

THE IMPACT OF 401(K) PARTICIPATION ON THE WEALTH DISTRIBUTION

VICTOR CHERNOZHUKOV AND CHRISTIAN HANSEN

ABSTRACT. In this paper, we use the instrumental quantile regression model developed in Chernozhukov and Hansen (2001) to examine the effects of 401(k) plans on wealth using data from the Survey of Income and Program Participation. Using 401(k) eligibility as an instrument for 401(k) participation, we estimate the quantile treatment effects of participation in a 401(k) plan on several measures of wealth. The results show the effects of 401(k) participation on net financial assets are positive and significant over the entire range of the asset distribution, and that the increase in the low tail of the assets distribution appears to translate completely into an increase in wealth. However, there is significant evidence of substitution from other forms of wealth in the upper tail of the distribution. We also provide estimates obtained using the estimator of Abadie, Angrist, and Imbens (2002).

JEL Classification: H31, C21, C31

1. INTRODUCTION

In the early 1980s, the United States introduced several tax deferred savings options in an effort to increase individual saving for retirement. The two options which have generated the most interest are Individual Retirement Accounts (IRAs) and 401(k) plans. Tax deferred IRAs and 401(k) plans are similar in that both allow the individual to deduct contributions from taxable income and allow tax-free accrual of interest on assets held within the plan. The key differences between the two savings options are that employers provide 401(k) plans, and employers may also match a certain percentage of an employee's contribution. Since 401(k) plans are provided by employers, only workers in firms offering plans are eligible for participation, while participation in IRAs is open to everyone.¹

While it is clear that 401(k) plans and, to a lesser extent, IRAs are widely used as vehicles for retirement saving, their impact on assets is less clear. The key problem in determining the effect of participation in IRA and 401(k) plans on accumulated assets is saver heterogeneity coupled with non-random selection into participation states. In particular, it is generally recognized that some people have a higher preference for saving than others. Thus, it seems likely that those individuals with the highest unobserved preference for saving would be most likely to choose to participate in tax-advantaged retirement savings plans and would also have higher savings in other assets than individuals with lower unobserved saving propensity. This implies that conventional estimates that do not account for saver heterogeneity and selection of the participation state will be biased upward, tending to overstate the actual savings effects of 401(k) and IRA participation.

This problem has long been recognized in the savings literature and has led to numerous important studies which attempt to overcome this problem. In a series of articles, Poterba, Venti, and Wise (1994, 1995, 1996) use comparisons between groups based on eligibility for 401(k) participation. They argue that 401(k) eligibility can be taken as exogenous given income. The argument is motivated by the fact that eligibility is determined by the employer, and so may be taken as exogenous conditional on covariates. Poterba, Venti, and Wise (1996) contains an overview of results used to substantiate this claim and mean and median regression estimates of the impact of 401(k) eligibility on household net financial assets. The results show that 401(k) eligibility has significant and positive effects on net financial assets. Based on the assumed exogeneity of 401(k) eligibility, they attribute this difference to the causal effect of 401(k) eligibility on savings.² Recent work by Benjamin (2003) examines the effects of eligibility on savings using matching based on the propensity score and finds positive, although much more modest, effects of 401(k) eligibility on assets.³

¹A detailed description of regulations regarding retirement programs can be found in a recent publication of the Employee Benefit Research Institute (1997).

²For a differing viewpoint, see Engen, Gale, and Scholz (1996), which contends that eligibility should not be treated as exogenous.

³Benjamin uses a more inclusive definition of assets and makes adjustments to account for replacement or substitution of an existing defined contribution or defined benefit plan by a 401(k).

A similar approach, which we follow in this paper, is that of Abadie (2001). Abadie, assuming that eligibility for a 401(k) is exogenous given income (and other covariates), uses 401(k) eligibility as an instrument for 401(k) participation in order to estimate the effect of 401(k) participation, not eligibility, on net financial assets. Abadie uses a novel semiparametric estimator which estimates the average effect for compliers.⁴ Since only individuals eligible for a 401(k) can participate, the average effect for compliers also corresponds to the average effect for the treated. Abadie's results suggest that the average effect for the treated of 401(k) participation is significant and positive.

One drawback of all of these studies is that they focus the analysis on measures of central tendency: the mean or the median. While the mean and median impacts are interesting and important measures in determining a program's impact, they are not sufficient to fully characterize the impact of the treatment except under very restrictive conditions. In particular, they are uninformative about the impact of the treatment on other, perhaps more interesting, points in the outcome distribution when the treatment effect is heterogeneous. Understanding the distributional impact of 401(k) plans is especially interesting from a policy perspective since policy makers may be particularly concerned about the impact of 401(k) plans on the lower part of the wealth distribution. In addition, knowledge of the distributional impact of a program provides a clearer picture of what is driving the mean results.

As with estimates of the mean effect, the analysis of the distributional effect is complicated by the possibility that individuals choose whether or not to participate in a 401(k) based on their unobserved preferences for saving. One estimator which would allow a more full characterization of the effect of a heterogeneous treatment given treatment exogeneity is the quantile regression estimator of Koenker and Bassett (1978). However, the self-selection of the participation state makes the conventional quantile regression estimator inappropriate.⁵

In this paper, we contribute to the extensive set of existing literature of the impact of 401(k) plans on wealth by analyzing the impact of 401(k) participation on the entire wealth distribution. Using the reasoning of Poterba, Venti, and Wise (1994, 1995, 1996) and Abadie (2001) outlined above, we use 401(k) eligibility as an instrument for 401(k) participation in order to estimate the effect of participating in a 401(k) on various measures of wealth. To do this, we employ a model and an estimator developed in Chernozhukov and Hansen (2001). The model provides a set of assumptions under which the conditional quantiles of the outcome distribution may be recovered from a set of statistical moment equations through

⁴In the context of 401(k) participation, the group of compliers is the group of individuals who would participate in a 401(k) if eligible but would not if ineligible. Non-compliers in this example are people who would not participate in the 401(k) regardless of their eligibility status.

⁵Also, treatment heterogeneity renders the two stage least absolute deviation estimator of Amemiya (1982) and its extension to quantile regression by Chen and Portnoy (1996) inconsistent. The inconsistency was first demonstrated in Chernozhukov and Hansen (2001). However, it should be noted that two stage quantile regression methods are useful for estimation of constant effect models.

the use of instrumental variables. The estimator we use is computationally convenient for linear quantile models and can be computed through a series of conventional linear quantile regressions. Chernozhukov and Hansen (2001) demonstrate that the estimator is consistent under endogeneity and treatment effect heterogeneity. Thus, this paper provides an important complement to the work discussed above which focuses on estimating the impact of 401(k) plans on the center of the outcome distribution. Also, due to the binary nature of both the participation decision and the eligibility instrument, the estimator developed by Abadie, Angrist, and Imbens (2002) to estimate quantile effects for binary treatments under endogeneity also applies. We present estimates obtained through both procedure to provide both a robustness check and a comparison of the two estimators. We find that the results are very similar using either estimation procedure.

The instrumental quantile regression estimates indicate that there is considerable heterogeneity in the effect of 401(k) participation on net financial assets, with the treatment effect increasing monotonically as one moves from the lower to the upper tail of the asset distribution. The results are also uniformly positive and significant, suggesting that 401(k) participation positively impacts net financial assets across the entire distribution. The effect of participation on total wealth is positive and approximately constant for all quantiles. In addition, it is of the same magnitude as the effect of participation on net financial assets for low quantiles, but is substantially smaller than the effect of participation on the upper quantiles of net financial assets. These results suggest that the increase in net financial assets observed in the lower tail of the conditional assets distribution can be interpreted as an increase in wealth, while the increase in the upper tail of the distribution is mitigated by substitution with some other component of wealth. The effect of participation on net non-401(k) financial assets⁶ is uniformly insignificant, which suggests there is little substitution for 401(k) assets along this dimension of wealth.

The remainder of the paper is organized as follows. Section 2 reviews the model of quantile treatment effects of Chernozhukov and Hansen (2001) and demonstrates how an empirical model for assets may be embedded in the model. In Section 3, the data used in the empirical analysis are described. Section 4 presents the empirical results and compares the results from the estimator of Chernozhukov and Hansen (2001) to those obtained with the estimator of Abadie, Angrist, and Imbens (2002), and Section 5 concludes.

2. AN INSTRUMENTAL VARIABLE MODEL FOR QUANTILE TREATMENT EFFECTS

In the following, we briefly present the assumptions and main implication of the instrumental variables model of quantile treatment effects developed in Chernozhukov and Hansen (2001). We then show how an empirical model of savings decisions may be embedded in this framework. This discussion helps illustrate the interpretation of the estimates of the model, especially the interpretation of the quantile index τ , and isolates the key identifying assumptions.

⁶Net non-401(k) financial assets are net financial assets minus 401(k) balances. More details about the wealth measures are found in the description of the data in Section 3.

2.1. Potential Outcomes and the QTE. The model is developed within the conventional potential (latent) outcome framework. Potential real-valued outcomes are indexed against treatment d and denoted Y_d . For example, Y_d is an individual's outcome when $D = d$. Treatments d take values in a subset \mathcal{D} of \mathbb{R}^l . The potential outcomes $\{Y_d\}$ are latent because, given the selected treatment D , the observed outcome for each individual or state of the world is $Y \equiv Y_D$. That is, only one component of potential outcomes vector $\{Y_d\}$ is observed for each observational unit.

While there are many features of the distributions of potential outcomes that may be interesting, we focus on the quantiles of potential outcomes conditional on covariates X ,⁷

$$\left\{ Q_{Y_d}(\tau|x), \tau \in (0, 1) \right\},$$

and the quantile treatment effects (QTE) that summarize the difference between the quantiles under different treatments (e.g. Doksum (1974)):

$$Q_{Y_d}(\tau|x) - Q_{Y_{d'}}(\tau|x) \quad \text{or, if defined,} \quad \frac{\partial}{\partial d} Q_{Y_d}(\tau|x).$$

Quantile treatment effects represent a useful way of describing the effect of treatment d on different points of the marginal distribution of potential outcomes.

