Social Security and Trends in Wealth Inequality*

Sylvain Catherine[†] Max Miller[‡] Natasha Sarin[§]

July 30, 2022

Abstract

Recent influential work finds large increases in inequality in the U.S. based on measures of wealth concentration that notably exclude the value of social insurance programs. This paper shows that top wealth shares have not increased in the last three decades when Social Security is properly accounted for. This is because Social Security wealth increased substantially from $3.4 trillion in 1989 to $37.2 trillion in 2019 and is now 57% of the wealth of the bottom 90% of the wealth distribution. This finding is robust to potential changes to taxes and benefits in response to system financing concerns.

Keywords: Social Security, Inequality, Top Wealth Shares
JEL codes: D31, E21, G51, H55, N32

*We are thankful to our discussants Adrien Auclert, Alan Auerbach, Alberto Bisin, Francisco Gomes, Daniel Greenwald, Fatih Karahan, Greg Mankiw, and Owen Zidar. We also benefited from insightful comments from Steve Cecchetti, Jason Furman, Joao Gomes, Wojtek Kopczuk, Antoine Levy, Hanno Lustig, Olivia Mitchell, Adair Morse, Josh Rauh, Kim Schoenholtz, Andrei Shleifer, Lawrence Summers, Constantine Yannelis, and Eric Zwick. We are also thankful for helpful comments from seminar participants at Wharton, Stanford GSB, New York University Stern, Yale University, The Federal Reserve Board, the Federal Reserve of Chicago, George Washington University and conference participants at the NBER Summer Institute (Income and Wealth), the NBER meeting in Public Economics, the NBER Summer Institute (Inequality and the Macroeconomy), the Western Finance Association, SFS Calvacade, the Red Rock Finance Seminar, CEPR Household Finance Conference, Chicago Booth Household Finance Conference, ASU Sonoran Winter Finance Conference, Northern Finance Association. We are grateful to Maria Gelrud for her outstanding research assistance. All remaining errors are our own. Previous versions of this paper were circulated with the title "Social Security and Trends in Inequality."

[†]University of Pennsylvania, Wharton
[‡]University of Pennsylvania, Wharton
[§]University of Pennsylvania, Law School and Wharton

1 Introduction

As work dating back to Feldstein (1976) shows, total wealth—inclusive of Social Security—is much more evenly distributed than private wealth alone. This paper highlights that including Social Security does not just change the level of wealth inequality, but its evolution as well. Specifically, although it is widely believed that wealth inequality in the United States is on the rise, we find that top wealth shares have not increased in the last three decades when Social Security is properly taken into account.

Accounting for Social Security matters for two reasons. First, asset price increases drive recent trends in U.S. wealth inequality (Hubmer, Krusell and Smith, 2020; Moll, 2020). Indeed, shifts in interest rates have a disproportionate effect on the market value of long duration assets, which are primarily owned by wealthy households, leading to greater wealth inequality (Greenwald, Leombroni, Lustig and Nieuwerburgh, 2021). However, the focus on market wealth overlooks an important long duration asset which represents a disproportionate part of the balance sheet of low and middle-class households and has increased in value: their Social Security benefits.

Second, as the simple conceptual framework we begin with shows, because mandatory Social Security contributions represent a larger fraction of the balance sheets of those outside of the top of the distribution, modest variations in savings rates have larger consequences on their private wealth accumulation. Because of this mechanical leverage effect, an overall decline in saving rates translates into greater market wealth inequality, but does not necessarily affect the distribution of total wealth, inclusive of Social Security. This becomes especially important when one considers that contribution rates into Social Security have increased from less than 2% in the 1950s to roughly 11% since 1988 (Miron and Weil, 1998).

Excluding Social Security wealth from inequality measures has broad policy implications. Perversely, existing wealth concentration measures that ignore the substitution between private and public wealth could mistakenly conclude that progressive social programs increase inequality, rather than redress it. A more inclusive wealth concept, in contrast, enables proper evaluation of the role redistributive public programs play in curbing inequality.

