

The Global Decline of the Labor Share*

Loukas Karabarbounis

Brent Neiman

University of Chicago and NBER

University of Chicago and NBER

June 2013

Abstract

The stability of the labor share of income is a key foundation in macroeconomic models. We document, however, that the global labor share has significantly declined since the early 1980s, with the decline occurring within the large majority of countries and industries. We show that the decrease in the relative price of investment goods, often attributed to advances in information technology and the computer age, induced firms to shift away from labor and toward capital. The lower price of investment goods explains roughly half of the observed decline in the labor share, even when we allow for other mechanisms influencing factor shares such as increasing profits, capital-augmenting technology growth, and the changing skill composition of the labor force. We highlight the implications of this explanation for welfare and macroeconomic dynamics.

JEL-Codes: E21, E22, E25.

Keywords: Labor Share, Price of Investment Goods.

*We thank Robert Barro, Martin Berka, Bob Hall, Chang-Tai Hsieh, John Huizinga, Chad Jones, Pete Klenow, and three anonymous referees for valuable comments. Gabriela Antonie, Sophie Wang, Bowen Yang, Anny Zhong, Michael Marvin, and Victor Lin provided excellent research assistance. This research was funded in part by the Initiative on Global Markets and the Neubauer Family Foundation at the University of Chicago Booth School of Business.

1 Introduction

At least since the work of Kaldor (1957), the stability of the labor share of income has been a fundamental feature of macroeconomic models, with broad implications for the shape of the production function, inequality, and macroeconomic dynamics. We document that the global labor share has declined significantly since the early 1980s, with the decline occurring within the large majority of countries and industries. We demonstrate how the decline of the labor share can be explained by the decline in the relative price of investment goods. Efficiency gains in capital producing sectors, often attributed to advances in information technology and the computer age, induced firms to shift away from labor and toward capital to such a large extent that the labor share of income declined.

We start by documenting a 5 percentage point decline in the share of global corporate gross value added paid to labor over the last 35 years. We measure the labor share using a novel dataset we compile by combining country-specific data posted on the Internet with sector-level national income accounting data from multilateral organizations obtained digitally and collected from physical books. Our baseline analyses focus on the labor share within the corporate sector as this allows us to circumvent important measurement difficulties confronted by most of the labor share literature. As emphasized by Gollin (2002), aggregate labor share measures are influenced by the methods chosen by national statistical agencies to separate the labor and capital income earned by entrepreneurs, sole proprietors, and unincorporated businesses. The corporate labor share is not subject to this imputation. While previous analyses of U.S. data have sometimes focused on the corporate labor share, we are unaware of other research focusing on corporate labor shares in such a large sample of countries.

Of the 56 countries with at least 15 years of data between 1975 and 2012, 38 exhibited downward trends in their labor shares. Of the trend estimates that are statistically significant, 34 are negative while only 9 are positive. We complement our analysis with industry-level data and show that six of the ten major industries experienced significant labor share declines while only two experienced the opposite. Most of the global decline in the labor share is attributable to within-industry changes rather than to changes in industrial composition. The pervasiveness of the decline in the labor share is even present in regional data for the United States, where

two-thirds of the states experienced declines over this period.

The decline in the price of investment relative to consumption goods accelerated starting in the early 1980s. We develop a model that relates the decline in the labor share to this coincident decline in the relative price of investment goods. The economy produces two final goods (consumption and investment) using a continuum of intermediate inputs. Technology differences in the production of final goods cause shifts in the price of investment relative to the price of consumption goods and affect the rate at which households rent capital to the firms. Monopolistically competitive firms produce intermediate inputs with capital and labor using a constant elasticity of substitution (CES) technology and sell their output each period at a constant markup over marginal cost. Changes in the rental rate of capital induce producers to change their capital-to-labor ratios and, for non-unitary elasticities of substitution, the shares of each factor in production costs. Changes in price markups additionally change the shares of each factor in income.

In our model the labor share will only change in response to shocks that influence the rental rate of capital, markups, or capital-augmenting technology, with the magnitude of any response being a function of the elasticity of substitution between capital and labor and the levels of the labor share and markups. Given our focus on long-term trends, we treat the data as being generated from the model's transition from one steady state to another. Assuming a constant household discount factor and depreciation rate of capital, changes across steady states in the rental rate only reflect changes in the relative price of investment. Heterogeneity across countries in the level or growth of any variable other than the relative price of investment, markups, or capital-augmenting technology will therefore not matter for long-term trends in the labor share. This logic argues against the possibility that shocks to other macroeconomic objects such as labor income taxes or household labor supply are important for explaining the labor share decline.

To determine the implications of the declining relative price of investment for the labor share, we use our model to estimate the elasticity of substitution between capital and labor. Most prior estimates use time series variation within a country in factor shares and factor prices to identify the elasticity. By contrast, our estimates are identified from cross-country variation in trends in rental rates and labor shares. Therefore, our estimates are not influenced

by the global component of the labor share decline, the object that we intend to explain. Put differently, even if each individual country experienced a decline in both its relative price of investment and its labor share, there is nothing in our methodology that prevents us from associating a global decline in the price of investment with a global increase in the labor share.

The rental rate of capital can be influenced at high frequency by various factors such as short-run changes in interest rates, adjustment costs, or financial frictions. These factors, however, are unlikely to have a significant influence on long-run trends in the rental rate, particularly compared to the relative price of investment goods, which moves proportionately with the rental rate across steady states of our model. Therefore, our estimates focus on low-frequency variation and only include countries with at least 15, and as many as 37, years of data. Rather than having to use more volatile proxies of the rental rate, this allows us to exploit high quality and widely available data on the relative price of investment.

We start by assuming that capital-augmenting technology growth is orthogonal to the price of investment shock and that the economy has zero profits. In the data, countries and industries experiencing larger declines in the relative price of investment also experienced larger labor share declines. This leads to our baseline estimate of the elasticity of substitution between capital and labor of 1.25. When confronted with the 25 percent decline in the global relative price of investment that occurred since 1975, our model delivers roughly half of the 5 percentage point decline in the global labor share.

Next, we allow for the possibility that markups affect our estimated elasticity. Imagine that markups increased more in countries with larger declines in the relative price of investment. Even in the Cobb-Douglas case, which features a constant labor share of costs, this would produce a spurious association between declining labor shares of income and declining prices of investment. Our baseline procedure would incorrectly estimate an elasticity greater than one. To address this concern, we use long-term trends in nominal investment rates to approximate changes in the capital-output ratio and follow Rotemberg and Woodford (1995) in using this ratio to calculate capital shares and markups. We find that markups generally increased and therefore did play a role in the labor share decline. However, when we modify our empirical framework to take markups into account, the estimated elasticity, and thus the implied contribution of the price of investment to the labor share decline, is essentially unchanged relative

to our benchmark results.

Similarly, our elasticity estimate might be biased upward if capital-augmenting technology growth is greater in countries with larger declines in the relative price of investment. The size of the bias is a function of the covariance between capital-augmenting technology growth and changes in the relative price of investment in the cross section of countries. We show that if the pattern of capital-augmenting technology growth is similar to the pattern of estimated total factor productivity (TFP) growth, then our estimated elasticity is biased upward by less than 0.05. Therefore, allowing for capital-augmenting technology growth does not alter our assessment of the importance of declines in the relative price of investment for the decline in the labor share.

We also consider the possibility that changes in the skill composition of the labor force impact our estimates and explanation. We modify the production function to allow for two types of labor that are differentially substitutable with capital. We use this framework to estimate the sensitivity of the labor share with respect to the relative price of investment goods, controlling for changes in the stock of skill relative to the stock of capital. Our results show that the declining price of investment goods continues to account for roughly half of the decline in the labor share.

We conclude by using our model to evaluate the aggregate implications of our explanation for the decline in the labor share. We start by comparing the impact of the observed shock to the relative price of investment in a standard model with Cobb-Douglas production relative to our model with CES production and an elasticity of substitution equal to 1.25. Welfare gains resulting from the shock are nearly a quarter (or 4 percentage points) higher in the CES case. Next, we compare the consequences of two shocks, a decline in the relative price of investment and an increase in markups, each of which generates an equal reduction in the labor share. The differences are stark. A labor share decline due to reductions in the relative price of investment is associated with large welfare gains. The same labor share decline, but due to increases in markups, is associated with modest welfare losses.

Our work relates to several strands of literature. First, our findings are consistent with earlier work by Blanchard (1997), Blanchard and Giavazzi (2003), Jones (2003), and Bentolila and Saint-Paul (2003) that focuses on the variability of labor shares over the medium run,

including the large declines seen during the 1980s in Western Europe.¹ Harrison (2002) and Rodriguez and Jayadev (2010) use UN data and are the broadest studies of trends in labor shares. Harrison (2002) finds a decreasing trend in the labor share of poor countries but an increasing trend in rich countries for 1960-1997. Rodriguez and Jayadev (2010) estimate a declining average trend in labor shares using an equally weighted set of 129 countries.

Our results improve and expand upon this related literature. We capture significant movements in the labor share subsequent to 2000, include important non-OECD countries such as China, and use exchange rates to aggregate across countries and examine the global labor share. By focusing on the labor share in the corporate sector, rather than the overall labor share, our results are less subject to measurement problems caused by the imputation of labor earnings in unincorporated enterprises and by shifts in economic activity across sectors. And importantly, we offer novel evidence tying the decline in the labor share to the decline in the relative price of investment goods and compare our mechanism to other potential explanations.