Typically D is selected in relation to $\{Y_d\}$ inducing endogeneity, so that the conditional quantile of Y given the selected treatment $D = d$, denoted $Q_Y(\tau|d, x)$, is generally not equal to the quantile of potential or latent outcome $Q_{Y_d}(\tau|x)$. This makes the conventional quantile regression inappropriate for the estimation of $Q_{Y_d}(\tau|x)$. The model of Chernozhukov and Hansen (2001), briefly presented below, states the conditions under which we can recover the quantiles of latent outcomes through a set of conditional moment restrictions.

2.2. The Instrumental Quantile Treatment Model. We build the model from the basic Skorohod representation⁸ of latent outcomes Y_d , which yields for each d given $X = x$

$$Y_d = q(d, x, U_d), \quad \text{where } U_d \stackrel{d}{\sim} U(0, 1), \quad (2.1)$$

and $q(d, x, \tau) = Q_{Y_d}(\tau|x)$ is the conditional τ -quantile of latent outcome Y_d . This representation is essential to the rest of the analysis.

The variable U_d is responsible for heterogeneity of outcomes for individuals with the same observed characteristics x and treatment d . It also determines their relative ranking in terms of potential outcomes. Hence we will call U_d the rank variable, and may think of it as representing some innate ability, level of preference, or "proneness".⁹ This allows interpretation of the quantile treatment effect as the treatment effect for people with a given

⁷We use $Q_Y(\tau|x)$ and $f_Y(y|x)$ to denote the conditional τ -quantile and density of Y given $X = x$. Capitals such as Y denote random variables, and lower case letters such as y denote the values they take.

⁸The basic Skorohod representation states that, given a collection of variables $\{\zeta_j\}$, each variable ζ_j can be represented as, a.s.

$$\zeta_j = Q_{\zeta_j}(U_j), \quad \text{for some } U_j \stackrel{d}{\sim} U(0, 1).$$

Recall that $Q_{\zeta_j}(\tau)$ denotes the τ -quantile of variable ζ_j .

⁹Doksum (1974) uses the term proneness as in "prone to learn fast" or "prone to grow taller".

innate ability, making quantile analysis an interesting tool for describing and learning the structure of heterogeneous treatment effects.

The model consists of five main conditions (some are representations) that hold jointly.

The IVQT Model: Given a common probability space (Ω, F, P) , for P -almost every value of X, Z where X represents covariates and Z represents excluded instruments, the following conditions A1-A5 hold *jointly*:

A1 POTENTIAL OUTCOMES. Given $X = x$, for each d , for some $U_d \stackrel{d}{\sim} U(0, 1)$,

$$Y_d = q(d, x, U_d),$$

where $q(d, x, \tau)$ is strictly increasing and left-continuous in τ .

A2 INDEPENDENCE. Given $X = x$, $\{U_d\}$ is independent of Z .

A3 SELECTION. Given $X = x, Z = z$, for unknown function δ and random vector V ,

$$D \equiv \delta(z, x, V).$$

A4 RANK SIMILARITY. For each d and d' , given (V, X, Z)

$$U_d \text{ is equal in distribution to } U_{d'}.$$

A5 OBSERVED variables consist of (for $U_D \equiv \sum_{d \in \mathcal{D}} I(D = d) \cdot U_d$)

$$\begin{cases} Y \equiv q(D, X, U_D), \\ D \equiv \delta(Z, X, V), \\ X, Z. \end{cases}$$

Chernozhukov and Hansen (2001) demonstrate that the following result is an implication of the IVQT model.

Theorem 1 (*Main Statistical Implication*). *Suppose conditions A1-A5 hold. Then, for any $\tau \in (0, 1)$, a.s.*

$$P[Y \leq q(D, X, \tau) | X, Z] = \tau \text{ and } P[Y < q(D, X, \tau) | X, Z] = \tau. \quad (2.2)$$

This result provides an important link of the parameters of the IVQT model to a set of conditional moment equations which are used in Chernozhukov and Hansen (2001) to develop identification conditions for the IVQT model as well as for estimation and inference. In addition, Chernozhukov and Hansen (2001) give an extensive discussion of the IVQT model, its assumptions, and its identification. While this discussion will not be repeated here, it is important to note that the assumptions of the IVQT model differ from those in other models with endogeneity and heterogeneous treatment effects¹⁰ in two key respects. First, the IVQT model imposes a different set of independence conditions; in particular,

¹⁰See, for example, Amemiya (1982), Heckman and Robb (1986), Imbens and Angrist (1994), and Vytlacil (2000).

it does not require that the instruments, Z , are independent of the errors in the selection equation V . The independence of Z and V may be violated when Z is measured with error or related to V in other ways. Second, the IVQT model imposes rank similarity, Assumption A4, which will be discussed in the context of saving decisions below.

2.3. The Instrumental Quantile Regression Model and Saving Decisions. Assumptions A1-A5 represent a plausible framework within which to analyze the effects of participating in a 401(k) plan on an individual's accumulated wealth. First, wealth, Y_d , in the participation state $d \in \{0, 1\}$ can be represented as

$$Y_d = q(d, X, U_d), \quad U_d \sim U(0, 1)$$

by the Skorohod representation of random variables, where $\tau \mapsto q(d, X, \tau)$ is the conditional quantile function of Y_d and U_d is an unobserved random variable. Following the discussion in Section 2.2, we will refer to U_d as the preference for saving and thus interpret the quantile index τ as indexing rank in the preference for saving distribution.¹¹ The individual selects the 401(k) participation state to maximize expected utility:

$$D = \arg \max_{d \in \mathcal{D}} E \left[W\{Y_d, d\} \middle| X, Z, V \right] = \arg \max_{d \in \mathcal{D}} E \left[W\{q(d, x, U_d), d\} \middle| X, Z, V \right], \quad (2.3)$$

where $W\{y, d\}$ is the unobserved Bernoulli utility function. As a result, the selection is represented by

$$D = \delta(Z, X, V)$$

where Z and X are observed, V is an unobserved information component that depends on rank U_d , and includes other unobserved variables that affect the selection, and function δ is unknown. Thus this model is a special case of the IVQT model. In this model, the independence condition A2 only requires that U_d is independent of Z , conditional on X . The error V in the selection equation need not be independent of Z .

The simplest form of rank similarity is rank invariance, under which the preference for saving vector U_d may be collapsed to a single random variable:

$$U = U_0 = U_1.$$

In this case, a single preference for saving is responsible for an individual's ranking across all treatment states. It is important to note that U is defined relative to observationally identical people (individuals with the same X and Z). Rank invariance has been used in many interesting models without endogeneity,¹² which supports its plausibility. In addition, traditional simultaneous equations models are built assuming rank invariance. However, as noted in Heckman and Smith (1997), rank invariance may be implausible on logical grounds

¹¹Since the outcomes of interest are all measures of accumulated wealth, perhaps more appropriate, but more cumbersome, terminology would be preference for accumulated assets. In addition, if there are unobservable factors besides preferences then this interpretation of U_d and τ is incorrect, and τ should be only interpreted as indexing rank in the conditional distribution of Y_d given x . For simplicity and clarity, we will refer to U_d and τ as relating to preference for saving throughout the rest of the paper.

¹²For example, Doksum (1974) and Heckman and Smith (1997).

since it implies that the potential outcomes Y_d are not truly multivariate, but have a jointly degenerate distribution.

The similarity condition A4 is a more general form of rank invariance – it relaxes the exact invariance of ranks U_d across d by allowing noisy, unsystematic variations of U_d across d , conditional on (V, X, Z) . This relaxation allows for variation in the ranks across the treatment states, requiring only a “rank invariance in expectation”. Therefore, similarity accommodates general multivariate models of outcomes. It states that given the information in (V, X, Z) employed to make the selection of treatment D , the expectations of any function of rank U_d does not vary across the treatment states. That is, ex-ante, conditional on (V, X, Z) , the ranks may be considered to be the same across potential treatments, but the realized, ex-post, rank may be different across treatment states.

From an econometric perspective, the similarity assumption is nothing but a restriction on the evolution on the unobserved heterogeneity component which precludes systematic variation of U_d across the treatment states. Similarity allows interpretation of the quantile treatment effect as the treatment effect holding the level of unobserved heterogeneity constant across the treatment states:

$$q(d, x, \tau) - q(d', x, \tau) = q(d, x, U_d) - q(d', x, U_{d'}) \Big|_{U_d=U_{d'}=\tau}.$$

Since changes in U_d across d are assumed to be asystematic, the quantile treatment effect not only summarizes the distributional impact but also the actual likely treatment effect.

To be more concrete, consider the following simple example where

$$U_d = F_{V+\eta_d}(V + \eta_d),$$

where $F(\cdot)_{V+\eta_d}$ is the distribution function of $V + \eta_d$ and $\{\eta_d\}$ are mutually iid conditional on V , X , and Z . The variable V represents an individual’s “mean” saving preference, while η_d is a noisy adjustment.¹³ This more general assumption leaves the individual optimization problem (2.3) unaffected, while allowing variation in an individual’s rank across different potential outcomes.

While we feel that similarity is quite a reasonable assumption, imposing similarity is not innocuous. In the context of 401(k) participation, matching practices of employers could jeopardize the validity of the similarity assumption. This is because individuals in firms with high match rates may be expected to have a higher rank in the asset distribution than workers in firms with less generous match rates. This suggests that the distribution of U_d may be different across the treatment states. Indeed, while we may think that the expected rank should remain the same with matching, it seems less likely that the expectation of any function of the rank would be the same across the participation and non-participation states.

Similarity may still hold in the presence of the employer match if the rank, U_d , in the asset distribution is insensitive to the match rate. The rank may be insensitive if, for example, individuals follow simple rules of thumb such as target saving when they make their

¹³Clearly similarity holds in this case – $U_d \stackrel{d}{=} U_{d'}$ given V , X , and Z .

savings decisions. Also, if the variation of match rates is small relative to the variation of individual heterogeneity or if the covariates capture most of the variation in match rates, then similarity will be satisfied approximately. Since the model is just-identified in our data, specification tests based on the implications of Theorem 1 may not be used to perform over-identifying tests. However, the quantile treatment effects model and estimator of Abadie, Angrist, and Imbens (2002), which apply only to binary treatment variables, provide a useful robustness check. While the approach of Abadie, Angrist, and Imbens (2002) and the approach presented in this paper generally identify and estimate different quantities, they will estimate the same thing when the assumptions of both models, including similarity, are satisfied and the set of compliers is representative of the population. If these conditions are not met, then the two estimators will in general have different probability limits, suggesting that a comparison of results based on the two models will provide evidence on the plausibility of these assumptions. Comparisons between the two estimators are presented with the empirical results in Section 4. The results from the two estimators are very similar, suggesting that employer matching does not result in a serious violation of rank similarity.

