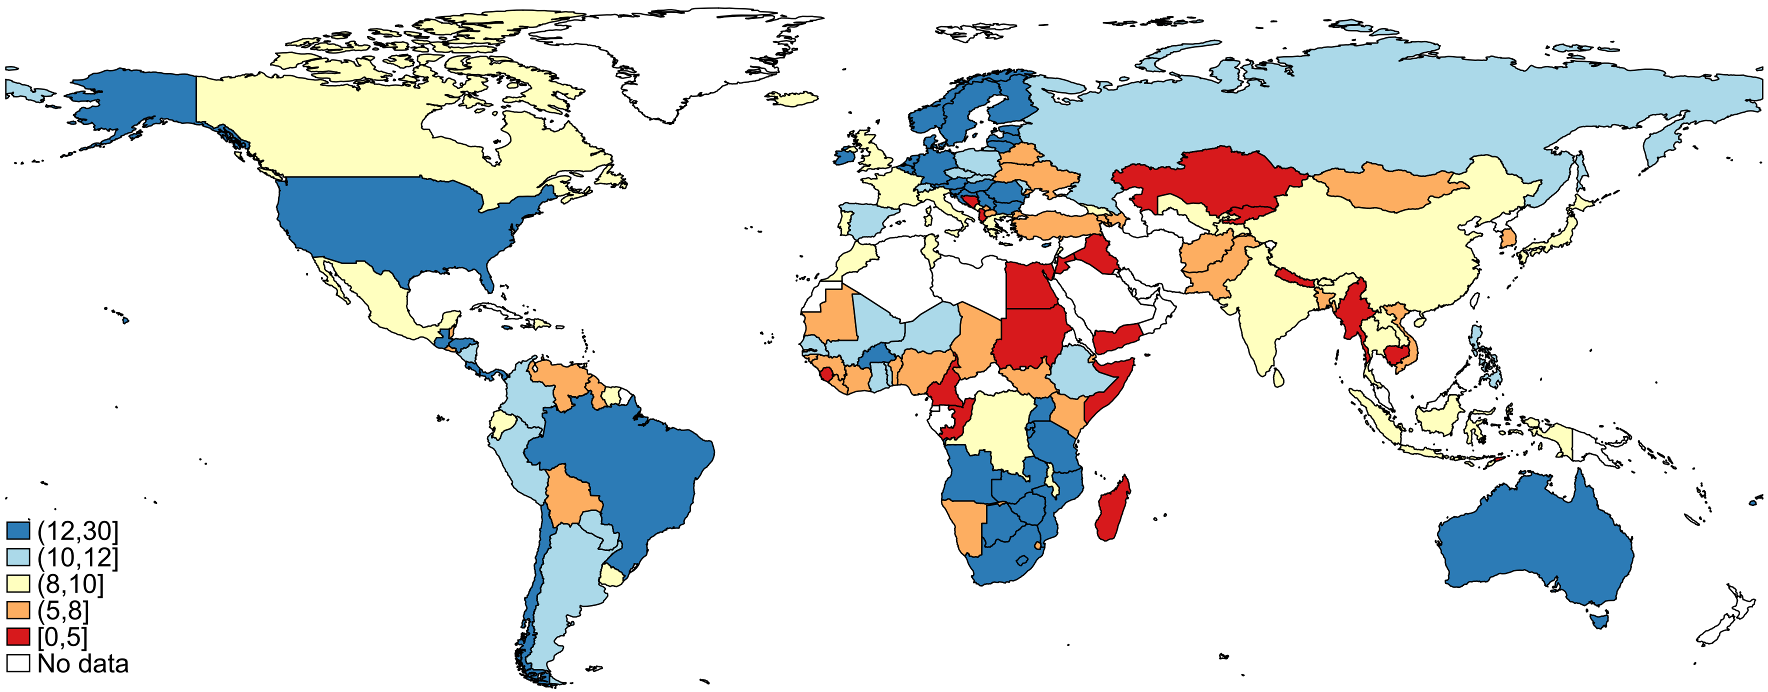

Figure A.52: Returns to an Additional Year of Primary Education

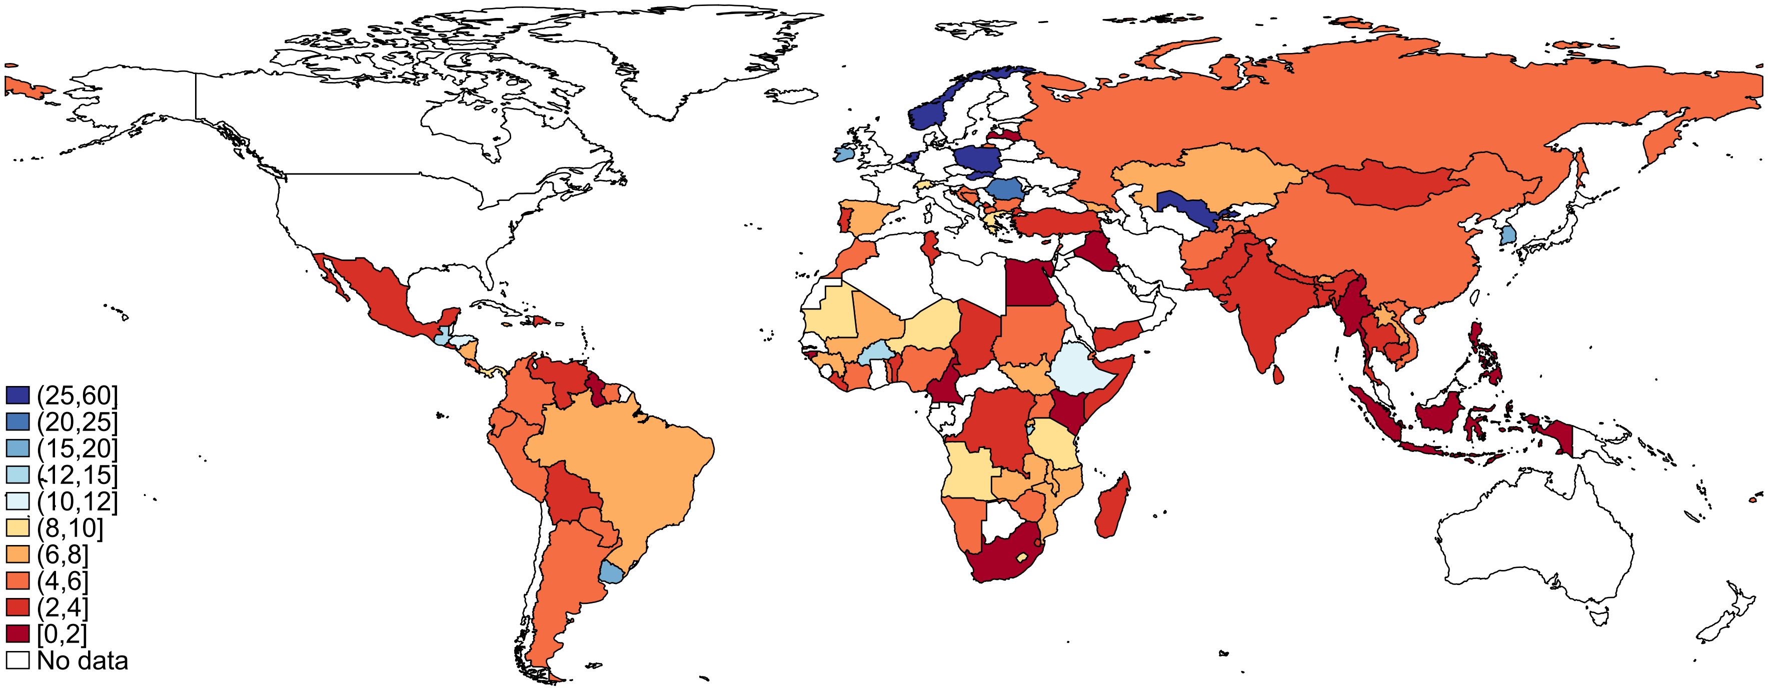

Figure A.53: Returns to an Additional Year of Secondary Education

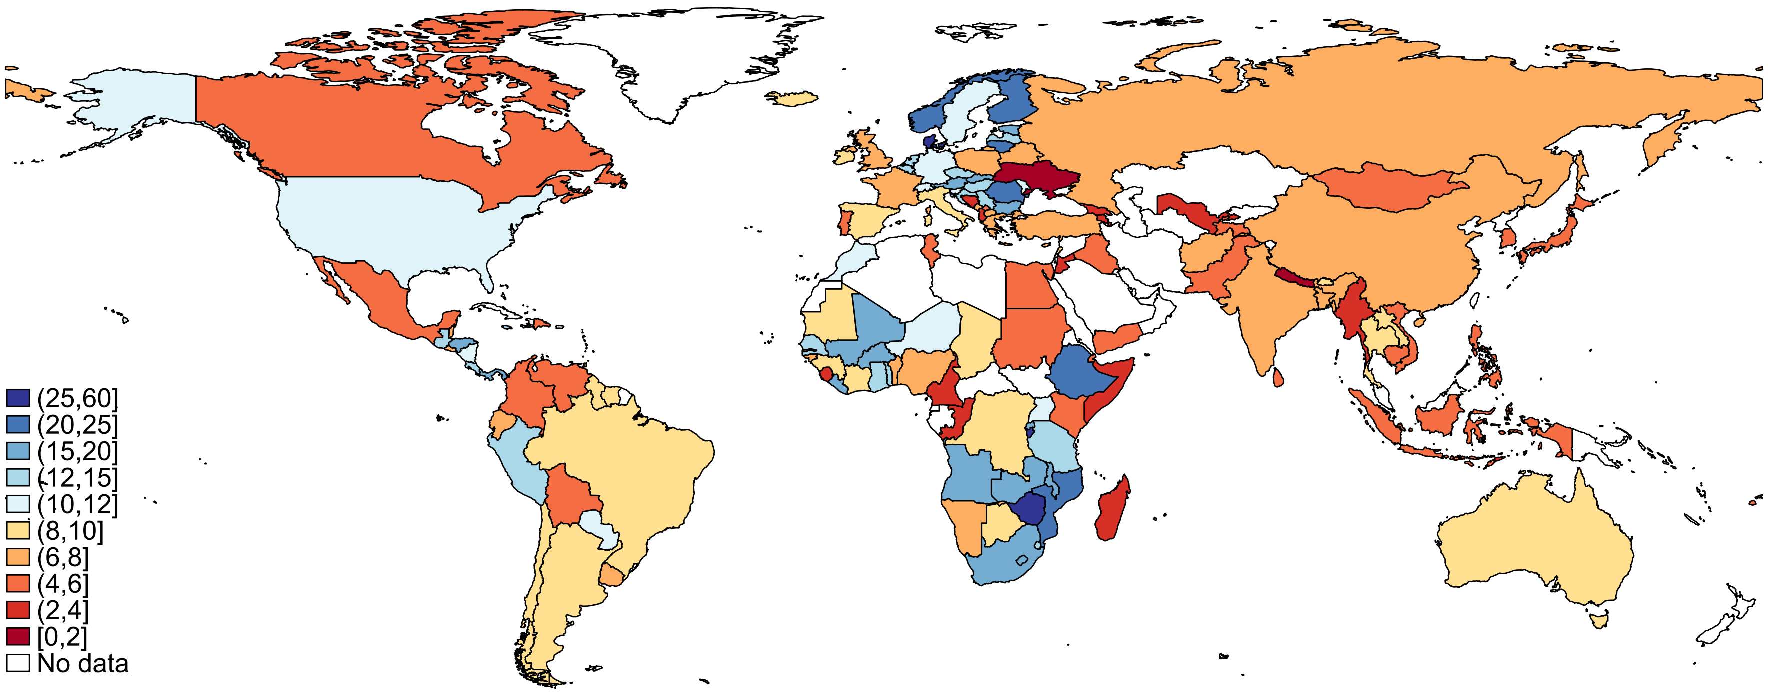

Figure A.54: Returns to an Additional Year of Tertiary Education

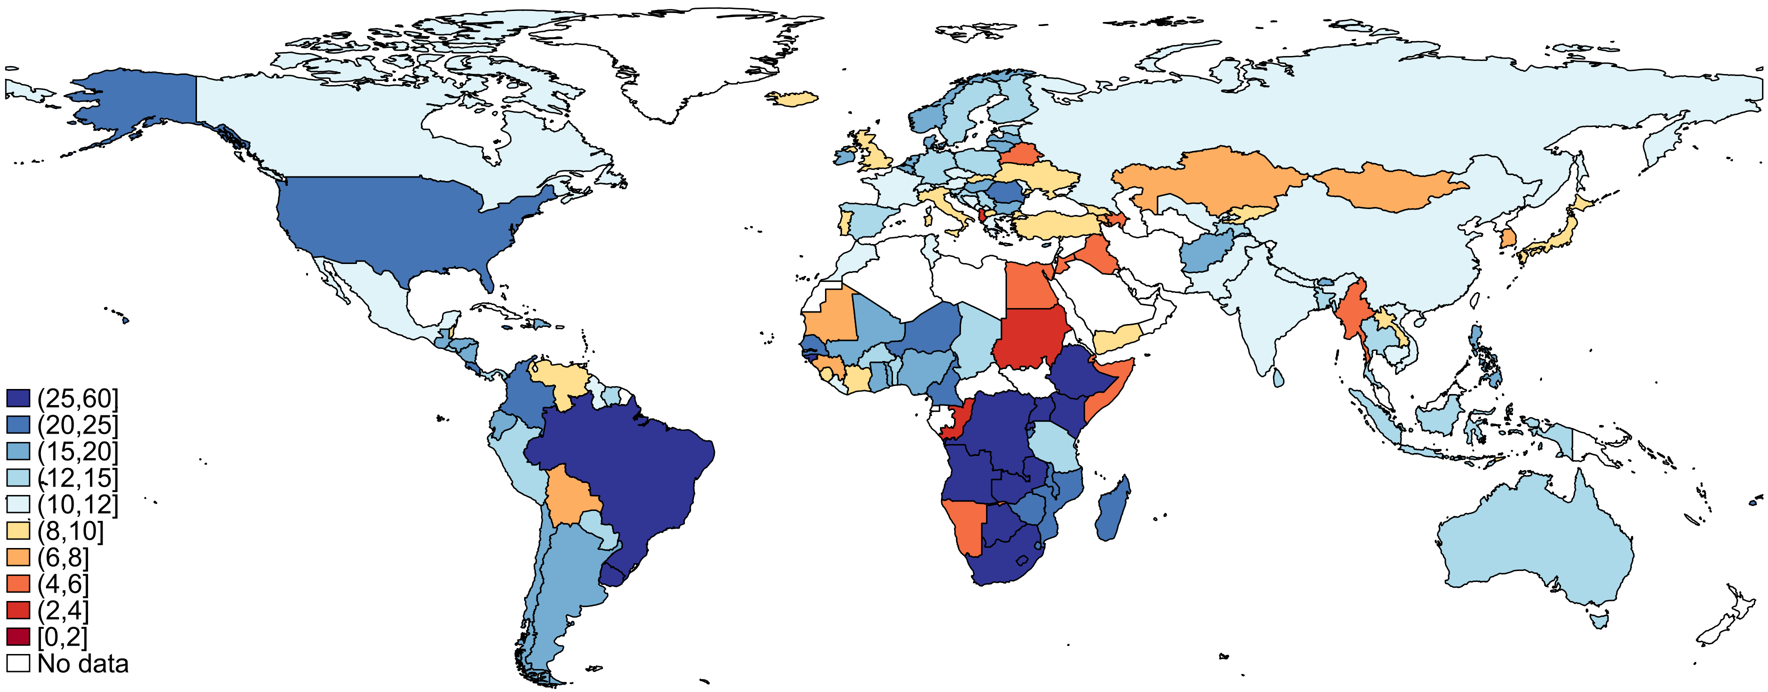

Figure A.55: Returns to Schooling: Ratio of IV to OLS Estimates

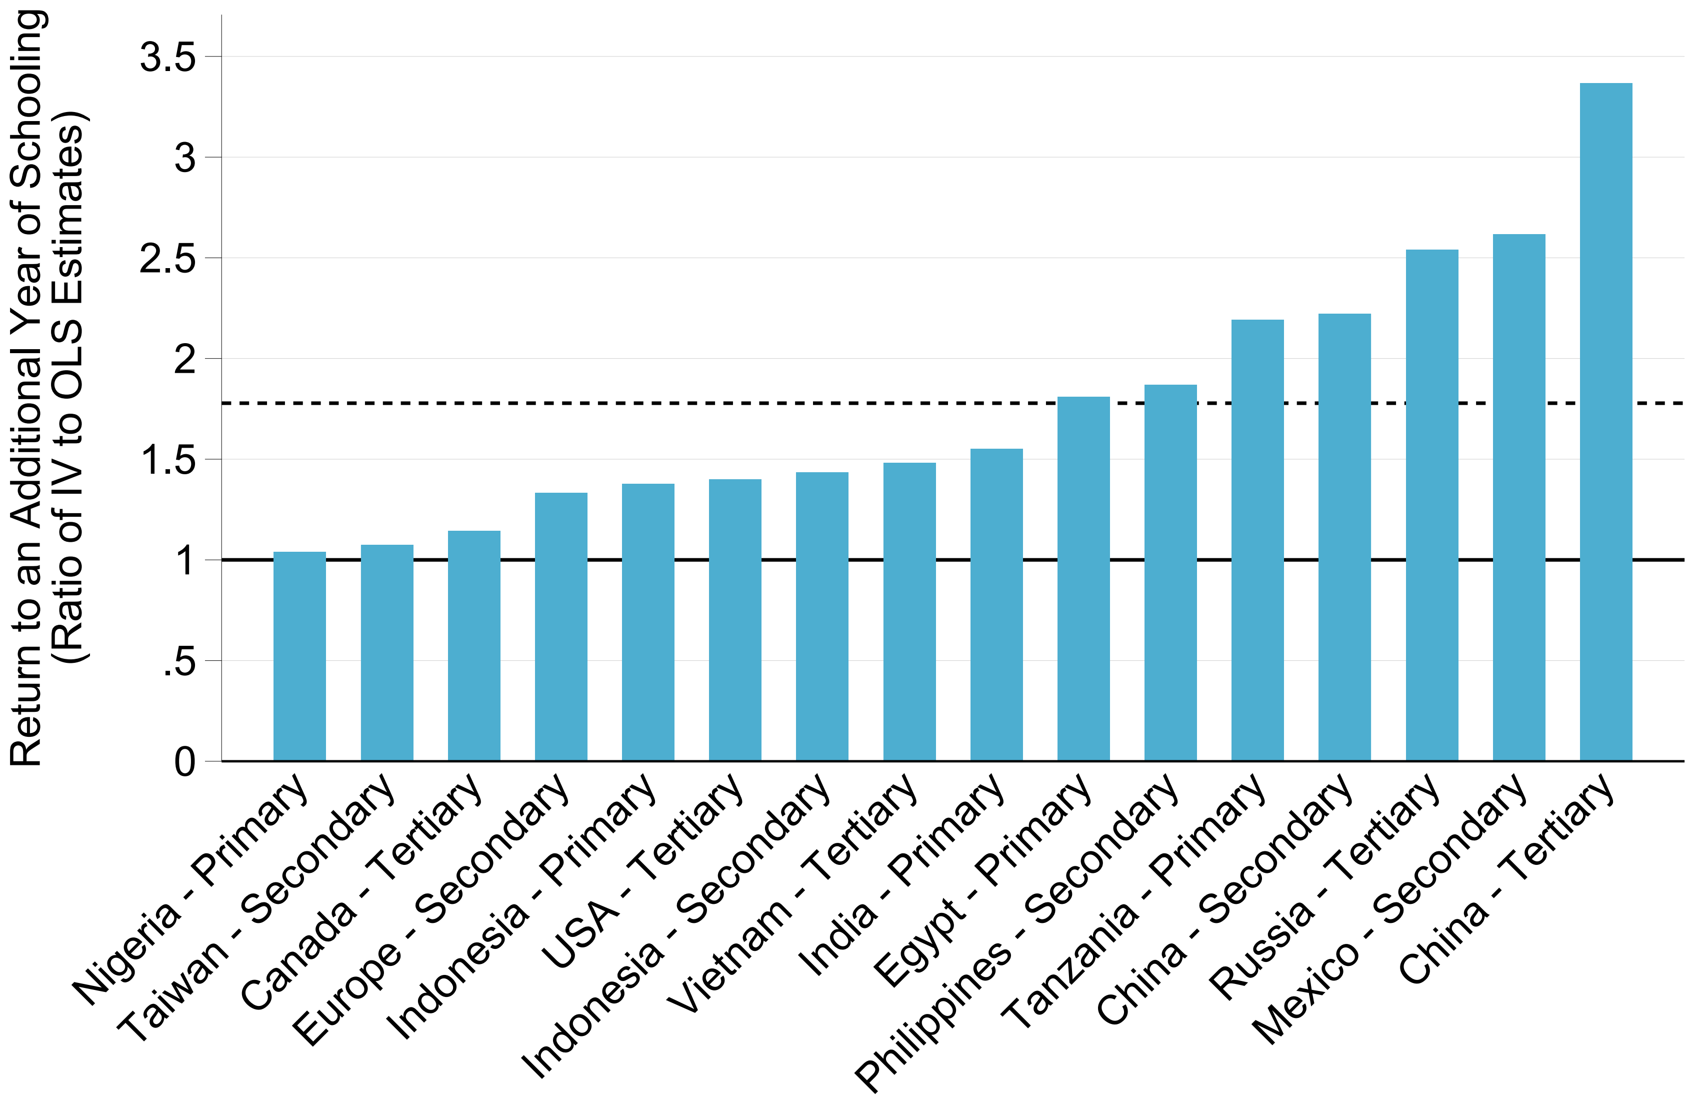

Notes. Author's elaboration compiling estimates from a number of published studies. Pri/Sec/Ter: returns to a year of primary/secondary/tertiary education. Dashed line: average ratio across all estimates.

Table A.28: Returns to Schooling: Standard versus Extended Mincer Equation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Years of Schooling	0.089***	0.090***	0.045***	0.046***	0.084***	0.085***	0.150***	0.150***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Level	All	All	Primary	Primary	Secondary	Secondary	Tertiary	Tertiary
Extended Model	No	Yes	No	Yes	No	Yes	No	Yes
N	4,912,763	4,912,763	1,493,033	1,493,033	3,184,353	3,184,353	2,772,992	2,772,992
Adj. R-squared	0.80	0.80	0.83	0.83	0.82	0.82	0.79	0.79

Notes. The table reports estimates of Mincerian returns, comparing "standard" and "extended" versions of the model by education level. Standard version: controls for gender, potential experience, and potential experience squared. Extended version: controls for gender, an experience quartic, and interactions between gender and the experience quartic, as in Lemieux (2006). Pooled regression across the full micro dataset. All estimates include country fixed effects. Observations are weighted to match each country's total population. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.29: Returns to Schooling: Total Personal Income Versus Wages

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Years of Schooling	0.090***	0.097***	0.046***	0.047***	0.085***	0.088***	0.150***	0.144***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Level	All	All	Primary	Primary	Secondary	Secondary	Tertiary	Tertiary
Income Concept	Total	Wages	Total	Wages	Total	Wages	Total	Wages
N	4,912,763	3,677,689	1,493,033	926,343	3,184,353	2,319,863	2,772,992	2,265,470
Adj. R-squared	0.80	0.88	0.83	0.90	0.82	0.87	0.79	0.88

Notes. The table reports estimates of Mincerian returns, comparing models including total personal income (wages + self-employment income) to models restricting the sample to wage earners. All models include controls for gender, an experience quartic, interactions between gender and the experience quartic, and country fixed effects. Observations are weighted to match each country's total population. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.30: Returns to Schooling: Personal Income Versus Per-Capita Consumption

	Individual Income	Consumption All Households	Consumption Households With Income Only
India	0.070***	0.063***	0.063***
	(0.001)	(0.001)	(0.001)
Pakistan	0.074***	0.071***	0.071***
	(0.001)	(0.001)	(0.001)
DR Congo	0.075***	0.071***	0.070***
	(0.002)	(0.002)	(0.002)
Burkina Faso	0.089***	0.106***	0.086***
	(0.006)	(0.003)	(0.005)
Benin	0.078***	0.068***	0.071***
	(0.003)	(0.002)	(0.003)
Côte d'Ivoire	0.073***	0.058***	0.052***
	(0.003)	(0.002)	(0.002)
Guinea-Bissau	0.041***	0.061***	0.052***
	(0.004)	(0.002)	(0.003)
Mali	0.072***	0.080***	0.061***
	(0.005)	(0.002)	(0.003)
Niger	0.108***	0.103***	0.101***
	(0.004)	(0.003)	(0.003)
Sénégal	0.078***	0.074***	0.075***
	(0.003)	(0.002)	(0.002)
Togo	0.100***	0.078***	0.076***
	(0.006)	(0.002)	(0.005)

Notes. The table compares returns to schooling estimated with three specifications. The first specification regresses individual income on individual years of schooling, controlling for age, gender, and their interaction. The second specification regresses per-capita consumption on average years of schooling of working-age adults at the household level, controlling for household size, average age, and the share of women. The third specification does the same, but after restricting the sample to households with at least one adult declaring positive personal income. India: 2019 PLFS survey. Pakistan: 2018 HIES survey. DR Congo: 2012 ECM survey. Other countries: 2018 EHCVM surveys. Robust standard errors in parenthesis. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.31: IV Estimates of Returns to Schooling

Source	Country	Level	OLS β	IV β	OLS SE	IV SE
Lemieux and Card (2001)	Canada	Tertiary	7	8	.2	4.4
Fang et al. (2012)	China	Secondary	9	20	.4	.6
Huang and Zhu (2022)	China	Tertiary	4.9	16.5		
Assaad et al. (2023)	Egypt	Primary	2.1	3.8	.3	4.5
Brunello, Weber, and Weiss (2015)	Europe	Secondary	4.2	5.6	.3	2.6
Khanna (2023)	India	Primary	10	15.5		
Carneiro, Lokshin, and Umapathi (2017)	Indonesia	Secondary	9	12.9	.5	4.8
Duflo (2001)	Indonesia	Primary	7.7	10.6	.06	2.2
Navarro-Sola (2021)	Mexico	Secondary	4.7	12.3	.1	1.6
Oyelere (2010)	Nigeria	Primary	2.6	2.7	.1	1.3
Sakellariou (2006)	Philippines	Secondary	6.1	11.4		
Kyui (2016)	Russian Federation	Tertiary	6.1	15.5	.25	1.1
Spohr (2003)	Taiwan	Secondary	5.4	5.8		
Delesalle (2021)	Tanzania	Primary	2.6	5.7	.01	2.1
Zimmerman (2014)	USA	Tertiary	10	14		
Vu and Vu-Thanh (2022)	Viet Nam	Tertiary	16	23.7		

Notes. The table reports instrumental variable estimates of returns to schooling from selected articles. OLS: return to schooling estimated by OLS. β: return to a year of schooling. SE: standard error associated with the estimate.

Appendix B

Appendix to "Revisiting Global Poverty Reduction: Public Goods and the World Distribution of Income, 1980-2019"

B.1 Additional Methodological Details

This section presents the methodology used to estimate the distribution of global pretax and posttax incomes. Section B.1.1 outlines the data sources used. Section B.1.2 explains the methodology used to construct aggregate government revenue and expenditure series. Section B.1.3 covers the distribution of transfers.

B.1.1 Data Sources

B.1.1.1 Macroeconomic Aggregates

My main source for macroeconomic aggregates is the World Inequality Database (WID, see `http://wid.world`), which combines various data sources to provide harmonized national accounts series and population totals in all countries in the world from 1950 to 2021 (Blanchet and Chancel, 2016). I use five main variables from the WID database in my analysis: gross domestic products, net national incomes, total populations, national income deflators, and PPP conversion factors to 2021 US dollars.

B.1.1.2 Government Revenue Aggregates

For government revenue aggregates, I rely on Bachas et al. (2022), who build a new database on the level and composition of tax revenue in 150 countries since 1965. Their database provides information on total tax revenue as a share of net domestic product, together with a breakdown by type of tax (personal income taxes, corporate income taxes, social contributions, property and wealth taxes, indirect taxes, and other taxes).

B.1.1.3 Government Expenditure Aggregates

Estimating the evolution of consolidated government expenditure and its composition is challenging, and there exists no single data source providing harmonized information on spending on different policies across countries. Accordingly, I combine various data sources to build a new database on government expenditure by function.

My primary data source for total expenditure is Mauro et al. (2015), who draw on historical data from the IMF and other sources to construct a new database on total consolidated government expenditure as a share of GDP in 170 countries from 1800 to 2011 (59 countries are covered in 1980, 91 in 1990, and 157 after 2000).[1] The main advantage of this database is its historical coverage and conceptual consistency: total expenditure covers consolidated government, incorporating both central and local government expenditure. Its main limitation is that it does not provide any information on the composition of expenditure.

The main data source used to cover the composition of expenditure (as well as total expenditure after 2011) is the IMF, which provides data on spending by Classification of the Functions of Government (COFOG) in 172 countries. Depending on the country and year, the series cover either the general government, or only unconsolidated central, state, and local government expenditure.

I use other data sources on specific types of expenditure to complement and further decompose IMF data.

The IFPRI-SPEED database (Yu, Magalhaes, and Benin, 2015) covers total central government expenditure in 147 countries, incorporating some country-specific sources absent from IMF series.

The World Bank's World Development Indicators (WDI) database provides series on total education and health expenditure as a share of government spending in 208

[1] See https://www.imf.org/external/datamapper/exp@FPP/USA.

countries.

For decomposing social protection expenditure into social insurance and social assistance, I rely on three sources: the OECD's Social Expenditure (SOCX) database, the United Nations Economic Commission for Latin America and the Caribbean's Social Expenditure database, and the World Bank's Atlas of Social Protection Indicators of Resilience and Equity (ASPIRE).[2] All three datasets provide data on total social protection expenditure as a share of GDP, as well as its decomposition by type of program.

B.1.1.4 Pretax Income Distribution Data

Data on the distribution of pretax income by country since 1980 come from the World Inequality Database, which brings together country-specific studies (e.g., Piketty, Saez, and Zucman (2018) for the US and Blanchet, Chancel, and Gethin (2022) for Europe) and other data sources to provide estimates of average pretax income by generalized percentile in all countries around the world since 1980 (see Chancel and Piketty, 2021).

The income concept covered in pretax national income, that is, the sum of all personal income flows before taking into account the operation of the tax-and-transfer system, but after taking into account the operation of pension and unemployment systems. By construction, average pretax income matches average net national income in each country.

B.1.1.5 Tax Incidence Data

For the distributional incidence of taxes, I rely on estimates from a companion paper (Fisher-Post and Gethin, 2023).

B.1.1.6 Transfer Incidence Data

For the distributional incidence of government expenditure, I rely on five data sources: Piketty, Saez, and Zucman (2018), Blanchet, Chancel, and Gethin (2022), Gethin, Kofi Tetteh Baah, and Lakner (forthcoming), the CEQ database, and the World Bank's ASPIRE database.

Piketty, Saez, and Zucman (2018) provide in their microfile data on all cash and health transfers received by US individuals from 1962 to 2021. I use this information

[2] See `https://www.oecd.org/social/expenditure.htm`; `https://statistics.cepal.org/portal/databank/index.html?lang=en&indicator_id=4407&area_id=`; `https://www.worldbank.org/en/data/datatopics/aspire`.

to compute the share of total cash and health transfers received by pretax income decile.

Blanchet, Chancel, and Gethin (2022) provide in their microfile data on family and social assistance transfers received by individuals in 32 European countries. I use it to compute the share of cash transfers received by pretax income decile in each country.

Gethin, Kofi Tetteh Baah, and Lakner (forthcoming) provides unique information on school attendance by age and household income in 155 countries since 1980. I use it to compute the share of education expenditure received by income decile in each country.

The CEQ database provides estimates of the share of cash transfers, total education expenditure, and total health expenditure received by pretax income decile in 45 countries.

Finally, the World Bank's ASPIRE database draws on harmonized survey microdata to compute the share of social assistance transfers received by pretax income quintile in 108 countries over the 1998-2019 period (most countries are covered since the mid-2000s).

B.1.2 Harmonization of Government Expenditure by Function: G

I combine all available data sources to build a harmonized database on the level and composition of government expenditure since 1980. I proceed in two steps. First, I combine existing sources to estimate total consolidated government expenditure in all countries and years. Second, I estimate the composition of consolidated expenditure by function.

B.1.2.1 Total Government Expenditure

My primary data source to measure total consolidated government expenditure is Mauro et al. (2015), which I use for all country-years in which data is available. In countries not covered at all by Mauro et al. (2015), I use available IMF general government series. In countries not covered at all by any of these two sources, I use the sum of central, state, and local government expenditure reported in IMF series.

To cover all countries from 1980 to 2019, I then combine all data sources to carry these combined series backward and forward. First, I carry Mauro et al. (2015) series backward and forward using growth rates in IMF general government series as a share of GDP. When data is still missing, I use growth rates in IMF central, state, and local government. When data is still missing, I use growth rates in total tax revenue as a share of GDP from Bachas et al. (2022). When data is still missing, I use growth rates in central government expenditure as a share of GDP from the IFPRI-SPEED database. When data is still missing, I extrapolate total expenditure backwards and forwards as a constant share of GDP. Finally, in the 13 small countries with no data on total government expenditure at all, I take continental averages of total expenditure as a share of GDP.

B.1.2.2 Composition of Government Expenditure

As for total government expenditure, I combine available data sources to estimate the composition of expenditure by function. My primary data source is the IMF series, which decompose expenditure into 10 large COFOG categories: social protection, education, health, recreation and culture, housing and community amenities, environmental protection, economic affairs, public order and safety, defense, and general public services.

I give priority to general government expenditure series, and use the sum of central, state, and local government expenditure series only when general government data is not available at all in a given country. I then extrapolate the composition of expenditure backward and forward so as to cover the entire 1980-2019 period. For countries with no data on the composition of expenditure, I take continental averages.

World Bank education and health expenditure series tend to be more consistent and cover more countries and years, so I incorporate them directly into these estimates. To do so, I simply replace education and health expenditure as a share of the general government budget by World Bank series when available. I then proportionally adjust other components of the general government budget so that the share of expenditure going to each function sums up to 1. This ensures that the resulting education and health expenditure series are fully consistent with World Bank data, while preserving the relative shares of other functions of government reported in IMF data.

Following the same principle, I then further decompose general public services and economic affairs into their subcomponents. As above, I use IMF series to split general public services into administration and debt service expenditure, extrapolating their respective ratios when data is missing. In countries with no data on these subcomponents, I assume that debt service absorbs one-third of general public services expenditure, which corresponds to the average observed across all country-years. I follow the same process to decompose economic affairs into transport expenditure and expenditure on other economic affairs.

Lastly, given that pretax income already includes pensions and unemployment benefits, I remove spending on social insurance transfers from social protection expenditure. To do so, I use the OECD's and the CEPAL's datasets to estimate a split between social insurance and social assistance transfers, and reduce social protection expenditure by the corresponding amount in the harmonized database. For countries not covered by these two datasets (all non-OECD, non-Latin American countries), I use the World Bank's ASPIRE database, which provides an estimate of total social assistance expenditure as a share of GDP in 124 countries. I take the ratio of this estimate to total social protection expenditure in my harmonized series, so as to reduce social protection expenditure to only cover social assistance. Finally, in countries with no data from either the OECD, the CEPAL or the World Bank, I make the conservative assumption that social protection expenditure matches social assistance expenditure (in other words, that the share of social insurance expenditure in social protection expenditure is zero).

B.1.3 Distribution of Transfers: $\gamma(m_i)$

I combine available data sources to estimate transfer incidence profiles by income group. My measure of interest consists in concentration curves, that is, the share of a specific type of transfer received by income decile.[3] I then distribute transfers by combining these profiles with government expenditure by function in each country.

In each case, I consider three scenarios for countries with missing data: one benchmark scenario corresponding to the average profile observed across all country-years; an upper bound in which missing countries are attributed the average transfer incidence profile of the five countries with the most progressive profiles; and a lower bound in which missing countries are attributed the average transfer incidence profile of the five countries with the most regressive profiles. In the absence of consistent data on the evolution of transfer progressivity over time (with the exception of the United States), I assume that is has remained constant in each country.

Social Assistance I combine concentration curves of social assistance expenditure by pretax income decile or quintile from Piketty, Saez, and Zucman (2018), Blanchet,

[3]Concentration curves are more meaningful to distribute transfers than incidence curves, given that unlike taxes paid, transfers received are not generally proportional to income or consumption.

Chancel, and Gethin (2022), the World Bank’s ASPIRE database, and the CEQ Institute, by order of priority. I then allocate total social assistance expenditure in each country-year based on these profiles.

Education For education, I derive concentration curves of education spending by combining data from Gethin, Kofi Tetteh Baah, and Lakner (forthcoming) with series of public education expenditure per child by level from the UNESCO, as explained in section 2.2.3.

Health The CEQ database (and Piketty, Saez, and Zucman, 2018 for health in the US) is, to the best of my knowledge, the only available data source providing consistent information on the distributional incidence of health expenditure. I allocate total health expenditure in each country-year based on the corresponding concentration curves by income decile.

Economic Affairs I assume that expenditure on economic affairs is received proportionally to consumption. I use incidence curves on the relationship between income and consumption from Chancel et al. (2022a), who combine a number of microdata sources to derive typical lower and upper bounds on savings rates by pretax income percentile. In my benchmark scenario, I apply the same consumption-income profile in each country, corresponding to the typical profile estimated in Chancel et al. (2022a). I then use their lower and upper bounds as lower and upper bounds on the progressivity of expenditure on economic affairs.

Other Government Expenditure Other components of the government budget include expenditure on economic affairs, public order and safety, housing and community amenities, administration, recreation and culture, defense, and environmental protection. I distribute them proportionally to posttax disposable income.

Distribution of Taxes I borrow estimates of the distribution of taxes by generalized percentile directly from Fisher-Post and Gethin (2023).

Debt Service, Budget Balance, and Local Taxes Finally, to reach a concept of posttax income consistent with the distributional national accounts framework (Piketty, Saez, and Zucman, 2018), I distribute debt service expenditure, the budget balance, and local taxes to individuals. This ensure that average income is consistent with the net national income. The main issue is that data on tax revenue from Bachas et al. (2022) only covers taxes collected by the central government. As a result, the gap between total consolidated government expenditure and central government revenue incorporates both local taxes and the government deficit, which available data do not allow to distinguish. In the absence of better information, I distribute the gap between total revenue and total expenditure proportionally to pretax income in each country.

B.2 Accounting for Public Sector Productivity

B.2.1 Conceptual Framework

I consider an extension in which the value of public goods is allowed to differ from cost of provision. The value of public goods received by individuals can theoretically be broken down into three components:

$$g(m_i) = \sum_j G^j \times \gamma^j(m_i) \times \theta^j(m_i) \tag{B.1}$$

With G^j government expenditure and $\gamma^j(m_i)$ the share of expenditure received by i. $\theta^j(m_i)$ captures the fact that for a given cost of provision, individuals may receive services of different quality. Empirically, it is useful to make a distinction between two notions of productivity:

$$\theta^j(m_i) = \Theta^j \times q^j(m_i) \tag{B.2}$$

Θ^j is the *aggregate productivity* of expenditure on function j, which does not depend on m_i. It captures the fact that the government may be more or less efficient at providing a given service than a benchmark production unit. For instance, public schools in country A may be on average less cost-efficient than public schools in country B, which implies that all public education transfers should be reduced by a constant factor in country A.

$q^j(m_i)$ is a *heterogeneous productivity* parameter. It captures the fact that the quality of services provided, holding cost constant, may differ between income groups. For instance, teachers teaching in poorer areas may be more or less qualified than those teaching in richer areas, independently from the wages they receive.

Consider for example a government providing free public education at a cost of $G^j = \$1000 \times N$, with N the size of the population. Because of inequalities in access to public education, however, the poorest 20% only receive \$500 per capita of funding: $\gamma^j(m_i) = 0.1$. Furthermore, the government appears to be particularly inefficient at providing public education: it under-performs by 50% relative to what it could do if it was at the production possibility frontier, which implies that $\Theta^j = 0.5$. Finally, schools attended by children belonging to the bottom quintile appear to be 20% less efficient at providing education than the average school in the country: $q^j(m_i) = 0.8$. Combining the different parameters, we get: $g^j(Q_1) = \$500 \times 0.5 \times 0.8 = \$200$.

B.2.2 Aggregate Productivity Θ^j

I start with the estimation of aggregate productivity Θ^j, corresponding to the overall efficiency of the government at providing public services.

B.2.2.1 Methodology

Following the existing literature measuring the productivity of governments by combining data on outcomes with data on government expenditure (e.g., Adam, Delis, and Kammas, 2011; Afonso, Schuknecht, and Tanzi, 2005; Herrera and Ouedraogo, 2018), I propose to estimate Θ^j by benchmarking the productivity of governments around the world to one another. If a government produces more output than any other for a given cost, then its efficiency is set to 1, and the productivity of other governments with comparable costs is estimated based on the outputs they deliver. The advantage of this approach is its simplicity and transparency: governments delivering better education and health outcomes are considered to be more productive.

I estimate simple models of public sector productivity based on international data covering government expenditure and outcomes. In broad strokes, I choose a function of government (e.g., health) and collect cross-country data on expenditure (public health spending per capita), other inputs (e.g., GDP per capita), and an outcome of interest (mortality). I then use data envelopment analysis to non-parametrically estimate the technical frontier, defined as the maximum output ever achieved in any country-year for a given level of expenditure and other inputs (e.g., Herrera and Ouedraogo, 2018). Finally, I use the estimated frontier to estimate Θ^j, based on the extent to which output could be improved without changing costs in a given country-year.[4] This yields measures of technical efficiency ranging from 0 to 1 for each country-year covered by the data.

I apply this methodology to estimate the productivity of public education and public healthcare. For each of these two functions of government, I estimate two alternative

[4] I use the `teradial` command in Stata.

production frontiers: one based on a single input and a single output, and one that incorporates additional inputs to account for the fact that, for instance, education outcomes might be higher because of higher GDP per capita rather than greater education spending. For other public goods, given the absence of high-quality data on service delivery, I take the average of public education and public healthcare measures. Finally, I interpolate between years and extrapolate backwards and forwards measures of productivity by function in each country, so as to cover the 1980-2019 period. For countries with no data at all, I take the global average observed in each year.

B.2.2.2 Productivity of Public Education

Inputs The first element required to estimate public education productivity is a measure of cost of provision. I take public education expenditure per child, expressed in 2021 PPP US dollars, estimated from the public spending database compiled in this paper. For the estimation of multiple-input efficiency, I add three auxiliary inputs to the model: the log of GDP per capita in 2021 PPP USD (available from the WID), the log of the adult literacy rate, and the log of the share of children enrolled in private schools (both available from the World Bank's WDI)

Output The second element needed is a measure of government performance. Following the large literature in macroeconomics investigating the role of education in explaining differences in economic development (e.g., Hanushek, Ruhose, and Woessman, 2017), I propose to measure the output of the education system as the expected human capital that a child can hope to obtain at age 5:

$$Y^{\text{education}} = \exp(r_S S + r_Q Q) \tag{B.3}$$

With S expected years of schooling at age 5, r_S the return to a year of schooling, Q a measure of education quality, and r_Q the return to education quality. Data on expected years of schooling come from the UNESCO and covers 202 countries over the 1970-2020 period. Education quality is taken from Altinok, Angrist, and Patrinos (2018), who compile data from various international test scores to construct a new database of education quality in 134 countries. The return to schooling is set to 10% per year and the return to quality to 15% per standard deviation, following the existing literature.

Results Figure B.32 plots the resulting relationship between performance and cost of provision for all country-years. There is a very strong correlation between the two variables ($\rho = 0.9$, $R^2 = 0.82$): countries spending more on education display education systems of substantially better quality. Yet, there is also significant dispersion in the expected human capital stock achieved for a given level of government expenditure. The upper dashed line represents the efficient frontier, estimated using data envelopment analysis with variable returns to scale. This corresponds to a piecewise linear estimate of the maximum achievable output by level of expenditure.

The trajectories of Niger, Indonesia, and South Korea are represented as examples. Education expenditure and schooling outcomes have significantly increased during this period in all three countries. Niger stands quite far below the frontier, while South Korea has remained one of the most cost-efficient countries in the database throughout the period. Indonesia falls somewhat in-between. The corresponding measures of education productivity are about $\Theta^{\text{education}} = 0.5$ for Niger, 0.85 for Indonesia, and 0.9 for South Korea in the last year available.[5]

B.2.2.3 Productivity of Public Healthcare

Inputs As for public education, the first step is to collect data on cost of provision. Given the particular role that private healthcare can play in some countries, I focus on total healthcare expenditure per capita (private and public combined). For the estimation of multiple-input efficiency, I add two auxiliary inputs to the model: the log of GDP per capita in 2021 PPP USD and the share of private health expenditure in total current health expenditure (both available from the WDI). All data series come from the World Bank's World Development Indicators.

Output Finding an accurate measure of the quality of healthcare provision is more challenging than for education. Indeed, unlike the human capital stock, which has a clear cardinal (monetary) interpretation, there is no obvious measure of healthcare performance whose units are directly comparable to cost of provision. Quality-adjusted life expectancy is often taken as a measure of interest (e.g., Cutler et al., 2022), yet this indicator is, by itself, arguably a poor measure of the performance of the healthcare system. Given these limitations, I turn instead to the healthcare access and quality (HAQ) index estimated in the context of the global burden of disease study (GBD, 2022). This indicator ranks healthcare systems from 0 to 100, based on death rates from 32 causes of death that could be avoided by timely

[5] Notice that as shown in figure B.32, I fit the efficiency frontier using the log of the human capital stock. To get correct efficiency measures, one then needs to convert the ratio of logs into the ratio of actual human capital stocks. More precisely, we have a measure of efficiency θ^{log} such that: $\theta^{log} = \frac{\log x}{\log \bar{f}(x)}$, with $\bar{x}$ the technical frontier evaluated at x. The objective is to convert θ^{log} into $\theta = \frac{x}{f(x)}$. Rearranging yields $f(x) = \exp(\frac{\log x}{\theta^{log}})$ and hence $\theta = \frac{x}{\exp(\frac{\log x}{\theta^{log}})}$.

and effective medical care. The main advantage of the HAQ index is that it was specifically created by health experts to measure the ability of healthcare systems to cure preventable diseases: it is explicitly a measure of performance. It also has the advantage of covering nearly all countries in the world since 1990. The disadvantage is that it is normalized from 0 to 100, so it has no cardinal interpretation. In the absence of better solution, I re-express the HAQ index in units of life expectancy by first regressing it on life expectancy at birth, and then normalizing it using the coefficient obtained. Reassuringly, this correction only marginally affects efficiency scores.[6]

Results Figure B.33 plots the resulting relationship between healthcare performance and cost of provision for all country-years. As for education, there is a very strong correlation between the two variables ($\rho = 0.93$, $R^2 = 0.87$): countries spending more on healthcare are much more able to limit deaths from curable diseases. The upper dashed line represents the efficient frontier, while the trajectories of Sweden, China, and India are represented for the sake of illustration. India is significantly below the frontier (with an implied Θ^{health} below 0.6 in all years), while China and Sweden have remained among the best-performing countries throughout the period.

B.2.2.4 Discussion: Estimates of Θ^j as Lower Bounds on Government Productivity

I view these estimates as providing a *lower bound* on government productivity, especially in poor countries, for three main reasons.

First, national income purchasing power parity conversion factors do already account for government productivity (World Bank, 2013). Indeed, public sector productivity is adjusted for all government services in the Asia-Pacific, Western Asia, and Africa regions, using a Cobb-Douglas function that assumes that government employees are less productive in poor countries because of a lower and less efficient stock of capital equipment (Heston, 2013). In OECD countries and the European Union, further adjustments are made for health and education, combining indicators on the quantity and quality of services provided (Blades, 2013). Hence, the correction made here to account for aggregate productivity implies adjusting transfers downwards twice, once when using PPP conversion factors to correct for price differences across countries, and once when multiplying transfers received by Θ^j.

[6] More specifically, I run a linear regression of life expectancy on the HAQ index, controlling for the log of GDP per capita, years of schooling of the working-age population, and country fixed effects. I then multiply the HAQ index by the coefficient obtained, so as to re-express it in "units of life expectancy."

Second, the frontier approach implies by construction that Θ^j cannot be greater than 1, given that the maximum input-output combination ever observed in any country-year is given a score of 1. As a result, governments are assumed to never be more productive than the private sector for any kind of service provided ($\Theta^j = 1$ corresponds to a government exactly as cost efficient as the private sector).

Third, omitted variable bias is likely to drive estimates of Θ^j in poor countries significantly *downwards*. Indeed, poor countries are likely to have worse outcomes for a given level of government expenditure not only because of inefficiencies, but also because of a number of other confounding factors. These include lower incomes, greater inequality, more extreme weather conditions, or lower basic knowledge, which directly affect education and health outcomes independently from government investment. For all these reasons, overall government expenditure is likely to be more efficient in these countries than what the model suggests.

B.2.2.5 Validation: Correlates of Government Efficiency

Finally, a useful way of checking the reliability of my measures of government productivity is to compare them to existing indicators. Appendix table B.7 shows that education and healthcare productivity are positively correlated with a number of indicators of government efficiency available from international sources and the literature. This is especially true of healthcare productivity, which is positively associated with a composite index of government effectiveness ($\rho = 0.57$ for single-input estimates), lower corruption ($\rho = 0.43$), and more transparent policy-making ($\rho = 0.34$). I also find a positive correlation between my measures of healthcare efficiency and the index of public sector productivity of Chong et al. (2014) ($\rho = 0.29$), who mail letters to 159 countries and argue that the rate of return of these letters to their original sender provides a simple and transparent measure of government productivity.

All four of my measures of productivity are also highly correlated with one another. In particular, the cross-country correlation between single-input and multiple-input estimates is 0.94 for education and 0.97 for healthcare.[7] In other words, accounting for other factors affecting the relationship between government expenditure and outcomes does not appear to significantly alter rankings of which countries are more or less efficient. I view these results as additional reassuring evidence that my estimates capture broad differences in government productivity across countries relatively well.

[7] See appendix table B.8, which provides raw pairwise correlations between measures.

B.2.3 Heterogeneous Productivity $q^j(m_i)$

Heterogeneity in productivity refers to the fact that the quality of public goods provided may vary by income group independently from their cost of provision, because, for instance, poorer geographical areas in a given country may provide public services in a more or less cost efficient way. Estimating heterogeneous productivity at a global scale is extraordinarily challenging, given the lack of high-quality data on service delivery by income group. In the absence of better information, I investigate using subjective perceptions of public services from international survey data to derive estimates of heterogeneous productivity by income group around the world. The data source is the Gallup World Poll, a yearly survey conducted since 2005 in 165 countries, which asks respondents whether they are satisfied with different types of public services in their area. I aggregate average responses by income quintile to measure differences in satisfaction with local public education, healthcare, police, and transport services.[8] I then use relative responses as a scaling parameter, to increase or decrease the transfer received by each income group, for each of these four functions of government.

These subjective indicators have advantages and disadvantages. On the one hand, they are available for nearly all countries in the world and cover different types of public services, providing a simple and transparent measure of differences in the perceived quality of public services. On the other hand, they may suffer from significant measurement biases, in particular the fact that subjective perceptions may not be comparable across income groups because of differences in expectations of what "good" and "bad" public services might be. This could lead to underestimating inequalities in the quality of services received by income group, if richer respondents evaluate the quality of public services by comparing them to a higher benchmark than low-income households.

At the same time, existing studies suggest that heterogeneity in quality by income group remains relatively limited. Drawing on various data sources in the context of South Africa, Gethin (2023c) finds that inequalities in the quality of public services received by income group tend to be small, both for subjective or objective indicators. Subjective perceptions of public services also appear to track objective indicators of inequality in service delivery relatively well. Similarly, Walter (2020) provides

[8] Respondents are asked whether they are "Satisfied" or "Dissatisfied" with the public transportation system, the quality of roads and highways, the educational system or the schools, and the availability of quality health care. I use these four measures to derive estimates of heterogeneous productivity in the provision of transport, education, and health care. For police services, I rely on a question that asks whether respondents have "confidence in the local police force."

evidence that pupil-teacher ratios tend to vary substantially within countries, in particular in developing countries, but that differences in local economic development or remoteness only explain a very small fraction of these variations.

B.3 Additional Figures and Tables

B.3.1 Additional Key Results

Figure B.1: Validation of Methodology
In-Kind Transfers Received by the Bottom 50% in South Africa, Simplified versus Detailed Series

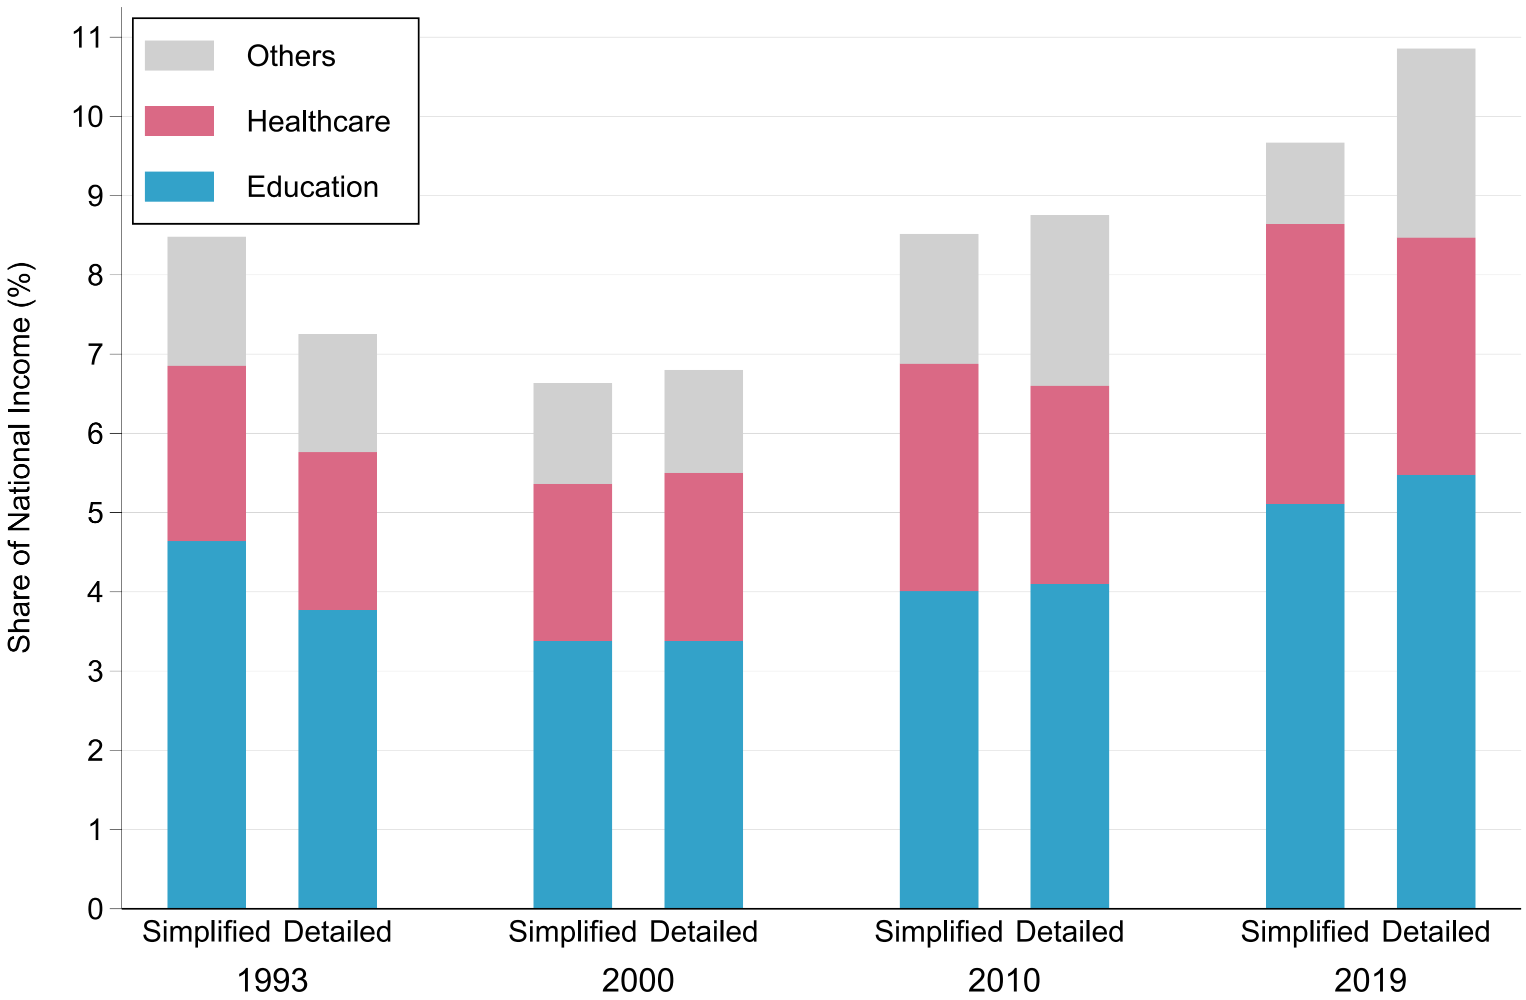

Notes. The figure plots the level and composition of public services received by the bottom 50% in South Africa, comparing simplified series (this paper) to detailed series constructed in Gethin (2023c).

Figure B.2: Validation of Methodology
Bottom 50% Share of Education Spending, Own Estimates Versus CEQ Database

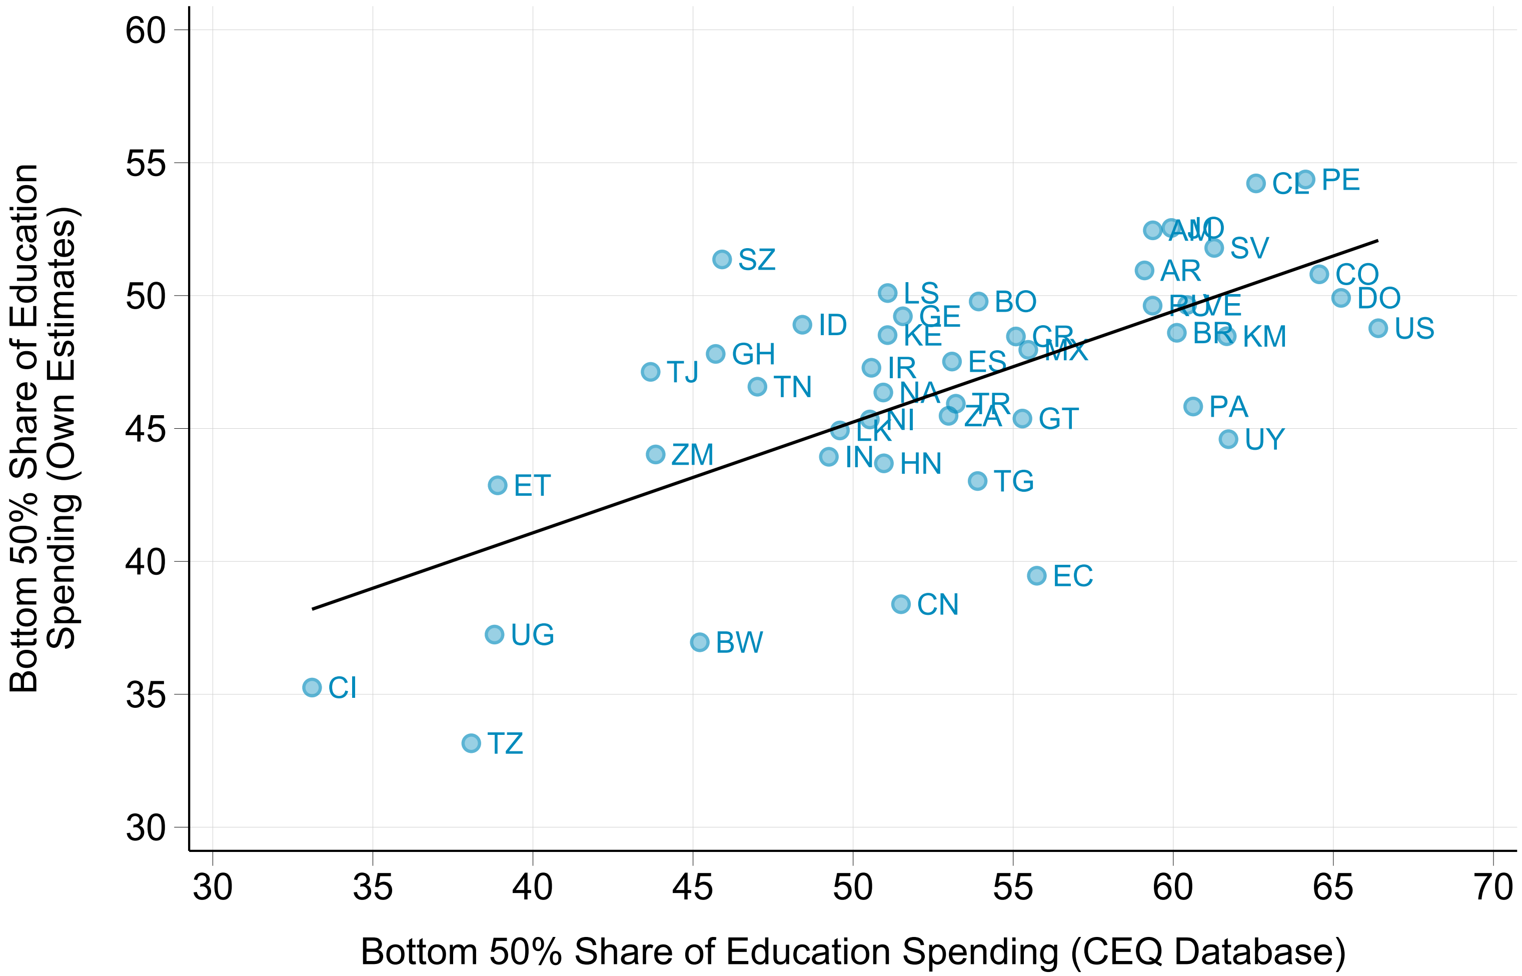

Notes. The figure compares the share of education spending received by the bottom 50% for selected countries in the Commitment to Equity (CEQ) Database and in the database of Gethin, Kofi Tetteh Baah, and Lakner (forthcoming).

Figure B.3: Global Poverty Headcount Ratio at $3.65 per day, 1980-2019

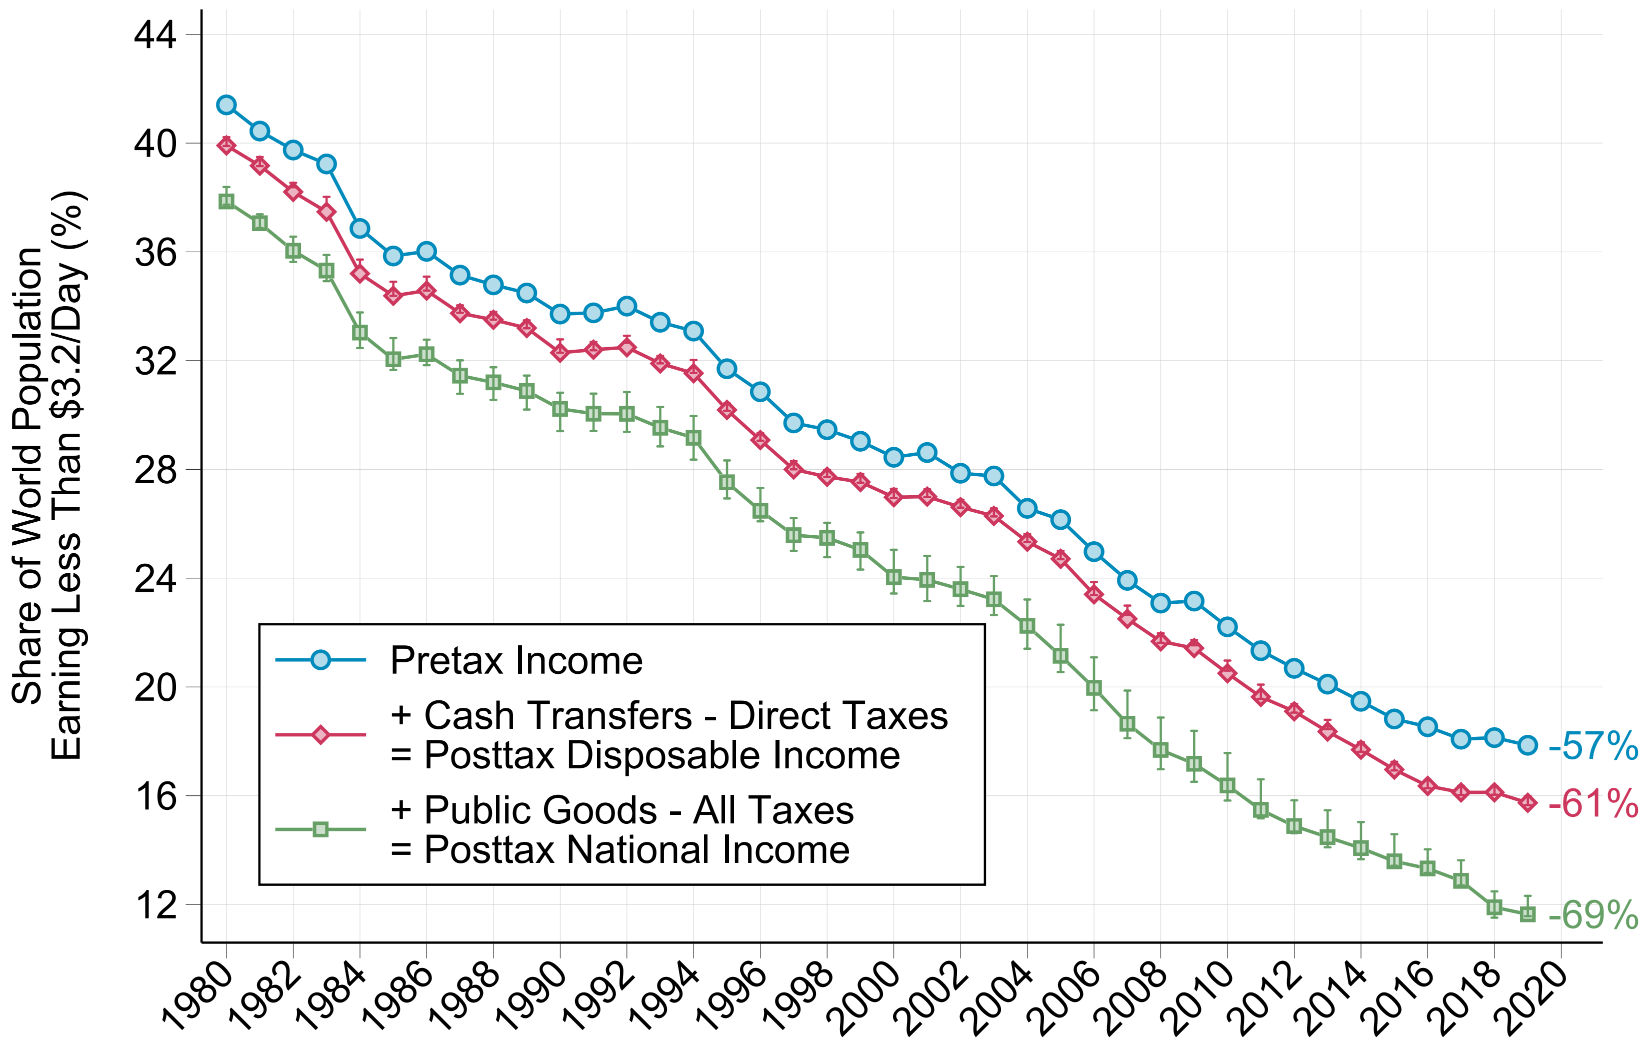

Notes. The figure plots the evolution of the poverty headcount ratio at $3.65 per day (2017 PPP USD) in the world as a whole, for different income concepts. Posttax disposable income removes direct taxes and adds cash transfers. Posttax national income removes all taxes and adds all cash and in-kind transfers. The unit of observation is the individual. Income is split equally between all household members.

Figure B.4: Global Poverty Headcount Ratio at $6.85 per day, 1980-2019

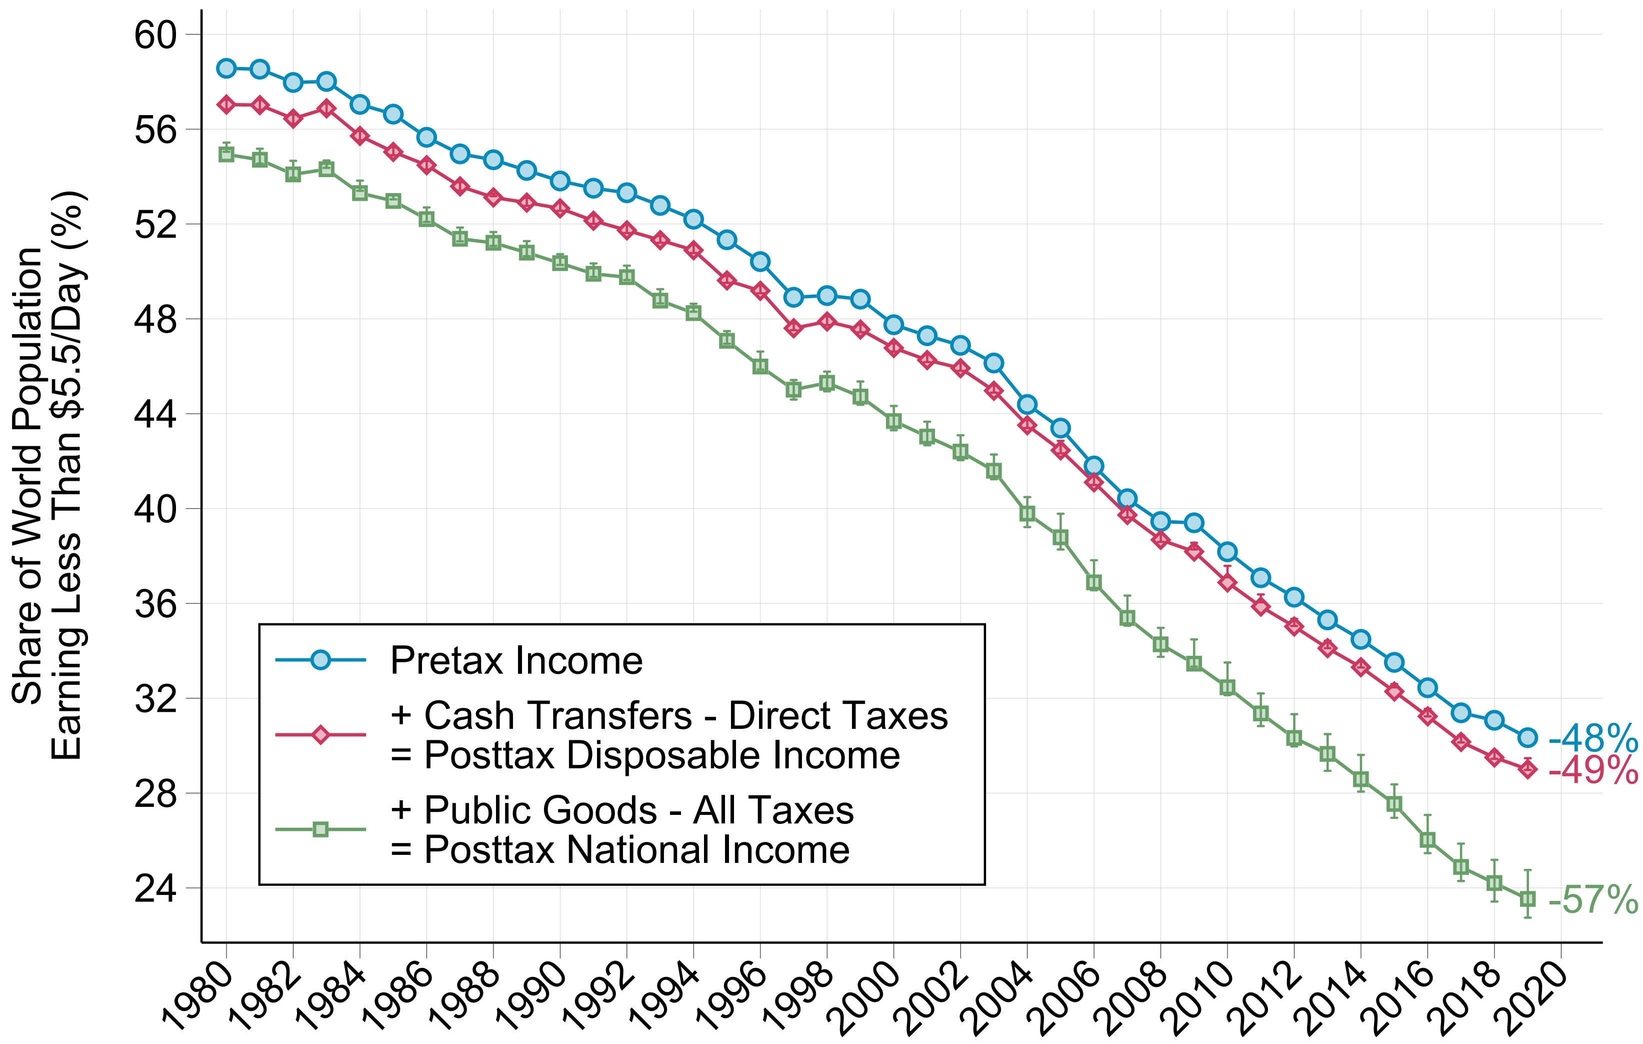

Notes. The figure plots the evolution of the poverty headcount ratio at $6.85 per day (2017 PPP USD) in the world as a whole, for different income concepts. Posttax disposable income removes direct taxes and adds cash transfers. Posttax national income removes all taxes and adds all cash and in-kind transfers. The unit of observation is the individual. Income is split equally between all household members.

Figure B.5: Global Poverty Headcount Ratio at $2.15 per day, 1980-2019 (World Bank Data)

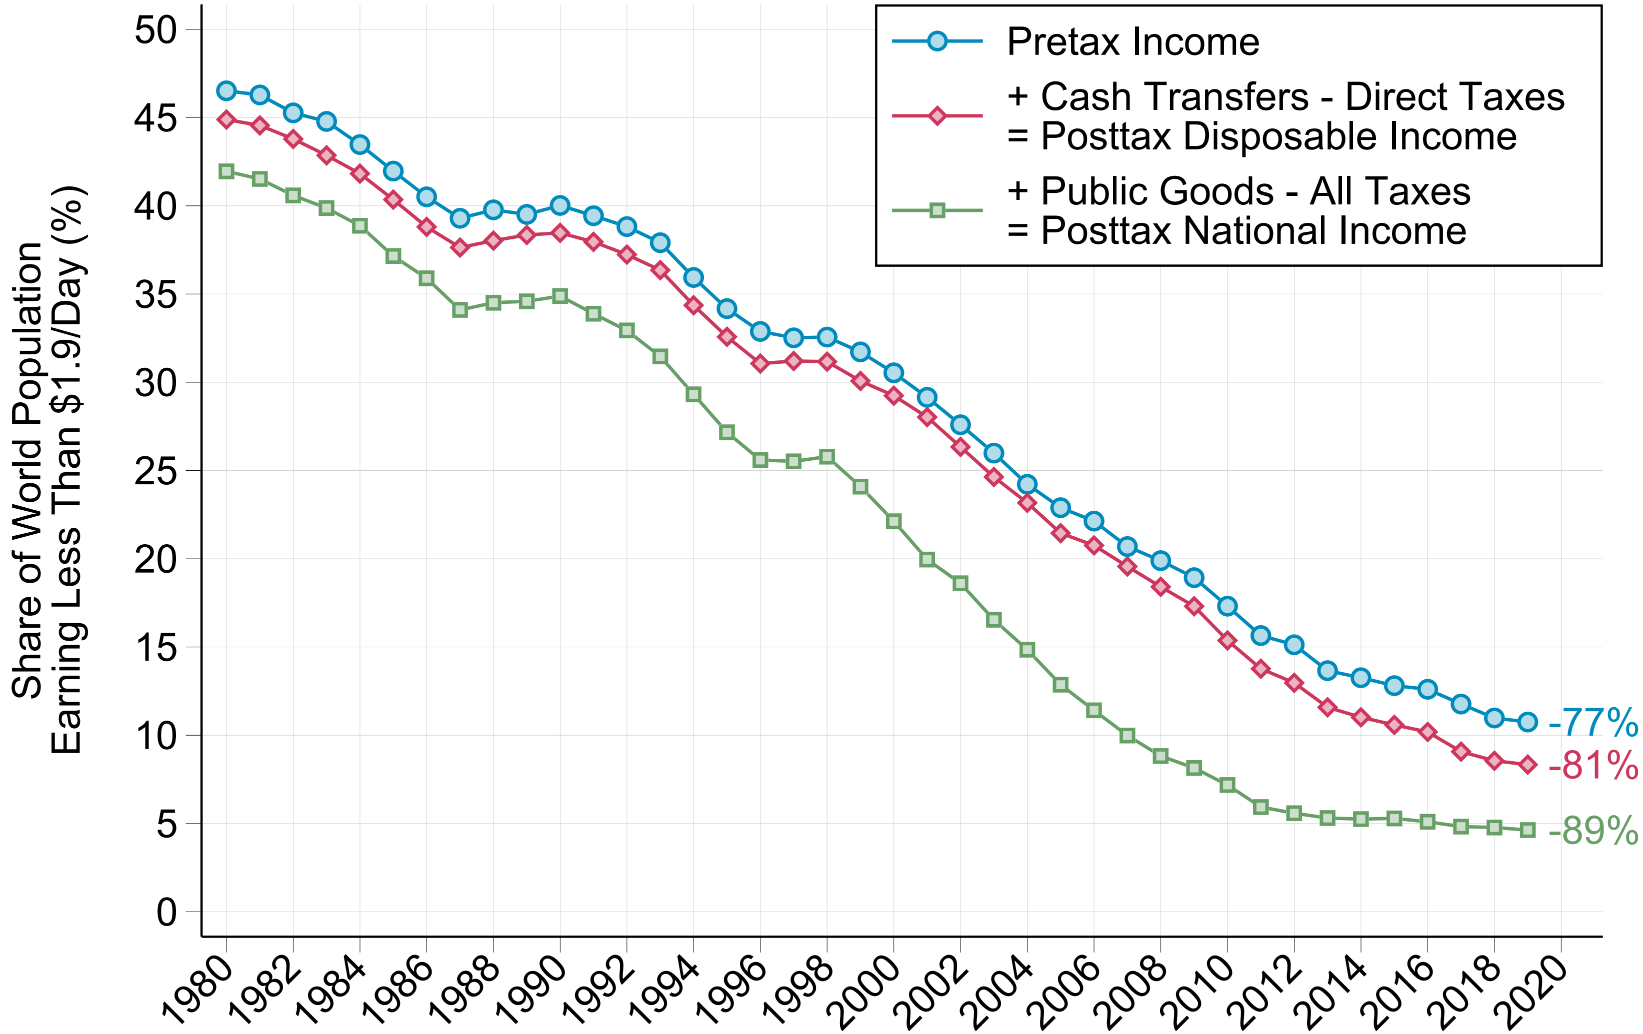

Notes. The figure plots the evolution of the poverty headcount ratio at $2.15 per day (2017 PPP USD) in the world as a whole, for different income concepts. Distributions of consumption or disposable income per capita from the World Bank. Pretax income is reconstructed as consumption minus social assistance plus direct taxes.

Figure B.6: Global Poverty Headcount Ratio at \$3.65 per day, 1980-2019 (World Bank Data)

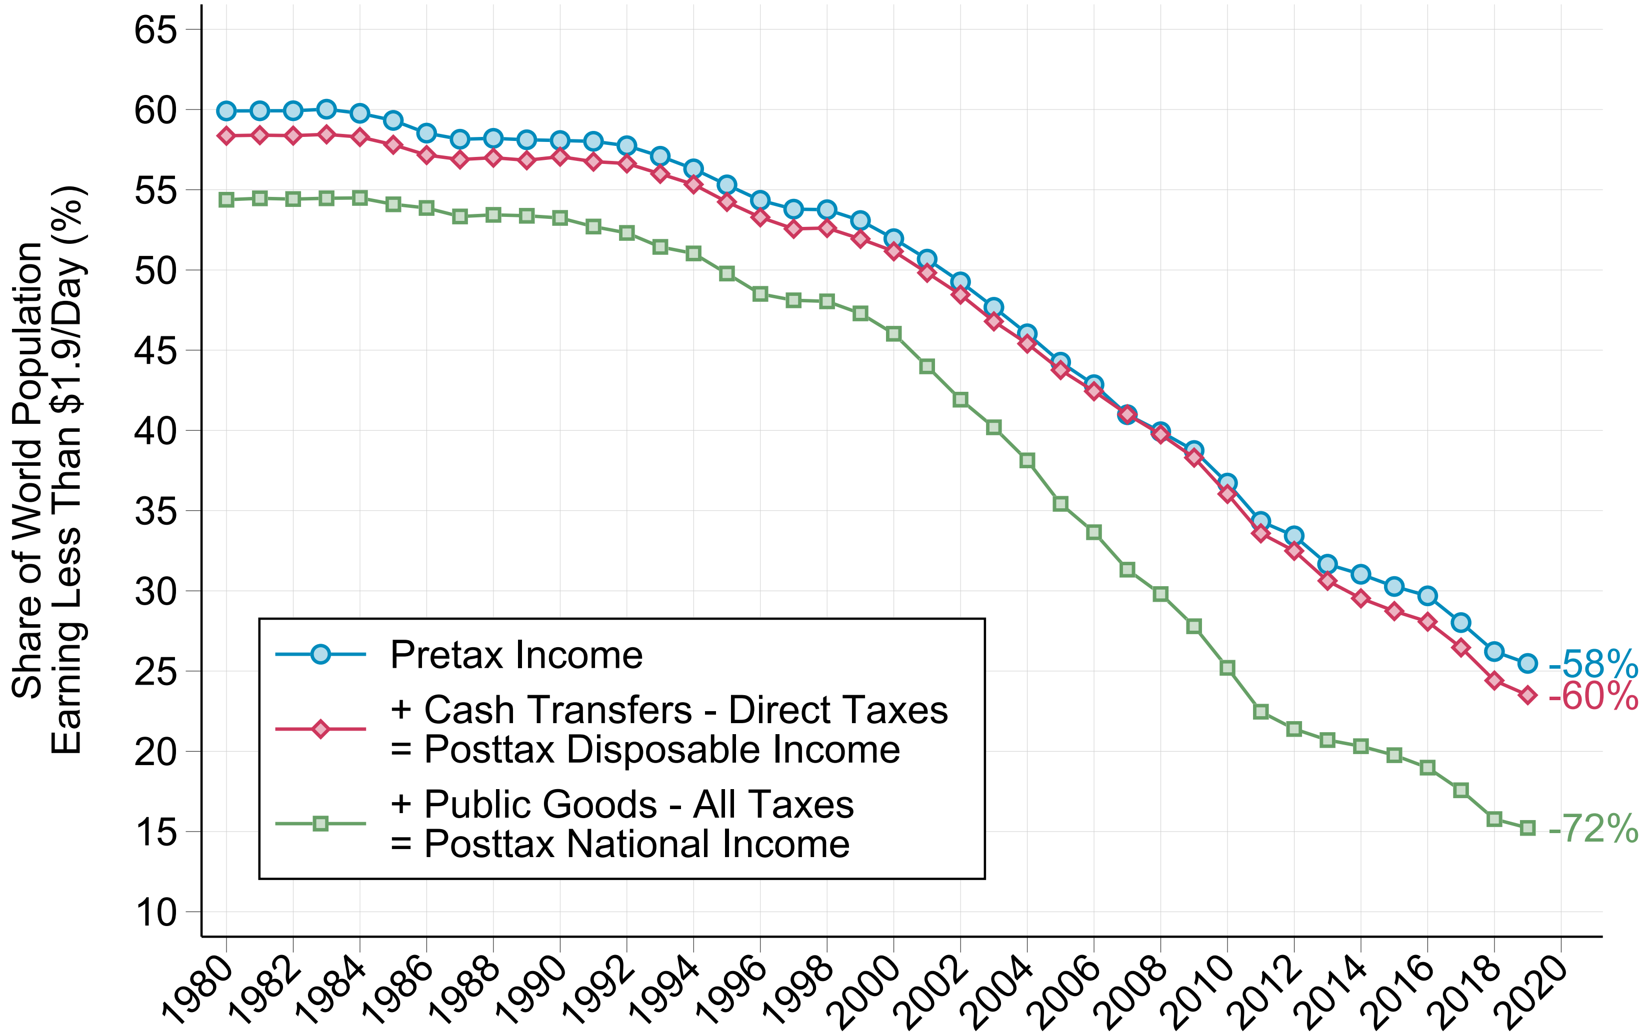

Notes. The figure plots the evolution of the poverty headcount ratio at \$3.65 per day (2017 PPP USD) in the world as a whole, for different income concepts. Distributions of consumption or disposable income per capita from the World Bank. Pretax income is reconstructed as consumption minus social assistance plus direct taxes.

Figure B.7: Global Poverty Headcount Ratio at $6.85 per day, 1980-2019 (World Bank Data)

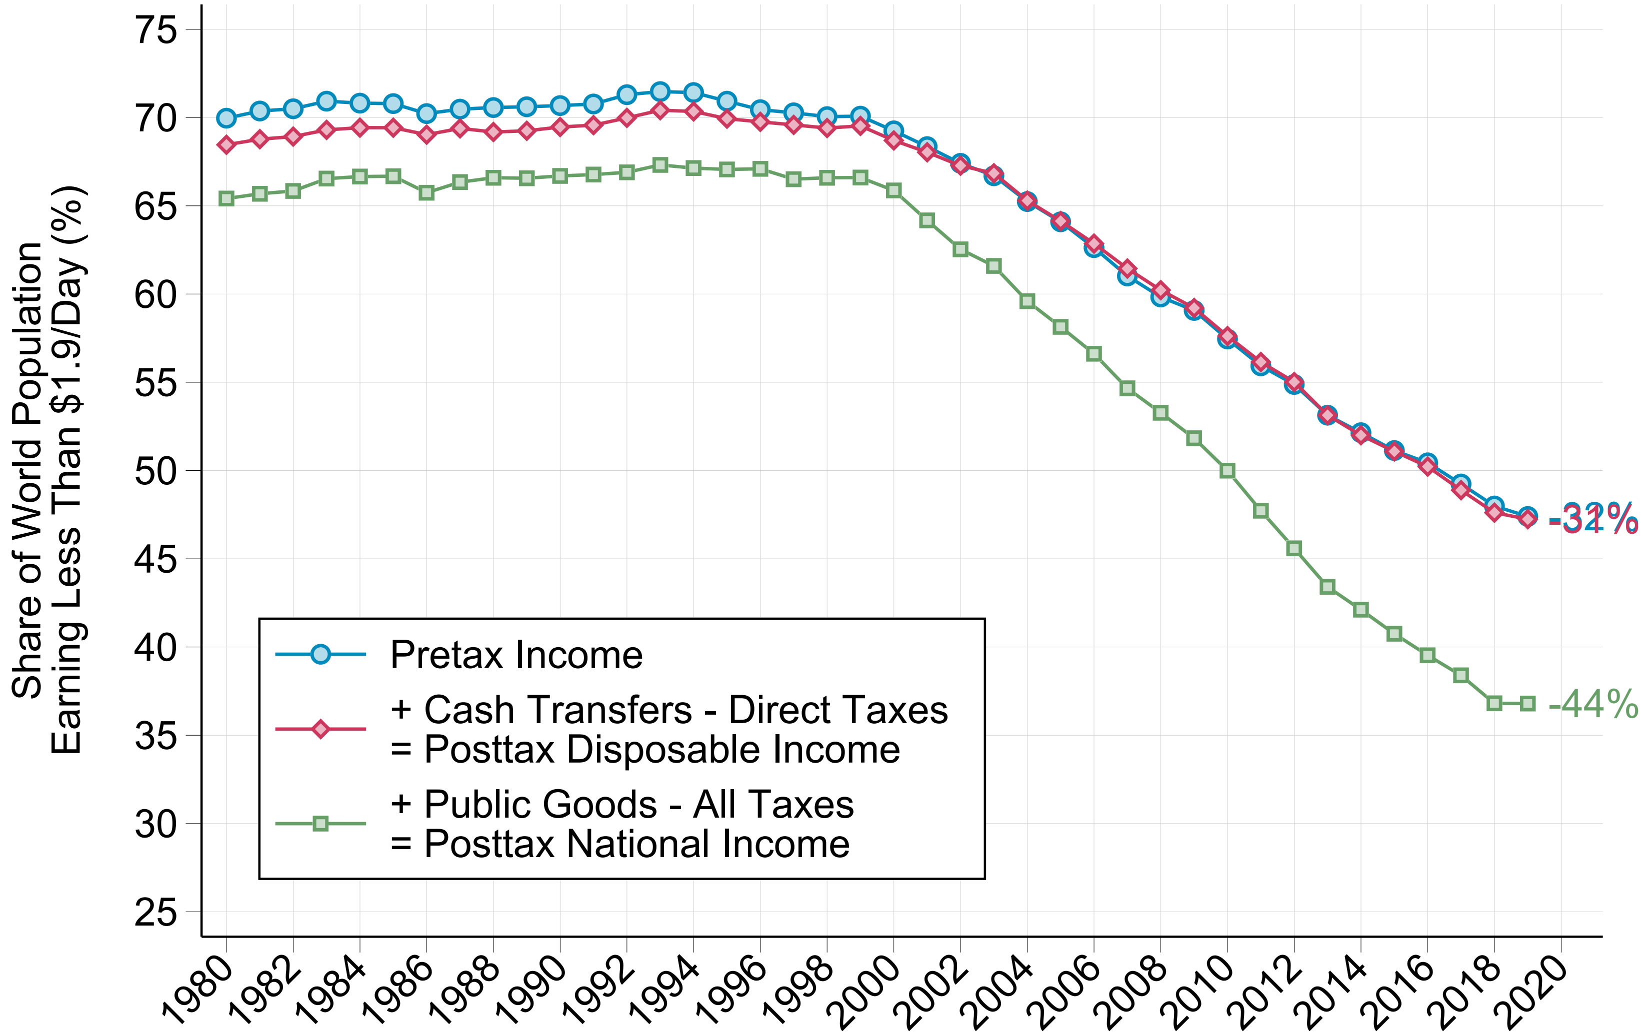

Notes. The figure plots the evolution of the poverty headcount ratio at $6.85 per day (2017 PPP USD) in the world as a whole, for different income concepts. Distributions of consumption or disposable income per capita from the World Bank. Pretax income is reconstructed as consumption minus social assistance plus direct taxes.

Figure B.8: Real Average Income of the Global Bottom 50%, 1980-2019

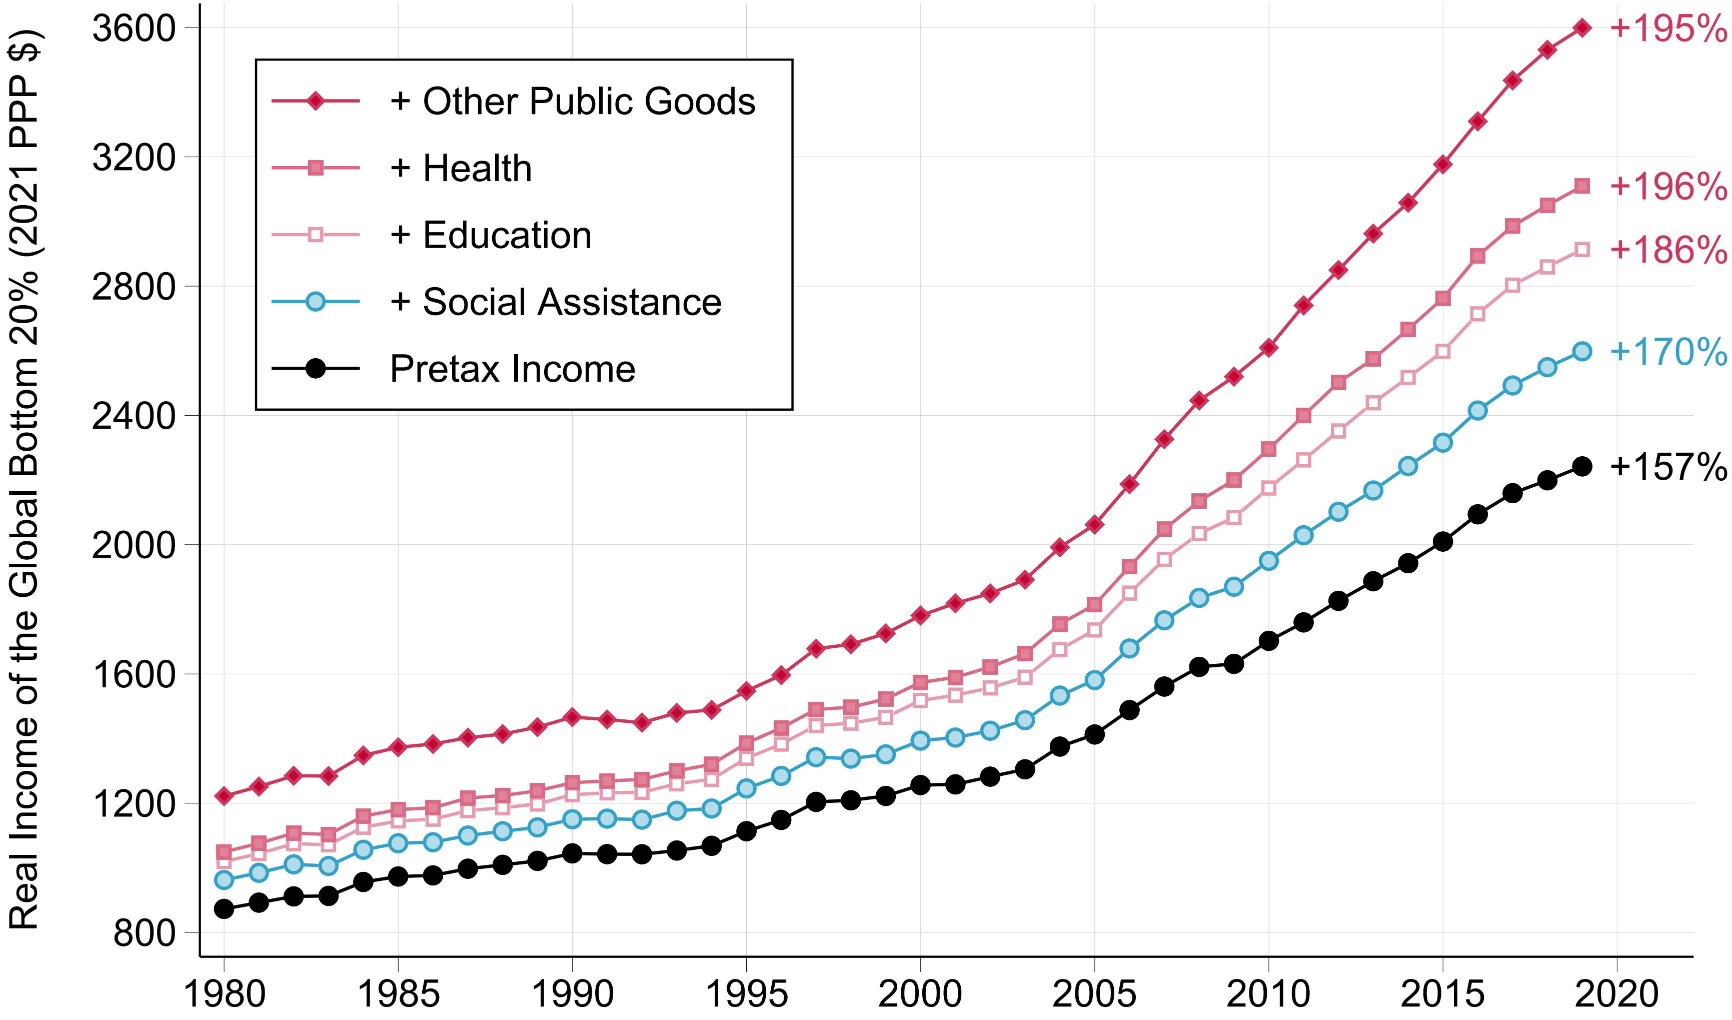

Notes. The figure plots the evolution of the global bottom 50% real average income from 1980 to 2019, before and after accounting for cash transfers and public goods. The unit of observation is the individual. Income is split equally between all household members.

Figure B.9: Total Growth in Posttax Disposable Income and Public Goods Received by Global Income Percentile, 1980-2019

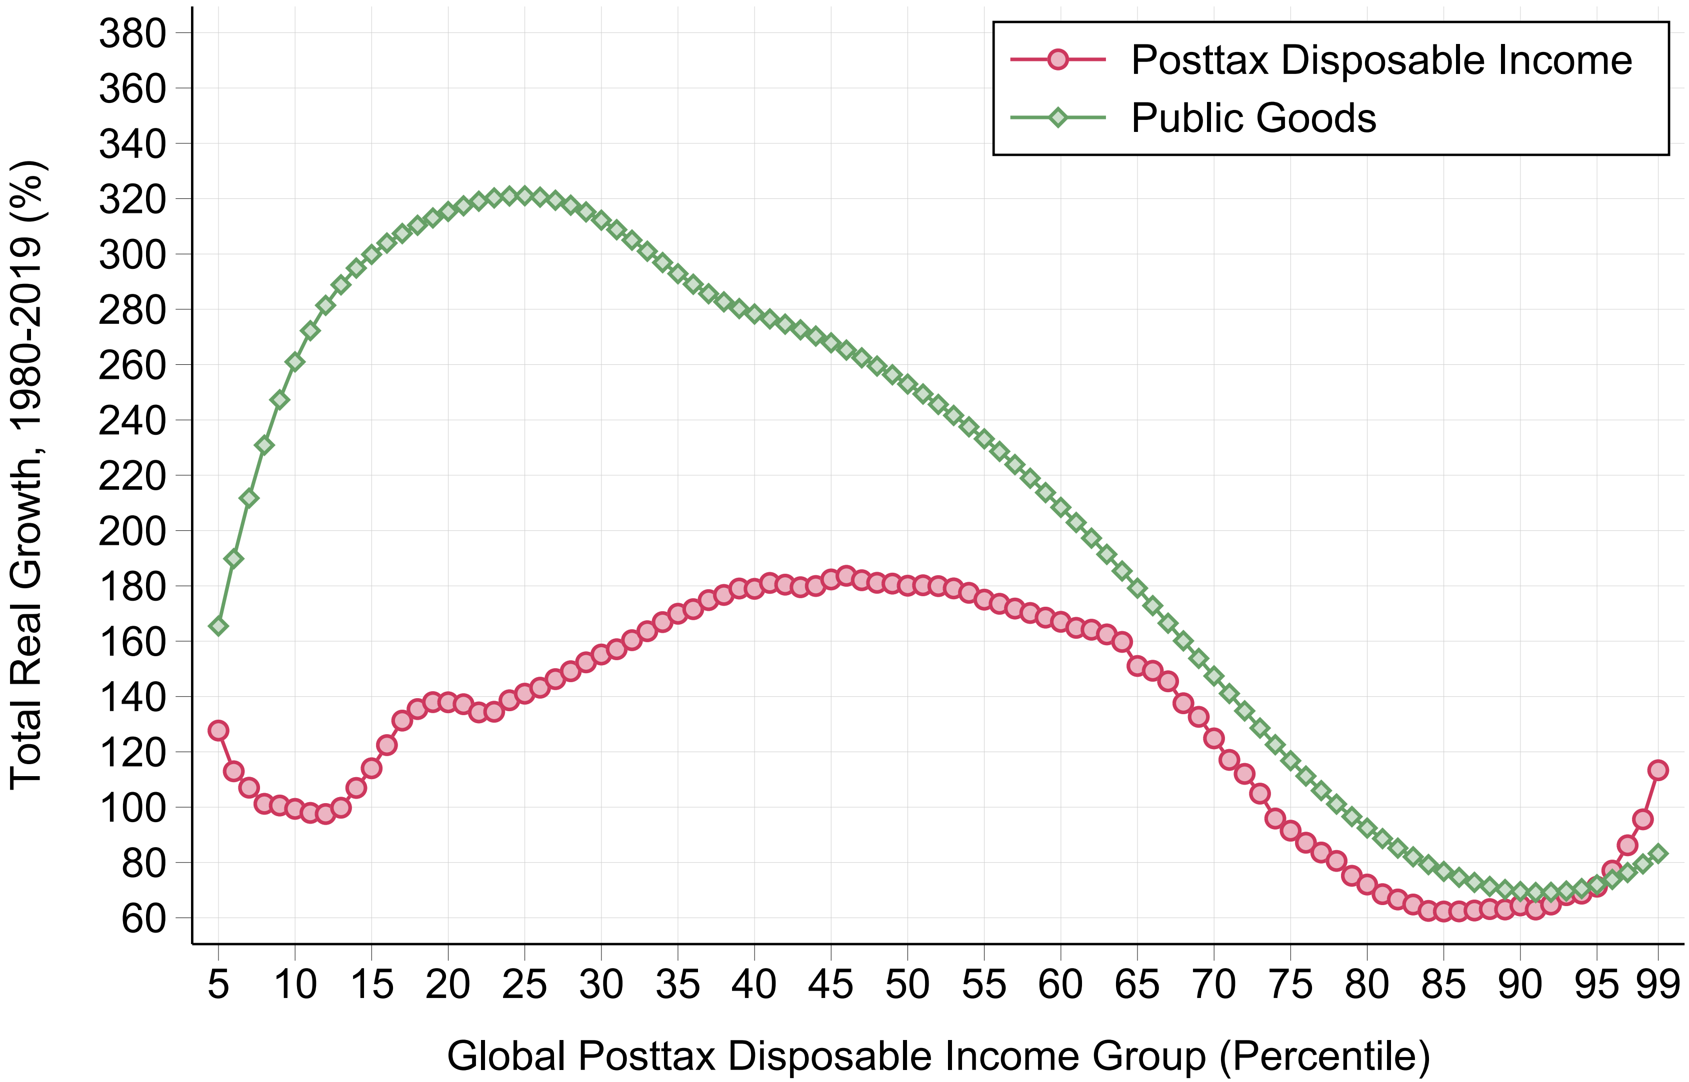

Notes. The figure plots the total growth rate in real posttax disposable income and in the real value of public goods received by global posttax disposable income percentile from 1980 to 2019. The unit of observation is the individual. Income is split equally between all household members. A LOWESS smoothing with 0.5 bandwidth is applied to the public goods curve.

Figure B.10: Total Expenditure on Public Goods Received by Global Income Decile, 1980-2019

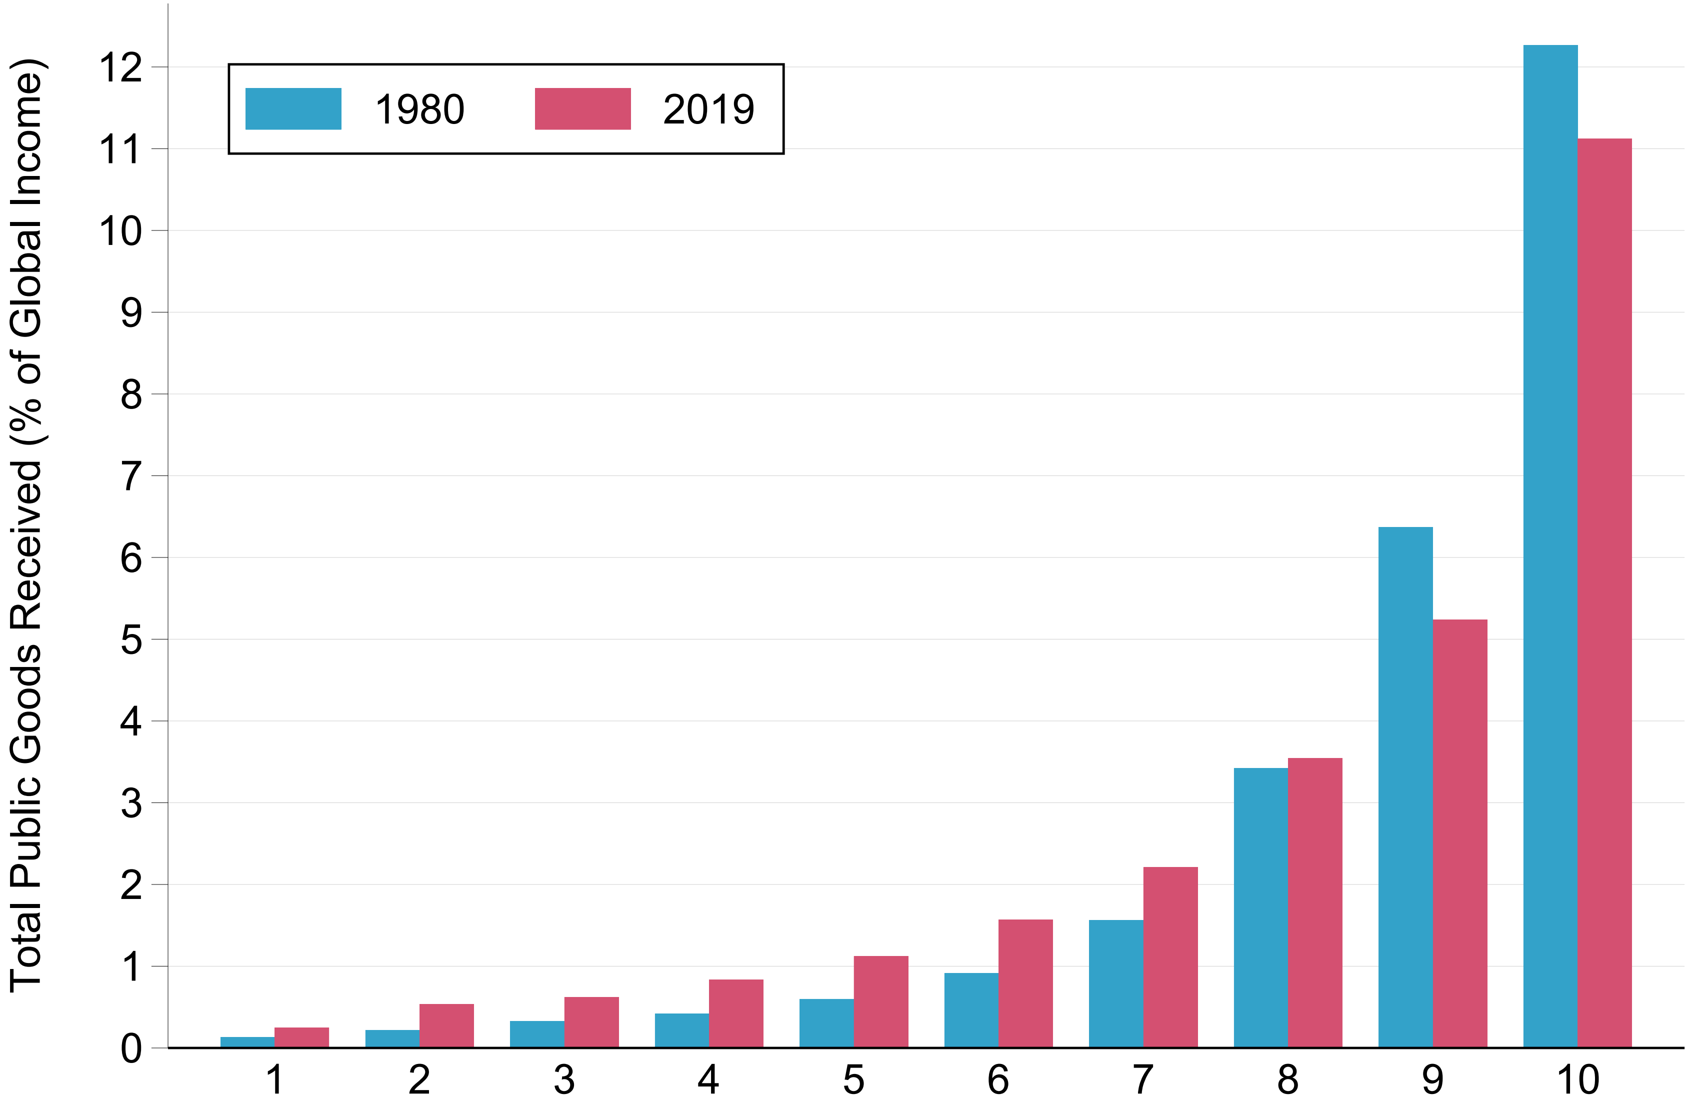

Notes. The figure plots the share of global expenditure on public goods received by global income decile. The unit of observation is the individual. Income is split equally between all household members.

Figure B.11: Total Expenditure on Public Goods Received by the Global Bottom 50% and Top 5%, 1980-2019

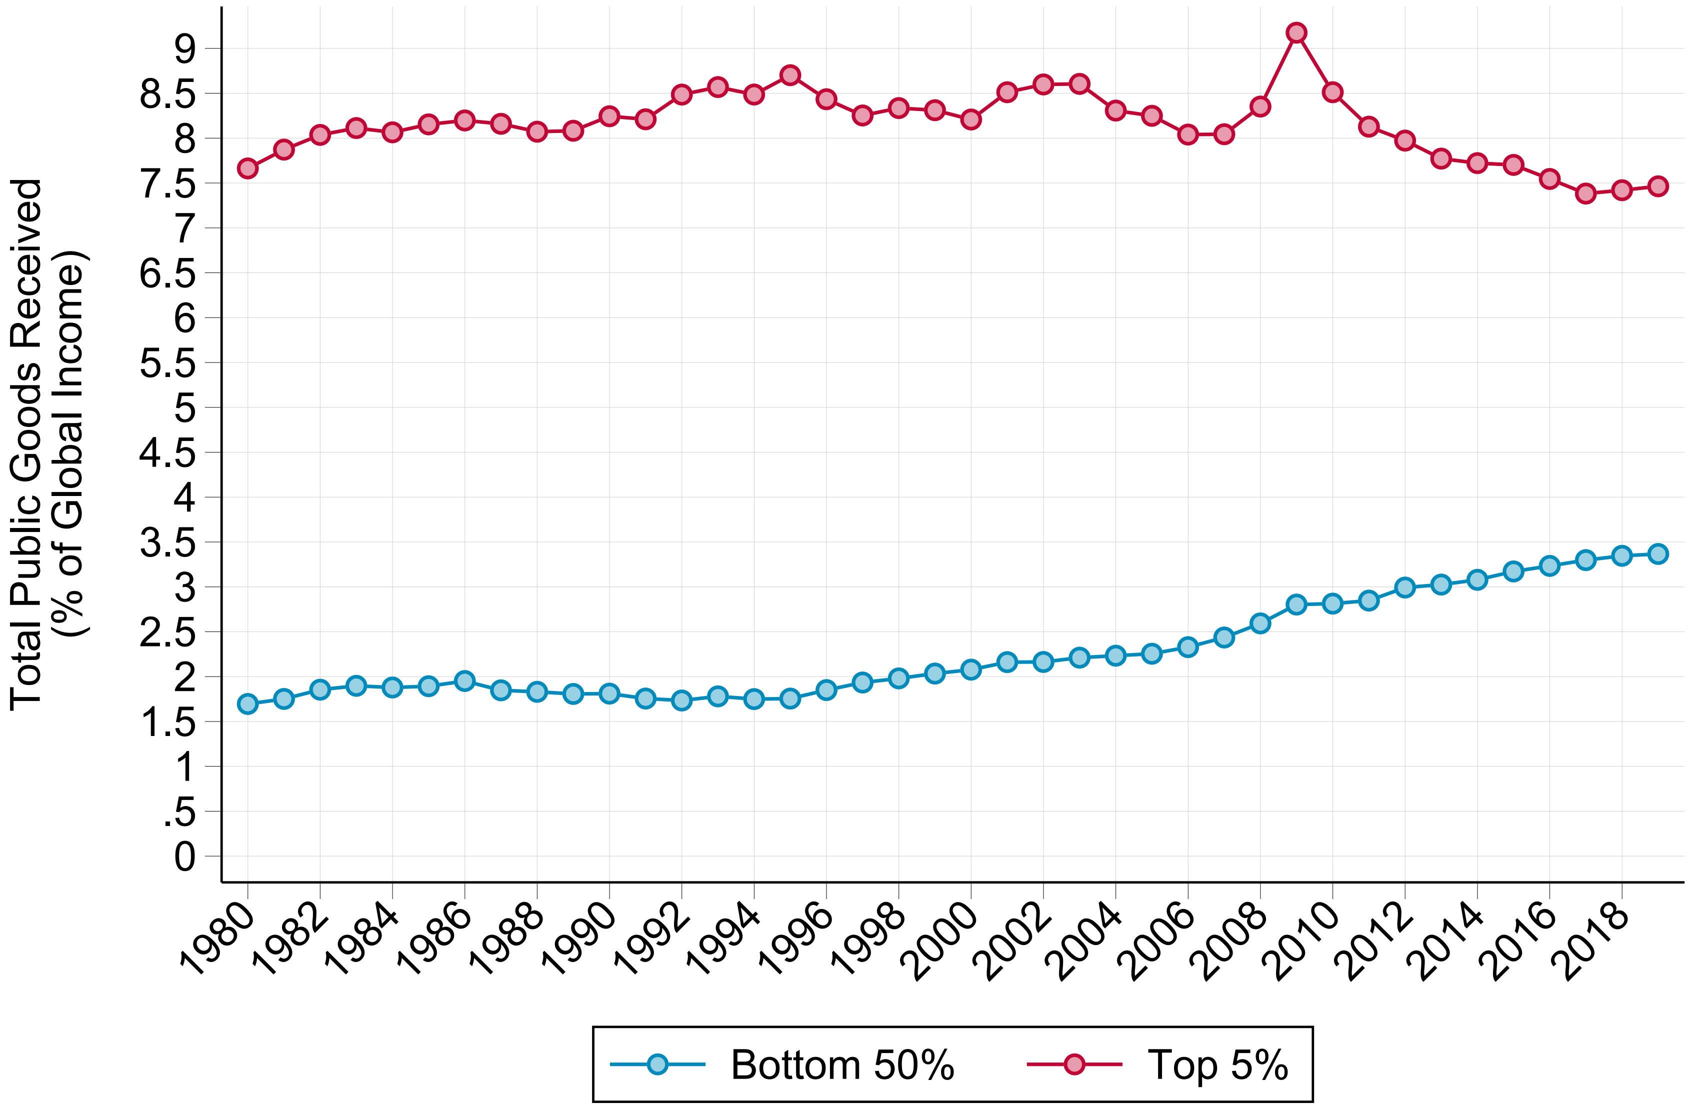

Notes. The figure plots total expenditure on public goods received by the bottom 50% and top 5% of earners in the world as a whole, expressed as a share of global income. The unit of observation is the individual. Income is split equally between all household members.

Figure B.12: Level and Composition of Public Services Received by the Global Bottom 50%, 1980-2019 (% of Global Income)

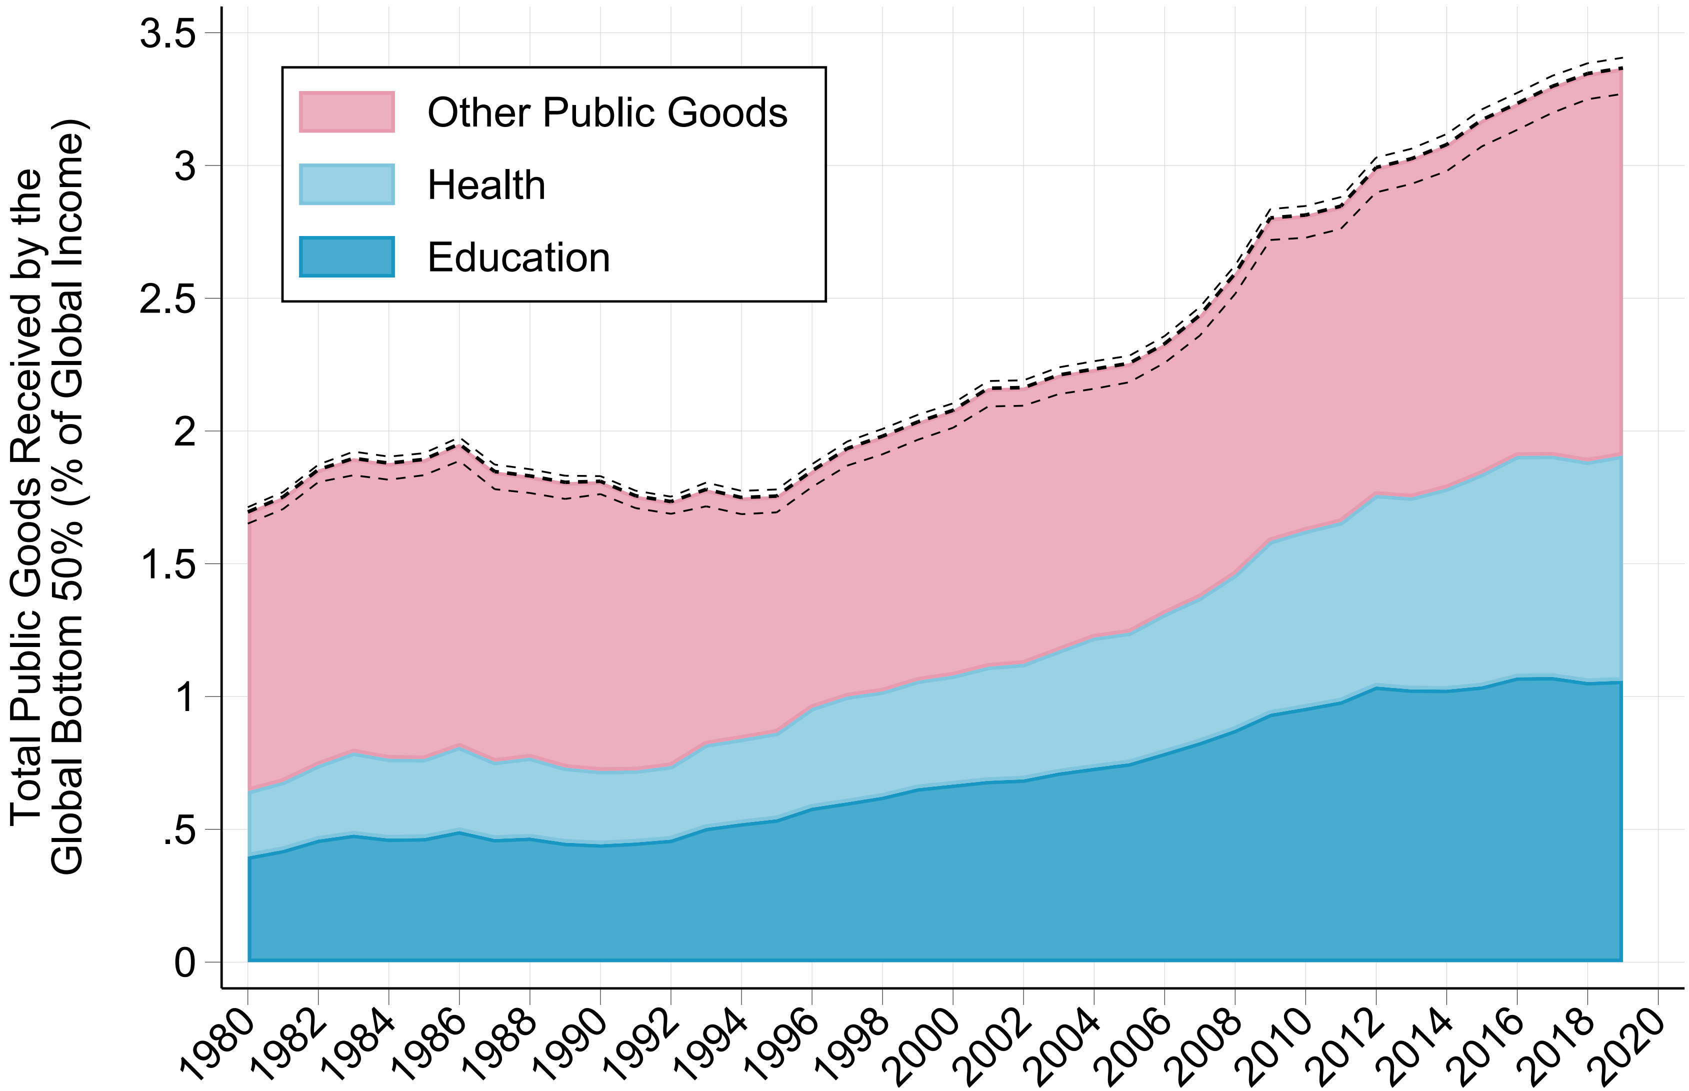

Notes. The figure plots the share of global income accruing to the global bottom 50% in the form of public goods. The unit of observation is the individual. Income is split equally between all household members.

Figure B.13: Level and Composition of Public Services Received by the Global Bottom 20% (Real 2021 PPP USD), 1980-2019

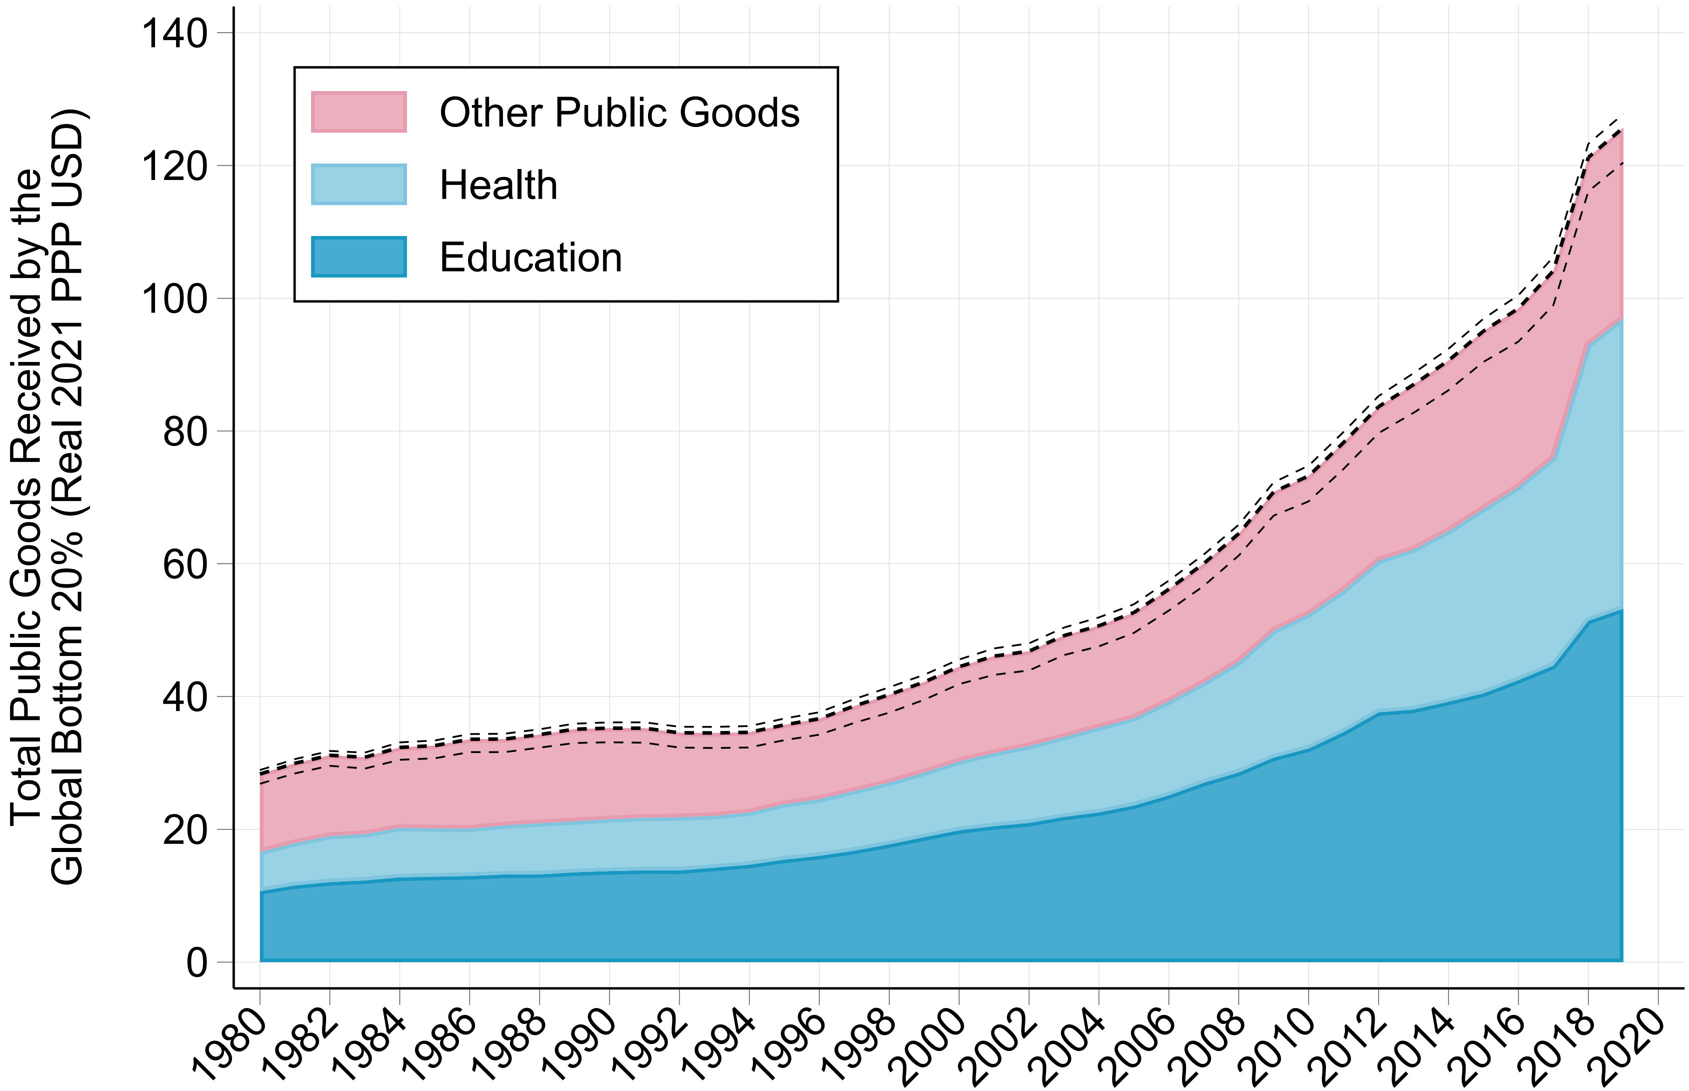

Notes. The figure plots the evolution of public services accruing to the global bottom 20%, expressed in real 2021 PPP US dollars. The unit of observation is the individual. Income is split equally between all household members.

Figure B.14: Level and Composition of Public Services Received by the Global Bottom 50% (Real 2021 PPP USD), 1980-2019

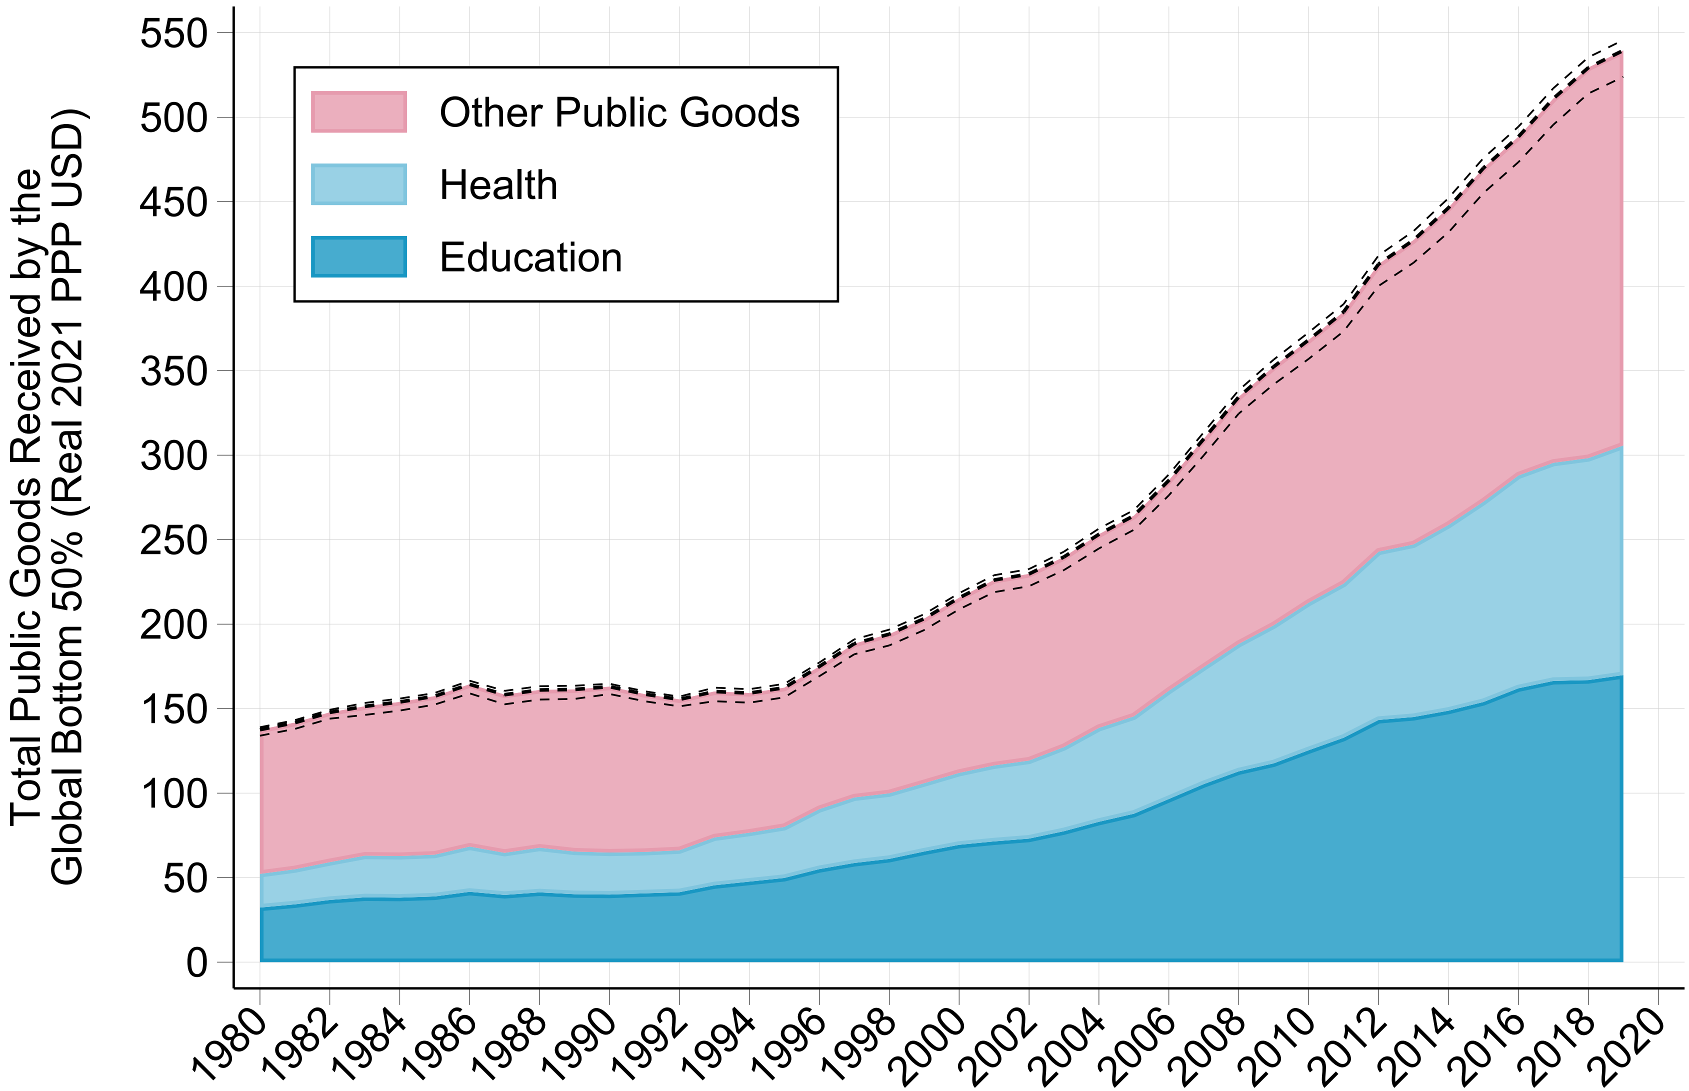

Notes. The figure plots the evolution of public services accruing to the global bottom 50%, expressed in real 2021 PPP US dollars. The unit of observation is the individual. Income is split equally between all household members.

Figure B.15: Gini Index of Global Income Inequality, 1980-2019

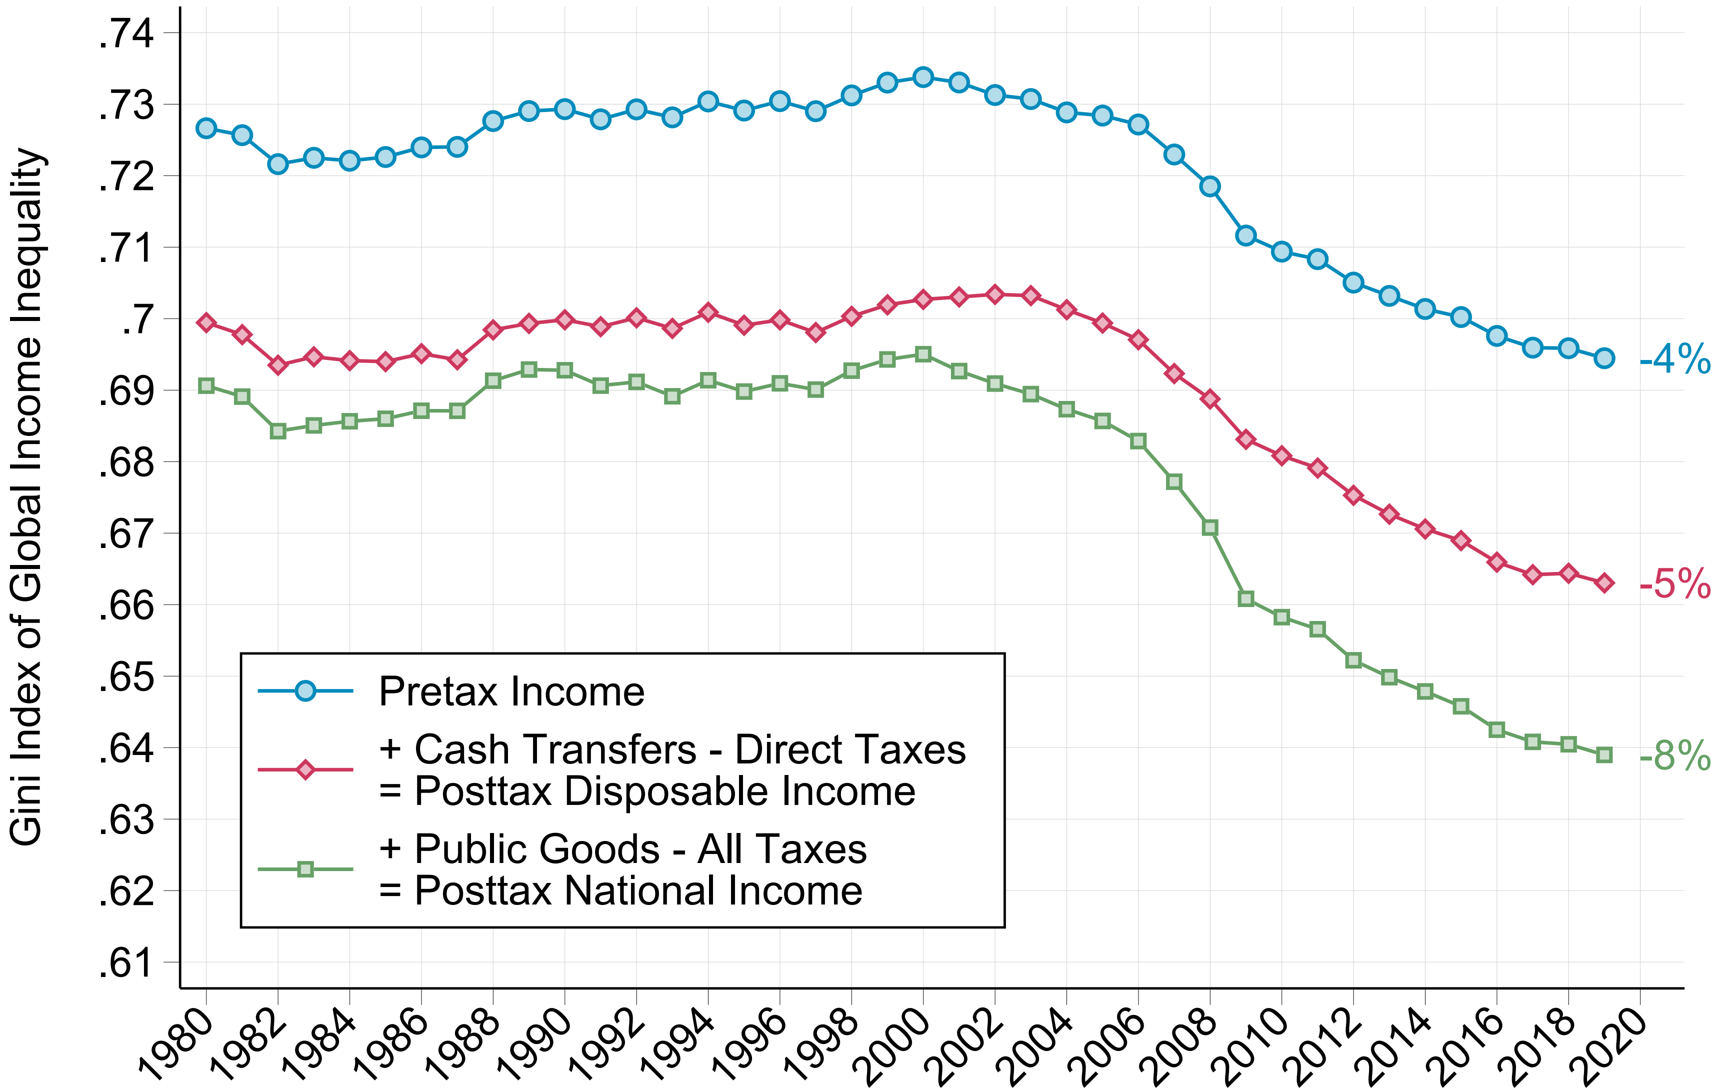

Notes. The figure plots the evolution of the Gini index of global income inequality for different income concepts. The unit of observation is the individual. Income is split equally between all household members.

Figure B.16: Theil Index of Global Income Inequality, 1980-2019

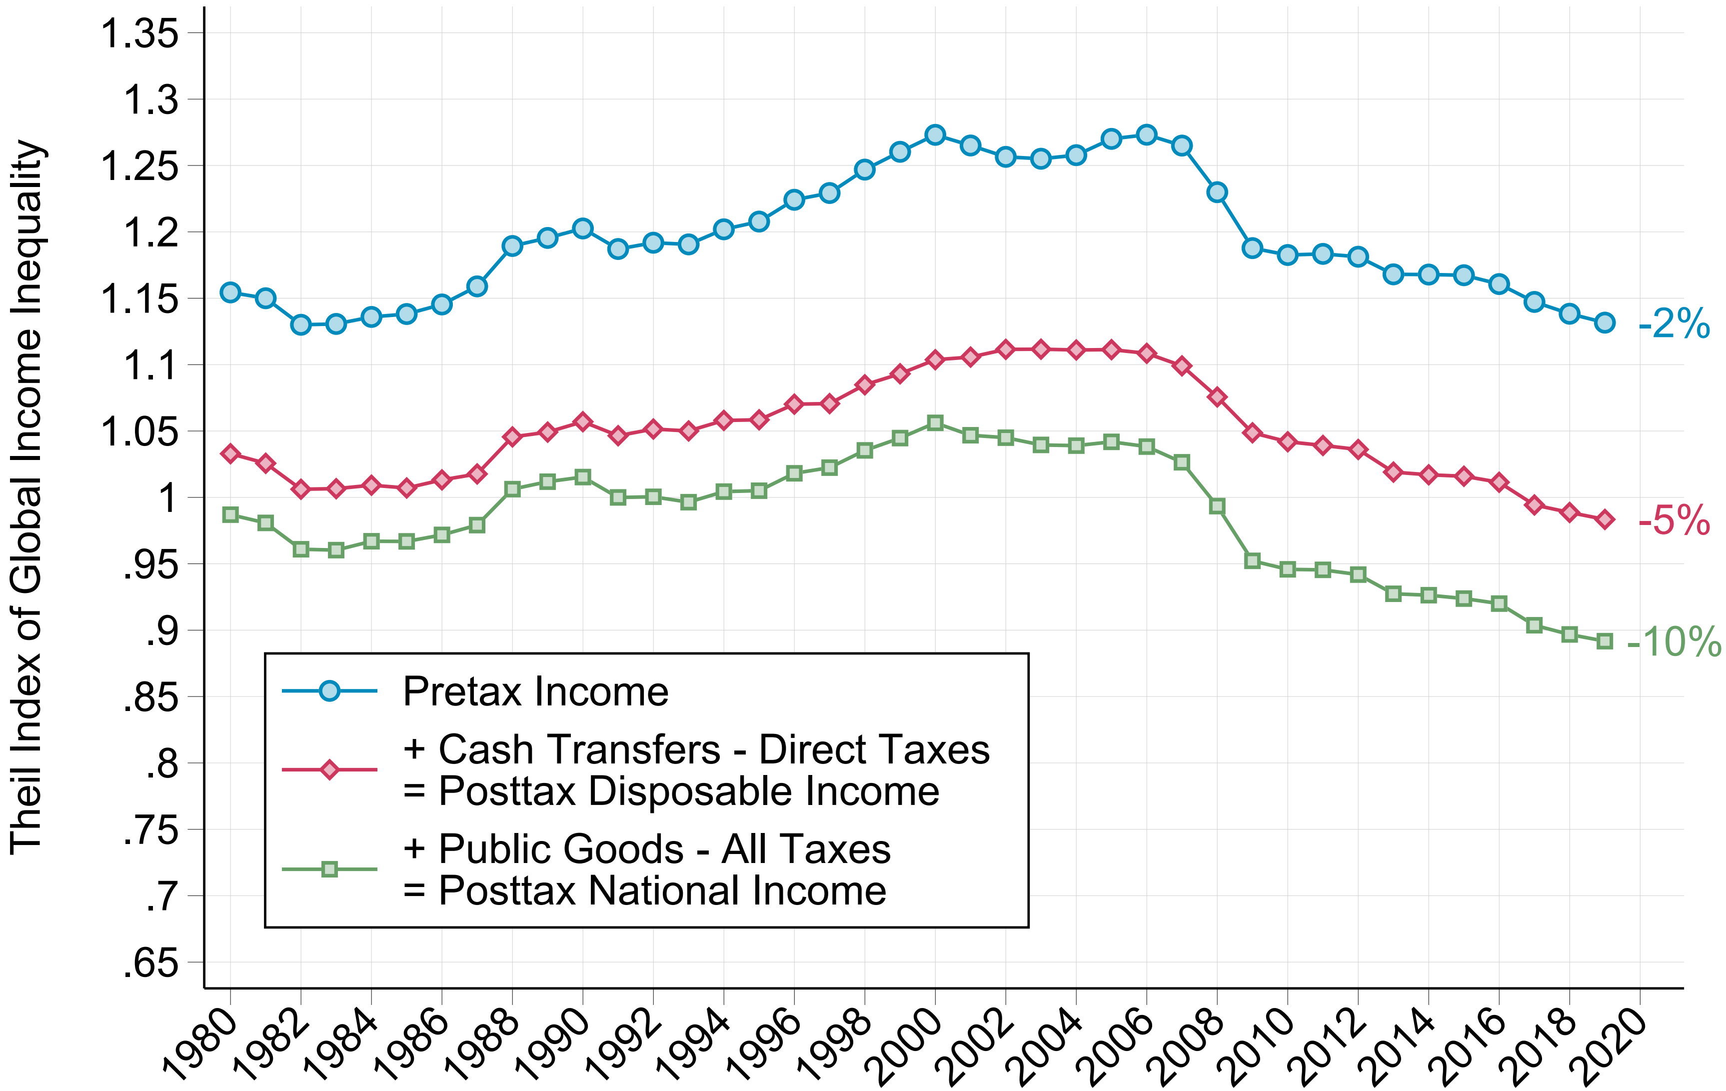

Notes. The figure plots the evolution of the Theil index of global income inequality for different income concepts. The unit of observation is the individual. Income is split equally between all household members.

Figure B.17: Share of Global Income Inequality Explained by Between-Country Inequalities, 1980-2019

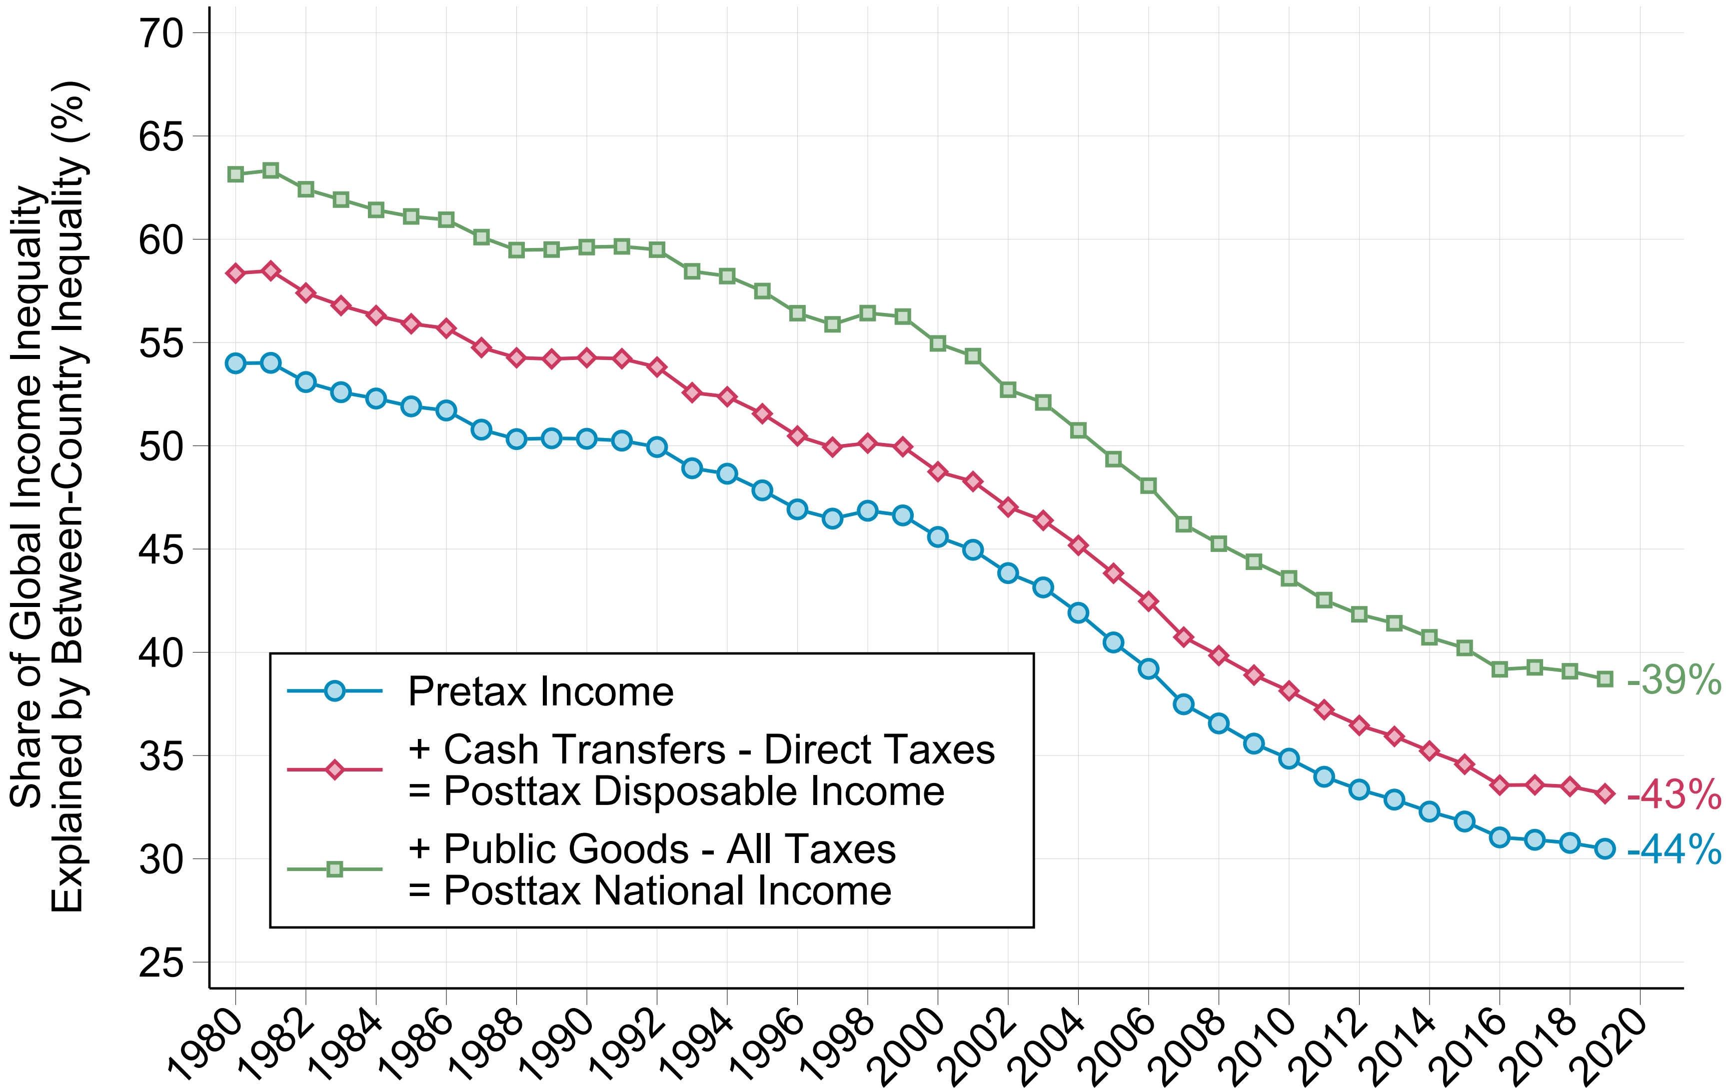

Notes. The figure plots the evolution of the share of global income inequality explained by differences in average incomes between countries, computed from a Theil decomposition of global inequality into a between-country component and a within-country component. The unit of observation is the individual. Income is split equally between all household members.

Figure B.18: Public Goods Received by the Global Bottom 20%: With Productivity-Adjusted Estimates

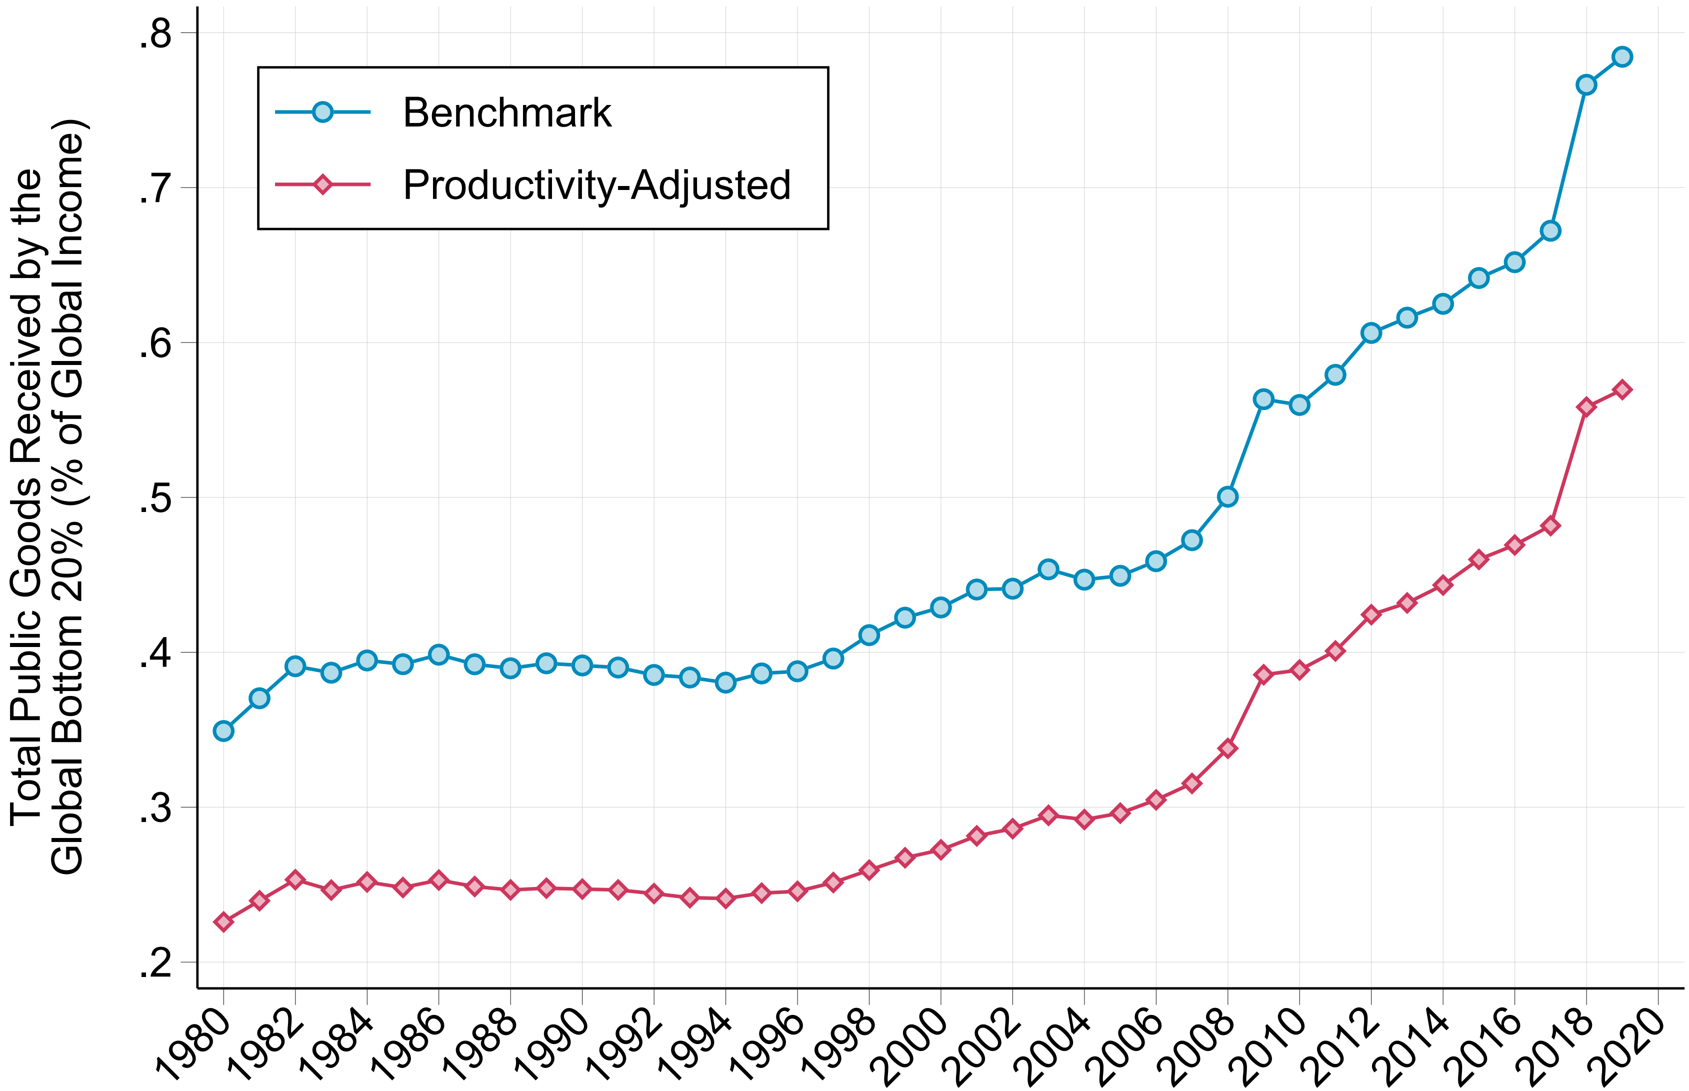

Notes. The figure plots the level of public goods accruing to the global bottom 20%, expressed as a share of global income, before and after adjusting for aggregate and heterogeneous productivity. The unit of observation is the individual. Income is split equally between all household members.

Figure B.19: Global Poverty Headcount Ratio, 1980-2019: With Productivity-Adjusted Estimates

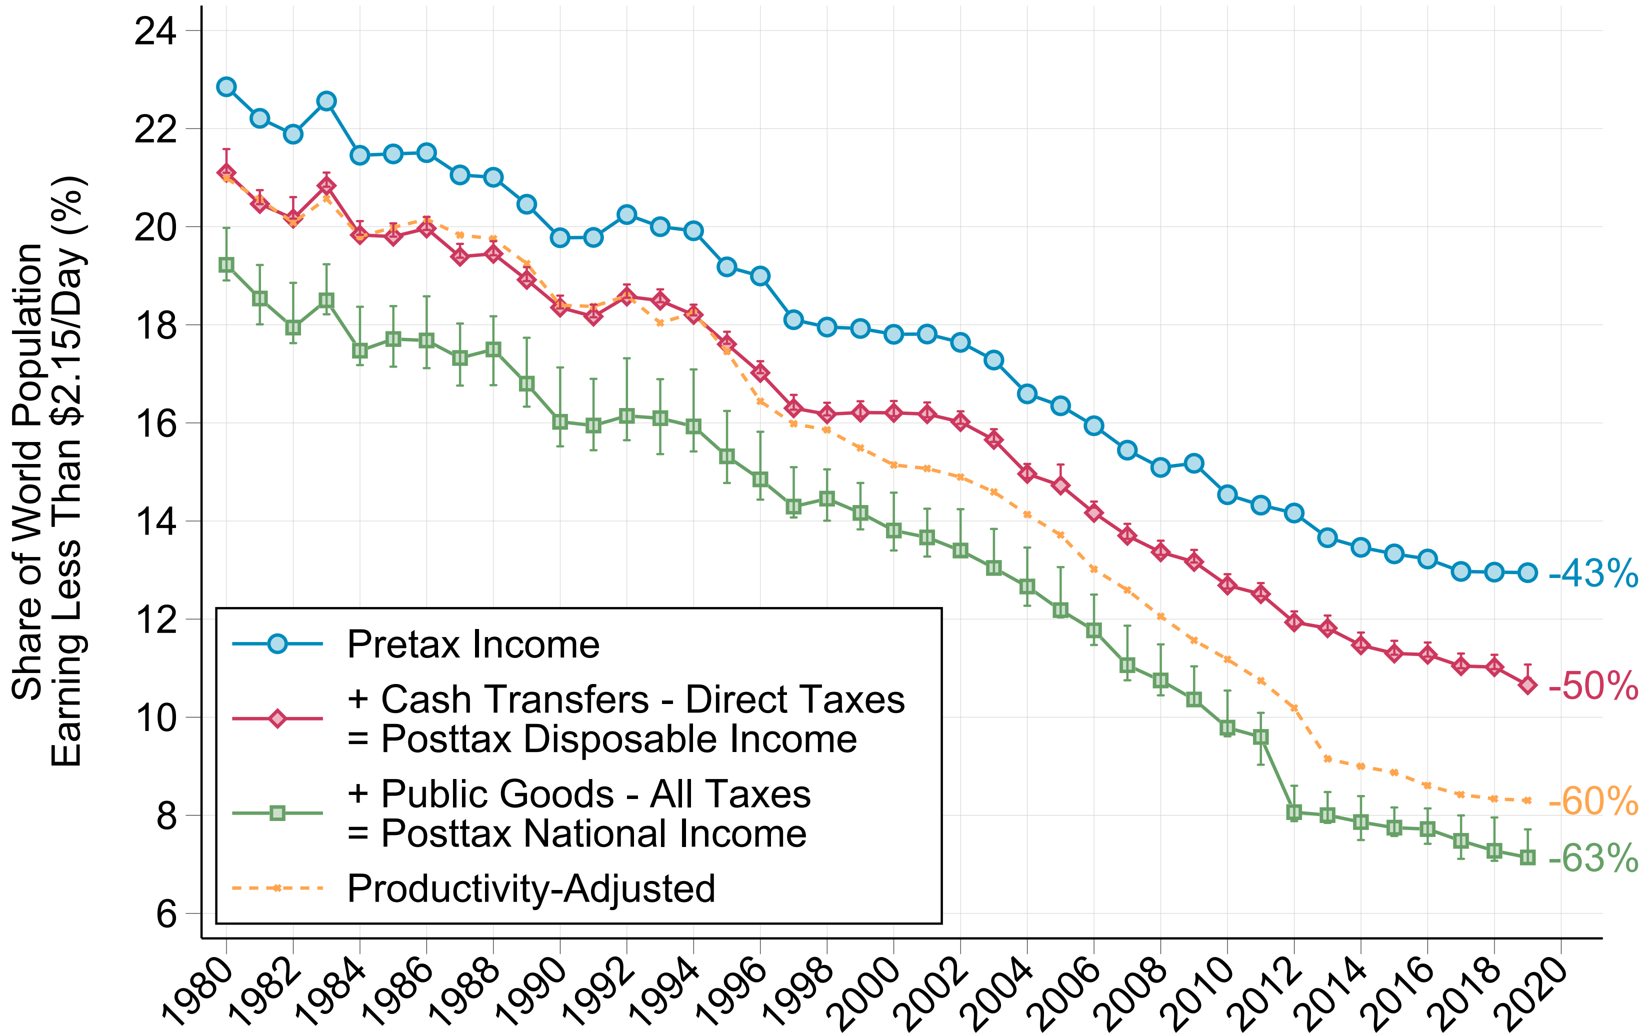

Notes. The figure plots the evolution of the poverty headcount ratio at $2.15 per day (2017 PPP USD) in the world as a whole, for different income concepts. Posttax disposable income removes direct taxes and adds cash transfers. Posttax national income removes all taxes and adds all cash and in-kind transfers. The unit of observation is the individual. Income is split equally between all household members.

Figure B.20: Real Income Growth Rate by Global Income Percentile, 1980-2019:
With Productivity-Adjusted Estimates

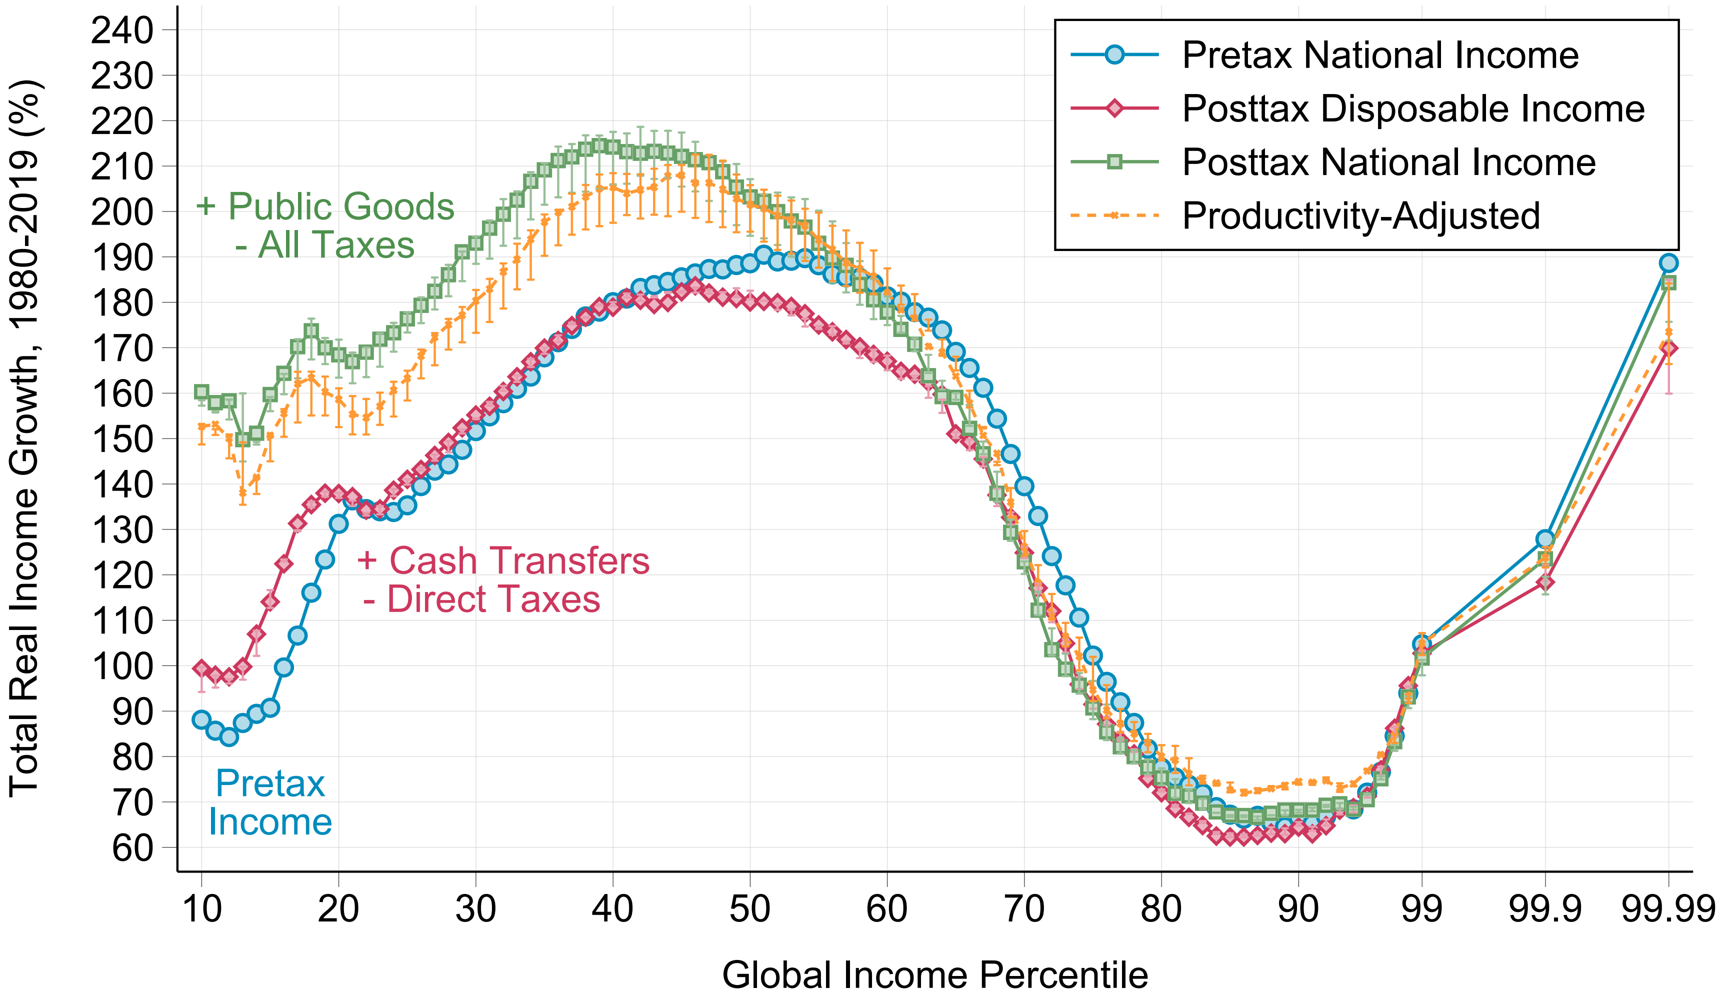

Notes. The figure plots total real income growth by global income percentile from 1980 to 2019 for different income concepts. Posttax disposable income removes direct taxes and adds cash transfers. Posttax national income removes all taxes and adds all cash and in-kind transfers. The unit of observation is the individual. Income is split equally between all household members.

Figure B.21: Decomposing the Incidence of Public Goods on Global Poverty:
With Productivity-Adjusted Estimates, \$2.15 Threshold

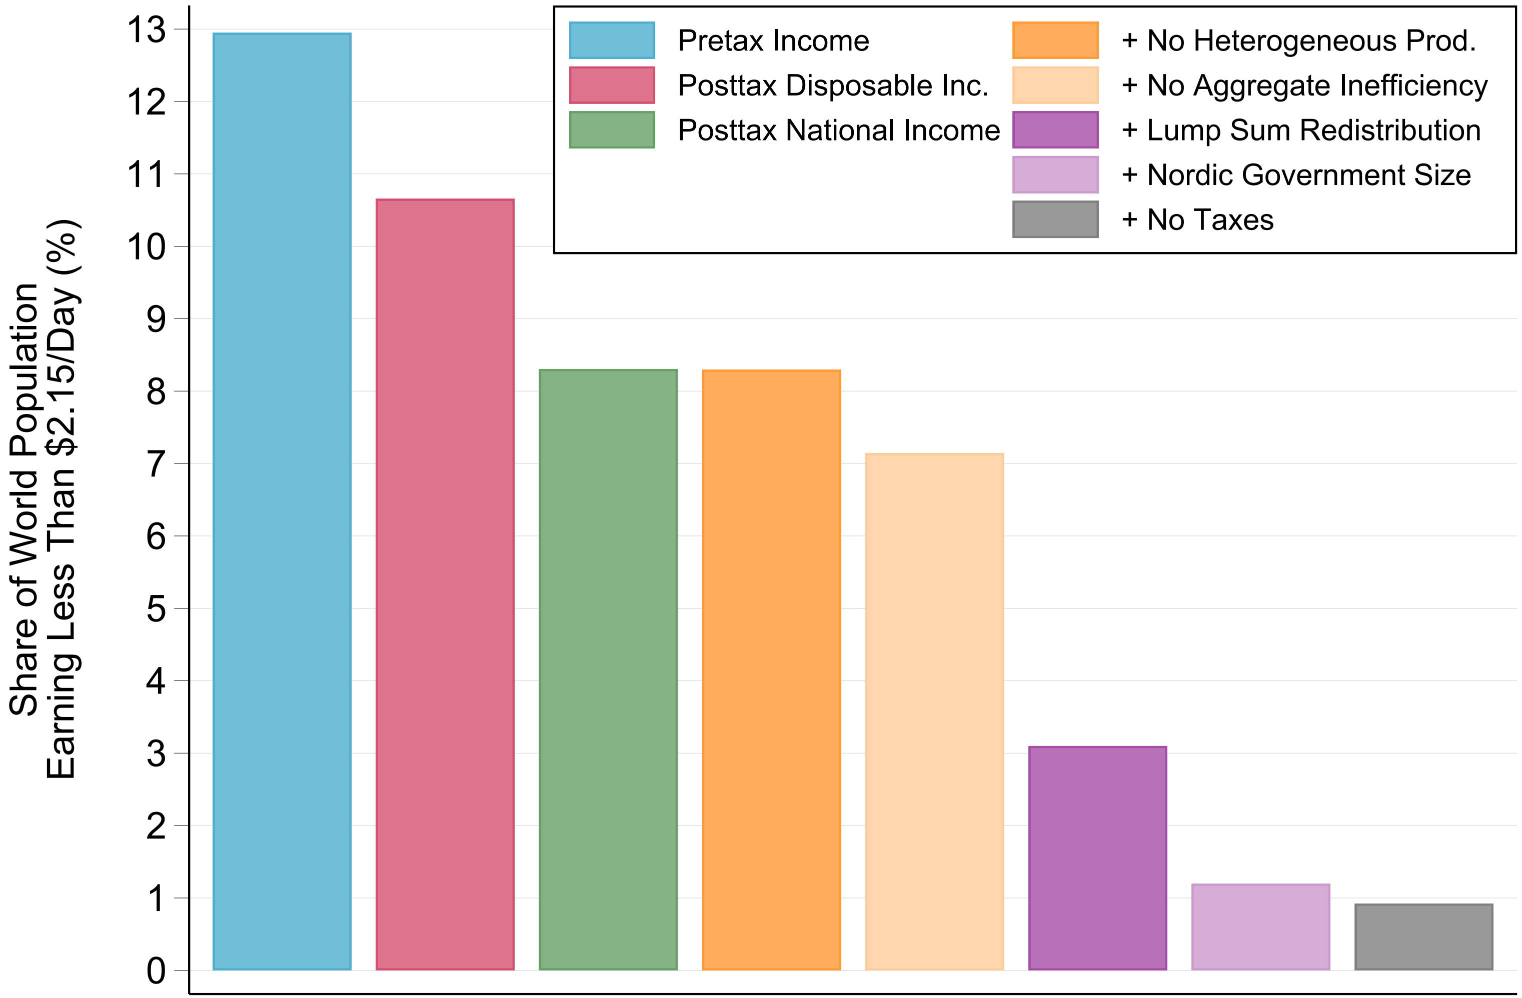

Notes. The figure plots the share of the world population living with less than \$2.15 per day in 2019, measured in 2017 PPP USD, by income concept. Posttax disposable income removes direct taxes and adds cash transfers. Posttax national income removes all taxes and adds all cash and in-kind transfers. The fourth bar assumes no heterogeneous productivity: $q^j(m_i) = 1$. The next bar further assumes no aggregate inefficiency: $\Theta^j = 1$. The next bar assumes that all transfers are received on a lump sum basis: $\gamma(m_i) = \gamma$. The next bar further considers that all countries have welfare states similar to that of Nordic countries, that is, general government expenditure is set at 50% of national income in each country. The last bar considers that no taxes are paid to finance transfers. The unit of observation is the individual. Income is split equally between all household members.

Figure B.22: Decomposing the Incidence of Public Goods on Global Poverty:
With Productivity-Adjusted Estimates, \$3.65 Threshold

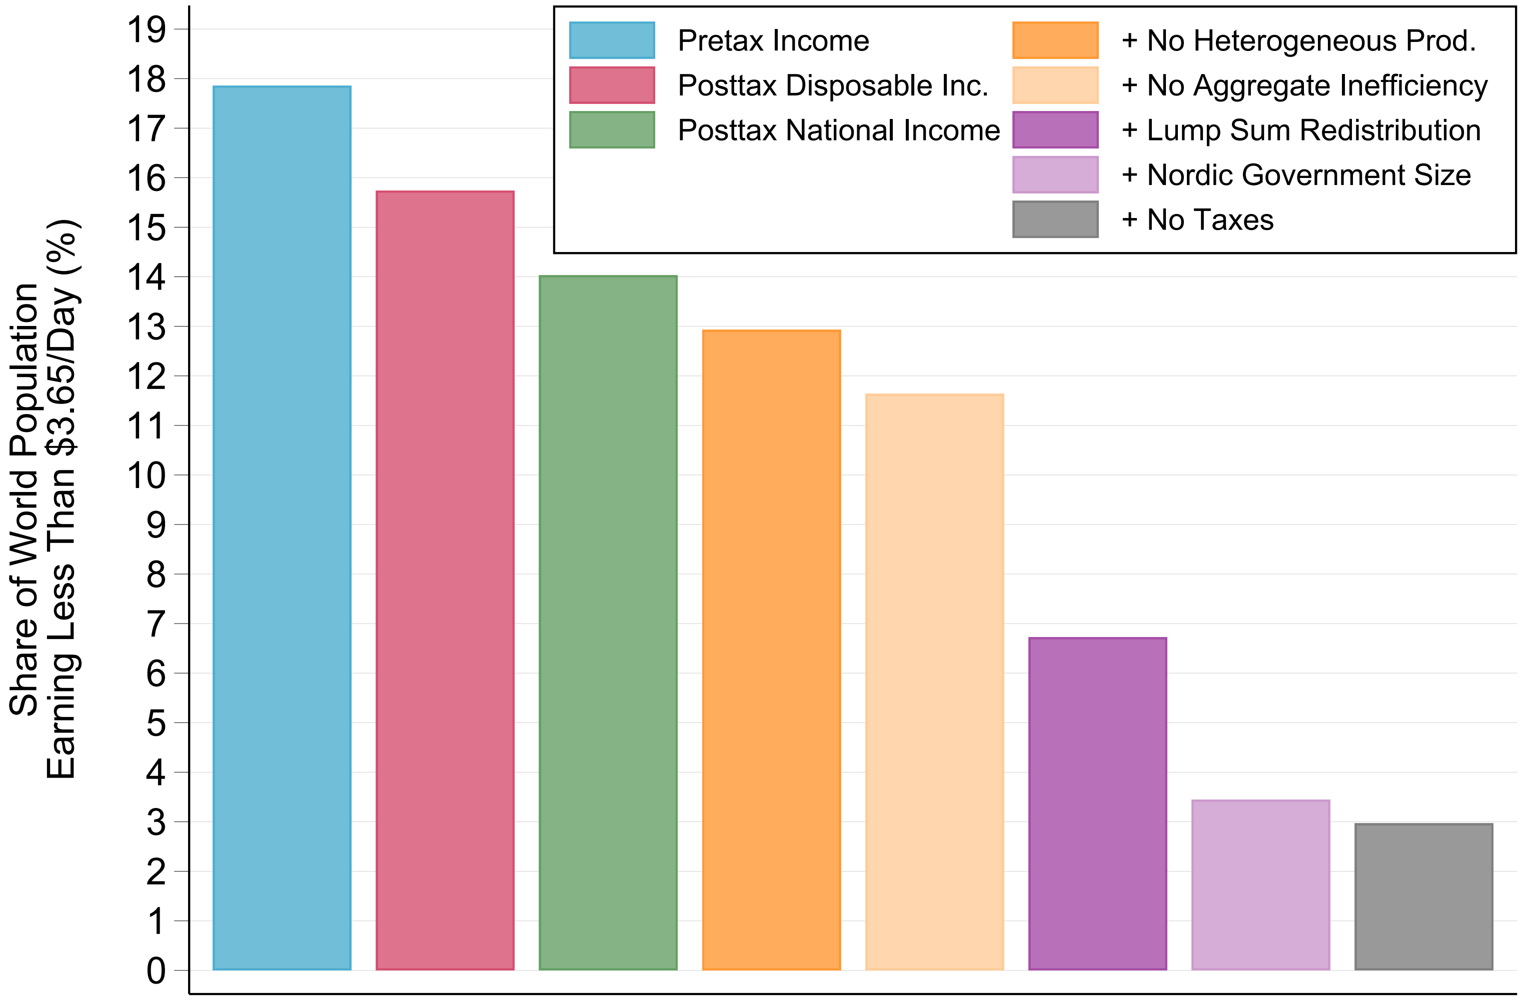

Notes. The figure plots the share of the world population living with less than \$3.65 per day in 2019, measured in 2017 PPP USD, by income concept. Posttax disposable income removes direct taxes and adds cash transfers. Posttax national income removes all taxes and adds all cash and in-kind transfers. The fourth bar assumes no heterogeneous productivity: $q^j(m_i) = 1$. The next bar further assumes no aggregate inefficiency: $\Theta^j = 1$. The next bar assumes that all transfers are received on a lump sum basis: $\gamma(m_i) = \gamma$. The next bar further considers that all countries have welfare states similar to that of Nordic countries, that is, general government expenditure is set at 50% of national income in each country. The last bar considers that no taxes are paid to finance transfers. The unit of observation is the individual. Income is split equally between all household members.

Figure B.23: Decomposing the Incidence of Public Goods on Global Poverty:
With Productivity-Adjusted Estimates, $6.85 Threshold

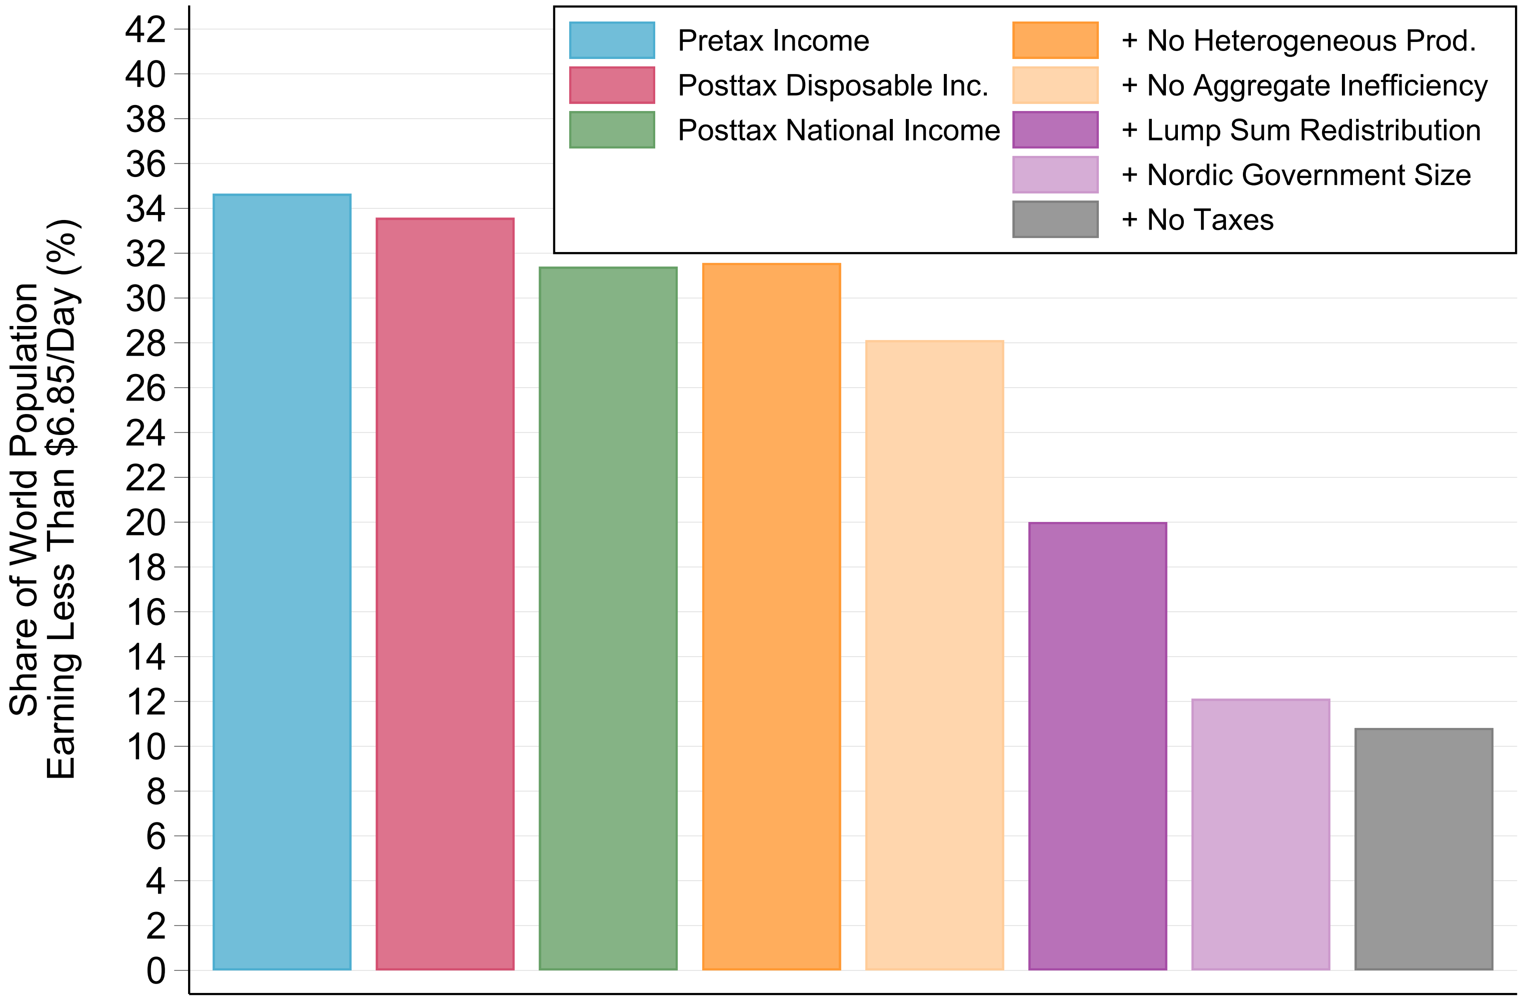

Notes. The figure plots the share of the world population living with less than $6.85 per day in 2019, measured in 2017 PPP USD, by income concept. Posttax disposable income removes direct taxes and adds cash transfers. Posttax national income removes all taxes and adds all cash and in-kind transfers. The fourth bar assumes no heterogeneous productivity: $q^j(m_i) = 1$. The next bar further assumes no aggregate inefficiency: $\Theta^j = 1$. The next bar assumes that all transfers are received on a lump sum basis: $\gamma(m_i) = \gamma$. The next bar further considers that all countries have welfare states similar to that of Nordic countries, that is, general government expenditure is set at 50% of national income in each country. The last bar considers that no taxes are paid to finance transfers. The unit of observation is the individual. Income is split equally between all household members.

Table B.1: Pairwise Correlations Between Dimensions of Government Redistribution Across Countries

	Cost	Progressivity	Aggregate Productivity	Heterogeneous Productivity	NNI per capita
Cost	1.00				
Progressivity	0.60***	1.00			
Aggregate Productivity	0.42***	0.59***	1.00		
Heterogeneous Productivity	0.08	0.49***	0.22***	1.00	
NNI per capita	0.56***	0.71***	0.63***	0.28***	1.00

Notes. The table reports raw correlation coefficients between different dimensions of government redistribution across countries. Cost (C^j) corresponds to total general government expenditure as a share of net national income. Progressivity ($\gamma^j(m_i)$) is measured as the share of total government expenditure received by the bottom 50% (excluding social security). Aggregate productivity (Θ^j) corresponds to single-input, output-oriented estimates for each function of government. Heterogeneous productivity is measured as the relative quality of public services received by the bottom 20% in each country. Statistics computed over all countries in the database ($N = 174$). * p¡0.10, ** p¡0.05, *** p¡0.01.

Table B.2: Public Goods and Global Poverty Reduction:
Sensitivity to Different Specifications and Geographical Restrictions

	Global Poverty Headcount Ratio at $1.9 Per Day		
	1980	2019	2019-1980
All Countries			
Posttax Disposable Income	21.1%	10.7%	-50%
Posttax National Income: Benchmark	19.2%	7.1%	-63%
Posttax National Income: Only Education & Health	26.1%	9.6%	-63%
Posttax National Income: Other Public Goods Lump Sum	15.6%	3.3%	-79%
Excluding China			
Posttax Disposable Income	16.9%	9.4%	-44%
Posttax National Income	15.2%	7.1%	-53%
Excluding India			
Posttax Disposable Income	14.5%	8.7%	-40%
Posttax National Income	12.9%	5.7%	-56%
Excluding China & India			
Posttax Disposable Income	10.2%	7.4%	-27%
Posttax National Income	8.9%	5.7%	-36%

Notes. The table reports how results on the incidence of public goods on global poverty reduction vary depending on assumptions regarding the progressivity of public goods and geographical restrictions. Only Education and Health: only allocate education and health expenditure. Other Public Goods Lump Sum: allocate all public goods other than education and health on a lump sum basis.

Table B.3: Public Goods Provision Over the Course of Development: Before and After Adjusting for Productivity

	Expenditure (% NNI) G	Share of Transfer Received (%) (γ, Bottom 50%)	Net Transfer Received (% NNI) (g, Bottom 50%)	Adjusted for Productivity (g, Bottom 50%)
Country Income Group				
Low-Income	23.3%	21.0%	4.9%	3.0%
Lower-Middle-Income	26.3%	23.3%	6.1%	4.0%
Upper-Middle-Income	25.6%	28.1%	7.1%	5.2%
High-Income	30.4%	33.0%	10.0%	8.3%
World Region				
Sub-Saharan Africa	25.9%	20.9%	5.4%	3.2%
Middle East and Northern Africa	28.6%	24.7%	7.0%	5.1%
China	23.3%	25.4%	5.9%	5.0%
India	31.4%	18.6%	5.8%	3.4%
Other Asia / Oceania	23.3%	27.1%	6.4%	4.8%
Latin America	25.8%	28.3%	7.2%	5.1%
US / Canada / Western Europe	30.3%	35.0%	10.6%	8.9%

Notes. The table reports statistics on dimensions of in-kind redistribution by country income group (defined based on the World Bank's classification) and world region. All figures focus on public goods, that is, total government expenditure excluding social protection spending. The last column adjusts estimates for differences in aggregate and heterogeneous productivity across countries.

Table B.4: Public Goods, Quality of Life, and the Gap Between Surveys and National Accounts

	Expected Years of Schooling		Youth Literacy		Secondary School Enrollment Rate		Infant Mortality		Life Expectancy	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A: No FE										
GDP-Survey Gap	0.16***	0.06***	0.09***	0.02	0.29***	0.10***	0.57***	-0.01	0.06***	0.00
	(0.02)	(0.01)	(0.02)	(0.02)	(0.03)	(0.02)	(0.07)	(0.03)	(0.01)	(0.00)
Educ./Health Spending		0.19***		0.11***		0.29***		0.60***		0.06***
		(0.00)		(0.01)		(0.01)		(0.01)		(0.00)
Panel B: Country FE										
GDP-Survey Gap	-0.04***	-0.01	0.08***	0.04	0.13***	-0.10***	0.33***	0.02	0.03***	0.00
	(0.01)	(0.02)	(0.03)	(0.03)	(0.03)	(0.03)	(0.05)	(0.04)	(0.01)	(0.01)
Educ./Health Spending		0.23***		0.04***		0.34***		0.67***		0.06***
		(0.01)		(0.01)		(0.01)		(0.01)		(0.00)
N	1193	1194	285	285	1409	1409	1760	1760	1772	1772
Adj. R-squared	0.93	0.88	0.86	0.87	0.82	0.89	0.88	0.95	0.87	0.92

Notes. Each column presents coefficients of a regression of a selected dependent variable on the gap between GDP and survey means, before and after controlling for education or health spending. GDP-Survey Gap: percentage difference between GDP per capita and survey mean income. Educ./Health Spending: log of public education spending (expected years of schooling, youth literacy, secondary school enrollment rate) or log of public health spending (infant mortality, life expectancy). Panel A runs simple OLS regressions. Panel B includes country fixed effects.

Table B.5: Political Correlates of Public Goods Redistribution

	(1)	(2)	(3)	(4)	(5)	(6)
Electoral Democracy Index (0-1)	1.212***	1.423***	0.975***	0.745**	1.155***	0.986***
	(0.280)	(0.295)	(0.331)	(0.317)	(0.338)	(0.350)
Political Competition Index (0-10)	-0.032*	-0.042**	0.011	-0.010	0.003	0.013
	(0.018)	(0.019)	(0.022)	(0.016)	(0.017)	(0.017)
Public Sector Corruption Index (0-1)	-0.583**	-0.412*	-0.581**	0.254	0.355	0.366
	(0.230)	(0.248)	(0.266)	(0.284)	(0.310)	(0.319)
Government Effectiveness (0-1)	-1.116***	-0.761*	-1.816***	-0.689	-0.601	0.413
	(0.395)	(0.424)	(0.491)	(0.470)	(0.502)	(0.551)
Log GDP Per Capita	0.784***	0.732***	0.761***	0.283**	0.159	0.092
	(0.061)	(0.065)	(0.071)	(0.126)	(0.136)	(0.143)
Additional Controls	X	X	X	X	X	X
Country FE				X	X	X
Excl. Western Democracies			X			X
Sample	1980-2019	2000-2019	2000-2019	1980-2019	2000-2019	2000-2019
N	2915	2637	2089	2915	2637	2089
Adj. R-squared	0.65	0.64	0.48	0.90	0.91	0.88

Notes. The table reports the results of a linear regression of redistribution on a number of political and economic variables. Redistribution is measured as the share of national income received by the bottom 50% in the form of public services. All estimates include country and year fixed effects and control for the following additional variables: bottom 50% pretax income share, log of total population, share of population aged 0-19, 20-39, and 40-59, and trade to GDP ratio. Country FE: country fixed effects. Excl. Western Democracies: excludes Western European countries, Canada, the United States, New Zealand, and Australia from the sample.

B.3.2 Macroeconomic Aggregates

Figure B.24: Data Coverage of Total Government Expenditure and Revenue by Source

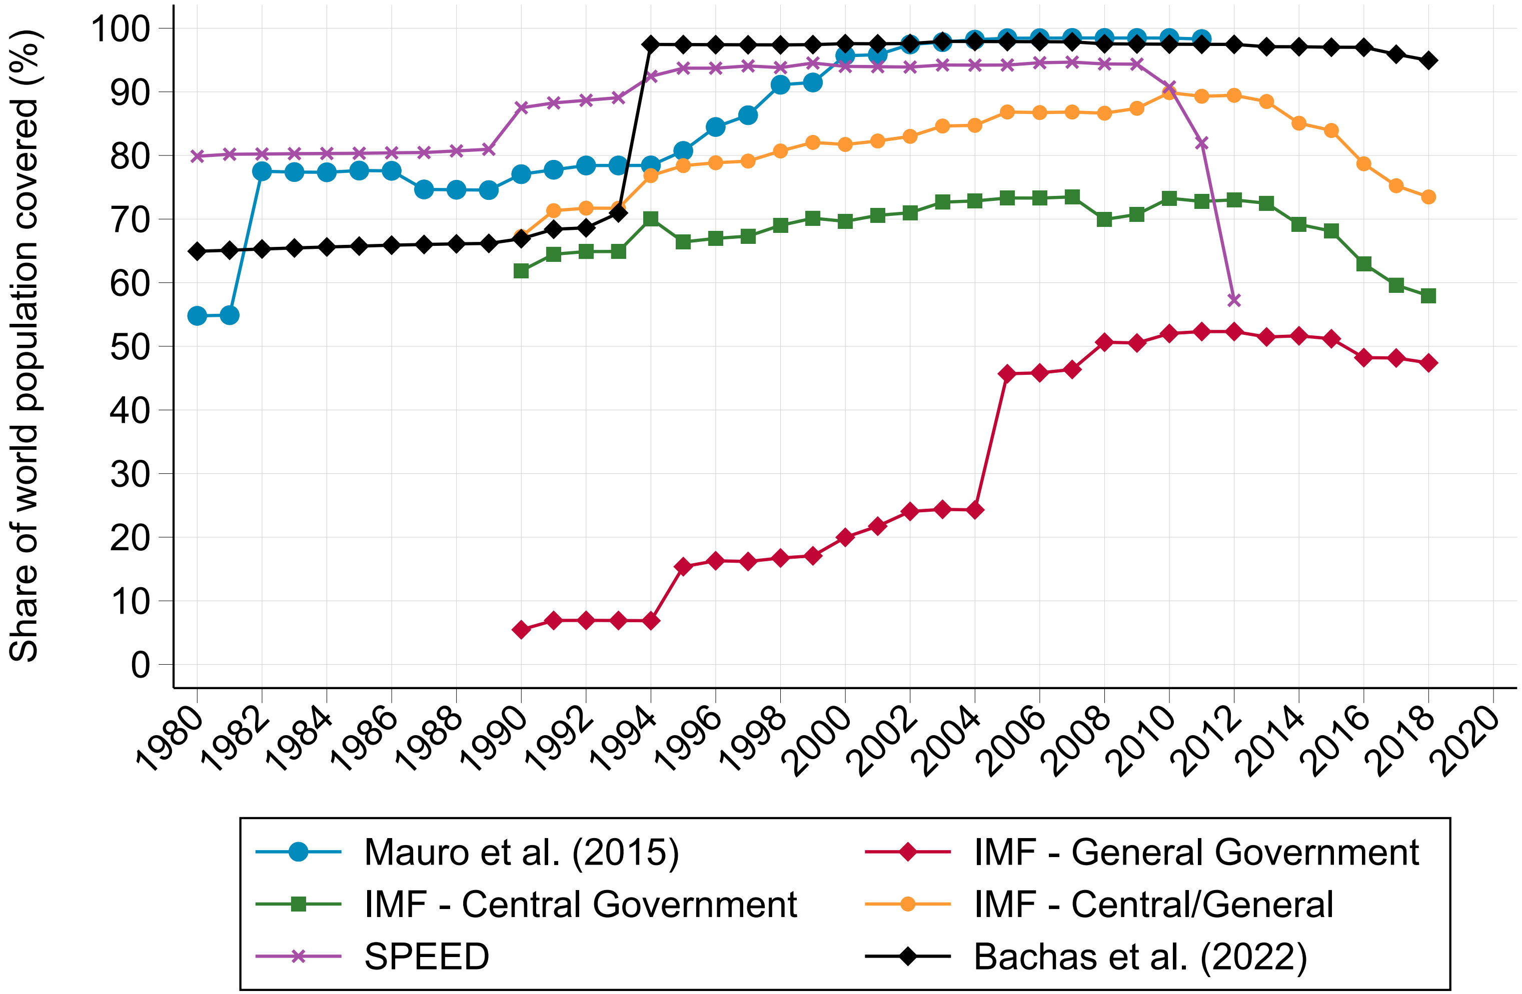

Notes. The figure shows the share of the world population covered by the different sources used to construct harmonized general government expenditure and central government revenue (in the case of Bachas et al. 2022) series.

Figure B.25: Global Government Expenditure, 1980-2019 (% of Global Income)

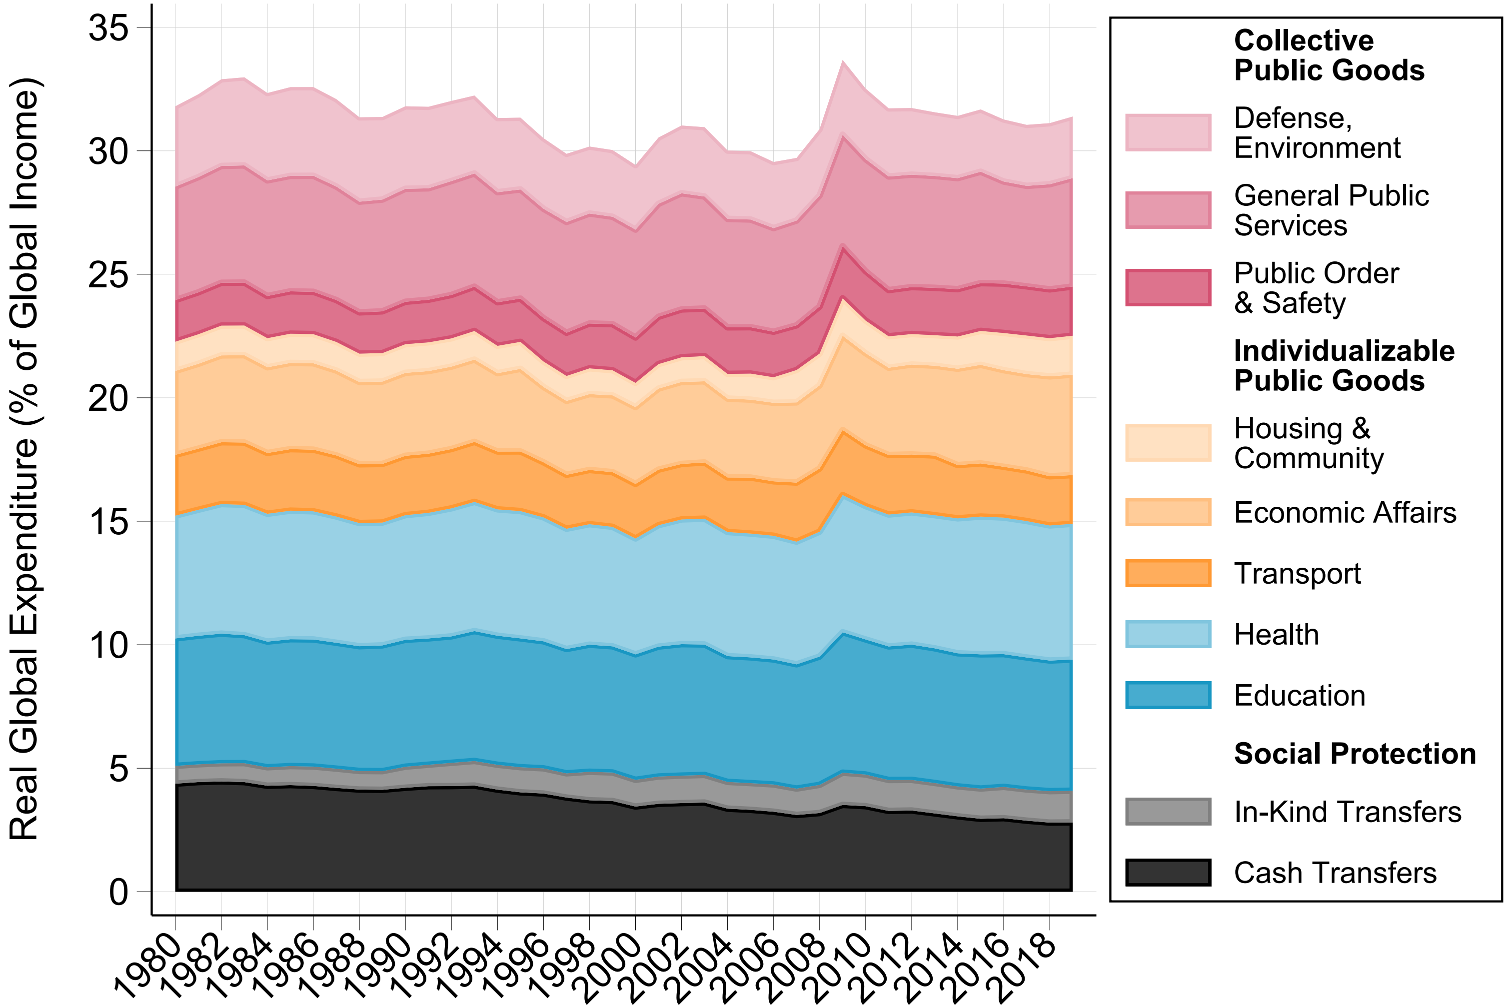

Notes. The figure shows the evolution of real average global general government expenditure, expressed as a share of total global national incomes.

Figure B.26: Government Expenditure on Public Goods Per Capita by Country Income Group, 1980-2019

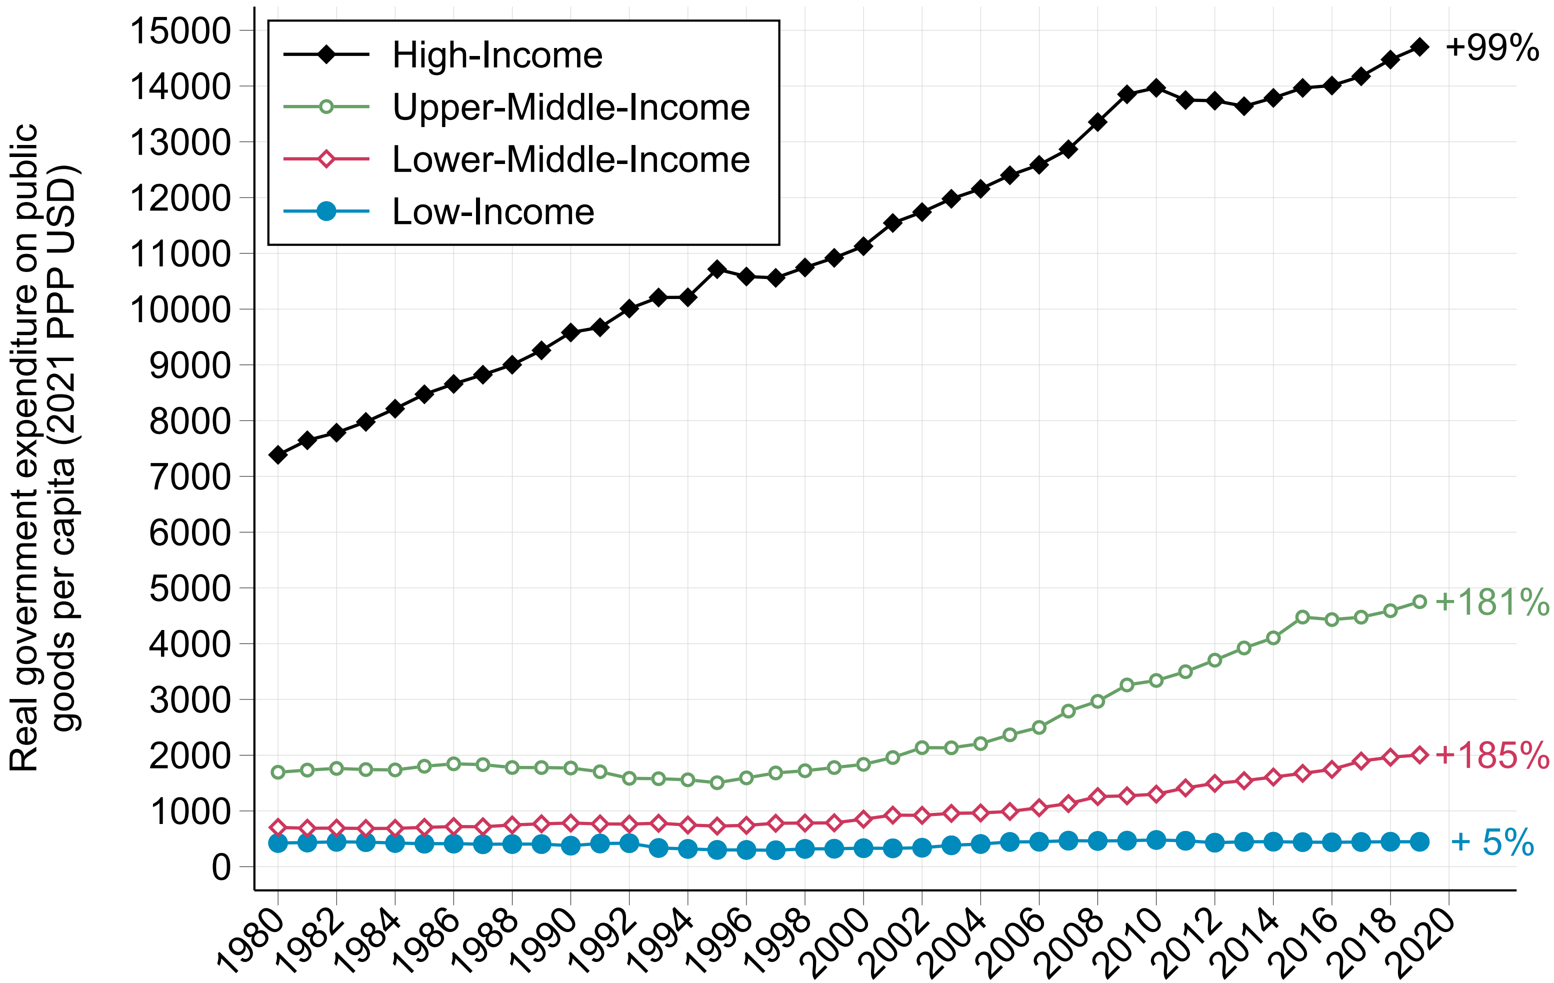

Notes. The figure shows the evolution of average real per capita general government expenditure on public goods by country income group, expressed in 2021 PPP USD. Population-weighted average across all countries in each group.

Figure B.27: Government Expenditure on Public Goods by Country Income Group, 1980-2019 (% of NNI)

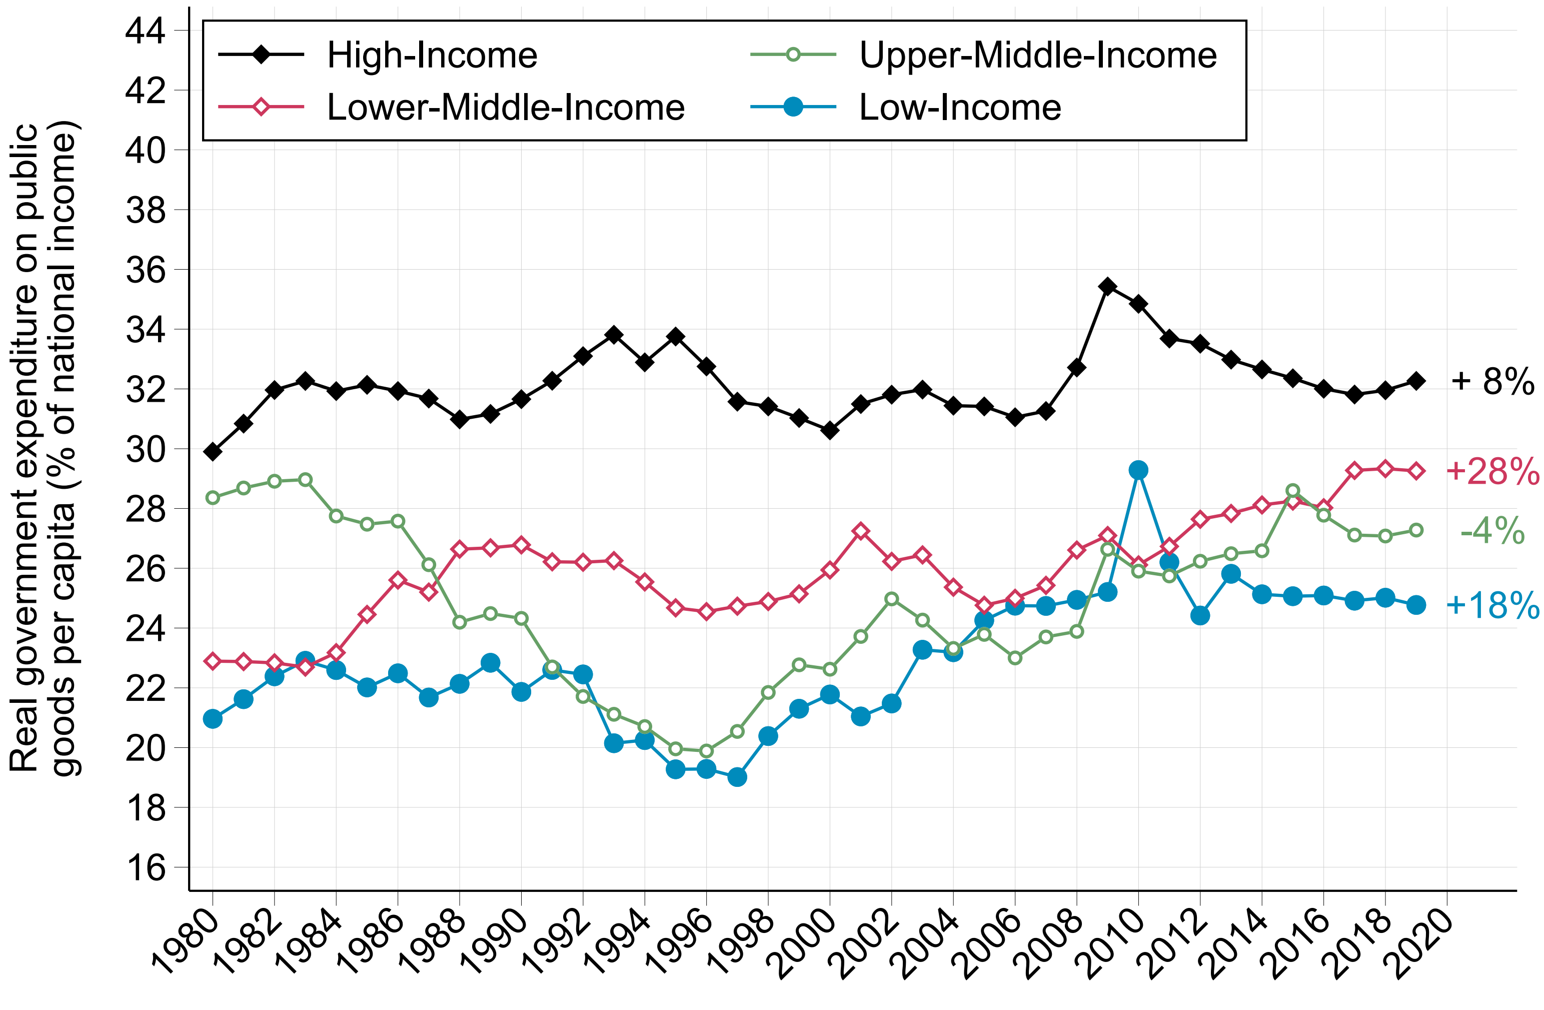

Notes. The figure shows the evolution of average general government expenditure on public goods by country income group, expressed as a share of national income. Population-weighted average across all countries in each group.

Figure B.28: Government Expenditure on Public Goods Per Capita by World Region, 1980-2019

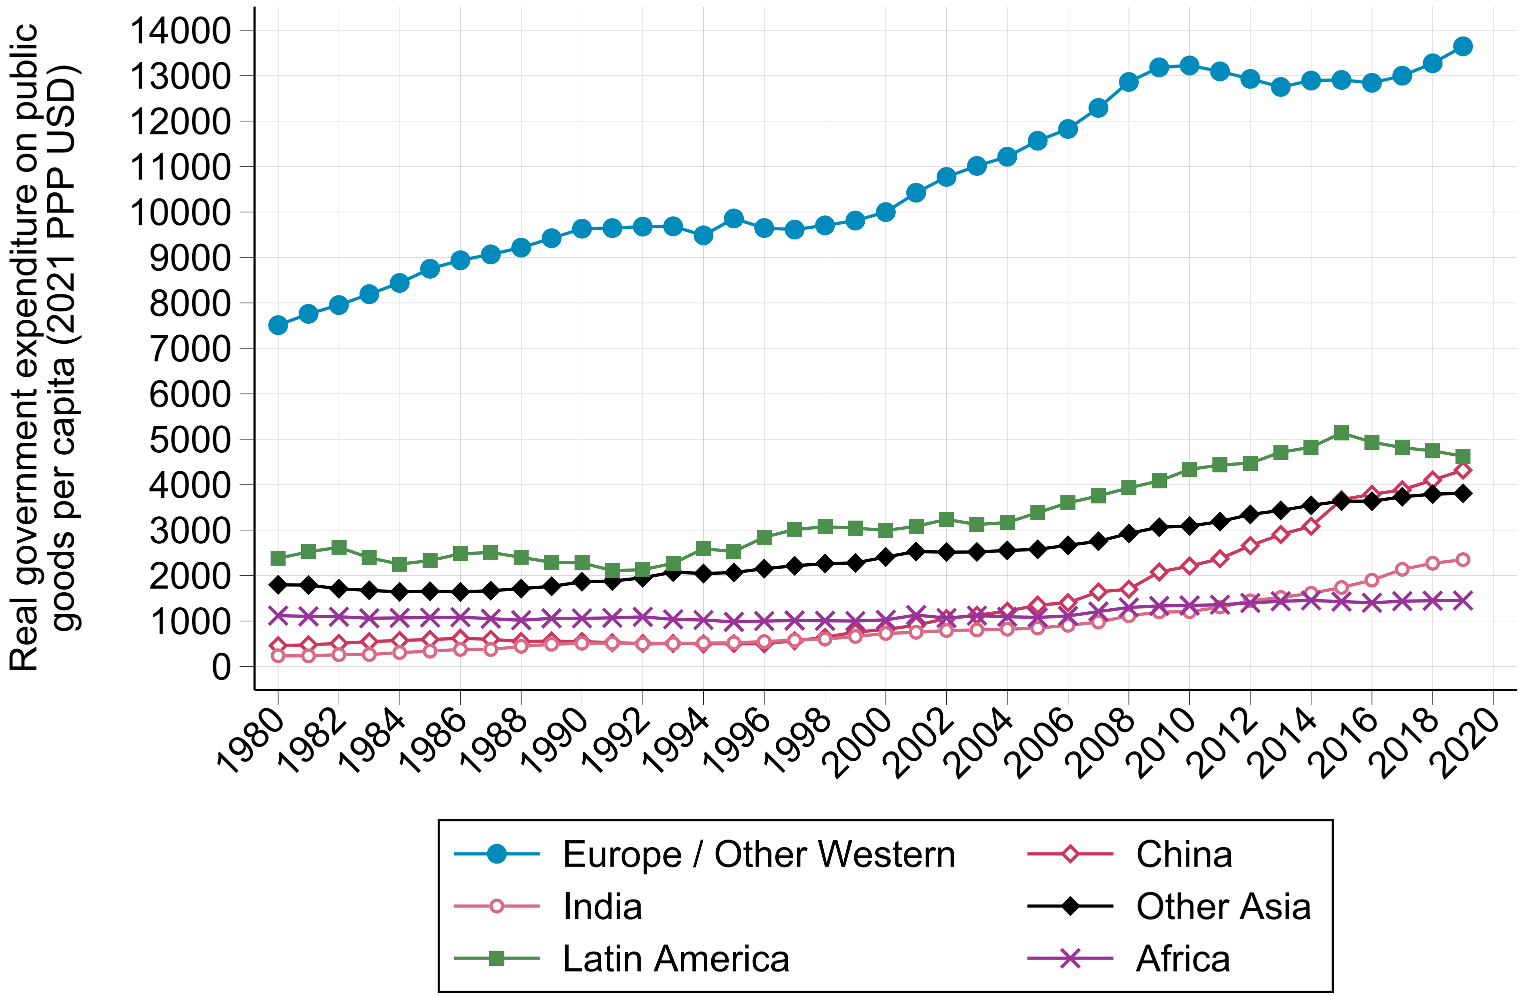

Notes. The figure shows the evolution of real per capita general government expenditure on public goods by world region, expressed in 2021 PPP USD. Other Western countries: United States, Canada, Australia, New Zealand. Population-weighted average across all countries in each group.

Figure B.29: Government Expenditure on Public Goods by World Region, 1980-2019 (% of NNI)

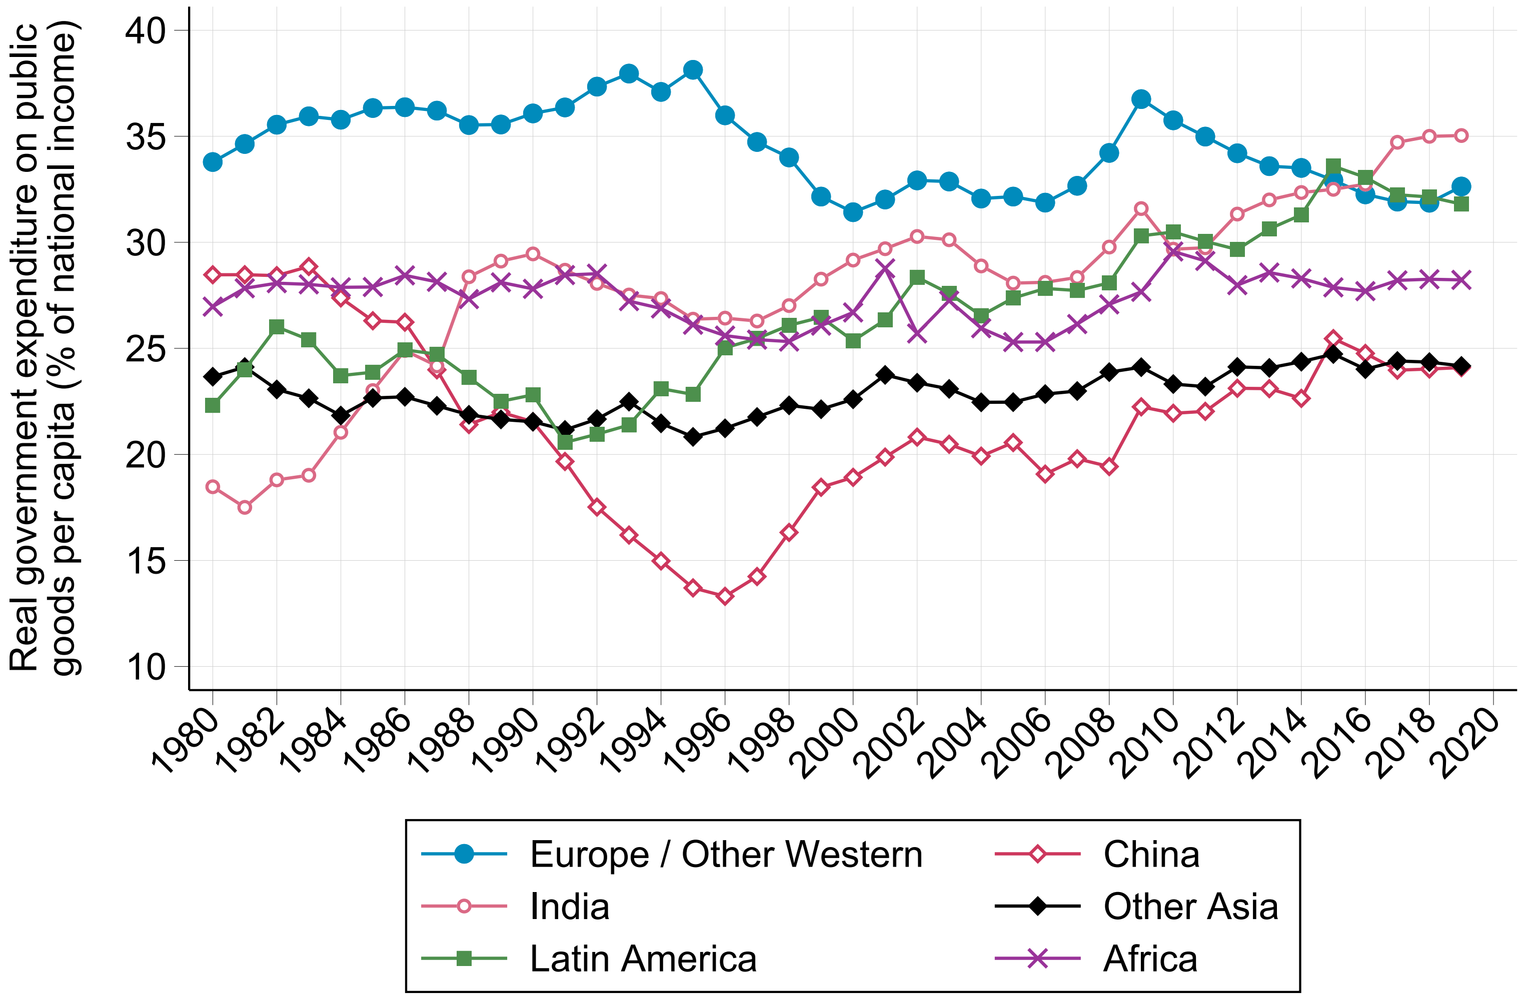

Notes. The figure shows the evolution of general government expenditure on public goods by world region, expressed as a share of national income. Other Western countries: United States, Canada, Australia, New Zealand. Population-weighted average across all countries in each group.

Figure B.30: Government Tax Revenue Per Capita by World Region, 1980-2019

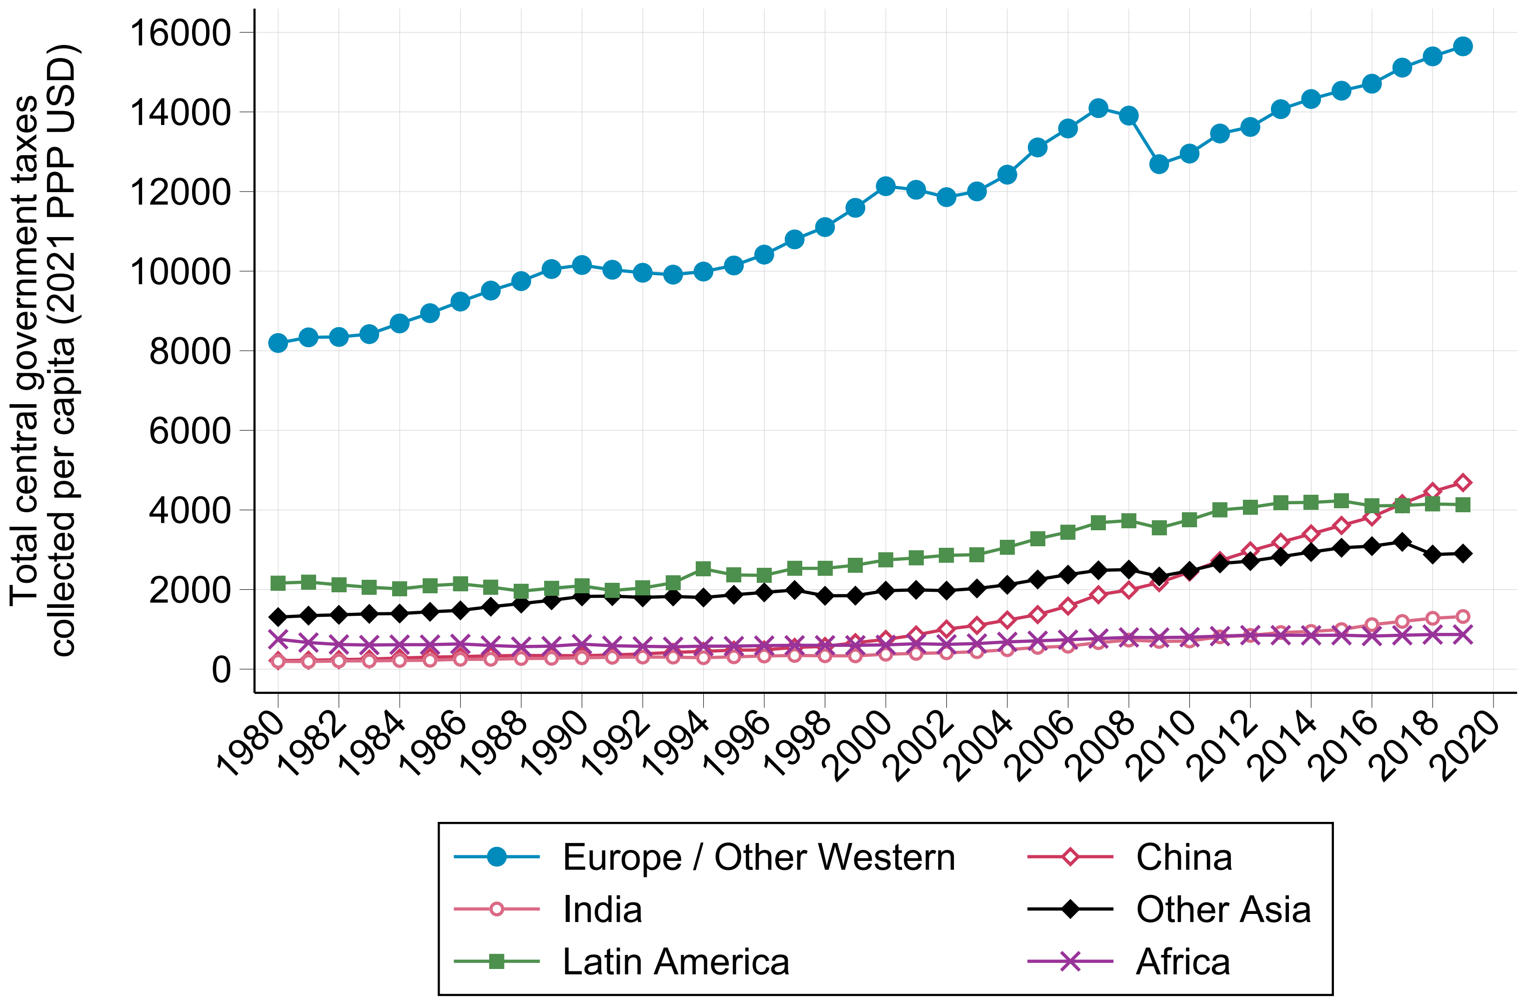

Notes. The figure shows the evolution of real per capita central government tax revenue by world region, expressed in 2021 PPP USD. Other Western countries: United States, Canada, Australia, New Zealand.

Figure B.31: Government Tax Revenue by World Region, 1980-2019 (% of Regional Income)

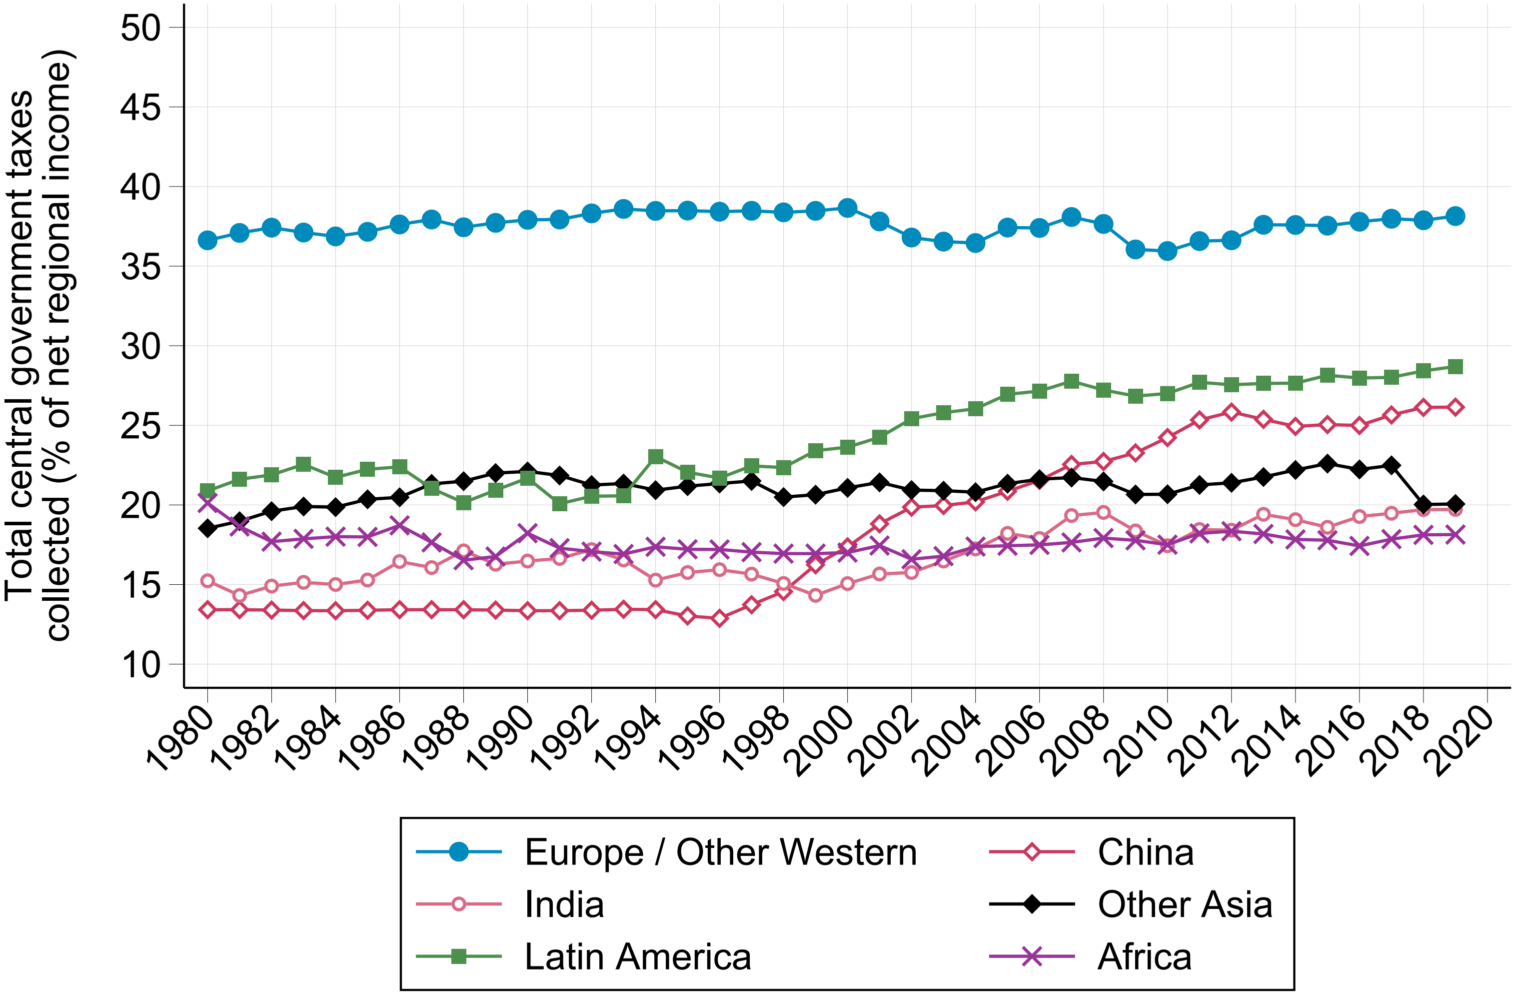

Notes. The figure shows the evolution of central government tax revenue by world region, expressed as a share of total regional income. Other Western countries: United States, Canada, Australia, New Zealand.

B.3.3 Public Sector Productivity: Aggregate Productivity

Figure B.32: Education Expenditure and Expected Human Capital at Age 5

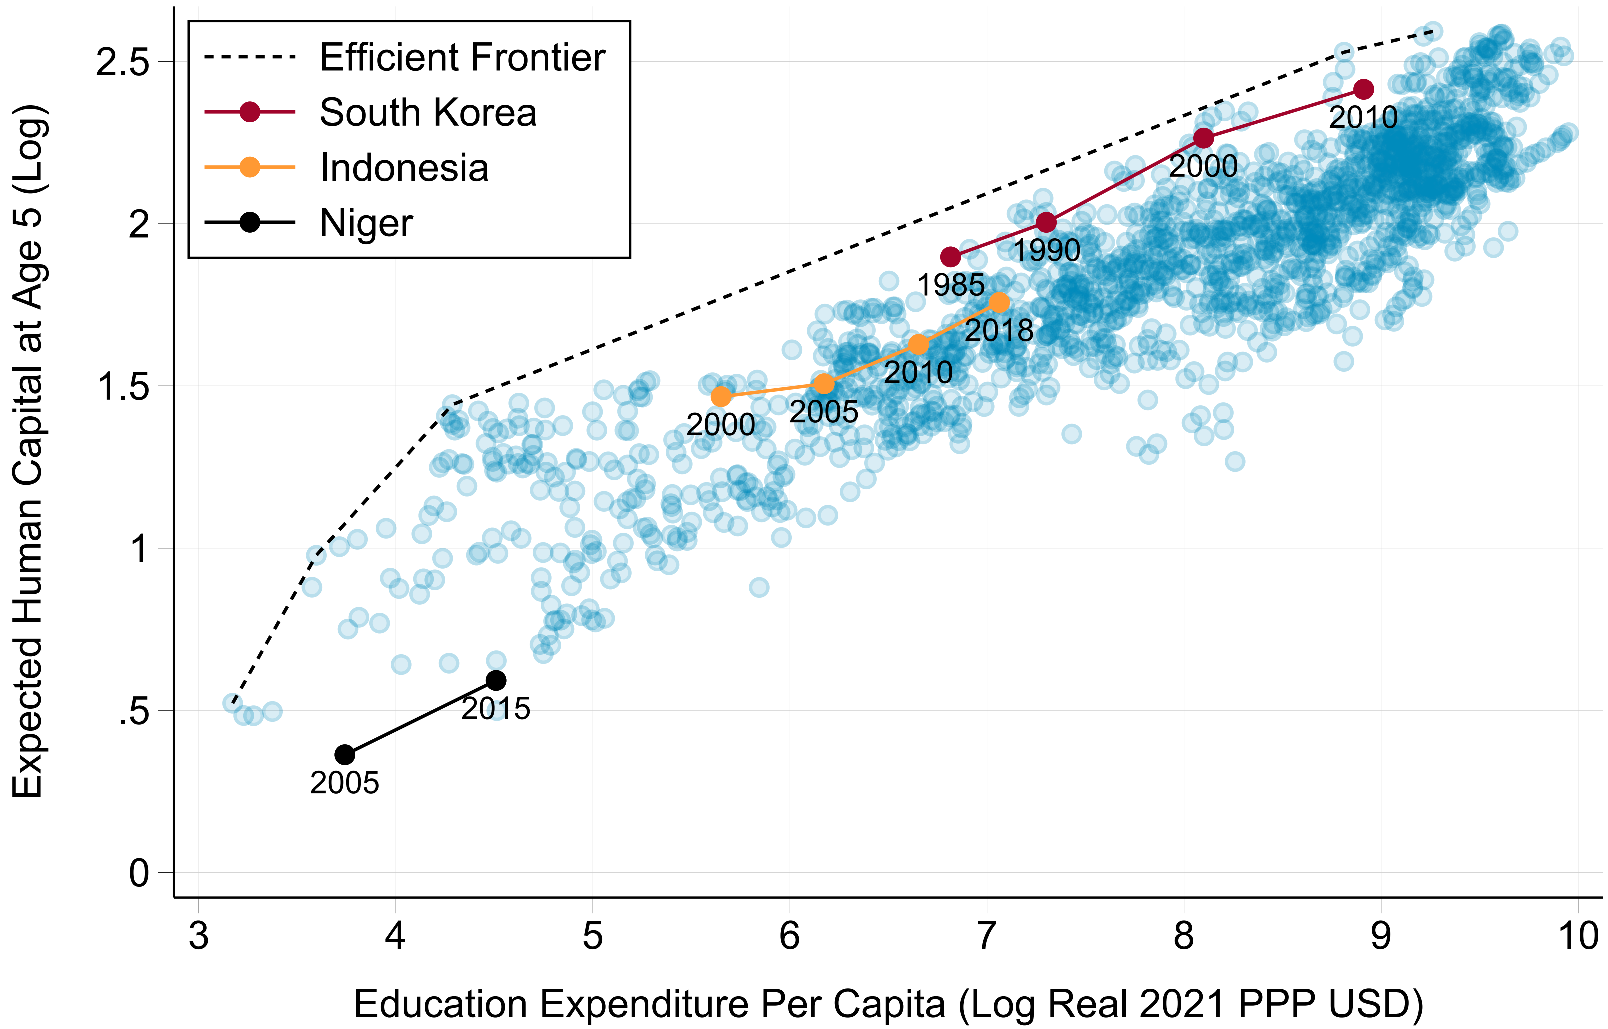

Notes. The unit of observation is the country-year. Data on expected years of schooling from the UNESCO. Data on education expenditure per child come from estimates presented in this paper.

Figure B.33: Health Expenditure and Quality of Healthcare

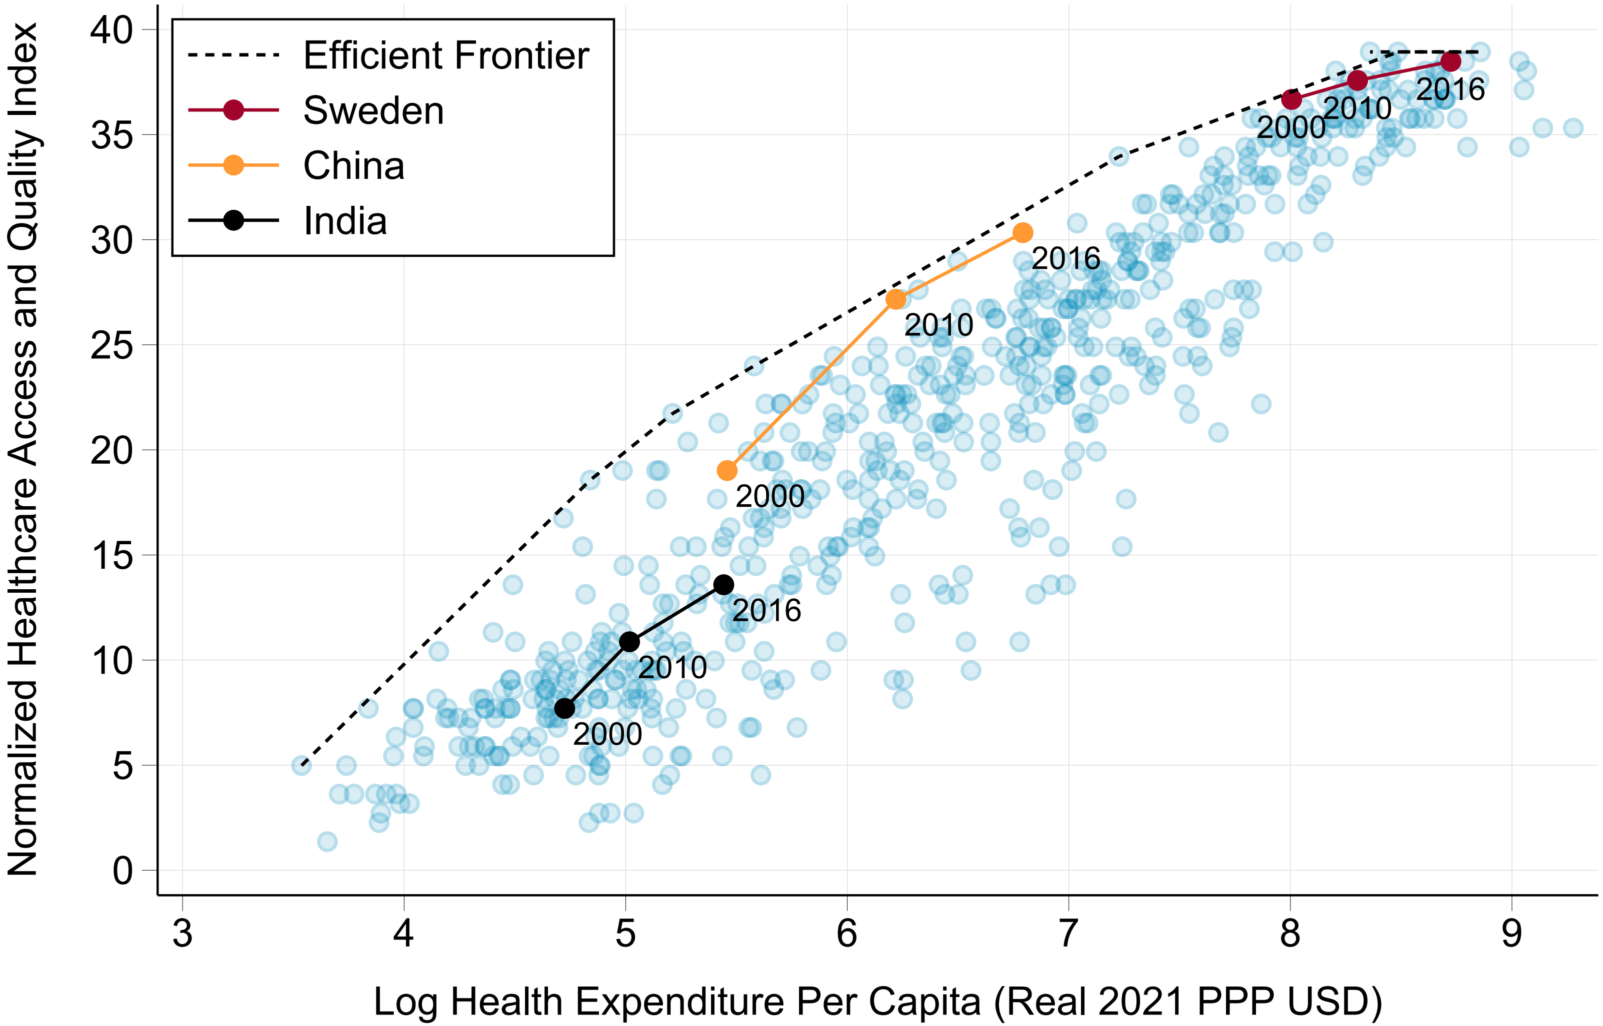

Notes. The unit of observation is the country-year. Data on healthcare access and quality index from the Global Burden of Disease Study. Data on health expenditure from the World Bank.

Figure B.34: Distribution of Aggregate Public Sector Productivity: Education

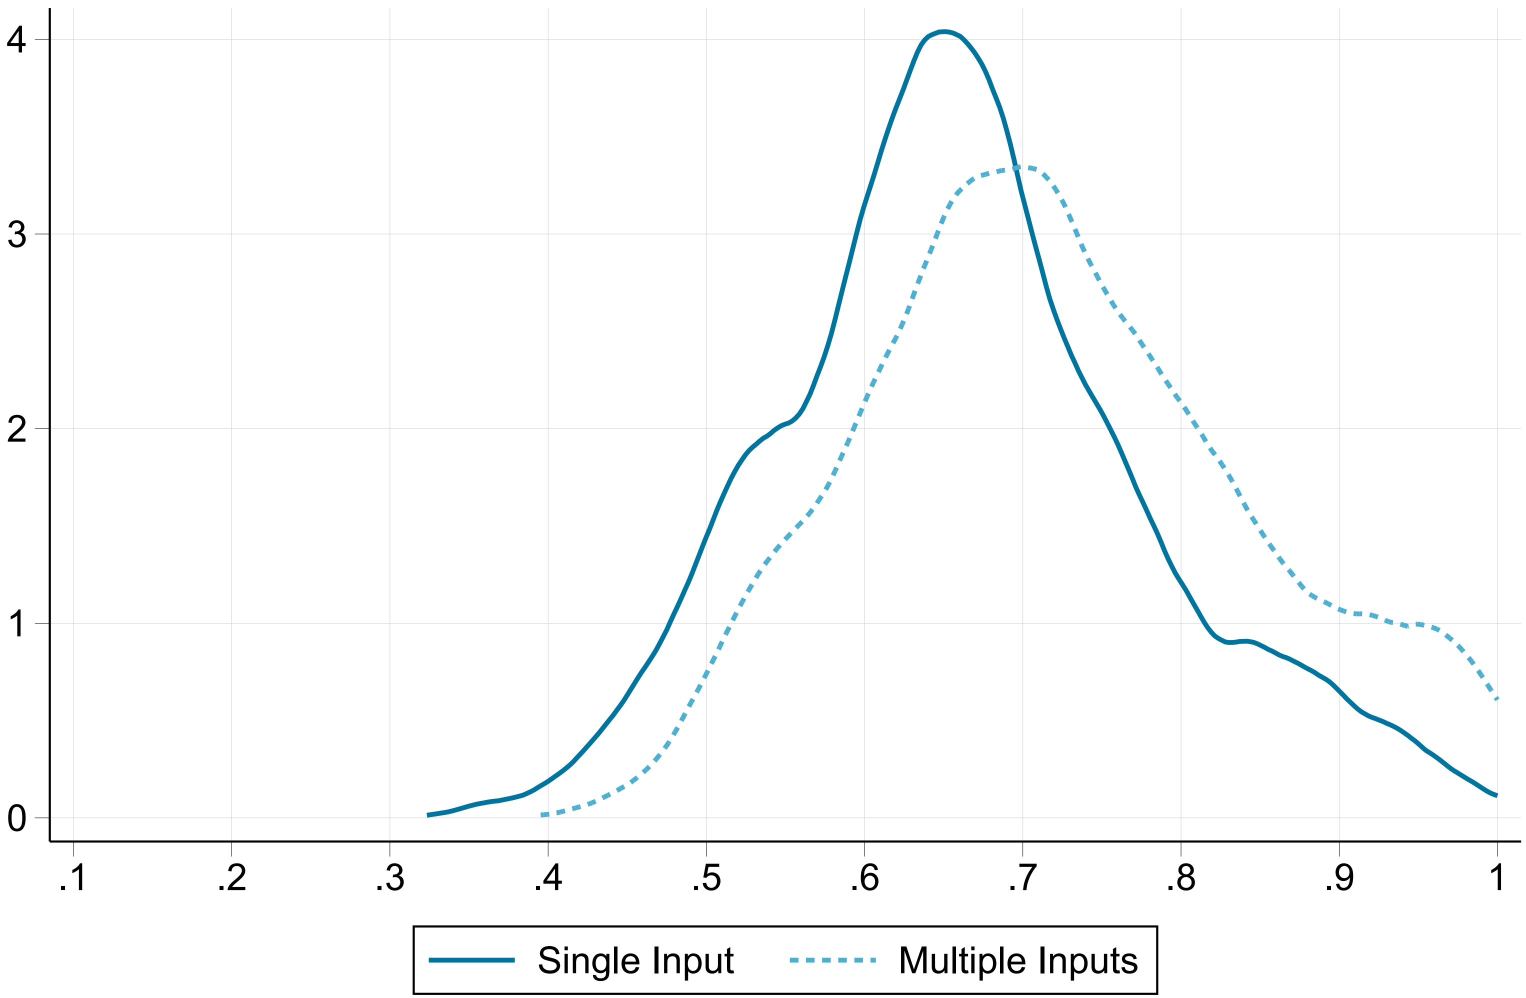

Notes. The figure plots the distribution of aggregate public sector productivity Θ^j for education expenditure, plotted across all country-years in the database, for each of the four models considered.

Figure B.35: Distribution of Aggregate Public Sector Productivity: Health

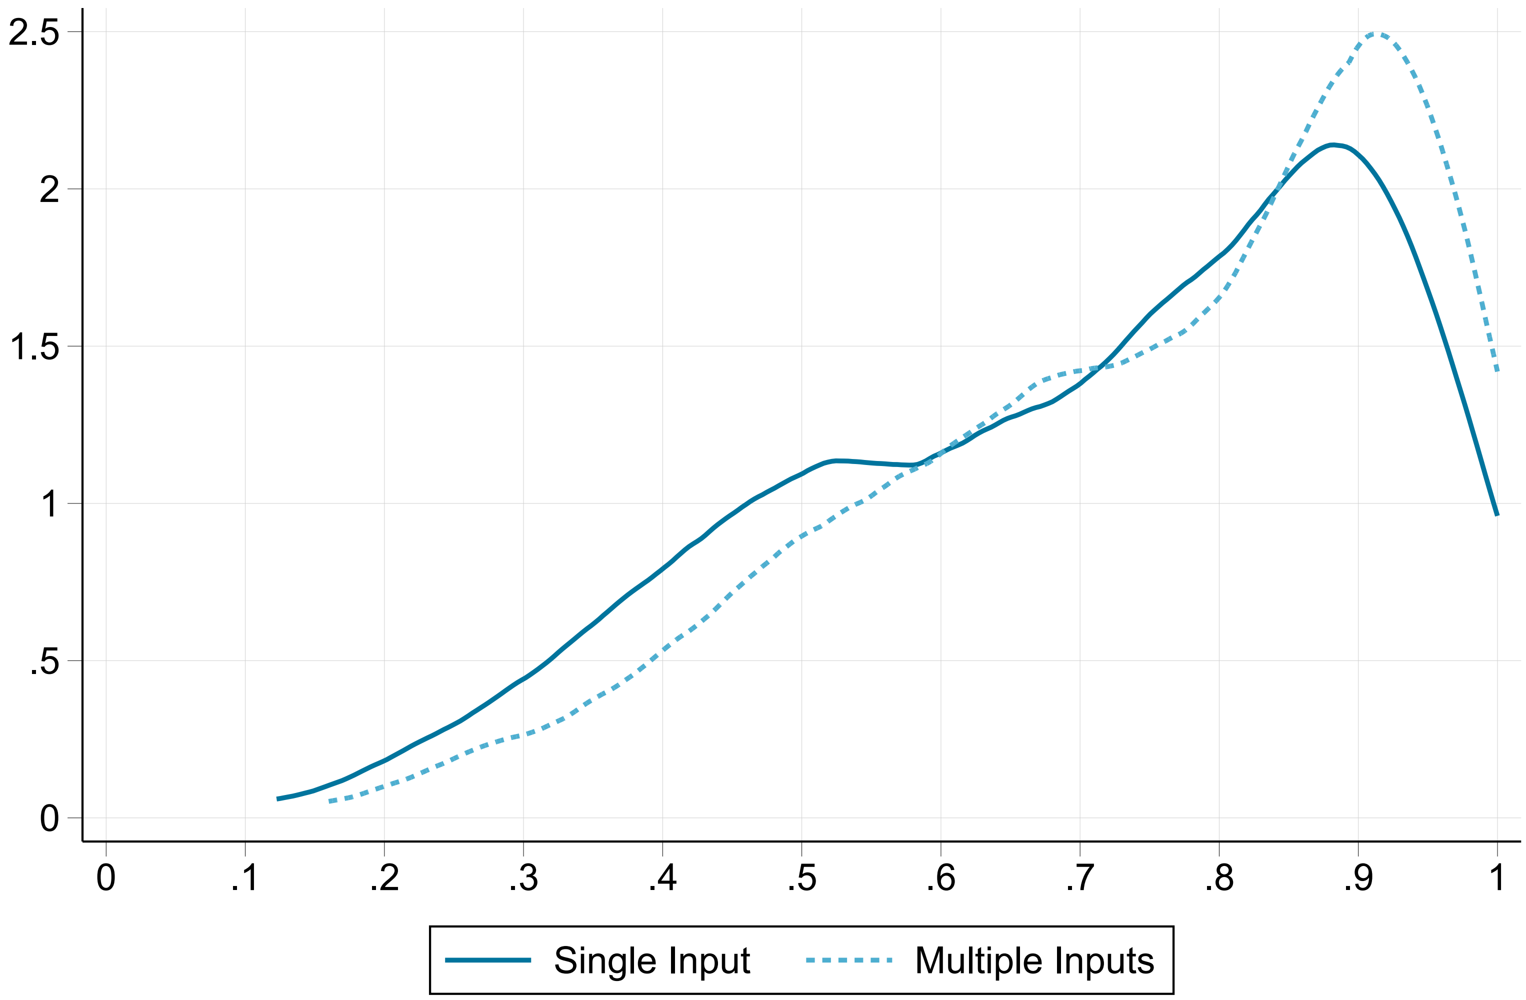

Notes. The figure plots the distribution of aggregate public sector productivity Θ^j for health expenditure, plotted across all country-years in the database, for each of the four models considered.

Table B.6: Summary Statistics on Cross-Country Government Aggregate Productivity Measures

	Mean	SD	Min	Max
Education				
One Input	0.68	0.09	0.32	1.00
Multiple Inputs	0.73	0.09	0.39	1.00
Health				
One Input	0.71	0.19	0.12	1.00
Multiple Inputs	0.74	0.19	0.16	1.00

Notes. Statistics computed over all country-years in the database and weighted by total population.

Table B.7: Correlates of Aggregate Government Productivity

	Education Single Input	Education Multiple Inputs	Health Single Input	Health Multiple Inputs	N
Chong et al. (2014) Mail Efficiency	0.08	0.00	0.29***	0.24***	159
Government Effectiveness	0.30***	0.13*	0.57***	0.49***	177
Control of Corruption	0.17**	0.04	0.43***	0.38***	177
Absence of Corruption	0.07	-0.05	0.27***	0.23***	160
Wastefulness of Government Spending	0.22***	0.14*	0.26***	0.24***	149
Irregular Payments and Bribes	0.24***	0.10	0.46***	0.41***	150
Favoritism in Government Decisions	0.15*	0.03	0.28***	0.22***	151
Transparency of Policymaking	0.20**	0.06	0.34***	0.29***	150
GDP per capita	0.35***	0.22***	0.62***	0.57***	177
Inequality in Public Service Delivery	0.34***	0.35***	0.30***	0.32***	160

Notes. The table reports raw pairwise correlations between the four measures of total technical efficiency and other qualitative indicators of government productivity. Correlations are computed over all countries with available data for each pair of indicators, for the last year available, and weighted by each country's total population. Chong et al. (2014) efficiency corresponds to the average number of days to get the letter back. GDP per capita data come from the World Inequality Database. Inequality in public service delivery is measured as the quality of public services received by the bottom quintile relative to the overall population ($q^j(Q1)$), estimated from the Gallup World Poll over the 2009-2021 period. Data on other indicators come from the World Bank. * p¡0.10, ** p¡0.05, *** p¡0.01.

Table B.8: Correlations Between Measures of Government Productivity

	Educ1	Educ2	Heal1	Heal2
Educ1		0.94***	0.57***	0.60***
Educ2	0.94***		0.54***	0.59***
Heal1	0.57***	0.54***		0.97***
Heal2	0.60***	0.59***	0.97***	

Notes. Correlations are computed over all countries with available data for each pair of indicators, for the last year available, and weighted by each country's total population. Educ: education; Heal: health. Numbers correspond to models with single (1) or multiple (2) inputs.

B.3.4 Public Sector Productivity: Heterogeneous Productivity

Figure B.36: Average Heterogeneous Productivity Profiles by Function, World

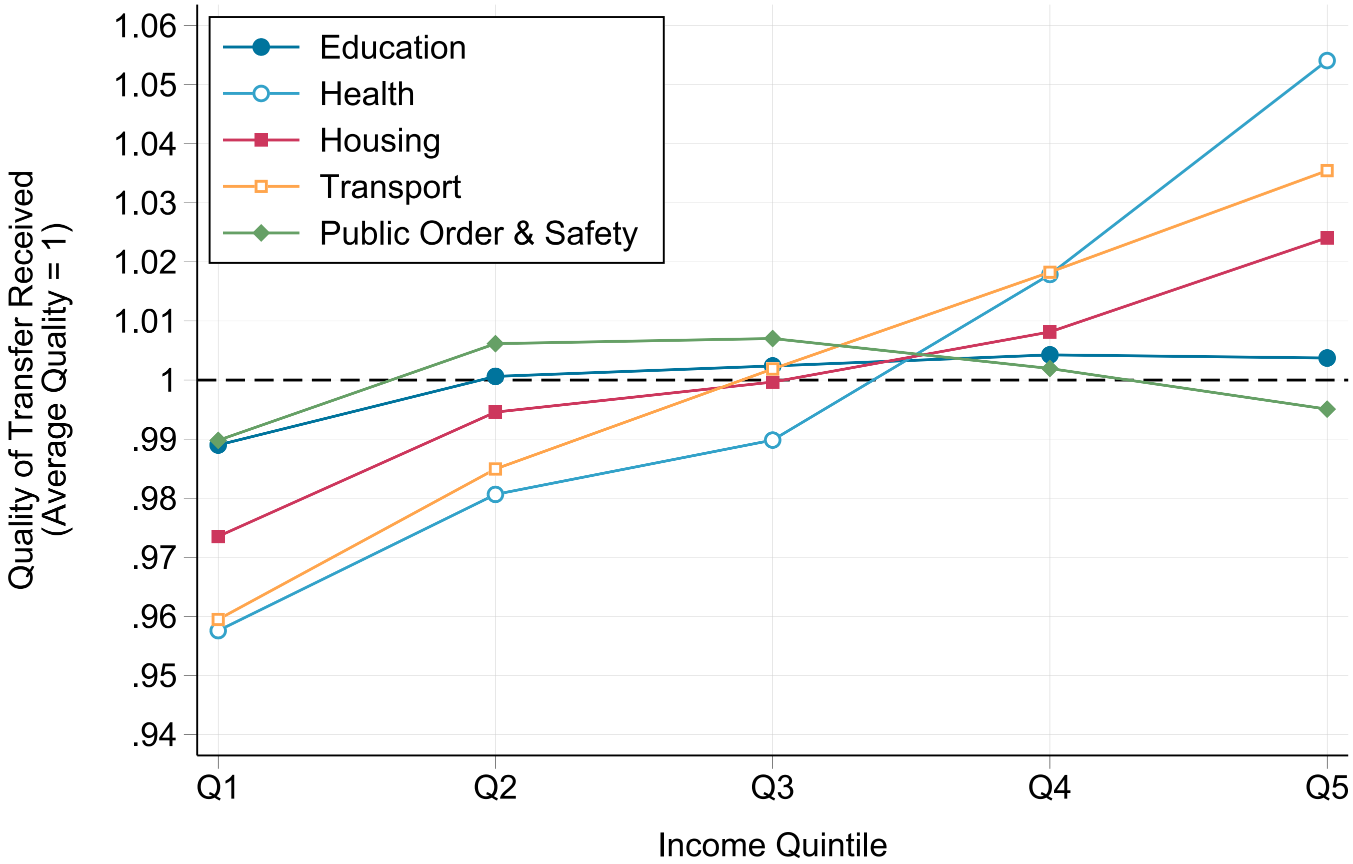

Notes. Author's computations using Gallup World Poll data. The figure represents the average of heterogeneous productivity profiles $q^j(m_i)$ applied to correct in-kind transfers received by income quintile, computed over all countries over the entire 2009-2021 period. Numbers correspond to the ratio of the quality of the transfer received to average quality. Quality is measured as the share of respondents who declare being satisfied with public services in the city or area where they live, for the following services: public transportation systems, roads and highways, the educational system or the schools, the quality of water, and the availability of quality health care. The quality of police services is measured as the share of respondents who declare having confidence in the local police force.

Figure B.37: Distribution of Heterogeneous Productivity Scores by Function

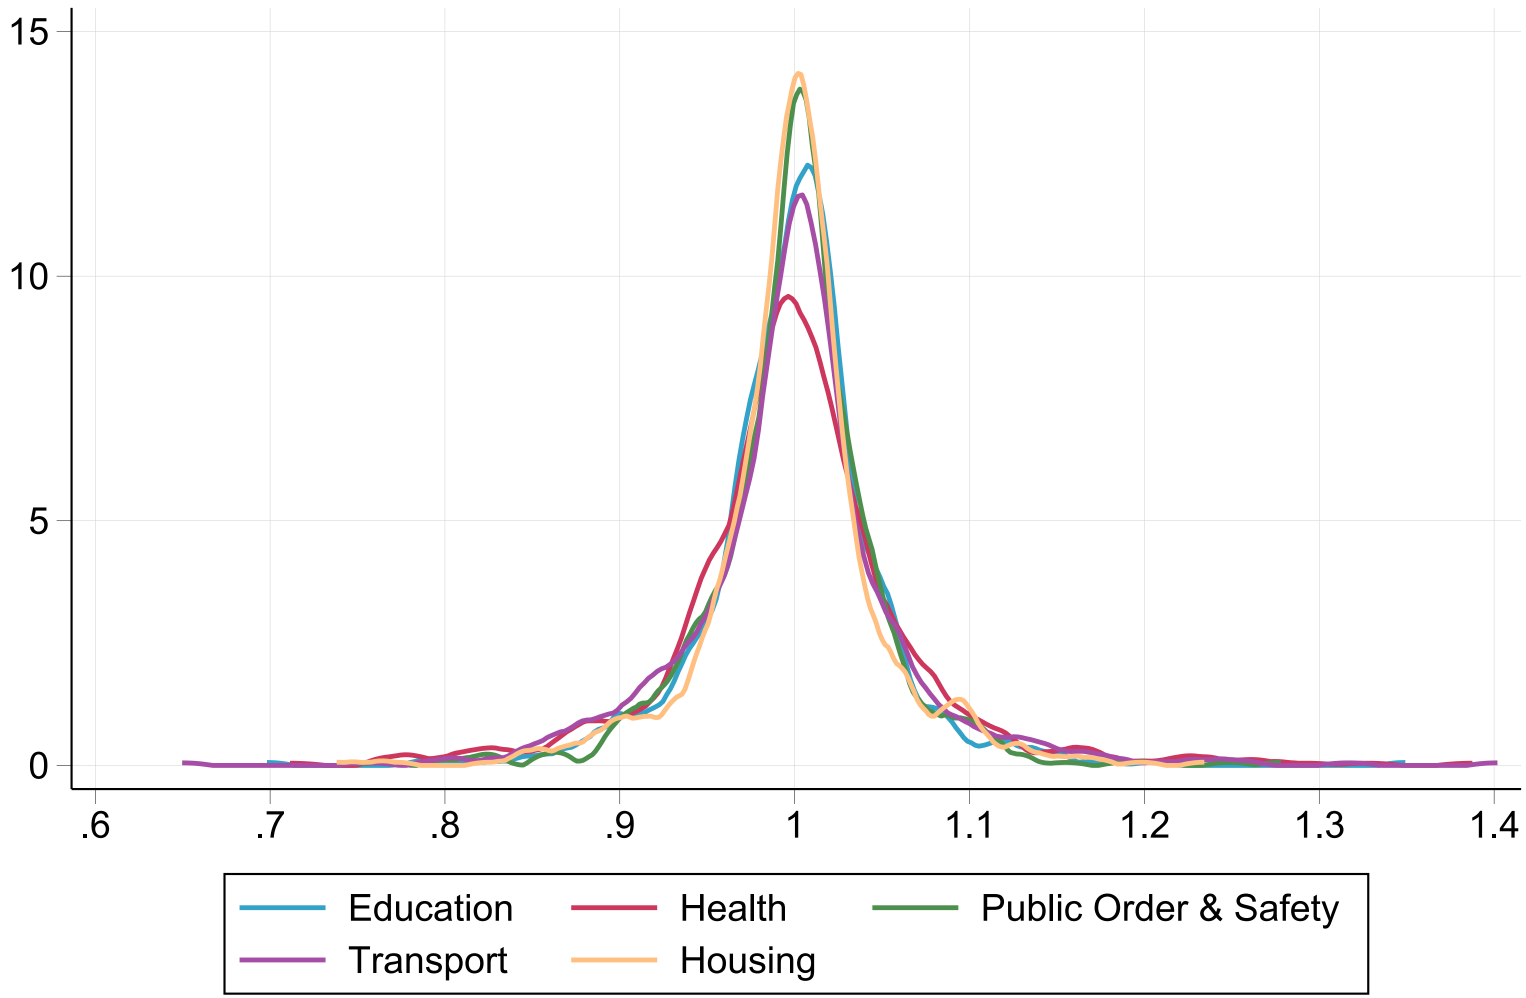

Notes. The figure represents the distribution of heterogeneous productivity scores by function, estimated from the Gallup World Poll data, across all countries with available data, for the bottom 20%. Figures correspond to the ratio of the quality of the transfer received by the bottom 20% to average quality. Quality is measured as the share of respondents who declare being satisfied with public services in the city or area where they live, for the following services: public transportation systems, roads and highways, the educational system or the schools, the quality of water, and the availability of quality health care. The quality of police services is measured as the share of respondents who declare having confidence in the local police force.

Table B.9: Indicators of Heterogeneous Public Service Delivery by Income Quintile in South Africa

	Q1	Q2	Q3	Q4	Q5	$q^j(Q1)$	Source
Subjective Indicators (% Positively Rating)							
Local public school	69%	69%	69%	68%	69%	1.01***	Census
Local public clinic	46%	45%	46%	46%	50%	0.98***	Census
Local public hospital	47%	47%	47%	48%	51%	0.97***	Census
Local police services	43%	43%	44%	45%	48%	0.97***	Census
Electricity supply	63%	63%	63%	64%	67%	0.99***	Census
Water supply	50%	54%	58%	62%	68%	0.85***	Census
Refuse removal services	49%	54%	57%	60%	66%	0.85***	Census
Sanitation services	52%	56%	59%	64%	74%	0.85***	Census
Government-subsidized dwelling	48%	49%	50%	51%	53%	0.96***	Census
Police response to reported crime	52%	53%	52%	53%	56%	0.98	VCS
Objective Indicators							
School teacher mathematics test success rate	38%	40%	40%	47%	67%	0.82***	SACMEQ
Share of reported crimes leading to arrest	24%	20%	21%	18%	20%	1.15	VCS
Asked to pay a bribe in past 12 months	5%	9%	8%	11%	15%	1.78***	VCS
Water interruption in past 3 months	19%	19%	17%	16%	14%	0.90***	Census
Electricity interruption in past 3 months	32%	28%	25%	21%	16%	0.76***	Census
Value of subsidized dwelling (R 1,000)	177	178	267	308	305	0.72***	GHS
Distance to Nearest Public Services (km)							
Primary school	1.5	1.5	1.6	1.8	2.0	1.12***	LCS
Secondary school	2.9	2.8	2.6	2.4	2.8	0.93***	LCS
Clinic	4.7	4.5	3.8	3.5	3.8	0.86***	LCS
Hospital	13.2	12.6	10.2	8.6	7.3	0.79***	LCS
Police station	8.6	8.1	6.1	4.9	4.6	0.75***	LCS
Public transport	1.1	1.0	1.1	1.0	1.3	1.04*	LCS

Notes. The table reports estimates of heterogeneous government productivity by income group, based on a number of subjective and objective indicators of public service delivery. Q1 to Q5 refer to income quintiles. $q^j(Q_1)$ is the corresponding measure of the relative quality of services received by the bottom quintile, equal to the ratio of the value of the indicator for Q1 to the overall sample mean (or its inverse when the scale of the variable is inverted). Statistical significance stars correspond to a regression of the indicator of interest on a dummy taking one if the individual belongs to the bottom quintile. *** p<0.01, ** p<0.05, * p<0.1. Census: 2016 national census. GHS: 2019 General Household Survey. VCS: 2017 Victims of Crime Survey. LCS: 2014-2015 Living Conditions Survey. SACMEQ: The Southern and Eastern Africa Consortium for Monitoring Educational Quality (estimates from Venkat and Spaull, 2015).

B.3.5 Maps

Figure B.38: Country Income Groups

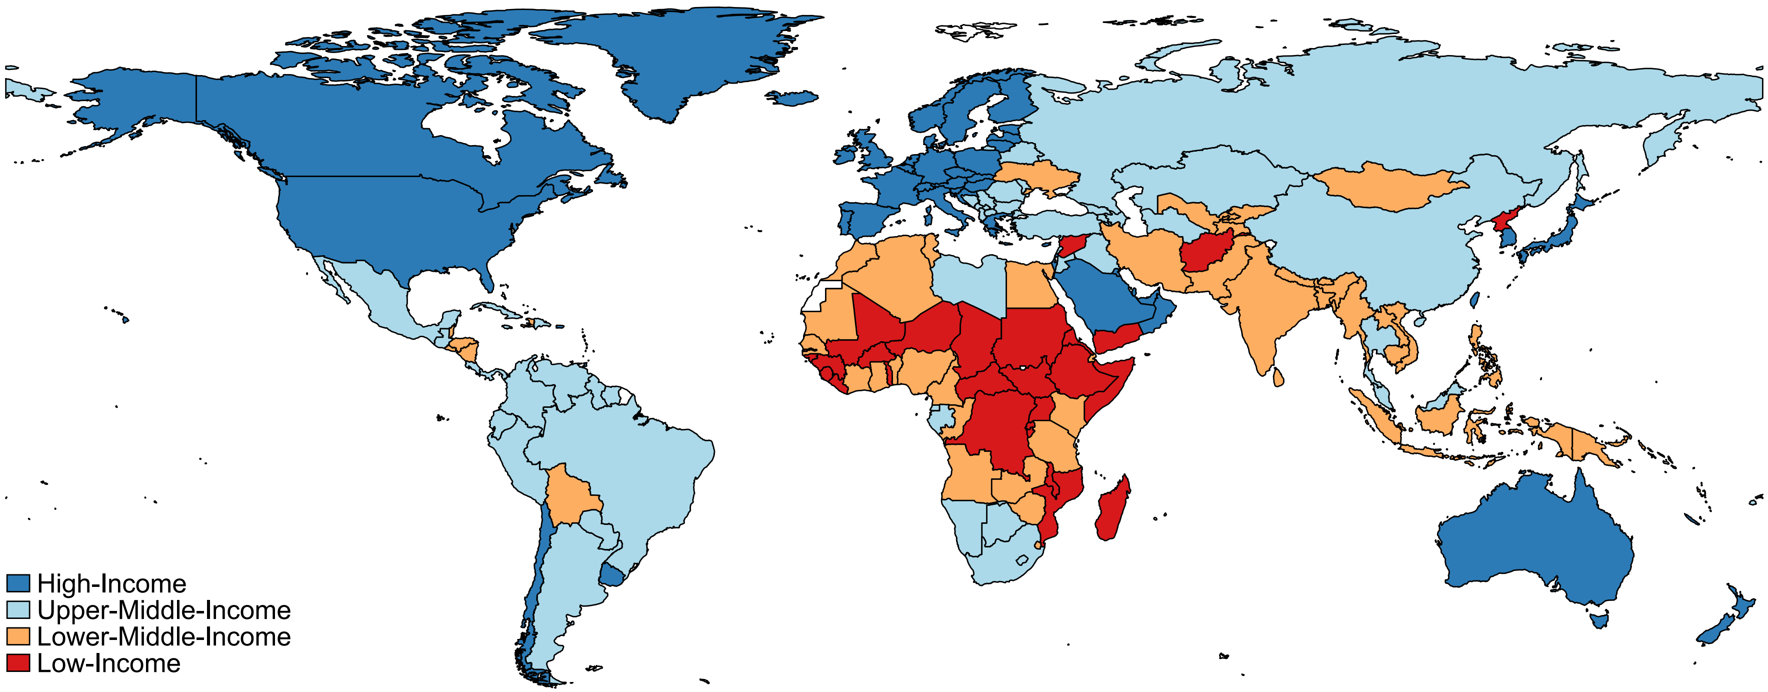

Notes. Authors' elaboration based on World Bank classification of country income groups.

Figure B.39: General Government Expenditure, 2019 (% of National Income)

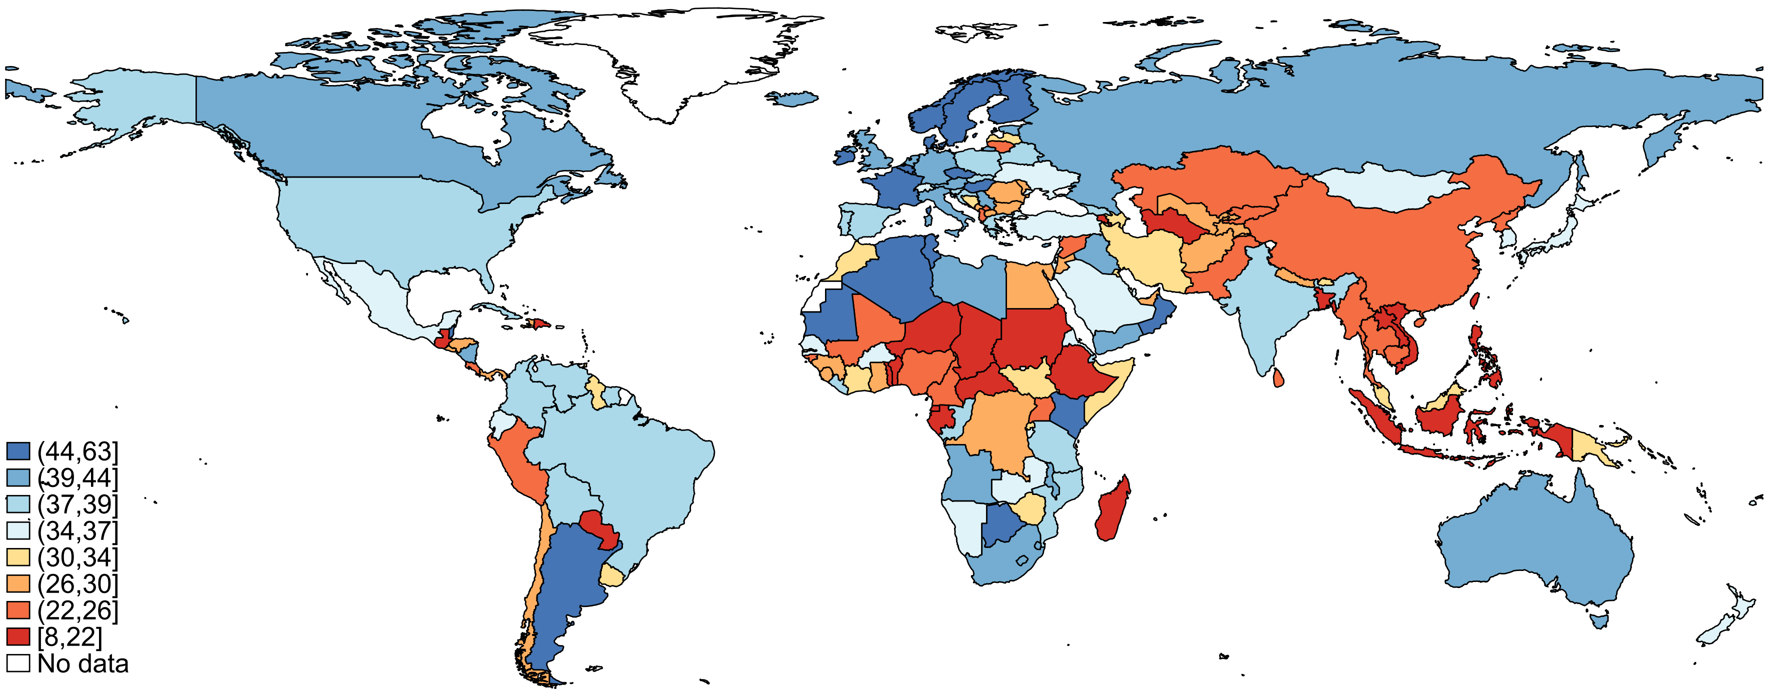

Notes. Authors' computations using national budget data.

Figure B.40: Change in General Government Expenditure, 1980-2019 (% of National Income)

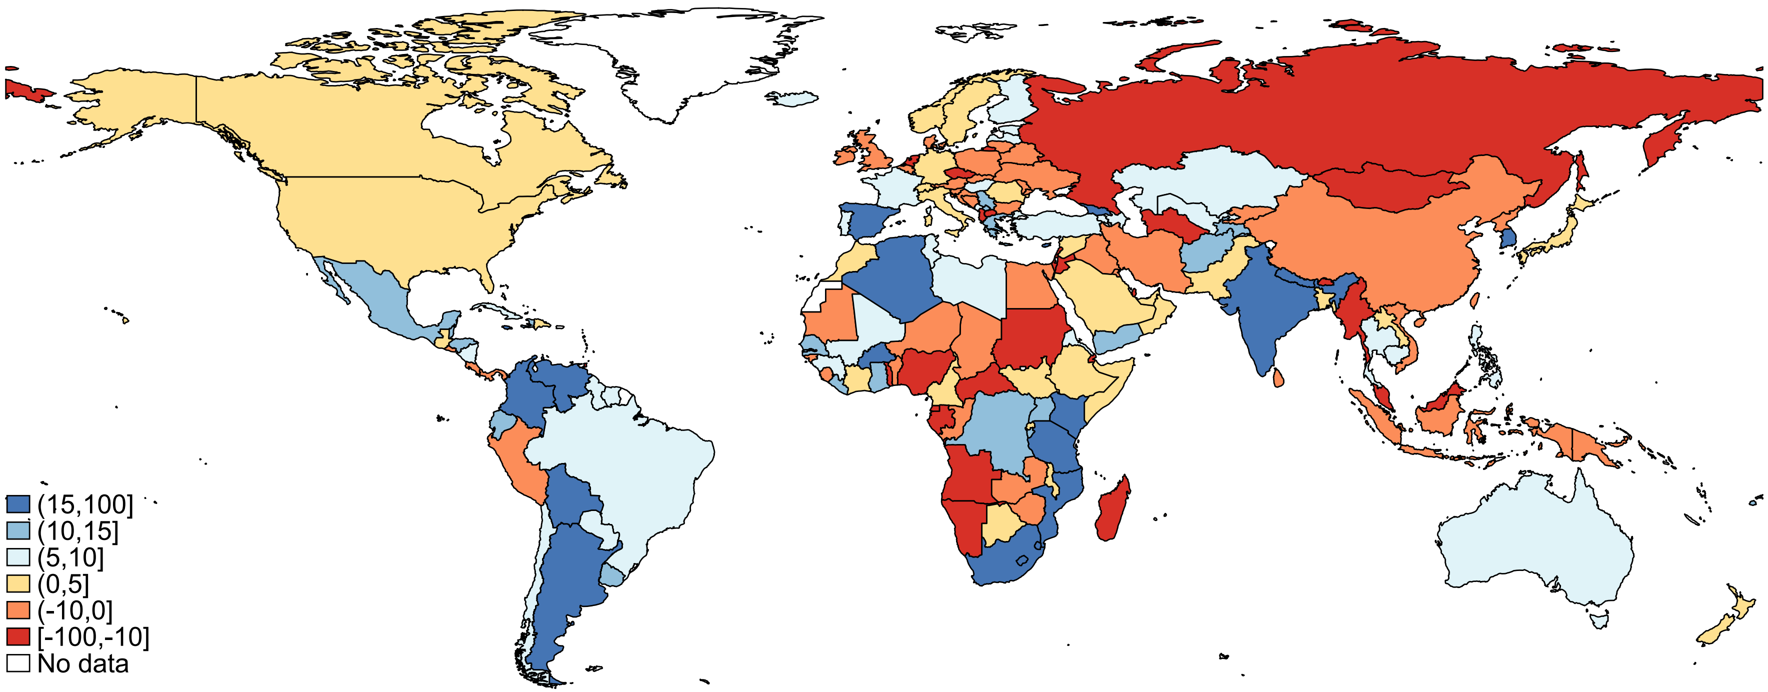

Notes. Authors' computations using national budget data.

Figure B.41: General Government Expenditure on Education and Health, 2019 (% of Total Expenditure)

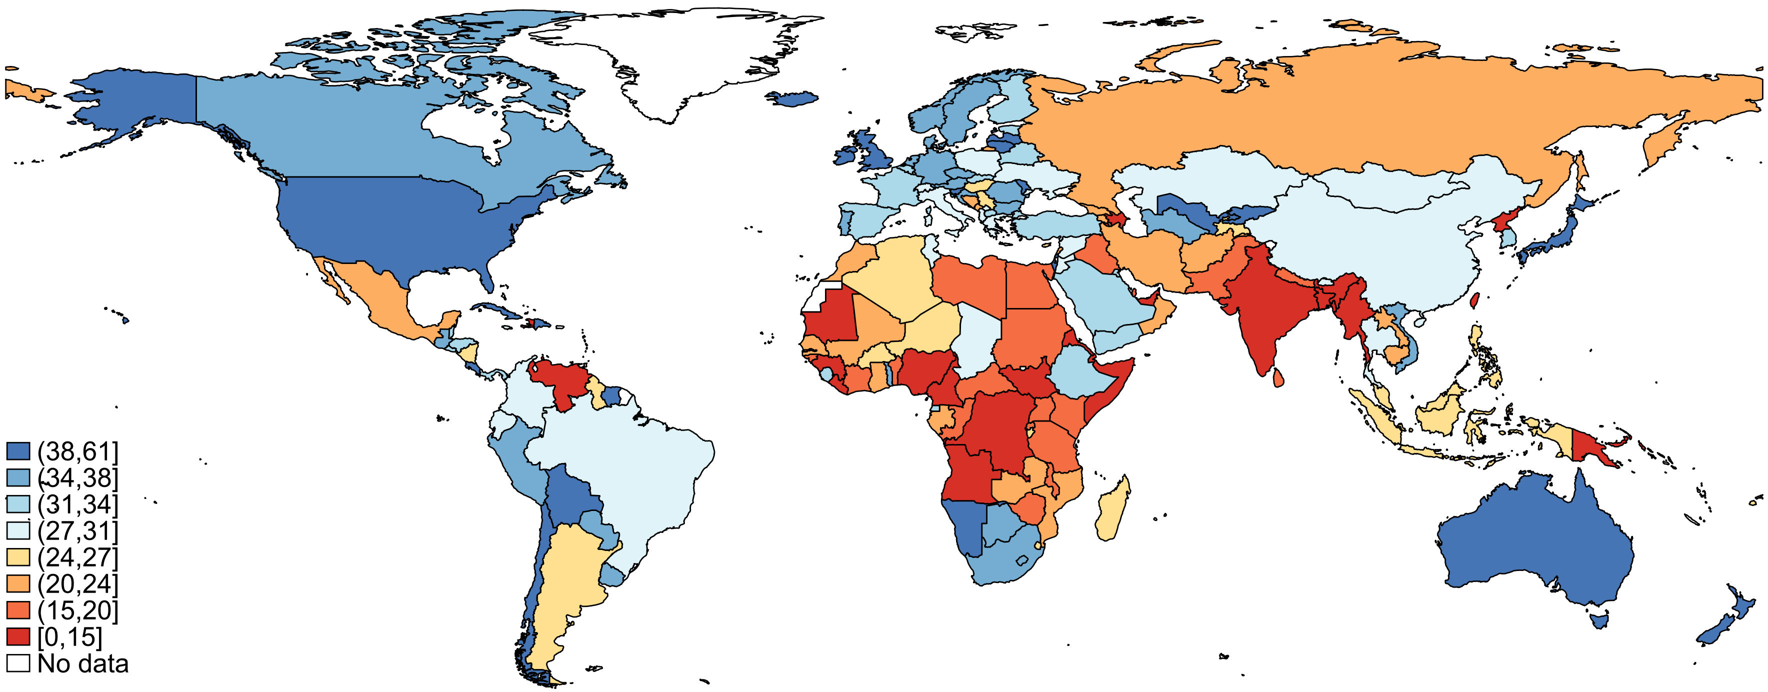

Notes. Authors' computations using national budget data.

Figure B.42: Share of Expenditure on Public Goods Received by the Bottom 50%

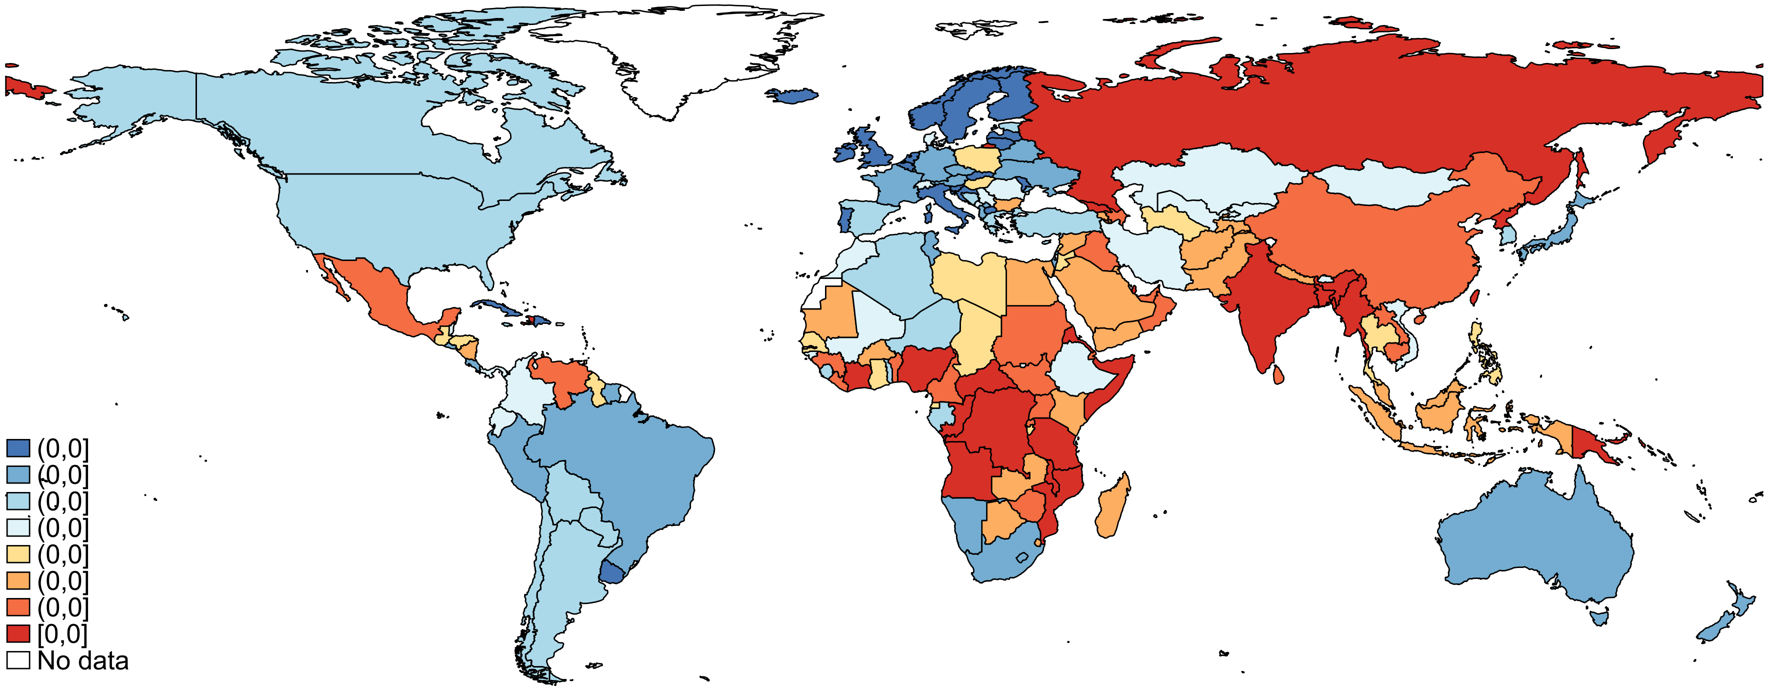

Notes. The map represents the share of total government expenditure on public goods received by the bottom 50% in each country.

Figure B.43: Aggregate Public Education Productivity Around the World, 2019

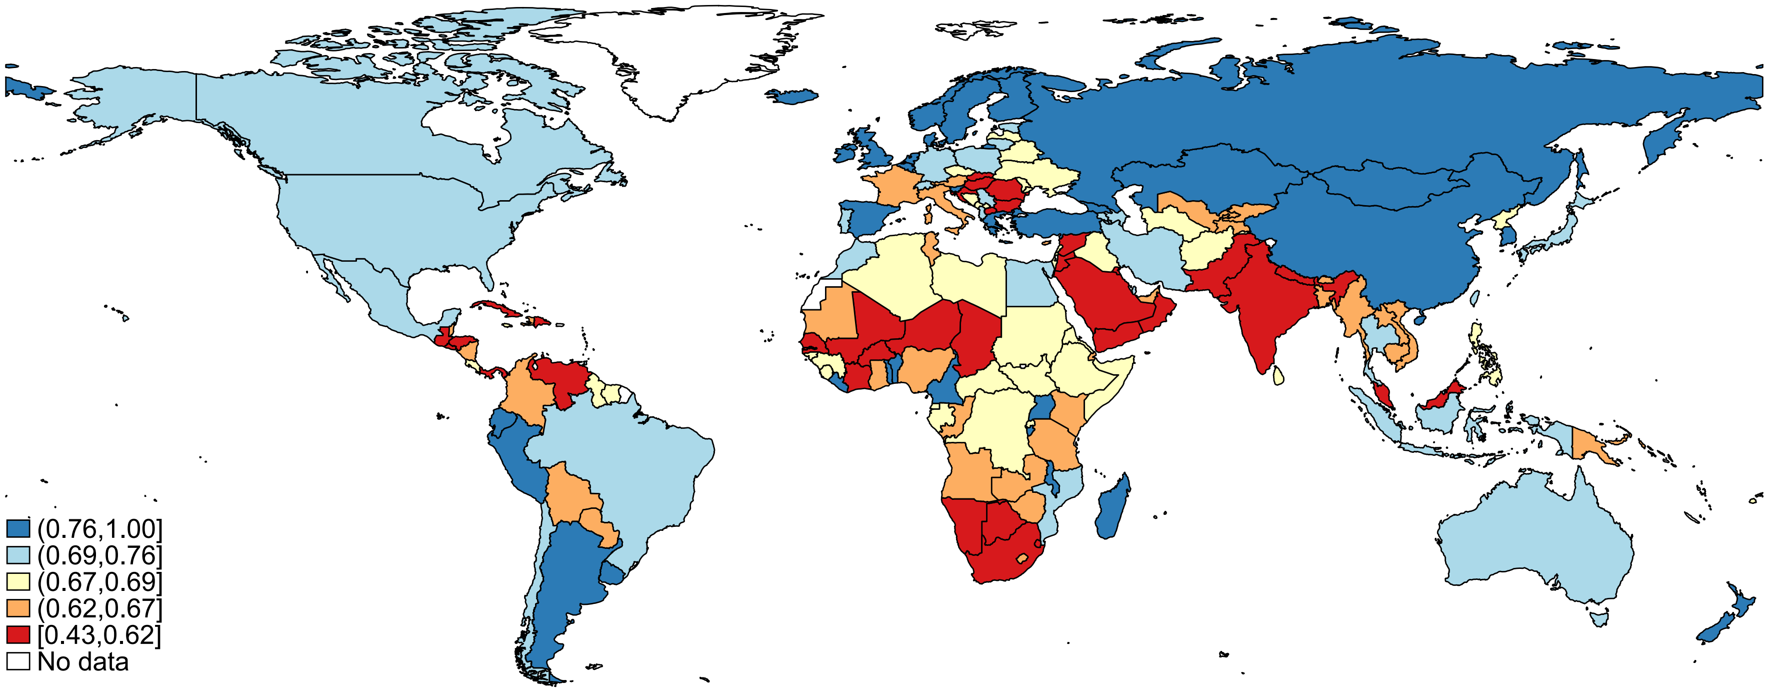

Notes. The map represents estimates of aggregate public education productivity Θ^j in 2019, estimated using public education spending as the only input.

Figure B.44: Aggregate Public Healthcare Productivity Around the World, 2019

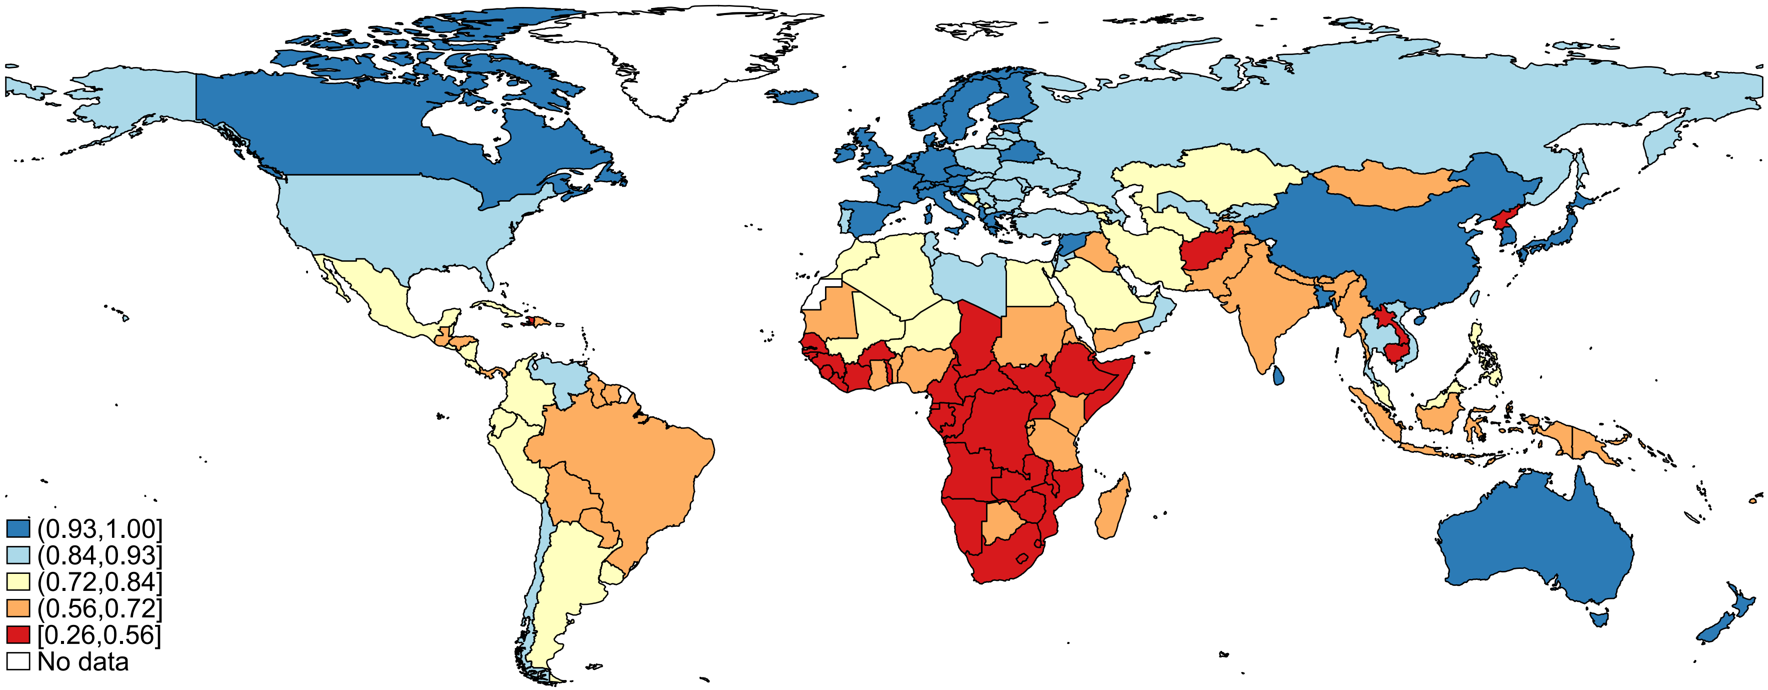

Notes. The map represents estimates of aggregate public healthcare productivity Θ^j in 2019, estimated using public health spending as the only input.

Figure B.45: Change in Aggregate Public Education Productivity Around the World, 1980-2019

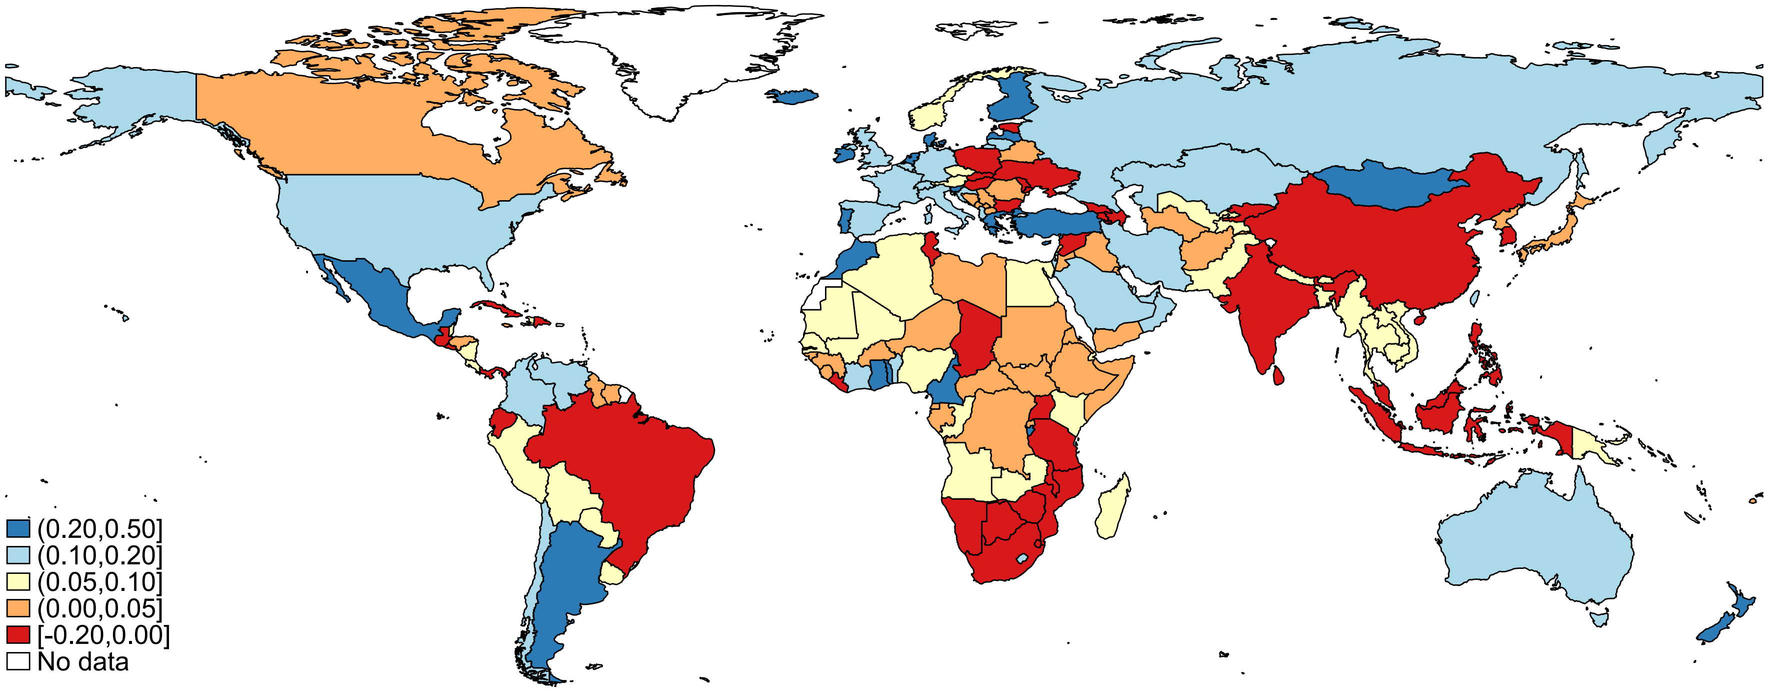

Notes. The map represents the percentage point change in aggregate public education productivity Θ^j between 1980 and 2019 around the world, estimated using a single-input estimate for each function of government.

Figure B.46: Change in Aggregate Public Healthcare Productivity Around the World, 1980-2019

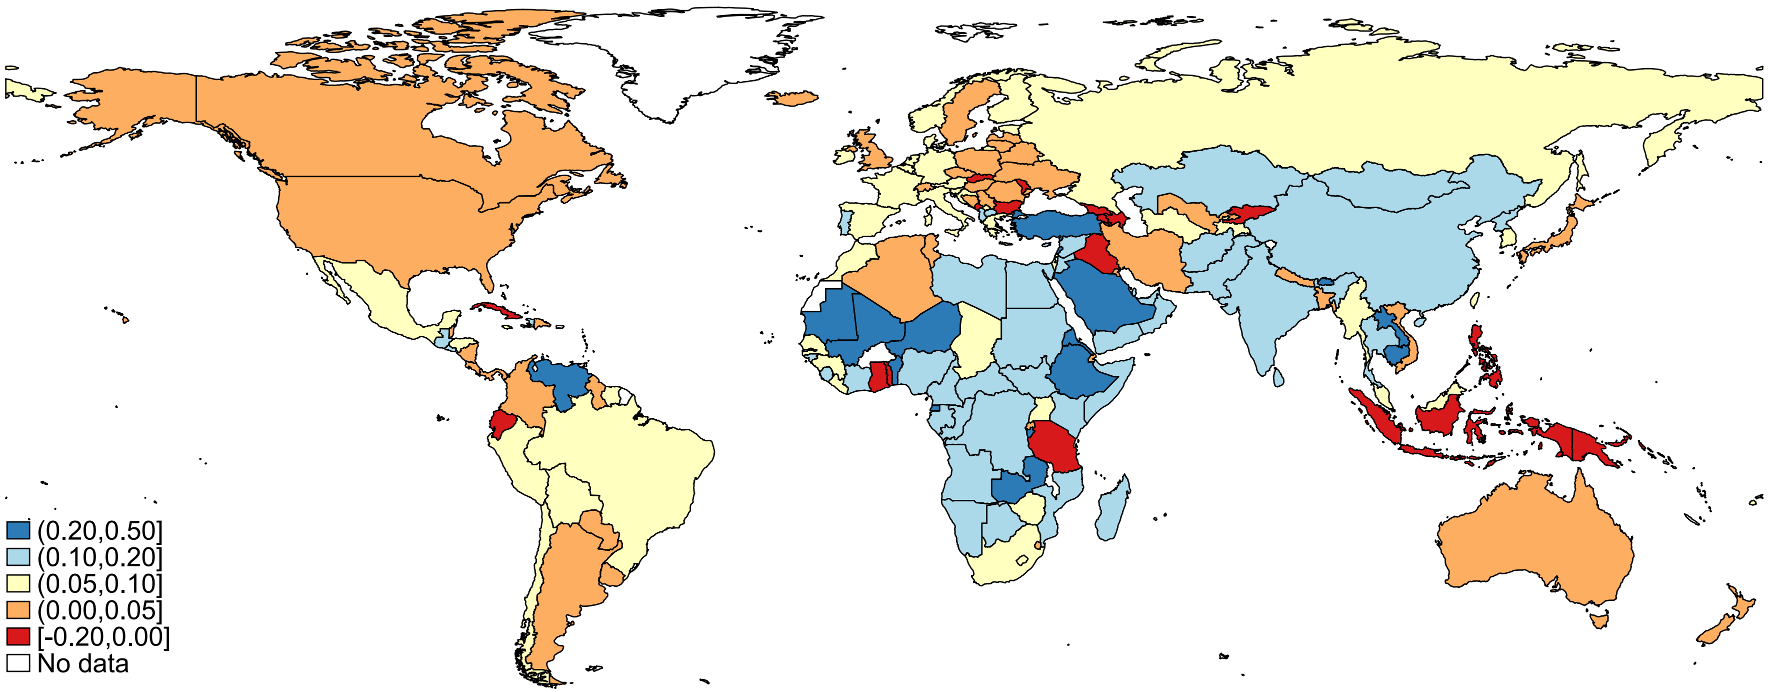

Notes. The map represents the percentage point change in aggregate public healthcare productivity Θ^j between 1980 and 2019 around the world, estimated using a single-input estimate for each function of government.

Figure B.47: Inequality in Public Service Delivery Around the World

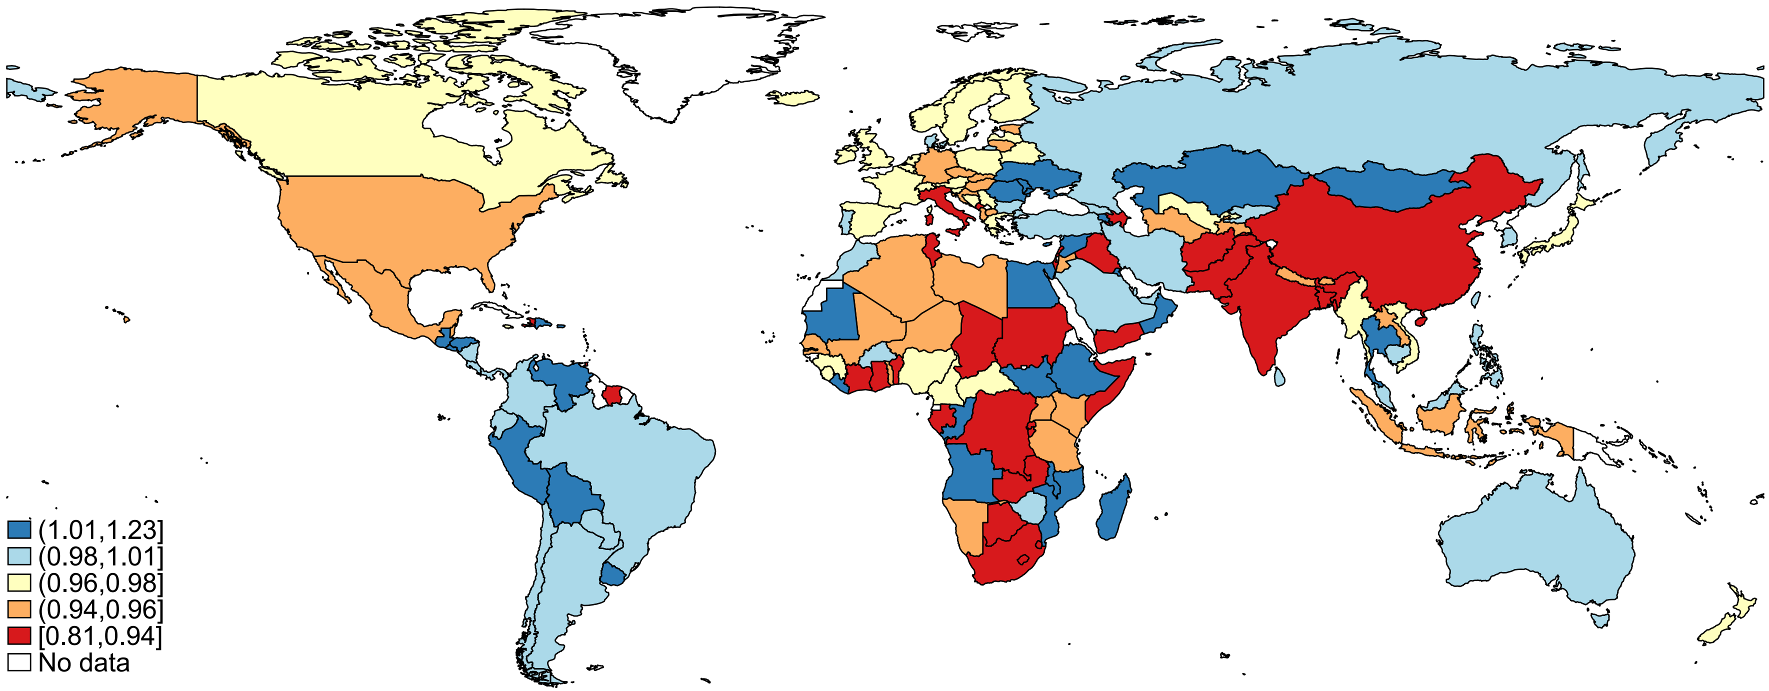

Notes. Author's computations using Gallup World Poll data. The figure represents the relative quality of public services received by the bottom 20% of income earners in comparison to the overall population. Values lower than 1 mean that the bottom quintile receive services of lower quality; values higher than 1 mean that they receive services of better quality. Quality is measured as the share of respondents who declare being satisfied with public services in the city or area where they live, for the following services: public transportation systems, roads and highways, the educational system or the schools, the quality of water, and the availability of quality health care. The quality of police services is measured as the share of respondents who declare having confidence in the local police force. These indicators are then aggregated by income quintile, and the ratio of the bottom quintile to the overall average is computed. Finally, the average of this indicator over all public services is calculated, over the entire 2009-2019 period, and represented in the figure.