Our work relates the decline in labor share to the decline in the relative price of investment by estimating an elasticity of substitution between capital and labor that exceeds unity.² As reviewed in Antras (2004) and Chirinko (2008) among others, there is a large literature estimating this elasticity. Though the range of estimates is very wide, most estimates are below one.³ As discussed above, an important difference between our approach and that taken by most papers in the literature is that we estimate the elasticity from cross-sectional variation using many countries and that, by focusing only on long-run trends, we take advantage of cross-country variation in the relative price of investment.

Finally, our paper also relates to the literature on investment-specific technical change.

¹Blanchard and Giavazzi (2003) argue that deregulation in product and labor markets decreased labor shares and increased unemployment in Europe in the 1980s. Azmat, Manning, and Van Reenen (2012) explore deregulations in the network industry to further advance this argument.

²As discussed in Barro and Sala-i-Martin (1995) and Jones (2003), a balanced growth path with non-zero factor shares will only emerge if technology growth is labor-augmenting, regardless of the production function, or if the production function is Cobb-Douglas, even if technology growth is capital-augmenting. Acemoglu (2003) and Jones (2005) develop models in which firms choose technical progress to be labor-augmenting along the balanced growth path. If real wage growth or increases in the capital-to-labor ratio are caused by labor-augmenting technology growth, there need not be movement in the labor share.

³A notable exception is Duffy and Papageorgiou (2000). Leon-Ledesma, McAdam, and Willman (2010) attribute some of the large variation in the estimates of the elasticity to the use of a single equation first-order condition rather than a joint estimation with the production function. Data limitations prevent us from employing their methodology. We note that they find a large downward bias in estimated elasticities from simulated time-series when the true underlying elasticity is greater than one.

Following Greenwood, Hercowitz, and Huffman (1988), and consistent with Hsieh and Klenow (2007), we interpret innovations in the relative price of investment as reflecting investment-specific technology shocks. Greenwood, Hercowitz, and Krusell (1997) study the relevance of this type of shock for growth in the United States. Fisher (2006) documents an acceleration in the decline in the relative price of investment since 1982 and evaluates the importance of investment shocks for business cycles. Perhaps closest in spirit to our narrative, Krusell, Ohanian, Rios-Rull, and Violante (2000) study the evolution of the U.S. skill premium in a model with price of investment shocks and CES production between skilled labor and capital equipment.

2 Trends in Labor Shares and Investment Prices

We start by documenting the pervasive decline in labor shares around the world at the country, U.S. state, and industry levels. Next, we document a decline in the relative price of investment goods, which we later show to be the key factor explaining the global trend in the labor share. Each subsection first describes the data sources used and then summarizes the relevant trends.

2.1 Declining Labor Shares

Our baseline results come from analysis of a new dataset we construct using country-level statistics on labor share in the corporate sector. We generate these data by combining five broad sources: (i) country-specific Internet web pages (such as that managed by the Bureau of Economic Analysis (BEA) for the United States); (ii) digital files obtained from the United Nations (UN); (iii) digital files obtained from the Organization for Economic Cooperation and Development (OECD); (iv) physical books published by the UN; and (v) physical books published by the OECD. Over time and across countries there are some differences in methodologies, but our data generally conform to System of National Accounts (SNA) standards.⁴

⁴To merge the data, we begin by using any values we are able to obtain from the Internet. This is our preferred source as it is the most likely to include data revisions. We then rank the digital files by the number of available years of data for each country and use these sources (in order) when the preferred sources lack data. Lastly, we similarly rank and use the printed sources. We only connect series together from multiple sources if they satisfy a “smooth pasting” condition whereby measures for any given variable from overlapping years for both sources are sufficiently close to each other. While there are some exceptions, this procedure typically implies that one or

Economic activity is divided in the SNA into the corporate (C), household (H), and government (G) sectors. The household sector includes unincorporated businesses, sole proprietors, non-profits serving households, and the actual and imputed rental income accruing to non-corporate owners of housing. The corporate sector includes financial and non-financial corporations. Nominal GDP Y equals the sum of sectoral gross value added Q (final output less intermediate consumption) and taxes net of subsidies on products:⁵

$$Y = Q_C + Q_H + Q_G + \text{Tax}_{\text{products}}. \quad (1)$$

The aggregate labor share equals total compensation of labor across all three sectors divided by GDP, or WN/Y , where W equals the average wage and N equals hours worked. Corporate gross value added Q_C equals the sum of compensation paid to labor $W_C N_C$, taxes net of subsidies on production (including items such as corporate income and property taxes), and gross operating surplus (including items such as interests on loans, retained earnings, and dividend payments).

For most of our analyses, we focus on the labor share in the corporate sector, $W_C N_C / Q_C$.⁶ The labor share measured within the corporate sector is not impacted by the statistical imputation of wages from the combined capital and labor income earned by sole proprietors and unincorporated enterprises, highlighted by Gollin (2002) as problematic for the consistent measurement of the labor share. Additionally, we find the focus on the corporate labor share theoretically appealing given difficulties in specifying a production function and optimization problem for the government.

For all our analyses, we start in 1975 and only include countries that have at least 15 two sources contribute the bulk of the data for any given country. These key sources do, however, differ across countries. We refer the reader to the SNA Section of the United Nations Statistics Division and to Lequiller and Blades (2006) for the most detailed descriptions of how national accounts are constructed and harmonized to meet SNA standards.

⁵This is true for a large majority of countries but there are several exceptions. For example, BEA accounts for the United States incorporate net taxes on products into the gross value added of each of the three sectors and also include unincorporated enterprises in what they call the business sector (rather than in the household sector). National accounts for Germany and China also include net taxes in sectoral gross value added.

⁶According to the SNA, compensation of employees includes wages and salaries in cash, wages and salaries in kind, and employers' social contributions for sickness, accidents, and retirement (to social security funds and insurance enterprises). Though the treatment of gains associated with the exercise of stock options is subject to data availability and is not uniform across countries, most developed countries try to account for the value of stock options granted to employees as part of labor compensation (Lequiller, 2002).

years of data. The resulting dataset contains corporate-sector-level information on the income structure of 56 countries for various years between 1975 and 2012. This is a significant increase in coverage relative to what is readily downloadable from the UN and OECD, but below we also report similar results when only using those standard data sources.

The solid line in Figure 1 shows the evolution of the global corporate labor share in our data by plotting the year fixed effects from a least-squares regression of the corporate labor share on country and year fixed effects. The regression includes country fixed effects to eliminate the influence of countries entering and exiting our dataset. We weight observations by corporate gross value added measured in U.S. dollars at market exchange rates. We normalize the fixed effects such that they equal the level of the corporate labor share in our dataset in 1975. From a level of roughly 65 percent, the global corporate labor share has exhibited a relatively steady downward trend, reaching about 60 percent at the end of the sample. The dashed line plots the fixed effects from an equivalent regression and shows that labor's share of the overall economy also declined globally.⁷ Unless otherwise noted, we refer to measures taken from the corporate sector when referring to the labor share below.

Figure 2 shows that the decline in the labor share occurred in each of the four largest economies in the world. The dashed lines plot linear trends estimated using all available data since 1975. All four trends are downward sloping and statistically significant at the 1 percent level. In fact, most countries in the world experienced this decline. Figure 3 shows the slope of these linear trends for all 56 countries with data available for at least 15 years. The coefficients are scaled such that the units represent the percentage point change in the labor share every 10 years. 38 countries experienced labor share declines compared to 18 which experienced increases. Of those 43 countries where the trends were statistically significant at the 5 percent level, the labor share declined in 34 of them. The largest eight economies are shaded, and with the United Kingdom as the only exception, they all experienced statistically significant declines.

The decline in the global labor share reflects declines in the large majority of countries

⁷The level of overall labor share is about 10 percentage points lower than the corporate labor share in part due to the inclusion of taxes in the denominator. There are a number of countries which lack the data required to calculate the corporate labor share, but which have data on the overall labor share. In such cases, we use the aggregate figures but scale them up by the average global ratio of corporate to overall labor share found in the dataset.

around the world and is not simply a reflection of trends in a few big countries. In fact, even looking at regional data within the United States, we find that the decline is similarly broad-based. We calculate the labor share for all U.S. states plus the District of Columbia by dividing total compensation by value added in the BEA's state-level GDP data (these data do not isolate the corporate sector).⁸ In parallel to the country trends plotted in Figure 3, Figure 4 shows labor share declines in the majority of U.S. states with 34 states experiencing labor share declines compared to 17 which experienced increases. Of those 38 states where the trends were statistically significant at the 5 percent level, the labor share declined in 27 of them.

Returning to the global analysis, we now ask how much of the global labor share decline reflects declines within industries and how much reflects changes in industrial composition. For example, the labor share in manufacturing is typically higher than in finance and business services. Does the decline in the labor share simply reflect the fact that manufacturing's share of economic activity has fallen while the share of economic activity in services has risen?

To answer this question, we use the EU KLEMS (KLEMS) dataset. It is available for far fewer countries and does not allow for a focus on the corporate sector, but it does allow us to construct labor shares for each country in commonly defined industries of varying granularity.⁹ Figure 5 plots labor share trends, scaled to equal the percentage point change per decade, for 10 non-overlapping industries that aggregate to the overall economy. For each industry we estimate a linear trend from a regression of the labor share that includes country fixed effects and weights countries by their value added. The labor share significantly declines in 6 of the 8 industries with statistically significant trends. The remaining two exhibit statistically insignificant declines in their labor shares.