3. THE DATA

To estimate the QTE, we use data on a sample of households from Wave 4 of the 1990 Survey of Income and Program Participation (SIPP).¹⁴ The sample is limited to households in which the reference person is 25-64 years old, in which at least one person is employed, and in which no one is self-employed.¹⁵ The sample consists of 9,915 households and all dollar amounts are in 1991 dollars.

The 1991 SIPP reports household financial data across a range of asset categories. These data include a variable for whether a person works for a firm that offers a 401(k) plan. Households in which a member works for such a firm are classified as eligible for a 401(k). In addition, the survey also records the amount of 401(k) assets. Households with a positive 401(k) balance are classified as participants, while eligible households with a zero balance are considered non-participants.

While there are several potential measures of wealth in the 1991 SIPP, we choose to focus our analysis on total wealth, net financial assets, and net non-401(k) financial assets. Net non-401(k) assets are defined as the sum of checking accounts, U.S. saving bonds, other interest-earning accounts in banks and other financial institutions, other interest-earning assets (such as bonds held personally), stocks and mutual funds less nonmortgage debt, and

¹⁴This sample has been extensively to study the effect of 401(k) plans on wealth. See, for example, Benjamin (2003), Abadie (2001), Engen and Gale (2000), Engen, Gale, and Scholz (1996), and Poterba, Venti, and Wise (1994, 1995, 1996). The sample is often referred to as the 1991 SIPP because data were collected between February and May of 1991.

¹⁵Analyses are restricted to this sample because the SIPP only asks 401(k) questions to people 25 and older, because retirement and saving behavior of people over 65 would complicate the analysis, and because the self-employed and unemployed do not have access to 401(k)s. The household reference person is the person in whose name the family's home is owned or rented.

IRA balances. Net financial assets are net non-401(k) financial assets plus 401(k) balances, and total wealth is net financial assets plus housing equity¹⁶ and the value of business, property, and motor vehicles.

We use the same set of covariates as Benjamin (2003). Specifically, we use age, income, family size, education, marital status, two-earner status, defined benefit (DB) pension status, IRA participation status, and home ownership status. Marital status, two-earner status, DB pension status, IRA participation status, and home ownership status are binary variables, where two-earner status indicates whether both household heads, where present, contribute to household income and DB pension status indicates whether the household's employer offers a DB pension plan. The education variable measures the number of years of school completed by the household reference person, and for the analysis we have categorized this variable into four groups: less than 12 years of education, 12 years of education, 13-15 years of education, and 16 or more years of education. Households are classified as IRA participants if they have positive IRA asset balances, and households are classified as home owners if the household has a positive home value. In addition, in the estimates reported below, we control for age using categorical variables: less than 30 years old, 30-35 years old, 36-44 years old, 45-54 years old, and 55 years old or older. Following Poterba, Venti, and Wise (1995), we control for income through the use of seven categorical variables. The income intervals are as follows: <\$10K, \$10-20K, \$20-30K, \$30-40K, \$40-50K, \$50-75K, and \$75K+.

Table 1 contains descriptive statistics for the full sample as well as by eligibility and participation status. 37% of the sample is eligible for a 401(k) plan while 26% choose to participate. Among those eligible for a 401(k) account, the participation rate is 70%. The descriptive statistics indicate that participants have larger holdings of all measures of wealth that we consider. As expected, the means of all of the wealth variables are substantially larger than their medians, indicating the high degree of skewness in wealth. The means also show that 401(k) participants have more income, are more likely to be married, are more likely to have IRAs and defined benefit pensions, are more likely to be home-owners, and are more educated than non-participants. Average age and family size are similar between the two groups. Descriptive statistics for the dependent variables by income category are also provided in Table 2.

4. EMPIRICAL RESULTS

4.1. Estimation and Inference Procedures. To capture the effects of 401(k) participation on net financial assets, we estimate linear quantile models¹⁷ of the form

$$Q_{Y_d|X}(\tau) = d\alpha(\tau) + X'\beta(\tau)$$

where d indicates 401(k) participation status and is instrumented for by 401(k) eligibility, following Abadie (2001) and Poterba, Venti, and Wise (1994, 1995, 1996). The potential

¹⁶Housing equity is defined as housing value less mortgage.

¹⁷The OLS and 2SLS estimates are based on analogous specifications.

outcomes Y_d are the three previously mentioned measures of wealth (total wealth, net financial assets, and net non-401(k) financial assets), and X consists of dummies for income category, dummies for age category, dummies for education category, a marital status indicator, family size, two-earner status, DB pension status, IRA participation status, home ownership status, and a constant.¹⁸ To more fully control for income, we also consider estimates obtained within each income category. In these cases, the income category dummies are omitted and a linear term in income is included to account for any remaining variation within income category.

The main results reported below are for the standard quantile regression (QR) estimator and the instrumental quantile regression (IQR) estimator of Chernozhukov and Hansen (2001) which corrects for the endogeneity of 401(k) participation under the assumptions of the model presented in Section 2 of this paper. The IQR estimator may be viewed as a convenient method of approximately solving the sample analog of moment equations (2.2):¹⁹

$$\frac{1}{\sqrt{n}} \sum_{i=1}^n (1(Y_i \leq D_i' \hat{\alpha} + X_i' \hat{\beta}) - \tau)(X_i', Z_i')' = o_p(1). \quad (4.1)$$

In particular, when the model is just identified and for a given quantile, the IQR estimator may be computed by running a series of standard quantile regressions of $Y - d\alpha_j$ on covariates X and instrument Z where $\{\alpha_j\}$ is a grid over α . The estimate of α may then be taken as the α_j that minimizes the absolute value of the coefficient on Z . Estimates of β are the corresponding coefficients on X . Chernozhukov and Hansen (2001) contains further details of the estimation as well as asymptotic theory. We also estimate all models using the estimator of Abadie, Angrist, and Imbens (2002).

Estimates of the QTE, $\alpha(\tau)$ ²⁰, for many different points τ also provide an estimate of the QTE process $\alpha(\cdot)$ which treats α as a function of τ . Knowledge of the QTE process allows formal testing of a number of interesting hypotheses. These include the constant effect hypothesis ($\alpha(\cdot) = \alpha$), of which the hypothesis of no effect ($\alpha(\cdot) = 0$) is a special case, and the hypothesis of no endogeneity ($\alpha(\cdot) = \alpha_{QR}(\cdot)$ where α_{QR} denotes the ordinary quantile regression estimate). If the constant effect hypothesis is not rejected, the distributional impact of the treatment may be captured by a single statistic, such as the mean or the median treatment effect. Also, failure to reject the hypothesis of no endogeneity suggests that the endogeneity bias is not statistically important and that standard QR estimates may be used. Chernozhukov and Hansen (2002) provides asymptotic theory for the IQR

¹⁸We also considered alternate specifications of the covariate vector. However, the estimate of the treatment effect was found to be largely insensitive to the specification. The most substantial difference is that when the home ownership dummy was excluded the results for total wealth closely tracked those of net financial assets across the entire distribution, indicating little or no substitution between 401(k) assets and other forms of wealth. All other results were very similar.

¹⁹Estimation using this set of moment equations was considered by Abadie (1995) who noted the computational difficulty in obtaining their solution.

²⁰The following discussion also applies to the coefficients of the covariates, $\beta(\tau)$.

process and suggests a computationally attractive method for performing inference on the IQR process.²¹

4.2. OLS and 2SLS Results. Table 3 provides OLS and 2SLS results of the participation effect. These estimates serve as a benchmark for the quantile and instrumental quantile regression estimates presented later. In addition, they estimate the impact of 401(k) participation on the center of the distribution and so are interesting in their own right. Indeed, in the case of a constant treatment effect, these estimates would be sufficient to fully characterize the distributional impact of the treatment.²²

The first stage estimates, reported in the third column of Table 2, confirm that eligibility for a 401(k) is highly correlated with participation. In the full sample and within each income category, the first stage estimate is large, positive, and highly significant. Indeed, conditional on eligibility, the rest of the covariates have very small effects on 401(k) participation.

In the full sample, the 2SLS estimates are uniformly smaller than the OLS estimates, confirming the intuition that the OLS estimates should be upward biased. However, the biases appear to be modest, especially when compared to the standard errors of the estimates. After accounting for endogeneity, the mean impact of 401(k) participation on both total wealth and net financial assets remains large and significant. Relative to the means, 401(k) participation increases net financial assets by approximately 70% and total wealth by approximately 14%. The magnitude of both effects is also quite similar, though slightly larger for net financial assets, suggesting little substitution between 401(k) assets and other forms of wealth. On the other hand, 401(k) participation has relatively little impact on net non-401(k) financial assets. Neither the OLS nor the 2SLS estimate of the effect of participation on net non-401(k) financial assets is significantly different from zero, and both are quite small in magnitude. Overall, these results suggest that the majority of the increase in net financial assets may be attributed to new saving due to 401(k) plans and not to substitution.

The results by income category provide additional evidence on substitution patterns. The loss of precision resulting from estimating the treatment effect within income categories makes drawing any firm conclusions difficult, but the patterns of the estimates are still quite interesting.²³ The impact of 401(k) participation on net financial assets is uniformly positive and significant and tends to increase as one moves from lower to higher income categories. This result appears to be consistent with the resource constraints of the different income groups. The results for net non-401(k) financial assets are never significantly different from

²¹The lack of formal theory for the empirical process of the estimates obtained through Abadie, Angrist, and Imbens (2002) is one reason that we focus on the IQR approach in this paper.

²²The process tests reported below suggest that this is the case when the dependent variable is total wealth or net non-401(k) financial assets.

²³In the following, we ignore estimates in the lowest income category which are greatly influenced by outliers in the upper tail of the distribution and the small sample size. The influence of the upper tail is seen clearly in the quantile regression results presented below.

zero. However, in all cases but one, the point estimate is negative and non-negligible, which provides weak evidence that there is financial asset substitution that was obscured in the results obtained in the full sample. While the results for total wealth show much less of a pattern as one looks across income categories, it can be seen that in no case is the effect significantly different from zero. The point estimates are uniformly positive and, in the majority of cases, are reasonably large. This again provides weak evidence that 401(k) participation increases total wealth by a modest amount, but that this increase is smaller than the increase to net financial assets, indicating substitution between 401(k) and other assets.