Figure B.48: Trends in Equal Access to Public Services Around the World, 2009-2019

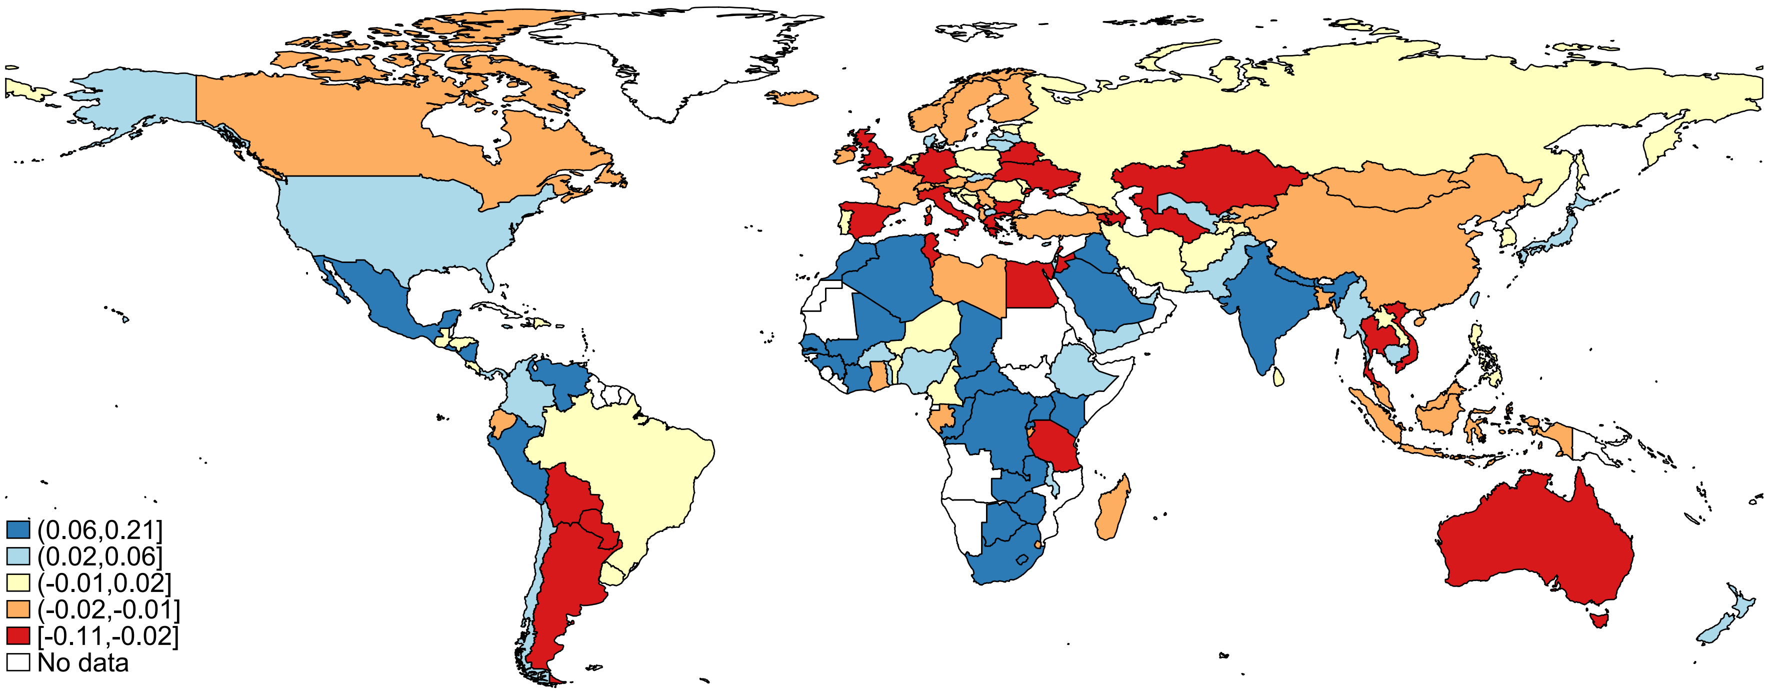

Notes. Author's computations using Gallup World Poll data. The figure represents the change in the relative quality of public services received by the bottom 20% of income earners, in comparison to the overall population, between 2009-2013 and 2016-2019. Values higher than zero mean that public services have become more progressive; values lower than zero mean that they have become more regressive. Quality is measured as the share of respondents who declare being satisfied with public services in the city or area where they live, for the following services: public transportation systems, roads and highways, the educational system or the schools, the quality of water, and the availability of quality health care. The quality of police services is measured as the share of respondents who declare having confidence in the local police force. These indicators are then aggregated by income quintile, and the ratio of the bottom quintile to the overall average is computed. Finally, the average of this indicator over all public services is calculated over the 2009-2013 and 2016-2019 periods, and the difference between the two periods is represented in the figure.

Appendix C

Appendix to "Government Redistribution and Development: Global Estimates of Tax-and-Transfer Progressivity, 1980-2019"

C.1 Distribution of Personal Income Taxes

In this appendix section, expanding on section 3.1.3 above, we provide more detail on methods and data used to estimate the distribution of personal income taxes.

In the case of the personal income tax (PIT), the only tax units that pay any PIT are those whose income places them above the personal income tax exemption threshold. We retrieve these exemption thresholds for more than 90 countries from Jensen (2022), and retrieve the missing country-years from Bachas et al. (2022). Bachas et al. (2022) impute the exemption threshold for country-years missing from Jensen (2022) in a way that is consistent with the findings of the latter study, which discovered that the PIT exemption threshold (expressed as a percentile of the income distribution) falls with rising per capita income, across countries and over time.

Starting from the PIT exemption threshold, we simulate the structure of personal income tax incidence using statutory rate schedules from the World Tax Indicators (WTI) database (see Peter, Buttrick, and Duncan, 2010). This database parameterizes the progressivity of the income tax structure. It observes the average and marginal statutory income tax rates at several levels of the pretax income distribution: at average income, then at two and three and four times that level, and finally the top marginal tax rate. While the WTI covers 189 countries, it does not observe years beyond 2005, so we extend the database with inputs from Strecker (2021) and Vegh and Vuletin (2015, updated 2019), the latter of which can also be used to corroborate top marginal tax rates from WTI. For the remaining country-years (and to check robustness) we retrieve statutory (marginal) rates schedules from Ernst & Young (2006-23) and PwC (2023) and similar sources online, including national tax authorities' legislative documents and independent scholarly accounts. From this basis, we can approximate a continuous schedule of statutory income tax incidence. We assign the statutory tax rate as zero at the exemption threshold K, rising to the top marginal tax rate at *p99.999p100* (the highest g-percentile), with kink points at the rates observed in WTI. Rates are interpolated linearly between each observed value.

Note that we also distinguish between individualized and joint personal income taxation systems: Some countries tax married couples together (or allow tax units this option), and some countries tax individual incomes separately. The former, joint taxation, conforms naturally to the benchmark WID pretax DINA income concept, as these distributions are estimated for "equal-split adults" (where households' total income is split equally among all adult members). However, where PIT systems tax individual incomes, we must transform the WID pretax income distribution from that of "equal-split" adults to that of "individualized" adults.[1] We do this by way of microdata from the International Labour Organization (2020), whose universe of labor force survey microdata represents more than 100 countries since the 1990s. For countries whose PIT systems are individual but for which no (household-identified, individual) income survey microdata exists, we use "nearest-neighbor matching" to simulate the effect, matching the microdata from a handpicked neighboring country. For tractability and reliability of the estimate, we implicitly assume a generalized country fixed-effect [rather than by country and year, i.e., we do not allow each country's distribution-wide correlations between individualized and equal-split incomes to vary over time] and only use the latest-year survey.[2] In this way, we are

[1]Note that individualized income distributions are more unequal than equal-split income distributions. This is so by construction among top earners (only if all top earners were married to each other would their equal-split incomes equal their individualized incomes), and generally true throughout the distribution. The left tail of the individualized distribution contains many more observations with zero incomes (non-working spouses).

[2]It is true that labor force participation—including among spouses—and assortative matching of high-income earners may change over time, but it is also true that this survey data is not the most reliable source to capture the entire effect, as it has little to say about capital income nor about

able to estimate the ratio of individualized income to equal-split income, across the g-percentile distribution, and to easily move back-and-forth between equal-split and individualized income distributions.

After we assign taxes to individuals, we can transform the taxes paid by each household—from an effective rate on individualized income, to an effective rate on equal-split income. For example, for a married couple in an individualized tax system, earning two different levels of income and being taxed at two different rates, this transformation adds up both the incomes earned and the taxes paid by the couple, then divides these by two for the uniform effective rate on their (identical, by construction) equal-split incomes. For countries whose PIT system is on individuals' incomes rather than taxing married couples jointly, this ILO-microdata transformation effectively moves an individualized income tax schedule onto the equal-split income distribution, with effective tax rates transformed accordingly.

Finally, we account for the empirical regularity that capital income is taxed less than labor income in PIT systems worldwide.[3] For each country for which we observe tax revenue aggregates (and statutory PIT rates on taxable income), we also tabulate the country's tax rates toward dividends and capital gains. While there are nuances within many tax administrations' policies on the taxability of dividends and capital gains [and other types of capital incomes], we simplify concepts for tractability on tax rates and tax bases in a DINA framework: Our benchmark concept for the rate of dividend taxation is the rate at which a resident is taxed on dividends from domestic companies. Similarly, our benchmark concept for the rate of capital gains taxation is the rate at which a resident is taxed on gains from selling shares in domestic companies. In the latter case, we also acknowledge reduced rates, or exemptions on short- vs. long-term capital gains, or other nuances in the treatment of this type of income.

These are not the only types of capital income that are taxed by PIT systems, but in our view they are the most significant and telling. Among other types of capital income that may be subject to tax in a PIT system: Mixed income comprises a capital share and a labor share; however, in most countries *all* self-employment income is taxed similarly to labor income from salaries and wages. Beyond that, many PIT systems cover (capital) income from rentals, from interest, from royalties,

household enterprise. In detailed DINA studies of the United States and France, the disparity between individualized and equal-split distributions has remained relatively stable over the past 40 years.

[3] Globally, we find that only 36% of corporate operating surplus (profits) is distributed in the form of dividends.

etc. It is perhaps worth noting that, from a DINA perspective, these are not in the "primary generation of income" account and would actually be double-counting part of national income if they were counted as individual's income without subtracting the corresponding part from, e.g., corporate profits (which indeed they would be in any fiscal system). In this sense, it seems reasonable to leave our simplified PIT simulation as taxing distributed (dividend) and undistributed (capital gain) corporate profits, with these two elements (summing to total corporate profits in the national accounts) also serving as a proxy for the tax treatment of other capital incomes. In any case, the tax treatment of interest, rents and royalties is usually very similar to that of dividends and capital gains. In our view, dividends and capital gains taxation represent emblematic proxies which together serve to cover what is *taxable* in pretax DINA capital incomes.

To assign dividend and capital gains rates that vary by country allows us a treatment of capital incomes under PIT systems that matches the rigor of the above-mentioned statutory rates on labor (salary, wage, and self-employment) income. The upshot is that much of capital income is untaxed, or taxed at a lower rate.Taxable income (in this concept) is less than total pretax income (in the DINA sense), and particularly so for the top g-percentiles where capital income is concentrated. Among DINA income concepts, we also exclude from the PIT tax base: imputed rent, government operating surplus, and indirect taxes. Social insurance benefits received are taxed as (deferred) labor income.

The elements of the PIT system, in this simplified simulation, can be summarized as follows, to estimate the tax rate τ for any g-percentile p and its corresponding income level z:

$$\tau(z)_{PIT} = \sum_{j=1}^{3} \frac{\tau_j z_j}{z}$$

where j refers to three types of PIT taxes (with taxable incomes z_j taxed at rate τ_j):

- labor income (employee compensation and mixed income[4]);
- dividend income (distributed corporate profits); and
- capital gains income (undistributed corporate profits).

[4]All of self-employment (*viz.* mixed) income is treated as labor income, for the purposes of this PIT simulation—as is the case in most PIT systems.

After building this statutory rate schedule, we fit its "predicted" revenues to actual PIT revenues received, observed in Bachas et al. (2022) and corresponding to T_{PIT} in equation (3.2) above. In this way, we simulate statutory rates in order to estimate effective tax rates throughout the distribution. It is important to note that the "predicted" statutory rates above do not match—but rather are proportional to—the effective rates we estimate. This mismatch between statutory and effective rates is to be expected, and can be true for a number of reasons that we do not observe in aggregate data (e.g., tax evasion or avoidance; unobserved deductions, allowances, exemptions and tax breaks that vary with income; differences within the rate schedule according to different types of [non-]taxable income, etc.).

Since we do not necessarily observe all the nuances by which an effective tax rate may differ from the statutory rate (even if we think that we have captured the main drivers above), we are *almost* forced to assume that the effective rate schedule retrieved from our statutory rate schedule is the correct one (i.e., that the "true" effective rate schedule is proportional to our estimated statutory rate schedule)—and holds as valid for the distribution of personal income tax rates along the pretax income distribution.

However, we do not have to leave this as an assumption, and can instead test its robustness (as a goodness-of-fit) against the existing DINA studies mentioned above. For reference, see Appendix Figure C.23 to compare the time series of US personal income tax rates between the benchmark estimates of Piketty, Saez, and Zucman (2018) and those of the present simulation—comparing the benchmark to our simulation at each of three representative points on the income distribution: p50, p90, and p99. As can be readily seen in the graph, the fit is excellent, and our simulated effective PIT rates rarely differ by more than half of a percentage point, matching on both levels and trends.

Given the goodness-of-fit of our simulation against the training sample of microdata-founded (DINA) estimates of PIT incidence, we are confident to extend our estimates to the worldwide sample of countries for whom we have collected precise data on the set of parameters listed above (the minimum from which we can estimate PIT incidence, as discussed here).

C.2 Measures of Fiscal Progressivity

In our main analysis, we summarize the progressivity of taxes (and/or transfers) with the percent difference in inequality, measured as the top 10% to bottom 50%

average income ratio, before and after removing taxes from (and/or adding transfers to) individual incomes. This is equation (3.6):

$$\gamma_\tau = \frac{r_{pre} - r_{net}}{r_{pre}}$$

After some algebra, this absolute progressivity statistic γ_τ—the redistribution ratio representing the percent reduction in inequality from fiscal policy—reduces to:[5]

$$\gamma_\tau = \frac{\overline{ETR}_{p90p100} - \overline{ETR}_{p0p50}}{1 - \overline{ETR}_{p0p50}} \tag{C.1}$$

Since γ_τ is a function only of the ETR profile (i.e., of the bracket average ETRs at the top and bottom of the income distribution), it is independent of the pretax inequality ratio r_{pre}. For the same ETR profile, γ_τ highlights the same percentage of redistribution, regardless of the overall level of inequality.

We note, however, that the "naive" γ_τ is sensitive to variations in the pretax income distribution *within* the top 10% or bottom 50% shares, i.e., different distributions of *p90p100* or *p0p50* incomes that would still deliver the same *average* income for the top 10% or bottom 50% shares, respectively.

To see why, imagine a monotonically increasing ETR profile within the bottom 50% of earners, e.g., from $ETR = 0\%$ at *p0* to $ETR = 10\%$ at *p50*, and a steeply increasing income profile within the same bottom 50% of earners, such that most of the income of the bottom 50% is near *p50*. In this case, the average ETR of the bottom 50% of earners would be close to 10% (the ETR at *p50*). By contrast, if the income distribution were closer to flat within the bottom 50%, the same ETR profile would deliver an average ETR closer to 5%. The redistribution ratio would be higher in latter case (where the average ETR of the bottom 50% is lower). The same idea holds for the top of the distribution *p90p100*. Intuitively, we would prefer a progressivity statistic that delivers the same results when applying a given ETR

[5]To arrive at this equation, we put r_{net} in terms of r_{pre} and plug into equation (3.6):

$$r_{net} = \frac{\bar{y}^{net}_{p90p100}}{\bar{y}^{net}_{p0p50}} = \frac{\bar{y}^{pre}_{p90p100} - \bar{y}^{pre}_{p90p100} \cdot \overline{ETR}_{p90p100}}{\bar{y}^{pre}_{p0p50} - \bar{y}^{pre}_{p0p50} \cdot \overline{ETR}_{p0p50}} = \frac{\bar{y}^{pre}_{p90p100}(1 - \overline{ETR}_{p90p100})}{\bar{y}^{pre}_{p0p50}(1 - \overline{ETR}_{p0p50})} = r_{pre} \cdot \frac{1 - \overline{ETR}_{p90p100}}{1 - \overline{ETR}_{p0p50}}$$

$$\gamma_\tau = \frac{r_{pre} - r_{pre} \cdot \frac{1-\overline{ETR}_{p90p100}}{1-\overline{ETR}_{p0p50}}}{r_{pre}} = 1 - \frac{1 - \overline{ETR}_{p90p100}}{1 - \overline{ETR}_{p0p50}} = \frac{(1 - \overline{ETR}_{p0p50}) - (1 - \overline{ETR}_{p90p100})}{1 - \overline{ETR}_{p0p50}}$$

profile to any pretax income distribution—and even robust to distributional variance within *p0p50* or *p90p100* (at the same $\bar{y}_{p0p50}$ and $\bar{y}_{p90p100}$).

To test sensitivity and resolve this potential source of bias, we normalize pretax income distributions across all countries and years. Following the literature from Kakwani (1977) through Gerber et al. (2020), we assign as constant the arbitrary income distribution $y_p = p^2$, a distribution whose inequality ratio r_{pre} happens to be close to the median value observed in our data. From this normalized pretax distribution, we calculate the net-of-tax distribution, as always, by subtracting taxes according to each country-year's observed ETR profile. Results of this exercise, in Figure C.25, are visibly similar to those of the earlier Figure 3.4 (above).

In this way, we generate a statistic that is independent of all variation in pretax income distributions, while still capturing qualities of both relative and absolute progressivity.

By *relative* progressivity, we refer to the comparison of the ETR on top 10 percent earners vs. on bottom 50 percent earners (expressed as the percent difference $\frac{\overline{ETR}_{p90p100} - \overline{ETR}_{p0p50}}{\overline{ETR}_{p0p50}}$, visible in, e.g., Figure C.24). A higher ratio between the two would be more progressive, by construction. Other, similar measures of relative progressivity could include the regression coefficient (slope) of the tax rate profile (see Peter, Buttrick, and Duncan, 2010, and section 3.1.4 above)—but of course one would also want to know the y-axis intercept and not only the slope of the profile. These measures, then, while relatively informative, do not necessarily account for the total level of taxation.

With an *absolute* progressivity statistic, we do account for the total level of taxation. If the slope of the ETR profile is greater than (less than) zero, an increase in total taxation is an increase (decrease) in absolute progressivity, even with no change in the slope of the ETR profile (see Kakwani, 1977). For this reason, our benchmark measure of fiscal progressivity is the one in the two equations above.

C.3 Additional Figures and Tables

Figure C.1: Corporate Income Tax: Selected Estimates of Corporate Income Tax Progressivity

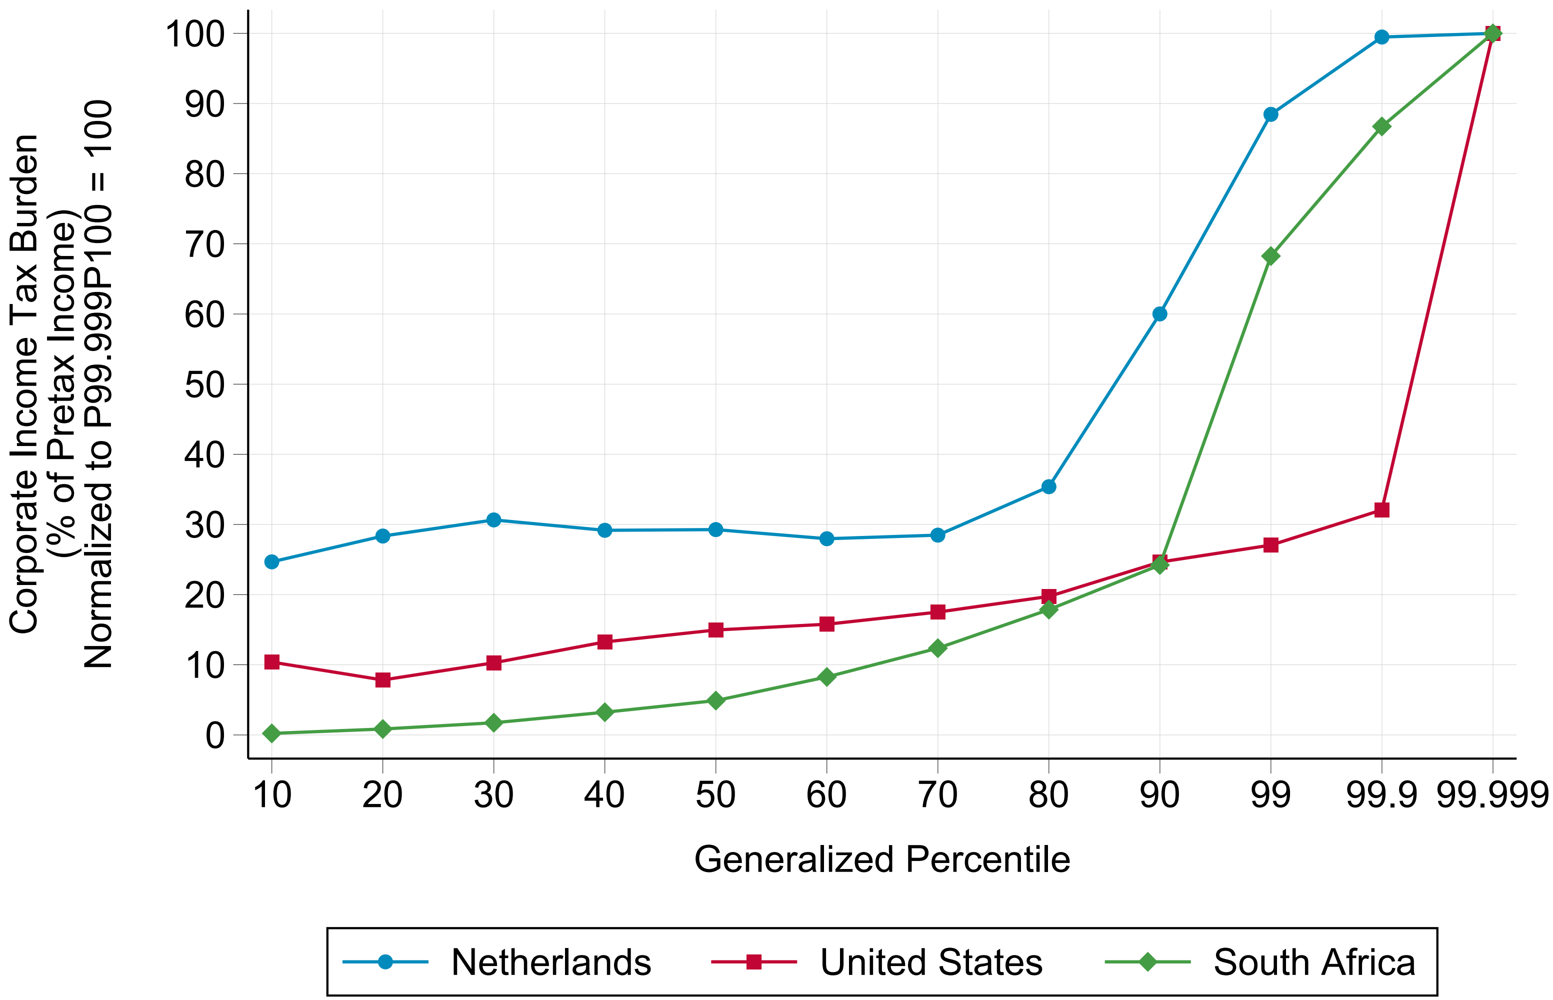

Notes. Netherlands: data from Bruil et al. (2022), 2016. United States: data from Piketty, Saez, and Zucman (2018), 2019. South Africa: data from Chatterjee, Czajka, and Gethin (2023), 2010-2019 average.

Figure C.2: Distributional Incidence Profiles: Income to Consumption Ratio

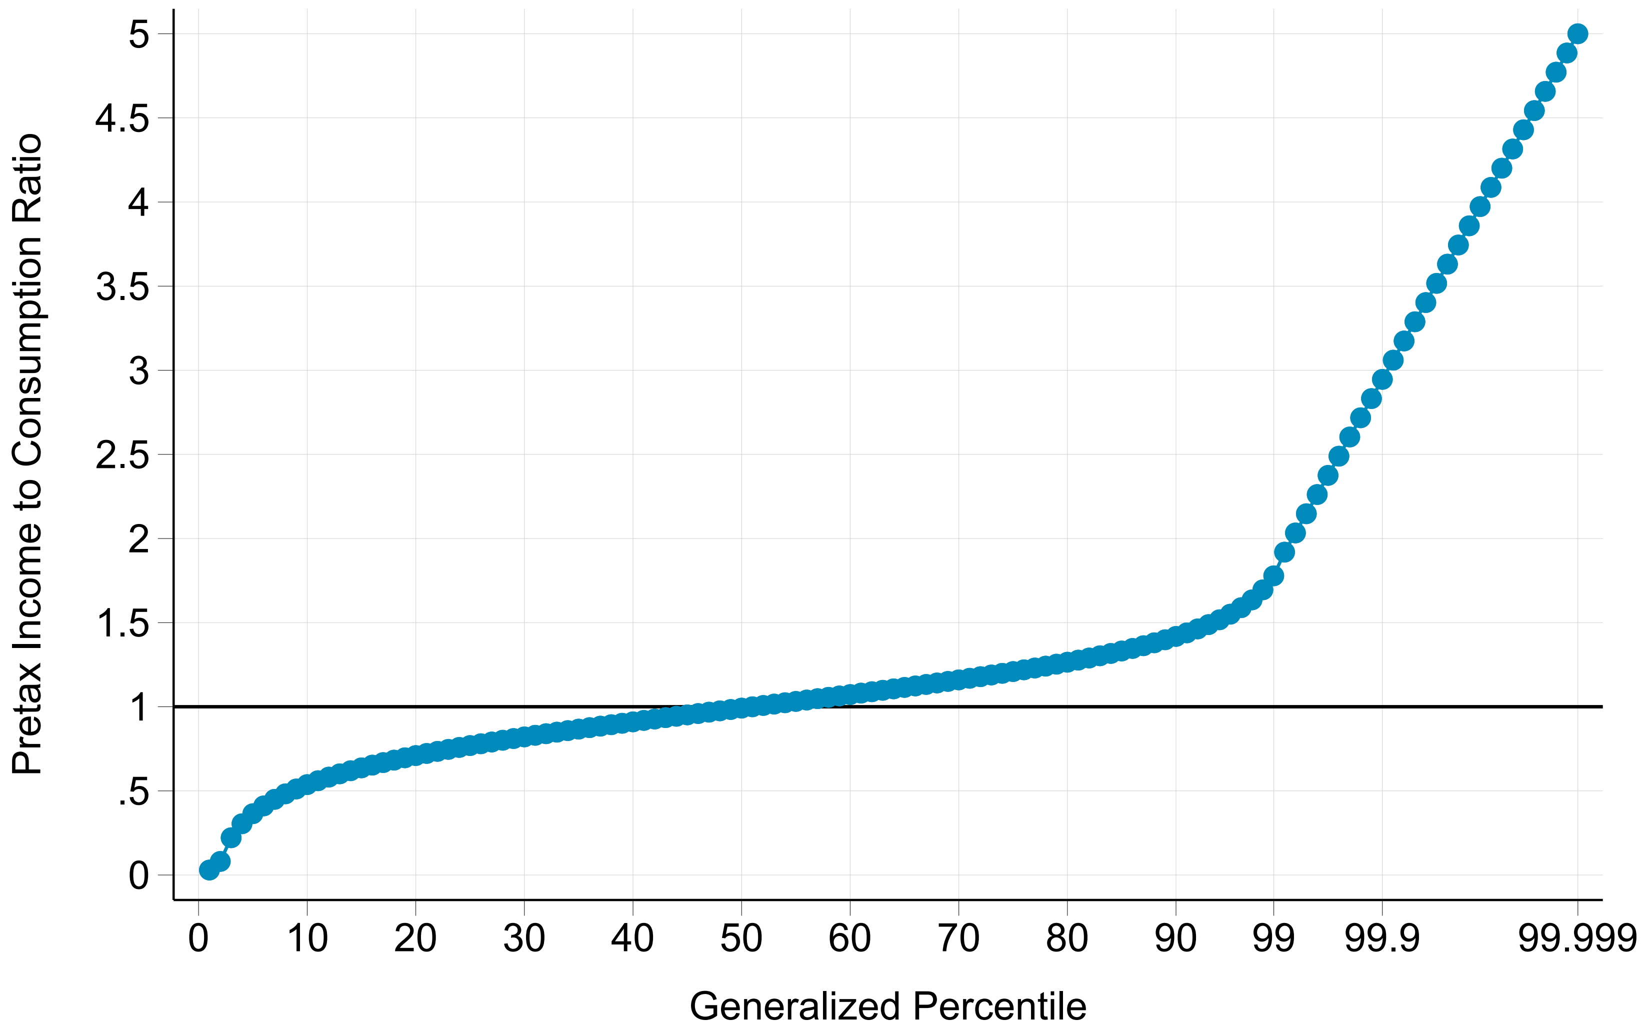

Notes. Authors' elaboration. The figure plots the stylized profile used to estimate consumption from pretax income in each country. See Chancel et al. (2023) for more details.

Figure C.3: Informal Consumption Elasticity and Economic Development

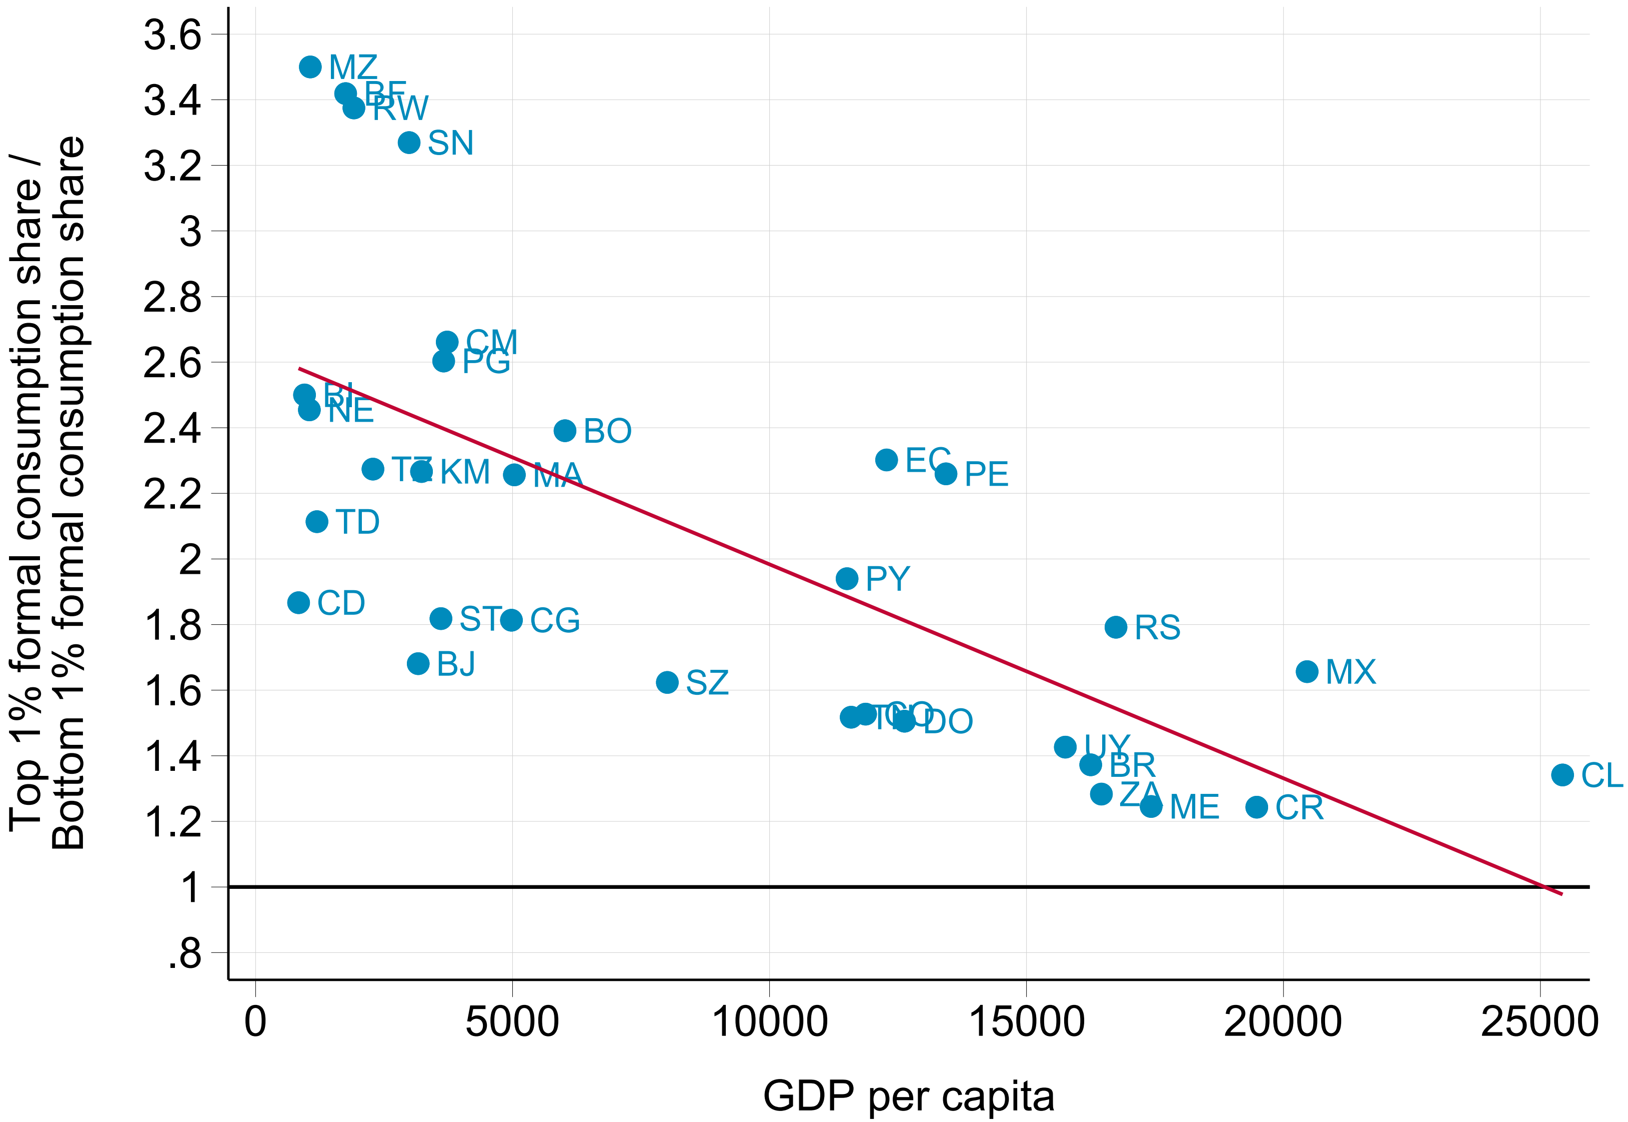

Notes. Authors' elaboration combining data from the World Inequality Database (GDP per capita) and Bachas, Gadenne, and Jensen, 2022 (informality). The figure plots the relationship between GDP per capita expressed in 2021 PPP USD and the gap in informal consumption between top and bottom income groups. In poorer countries, low-income households purchase more goods and services in informal markets than high-income households to a greater extent than in high-income countries.

Figure C.4: Incidence of Indirect Taxes and Informality: Niger, 2019

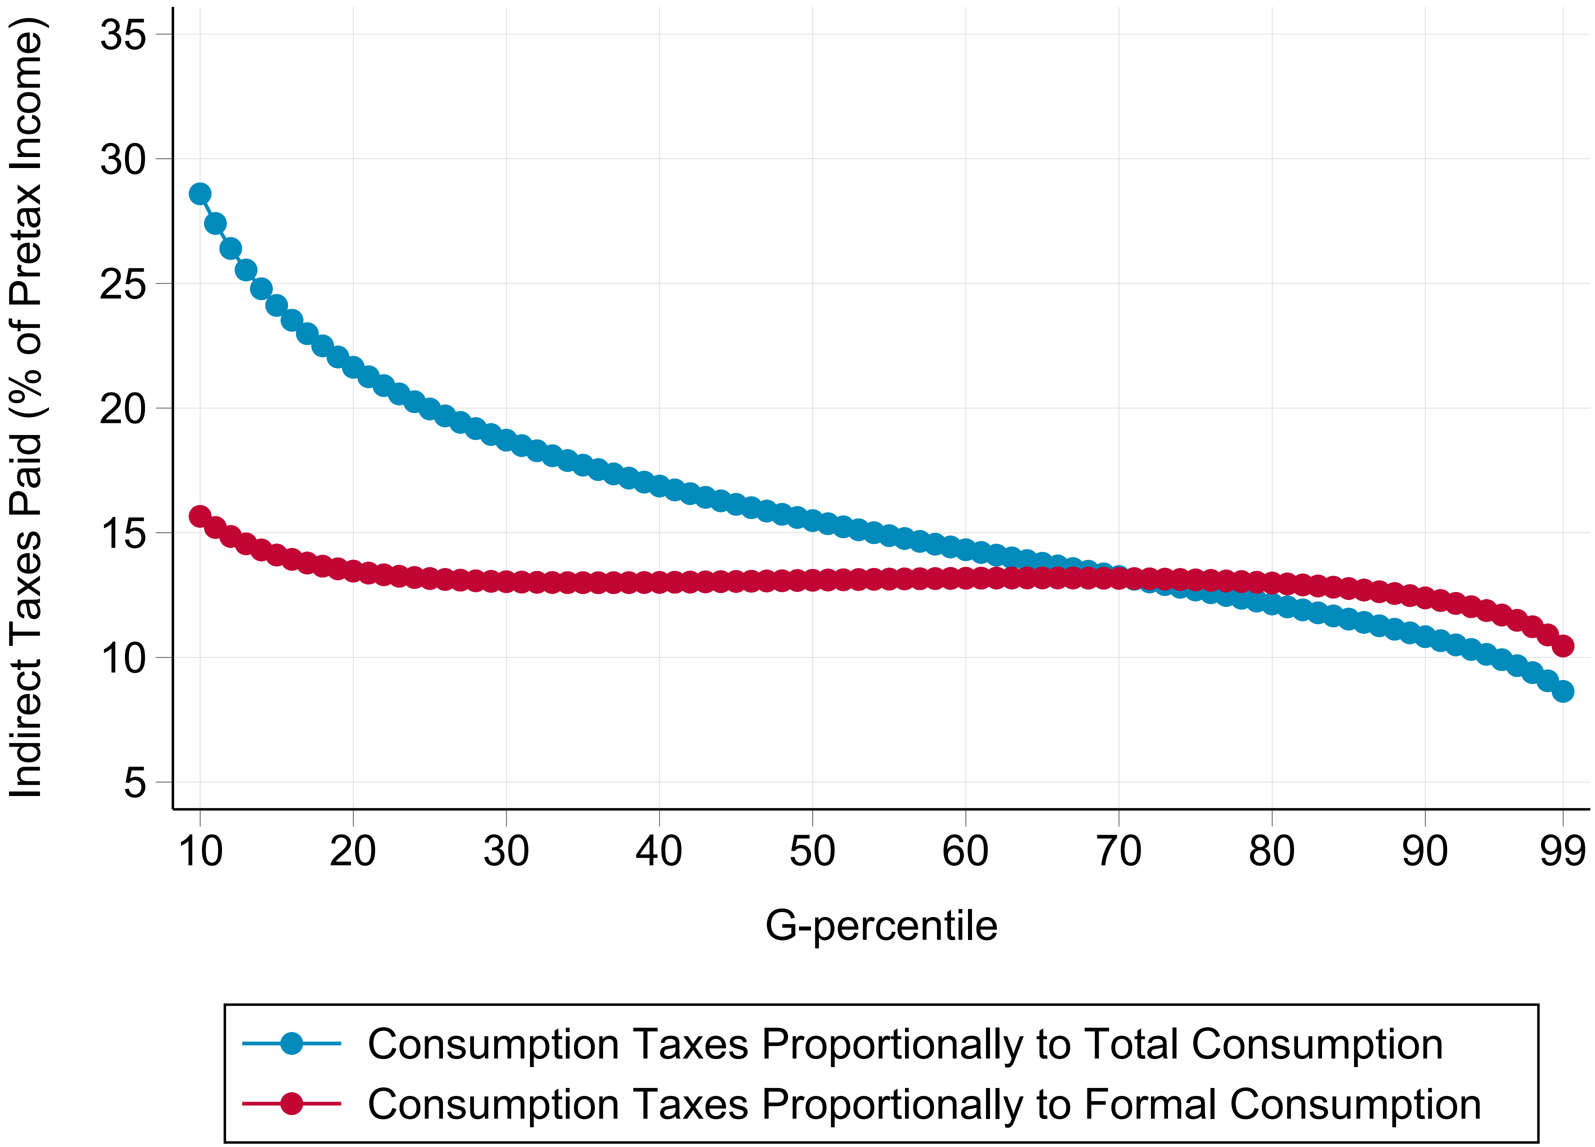

Notes. Authors' elaboration. The figure plots estimates of the distributional incidence of indirect taxes in Niger in 2019, before and after accounting for informal consumption. Before accounting for informal consumption, consumption taxes are very regressive, because low-income households tend to dissave, while high-income households display large positive savings. After accounting for the fact that low-income households tend to more intensively consume in informal markets, however, consumption taxes appear to only be mildly regressive.

Figure C.5: Incidence of Social Contributions and Informality: Argentina, 2019

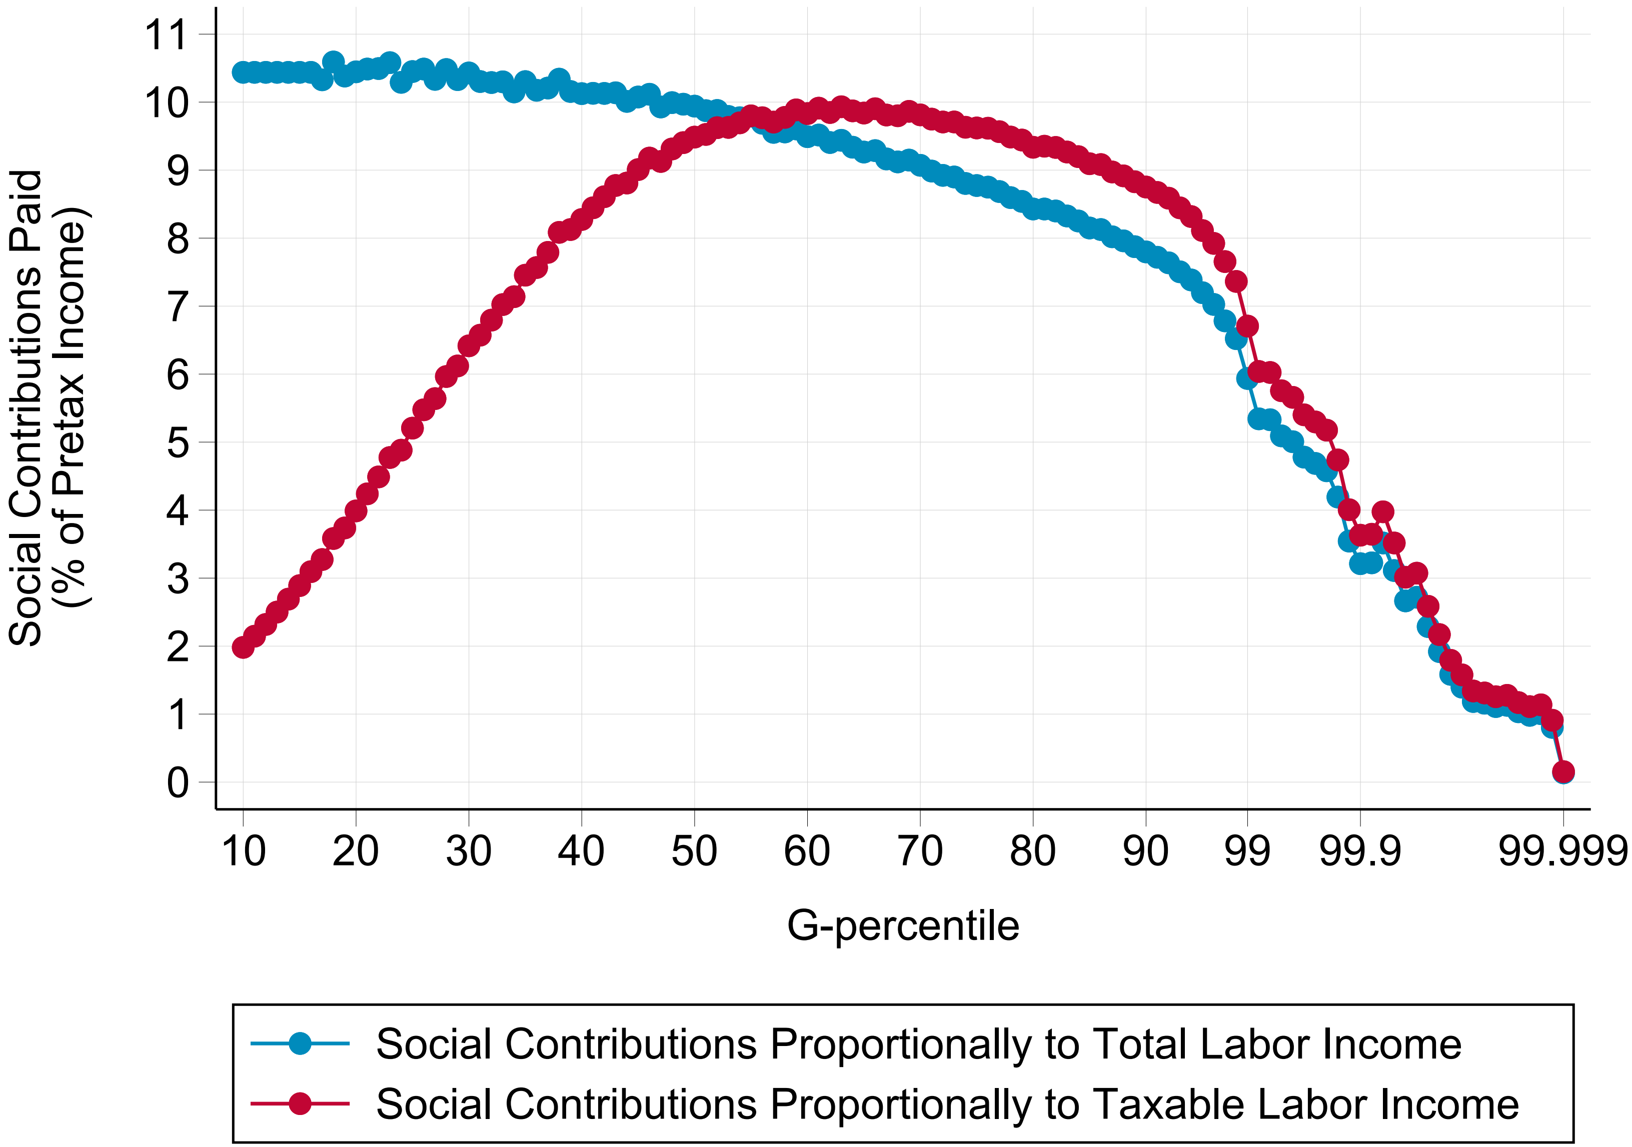

Notes. Authors' elaboration. The figure compares the distributional incidence of social contributions in Argentina before and after accounting for the fact that contribution payments different alongside the wage distribution. Distributing contributions proportionally to total labor income (blue line) implies a much more regressive profile than when distributing them proportionally to taxable labor income (red line), that is, accounting for the fact that a large share of low-wage earners do not pay social contributions.

Figure C.6: Validation: Comparison of Distributional Tax Incidence, by Type of Tax

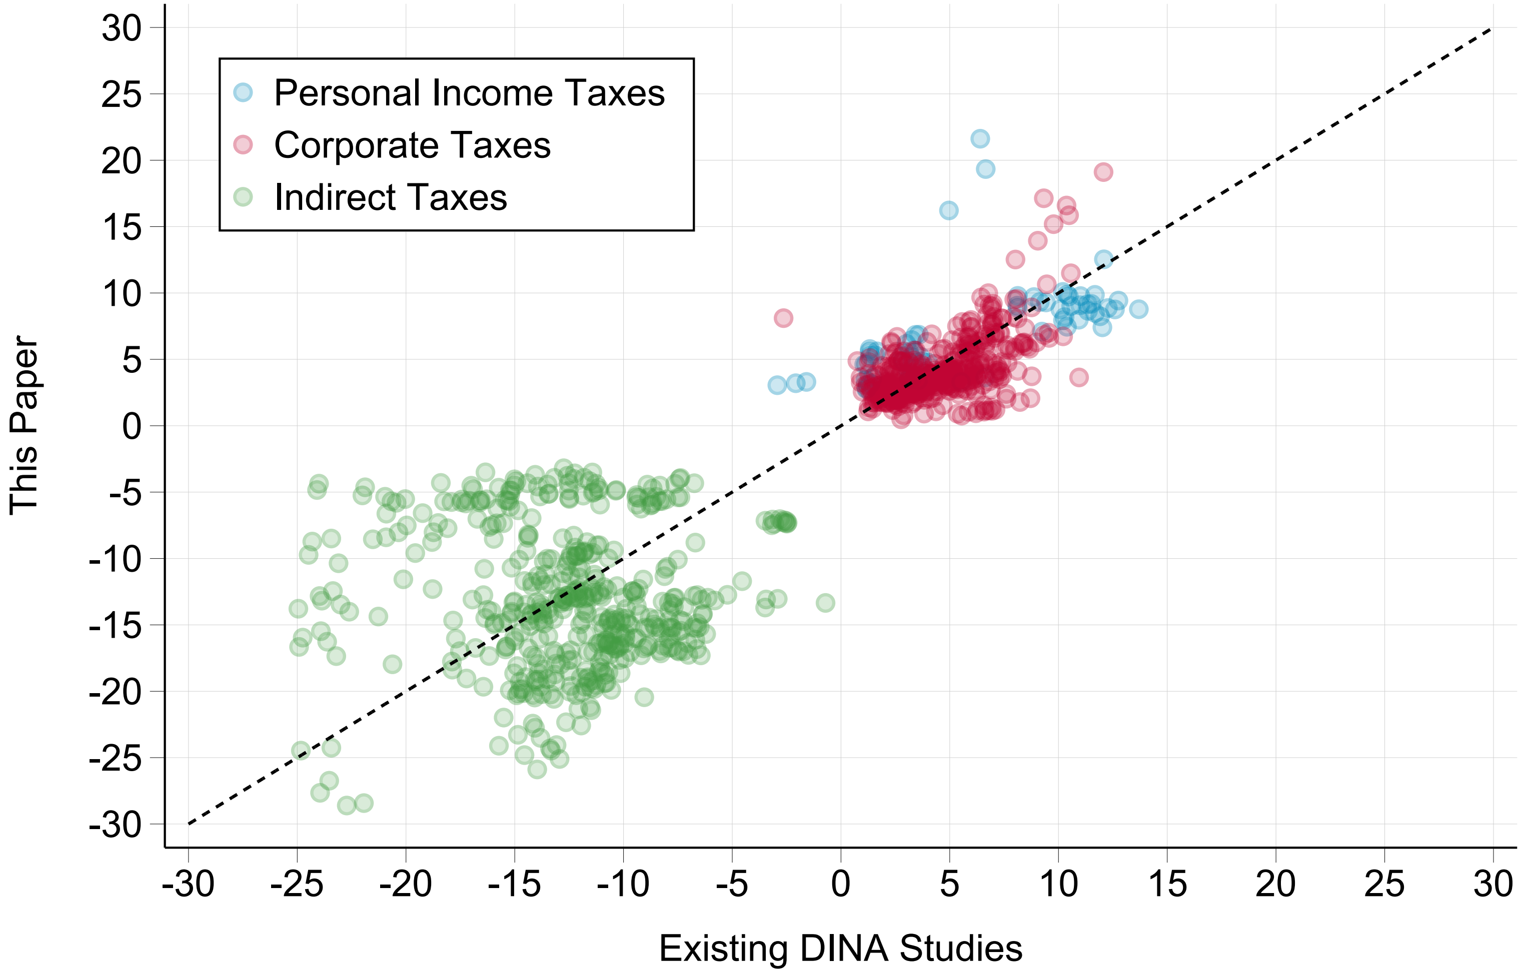

Notes. Axes represent tax progressivity summary statistic γ_τ.

Figure C.7: Validation: United States
Level and Composition of Taxes Paid by Generalized Percentile

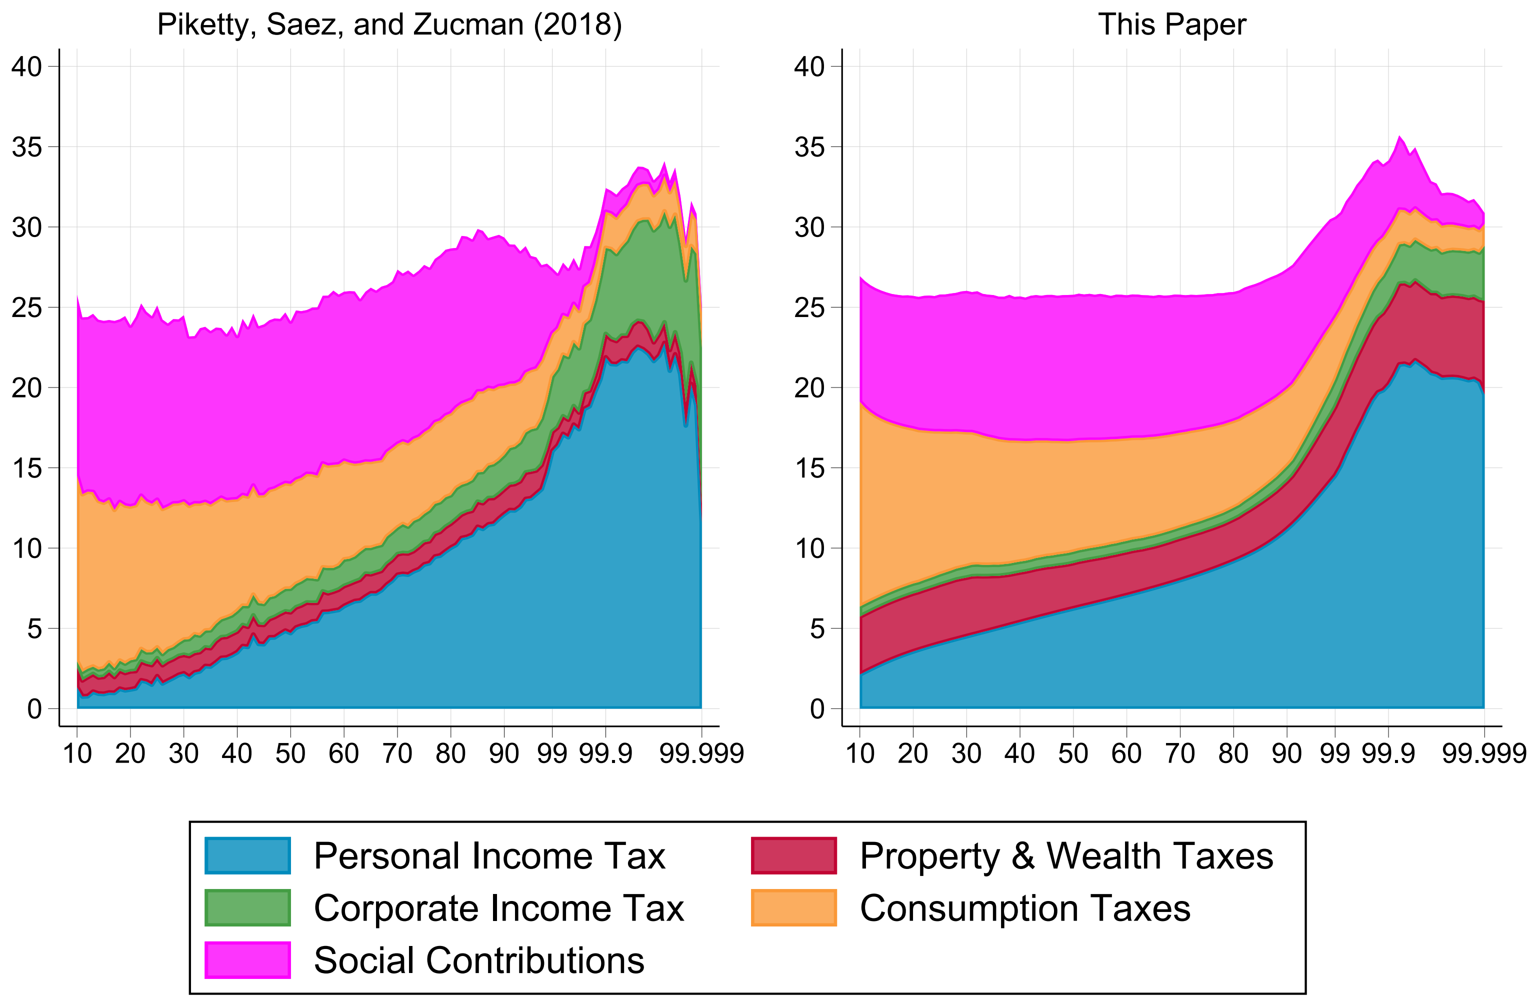

Figure C.8: Validation: Netherlands
Level and Composition of Taxes Paid by Generalized Percentile

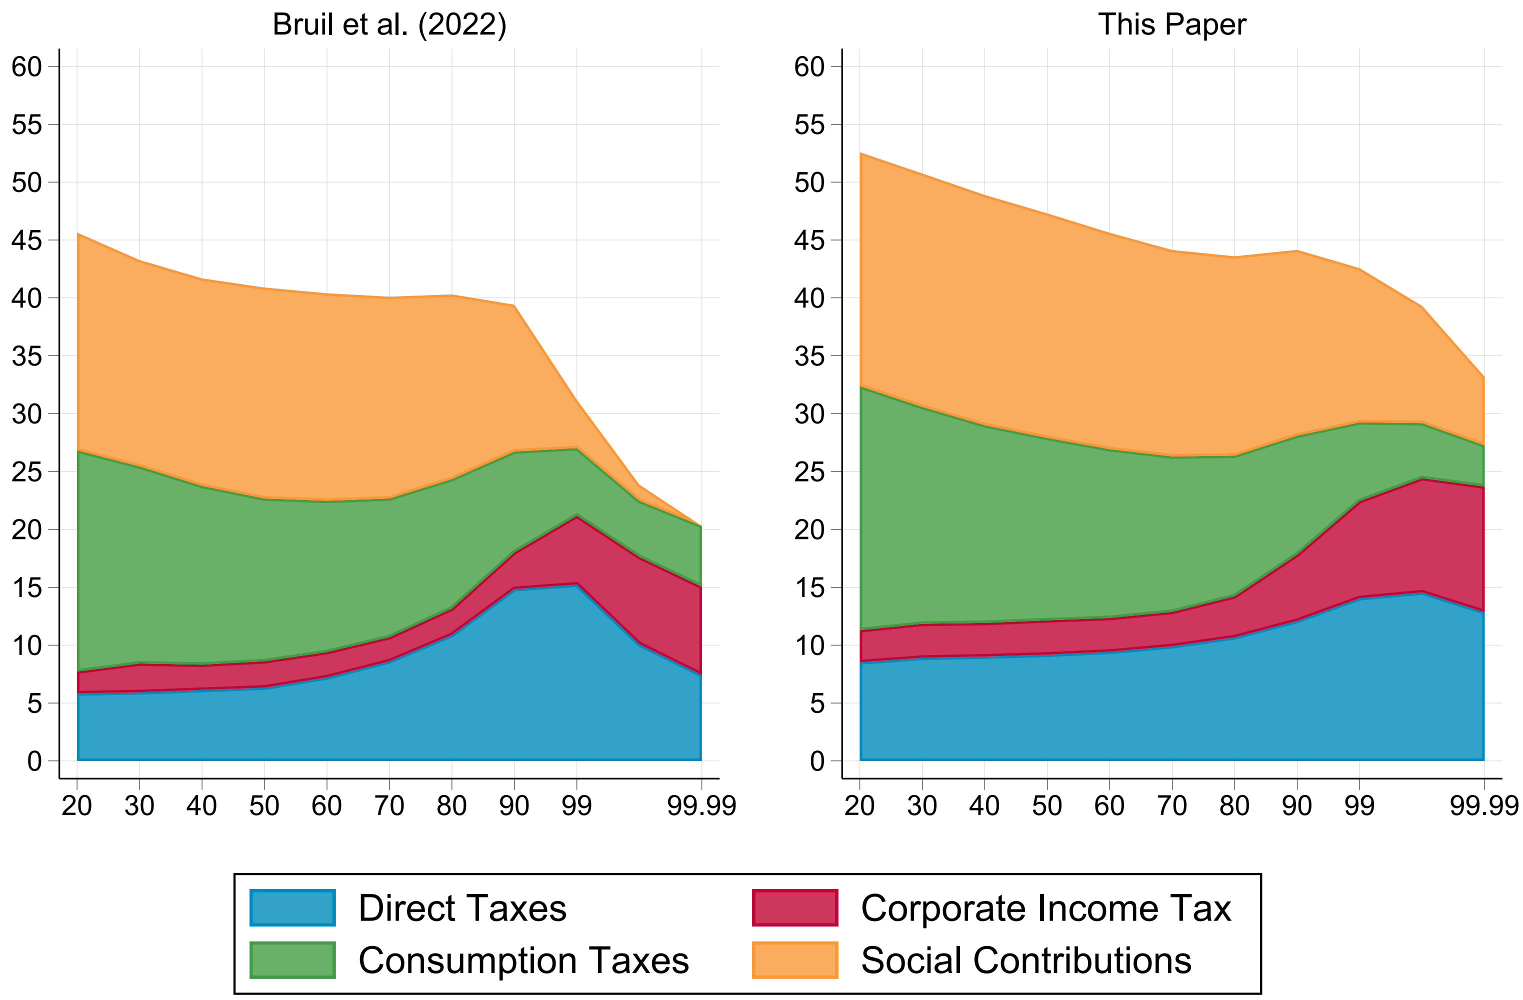

Figure C.9: Validation: South Africa, 2019
Level and Composition of Taxes Paid by Generalized Percentile

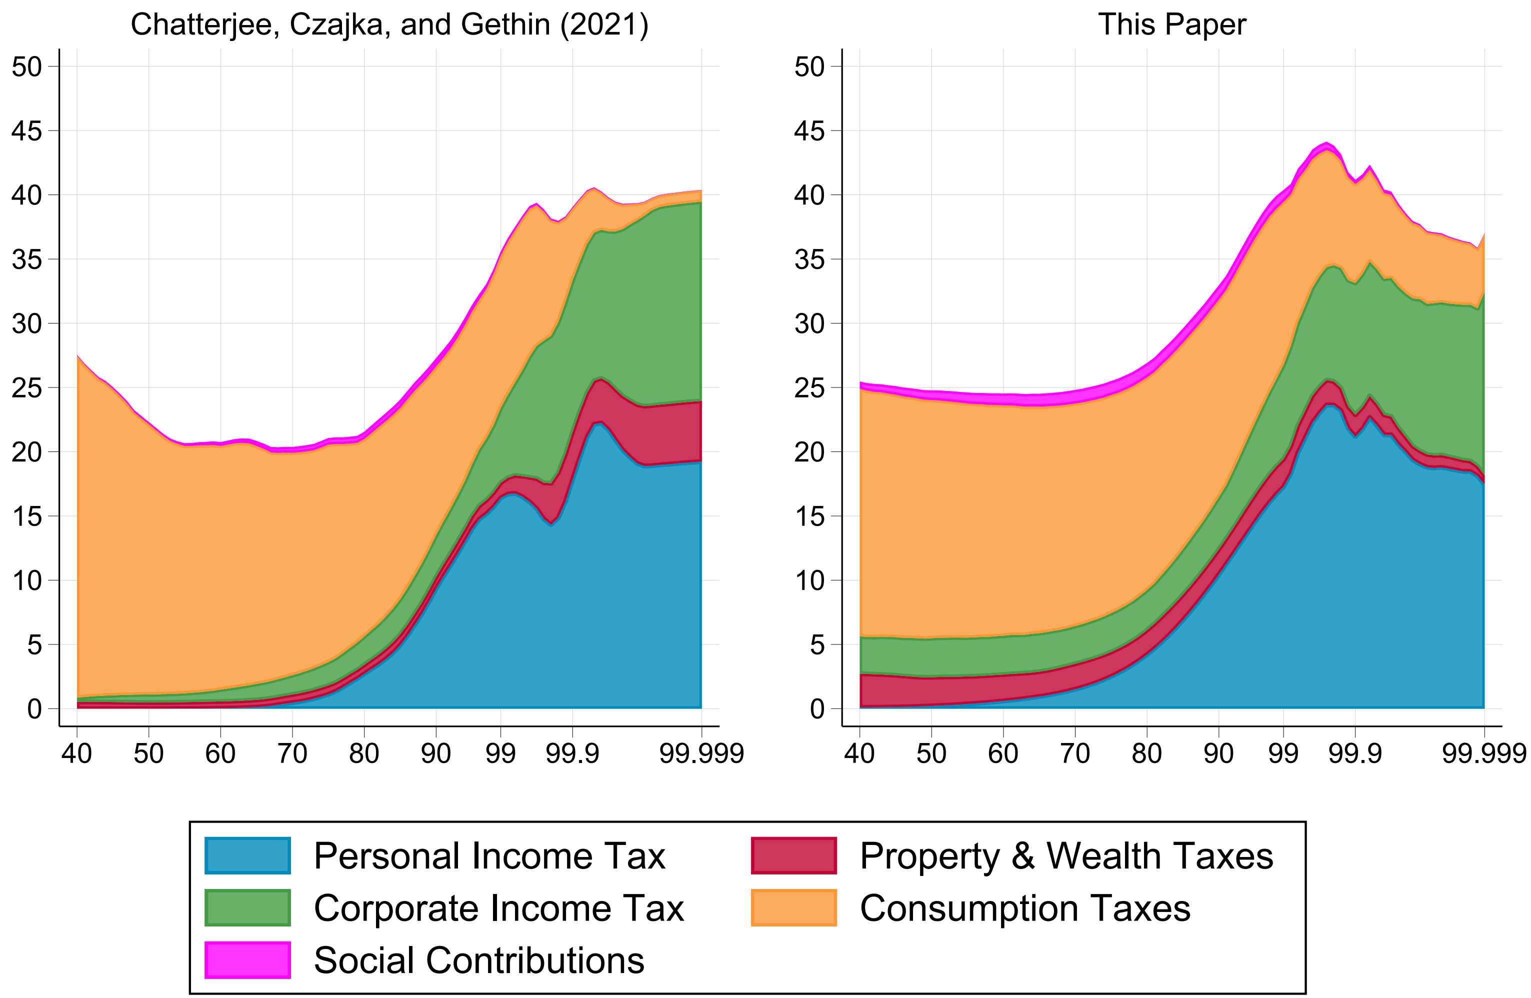

Figure C.10: Validation: Cross-Country Differences in Tax Progressivity

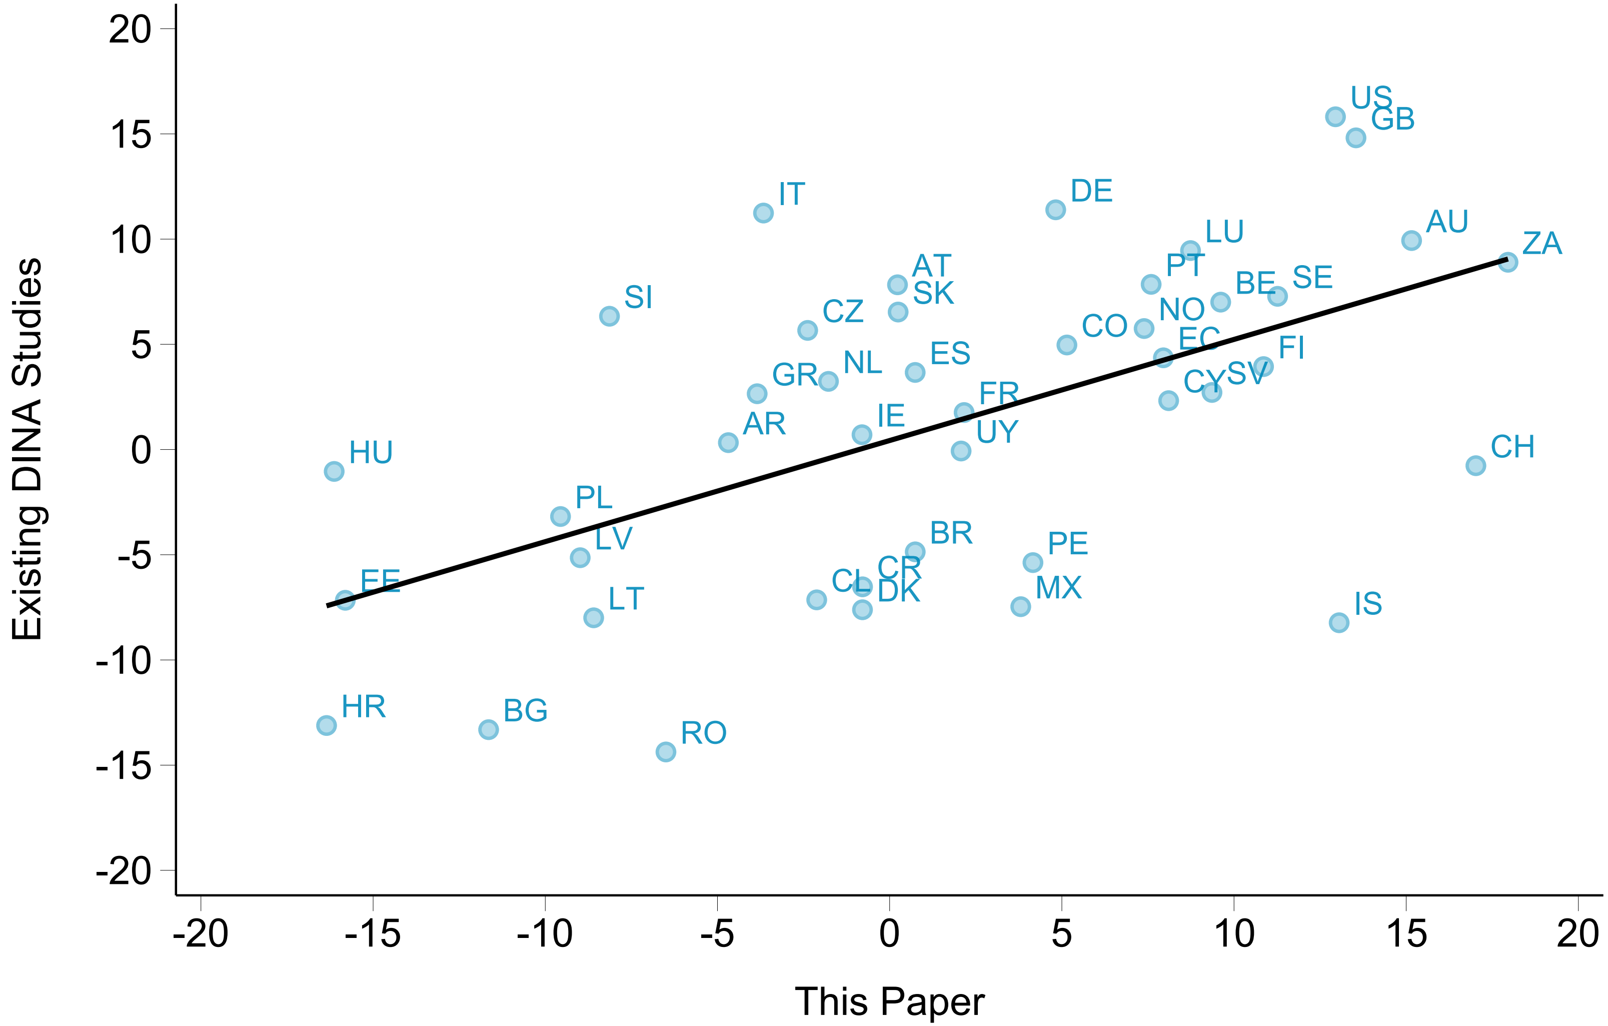

Notes. Authors' elaboration. The figure compares our estimates of tax progressivity to that of DINA papers across all country-years available. Tax progressivity is measured as the percent difference in the top 10% to bottom 50% average income ratio before and after removing taxes from pretax income.

Figure C.11: Top 1% Effective Tax Rate

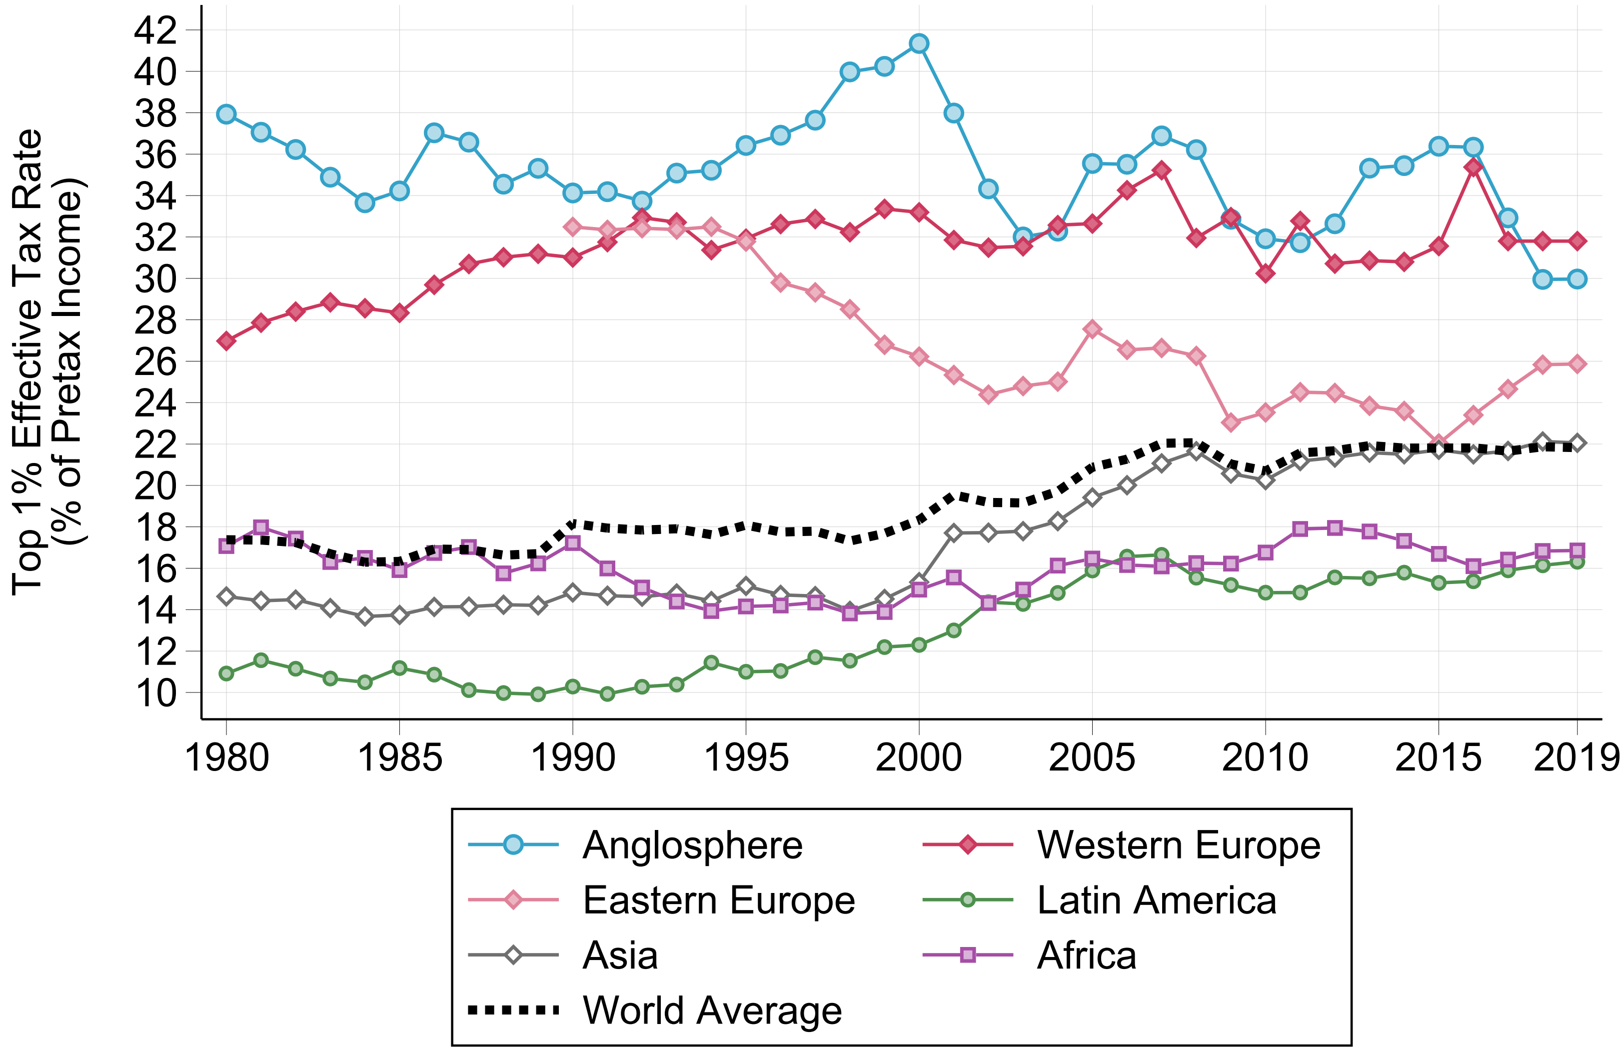

Notes. Authors' elaboration. Excludes social contributions.

Figure C.12: Top 10% Effective Tax Rate

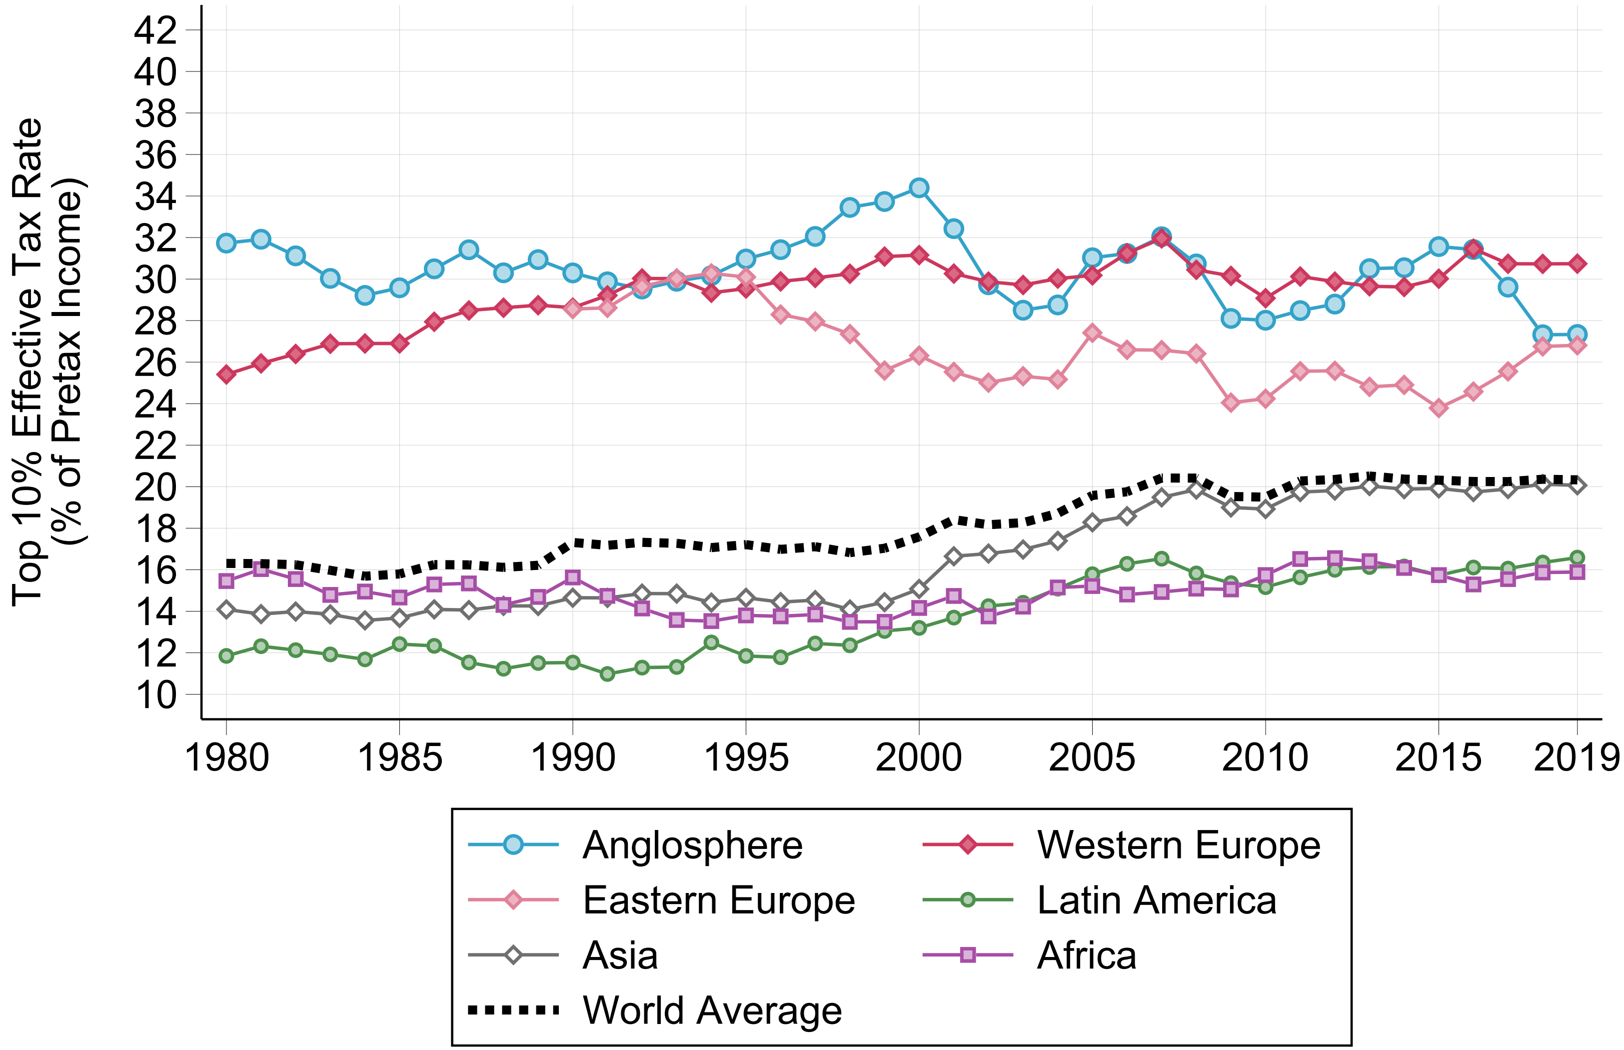

Notes. Authors' elaboration. Excludes social contributions.

Figure C.13: Bottom 50% Effective Tax Rate

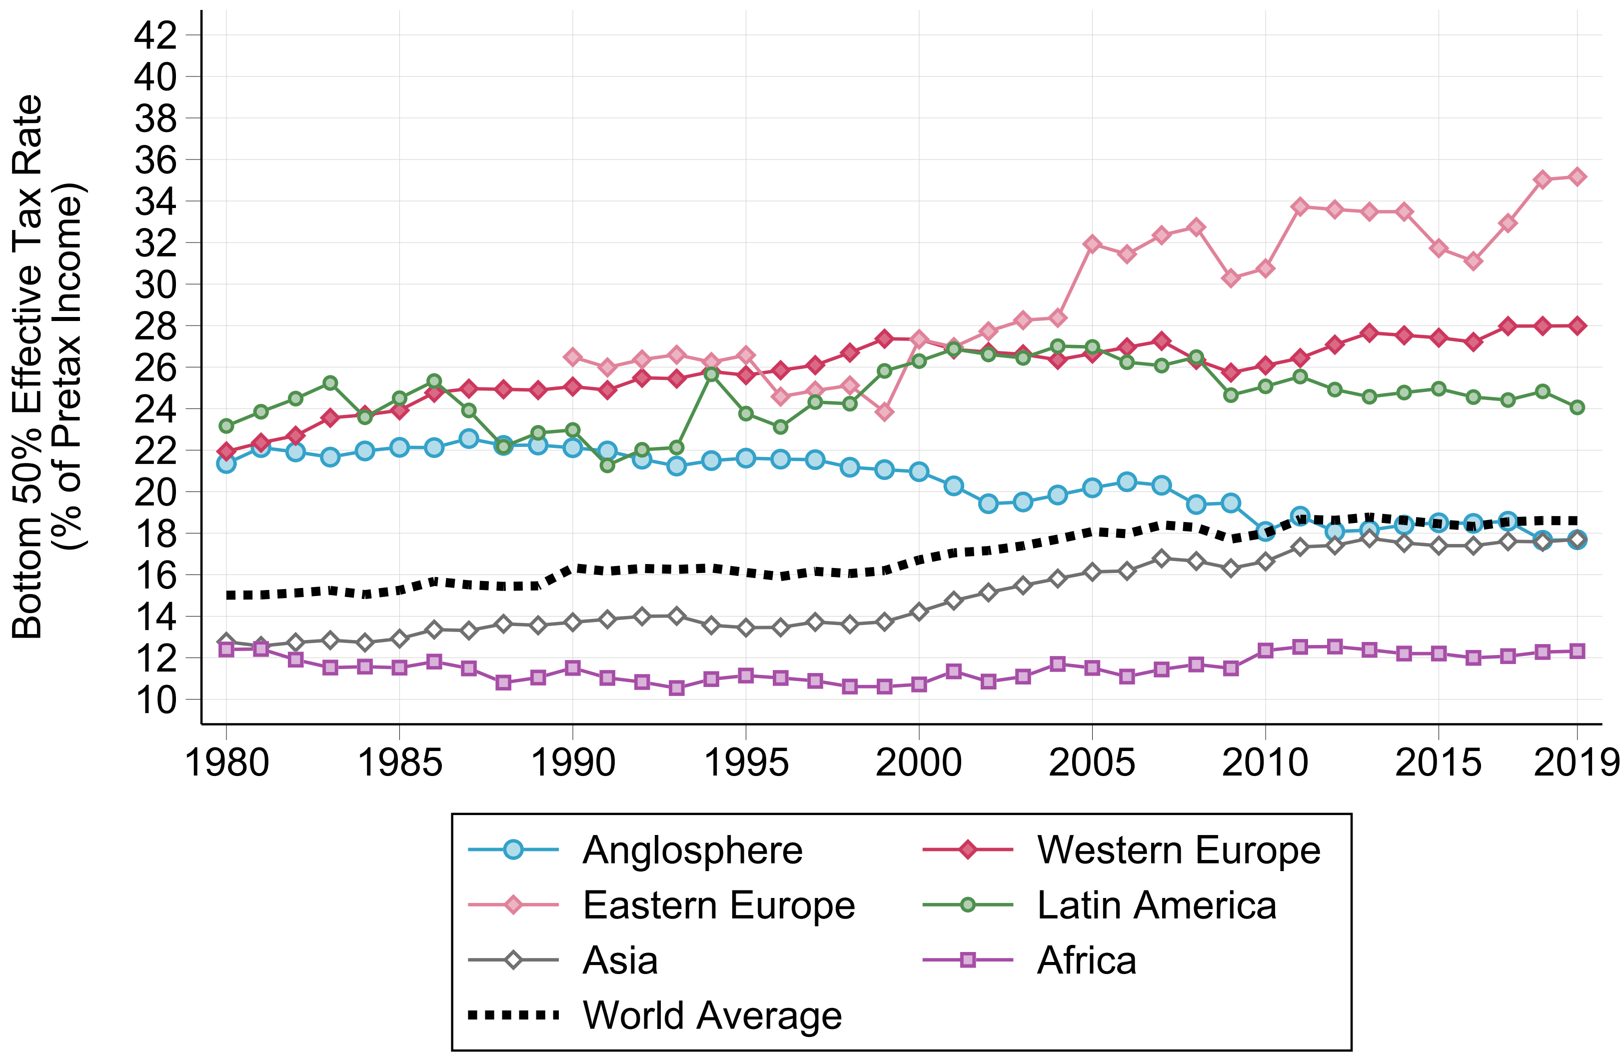

Notes. Authors' elaboration. Excludes social contributions.

Figure C.14: Transfer Progressivity Over the Course of Development:
Total Transfer Received by the Bottom 50% (% of Total Public Spending)

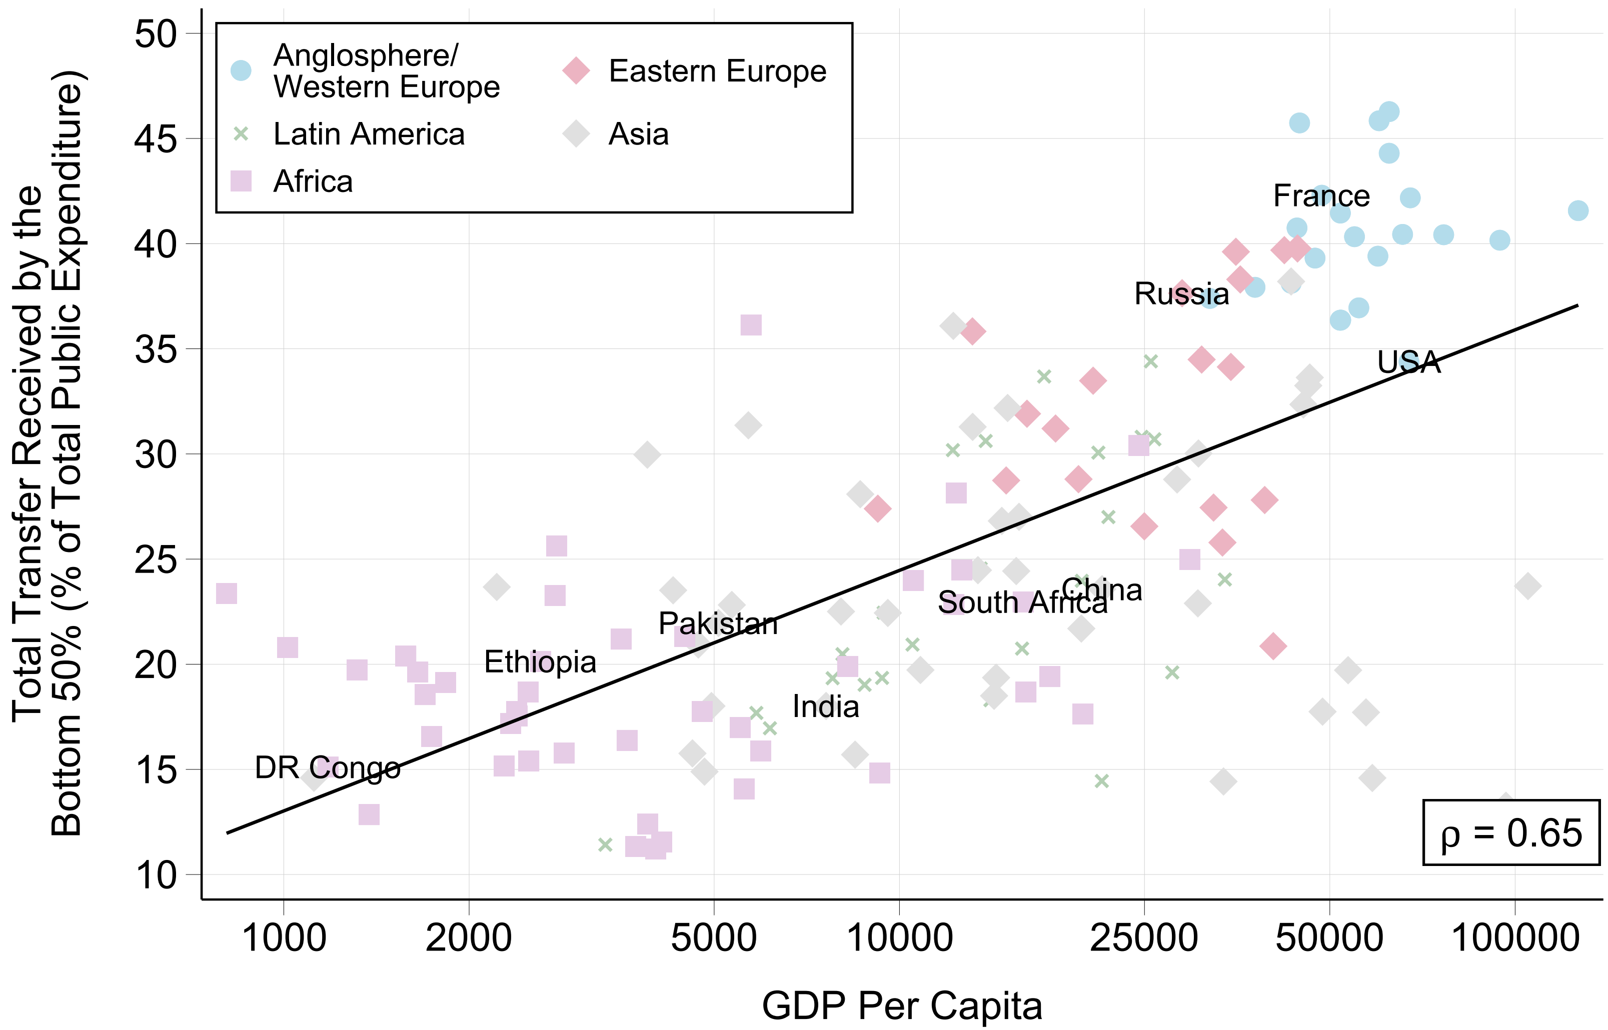

Notes. Total transfer received: sum of all transfers received (before paying taxes), expressed as a share of total government expenditure.

Figure C.15: Extent of Redistribution by Country Income Group, 1980-2019:
Percent Reduction in Top 10% to Bottom 50% Income Ratio, Pretax - Posttax

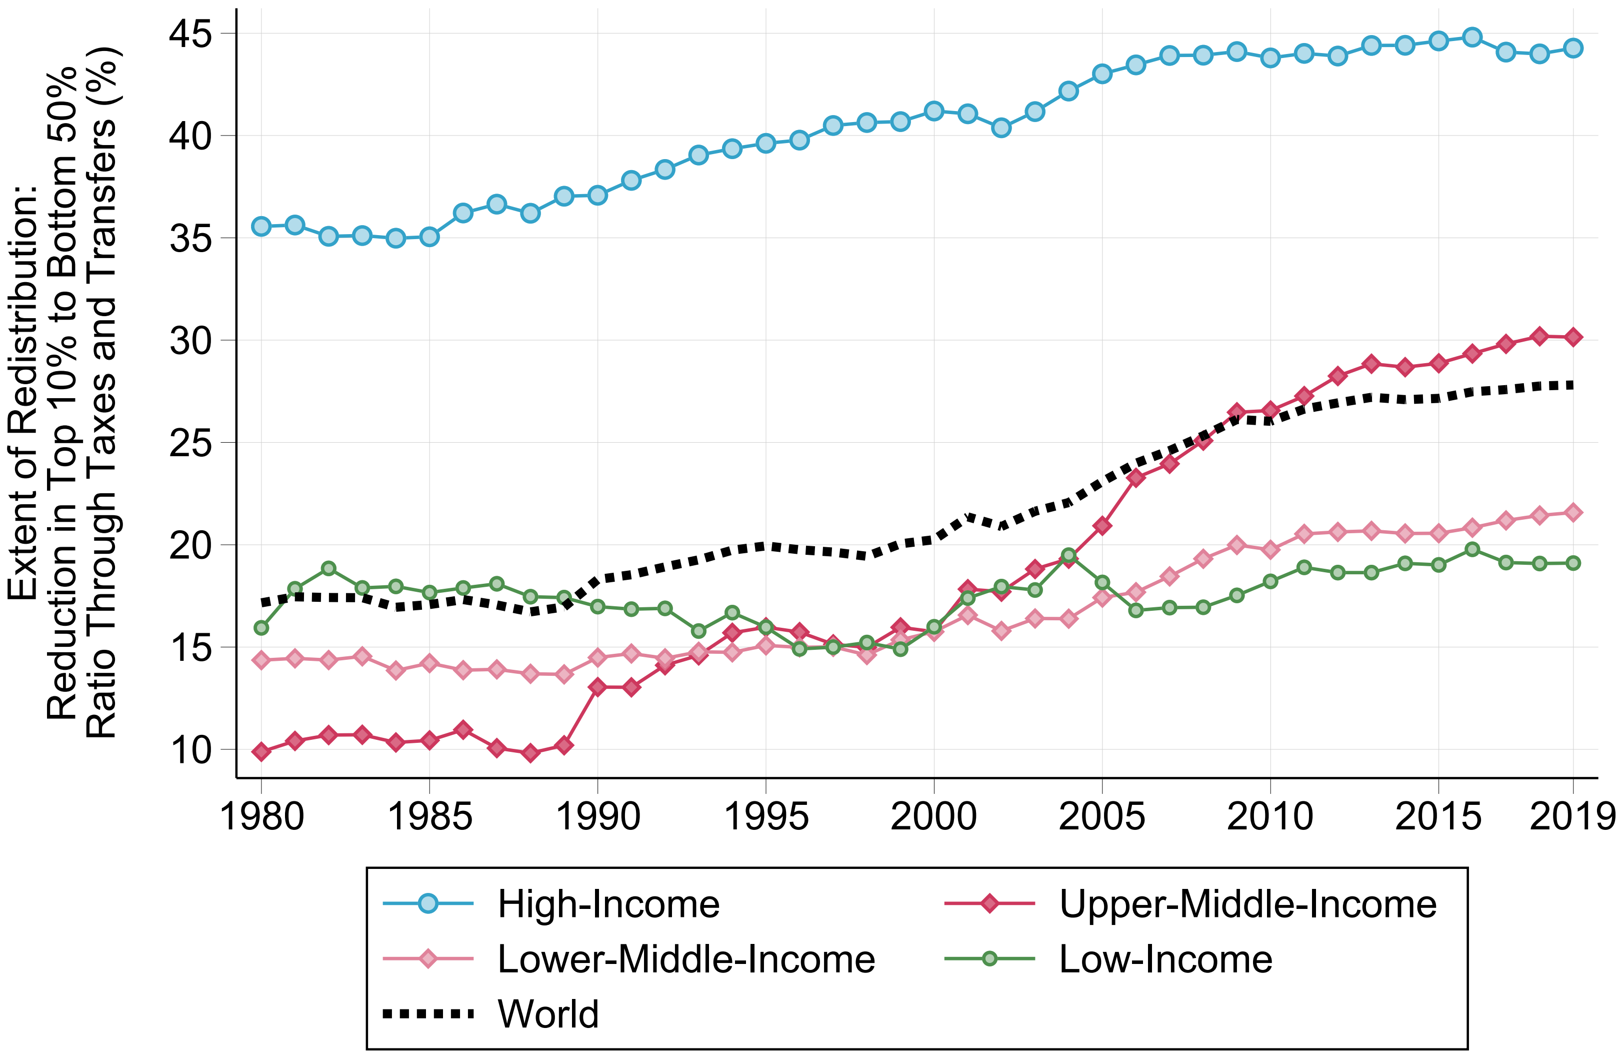

Notes. Country income groups from the World Bank.

Figure C.16: Extent of Redistribution by Country Income Group, 1980-2019:
Net Transfer Received by the Bottom 50% (% of National Income)

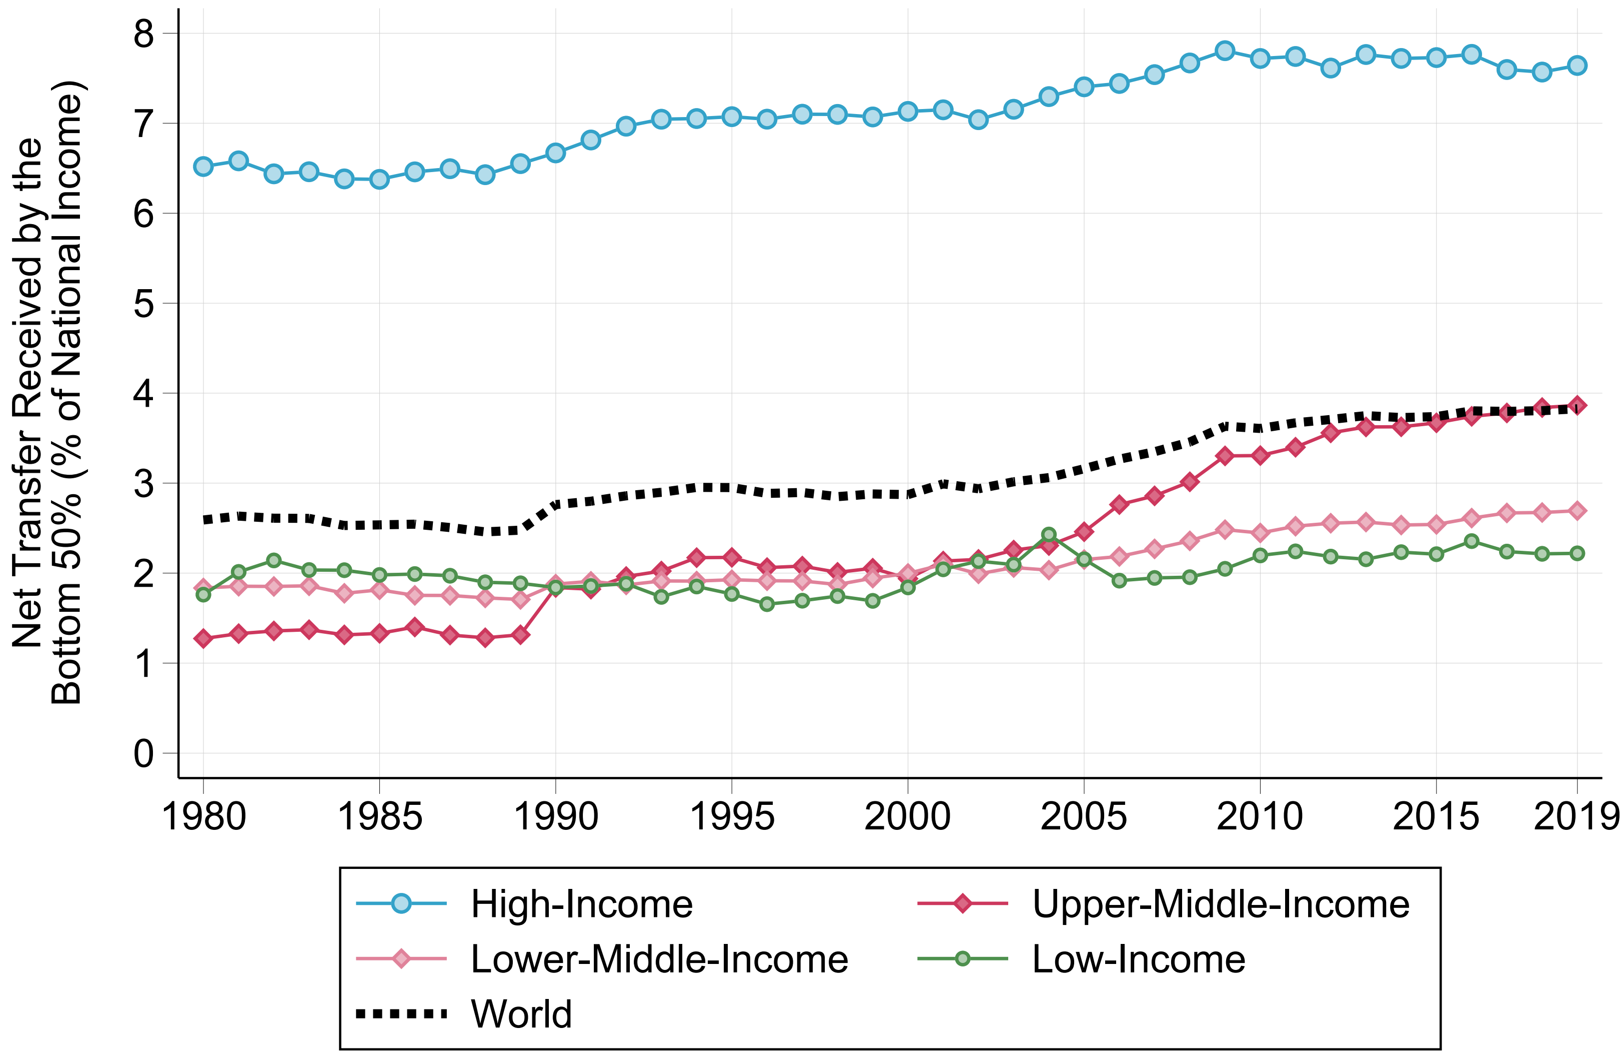

Notes. Country income groups from the World Bank.

Figure C.17: Top 10% to Bottom 50% Income Ratio: Pretax Versus Posttax

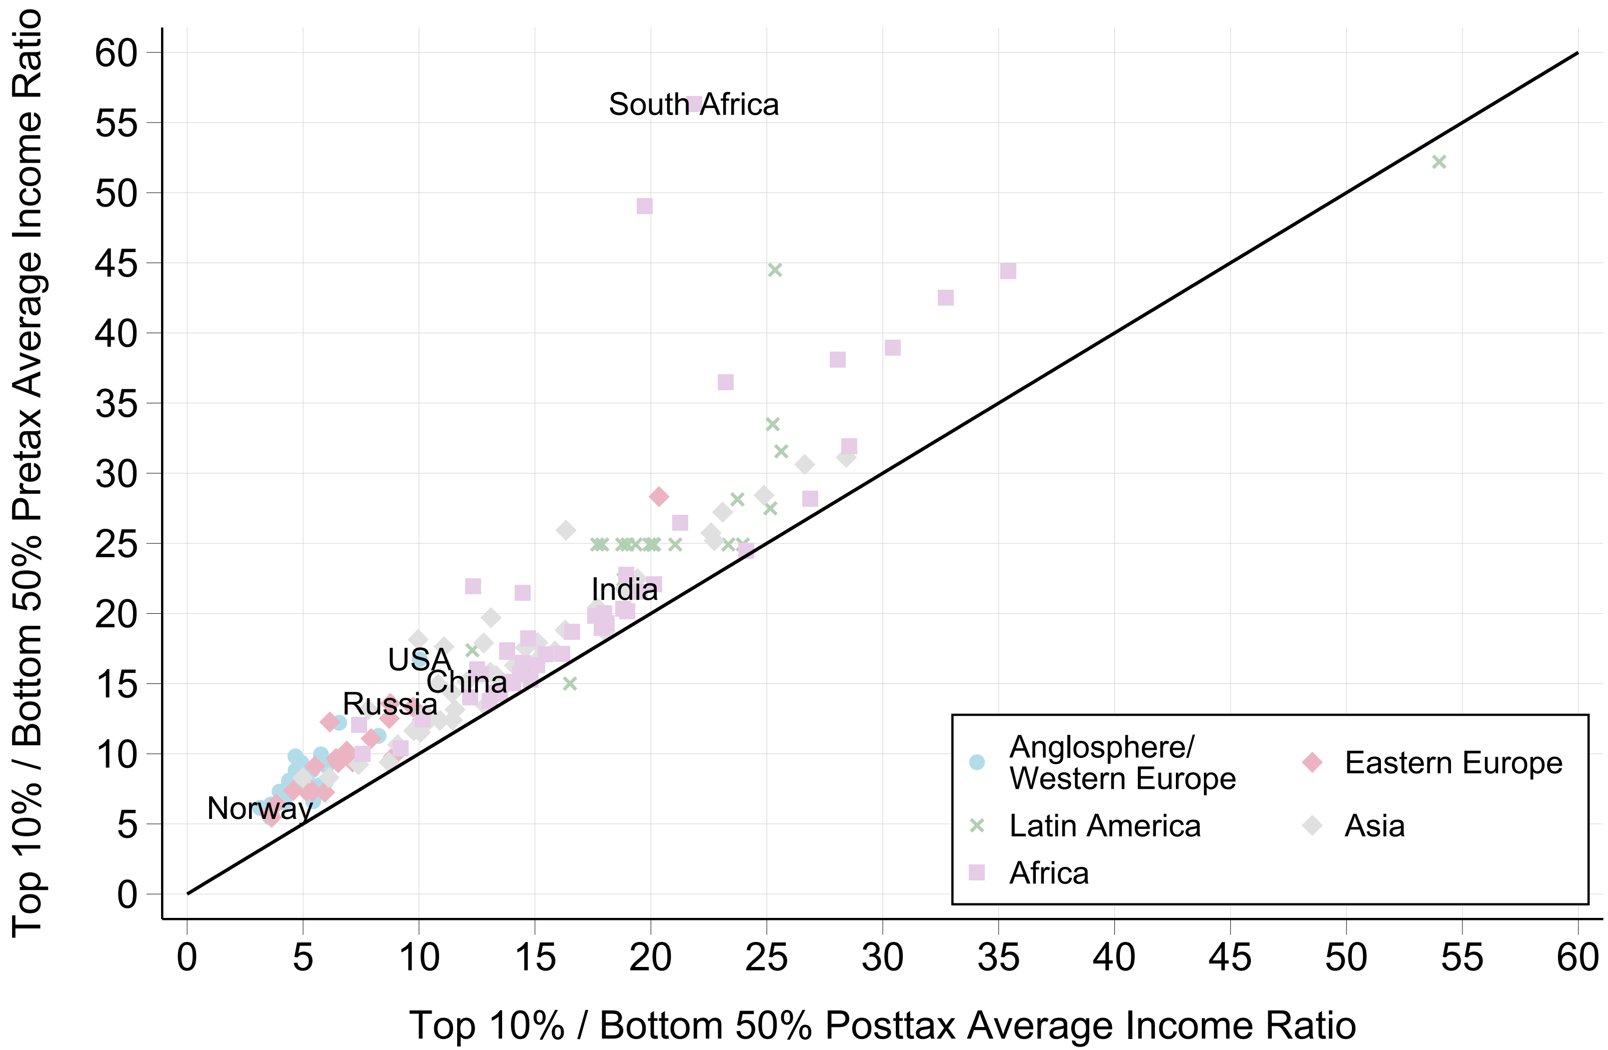

Figure C.18: Predistribution versus Redistribution:
Bottom 50% Pretax versus Posttax National Income Shares by World Region, 2019

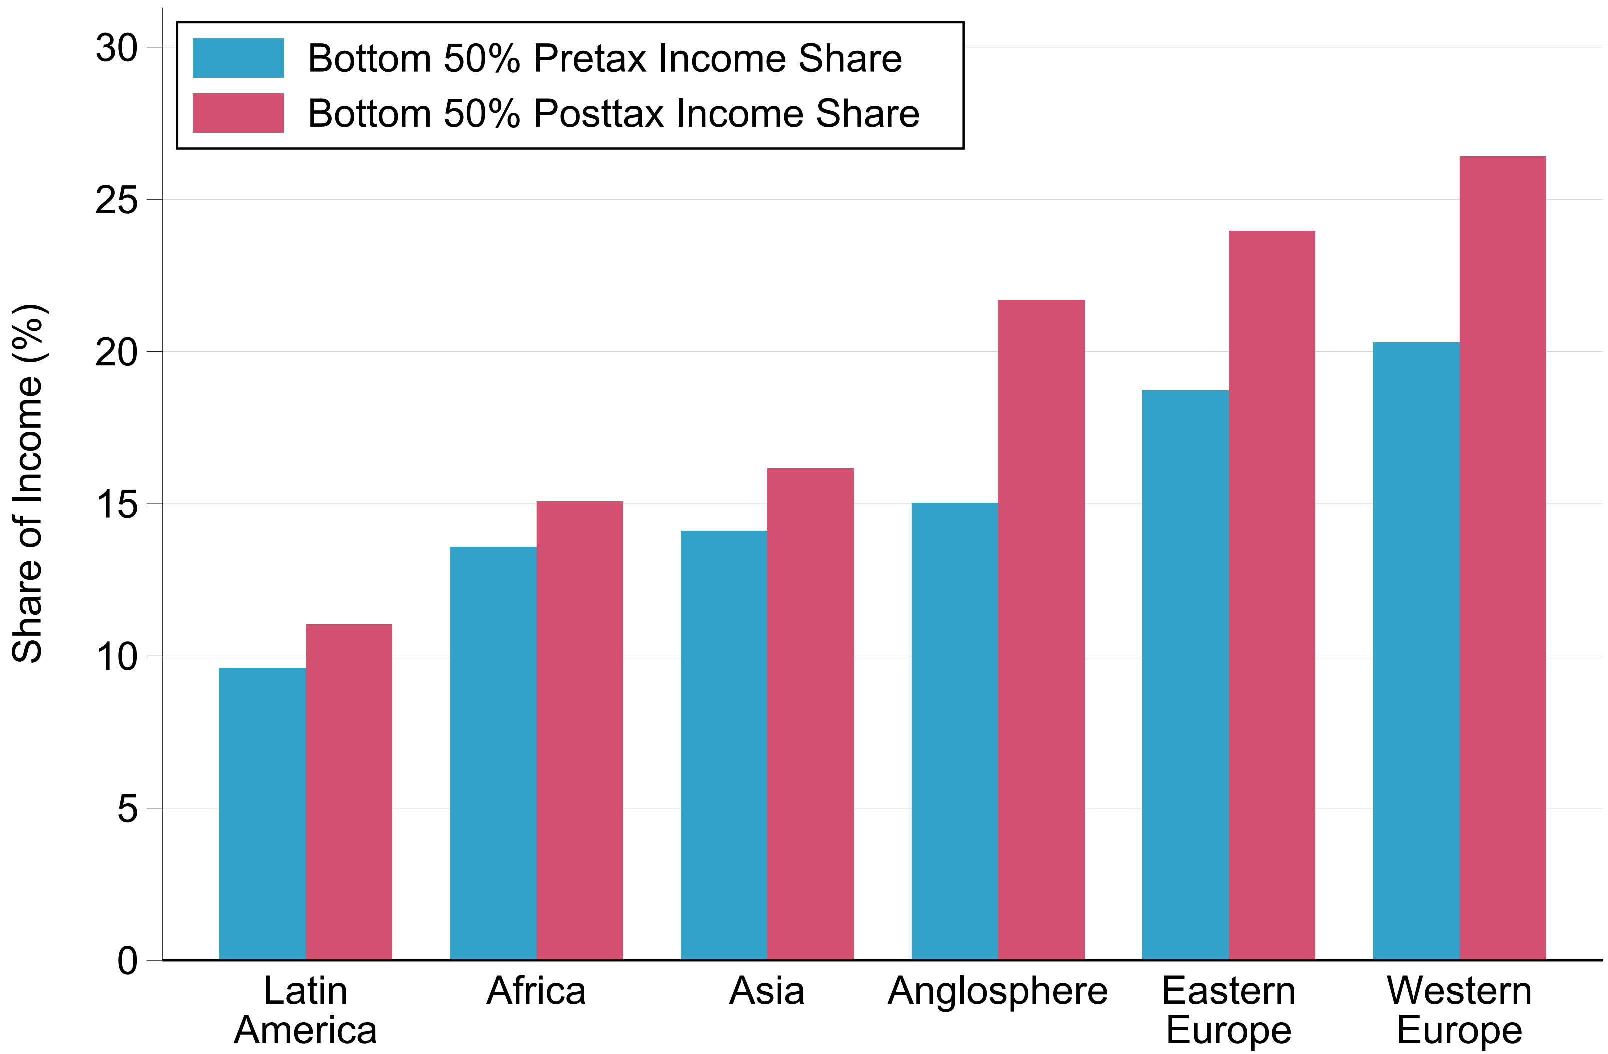

Figure C.19: Top 10% Pretax versus Posttax National Income Shares by World Region

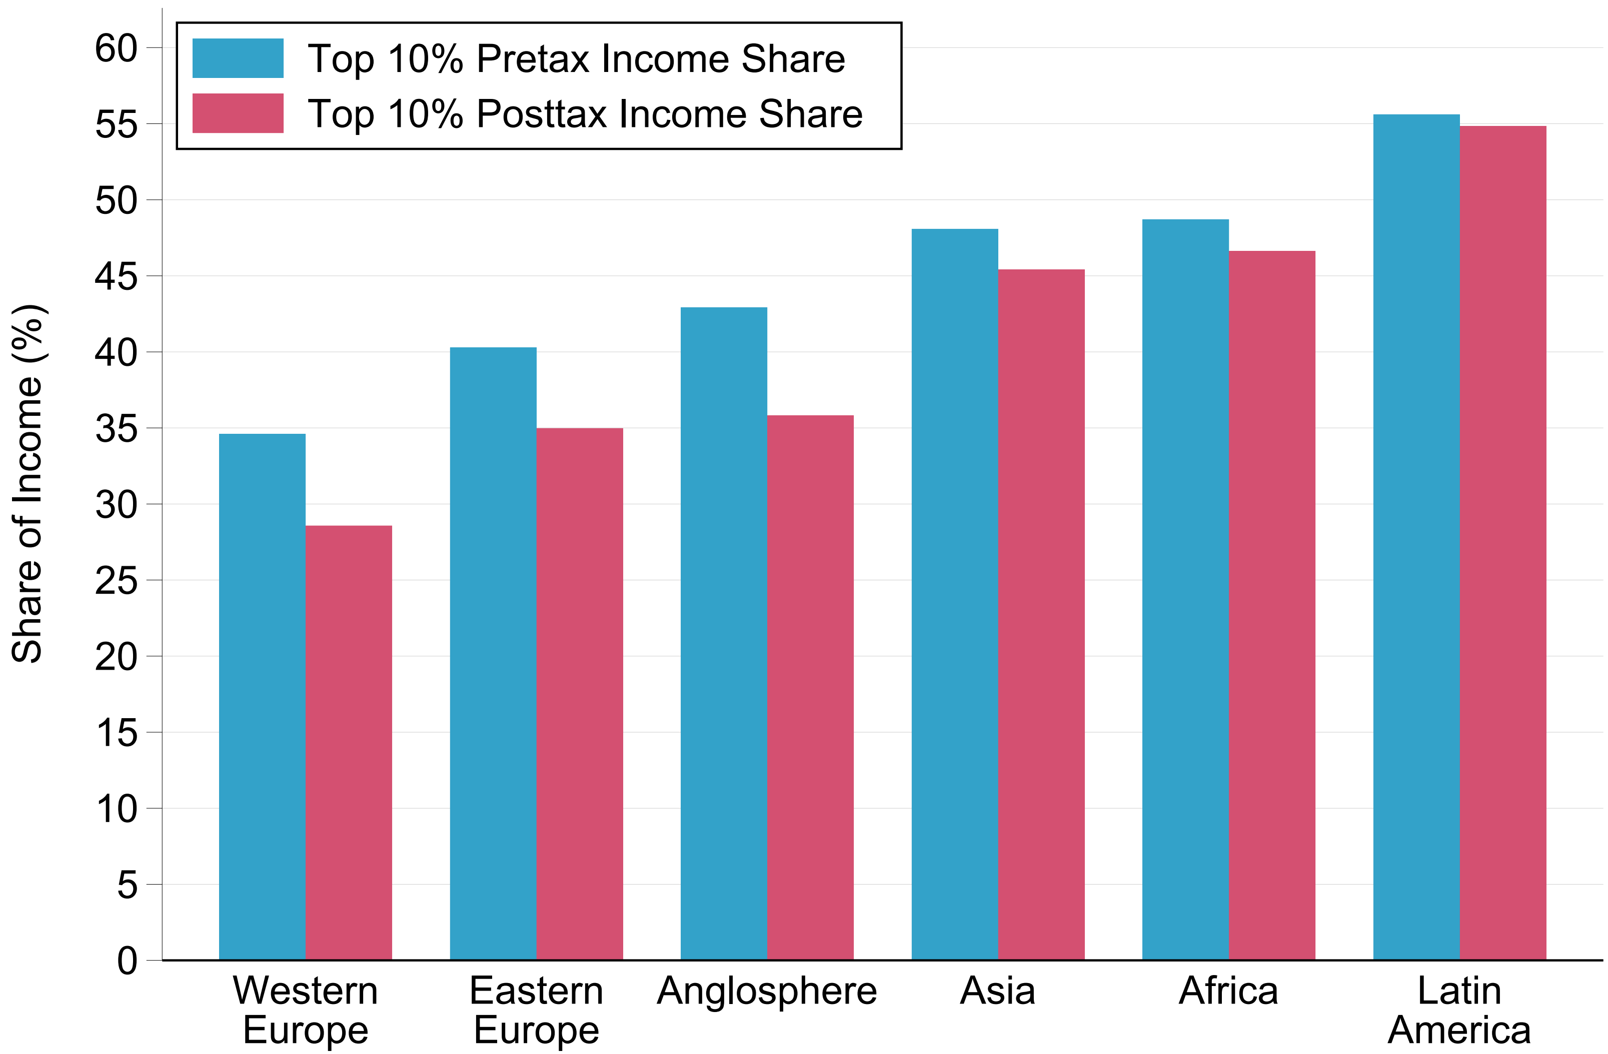

Figure C.20: Top 1% Pretax versus Posttax National Income Shares by World Region

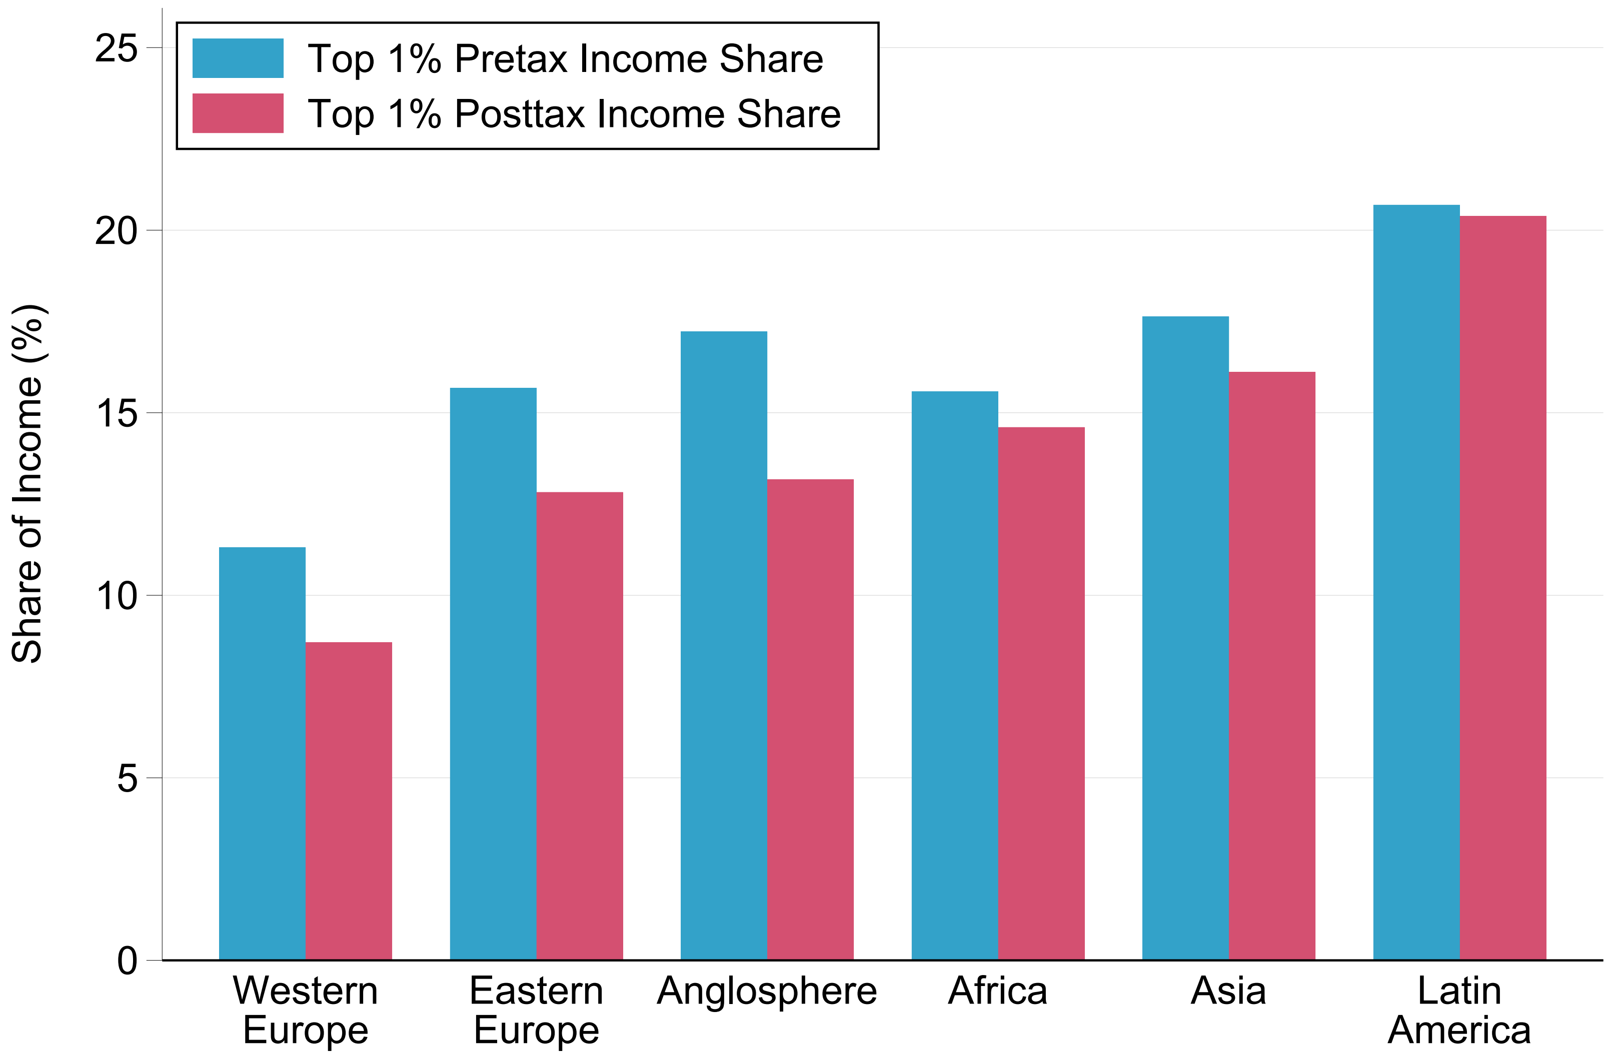

Figure C.21: Top 10% to Bottom 50% Pretax Income Ratio Versus Extent of Redistribution

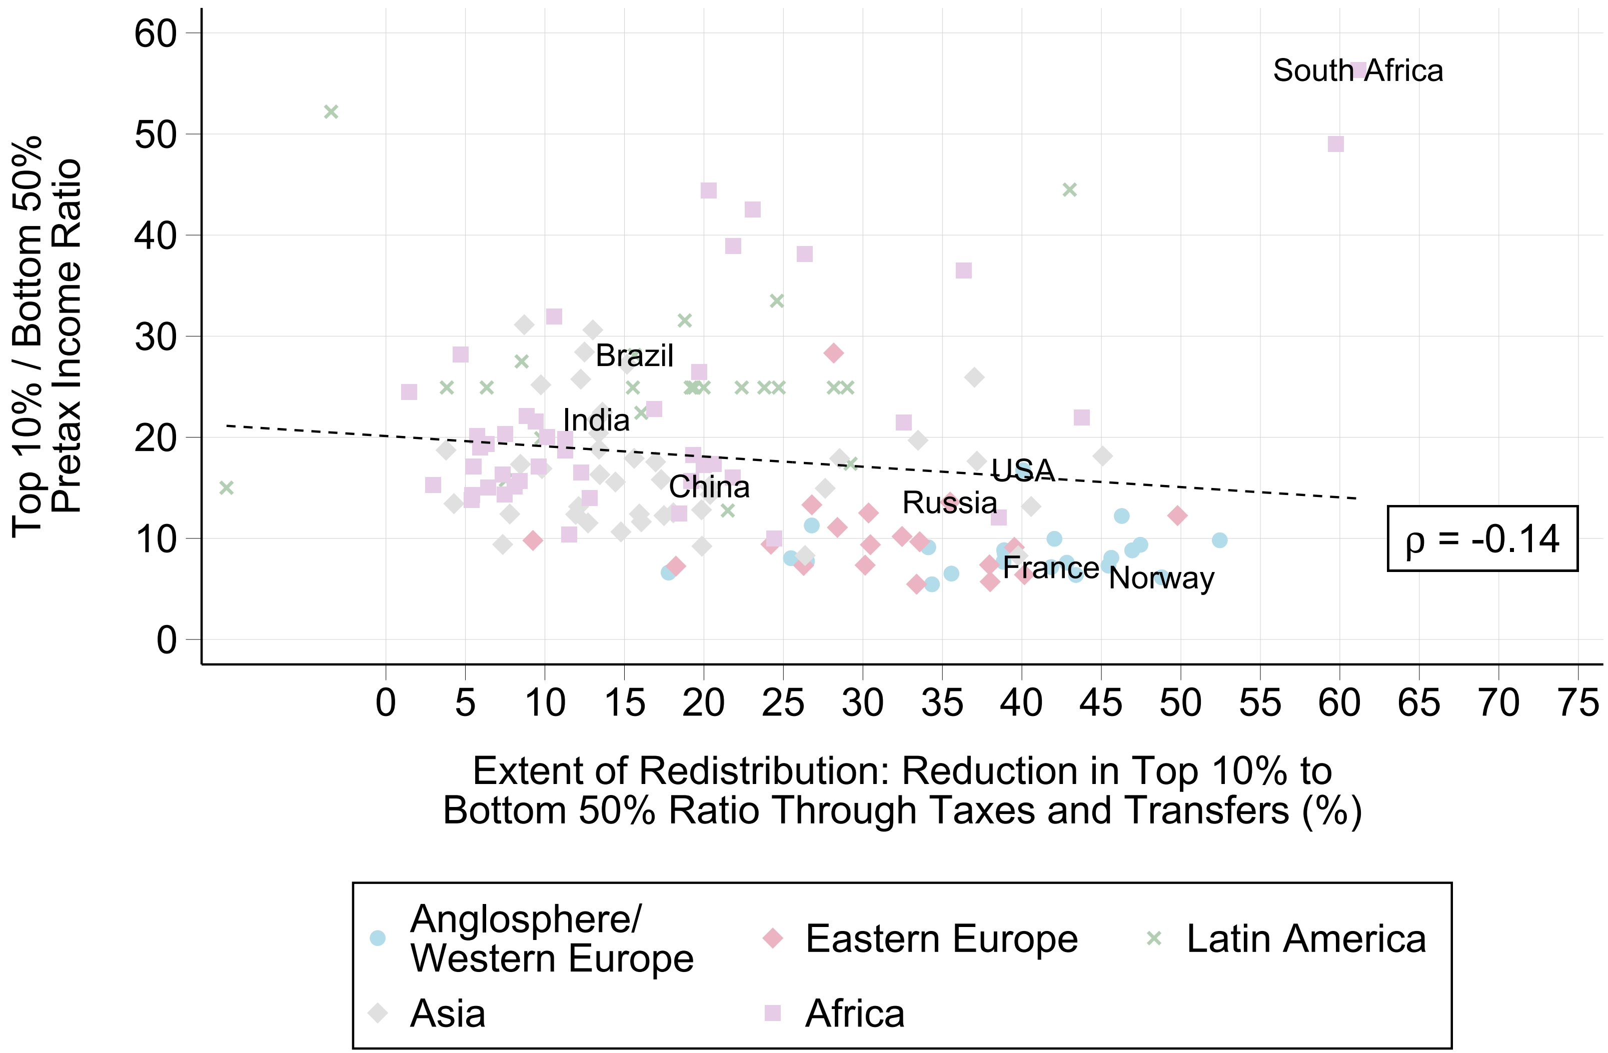

Figure C.22: Top 10% to Bottom 50% Pretax Income Ratio Versus Net Transfer Received by the Bottom 50%

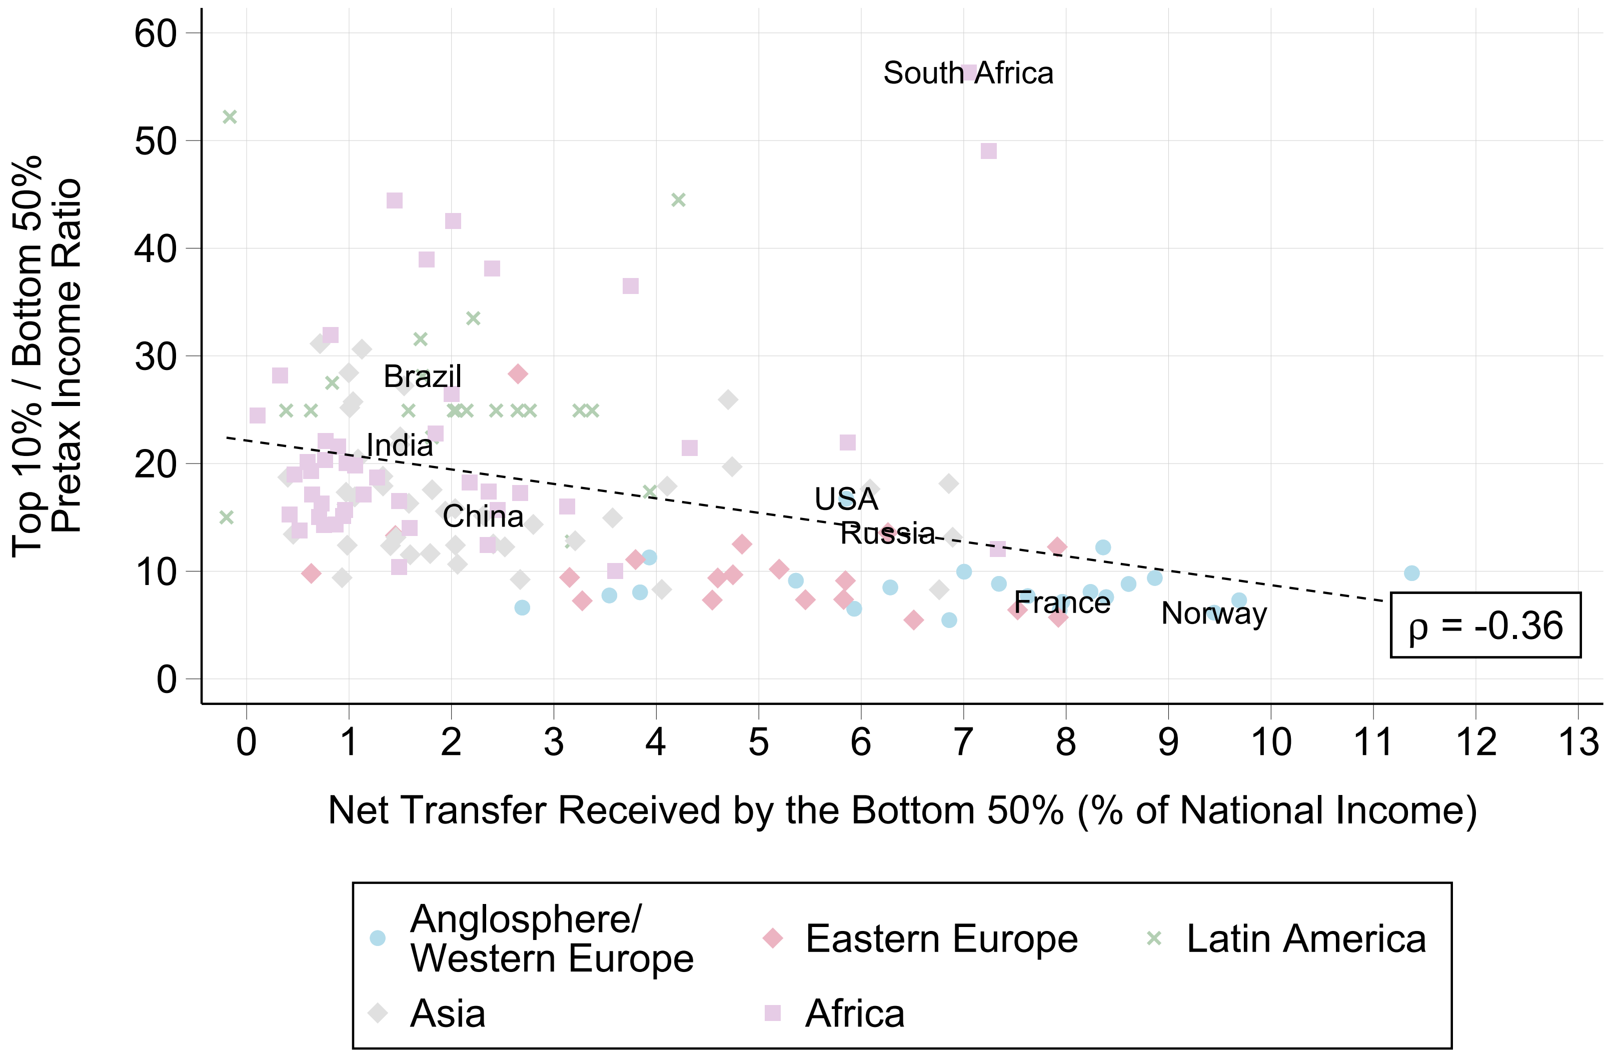

Figure C.23: Validation: Distributional Incidence of Personal Income Tax, United States, 1980-2018)

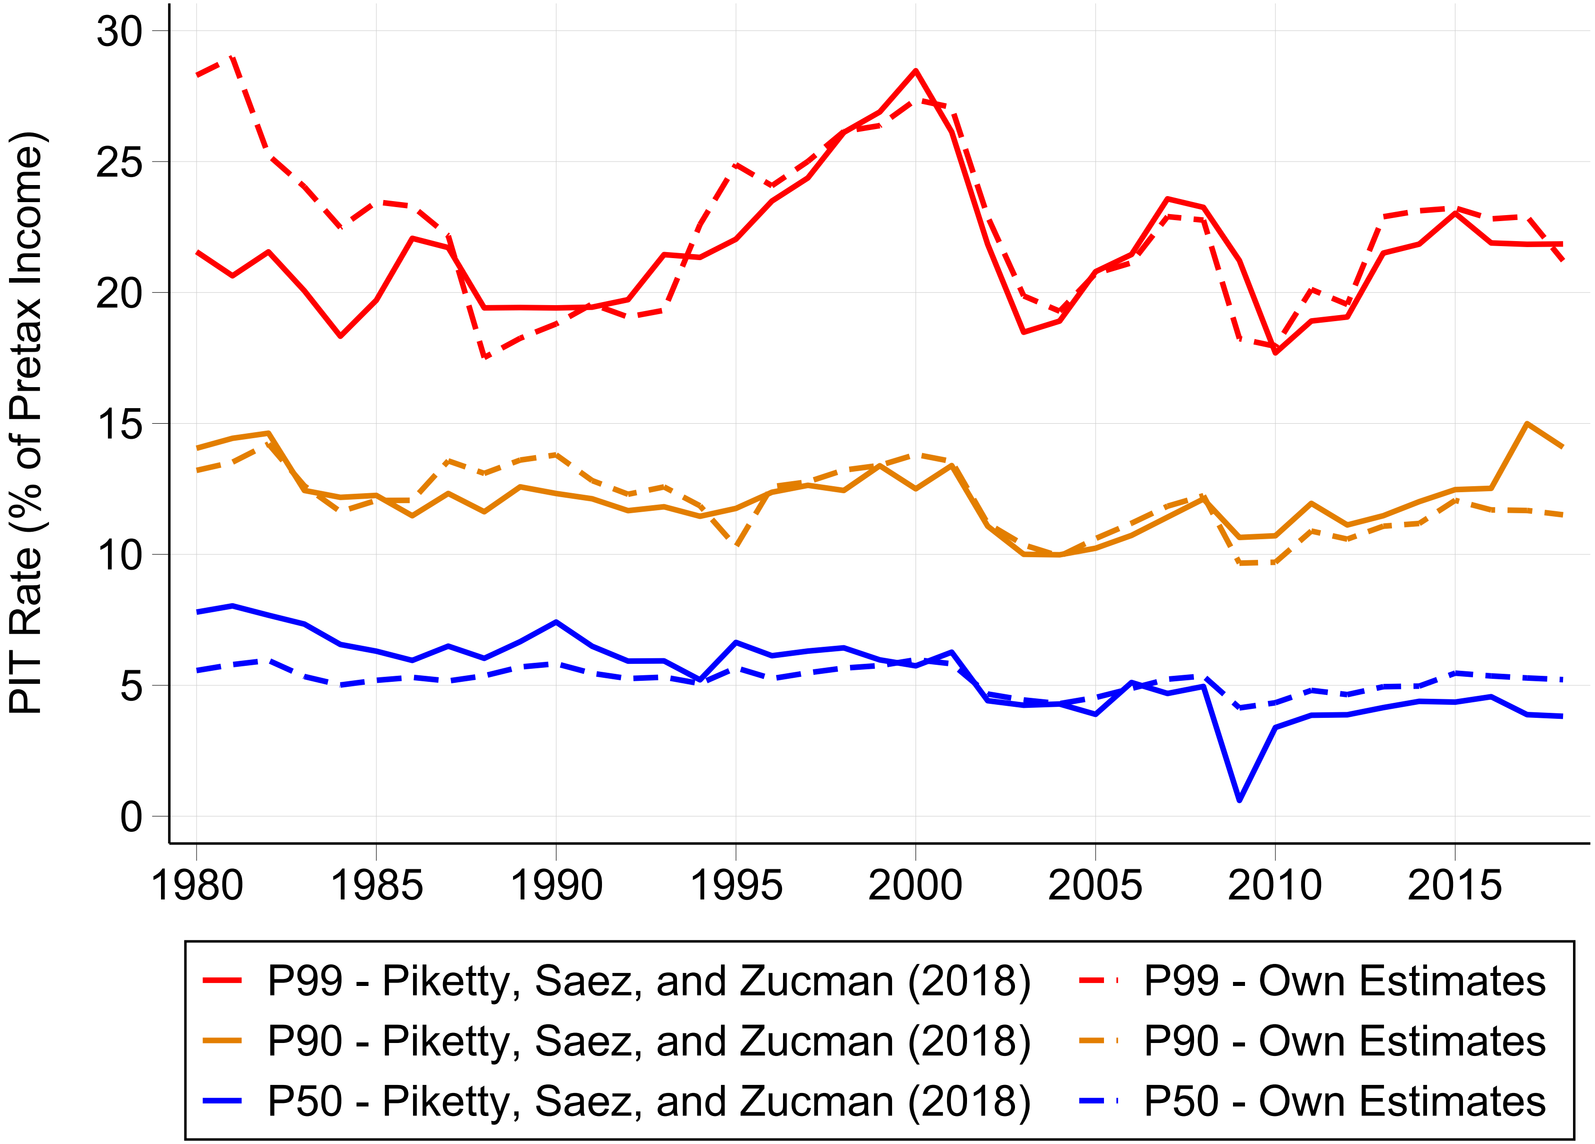

Notes. Authors' elaboration combining own estimates and data from Piketty, Saez, and Zucman (2018).

Table C.1: Extent of Redistribution by World Region: Decomposition by Tax and Transfer, 1980

	World Average	Anglosphere	Western Europe	Eastern Europe	Latin America	Asia	Africa
Personal Income Taxes	3.4%	11.4%	10.4%	3.3%	2.3%	1.9%	2.4%
Corporate Taxes	2.9%	3.4%	2.7%	8.5%	2.7%	2.2%	3.5%
Property & Wealth Taxes	0.3%	1.1%	0.7%	0.5%	0.2%	0.2%	0.1%
Indirect Taxes	-4.7%	-6.0%	-10.5%	-9.9%	-6.2%	-3.2%	-3.4%
Social Contributions	-0.4%	-2.0%	-0.3%	-1.1%	0.6%	-0.2%	-0.8%
All Taxes	3.0%	12.4%	6.0%	5.1%	-0.2%	1.6%	3.5%
Social Assistance	7.4%	14.5%	16.5%	15.4%	19.1%	3.6%	3.7%
Healthcare	6.3%	13.9%	11.1%	6.9%	16.4%	3.4%	6.0%
All Transfers	16.9%	32.1%	28.4%	24.8%	35.8%	10.3%	15.7%

Notes. Population-weighted averages of indicators in each country. The table reports the negative of the percent change in the top 10% to bottom 50% income ratio before and after removing the corresponding tax or adding to corresponding transfer to pretax income. For instance, the top row reports the percent reduction in inequality resulting from removing personal income taxes from individual incomes. Positive values indicate that the corresponding tax or transfer reduces inequality.

C.4 Alternative Measures of Tax Progressivity

Figure C.24: Relative Tax Progressivity Around the World:
Ratio of Top 10% to Bottom 50% Effective Tax Rates, 2019

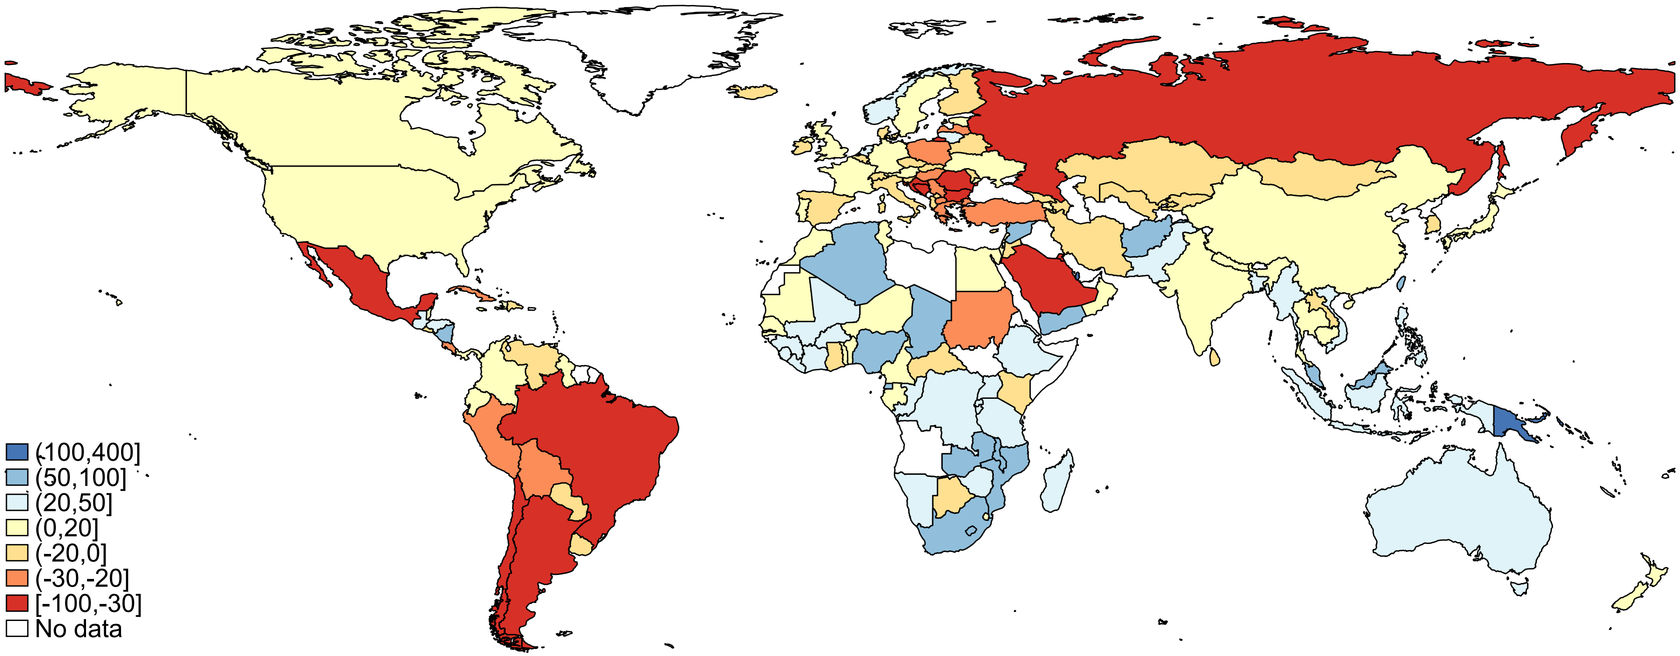

Notes. Includes social contributions.

Figure C.25: Normalized Tax Progressivity Around the World:
Percent Reduction in Top 10% to Bottom 50% Average Income Ratio (Pretax versus Net-of-tax Income)

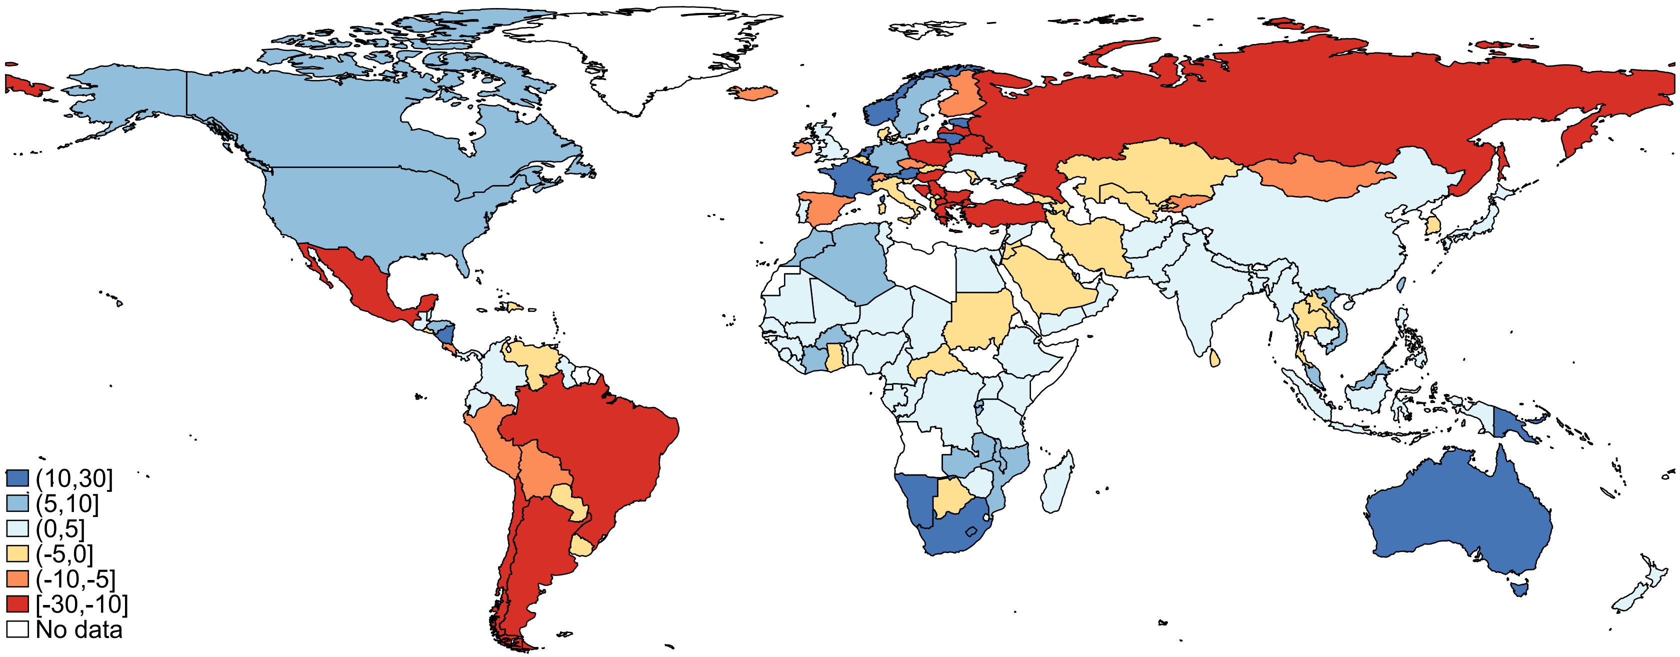

Notes. Includes social contributions.

Figure C.26: Relative Tax Progressivity Over the Course of Development:
Percent Difference Between Top 10% and Bottom 50% Effective Tax Rates, 2019

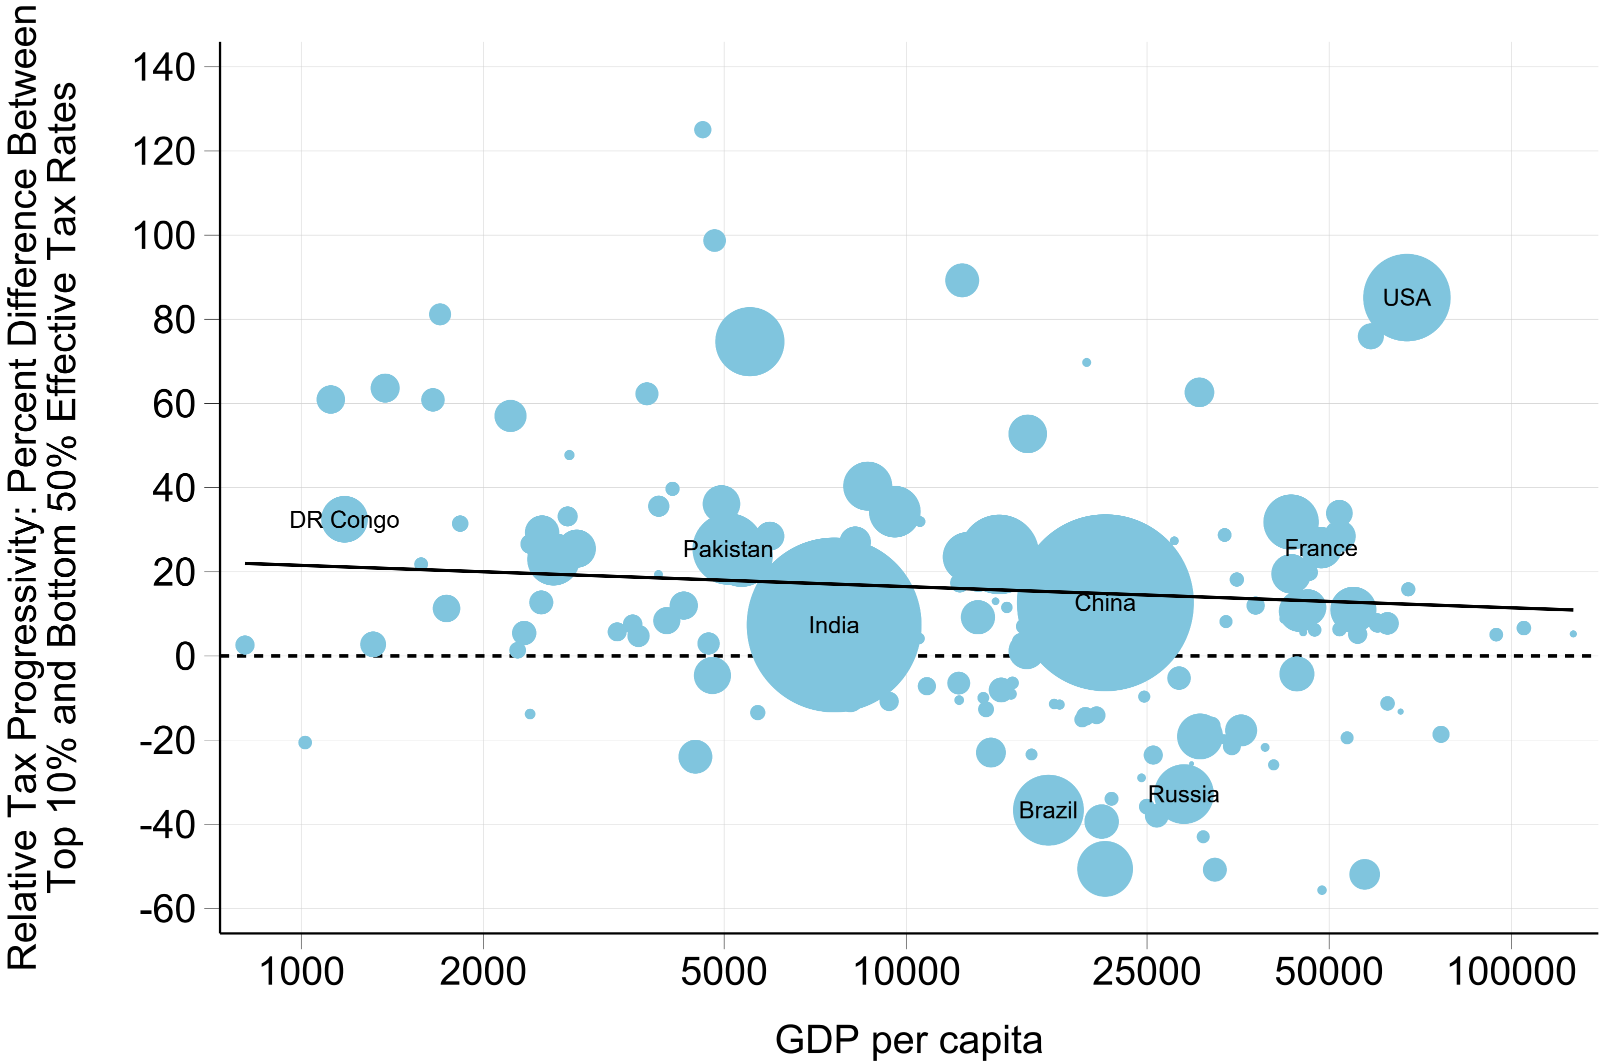

Notes. Excludes social contributions.

Figure C.27: Relative Tax Progressivity (Including Social Contributions) Over the Course of Development: Percent Difference Between Top 10% and Bottom 50% Effective Tax Rates, 2019

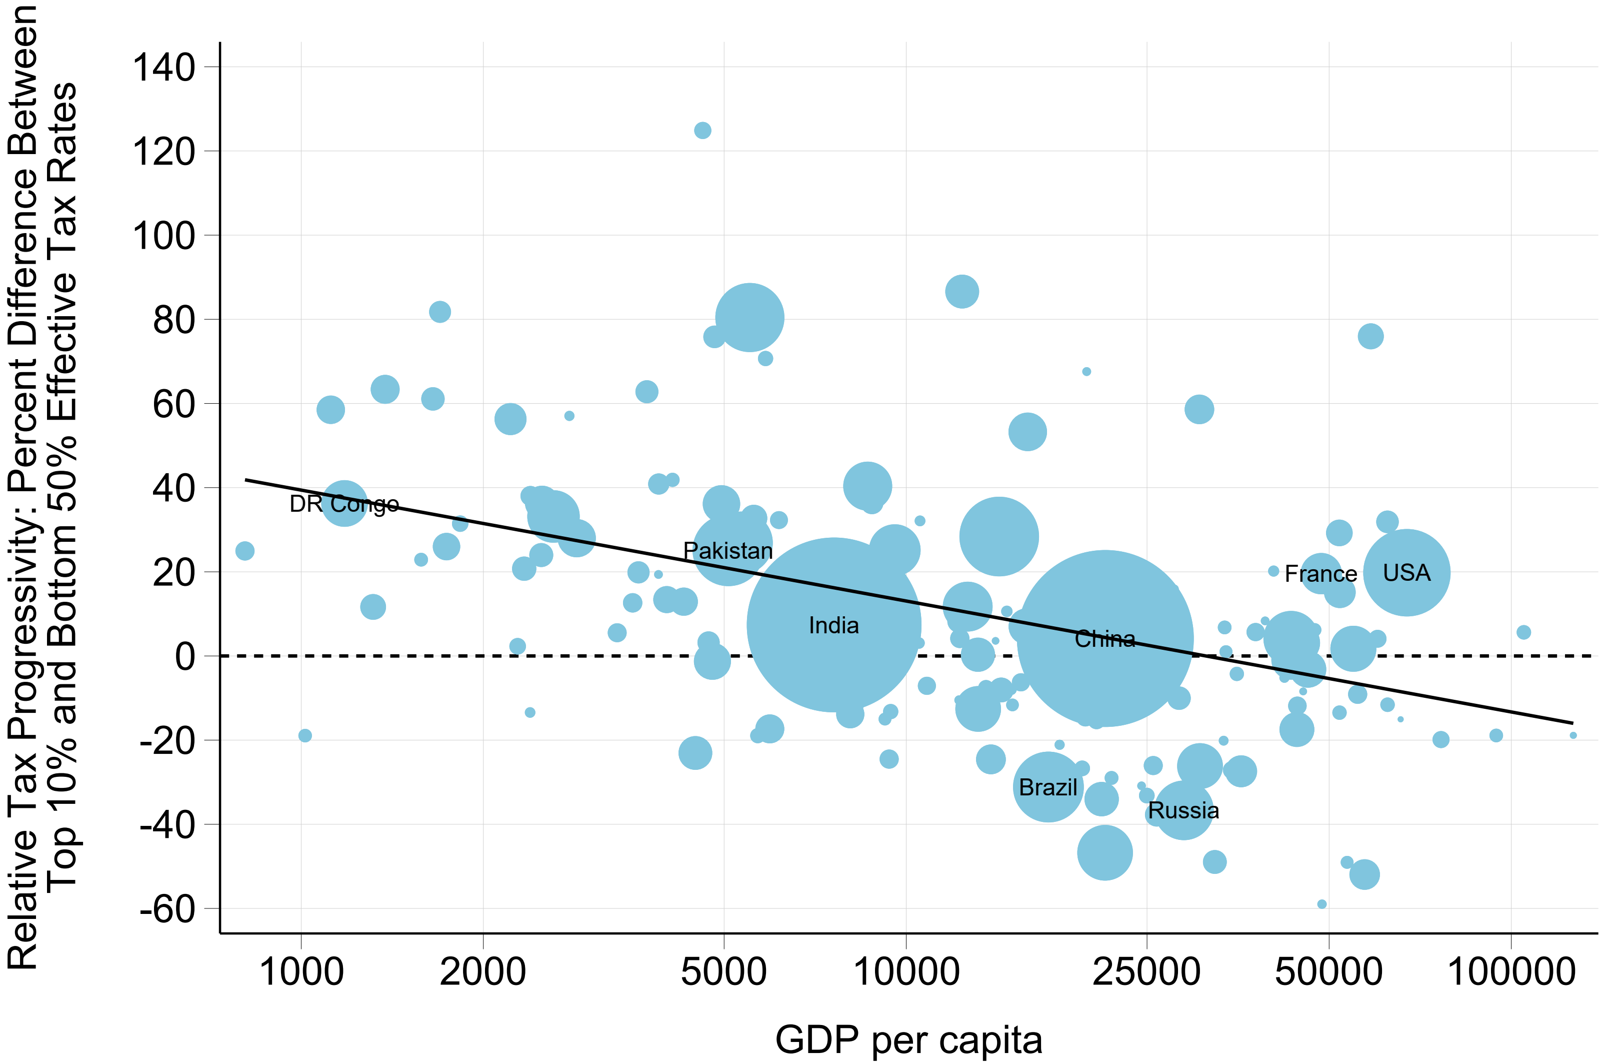

Notes. Includes social contributions.

Figure C.28: Absolute Tax Progressivity (Including Social Contributions) Over the Course of Development: Percent Reduction in Top 10% to Bottom 50% Average Income Ratio (Pretax versus Net-of-tax Income)

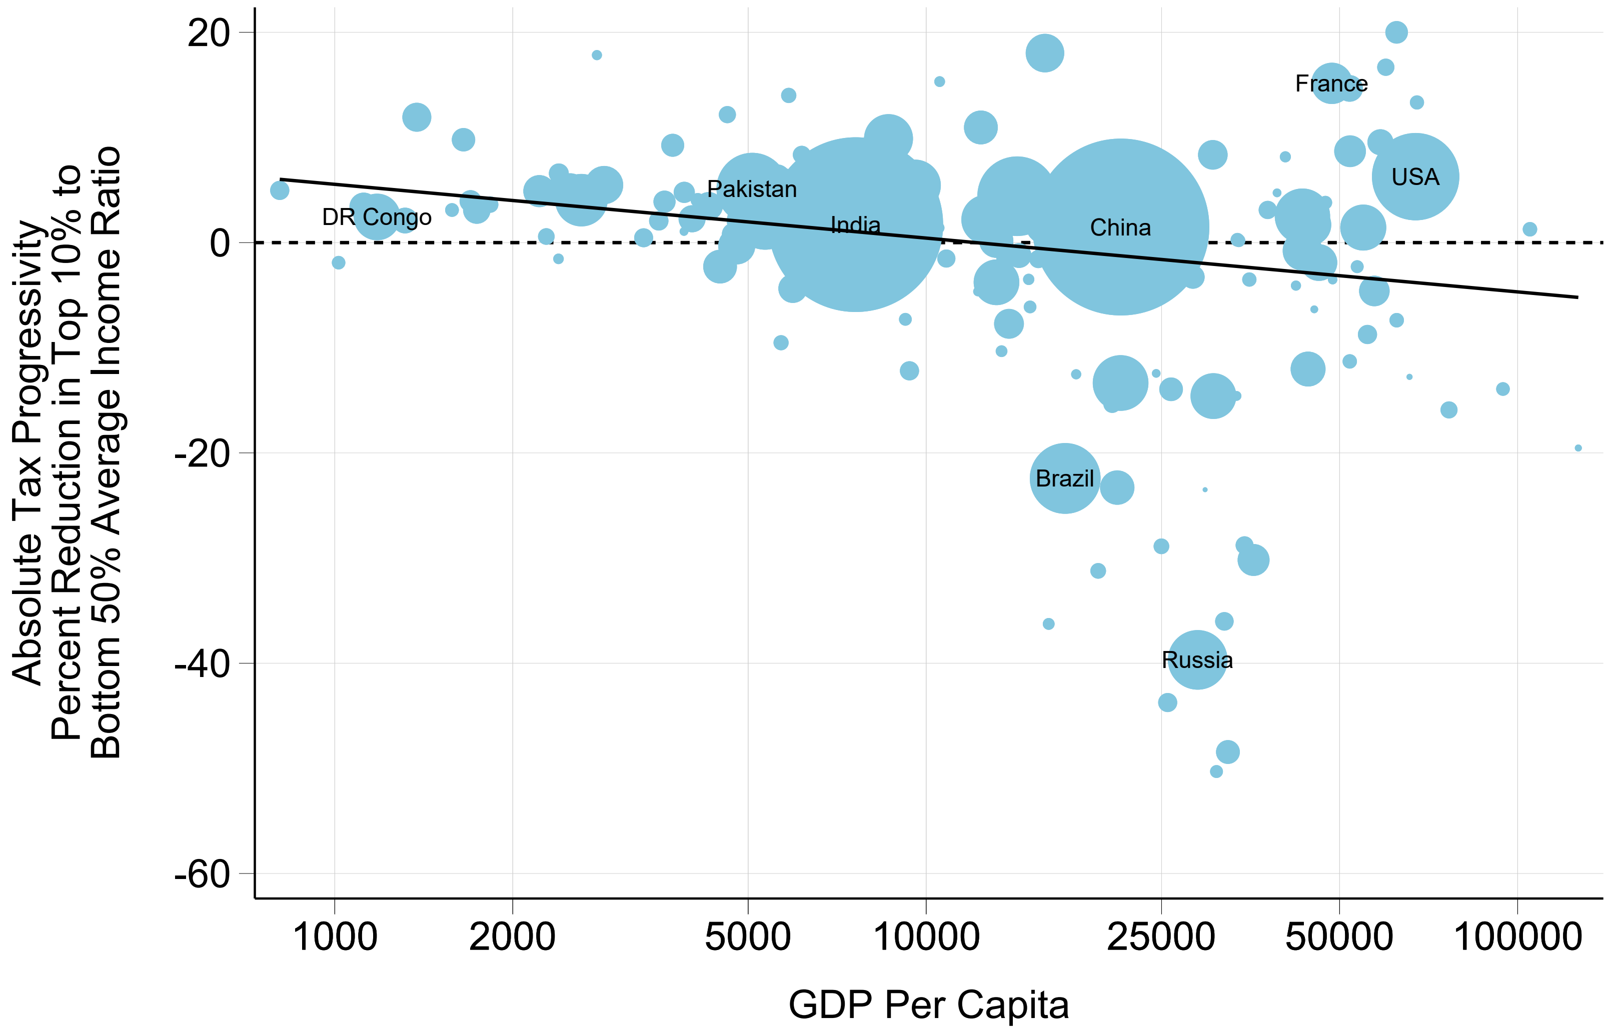

Notes. Includes social contributions.

Figure C.29: Normalized Tax Progressivity Over the Course of Development:
Percent Reduction in Top 10% to Bottom 50% Average Income Ratio (Pretax versus Net-of-tax Income)

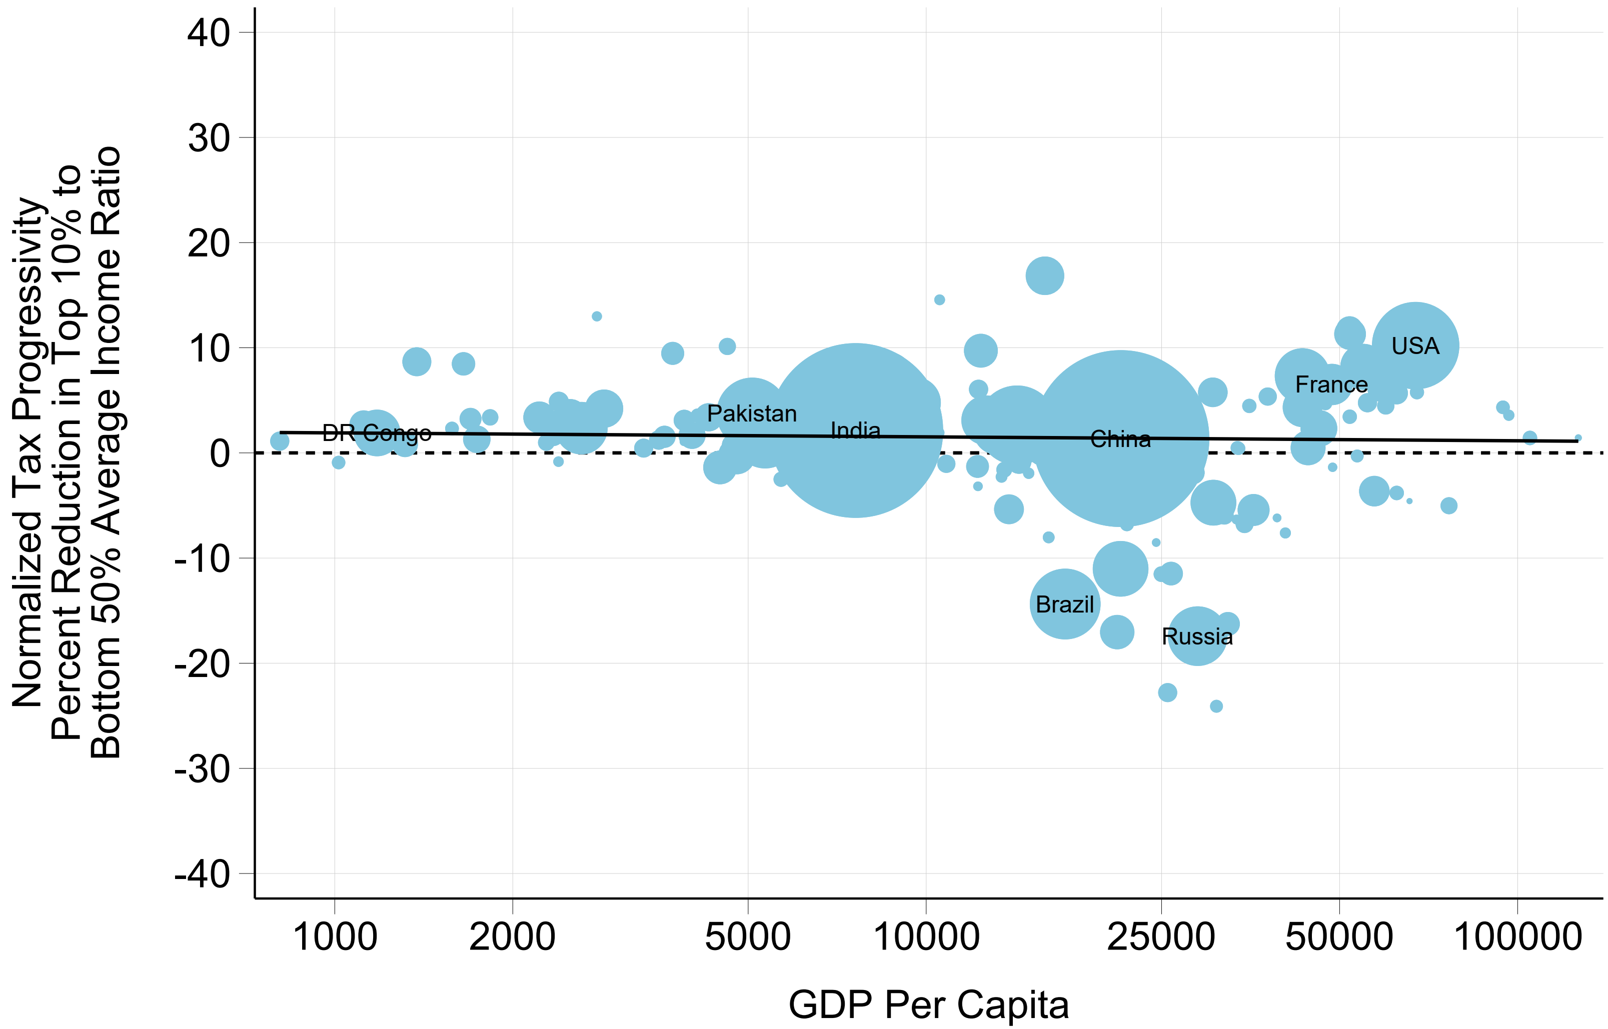

Notes. Excludes social contributions.

Figure C.30: Relative Tax Progressivity by World Region, 1980-2019

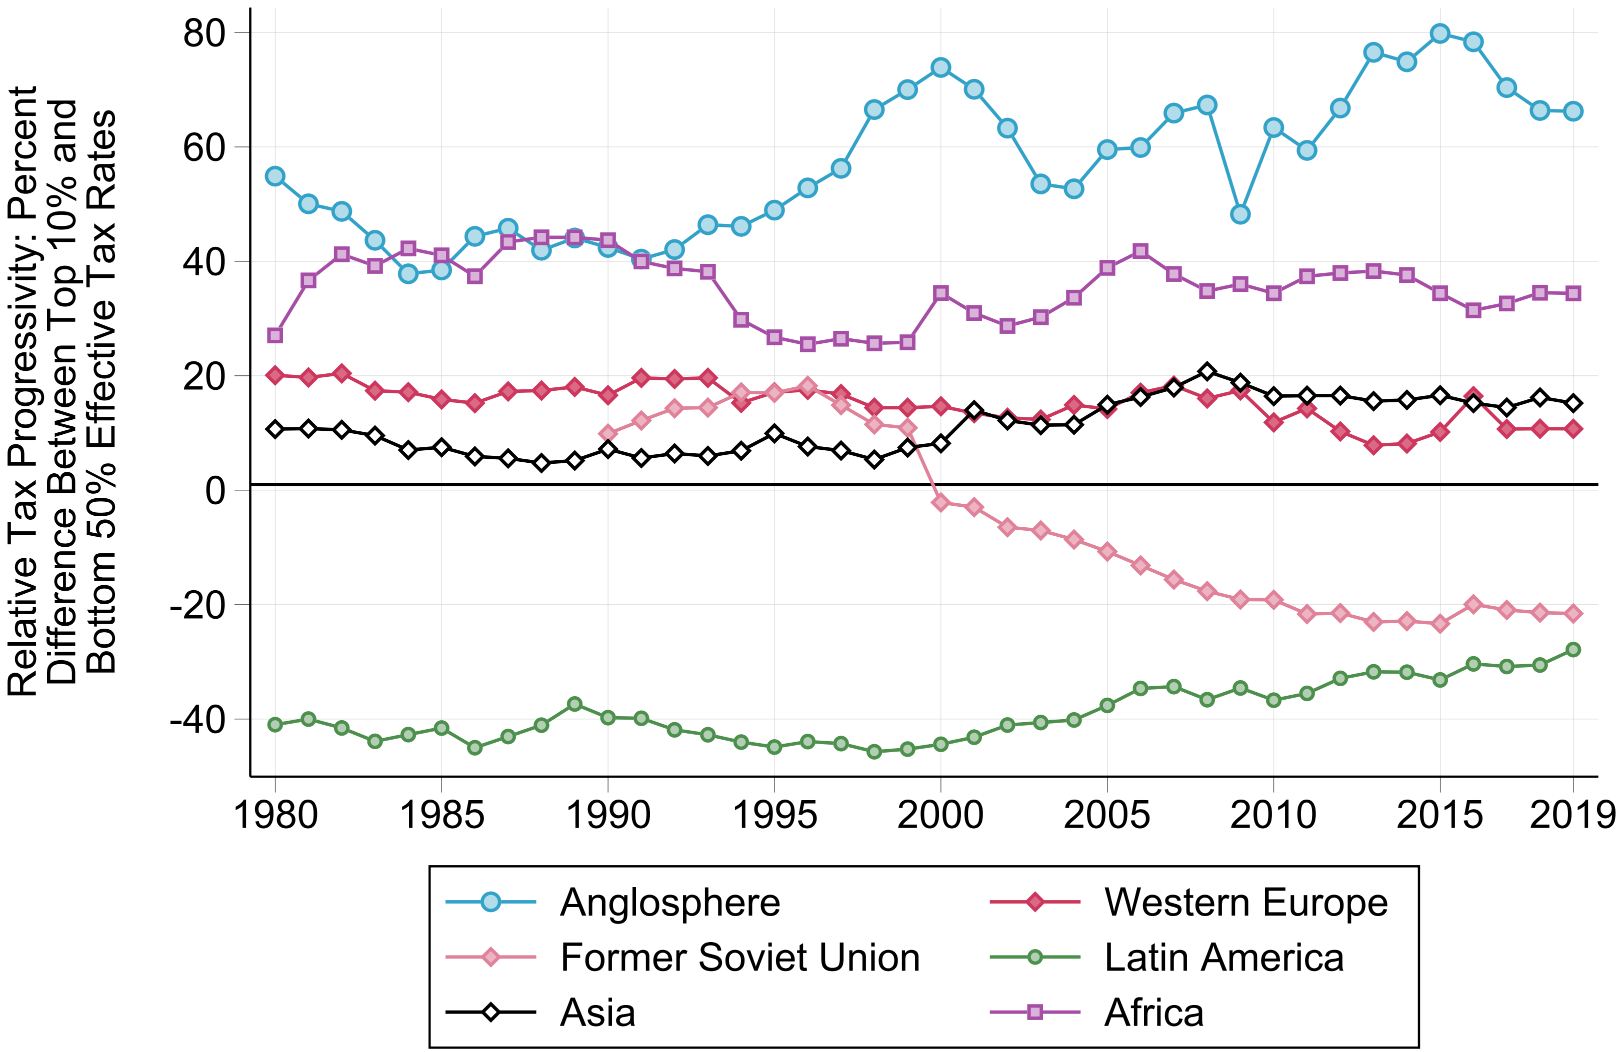

Notes. Excludes social contributions.

Figure C.31: Relative Tax Progressivity (Including Social Contributions) by World Region, 1980-2019

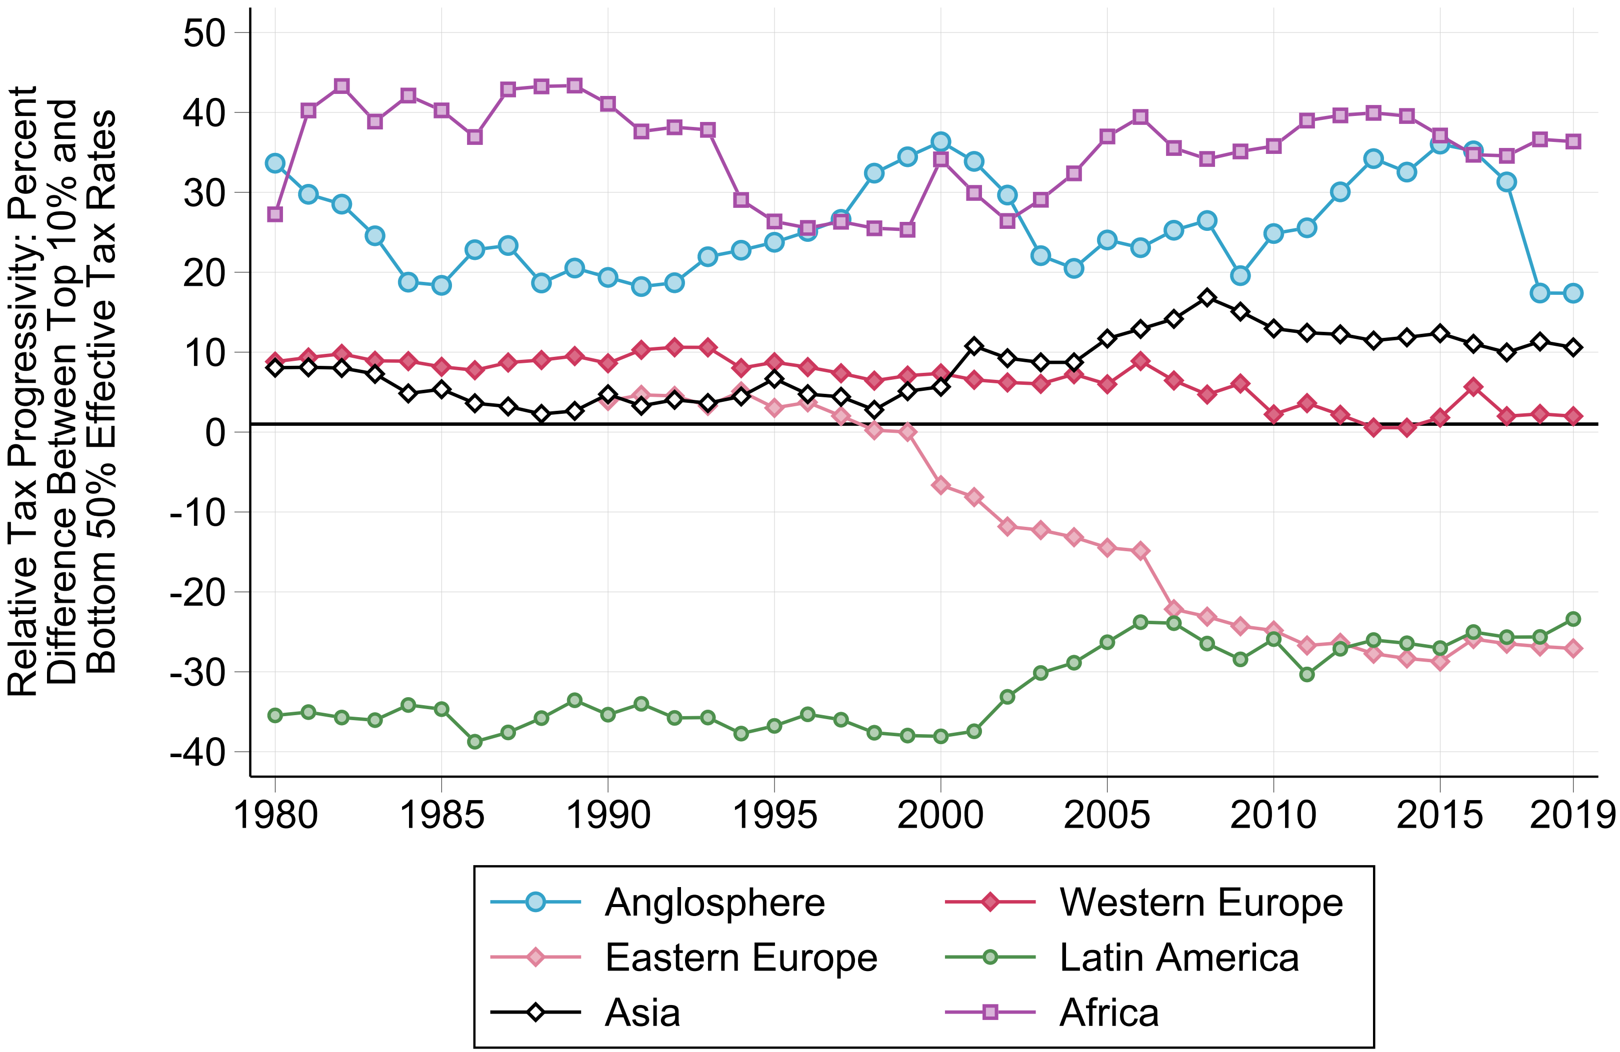

Notes. Includes social contributions.

Figure C.32: Normalized Tax Progressivity by World Region, 1980-2019

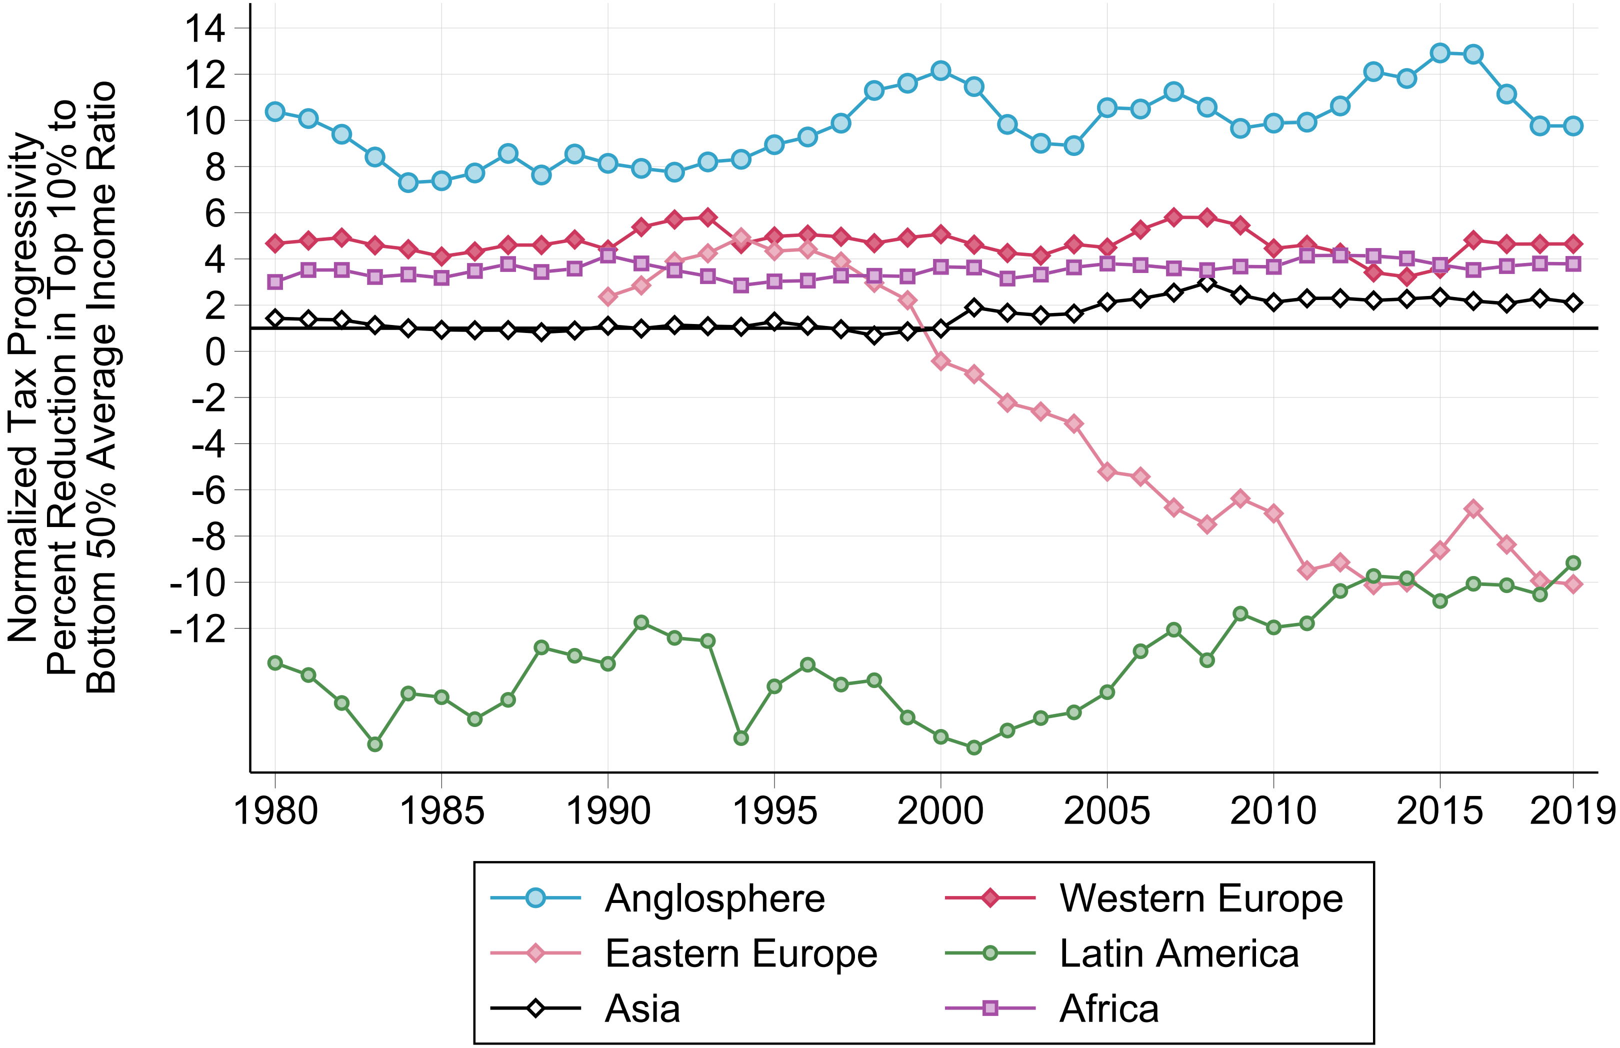

Notes. Excludes social contributions.

C.5 Results With Education Distributed Based on School Attendance

Figure C.33: Government Transfers Received by Income Group and World Region, 2019
Education Distributed Based on School Attendance

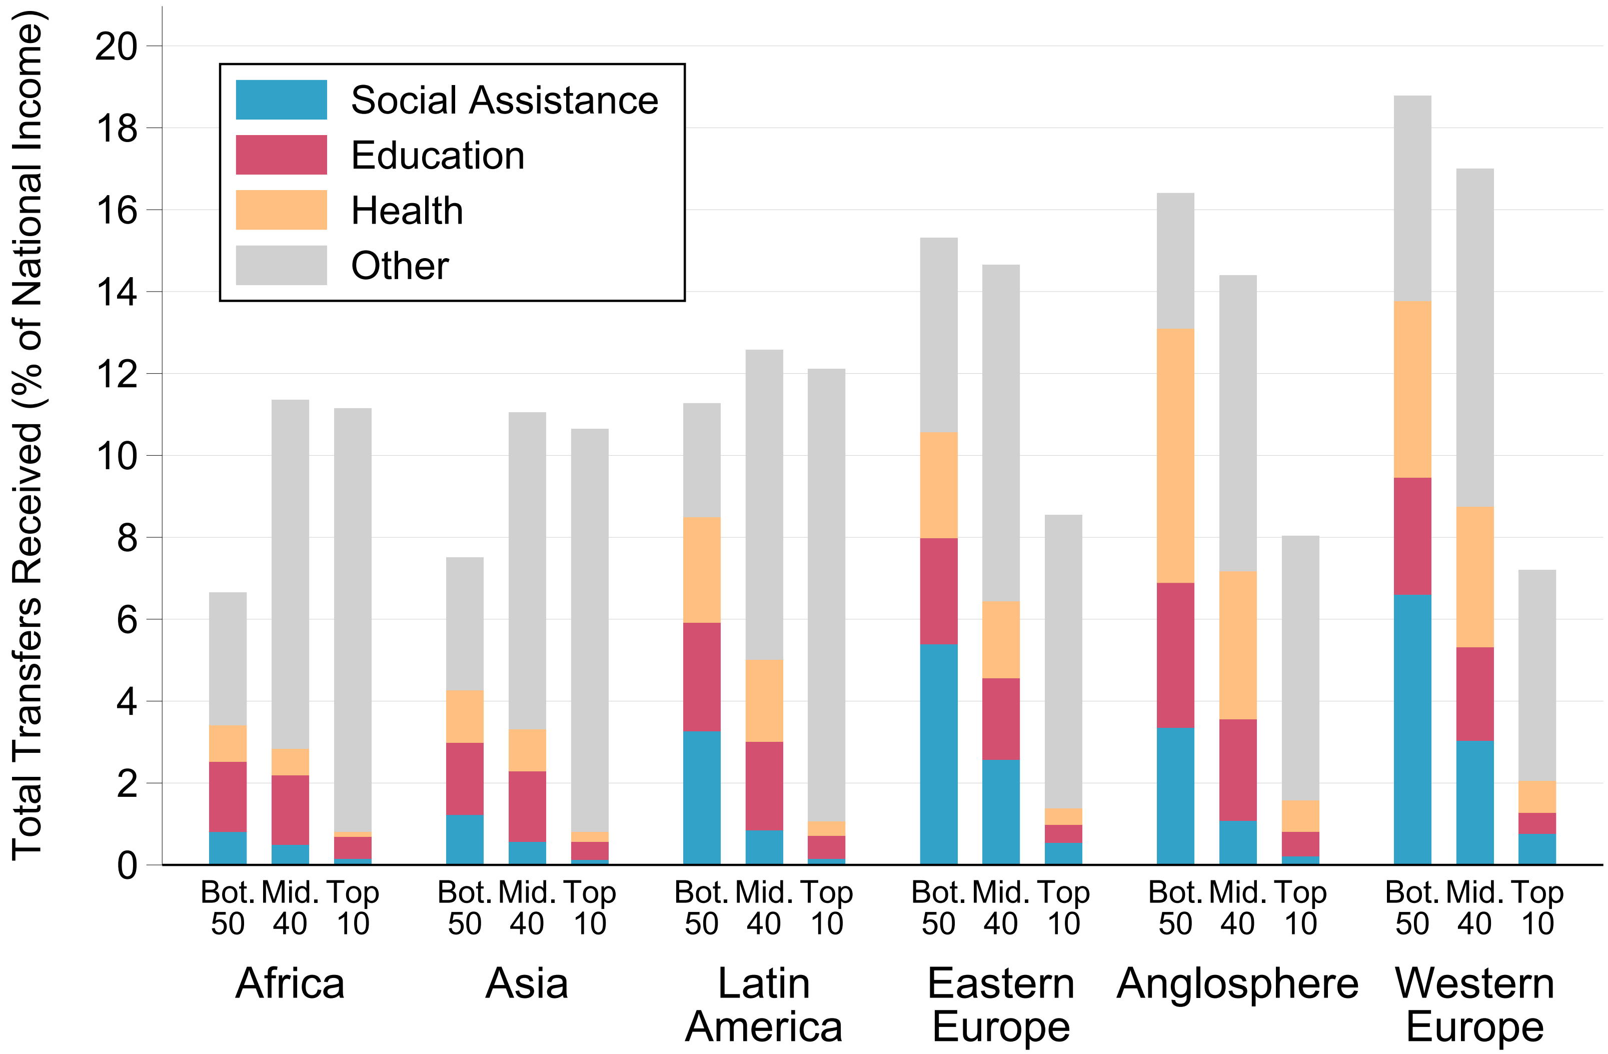

Notes. Population-weighted average of transfers received by income group in each country. Bot. 50: bottom 50% (p0p50); Mid. 40: middle 40% (p50p90); top 10: top 10% (p90p100).

Figure C.34: A Global Map of Redistribution
Percent Reduction in Top 10% to Bottom 50% Income Ratio, Pretax - Posttax
Education Distributed Based on School Attendance

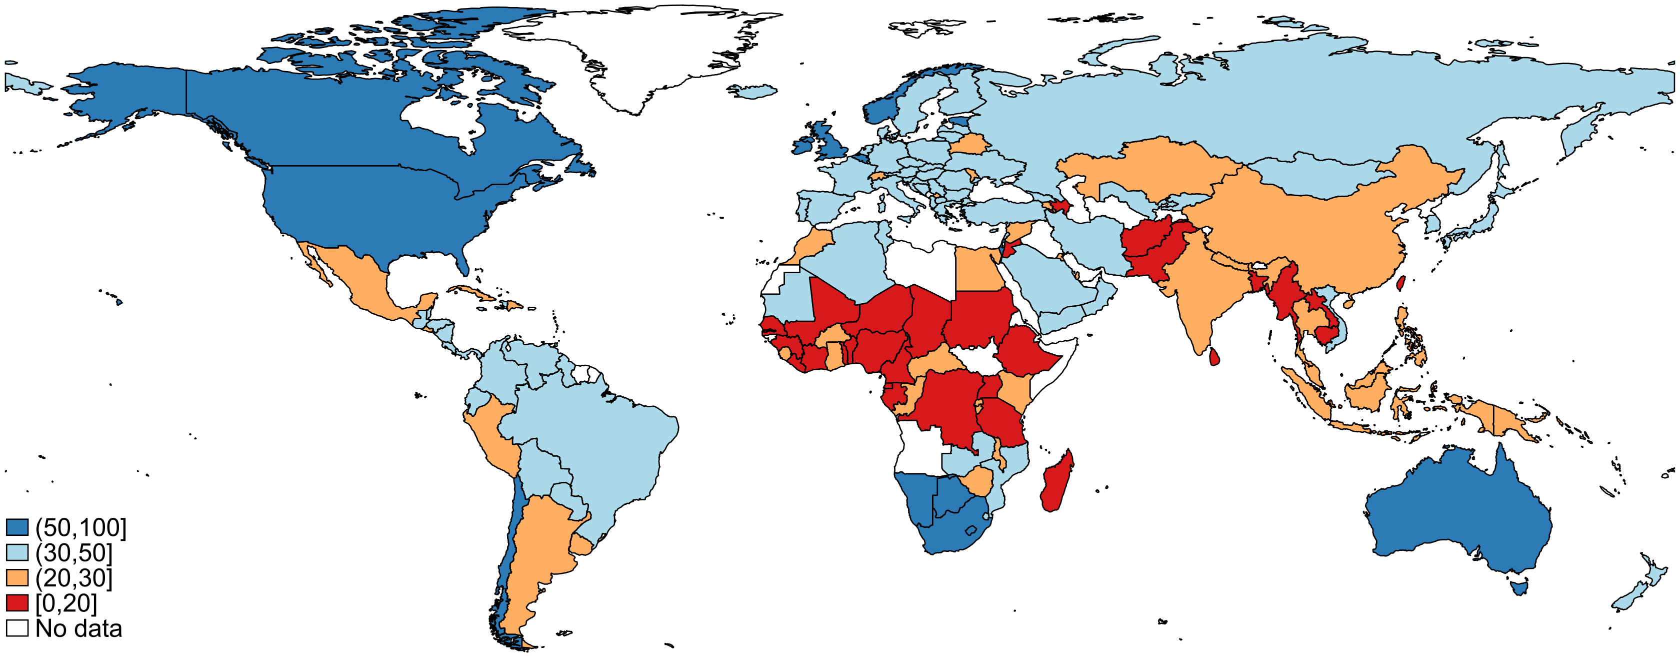

Notes. Posttax income: pretax income, minus all taxes, plus all transfers. Taxes exclude social contributions.

Figure C.35: A Global Map of Redistribution: Net Transfers Operated by the Tax-and-Transfer System Between Pretax Income Groups, 2019
Education Distributed Based on School Attendance

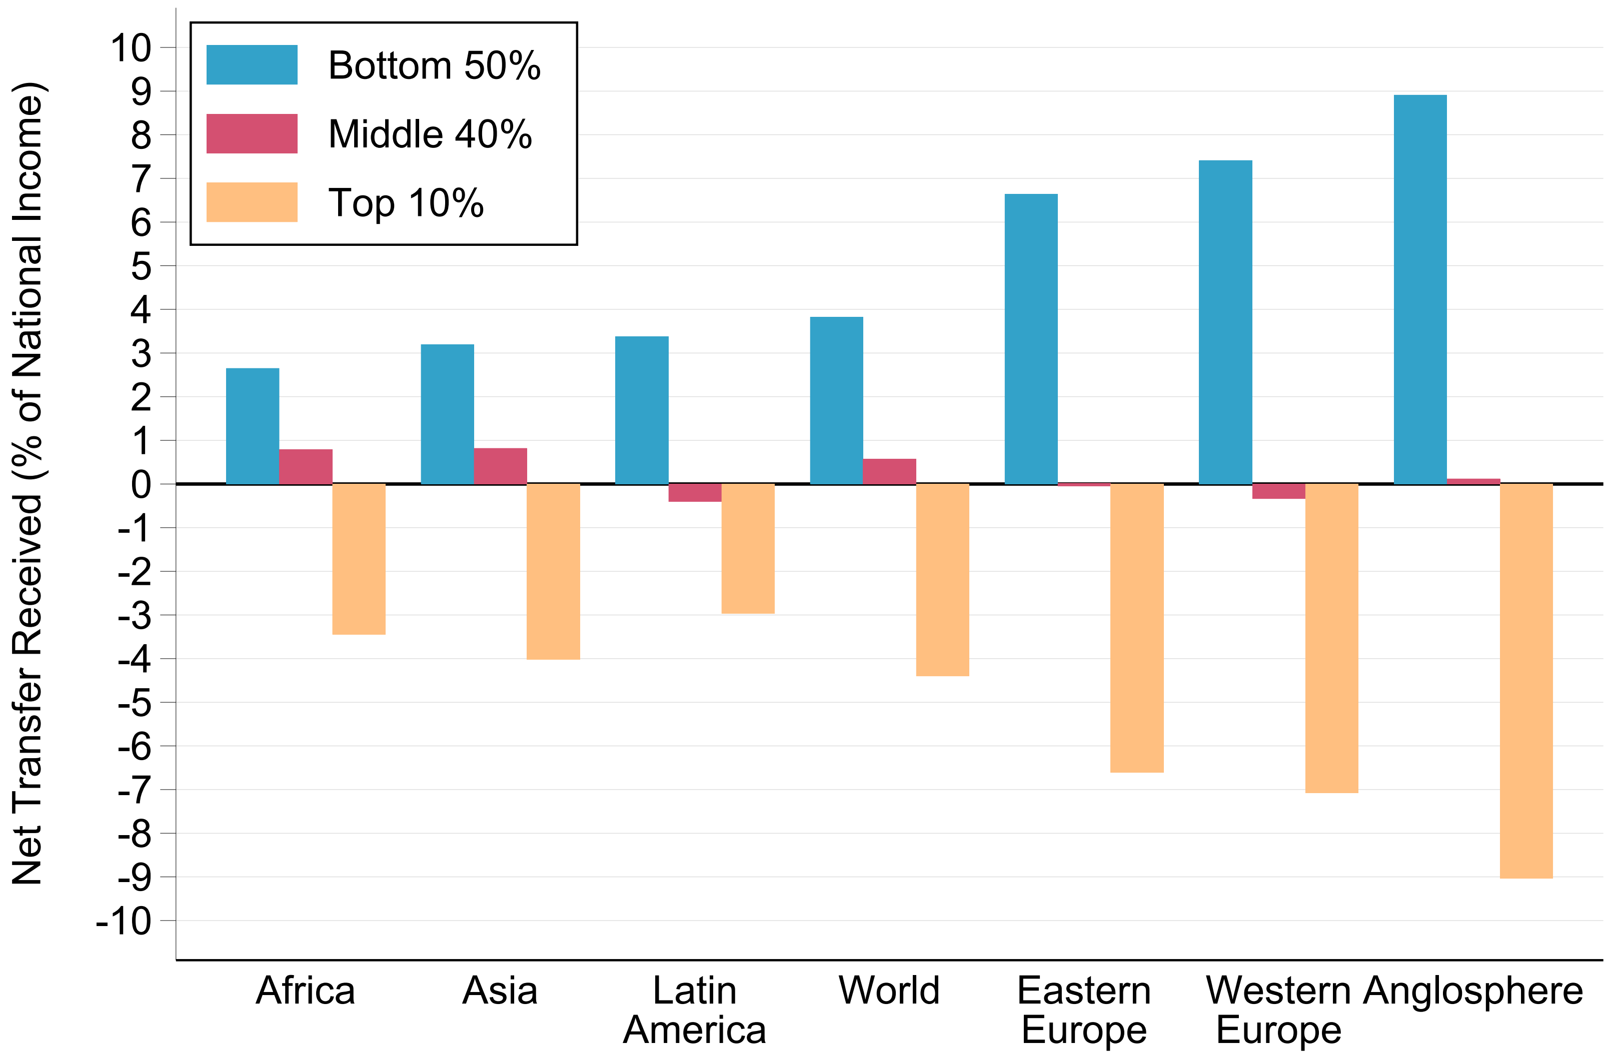

Notes. Net transfer: all transfers received minus all taxes paid, expressed as a share of national income. Taxes exclude social contributions. Population-weighted averages of net transfers received by income group in each country.

Figure C.36: Extent of Redistribution by World Region, 1980-2019:
Percent Reduction in Top 10% to Bottom 50% Income Ratio, Pretax - Posttax
Education Distributed Based on School Attendance

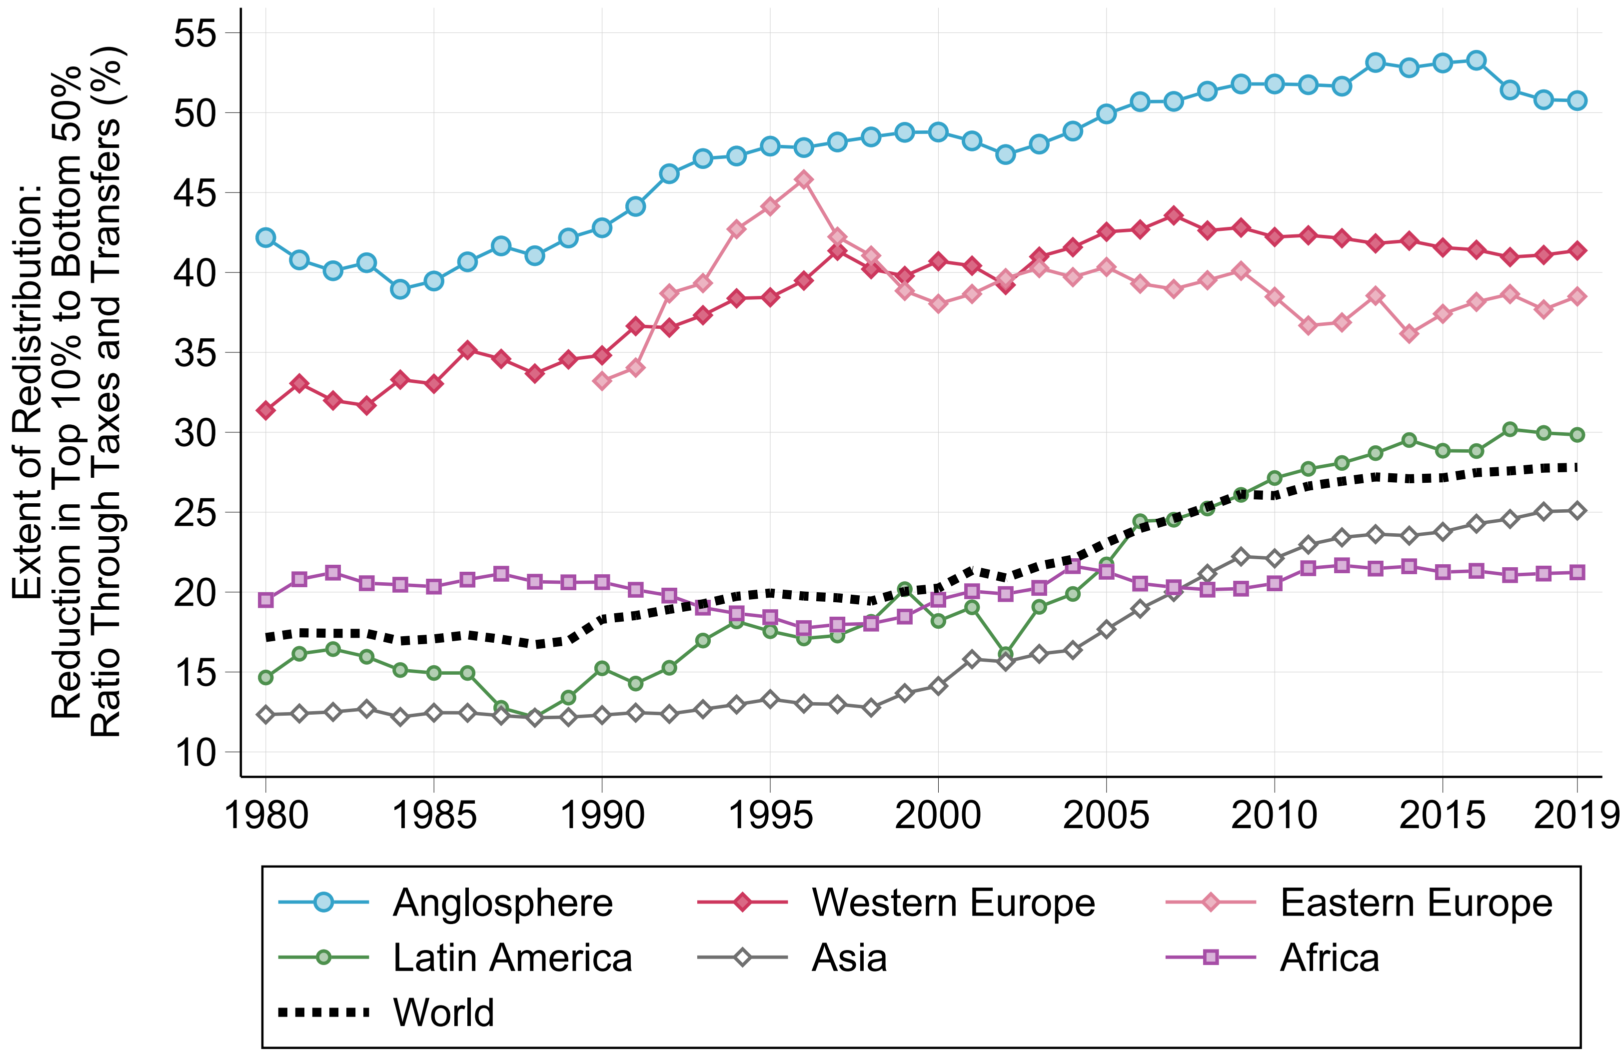

Notes. Population-weighted averages of the extent of redistribution in each country.

Figure C.37: Extent of Redistribution by World Region, 1980-2019:
Net Transfer Received by the Bottom 50% (% of National Income)
Education Distributed Based on School Attendance

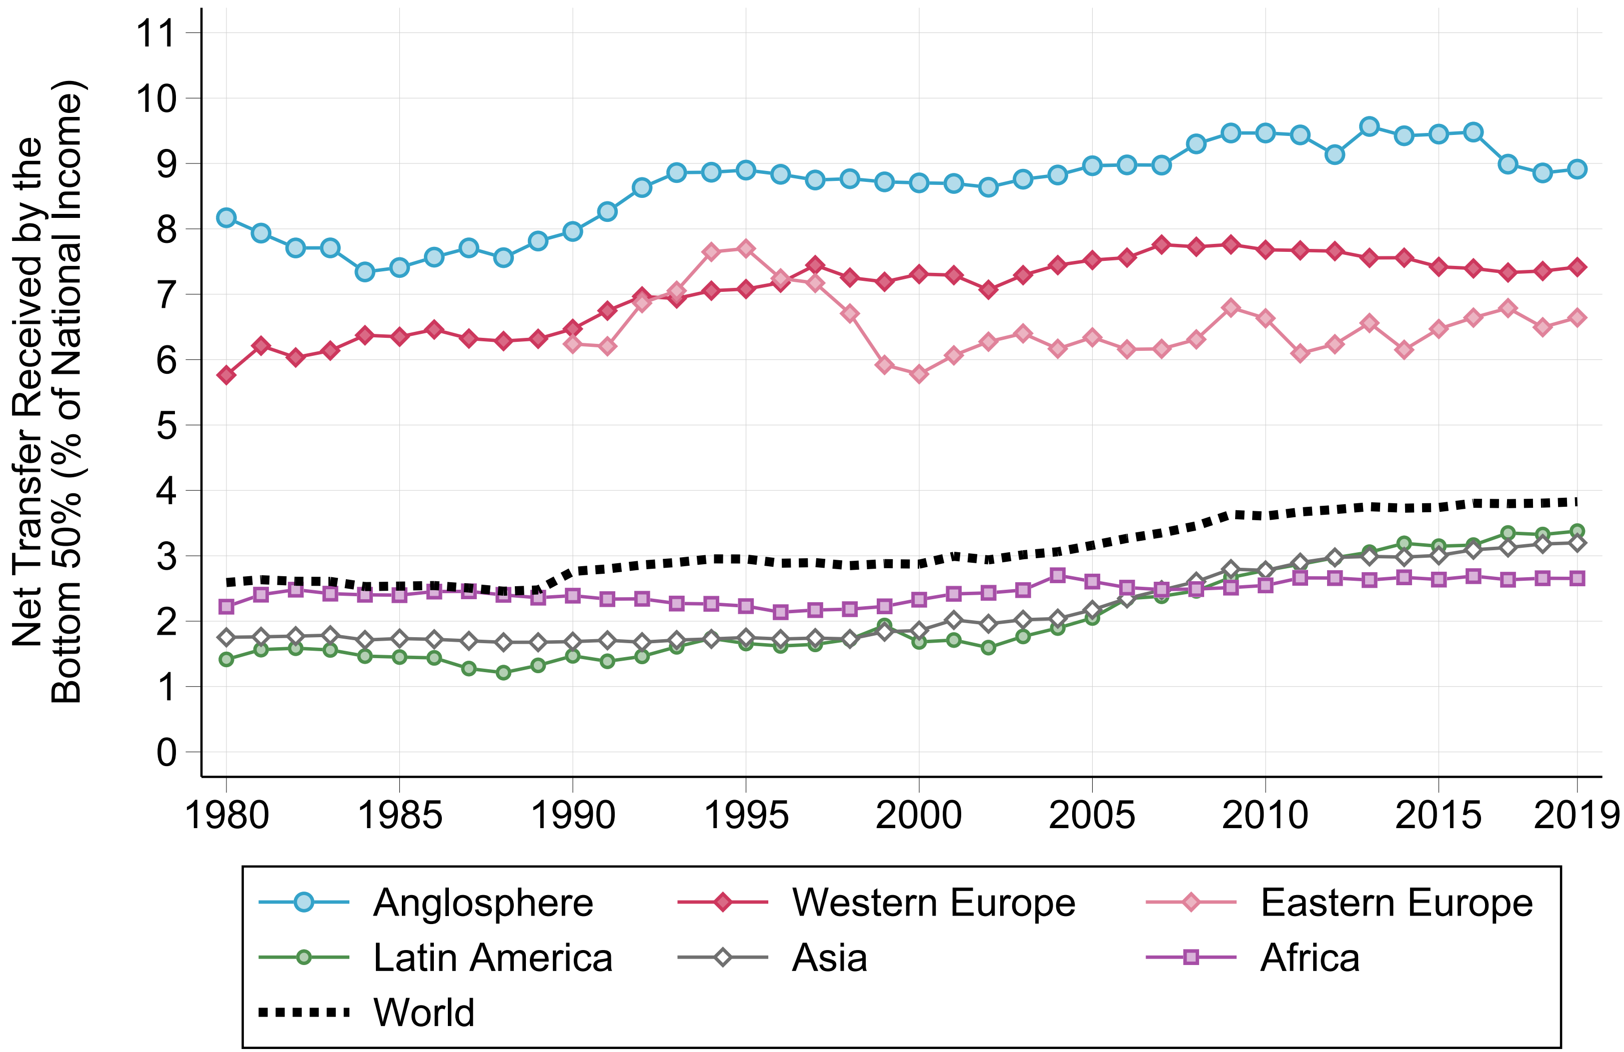

Notes. Net transfer: all transfers received minus all taxes paid, expressed as a share of national income. Population-weighted averages of net transfers received in each country.

Figure C.38: Transfer Progressivity Over the Course of Development:
Total Transfer Received by the Bottom 50% (% of National Income)
Education Distributed Based on School Attendance

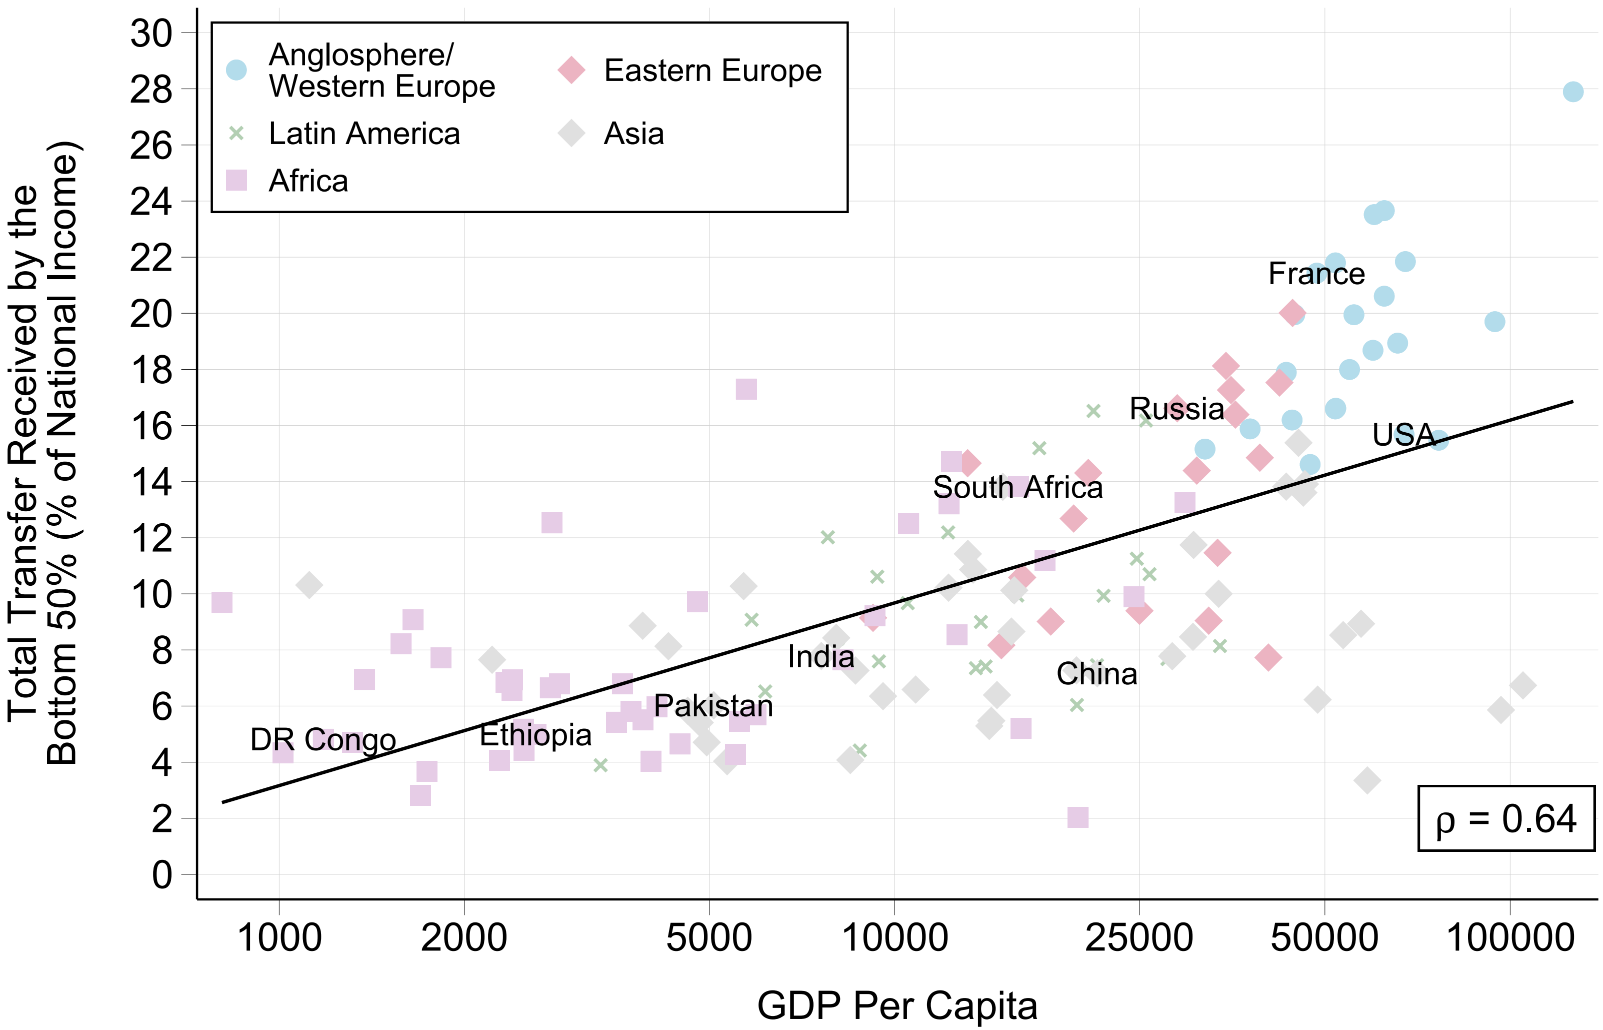

Notes. Total transfer received: sum of all transfers received (before paying taxes), expressed as a share of national income.

Figure C.39: Net Redistribution Over the Course of Development:
Percent Reduction in Top 10% to Bottom 50% Income Ratio, Pretax - Posttax
Education Distributed Based on School Attendance

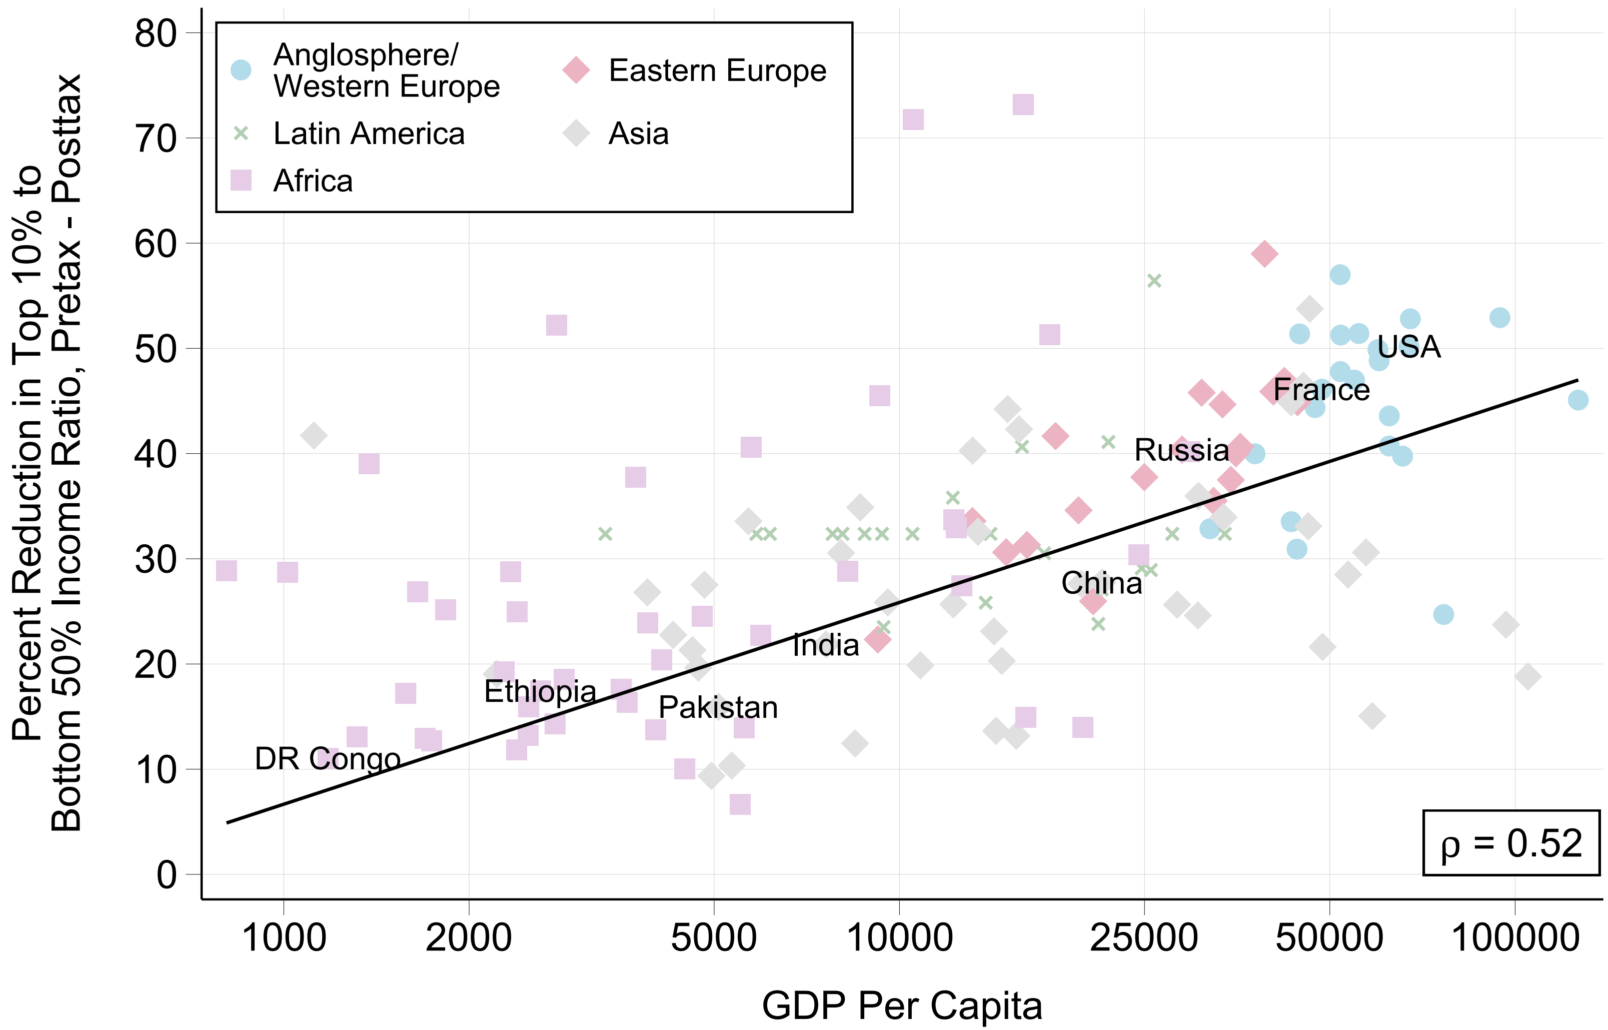

Figure C.40: Predistribution versus Redistribution:
Bottom 50% Pretax versus Posttax National Income Shares by Country, 2019
Education Distributed Based on School Attendance

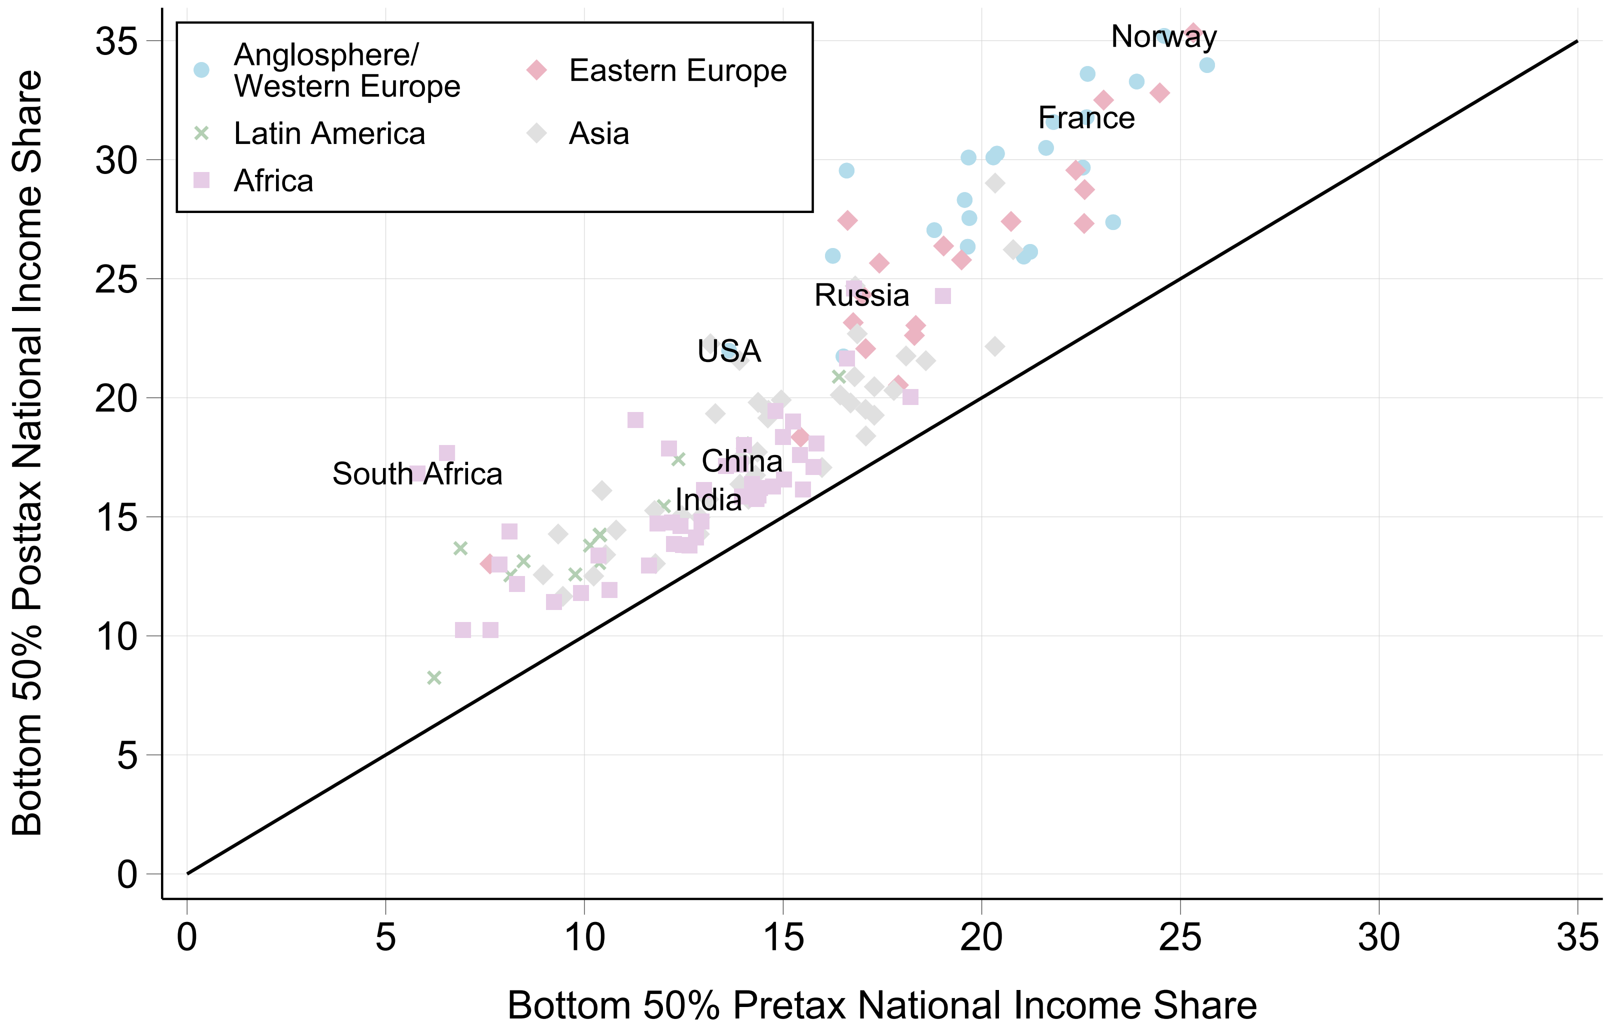

Figure C.41: Predistribution versus Redistribution:
Bottom 50% Pretax Income Share versus Extent of Redistribution, 2019
Education Distributed Based on School Attendance

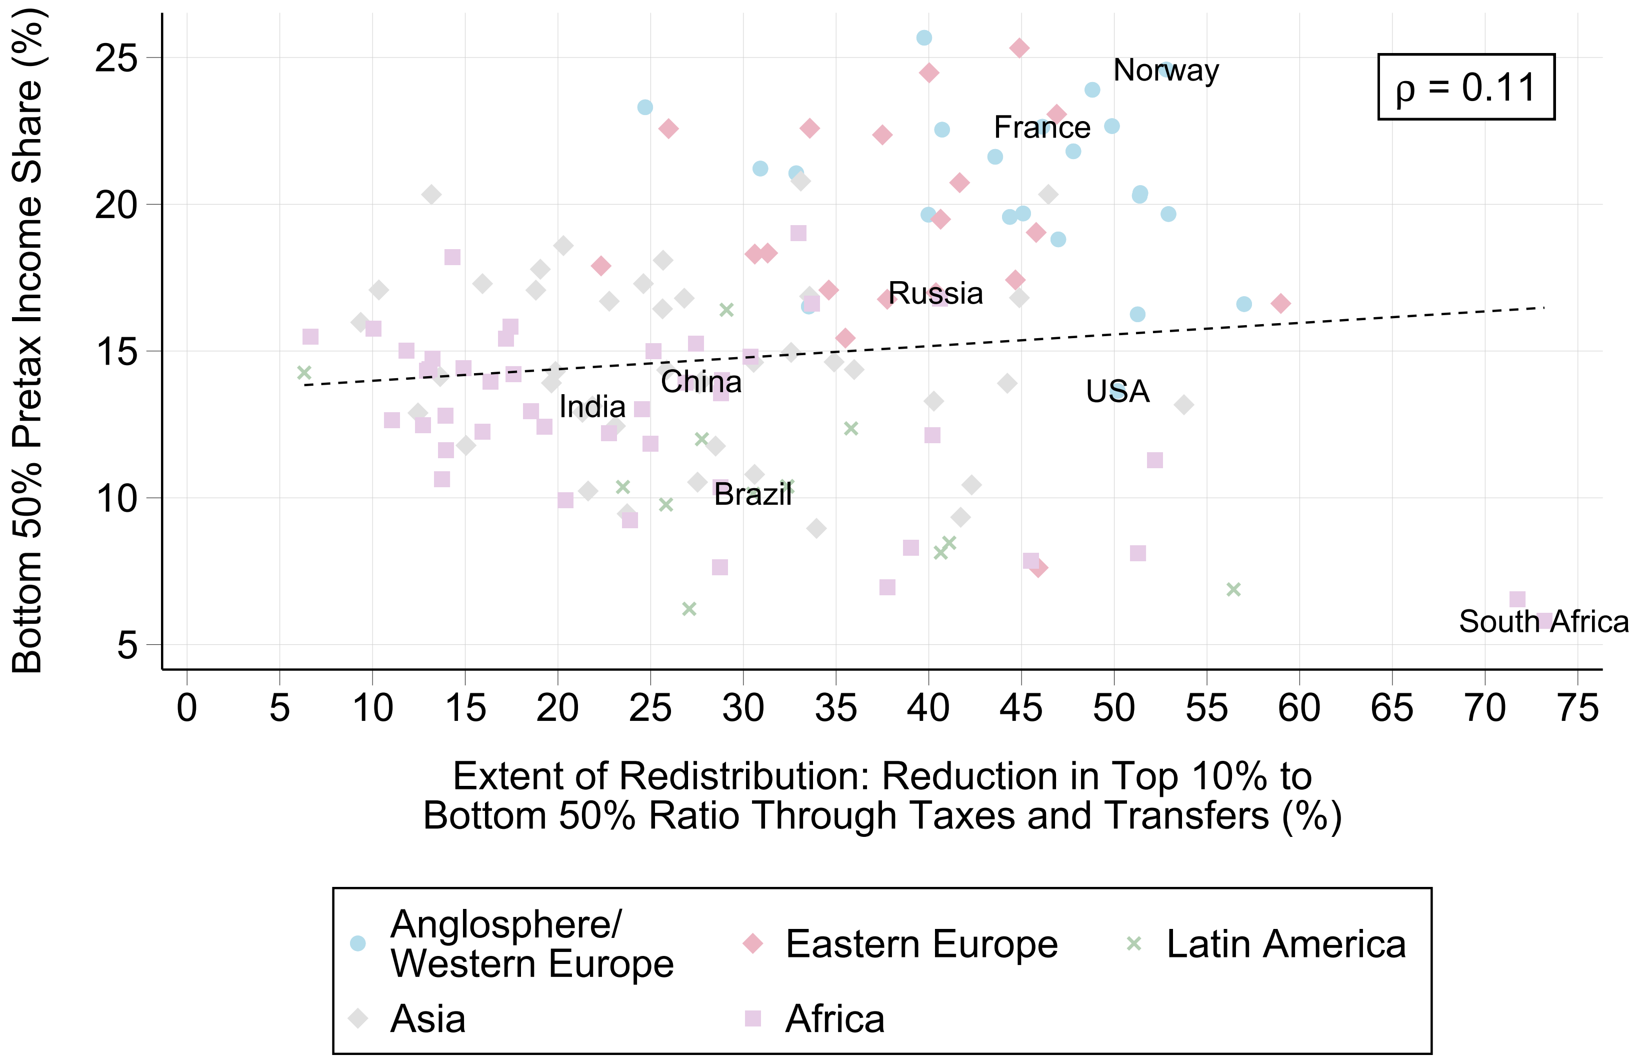

Figure C.42: Predistribution versus Redistribution:
Bottom 50% Pretax Income Share versus Net Transfer Received by the Bottom 50%, 2019
Education Distributed Based on School Attendance

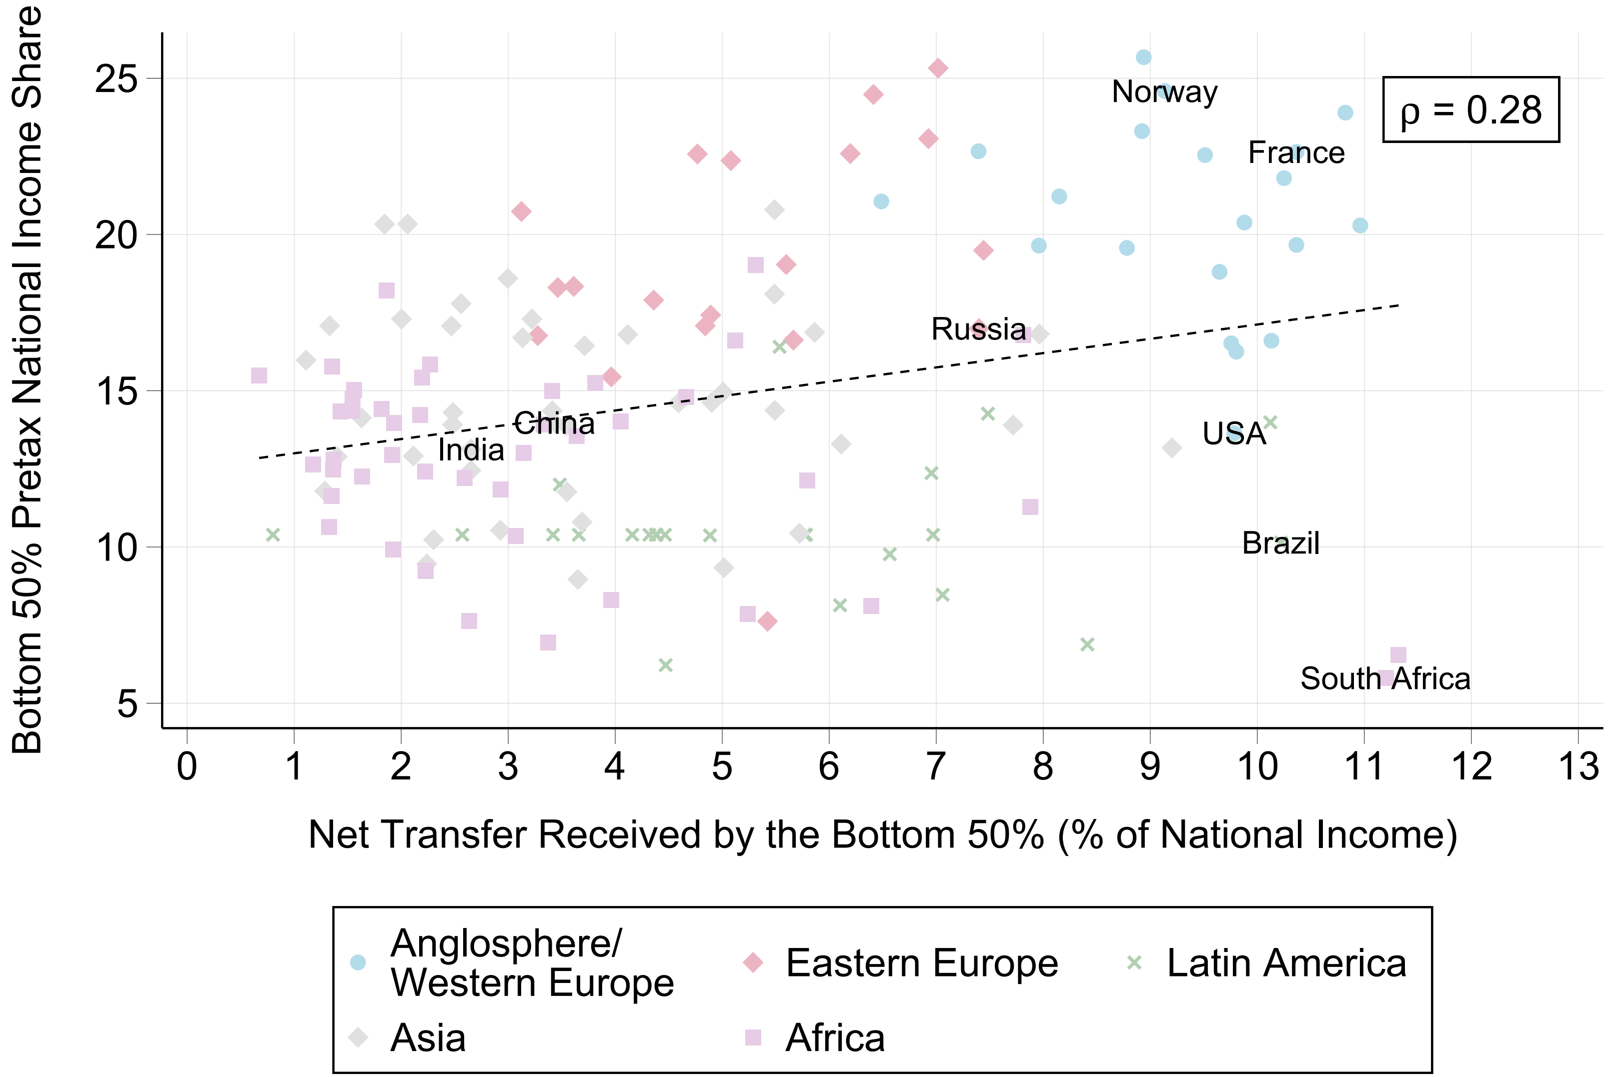

Table C.2: Extent of Redistribution by World Region: the Dominant Role of Transfers
Education Distributed Based on School Attendance

	Top 10% / Bottom 50% Average Income Ratio			Extent of Redistribution: Percent Reduction in Inequality		
	Pretax Income	After Taxes	After Taxes & Transfers	Through Taxes	Through Taxes & Transfers	Tax Share of Redistribution
Africa	20.0	18.9	14.5	4.2%	21.2%	19.7%
Anglosphere	14.8	13.0	7.3	11.6%	50.8%	22.8%
Asia	17.4	17.0	13.1	2.9%	25.1%	11.4%
Eastern Europe	11.2	13.0	6.9	-13.7%	38.5%	-35.6%
Latin America	31.6	35.0	21.9	-10.6%	29.8%	-35.4%
Western Europe	8.7	8.4	5.1	3.8%	41.4%	9.3%
World Average	18.2	18.0	13.1	1.8%	27.8%	6.5%

Notes. Population-weighted averages of indicators in each country. After taxes: top 10% to bottom 50% average income ratio in terms of net-of-tax income (pretax income minus all taxes). After taxes and transfers: top 10% to bottom 50% average income ratio in terms of posttax income (pretax income minus all taxes plus all transfers). Tax share of redistribution: ratio of extent of redistribution through taxes over extent of redistribution through taxes and transfers.

Table C.3: Extent of Redistribution by World Region: Decomposition by Tax and Transfer, 2019
Education Distributed Based on School Attendance

	World Average	Anglosphere	Western Europe	Eastern Europe	Latin America	Asia	Africa
Personal Income Taxes	4.4%	12.4%	14.0%	3.7%	4.6%	3.1%	3.2%
Corporate Taxes	4.2%	3.7%	3.7%	4.4%	4.0%	4.6%	3.3%
Property & Wealth Taxes	0.6%	0.8%	1.3%	0.6%	0.4%	0.6%	0.0%
Indirect Taxes	-7.7%	-7.3%	-14.7%	-23.4%	-10.2%	-6.9%	-3.3%
Social Contributions	-1.3%	-5.7%	-2.5%	-6.6%	-0.7%	-0.9%	0.2%
All Taxes	3.1%	12.1%	9.5%	-12.3%	0.9%	2.9%	4.2%
Social Assistance	10.4%	16.6%	22.9%	20.7%	23.5%	7.5%	5.5%
Education	12.0%	18.3%	11.0%	11.1%	21.4%	10.4%	10.9%
Healthcare	10.3%	28.4%	15.8%	11.2%	20.3%	7.5%	6.5%
All Transfers	24.7%	43.5%	37.3%	33.1%	43.0%	20.6%	17.3%

Notes. Population-weighted averages of indicators in each country. The table reports the negative of the percent change in the top 10% to bottom 50% income ratio before and after removing the corresponding tax or adding to corresponding transfer to pretax income. For instance, the top row reports the percent reduction in inequality resulting from removing personal income taxes from individual incomes. Positive values indicate that the corresponding tax or transfer reduces inequality. All series from this paper (existing DINA studies do not provide comparable, detailed decompositions by type of tax).

Appendix D

Appendix to "Why Is Europe More Equal than the United States?"

D.1 Detailed Methodology

This section describes in details the different steps of our methodology. We primarily focus on methodological questions. For detailed information on the availability of sources by country and the effect of the different adjustments, see the extended online appendix.

D.1.1 Aggregate Income Data

We collect data on key income aggregates, primarily from the system of national accounts, but also using auxiliary data sources when necessary.

Aggregate National Income, PPP and Market Exchange Rates We use estimates of national income, purchasing power parities (PPP) and market exchange rates from the World Inequality Database (`https://wid.world`). GDP estimates for former Eastern European countries come from the Maddison database (Bolt and van Zanden, 2020).

Decomposition of National Income We retrieve the decomposition of national income by institutional sector from three main official sources: Eurostat, the OECD and the UN SNA. Eurostat and the OECD arguably provide the highest quality data, so we use them in priority. However they have limited coverage before 1995 or in certain Eastern European countries. We fill these gaps using the UN SNA data, which are more complete, in particular because they include more countries and also historical series from earlier iterations of the system of national accounts.

When combining these series together, we apply a systematic splicing procedure that looks at the gap between two sources in the first year they overlap, and apply that same gap to the less recent data series (i.e., we adjust its level but preserve its trend).

Imputed Rents In practice, the treatment of the imputed rents of owner-occupied dwellings is not homogenous between countries in their current implementation of the SNA. In some countries, the net operating surplus of the household sector is entirely made up of imputed rents, while in other countries it includes both imputed and non-imputed rents. To fix that issue, we use the supply-and-use tables published by the OECD, which explicitly identifies the imputed rents of owner-occupied dwellings, to split the net operating surplus of the housing sector into imputed and non-imputed rents when necessary.

Separation of Retained Earnings between Shareholders, Pension funds and Government The income of the corporate sector can ultimately accrue to shareholder households, to pension funds or to the government. To estimate that split, we rely on the OECD's financial balance sheets and pension fund statistics. The OECD pension funds statistics include the value of funded pensions, and the share of these pensions that is invested in stocks. The financial balance sheets contain the value of equity that is held by the household and general government sectors. We split retained earnings between shareholder households, pension funds and the government in proportion to their respective equity holdings.

We make one adjustment in Norway, where public shareholdings are very large due to its sovereign wealth fund, but represent profits that are essentially made abroad and therefore are not included in its domestic corporate income. For this reason, we subtract the value of Norway's wealth fund from its public shareholding before we do the computation. In other countries we assume that government shareholdings are essentially made up of domestic companies.

Social Expenditures In the SNA, all social expenditures in cash (including social insurance such as pension and unemployment on the one hand, and social assistance benefits in the other) are pooled into item D62 ("Social benefits other than social transfers in kind"). In principle, this item is meant to be broken down further into the different types of benefits in the SNA nomenclature, but in practice that level of detail in not available directly in most countries. To overcome that issue, we use the

OECD social expenditure database, which breaks down social benefits by type, to split item D62 into pension, unemployment and other.

Health Expenditures Public health expenditure are part of government final consumption expenditure (item P3 in the SNA). In the main SNA tables, this item is broken down into individual expenditures (P31) and collective expenditures (P32). Health is generally included in individual expenditures (P31) alongside other types of spending (e.g., education), and this item is not broken down further.

To get an estimate of public health spending, we rely on two other databases. One is a satellite account of the SNA, the "Government final consumption expenditure by function," (COFOG) which is published by the OECD, Eurostat, and the UN SNA, and breaks down government final expenditures by function, including a separate item for health. The other is the OECD health database, which also provides data on government health spending.

Switzerland is the one country that requires a special treatment. The health system in Switzerland rests on private health insurance with public subsidies and a strict individual mandate. Other European countries have similar system but nonetheless classify their health subsidies as public expenditure (P3) in the national accounts. Switzerland, on the other hand, has virtually no final consumption expenditures on health in the SNA and classifies most of its public spending as subsidies (D3). For more comparable results, we reclassify these health subsidies a public health expenditures.

Imputations Data coverage of aggregate data is quite good, especially after 1995. For the remaining missing data, we extrapolate backward in time the first available value as a fraction of national income, and when a piece of information is entirely missing for a country, we rely on a European average. We systematically rescale the subcomponents of income to match accounting identities.[1]

D.1.2 Estimation of Incomes from Survey Microdata

D.1.2.1 Construction of Factor, Pretax and Posttax Income from EU-SILC

We use the EU-SILC survey as our key source for microdata on the distribution of income. The EU-SILC is a pan-European survey managed by Eurostat, which covers

[1]To extrapolate the first available value backward we use simple exponential smoothing with a coefficient of 0.9, to somewhat limit the impact of having an atypical first value on the whole series.

most European countries with detailed information on income. The first wave of the survey was 2004, with more countries and more detailed income information being progressively added over time. In particular, most pretax income information started being added with the 2007 wave in most countries.

The EU-SILC records wages of employees and the self-employed, distributed capital incomes, and government taxes and transfers. We use these data to construct factor, pretax and posttax incomes according to our definitions, with the exclusion of incomes not included in surveys (retained earnings, taxes on products, etc.), which are included in further steps. In general, incomes recorded in EU-SILC data for year N refer to the year $N-1$, with two exceptions: in Ireland the income reference period is the last twelve months, and in the United Kingdom current income is annualized and aims to refer to the current calendar year. We accordingly adjust income years.

The EU-SILC also records basic demographic information (age, household structure, etc.) that we use to calculate income according to various equivalence scales. Importantly, it also allows us to identify couples within households (defined as married people and partners in a consensual union, with or without a legal basis), in cases where multiple couples live within the same household. This allows us to estimate the distribution of incomes both according to the "broad equal-split" convention (income split equally among all household members) and the "narrow equal-split" convention (income split equally among members of couples).

D.1.2.2 Estimation of Social Contributions

One limitation of EU-SILC is that it does not record separately employee social contributions from taxes on income and wealth. Following the recommendations of the Canberra Group (Canberra Group, 2011), the EU-SILC pools those two items together, even as it separates employee social contributions from employer social contributions in cases where the latter are recorded. To overcome that issue, we use the social contribution schedules published by the OECD to simulate social contributions at the individual level. Note that these imputations may impact the distribution of pretax income, but have no impact on posttax incomes, because posttax incomes deduct both taxes and contributions.

We separately impute for each individual (i) social contributions of employees, (ii) social contributions of the self-employed and (iii) employer social contributions. Employer contributions have started to be recorded in EU-SILC directly in recent years, in which case we use the EU-SILC value directly. In other cases we rely on our estimation. At every step, we ensure the plausibility of our results by making sure that (i) our estimated social contributions are smaller than the combined value of taxes and employee social contributions from EU-SILC and that (ii) our estimates of employer social contributions are consistent with the value recorded in EU-SILC whenever the latter is available. We found the two sources (OECD and EU-SILC) to be largely consistent. There are only three countries (Croatia, Romania and Serbia) that have EU-SILC data but no OECD data on social contributions. For those three countries, and absent better information, we assume that social contributions are proportional to factor income.

Having estimated social contributions (both employer and employee), we separate them into a “contributory” and a “non-contributory” component. The contributory component pays for social insurance (i.e., pension and unemployment benefits) while the non-contributory component pays for other benefits (e.g., family benefits). One solution would be to separate which contribution is meant to pay for which type of benefit in the social contribution schedule directly, but on top of being very demanding, this approach would not yield useful results. Indeed, due to the fungible nature of public funds, social contributions that are supposed to pay for a given benefit can often exceed or fall short of the benefit amount for spurious reasons. Hence, we follow a more simple and robust first-order approach, which is to split contributory and non-contributory contributions proportionally, so that contributory contributions match the overall amount of pension and unemployment benefits paid. By construction, this approach ensures equilibrium between contributions and benefits, by implicitly distributing the surplus or deficit of the social insurance system proportionally to social contributions. In some countries, pension and unemployment benefits exceed the total amount of social contribution. The most notable example is Denmark, where social contribution are virtually nonexistent because social insurance is primarily financed by regular taxes. In such cases, we consider that a fraction of the income tax pays for social insurance, and we treat that fraction like social contributions.

D.1.3 Harmonization of Other Survey Data Sources

D.1.3.1 Data Collection and Interpolation

To extend our coverage of survey data, we gather a large collection of survey tabulations from a variety of sources. Some of them take the form of survey tabulations, coming from PovcalNet (World Bank, n.d.[b]), the World Income Inequality Database (UNU-WIDER), and Eastern European estimates published by Milanović (1998).

These tabulations describe distributions of income by giving income shares of various brackets, whose number and location vary. We construct complete tabulations by g-percentile using the generalized Pareto interpolation method introduced by Blanchet, Fournier, and Piketty (2021).[2] Most of these tabulations refer to either post-tax income or consumption.

We also use survey microdata from a variety of sources, from which we calculate all income concepts and equivalence scales possible, and collapse them into tabulated distributions. These include distributions from the Luxembourg Income Study (LIS), a database that collects, harmonizes, and makes available to researchers a wide range of survey microdata from many countries across the world. They also include the European Community Household Panel (ECHP), the precursor to EU-SILC, and two surveys from the World Bank covering Serbia in 2002, 2003 and 2007 (as well as Kosovo in 2000). In all cases these surveys cover posttax income, but in many cases they also cover pretax income.[3]

D.1.3.2 Harmonization of Income Concepts

The set of income distributions that we collect is very heterogeneous. It uses various various income concepts (pretax income, posttax income, consumption), various statistical units (individual, household), and various equivalence scales (square root, OECD, equal-split per capita, equal-split per adult). We harmonize this dataset to retrieve our concepts of interest: equal-split per adult, both at the household level (broad equal-split) and at the couple level (narrow equal-split). To that end, we notably take advantage of our access to survey microdata, which makes it possible to calculate variants of the income distribution for a wide array of income concepts, and therefore lets us observe how they tend to relate to one another.

Indeed, distributions for the different income concepts across country-years are correlated: therefore, we can use the distribution for one income concept to impute the distribution for another whenever the former is observed but not the latter. To do so, we use all the cases where the income distribution is simultaneously observed for two different concepts to learn how one tends to relate to another.

We can observe the p-th quantile of both the source and the target distributions for

[2]What we call g-percentiles refer to every percentile from $p = 0\%$ to $p = 99\%$, then $p = 99.1\%$ to $p = 99.9\%$, then $p = 99.91\%$ to $p = 99.99\%$, and finally $p = 99.991\%$ to $p = 99.999\%$.

[3]The treatment of social contributions in these surveys is not always as satisfying as what we were able to do for EU-SILC. However, to the extent that the deduction of social contributions makes little difference to the distribution of pretax incomes in EU-SILC—which is usually the case—we used pretax income from these surveys as a proxy for true pretax income for the historical period.

Table D.1: 5-fold cross validation mean relative error on the average by percentile when imputing pretax and posttax incomes from different concepts using our benchmark machine learning algorithm

	predictor	predicted concept			
		pretax income (broad equal-split)	pretax income (narrow equal-split)	posttax income (broad equal-split)	posttax income (narrow equal-split)
consumption	equal-split (broad)	9.9%	11.0%	8.4%	11.1%
	per capita	8.7%	11.1%	9.5%	12.0%
	households	9.2%	10.8%	7.9%	10.2%
	OECD scale	9.7%	10.4%	8.8%	11.7%
	square root scale	9.3%	10.7%	8.2%	11.7%
pretax income	equal-split (broad)	n/a	3.3%	5.8%	6.0%
	equal-split (narrow)	2.9%	n/a	5.6%	4.7%
	per capita	3.7%	5.1%	6.3%	6.4%
	households	3.9%	4.8%	7.2%	6.7%
	OECD scale	2.4%	3.8%	6.2%	6.2%
	square root scale	2.7%	4.1%	6.4%	6.5%
posttax income	equal-split (broad)	5.6%	6.4%	n/a	4.3%
	equal-split (narrow)	5.3%	4.8%	3.9%	n/a
	per capita	6.8%	7.6%	3.6%	5.5%
	households	6.4%	7.0%	3.9%	5.5%
	OECD scale	5.7%	6.5%	2.2%	4.5%
	square root scale	5.6%	6.5%	2.7%	4.7%

Source: authors' computations. *Note*: Error calculated only for the top 80% of distributions to avoid problems of denominator near zero. The algorithm is XGBoost's implementation of boosted regression trees using $\eta = 0.1$ (Chen and Guestrin, 2016). Auxiliary variables included in the model are: regional dummies, average national income per adult (PPP), share of households with size 1 to 6, gross saving rate (% of GDP), overall social expenditures (% of GDP), top marginal income tax rate, income tax revenue (% of GDP), overall tax revenue (% of GDP), share of population by 10-year age bands and sex, corporate tax rate, VAT tax rate. *Interpretation*: When imputing pretax income per equal-split adult (broad) from consumption per household, the mean relative error for the average income of a given percentile is 9.2%.

a variety of countries i and a variety of years t: write them $Q_{it}^{\text{target}}(p)$ and $Q_{it}^{\text{source}}(p)$. To construct the best mappings φ between the different concepts, we consider a very general model. In that model, each percentile of the target distribution is an arbitrary function of every percentile of the source distribution, and of additional covariates. We write:

$$\mathbb{E}[Q_{it}^{\text{target}}(p)] = \varphi(Q_{it}^{\text{source}}(p_1), \dots, Q_{it}^{\text{source}}(p_m), p, t, Z_{it})$$

for a grid $0 \leq p_1 < \cdots < p_m < 1$ of fractiles, and for auxiliary variables Z_{it}. Estimating such a model raises some challenges. Linear regression will not be flexible enough due to its parametric assumptions and will tend to overfit the data if m is large.

To estimate this model, we therefore rely on more recent advances in high-dimensional, nonparametric regression, also known as *machine learning* methods. The algorithm we use is known as *boosted regression trees*, a powerful and commonly used method introduced by Friedman (2001). We rely on an implementation known as XGBoost (Chen and Guestrin, 2016), which has enjoyed great success due to its speed and performance, to the point that is has earned a reputation for "winning every machine

Table D.2: 5-fold cross validation mean relative error on the average by percentile when imputing pretax and posttax incomes from different concepts using a machine learning algorithm without auxiliary variables

	predictor	predicted concept			
		pretax income (broad equal-split)	pretax income (narrow equal-split)	posttax income (broad equal-split)	posttax income (narrow equal-split)
consumption	equal-split (broad)	11.1%	12.2%	10.7%	11.8%
	per capita	11.0%	12.7%	9.2%	12.1%
	households	9.9%	11.8%	9.2%	11.7%
	OECD scale	10.8%	12.5%	9.9%	12.3%
	square root scale	10.6%	12.3%	9.3%	11.9%
pretax income	equal-split (broad)	n/a	3.7%	6.3%	6.5%
	equal-split (narrow)	3.1%	n/a	5.5%	4.5%
	per capita	3.9%	5.5%	6.8%	7.6%
	households	3.7%	5.4%	7.5%	7.5%
	OECD scale	2.4%	4.2%	6.4%	6.6%
	square root scale	2.6%	4.3%	6.6%	6.7%
posttax income	equal-split (broad)	5.8%	6.4%	n/a	4.4%
	equal-split (narrow)	5.4%	4.8%	4.0%	n/a
	per capita	7.3%	7.8%	3.8%	5.8%
	households	6.6%	6.7%	3.8%	5.7%
	OECD scale	6.2%	6.5%	2.3%	4.6%
	square root scale	6.2%	6.5%	2.7%	5.0%

Source: authors' computations. *Note*: Error calculated only for the top 80% of distributions to avoid problems of denominator near zero. The algorithm is XGBoost's implementation of boosted regression trees using $\eta = 0.1$ (Chen and Guestrin, 2016). No auxiliary variables are included in this model. *Interpretation*: When trying to impute pretax income per equal-split adult from consumption per household, the mean relative error for the average income of a given percentile is 9.9%.

learning competition" (Nielsen, 2016). On top of its performance, boosted regression makes it easy to deal with missing values, or to impose certain constraints, such as the fact that the quantile function $Q(p)$ must be increasing with p.

We use five-fold cross-validation to determine the optimal number of "boosting rounds" that the algorithm performs, which determines the trade-off between bias and variance. Since our dataset is made up of countries that we follow over the years, it has a panel dimension, which we take into account as follows. We assume that the country-specific prediction error is independent conditional on all observed variables (i.e., that it is a *random* rather than a *fixed* effect.) Under that assumption, the imputation method remains valid because the error term remains exogenous. However, there is a risk of over-fitting if we do not make sure that the different subsamples used in the cross-validation are not independent, because then we would force the algorithm to try to predict the country random effect. To avoid that problem, we perform the cross-validation by making sure that all observations for one country are in the same cross-validation fold, which is known as leave-one-cluster-out cross validation (Fang, 2011). When possible, we also estimate and include the country random effect into our imputation. The random effect is estimated as a function of the percentile using the mean prediction error by country and percentile.

Table D.3: 5-fold cross validation mean relative error on the average by percentile when imputing pretax and posttax incomes from different concepts using a single correction coefficient by percentile

	predictor	predicted concept			
		pretax income (broad equal-split)	pretax income (narrow equal-split)	posttax income (broad equal-split)	posttax income (narrow equal-split)
consumption	equal-split (broad)	15.2%	17.2%	10.5%	15.2%
	per capita	20.3%	23.7%	11.0%	19.3%
	households	15.9%	18.2%	11.7%	16.2%
	OECD scale	16.7%	19.1%	11.0%	16.6%
	square root scale	14.9%	17.3%	11.1%	15.3%
pretax income	equal-split (broad)	n/a	3.7%	5.9%	6.1%
	equal-split (narrow)	3.7%	n/a	6.3%	4.5%
	per capita	3.9%	5.7%	6.7%	7.2%
	households	4.6%	5.9%	8.1%	8.0%
	OECD scale	2.4%	4.5%	6.3%	6.5%
	square root scale	2.8%	4.7%	6.6%	6.8%
posttax income	equal-split (broad)	5.8%	6.4%	n/a	4.9%
	equal-split (narrow)	6.1%	4.6%	4.8%	n/a
	per capita	6.7%	7.5%	3.9%	6.2%
	households	7.3%	7.6%	4.7%	6.6%
	OECD scale	6.1%	6.6%	2.2%	5.1%
	square root scale	6.2%	6.8%	2.7%	5.5%

Source: authors' computations. *Note*: Error calculated only for the top 80% of distributions to avoid problems of denominator near zero. *Interpretation*: When trying to impute pretax income per equal-split adult from consumption per household, the mean relative error for the average income of a given percentile is 15.9%.

In the end, for any target concept of interest, we get as many predictions as there are sources available. Let $\boldsymbol{y} = (\hat{Q}_{it}^{\text{target,1}}, \dots, \hat{Q}_{it}^{\text{target,n}})'$ be the n different predictions. Using the cross-validation estimation of the prediction error, we can estimate the variance-covariance matrix $\boldsymbol{\Sigma}$ between the different predictions. Following the logic of generalized least squares, the optimal way of combining the n predictions into one is to average them, weighted by the row or column sums of the symmetric matrix $\boldsymbol{\Sigma}^{-1}$. This yields our harmonized estimate of the distribution, taking into account observed regularities across concepts and percentile groups.

As table D.1 shows, the mean (cross-validation) prediction error for the value of the average of a percentile is between 2% and 11% depending on the concept that was used for the prediction.[4] Adjusting for the statistical unit while keeping the income concept identical creates the least difficulties. Consumption, on the other hand, is a rather poor predictor of income. Moving from posttax to pretax income is a somewhat intermediary situation. The auxiliary variables that we use to improve

[4]Before training the model, we transform the data using the transform $y \mapsto \text{asinh}(y)$ for the value of the quantiles and $x \mapsto -\log(1-x)$ for the corresponding rank. This stabilizes the mode without changing the nature of the data. The use of asinh rather than the logarithm avoids issues with having zero or near-zero incomes at the bottom of the distribution. All distributions are normalized by their average since we are only concerned with the distribution of income. When we report prediction errors, these are computed for distributions that have been properly transformed back to their original value.

the performance of the prediction are: regional dummies, average national income per adult (PPP), share of households with size 1 to 6, gross saving rate (% of GDP), overall social expenditures (% of GDP), top marginal income tax rate, income tax revenue (% of GDP), overall tax revenue (% of GDP), share of population by 10-year age bands and sex, corporate tax rate, and VAT tax rate. Table D.2 shows the performance of a model that does not include these variables. While their inclusion has only second-order effects on our harmonized series, they do improve the prediction error, especially when trying to impute based on consumption: we improve the mean relative error by up to 2 pp.

Table D.3 shows the performance of a much more simple imputation method, namely using a single correction coefficient by percentile to move from one concept to another. This coefficient is computed as the mean ratio between two concepts for a given percentile. While this method performs reasonably well for concepts that are close to one another, it exhibit much worse performance when using a poor predictor such as consumption. In such cases, the prediction can be 50% or even 100% worse than our benchmark algorithm.

D.1.4 Calibration of Survey Sources to Tax Data

D.1.4.1 Tax Data Sources

We collect a large set of top income shares estimated from tax data, and use it to adjust our survey estimates. Most of our data comes from the World Inequality Database, from which we extracted "fiscal" top income shares excluding capital gains (which are excluded from national income and from surveys). We also extend series to the latest available year when necessary, by going back to the original source, and add new tax tabulations that we were able to find. These new data series are described country by country in section D.1.7.

D.1.4.2 Calibration Algorithm

We correct survey data for non-sampling error using known top income shares estimated from administrative tax data. We do so by adjusting survey weights using survey calibration methods (Deville and Särndal, 1992). Statistical institutes already routinely use these methods to ensure that their surveys are representative, typically in terms of age and gender. Our approach is a natural extension of theirs, in the sense that we enforce representativity in terms of taxable income in addition to age and gender.

We apply a standard linear calibration algorithm (Deville and Särndal, 1992) to make the survey match the top income shares estimated from the tax data, while minimizing distortions from the original survey data. Because surveys tend to underrepresent top incomes, in practice this means that we inflate the weights of the survey data at the top of the distribution.

One notable difficulty of our setting is that the statistics we calibrate the survey on (top income shares) are not linear statistics of the data, and therefore the most standard calibration framework does not apply. To overcome that issue we apply a two-step calibration procedure following Lesage (2009).

First Step. In the first step, we linearize the top share statistics so that we can apply the standard calibration algorithm. To do that, we need to calculate the *influence function* (Cook and Weisberg, 1980) of top income shares. Let y_k be the income of observation $k \in \{1, \dots, N\}$ in the survey. Let S_α be the top $100(1-\alpha)\%$ income share from the tax data, and let $\hat{Q}_\alpha$ be the α-th quantile in the survey data. Langel and Tillé (2011) showed that the centered influence of observation k on the top $100(1-\alpha)\%$ income share from the survey is:

$$z_k = y_k H\left(\frac{\alpha N - W_{k-1}}{w_k}\right) + (\alpha - \mathbb{1}_{y_k < \hat{Q}_\alpha})\hat{Q}_\alpha - (1 - S_\alpha) y_k$$

where $H(x) = 0$ if $x < 0$, $H(x) = x$ if $0 \leq x < 1$ and $H(x) = 1$ if $x \geq 1$, $W_k = \sum_{k \in s} w_l \mathbb{1}_{y_l \leq y_k}$, $N = W_n$, and $\hat{Q}_\alpha = y_i$ with $W_{i-1} < \alpha W_n \leq W_i$. As explained by Lesage (2009), to calibrate the survey we can enforce that z_k sums to zero using the standard calibration algorithm (Deville and Särndal, 1992).

Second Step. As explained by Lesage (2009), the first step described above works well, but because it relies on a linear approximation of the top share statistics, it only provides a first-order approximation of the solution. To get rid of the remaining discrepancy, we introduce a nuisance parameter: we set the value of the α-th quantile in the survey, and then apply the calibration algorithm to enforce the proper number of people and their proper amount of income on both sides of the quantile. Once $\hat{Q}_\alpha$ is fixed as such, the problem once again becomes linear so we can apply the standard version of the algorithm described by Deville and Särndal (1992).

We apply this two-step calibration method using the top 10% and the top 1% income shares measured from the tax data. In every case, we carefully match the statistical unit and the income concept in the survey to that of the tax data before we apply the method. Having applied the calibration with the right income concept, we can retrieve the corrected version of other income concepts using the microdata with the calibrated weights, most importantly for us pretax and posttax income per equal-split adults.

The key assumption for us to get an appropriate estimate of pretax and posttax inequality via this calibration approach is that, conditional on their fiscal income, the probability that people are included in the survey is not correlated to their pretax or posttax income. Put differently, the fiscal income concept that serves as the basis for calibration must be sufficiently comprehensive to capture what drives the underrepresentation of the rich in the survey. Given that income taxes in Europe are relatively comprehensive we think this is reasonable as a first-order assumption. (The situation would arguably be different in developing countries with very large informal sectors.)

D.1.4.3 Extrapolation of the Tax Data Correction to All Tabulations

To apply the survey calibration method described above, we need access to survey microdata so that we can match income concepts and statistical units to that of the tax data. When we have access to such microdata, this is a very powerful way of harmonizing top income share series that are otherwise difficult to compare.

Unfortunately, adequate microdata is rare before the start of the EU-SILC survey (i.e., 2007 in many cases). Therefore, for the historical period, we retropolate the adjustment. That is, we observe the gap between the distribution of tax-based top income shares (which correspond to fiscal income per tax unit) and the top income shares from the calibrated surveys (which correspond to pretax and posttax income per equal-split adult) over the years with microdata available. We notice that this gap is very stable over time, meaning that our adjustment of the tax-based top income share series affects the levels but has only second-order effects on the trends. Therefore, we retropolate the adjustment to the top income share series as follows.

We calculate the average income of each g-percentile in (i) the tax-based series and (ii) the series based on the calibrated tax data, with the overall income distribution normalized to one in both cases. For each g-percentile, we calculate the ratio between the average of (i) and (ii). We carry that coefficient backward in time and use it to adjust the rest of the tax-based top income share series.

Using the adjusted tax-based series, which now cover the same period of time as the original tax-based series but correspond to our income concepts and statistical units of interest, we re-run our calibration algorithm directly on the harmonized survey tabulations from section D.1.3 using the same algorithm as section D.1.4.

D.1.4.4 Adjustment Within the Top 10%

One issue with using survey data to adjust the tax-based income shares is that surveys have limited granularity at the very top, because of limited sample sizes. Therefore, to improve the quality of our estimates within the very top, we apply one last adjustment. We stress that, by construction, that adjustment has no impact on the top 10% share, and only affects the distribution of income within the top 10%.

This adjustment involves modeling the top 10% of the distribution with a generalized Pareto distribution, which has the cumulative distribution function:

$$F(x) = 1 - \left\{1 + \xi\left(\frac{x-\mu}{\sigma}\right)\right\}^{-1/\xi}$$

This distribution is known in extreme value theory to work as a quasi-universal model of top tails (Ferreira and Haan, 2006). We estimate its parameters using the method of probability-weighted moments (Hosking and Wallis, 1987), a more robust alternative to other methods, which also lets us preserve the average income of the top 10%. For X following a generalized Pareto distribution, define $a = \mathbb{E}[X]$ and $b = \mathbb{E}[X(1 - F(x))]$. Then we have $\xi = (a - 4b + \mu)/(a - 2b)$ and $\sigma = (a - \mu)(2b - \mu)/(a - 2b)$, while μ is determined a priori from the threshold from which we start to use the model. We obtain the complete distribution by combining the empirical distribution for the bottom 90% with the generalized Pareto model for the top 10%.

D.1.5 Distribution of Additional Income Components

D.1.5.1 Data Sources

There are three components of national income that require additional data sources to be distributed: imputed rents, taxes on products and retained earnings (and the corporate tax). We use specific sources for these three components.

Imputed Rents We use imputed rents from EU-SILC. The EU-SILC survey has started to incorporate an imputed rent variable from EU-SILC in recent years, although it is not included in the headline income statistics published by Eurostat.

Taxes on Products Taxes on products are distributed proportionally to consumption. We measure consumption using the household budget surveys (HBS) collected by Eurostat.

Retained Earnings and the Corporate Tax Retained earnings and the corporate tax are split up into three subcomponents: the share that accrues to the general government, the share that accrues to shareholder households, and the share that accrues to pension funds. The government share does not require additional data since it is distributed like the rest of government income (proportionally). For the rest, we rely on the Household Finance and Consumption Survey (HFCS) (a European wealth survey spearheaded by the ECB) and on the Wealth and Asset Survey (WAS) in the United Kingdom. We identify the shareholdings of households in these surveys, be they public or private stock, held directly or via mutual funds, as long as they correspond to incorporated entities (that is, we exclude unincorporated businesses, which in the SNA are not part of the corporate sector and in the surveys would be recorded as self-employment income). Retained earnings that correspond to household shareholdings are distributed proportionally to this value. Retained earnings that correspond to pension funds are distributed proportionally to labor and pension income.

D.1.5.2 Matching of Additional Income Components to Tabulations

To incorporate additional sources of income to our tabulations, we apply the following procedure. First, we calibrate the surveys from section D.1.5.1 above using the procedure from section D.1.4 to correct for the underrepresentation of the rich.

Second, we create a synthetic dataset by matching the three sources in D.1.5.1 to the calibrated survey microdata. Our statistical matching procedure is straightforward: we rank the sources according to their own internal pretax income variable, and then match observations one-by-one according to their income rank.[5]

Third, we take a tabulation of pretax or posttax income excluding additional income components (i.e., from section D.1.4). To each observation of the synthetic dataset, we attribute the income of the corresponding rank in the tabulation. Then, we rescale the different components to their macroeconomic totals, add them up, and

[5]In practice, because different datasets have different weights and different sample sizes, observations have to be partially matched with one another. For example, imagine that the first (sorted) dataset has the weights $\{3, 1, \dots\}$ and the second one the weights $\{2, 4, \dots\}$. The matched dataset starts with one observation with weight 2 that has the characteristics of the first observation of each dataset. However, the first observation of the first dataset cannot be fully matched because its weight (3) is larger than the weight of the first observation from the second dataset (2). So we keep the first observation in the first dataset with its remaining weight (1), and match it to the second observation of the second dataset. That observation's weight (4) is in turn larger than 1, so we follow the same procedure. We continue the process until all the probability mass from both datasets has been matched. One can show that, if the initial datasets have sizes N and M, the matched dataset will at most have size $N + M - 1$.

calculate the complete distribution of income. When data sources are not available for a given year, we use the value from the closest available year. When they are not available at all for a given country, we use the European average.

D.1.6 Auxiliary Data

D.1.6.1 Income Distribution in the United States

To compare the geography of inequality in Europe with that of the United States, we use distributional national accounts data from Piketty, Saez, and Zucman (2018) and national accounts data by US state.

We attribute national income to each state based on their share of GDP (the only national account aggregate available at the state level). To that end, we use data on total state domestic products from the Bureau of Economic Analysis, along with state adult population series from the National Cancer Institute "Surveillance, Epidemiology and End Results Program".[6]

This provides us with an estimate of national income by state, which lets us compute between-state inequality in the United States. Using the data from Piketty, Saez, and Zucman (2018), we can calculate the overall Theil index for the United States. Using the decomposability of the Theil index, we can then estimate the within-state component of inequality for the United States as a residual.

D.1.6.2 Top Marginal Tax Rates

We construct a database of comprehensive top marginal tax rates that cover 30 countries from 1981 to 2019 (29 European countries plus the United States). Of these 30 countries, 27 are continuously covered from 1981 onward, and the three remaining countries (Bulgaria, Croatia, and Romania) are covered from 2009 onward.

This database is an extension of Kleven et al. (2020), which was itself an extension of data collected by Kleven, Landais, and Saez (2013), Piketty, Saez, and Stantcheva (2014) and Roine, Vlachos, and Waldenström (2009). We extend that database in two ways: we improve the time coverage of countries (in particular Eastern European countries) that were only included for recent years in Kleven et al. (2020). We also collect data on the corporate income tax rate to get a more comprehensive measure of the top marginal tax rate for robustness checks, in line with our inclusion of

[6]State domestic products provided by the Bureau of Economic Analysis go back as far as 1967. We extrapolate these series back to 1929 by using the growth rates in personal income per capita available from Barro and Sala-i-Martin (1992).

undistributed profits in our measure of personal income.

Definition of the Top Marginal Tax Rate Our formula for the top marginal tax rate combines the top personal income tax rate τ_{i}, the payroll tax rate on employees (τ_{pw}) and employers (τ_{pf}) and the VAT or sales tax rate (τ_{c}). This measure combines all marginal tax rates as:

$$1 - \tau = \frac{(1 - \tau_{\mathrm{i}})(1 - \tau_{\mathrm{pw}})}{(1 + \tau_{\mathrm{pf}})(1 + \tau_{\mathrm{c}})}$$

If an individual at the top of the income distribution increases their output by one unit, then they can increase their consumption by $1 - \tau$. We can consider a variant of the formula, which also includes the corporate tax rate $(1 + \tau_{\mathrm{f}})$ at the denominator. This inclusion is a departure from Kleven et al. (2020) and earlier works, and while it makes sense in light of our inclusion of undistributed profits in personal income, there is room for debate. The rationale for including the corporate tax in the formula is that higher corporate tax rates may discourage shareholders from bargaining for a higher share of the company's surplus, and therefore reduce the share of top incomes. Yet the proper measure of the marginal tax rate would ideally depend on the characteristic of each individual top earner (employee or self-employed, via an incorporated business or not, earning mostly labor or capital income, etc.). The inclusion of the corporate tax would be justified in some cases but not others, or at a varying intensity. Moreover, we stress that the way it is included in the formula is *ad hoc* and should be viewed as a pure reduced-form specification. For all these reasons, we report results both including and excluding the corporate tax from the formula.

Top Income Tax and Payroll Tax Rates For top income tax and payroll rates, we extend the database of Kleven et al. (2020) with the OECD tax database (available from 1981 to 2019). The data includes both central and subcentral government tax rates. We cross-check the OECD data with Kleven et al. (2020) to ensure consistent results and conventions.

Value-Added Taxes We extend the data of Kleven et al. (2020) using the OECD's data on Value Added Tax/Goods and Services Tax (VAT/GST), which covers the years 1976 to 2020. We use the standard rate (i.e. we ignore reduced rates on certain products or specific regional rates).

Corporate Income tax Rate Our corporate income tax rate is the "Combined corporate income tax rate" estimated by the OECD.

D.1.7 Country-Specific Estimations of Top Shares from Tabulated Tax Returns

D.1.7.1 Austria

Our data for Austria comes from Altzinger et al. (2010), who use tax data from the *Integrierte Lohn und Einkommensteuerstatistik* (LUE) to study the evolution of top income shares between 1976 and 2006. We complete their series by gathering more recent LUE tabulations from Statistics Austria (2008-2015). These tabulations cover the entire Austrian population and can therefore be directly used to compute top income shares. We turn the tabulations into complete distributions by using generalized Pareto interpolation (Blanchet, Fournier, and Piketty, 2021). Our results for more recent years are very consistent with those found by Altzinger et al. (2010): before 2010, the top 10% income share remained very stable around 33% and the top 1% share decreased from 10% to 9%.

D.1.7.2 East Germany

Our data for the distribution of East German income comes from a yearly publication of official statistics on the economy of East Germany (*Statistisches Jahrbuch der deutschen Demokratischen Republik*). The 1990 edition of that publication provides estimates of the population by income bracket and by type of household over the period 1980–1990. We interpolate the distribution for each type of household (Blanchet, Fournier, and Piketty, 2021), and then merge them into a single distribution after having multiplied the number of observations corresponding to each type of household by the number of adults in the corresponding type of household. That way, we get a distribution for equal-split adults.

That data relate to the distribution of posttax income only. As an approximation, we use the same distribution for pretax income. The distinction between pretax and posttax income in socialist economies was indeed less salient than it is today: see Bukowski and Novokmet (2017a) for a detailed discussion of that issue in the case of Poland.

D.1.7.3 Estonia

We estimate top income shares for Estonia by exploiting tabulated tax returns from various reports of the Tax and Customs Board. Tabulations are available from 2002 to 2017. For each year, they provide information on the total number of taxpayers and total taxable income for various income brackets. The income tax in Estonia

Figure D.1: Top 10% income share in Estonia:
survey data vs. tax data vs. corrected survey

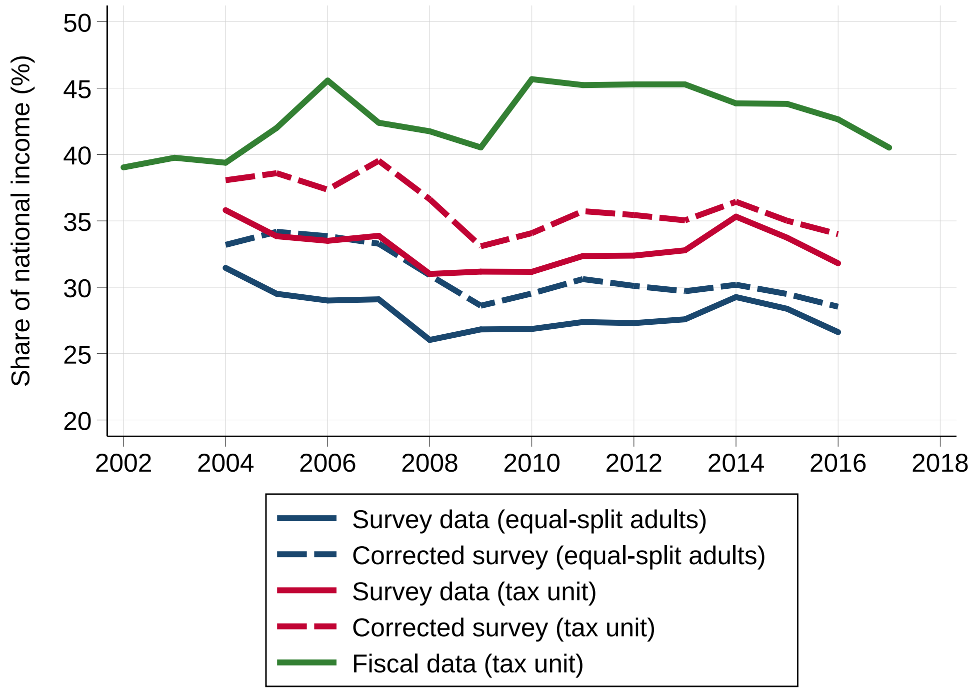

is a flat tax, collected on individual earnings. It applies to most sources of income (income from work, interest income, royalties, dividends...), which are taxed on a gross basis.

We use these tabulations to estimate top income shares by matching them with survey microdata from EU-SILC in the following way. We first use generalized Pareto interpolation techniques (Blanchet, Fournier, and Piketty, 2021) to compute thresholds and average incomes for various quantiles of the fiscal income distribution. We then correct the EU-SILC survey by using the Blanchet, Flores, and Morgan (2018) method (BFM), which exploits the fiscal data to reweigh survey observations so that top incomes are properly represented. Since the BFM method preserves the survey microdata, and in particular other covariates, it allows us to directly account for the fact that (1) the unit of observation in the tax data is the individual, not the equal-split adult and (2) taxable income includes gross components that must be deducted to obtain pretax income estimates. We can therefore directly compute the share of pretax income accruing to top earners in the corrected survey by changing the unit of observation and the income concept *after* having reweighed survey observations.

Figure D.1 compares the top 10% income share estimated from survey data, tax data and corrected survey data. Inequality is highest when measured directly from tax tabulations since many individuals have zero taxable income, mainly due to the possibility to deduct some expenditures. Correcting the survey for the under-representation of top incomes increases significantly the top 10% income share, even if the overall trend is not substantially affected. Unsurprisingly, inequality is lower between equal-split adults than between tax units (here, individuals) since the former does not account for within-household heterogeneity. Our final estimates show a decrease in the top 10% income share from 35% in 2004 to 30% in 2016. Since survey microdata is not available for 2002, 2003 and 2017, we extrapolate top income shares to these years by using the average ratio of pretax income between fiscal data series and corrected survey estimates over the 2004-2016 period, by generalized percentile.

D.1.7.4 Greece

Our data for Greece comes from Chrissis and Koutentakis (2017), who used published tax tabulations to measure the evolution of top income shares from 1967 to 2017. By combining these tabulations with control totals for income and the adult population, they estimate that the top 10% fiscal income share varied between 23% and 29% over the period. This appears surprisingly low compared to results from other European countries, especially given that the unit of observation is the individual.

One specific concern with the Greek case has to do with tax evasion, which has previously been found to be particularly pronounced at the top of the distribution. Based on a matched samples of income taxpayers and respondents from the household budget survey, Matsaganis and Flevotomou (2010) find that top 1% earners report incomes which are 23.6% *lower* in the tax data than in the survey. This result is consistent with our own results obtained from the EU-SILC survey, where we find the top 10% pretax income share (among individual adults) to fluctuate around 35% between 2006 and 2015. The under-representation of top incomes if Greek tax data therefore threatens the comparability of our estimates and calls for a specific adjustment.

In order to correct Greek top income shares, we proceed as follows. First, we define a new "taxable income" concept in the EU-SILC survey such that we artificially reduce the pretax incomes of individuals based on the coefficients provided by Matsaganis and Flevotomou (2010) on underreporting by income decile and the top 1%. Then, we interpolate the fiscal income averages of Chrissis and Koutentakis (2017) using generalized Pareto interpolation (Blanchet, Fournier, and Piketty, 2021) and we apply the Blanchet, Flores, and Morgan (2018) method to rescale our new taxable income concept to the fiscal data. Finally, we use the reweighed survey to compute top income shares in our concept of interest, that is pretax income splitted equally among spouses, and we correct top income shares before 2008 by extrapolating the correction coefficient by percentile that we obtained from the correction. This method has the advantage of fully exploiting the tax data, which is more granular at the very top of the distribution and covers every year from 1980 to 2017, while at the same time accounting for tax evasion in a simple way. That being said, we stress that this adjustment is far from being sufficient, so that distributional data for Greece should be interpreted with care. As tax evasion is increasingly tackled by tax authorities, future research will hopefully be able to obtain more reliable estimates.

D.1.7.5 Iceland

For Iceland, we directly use tax data available online since 1990 from Statistics Iceland. Given that Iceland has had a flat—or nearly flat—comprehensive income tax over the entire period, the entire distribution is covered, so we use it to directly compute top income shares.

D.1.7.6 Italy

Top income shares for Italy are available from the World Inequality Database from 1980 to 2009 thanks to previous work done by Alvaredo and Pisano (2010). We update their series by collecting tax tabulations available from the data portal of the Italian ministry of Finance.[7] These tabulations are available over the 2008-2016 period and provide information on the number of taxpayers and total taxable income for different income brackets.

The income tax in Italy applies to individuals and includes most income components on a gross basis, except for interest income, which is not taxed. We compute top income shares over the 2008-2016 period by using the exact same methodology as the one used for Estonia (see above). That is, we use the method developed by Blanchet, Flores, and Morgan (2018) to reweigh the survey and compute income shares that are both representative of top incomes and consistent with the benchmark income concept and population unit used in this paper.

Figure D.2 compares the top 10% income share estimated from survey data, tax data and corrected survey data. Tax data leads to increasing inequality less than in Estonia, perhaps because some components of capital income are not reported in the tabulated tax returns. For the two years for which we can compare our estimates

[7] See `http://www1.finanze.gov.it/finanze3/pagina_dichiarazioni/dichiarazioni.php`.

Figure D.2: Top 10% income share in Italy:
survey data vs. tax data vs. corrected survey

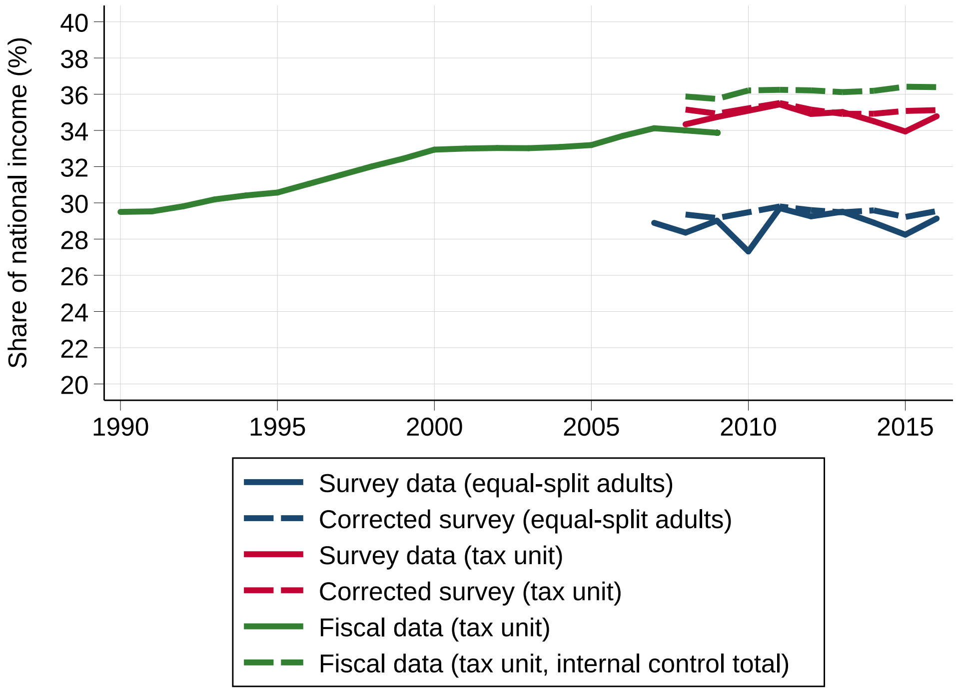

with that of Alvaredo and Pisano (2010), 2008 and 2009, the top 10% income shares coincide almost perfectly, which suggests that both methods are alternative and complementary ways of obtaining robust estimates of the evolution of top incomes. Changing the welfare concept from individual taxable income to pretax income per adult decreased the top 10% share by about 4 percentage points. We use this relationship to correct conceptual discrepancies in Italian top income shares over the 1980-2009 period. For each generalized percentile among the top decile, we compute the ratio of average taxable individual income to pretax income per adult over the 2009-2016 period. We then use the average ratio over this period to harmonize top income share series in previous years.

D.1.7.7 Luxembourg

For Luxembourg, we use two years of tax data that were published as part of reports by the Conseil Économique et Social (*Analyse des données fiscales au Luxembourg, 2015* and *Analyse des données fiscales au Luxembourg, 2018*) (Conseil Economique et Social, 2015, 2018). These contain detailed tabulations that cover the income of resident households for two years, 2010 and 2012.

We interpolate these two distributions using generalized Pareto interpolation (Blanchet,

Figure D.3: Top 10% income share in Luxembourg:
survey data vs. tax data vs. corrected survey

(a) Correction for the two years with tax data

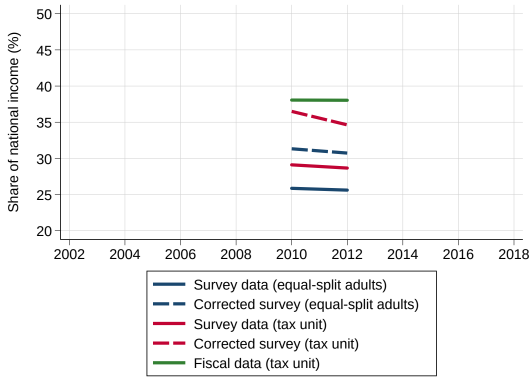

(b) Extrapolation of the correction

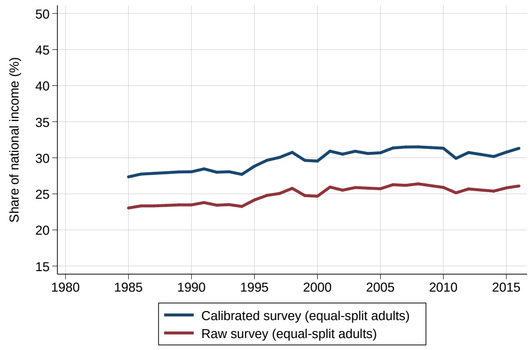

Fournier, and Piketty, 2021) and then correct the EU-SILC data in the two corresponding years using the method of Blanchet, Flores, and Morgan (2018). The correction is very similar for both years, with the top 10% share increasing by roughly 5pp (see figure D.3a). We then extrapolate that correction to previous years by extrapolating the correction coefficient by percentile that we obtained from the tax data correction (see figure D.3b).

D.1.7.8 Portugal

Top income shares for Portugal are available from the World Inequality Database from 1980 to 2009 thanks to the work done by Alvaredo (2009). We update these series by collecting tax tabulations available from the data portal Pordata.[8] These tabulations are available over the 1990-2016 period and provide information on the number of taxpayers and total taxable income for different income brackets.

The income tax in Portugal applies to most income components on a gross basis, except for most capital gains and all interest income, which are not taxed. The unit observed in the tax data is the married couple, or single adult. We compute top income shares over the 2007-2016 period by using the exact same methodology as the one used for Estonia (see above). That is, we use the method developed by Blanchet, Flores, and Morgan (2018) to reweigh the survey and compute income shares that are both representative of top incomes and consistent with the benchmark income concept and population unit used in this paper. In the case of Portugal, since tax

[8] See `https://www.pordata.pt`.

Figure D.4: Top 10% income share in Portugal:
survey data vs. tax data vs. corrected survey

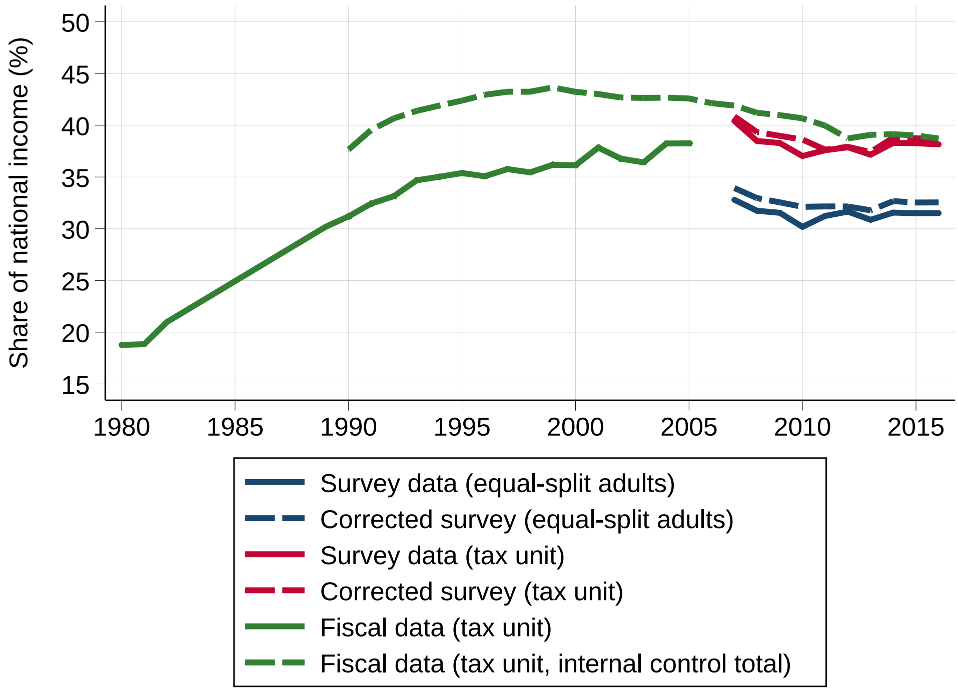

units are either individuals or married couples, we first match couples in the EU-SILC survey and aggregate their incomes. We are then able to use tax tabulations to correct for the under-representation of "top tax units" in the survey.

Figure D.4 compares the top 10% income share estimated from survey data, tax data and corrected survey data. Using tax data leads to only moderately higher inequality, perhaps because some components of capital income are not taxed. While there is a gap in the Alvaredo (2009) series and our series between 2005 and 2007, comparing the two estimates suggests that using the BFM methodology leads to a slightly higher top 10% income share, which might be due to the income control being too high in previous estimates. We use our estimates to correct conceptual discrepancies in Portuguese top income shares in previous years. First, we extrapolate our series back to 2005 by using the trends observed in the fiscal data (with internal income control) over the 2005-2007 period. For each generalized percentile among the top decile, we then use the ratio of average taxable income per tax unit to pretax income per adult in 2005 to harmonize top income shares before 2005.

Figure D.5: Top Marginal Tax Rate and Inequality in Europe: Time Series

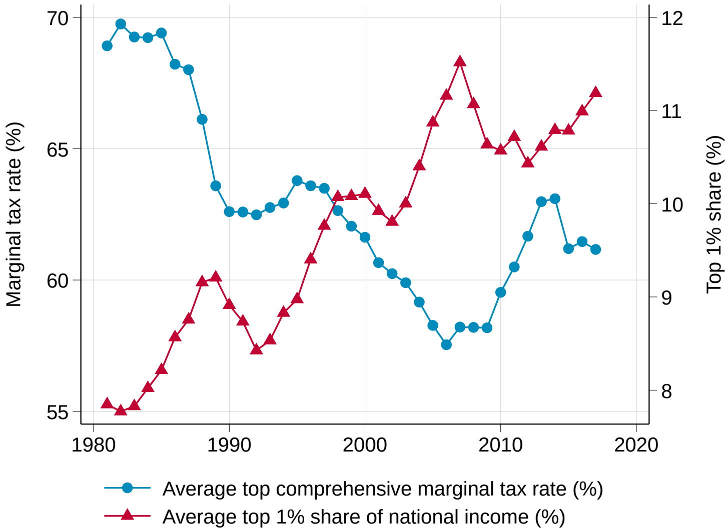

Source: Authors' estimation, see main text. Marginal tax rate does not include the corporate tax. *Note:* Estimates refer to population-weighted averages of European countries with data available since 1981 (Austria, Belgium, Denmark, Finland, France, Germany, Greece, Ireland, Italy, Luxembourg, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, United Kingdom).

Figure D.6: Top Marginal Tax Rate and Inequality in Europe: Cross-country Evidence

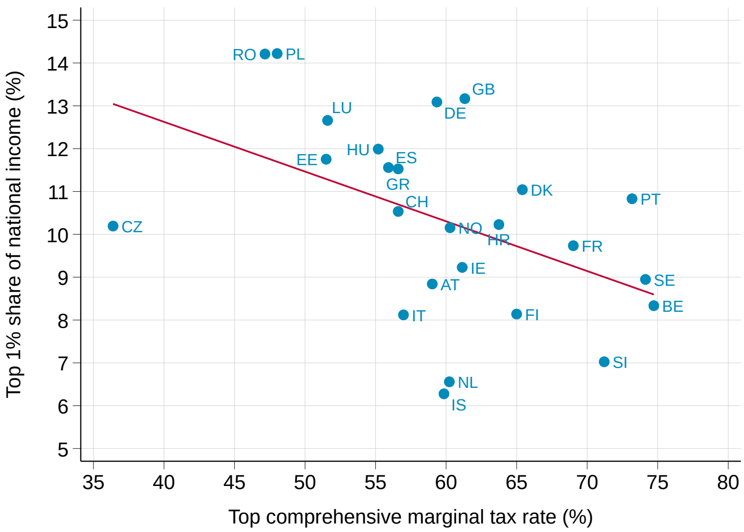

Source: Authors' estimation, see main text. Marginal tax rate does not include the corporate tax.

Table D.4: Elasticity of the Top 1% Share
With Respect to the Top Marginal Tax Rate

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
elasticity	0.45	0.41	0.38	0.33	0.40	0.39	0.13	0.12
95% CI	[0.24, 0.86]	[0.23, 0.74]	[0.13, 0.83]	[0.12, 0.70]	[0.25, 0.71]	[0.25, 0.67]	[−0.01, 0.33]	[−0.03, 0.29]
observations	827	827	827	827	827	827	827	827
clusters	26	26	26	26	26	26	26	26
R^2	0.26	0.24	0.35	0.32	0.68	0.69	0.80	0.80
incl. corporate tax		×		×		×		×
year fixed effects			×	×			×	×
country fixed effects					×	×	×	×

Confidence intervals adjusted for heteroscedasticity and clustered by country using the wild bootstrap test (Roodman et al., 2019).

D.1.7.9 Romania

Our data for Romania comes from Oancea, Andrei, and Pirjol (2017). The authors had access to the universe of individual income tax returns for 2013 and provide detailed information on the distribution of taxable income. The income tax data covers about 45% of the adult population. We correct the EU-SILC data in 2013 using the method of Blanchet, Flores, and Morgan (2018). The comparison of the survey with the tax data reveals that top earners are strongly underrepresented in EU-SILC: the average income of the top 1% is below 70,000 lei in the surveys compared to 150,000 lei in the tax data. The correction increases the top 10% income share from 26% to 31% and the top 1% share from 5% to 8%. We extend that correction to previous years by extrapolating the coefficient by percentile that we obtained from the correction.

D.1.7.10 Serbia

Our data for Serbia comes from the Statistical Office of the Republic of Serbia, which provided us with detailed tabulations on the pretax income of Serbian taxpayers in 2017 and 2018. Income shown in the tables are taken over from the Individual tax return form on accrued taxes and contributions (PPP-PD form), which is submitted to the Tax Administration. The data covers employees, founders and members of companies employed in their company, persons insured on the basis of independent activity including independent artists, persons insured on the basis of agricultural activities, persons not provided on other grounds, non-residents, disabled persons, military insured persons, pensioners self-employed, pensioners on the basis of employment, military pensioners and agricultural retirees. As a simple approximation, we use the 2017 tabulation to directly calibrate the 2016 EU-SILC survey with the Blanchet, Flores, and Morgan (2018) method.

Table D.5: Elasticity of the Bottom 50% Share
With Respect to Redistribution to the Bottom 50%

	(1)	(2)	(3)	(4)
elasticity	0.10	0.10	0.01	0.01
95% CI	$[-0.01, 0.23]$	$[-0.01, 0.23]$	$[-0.02, 0.05]$	$[-0.02, 0.04]$
observations	271	271	271	271
clusters	26	26	26	26
R^2	0.20	0.20	0.96	0.96
year fixed effects		×		×
country fixed effects			×	×

Confidence intervals adjusted for heteroscedasticity and clustered by country using a wild bootstrap test (Roodman et al., 2019).

D.1.8 Indirect Effect of Top Marginal tax Rates on Pretax Inequality

As shown in figure D.5, the rise of the top 1% pretax income share in Europe has been concomitant to a decrease in the top marginal tax rate. A similar pattern can be found across countries, as shown by figure D.6.

Following Piketty, Saez, and Stantcheva (2014), we estimate an elasticity of the top 1% share with respect to (one minus) the top marginal tax rate using the following model:

$$\log(\text{top } 1\% \text{ pretax income share}) = \beta + \sigma \log(1 - \text{top marginal tax rate})$$

where σ is our estimate of the elasticity. Table D.4 shows estimates of σ across a range of specifications. The inclusion of country fixed effects attenuates the estimate of the elasticity most significantly, which shows that the effect is mostly estimated from cross-country variations. The inclusion or exclusion of the corporate tax from our measure of the top marginal tax rate makes little difference.

D.1.9 Indirect Effect of Transfers on Pretax Inequality

We measure the elasticity between the pretax income share of the bottom 50% and redistribution to the bottom 50% by running the regression:

$$\log\left(\text{share}^{\text{pretax}}_{\text{bottom } 50\%}\right) = \beta + \sigma \log\left(\text{share}^{\text{posttax}}_{\text{bottom } 50\%} - \text{share}^{\text{pretax}}_{\text{bottom } 50\%}\right)$$

and use σ as our estimate of the elasticity. Table D.5 reports estimates across several specifications, which include different sets of fixed effects.