Finally, we more formally address the question of how much of the change in the labor share is due to changing sizes of industries with different levels of labor shares and how much is due to changes in labor shares within those industries. We write the standard within-between

⁸The data are available prior to and including 1997 using the SIC classification system and from 1997 onward using the NAICS classification system. We match the 1997 values at the most disaggregated level and scale the SIC values by the discrepancy. This effectively shifts the earlier values of overall U.S. labor share up by about 1 percentage point and results in a continuous series which we then use.

⁹These data end in 2007, several years earlier than our baseline dataset. When we overlap the KLEMS data with our dataset, we find that the country labor share trends are quite similar for the years in which both sources exist. The correlation of country trends is over 90 percent if one excludes the two large outliers of Luxembourg and Ireland.

accounting decomposition for each country i across 10 industries k :

$$\Delta s_{Li} = \underbrace{\sum_k \bar{\omega}_{i,k} \Delta s_{Li,k}}_{\text{Within-Industry}} + \underbrace{\sum_k \bar{s}_{Li,k} \Delta \omega_{i,k}}_{\text{Between-Industry}}, \quad (2)$$

where $\omega_{i,k}$ denotes industry k 's share in country i 's value added, \bar{x} denotes the arithmetic mean of the variable x , and Δx denotes the estimated linear trend in x .

Figure 6 plots labor share trends, the left-hand side of equation (2), in the horizontal axis against the within-industry component, the first term on the right-hand-side of equation (2), in the vertical axis. With a few exceptions, countries in Figure 6 are aligned along the 45 degree line, implying that labor share declines are predominantly driven by the within-industry component. Further, critical for our cross-sectional analyses below, cross-country variation in labor share trends is largely explained by cross-country variation in the within-industry component. When we add the within-industry component across all countries and divide by the sum of the total components, we conclude that more than 90 percent of the labor share decline reflects within-industry declines.¹⁰

The prominence of the within-industry component is also interesting as it rules out otherwise plausible stories related to the increasing trade integration of China or globalization more generally. For example, imagine a simple two-country Heckscher-Ohlin model with Cobb-Douglas production in two sectors with different labor shares. Compared to autarky, the relatively capital-abundant economy will in the free-trade equilibrium allocate a larger share of its inputs to the production of the lower labor share industry. If we think of China as the relatively labor-abundant economy opening up to trade, one might predict a decline in the rest of the world's labor share due to this mechanism. In addition to the fact that we document a labor share decline in China itself, this trade-induced compositional change is unlikely to explain the labor share decline in other countries because it counterfactually implies an important role for the between component in equation (2).¹¹

¹⁰This statistic will vary with the fineness of the industry definition. If we calculate this statistic using an alternative decomposition available in KLEMS containing 23, rather than 10, industries, we attribute roughly 85 percent of the decline to the within-industry component.

¹¹We calculated each country's trends in imports, exports, and overall trade both bilaterally with China and multilaterally, where the change in flows are expressed relative the country's GDP. These measures of increasing

2.2 Declining Prices of Investment Goods

In parallel with these large and broad trends in the labor share, the price of investment relative to consumption goods has also experienced a pervasive decline. For our cross-country analyses, we measure the relative price of investment goods in two datasets which offer different costs and benefits. Our first source is the Penn World Tables (PWT, Mark 7.1), which offers measures at a point in time of the relative price levels of investment and consumption goods for many countries around the world. The PWT data are translated using investment-specific and consumption-specific purchasing power parity exchange rates, which is undesirable for our exercise because we wish to know the price of investment relative to consumption that a domestic producer faces. We therefore follow Restuccia and Urrutia (2001) and divide the PWT relative price of investment of each country by the PWT relative price of investment in the United States. We then multiply this ratio by the ratio of the investment price deflator to the personal consumption expenditure deflator for the United States, obtained from the BEA. This procedure yields for each country the relative price of investment measured at domestic prices.

The PWT data cover a large set of countries and in some cases extend back to 1950. Further, by combining the PWT's information on the cross-section of international prices with time-series information on the relative price of investment from the United States, the constructed series are insensitive to cross-country differences in methodologies used to construct investment price deflators. If the U.S. BEA employs hedonic adjustments to properly capture changes in the quality of computers, for example, then our methodology will imply that this same adjustment is implicitly captured for all countries in the data.

The solid line in Figure 7 plots year fixed effects from a regression of the log relative price of investment in the PWT dataset after absorbing country dummies. The regressions are weighted by GDP and the fixed effects are normalized to equal 0 in 1980. The series exhibits a mild decline from 1950 to 1980, trending downward about 0.02 log point per decade. Consistent with Fisher (2006), however, the series exhibits a clear break around 1980 and declines at a rate closer to 0.1 log point per decade after 1980. This steep downward trend occurred all around the world. Of the countries with at least 15 years of data on both the PWT relative price of exposure to China and to the rest of the world do not generally correlate with declines in the labor share.

investment and on labor shares, 42 experienced declines in the relative price of investment since 1975 compared to only 13 which experienced an increase.

As a second measure of the relative price of investment goods, we take the ratio in each country of the fixed investment deflator to the consumer price index, obtained from the Economist Intelligence Unit (EIU). These data rely more on the individual statistical agencies in each country but offer the benefit of properly capturing differences in the composition of investment spending across countries. The short-dashed line in Figure 7 plots the equivalent series of year fixed effects as the solid line but estimated using the EIU data, which is only available from 1980. The global trends are highly similar, with the decline in the relative price of investment being slightly less steep in the EIU data. The decline in the EIU data is also widespread across countries, with 38 experiencing declines compared to 15 with increases. There are differences in country and time coverage between the PWT and EIU sources, but when we consider overlapping country-year observations, we measure a cross-country correlation of about 0.8 between the trends found in the two datasets.

Finally, our industry-level analyses use KLEMS data on investment and output prices in each industry. Though the key variation we will use from this dataset will be cross-industry differences in declines in the relative price of investment, we demonstrate comparability with the other sources by plotting time fixed effects from regressions using the country-level relative price of investment in KLEMS. The long-dashed line in Figure 7 shows an increasing trend which sharply reverses in the early 1980s, consistent with the timing of the decline captured in the other data sources. All countries in KLEMS with sufficient data for this analysis exhibit a declining relative price of investment and the correlation of these trends with the trends found in PWT and EIU across countries is nearly 0.6.

3 A Model of the Labor Share

We now develop a model that relates the labor share to the relative price of investment goods as well as to other macroeconomic variables such as price markups and factor-augmenting technology. We consider a two-sector economic environment in which final consumption and investment goods are produced by combining intermediate inputs using a CES technology.

Time is discrete and the horizon is infinite, $t = 0, 1, 2, \dots$. There is no uncertainty and all economic agents have perfect foresight. All payments in this economy are made in terms of the final consumption good, which is the numeraire.

3.1 Final Consumption Good

Competitive producers assemble the final consumption good C_t from a continuum of intermediate inputs $z \in [0, 1]$ and sell it to the household at a price P_t^c . They produce final consumption with the technology:

$$C_t = \left(\int_0^1 c_t(z)^{\frac{\epsilon_t-1}{\epsilon_t}} dz \right)^{\frac{\epsilon_t}{\epsilon_t-1}}, \quad (3)$$

where $c_t(z)$ denotes the quantity of input z used in production of the final consumption good and $\epsilon_t > 1$ denotes the elasticity of substitution between input varieties. The consumption good producers purchase these inputs at prices $p_t(z)$ from monopolistically competitive firms that charge a markup over marginal cost that depends on ϵ_t . To capture changes in markups over time, we allow the elasticity of substitution across varieties to vary over time.

Cost-minimization implies that the demand for input variety z for use in producing the consumption good is $c_t(z) = (p_t(z)/P_t^c)^{-\epsilon_t} C_t$. The final consumption good is the numeraire and has a price of one. It is competitively produced, so its price equals the marginal cost of production:

$$P_t^c = \left(\int_0^1 p_t(z)^{1-\epsilon_t} dz \right)^{\frac{1}{1-\epsilon_t}} = 1. \quad (4)$$

3.2 Final Investment Good

Competitive producers assemble the final investment good X_t from the same continuum of intermediate inputs z :

$$X_t = \left(\frac{1}{\xi_t} \right) \left(\int_0^1 x_t(z)^{\frac{\epsilon_t-1}{\epsilon_t}} dz \right)^{\frac{\epsilon_t}{\epsilon_t-1}}. \quad (5)$$

The exogenous variable ξ_t denotes the technology level in the production of the consumption good relative to the investment good. A decline of ξ_t implies an improvement in the technology of producing the investment good relative to the consumption good.

Since firms in the final investment good sector are competitive, the price of the final invest-

ment good equals the marginal cost of production, $P_t^x = \xi_t \left(\int_0^1 p_t(z)^{1-\epsilon_t} dz \right)^{\frac{1}{1-\epsilon_t}} = \xi_t$. We refer to $\xi_t = P_t^x/P_t^c$ as the relative price of investment, which declines whenever technology in the investment good sector improves relative to the consumption good sector. Finally, demand for input variety z for use in production of the investment good is given by $x_t(z) = \xi_t p_t(z)^{-\epsilon_t} X_t$.