4.3. Quantile Regression and Instrumental Quantile Regression Results: Full Sample. While the OLS and 2SLS results presented above serve to characterize the impact of treatment on the mean of the distribution, they fail to capture the distributional impact of 401(k) participation on wealth. To further explore the effect of 401(k) participation on wealth, we report results obtained from both standard quantile regression and the instrumental quantile regression of Chernozhukov and Hansen (2001) in Figure 1.

The left column of Figure 1 contains QR estimates of the effect of 401(k) participation on the wealth measures, and the right column of Figure 1 presents the IQR estimates of the QTE. The shaded region in the first six panels represents the 95% confidence interval.²⁴ The last two panels plot the estimated effects for each of the dependent variables together to provide a comparison of the magnitudes and to facilitate the discussion of substitution between the different wealth measures.

The results exhibit a number of striking features. First, the difference between the QR and IQR estimators is not dramatic. Both exhibit the same pattern of results, though there is some upward bias evident in the QR estimates. This bias is most evident in the estimates for net financial assets and net non-401(k) financial assets, but is hardly noticeable in the total wealth results.

Another interesting feature of the results is that the effect of participation on net financial assets is highly non-constant, appearing to increase monotonically in the quantile index. This result suggests that, conditional on income and other observables, people who rank higher in the conditional wealth distribution are impacted far more than those ranking lower in the conditional distribution. In addition, the effect is strongly positive across the entire distribution. While these results correspond to our intuition, there is actually no other a priori reason to believe that net financial assets must react in this way. In particular, if people were simply substituting financial assets held in 401(k)s for other forms of financial assets, the effect of 401(k) participation on net financial assets would be zero. These results provide strong evidence against this hypothesis at all quantiles.

The impact of 401(k) participation on total wealth relative to its impact on net financial assets also provides interesting insights. As with net non-401(k) financial assets, the effect

²⁴Standard errors were estimated using heteroskedasticity consistent standard errors as in Powell (1984, 1986) and Buchinsky (1995) using the methods outlined in Chernozhukov and Hansen (2001).

of participating in a 401(k) on total wealth is roughly constant, though in this case it is uniformly positive. The most interesting feature of the effect on total wealth is that for low quantiles it is of almost the same magnitude as the effect on net financial assets, while it is substantially smaller than the effect on net financial assets in the upper tail of the distribution. Taken together, these findings suggest that the increase in net financial assets observed in the lower tail of the conditional assets distribution can be interpreted as an increase in wealth, while the increase in the upper tail of the distribution is being mitigated by substitution with some other (non-financial) component of wealth. However, even for the highest quantiles, the substitution does not appear to be complete.

A final outstanding feature of the results is the indication that 2SLS estimates substantially overstate the treatment effect across a large range of the net financial asset distribution. In fact, the 2SLS estimates of the treatment effect on net financial assets correspond much more closely to the treatment effect at the 75th percentile of the distribution than to that of the median.

In order to strengthen and further develop our conclusions, we present test results based on the empirical instrumental quantile regression process computed using the methods of Chernozhukov and Hansen (2002). Kolmogorov-Smirnov (KS) test statistics and 95% critical values are given in Table 4. The test results lend further support to the conclusions already drawn. The tests strongly reject the hypothesis that the impact of 401(k) participation on net financial assets is constant and confirm that the impact is significantly different from zero. In addition, we see that the hypothesis of exogeneity of treatment is rejected for net financial assets. However, the tests fail to reject both the hypothesis of a constant treatment effect (equal to the median effect) and the hypothesis of exogeneity for total wealth and net non-401(k) financial assets. That the treatment effect for both total wealth and net non-401(k) financial assets is statistically constant adds further credibility to the conclusion that there is little substitution between 401(k) assets and other forms of wealth in the low tail of the assets distribution but that there is substantial substitution in the upper tail. In addition, the results of the exogeneity tests provide some evidence that there is endogeneity bias in the conventional QR estimates of the treatment effects.

4.4. Quantile Regression and Instrumental Quantile Regression Results: By Income Interval. As with analysis of the mean effect presented above, additional insights about the QTE may be gained by examining the effect of 401(k) participation on our chosen wealth measures within income interval. The independence assumption, A2, may also be more plausible within income categories due to the finer conditioning on income, since the arguments of Poterba, Venti, and Wise (1995) suggest that 401(k) eligibility is as good as randomly assigned once income is conditioned upon. Of course, the estimates within income category do suffer from a loss of precision relative to estimates obtained with a coarser income control, which makes drawing firm inferences more difficult.

IQR estimation results by income category are reported in Figures 2-5.²⁵ The figures are arranged by dependent variable, with Figure 2 corresponding to net financial assets, Figure 3 to net non-401(k) financial assets, and Figure 4 to total wealth. In all cases, the shaded region represents the 95% confidence interval.²⁶ Figure 5 contains plots of the estimated effects for each of the dependent variables together to facilitate comparison of the magnitudes. Table 5 reports process test results.

Within income categories, the results for net financial assets follow roughly the same pattern as the results in the full sample. In all categories, the results are generally increasing in the quantile index, and in all but the first income category, the process tests reveal that the treatment effect is different from zero. In addition, the hypothesis of a constant effect is rejected in all but the first and last income categories. As would be expected, the magnitudes of the results increases as income increases (as one progresses from Interval 1 to Interval 7). The point estimates in the first category are close to zero across the majority of the quantiles, suggesting that participation in a 401(k) has little effect on those with incomes of less than \$10,000. Also, in each income interval, the results are fairly constant and quite modest for quantiles below the median. Overall, these results indicate that 401(k) participation increases accumulated net financial assets in all, excepting possibly the first, income categories, but that these effects may be quite modest through much of the distribution.

As with the results in the full sample, the estimated treatment effect of 401(k) participation on net non-401(k) financial assets is not significantly different from zero in any case. The point estimates are also generally quite small, though they do exhibit some tendency to be negative more often than positive. This negative tendency provides weak evidence for some substitution between financial assets held in 401(k)s and other forms of financial assets. That this negative tendency appears to be most pronounced for low quantiles also suggests that those with low preferences for saving, who probably have relatively little in the form of financial assets, are choosing to accumulate assets within 401(k)s instead of elsewhere, whereas those with higher preferences for saving are saving in both locations.

The results for the effect of 401(k) participation on total wealth are the most varied across income categories, though the lack of precision makes comparison difficult. One result which is quite interesting is that within the lowest income category, there appear to be extreme outliers in the upper tail of the distribution. Examining the quantile results suggests that for the majority of the wealth distribution there is little effect of 401(k) participation on wealth within the first income category. However, at approximately the 60th percentile, the effects increase dramatically. These large effects in the upper tail also explain the anomalous OLS and 2SLS results within the first income category illustrated in Table 2. The process test of no effect does not reject, which seems to be a plausible conclusion given the small effect for most quantiles. Interestingly, in intervals 2, 5, and 6, the effect of participation

²⁵QR results are not reported but are quite similar to the IQR results and are available upon request.

²⁶Standard errors were estimated using heteroskedasticity consistent standard errors as in Powell (1984, 1986) and Buchinsky (1995) using the methods outlined in Chernozhukov and Hansen (2001).

on total wealth appears to track the effect on net financial assets quite closely, suggesting that little of the gain to assets due to 401(k) participation is substituting for other forms of wealth. However, in intervals 3, 4, and 7, we see that in the upper quantiles, the wealth effect is smaller than the net financial assets effect. This effect is especially pronounced in the highest income category, where the estimated participation effect on total wealth is close to zero in the upper quantiles of the wealth distribution. Overall, it is difficult to draw any firm conclusions due to large estimated standard errors of the effects. However, one robust finding seems to be that the estimated effect of participation on total wealth and the estimated effect of participation on net financial assets are quite similar in the lower tail of the wealth distribution, which suggests that participation in 401(k) plans stimulates asset accumulation of those with low preferences for saving.

A final interesting note is that, within income categories, the hypothesis of the exogeneity of 401(k) participation is never rejected. This could be because, conditional on income and other covariates, 401(k) participation is as good as randomly assigned, or it could be driven by small sample size and the lack of precision of the estimates. We choose to focus on the IQR estimates because they are robust to endogeneity, but there is no statistical evidence that endogeneity is present.

4.5. Comparison with Abadie, Angrist, and Imbens (2002) Estimator. One key criticism of the approach pursued thus far in this paper is that employer matching practices may invalidate the similarity assumption required in the model in Section 2. However, since both the instrument and endogenous variable are binary, the estimator of Abadie, Angrist, and Imbens (2002) applies. A comparison between the results from the two estimators then provides a specification check of the developed results. In addition, the comparison serves as an interesting illustration of two different methods which may be used to recover the quantile treatment effects.

The estimator of Abadie, Angrist, and Imbens (2002) is developed within the LATE framework of Imbens and Angrist (1994). In particular, Abadie, Angrist, and Imbens (2002) show that if

1. the instrument Z is independent of the outcome error, U_d in our notation, and the error in the selection equation, V in our notation,
2. monotonicity, $P(D_1 \geq D_0|X) = 1$ where D_1 is the treatment state of an individual when $Z = 1$ and D_0 is defined similarly, holds,
3. and other standard conditions are met,

then the QTE for compliers, those individuals with $D_1 > D_0$, is identified and develop an estimator for the QTE for compliers.²⁷ Since only individuals eligible for a 401(k) can

²⁷In other words, the model of Abadie, Angrist, and Imbens (2002) replaces similarity and our independence assumption with monotonicity and a stronger independence assumption and identifies a different quantity: the QTE for compliers. In the context of 401(k) participation, the group of compliers is the group of individuals who would participate in a 401(k) if eligible but would not if ineligible. Non-compliers in this example are people who would not participate in the 401(k) regardless of their eligibility status.

participate, monotonicity holds trivially, and the QTE for compliers corresponds to the QTE for the treated, which will correspond to the quantity identified by the IVQT model of Section 2 if the treated are representative of the population and the assumptions of the IVQT model are satisfied.