D.2 Additional figures and tables

D.2.1 Methodology and national accounts

Figure D.7: Level and composition of capital income in Europe, 1980-2017

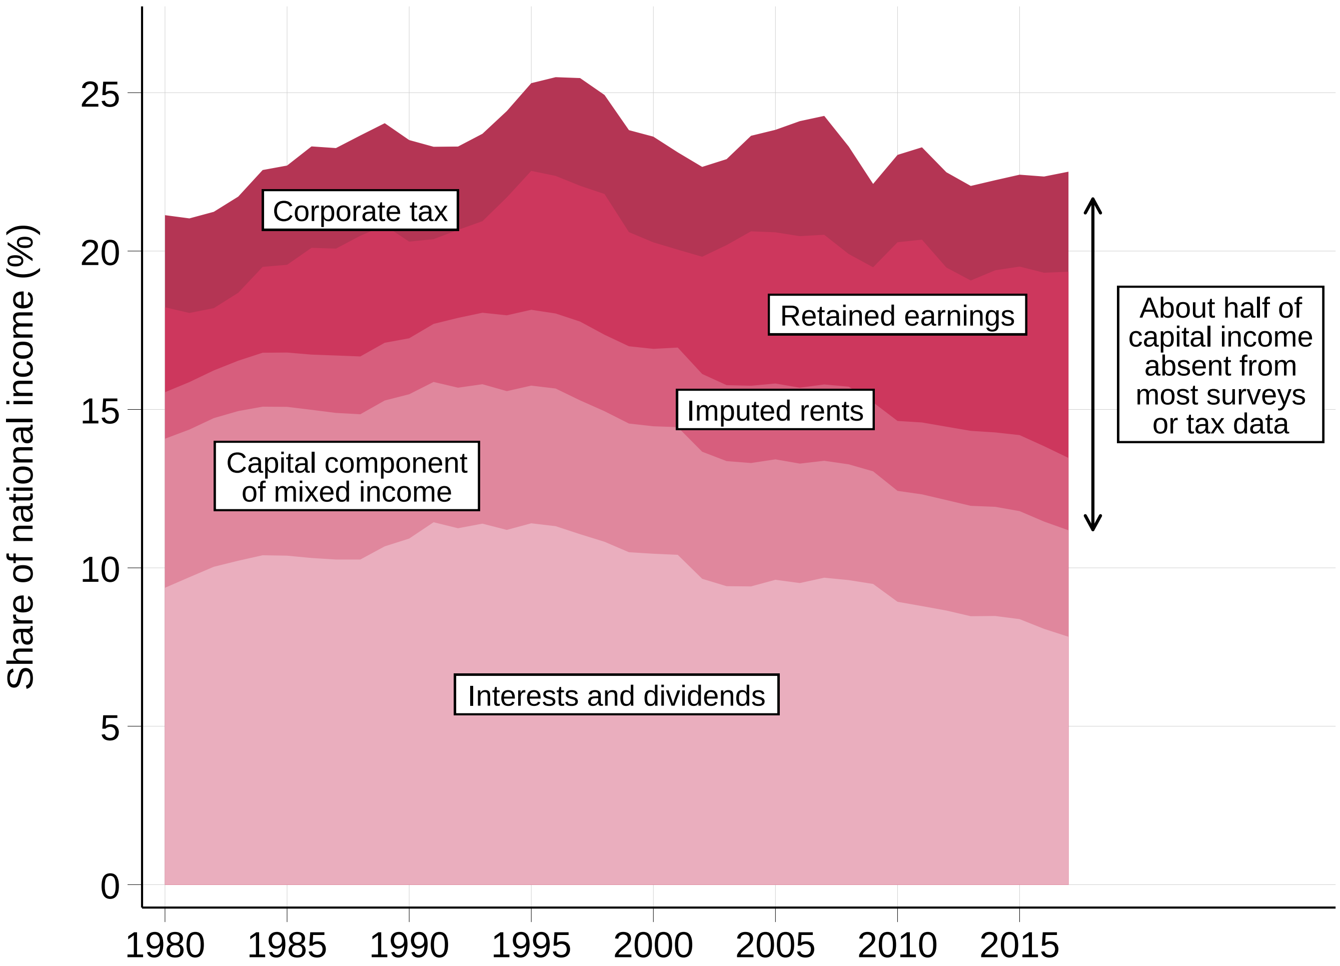

Notes. The figure plots the share of capital income in overall European income – equal to the sum of all European national incomes – between 1980 and 2015. The capital component of mixed income is assumed to be equal to one third of mixed income.

Figure D.8: Level and composition of government final expenditures in Europe, 1980–2017

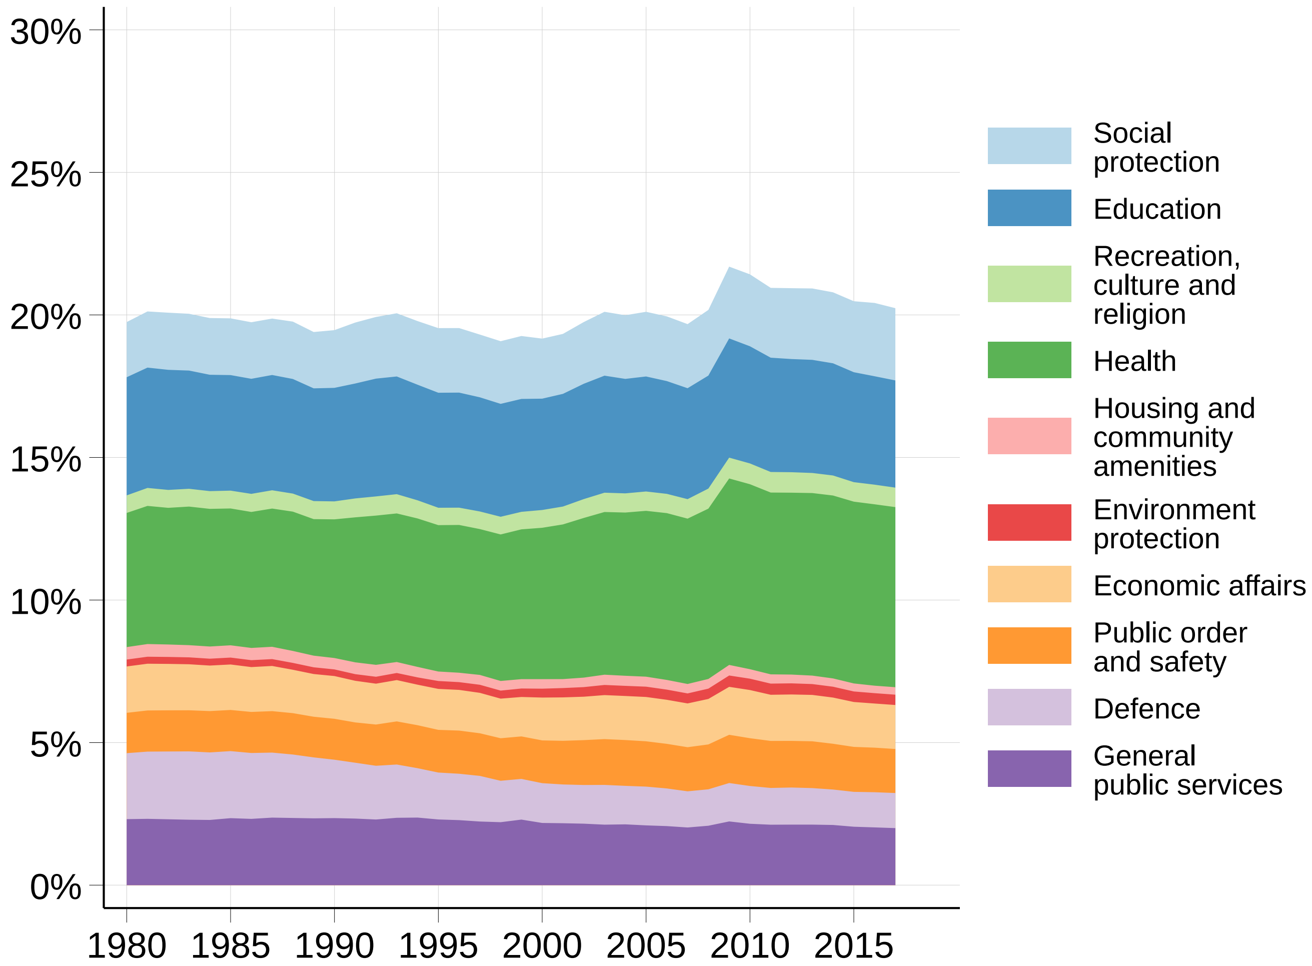

Sources: Government expenditures by function (COFOG) tables from the OECD and the UN SNA. OECD health database for health spending. *Notes.* The figure plots the total value of government final consumption expenditures as a share of national income, and its decomposition into the different functions of government.

Figure D.9: Average regional incomes per adult relative to European-wide average, 1980-2017

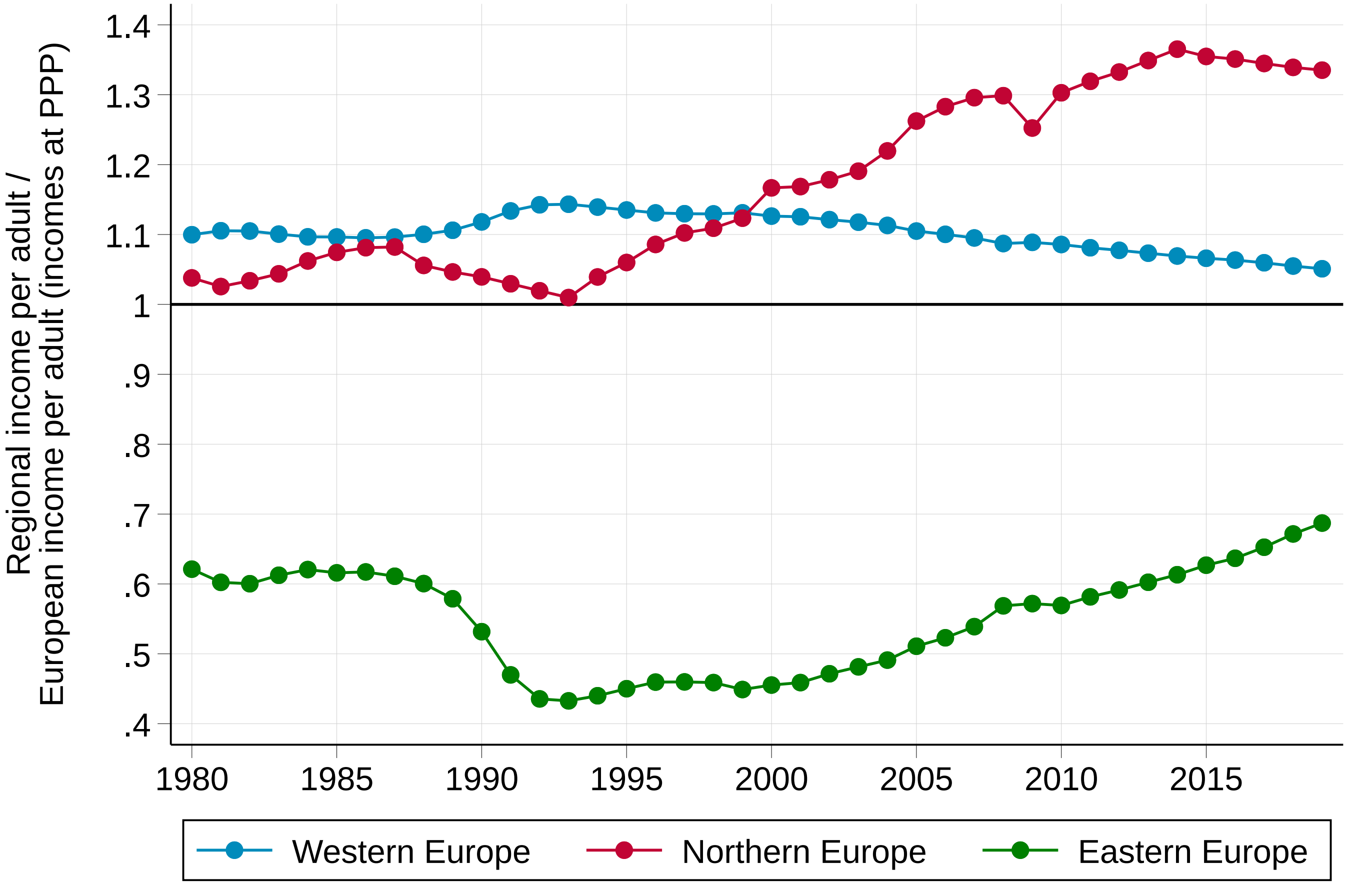

Notes. Western Europe includes Austria, Belgium, Cyprus, France, Germany, Greece, Ireland, Italy, Luxembourg, Malta, the Netherlands, Portugal, Spain, Switzerland, and the United Kingdom. Northern Europe includes Sweden, Norway, Finland, Denmark, and Iceland. Eastern Europe includes the remaining European countries.

Figure D.10: The level and composition of taxes in Europe and the United States, 2007-2017

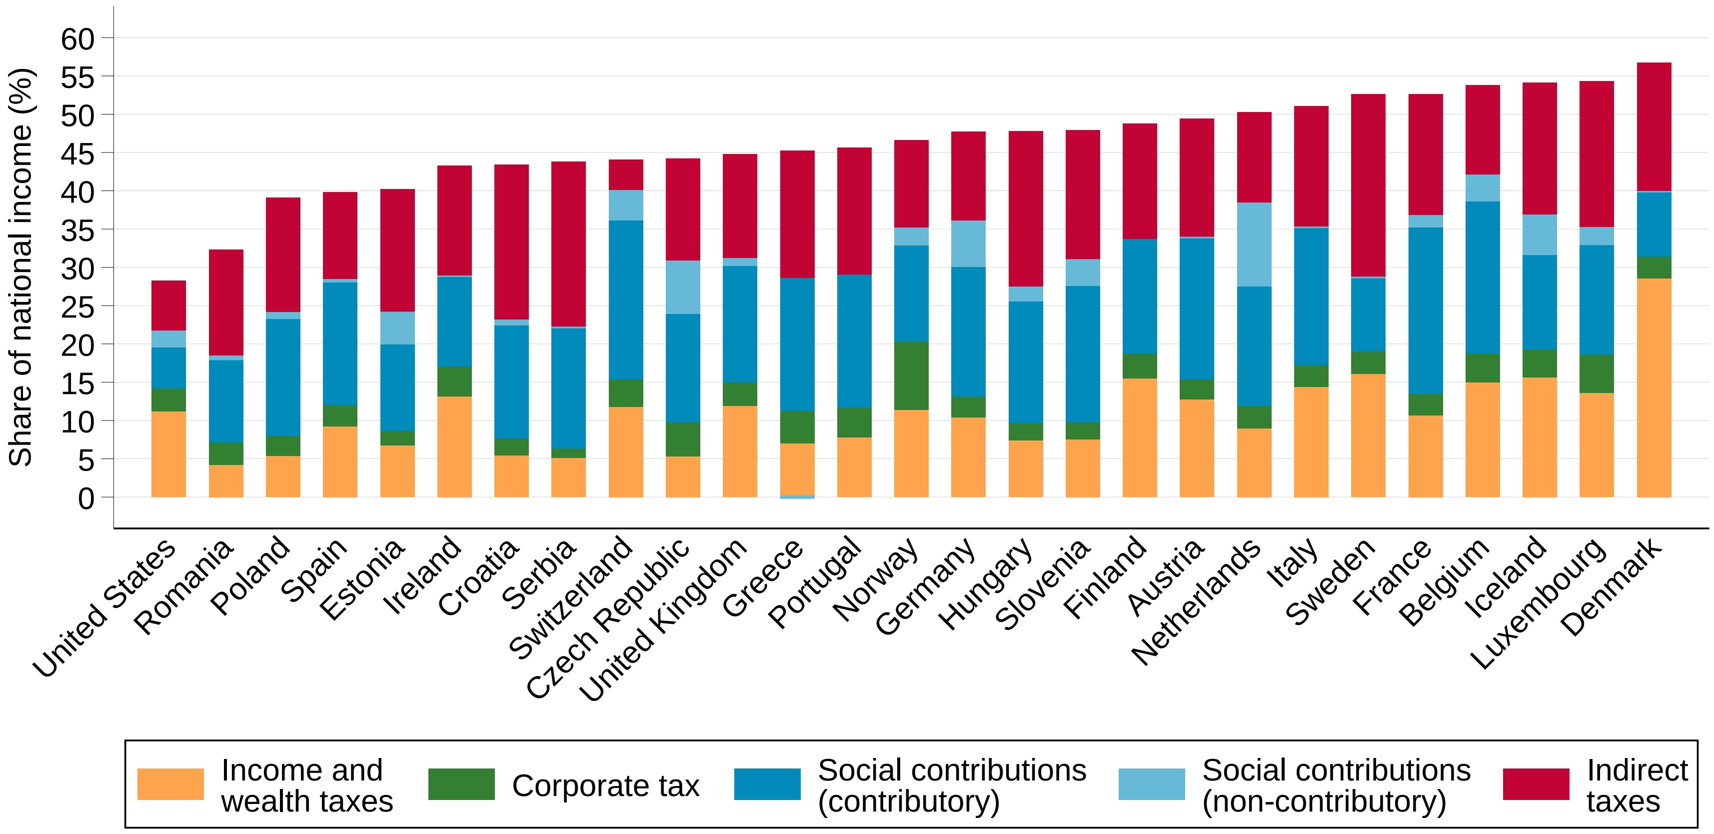

Source. Authors' computations using national accounts.

Figure D.11: The level and composition of taxes in Europe and the United States, 2007-2017 (non-contributory taxes)

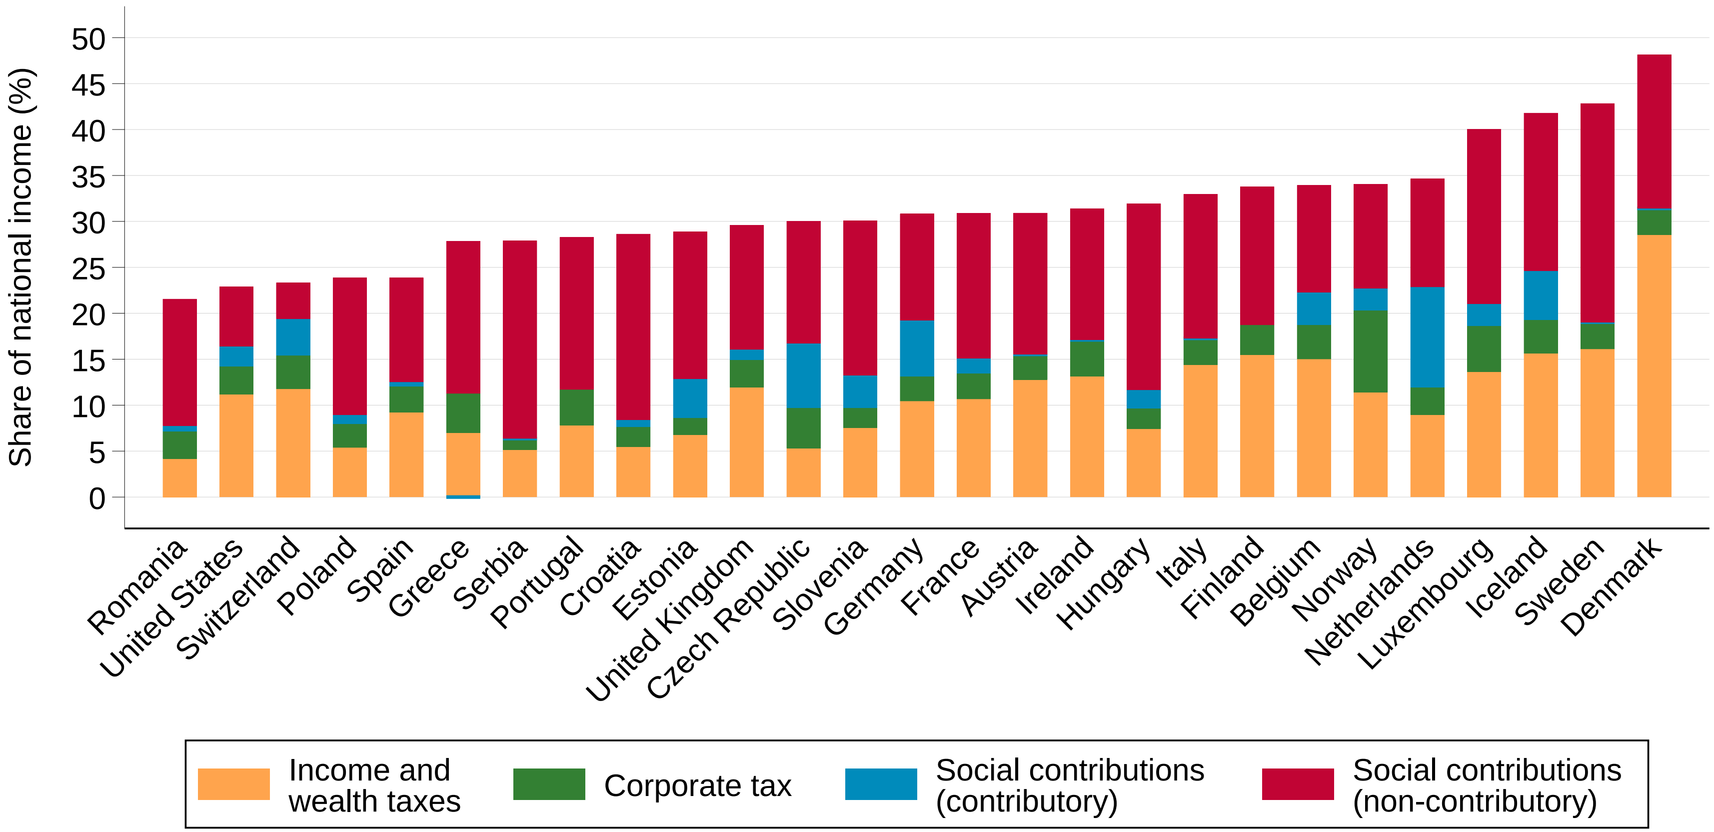

Source. Authors' computations using national accounts.

Figure D.12: The level and composition of transfers in Europe and the United States, 2007-2017

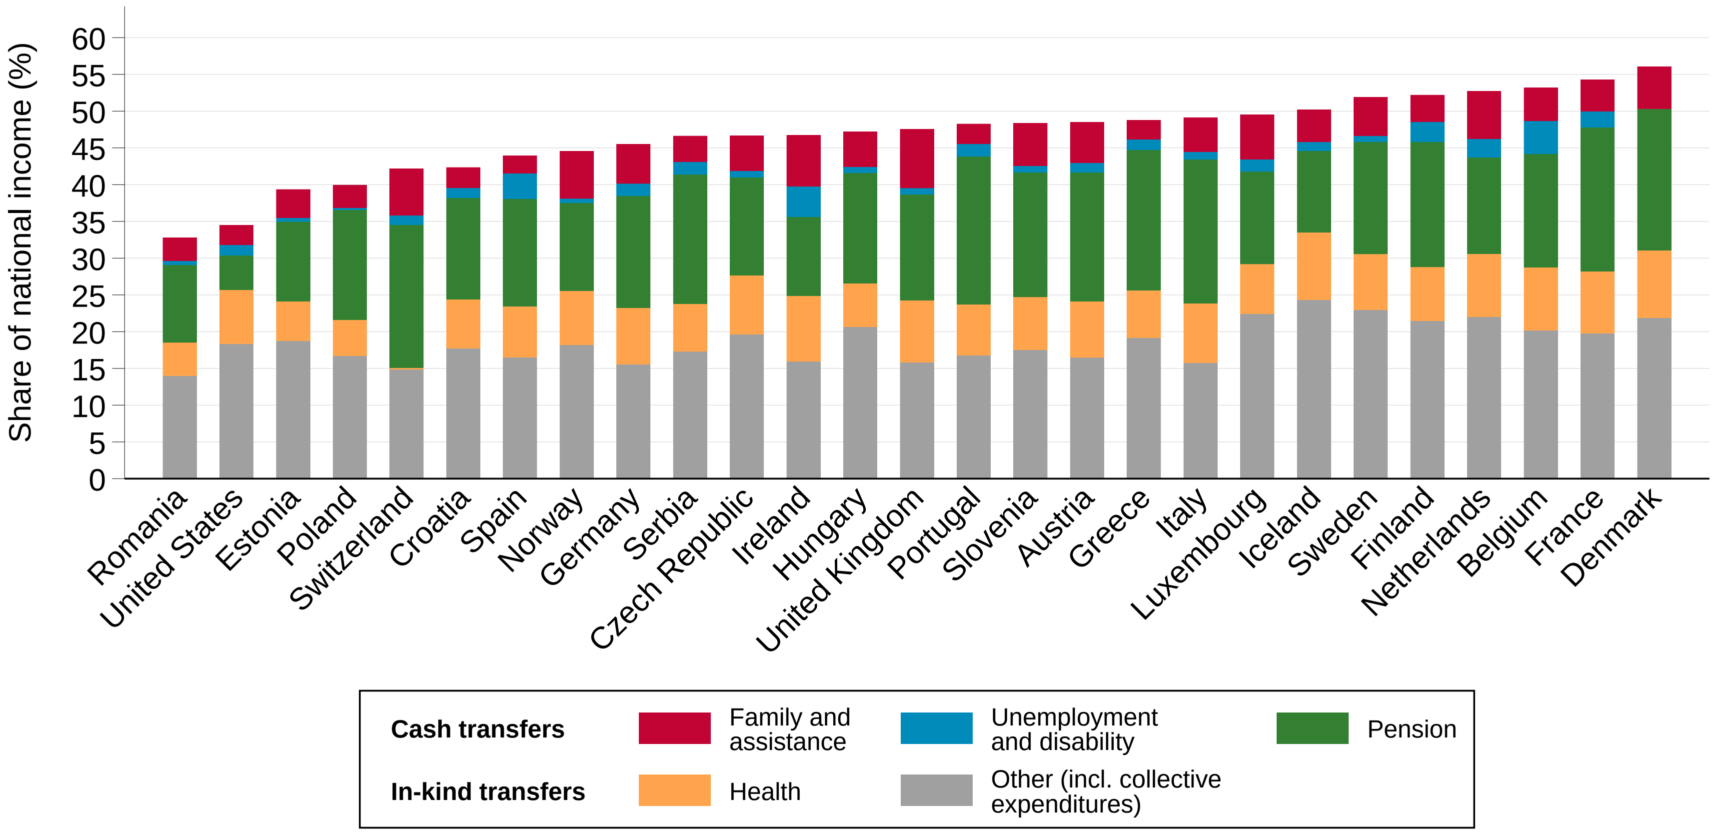

Source. Authors' computations using national accounts.

Figure D.13: Average posttax income quintile share ratio in the European Union: Eurostat vs. posttax disposable income vs. posttax national income

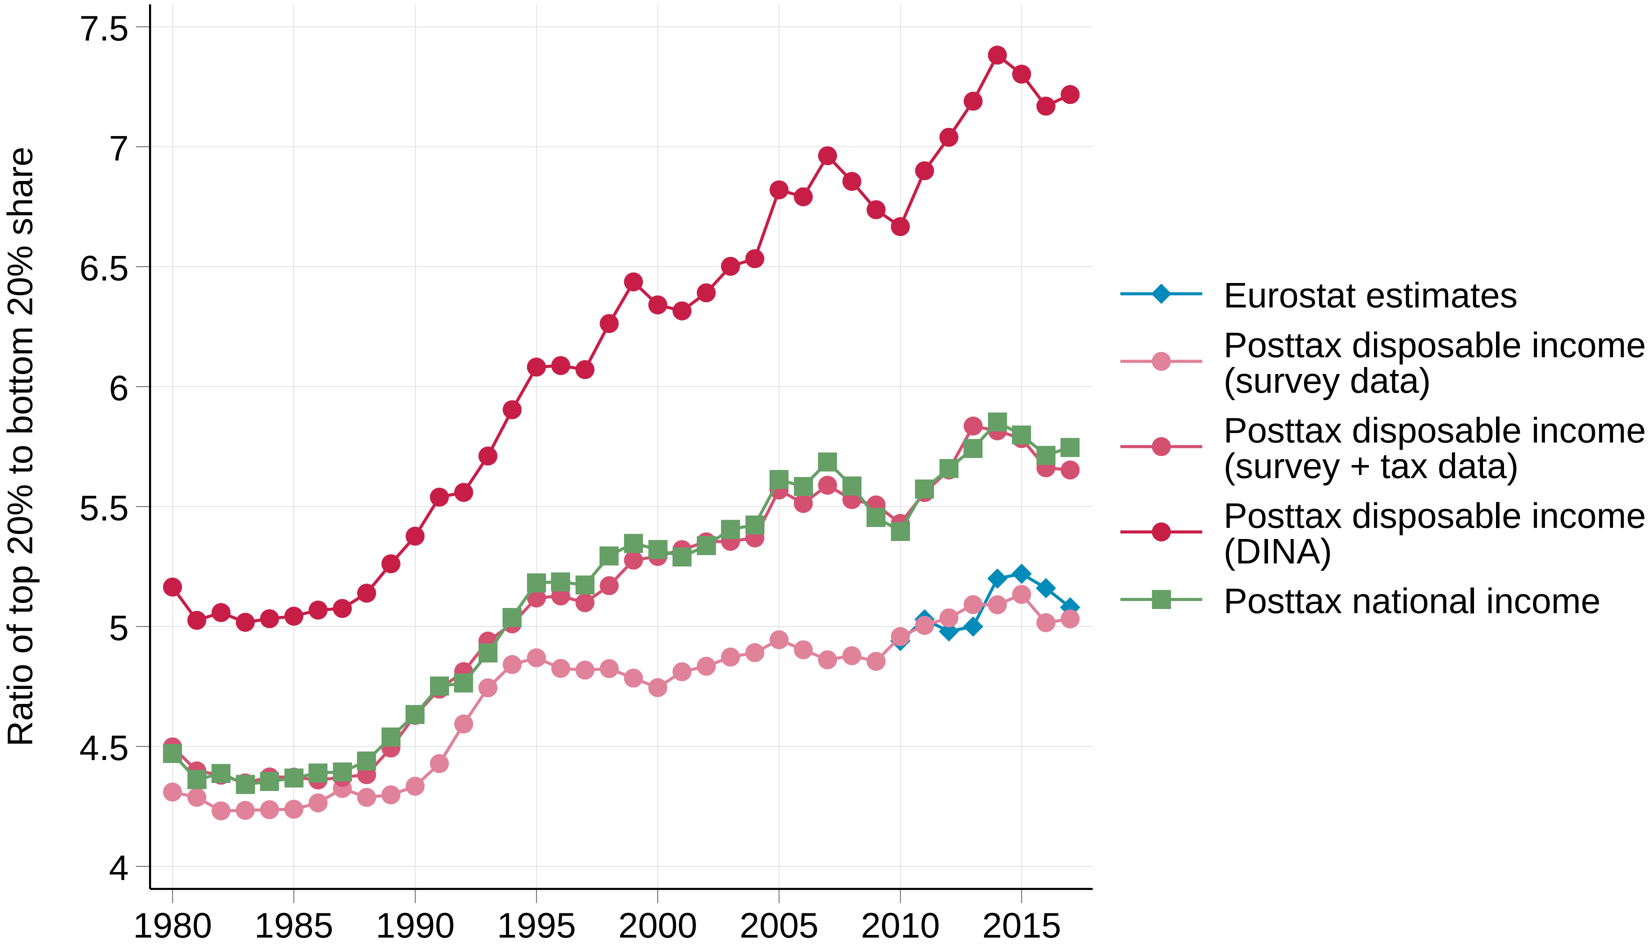

Notes. The figure compares the evolution of the average posttax income quintile share ratio (the share of the top 20% over the share of the bottom 20%), in the European Union (28 countries) between 1980 and 2017. The figure corresponds to population-weighed averages of the indicator. Posttax disposable income corresponds to income after taxes and transfers, but excluding collective government expenditures. Posttax national income includes collective government expenditures (see methodology).

Figure D.14: Posttax income quintile share ratio in Europe: DINA vs. Eurostat

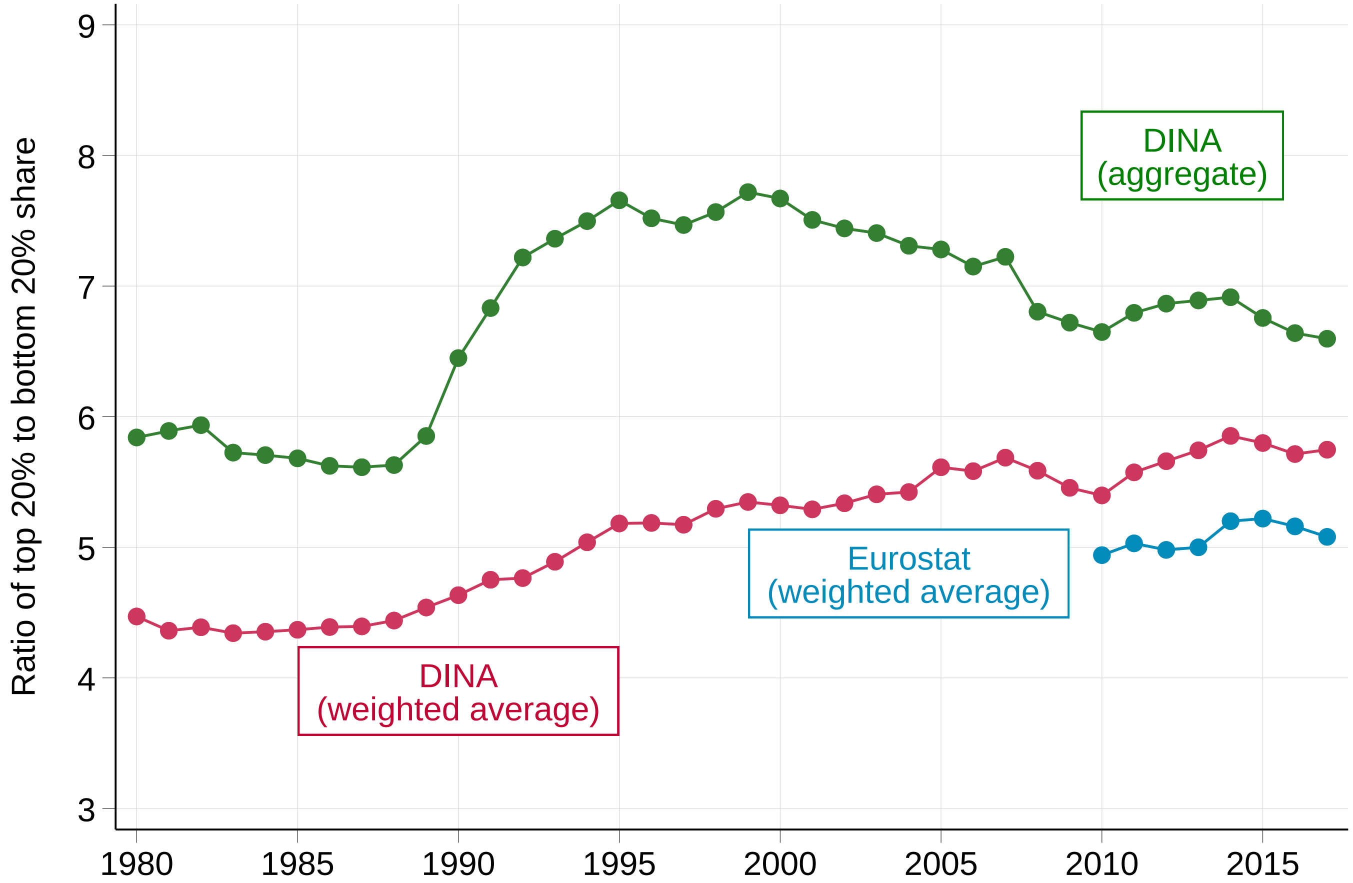

Notes. The figure plots the ratio of the top 20% posttax income share to the bottom 20% posttax income share in the European Union (28 countries) between 1980 and 2017. Eurostat estimates correspond to population-weighed averages of posttax disposable income quintile share ratios. DINA estimates correspond to posttax national income series (see methodology).

Figure D.15: Comparison of our Results with Other DINA Studies in France: Bottom 50% Share

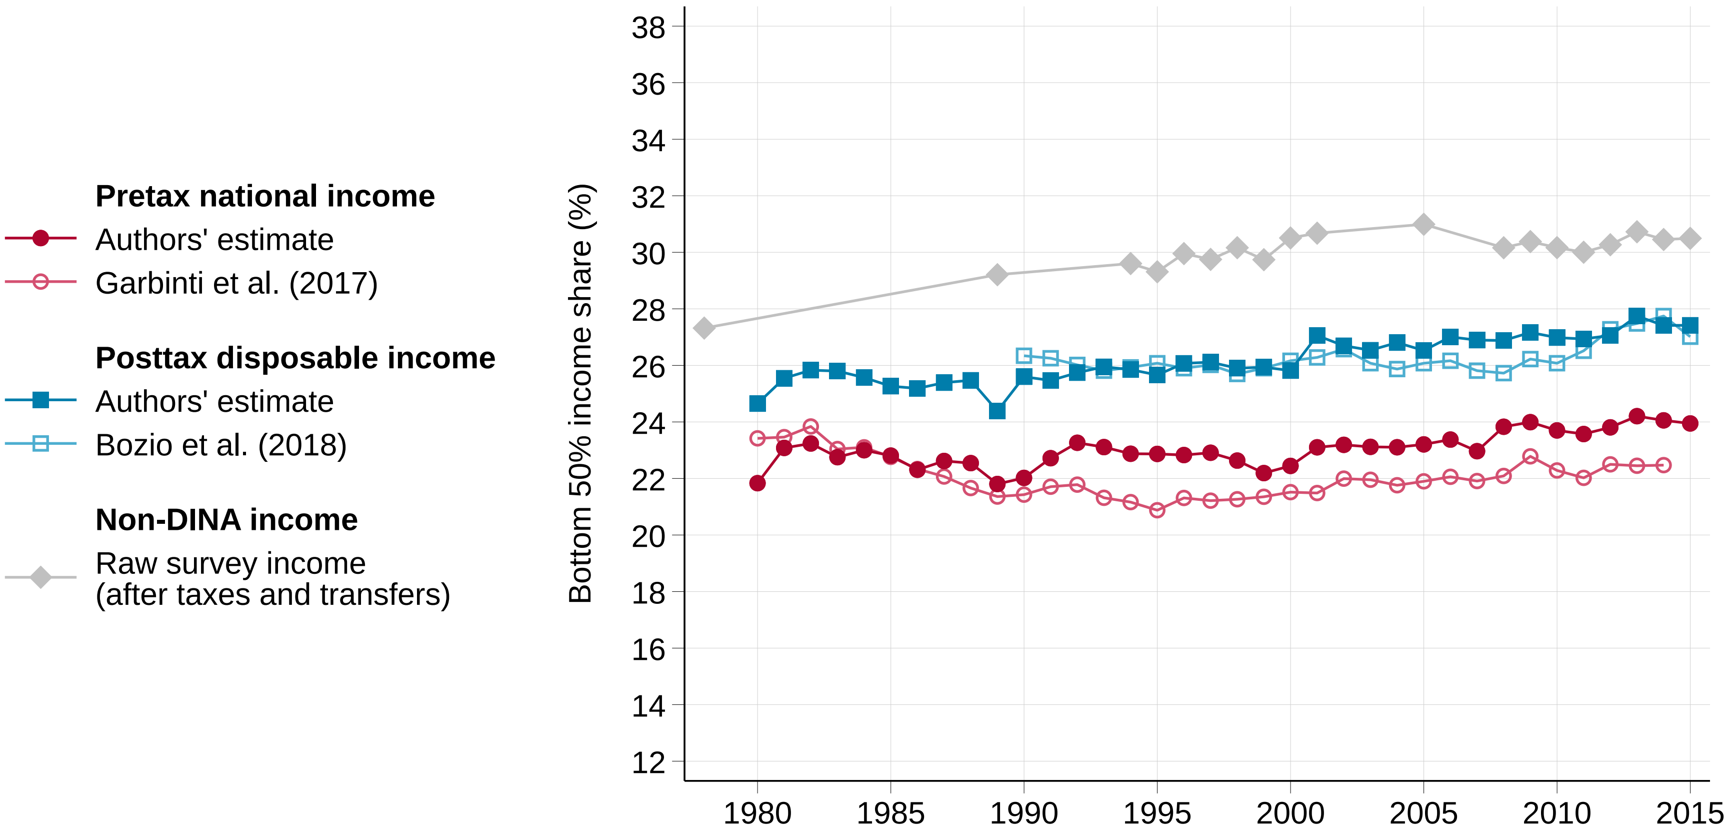

Source: Authors' computations combining surveys, tax data and national accounts. *Note*: The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses, except for the "raw survey income" series in the bottom panel for which income is split equally among all adult household members.

Figure D.16: From surveys to DINA: top 10% pretax income share by country, 2017

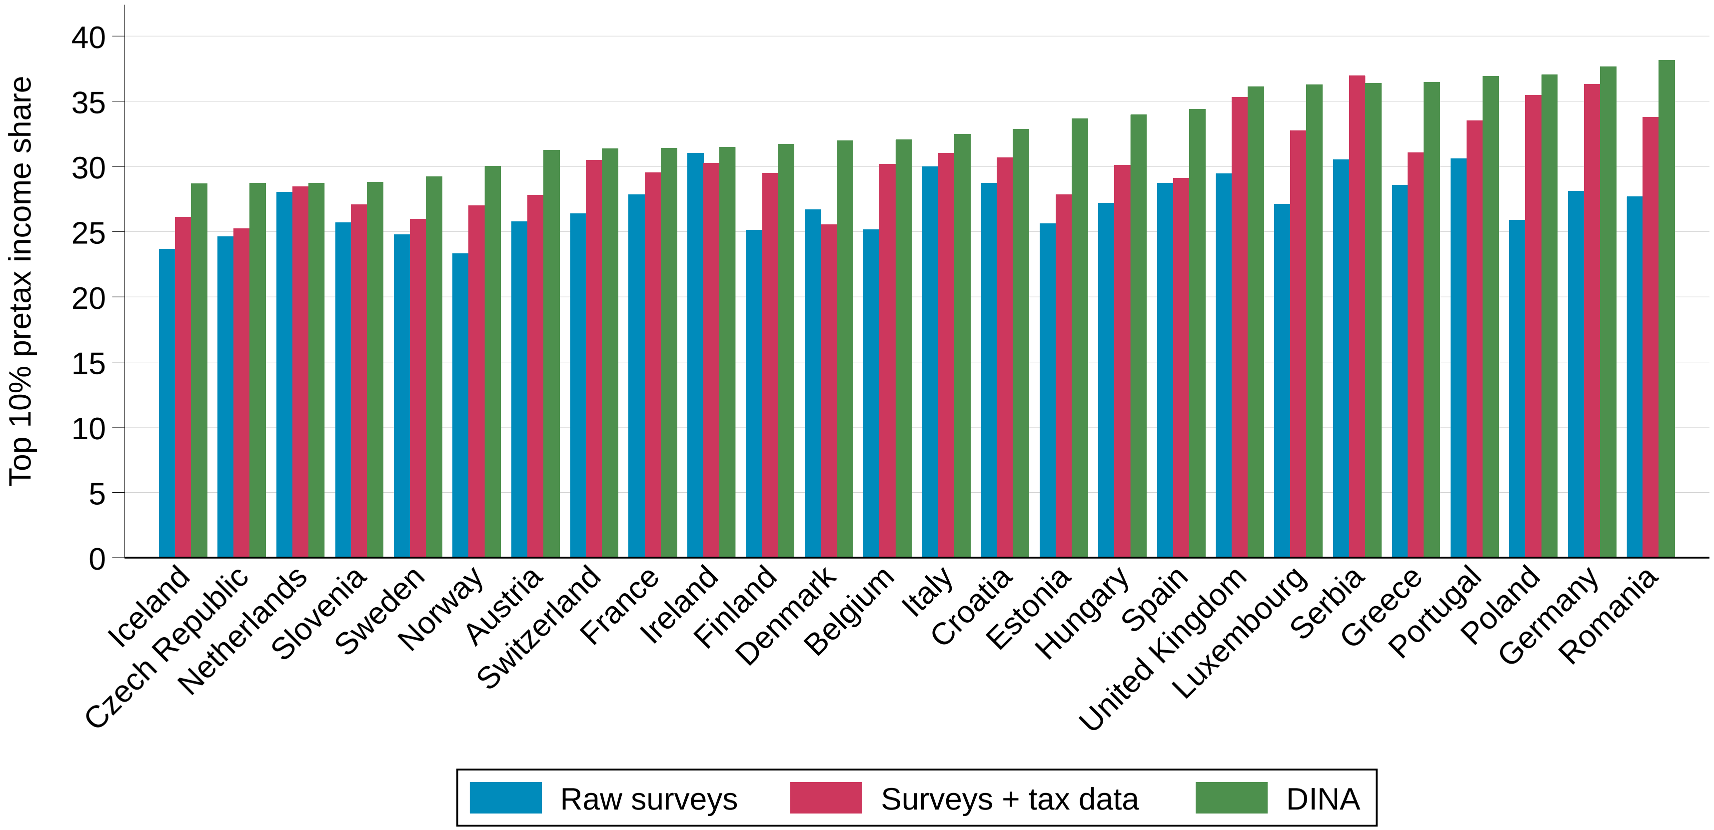

Source: Authors' computations combining surveys, tax data and national accounts. *Note*: The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure D.17: From surveys to DINA: percentage point change in estimated top 10% pretax income share by country, 2017

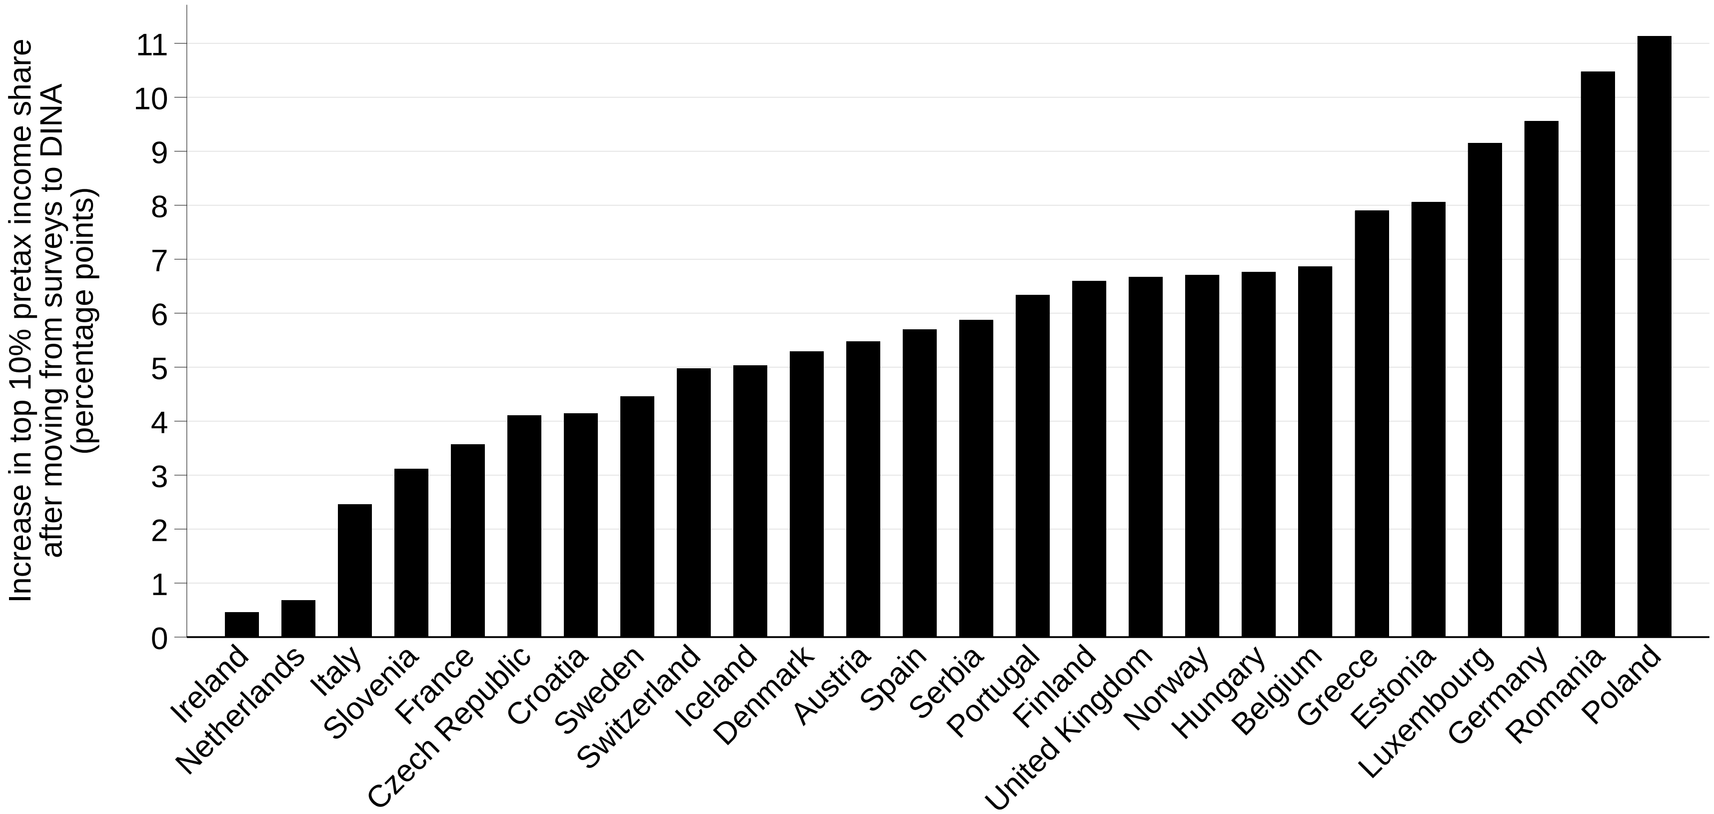

Source: Authors' computations combining surveys, tax data and national accounts. *Note*: The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure D.18: From surveys to DINA: top 1% pretax income share by country, 2017

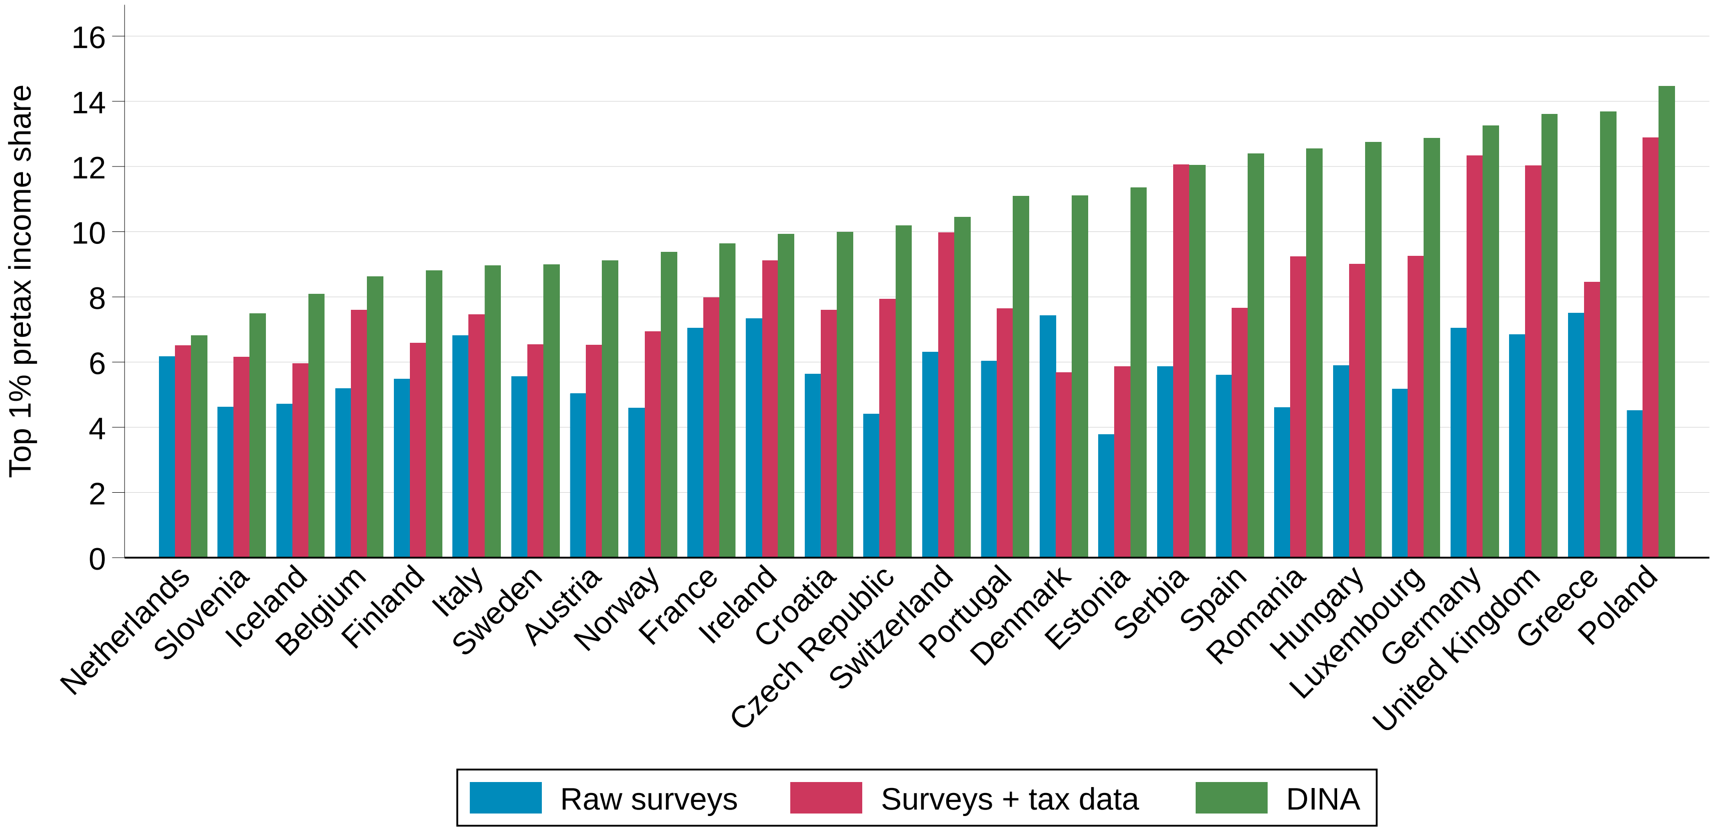

Source: Authors' computations combining surveys, tax data and national accounts. *Note*: The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure D.19: From surveys to DINA: percentage point change in estimated top 1% pretax income share by country, 2017

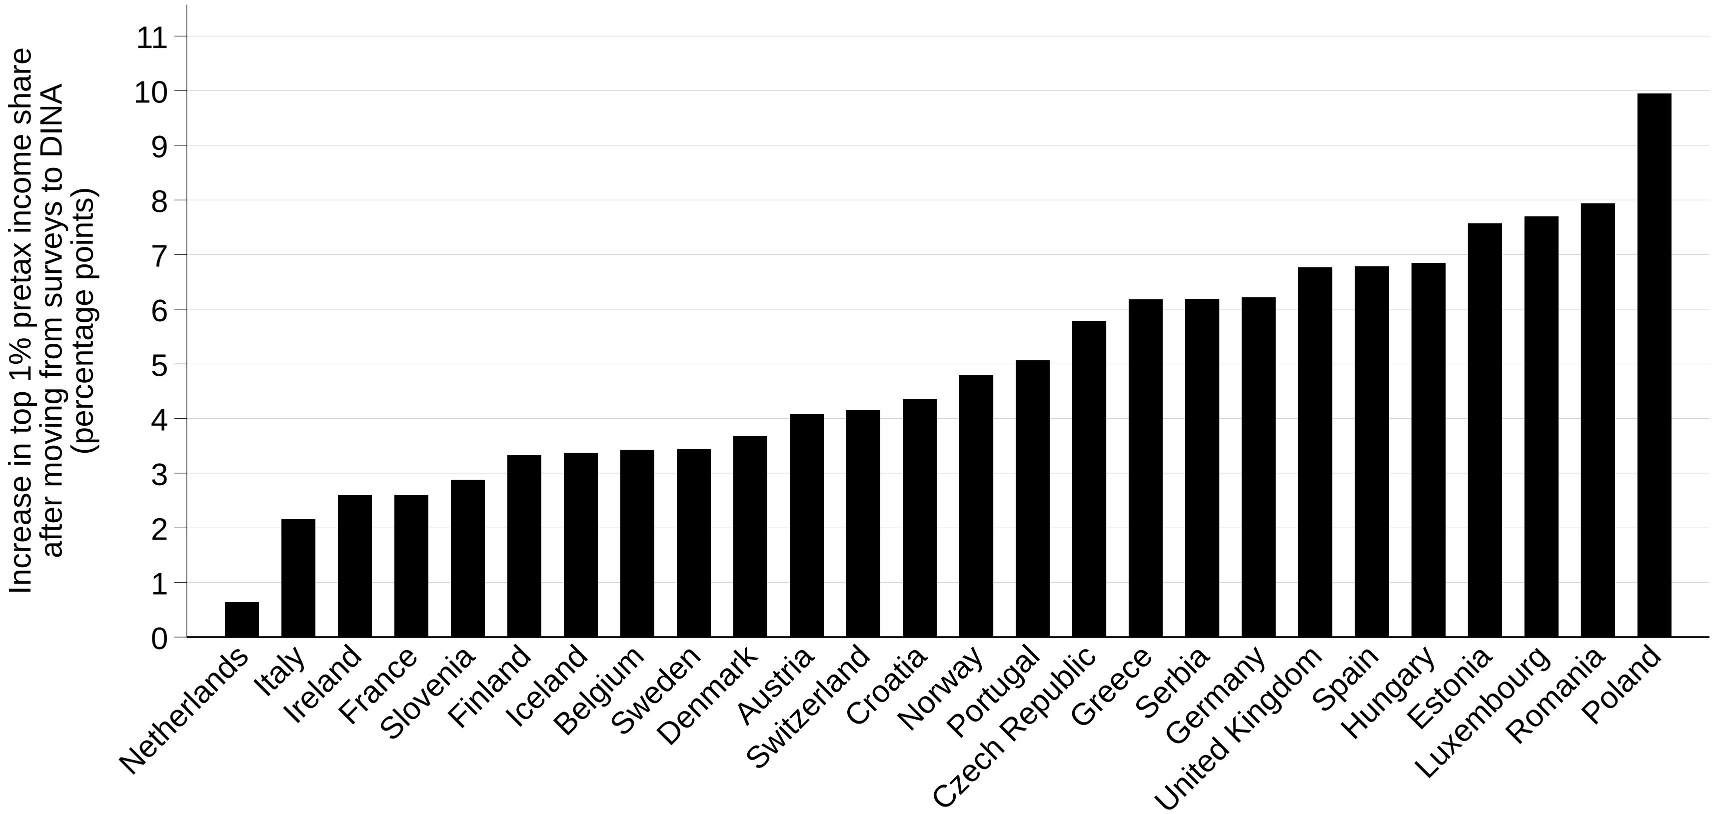

Source: Authors' computations combining surveys, tax data and national accounts. *Note*: The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure D.20: Robustness Check: Exclusion of Countries with Imputed Nonresponse instead of Tax Data (pretax income inequality)

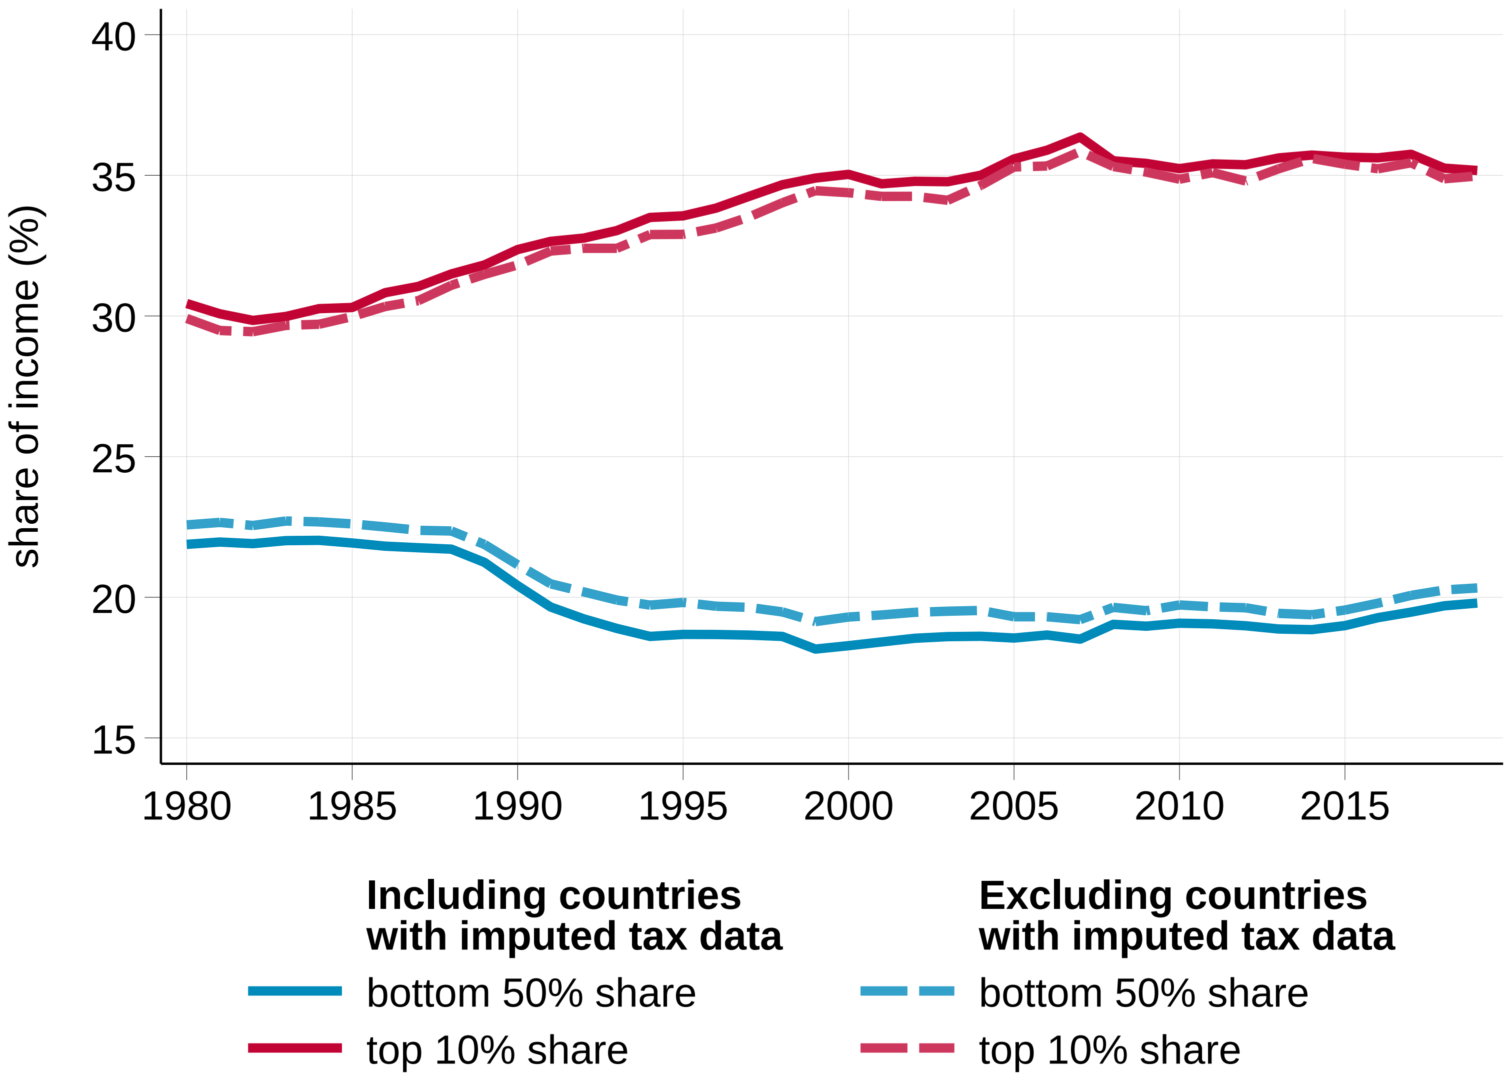

Source: Authors' computations combining surveys, tax data and national accounts. *Note*: Incomes measured at purchasing power parity. The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure D.21: Robustness Check: Exclusion of Countries with Imputed Nonresponse instead of Tax Data (posttax income inequality)

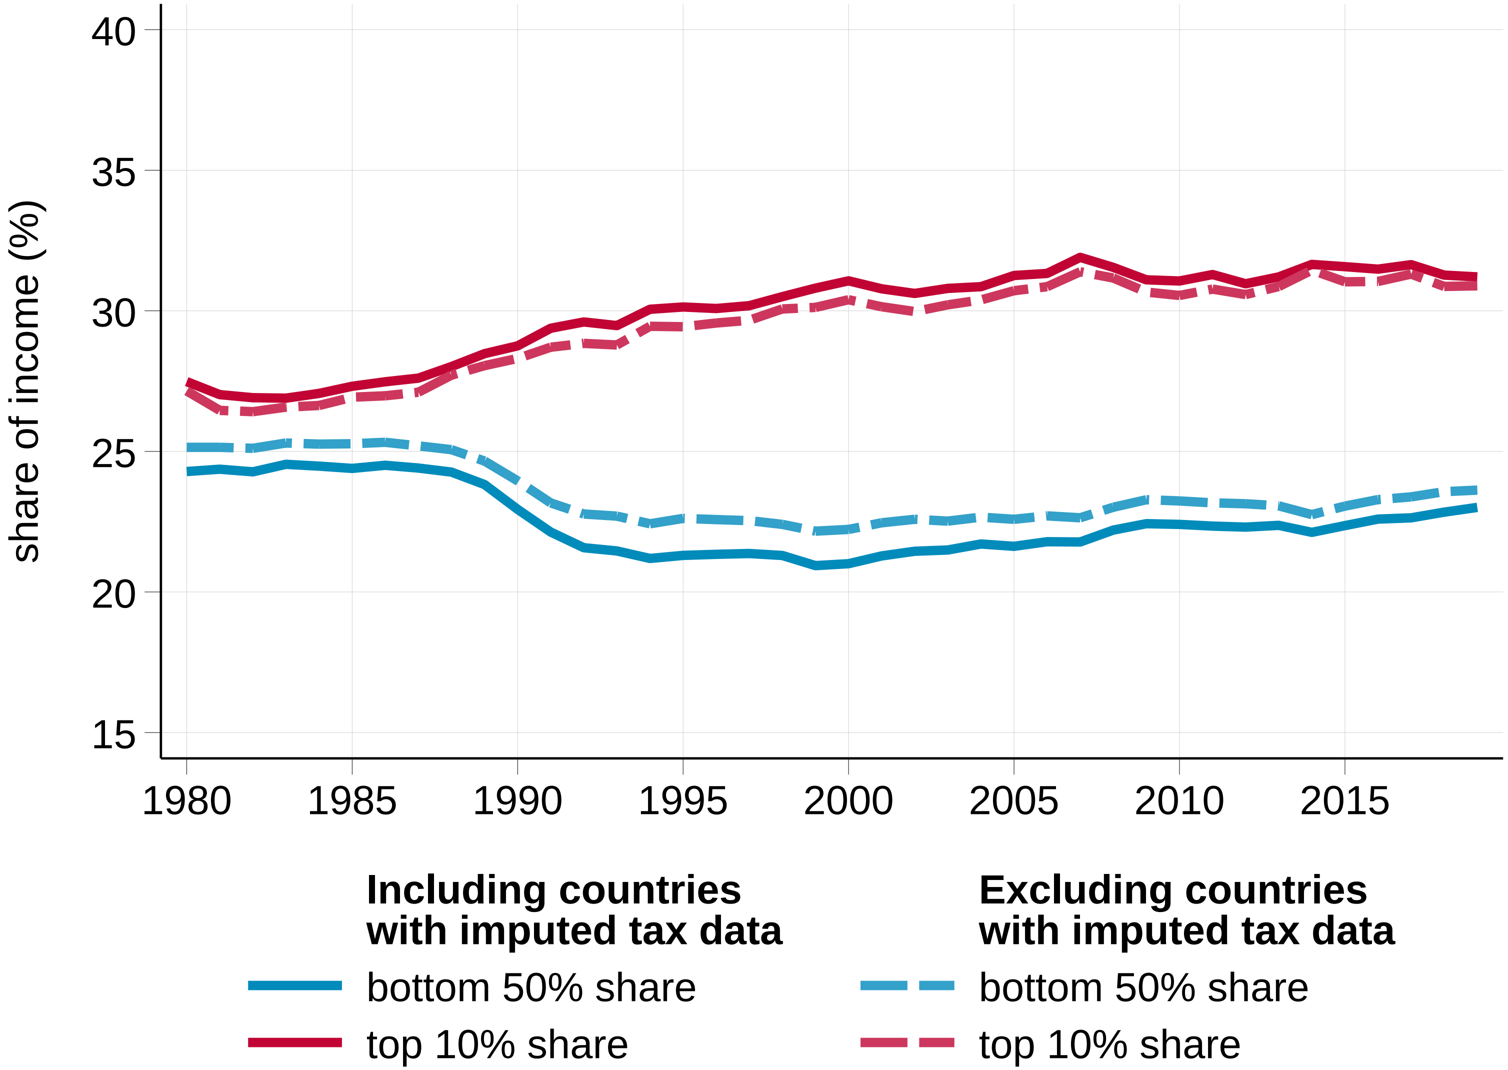

Source: Authors' computations combining surveys, tax data and national accounts. *Note*: Incomes measured at purchasing power parity. The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure D.22: Pretax income shares in Europe: distribution of taxes on products

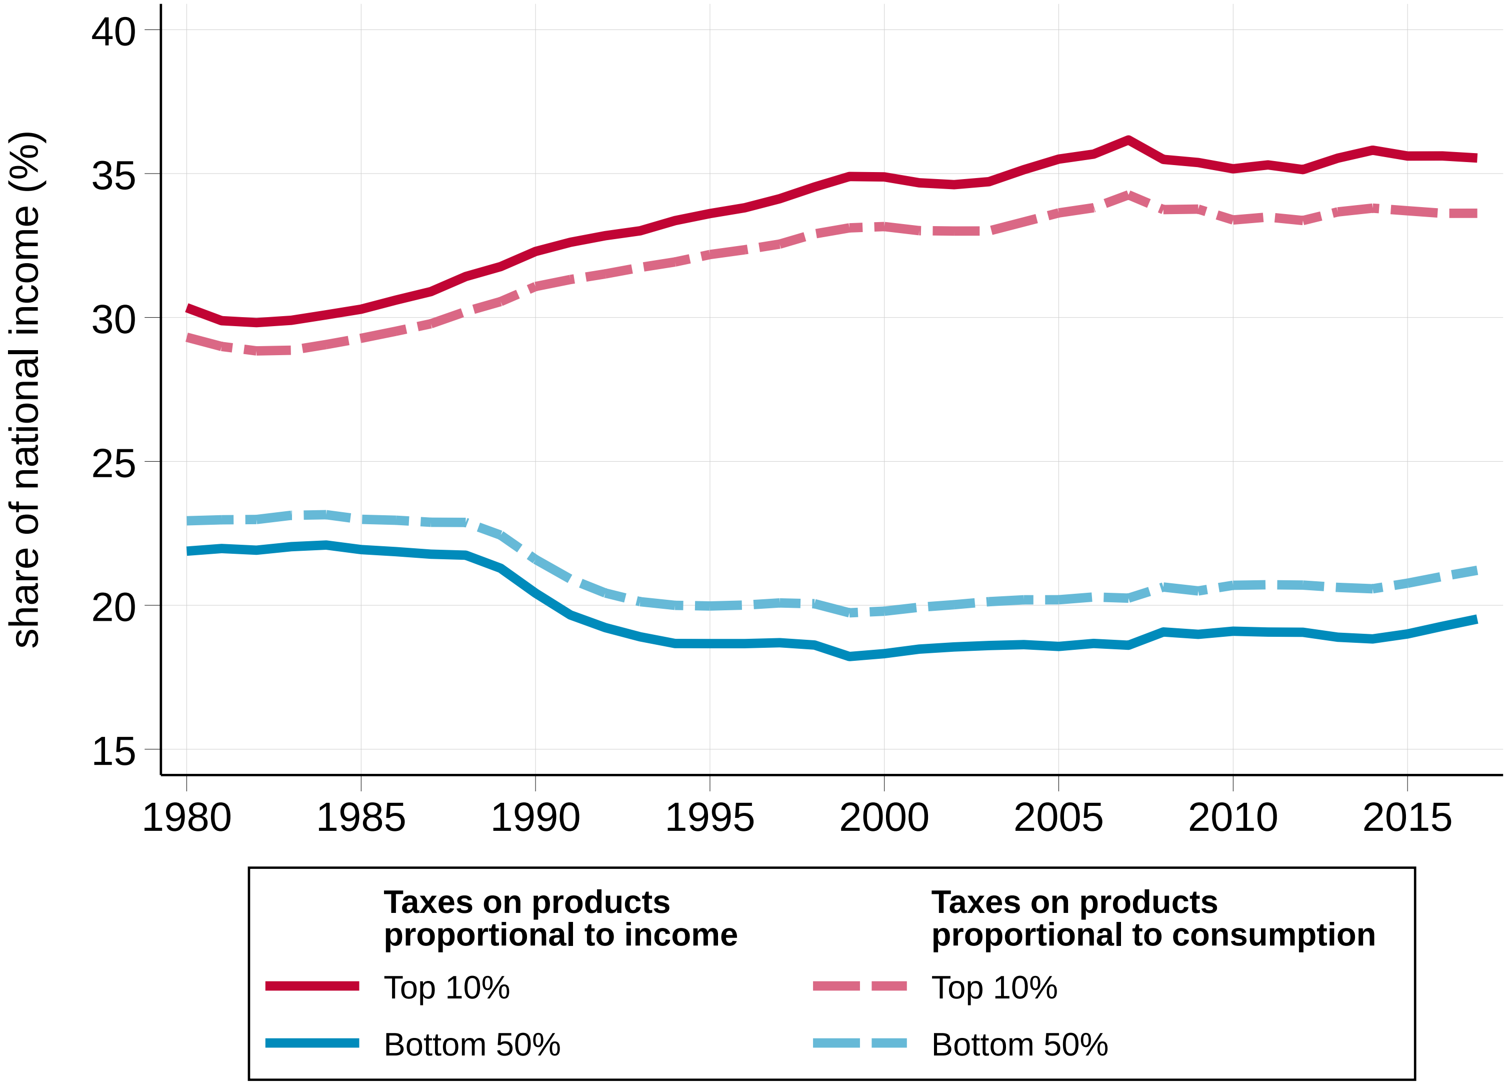

Notes. The figure compares the top 10% and bottom 50% European income shares in two scenarios: one in which taxes on products are distributed proportionally to income, and one in which they are distributed proportionally to consumption.

Figure D.23: Pretax income shares in Europe: broad equal-split vs. narrow equal-split

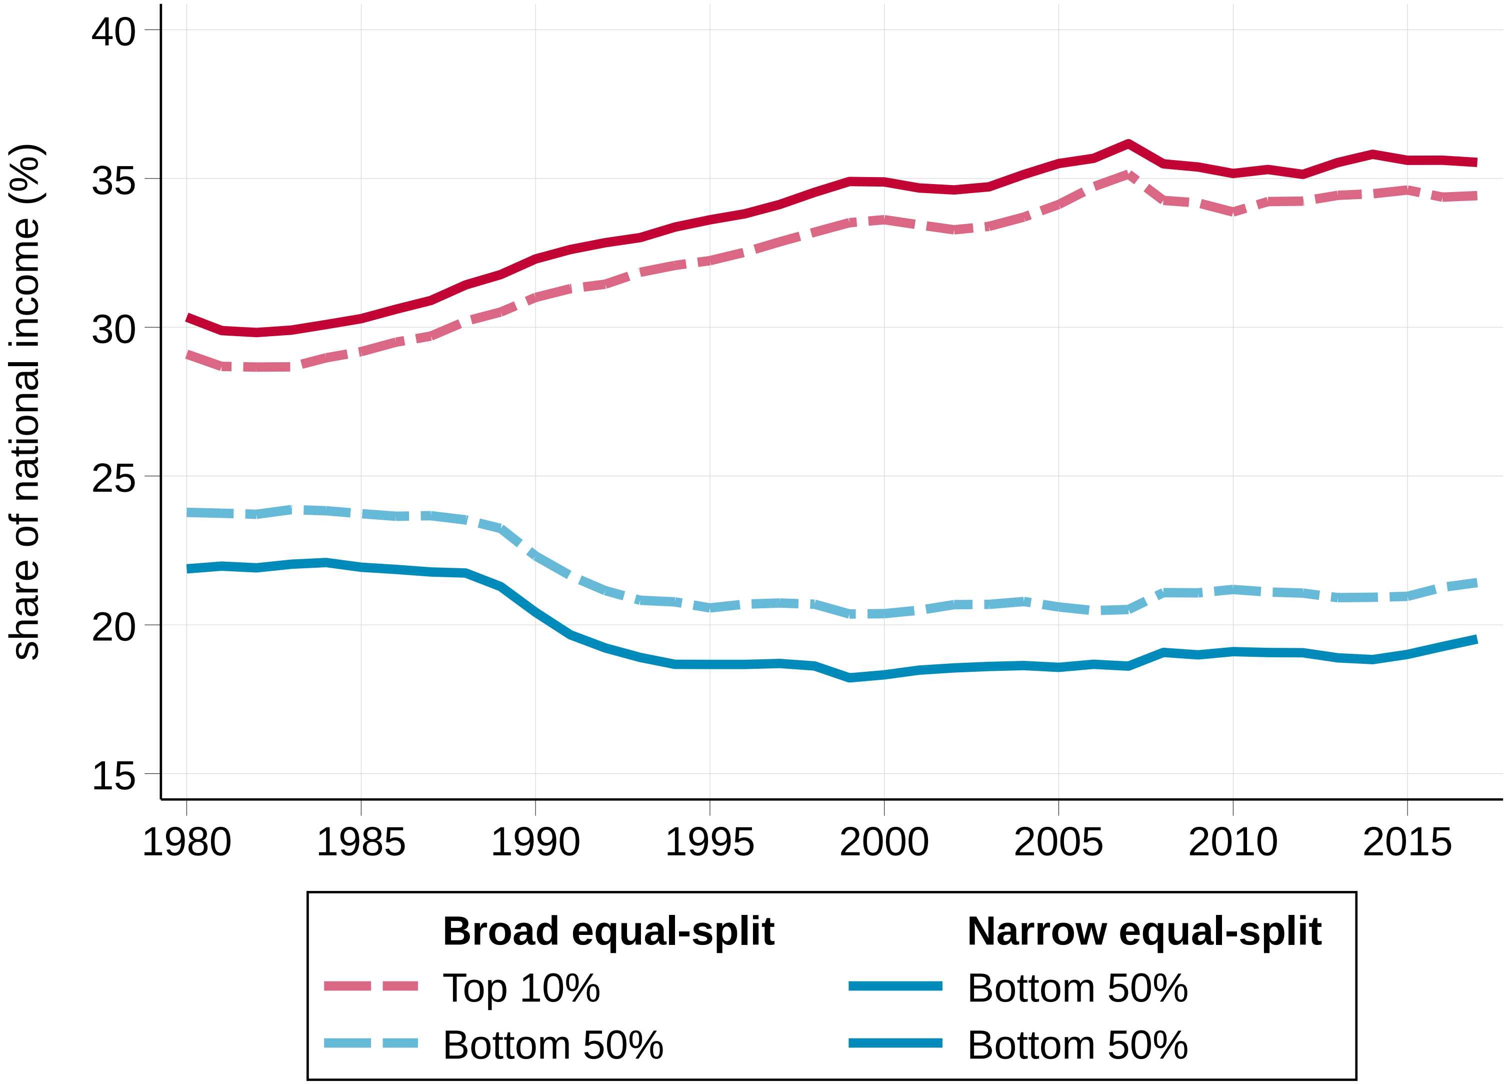

Notes. The figure compares the top 10% and bottom 50% European income shares in two scenarios: one in which income is split equally among all members of the household (broad equal-split), and one in which income is split equally among spouses (narrow equal-split).

Figure D.24: Top 1% income share in Europe and the United States: comparison of estimates

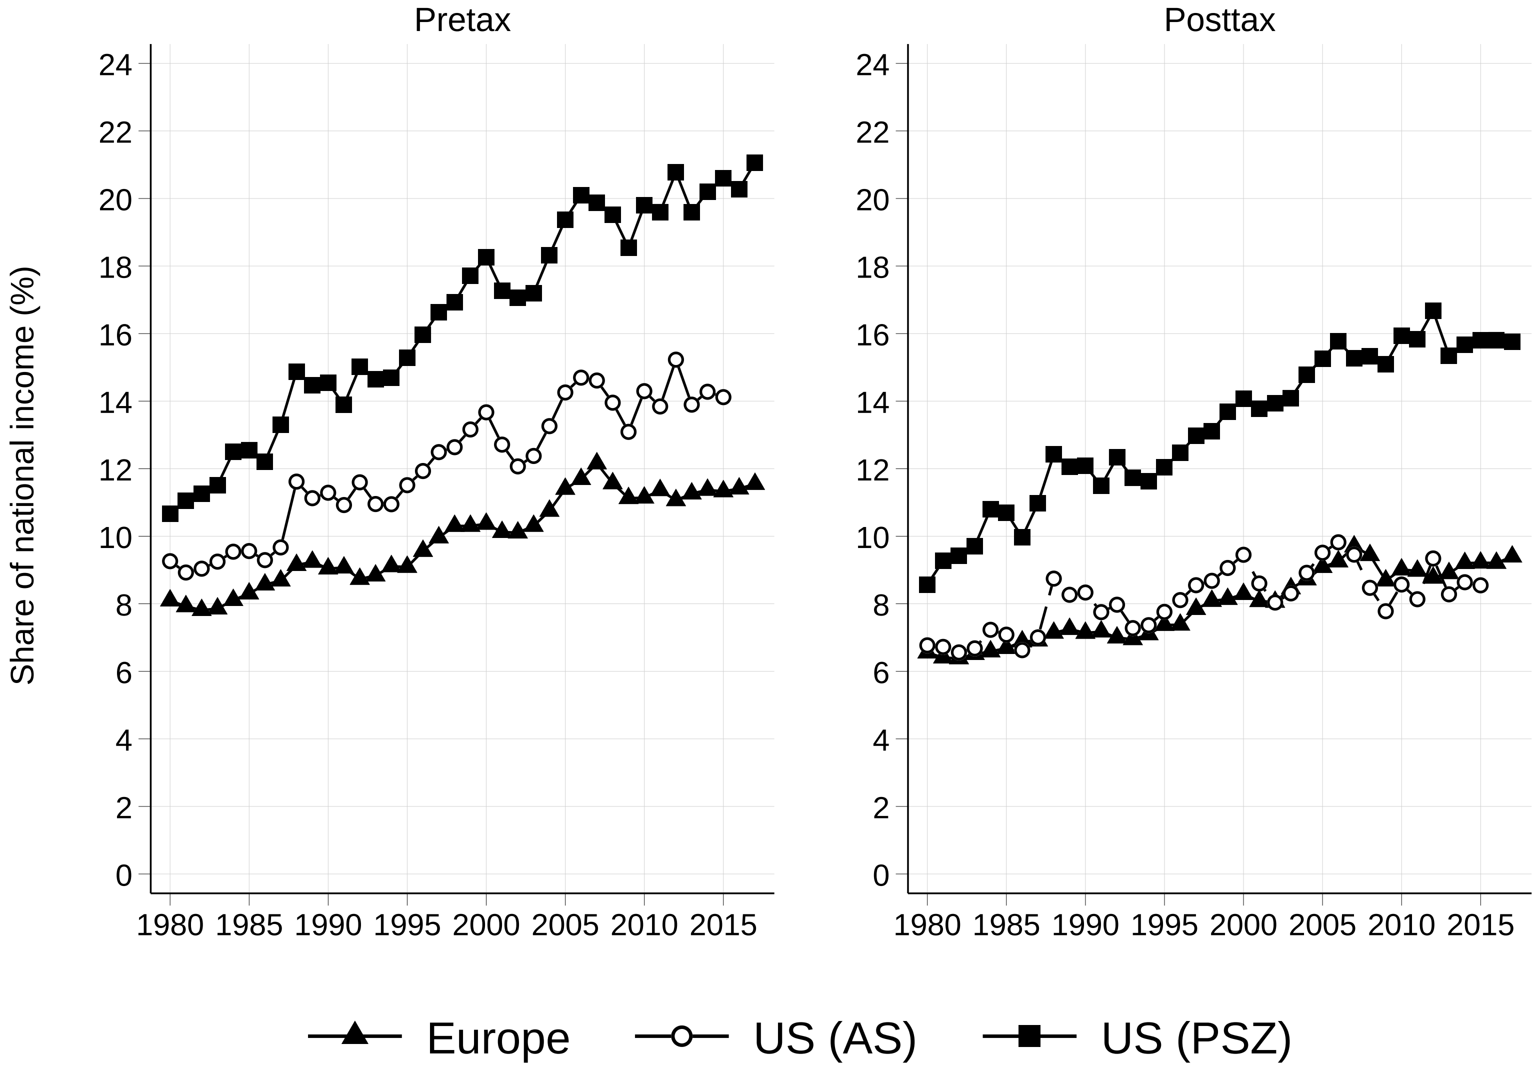

Source: Authors' computations combining surveys, tax data and national accounts for European countries and Piketty, Saez, and Zucman, 2018 (US-PSZ) as well as Auten and Splinter, 2019 (US-AS) for the US.

D.2.2 Distribution of pretax income

Figure D.25: Average annual pretax income growth by percentile in Europe and the United States, 1980-2017

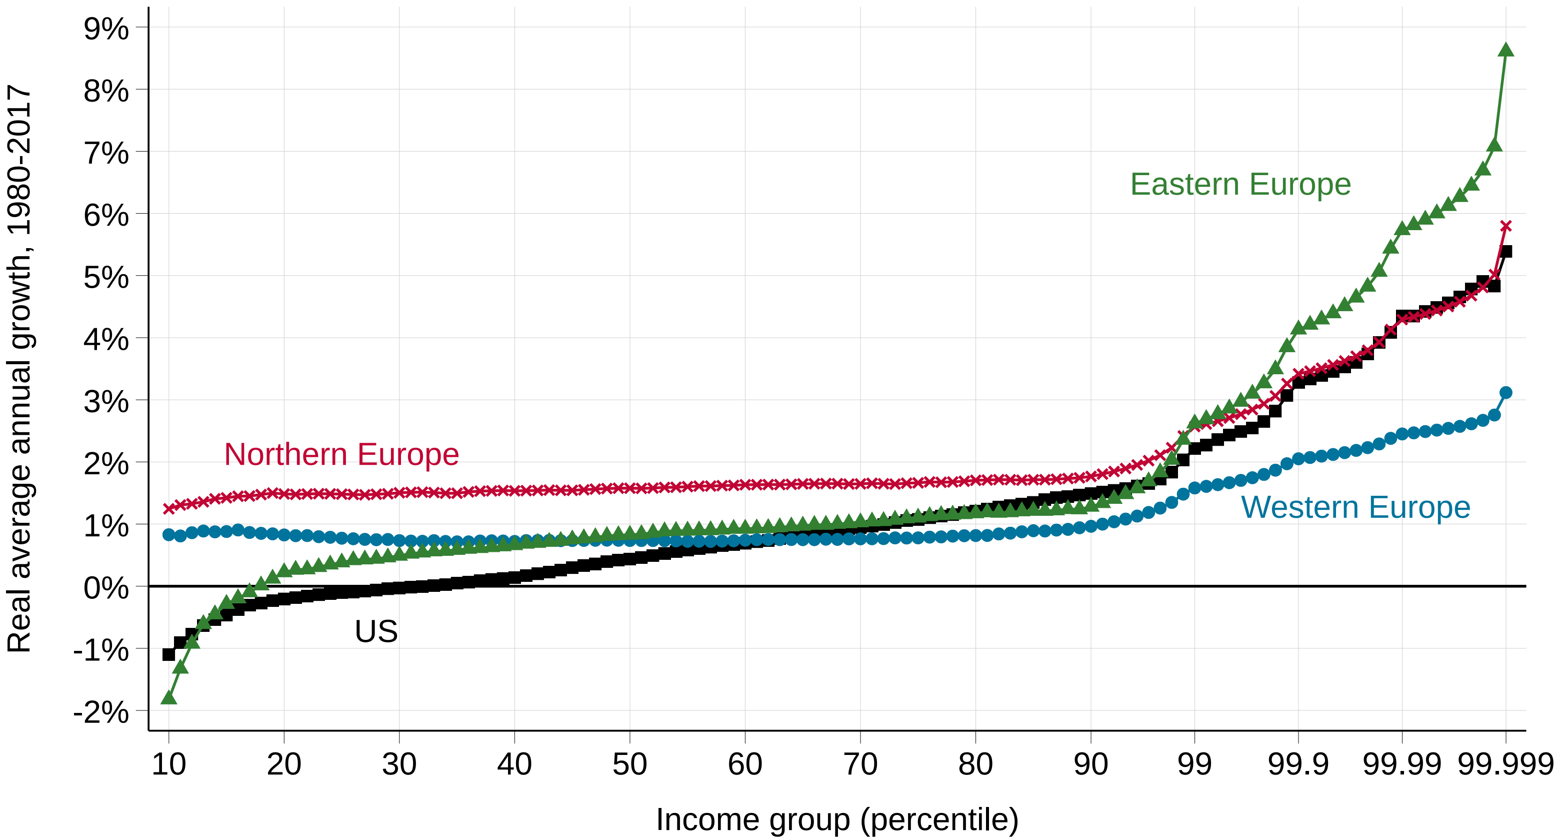

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The figure shows the average annual growth rate of pretax national income by percentile in Western Europe, Northern Europe, Eastern Europe, and the United States, with a further decomposition of the top percentile, between 1980 and 2017. The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure D.26: Cumulated growth by pretax income group: Western Europe

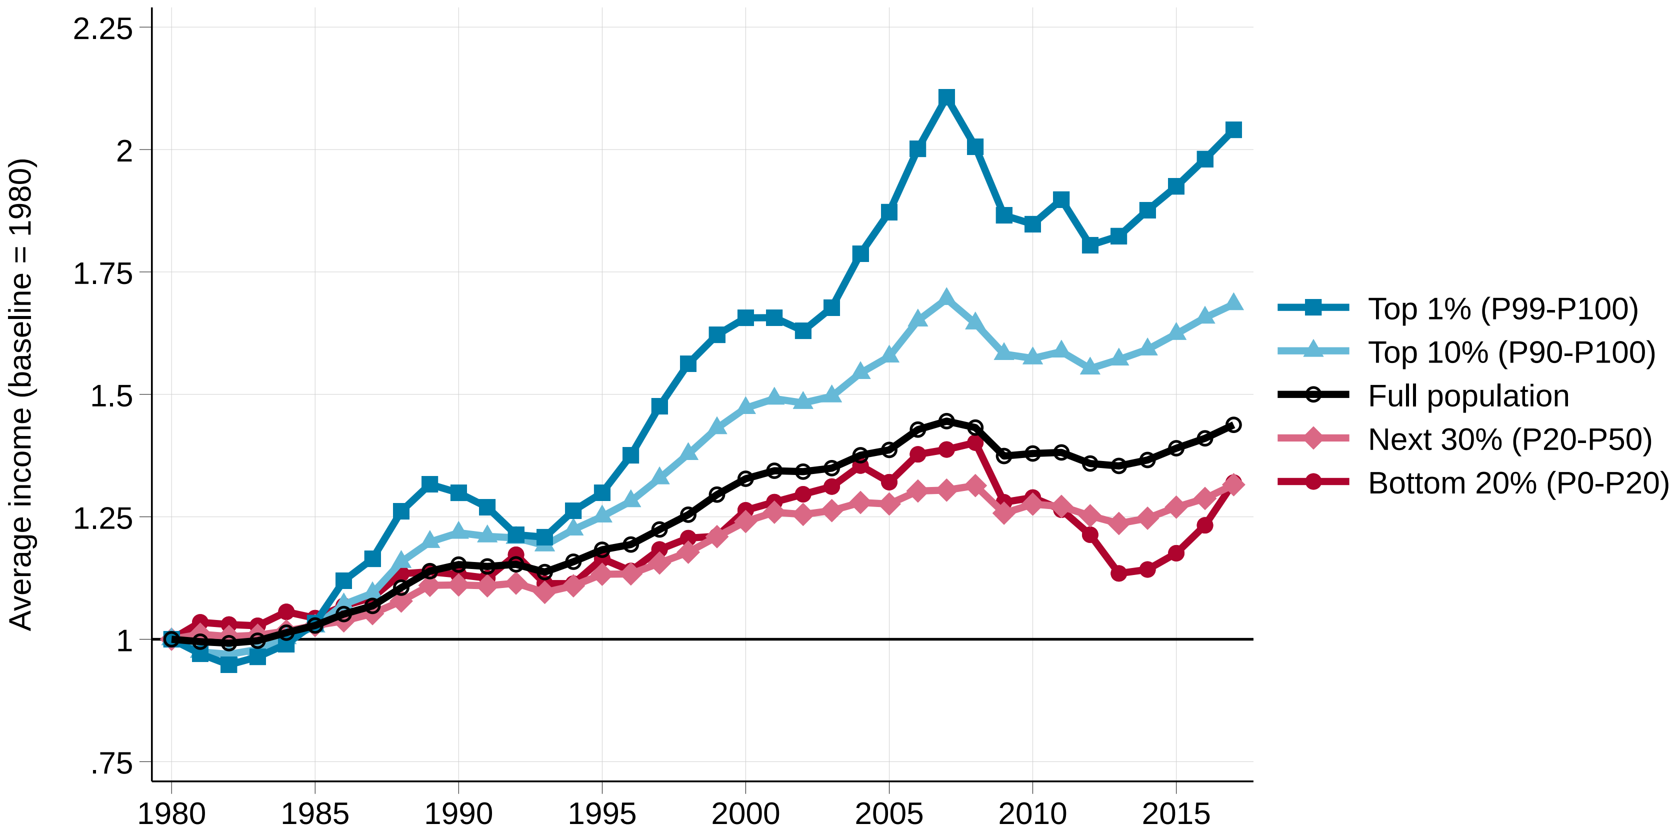

Source: Authors' computations combining surveys, tax data and national accounts for Europe. Piketty, Saez, and Zucman, 2018 for the US. *Notes:* This figure shows the evolution of the average pretax income of the top 1% (p99p100), the top 10% (p90p100), the bottom 20% (p0p20), the next 30% (p20p50) and the average regional income relative to 1980. The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. Incomes measured at purchasing power parity.

Figure D.27: Cumulated growth by pretax income group: Northern Europe

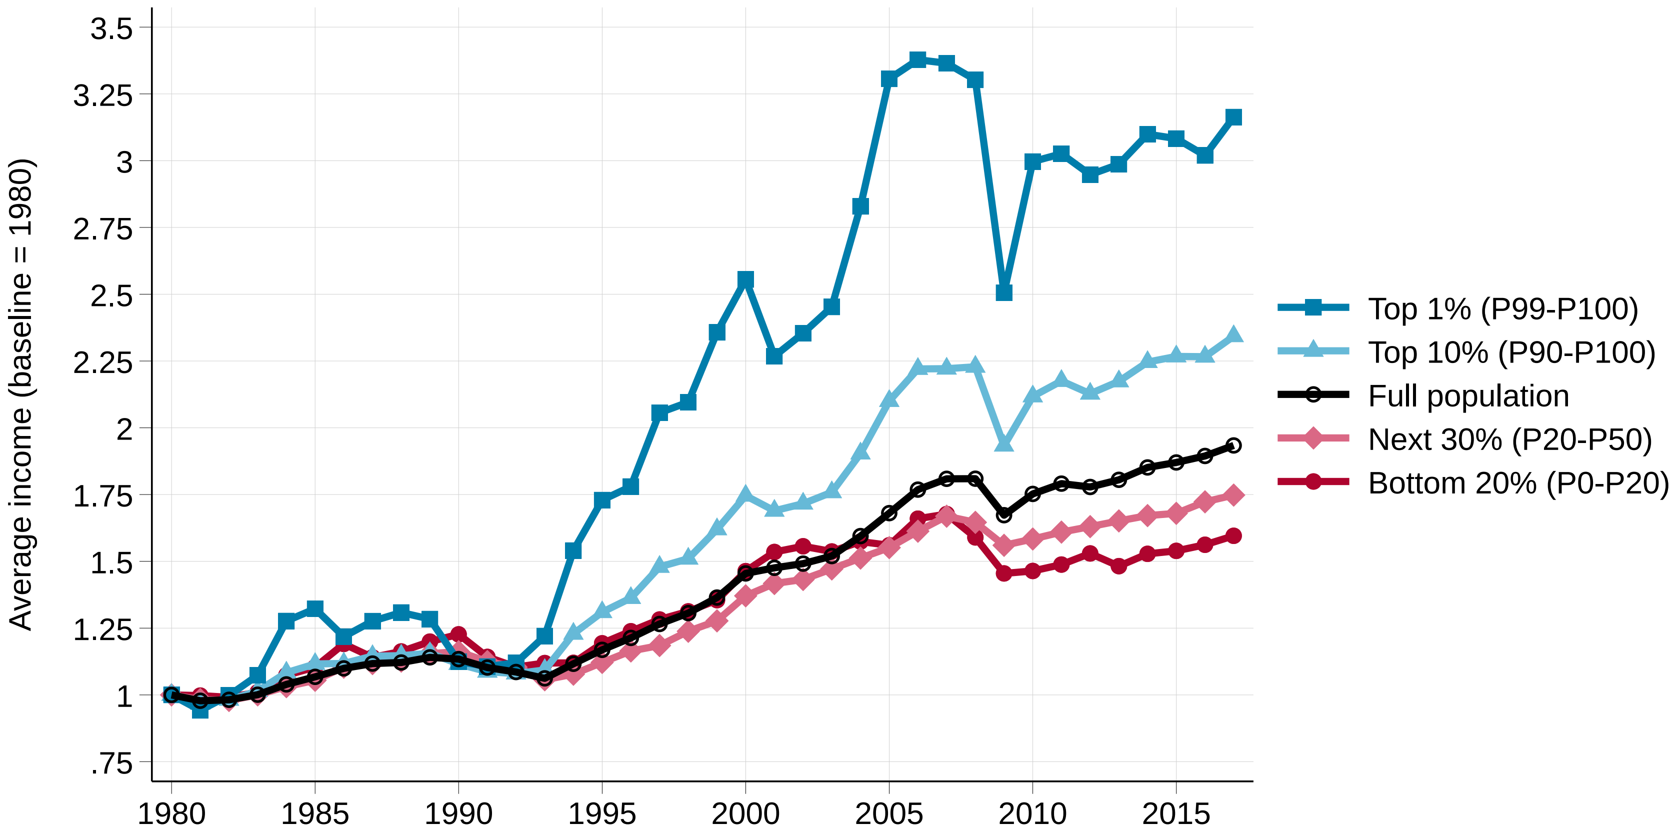

Source: Authors' computations combining surveys, tax data and national accounts for Europe. Piketty, Saez, and Zucman, 2018 for the US. *Notes:* This figure shows the evolution of the average pretax income of the top 1% (p99p100), the top 10% (p90p100), the bottom 20% (p0p20), the next 30% (p20p50) and the average regional income relative to 1980. The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. Incomes measured at purchasing power parity.

Figure D.28: Cumulated growth by pretax income group: Eastern Europe

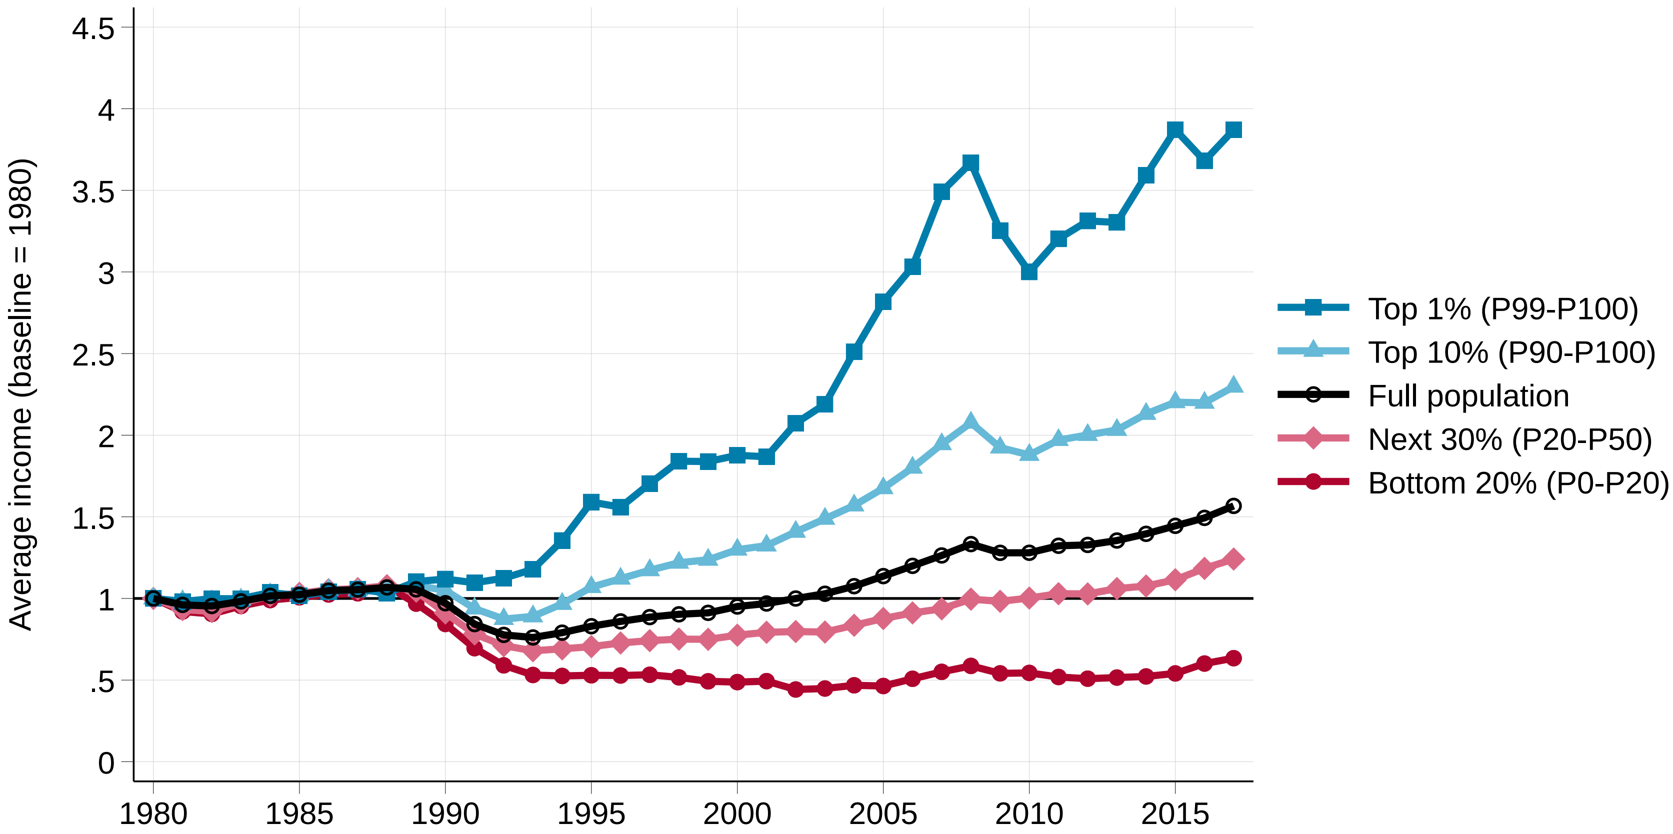

Source: Authors' computations combining surveys, tax data and national accounts for Europe. Piketty, Saez, and Zucman, 2018 for the US. *Notes:* This figure shows the evolution of the average pretax income of the top 1% (p99p100), the top 10% (p90p100), the bottom 20% (p0p20), the next 30% (p20p50) and the average regional income relative to 1980. The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. Incomes measured at purchasing power parity.

Figure D.29: Cumulated growth by pretax income group: United States

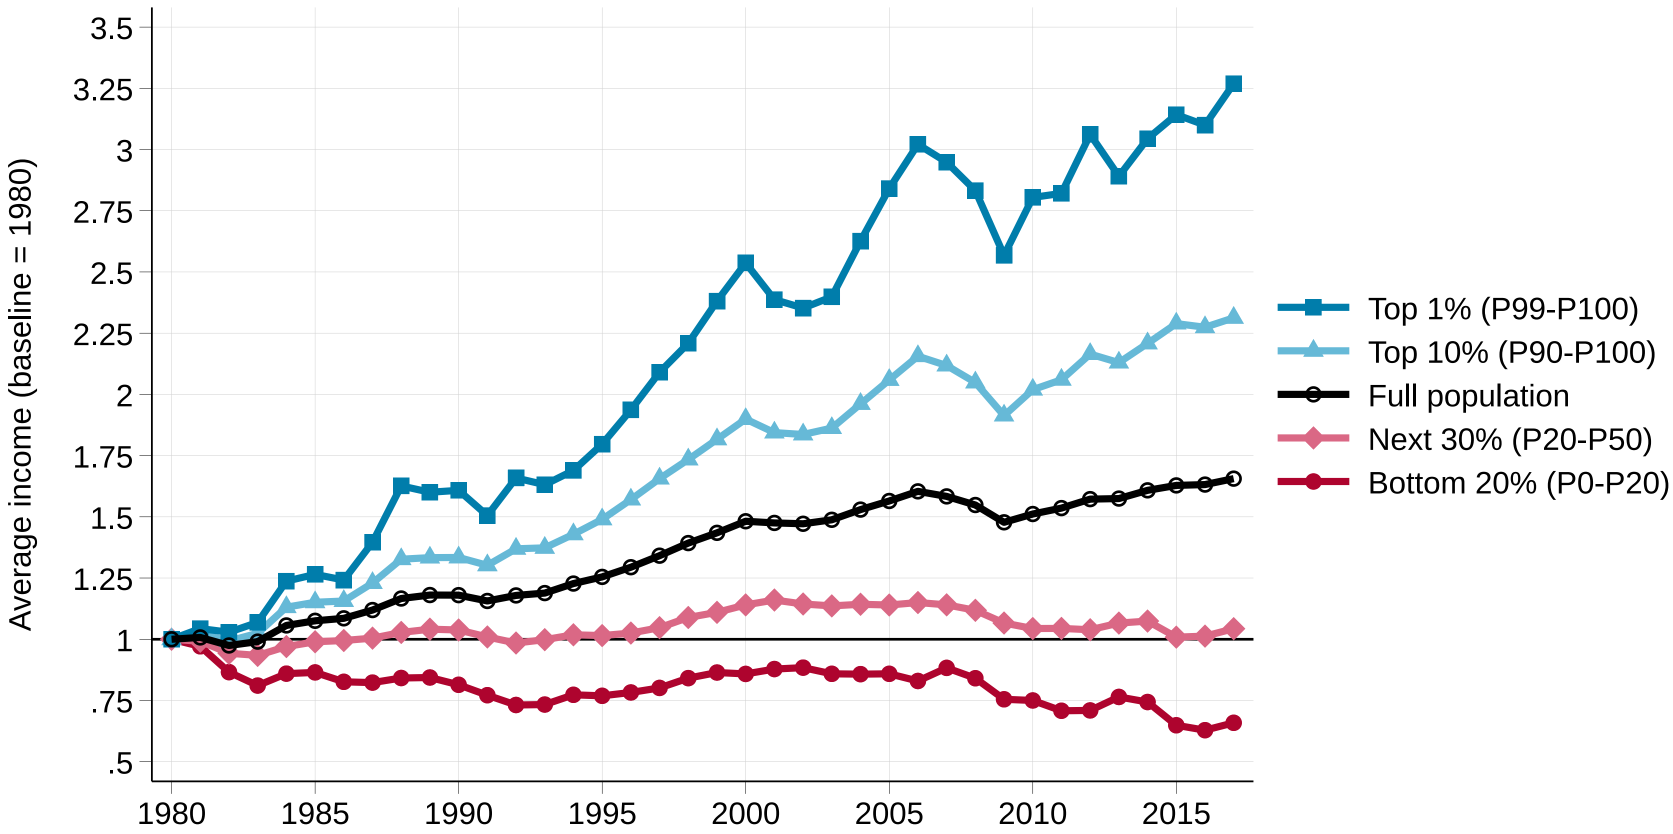

Source: Authors' computations combining surveys, tax data and national accounts for Europe. Piketty, Saez, and Zucman, 2018 for the US. *Notes:* This figure shows the evolution of the average pretax income of the top 1% (p99p100), the top 10% (p90p100), the bottom 20% (p0p20), the next 30% (p20p50) and the average regional income relative to 1980. The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. Incomes measured at purchasing power parity.

Figure D.30: Top 10% pretax income share in Europe: Geographical decomposition

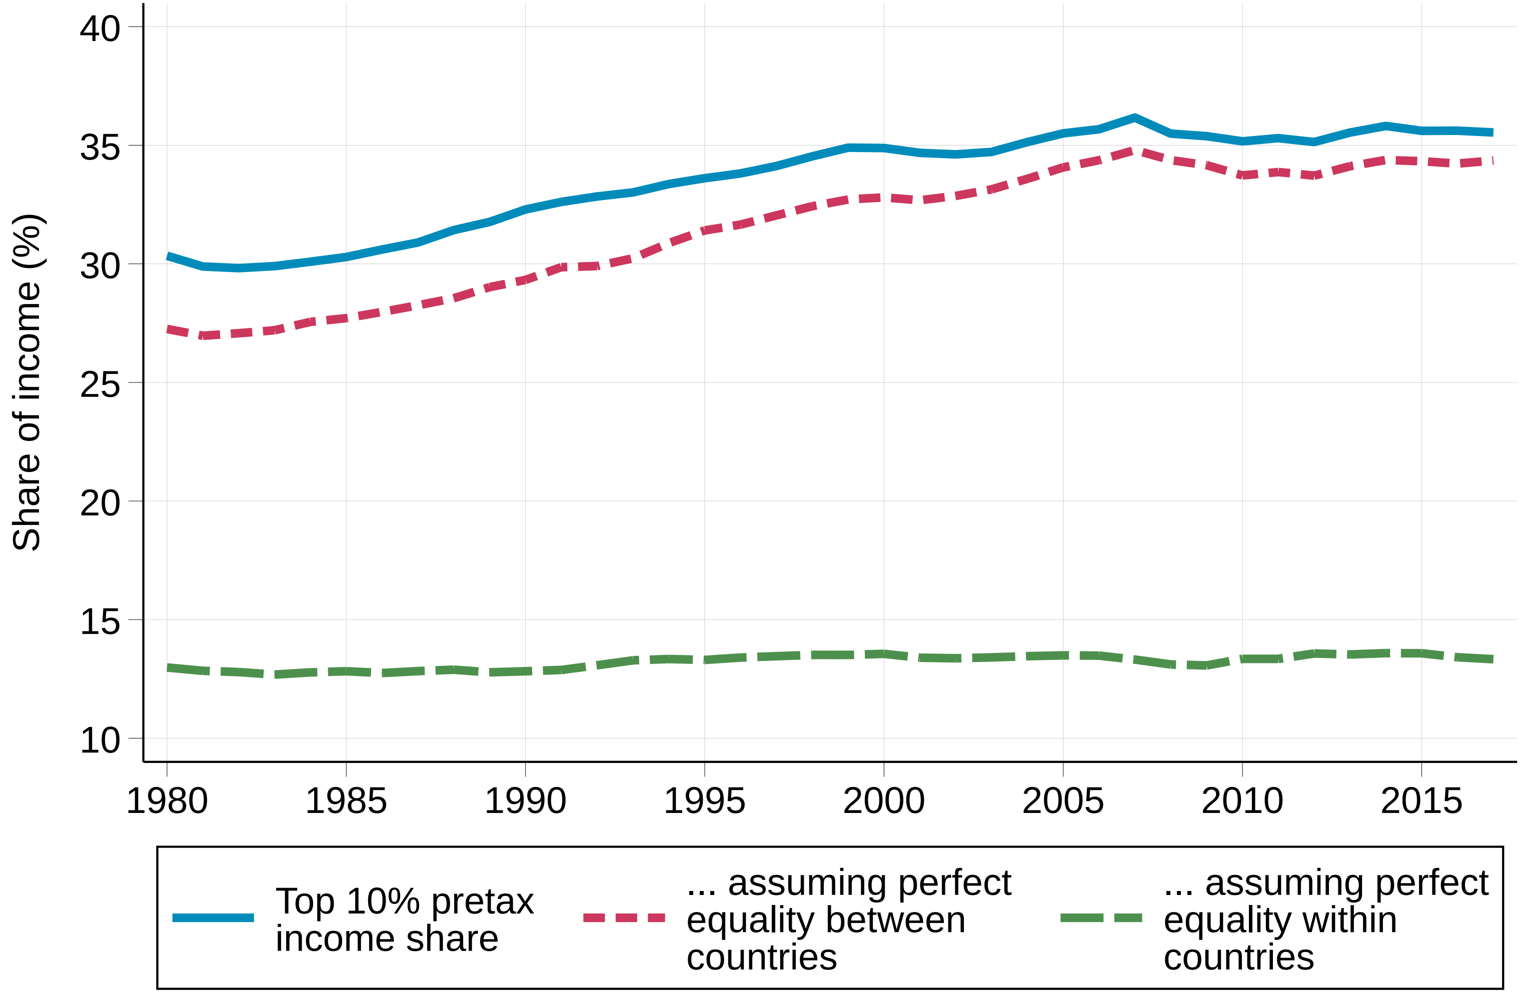

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* Incomes are measured at Purchasing Power Parity in real 2017 Euros. PPP Euro 1 = PPP dollar 1.3. The unit of observation is the adult individual aged 20 or above. See Table D.6 for the composition of European regions.

Figure D.31: Bottom 50% pretax income share in Europe: counterfactual decomposition

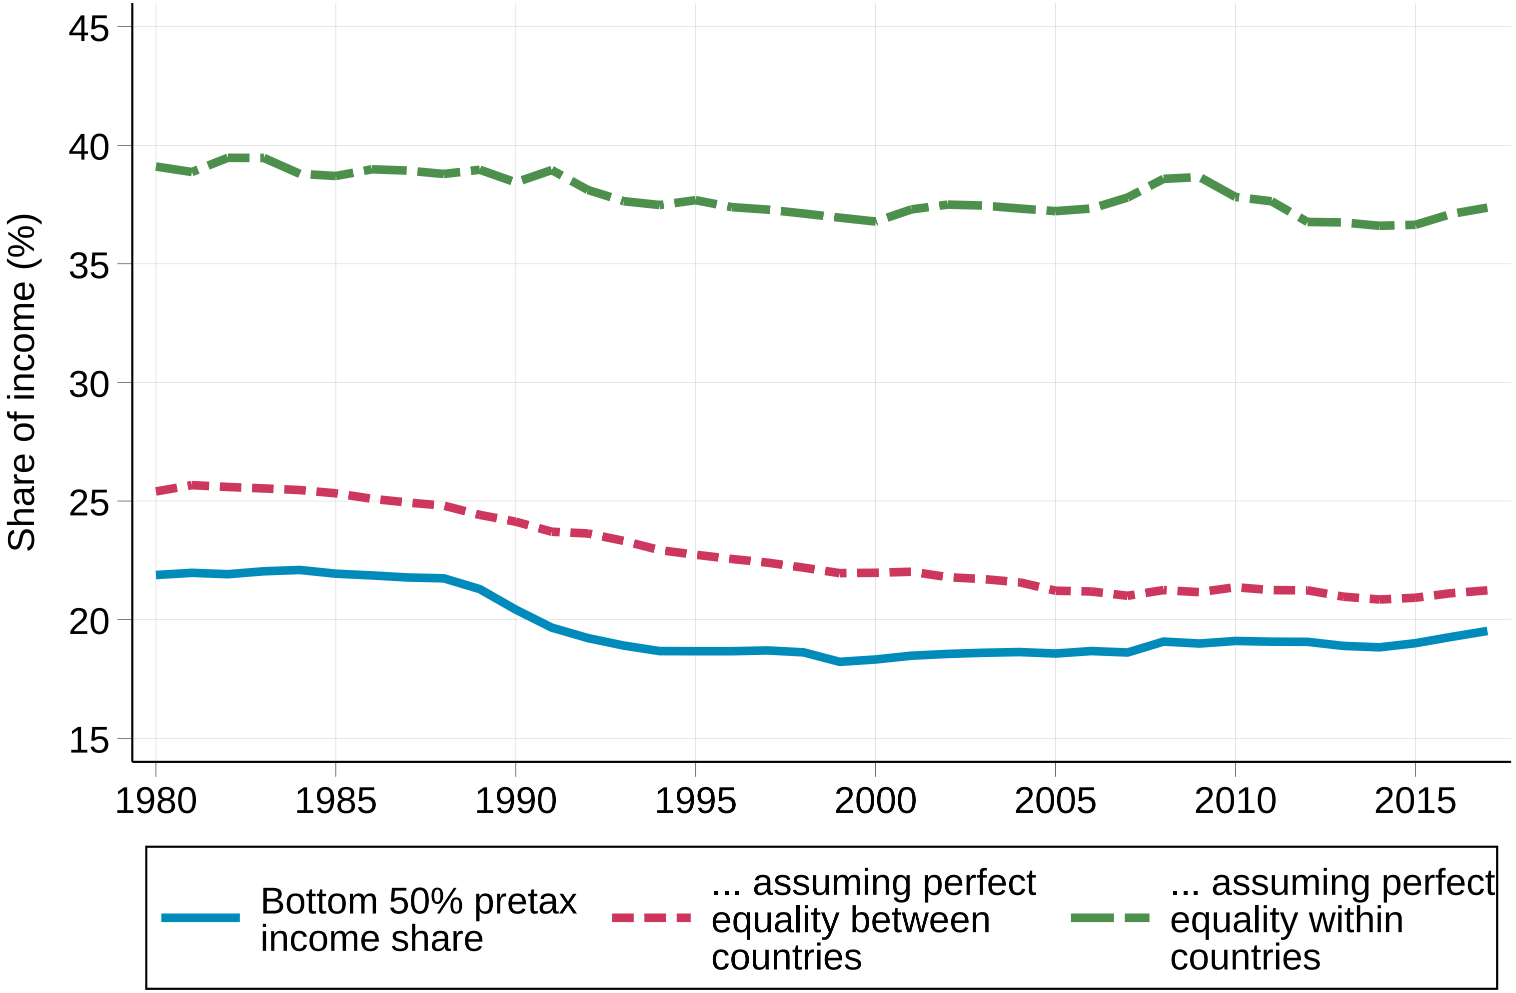

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* Incomes are measured at Purchasing Power Parity in real 2017 Euros. PPP Euro 1 = PPP dollar 1.3. The unit of observation is the adult individual aged 20 or above. See Table D.6 for the composition of European regions.

Figure D.32: Top 10% pretax income share by country: Western Europe

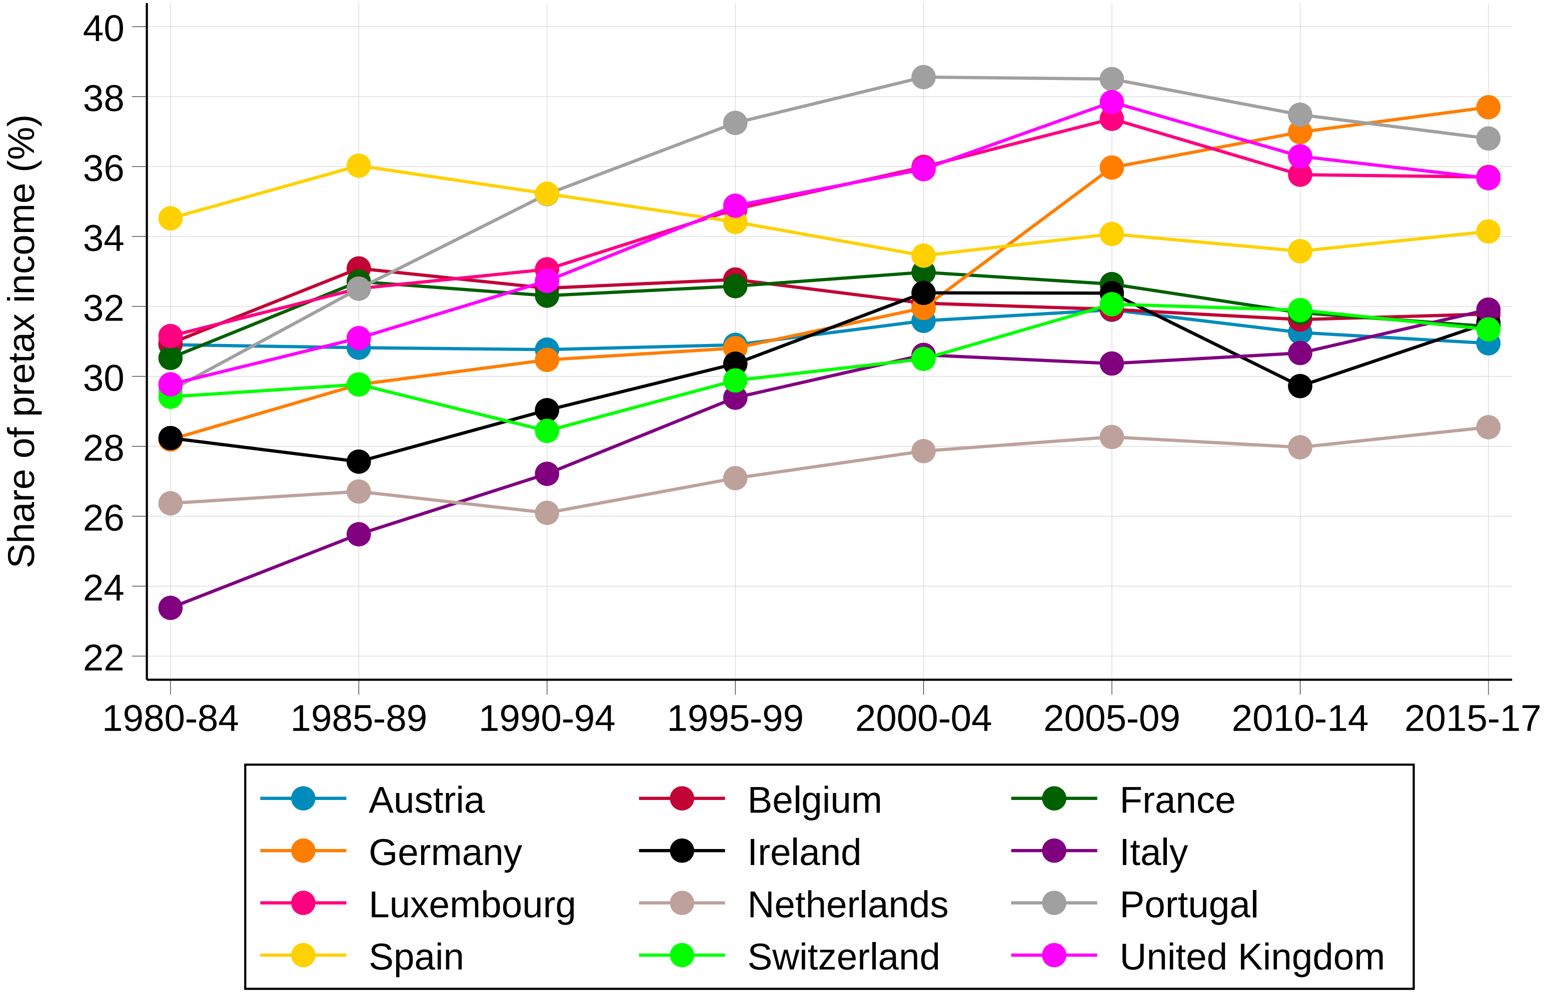

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure D.33: Top 10% pretax income share by country: Northern Europe

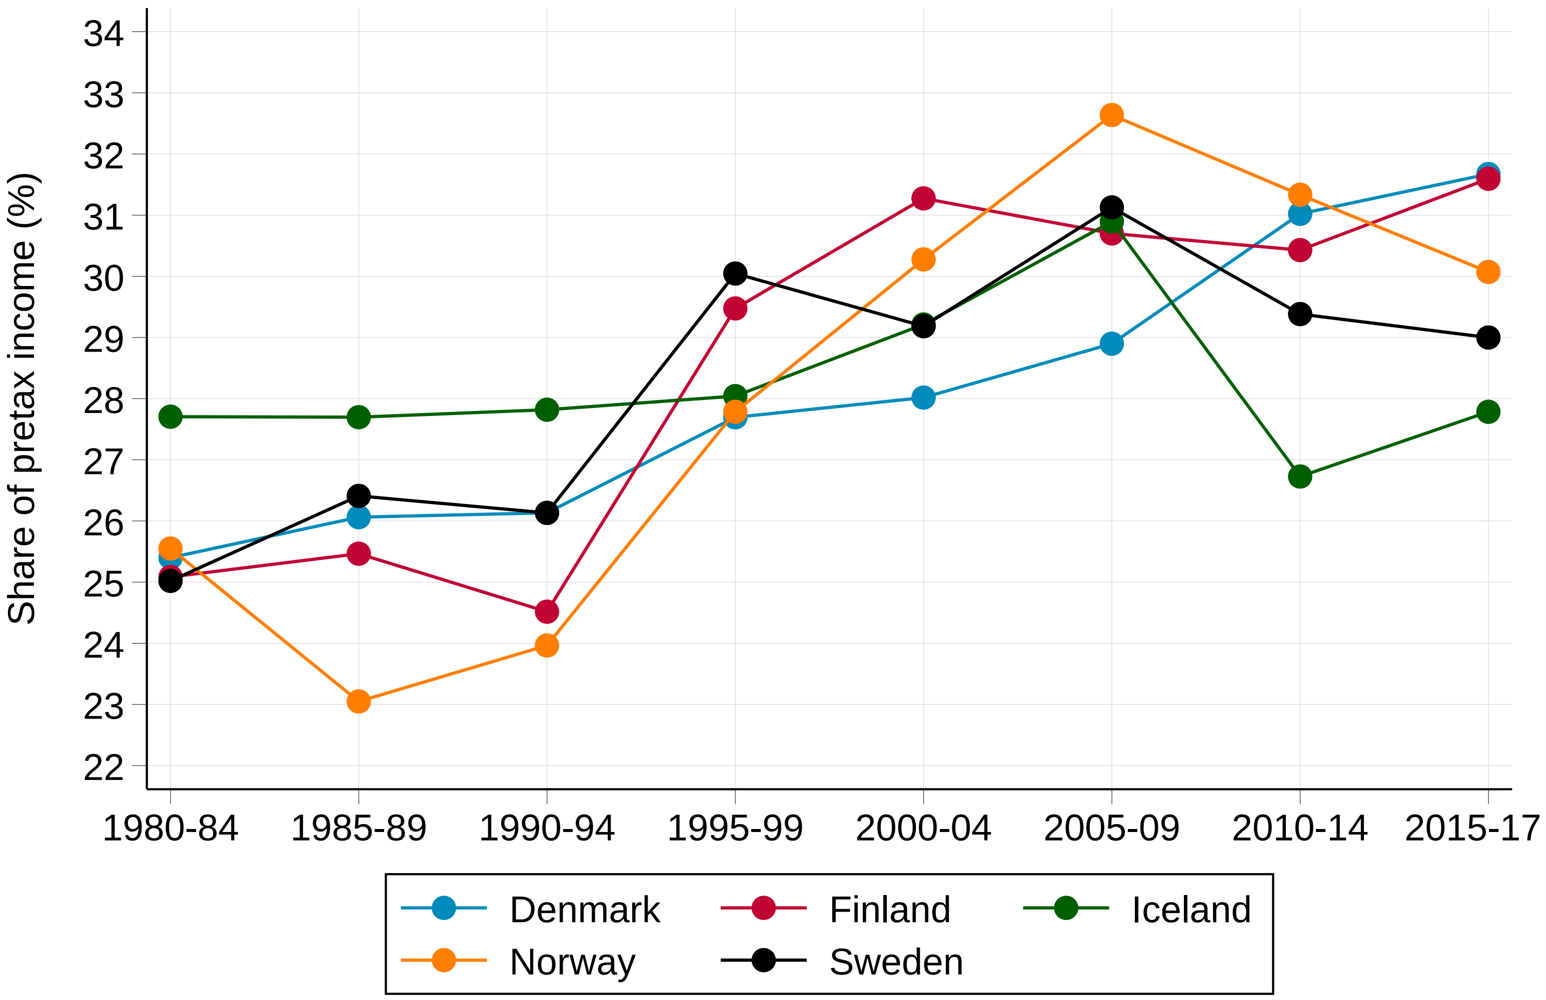

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure D.34: Top 10% pretax income share by country: Eastern Europe

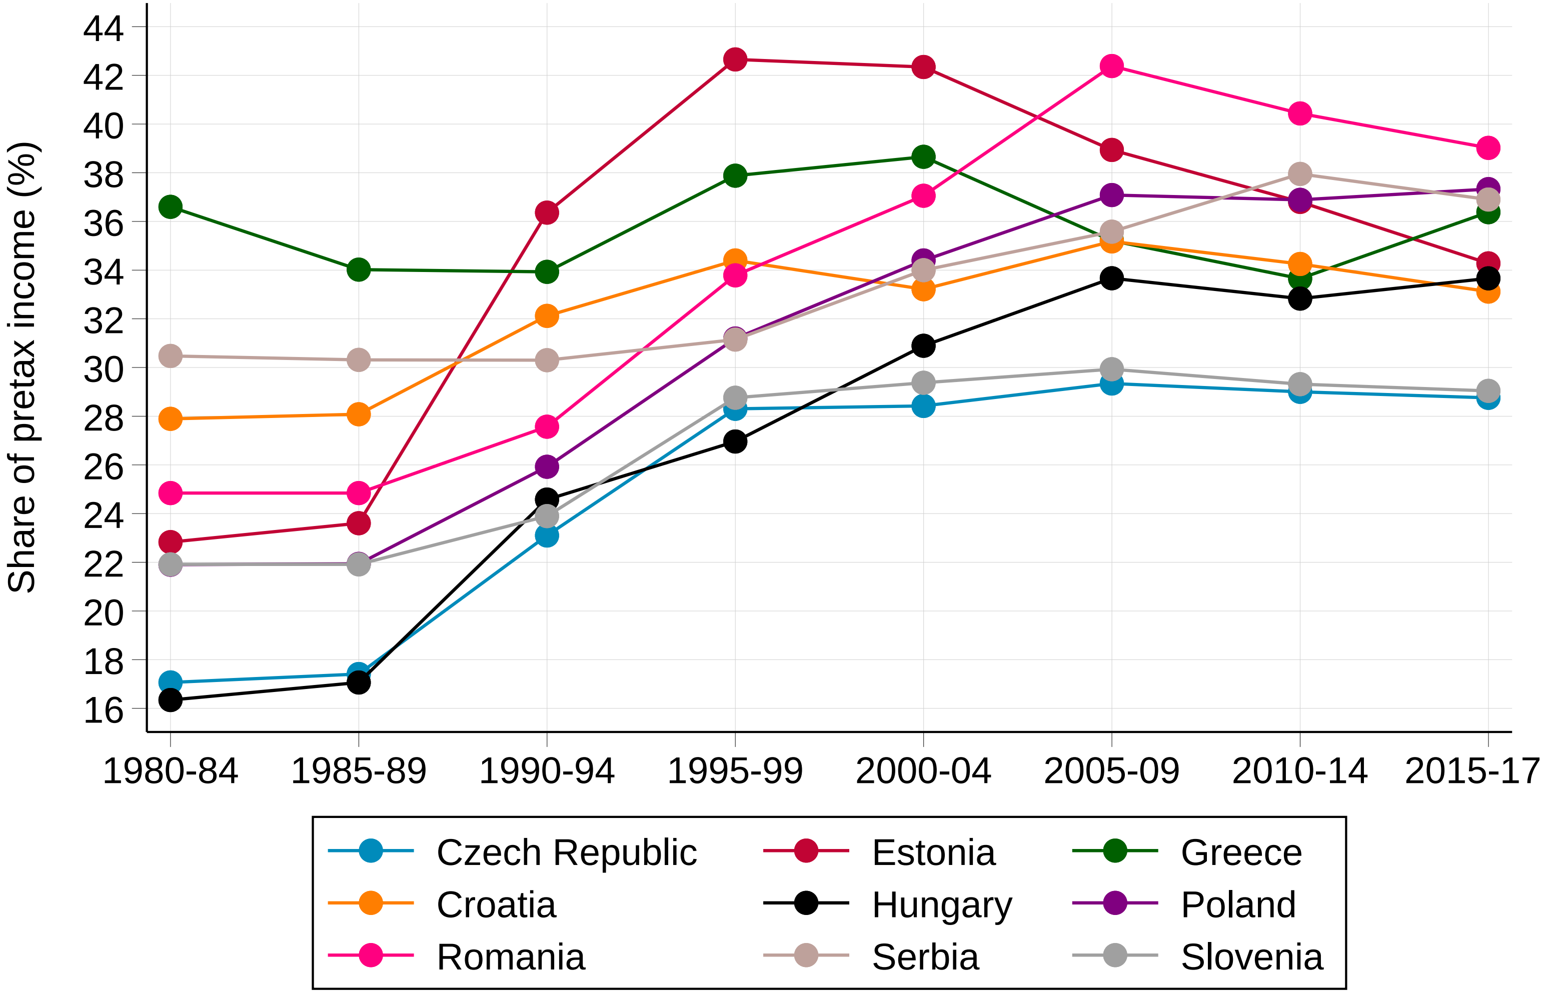

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure D.35: Top 1% pretax income share by country: Western Europe

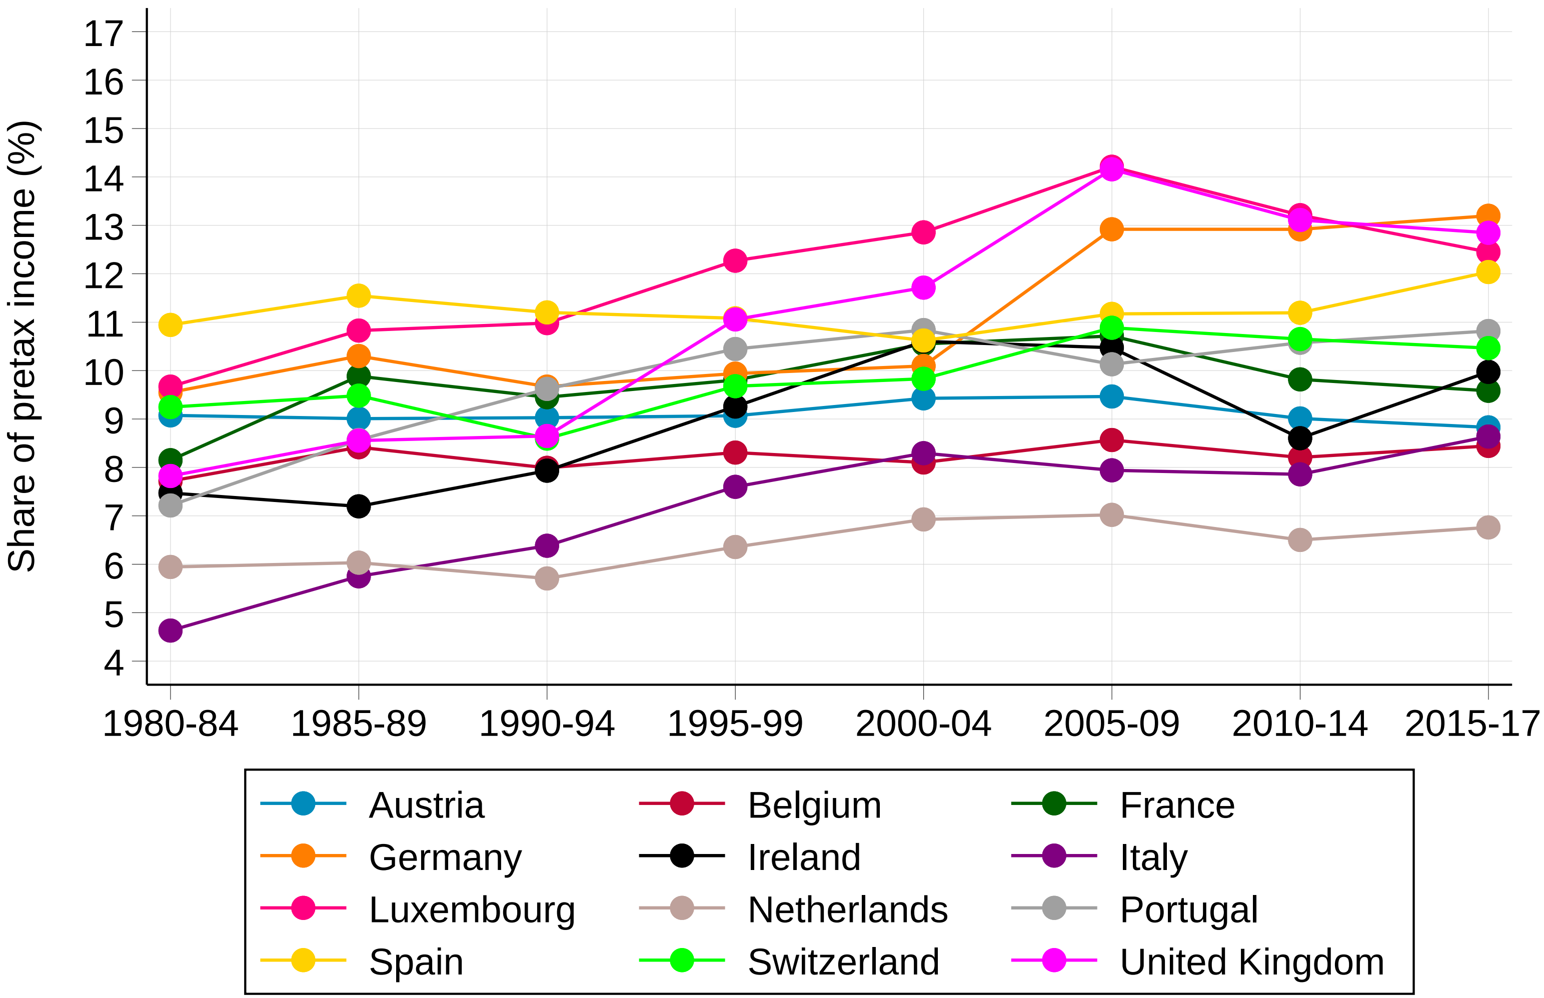

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure D.36: Top 1% pretax income share by country: Northern Europe

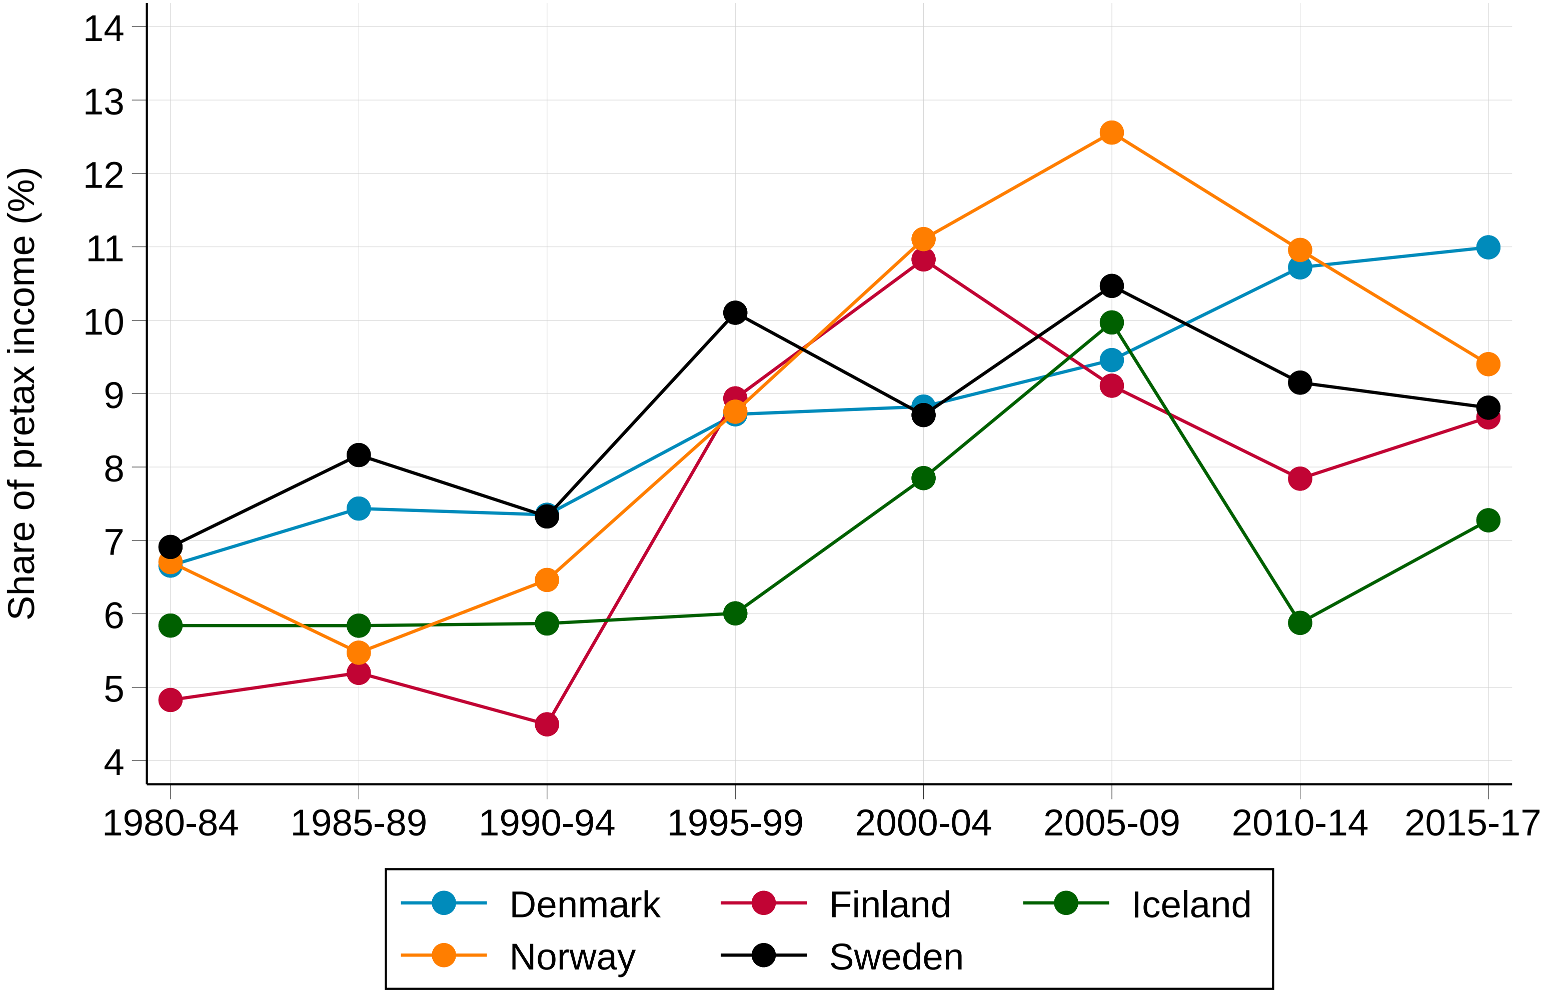

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure D.37: Top 1% pretax income share by country: Eastern Europe

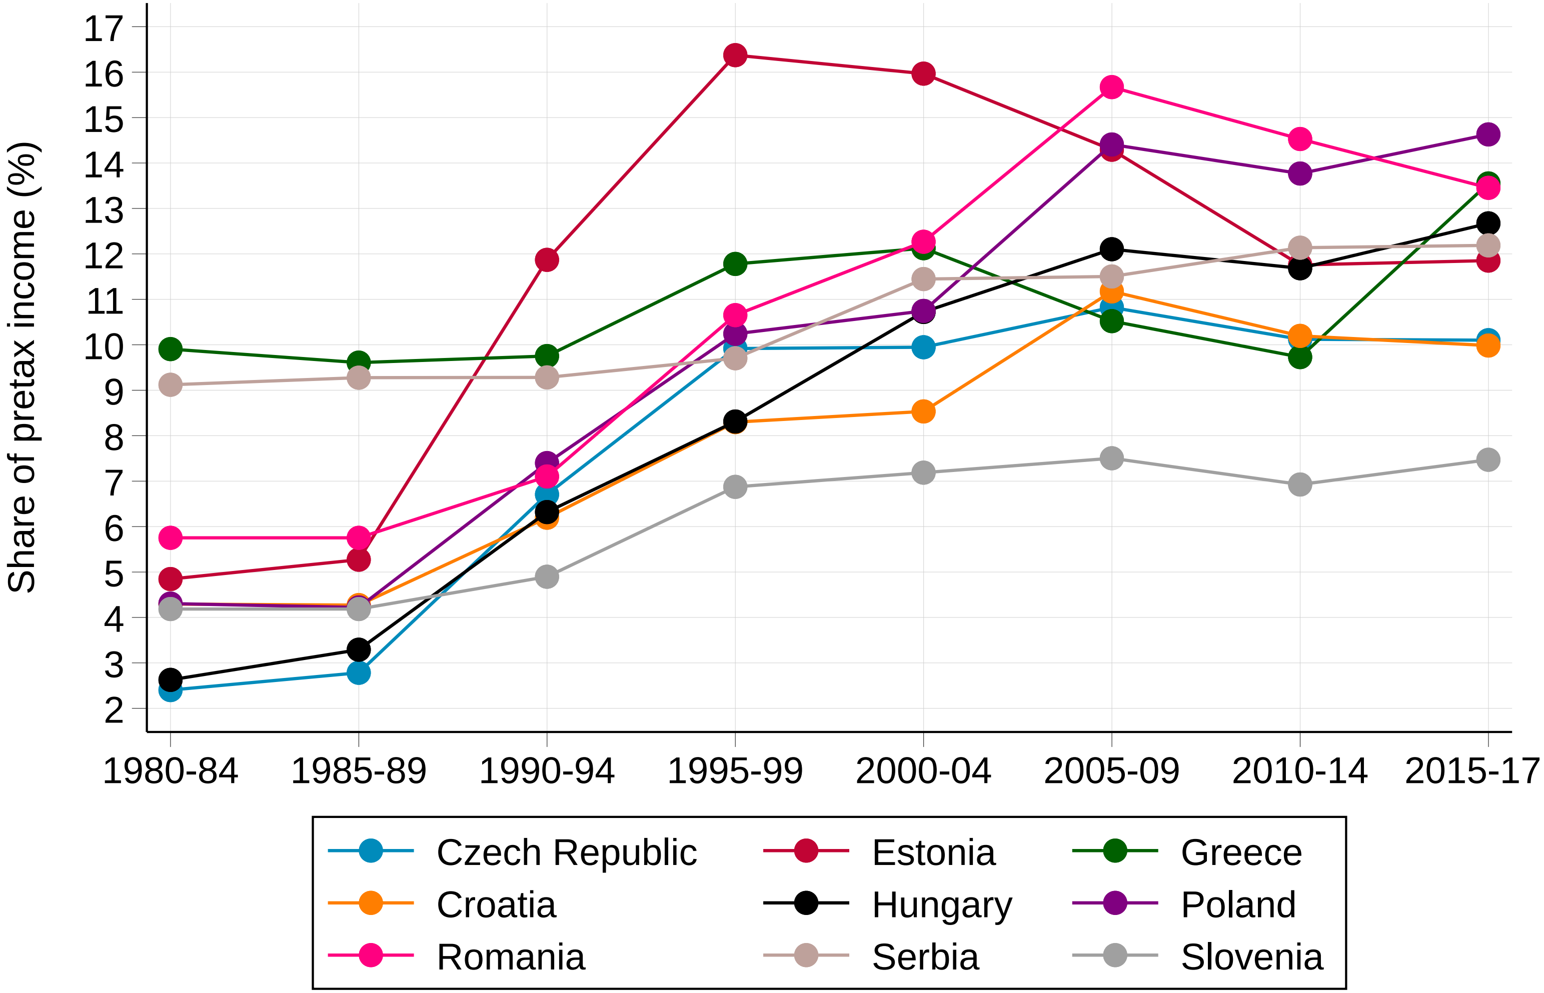

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure D.38: Bottom 50% pretax income share by country: Western Europe

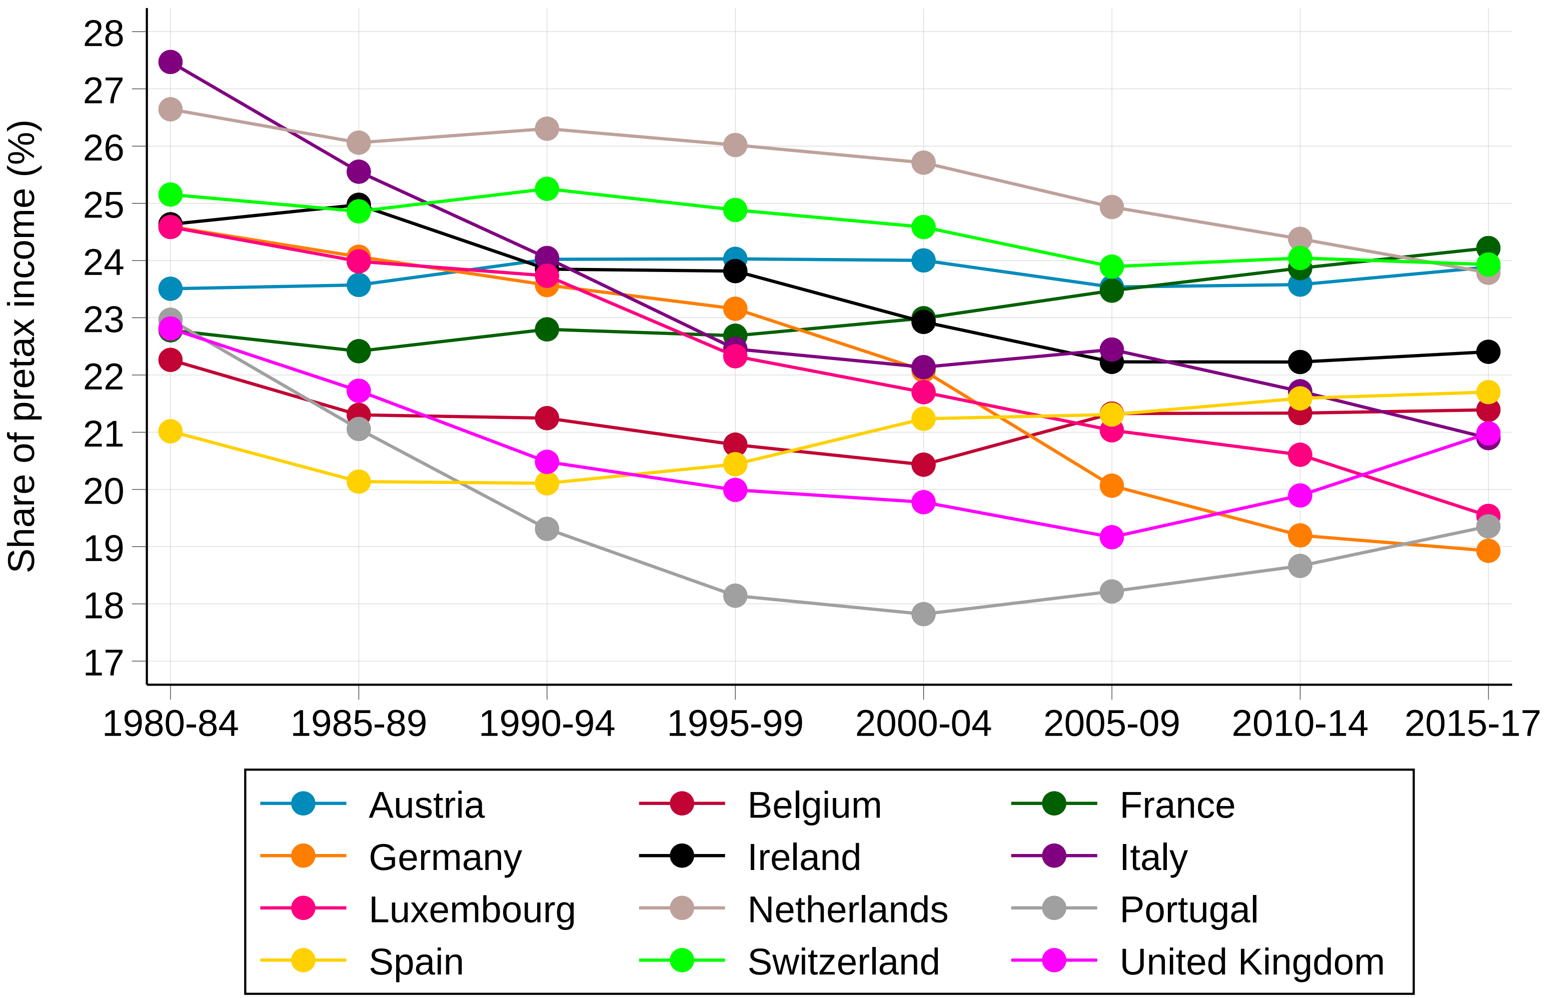

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure D.39: Bottom 50% pretax income share by country: Northern Europe

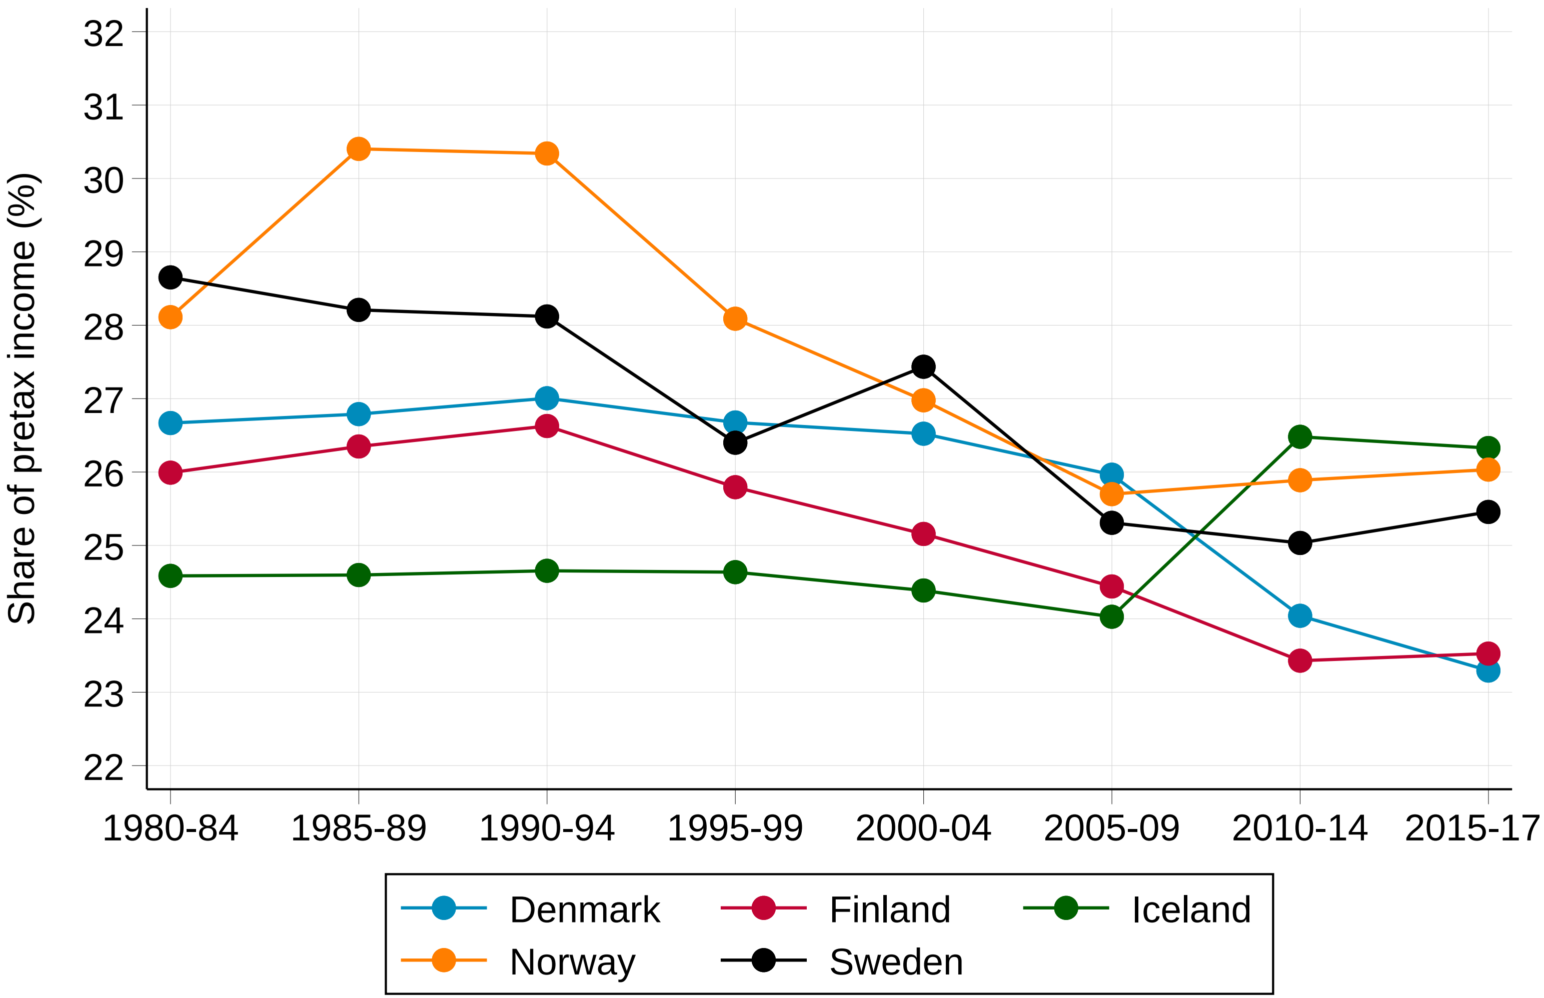

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure D.40: Bottom 50% pretax income share by country: Eastern Europe

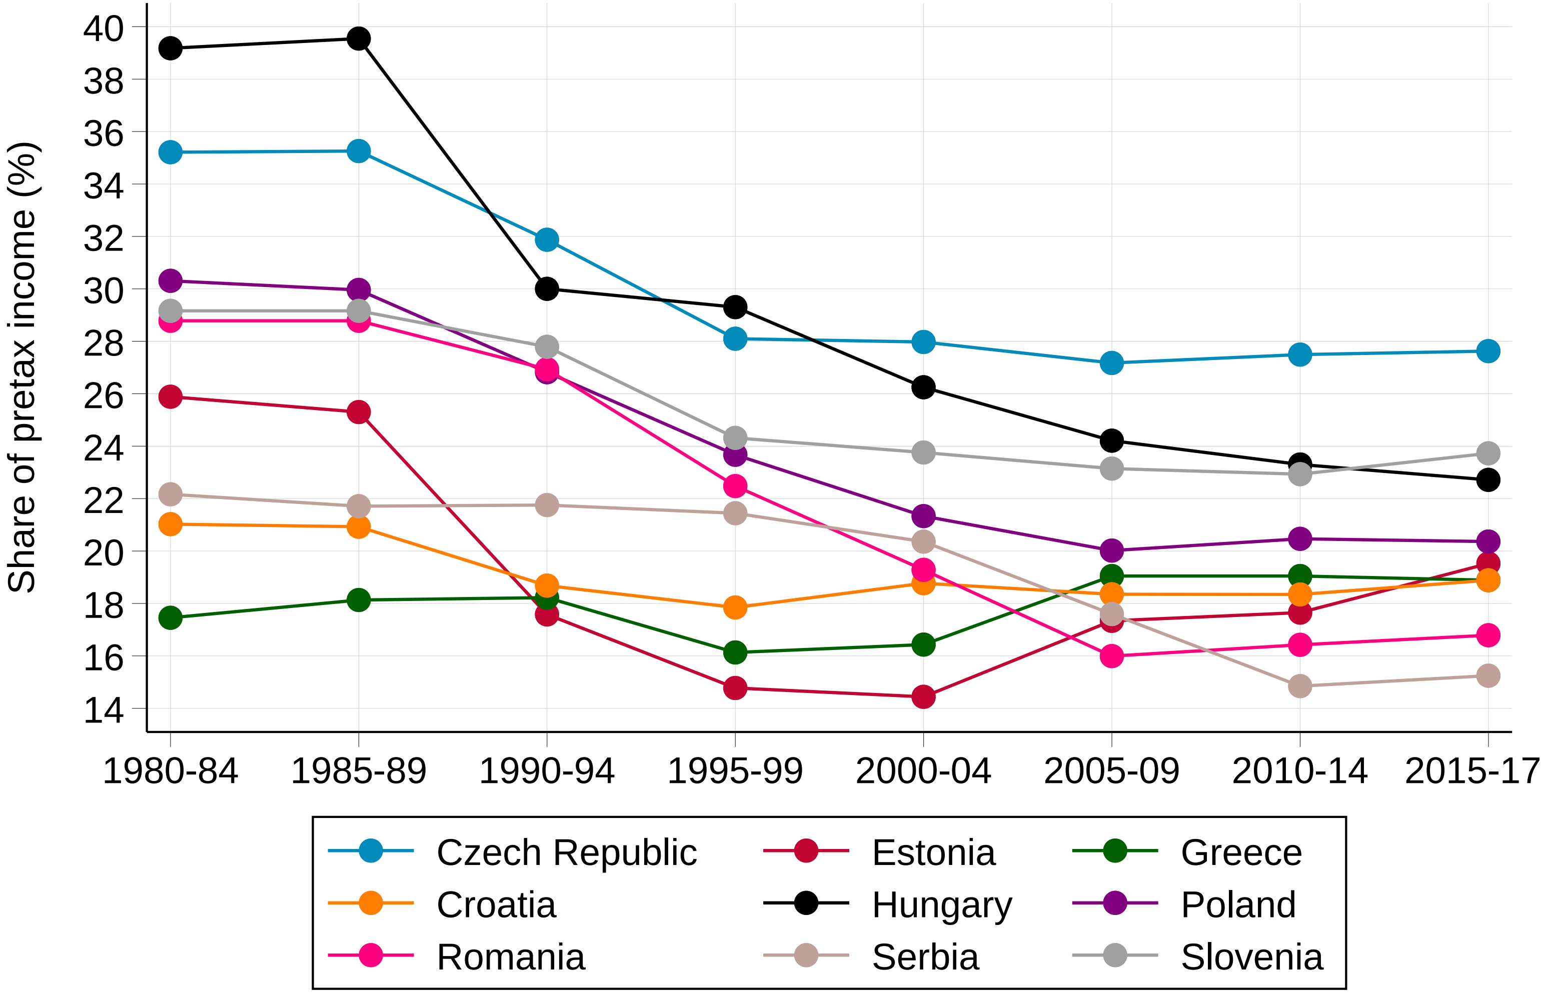

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure D.41: Top 10% pretax income share by country: 1980 versus 2017

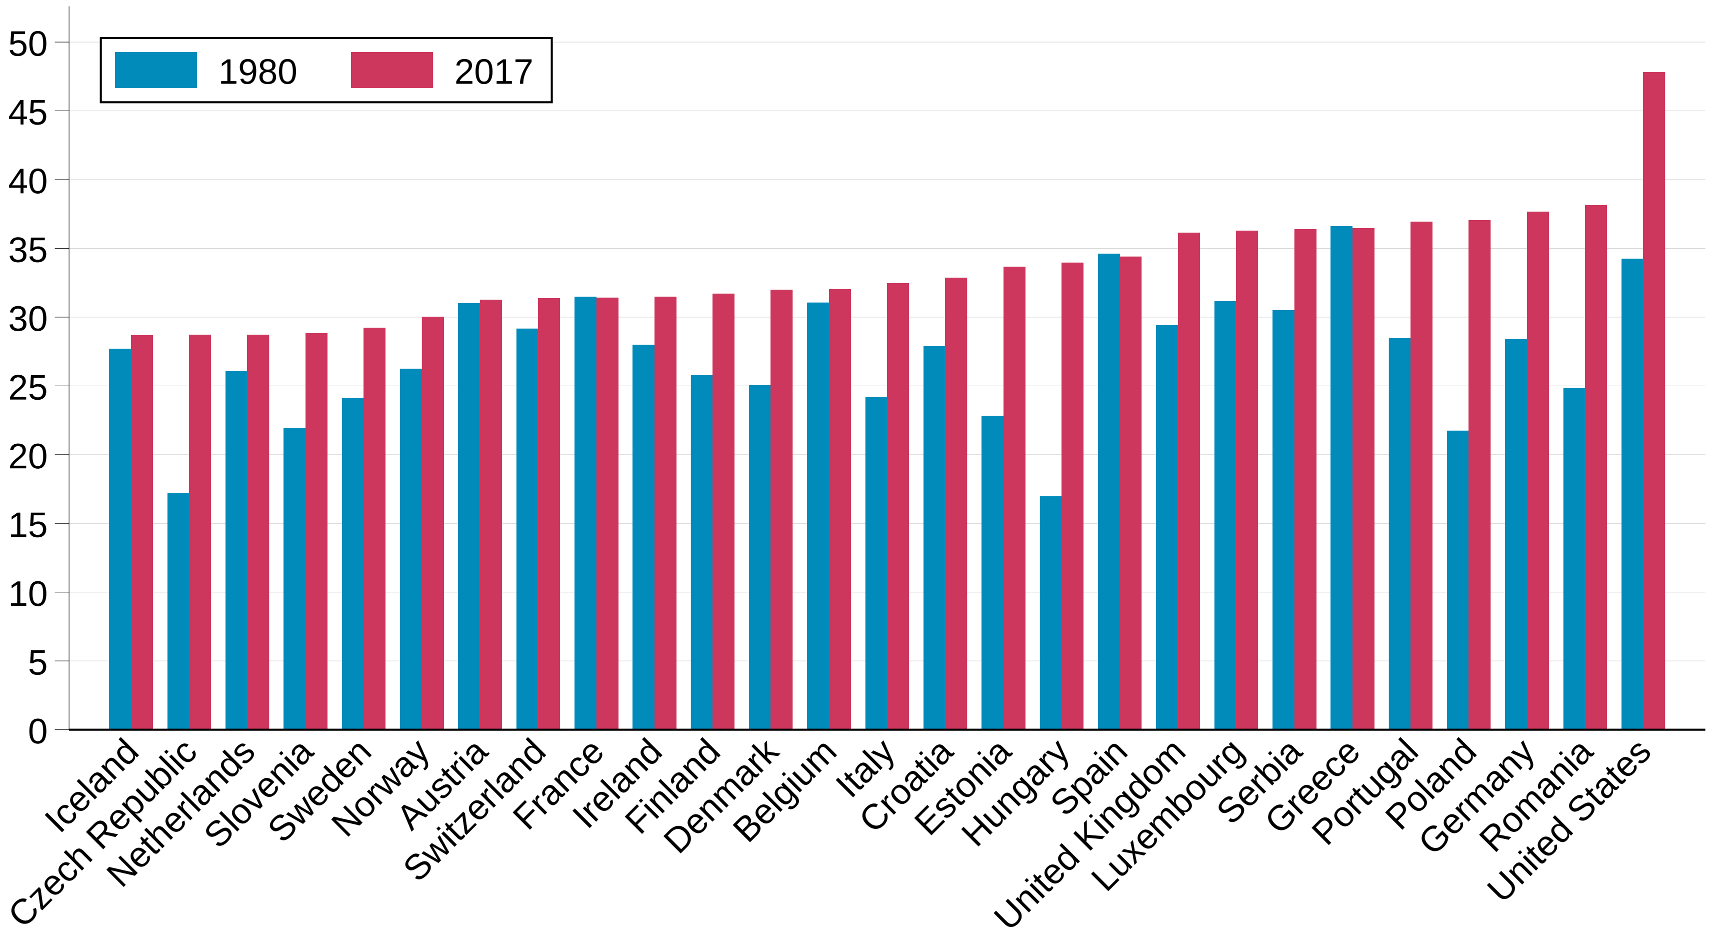

Source: Authors' computations combining surveys, tax data and national accounts for Europe; Piketty, Saez, and Zucman, 2018 for the US. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure D.42: Top 1% pretax income share by country: 1980 versus 2017

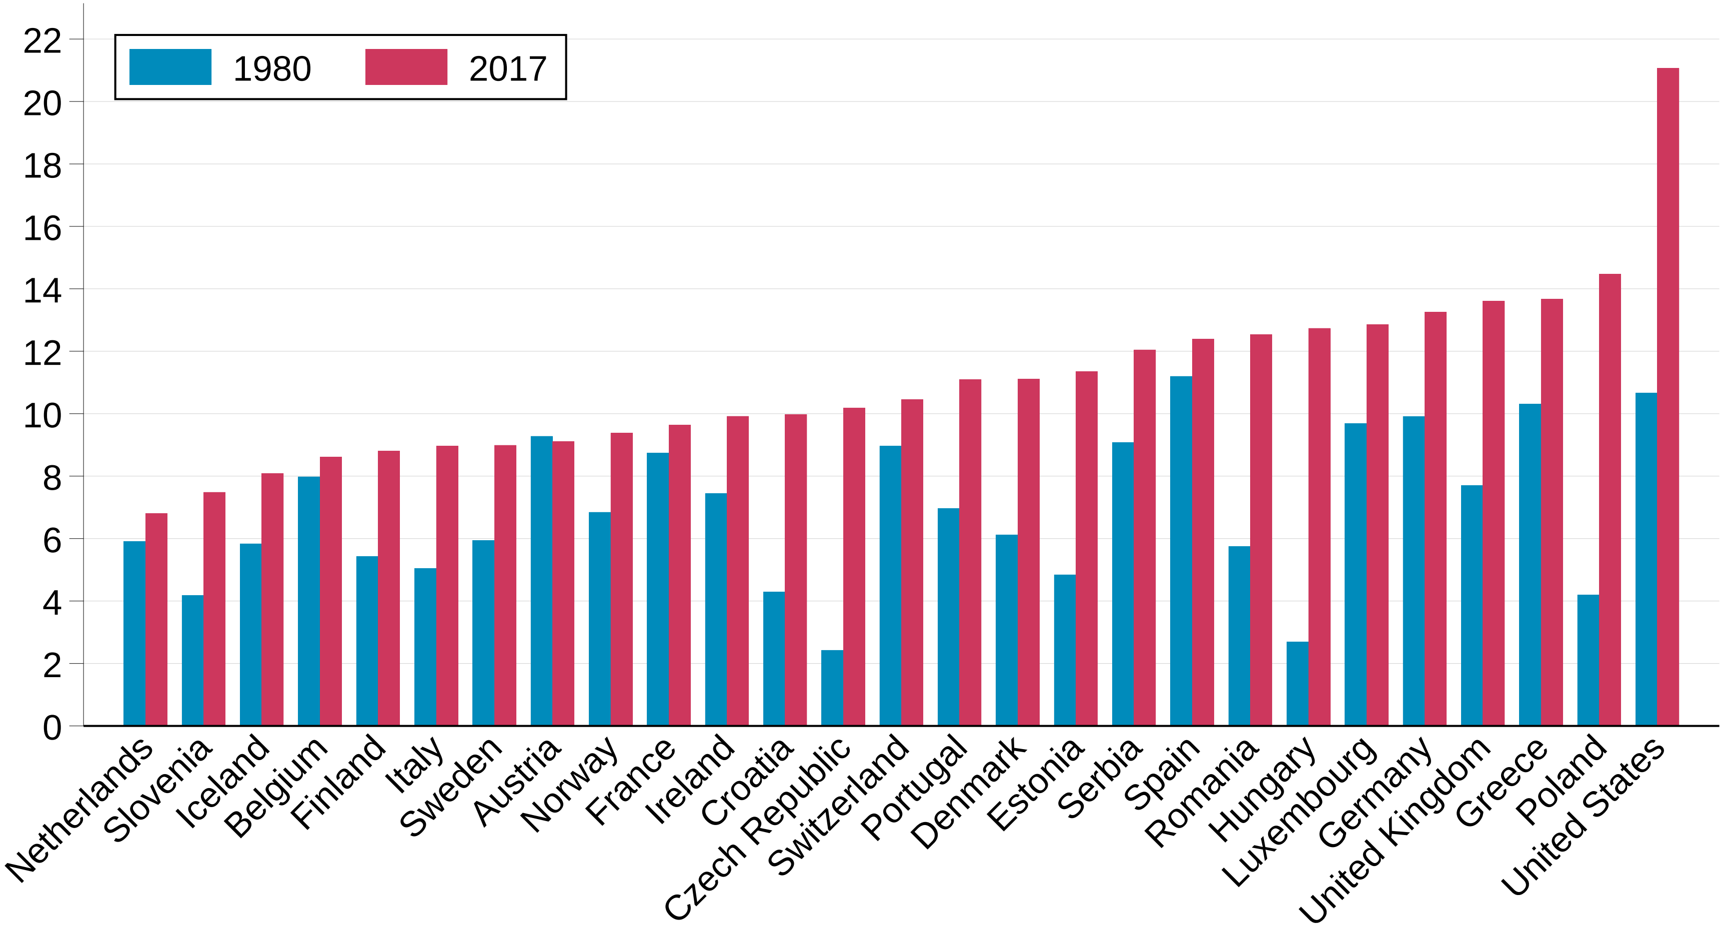

Source: Authors' computations combining surveys, tax data and national accounts for Europe; Piketty, Saez, and Zucman, 2018 for the US. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure D.43: Bottom 50% pretax income share by country: 1980 versus 2017

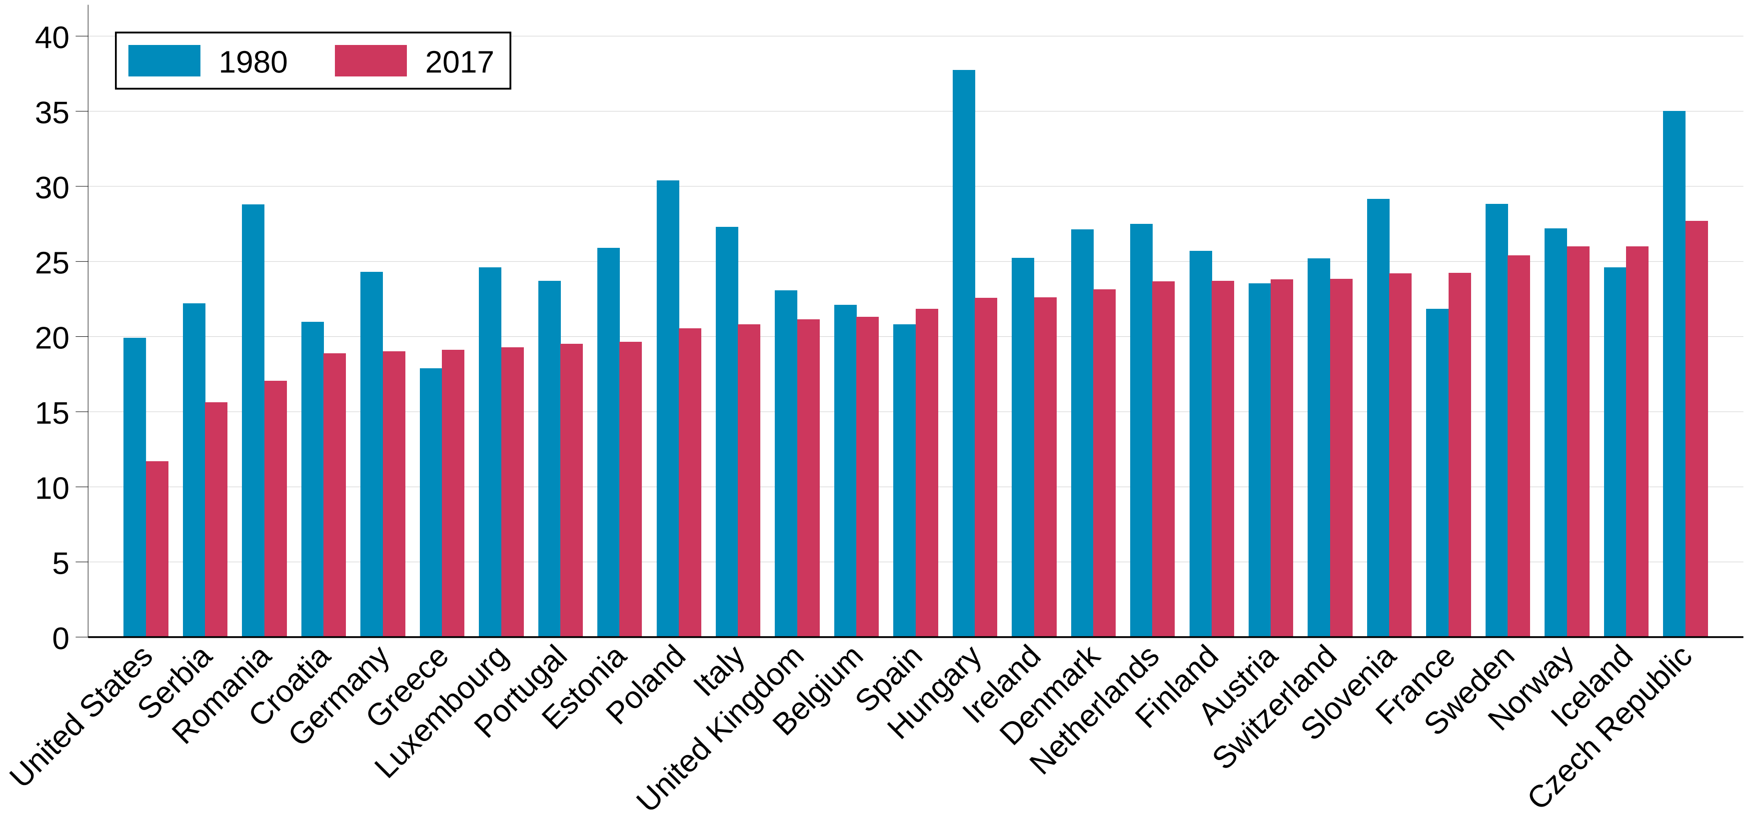

Source: Authors' computations combining surveys, tax data and national accounts for Europe; Piketty, Saez, and Zucman, 2018 for the US. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure D.44: Change in top 10% pretax income share by country, 1980-2017

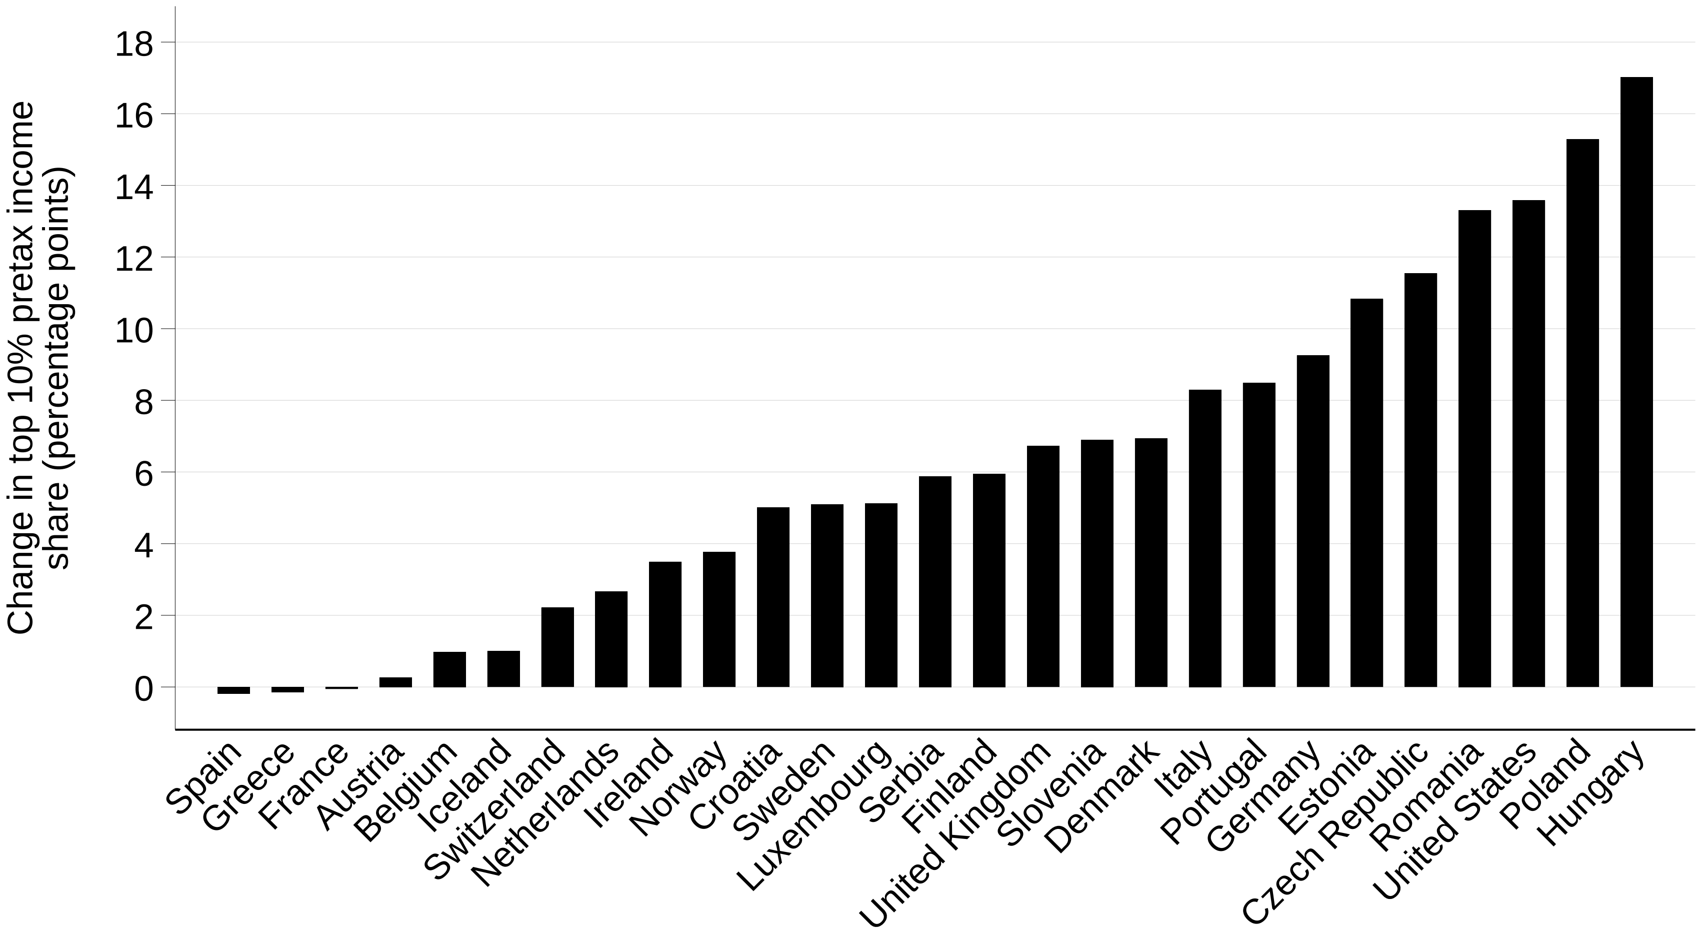

Source: Authors' computations combining surveys, tax data and national accounts for Europe; Piketty, Saez, and Zucman, 2018 for the US. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure D.45: Change in top 1% pretax income share by country, 1980-2017

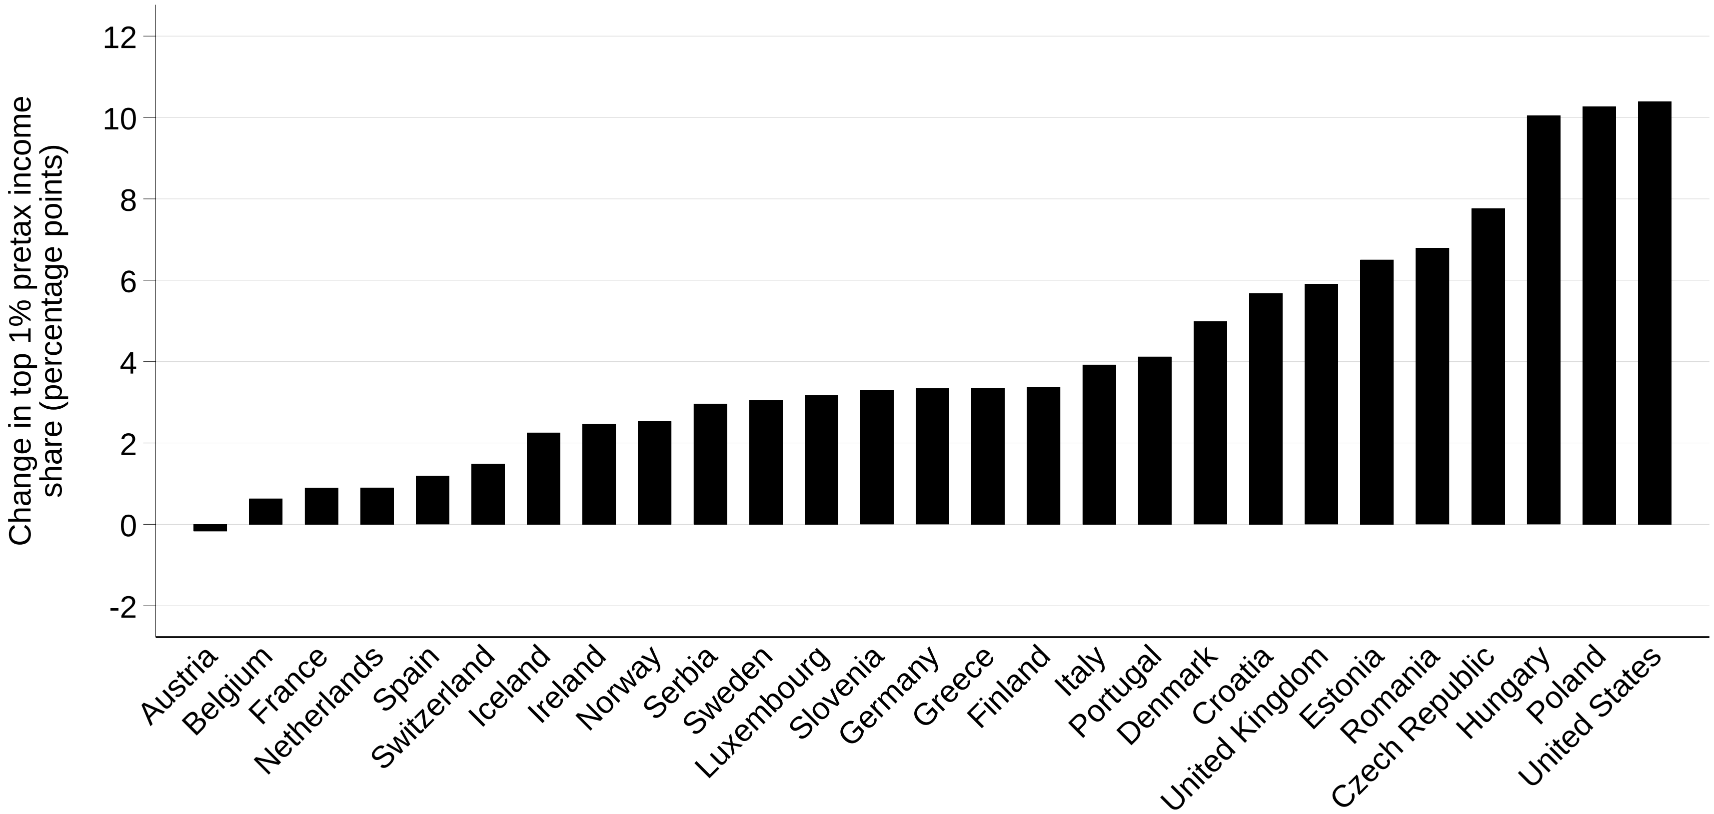

Source: Authors' computations combining surveys, tax data and national accounts for Europe; Piketty, Saez, and Zucman, 2018 for the US. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure D.46: Change in top 50% pretax income share by country, 1980-2017

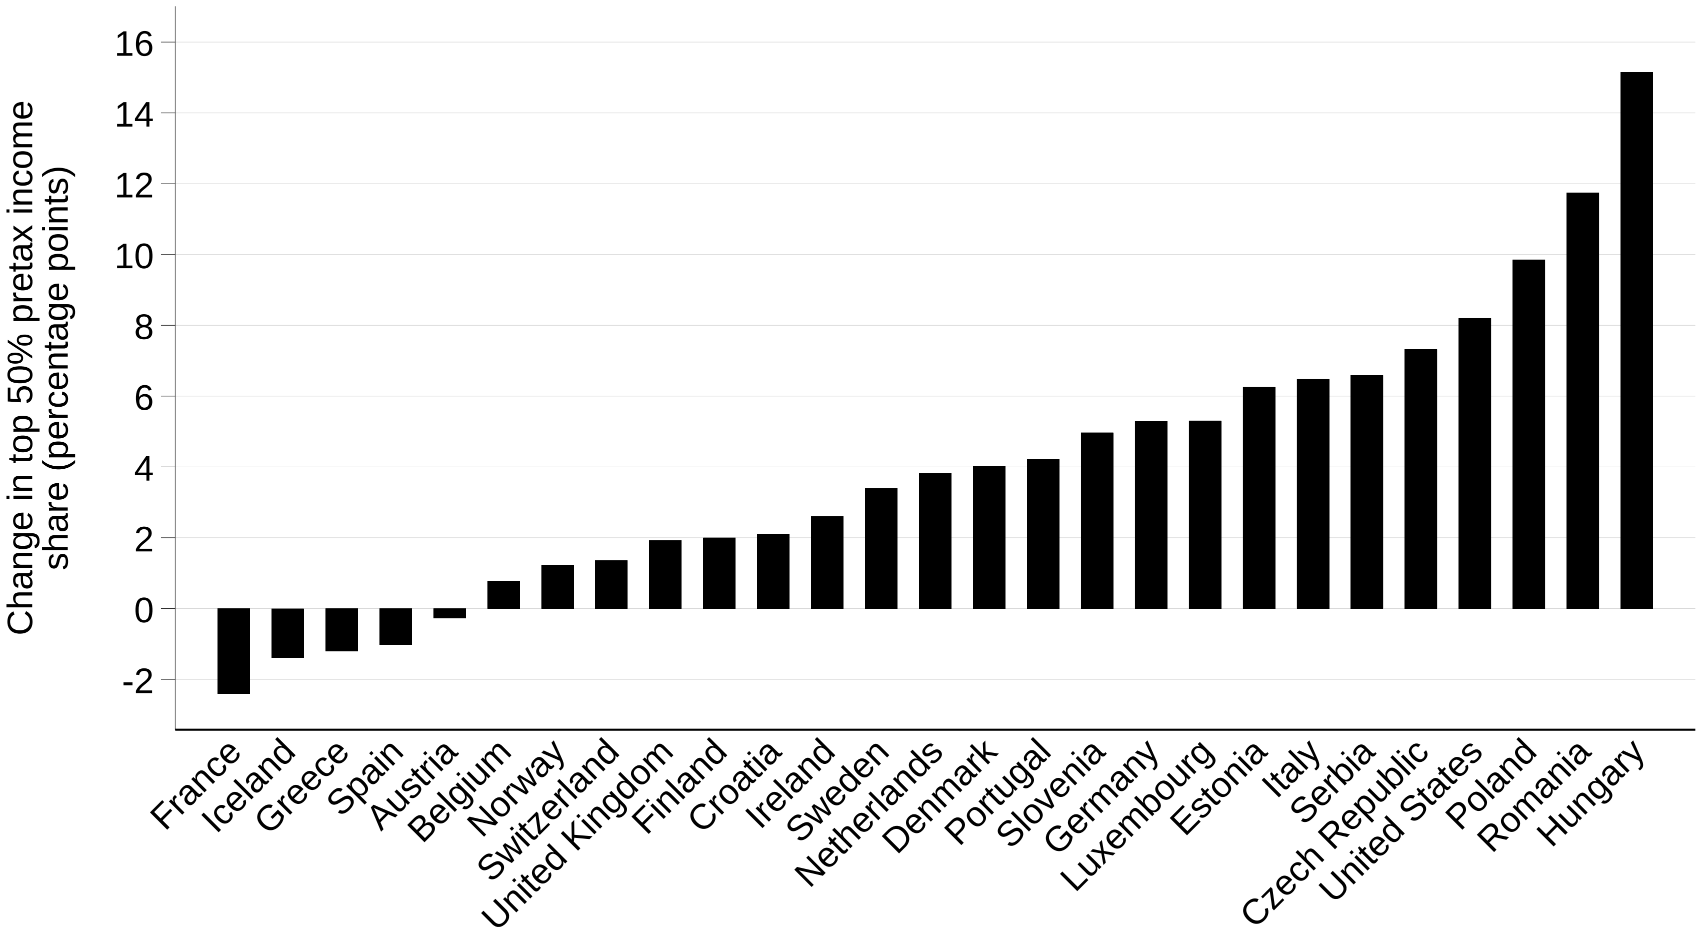

Source: Authors' computations combining surveys, tax data and national accounts for Europe; Piketty, Saez, and Zucman, 2018 for the US. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses.

Figure D.47: Average national incomes in Europe and the United States, 1980

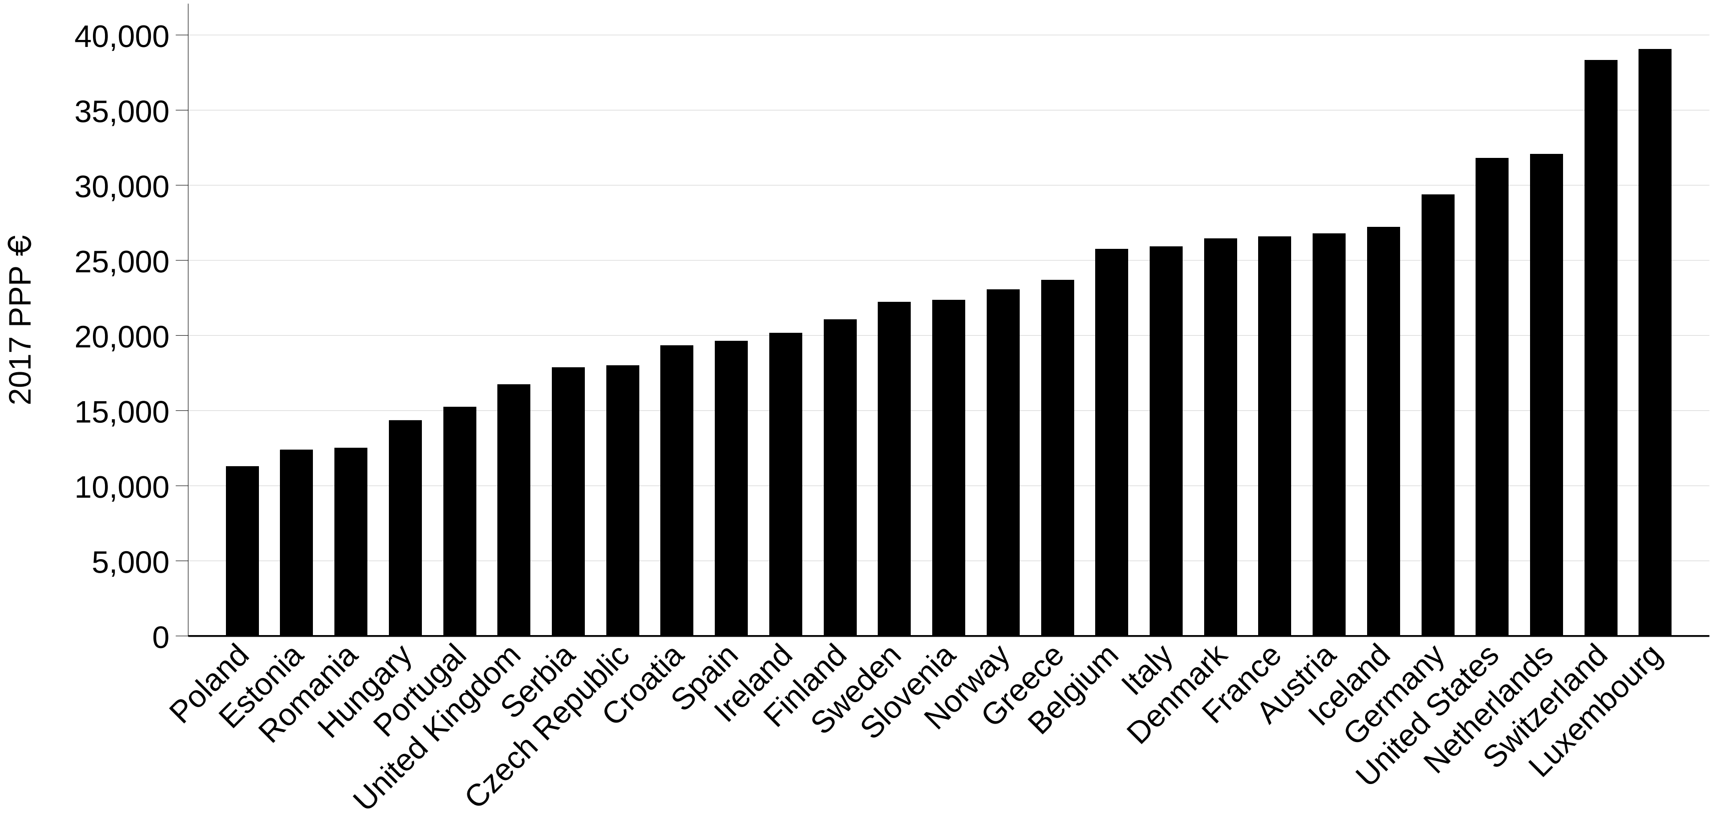

Source: Authors' computations combining surveys, tax data and national accounts for Europe; Piketty, Saez, and Zucman, 2018 for the US.

Figure D.48: Average national incomes in Europe and the United States, 2017

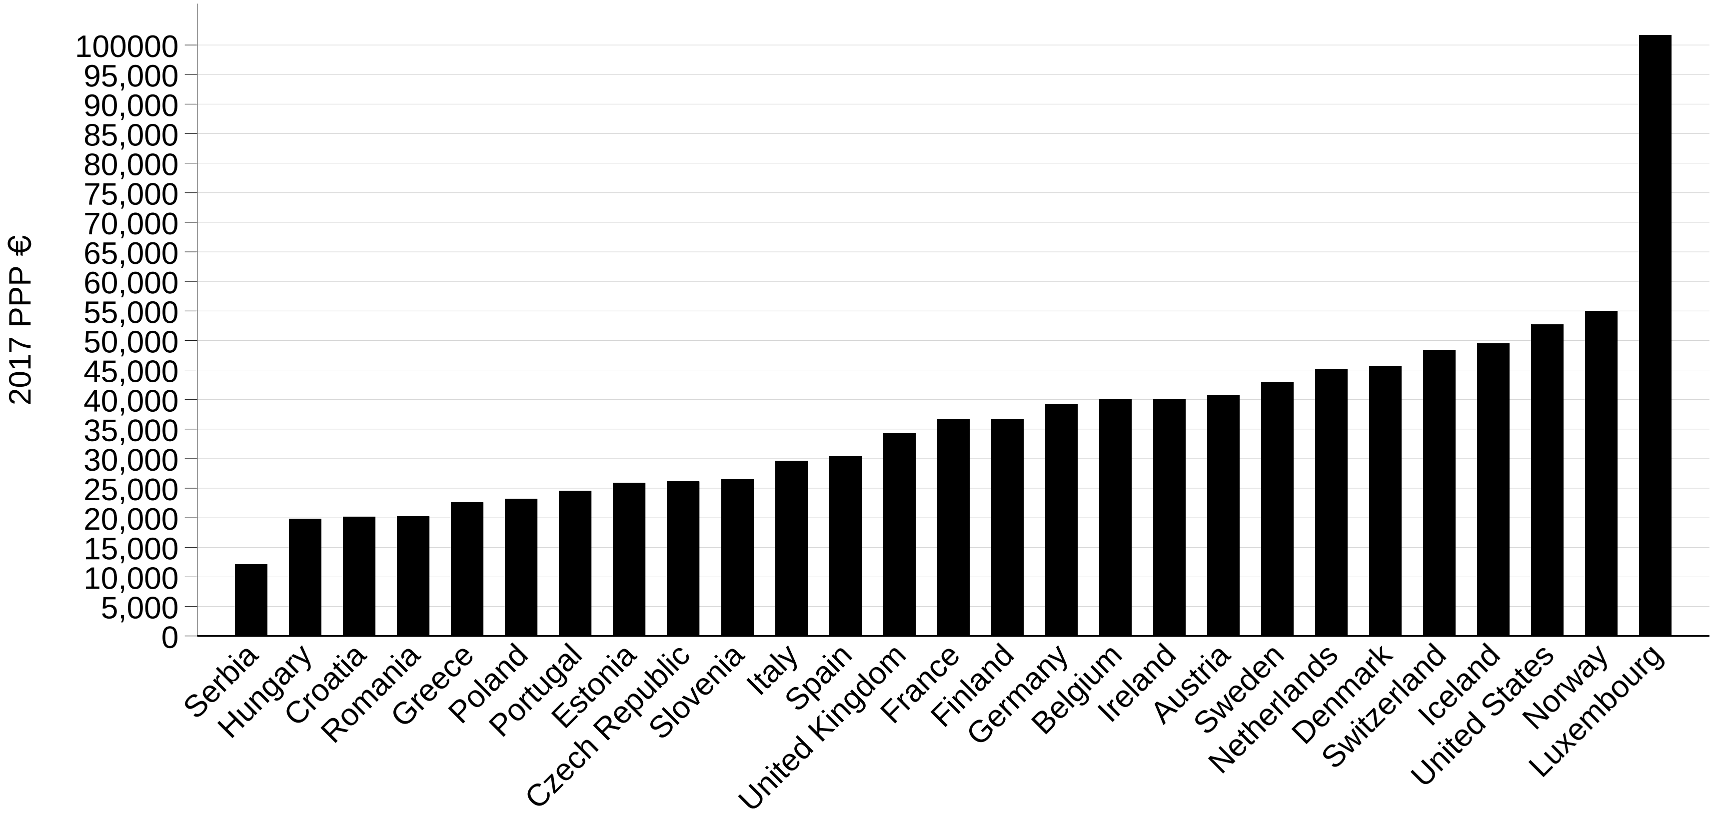

Source: Authors' computations combining surveys, tax data and national accounts for Europe; Piketty, Saez, and Zucman, 2018 for the US.

Figure D.49: Average bottom 50% pretax incomes in Europe and the United States, 1980

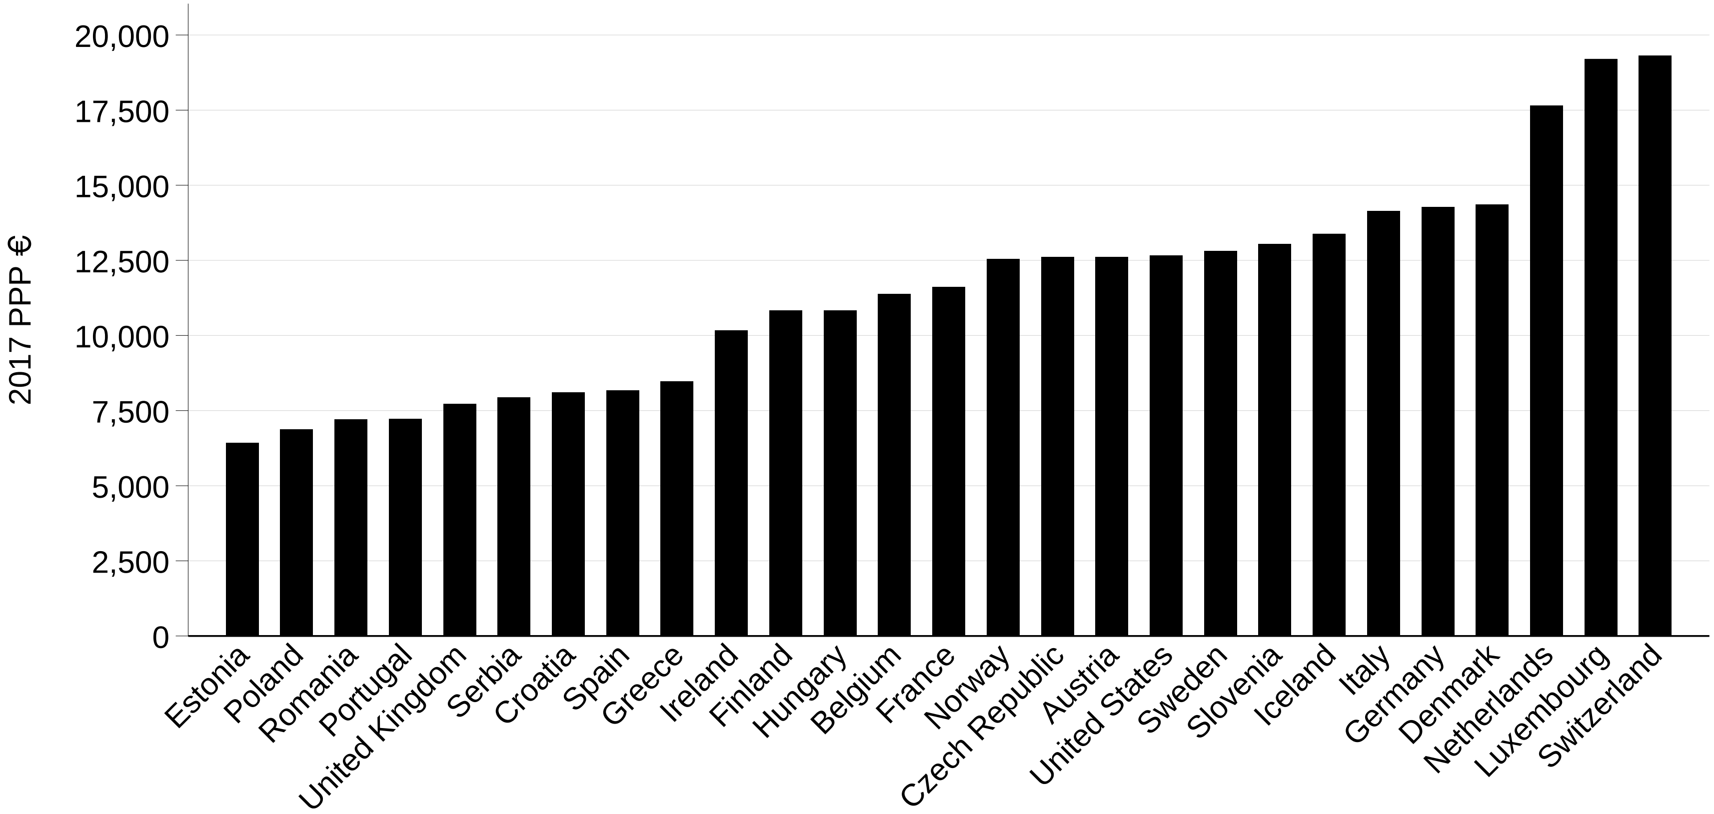

Source: Authors' computations combining surveys, tax data and national accounts for Europe; Piketty, Saez, and Zucman, 2018 for the US.

Figure D.50: Average bottom 50% pretax incomes in Europe and the United States, 2017

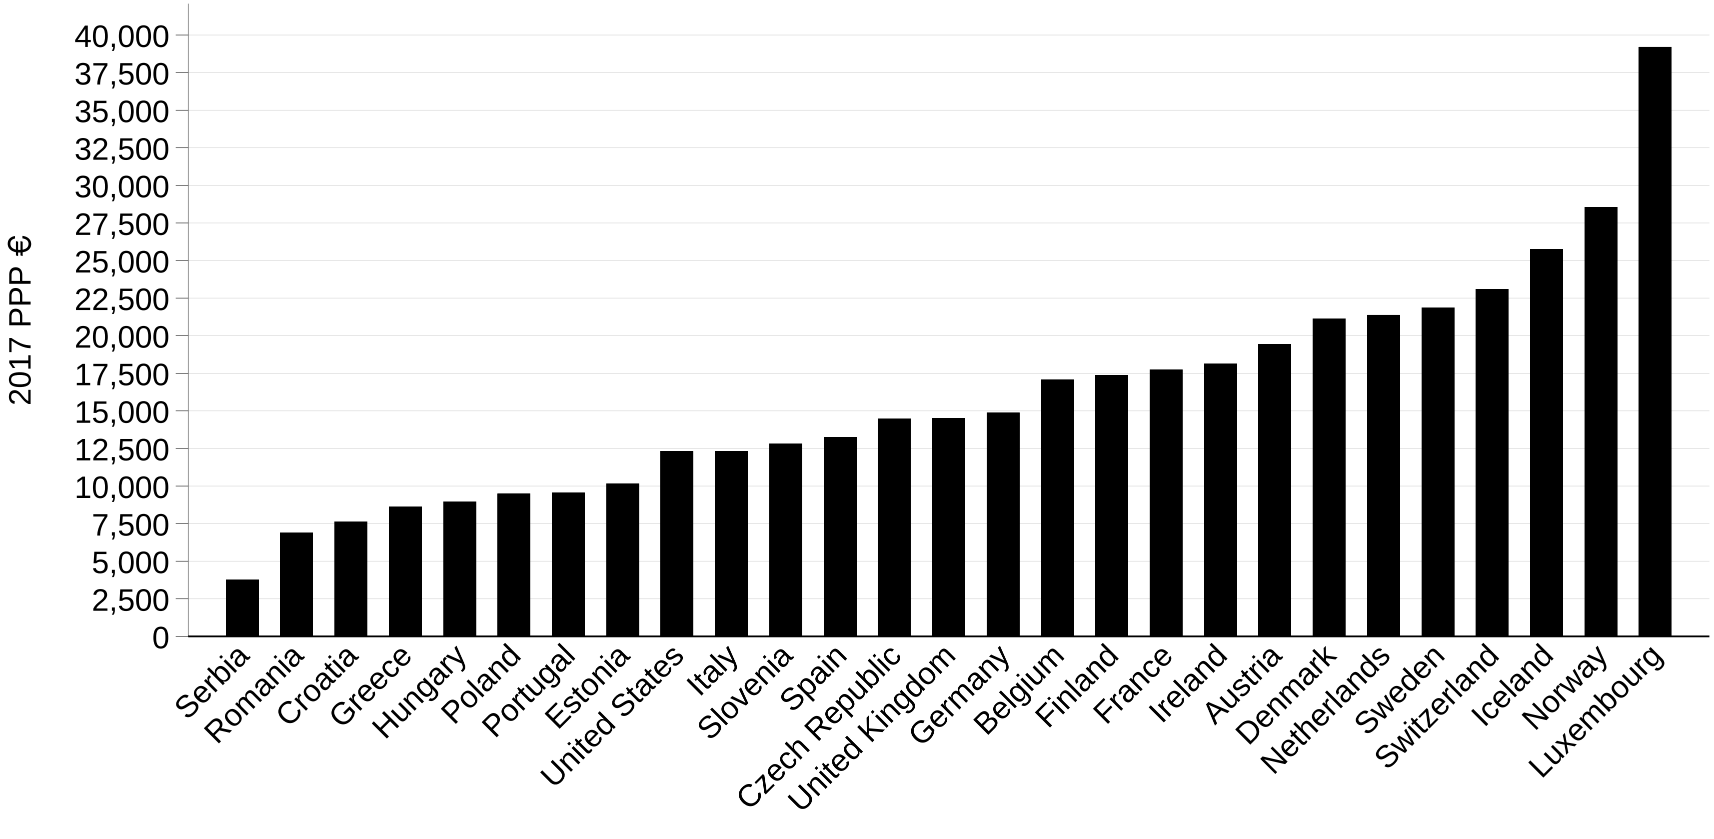

Source: Authors' computations combining surveys, tax data and national accounts for Europe; Piketty, Saez, and Zucman, 2018 for the US.

Figure D.51: Average top 10% pretax incomes in Europe and the United States, 1980

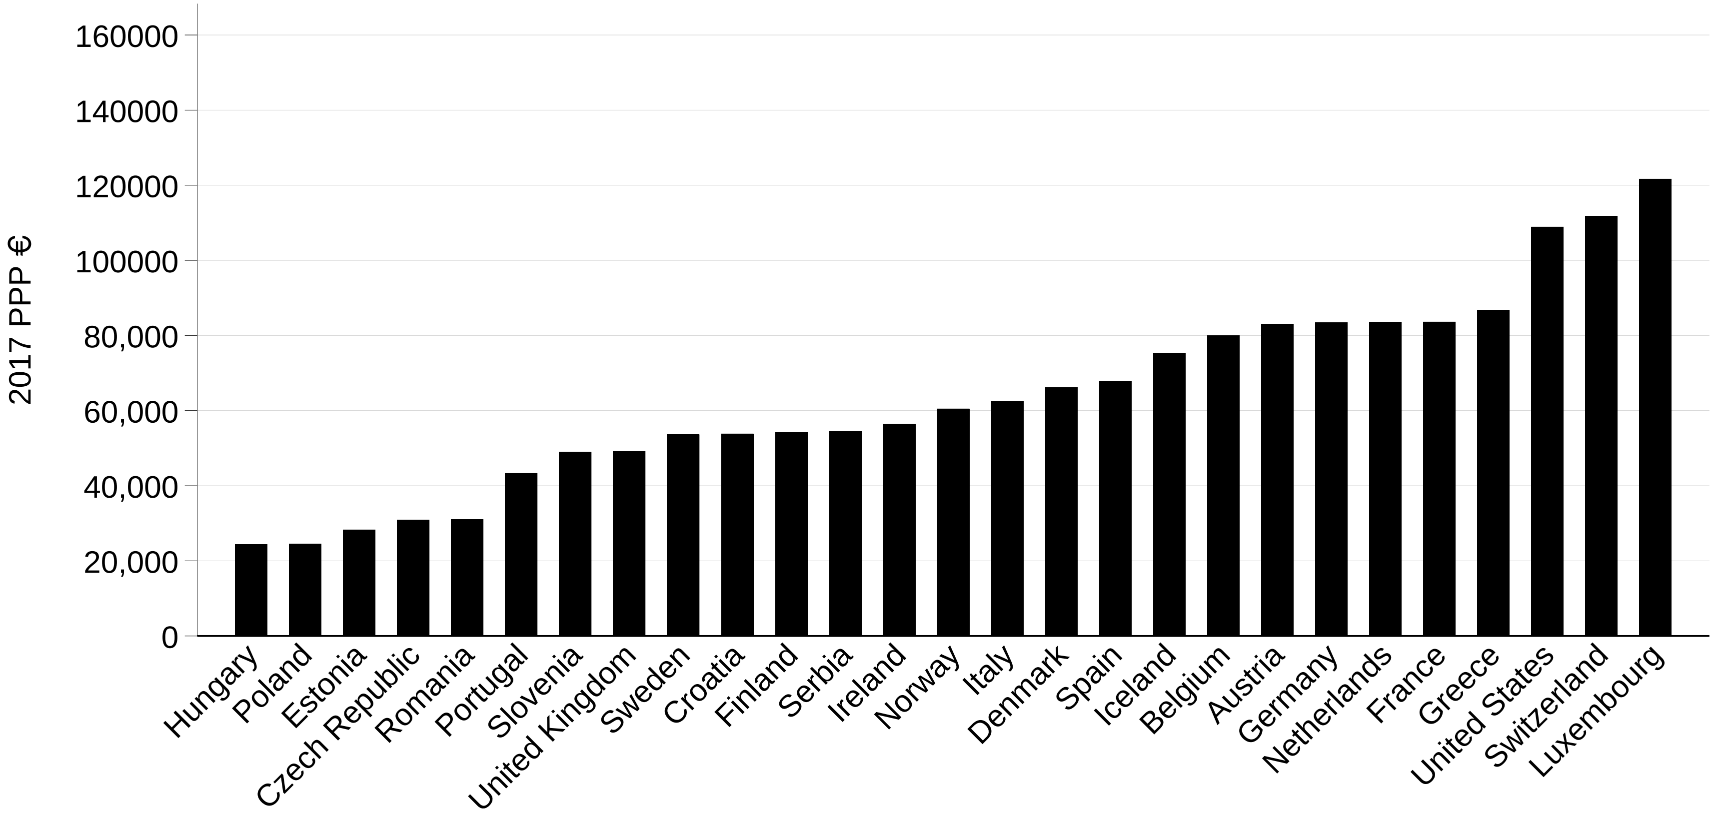

Source: Authors' computations combining surveys, tax data and national accounts for Europe; Piketty, Saez, and Zucman, 2018 for the US.

Figure D.52: Average top 10% pretax incomes in Europe and the United States, 2017

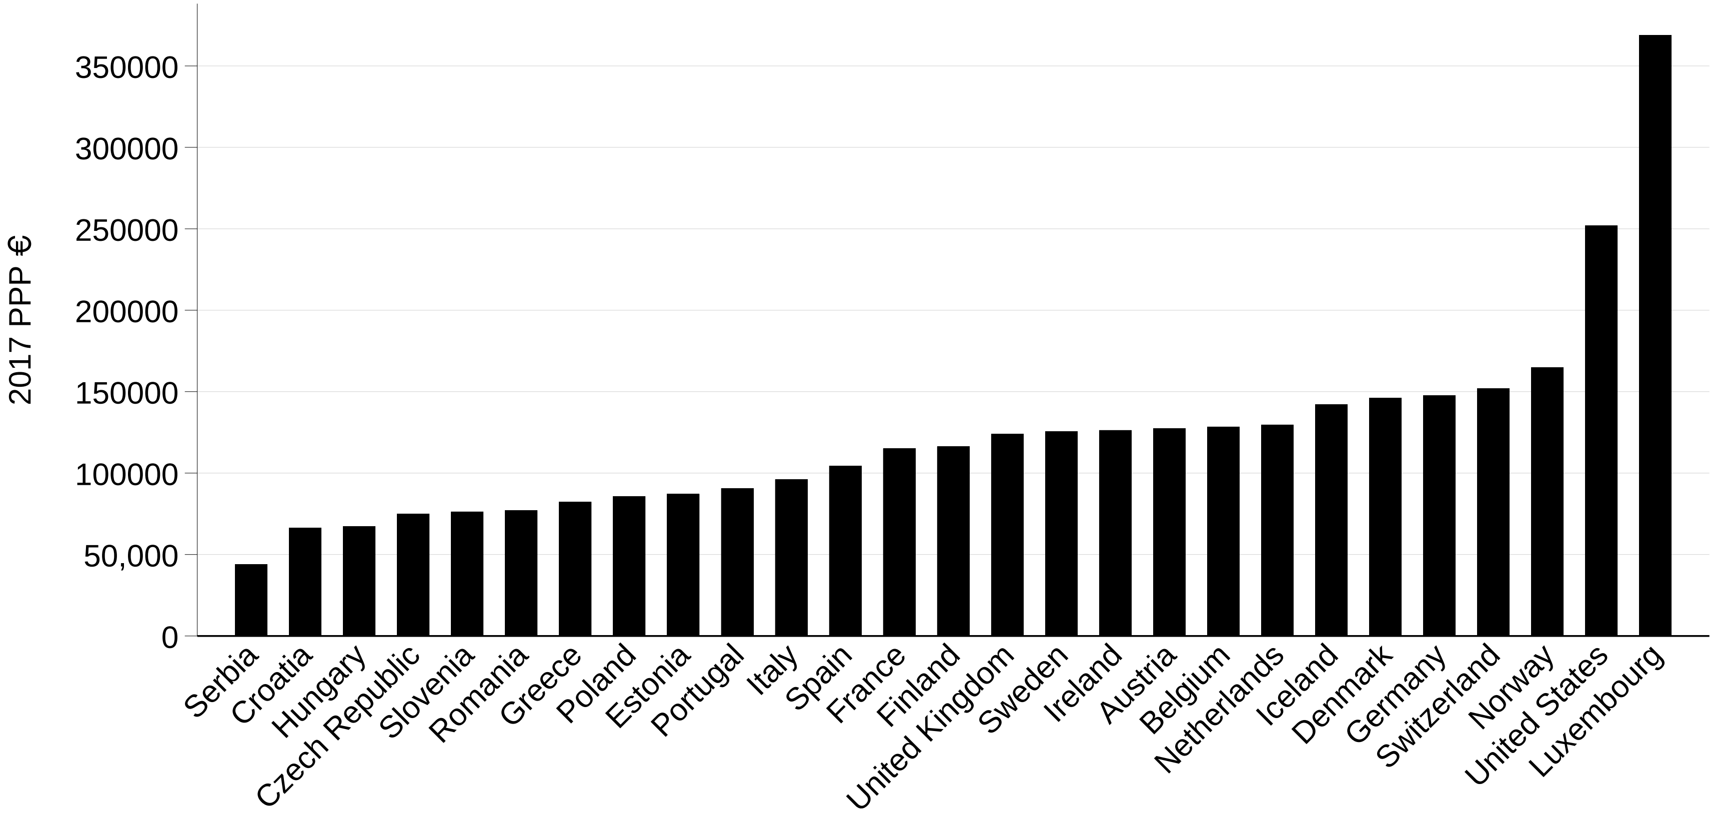

Source: Authors' computations combining surveys, tax data and national accounts for Europe; Piketty, Saez, and Zucman, 2018 for the US.

D.2.3 Distribution of taxes

Figure D.53: Total taxes paid as a share of factor income (working-age population) in Europe and the United States

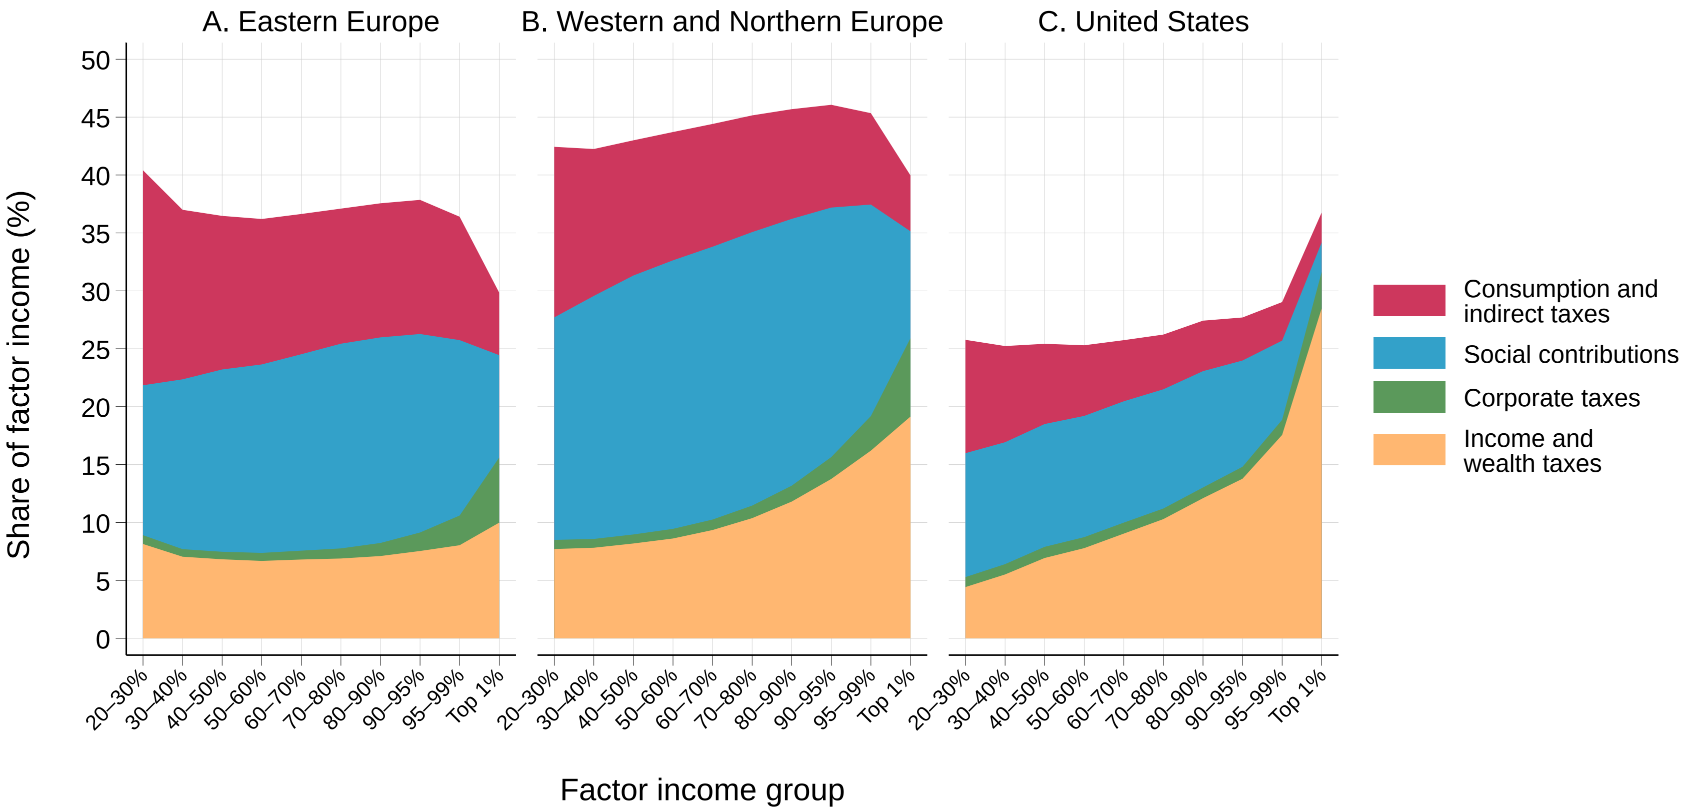

Source: Authors' computations combining surveys, tax data and national accounts. *Notes.* The figure represents total taxes (including all direct and indirect taxes, as well as all social contributions) paid by factor income group, expressed as a share of factor income. The data correspond to population-weighted averages over the period 2007–2017 for Europe, and to 2017–2018 for the US. The unit of observation is the adult individual aged between 25 and 59 (working-age population). Income is split equally among spouses.

Figure D.54: Effective tax rate of the top 10% by country (non-contributory taxes, % of pretax income)

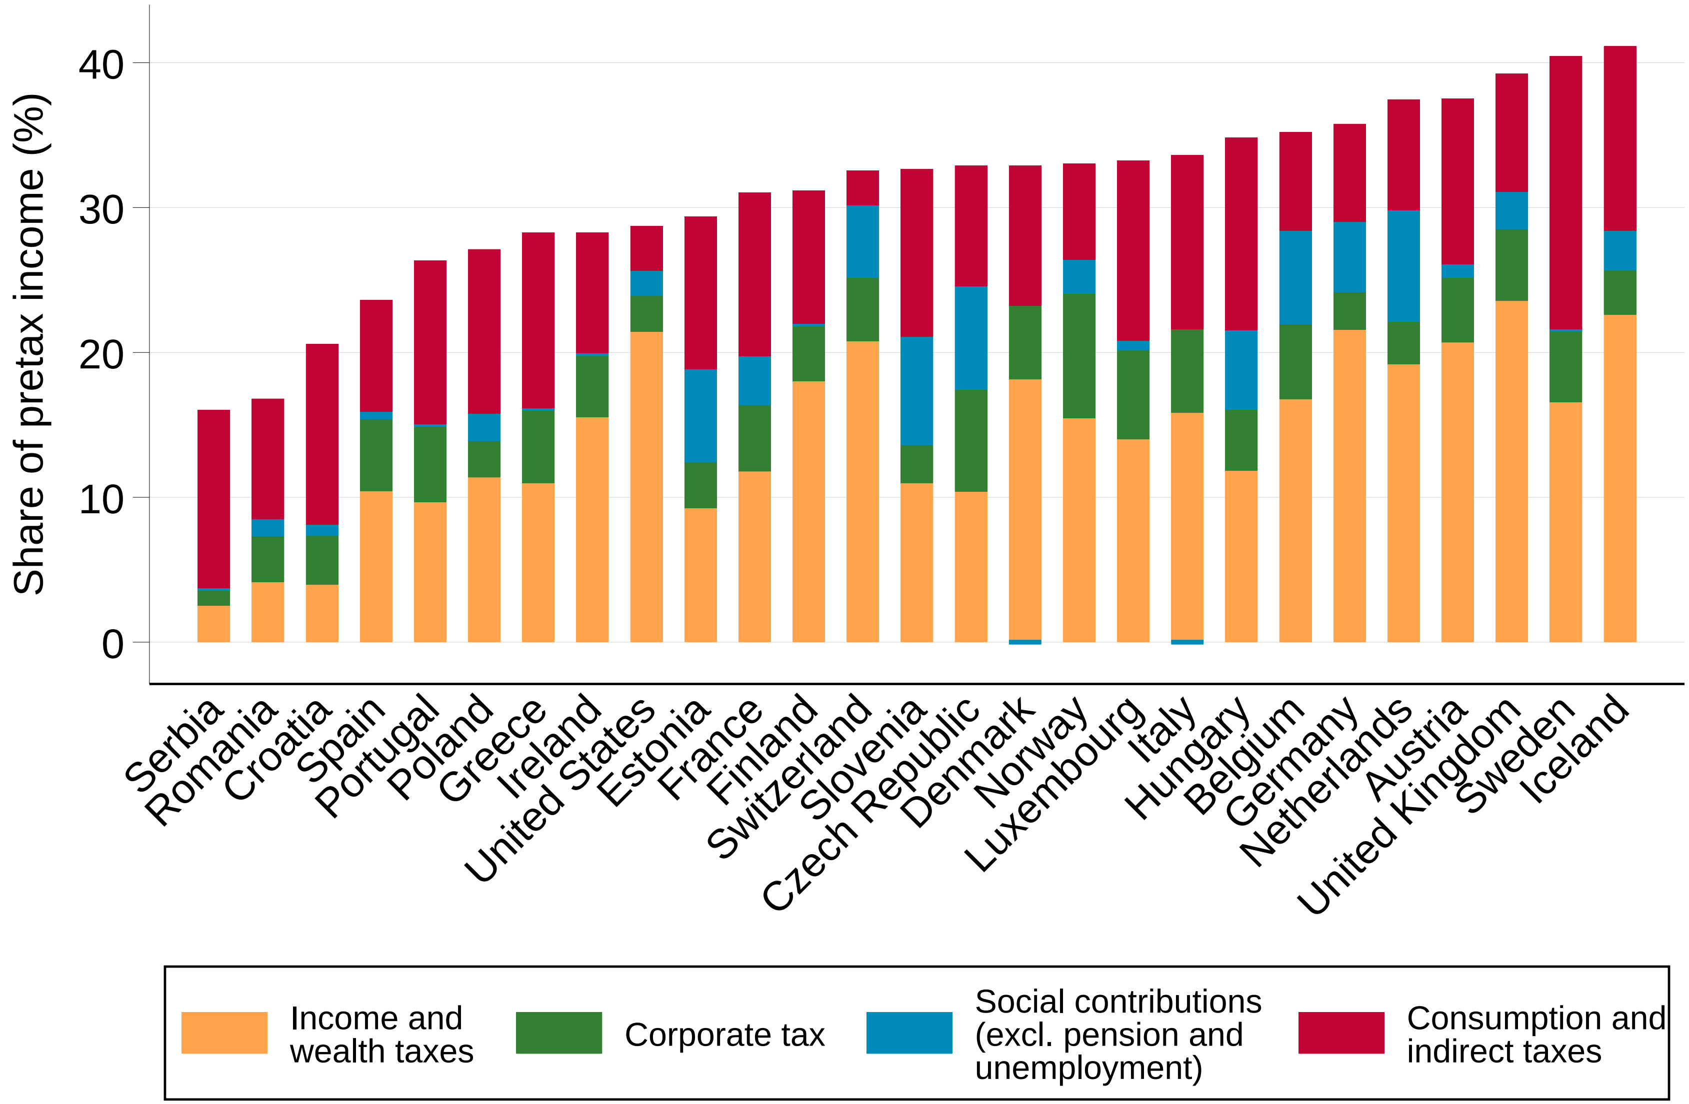

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. Average over the 2007-2017 period.

Figure D.55: Effective tax rate of the bottom 50% by country (non-contributory taxes, % of pretax income)

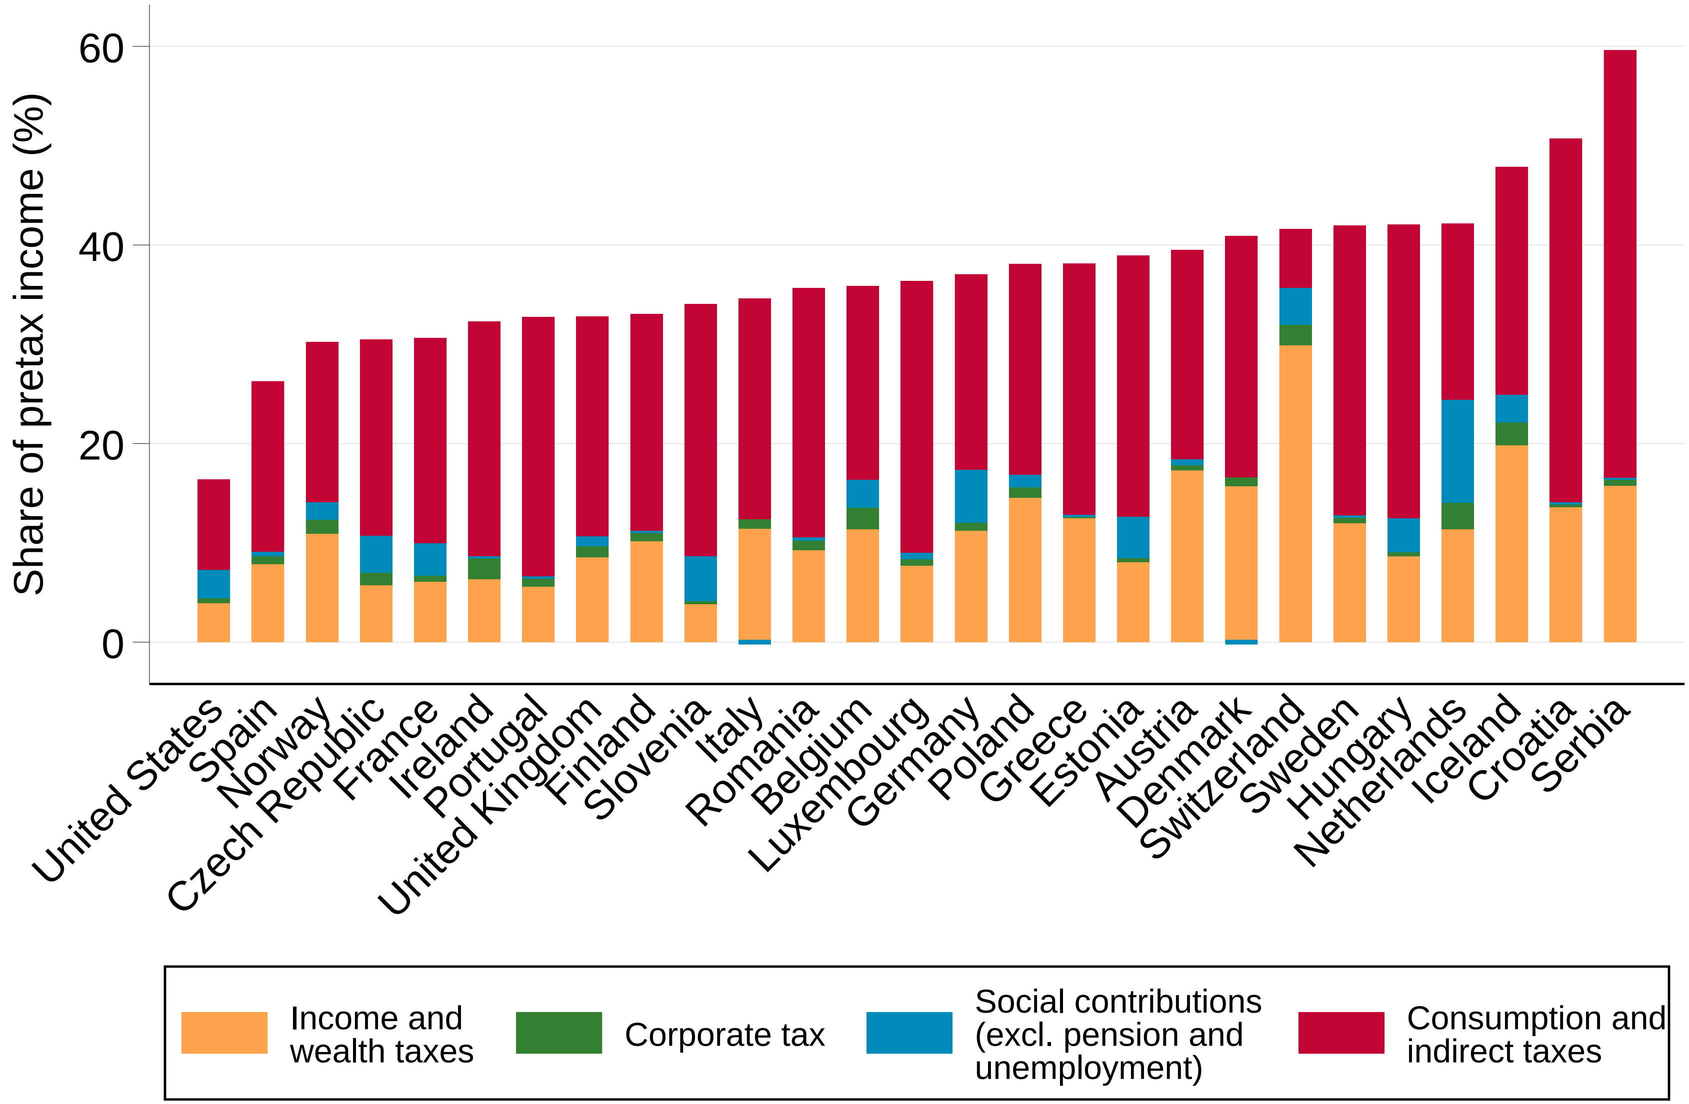

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. Average over the 2007-2017 period.

Figure D.56: Ratio of top 10% to bottom 50% effective tax rates by country (non-contributory taxes, % of pretax income)

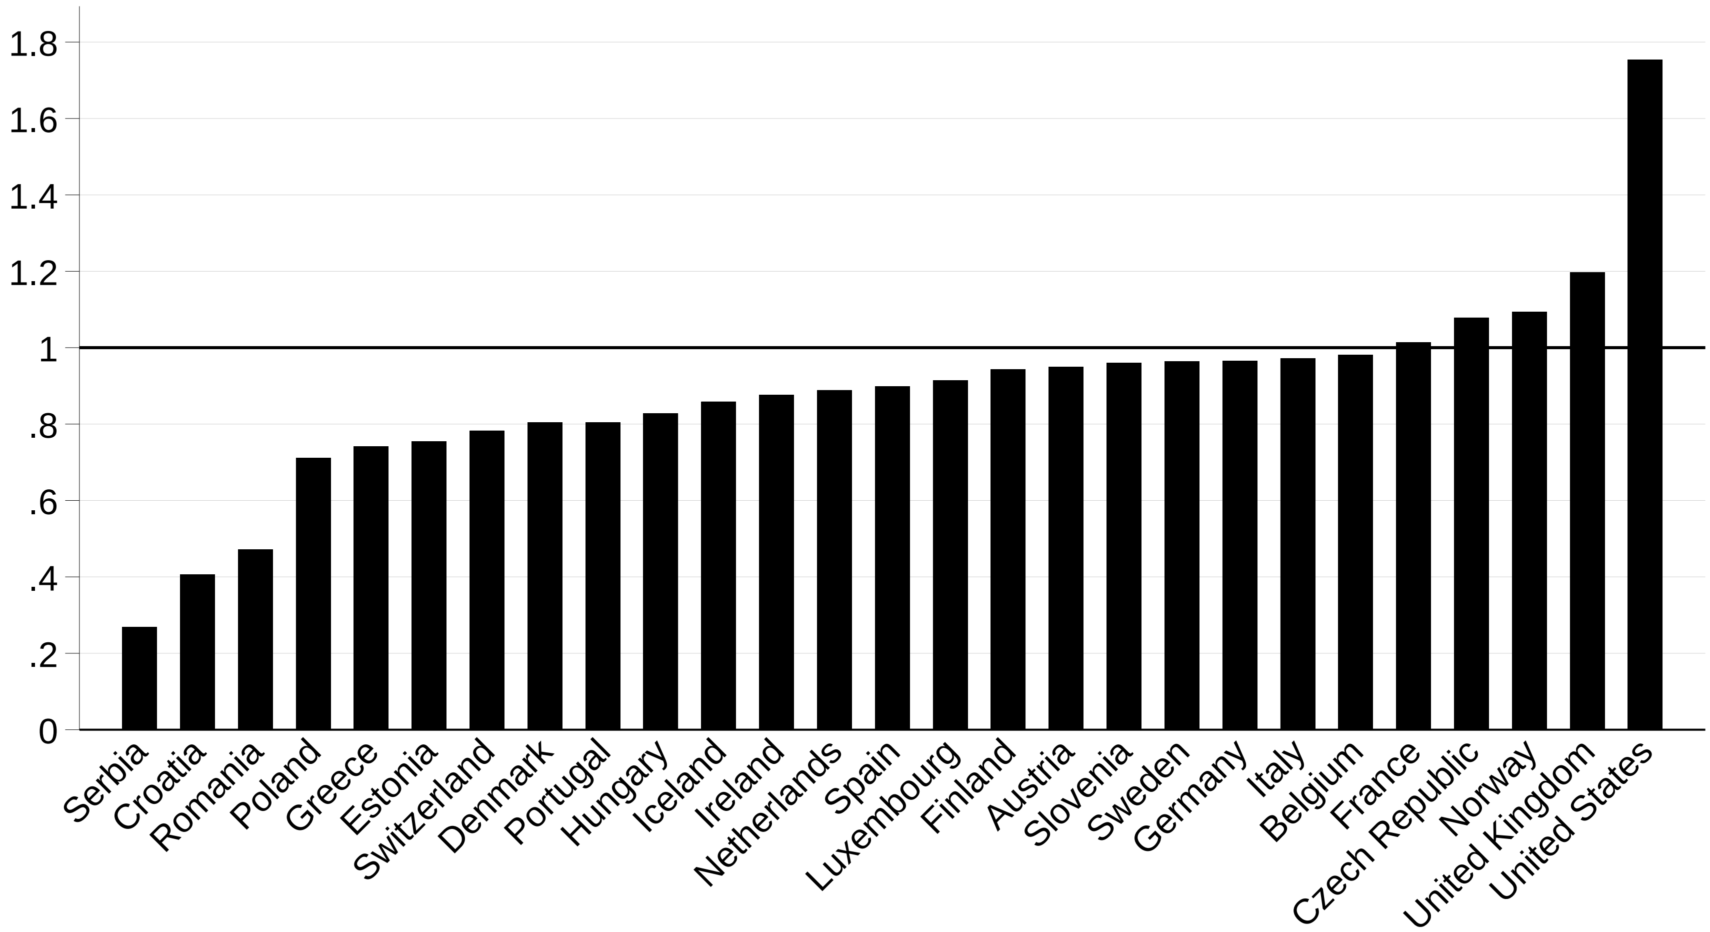

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. Average over the 2007-2017 period.

Figure D.57: Effective tax rate of the top 10% by country (all taxes, % of factor income, working-age population)

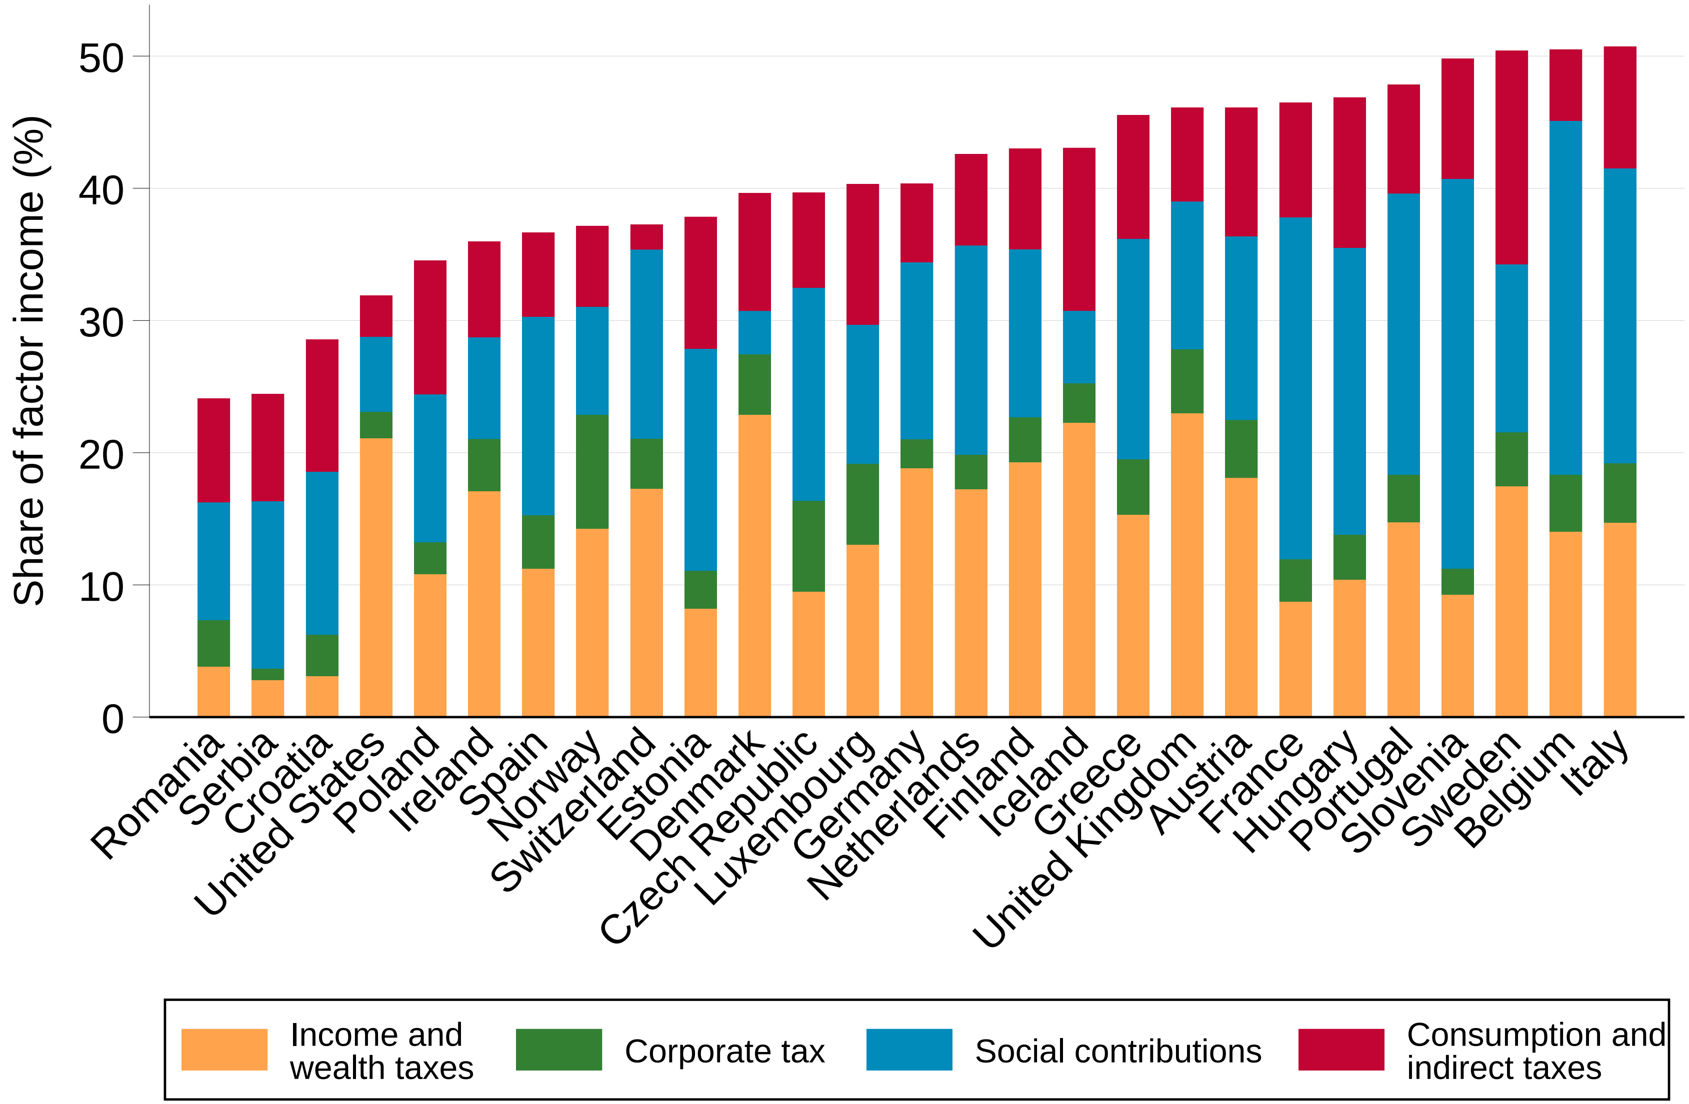

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged between 25 and 59. Income is split equally among spouses. Average over the 2007-2017 period.

Figure D.58: Effective tax rate of the bottom 50% by country (all taxes, % of factor income, working-age population)

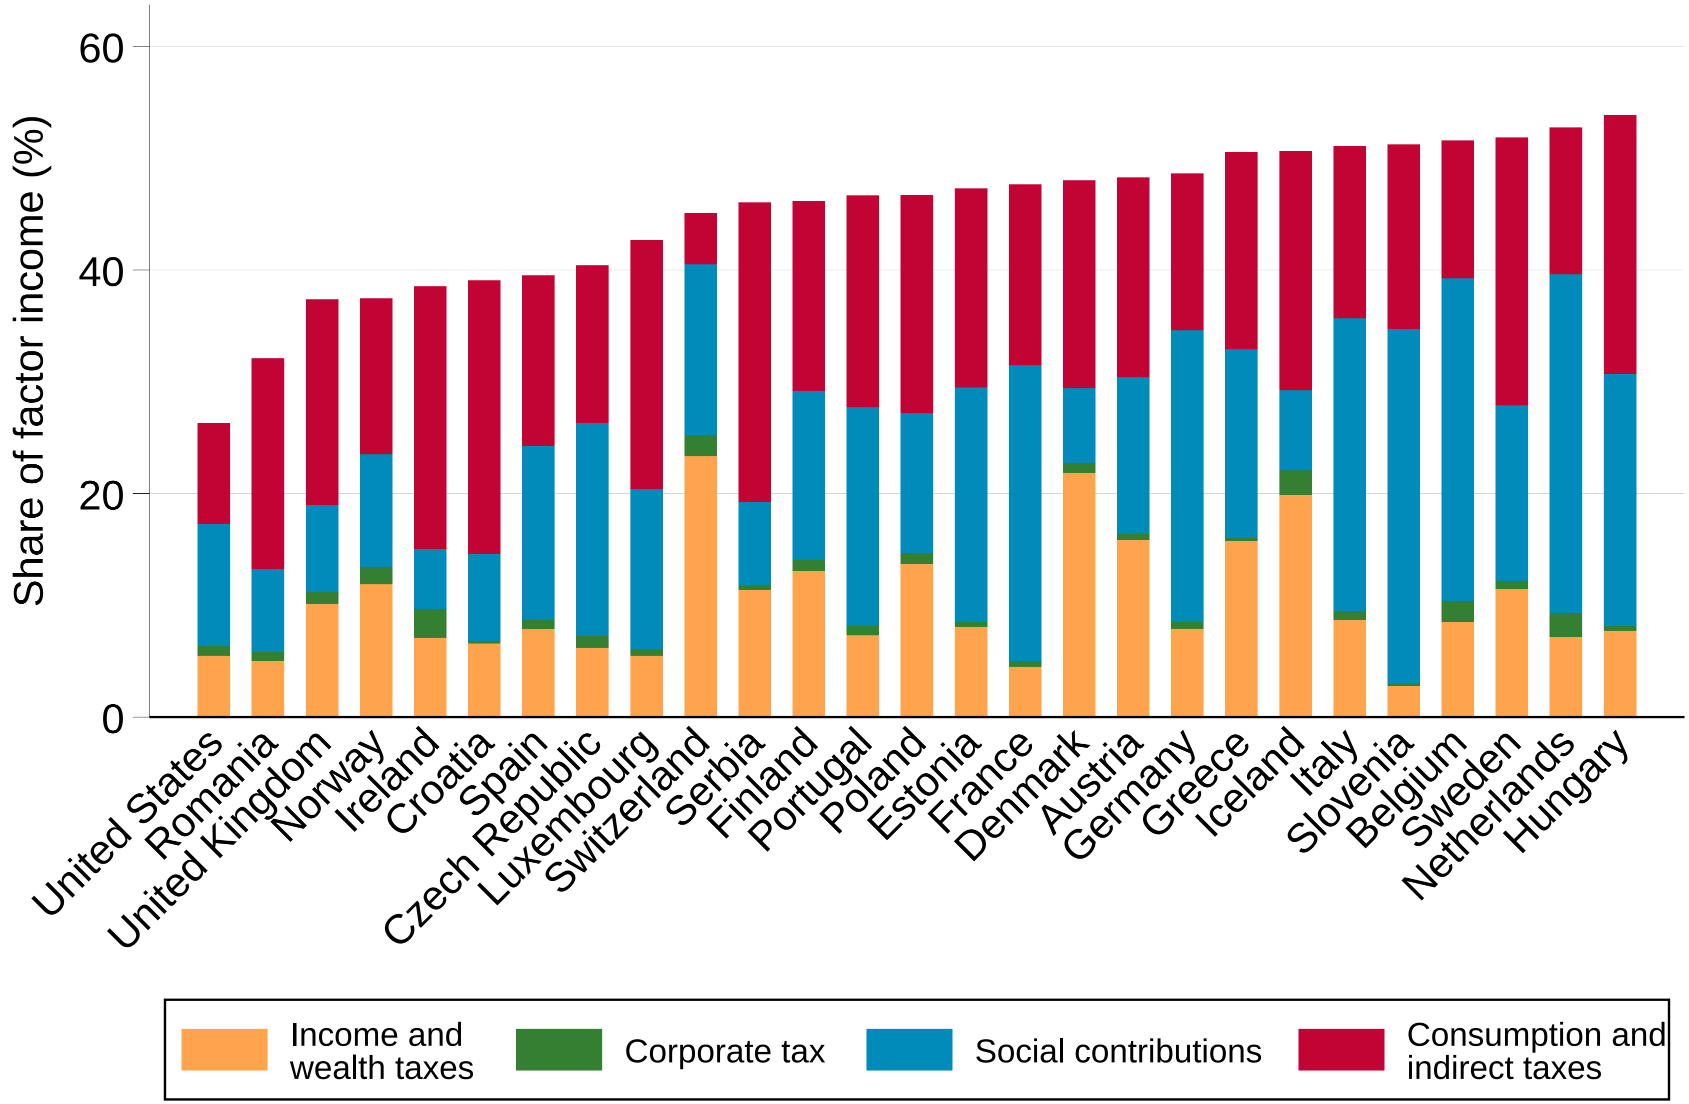

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged between 25 and 59. Income is split equally among spouses. Average over the 2007-2017 period.

Figure D.59: Ratio of top 10% to bottom 50% effective tax rates by country (all taxes, % of factor income, working-age population)

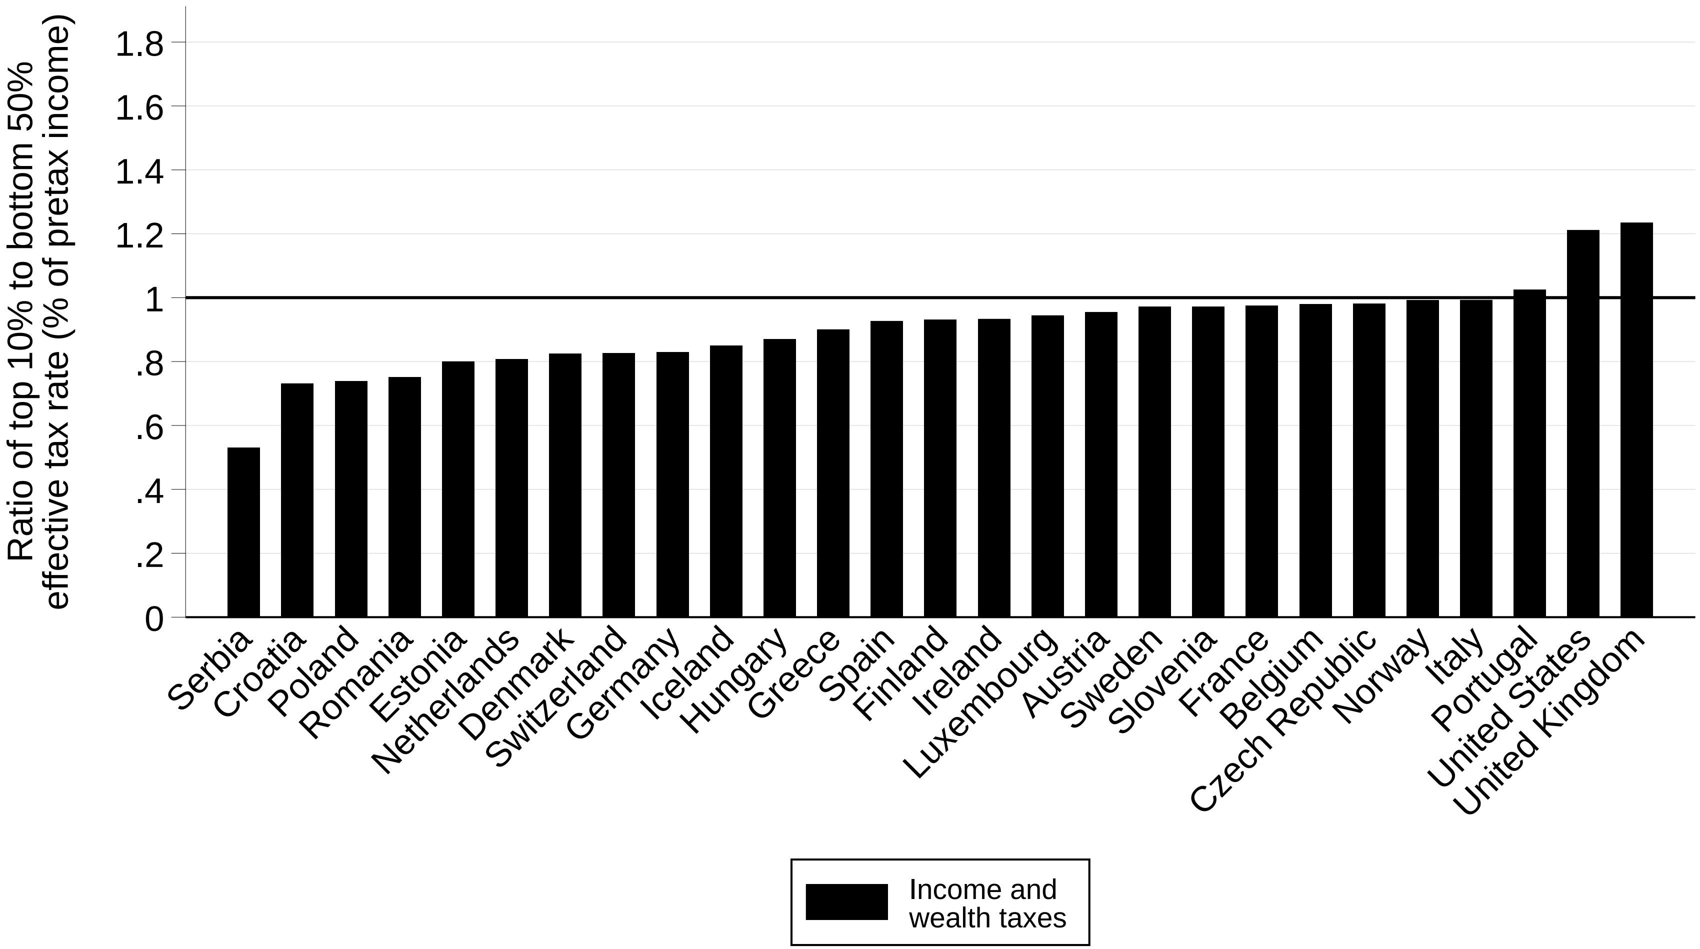

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged between 25 and 59. Income is split equally among spouses. Average over the 2007-2017 period.

D.2.4 Distribution of transfers

Figure D.60: Total transfers received by the bottom 50% by country (% of posttax income)

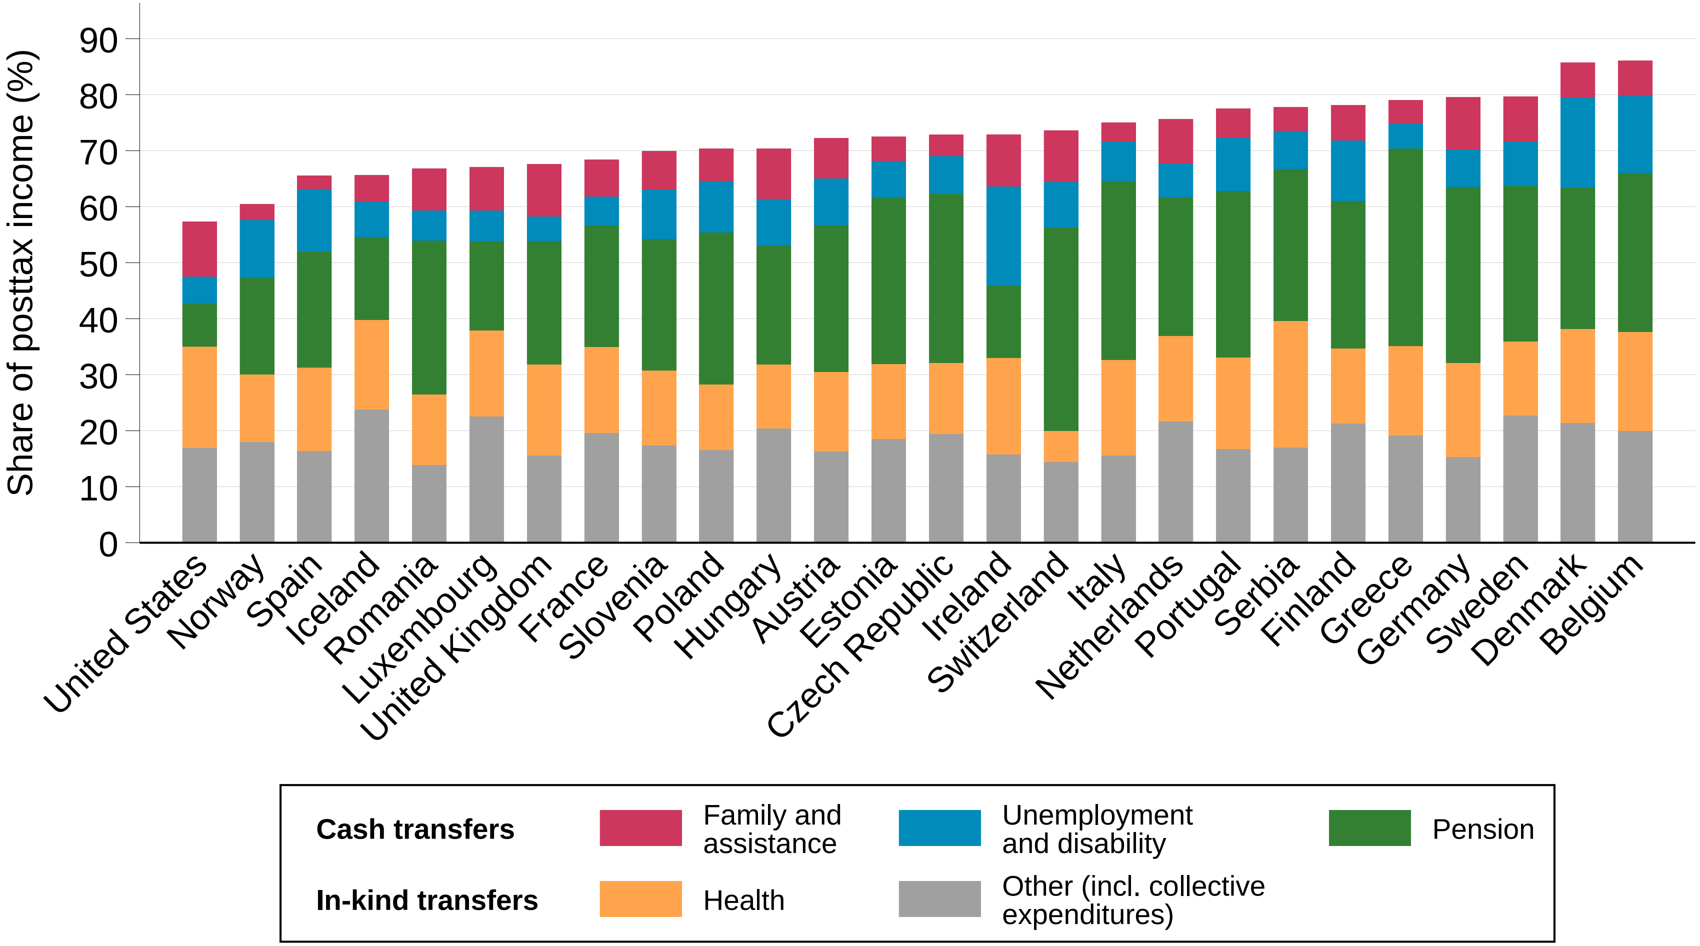

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. Average over the 2007-2017 period.

Figure D.61: Total transfers received by the middle 40% by country (% of posttax income)

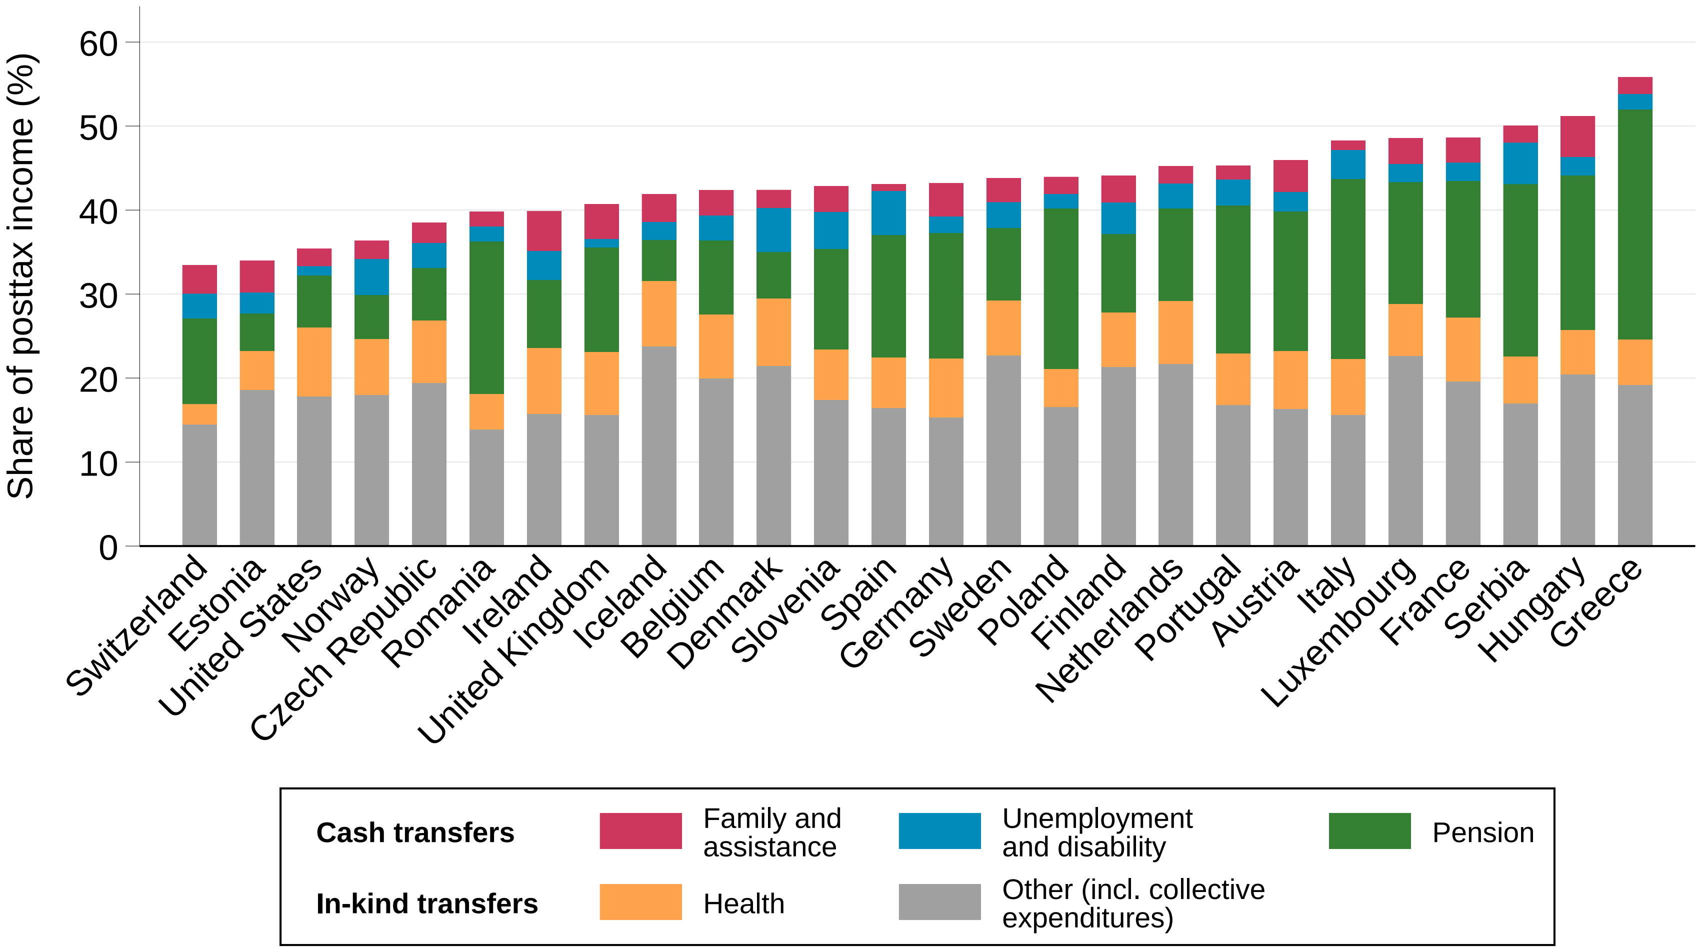

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. Average over the 2007-2017 period.

D.2.5 Distribution of posttax income

Figure D.62: Top 1% and Bottom 50% posttax income shares in Europe and the US

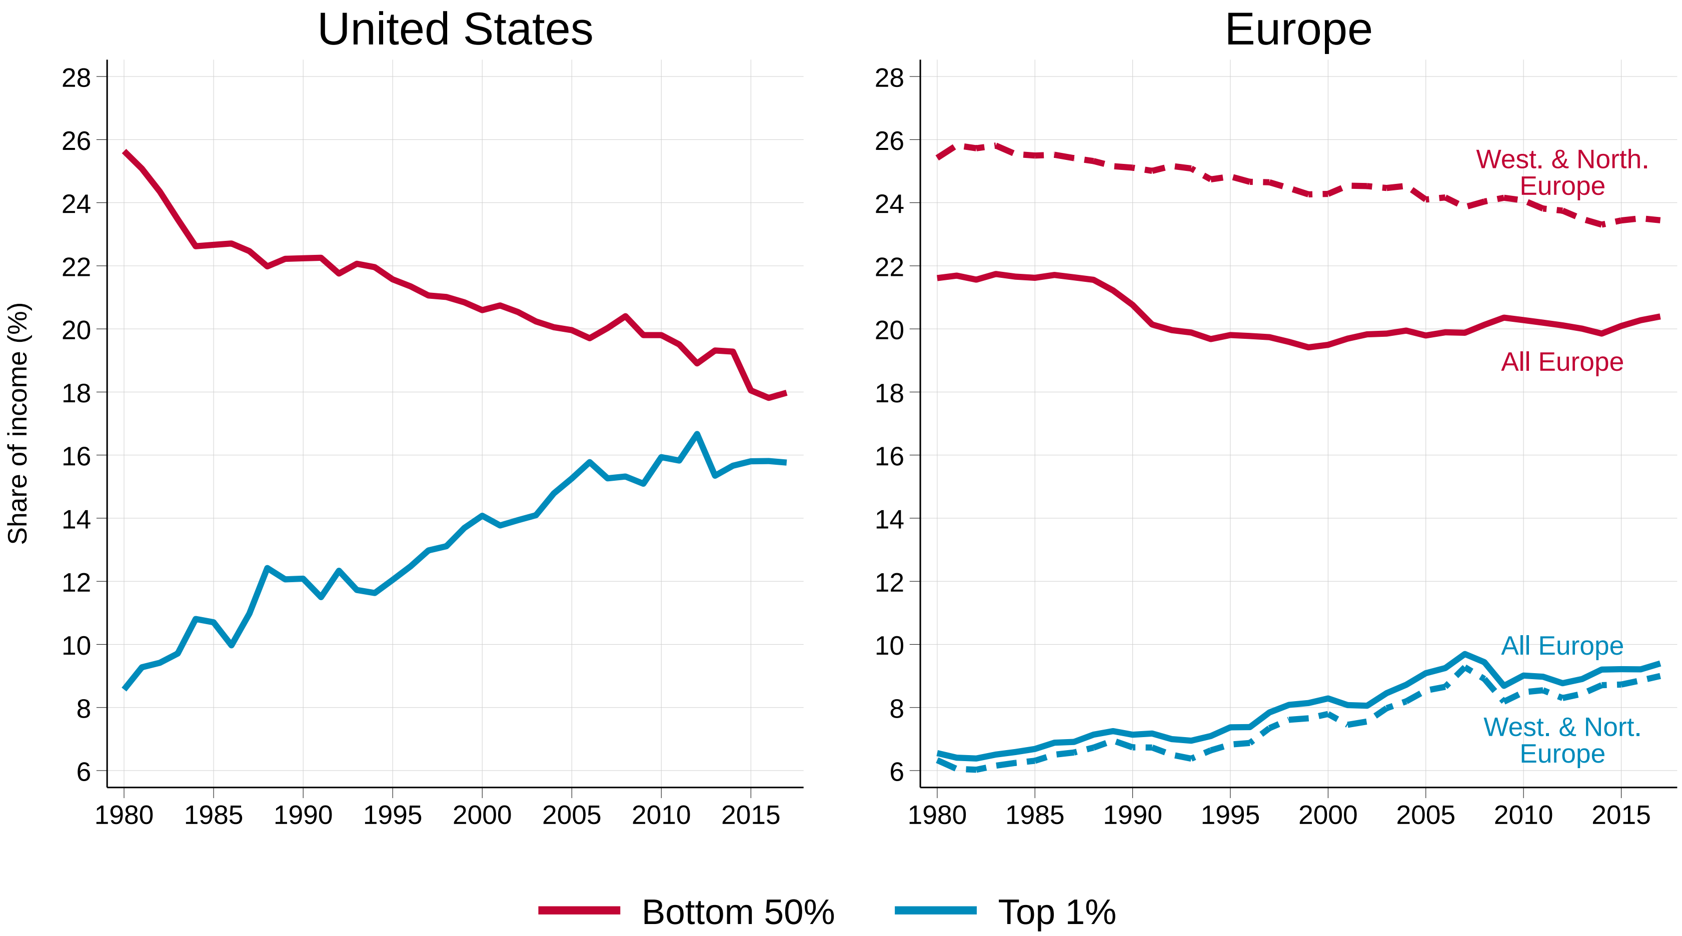

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The figure compares the share of posttax income received by the bottom 50% to that received by the top 1% of the regional population. Figures for the US come from Piketty, Saez, and Zucman (2018). Figures for Europe are aggregated using market exchange rates. The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. See Table D.6 for the composition of European regions.

Figure D.63: Average annual posttax income growth by percentile, 1980-2017

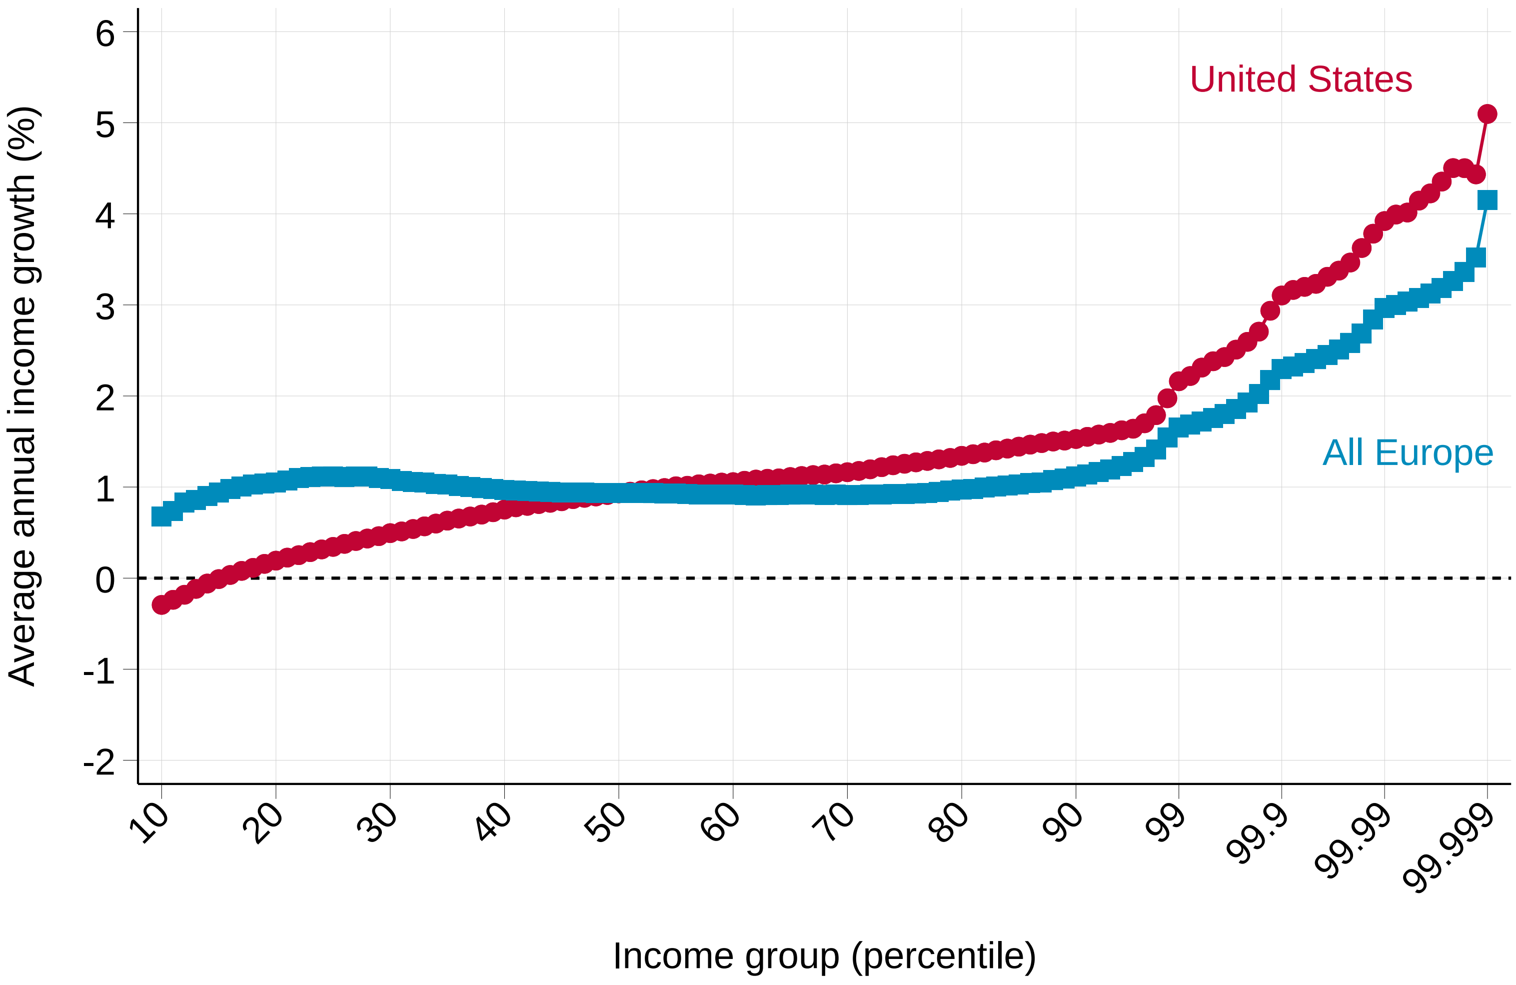

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The figure plots the average annual posttax income growth rate by percentile, with a further decomposition of the top percentile. Figures for the US come from Piketty, Saez, and Zucman (2018). Figures for Europe are aggregated using market exchange rates. The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. See Table D.6 for the composition of European regions.

Figure D.64: Bottom 50% incomes in Europe, 1980-2017

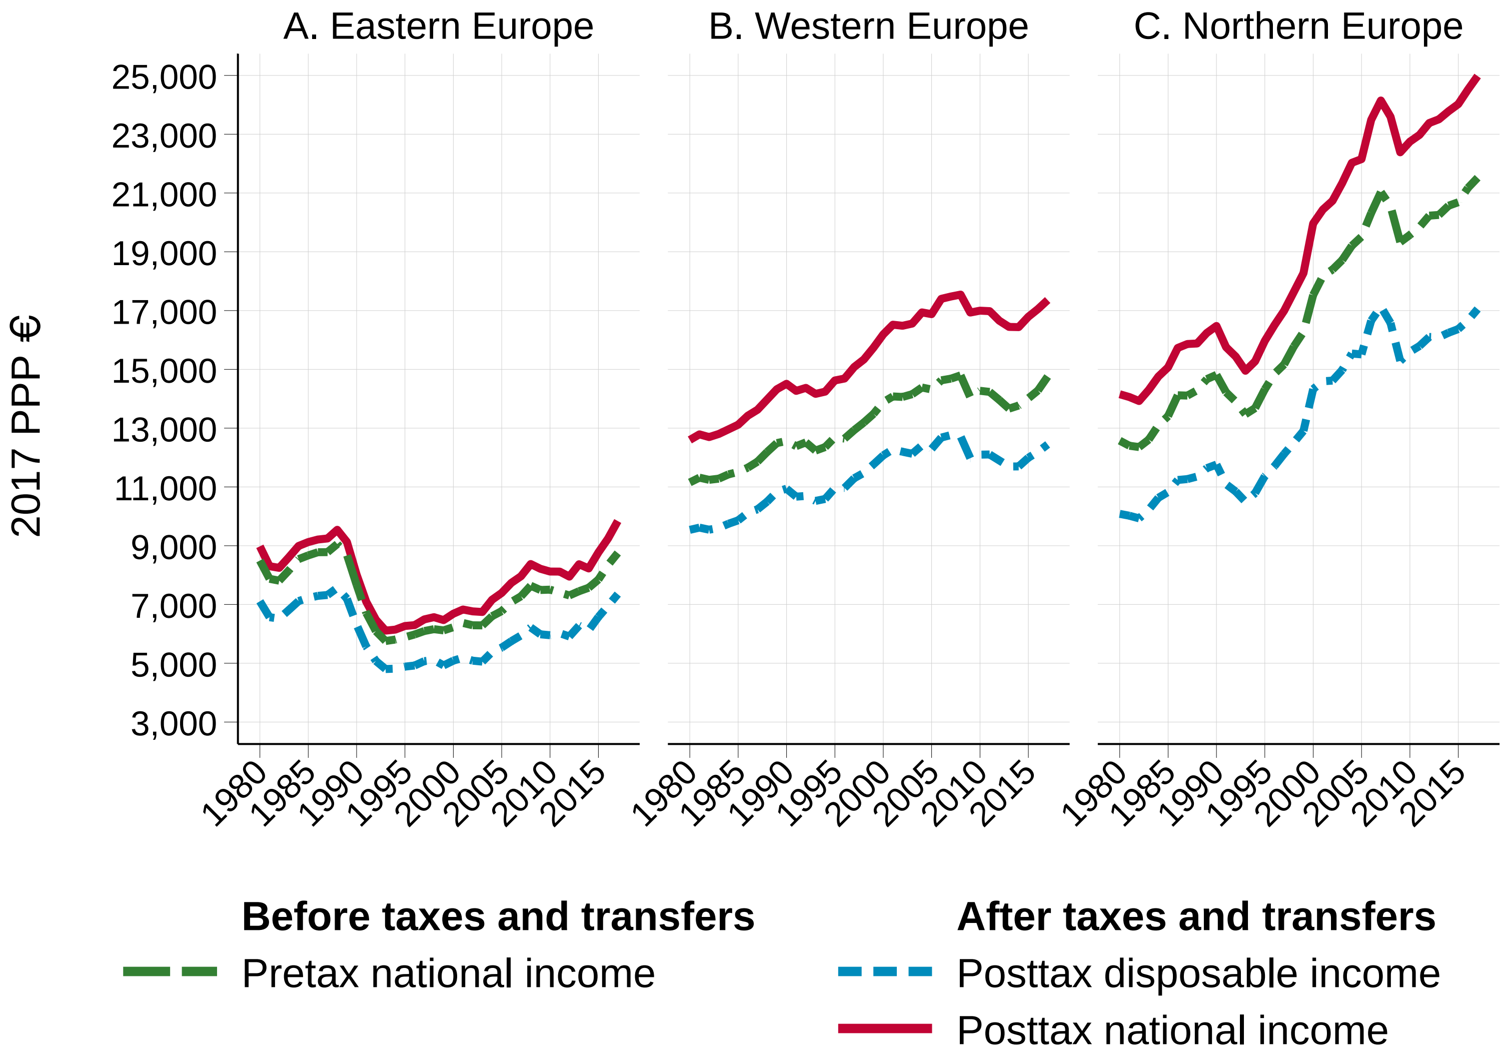

Source: Authors' computations combining surveys, tax data and national accounts for European countries and Piketty, Saez, and Zucman, 2018 for the US.
Notes: Incomes are measured at Purchasing Power Parity in real 2017 Euros. PPP Euro 1 = PPP dollar 1.3. The unit of observation is the adult individual aged 20. See Table D.6 for the composition of European regions.

Figure D.65: Bottom 50% and Top 10% real incomes in Europe and the US, 1980-2017

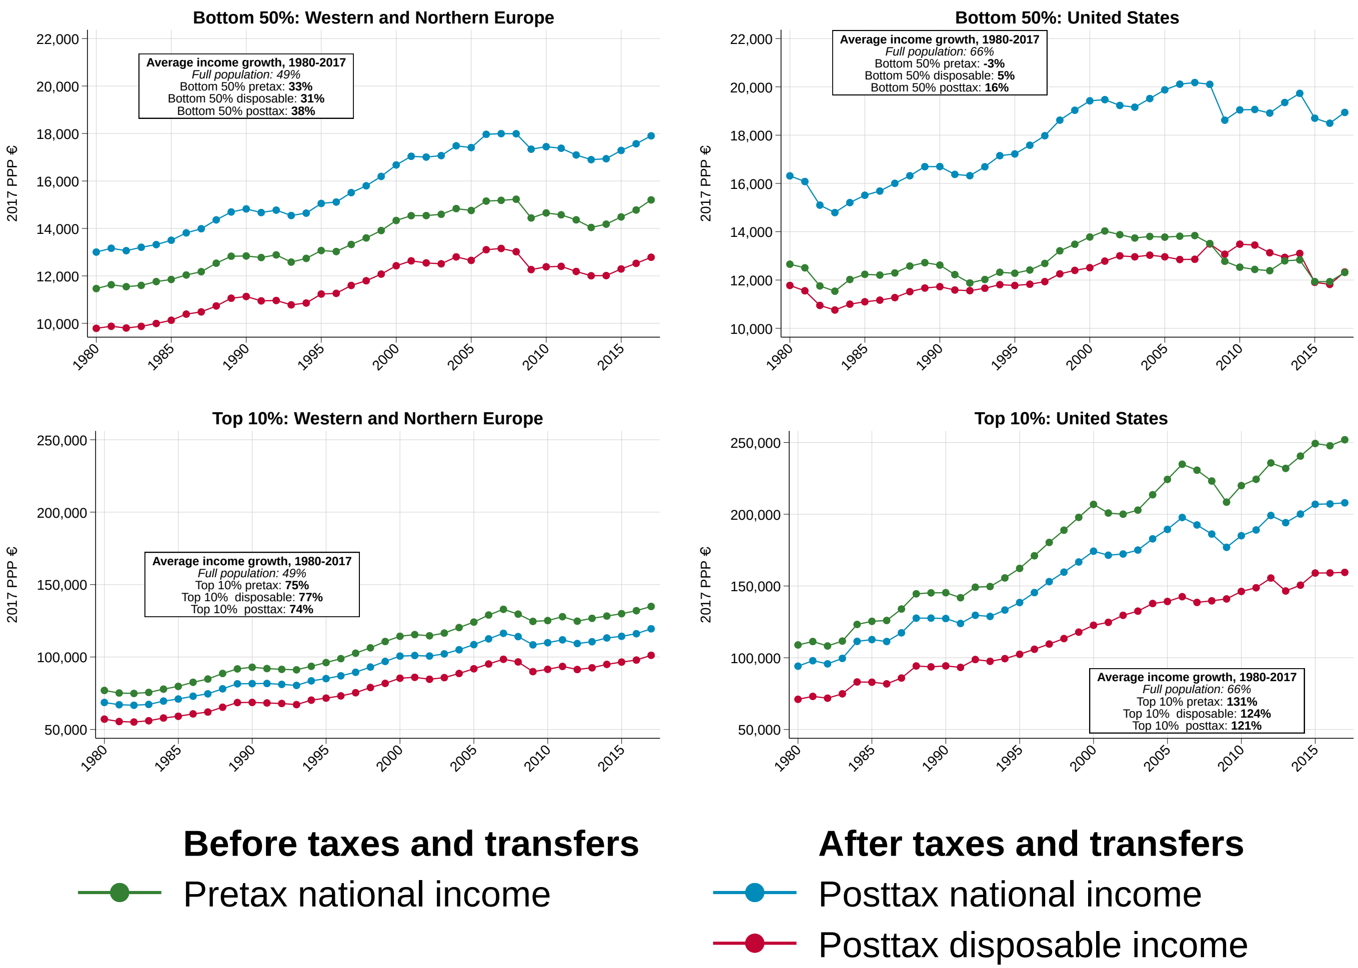

Source: Authors' computations combining surveys, tax data and national accounts for European countries and Piketty, Saez, and Zucman, 2018 for the US. *Notes:* Incomes are measured at Purchasing Power Parity in real 2017 Euros. PPP €1 = PPP dollar 1.3. The unit of observation is the adult individual aged 20. See Table D.6 for the composition of European regions.

Figure D.66: Middle 40% and Bottom 20% incomes in Europe and the US, 1980-2017

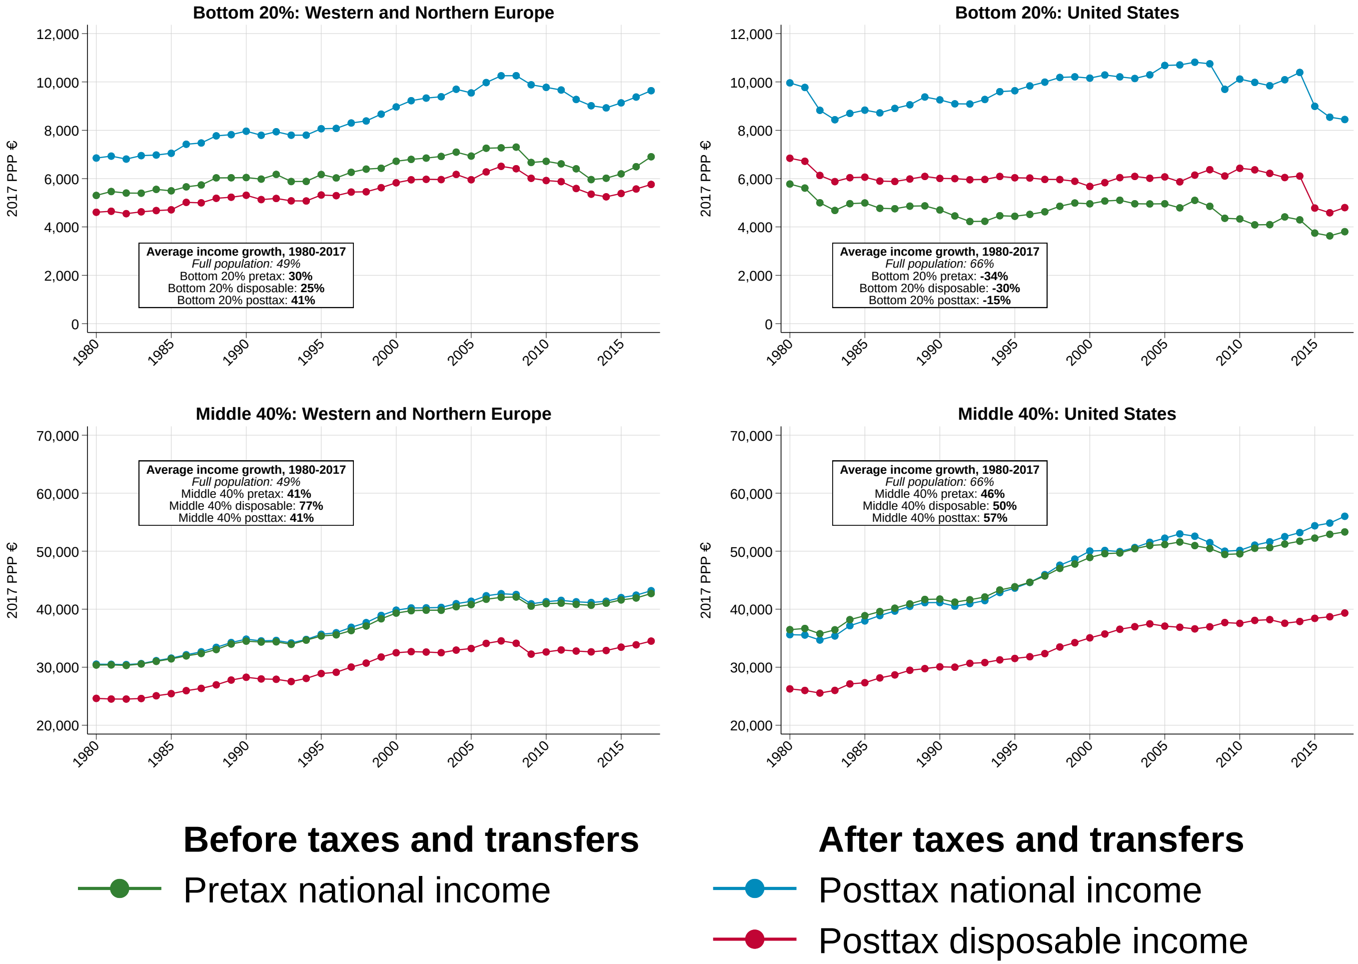

Source: Authors' computations combining surveys, tax data and national accounts for European countries and Piketty, Saez, and Zucman, 2018 for the US. *Notes:* Incomes are measured at Purchasing Power Parity in real 2017 Euros. PPP Euro 1 = PPP dollar 1.3. The unit of observation is the adult individual aged 20. See Table D.6 for the composition of European regions.

Figure D.67: Redistribution in Europe and the United States, 1980-2017:
Ratio of top 10% to bottom 50% average incomes

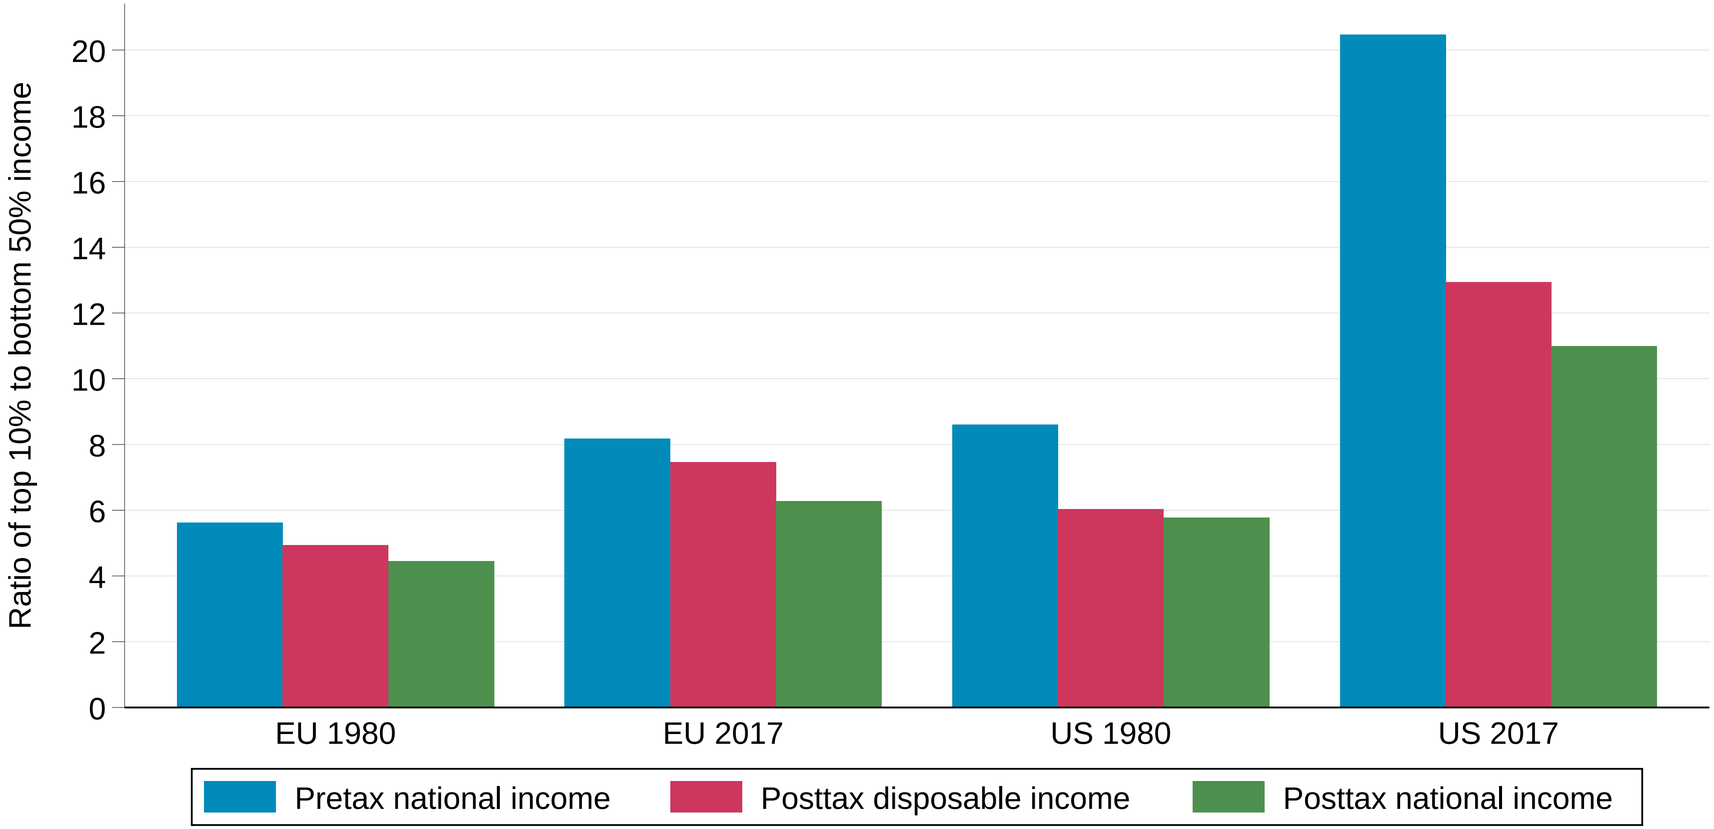

Source: Authors' computations combining surveys, tax data and national accounts for European countries and Piketty, Saez, and Zucman, 2018 for the US.
Notes: The unit of observation is the adult individual aged 20. Indicators are population weighted. European inequality estimates contain all Western, Northern and Eastern European countries. See Appendix Table D.6 for the composition of European regions.

Figure D.68: Net redistribution in Europe and the US (% of pretax income)

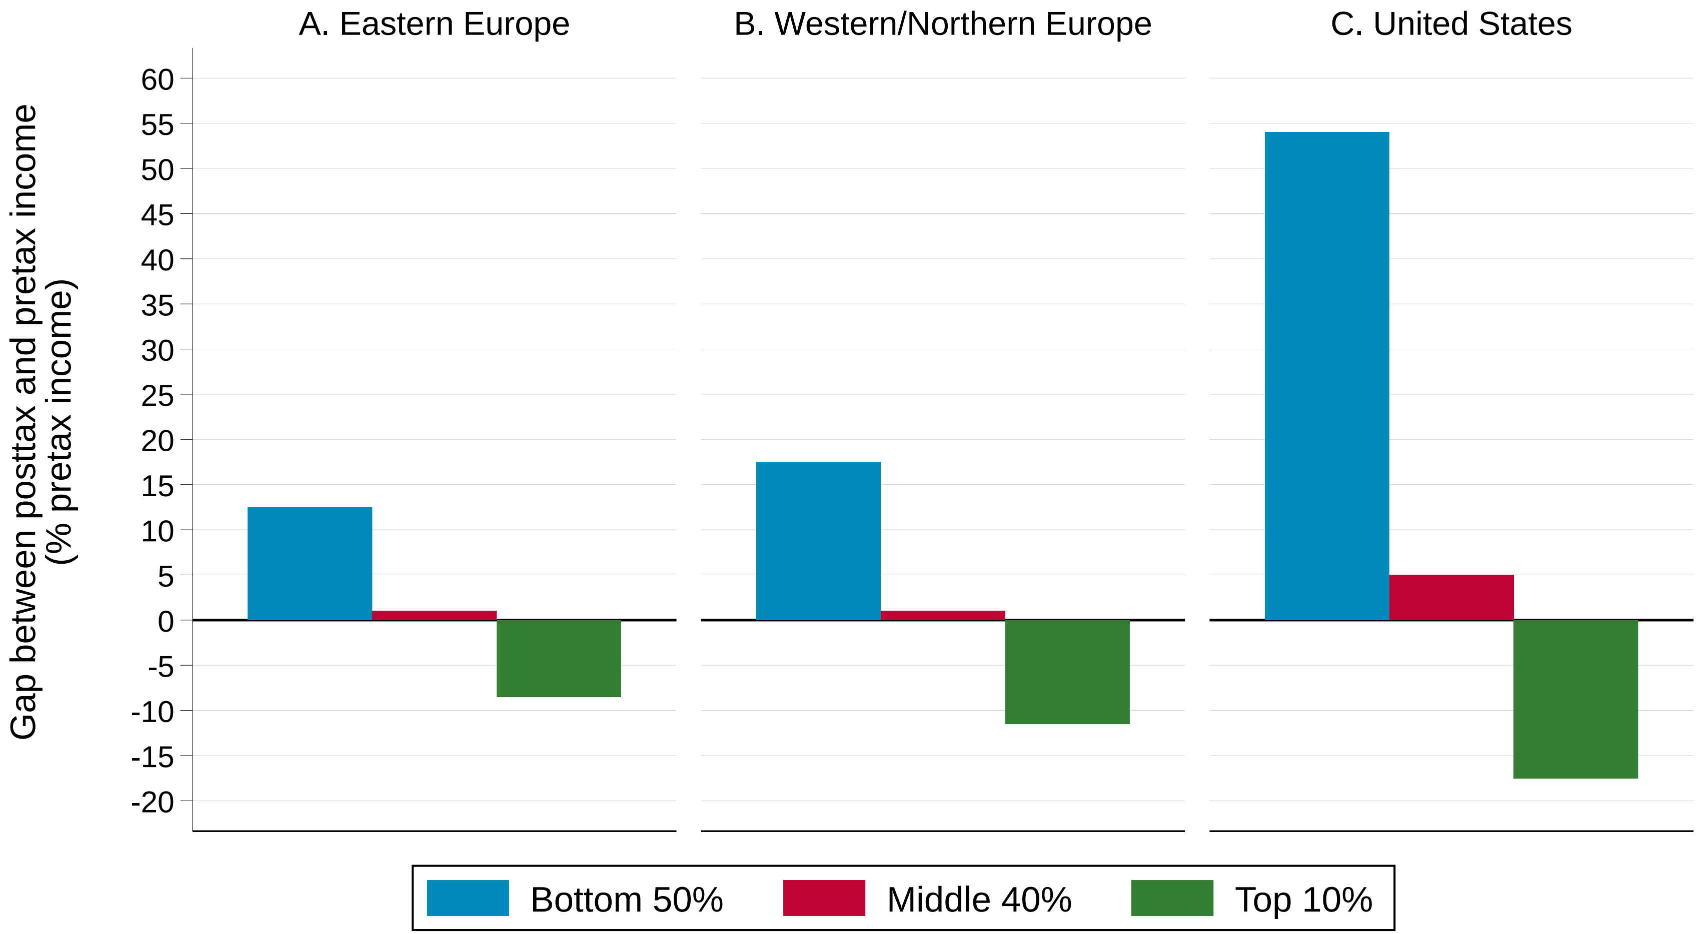

Source: Authors' computations combining surveys, tax data and national accounts for European countries and Piketty, Saez, and Zucman, 2018 for the US.
Notes: The unit of observation is the adult individual aged 20. Indicators are population weighted. European inequality estimates contain all Western, Northern and Eastern European countries.

Figure D.69: Net redistribution in Europe and the US (decomposing the bottom 50%)

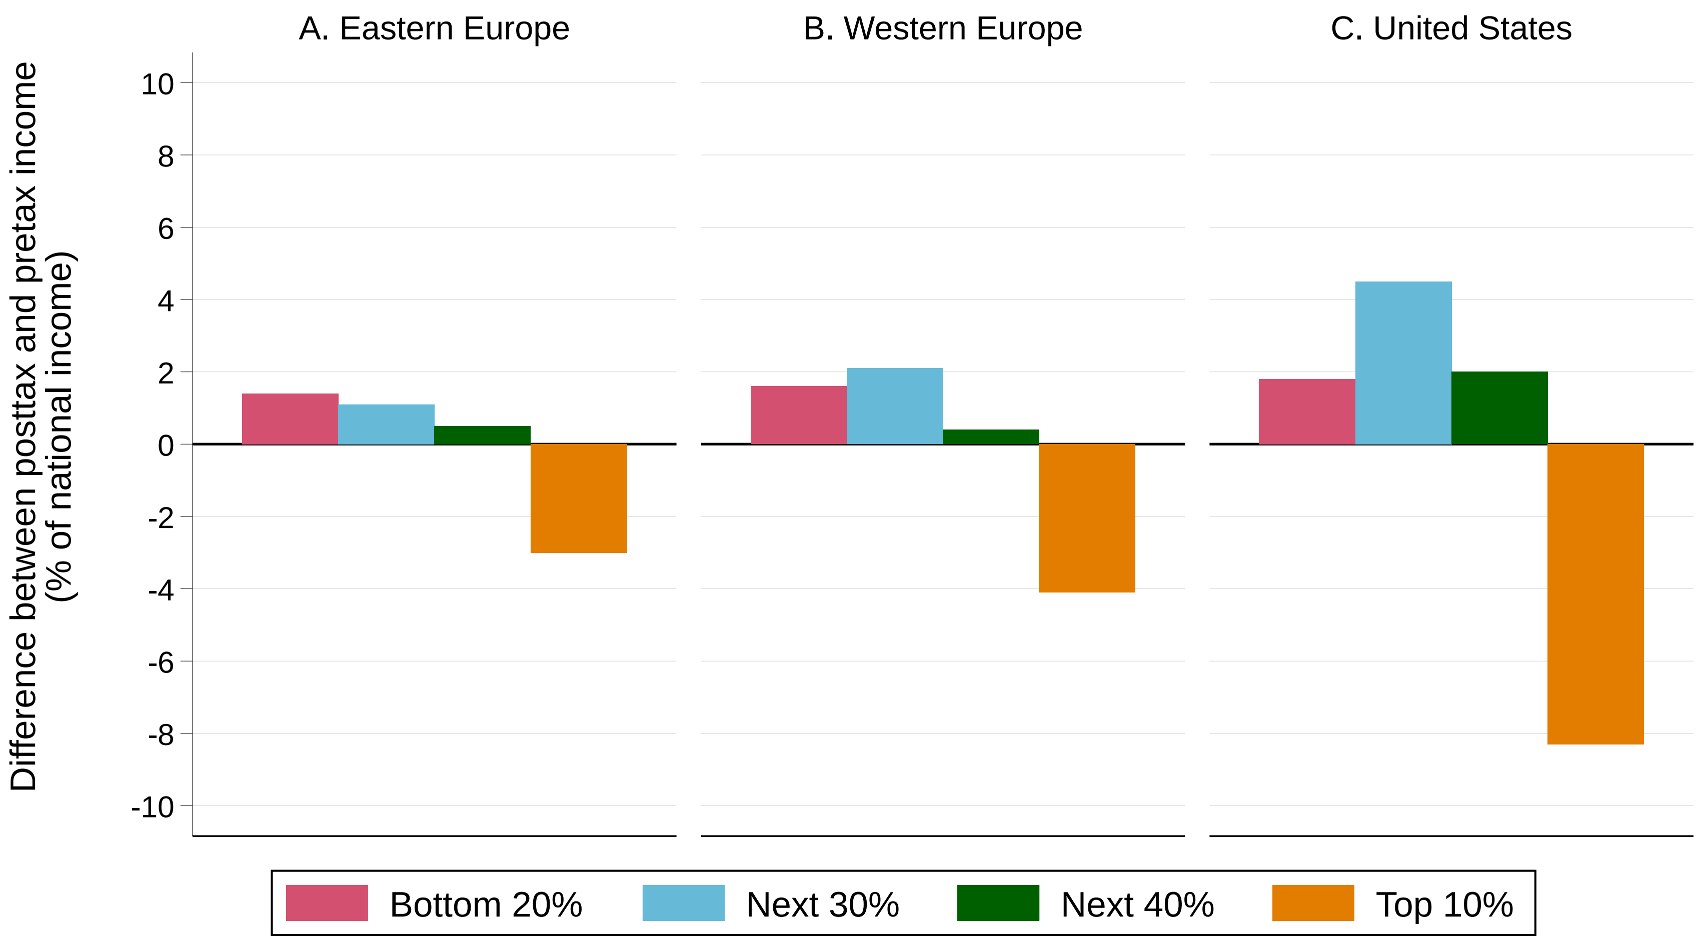

Source: Authors' computations combining surveys, tax data and national accounts for European countries and Piketty, Saez, and Zucman, 2018 for the US.
Notes: The unit of observation is the adult individual aged 20. Indicators are population weighted. European inequality estimates contain all Western, Northern and Eastern European countries.