To incorporate Social Security into top wealth estimates, we must know both the aggregate size of the Social Security program and how Social Security wealth is distributed across the marketable wealth distribution. This paper derives estimates of the stock and distribution of Social Security wealth by simulating households' future benefits and payroll taxes, relying on data from the Survey of Consumer Finances (SCF). Our estimates are conservative since we focus on Social Security's old-age retirement program, excluding disability insurance which would lead to an even larger reduction in top wealth shares if included.

For retirees, we can calculate Social Security wealth from the SCF directly using reported benefits. For workers who are still in the labor force, we simulate earnings trajectories by relying on previous empirical work that provides a labor income process that matches many moments of the SSA administrative panel data (Guvenen, Karahan, Ozkan and Song, 2021). We then apply the Social Security benefit and tax formulas to construct estimates of future retirement benefits that these households will accrue, net of the taxes that they will pay. We validate these estimates by comparing them to aggregate wealth estimates reported by the SSA and to benefits reported for retirees in the SCF. Finally, we determine the share of Social Security wealth going to the top of the wealth distribution based on the relationship between Social Security and marketable wealth for retired workers, readily observable in the SCF.

Computing the present value of Social Security wealth also requires choosing an appropriate discount rate. For example, Novy-Marx and Rauh (2011) discount public employee pension promises using state yield curves. Similarly, we first offer a risk-free valuation of Social Security wealth using the Treasury market yield curve. We find that the share of "marketable wealth" owned by the top 10% and top 1% grew by approximately 10 percentage points (pp) between 1989 and 2019, in line with existing estimates (Smith, Zidar and Zwick, 2020). Once Social Security wealth is included, these trends disappear: the share of the top 10% and top 1% actually dropped by 5.4 and 1.3 pp, respectively.

However, discounting should reflect the risks associated with the Social Security program. As such, our second set of results accounts for the labor market risk inherent in pay-as-you-go systems. Social Security is wage-indexed, so future benefits are directly tied to economic growth. Given the cointegration between the labor and stock markets (Benzoni, Collin-Dufresne and Goldstein, 2007), it is important to adjust for the market beta of future Social Security payouts (Geanakoplos and Zeldes, 2010). Our risk-adjustment decreases the stock of Social Security wealth by nearly 20%, with a disproportionate effect on the bottom 90%. Even after this correction, inequality trends are substantially attenuated relative to past estimates that exclude Social Security. From 1989 to 2019, the top 10% wealth share decreases by 2.4 pp. The top 1% share remained flat, increasing by 0.2 pp.

Our findings are also robust to alternative assumptions that incorporate the risk of future benefit cuts (or tax hikes), differences in life expectancy among the rich and the poor, adding a substantial liquidity premium to discount rates, and the possibility of persistently low economic growth.

Overall, Social Security dramatically impacts inequality trends because the growth in Social Security wealth has outpaced the growth in marketable wealth over the last three decades. This increase can be attributed to three factors. First, Social Security expanded in scope over our sample period, as the share of earnings subject to Social Security payroll taxes increased from a maximum of 1.25 times average annual earnings to 2.5 times. Second, there have been demographic shifts: the U.S. population is aging and living longer. The share of workers near retirement age and for whom Social Security wealth is at its peak (because they have paid in fully to the fund, but have yet to receive any benefits) grew by nearly 50%. Moreover, life expectancy increased by nearly 4 years.

Finally, and most importantly, real interest rates have fallen. This means that, in order to fund the same level of consumption during retirement, an investor in 2019 has to save considerably more (since less interest will accrue) or buy a higher-priced annuity than an investor in 1989. As such, Social Security's value rises, since the future purchasing power of contributions corresponds to more marketable wealth when workers face low rates of returns on their private savings.