Given that the two models are mutually compatible under the conditions outlined above and the monotonicity assumption of Abadie, Angrist, and Imbens (2002) holds in the case of 401(k) participation, a comparison of the previous results obtained via IQR and results obtained via the estimator of Abadie, Angrist, and Imbens (2002) provides a useful robustness check of the previous results and the assumptions which underlie their interpretation. Figure 6 reports results from the estimator of Abadie, Angrist, and Imbens (2002) in the full sample and comparisons with corresponding IQR estimates.²⁸ From this exercise, we see that the pattern of results obtained from the two estimators are quite similar, with the major differences being that the Abadie, Angrist, and Imbens (2002) estimates of the effects of 401(k) participation on total wealth and on net non-401(k) financial assets appear to be even more constant than those obtained through IQR.²⁹

Overall, it appears that the difference in the estimates is small relative to sampling variation and that one would not draw substantively different conclusions from either set of estimates, though formal testing of these statements would require developing theory for the empirical process of the Abadie, Angrist, and Imbens (2002) estimator and deriving the joint distribution of the two estimators. The striking similarity between the estimates provides further support for the IQR results discussed above and strongly suggests that employer matching of 401(k) contributions does not result in failure of rank similarity.

4.6. Overall Conclusions and Cautions. Overall, the results indicate that 401(k) participation has a positive impact on the accumulation of net financial assets. The results suggest that the effect on net financial assets is increasing as one approaches the upper tail of the net financial asset distribution. Estimates for the effect of 401(k) participation on total wealth and net non-401(k) financial assets are approximately constant and indicate that 401(k) participation generally increases total wealth but has little effect on net non-401(k) financial assets. We interpret these results as indicating that participation in 401(k)s increases total wealth and that there is little substitution between financial assets in 401(k)s and other financial assets. In addition, the results suggest that there is substitution between assets held in 401(k)s and other components of wealth in the upper tail of the wealth distribution, but that most financial assets held in 401(k)s in the lower tail of the distribution represent new savings. This has important policy implications, as the people

²⁸The Abadie, Angrist, and Imbens (2002) estimator may be computed by running weighted quantile regression, where the weights are nonparametrically estimated. In our analysis, we used series methods to estimate the weights. The exact parameterization used to estimate the weights is available upon request. We also found that the overall results were not sensitive to the exact specification used to estimate the weights.

²⁹Estimates using the estimator of Abadie, Angrist, and Imbens (2002) within income categories were also very similar to the IQR estimates previously reported.

in the low tail of the net financial asset distribution are also likely to be the people with the lowest retirement savings.

The estimates also clearly indicate the inability of a single summary statistic, such as the 2SLS regression estimate of the treatment impact, to provide a clear picture of the impact of a program on the distribution of the outcomes of interest. The 2SLS estimate for the effect of 401(k) participation on net financial assets appears to overstate the actual treatment across much of the distribution, corresponding most closely to the estimates for the upper tail of the asset distribution. In addition, the single summary statistic provided by 2SLS or OLS obscures the regions where divergences between the effect of 401(k) participation on the different wealth measures occur and thus do not provide as full a description of the program impact as the quantile-based methods.

While we feel that this paper provides insight into the effect of 401(k) participation on wealth, it does suffer from limitations. First, all of the dependent variables used in this analysis represent stocks of assets rather than the flow of savings. The accumulated level of assets is interesting because it provides a summary of a person's wealth and the resources that are available to the individual. However, they are not sufficient to capture the effect of the program on savings. In particular, given employer matching and the tax-advantaged nature of 401(k) saving, it may be possible to have a large increase in accumulated assets with little change in the individual's flow of savings. Second, the data available in the SIPP do not report all sources of pension wealth. In particular, the SIPP does not contain information on assets held in DB plans or defined contribution plans other than IRAs and 401(k)s. The lack of these data could potentially bias the results upward if 401(k) assets are substituting for these other forms of assets. While evidence from Poterba, Venti, and Wise (2001) and Papke, Peterson, and Poterba (1996) is consistent with the view that 401(k)s rarely cause DB termination, it does not preclude substantial substitution between the different forms of pensions.

5. CONCLUSION

In this paper, we apply the instrumental quantile regression model and estimators developed in Chernozhukov and Hansen (2001) and Chernozhukov and Hansen (2002) to data from the SIPP which has previously been used by Poterba, Venti, and Wise (1996), Abadie (2001), Benjamin (2003), and Engen, Gale, and Scholz (1996) to examine the effects of 401(k) plans on savings. Following Poterba, Venti, and Wise (1996), Abadie (2001), and Benjamin (2003), we use 401(k) eligibility as an instrument for 401(k) participation to estimate the QTE of participation in a 401(k) plan on various wealth measures. The QTE provide a more full characterization of the effect of 401(k) participation on savings than do conventional IV methods and supplement these methods by providing a more detailed description of the distributional impact of 401(k) program participation.

The IQR estimates suggest that the effect of 401(k) participation on net financial assets is quite heterogeneous, with the largest returns accruing to those who are highest in the assets distribution. The results also indicate that the effect of 401(k) participation on

net financial assets is positive and significant over the entire range of the asset distribution and that the effect is roughly monotonically increasing in the quantile index. Effects on total wealth and net non-401(k) financial assets, on the other hand, appear to be roughly constant, and the effect on net non-401(k) financial assets is not significantly different from zero while the effect on total wealth is positive and significant. Overall, the results suggest that participation in 401(k)s increases net financial assets across the asset distribution, but that this effect is mitigated by substitution with other forms of wealth in the upper tail of the distribution. They also demonstrate that estimates of treatment effects which focus on a single feature of the outcome distribution may fail to capture the full impact of the treatment and that examining additional features may enhance our understanding of the economic relationships involved.

REFERENCES

- ABADIE, A. (1995): "Changes in Spanish Labor Income Structure During the 1980s: A Quantile Regression Approach," CEMFI Working Paper.
- ABADIE, A. (2001): "Semiparametric instrumental variable estimation of treatment response models," Harvard, mimeo.
- ABADIE, A., J. ANGRIST, AND G. IMBENS (2002): "Instrumental Variables Estimates of the Effect of Subsidized Training on the Quantiles of Trainee Earnings," *Econometrica*, 70(1), 91–117.
- AMEMIYA, T. (1982): "Two Stage Least Absolute Deviations Estimators," *Econometrica*, 50, 689–711.
- BENJAMIN, D. J. (2003): "Does 401(k) eligibility increase saving? Evidence from propensity score subclassification," *Journal of Public Economics*, 87(5-6), 1259–1290.
- BUCHINSKY, M. (1995): "Estimating the asymptotic covariance matrix for quantile regression models: A Monte Carlo study," *Journal of Econometrics*, 68, 303–338.
- CHEN, L., AND S. PORTNOY (1996): "Two-stage regression quantiles and two-stage trimmed least squares estimators for structural equation models," *Comm. Statist. Theory Methods*, 25(5), 1005–1032.
- CHERNOZHUKOV, V., AND C. B. HANSEN (2001): "An IV Model of Quantile Treatment Effects," September, Department of Economics Working Paper (download at www.ssrn.com).
- (2002): "Inference for Distributional Effects using Instrumental Quantile Regression," March, MIT Department of Economics Working Paper (download at www.ssrn.com).
- DOKSUM, K. (1974): "Empirical probability plots and statistical inference for nonlinear models in the two-sample case," *Ann. Statist.*, 2, 267–277.
- EMPLOYEE BENEFIT RESEARCH INSTITUTE (1997): *Fundamentals of Employee Benefit Programs*. EBRI, Washington, DC.
- ENGEN, E. M., AND W. G. GALE (2000): "The effects of 401(k) plans on household wealth: differences across earnings groups," NBER Working Paper 8032.
- ENGEN, E. M., W. G. GALE, AND J. K. SCHOLZ (1996): "The illusory effects of saving incentives on saving," *Journal of Economic Perspectives*, 10(4), 113–138.
- HECKMAN, J., AND R. ROBB (1986): "Alternative Methods for Solving the Problem of Selection Bias in Evaluating the Impact of Treatments on Outcomes," in *Drawing Inference from Self-Selected Samples*, ed. by H. Wainer, pp. 63–107. Springer-Verlag, New York.
- HECKMAN, J. J., AND J. SMITH (1997): "Making the most out of programme evaluations and social experiments: accounting for heterogeneity in programme impacts," *Rev. Econom. Stud.*, 64(4), 487–535, With the assistance of Nancy Clements, Evaluation of training and other social programmes (Madrid, 1993).

- IMBENS, G. W., AND J. D. ANGRIST (1994): "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 62(2), 467–475.
- KOENKER, R., AND G. S. BASSETT (1978): "Regression Quantiles," *Econometrica*, 46, 33–50.
- PAPKE, L. E., M. PETERSON, AND J. M. POTERBA (1996): "401(k) plans and tax-deferred savings," in *Advances in the Economics of Aging*, ed. by D. A. Wise. University of Chicago Press: Chicago.
- POTERBA, J. M., S. F. VENTI, AND D. A. WISE (1994): "401(k) plans and tax-deferred savings," in *Studies in the Economics of Aging*, ed. by D. A. Wise. University of Chicago Press: Chicago.
- (1995): "Do 401(k) contributions crowd out other personal saving?," *Journal of Public Economics*, 58, 1–32.
- (1996): "Personal retirement saving programs and asset accumulation: Reconciling the evidence," NBER Working Paper 5599.
- (2001): "The transition to personal accounts and increasing retirement wealth: Macro and micro evidence," NBER Working Paper 8610.
- POWELL, J. L. (1986a): "Censored regression quantiles," *Journal of Econometrics*, 32, 143–155.
- (1986b): "Least absolute deviations estimation for the censored regression model," *Journal of Econometrics*, 25, 303–325.
- VYTLACIL, E. (2000): "Semiparametric Identification of the Average Treatment Effect in Nonseparable Models," mimeo.