Figure D.70: Top 10% and bottom 50% posttax income shares in Europe and the United States: lump-sum vs. proportional allocation of collective expenditure

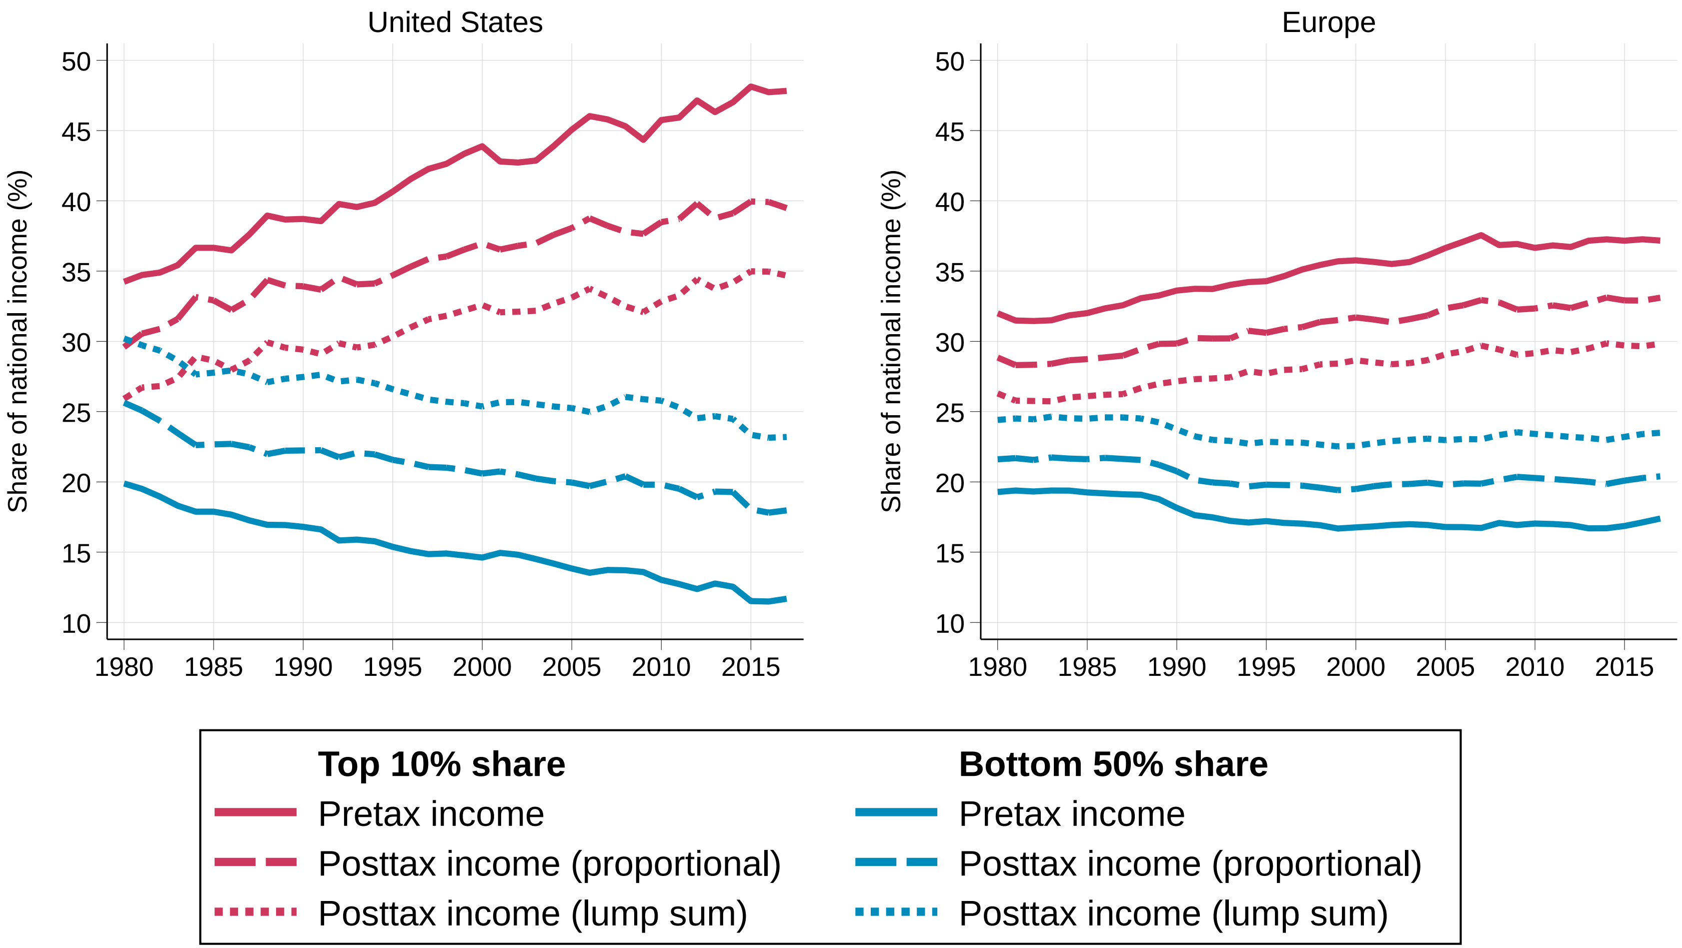

Source: Authors' computations combining surveys, tax data and national accounts for European countries and Piketty, Saez, and Zucman, 2018 for the US.
Notes: The figure represents the top 10% and bottom 50% shares in Europe and the United States in terms of pretax income, posttax national income assuming that all non-health collective government expenditure is distributed proportionally to posttax disposable income, and posttax national income assuming that all non-health collective government expenditure is distributed on a lump sum basis. The unit of observation is the adult individual aged 20. Income is split equally among spouses. See Table D.6 for the composition of European regions.

Figure D.71: Net redistribution (% of group average income): lump-sum vs. proportional allocation of collective expenditures

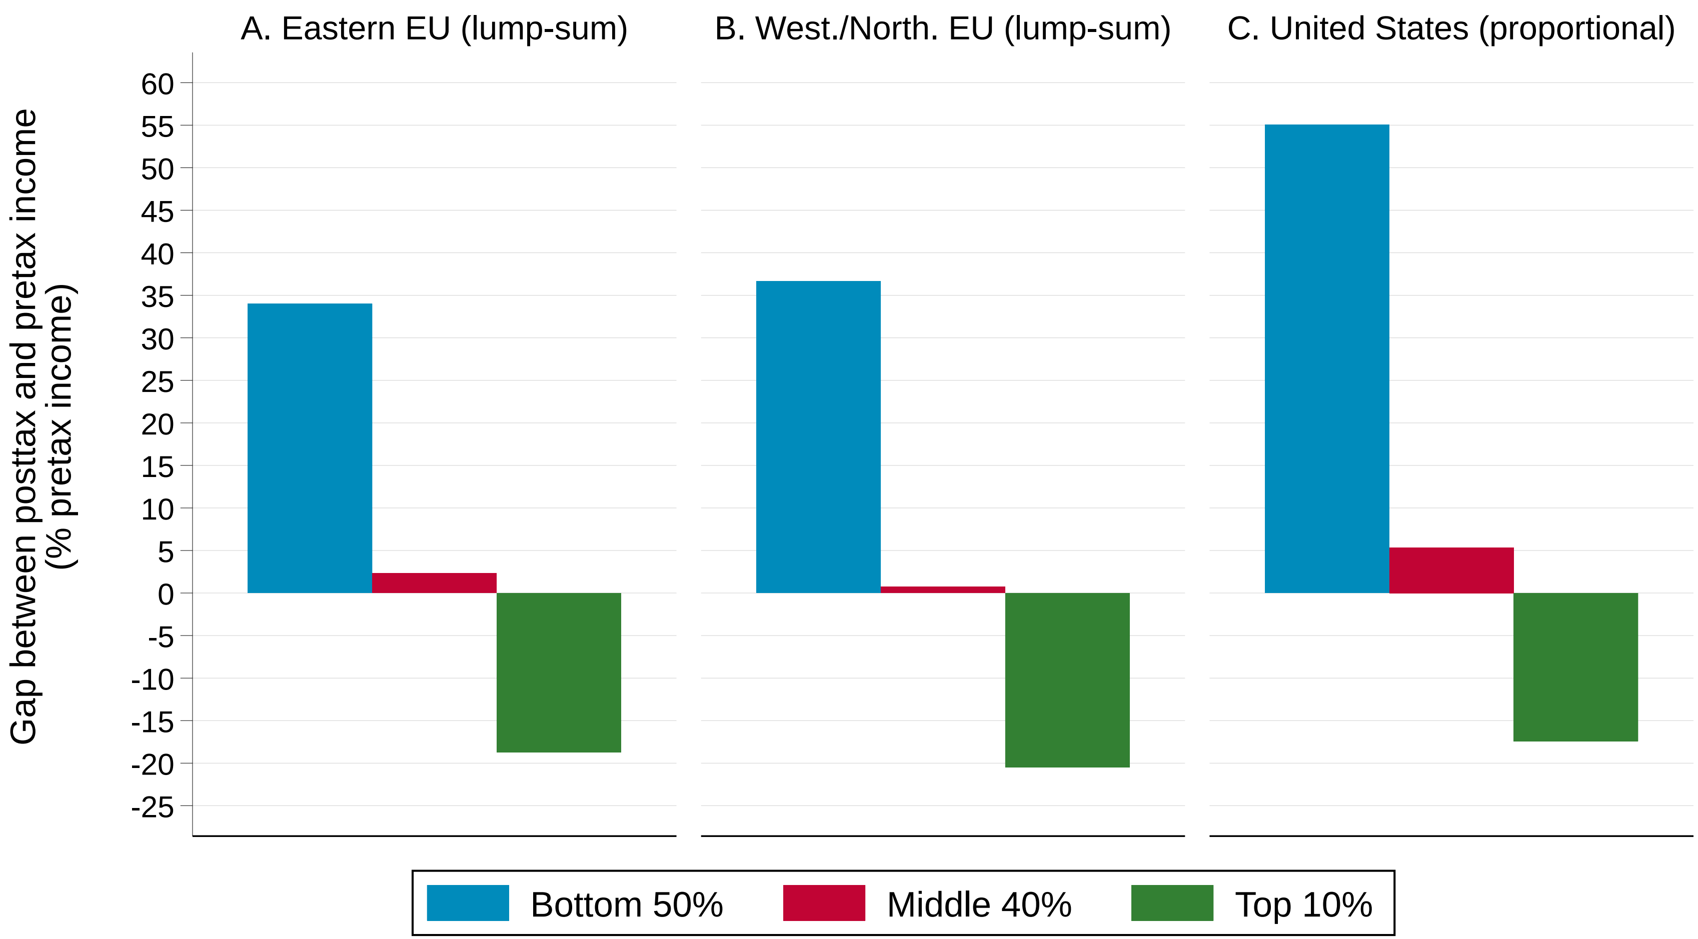

Source: Authors' computations combining surveys, tax data and national accounts for European countries and Piketty, Saez, and Zucman, 2018 for the US.
Notes: The figure represents the net transfer operated between pretax income groups, expressed as a share of national income, assuming that all non-health collective expenditures are allocated on a lump-sum basis in Europe, and proportionally to income in the United States. The unit of observation is the adult individual aged 20. Income is split equally among spouses. See Table D.6 for the composition of European regions.

Figure D.72: Net redistribution (% of national income): lump-sum vs. proportional allocation of collective expenditures

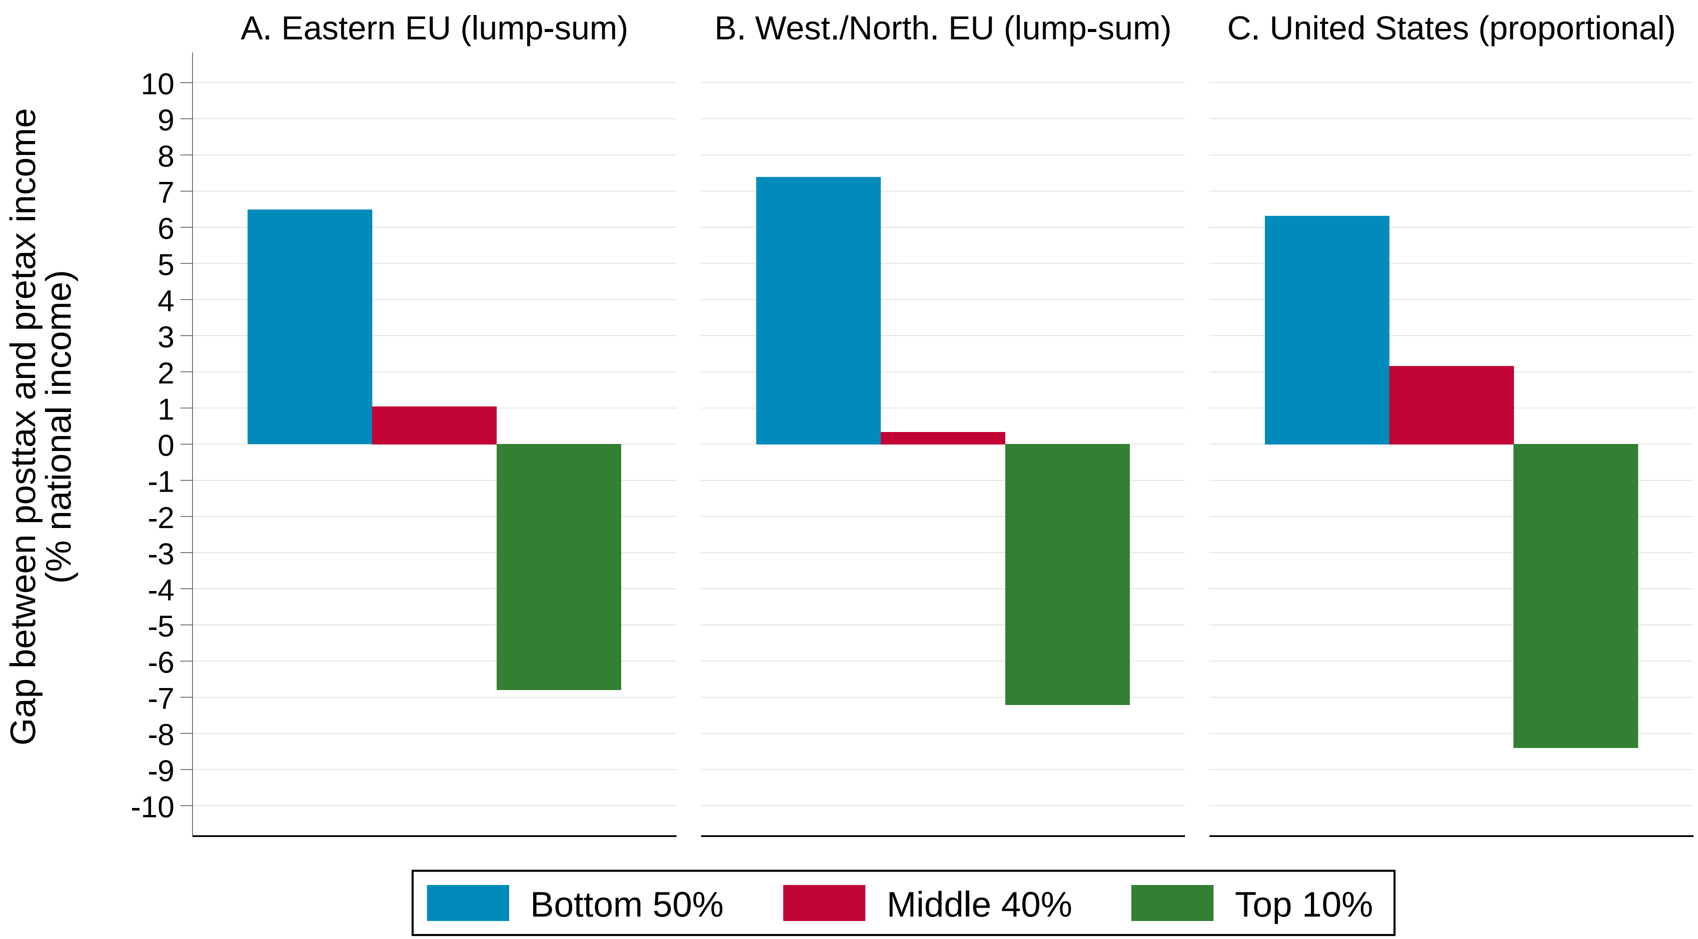

Source: Authors' computations combining surveys, tax data and national accounts for European countries and Piketty, Saez, and Zucman, 2018 for the US.
Notes: The figure represents the net transfer operated between pretax income groups, expressed as a share of national income, assuming that all non-health collective expenditures are allocated on a lump-sum basis in Europe, and proportionally to income in the United States. The unit of observation is the adult individual aged 20. Income is split equally among spouses. See Table D.6 for the composition of European regions.

Figure D.73: Bottom 50% factor income share, working-age population, Europe vs. US, 2007-2015

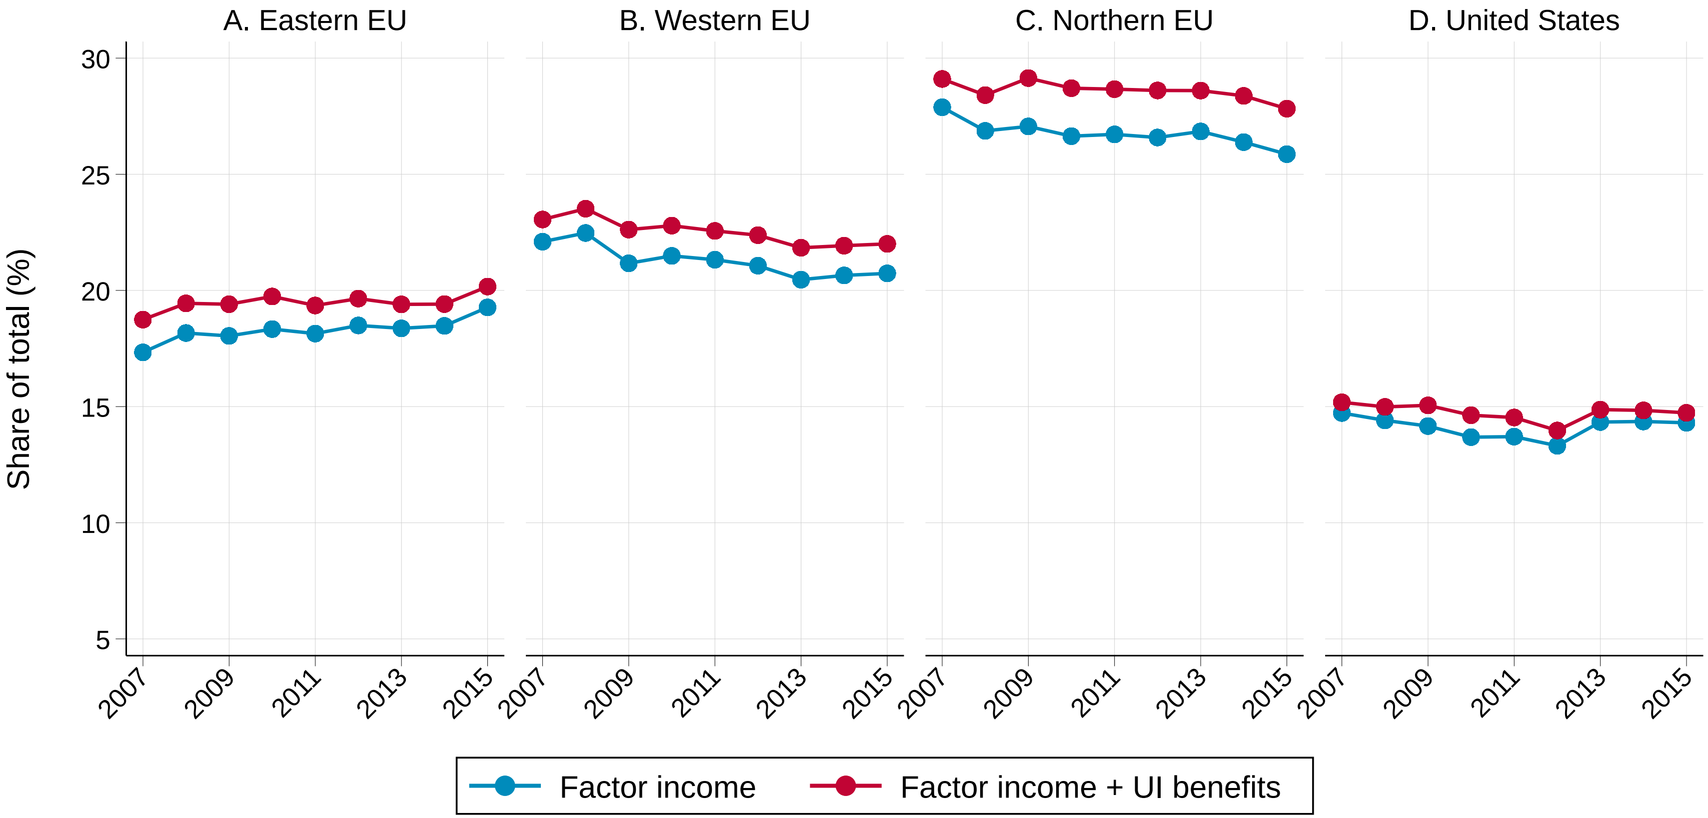

Source: Authors' computations combining surveys, tax data and national accounts for European countries and Piketty, Saez, and Zucman, 2018 for the US. *Notes:* Distribution of factor income among the working age population. The unit of observation is the adult individual aged 25-59 in European countries and 20-64 in the US. Available microdata does not allow for a detailed decomposition of factor income and UI benefits in Europe before 2007, see methodology section.

Figure D.74: Net transfer received by the bottom 50% by country (% of national income)

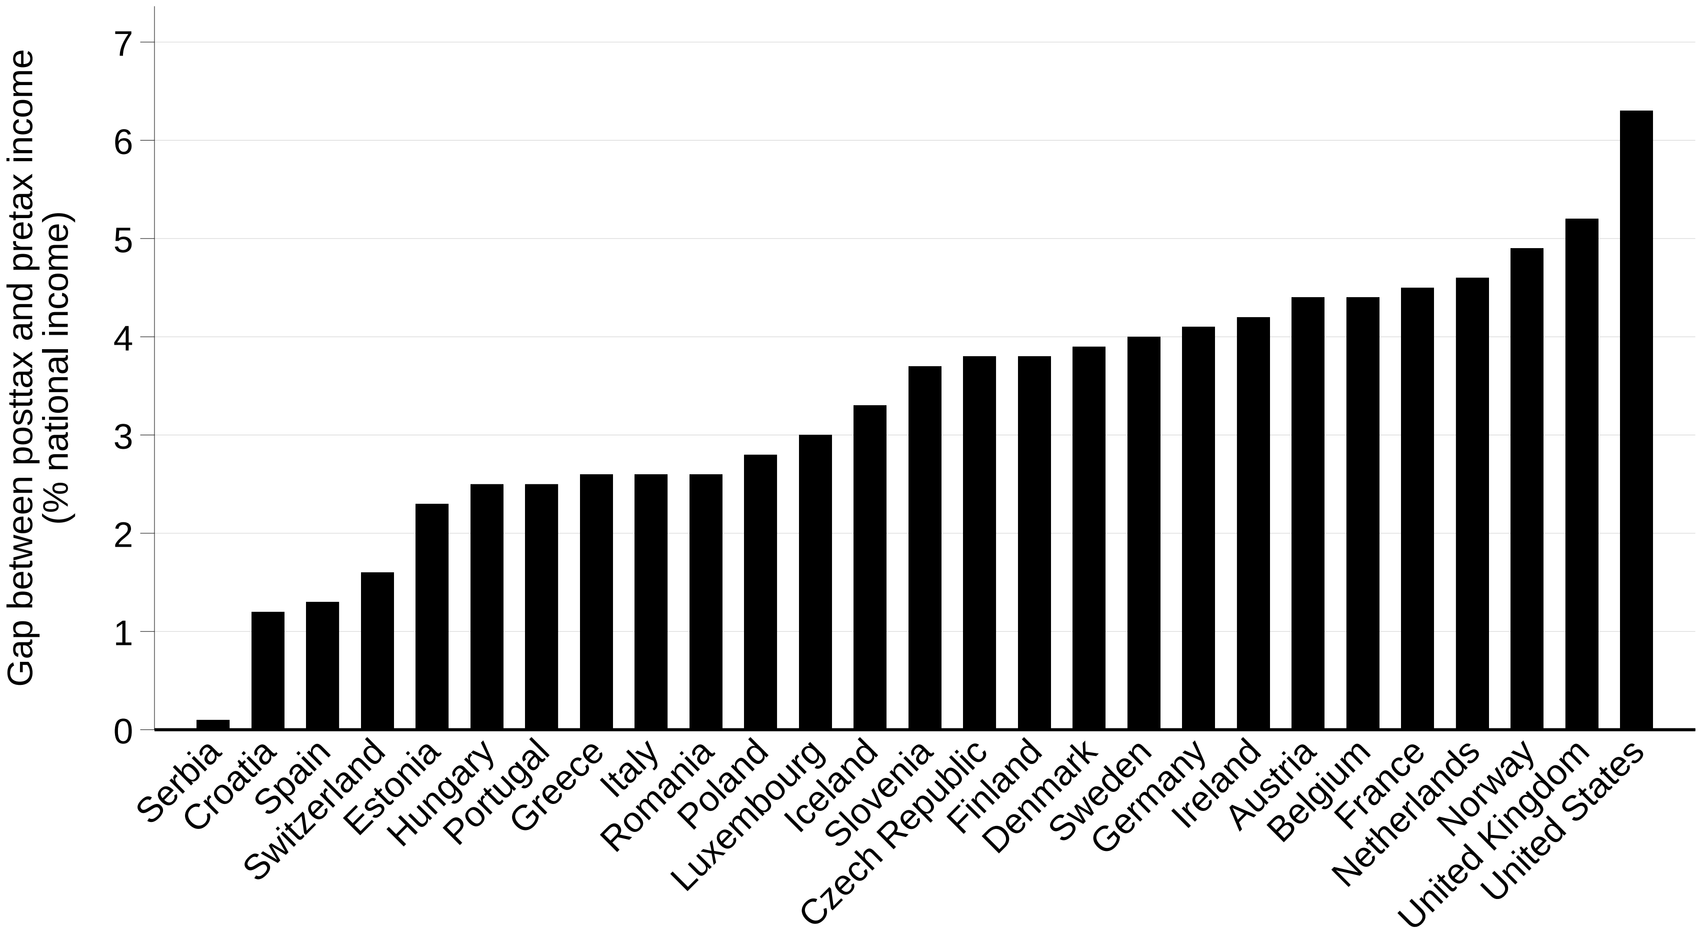

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. Results reported for the year 2017. Non-health collective government expenditures are assumed to be distributed proportionally to posttax disposable income.

Figure D.75: Net transfer received by the middle 40% by country (% of national income)

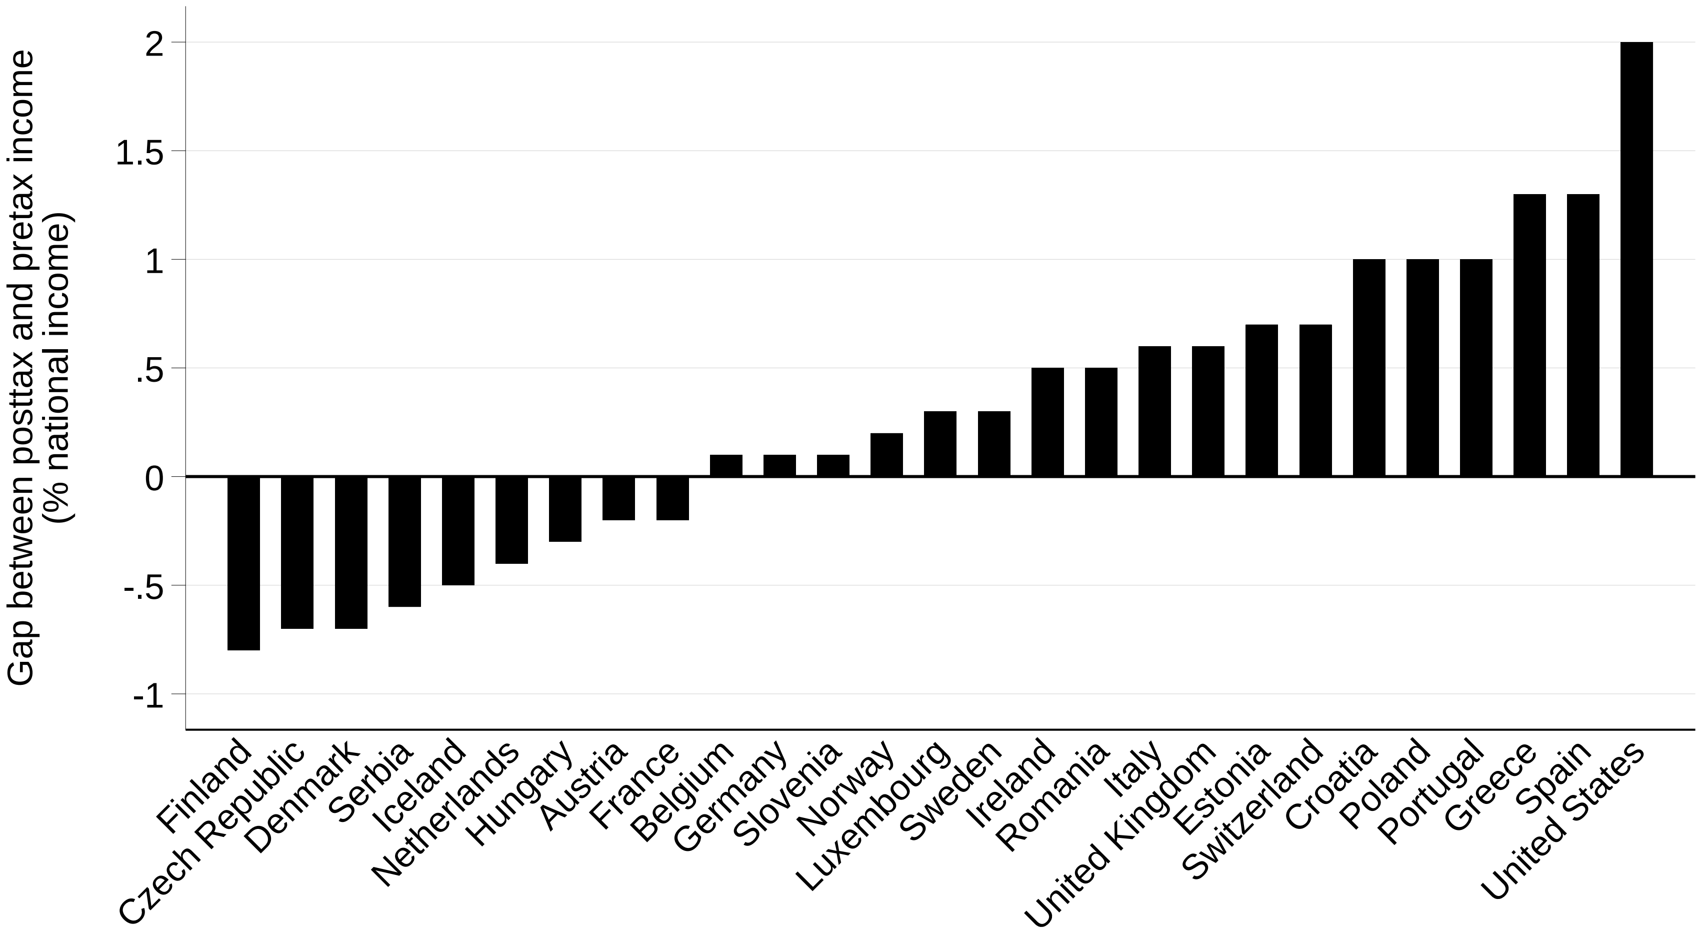

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. Results reported for the year 2017. Non-health collective government expenditures are assumed to be distributed proportionally to posttax disposable income.

Figure D.76: Net transfer received by the top 10% by country (% of national income)

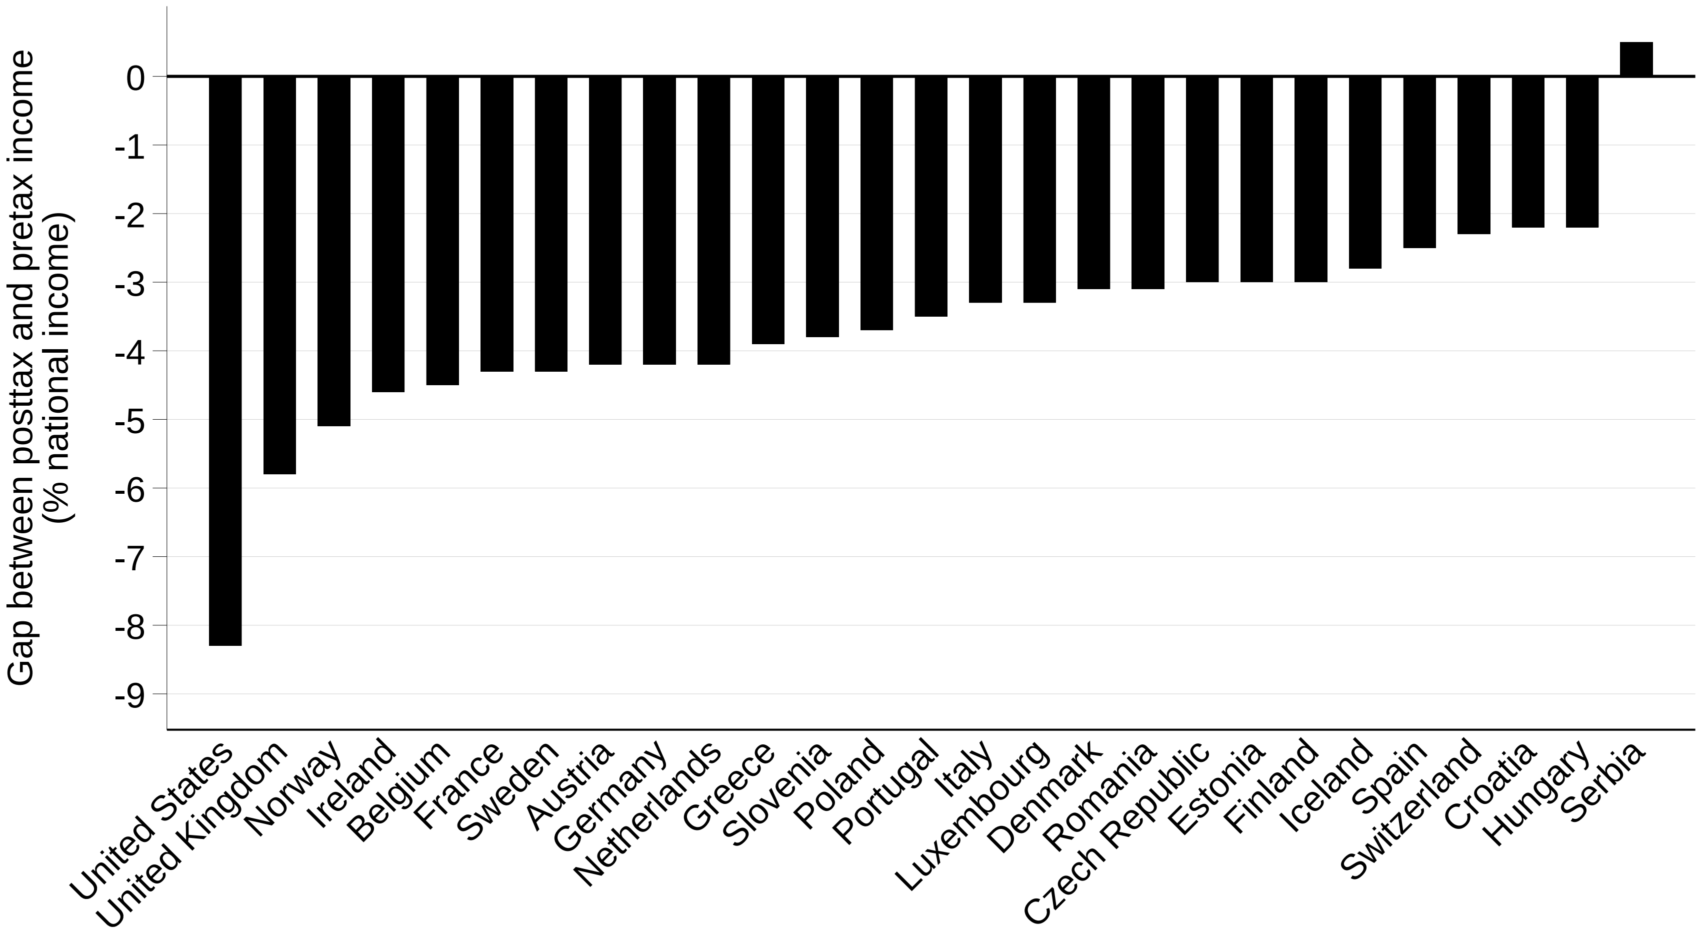

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. Results reported for the year 2017. Non-health collective government expenditures are assumed to be distributed proportionally to posttax disposable income.

Figure D.77: Net transfer received by the bottom 50% by country (% of pretax income)

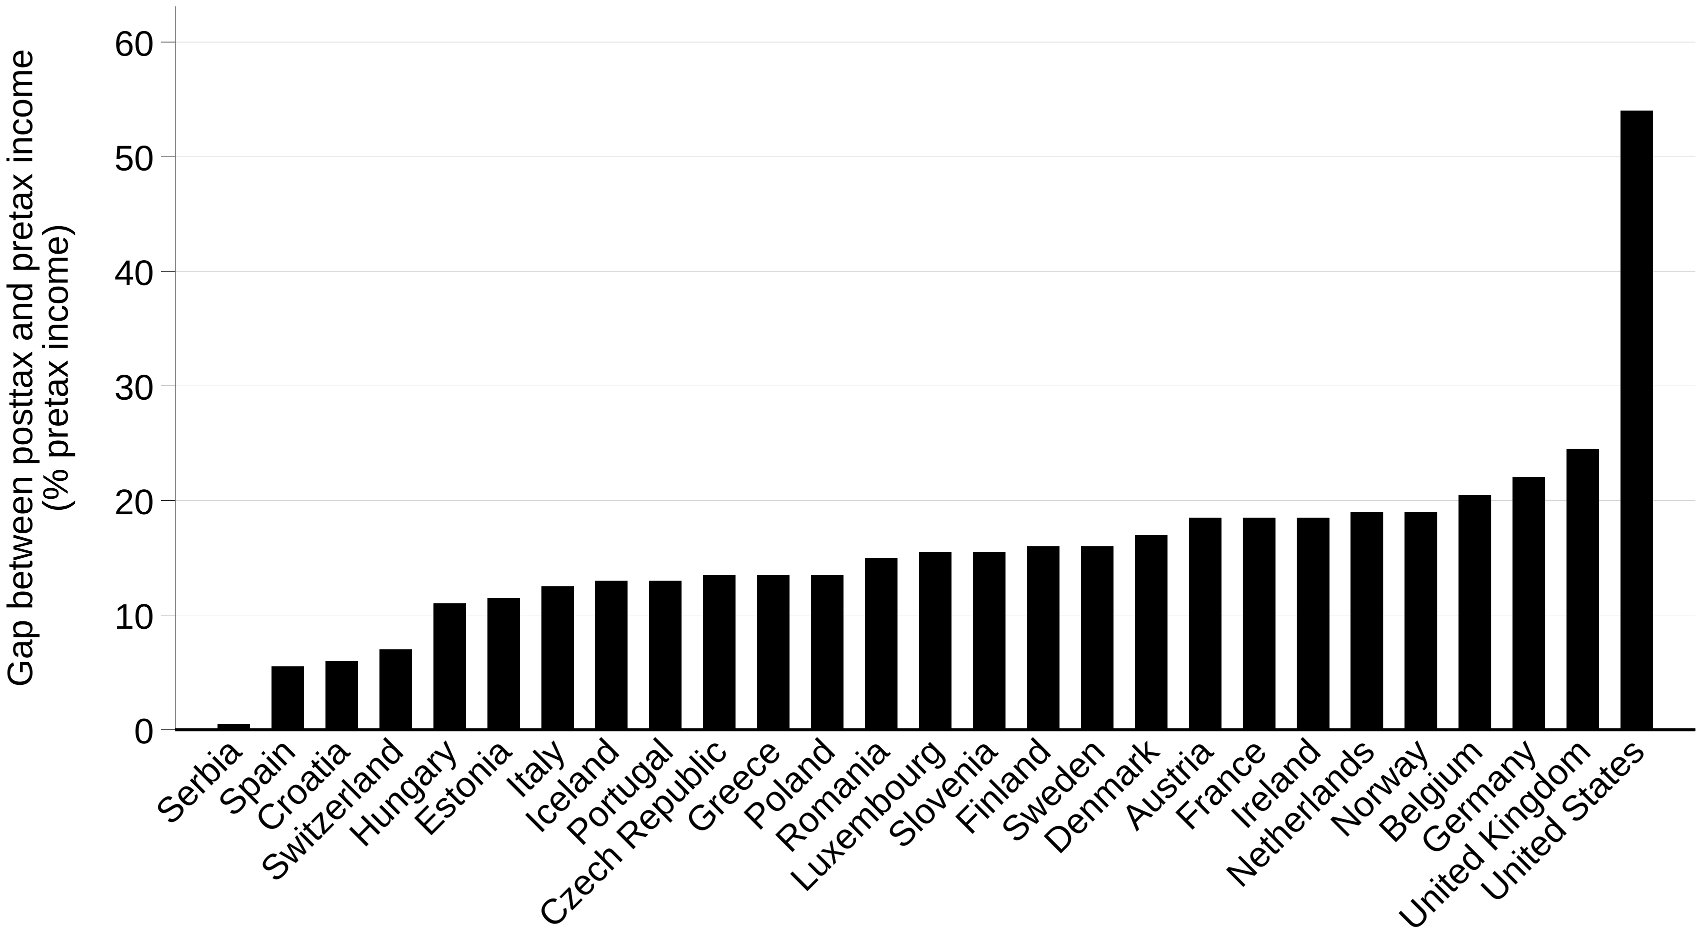

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. Results reported for the year 2017. Non-health collective government expenditures are assumed to be distributed proportionally to posttax disposable income.

Figure D.78: Net transfer received by the middle 40% by country (% of pretax income)

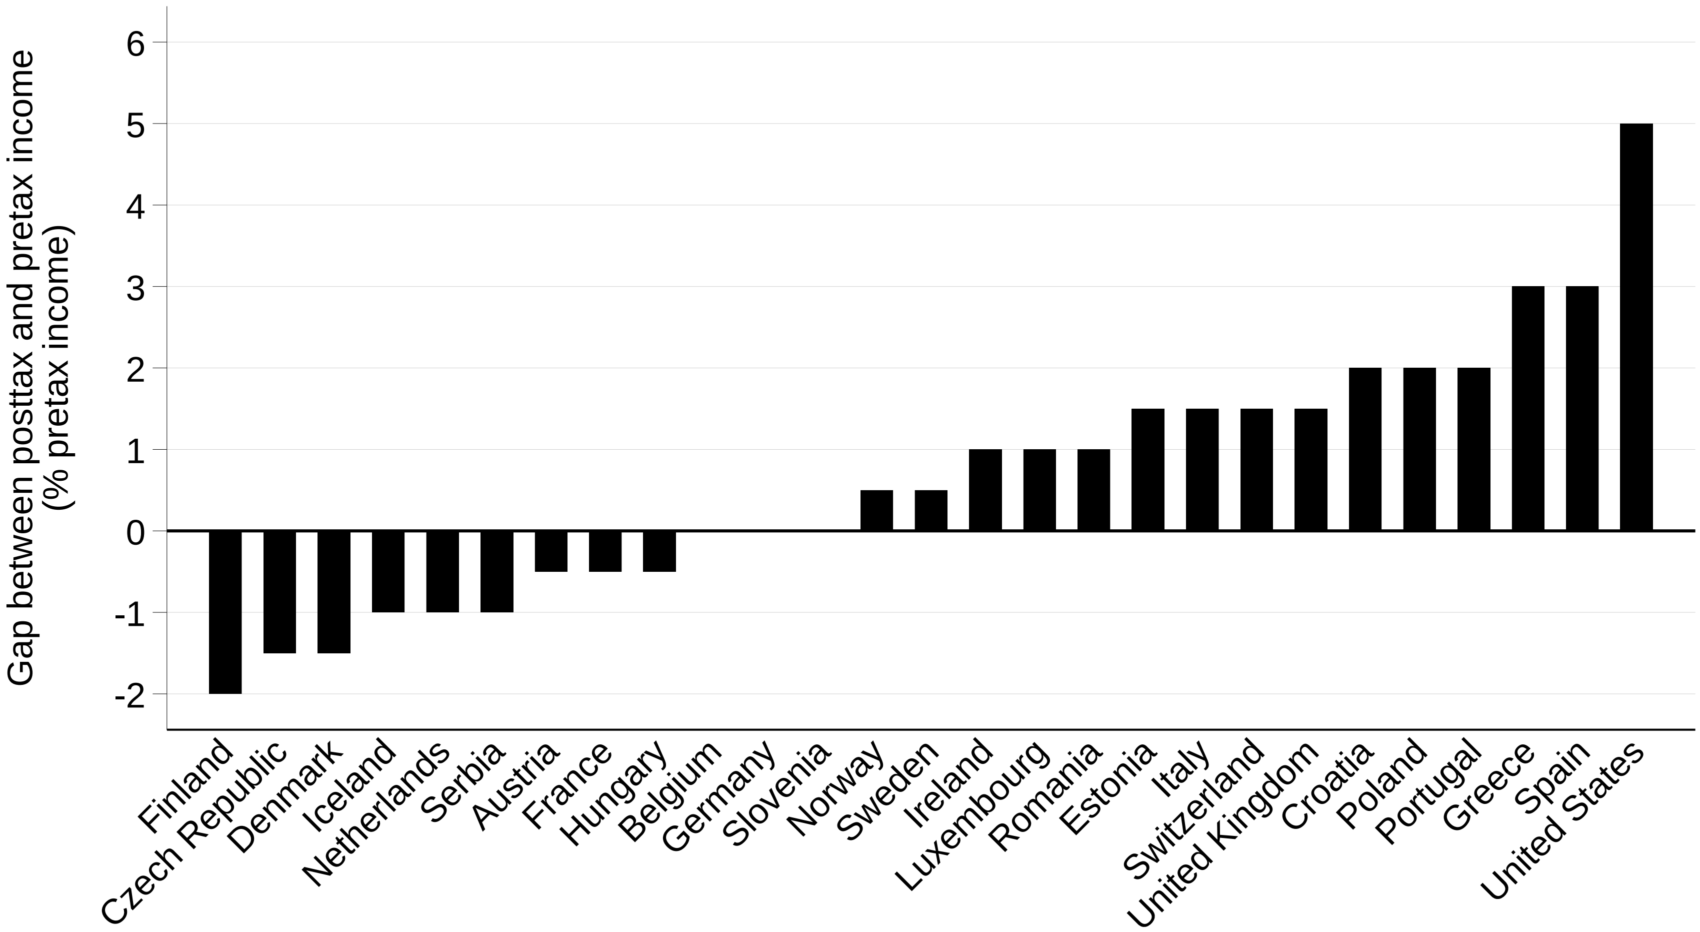

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. Results reported for the year 2017. Non-health collective government expenditures are assumed to be distributed proportionally to posttax disposable income.

Figure D.79: Net transfer received by the top 10% by country (% of pretax income)

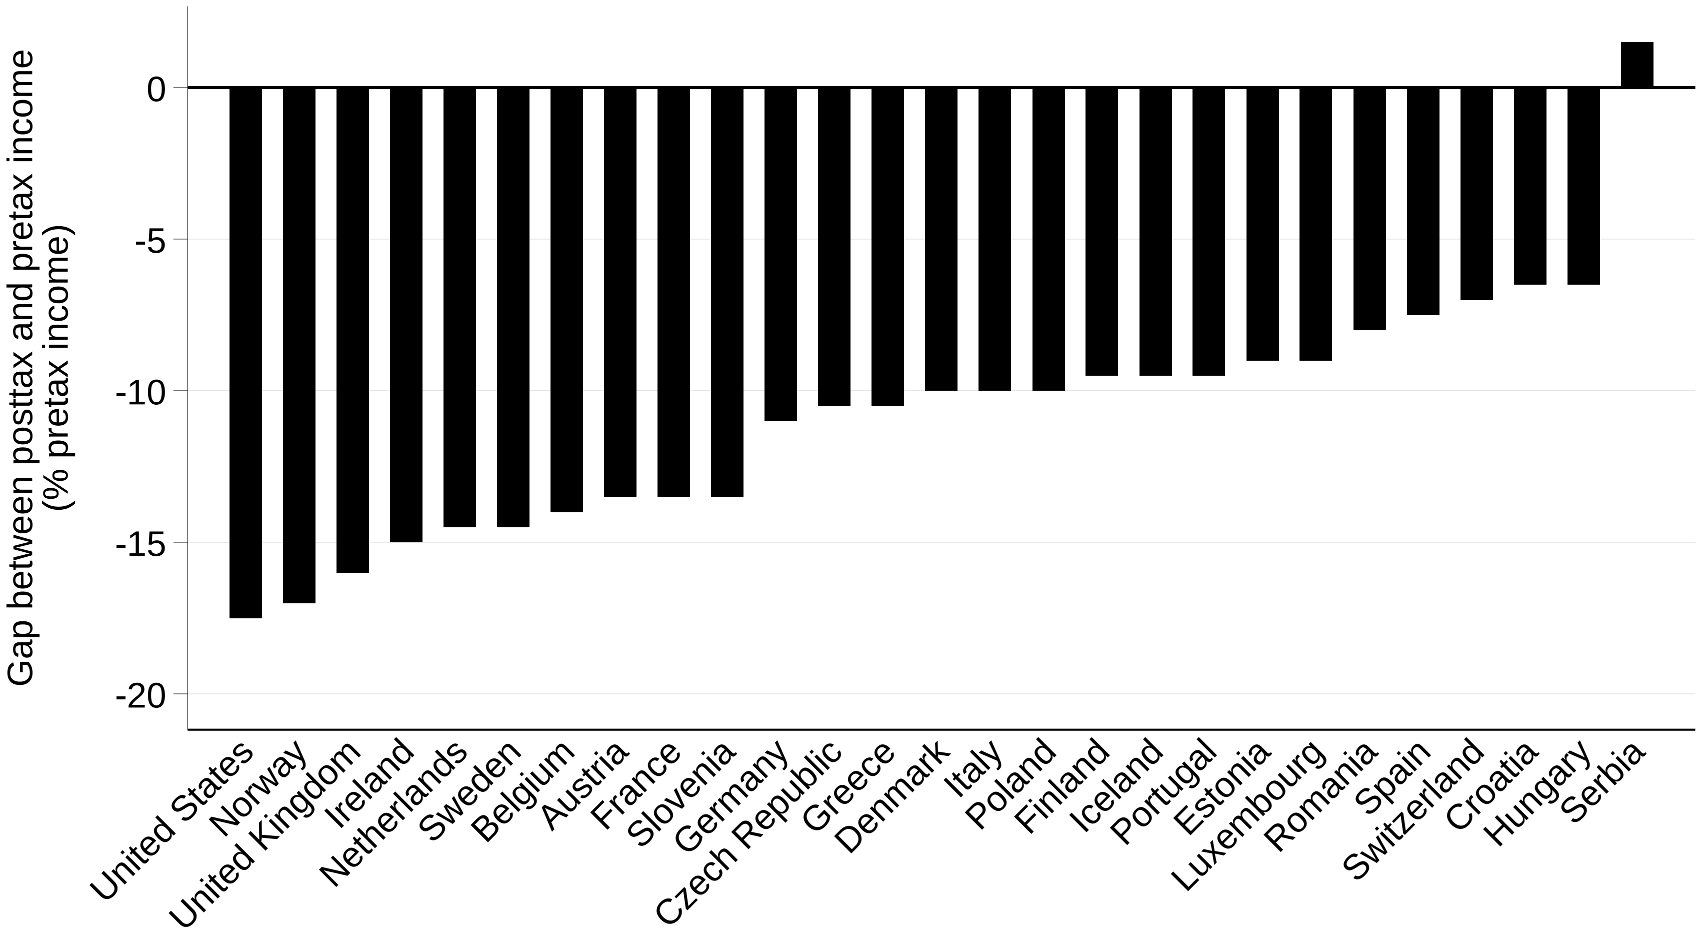

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. Results reported for the year 2017. Non-health collective government expenditures are assumed to be distributed proportionally to posttax disposable income.

Figure D.80: Net transfer received by the bottom 50% by country
(% of national income, lump sum allocation of collective expenditure)

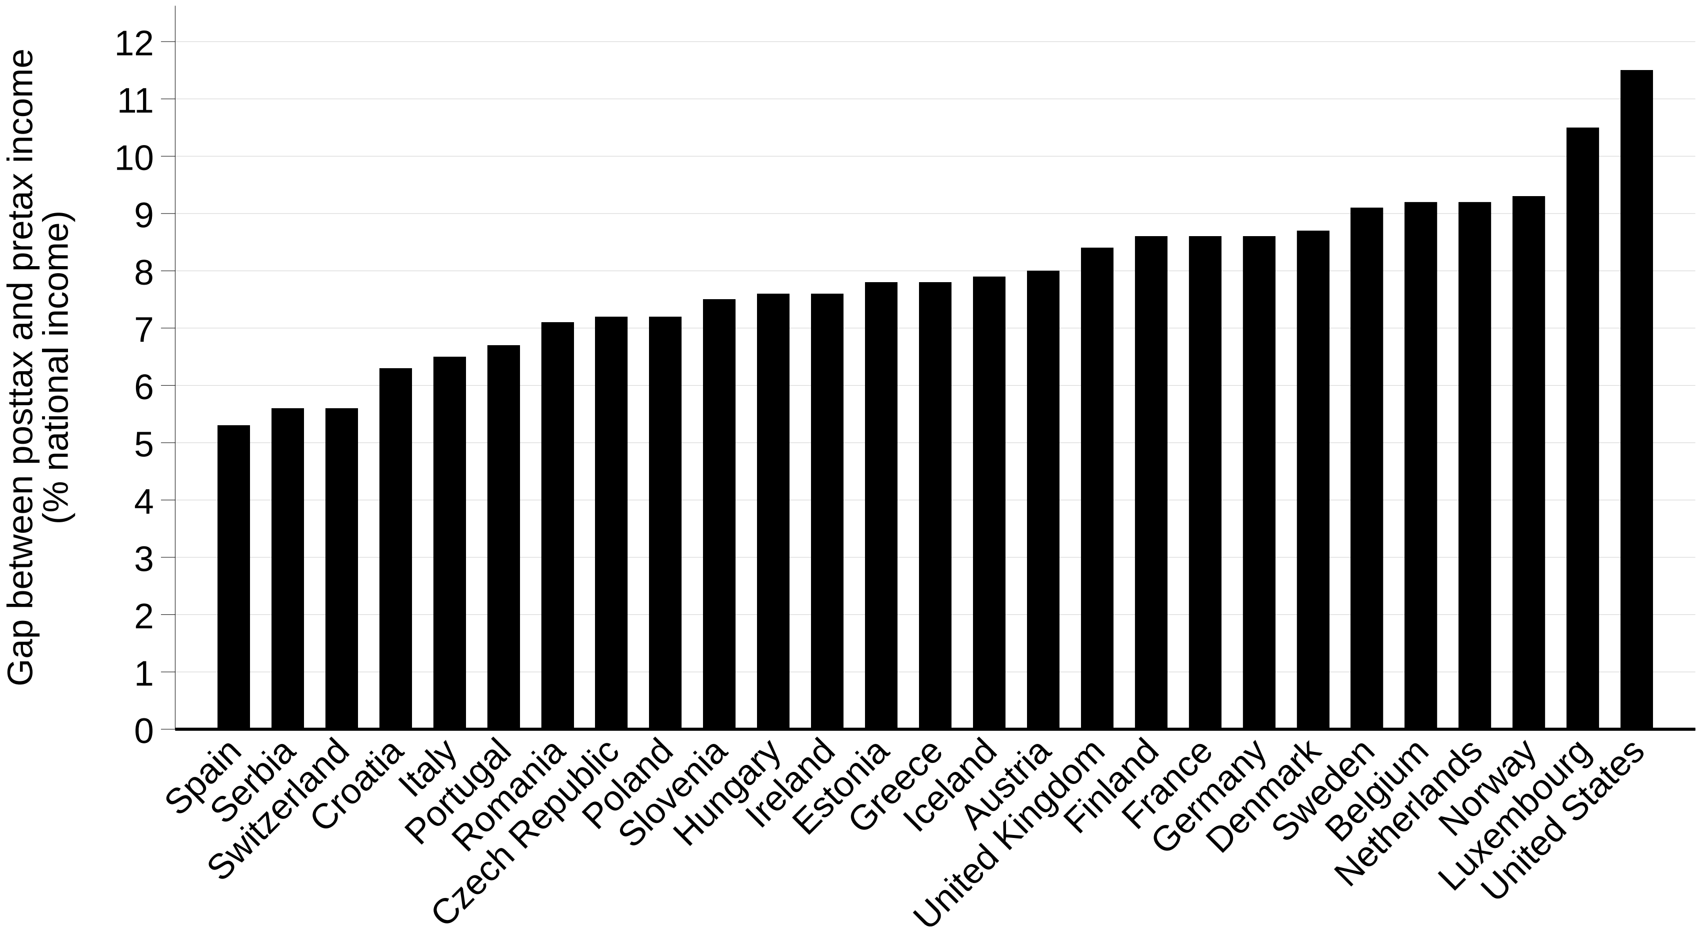

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. Results reported for the year 2017. Non-health collective government expenditures are assumed to be distributed on a lump sum basis.

Figure D.81: Net transfer received by the middle 40% by country
(% of national income, lump sum allocation of collective expenditure)

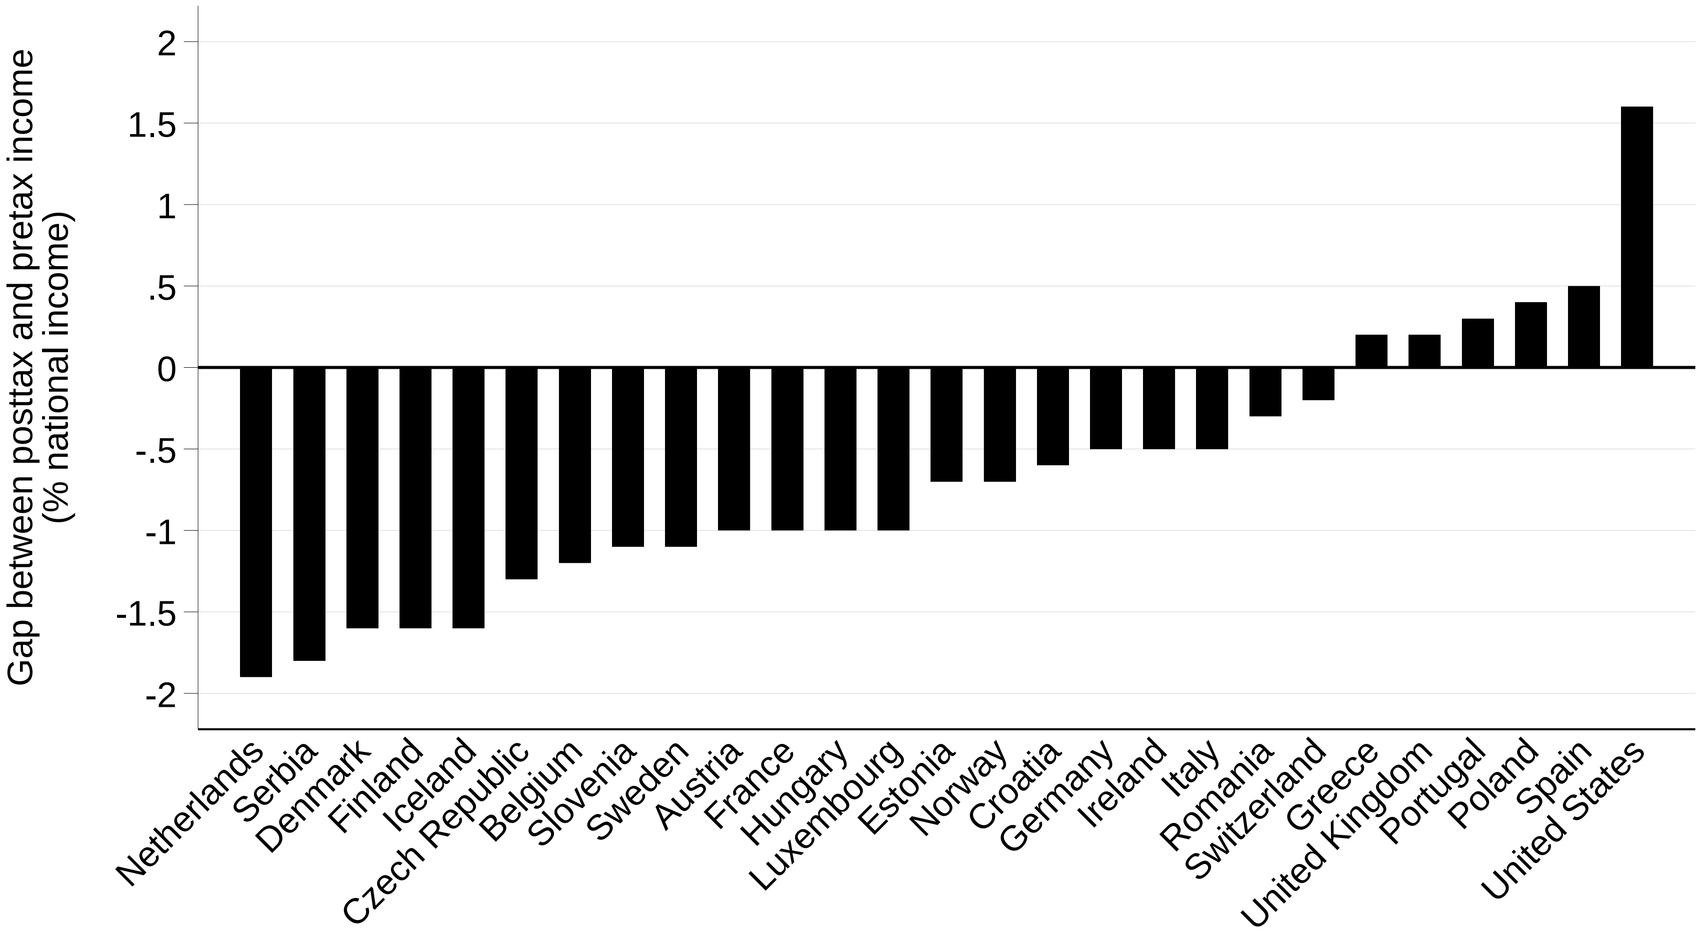

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. Results reported for the year 2017. Non-health collective government expenditures are assumed to be distributed on a lump sum basis.

Figure D.82: Net transfer received by the top 10% by country
(% of national income, lump sum allocation of collective expenditure)

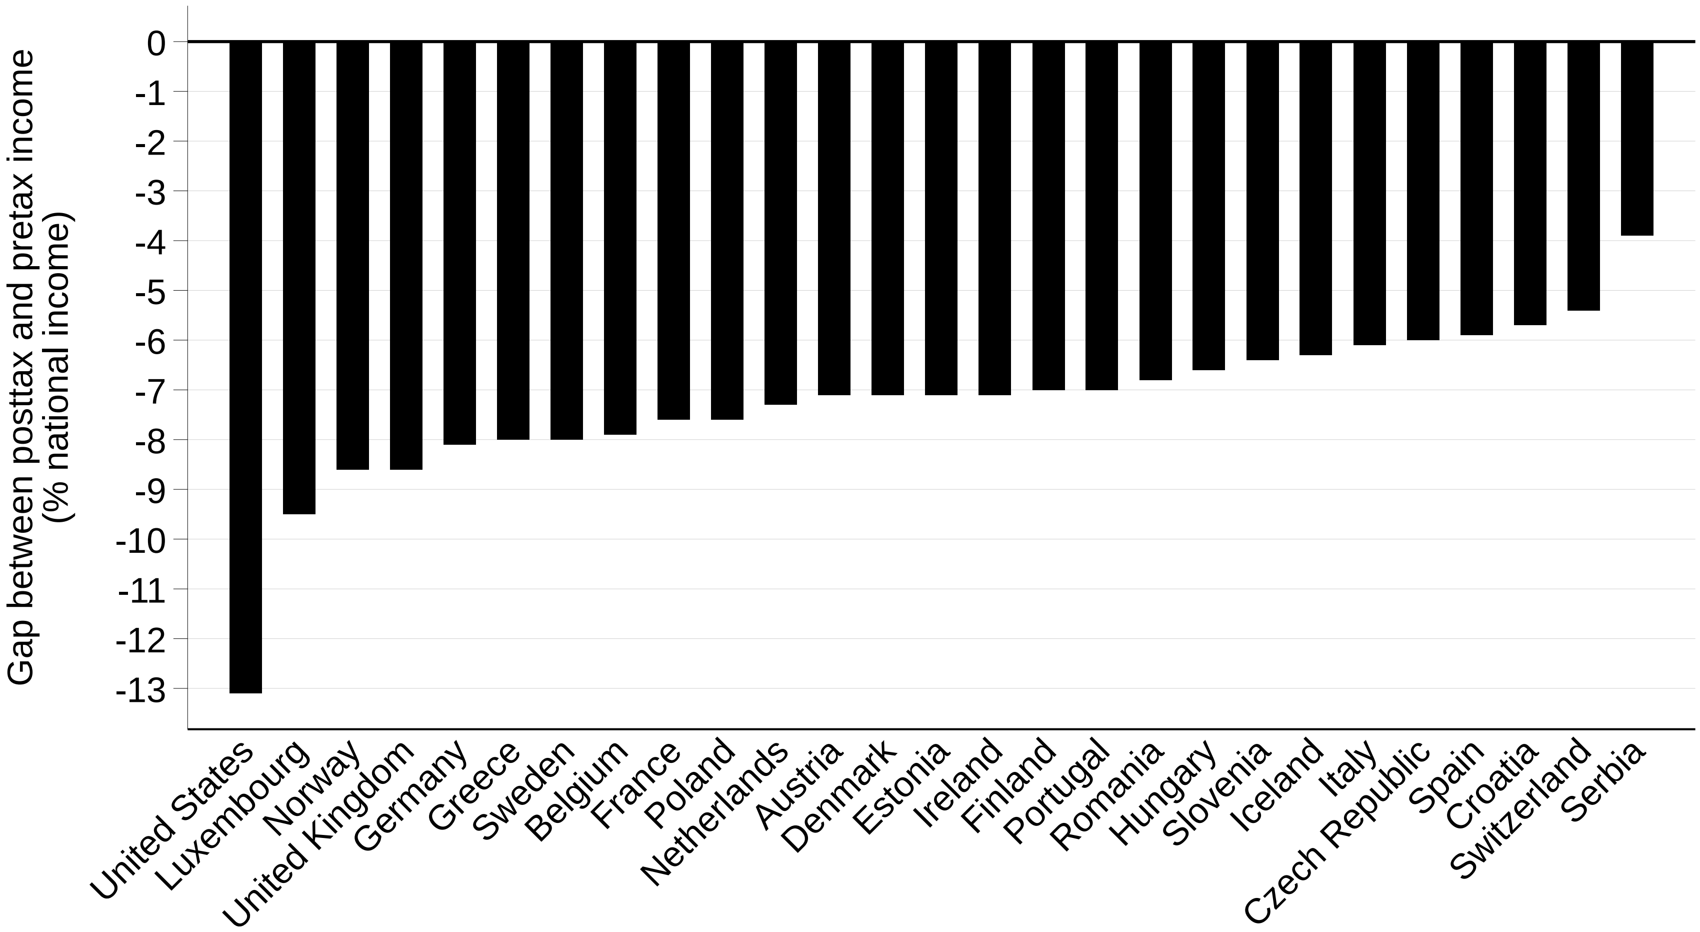

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. Results reported for the year 2017. Non-health collective government expenditures are assumed to be distributed on a lump sum basis.

D.2.6 Maps

Figure D.83: Map of top 10% pretax income share in Europe, 1980

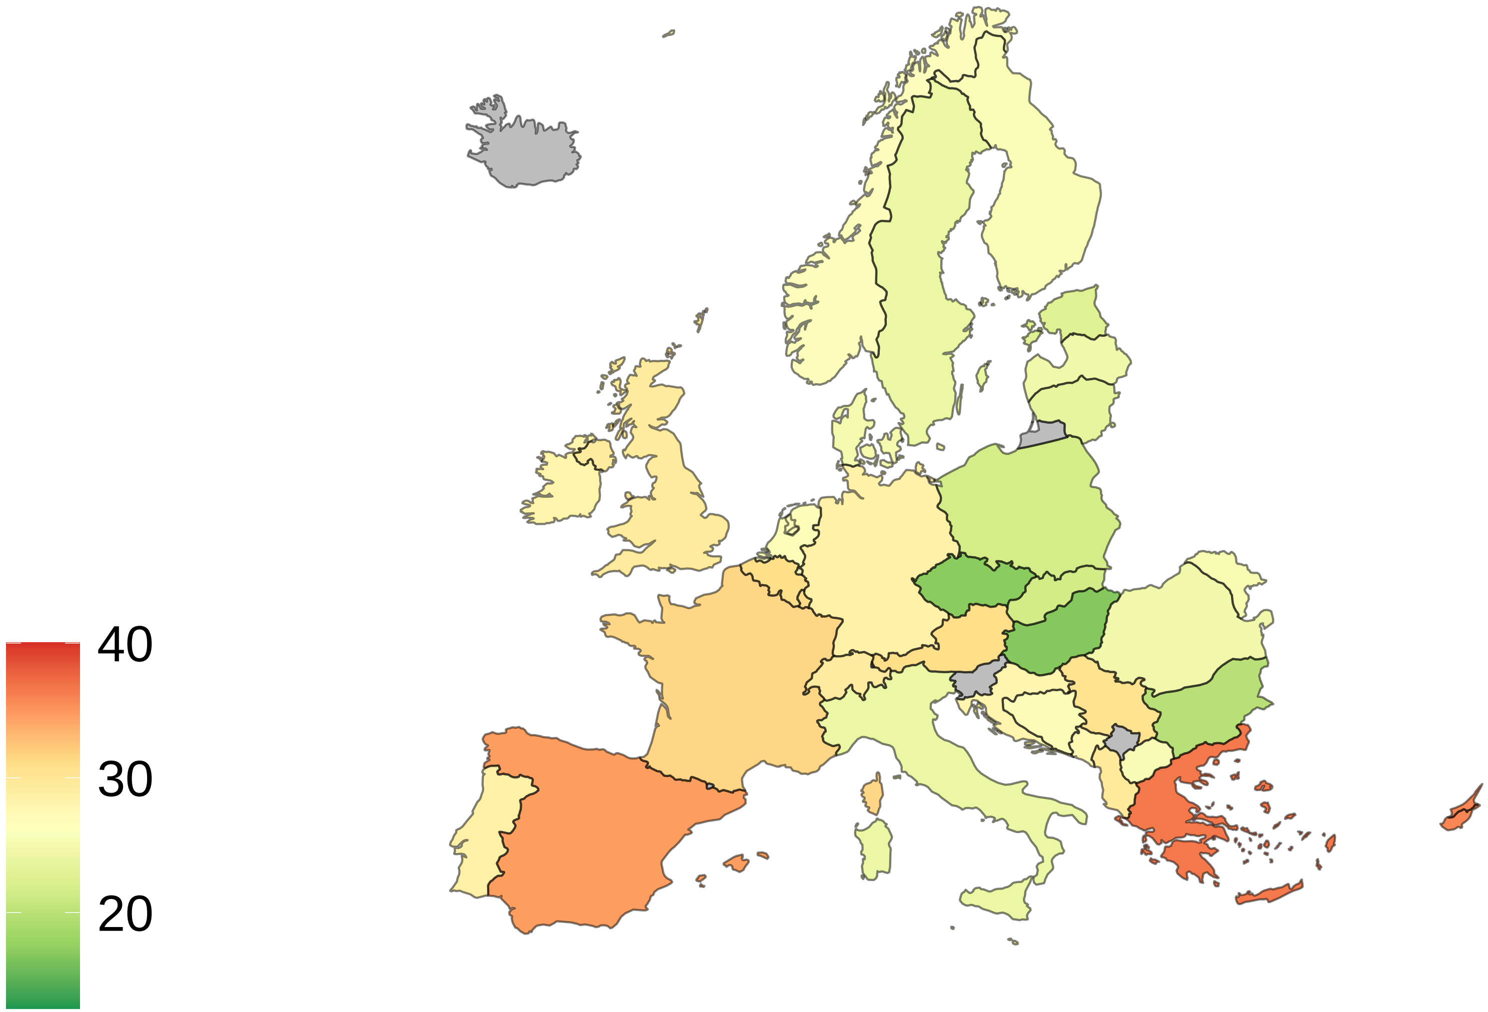

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.84: Map of top 10% pretax income share in Europe, 1990

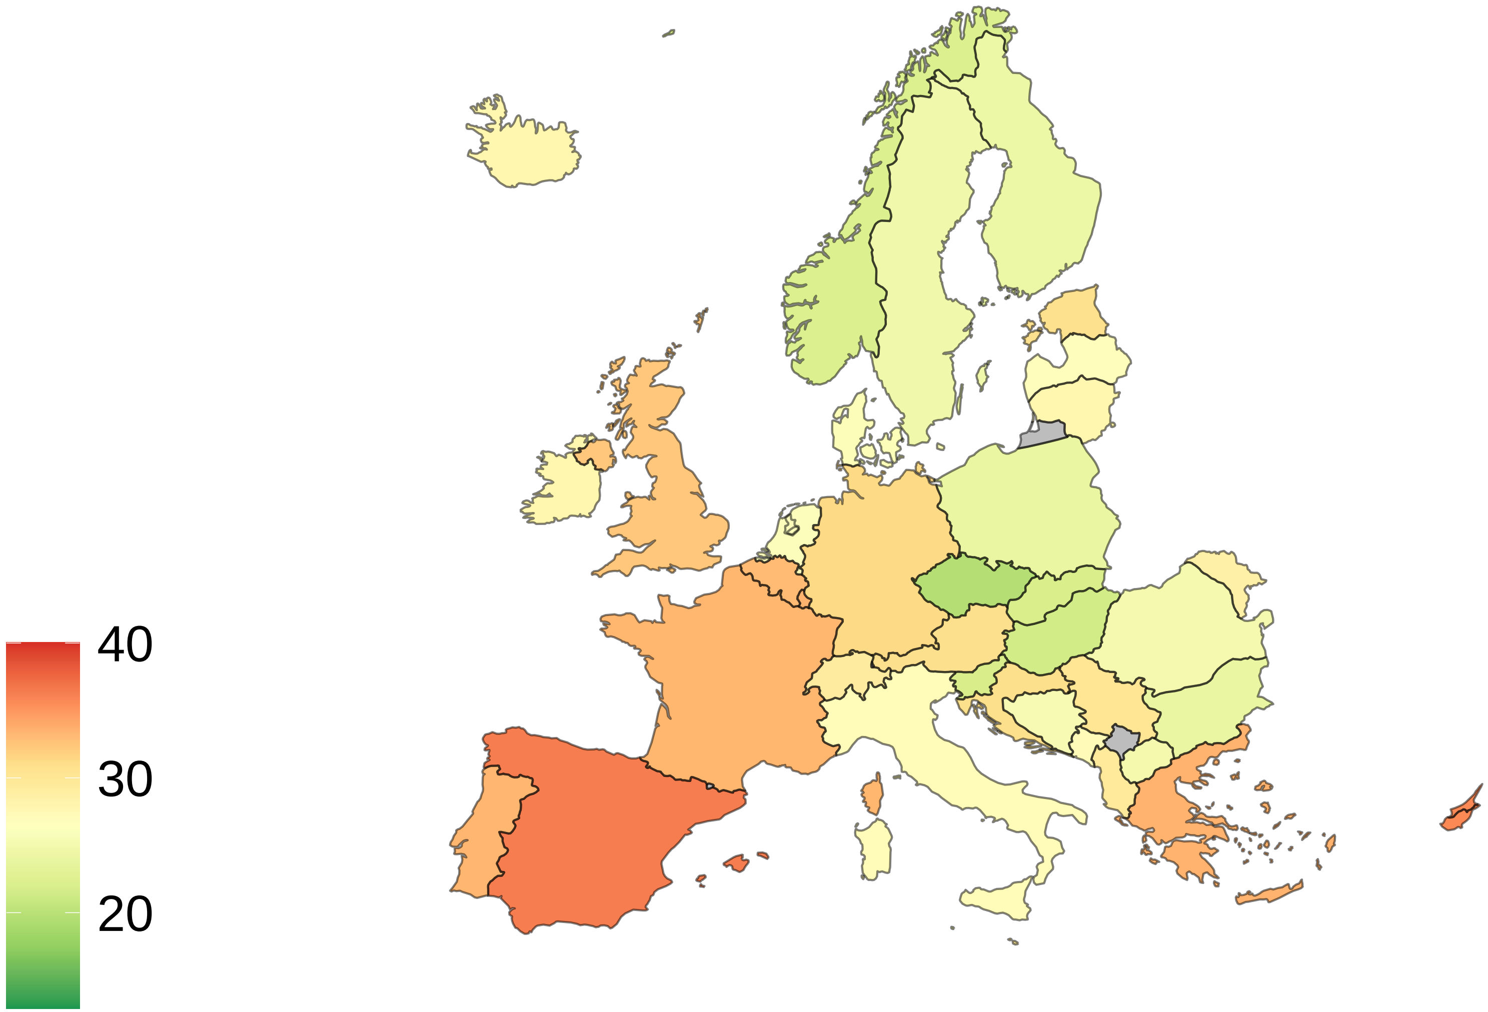

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.85: Map of top 10% pretax income share in Europe, 2000

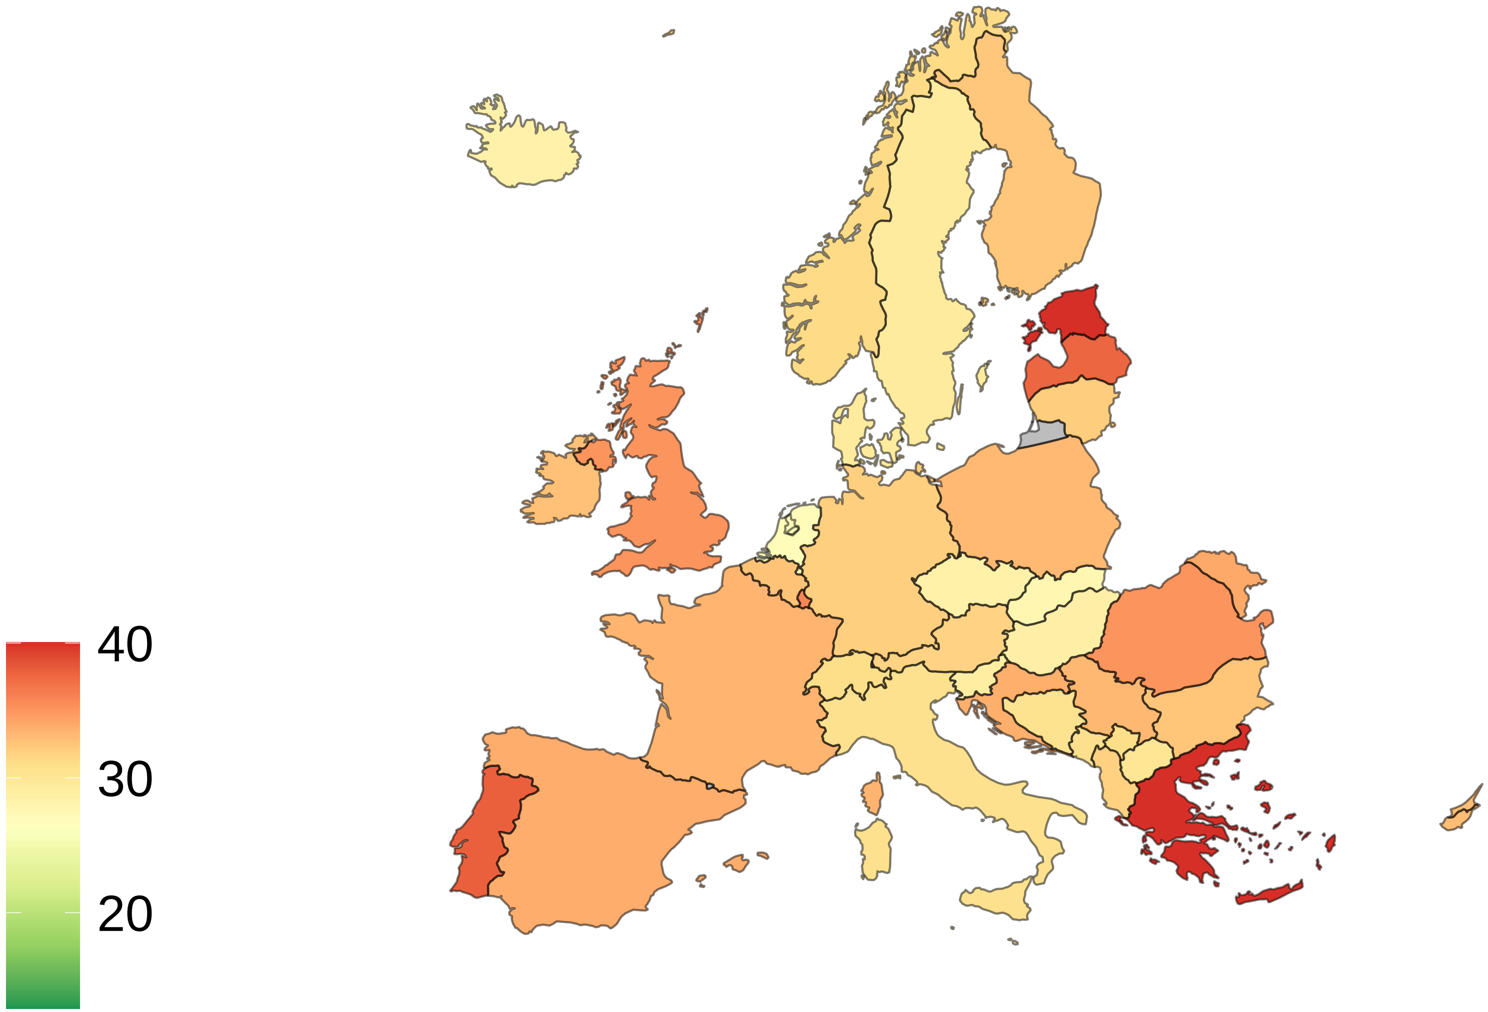

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.86: Map of top 10% pretax income share in Europe, 2007

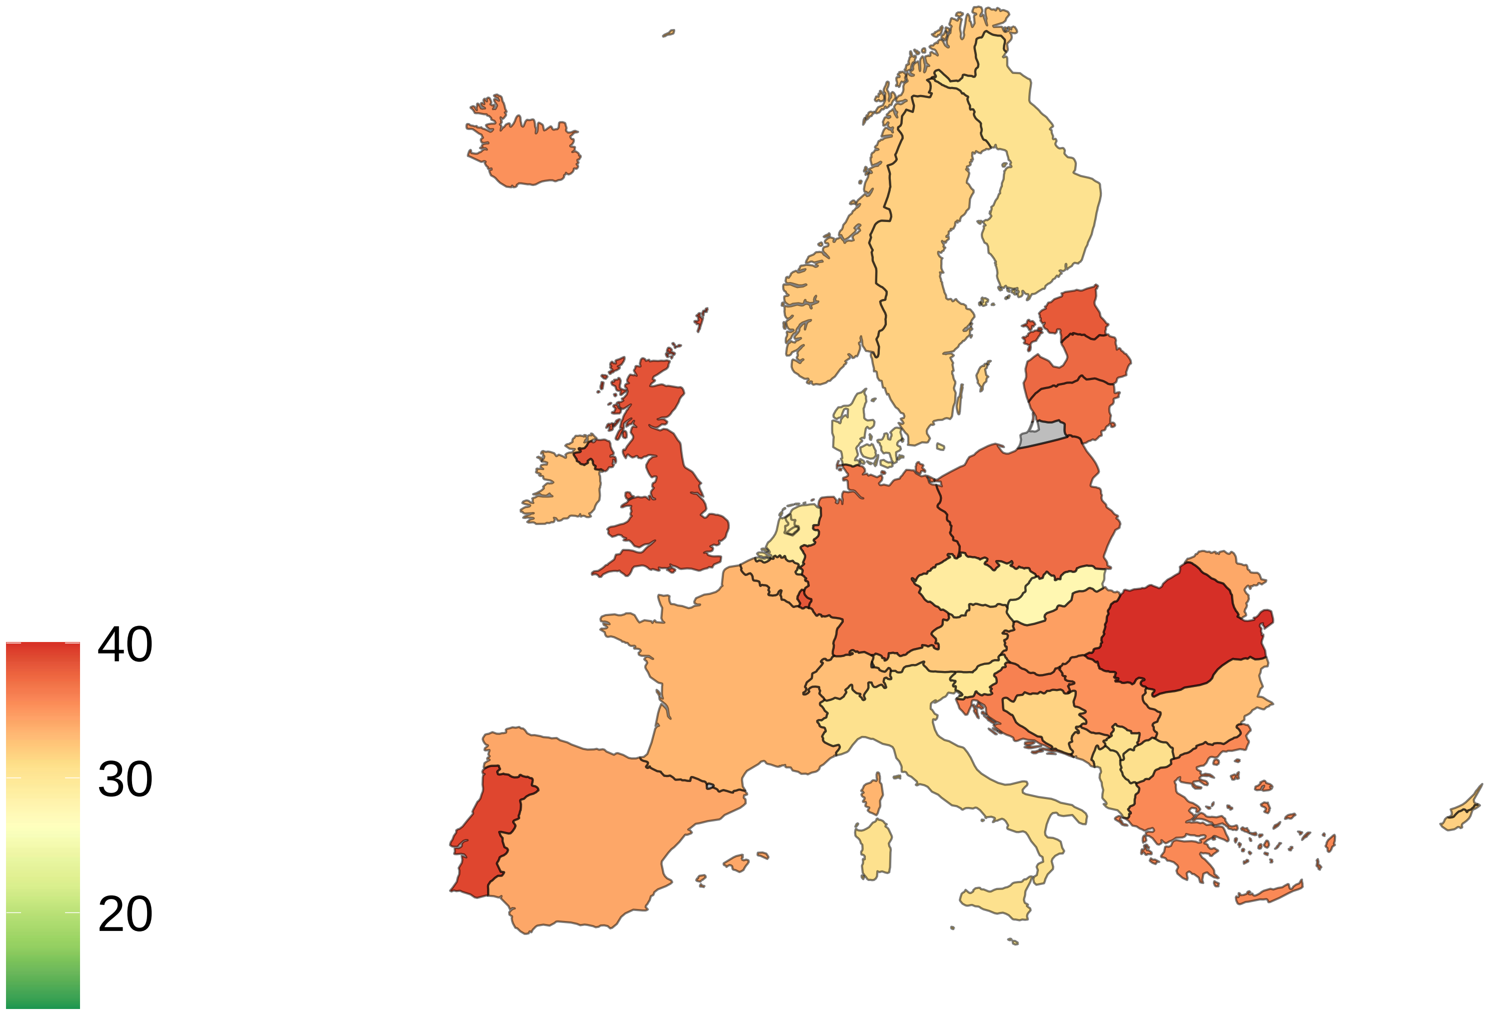

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.87: Map of top 10% pretax income share in Europe, 2017

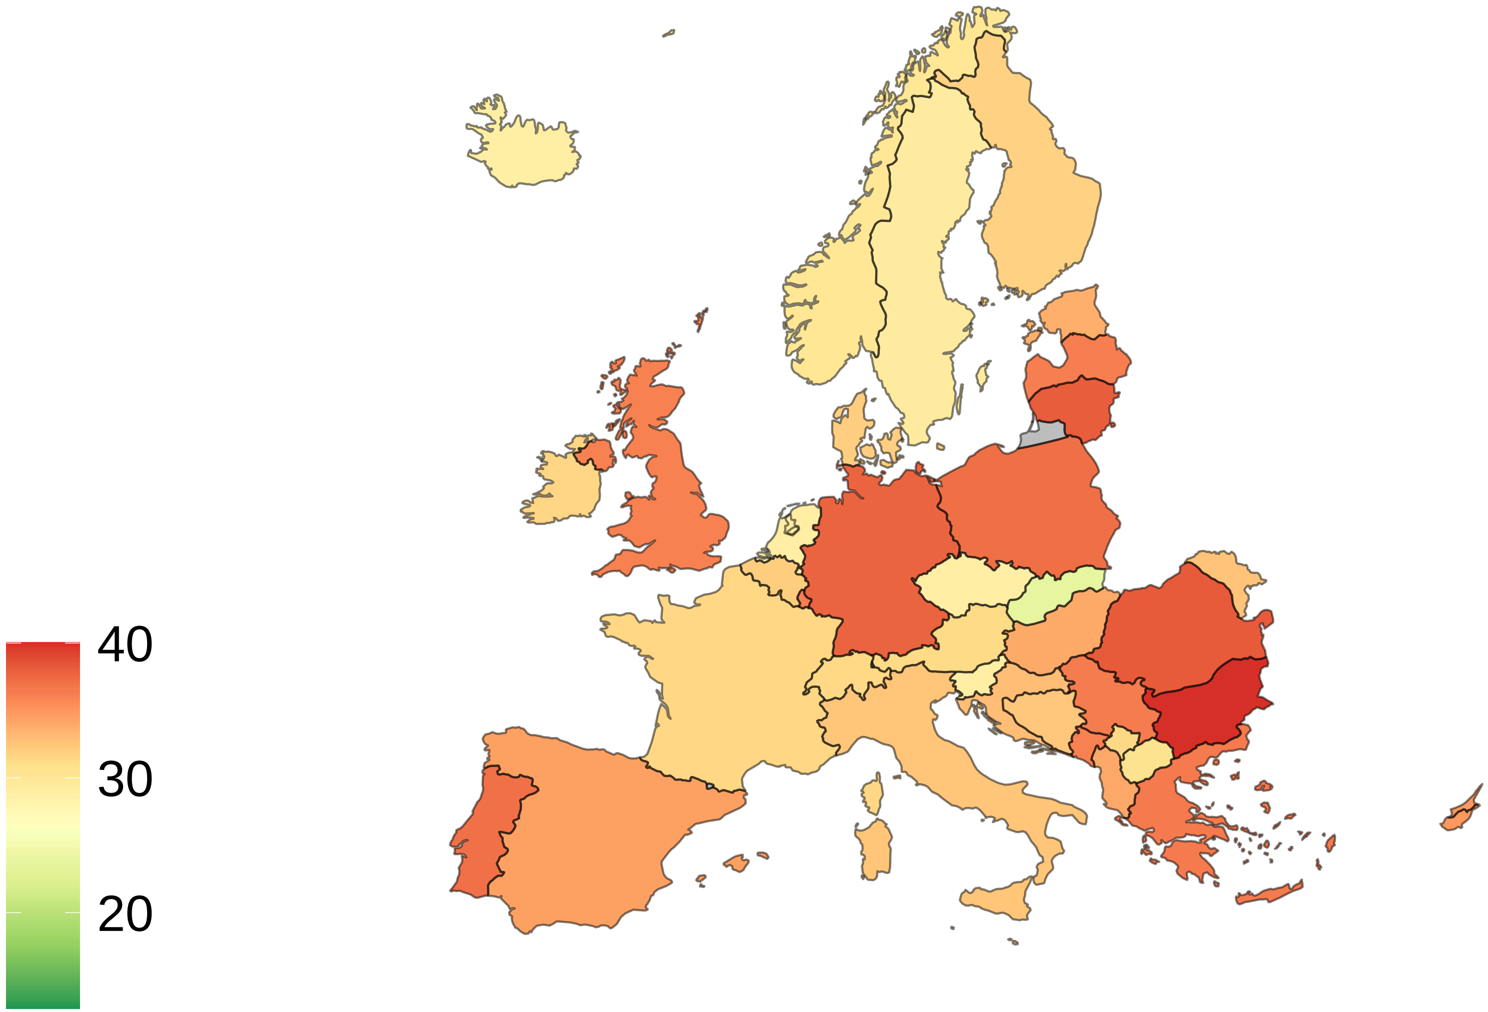

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.88: Map of top 10% posttax income share in Europe, 1980

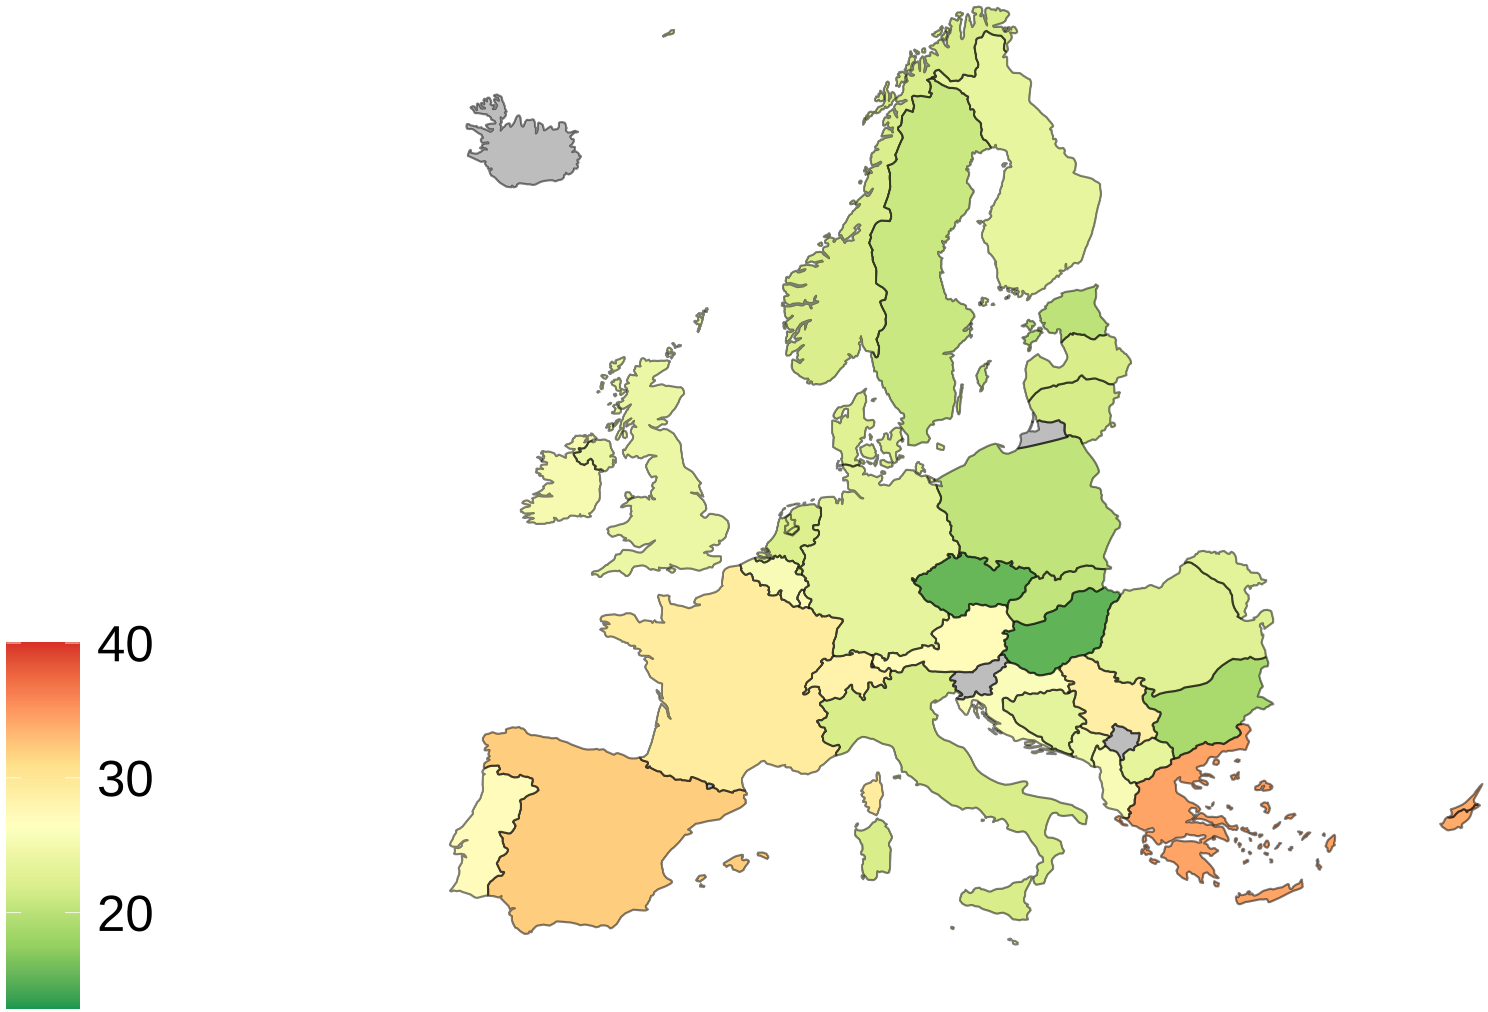

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.89: Map of top 10% posttax income share in Europe, 1990

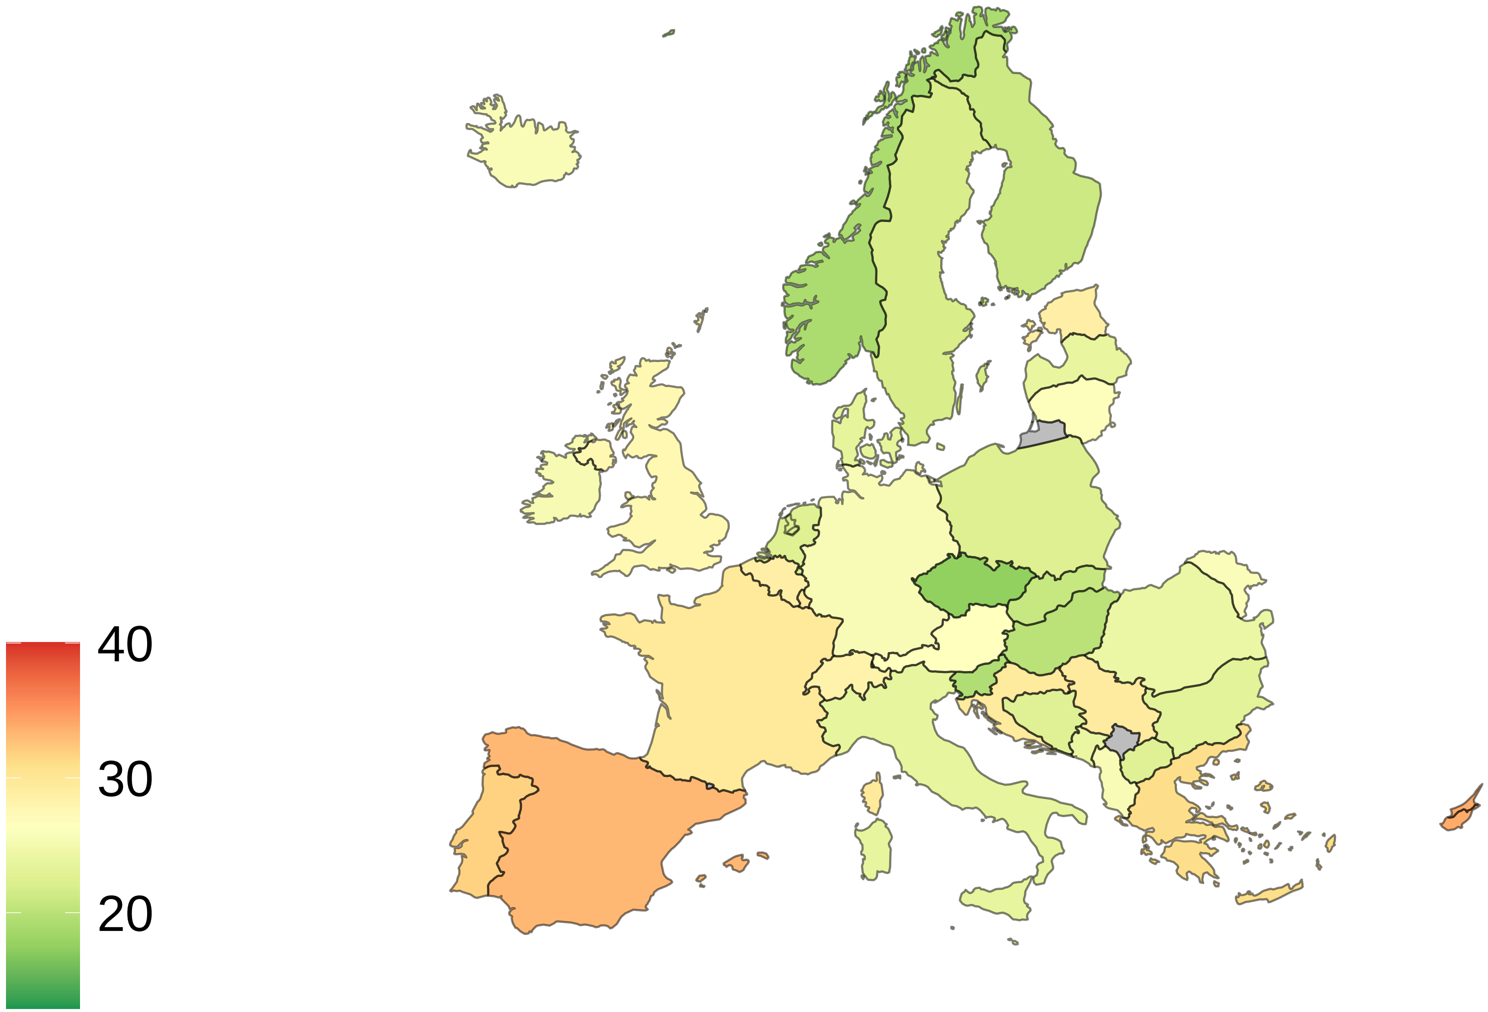

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.90: Map of top 10% posttax income share in Europe, 2000

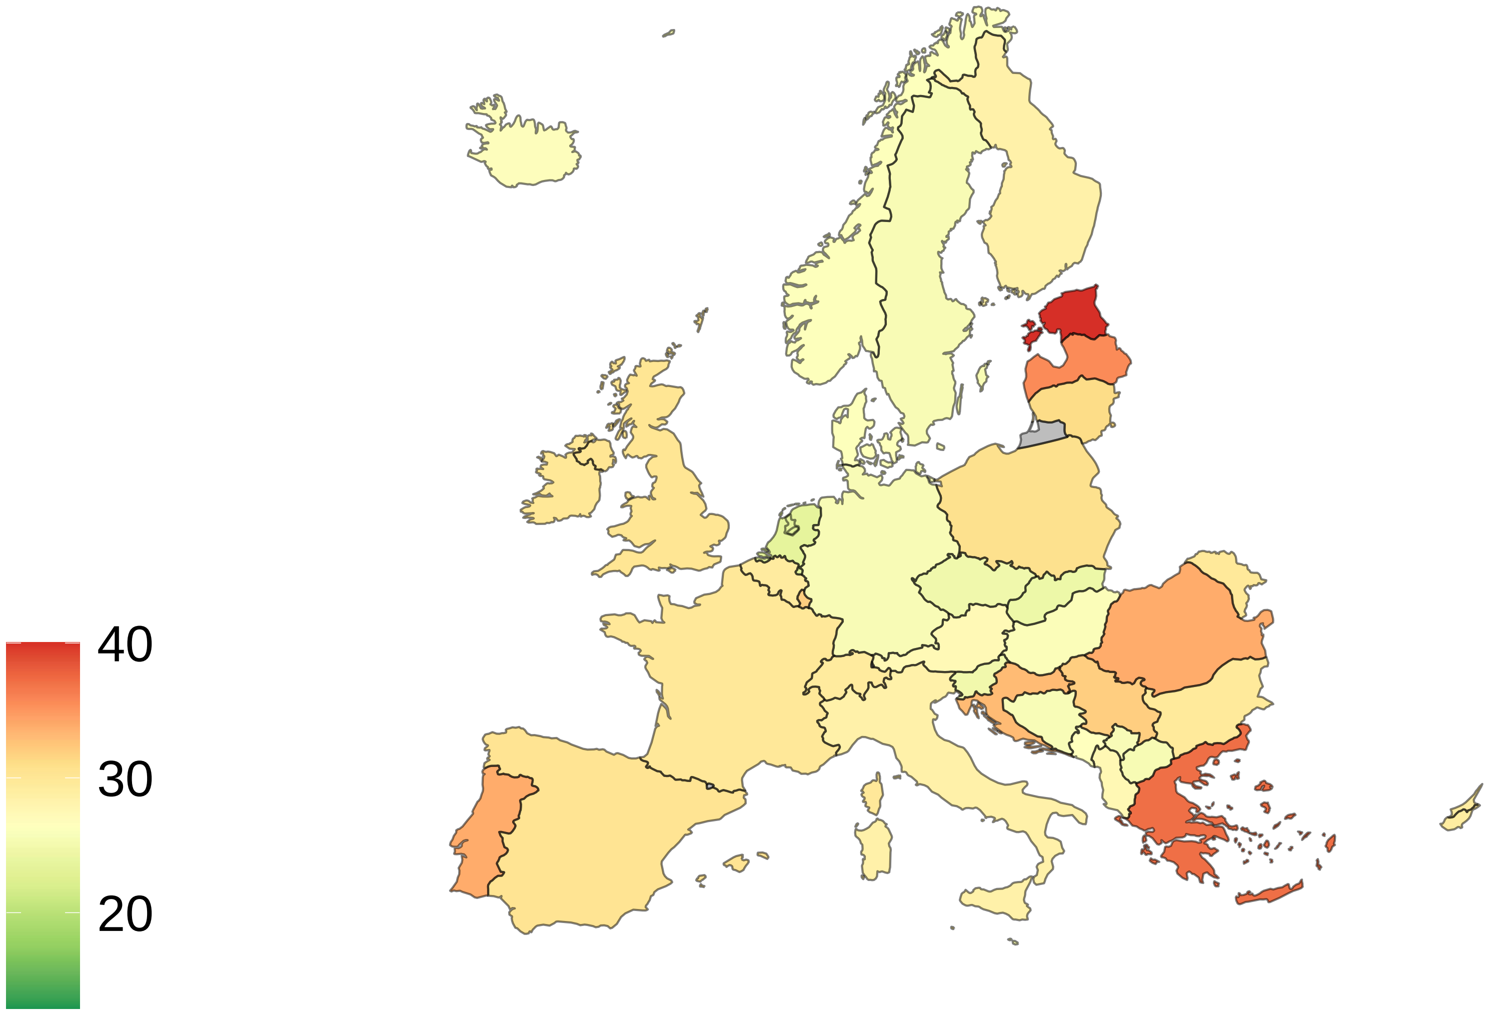

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.91: Map of top 10% posttax income share in Europe, 2007

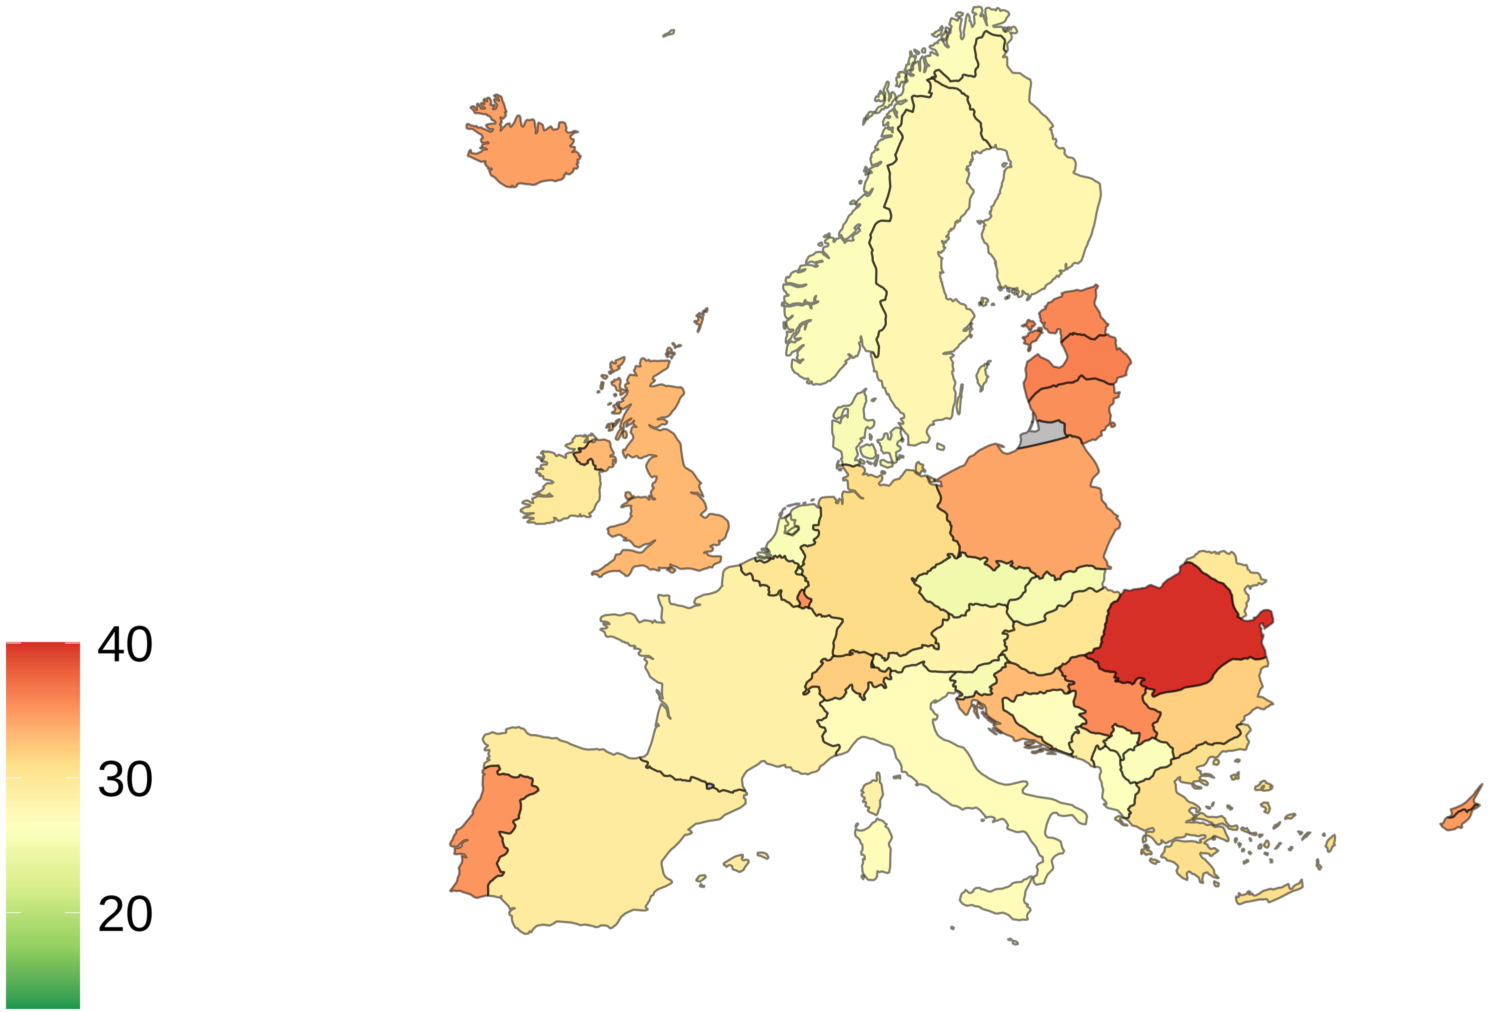

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.92: Map of top 10% posttax income share in Europe, 2017

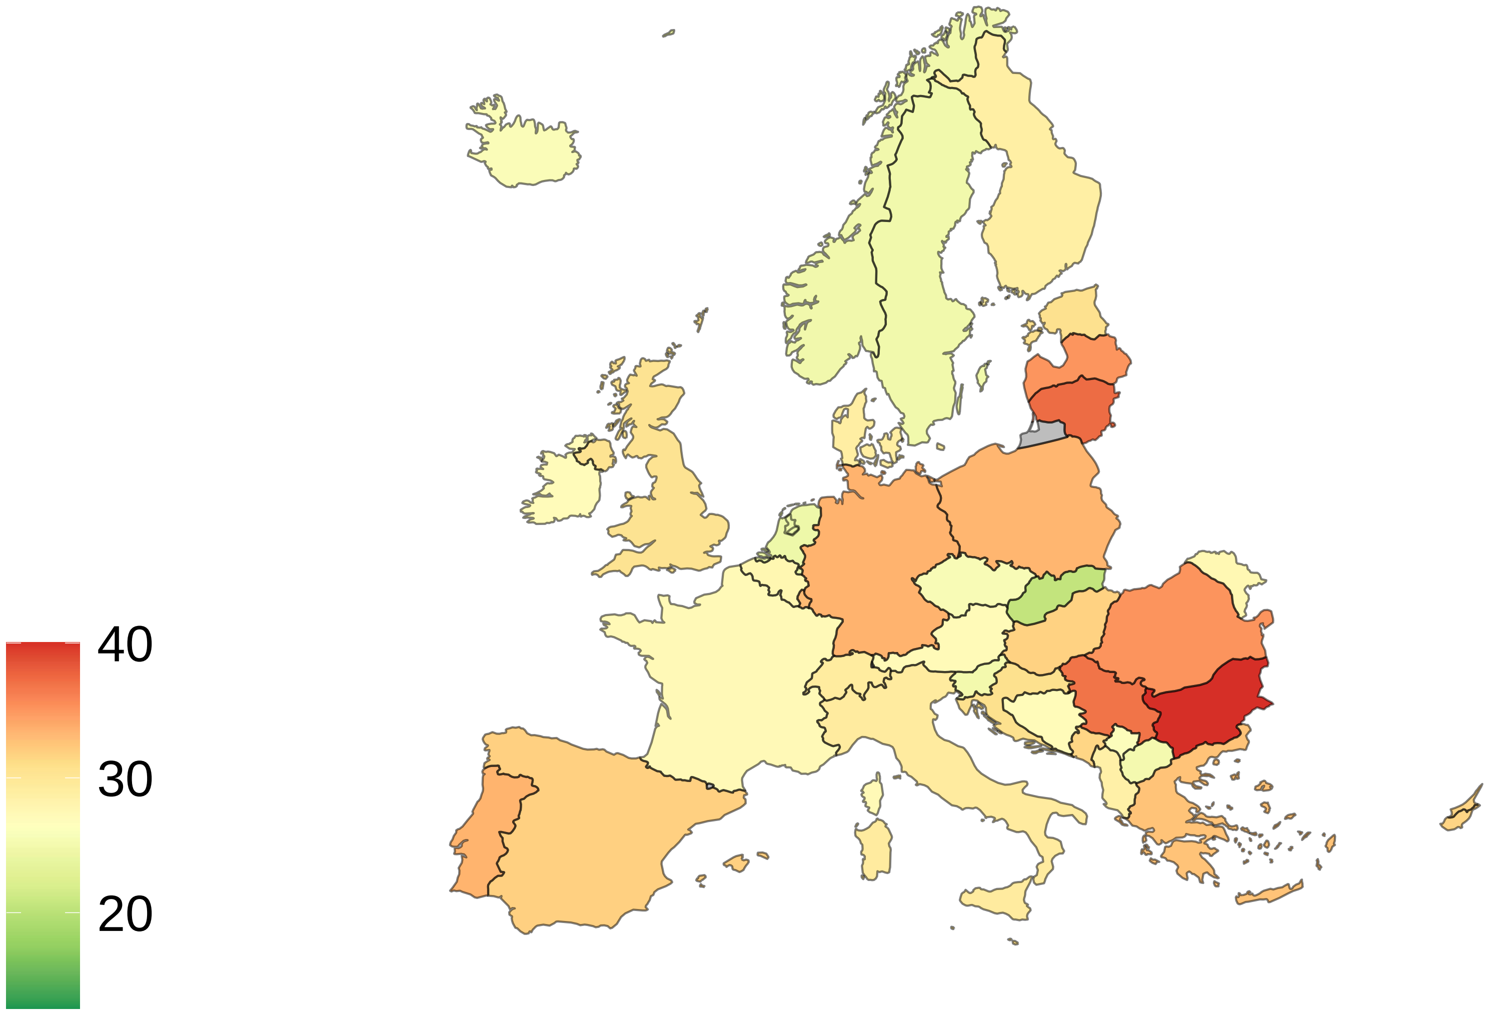

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.93: Map of top 1% pretax income share in Europe, 1980

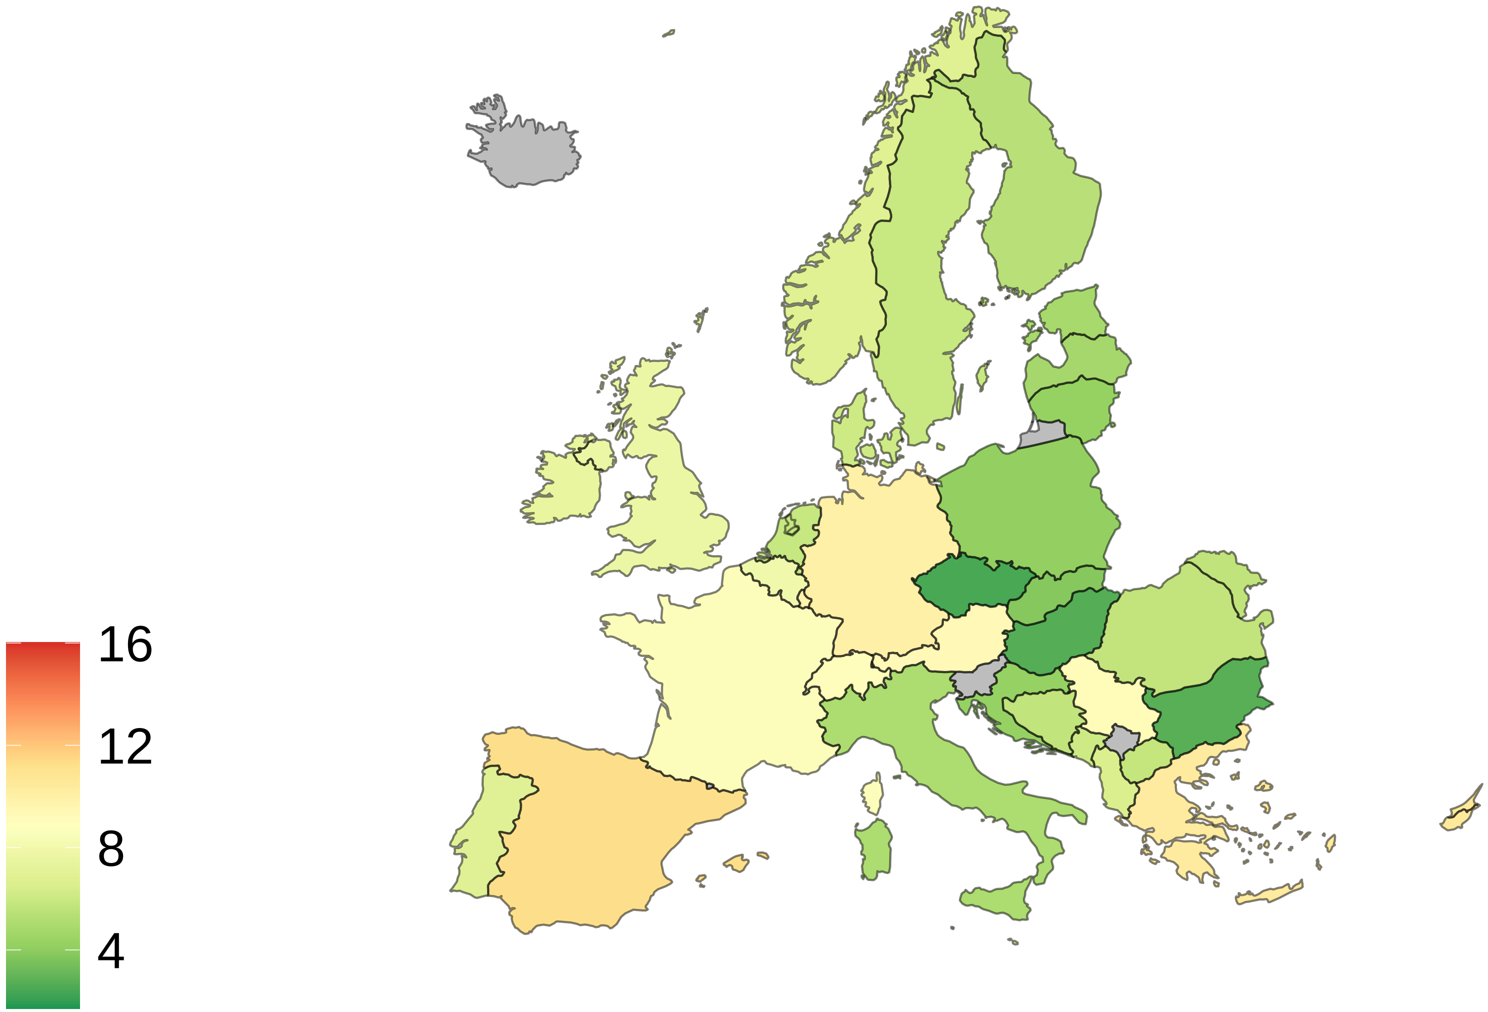

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.94: Map of top 1% pretax income share in Europe, 1990

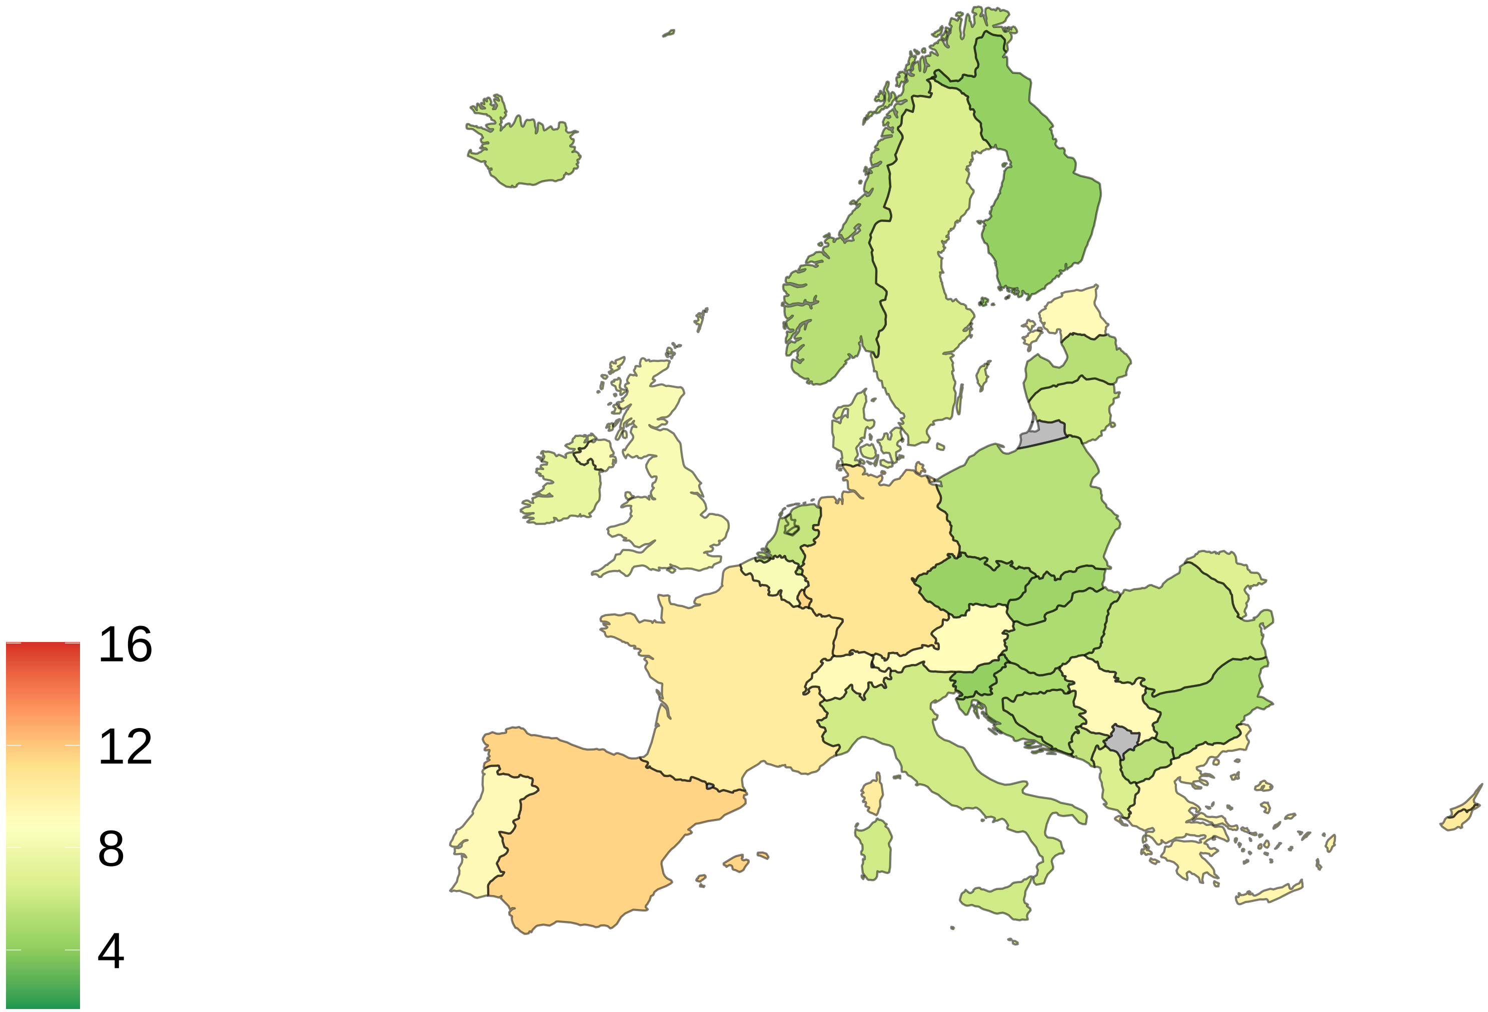

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.95: Map of top 1% pretax income share in Europe, 2000

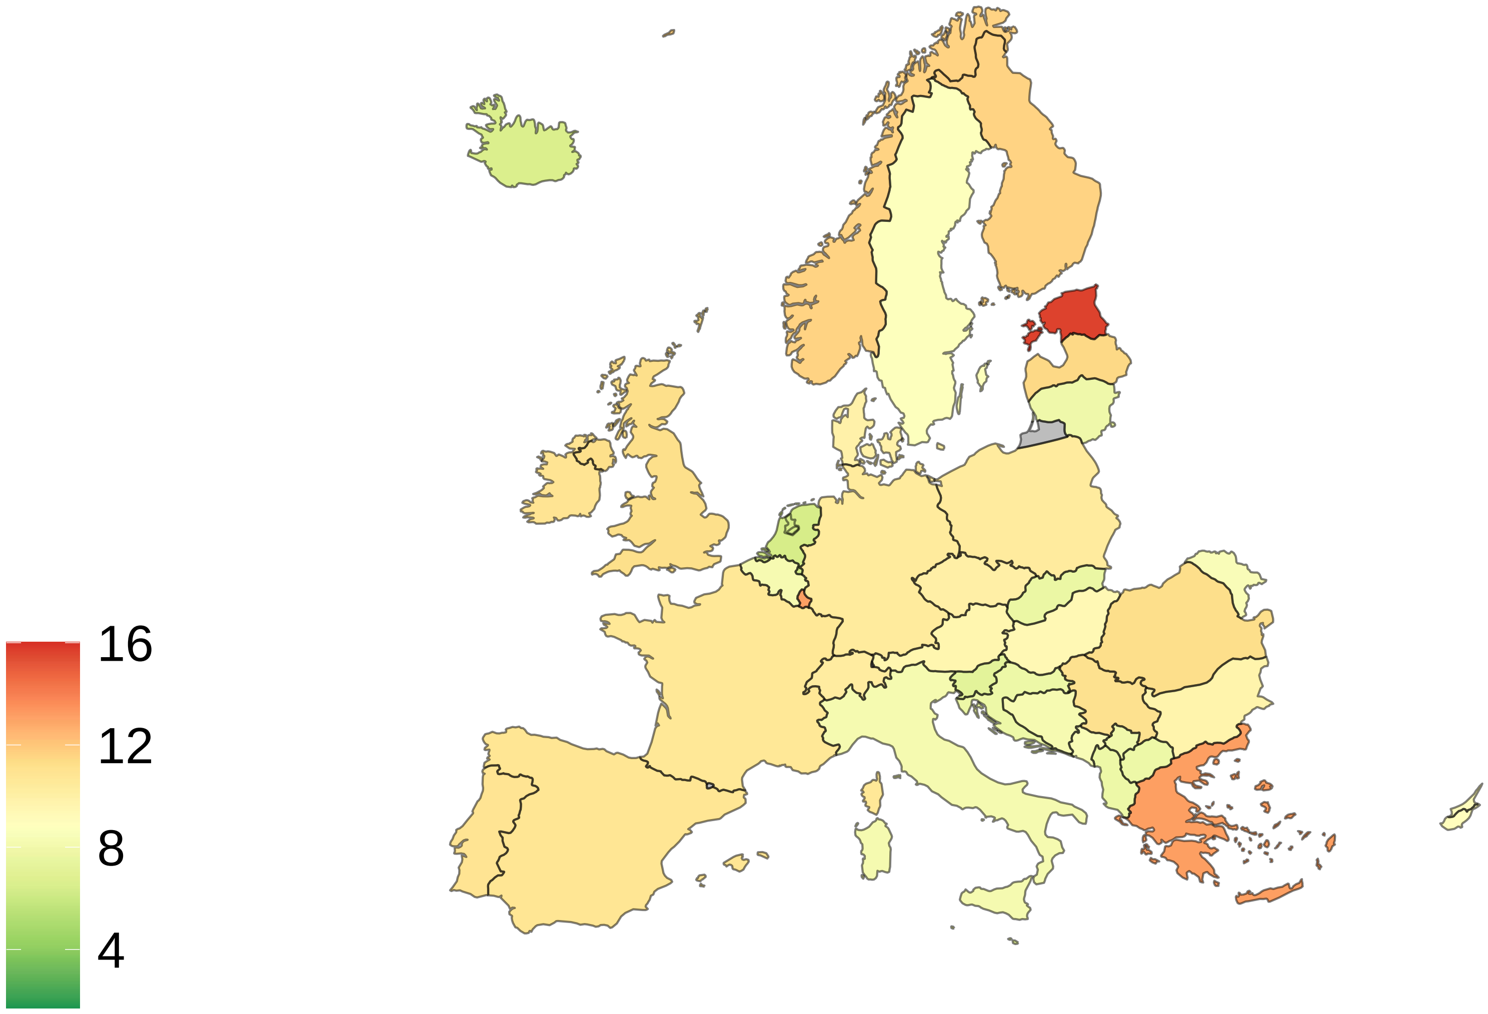

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.96: Map of top 1% pretax income share in Europe, 2007

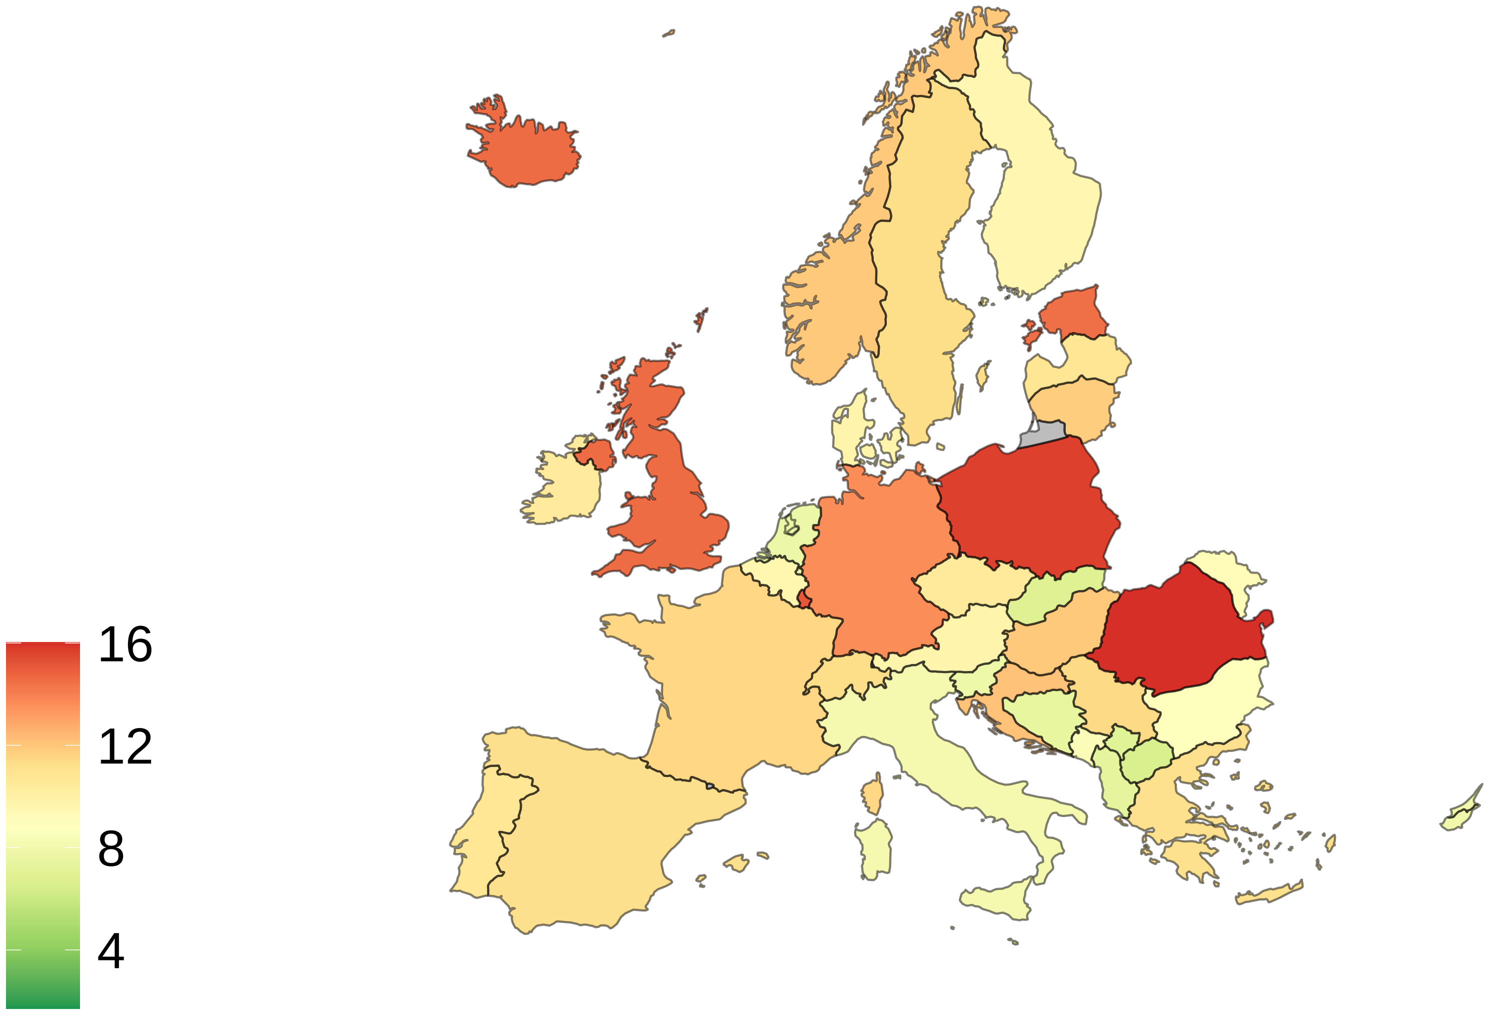

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.97: Map of top 1% pretax income share in Europe, 2017

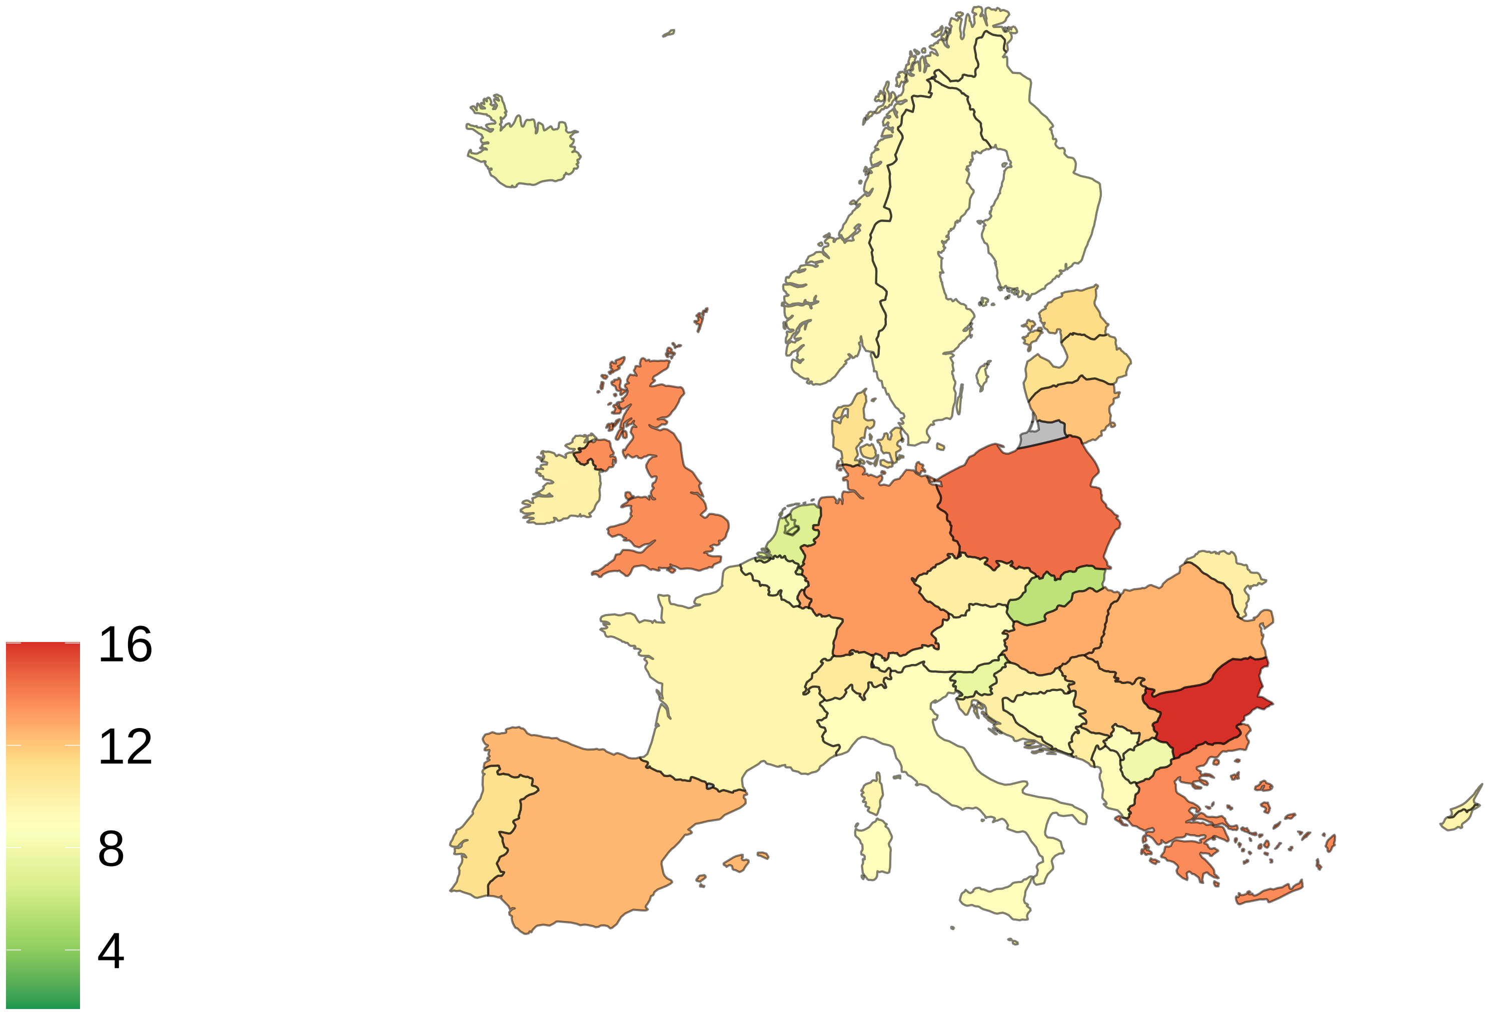

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.98: Map of top 1% posttax income share in Europe, 1980

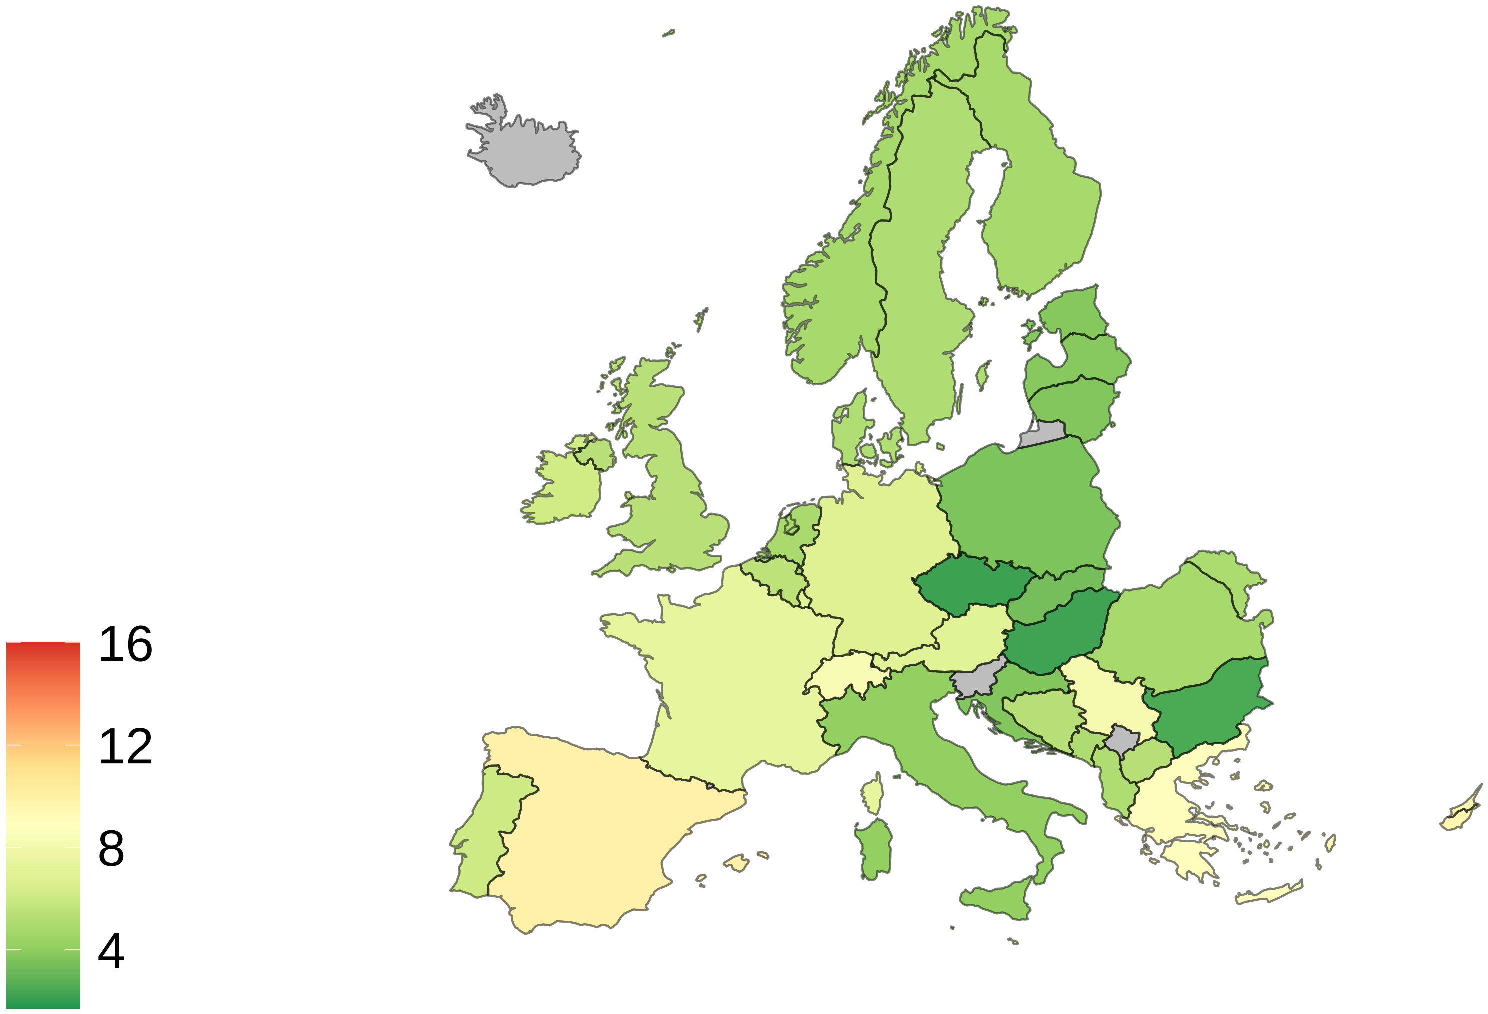

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.99: Map of top 1% posttax income share in Europe, 1990

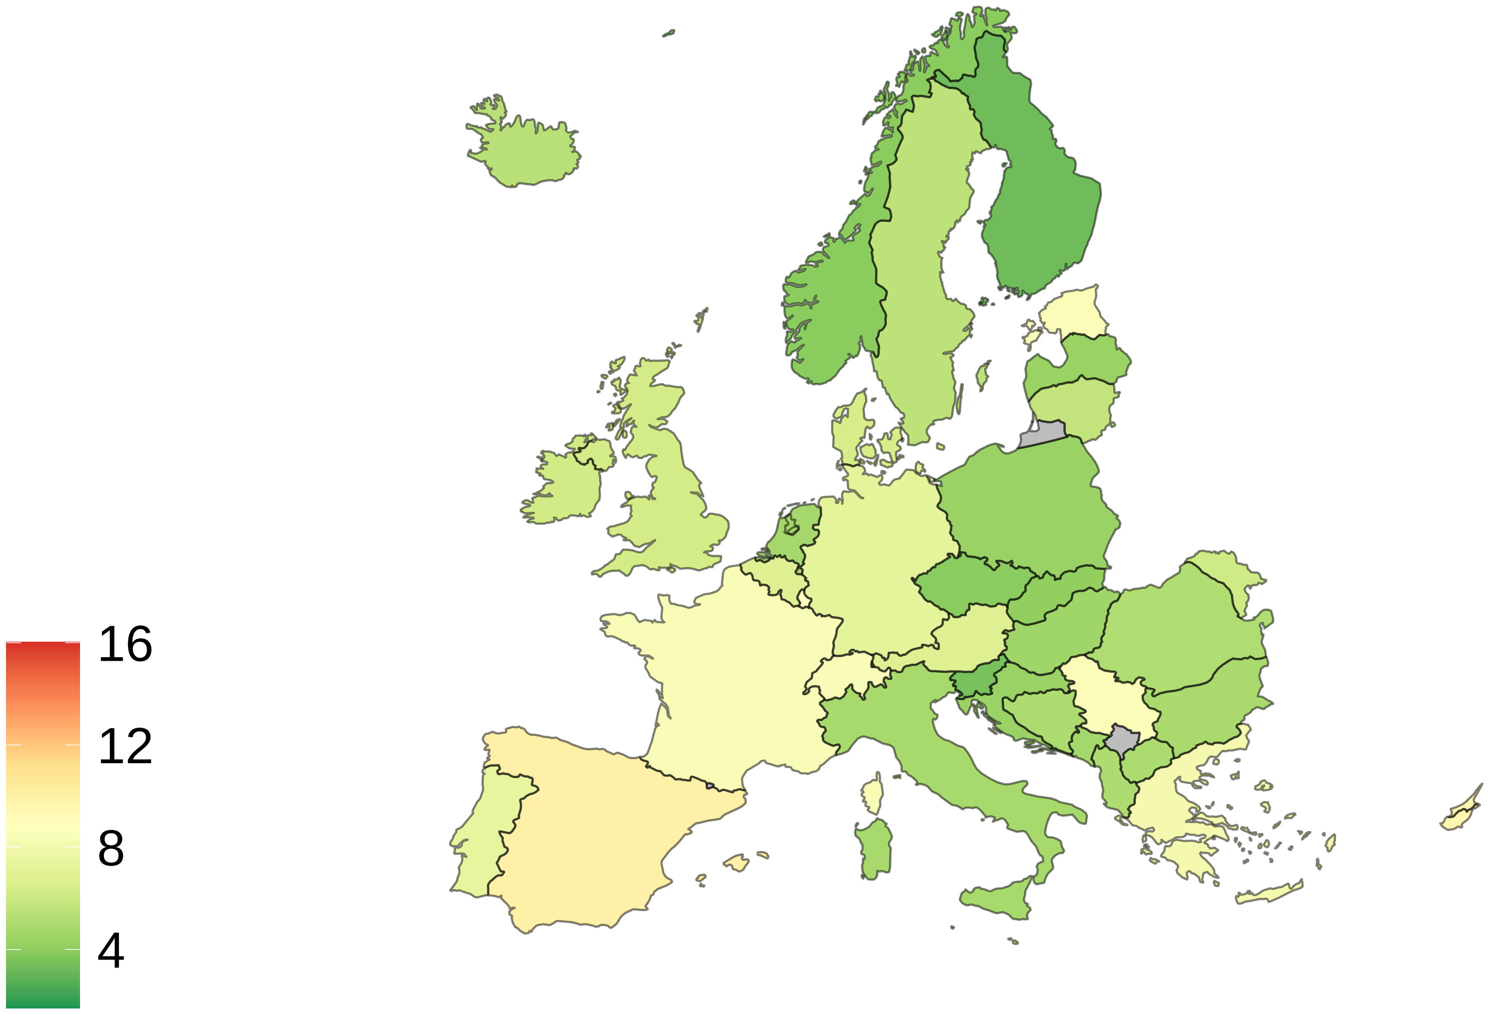

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.100: Map of top 1% posttax income share in Europe, 2000

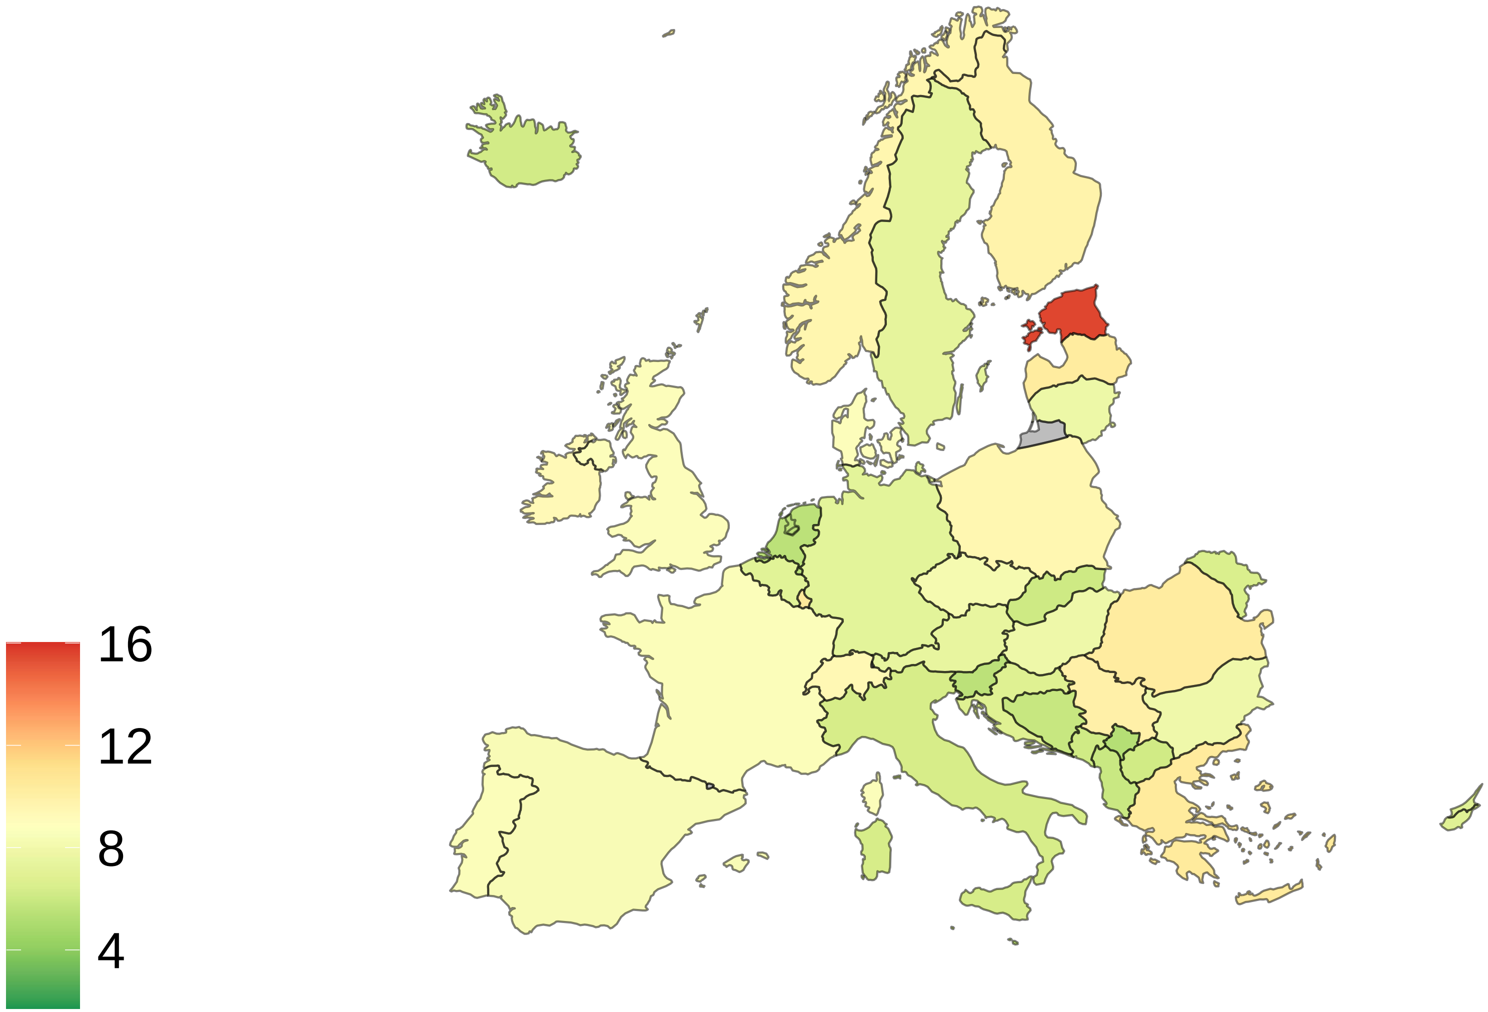

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.101: Map of top 1% posttax income share in Europe, 2007

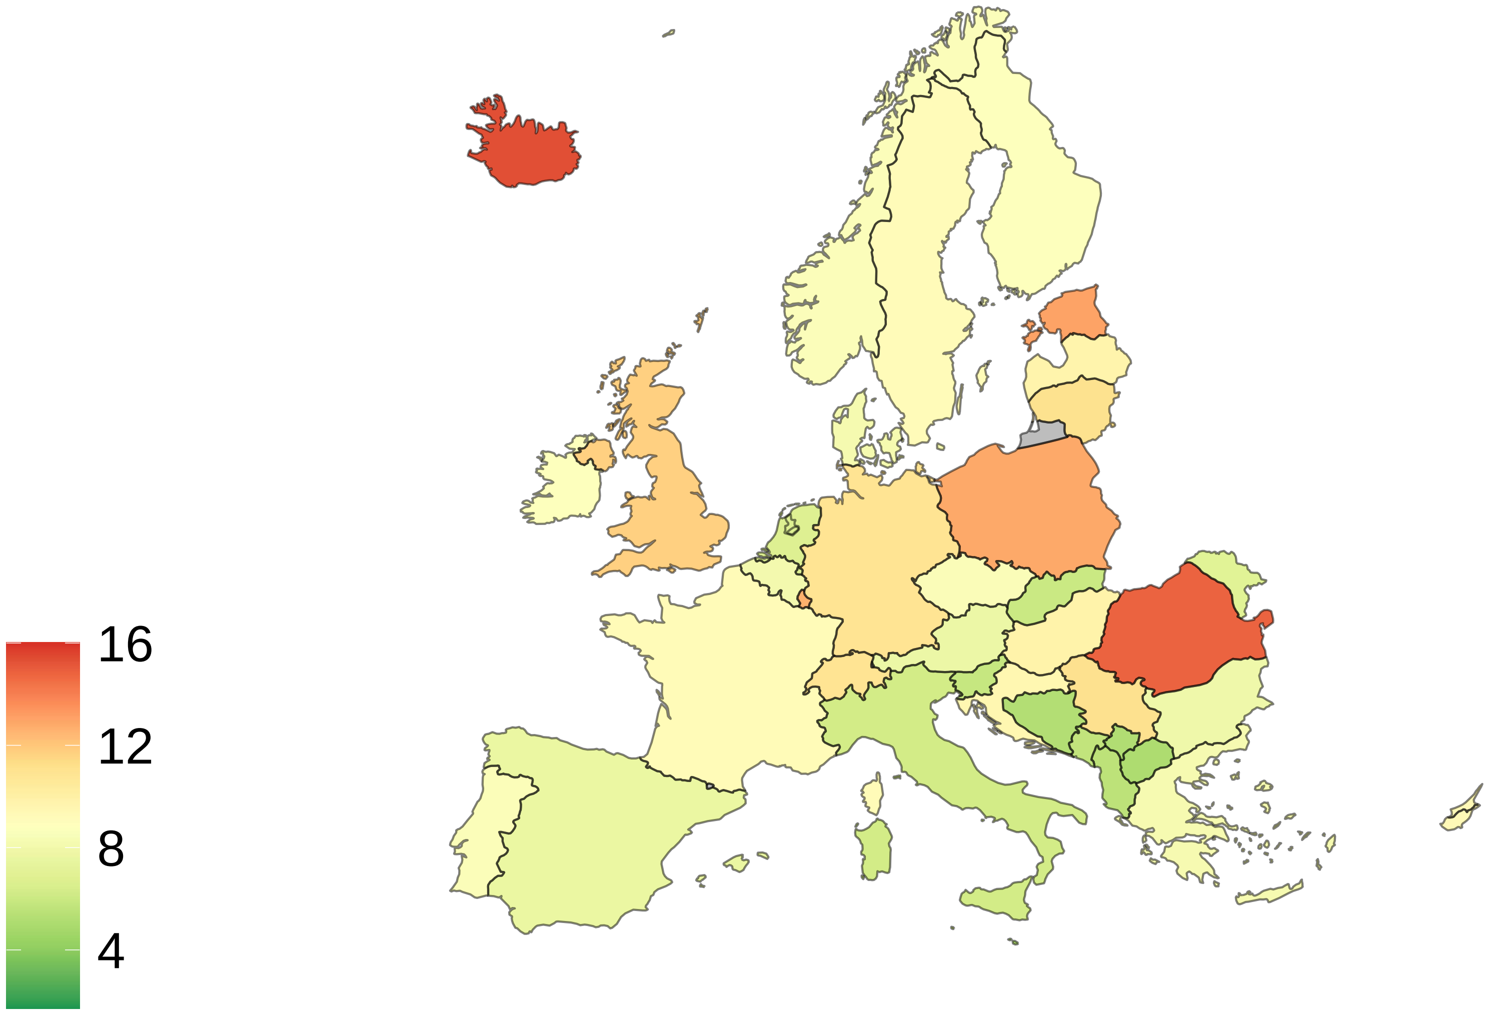

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.102: Map of top 1% posttax income share in Europe, 2017

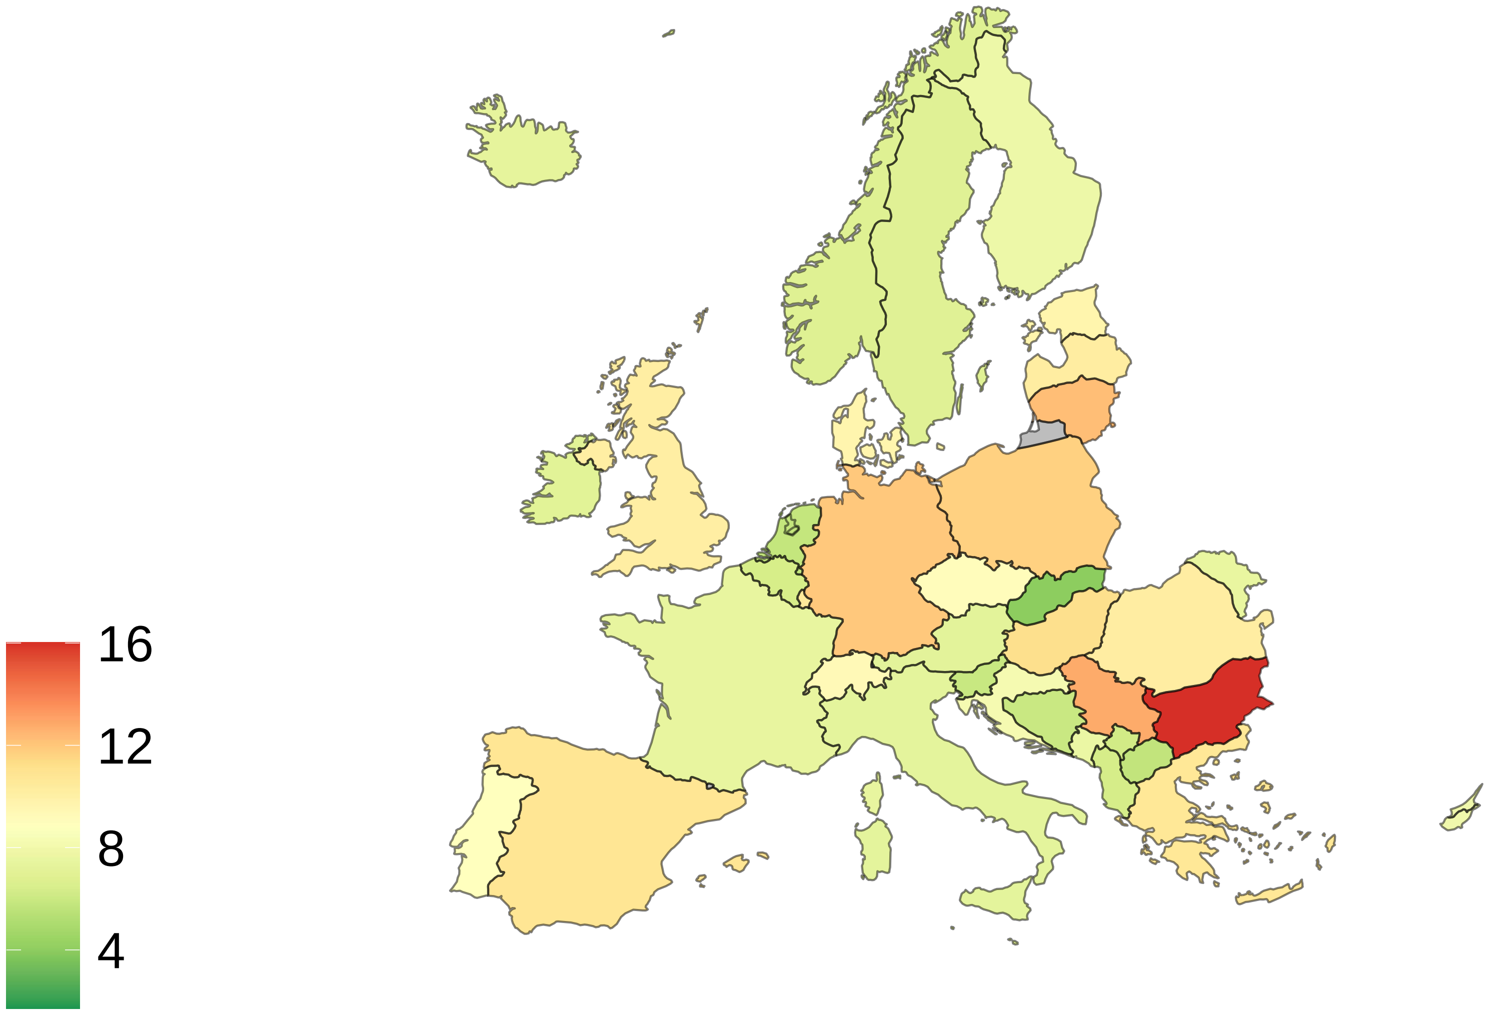

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.103: Map of bottom 50% pretax income share in Europe, 1980

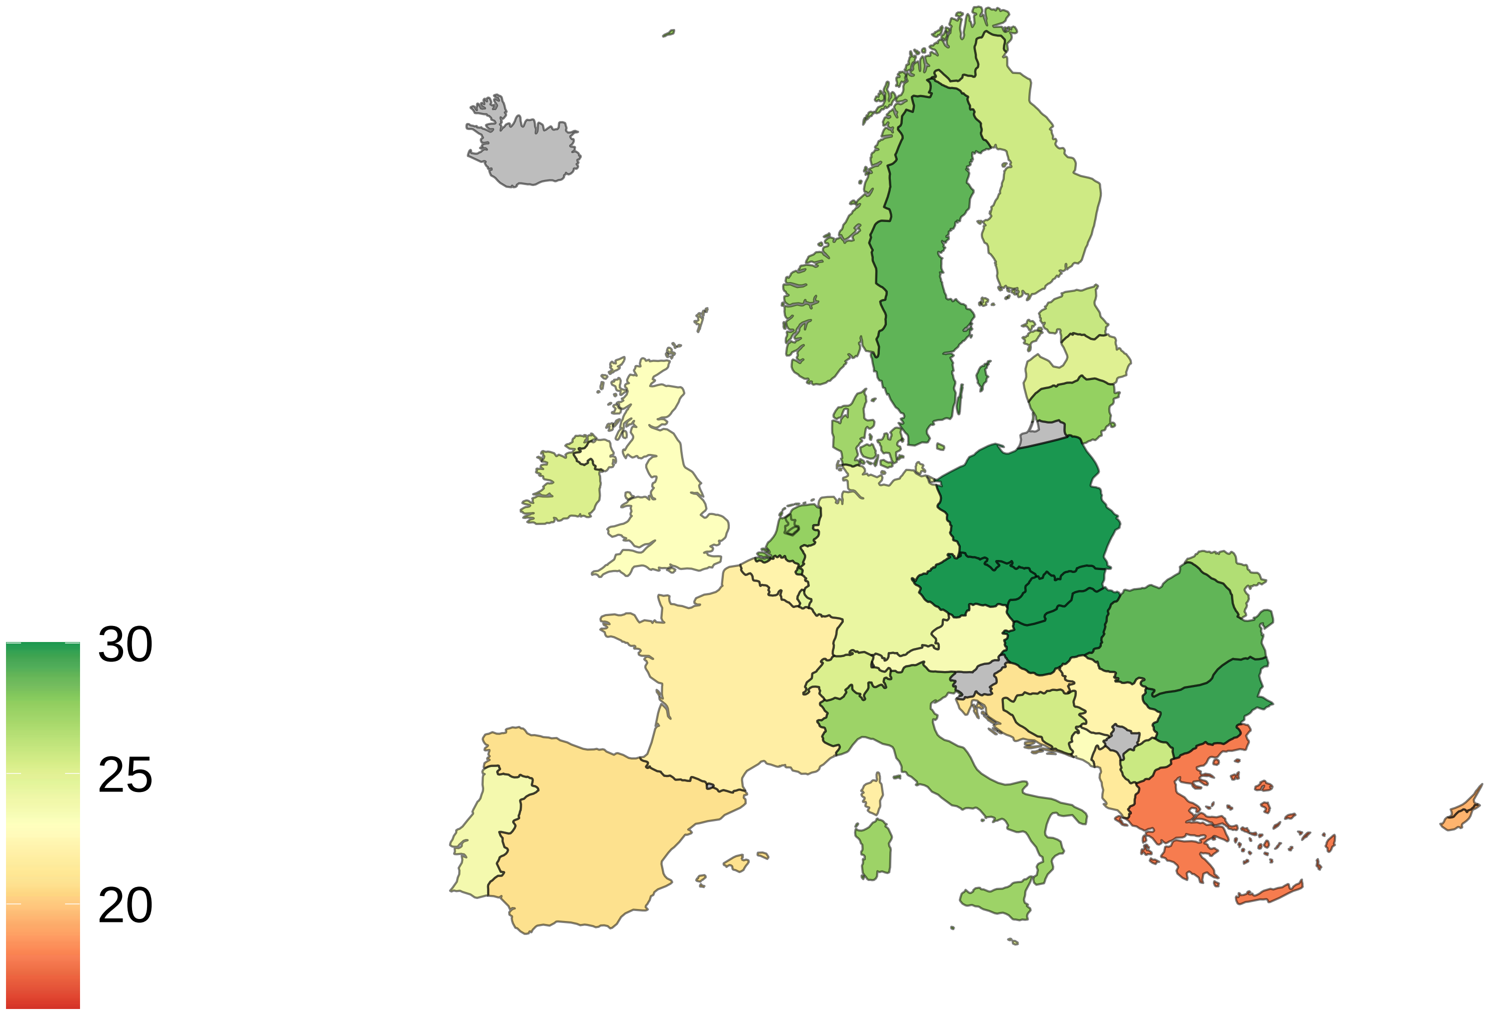

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.104: Map of bottom 50% pretax income share in Europe, 1990

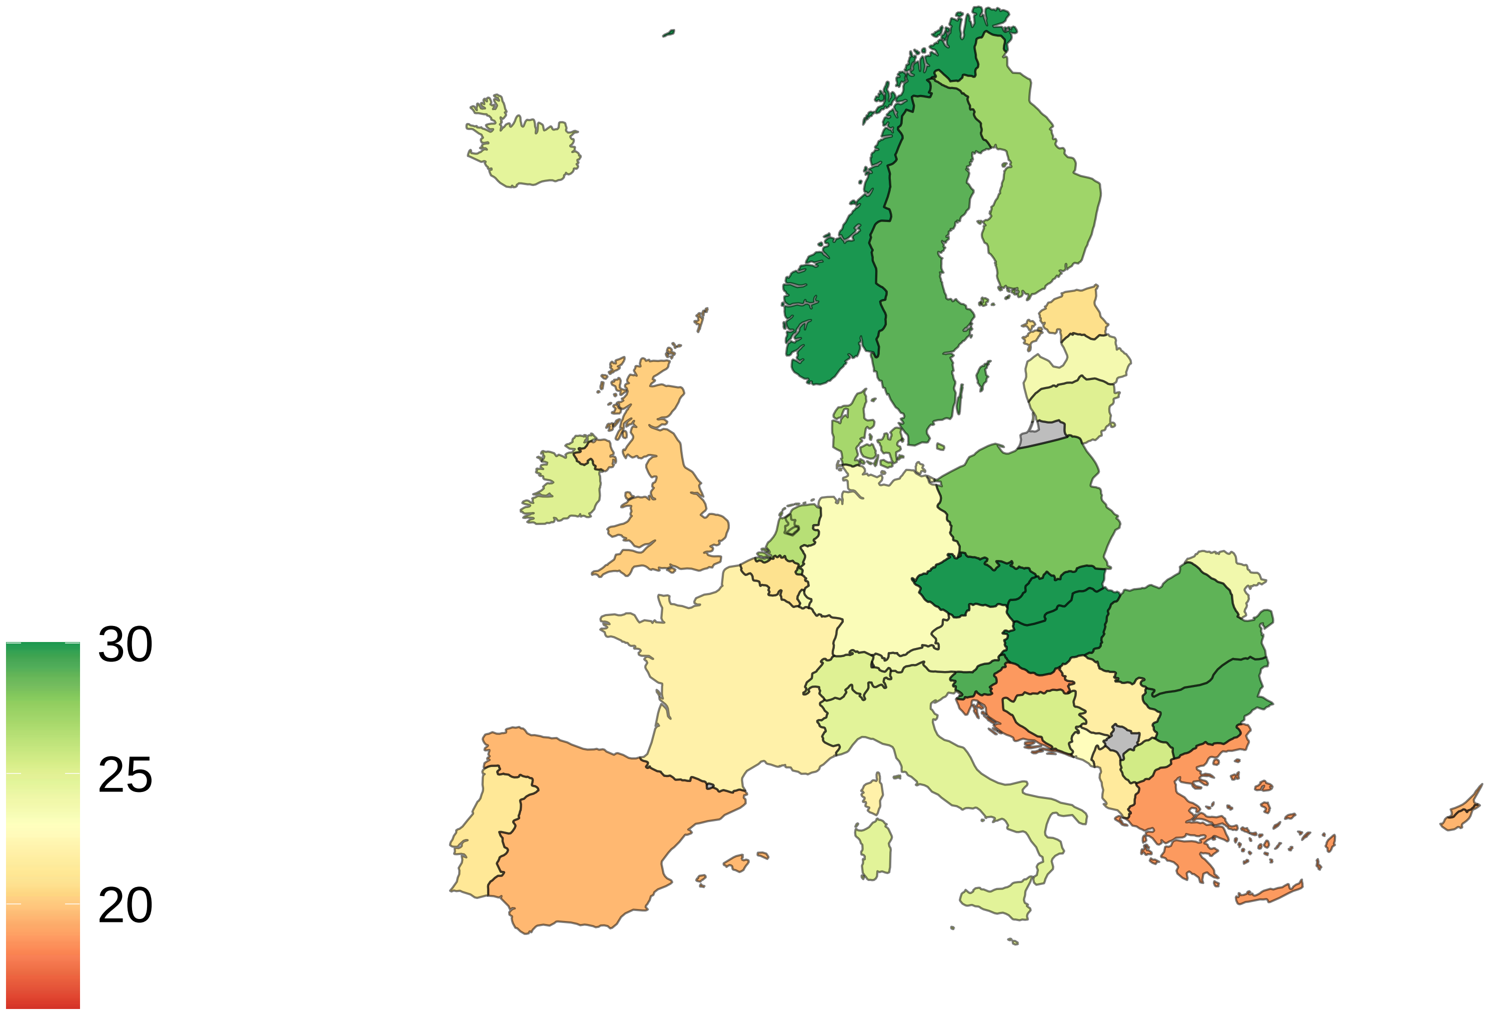

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.105: Map of bottom 50% pretax income share in Europe, 2000

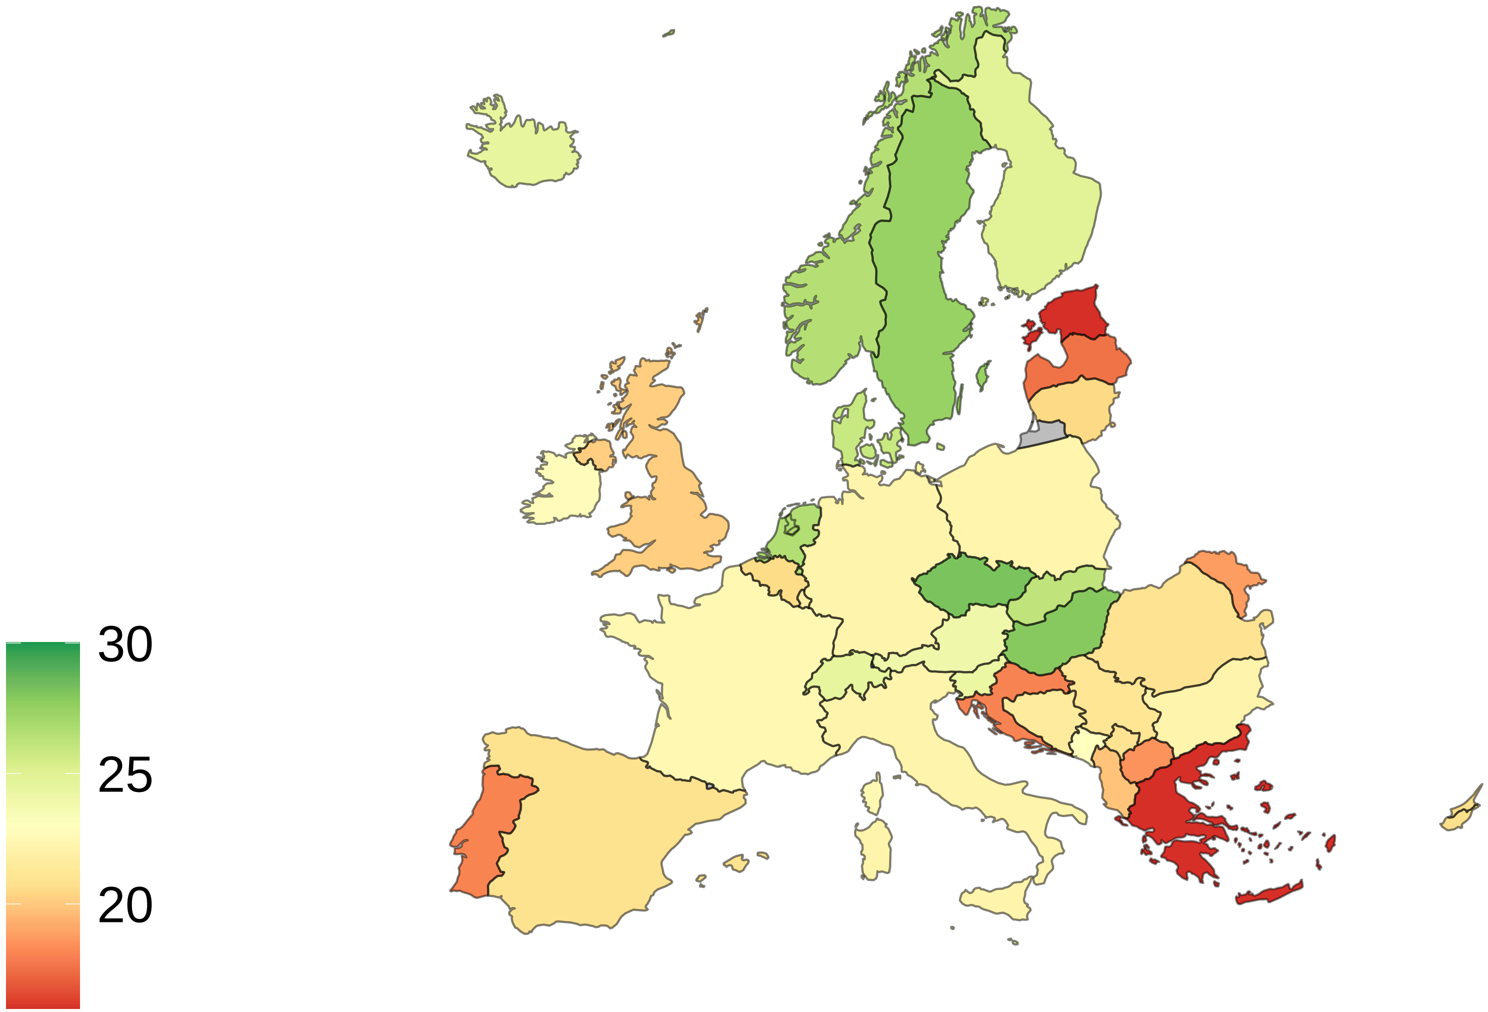

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.106: Map of bottom 50% pretax income share in Europe, 2007

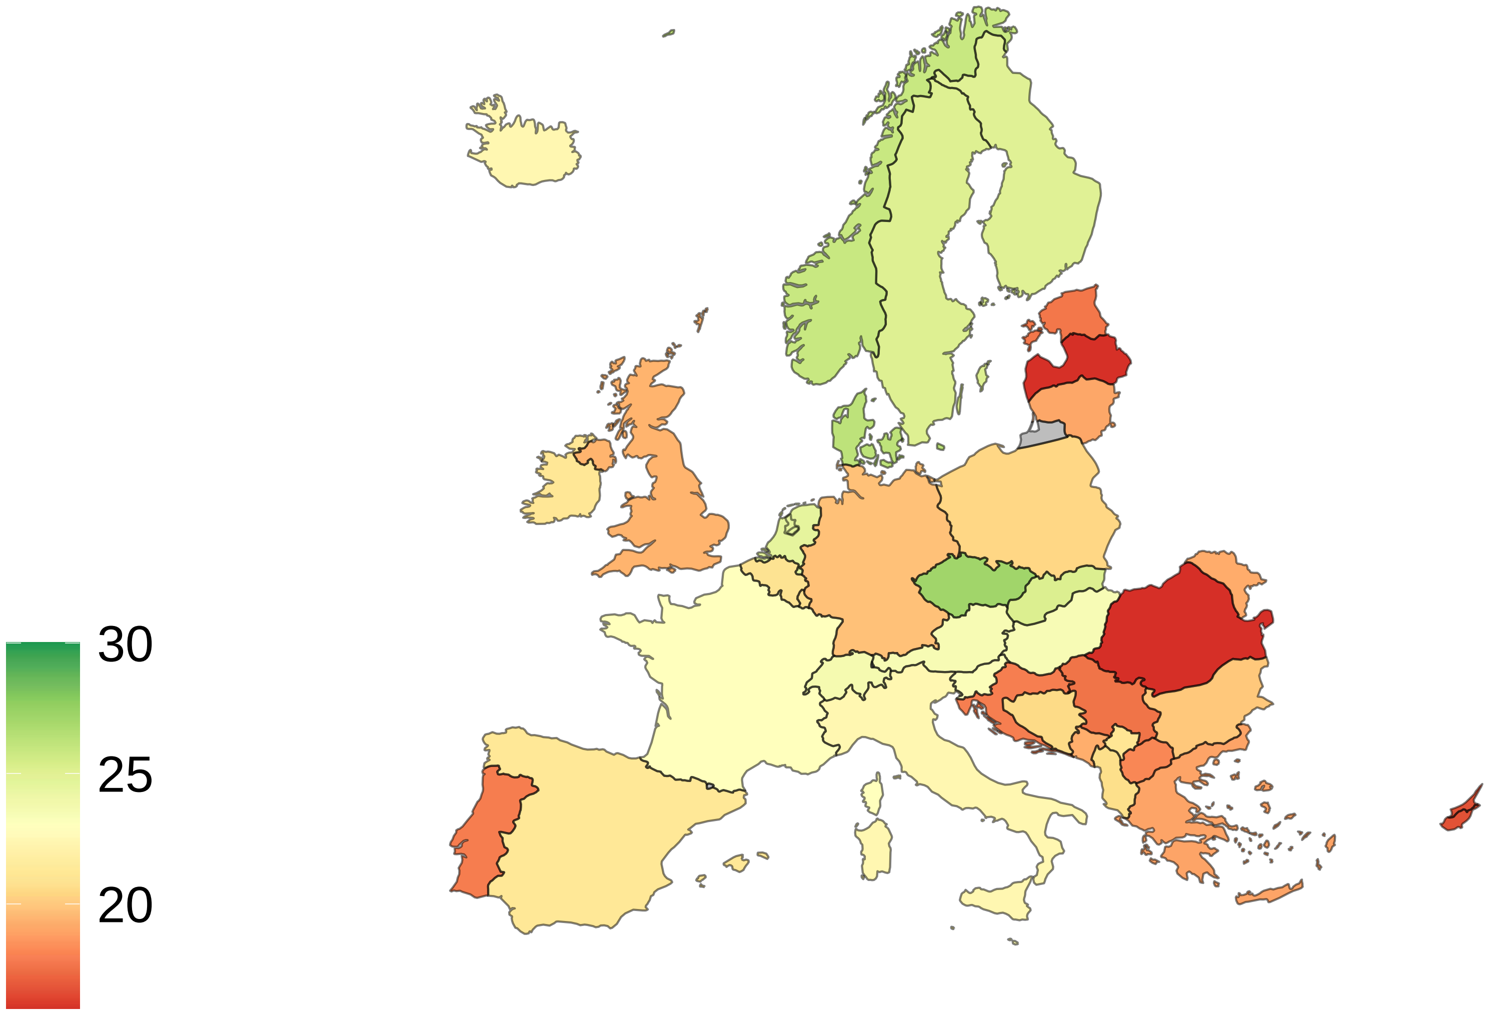

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.107: Map of bottom 50% pretax income share in Europe, 2017

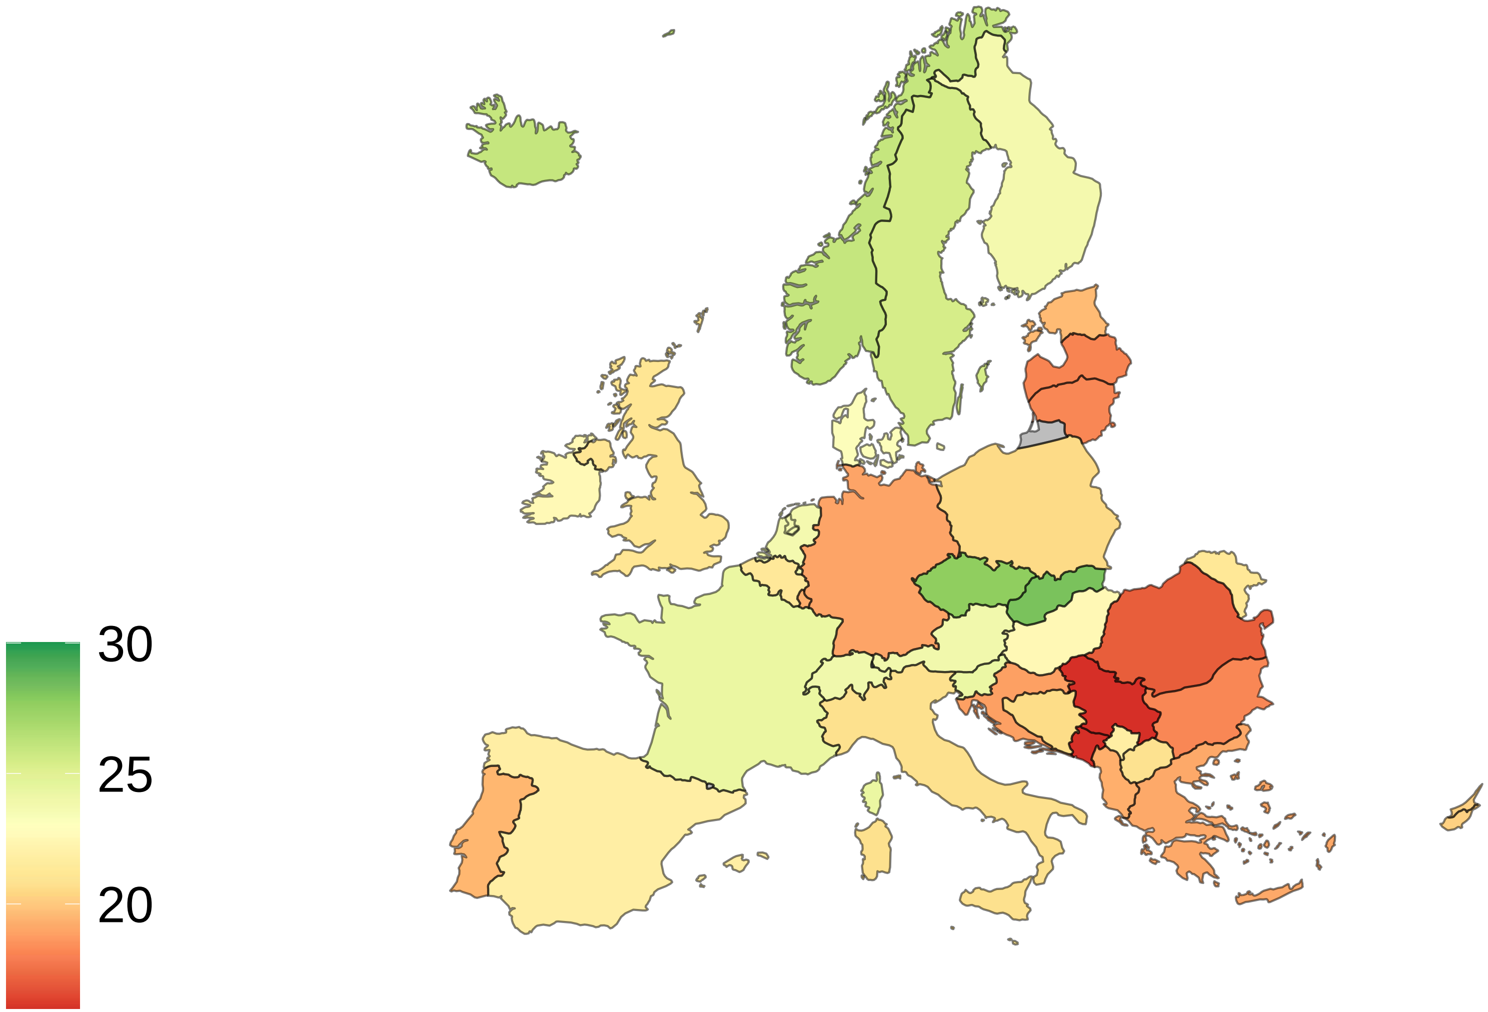

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.108: Map of bottom 50% posttax income share in Europe, 1980

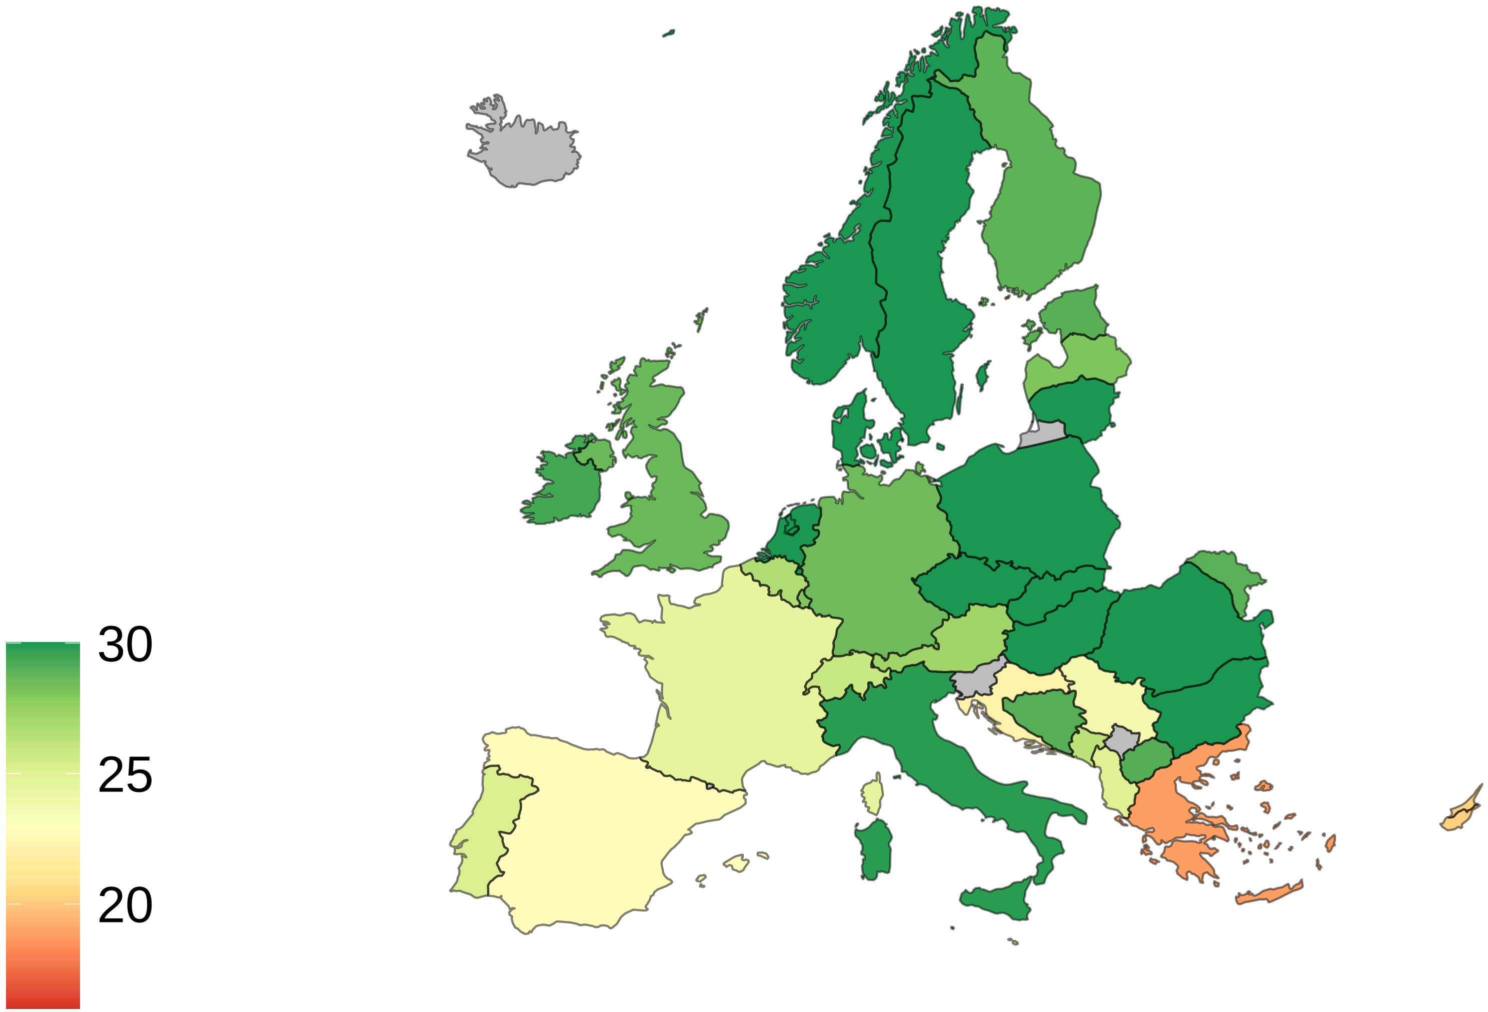

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.109: Map of bottom 50% posttax income share in Europe, 1990

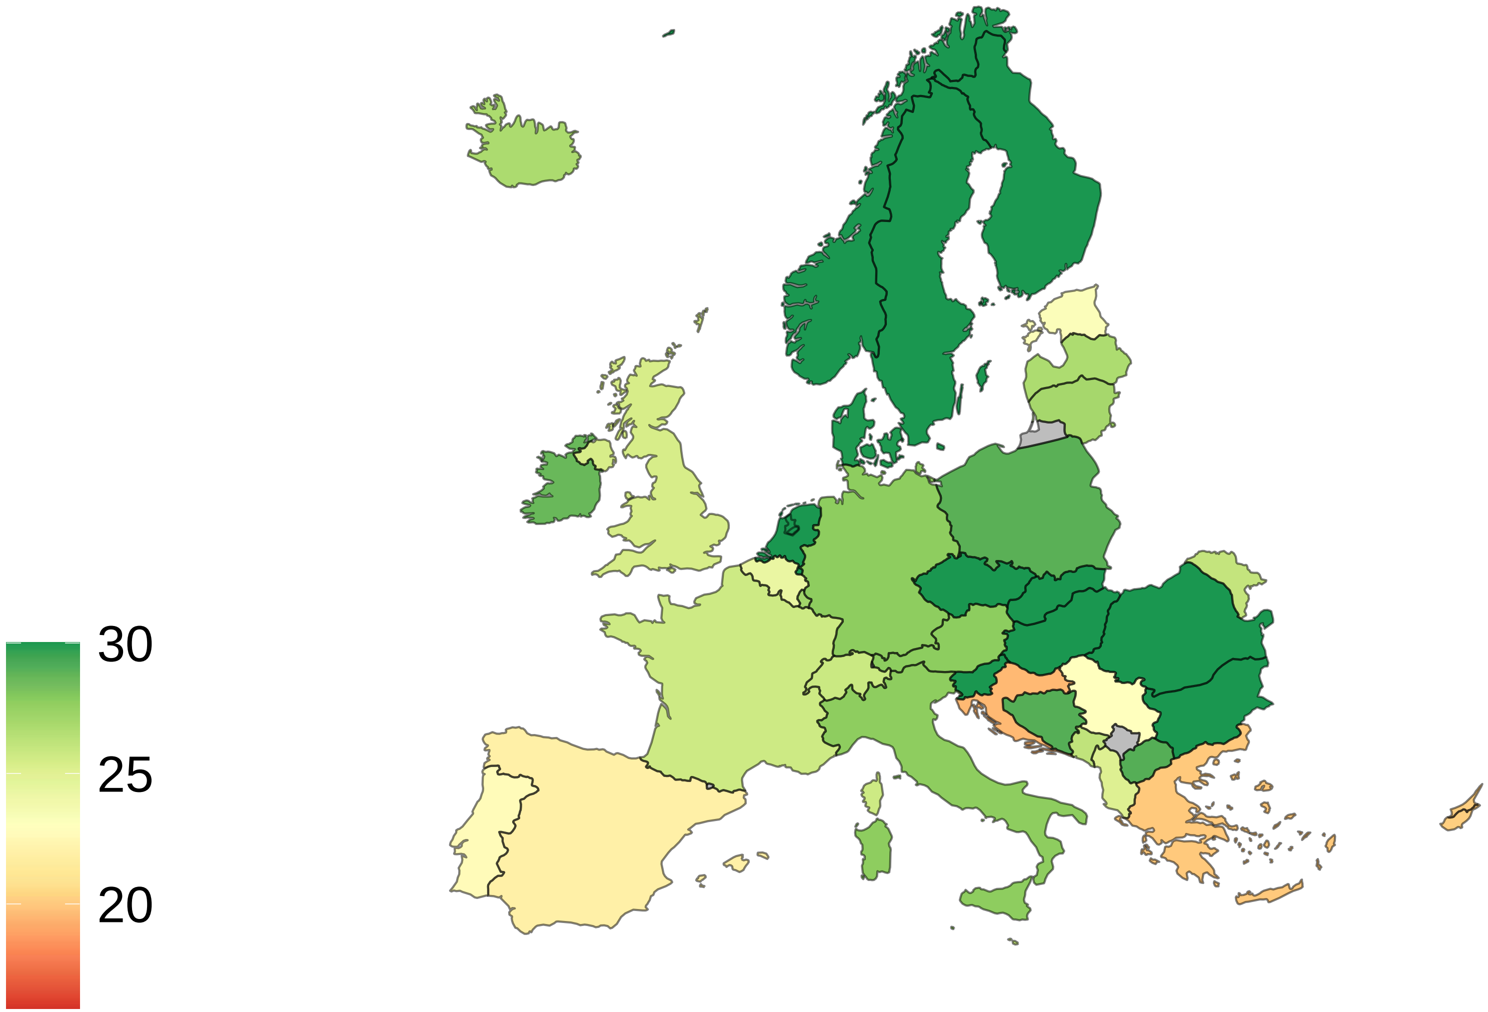

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.110: Map of bottom 50% posttax income share in Europe, 2000

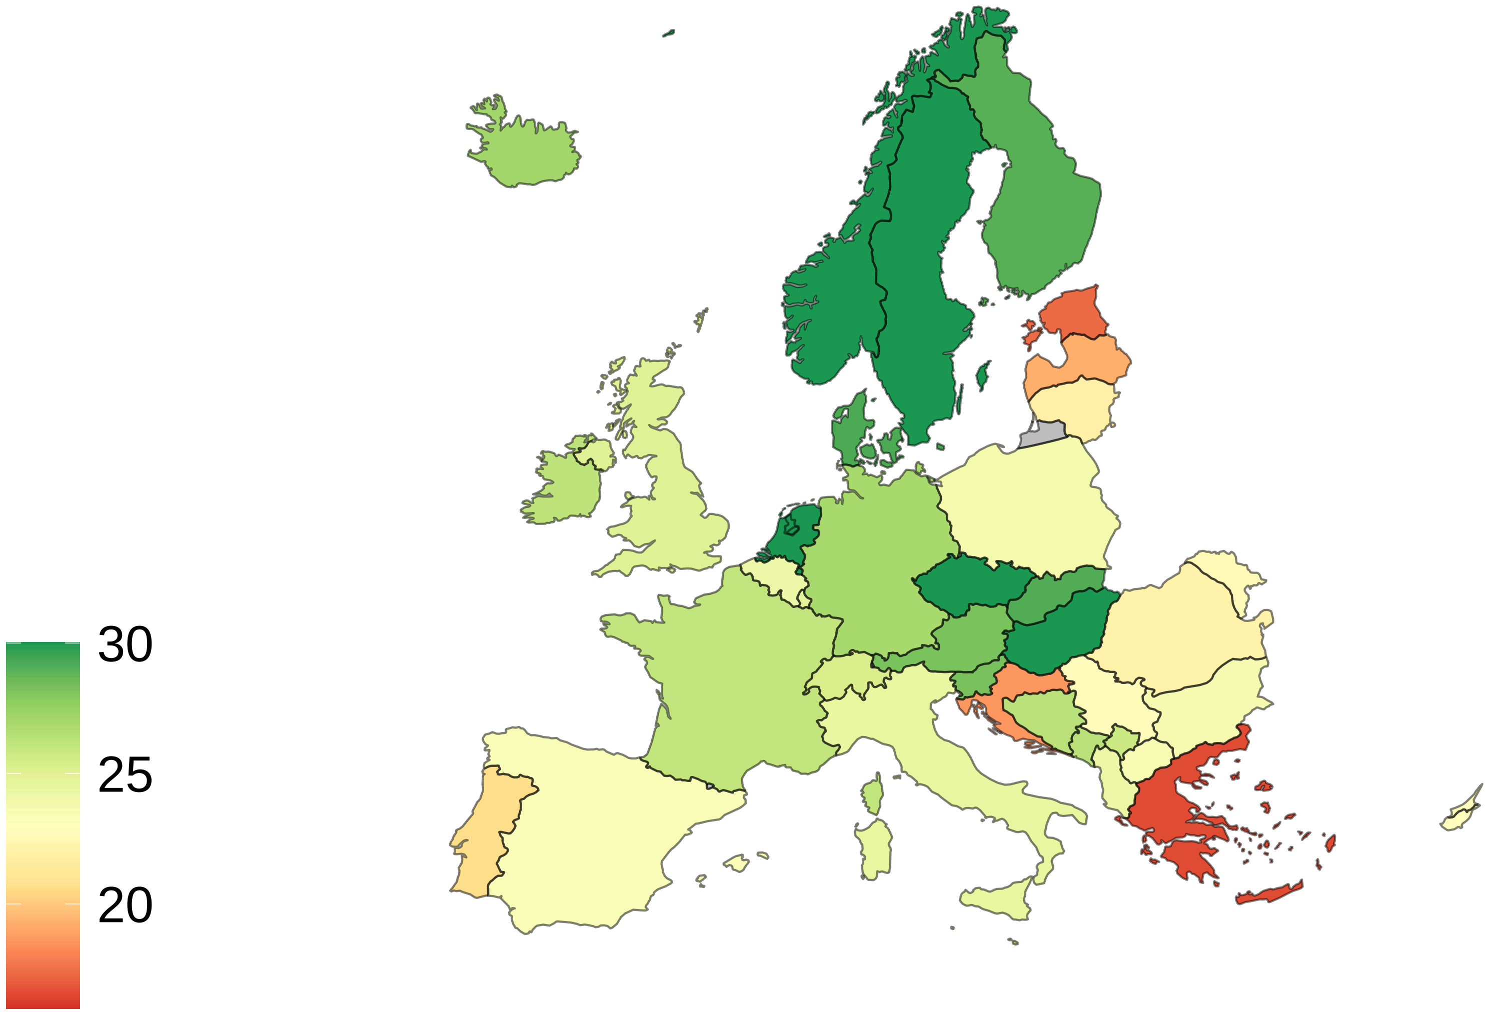

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.111: Map of bottom 50% posttax income share in Europe, 2007

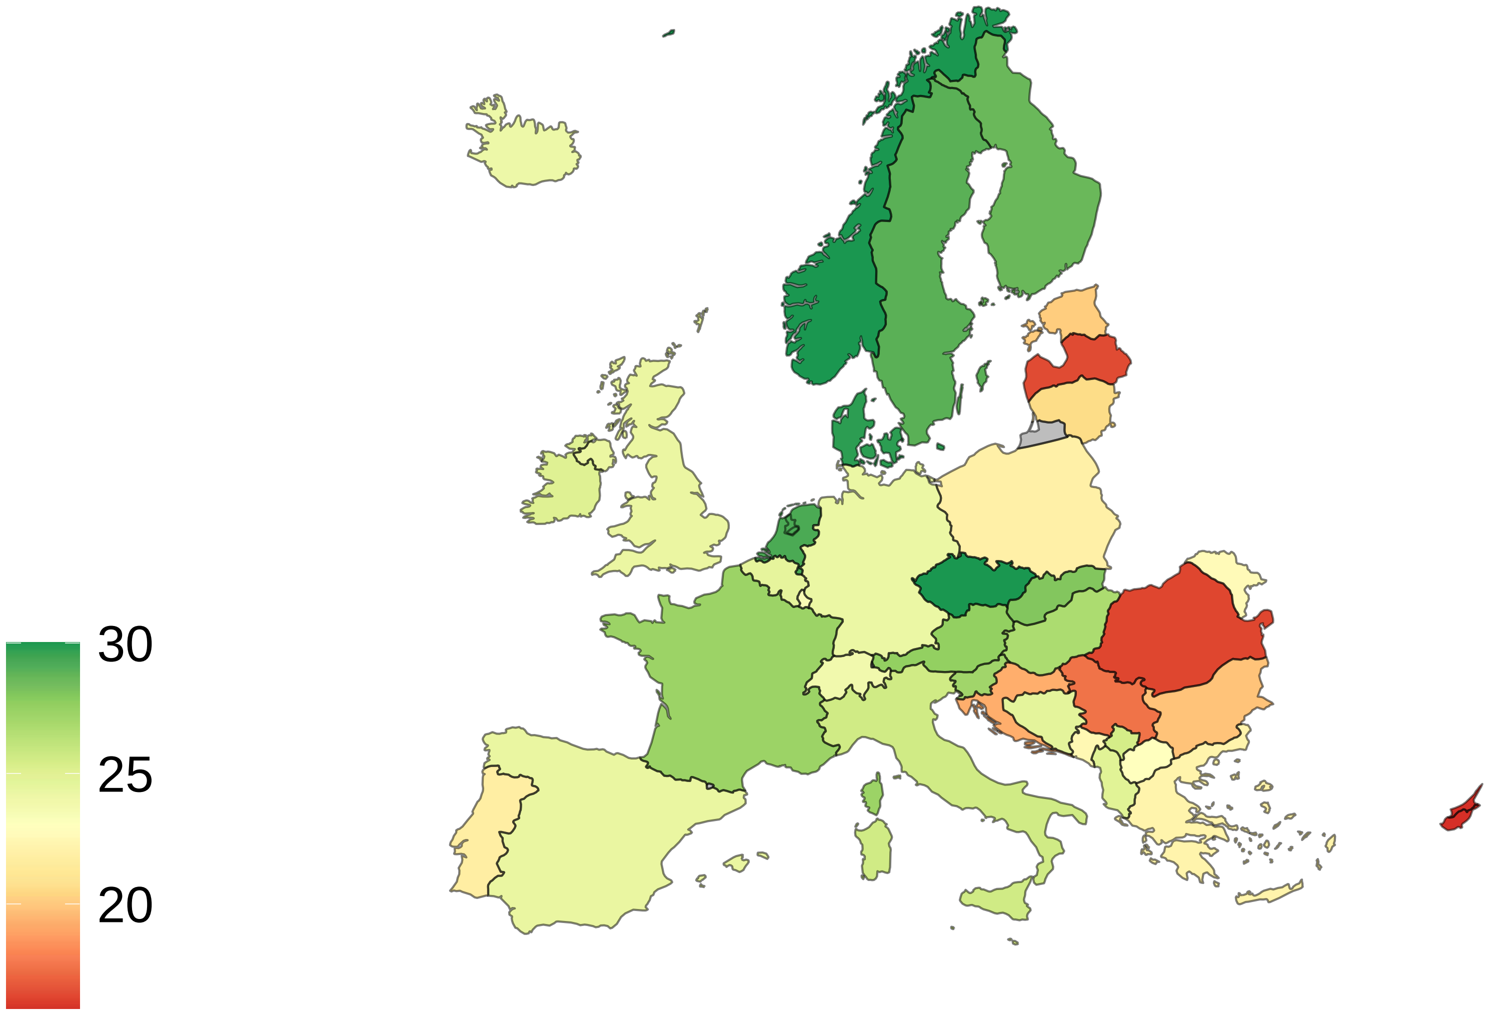

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

Figure D.112: Map of bottom 50% posttax income share in Europe, 2017

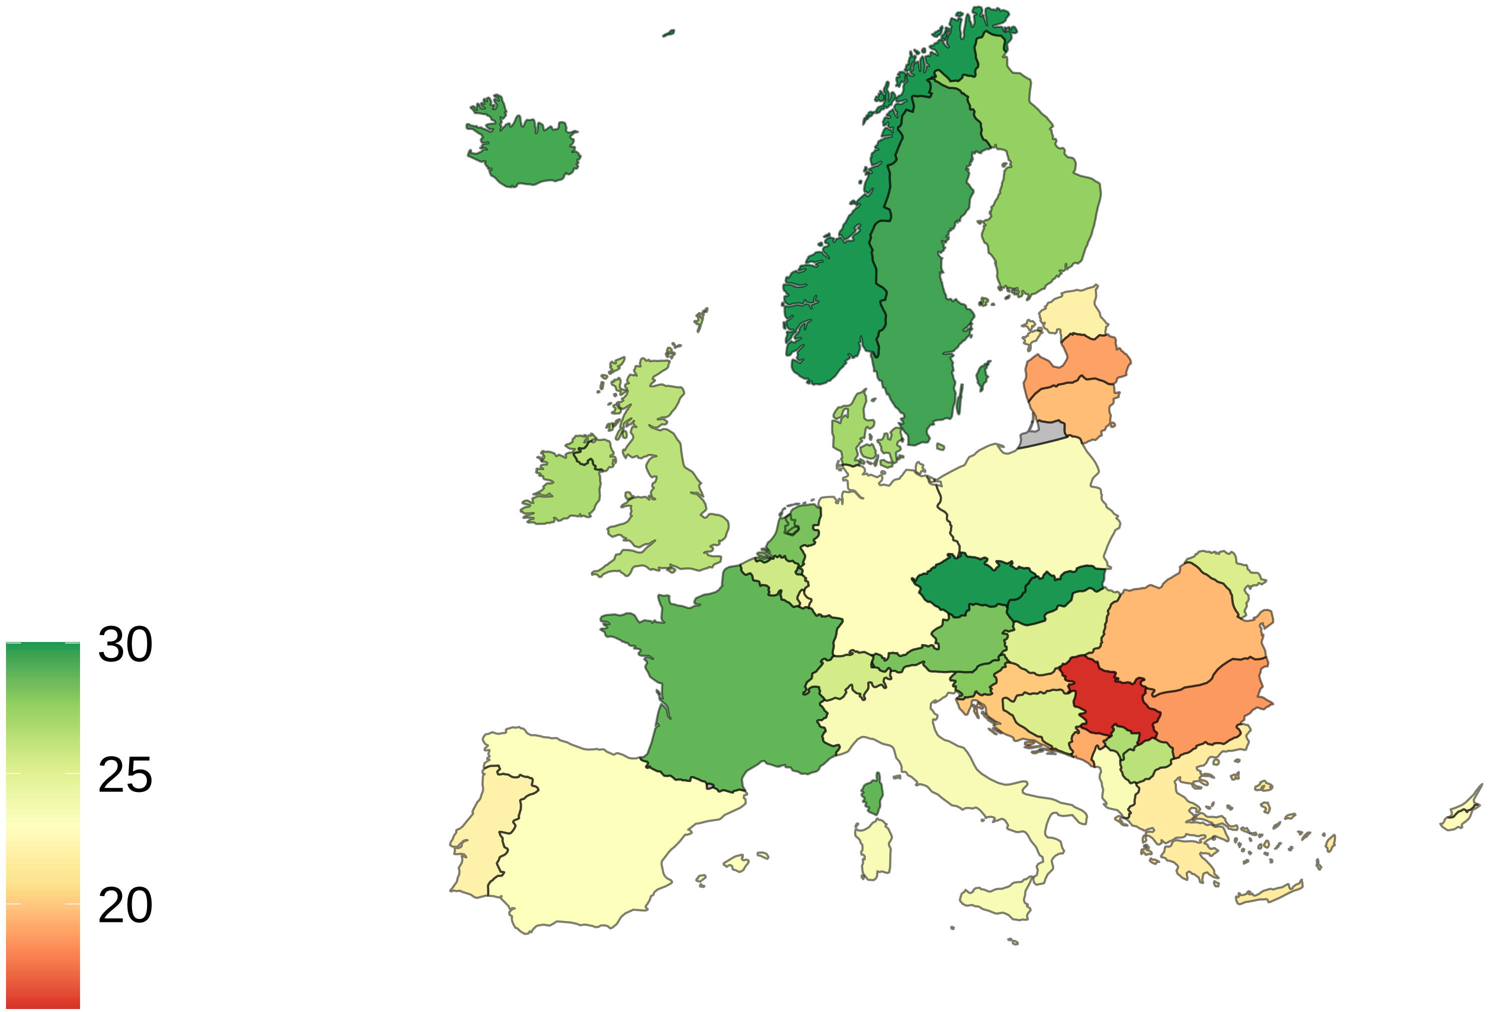

Source: Authors' computations combining surveys, tax data and national accounts. *Notes:* The unit of observation is the adult individual aged 20 or above. Income is split equally among spouses. The map includes countries with no tax data (see appendix table D.6).

D.2.7 Supplementary tables

Table D.6: Coverage of data sources (all European countries)

Country	Surveys	Tax data	Undistrib. prof.	Imp. rents	Tax data source	Quality score
Western Europe						
Austria	1987-2017	1976-2015	1995-2018	1995-2018	Altzinger et al. (2010)	Medium
Belgium	1985-2017	1990-2016	1985-2018	1985-2018	Decoster, Dobbeleer, and Maes (2017)	High
France	1989-2017	1980-2014	1980-2018	1980-2018	Garbinti, Goupille-Lebret, and Piketty (2018)	Very high
Germany	1981-2017	1980-2013	1991-2018	1991-2018	Bartels (2017a)	High
Ireland	1980-2018	1980-2015	1995-2018	1995-2018	Jäntti et al. (2007)	High
Italy	1981-2017	1980-2009	1980-2019	1980-2019	Alvaredo and Pisano (2010)	High
Luxembourg	1985-2017	2010-2012		1995-2018	Authors, from Conseil Economique et Social (2015)	High
Netherlands	1983-2017	1981-2012	1980-2018	1980-2019	Salverda and Atkinson (2007)	High
Portugal	1980-2017	1980-2005	1995-2019	1995-2019	Alvaredo (2009)	High
Spain	1980-2017	1981-2012	1995-2018	1995-2018	Alvaredo and Saez (2010)	High
Switzerland	1982-2017	1981-2014	1990-2018	1990-2018	Foellmi and Martínez (2017)	High
United Kingdom	1986-2018	1981-2014	1987-2018	1990-2018	Atkinson and Piketty (2007)	High
Northern Europe						
Denmark	1981-2017	1980-2010	1981-2018	1990-2018	Atkinson and Søgaard (2013)	High
Finland	1981-2017	1980-2009	1980-2019	1980-2019	Jäntti et al. (2010)	High
Iceland	2003-2015	1990-2016	2000-2014	2000-2014	Authors, from Statistics Iceland (2020)	High
Norway	1986-2017	1981-2011	1980-2018	1980-2018	Aaberge and Atkinson (2010)	High
Sweden	1981-2017	1980-2013	1980-2019	1980-2019	Roine and Waldenström (2010)	High
Eastern Europe						
Croatia	1983-2017	1983-2013	1997-2014	2002-2018	Kump and Novokmet (2018)	High
Czech Republic	1980-2017	1980-2015	1995-2018	1995-2018	Novokmet (2018)	High
Estonia	1988-2017	2002-2017	1994-2018	1994-2018	Authors, from Tax and Customs Board (2020)	High
Greece	1981-2017	2004-2011	1995-2018	1995-2018	Chrissis and Koutentakis (2017)	High
Hungary	1982-2017	1980-2008	1995-2018	1995-2018	Mavridis and Mosberger (2017)	High
Poland	1983-2017	1983-2015	1996-2018	1996-2018	Bukowski and Novokmet (2017b)	High
Romania	1989-2017	2013	1995-2017	1995-2019	Oancea, Andrei, and Pirjol (2017)	Medium
Serbia	1983-2017	2017	2000-2011	1997-2011	Authors, data provided by Statistical Office	Medium
Slovenia	1987-2017	1991-2012	1995-2018	1995-2018	Kump and Novokmet (2018)	High
Other Eastern						
Albania	1996-2017					Low
Bosn. & Herz.	1983-2015					Medium Low
Bulgaria	1980-2017		1999-2017	1999-2017		Medium
Cyprus	1990-2017		1995-2017	1995-2018		Medium Low
Kosovo	2003-2017					Medium Low
Latvia	1988-2017		2001-2018	1995-2018		Medium
Lithuania	1988-2017		1995-2018	1995-2018		Medium
Malta	2006-2017			2000-2018		Medium Low
Moldova	1988-2018					Low
Montenegro	1983-2015					Medium Low
Macedonia	1983-2017					Medium Low
Slovakia	1980-2017		1995-2019	1995-2019		Medium

Notes: The table shows the time coverage of the main data sources used to estimate distributional national accounts by country. Other Eastern correspond to countries not included in the main paper (countries for which no tax data was available at the time of writing).

Table D.7: Total taxes and transfers in Europe and the United States, 2007-2017
(% of national income)

	Western Europe	Northern Europe	Eastern Europe	All Europe	United States
All taxes & social contributions	**47.8%**	**51.6%**	**40.5%**	**46.5%**	**28.2%**
Social contributions	20.2%	11.7%	16.1%	18.9%	7.6%
Inc. contributory contributions	*17.6%*	*11.2%*	*14.6%*	*16.7%*	*5.3%*
Inc. non-contributory contributions	*2.5%*	*0.5%*	*1.5%*	*2.2%*	*2.2%*
Taxes	27.6%	39.9%	24.4%	27.6%	20.7%
Inc. Income & wealth taxes	*11.3%*	*17.7%*	*5.6%*	*10.4%*	*11.2%*
Inc. Corporate tax	*2.9%*	*4.1%*	*2.9%*	*3.0%*	*3.0%*
Inc. Indirect & consumption taxes	*13.4%*	*18.0%*	*16.0%*	*14.1%*	*6.5%*
All non-contributory taxes & contributions	30.2%	40.4%	25.9%	29.8%	22.9%
All transfers	**48.3%**	**51.4%**	**41.6%**	**47.1%**	**34.5%**
Cash transfers	23.6%	22.1%	18.7%	22.5%	8.8%
Inc. Pensions	*16.6%*	*15.8%*	*14.5%*	*16.1%*	*4.7%*
Inc. Unemployment & disability	*1.9%*	*1.0%*	*0.7%*	*1.6%*	*1.4%*
Inc. Other cash transfers	*5.1%*	*5.2%*	*3.5%*	*4.8%*	*2.7%*
In-kind transfers	24.7%	29.4%	22.9%	24.6%	25.7%
Inc. Health	*7.8%*	*7.8%*	*5.6%*	*7.4%*	*7.3%*
Inc. Other in-kind transfers	*16.9%*	*21.5%*	*17.3%*	*17.2%*	*18.3%*

Source: Authors' computations based on national accounts data. *Notes:* The table shows the structure of taxes and transfers in the United States and Europe, expressed as a share of national income. Values are population-weighted and averaged over the 2007-2017 period. See Appendix Table D.6 for the composition of European regions.

Table D.8: Summary measures of inequality in Europe and the US, 1980-2017

	Region	**Gini**		**Theil**		**Atkinson**	
		Pretax	*Posttax*	*Pretax*	*Posttax*	*Pretax*	*Posttax*
1980	United States	.452	.373	.441	.307	.332	.232
	Eastern Europe	.316	.291	.182	.154	.194	.161
	Northern Europe	.339	.286	.212	.151	.249	.154
	Western Europe	.403	.356	.317	.238	.295	.222
1981	United States	.459	.384	.461	.336	.346	.248
	Eastern Europe	.329	.305	.194	.167	.205	.173
	Northern Europe	.328	.278	.195	.14	.19	.139
	Western Europe	.395	.346	.306	.224	.286	.213
1982	United States	.46	.388	.468	.341	.342	.237
	Eastern Europe	.331	.307	.198	.169	.207	.175
	Northern Europe	.332	.287	.204	.154	.2	.162
	Western Europe	.395	.346	.302	.223	.287	.214
1983	United States	.47	.4	.484	.358	.363	.251
	Eastern Europe	.322	.298	.189	.16	.201	.168
	Northern Europe	.333	.283	.21	.152	.199	.147
	Western Europe	.396	.347	.306	.225	.293	.211
1984	United States	.48	.417	.525	.404	.368	.267
	Eastern Europe	.319	.296	.186	.16	.193	.169

Table D.8: Summary measures of inequality in Europe and the US, 1980-2017

	Region	**Gini**		**Theil**		**Atkinson**	
		Pretax	*Posttax*	*Pretax*	*Posttax*	*Pretax*	*Posttax*
	Northern Europe	.336	.289	.226	.172	.194	.154
	Western Europe	.398	.349	.308	.228	.289	.216
1985	United States	.48	.415	.523	.397	.369	.267
	Eastern Europe	.316	.293	.181	.156	.188	.164
	Northern Europe	.337	.292	.229	.172	.196	.149
	Western Europe	.402	.351	.317	.23	.297	.218
1986	United States	.482	.412	.515	.372	.378	.264
	Eastern Europe	.318	.296	.184	.159	.193	.174
	Northern Europe	.326	.284	.205	.156	.183	.149
	Western Europe	.407	.359	.332	.245	.302	.233
1987	United States	.492	.417	.554	.403	.388	.271
	Eastern Europe	.31	.293	.174	.156	.165	.164
	Northern Europe	.34	.287	.224	.162	.232	.141
	Western Europe	.408	.354	.338	.24	.303	.22
1988	United States	.502	.43	.619	.461	.398	.285
	Eastern Europe	.312	.288	.176	.15	.206	.165
	Northern Europe	.33	.29	.214	.167	.19	.148
	Western Europe	.413	.359	.35	.249	.305	.221
1989	United States	.501	.425	.599	.441	.397	.278

Table D.8: Summary measures of inequality in Europe and the US, 1980-2017

	Region	Gini		Theil		Atkinson	
		Pretax	*Posttax*	*Pretax*	*Posttax*	*Pretax*	*Posttax*
	Eastern Europe	.322	.306	.189	.172	.184	.185
	Northern Europe	.333	.289	.217	.165	.22	.158
	Western Europe	.415	.362	.356	.256	.32	.226
1990	United States	.502	.424	.603	.44	.402	.278
	Eastern Europe	.35	.33	.225	.197	.228	.204
	Northern Europe	.324	.276	.196	.142	.193	.139
	Western Europe	.419	.363	.354	.252	.311	.224
1991	United States	.503	.423	.588	.425	.416	.277
	Eastern Europe	.354	.335	.237	.21	.218	.224
	Northern Europe	.322	.275	.192	.139	.182	.129
	Western Europe	.418	.364	.349	.253	.311	.228
1992	United States	.512	.432	.631	.459	.412	.287
	Eastern Europe	.364	.342	.263	.224	.237	.232
	Northern Europe	.326	.278	.199	.146	.189	.132
	Western Europe	.414	.368	.334	.257	.3	.249
1993	United States	.511	.427	.621	.44	.411	.282
	Eastern Europe	.38	.357	.29	.244	.261	.248
	Northern Europe	.332	.292	.213	.166	.187	.159
	Western Europe	.417	.368	.341	.255	.311	.268

Table D.8: Summary measures of inequality in Europe and the US, 1980-2017

	Region	Gini		Theil		Atkinson	
		Pretax	*Posttax*	*Pretax*	*Posttax*	*Pretax*	*Posttax*
1994	United States	.513	.428	.624	.438	.41	.283
	Eastern Europe	.393	.373	.322	.273	.282	.265
	Northern Europe	.358	.31	.266	.199	.244	.164
	Western Europe	.421	.374	.35	.261	.318	.248
1995	United States	.521	.435	.652	.456	.419	.291
	Eastern Europe	.408	.386	.361	.301	.298	.279
	Northern Europe	.354	.312	.269	.211	.209	.171
	Western Europe	.419	.373	.349	.266	.312	.245
1996	United States	.534	.441	.696	.475	.457	.296
	Eastern Europe	.412	.389	.352	.3	.305	.262
	Northern Europe	.361	.313	.278	.21	.253	.172
	Western Europe	.425	.375	.365	.268	.322	.249
1997	United States	.54	.446	.725	.498	.456	.303
	Eastern Europe	.418	.396	.37	.325	.309	.282
	Northern Europe	.371	.32	.306	.233	.24	.177
	Western Europe	.428	.375	.378	.278	.319	.254
1998	United States	.541	.448	.734	.5	.456	.306
	Eastern Europe	.425	.395	.391	.314	.322	.271
	Northern Europe	.361	.318	.292	.228	.216	.173

Table D.8: Summary measures of inequality in Europe and the US, 1980-2017

	Region	**Gini**		**Theil**		**Atkinson**	
		Pretax	*Posttax*	*Pretax*	*Posttax*	*Pretax*	*Posttax*
	Western Europe	.431	.38	.388	.287	.317	.258
1999	United States	.547	.452	.77	.525	.456	.313
	Eastern Europe	.43	.405	.393	.341	.324	.288
	Northern Europe	.368	.325	.315	.252	.225	.183
	Western Europe	.434	.382	.395	.289	.337	.254
2000	United States	.551	.456	.797	.542	.466	.32
	Eastern Europe	.435	.412	.394	.339	.345	.302
	Northern Europe	.367	.315	.317	.239	.222	.168
	Western Europe	.433	.382	.392	.291	.33	.26
2001	United States	.542	.453	.752	.53	.456	.316
	Eastern Europe	.434	.412	.388	.337	.34	.313
	Northern Europe	.361	.308	.284	.211	.236	.165
	Western Europe	.433	.379	.388	.282	.327	.257
2002	United States	.543	.457	.745	.54	.465	.322
	Eastern Europe	.444	.425	.415	.358	.363	.309
	Northern Europe	.361	.307	.29	.217	.231	.162
	Western Europe	.432	.379	.383	.283	.321	.249
2003	United States	.542	.461	.749	.552	.444	.329
	Eastern Europe	.454	.437	.433	.377	.38	.328

Table D.8: Summary measures of inequality in Europe and the US, 1980-2017

	Region	Gini		Theil		Atkinson	
		Pretax	*Posttax*	*Pretax*	*Posttax*	*Pretax*	*Posttax*
	Northern Europe	.363	.305	.298	.226	.268	.163
	Western Europe	.432	.379	.389	.292	.325	.248
2004	United States	.551	.466	.8	.581	.456	.334
	Eastern Europe	.461	.432	.471	.385	.4	.327
	Northern Europe	.374	.314	.332	.267	.285	.169
	Western Europe	.434	.379	.401	.296	.32	.246
2005	United States	.56	.469	.853	.604	.467	.335
	Eastern Europe	.472	.437	.512	.4	.438	.329
	Northern Europe	.384	.338	.367	.334	.256	.2
	Western Europe	.44	.385	.418	.303	.33	.252
2006	United States	.569	.475	.888	.625	.482	.344
	Eastern Europe	.477	.441	.52	.42	.428	.33
	Northern Europe	.386	.327	.356	.255	.249	.183
	Western Europe	.444	.384	.432	.31	.332	.252
2007	United States	.57	.469	.891	.606	.482	.332
	Eastern Europe	.484	.447	.565	.447	.437	.337
	Northern Europe	.382	.326	.347	.257	.248	.181
	Western Europe	.448	.388	.449	.328	.339	.251
2008	United States	.563	.463	.871	.608	.473	.332

Table D.8: Summary measures of inequality in Europe and the US, 1980-2017

	Region	**Gini**		**Theil**		**Atkinson**	
		Pretax	*Posttax*	*Pretax*	*Posttax*	*Pretax*	*Posttax*
	Eastern Europe	.479	.449	.55	.461	.424	.331
	Northern Europe	.389	.333	.349	.253	.262	.191
	Western Europe	.441	.384	.431	.319	.336	.249
2009	United States	.554	.469	.835	.621	.456	.347
	Eastern Europe	.47	.44	.511	.425	.418	.326
	Northern Europe	.375	.317	.294	.205	.259	.17
	Western Europe	.444	.38	.424	.299	.348	.241
2010	United States	.567	.473	.902	.658	.469	.35
	Eastern Europe	.461	.436	.471	.399	.402	.326
	Northern Europe	.397	.331	.347	.235	.317	.182
	Western Europe	.441	.382	.417	.305	.346	.246
2011	United States	.571	.477	.886	.646	.481	.356
	Eastern Europe	.468	.437	.49	.402	.418	.327
	Northern Europe	.393	.334	.335	.234	.272	.189
	Western Europe	.444	.384	.428	.31	.354	.25
2012	United States	.581	.489	.946	.689	.494	.369
	Eastern Europe	.471	.444	.504	.414	.423	.342
	Northern Europe	.39	.326	.333	.222	.312	.179
	Western Europe	.444	.385	.419	.305	.357	.252

Table D.8: Summary measures of inequality in Europe and the US, 1980-2017

	Region	**Gini**		**Theil**		**Atkinson**	
		Pretax	*Posttax*	*Pretax*	*Posttax*	*Pretax*	*Posttax*
2013	United States	.573	.479	.888	.628	.477	.356
	Eastern Europe	.469	.424	.494	.382	.42	.306
	Northern Europe	.395	.33	.337	.23	.318	.183
	Western Europe	.451	.39	.431	.312	.38	.259
2014	United States	.579	.482	.918	.64	.49	.359
	Eastern Europe	.475	.45	.519	.44	.441	.347
	Northern Europe	.398	.335	.342	.233	.317	.19
	Western Europe	.452	.394	.438	.327	.377	.27
2015	United States	.596	.5	.942	.655	.517	.388
	Eastern Europe	.475	.433	.538	.416	.44	.318
	Northern Europe	.399	.334	.339	.235	.314	.187
	Western Europe	.451	.392	.44	.323	.373	.267
2016	United States	.594	.501	.933	.661	.523	.397
	Eastern Europe	.461	.429	.49	.398	.408	.317
	Northern Europe	.394	.331	.328	.228	.311	.184
	Western Europe	.451	.392	.442	.33	.368	.267
2017	United States	.593	.499	.963	.66	.523	.408
	Eastern Europe	.461	.424	.492	.388	.407	.305
	Northern Europe	.397	.333	.336	.233	.319	.186

Table D.8: Summary measures of inequality in Europe and the US, 1980-2017

Region	**Gini**		**Theil**		**Atkinson**	
	Pretax	*Posttax*	*Pretax*	*Posttax*	*Pretax*	*Posttax*
Western Europe	.448	.394	.441	.334	.364	.271

Sources. Authors' computations combining surveys, tax data and national accounts for European countries and Piketty, Saez, and Zucman, 2018 for the US. *Notes.* The table presents summary pretax and posttax income inequality statistics in Eastern Europe, Northern Europe, Western Europe, and the United States. See Table D.6 for the composition of European regions. The Atkinson parameter is set to 1.

Table D.9: Performance of European countries and the United States in reaching SDG 10.1, 1980-2017

	1980-2017			2007-2017		
	Bottom 40%	Average	Difference	Bottom 40%	Average	Difference
Austria	1.2 %	1.1 %	0.1 p.p.	-0.1 %	-0.2 %	0.1 p.p.
Belgium	1.1 %	1.2 %	-0.1 p.p.	0.2 %	0.2 %	0.0 p.p.
Switzerland	0.5 %	0.6 %	-0.2 p.p.	0.4 %	0.3 %	0.0 p.p.
Czech Republic	0.3 %	1.0 %	-0.7 p.p.	1.4 %	1.2 %	0.2 p.p.
Germany	0.0 %	0.8 %	-0.8 p.p.	0.2 %	0.7 %	-0.6 p.p.
Denmark	1.0 %	1.5 %	-0.5 p.p.	-1.0 %	0.4 %	-1.4 p.p.
Estonia	1.2 %	2.0 %	-0.8 p.p.	2.1 %	1.0 %	1.2 p.p.
Spain	1.4 %	1.2 %	0.2 p.p.	0.7 %	0.4 %	0.3 p.p.
Finland	1.3 %	1.5 %	-0.2 p.p.	-0.9 %	-0.5 %	-0.4 p.p.
France	1.3 %	0.9 %	0.4 p.p.	0.4 %	-0.2 %	0.6 p.p.
United Kingdom	1.8 %	2.0 %	-0.2 p.p.	1.0 %	0.0 %	0.9 p.p.
Greece	0.0 %	-0.1 %	0.1 p.p.	-3.6 %	-3.4 %	-0.2 p.p.
Croatia	-0.2 %	0.1 %	-0.4 p.p.	0.5 %	0.1 %	0.4 p.p.
Hungary	-0.8 %	0.9 %	-1.7 p.p.	0.7 %	1.5 %	-0.8 p.p.
Ireland	1.6 %	1.9 %	-0.3 p.p.	0.5 %	-0.5 %	0.9 p.p.
Iceland	1.8 %	1.6 %	0.2 p.p.	2.3 %	0.6 %	1.7 p.p.
Italy	-0.5 %	0.4 %	-0.9 p.p.	-2.2 %	-1.3 %	-0.9 p.p.
Luxembourg	1.8 %	2.6 %	-0.8 p.p.	-4.0 %	-2.9 %	-1.2 p.p.

Table D.9: Performance of European countries and the United States in reaching SDG 10.1, 1980-2017

	1980-2017			2007-2017		
	Bottom 40%	Average	Difference	Bottom 40%	Average	Difference
Netherlands	0.4 %	0.9 %	-0.5 p.p.	-0.3 %	0.2 %	-0.5 p.p.
Norway	2.3 %	2.4 %	-0.1 p.p.	0.9 %	0.9 %	-0.0 p.p.
Poland	0.7 %	2.0 %	-1.3 p.p.	3.0 %	3.0 %	-0.1 p.p.
Portugal	0.7 %	1.3 %	-0.6 p.p.	0.8 %	-0.1 %	0.9 p.p.
Romania	-0.4 %	1.3 %	-1.7 p.p.	4.1 %	2.8 %	1.3 p.p.
Serbia	-2.3 %	-1.0 %	-1.3 p.p.	-1.0 %	0.9 %	-1.9 p.p.
Sweden	1.4 %	1.8 %	-0.4 p.p.	1.1 %	1.1 %	0.0 p.p.
Slovenia	-0.1 %	0.5 %	-0.5 p.p.	0.8 %	0.2 %	0.6 p.p.
United States	-0.3 %	1.4 %	-1.6 p.p.	-1.4 %	0.4 %	-1.9 p.p.

Source. Authors' computations combining surveys, tax data and national accounts. *Notes.* The table shows the average annual real growth of the pretax income of the bottom 40%, the average annual real growth of the average national income per adult, and the percentage points difference between the two growth rates over the 1980-2017 and 2007-2017 periods. Negative differences imply that the income of the bottom 40% grew slower than the average national income. The unit of observation is the adult individual aged 20 or above.

Table D.10: Average national incomes in Europe, 1980-2017

	Average income (2017 PPP €)					% of European average income				
	1980	1990	2000	2007	2017	1980	1990	2000	2007	2017
European regions										
Europe	21380	24320	27640	31170	32250	90	88	83	84	82
EU-15 (West)	24230	28150	32260	35380	35260	102	102	97	95	90
EU-13 (East)	12960	13030	13100	17770	22170	55	47	39	48	57
Other West	32310	34970	42550	47990	50850	136	127	127	129	130
Other East	10980	9710	6630	9170	10600	46	35	20	25	27
Eastern Europe										
Albania	6690	5520	6530	9180	11080	28	20	20	25	28
Bosnia and Herzegovina	2030	1650	7480	9540	11400	9	6	22	26	29
Bulgaria	7040	8780	8330	11890	15630	30	32	25	32	40
Croatia	19330	17040	14640	20030	20200	82	62	44	54	52
Czech Republic	18000	20670	18130	23310	26140	76	75	54	63	67
Estonia	12400	14280	14200	23470	25900	52	52	43	63	66
Hungary	14360	15390	13380	17070	19840	61	56	40	46	51
Latvia	13730	15910	9050	18050	20220	58	58	27	49	52
Lithuania	13930	14690	11890	21040	25930	59	53	36	57	66
Moldova	7040	7650	2750	4130	5390	30	28	8	11	14
Montenegro	21540	16570	11590	14820	17430	91	60	35	40	45

Table D.10: Average national incomes in Europe, 1980-2017

	Average income (2017 PPP €)					% of European average income				
	1980	1990	2000	2007	2017	1980	1990	2000	2007	2017
North Macedonia	12940	11160	9210	9840	11850	55	40	28	26	30
Poland	11300	10090	14170	17180	23160	48	37	42	46	59
Romania	12510	12260	9780	15360	20210	53	44	29	41	52
Serbia	17870	16220	6540	10470	11600	75	59	20	28	30
Slovakia	12550	13510	11720	18740	23180	53	49	35	50	59
Slovenia	22360	18190	20340	25980	26500	94	66	61	70	68
Southern Europe										
Cyprus	15860	24110	29300	36810	31580	67	87	88	99	81
Greece	23690	23910	26680	31970	22590	100	87	80	86	58
Italy	25910	29440	32620	33780	29610	109	107	98	91	76
Malta	14130	18160	23030	25030	32290	60	66	69	67	83
Portugal	15240	20200	24280	24800	24550	64	73	73	67	63
Spain	19630	23630	27050	29170	30360	83	86	81	78	78
Western Europe										
Austria	26790	30790	37000	41800	40800	113	112	111	112	104
Belgium	25760	29980	36320	39410	40110	109	109	109	106	103
France	26580	30410	35010	37260	36620	112	110	105	100	94
Germany	28030	31350	33030	36480	39210	118	114	99	98	100

Table D.10: Average national incomes in Europe, 1980-2017

	Average income (2017 PPP €)					% of European average income				
	1980	1990	2000	2007	2017	1980	1990	2000	2007	2017
Ireland	20170	24280	35450	42060	40130	85	88	106	113	103
Luxembourg	39060	47710	76720	135870	101690	165	173	230	365	260
Netherlands	32090	31690	40260	44070	45170	135	115	121	119	115
Switzerland	38330	42400	46080	46820	48430	162	154	138	126	124
United Kingdom	16730	22140	29890	34220	34300	71	80	90	92	88
Northern Europe										
Denmark	26450	29870	37880	43920	45680	112	108	113	118	117
Finland	21060	25560	31520	38400	36660	89	93	94	103	94
Iceland	27210	30480	37450	46710	49540	115	110	112	126	127
Norway	23050	23130	37050	50020	54980	97	84	111	135	141
Sweden	22240	25860	30670	38530	43000	94	94	92	104	110

Notes. The table shows the average national income per adult of European countries in 2017 PPP euros (five first columns) and relative to the European average income per adult (five last columns). Serbia includes Kosovo.

Table D.11: Average state incomes in the United States, 1980-2017

	Average income (2017 PPP €)					% of US average income				
	1980	1990	2000	2007	2017	1980	1990	2000	2007	2017
Alabama	25350	29540	34630	38170	38420	80	79	73	76	73
Alaska	107310	81880	59480	79110	65000	337	218	126	157	123
Arizona	30360	32490	42450	45730	40800	95	87	90	91	77
Arkansas	23980	27480	33370	36060	37500	75	73	71	72	71
California	36080	43330	53420	57930	62130	113	115	113	115	118
Colorado	35090	38290	54530	54050	54020	110	102	116	107	103
Connecticut	33920	48350	63240	69700	63780	107	129	134	138	121
Delaware	35470	48760	73480	68090	66970	112	130	156	135	127
District of Columbia	77000	99640	127890	158180	159120	242	265	271	314	302
Florida	24890	31220	38600	42520	39420	78	83	82	84	75
Georgia	28170	36490	49120	48250	48390	89	97	104	96	92
Hawaii	37260	47850	43500	50780	53580	117	127	92	101	102
Idaho	29800	31880	40210	39910	38970	94	85	85	79	74
Illinois	33950	40370	51800	53970	57190	107	108	110	107	109
Indiana	29170	33270	44170	45140	48770	92	89	94	90	93
Iowa	31760	33540	41960	48180	54460	100	89	89	96	103
Kansas	31560	35090	42660	47110	49610	99	93	90	94	94
Kentucky	27540	30760	36310	38060	40520	87	82	77	76	77

Table D.11: Average state incomes in the United States, 1980-2017

	Average income (2017 PPP €)					% of US average income				
	1980	1990	2000	2007	2017	1980	1990	2000	2007	2017
Louisiana	42730	39080	39890	50930	47270	134	104	85	101	90
Maine	24390	30990	36630	37660	38990	77	83	78	75	74
Maryland	29720	38230	46920	53150	57740	93	102	100	105	110
Massachusetts	31350	42080	57510	60440	66680	99	112	122	120	127
Michigan	30720	34780	46270	42440	44820	97	93	98	84	85
Minnesota	33110	39300	50850	52480	56640	104	105	108	104	108
Mississippi	24130	26170	31030	33680	33950	76	70	66	67	64
Missouri	28210	33430	43750	42950	44390	89	89	93	85	84
Montana	30750	28210	31600	38520	40090	97	75	67	76	76
Nebraska	31190	36060	44020	49010	58230	98	96	93	97	111
Nevada	37430	41880	49310	53690	46280	118	112	105	107	88
New Hampshire	26680	35050	47310	48240	51240	84	93	100	96	97
New Jersey	31220	44160	54980	57770	58190	98	118	117	115	110
New Mexico	34090	30720	40890	43520	41890	107	82	87	86	80
New York	34900	44130	56180	60180	68250	110	118	119	119	130
North Carolina	26660	34260	43800	45750	46630	84	91	93	91	89
North Dakota	31830	30420	36430	45750	66070	100	81	77	91	125
Ohio	30090	34660	45100	44980	49510	95	92	96	89	94
Oklahoma	33030	30900	34670	42490	43820	104	82	74	84	83

Table D.11: Average state incomes in the United States, 1980-2017

	Average income (2017 PPP €)					% of US average income				
	1980	1990	2000	2007	2017	1980	1990	2000	2007	2017
Oregon	29830	32740	44010	47400	49260	94	87	93	94	94
Pennsylvania	27750	33120	42140	45820	51190	87	88	89	91	97
Rhode Island	26140	34330	41950	46730	48500	82	91	89	93	92
South Carolina	24330	31300	37310	37760	38440	76	83	79	75	73
South Dakota	26820	31610	40280	47300	52320	84	84	85	94	99
Tennessee	26100	31590	40700	41090	45580	82	84	86	82	87
Texas	39030	38570	47950	54480	55970	123	103	102	108	106
Utah	32040	35420	45640	52370	52740	101	94	97	104	100
Vermont	25500	34280	38590	40650	44110	80	91	82	81	84
Virginia	28900	38050	48320	52730	53190	91	101	102	105	101
Washington	33440	40420	52270	55920	60200	105	108	111	111	114
West Virginia	25240	25290	29340	32300	36340	79	67	62	64	69
Wisconsin	30310	34250	44180	45470	49510	95	91	94	90	94
Wyoming	61270	49350	46230	71160	62160	193	131	98	141	118

Notes. The table shows the average income per adult of US states in 2017 PPP euros (five first columns) and relative to the US average national income per adult (five last columns). *Sources.* Bureau of Economic Analysis (GDP) and Census Bureau (adult population estimates).

Table D.12: Predistribution versus redistribution in Europe and the United States: estimates of the top 1% share and of Gini and Theil indices using different concepts and data sources

	Top 1%			Gini			Theil		
	United States	Europe	Difference	United States	Europe	Difference	United States	Europe	Difference
Surveys									
Factor income	9.6%	7.2%	+2.4 pp.	52.1%	56.1%	−4.1 pp.	51.3%	54.2%	−2.9 pp.
Pretax income	8.5%	6.1%	+2.4 pp.	44.9%	36.7%	+8.2 pp.	38.1%	24.7%	+13.4 pp.
Posttax income	6.5%	5.3%	+1.3 pp.	39.2%	32.4%	+6.8 pp.	28.2%	18.6%	+9.7 pp.
Surveys + Tax data									
Factor income	15.8%	10.1%	+5.7 pp.	59.7%	57.5%	+2.3 pp.	81.7%	68.1%	+13.6 pp.
Pretax income	16.5%	9.4%	+7.1 pp.	54.7%	43.1%	+11.6 pp.	74.8%	41.2%	+33.6 pp.
Posttax income	13.1%	7.8%	+5.3 pp.	48.6%	38.4%	+10.2 pp.	54.9%	31.6%	+23.4 pp.
DINA									
Factor income	18.9%	11.8%	+7.1 pp.	61.0%	52.1%	+8.8 pp.	95.4%	60.8%	+34.5 pp.
Pretax income	19.5%	11.5%	+8.0 pp.	59.0%	44.5%	+14.4 pp.	92.6%	47.2%	+45.4 pp.
Posttax income	14.3%	9.4%	+4.9 pp.	47.3%	38.8%	+8.6 pp.	59.3%	35.5%	+23.8 pp.

Source: Authors' computations combining surveys, tax data and national accounts for Europe (population-weighted average). Survey-based estimates for the United States come from the Luxembourg Income Study. Surveys + Tax data and DINA estimates for the United States come from Piketty, Saez, and Zucman (2018). *Notes:* The table shows how estimates of top 1% factor income, pretax income, and posttax income shares in Europe and the United States in 2017, as well as Gini and Theil indices, vary depending on whether they are observed in household surveys, computed by combining surveys and tax data, or estimated using the distributional national accounts methodology.

Appendix E

Appendix to "Who Benefits from Public Goods? Public Services and Inequality in Post-Apartheid South Africa"

E.1 Additional Key Results

Figure E.1: GDP and National Income Per Capita in South Africa, 1993-2019

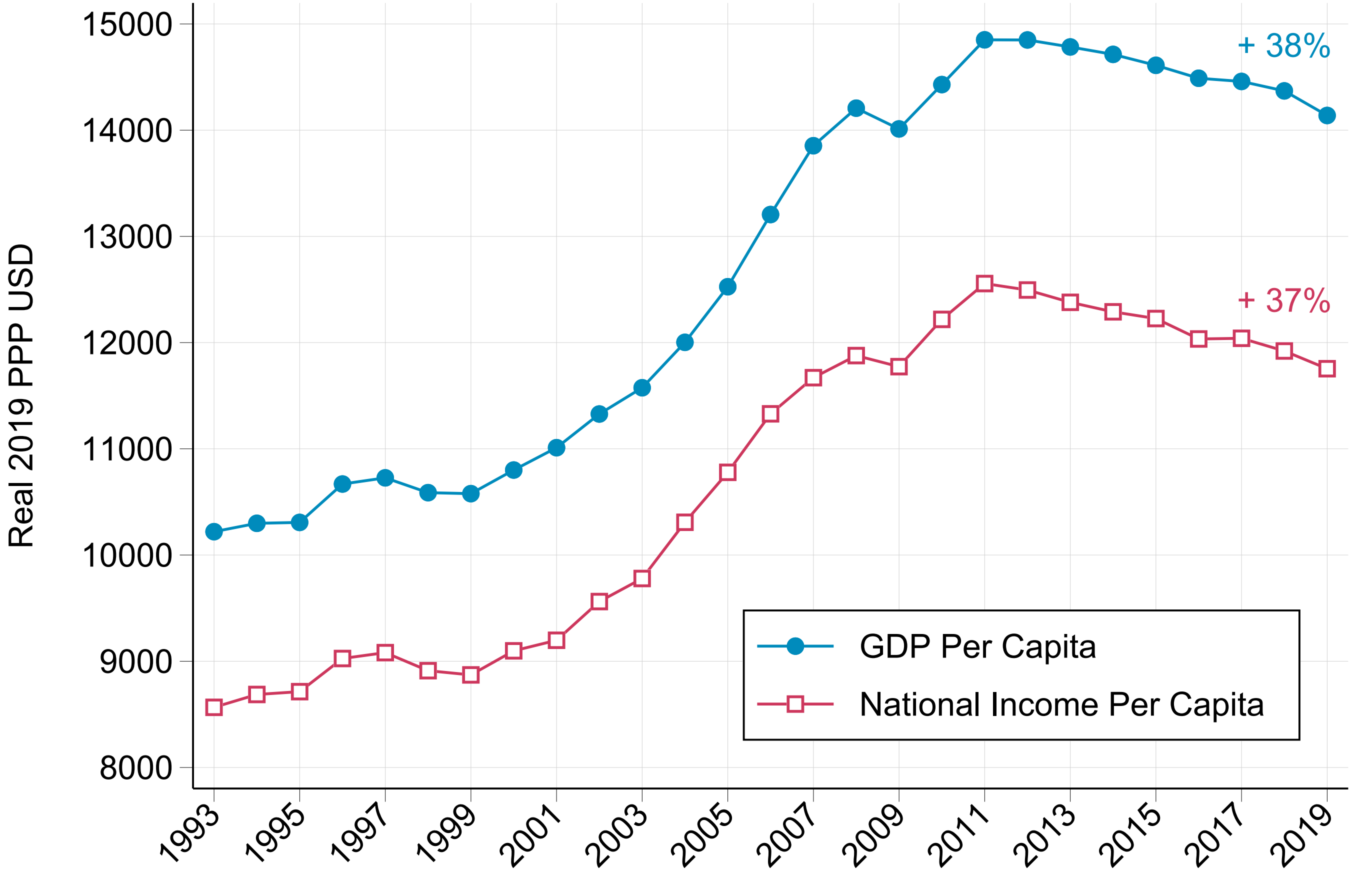

Notes. Author's elaboration using data from the South African Reserve Bank. Growth figures correspond to total real growth rates between 1993 and 2019.

Figure E.2: Bottom 50% Average Income Before and After Transfers, 1993-2019: Productivity-Adjusted

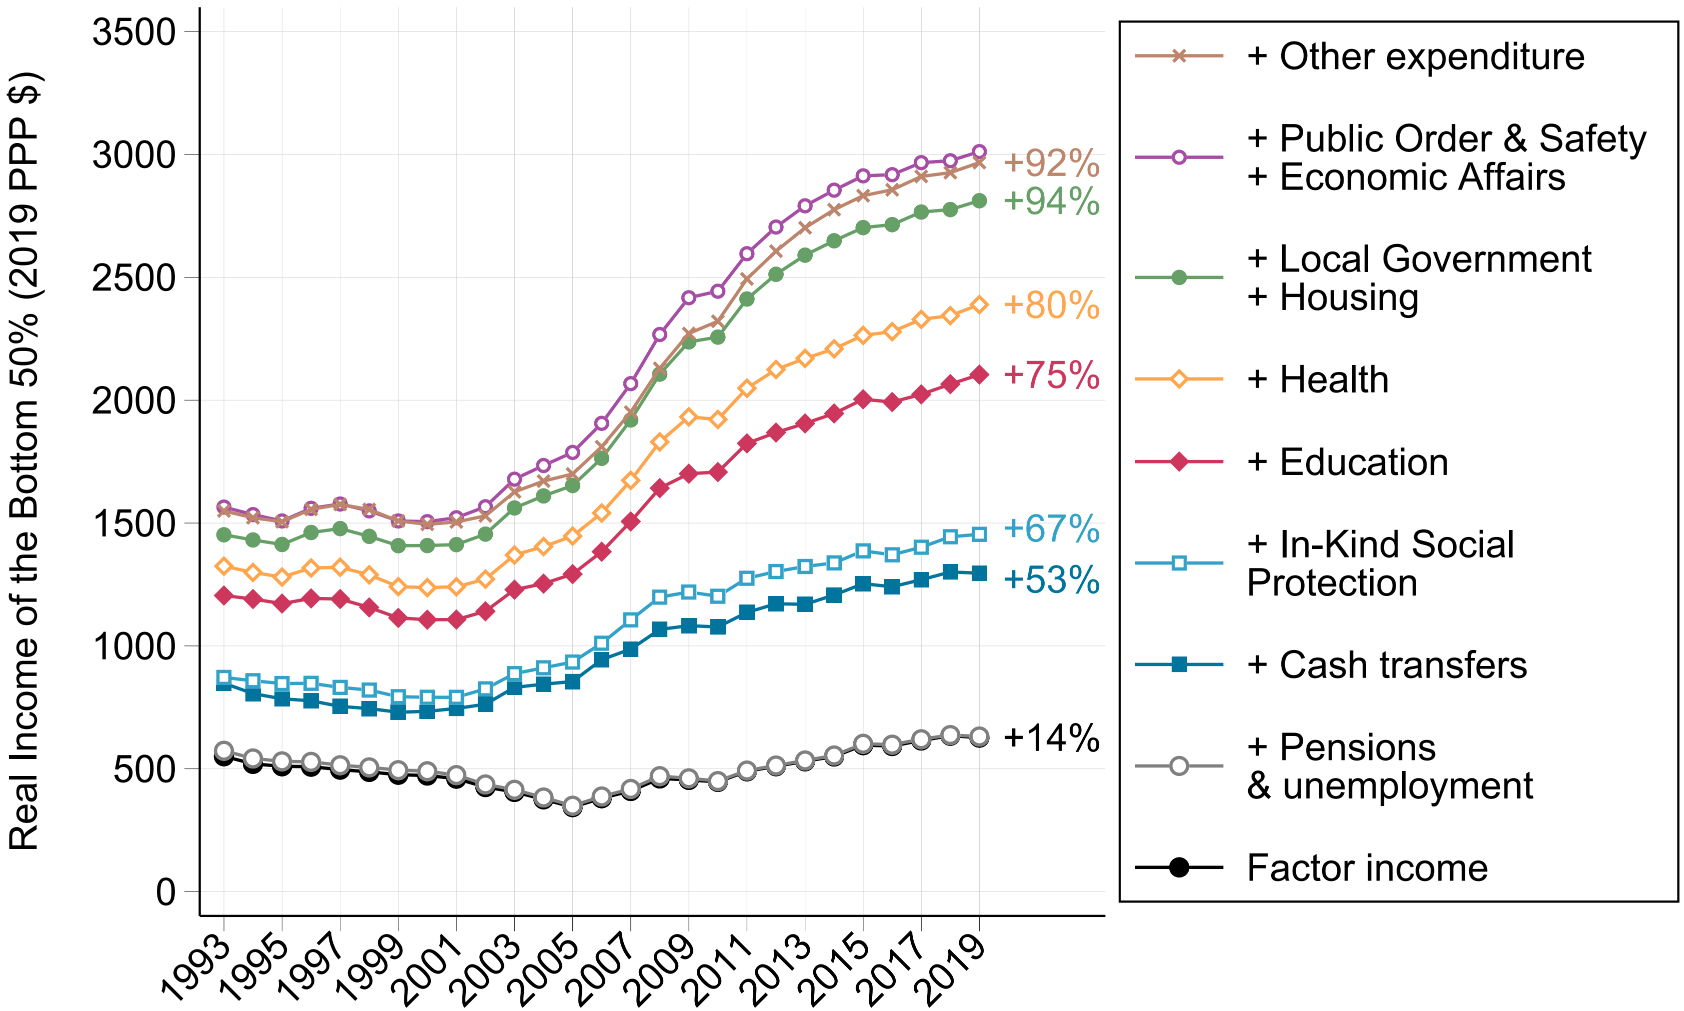

Notes. The figure represents the evolution of the real average income of the bottom 50%, before and after adding cash and in-kind transfers one by one to the analysis. All in-kind transfers are adjusted for aggregate and heterogeneous productivity. Other expenditure corresponds to general public services and defense, distributed proportionally to posttax disposable income. The unit of observation is the individual. Income is split equally among all household members.

Table E.1: The Distribution of Income in South Africa in 2019: Productivity-Adjusted

	Pretax National Income		Posttax Disposable Income		Posttax National Income	
	Average Income	Income Share	Average Income	Income Share	Average Income	Income Share
Full population	\$ 11,800	100%	\$ 7,780	100%	\$ 9,480	100%
Bottom 50%	\$ 630	2.7%	\$ 1,020	6.5%	\$ 1,690	8.9%
Bottom 20%	\$ 45	0.1%	\$ 410	1.1%	\$ 400	0.8%
Next 30%	\$ 1,020	2.6%	\$ 1,420	5.5%	\$ 2,560	8.1%
Middle 40%	\$ 8,410	28.6%	\$ 6,530	33.6%	\$ 8,020	33.8%
Top 10%	\$ 80,700	68.7%	\$ 46,600	59.9%	\$ 54,300	57.2%
Top 1%	\$ 329,000	28.0%	\$ 170,000	21.9%	\$ 206,000	21.7%
Top 0.1%	\$ 970,000	8.3%	\$ 519,000	6.7%	\$ 609,000	6.4%

Notes. The table reports statistics on the distribution of income in South Africa in 2019 for different income concepts. Posttax disposable income is the sum of primary incomes, minus direct taxes, plus cash transfers. Posttax national income deducts all taxes and adds all transfers. General public services and defense are distributed proportionally to posttax disposable income. In-kind transfers are adjusted for aggregate and heterogeneous productivity. The unit of observation is the individual. Income is split equally between all household members.

E.2 Pension and Unemployment Systems

Figure E.3: Pension Contributions and Benefits Paid/Received by Income Decile, 2019

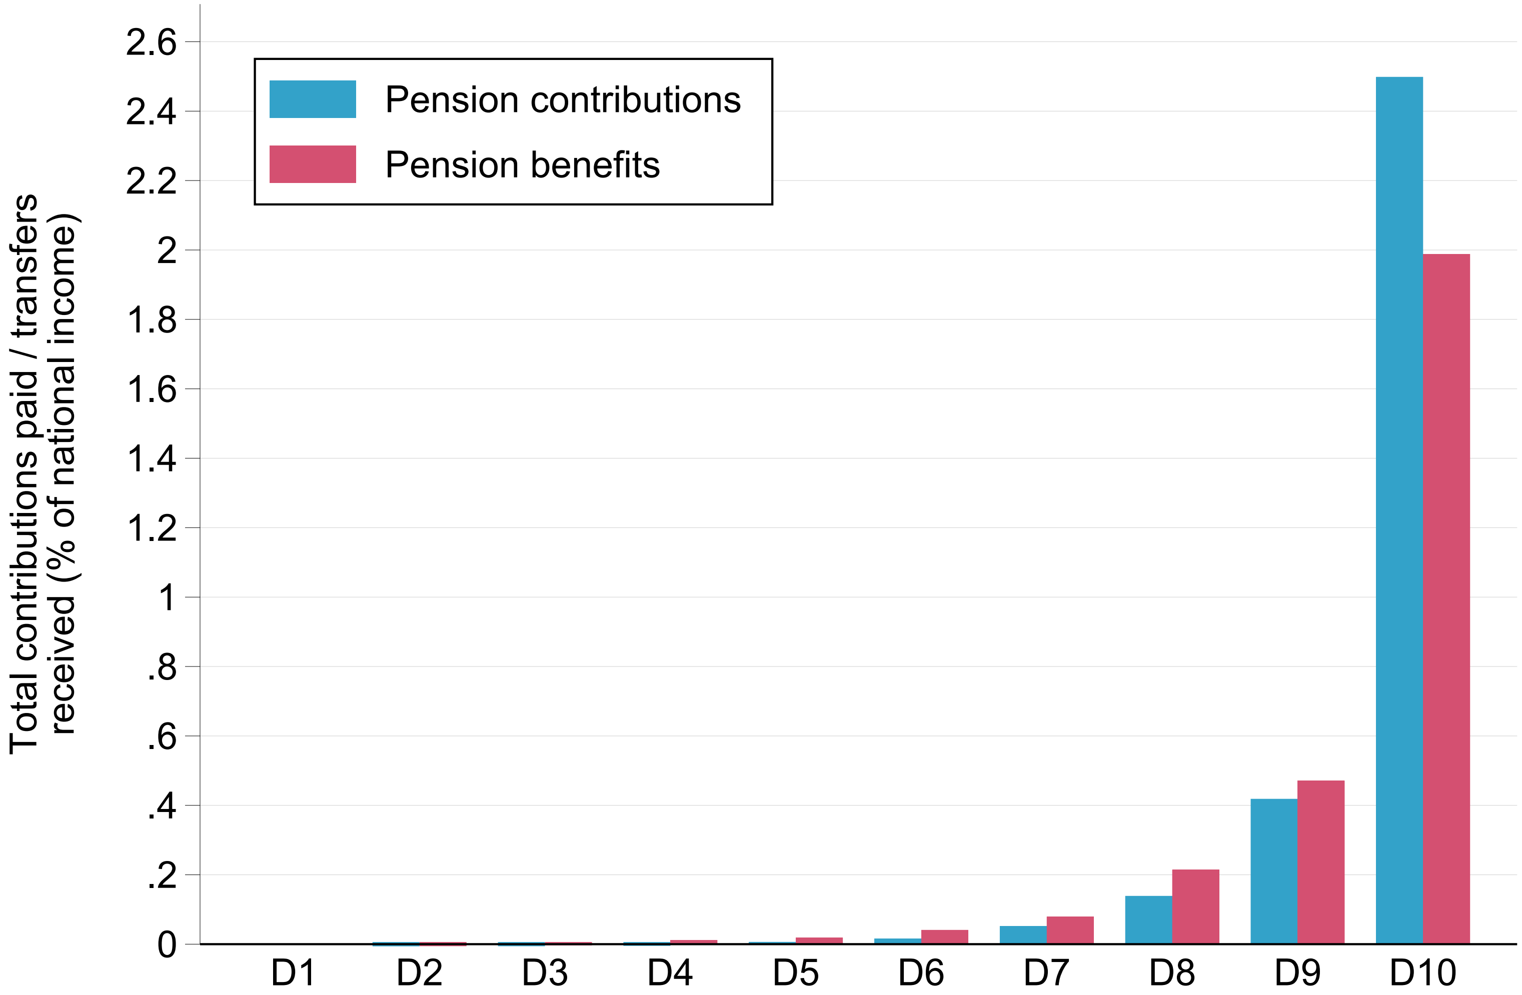

Notes. Author's computations combining surveys, tax, and national accounts data.

Figure E.4: Unemployment Insurance Contributions and Benefits Paid/Received by Income Decile, 2019

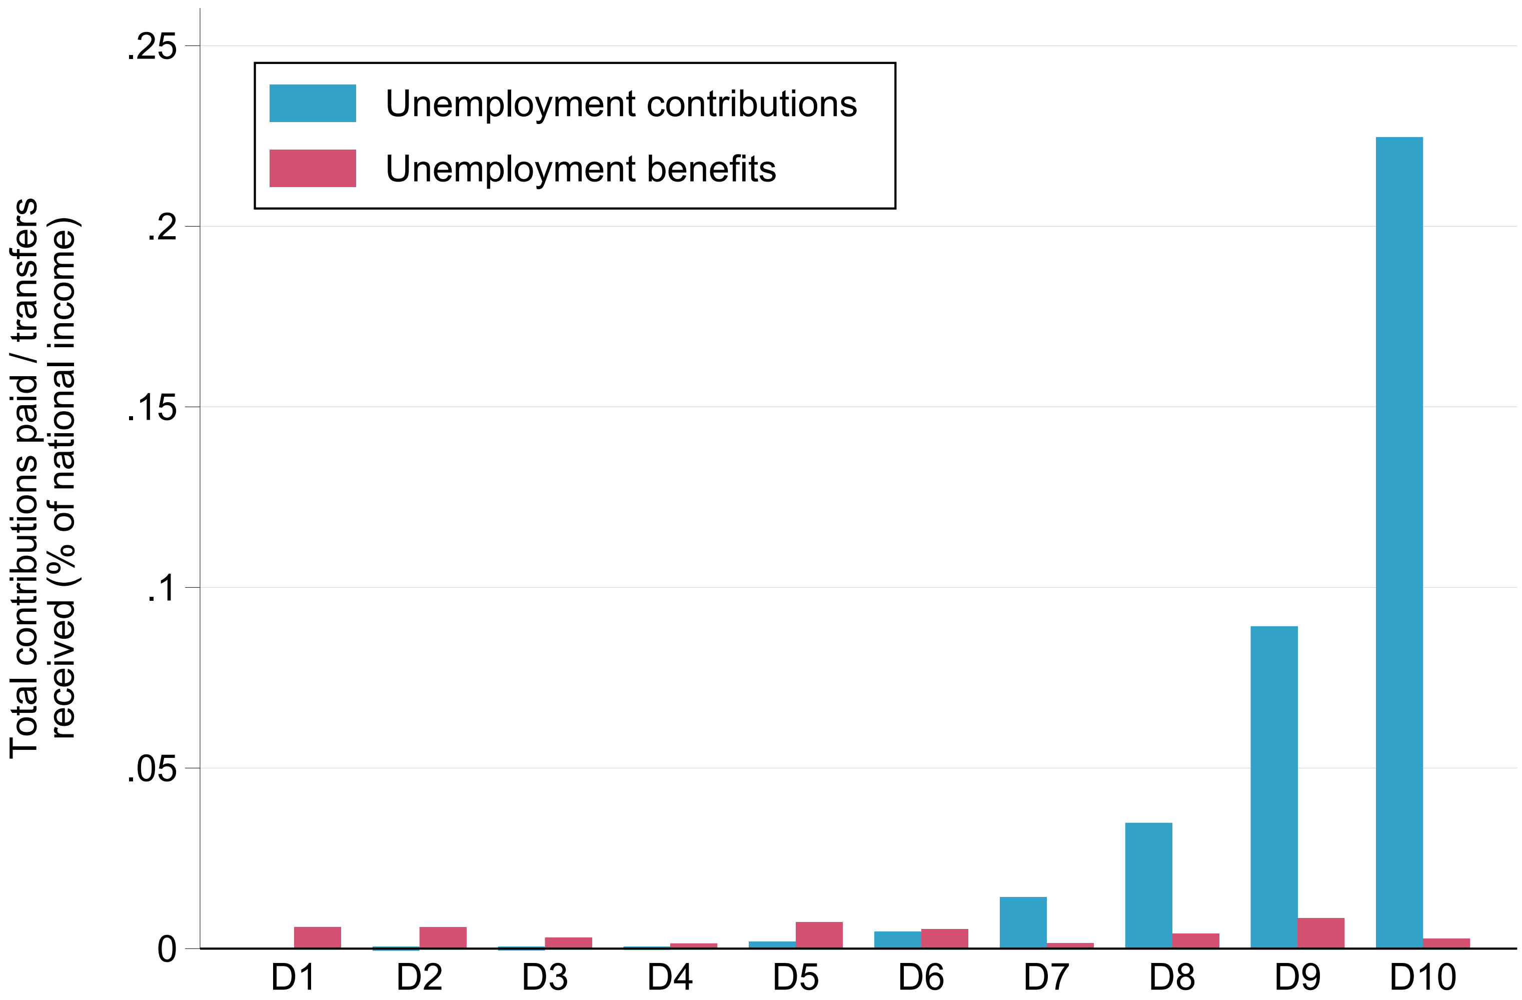

Notes. Author's computations combining surveys, tax, and national accounts data.

Figure E.5: Net Transfers Operated by the Pension and Unemployment Insurance Systems Between Income Deciles, 2019

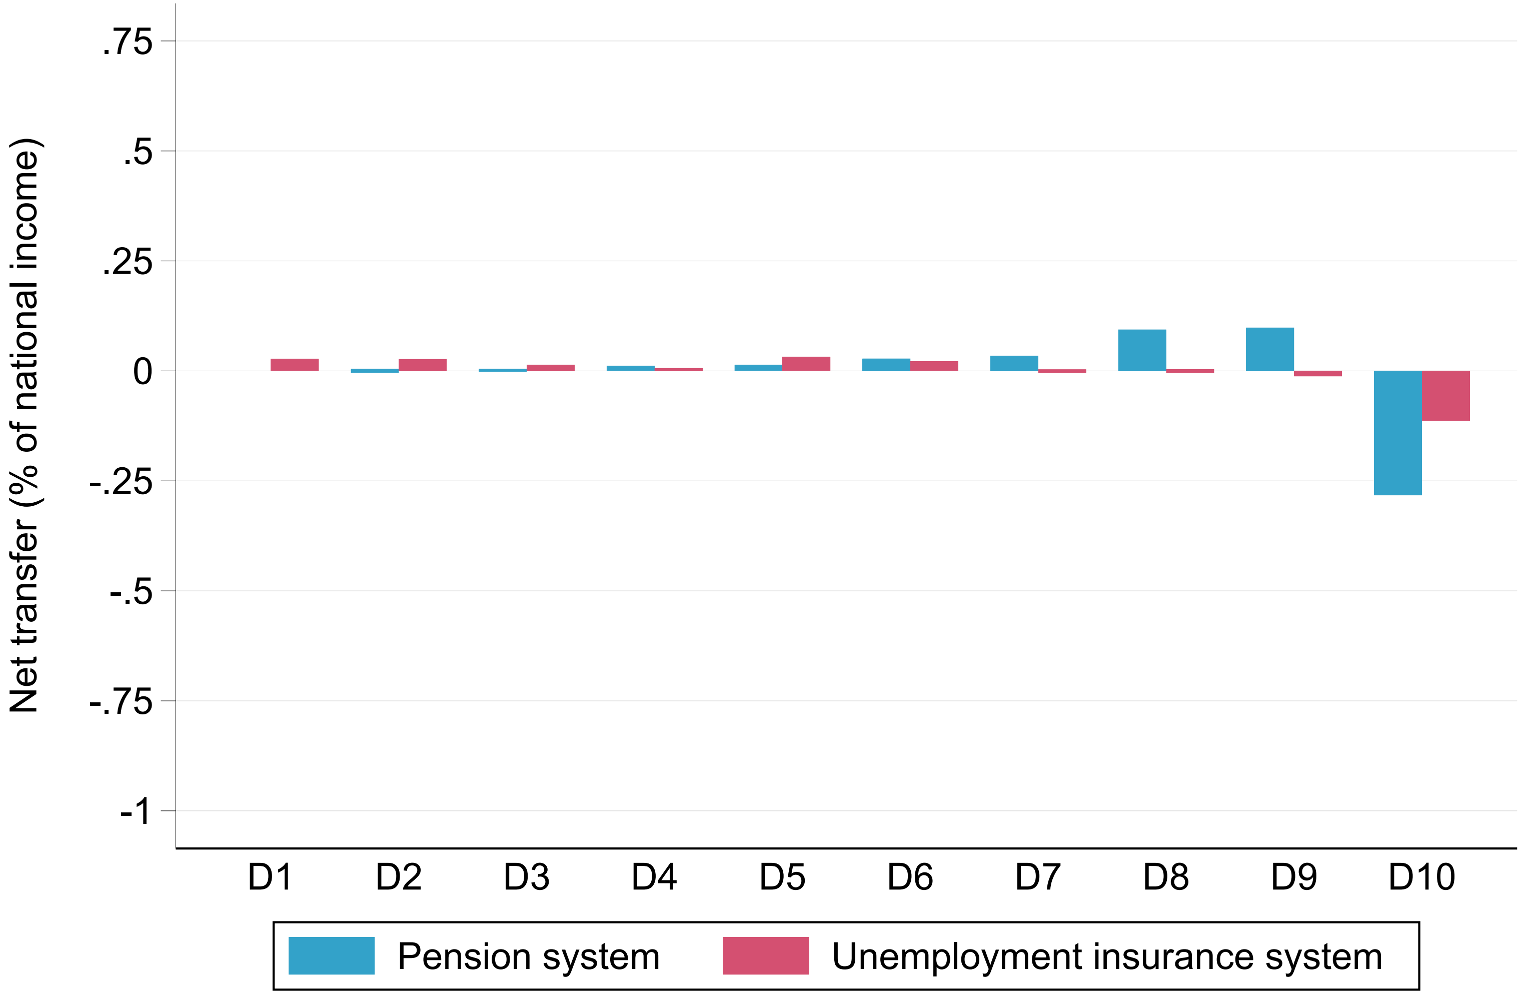

Notes. Author's computations combining surveys, tax, and national accounts data.

E.3 Social Protection

Figure E.6: Level and Composition of Social Protection Expenditure in South Africa, 1993-2019

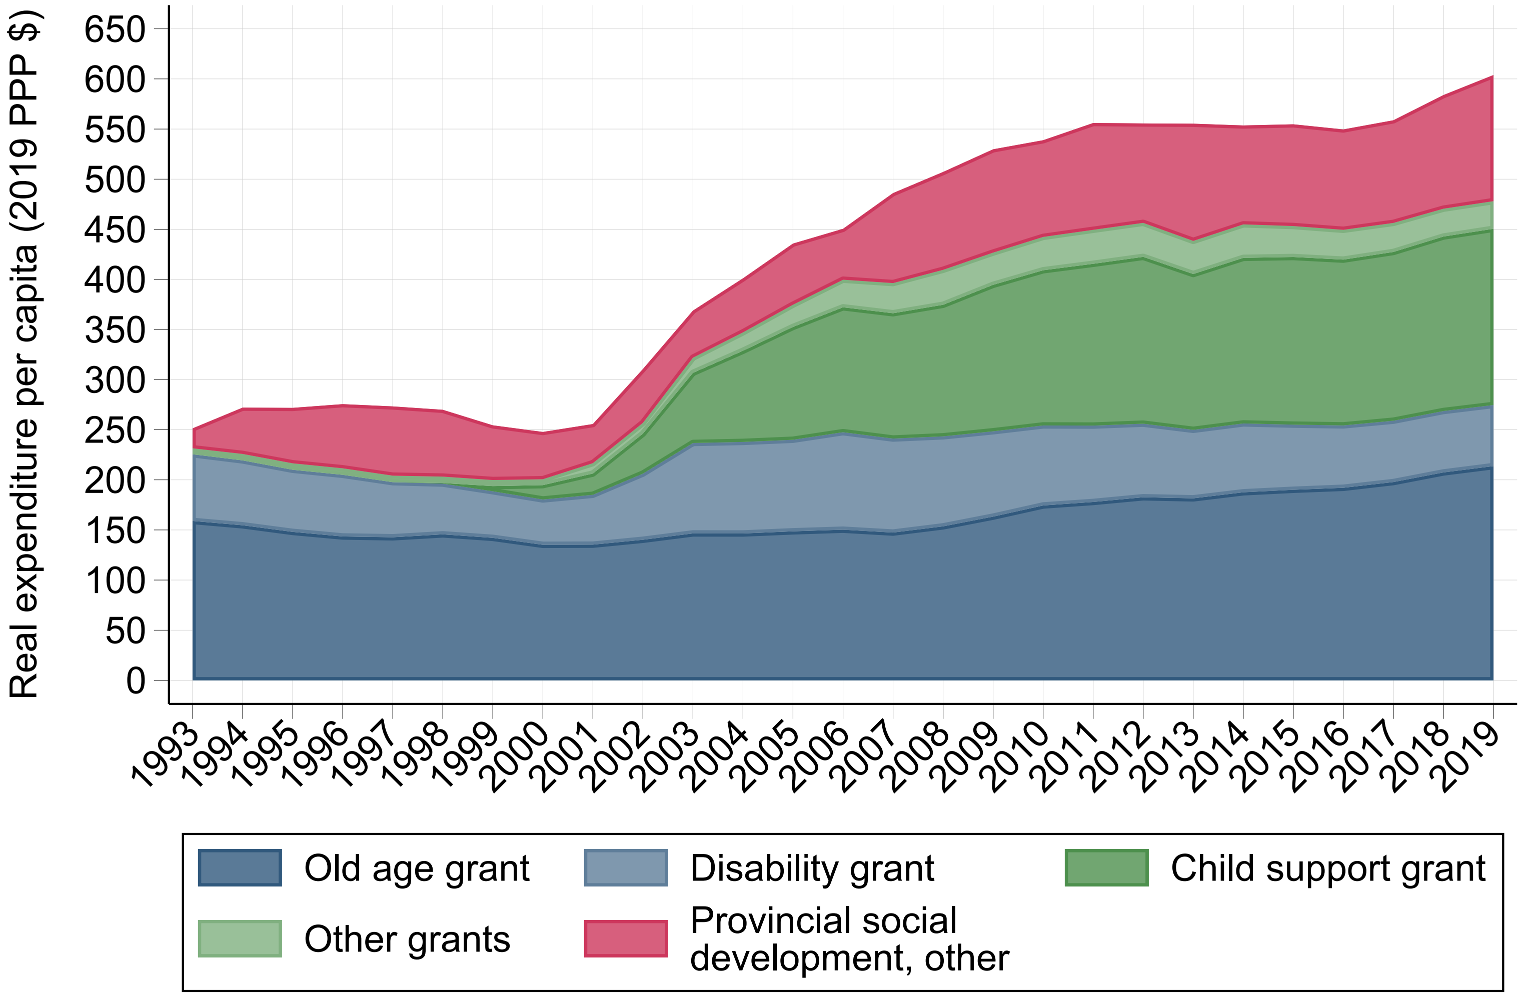

Notes. Author's computations combining data from South African National Treasury Budget Reports (1994-2020).

Figure E.7: Per Capita Expenditure on Social Grants in South Africa, 1993-2019

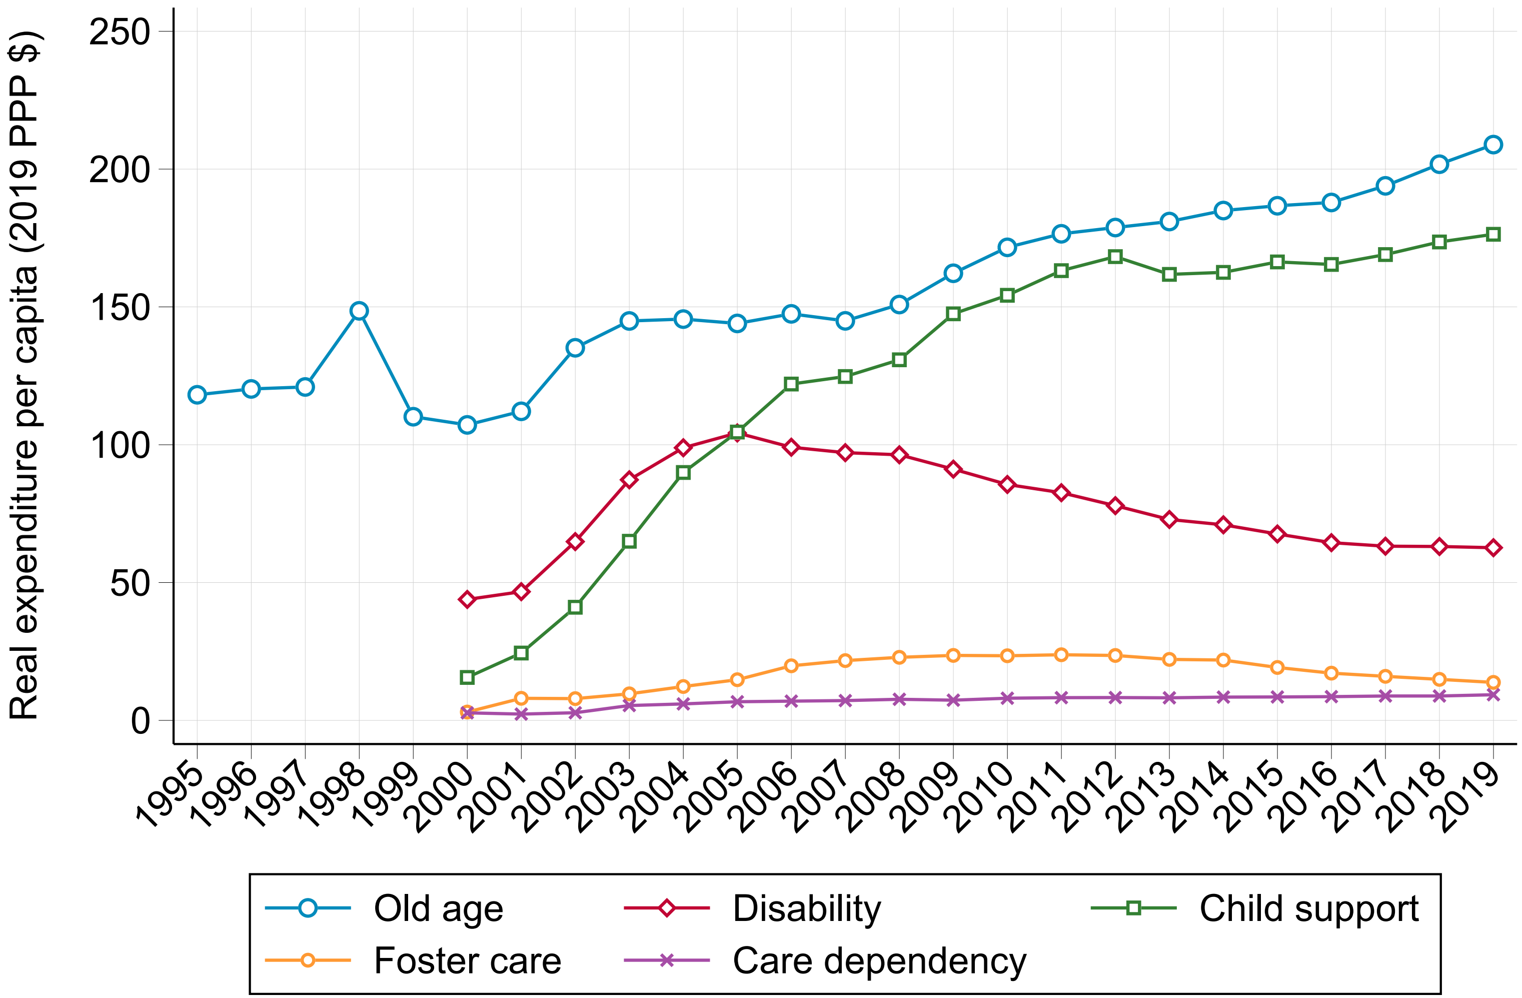

Notes. Author's computations combining data from South African National Treasury Budget Reports (1994-2020).

Figure E.8: Real Monthly Value of Social Grants in South Africa, 1993-2019

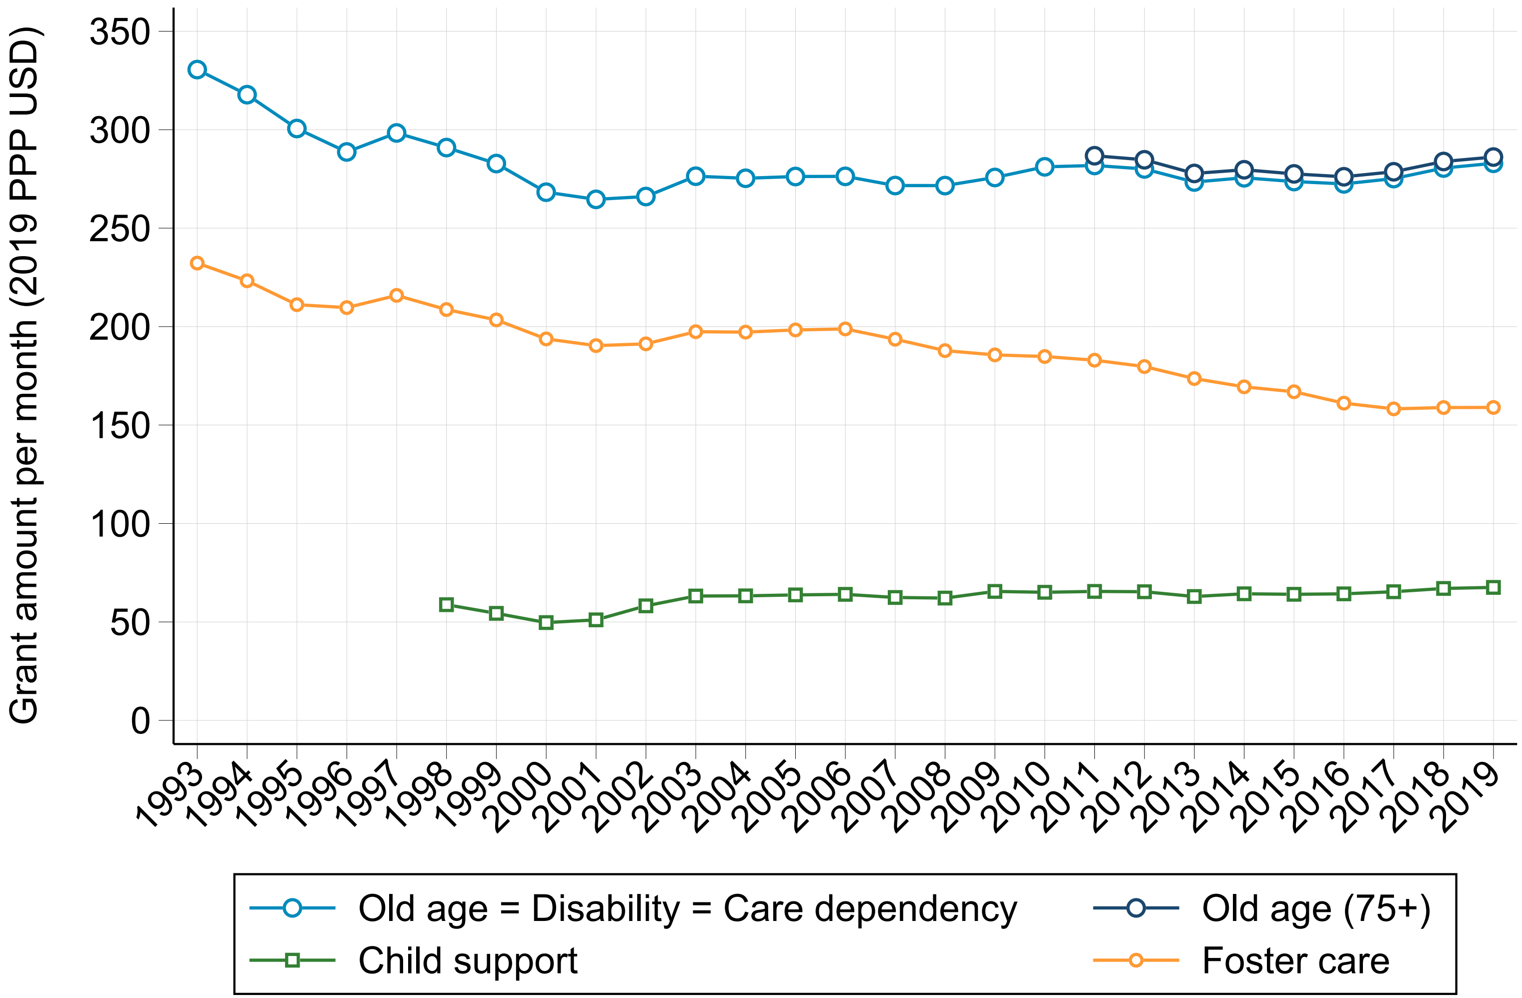

Notes. Author's computations combining data from South African National Treasury Budget Reports (1994-2020).

Figure E.9: Share of Population Receiving Social Grants in South Africa, 1993-2019

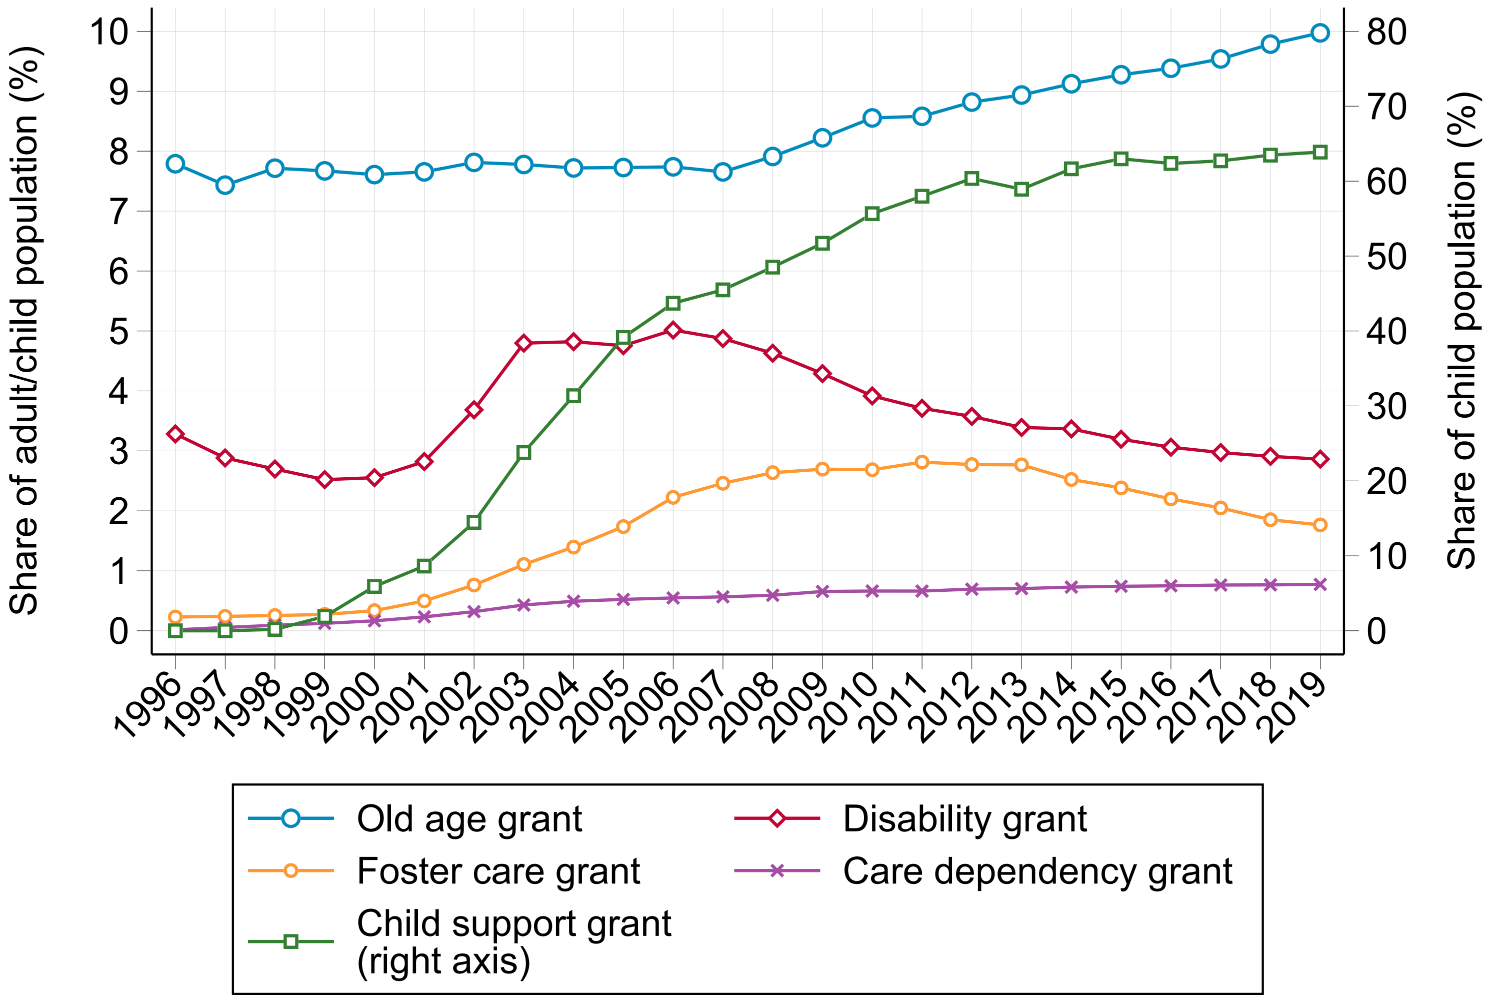

Notes. Author's computations combining data from South African National Treasury Budget Reports (1994-2020).

Figure E.10: Average Social Protection Transfer Received by Income Group, 1993-2019

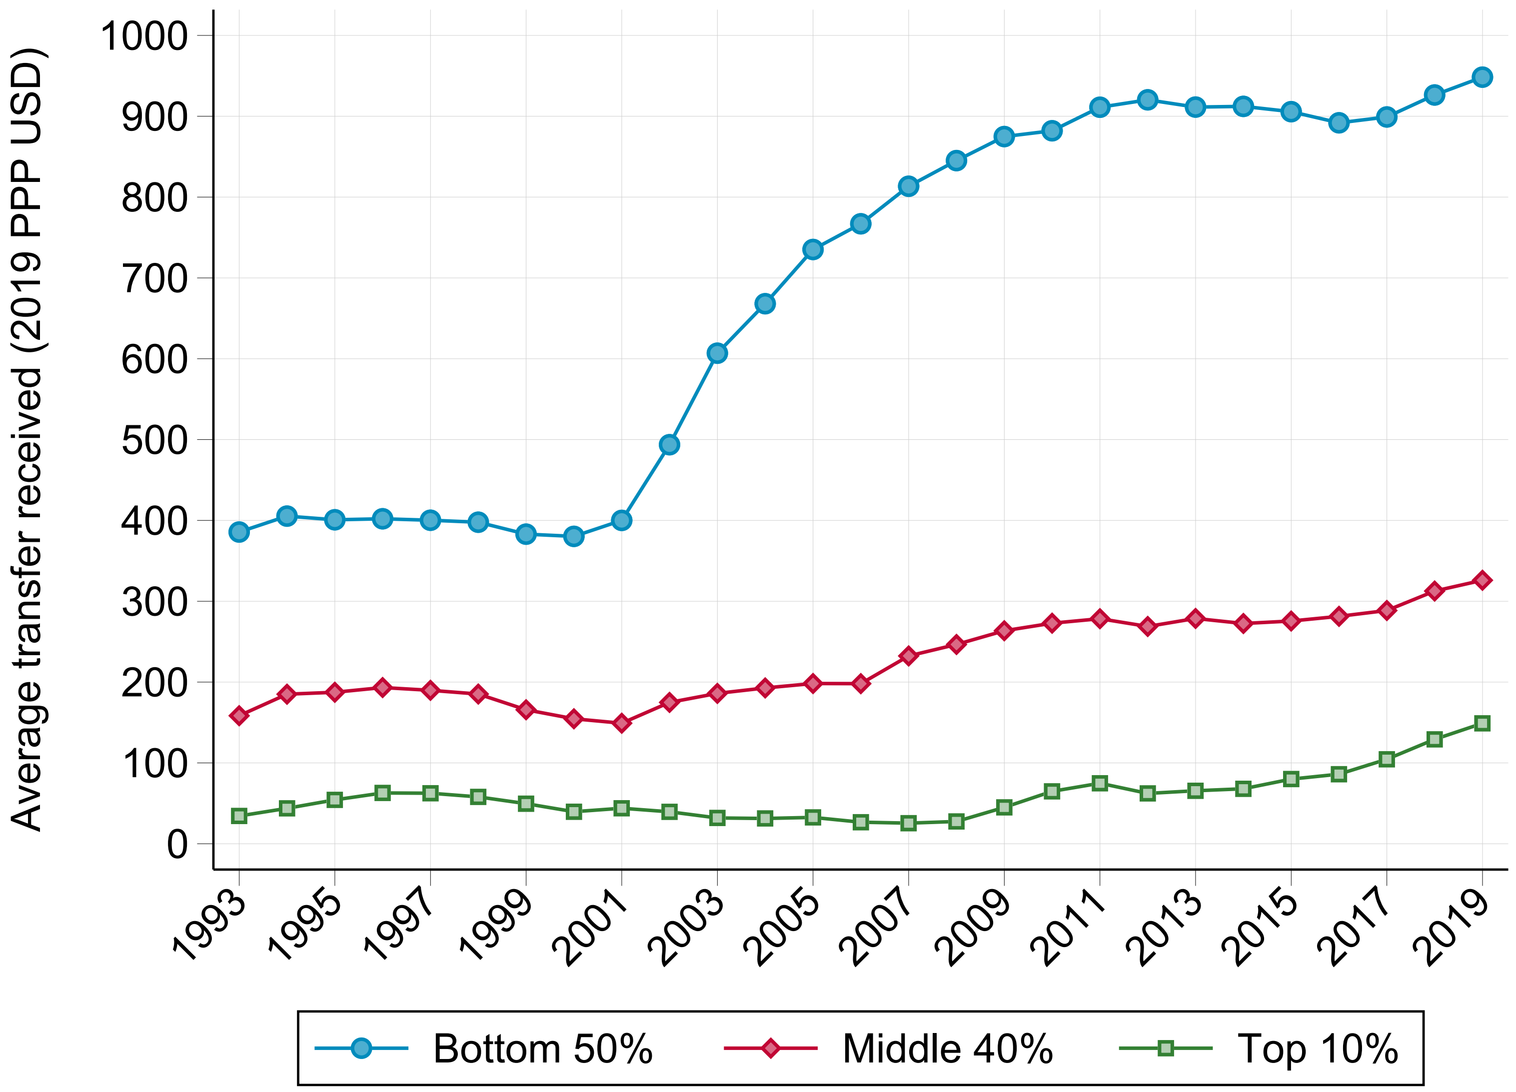

Notes. Author's computations combining surveys, tax, and national accounts data.

Figure E.11: Average Social Protection Transfer Received by Income Decile, 1993-2019

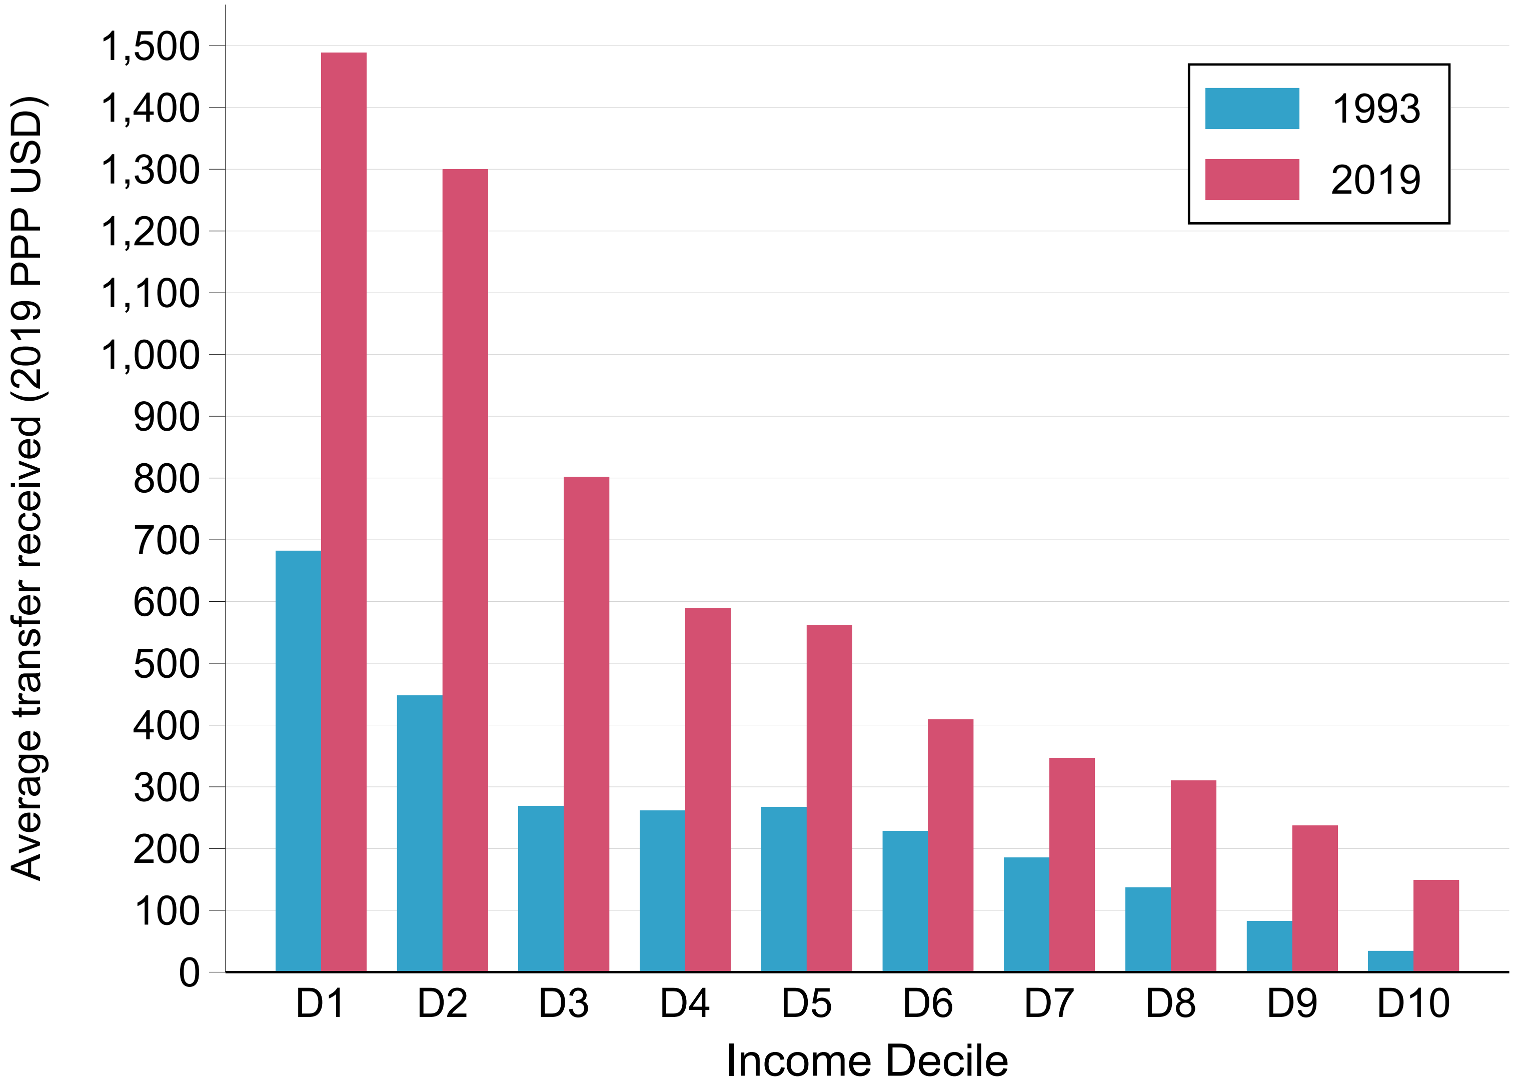

Notes. Author's computations combining surveys, tax, and national accounts data.

E.4 Education

Figure E.12: Level and Composition of Education Expenditure in South Africa, 1993-2019

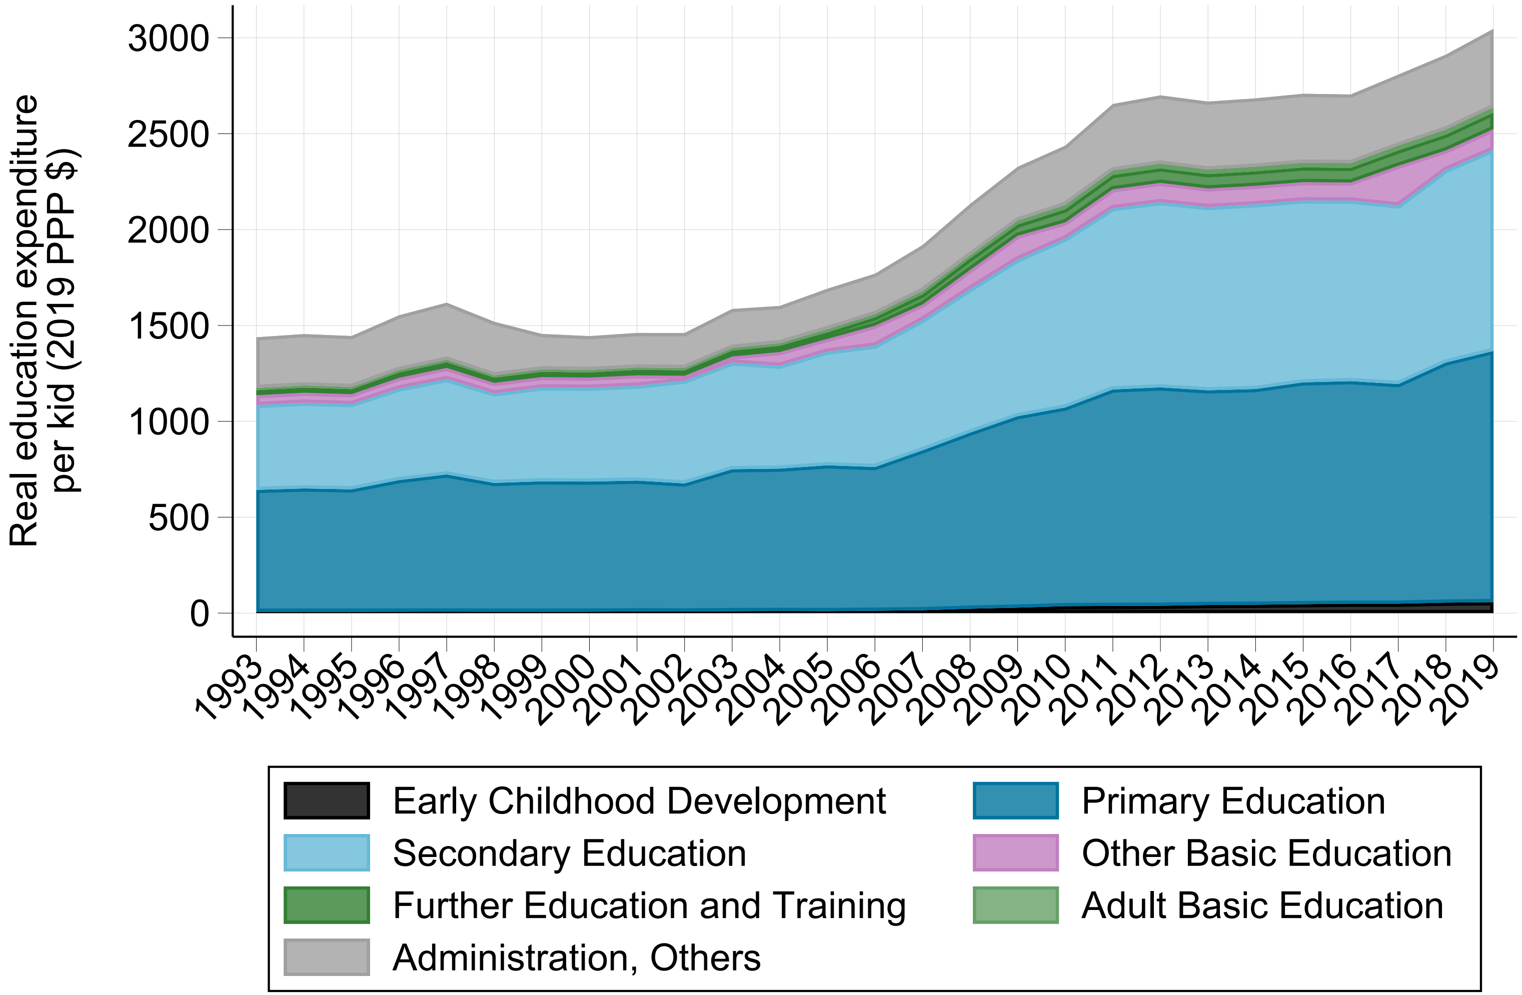

Notes. Author's computations combining data from South African National Treasury Budget Reports (1994-2020) and Provincial Budget Reports (2002-2020).