3.3 Producers of Intermediate Inputs

The producer of intermediate input variety z operates a constant returns to scale technology in capital and labor inputs to produce output sold to both consumption and investment good producers, $y_t(z) = F(k_t(z), n_t(z))$. Capital is rented at rate R_t and labor is rented at a price W_t from the household. Producers of intermediate inputs take input prices and aggregate demand, $Y_t = C_t + \xi_t X_t$, as given.

The profit-maximization problem of the producer of intermediate input z is:

$$\max_{p_t(z), y_t(z), k_t(z), n_t(z)} \Pi_t(z) = p_t(z)y_t(z) - R_t k_t(z) - W_t n_t(z), \quad (6)$$

subject to:

$$y_t(z) = c_t(z) + x_t(z) = p_t(z)^{-\epsilon_t} (C_t + \xi_t X_t) = p_t(z)^{-\epsilon_t} Y_t. \quad (7)$$

The first-order condition with respect to capital is $p_t(z)F_{k,t}(z) = \mu_t R_t$ and with respect to labor is $p_t(z)F_{n,t}(z) = \mu_t W_t$. Firms set the marginal revenue product of factors as a markup $\mu_t = \epsilon_t / (\epsilon_t - 1)$ over factor prices.

3.4 Household

The household derives utility from consumption goods and disutility from supplying labor. It purchases consumption and investment goods from final good producers at prices one and ξ_t respectively. The household uses the investment good to augment the physical capital stock and rents capital to producers of intermediate goods at a rental rate R_t . The household owns all firms in the economy and receives their profits as dividends in every period. The household supplies labor to intermediate input producers at a wage W_t . It can also hold some asset B_t that pays a real interest rate r_t and is in zero net supply. Denoting by χ_t a household preference

shifter and by β the discount factor, the problem of the household in some period t_0 is:

$$\max_{\{C_t, \{n_t(z)\}, X_t, K_{t+1}, B_{t+1}\}_{t=t_0}^{\infty}} \sum_{t=t_0}^{\infty} \beta^{t-t_0} V(C_t, N_t; \chi_t), \quad (8)$$

subject to initial capital K_0 and assets B_0 , the capital accumulation equation, $K_{t+1} = (1 - \delta)K_t + X_t$, and the household budget constraint:

$$C_t + \xi_t X_t + B_{t+1} - (1 + r_t)B_t = \int_0^1 (W_t n_t(z) + R_t k_t(z) + \Pi_t(z)) dz. \quad (9)$$

Aggregate labor supplied by the household is $N_t = \int_0^1 n_t(z) dz$ and the aggregate capital stock is $K_t = \int_0^1 k_t(z) dz$.

Household optimization implies a standard Euler equation for consumption across time and a standard intraperiod condition for leisure and consumption. Finally, the first-order condition with respect to capital is given by:

$$R_{t+1} = \xi_t (1 + r_{t+1}) - \xi_{t+1} (1 - \delta), \quad (10)$$

where $1 + r_{t+1} = (1/\beta) (V_C(C_t, N_t)/V_C(C_{t+1}, N_{t+1}))$ denotes the gross real interest rate. This condition says that the household invests in physical capital up to the point where the marginal benefit of investing in capital (the rental rate) equals the marginal cost of investing in capital.

3.5 Equilibrium

We define an equilibrium for this economy as a sequence of prices and quantities such that, given a sequence of exogenous variables: (i) the household maximizes its utility; (ii) final producers of the consumption good minimize their costs; (iii) final producers of the investment good minimize their costs; (iv) each producer of input variety z maximizes profits; and (v) markets for labor, capital, assets, consumption, investment, and intermediate inputs clear in every date. We define a steady state as an equilibrium in which all variables are constant over time.

The equilibrium of the model is symmetric, with $p_t(z) = P_t^c = 1$, $k_t(z) = K_t$, $n_t(z) = N_t$,

$c_t(z) = C_t$, $x_t(z) = \xi_t X_t$, $y_t(z) = Y_t = C_t + \xi_t X_t$, and $Y_t = F(K_t, N_t)$, where N_t and K_t are total labor and total capital. The share of income paid as wages for labor services, rentals for capital, and profits are given by:

$$s_{L,t} = \frac{W_t N_t}{Y_t} = \left(\frac{1}{\mu_t} \right) \left(\frac{W_t N_t}{W_t N_t + R_t K_t} \right), \quad (11)$$

$$s_{K,t} = \frac{R_t K_t}{Y_t} = \left(\frac{1}{\mu_t} \right) \left(\frac{R_t K_t}{W_t N_t + R_t K_t} \right), \quad (12)$$

$$s_{\Pi,t} = \frac{\Pi_t}{Y_t} = 1 - \frac{1}{\mu_t}, \quad (13)$$

where $s_{L,t} + s_{K,t} + s_{\Pi,t} = 1$.

3.6 The Production Function

We assume intermediate inputs are produced with a CES production function:

$$Y_t = F(K_t, N_t) = \left(\alpha_k (A_{K,t} K_t)^{\frac{\sigma-1}{\sigma}} + (1 - \alpha_k) (A_{N,t} N_t)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \quad (14)$$

where σ denotes the elasticity of substitution between capital and labor in production and α_k is a distribution parameter. We let $A_{K,t}$ and $A_{N,t}$ denote capital-augmenting and labor-augmenting technology respectively. The limit of the CES production function as σ approaches 1 is the Cobb-Douglas production function, $F(K_t, N_t) = (A_{K,t} K_t)^{\alpha_k} (A_{N,t} N_t)^{1-\alpha_k}$. With the production function (14), the firm's first-order conditions with respect to capital and labor are:

$$F_{K,t} = \alpha_k A_{K,t}^{\frac{\sigma-1}{\sigma}} \left(\frac{Y_t}{K_t} \right)^{\frac{1}{\sigma}} = \mu_t R_t, \quad (15)$$

$$F_{N,t} = (1 - \alpha_k) A_{N,t}^{\frac{\sigma-1}{\sigma}} \left(\frac{Y_t}{N_t} \right)^{\frac{1}{\sigma}} = \mu_t W_t. \quad (16)$$

3.7 The Labor Share

We can now discuss the determinants of the labor share of income $s_{L,t}$. Using the first-order condition for capital in equation (15) and the definitions of the income shares in equations

(11)-(13), we derive an expression that relates the labor share to markups, capital-augmenting technology, and the rental rate of capital for some value of the elasticity of substitution σ and the distribution parameter α_k :

$$1 - s_{L,t}\mu_t = \alpha_k^\sigma \left(\frac{A_{K,t}}{\mu_t R_t} \right)^{\sigma-1}. \quad (17)$$

In the limiting case of a Cobb-Douglas production function where $\sigma = 1$, the labor share of income simply becomes $s_{L,t} = (1 - \alpha_k)/\mu_t$. Therefore, with Cobb-Douglas production the labor share of income varies over time only when markups vary over time and the labor share of cost, $s_{L,t}\mu_t = W_t N_t / (W_t N_t + R_t K_t)$, is constant.

As we discuss below, our estimation strategy focuses on labor share trends as this allows us to control for substantial cross-country heterogeneity both in economic parameters (e.g. initial level of technology) and in measurement practices. We write equation (17) in changes between two arbitrary periods t and $s > t$ as:

$$\left(\frac{1}{1 - s_L \mu} \right) (1 - s_L (1 + \hat{s}_L) \mu (1 + \hat{\mu})) = \left(\frac{1 + \hat{A}_K}{(1 + \hat{\mu})(1 + \hat{R})} \right)^{\sigma-1}, \quad (18)$$

where $\hat{Z} = Z_s/Z_t - 1$ denotes the percent change of some variable Z between periods t and s and where for notational convenience we drop subscripts from values corresponding to the initial period t .

For any value of the elasticity of substitution σ , given initial conditions for the labor share s_L and markups μ , equation (18) implies that information on changes in markups $\hat{\mu}$, in capital-augmenting technology \hat{A}_K , and in the rental rate of capital \hat{R} is sufficient to pin down changes in the labor share \hat{s}_L . Additional knowledge of wage and labor-augmenting technology is not informative for understanding changes in the labor share.

We derive (18) from the first-order condition with respect to capital, equation (15), rather than the first-order condition with respect to labor, equation (16). We choose this approach because we estimate equation (18) using long-run trends spanning 15 to 37 years of data and treat these trends as steady state to steady state transitions. In steady state, equation (10)

for capital becomes $R = \xi(1/\beta - 1 + \delta)$. Assuming constant discount factors β and constant depreciation rates δ over time (but not necessarily across countries), this allows us to directly equate the growth in the rental rate of capital to the growth in the relative price of investment, $\hat{R} = \hat{\xi}$. Internationally comparable and high quality data on $\hat{\xi}$ is more readily available than data on wage growth.

To summarize, holding fixed the discount factor and the depreciation rate of capital, the labor share will only change in the steady state of our model if A_K , μ , or ξ change. This general result argues against the relevance for the long-run decline in labor share of a large set of factors such as wage markups, labor income taxes and other labor supply shocks, and government spending shocks that do not directly affect the production function.¹²

4 The Elasticity of Substitution

In this section we confront equation (18) with our data to estimate the elasticity of substitution between capital and labor σ . We start by focusing only on trends in the relative price of investment and abstract from markups and capital-augmenting technological progress because we lack direct measurements on them. Next, we introduce assumptions that allow us to impute time-varying markups from data on investment spending and we quantify the sensitivity of our estimates to capital-augmenting technological progress. Finally, we estimate σ with production functions that allow for differential substitutability between capital and two types of labor. Equipped with our estimates of σ , in Section 5 we quantify the effect of the decline in the relative price of investment goods on the labor share and explore the broader macroeconomic and welfare implications of our findings.