TABLE 1. Means and Standard Deviations

	Entire Sample	By 401(k) Participation		By 401(k) Eligibility	
		Participants	Non- Participants	Eligibles	Non- Eligibles
<i>Treatment:</i>					
Participation in 401(k)	0.26 (0.44)			0.70 (0.46)	0.00 (0.00)
<i>Instrument:</i>					
Eligibility for 401(k)	0.37 (0.48)	1.00 (0.00)	0.15 (0.36)		
<i>Outcome Variables:</i>					
Net Financial Assets	18,051 (63,523)	38,262 (79,088)	10,890 (55,257)	30,347 (74,800)	10,788 (54,518)
Median	[1,499]	[15,249]	[200]	[9,122]	[145]
Net Non-401(k) Assets	13,877 (59,605)	22,775 (70,415)	10,724 (54,930)	19,396 (67,439)	10,617 (54,192)
Median	[542]	[3,830]	[200]	[2,711]	[130]
Total Wealth	63,817 (111,530)	96,920 (127,790)	52,088 (102,646)	86,240 (124,006)	50,571 (101,155)
Median	[25,100]	[53,441]	[16,645]	[45,356]	[14,640]
<i>Covariates:</i>					
Income	37,201 (24,774)	49,367 (27,208)	32,890 (22,316)	46,862 (25,958)	31,494 (22,151)
Age	41.06 (10.34)	41.51 (9.66)	40.90 (10.57)	41.48 (9.61)	40.81 (10.75)
Family Size	2.87 (1.54)	2.92 (1.47)	2.85 (1.56)	2.90 (1.48)	2.84 (1.57)
Married	0.60 (0.49)	0.69 (0.46)	0.57 (0.49)	0.67 (0.47)	0.56 (0.50)
Participation in IRA	0.24 (0.43)	0.36 (0.48)	0.20 (0.40)	0.32 (0.47)	0.20 (0.40)
Defined Benefit Pension	0.27 (0.44)	0.39 (0.49)	0.23 (0.42)	0.42 (0.49)	0.19 (0.39)
Home Owner	0.64 (0.48)	0.77 (0.42)	0.59 (0.49)	0.74 (0.44)	0.57 (0.49)
<i>Years Education:</i>					
< 12	0.13 (0.33)	0.07 (0.25)	0.15 (0.36)	0.07 (0.26)	0.16 (0.37)
12	0.38 (0.48)	0.35 (0.48)	0.39 (0.49)	0.35 (0.48)	0.39 (0.49)
> 12 and < 16	0.24 (0.43)	0.26 (0.44)	0.24 (0.43)	0.26 (0.44)	0.24 (0.42)
> 16	0.25 (0.43)	0.33 (0.47)	0.22 (0.42)	0.32 (0.47)	0.21 (0.41)

The sample is drawn from the 1991 SIPP and consists of 9915 observations. The observational units are household reference persons aged 25-64 and spouse if present. Households are included in the sample if at least one person is employed and no one is self-employed. Standard deviations are in parentheses, and medians are in brackets.

TABLE 2. Means and Standard Deviations of Asset Measures by Income Interval

	Income						
	<\$10K	\$10-20K	\$20-30K	\$30-40K	\$40-50K	\$50-75K	\$75K+
Net Financial Assets	735 (10,827)	2,308 (15,498)	6,311 (30,615)	11,938 (43,519)	19,348 (54,773)	33,708 (66,894)	83,709 (157,168)
Median	[0]	[0]	[400]	[2,053]	[5,761]	[14,500]	[43,779]
Net Non-401(k) Financial Assets	431 (9,143)	1,543 (14,699)	4,979 (29,525)	8,775 (40,991)	14,942 (52,718)	25,179 (62,438)	66,999 (151,627)
Median	[0]	[0]	[110]	[651]	[3,437]	[8,676]	[29,800]
Total Wealth	16,235 (40,772)	21,620 (43,631)	36,730 (67,659)	55,119 (83,203)	74,006 (97,913)	105,285 (119,531)	202,240 (226,077)
Median	[1,258]	[4,225]	[12,500]	[29,224]	[44,197]	[71,025]	[152,500]
N	638	1,948	2,074	1,712	1,204	1,572	767

The sample is drawn from the 1991 SIPP and consists of 9915 observations. The observational units are household reference persons aged 25-64 and spouse if present. Households are included in the sample if at least one person is employed and no one is self-employed. Standard deviations are in parentheses, and medians are in brackets.

TABLE 3. OLS and 2SLS Estimates of Effect of 401(k) Participation

Sample	N	First Stage	Net Financial Assets		Net Non-401(k) Financial Assets		Total Wealth	
			OLS	2SLS	OLS	2SLS	OLS	2SLS
A. Full Sample								
Full Sample	9,915	0.697 (0.006)	14,250 (1,551)	13,087 (1,922)	778 (1,477)	-355 (1,855)	10,694 (2,388)	9,259 (3,035)
B. By Income Interval								
Interval 1	638	0.711 (0.020)	9,843 (4,921)	9,149 (4,914)	4,093 (3,447)	3,443 (3,527)	20,464 (11,311)	17,224 (11,518)
Interval 2	1,948	0.650 (0.013)	5,591 (1,463)	5,352 (1,629)	-759 (1,227)	-917 (1,427)	4,729 (2,265)	6,138 (3,218)
Interval 3	2,074	0.627 (0.013)	7,083 (1,315)	4,143 (2,268)	448 (1,124)	-2,518 (2,152)	5,462 (3,119)	0.183 (4,502)
Interval 4	1,712	0.672 (0.015)	12,136 (2,513)	10,273 (2,880)	1,077 (2,297)	-909 (2,677)	10,683 (3,891)	4,881 (5,103)
Interval 5	1,204	0.723 (0.018)	12,858 (2,470)	9,980 (3,741)	500 (2,330)	-2,479 (3,646)	13,470 (4,905)	13,205 (6,675)
Interval 6	1,572	0.744 (0.017)	20,800 (3,010)	21,920 (3,444)	1,803 (2,876)	2,985 (3,310)	12,881 (5,132)	12,202 (6,718)
Interval 7	767	0.831 (0.022)	23,103 (10,417)	24,013 (12,895)	-6,735 (10,228)	-5,252 (12,713)	5,514 (13,645)	10,470 (17,174)

IMPACT OF 401(K) ON ASSETS

The table reports OLS and 2SLS estimates of the effect of 401(k) participation on wealth. The second column contains the sample size used for the estimates in each row. The third column reports the first stage coefficient estimate from a regression of 401(k) participation on 401(k) eligibility and covariates. The income intervals are as follows: <\$10K, \$10-20K, \$20-30K, \$30-40K, \$40-50K, \$50-75K, \$75K+. Covariates are as described in the main text. Heteroskedasticity robust standard errors are given in parentheses.

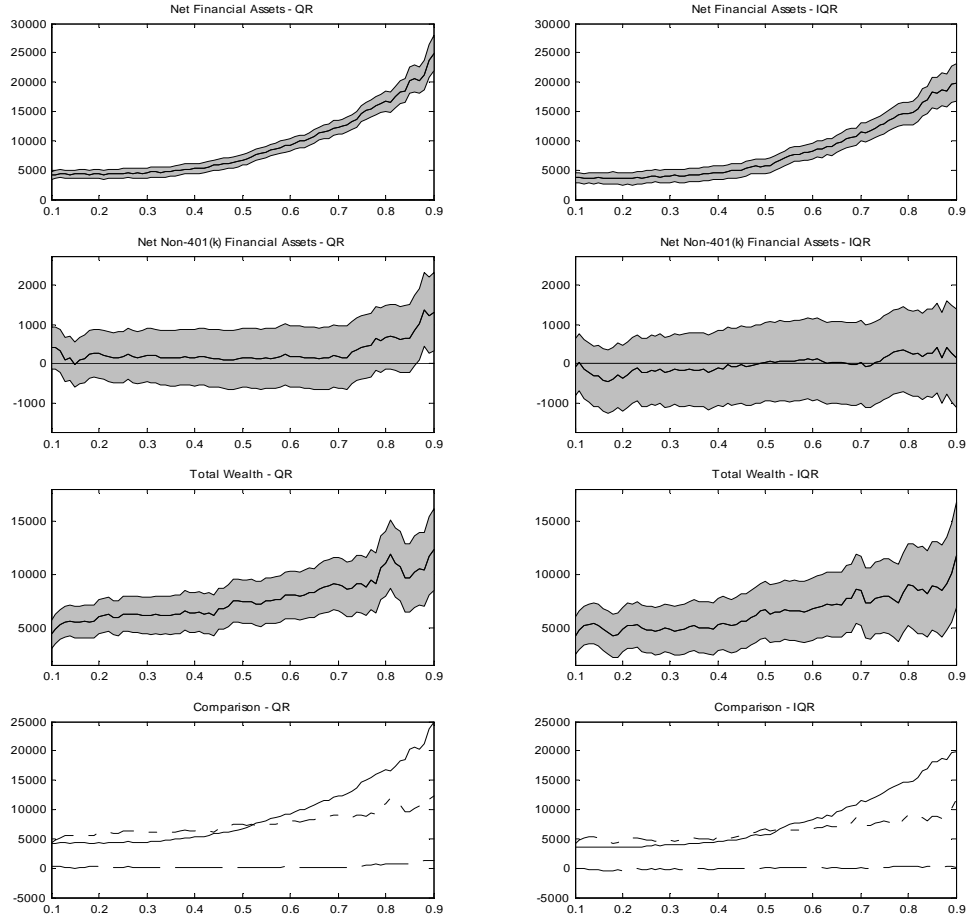


FIGURE 1. The sample size is 9,915. The left column contains standard quantile regression estimates, and the right column contains instrumental quantile regression. Each panel is labeled with the dependent variable used in estimation of the presented results. The bottom panel in each column compares the point estimates for each wealth measure. The solid line corresponds to net financial assets, the dashed line to net non-401(k) financial assets, and the dash-dot line to total wealth. The vertical axis measures the dollar increase in the wealth measure due to 401(k) participation. The quantile of the conditional wealth distribution is on the horizontal axis. Covariates are as described in the main text. The shaded region is the 95% confidence band using robust standard errors. Estimates are reported for $\tau \in [.10, .90]$ at .01 unit intervals.

TABLE 4. Tests on the Instrumental Quantile Regression Process in the Full Sample

Null Hypothesis	Net		Net Non-401(k)		Total Wealth	
	Financial Assets		Financial Assets			
	Statistic	$c_{.95}$	Statistic	$c_{.95}$	Statistic	$c_{.95}$
No Effect	12.875	3.009	0.921	2.882	4.538	3.003
Constant Effect	9.093	3.321	0.843	3.452	1.850	3.213
Exogeneity	3.851	3.209	2.287	3.056	1.899	3.086

The table report inference results on the inverse quantile regression process. Reported are Kolmogorov-Smirnov statistics and 95% critical values. The statistics and critical values are computed using the methods in Chernozhukov and Hansen (2002). The null hypotheses tested are as follows: no effect - $\alpha(\cdot) = 0$, constant effect - $\alpha(\cdot) = \alpha(.5)$, exogeneity - $\alpha(\cdot) = \alpha_{QR}(\cdot)$, where $\alpha(\cdot)$ denotes the instrumental quantile regression process and $\alpha_{QR}(\cdot)$ denotes the quantile regression process.