Figure E.13: The Rise of Education Expenditure in South Africa, 1993-2019:
The Role of Basic Education

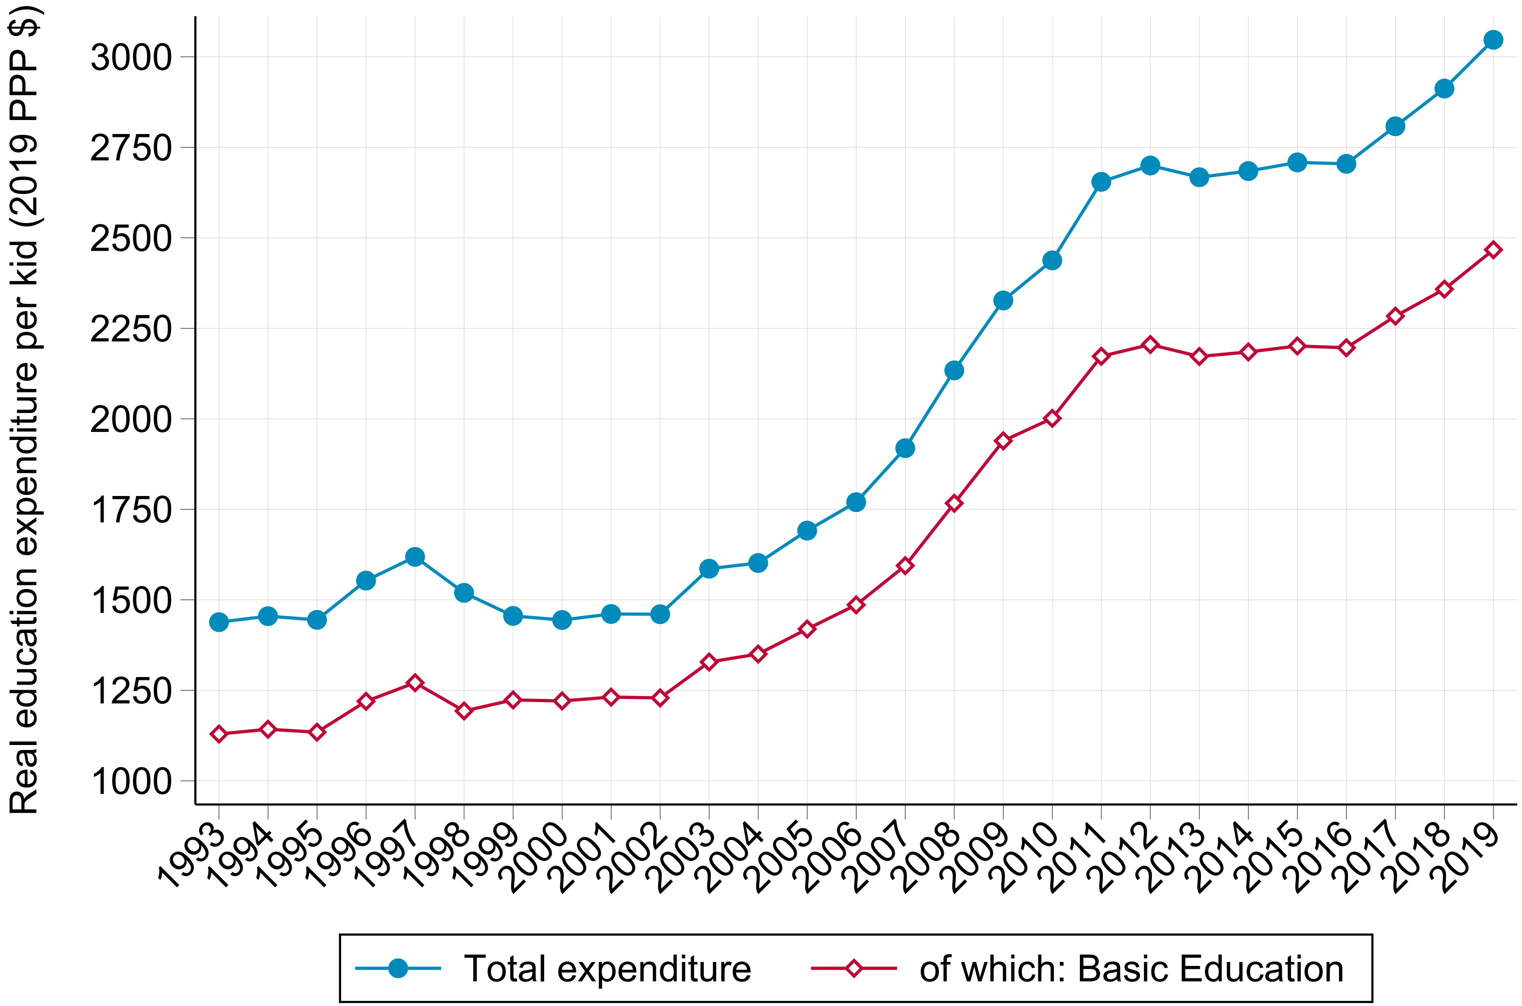

Notes. Author's computations combining data from South African National Treasury Budget Reports (1994-2020).

Figure E.14: Real Education Expenditure Per Kid by South African Province, 1993-2019

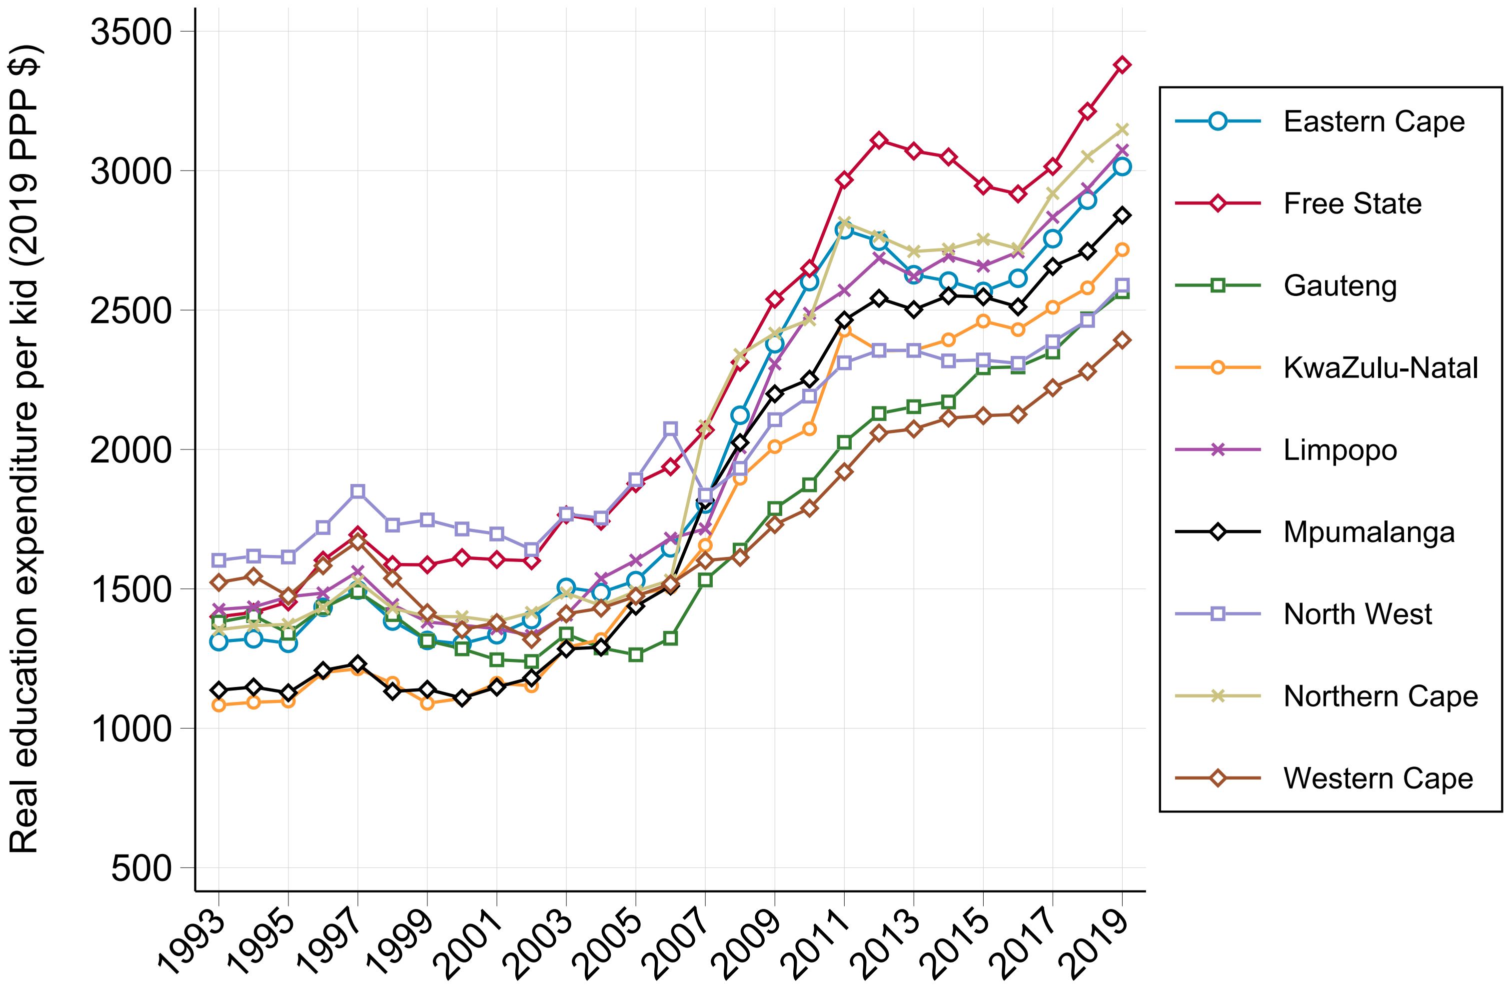

Notes. Author's computations combining data from South African National Treasury Budget Reports (1994-2020) and Provincial Budget Reports (2002-2020).

Figure E.15: Average Number of Children Attending Public Schools by Income Group, 1996-2016

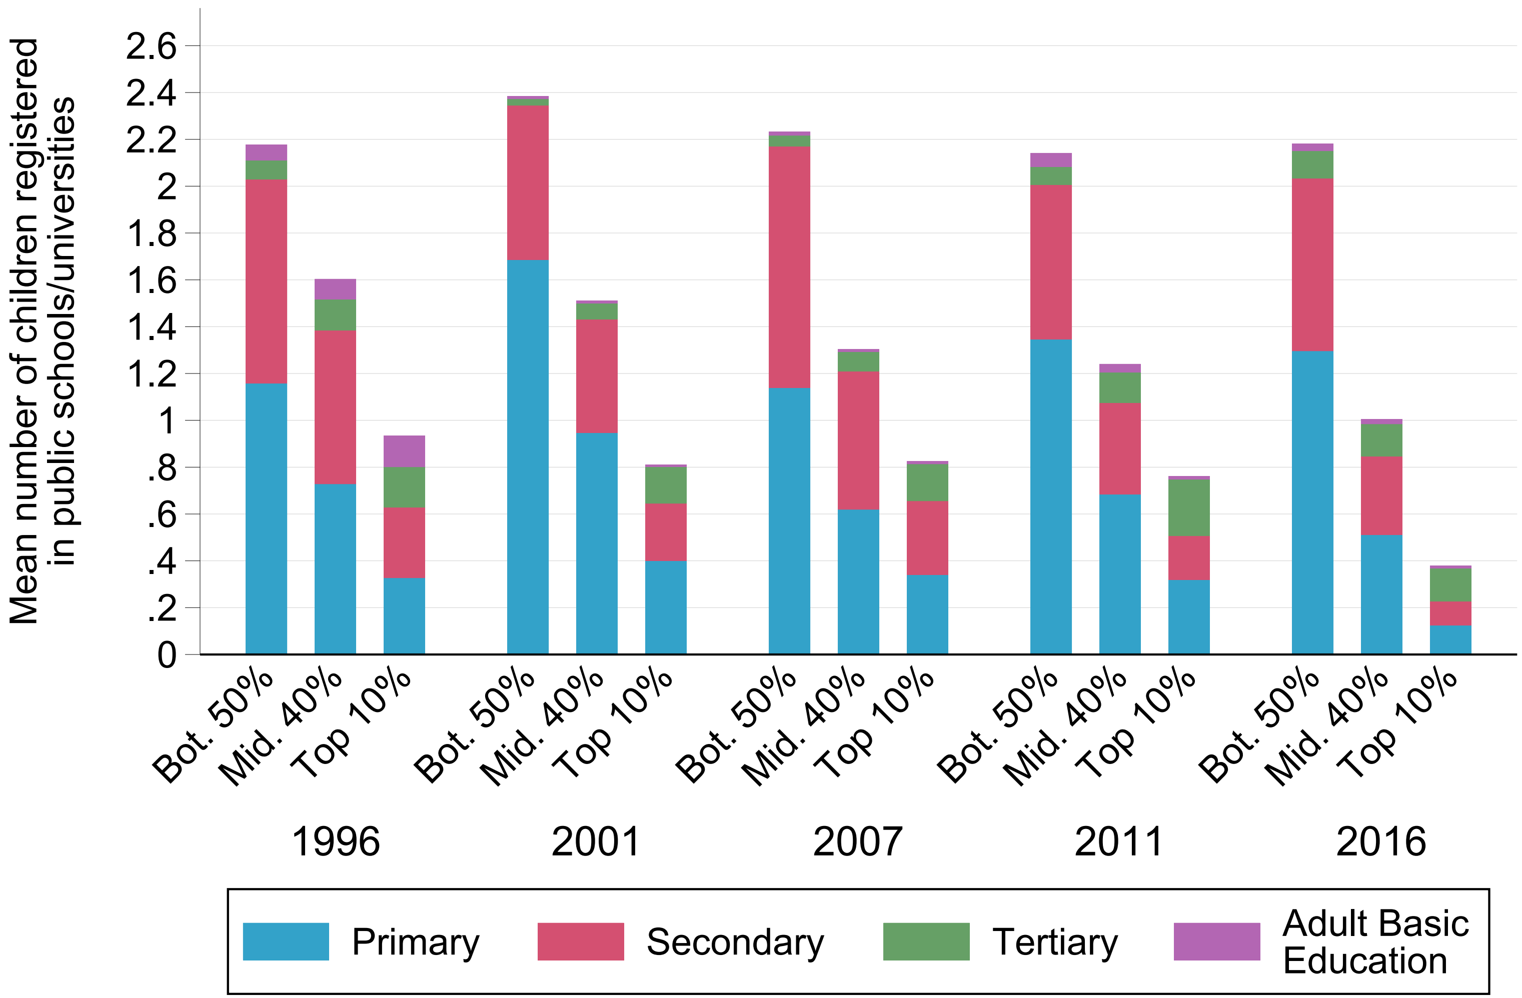

Notes. Author's computations using census sample microdata.

Figure E.16: Average Share of Children Attending Private Schools by Income Group, 2001-2016

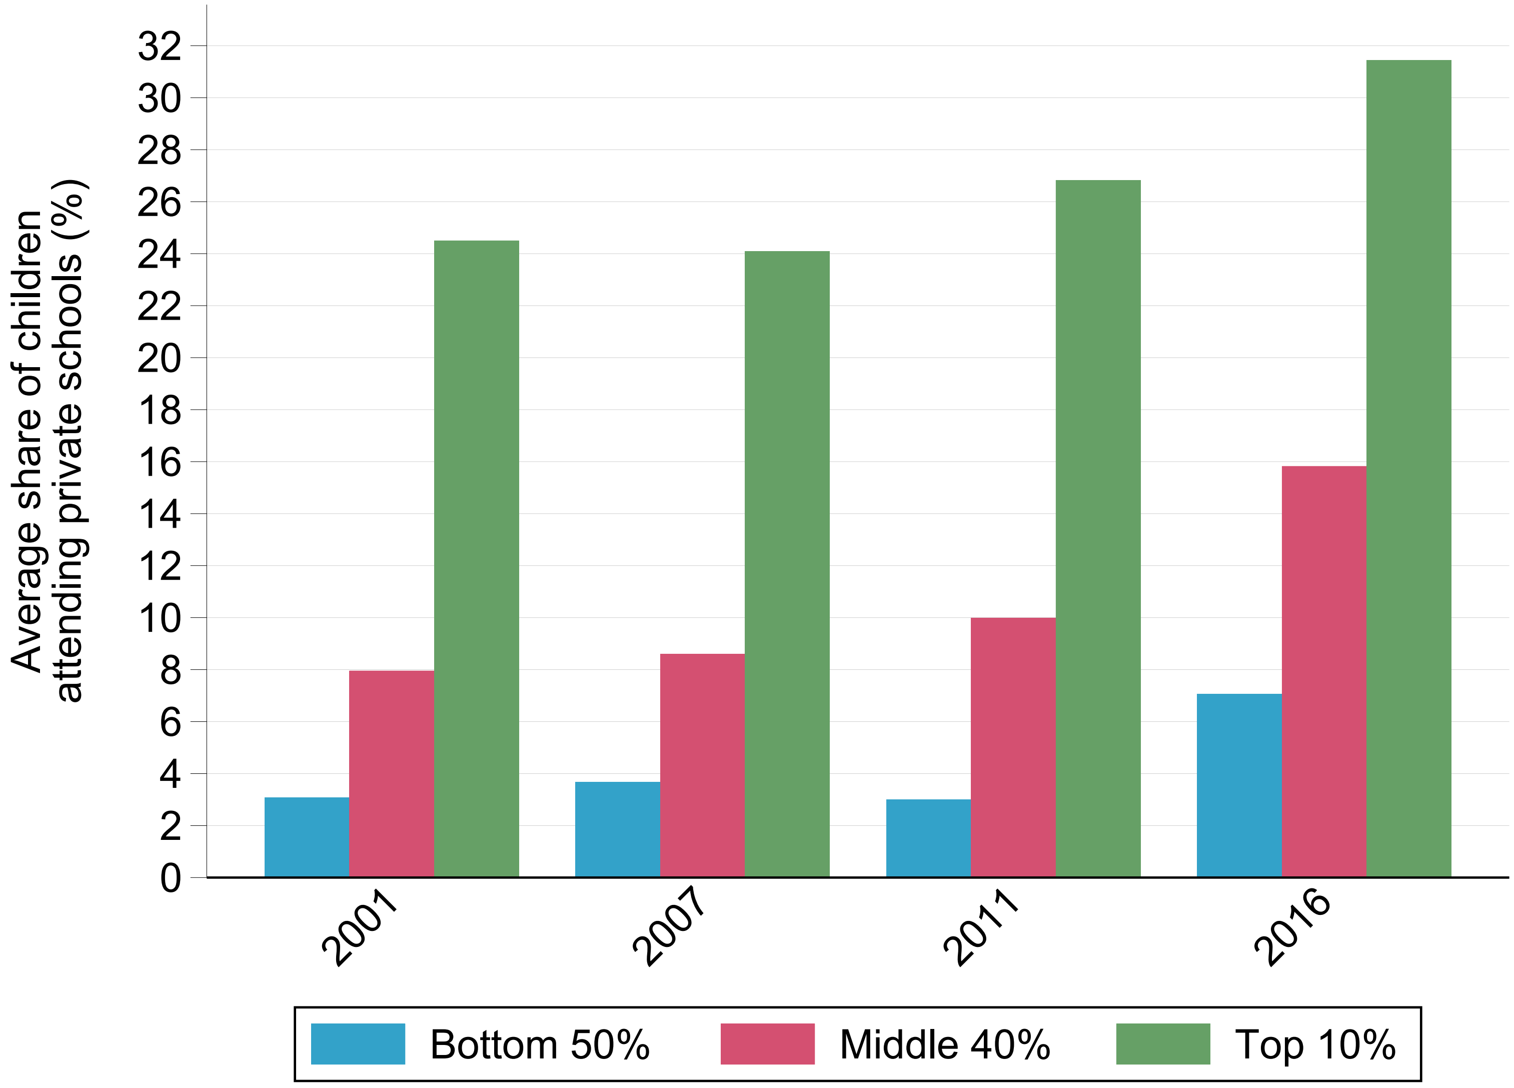

Notes. Author's computations using census sample microdata.

Figure E.17: Average Education Transfer Received by Income Group, 1993-2019

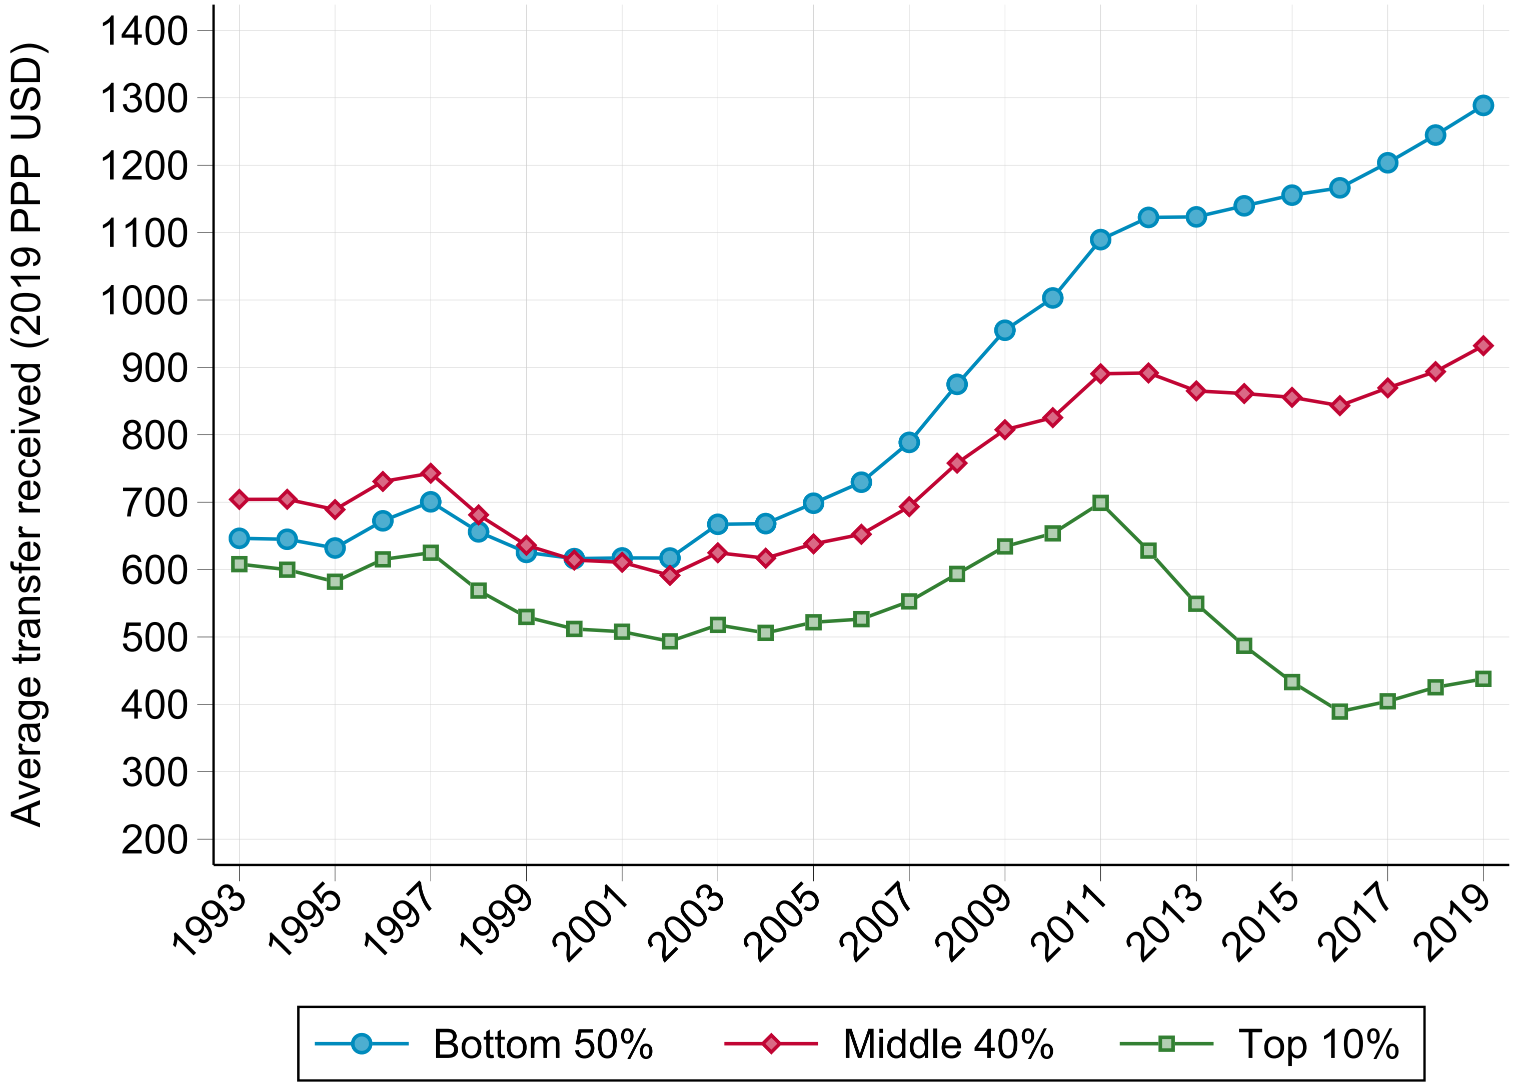

Notes. Author's computations combining surveys, tax, and national accounts data.

Figure E.18: Average Education Transfer Received by Income Decile, 1993-2019

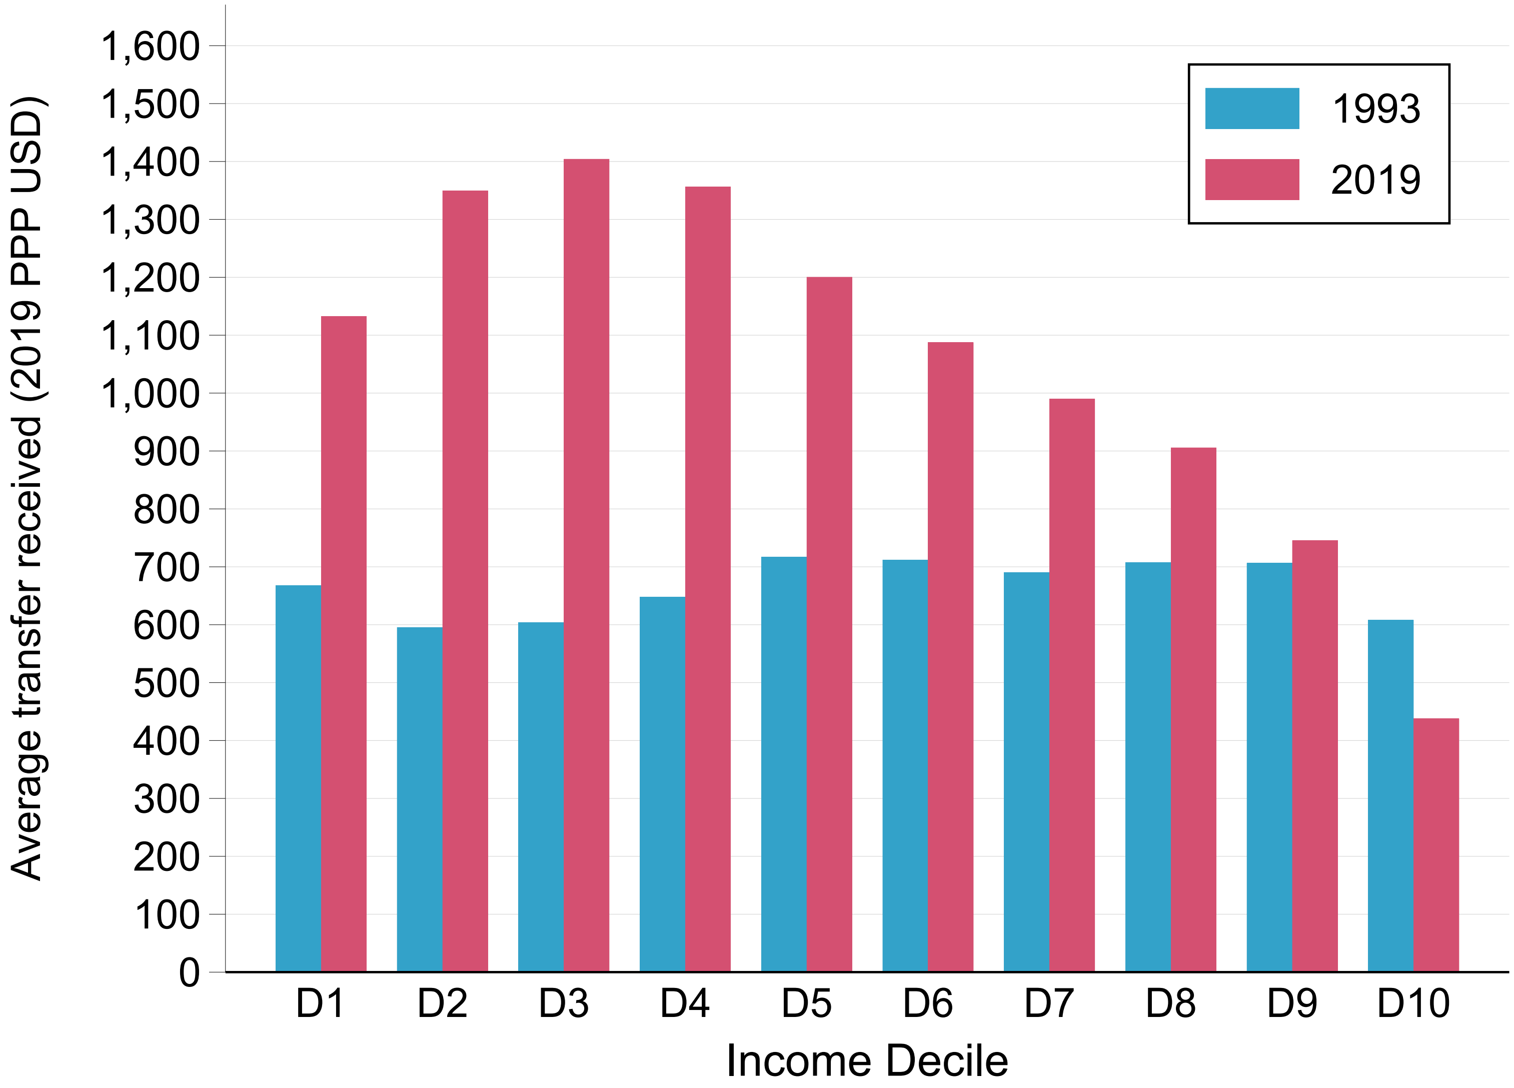

Notes. Author's computations combining surveys, tax, and national accounts data.

E.5 Health

Figure E.19: Level and Composition of Health Expenditure in South Africa, 1993-2019

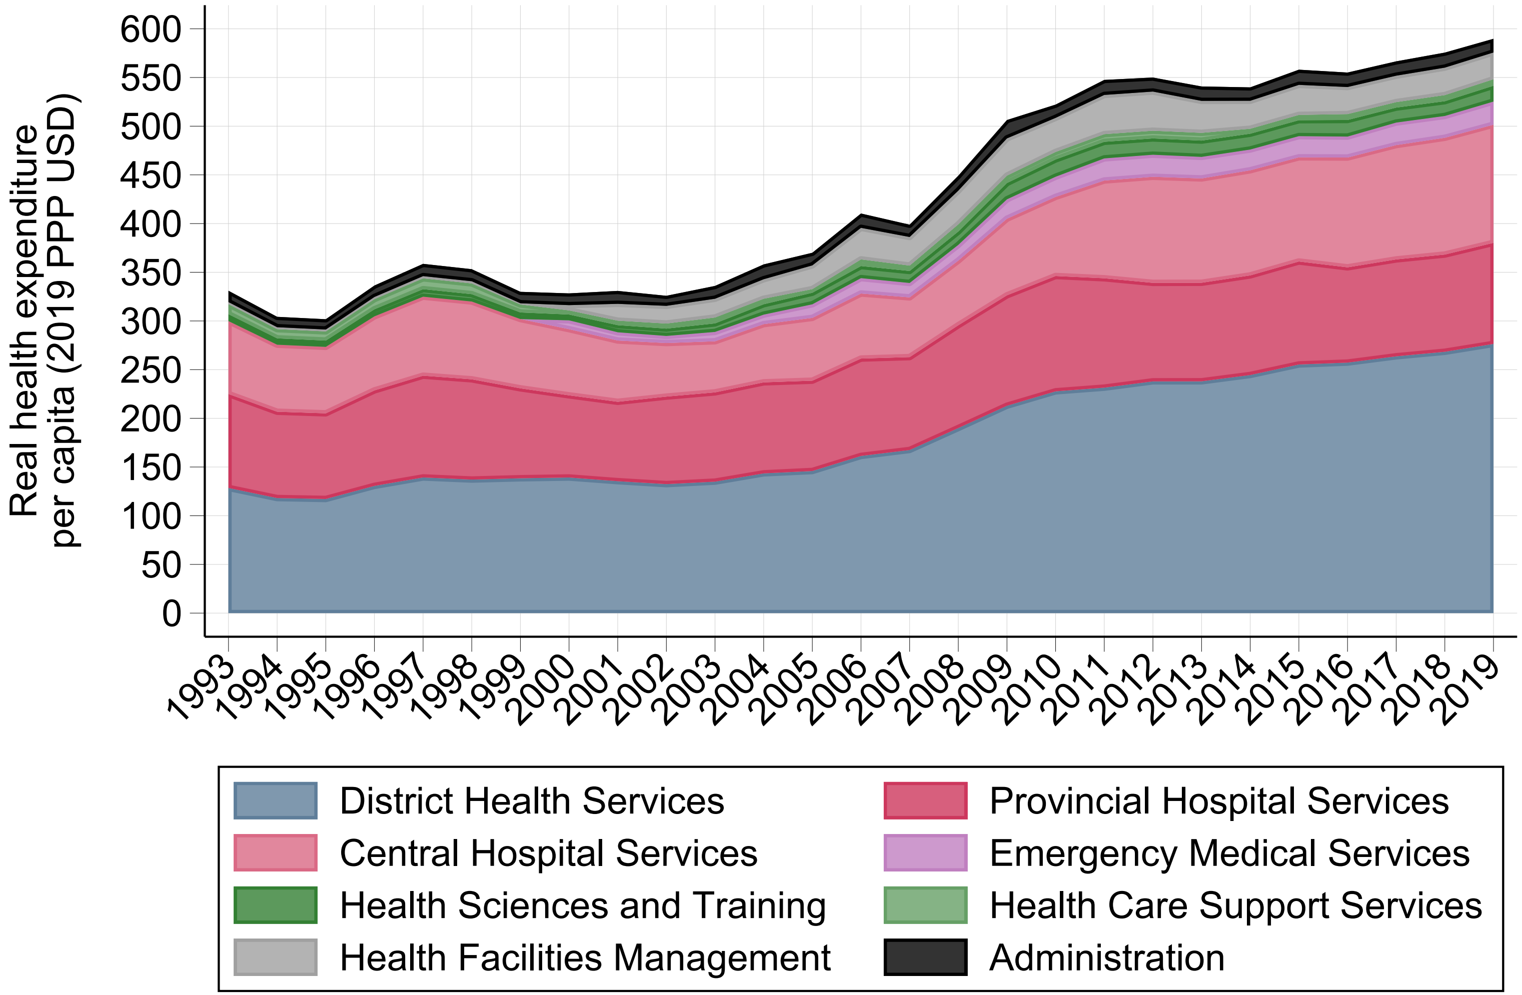

Notes. Author's computations combining data from South African National Treasury Budget Reports (1994-2020) and Provincial Budget Reports (2002-2020).

Figure E.20: Level and Composition of Health Expenditure in South Africa, 1993-2019: Clinics Versus Hospitals

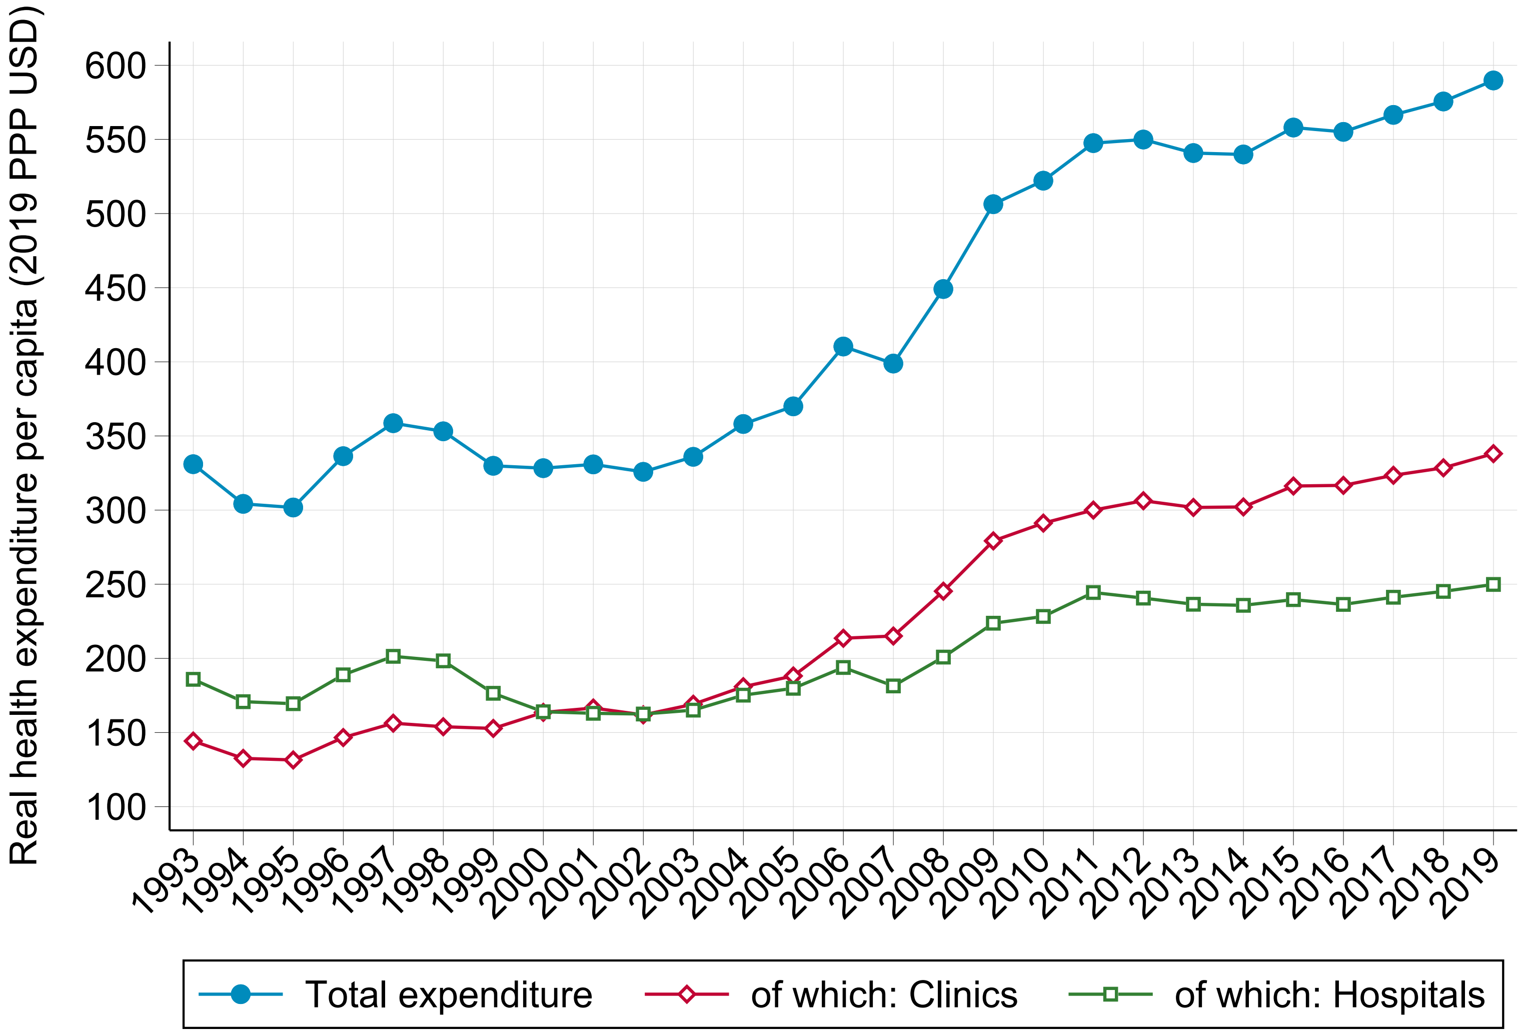

Notes. Author's computations combining data from South African National Treasury Budget Reports (1994-2020) and Provincial Budget Reports (2002-2020).

Figure E.21: Real Health Expenditure Per Capita by South African Province, 1993-2019

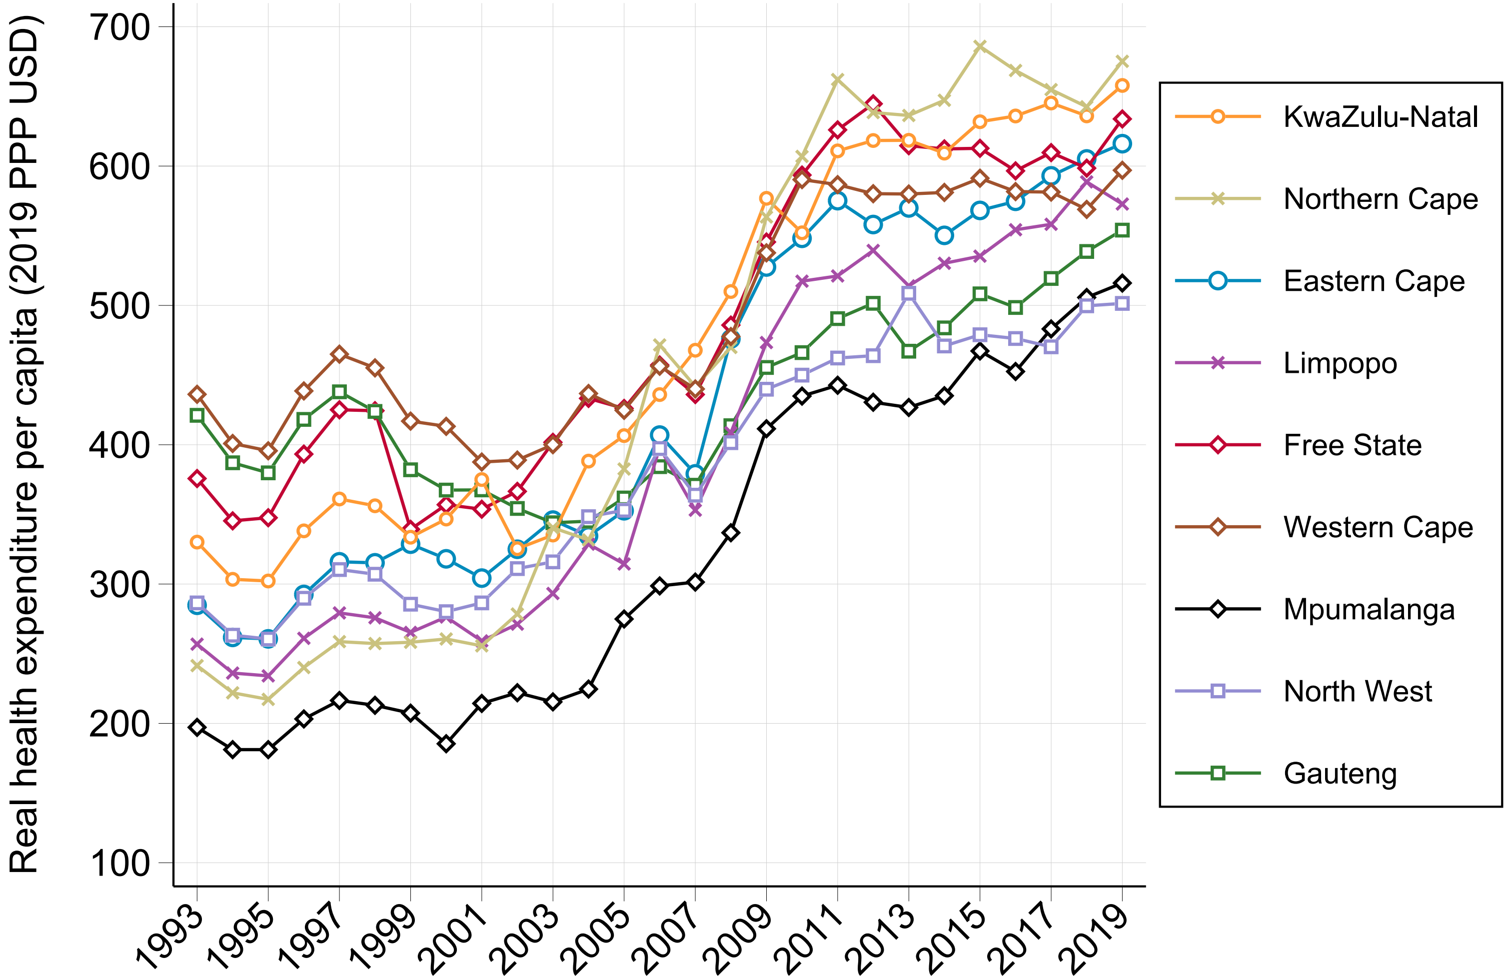

Notes. Author's computations combining data from South African National Treasury Budget Reports (1994-2020) and Provincial Budget Reports (2002-2020).

Figure E.22: Intensity of Use of the Public Healthcare System by Income Quintile, 1995-2019

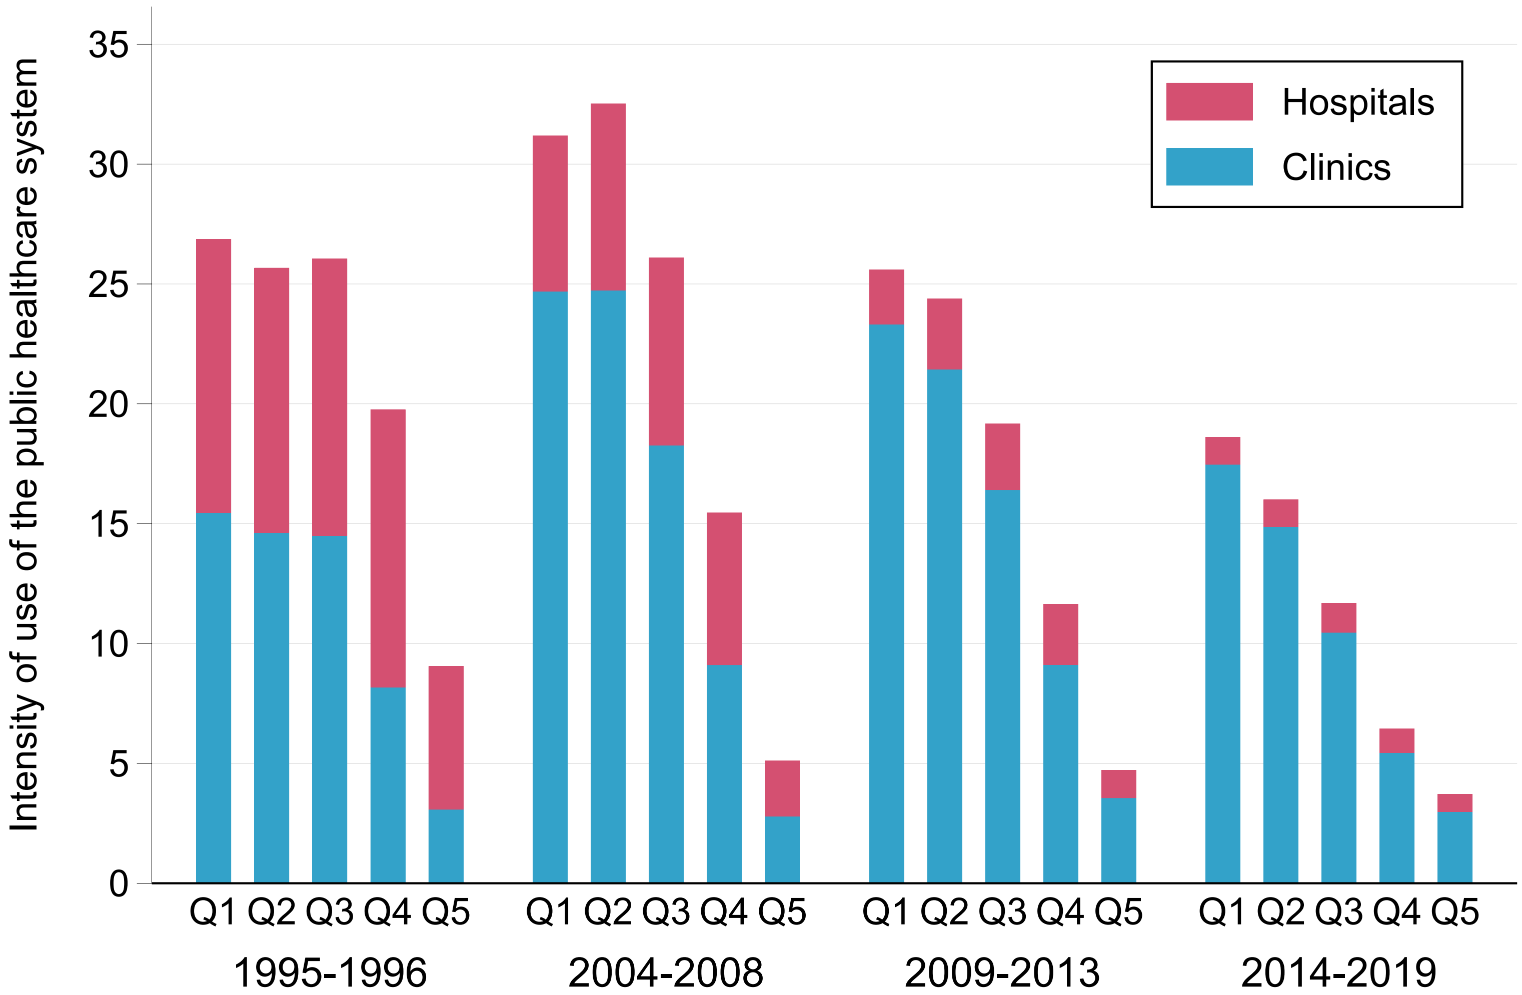

Notes. Author's computations using General Household Surveys (GHS, 2004-2019) and October Household Surveys (OHS, 1995-1996). GHS figures correspond to the share of individuals who consulted a health worker in the past three months and declare going most often to public institutions to do so. OHS figures correspond to the share of individuals who either went to the hospital, or consulted a health worker in the past month, and declare going most often to public institutions to do so.

Figure E.23: Private Healthcare Use by Income Quintile, 1995-2019

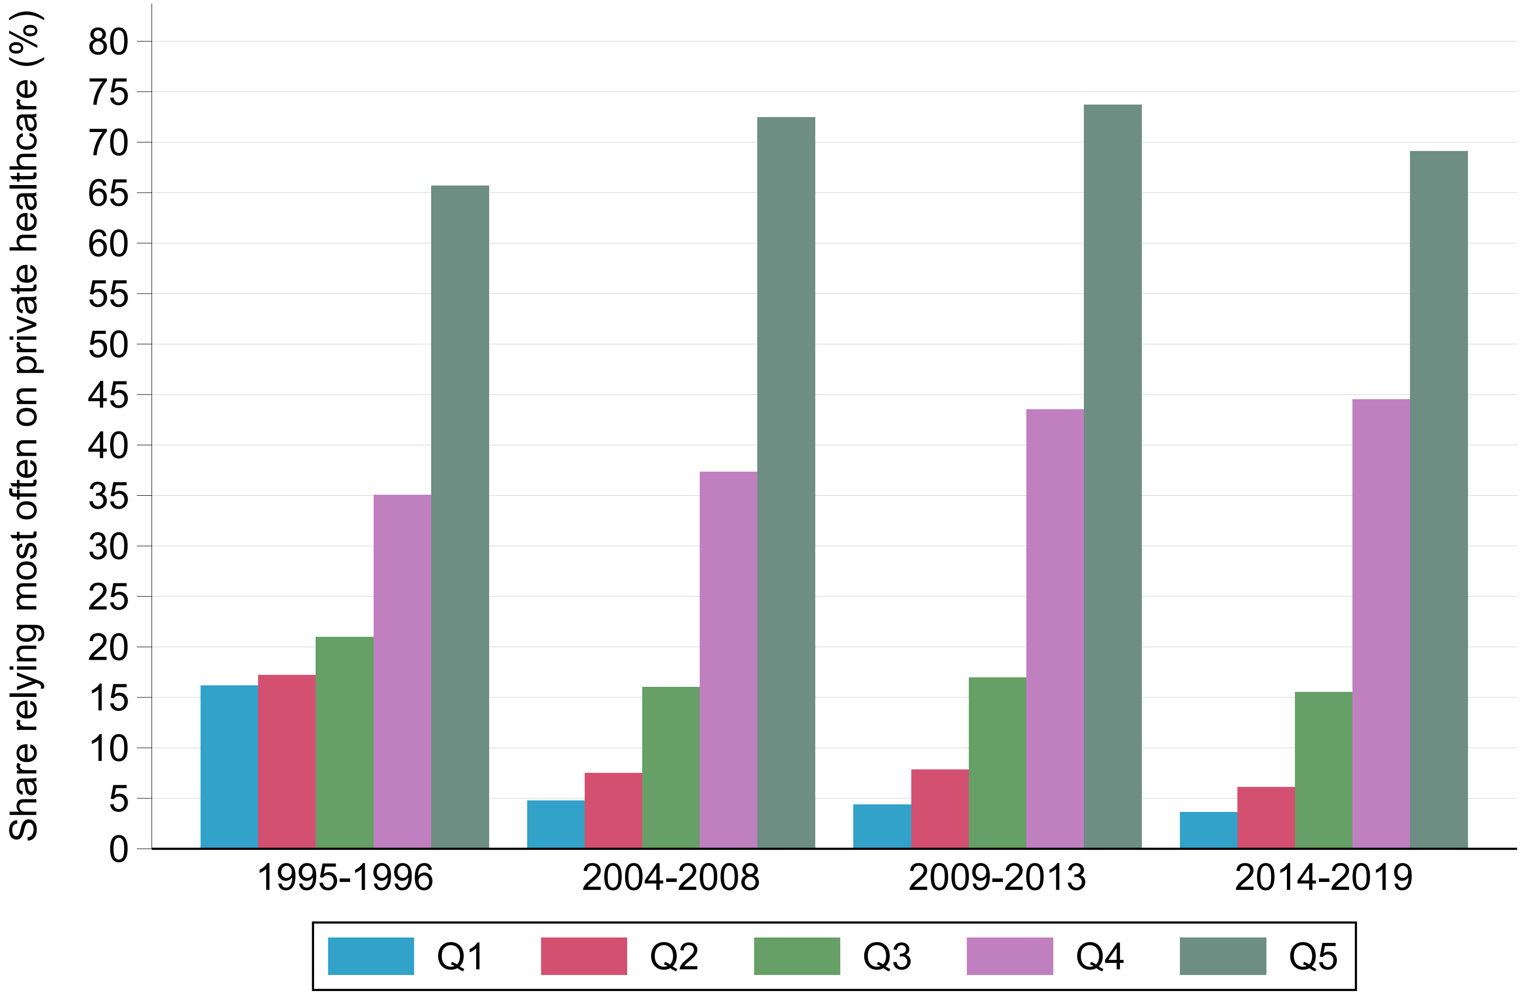

Notes. Author's computations using General Household Surveys (GHS, 2004-2019) and October Household Surveys (OHS, 1995-1996). The figure shows the share of individuals declaring going most often to private clinics or private hospitals for healthcare by income quintile.

Figure E.24: Private Health Insurance Coverage by Income Quintile, 1995-2019

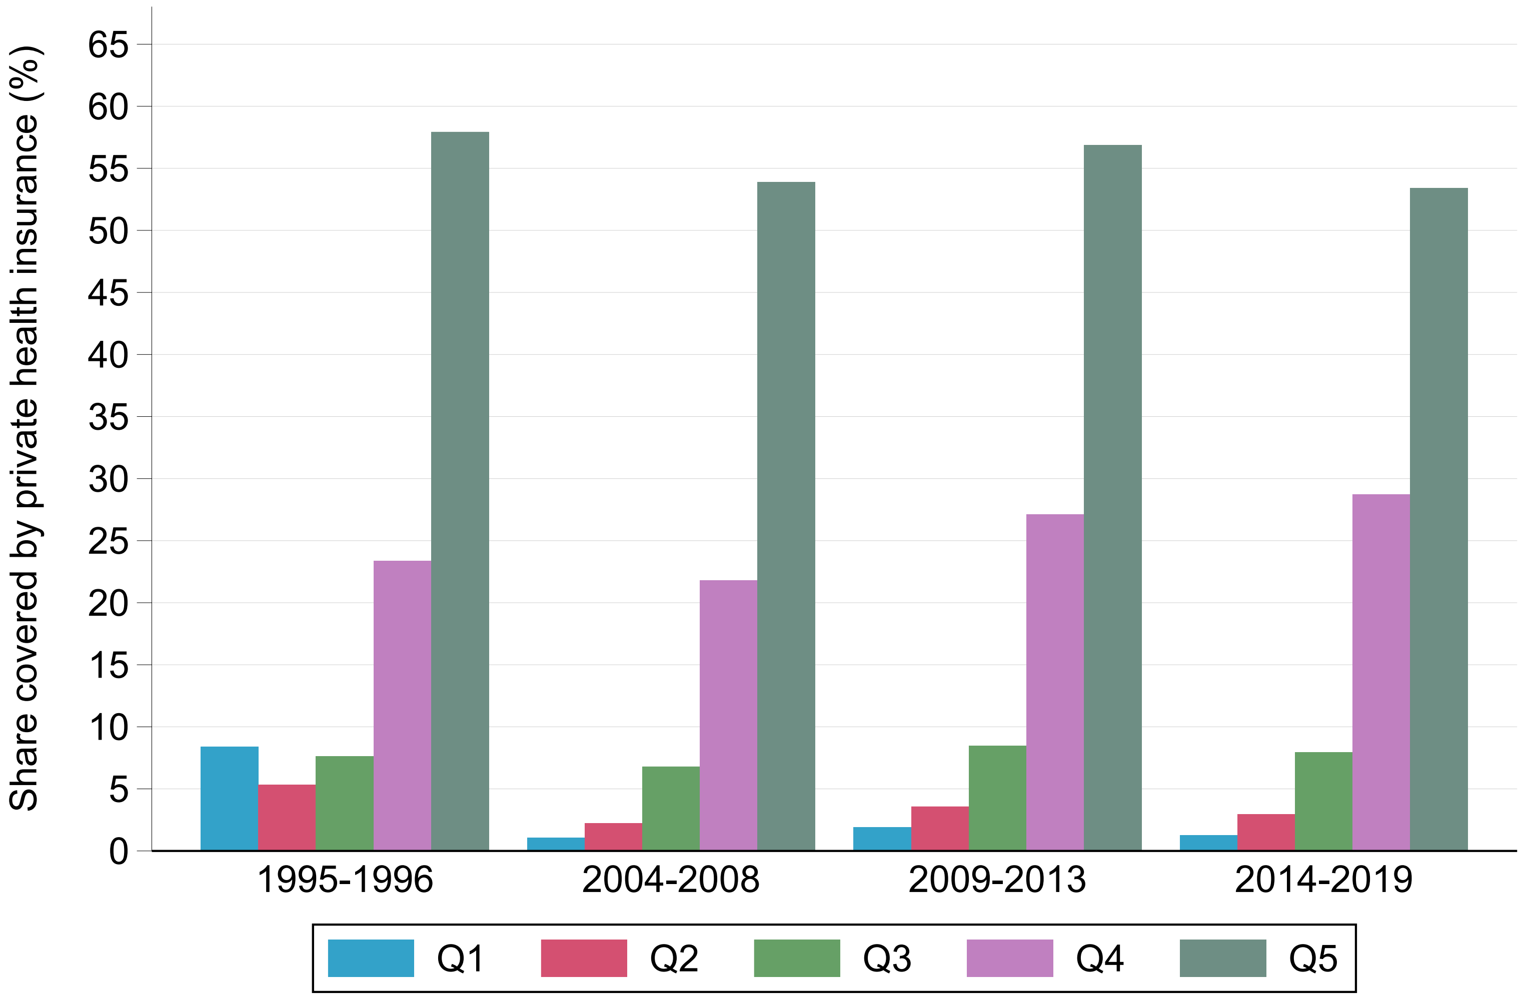

Notes. Author's computations using General Household Surveys (GHS, 2004-2019) and October Household Surveys (OHS, 1995-1996). The figure shows the share of individuals declaring being covered by a medical aid, a medical benefit scheme, or any other form of private insurance by income quintile.

Figure E.25: Average Health Transfer Received by Income Group, 1993-2019

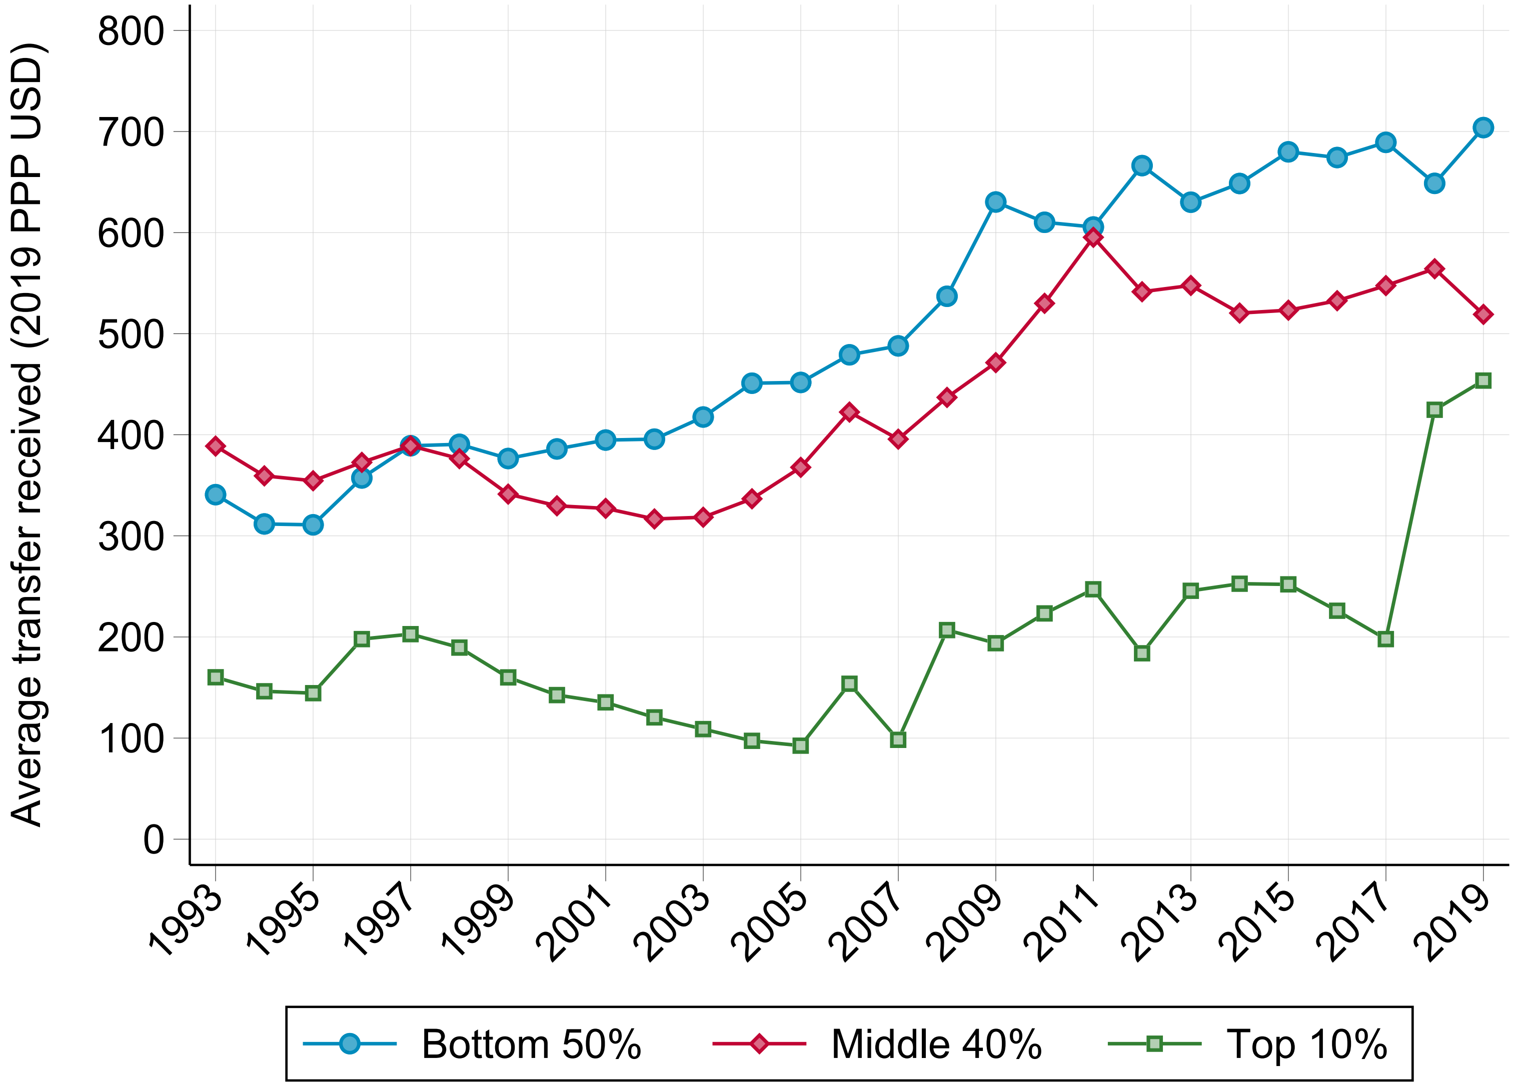

Notes. Author's computations combining surveys, tax, and national accounts data.

Figure E.26: Average Health Transfer Received by Income Decile, 1993-2019

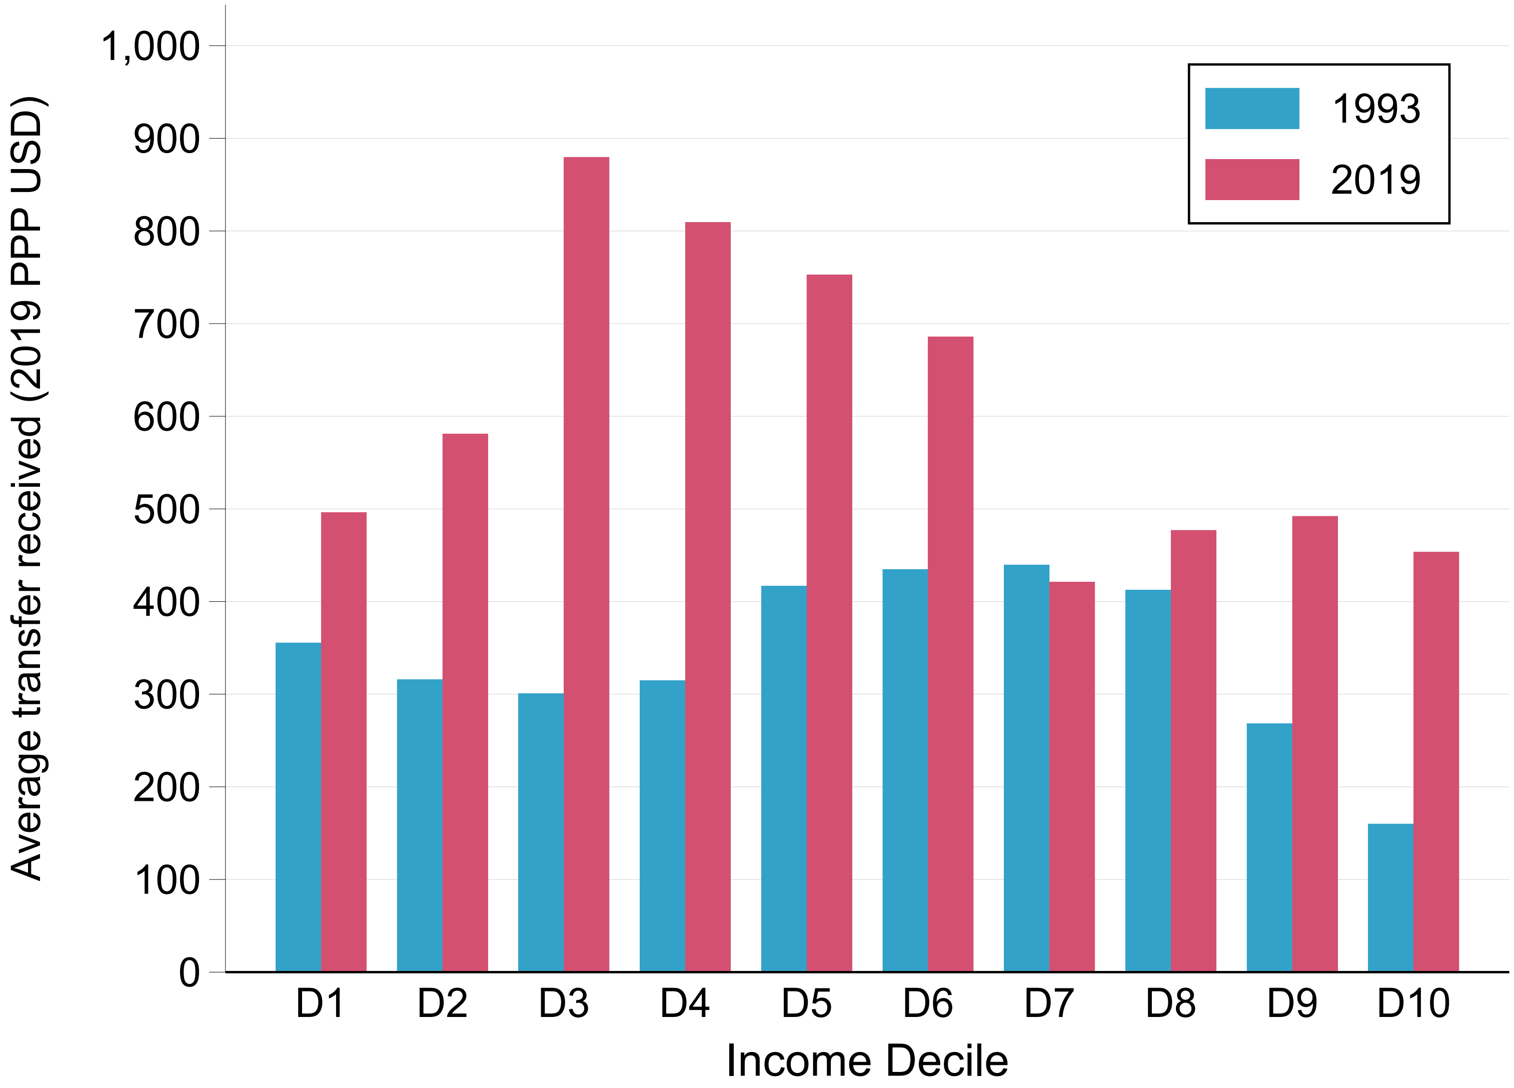

Notes. Author's computations combining surveys, tax, and national accounts data.

E.6 Local Government

Figure E.27: Level and Composition of Local Government Expenditure, 2001-2019

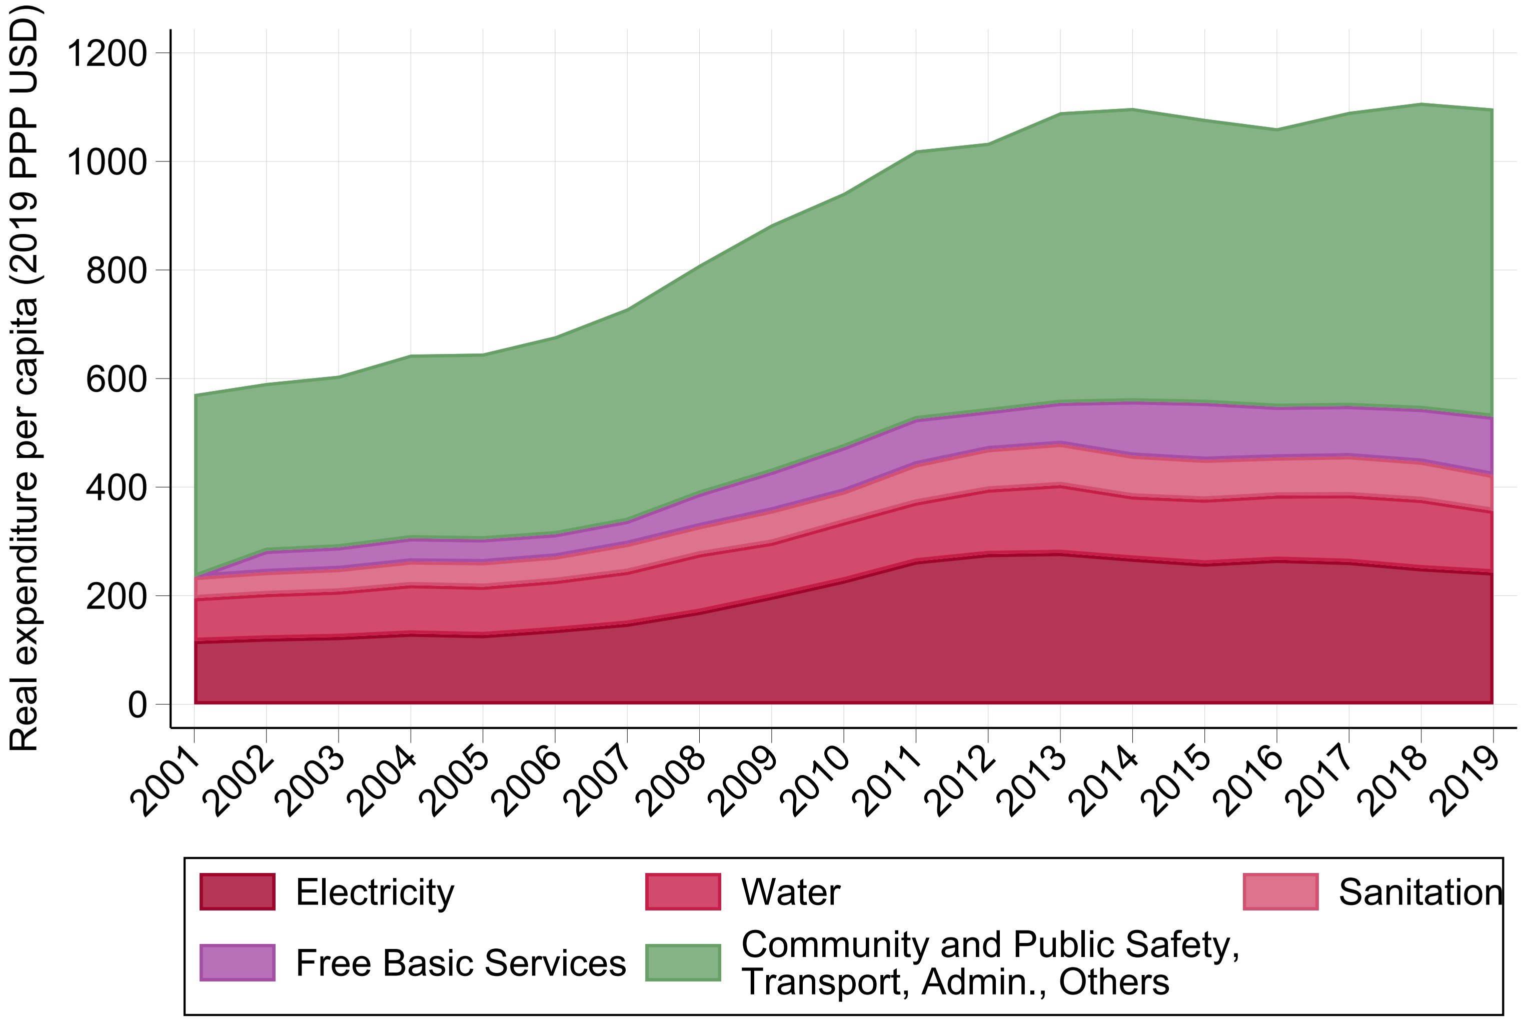

Notes. Author's computations combining data from Local Government Budget Reports.

Figure E.28: Local Government Expenditure in South Africa by Type of Municipality, 2003-2019

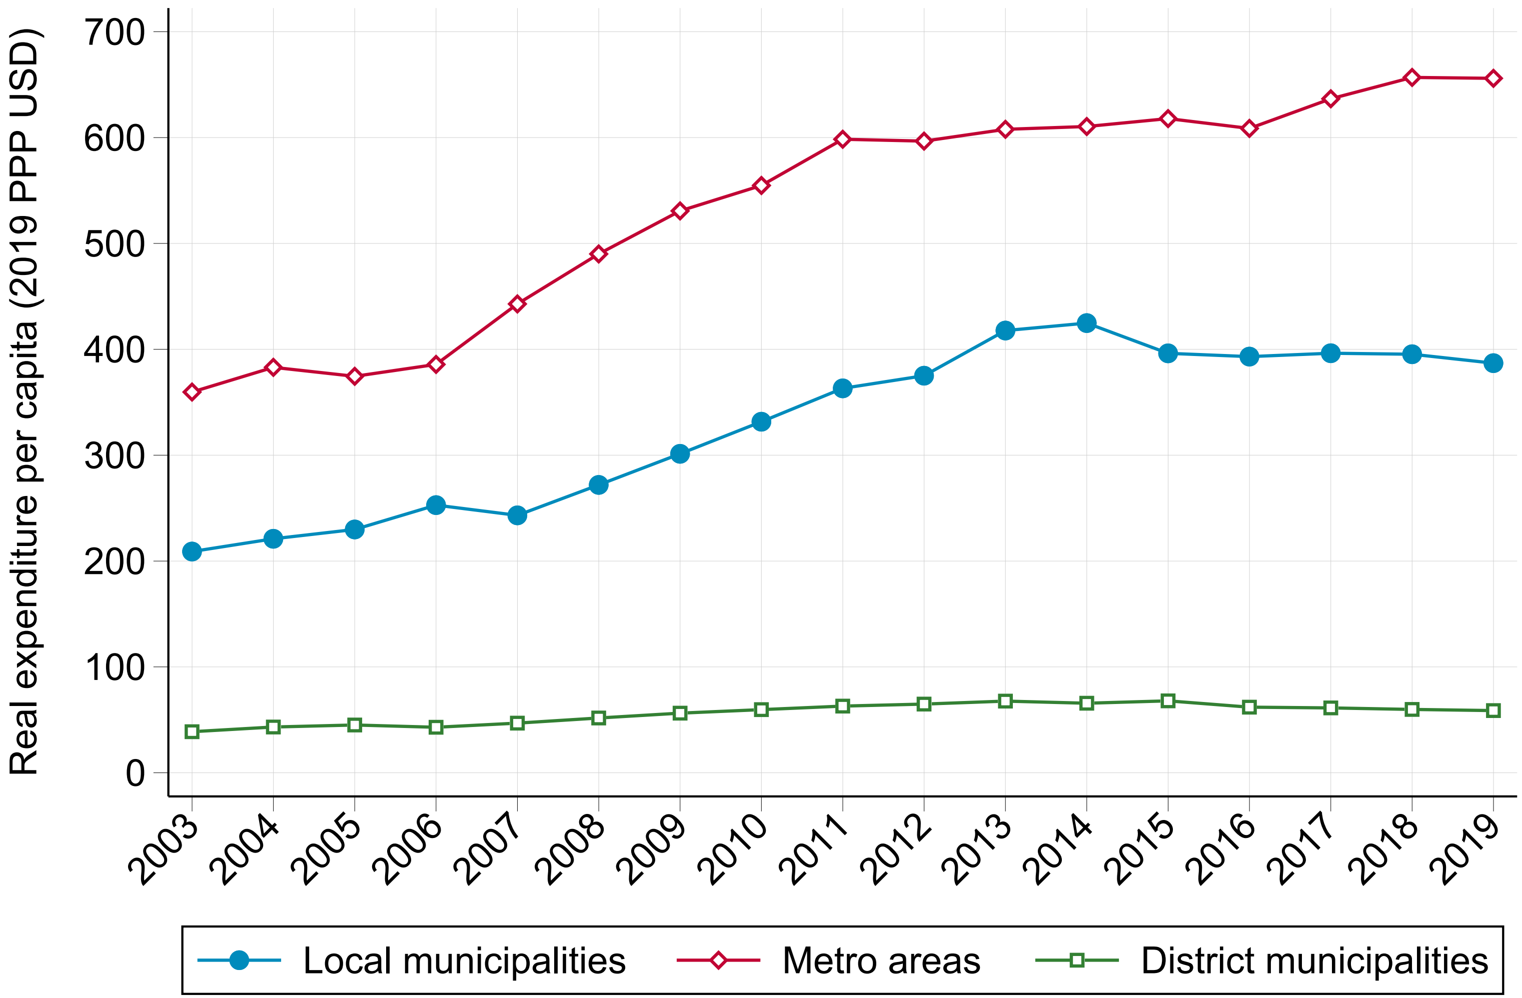

Notes. Author's computations combining data from Local Government Budget Reports.

Figure E.29: The Decline of Spatial Inequalities in Local Public Goods:
Total Expenditure in 2003 Versus 2003-2019 Growth Rate

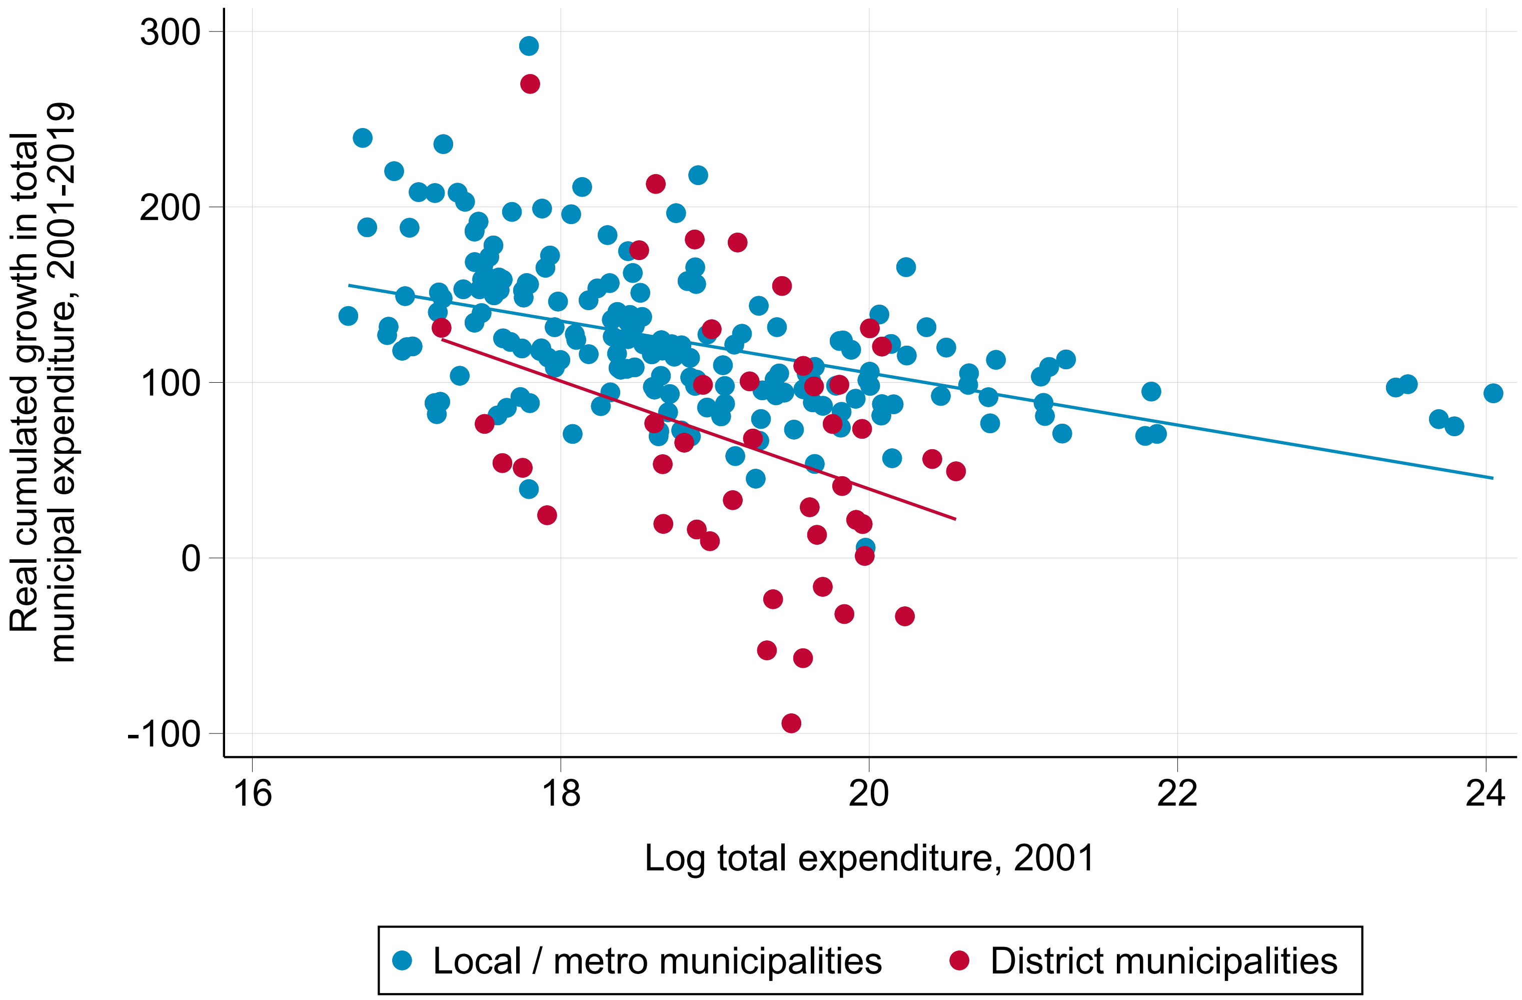

Notes. Author's computations combining data from Local Government Budget Reports.

Figure E.30: The Decline of Spatial Inequalities in Local Public Goods:
Kernel Density of Local Municipality Total Expenditure, 2001-2019

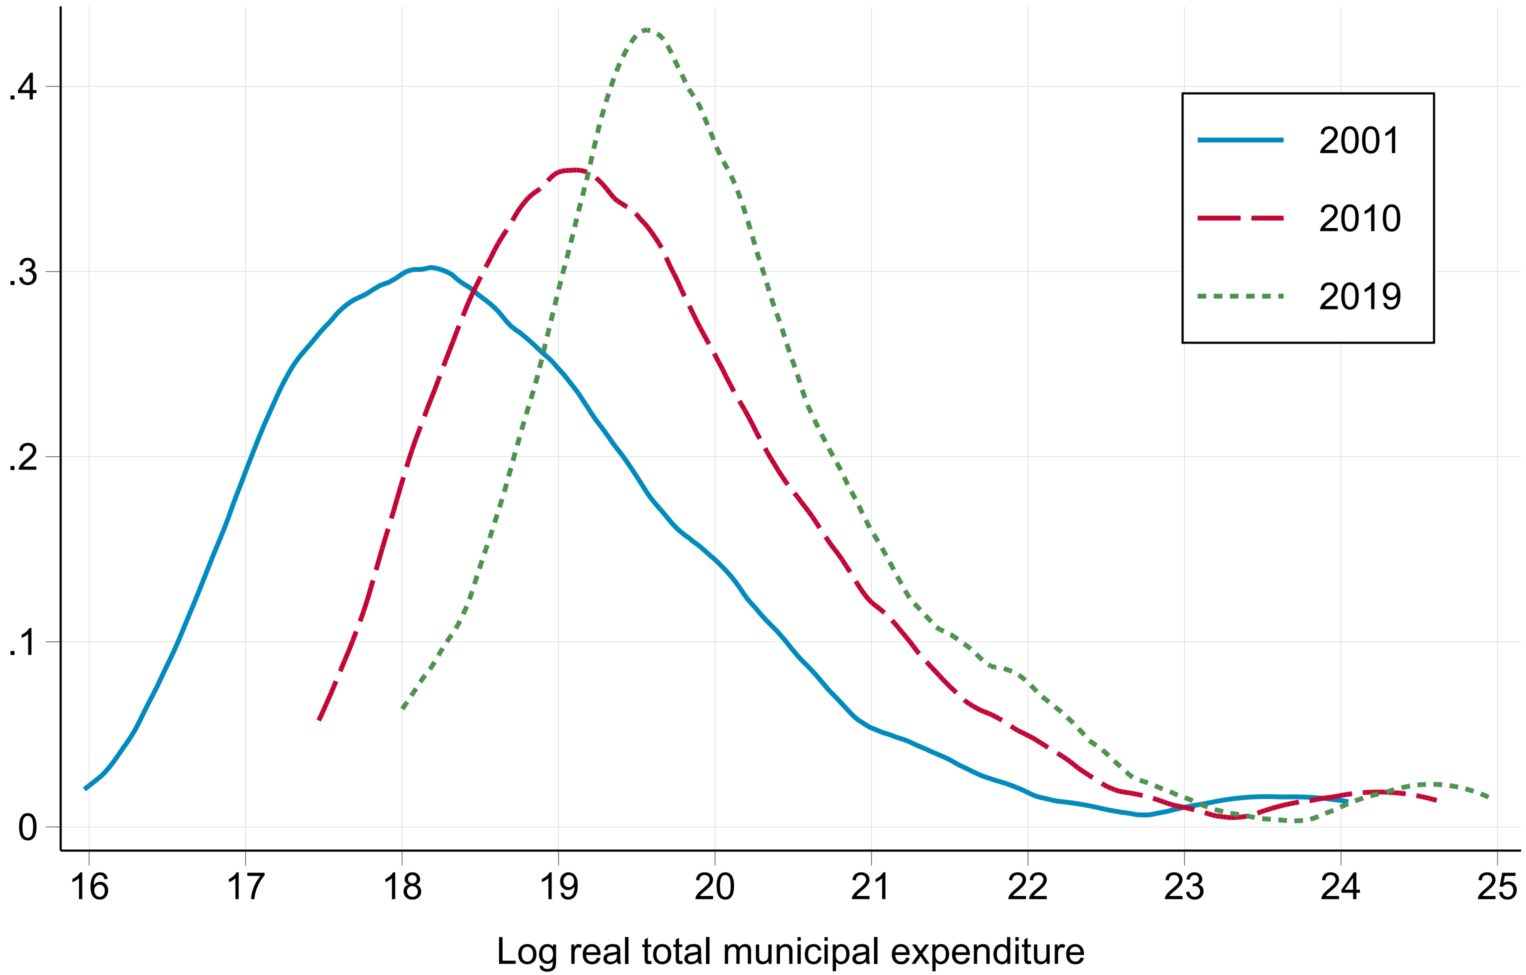

Notes. Author's computations combining data from Local Government Budget Reports.

Figure E.31: Average Local Government Transfer Received by Income Group, 1993-2019

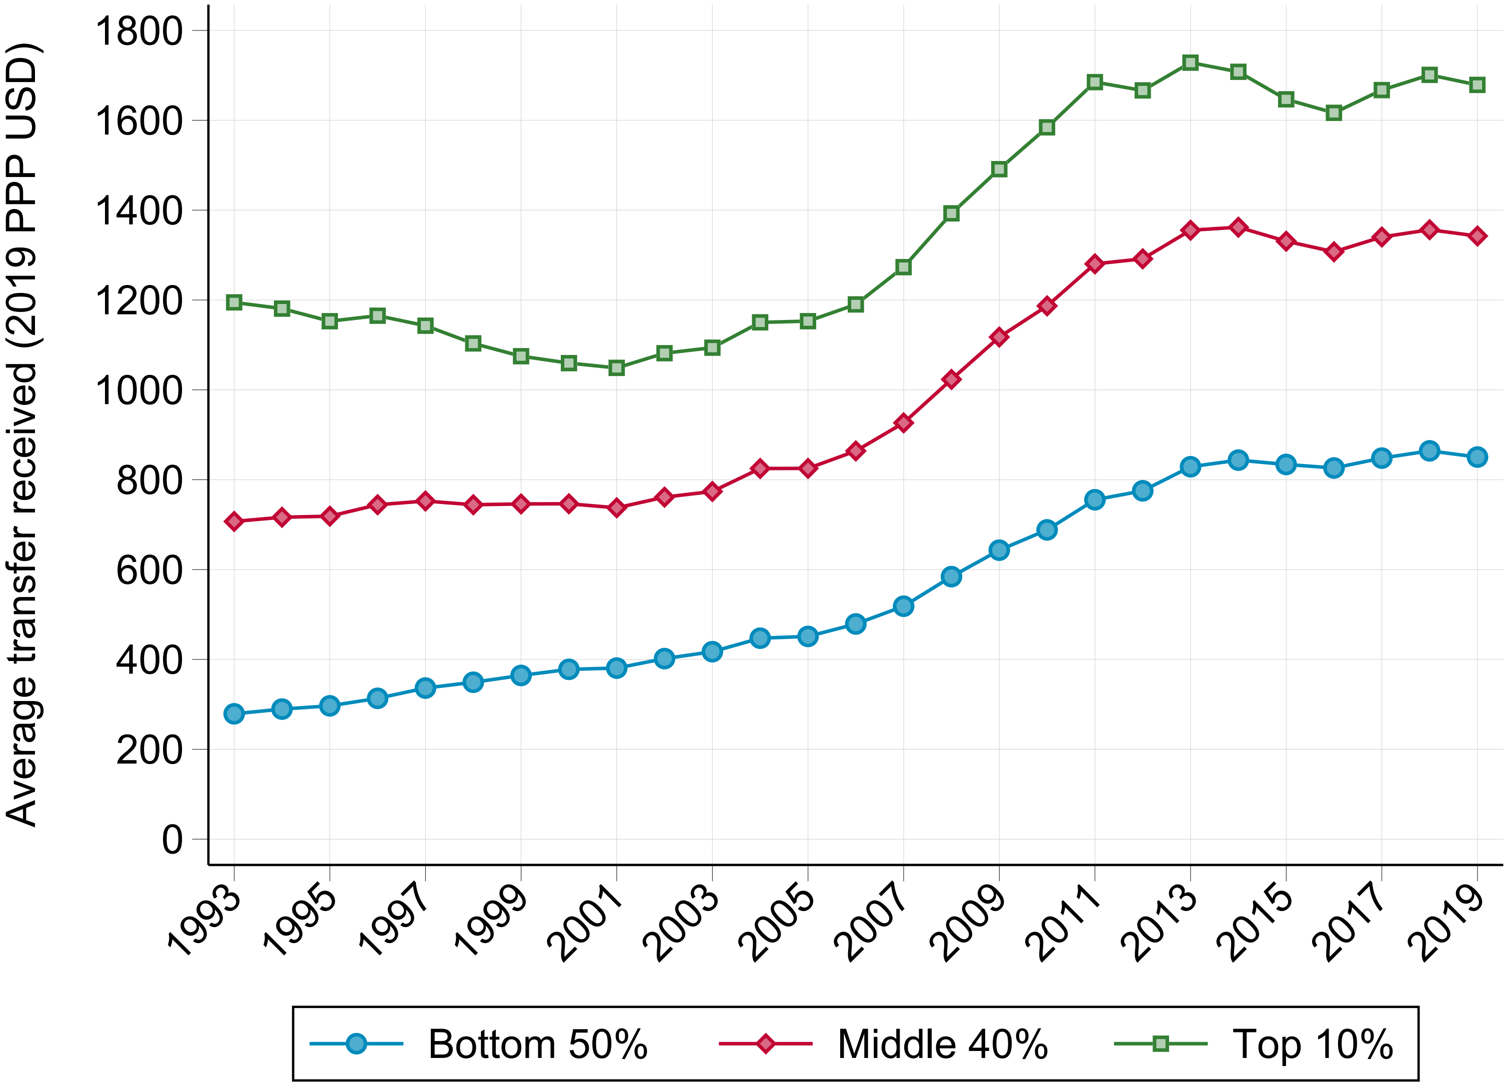

Notes. Author's computations combining surveys, tax, and national accounts data.

Figure E.32: Average Local Government Transfer Received by Income Decile, 1993-2019

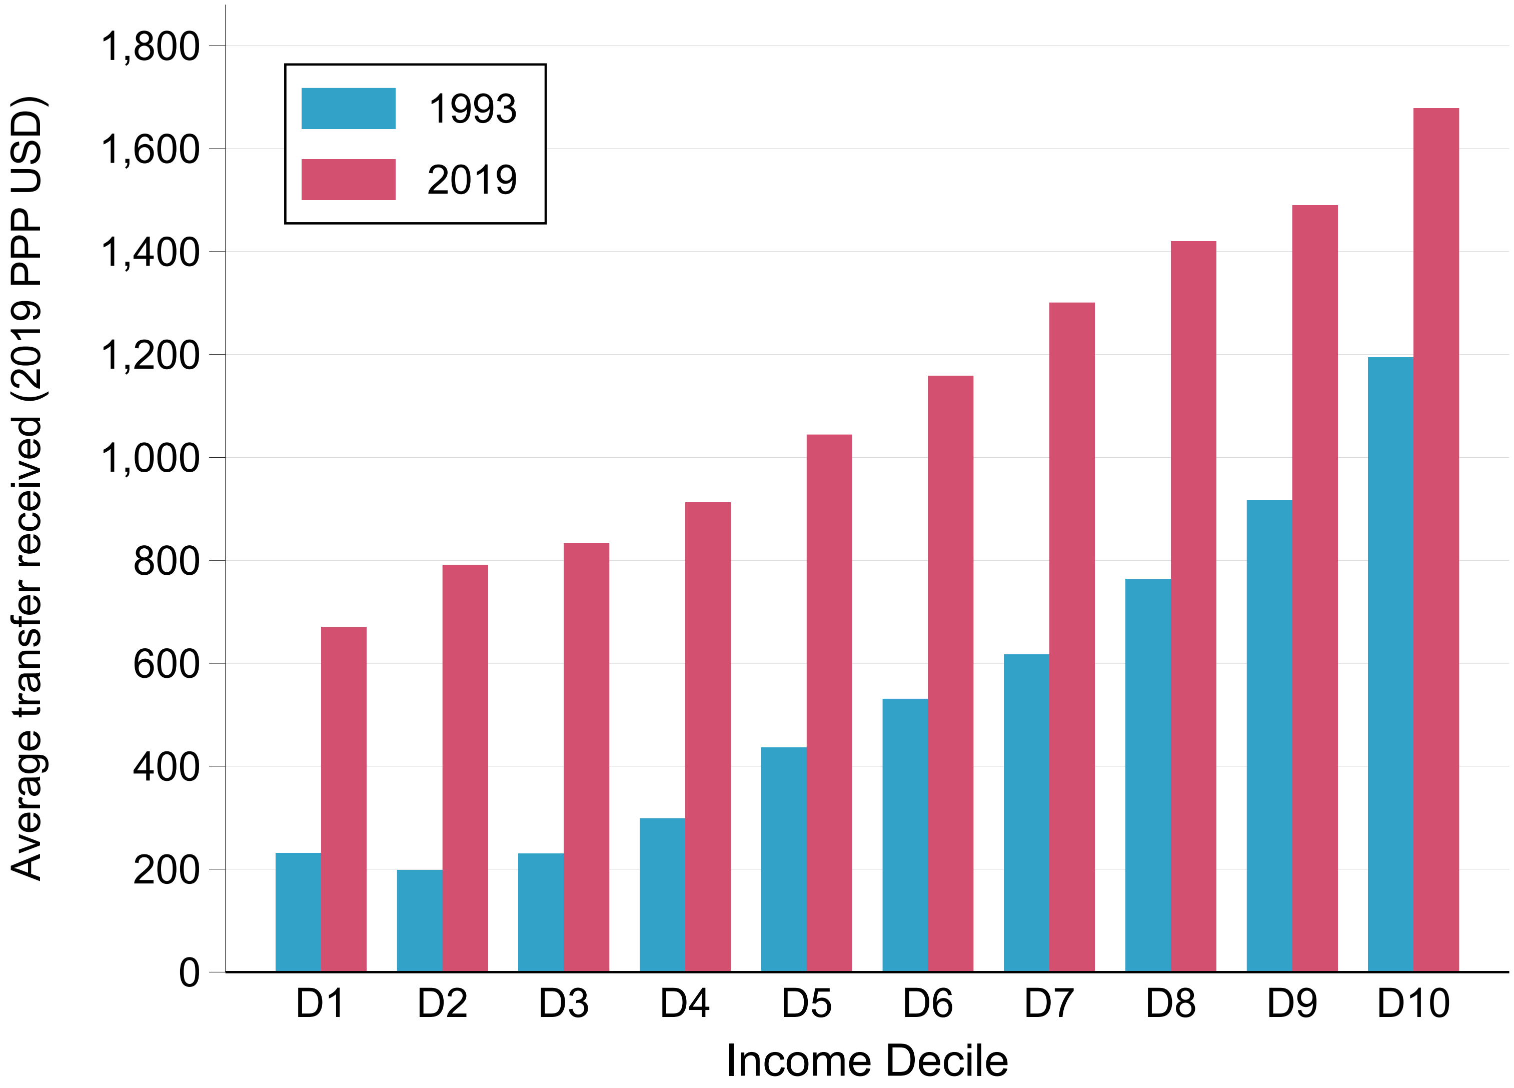

Notes. Author's computations combining surveys, tax, and national accounts data.

Figure E.33: Access to Free Basic Electricity by Income Group, 2004-2019

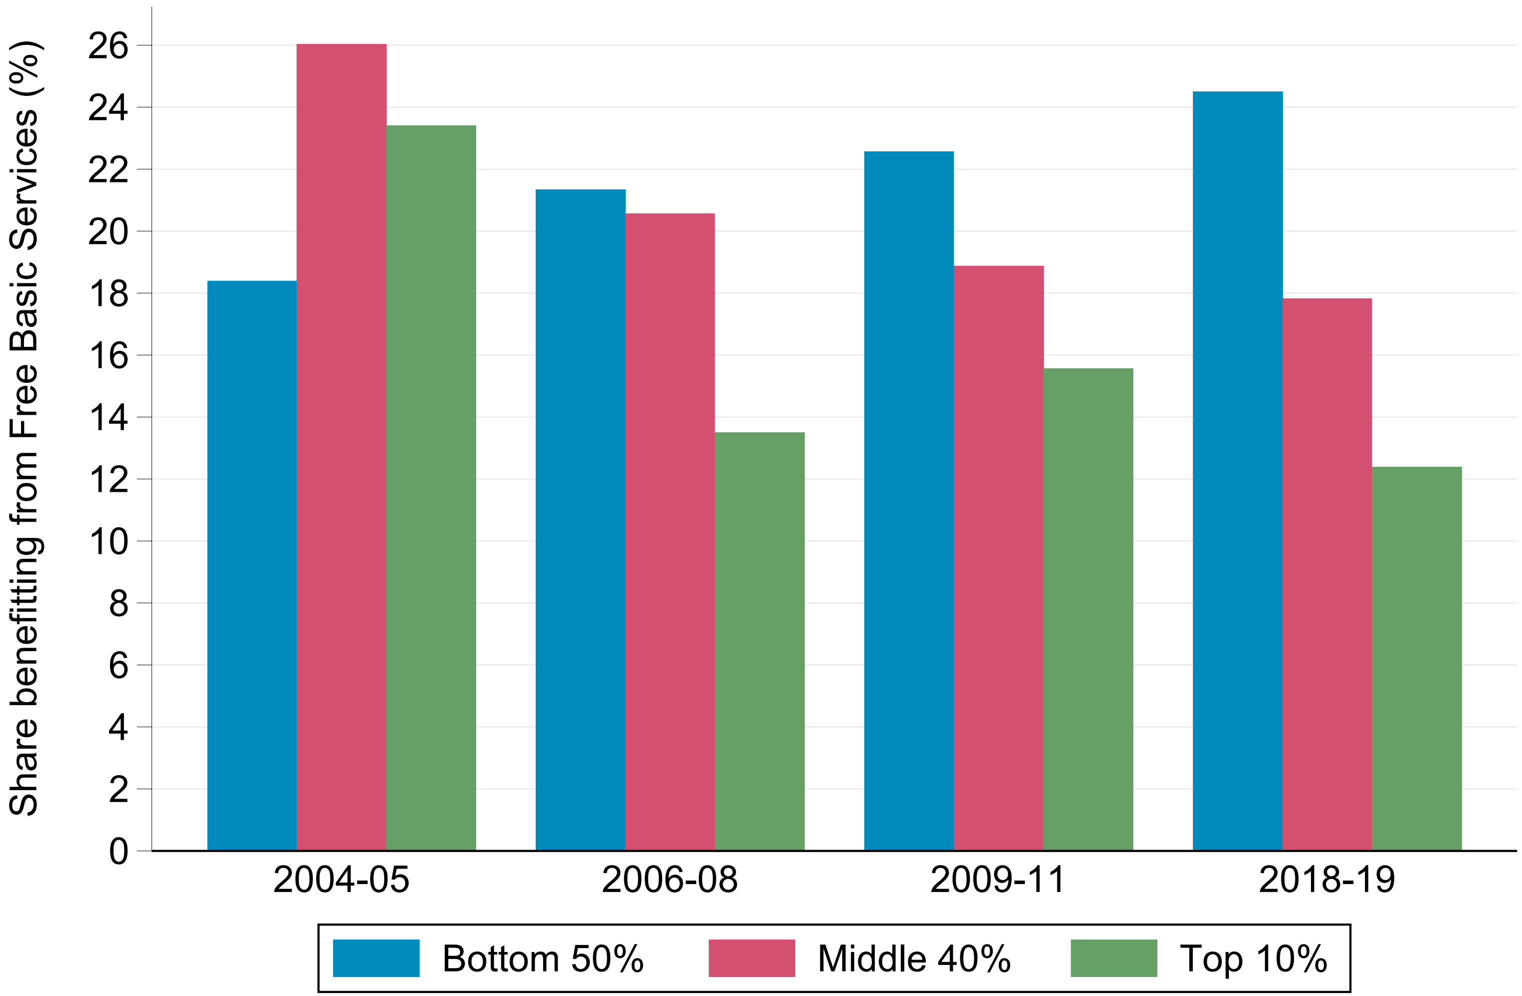

Notes. Author's computations combining data from General Household Surveys. The figure represents the share of individuals who declare benefiting from free basic electricity in their municipality of residence. Income groups are defined based on household expenditure per capita.

E.7 Housing

Figure E.34: Share of Individuals Living in Government-Subsidized Dwelling by Income Group, 2008-2019

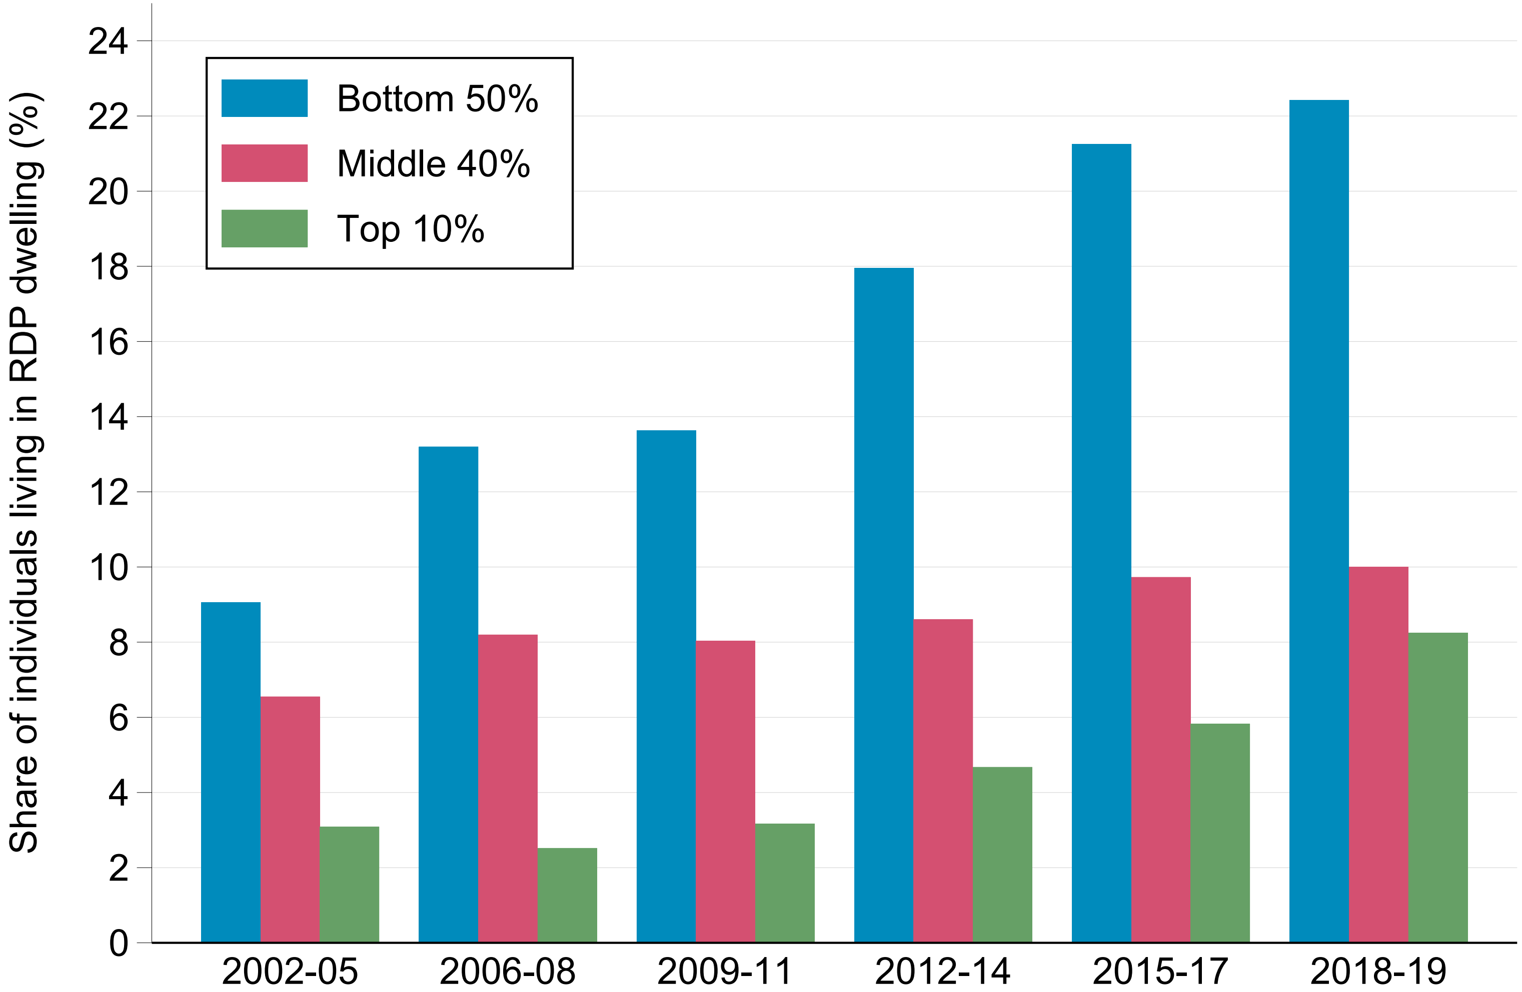

Notes. Author's computations combining General Household Surveys. The figure shows the share of individuals living in households with at least one person who declared receiving "assistance from government to obtain this, or any other dwelling."

Figure E.35: Average Housing Transfer Received by Income Group, 1993-2019

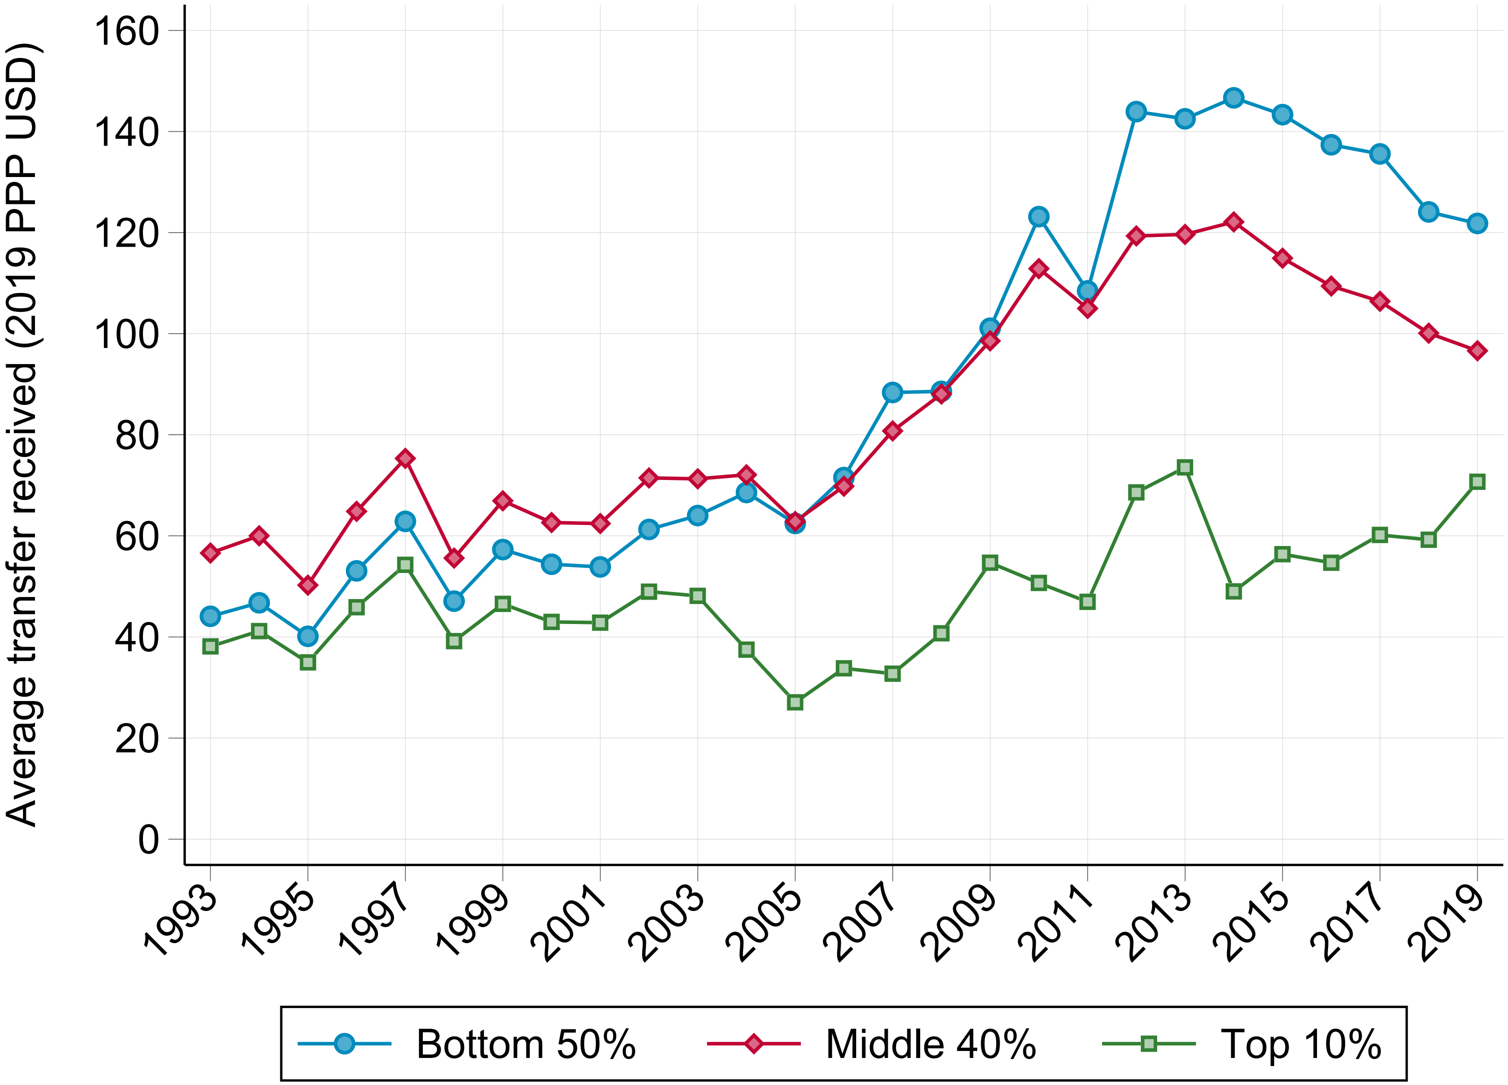

Notes. Author's computations combining surveys, tax, and national accounts data.

Figure E.36: Average Housing Transfer Received by Income Decile, 1993-2019

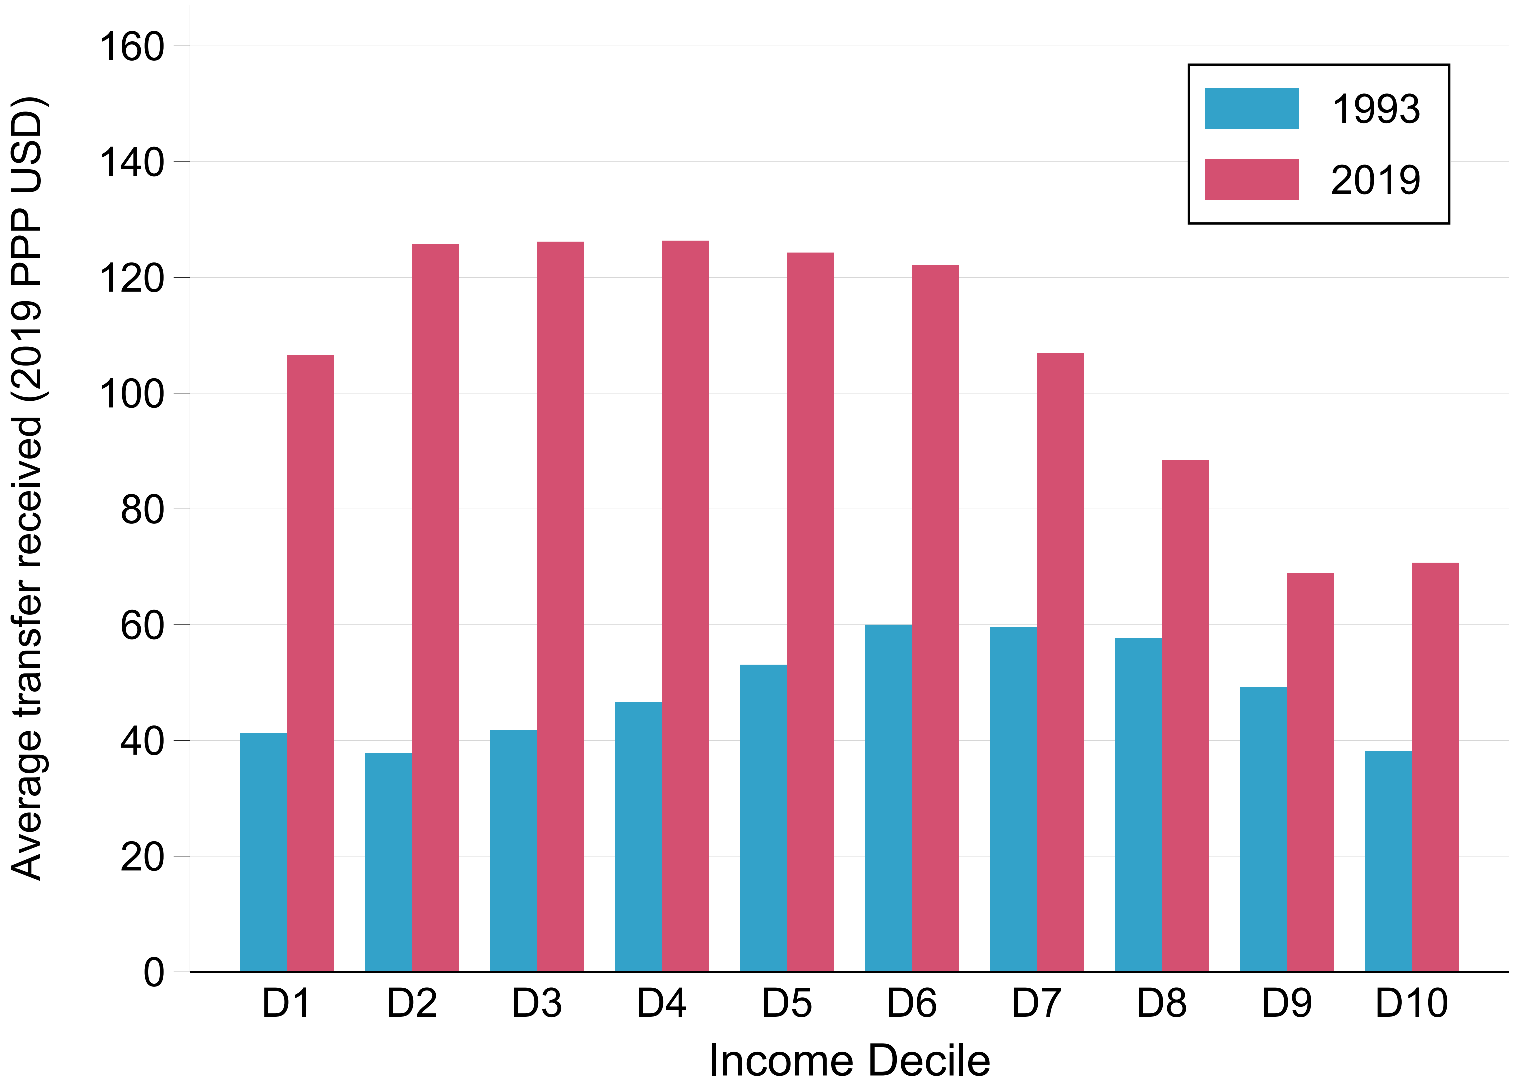

Notes. Author's computations combining surveys, tax, and national accounts data.

E.8 Public Order and Safety

Figure E.37: Level and Composition of Public Order and Safety Expenditure in South Africa, 1993-2019

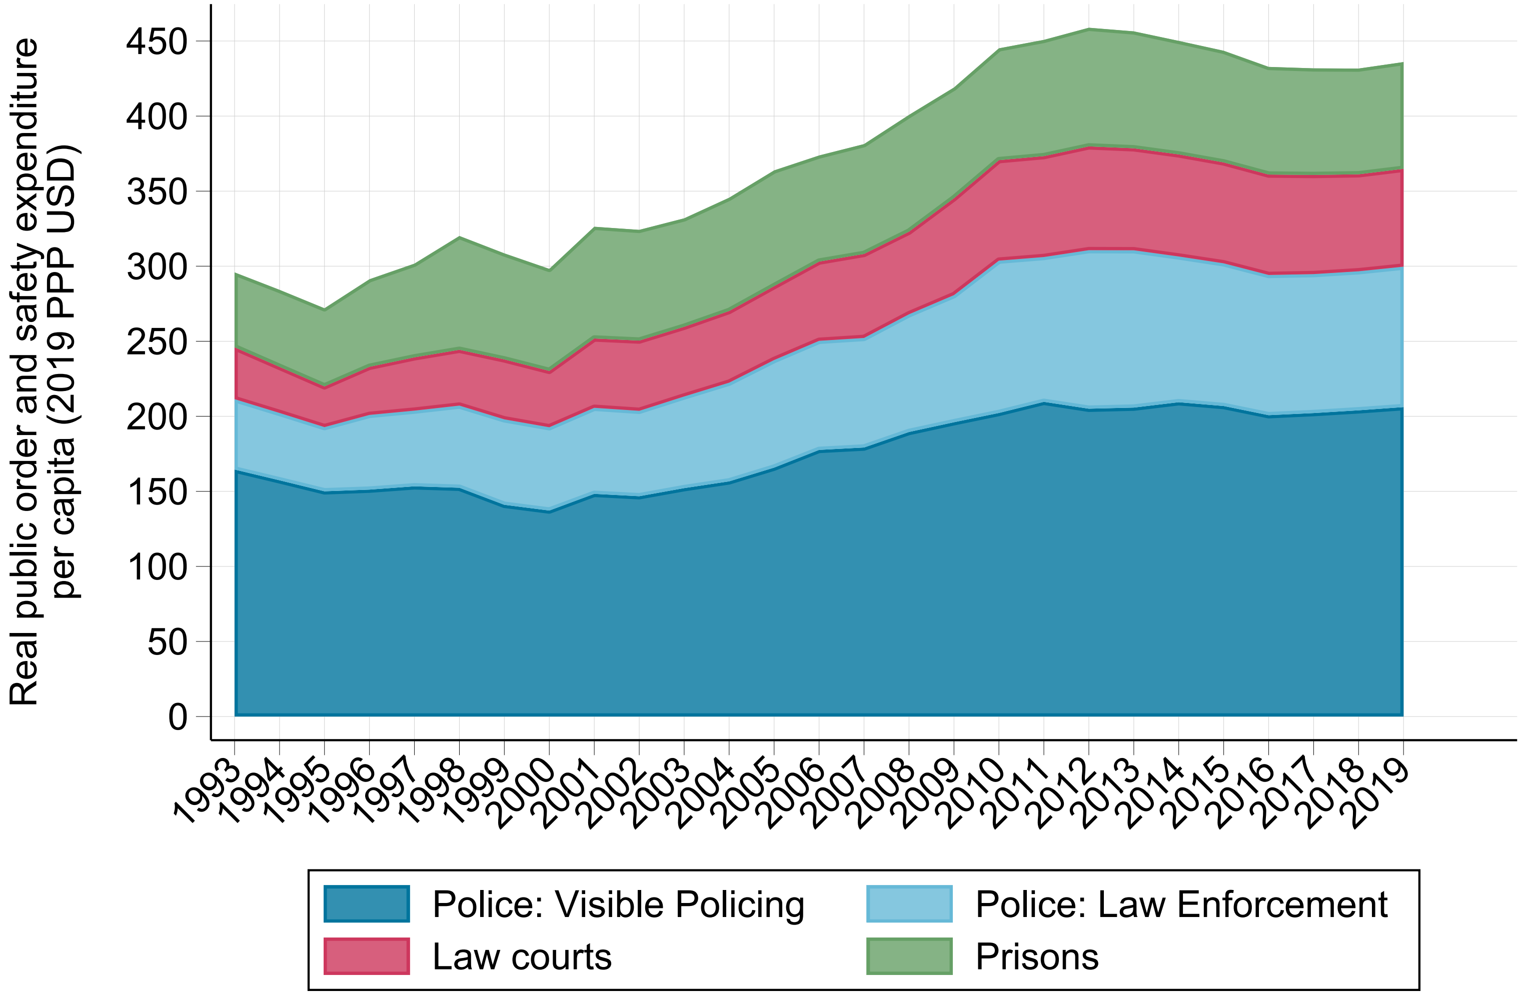

Notes. Author's computations combining data from South African National Treasury Budget Reports (1994-2020).

Figure E.38: Level and Composition of Public Order and Safety Expenditure in South Africa, 1993-2019: Insurance Versus Use

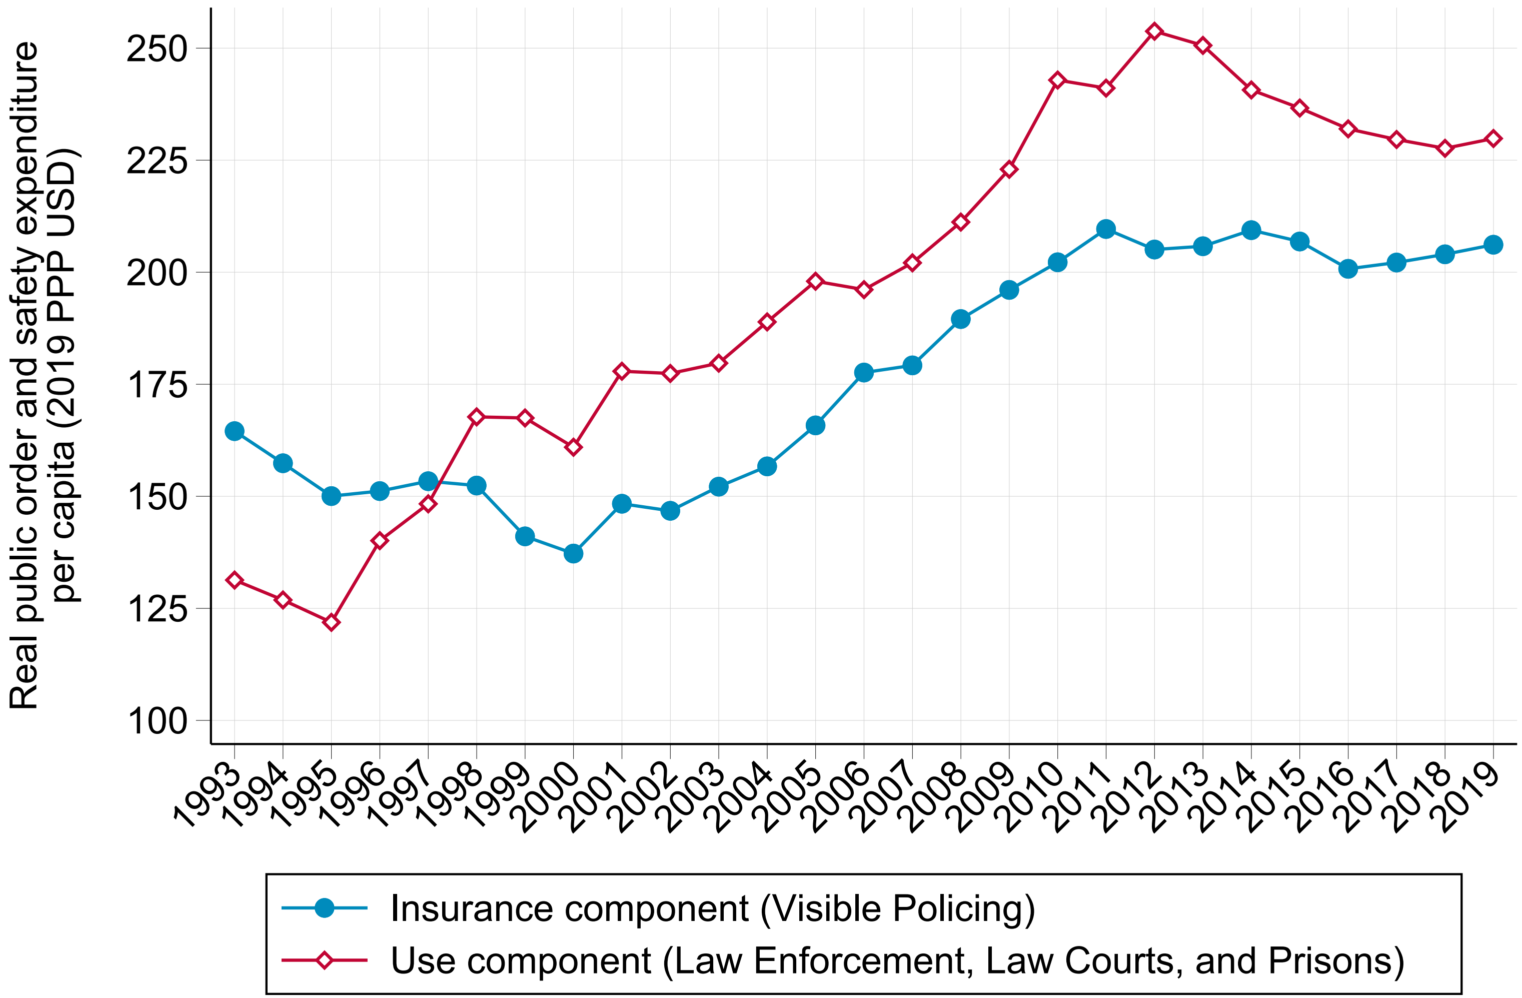

Notes. Author's computations combining data from South African National Treasury Budget Reports (1994-2020).

Figure E.39: Number of Crimes Reported to the Police by Income Quintile, 1998-2017

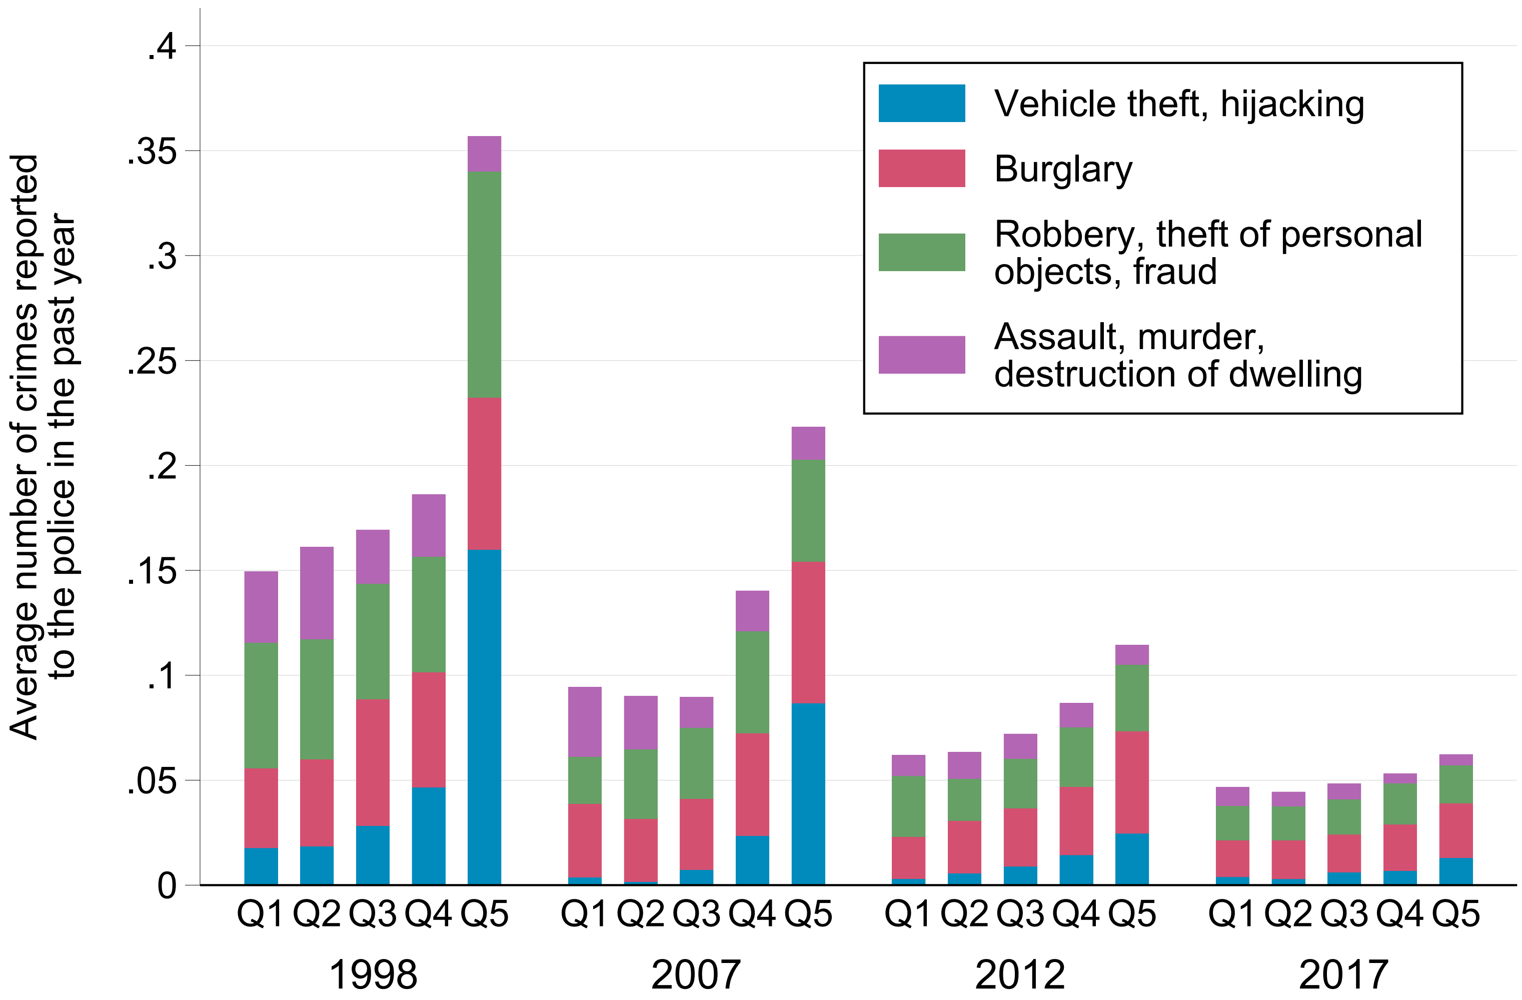

Notes. Author's computations combining data from Victims of Crime Surveys.

Figure E.40: Intensity of Local Police Presence by Income Quintile, 1998-2017

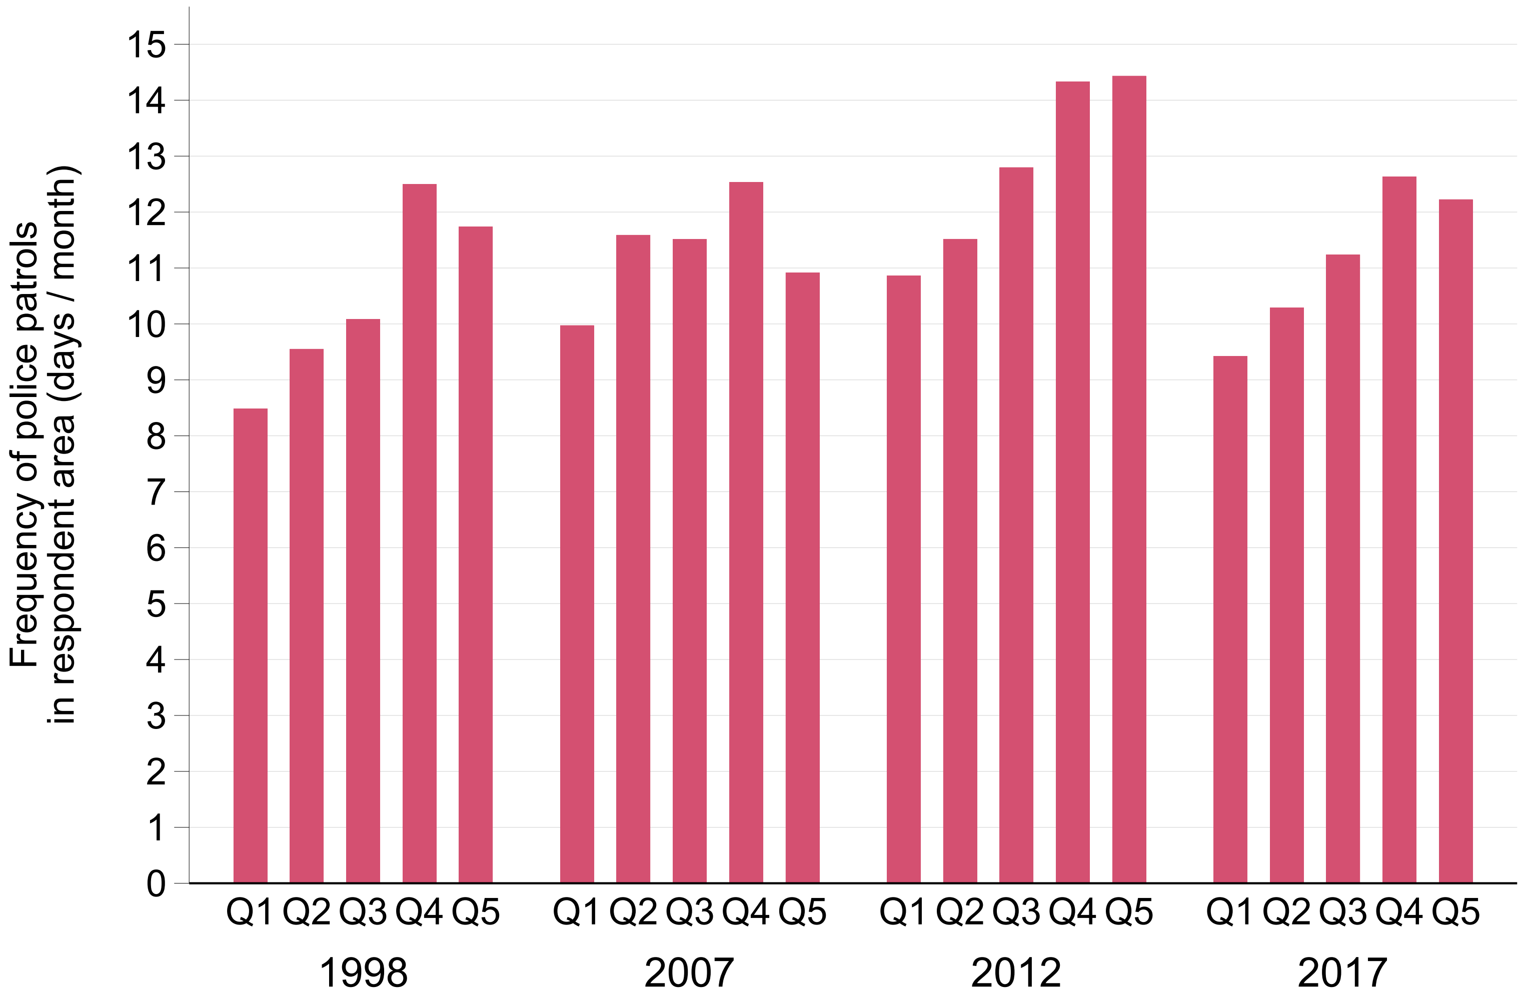

Notes. Author's computations combining data from Victims of Crime Surveys. Figures correspond to the average number of days per month that the respondent declares seeing a police officer in uniform or a police vehicle in her area.

Figure E.41: Average Public Order and Safety Transfer Received by Income Group, 1993-2019

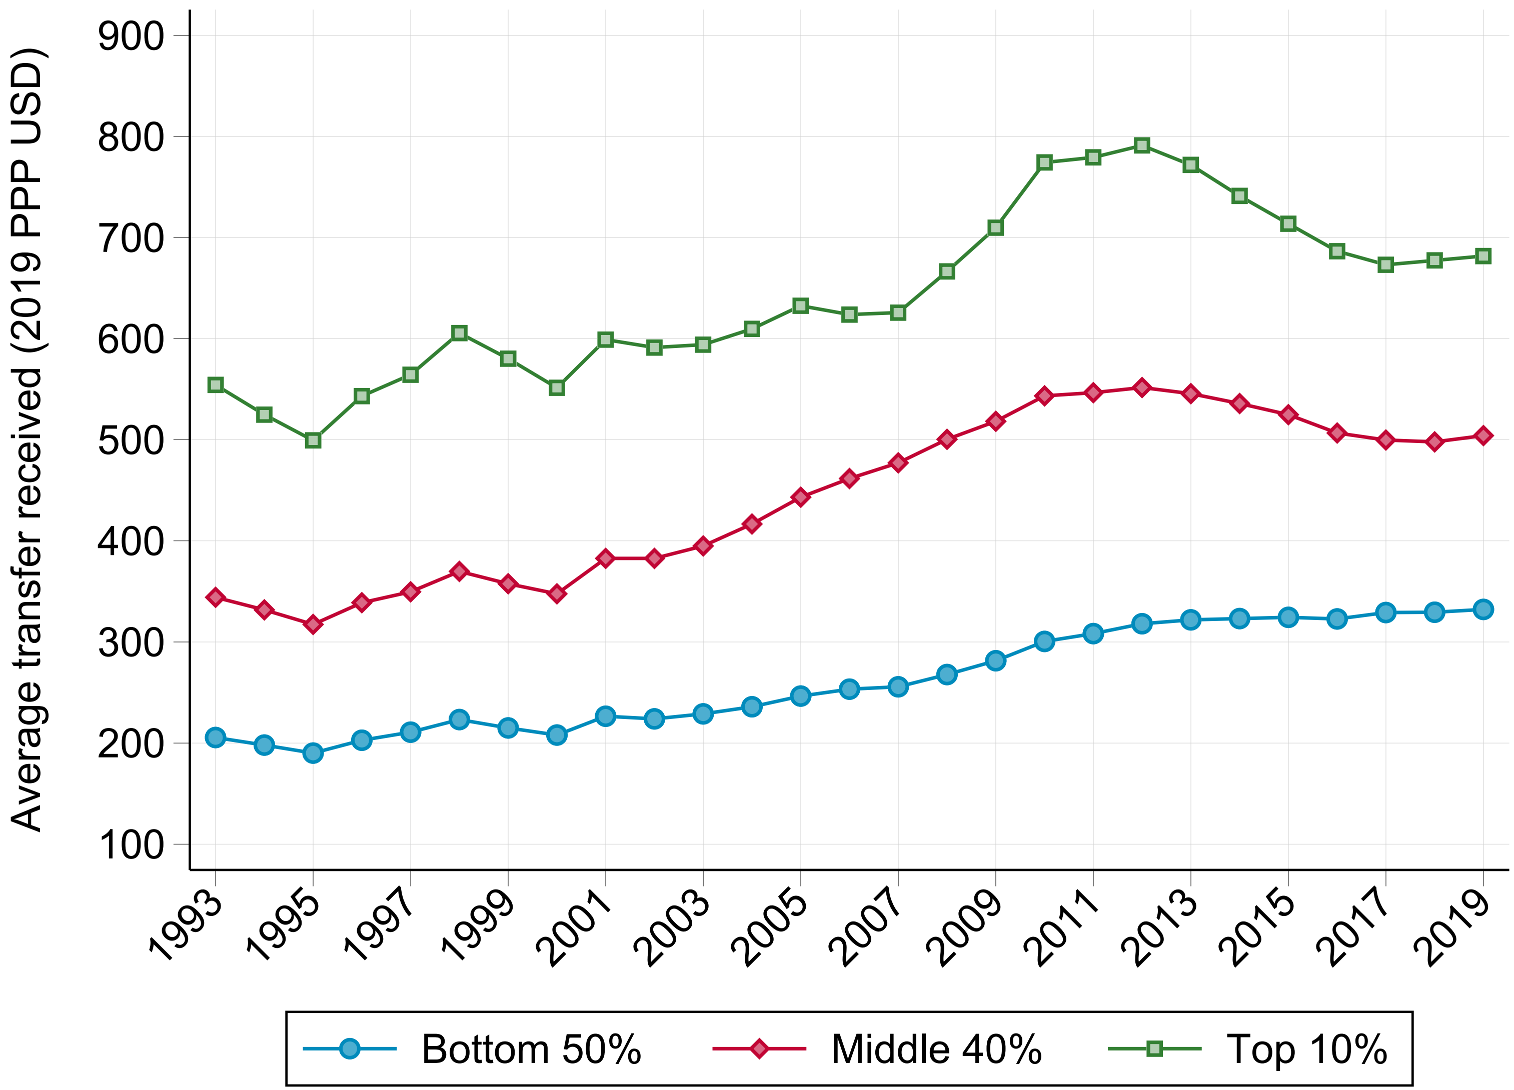

Notes. Author's computations combining surveys, tax, and national accounts data.

Figure E.42: Average Public Order and Safety Transfer Received by Income Decile, 1993-2019

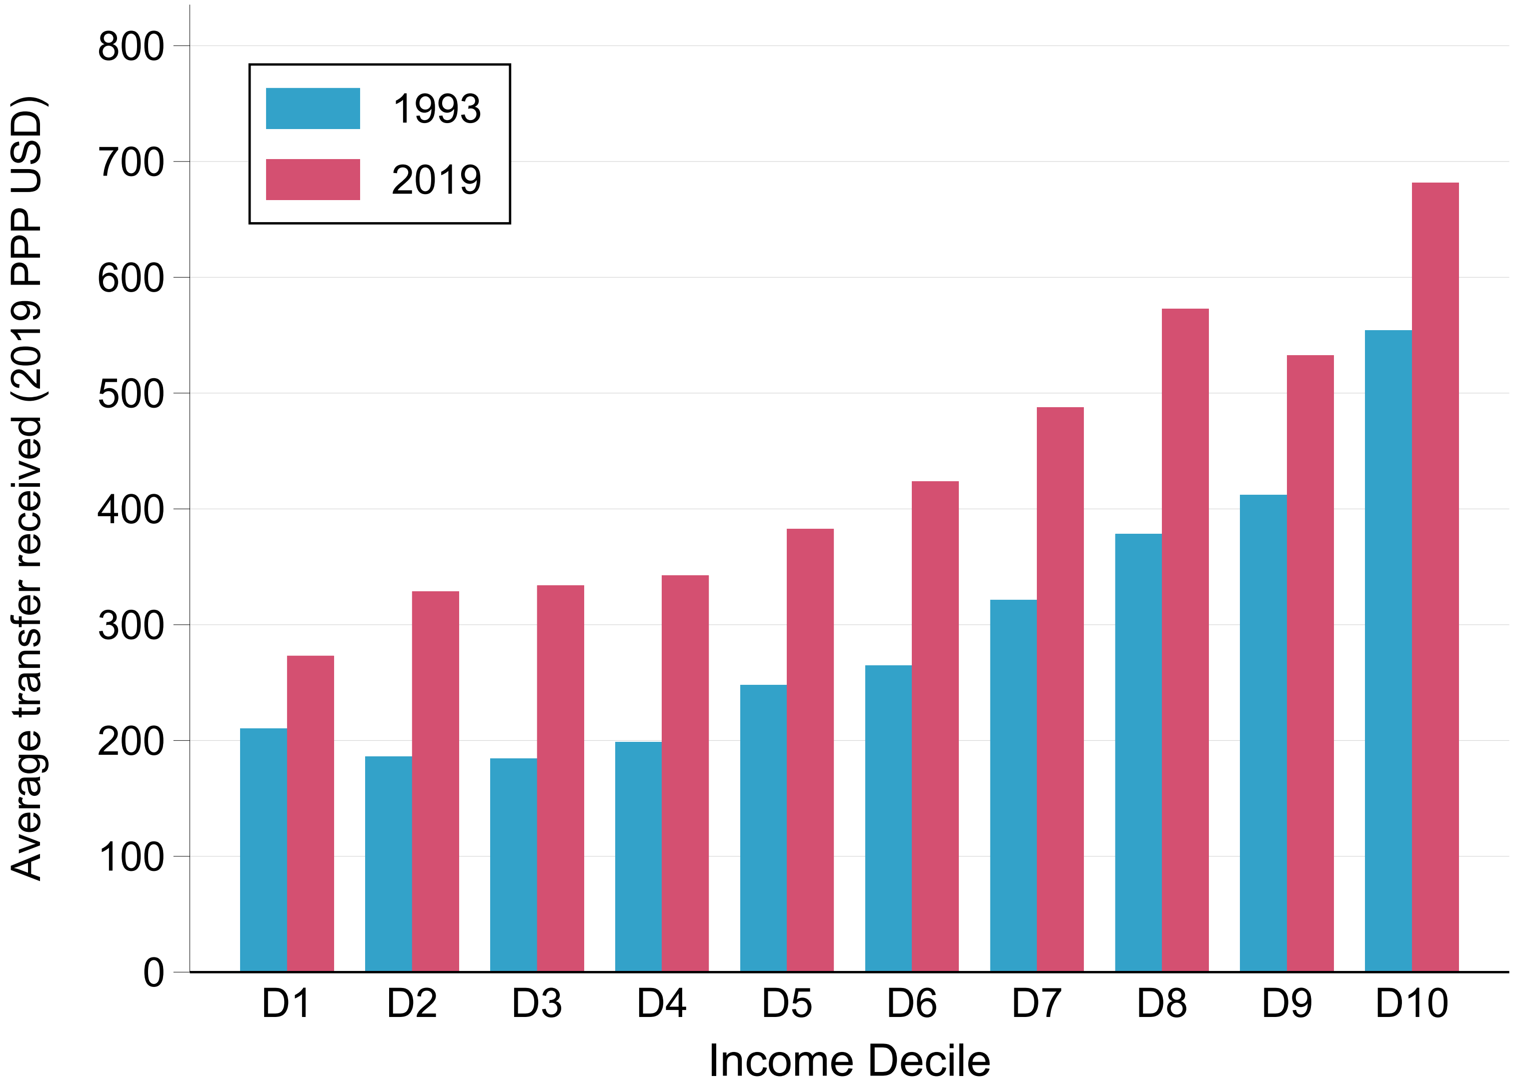

Notes. Author's computations combining surveys, tax, and national accounts data.

E.9 Transport and Other Economic Affairs

Figure E.43: Level and Composition of Expenditure on Economic Affairs in South Africa, 1993-2019

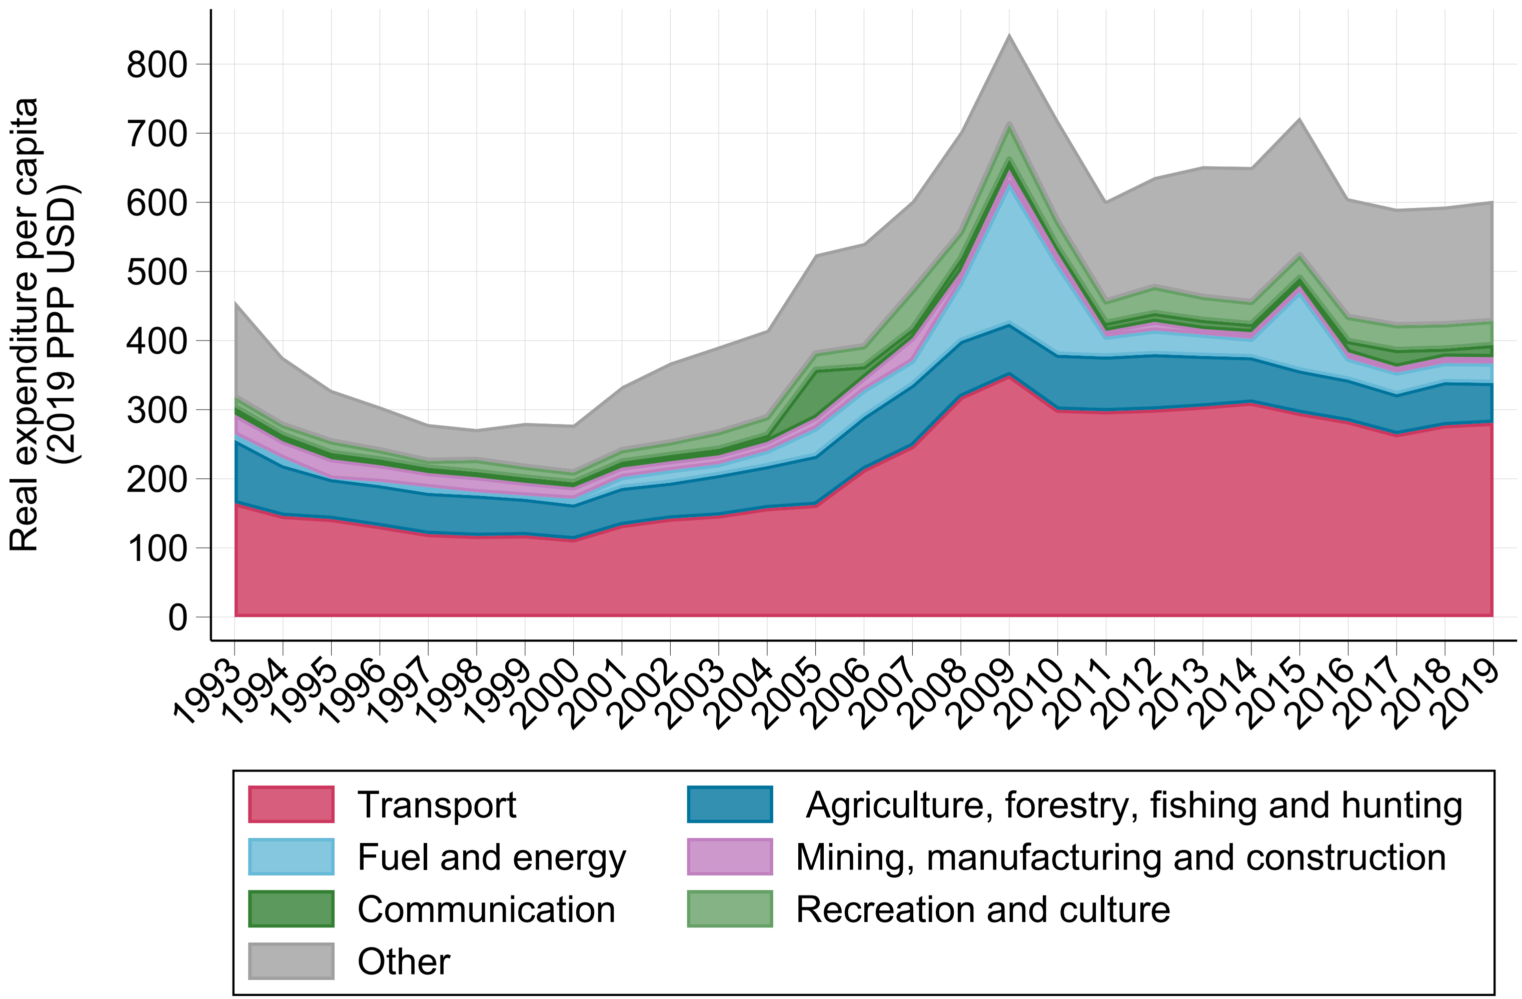

Notes. Author's computations combining data from South African National Treasury Budget Reports (1994-2020).

Figure E.44: Level and Composition of Transport Expenditure in South Africa, 1993-2019

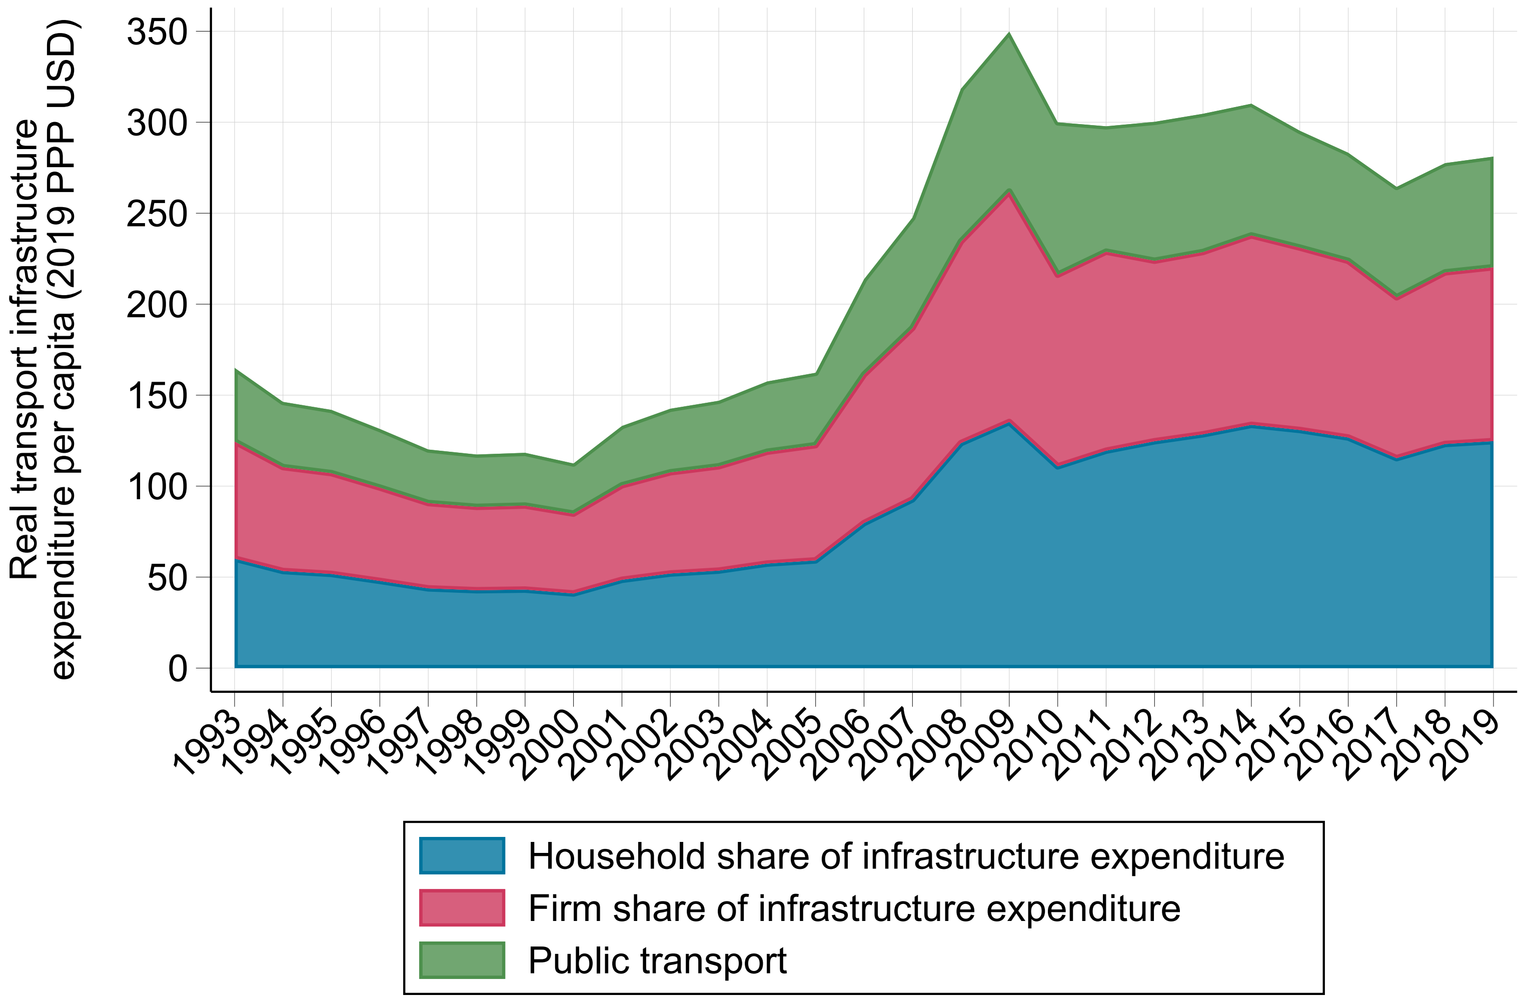

Notes. Author's computations combining data from South African National Treasury Budget Reports (1994-2020).

Figure E.45: Public Transport Use Intensity by Income Quintile: Buses

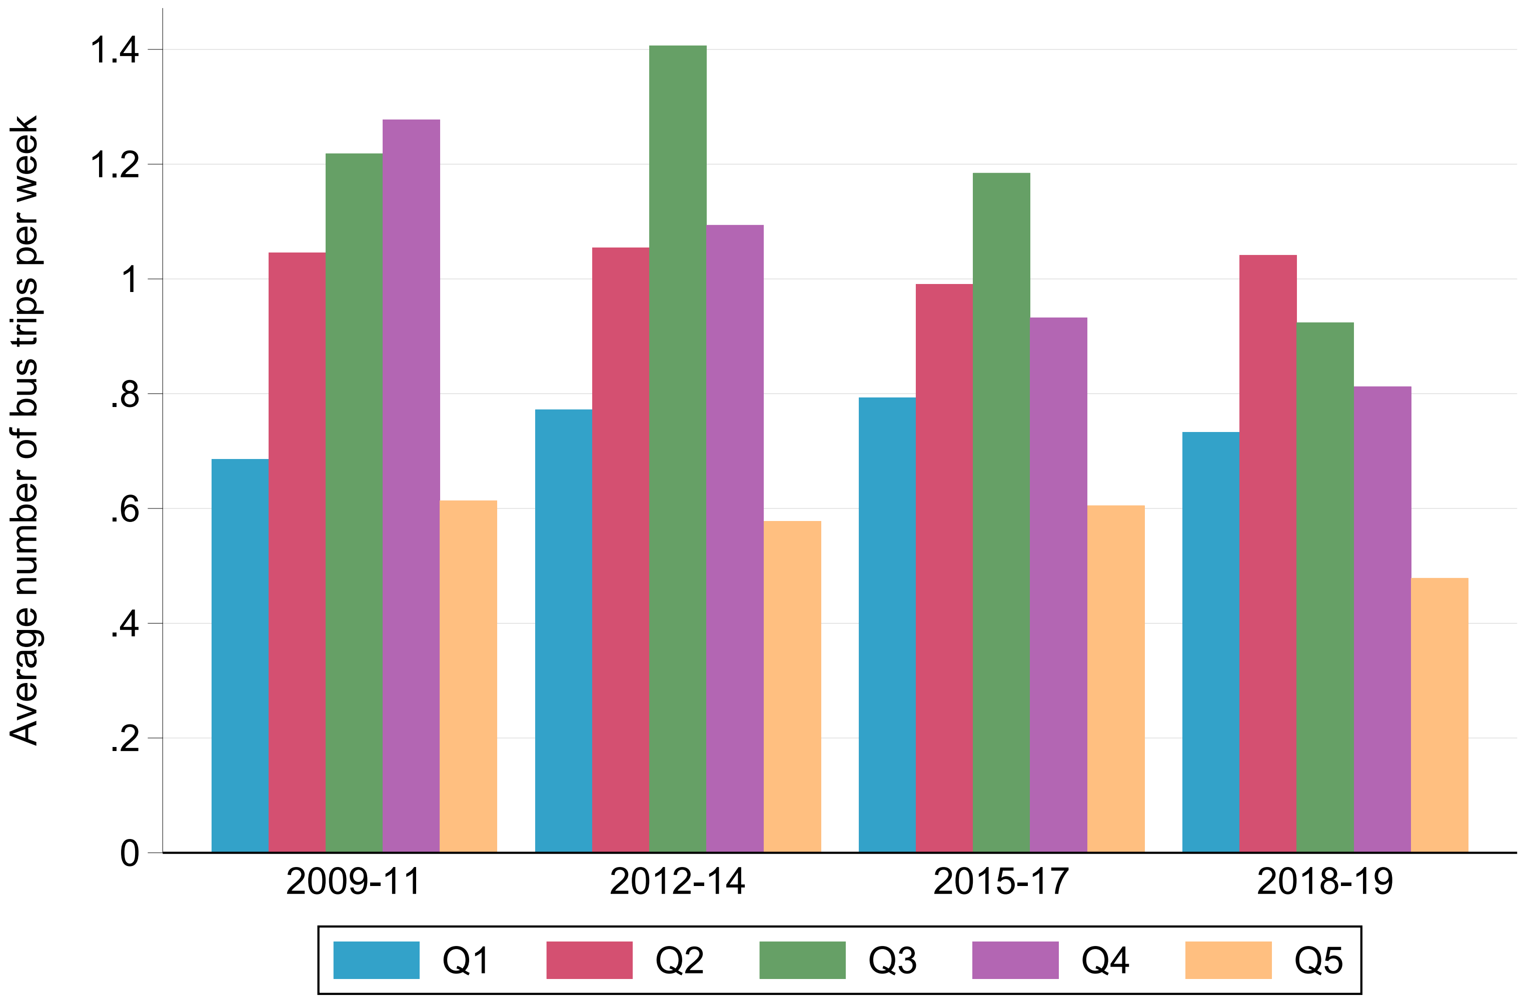

Notes. Author's computations combining General Household Surveys.

Figure E.46: Public Transport Use Intensity by Income Quintile: Trains

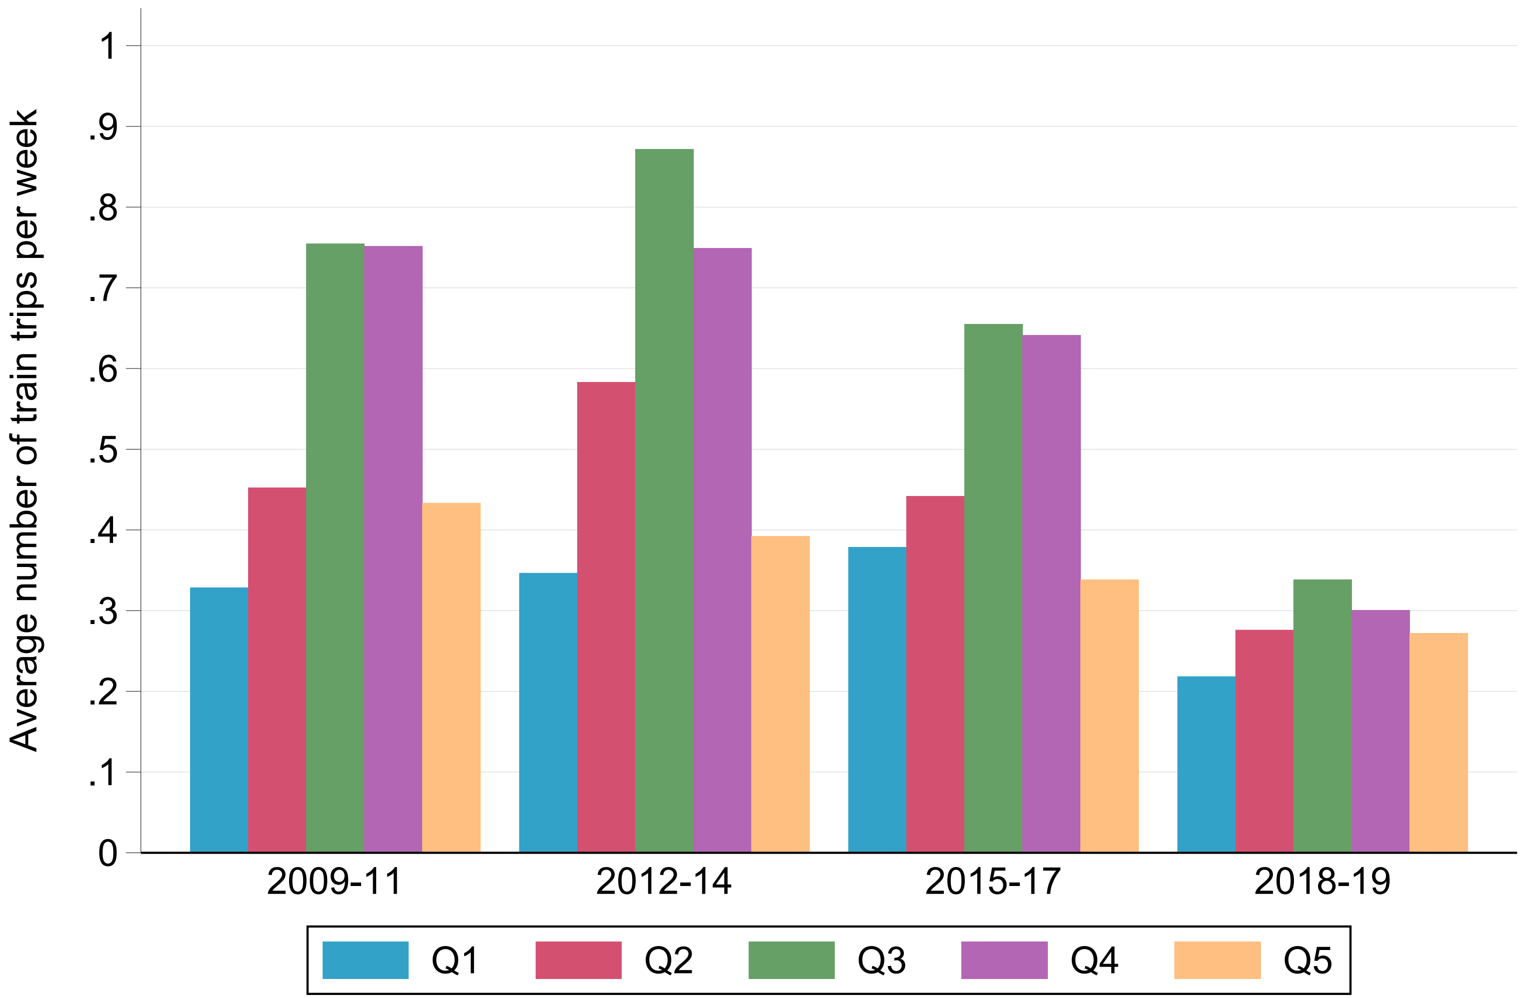

Notes. Author's computations combining General Household Surveys.

Figure E.47: Average Transport Transfer Received by Income Group, 1993-2019

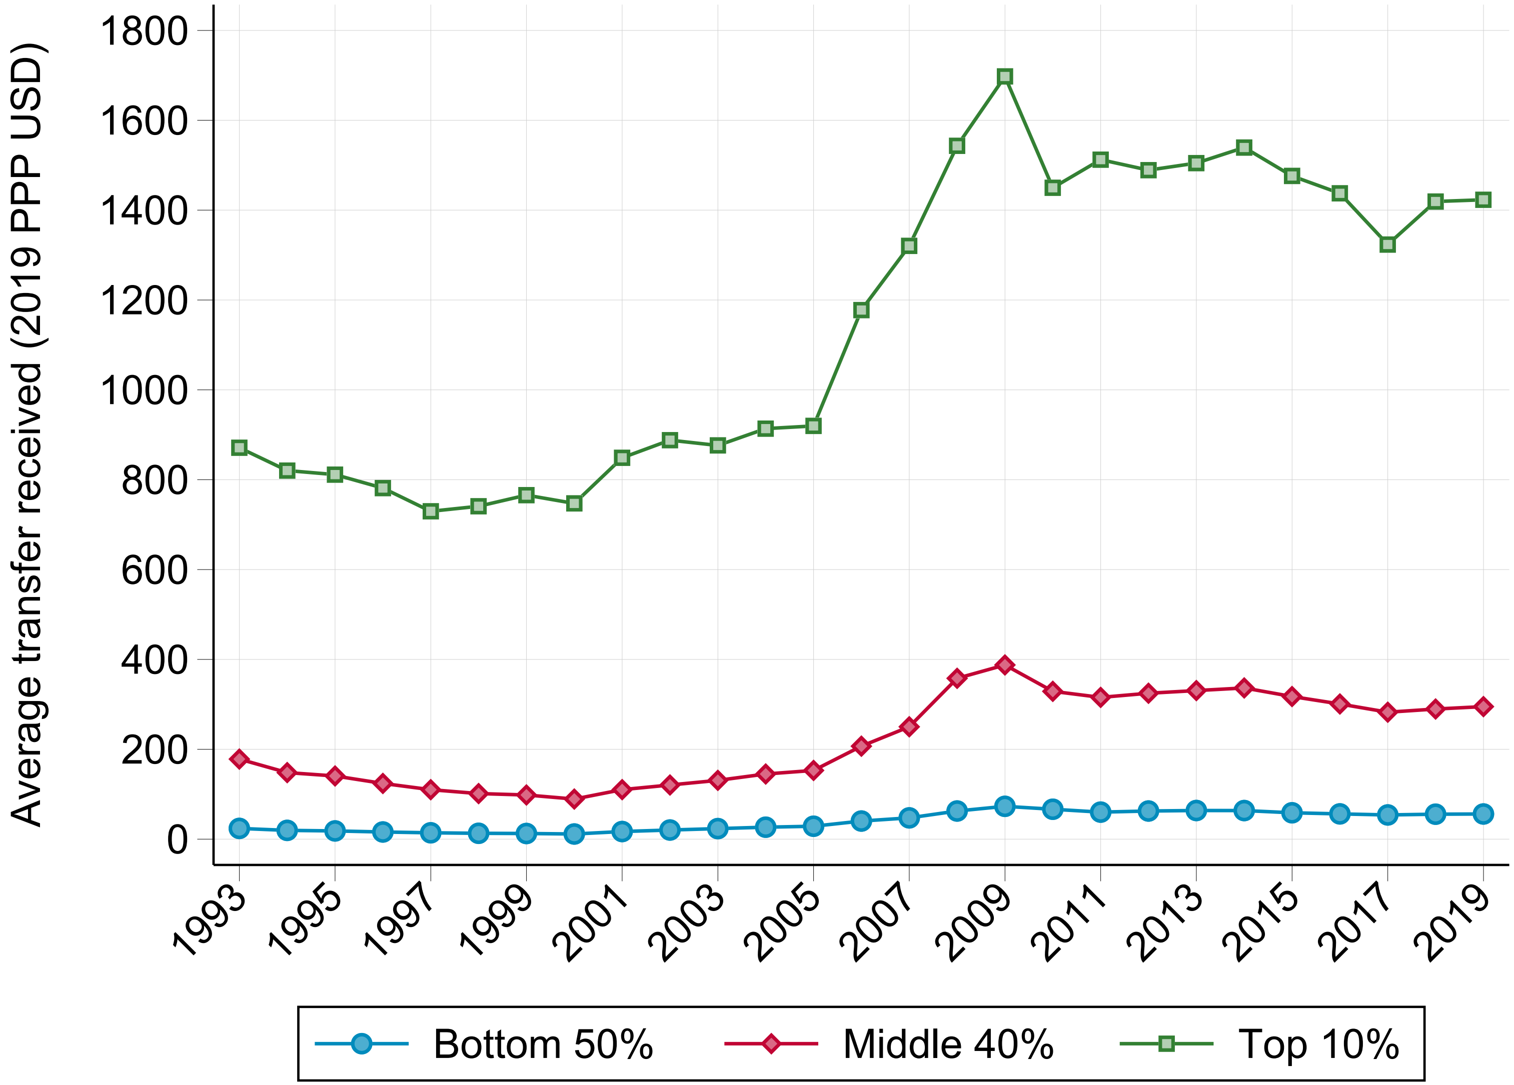

Notes. Author's computations combining surveys, tax, and national accounts data.

Figure E.48: Average Transport Transfer Received by Income Decile, 1993-2019

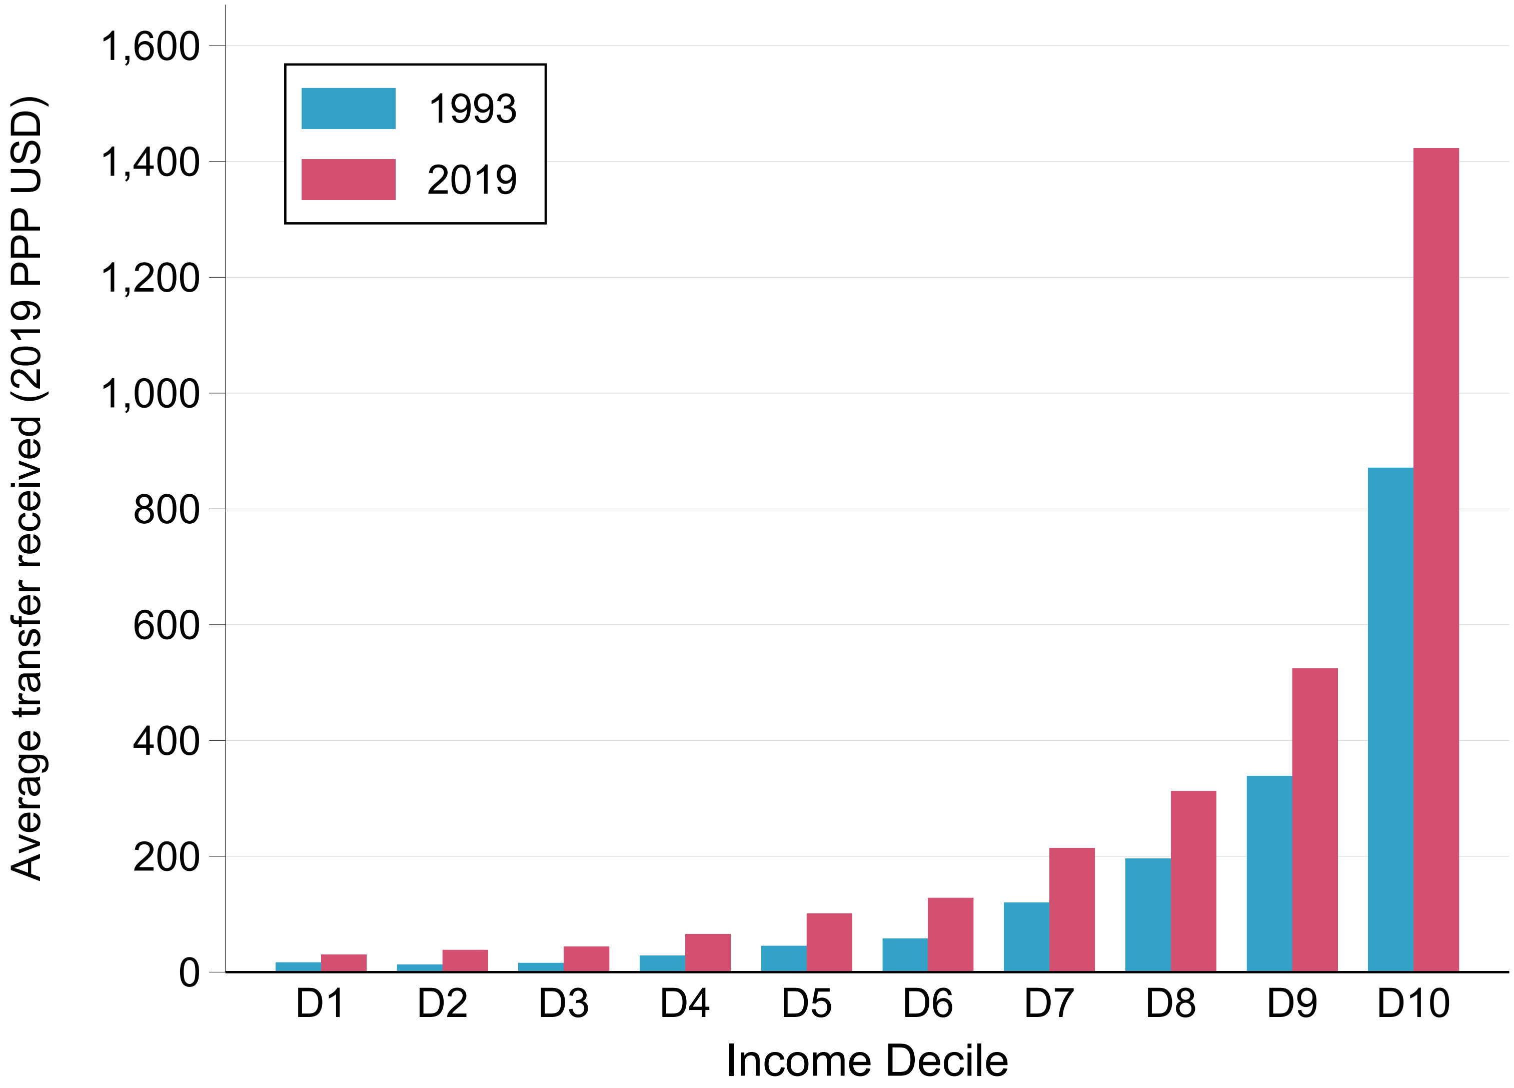

Notes. Author's computations combining surveys, tax, and national accounts data.

Figure E.49: Average Transfer on Economic Affairs Received by Income Decile, 1993-2019

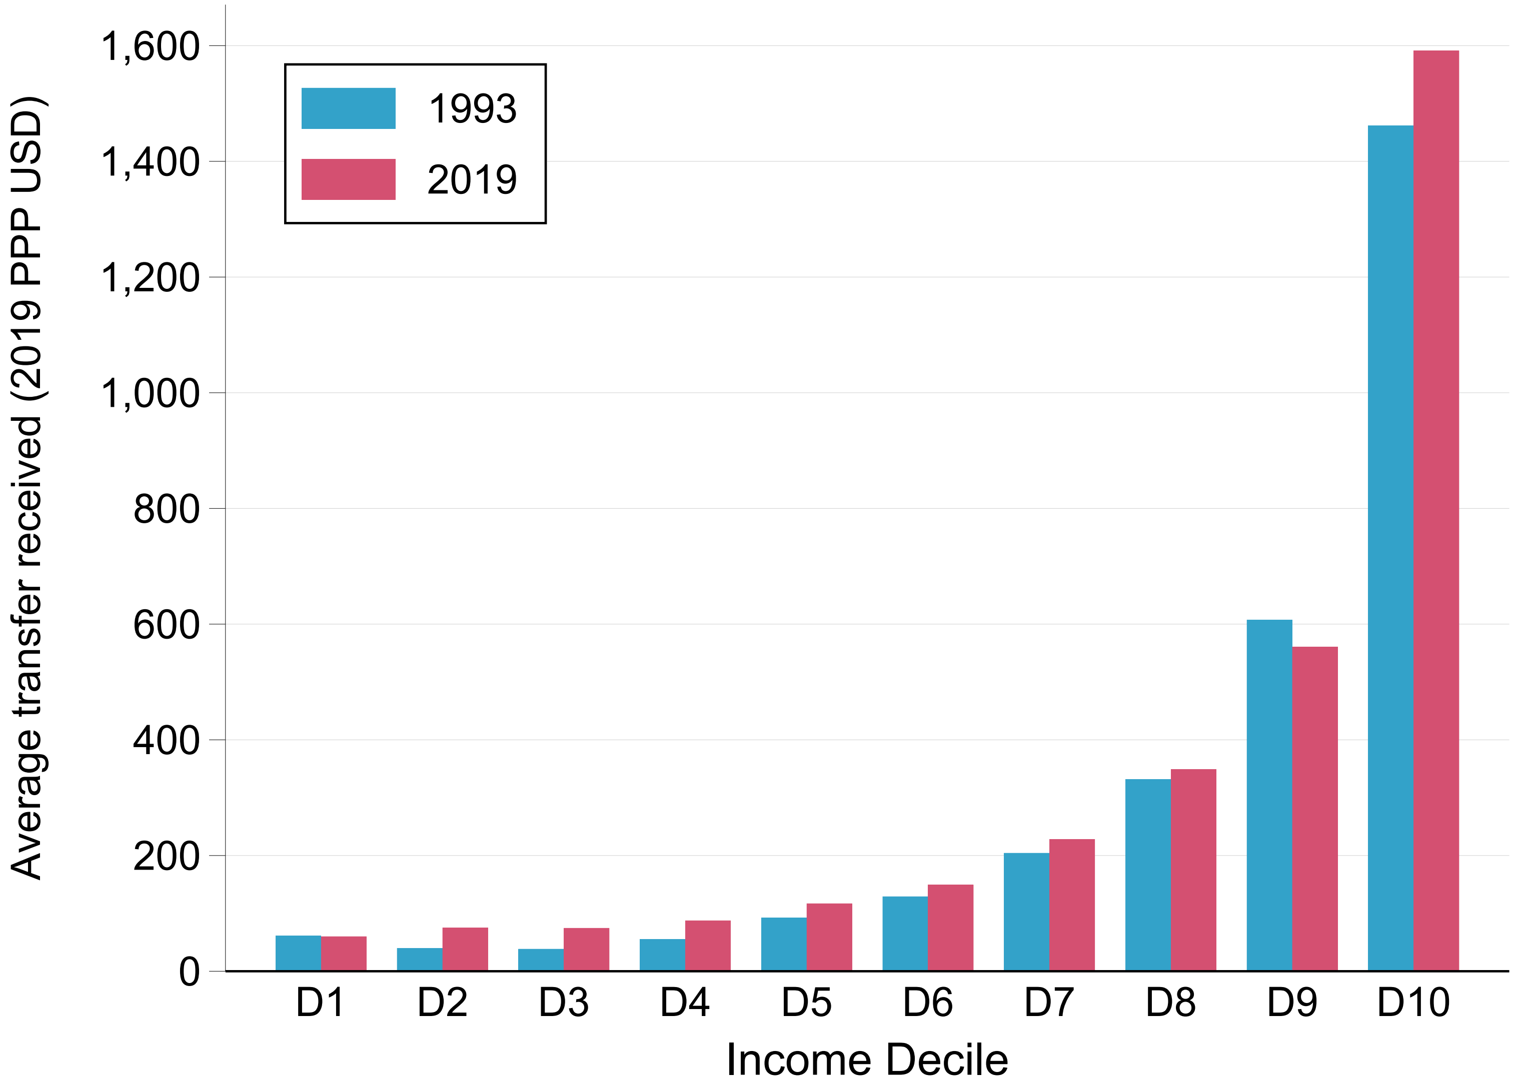

Notes. Author's computations combining surveys, tax, and national accounts data.

Appendix F

Appendix to "Redistribution without Inclusion? Inequality in South Africa Since the End of Apartheid"

F.1 Construction of Distributional National Accounts Microfile

This section provides additional details on the methodology used to build South African Distributional National Accounts. Section F.1.1 lists the data sources used to estimate macroeconomic aggregates, including national accounts, population estimates, and other government budget and administrative data. Section F.1.2 describes the combination of available survey and tax data to build a microfile covering the distributions of factor national income every year from 1993 to 2019. Section F.1.3 explains how taxes and transfers are allocated to reach posttax national income.

F.1.1 Harmonization of Macroeconomic Aggregates

F.1.1.1 National Accounts Data

Main Aggregates Estimates of national income, wealth, and expenditure aggregates come from the South African Reserve Bank (SARB) quarterly bulletin.[1] The published files provide detailed decompositions of national accounts components, which we directly match with the microfile to estimate distributional national accounts. The exceptions are mixed income and corporate undistributed profits, which we decompose further to refine the imputation.

Decomposition of Mixed Income and Imputed Rents The SARB data does not publish separate series for mixed income, rental income, and imputed rents, instead providing a single aggregate for [B2N + B3N, S14]. To derive an estimate of total rental income received by households, we combine all income surveys (1993, 1995, 2000, 2005, 2008, 2010, 2015: see section F.1.2) and General Household Surveys (GHS, 2016-2019), which have collected information on rents paid by South African tenants.[2] The resulting total rental income represented 1.9% of national income (14% of [B2N + B3N, S14]) in 2019, up from 1.4% (12%) in 1993. Following recommendations by the South Africa Reserve Bank, we assume that imputed rents represent a fixed 20% of the total, and we compute mixed income (*i.e.*, self-employment income excluding rental income) as the residual of these two categories.

Decomposition of Corporate Undistributed Profits To allocate corporate retained earnings to individuals, one has to decompose them between the part that belongs to households (distributed proportionally to equity ownership) and the part that belongs to the government (distributed proportionally to factor income). We do so by relying on a preliminary estimate published by the SARB on the equity assets and liabilities of the household and government sectors in 2011 (see Beer and Kock, 2017). Dividing the sum of the equity assets held by the government by the total equity liabilities of the corporate sector, we estimate that about 93% of retained earnings can be attributed to households. In the absence of better data, we assume that this share has remained stable over the 1993-2019 period.

[1] See `https://www.resbank.co.za/en/home/publications/quarterly-bulletin1/download-information-from-xlsx-data-files`.

[2] I first aggregate all rent payments recorded in income surveys. We then interpolate the series linearly between years to cover the entire 1993-2015 period. Finally, we use GHS growth rates in rent payments to extrapolate series forward to 2019.

F.1.1.2 General Government Revenue and Expenditure Data

To move from factor income to pretax income and then posttax income, we collect data on general government revenue and expenditure from three main sources: the SARB, the OECD, and the South African National Treasury.

Government Revenue Yearly data on consolidated government revenue and its decomposition are available from the public finance series published in the SARB Quarterly Bulettin. We complement these harmonized series with OECD public revenue data to further decompose revenue from direct taxes into the personal income tax, the corporate income tax, and other taxes on income and wealth.[3]

Government Expenditure Data on the composition of general government expenditure by function are available from the Treasury Budget Reviews.[4]

Social Security Data To make the DINA microfile more representative of Unemployment Insurance Fund (UIF) and private pension contributions and benefits, we collect data on total contributions/benefits and number of contributors/recipients to the UIF and private pension funds in South Africa. Data on total UIF revenue and expenditure (2001-2019) and on the number of UIF recipients (2008-2012) are reported in various issues of the Treasury Budget Review. The number of individuals earning private pension income is estimated from the income tax panel microdata (2011-2017) available from the South African Revenue Service (see Ebrahim and Axelson, 2019), and extrapolated to 1993 assuming that it has remained a constant share of the adult population.[5] Total contributions to private pension funds and total private pension income are also estimated from the income tax panel, and extrapolated to 1993 using the growth rates of social contributions received by financial corporations and social benefits paid by financial corporations, respectively (both available from SARB national accounts data).

Social Protection Data We also collect data on the number of recipients and the monthly values of social grants from various issues of the Treasury Budget Reviews. Data on grant values are available every year since 1993 (or since the year the grant was implemented) for all major cash transfers in South Africa (including the old age grant, the disability grant, the child support grant, the foster care grant, and the

[3] See `https://stats.oecd.org/Index.aspx?DataSetCode=REVZAF`.

[4] See `http://www.treasury.gov.za/documents/national%20budget/default.aspx`.

[5] This is a reasonable assumption to the extent that the number of pension recipients has also remained stable in income surveys, although at a lower level than in the tax microdata.

care dependency grant). Data on the number of recipients of each grant are available since 1996.

F.1.2 Construction of DINA Microfile

F.1.2.1 Combination of Survey Data Sources

The main data source used to estimate the distributions of income, consumption, and wealth at the micro level are household surveys that have collected detailed information on the earnings and expenditure of households in South Africa. Seven such surveys, which we refer to as "income surveys" in what follows, have been conducted since 1993: the Project for Statistics on Living Standards and Development (1993), the Income and Expenditure Surveys (1995, 2000, 2005, 2010), and the Living Conditions Surveys (2008, 2015). Drawing on representative samples of households, they ask individuals to report earnings from various sources (such as wages, self-employment income, and property income), as well as other information such as contributions to private pension funds, taxes and transfers received, the market value of the home individuals live in, or expenditure on specific goods and services.

I create a harmonized microfile covering the entire 1993-2019 period by combining all available surveys (1993, 1995, 2000, 2005, 2008, 2010, and 2015) and filling missing years in the following way. For a given missing year (for instance 1997), we create a new dataset by appending all observations from the two surveys available in surrounding years (1995 and 2000), and then reweight observations so as to give a weight to each survey that is proportional to the distance from the year considered. To approximate the distribution of income in 1997, for instance, we append the 1995 and 2000 IES surveys, and then multiply existing sample weights by 1/2 in the former and 1/3 in the latter. This is similar to a linear interpolation strategy: it amounts to considering that in 1997 the distribution of income was somewhere between that of 1995 and that of 2000, and was closer to that of 1995. The resulting microfile thus combines all available surveys to cover individual-level data every year from 1993 to 2019.

F.1.2.2 Combination of Surveys with Tax Data

I correct surveys for misreporting of income at the top of the distribution by combining them with tabulated income tax returns. This correction is performed in three steps, following the methodology developed by Blanchet, Flores, and Morgan (2022).

First, we define an income concept, "merging income", that can be consistently measured in both survey data and the income tax panel microdata (2011-2017). This income concept is equal to the sum of gross wages, business income, interest, rental income, and private pension income.

Second, we generate a "taxable income" variable in the survey microfile by multiplying merging income by the ratio of taxable income to merging income by percentile observed in the tax microdata. This effectively amounts to incorporating deductions (that is, the gap between merging income and taxable income) in the survey microdata.[6]

Third, we calibrate the survey microfile on the tabulated income tax returns available from SARS, which report the number of taxpayers and total taxable income by income tax bracket every year since 2002 (as well as in 1993). We first recover full distributions from the tax tabulations using Generalized Pareto Interpolation (Blanchet, Fournier, and Piketty, 2021).[7] We then calibrate the survey microdata on the tax tabulations using the algorithm developed by Blanchet, Flores, and Morgan (2022), which reweights survey observations so as to match the distribution of top taxable incomes reported in the tax data. The resulting survey microfile is thus perfectly representative of the distribution of taxable income reported in the income tax tabulations.

F.1.2.3 Combination of Survey and Tax Data with Macroeconomic Aggregates

After combining surveys with tax data, we rescale reported household income components to macro totals, and distribute components of the net national income that are not directly received by individuals.

First, we proportionally scale up household income components to their corresponding totals reported in the national accounts:

- Gross wages proportionally to compensation of employees.
- Self-employment/business income proportionally to mixed income (excluding rental income, see section F.1.1)
- Rental income proportionally to total rents paid by households

[6] For simplicity, we take the overall average of this ratio by percentile observed in 2011-2017 and apply it to the entire period. This corresponds to assuming that the profile of deductions has remained relatively stable between 1993 and 2019.

[7] For missing years (1994-2001), we assume that the extent of the under-representation of top incomes in survey data has evolved linearly, that is, we create synthetic income tax tabulations by linearly interpolating the correction by percentile observed in 1993 and 2002.

- Interest income proportionally to total interest received by households
- Dividends proportionally to total dividends received by households

Second, we distribute unreported income components proportionally to proxy variables available in surveys:

- Imputed rents proportionally to the reported market value of the home of owner-occupiers
- Property income attributed to insurance holders and pension entitlements proportionally to the value of pension and life insurance assets
- Interest paid by households proportionally to the factor income of debtors
- Private corporate undistributed profits proportionally to directly and indirectly held stock ownership
- Government primary income and other remaining national income components proportionally to factor income

F.1.3 Distribution of Taxes and Transfers

F.1.3.1 Pension and Unemployment Systems

Pension and unemployment contributions and benefits are recorded in income surveys, so we distribute macro aggregates proportionally to values reported by respondents. In order to reach pretax national income, we distribute 50% of the deficit or surplus of each system proportionally to contributions paid, and 50% proportionally to benefits received.

F.1.3.2 Taxes

Personal Income Tax We microsimulate the personal income tax every year from 1993 to 2019. To do so, we first collect data on taxable income thresholds, marginal tax rates, and rebates at each income level from various reports published by the South African Revenue Service. We then apply the corresponding rules in the microdata to calculate the tax burden of each individual. Because we have calibrated top taxable incomes directly on income tax tabulations (see section F.1.2.2), the estimates of total personal income tax revenue derived from microsimulation match almost perfectly actual revenue statistics. We close the residual gap between micro and macro estimates by proportionally rescaling the income tax burden of each individual.

Dividends Tax we distribute the dividends tax proportionally to dividends reported in income surveys.

Corporate Income Tax we distribute the corporate income tax proportionally to equity ownership, including both directly held equity and equity held indirectly through pension funds (see Chatterjee, Czajka, and Gethin, 2022).

Skills Development Levy The Skills Development Levy (SDL) is a 1% additional levy paid by wage earners who already contribute to the Unemployment Insurance Fund. We simulate it following this rule, and proportionally rescale the total to match total SDL revenue throughout the period.

Other Direct Taxes Other direct taxes include a number of minor taxes and levies, which have represented less than 1% of national income from 1993 to 2019. We distribute them proportionally to pretax income.

Transfer Duties The Transfer Duty is a tax levied on the value of properties acquired by individuals in South Africa. In the absence of information on property transactions, we distribute it proportionally to housing wealth (including both owner-occupied and tenant-occupied housing: see Chatterjee, Czajka, and Gethin, 2022).

Securities Transfer Tax The Securities Transfer Tax is a small tax that applies to the purchase and transfers of listed and unlisted securities. We distribute it proportionally to equity ownership.

Estate Duty and Donations Tax The Estate Duty and the Donations Tax are taxes on inheritance. In the absence of data on these transactions, we distribute them proportionally to total household wealth.

Value Added Tax we distribute total VAT revenue proportionally to household consumption expenditure, excluding both VAT-exempt goods and goods purchased on the informal market. Following the tax legislation, we directly identify VAT-exempt goods in income surveys and exclude them from taxable consumption. To identify goods purchased on the informal market, we derive a profile of informal consumption by income rank using the 2010 Income and Expenditure Survey, which reports the type of store at which the household purchased different kinds of goods. We extrapolate this profile to all years, assuming it has remained constant over the period. Expenditure in the informal sector is very small in South Africa, so that accounting for informality only has a negligible impact on the estimated distributional incidence of indirect taxation.

General Fuel Levy The General Fuel Levy is an excise tax charged on petroleum products. We distribute it proportionally to total transport expenditure reported by households in income surveys.

Other Excise Taxes Other excise duties mainly consist in excises applied to alcohol and tobacco products. In the absence of data on the decomposition of these taxes category by type of product, we distribute total revenue from non-GFL excises proportionally to combined alcohol and tobacco expenditure, as reported in income surveys.

Other Taxes on Goods and Services Other taxes on goods and services include a number of other small taxes, which have represented less than 0.5% of national income from 1993 to 2019. We distribute them proportionally to overall consumption expenditure.

Taxes on International Trade Import duties are effectively paid by households consuming a greater proportion of goods imported from abroad. Accordingly, we distribute taxes on international trade proportionally to import-intensive household expenditure, which we estimate using input-output tables available from the OECD (2005-2015).

Other Taxes Other taxes consist in a number of other small taxes and levies such as stamp duties. They have represented less than 0.5% of national income since 1993. We distribute them proportionally to pretax income.

Other Government Revenue we distribute all other government revenue, including non-tax revenue, proportionally to pretax income, so as to match total consolidated general government revenue in South Africa throughout the 1993-2019 period.

F.1.3.3 Social Protection

Social protection expenditure in South Africa mainly consists in the old age grant, the disability grant, the child support grant, other small cash transfers, and other social protection expenditure.

Old Age Grant The old age grant is a means-tested benefit paid to South African citizens who are 60 years or older. Old age grant beneficiaries are directly reported in income surveys, but their number is slightly below that reported in administrative data sources, suggesting a tendency to under-report. To correct this bias and ensure that my microfile matches both the true number of beneficiaries and total expenditure on the grant as reported in government budgets, we impute additional beneficiaries in two steps. First, we estimate the probability of surveyed individuals to receive the grant using a saturated linear probability model with the following explanatory variables: pretax income percentile, household expenditure percentile, gender, age, race, province or residence, and rural-urban location. Second, we rank individuals according to the predicted probability to receive the grant, and recursively allocate additional grants to those individuals most likely to receive it, until reaching the true number of beneficiaries every year from 1993 to 2019.

Disability Grant The disability grant is a means-tested benefit given to South African citizens who have a physical or mental disability that makes them unfit to work for a period of longer than six months. As in the case of the old age grant, it is reported in income surveys. We follow the same two-step strategy to impute additional beneficiaries when necessary, so as to match administrative statistics on both number of beneficiaries and total grant expenditure.

Child Support Grant The child support grant is a means-tested benefit given to low-income South African families to assist parents with the costs of the basic needs of their children. As in the case of the old age and disability grants, it is reported in income surveys. We follow the same imputation strategy as for these two grants, so as to match administrative statistics on both number of beneficiaries and total grant expenditure. The child support grant was first implemented in 1998, so we set grant expenditure and beneficiaries to zero before that year.

Other Social Grants Other small cash grants in South Africa include the foster care grant, the care dependency grant, the grant-in-aid, and social relief. We distribute them proportionally to their values reported in income surveys.[8]

Other Social Protection Expenditure Other social protection expenditure mainly consists in "provincial social development" expenditure, which brings together

[8]Most income surveys do not report receipts from these grants separately, so we derive an aggregate for "other social grants" in each survey and distribute total expenditure on these grants proportionally to this aggregate.

a large number of heterogeneous subnational policies targeted to poor households. These include, for instance, projects dedicated to reducing HIV prevalence, supporting disabled persons, providing centers for the treatment and prevention of drug abuse, or developing services aimed to prevent violence against women and children. In the absence of precise information on who benefits from each of these policies, we distribute other social protection expenditure proportionally to total social grants received.

F.1.3.4 Other Government Transfers

See Gethin (2023c).

Appendix G

Appendix to "Wealth Inequality in South Africa, 1993-2017"

G.1 Harmonization of macrodata sources

The objective of our study is to estimate the distribution of household wealth by matching macrodata on wealth with microdata on reported assets and capital income flows. In order to improve our estimates of the wealth distribution and obtain a better mapping of macrodata and microdata components, we address five shortcomings of available household balance sheets published by the SARB: the decomposition of non-financial assets, the decomposition of housing wealth into tenant-occupied and owner-occupied, the decomposition of financial assets, the decomposition of pension and life insurance assets, and the inclusion of wealth held offshore in tax havens.

The SARB currently publishes decompositions of household wealth into its financial and non-financial components, along with broad decompositions by asset class and information on household debt (see figure 7.1). Non-financial assets are divided into two components: residential buildings (the market value of residential properties owned by household, excluding land) and other non-financial assets (including land and unincorporated business assets). Financial assets are divided into three components: interest in pension funds and long-term insurers, assets with monetary institutions, and other financial assets. Interest in pension funds and long-term insurers corresponds to all pension assets and life insurance holdings of the household sector.[1] Assets with monetary institutions include all forms of currency and deposits

[1]This corresponds to the sum of the total assets of official pension and provident funds (series KBP2215 in Capital Markets Statistics), the total liabilities of private self-administered pension and provident funds (KBP2339), and the liabilities of long-term insurers under unmatured policies

with banks, as well as notes and coins held by households. Other financial assets include investment in government and public entities stock, collective investment schemes, corporate bonds and equities, other long-term deposits and households' investment in foreign assets. Finally, the SARB decomposes household debt into two components: mortgage advances, corresponding to loans provided by the commercial banking sector, and other debt (including trade credit, personal bank loans, credit card debt, instalment sales and lease agreements, and other formal and informal loans).

Starting from these broad categories, we derive further decompositions of macroeconomic household balance sheets to match specific types of assets with their corresponding income flows.

Land underlying dwellings The "Other non-financial assets" category provided by the SARB includes both land underlying dwellings and business assets. These two components are arguably distributed very differently. In particular, it is reasonable to assume that land underlying dwellings is distributed similarly to residential buildings (therefore defining total housing assets as the sum of land and residential buildings), while the distribution of unincorporated business assets is better approximated by that of mixed income. Given our income capitalization methodology, we therefore need to split "Other non-financial assets" into the two sub-aggregates. Based on complementary evidence from SARB, we assume that 70% of other non-financial assets correspond to land underlying dwellings, the remaining 30% amounting to the assets held by unincorporated businesses. This implies that total housing wealth (including land) was equal to 38% of net wealth in 2018, while business assets (machinery and equipment, excluding land) amounted to about 5% of net wealth.

Tenant- versus owner-occupied housing Housing wealth can be decomposed into tenant-occupied housing and owner-occupied housing. Available studies combining surveys with tax microdata typically assume that the distribution of tenant-occupied housing can be well approximated by the distribution of rental income, while owner-occupied housing assets are better captured using direct measurement available from surveys or administrative data (Garbinti, Goupille-Lebret, and Piketty, 2017; Saez and Zucman, 2016). Unfortunately, the "Residential buildings" category published by the SARB does not provide this decomposition, so we choose to derive

from the pension business (KBP2215). Notice that the original estimates of the South African household balance sheets done by Muellbauer and Aron (1999) excluded life insurance assets and all other assets associated with the non-pension business of long-term insurers. However, these items are now included by the SARB in line with the SNA guidelines.

the proportions from survey data (General Household Survey). To the best of our knowledge, the only available surveys collecting information on housing values for both tenants and owner-occupiers are the IES and LCS (1995, 2005, 2008, 2010) as well as the GHS since 2008. These surveys suggest that the share of tenant-occupied housing assets in total housing assets amounts to about 20% in recent years, down from some 25% in 1995. Notice however that we are considering all housing assets, including those owned by the government, corporations and other institutions in the denominator, as well as houses which are rented for free. In order to reach an aggregate closer to households' housing assets, we exclude tenants living in their dwelling without paying rents, as well as those declaring that they are renting from entities other than individuals. This leaves us with a clear distinction between tenants paying income to individual landlords, and formal owners of their houses, which is the concept we are interested in. This decomposition only exists in the GHS from 2013 onwards. The results show a decrease in owner-occupied housing wealth from above 75% in 2008 to 71% in 2013. We extrapolate this share to earlier years and apply it to the total reported in the households balance sheets.

Non-pension financial wealth The "assets with monetary institutions" and "other financial assets" categories published by the SARB gather together very different forms of financial assets, with arguably very heterogeneous distributions at the micro level, and thus must be split as well. "Assets with monetary institutions" include both non-interest bearing deposits such as cheque accounts, which do not generate any income flow, and interest bearing deposits, which generate interest income. "Other financial assets" include both bonds and corporate shares, which generate interest and dividends respectively. We follow Orthofer (2015) and assume that the composition of other financial assets held by households is similar to that reported by unit trusts as per SARB capital markets statistics. This implies that between 80% and 95% of other financial assets consist in corporate shares over the 1975-2018 period, the remaining being classified as bonds.[2] Finally, we separate currency, notes and coins (0.8% of net wealth) from interest-bearing deposits (17% of net wealth) using SARB capital markets statistics.[3]

[2] More precisely, we estimate the share of corporate shares in other financial assets by comparing the market value of ordinary shares held by unit trusts (KBP 2412) to the sum of the market values of security holdings of public sector entities, stocks and debentures held by unit trusts (KBP 2410 + KBP 2411) in the capital market statistics published by the SARB.

[3] The variable "Monetary sector liabilities: banknotes and coins in circulation" (series KBP1312) corresponds to currency, notes and coins held by all institutions. We assume that 70% of the total can be attributed to households. Given the small share of this component in total wealth, especially at the top of the wealth distribution, our results are not affected by alternative scenarios.

Pension assets and life insurance Pension assets correspond to the assets accumulated by wage earners through contributions to pension funds throughout their career, so they should in large part be distributed to wage earners and pensioners receiving pension income or annuities. Life insurance assets, by contrast, better correspond to a form of savings device, but they do not directly generate interest income, so they cannot be categorised with interest deposits or bonds and have to be distributed differently. Accordingly, we use available SARB capital markets data to decompose the “Interest in pension funds and long-term insurers” item into these two components.[4] In 2018, pension and life insurance assets amounted to about 28% and 13% of net wealth respectively.

Offshore wealth Offshore wealth corresponds to the assets held abroad by South African residents, mainly for tax avoidance purposes. By definition, these assets are not recorded in official records and are therefore not included in the household balance sheets. Alstadsæter, Johannesen, and Zucman (2018) combine a number of macroeconomic data sources to measure the total amount of financial assets held in offshore tax havens and distribute it to specific countries. They estimate that the equivalent of about 11.8% of South African GDP was held offshore in 2007, corresponding to about 5% of net wealth. We add this quantity to total household wealth in 2007 and extrapolate it to other years by assuming that it has remained a constant fraction of GDP.[5]

G.2 Harmonization of microdata sources

G.2.1 Harmonisation of household survey data, 1993-2018

Broadly speaking, two main types of nationally representative surveys covering the distribution of income and wealth have been conducted in South Africa since 1993: surveys covering all main types of income sources (such as wages, mixed income, rental income, interest, dividends or pension income) and labour force surveys covering only wages and mixed income. The first type of survey includes the 1993 Project for Statistics on Living Standards and Development (PSLSD); the Income Expenditure

[4]The share of interest in pension funds and long-term insurers corresponding to assets held by long-term insurers is recorded in the Capital Markets Statistics published by the SARB under series KBP2215, ”liabilities of long-term insurers under unmatured policies from the pension business”.

[5]Given that offshore wealth is known to have grown globally, this is a relatively conservative assumption for the period after 2007. If anything, wealth inequality could have increased more since 1993 than what our series suggest, as offshore wealth is well-known for been concentrated at the very top end of the distribution (Alstadsæter, Johannesen, and Zucman, 2019).

Surveys (IES) conducted in 1995, 2000, 2005, 2010; the Living Conditions Surveys (LCS) conducted in 2008 and 2015; and the National Income Dynamics Study (NIDS) conducted five times between 2008 and 2017. Labour force surveys include the October Household Surveys (OHS) conducted once a year between 1994 and 1999; the Labour Force Surveys (LFS) conducted twice a year between 2000 and 2007; and the Quarterly Labour Force Surveys (QLFS) conducted every three months since 2008.

In order to get yearly estimates of the wealth distribution between 1993 and 2018, we build a harmonised survey microfile by combining all these surveys in two steps. In a first step, we create a microfile covering the entire 1993-2017 period by combining income surveys (available in 1993, 1995, 2000, 2005, 2008, 2010, and 2015) in the following way: for a given year (for instance 1997), we create a new data set containing all observations from the two surveys available in surrounding years (1995 and 2000), and reweigh the data to give a weight to each survey that is proportional to the distance from the year considered. For 1997, for instance, we combine the 1995 IES survey and the 2000 IES survey, and we multiply existing sample weights by 1/2 for the former and 1/3 for the latter. This is similar to a linear interpolation strategy: it corresponds to considering that in 1997 the distribution of income was somewhere between that of 1995 and that of 2000, and was closer to that of 1995 if inequality evolved linearly. Given issues of comparability in income variables and sampling methods, we rely solely on the PSLSD, the IES and the LCS and we do not incorporate the NIDS into our harmonised file.

In a second step, we take advantage of the fact that while income surveys do provide information on the distribution of wages and mixed income, labour force surveys are more reliable for that very purpose and are available on a yearly basis. We therefore rank observations in the income surveys according to wages and mixed income and force the distribution of these two variables in our surveys (including interpolated years) to match that observed in the LFS or QLFS during the corresponding years by rescaling average incomes by rank. Due to difficulties in creating consistent inequality series from the OHS series, especially regarding mixed income, we choose to not exploit this data source and keep the PSLSD 1993 and the IES 1995 as our only survey data sources for the 1990s.

Finally, we extract yearly data on the distribution of the South African population by age, gender, province and population groups from the PALMS dataset and use simple linear calibration to calibrate the survey weights on the distribution of these sociodemographic variables. This ensures that the entire dataset is representative of the South Africa population in terms of these variables throughout the 1993-2017 period.

G.2.2 Comparing survey wealth aggregates to macroeconomic balance sheet totals

In this section, we briefly compare estimates of total wealth derived from existing surveys to macroeconomic balance sheets totals. The main finding that arises from this comparison is the existence of large differences between the two sources, due in particular to strong underreporting of financial assets in surveys. This motivates our mixed method of mapping micro wealth components with macro sources and capitalizing relevant income flows.

The only available surveys to directly measure wealth inequality in South Africa are waves 4 and 5 of the National Income Dynamics Study (NIDS). The comparison of household assets and liabilities reported in the NIDS surveys to macroeconomic statistics shows important inconsistencies (see table G.2). The market value of owner-occupied housing wealth is between 50% and 120% higher in the NIDS surveys than in the balance sheets, while tenant-occupied housing is closer to the macro aggregate. This most likely reflects the different methods in measuring market values.[6] Business assets are covered very differently in the two waves: they are overestimated in wave 4 and underestimated in wave 5. Pension and life insurance assets, after correction[7], seem to be relatively close to balance sheets figures, and they even slightly overestimate them. Other financial assets are extremely badly covered: the total reported in the NIDS surveys does not exceed 4% of households' bonds and stock reported in the balance sheets by the SARB. Considering that the underlying sources of SARB's series consist of financial statements submitted by all financial intermediaries[8] and several capital markets data, we interpret these discrepancies as

[6]It is beyond the scope of this paper to discuss and evaluate these methods. However, this issue is not one specific to South Africa - in the US, survey values have also been found to be higher than in balance sheets figures, and which source of information provides the more accurate estimate of market values is contested (Blanchet, 2016; Dettling et al., 2015; Henriques and Hsu, 2014). Another potential issue is how to treat RDP housing, a government-funded social housing project in South Africa, due to complexities around ownership. However, given the typical low market value of these properties, it is unlikely to affect our distributional estimates.

[7]There are important inconsistencies in data on pensions and other retirement funds in the NIDS survey. For example, in wave 5 of the survey, 61% of individuals declaring contributions to pensions funds declare having no "pension or retirement annuity", while 77% of individuals declaring income from a pension or provident fund declare no "pension or retirement annuity". We correct for these gaps by imputing all missing values using predictive mean matching.

[8]Monetary authority, banks, insurers, retirement funds, trusts and other types of finance companies. For more details about how the Flow of Funds data is compiled, see Beer, Nhlapo, Nhleko, et al. (2010)

a sign of the weakness of the NIDS surveys resulting from the difficulty to survey the wealthiest individuals. Household debts are slightly better covered, but still fall significantly below macroeconomic statistics.

The other surveys we use in this study (PSLSD, IES, and LCS) also contain some information on owner-occupied housing and debts. Owner-occupied housing seems to be over-stated relative to the balance sheets in these surveys as in the NIDS surveys (see table G.3). Debts are always below balance sheets totals, but with important fluctuations across surveys. All these limitations justify our approach to correct for discrepancies between micro and macro totals. Indeed, the households balance sheets have the advantage of tracking the evolution of aggregate wealth consistently, in contrast with surveys, which show much greater fluctuations in reported aggregates. By mapping the surveys with macroeconomic statistics, we are at least able to get estimates of the wealth distribution that are consistent with what we know of the level of aggregate wealth and its composition over time.

G.2.3 Comparing survey income aggregates to national accounts totals

As more surveys and available tax microdata deal with incomes, and generally income reporting is seen as more credible, capital related income provides alternate sources of information for the wealth distribution. In this section, we compare incomes from surveys to the corresponding totals recorded in the national accounts. For our purposes, the components we consider are gross wages (to capitalise pension wealth), mixed income (income from unincorporated enterprises, to capitalise unincorporated business assets), rental income (to capitalise tenant-occupied housing) and interest and dividends (for equity and bonds). The surveys we consider were designed to capture information about consumption, expenditure and earnings: these are the Project for Statistics on Living Standards and Development (PSLSD) conducted in 1993, the Income and Expenditure Surveys (IES) from 1995 to 2010, the Living Conditions Surveys (LCS) of 2008 and 2015, and the NIDS surveys.

As table G.4 shows, gross wages and mixed income are much better covered than capital incomes, and are better covered in the PSLSD, IES, and LCS than in the NIDS surveys. Rental income, interest and dividends are unfortunately poorly covered in all household surveys. This is due to this sort of income being concentrated by those at the upper end of the income distribution, who are typically underrepresented in surveys due to issues of sampling and non-response. This motivates our use of the tax microdata to better cover top incomes.

G.2.4 Extrapolation of tax data series back to 1993

Our wealth inequality series based on tax data cover the 2010-2017 period, while we can go back to 1993 by capitalising the income flows reported in household surveys. Series based on tax data typically show slightly higher levels of wealth concentration at the very top, so one meaningful way to extrapolate the tax data series back to 1993 is to assume that the underrepresentation of top wealth groups in surveys has remained constant before 2010.

We correct the survey series before 2010 by following the methodology developed by Blanchet, Chancel, and Gethin (2022) to correct a distribution based on observed relationships between quantile functions covering different concepts and data sources. Formally, consider for a given quantile $p \in [0; 1]$ the quantile function of the wealth survey series $Q_S(p)$ and the quantile function of the tax data series $Q_T(p)$. To impute the tax data series from the survey series, one can write:

$$Q_T(p) = Q_S(p) \times \beta(p)$$

Where $\beta(p) = Q_T(p)/Q_S(p)$. Therefore, it suffices in our case to estimate $\hat{\beta}(p)$ over the 2010-2017 period (where both survey and tax data series are available) and to then multiply $Q_S(p)$ by $\hat{\beta}(p)$ before 2011 to get a corrected survey series. This will be an efficient method, however, only in the case where both $Q_T(p)$ and $Q_S(p)$ are strictly positive, which is not true in our case since our wealth quantile functions include a significant share of zero and negative values. Blanchet, Chancel, and Gethin (2022) show that a good way of accounting for zeros and negative values is instead to work with the following transformation:

$$Q_T(p) = sinh\Big(asinh[Q_S(p)] + \beta'(p)\Big)$$

With $\beta'(p) = asinh(Q_T(p)) - asinh(Q_S(p))$, and where *sinh* is the hyperbolic sine and *asinh* is the inverse hyperbolic sine. We apply this method to get consistent series covering the 1993-2017 period.

G.3 Other issues

G.3.1 Negative net worth and the measurement of household wealth at the bottom end

Household debts are among the most difficult components of personal wealth to estimate, in part due to the difficulty for survey respondents to properly assess their remaining debt balances. The coverage of debt is very erratic in South African surveys, who cover from 14% to 87% of mortgage debt, and from 17% to 57% of other forms of debt. These difficulties are not specific to South Africa: in France, for instance, Garbinti, Goupille-Lebret, and Piketty (2017) choose to set negative net wealth values to zero, given the unavailability of proper information on the net worth of poorest households. Other recent comparable studies on India (Bharti, 2018), China (Piketty, Yang, and Zucman, 2019), Russia (Novokmet, Piketty, and Zucman, 2018) or the United States (Saez and Zucman, 2016) have generally found negative net worth to be restricted to the bottom 5% or 10% of the population, with the exception of the United States where households are highly leveraged.

In South Africa, in spite of the lack of high-quality data, there is considerable evidence that a substantial share of households have either zero or negative net worth. The National Income Dynamics Survey, for instance, asks specifically to adults: "Suppose you (and your household members living here) were to sell off everything that you have (including your home and vehicles), cash in your investments and pay all your debts, would you have money left over, break even or be in debt?" In 2017, 50% of households declared they would have something left over, 24% declared they would more or less break even, and 4% declared that they would still be in debt. The remaining 22% declared not knowing whether they would still have something left, which is a relatively clear indication of net wealth being very close to zero. Notice in particular that this question includes household durables, which are excluded from our SNA definitions of household wealth, so that the share of negative-net-worth households is clearly underestimated in this question. In any case, the evidence is suggestive of a substantial share of the population (at least between 30% and 50%) having either negative wealth, or wealth very close to zero.

Other evidence points to the concentration of debts among the bottom of the wealth distribution, and the lack of assets covering these debts. According to the 2019 Eighty 20 and XDS Credit Stress Report, the number of unsecured credit products – that is, debt which is not backed by any form of asset – far outweighed those holding secured accounts (Eighty 20 and XDS, 2019). In terms of values, unsecured debts amounted to 28% of total consumer credit products in South Africa in the third quarter of 2019. At the same period, the default rate was as high as 20% among consumers aged 18 to 24. These figures clearly indicate that a very large share of the South African population is highly leveraged, contracting consumer credits with no corresponding assets to back them – which means that they are by definition in negative net worth.

Our benchmark method for allocating debt to households is to rely on the share of households declaring debt and on a proxy variable of ability to pay rather than on direct measurement of debt balances. This avoids having too many households with unsustainable debt levels, while at the same time allowing us to fully close the micro-macro gap and distribute all debts recorded in households' balance sheets. For mortgages, we rely on the reported market value of the house, which is arguably a reasonable proxy for the average size of the mortgage balance across the wealth distribution. This method is comparable to that used by Saez and Zucman (2016), who distribute US mortgages proportionally to reported mortgage payments. For other debts, given the lack of other data, we rely on consumption, which is less unequally distributed than incomes and therefore evens out debts across the wealth distribution. After splitting wealth equally among adult members of the household, our estimates imply that 10% of the adult population has negative net worth; the entry thresholds for the next deciles are R 0, R 1700, R 10,000 and R 18,000. Median wealth amounts to R 30,000 (about 4800 dollars at purchasing power parity, or about a quarter of the average national income per adult). These low levels are consistent with the descriptive evidence above suggesting that some 30% to 50% of South Africans have close to zero wealth. In any case, as we show in figure G.14, top wealth shares are only moderately affected by the exclusion of debts from our framework: assets are extremely concentrated, with the top 10% owning 80% of the total.

That being said, it is important to note that durable goods are not included in the SNA definition of wealth, but that debts associated to the purchase of durable goods are. Given the importance of consumer credits and their use to buy cars or furniture among poorer households in South Africa, this may explain in large part why wealth is so negative at the bottom of the distribution. Whether durable goods should be included in wealth or not is a subject of debate. On the one hand, one might argue that the goods purchased with household debt should be included in households' net worth for consistency with individuals' experiences of what they own. On the other hand, debts are a form of stock generating an income flow, while consumer durables are not - they are consumed in a relatively short time, or depreciate at a comparatively high rate, and they do not generally generate any income flow -, so that one could argue that all consumer credits should be included in net worth, while consumer durables should not. Finally, let us also stress that survey data does not allow us to capture other forms of collective ownership – such as rights to land or cattle, which are particularly important in rural areas, both economically and symbolically – as surveys are restricted to wealth held at the household level. The inclusion of these components in household wealth can also be debated and should in any case be the subject of future research.

G.3.2 Limitations of the personal income tax data

G.3.2.1 General Comments

The fact that the ITR12 forms are self-assessed implies that there may be tax evasion or under-reporting of income flows, especially if the likelihood of being controlled by tax authorities is low. More importantly, tax microdata only covers forms of incomes which are useful for tax collection and deductions purposes, which implies that other forms of non-taxable incomes are not reported in the data. This, as we show below, is particular problematic for the measurement of capital incomes.

Table G.7 shows that when looking specifically at capital incomes in the tax data, the reported totals fall significantly below the national accounts. Interest income is better measured than rental income and dividends, reaching between 25 % and 30 % of total interest received by households in the national accounts. Rental income and dividends are significantly lower and inconsistent, covering between 2% and 25% of national accounts totals. [9]

This under-representation of capital incomes in the tax data is due to three main factors. First, the taxable incomes are different from incomes reported in the national accounts, due to filing rules and tax base. This is particularly problematic for dividends, which in the ITR12 relate to dividends from equities that form part of compensation packages, such as equity share plans. These sort of dividends are subject to income tax, and so part of this data set, whereas dividends from regular ownership of equity is subject to a separate dividend tax. Approximately 80 % of dividend information would be recorded through this dividend tax returns (DTR01/2 forms), and this information would be useful to make our estimate more reliable.

Secondly, there may be issues of misreporting of incomes by individual taxpayers. Interest income seems to be poorly covered as a result of incomplete tax filing by

[9]The particularly low figures obtained in 2017 (fiscal year 2018) are mainly due to the fact that assessment was incomplete at the time of writing.

taxpayers. In principle, the South African Reserve Bank receives direct information from banks and financial services that they provide about interest. Banks and financial service providers separately supply customers with a tax certificate (IT3(b) certificate), which is meant to inform the interest income declared by the taxpayer. At the same time, the bank sends the South African Revenue Service a third-party submission about incomes its customers' receive. However, given that interest income is typically low relative to total taxable income, it is possible that small interest income received go unreported. The misreporting of rental income received by individual taxpayers is likely to be more significant, given that rental income is self-reported and that there may be a significant amount of informal letting of fixed property. [10]

Despite of all this, tax microdata remains much better at capturing dividend and interest income than household surveys.

G.3.2.2 Trust income

The most important issue regarding the coverage of capital incomes in the tax microdata is likely to be due to the definition of the taxpayer. The tax data covers only individuals and does not account forms of capital incomes received through units trusts or investment funds. This is particularly problematic in the case of South Africa, both because wealth is highly concentrated at the top of the distribution and because top wealth groups rely extensively on unit trusts. As shown in figure G.18, the share of financial assets held through trusts exploded around the time of, politically, the end of apartheid, and economically, liberalisation and financialisation. Over half of specifically interest bearing and dividend earnings financial assets are held in trusts. Trusts in South Africa are used more extensively, including housing mutual funds, as well as tax avoidance vehicles, and one mechanism of several to protect against wealth dilation (wealth loss across generations) (Ytterberg and Weller, 2010). There is therefore a clear need to access data on trusts to gain more complete and precise information on the distribution of capital incomes (and their corresponding assets) at the top of the distribution, as well as to understand the mechanisms that results in the persistence of wealth concentration. However, the fact that we could not have access to sufficiently detailed data on trust does not imply that we did not distribute wealth held by households through trusts. Indeed, our

[10] Notice here that total rental income paid to individuals in the economy is estimated by the authors based on data from the PSLSD, the IES and the GHS surveys on total rental income paid by households to individual landlords. Therefore, this includes informal rents paid, which may explain why the rental income the tax data is so low compared to the macro aggregate.

methodology takes this share of wealth into account as it is by definition included into the macro aggregate we distribute over our microfiles. Access to better micro data on trust would only have allowed more precise allocation of wealth at the extreme top of the wealth distribution. In the following paragraphs we further document our exploration of the issue.

Just like individuals, all unit trusts in South Africa are required to file an ITR12T form covering all non-dividend sources of income, as well as a dividends tax form separately. The ITR12T form also contains information on taxpayer reference numbers and passport numbers of the beneficiary to whom income, capital or assets were distributed or vested with the highest monetary value. In parallel, individuals filing ITR12 returns are asked to provide detailed information on all forms of income distributed or vested to them as a beneficiary of a trust, as well as the trust name, the trust registration number and the trust tax reference number. In theory, this provides largely sufficient information to link trusts to their beneficiaries and accordingly distribute trust income and trust wealth. Unfortunately, the tax microdata provided by SARS does not include these entries, which were not extracted during the process of making the data accessible to researchers. In the ITR12 data, there is no trust information at all. SARS does provide researchers with the ITR12T data, but available variables are very limited, being restricted to the sources of income received by the different trusts, without any information on who owns them. This makes it impossible to distribute non-dividend trust income in any meaningful way, since individuals may have accounts in multiple trusts, and accounts may belong to multiple individuals. Furthermore, given that about 90% of trust assets correspond to corporate shares, the ITR12T data is only of very limited use at it excludes dividends from ownership of regular shares.

Table G.6 shows descriptive statistics computed from the ITR12T data. The number of tax returns has decreased from about 140,000 to 94,000 between 2014 and 2018, probably due to incomplete assessments at the time of writing. This implies that there was one trust for about 2400 adults in South Africa in 2014, which shows how the use of trusts is widespread in the country. When it comes to sources of incomes assessed however, the quantities observed appear to be extremely low compared to macro figures, in particular knowing that trusts hold a substantial share of financial wealth. Interest income received by trusts amounts to only 3% of total interest received by households in the national accounts. The corresponding figures are 2% of rental income and less than 2% of business income. Less than 0.5% of dividends are covered, which is consistent with the fact that only very specific types of dividends are covered in this data, the bulk of them being filed separately through the dividends tax form. Capital gains are among the biggest components of trust income, amounting to between 1% and 2% of total property income received by households (the sum of interest, rental income and dividends). Overall, summing up all forms of trust income – including other receipts and accruals, and excluding losses –, we only reach between 4.5% and 6% of total property income received by households, or 0.3% to 0.45% of the national income. This is very puzzling, and points to potentially large under-reporting, evasion or exemptions.

Bibliography

Alstadsæter, Annette, Niels Johannesen, and Gabriel Zucman (2019). "Tax evasion and inequality". In: *American Economic Review* 109.6, pp. 2073–2103.

Alstadsæter, Annette, Niels Johannesen, and Gabriel Zucman (2018). "Who owns the wealth in tax havens? Macro evidence and implications for global inequality". In: *Journal of Public Economics* 162.1, pp. 89–100.

Beer, Barend de, Nonhlanhla Nhlapo, Zeph Nhleko, et al. (2010). "A perspective on the South African flow of funds compilation–theory and analysis". In: *IFC Bulletin* 33, pp. 239–48.

Bharti, N. K. (2018). "Wealth inequality, class and caste in India, 1961-2012". WID.world Working Paper 2018/14.

Blanchet, Thomas (2016). "Wealth inequality in Europe and in the United States: estimations from surveys, national accounts and wealth rankings". Paris School of Economics Master Thesis.

Blanchet, Thomas, Lucas Chancel, and Amory Gethin (2022). "Why Is Europe More Equal Than the United States?" In: *American Economic Journal: Applied Economics* 14.4, pp. 1–40.

Dettling, L. J. et al. (2015). "Comparing Micro and Macro Sources for Household Accounts in the United States: Evidence from the Survey of Consumer Finances". Finance and Economics Discussion Series, Divisions of Research & Statistics and Monetary Affairs of the Federal Reserve Board.

Ebrahim, Amina and Chris Axelson (2019). "The creation of an individual level panel using administrative tax microdata in South Africa". UNU-WIDER Working Paper 2019/27.

Eighty 20 and XDS (2019). *Eighty20/XDS Credit Stress Report 2019*. Technical Report.

Garbinti, Bertrand, Jonathan Goupille-Lebret, and Thomas Piketty (2017). "Accounting for wealth inequality dynamics: Methods, estimates and simulations for France (1800-2014)". WID.world Working Paper 2016/5.

Henriques, A. M. and J. W. Hsu (2014). “Analysis of Wealth Using Micro- and Macrodata: A Comparison of the Survey of Consumer Finances and Flow of Funds Accounts”. In: *Measuring Economic Sustainability and Progress*. Ed. by D. W. Jorgenson, J. S. Landefeld, and P. Schreyer. University of Chicago Press. Chap. 9, pp. 245–274.

Muellbauer, John and Janine Aron (1999). “Estimates of personal sector wealth for South Africa”. CEPR Discussion Paper No. 4646.

Novokmet, Filip, Thomas Piketty, and Gabriel Zucman (2018). “From Soviets to oligarchs: inequality and property in Russia 1905-2016”. In: *Journal of Economic Inequality* 16.1, 189–223.

Orthofer, Anna (2015). “Private wealth in a developing country: A South African perspective on Piketty”. In: *ERSA [Economic Research Southern Africa] Working Paper* 564.

Piketty, Thomas, Li Yang, and Gabriel Zucman (2019). “Capital Accumulation, Private Property, and Rising Inequality in China, 1978-2015”. In: *American Economic Review* 109.7, pp. 2469–2496.

Saez, Emmanuel and Gabriel Zucman (2016). “Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data”. In: *Quarterly Journal of Economics* 131.May, pp. 519–578.

Ytterberg, Alan V and James P Weller (2010). “Managing Family Wealth Through a Private Trust Company”. In: *ACTEC Law Journal* 36.1, pp. 501–518.

G.4 Additional figures and tables

Figure G.1: The evolution of household wealth in South Africa, 1975-2018

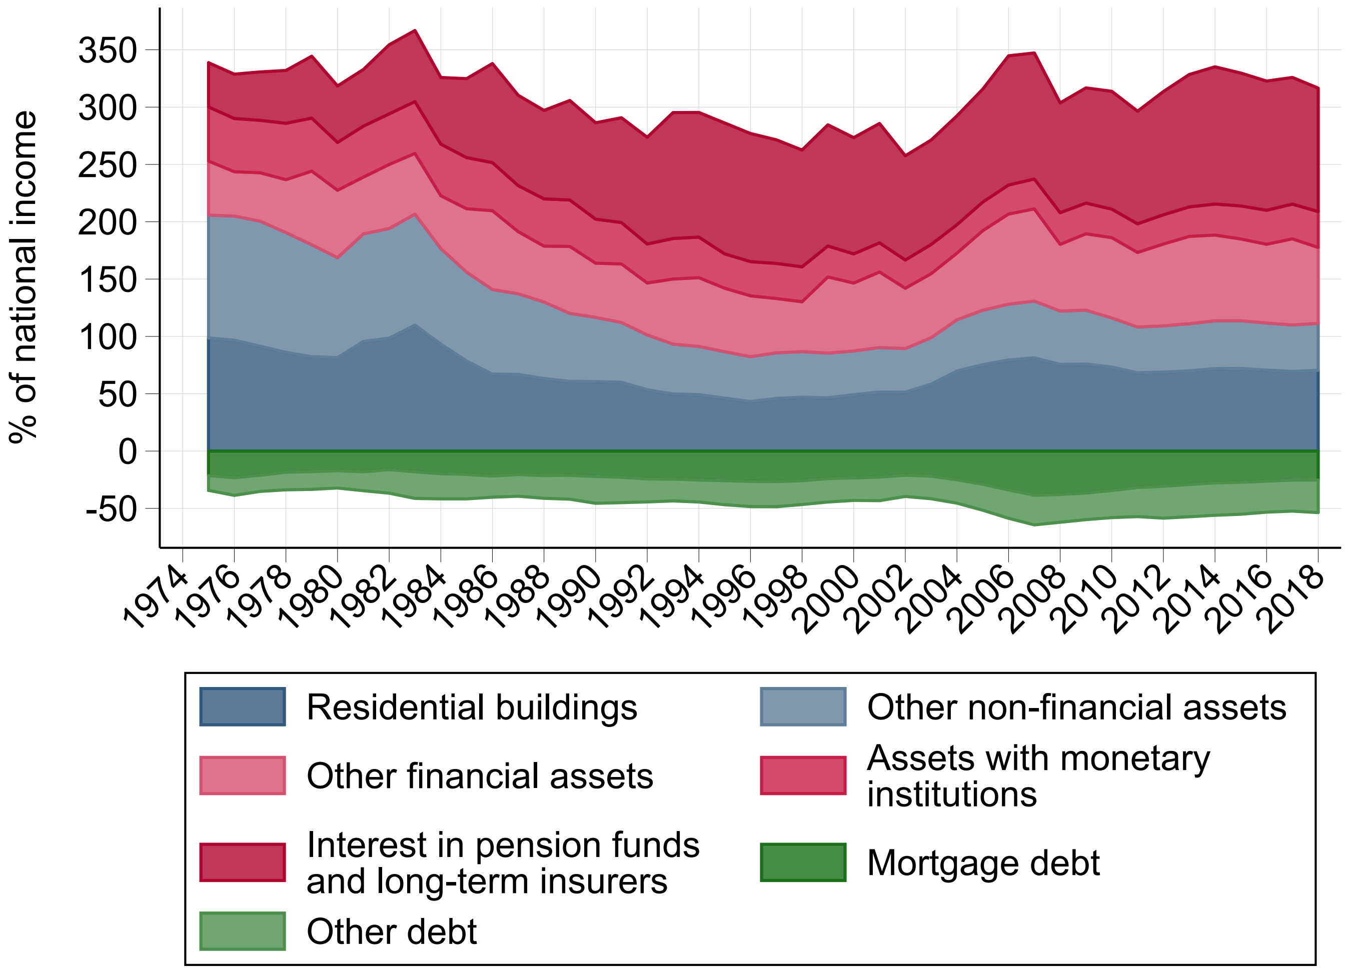

Notes: This figure shows the level and composition of household wealth in South Africa between 1975 and 2018, expressed as a share of the net national income.

Source: authors' compilation based on data from the South African Reserve Bank.

Figure G.2: South African wealth inequality in comparative perspective: Middle 40% wealth share

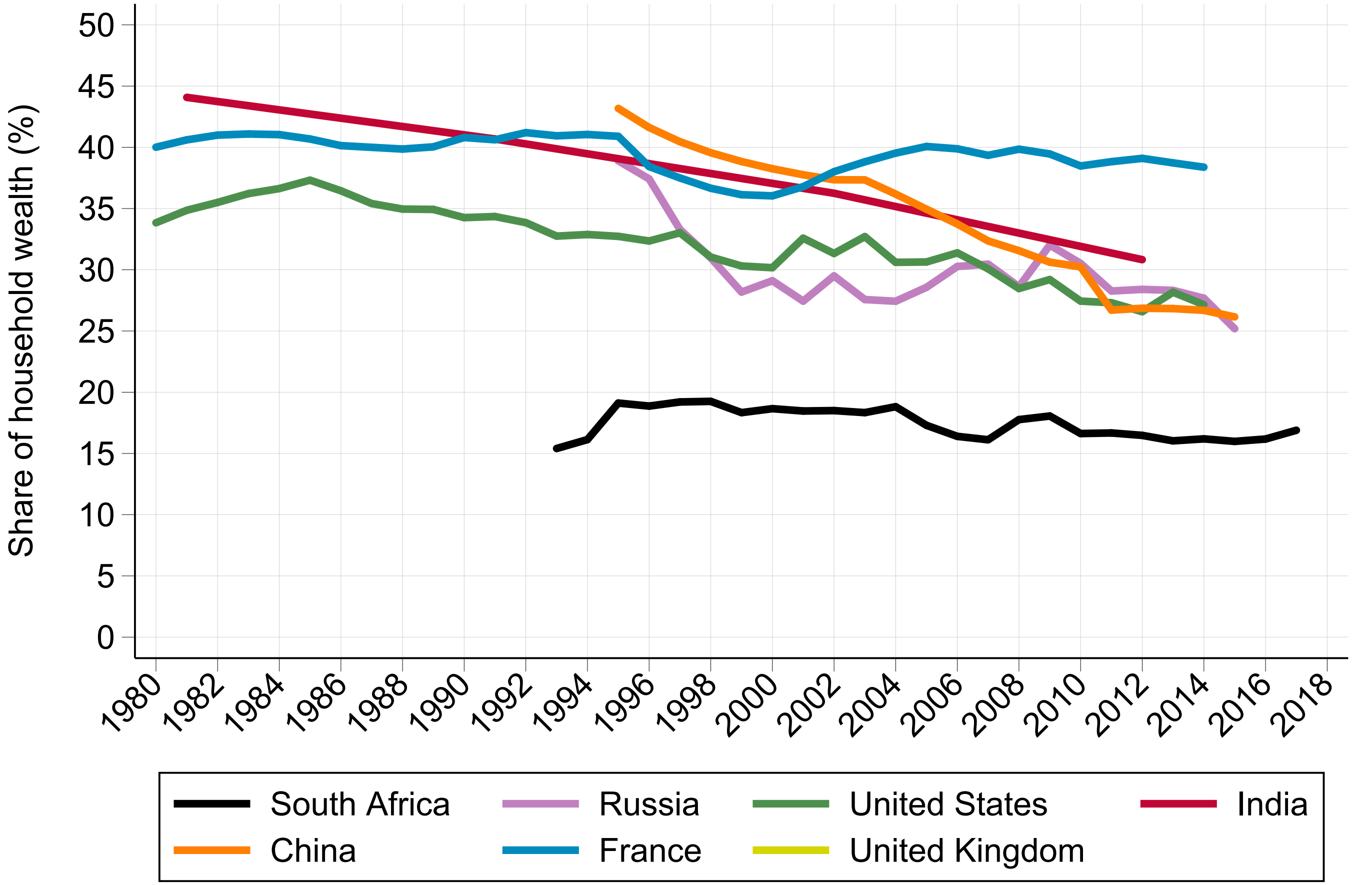

Notes: The figure compares the middle 40% wealth share in South Africa to that of other countries. The unit of observation is the individual adult aged 20 or above. Wealth is individualised (South Africa) or split equally among adult household members (other countries).

Source: authors' computations based on data for South Africa; World Inequality Database (`http://wid.world`) for other countries.

Figure G.3: South African wealth inequality in comparative perspective: Top 0.1% wealth share

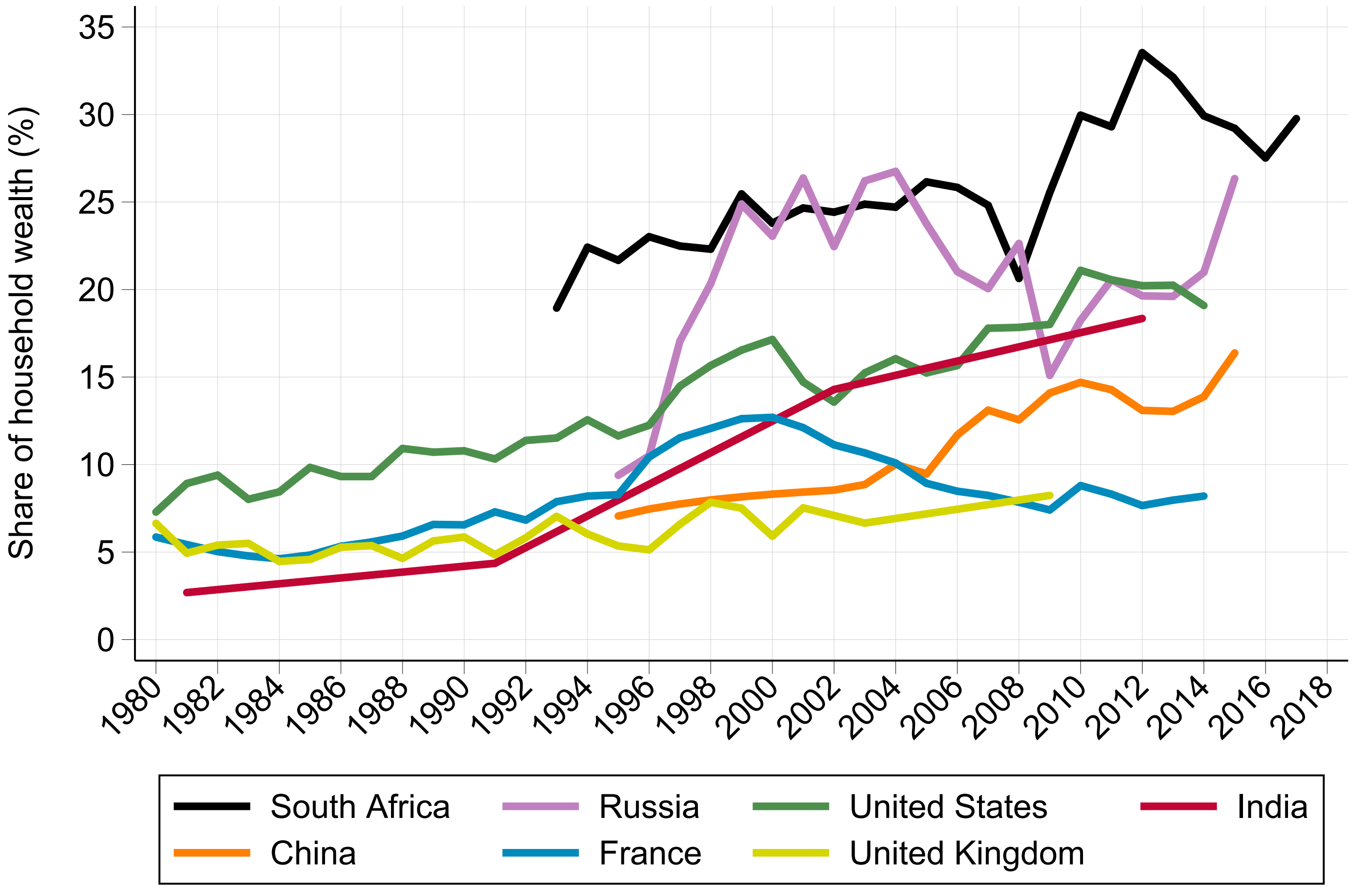

Notes: The figure compares the top 0.1% wealth share in South Africa to that of other countries. The unit of observation is the individual adult aged 20 or above. Wealth is individualised (South Africa) or split equally among adult household members (other countries).

Source: authors' computations based on data for South Africa; World Inequality Database (`http://wid.world`) for other countries.

Figure G.4: South African wealth inequality in comparative perspective: Bottom 50% wealth share

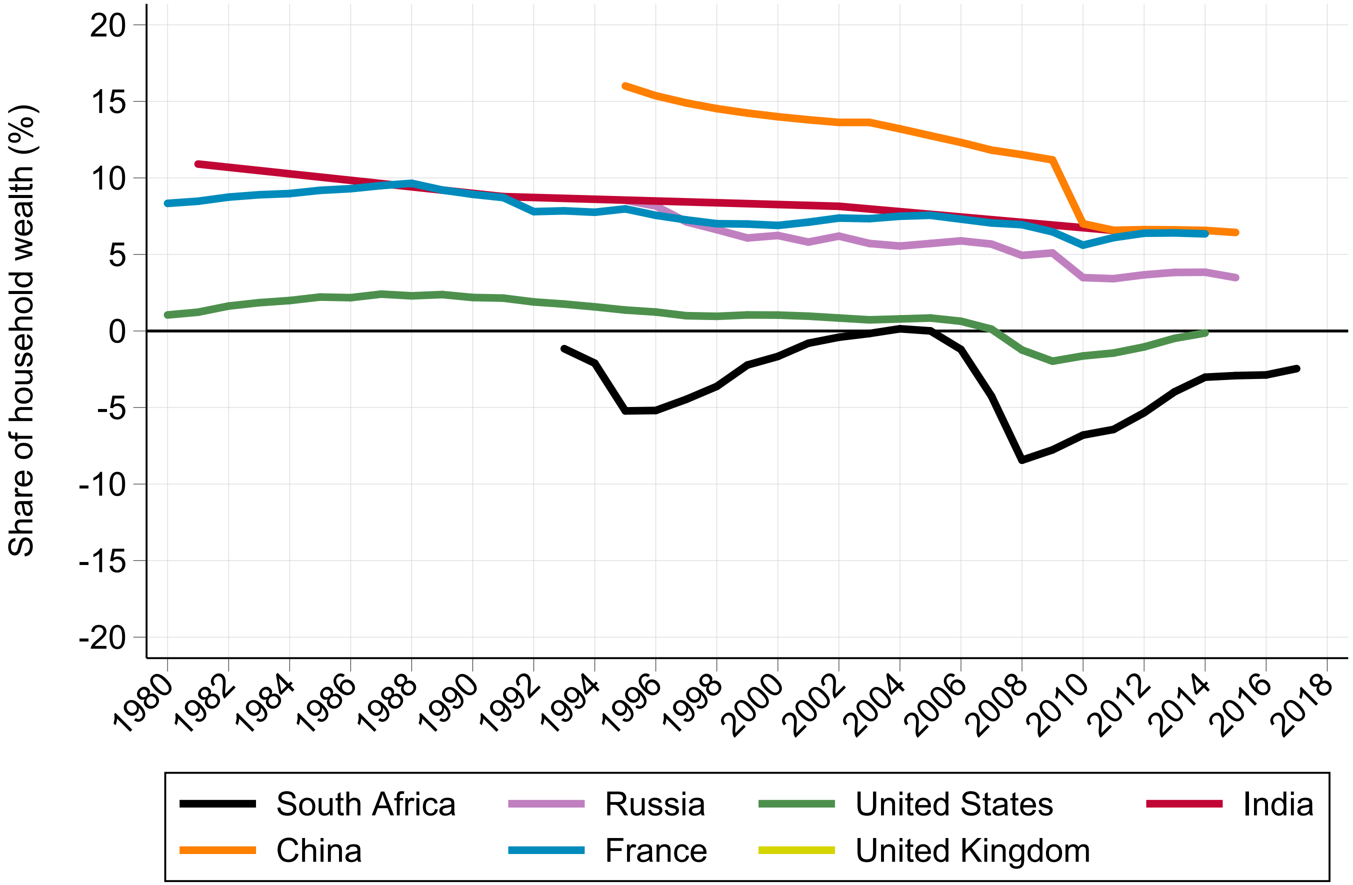

Notes: The figure compares the bottom 50% wealth share in South Africa to that of other countries. The unit of observation is the individual adult aged 20 or above. Wealth is individualised (South Africa) or split equally among adult household members (other countries).

Source: authors' computations based on data for South Africa; World Inequality Database (`http://wid.world`) for other countries.

Figure G.5: The evolution of household debt in South Africa, 1992-2018: the boom and bust of mortgage debt

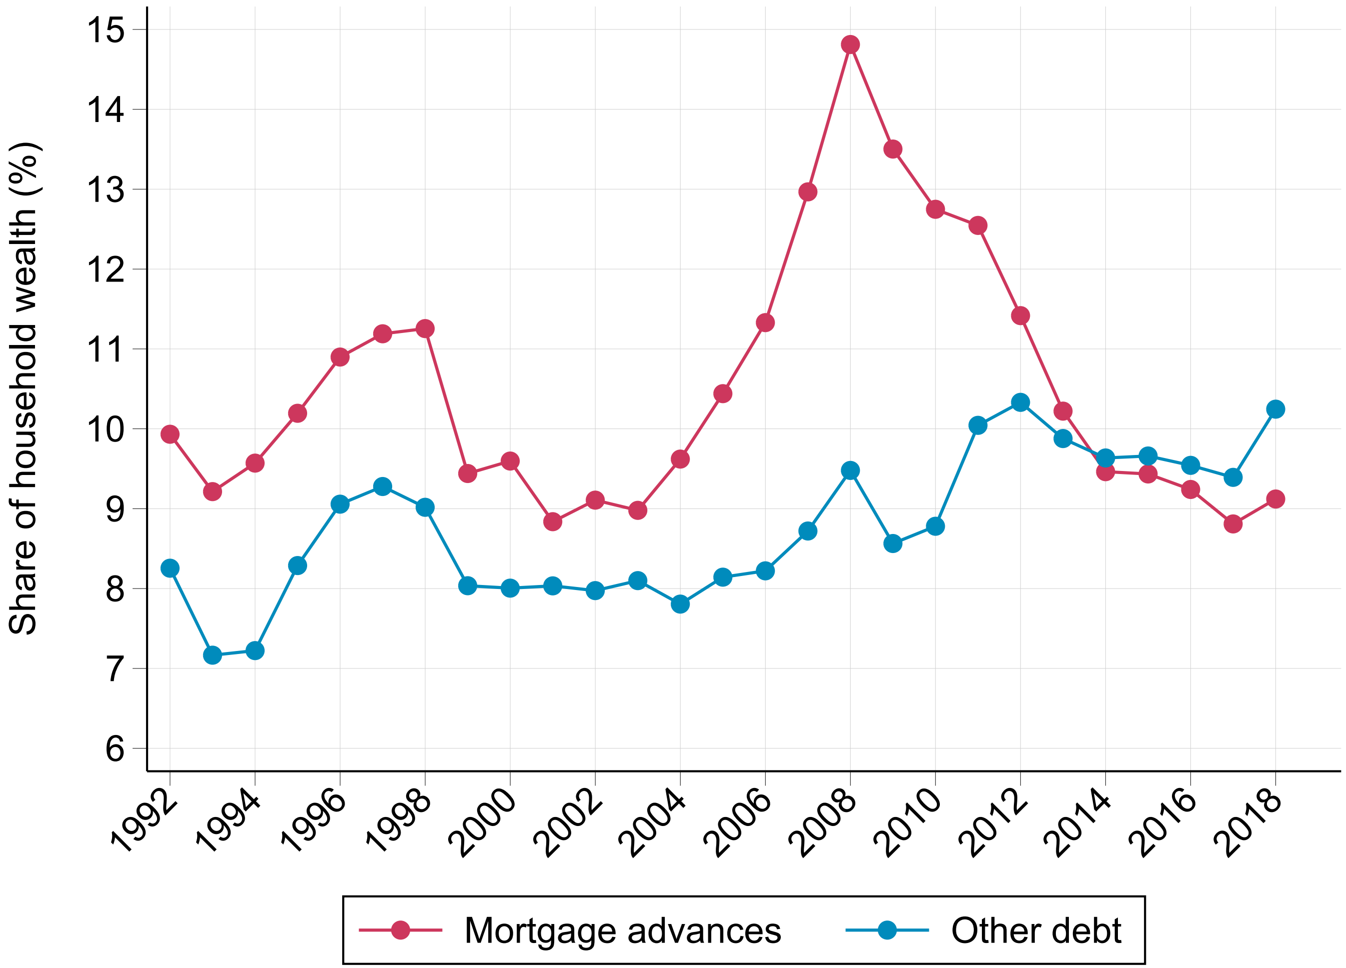

Notes: The figure shows the evolution of total household mortgage advances and total other household debts between 1992 and 2018, expressed as a share of household net wealth.

Source: authors' computations based on data from the SARB.

Figure G.6: Wealth inequality in NIDS: reported vs. capitalised pension wealth

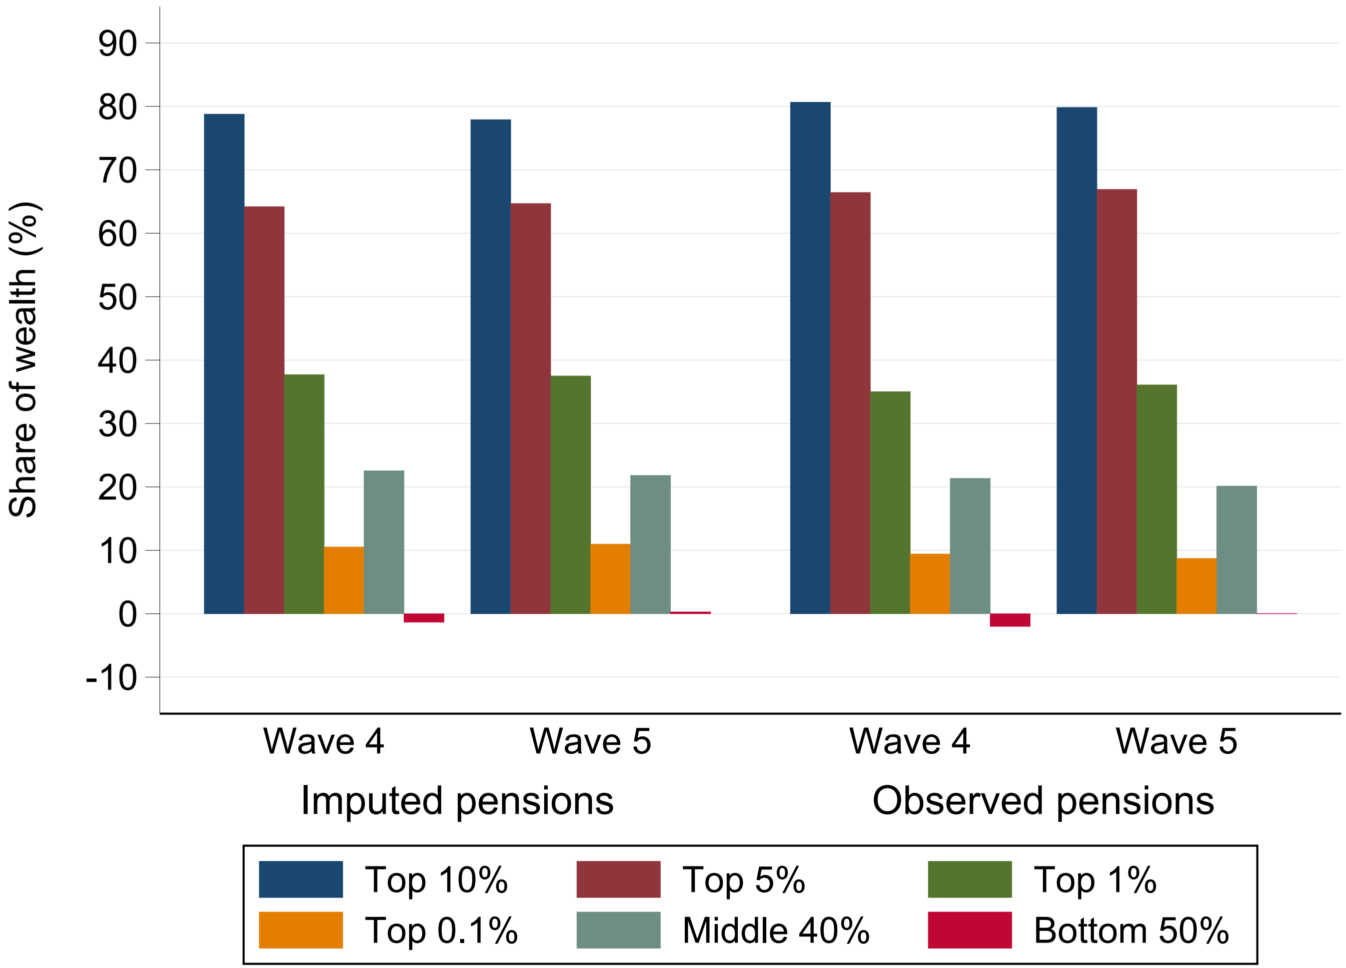

Notes: The figure compares the wealth shares estimated after capitalising pension wealth in NIDS (assuming that 75% of pension assets go to wage earners proportionally to pension contributions, and 25% belong to pensioners proportionally to pension income) to the wealth shares estimated by direct measurement of pension assets in NIDS.

Source: authors' computations based on data.

Figure G.7: Wealth inequality in NIDS: reported vs. capitalised life insurance assets

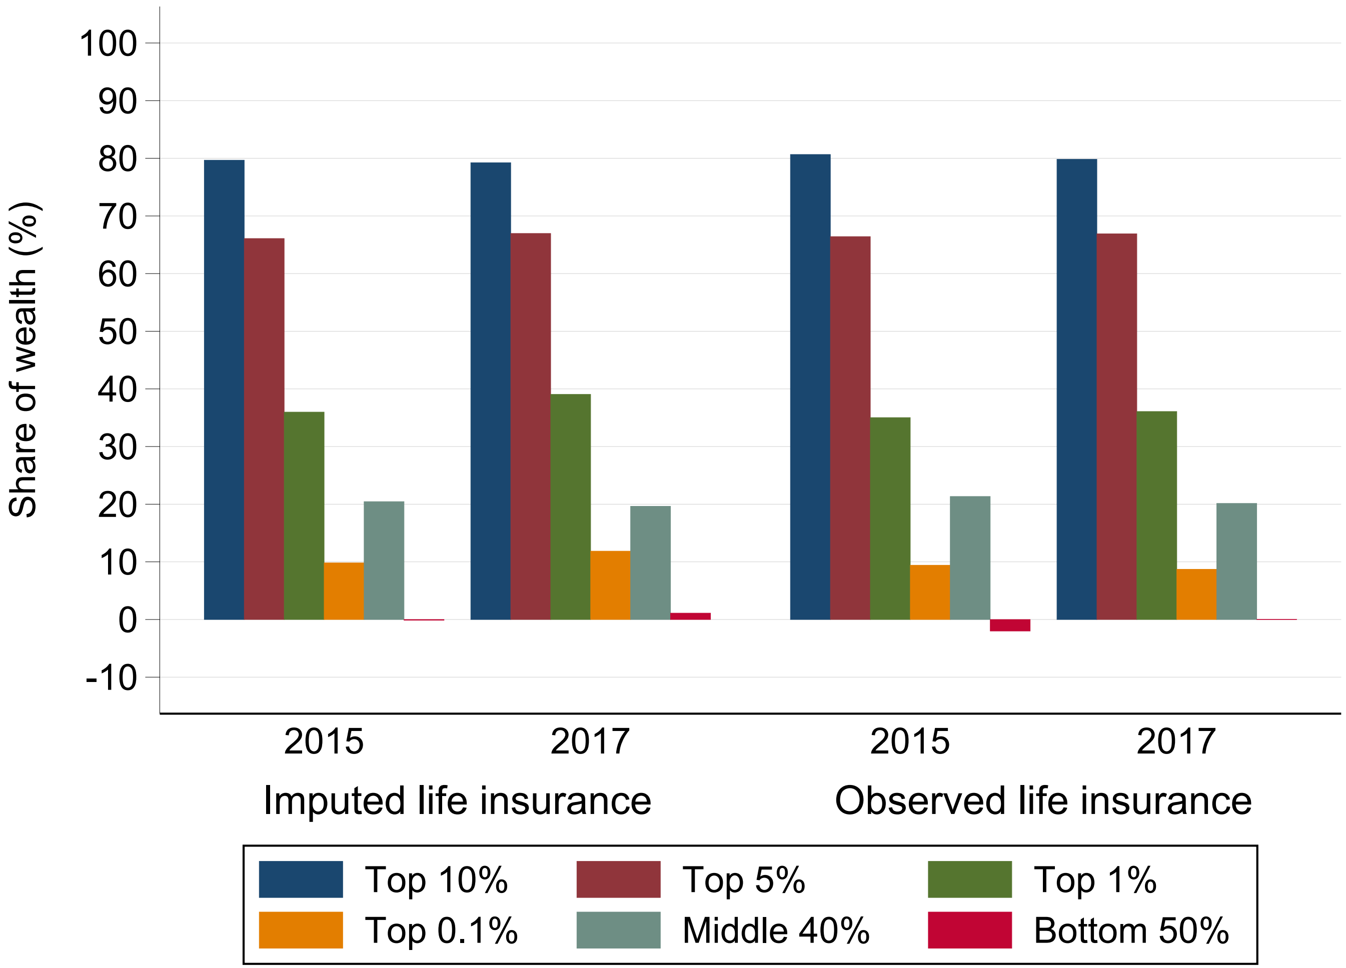

Notes: The figure compares the wealth shares estimated after capitalising life insurance assets in NIDS (assuming that 50% go to wage earners proportionally to factor income, and 50% to other earners proportionally to factor income) to the wealth shares estimated by direct measurement of life insurance assets in NIDS.

Source: authors' computations based on data.

Figure G.8: Combination of survey and tax data: quantile functions of merging income, 2017

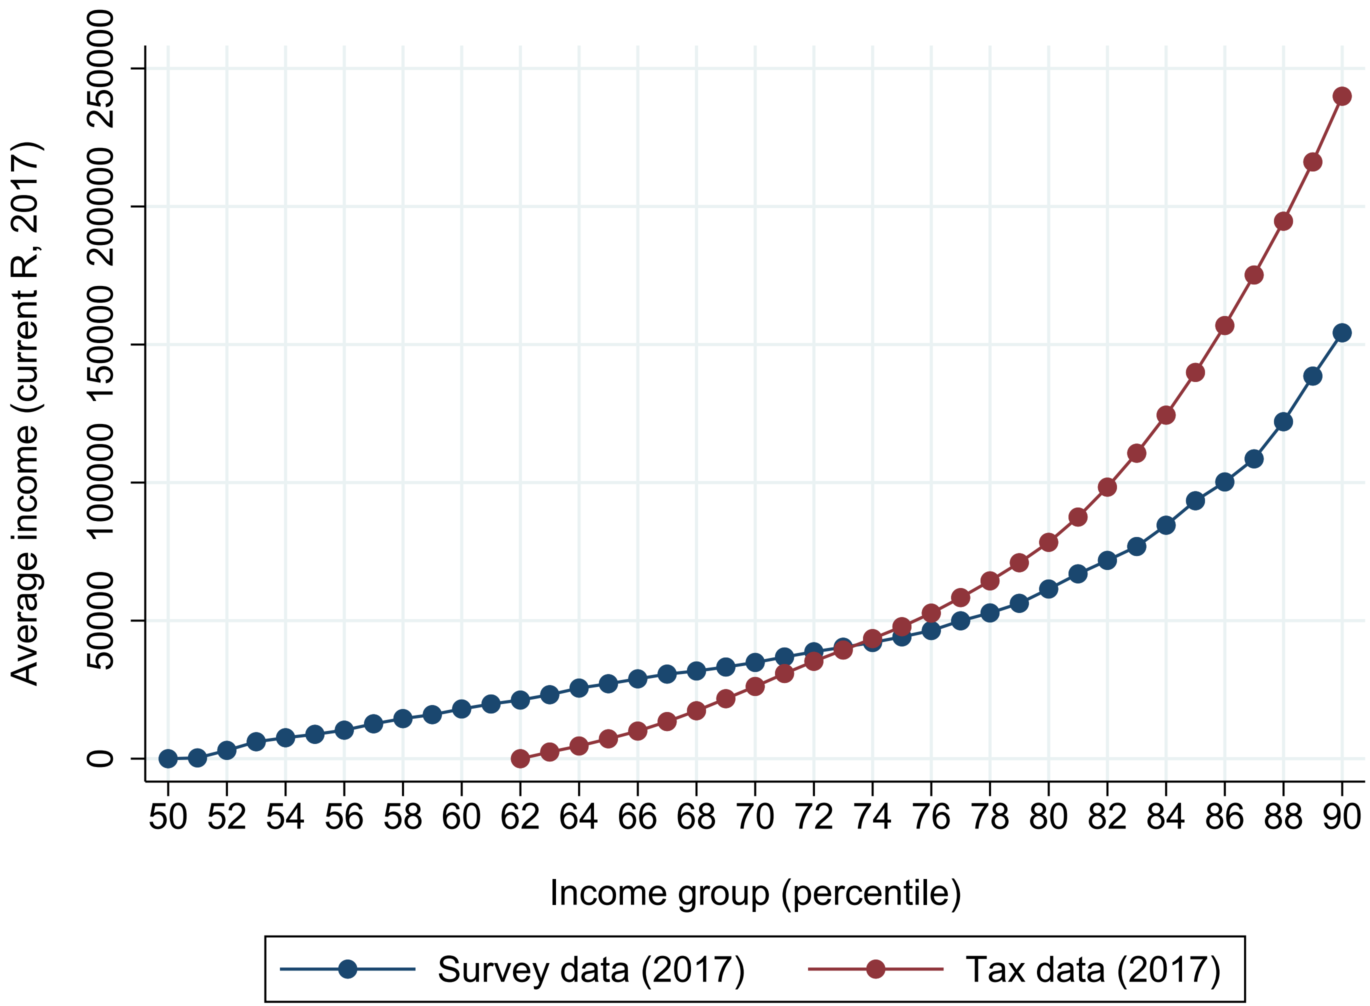

Notes: The figure compares the average merging income by percentile in the survey and in the tax microdata in 2017. Merging income is the sum of gross wages, business income, rental income, interest income and private pension income.

Source: authors' computations based on data.

Figure G.9: Combination of survey and tax data: ratio of quantile functions of merging income, 2010-2017

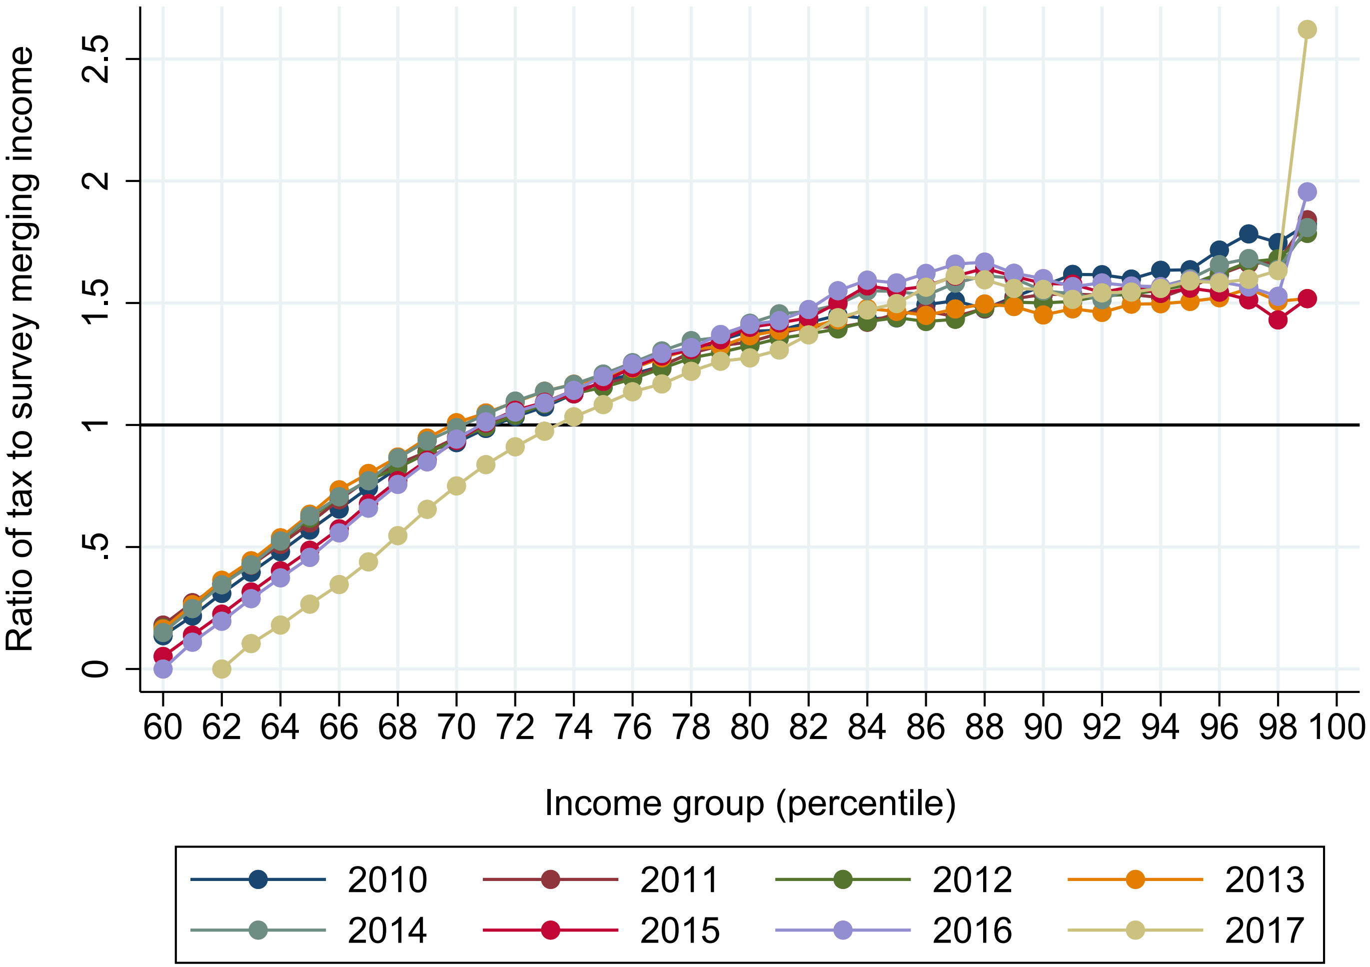

Notes: The figure plots the ratio of average merging income by percentile in the tax microdata to the harmonised survey data between 2010 and 2017. Merging income is the sum of gross wages, business income, rental income, interest income and private pension income.

Source: authors' computations based on data.

Figure G.10: Impact of changes in equivalence scales on wealth inequality: Top 10% and Top 1% shares

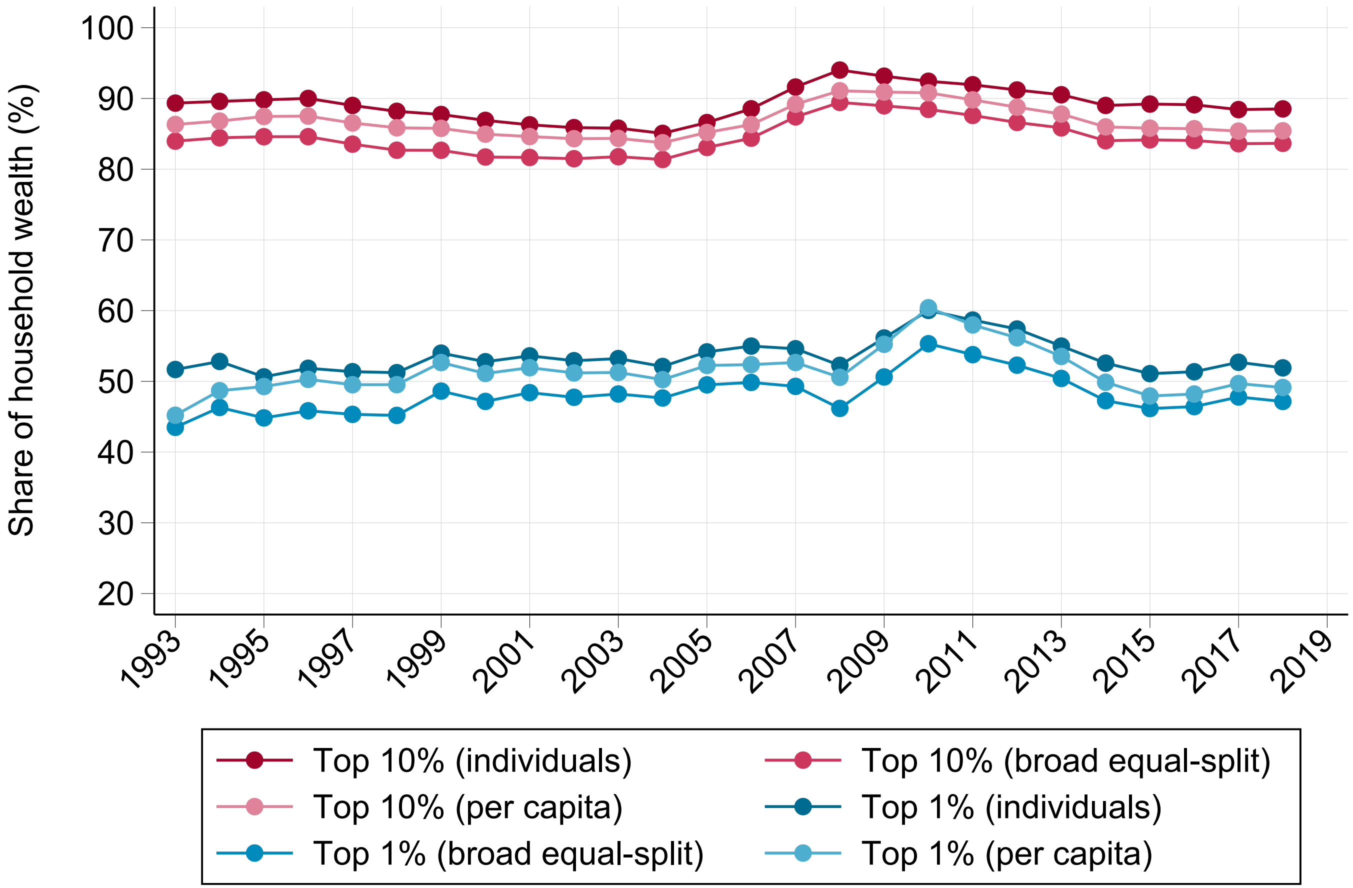

Notes: The figure compares the wealth shares estimated from the mixed method applied to household surveys depending on three different equivalence scales: individual series, broad equal-split series and per capita series.

Source: authors' computations based on data.

Figure G.11: Impact of changes in equivalence scales on wealth inequality: Middle 40% and Bottom 50% wealth shares

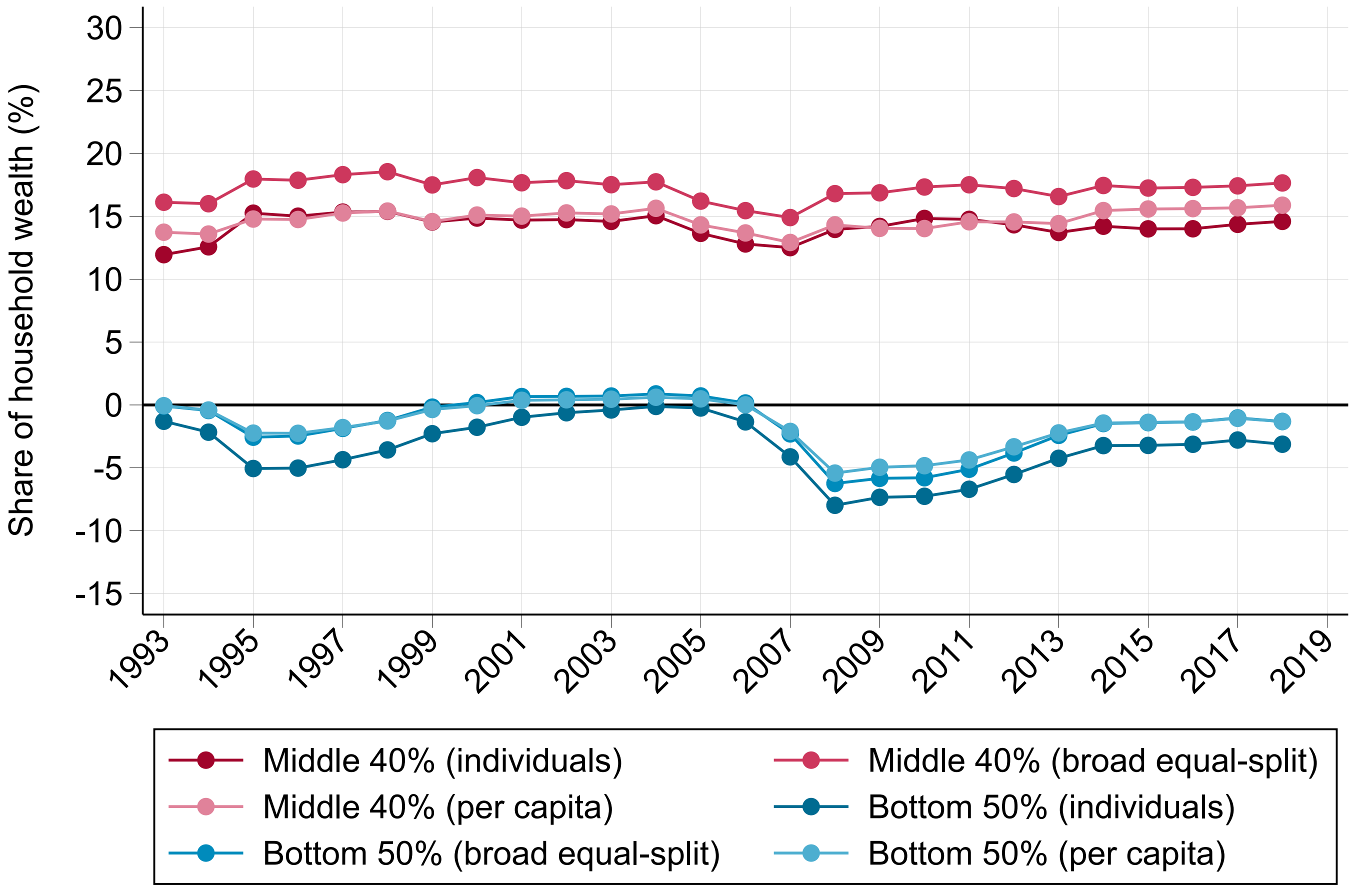

Notes: The figure compares the wealth shares estimated from the mixed method applied to household surveys depending on three different equivalence scales: individual series, broad equal-split series and per capita series.

Source: authors' computations based on data.

Figure G.12: Impact of changes in aggregate housing wealth on wealth inequality: Top 10% and top 1% wealth shares

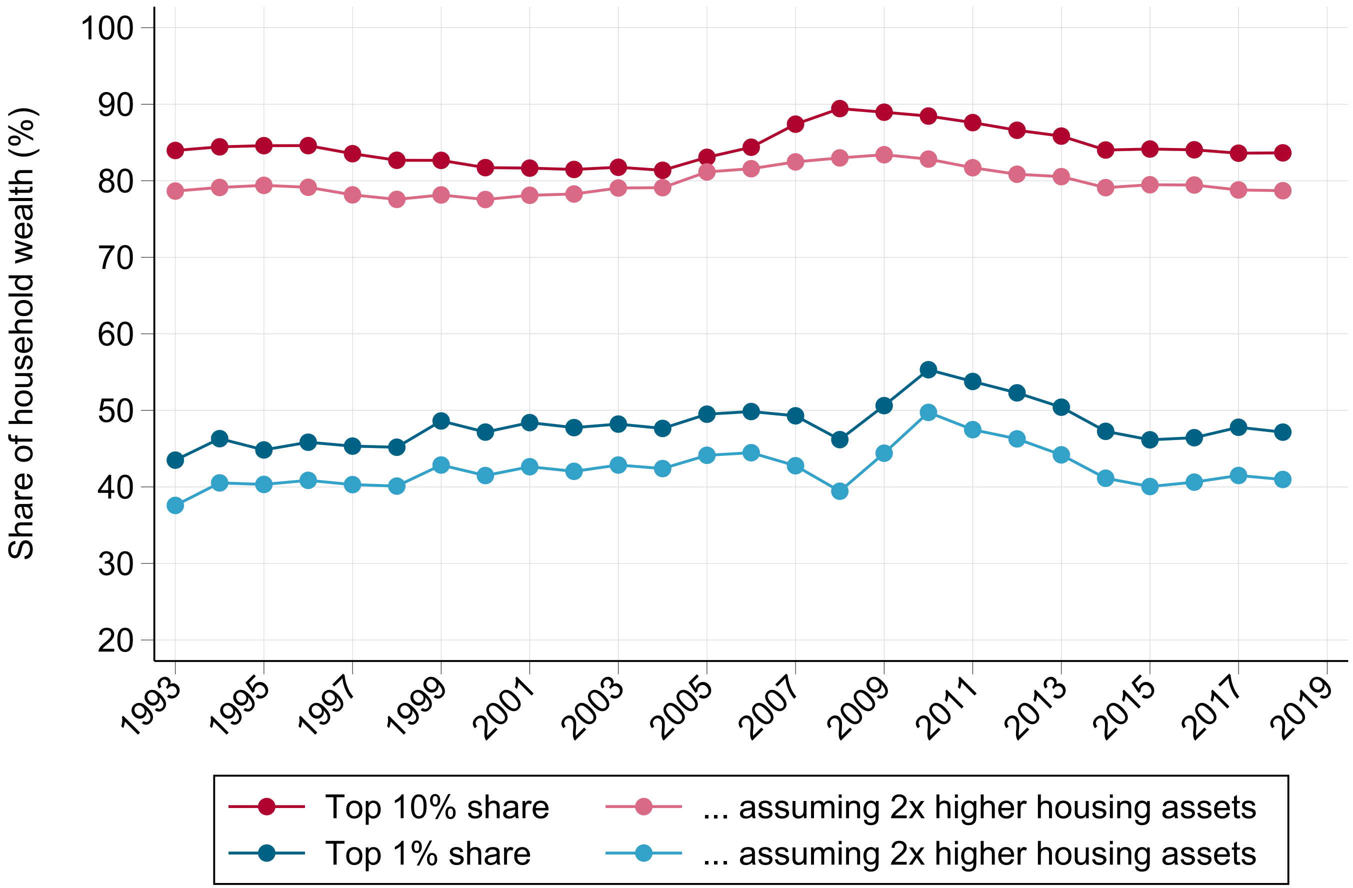

Notes: The figure compares the wealth shares estimated from the mixed method applied to household surveys under two scenarios: one in which total aggregated housing wealth corresponds to official balance sheets figures, and one in which it is estimated to be twice that amount.

Source: authors' computations based on data.

Figure G.13: Impact of changes in aggregate housing wealth on wealth inequality: Middle 40% and bottom 50% wealth shares

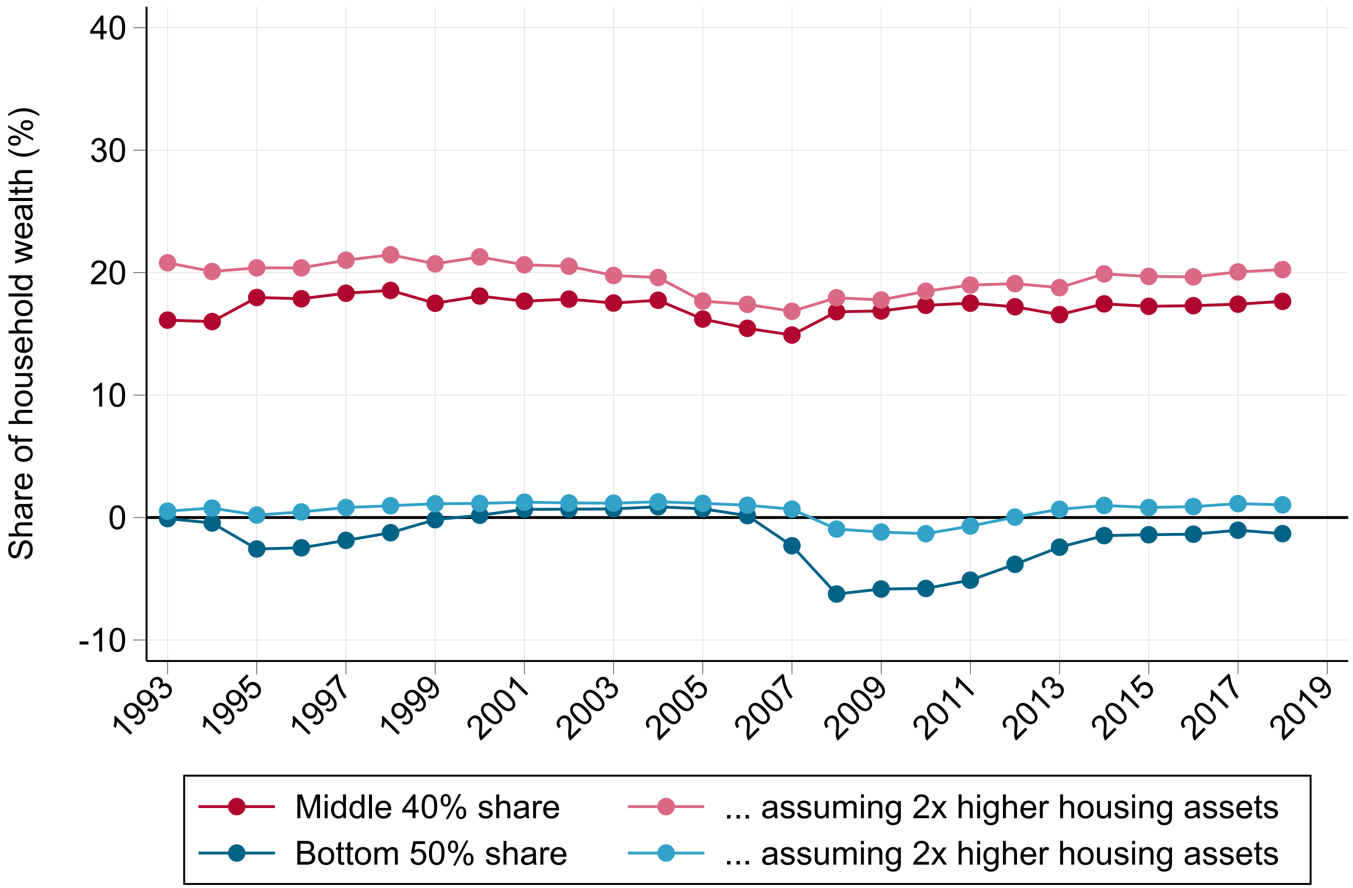

Notes: The figure compares the wealth shares estimated from the mixed method applied to household surveys under two scenarios: one in which total aggregated housing wealth corresponds to official balance sheets figures, and one in which it is estimated to be twice that amount.

Source: authors' computations based on data.

Figure G.14: Distribution of wealth vs. distribution of assets: top 10% and top 1% shares

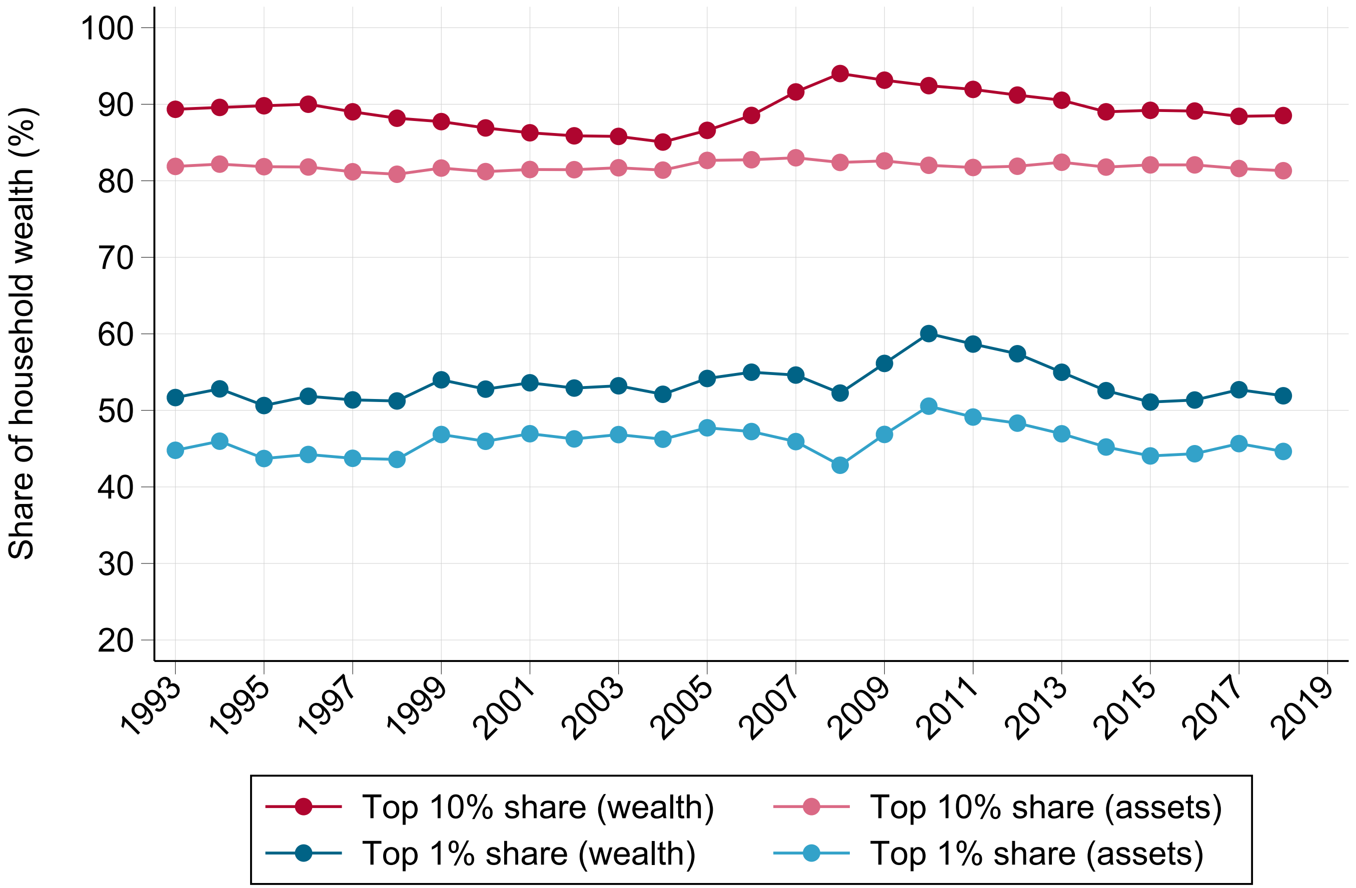

Notes: The figure compares the distribution of wealth and the distribution of assets (that is, excluding debt) in South Africa, estimated from surveys using the mixed method.

Source: authors' computations based on data.

Figure G.15: Comparison of methodologies: top 10% share

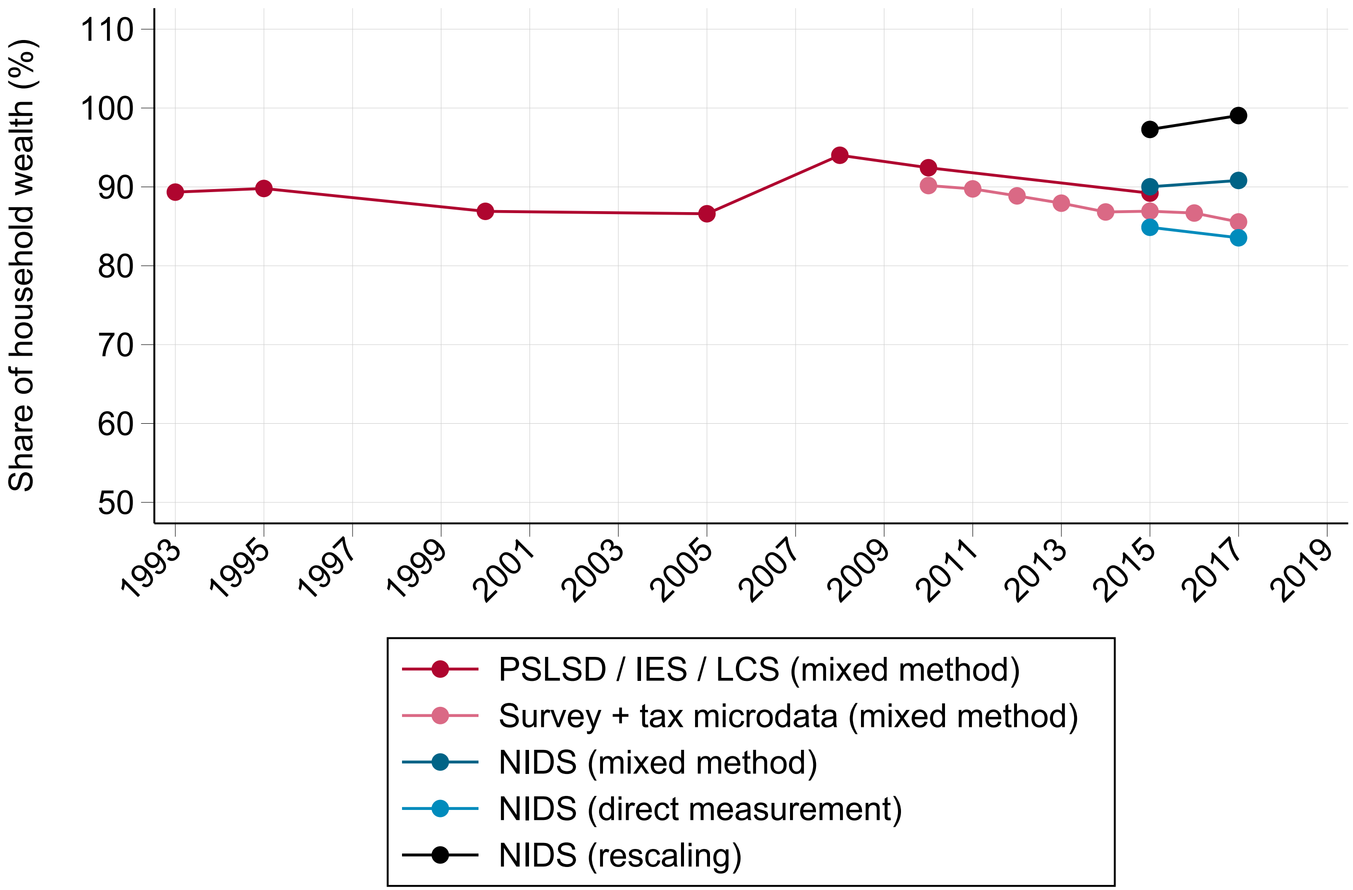

Notes: The figure compares the wealth shares estimated from the mixed method, direct measurement and rescaling of reported wealth components. Source: authors' computations based on data.

Figure G.16: Comparison of methodologies: top 1% share

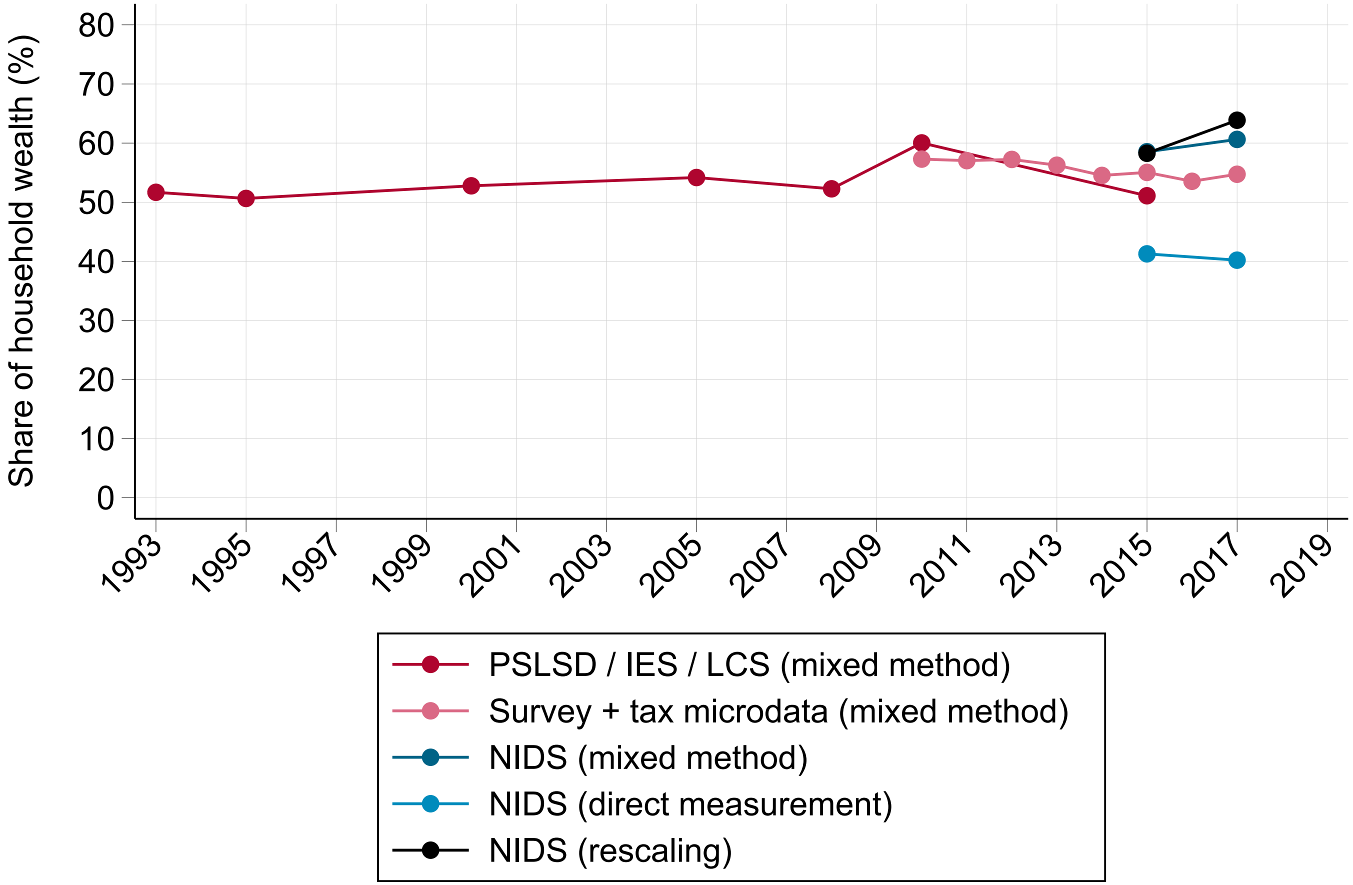

Notes: The figure compares the wealth shares estimated from the mixed method, direct measurement and rescaling of reported wealth components.
Source: authors' computations based on data.

Figure G.17: Comparison of methodologies: top 0.1% share

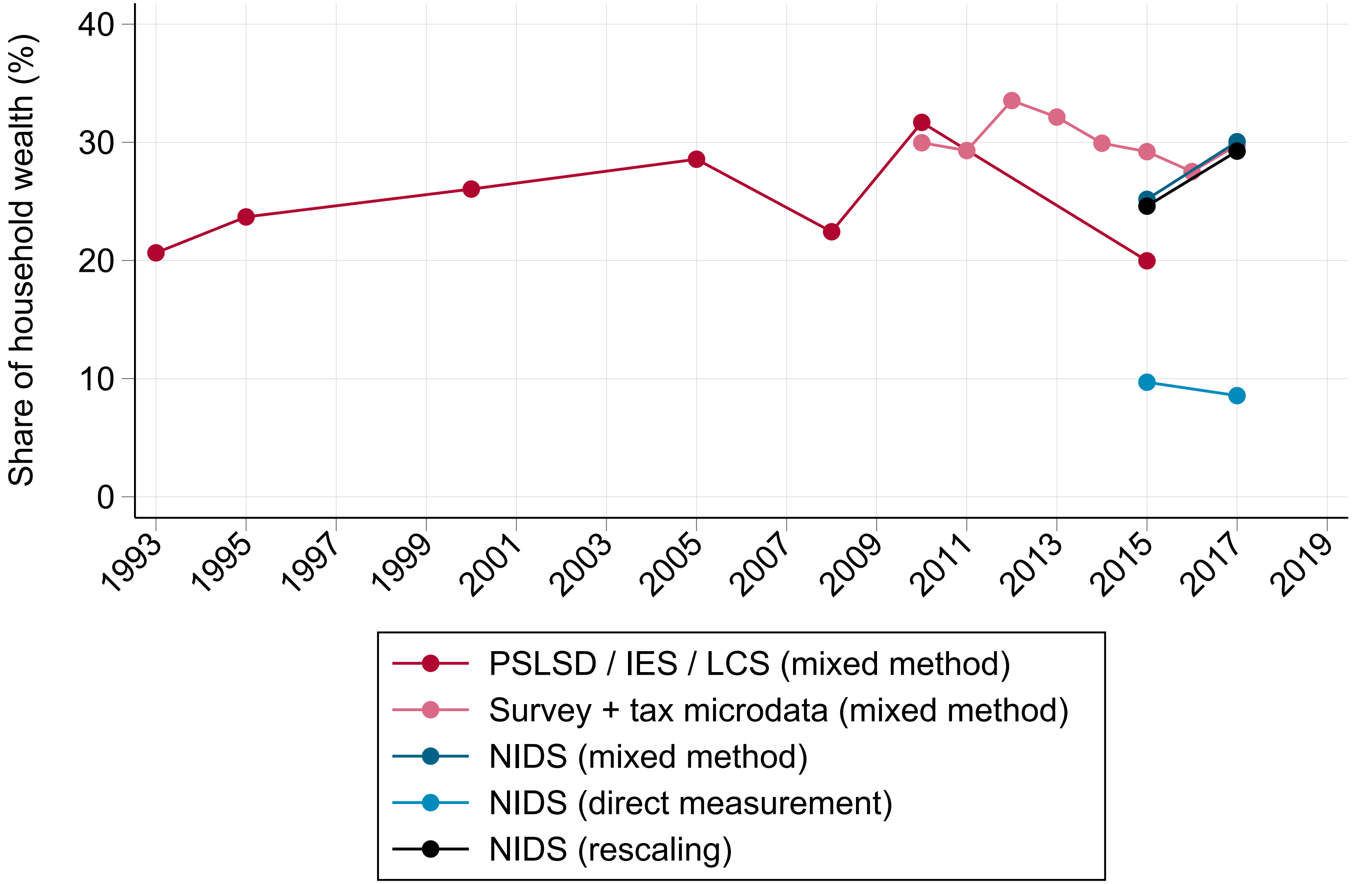

Notes: The figure compares the wealth shares estimated from the mixed method, direct measurement and rescaling of reported wealth components. Source: authors' computations based on data.

Figure G.18: Share of financial assets held through trusts, 1975-2018

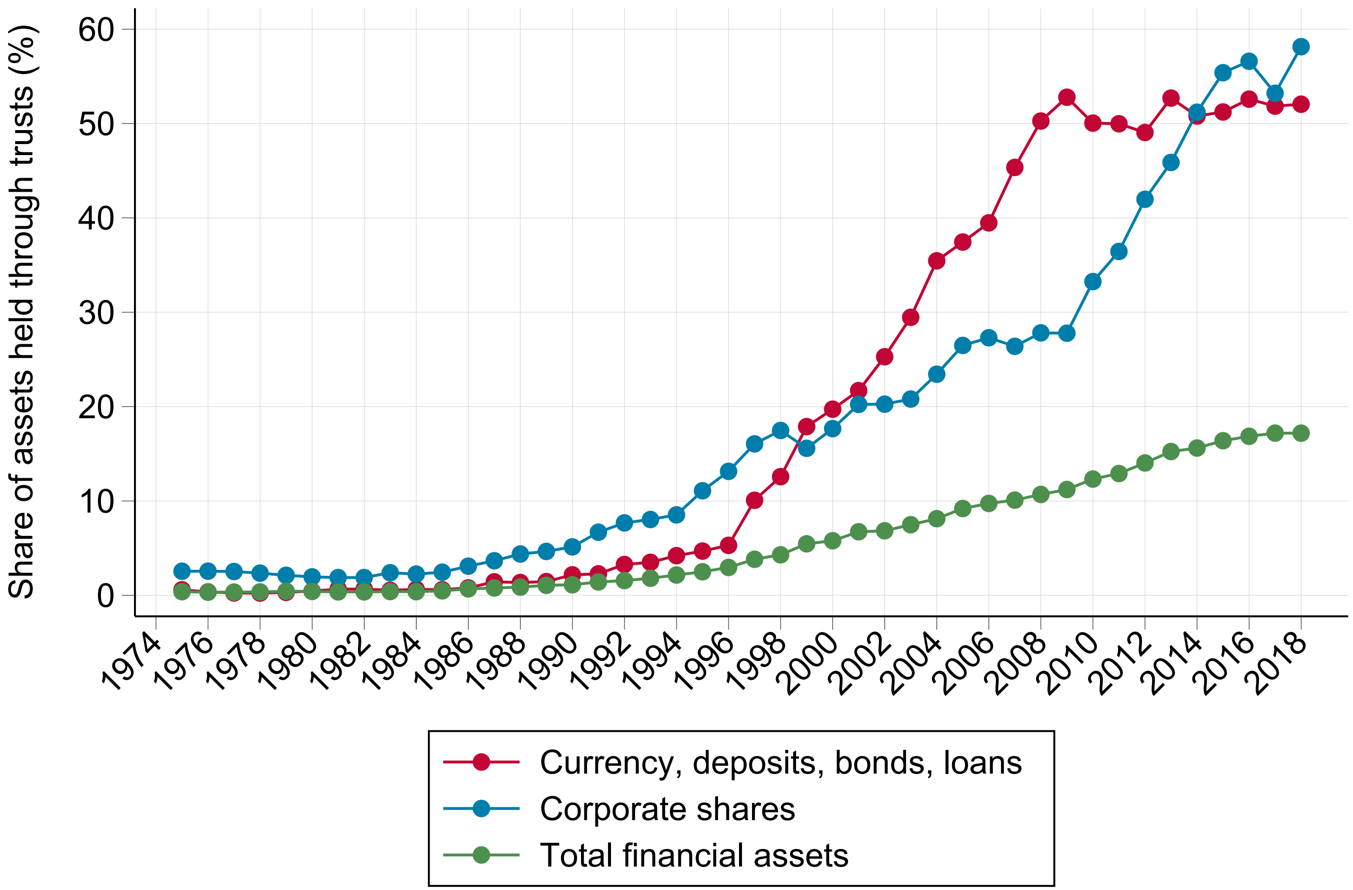

Notes: The figure shows the share of total household assets in the economy held by unit trusts.