We measure the percent change of all variables (corresponding to our “hat” notation) as the linear trend in the log of the variable using all available data. We replace the variables in levels with their average values in our sample. As discussed above, we focus on long-run trends and below will think of them as capturing movements from an initial to a final steady

¹²For this reason our model abstracts from wage markups, labor income taxes, and household shocks to the utility of leisure relative to consumption. All these factors are isomorphic to a change in χ in household preferences in the sense that none can affect the labor share of income, the capital-to-labor ratio, and the wage in the steady state of the model. Steady state wages do not depend on χ because with constant returns to scale the de-accumulation of capital implies a shift in labor demand that exactly offsets any shifts in labor supply.

state. Assuming a constant household discount factor β and depreciation rate δ , this allows us to substitute the percent change across steady states in the rental rate of capital with the percent change in the relative price of investment goods, $\hat{R} = \hat{\xi}$. However, we also show that our results are robust when we allow for trends in depreciation rates.

4.1 Relative Price of Investment

We start with equation (18) and set $\mu = 1$, $\hat{\mu} = 0$, and $\hat{A}_K = 0$. We take a linear approximation around $\hat{\xi} = 0$ and add a constant and an idiosyncratic error term to obtain our estimating equation:

$$\frac{s_{L,j}}{1 - s_{L,j}} \hat{s}_{L,j} = \gamma + (\sigma - 1) \hat{\xi}_j + u_j, \quad (19)$$

where j denotes observations. The intuition behind equation (19) is simple. Absent economic profits and capital-augmenting technology growth, a positive relationship between trends in the relative price of investment $\hat{\xi}_j$ and trends in the labor share $\hat{s}_{L,j}$ is possible only when the elasticity of substitution between capital and labor σ exceeds one. In that case, a decrease of the cost of capital due to decreases in the relative price of investment induces firms to substitute away from labor and toward capital to such an extent that it drives down the labor share. If trends in the relative price of investment were unrelated to trends in the labor share this would imply a Cobb-Douglas production function with $\sigma = 1$. A negative relationship between trends in the relative price of investment and trends in the labor share would imply $\sigma < 1$.

We emphasize that the identification of σ in equation (19) comes from the cross-sectional variation of trends in labor shares and trends in the price of investment. Specifically, adding a constant γ to the regression allows us to control for global factors that affect all countries. For example, imagine that all countries experienced declining trends both in labor shares and in the price of investment goods, but the extent of the labor share decline across countries was unrelated to the extent of the decline in the price of investment across countries. In such a case, we would estimate $\sigma = 1$. Put differently, even though both the labor share and the relative price of investment declined over time for the typical country in our sample, our estimation using cross-sectional variation could hypothetically produce an elasticity estimate less than one. In this sense, our methodology for estimating the elasticity does not incorporate information

from the global trend that we hope to explain.

Given the small number of observations, estimation of equation (19) is particularly sensitive to outliers. We standardize our selection and treatment of outliers by generating “robust regression” estimates, which place less weight on extreme values that are identified endogenously during the estimation.¹³ In practice, the primary difference compared with OLS estimates for most of our results is that the robust regressions endogenously assign very little weight to Kazakhstan, Kyrgyzstan, and Niger.

Table 1 presents our baseline estimates of σ from equation (19). In the first four rows we estimate σ using our country dataset, with j indexing the country observations. The first two rows under the column labeled “Labor Share” list the source as “KN Merged” to refer to our full dataset described above and rows (iii)-(iv) list the source as “OECD and UN” to refer to a similarly constructed dataset which only uses data easily downloadable online from the UN and OECD (i.e. it discards the data we collected ourselves from country-specific Internet sources and physical books). The column labeled “Investment Price” alternates in these rows between PWT and EIU to indicate the data source used for $\hat{\xi}$. These first four specifications produce highly similar results, with the estimated σ always being greater than 1.

To visualize these results, we plot in Figure 8 the cross-country relationship between the left-hand side of equation (19) and the trends in the price of investment goods.¹⁴ We use the “KN Merged” dataset for the labor share, the PWT for the price of investment goods, and drop the three outliers discussed above to ensure that the plotted least-squares line closely corresponds to the estimate presented in the first row of Table 1. Countries with larger relative price of investment declines also experienced larger labor share declines, which results in a statistically significant positive slope of 0.30 and an implied elasticity of 1.30.

¹³This regression is implemented with the command *rreg* in the statistical package STATA. The idea behind a robust regression is to weight less observations which lie further from the regression line. The method starts by dropping observations with Cook’s distance greater than one. Then an iterative process calculates weights based on absolute residuals. The process stops when the maximum change between the weights from one iteration to the next is below some tolerance level. All regressions reported in this paper use this procedure.

¹⁴In our closed economy model, the variation in the relative price of investment across countries is tied to different sectoral productivity shocks in each country. This intuition can be extended to the open economy if one thinks of investment goods as freely tradable across countries and consumption goods as being at least partly non-tradable. As in Balassa-Samuelson, prices of investment goods would be equalized across countries but differential technology growth in traded relative to non-traded sectors will shift the price of investment relative to consumption in each country. Aside from productivity shocks, cross-country differences in the relative price of investment may also come from differences in the scale and timing of reductions in tariffs and other trade frictions.

In the last two rows of the table, we estimate σ using the KLEMS dataset, with j indexing country-industry observations for 10 major industries and 14 developed economies. We include both country and industry fixed effects in the regression. Unfortunately, we cannot isolate the corporate labor share in the KLEMS data but we can measure the labor share in two alternative ways, which we label as “KLEMS 1” and “KLEMS 2.”¹⁵ As shown in rows (v) and (vi) of Table 1, both elasticity estimates from these data are also significantly greater than 1.

Our estimates allow for substantial heterogeneity across countries and industries. Initial differences in technology, wages, relative prices of investment, preferences, and depreciation rates are all captured by the initial level of the labor share $s_{L,j}$, which is allowed to vary across observations in the left-hand side of equation (19). If our trends capture a steady state to steady state transition, and assuming constant discount factor and depreciation rates over time (but not necessarily across countries), our estimates of σ allow for differences in the growth of wages, labor-augmenting technology, and anything other than capital-augmenting technology and markups, which we address below. The only substantial restriction we are imposing is a common elasticity of substitution σ between capital and labor across countries or industries.¹⁶

Up to now our analysis imposes constant depreciation rates over time. We note that if industries or countries which experienced larger declines in their relative price of investment are systematically shifting the composition of their capital stock towards capital goods with higher depreciation rates (e.g. computers), then our estimated σ is generally biased downward. This is because a given labor share decline would be associated with smaller declines in the rental rate of capital, therefore increasing the elasticity of substitution between capital and labor necessary to generate the positive relationship between labor share trends and trends in

¹⁵In KLEMS 1 the labor share at the country-industry level is defined as compensation of employees divided by gross value added, whereas in KLEMS 2 the labor share is defined as labor compensation divided by gross value added. Labor compensation equals compensation of employees plus a fraction of other taxes on production and an imputation for the income of self-employed using the category operating surplus / mixed income. Our preferred notion of corporate labor compensation falls in between these two as it includes other taxes on production but excludes the self-employed income imputation. The results reported in Section 2.1 use the KLEMS 1 definition of the labor share. They are qualitatively similar if we use KLEMS 2.

¹⁶This restriction is tied to our strategy of using cross-sectional variation to identify σ . We prefer using cross-sectional variation as it eliminates the influence of global trends, but to check the robustness of our results, we have also estimated country or country-industry specific σ 's based on long-term time-series variation. We generate an estimate $\tilde{\sigma}_j$ for each observation by dividing the left-hand side of equation (19) by the right-hand side of equation (19) after dropping the constant and the error term, $\tilde{\sigma}_j = 1 + (s_{L,j}\hat{s}_{L,j}) / \left((1 - s_{L,j})\hat{\xi}_j \right)$. The median elasticity estimate across countries is close to our baseline value of 1.25.

the rental rate.¹⁷

We empirically assess the extent of this bias in the KLEMS dataset because it includes estimates of depreciation and capital stocks at the country-industry level. We modify the expression for the rental rate by adding a subscript to the depreciation rate, $R_j = \xi_j(1/\beta - 1 + \delta_j)$, which means we can no longer equate growth in the rental rate with growth in the relative price of investment. We assume $\beta = 1/(1 + 0.10)$, measure δ_j and $\hat{\delta}_j$ in the KLEMS data, and calculate alternative values for \hat{R}_j . Our estimated elasticity does not change meaningfully. The values in rows (v) and (vi) of Table 1, 1.17 and 1.49, increase to 1.19 and 1.51 respectively when taking into account heterogeneous and time-varying depreciation rates.

To summarize, our six baseline estimates of σ average roughly 1.25 and are all statistically different from one at the 5 percent significance level. Section 5 analyzes in greater detail the implications of this value. Here we note that using $\sigma = 1.25$ together with calibrated global values for s_L and $\hat{\xi}$ as inputs in equation (19) implies that roughly half of the global decline in the labor share is explained by the decline in the relative price of investment.