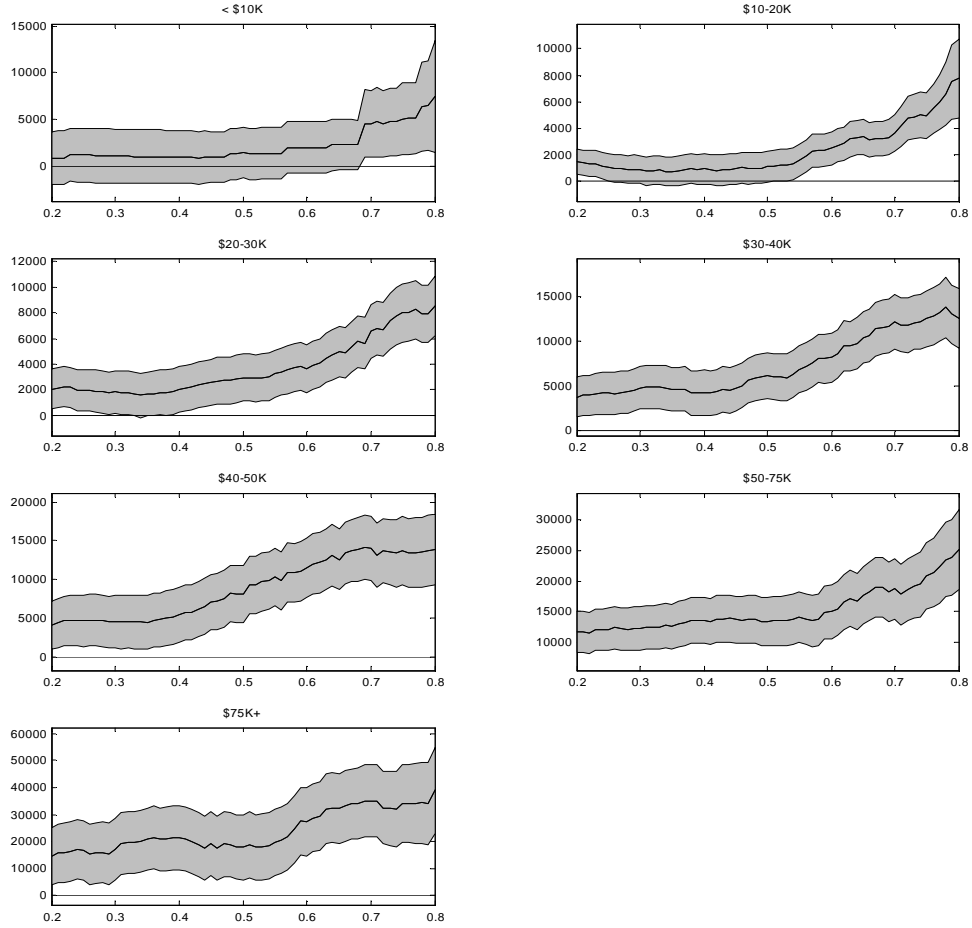


FIGURE 2. The figure reports the effect of 401(k) participation on net financial assets by income interval. Each panel is labeled with the income interval to which it corresponds. The vertical axis measures the dollar increase in net financial assets due to 401(k) participation. The quantile of the conditional net financial assets distribution is on the horizontal axis. Covariates are as described in the main text. The shaded region is the 95% confidence band using robust standard errors. Estimates are reported for $\tau \in [.20, .80]$ at .01 unit intervals.

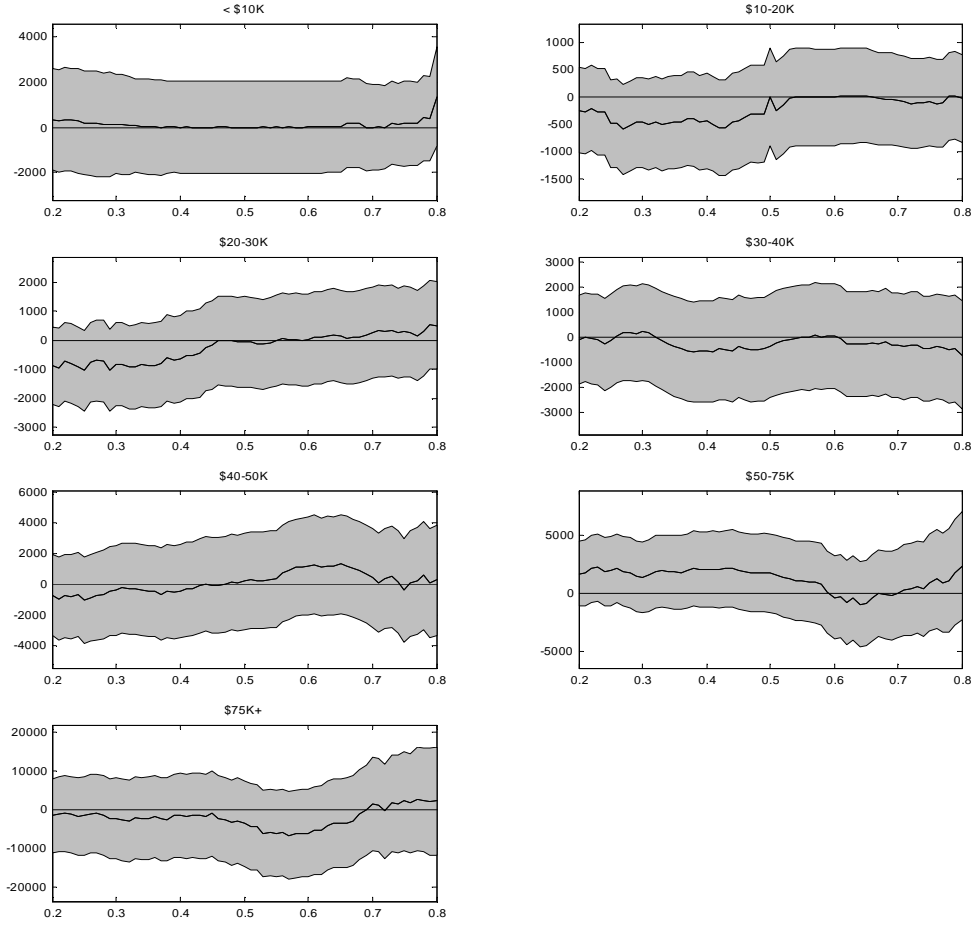


FIGURE 3. The figure reports the effect of 401(k) participation on net non-401(k) financial assets by income interval. Each panel is labeled with the income interval to which it corresponds. The vertical axis measures the dollar increase in net non-401(k) financial assets due to 401(k) participation. The quantile of the conditional net non-401(k) financial assets distribution is on the horizontal axis. Covariates are as described in the main text. The shaded region is the 95% confidence band using robust standard errors. Estimates are reported for $\tau \in [-.20, .80]$ at .01 unit intervals.

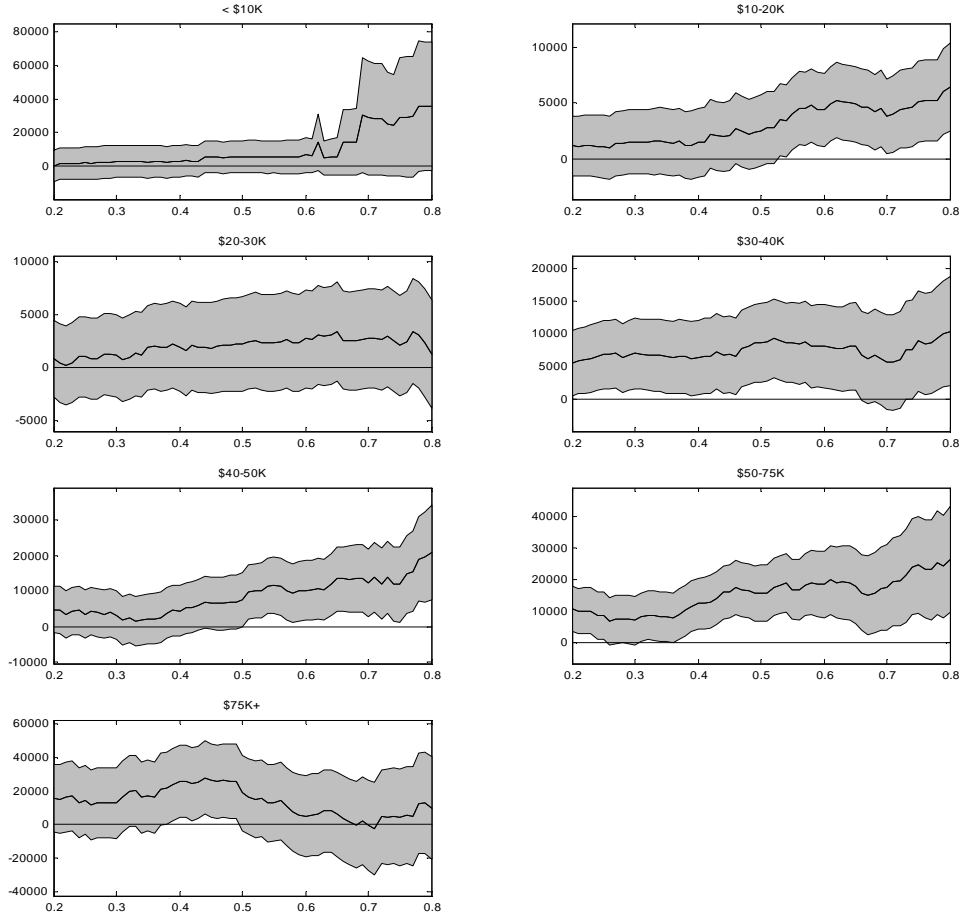


FIGURE 4. The figure reports the effect of 401(k) participation on total wealth by income interval. Each panel is labeled with the income interval to which it corresponds. The income intervals are as follows: <\$10K, \$10-20K, \$20-30K, \$30-40K, \$40-50K, \$50-75K, \$75K+. The vertical axis measures the dollar increase in total wealth due to 401(k) participation. The quantile of the conditional total wealth distribution is on the horizontal axis. Covariates are as described in the main text. The shaded region is the 95% confidence band using robust standard errors. Estimates are reported for $\tau \in [.20, .80]$ at .01 unit intervals.