Source: authors' compilation based on data from the SARB.

Table G.1: The level and composition of household wealth in South Africa in 2018

	Market value (R billion)	% of national income	% of net wealth
Non-financial assets	4504	111.4 %	42.4 %
Owner-occupied housing	3020	74.7 %	28.4 %
Tenant-occupied housing	988	24.4 %	9.3 %
Business assets	497	12.3 %	4.7 %
Financial assets	8294	205.1 %	78.0 %
Pension assets	2944	72.8 %	27.7 %
Life insurance assets	1412	34.9 %	13.3 %
Bonds and interest deposits	1798	44.5 %	16.9 %
Currency, notes and coins	87	2.2 %	0.8 %
Corporate shares	2053	50.8 %	19.3 %
Total liabilities	2170	53.7 %	20.4 %
Mortgage debt	1022	25.3 %	9.6 %
Non-mortgage debt	1148	28.4 %	10.8 %
Net household wealth	10629	262.9 %	100.0 %
Offshore wealth	575	14.2 %	5.4 %
Net wealth incl. offshore wealth	11204	277.1 %	105.4 %

Notes: The table shows the level and composition of household wealth in South Africa in 2018. The market value of each component is expressed in current billion rands. Source: Own estimates combining available data sources from the SARB.

Table G.2: Ownership rates and coverage of household balance sheets by asset class in NIDS

	% of adults with asset or debt		% of balance sheets covered	
	Wave 4	Wave 5	Wave 4	Wave 5
Household assets				
Owner-occupied housing	72.3 %	65.2 %	151.7 %	220.8 %
Tenant-occupied housing	3.3 %	3.5 %	122.4 %	97.2 %
Business assets	5.6 %	5.0 %	135.4 %	59.6 %
Pension and life insurance	25.7 %	24.4 %	110.0 %	104.3 %
Bonds and stock	1.5 %	1.3 %	3.9 %	3.8 %
Household debts				
Mortgage debt	8.0 %	7.0 %	71.0 %	56.8 %
Other debts	36.3 %	33.7 %	54.5 %	37.0 %

Notes: The table shows the share of South Africans who declare having a particular type of asset or debt, along with the share of the total value of this asset or debt in the economy captured by the NIDS survey. Source: authors' computations based on data. The unit of observation is the adult individual aged 20 or above. Calculations based on weighted sample using design weights.

Table G.3: The coverage of owner-occupied housing, mortgage debt and other debt in South African surveys

	Owner-occupied housing	Mortgage debt	Other debt
PSLSD, 1993	143.5 %	86.5 %	37.4 %
IES, 1995	121.7 %	27.2 %	16.5 %
IES, 2000		40.3 %	34.9 %
IES, 2005	105.9 %	67.9 %	41.5 %
IES, 2010	193.9 %	16.4 %	20.5 %
LCS, 2008	145.4 %	13.9 %	18.4 %
LCS, 2015	179.5 %	51.0 %	22.2 %
NIDS, wave 4	122.3 %	74.3 %	57.4 %
NIDS, wave 5	258.8 %	56.8 %	37.0 %

Notes: The table shows the ratio of total assets or debts reported in surveys to the corresponding totals reported in the household balance sheets. PSLSD: Project for Statistics on Living Standards and Development. IES: Income and Expenditure Survey. LCS: Living Conditions Survey. NIDS: National Income Dynamics Study.
Source: authors' computations based on data. The unit of observation is the adult individual aged 20 or above. Calculations based on weighted samples using weights calibrated by the authors' (see appendix).

Table G.4: The coverage of selected national accounts components in South African surveys

	Gross wages	Mixed income	Rental income	Interest and dividends
PSLSD, 1993	87.7 %	51.7 %	38.4 %	11.5 %
IES, 1995	76.9 %	55.0 %	9.9 %	8.8 %
IES, 2000	70.9 %	37.2 %	23.1 %	3.4 %
IES, 2005	80.5 %	64.2 %	21.7 %	3.8 %
IES, 2010	80.2 %	71.9 %	13.5 %	4.5 %
LCS, 2008	77.7 %	75.8 %	16.3 %	8.4 %
LCS, 2015	74.6 %	86.8 %	21.6 %	12.6 %
NIDS, wave 1	62.7 %	12.0 %	65.4 %	7.3 %
NIDS, wave 2	67.6 %	4.1 %	13.0 %	0.8 %
NIDS, wave 3	65.7 %	20.6 %	20.7 %	7.3 %
NIDS, wave 4	73.5 %	12.9 %	43.9 %	2.5 %
NIDS, wave 5	72.1 %	14.1 %	41.0 %	5.5 %

Notes: The table shows the ratio of total income reported in surveys to the total corresponding income reported in the national accounts published by the SARB. PSLSD: Project for Statistics on Living Standards and Development. IES: Income and Expenditure Survey. LCS: Living Conditions Survey. NIDS: National Income Dynamics Study. Source: authors' computations based on data. The unit of observation is the adult individual aged 20 or above. Calculations based on weighted samples using weights calibrated by the authors' (see appendix).

Table G.5: Shares of household wealth held by groups in South Africa: results from tax microdata and survey combined

	Bottom 50%	Middle 40%	Top 10%	Top 1%	Top 0.1%
2010	-6.8 %	16.6 %	90.2 %	57.3 %	30.0 %
2011	-6.4 %	16.7 %	89.8 %	57.0 %	29.3 %
2012	-5.3 %	16.5 %	88.9 %	57.2 %	33.5 %
2013	-4.0 %	16.0 %	87.9 %	56.3 %	32.1 %
2014	-3.0 %	16.2 %	86.8 %	54.5 %	29.9 %
2015	-2.9 %	16.0 %	86.9 %	55.0 %	29.2 %
2016	-2.9 %	16.2 %	86.7 %	53.5 %	27.5 %
2017	-2.5 %	16.9 %	85.6 %	54.7 %	29.8 %

Notes: The table shows estimates of the share of household wealth owned by the bottom 50% (p0p50), the middle 40% (p50p90), the top 10% (p90p100), the top 1% (p99p100 and the top 0.1% (p99.9p100) obtained from the income capitalisation method combining surveys and tax microdata. The unit of observation is the individual adult aged 20 or above. Source: authors' computations based on data.

Table G.6: Trust data (ITR12T) descriptive statistics

	2014	2015	2016	2017	2018
Number of trusts	138859	134106	127457	115825	93379
Dividends (% of household dividends)	0.0%	0.3%	0.5%	0.5%	0.3%
Interest income (% of household interest)	3.1%	2.9%	2.5%	2.6%	1.7%
Capital gain (% of property income)	1.3%	1.6%	2.4%	1.4%	0.6%
Rental income (% of household rental income)	2.4%	2.4%	2.1%	1.9%	1.4%
Business income (% of mixed income)	1.7%	1.6%	1.6%	1.4%	1.0%
Total trust income (% of property income)	4.6%	5.2%	5.9%	4.7%	2.9%

Notes: The table provides information on the number of trusts filing ITR12T forms in South Africa, as well as coverage of selected national income components. Source: authors' computations based on data.

Table G.7: The coverage of capital income in the tax microdata

	Rental income	Interest income	Dividends
2010	9.5 %	25.4 %	2.4 %
2011	11.7 %	25.0 %	5.3 %
2012	12.3 %	28.3 %	3.9 %
2013	13.4 %	28.8 %	5.2 %
2014	12.1 %	27.8 %	25.1 %
2015	12.3 %	27.8 %	10.6 %
2016	13.7 %	31.0 %	13.1 %
2017	6.9 %	18.3 %	15.8 %

Notes: The table shows the ratio of total income reported in the tax microdata to the corresponding total reported in the national accounts published by the SARB. Source: authors' computations based on data.

Table G.8: Source codes categories used in tax microdata

Income concept	**Source code**	**Description**
Gross wage	3601	Income (subject to PAYE)
Gross wage	3602	Income (non-taxable)
Gross wage	3605	Annual payment (subject to PAYE)
Gross wage	3606	Commission (subject to PAYE)
Gross wage	3607	Overtime (subject to PAYE)
Gross wage	3608	Arbitration award (subject to PAYE)
Gross wage	3609	Arbitration award (non-taxable)
Gross wage	3611	Purchased annuity (subject to PAYE)
Gross wage	3612	Purchased annuity (non-taxable)
Gross wage	3613	Restraint of trade (subject to PAYE)
Gross wage	3615	Director's remuneration (subject to PAYE)
Gross wage	3616	Independent contractors (subject to PAYE)
Gross wage	3617	Labour Brokers (subject to PAYE)
Gross wage	3619	Labour Brokers (IT)
Gross wage	3620	Directors fees RSA resident
Gross wage	3621	Directors fees non-resident
Gross wage	3651	Foreign income (subject to paye)
Gross wage	3652	Foreign income (non-taxable)
Gross wage	3655	Foreign annual payment (subject to paye)
Gross wage	3656	Foreign commission (subject to paye)
Gross wage	3657	Foreign overtime (subject to paye)
Gross wage	3658	Foreign arbitration award (subject to paye)
Gross wage	3659	Foreign arbitration award (non-taxable)
Gross wage	3661	Foreign purchased annuity (subject to paye)
Gross wage	3662	Foreign purchased annuity (non-taxable)
Gross wage	3663	Foreign restraint of trade (subject to paye)
Gross wage	3665	Foreign director's remuneration (subject to paye)
Gross wage	3666	Foreign independent contractors (subject to paye)
Gross wage	3667	Foreign labour brokers (subject to paye)
Gross wage	3669	Foreign labour brokers (it)
Gross wage	3670	Foreign directors fees rsa resident
Gross wage	3701	Travel allowance (subject to PAYE)

Table G.8: Source codes categories used in tax microdata

Income concept	Source code	Description
Gross wage	3702	Reimbursive travel allowance (IT)
Gross wage	3703	Reimbursive travel allowance (non-taxable)
Gross wage	3704	Subsistence allowance local travel (IT)
Gross wage	3705	Subsistence allowance local travel (non-taxable)
Gross wage	3706	Entertainment allowance (subject to PAYE)
Gross wage	3707	Share options exercised (subject to PAYE)
Gross wage	3708	Public office allowance (subject to PAYE)
Gross wage	3709	Uniform allowance (non-taxable)
Gross wage	3710	Tool allowance (subject to PAYE)
Gross wage	3711	Computer allowance (subject to PAYE)
Gross wage	3712	Telephone allowance (subject to PAYE)
Gross wage	3713	Other allowances (subject to PAYE)
Gross wage	3714	Other allowances (non-taxable)
Gross wage	3715	Subsistence allowance foreign travel (IT)
Gross wage	3716	Subsistence allowance foreign travel (non-taxable)
Gross wage	3722	Reimbursive travel allowance
Gross wage	3751	Foreign travel allowance (subject to paye)
Gross wage	3752	Foreign reimbursive travel allowance (it)
Gross wage	3753	Foreign reimbursive travel allowance (non-taxable)
Gross wage	3754	Foreign subsistence allowance local travel (it)
Gross wage	3755	Foreign subsistence allowance local travel (non-taxable)
Gross wage	3756	Foreign entertainment allowance (subject to paye)
Gross wage	3757	Foreign share options exercised (subject to paye)
Gross wage	3758	Foreign public office allowance (subject to paye)
Gross wage	3759	Foreign uniform allowance (non-taxable)
Gross wage	3760	Foreign tool allowance (subject to paye)
Gross wage	3761	Foreign computer allowance (subject to paye)
Gross wage	3762	Foreign telephone allowance (subject to paye)
Gross wage	3763	Foreign other allowances (subject to paye)
Gross wage	3764	Foreign other allowances (non-taxable)
Gross wage	3765	Foreign subsistence allowance foreign travel (it)
Gross wage	3766	Foreign subsistence allowance foreign travel (non-taxable)

Table G.8: Source codes categories used in tax microdata

Income concept	Source code	Description
Gross wage	3772	Foreign reimbursive travel allowance
Gross wage	3801	General fringe benefits (subject to PAYE)
Gross wage	3802	Use of motor acquired by employer not via operating lease (subject to PAYE)
Gross wage	3803	Use of asset (subject to PAYE)
Gross wage	3804	Meals etc (subject to PAYE)
Gross wage	3805	Accomodation (subject to PAYE)
Gross wage	3806	Services (subject to PAYE)
Gross wage	3807	Loans or subsidy (subject to PAYE)
Gross wage	3809	Taxable bursaries or scholarships to a non-disabled person basic education (subject to PAYE)
Gross wage	3810	Medical aid contributions (subject to PAYE)
Gross wage	3813	Medical services costs (subject to PAYE)
Gross wage	3815	Non-taxable bursaries or scholarships to non-disabled person basic education
Gross wage	3816	Use of motor vehicle acquired by employers via operating lease (subject to PAYE)
Gross wage	3820	Taxable bursaries or scholarships to a non-disabled person further education (subject to PAYE)
Gross wage	3821	Non-taxable bursaries or scholarships to non-disabled person further education
Gross wage	3822	Non-taxable benefit on acquisition of immovable property
Gross wage	3829	Taxable bursaries or scholarships to a disabled person basic education (subject to PAYE)
Gross wage	3830	Non-taxable bursaries or scholarships to a disabled person basic education
Gross wage	3831	Taxable bursaries or scholarships to a disabled person further education (subject to PAYE)
Gross wage	3832	Non-taxable bursaries or scholarships to a disabled person further education
Gross wage	3851	Foreign general fringe benefits (subject to paye)
Gross wage	3852	Foreign use of motor acquired by employer not via operating lease (subject to paye)
Gross wage	3853	Foreign use of asset (subject to paye)
Gross wage	3854	Foreign meals etc (subject to paye)
Gross wage	3855	Foreign accomodation (subject to paye)
Gross wage	3856	Foreign services (subject to paye)
Gross wage	3857	Foreign loans or subsidy (subject to paye)
Gross wage	3859	Foreign taxable bursaries or scholarships to a non-disabled person basic education (subject to paye)
Gross wage	3860	Foreign medical aid contributions (subject to paye)
Gross wage	3863	Foreign medical services costs (subject to paye)
Gross wage	3865	Foreign non-taxable bursaries or scholarships to non-disabled person basic education
Gross wage	3866	Foreign use of motor vehicle acquired by employers via operating lease (subject to paye)

Table G.8: Source codes categories used in tax microdata

Income concept	Source code	Description
Gross wage	3870	Foreign taxable bursaries or scholarships to a non-disabled person further education (subject to paye)
Gross wage	3871	Foreign non-taxable bursaries or scholarships to non-disabled person further education
Gross wage	3872	Foreign non-taxable benefit on acquisition of immovable property
Gross wage	3879	Foreign taxable bursaries or scholarships to a disabled person basic education (subject to paye)
Gross wage	3880	Foreign non-taxable bursaries or scholarships to a disabled person basic education
Gross wage	3881	Foreign taxable bursaries or scholarships to a disabled person further education (subject to paye)
Gross wage	3882	Foreign non-taxable bursaries or scholarships to a disabled person further education
Gross wage	4236	Remuneration from foreign employer for services rendered in South Africa
Business income	102-4222	Business income (gains and losses)
Pension contributions	4001	Total pension fund contributions paid and deemed paid by employee
Pension contributions	4002	Arrear pension fund contributions paid by employee
Pension contributions	4003	Total provident fund contributions paid and deemed paid by employee
Pension contributions	4004	Arrear provident fund contributions paid by employee
Pension contributions	4006	Total retirement annuity fund contributions paid and deemed paid by employee
Pension contributions	4007	Arrear retirement annuity fund contributions paid by employee
Pension income	3603	Pension (subject to PAYE)
Pension income	3604	Pension (non-taxable)
Pension income	3610	Annuity from a RAF (subject to PAYE)
Pension income	3614	Other retirement lump sums (subject to PAYE)
Pension income	3653	Foreign pension (subject to paye)
Pension income	3654	Foreign pension (non-taxable)
Pension income	3660	Foreign annuity from a raf (subject to paye)
Pension income	3664	Foreign other retirement lump sums (subject to paye)
Pension income	3902	Pension or RAF in respect of withdrawal (subject to PAYE)
Pension income	3903	Pension or RAF in respect of retirement (subject to PAYE)
Pension income	3904	Provident in respect of withdrawal (subject to PAYE)
Pension income	3905	Provident in respect of retirement (subject to PAYE)
Pension income	3908	Surplus apportionments and exempt policy proceeds (non-taxable)
Pension income	3909	Unclaimed benefits
Pension income	3915	Retirement or termination of employment lump sum benefits or commutation of annuities
Pension income	3920	Lump sum withdrawal benefits (subject to PAYE)
Pension income	3921	Living annuity and section 15C of the pension funds act, surplus apportionments (subject to PAYE)

Table G.8: Source codes categories used in tax microdata

Income concept	Source code	Description
Pension income	3923	Transfer of unclaimed benefits
Pension income	3924	Transfer on retirement (subject to PAYE)
Pension income	3952	Foreign pension or raf in respect of withdrawal (subject to paye)
Pension income	3953	Foreign pension or raf in respect of retirement (subject to paye)
Pension income	3954	Foreign provident in respect of withdrawal (subject to paye)
Pension income	3955	Foreign provident in respect of retirement (subject to paye)
Interest income	4201	Local interest excluding SARS
Interest income	4218	Foreign interest
Interest income	4237	SARS interest received
Interest income	4241	Tax free investment account interest
Rental income	2532	Business income component: property letting income, residential accomodation
Rental income	2533	Business income component: property letting loss, residential accomodation
Rental income	4210	Local rental from letting of fixed property
Rental income	4288	Foreign rental gain
Dividends	3717	Broad-based employee share plan (subject to PAYE)
Dividends	3718	Vesting of equity instruments or return of capital iro restricted instruments (PAYE)
Dividends	3719	Dividends not exempt ito para (dd) of the proviso to s10(1)(k)(i) (PAYE)
Dividends	3720	Dividends not exempt ito para (ii) of the proviso to s10(1)(k)(i) (PAYE)
Dividends	3721	Dividends not exempt ito para (jj) of the proviso to s10(1)(k)(i) (PAYE)
Dividends	3723	Dividends not exempt ito para (kk) of the proviso to s10(1)(k)(i) (PAYE)
Dividends	3767	Foreign broad-based employee share plan (subject to paye)
Dividends	3768	Foreign vesting of equity instruments or return of capital iro restricted instruments (paye)
Dividends	3769	Foreign dividends not exempt ito para (dd) of the proviso to s10(1)(k)(i) (paye)
Dividends	3770	Foreign dividends not exempt ito para (ii) of the proviso to s10(1)(k)(i) (paye)
Dividends	3771	Foreign dividends not exempt ito para (jj) of the proviso to s10(1)(k)(i) (paye)
Dividends	3773	Foreign dividends not exempt ito para (kk) of the proviso to s10(1)(k)(i) (paye)
Dividends	4216	Foreign dividends
Dividends	4230	Controlled foreign company share of profit
Dividends	4238	Taxable local dividends ie REIT
Dividends	4242	Tax free investment account dividends
Dividends	4257	Tax free investments other
Dividends	4292	Dividends deemed to be income in terms of s8E and s8EA

Table G.8: Source codes categories used in tax microdata

Income concept	Source code	Description
Not used	3618	Misclassification or undefined
Not used	3695	Misclassification or undefined
Not used	3696	Gross non-taxable income
Not used	3697	Gross retirement funding employment income
Not used	3698	Gross non-retirement funding employment income
Not used	3699	Gross employment income taxable
Not used	3808	Employee's debt (subject to PAYE)
Not used	3817	Benefit employer pension fund contributions (subject to PAYE)
Not used	3818	Misclassification or undefined
Not used	3819	Misclassification or undefined
Not used	3825	Benefit employer provident fund contributions (subject to PAYE)
Not used	3826	Misclassification or undefined
Not used	3827	Misclassification or undefined
Not used	3828	Benefit retirement annuity fund contributions (subject to PAYE)
Not used	3858	Foreign employee's debt (subject to paye)
Not used	3867	Foreign benefit employer pension fund contributions (subject to paye)
Not used	3875	Foreign benefit employer provident fund contributions (subject to paye)
Not used	3876	Misclassification or undefined
Not used	3877	Misclassification or undefined
Not used	3878	Foreign benefit retirement annuity fund contributions (subject to paye)
Not used	3901	Gratuities and severance benefits (subject to PAYE)
Not used	3906	Special remuneration (subject to PAYE)
Not used	3907	Other lump sums (subject to PAYE)
Not used	3922	Compensation iro of death during employment (non-taxable)
Not used	3951	Foreign gratuities and severance benefits (subject to paye)
Not used	3956	Foreign special remuneration (subject to paye)
Not used	3957	Foreign other lump sums (subject to paye)
Not used	4005	Medical scheme fees paid and deemed paid by employee
Not used	4008	Misclassification or undefined
Not used	4009	Misclassification or undefined
Not used	4011	Donations allowable in terms of section 18a to an approved public benefit organisation
Not used	4014	Misclassification or undefined

Table G.8: Source codes categories used in tax microdata

Income concept	Source code	Description
Not used	4015	Travel expenses (no allowance, commission income)
Not used	4016	Other deductions
Not used	4017	Expenses against local taxable subsistence allowance
Not used	4018	Premiums paid for loss of income policies
Not used	4019	Expenses against foreign taxable subsistence allowance
Not used	4024	Medical services costs deemed to be paid by the employee
Not used	4025	Medical contribution paid by employee allowed as a deduction for employees tax purposes
Not used	4026	Arrear pension fund contributions non-statutory forces
Not used	4027	Depreciation
Not used	4028	Home office expenses
Not used	4029	Retirement fund contributions total
Not used	4030	Donations deducted from the employee remuneration and paid by employer to organisation
Not used	4031	Section 8C losses
Not used	4032	Remuneration (s8A/8C gains) taxed on IRP5 but comply with exemption in terms of s10(i)(o)(ii)
Not used	4033	Remuneration taxed on IRP5 but comply with exemption in terms of s10(1)(o)(i)
Not used	4041	Remuneration taxed on IRP5 but comply with exemption in terms of s10(1)(o)(ii) (excluding s 8A/8C gains)
Not used	4042	Amounts refunded ito section 11(nA) and 11(nB)
Not used	4043	Allowable accountancy or administration expenses
Not used	4044	Legal expenses
Not used	4045	Bad debt
Not used	4046	Use of motor vehicle
Not used	4047	Holders of public office deduction
Not used	4048	Misclassification or undefined
Not used	4050	Misclassification or undefined
Not used	4051	Misclassification or undefined
Not used	4101	SITE
Not used	4102	PAYE
Not used	4103	Misclassification or undefined
Not used	4104	Misclassification or undefined
Not used	4110	Misclassification or undefined
Not used	4111	Other foreign tax credits individuals
Not used	4112	Foreign tax credits on such foreign dividends

Table G.8: Source codes categories used in tax microdata

Income concept	Source code	Description
Not used	4113	Foreign tax credits on foreign interest
Not used	4114	Foreign tax credits in respect of foreign capital gain or loss
Not used	4115	Tax on retirement lump sum and severance benefits
Not used	4116	Medical scheme fees tax credit
Not used	4117	Foreign tax credits in respect of S6quin
Not used	4118	Sum of ETI amounts
Not used	4120	Additional medical expenses tax credit
Not used	4121	Foreign tax credits on foreign rental income
Not used	4141	UIF contribution
Not used	4142	SDL contribution
Not used	4149	Total tax
Not used	4150	Metadata
Not used	4211	Local rental loss from letting of fixed property
Not used	4212	Royalties
Not used	4213	Loss royalties
Not used	4214	Other receipts and accruals
Not used	4215	Misclassification or undefined
Not used	4219	Tax free investment account contribution
Not used	4220	Misclassification or undefined
Not used	4221	Misclassification or undefined
Not used	4223	Loss foreign business or trading
Not used	4228	Other foreign income
Not used	4229	Loss other foreign income
Not used	4235	Income reflected on a South African IRP5 or IT3a that was subject to tax outside SA
Not used	4239	Tax free investment account net return on investment profit
Not used	4240	Tax free investment account net return on investment loss
Not used	4243	Tax free investment account capital gain
Not used	4244	Tax free investment account capital loss
Not used	4245	Misclassification or undefined
Not used	4246	Tax free investment account transfer in
Not used	4247	Tax free investment account transfer out
Not used	4248	Tax free investment account withdrawal

Table G.8: Source codes categories used in tax microdata

Income concept	Source code	Description
Not used	4249	Foreign tax credits refunded or discharged in terms of S6quat(1C)
Not used	4250	Local capital gain
Not used	4251	Loss local capital
Not used	4252	Foreign capital gain
Not used	4253	Loss foreign capital
Not used	4278	Foreign royalties
Not used	4279	Loss foreign royalties
Not used	4280	Sporting
Not used	4281	Loss sporting
Not used	4282	Collectables
Not used	4283	Loss collectables
Not used	4284	Animal showing
Not used	4285	Loss animal showing
Not used	4286	Gambling
Not used	4287	Loss gambling
Not used	4289	Foreign rental loss
Not used	4291	Foreign income in terms of s6quat(1C)
Not used	4301	Misclassification or undefined
Not used	4302	Misclassification or undefined
Not used	4472	Employer pension fund contributions paid for the benefit of employee
Not used	4473	Employer provident fund contributions paid for the benefit of employee
Not used	4474	Employer medical scheme fees paid for the benefit of employee
Not used	4475	Employer retirement annuity fund contributions paid for the benefit of employee
Not used	4476	Misclassification or undefined
Not used	4485	Medical services costs deemed to be paid by the employee for other relatives
Not used	4486	Capped amount determined by employer in terms of section 18(2)(c)(i)
Not used	4487	No value benefits in respect of medical services provided or incurred by the employer
Not used	4493	Employer's medical scheme fees paid for the benefit of a retired/former of the Seventh Schedule
Not used	4497	Total deductions and contributions
Not used	4582	The portion of the allowances and benefits which represents remuneration
Not used	4583	The portion of other allowances and benefits which represents remuneration

Table G.8: Source codes categories used in tax microdata

Income concept	Source code	Description

Source. Authors' elaboration. The tax microdata used in this paper refers to the "Individual Panel" dataset (see Ebrahim and Axelson, 2019). The data was accessed from August 2019 to March 2020. The version of the dataset used in this paper is 2019_1. The table below shows all the source codes used, along with the corresponding income category attributed to each source code.

Appendix H

Appendix to “Income Inequality in Africa, 1990-2019: Measurement, Patterns, Determinants”

Figure H.1: Evolution of the Pan-African Income Distribution
(Survey-Based versus DINA Estimates)

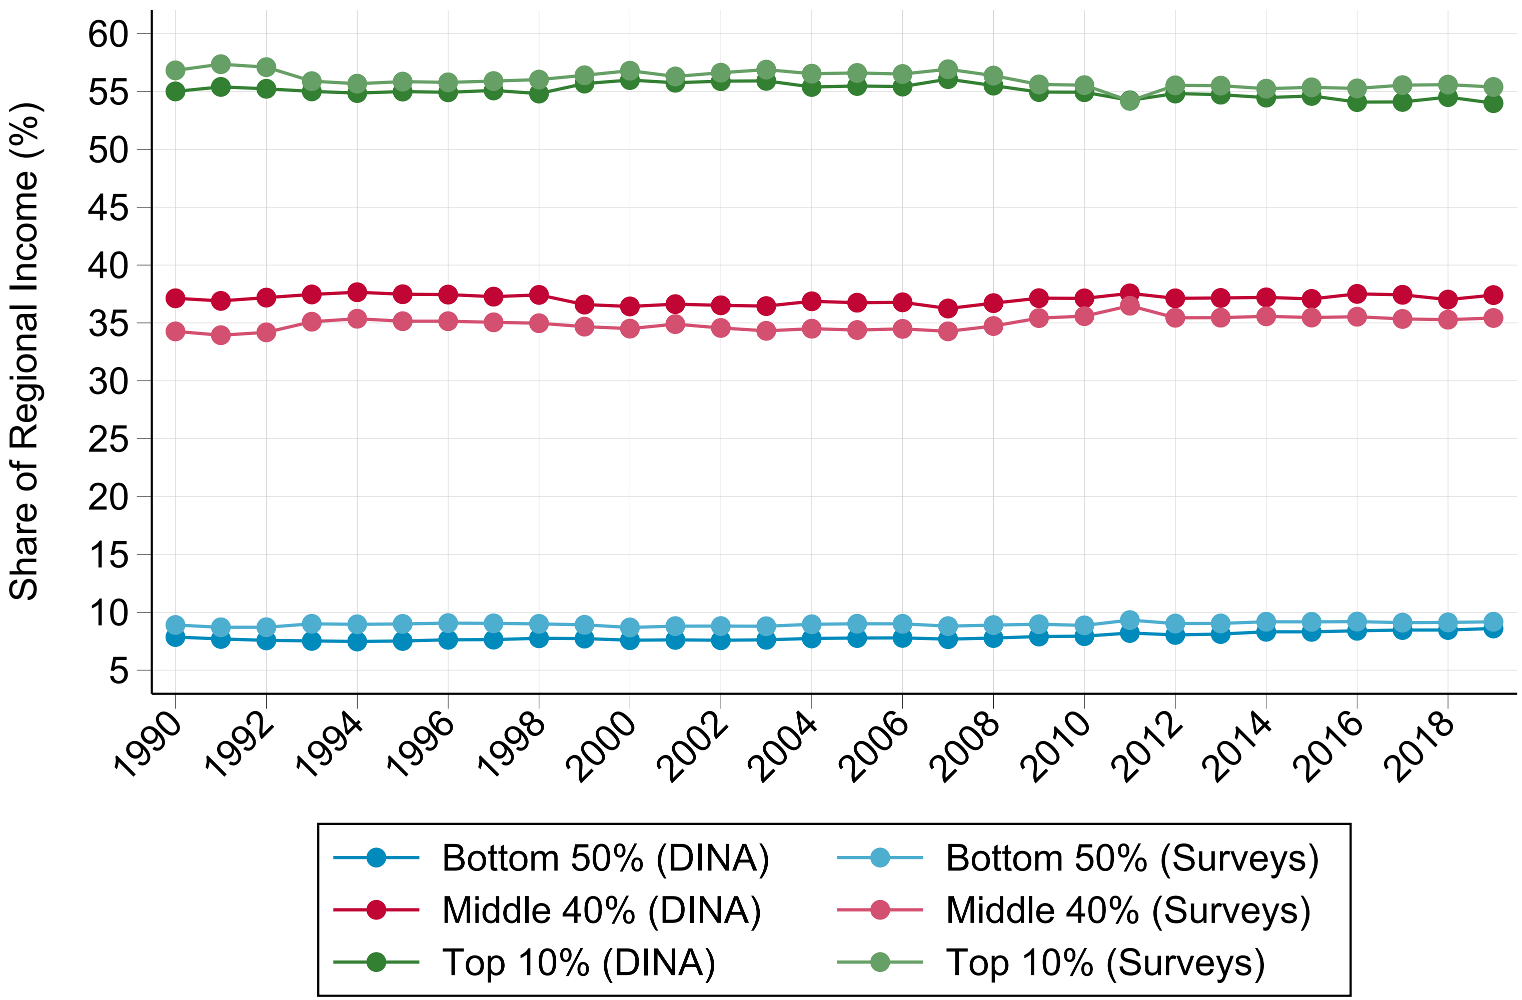

Notes. Authors' computations combining survey, tax, and national accounts data. The figure compares DINA estimates, rescaling each distribution to net national income, to survey-based estimates, which rely on survey estimates of average income.

Figure H.2: Evolution of the Pan-African Income Distribution
(Pretax versus Posttax)

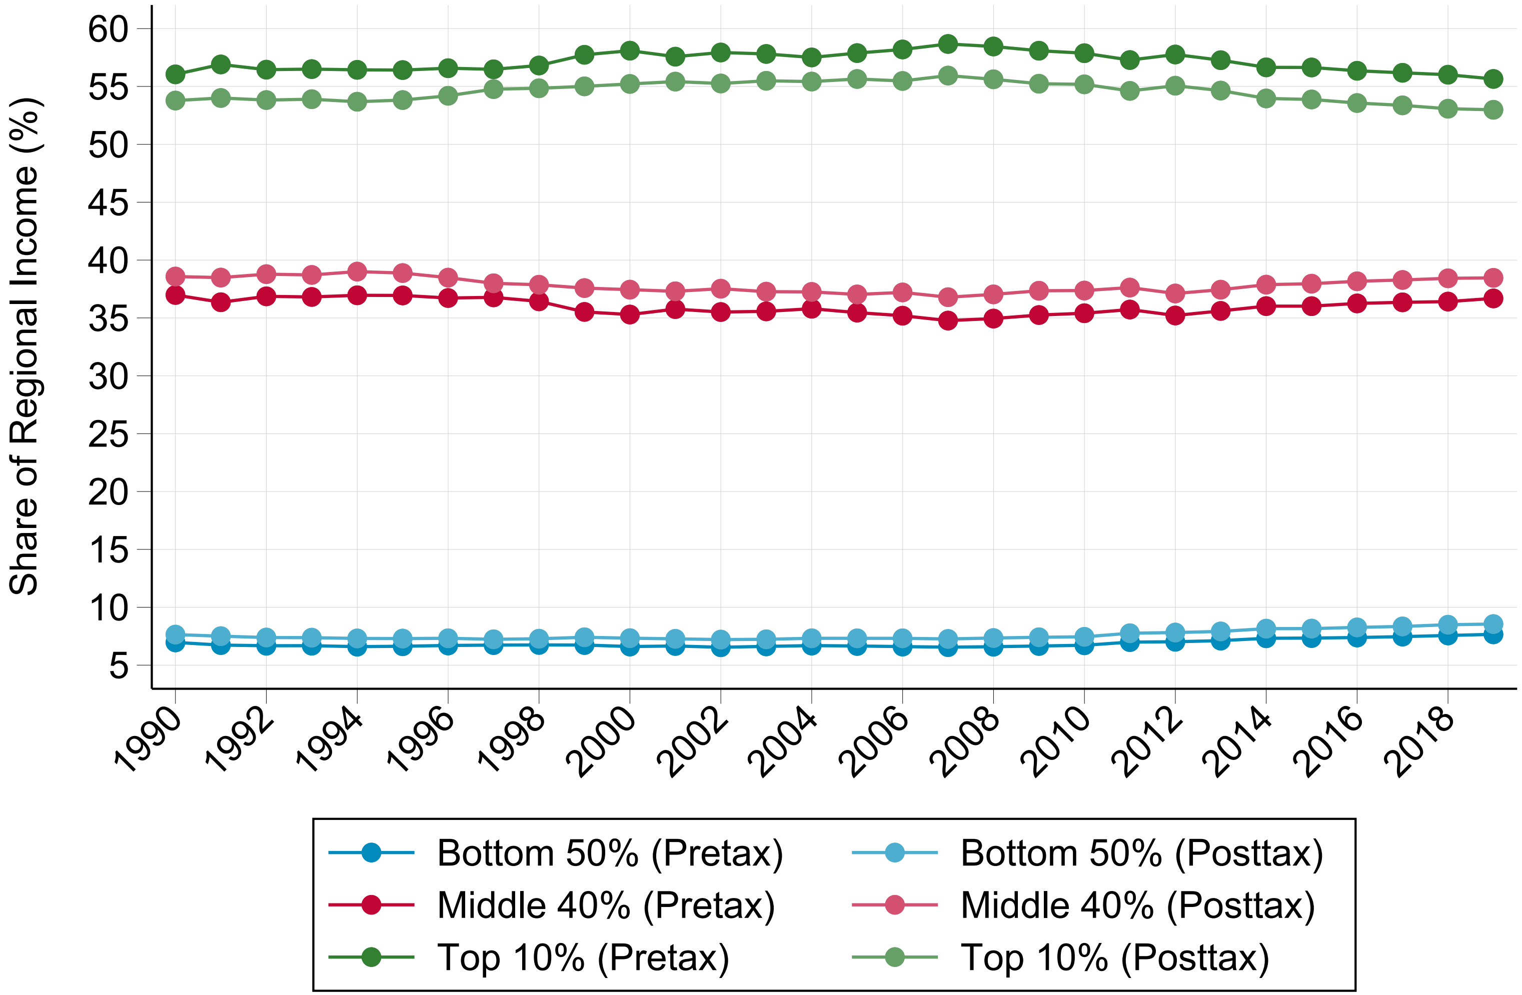

Notes. Authors' computations combining survey, tax, and national accounts data (pretax); preliminary estimates from Gethin (2023b) (posttax). The figure compares the evolution of top 10%, middle 40%, and bottom 50% shares in Africa as a whole in terms of pretax income and posttax disposable income (pretax income, minus direct taxes, plus social assistance transfers).

Table H.1: Logistic Fit of
Income-Consumption Profiles

Survey	$\hat{\alpha}$	$\hat{\beta}$	Adj. R^2
African countries			
Cote d'Ivoire, 1998	0.85	0.10	0.96
Cote d'Ivoire, 2002	0.81	0.13	0.99
Cote d'Ivoire, 2008	0.84	0.11	0.96
Cote d'Ivoire, 2015	0.85	0.12	0.99
Ghana, 1988	0.87	0.13	0.99
Ghana, 1998	0.81	0.14	0.96
Guinea, 1994	0.85	0.07	0.74
Madagascar, 1993	0.91	0.06	0.94
Uganda, 1992	0.92	0.04	0.72
Other countries			
India, 2005	0.82	0.14	0.99
India, 2011	0.86	0.14	0.98
Thailand, 2000	0.78	0.16	0.97
Thailand, 2001	0.77	0.16	0.98
Thailand, 2002	0.82	0.14	0.98
Thailand, 2004	0.86	0.11	0.96
Thailand, 2006	0.82	0.13	0.95
Thailand, 2007	0.83	0.13	0.95
Thailand, 2009	0.84	0.12	0.96
Thailand, 2011	0.84	0.12	0.93
Thailand, 2013	0.87	0.11	0.91
Thailand, 2015	0.89	0.10	0.93

Source: authors' computations. *Interpretation*: the best logistic fit for the ratio of consumption to income by percentile in 1998 Cote d'Ivoire yields a functional form of $Q(p) = 0.86 + 0.11 \log \frac{p}{1-p}$, with an adjusted R-squared of 97%.

Table H.2: Top 10% and Bottom 50% Income Shares Before and After Correction, 2019

Country	Top 10%				Bottom 50%			
	Original Survey	Corrected Survey	Lower Bound	Upper Bound	Original Survey	Corrected Survey	Lower Bound	Upper Bound
Algeria	22.8%	37.3%	33.8%	40.6%	31.3%	20.7%	22.7%	18.9%
Angola	39.6%	57.7%	53.7%	61.2%	17.0%	9.5%	10.8%	8.4%
Benin	37.6%	54.7%	50.9%	58.2%	19.3%	11.5%	12.9%	10.2%
Botswana	41.5%	58.9%	55.1%	62.2%	15.8%	8.7%	9.9%	7.6%
Burkina Faso	29.6%	46.4%	42.5%	50.0%	27.0%	16.5%	18.5%	14.8%
Burundi	31.0%	47.8%	43.9%	51.3%	24.8%	15.1%	16.9%	13.5%
Cabo Verde	32.3%	48.6%	44.9%	52.0%	21.9%	13.3%	14.8%	11.9%
Cameroon	35.0%	51.7%	47.9%	55.1%	19.1%	11.3%	12.7%	10.1%
Central African Republic	46.2%	64.6%	60.8%	68.0%	15.3%	8.0%	9.3%	7.0%
Chad	32.4%	48.9%	45.1%	52.4%	21.3%	13.0%	14.5%	11.6%
Comoros	33.7%	49.9%	46.2%	53.3%	19.8%	12.0%	13.4%	10.7%
Congo	37.9%	55.6%	51.7%	59.1%	18.3%	10.5%	11.9%	9.3%
Cote d'Ivoire	36.1%	53.5%	49.6%	57.0%	20.1%	11.7%	13.2%	10.4%
DR Congo	32.0%	48.4%	44.6%	51.8%	22.1%	13.5%	15.0%	12.1%
Djibouti	32.3%	49.1%	45.3%	52.7%	22.7%	13.8%	15.4%	12.3%
Egypt	26.9%	43.4%	39.5%	47.0%	29.4%	18.5%	20.5%	16.6%
Equatorial Guinea	34.4%	51.2%	47.4%	54.7%	20.4%	12.2%	13.6%	10.8%
Eritrea	28.5%	44.9%	41.0%	48.4%	27.1%	17.0%	18.9%	15.3%
Eswatini	42.7%	59.5%	55.9%	62.8%	15.2%	8.4%	9.6%	7.4%

Ethiopia	28.5%	44.9%	41.0%	48.4%	27.1%	17.0%	18.9%	15.3%
Gabon	27.7%	42.8%	39.3%	46.2%	24.1%	15.4%	17.0%	13.9%
Gambia	28.7%	45.2%	41.4%	48.8%	26.2%	16.2%	18.1%	14.6%
Ghana	32.2%	48.6%	44.8%	52.0%	21.0%	12.8%	14.3%	11.5%
Guinea	26.4%	42.1%	38.4%	45.5%	27.3%	17.4%	19.3%	15.8%
Guinea-Bissau	42.0%	59.7%	55.9%	63.1%	18.4%	10.2%	11.6%	9.0%
Kenya	31.6%	48.2%	44.4%	51.7%	23.1%	14.0%	15.6%	12.5%
Lesotho	32.9%	49.1%	45.4%	52.5%	19.8%	12.0%	13.4%	10.7%
Liberia	27.1%	42.6%	39.0%	46.0%	26.2%	16.6%	18.4%	15.0%
Libya	27.3%	43.4%	39.7%	47.0%	28.1%	17.8%	19.7%	16.0%
Madagascar	33.5%	50.3%	46.4%	53.7%	22.2%	13.3%	15.0%	11.9%
Malawi	36.5%	55.8%	51.5%	59.6%	22.9%	12.8%	14.7%	11.3%
Mali	25.7%	40.6%	37.0%	43.9%	27.5%	17.7%	19.6%	16.1%
Mauritania	24.9%	39.9%	36.3%	43.2%	27.7%	18.0%	19.9%	16.4%
Mauritius	29.9%	46.7%	42.8%	50.3%	26.0%	16.0%	17.8%	14.3%
Morocco	31.9%	48.8%	44.9%	52.4%	24.2%	14.6%	16.4%	13.1%
Mozambique	45.5%	64.2%	60.3%	67.6%	17.0%	8.9%	10.3%	7.7%
Namibia	47.2%	64.0%	60.4%	67.1%	12.8%	6.9%	7.9%	6.0%
Niger	27.0%	42.6%	38.9%	46.0%	26.9%	17.1%	18.9%	15.4%
Nigeria	26.7%	42.1%	38.5%	45.5%	26.2%	16.7%	18.4%	15.1%
Rwanda	35.6%	53.4%	49.4%	56.9%	22.1%	12.8%	14.5%	11.3%
Sao Tome and Principe	24.2%	38.7%	35.2%	41.9%	29.0%	19.0%	20.9%	17.3%
Senegal	31.0%	47.2%	43.5%	50.6%	23.3%	14.3%	16.0%	12.9%
Seychelles	33.7%	51.6%	47.5%	55.2%	22.2%	13.0%	14.7%	11.6%

Sierra Leone	29.4%	46.2%	42.4%	49.8%	26.7%	16.4%	18.3%	14.7%
Somalia	27.9%	43.5%	39.9%	47.0%	25.2%	16.0%	17.7%	14.4%
South Africa	50.5%	65.1%	65.1%	65.1%	10.7%	6.3%	5.3%	7.2%
South Sudan	33.2%	49.6%	45.9%	53.1%	20.8%	12.6%	14.0%	11.2%
Sudan	27.8%	44.3%	40.5%	47.9%	27.4%	17.1%	19.0%	15.4%
Tanzania	33.1%	50.7%	46.7%	54.3%	23.9%	14.1%	15.9%	12.5%
Togo	31.6%	47.6%	43.9%	51.0%	21.1%	12.9%	14.4%	11.6%
Tunisia	25.6%	40.7%	37.1%	44.1%	27.8%	17.9%	19.8%	16.2%
Uganda	34.2%	51.5%	47.6%	55.0%	22.3%	13.1%	14.8%	11.7%
Zambia	44.4%	61.5%	57.8%	64.7%	13.4%	7.3%	8.4%	6.4%
Zimbabwe	40.8%	58.5%	54.6%	61.9%	18.0%	10.0%	11.4%	8.7%
Africa	41.0%	54.4%	51.6%	57.0%	13.5%	8.7%	9.5%	7.9%
Eastern Africa	37.9%	53.5%	49.9%	56.7%	18.7%	11.2%	12.6%	10.0%
Middle Africa	48.6%	60.9%	58.1%	63.6%	12.1%	7.6%	8.4%	6.8%
Northern Africa	28.9%	44.5%	40.8%	47.9%	24.4%	15.5%	17.2%	14.0%
Southern Africa	50.4%	65.0%	64.7%	65.3%	10.9%	6.4%	5.6%	7.1%
Subsaharan Africa	42.1%	55.6%	52.9%	58.2%	14.8%	9.3%	10.2%	8.5%
Western Africa	31.4%	46.3%	42.9%	49.6%	20.4%	13.0%	14.4%	11.7%

Table H.3: Data Sources

Country	Distributional data	National accounts data	Method
Angola	HH consumption surveys: 1995, 2000, 2008, 2018	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Burkina Faso	HH consumption surveys: 1994, 1998, 2003, 2009, 2014	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Burundi	HH consumption surveys: 1992, 1998, 2006, 2013	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Benin	HH consumption surveys: 2003, 2011, 2015	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Botswana	HH consumption surveys: 1985, 1993, 2002, 2009, 2015	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
DR Congo	HH consumption surveys: 2004, 2005, 2008, 2012	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts

Central African Republic	HH consumption surveys: 1992, 2003, 2008	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Congo	HH consumption surveys: 2005, 2011	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Cote d'Ivoire	HH consumption surveys: 1985, 1986, 1987, 1988, 1992, 1993, 1995, 1998, 2002, 2008, 2015; Tax data:2014	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using tax data and national accounts
Cameroon	HH consumption surveys: 1996, 2001, 2007, 2014	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Cabo Verde	HH consumption surveys: 2001, 2007, 2015	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Djibouti	HH consumption surveys: 2002, 2012, 2013, 2017	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Algeria	HH consumption surveys: 1988, 1995, 2011	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts

Egypt	HH consumption surveys: 1990, 1995, 1999, 2004, 2008, 2010, 2012, 2015, 2017	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Ethiopia	HH consumption surveys: 1981, 1995, 1999, 2004, 2005, 2010, 2015	IMF levels (2019) and World Bank growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Gabon	HH consumption surveys: 2005, 2017	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Ghana	HH consumption surveys: 1987, 1991, 2005, 2012, 2016	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Gambia	HH consumption surveys: 1992, 1998, 2003, 2010, 2015	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Guinea	HH consumption surveys: 1991, 2002, 2003, 2007, 2012	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Guinea-Bissau	HH consumption surveys: 1991, 1993, 2002, 2010	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts

Kenya	HH consumption surveys: 1992, 1994, 1997, 2005, 2015	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Comoros	HH consumption surveys: 1995, 2004, 2014	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Liberia	HH consumption surveys: 2007, 2014, 2016	World Bank levels (2019) and UN SNA growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Lesotho	HH consumption surveys: 1986, 1993, 1994, 2002, 2010, 2017	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Morocco	HH consumption surveys: 1984, 2000, 2006, 2013	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Madagascar	HH consumption surveys: 1980, 1997, 1999, 2001, 2005, 2010, 2012	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Mali	HH consumption surveys: 1994, 2001, 2006, 2009	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Mauritania	HH consumption surveys: 1987, 1993, 1995, 2000, 2004, 2008, 2014	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts

Mauritius	HH consumption surveys: 2006, 2012, 2017	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Malawi	HH consumption surveys: 1997, 2004, 2010, 2016	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Mozambique	HH consumption surveys: 1996, 2002, 2007, 2008, 2014	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Namibia	HH consumption surveys: 2003, 2009, 2015	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Niger	HH consumption surveys: 1992, 1994, 2005, 2007, 2011, 2014	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Nigeria	HH consumption surveys: 1985, 1992, 2003, 2009, 2018	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Rwanda	HH consumption surveys: 1984, 2000, 2005, 2010, 2013, 2016	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Seychelles	HH consumption surveys: 1999, 2006	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts

Sudan	HH consumption surveys: 2009, 2014	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Sierra Leone	HH consumption surveys: 1989, 2003, 2011, 2018	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Senegal	HH consumption surveys: 1991, 1994, 2001, 2005, 2011	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Somalia	HH consumption surveys: 2017	UN SNA levels (2019) and World Bank growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
South Sudan	HH consumption surveys: 2009, 2016	UN SNA levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Sao Tome and Principe	HH consumption surveys: 2000, 2010	World Bank levels (2019) and UN SNA growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Eswatini	HH consumption surveys: 1994, 2000, 2009, 2016	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Chad	HH consumption surveys: 2002, 2003, 2011	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts

Togo	HH consumption surveys: 2006, 2011, 2015	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Tunisia	HH consumption surveys: 1985, 1990, 1995, 2000, 2005, 2010, 2015	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Tanzania	HH consumption surveys: 1991, 2000, 2007, 2011, 2014, 2017	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
Uganda	HH consumption surveys: 1989, 1996, 1999, 2002, 2005, 2009, 2012, 2016	UN SNA levels (2019) and World Bank growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts
South Africa	HH consumption surveys: 1993, 1996, 2000, 2005, 2008, 2010, 2014; Tax data:1990-1993, 2002-2012	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using tax data and national accounts
Zambia	HH consumption surveys: 1991, 1993, 1996, 1998, 2002, 2004, 2006, 2010, 2015	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts

Zimbabwe	HH consumption surveys: 1991, 1996, 2011, 2017, 2019	World Bank levels (2019) and growth rates (1950–2021).	Correction of surveys using stylized correction profile (see section 3.2 and 3.3) and national accounts

Table H.4: Average Incomes: Surveys versus National Accounts

Country	Year	Survey Mean	NNI Per Capita	Ratio of Survey Mean to NNI Per Capita
Angola	2018	2366	6139	0.39
Benin	2018	1839	3146	0.58
Botswana	2015	3835	12248	0.31
Burkina Faso	2018	1854	2011	0.92
Burundi	2013	977	800	1.22
Cabo Verde	2015	3699	5780	0.64
Cameroon	2014	2246	3418	0.66
Central African Republic	2008	1248	1071	1.16
Chad	2018	1441	1669	0.86
Comoros	2014	2822	2942	0.96
Congo	2011	1373	4047	0.34
Cote d'Ivoire	2018	2148	5043	0.43
DR Congo	2012	833	917	0.91
Djibouti	2017	2109	4473	0.47
Egypt	2017	4058	10944	0.37
Ethiopia	2015	1500	1949	0.77
Gabon	2017	5518	13854	0.40
Gambia	2015	1955	1908	1.02
Ghana	2016	2132	4735	0.45
Guinea	2018	1986	2006	0.99

Guinea-Bissau	2018	1641	1824	0.90
Kenya	2015	1783	3808	0.47
Lesotho	2017	1743	3262	0.53
Liberia	2016	376	1402	0.27
Madagascar	2012	610	1481	0.41
Malawi	2019	863	1458	0.59
Mali	2018	1829	2259	0.81
Mauritania	2014	2313	4560	0.51
Mauritius	2017	6904	25687	0.27
Morocco	2013	4437	7007	0.63
Mozambique	2014	1034	1165	0.89
Namibia	2015	4788	10178	0.47
Niger	2018	1053	1204	0.87
Nigeria	2018	1676	5170	0.32
Rwanda	2016	1155	1903	0.61
Sao Tome and Principe	2010	1657	3449	0.48
Senegal	2018	2310	3050	0.76
Seychelles	2018	10120	25098	0.40
Sierra Leone	2018	1470	1695	0.87
Somalia	2017	1530	1070	1.43
South Africa	2014	5382	13590	0.40
South Sudan	2016	101	675	0.15
Sudan	2014	2744	4605	0.60
Swaziland	2016	2122	7709	0.28

Tanzania	2017	1137	2273	0.50
Togo	2018	1826	2088	0.87
Tunisia	2015	5160	9832	0.52
Uganda	2019	1342	2216	0.61
Zambia	2015	1184	3155	0.38
Zimbabwe	2019	2913	3364	0.87

Table H.5: European settlement and Islam correlates versus regional differences. All Africa

	Top 10% income share			Bottom 50% income share		
	A	B	C	A	B	C
European settlement	+0.052***		+0.063**	-0.020**		-0.025**
	(0.018)		(0.021)	(0.009)		(0.010)
Muslims share	-0.099***		-0.059**	+0.052***		+0.030**
	(0.019)		(0.026)	(0.009)		(0.012)
Northern		-0.099***	-0.091**		+0.056***	+0.048**
		(0.028)	(0.035)		(0.012)	(0.016)
North-Eastern		-0.066**	-0.027		+0.033***	+0.015
		(0.026)	(0.026)		(0.011)	(0.012)
Western		-0.056***	-0.022		+0.025***	+0.009
		(0.018)	(0.020)		(0.008)	(0.009)
Southern		+0.065**	+0.024		-0.035***	-0.018
		(0.026)	(0.026)		(0.011)	(0.011)
Small islands		-0.048	-0.070**		+0.024*	+0.034***
		(0.029)	(0.027)		(0.013)	(0.012)
F-test regional variables (p-value)		0.000	0.081		0.000	0.012
N	54	54	54	54	54	54
Adj. R^2	0.397	0.395	0.518	0.41	0.486	0.587

Source: authors' computations. Standard errors in parentheses; *: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$.
European settlement and Muslims share: see Table 8.1 and text.
Northern: Algeria, Egypt, Libya, Morocco, Tunisia. North-Eastern: Djibouti, Eritrea, Ethiopia, Somalia, Sudan, South Sudan. Western: Benin, Burkina Faso, Cabo Verde, Côte d'Ivoire, Gambia, Ghana, Guinea, Guinea Bissau, Liberia, Mali, Mauritania, Niger, Nigeria, Senegal, Sierra Leone, Togo. Eastern (omitted): Burundi, Comoros, Kenya, Madagascar, Mauritius, Mozambique, Malawi, Rwanda, Seychelles, Tanzania, Uganda, Zambia.
Southern: Botswana, Eswatini, Lesotho, Namibia, South Africa, Zimbabwe. Small islands: Islands that were uninhabited before slave trade and colonization: C. Verde, Mauritius, São Tome & P., Seychelles.
F-test for regional variables does not include the small islands dummy.

Table H.6: European settlement and Islam correlates versus geography, precolonial history, and colonizers' identity. All Africa

	Top 10% income share					Bottom 50% income share				
	A	B	C	D	E	A	B	C	D	E
European settlement	+0.053***	+0.057**	+0.059***	+0.042**	+0.041	-0.020**	-0.023*	-0.023**	-0.016**	-0.016
	(0.017)	(0.026)	(0.019)	(0.017)	(0.028)	(0.008)	(0.012)	(0.009)	(0.008)	(0.012)
Muslims share	-0.112***	-0.116***	-0.106***	-0.113***	-0.111***	+0.058***	+0.062***	+0.054***	+0.63***	+0.064***
	(0.018)	(0.024)	(0.024)	(0.019)	(0.031)	(0.009)	(0.011)	(0.011)	(0.009)	(0.014)
Controls:		p-value	p-value	p-value	p-value	p-value	p-value	p-value	p-value	p-value
Geography		0.724			0.686		0.487			0.615
Slave exports			0.119		0.066			0.096		0.039
Precolonial pol.			0.668		0.722			0.685		0.776
Ethnic fract.			0.268		0.210			0.302		0.253
Colonizer ident.				0.093	0.170				0.035	0.073
N	54	54	54	54	54	54	54	54	54	54
Adj. R^2	0.471	0.444	0.450	0.520	0.483	0.501	0.496	0.483	0.572	0.557

Source: authors' computations. Standard errors in parentheses; *: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$. European settlement: Dummy for whether European settlers went above 2.5% of total population between 1870 and 1970 (Easterly and Levine, 2016). Eur. settlement: Algeria, Angola, Eswatini, Libya, Morocco, Mozambique, Mauritius, Namibia, South Africa, Tunisia, Zambia, Zimbabwe. Muslim share: proportion of Muslims in total population circa 2010. Muslims > 50%: Algeria, B. Faso, Chad, Comoros, Djib., Egypt, Guinea, G. Bissau, Libya, Morocco, Mali, Mauritania, Niger, Sudan, Senegal, S. Leone, Somalia, Tunisia. Geography: Abs. latitude, longitude, min month. avg rainfall, max month. afternoon avg humidity, min avg month. low temp, log(coastline/area). (Nunn, 2008). Slave exports: Log total slave exports normalized by historic population (Nunn, 2008); results are similar with slave exports normalized by land area. Precolonial polities: Percentages of population from Centralized Stratified, Centr. Egalitarian, and Fragmented Strat.groups; Frag. and Egal. being omitted (Gennaioli and Rainer, 2007). The variables were constructed using the dataset from Michalopoulos and Papaioannou (2013). Ethnic fractionalization: Alesina et al. (2003). Sao Tome and Principe was set at the value for Cabo Verde. Colonizer identity: Dummy variables for the last colonizer being either Belgian, British, French, or Portuguese (Somalia has 0.5 for British as it was shared with Italy), and for non-colonized (Ethiopia and Liberia). In all regressions, a "small island" dummy is included: Cabo Verde, Mauritius, São Tome & P., Seychelles. These islands were uninhabited before slavery and colonization. The precolonial dummies were set at zero (meaning 100% was fragmented and egalitarian); given the small island dummy, this has no impact on reported point estimates.

Appendix I

Appendix to "Brahmin Left Versus Merchant Right: Changing Political Cleavages in 21 Western Democracies, 1948–2020"

I.1 Estimation of quantile groups from discrete categories

One of the contribution of this paper is to provide data on the vote share received by specific parties and coalitions by income and education groups, decomposing for instance the population into its poorest or least educated half (the bottom 50%), the next 40% (the middle 40%), and the highest decile (the top 10%). Such groups are key to track political cleavages over time and compare them across countries. The problem is that existing surveys do not provide continuous values for income or education: these variables are most often coded in discrete categories (educational levels in the case of education, income brackets in the case of income).

To partially overcome this issue, we introduce a simple reweighing method, which exploits the distribution of individuals in each bracket or category to approximate quantiles. Consider for example the 2015 Canadian Election Study, which contains an income variable coded in eighteen brackets (see table 1). One is interested in computing the proportion of individuals belonging to the lowest income decile voting for the New Democratic Party $\bar{y}_{d=1}$, where y is a binary variable taking 1 is the respondent voted for the NDP and 0 otherwise, and where d refers to the income decile to which the respondents belong. Unfortunately, this is not directly possible with this income variable since only 5% of individuals belong to the first income bracket ($b = 1$), and 15.5% of them belong to the lowest two brackets ($b \in [1, 2]$). If support for the NDP decreases linearly with income, then $\bar{y}_{b=1}$ will strongly overestimate $\bar{y}_{d=1}$, while $\bar{y}_{b=2}$ will strongly underestimate it since we are looking at individuals who are on average too poor in the first case and too rich in the second. However, it is easy to see that since individuals within the second bracket range from quantiles 0.05 to 0.155, this means that $0.05/(0.155 - 0.05) \approx 48\%$ of them belong to the bottom 10%, while 52% of them belong to the rest of the population, assuming for simplicity that individuals within brackets are uniformly distributed.

Therefore, a reasonable approximation of the vote share received by the NDP among bottom 10% earners is a weighed average of vote shares in the two brackets:

$$\bar{y}_{d=1} = \frac{1 \times \bar{y}_{b=1} + 0.48 \times \bar{y}_{b=2}}{1 + 0.48} \tag{I.1}$$

This estimator is consistent, assuming that the average value taken by the dependent variable is constant within brackets. In practice, however, it does make sense to believe that the vote shares vary also within brackets in the same direction as observed between them. Therefore, this approximation should be considered as a lower bound of the true effect. Still, this method clearly does much better than computing deciles or quintiles directly from brackets – which could in fact not be quantile groups given that frequencies would necessarily be imbalanced.

Figure 1 shows the results obtained when computing vote shares for the New Democratic Party in the 2015 Canadian national election. Unsurprisingly, the two pictures look very similar, since computing vote shares by decile amounts to computing weighed averages across income brackets. Another interesting aspect of this method is that it enables us to control for structural changes not only in income, but also in other ordered variables such as education, wealth or even rural-urban scales. If university graduates were originally 5% in the 1960s and increased up to 30% in the 2010s, for instance, then one can exploit detailed educational categories to approximate "top 10% educated voters". In the 1960s, this category is composed of both university graduates and some secondary educated voters; in the 2010s, it gives more weight to individuals with masters or PhDs. This is what we do throughout the paper.

Finally, one issue is that 'splitting' brackets into deciles implies that a single individual may belong to different quantile groups: in the example above, individuals in bracket

2 belong both to the first and the second deciles. While this is not problematic when computing averages, it makes regression models impossible to solve: without changing the dataset, one cannot compare the vote shares of the first and second decile with control variables.

To solve this problem, we expand the entire dataset as many times as the number of quantile groups required. In the case of deciles, for instance, the procedure consists in duplicating all observations ten times. Then, one simply needs to attribute the corresponding weights to duplicated individuals: individuals belonging to bracket 2 see their sample weight multiplied by 0.48 in their first observation, 0.52 in the second time they appear in the dataset, and 0 in all other instances. Since this process only reweighs individuals, it leaves the effect of other explanatory variables perfectly unchanged. Finally, to account for correlation of the outcome variable of interest across duplicated observations, we cluster standard errors by individual.

I.2 Supplementary figures and tables

Table A1 - Data sources		
Country	**Election**	**Source**
Australia	1966	International Social Mobility and Politics File (Franklin et al. 1992)
Australia	1972	International Social Mobility and Politics File (Franklin et al. 1992)
Australia	1977	International Social Mobility and Politics File (Franklin et al. 1992)
Australia	1983	International Social Mobility and Politics File (Franklin et al. 1992)
Australia	1984	International Social Mobility and Politics File (Franklin et al. 1992)
Australia	1987	Australian Election Study
Australia	1990	Australian Election Study
Australia	1993	Australian Election Study
Australia	1996	Australian Election Study
Australia	1998	Australian Election Study
Australia	2001	Australian Election Study
Australia	2004	Australian Election Study
Australia	2007	Australian Election Study
Australia	2010	Australian Election Study
Australia	2013	Australian Election Study
Australia	2016	Australian Election Study
Australia	2019	Australian Election Study
Austria	1971	International Social Mobility and Politics File (Franklin et al. 1992)
Austria	1983	International Social Mobility and Politics File (Franklin et al. 1992)
Austria	1986	International Social Mobility and Politics File (Franklin et al. 1992)
Austria	1994	Eurobarometers
Austria	1995	Eurobarometers
Austria	1999	Eurobarometers
Austria	2002	European Social Survey
Austria	2006	European Social Survey
Austria	2013	Comparative Study of Electoral Systems (CSES)
Austria	2017	Comparative Study of Electoral Systems (CSES)
Belgium	1971	Eurobarometers
Belgium	1974	Eurobarometers
Belgium	1977	Eurobarometers
Belgium	1978	Eurobarometers
Belgium	1981	Eurobarometers
Belgium	1985	Eurobarometers
Belgium	1987	Eurobarometers
Belgium	1991	Belgium General Election Study
Belgium	1995	Belgium General Election Study
Belgium	1999	Belgium General Election Study
Belgium	2003	European Social Survey
Belgium	2007	European Social Survey
Belgium	2010	European Social Survey
Belgium	2014	European Social Survey
Canada	1963	Canadian Election Studies
Canada	1965	Canadian Election Studies
Canada	1968	Canadian Election Studies
Canada	1974	Canadian Election Studies

Canada	1979	Canadian Election Studies
Canada	1980	Canadian Election Studies
Canada	1984	Canadian Election Studies
Canada	1988	Canadian Election Studies
Canada	1993	Canadian Election Studies
Canada	1997	Canadian Election Studies
Canada	2000	Canadian Election Studies
Canada	2004	Canadian Election Studies
Canada	2006	Canadian Election Studies
Canada	2008	Canadian Election Studies
Canada	2011	Canadian Election Studies
Canada	2015	Canadian Election Studies
Canada	2019	Canadian Election Studies
Denmark	1960	Danish Election Study
Denmark	1964	Danish Election Study
Denmark	1966	Danish Election Study
Denmark	1968	Danish Election Study
Denmark	1971	Danish Election Study
Denmark	1973	Danish Election Study
Denmark	1975	Danish Election Study
Denmark	1977	Danish Election Study
Denmark	1979	Danish Election Study
Denmark	1981	Danish Election Study
Denmark	1984	Danish Election Study
Denmark	1987	Danish Election Study
Denmark	1988	Danish Election Study
Denmark	1990	Danish Election Study
Denmark	1994	Danish Election Study
Denmark	1998	Danish Election Study
Denmark	2001	Danish Election Study
Denmark	2005	Danish Election Study
Denmark	2007	Danish Election Study
Denmark	2011	Danish Election Study
Denmark	2015	Danish Election Study
Finland	1972	Finnish Voter Barometers
Finland	1975	Finnish Voter Barometers
Finland	1979	Finnish Voter Barometers
Finland	1983	Finnish Voter Barometers
Finland	1987	Finnish Voter Barometers
Finland	1995	Finnish Voter Barometers
Finland	1999	Finnish Voter Barometers
Finland	2003	Finnish Voter Barometers
Finland	2007	Finnish National Election Studies
Finland	2011	Finnish National Election Studies
Finland	2015	Finnish National Election Studies
France	1956	French Election Studies
France	1958	French Election Studies
France	1962	French Election Studies
France	1965	French Election Studies
France	1967	French Election Studies

France	1973	French Election Studies
France	1974	French Election Studies
France	1978	French Election Studies
France	1986	French Election Studies
France	1988	French Election Studies
France	1993	French Election Studies
France	1995	French Election Studies
France	1997	French Election Studies
France	2002	French Election Studies
France	2007	French Election Studies
France	2012	French Election Studies
France	2017	French election studies
Germany	1949	German Federal Election Studies
Germany	1953	German Federal Election Studies
Germany	1957	German Federal Election Studies
Germany	1961	German Federal Election Studies
Germany	1965	German Federal Election Studies
Germany	1969	German Federal Election Studies
Germany	1972	German Federal Election Studies
Germany	1976	German Federal Election Studies
Germany	1980	German Federal Election Studies
Germany	1983	German Federal Election Studies
Germany	1987	German Federal Election Studies
Germany	1990	German Federal Election Studies
Germany	1994	German Federal Election Studies
Germany	1998	German Federal Election Studies
Germany	2002	German Federal Election Studies
Germany	2005	German Federal Election Studies
Germany	2009	German Federal Election Studies
Germany	2013	German Federal Election Studies
Germany	2017	German Federal Election Studies
Iceland	1978	Icelandic National Election Studies
Iceland	1983	Icelandic National Election Studies
Iceland	1987	Icelandic National Election Studies
Iceland	1991	Icelandic National Election Studies
Iceland	1995	Icelandic National Election Studies
Iceland	1999	Icelandic National Election Studies
Iceland	2003	Icelandic National Election Studies
Iceland	2007	Icelandic National Election Studies
Iceland	2009	Icelandic National Election Studies
Iceland	2013	Icelandic National Election Studies
Iceland	2016	Icelandic National Election Studies
Iceland	2017	Icelandic National Election Studies
Ireland	1973	Eurobarometers
Ireland	1977	Eurobarometers
Ireland	1981	Eurobarometers
Ireland	1982	Eurobarometers
Ireland	1987	Eurobarometers
Ireland	1989	Eurobarometers
Ireland	1992	Eurobarometers

Ireland	1997	Eurobarometers
Ireland	2002	European Social Survey
Ireland	2007	European Social Survey
Ireland	2011	European Social Survey
Ireland	2016	European Social Survey
Ireland	2020	UCD Online Election Poll
Italy	1953	Inter-university Consortium for Political and Social Research (ICPSR)
Italy	1958	Inter-university Consortium for Political and Social Research (ICPSR)
Italy	1968	Italian National Election Studies
Italy	1972	Italian National Election Studies
Italy	1983	Italian National Election Studies
Italy	1987	Italian National Election Studies
Italy	1992	Italian National Election Studies
Italy	1994	Italian National Election Studies
Italy	1996	Italian National Election Studies
Italy	2001	Italian National Election Studies
Italy	2006	Comparative Study of Electoral Systems (CSES)
Italy	2008	Italian National Election Studies
Italy	2013	Italian National Election Studies
Italy	2018	Italian National Election Studies
Luxembourg	1974	Eurobarometers
Luxembourg	1979	Eurobarometers
Luxembourg	1984	Eurobarometers
Luxembourg	1989	Eurobarometers
Luxembourg	1994	Eurobarometers
Luxembourg	1999	Eurobarometers
Luxembourg	2004	European Social Survey
Luxembourg	2013	European Election Studies (EES)
Luxembourg	2018	European Election Studies (EES)
Netherlands	1967	Dutch Parliamentary Election Studies
Netherlands	1971	Dutch Parliamentary Election Studies
Netherlands	1972	Dutch Parliamentary Election Studies
Netherlands	1977	Dutch Parliamentary Election Studies
Netherlands	1981	Dutch Parliamentary Election Studies
Netherlands	1982	Dutch Parliamentary Election Studies
Netherlands	1986	Dutch Parliamentary Election Studies
Netherlands	1989	Dutch Parliamentary Election Studies
Netherlands	1994	Dutch Parliamentary Election Studies
Netherlands	1998	Dutch Parliamentary Election Studies
Netherlands	2002	Dutch Parliamentary Election Studies
Netherlands	2006	Dutch Parliamentary Election Studies
Netherlands	2010	Dutch Parliamentary Election Studies
Netherlands	2012	Dutch Parliamentary Election Studies
Netherlands	2017	Dutch Parliamentary Election Studies
New Zealand	1972	New Zealand Election Studies
New Zealand	1975	New Zealand Election Studies
New Zealand	1978	New Zealand Election Studies
New Zealand	1981	New Zealand Election Studies
New Zealand	1984	New Zealand Election Studies
New Zealand	1987	New Zealand Election Studies

New Zealand	1990	New Zealand Election Studies
New Zealand	1993	New Zealand Election Studies
New Zealand	1996	New Zealand Election Studies
New Zealand	1999	New Zealand Election Studies
New Zealand	2002	New Zealand Election Studies
New Zealand	2005	New Zealand Election Studies
New Zealand	2008	New Zealand Election Studies
New Zealand	2011	New Zealand Election Studies
New Zealand	2014	New Zealand Election Studies
New Zealand	2017	New Zealand Election Studies
Norway	1957	Norwegian National Election Studies
Norway	1965	Norwegian National Election Studies
Norway	1969	Norwegian National Election Studies
Norway	1973	Norwegian National Election Studies
Norway	1977	Norwegian National Election Studies
Norway	1981	Norwegian National Election Studies
Norway	1985	Norwegian National Election Studies
Norway	1989	Norwegian National Election Studies
Norway	1993	Norwegian National Election Studies
Norway	1997	Norwegian National Election Studies
Norway	2001	Norwegian National Election Studies
Norway	2005	Norwegian National Election Studies
Norway	2009	Norwegian National Election Studies
Norway	2013	Norwegian National Election Studies
Norway	2017	Norwegian National Election Studies
Portugal	1983	ESEO
Portugal	1985	ESEO
Portugal	1987	ESEO
Portugal	1991	ESEO
Portugal	1995	European Election Studies (EES)
Portugal	2002	Comparative Study of Electoral Systems (CSES)
Portugal	2005	Comparative Study of Electoral Systems (CSES)
Portugal	2009	Comparative Study of Electoral Systems (CSES)
Portugal	2015	Comparative Study of Electoral Systems (CSES)
Portugal	2019	Portuguese Election Study
Spain	1982	Centro de Investigaciones Sociológicas
Spain	1986	Centro de Investigaciones Sociológicas
Spain	1989	Centro de Investigaciones Sociológicas
Spain	1993	Centro de Investigaciones Sociológicas
Spain	1996	Centro de Investigaciones Sociológicas
Spain	2000	Centro de Investigaciones Sociológicas
Spain	2004	Centro de Investigaciones Sociológicas
Spain	2008	Centro de Investigaciones Sociológicas
Spain	2011	Centro de Investigaciones Sociológicas
Spain	2015	Centro de Investigaciones Sociológicas
Spain	2016	Centro de Investigaciones Sociológicas
Spain	2019	Centro de Investigaciones Sociológicas
Spain	2020	Centro de Investigaciones Sociológicas
Sweden	1956	Swedish National Election Studies
Sweden	1958	Swedish National Election Studies

Sweden	1960	Swedish National Election Studies
Sweden	1964	Swedish National Election Studies
Sweden	1968	Swedish National Election Studies
Sweden	1970	Swedish National Election Studies
Sweden	1973	Swedish National Election Studies
Sweden	1976	Swedish National Election Studies
Sweden	1979	Swedish National Election Studies
Sweden	1982	Swedish National Election Studies
Sweden	1985	Swedish National Election Studies
Sweden	1988	Swedish National Election Studies
Sweden	1991	Swedish National Election Studies
Sweden	1994	Swedish National Election Studies
Sweden	1998	Swedish National Election Studies
Sweden	2002	Swedish National Election Studies
Sweden	2006	Swedish National Election Studies
Sweden	2010	Swedish National Election Studies
Sweden	2014	Comparative Study of Electoral Systems (CSES)
Switzerland	1967	Swiss National Election Studies
Switzerland	1971	Swiss National Election Studies
Switzerland	1975	Swiss National Election Studies
Switzerland	1979	Swiss National Election Studies
Switzerland	1983	Swiss National Election Studies
Switzerland	1987	Swiss National Election Studies
Switzerland	1991	Swiss National Election Studies
Switzerland	1995	Swiss National Election Studies
Switzerland	1999	Swiss National Election Studies
Switzerland	2003	Swiss National Election Studies
Switzerland	2007	Swiss National Election Studies
Switzerland	2011	Swiss National Election Studies
Switzerland	2015	Swiss National Election Studies
Switzerland	2019	Swiss National Election Studies
UK	1955	British Election Studies
UK	1959	British Election Studies
UK	1964	British Election Studies
UK	1966	British Election Studies
UK	1970	British Election Studies
UK	1974	British Election Studies
UK	1979	British Election Studies
UK	1983	British Election Studies
UK	1987	British Election Studies
UK	1992	British Election Studies
UK	1997	British Election Studies
UK	2001	British Election Studies
UK	2005	British Election Studies
UK	2010	British Election Studies
UK	2015	British Election Studies
UK	2017	British Election Studies
US	1948	American National Election Studies
US	1952	American National Election Studies
US	1956	American National Election Studies

US	1960	American National Election Studies
US	1964	American National Election Studies
US	1968	American National Election Studies
US	1972	American National Election Studies
US	1976	American National Election Studies
US	1980	American National Election Studies
US	1984	American National Election Studies
US	1988	American National Election Studies
US	1992	American National Election Studies
US	1996	American National Election Studies
US	2000	American National Election Studies
US	2004	American National Election Studies
US	2008	American National Election Studies
US	2012	American National Election Studies
US	2016	American National Election Studies
US	2020	American National Election Studies

Source: authors' elaboration.

Table A2 - Main classification of political parties

	Social Democratic / Socialist / Communist / Green / Other left-wing parties
Australia	Labor Party, Greens
Austria	Social Democratic Party, KPÖ, Greens, NEOS, Other left
Belgium	Socialist Party, Socialist Party Differently, Ecolo, Groen, PTB
Canada	Liberal Party, Green Party, New Democratic Party
Denmark	Social Democrats, Socialist People's Party, Social Liberal Party, Red-Green Alliance
Finland	Social Democratic Party, Green League, Left Alliance, Other left
France	Socialist Party, Communist Party, Other left
Germany	Social Democratic Party, Alliance 90/The Greens, Die Linke
Iceland	Left-Green Movement, Social Democratic Alliance, People's Party
Ireland	Fianna Fáil, Sinn Féin, Labour Party, Green Party, Other left
Italy	Democratic Party, Free and Equal, Other left
Luxembourg	Socialist Workers' Party, Greens, Other left
Netherlands	Labour Party, Socialist Party, D66, Greens, Other left
New Zealand	Labour Party, Greens, Other left
Norway	Labour Party, Green Party, Socialist Left Party
Portugal	Socialist Party, Left Bloc, Unitary Democratic Coalition
Spain	Socialist Workers' Party, Podemos, United Left, Other left
Sweden	Social Democratic Party, Left Party, Green Party
Switzerland	Social Democrats, Party of Labour, Green Party, Green Liberal Party
United Kingdom	Labour Party
United States	Democratic Party

Source: authors' elaboration.

Table A3 - Detailed classification of political parties

Country	Party	Family	Left-right score (voters)	Left-right score (manifestos)
Australia	Labor Party	Social Democrats / Socialists / Other left	-0,7	-17,0
Australia	Liberal Party	Conservatives / Christian Democrats	0,8	18,2
Australia	Australian Greens	Greens	-1,5	-30,5
Australia	National Party	Conservatives / Christian Democrats	0,8	16,6
Australia	Australian Democrats	Conservatives / Christian Democrats	-0,6	-17,1
Australia	Palmer United Party	Anti-immigration		7,4
Australia	One Nation Party	Anti-immigration	0,5	
Austria	Social Democratic Party of Austria (SPÖ)	Social Democrats / Socialists / Other left	-0,6	-15,8
Austria	Austrian People's Party (ÖVP)	Conservatives / Christian Democrats	0,4	12,2
Austria	Freedom Party of Austria (FPÖ)	Anti-immigration	1,0	4,2
Austria	Greens	Greens	-1,1	-11,2
Austria	NEOS / Liberal Forum	Liberals / Social-liberals	-0,1	9,0
Belgium	Christian People's Party (CVP)	Conservatives / Christian Democrats	0,7	5,5
Belgium	Belgian Socialist Party (PSB)	Social Democrats / Socialists / Other left	-1,8	-15,2
Belgium	Socialist Party (PS)	Social Democrats / Socialists / Other left	-1,3	-16,0
Belgium	New Flemish Alliance (N-VA)	Other	0,9	9,6
Belgium	Party for Freedom and Progress (PLP/PVV)	Liberals / Social-liberals	0,4	21,1
Belgium	Open Flemish Liberals and Democrats (VLD)	Liberals / Social-liberals	0,5	7,8
Belgium	Socialist Party (SP / sp.a)	Social Democrats / Socialists / Other left	-1,3	-12,8
Belgium	Reformist movement (MR)	Liberals / Social-liberals	1,1	-12,9
Belgium	Christian Democratic and Flemish (CD&V)	Conservatives / Christian Democrats	0,5	9,8
Belgium	PL	Liberals / Social-liberals		21,9
Belgium	Christian Social Party (PSC)	Conservatives / Christian Democrats	0,5	-2,9
Belgium	Liberal Reformist Party (PRL)	Liberals / Social-liberals	0,3	7,1
Belgium	Volksunie (VU)	Other	0,3	3,3
Belgium	Vlaams Blok	Anti-immigration	1,1	8,7
Belgium	Workers' Party of Belgium (PTB)	Social Democrats / Socialists / Other left	-2,1	-29,3

Belgium	Communist Party (PCB)	Social Democrats / Socialists / Other left	-2,7	
Canada	Liberal Party	Social Democrats / Socialists / Other left	-0,1	-1,1
Canada	Conservative Party	Conservatives / Christian Democrats	0,7	10,5
Canada	Canadian Alliance	Conservatives / Christian Democrats		18,8
Canada	Reform Party	Conservatives / Christian Democrats	0,6	29,2
Canada	New Democratic Party	Social Democrats / Socialists / Other left	-0,9	-26,9
Canada	Bloc Québécois	Other	-0,7	-5,9
Canada	Social Credit Party	Conservatives / Christian Democrats	-0,5	7,5
Denmark	Social Democratic Party	Social Democrats / Socialists / Other left	-1,0	-15,5
Denmark	Liberal Party of Denmark (Venstre)	Liberals / Social-liberals	1,5	17,0
Denmark	Conservative People's Party	Conservatives / Christian Democrats	1,8	23,9
Denmark	Danish People's Party	Anti-immigration	1,3	27,9
Denmark	Progress Party	Anti-immigration	1,5	25,5
Denmark	Socialist People's Party	Social Democrats / Socialists / Other left	-2,2	-34,9
Denmark	Danish Social-Liberal Party (Radikale Venstre)	Liberals / Social-liberals	-0,6	-7,8
Finland	Social Democratic Party	Social Democrats / Socialists / Other left	-1,1	-4,2
Finland	Agrarian Union	Other		8,9
Finland	Centre Party	Other	0,6	2,2
Finland	Finnish People's Democratic League	Communists	-2,1	-22,4
Finland	National Coalition Party	Conservatives / Christian Democrats	1,5	10,4
Finland	True Finns	Anti-immigration	-0,2	7,4
Finland	Left Alliance	Social Democrats / Socialists / Other left	-2,2	-27,1
Finland	Greens	Greens	-0,8	-17,2
Finland	Finnish People’s Party	Liberals / Social-liberals		27,0
Finland	Finnish Rural Party	Conservatives / Christian Democrats	-0,1	26,3
Finland	Swedish People's Party	Other	0,9	0,7
France	UDR/UNR	Conservatives / Christian Democrats		25,6
France	La République En Marche! (LRM)	Liberals / Social-liberals	-0,4	4,8
France	UDF/MoDem	Conservatives / Christian Democrats	0,1	14,8
France	LR/UMP/RPR	Conservatives / Christian Democrats	1,5	9,7
France	PS/SFIO	Social Democrats / Socialists / Other left	-1,7	-23,0
France	Communist Party (PCF)	Communists		-24,4
France	MRP/CD	Conservatives / Christian Democrats		10,3
France	Reforming Movement (MR, 1973)	Conservatives / Christian Democrats		3,8

France	Republican Party of Liberty - Conservatives	Conservatives / Christian Democrats		1,5
France	National Front (FN)	Anti-immigration	1,5	32,2
France	Progress and Modern Democracy	Other		1,2
France	Rally for the French People - Gaullists	Conservatives / Christian Democrats		12,0
France	La France Insoumise (FI) / Front de gauche (FDG)	Social Democrats / Socialists / Other left	-2,2	-27,6
France	National Centre of Independents and Peasants (CNIP)	Conservatives / Christian Democrats		23,1
France	Radical Party	Social Democrats / Socialists / Other left		-6,3
Germany	CDU/CSU	Conservatives / Christian Democrats		12,6
Germany	Social Democratic Party of Germany (SPD)	Social Democrats / Socialists / Other left		-13,0
Germany	Die Linke	Social Democrats / Socialists / Other left		-29,1
Germany	Free Democratic Party (FDP)	Liberals / Social-liberals		4,5
Germany	Alternative for Germany (AfD)	Anti-immigration		15,9
Germany	Greens	Greens		-17,2
Germany	All-German Bloc (GB/BHE)	Conservatives / Christian Democrats		-1,3
Iceland	Independence Party	Conservatives / Christian Democrats	1,6	15,4
Iceland	Social Democratic Alliance	Social Democrats / Socialists / Other left	-1,3	-12,2
Iceland	Progressive Party	Conservatives / Christian Democrats	0,0	6,5
Iceland	United Socialist Party	Social Democrats / Socialists / Other left		-13,4
Iceland	People's Alliance	Social Democrats / Socialists / Other left	-1,9	-26,3
Iceland	Social Democratic Party	Social Democrats / Socialists / Other left	-0,2	-24,3
Iceland	Left-Green Movement	Greens	-2,2	-15,6
Iceland	Centre Party	Conservatives / Christian Democrats	0,9	
Iceland	Pirate Party	Other	-1,0	-15,3
Iceland	Reform Party	Liberals / Social-liberals	0,7	5,7
Iceland	Women's Alliance	Social Democrats / Socialists / Other left	-1,1	-33,5
Iceland	People’s Party	Other		-18,0
Iceland	Liberal Party	Conservatives / Christian Democrats	-0,1	13,9
Iceland	National Preservation Party	Other		-38,5
Iceland	Bright Future	Liberals / Social-liberals	-0,7	2,2
Ireland	Fianna Fáil	Social Democrats / Socialists / Other left	0,4	2,8
Ireland	Fine Gael	Conservatives / Christian Democrats	0,3	6,7
Ireland	Labour Party	Social Democrats / Socialists / Other left	-1,1	-21,9
Ireland	Sinn Féin	Social Democrats / Socialists / Other left	-1,3	-9,4
Ireland	Progressive Democrats	Conservatives / Christian Democrats	0,3	11,1

Italy	Christian Democracy (DC)	Conservatives / Christian Democrats	1,3	6,6
Italy	Olive Tree	Social Democrats / Socialists / Other left	-2,1	-32,9
Italy	People of Freedom (PDL)	Conservatives / Christian Democrats	2,5	14,7
Italy	Five Star Movement (M5S)	Social Democrats / Socialists / Other left	-0,6	-20,5
Italy	Italian Communist Party (PCI)	Communists	-2,2	-10,2
Italy	Democratic Party (PD)	Social Democrats / Socialists / Other left	-2,1	-3,2
Italy	Forza Italia (FI)	Conservatives / Christian Democrats	2,2	25,4
Italy	Democratic Party of the Left (PDS)	Social Democrats / Socialists / Other left	-2,9	-2,8
Italy	Democrats of the Left (DS) / Margherita / Ulivo	Social Democrats / Socialists / Other left	-1,8	-12,8
Italy	Italian Socialist Party of Proletarian Unity (PSIUP)	Social Democrats / Socialists / Other left	-1,3	-1,5
Italy	National Alliance (AN)	Conservatives / Christian Democrats	3,1	6,5
Italy	Populars for Italy (PPI)	Social Democrats / Socialists / Other left	0,1	-2,2
Italy	Italian Socialist Party (PSI)	Social Democrats / Socialists / Other left	-0,4	-9,9
Italy	Civic Choice	Conservatives / Christian Democrats	0,3	15,3
Italy	Lega	Anti-immigration	1,8	7,0
Italy	Socialist Party of Italian Workers	Social Democrats / Socialists / Other left		-34,7
Italy	Communist Refoundation Party (PRC)	Social Democrats / Socialists / Other left	-3,1	-32,9
Italy	Italian Social Movement (MSI, MSI-DN)	Anti-immigration	3,6	16,0
Luxembourg	Christian Social People's Party	Conservatives / Christian Democrats	1,1	8,4
Luxembourg	Luxembourg Socialist Workers' Party	Social Democrats / Socialists / Other left	-1,3	-13,9
Luxembourg	Democratic Party	Liberals / Social-liberals	0,2	11,8
Luxembourg	Democratic Group	Liberals / Social-liberals		1,5
Luxembourg	Patriotic and Democratic Group	Liberals / Social-liberals		9,5
Luxembourg	Action Committee	Conservatives / Christian Democrats	0,2	7,7
Luxembourg	The Greens	Greens	-1,4	-11,1
Luxembourg	Communist Party of Luxembourg	Communists	-2,0	-25,3
Luxembourg	Green List Ecological Initiative	Greens	-1,2	-10,1
Luxembourg	Alternative Democratic Reform Party	Anti-immigration		14,9
Netherlands	Catholic People's Party (KVP)	Conservatives / Christian Democrats		5,0
Netherlands	Labour Party (PvdA)	Social Democrats / Socialists / Other left	-1,8	-15,1
Netherlands	Christian Democratic Appeal (CDA)	Conservatives / Christian Democrats	1,2	1,2
Netherlands	People's Party for Freedom and Democracy (VVD)	Liberals / Social-liberals	1,4	19,6
Netherlands	Pim Fortuyn List (LPF)	Anti-immigration	1,0	4,2
Netherlands	Party for Freedom (PVV)	Anti-immigration	1,3	17,2

Netherlands	Anti-Revolutionary Party (ARP)	Conservatives / Christian Democrats		11,9
Netherlands	Christian Historical Union (CHU)	Conservatives / Christian Democrats		15,8
Netherlands	Socialist Party (SP)	Social Democrats / Socialists / Other left	-1,4	-20,3
Netherlands	Democrats 66 (D66)	Liberals / Social-liberals	-0,7	-6,5
Netherlands	Communist Party of the Netherlands	Communists		-29,3
Netherlands	PvdV	Conservatives / Christian Democrats		20,7
Netherlands	GroenLinks (GL)	Greens	-2,3	-9,6
New Zealand	National Party	Conservatives / Christian Democrats	1,2	15,1
New Zealand	Labour Party	Social Democrats / Socialists / Other left	-1,1	-15,0
New Zealand	Alliance	Greens	-1,5	-14,1
New Zealand	Social Credit Party	Social Democrats / Socialists / Other left	-0,9	-8,7
New Zealand	New Zealand First	Anti-immigration	0,0	0,9
New Zealand	Green Party of Aotearoa	Greens	-2,0	-2,9
Norway	Labour Party	Social Democrats / Socialists / Other left	-1,2	-15,4
Norway	Conservative Party	Conservatives / Christian Democrats	1,8	17,1
Norway	Progress Party	Anti-immigration	1,8	35,2
Norway	Christian Democratic Party	Conservatives / Christian Democrats	0,6	10,3
Norway	Centre Party	Conservatives / Christian Democrats	-0,3	6,1
Norway	Socialist Left Party / Socialist Electoral League	Social Democrats / Socialists / Other left	-2,4	-20,4
Norway	Liberal Party	Liberals / Social-liberals	-0,3	-3,3
Portugal	Socialist Party (PS)	Social Democrats / Socialists / Other left	-0,9	-6,4
Portugal	PPD/PSD	Conservatives / Christian Democrats	1,6	5,2
Portugal	United People Alliance (APU)	Greens	-2,7	9,8
Portugal	PCTP/MRPP	Communists		-7,0
Portugal	CDS / People's Party (PP)	Conservatives / Christian Democrats	1,8	9,7
Portugal	Unitary Democratic Coalition (CDU, PCP-PEV)	Greens	-3,1	-8,1
Portugal	Left Bloc (BE)	Social Democrats / Socialists / Other left	-2,4	-23,1
Spain	Spanish Socialist Workers' Party (PSOE)	Social Democrats / Socialists / Other left	-1,1	-8,2
Spain	People's Party (PP)	Conservatives / Christian Democrats	1,9	13,2
Spain	Union of the Democratic Centre (UCD)	Other	-1,3	2,6
Spain	AP-PDP	Conservatives / Christian Democrats	2,2	16,9
Spain	VOX	Anti-immigration	2,5	45,9
Spain	Ciudadanos	Liberals / Social-liberals	0,8	-1,2
Spain	Podemos	Social Democrats / Socialists / Other left	-1,9	-20,8

Spain	Communist Party of Spain (PCE)	Communists	-2,0	-17,1
Spain	United Left (IU)	Social Democrats / Socialists / Other left	-2,1	-20,0
Spain	Democratic and Social Centre (CDS)	Other	0,5	-3,9
Sweden	Swedish Social Democratic Party	Social Democrats / Socialists / Other left	-1,4	-15,6
Sweden	Moderate/Right Party	Conservatives / Christian Democrats	2,1	39,1
Sweden	Liberal People's Party	Liberals / Social-liberals	1,1	6,0
Sweden	Centre Party	Liberals / Social-liberals	0,9	7,4
Sweden	Sweden Democrats	Anti-immigration	0,5	15,0
Sweden	Left Party	Social Democrats / Socialists / Other left	-2,4	-29,6
Sweden	Christian Democrats	Conservatives / Christian Democrats	1,1	5,9
Sweden	New Democracy	Anti-immigration	1,0	34,4
Sweden	Green Party	Greens	-0,9	-14,2
Sweden	Left Party/Communists	Communists	-2,7	-28,6
Switzerland	Social Democratic Party of Switzerland (SPS/PSS)	Social Democrats / Socialists / Other left	-2,0	-30,0
Switzerland	Free Democratic Party of Switzerland (FDP/PLR)	Liberals / Social-liberals	0,8	16,1
Switzerland	CVP/PDC	Conservatives / Christian Democrats	0,6	5,0
Switzerland	Swiss People's Party (SVP/UDC)	Anti-immigration	1,3	13,9
Switzerland	Green Party of Switzerland (GPS/PES)	Greens	-2,0	-26,3
Switzerland	Green Liberal Party of Switzerland (GLP/PVL)	Greens	-1,0	-5,2
USA	Democratic Party	Social Democrats / Socialists / Other left	-0,9	-13,3
USA	Republican Party	Conservatives / Christian Democrats	1,0	14,6
UK	Conservative Party	Conservatives / Christian Democrats	1,2	15,5
UK	Labour Party	Social Democrats / Socialists / Other left	-0,9	-14,7
UK	Liberal Democrats	Liberals / Social-liberals	-0,4	-0,8
UK	Social Democratic Party	Social Democrats / Socialists / Other left		-10,4
UK	UK Independence Party (UKIP)	Anti-immigration	0,3	16,5

Source: authors' elaboration.

Note: the table provides information on the categorization of political parties by family in the survey dataset (see Figure 4 on election results). Parties are sorted by decreasing order of their average vote share in all elections to which they participated. Excludes small parties (average vote share lower than 5% across elections in which the party participated). The left-right score (voters) corresponds to the difference between the average self-placement on a left-right scale (0 to 10) of voters of the corresponding party and the overall average of this variable across all voters. Negative values mean that voters supporting the party are on average more left-wing than the rest of the electorate. The left-right score (manifestos) corresponds to the difference between the average left-right ideological index of the corresponding party in the Comparative Manifesto Project database (-100 to 100) and the overall average of this variable across all parties. Averages over the entire dataset.

Figure A1 - The disconnection of income and education cleavages in Western democracies (before and after controls)

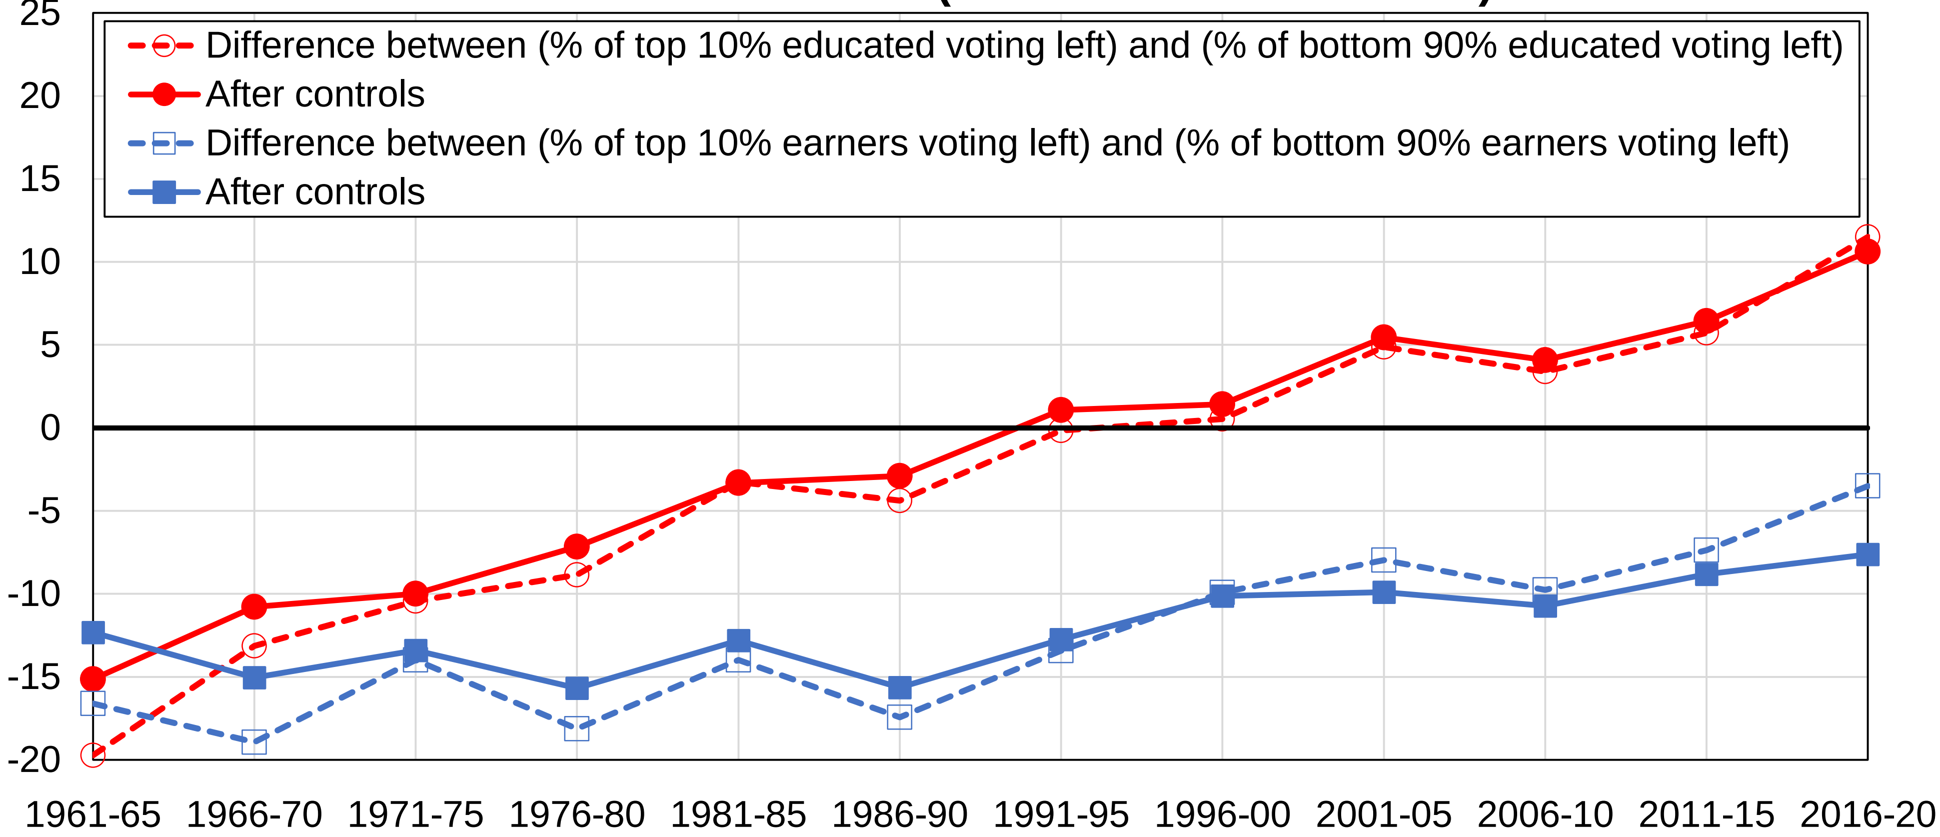

Source: authors' computations using the World Political Cleavages and Inequality Database.

Note: in the 1960s, both higher-educated and high-income voters were less likely to vote for left-wing (social democratic / socialist / communist / green / other left-wing) parties than lower-educated and low-income voters by more than 10 percentage points. The left vote has gradually become associated with higher education voters, giving rising to a remarkable divergence of the effects of income and education on the vote. Figures correspond to five-year averages for Australia, Britain, Canada, Denmark, France, Germany, Italy, the Netherlands, Norway, Sweden, Switzerland, and the US. The estimates are presented before and after controlling for income/education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure A2 - The disconnection of income and education in Western democracies, unbalanced panel

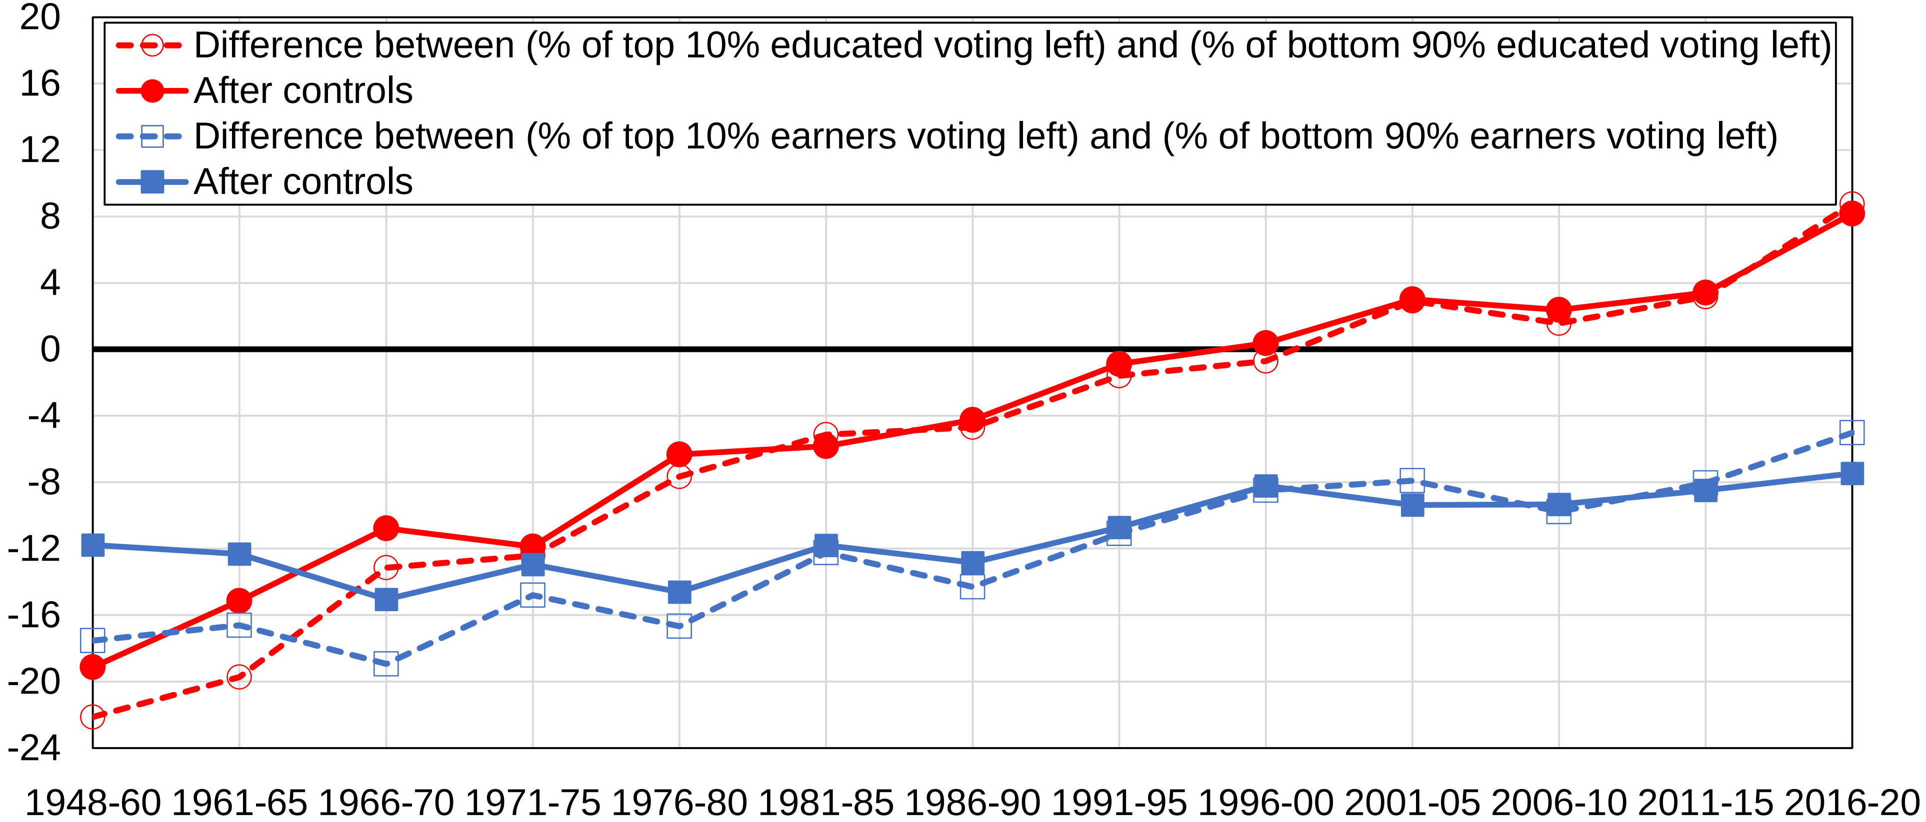

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: in the 1960s, both higher-educated and high-income voters were less likely to vote for left-wing (social democratic / socialist / communist / green / other left-wing) parties than lower-educated and low-income voters by more than 10 percentage points. The left vote has gradually become associated with higher education voters, giving rising to a remarkable divergence of the effects of income and education on the vote. Figures correspond to five-year averages over all countries available for a given time period (unbalanced panel of all 21 Western democracies). The estimates are presented before and after controlling for income/education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure A3 - The disconnection of income and education cleavages in Western democracies (top 50% vs. bottom 50%)

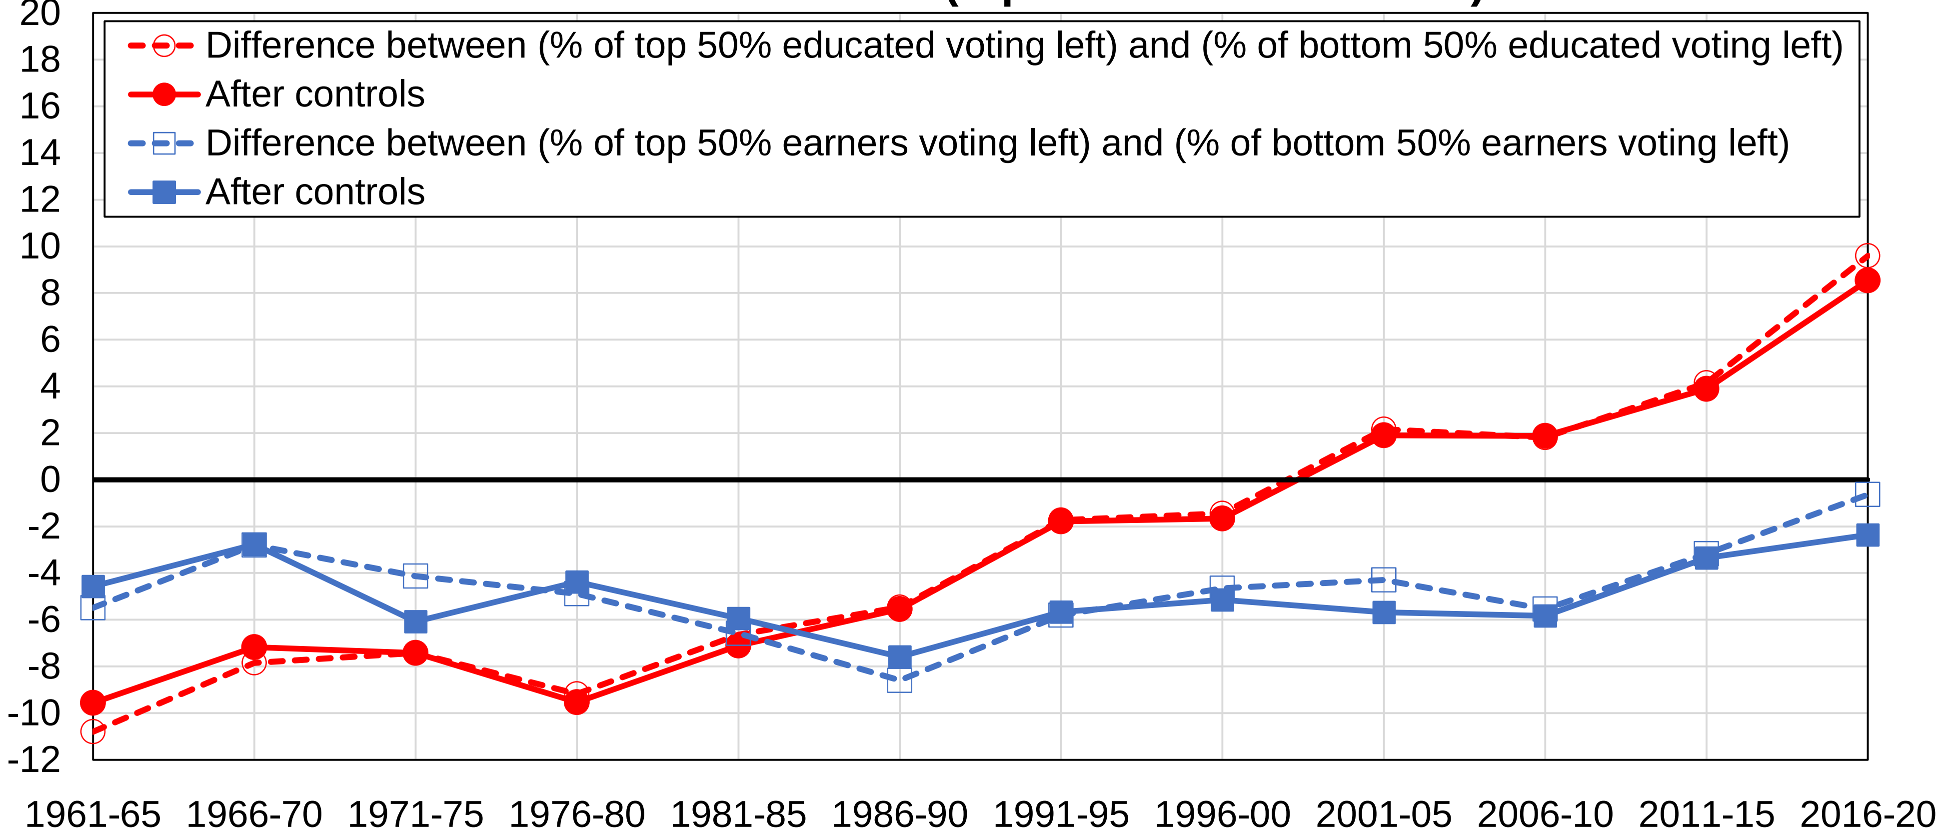

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: in the 1960s, both higher-educated and high-income voters were less likely to vote for left-wing (social democratic / socialist / communist / green / other left-wing) parties than lower-educated and low-income voters. The left vote has gradually become associated with higher education voters, giving rising to a remarkable divergence of the effects of income and education on the vote. Figures correspond to five-year averages for Australia, Britain, Canada, Denmark, France, Germany, Italy, the Netherlands, Norway, Sweden, Switzerland, and the US. The estimates are presented before and after controlling for income/education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure A4 - The disconnection of income and education in Western democracies, unbalanced panel (top 50% vs. bottom 50%)

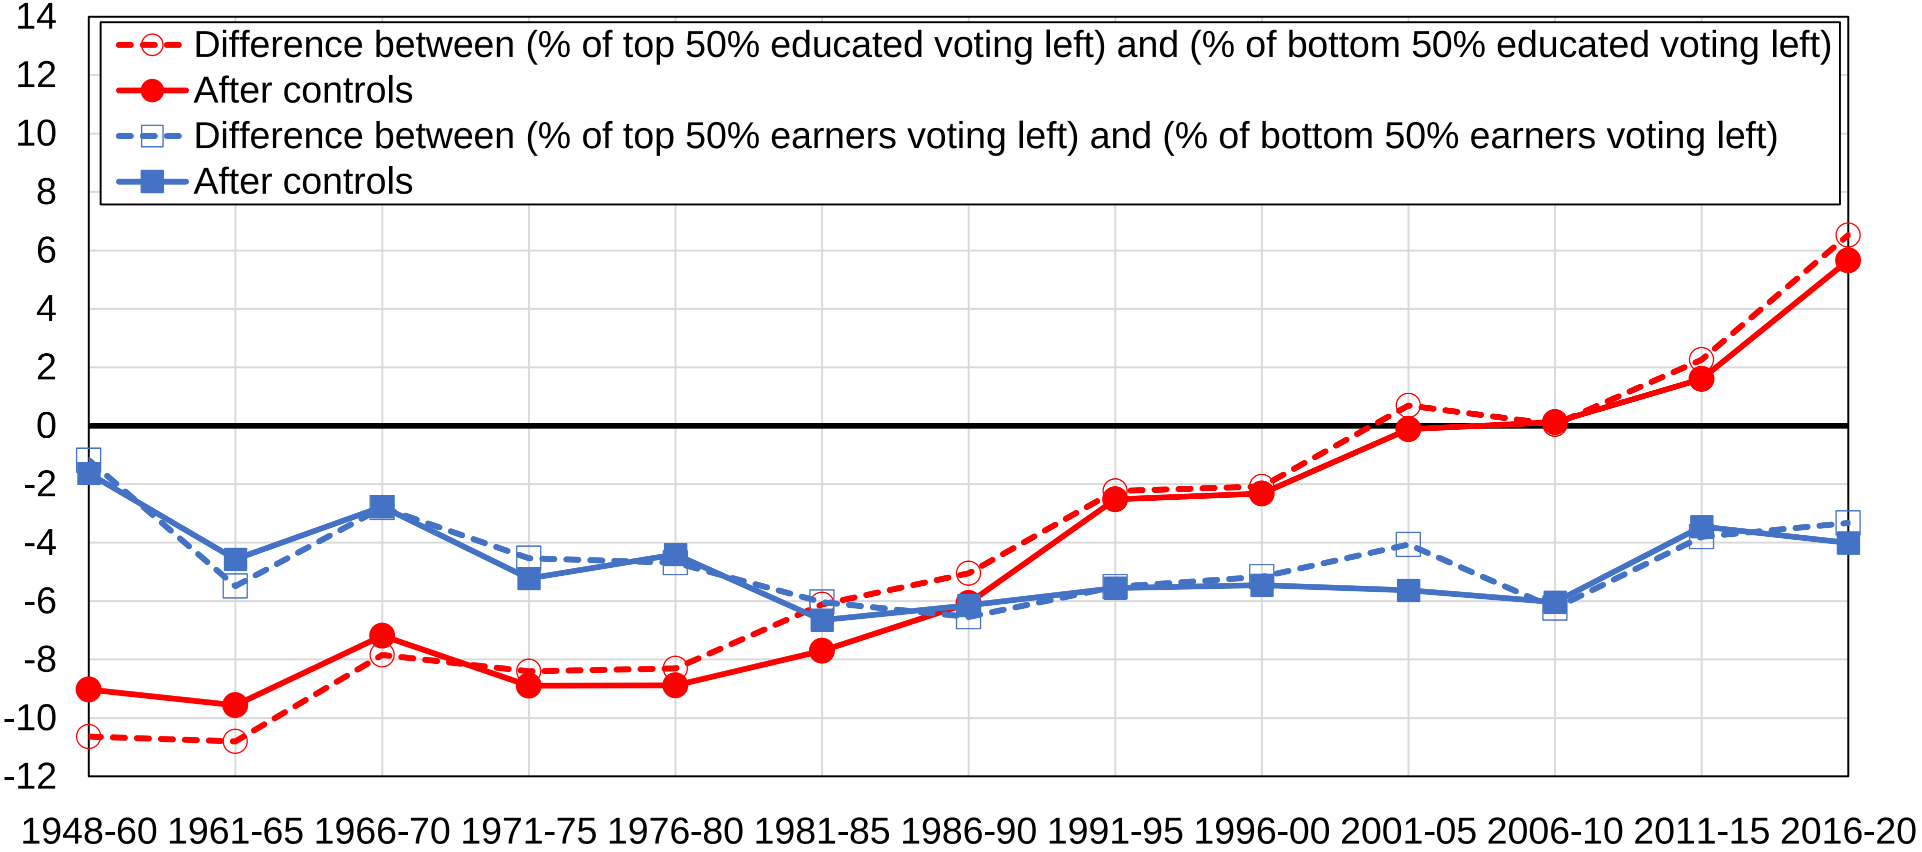

Source: authors' computations using the World Political Cleavages and Inequality Database.

Note: in the 1960s, both higher-educated and high-income voters were less likely to vote for left-wing (social democratic / socialist / communist / green / other left-wing) parties than lower-educated and low-income voters. The left vote has gradually become associated with higher education voters, giving rising to a remarkable divergence of the effects of income and education on the vote. Figures correspond to five-year averages over all countries available for a given time period (unbalanced panel of all 21 Western democracies). The estimates are presented before and after controlling for income/education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure A5 - The reversal of educational divides in Western democracies (top 10%)

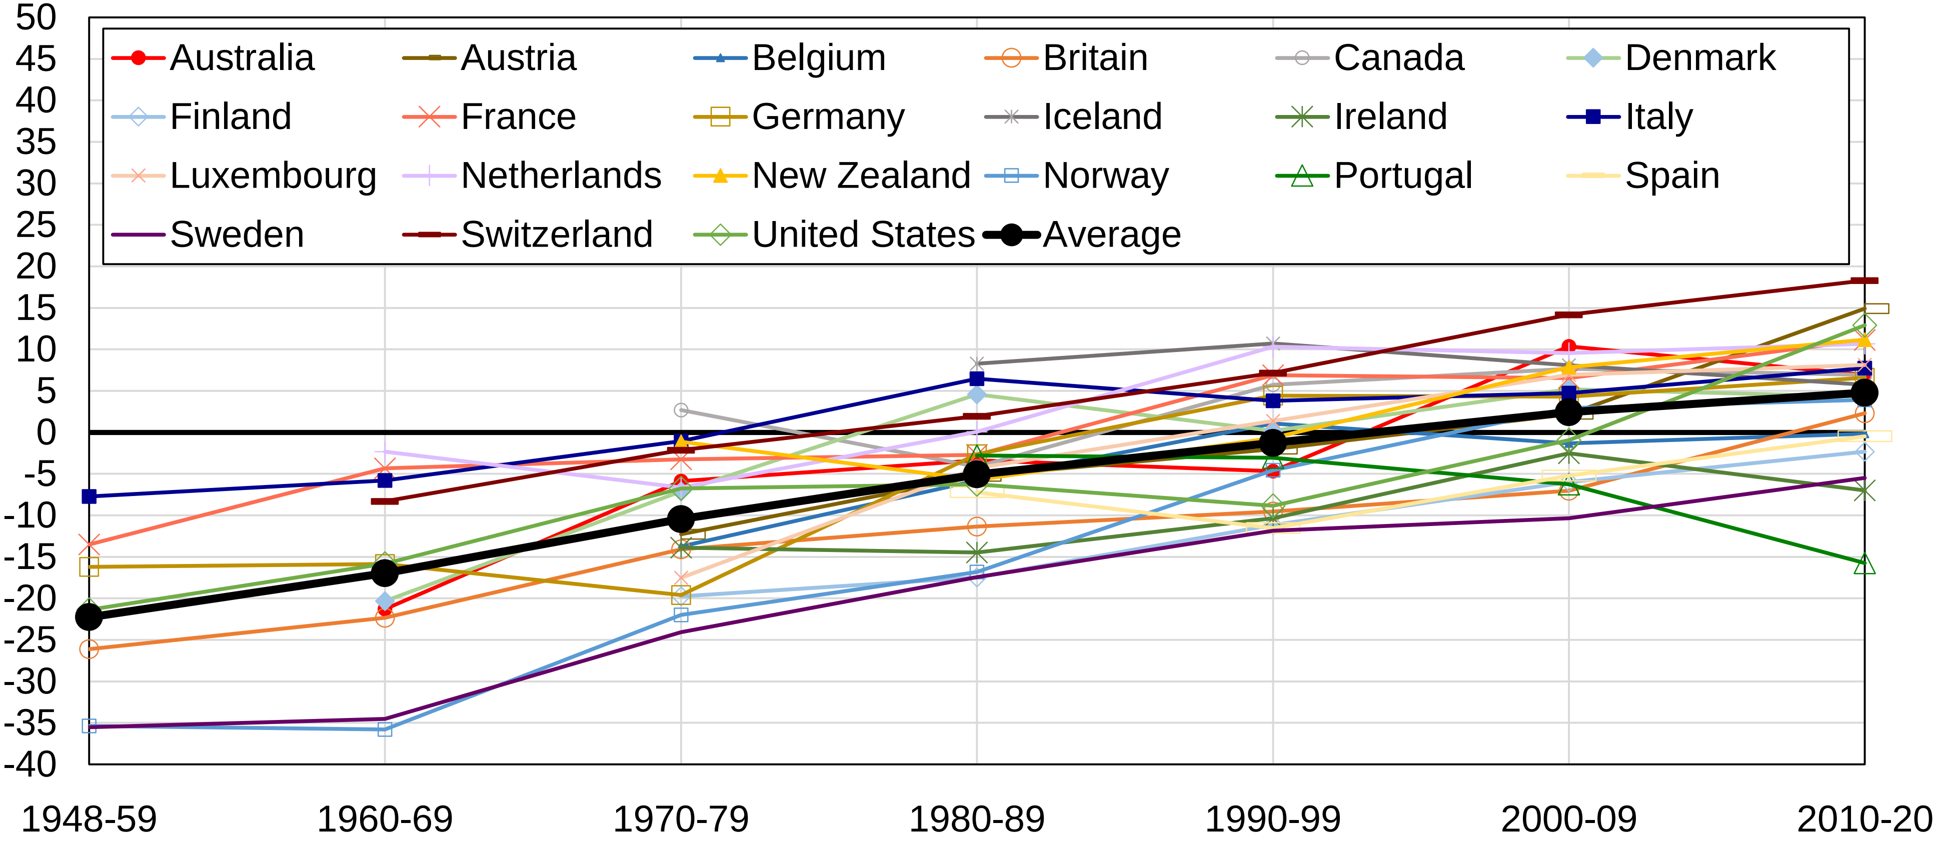

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of higher-educated (top 10%) and lower-educated (bottom 90%) voters voting for left-wing (socialist, social democratic, communist, and green) parties in Western countries. In nearly all countries, higher-educated voters used to be significantly more likely to vote for right-wing parties and have gradually become more likely to vote for left-wing parties.

Figure A6 - The reversal of educational divides in Western democracies (top 10%), after controls

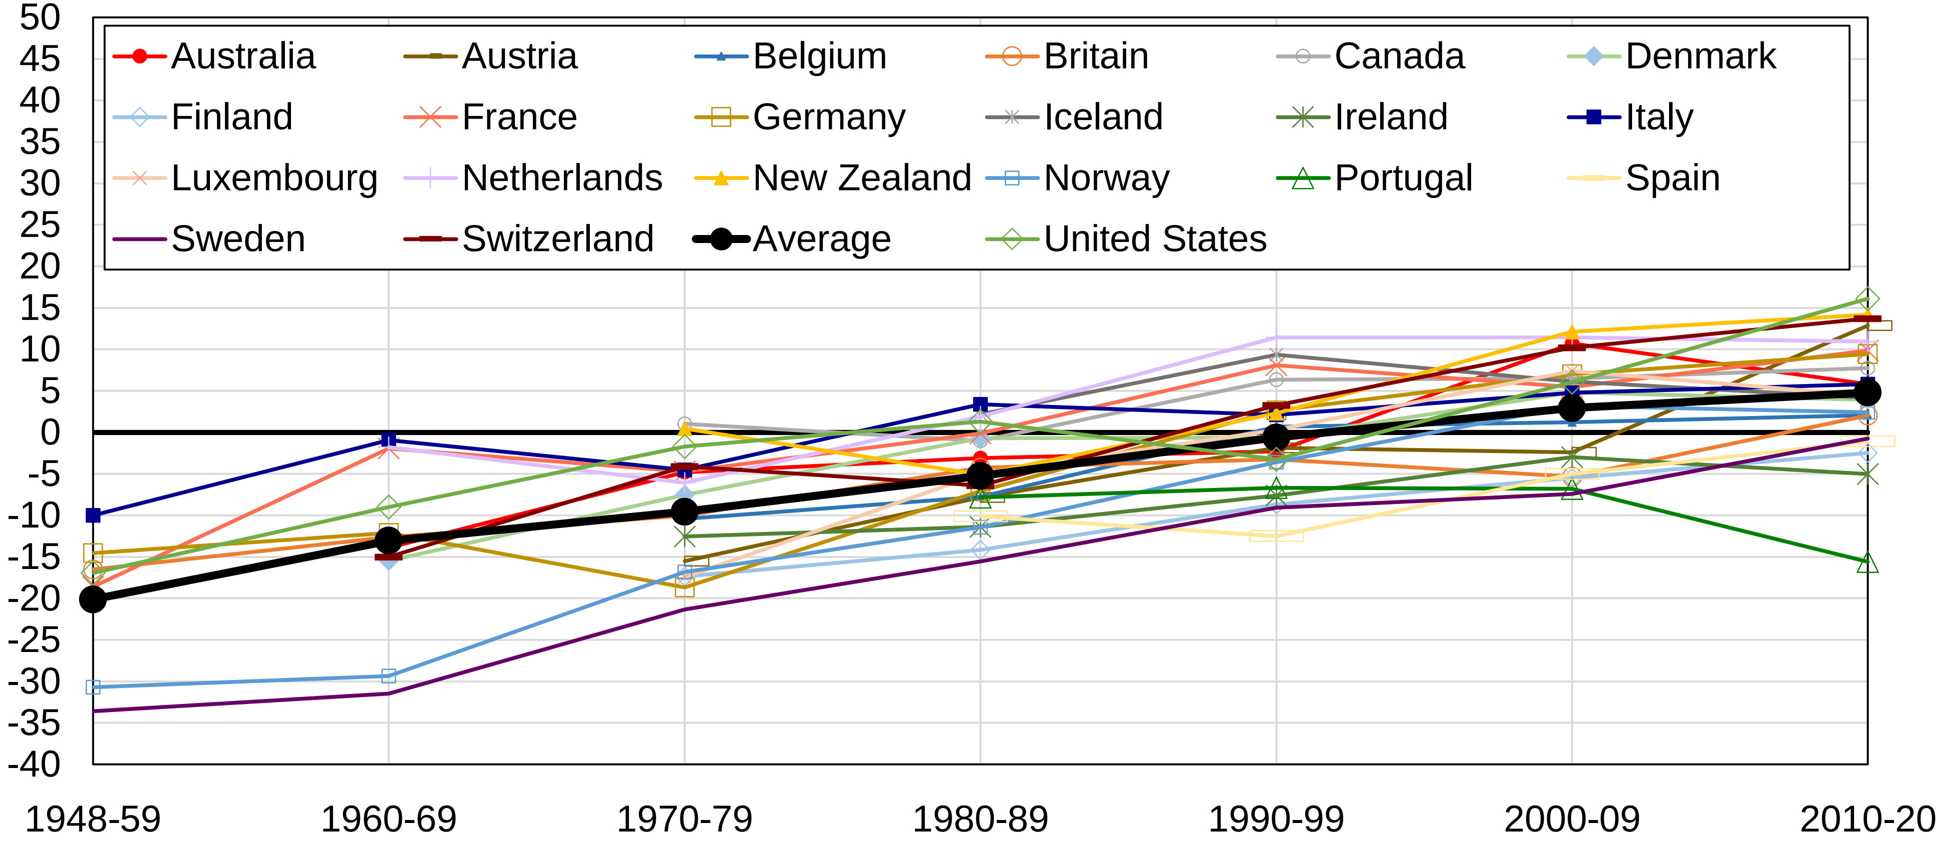

Source: authors' computations using the World Political Cleavages and Inequality Database.

Note: the figure represents the difference between the share of higher-educated (top 10%) and lower-educated (bottom 90%) voters voting for left-wing (socialist, social democratic, communist, and green) parties in Western countries, after controlling for income, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available). In nearly all countries, higher-educated voters used to be significantly more likely to vote for right-wing parties and have gradually become more likely to vote for left-wing parties.

Figure A7 - The reversal of educational divides in Western democracies (university graduates)

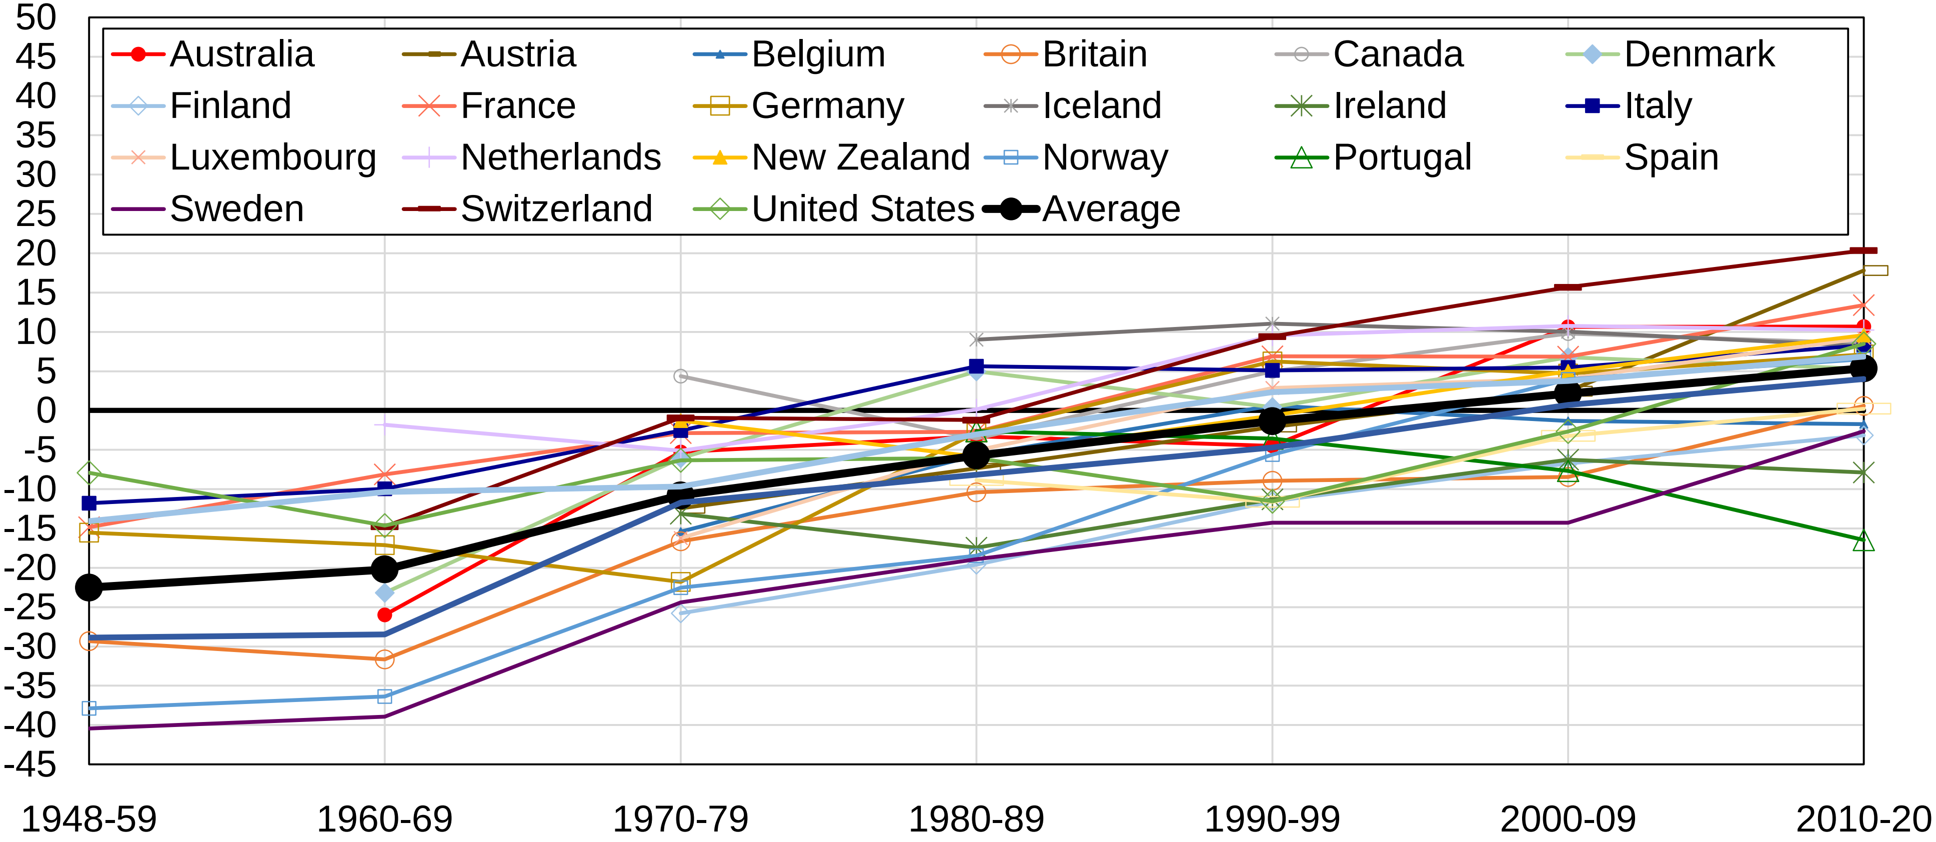

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of university graduates and the share of non-university graduates voting for left-wing (socialist, social democratic, communist, and green) parties in Western countries. In nearly all countries, higher-educated voters used to be significantly more likely to vote for right-wing parties and have gradually become more likely to vote for left-wing parties.

Figure A8 - The reversal of educational divides in Western democracies (university graduates), after controls

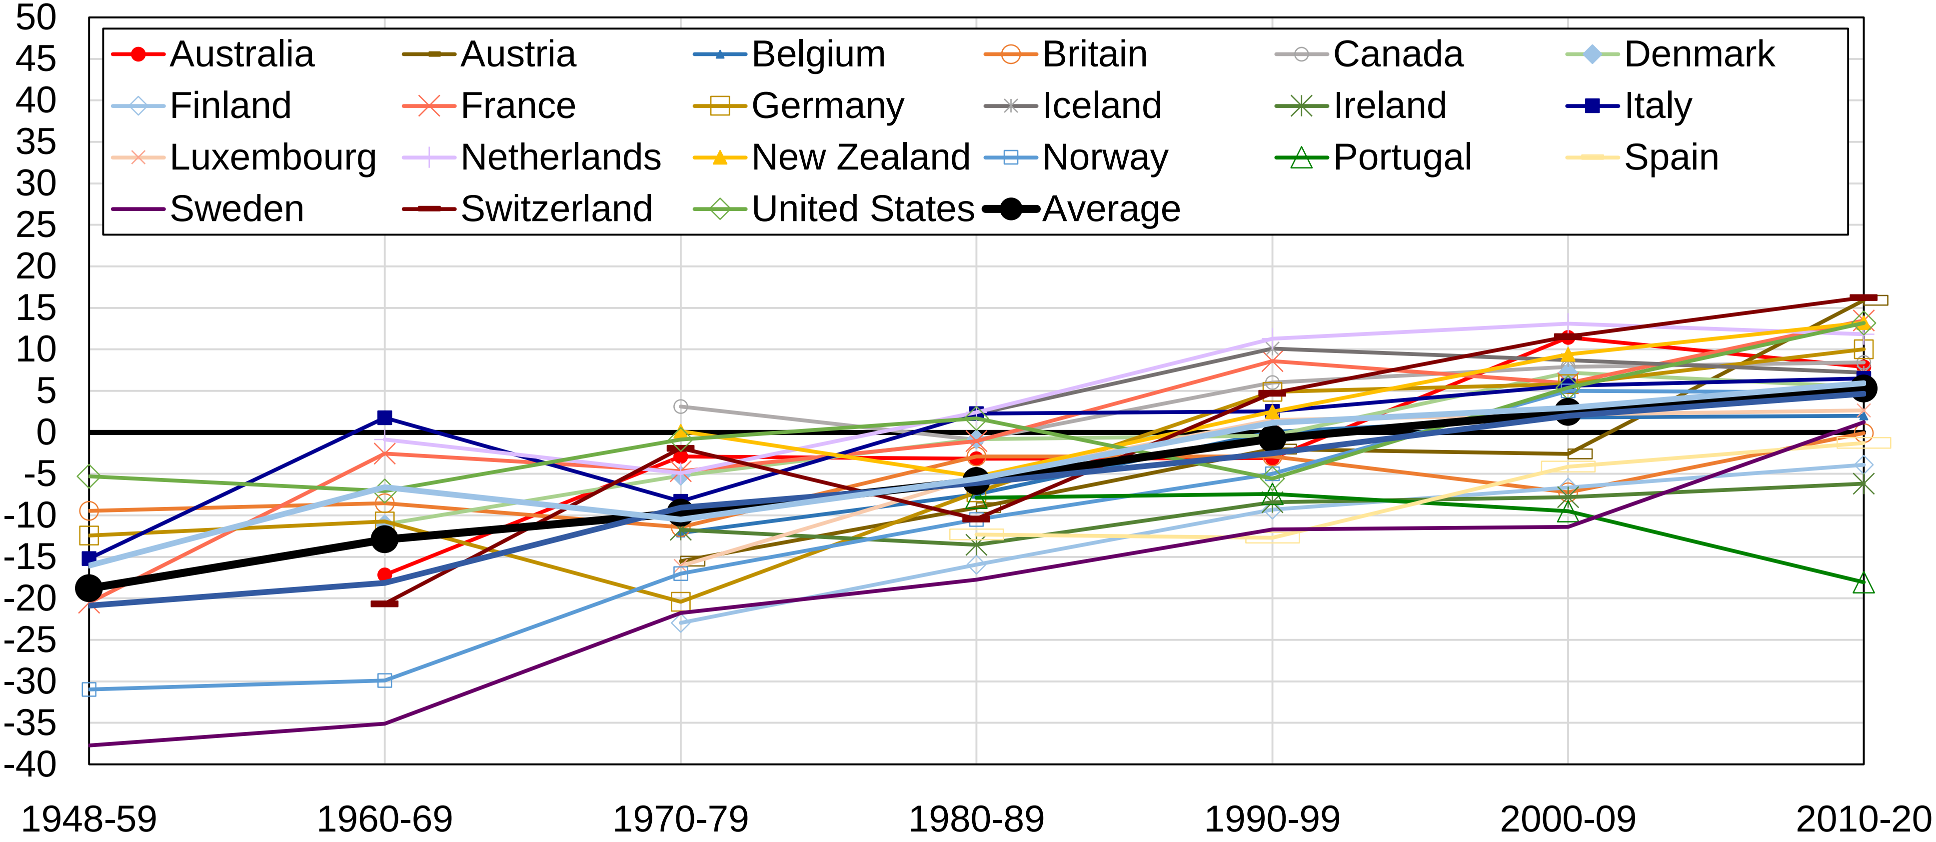

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of university graduates and the share of non-university graduates voting for left-wing (socialist, social democratic, communist, and green) parties in Western countries, after controlling for income, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available). In nearly all countries, higher-educated voters used to be significantly more likely to vote for right-wing parties and have gradually become more likely to vote for left-wing parties.

Figure A9 - The reversal of educational divides in Western democracies (bottom 50%)

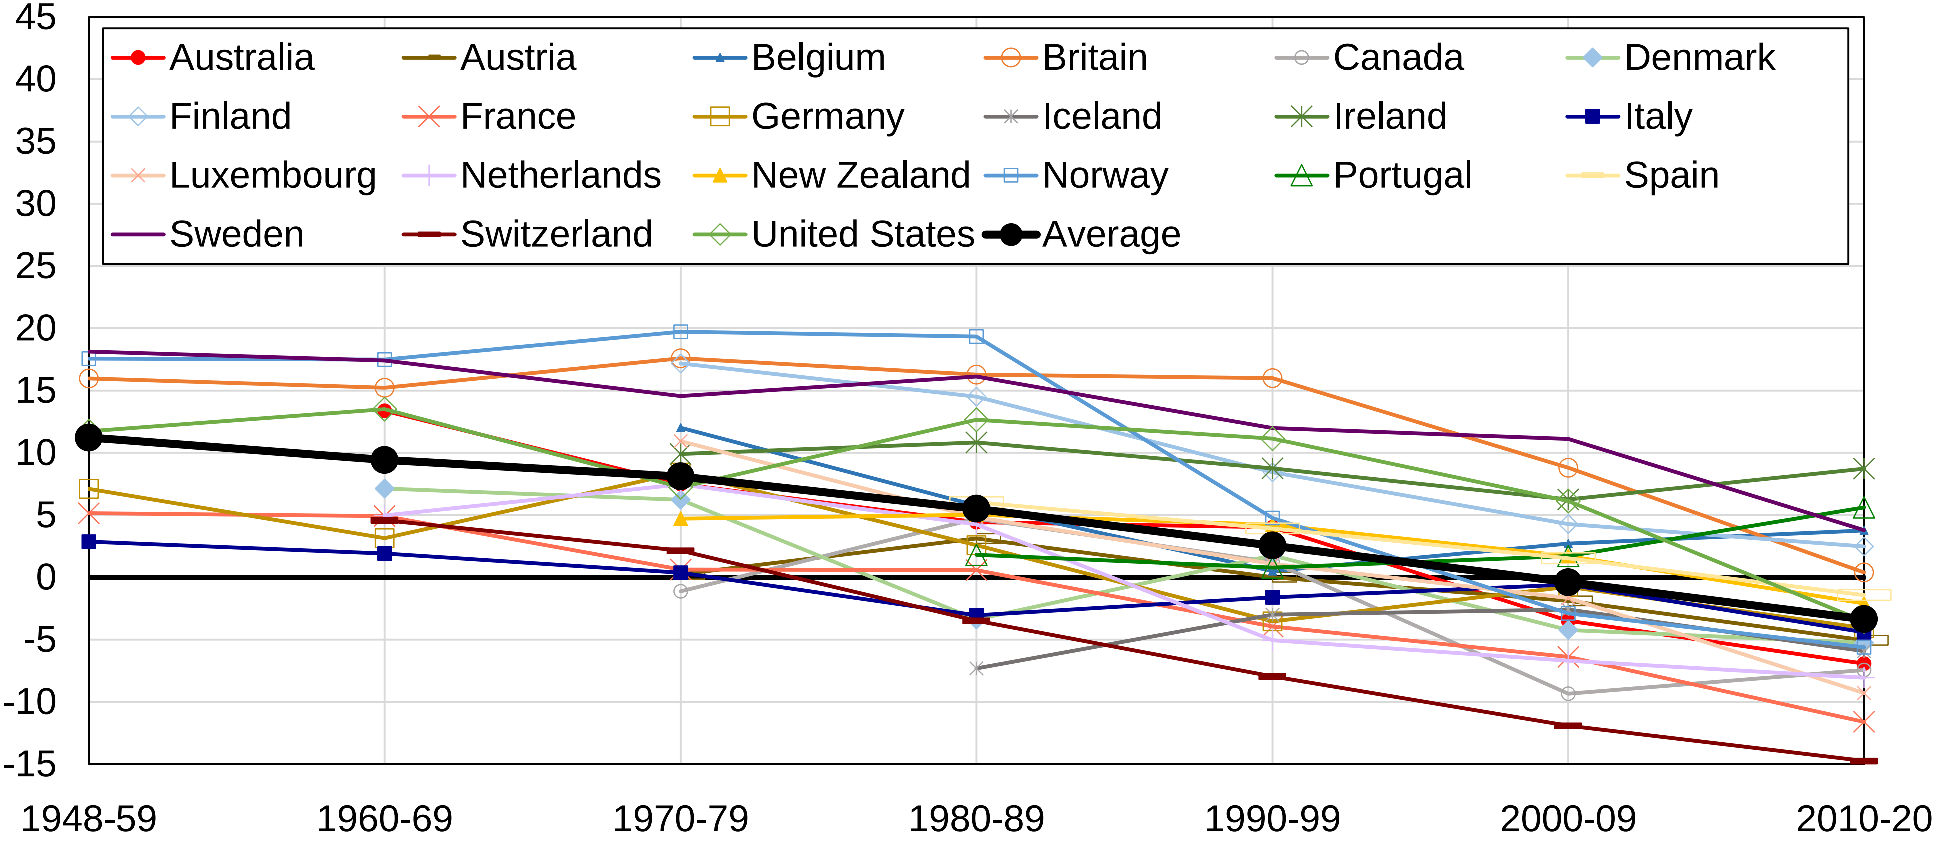

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of lower-educated (bottom 50%) and higher-educated (top 50%) voters voting for left-wing (socialist, social democratic, communist, and green) parties in Western countries. In nearly all countries, higher-educated voters used to be significantly more likely to vote for right-wing parties and have gradually become more likely to vote for left-wing parties.

Figure A10 - The reversal of educational divides in Western democracies (bottom 50%), after controls

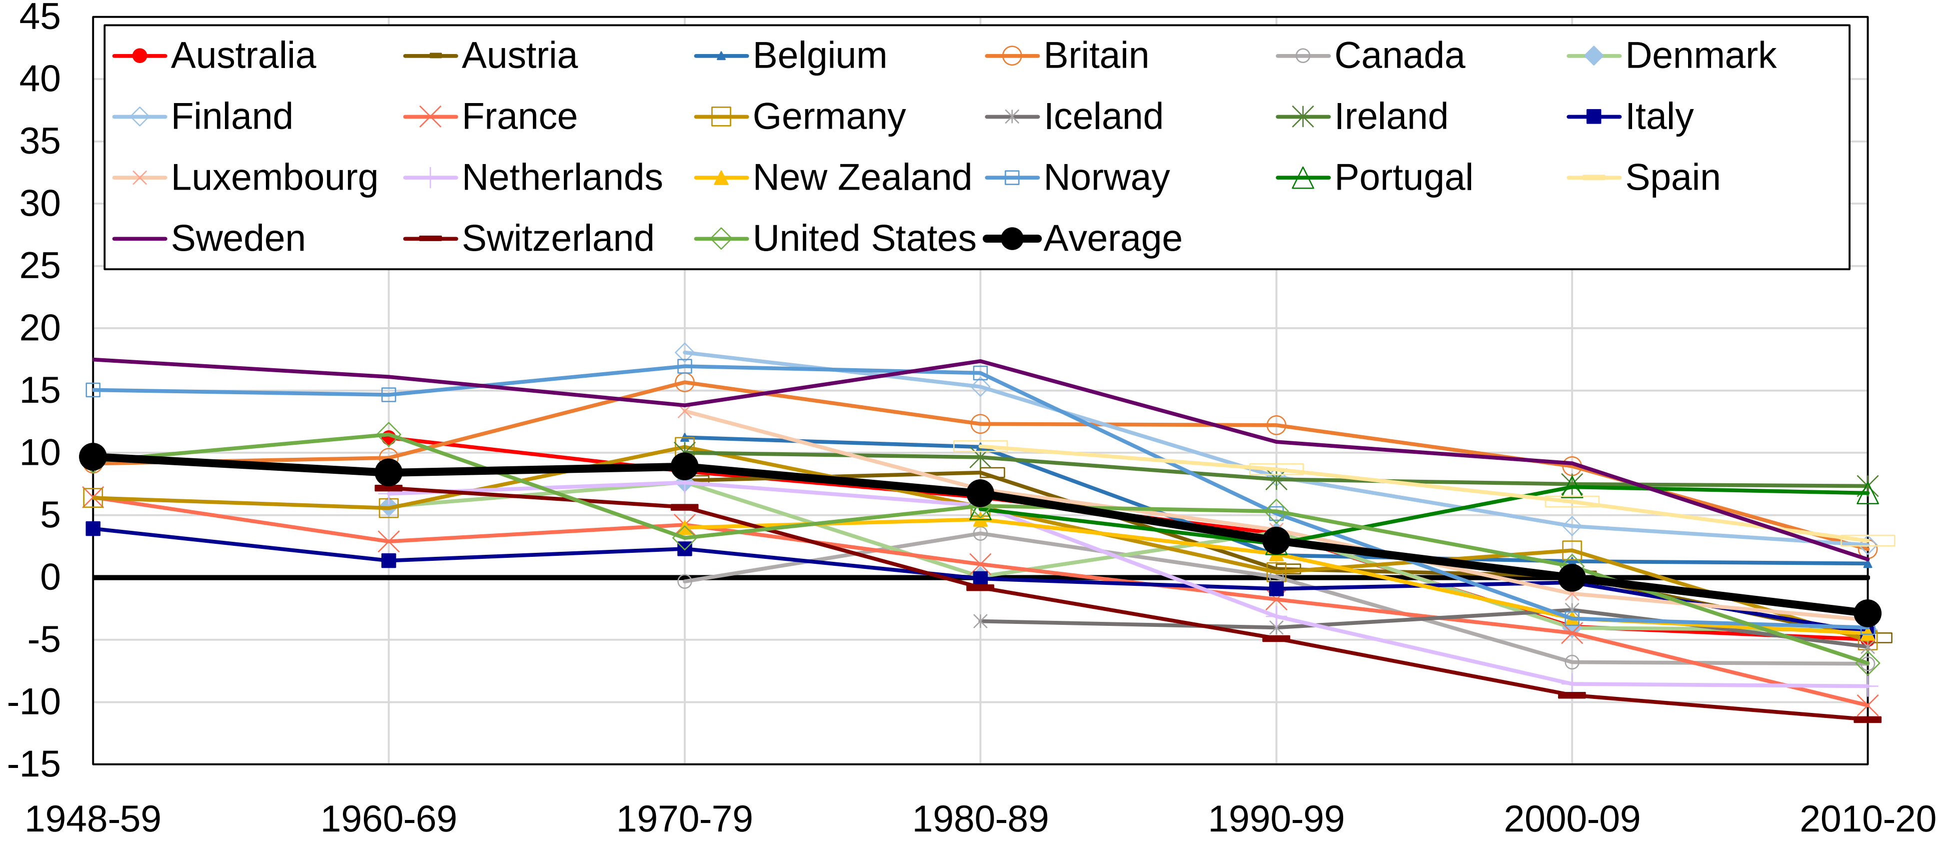

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of lower-educated (bottom 50%) and higher-educated (top 50%) voters voting for left-wing (socialist, social democratic, communist, and green) parties in Western countries, after controlling for income, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available). In nearly all countries, higher-educated voters used to be significantly more likely to vote for right-wing parties and have gradually become more likely to vote for left-wing parties.

Figure A11 - The reversal of educational divides in Western democracies (primary-educated voters)

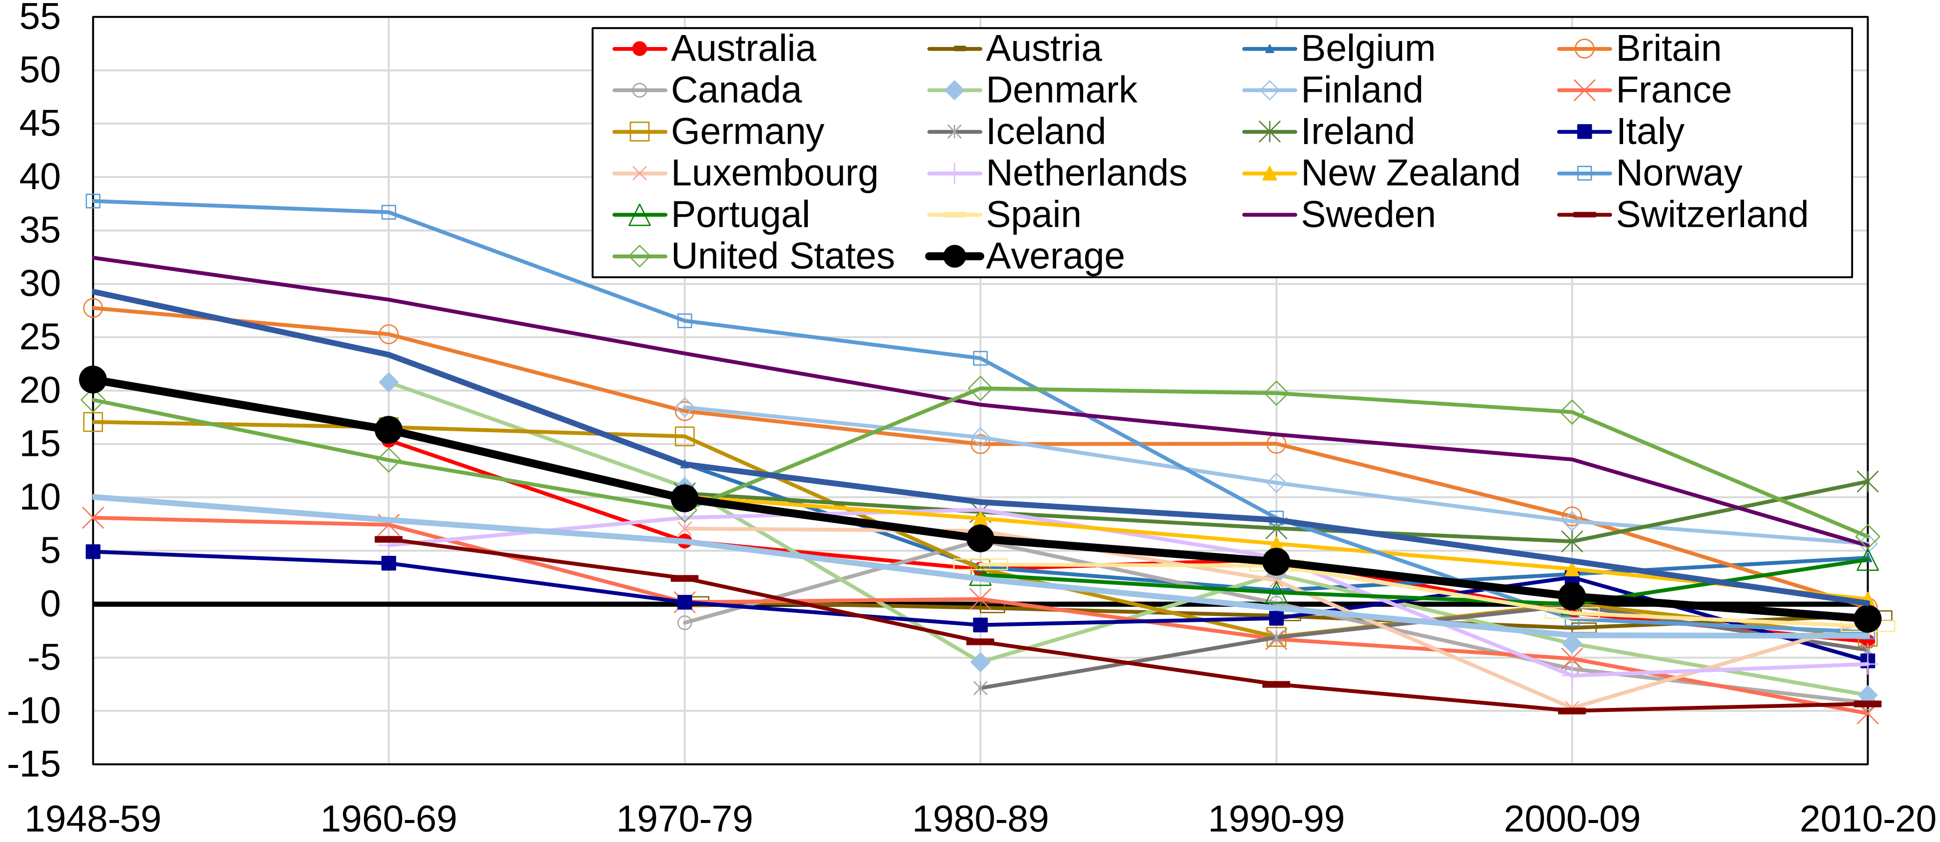

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of primary-educated voters and the share of other voters voting for left-wing (socialist, social democratic, communist, and green) parties in Western countries. In nearly all countries, higher-educated voters used to be significantly more likely to vote for right-wing parties and have gradually become more likely to vote for left-wing parties.

Figure A12 - The reversal of educational divides in Western democracies (primary-educated voters), after controls

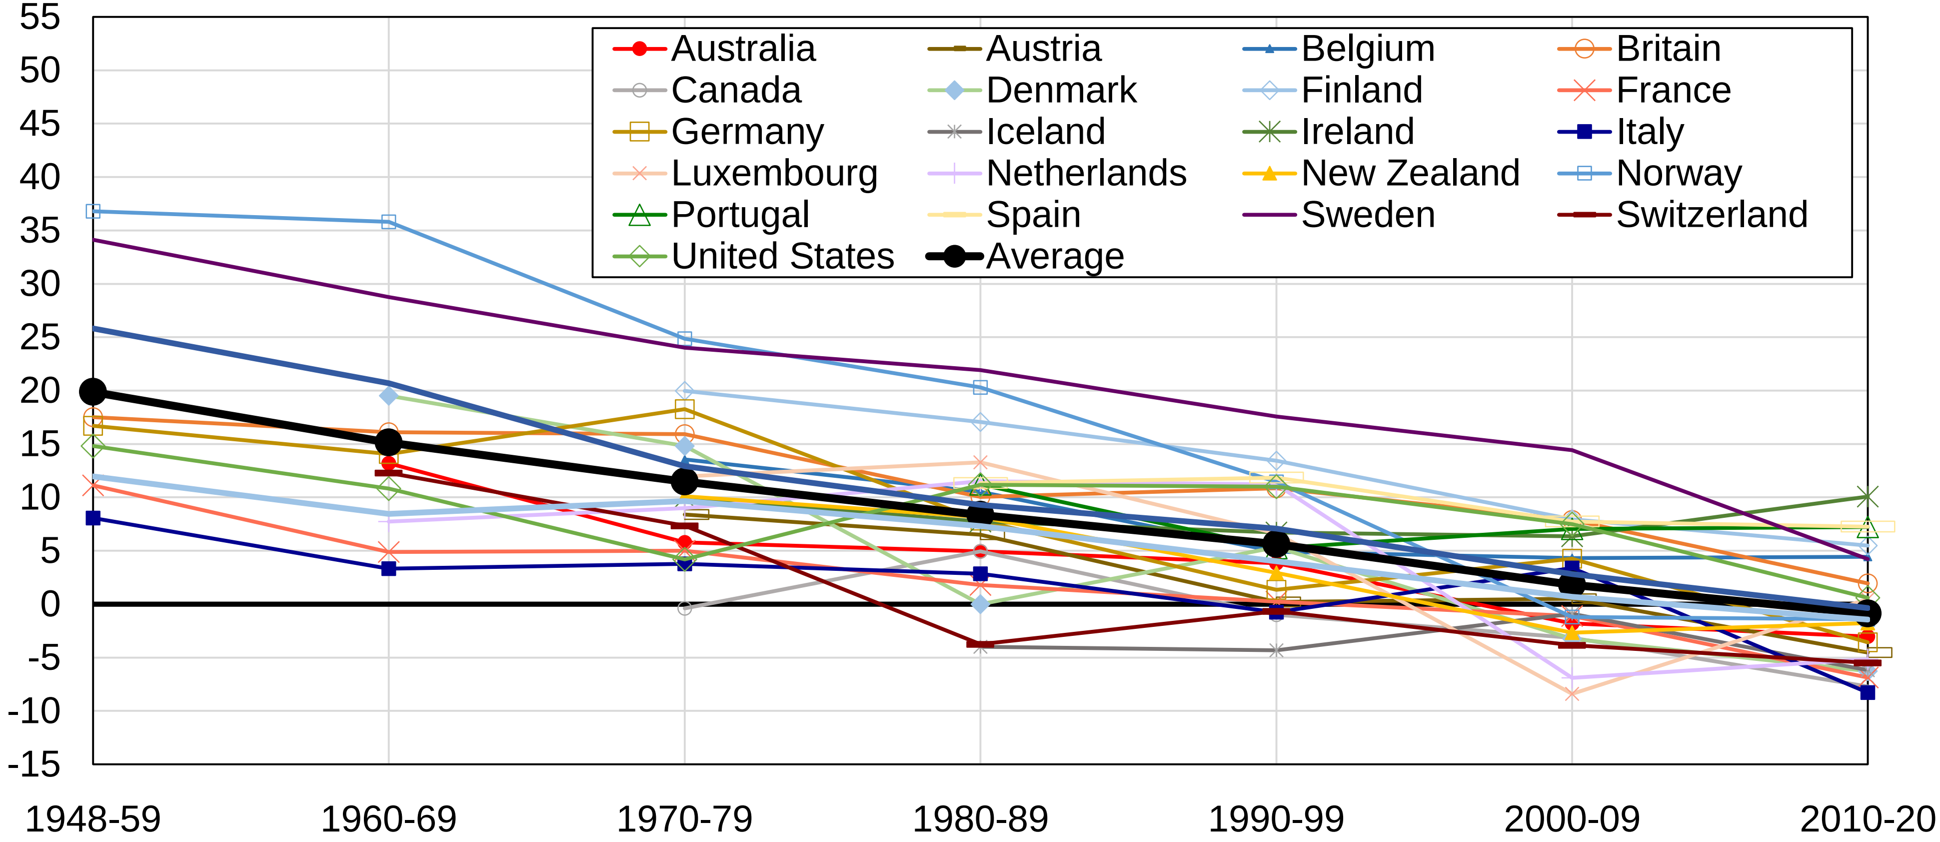

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of primary-educated voters and the share of other voters voting for left-wing (socialist, social democratic, communist, and green) parties in Western countries, after controlling for income, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available). In nearly all countries, higher-educated voters used to be significantly more likely to vote for right-wing parties and have gradually become more likely to vote for left-wing parties.

Figure A13 - The reversal of educational divides in Western democracies (continuous variable)

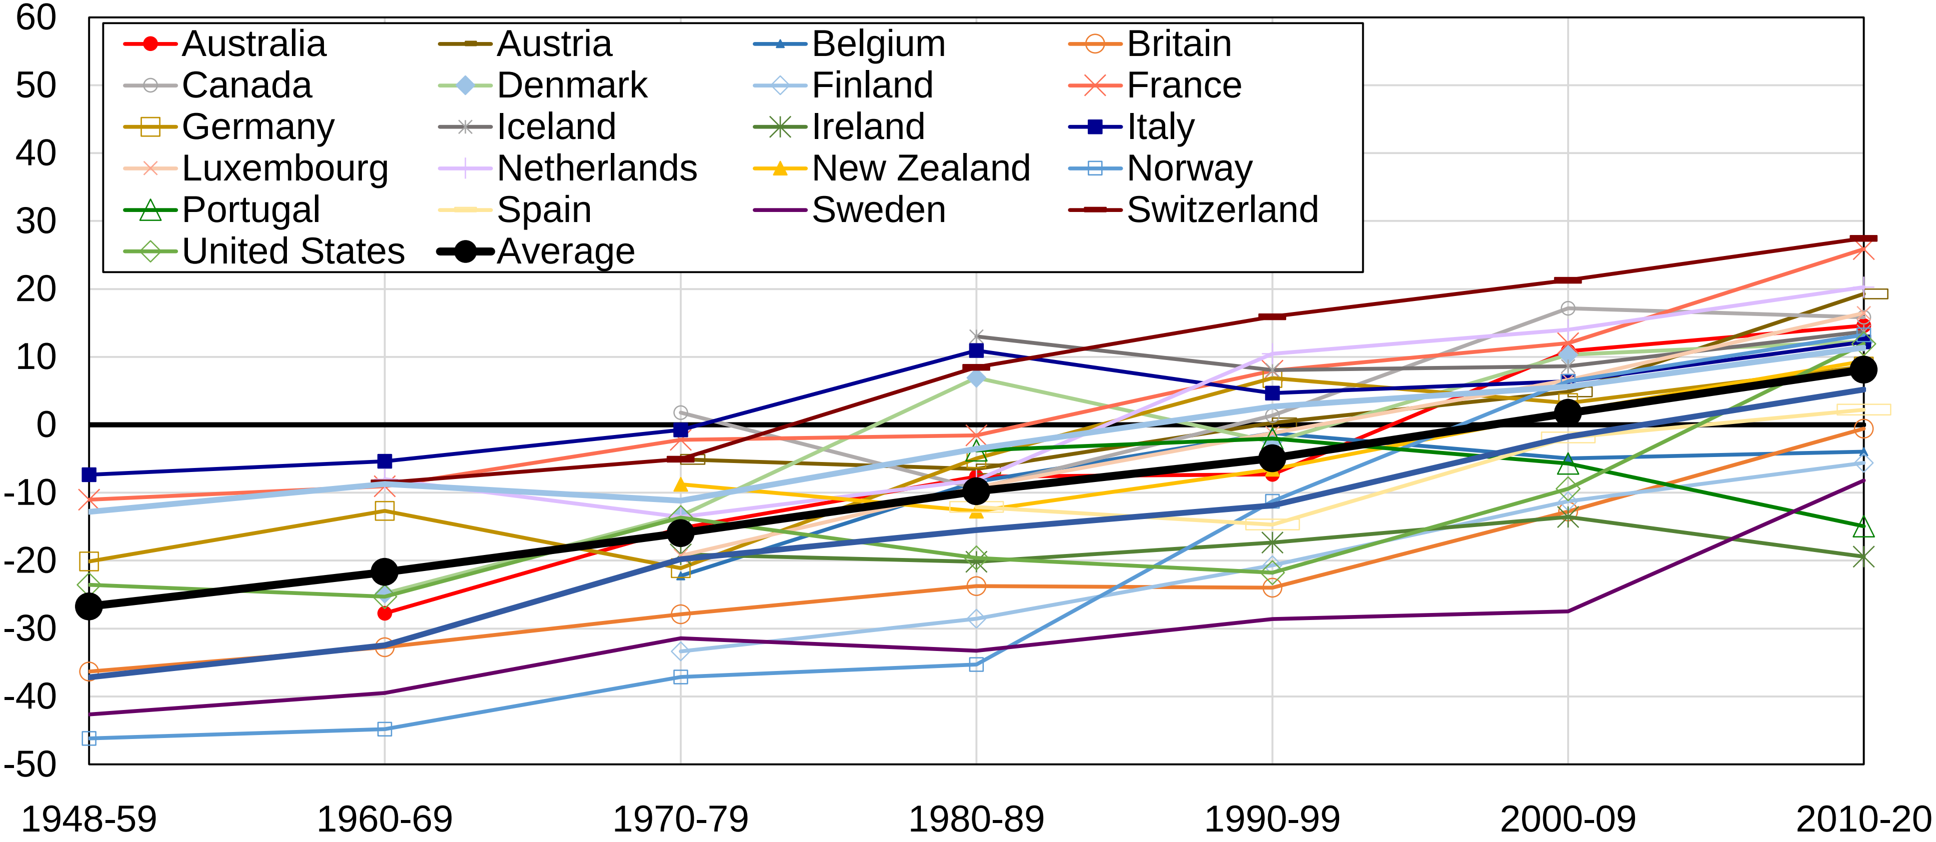

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the marginal effect of the education rank (quantile) of voters on support for left-wing (socialist, social democratic, communist, and green) parties in Western countries. In nearly all countries, higher-educated voters used to be significantly more likely to vote for right-wing parties and have gradually become more likely to vote for left-wing parties.

Figure A14 - The reversal of educational divides in Western democracies (continuous variable), after controls

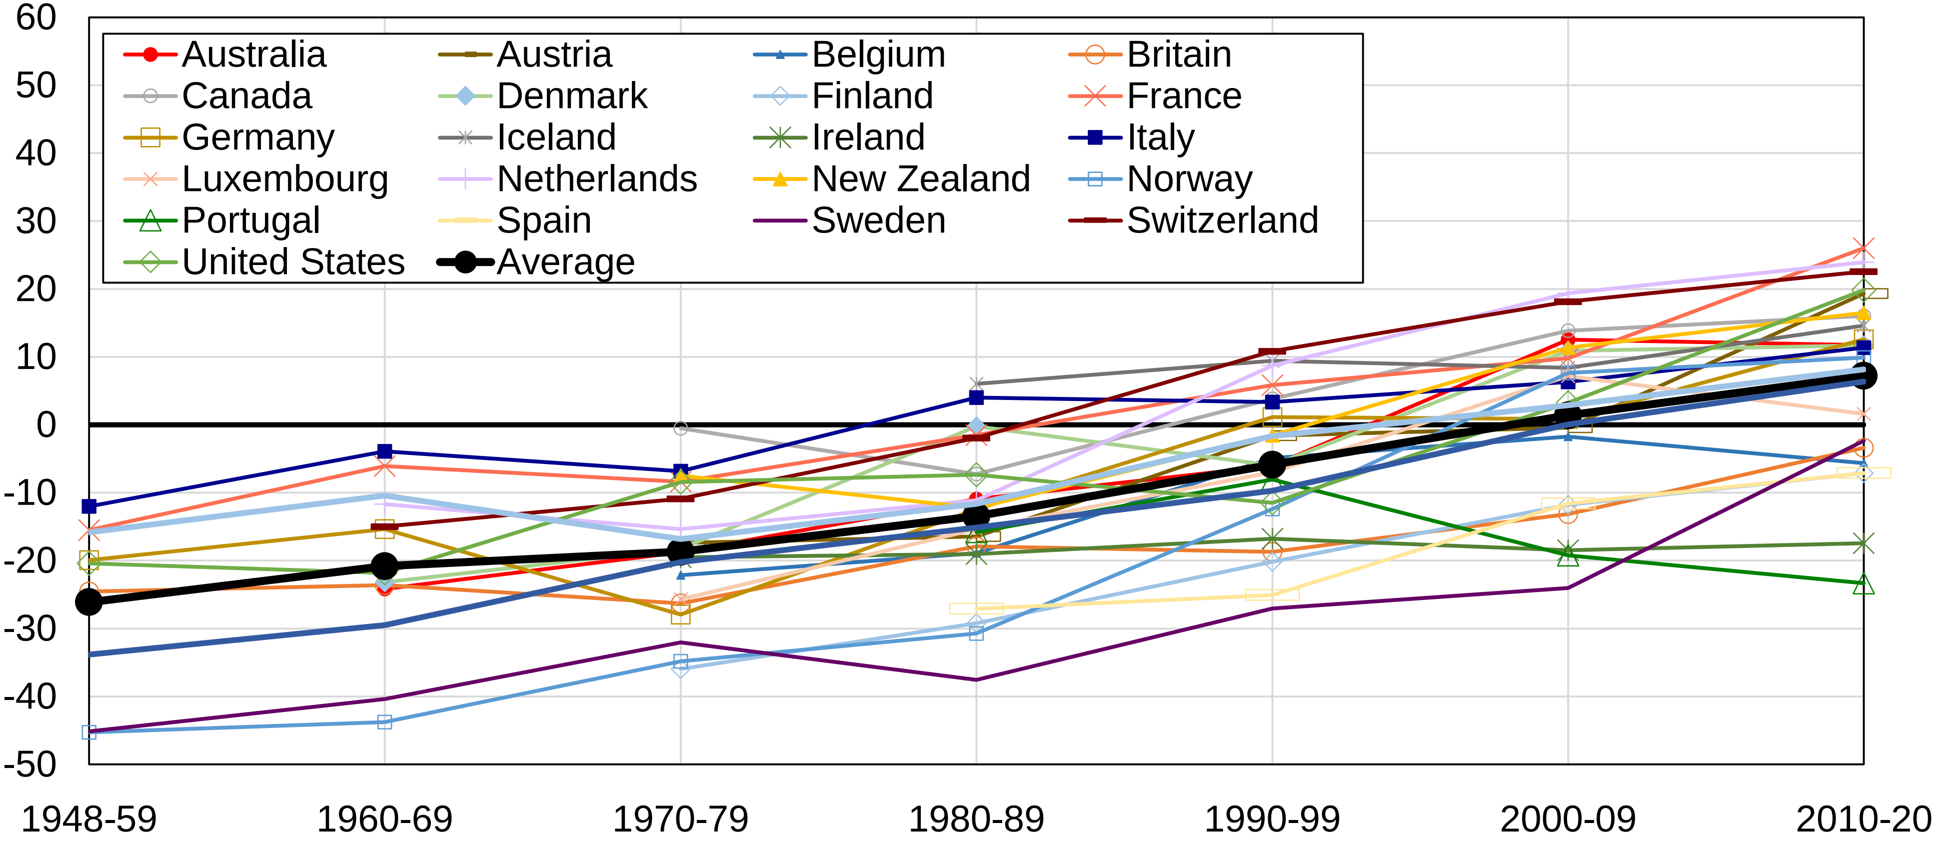

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the marginal effect of the education rank (quantile) of voters on support for left-wing (socialist, social democratic, communist, and green) parties in Western countries, after controlling for income, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available). In nearly all countries, higher-educated voters used to be significantly more likely to vote for right-wing parties and have gradually become more likely to vote for left-wing parties.

Figure A15 - The decline/stability of income divides in Western democracies (top 10%)

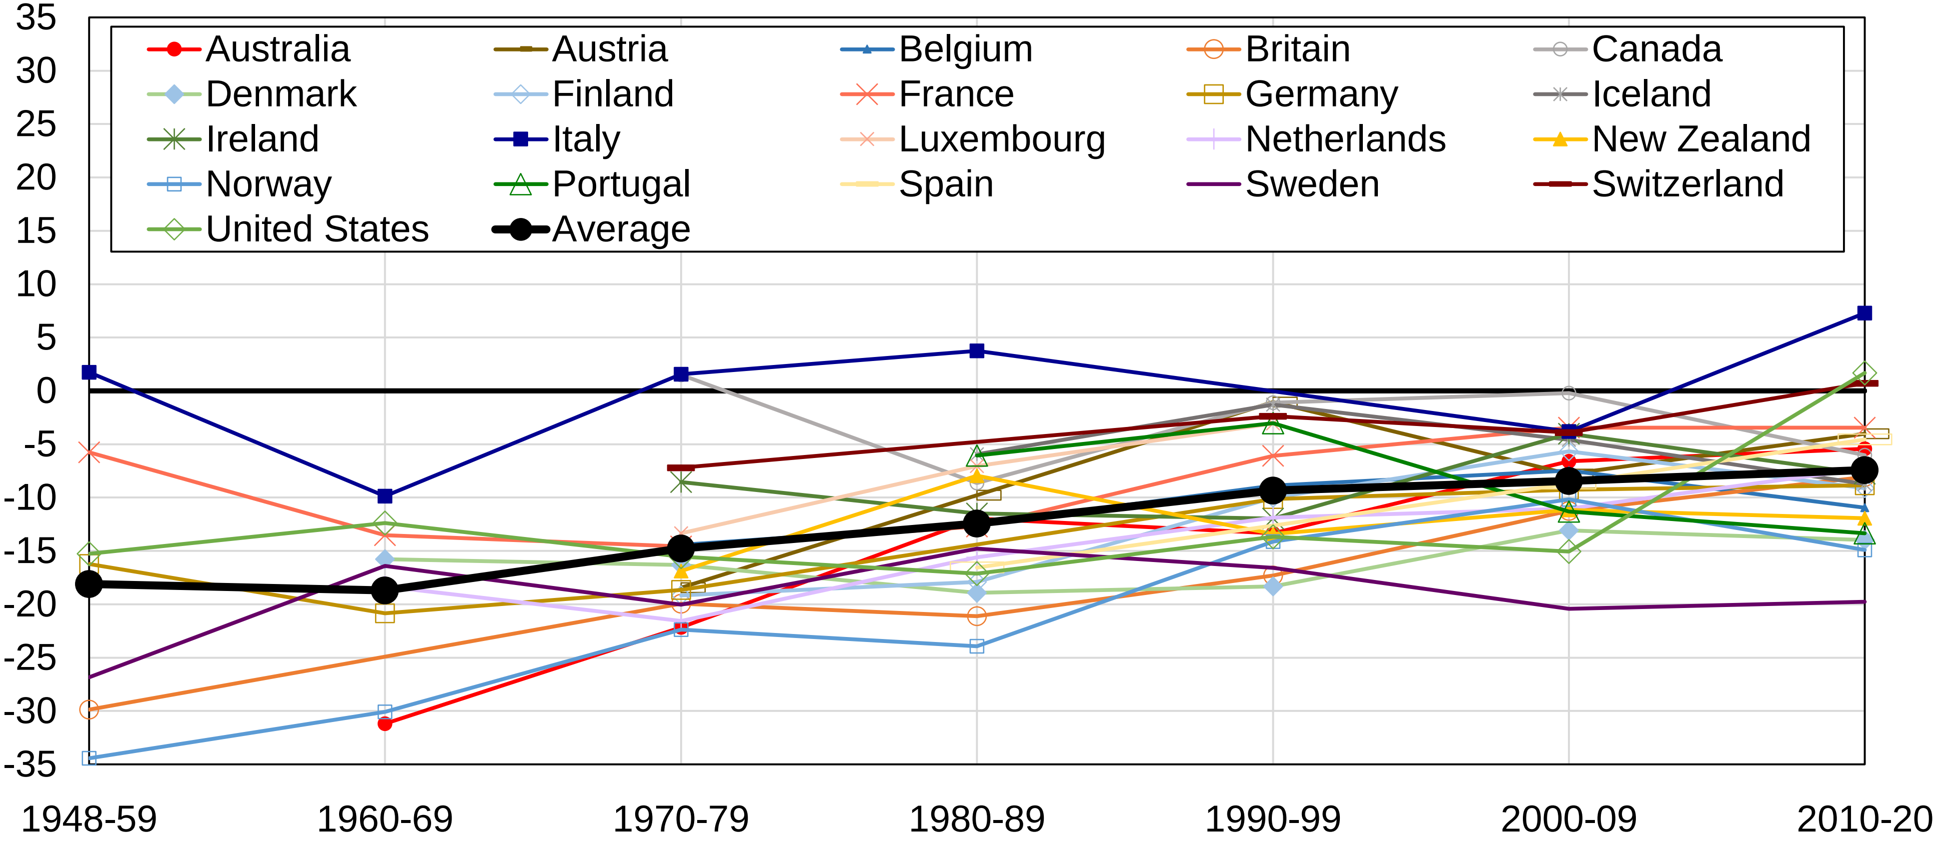

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of high-income (top 10%) and low-income (bottom 90%) voters voting for left-wing (socialist, social democratic, communist, and green) parties in Western countries. In nearly all countries, top-income voters have remained significantly less likely to vote for left-wing parties than low-income voters.

Figure A16 - The decline/stability of income divides in Western democracies (top 10%), after controls

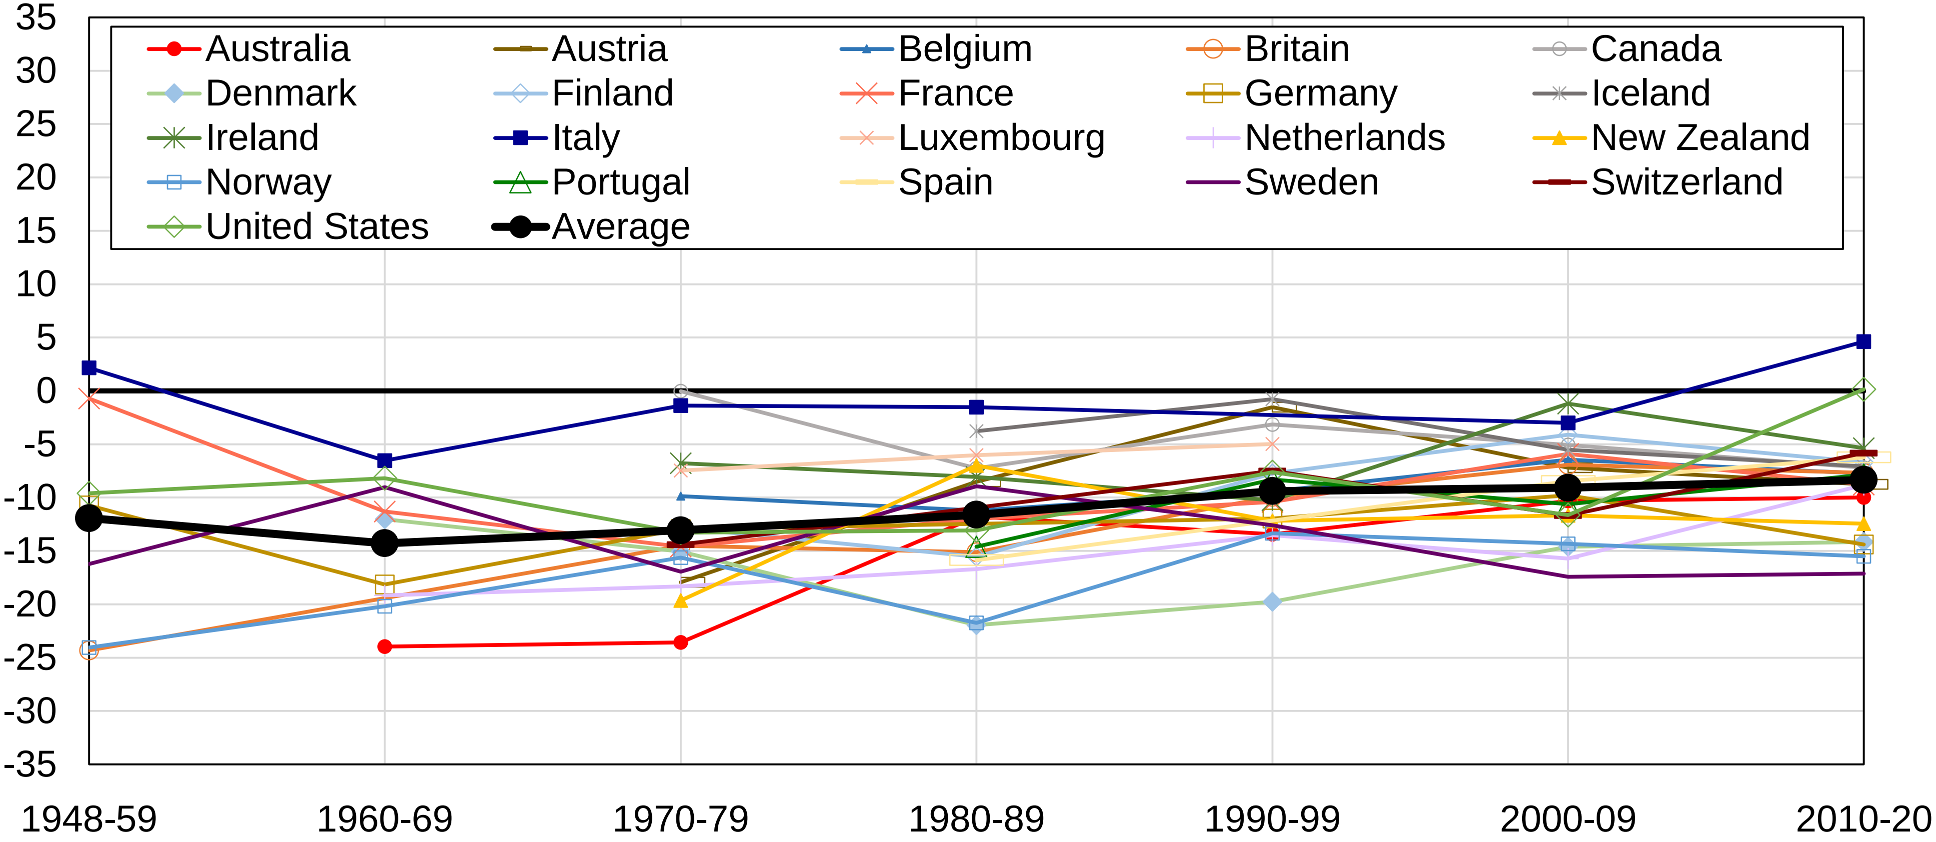

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of high-income (top 10%) and low-income (bottom 90%) voters voting for left-wing (socialist, social democratic, communist, and green) parties in Western countries. In nearly all countries, top-income voters have remained significantly less likely to vote for left-wing parties than low-income voters. Estimates control for education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure A17 - The decline/stability of income divides in Western democracies (bottom 50%)

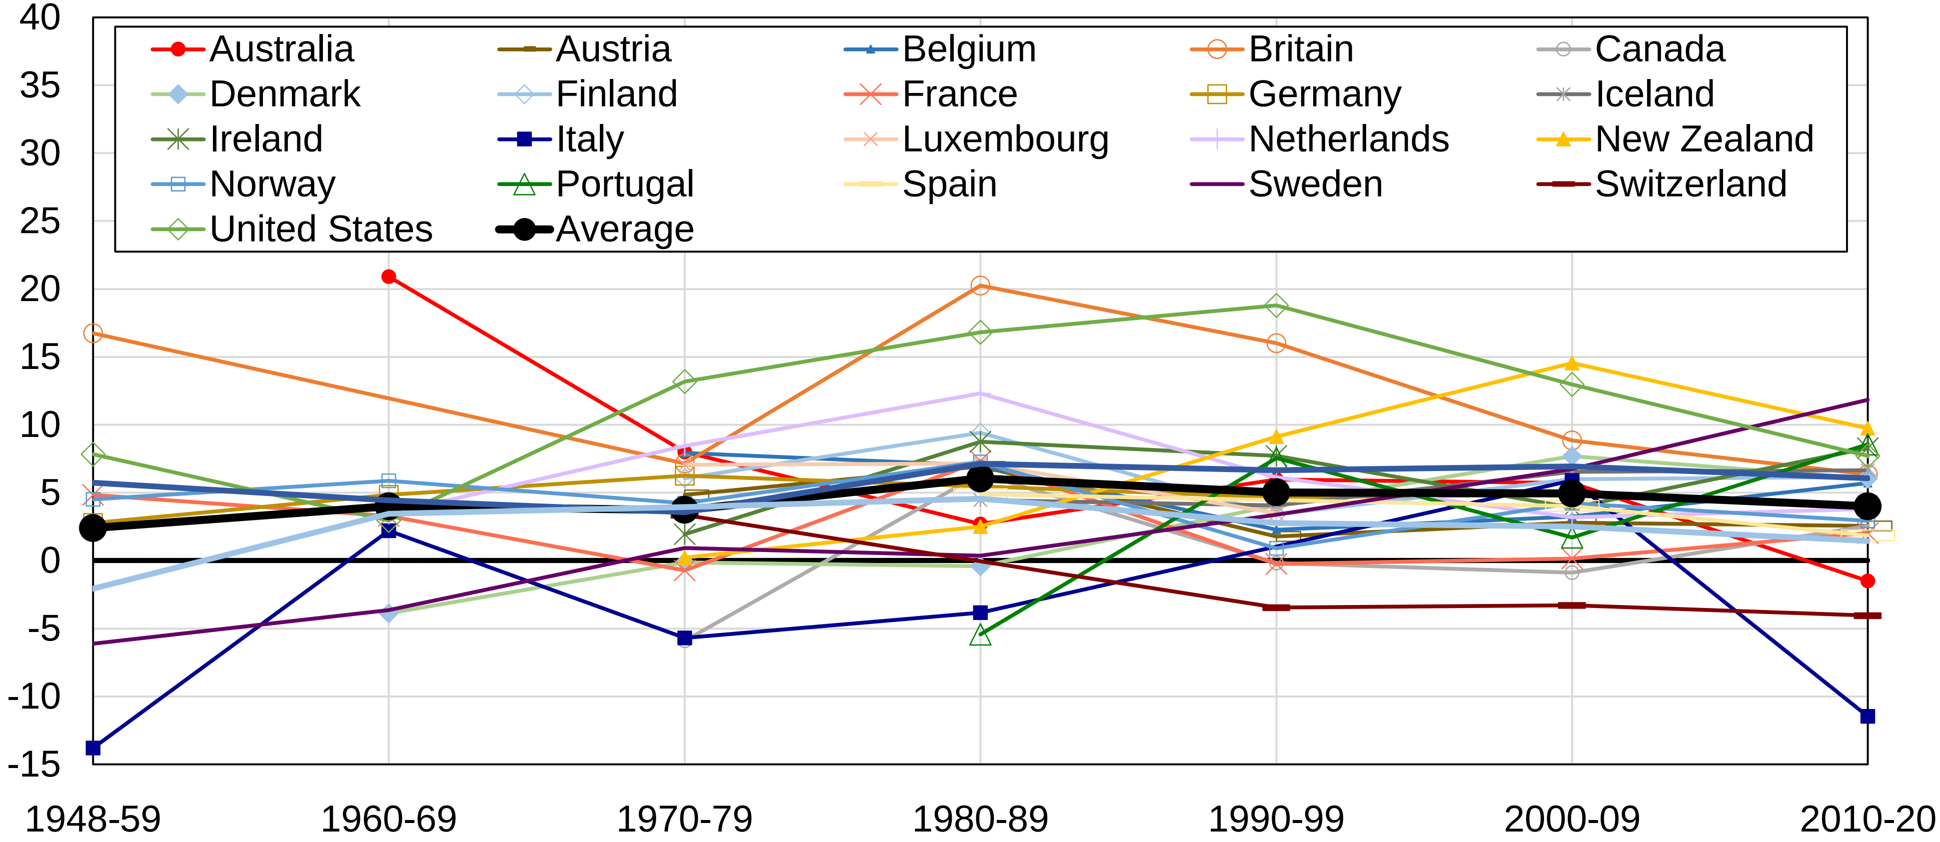

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of low-income (bottom 50%) and top-income (top 50%) voters voting for left-wing (socialist, social democratic, communist, and green) parties in Western countries. In nearly all countries, top-income voters have remained significantly less likely to vote for left-wing parties than low-income voters.

Figure A18 - The decline/stability of income divides in Western democracies (bottom 50%), after controls

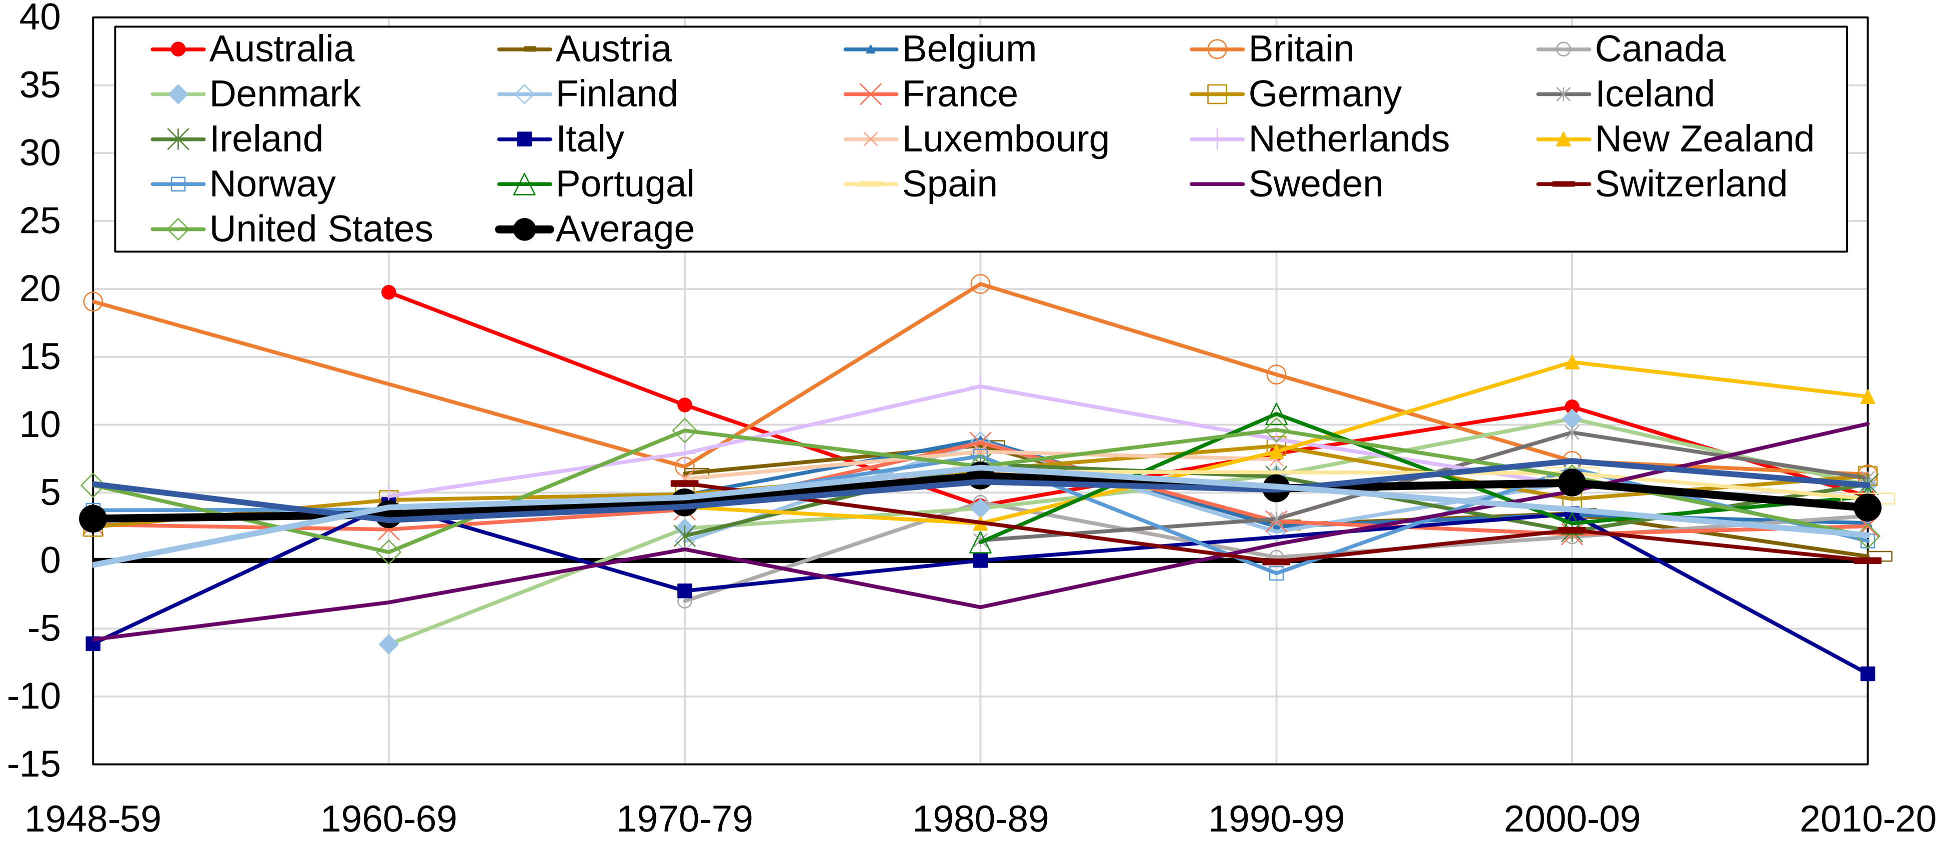

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of low-income (bottom 50%) and top-income (top 50%) voters voting for left-wing (socialist, social democratic, communist, and green) parties in Western countries. In nearly all countries, top-income voters have remained significantly less likely to vote for left-wing parties than low-income voters. Estimates control for education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure A19 - The decline/stability of income divides in Western democracies (continuous variable)

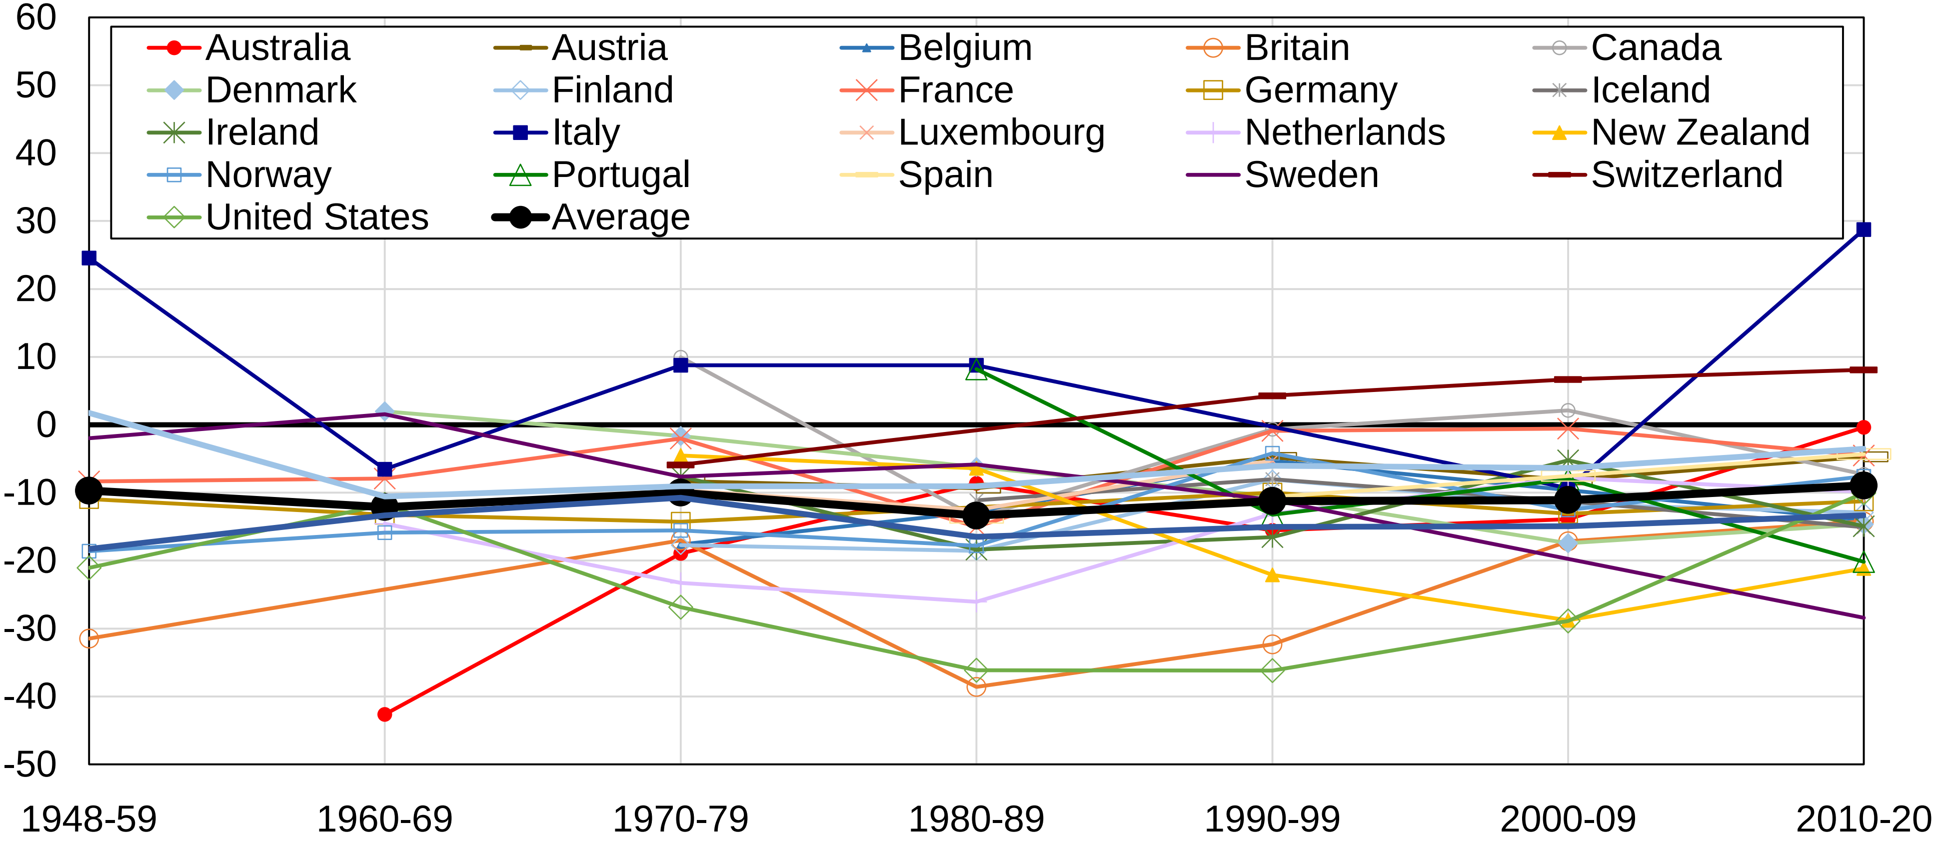

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the marginal effect of the income rank (quantile) of voters on support for left-wing (socialist, social democratic, communist, and green) parties in Western countries. In nearly all countries, top-income voters have remained significantly less likely to vote for left-wing parties than low-income voters.

Figure A20 - The decline/stability of income divides in Western democracies (continuous variable), after controls

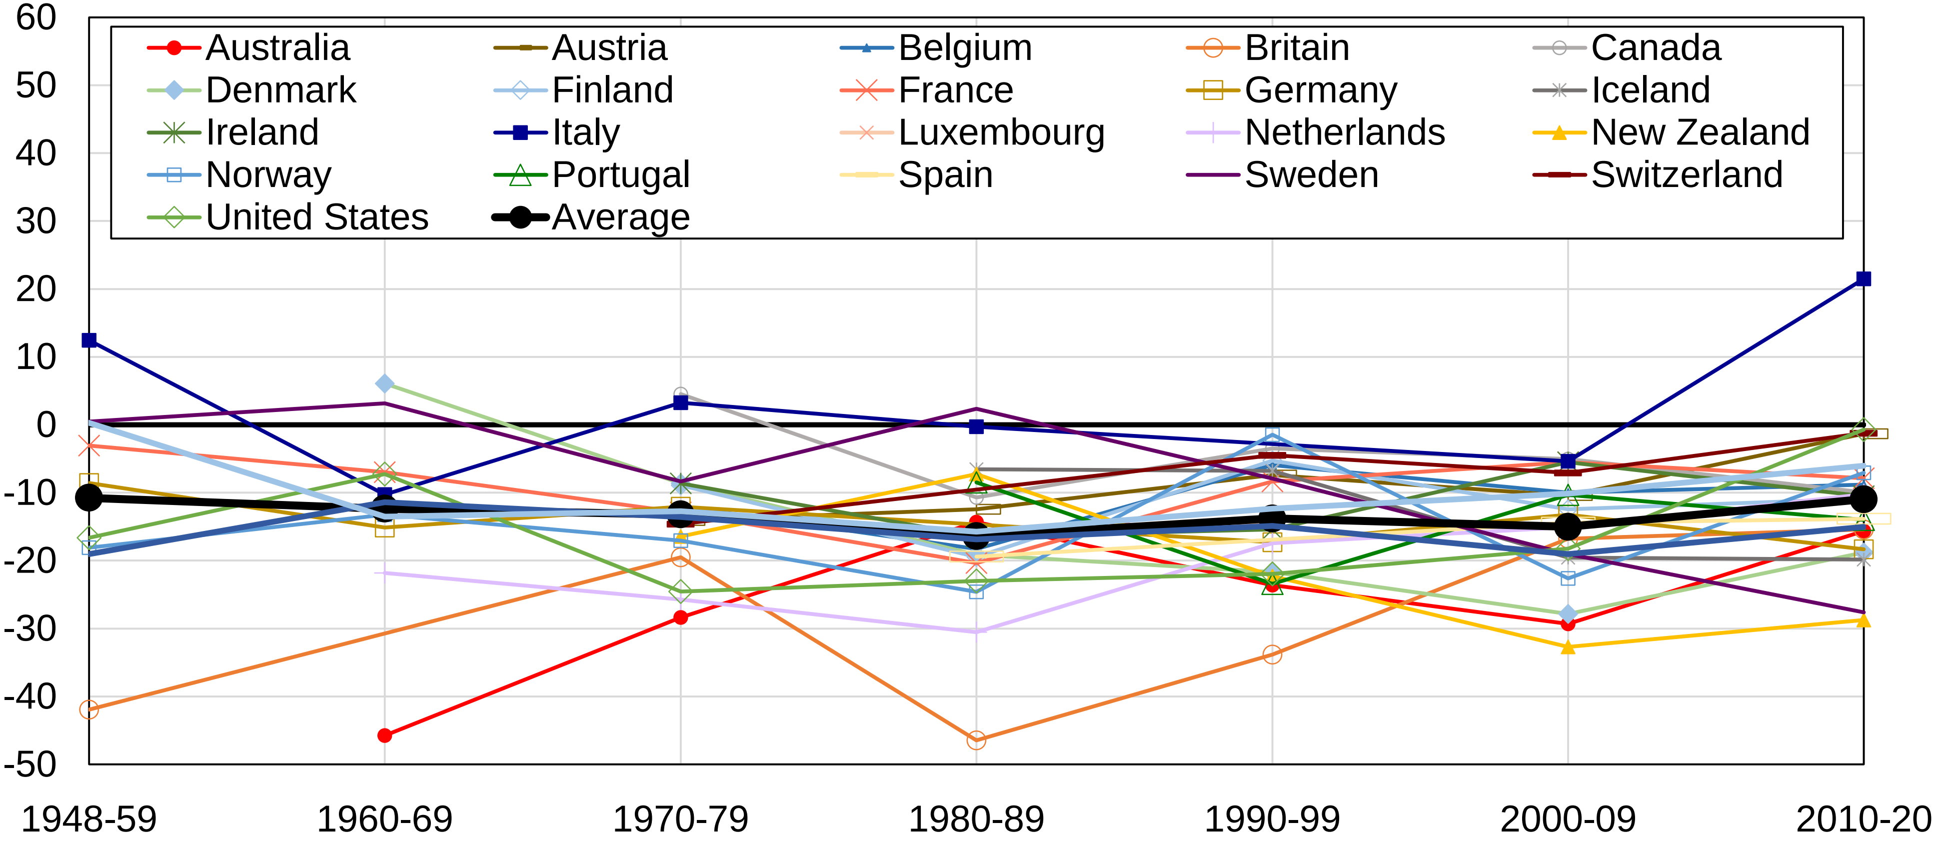

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the marginal effect of the income rank (quantile) of voters on support for left-wing (socialist, social democratic, communist, and green) parties in Western countries. In nearly all countries, top-income voters have remained significantly less likely to vote for left-wing parties than low-income voters. Estimates control for education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure A21 - The disconnection of income and education in Western democracies (including/excluding green parties)

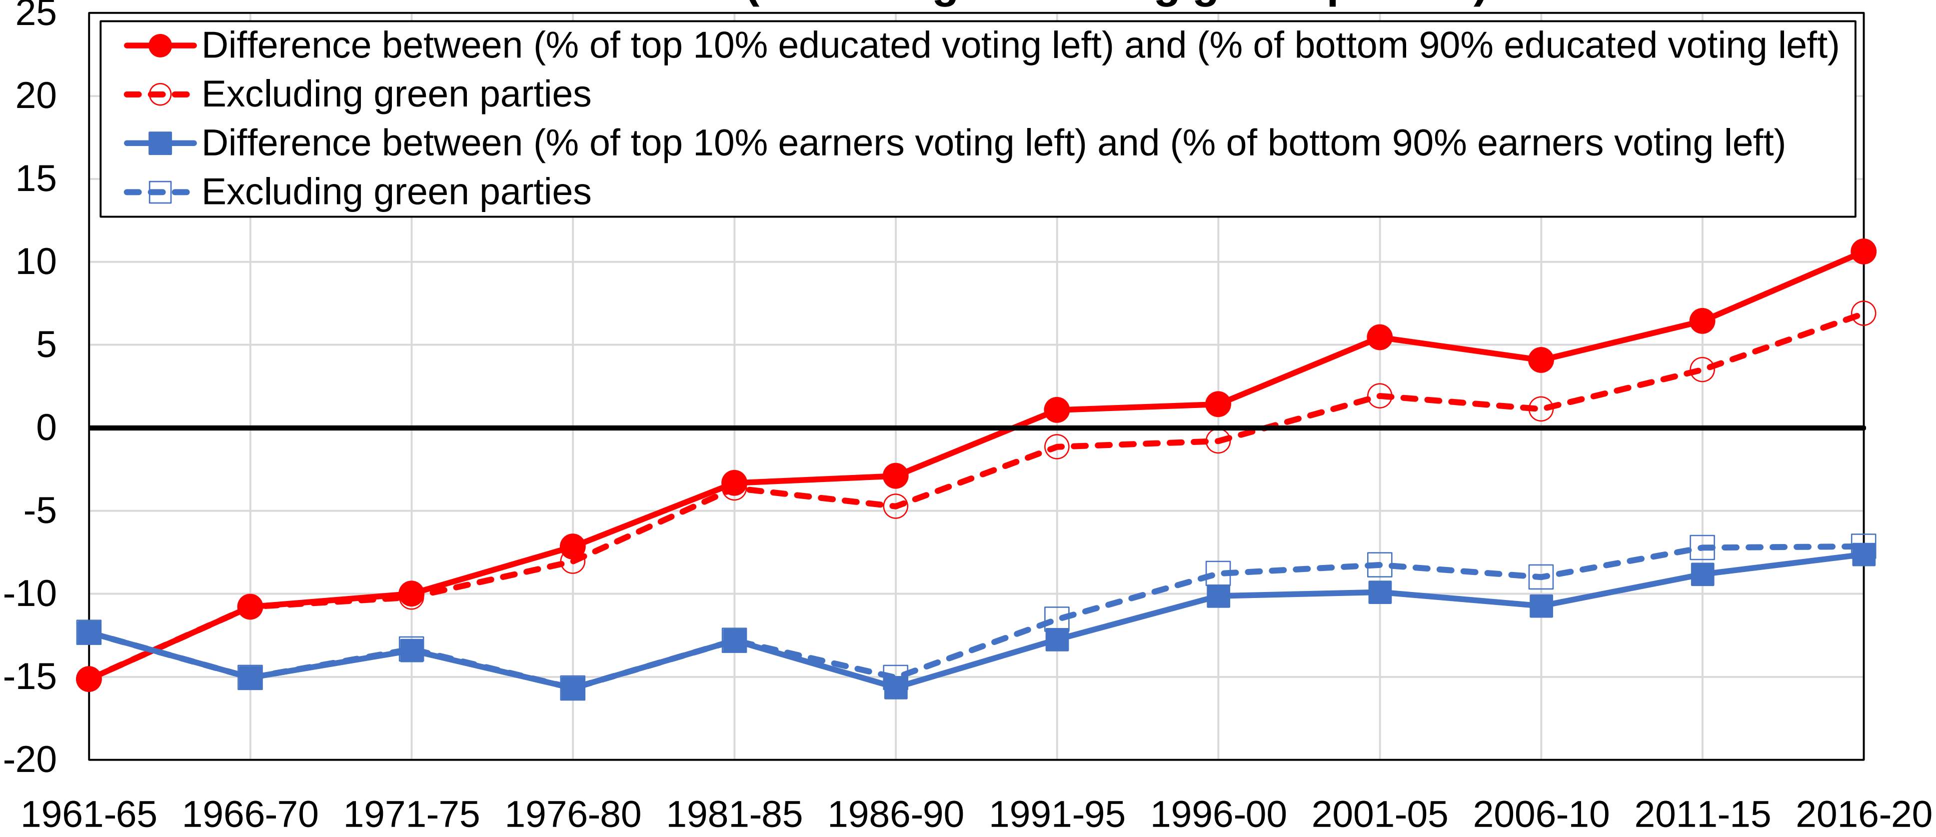

Source: authors' computations using the World Political Cleavages and Inequality Database.

Note: in the 1960s, both higher-educated and high-income voters were less likely to vote for left-wing (social democratic / socialist / communist / green / other left-wing) parties than lower-educated and low-income voters by more than 10 percentage points. The left vote has gradually become associated with higher education voters, giving rising to a remarkable divergence of the effects of income and education on the vote. Figures correspond to five-year averages for Australia, Britain, Canada, Denmark, France, Germany, Italy, the Netherlands, Norway, Sweden, Switzerland, and the US. The estimates are presented after controlling for income/education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure A22 - The disconnection of income and education cleavages in Western democracies (incl./excl. green parties), unbalanced panel

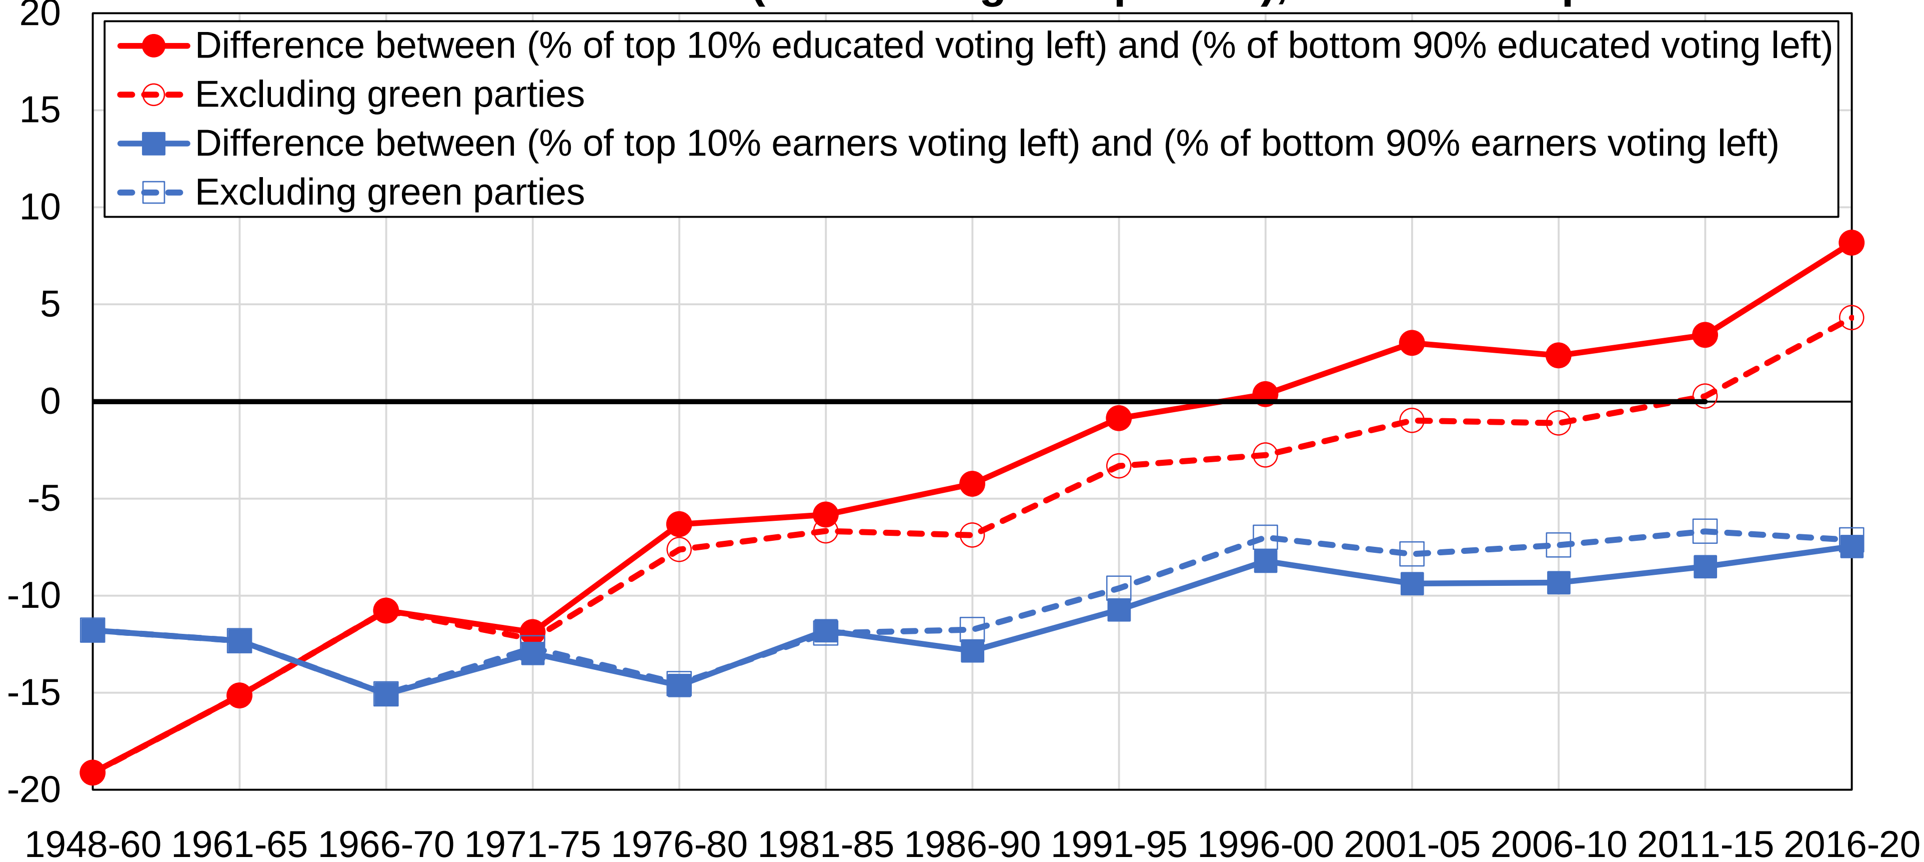

Source: authors' computations using the World Political Cleavages and Inequality Database.

Note: in the 1960s, both higher-educated and high-income voters were less likely to vote for left-wing (social democratic / socialist / communist / green / other left-wing) parties than lower-educated and low-income voters by more than 10 percentage points. The left vote has gradually become associated with higher education voters, giving rising to a remarkable divergence of the effects of income and education on the vote. Figures correspond to five-year averages over all countries available for a given time period (unbalanced panel of all 21 Western democracies). The estimates are presented before and after controlling for income/education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure A23 - Support for left-wing parties (excluding Greens) among top 10% educated voters, after controls

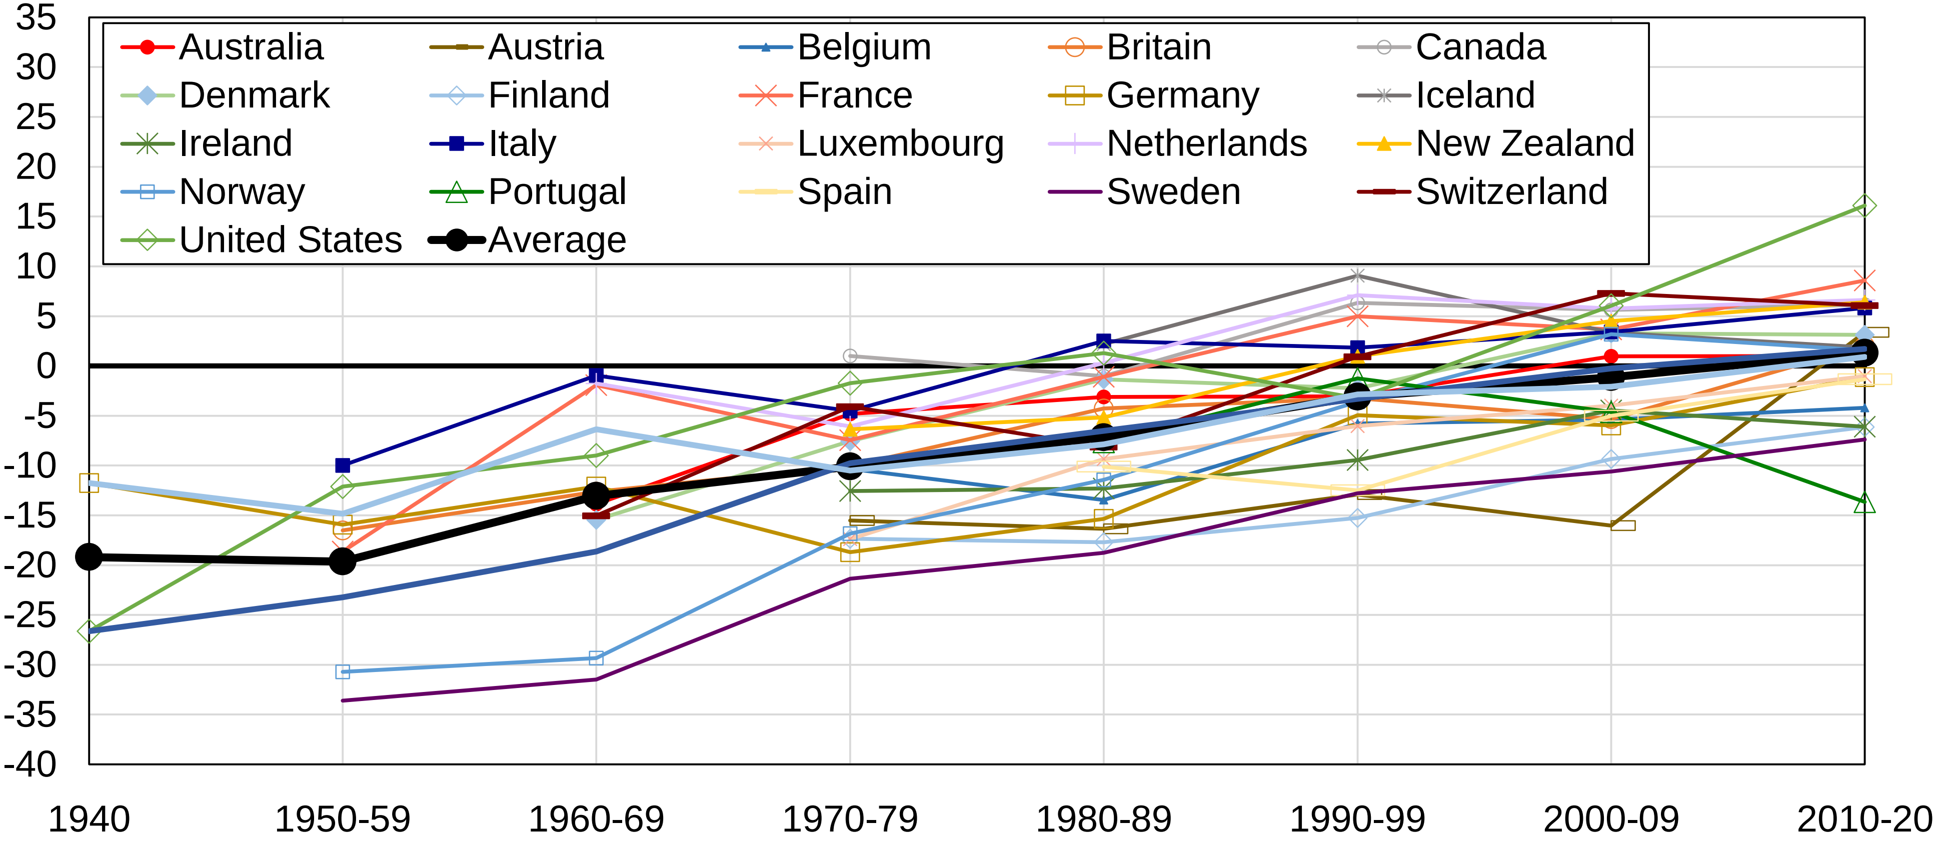

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of higher-educated (top 10%) and lower-educated (bottom 90%) voters voting for left-wing parties (excluding Greens) in Western countries, after controlling for income, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available). In nearly all countries, higher-educated voters used to be significantly more likely to vote for right-wing parties and have gradually become more likely to vote for left-wing parties.

Figure A24 - Support for left-wing parties (excluding Greens) among top 10% income voters, after controls

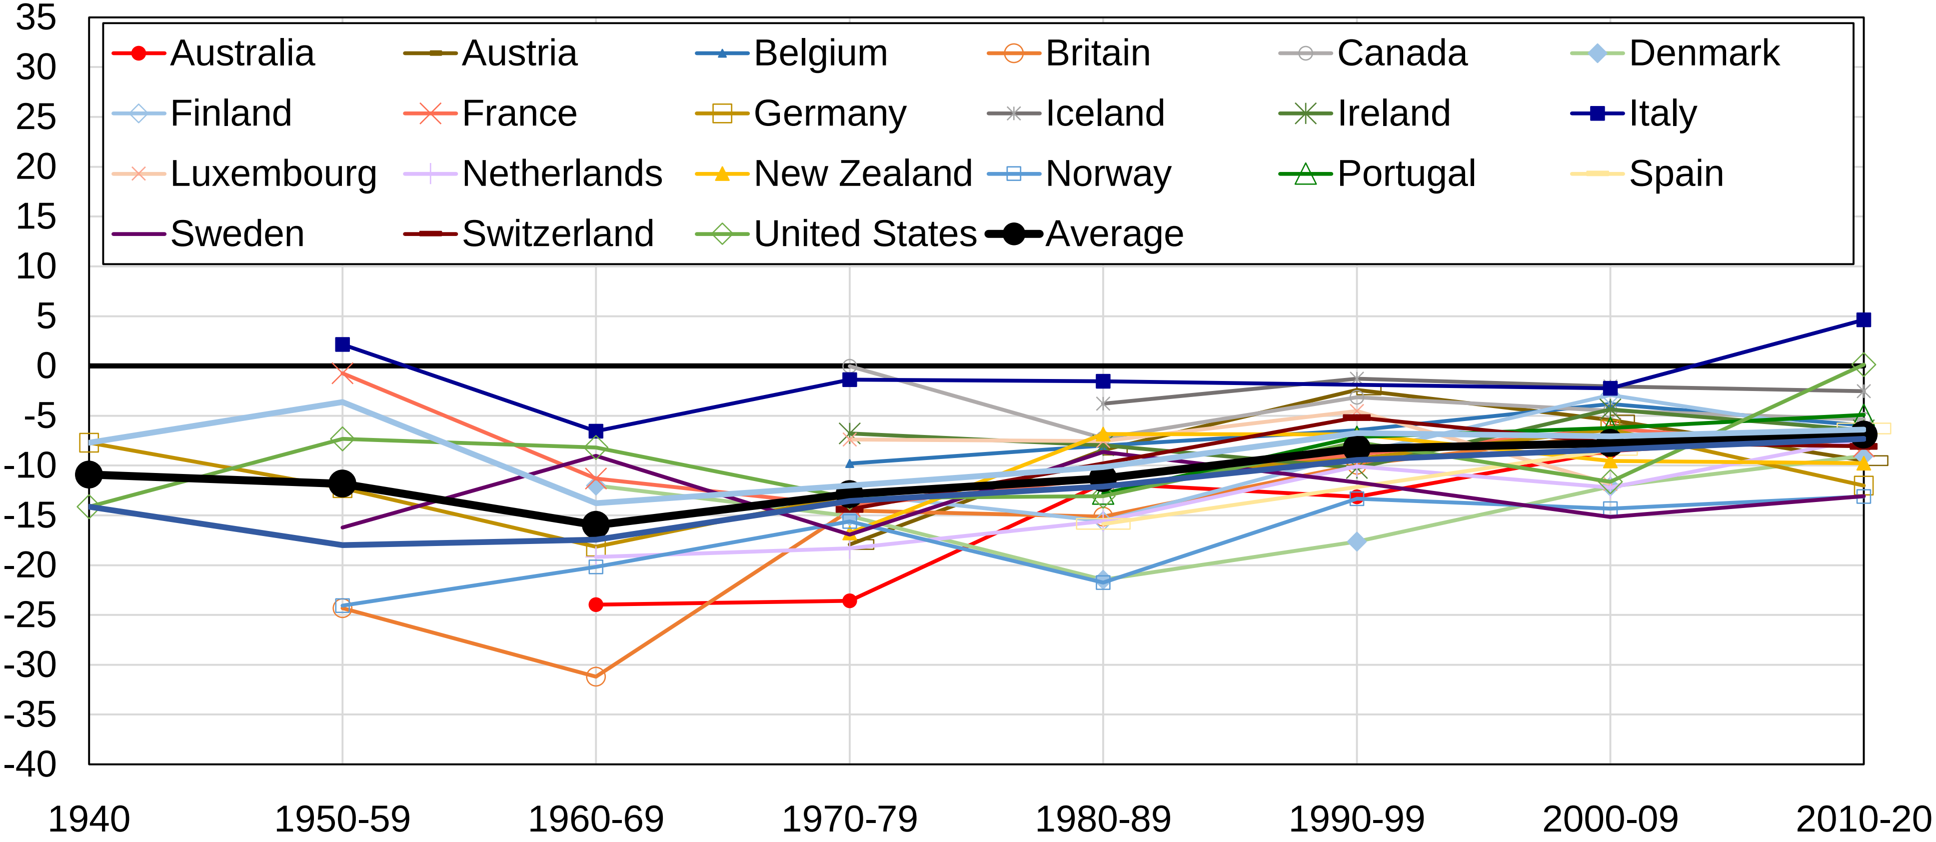

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of high-income (top 10%) and low-income (bottom 90%) voters voting for left-wing parties (excluding Greens) in Western countries. In nearly all countries, top-income voters have remained significantly less likely to vote for left-wing parties than low-income voters. Estimates control for education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure A25 - The reversal of educational divides in Western democracies: decomposition by party family (left-wing parties)

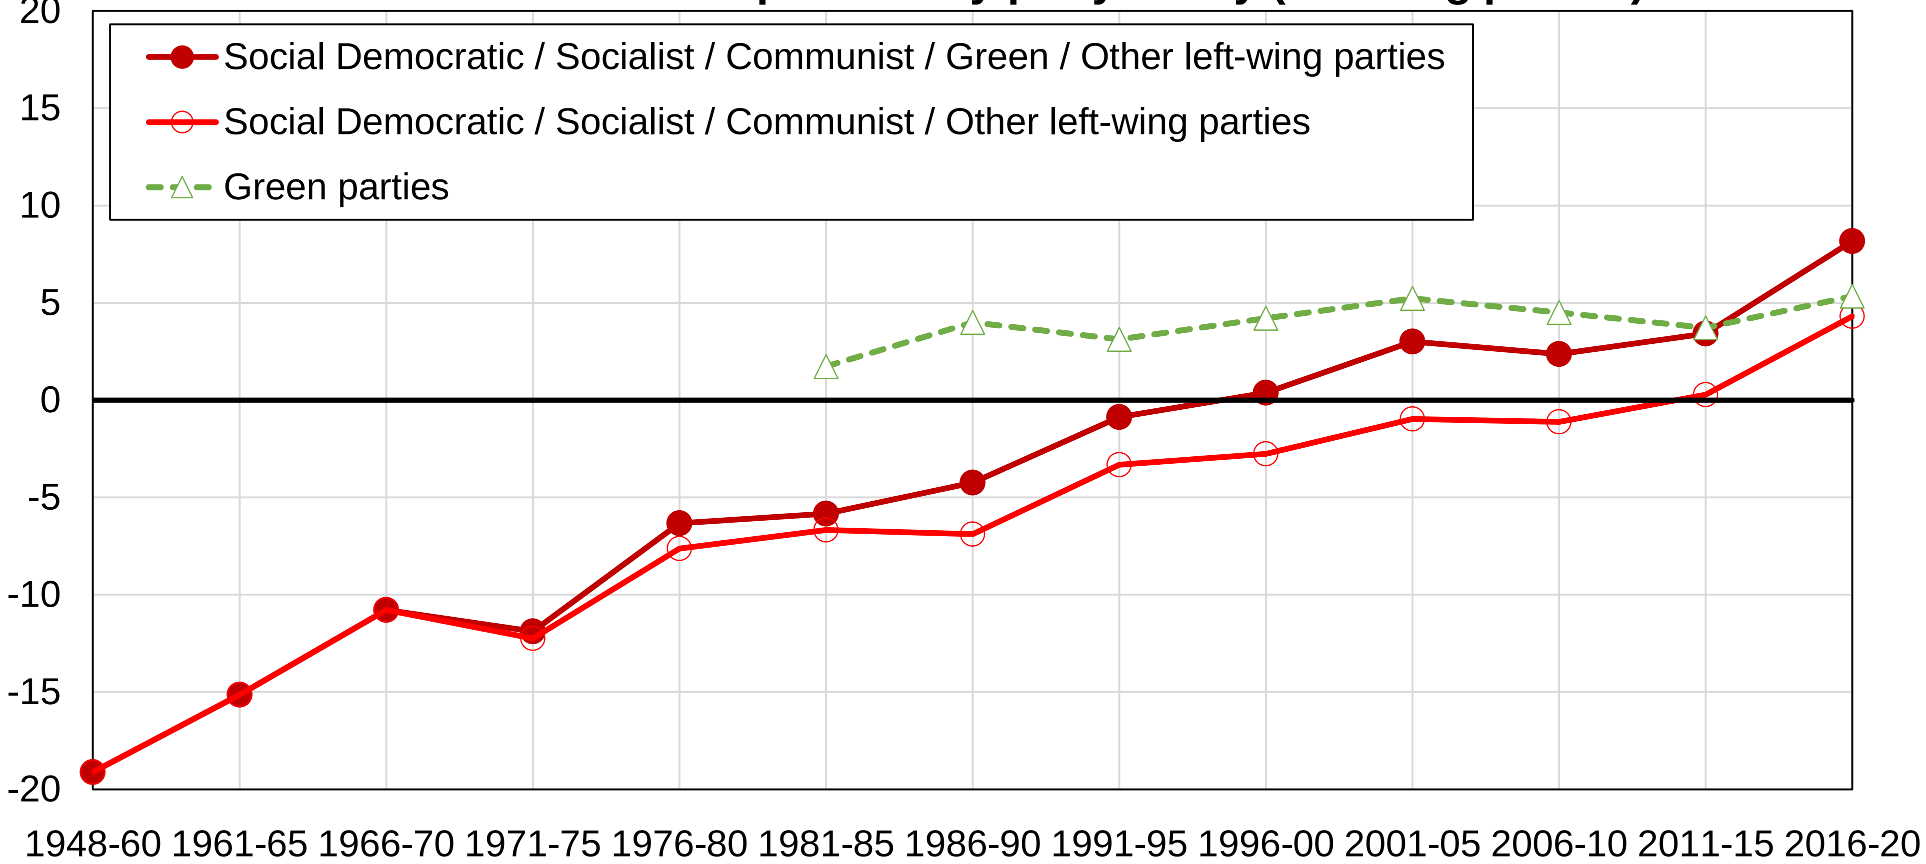

Source: authors' computations using the World Political Cleavages and Inequality Database.

Note: the figure represents the difference between the share of top 10% educated voters and the share of bottom 90% educated voters voting for specific families of parties. Figures correspond to five-year averages over all countries available for a given time period (unbalanced panel of all 21 Western democracies). The estimates are presented after controlling for income, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure A26 - The reversal of educational divides in Western democracies: decomposition by party family (right-wing parties)

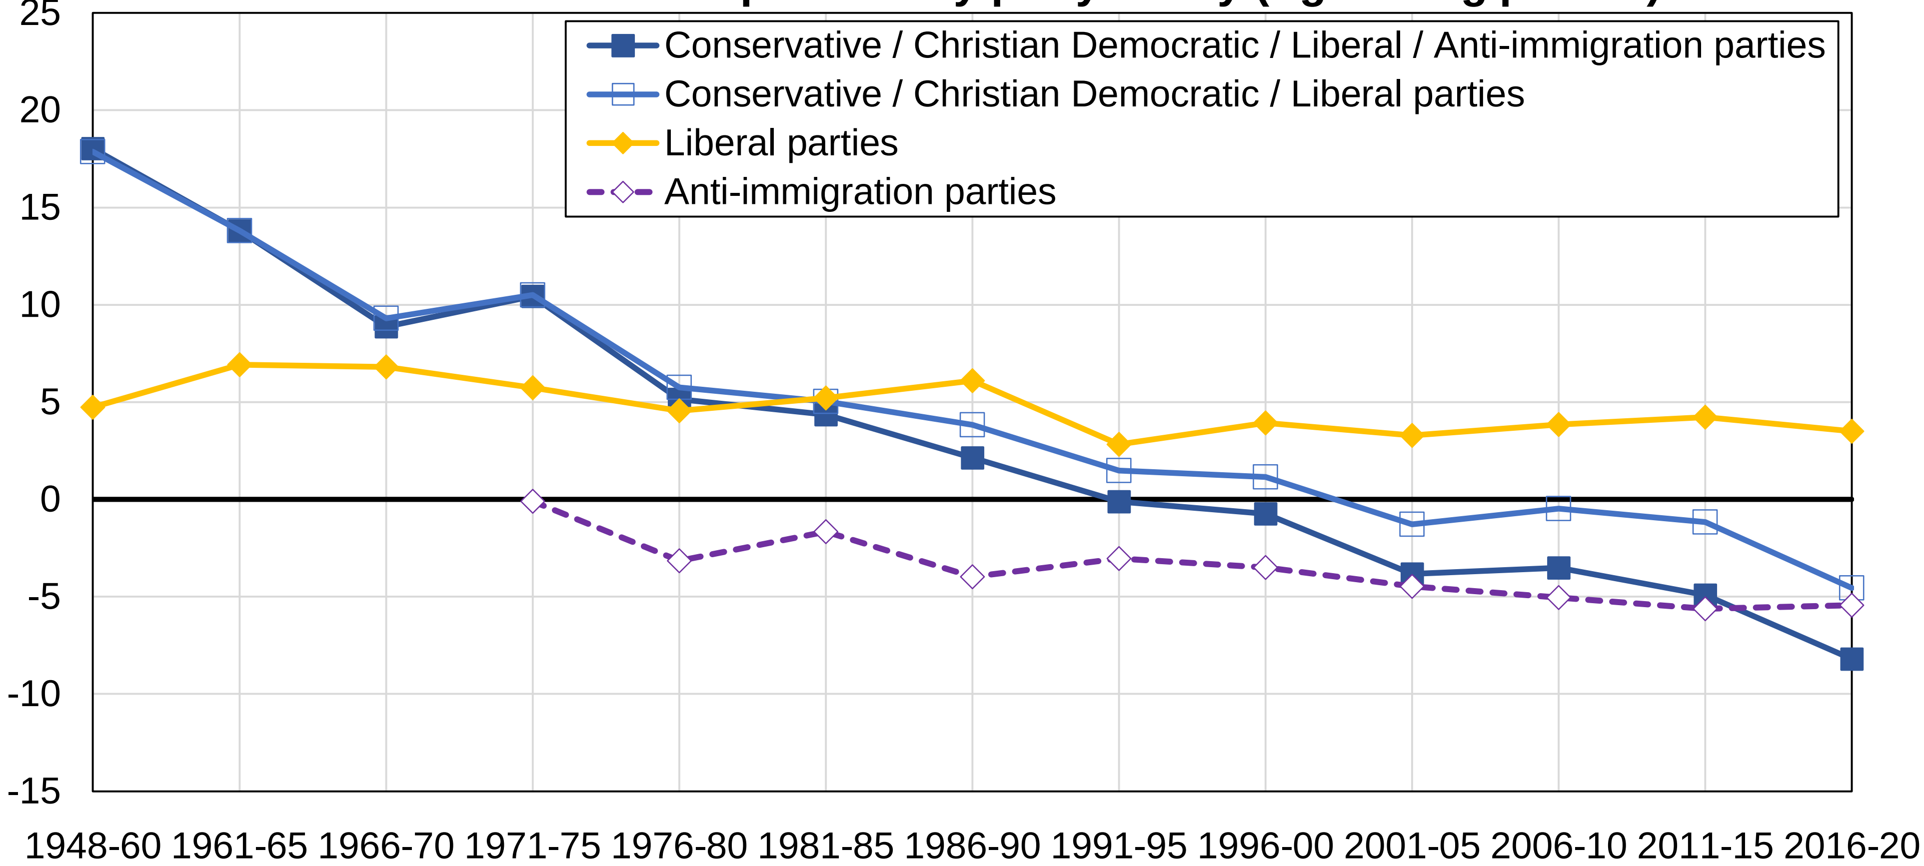

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of top 10% educated voters and the share of bottom 90% educated voters voting for specific families of parties. Figures correspond to five-year averages over all countries available for a given time period (unbalanced panel of all 21 Western democracies). The estimates are presented after controlling for income, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure A27 - The decline/stability of income divides in Western democracies: decomposition by party family (left-wing parties)

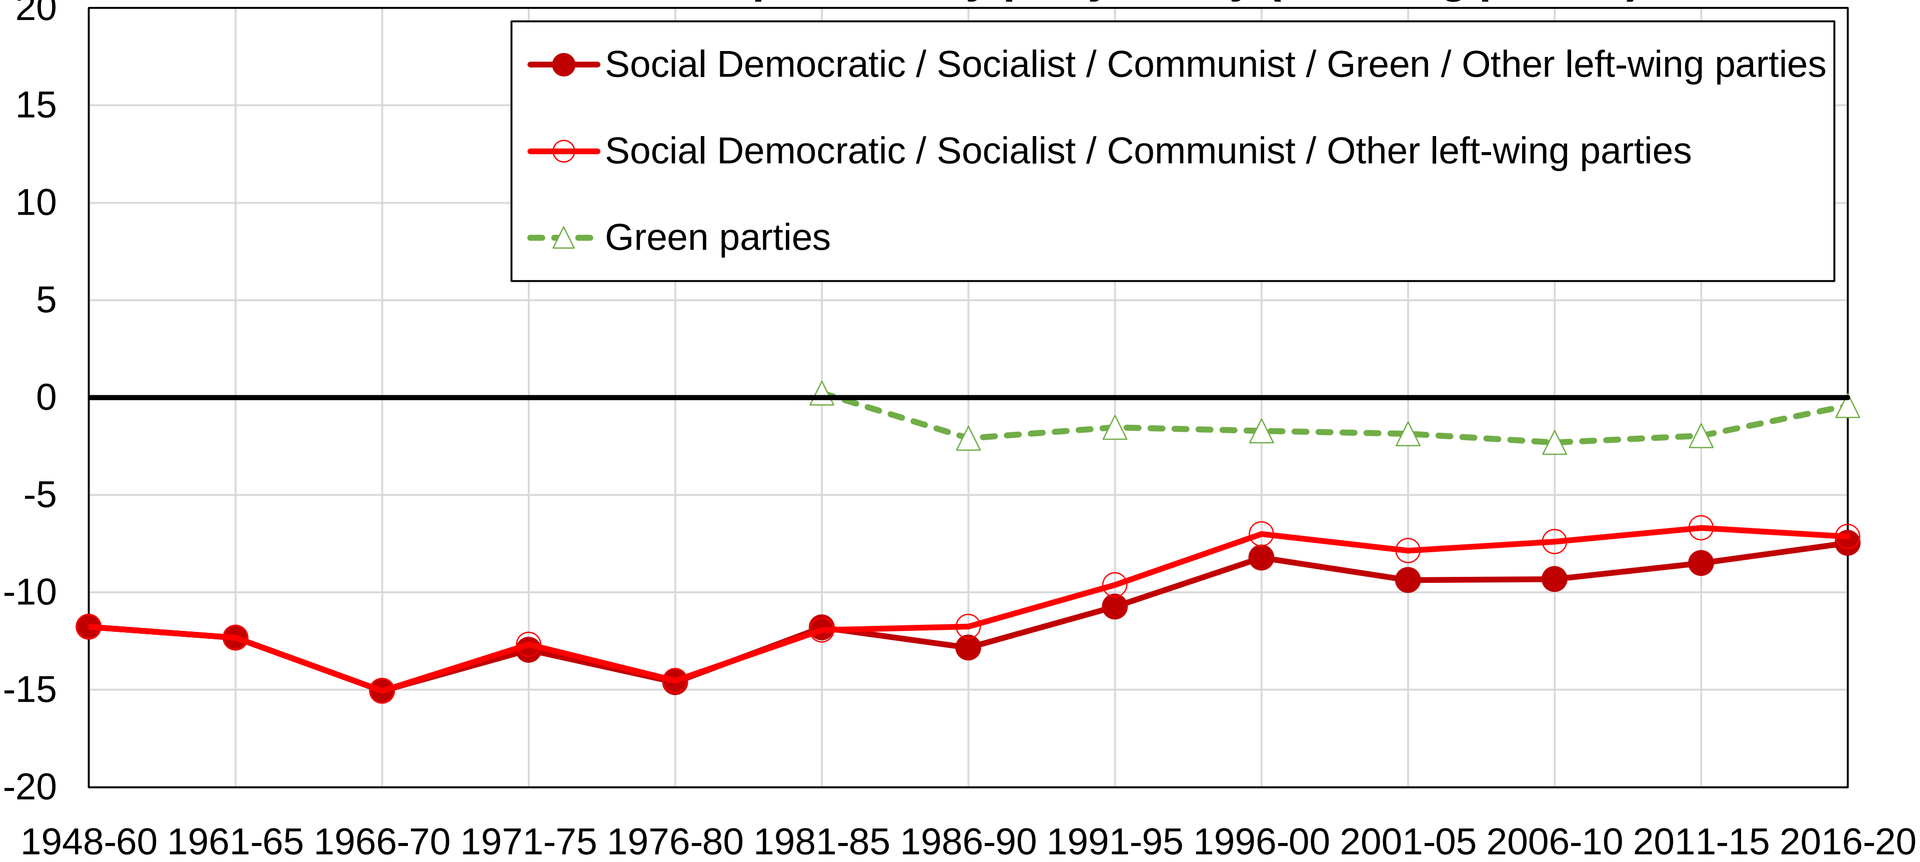

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of top 10% income voters and the share of bottom 90% income voters voting for specific families of parties. Figures correspond to five-year averages over all countries available for a given time period (unbalanced panel of all 21 Western democracies). The estimates are presented after controlling for education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure A28 - The decline/stability of income divides in Western democracies: decomposition by party family (left-wing parties)

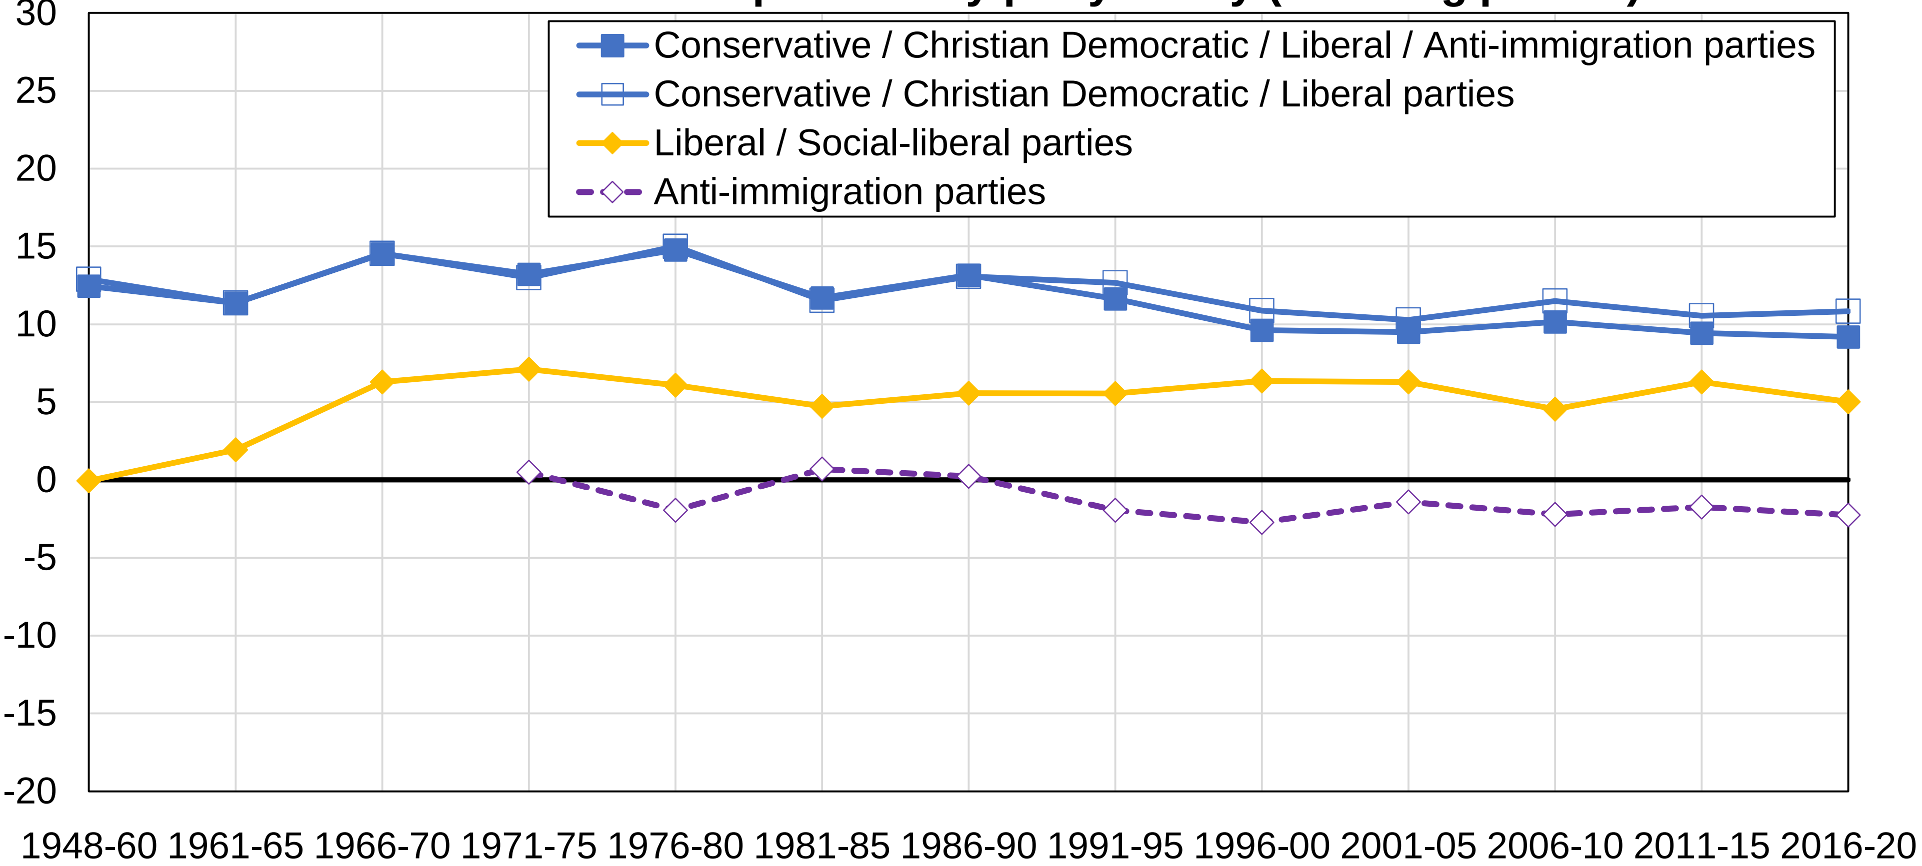

Source: authors' computations using the World Political Cleavages and Inequality Database.

Note: the figure represents the difference between the share of top 10% income voters and the share of bottom 90% income voters voting for specific families of parties. Figures correspond to five-year averages over all countries available for a given time period (unbalanced panel of all 21 Western democracies). The estimates are presented after controlling for education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure A29 - Vote for left-wing parties by education group: decomposition by party family

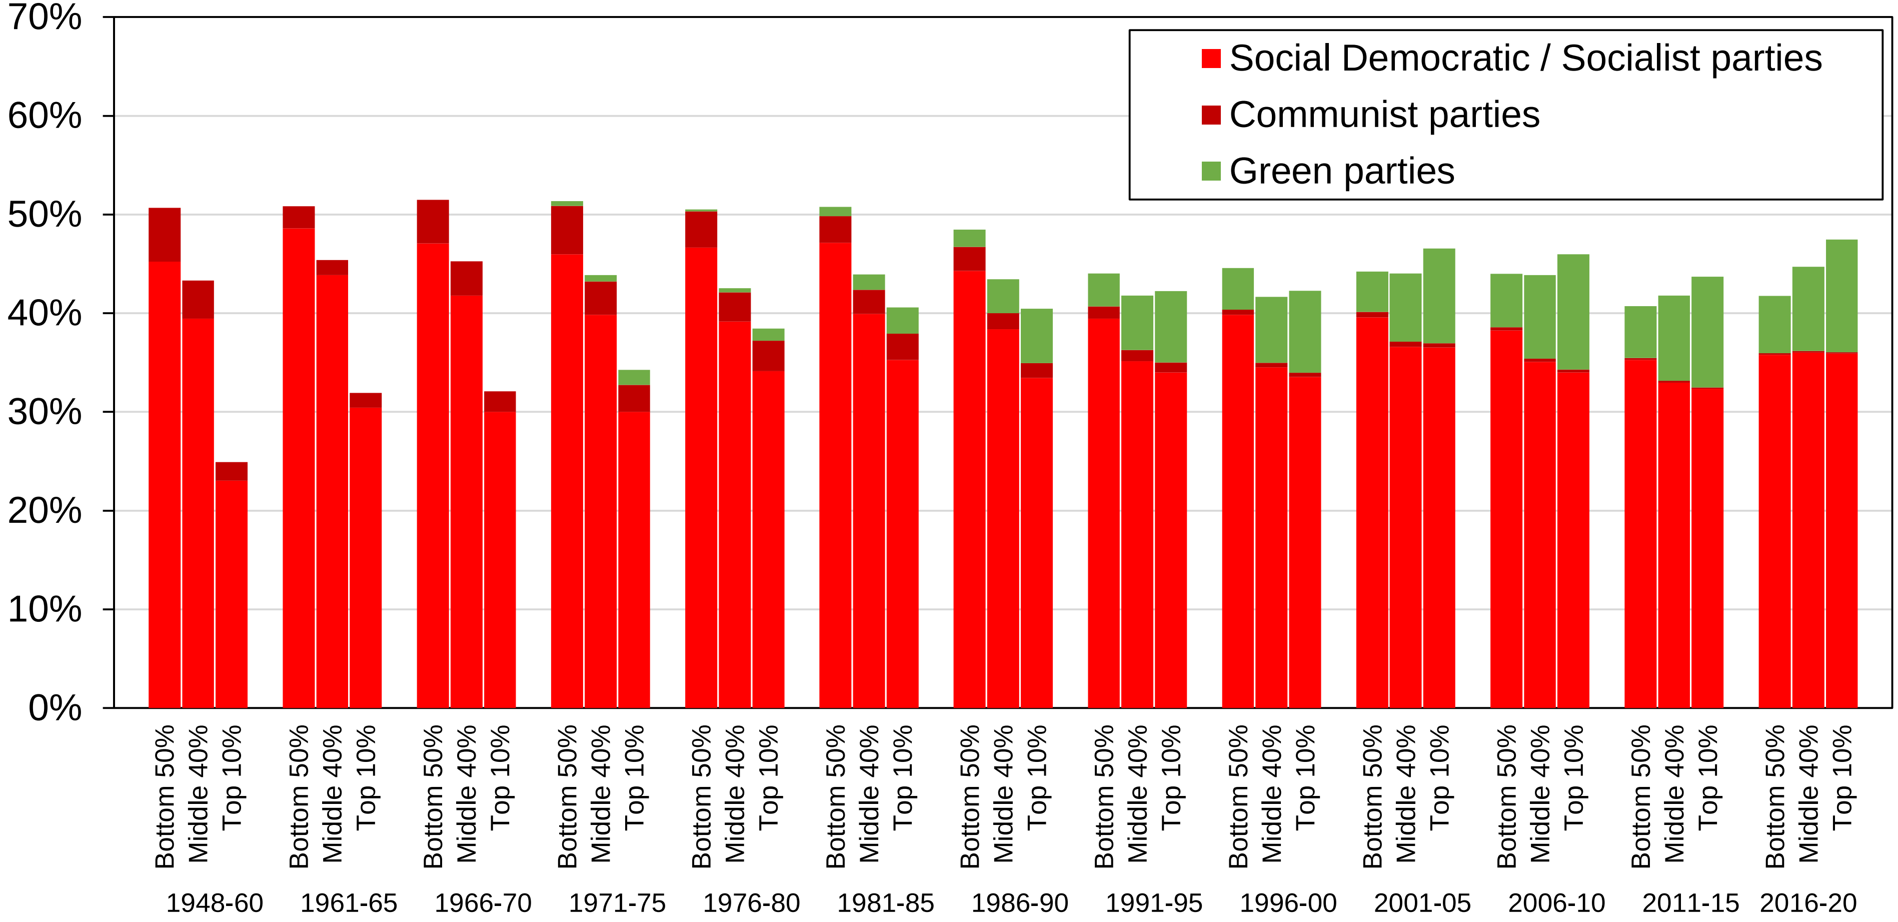

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by each family of parties by education group between 1955 and 2020. Average over all Western democracies.

Figure A30 - Vote for right-wing parties by education group: decomposition by party family

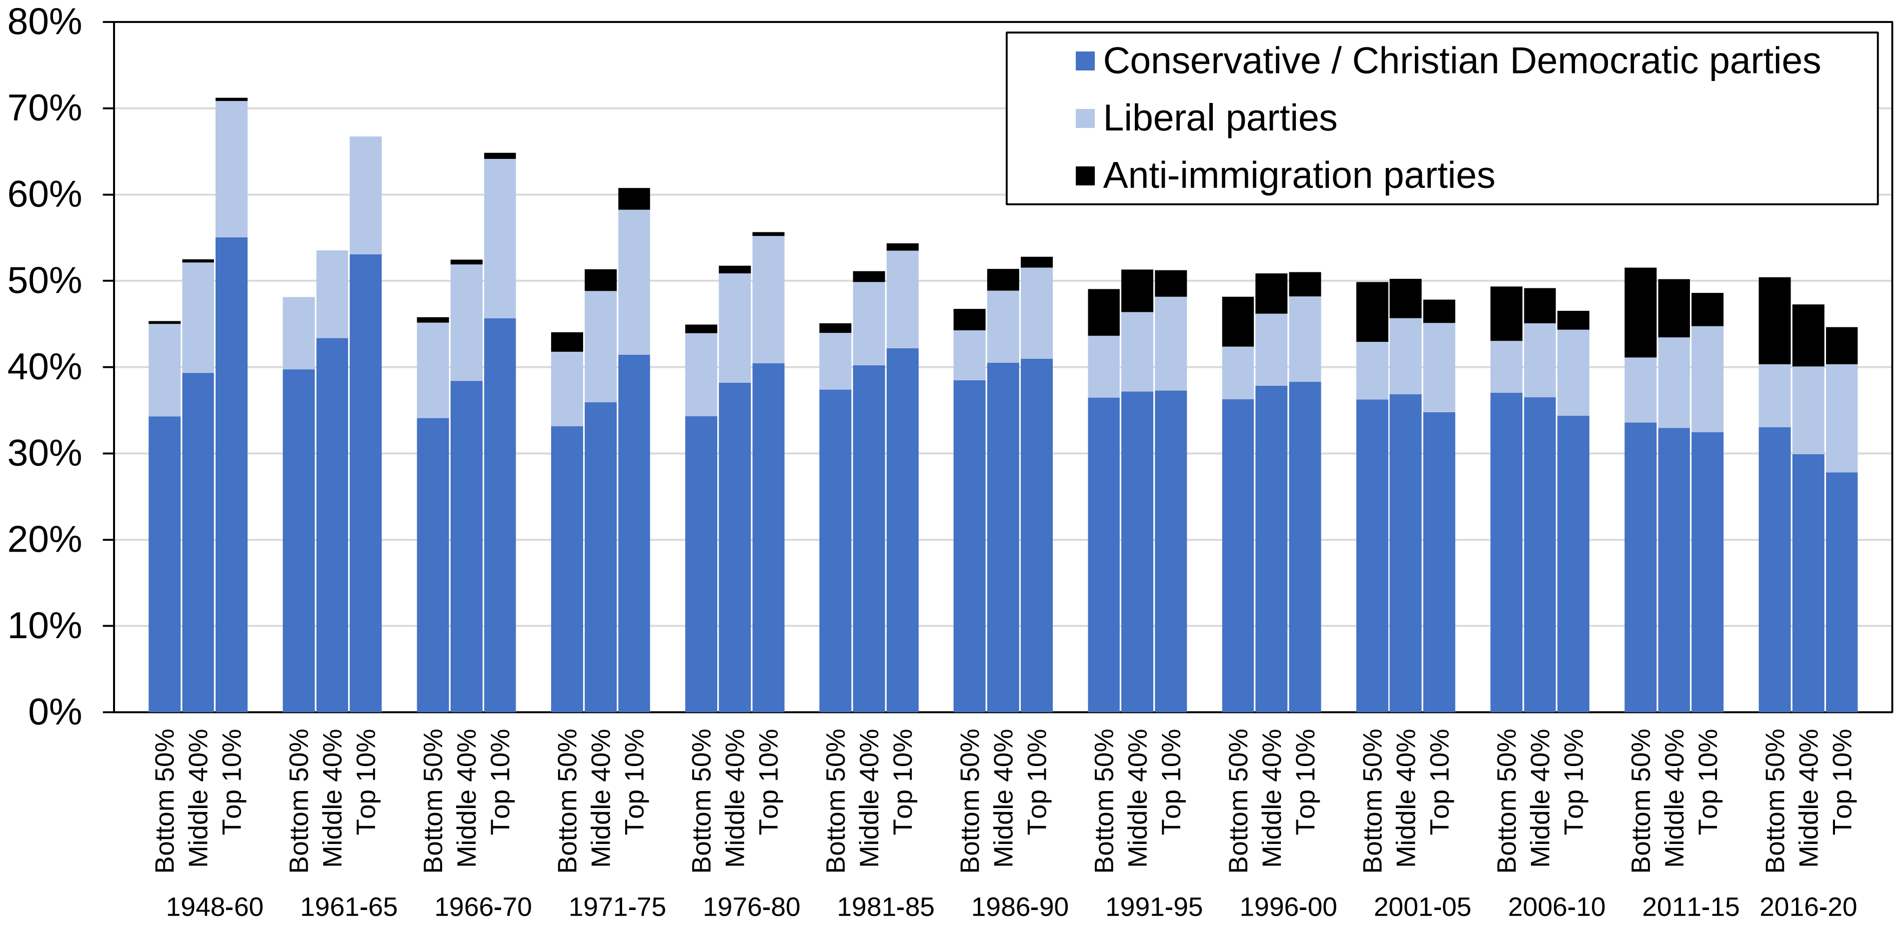

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by each family of parties by education group between 1955 and 2020. Average over all Western democracies.

Figure A31 - Vote for left-wing parties by income group: decomposition by party family

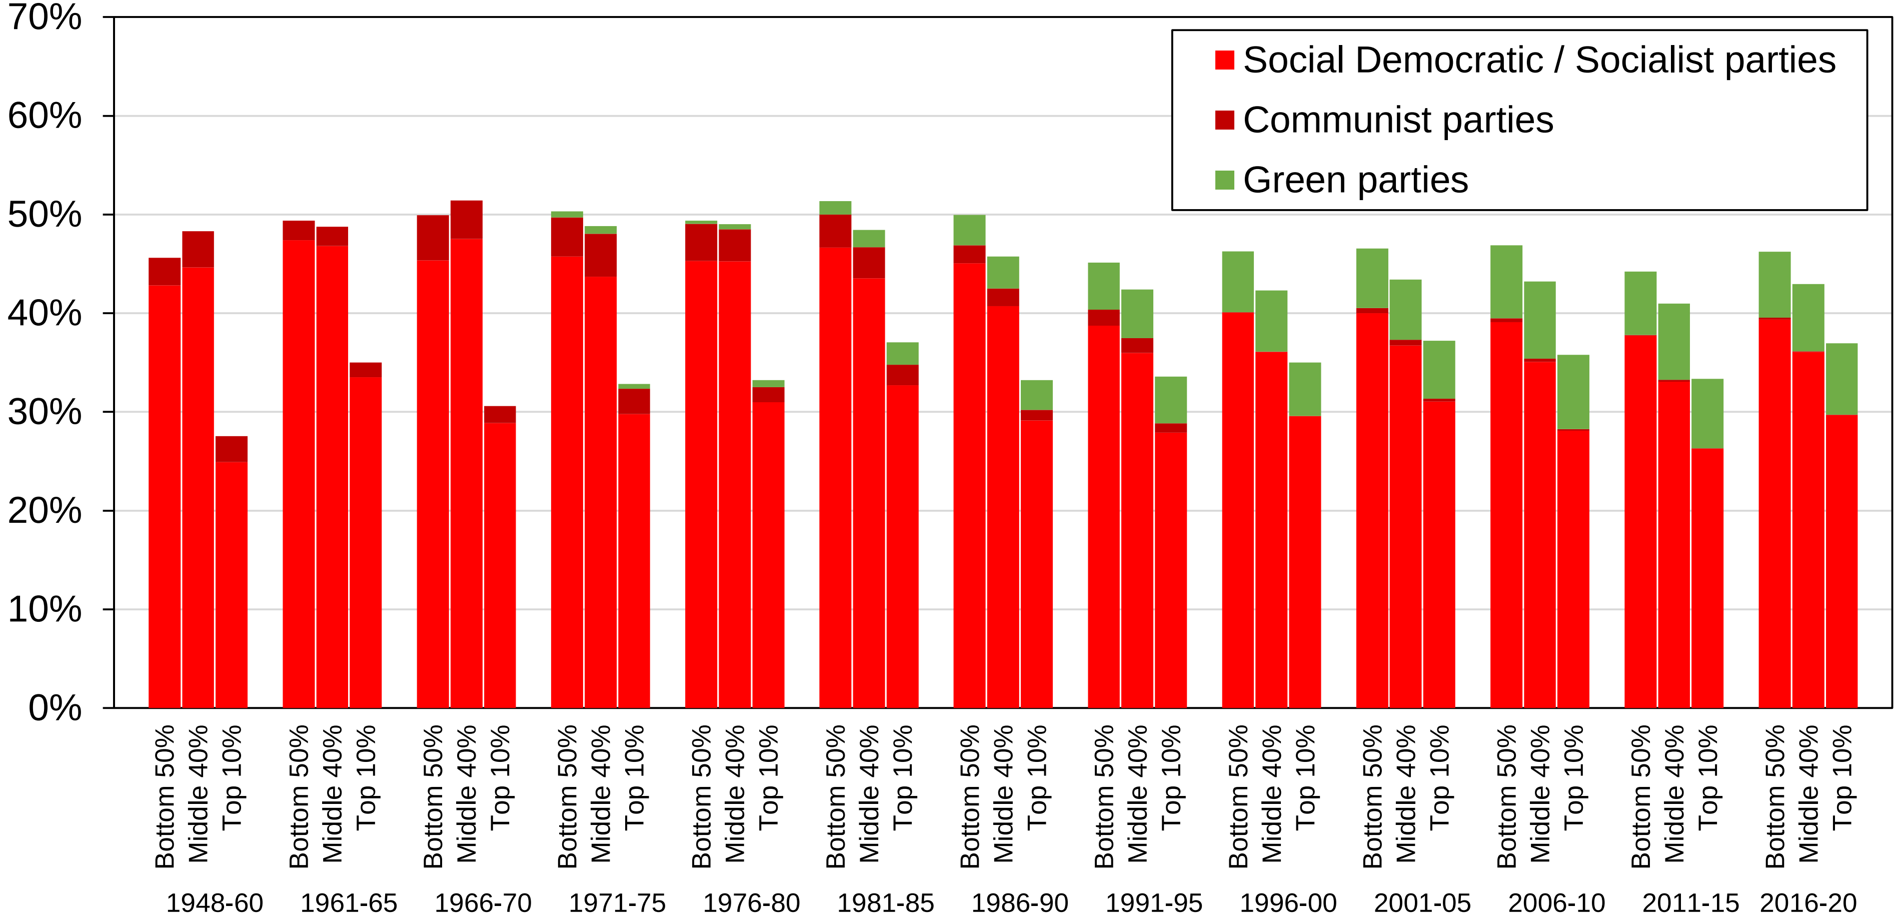

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by each family of parties by income group between 1955 and 2020. Average over all Western democracies.

Figure A32 - Vote for right-wing parties by income group: decomposition by party family

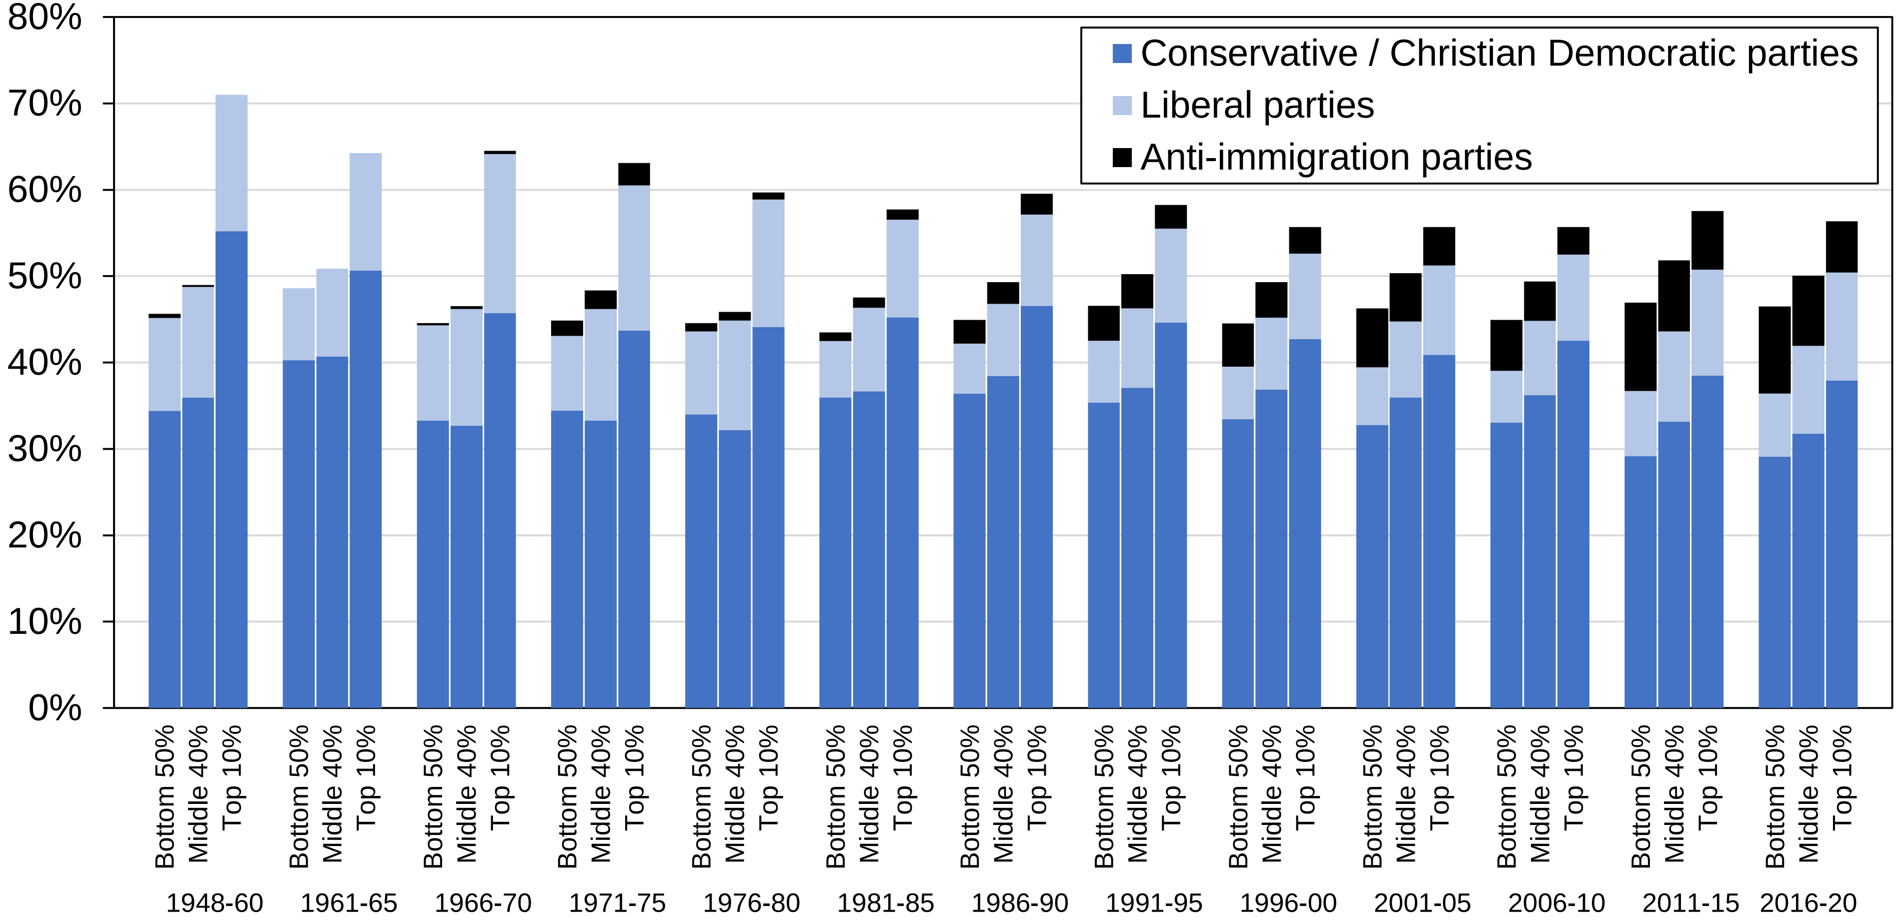

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by each family of parties by income group between 1955 and 2020. Average over all Western democracies.

Figure A33 - Vote for Green parties by education group

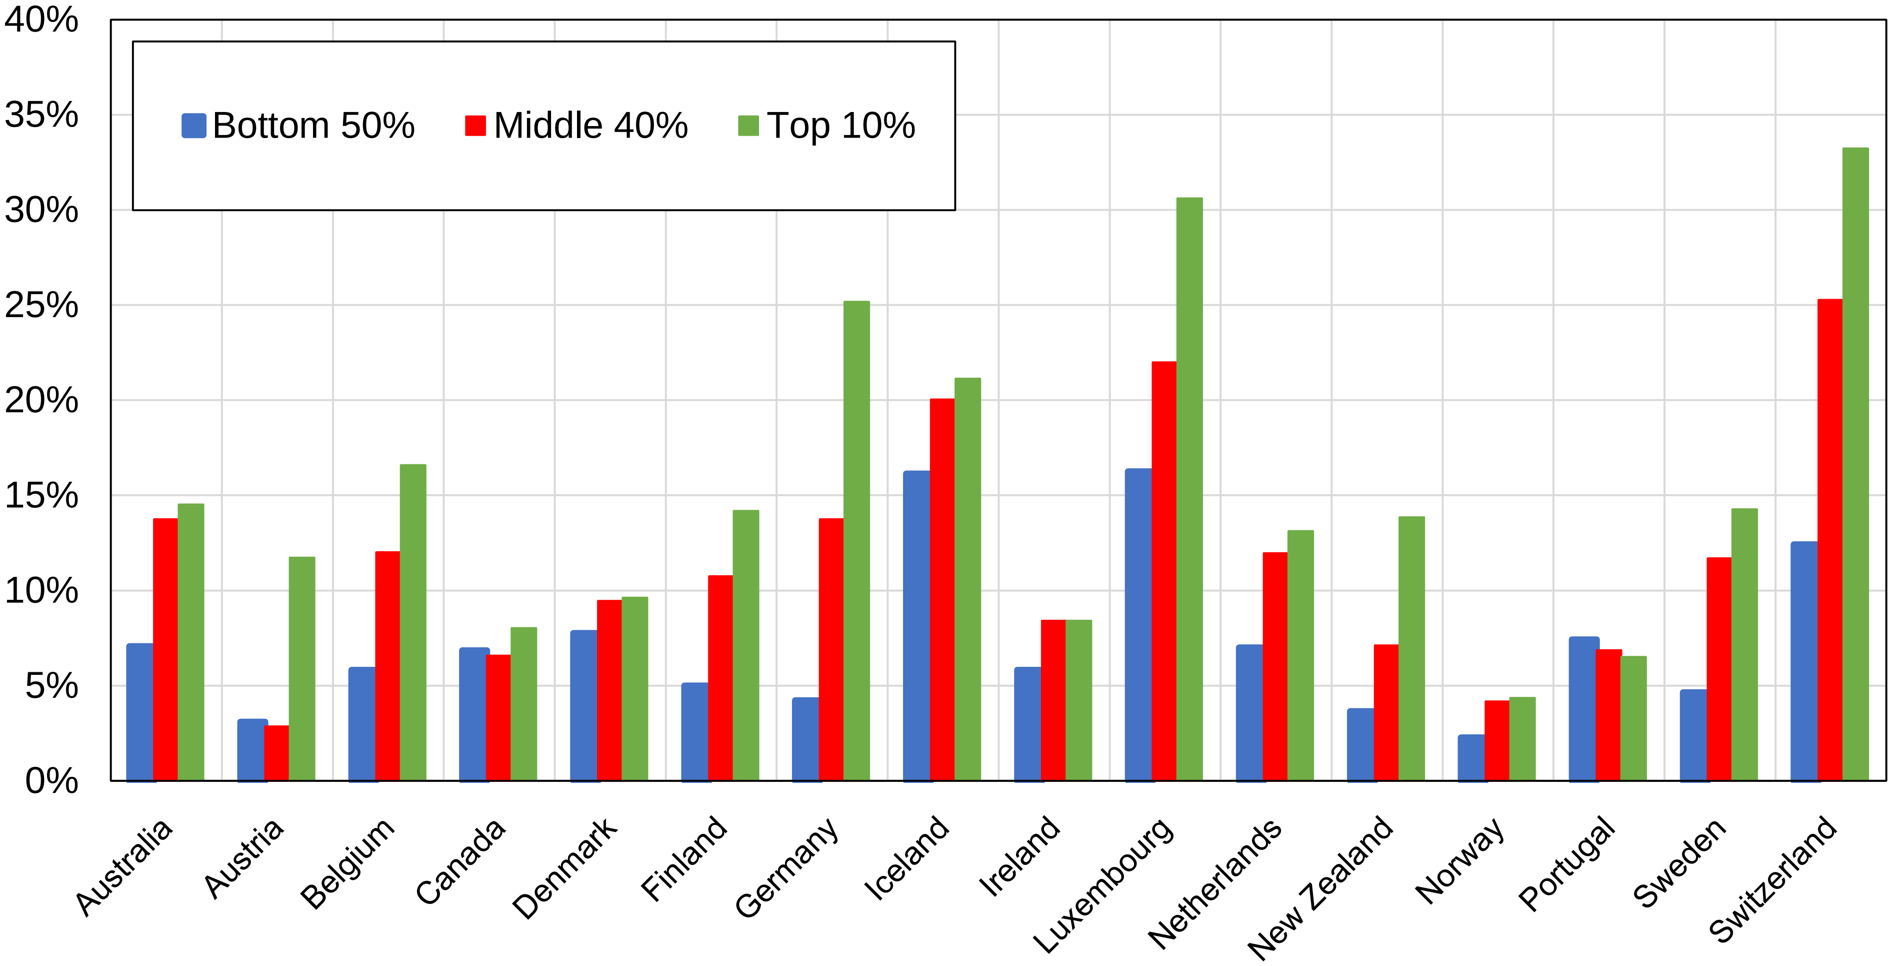

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by Green parties in Western democracies in the last election available by education group.

Figure A34 - Vote for Green parties by income group

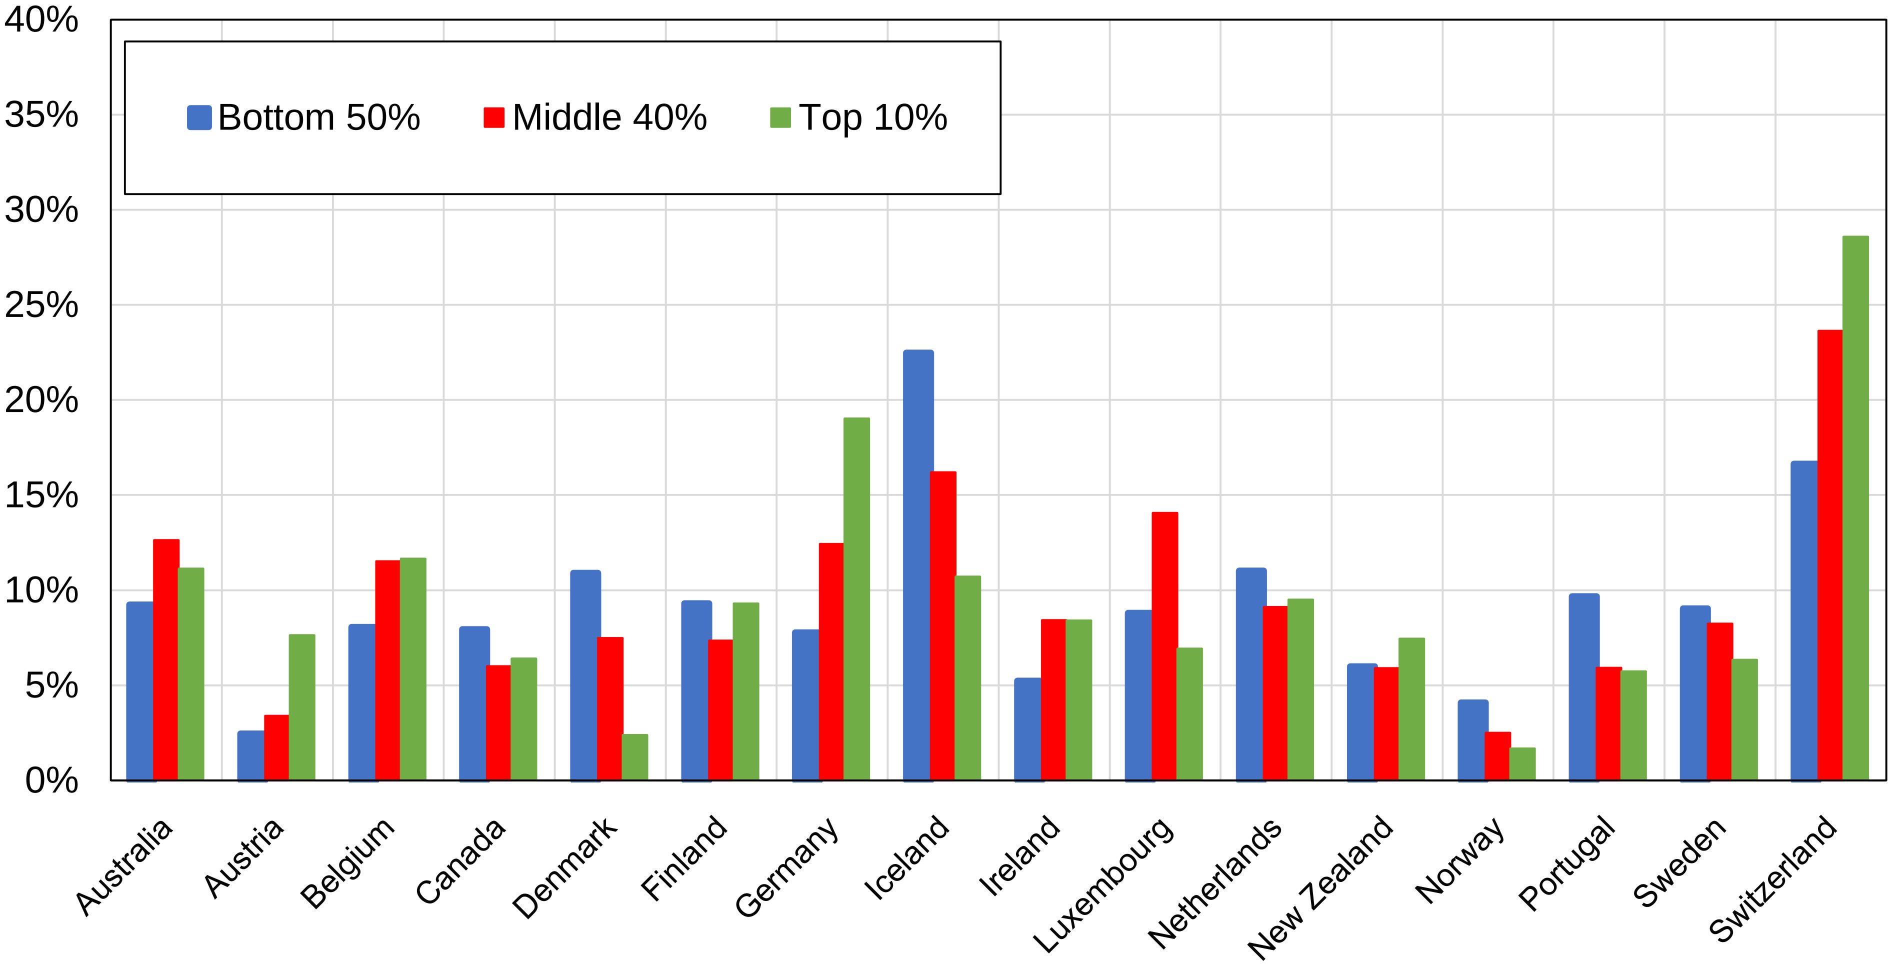

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by Green parties in Western democracies in the last election available by income group.

Figure A35 - Vote for anti-immigration parties by education group

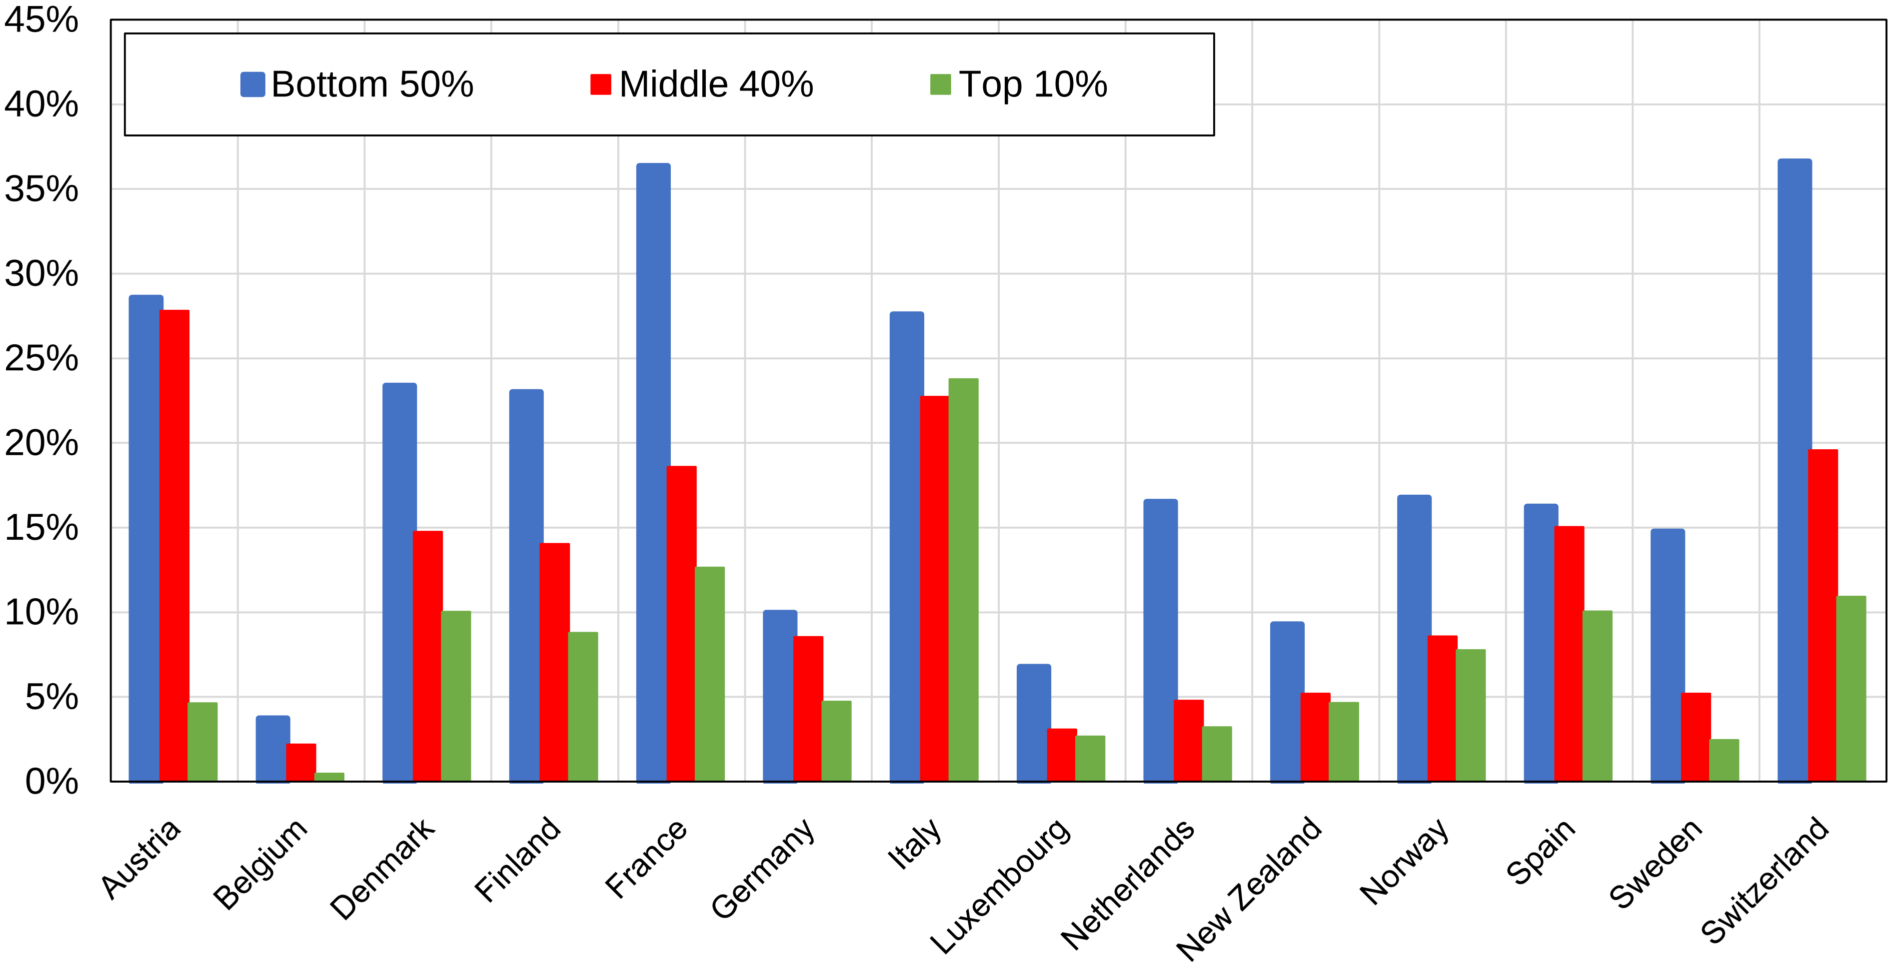

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by anti-immigration parties in Western democracies in the last election available by education group.

Figure A36 - Vote for anti-immigration parties by income group

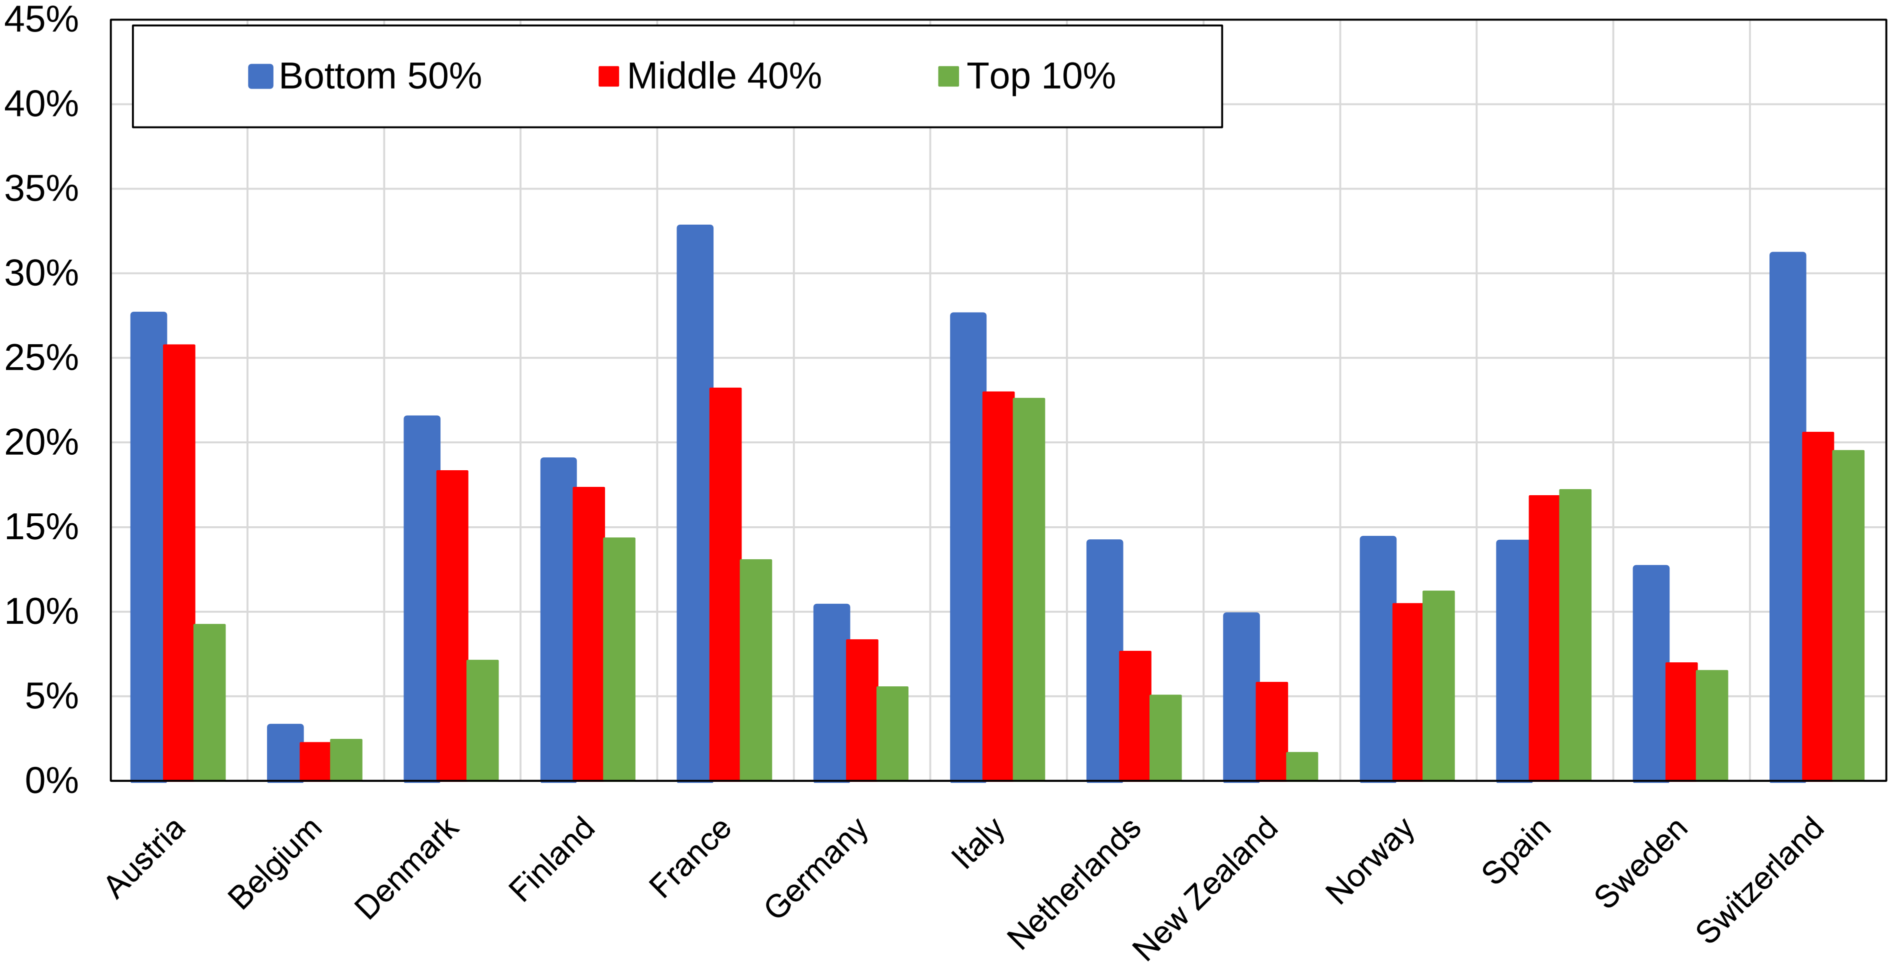

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by anti-immigration parties in Western democracies in the last election available by income group.

Figure A37 - Composition of parties voted for by top 10% educated voters

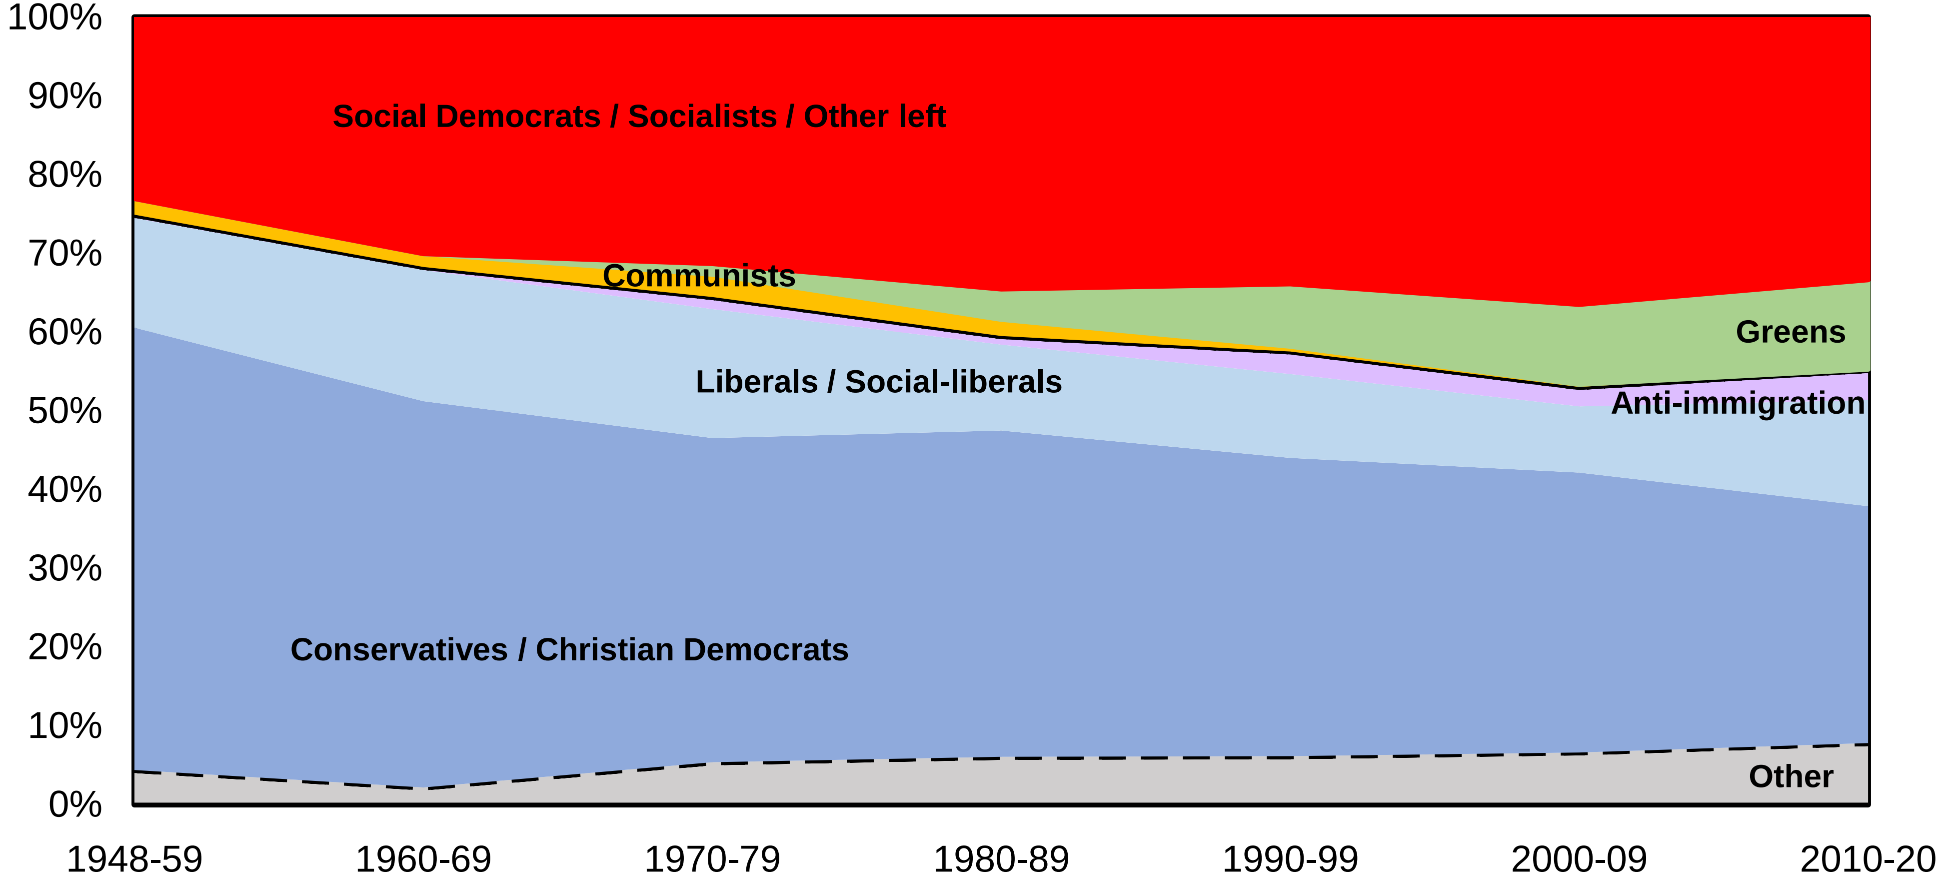

Source: authors' computations using electoral surveys.
Note: the figure represents the average share of votes received by selected families of political parties in Western democracies between the 1940s and the 2010s within the top 10% group of highest educated voters. Decennial averages over all Western democracies. The dashed lines delimit the categorization of parties considered in the main specification (social democrats and affiliated, conservatives and affiliated, and other parties).

Figure A38 - Composition of parties voted for by top 10% income voters

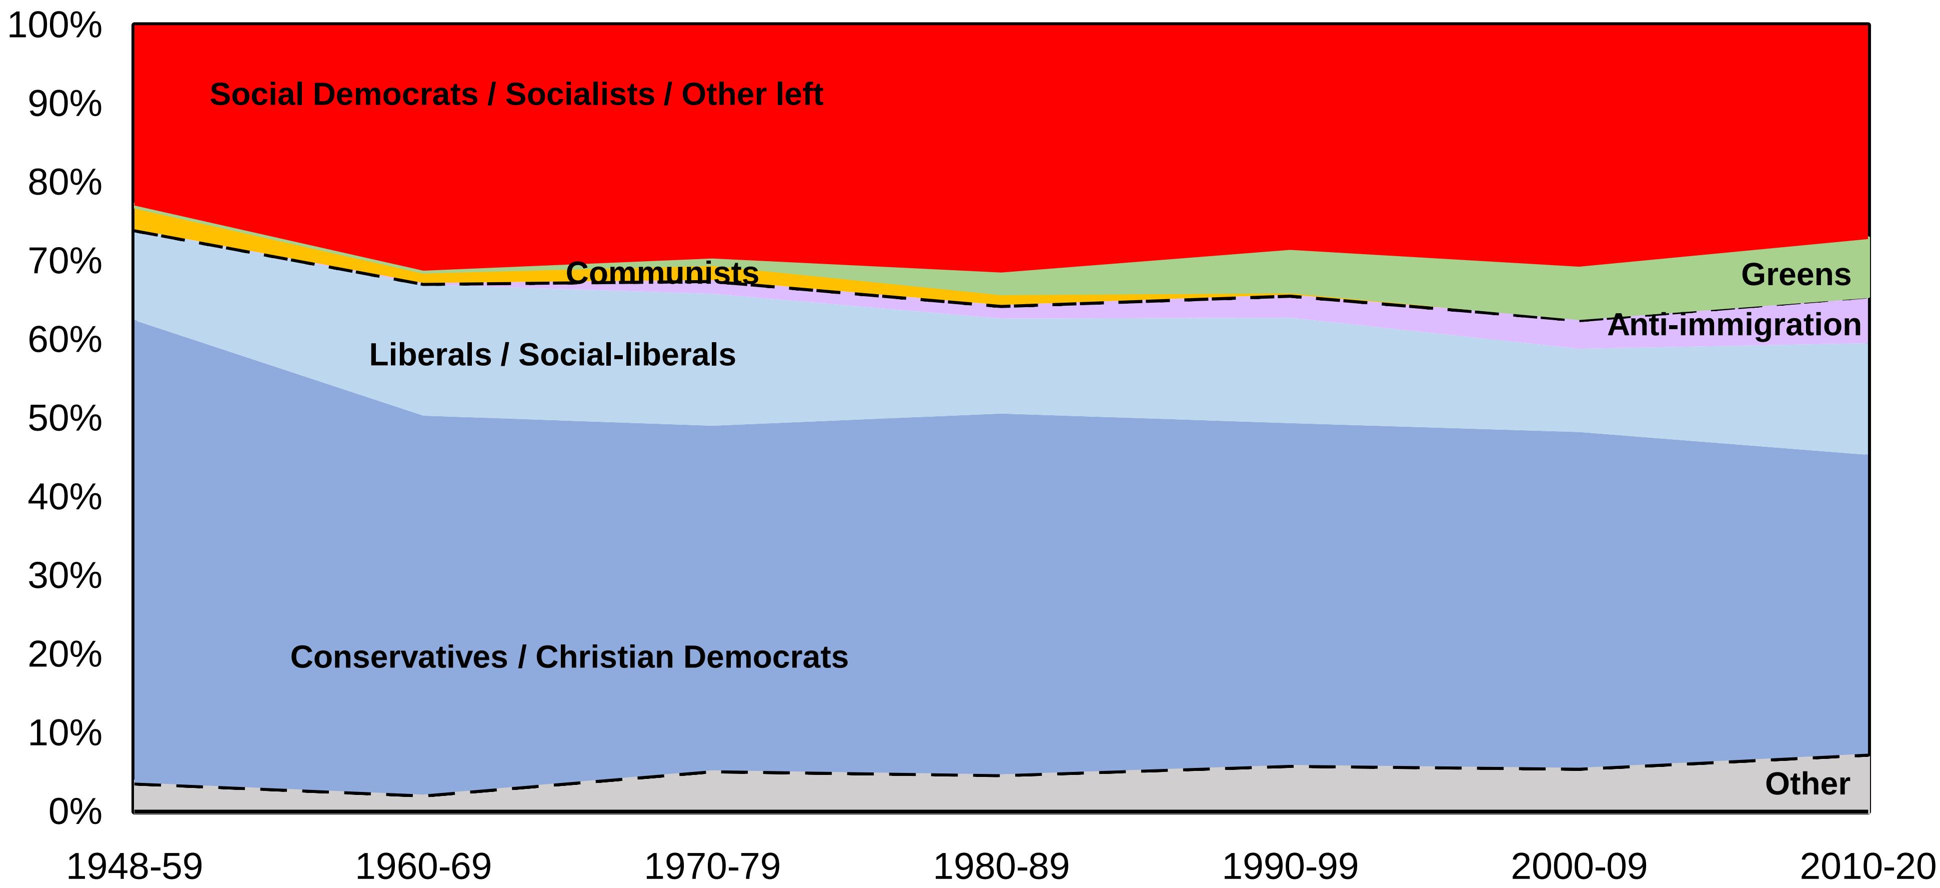

Source: authors' computations using electoral surveys.
Note: the figure represents the average share of votes received by selected families of political parties in Western democracies between the 1940s and the 2010s within the top 10% group of highest income voters. Decennial averages over all Western democracies. The dashed lines delimit the categorization of parties considered in the main specification (social democrats and affiliated, conservatives and affiliated, and other parties).