4.2 Markups

We now allow for the possibility that our estimated elasticity is impacted by markups μ_j . Since the labor, capital, and profit shares add up to one, measures of μ_j and $\hat{\mu}_j$ can be obtained with the additional knowledge of capital share levels $s_{K,j}$ and changes $\hat{s}_{K,j}$. Under the assumption that the trends we are interested in reflect movements from one steady state to another, and holding constant the discount factor β and depreciation δ over time, we can simply calculate the trend in the capital share as the trend in the nominal investment rate, $\hat{s}_{K,j} = (\widehat{R_j K_j / Y_j}) = (\widehat{\xi_j X_j / Y_j})$. We calculate the level of the capital share from the steady state condition $s_{K,j} = \left[\left(\frac{1}{\beta} - 1 + \delta \right) / \delta \right] [\xi_j X_j / Y_j]$ using data on the average nominal investment rate of countries in the sample and assuming homogeneous discount factors and depreciation rates.¹⁸ We then use

¹⁷This argument assumes that the increase of depreciation does not more than offset the decline in the price of investment, which is empirically true for the great majority of observations.

¹⁸We set β as above and let $\delta = 0.10$. The average initial profit share varies across data sources but generally is less than 5 percent. For some countries this methodology implies negative profit shares. We have considered various ad hoc treatments for such observations that bound profit shares above zero. Results vary slightly depending on the methodology, but consistently produce average estimates of σ that are close to our baseline estimates. We only do this analysis at the country level because when applied to the more disaggregated industry data, our procedure often imputes implausibly large capital shares.

$\mu_j = 1/(s_{L,j} + s_{K,j})$ to calculate the levels and changes in markups.

As a simple example of how markups could bias our elasticity estimates from Section 4.1, consider the Cobb-Douglas case which has an elasticity of substitution between capital and labor equal to one. Though the labor share of costs is constant, markups generate a wedge between costs and revenues and can cause movements in the labor share of income. If markups increased more in countries with larger declines in the relative price of investment, then labor shares would decline more in these countries. Without taking into account markup variation, our baseline procedure would incorrectly estimate an elasticity greater than one.

Using our estimated trends in capital shares, we can rule out this possibility. If the true elasticity is one and markup growth drives all labor share movements, then labor and capital shares of income change by the same percent. Figure 9 plots the percent change in labor shares against the percent change in capital shares. Countries do not lie along the 45 degree line. Rather, the best-fit line has essentially zero slope. This finding provides strong evidence against the possibility that we estimate a non-unitary elasticity purely due to the bias from markups.

To more generally assess the impact of markups on our elasticity estimate, we derive a modified estimating equation. We continue to assume no capital-augmenting technological progress and set $\hat{A}_K = 0$ in equation (18). Taking a linear approximation around $\hat{\xi} = 0$ and $\hat{\mu} = 0$, adding a constant and an idiosyncratic error term, we obtain:

$$\left(\frac{s_{L,j}\mu_j}{1 - s_{L,j}\mu_j} \right) ((1 + \hat{s}_{L,j})(1 + \hat{\mu}_j) - 1) = \gamma + (\sigma - 1) (\hat{\xi}_j + \hat{\mu}_j) + u_j. \quad (20)$$

Table 2 reports our estimates of σ from equation (20). As before, we report values across multiple data sources on the labor share and the relative price of investment, but here we additionally report whether we use total or corporate nominal investment data in constructing markups. We consider these estimates less reliable than our baseline estimates as they require richer assumptions and make use of imputed values as opposed to direct measurements. Nonetheless, our results in Table 2 are highly similar to our baseline results in Table 1.

The similarity in our estimated σ when including or excluding information on markups does not imply that markups played no role in labor share movements. In fact, since we

do not generally find significant increases in capital shares, it must be the case that some of labor share's decline is attributable to markup growth. Given that our elasticity estimates remain unchanged, however, we maintain our conclusion that the decline in the relative price of investment explains roughly half of the labor share decline.

4.3 Capital-Augmenting Technological Progress

Given the difficulty of properly measuring capital-augmenting technology growth, our baseline analysis assumed that it is orthogonal to changes in the relative price of investment.¹⁹ In this section, we gauge how large a bias this assumption could plausibly create. We conclude that while the cross-country pattern of capital-augmenting technology change may bias upward our elasticity estimates, the bias is unlikely to be quantitatively large.

Consider the baseline estimating equation for the labor share, modified to allow for capital-augmenting technology growth:

$$\frac{s_{L,j}}{1 - s_{L,j}} \hat{s}_{L,j} = \gamma + (\sigma - 1) \hat{\xi}_j + (1 - \sigma) \hat{A}_{K,j} + u_j. \quad (21)$$

Let $\hat{\sigma}$ denote our estimate of the elasticity of substitution when omitting capital-augmenting technology growth from the regression and let σ denote the true elasticity of substitution. Then the bias is given by:

$$\hat{\sigma} - \sigma = (1 - \sigma) \text{corr} \left(\hat{A}_K, \hat{\xi} \right) \frac{\text{sd} \left(\hat{A}_K \right)}{\text{sd} \left(\hat{\xi} \right)}, \quad (22)$$

where $\text{corr} \left(\hat{A}_K, \hat{\xi} \right)$ denotes the correlation between capital-augmenting technology growth and changes in the relative price of investment, and $\text{sd} \left(\hat{A}_K \right)$ and $\text{sd} \left(\hat{\xi} \right)$ denote their respective standard deviations, in the cross section of countries. The bias tends to zero as the true elasticity σ approaches one.

Consider the possibility that countries that experienced the greatest declines in their relative prices of investment also experienced the highest capital-augmenting technology growth,

¹⁹Many estimates of capital-augmenting technology come from equation (18) which can be used to infer \hat{A}_K assuming no profits and a given value of the elasticity of substitution. Given that we estimate the elasticity and allow for markups, this procedure is infeasible.

$\text{corr}(\hat{A}_K, \hat{\xi}) < 0$.²⁰ Equation (22) shows that if the true elasticity of substitution is greater than one, then our estimate is upward biased ($\hat{\sigma} > \sigma$). If the true elasticity is lower than one, then our estimate is downward biased. This logic implies that the bias from capital-augmenting technology growth would never cause us to mistakenly estimate an elasticity of substitution σ that exceeds one if the true elasticity was smaller than one.

To quantify the size of the bias using equation (22), we need to specify values for the standard deviations of the relative price of investment and capital-augmenting technology growth and for the correlation between these two variables. To get a sense for these moments, we combine our PWT and EIU data on $\hat{\xi}$ with cross-country estimates of TFP growth that we use as proxies for capital-augmenting technology growth \hat{A}_K .²¹ While imperfect, this is a reasonable exercise as \hat{A}_K is the product of capital-augmenting and Hicks-neutral technology growth. We estimate $\text{corr}(\hat{A}_K, \hat{\xi}) = 0.23$, $\text{sd}(\hat{A}_K) = 0.093$ and $\text{sd}(\hat{\xi}) = 0.120$.

Given these values, and given our estimate $\hat{\sigma} = 1.25$, equation (22) implies a true elasticity of $\sigma = 1.21$. We conclude that the upward bias from capital-augmenting technology growth is small and unlikely to alter our conclusions. With alternative estimates of the covariance between \hat{A}_K and $\hat{\xi}$, one can use equation (22) to obtain different magnitudes of the bias. For the results to differ significantly from ours, however, the cross-country pattern of capital-augmenting technology growth would need to be significantly different than the pattern in these estimates of TFP growth.

4.4 Skilled vs. Unskilled Labor

Our analyses thus far assume that all labor types are equally substitutable with capital. Influential work such as Krusell, Ohanian, Rios-Rull, and Violante (2000), however, has suggested the importance of the differential substitutability of capital with different types of skill. Additionally, there have been significant trends in recent decades in the skill composition of the labor force related to changing skill premia, specialization, and international trade. We use KLEMS

²⁰We emphasize this possibility as we find it far more plausible than the alternative in which countries experiencing the greatest declines in the price of investment have the slowest capital-augmenting technology growth.

²¹We use TFP estimates from The Conference Board's Total Economy Database, which are available from 1990-2012 for about 50 countries overlapping with our dataset. The estimated moments account for the existence of outliers and represent averages between the PWT and the EIU datasets.

data to evaluate whether changes in the skill composition of the labor force in a production function with differential capital-skill substitutability alter our conclusion that the decline in the relative price of investment goods accounts for half of the decline in the global labor share.

We maintain the assumption of a homogeneous capital stock K_t but now distinguish between two types of labor, skilled S_t and unskilled U_t . Within the CES framework, there are three ways in which skilled labor, unskilled labor, and the capital stock can be nested. The first way is as in the production function (14), in which the aggregate labor input is a function of different skills, $N_t = N_t(S_t, U_t)$, and N_t and the capital stock K_t combine with a constant elasticity of substitution σ . In this case all our previous results continue to apply.