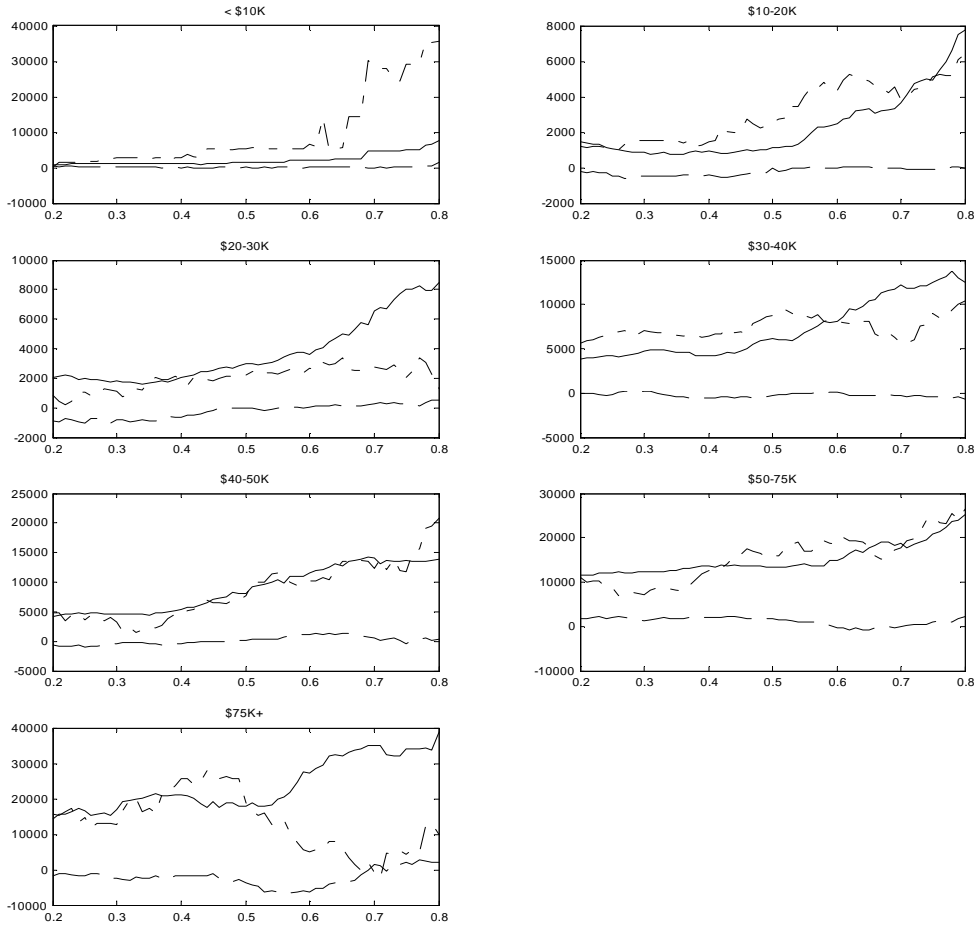


FIGURE 5. The figure compares the effect of 401(k) participation on net financial assets, net non-401(k) financial assets, and total wealth by income interval. Each panel is labeled with the income interval to which it corresponds. The vertical axis measures the dollar increase in total wealth due to 401(k) participation. The quantile of the conditional total wealth distribution is on the horizontal axis. Covariates are as described in the main text. The solid line corresponds to net financial assets, the dashed line to net non-401(k) financial assets, and the dash-dot line to total wealth. Estimates are reported for $\tau \in [.20, .80]$ at .01 unit intervals.

TABLE 4. Tests on the Instrumental Quantile Regression Process by Income Interval

Null Hypothesis	Net					
	Net		Net-401(k)		Total Wealth	
	Statistic	$c_{.95}$	Statistic	$c_{.95}$	Statistic	$c_{.95}$
A. Interval 1						
No Effect	2.231	2.909	0.982	3.209	1.744	2.955
Constant Effect	2.021	3.786	0.982	3.947	1.636	3.894
Exogeneity	1.981	3.801	1.469	4.032	1.243	3.722
B. Interval 2						
No Effect	4.811	2.754	1.177	2.725	2.723	2.758
Constant Effect	4.649	3.169	1.633	3.361	2.114	3.455
Exogeneity	1.927	3.012	1.756	3.073	0.783	3.058
C. Interval 3						
No Effect	6.133	2.603	1.217	2.817	1.204	2.786
Constant Effect	4.418	3.140	1.260	3.262	0.807	3.145
Exogeneity	1.367	2.930	1.292	2.909	2.472	3.132
D. Interval 4						
No Effect	6.746	2.895	0.550	2.738	2.538	3.001
Constant Effect	4.243	3.186	0.707	3.127	1.367	3.340
Exogeneity	1.451	2.967	0.964	3.087	1.863	2.831
E. Interval 5						
No Effect	5.678	2.760	0.673	3.000	2.648	2.724
Constant Effect	3.139	3.124	0.937	3.298	2.286	3.349
Exogeneity	0.980	2.972	1.127	2.959	1.909	3.134
F. Interval 6						
No Effect	6.475	2.817	1.253	2.612	3.402	2.831
Constant Effect	3.410	3.202	1.820	3.135	2.112	3.109
Exogeneity	1.009	2.880	0.947	2.878	1.674	2.859
G. Interval 7						
No Effect	4.396	2.829	0.990	2.853	2.112	2.964
Constant Effect	2.721	3.266	1.017	3.339	1.827	3.410
Exogeneity	1.359	3.321	0.823	2.945	1.098	3.089

The table report inference results on the inverse quantile regression process computed by income interval. The income intervals are as follows: <\$10K, \$10-20K, \$20-30K, \$30-40K, \$40-50K, \$50-75K, \$75K+. Reported are Kolmogorov-Smirnov statistics and 95% critical values. The statistics and critical values are computed using the methods in Chernozhukov and Hansen (2002). The null hypotheses tested are as follows: no effect - $\alpha(\cdot) = 0$, constant effect - $\alpha(\cdot) = \alpha(.5)$, exogeneity - $\alpha(\cdot) = \alpha_{QR}(\cdot)$, where $\alpha(\cdot)$ denotes the instrumental quantile regression process and $\alpha_{QR}(\cdot)$ denotes the quantile regression process.

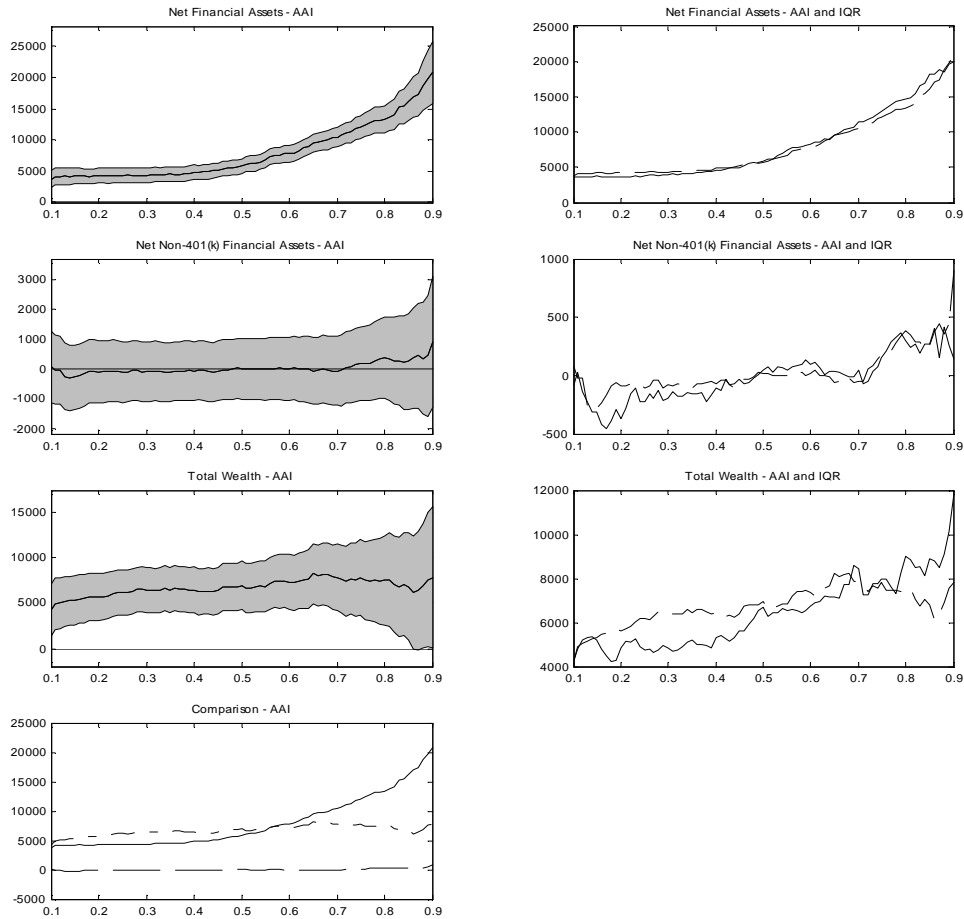


FIGURE 6. The sample size is 9,915. The left column contains estimates obtained using the estimator of Abadie, Angrist, and Imbens (2002) (AAI), and the right column compares them with the corresponding estimates obtained through the IQR estimator of Chernozhukov and Hansen (2001). The solid line corresponds to the IQR estimator, and the dashed line corresponds to the AAI estimator. Each panel is labeled with the dependent variable used in estimation of the presented results. The bottom panel in the left column compares the AAI point estimates for each wealth measure. The solid line corresponds to net financial assets, the dashed line to net non-401(k) financial assets, and the dash-dot line to total wealth. The vertical axis measures the dollar increase in the wealth measure due to 401(k) participation. The quantile of the conditional wealth distribution is on the horizontal axis. Covariates are as described in the main text. The shaded region is the 95% confidence band using robust standard errors. Estimates are reported for $\tau \in [.10, .90]$ at .01 unit intervals.