Figure A39 - The fragmentation of political cleavage structures.
Panel A. 1960s-1980s

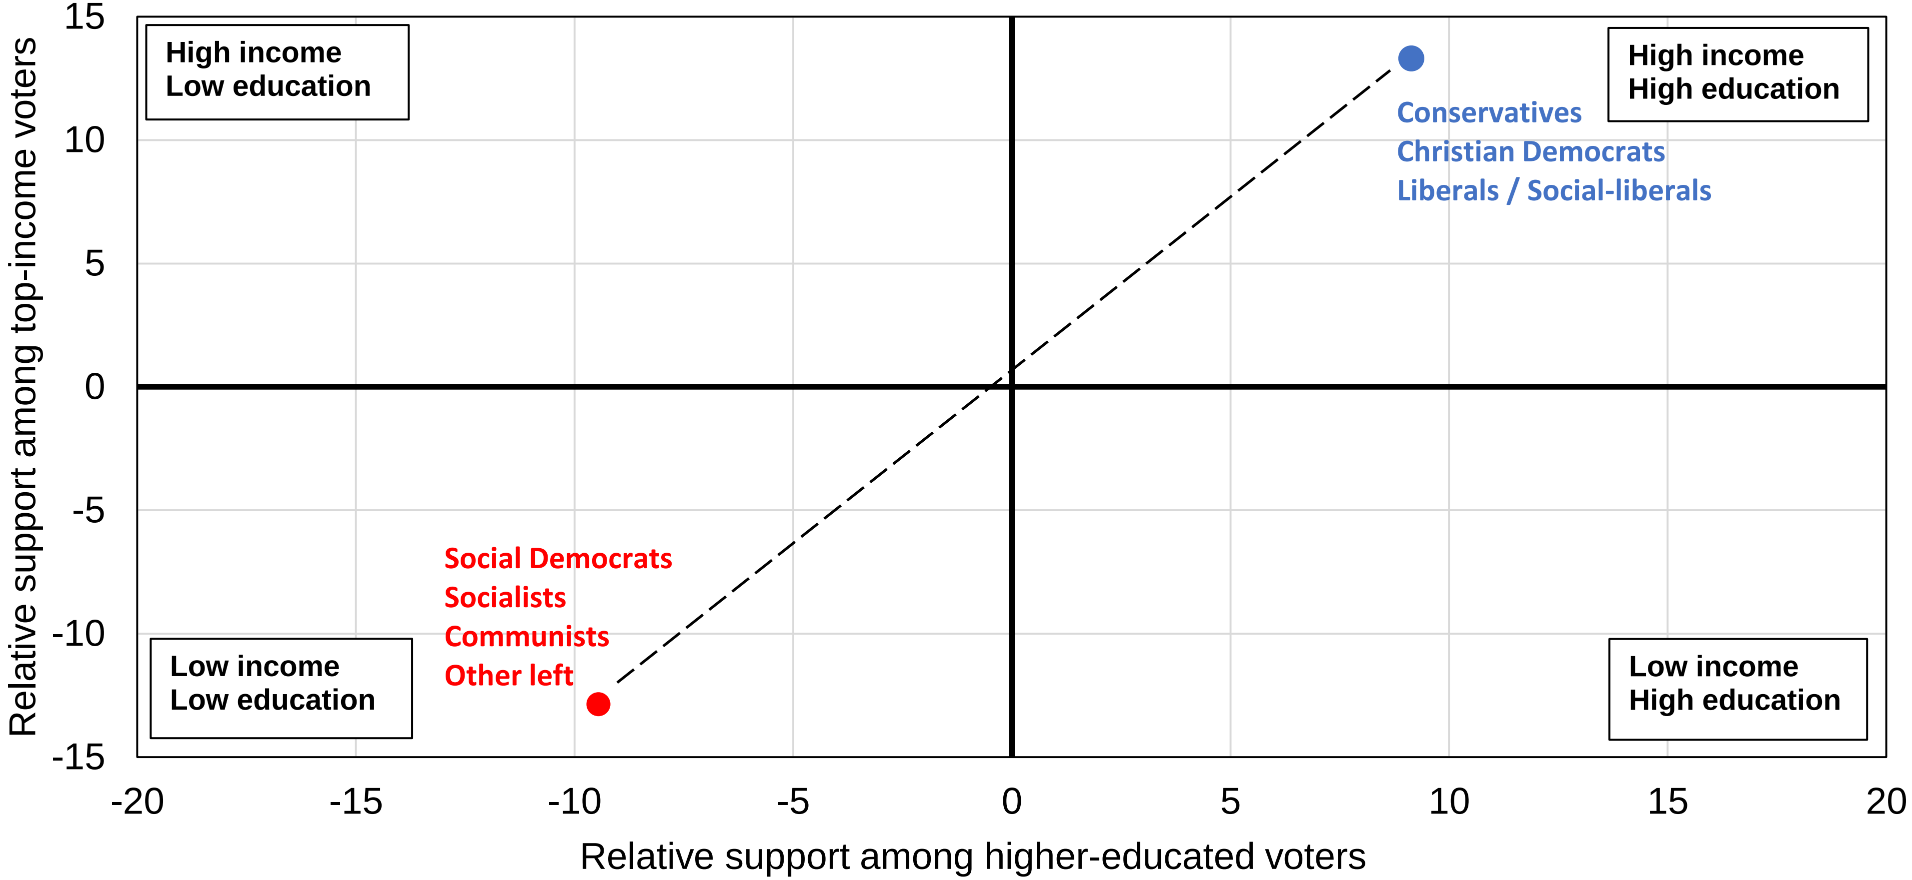

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of high-income (top 10%) and low-income (bottom 90%) voters voting for selected groups of parties on the y-axis, and the same difference between higher-educated (top 10%) and lower-educated (bottom 90%) voters on the x-axis. In the 1960s, social democratic, socialist, and communist parties were supported by both low-income and lower-educated voters, while conservative, Christian, and liberal parties were supported by both high-income and higher-educated voters. Averages over all Western democracies. Estimates control for income/education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure A40 - The fragmentation of political cleavage structures.
Panel B. 2000-2020

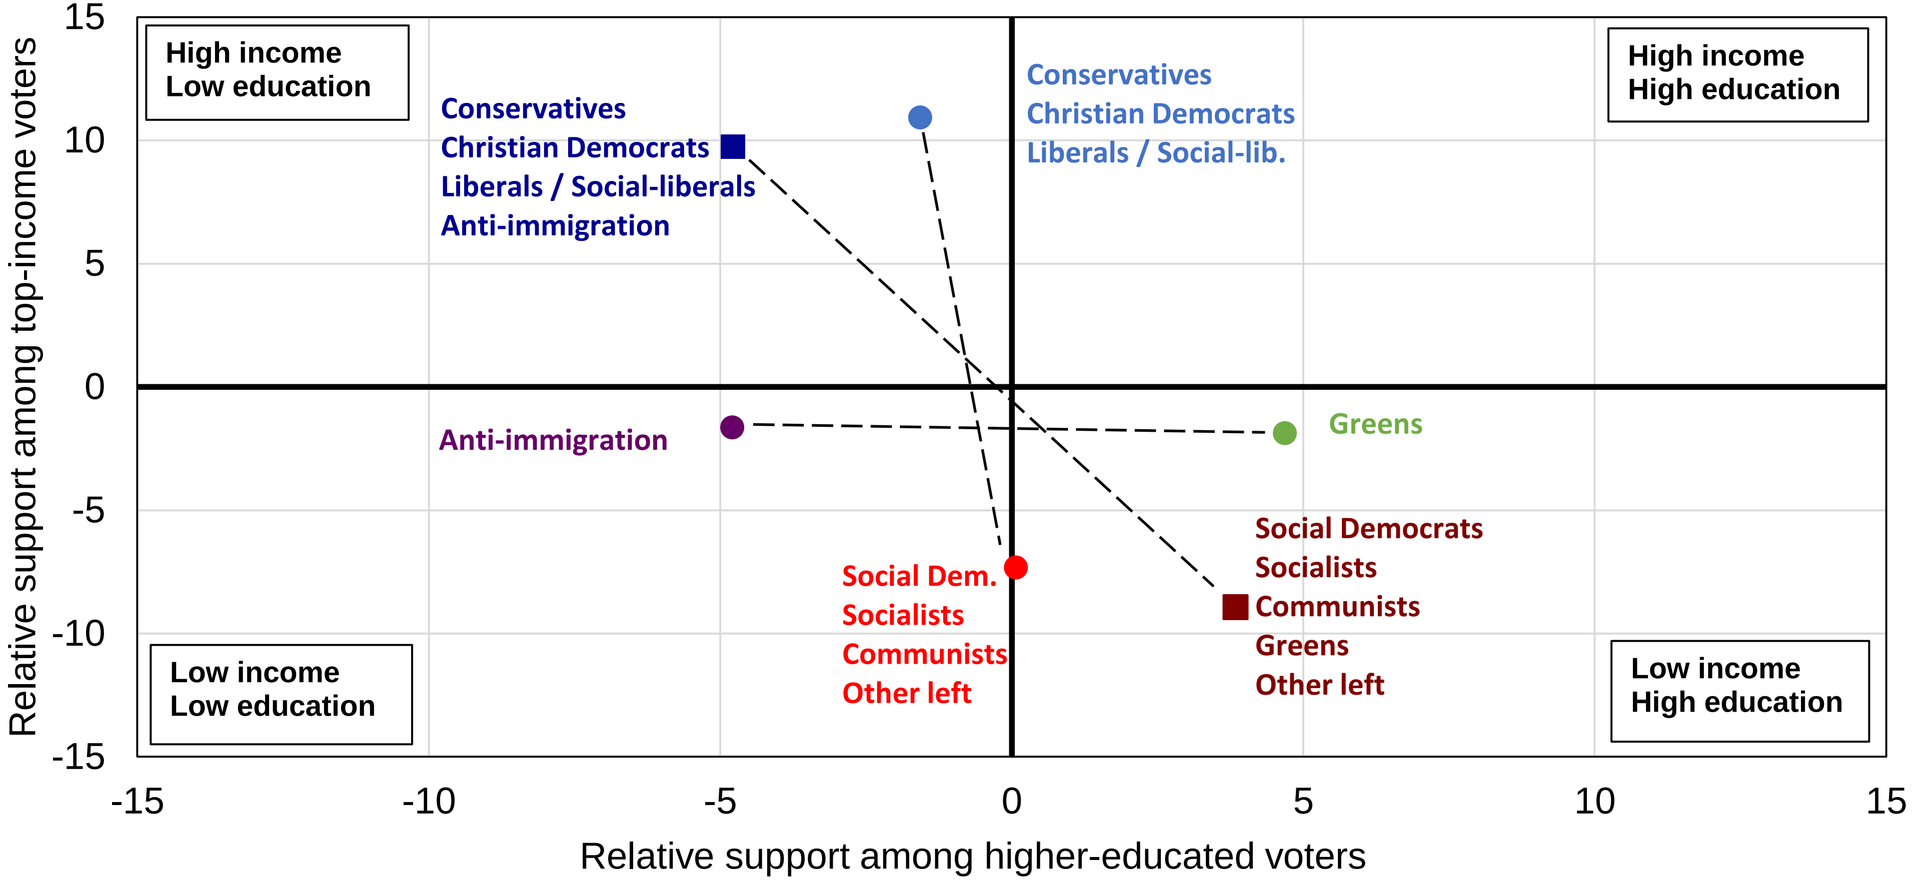

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of high-income (top 10%) and low-income (bottom 90%) voters voting for selected groups of parties on the y-axis, and the same difference between higher-educated (top 10%) and lower-educated (bottom 90%) voters on the x-axis. In 2000-2020, education most clearly distinguishes anti-immigration from green parties, while income most clearly distinguishes conservative and Christian parties from social democratic, socialist, and communist parties. Averages over all Western democracies. Estimates control for income/education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure A41 - Educational and income divides: Detailed party families

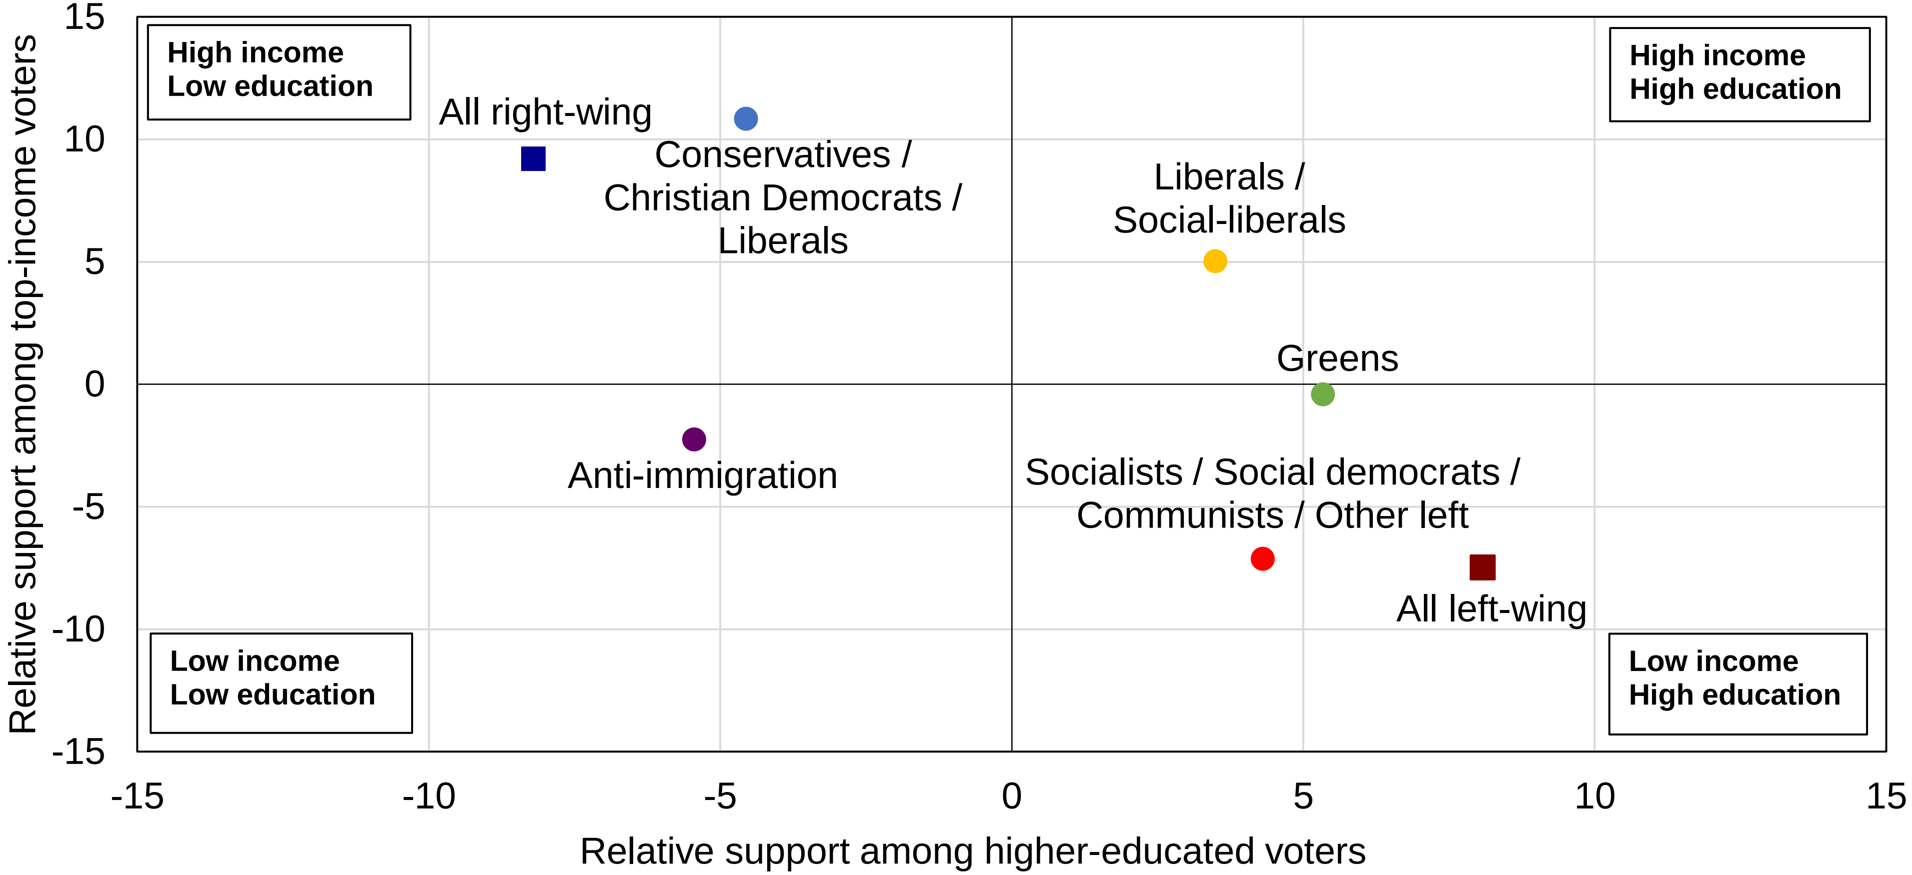

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of high-income (top 10%) and low-income (bottom 90%) voters voting for selected groups of parties on the y-axis, and the same difference between higher-educated (top 10%) and lower-educated (bottom 90%) voters on the x-axis. Education most clearly distinguishes anti-immigration from green parties, while income distinguishes most clearly conservative and Christian parties from socialist, social democratic and communist parties. Averages over all Western democracies over the 2000-2020 period. Estimates control for income/education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure A42 - The disconnection of income and education cleavages in Western democracies (quadrant representation), all countries

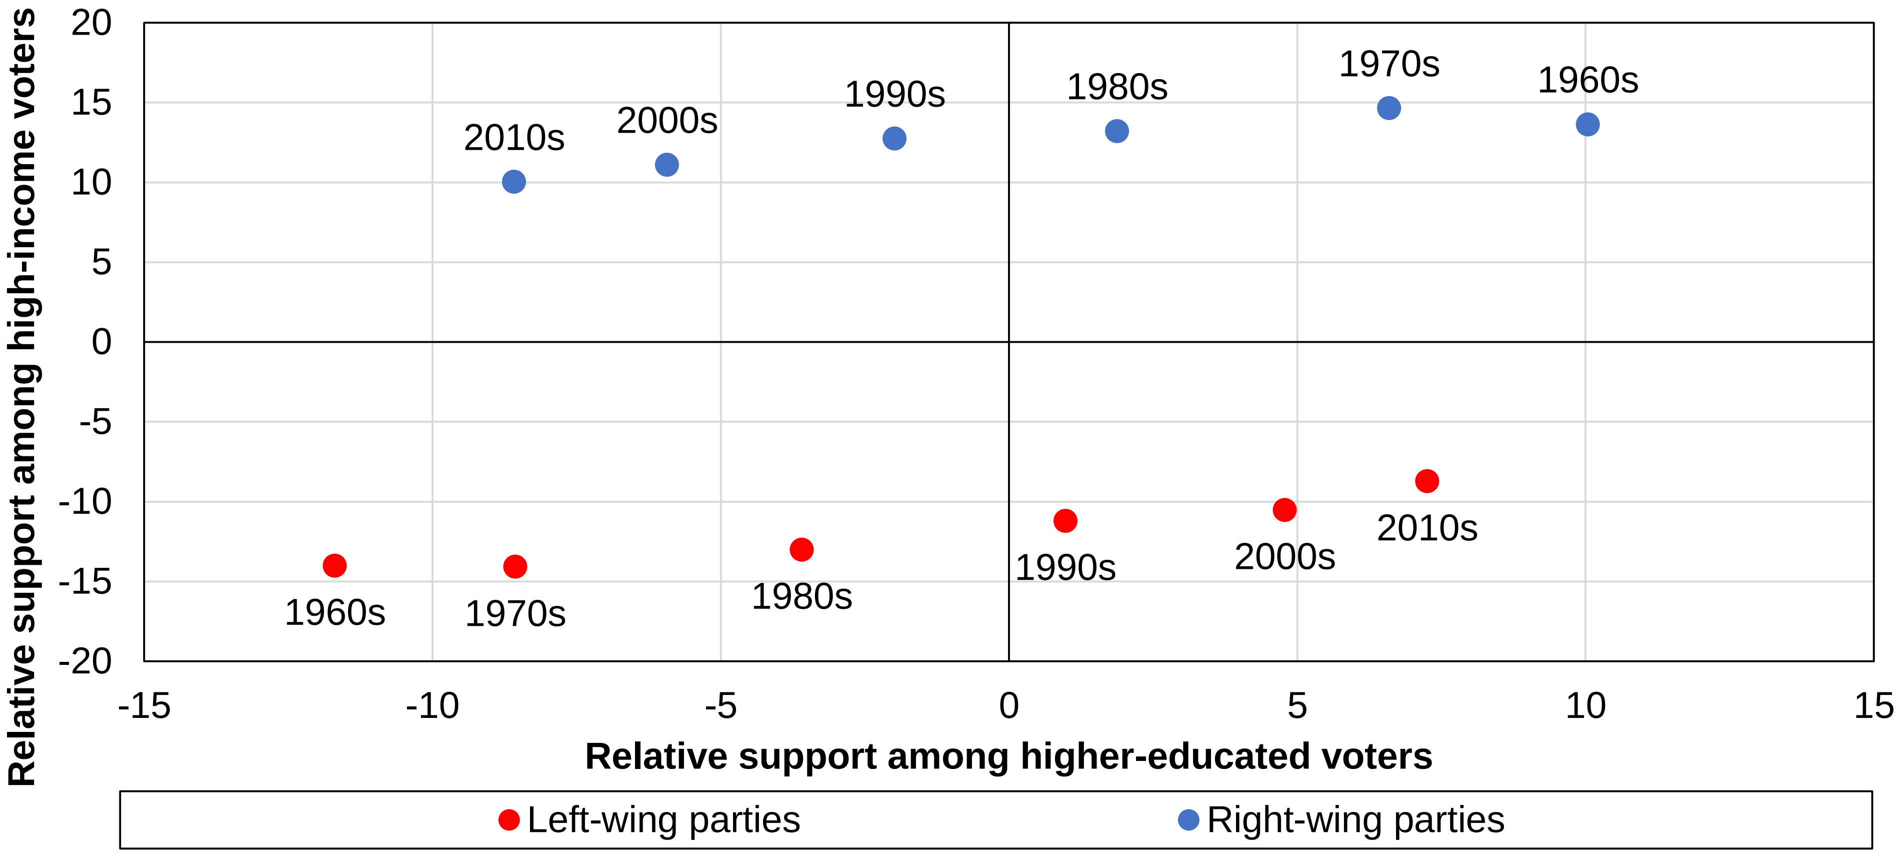

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of high-income (top 10%) and low-income (bottom 90%) voters voting for selected groups of parties on the y-axis, and the same difference between higher-educated (top 10%) and lower-educated (bottom 90%) voters on the x-axis. Estimates control for income/education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available). Figures correspond to ten-year averages for Australia, Britain, Canada, Denmark, France, Germany, Italy, the Netherlands, Norway, Sweden, Switzerland, and the US.

Figure A43 - Income and educational divides in Western democracies, 1950s

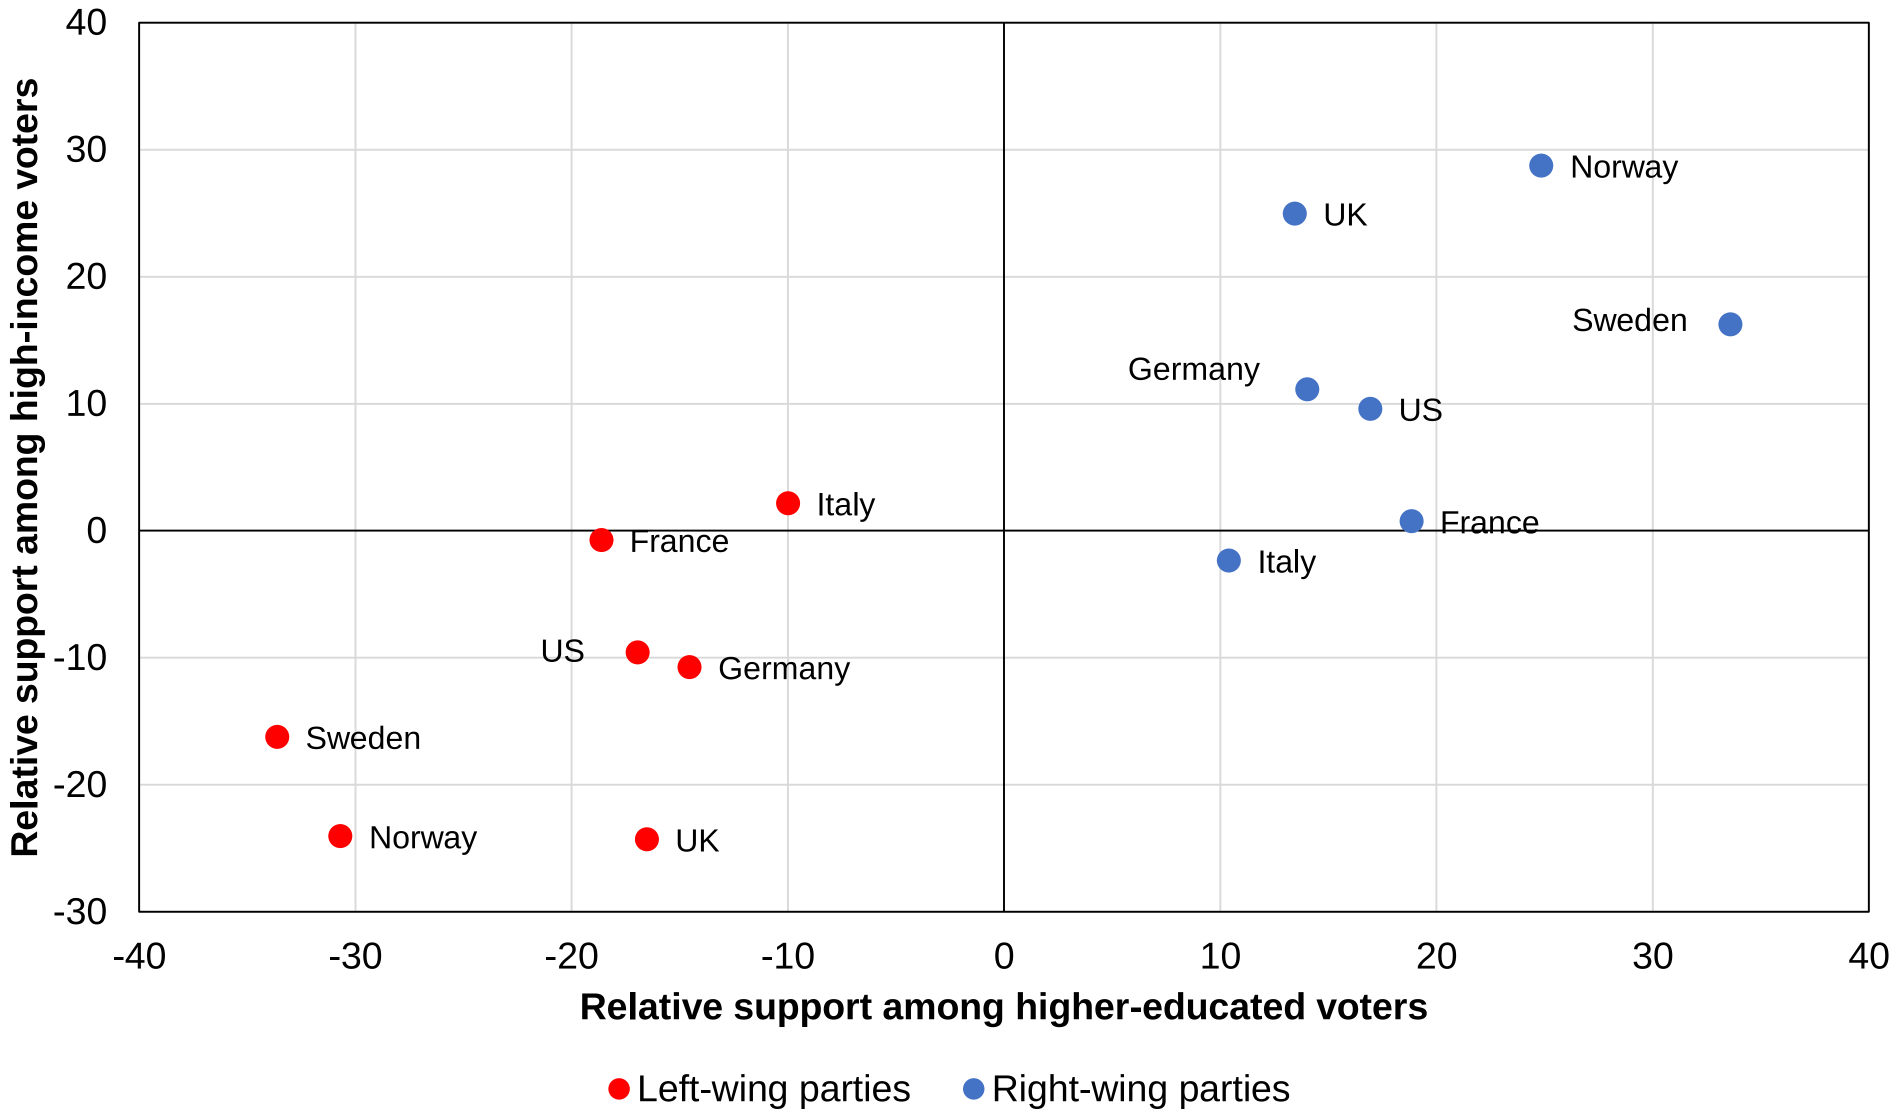

Figure A44 - Income and educational divides in Western democracies, 1960s

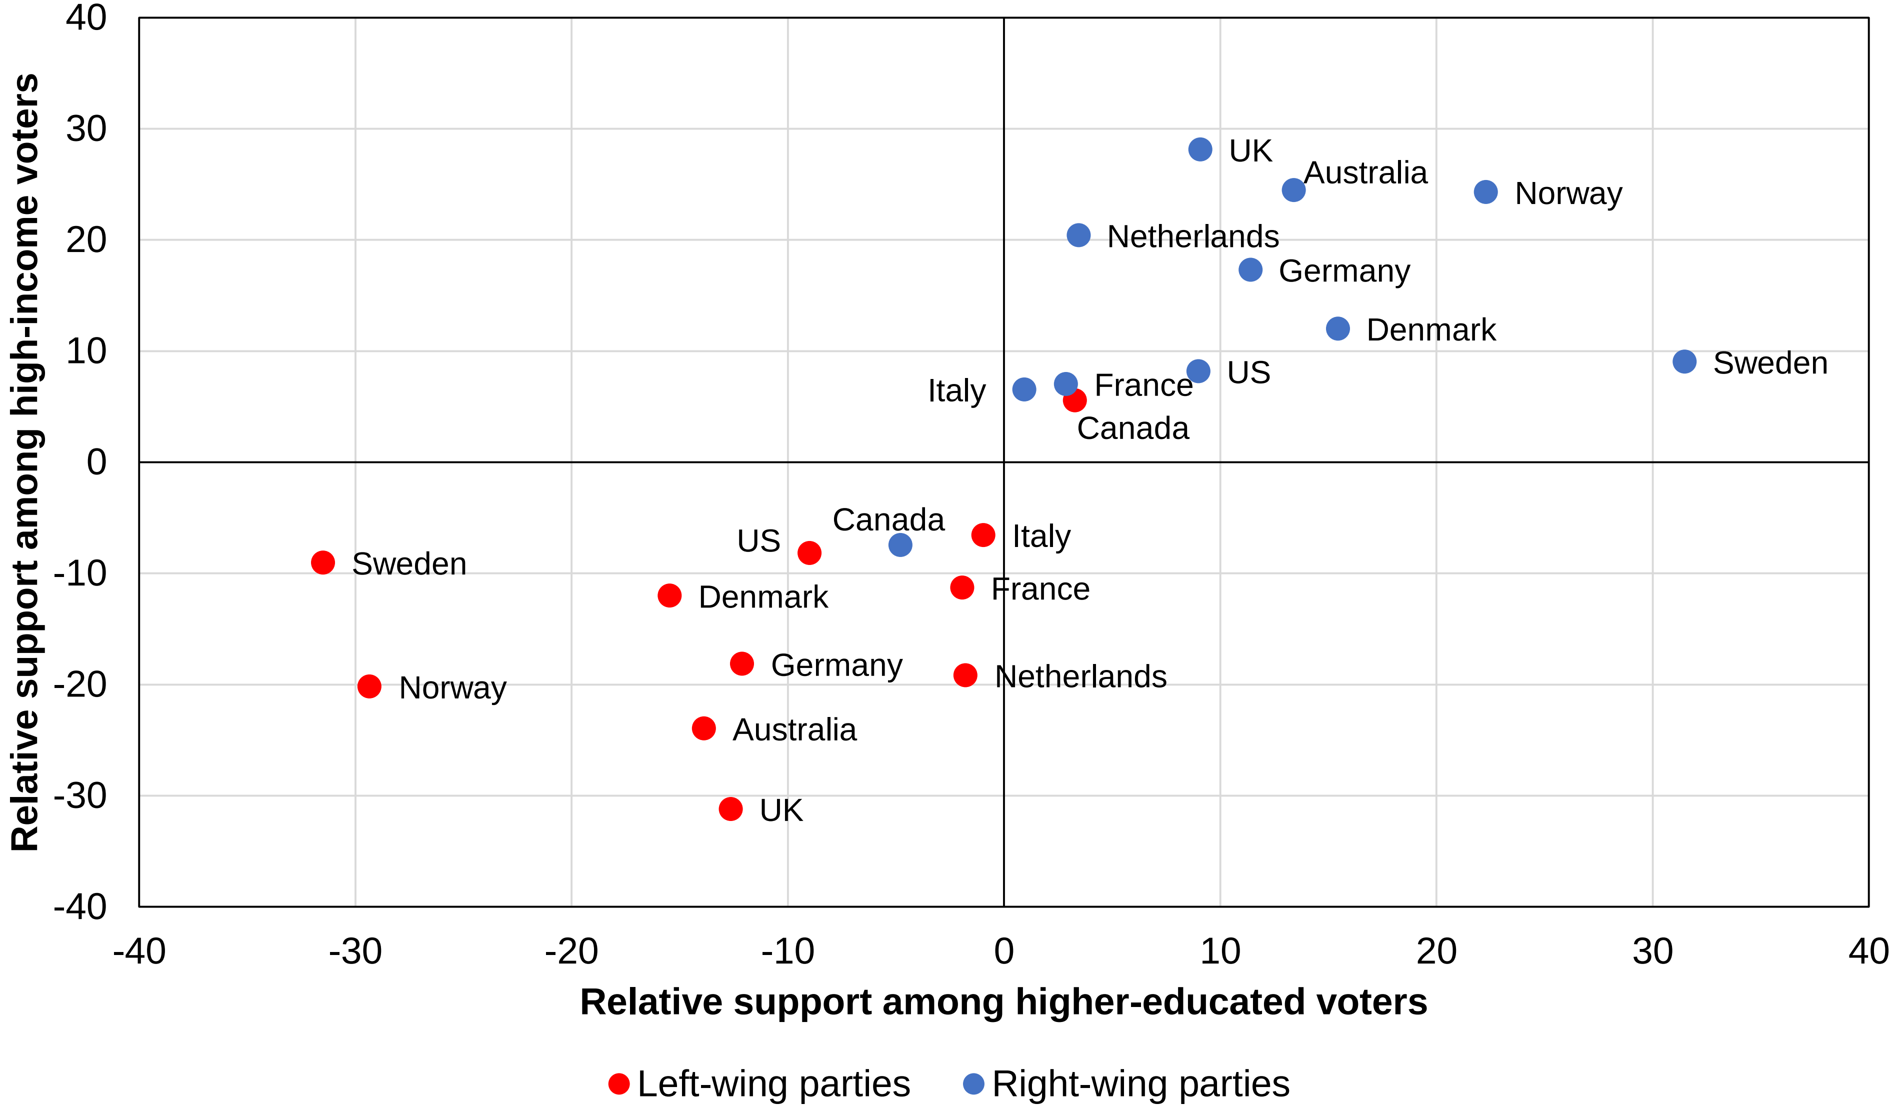

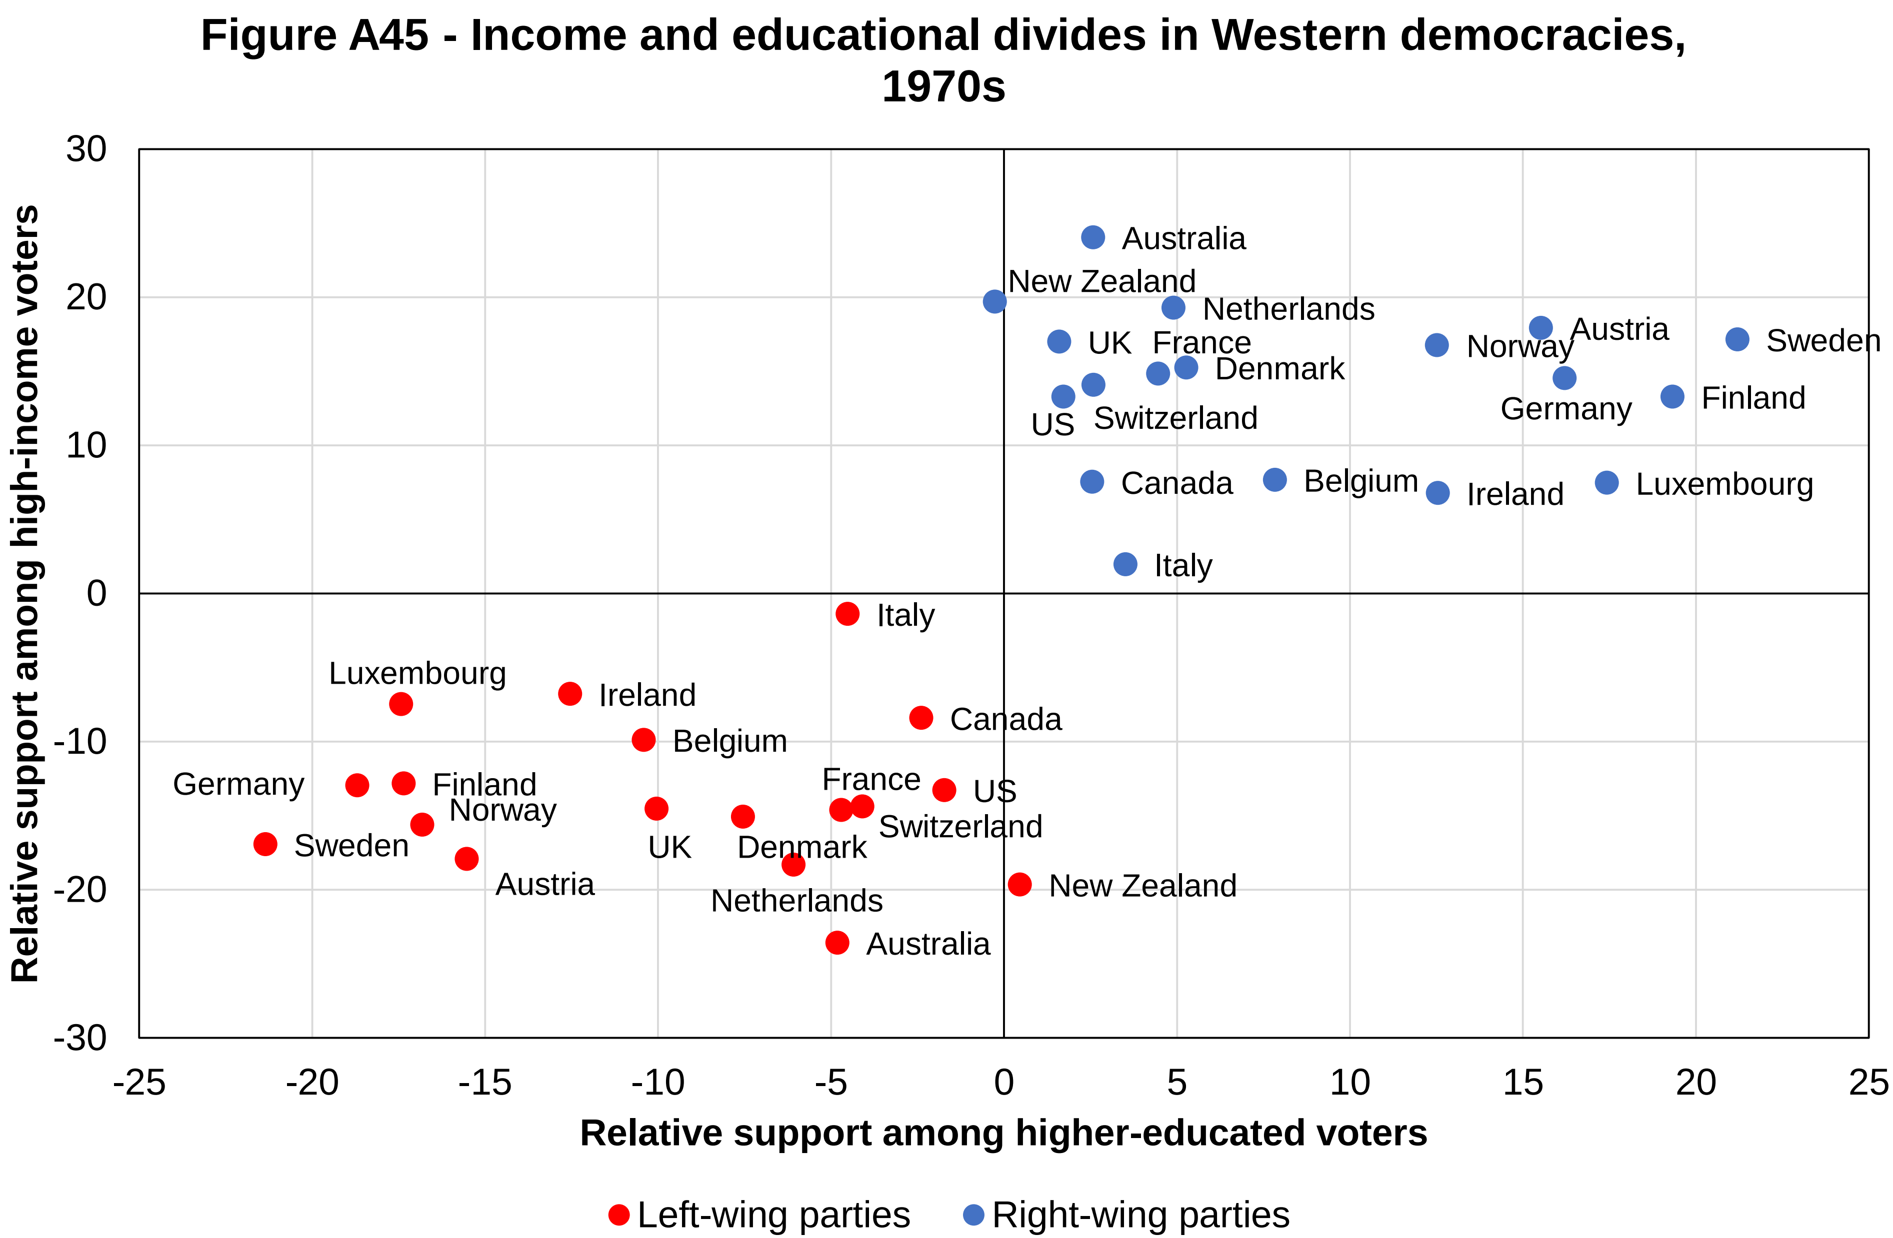
Figure A45 - Income and educational divides in Western democracies, 1970s
Relative support among high-income voters
Relative support among higher-educated voters
30
20
10
0
-10
-20
-30
-25
-20
-15
-10
-5
0
5
10
15
20
25
Australia
New Zealand
Netherlands
UK
France
Denmark
Norway
Austria
Sweden
Germany
Finland
US
Switzerland
Canada
Belgium
Ireland
Luxembourg
Italy
Italy
Luxembourg
Ireland
Canada
Belgium
Germany
Finland
Norway
France
US
Switzerland
Sweden
Austria
UK
Denmark
Netherlands
New Zealand
Australia
Left-wing parties
Right-wing parties

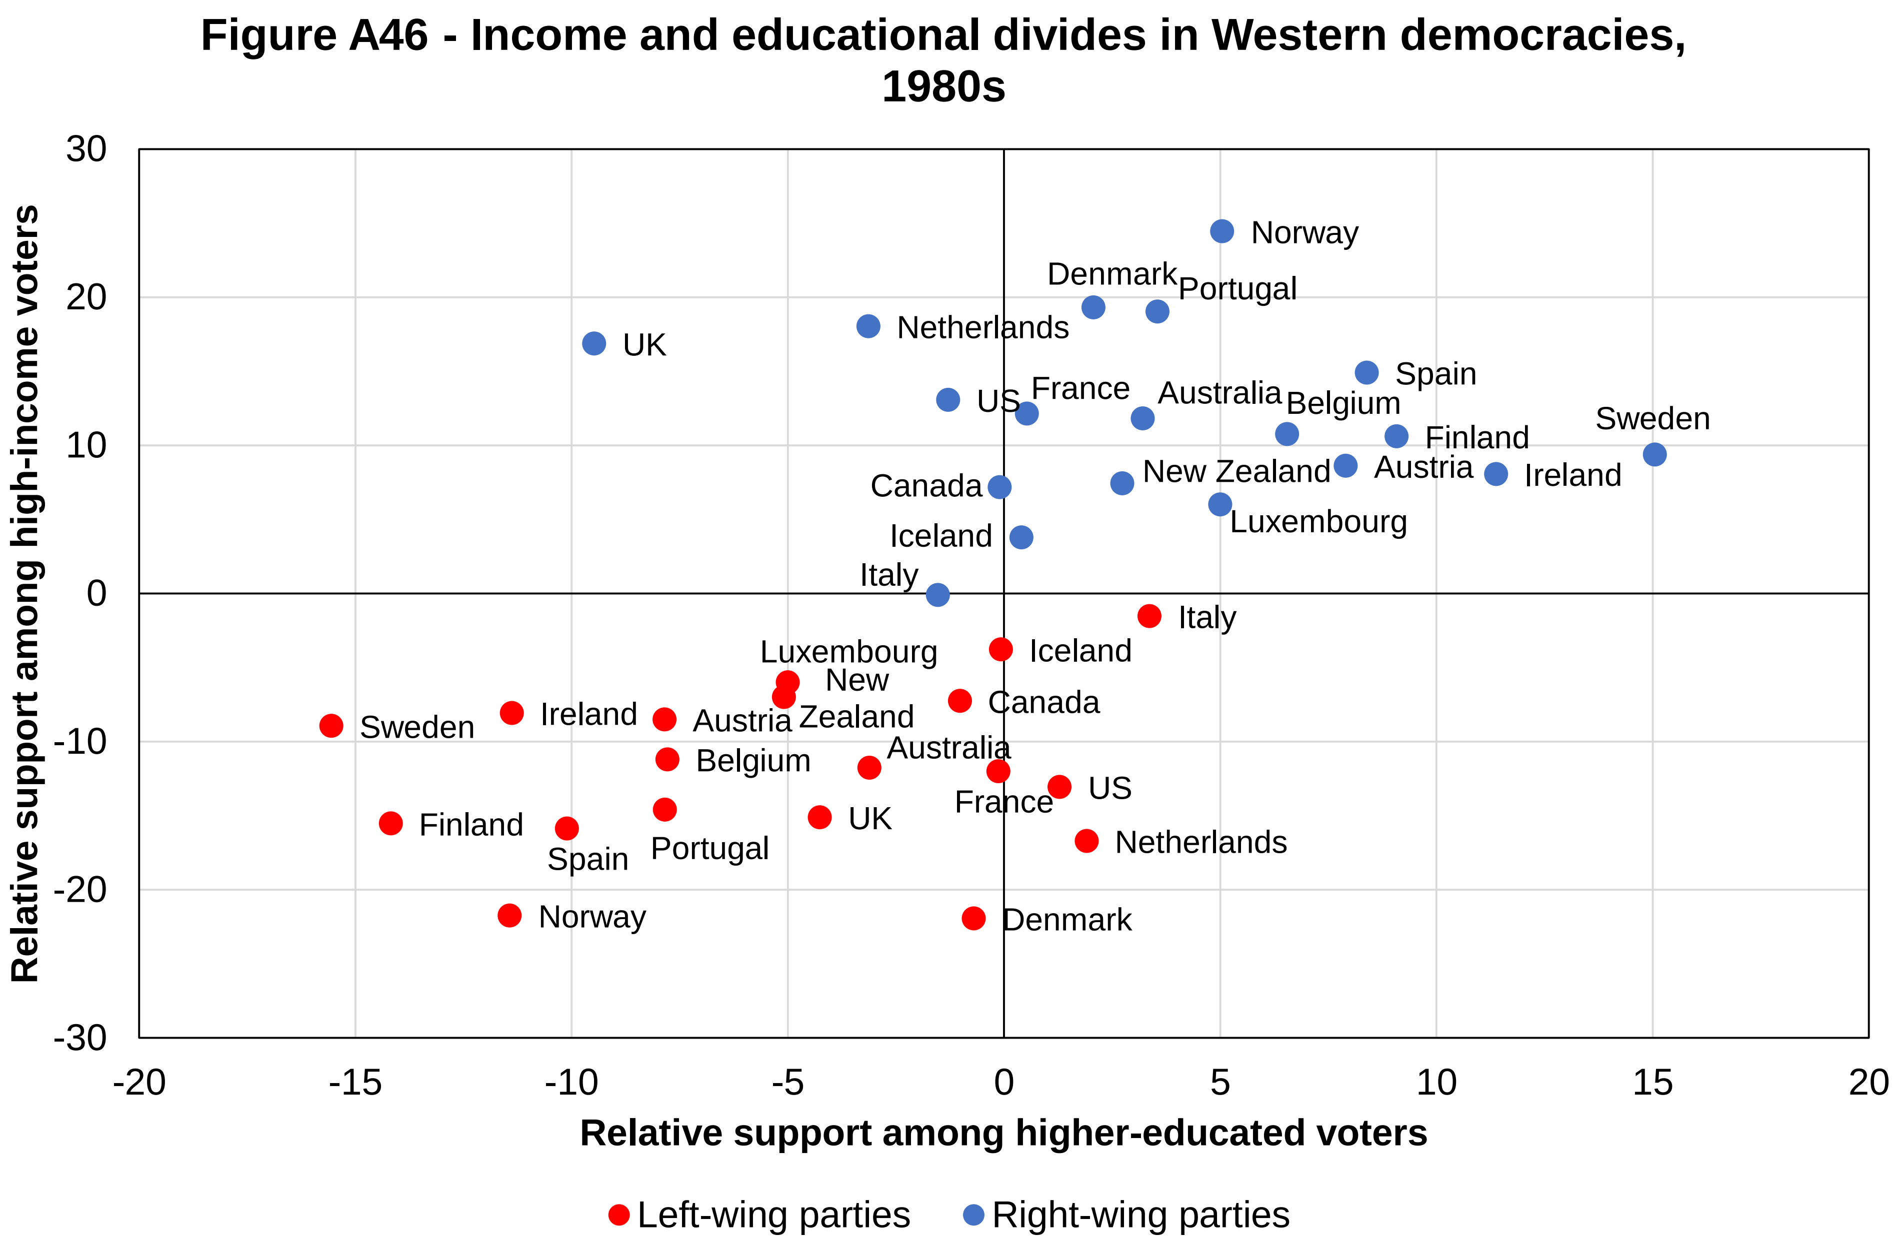
Figure A46 - Income and educational divides in Western democracies, 1980s
Relative support among high-income voters
Relative support among higher-educated voters
30
20
10
0
-10
-20
-30
-20
-15
-10
-5
0
5
10
15
20
Norway
Denmark
Portugal
UK
Netherlands
Spain
US
France
Australia
Belgium
Sweden
Finland
Canada
New Zealand
Austria
Ireland
Luxembourg
Iceland
Italy
Italy
Iceland
Luxembourg
New
Zealand
Canada
Sweden
Ireland
Austria
Belgium
Australia
US
France
Finland
UK
Portugal
Spain
Netherlands
Norway
Denmark
Left-wing parties
Right-wing parties

Figure A47 - Income and educational divides in Western democracies, 1990s

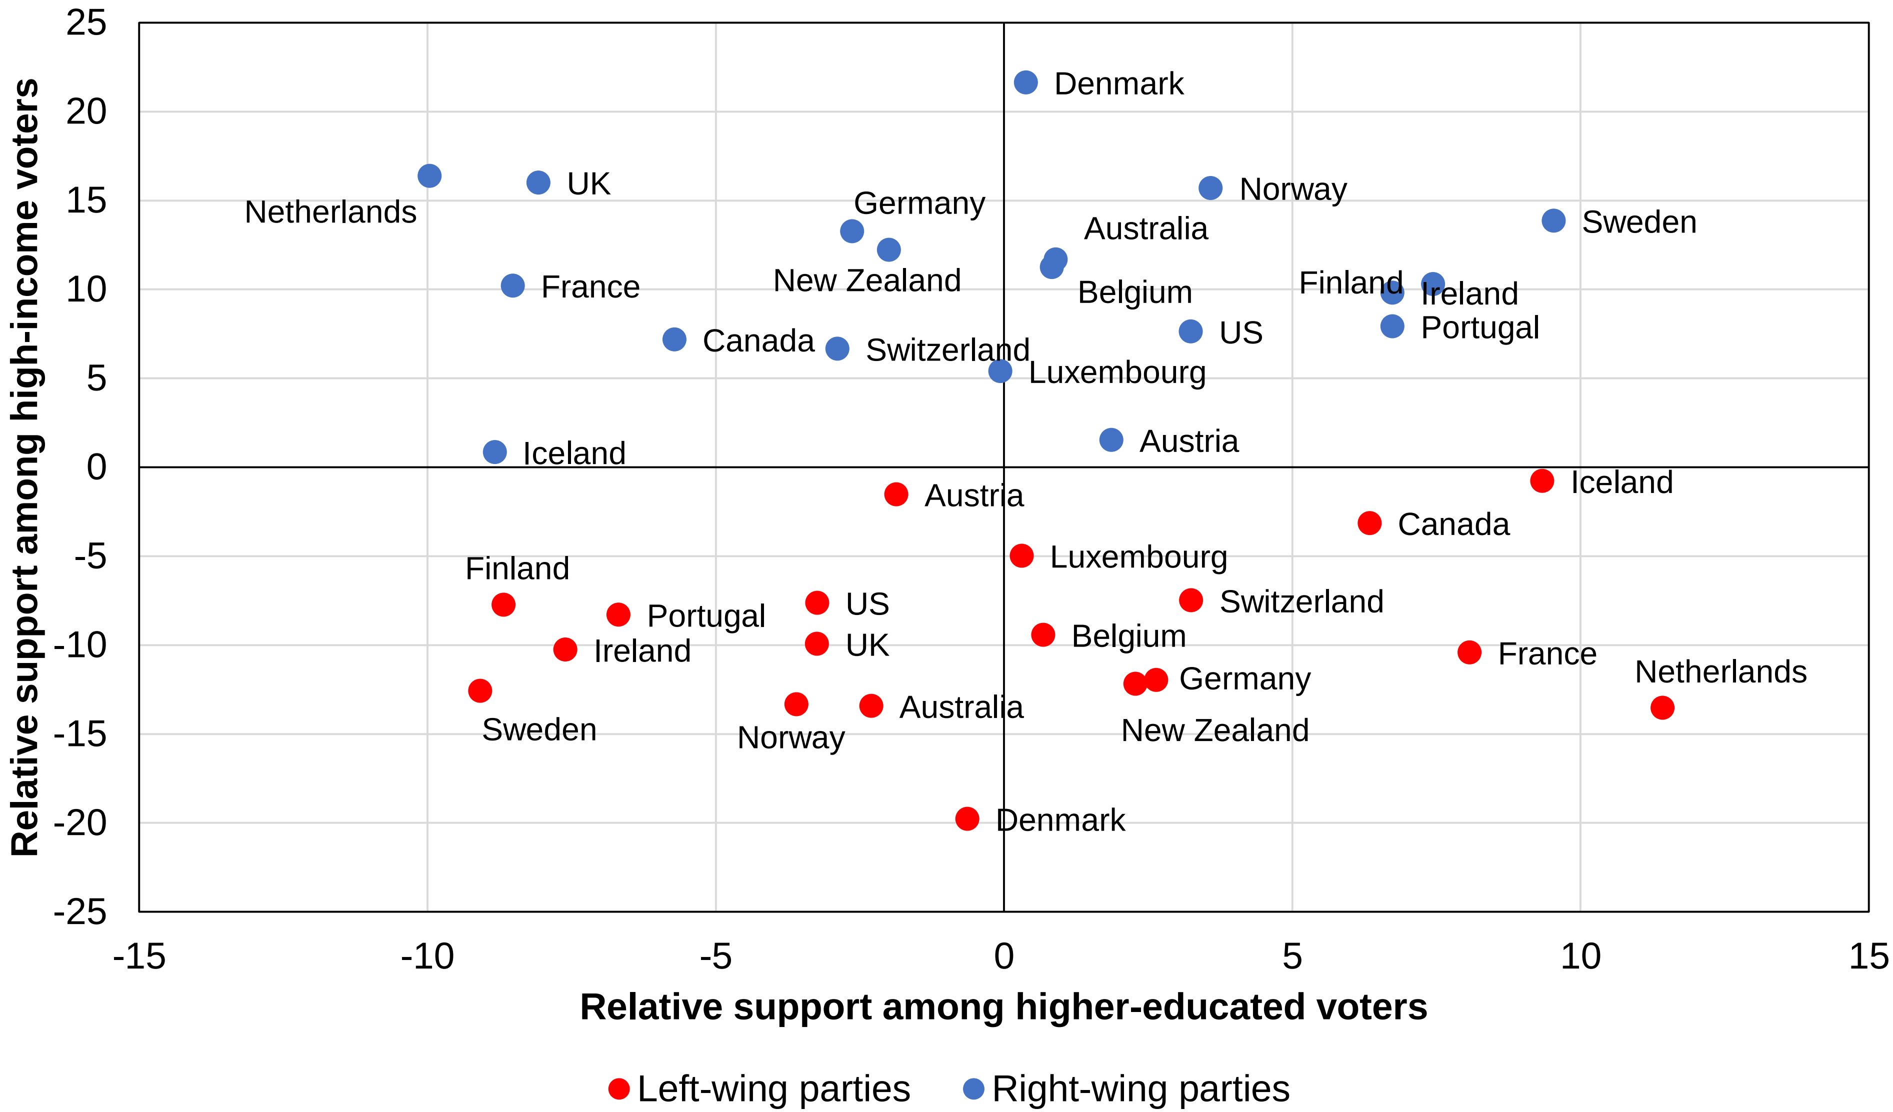

Figure A48 - Income and educational divides in Western democracies, 2000s

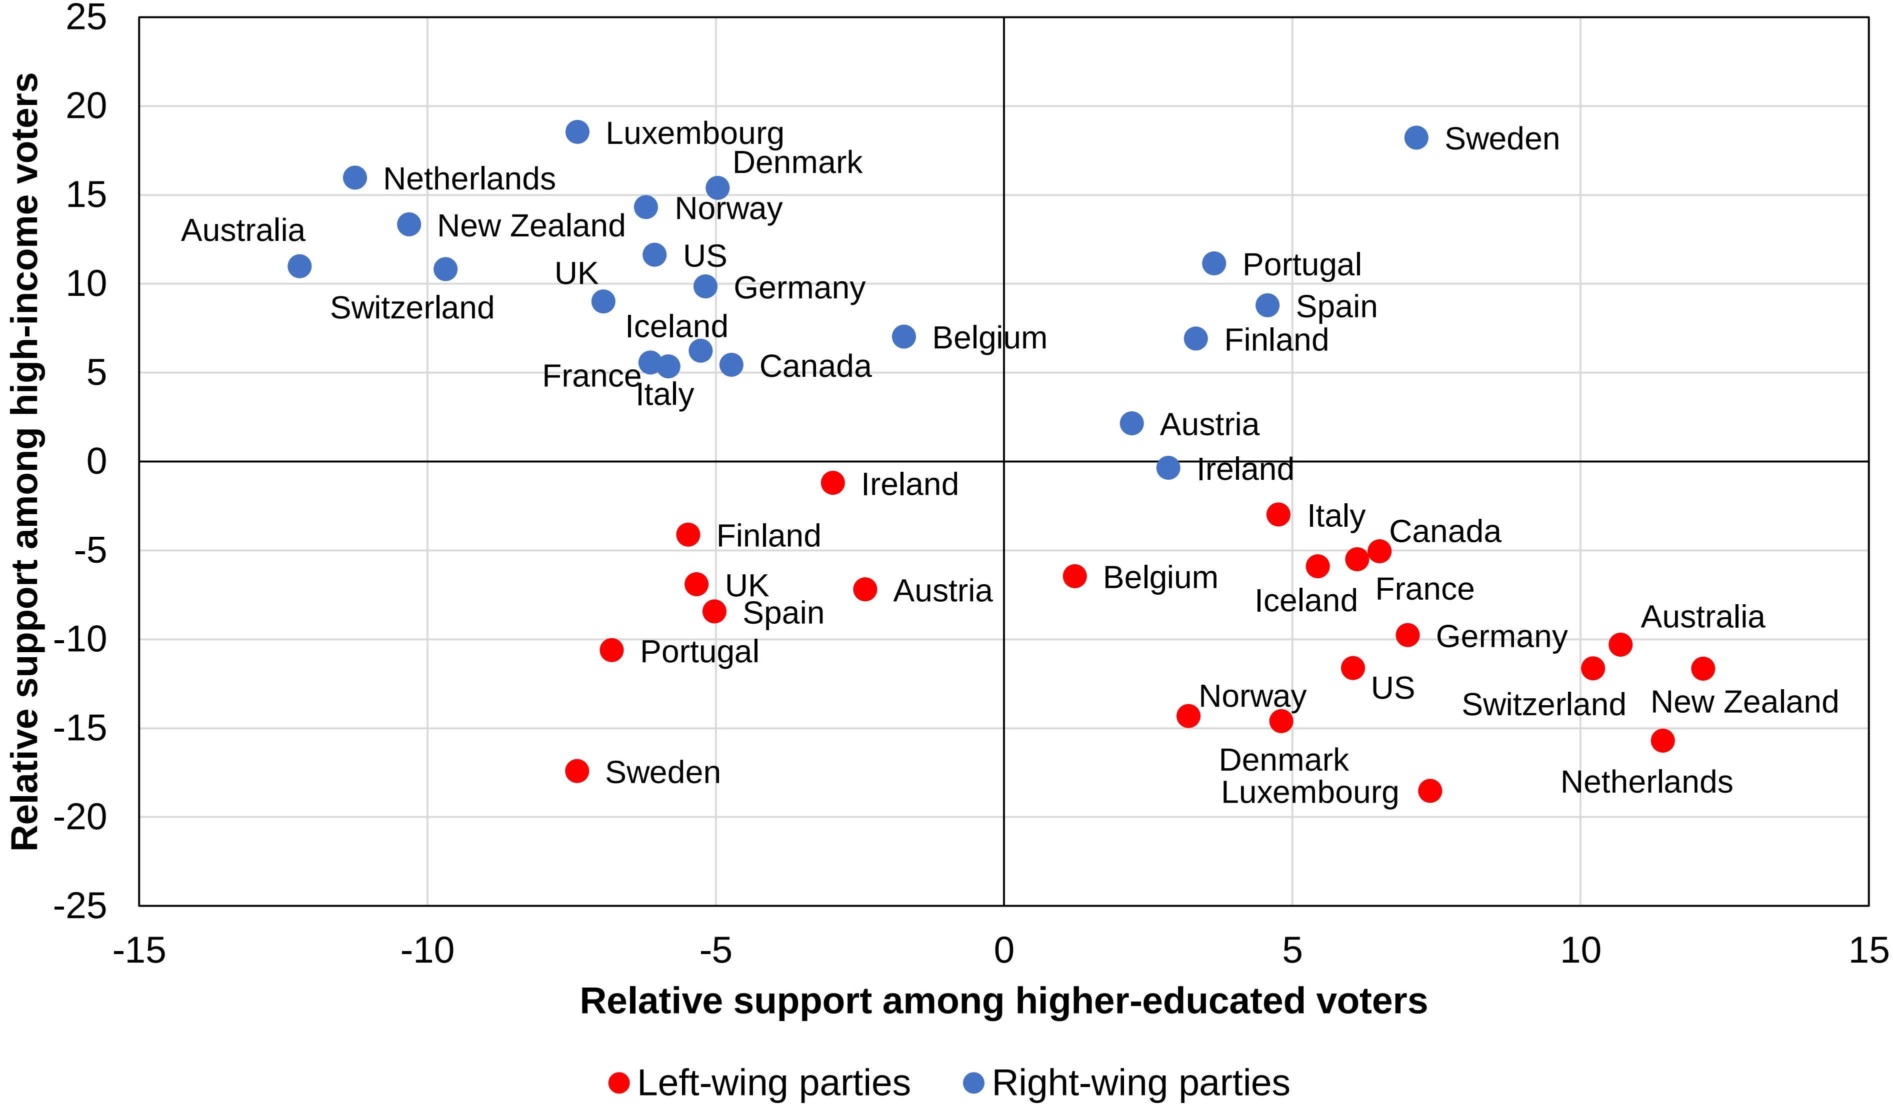

Figure A49 - Income and educational divides in Western democracies, 2010s

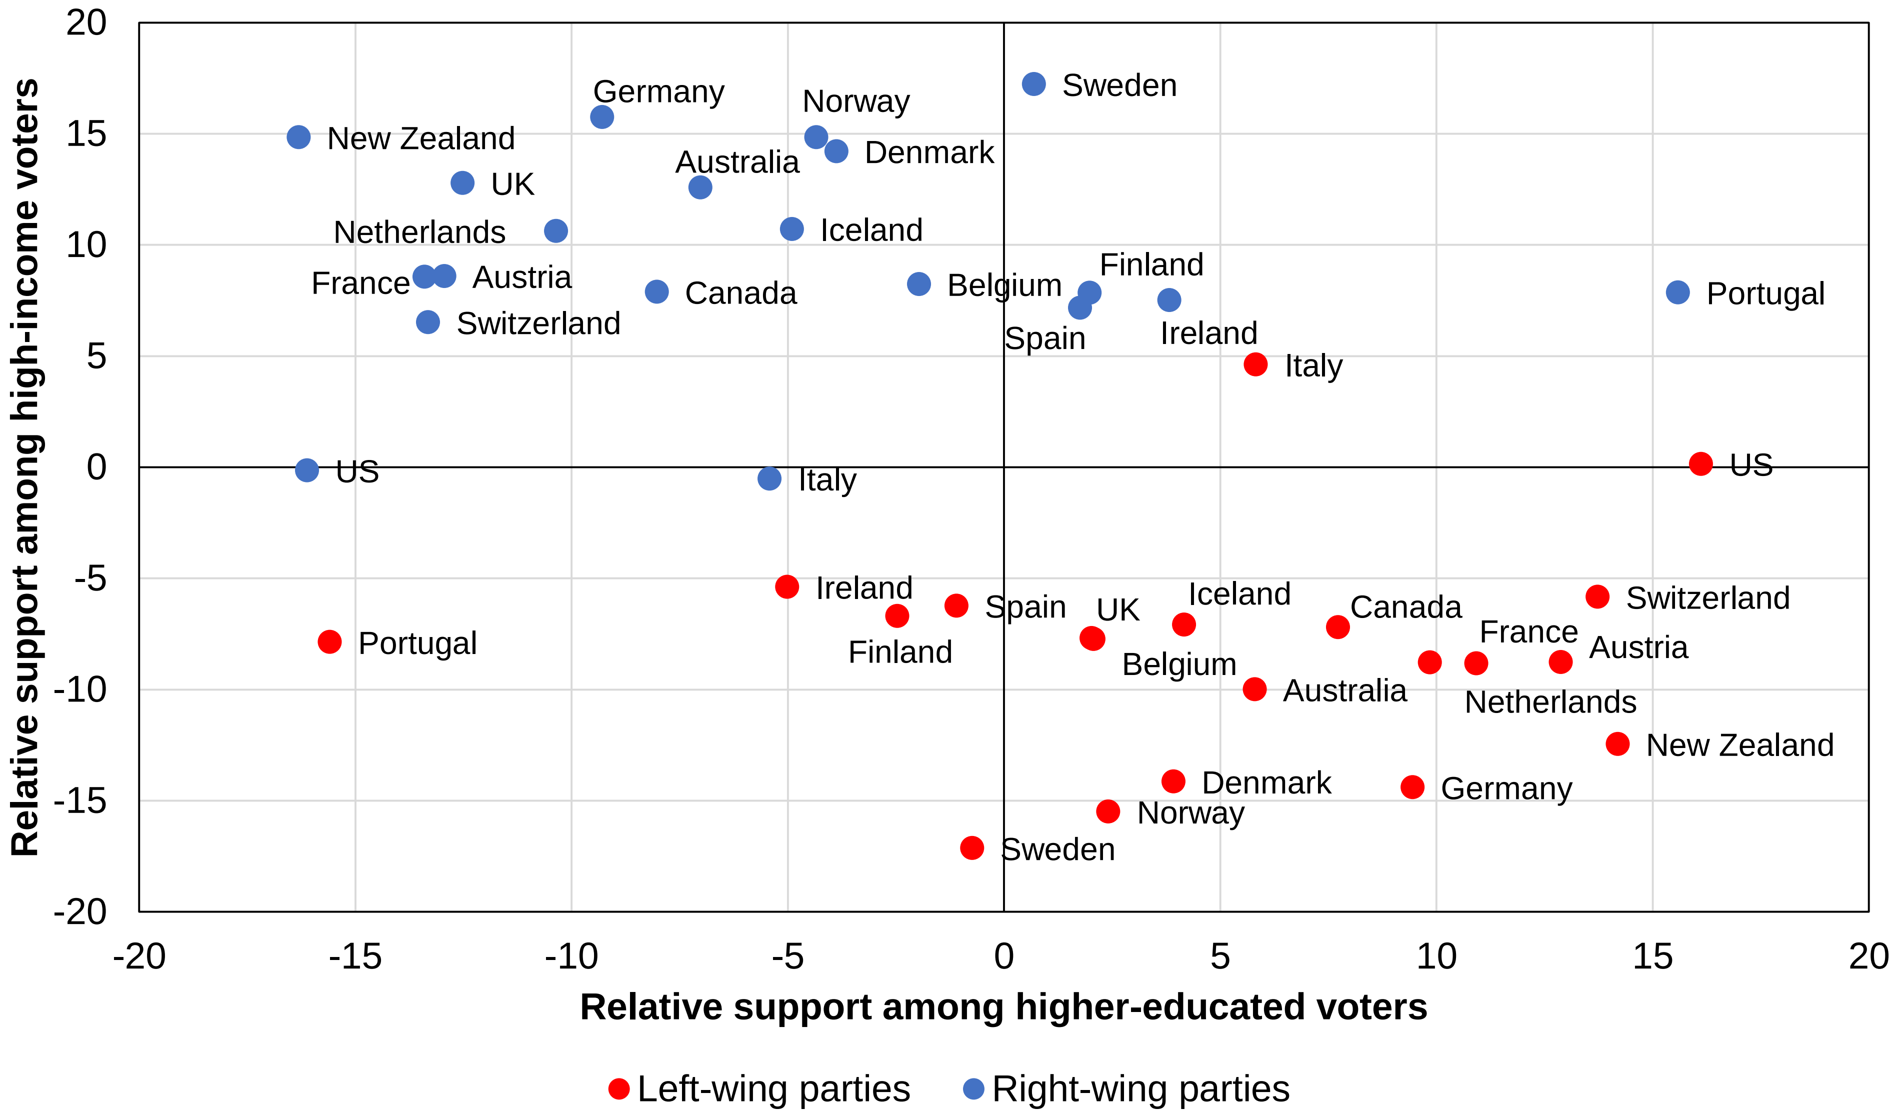

Figure A50 - Correlation between income and education

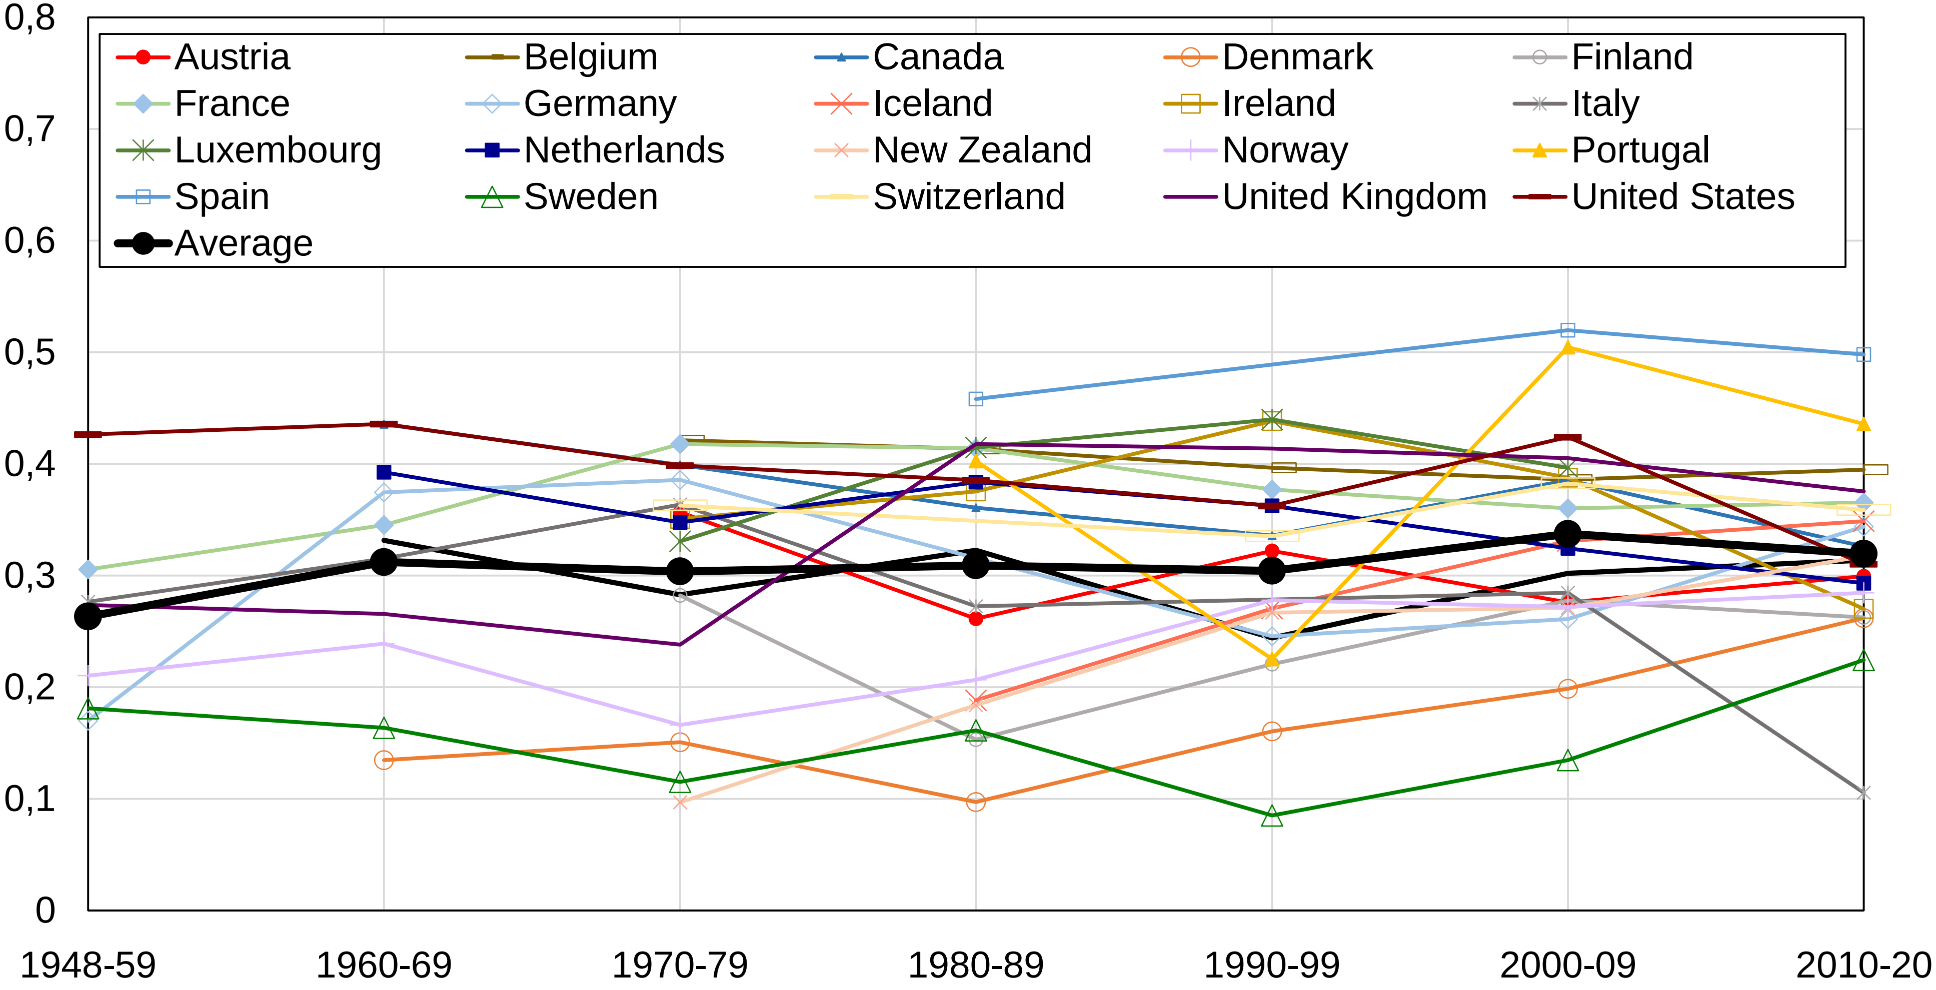

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the correlation between income and education in post-electoral surveys in all Western democracies. Income is defined as the rank (quantile group) to which individuals belong, computed directly from raw income brackets. Education is defined as education deciles, computed from available educational categories (see methodology).

Figure A51 - Kitagawa-Oaxaca-Blinder decomposition of the education cleavage

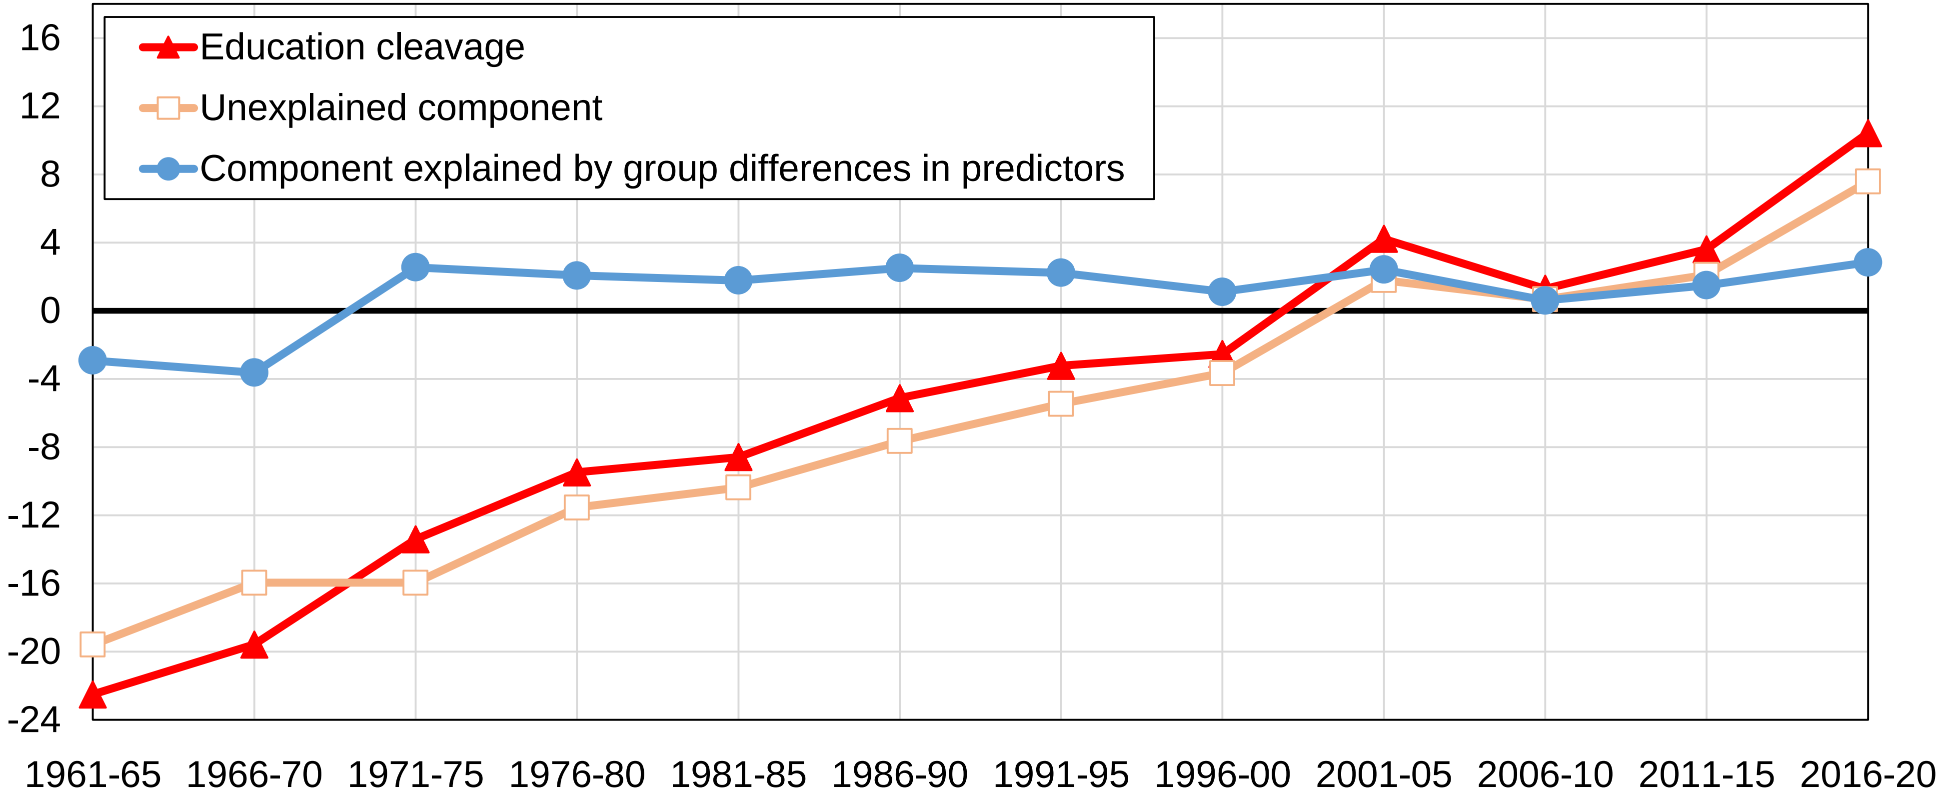

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents a two-way Kitagawa-Oaxaca-Blinder decomposition of the educational cleavage by five-year interval, separating it into a component explained by group differences in predictors (that is, differences in the composition of educational groups in terms of income, gender, age, religion, religious practice, rural/urban location, region, employment and marital status, private/public sector of employment, union membership, and home ownership) and an unexplained component. The unexplained component is very close to the actual indicator, revealing that the reversal of educational divides cannot be accounted for by changes in the composition of education groups. The decomposition is computed after pooling surveys covering the following countries: Australia, Denmark, Finland, France, the Netherlands, Norway, New Zealand, Sweden, the United Kingdom, and the United States. All estimates include election (country-year) fixed effects.

Table B1 - Bakker-Hobolt modified Comparative Manifesto Project measures

A. Economic-distributive dimension	
Pro-free-market emphases	**Pro-redistribution emphases**
Free entreprise	Regulate capitalism
Economic incentives	Economic planning
Anti-protectionism	Pro-protectionism
Social services limitation	Social services expansion
Education limitation	Education expansion
Productivity: positive	Nationalization
Economic orthodoxy: positive	Controlled economy
Labour groups: negative	Labour groups: positive
	Corporatism: positive
	Keynesian demand management: positive
	Marxist analysis: positive
	Social justice
B. Sociocultural dimension	
Conservative emphases	**Progressive emphases**
Political authority	Environmental protection
National way of life: positive	National way of life: negative
Traditional morality: positive	Traditional morality: negative
Law and order	Culture
Multiculturalism: negative	Multiculturalism: positive
Social harmony	Anti-growth
	Underprivileged minority groups
	Non-economic demographic groups: positive
	Freedom-human rights
	Democracy

Source: adapted from R. Bakker and S. B. Hobolt, "Measuring Party Positions," in G. Evans and N. D. de Graaf (ed.), *Political Choice Matters: Explaining the Strength of Class and Religious Cleavages in Cross-National Perspective* , Oxford University Press, 2013, 38. For more detail on the content of each category and the Manifesto Project methodology, see https://manifesto-project.wzb.eu/.

Table B2 - Ideological polarization in Western democracies, 1945-2020

	Economic-distributive index				Sociocultural index			
	Social Democrats	Conservatives	Anti-immigration	Greens	Social Democrats	Conservatives	Anti-immigration	Greens
1945-59	-12,3	11,2			-2,2	2,2		
1960-69	-9,1	9,2			-1,1	0,9		
1970-79	-9,3	8,8	17,6		-0,6	0,6	3,9	
1980-89	-10,9	10,9	15,8	-8,5	-1,9	2,5	3,4	-24,1
1990-99	-9,9	8,2	11,6	-11,5	-3,6	5,2	7,1	-25,4
2000-09	-9,4	8,1	10,4	-6,8	-4,9	6,3	11,2	-24,8
2010-20	-13,5	11,2	8,7	-11,2	-5,4	4,4	20,4	-25,1

Source: authors' computations using the Comparative Manifesto Project database.
Note: the table displays the average economic-distributive and sociocultural scores by decade for four families of parties across all Western democracies: social democratic, socialist and other left-wing parties; conservative, Christian democratic, and liberal parties; anti-immigration parties; and green parties. Negative values on the economic-distributive index correspond to greater proportions of pro-redistribution emphases relatively to pro-free-market emphases in party manifestos. Negative values on the sociocultural index correspond to greater proportions of progressive emphases relatively to conservative emphases. Indices are normalized by the average score by decade so as to better highlight the dynamics of polarization.

Table B3 - Sources of ideological polarization in Western democracies in the 2010s				
	Greens	Social Democrats	Conservatives	Anti-immigration
Sociocultural dimension				
Conservative emphases				
Law and order +	1,4	3,0	5,2	8,5
Political authority	1,4	2,9	2,9	3,1
Civic mindedness +	1,2	1,3	1,7	0,8
National way of life +	0,8	1,1	2,4	9,0
Traditional morality +	0,3	0,5	1,4	2,4
Multiculturalism -	0,2	0,3	1,0	5,0
Progressive emphases				
Environmentalism +	13,4	5,8	4,3	3,0
Democracy	3,2	3,2	2,0	4,4
Anti-growth economy +	6,9	2,8	1,9	0,8
Culture +	2,5	2,4	2,1	1,6
Freedom & human rights	3,7	1,8	2,4	2,2
Non-economic demographic groups	1,1	1,4	1,3	1,1
Multiculturalism +	1,5	1,1	0,8	0,3
Minority groups	0,7	0,7	0,5	0,4
Traditional morality -	0,9	0,6	0,3	0,1
National way of life -	1,1	0,4	0,5	0,1
Economic-distributive dimension				
Pro-free-market emphases				
Incentives	1,2	2,1	3,7	2,0
Economic growth +	0,6	1,8	3,0	0,6
Economic orthodoxy	0,6	1,2	2,9	1,1
Protectionism -	0,1	0,3	0,6	0,2
Free market economy	0,5	0,3	2,7	2,5
Welfare -	0,2	0,2	1,5	1,7
Labour groups -	0,0	0,0	0,4	0,2
Education -	0,0	0,0	0,2	0,3
Pro-redistribution emphases				
Welfare +	11,1	12,8	9,0	8,4
Equality +	9,6	8,8	4,4	3,1
Education +	6,1	6,5	5,4	3,9
Labour groups +	4,5	5,8	3,0	2,3
Market regulation	3,4	4,9	3,0	2,8
Controlled economy	0,9	1,0	0,3	0,5
Nationalisation	0,7	0,8	0,2	0,3
Keynesian demand management	0,2	0,6	0,2	0,2
Economic planning	0,3	0,4	0,5	0,2
Corporatism/mixed economy	0,3	0,3	0,3	0,1
Protectionism +	0,3	0,3	0,3	0,7
Other categories				
Technology & infrastructure	6,0	6,9	7,6	4,2
Gov-admin efficiency	1,5	2,8	4,7	3,3
Internationalism +	2,5	2,6	2,3	1,4
Decentralisation	1,3	1,5	1,8	1,3
Europe +	1,1	1,3	1,6	0,2

Agriculture +	1,2	1,2	2,0	2,2
Military +	0,2	1,1	2,1	2,5
Economic goals	0,5	1,1	1,2	0,6
Political corruption	0,8	0,8	0,5	0,5
Military -	0,9	0,5	0,1	0,2
Peace	0,4	0,4	0,2	0,0
Europe -	0,4	0,3	0,6	6,3
Foreign special +	0,1	0,3	0,3	0,0
Constitution -	0,2	0,3	0,2	0,2
Middle class and prof. groups	0,1	0,2	0,4	0,2
Constitution +	0,2	0,2	0,4	0,1
Internationalism -	0,1	0,1	0,2	1,5
Anti-imperalism	0,1	0,1	0,0	0,1
Marxist analysis +	0,0	0,1	0,0	0,0
Centralisation	0,3	0,1	0,3	0,4
Foreign special -	0,0	0,0	0,0	0,0

Note: The table reports the scores of green parties, social democratic / socialist / communist / other left-wing parties, conservative / Christian democratic / liberal parties, and anti-immigration parties on all the items available in the Comparative Manifesto Project database over the 2010-2020 period. Values correspond to the share of "quasi-sentences" dedicated to emphasizing each category of issues in parties' manifestos. Vote-share-weighted average over all parties with available data in the corresponding decade.

Table B4 - Manifesto scores of anti-immigration parties

	1970s	1980s	1990s	2000s	2010s
Sociocultural dimension					
Conservative emphases					
National way of life +	0,6	2,0	4,2	4,7	9,0
Law and order +	1,2	3,3	5,4	7,5	8,5
Multiculturalism -	0,0	0,9	0,9	4,0	5,0
Political authority	2,7	2,8	4,7	3,9	3,1
Traditional morality +	1,7	2,6	3,5	2,3	2,4
Civic mindedness +	1,0	0,7	1,2	1,3	0,8
Progressive emphases					
Democracy	2,6	2,6	3,2	2,2	4,4
Environmentalism +	3,8	4,6	4,3	4,2	3,0
Freedom & human rights	2,5	2,8	4,8	2,7	2,2
Culture +	0,9	2,3	2,1	2,1	1,6
Non-economic demographic groups	3,5	5,0	2,3	1,9	1,1
Anti-growth economy +	0,0	0,0	0,1	0,3	0,8
Minority groups	0,7	1,4	0,7	0,8	0,4
Multiculturalism +	0,1	0,2	0,8	0,2	0,3
Traditional morality -	0,4	0,1	0,0	0,2	0,1
National way of life -	0,0	0,2	0,0	0,0	0,1
Economic-distributive dimension					
Pro-free-market emphases					
Free market economy	5,7	6,3	5,2	3,9	2,5
Incentives	1,4	2,5	3,1	2,9	2,0
Welfare -	1,2	2,8	2,1	1,6	1,7
Economic orthodoxy	5,2	4,9	2,8	2,4	1,1
Economic growth +	1,6	1,0	1,1	1,1	0,6
Education -	0,8	0,2	0,1	0,1	0,3
Protectionism -	0,1	0,6	0,1	0,5	0,2
Labour groups -	0,1	0,5	0,2	0,2	0,2
Pro-redistribution emphases					
Welfare +	4,0	3,1	4,5	6,9	8,4
Education +	2,1	3,0	3,3	4,3	3,9
Equality +	3,6	1,2	2,3	3,3	3,1
Market regulation	0,8	1,8	1,8	1,1	2,8
Labour groups +	1,0	1,0	1,3	1,9	2,3
Protectionism +	0,1	0,4	0,5	0,9	0,7
Controlled economy	0,4	0,7	0,2	0,5	0,5
Nationalisation	0,1	0,0	0,1	0,2	0,3
Keynesian demand management	0,1	0,0	0,0	0,2	0,2
Economic planning	0,4	0,2	0,2	0,1	0,2
Corporatism/mixed economy	0,7	0,5	0,0	0,0	0,1
Other categories					
Europe -	0,1	0,3	1,1	2,7	6,3
Technology & infrastructure	2,5	3,0	3,5	5,7	4,2
Gov-admin efficiency	5,7	4,9	6,7	4,5	3,3
Military +	1,1	2,7	2,4	2,7	2,5
Agriculture +	2,3	2,0	2,7	2,1	2,2

Internationalism -	1,3	1,4	0,7	1,6	1,5
Internationalism +	2,5	1,8	1,8	1,1	1,4
Decentralisation	0,8	1,7	2,5	2,5	1,3
Economic goals	2,7	2,6	1,6	1,0	0,6
Political corruption	0,2	2,0	3,2	0,8	0,5
Centralisation	0,6	0,5	0,1	0,1	0,4
Middle class and prof. groups	0,6	0,3	0,5	0,4	0,2
Constitution -	0,6	0,2	0,2	0,1	0,2
Military -	0,3	0,1	0,2	0,1	0,2
Europe +	0,5	0,7	0,9	0,6	0,2
Constitution +	1,3	0,3	1,0	0,3	0,1
Anti-imperalism	0,0	0,0	0,0	0,0	0,1
Peace	0,0	0,3	0,2	0,1	0,0
Foreign special +	0,1	0,2	0,2	0,3	0,0
Marxist analysis +	0,0	0,0	0,0	0,0	0,0
Foreign special -	0,1	0,0	0,0	0,0	0,0

Note: The table reports the scores of anti-immigration parties on all the items available in the Comparative Manifesto Project database. Values correspond to the share of "quasi-sentences" dedicated to emphasizing each category of issues in parties' manifestos. Vote-share-weighted average over all parties with available data in the corresponding decade. Figure are ranked in decreasing order of their magnitude in the 2010s.

Table B5 - Manifesto scores of green parties				
	1980s	**1990s**	**2000s**	**2010s**
Sociocultural dimension				
Conservative emphases				
Law and order +	1,1	1,1	2,0	1,4
Political authority	18,6	4,1	2,7	1,4
Civic mindedness +	1,0	1,4	1,4	1,2
National way of life +	0,1	0,7	0,6	0,8
Traditional morality +	0,3	0,4	0,3	0,3
Multiculturalism -	0,1	0,1	0,2	0,2
Progressive emphases				
Environmentalism +	12,2	16,7	13,3	13,4
Anti-growth economy +	2,9	2,7	3,9	6,9
Freedom & human rights	2,3	2,1	2,5	3,7
Democracy	6,7	6,2	4,0	3,2
Culture +	1,8	2,2	2,6	2,5
Multiculturalism +	0,3	0,8	1,2	1,5
Non-economic demographic groups	3,3	4,4	2,7	1,1
National way of life -	0,1	0,3	0,1	1,1
Traditional morality -	0,4	0,7	0,8	0,9
Minority groups	1,7	2,4	2,0	0,7
Economic-distributive dimension				
Pro-free-market emphases				
Incentives	0,7	0,7	2,4	1,2
Economic growth +	1,2	0,9	2,2	0,6
Economic orthodoxy	0,5	1,0	0,5	0,6
Free market economy	0,2	0,2	0,2	0,5
Welfare -	0,0	0,0	0,1	0,2
Protectionism -	0,0	0,1	0,1	0,1
Education -	0,0	0,0	0,0	0,0
Labour groups -	0,0	0,0	0,0	0,0
Pro-redistribution emphases				
Welfare +	5,3	8,1	8,1	11,1
Equality +	2,7	7,7	9,7	9,6
Education +	1,5	3,2	4,5	6,1
Labour groups +	4,2	2,6	3,4	4,5
Market regulation	1,2	2,4	2,4	3,4
Controlled economy	0,3	0,6	0,7	0,9
Nationalisation	0,4	0,6	0,8	0,7
Protectionism +	0,2	0,3	0,2	0,3
Corporatism/mixed economy	0,0	0,1	0,2	0,3
Economic planning	0,3	0,4	0,1	0,3
Keynesian demand management	0,1	0,0	0,2	0,2
Other categories				
Technology & infrastructure	2,3	3,7	4,8	6,0
Internationalism +	3,4	3,6	3,2	2,5
Gov-admin efficiency	2,0	2,7	2,3	1,5
Decentralisation	2,2	1,3	1,3	1,3
Agriculture +	1,4	1,7	2,1	1,2

Europe +	0,5	1,2	1,5	1,1
Military -	3,0	1,6	1,5	0,9
Political corruption	1,9	0,5	0,8	0,8
Economic goals	1,8	2,7	1,5	0,5
Europe -	0,7	1,1	0,7	0,4
Peace	1,7	0,6	0,8	0,4
Centralisation	0,0	0,0	0,2	0,3
Constitution -	0,1	0,2	0,0	0,2
Military +	0,4	0,2	0,2	0,2
Constitution +	1,4	0,4	0,2	0,2
Middle class and prof. groups	0,3	0,3	0,5	0,1
Internationalism -	0,1	0,2	0,4	0,1
Foreign special +	0,4	0,1	0,1	0,1
Anti-imperialism	0,7	0,1	0,1	0,1
Marxist analysis +	0,7	0,0	0,2	0,0
Foreign special -	0,0	0,0	0,0	0,0

Note: The table reports the scores of green parties on all the items available in the Comparative Manifesto Project database. Values correspond to the share of "quasi-sentences" dedicated to emphasizing each category of issues in parties' manifestos. Vote-share-weighted average over all parties with available data in the corresponding decade. Figure are ranked in decreasing order of their magnitude in the 2010s.

Table B6 - Manifesto scores of Social Democratic / Socialist / Other left-wing parties							
	1950s	**1960s**	**1970s**	**1980s**	**1990s**	**2000s**	**2010s**
Sociocultural dimension							
Conservative emphases							
Law and order +	0,2	0,6	1,4	1,9	2,7	4,2	3,0
Political authority	3,5	2,3	3,4	2,9	4,3	4,0	2,9
Civic mindedness +	2,3	1,6	1,7	2,1	2,1	1,6	1,3
National way of life +	0,8	0,9	0,4	0,6	0,8	1,5	1,1
Traditional morality +	0,4	0,3	0,3	0,4	1,0	0,7	0,5
Multiculturalism -	0,6	0,3	0,3	0,1	0,1	0,4	0,3
Progressive emphases							
Environmentalism +	0,4	1,1	3,0	4,7	5,9	5,6	5,8
Democracy	2,8	2,2	5,9	3,0	2,9	2,5	3,2
Anti-growth economy +	0,0	0,0	0,0	0,2	0,5	1,0	2,8
Culture +	1,2	2,2	1,8	2,9	2,6	3,0	2,4
Freedom & human rights	2,2	1,5	2,0	2,1	1,6	1,2	1,8
Non-economic demographic groups	4,4	5,1	4,4	5,4	4,4	3,2	1,4
Multiculturalism +	0,4	0,5	0,5	0,5	0,6	0,9	1,1
Minority groups	0,6	0,9	1,0	1,4	1,4	1,7	0,7
Traditional morality -	0,1	0,1	0,2	0,2	0,3	0,3	0,6
National way of life -	0,2	0,2	0,2	0,1	0,1	0,1	0,4
Economic-distributive dimension							
Pro-free-market emphases							
Incentives	1,7	2,3	2,0	2,4	2,9	2,2	2,1
Economic growth +	3,6	4,0	2,3	2,7	2,2	2,2	1,8
Economic orthodoxy	1,2	1,2	1,7	1,8	2,3	1,3	1,2
Protectionism -	0,3	0,4	0,1	0,1	0,3	0,4	0,3
Free market economy	1,0	0,6	0,4	0,6	0,8	0,7	0,3
Welfare -	0,0	0,1	0,1	0,1	0,3	0,1	0,2
Labour groups -	0,0	0,0	0,1	0,0	0,0	0,0	0,0
Education -	0,0	0,3	0,0	0,0	0,0	0,0	0,0
Pro-redistribution emphases							
Welfare +	8,2	9,6	7,2	8,3	9,3	11,6	12,8
Equality +	4,7	3,9	6,3	5,7	7,1	6,4	8,8
Education +	3,4	4,9	3,9	3,9	5,0	7,2	6,5
Labour groups +	5,0	4,0	4,1	3,8	3,8	3,3	5,8
Market regulation	2,7	2,3	2,6	1,9	2,1	2,5	4,9
Controlled economy	1,9	1,5	1,9	1,1	0,7	0,6	1,0
Nationalisation	1,3	0,8	1,1	0,8	0,5	0,5	0,8
Keynesian demand management	0,4	0,4	0,3	0,5	0,3	0,2	0,6
Economic planning	2,6	2,4	2,9	1,3	0,5	0,7	0,4
Corporatism/mixed economy	0,2	0,4	0,4	0,4	0,5	0,4	0,3
Protectionism +	0,5	0,4	0,6	0,3	0,2	0,1	0,3
Other categories							
Technology & infrastructure	4,2	5,0	4,6	5,5	6,1	6,9	6,9
Gov-admin efficiency	1,3	1,8	2,0	3,1	4,2	3,9	2,8
Internationalism +	2,3	2,7	2,4	2,8	3,2	3,7	2,6
Decentralisation	1,0	1,7	1,7	1,8	1,9	1,8	1,5
Europe +	0,4	1,0	0,7	1,2	2,0	1,7	1,3

Agriculture +	5,9	4,9	2,8	2,8	2,0	1,4	1,2
Military +	1,2	1,9	0,9	0,9	0,6	1,2	1,1
Economic goals	3,7	2,8	4,8	3,3	2,7	1,8	1,1
Political corruption	0,5	0,3	0,4	0,5	0,6	0,4	0,8
Military -	1,6	1,5	0,8	1,4	0,6	0,6	0,5
Peace	2,1	1,7	1,2	2,2	0,5	0,7	0,4
Europe -	0,0	0,2	0,4	0,2	0,3	0,2	0,3
Foreign special +	1,7	1,3	0,9	1,0	0,4	0,6	0,3
Constitution -	0,7	0,6	0,2	0,3	0,0	0,0	0,3
Middle class and prof. groups	1,5	1,1	0,7	0,7	0,2	0,5	0,2
Constitution +	1,0	0,5	0,6	0,7	0,4	0,4	0,2
Internationalism -	0,5	0,4	0,4	0,1	0,1	0,2	0,1
Anti-imperalism	0,4	0,4	0,6	0,3	0,1	0,1	0,1
Marxist analysis +	0,1	0,0	0,2	0,1	0,1	0,1	0,1
Centralisation	0,2	0,5	0,5	0,1	0,1	0,2	0,1
Foreign special -	0,7	0,5	0,4	0,3	0,0	0,1	0,0

Note: The table reports the scores of social democratic, socialist, and other left-wing parties (excluding Greens) on all the items available in the Comparative Manifesto Project database. Values correspond to the share of "quasi-sentences" dedicated to emphasizing each category of issues in parties' manifestos. Vote-share-weighted average over all parties with available data in the corresponding decade. Figure are ranked in decreasing order of their magnitude in the 2010s.

Table B7 - Manifesto scores of Conservative / Christian Democratic / Liberal parties							
	1950s	**1960s**	**1970s**	**1980s**	**1990s**	**2000s**	**2010s**
Sociocultural dimension							
Conservative emphases							
Law and order +	0,7	0,9	1,9	2,4	4,7	6,4	5,2
Political authority	4,4	3,2	3,4	3,4	5,4	6,0	2,9
National way of life +	2,0	1,0	0,8	0,9	1,4	2,3	2,4
Civic mindedness +	3,1	1,4	1,5	1,5	1,7	1,8	1,7
Traditional morality +	3,1	1,6	1,8	2,3	3,1	2,5	1,4
Multiculturalism -	0,2	0,3	0,2	0,1	0,4	1,0	1,0
Progressive emphases							
Environmentalism +	0,4	1,5	3,8	4,5	5,1	4,2	4,3
Freedom & human rights	3,3	2,0	3,0	3,0	2,3	2,2	2,4
Culture +	1,0	2,2	2,8	2,5	2,0	2,0	2,1
Democracy	2,8	2,4	3,8	2,1	2,3	2,0	2,0
Anti-growth economy +	0,0	0,0	0,0	0,1	0,3	0,5	1,9
Non-economic demographic groups	3,9	4,9	4,9	4,0	3,4	2,9	1,3
Multiculturalism +	1,2	0,8	0,6	0,7	0,5	0,7	0,8
Minority groups	0,5	0,7	1,1	1,1	1,1	1,3	0,5
National way of life -	0,1	0,1	0,0	0,0	0,0	0,0	0,5
Traditional morality -	0,1	0,0	0,1	0,1	0,1	0,2	0,3
Economic-distributive dimension							
Pro-free-market emphases							
Incentives	3,0	3,7	3,9	3,9	4,1	3,9	3,7
Economic growth +	3,1	3,9	2,3	3,0	2,2	2,4	3,0
Economic orthodoxy	5,0	3,9	4,3	5,4	4,4	2,5	2,9
Free market economy	5,0	3,8	3,1	4,4	3,5	2,8	2,7
Welfare -	0,4	0,3	0,6	0,8	1,1	0,8	1,5
Protectionism -	0,4	0,5	0,2	0,6	0,3	0,4	0,6
Labour groups -	0,4	0,2	0,5	0,5	0,3	0,1	0,4
Education -	0,0	0,0	0,2	0,0	0,1	0,1	0,2
Pro-redistribution emphases							
Welfare +	4,9	6,0	6,6	5,6	6,1	8,5	9,0
Education +	2,2	4,6	4,0	3,4	4,8	5,7	5,4
Equality +	3,1	3,1	3,9	2,9	3,5	3,1	4,4
Labour groups +	2,4	2,0	1,9	1,7	1,2	2,0	3,0
Market regulation	1,4	1,0	1,5	1,4	2,0	2,0	3,0
Economic planning	0,9	1,5	1,1	0,4	0,3	0,4	0,5
Protectionism +	0,5	0,5	0,5	0,2	0,2	0,2	0,3
Controlled economy	0,3	0,4	0,9	0,4	0,4	0,3	0,3
Corporatism/mixed economy	0,6	0,2	0,3	0,3	0,3	0,1	0,3
Keynesian demand management	0,3	0,3	0,6	0,3	0,1	0,1	0,2
Nationalisation	0,2	0,0	0,1	0,0	0,0	0,1	0,2
Other categories							
Technology & infrastructure	3,2	5,8	4,2	5,5	5,2	6,3	7,6
Gov-admin efficiency	1,6	2,5	2,7	4,8	5,2	5,8	4,7
Internationalism +	1,6	2,6	1,8	2,2	3,0	2,4	2,3
Military +	2,1	2,6	1,7	2,6	1,4	1,7	2,1
Agriculture +	5,3	4,5	3,4	4,2	2,7	1,9	2,0

Decentralisation	2,0	2,2	3,2	2,1	1,9	1,9	1,8
Europe +	0,6	1,4	1,3	1,8	2,7	1,7	1,6
Economic goals	2,4	2,2	3,6	3,2	2,5	2,5	1,2
Europe -	0,1	0,1	0,0	0,1	0,5	0,4	0,6
Political corruption	0,4	0,3	0,4	0,3	0,6	0,6	0,5
Middle class and prof. groups	2,3	1,4	1,1	0,7	0,4	0,6	0,4
Constitution +	1,1	0,9	0,7	0,5	0,4	0,3	0,4
Centralisation	0,3	0,3	0,3	0,2	0,1	0,2	0,3
Foreign special +	2,1	1,8	0,9	1,1	0,6	0,7	0,3
Internationalism -	0,4	0,8	0,4	0,2	0,2	0,2	0,2
Constitution -	0,4	0,3	0,2	0,2	0,1	0,0	0,2
Peace	0,9	1,2	0,7	0,9	0,4	0,3	0,2
Military -	0,3	0,3	0,2	0,5	0,2	0,1	0,1
Foreign special -	0,4	0,3	0,1	0,3	0,0	0,1	0,0
Anti-imperalism	0,4	0,5	0,2	0,2	0,0	0,0	0,0
Marxist analysis +	0,0	0,0	0,0	0,0	0,0	0,0	0,0

Note: The table reports the scores of conservative, Christian democratic, and liberal parties on all the items available in the Comparative Manifesto Project database. Values correspond to the share of "quasi-sentences" dedicated to emphasizing each category of issues in parties' manifestos. Vote-share-weighted average over all parties with available data in the corresponding decade. Figure are ranked in decreasing order of their magnitude in the 2010s.

Table B8 - Sociocultural polarization and educational divides: regression results

	Raw coefficient	After controls and country/year fixed effects	After controls and election fixed effects
1948-1979	-0.13*	0.12	0.11
1980-1999	-0.68***	-0.13	-0.21
2000-2020	-1.21***	-0.65***	-0.73***

Source: authors' computations combining the World Political Cleavages and Inequality Datatabase with Manifesto Project data.

Note: the table reports the coefficient associated to a regression of the sociocultural index on the education gradient (the share of top 10% educated voters within a given party's electorate) at the party level, decomposing the dataset into three time periods: 1948-1979, 1980-1999, and 2000-2020. The first column reports the raw coefficient (without controls). The second column reports the coefficient after controlling for country and year fixed effects and for the composition of the electorate of each party in terms of income, age, gender, rural-urban location, and religion. The third column reports the same coefficient after controlling for the same variables and for election fixed effects (that is, interacting country and year fixed effects). * p<0.10, ** p<0.05, *** p<0.01.

Interpretation: in 1948-1979, the link between a party's position on the sociocultural axis and the composition of its electorate in terms of education was small and not statistically significant; in 2000-2020, it has become strongly negative and statistically significant at the 1% level, so that parties strongly emphasizing progressive issues in their manifestos receive much greater support from higher-educated voters.

Table B9 - Sociocultural polarization and educational divides: complete regression results

	Raw coefficient			After controls and country/year fixed effects			After controls and election fixed effects		
	1948-1979	1980-1999	2000-2020	1948-1979	1980-1999	2000-2020	1948-1979	1980-1999	2000-2020
Share of top 10% educated voters in party's electorate	-0.134* (0.079)	-0.681*** (0.103)	-1.208*** (0.118)	0.122 (0.174)	-0.133 (0.200)	-0.651*** (0.205)	0.114 (0.176)	-0.208 (0.193)	-0.733*** (0.207)
R-squared	0.01	0.06	0.14	0.59	0.52	0.43	0.61	0.59	0.47
Observations	444	661	640	159	266	341	159	266	341

Source: authors' computations combining the World Political Cleavages and Inequality Datatabase with Manifesto Project data.
Note: the table reports the results of a regression of the sociocultural index on the education gradient (the share of top 10% educated voters within a given party's electorate) at the party level, decomposing the dataset into three time periods: 1948-1979, 1980-1999, and 2000-2020. The first panel reports the raw coefficient (without controls). The second panel reports the coefficient after controlling for country and year fixed effects and for the composition of the electorate of each party in terms of income, age, gender, rural-urban location, and religion. The third panel reports the same coefficient after controlling for the same variables and for election fixed effects (that is, interacting country and year fixed effects). * p<0.10, ** p<0.05, *** p<0.01.
Interpretation: in 1948-1979, the link between a party's position on the sociocultural axis and the composition of its electorate in terms of education was small and not statistically significant; in 2000-2020, it has become strongly negative and statistically significant at the 1% level, so that parties strongly emphasizing progressive issues in their manifestos receive much greater support from higher-educated voters.

Table B10 - Correlation between income and education gradients and all Manifesto items, 1960s-2010s												
	Relative support among top 10% educated voters						Relative support among top 10% income voters					
	1960s	1970s	1980s	1990s	2000s	2010s	1960s	1970s	1980s	1990s	2000s	2010s
Sociocultural dimension												
Conservative emphases												
Law and order +	0.11	0.11*	-0.08	-0.09*	-0.07	-0.27***	-0.03	0.12*	0.13**	0.25***	0.14**	-0.07
National way of life +	-0.05	-0.11*	-0.05	-0.24***	-0.20***	-0.28***	-0.07	0.04	0.06	-0.15**	-0.10*	-0.09*
Multiculturalism -	-0.09	-0.03	-0.12**	-0.14***	-0.19***	-0.30***	-0.13	-0.01	-0.07	-0.07	-0.08	-0.09*
Traditional morality +	-0.02	-0.04	-0.04	-0.14**	-0.18***	-0.11**	-0.06	-0.02	-0.07	-0.10	-0.19***	-0.11**
Political authority	-0.05	-0.09	0.02	0.04	-0.04	-0.08	-0.08	-0.01	-0.09	0.06	0.07	0.03
Civic mindedness +	0.02	-0.07	-0.11**	-0.04	0.01	0.02	-0.08	0.05	-0.07	-0.02	-0.07	0.07
Progressive emphases												
Culture +	-0.11	-0.09	-0.11**	0.02	-0.03	0.15***	-0.06	-0.04	-0.09	-0.08	0.04	0.10*
Freedom & human rights	0.09	0.11*	0.11**	0.13**	0.16***	0.23***	0.14	0.17***	0.24***	0.19***	0.12**	0.17***
Anti-growth economy +	-0.08	-0.00	0.13**	0.23***	0.22***	0.23***	-0.07	0.02	0.02	-0.00	0.06	-0.06
Environmentalism +	0.02	0.15**	0.28***	0.23***	0.20***	0.27***	-0.10	0.01	-0.02	-0.12*	-0.05	-0.02
Traditional morality -	0.07	0.06	0.09	0.13**	0.21***	0.27***	0.00	0.16**	0.02	0.02	0.04	0.14**
Multiculturalism +	-0.06	0.01	0.07	0.13**	0.13**	0.11**	-0.09	0.12*	0.01	-0.02	0.01	0.03
National way of life -	-0.10	-0.07	0.01	0.01	0.01	0.11**	-0.13	-0.08	-0.03	-0.07	-0.12**	0.02
Non-economic demographic groups	-0.09	-0.07	0.01	0.04	0.01	-0.06	-0.15*	-0.11	-0.06	-0.06	-0.01	-0.05
Minority groups	0.06	0.01	0.03	0.18***	0.16***	0.02	0.01	-0.11*	-0.02	0.05	0.08	-0.04
Democracy	-0.03	0.10	0.06	0.16***	0.08	0.05	-0.08	0.02	-0.15**	-0.06	-0.10*	-0.08
Economic-distributive dimension												
Pro-free-market emphases												
Incentives	-0.05	0.05	-0.11**	-0.11**	-0.08	-0.18***	-0.10	0.17***	0.24***	0.12*	0.17***	0.08
Economic growth +	0.01	-0.12*	-0.08	-0.04	-0.13**	-0.09*	-0.12	-0.09	0.10	0.09	-0.00	0.16***
Economic orthodoxy	0.38***	-0.01	-0.06	-0.06	0.01	-0.10*	0.42***	0.22***	0.30***	0.19***	0.14**	0.09*
Labour groups -	0.08	0.03	-0.01	-0.07	-0.08	-0.00	0.09	0.10	0.12*	0.05	0.08	0.13**
Education -	0.01	0.08	-0.09	-0.07	-0.00	-0.01	-0.07	0.08	0.05	0.03	0.04	0.11**
Protectionism -	0.15*	0.04	-0.02	0.00	0.00	-0.03	-0.10	0.08	0.05	0.06	0.08	0.20***
Welfare -	0.19**	-0.05	-0.11*	-0.04	-0.02	-0.05	0.39***	0.06	0.08	0.09	0.15***	0.15***
Free market economy	0.30***	0.05	-0.04	0.01	-0.09*	-0.07	0.25***	0.25***	0.32***	0.28***	0.20***	0.30***
Pro-redistribution emphases												
Equality +	-0.02	-0.12*	-0.04	0.07	0.12**	0.21***	-0.05	-0.16**	-0.29***	-0.18***	-0.17***	-0.05
Keynesian demand management	0.02	-0.10	-0.01	-0.04	-0.03	-0.11**	-0.11	-0.04	0.04	-0.01	-0.10*	-0.07
Labour groups +	-0.16*	0.15**	0.01	0.06	0.09	0.13**	-0.11	-0.19***	-0.25***	-0.13**	-0.12**	-0.06

Protectionism +	-0.17*	-0.09	-0.03	-0.18***	-0.10*	-0.10*	-0.17*	-0.14**	-0.07	-0.10*	-0.11*	-0.18***
Education +	-0.01	-0.06	-0.12**	0.01	0.03	0.11*	0.01	-0.01	-0.01	0.13**	-0.01	0.03
Welfare +	-0.22**	-0.09	-0.10*	-0.15***	0.00	-0.05	-0.22**	-0.06	-0.20***	-0.24***	-0.11*	-0.16***
Economic planning	-0.15*	-0.17***	-0.13**	-0.01	-0.09	-0.06	-0.19**	-0.05	-0.15**	-0.09	-0.08	0.02
Corporatism/mixed economy	-0.00	0.09	-0.11*	-0.05	-0.11*	-0.08	-0.03	0.06	-0.02	0.01	-0.11*	-0.08
Controlled economy	-0.15*	-0.14**	-0.00	-0.07	0.04	0.00	-0.05	-0.11*	-0.22***	-0.11*	-0.01	-0.18***
Market regulation	-0.11	-0.00	0.01	0.01	0.08	0.01	-0.19**	-0.12*	-0.02	-0.02	0.02	-0.09
Nationalisation	-0.05	0.08	0.06	-0.09	-0.03	0.01	-0.03	-0.21***	-0.26***	-0.16***	-0.23***	-0.22***
Other categories												
Agriculture +	-0.06	-0.14**	-0.15***	-0.11**	-0.20***	-0.17***	-0.02	-0.03	0.01	-0.06	0.01	-0.04
Military +	0.11	0.02	-0.00	-0.10*	-0.10*	-0.18***	0.06	0.21***	0.20***	0.15**	0.12**	0.05
Europe -	0.00	0.05	0.05	0.01	-0.07	-0.27***	0.02	-0.04	-0.11*	-0.15**	-0.18***	-0.16***
Political corruption	0.02	0.04	0.05	-0.05	-0.05	0.18***	-0.01	0.06	-0.03	0.02	-0.02	0.08
Europe +	0.04	-0.01	-0.05	0.04	0.11**	0.19***	0.03	0.06	0.10*	0.17***	0.22***	0.23***
Military -	-0.11	0.13**	0.21***	0.26***	0.25***	0.20***	-0.05	-0.17***	-0.15**	-0.10*	-0.13**	-0.17***
Internationalism -	-0.07	0.00	-0.13**	-0.11**	-0.17***	-0.13**	-0.07	-0.03	-0.02	-0.18***	-0.17***	-0.10*
Internationalism +	0.04	0.10	0.13**	0.09	0.28***	0.12**	-0.12	0.08	-0.04	-0.03	-0.01	0.03
Centralisation	-0.09	-0.07	-0.04	-0.03	0.05	0.10*	-0.14	-0.12*	0.05	-0.01	0.05	0.09
Constitution +	0.08	-0.05	0.05	-0.02	-0.00	-0.00	0.16*	0.03	0.02	0.02	0.05	0.09
Gov-admin efficiency	0.18**	-0.03	-0.02	-0.08	-0.01	-0.04	0.17*	0.10	0.22***	0.16***	0.19***	0.16***
Constitution -	-0.04	0.10	0.02	-0.06	-0.06	-0.04	-0.08	-0.04	0.10*	-0.04	-0.04	0.00
Decentralisation	-0.06	-0.11*	-0.06	-0.04	-0.08	-0.06	-0.17*	-0.06	-0.04	0.03	0.02	-0.00
Middle class and prof. groups	-0.03	0.01	0.10*	-0.08	0.03	-0.06	-0.02	0.03	0.10*	0.07	0.04	0.06
Technology & infrastructure	-0.04	-0.10	-0.12**	-0.12**	-0.14**	-0.08	-0.09	-0.11*	-0.01	-0.04	0.05	0.04
Foreign special +	0.13	-0.03	-0.05	-0.03	0.12**	-0.08	0.08	0.02	-0.03	0.06	0.12**	0.00
Economic goals	0.06	-0.09	-0.06	0.03	-0.09	-0.09	0.05	-0.06	0.01	-0.02	0.09	0.00
Foreign special -	-0.13	-0.08	0.11*	-0.03	0.20***	0.03	0.08	-0.06	-0.06	0.02	0.09	0.05
Anti-imperalism	0.04	-0.03	0.05	0.05	0.05	0.06	-0.03	-0.09	-0.03	-0.17***	-0.08	-0.14**
Marxist analysis +	.	-0.07	0.01	0.13**	-0.03	0.07	.	-0.07	-0.12**	-0.13**	-0.04	-0.12**
Peace	-0.22**	-0.12*	0.14***	0.13**	0.17***	0.07	-0.13	-0.18***	-0.21***	-0.01	0.02	-0.05

Note: The table reports the correlation coefficient between all items available in the Comparative Manifesto Project database and (1) the education gradient (defined as the share of top 10% educated voters within the electorate of the corresponding party) and (2) the income gradient (defined as the share of top 10% income voters within the electorate of the corresponding party). The unit of observation is the political party. Manifesto items correspond to the share of "quasi-sentences" dedicated to emphasizing each category of issues in parties' manifestos. * p<0.10, ** p<0.05, *** p<0.01.

Figure B1 - Share of votes covered by the survey-manifesto dataset

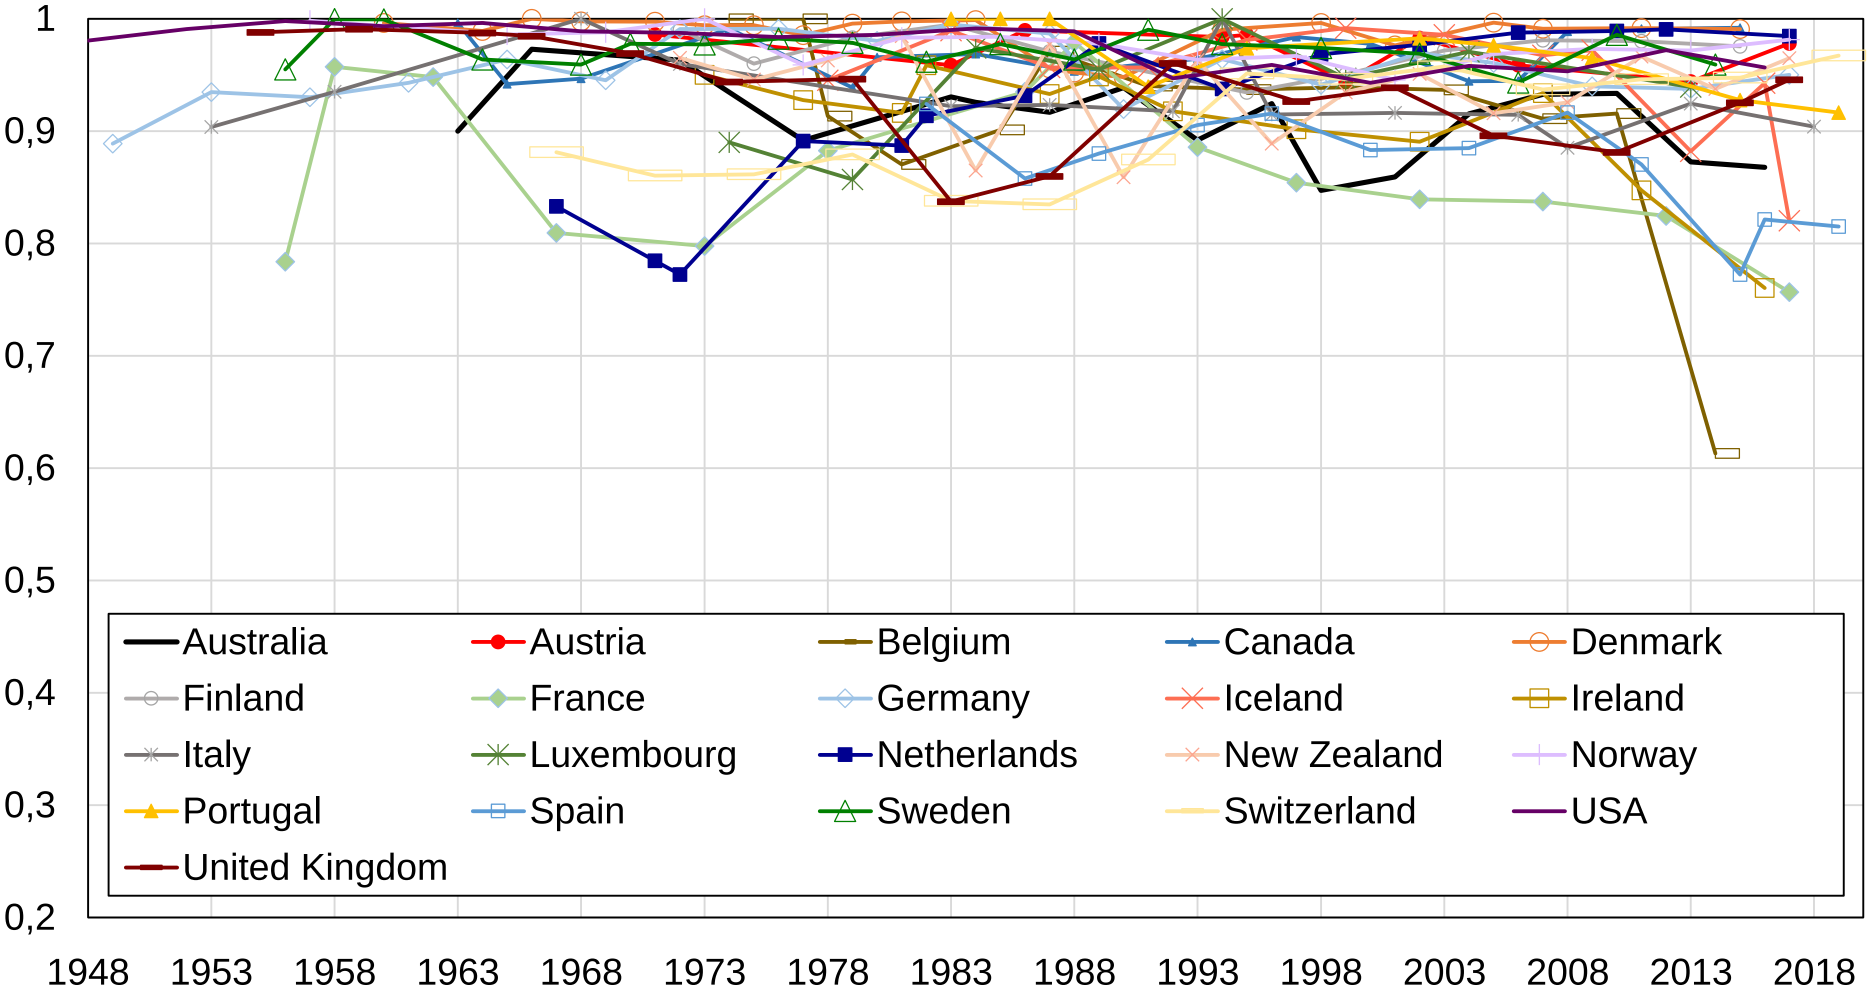

Source: authors' computations combining the World Political Cleavages and Inequality Database and Manifesto Project data.
Note: the figure represents the total share of votes captured by the merged survey-manifesto dataset by country for all elections available between 1945 and 2020.

Figure B2 - Ideological polarization in Western democracies, 1950s

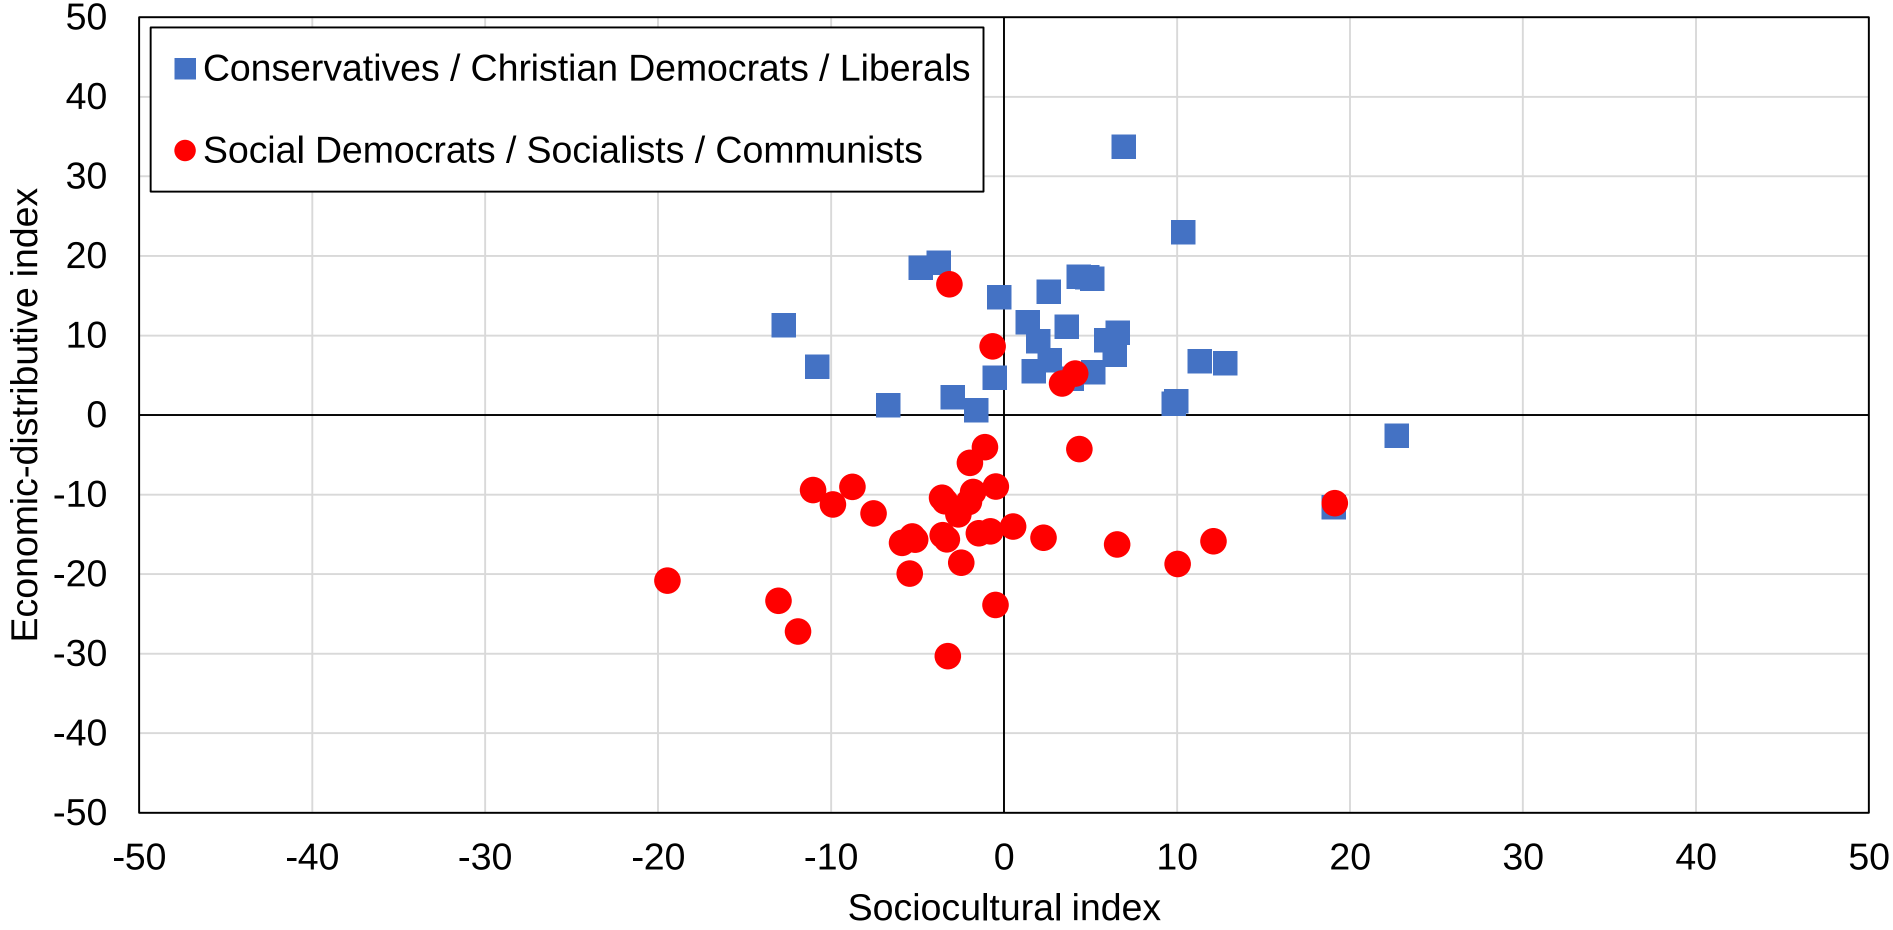

Source: authors' computations using the Comparative Manifesto Project database.
Note: the figure displays the average score of all parties available in the CMP dataset in the 1950s on the economic-distributive index (y-axis) and the sociocultural index (x-axis). Parties are categorized into conservative, Christian democratic, and liberal parties; social democratic, socialist, communist, and other left-wing parties, anti-immigration parties; and green parties.

Figure B3 - Ideological polarization in Western democracies, 1960s

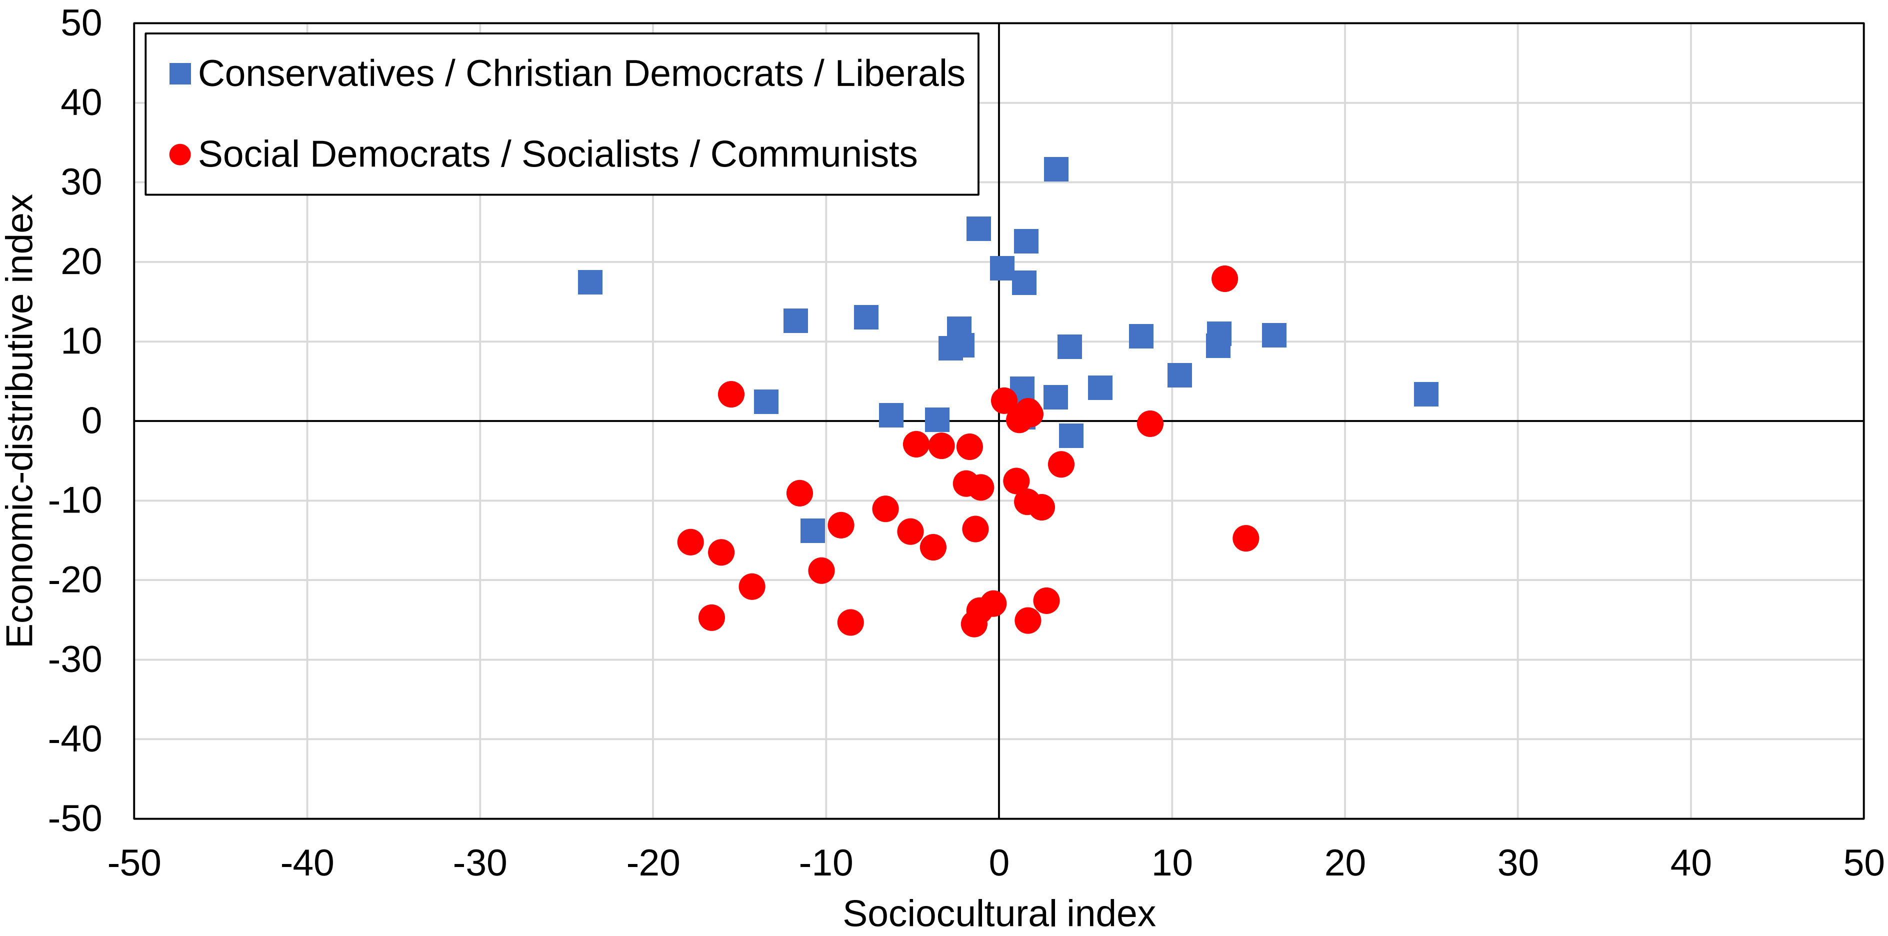

Source: authors' computations using the Comparative Manifesto Project database.
Note: the figure displays the average score of all parties available in the CMP dataset in the 1960s on the economic-distributive index (y-axis) and the sociocultural index (x-axis). Parties are categorized into conservative, Christian democratic, and liberal parties; social democratic, socialist, communist, and other left-wing parties, anti-immigration parties; and green parties.

Figure B4 - Ideological polarization in Western democracies, 1970s

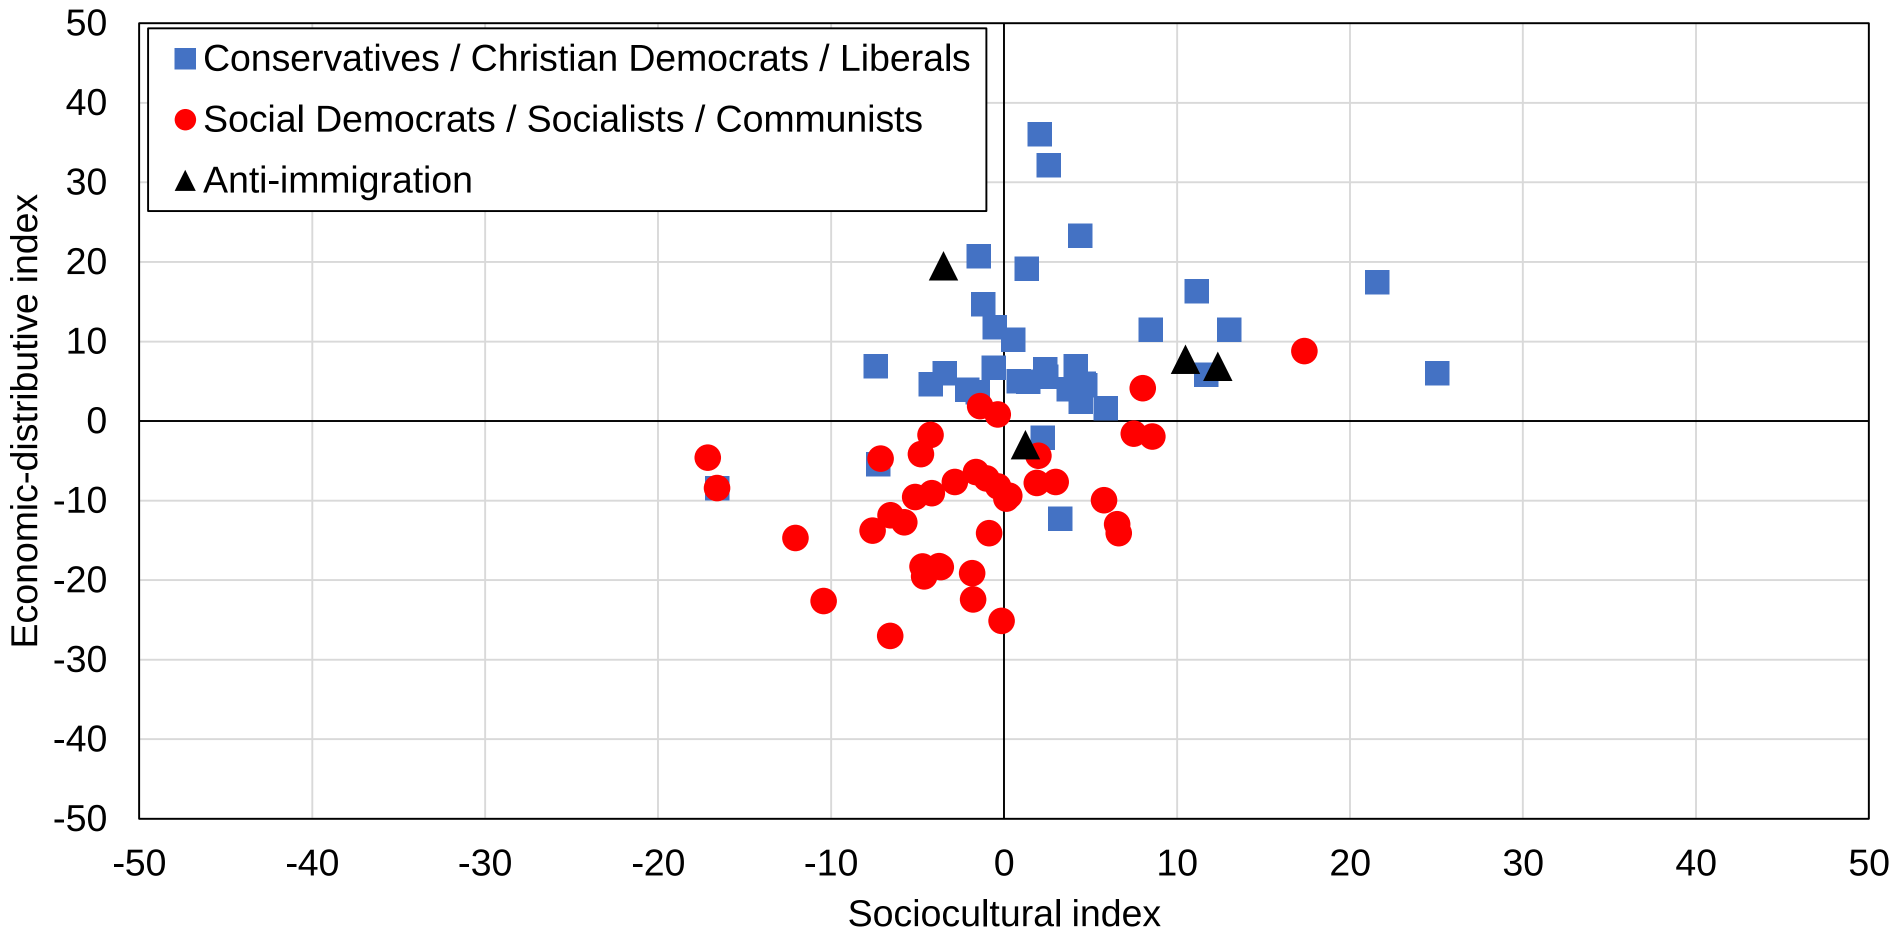

Source: authors' computations using the Comparative Manifesto Project database.
Note: the figure displays the average score of all parties available in the CMP dataset in the 1970s on the economic-distributive index (y-axis) and the sociocultural index (x-axis). Parties are categorized into conservative, Christian democratic, and liberal parties; social democratic, socialist, communist, and other left-wing parties, anti-immigration parties; and green parties.

Figure B5 - Ideological polarization in Western democracies, 1980s

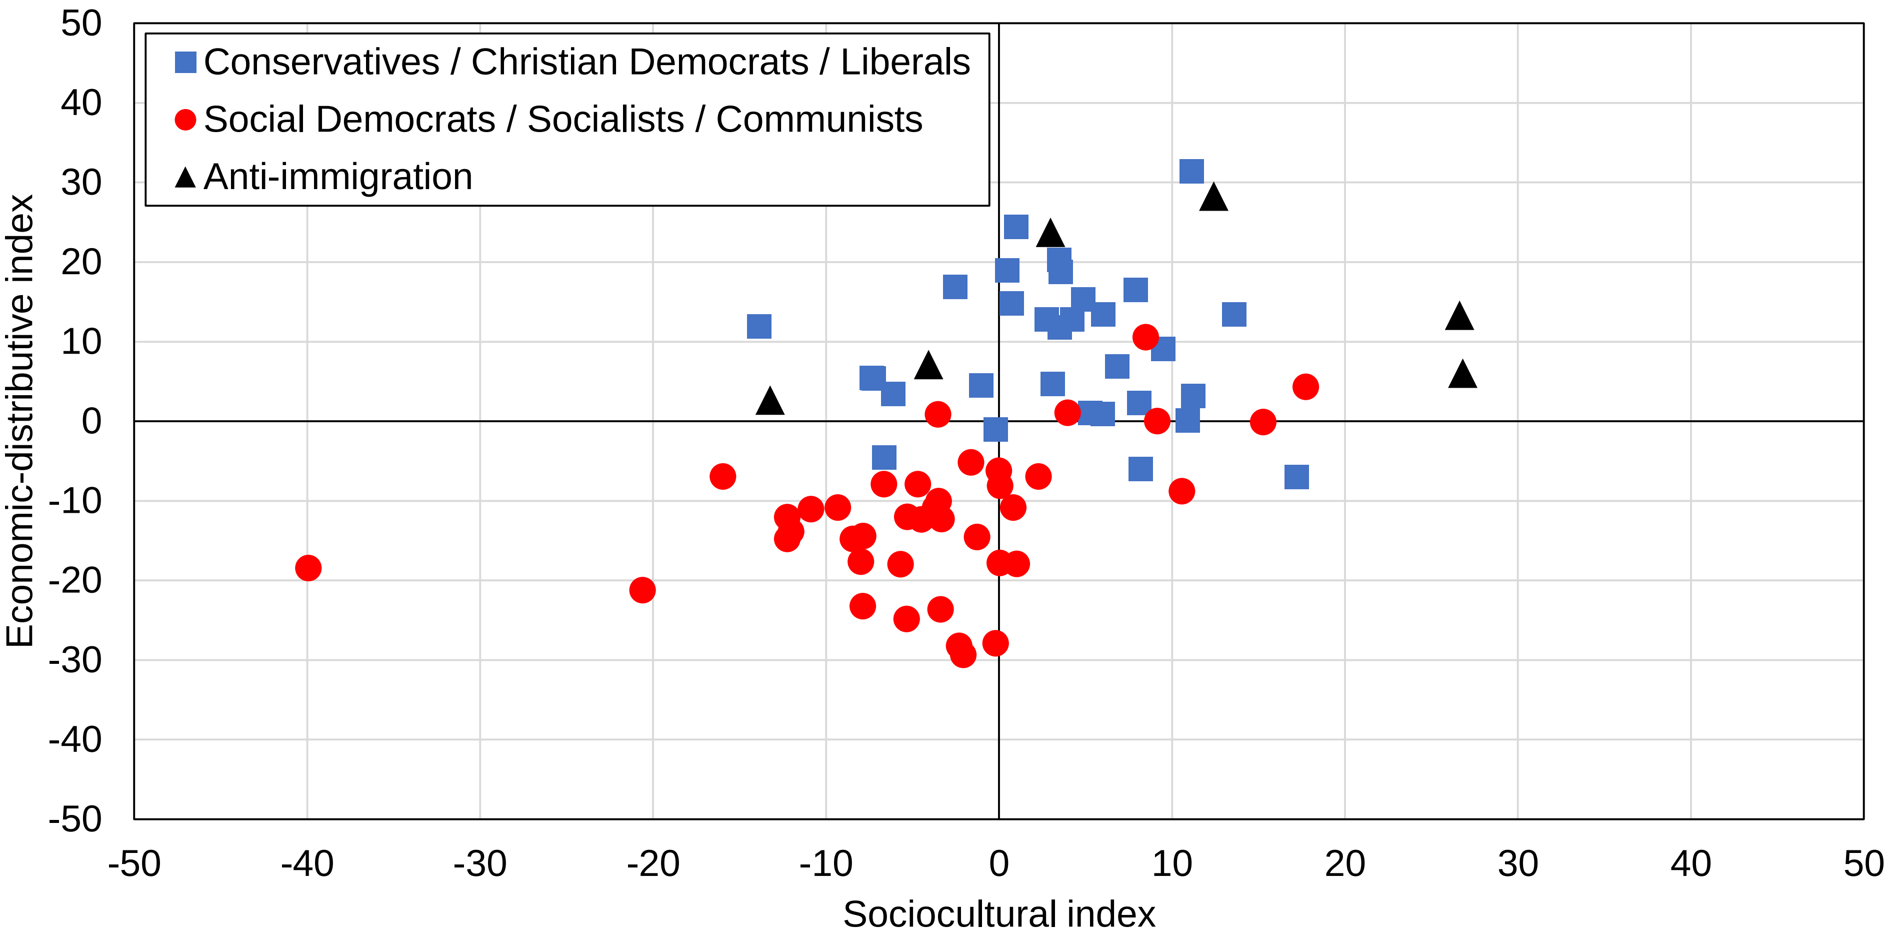

Source: authors' computations using the Comparative Manifesto Project database.
Note: the figure displays the average score of all parties available in the CMP dataset in the 1980s on the economic-distributive index (y-axis) and the sociocultural index (x-axis). Parties are categorized into conservative, Christian democratic, and liberal parties; social democratic, socialist, communist, and other left-wing parties, anti-immigration parties; and green parties.

Figure B6 - Ideological polarization in Western democracies, 1990s

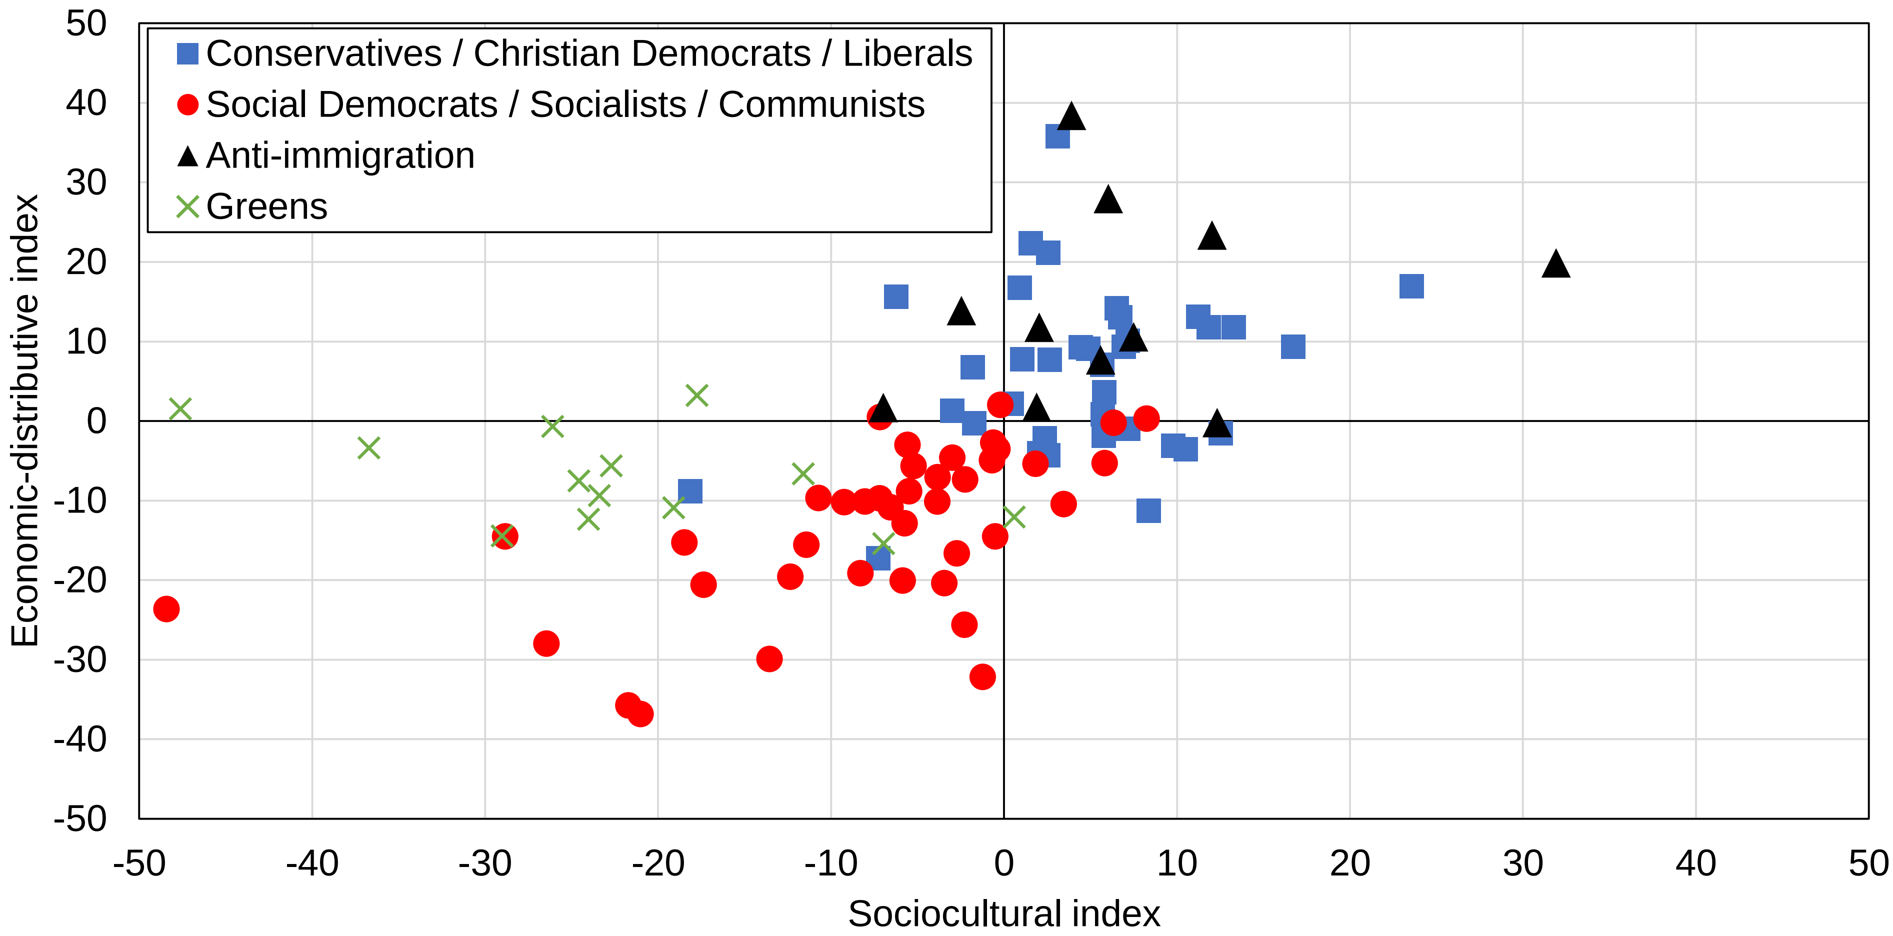

Source: authors' computations using the Comparative Manifesto Project database.
Note: the figure displays the average score of all parties available in the CMP dataset in the 1990s on the economic-distributive index (y-axis) and the sociocultural index (x-axis). Parties are categorized into conservative, Christian democratic, and liberal parties; social democratic, socialist, communist, and other left-wing parties, anti-immigration parties; and green parties.

Figure B7 - Ideological polarization in Western democracies, 2000s

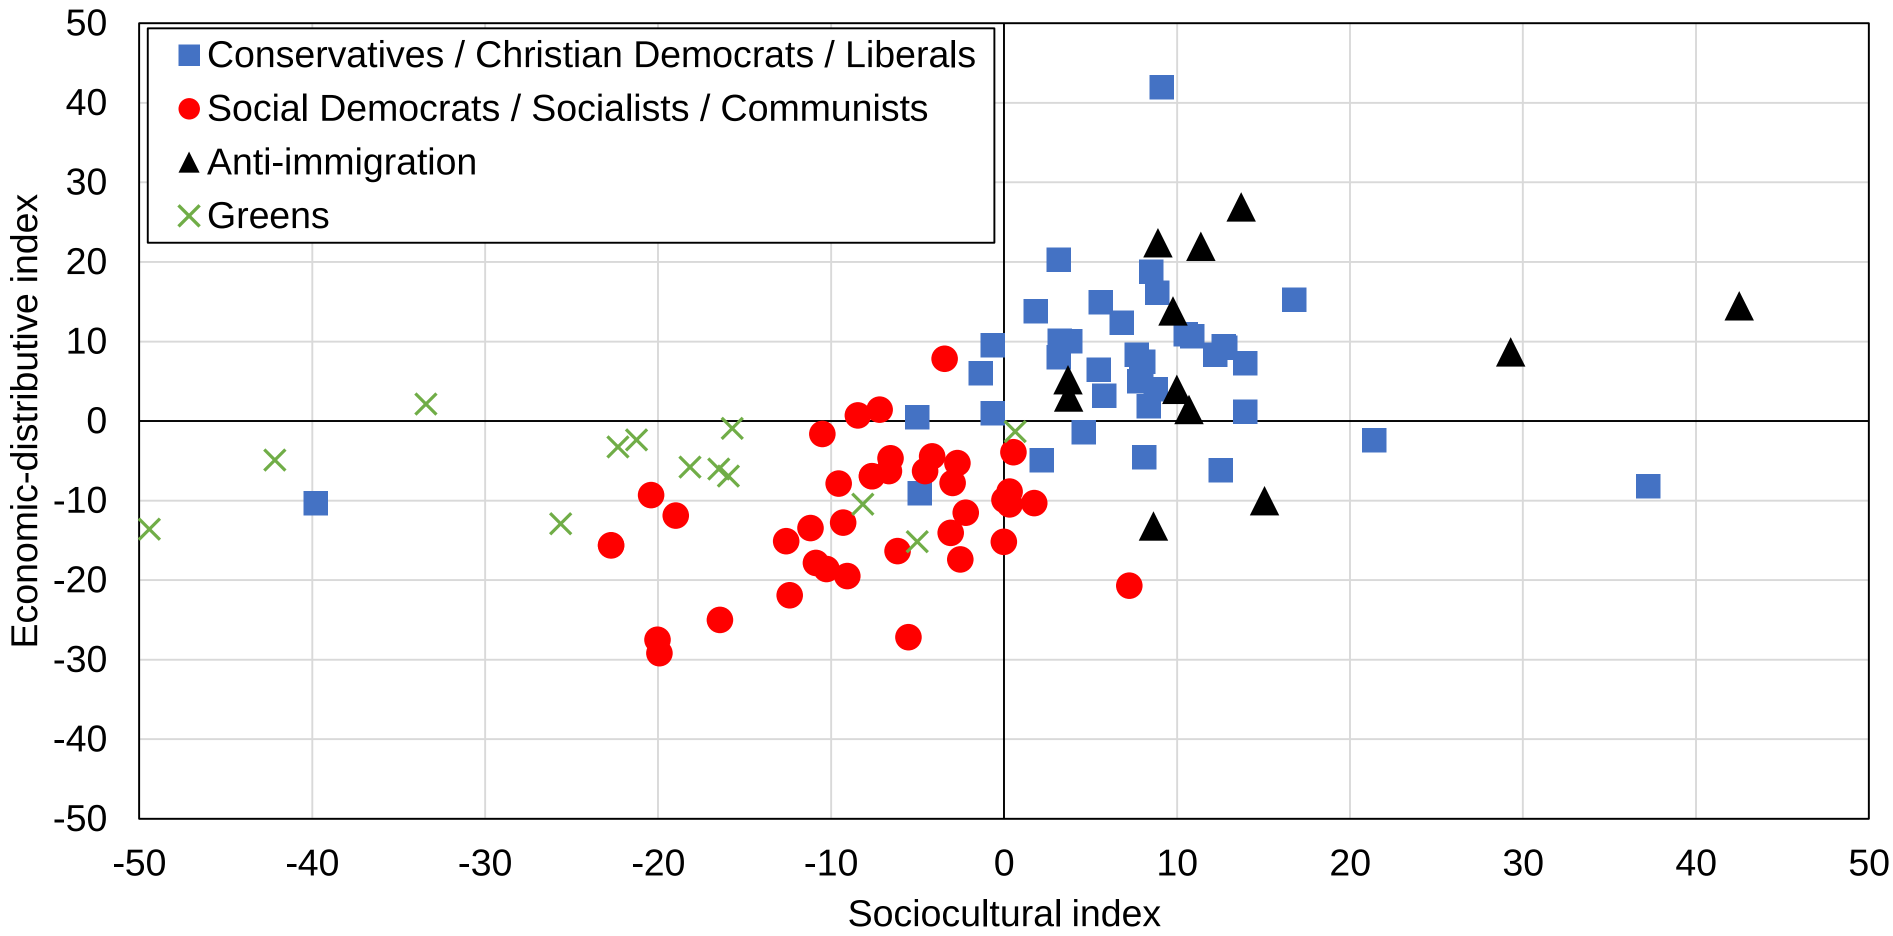

Source: authors' computations using the Comparative Manifesto Project database.
Note: the figure displays the average score of all parties available in the CMP dataset in the 2000s on the economic-distributive index (y-axis) and the sociocultural index (x-axis). Parties are categorized into conservative, Christian democratic, and liberal parties; social democratic, socialist, communist, and other left-wing parties, anti-immigration parties; and green parties.

Figure B8 - Ideological polarization in Western democracies, 2010s

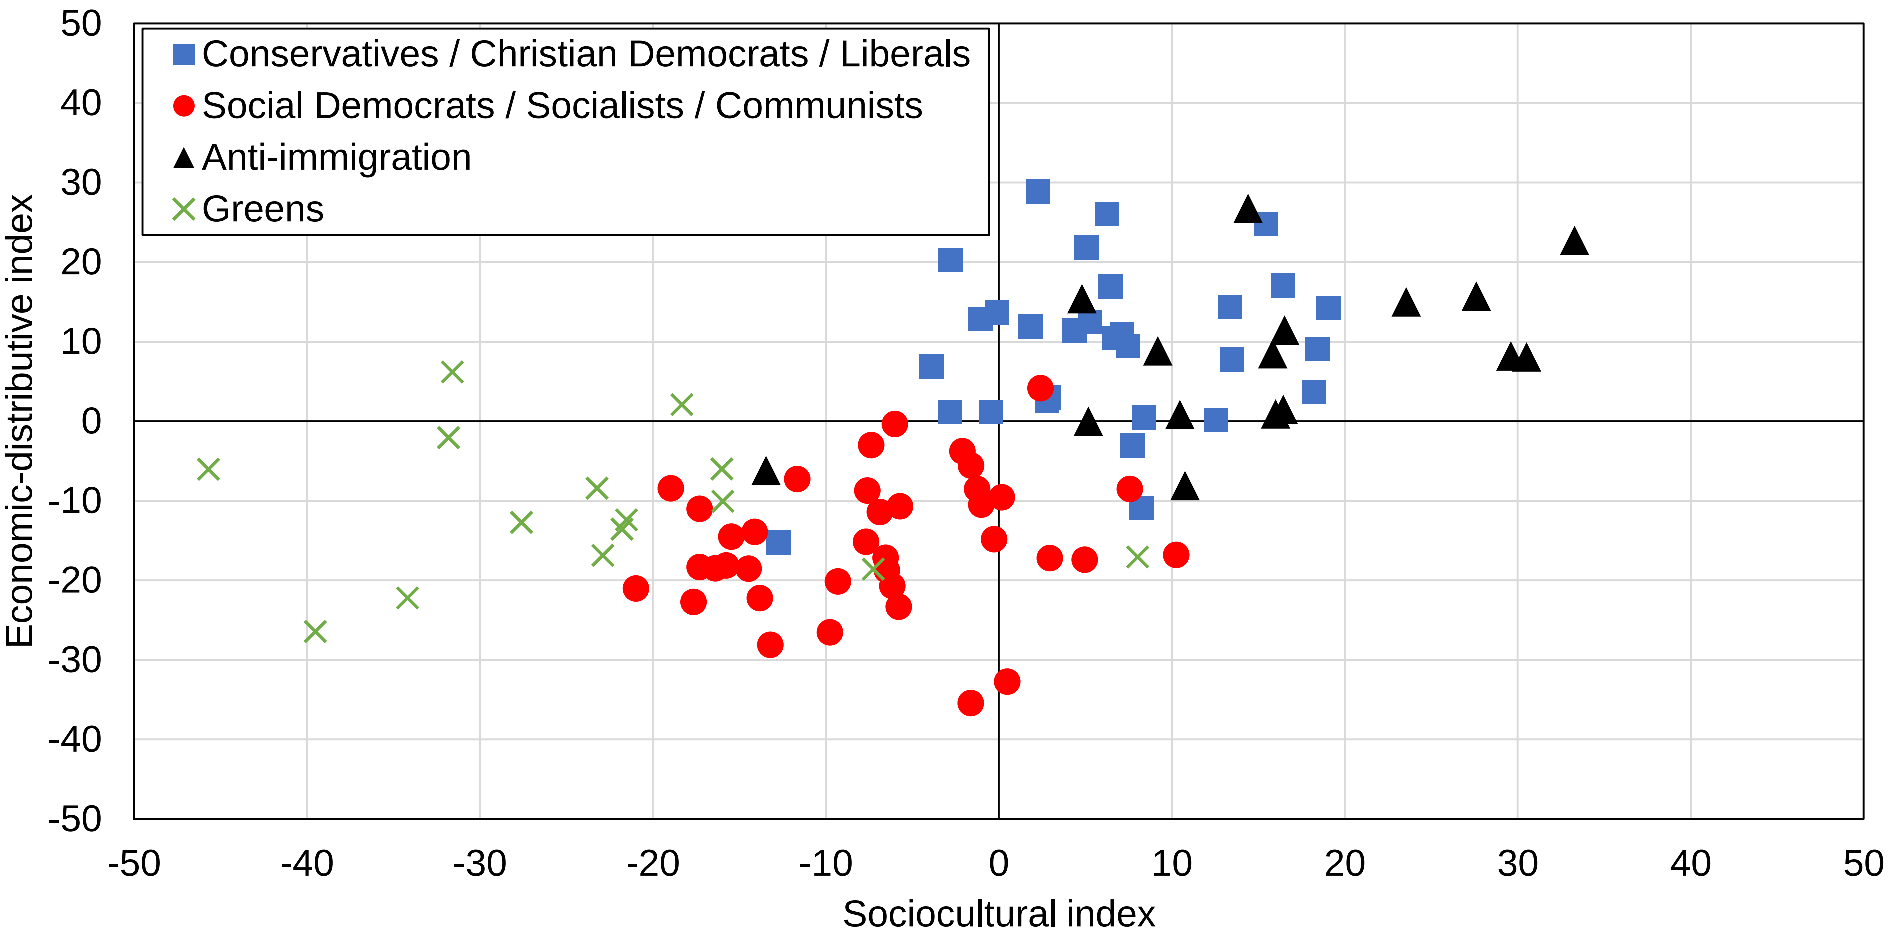

Source: authors' computations using the Comparative Manifesto Project database.
Note: the figure displays the average score of all parties available in the CMP dataset in the 2010s on the economic-distributive index (y-axis) and the sociocultural index (x-axis). Parties are categorized into conservative, Christian democratic, and liberal parties; social democratic, socialist, communist, and other left-wing parties, anti-immigration parties; and green parties.

Figure B9 - Multidimensional political conflict and the divergence of income and education (bottom 50%)

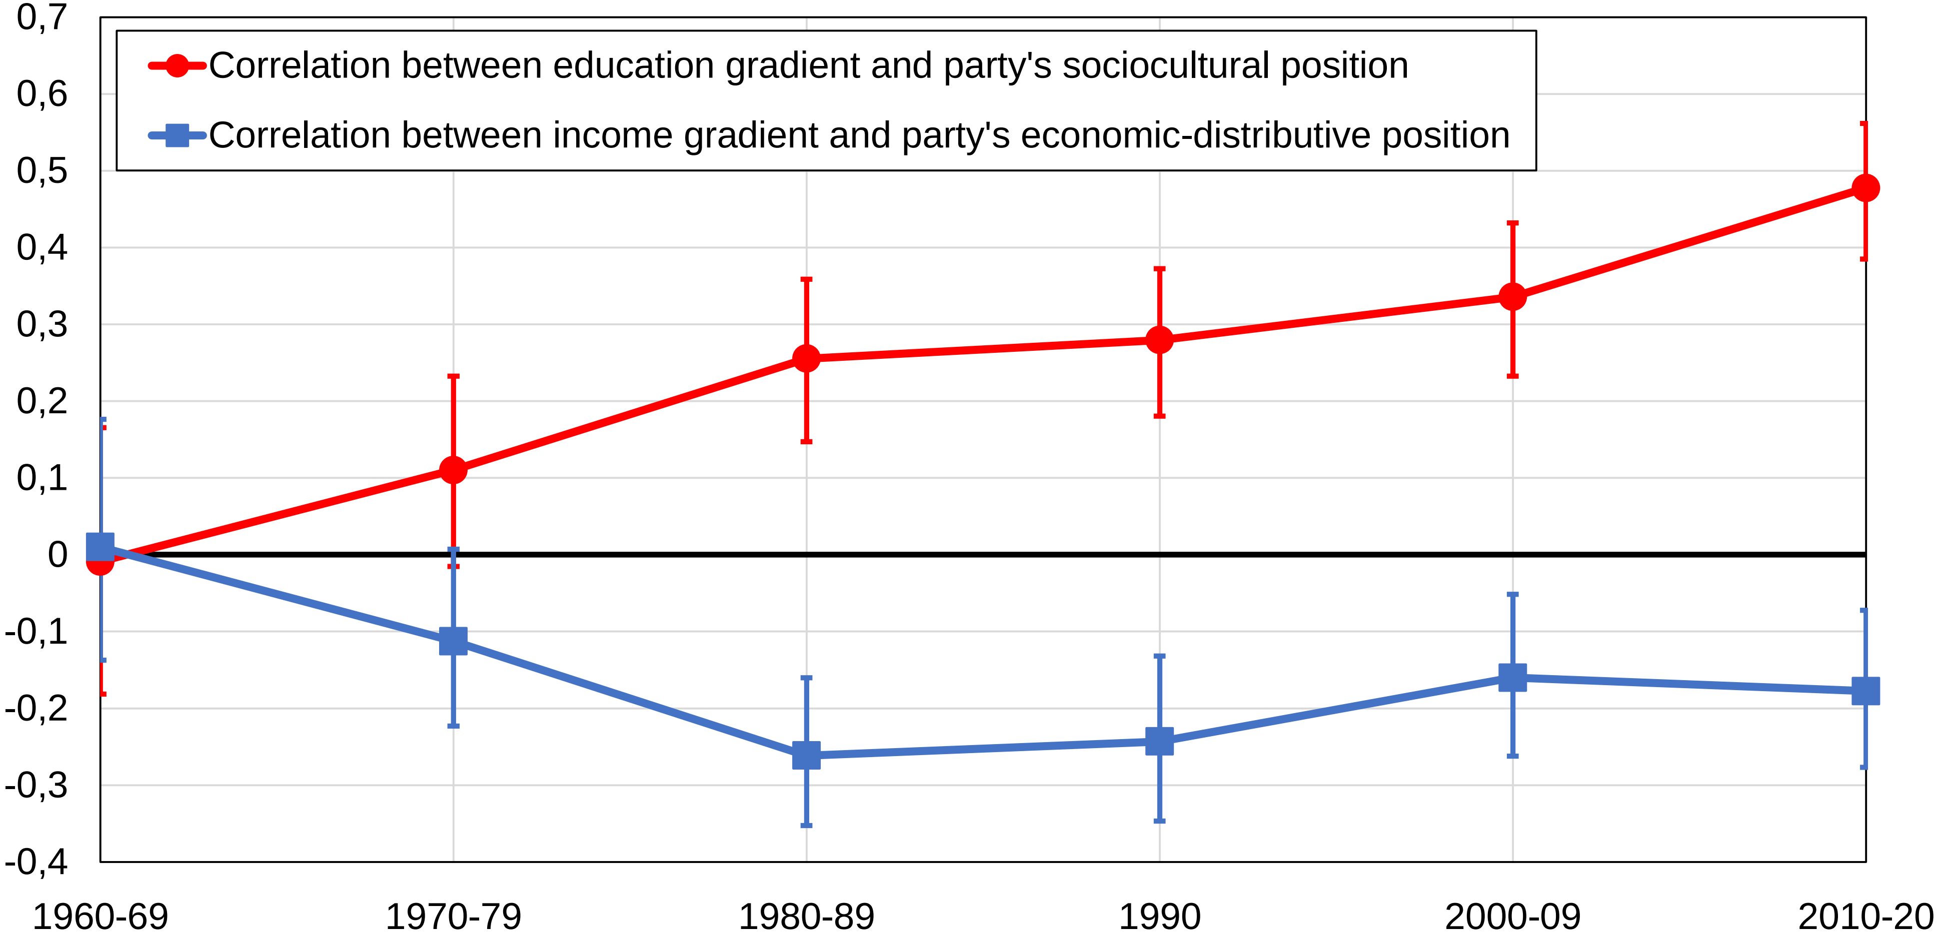

Source: authors' computations combining the World Political Cleavages and Inequality Database with Manifesto Project data.
Note: the upper lines plots the raw correlation between the education gradient (defined as the share of top 50% educated voters within the electorate of a given party) and the sociocultural index. The bottom line plots the raw correlation between the income gradient (defined as the share of top 50% income voters within the electorate of a given party) and the economic-distributive index (inverted, so that higher values correspond to greater pro-redistribution emphases). Error bars represent 95% confidence intervals.

Figure B10 - Sociocultural polarization and educational divides, 1970s

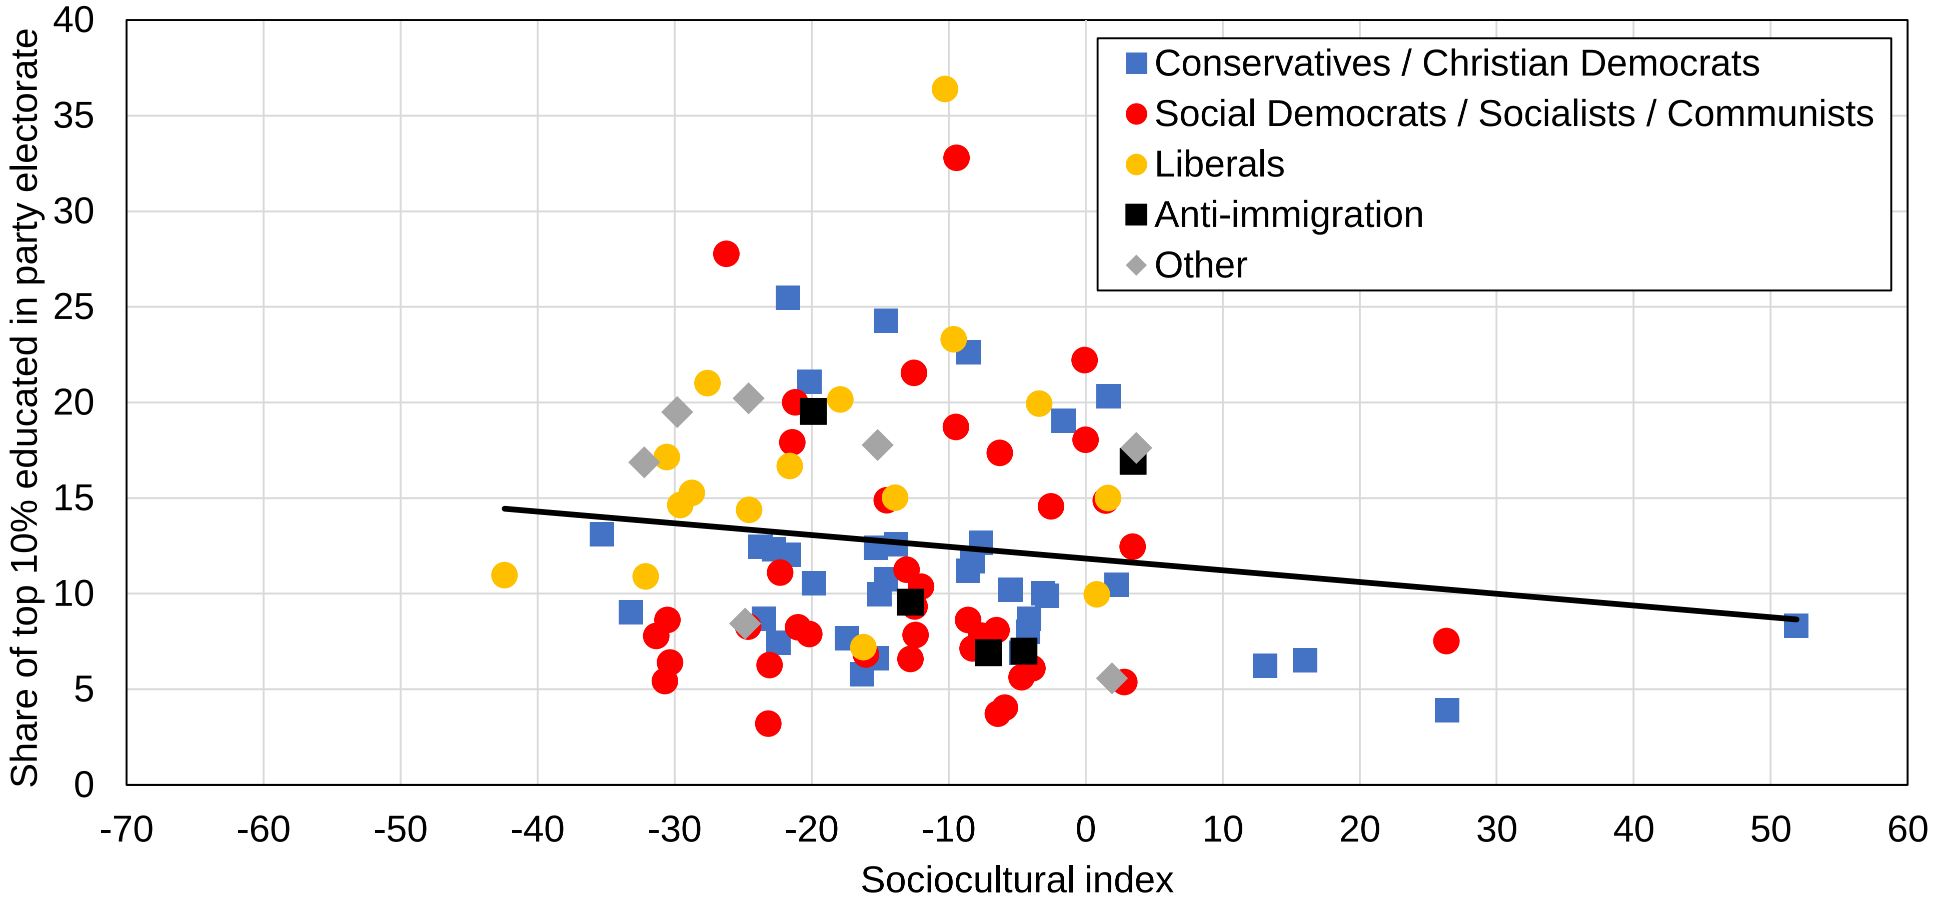

Source: authors' computations combining the World Political Cleavages and Inequality Database and Manifesto Project data.
Note: parties are categorized into conservative and Christian democratic parties; liberal and social-liberal parties; social democratic, socialist, communist and other left-wing parties, anti-immigration parties; green parties; and other unclassifiable parties.

Figure B11 - Sociocultural polarization and educational divides, 1980s

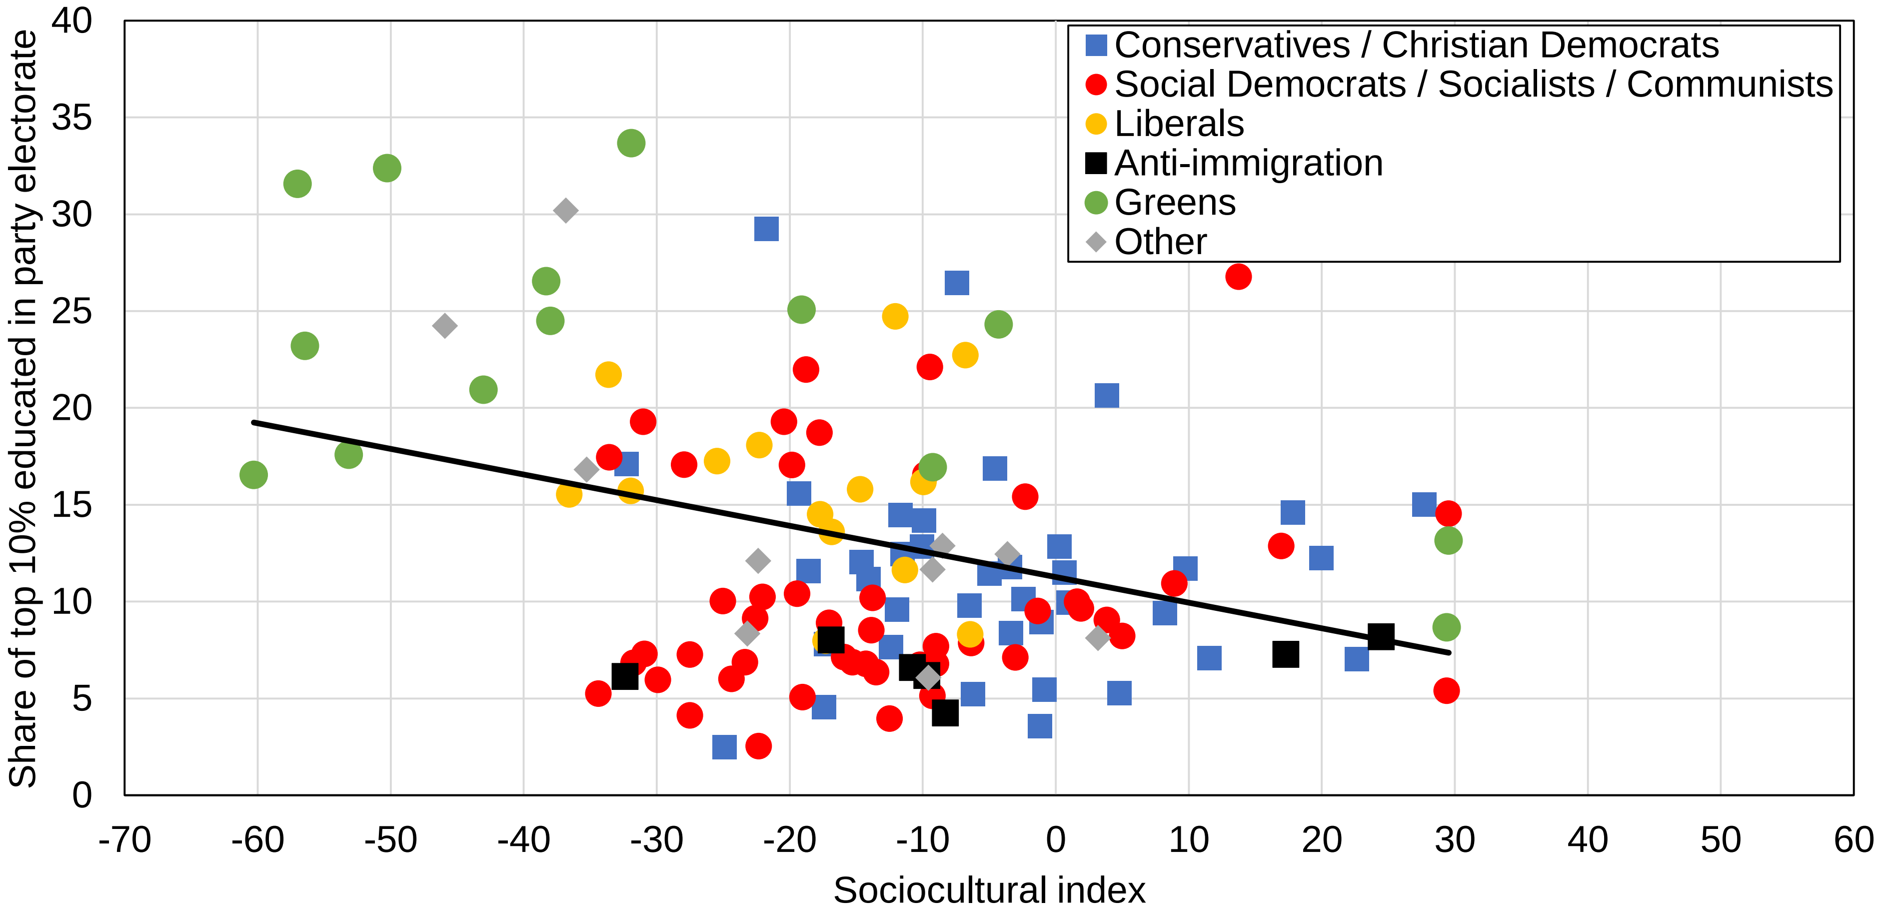

Source: authors' computations combining the World Political Cleavages and Inequality Database and Manifesto Project data.
Note: parties are categorized into conservative and Christian democratic parties; liberal and social-liberal parties; social democratic, socialist, communist and other left-wing parties, anti-immigration parties; green parties; and other unclassifiable parties.

Figure B12 - Sociocultural polarization and educational divides, 1990s

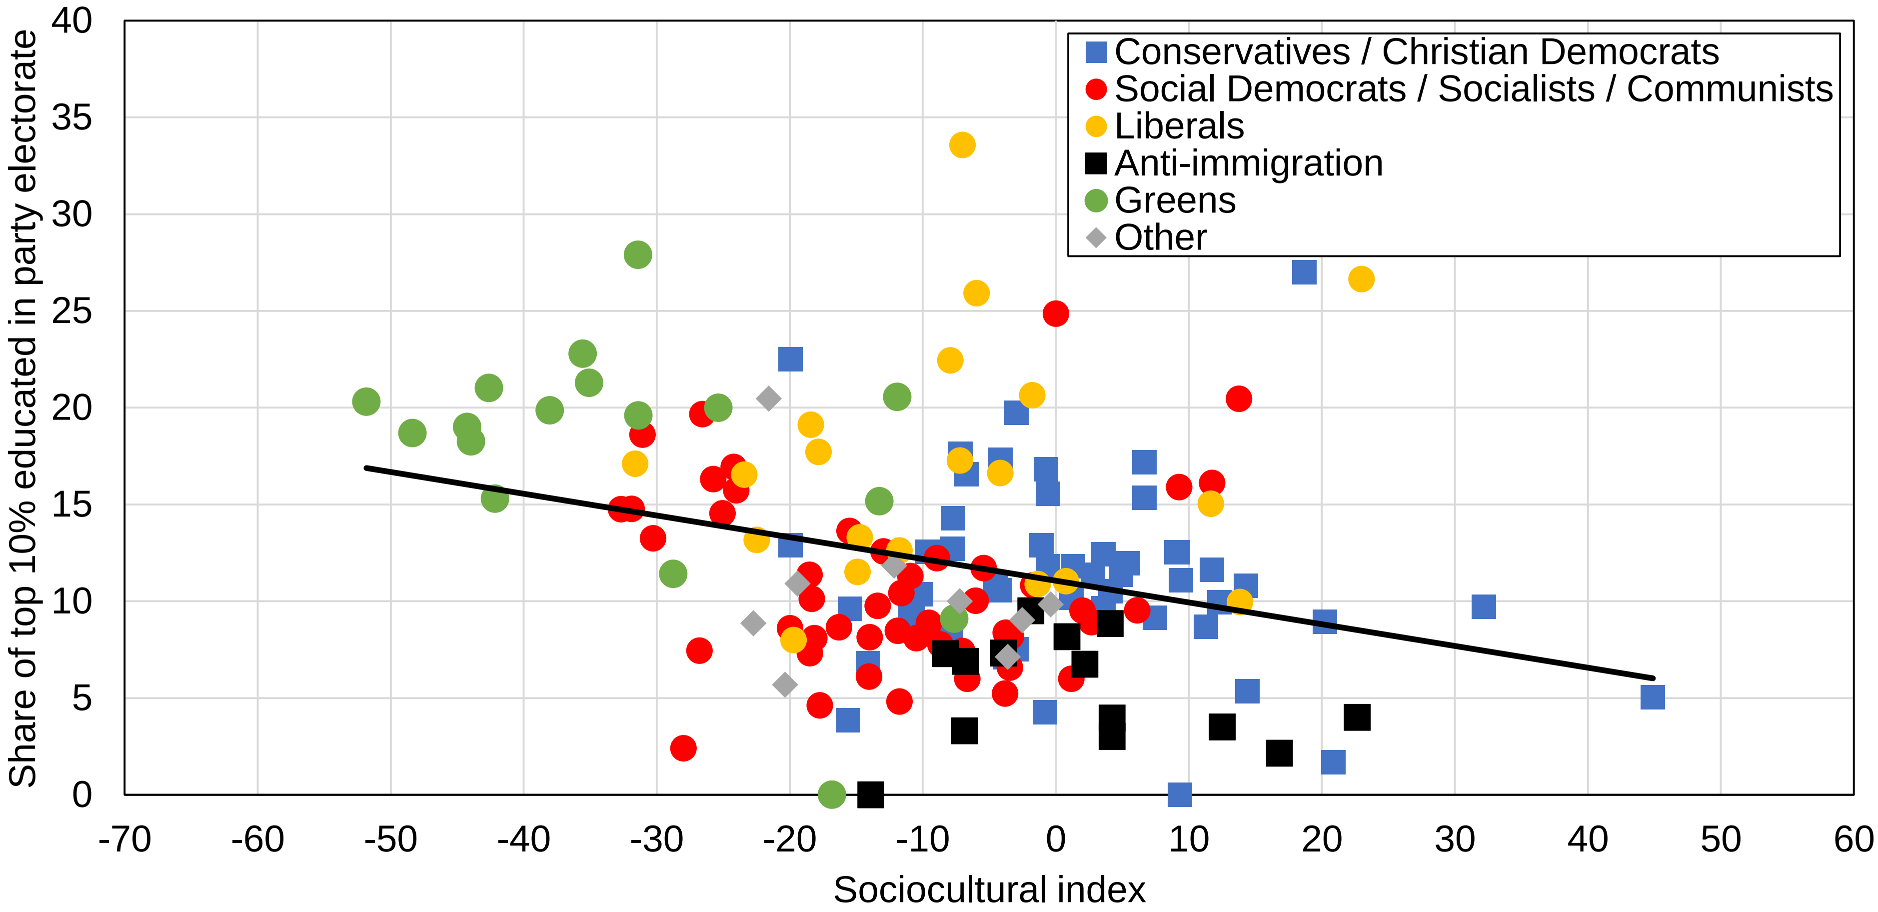

Source: authors' computations combining the World Political Cleavages and Inequality Database and Manifesto Project data.
Note: parties are categorized into conservative and Christian democratic parties; liberal and social-liberal parties; social democratic, socialist, communist and other left-wing parties, anti-immigration parties; green parties; and other unclassifiable parties.

Figure B13 - Sociocultural polarization and educational divides, 2000s

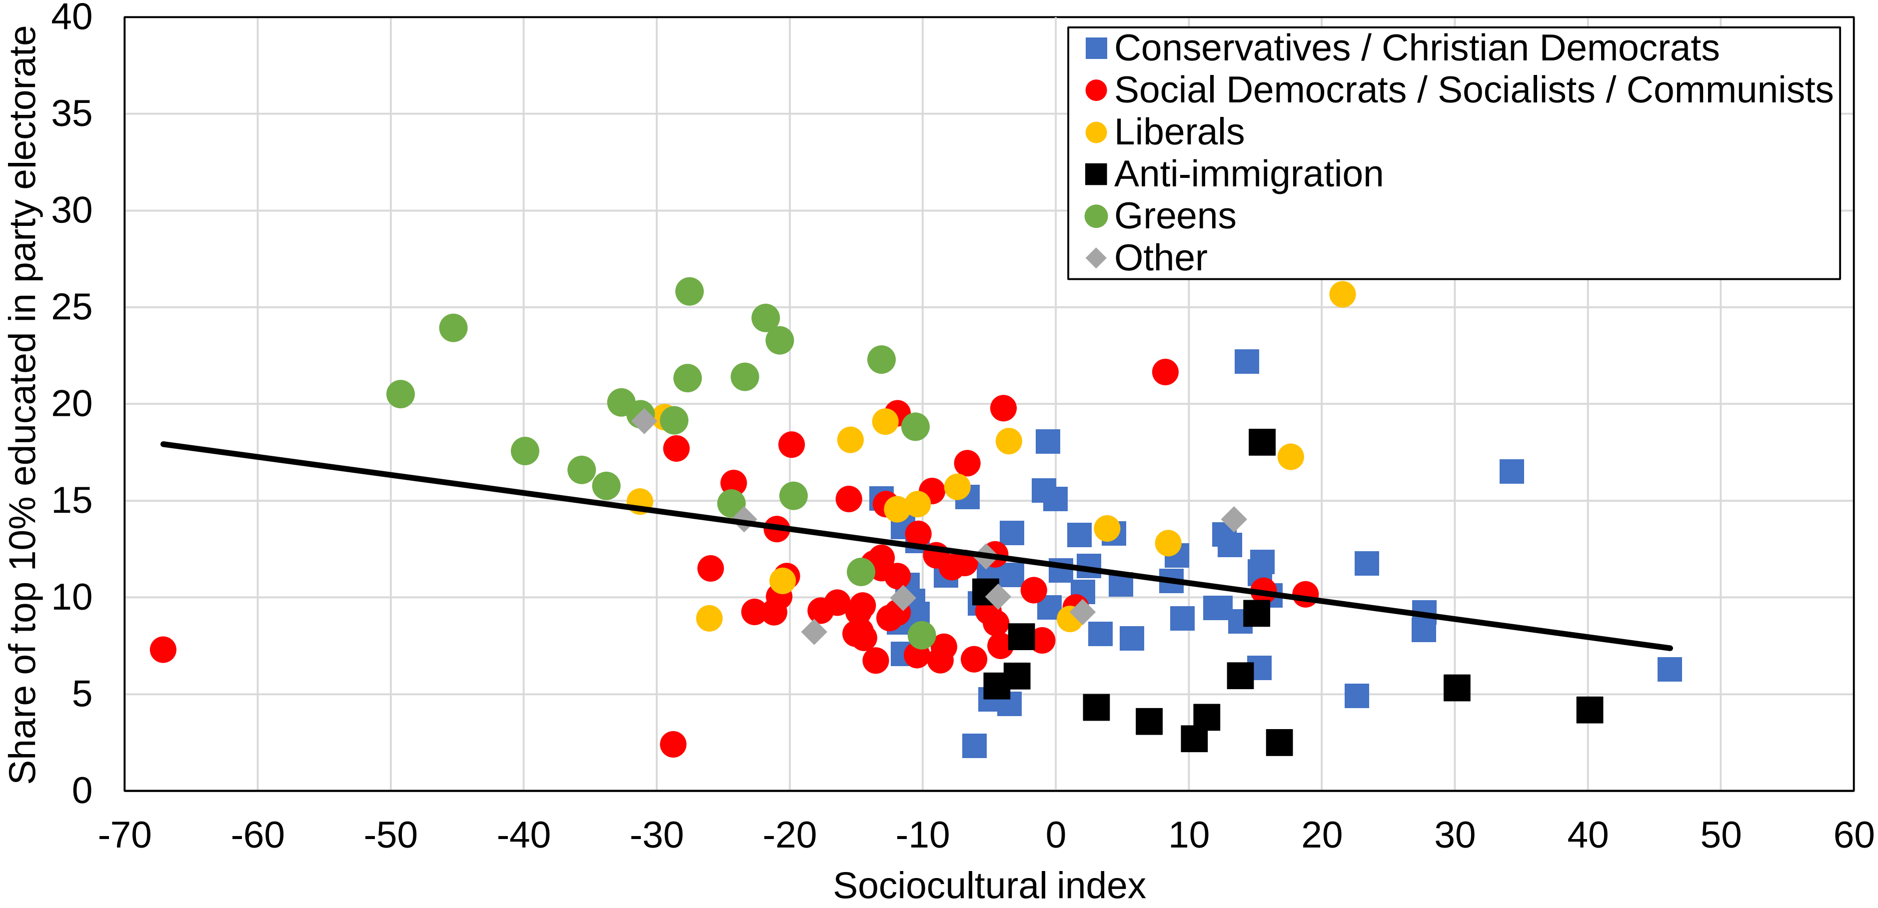

Source: authors' computations combining the World Political Cleavages and Inequality Database and Manifesto Project data.
Note: parties are categorized into conservative and Christian democratic parties; liberal and social-liberal parties; social democratic, socialist, communist and other left-wing parties, anti-immigration parties; green parties; and other unclassifiable parties.

Figure B14 - Sociocultural polarization and educational divides, 2010s

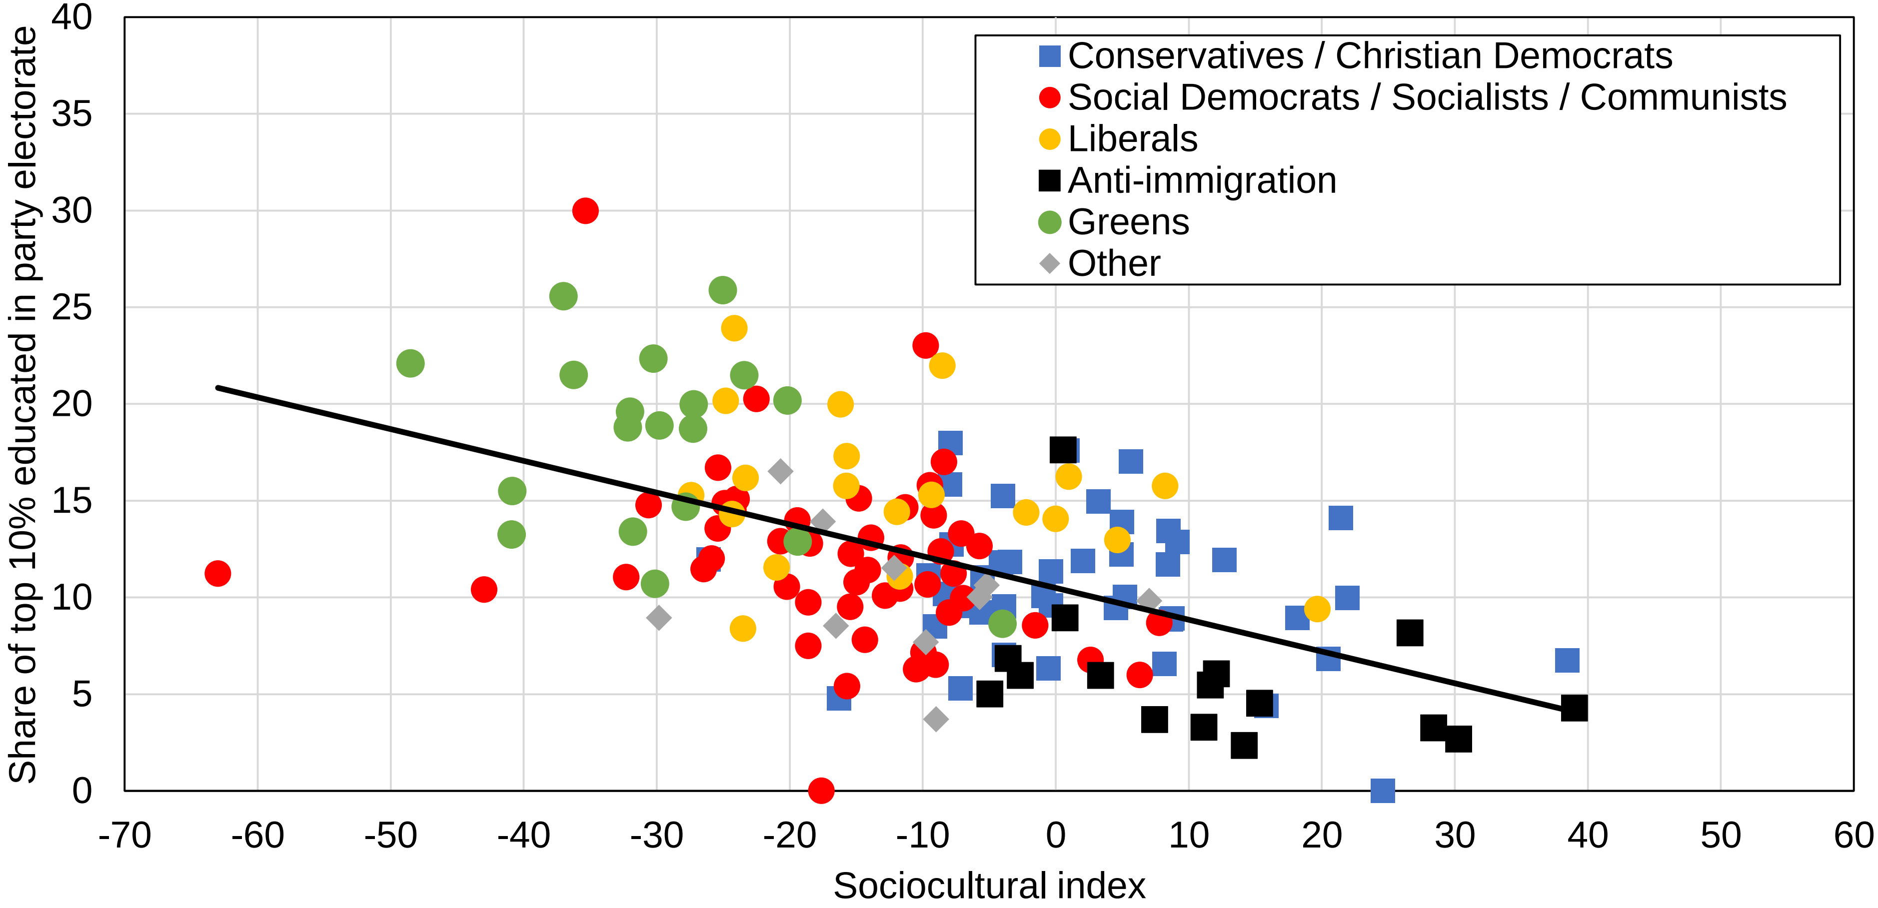

Source: authors' computations combining the World Political Cleavages and Inequality Database and Manifesto Project data.
Note: parties are categorized into conservative and Christian democratic parties; liberal and social-liberal parties; social democratic, socialist, communist and other left-wing parties, anti-immigration parties; green parties; and other unclassifiable parties.

Figure B15 - Multidimensional political conflict and the divergence of income and education (country-level analysis)

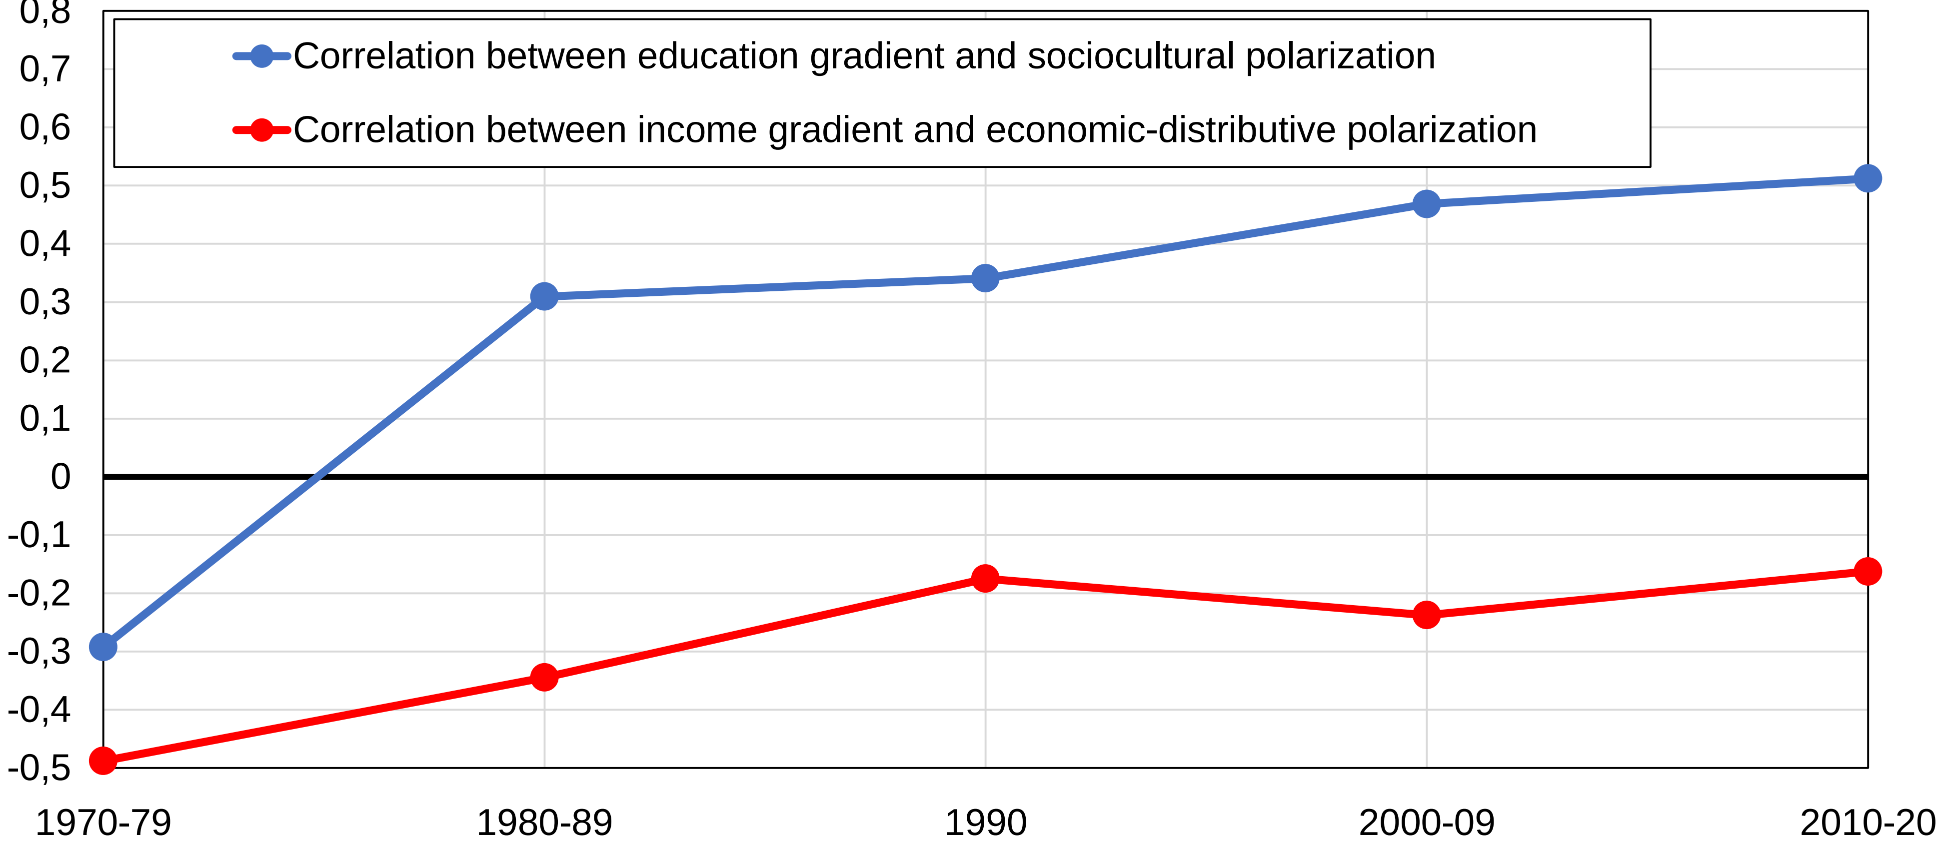

Source: authors' computations combining the World Political Cleavages and Inequality Database and Manifesto Project data.
Note: the upper lines plots the raw correlation between the education gradient (defined as the difference between the share of top 10% educated voters and the share of bottom 90% educated voters voting for left-wing parties) and sociocultural polarization (defined as the standard deviation of the sociocultural index across all parties in a given country). Conversely, the bottom line plots the raw correlation between the income gradient and economic-distributive polarization (inverted, so that higher values correspond to greater pro-redistribution emphases). Both polarization indices are normalized to the average standard deviation to highlight relative evolutions. The unit of observation is the country.

Figure B16 - Average left-right positions of political parties in Western democracies, 2000-2020: survey data vs. manifesto data

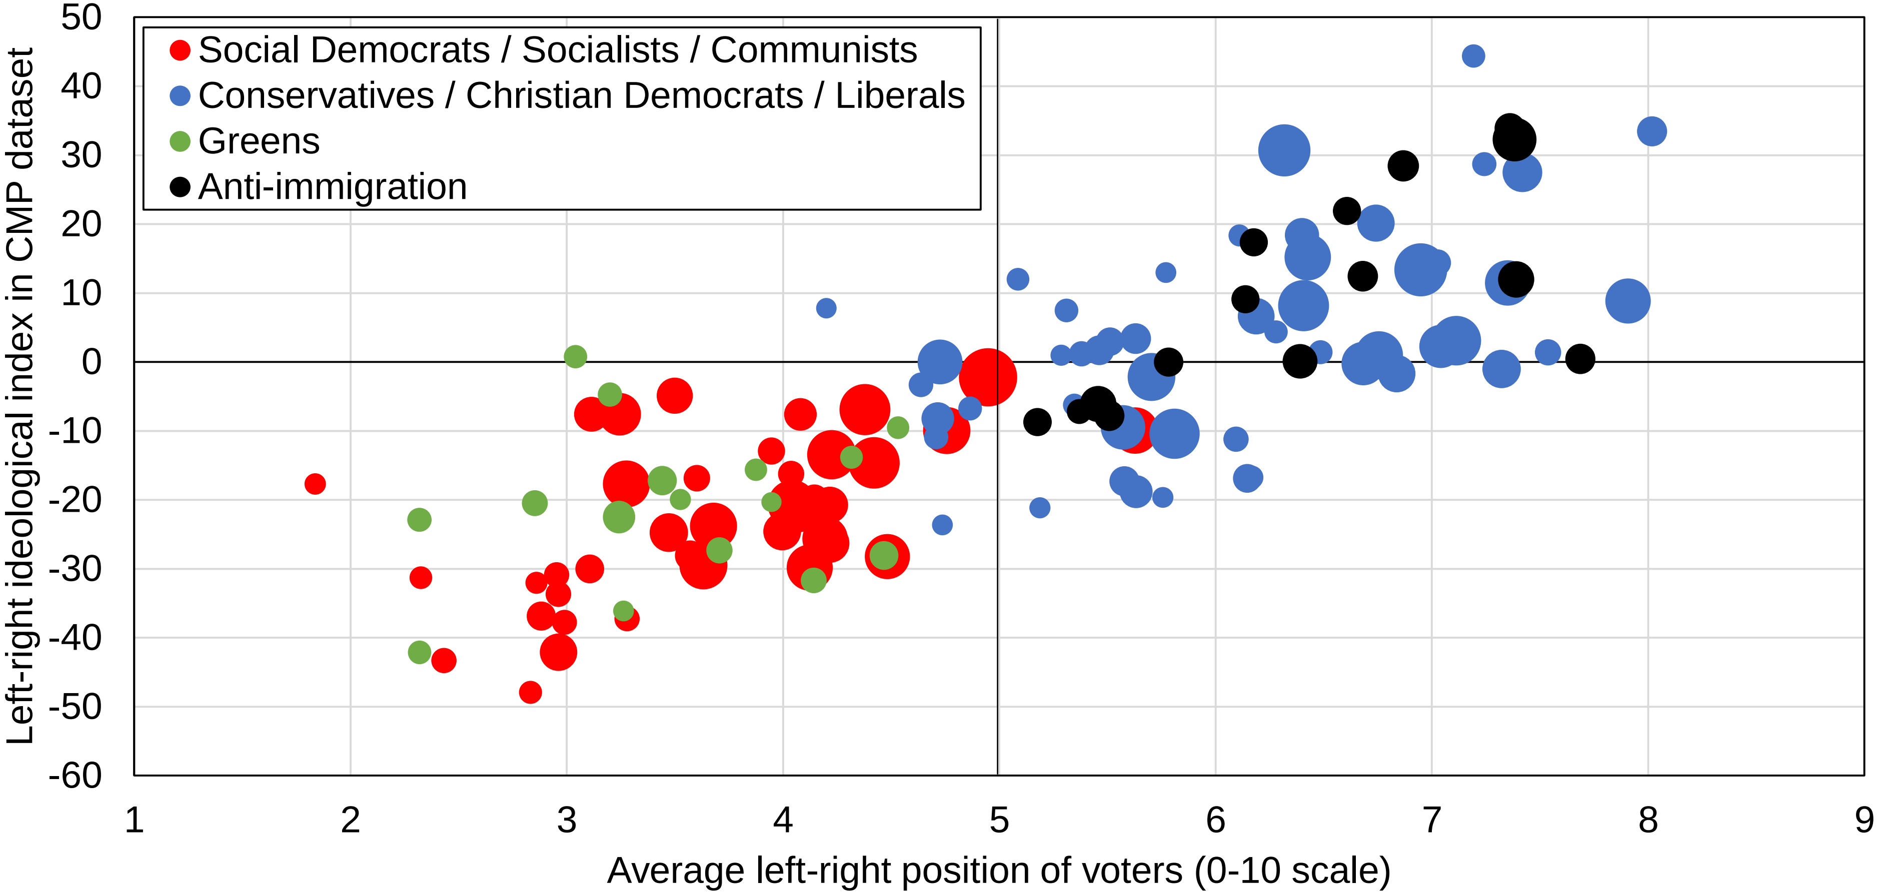

Source: authors' computations combining the World Political Cleavages and Inequality Database and the CMP database.
Note: the figure displays the average score of parties on the left-right ideological index in the Comparative Manifesto Project database (y-axis) and the average self-reported left-right placement of voters supporting these parties, as reported in survey data (x-axis). Average over the 2000-2020 period. Excludes parties that received less than 5% of the vote in a given election. Parties are categorized into conservative, Christian democratic, and liberal parties; social democratic, socialist, communist and other left-wing parties, anti-immigration parties; and green parties. The size of bubbles is proportional to the square root of the average vote share of each party.

Figure B17 - Average left-right positions of political parties in Western democracies, 2000-2020: survey data vs. manifesto data (normalized)

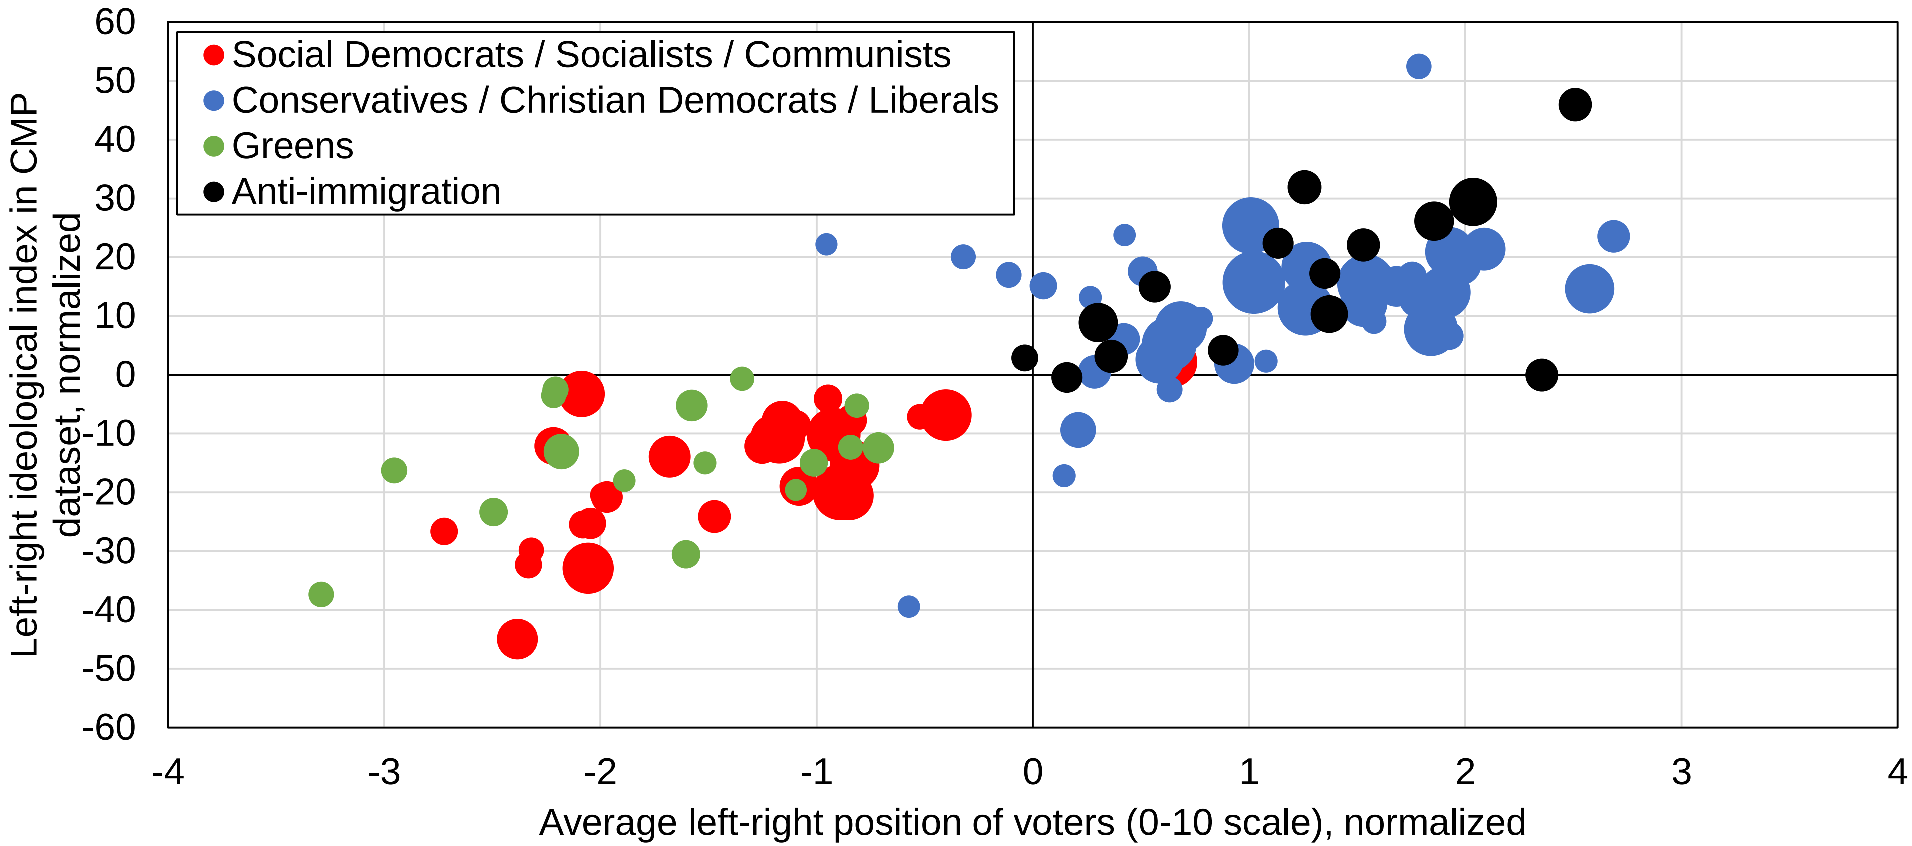

Source: authors' computations combining the World Political Cleavages and Inequality Database and the CMP database.
Note: the figure displays the average score of parties on the left-right ideological index in the Comparative Manifesto Project database (y-axis) and the average self-reported left-right placement of voters supporting these parties, as reported in survey data (x-axis). Both variables are normalized by taking the difference between the party's value and the vote-share-weighted average value in a given country-year. Average over the 2000-2020 period. Excludes parties that received less than 5% of the vote in a given election. Parties are categorized into conservative, Christian democratic, and liberal parties; social democratic, socialist, communist and other left-wing parties, anti-immigration parties; and green parties. The size of bubbles is proportional to the square root of the average vote share of each party.

Figure B18 - Average self-declared left-right position of voters supporting Conservative / Christian Democratic / Liberal parties

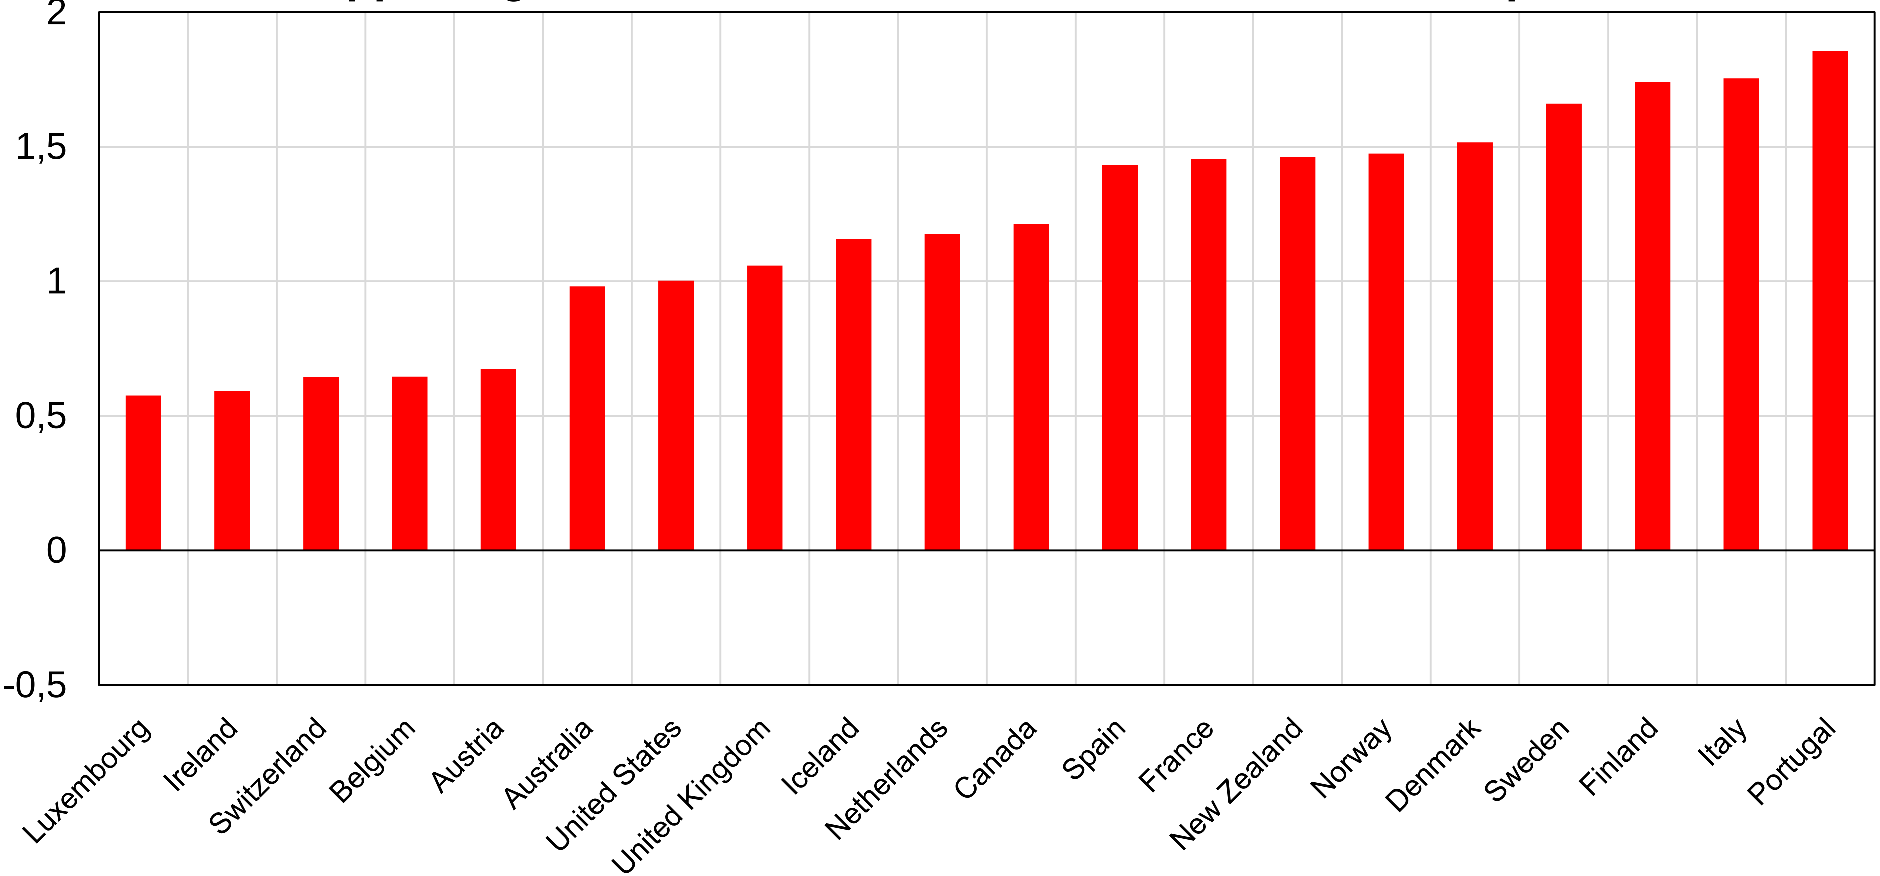

Source: authors' computations using the World Political Cleavages and Inequality Database.

Note: the figure represents the difference between the average self-declared left-right position of voters supporting Conservative, Christian Democratic, and Liberal parties and the average self-declared left-right position of all voters over the 2000-2020 period by country.

Interpretation: In all countries, voters supporting Conservative / Christian Democratic / Liberal parties are significantly more likely to declare being more right-wing than other voters.

Figure B19 - Average self-declared left-right position of voters supporting anti-immigration parties

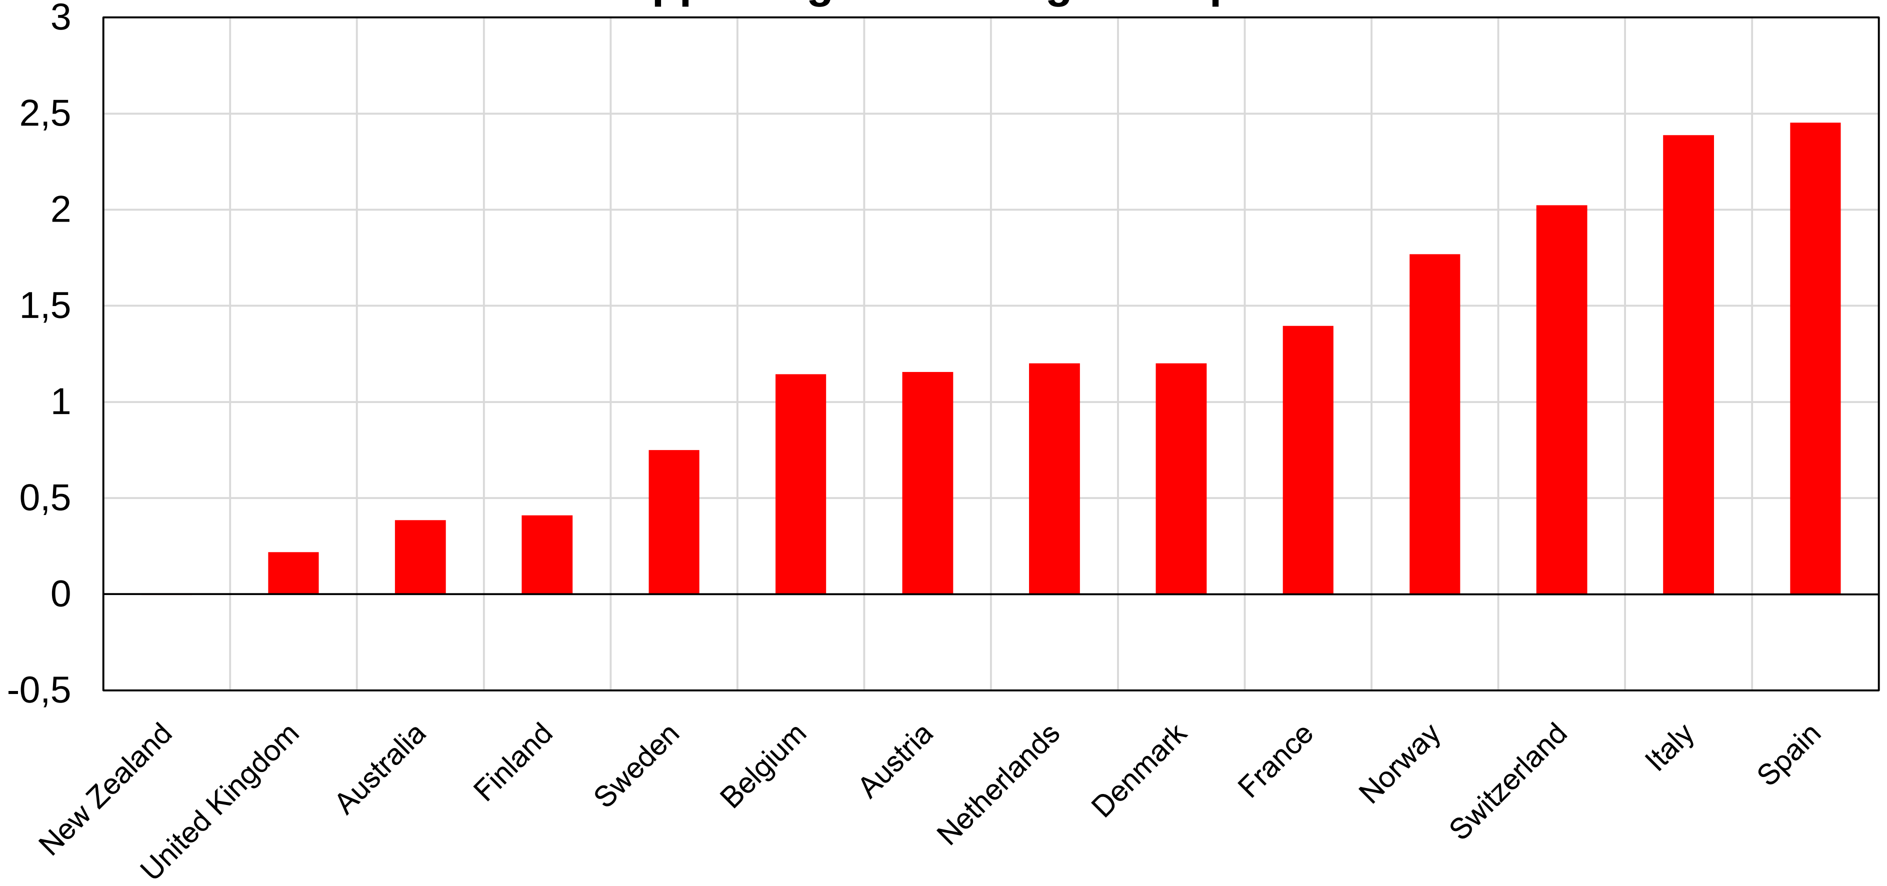

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the average self-declared left-right position of voters supporting anti-immigration parties and the average self-declared left-right position of all voters over the 2000-2020 period by country.
Interpretation: In nearly all countries, voters supporting anti-immigration parties are significantly more likely to declare being more right-wing than other voters.

Figure B20 - Average self-declared left-right position of voters supporting Social Democratic / Socialist / Communist / Other left-wing parties (excl. Greens)

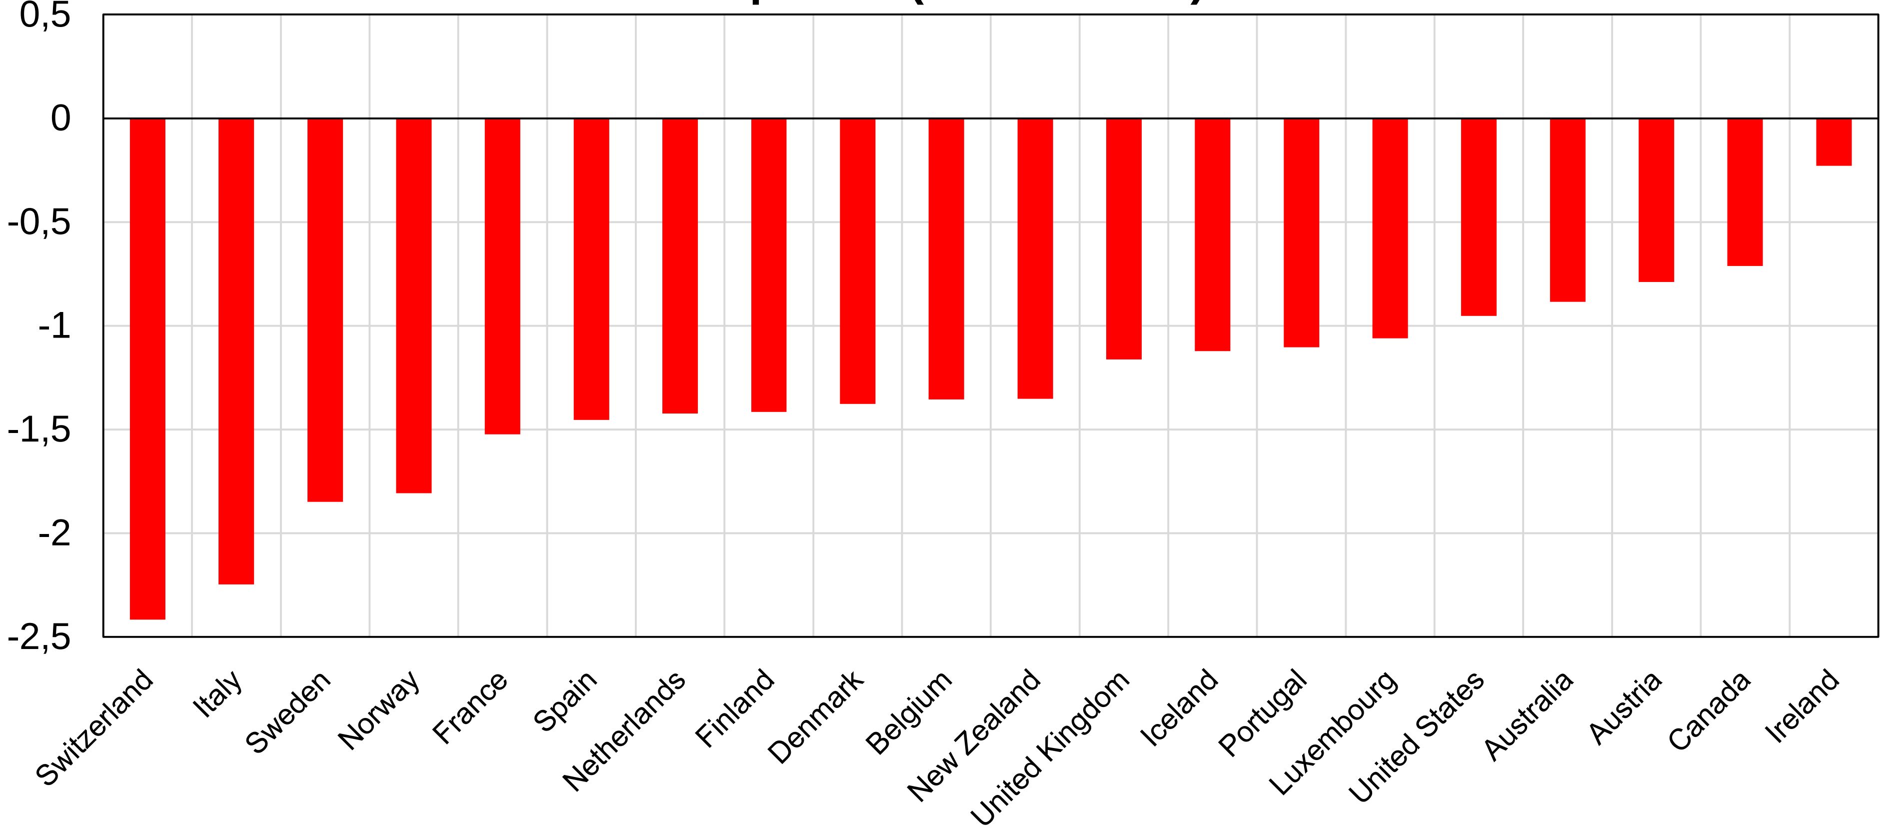

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the average self-declared left-right position of voters supporting Social Democratic, Socialist, communist and other left-wing parties (excluding Greens) and the average self-declared left-right position of all voters over the 2000-2020 period by country.
Interpretation: In all countries, voters supporting Social Democratic / Socialist / Communist / Other left-wing parties are significantly more likely to declare being more left-wing than other voters.

Figure B21 - Average self-declared left-right position of voters supporting green parties

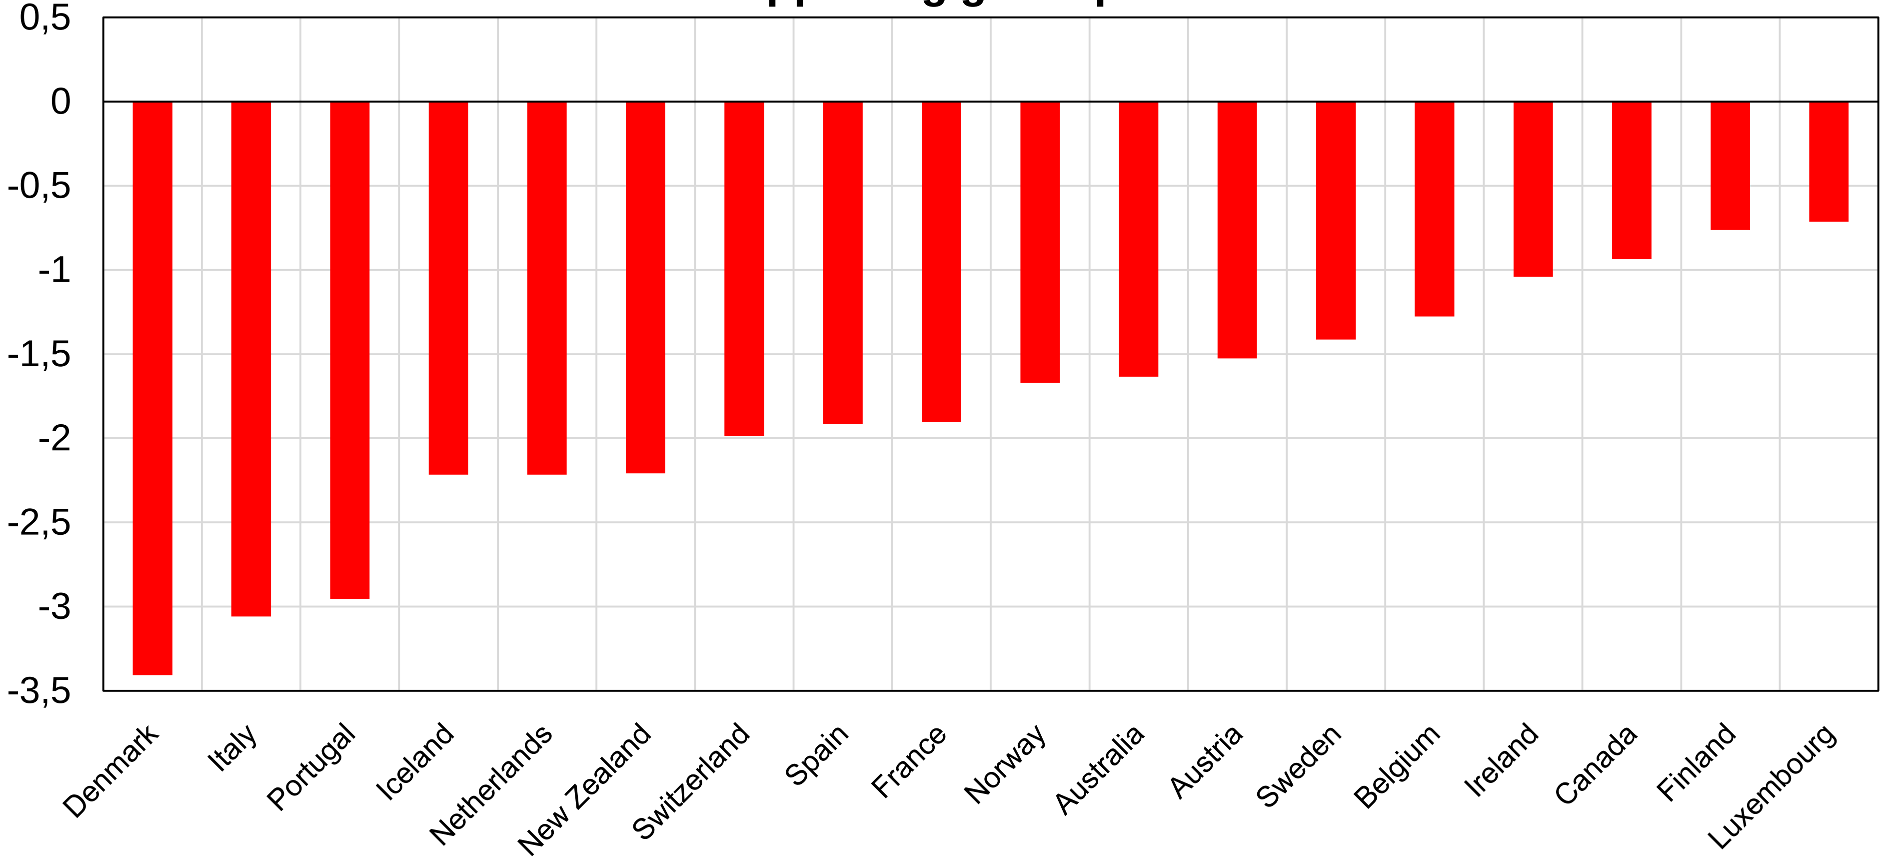

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the average self-declared left-right position of voters supporting green parties and the average self-declared left-right position of all voters over the 2000-2020 period by country.
Interpretation: In all countries, voters supporting green parties are significantly more likely to declare being more left-wing than other voters.

Figure B22 - Average CMP left-right ideological index of Conservative / Christian Democratic / Liberal parties in Manifesto Project Database

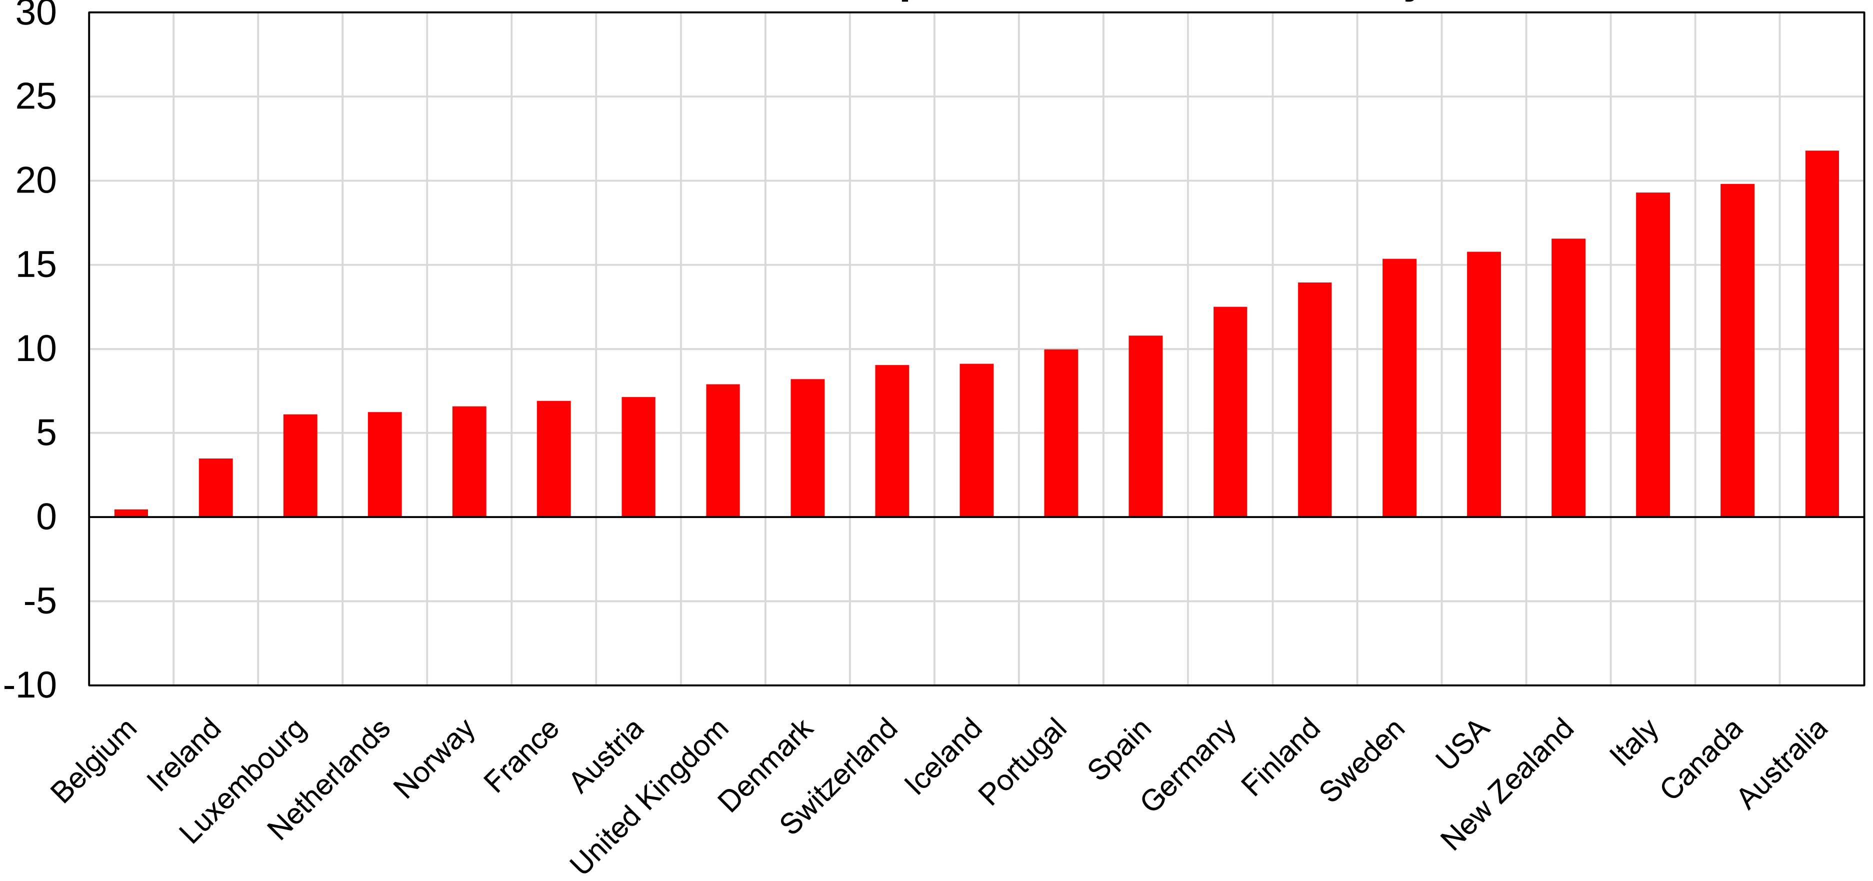

Source: authors' computations using the Comparative Manifesto Project Database.
Note: the figure represents the difference between the left-right ideological index of Conservative, Christian Democratic, and Liberal parties and the overall vote-share-weighted average of the same index (by country and election) over the 2000-2020 period by country.
Interpretation: In all countries, Conservative / Christian Democratic / Liberal parties have a left-right ideological index that is higher (that is, more right-wing) than that of other parties.

Figure B23 - Average CMP left-right ideological index of anti-immigration parties

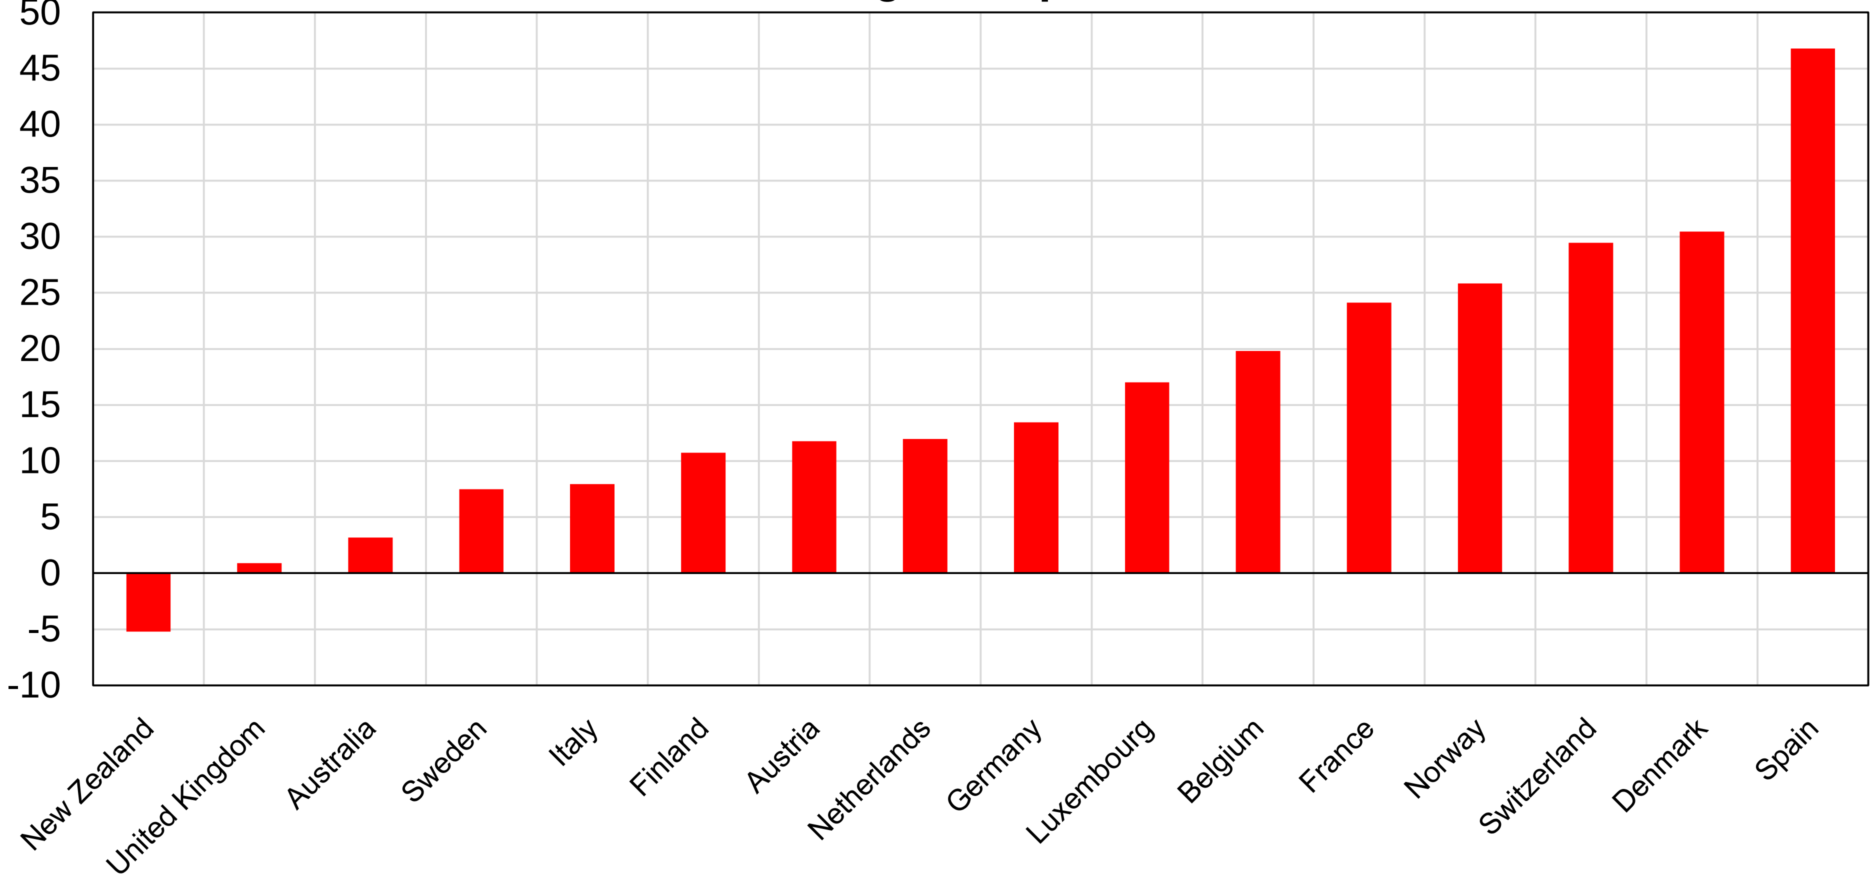

Source: authors' computations using the Comparative Manifesto Project Database.
Note: the figure represents the difference between the left-right ideological index of anti-immigration parties and the vote-share-weighted average of the same index (by country and election) over the 2000-2020 period by country.
Interpretation: In nearly all countries, anti-immigration parties have a left-right ideological index that is higher (that is, more right-wing) than that of other parties.

Figure B24 - Average CMP left-right ideological index of Social Democratic / Socialist / Communist / Other left-wing parties (excl. Greens)

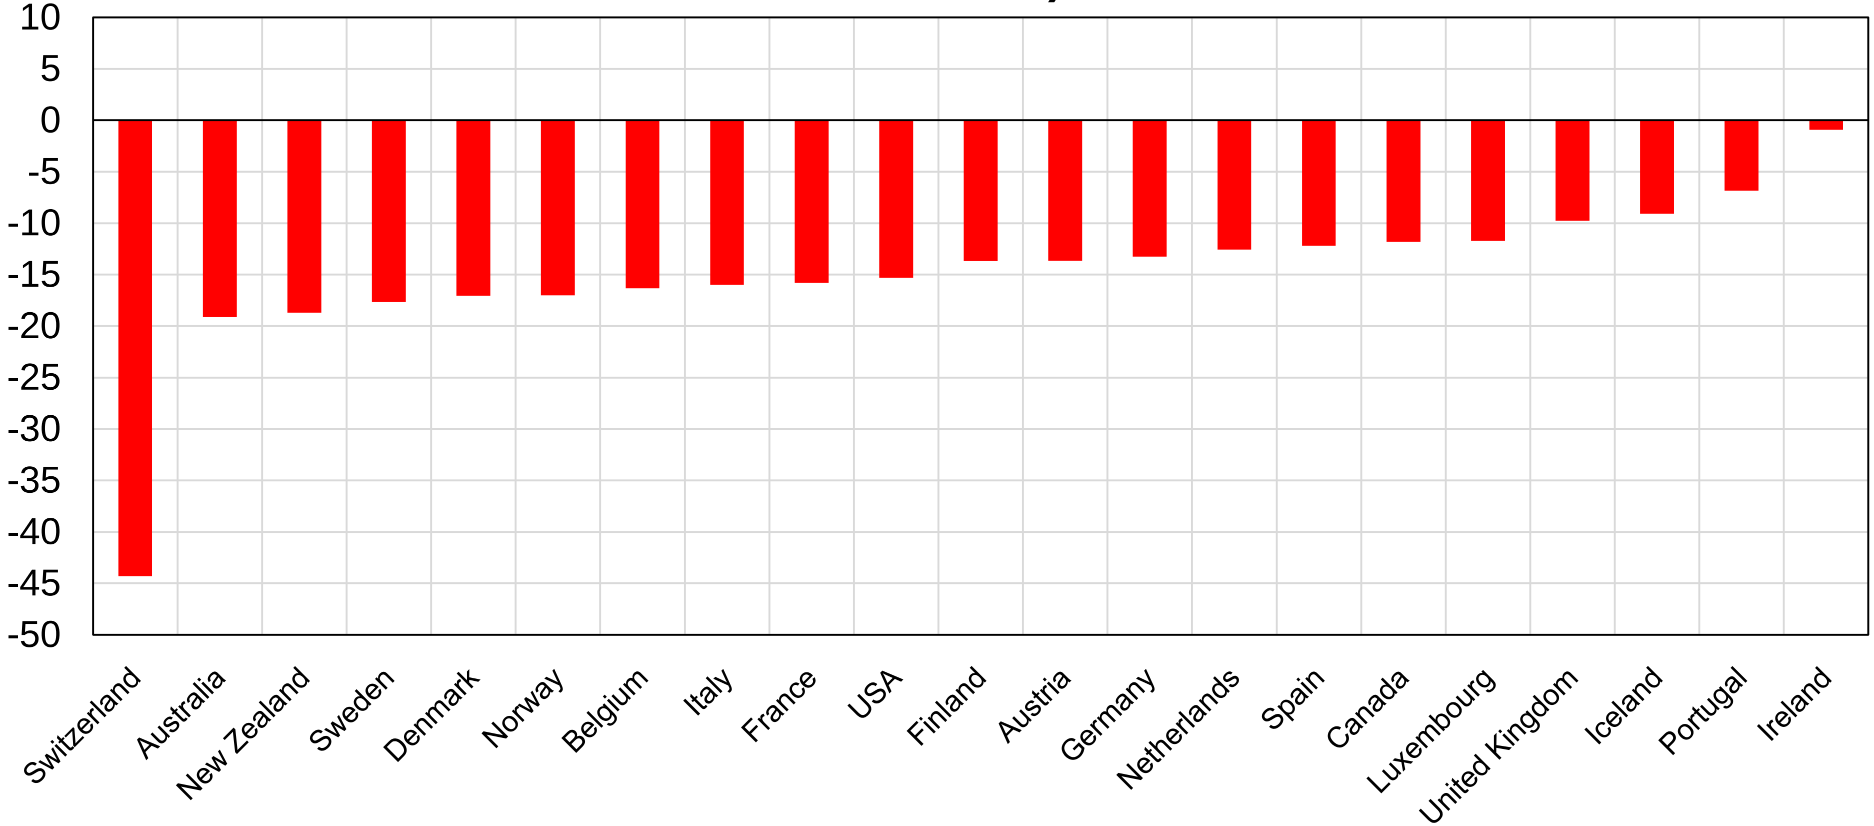

Source: authors' computations using the Comparative Manifesto Project Database.
Note: the figure represents the difference between the left-right ideological index of Social Democratic, Socialist, Communist and other left-wing parties (excluding Greens) and the vote-share-weighted average of the same index (by country and election) over the 2000-2020 period by country.
Interpretation: In all countries, Social Democratic / Socialist / Communist / Other left-wing parties parties have a left-right ideological index that is lower (that is, more left-wing) than that of other parties.

Figure B25 - Average CMP left-right ideological index of green parties

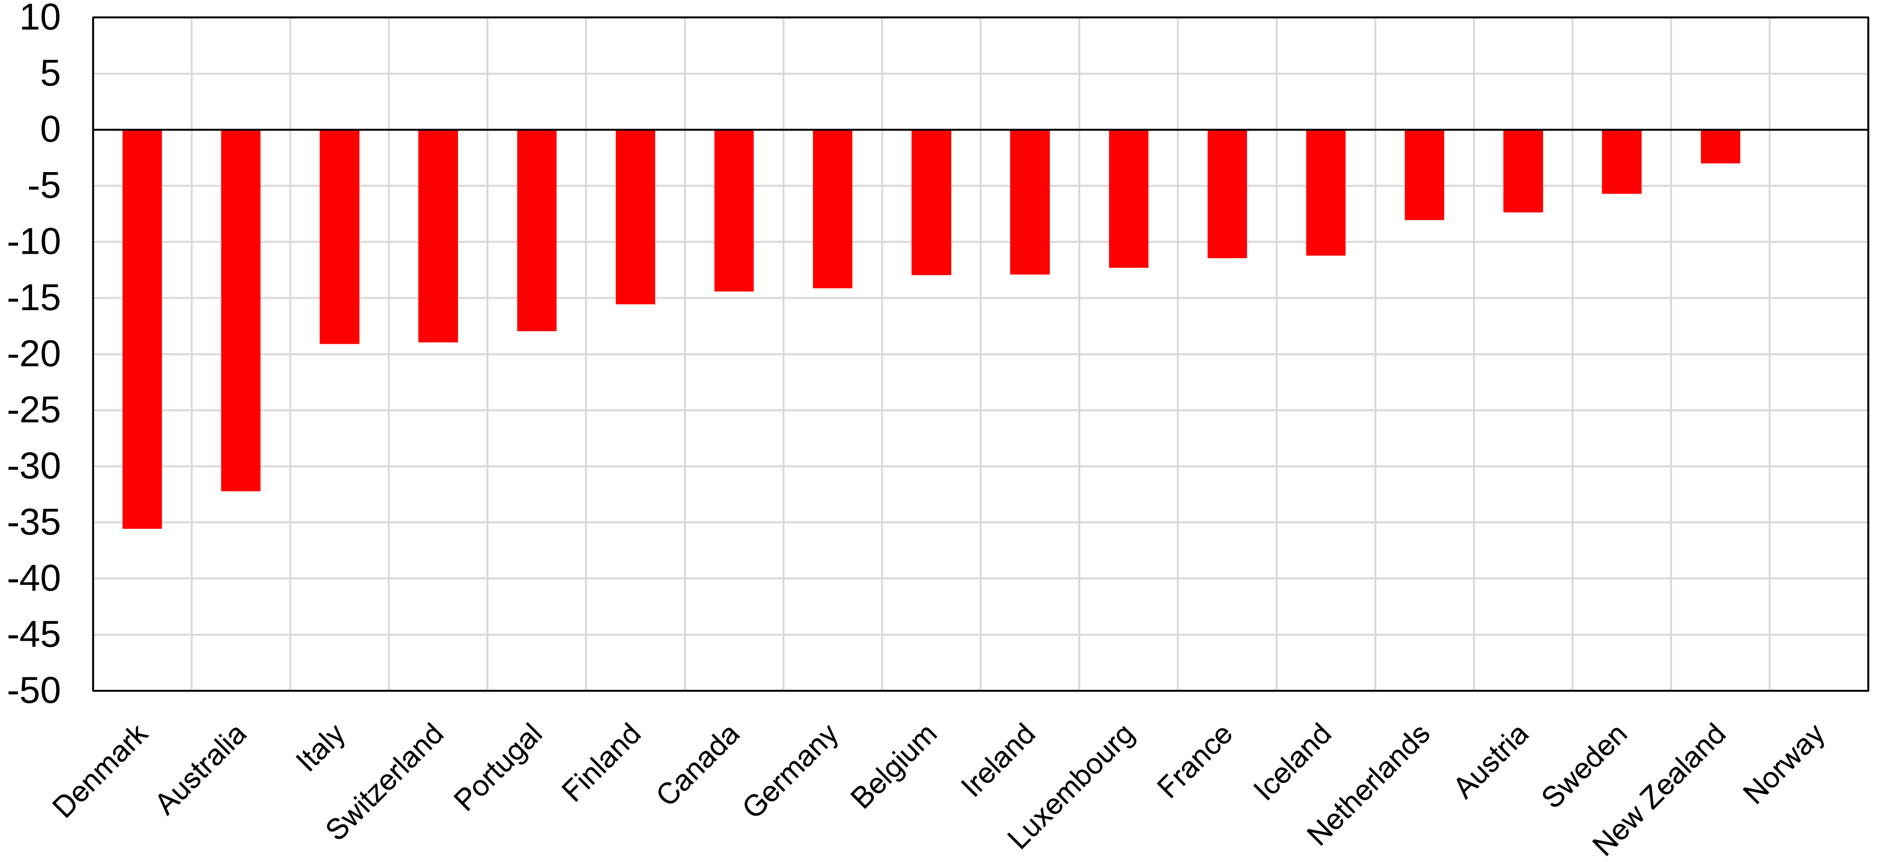

Source: authors' computations using the Comparative Manifesto Project Database.
Note: the figure represents the difference between the left-right ideological index of green parties and the vote-share-weighted average of the same index (by country and election) over the 2000-2020 period by country.
Interpretation: In all countries, green parties have a left-right ideological index that is lower (that is, more left-wing) than that of other parties.

Figure CA1 - Vote for left-wing parties among young voters in Western democracies

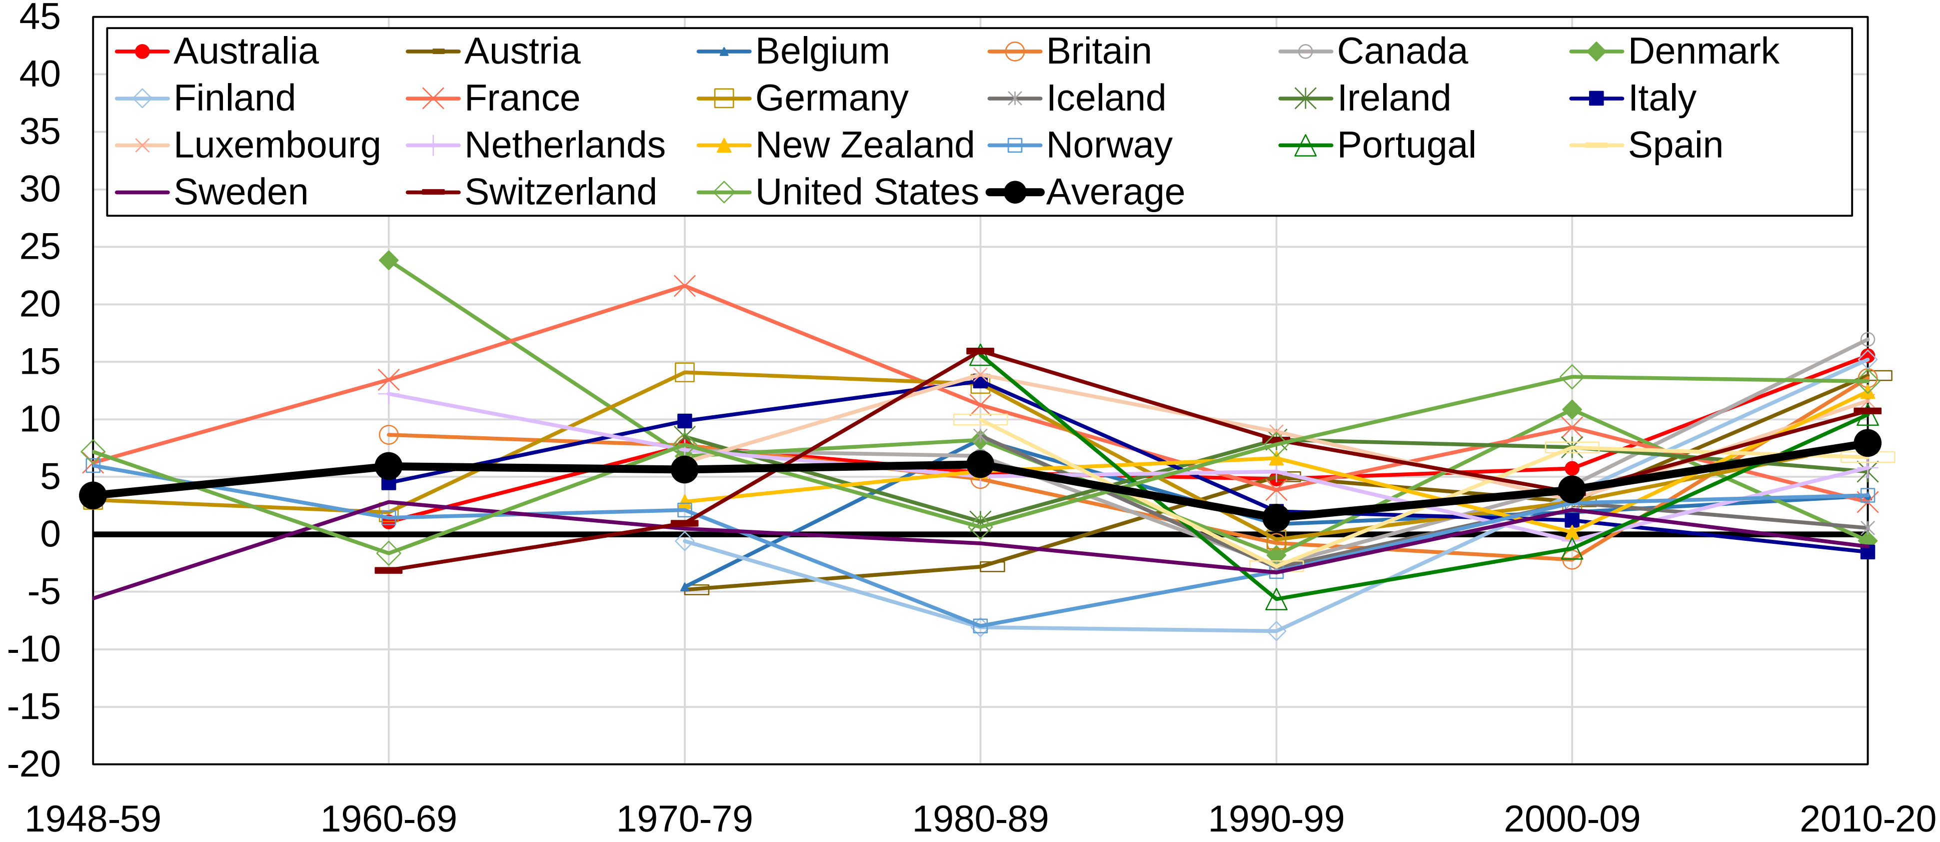

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the difference between the share of voters younger than 25 and the share of voters aged 25 or above voting for left-wing parties in Western democracies.

Figure CA2 - Vote for left-wing parties among old voters in Western democracies

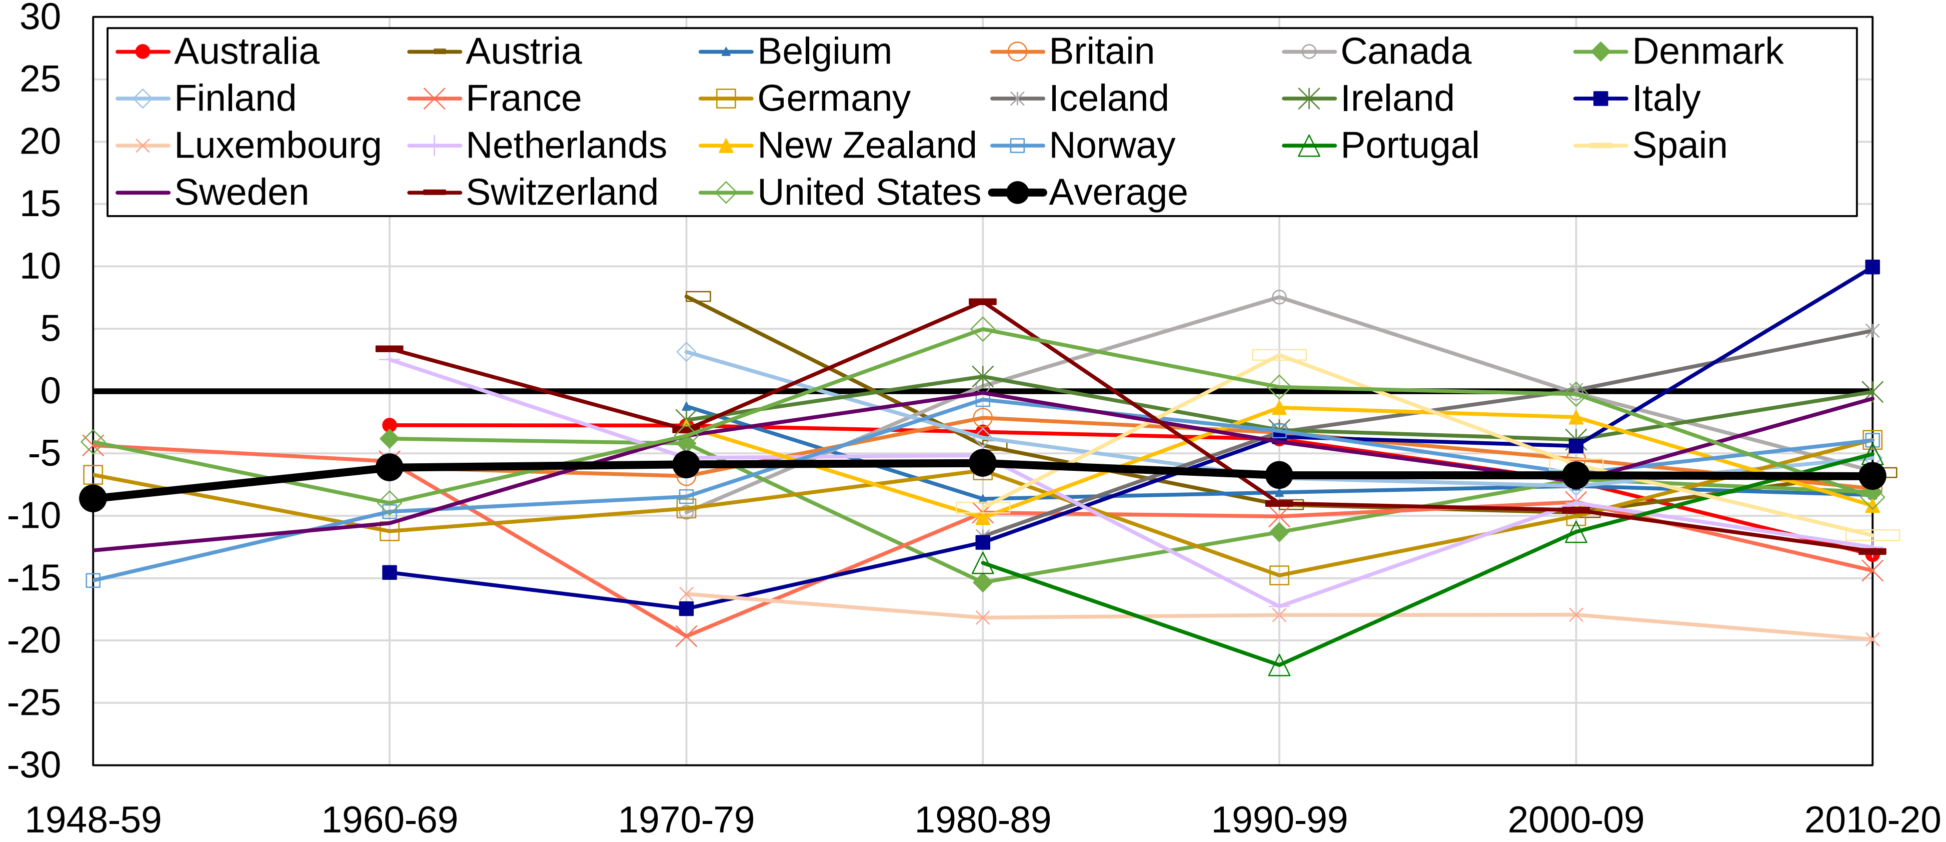

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the difference between the share of the 10% oldest voters and the share of the youngest 90% voters voting for left-wing parties in Western democracies.

Figure CA3 - Vote for left-wing parties among young voters in Western democracies, after controls

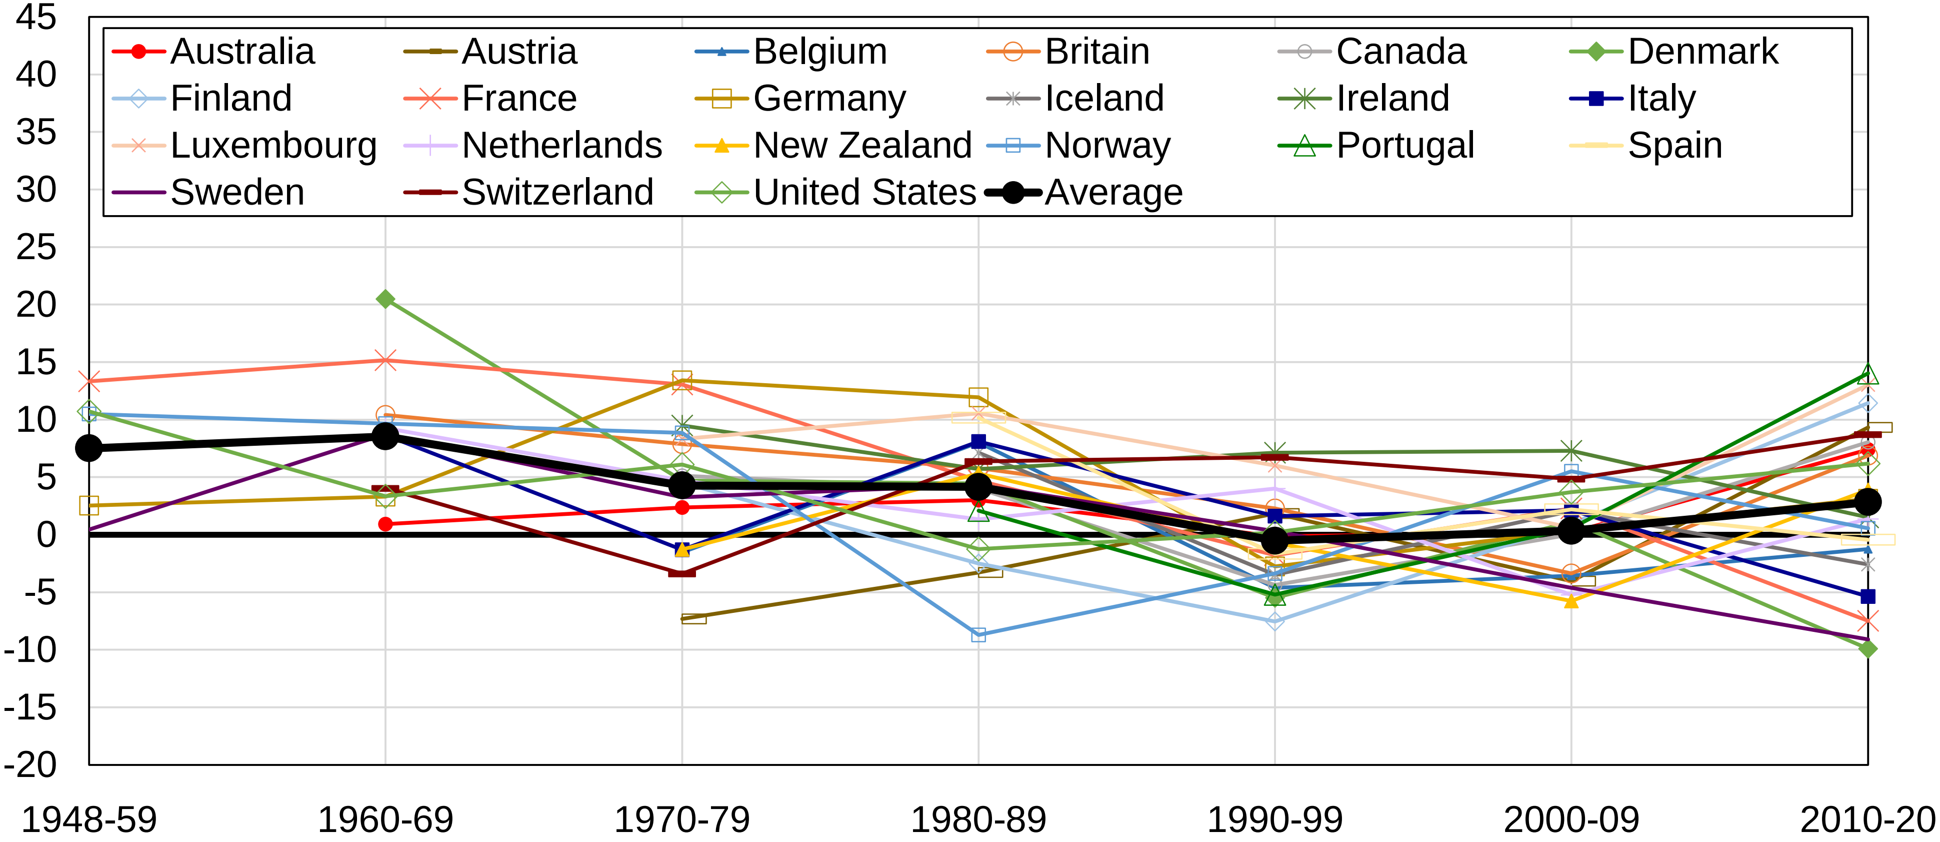

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the difference between the share of voters younger than 25 and the share of voters aged 25 or above voting for left-wing parties in Western democracies, after controlling for income, education, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status.

Figure CA4 - Vote for left-wing parties among old voters in Western democracies, after controls

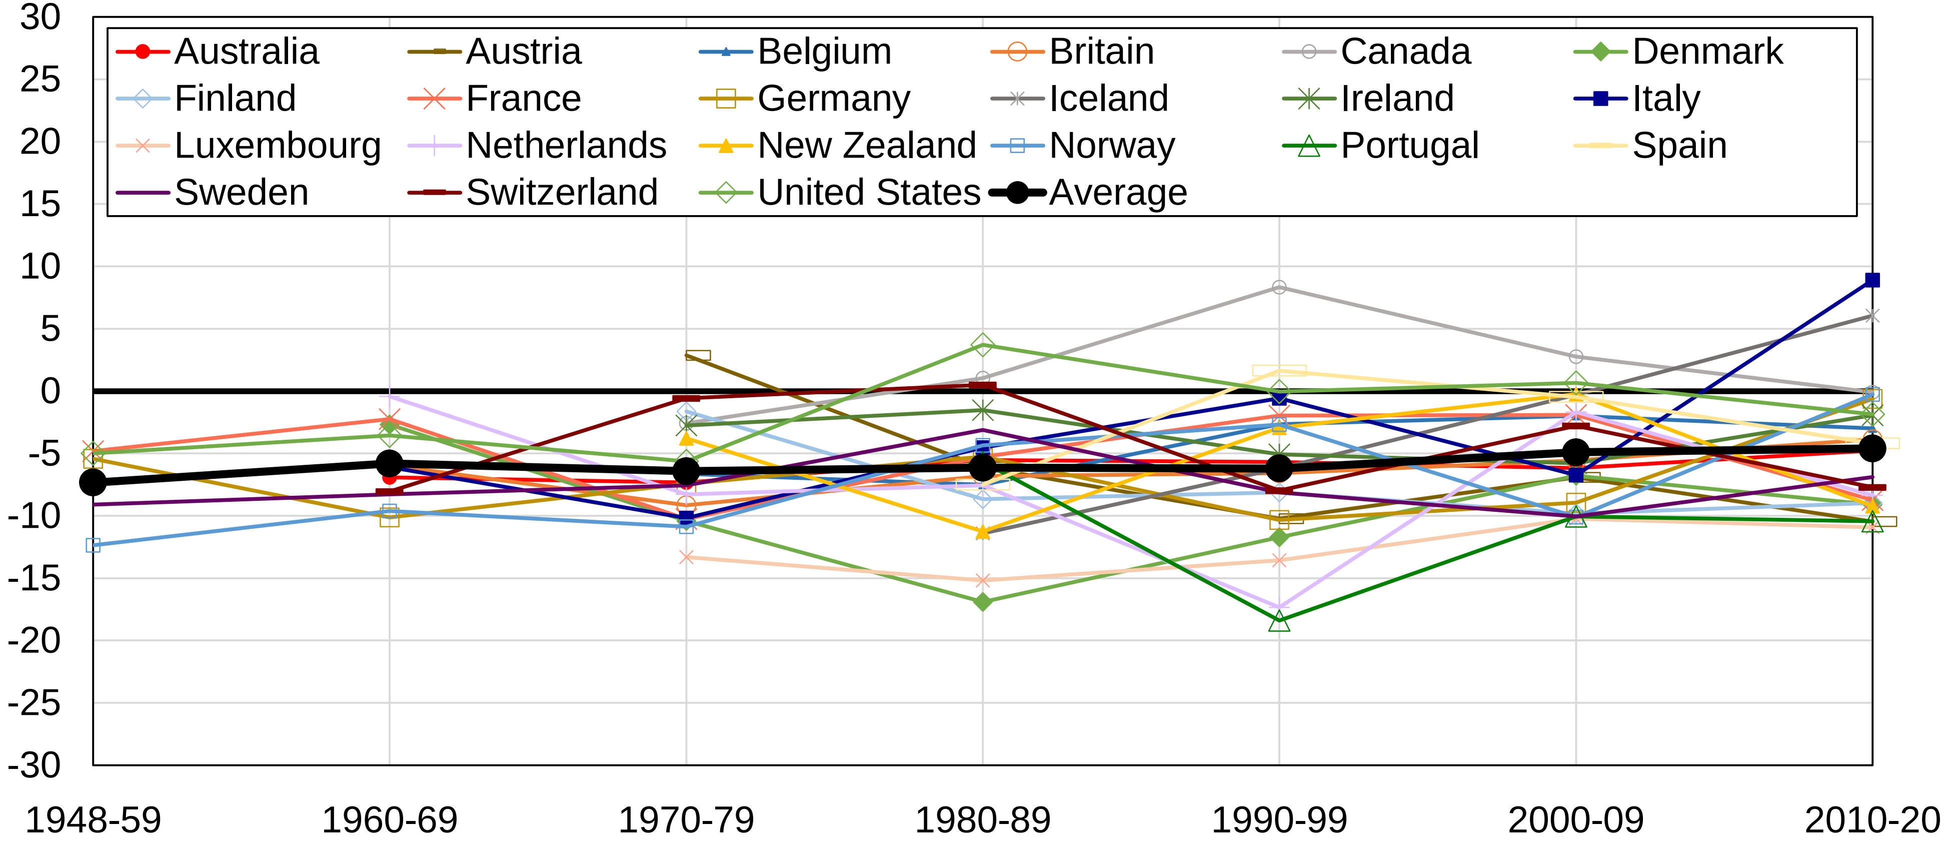

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the difference between the share of the 10% oldest voters and the share of the youngest 90% voters voting for left-wing parties in Western democracies, after controlling for income, education, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status.

Figure CA5 - Generational cleavages and party system fragmentation

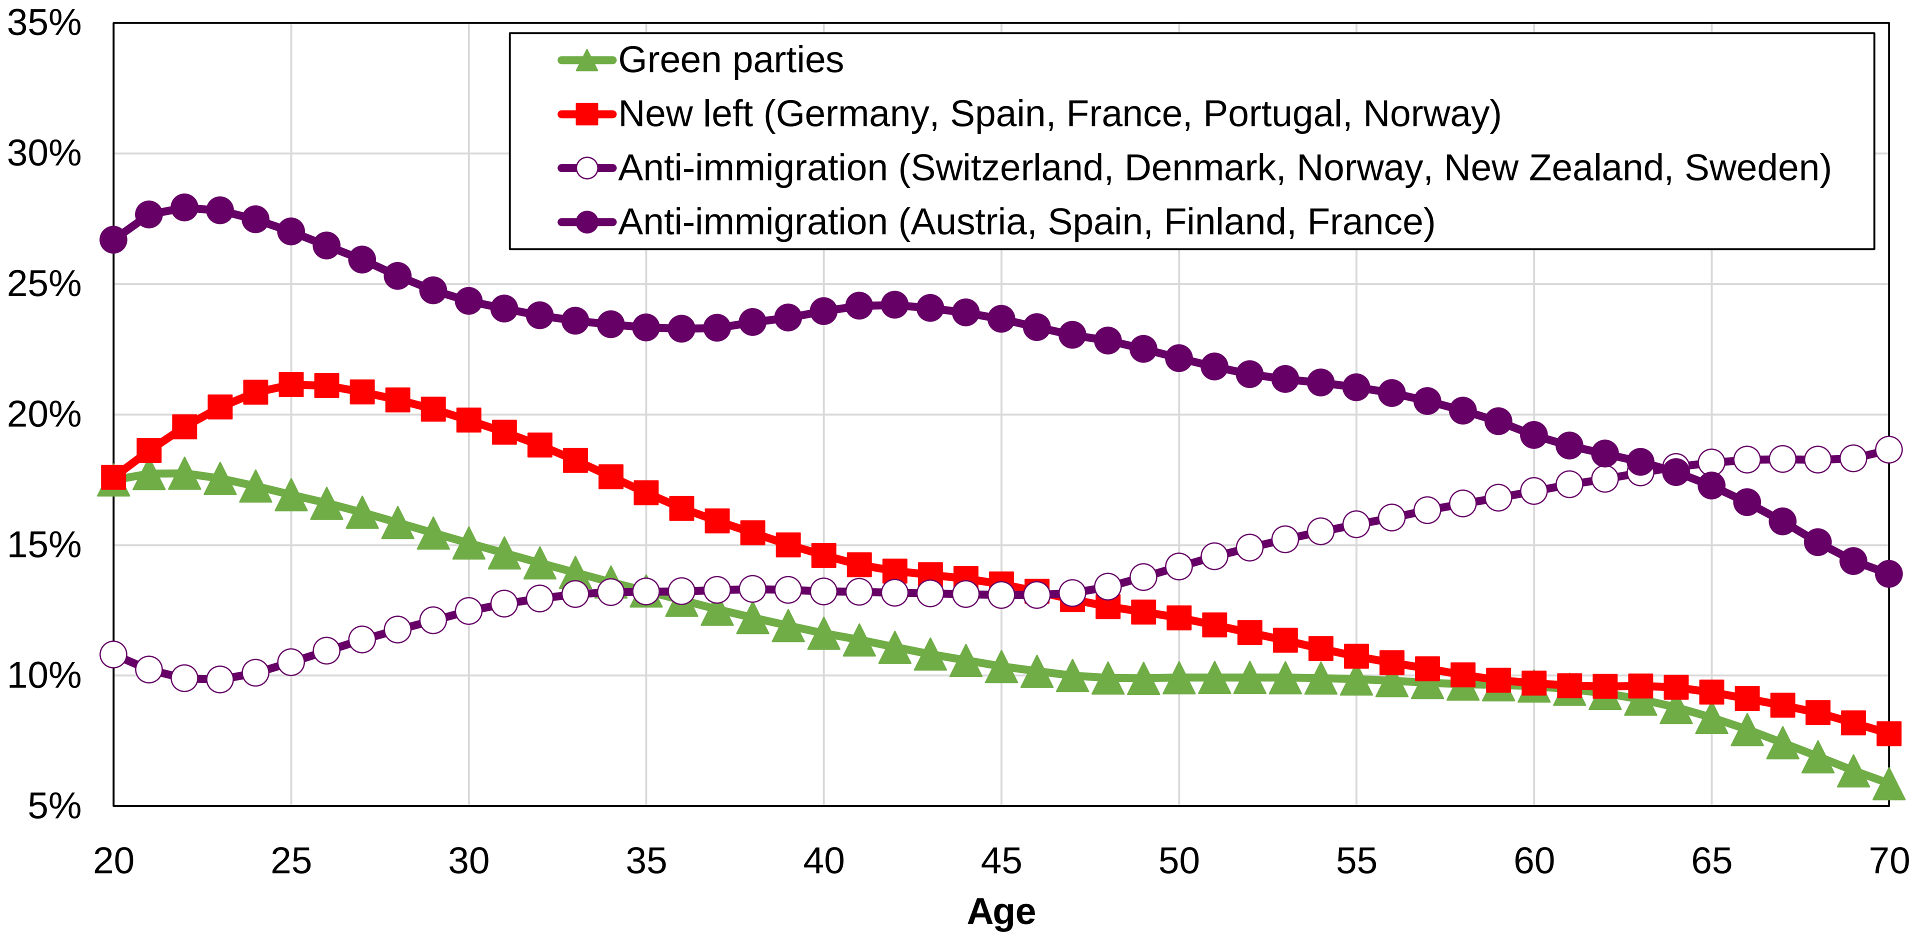

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the share of votes received by selected groups of parties in Western democracies by age in the last election available. Green parties and "New left" parties (Die Linke, Podemos, France Insoumise, Bloco de Esquerda, Norwegian Socialist Left Party) make much higher scores among the youth than among older generations. By contrast, there is no clear age profile in the case of far-right or anti-immigration parties. 20 correponds to voters aged 20 or younger; 70 corresponds to voters 70 or older.

Figure CA6 - Vote for Green parties by age group

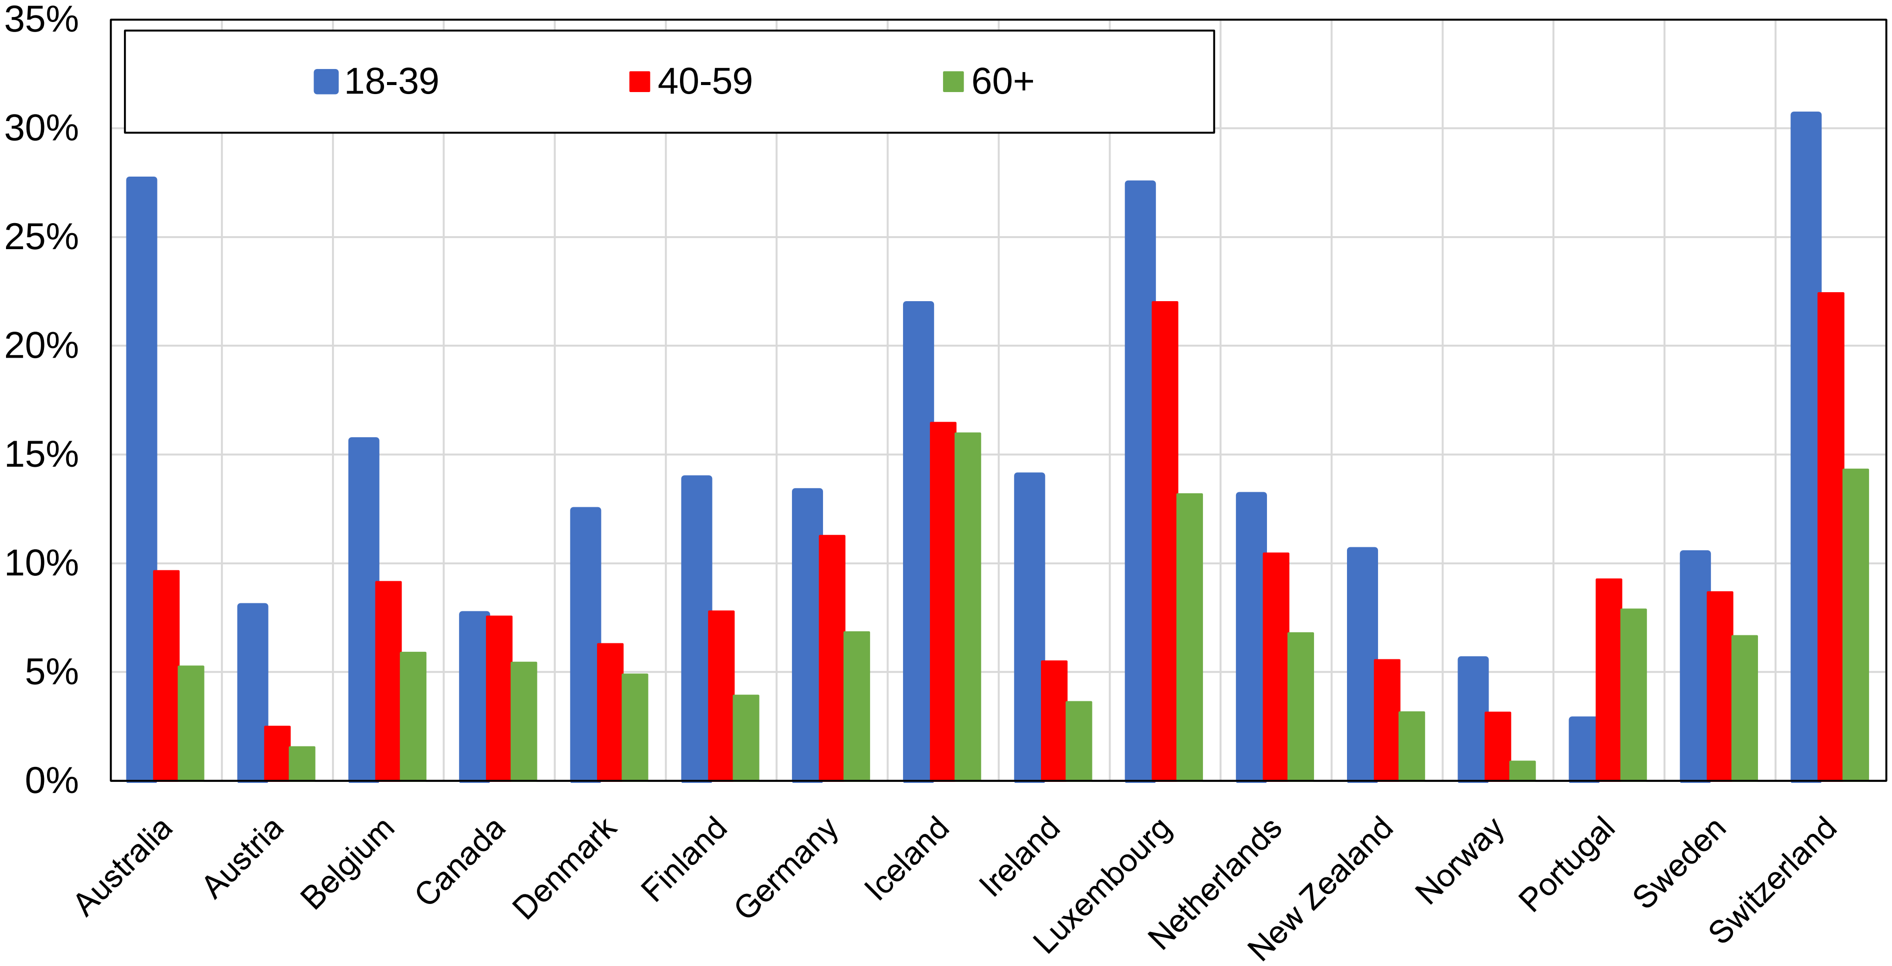

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by Green parties in Western democracies in the last election available by age group.

Figure CA7 - Vote for anti-immigration parties by age group

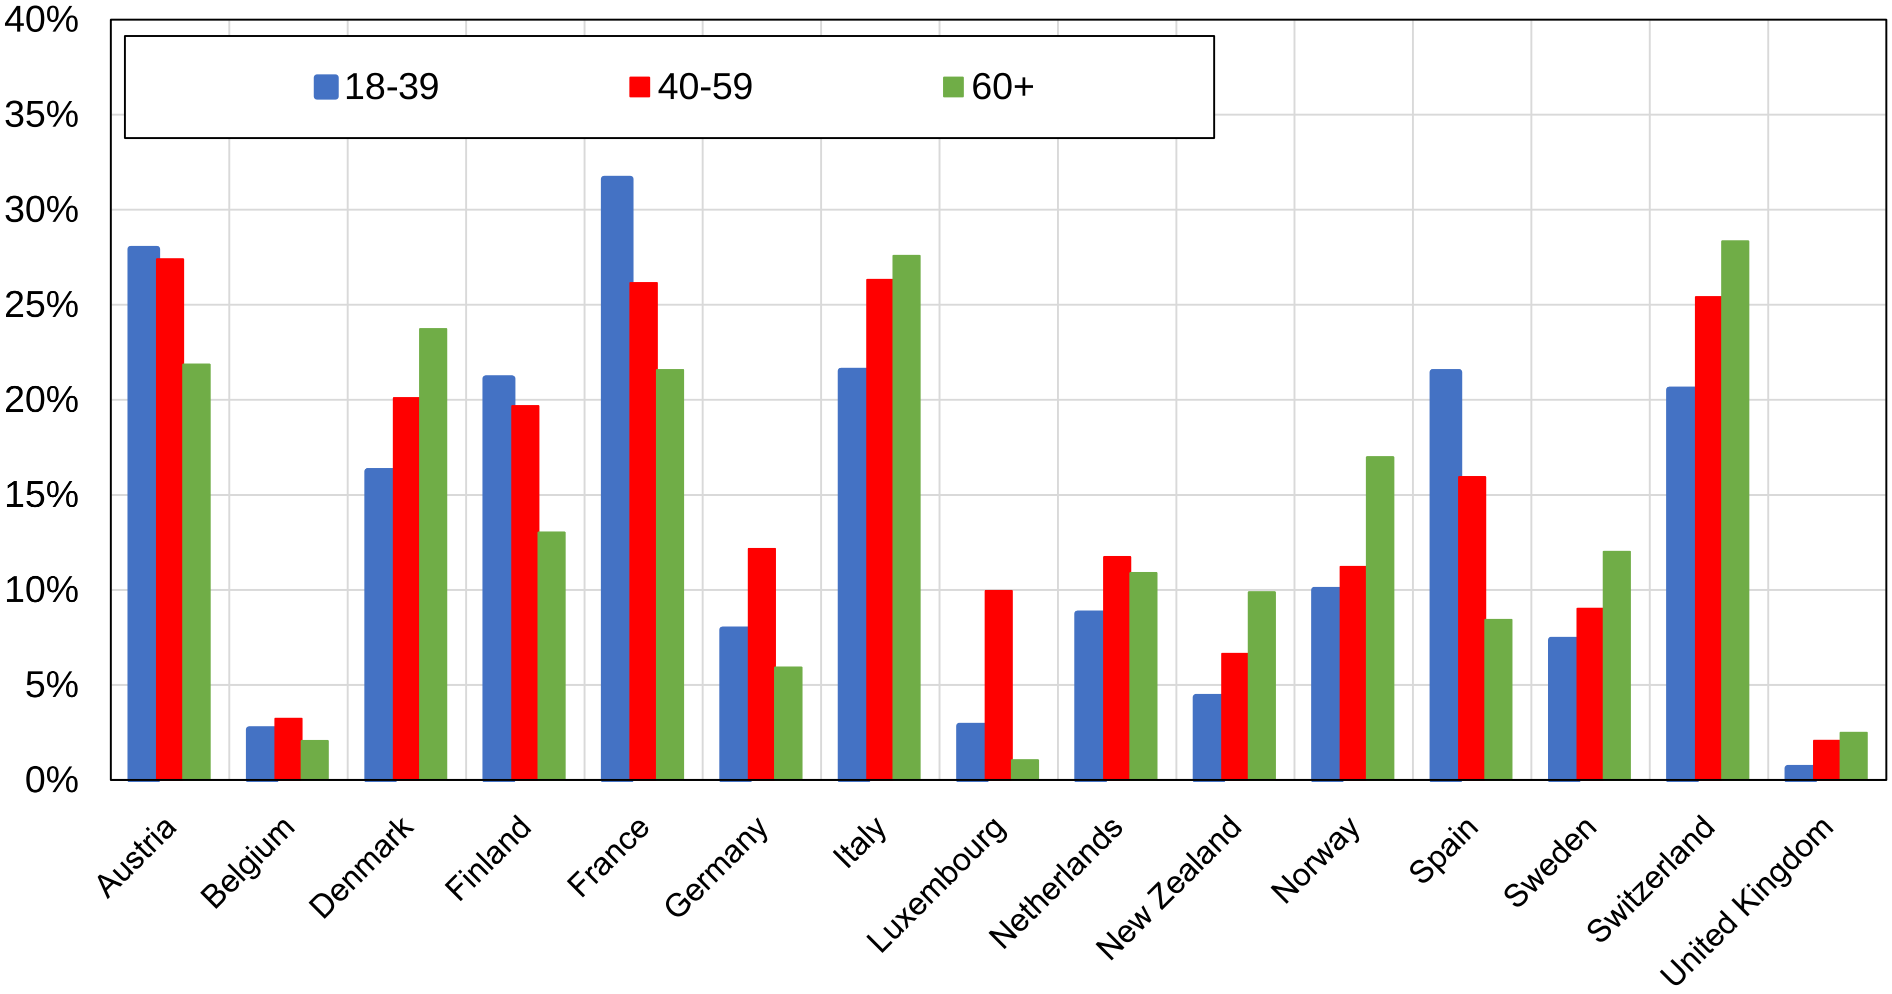

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by anti-immigration parties in Western democracies in the last election available by age group.

Figure CB1 - The rural-urban divide

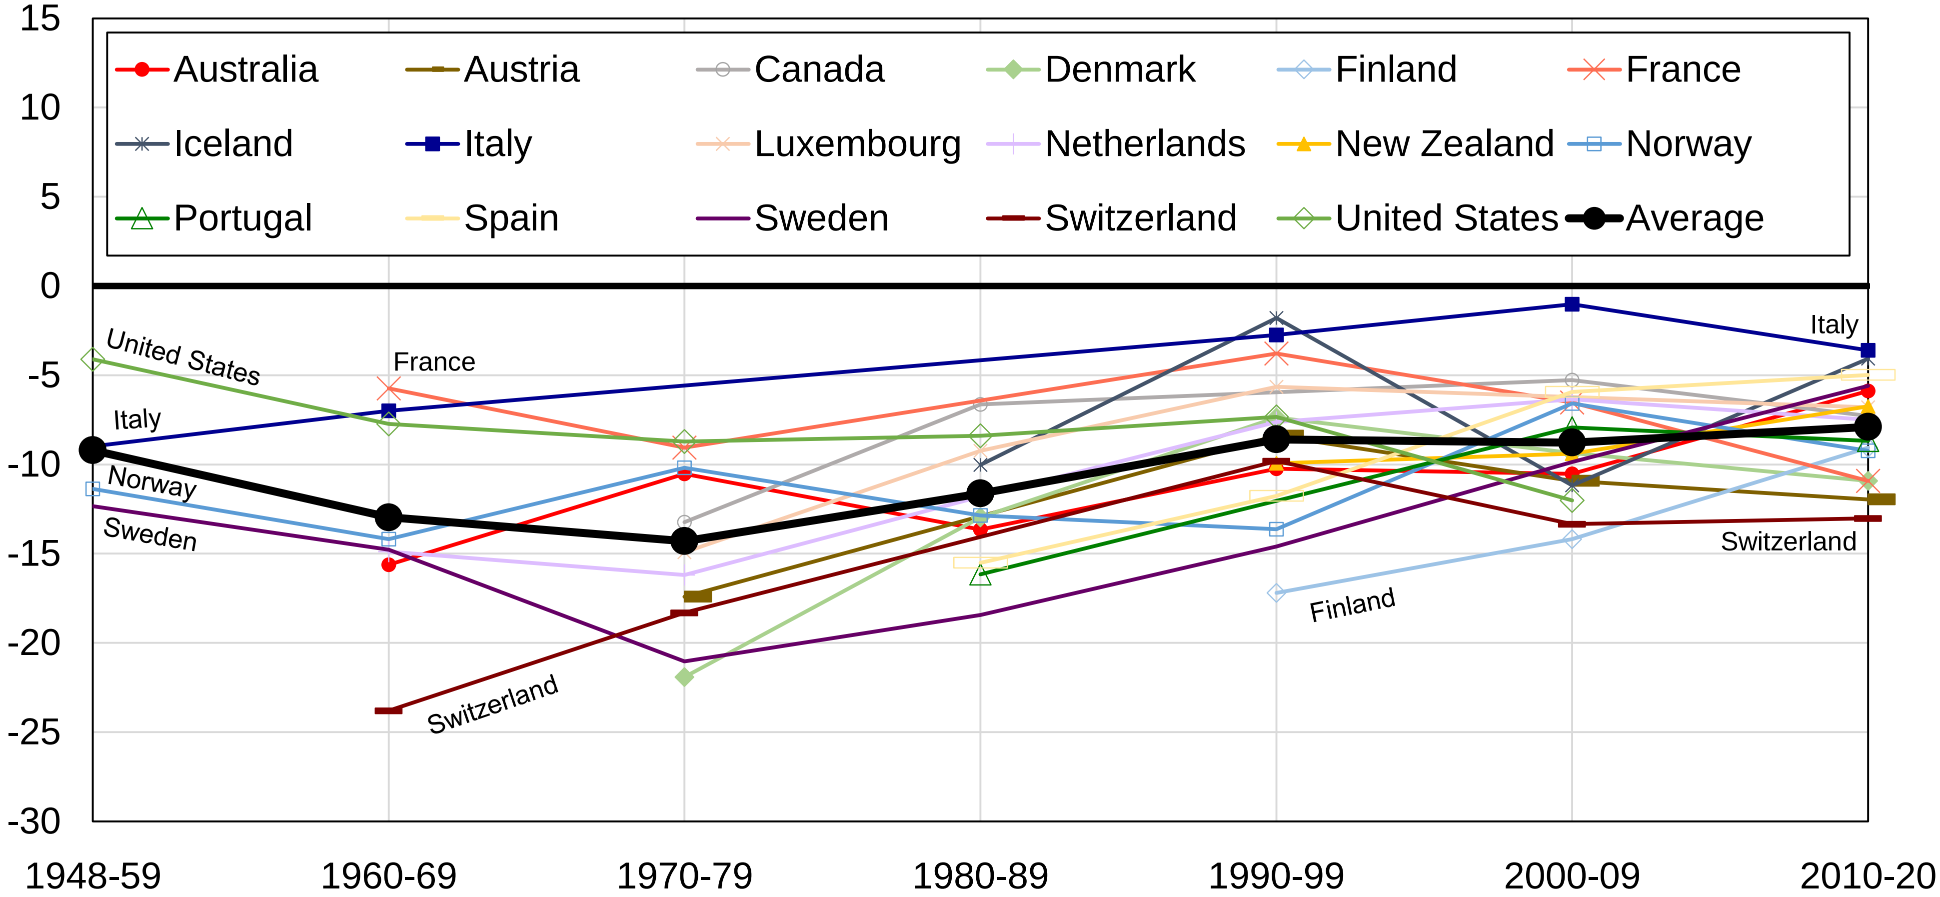

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure displays the difference between the share of rural areas and the share of urban areas voting for social democratic / socialist / communist / green parties. In all countries, rural areas have remained significantly less likely to vote for these parties than cities, with no clear trend over time. Estimates control for income, education, age, gender, employment status, and marital status (in country-years for which these variables are available).