The second way to nest the three inputs is through the production function:

$$Y_t = \left(\phi_1 \left(\phi_2 K_t^{\frac{\rho-1}{\rho}} + (1 - \phi_2) S_t^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1} \frac{\sigma-1}{\sigma}} + (1 - \phi_1) U_t^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \quad (23)$$

where ρ is the elasticity of substitution between capital and skilled labor and σ is the elasticity of those factors with unskilled labor. We follow the same steps as in the two-factor case to derive the corresponding estimating equation:

$$\frac{s_{L,j}}{1 - s_{L,j}} \hat{s}_{L,j} = \gamma_c + \gamma_s + (\sigma - 1) \hat{\xi}_j + \kappa \left(\widehat{S_j/K_j} \right) + u_j, \quad (24)$$

where we continue to define the labor share as the sum of all compensation to all labor types. The term $\widehat{S_j/K_j}$ denotes the change in the ratio of skilled labor to capital. The third way to nest the three inputs is to reverse the structure in (23), with capital and unskilled labor combining with each other with an elasticity of substitution ρ and the combined input aggregating with skilled labor with an elasticity σ . This alternative production function leads to an identical estimating equation as (24), but with “ S_j ” replaced by “ U_j ”.

We estimate σ for both of these nesting structures using the KLEMS industry data, where the change in either skilled or unskilled labor relative to capital is added as a covariate. We consider two definitions of unskilled labor, one that includes KLEMS’ definitions of both “middle” and “low” skill and the other just using “low.” As with the previous industry-level estimates, we use both KLEMS definitions of the labor share and include industry and country fixed ef-

fects. Table 3 presents our estimates. Across the six specifications, the estimates for σ average 1.26, the same as our benchmark value. In all cases, the estimated σ is significantly different from one at the 10 percent level.

As with the case of markups, the similarity in our estimated σ when including or excluding the possibility of capital-skill complementarity does not imply that changes in the stock of skill played no role in labor share movements.²² Rather, the results in Table 3 simply confirm that even with these alternative production functions and taking into account the changing skill composition of labor, the decline in the relative price of investment continues to account on its own for about half of the labor share decline.

5 The Decline in the Labor Share

Figure 1 documented a 5 percentage point global decline in the labor share. Figure 7 documented a global decline in the relative price of investment goods of about 25 percent. Using cross-country variation, we estimated an elasticity of substitution between capital and labor of about 1.25. This estimate proves stable when we take into account markup variation, capital-augmenting technology growth, and changes in the skill composition of the labor force. Using this elasticity estimate and setting the global labor share to the average level in our sample, we find that the 25 percent negative shock to the relative price of investment generates roughly half of the decline in the global labor share.

Our estimates of the elasticity σ , markup growth $\hat{\mu}$, and the shock in the relative price of investment $\hat{\xi}$ have additional implications for other macroeconomic aggregates and for welfare. In this section, we solve for the general equilibrium of our model in order to consider the broader importance of our findings. We highlight that the implications of our explanation of the decline in labor share can differ starkly from those of alternative explanations.

To assess the impact of our estimated elasticity we calibrate two economies, one with CES production and $\sigma = 1.25$ and the other an otherwise identical economy but with Cobb-Douglas

²²Depending on the specification, the coefficients on $\widehat{S/K}$ or $\widehat{U/K}$ suggest that skill composition may in fact have played some role in the declining labor share. The coefficient on these covariates is a function of the two elasticities σ and ρ , the level of S/K or U/K , and the distribution parameter ϕ_2 . So this regression alone cannot be used to identify the value of ρ .

production (i.e. $\sigma = 1$).²³ The first two columns of Table 4 report the results when we introduce into the Cobb-Douglas and CES economies a 25 percent negative shock to the relative price of investment. All changes in the table are across steady states.²⁴

The first three rows show the percentage point change in factor shares. As expected, the $\hat{\xi}$ shock has no impact on the labor share in the Cobb-Douglas economy while it generates a 2.6 percentage point decline in the CES economy. Given that markups do not change, the decline in the CES case is associated with an equal percentage point increase in the capital share. In addition to the implications for labor share, a comparison of these first two columns reveals important differences for output and welfare in the economies' responses to the shock. Given the greater substitutability between capital and labor, the CES economy adjusts more to the lower cost of capital, resulting in a larger increase in the capital-to-labor ratio than in the Cobb-Douglas economy. This implies that in response to the same decline in the price of investment, the CES economy experiences higher GDP, consumption, and investment growth. Welfare, in terms of equivalent consumption units, increases by 22 percent, or 4 percentage points, more in the CES economy relative to the Cobb-Douglas economy.

Columns three and four evaluate the response of the two economies to a markup shock which increases the profit share from an initial level of 3 percent to a final level of 8 percent while holding ξ constant. Broadly in line with our results from Section 4.2, we choose the scale of this markup shock in order to generate an identical decline in the labor share as generated by the $\hat{\xi}$ shock in the CES case. In the Cobb-Douglas case, consistent with the logic presented earlier, the labor and capital shares decline by an equal percent (the values in rows (i) and (ii) are not equal as they are in percentage points). Comparing the second and fourth columns, we conclude that alternative explanations for an equal decline in the labor share entail different macroeconomic implications. If labor share declines result from declines in the relative price of investment with CES production, they are associated with significant output and welfare gains.

²³Period utility is given by $V_t = \log(C_t) - (1/2)N_t^2$. We normalize the relative price of investment goods in the initial steady state to $\xi = 1$ and set the depreciation rate at 10 percent to target a steady state investment rate of 20 percent. We set the discount factor as before to $\beta = 1/(1 + 0.10)$. Finally, we choose α_k in the production function (14) to target a 60 percent steady state labor share. By appropriately normalizing the levels of Hicks-neutral technology separately for the CES production function and for the Cobb-Douglas production function, the two economies share exactly the same initial steady state in all other variables.

²⁴Here we explore the long-run implications of changes in the labor share. Rios-Rull and Santaella-Llopis (2010) discuss the implications of variable labor shares for business cycle fluctuations.

In contrast, labor share declines associated with markup increases in fact reduce welfare.

Columns five and six then consider the simultaneous introduction of both a decline in the relative price of investment and an increase in markups. In the CES case, these shocks together can produce the entire 5 percentage point decline in the labor share. The changes in output and welfare in the case with both shocks, however, far more closely resemble the outcomes with only the $\hat{\xi}$ shock than those with only the $\hat{\mu}$ shock. Markups may be of roughly equal importance as the relative price of investment for explaining the total global labor share decline, but this evidence suggests that the component attributable to the markup shock had far less important macroeconomic implications than the component attributable to the decline in the relative price of investment.

6 Conclusion

In this paper we do three things. We document a decline in the global labor share over the past 35 years, offer an explanation for the decline, and assess the resulting macroeconomic implications. We start by showing that the share of income accruing to labor has declined in the large majority of countries and industries. Larger labor share declines occurred in countries or industries with larger declines in their relative price of investment goods. Next, we use this cross-sectional variation to estimate the shape of the production function and conclude that the decline in the relative price of investment explains roughly half of the decline in the global labor share. Finally, we explore the macroeconomic and welfare implications of our results. We emphasize that our explanation for the labor share decline carries with it significantly different implications from alternative explanations.

Our conclusions suggest several paths for future research. For example, our results imply meaningful changes in the distribution of income when households have heterogeneous assets or when skills are differentially substitutable with capital and can accumulate endogenously. Further, as labor shares have declined, business earnings and corporate saving have increased. This large change in the flow of funds between households and firms may have important macroeconomic repercussions.²⁵

²⁵See Karabarbounis and Neiman (2012) for documentation of this global increase in saving by the corporate

Lastly, our results support the view that changes in technology, likely associated with the computer and information technology age, are key factors in understanding long-term changes in factor shares. This raises natural questions. What will be the future path of the relative price of investment? Will the elasticity of substitution between capital and labor change over time? Standard macroeconomic models do not allow for long-term trends in labor shares, a strong prediction which we show to be violated in the data since the early 1980s. We hope our results generate new frameworks and analyses useful for thinking about these future trends.

sector and for a model in which the sectoral distribution of saving matters for macroeconomic allocations.

References

- ACEMOGLU, D. (2003): “Labor- and Capital-Augmenting Technical Change,” *Journal of European Economic Association*, 1(1), 1–37.
- ANTRAS, P. (2004): “Is the U.S. Aggregate Production Function Cobb-Douglas? New Estimates of the Elasticity of Substitution,” *The B.E. Journal of Macroeconomics*, 4(1).
- AZMAT, G., A. MANNING, AND J. VAN REENEN (2012): “Privatization and the Decline of the Labour’s Share: International Evidence from Network Industries,” *Economica*, 79, 470–92.
- BARRO, R., AND X. SALA-I-MARTIN (1995): *Economic Growth*. McGraw Hill.
- BENTOLILA, S., AND G. SAINT-PAUL (2003): “Explaining Movements in the Labor Share,” *The B.E. Journal of Macroeconomics*, 3(1).
- BLANCHARD, O. (1997): “The Medium Run,” *Brookings Papers on Economic Activity*, 2, 89–158.
- BLANCHARD, O., AND F. GIAVAZZI (2003): “Macroeconomic Effects of Regulation And Deregulation In Goods and Labor Markets,” *The Quarterly Journal of Economics*, 118(3), 879–907.
- CHIRINKO, R. (2008): “ σ : The Long and Short Of It,” *Journal of Macroeconomics*, 30, 671–86.
- DUFFY, J., AND C. PAPAGEORGIOU (2000): “A Cross-Country Empirical Investigation of the Aggregate Production Function Specification,” *Journal of Economic Growth*, 5, 87–120.
- FISHER, J. D. (2006): “The Dynamic Effects of Neutral and Investment-Specific Technology Shocks,” *Journal of Political Economy*, 114(3), 413–51.
- GOLLIN, D. (2002): “Getting Income Shares Right,” *Journal of Political Economy*, 110(2), 458–74.
- GREENWOOD, J., Z. HERCOWITZ, AND G. W. HUFFMAN (1988): “Investment, Capacity Utilization, and the Real Business Cycle,” *The American Economic Review*, 78(3), 402–17.
- GREENWOOD, J., Z. HERCOWITZ, AND P. KRUSELL (1997): “Long-Run Implications of Investment-Specific Technological Change,” *The American Economic Review*, 87(3), 342–62.
- HARRISON, A. (2002): “Has Globalization Eroded Labor’s Share? Some Cross-Country Evidence,” Working Paper, University of California at Berkeley.

- HSIEH, C.-T., AND P. KLENOW (2007): “Relative Prices and Relative Prosperity,” *The American Economic Review*, 97(3), 562–85.
- JONES, C. (2003): “Growth, Capital Shares, and a New Perspective on Production Functions,” Working Paper, Stanford.
- (2005): “The Shape of Production Functions and the Direction of Technical Change,” *Quarterly Journal of Economics*, 120(2), 517–49.
- KALDOR, N. (1957): “A Model of Economic Growth,” *The Economic Journal*, 67(268), 591–624.
- KARABARBOUNIS, L., AND B. NEIMAN (2012): “Declining Labor Shares and the Global Rise in Corporate Saving,” NBER Working Paper No. 18154.
- KRUSELL, P., L. E. OHANIAN, J.-V. RIOS-RULL, AND G. L. VIOLANTE (2000): “Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis,” *Econometrica*, 68(5), 1029–53.
- LEON-LEDESMA, M., P. MCADAM, AND A. WILLMAN (2010): “Identifying the Elasticity of Substitution with Biased Technical Change,” *The American Economic Review*, 100, 1330.
- LEQUILLER, F. (2002): “Treatment of Stock Options in National Accounts of Non-European OECD Member Countries,” Working Paper, OECD.
- LEQUILLER, F., AND D. BLADES (2006): *Understanding National Accounts*. OECD.
- RESTUCCIA, D., AND C. URRUTIA (2001): “Relative Prices and Investment Rates,” *Journal of Monetary Economics*, 47, 93–121.
- RIOS-RULL, J.-V., AND R. SANTAEULALIA-LLOPIS (2010): “Redistributive Shocks and Productivity Shocks,” *Journal of Monetary Economics*, 57, 931–48.
- RODRIGUEZ, F., AND A. JAYADEV (2010): “The Declining Labor Share of Income,” Human Development Reports Research Paper 2010/36.
- ROTEMBERG, J., AND M. WOODFORD (1995): “Dynamic General Equilibrium Models With Imperfectly Competitive Product Markets,” in *Frontiers of Business Cycle Research*. Princeton University Press.




Figure 1: Declining Global Labor Share

Notes: The figure shows year fixed effects from a regression of corporate and overall labor shares that also include country fixed effects to account for entry and exit during the sample. The regressions are weighted by corporate gross value added and GDP measured in U.S. dollars at market exchange rates. We normalize the fixed effects to equal the level of the global labor share in our dataset in 1975.







Figure 2: Declining Labor Share for the Largest Countries

Notes: The figure shows the labor share and its linear trend for the four largest economies in the world from 1975.




Figure 3: Estimated Trends in Country Labor Shares

Notes: The figure shows estimated trends in the labor share for all countries in our dataset with at least 15 years of data starting in 1975. Trend coefficients are reported in units per 10 years (i.e. a value of -5 means a 5 percentage point decline every 10 years). The largest 8 economies are shaded.




Figure 4: Estimated Trends in U.S. State Labor Shares

Notes: The figure shows estimated trends in the labor share for 51 U.S. states plus District of Columbia in BEA data starting in 1975. Trend coefficients are reported in units per 10 years (i.e. a value of -5 means a 5 percentage point decline every 10 years).




Figure 5: Estimated Trends in Industry Labor Shares

Notes: The figure shows estimated trends in the labor share for 10 non-overlapping industries in the KLEMS data starting in 1975. Trend coefficients are reported in units per 10 years (i.e. a value of -5 means a 5 percentage point decline every 10 years).




Figure 6: Within Component and Total Trends in Country Labor Shares

Notes: The figure plots the trend in the labor share against the within-industry component as defined in equation (2) using the KLEMS data.




Figure 7: Declining Global Price of Investment Goods

Notes: The figure shows year fixed effects from regressions of the log relative price of investment that absorb country fixed effects to account for entry and exit during the sample. The regressions are weighted by GDP measured in U.S. dollars at market exchange rates.




Figure 9: Capital Share and Labor Share

Notes: The figure plots the trend in the log capital share against the trend in the log labor share. All values are scaled to denote percent changes per 10 years. For example, a value of -10 for the trend in the log labor share means a 10 percent decline of the labor share every 10 years. For illustrative reasons, in this figure we drop three observations (Kazakhstan, Kyrgyzstan, and Macao) with extremely low weights in the regression of the first row of Table 2 and we winsorize one observation in each dimension for both variables. The solid line represents the fitted relationship between trends in capital share and trends in the labor share (slope 0.22 with a standard error of 0.23), whereas the dashed line represents the 45 degree line.

	Labor Share	Investment Price	$\hat{\sigma}$	Std. Error	90% Conf. Interval	Obs.
(i)	KN Merged	PWT	1.26	0.08	[1.12,1.40]	55
(ii)	KN Merged	EIU	1.21	0.07	[1.10,1.32]	53
(iii)	OECD and UN	PWT	1.25	0.08	[1.12,1.39]	48
(iv)	OECD and UN	EIU	1.16	0.06	[1.06,1.26]	46
(v)	KLEMS 1	KLEMS	1.17	0.06	[1.06,1.27]	129
(vi)	KLEMS 2	KLEMS	1.49	0.13	[1.29,1.70]	129
	Average		1.26			

Table 1: Baseline Estimates of Elasticity of Substitution

	Labor Share	Investment Price	Investment Rate	$\hat{\sigma}$	Std. Error	90% Conf. Interval	Obs.
(i)	KN Merged	PWT	Total	1.17	0.10	[1.00,1.35]	52
(ii)	KN Merged	PWT	Corporate	1.02	0.09	[0.86,1.17]	53
(iii)	KN Merged	EIU	Total	1.22	0.09	[1.07,1.35]	51
(iv)	KN Merged	EIU	Corporate	1.28	0.09	[1.13,1.43]	52
(v)	OECD and UN	PWT	Total	1.24	0.10	[1.07,1.41]	44
(vi)	OECD and UN	PWT	Corporate	1.21	0.11	[1.02,1.39]	44
(vii)	OECD and UN	EIU	Total	1.21	0.09	[1.07,1.35]	43
(viii)	OECD and UN	EIU	Corporate	1.32	0.10	[1.16,1.49]	43
Average				1.21			

Table 2: Estimates of Elasticity of Substitution Allowing for Markups

	Labor Share	Nested Input with Capital	$\hat{\sigma}$	Std. Error	90% Conf. Interval	Obs.
(i)	KLEMS 1	High Skill	1.23	0.08	[1.11,1.36]	100
(ii)	KLEMS 1	Middle and Low Skill	1.19	0.08	[1.05,1.33]	100
(iii)	KLEMS 1	Low Skill	1.19	0.09	[1.04,1.34]	100
(iv)	KLEMS 2	High Skill	1.34	0.16	[1.07,1.60]	100
(v)	KLEMS 2	Middle and Low Skill	1.31	0.17	[1.03,1.60]	100
(vi)	KLEMS 2	Low Skill	1.31	0.18	[1.02,1.61]	100
	Average		1.26			

Table 3: Estimates of Elasticity of Substitution with Different Production Functions

		CD	CES	CD	CES	CD	CES
		$\hat{\xi}$	$\hat{\xi}$	$\hat{\mu}$	$\hat{\mu}$	$(\hat{\xi}, \hat{\mu})$	$(\hat{\xi}, \hat{\mu})$
(i)	Labor Share (Percentage Points)	0.0	-2.6	-3.1	-2.6	-3.1	-4.9
(ii)	Capital Share (Percentage Points)	0.0	2.6	-1.9	-2.4	-1.9	-0.1
(iii)	Profit Share (Percentage Points)	0.0	0.0	5.0	5.0	5.0	5.0
(iv)	Consumption	18.1	20.1	-5.2	-5.4	10.7	12.4
(v)	Nominal Investment	18.1	30.8	-11.1	-12.7	3.7	11.9
(vi)	Labor Input	0.0	-1.4	-3.2	-2.9	-3.2	-4.2
(vii)	Capital Input	51.6	67.8	-11.1	-12.7	33.2	43.6
(viii)	Output	18.1	22.8	-6.3	-6.8	9.4	12.3
(ix)	Wage	18.1	19.2	-8.2	-8.2	7.1	7.7
(x)	Rental Rate	-22.1	-22.1	0.0	0.0	-22.1	-22.1
(xi)	Capital-to-Output	28.4	36.6	-5.2	-6.4	21.8	27.9
(xii)	Welfare Equivalent Consumption	18.1	22.1	-3.0	-3.4	13.2	15.8

Table 4: Evaluating Labor Share's Decline (Percent Changes Across Steady States)