Figure CB2 - Vote for Green parties by rural-urban location in Western democracies

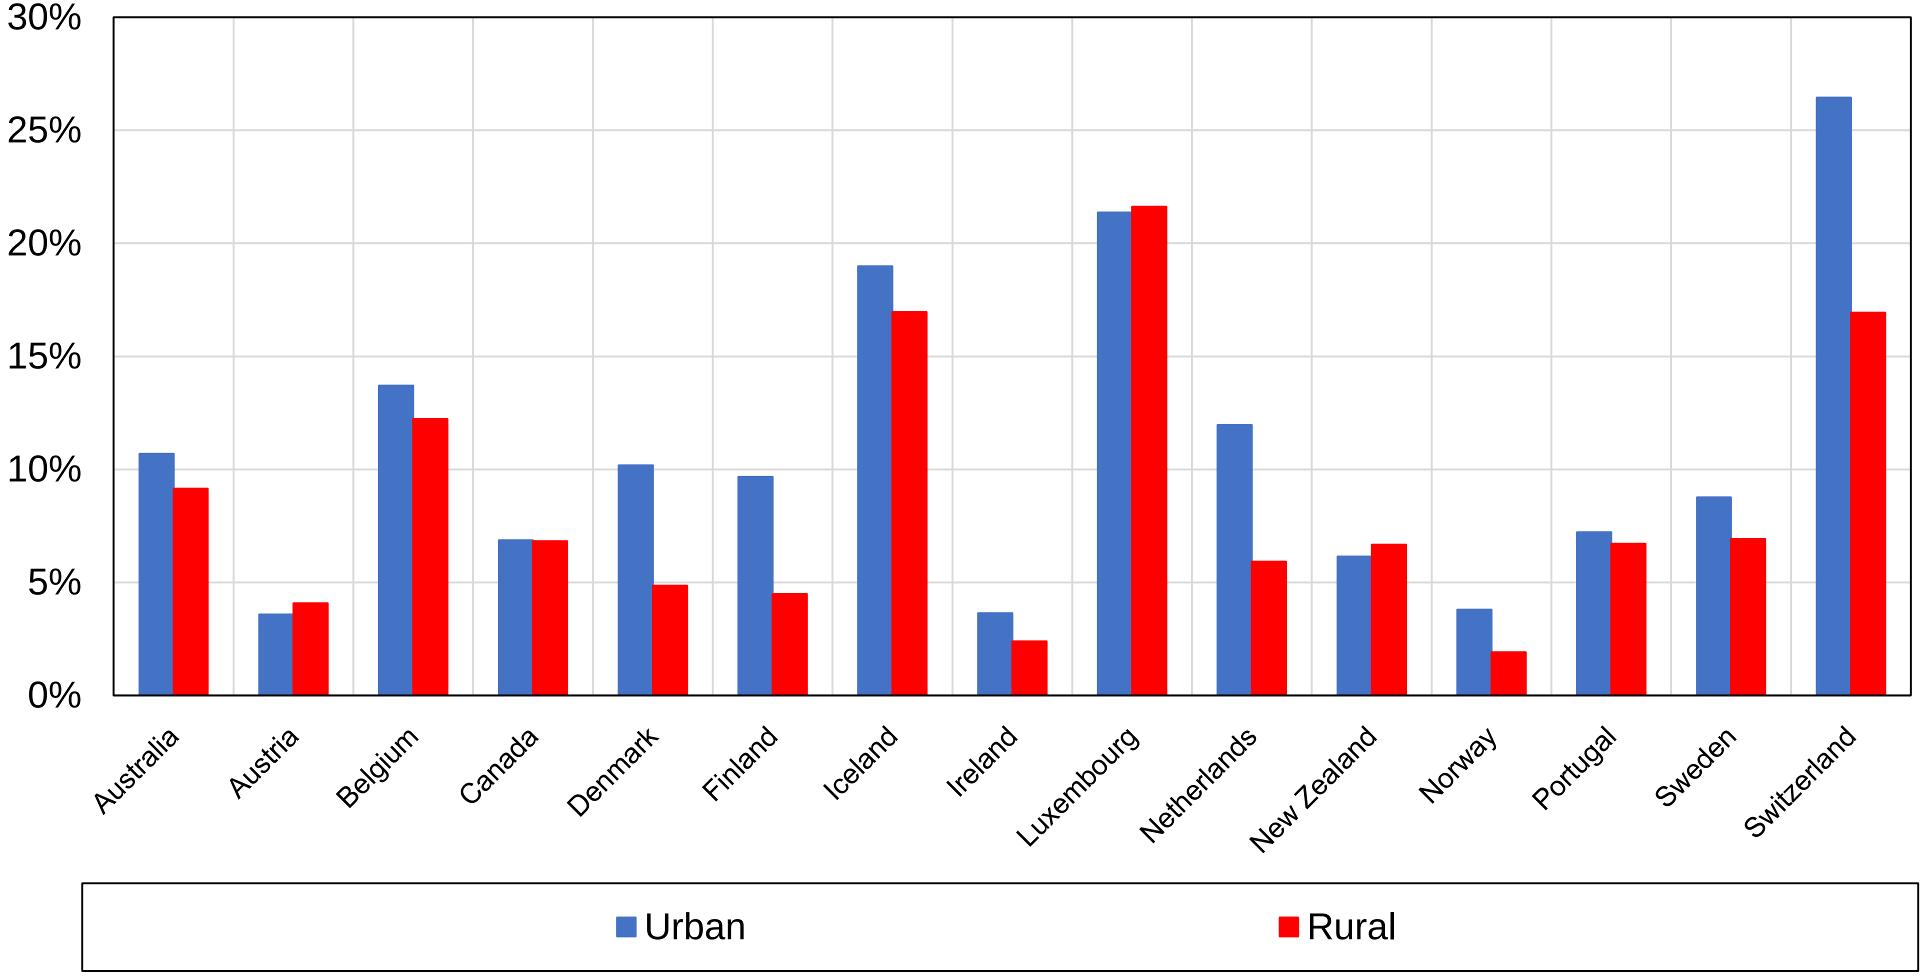

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by Green parties by rural-urban location in Western democracies.

Figure CB3 - Vote for anti-immigration parties by rural-urban location in Western democracies

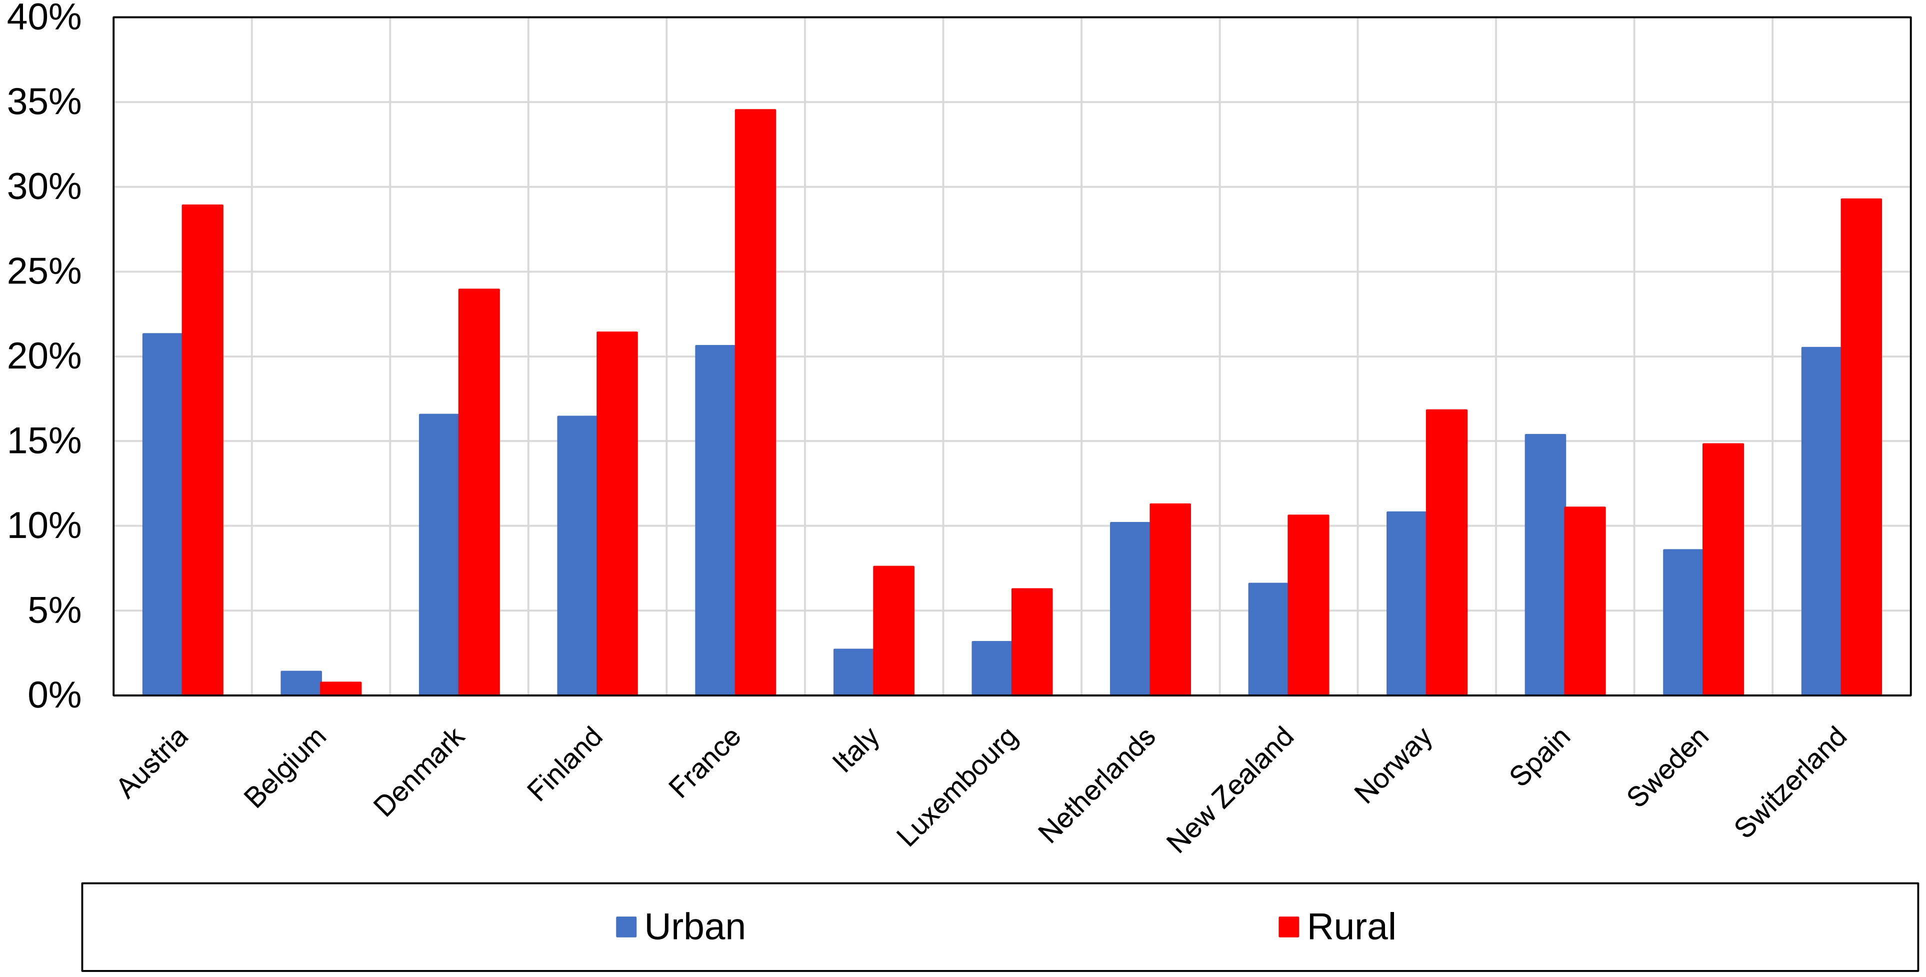

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by anti-immigration parties by rural-urban location in Western democracies.

Figure CB4 - Vote for left-wing parties by center-periphery location in Western democracies

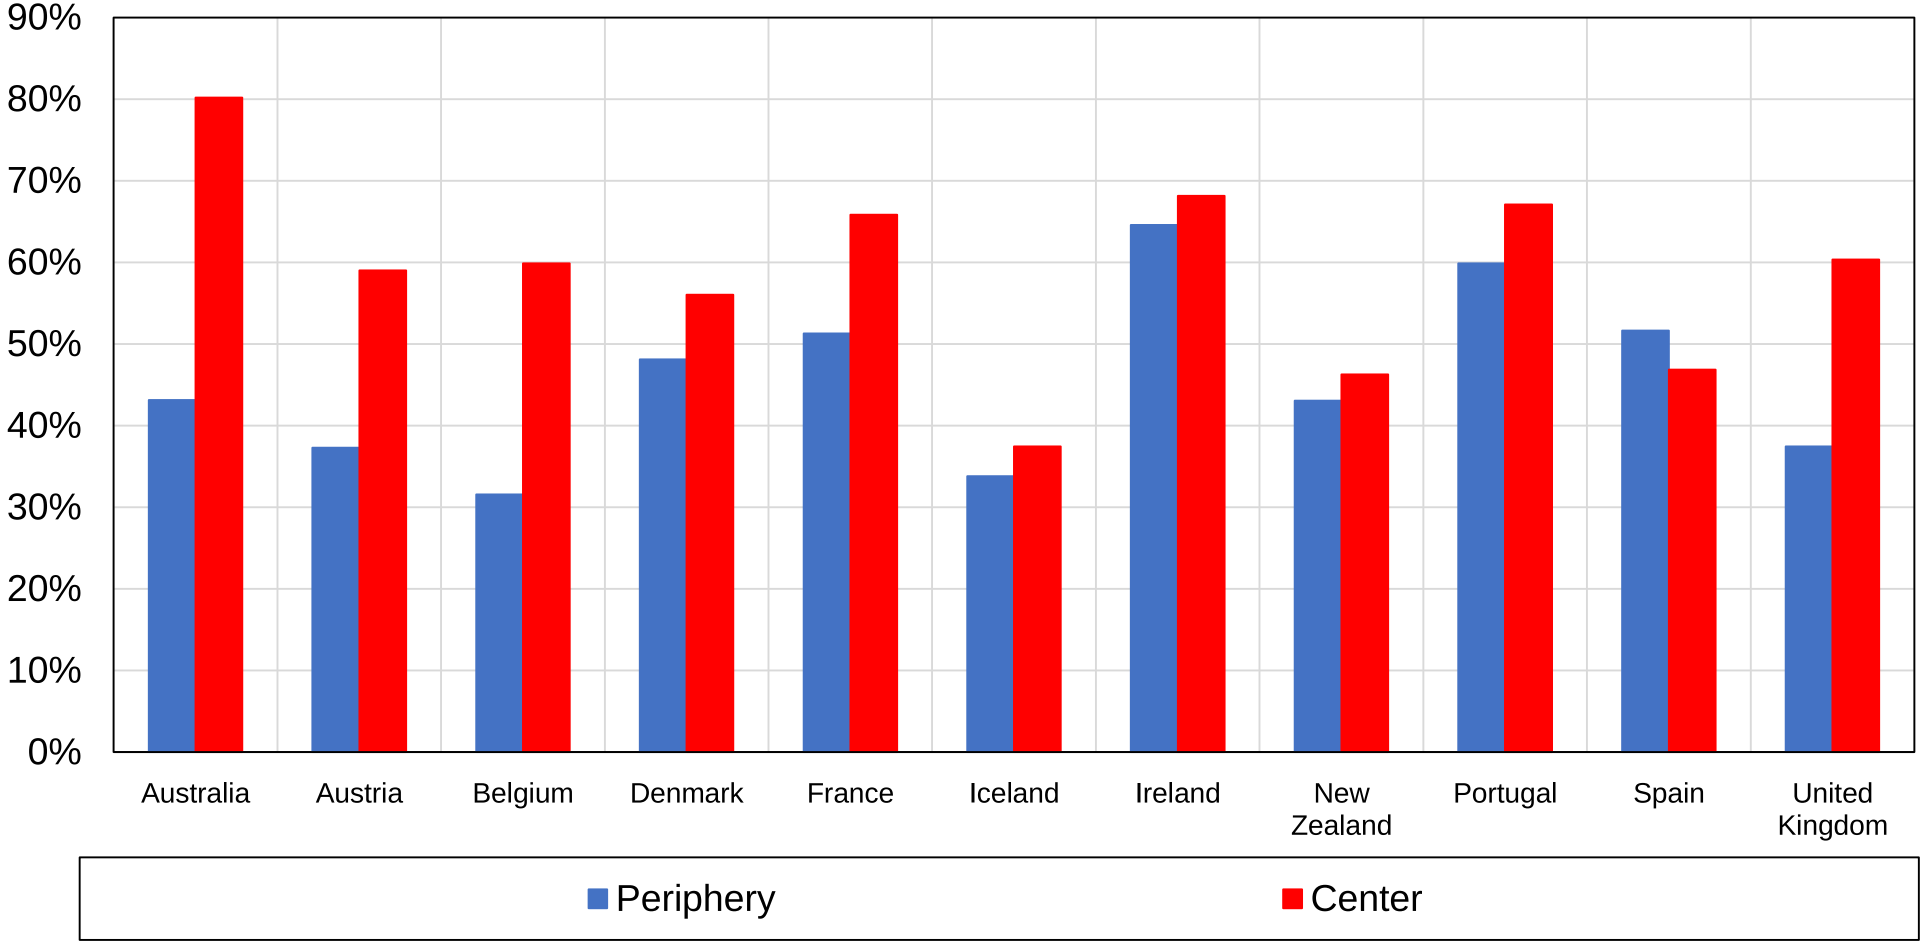

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by left-wing parties by center-periphery location in Western democracies. Centers correspond to the Australian Capital Territory (Australia), Vienna (Austria), Brussels (Belgium), Copenhagen (Denmark), Paris (France), Reykjavík (Iceland), Dublin (Ireland), Auckland and Wellington (New Zealand), Lisbon (Portugal), Madrid (Spain), and London (United Kingdom).

Figure CB5 - Vote for Green parties by center-periphery location in Western democracies

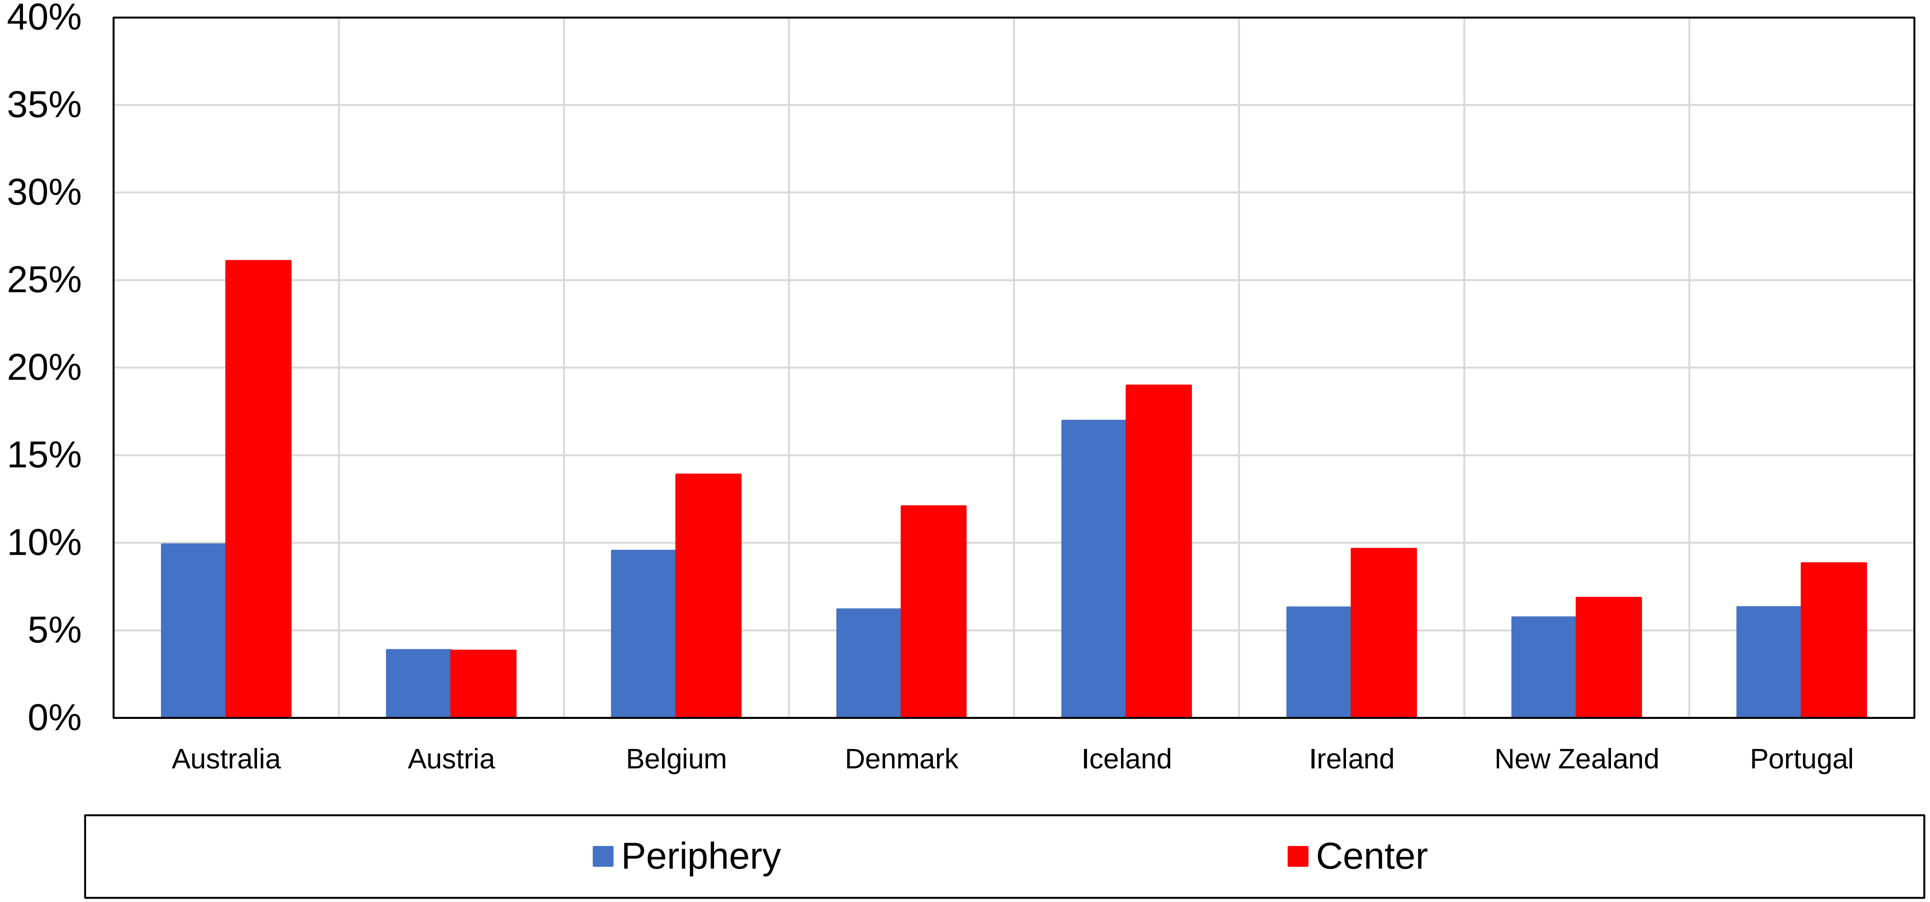

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by Green parties by center-periphery location in Western democracies. Centers correspond to the Australian Capital Territory (Australia), Vienna (Austria), Brussels (Belgium), Copenhagen (Denmark), Paris (France), Reykjavík (Iceland), Dublin (Ireland), Auckland and Wellington (New Zealand), Lisbon (Portugal), Madrid (Spain), and London (United Kingdom).

Figure CB6 - Vote for anti-immigration parties by center-periphery location in Western democracies

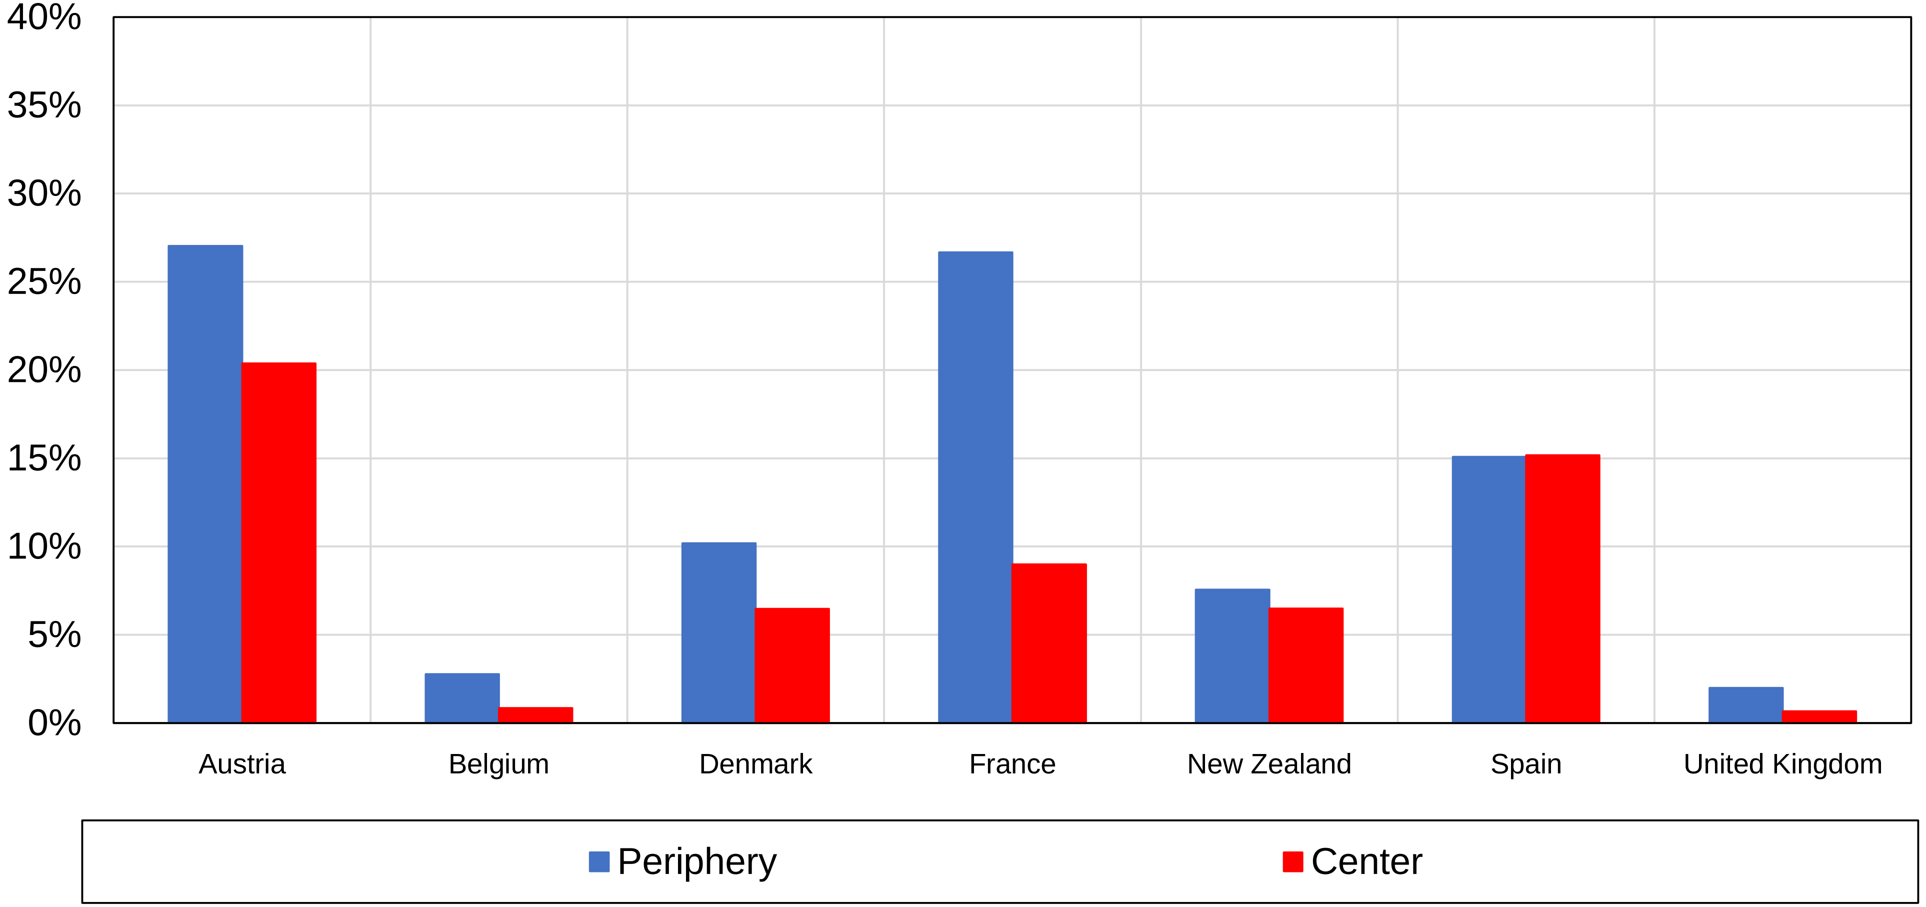

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by anti-immigration parties by center-periphery location in Western democracies. Centers correspond to the Australian Capital Territory (Australia), Vienna (Austria), Brussels (Belgium), Copenhagen (Denmark), Paris (France), Reykjavík (Iceland), Dublin (Ireland), Auckland and Wellington (New Zealand), Lisbon (Portugal), Madrid (Spain), and London (United Kingdom).

Figure CB7 - Vote for left-wing parties among capital cities in Western democracies

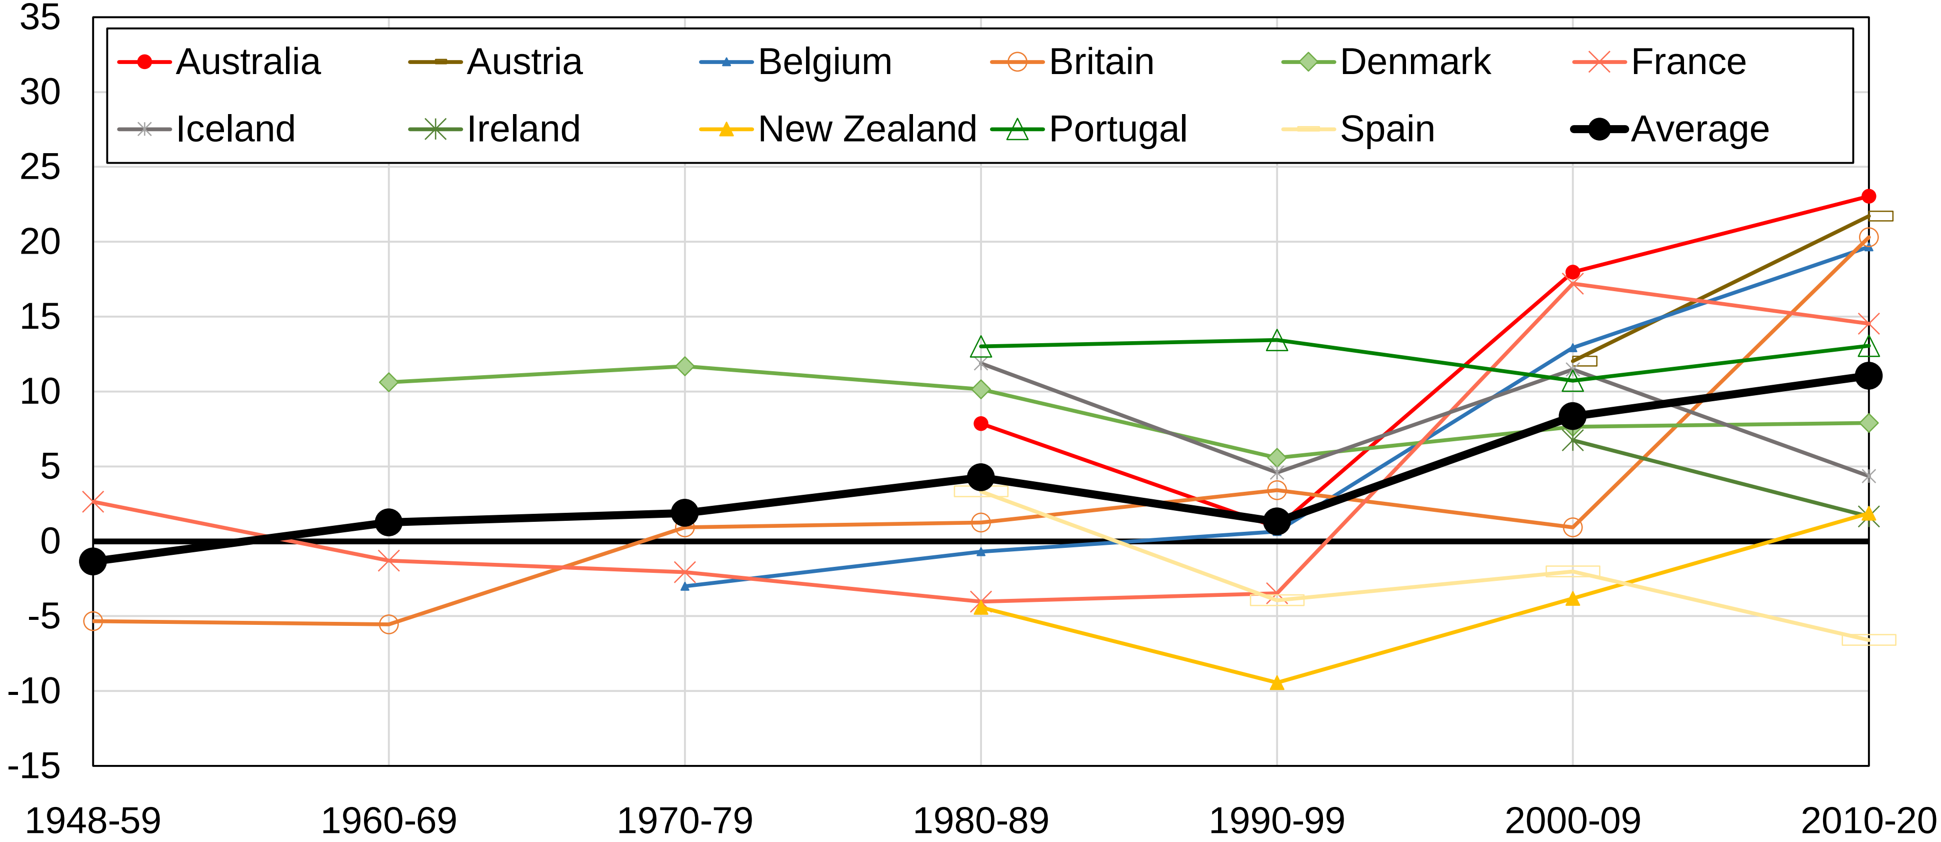

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the difference between the share of voters living in the capital city and the share of other voters voting for left-wing parties in Western democracies. Centers correspond to the Australian Capital Territory (Australia), Vienna (Austria), Brussels (Belgium), Copenhagen (Denmark), Paris (France), Reykjavík (Iceland), Dublin (Ireland), Auckland and Wellington (New Zealand), Lisbon (Portugal), Madrid (Spain), and London (United Kingdom).

Figure CB8 - Vote for left-wing parties among capital cities in Western democracies, after controls

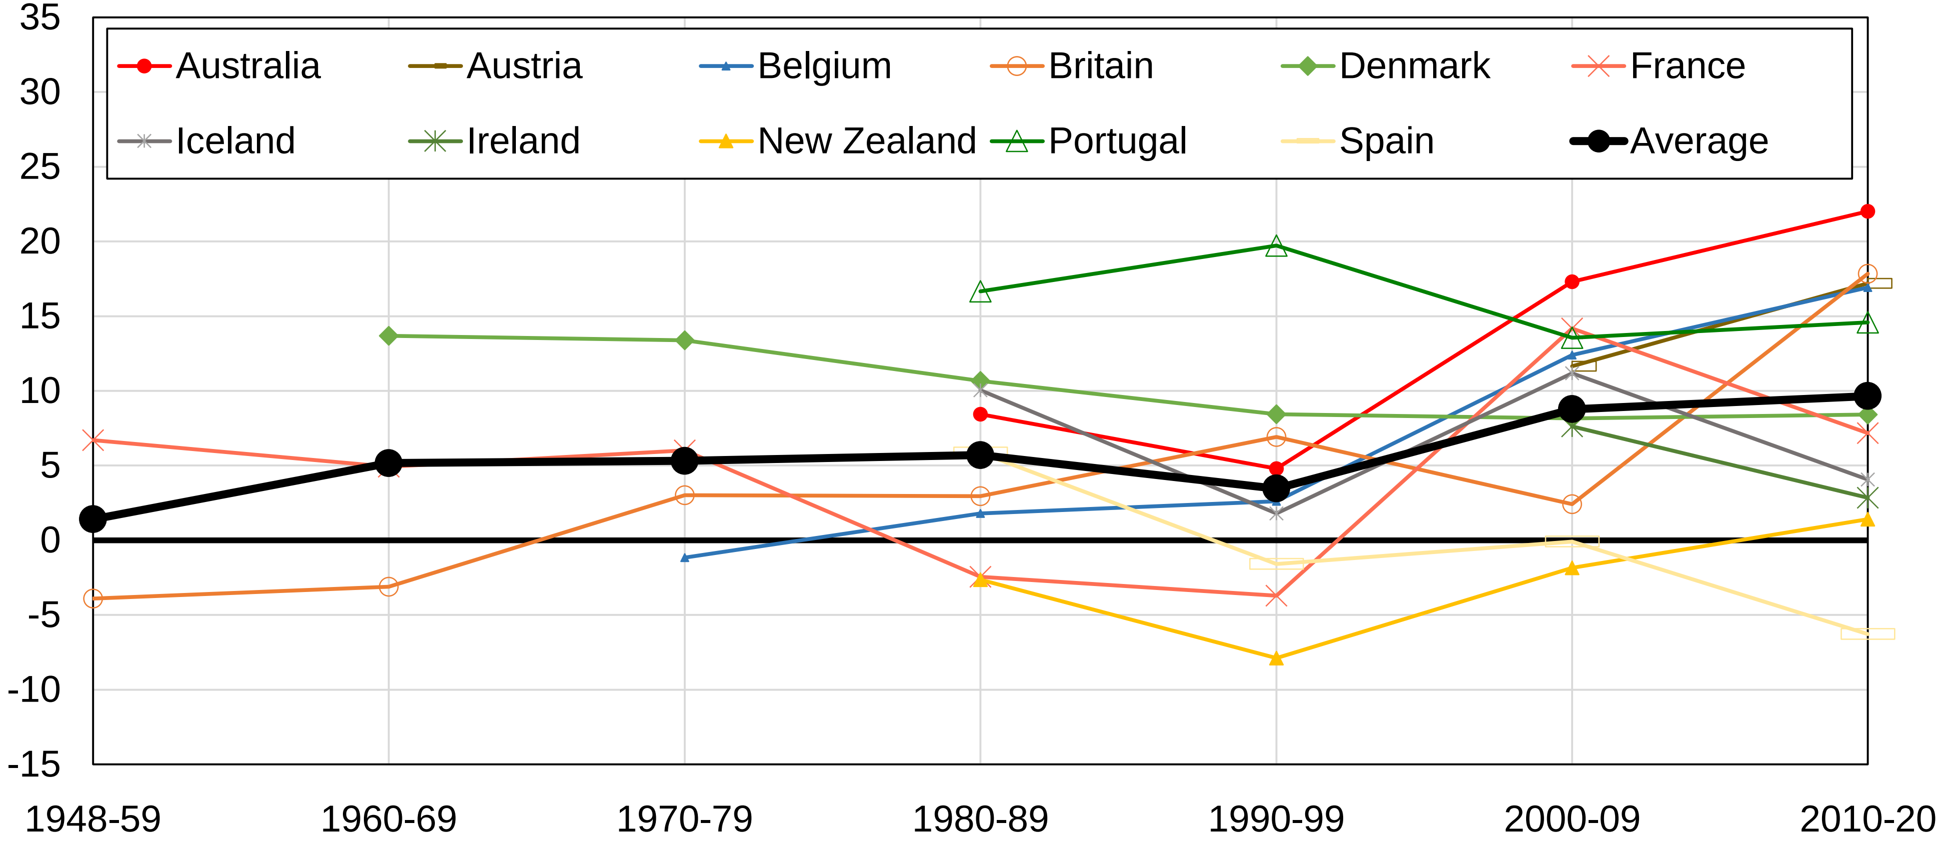

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the difference between the share of voters living in the capital city and the share of other voters voting for left-wing parties in Western democracies, after controlling for income, education, age, gender, employment status, and marital status. Centers correspond to the Australian Capital Territory (Australia), Vienna (Austria), Brussels (Belgium), Copenhagen (Denmark), Paris (France), Reykjavík (Iceland), Dublin (Ireland), Auckland and Wellington (New Zealand), Lisbon (Portugal), Madrid (Spain), and London (United Kingdom).

Figure CC1 - The religious divide

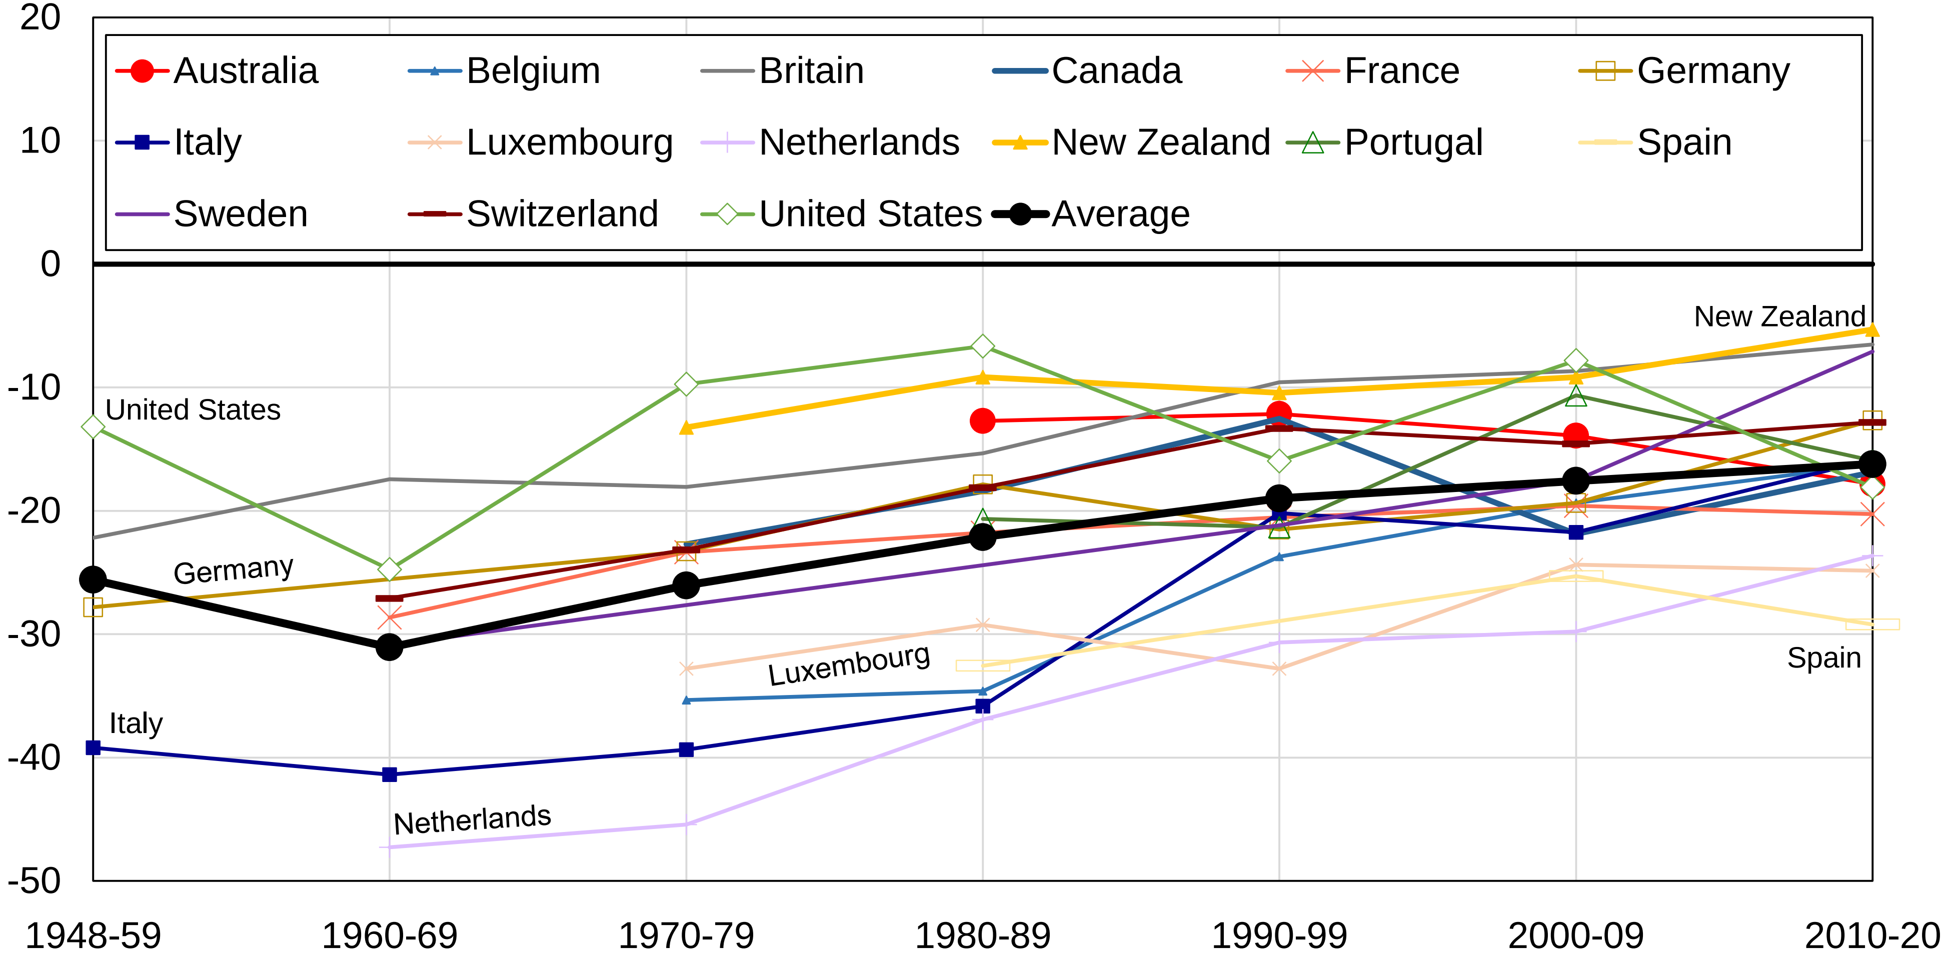

Source: authors' computations using the World Political Cleavages and Inequality Database.

Note: the figure displays the difference between the share of Catholics (or Catholics and Protestants in mixed countries) declaring going to church at least once a year and the share of other voters voting for social democratic / socialist / communist / green parties. In all countries, religious voters have remained significantly less likely to vote for these parties than other voters.

Figure CC2 - Vote for left-wing parties by religion in Western democracies, 1970s

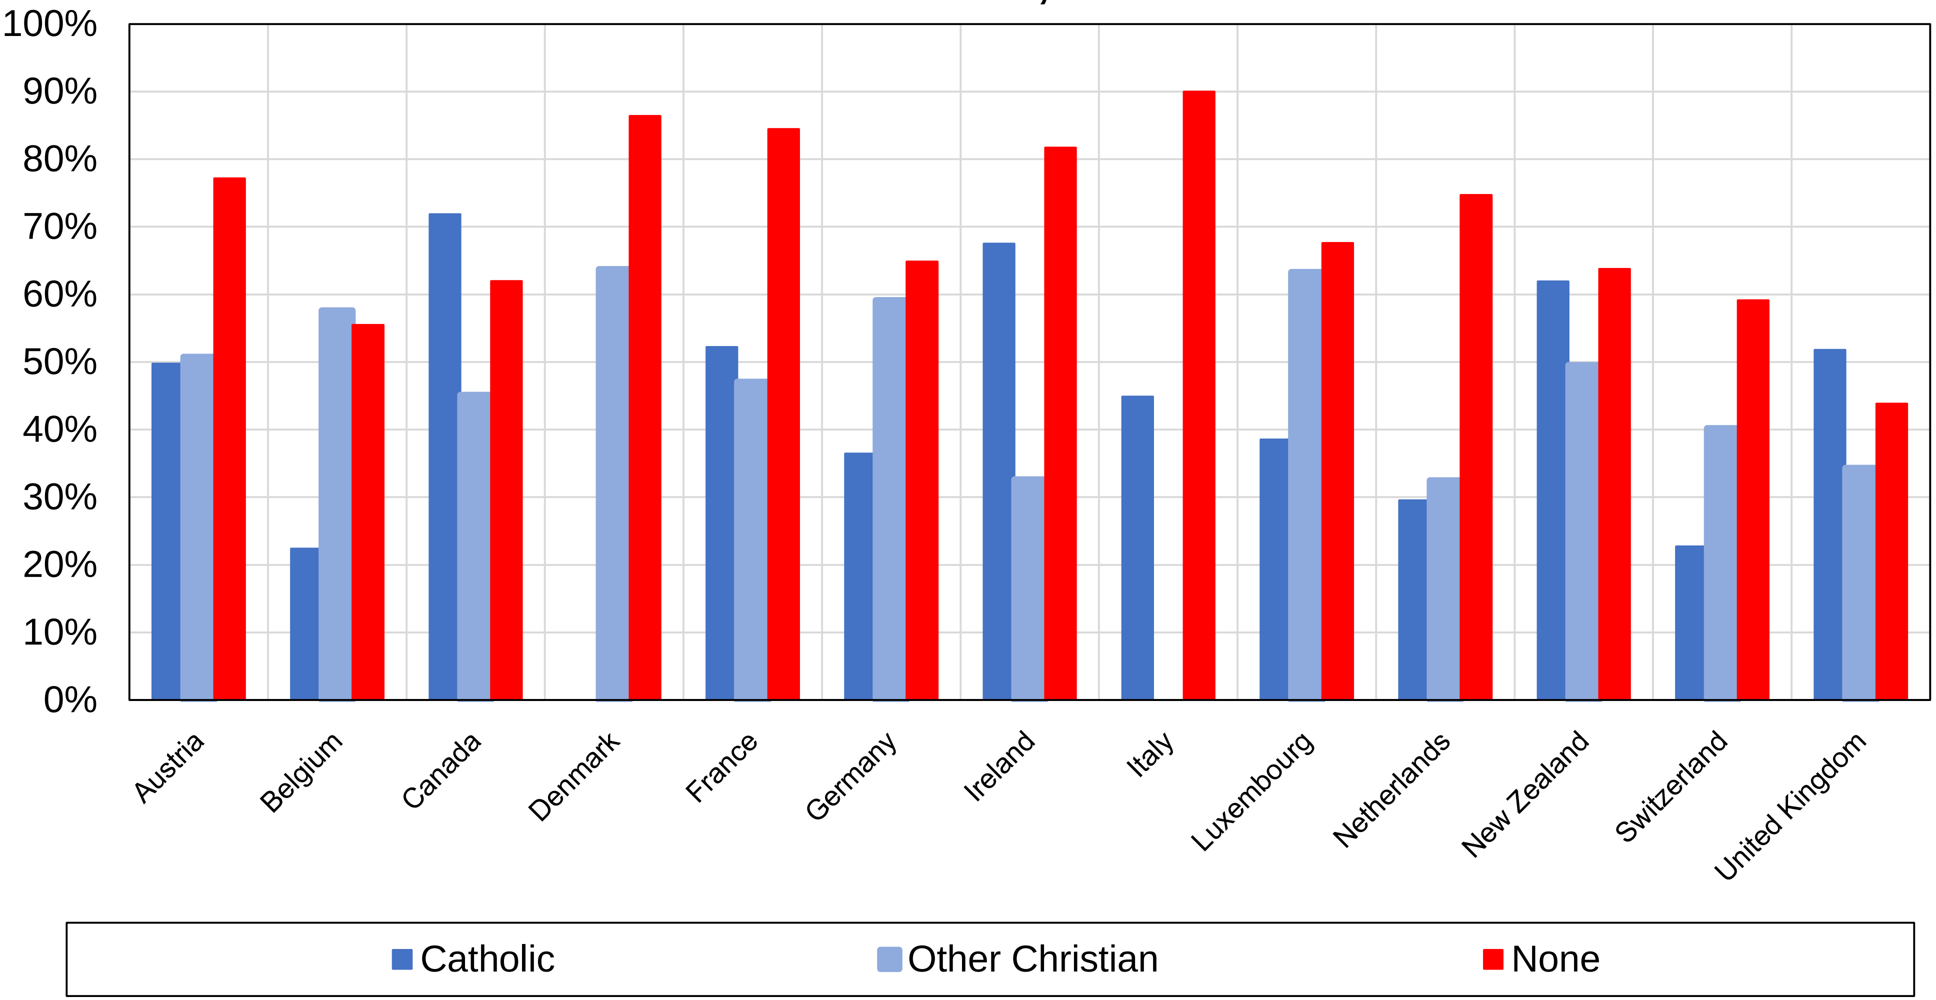

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by left-wing parties by religion in the 1970s in Western democracies.

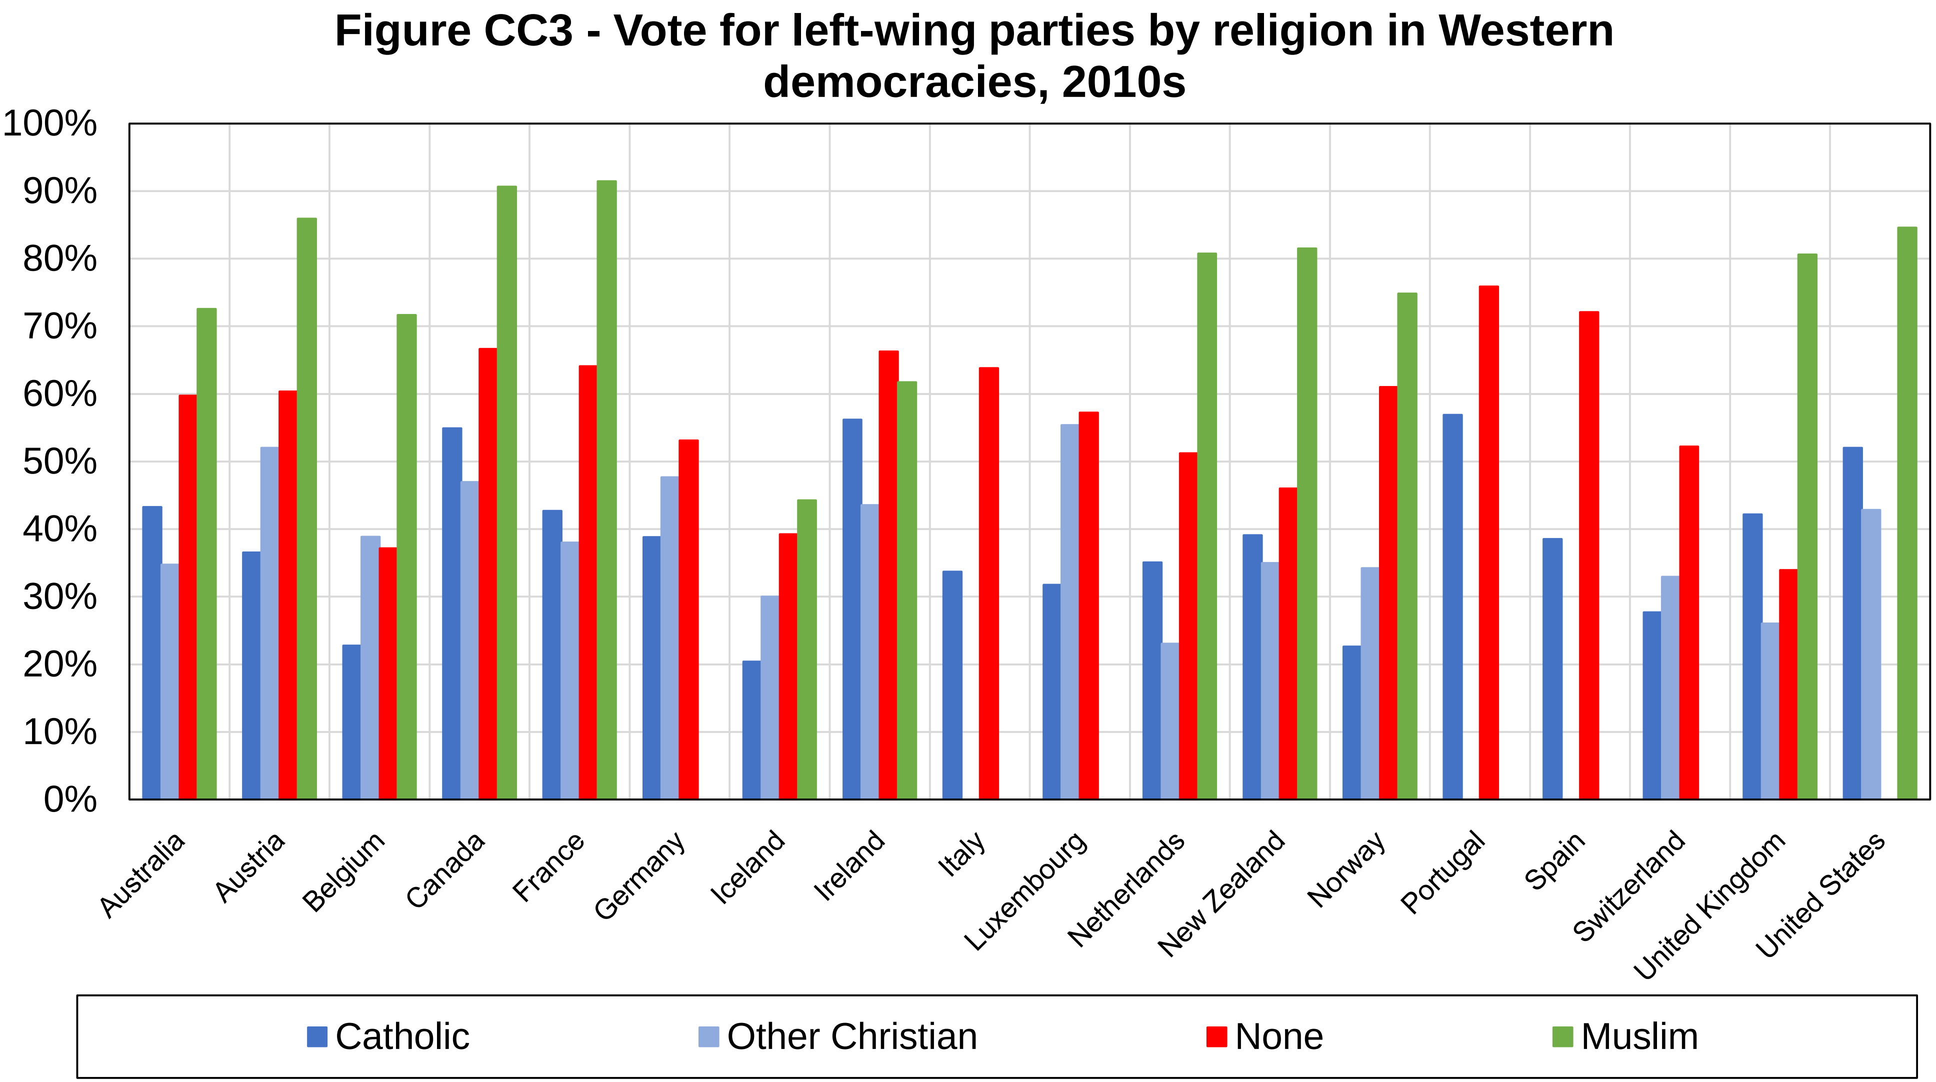

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by left-wing parties by religion in the 2010s in Western democracies.

Figure CC4 - Vote for left-wing parties among voters with no religion in Western democracies

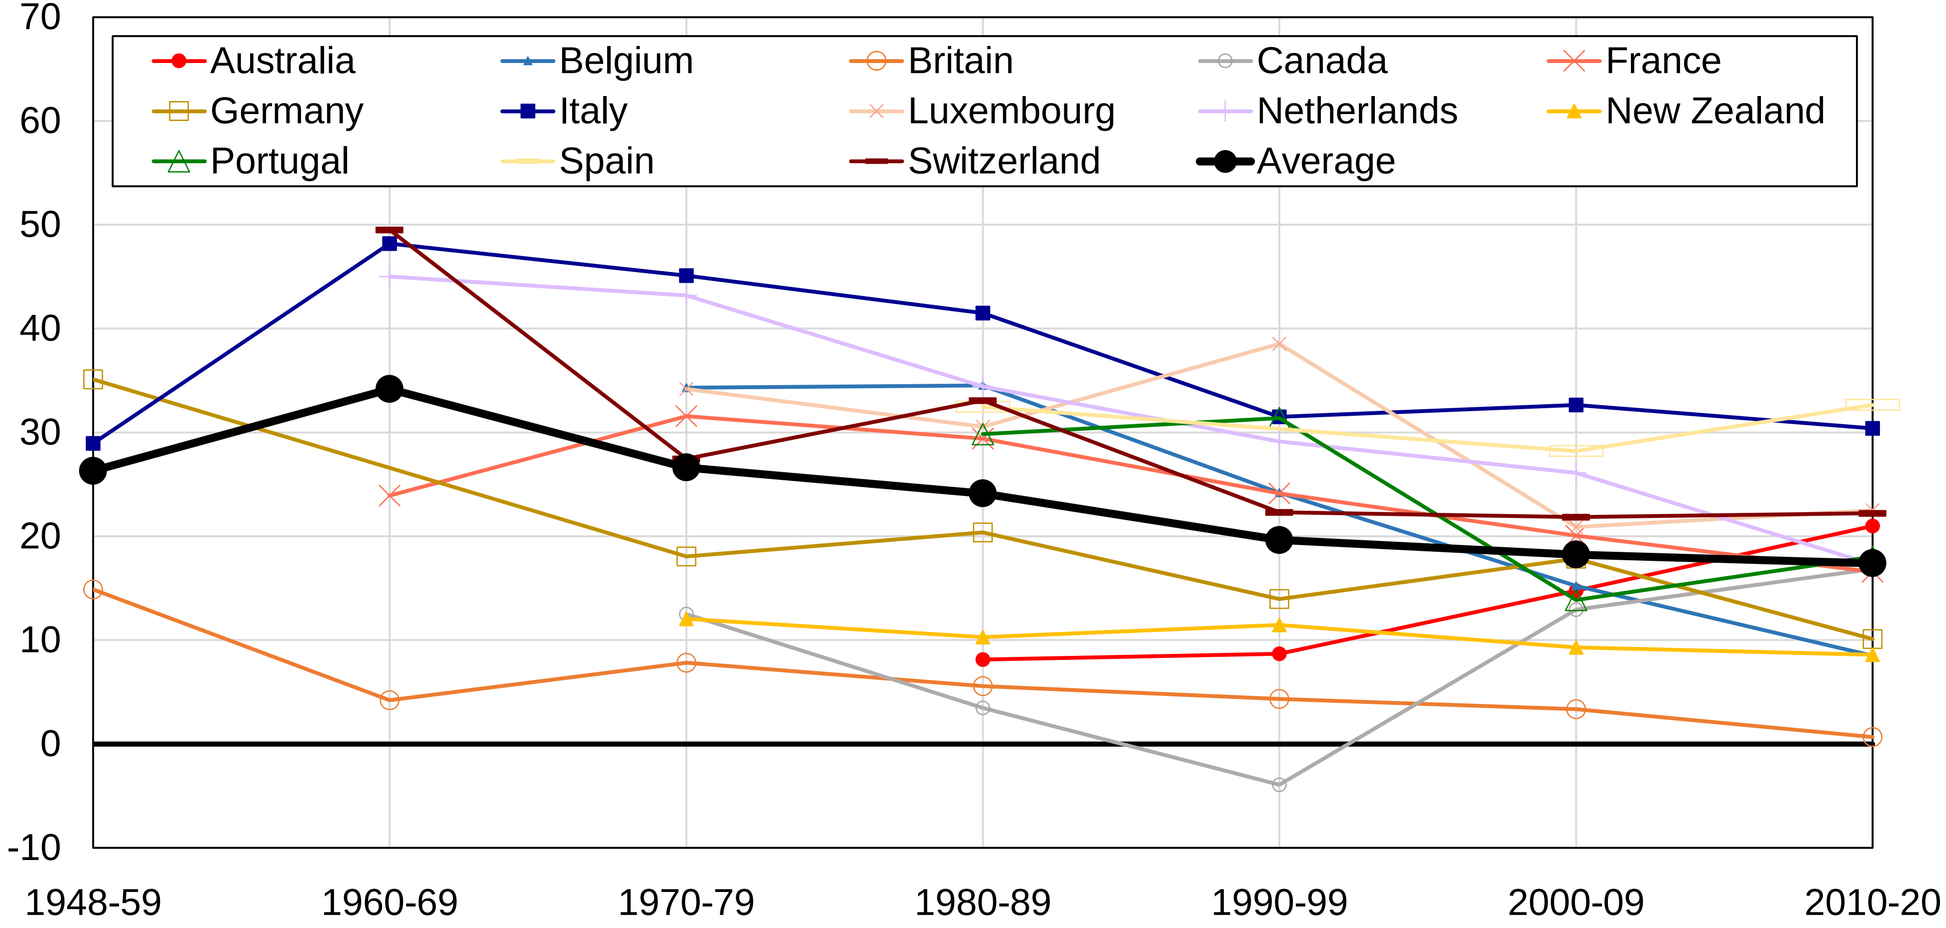

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of voters belonging to no religion and the share of other voters voting for left-wing parties in Western democracies. Non-religious voters have remained significantly more left-wing than the rest of the electorate since the 1950s.

Figure CC5 - Vote for Green parties by religion, 2010s

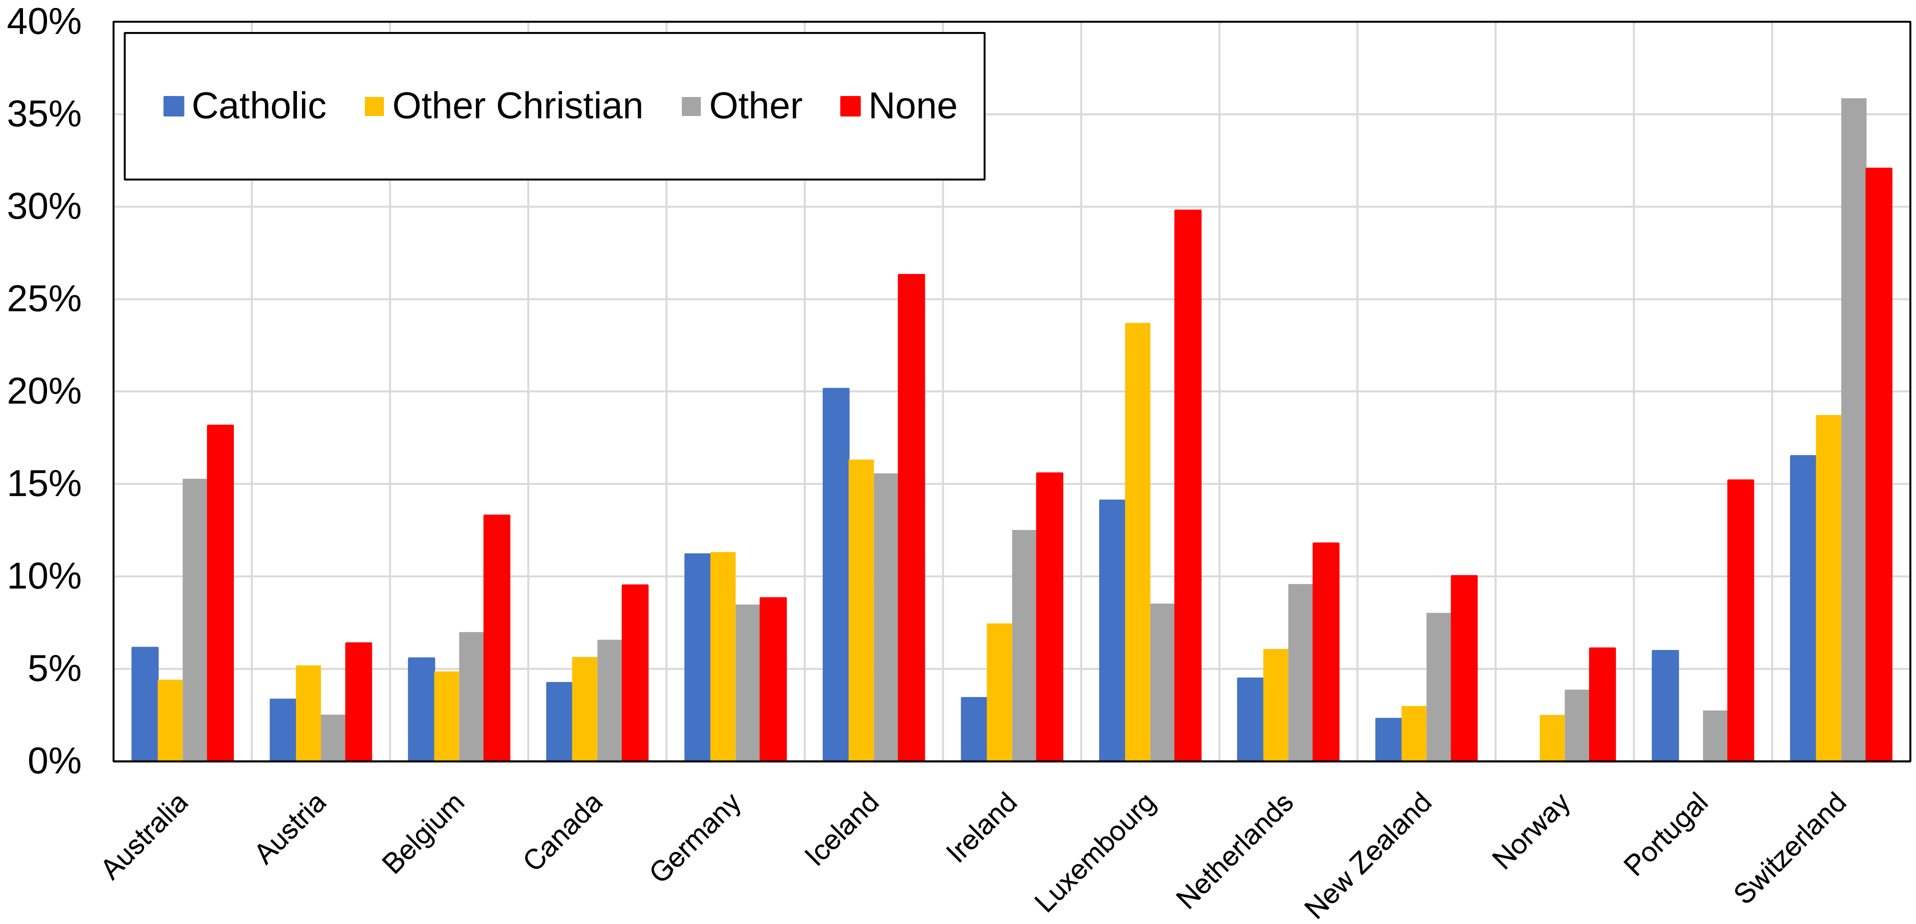

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by Green parties by religious affiliation.

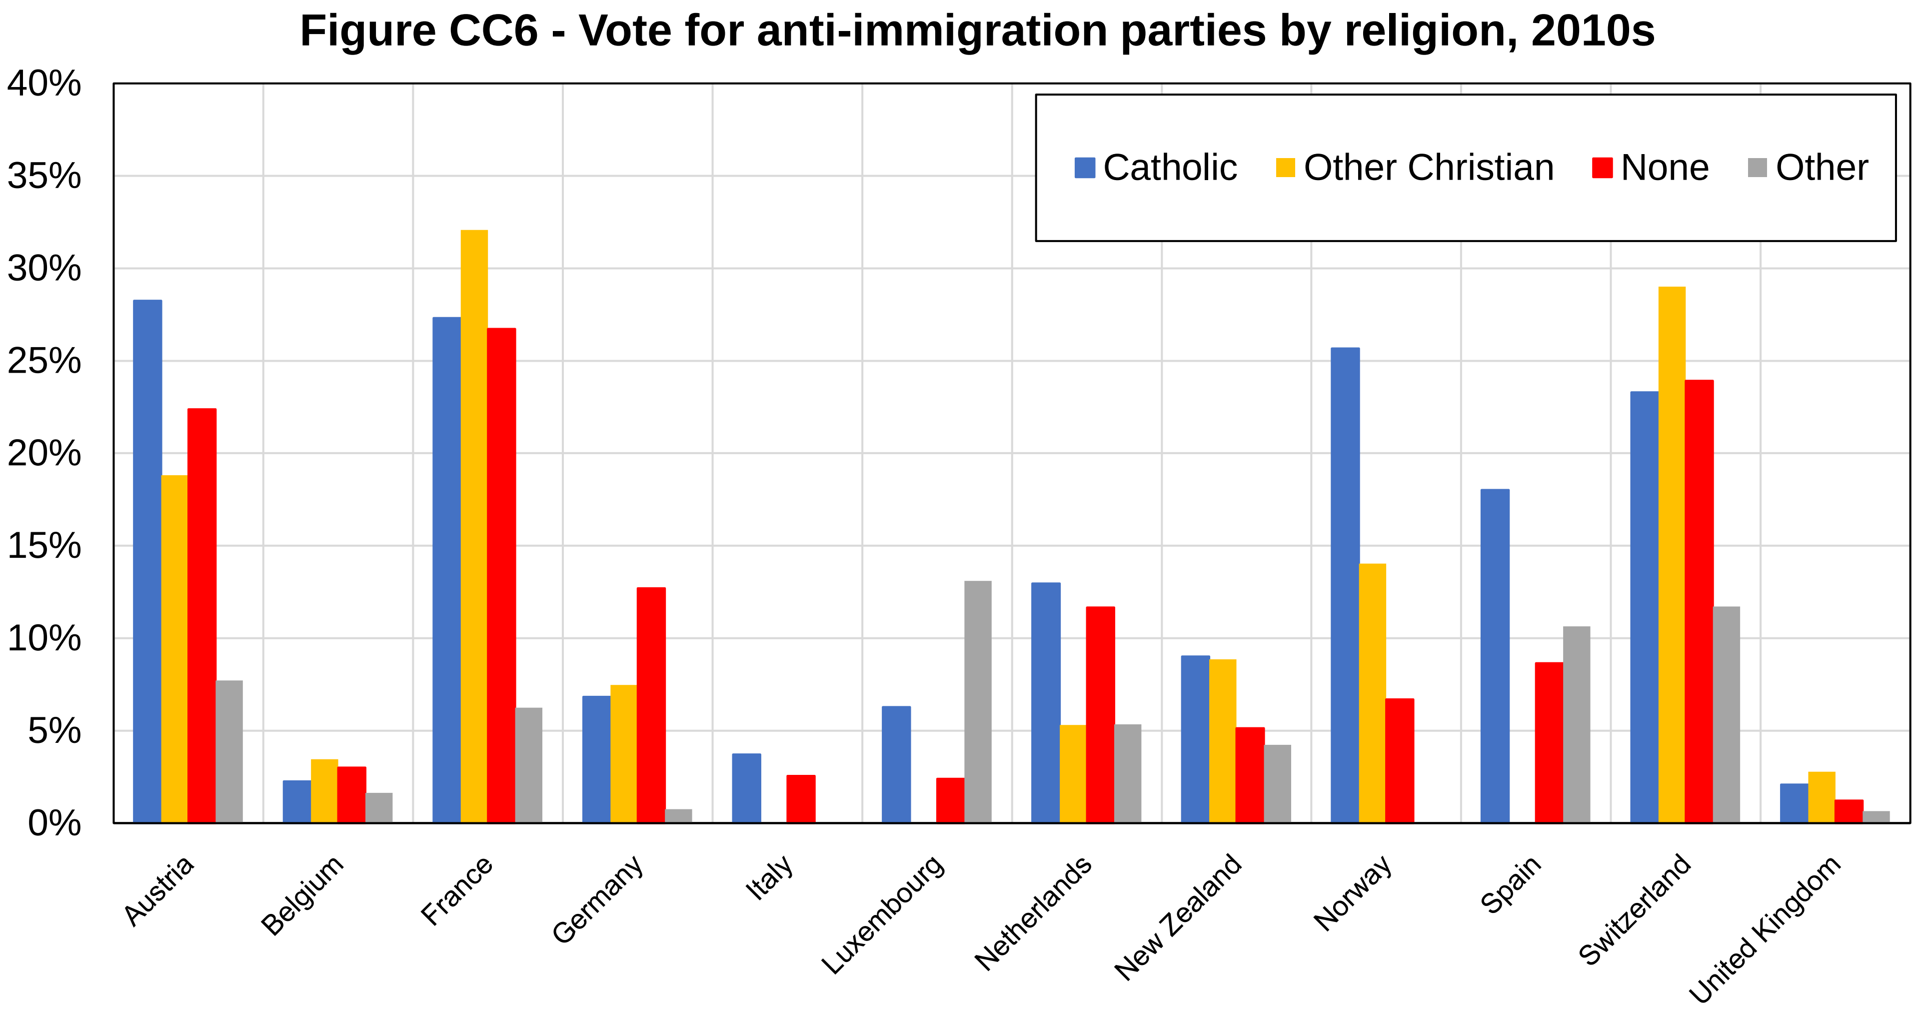

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by anti-immigration parties by religious affiliation.

Figure CD1 - The reversal of the gender cleavage in Western democracies (before and after controls)

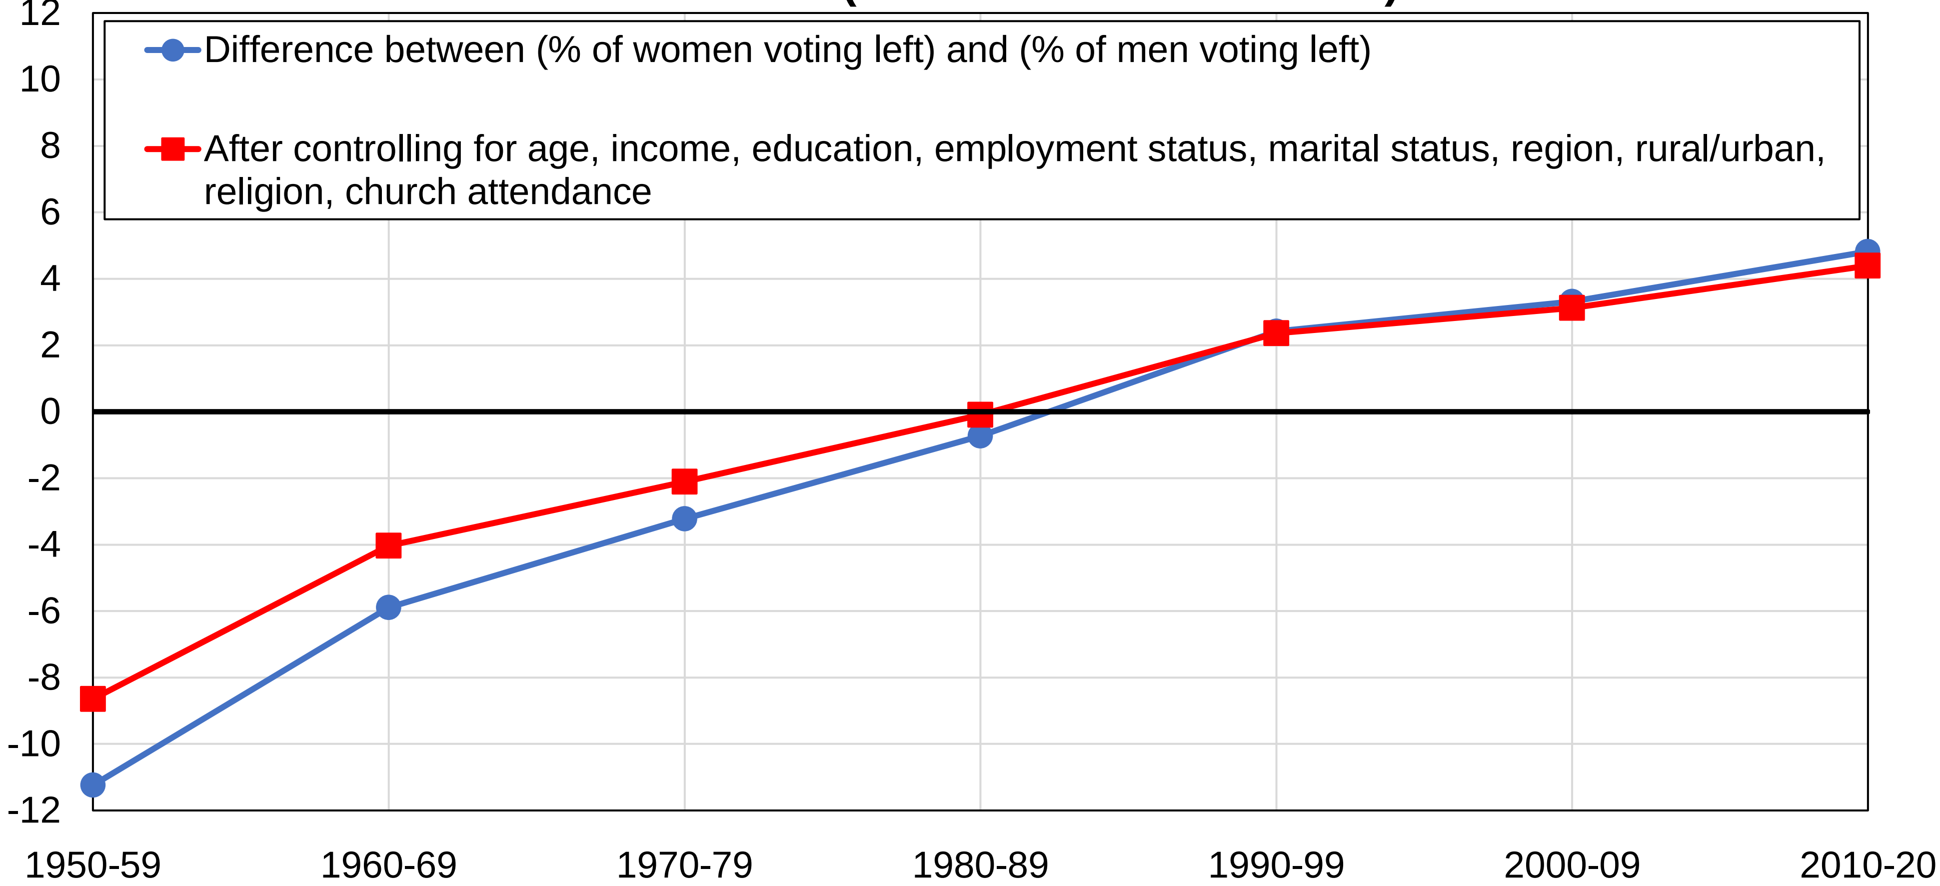

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure displays the difference between the share of women and the share of men voting for left-wing (social democratic, socialist, communist, and green) parties in Western democracies, before and after controlling for other covariates (for country-years in which these variables are available). Women have gradually shifted from being significantly more right-wing to being significantly more left-wing than men, both before and after controls. Average over all Western democracies.

Figure CD2 - The reversal of the gender cleavage

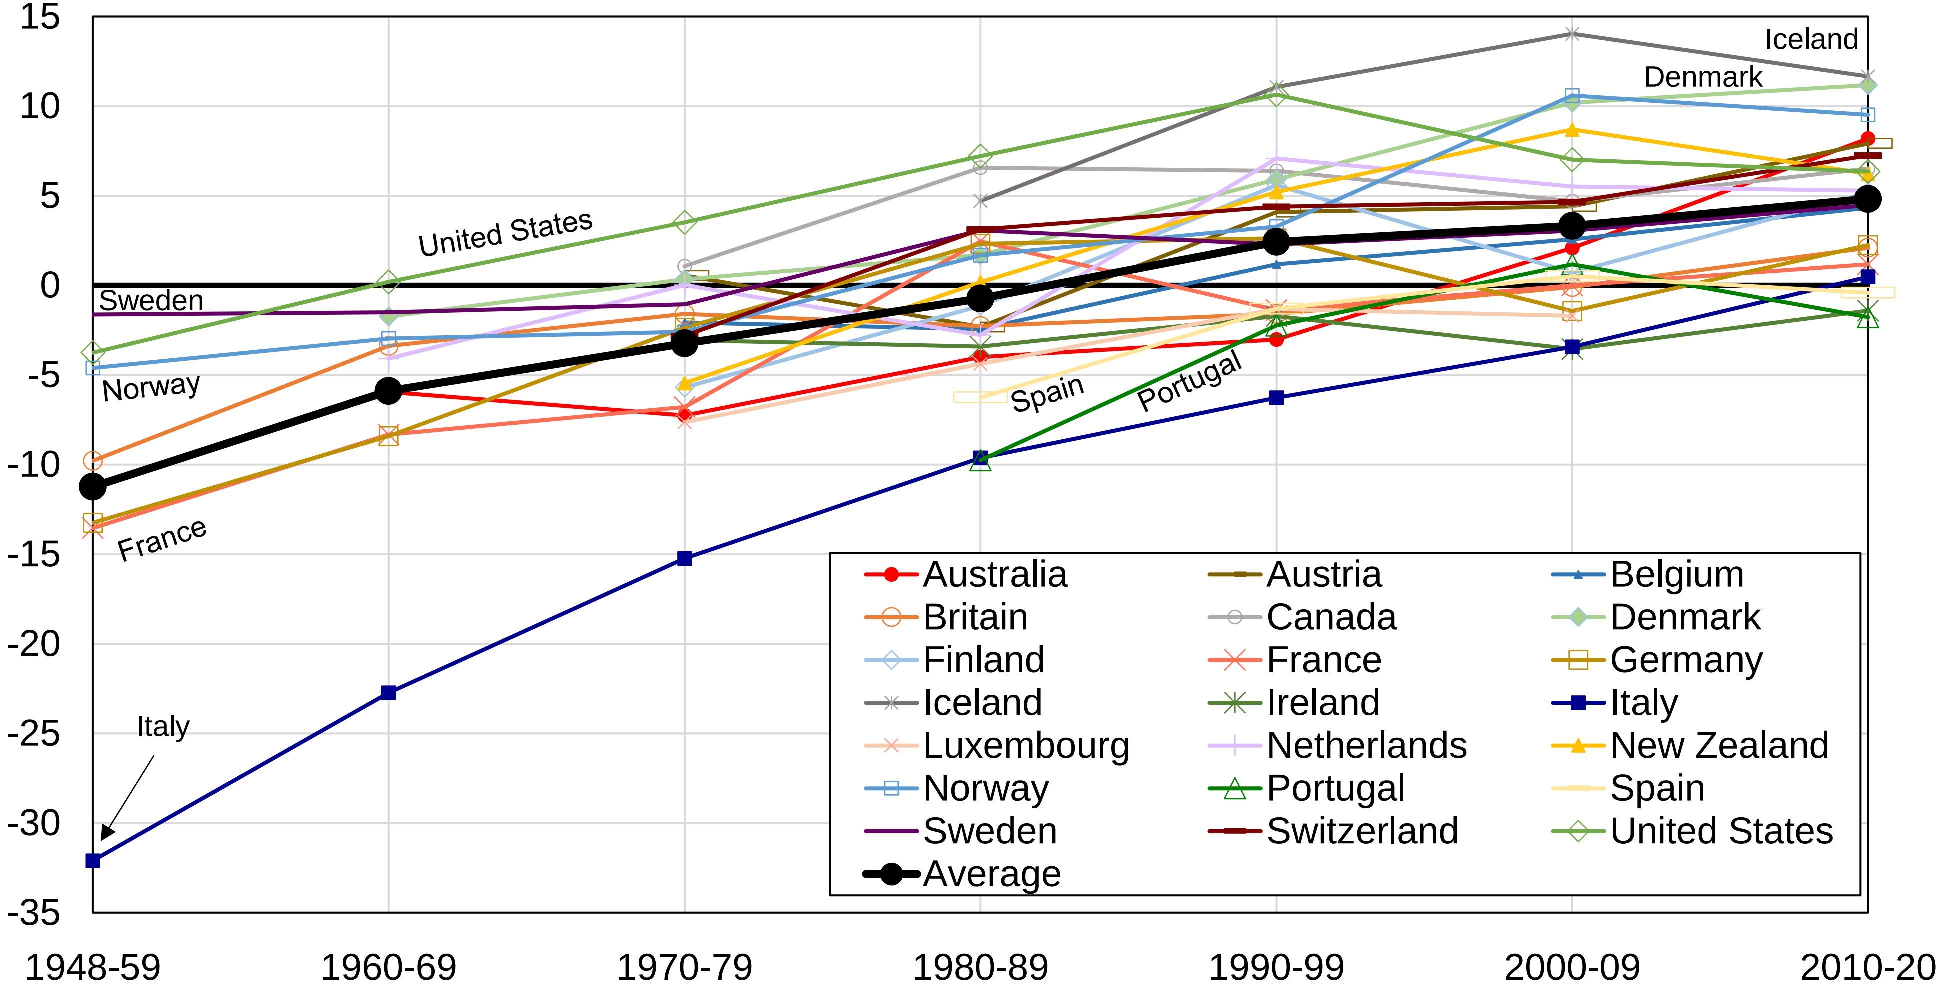

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure displays the difference between the share of women and the share of men voting for social democratic / socialist / communist / green parties in Western democracies. In the majority of countries, women have gradually shifted from being significantly more conservative than men in the 1950s-1960s to being significantly more left-wing in the 2000s-2010s.

Figure CD3 - Vote for left-wing parties among women in Western democracies (after controlling for religion)

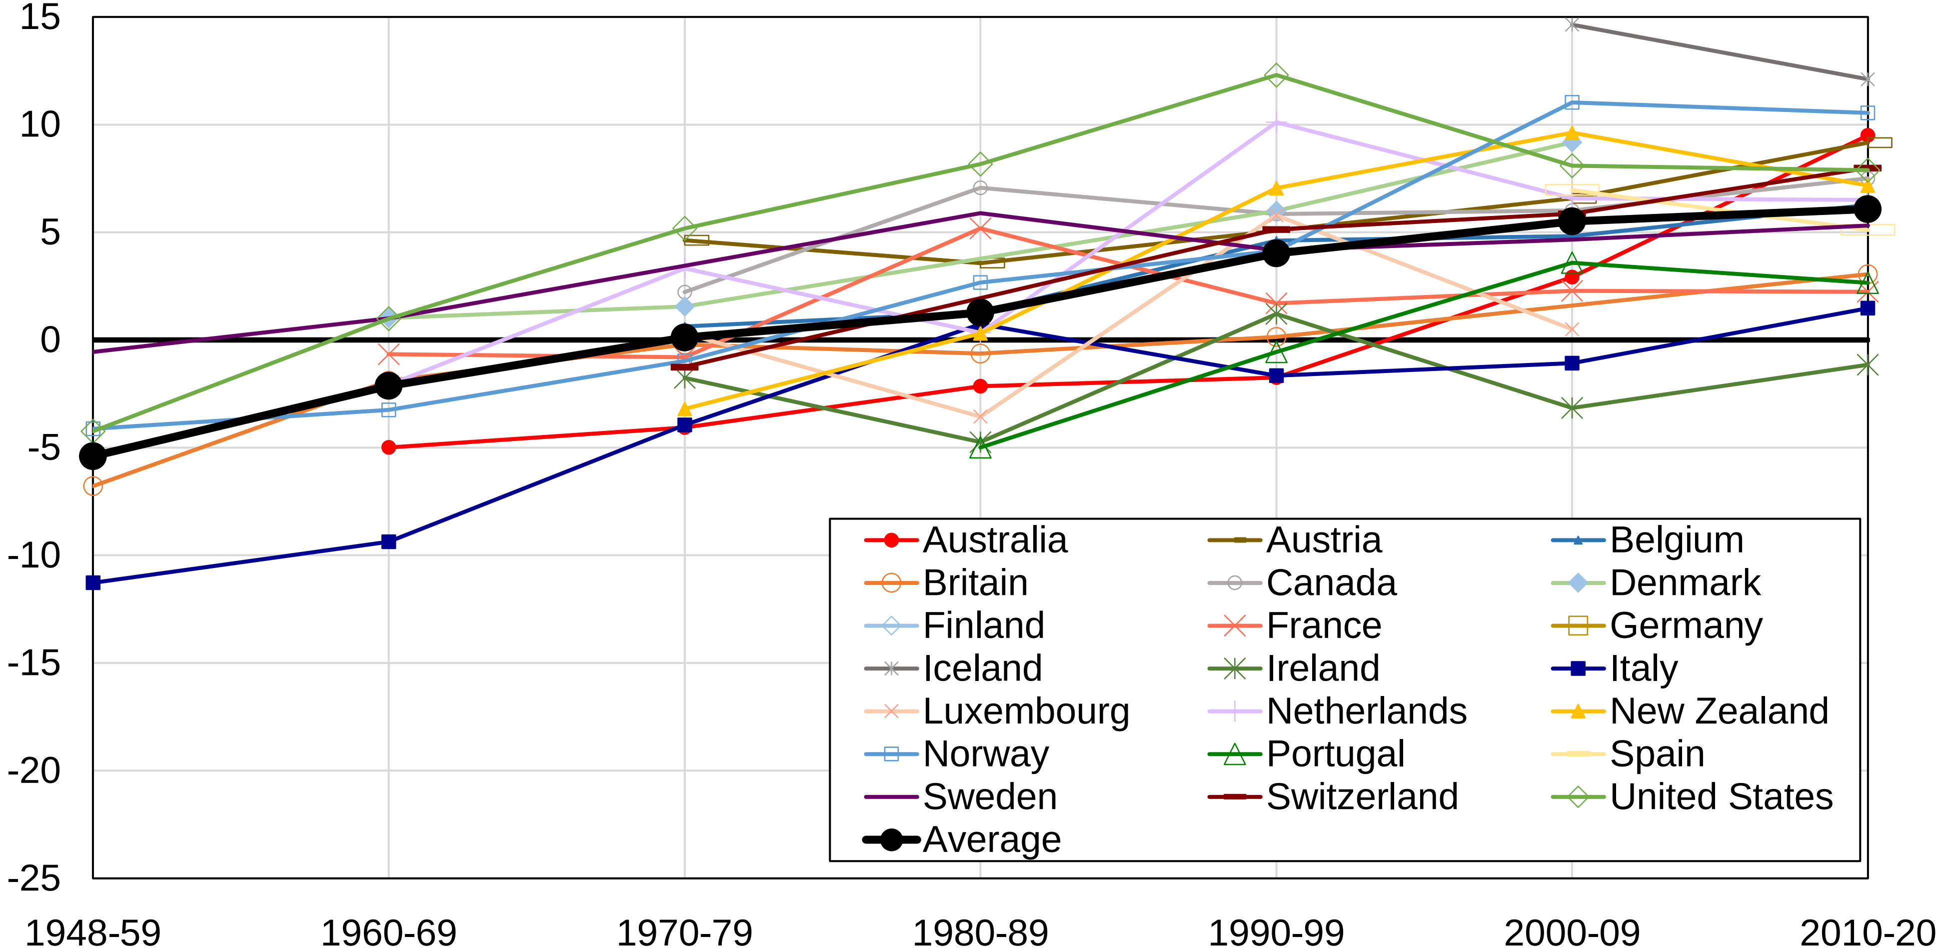

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure displays the difference between the share of women and the share of men voting for left-wing (socialist, social democratic, communist, and green) parties in Western democracies, after controlling for religion and church attendance. In the majority of countries, women have gradually shifted from being significantly more right-wing to being significantly more left-wing than men.

Figure CD4 - Gender cleavages and sectoral specialization in Western democracies

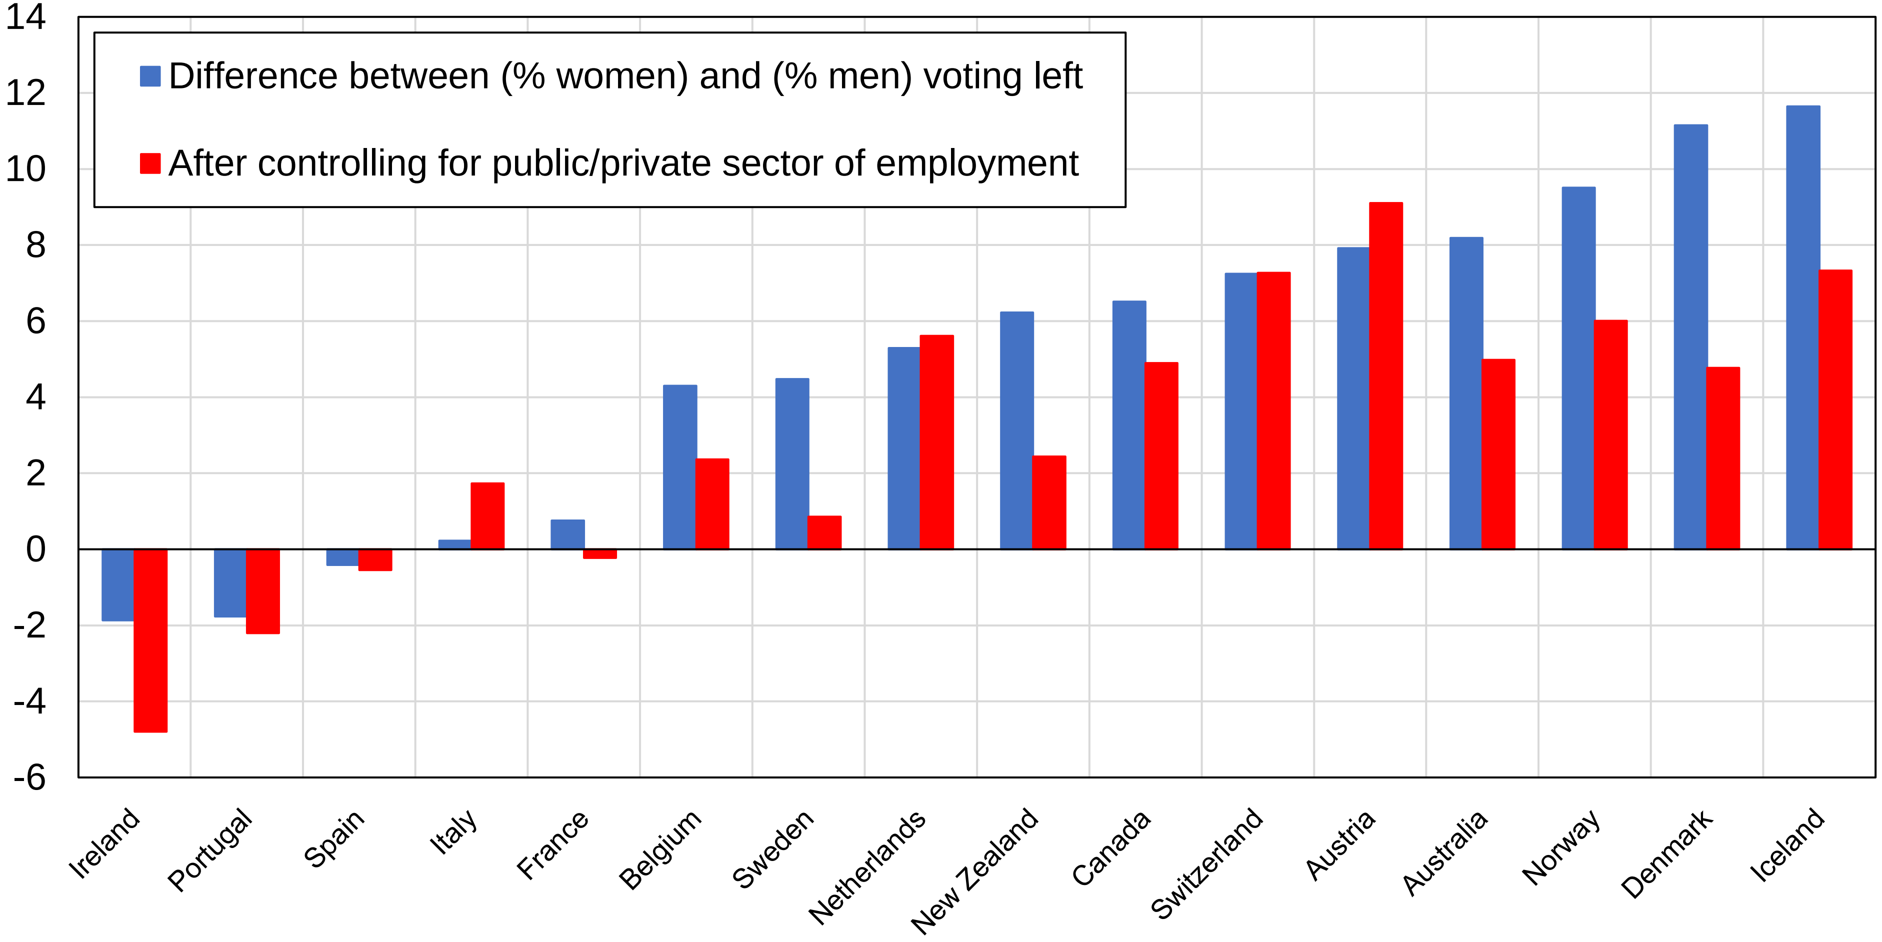

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the difference between the share of women and the share of men voting for left-wing parties in Western democracies in the last election available, before and after controlling for occupation (employment status + private/public sector of employment).

Figure CD5 - Vote for green parties by gender in Western democracies

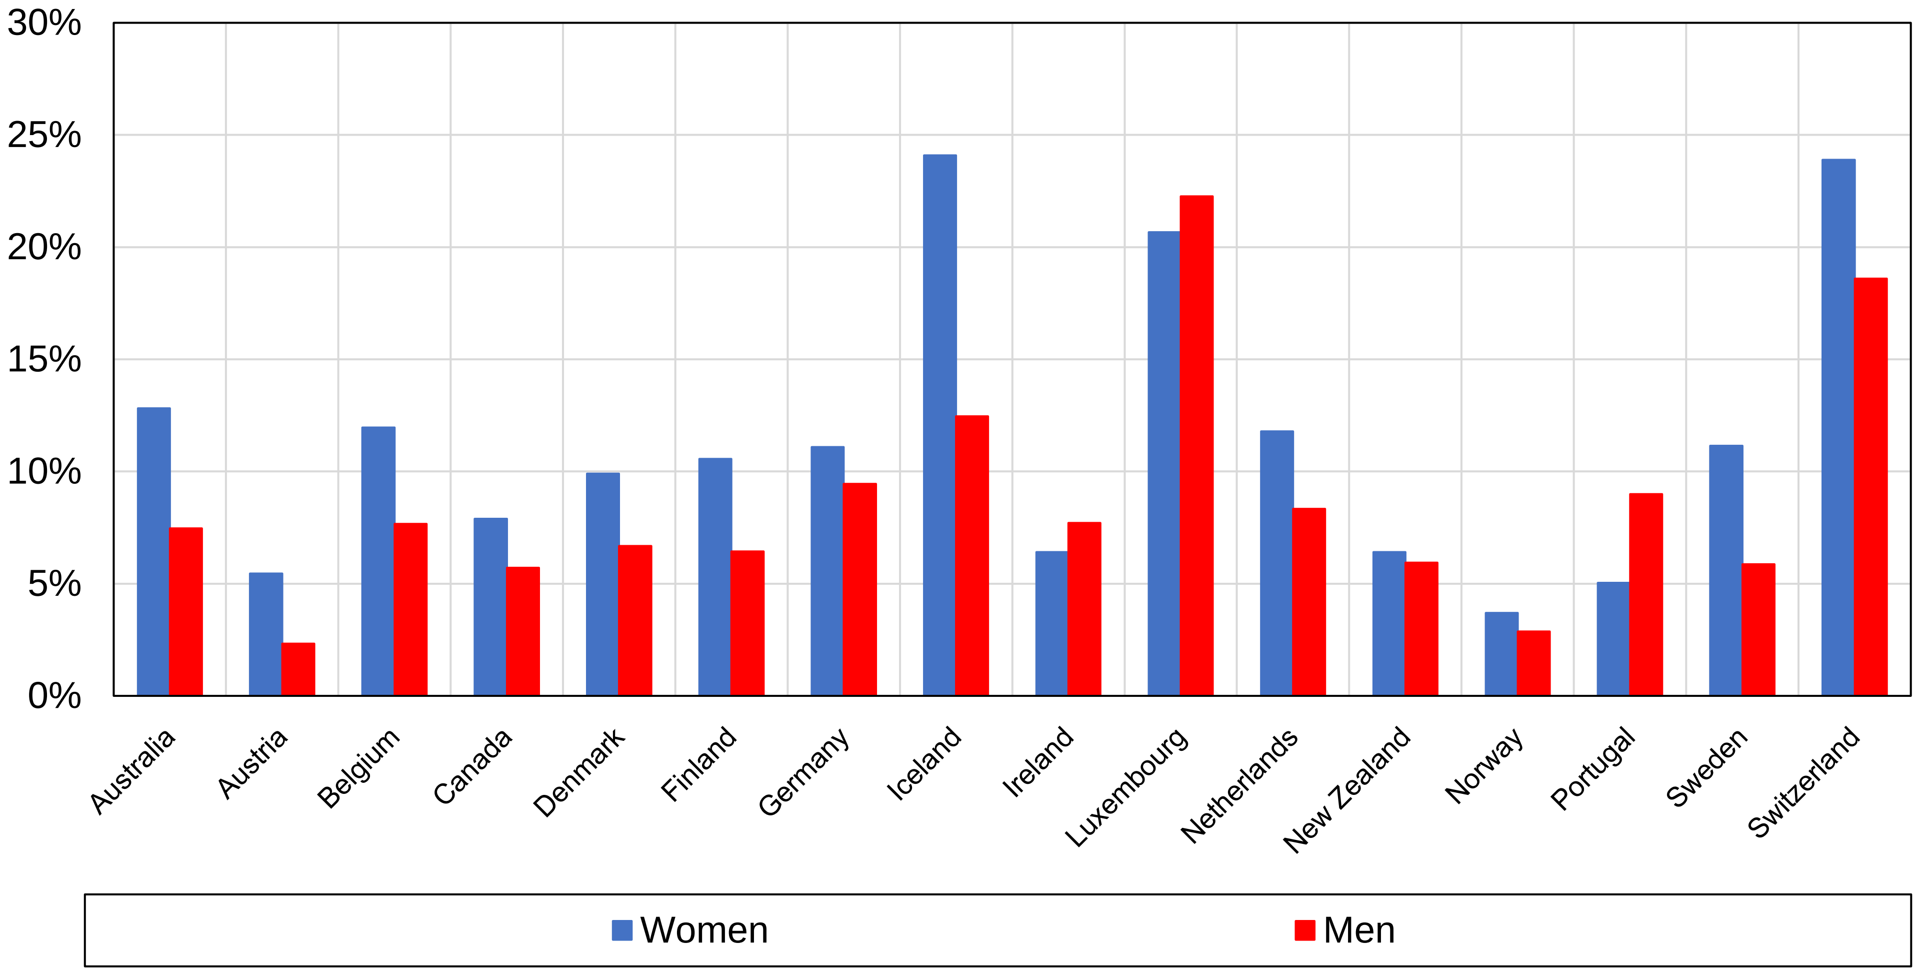

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by Green parties by gender in Western democracies in the last election available.

Figure CD6 - Vote for anti-immigration parties by gender in Western democracies

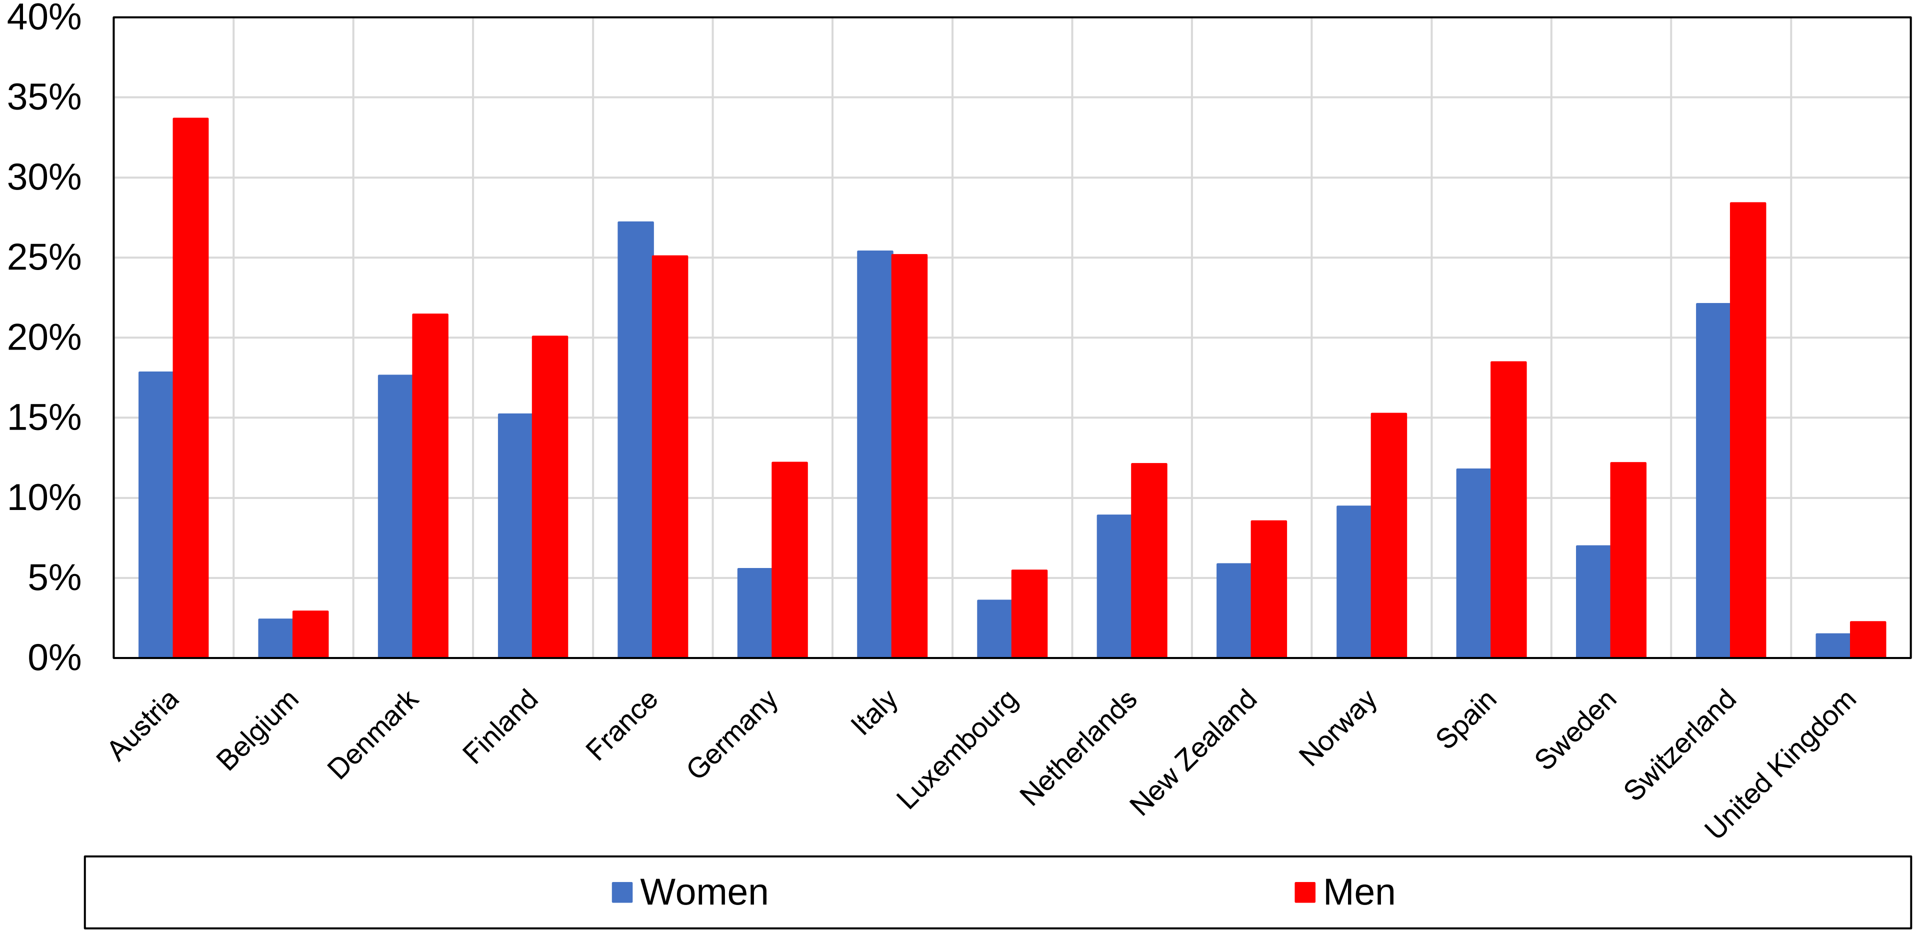

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by anti-immigration parties by gender in Western democracies in the last election available.

Figure CE1 - The nativist cleavage
The native-immigrant cleavage

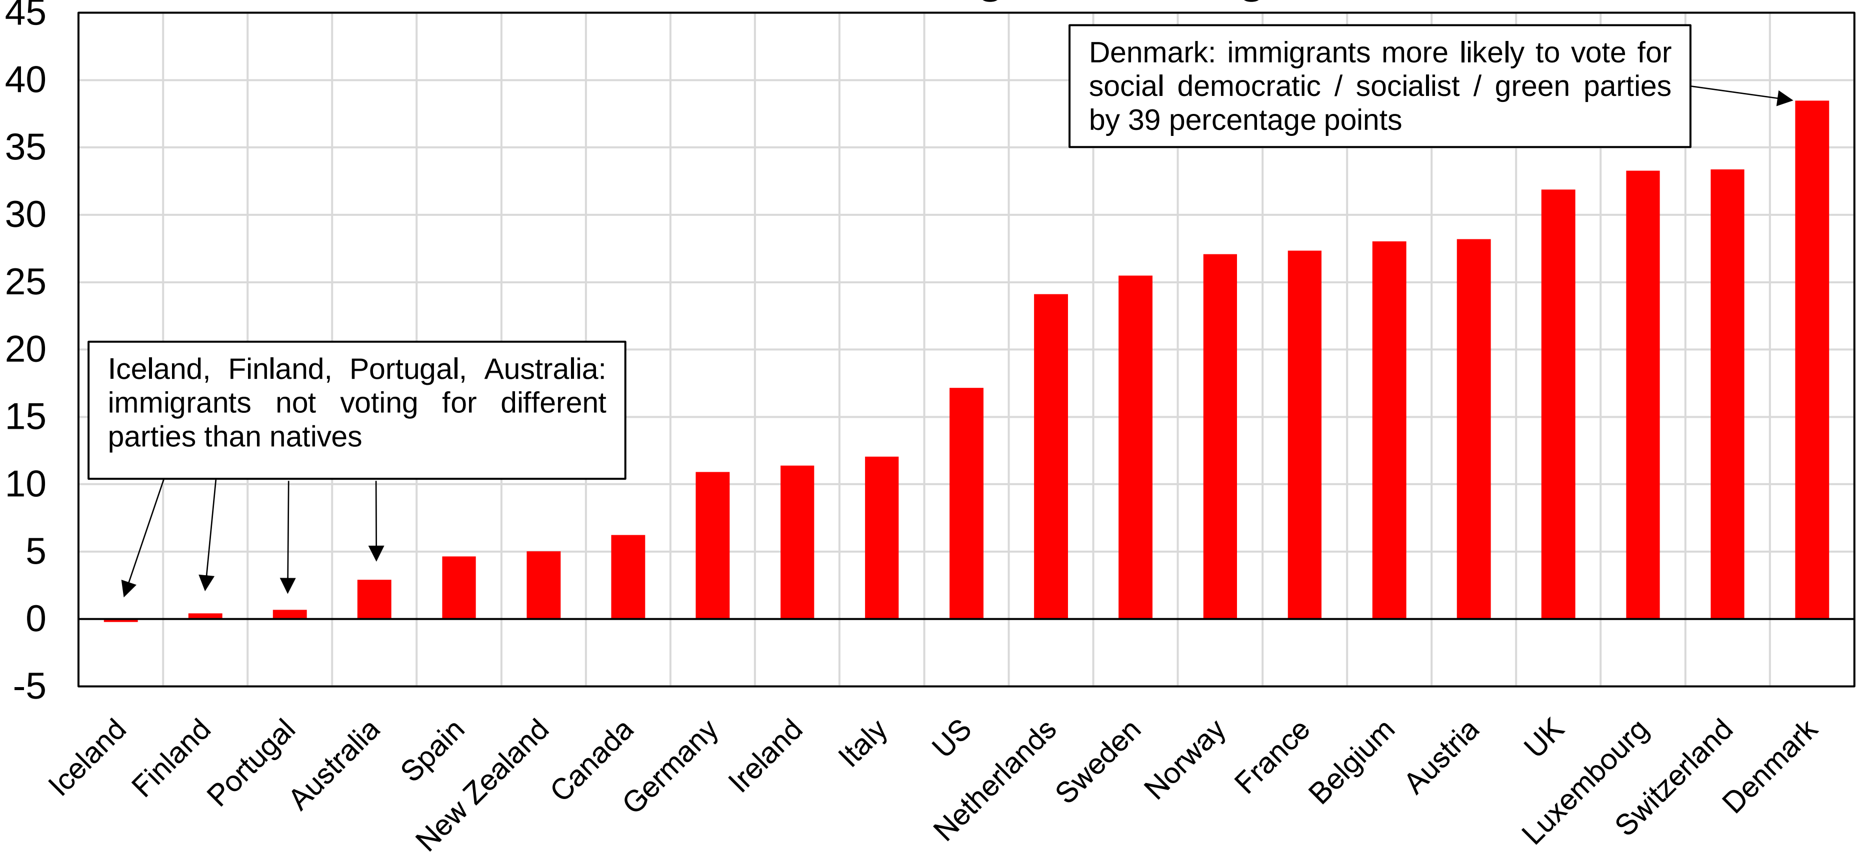

Source: authors' computations using the World Political Cleavages and Inequality Database and the European Social Survey for Denmark, Finland, Germany, Italy, Norway, Spain, Sweden, Switzerland, and the United Kingdom.
Note: the figure represents the difference between the share of voters born in non-Western countries (all countries excluding Europe, Australia, New Zealand, Canada, and the United States) and the share of natives (voters born in the country considered) voting for social democratic / socialist / communist / green parties over the 2010-2020 period. In nearly all Western countries, immigrants are much more likely to vote for these parties than natives. US and Iceland figures include voters born in Western countries given lack of data on exact country of origin. Excludes Fianna Fáil in Ireland.

Figure CE2 - The nativist cleavage
The Muslim vote

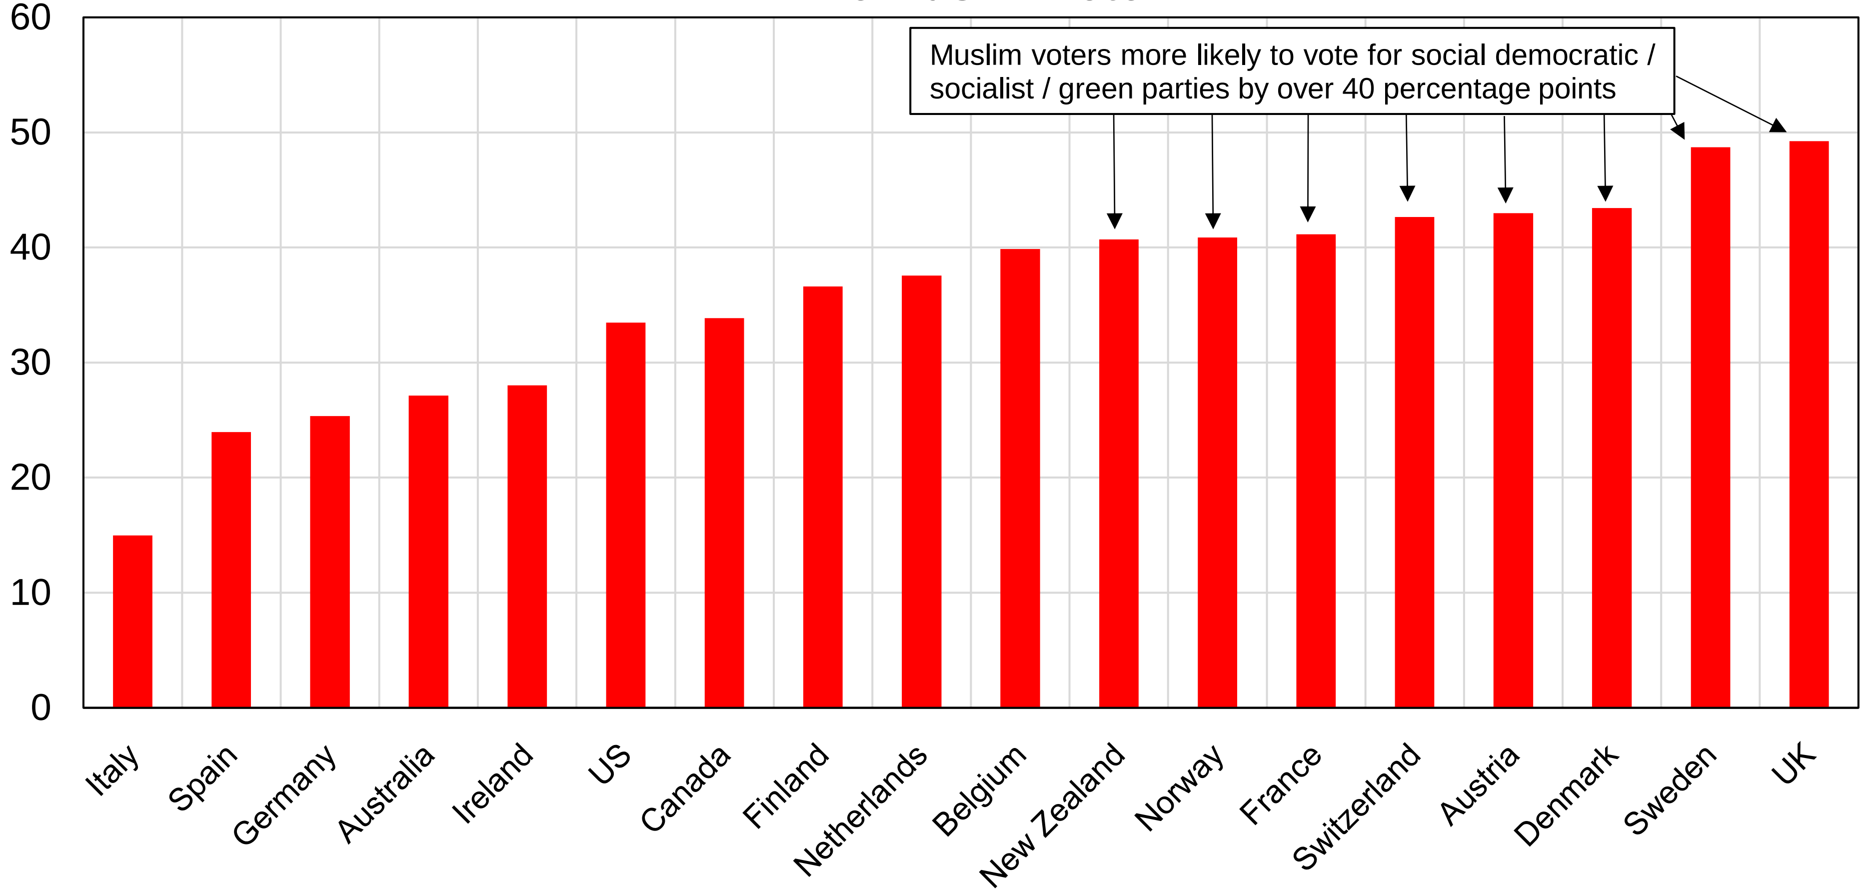

Source: authors' computations using the World Political Cleavages and Inequality Database and the European Social Survey for Denmark, Finland, Germany, Italy, Norway, Spain, Sweden, Switzerland, and the United Kingdom.
Note: the figure represents the difference between the share of Muslim voters and the share of non-Muslims voting for social democratic / socialist / communist / green parties over the 2010-2020 period. In all Western countries, Muslims are substantially more likely to vote for these parties than non-Muslims. This cleavage is stronger in countries with strong far-right parties (e.g. Sweden, Denmark, Austria, Switzerland, France). Excludes Fianna Fáil in Ireland.

Figure CE3 - Vote for left-wing parties by country of birth in Western democracies, 2010s

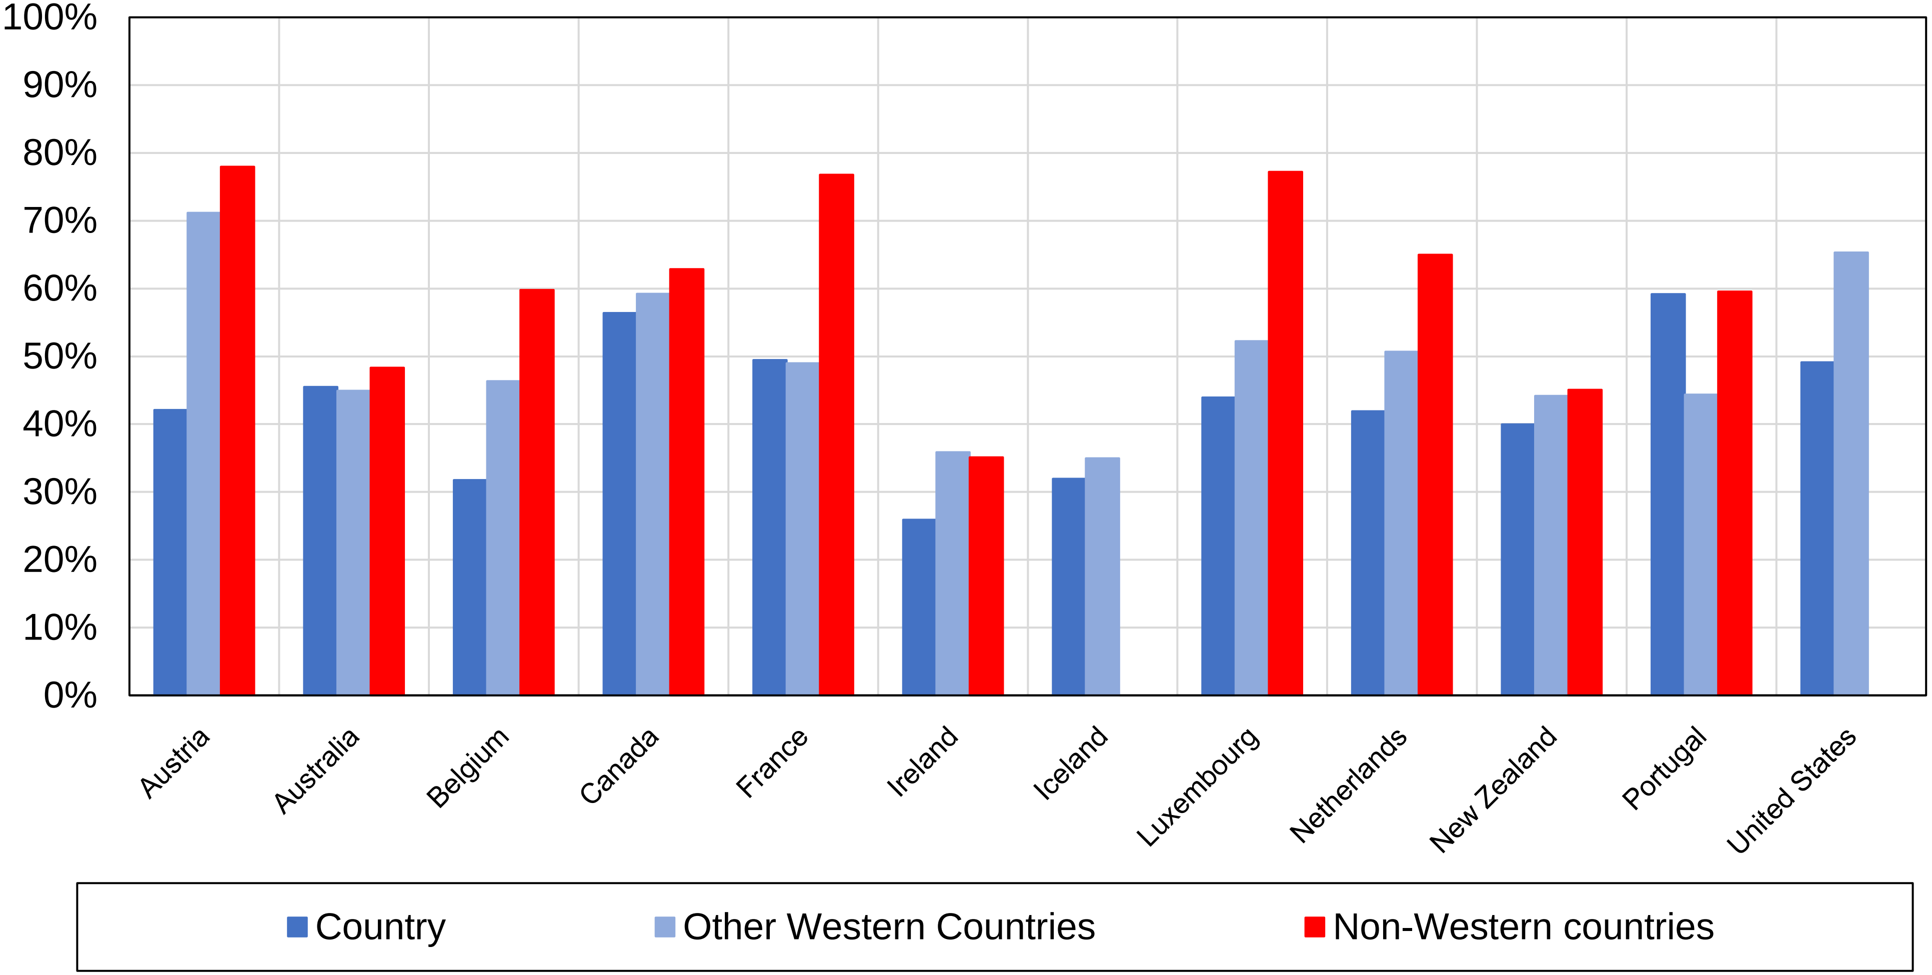

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by left-wing parties by country of birth in Western democracies in the 2010s. Excludes Fianna Fáil in Ireland. Covers 2007 and 2012 elections in France (no data in 2017).

Figure CF1 - The decline of self-perceived class cleavages in Western democracies (before controls)

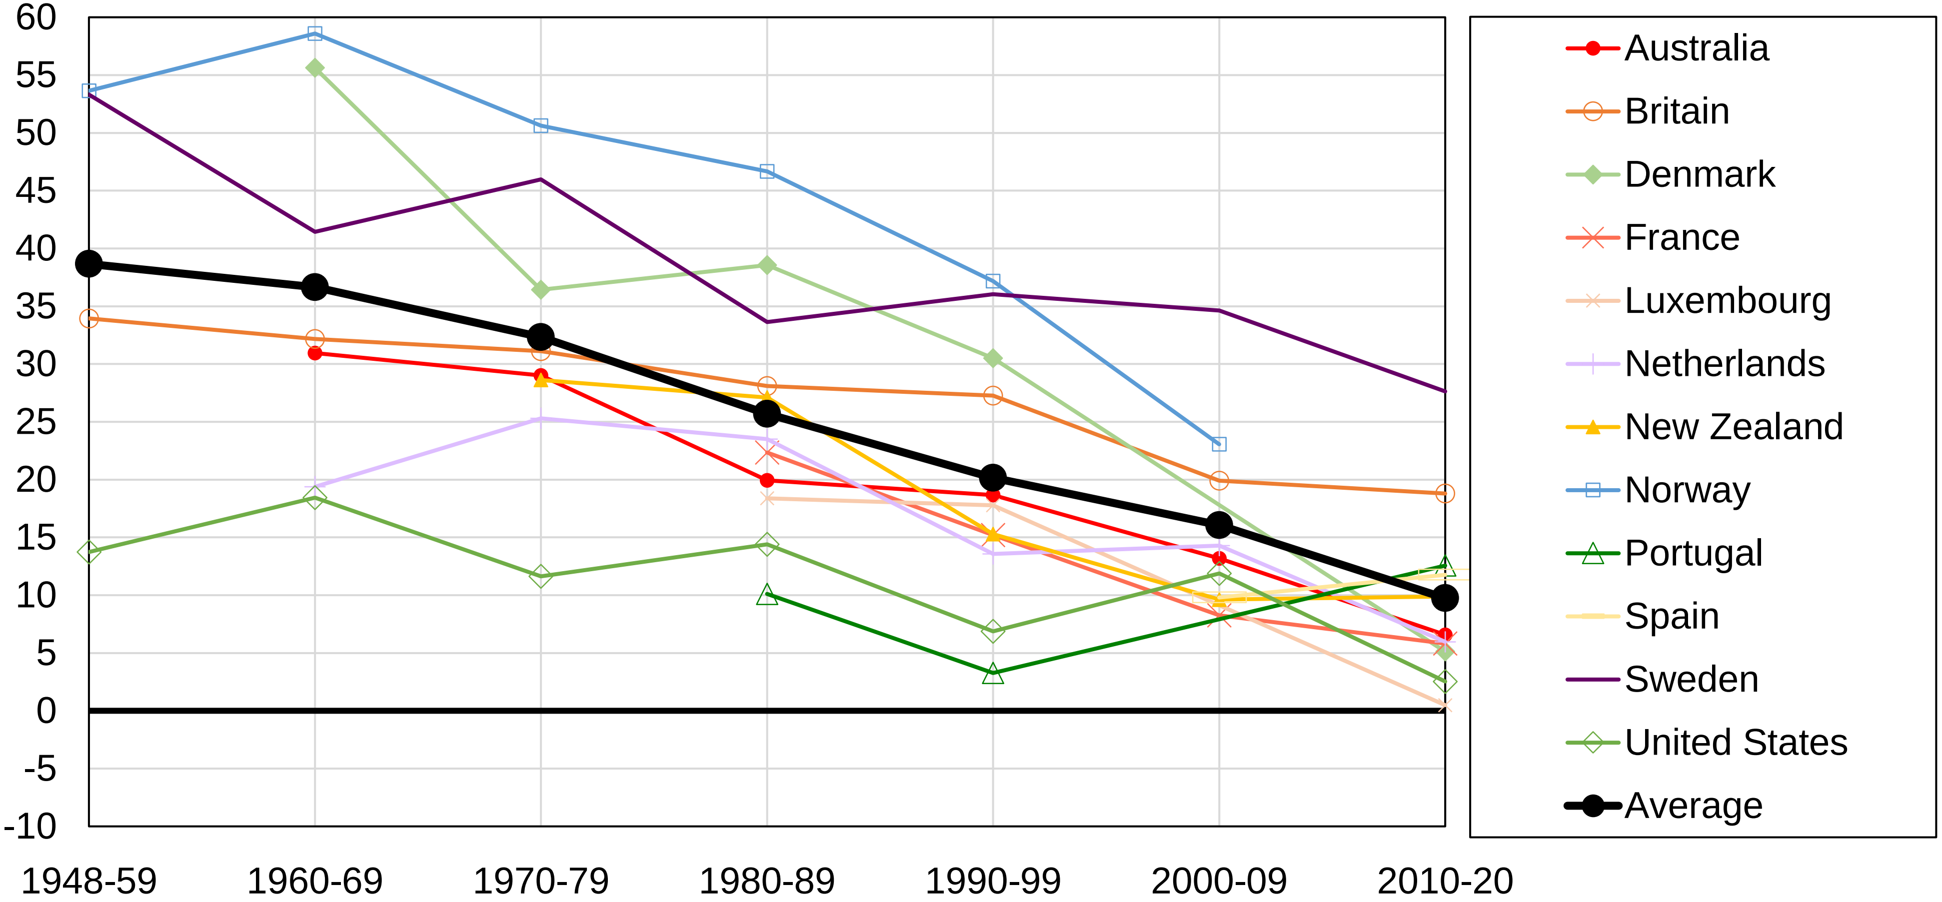

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the difference between the share of voters self-identifying as belonging to the "working class" or the "lower class" and the share of voters identifying with the "middle class", the "upper class" or "no class" voting for left-wing (socialist, social democratic, communist, and green) parties.

Figure CF2 - The decline of self-perceived class cleavages in Western democracies (after controls)

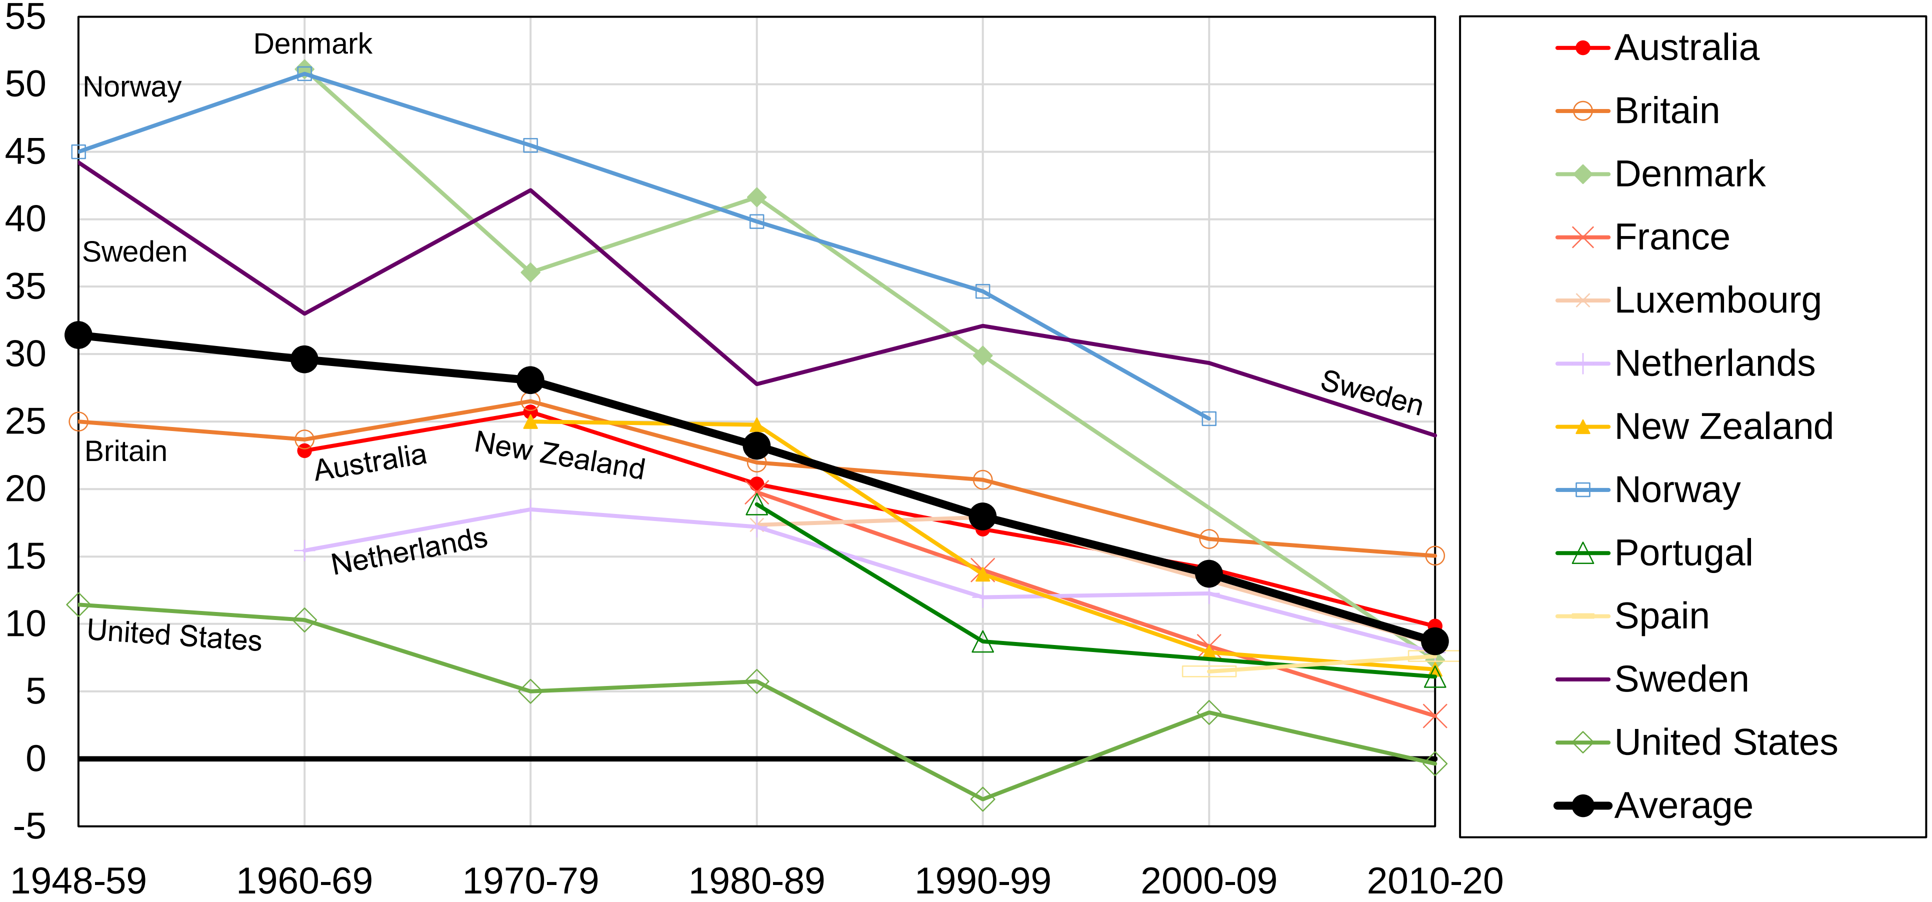

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the difference between the share of voters self-identifying as belonging to the "working class" or the "lower class" and the share of voters identifying with the "middle class", the "upper class" or "no class" voting for social democratic / socialist / communist / green parties. Self-perceived class cleavages have declined significantly over the past decades. Estimates control for income, education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available).

Figure CF3 - Vote for green parties by self-perceived class

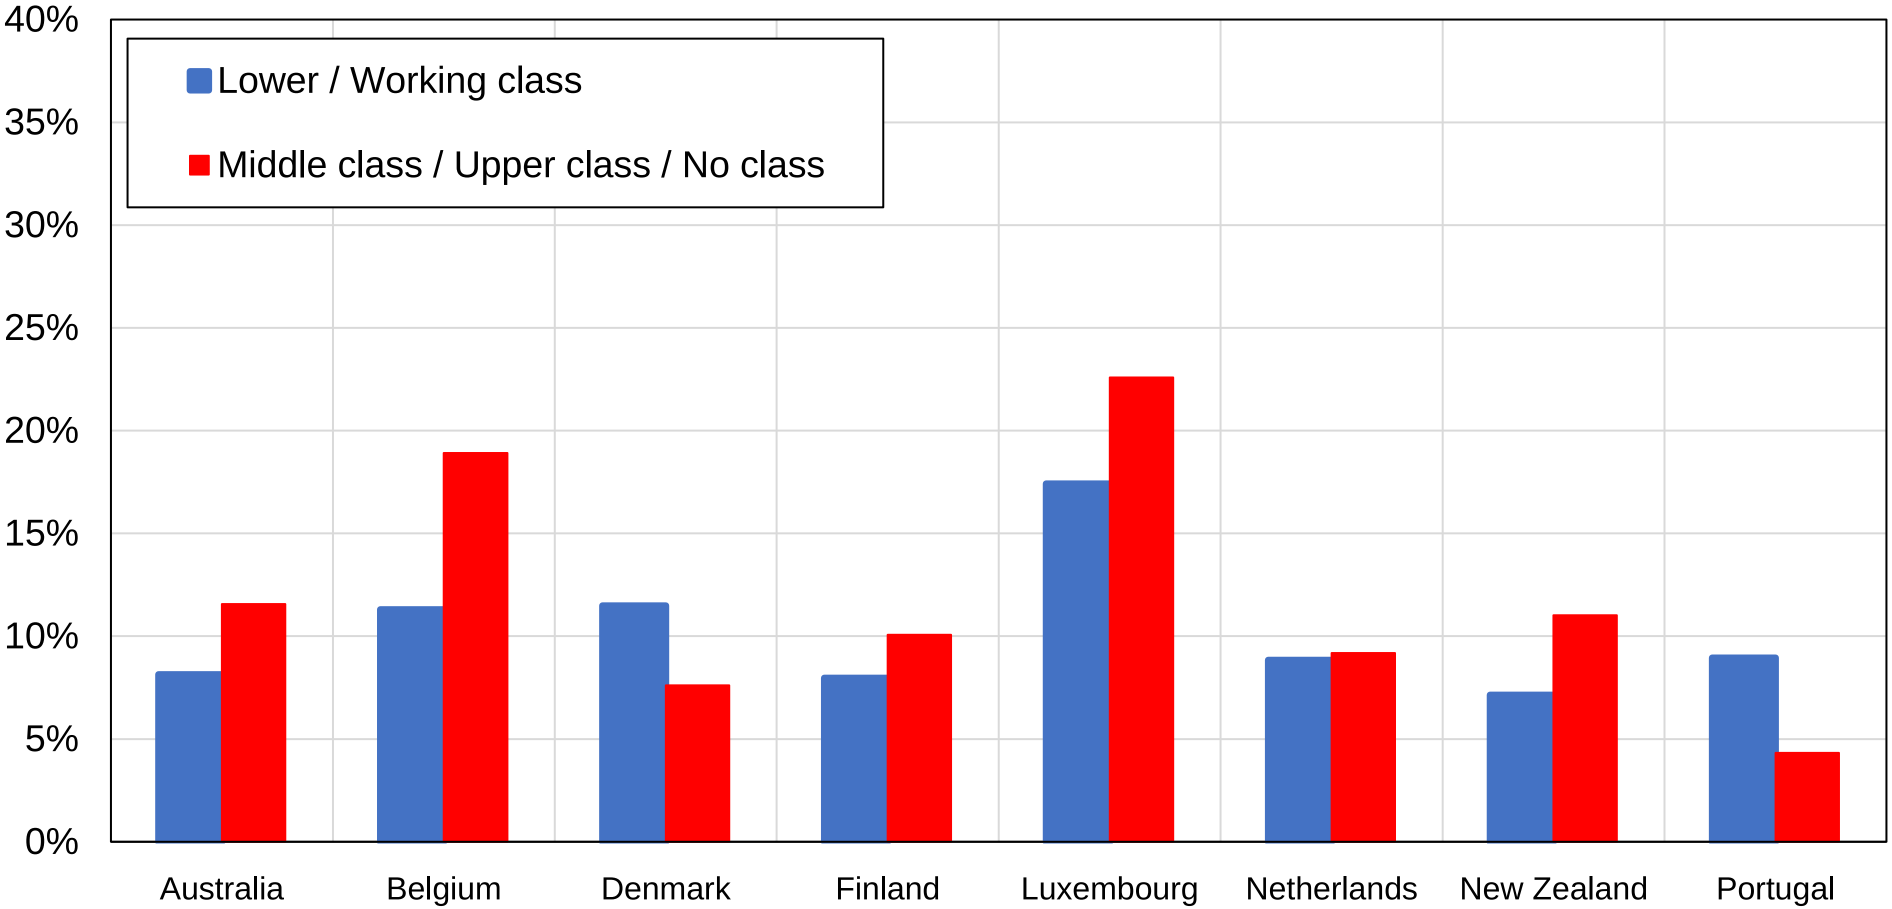

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by Green parties in Western democracies in the last election available by self-perceived social class.

Figure CF4 - Vote for anti-immigration parties by self-perceived class

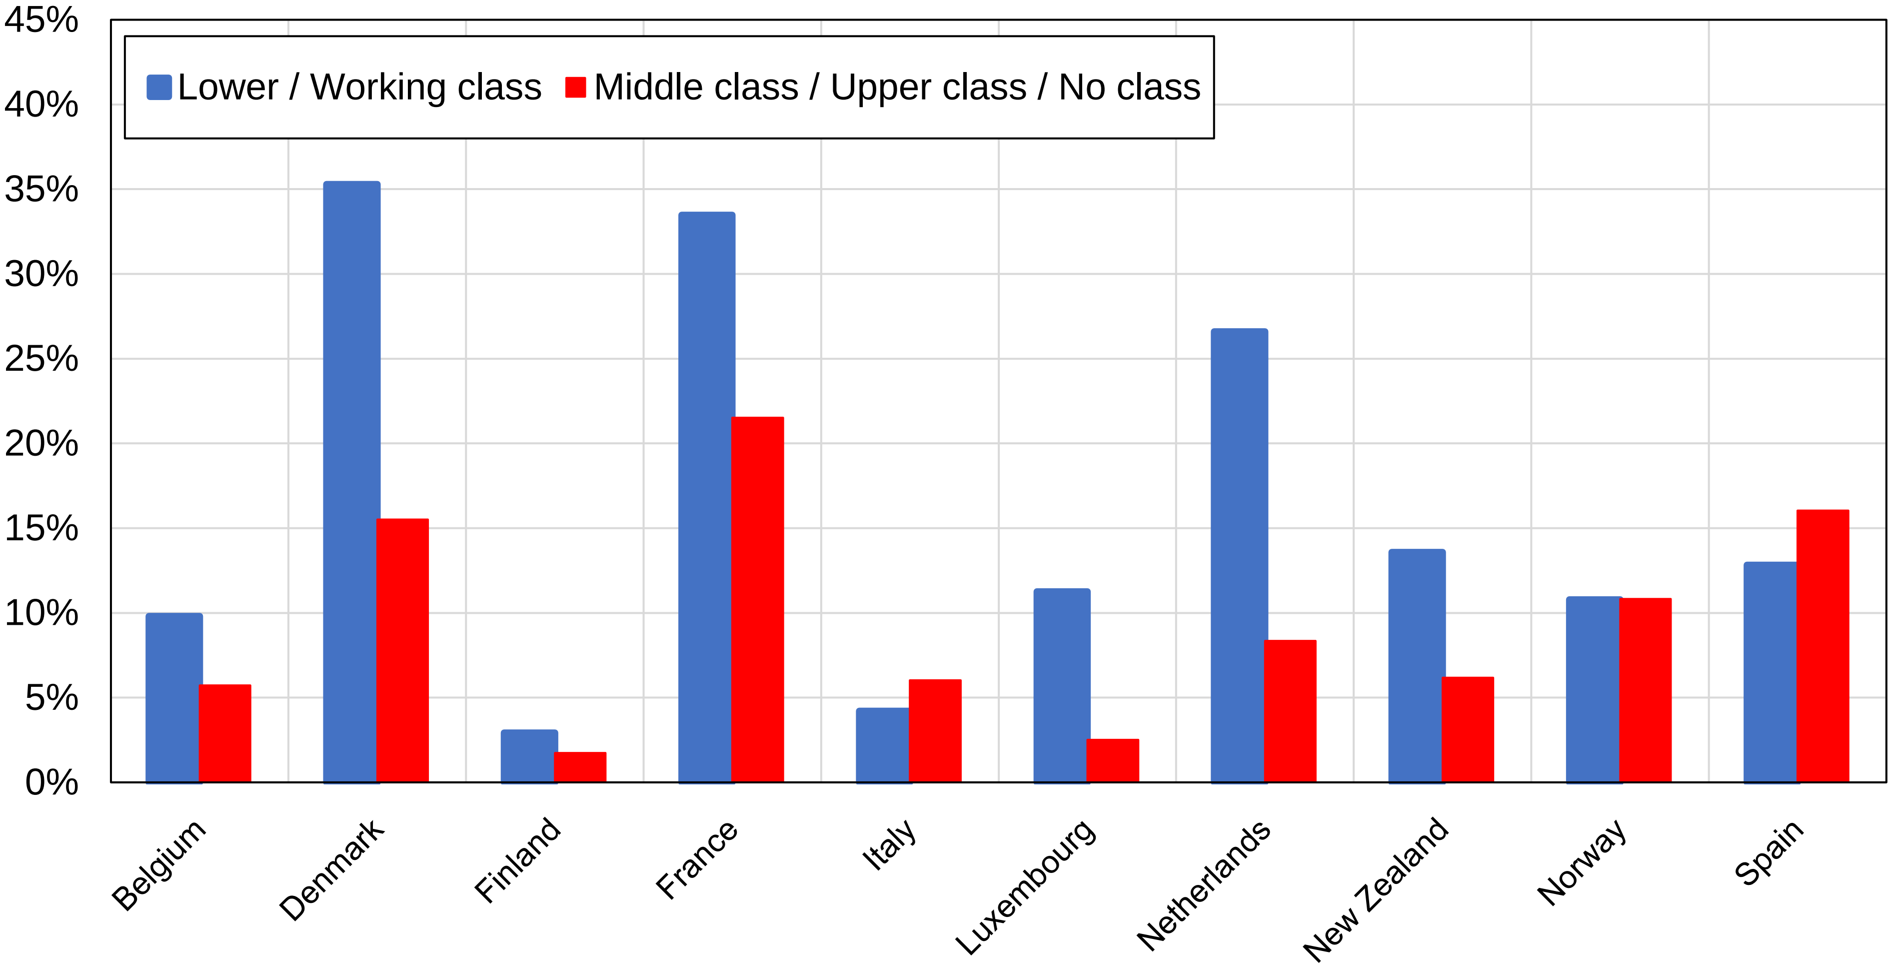

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure shows the share of votes received by anti-immigration parties in Western democracies in the last election available by self-perceived social class.

Figure CF5 - Vote for left-wing parties among union members in Western democracies

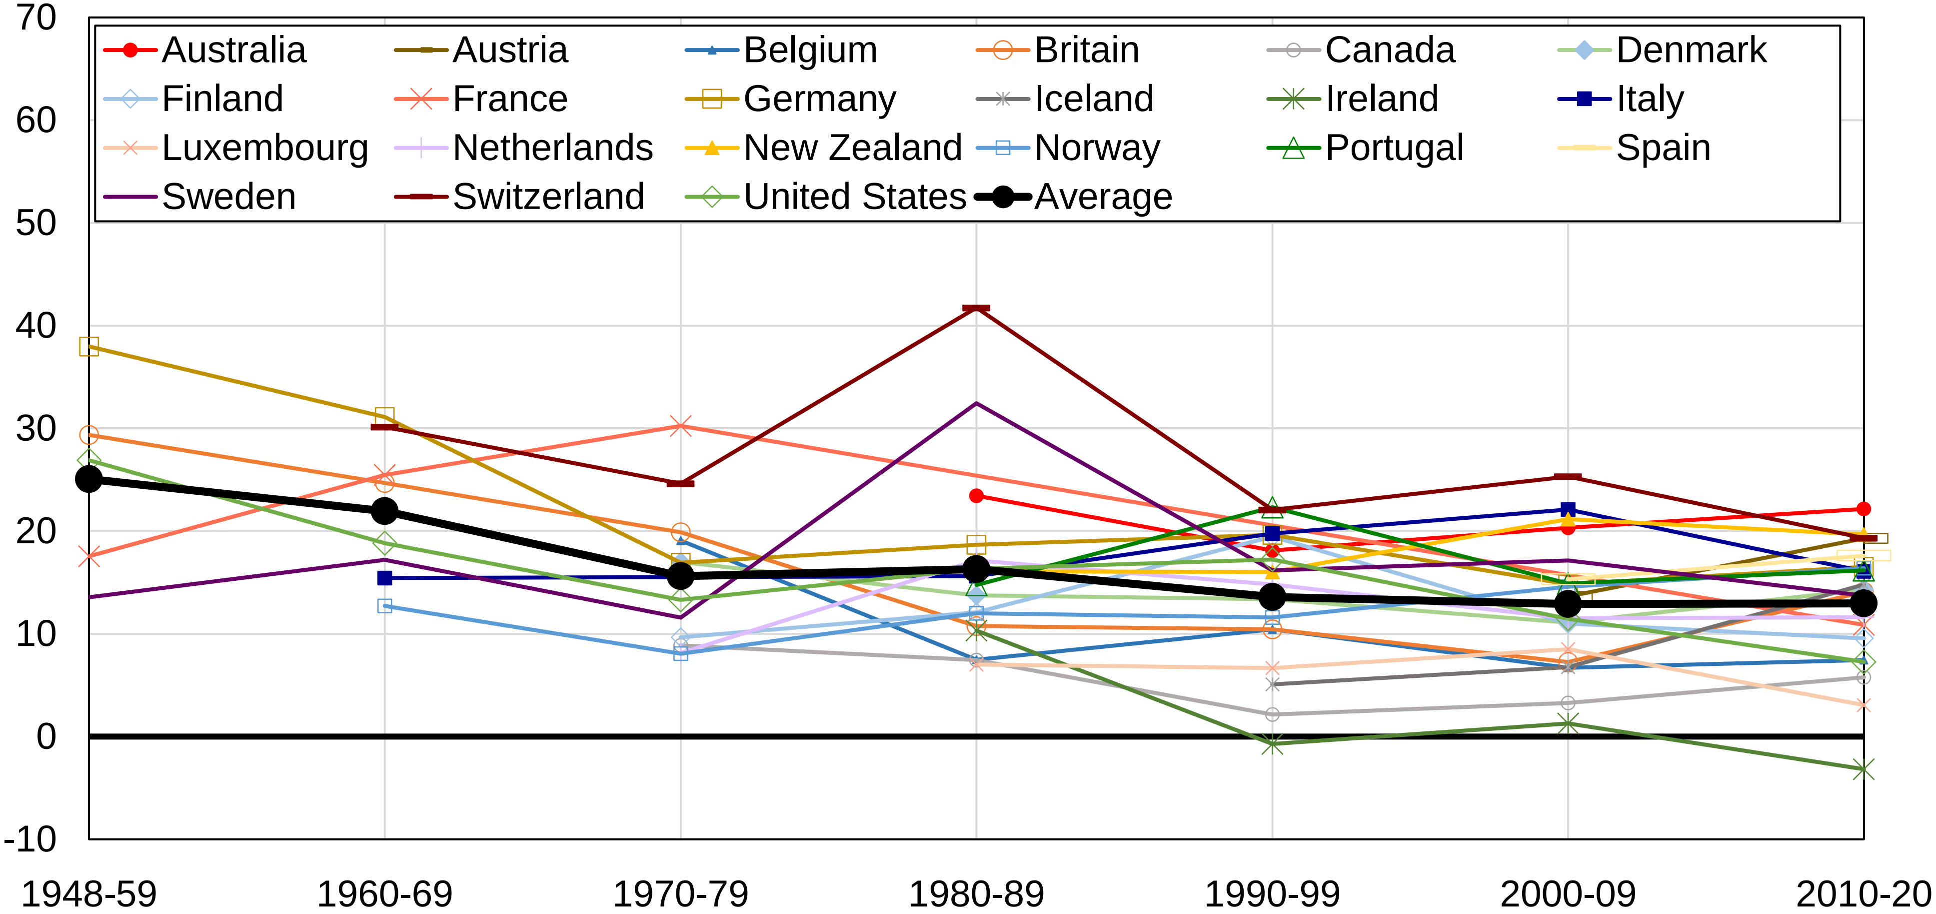

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of union members and the share of non-union members voting for social democratic, socialist, communist, and green parties in Western democracies.

Figure CF6 - Vote for left-wing parties among union members in Western democracies, after controls

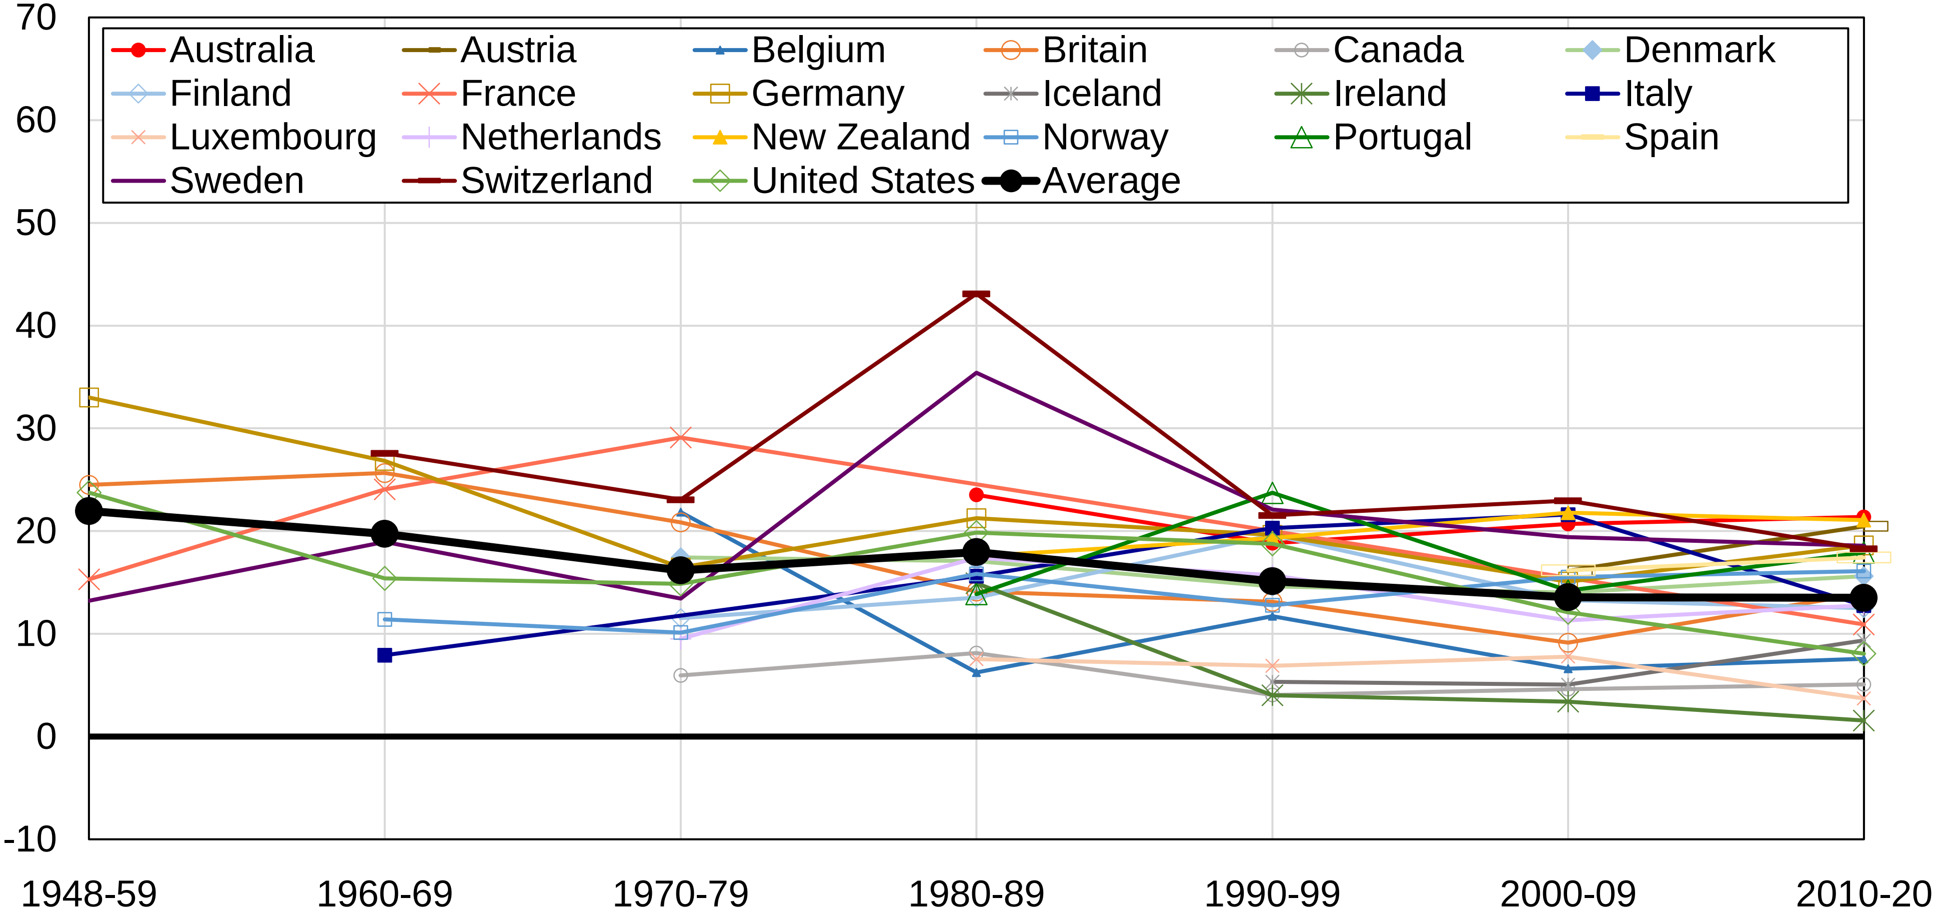

Source: authors' computations using the World Political Cleavages and Inequality Database.

Note: the figure represents the difference between the share of union members and the share of non-union members voting for social democratic, socialist, communist, and green parties in Western democracies. Estimates control for education, income, age, gender, religion, church attendance, rural/urban, region, employment status, and marital status (in country-years for which these variables are available).

Figure CF7 - Vote for left-wing parties among public sector workers in Western democracies

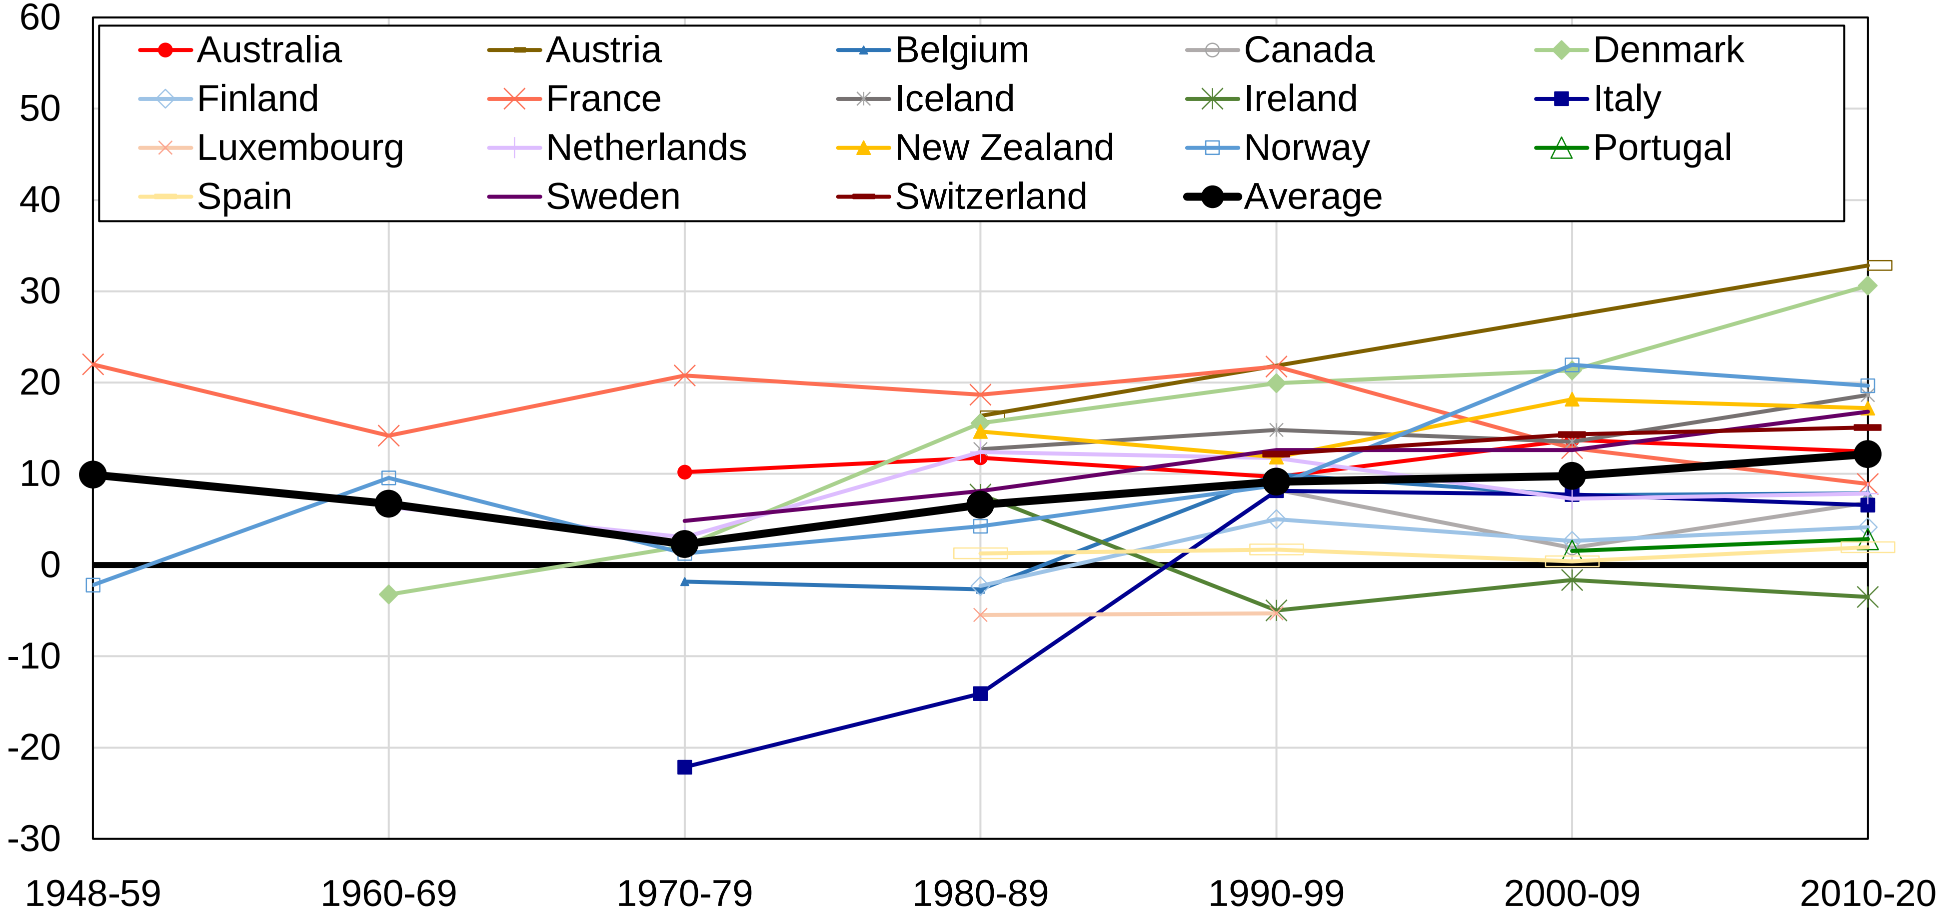

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of public sector workers and the share of private sector workers voting for social democratic, socialist, communist, and green parties in Western democracies.

Figure CF8 - Vote for left-wing parties among public sector workers in Western democracies, after controls

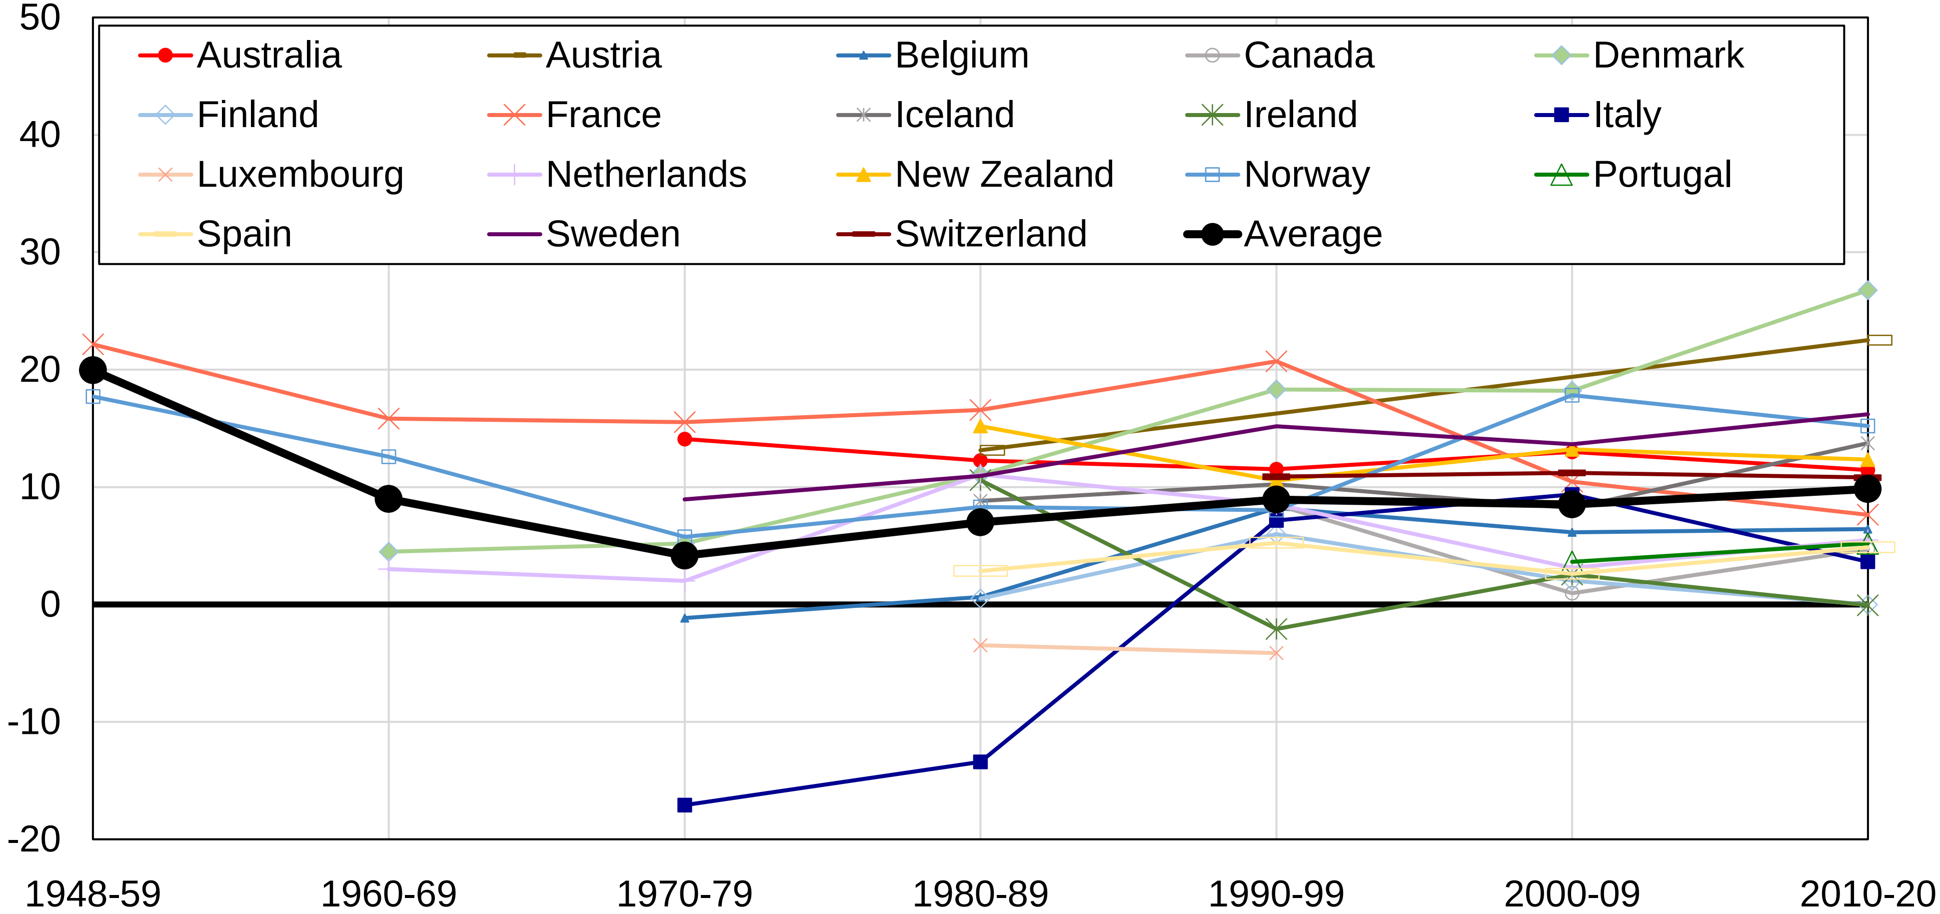

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of public sector workers and the share of private sector workers voting for social democratic, socialist, communist, and green parties in Western democracies. Estimates control for education, income, age, gender, religion, church attendance, rural/urban, region, employment status, and marital status (in country-years for which these variables are available).

Figure CF9 - Vote for left-wing parties among homeowners in Western democracies

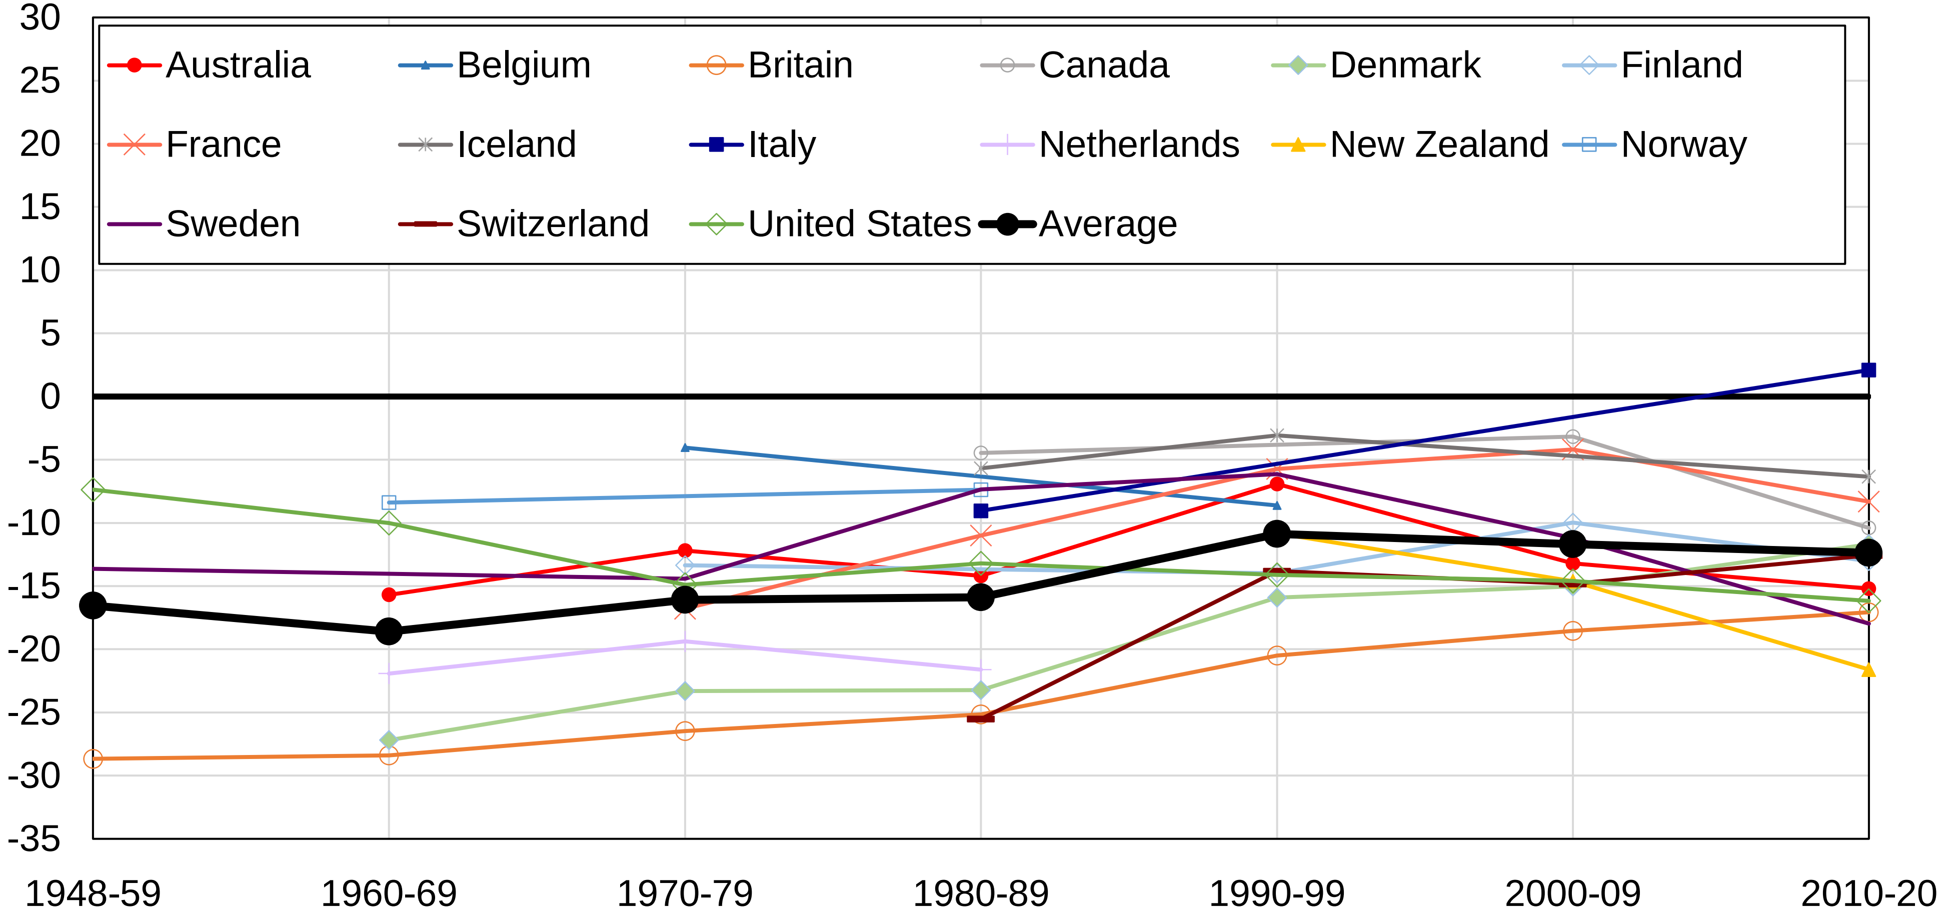

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of homeowners and the share of renters voting for social democratic, socialist, communist, and green parties in Western democracies.

Figure CF10 - Vote for left-wing parties among homeowners in Western democracies, after controls

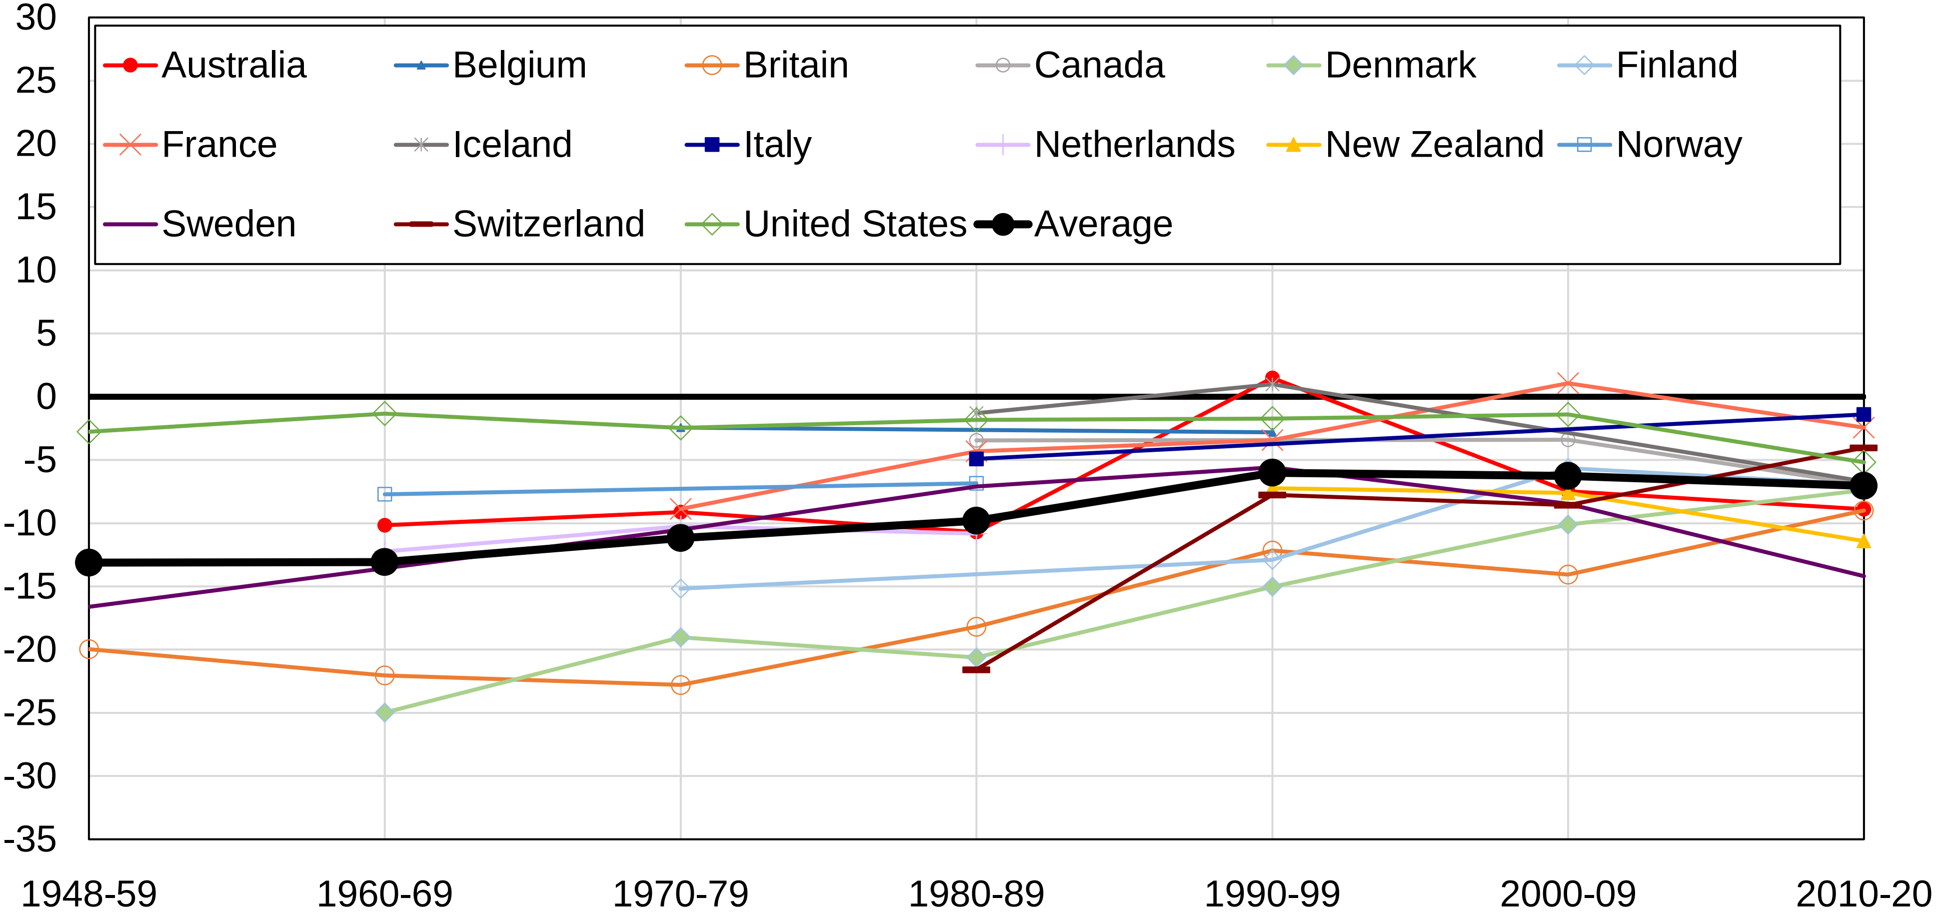

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the figure represents the difference between the share of homeowners and the share of renters voting for social democratic, socialist, communist, and green parties in Western democracies. Estimates control for education, income, age, gender, religion, church attendance, rural/urban, region, employment status, and marital status (in country-years for which these variables are available).

Table D1 - Marginal effect of belonging to top 10% educated voters on support for social democratic and affiliated parties by country and decade, after controls

	1948-59	1960-69	1970-79	1980-89	1990-99	2000-09	2010-2020
Australia		-13.3***	-4.8	-3.2	-3.0*	10.6***	5.9***
		(2.4)	(3.3)	(2.3)	(1.7)	(2.3)	(1.4)
Austria			-15.5***	-10.4***	-2.4	-2.4	12.8***
			(4.3)	(3.3)	(2.6)	(2.7)	(3.3)
Belgium			-10.6***	-8.2***	0.4	1.2	2.1*
			(1.5)	(1.5)	(1.5)	(1.1)	(1.3)
Canada		3.1	-2.0	-0.7	5.8***	6.9***	7.7***
		(3.0)	(2.9)	(1.9)	(2.0)	(1.9)	(1.7)
Denmark		-15.3***	-7.9***	-1.2	-0.8	5.0***	4.1***
		(3.5)	(2.0)	(1.7)	(1.2)	(0.9)	(1.5)
Finland			-19.0***	-14.6***	-8.4***	-5.6***	-2.6*
			(2.2)	(2.1)	(1.8)	(1.6)	(1.4)
France	-18.2***	-2.5	-4.8***	-0.0	8.1***	4.8***	10.2***
	(3.5)	(2.2)	(1.4)	(1.7)	(1.6)	(1.6)	(1.6)
Germany	-15.1***	-12.9***	-18.4***	-6.8*	3.1	4.2*	10.3***
	(2.2)	(2.7)	(3.8)	(3.5)	(2.4)	(2.4)	(2.7)
Iceland			6.7	-0.1	9.7***	5.6***	4.4***
			(5.7)	(3.1)	(2.4)	(1.5)	(1.0)
Ireland			-13.3***	-11.4***	-6.6***	-3.1*	-5.2***
			(2.3)	(1.4)	(2.0)	(1.7)	(1.4)
Italy	-9.9**	-0.9	-4.5	3.2	2.0	4.8**	5.2*
	(4.9)	(2.4)	(2.9)	(2.4)	(1.8)	(2.3)	(2.8)
Luxembourg			-16.0***	-5.1**	-0.6	7.4	5.5
			(3.1)	(2.5)	(1.8)	(4.9)	(5.0)
Netherlands		-1.8	-7.3***	1.9	11.5***	10.8***	10.7***
		(3.4)	(2.1)	(1.6)	(2.0)	(2.1)	(1.4)
New Zealand			0.6	-5.8*	2.3	12.7***	14.1***
			(2.3)	(3.3)	(1.5)	(1.8)	(1.7)
Norway	-30.7***	-33.1***	-16.8***	-12.0***	-3.6**	3.2***	2.7***

	(3.3)	(2.6)	(2.2)	(1.5)	(1.5)	(1.0)	(1.0)
Portugal				-8.9	-5.9	-8.1***	-16.4***
				(5.4)	(5.2)	(2.5)	(3.8)
Spain				-9.9***	-12.5***	-6.1***	-1.8**
				(1.4)	(1.9)	(1.3)	(0.7)
Sweden	-35.5***	-33.2***	-23.4***	-17.0***	-9.4***	-7.3***	-0.9
	(2.5)	(1.6)	(1.6)	(1.4)	(1.2)	(1.1)	(3.2)
Switzerland		-15.0***	-4.5*	-4.5	4.6**	10.1***	14.1***
		(5.3)	(2.7)	(4.4)	(2.3)	(2.0)	(1.2)
United Kingdom	-16.6***	-12.2***	-10.5***	-4.7***	-3.2**	-5.4***	2.1
	(2.6)	(2.2)	(1.0)	(1.1)	(1.4)	(1.5)	(1.6)
United States	-15.1***	-10.4***	-2.5	2.0	-3.1	4.6**	17.6***
	(2.0)	(2.3)	(2.1)	(1.8)	(2.1)	(1.9)	(1.1)

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the table reports the marginal effect of belonging to top 10% educated voters on the probability to support Social Democratic / Socialist / Green / Communist / Other left-wing parties, after controlling for income, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available). The original survey dataset is duplicated for each education category to approximate education deciles (see methodology). Robust standard errors clustered at the individual level. Coefficient standard errors in parenthesis. * p<0.10, ** p<0.05, *** p<0.01.

Table D2 - Marginal effect of belonging to top 10% income voters on support for social democratic and affiliated parties by country and decade, after controls

	1948-59	1960-69	1970-79	1980-89	1990-99	2000-09	2010-2020
Australia		-24.2***	-20.6***	-12.2***	-13.4***	-10.5***	-10.5***
		(2.7)	(3.8)	(3.7)	(2.0)	(2.2)	(1.8)
Austria			-17.9***	-7.4**	-2.4	-8.4*	-8.5**
			(4.0)	(3.6)	(3.0)	(4.9)	(3.4)
Belgium			-5.7***	-9.8***	-9.3***	-6.2***	-7.8***
			(1.6)	(1.6)	(1.5)	(1.7)	(1.7)
Canada		5.3*	-8.6***	-7.0***	-3.1	-5.4***	-7.0***
		(2.9)	(2.5)	(2.1)	(2.7)	(2.0)	(1.8)
Denmark		-12.6***	-14.9***	-22.2***	-19.8***	-14.5***	-14.6***
		(4.7)	(2.3)	(1.9)	(2.0)	(1.9)	(2.8)
Finland			-12.0***	-15.0***	-7.3***	-4.1*	-6.7***
			(2.4)	(2.1)	(2.1)	(2.2)	(1.9)
France	-0.8	-11.2***	-14.7***	-12.0***	-10.4***	-6.1***	-8.8***
	(5.7)	(2.8)	(1.6)	(1.8)	(1.8)	(1.5)	(2.6)
Germany	-11.4***	-17.7***	-12.1***		-11.8***	-10.1***	-13.8***
	(2.0)	(2.4)	(3.9)		(4.2)	(2.8)	(3.4)
Iceland				-4.0	-0.7	-6.2***	-7.1***
				(3.1)	(1.9)	(1.9)	(1.6)
Ireland			-6.7***	-8.1***	-10.6***	-1.3	-7.0***
			(2.5)	(1.3)	(2.7)	(3.1)	(2.4)
Italy	2.2	-6.6**	-1.4	-1.5		-3.0	4.6***
	(8.7)	(3.3)	(4.4)	(3.8)		(5.6)	(1.5)
Luxembourg			-7.8***	-7.6***	-5.0***	-18.2***	
			(2.9)	(2.4)	(1.6)	(6.1)	
Netherlands		-18.0***	-17.6***	-16.0***	-13.8***	-15.2***	-8.7***
		(3.5)	(2.8)	(2.0)	(2.4)	(2.4)	(1.9)
New Zealand			-19.9***	-6.4**	-11.8***	-11.3***	-12.2***
			(2.8)	(3.2)	(1.8)	(2.5)	(2.4)
Norway	-22.6***	-20.5***	-15.9***	-22.0***	-12.9***	-13.4***	-15.6***

	(4.1)	(2.6)	(2.0)	(2.0)	(2.8)	(2.2)	(2.5)
Portugal				-14.6*	-11.6*	-11.0***	-7.7
				(7.6)	(6.0)	(2.6)	(5.6)
Spain				-15.6***		-8.6***	-5.9***
				(3.2)		(1.5)	(1.4)
Sweden	-16.3***	-8.4***	-17.2***	-8.2***	-12.0***	-15.7***	-17.4***
	(3.7)	(1.6)	(1.9)	(1.3)	(1.9)	(1.8)	(2.3)
Switzerland			-11.9***		-7.2***	-11.7***	-5.6***
			(4.1)		(2.3)	(2.1)	(1.5)
United Kingdom	-23.7***	-31.3***	-15.6***	-15.3***	-10.0***	-6.8***	-7.6***
	(2.9)	(2.1)	(1.2)	(1.2)	(1.6)	(1.8)	(1.9)
United States	-9.6***	-8.2***	-12.8***	-13.1***	-7.7***	-11.1***	-0.0
	(2.2)	(2.4)	(2.2)	(2.1)	(2.7)	(2.8)	(1.9)

Source: authors' computations using the World Political Cleavages and Inequality Database.

Note: the table reports the marginal effect of belonging to top 10% income voters on the probability to support Social Democratic / Socialist / Communist / Green / Other left-wing parties, after controlling for education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available). The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). Robust standard errors clustered at the individual level. Coefficient standard errors in parenthesis. * p<0.10, ** p<0.05, *** p<0.01.

Table D3 - Marginal effect of belonging to top 10% educated voters on support for specific families of parties by country, 2010-2020, after controls

	Social Democratic / Socialist / Communist / Other left	Conservative / Christian Democratic / Liberal	Green	Anti-immmigration
Australia	1.1	-7.0***	4.8***	-0.2**
	(1.4)	(1.4)	(1.1)	(0.1)
Austria	3.2	3.8	9.6***	-16.4***
	(3.1)	(3.2)	(2.4)	(2.4)
Belgium	-4.3***	0.1	6.3***	-2.0***
	(1.1)	(1.3)	(0.9)	(0.4)
Canada	6.1***	-8.2***	1.7**	
	(1.7)	(1.6)	(0.7)	
Denmark	3.5**	1.0	0.6	-5.0***
	(1.4)	(1.4)	(0.8)	(0.9)
Finland	-6.3***	9.7***	3.7***	-7.9***
	(1.3)	(1.5)	(0.9)	(1.2)
France	9.0***	-2.0	1.2**	-11.7***
	(1.6)	(1.5)	(0.5)	(1.3)
Germany	-0.6	-7.7***	11.1***	-1.6
	(2.4)	(2.6)	(2.2)	(1.1)
Iceland	2.0**	-5.1***	2.4***	
	(0.8)	(1.0)	(0.7)	
Ireland	-6.0***	4.0***	0.8	
	(1.5)	(1.4)	(0.6)	
Italy	5.2*	-4.9*		
	(2.8)	(2.9)		(2.1)
Luxembourg	-0.7	-4.0	6.1	-1.3
	(5.0)	(5.0)	(4.7)	(1.7)
Netherlands	6.6***	-2.3*	4.1***	-8.0***
	(1.4)	(1.4)	(0.8)	(0.7)
New Zealand	6.0***	-14.7***	8.1***	-1.2

	(1.6)	(1.7)	(1.1)	(0.7)
Norway	1.8*	-0.8	0.9**	-3.6***
	(1.0)	(1.0)	(0.4)	(0.7)
Portugal	-14.1***	16.4***	-2.2	
	(3.6)	(3.8)	(1.7)	
Spain	-2.1***	4.8***	0.3***	-2.6***
	(0.7)	(0.7)	(0.1)	(0.4)
Sweden	-7.5**	5.2	6.6***	-4.3***
	(3.1)	(3.3)	(2.2)	(1.2)
Switzerland	6.2***	-0.6	7.9***	-13.0***
	(1.0)	(1.2)	(1.0)	(1.0)
United Kingdom	2.1	-10.2***		-2.3***
	(1.6)	(1.6)		(0.4)
United States	17.6***	-17.6***		
	(1.1)	(1.1)		

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the table reports the marginal effect of belonging to top 10% educated voters on the probability to support specific families of parties in the 2010-2020 period, after controlling for income, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available). The original survey dataset is duplicated for each education category to approximate education deciles (see methodology). Robust standard errors clustered at the individual level. Coefficient standard errors in parenthesis.
* p<0.10, ** p<0.05, *** p<0.01.

Table D4 - Marginal effect of belonging to top 10% income voters on support for specific families of parties by country, 2010-2020, after controls

	Social Democratic / Socialist / Communist / Other left	Conservative / Christian Democratic / Liberal	Green	Anti-immmigration
Australia	-7.5***	13.1***	-3.0**	-0.2
	(1.7)	(1.8)	(1.2)	(0.1)
Austria	-7.3**	12.9***	-1.2	-4.4
	(3.2)	(3.4)	(2.7)	(2.8)
Belgium	-5.9***	8.6***	-2.0*	-0.1
	(1.4)	(1.9)	(1.2)	(0.7)
Canada	-5.4***	7.6***	-1.6***	
	(1.8)	(1.7)	(0.6)	
Denmark	-9.3***	20.7***	-5.3***	-6.0***
	(2.8)	(2.9)	(1.1)	(1.6)
Finland	-7.0***	9.0***	0.3	-1.2
	(1.7)	(1.9)	(1.1)	(1.5)
France	-8.2***	13.5***	-0.5	-5.1**
	(2.6)	(2.9)	(0.6)	(2.1)
Germany	-12.0***	15.5***	-0.3	-0.2
	(3.0)	(3.5)	(2.7)	(1.8)
Iceland	-2.4*	10.5***	-4.7***	
	(1.3)	(1.8)	(1.3)	
Ireland	-7.9***	8.1***	0.9	
	(2.4)	(2.4)	(0.9)	
Italy	4.6***	0.9		-2.4*
	(1.5)	(1.6)		(1.3)
Netherlands	-7.1***	12.2***	-1.6*	-1.8*
	(1.8)	(1.9)	(0.9)	(1.0)
New Zealand	-9.9***	16.6***	-2.3*	-1.8*
	(2.3)	(2.5)	(1.2)	(1.1)
Norway	-13.0***	13.8***	-2.6***	1.0

	(2.5)	(2.9)	(0.7)	(2.0)
Portugal	-4.3	7.7	-3.5	
	(5.5)	(5.6)	(2.4)	
Spain	-6.0***	5.3***	0.1	1.5*
	(1.4)	(1.3)	(0.1)	(0.8)
Sweden	-13.4***	19.2***	-4.0***	-1.8
	(2.2)	(2.4)	(1.2)	(1.3)
Switzerland	-8.1***	9.6***	2.5**	-3.2**
	(1.3)	(1.7)	(1.2)	(1.4)
United Kingdom	-7.6***	15.0***		-1.9***
	(1.9)	(2.1)		(0.6)
United States	-0.0	0.0		
	(1.9)	(1.9)		

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the table reports the marginal effect of belonging to top 10% income voters on the probability to support specific families of parties in the 2010-2020 period, after controlling for education, age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status (in country-years for which these variables are available). The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). Robust standard errors clustered at the individual level. Coefficient standard errors in parenthesis. * p<0.10, ** p<0.05, *** p<0.01.

Table D5 - Effect of income and education on support for more left-wing parties (dummy income and education variables, continuous left-right ideological index)

	(1) 1948-1959	(2) 1960-69	(3) 1970-79	(4) 1980-89	(5) 1990-99	(6) 2000-09	(7) 2010-20
Income: Top 10%	-5.700***	-4.505***	-5.066***	-4.213***	-3.809***	-3.255***	-3.829***
	(0.673)	(0.430)	(0.327)	(0.267)	(0.207)	(0.191)	(0.200)
Education: University graduate	-10.880***	-6.278***	-2.158***	-1.060***	1.055***	2.212***	2.264***
	(1.211)	(0.614)	(0.388)	(0.251)	(0.195)	(0.174)	(0.165)
R-squared	0.35	0.24	0.24	0.23	0.34	0.27	0.17
Observations	35196	82331	158203	210450	170789	212937	208247

Note: The table reports the effect of income and education on support for more left-wing parties by decade across all Western democracies with available data. All estimates include election fixed effects. The dependent variable is the (inverted) left-right ideological index available from the Comparative Manifesto Project database, which theoretically ranges from -100 (most right-wing) to 100 (most left-wing). *p<0.10, **p<0.05, ***p<0.01.

Interpretation: in 1948-1959, higher income and higher education were both associated with support for more right-wing parties. By 2010-2020, higher income is still associated with support for more right-wing parties, but higher education is now associated with higher support for more left-wing parties.

Table D6 - Effect of income and education on support for more left-wing parties, after controls (dummy income and education variables, continuous left-right ideological index)

	(1) 1948-1959	(2) 1960-69	(3) 1970-79	(4) 1980-89	(5) 1990-99	(6) 2000-09	(7) 2010-20
Income: Top 10%	-6.445***	-5.398***	-5.618***	-4.370***	-3.523***	-3.135***	-3.462***
	(0.673)	(0.420)	(0.331)	(0.255)	(0.208)	(0.189)	(0.197)
Education: University graduate	-11.640***	-7.119***	-2.830***	-1.558***	0.700***	1.734***	1.667***
	(1.230)	(0.614)	(0.391)	(0.250)	(0.195)	(0.174)	(0.169)
R-squared	0.37	0.26	0.26	0.25	0.36	0.32	0.22
Observations	35196	82331	158203	210450	170789	212937	208247

Note: The table reports the effect of income and education on support for more left-wing parties by decade across all Western democracies with available data. The dependent variable is the (inverted) left-right ideological index available from the Comparative Manifesto Project database, theoretically ranging from -100 (most right-wing) to 100 (most left-wing). All estimates include election fixed effects and control for the following variables (in country-years for which they are available): age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status. *p<0.10, **p<0.05, ***p<0.01.

Interpretation: in 1948-1959, higher income and higher education were both associated with support for more right-wing parties. By 2010-2020, higher income is still associated with support for more right-wing parties, but higher education is now associated with higher support for more left-wing parties.

Table D7 - Effect of income and education on support for more left-wing parties
(continuous income and education variables, continuous left-right ideological index)

	(1) 1948-1959	(2) 1960-69	(3) 1970-79	(4) 1980-89	(5) 1990-99	(6) 2000-09	(7) 2010-20
Income rank	0.422	1.957***	-1.569***	-3.589***	-4.882***	-4.909***	-5.688***
	(0.722)	(0.408)	(0.387)	(0.294)	(0.269)	(0.262)	(0.287)
Education rank	-10.214***	-9.269***	-5.295***	-2.413***	1.142***	3.993***	4.649***
	(0.590)	(0.382)	(0.351)	(0.266)	(0.261)	(0.255)	(0.281)
R-squared	0.35	0.24	0.24	0.23	0.34	0.27	0.18
Observations	13025	34028	70328	91076	86594	97681	100116

Note: The table reports the effect of income and education on support for more left-wing parties by decade across all Western democracies with available data. All estimates include election fixed effects. Income and education ranks/quantiles (ranging from 0 to 1) are defined discretely based on all income and education categories available in each survey. The dependent variable is the (inverted) left-right ideological index available from the Comparative Manifesto Project database, which theoretically ranges from -100 (most right-wing) to 100 (most left-wing). *p<0.10, **p<0.05, ***p<0.01.

Interpretation: in 1948-1969, higher income and higher education were both associated with support for more right-wing parties. By 2010-2020, higher income is still associated with support for more right-wing parties, but higher education is now associated with higher support for more left-wing parties.

Table D8 - Effect of income and education on support for more left-wing parties, after controls (continuous income and education variables, continuous left-right ideological index)

	(1) 1948-1959	(2) 1960-69	(3) 1970-79	(4) 1980-89	(5) 1990-99	(6) 2000-09	(7) 2010-20
Income rank	-2.719***	-0.766*	-3.918***	-4.838***	-5.025***	-5.312***	-5.139***
	(0.775)	(0.437)	(0.440)	(0.315)	(0.299)	(0.290)	(0.312)
Education rank	-10.066***	-10.112***	-6.448***	-3.841***	0.291	3.099***	3.119***
	(0.607)	(0.389)	(0.360)	(0.272)	(0.271)	(0.264)	(0.293)
R-squared	0.37	0.27	0.27	0.25	0.36	0.32	0.22
Observations	13025	34028	70328	91076	86594	97681	100116

Note: The table reports the effect of income and education on support for more left-wing parties by decade across all Western democracies with available data. Income and education ranks/quantiles (ranging from 0 to 1) are defined discretely based on all income and education categories available in each survey. The dependent variable is the (inverted) left-right ideological index available from the Comparative Manifesto Project database, theoretically ranging from -100 (most right-wing) to 100 (most left-wing). All estimates include election fixed effects and control for the following variables (in country-years for which they are available): age, gender, religion, church attendance, rural/urban, region, race/ethnicity, employment status, and marital status. *p<0.10, **p<0.05, ***p<0.01.

Interpretation: in 1948-1959, higher income and higher education were both associated with support for more right-wing parties. By 2010-2020, higher income is still associated with support for more right-wing parties, but higher education is now associated with higher support for more left-wing parties.

Table D9 - The reversal of educational divides, 1960-2020: before and after controls

	1960-69	1970-79	1980-89	1990-99	2000-09	2010-20	Difference 2010s-1960s
Raw coefficient	-21.6***	-11.8***	-7.3***	-2.7***	3.4***	5.3***	26,9
	(1.0)	(0.7)	(0.7)	(0.6)	(0.6)	(0.6)	
After controlling for income	-18.0***	-9.8***	-4.9***	-0.8	5.1***	6.6***	24,6
	(1.0)	(0.7)	(0.7)	(0.6)	(0.6)	(0.6)	
After controlling for the above and: Gender	-18.3***	-10.1***	-4.9***	-0.8	5.0***	6.5***	24,8
	(1.0)	(0.7)	(0.7)	(0.6)	(0.6)	(0.6)	
After controlling for the above and: Age	-18.9***	-11.0***	-5.9***	-1.5**	4.6***	5.7***	24,6
	(1.0)	(0.7)	(0.7)	(0.6)	(0.6)	(0.5)	
After controlling for the above and: Religion	-19.1***	-11.4***	-6.5***	-2.3***	4.1***	4.9***	24,0
	(1.0)	(0.7)	(0.7)	(0.6)	(0.6)	(0.5)	
After controlling for the above and: Religious practice	-18.5***	-10.9***	-5.9***	-1.8***	4.3***	5.1***	23,6
	(1.0)	(0.7)	(0.7)	(0.6)	(0.6)	(0.5)	
After controlling for the above and: Rural/urban	-19.2***	-11.6***	-6.5***	-2.2***	3.8***	4.6***	23,8
	(1.0)	(0.7)	(0.7)	(0.6)	(0.6)	(0.5)	
After controlling for the above and: Region	-19.9***	-11.9***	-6.6***	-2.2***	3.6***	4.5***	24,4
	(1.0)	(0.7)	(0.7)	(0.6)	(0.6)	(0.5)	
After controlling for the above and: Employment/marital status	-19.4***	-11.7***	-6.5***	-2.3***	3.6***	4.6***	24,0
	(1.0)	(0.7)	(0.7)	(0.6)	(0.6)	(0.5)	
After controlling for the above and: Sector of employment	-19.7***	-12.5***	-7.7***	-3.7***	2.1***	3.6***	23,3
	(1.0)	(0.7)	(0.7)	(0.6)	(0.6)	(0.5)	
After controlling for the above and: Union membership	-19.4***	-12.5***	-7.9***	-3.7***	1.7***	3.3***	22,7
	(1.0)	(0.7)	(0.7)	(0.6)	(0.6)	(0.5)	
After controlling for the above and: Home ownership	-18.8***	-12.1***	-7.7***	-3.7***	1.9***	3.6***	22,4
	(1.0)	(0.7)	(0.7)	(0.6)	(0.6)	(0.5)	

Note: The table reports the marginal effect of belonging to top 10% educated voters on the probability to support Social Democratic / Socialist / Communist / Green / Other left-wing parties, before and after controlling for a set of covariates. The regressions are run on the restricted number of countries for which these covariates are available in most decades: Australia, Denmark, Finland, France, the Netherlands, Norway, New Zealand, Sweden, the United Kingdom, and the United States. All estimates include election (country-year) fixed effects.

Table D10 - The reversal of educational divides by subgroup

	1948-59	1960-69	1970-79	1980-89	1990-99	2000-09	2010-20	2010s - 1950s
Gender								
Men	-25.1***	-17.0***	-7.8***	-3.8***	3.8***	5.0***	6.9***	32,0
	(1.6)	(1.4)	(1.0)	(0.9)	(0.7)	(0.7)	(0.7)	
Women	-24.9***	-16.7***	-13.2***	-7.4***	-3.3***	0.8	2.0***	26,9
	(1.9)	(1.1)	(0.8)	(0.7)	(0.7)	(0.7)	(0.7)	
Location								
Urban areas	-25.2***	-16.2***	-11.9***	-8.6***	-2.5***	1.7***	3.3***	28,5
	(2.6)	(1.2)	(1.0)	(0.8)	(0.8)	(0.6)	(0.6)	
Rural areas	-18.1***	-13.4***	-2.9*	-3.7***	3.5**	5.9***	9.7***	27,8
	(3.2)	(1.9)	(1.7)	(1.4)	(1.7)	(1.4)	(1.6)	
Religion								
No religion	-24.0***	-24.0***	-6.1***	-0.6	4.6***	5.2***	7.8***	31,8
	(8.9)	(3.0)	(1.5)	(1.4)	(1.2)	(1.1)	(0.9)	
Christian / Other	-18.1***	-13.6***	-11.6***	-8.1***	-3.0***	-0.2	1.3**	19,4
	(2.4)	(1.1)	(0.9)	(0.9)	(0.7)	(0.7)	(0.7)	
Sector of employment								
Private sector	-25.2***	-20.7***	-14.9***	-8.1***	-4.2***	-0.9	1.5**	26,7
	(3.5)	(3.0)	(1.6)	(1.2)	(0.8)	(0.9)	(0.7)	
Public sector	-12.3**	-22.4***	-3.4*	-3.1***	1.1	5.4***	5.8***	18,1
	(6.0)	(4.7)	(1.8)	(1.2)	(0.9)	(0.9)	(0.8)	
Subjective social class								
Working/Lower class	-13.4***	-6.7***	0.8	-3.7*	-1.7	2.6	4.9*	18,3
	(4.3)	(2.6)	(2.0)	(2.0)	(2.2)	(2.4)	(2.8)	
Middle/Upper class	-11.0***	-6.0***	0.2	0.9	3.2***	5.3***	7.3***	18,3
	(2.4)	(1.2)	(1.2)	(1.0)	(0.7)	(0.8)	(1.0)	

Source: authors' computations using the World Political Cleavages and Inequality Database.
Note: the table reports the unconditional effect of belonging to top 10% educated voters on the probability to support Social Democratic / Socialist / Green / Other left-wing parties, decomposed by subgroup of voters. Within nearly all groups, most educated voters used to be significantly less likely to vote for these parties in the 1950s and 1960s. By the 2010s, they had become significantly more likely to do so. Figures correspond to regression results on all countries with available data for each decade. All estimates include election fixed effects. The original survey dataset is duplicated for each education category to approximate education deciles (see methodology). Robust standard errors clustered at the individual level. Coefficient standard errors in parenthesis. * p<0.10, ** p<0.05, *** p<0.01.

Table E1 - Determinants of support for Labor / Greens in Australia

	(1) 1960-69	(2) 1970-79	(3) 1980-89	(4) 1990-99	(5) 2000-09	(6) 2010-20
Education: None/Primary	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Education: Secondary	-0.118***	-0.054**	-0.045***	-0.040***	-0.019	-0.010
	(0.018)	(0.024)	(0.015)	(0.012)	(0.017)	(0.013)
Education: University	-0.238***	-0.047	-0.046*	-0.070***	0.086***	0.061***
	(0.031)	(0.040)	(0.025)	(0.021)	(0.023)	(0.017)
Education: Postgraduate			-0.125	-0.021	0.128***	0.077***
			(0.082)	(0.025)	(0.026)	(0.017)
Income group: Bottom 50%	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Income group: Middle 40%	-0.165***	-0.071***	-0.012	-0.038***	-0.104***	-0.034***
	(0.018)	(0.015)	(0.023)	(0.014)	(0.016)	(0.013)
Income group: Top 10%	-0.333***	-0.275***	-0.110***	-0.149***	-0.179***	-0.126***
	(0.029)	(0.024)	(0.037)	(0.021)	(0.024)	(0.020)
Age: 20-39	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Age: 40-59	0.020	0.009	-0.033**	-0.067***	0.028	-0.001
	(0.018)	(0.017)	(0.016)	(0.013)	(0.017)	(0.015)
Age: 60+	-0.067**	-0.059***	-0.074***	-0.098***	-0.071***	-0.112***
	(0.028)	(0.022)	(0.020)	(0.016)	(0.021)	(0.016)
Gender: Woman	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Gender: Man	0.068***	0.087***	0.048***	0.039***	-0.012	-0.073***
	(0.019)	(0.017)	(0.014)	(0.011)	(0.014)	(0.011)
Religion: None			(baseline)	(baseline)	(baseline)	(baseline)
			(.)	(.)	(.)	(.)
Religion: Catholic			0.084**	0.064***	-0.049**	-0.090***
			(0.041)	(0.022)	(0.024)	(0.018)
Religion: Other Christian			-0.024	-0.059***	-0.146***	-0.165***
			(0.037)	(0.019)	(0.021)	(0.015)
Religion: Other			0.099	-0.029	0.040	-0.043
			(0.113)	(0.027)	(0.044)	(0.033)
Religion: Muslim				0.274*	0.307***	0.193***
				(0.152)	(0.069)	(0.058)
Religious practice: Never	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Religious practice: Less than monthly	-0.083***	-0.034	-0.094***	-0.068***	-0.026	-0.053***
	(0.025)	(0.034)	(0.016)	(0.015)	(0.018)	(0.015)
Religious practice: Monthly or more	-0.095***	-0.096***	-0.210***	-0.136***	-0.092***	-0.124***
	(0.026)	(0.035)	(0.019)	(0.017)	(0.021)	(0.017)
Location: Urban	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Location: Rural	-0.132***	-0.105***	-0.130***	-0.090***	-0.093***	-0.063***
	(0.018)	(0.023)	(0.015)	(0.020)	(0.017)	(0.013)
Employment status: Employed	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Employment status: Unemployed/Inactive	-0.017	0.018	0.045***	0.036**	0.026	0.013
	(0.039)	(0.019)	(0.016)	(0.014)	(0.017)	(0.013)

Marital status: Single	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Marital status: Married/With partner	0.018	-0.014	-0.023	-0.039***	-0.029*	-0.037***
	(0.024)	(0.018)	(0.015)	(0.013)	(0.015)	(0.012)
Region: Australian Capital Territory			(baseline)	(baseline)	(baseline)	(baseline)
			(.)	(.)	(.)	(.)
Region: New South Wales			-0.044	0.009	-0.142***	-0.177***
			(0.094)	(0.044)	(0.044)	(0.032)
Region: Northern Territory			-0.085	0.093	-0.285***	-0.130*
			(0.202)	(0.082)	(0.101)	(0.069)
Region: Queensland			-0.018	-0.050	-0.181***	-0.234***
			(0.097)	(0.045)	(0.045)	(0.033)
Region: South Australia			-0.091	-0.074	-0.188***	-0.198***
			(0.098)	(0.045)	(0.049)	(0.035)
Region: Tasmania			-0.000	0.067	-0.051	-0.145***
			(0.113)	(0.049)	(0.058)	(0.043)
Region: Victoria			-0.081	0.015	-0.122***	-0.164***
			(0.095)	(0.044)	(0.045)	(0.032)
Region: Western Australia			-0.062	-0.032	-0.137***	-0.232***
			(0.100)	(0.046)	(0.047)	(0.034)
Constant	0.728***	0.510***	0.705***	0.623***	0.804***	0.886***
	(0.038)	(0.036)	(0.100)	(0.048)	(0.048)	(0.035)
R-squared	0.10	0.07	0.05	0.05	0.07	0.09
Observations	9787	10182	7064	12457	8151	14875
Clusters	2039	4066	2934	2997	2001	3932

Note: The table reports the effect of a set of individual characteristics on the probability to vote for Labor / Greens by decade in Australia. The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). The number of clusters corresponds to the number of surveyed individuals in each decade. Robust standard errors clustered at the individual level.

Table E2 - Determinants of support for SPÖ / KPÖ / Greens / NEOS in Austria

	(1) 1970-79	(2) 1980-89	(3) 1990-99	(4) 2000-09	(5) 2010-20
Education: None/Primary	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Education: Secondary	-0.056**	-0.033	0.001	0.001	0.001
	(0.028)	(0.024)	(0.019)	(0.025)	(0.035)
Education: University	-0.197***	-0.133***	-0.029	-0.028	0.170***
	(0.050)	(0.036)	(0.029)	(0.037)	(0.044)
Income group: Bottom 50%	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Income group: Middle 40%	-0.039	-0.085***	-0.022	-0.021	0.005
	(0.031)	(0.023)	(0.016)	(0.024)	(0.025)
Income group: Top 10%	-0.201***	-0.123***	-0.024	-0.101*	-0.076*
	(0.045)	(0.038)	(0.027)	(0.053)	(0.039)
Age: 20-39	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Age: 40-59	0.010	-0.010	-0.031*	0.067**	-0.053*
	(0.029)	(0.027)	(0.016)	(0.027)	(0.030)
Age: 60+	0.058	-0.075**	-0.072***	0.009	-0.125***
	(0.036)	(0.032)	(0.021)	(0.032)	(0.039)
Gender: Woman	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Gender: Man	0.029	-0.007	-0.043***	-0.057***	-0.081***
	(0.028)	(0.022)	(0.014)	(0.021)	(0.023)
Religion: None	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Religion: Catholic	0.002	-0.196***	-0.086*	-0.009	-0.135***
	(0.058)	(0.032)	(0.050)	(0.042)	(0.042)
Religion: Other Christian	-0.119	-0.205***	0.013	0.245***	-0.021
	(0.087)	(0.055)	(0.073)	(0.064)	(0.074)
Religion: Other	0.131	0.032	0.124	0.180	0.150*
	(0.172)	(0.102)	(0.129)	(0.132)	(0.087)
Religion: Muslim			-0.643***	0.275**	0.290***
			(0.055)	(0.128)	(0.108)
Religious practice: Never	(baseline)	(baseline)		(baseline)	(baseline)
	(.)	(.)		(.)	(.)
Religious practice: Less than monthly	-0.045	-0.148***		-0.212***	-0.056
	(0.037)	(0.030)		(0.038)	(0.037)
Religious practice: Monthly or more	-0.469***	-0.388***		-0.436***	-0.116***
	(0.036)	(0.028)		(0.042)	(0.041)
Location: Urban	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Location: Rural	-0.055*	-0.092***	-0.073***	-0.102**	-0.071**
	(0.029)	(0.024)	(0.024)	(0.042)	(0.028)
Employment status: Employed	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Employment status: Unemployed/Inactive	0.068**	0.032	0.004	0.032	0.107***
	(0.030)	(0.026)	(0.016)	(0.024)	(0.034)
Marital status: Single	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)

Marital status: Married/With partner	-0.042	0.011	-0.013	-0.010	-0.029
	(0.028)	(0.022)	(0.015)	(0.023)	(0.027)
Region: Burgenland				(baseline)	(baseline)
				(.)	(.)
Region: Carinthia				-0.110	-0.255**
				(0.067)	(0.105)
Region: Lower Austria				-0.228***	-0.245***
				(0.053)	(0.094)
Region: Salzburg				-0.210***	-0.349***
				(0.066)	(0.105)
Region: Styria				-0.213***	-0.193*
				(0.057)	(0.100)
Region: Tyrol				-0.218***	-0.334***
				(0.059)	(0.099)
Region: Upper Austria				-0.193***	-0.211**
				(0.054)	(0.099)
Region: Vienna				-0.153***	-0.168*
				(0.055)	(0.098)
Region: Vorarlberg				-0.221***	-0.065
				(0.065)	(0.117)
Constant	0.861***	1.003***	0.647***	1.008***	0.881***
	(0.069)	(0.042)	(0.053)	(0.268)	(0.105)
R-squared	0.23	0.20	0.02	0.12	0.12
Observations	2137	4158	11336	8514	3559
Clusters	1336	2688	6468	2731	1162

Note: The table reports the effect of a set of individual characteristics on the probability to vote for SPÖ / KPÖ / Greens / NEOS by decade in Austria. The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). The number of clusters corresponds to the number of surveyed individuals in each decade. Robust standard errors clustered at the individual level.

Table E3 - Determinants of support for Socialists / Greens in Belgium

	(1) 1970-79	(2) 1980-89	(3) 1990-99	(4) 2000-09	(5) 2010-20
Education: None/Primary	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Education: Secondary	-0.114***	-0.090***	-0.057***	-0.049**	-0.053***
	(0.015)	(0.013)	(0.014)	(0.023)	(0.020)
Education: University	-0.196***	-0.143***	-0.051***	-0.035	-0.028
	(0.021)	(0.017)	(0.017)	(0.026)	(0.022)
Income group: Bottom 50%	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Income group: Middle 40%	-0.032***	-0.063***	-0.004	-0.022	-0.012
	(0.012)	(0.011)	(0.013)	(0.015)	(0.014)
Income group: Top 10%	-0.102***	-0.144***	-0.093***	-0.076***	-0.083***
	(0.018)	(0.017)	(0.017)	(0.020)	(0.020)
Age: 20-39	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Age: 40-59	0.046***	-0.086***	0.006	0.024	0.012
	(0.016)	(0.012)	(0.012)	(0.016)	(0.016)
Age: 60+	-0.001	-0.197***	-0.059***	-0.043*	-0.032
	(0.020)	(0.015)	(0.017)	(0.023)	(0.021)
Gender: Woman	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Gender: Man	0.017	0.025**	-0.037***	-0.035***	-0.044***
	(0.011)	(0.010)	(0.011)	(0.013)	(0.012)
Religion: None	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Religion: Catholic	-0.308***	-0.212***	-0.148***	-0.126***	-0.126***
	(0.028)	(0.022)	(0.021)	(0.022)	(0.019)
Religion: Other Christian	-0.087	-0.143**	0.046	-0.087	0.014
	(0.083)	(0.064)	(0.058)	(0.069)	(0.054)
Religion: Other	-0.179*	0.088	0.033	(baseline)	-0.005
	(0.092)	(0.063)	(0.043)	(0.080)	(0.073)
Religion: Muslim				0.364***	0.320***
				(0.064)	(0.042)
Religious practice: Never	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Religious practice: Less than monthly	-0.114***	-0.054**	-0.054***	-0.039*	-0.027
	(0.026)	(0.026)	(0.019)	(0.023)	(0.021)
Religious practice: Monthly or more	-0.342***	-0.258***	-0.198***	-0.150***	-0.061**
	(0.021)	(0.025)	(0.019)	(0.027)	(0.028)
Location: Urban	(baseline)	(baseline)			
	(.)	(.)			
Location: Rural	-0.073***	-0.051***			
	(0.016)	(0.010)			
Employment status: Employed	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Employment status: Unemployed/Inactive	-0.021	0.007	0.020	0.044**	0.042**
	(0.017)	(0.012)	(0.013)	(0.019)	(0.017)
Marital status: Single	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)

Marital status: Married/With partner	0.022	-0.002	0.015	-0.003	-0.034**
	(0.016)	(0.011)	(0.012)	(0.015)	(0.014)
Race/ethnicity/language: Dutch	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Race/ethnicity/language: French	0.079***	0.055***	0.135***	-0.053*	-0.010
	(0.023)	(0.020)	(0.035)	(0.029)	(0.032)
Race/ethnicity/language: Other				-0.023	0.105
				(0.076)	(0.070)
Region: Brussels	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Region: Flanders	0.101***	-0.048***	0.068**	-0.214***	-0.187***
	(0.029)	(0.018)	(0.029)	(0.038)	(0.030)
Region: Wallonia	0.240***	0.161***	0.124***	0.054*	0.090***
	(0.028)	(0.018)	(0.023)	(0.033)	(0.030)
Constant	0.622***	0.813***	0.468***	0.588***	0.557***
	(0.044)	(0.032)	(0.036)	(0.049)	(0.042)
R-squared	0.15	0.13	0.12	0.12	0.14
Observations	22962	25787	11737	10767	10034
Clusters	11054	12947	4411	1777	1825

Note: The table reports the effect of a set of individual characteristics on the probability to vote for Socialists / Greens by decade in Belgium. The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). The number of clusters corresponds to the number of surveyed individuals in each decade. Robust standard errors clustered at the individual level.

Table E4 - Determinants of support for Liberal / NDP / Green in Canada

	(1)	(2)	(3)	(4)	(5)	(6)
	1960-69	1970-79	1980-89	1990-99	2000-09	2010-20
Education: None/Primary	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Education: Secondary	0.011	-0.016	-0.057***	-0.001	0.001	0.051**
	(0.021)	(0.025)	(0.017)	(0.022)	(0.021)	(0.021)
Education: University	0.044	0.013	-0.055**	0.046*	0.081***	0.117***
	(0.036)	(0.035)	(0.024)	(0.027)	(0.024)	(0.024)
Education: Postgraduate				0.080**	0.103***	0.145***
				(0.037)	(0.029)	(0.027)
Income group: Bottom 50%	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Income group: Middle 40%	0.072***	-0.022	-0.034**	0.002	-0.009	-0.022*
	(0.018)	(0.019)	(0.016)	(0.019)	(0.015)	(0.012)
Income group: Top 10%	0.097***	-0.096***	-0.088***	-0.031	-0.053**	-0.082***
	(0.032)	(0.030)	(0.024)	(0.030)	(0.023)	(0.020)
Age: 20-39	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Age: 40-59	-0.003	-0.039	-0.003	0.041**	0.015	-0.034**
	(0.018)	(0.025)	(0.017)	(0.018)	(0.015)	(0.014)
Age: 60+	-0.032	-0.073**	-0.021	0.073***	0.005	-0.051***
	(0.025)	(0.032)	(0.022)	(0.026)	(0.020)	(0.016)
Gender: Woman	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Gender: Man	-0.011	-0.036	-0.069***	-0.060***	-0.062***	-0.071***
	(0.017)	(0.023)	(0.015)	(0.016)	(0.013)	(0.011)
Religion: None	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Religion: Catholic	0.142**	0.139**	0.114***	0.046	-0.065***	-0.067***
	(0.061)	(0.058)	(0.036)	(0.033)	(0.025)	(0.017)
Religion: Other Christian	-0.158***	-0.105*	-0.057*	-0.085***	-0.213***	-0.151***
	(0.059)	(0.056)	(0.034)	(0.032)	(0.024)	(0.018)
Religion: Other	0.042	0.011	0.087*	0.070	-0.007	-0.030
	(0.066)	(0.063)	(0.046)	(0.048)	(0.035)	(0.027)
Religion: Muslim					0.312***	0.283***
					(0.040)	(0.047)
Religious practice: Never	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Religious practice: Less than monthly	-0.032	-0.061*	-0.056**	0.025	0.015	-0.016
	(0.039)	(0.032)	(0.024)	(0.025)	(0.019)	(0.016)
Religious practice: Monthly or more	-0.112***	-0.088***	-0.067***	0.033	-0.036*	-0.101***
	(0.039)	(0.034)	(0.025)	(0.026)	(0.021)	(0.018)
Location: Urban	(baseline)	(baseline)	(baseline)		(baseline)	(baseline)
	(.)	(.)	(.)		(.)	(.)
Location: Rural	-0.105***	-0.091***	-0.045**		-0.042	-0.078***
	(0.020)	(0.027)	(0.021)		(0.028)	(0.017)
Employment status: Employed	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Employment status: Unemployed/Inactive	0.005	-0.021	0.001	-0.000	0.026	0.028**
	(0.021)	(0.023)	(0.018)	(0.020)	(0.016)	(0.013)

Marital status: Single	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Marital status: Married/With partner	0.010	0.025	-0.058***	-0.050***	-0.056***	-0.047***
	(0.022)	(0.024)	(0.016)	(0.018)	(0.015)	(0.012)
Race/ethnicity/language: English	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Race/ethnicity/language: French	-0.078**	-0.118***	-0.023	-0.164***	-0.241***	-0.100***
	(0.033)	(0.045)	(0.031)	(0.032)	(0.031)	(0.021)
Race/ethnicity/language: Other	0.026	0.054	0.053*	0.188***	0.041	-0.041*
	(0.042)	(0.041)	(0.030)	(0.045)	(0.027)	(0.022)
Region: Eastern	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Region: Ontario	0.054*	0.007	0.012	-0.050*	-0.095***	-0.086***
	(0.029)	(0.034)	(0.023)	(0.029)	(0.025)	(0.020)
Region: Quebec	0.017	0.098**	-0.060*	-0.268***	-0.198***	-0.095***
	(0.037)	(0.044)	(0.032)	(0.035)	(0.034)	(0.024)
Region: Western	-0.018	-0.107***	-0.076***	-0.200***	-0.235***	-0.214***
	(0.031)	(0.035)	(0.024)	(0.028)	(0.025)	(0.020)
Constant	0.666***	0.741***	0.764***	0.674***	0.860***	0.838***
	(0.059)	(0.062)	(0.040)	(0.043)	(0.039)	(0.031)
R-squared	0.09	0.10	0.05	0.09	0.11	0.08
Observations	11112	7188	13319	7025	11959	20018
Clusters	2642	2381	3368	3646	5872	12260

Note: The table reports the effect of a set of individual characteristics on the probability to vote for Liberal / NDP / Green by decade in Canada. The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). The number of clusters corresponds to the number of surveyed individuals in each decade. Robust standard errors clustered at the individual level.

Table E5 - Determinants of support for Social Democratic Party / Socialist People's Party / Social Liberal Party / Red-Green Alliance in Denmark

	(1) 1960-69	(2) 1970-79	(3) 1980-89	(4) 1990-99	(5) 2000-09	(6) 2010-20
Education: None/Primary	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Education: Secondary	-0.200***	-0.163***	0.004	-0.063***	-0.002	0.057***
	(0.043)	(0.018)	(0.017)	(0.016)	(0.015)	(0.022)
Education: University	-0.249***	-0.115***	-0.018	-0.037**	0.070***	0.095***
	(0.072)	(0.024)	(0.019)	(0.018)	(0.016)	(0.025)
Income group: Bottom 50%	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Income group: Middle 40%	0.084**	0.005	-0.003	-0.031**	-0.083***	-0.029
	(0.039)	(0.016)	(0.013)	(0.013)	(0.013)	(0.020)
Income group: Top 10%	-0.041	-0.150***	-0.223***	-0.216***	-0.201***	-0.161***
	(0.058)	(0.025)	(0.022)	(0.022)	(0.021)	(0.031)
Age: 20-39	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Age: 40-59	-0.007	-0.061***	-0.113***	0.064***	0.099***	0.046**
	(0.033)	(0.017)	(0.015)	(0.015)	(0.016)	(0.019)
Age: 60+	-0.034	-0.086***	-0.156***	-0.093***	-0.001	-0.014
	(0.043)	(0.021)	(0.018)	(0.020)	(0.019)	(0.025)
Gender: Woman	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Gender: Man	0.005	0.006	0.010	-0.034**	-0.068***	-0.087***
	(0.037)	(0.013)	(0.013)	(0.014)	(0.013)	(0.016)
Religious practice: Never	(baseline)	(baseline)		(baseline)	(baseline)	
	(.)	(.)		(.)	(.)	
Religious practice: Less than monthly	-0.016	-0.031		-0.065***	-0.073***	
	(0.057)	(0.048)		(0.024)	(0.018)	
Religious practice: Monthly or more	-0.154**	-0.143**		-0.151***	-0.175***	
	(0.067)	(0.060)		(0.050)	(0.035)	
Location: Urban		(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
		(.)	(.)	(.)	(.)	(.)
Location: Rural		-0.213***	-0.111***	-0.069***	-0.077***	-0.098***
		(0.016)	(0.015)	(0.016)	(0.014)	(0.017)
Employment status: Employed	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Employment status: Unemployed/Inactive	0.057	-0.049***	-0.005	-0.008	-0.020	0.067***
	(0.040)	(0.015)	(0.018)	(0.018)	(0.015)	(0.025)
Marital status: Single	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Marital status: Married/With partner	-0.039	-0.078***	-0.002	-0.023	-0.043***	0.013
	(0.038)	(0.016)	(0.014)	(0.014)	(0.013)	(0.017)
Region: Capital	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Region: Central Jutland	-0.139***	-0.076***	-0.088***	-0.050**	-0.071***	-0.047
	(0.041)	(0.024)	(0.016)	(0.022)	(0.018)	(0.031)
Region: Northern Jutland	-0.130***	-0.103***	-0.080***	-0.036	-0.010	0.005
	(0.049)	(0.031)	(0.022)	(0.030)	(0.024)	(0.045)
Region: Southern Denmark	-0.036	-0.067***	-0.034	-0.106***	-0.068***	-0.059*
	(0.048)	(0.024)	(0.022)	(0.022)	(0.018)	(0.032)

Region: Zealand	-0.037	-0.019	-0.017	-0.041*	-0.061***	-0.038
	(0.044)	(0.026)	(0.021)	(0.024)	(0.021)	(0.038)
Constant	0.752***	1.072***	0.754***	0.742***	0.612***	0.556***
	(0.078)	(0.055)	(0.021)	(0.030)	(0.026)	(0.038)
R-squared	0.10	0.09	0.07	0.05	0.05	0.04
Observations	11059	22837	24186	23048	20258	7069
Clusters	1137	1923	3809	2028	3987	2174

Note: The table reports the effect of a set of individual characteristics on the probability to vote for Social Democratic Party / Socialist People's Party / Social Liberal Party / Red-Green Alliance by decade in Denmark. The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). The number of clusters corresponds to the number of surveyed individuals in each decade. Robust standard errors clustered at the individual level.

Table E6 - Determinants of support for Social Democratic Party / Finnish People's Democratic League / Left Alliance / Green League in Finland

	(1)	(2)	(3)	(4)	(5)
	1970-79	1980-89	1990-99	2000-09	2010-20
Education: None/Primary	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Education: Secondary	-0.187***	-0.159***	-0.113***	-0.070***	-0.041
	(0.018)	(0.018)	(0.016)	(0.025)	(0.038)
Education: University	-0.337***	-0.261***	-0.188***	-0.131***	-0.086**
	(0.032)	(0.026)	(0.021)	(0.029)	(0.041)
Income group: Bottom 50%	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Income group: Middle 40%	0.022	-0.067***	-0.006	-0.057**	-0.021
	(0.021)	(0.017)	(0.014)	(0.022)	(0.023)
Income group: Top 10%	-0.115***	-0.193***	-0.089***	-0.087***	-0.077***
	(0.028)	(0.024)	(0.021)	(0.028)	(0.027)
Age: 20-39	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Age: 40-59	-0.050***	-0.044***	0.052***	0.027	-0.046*
	(0.019)	(0.017)	(0.014)	(0.021)	(0.027)
Age: 60+	-0.081***	-0.086***	-0.079***	-0.055**	-0.092***
	(0.025)	(0.022)	(0.019)	(0.026)	(0.026)
Gender: Woman	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Gender: Man	0.071***	0.016	-0.049***	0.003	-0.042**
	(0.016)	(0.015)	(0.012)	(0.018)	(0.021)
Location: Urban			(baseline)	(baseline)	(baseline)
			(.)	(.)	(.)
Location: Rural			-0.141***	-0.121***	-0.086***
			(0.017)	(0.019)	(0.023)
Employment status: Employed		(baseline)	(baseline)	(baseline)	(baseline)
		(.)	(.)	(.)	(.)
Employment status: Unemployed/Inactive		-0.049***	0.020	-0.004	0.043*
		(0.019)	(0.015)	(0.021)	(0.024)
Marital status: Single	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Marital status: Married/With partner	0.039	0.023	-0.005	-0.017	-0.023
	(0.037)	(0.019)	(0.015)	(0.022)	(0.023)
Region: Central Finland	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Region: Northern Finland	-0.036	0.001	0.055***	-0.078**	-0.038
	(0.027)	(0.029)	(0.020)	(0.031)	(0.042)
Region: Southern Finland	0.094***	0.078***	0.145***	0.059***	0.062*
	(0.019)	(0.018)	(0.015)	(0.019)	(0.032)
Constant	0.501***	0.534***	0.517***	0.553***	0.502***
	(0.046)	(0.026)	(0.027)	(0.035)	(0.049)
R-squared	0.07	0.05	0.06	0.04	0.03
Observations	7403	9839	11737	7665	5175
Clusters	1358	1196	2480	1562	1442

Note: The table reports the effect of a set of individual characteristics on the probability to vote for left-wing parties by decade in Finland. The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). The number of clusters corresponds to the number of surveyed individuals in each decade. Robust standard errors clustered at the individual level.

Table E7 - Determinants of support for PS / PCF / Radicaux / Other left in France

	(1) 1950-59	(2) 1960-69	(3) 1970-79	(4) 1980-89	(5) 1990-99	(6) 2000-09	(7) 2010-20
Education: None/Primary	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Education: Secondary	-0.089***	-0.054***	-0.048***	-0.017	-0.009	0.002	0.046*
	(0.026)	(0.018)	(0.011)	(0.014)	(0.013)	(0.015)	(0.025)
Education: University	-0.238***	-0.078**	-0.094***	-0.031	0.081***	0.062***	0.177***
	(0.055)	(0.037)	(0.017)	(0.020)	(0.020)	(0.021)	(0.030)
Income group: Bottom 50%	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Income group: Middle 40%	-0.030	-0.022	-0.014	-0.071***	-0.007	-0.011	-0.015
	(0.030)	(0.016)	(0.011)	(0.012)	(0.014)	(0.011)	(0.019)
Income group: Top 10%	-0.015	-0.133***	-0.153***	-0.160***	-0.090***	-0.065***	-0.098***
	(0.056)	(0.031)	(0.017)	(0.020)	(0.020)	(0.017)	(0.030)
Age: 20-39	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Age: 40-59	-0.033	-0.029	-0.083***	-0.064***	0.022	0.036***	0.058***
	(0.025)	(0.018)	(0.011)	(0.015)	(0.014)	(0.014)	(0.021)
Age: 60+	-0.092***	-0.069***	-0.157***	-0.068***	-0.012	0.002	0.020
	(0.033)	(0.023)	(0.014)	(0.019)	(0.017)	(0.018)	(0.028)
Gender: Woman	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Gender: Man	0.168***	0.070***	0.014	-0.040***	-0.019*	-0.013	-0.016
	(0.026)	(0.017)	(0.010)	(0.013)	(0.011)	(0.011)	(0.017)
Religion: None		(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
		(.)	(.)	(.)	(.)	(.)	(.)
Religion: Catholic		-0.227***	-0.299***	-0.305***	-0.243***	-0.207***	-0.189***
		(0.041)	(0.012)	(0.019)	(0.015)	(0.014)	(0.020)
Religion: Other Christian		-0.033	-0.323***	-0.357***	-0.320***	-0.308***	-0.214***
		(0.073)	(0.033)	(0.046)	(0.037)	(0.037)	(0.062)
Religion: Other			-0.273***	-0.143***	-0.121**	-0.096**	-0.026
			(0.044)	(0.055)	(0.051)	(0.046)	(0.059)
Religion: Muslim				-0.135	0.261***	0.207***	0.281***
				(0.111)	(0.052)	(0.043)	(0.037)
Location: Urban	(baseline)	(baseline)	(baseline)		(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)		(.)	(.)	(.)
Location: Rural	0.012	-0.024	-0.052***		-0.024	-0.052***	-0.095***
	(0.025)	(0.017)	(0.011)		(0.025)	(0.012)	(0.019)
Employment status: Employed	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Employment status: Unemployed/Inactive	0.035	-0.009	-0.003	-0.008	-0.013	0.003	0.034
	(0.029)	(0.020)	(0.011)	(0.015)	(0.014)	(0.015)	(0.022)
Marital status: Single	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Marital status: Married/With partner	0.102***	-0.001	0.033***	0.002	0.011	-0.028**	-0.038*
	(0.034)	(0.022)	(0.012)	(0.015)	(0.013)	(0.012)	(0.019)
Region: Auvergne-Rhone-Alpes	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Region: Bourgogne-Franche-Comte	0.054	-0.006	-0.073***	0.031	0.051*	-0.003	-0.028
	(0.066)	(0.037)	(0.027)	(0.031)	(0.028)	(0.049)	(0.078)
Region: Bretagne	-0.058	-0.062*	-0.058**	0.129***	0.046	0.041	0.009
	(0.058)	(0.037)	(0.026)	(0.033)	(0.029)	(0.048)	(0.062)
Region: Centre-Val de Loire	-0.057	-0.007	-0.103***	-0.032	0.017	-0.072	-0.022
	(0.060)	(0.047)	(0.024)	(0.034)	(0.032)	(0.052)	(0.076)
Region: Grand Est	-0.162***	-0.133***	-0.035*	-0.010	0.023	0.008	-0.154***
	(0.052)	(0.036)	(0.020)	(0.028)	(0.022)	(0.039)	(0.058)
Region: Hauts-de-France	0.035	0.014	-0.043**	0.060**	0.036*	-0.042	-0.135**
	(0.045)	(0.034)	(0.020)	(0.027)	(0.022)	(0.038)	(0.056)

Region: Ile-de-France	0.051	0.003	0.021	-0.022	-0.016	-0.020	-0.063
	(0.068)	(0.038)	(0.018)	(0.025)	(0.021)	(0.037)	(0.052)
Region: Normandie	-0.138**	-0.147***	-0.050**	-0.041	0.044*	-0.033	0.007
	(0.055)	(0.038)	(0.021)	(0.032)	(0.025)	(0.047)	(0.074)
Region: Nouvelle-Aquitaine	0.046	0.058	0.025	0.066**	0.018	-0.039	-0.037
	(0.047)	(0.036)	(0.020)	(0.028)	(0.023)	(0.040)	(0.059)
Region: Occitanie	0.155***	0.149***	0.112***	0.034	0.015	-0.011	0.036
	(0.052)	(0.039)	(0.021)	(0.027)	(0.023)	(0.040)	(0.056)
Region: PACA	0.148**	0.064	0.014	-0.060**	-0.094***	-0.081*	-0.170***
	(0.064)	(0.045)	(0.023)	(0.029)	(0.025)	(0.042)	(0.063)
Region: Paris	0.071	0.060	-0.007	-0.028	-0.084**	0.100	-0.066
	(0.051)	(0.040)	(0.024)	(0.037)	(0.034)	(0.062)	(0.093)
Region: Pays de la Loire	-0.152***	-0.181***	-0.045*	0.001	0.003	0.076*	-0.039
	(0.056)	(0.039)	(0.024)	(0.029)	(0.025)	(0.044)	(0.064)
Constant	0.550***	0.732***	0.948***	0.874***	0.793***	0.645***	0.646***
	(0.072)	(0.053)	(0.022)	(0.031)	(0.037)	(0.036)	(0.051)
R-squared	0.09	0.12	0.18	0.11	0.09	0.11	0.12
Observations	3650	9522	20668	15563	17578	18054	7122
Clusters	1339	1936	4474	3819	3964	3953	2457

Note: The table reports the effect of a set of individual characteristics on the probability to vote for left-wing parties (PS, PCF, Radicaux, etc.) by decade in France. The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). The number of clusters corresponds to the number of surveyed individuals in each decade. Robust standard errors clustered at the individual level.

Table E8 - Determinants of support for SPD / Die Grünen / Die Linke in Germany

	(1) 1950-59	(2) 1960-69	(3) 1970-79	(4) 1980-89	(5) 1990-99	(6) 2000-09	(7) 2010-20
Education: None/Primary	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Education: Secondary	-0.155***	-0.139***	-0.166***	-0.066**	-0.043*	-0.090***	-0.008
	(0.021)	(0.025)	(0.028)	(0.026)	(0.023)	(0.021)	(0.028)
Education: University	-0.172***	-0.189***	-0.263***	-0.090**	0.032	-0.005	0.072**
	(0.031)	(0.038)	(0.042)	(0.039)	(0.028)	(0.024)	(0.033)
Education: Postgraduate						-0.056	0.115***
						(0.042)	(0.034)
Income group: Bottom 50%	(baseline)	(baseline)	(baseline)		(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)		(.)	(.)	(.)
Income group: Middle 40%	-0.009	-0.013	-0.026		-0.061**	-0.032*	-0.045*
	(0.014)	(0.019)	(0.023)		(0.025)	(0.018)	(0.023)
Income group: Top 10%	-0.119***	-0.187***	-0.136***		-0.143***	-0.118***	-0.167***
	(0.022)	(0.027)	(0.042)		(0.044)	(0.030)	(0.037)
Age: 20-39	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Age: 40-59	-0.013	-0.030	-0.059**	-0.066***	-0.044*	0.031	0.041
	(0.018)	(0.019)	(0.025)	(0.026)	(0.023)	(0.022)	(0.029)
Age: 60+	-0.071***	-0.096***	-0.088***	-0.097***	-0.132***	-0.064***	0.015
	(0.022)	(0.025)	(0.029)	(0.028)	(0.027)	(0.023)	(0.030)
Gender: Woman	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Gender: Man	0.122***	0.077***	0.003	-0.041*	-0.053***	-0.013	-0.029
	(0.016)	(0.018)	(0.022)	(0.021)	(0.020)	(0.017)	(0.021)
Religion: None	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Religion: Catholic	-0.298***	-0.116	-0.191***	-0.170***	-0.131***	-0.207***	-0.091***
	(0.044)	(0.139)	(0.053)	(0.052)	(0.035)	(0.030)	(0.034)
Religion: Other Christian	-0.196***	0.063	-0.016	-0.077	-0.053*	-0.047*	-0.022
	(0.043)	(0.139)	(0.052)	(0.050)	(0.031)	(0.026)	(0.030)
Religion: Other	-0.157*	-0.443***	-0.166	0.024	0.103	0.106*	0.129
	(0.083)	(0.139)	(0.107)	(0.128)	(0.100)	(0.062)	(0.080)
Religious practice: Never	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Religious practice: Less than monthly	-0.138***	-0.155***	-0.201***	-0.092***	-0.130***	-0.087***	-0.052*
	(0.022)	(0.036)	(0.038)	(0.026)	(0.026)	(0.022)	(0.029)
Religious practice: Monthly or more	-0.301***	-0.270***	-0.381***	-0.288***	-0.301***	-0.201***	-0.190***
	(0.021)	(0.037)	(0.042)	(0.032)	(0.033)	(0.034)	(0.037)
Region: East					(baseline)	(baseline)	(baseline)
					(.)	(.)	(.)
Region: West					0.093***	-0.008	0.045*
					(0.025)	(0.022)	(0.024)
Constant	0.738***	0.605***	0.851***	0.788***	0.777***	0.728***	0.514***
	(0.045)	(0.143)	(0.054)	(0.051)	(0.032)	(0.029)	(0.039)
R-squared	0.13	0.11	0.16	0.08	0.10	0.08	0.04
Observations	15983	5837	4993	3034	5849	9169	6293
Clusters	4705	2958	2155	3034	3937	4726	3131

Note: The table reports the effect of a set of individual characteristics on the probability to vote for SPD / Die Grünen / Die Linke by decade in Germany. The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). The number of clusters corresponds to the number of surveyed individuals in each decade. Robust standard errors clustered at the individual level.

Table E9 - Determinants of support for Social Democratic Alliance / Left-Green movement in Iceland

	(1) 1970-79	(2) 1980-89	(3) 1990-99	(4) 2000-09	(5) 2010-20
Education: None/Primary	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Education: Secondary	0.084*	0.028	0.022	-0.022	0.029
	(0.048)	(0.024)	(0.018)	(0.020)	(0.019)
Education: University	0.125*	0.014	0.111***	0.075***	0.084***
	(0.073)	(0.043)	(0.027)	(0.024)	(0.020)
Income group: Bottom 50%		(baseline)	(baseline)	(baseline)	(baseline)
		(.)	(.)	(.)	(.)
Income group: Middle 40%		-0.045*	-0.040**	-0.091***	-0.042**
		(0.027)	(0.017)	(0.018)	(0.016)
Income group: Top 10%		-0.087**	-0.035	-0.122***	-0.097***
		(0.039)	(0.025)	(0.025)	(0.022)
Age: 20-39	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Age: 40-59	-0.107**	-0.075***	-0.021	0.013	0.010
	(0.043)	(0.024)	(0.018)	(0.019)	(0.017)
Age: 60+	-0.105**	-0.142***	0.011	-0.002	0.074***
	(0.053)	(0.029)	(0.023)	(0.024)	(0.019)
Gender: Woman	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Gender: Man	-0.041	-0.067***	-0.110***	-0.112***	-0.083***
	(0.044)	(0.022)	(0.016)	(0.017)	(0.016)
Location: Urban	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Location: Rural	0.007	-0.014	0.003	-0.140***	-0.059**
	(0.065)	(0.036)	(0.029)	(0.029)	(0.027)
Employment status: Employed	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Employment status: Unemployed/Inactive	-0.020	-0.109***	-0.090***	-0.082***	-0.125***
	(0.047)	(0.026)	(0.020)	(0.020)	(0.019)
Marital status: Single	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Marital status: Married/With partner	0.085*	-0.024	-0.041**	0.008	-0.053***
	(0.048)	(0.026)	(0.019)	(0.021)	(0.018)
Region: Capital area	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Region: East	-0.148	-0.081	-0.018	-0.043	0.012
	(0.105)	(0.056)	(0.045)	(0.046)	(0.047)
Region: Northeast	-0.030	-0.053	-0.007	0.108***	0.079**
	(0.082)	(0.046)	(0.034)	(0.036)	(0.034)
Region: Northwest	-0.236**	-0.232***	-0.142***	-0.021	-0.017
	(0.108)	(0.057)	(0.045)	(0.050)	(0.049)
Region: South	-0.207**	-0.171***	-0.059	0.026	0.005
	(0.085)	(0.044)	(0.036)	(0.039)	(0.033)
Region: Sudurnes	-0.007	-0.145***	0.036	0.043	-0.027
	(0.098)	(0.053)	(0.042)	(0.046)	(0.038)
Region: West	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)

Constant	0.505***	0.688***	0.568***	0.683***	0.474***
	(0.066)	(0.038)	(0.026)	(0.028)	(0.026)
R-squared	0.06	0.06	0.05	0.06	0.05
Observations	716	4498	9618	9245	9516
Clusters	716	1598	1688	1550	1981

Note: The table reports the effect of a set of individual characteristics on the probability to vote for Social Democratic Alliance / Left-Green movement by decade in Iceland. The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). The number of clusters corresponds to the number of surveyed individuals in each decade. Robust standard errors clustered at the individual level.

Table E10 - Determinants of support for Fianna Fáil / Sinn Féin / Other left-wing parties in Ireland

	(1) 1970-79	(2) 1980-89	(3) 1990-99	(4) 2000-09	(5) 2010-20
Education: None/Primary	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Education: Secondary	-0.099***	-0.073***	-0.050***	-0.064***	-0.095***
	(0.015)	(0.010)	(0.016)	(0.018)	(0.026)
Education: University	-0.188***	-0.205***	-0.120***	-0.128***	-0.151***
	(0.035)	(0.020)	(0.026)	(0.022)	(0.028)
Income group: Bottom 50%	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Income group: Middle 40%	-0.004	-0.061***	-0.044***	-0.029	-0.048***
	(0.015)	(0.010)	(0.016)	(0.017)	(0.017)
Income group: Top 10%	-0.069***	-0.118***	-0.125***	-0.041	-0.095***
	(0.025)	(0.015)	(0.026)	(0.032)	(0.026)
Age: 20-39	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Age: 40-59	-0.085***	-0.056***	-0.042***	-0.043***	-0.034*
	(0.016)	(0.010)	(0.013)	(0.017)	(0.018)
Age: 60+	-0.092***	-0.071***	-0.074***	-0.090***	-0.056***
	(0.019)	(0.012)	(0.016)	(0.020)	(0.021)
Gender: Woman	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Gender: Man	0.027*	0.056***	0.027**	0.036***	0.004
	(0.015)	(0.009)	(0.011)	(0.014)	(0.014)
Religion: None	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Religion: Catholic	-0.102	0.063	0.061	0.010	-0.046*
	(0.099)	(0.048)	(0.046)	(0.042)	(0.027)
Religion: Other Christian	-0.416***	-0.269***	-0.301***	-0.185***	-0.203***
	(0.105)	(0.056)	(0.057)	(0.055)	(0.045)
Religion: Other	-0.346***	-0.147*	0.066	0.217**	0.008
	(0.133)	(0.083)	(0.086)	(0.101)	(0.057)
Religion: Muslim				0.131	0.048
				(0.124)	(0.168)
Religious practice: Never	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Religious practice: Less than monthly	-0.011	-0.064	-0.065*	-0.047	-0.045
	(0.089)	(0.043)	(0.037)	(0.040)	(0.028)
Religious practice: Monthly or more	-0.013	-0.046	-0.112***	-0.061	-0.073***
	(0.085)	(0.039)	(0.034)	(0.040)	(0.028)
Location: Urban	(baseline)	(baseline)	(baseline)		
	(.)	(.)	(.)		
Location: Rural	-0.103***	-0.058***	-0.043***		
	(0.014)	(0.008)	(0.013)		
Employment status: Employed	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Employment status: Unemployed/Inactive	0.010	0.048***	0.061***	0.025	0.036**
	(0.017)	(0.010)	(0.013)	(0.015)	(0.016)
Marital status: Single	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)

Marital status: Married/With partner	-0.006	0.008	-0.007	-0.013	-0.033**
	(0.014)	(0.009)	(0.012)	(0.014)	(0.016)
Region: Border				(baseline)	(baseline)
				(.)	(.)
Region: Dublin				0.062**	-0.038
				(0.029)	(0.026)
Region: Mid-East				-0.023	-0.046
				(0.036)	(0.030)
Region: Mid-West				-0.049	-0.069**
				(0.035)	(0.032)
Region: Midlands				-0.020	-0.081**
				(0.040)	(0.036)
Region: South-East				-0.001	-0.049
				(0.035)	(0.030)
Region: South-West				0.002	-0.026
				(0.031)	(0.028)
Region: West				-0.129***	-0.151***
				(0.034)	(0.030)
Constant	0.959***	0.730***	0.783***	0.776***	0.849***
	(0.058)	(0.043)	(0.043)	(0.039)	(0.040)
R-squared	0.04	0.03	0.03	0.03	0.06
Observations	17708	31395	18108	12435	16099
Clusters	8254	18359	12790	2384	2678

Note: The table reports the effect of a set of individual characteristics on the probability to vote for Fianna Fáil / Sinn Féin / Labour / Other left by decade in Ireland. The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). The number of clusters corresponds to the number of surveyed individuals in each decade. Robust standard errors clustered at the individual level.

Table E11 - Determinants of support for Social Democrats / Socialists / Communists / Greens in Italy

	(1) 1950-59	(2) 1960-69	(3) 1970-79	(4) 1980-89	(5) 1990-99	(6) 2000-09	(7) 2010-20
Education: None/Primary	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Education: Secondary	-0.123**	-0.050**	-0.034	-0.026	-0.014	-0.011	0.070**
	(0.048)	(0.024)	(0.034)	(0.030)	(0.028)	(0.023)	(0.032)
Education: University	-0.243**	-0.030	-0.135**	0.012	0.043	0.041	0.137***
	(0.098)	(0.050)	(0.064)	(0.042)	(0.033)	(0.033)	(0.036)
Income group: Bottom 50%	(baseline)	(baseline)	(baseline)	(baseline)		(baseline)	(baseline)
	(.)	(.)	(.)	(.)		(.)	(.)
Income group: Middle 40%	0.075*	-0.040**	0.016	0.013		-0.059	0.089***
	(0.044)	(0.017)	(0.028)	(0.024)		(0.057)	(0.018)
Income group: Top 10%	0.063	-0.096***	-0.008	0.012		-0.076	0.097***
	(0.091)	(0.035)	(0.049)	(0.042)		(0.070)	(0.020)
Age: 20-39	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Age: 40-59	-0.124**	-0.036*	-0.011	-0.026	0.020	0.018	0.047**
	(0.050)	(0.022)	(0.032)	(0.028)	(0.027)	(0.021)	(0.021)
Age: 60+	-0.080	-0.112***	-0.079**	-0.055	-0.023	0.017	0.150***
	(0.065)	(0.027)	(0.038)	(0.036)	(0.038)	(0.025)	(0.026)
Gender: Woman	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Gender: Man	0.293***	0.104***	0.023	0.014	0.014	0.010	-0.025
	(0.052)	(0.023)	(0.032)	(0.023)	(0.021)	(0.017)	(0.018)
Religion: None	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Religion: Catholic	0.047	-0.161***	-0.159***	-0.106*	-0.165***	-0.178***	-0.222***
	(0.179)	(0.045)	(0.055)	(0.065)	(0.043)	(0.043)	(0.065)
Religion: Other	-0.186			0.071	-0.061	-0.081	-0.419***
	(0.197)			(0.060)	(0.045)	(0.073)	(0.106)
Religious practice: Never	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Religious practice: Less than monthly	-0.027	-0.067	-0.073	-0.130***	-0.137***	-0.140***	-0.053**
	(0.082)	(0.042)	(0.049)	(0.040)	(0.035)	(0.026)	(0.022)
Religious practice: Monthly or more	-0.425***	-0.377***	-0.424***	-0.433***	-0.208***	-0.183***	-0.107***
	(0.074)	(0.042)	(0.048)	(0.039)	(0.033)	(0.024)	(0.024)
Location: Urban	(baseline)	(baseline)			(baseline)	(baseline)	(baseline)
	(.)	(.)			(.)	(.)	(.)
Location: Rural	-0.141***	-0.029			0.006	-0.007	-0.024
	(0.046)	(0.023)			(0.040)	(0.021)	(0.042)
Employment status: Employed	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Employment status: Unemployed/Inactive	0.109**	-0.062***	-0.069**	0.014	-0.021	0.004	-0.013
	(0.050)	(0.024)	(0.033)	(0.025)	(0.022)	(0.019)	(0.020)
Marital status: Single	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Marital status: Married/With partner	-0.016	0.053**	0.007	0.016	0.010	0.020	-0.001
	(0.050)	(0.022)	(0.030)	(0.032)	(0.027)	(0.018)	(0.021)
Region: Center	(baseline)	(baseline)		(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)		(.)	(.)	(.)	(.)
Region: Islands	-0.270***	-0.139***		-0.315***	-0.066	-0.091**	0.005
	(0.082)	(0.035)		(0.073)	(0.043)	(0.040)	(0.034)
Region: North	0.021	-0.031		-0.044	-0.114***	-0.049*	0.021
	(0.059)	(0.026)		(0.046)	(0.031)	(0.029)	(0.025)
Region: South	-0.219***	-0.136***		-0.180***	-0.072**	-0.031	0.004
	(0.065)	(0.029)		(0.054)	(0.035)	(0.033)	(0.029)
Constant	0.696***	0.985***	0.948***	0.954***	0.838***	0.799***	0.478***
	(0.179)	(0.042)	(0.052)	(0.068)	(0.058)	(0.074)	(0.074)

R-squared	0.34	0.21	0.21	0.18	0.07	0.04	0.08
Observations	2197	7780	3333	5608	4243	5268	12033
Clusters	523	2422	1238	1602	2406	2867	2045

Note: The table reports the effect of a set of individual characteristics on the probability to vote for Social Democrats / Socialists / Communists / Greens by decade in Italy. The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). The number of clusters corresponds to the number of surveyed individuals in each decade. Robust standard errors clustered at the individual level.

Table E12 - Determinants of support for LSAP / Greens / Other left in Luxembourg

	(1) 1970-79	(2) 1980-89	(3) 1990-99	(4) 2000-09	(5) 2010-20
Education: None/Primary	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Education: Secondary	-0.103***	-0.135***	-0.077***	0.099	-0.007
	(0.027)	(0.026)	(0.027)	(0.064)	(0.071)
Education: University	-0.228***	-0.190***	-0.061**	0.082	0.050
	(0.040)	(0.033)	(0.031)	(0.062)	(0.070)
Income group: Bottom 50%	(baseline)	(baseline)	(baseline)	(baseline)	
	(.)	(.)	(.)	(.)	
Income group: Middle 40%	-0.051**	-0.076***	-0.060***	-0.028	
	(0.022)	(0.018)	(0.018)	(0.047)	
Income group: Top 10%	-0.122***	-0.103***	-0.079***	-0.185***	
	(0.033)	(0.026)	(0.021)	(0.067)	
Age: 20-39	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Age: 40-59	-0.031	-0.117***	-0.095***	0.104*	0.010
	(0.024)	(0.018)	(0.017)	(0.055)	(0.058)
Age: 60+	-0.110***	-0.228***	-0.202***	-0.019	-0.067
	(0.035)	(0.025)	(0.020)	(0.073)	(0.066)
Gender: Woman	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Gender: Man	0.089***	0.044***	-0.014	0.003	0.049
	(0.023)	(0.017)	(0.015)	(0.042)	(0.043)
Religion: None	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Religion: Catholic	-0.077	-0.181***	-0.204***	-0.032	-0.124
	(0.055)	(0.044)	(0.041)	(0.074)	(0.078)
Religion: Other Christian	0.102	0.019	0.038	0.091	0.113
	(0.137)	(0.147)	(0.107)	(0.085)	(0.130)
Religion: Other	-0.015	0.087	-0.121	0.377***	-0.072
	(0.204)	(0.109)	(0.101)	(0.086)	(0.134)
Religious practice: Never	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Religious practice: Less than monthly	-0.168***	-0.112***	-0.115***	-0.239***	-0.133*
	(0.040)	(0.032)	(0.034)	(0.067)	(0.071)
Religious practice: Monthly or more	-0.379***	-0.353***	-0.286***	-0.333***	-0.200**
	(0.040)	(0.033)	(0.035)	(0.078)	(0.080)
Location: Urban	(baseline)	(baseline)	(baseline)		(baseline)
	(.)	(.)	(.)		(.)
Location: Rural	-0.114***	-0.081***	-0.034*		-0.066
	(0.019)	(0.015)	(0.018)		(0.047)
Employment status: Employed	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Employment status: Unemployed/Inactive	0.044*	0.027	0.014	0.019	-0.028
	(0.026)	(0.020)	(0.017)	(0.059)	(0.068)
Marital status: Single	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Marital status: Married/With partner	0.006	0.011	0.038**	0.052	0.065
	(0.023)	(0.018)	(0.016)	(0.046)	(0.074)

Region: Centre					(baseline)
					(.)
Region: East					0.134*
					(0.075)
Region: North					0.104
					(0.073)
Region: South					0.091*
					(0.051)
Constant	0.821***	0.996***	0.847***	0.411***	0.518***
	(0.059)	(0.049)	(0.047)	(0.086)	(0.095)
R-squared	0.10	0.10	0.07	0.12	0.10
Observations	7744	8821	10633	1705	692
Clusters	3561	4761	6229	761	466

Note: The table reports the effect of a set of individual characteristics on the probability to vote for LSAP / Greens / Other left in Luxembourg. The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). The number of clusters corresponds to the number of surveyed individuals in each decade. Robust standard errors clustered at the individual level.

Table E13 - Determinants of support for PvdA / D66 / Greens / Other left in the Netherlands

	(1) 1960-69	(2) 1970-79	(3) 1980-89	(4) 1990-99	(5) 2000-09	(6) 2010-20
Education: None/Primary	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Education: Secondary	-0.088***	-0.093***	-0.136***	-0.136***	0.058***	-0.004
	(0.022)	(0.016)	(0.015)	(0.026)	(0.018)	(0.020)
Education: University	-0.077*	-0.127***	-0.099***	-0.008	0.154***	0.100***
	(0.042)	(0.026)	(0.021)	(0.031)	(0.024)	(0.022)
Education: Postgraduate					0.194***	0.165***
					(0.042)	(0.037)
Income group: Bottom 50%	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Income group: Middle 40%	-0.014	-0.068***	-0.110***	-0.071***	-0.028*	-0.028*
	(0.024)	(0.018)	(0.014)	(0.020)	(0.017)	(0.014)
Income group: Top 10%	-0.204***	-0.226***	-0.234***	-0.184***	-0.168***	-0.103***
	(0.037)	(0.029)	(0.023)	(0.028)	(0.026)	(0.020)
Age: 20-39	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Age: 40-59	-0.055**	-0.023	-0.047***	0.020	0.102***	0.065***
	(0.023)	(0.017)	(0.014)	(0.020)	(0.019)	(0.017)
Age: 60+	-0.064**	-0.077***	-0.101***	-0.131***	0.028	0.014
	(0.030)	(0.021)	(0.019)	(0.026)	(0.025)	(0.021)
Gender: Woman	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Gender: Man	0.021	-0.032*	0.007	-0.087***	-0.061***	-0.068***
	(0.026)	(0.017)	(0.014)	(0.018)	(0.015)	(0.013)
Religion: None	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Religion: Catholic	-0.229***	-0.131***	-0.133***	-0.131***	-0.154***	-0.108***
	(0.046)	(0.031)	(0.023)	(0.032)	(0.026)	(0.021)
Religion: Other Christian	-0.128***	-0.129***	-0.138***	-0.195***	-0.194***	-0.182***
	(0.042)	(0.027)	(0.022)	(0.030)	(0.025)	(0.023)
Religion: Other	0.008	-0.076	0.002	0.026	-0.050	-0.031
	(0.073)	(0.055)	(0.034)	(0.049)	(0.042)	(0.035)
Religion: Muslim					0.537***	0.364***
					(0.056)	(0.062)
Religious practice: Never	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Religious practice: Less than monthly	-0.087*	-0.168***	-0.158***	-0.071**	-0.076**	0.011
	(0.052)	(0.034)	(0.025)	(0.035)	(0.030)	(0.029)
Religious practice: Monthly or more	-0.388***	-0.440***	-0.345***	-0.242***	-0.206***	-0.150***
	(0.039)	(0.026)	(0.019)	(0.028)	(0.023)	(0.022)
Location: Urban	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Location: Rural	-0.023	-0.053***	-0.041***	-0.034*	-0.022	-0.060***
	(0.021)	(0.016)	(0.014)	(0.018)	(0.016)	(0.015)
Employment status: Employed	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Employment status: Unemployed/Inactive	-0.002	-0.022	0.010	0.027	0.011	0.044**
	(0.028)	(0.019)	(0.016)	(0.022)	(0.020)	(0.017)

Marital status: Single	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Marital status: Married/With partner	0.009	-0.004	0.023	-0.024	-0.073***	-0.059***
	(0.028)	(0.023)	(0.014)	(0.019)	(0.017)	(0.014)
Region: East	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Region: North	-0.044	0.017	0.033	0.073*	0.078***	0.077***
	(0.034)	(0.025)	(0.021)	(0.040)	(0.027)	(0.024)
Region: South	0.018	0.048**	0.024	0.043	0.012	0.012
	(0.029)	(0.023)	(0.019)	(0.035)	(0.023)	(0.020)
Region: West	-0.027	-0.019	-0.024	-0.038	-0.033*	-0.007
	(0.028)	(0.019)	(0.016)	(0.031)	(0.019)	(0.017)
Constant	0.818***	0.947***	0.846***	0.858***	0.526***	0.529***
	(0.053)	(0.033)	(0.025)	(0.040)	(0.029)	(0.028)
R-squared	0.31	0.29	0.21	0.17	0.16	0.09
Observations	3025	7479	10041	5401	6139	9884
Clusters	1753	2186	2110	1956	2590	3215

Note: The table reports the effect of a set of individual characteristics on the probability to vote for PvdA / D66 / Greens / Other left by decade in the Netherlands. The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). The number of clusters corresponds to the number of surveyed individuals in each decade. Robust standard errors clustered at the individual level.

Table E14 - Determinants of support for Labour / Greens / Other left in New Zealand

	(1) 1970-79	(2) 1980-89	(3) 1990-99	(4) 2000-09	(5) 2010-20
Education: None/Primary	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Education: Secondary	-0.096***	-0.080***	-0.029***	-0.020	-0.032*
	(0.021)	(0.025)	(0.011)	(0.016)	(0.018)
Education: University	-0.073**	-0.119***	-0.019	0.031	0.072***
	(0.030)	(0.039)	(0.017)	(0.022)	(0.023)
Education: Postgraduate				0.151***	0.159***
				(0.028)	(0.029)
Income group: Bottom 50%	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Income group: Middle 40%	-0.015	-0.012	-0.043***	-0.124***	-0.098***
	(0.023)	(0.021)	(0.013)	(0.015)	(0.017)
Income group: Top 10%	-0.208***	-0.072**	-0.132***	-0.169***	-0.183***
	(0.032)	(0.035)	(0.019)	(0.027)	(0.027)
Age: 20-39	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Age: 40-59	0.002	-0.067***	0.038***	0.013	-0.008
	(0.021)	(0.025)	(0.012)	(0.017)	(0.021)
Age: 60+	-0.016	-0.101***	0.029*	-0.010	-0.105***
	(0.027)	(0.032)	(0.015)	(0.021)	(0.021)
Gender: Woman	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Gender: Man	0.056***	0.001	-0.051***	-0.086***	-0.043***
	(0.021)	(0.023)	(0.010)	(0.014)	(0.015)
Religion: None	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Religion: Catholic	0.051	-0.024	-0.010	-0.038	-0.062**
	(0.038)	(0.060)	(0.019)	(0.027)	(0.029)
Religion: Other Christian	-0.090***	-0.082*	-0.105***	-0.111***	-0.101***
	(0.031)	(0.049)	(0.014)	(0.020)	(0.020)
Religion: Other	-0.055	0.173*	-0.014	0.007	0.034
	(0.049)	(0.091)	(0.029)	(0.035)	(0.036)
Religion: Muslim				0.154	0.305***
				(0.199)	(0.082)
Religious practice: Never	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Religious practice: Less than monthly	-0.059*	-0.025	-0.042***	-0.002	0.008
	(0.032)	(0.048)	(0.013)	(0.019)	(0.021)
Religious practice: Monthly or more	-0.079**	-0.048	-0.123***	-0.128***	-0.016
	(0.032)	(0.049)	(0.014)	(0.020)	(0.022)
Location: Urban			(baseline)	(baseline)	(baseline)
			(.)	(.)	(.)
Location: Rural			-0.089***	-0.098***	-0.077***
			(0.012)	(0.017)	(0.022)
Employment status: Employed	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Employment status: Unemployed/Inactive	-0.007	0.007	0.064***	0.030*	0.060***
	(0.023)	(0.026)	(0.013)	(0.017)	(0.017)

Marital status: Single	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Marital status: Married/With partner	0.020	0.029	-0.073***	-0.037**	-0.037**
	(0.022)	(0.023)	(0.011)	(0.017)	(0.019)
Race/ethnicity/language: European	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)
Race/ethnicity/language: Maori	0.293***	0.277***	0.173***	0.364***	0.339***
	(0.057)	(0.052)	(0.019)	(0.023)	(0.021)
Race/ethnicity/language: Other	0.214***	0.116	0.177***	0.212***	0.046
	(0.079)	(0.088)	(0.025)	(0.038)	(0.034)
Region: Auckland		(baseline)	(baseline)	(baseline)	(baseline)
		(.)	(.)	(.)	(.)
Region: Other		0.042	0.118***	0.093***	0.032*
		(0.033)	(0.024)	(0.019)	(0.017)
Region: Wellington		0.068	0.077**	0.117***	0.104***
		(0.048)	(0.032)	(0.028)	(0.024)
Constant	0.670***	0.688***	0.659***	0.550***	0.516***
	(0.035)	(0.132)	(0.025)	(0.029)	(0.031)
R-squared	0.05	0.04	0.06	0.11	0.10
Observations	6539	8027	26066	17102	17512
Clusters	1581	1482	5815	3680	3419

Note: The table reports the effect of a set of individual characteristics on the probability to vote for Labour / Greens / Other left by decade in New Zealand. The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). The number of clusters corresponds to the number of surveyed individuals in each decade. Robust standard errors clustered at the individual level.

Table E15 - Determinants of support for Labour Party / Socialist Left Party / Other left in Norway

	(1) 1950-59	(2) 1960-69	(3) 1970-79	(4) 1980-89	(5) 1990-99	(6) 2000-09	(7) 2010-20
Education: None/Primary	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Education: Secondary	-0.373***	-0.344***	-0.261***	-0.216***	-0.110***	-0.018	-0.014
	(0.035)	(0.025)	(0.019)	(0.014)	(0.023)	(0.018)	(0.033)
Education: University	-0.457***	-0.409***	-0.310***	-0.218***	-0.146***	0.029	0.051
	(0.053)	(0.028)	(0.027)	(0.026)	(0.026)	(0.019)	(0.035)
Income group: Bottom 50%	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Income group: Middle 40%	0.005	-0.001	-0.013	-0.045***	0.030	-0.035**	0.009
	(0.031)	(0.021)	(0.016)	(0.015)	(0.021)	(0.017)	(0.023)
Income group: Top 10%	-0.239***	-0.194***	-0.168***	-0.239***	-0.100***	-0.146***	-0.151***
	(0.045)	(0.029)	(0.024)	(0.023)	(0.031)	(0.024)	(0.030)
Age: 20-39	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Age: 40-59	-0.091***	-0.048**	-0.030	0.046***	-0.022	0.014	-0.034
	(0.030)	(0.021)	(0.019)	(0.016)	(0.019)	(0.016)	(0.022)
Age: 60+	-0.226***	-0.122***	-0.129***	-0.018	-0.099***	-0.058***	-0.051**
	(0.043)	(0.025)	(0.022)	(0.017)	(0.023)	(0.020)	(0.025)
Gender: Woman	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Gender: Man	0.013	0.059***	0.035*	-0.022*	-0.037**	-0.105***	-0.080***
	(0.039)	(0.019)	(0.018)	(0.013)	(0.016)	(0.014)	(0.017)
Location: Urban	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Location: Rural	-0.113***	-0.128***	-0.095***	-0.142***	-0.129***	-0.063***	-0.113***
	(0.029)	(0.021)	(0.027)	(0.016)	(0.028)	(0.015)	(0.020)
Employment status: Employed	(baseline)	(baseline)	(baseline)			(baseline)	(baseline)
	(.)	(.)	(.)			(.)	(.)
Employment status: Unemployed/Inactive	-0.034	0.059	0.016			0.015	-0.002
	(0.043)	(0.039)	(0.024)			(0.018)	(0.023)
Marital status: Single	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Marital status: Married/With partner	0.119***	0.104***	0.041*	0.093***	-0.016	-0.010	-0.031
	(0.037)	(0.024)	(0.021)	(0.017)	(0.022)	(0.017)	(0.021)
Region: East				(baseline)	(baseline)	(baseline)	(baseline)
				(.)	(.)	(.)	(.)
Region: North				-0.053*	-0.048	-0.004	-0.003
				(0.027)	(0.033)	(0.027)	(0.036)
Region: South and Oslo				-0.136***	-0.079***	-0.065***	-0.082***
				(0.020)	(0.025)	(0.020)	(0.028)
Region: Trondelag				-0.026	-0.003	-0.011	-0.014
				(0.027)	(0.034)	(0.027)	(0.038)
Region: West				-0.212***	-0.160***	-0.148***	-0.146***
				(0.021)	(0.025)	(0.020)	(0.029)
Constant	0.705***	0.533***	0.630***	0.678***	0.738***	0.594***	0.585***
	(0.047)	(0.039)	(0.028)	(0.024)	(0.035)	(0.028)	(0.047)
R-squared	0.18	0.15	0.09	0.11	0.04	0.04	0.05
Observations	2404	5125	8931	12608	5085	7359	4433
Clusters	1170	1598	2393	2184	2119	2082	1887

Note: The table reports the effect of a set of individual characteristics on the probability to vote for Labour Party / Socialist Left Party / Other left by decade in Norway. The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). The number of clusters corresponds to the number of surveyed individuals in each decade. Robust standard errors clustered at the individual level.

Table E16 - Determinants of support for Socialists / Communists / Greens / Left bloc in Portugal

	(1) 1980-89	(2) 1990-99	(3) 2000-09	(4) 2010-20
Education: None/Primary	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)
Education: Secondary	-0.108**	-0.038	-0.028	-0.023
	(0.048)	(0.046)	(0.022)	(0.037)
Education: University	-0.172***	-0.090	-0.120***	-0.175***
	(0.065)	(0.068)	(0.030)	(0.055)
Income group: Bottom 50%	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)
Income group: Middle 40%	0.002	-0.069*	-0.023	-0.042
	(0.037)	(0.040)	(0.018)	(0.035)
Income group: Top 10%	-0.140*	-0.150**	-0.130***	-0.124**
	(0.079)	(0.066)	(0.029)	(0.059)
Age: 20-39	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)
Age: 40-59	-0.006	-0.001	0.011	-0.007
	(0.043)	(0.043)	(0.023)	(0.042)
Age: 60+	-0.103**	-0.020	-0.103***	-0.024
	(0.049)	(0.048)	(0.028)	(0.050)
Gender: Woman	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)
Gender: Man	0.055	0.018	-0.019	-0.015
	(0.037)	(0.035)	(0.019)	(0.029)
Religion: None	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)
Religion: Catholic	-0.147	-0.093	-0.117	-0.116**
	(0.107)	(0.098)	(0.081)	(0.058)
Religion: Other	0.003	-0.029	-0.143	-0.030
	(0.161)	(0.159)	(0.098)	(0.099)
Religious practice: Never	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)
Religious practice: Less than monthly	-0.043	-0.156**	0.001	-0.062
	(0.066)	(0.066)	(0.030)	(0.047)
Religious practice: Monthly or more	-0.168***	-0.221***	-0.108***	-0.192***
	(0.065)	(0.068)	(0.030)	(0.049)
Location: Urban	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)
Location: Rural	-0.105***	0.021	-0.054***	-0.043
	(0.036)	(0.037)	(0.018)	(0.034)
Employment status: Employed	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)
Employment status: Unemployed/Inactive	-0.063	-0.101**	0.021	0.060*
	(0.041)	(0.040)	(0.020)	(0.035)
Marital status: Single	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)
Marital status: Married/With partner	0.031	0.095***	-0.008	-0.004
	(0.035)	(0.035)	(0.021)	(0.033)
Region: Alentejo	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)

Region: Algarve	-0.511***	-0.566***	-0.188***	-0.200**
	(0.061)	(0.063)	(0.054)	(0.096)
Region: Center	-0.242***	-0.283***	-0.234***	-0.367***
	(0.062)	(0.065)	(0.038)	(0.069)
Region: Lisbon	-0.206***	-0.227***	-0.094**	-0.156**
	(0.053)	(0.056)	(0.037)	(0.071)
Region: North	-0.334***	-0.404***	-0.167***	-0.229***
	(0.054)	(0.052)	(0.035)	(0.067)
Constant	1.109***	0.980***	0.978***	1.074***
	(0.111)	(0.106)	(0.091)	(0.099)
R-squared	0.14	0.11	0.07	0.11
Observations	7986	3442	12259	3759
Clusters	1223	1407	2459	1105

Note: The table reports the effect of a set of individual characteristics on the probability to vote for Socialists / Communists / Greens / Left bloc by decade. The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). The number of clusters corresponds to the number of surveyed individuals in each decade. Robust standard errors clustered at the individual level.

Table E17 - Determinants of support for PSOE / Podemos / IU / Other left in Spain

	(1) 1980-89	(2) 1990-99	(3) 2000-09	(4) 2010-20
Education: None/Primary	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)
Education: Secondary	-0.101***	-0.081***	-0.045***	-0.065***
	(0.013)	(0.017)	(0.011)	(0.011)
Education: University	-0.172***	-0.178***	-0.074***	-0.073***
	(0.021)	(0.023)	(0.014)	(0.013)
Education: Postgraduate	-0.375***	-0.297**	-0.124***	-0.080**
	(0.144)	(0.134)	(0.042)	(0.035)
Income group: Bottom 50%	(baseline)		(baseline)	(baseline)
	(.)		(.)	(.)
Income group: Middle 40%	-0.029		-0.055***	-0.034***
	(0.018)		(0.010)	(0.009)
Income group: Top 10%	-0.144***		-0.123***	-0.080***
	(0.032)		(0.016)	(0.015)
Age: 20-39	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)
Age: 40-59	-0.148***	-0.093***	0.022**	0.034***
	(0.016)	(0.016)	(0.010)	(0.010)
Age: 60+	-0.169***	-0.121***	-0.038***	0.027**
	(0.018)	(0.019)	(0.013)	(0.012)
Gender: Woman	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)
Gender: Man	0.045***	0.024*	-0.067***	-0.048***
	(0.014)	(0.013)	(0.008)	(0.007)
Religion: None	(baseline)		(baseline)	(baseline)
	(.)		(.)	(.)
Religion: Catholic	-0.303***		-0.220***	-0.280***
	(0.025)		(0.012)	(0.011)
Religion: Other	0.021		-0.047	-0.009
	(0.084)		(0.047)	(0.029)
Religious practice: Never	(baseline)		(baseline)	(baseline)
	(.)		(.)	(.)
Religious practice: Less than monthly	-0.139***		-0.132***	-0.136***
	(0.024)		(0.012)	(0.011)
Religious practice: Monthly or more	-0.398***		-0.259***	-0.267***
	(0.021)		(0.011)	(0.010)
Location: Urban	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)
Location: Rural	-0.093***	-0.096***	-0.036**	-0.008
	(0.018)	(0.028)	(0.015)	(0.014)
Employment status: Employed	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)
Employment status: Unemployed/Inactive	0.010	0.013	0.001	0.023**
	(0.015)	(0.014)	(0.010)	(0.010)
Marital status: Single	(baseline)		(baseline)	(baseline)
	(.)		(.)	(.)
Marital status: Married/With partner	0.019		-0.008	-0.007
	(0.018)		(0.010)	(0.008)

Region: Andalucia	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)
Region: Aragon	-0.101***	-0.185***	-0.014	-0.011
	(0.032)	(0.036)	(0.024)	(0.022)
Region: Asturias	-0.083**	-0.092**	(baseline)	0.009
	(0.034)	(0.038)	(0.026)	(0.021)
Region: Baleares	-0.230***	-0.239***	-0.187***	-0.043*
	(0.041)	(0.045)	(0.028)	(0.024)
Region: Basque Country	-0.107***	-0.173***	-0.185***	0.041*
	(0.028)	(0.033)	(0.022)	(0.021)
Region: Canarias	-0.203***	-0.269***	-0.251***	-0.042**
	(0.032)	(0.035)	(0.020)	(0.020)
Region: Cantabria	-0.189***	-0.076	-0.060*	0.140***
	(0.047)	(0.059)	(0.033)	(0.030)
Region: Castilla La Mancha	-0.233***	-0.230***	-0.079***	-0.067***
	(0.029)	(0.031)	(0.020)	(0.019)
Region: Castilla y Leon	-0.239***	-0.211***	-0.107***	-0.090***
	(0.025)	(0.027)	(0.017)	(0.018)
Region: Catalonia	-0.209***	-0.147***	-0.063***	0.017
	(0.019)	(0.020)	(0.015)	(0.015)
Region: Extremadura	-0.141***	(baseline)	0.014	-0.010
	(0.035)	(0.038)	(0.024)	(0.020)
Region: Galicia	-0.249***	-0.177***	-0.035**	(baseline)
	(0.020)	(0.028)	(0.017)	(0.017)
Region: Madrid	-0.110***	-0.154***	-0.120***	-0.081***
	(0.020)	(0.022)	(0.016)	(0.014)
Region: Murcia	-0.103***	-0.237***	-0.224***	-0.091***
	(0.037)	(0.039)	(0.023)	(0.019)
Region: Navarra	0.016	-0.025	0.023	0.108***
	(0.054)	(0.054)	(0.036)	(0.030)
Region: Rioja	-0.171**	-0.289***	-0.138**	-0.044
	(0.070)	(0.062)	(0.057)	(0.034)
Region: Valencia	-0.109***	-0.203***	-0.185***	-0.022
	(0.021)	(0.023)	(0.015)	(0.014)
Region: Ceuta				-0.278***
				(0.040)
Region: Melilla				-0.227***
				(0.063)
Constant	1.447***	0.749***	0.982***	0.836***
	(0.093)	(0.024)	(0.019)	(0.019)
R-squared	0.13	0.05	0.13	0.14
Observations	20532	11048	45305	74833
Clusters	4358	4925	6005	6216

Note: The table reports the effect of a set of individual characteristics on the probability to vote for PSOE / Podemos / IU / Other left by decade in Spain. The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). The number of clusters corresponds to the number of surveyed individuals in each decade. Robust standard errors clustered at the individual level.

Table E18 - Determinants of support for Social Democratic Party / Left Party / Green Party in Sweden

	(1) 1950-59	(2) 1960-69	(3) 1970-79	(4) 1980-89	(5) 1990-99	(6) 2000-09	(7) 2010-20
Education: None/Primary	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Education: Secondary	-0.312***	-0.244***	-0.218***	-0.199***	-0.163***	-0.108***	-0.036
	(0.025)	(0.015)	(0.013)	(0.013)	(0.016)	(0.023)	(0.033)
Education: University	-0.513***	-0.472***	-0.302***	-0.323***	-0.255***	-0.212***	-0.029
	(0.034)	(0.020)	(0.017)	(0.017)	(0.018)	(0.024)	(0.037)
Income group: Bottom 50%	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Income group: Middle 40%	0.121***	0.051***	0.026**	0.050***	0.010	-0.024**	-0.085***
	(0.027)	(0.013)	(0.011)	(0.011)	(0.013)	(0.012)	(0.022)
Income group: Top 10%	-0.048	-0.052***	-0.136***	-0.053***	-0.121***	-0.189***	-0.225***
	(0.039)	(0.019)	(0.017)	(0.016)	(0.021)	(0.020)	(0.029)
Age: 20-39	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Age: 40-59	-0.071***	-0.061***	-0.024**	-0.062***	0.017	0.050***	-0.016
	(0.024)	(0.014)	(0.012)	(0.014)	(0.014)	(0.017)	(0.029)
Age: 60+	-0.135***	-0.079***	-0.061***	-0.086***	-0.061***	-0.096***	-0.012
	(0.030)	(0.017)	(0.013)	(0.016)	(0.019)	(0.022)	(0.032)
Gender: Woman	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Gender: Man	0.027	0.018	0.021**	-0.040***	-0.036***	-0.031**	-0.030
	(0.021)	(0.014)	(0.010)	(0.012)	(0.012)	(0.014)	(0.023)
Religious practice: Never	(baseline)	(baseline)		(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)		(.)	(.)	(.)	(.)
Religious practice: Less than monthly	-0.190***	-0.270***		-0.229***	-0.193***	-0.169***	-0.113**
	(0.039)	(0.017)		(0.022)	(0.021)	(0.028)	(0.047)
Religious practice: Monthly or more	0.119***	0.128***		0.137***	0.106***	0.093***	0.046*
	(0.037)	(0.019)		(0.016)	(0.013)	(0.016)	(0.027)
Location: Urban	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Location: Rural	-0.097***	-0.122***	-0.214***	-0.167***	-0.133***	-0.090***	-0.072**
	(0.022)	(0.013)	(0.018)	(0.015)	(0.016)	(0.020)	(0.030)
Employment status: Employed				(baseline)	(baseline)	(baseline)	(baseline)
				(.)	(.)	(.)	(.)
Employment status: Unemployed/Inactive				0.024	0.030*	0.032	0.019
				(0.018)	(0.016)	(0.020)	(0.028)
Marital status: Single	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Marital status: Married/With partner	0.028	0.046***	-0.004	0.017	0.019	-0.019	-0.001
	(0.024)	(0.014)	(0.011)	(0.013)	(0.014)	(0.017)	(0.029)
Region: Gotland					(baseline)	(baseline)	(baseline)
					(.)	(.)	(.)
Region: Norrland					0.139***	0.110***	0.186***
					(0.017)	(0.021)	(0.037)
Region: Svealand					-0.002	-0.011	0.032
					(0.013)	(0.015)	(0.024)
Constant	0.855***	0.688***	0.634***	0.718***	0.690***	0.663***	0.536***
	(0.147)	(0.018)	(0.015)	(0.019)	(0.023)	(0.031)	(0.046)
R-squared	0.16	0.15	0.06	0.09	0.09	0.07	0.05
Observations	4441	18082	24545	22345	15299	14440	9405
Clusters	1414	3234	4536	3745	3450	3370	2684

Note: The table reports the effect of a set of individual characteristics on the probability to vote for Social Democratic Party / Left Party / Green Party by decade in Sweden. The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). The number of clusters corresponds to the number of surveyed individuals in each decade. Robust standard errors clustered at the individual level.

Table E19 - Determinants of support for Social Democrats / Greens / Other left in Switzerland

	(1) 1960-69	(2) 1970-79	(3) 1980-89	(4) 1990-99	(5) 2000-09	(6) 2010-20
Education: None/Primary	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Education: Secondary	-0.053	-0.059**	0.045	(baseline)	0.032	0.036*
	(0.068)	(0.024)	(0.043)	(0.026)	(0.021)	(0.020)
Education: University	-0.217***	-0.069	-0.033	0.062*	0.147***	0.193***
	(0.065)	(0.043)	(0.070)	(0.037)	(0.029)	(0.023)
Income group: Bottom 50%		(baseline)		(baseline)	(baseline)	(baseline)
		(.)		(.)	(.)	(.)
Income group: Middle 40%		-0.022		0.018	0.003	0.012
		(0.031)		(0.017)	(0.016)	(0.011)
Income group: Top 10%		-0.114***		-0.050*	-0.111***	-0.049***
		(0.044)		(0.026)	(0.023)	(0.017)
Age: 20-39	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Age: 40-59	0.020	0.024	-0.105**	-0.051**	0.007	-0.006
	(0.055)	(0.023)	(0.046)	(0.021)	(0.019)	(0.015)
Age: 60+	-0.042	-0.024	-0.122**	-0.117***	-0.082***	-0.056***
	(0.070)	(0.028)	(0.054)	(0.026)	(0.021)	(0.016)
Gender: Woman		(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
		(.)	(.)	(.)	(.)	(.)
Gender: Man		0.017	-0.007	-0.054***	-0.066***	-0.082***
		(0.025)	(0.040)	(0.018)	(0.014)	(0.011)
Religion: None	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Religion: Catholic	-0.383***	-0.203***	-0.376***	-0.136***	-0.136***	-0.090***
	(0.139)	(0.077)	(0.079)	(0.048)	(0.034)	(0.028)
Religion: Other Christian	-0.255*	-0.123	-0.245***	-0.047	-0.060*	-0.050*
	(0.132)	(0.075)	(0.080)	(0.047)	(0.034)	(0.028)
Religion: Other	-0.149	-0.044	-0.114	-0.019	-0.084**	-0.107***
	(0.328)	(0.136)	(0.148)	(0.068)	(0.042)	(0.036)
Religious practice: Never	(baseline)	(baseline)		(baseline)	(baseline)	(baseline)
	(.)	(.)		(.)	(.)	(.)
Religious practice: Less than monthly	-0.137*	-0.081**		-0.100***	-0.073**	-0.093***
	(0.078)	(0.036)		(0.035)	(0.029)	(0.026)
Religious practice: Monthly or more	-0.312***	-0.270***		-0.221***	-0.183***	-0.218***
	(0.087)	(0.039)		(0.038)	(0.032)	(0.028)
Location: Urban	(baseline)	(baseline)		(baseline)	(baseline)	(baseline)
	(.)	(.)		(.)	(.)	(.)
Location: Rural	-0.162***	-0.120***		-0.080***	-0.118***	-0.118***
	(0.050)	(0.025)		(0.016)	(0.014)	(0.011)
Employment status: Employed	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Employment status: Unemployed/Inactive	0.028	-0.015	0.094**	-0.000	-0.042**	0.001
	(0.081)	(0.027)	(0.048)	(0.023)	(0.018)	(0.014)
Marital status: Single	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Marital status: Married/With partner	-0.053	-0.016	-0.045	-0.049***	-0.018	-0.061***
	(0.056)	(0.025)	(0.043)	(0.019)	(0.014)	(0.012)

Region: French	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)
Region: German	-0.017	-0.016	-0.148***	-0.041*	-0.049***	-0.054***
	(0.060)	(0.029)	(0.055)	(0.021)	(0.017)	(0.014)
Region: Italian	-0.087	-0.019	-0.299***	-0.123***	-0.039*	-0.111***
	(0.105)	(0.063)	(0.070)	(0.030)	(0.023)	(0.017)
Constant	1.038***	0.831***	0.779***	0.675***	0.641***	0.644***
	(0.123)	(0.084)	(0.099)	(0.046)	(0.033)	(0.028)
R-squared	0.17	0.10	0.11	0.07	0.10	0.11
Observations	456	3294	900	11775	11681	18865
Clusters	456	2182	582	6599	6567	11127

Note: The table reports the effect of a set of individual characteristics on the probability to vote for Social Democrats / Greens / Other left by decade in Switzerland. The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). The number of clusters corresponds to the number of surveyed individuals in each decade. Robust standard errors clustered at the individual level.

Table E20 - Determinants of support for the Labour Party in the United Kingdom

	(1) 1950-59	(2) 1960-69	(3) 1970-79	(4) 1980-89	(5) 1990-99	(6) 2000-09	(7) 2010-20
Education: None/Primary	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Education: Secondary	-0.214***	-0.190***	-0.177***	-0.130***	-0.129***	-0.072***	-0.022*
	(0.025)	(0.021)	(0.009)	(0.009)	(0.011)	(0.013)	(0.013)
Education: University	-0.213***	-0.212***	-0.207***	-0.114***	-0.090***	-0.131***	-0.034**
	(0.065)	(0.053)	(0.015)	(0.013)	(0.016)	(0.018)	(0.017)
Education: Postgraduate					0.033	-0.057**	0.010
					(0.050)	(0.026)	(0.021)
Income group: Bottom 50%	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Income group: Middle 40%	-0.165***	-0.084***	-0.039***	-0.167***	-0.142***	-0.078***	-0.047***
	(0.022)	(0.019)	(0.010)	(0.011)	(0.012)	(0.014)	(0.014)
Income group: Top 10%	-0.346***	-0.371***	-0.180***	-0.259***	-0.222***	-0.131***	-0.106***
	(0.030)	(0.023)	(0.014)	(0.014)	(0.018)	(0.020)	(0.021)
Age: 20-39	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Age: 40-59	-0.039*	-0.048***	-0.012	-0.046***	-0.034***	-0.062***	-0.052***
	(0.023)	(0.018)	(0.009)	(0.009)	(0.012)	(0.015)	(0.016)
Age: 60+	-0.165***	-0.089***	-0.067***	-0.075***	-0.060***	-0.100***	-0.132***
	(0.028)	(0.022)	(0.010)	(0.012)	(0.014)	(0.017)	(0.019)
Gender: Woman	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Gender: Man	0.102***	0.053***	0.026***	0.033***	0.021**	0.004	-0.020*
	(0.024)	(0.020)	(0.009)	(0.008)	(0.010)	(0.011)	(0.011)
Religion: None	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Religion: Catholic	-0.131*	0.068	0.074***	0.096***	0.095***	0.110***	0.084***
	(0.070)	(0.055)	(0.017)	(0.016)	(0.018)	(0.021)	(0.022)
Religion: Other Christian	-0.241***	-0.131***	-0.093***	-0.093***	-0.091***	-0.056***	-0.068***
	(0.060)	(0.048)	(0.010)	(0.009)	(0.011)	(0.013)	(0.012)
Religion: Other	-0.290**	-0.134	0.096***	-0.090**	-0.069*	-0.093***	-0.026
	(0.116)	(0.084)	(0.036)	(0.041)	(0.036)	(0.026)	(0.044)
Religion: Muslim				0.043	-0.028	-0.058	0.220***
				(0.062)	(0.076)	(0.074)	(0.050)
Employment status: Employed	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Employment status: Unemployed/Inactive	0.007	0.037*	0.024**	0.002	-0.019	-0.012	0.023
	(0.026)	(0.022)	(0.010)	(0.010)	(0.012)	(0.012)	(0.015)
Marital status: Single	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Marital status: Married/With partner	0.072***	0.059***	0.029***	-0.004	0.013	0.004	-0.046***
	(0.025)	(0.020)	(0.010)	(0.010)	(0.011)	(0.012)	(0.012)
Race/ethnicity: African / Caribbean			(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
			(.)	(.)	(.)	(.)	(.)
Race/ethnicity: Indian / Pak. / Bang.			0.101	0.141***	-0.077	-0.028	-0.146**
			(0.134)	(0.054)	(0.064)	(0.082)	(0.061)
Race/ethnicity: Other			-0.310	-0.408***	-0.176**	-0.260***	-0.295***
			(0.231)	(0.069)	(0.076)	(0.074)	(0.060)
Race/ethnicity: White			-0.326***	-0.242***	-0.329***	-0.337***	-0.371***
			(0.081)	(0.039)	(0.036)	(0.055)	(0.043)
Constant	0.782***	0.647***	0.888***	0.751***	0.893***	0.886***	0.840***
	(0.067)	(0.053)	(0.082)	(0.040)	(0.038)	(0.057)	(0.046)
R-squared	0.12	0.10	0.06	0.09	0.08	0.05	0.09
Observations	5122	6732	26522	18740	17793	14053	15439
Clusters	2025	2377	8082	7409	6770	5957	5760

Note: The table reports the effect of a set of individual characteristics on the probability to vote for the Labour Party by decade in Britain. The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). The number of clusters corresponds to the number of surveyed individuals in each decade. Robust standard errors clustered at the individual level.

Table E21 - Determinants of support for the Democratic Party in the United States

	(1) 1948-59	(2) 1960-69	(3) 1970-79	(4) 1980-89	(5) 1990-99	(6) 2000-09	(7) 2010-20
Education: None/Primary	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Education: Secondary	-0.115***	-0.106***	-0.026	-0.126***	-0.098***	-0.100***	-0.031
	(0.021)	(0.023)	(0.022)	(0.024)	(0.035)	(0.037)	(0.026)
Education: University	-0.149***	-0.164***	-0.057*	-0.137***	-0.166***	-0.071*	0.077***
	(0.035)	(0.036)	(0.031)	(0.030)	(0.041)	(0.042)	(0.028)
Education: Postgraduate		-0.178***	0.041	0.020	-0.105**	0.005	0.217***
		(0.051)	(0.045)	(0.037)	(0.046)	(0.047)	(0.028)
Income group: Bottom 50%	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Income group: Middle 40%	-0.037*	0.004	-0.072***	-0.051**	-0.080***	-0.047**	-0.006
	(0.021)	(0.023)	(0.022)	(0.021)	(0.025)	(0.023)	(0.013)
Income group: Top 10%	-0.124***	-0.078**	-0.177***	-0.173***	-0.134***	-0.140***	-0.002
	(0.028)	(0.031)	(0.028)	(0.026)	(0.033)	(0.033)	(0.022)
Age: 20-39	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Age: 40-59	-0.049**	-0.026	-0.031	0.029	0.084***	-0.011	-0.040***
	(0.021)	(0.023)	(0.021)	(0.018)	(0.024)	(0.022)	(0.014)
Age: 60+	-0.118***	-0.107***	-0.051*	0.009	0.045	-0.011	-0.033**
	(0.028)	(0.030)	(0.026)	(0.024)	(0.030)	(0.027)	(0.015)
Gender: Woman	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Gender: Man	0.035*	-0.013	-0.046**	-0.070***	-0.100***	-0.081***	-0.062***
	(0.019)	(0.020)	(0.019)	(0.017)	(0.021)	(0.018)	(0.011)
Religion: Catholic	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Religion: Other Christian	-0.191***	-0.403***	-0.136***	-0.122***	-0.141***	-0.102***	-0.102***
	(0.025)	(0.024)	(0.022)	(0.021)	(0.027)	(0.025)	(0.015)
Religion: Other	-0.112*	-0.229***	0.068	0.018	-0.035	-0.012	0.034*
	(0.064)	(0.074)	(0.047)	(0.039)	(0.038)	(0.032)	(0.019)
Religious practice: Never	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Religious practice: Less than monthly	-0.038	-0.021	-0.043	-0.009	-0.079**	0.018	-0.030
	(0.045)	(0.046)	(0.030)	(0.029)	(0.035)	(0.034)	(0.020)
Religious practice: Monthly or more	-0.082*	-0.113**	-0.092***	-0.042	-0.212***	-0.125***	-0.136***
	(0.044)	(0.044)	(0.028)	(0.027)	(0.027)	(0.025)	(0.016)
Location: Urban	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	
	(.)	(.)	(.)	(.)	(.)	(.)	
Location: Rural	-0.012	0.050**	-0.018	-0.017	-0.003	-0.030	
	(0.022)	(0.021)	(0.020)	(0.018)	(0.023)	(0.047)	
Employment status: Employed			(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
			(.)	(.)	(.)	(.)	(.)
Employment status: Unemployed/Inactive			-0.026	-0.012	-0.060**	0.017	0.020
			(0.021)	(0.021)	(0.026)	(0.023)	(0.013)
Marital status: Single	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Marital status: Married/With partner	0.003	0.032	-0.021	-0.009	-0.036	-0.064***	-0.080***
	(0.034)	(0.026)	(0.021)	(0.017)	(0.023)	(0.021)	(0.013)
Race/ethnicity/language: Black	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)
Race/ethnicity/language: White	-0.248***	-0.458***	-0.521***	-0.543***	-0.507***	-0.575***	-0.497***
	(0.049)	(0.036)	(0.023)	(0.021)	(0.025)	(0.022)	(0.019)
Race/ethnicity/language: Other		-0.792***	-0.318***	-0.376***	-0.348***	-0.411***	-0.289***
		(0.210)	(0.065)	(0.038)	(0.040)	(0.032)	(0.024)
Region: North Central	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)	(baseline)
	(.)	(.)	(.)	(.)	(.)	(.)	(.)

Region: Northeast	-0.111***	-0.073***	-0.016	-0.043*	0.040	-0.022	-0.015
	(0.024)	(0.024)	(0.025)	(0.024)	(0.032)	(0.032)	(0.023)
Region: South	0.129***	0.036	-0.046**	-0.021	-0.009	-0.094***	-0.079***
	(0.029)	(0.027)	(0.023)	(0.020)	(0.026)	(0.024)	(0.018)
Region: West	0.001	0.009	-0.040	0.017	0.001	0.031	0.003
	(0.033)	(0.029)	(0.027)	(0.023)	(0.031)	(0.029)	(0.021)
Constant	0.997***	1.415***	1.244***	1.211***	1.407***	1.375***	1.131***
	(0.077)	(0.064)	(0.046)	(0.044)	(0.053)	(0.055)	(0.038)
R-squared	0.12	0.19	0.15	0.18	0.21	0.21	0.20
Observations	6532	5513	4498	4892	3491	5301	19023
Clusters	1718	1486	2138	1986	1821	2043	7765

Note: The table reports the effect of a set of individual characteristics on the probability to vote for the Democratic Party by decade in the United States. The original survey dataset is duplicated for each income bracket to approximate income deciles (see methodology). The number of clusters corresponds to the number of surveyed individuals in each decade. Robust standard errors clustered at the individual level.