Considering the role rates play in the evolution of Social Security wealth is imperative, since, in economic terms, the cashflows generated by Social Security and marketable assets must enter the intertemporal budget constraint of households in a consistent way. The prevailing low rate environment has driven up the valuation of marketable assets. As a simple life-cycle framework makes clear, when optimal savings rates decline for whatever reason, private savings fall disproportionately for low-earners. This is because mandatory Social Security contributions—which represent a disproportionate share of low-earning households' total savings—are not a lever for adjustment.

Falling interest rates affect wealth inequality by redistributing wealth away from holders of short-duration assets, favoring those with long-term investments (Auclert, 2019). Because long duration assets represent a greater share of the private wealth of those at the top of the distribution, marketable wealth inequality rises when interest rates fall (Greenwald et al., 2021). However, the relationship between households' wealth and the average duration of their assets is much weaker when we include Social Security, a long-term investment representing a disproportionate share of the total wealth of the bottom 90%. As Catherine, Miller, Paron and Sarin (2022) show, the large implicit Social Security wealth of poorer households reduces the optimal share of long-term assets in their market wealth, resulting in lower capital gains than wealthy households when interest rates fall. Overall, by focusing on marketable wealth alone, previous studies have taken into account the increased value of privately long-duration assets disproportionately owned by the rich but not the increased value of those owned by the rest of the population, mainly the increased market value of their Social Security benefits.

For essentially any question related to inequality trends that researchers seek to answer, a broader wealth concept that includes Social Security is valuable. For example, one reason to care about wealth inequality is that it is a measure of consumption or welfare inequality, since accumulated wealth funds retirement consumption. In this case, Social Security's inclusion is important because retirement benefits serve the same purpose. Additionally, to understand the evolution of inequality across countries or regimes, it is imperative to consider differences in pension systems that can encourage or discourage private savings. Another reason to study wealth inequality is that wealth brings political power. There too including Social Security is important. The more households can rely on public programs during retirement, the more they can afford to spend private wealth on political causes. Further, Social Security is revelatory on the counterfactual: private wealth would have evolved differently if individuals had to fund their later-life consumption through private savings.

Our findings contribute to the inequality literature in several ways. First, we provide new estimates of wealth concentration in the U.S., documenting the significant effect of the inclusion of Social Security. Specifically, we quantify how the value of Social Security has increased over the last thirty years, finding that it attenuates the upward trend in marketable wealth inequality. Our results illustrate how Social Security has decreased households' exposure to the low rate environment: In the absence of Social Security, with rates of return near zero, households would have to save more to finance a given level of consumption in retirement. But, in reality, the rate of return on Social Security contributions has not decreased as much as the return on private wealth.

Second, we shed new light on the recent divergence in savings behavior across the wealth distribution. Saez and Zucman (2016) credit a decrease in private savings for the bottom 90%, which has been 0% since 2000, as precipitating the rise in wealth inequality. This is true mechanically: If savings rates are 0% for the bottom 90%, then all the increase in wealth is captured by the top of the distribution. But private savings rates would not be zero if workers were not saving through the Social Security program. Strangely, existing estimates of savings rates and thus wealth concentration would arrive at different conclusions in the counterfactual world where Social Security contributions were invested in private accounts with the same rate of return and progressivity. Ignoring the known substitution between Social Security contribution and private savings (Attanasio and Brugiavini, 2003; Attanasio and Rohwedder, 2003; Scholz, Seshadri and Khitatrakun, 2006) is misleading.

Third, we build a measure of wealth inequality that is better suited for policy evaluation and cross-country comparisons than prior work. A focus on narrow marketable wealth inequality means that expansions of the social safety net may well register as increasing inequality, because public funds for retirement, or to cover healthcare costs or unemployment shocks, displace the need for private wealth accumulation. This is why a broader wealth concept is important. Our approach also enables comparison with other countries, as different retirement structures or the generosity of social insurance arrangements should not mechanically drive conclusions about how unequally distributed resources are.

To be sure, this is an incomplete undertaking: we too exclude important components of wealth from our estimates, for example, the provision of public healthcare benefits. It is our hope that this paper represents a first step toward a broader wealth concept that will enable accurate measurement and analysis of inequality trends.

Related Literature Narrowly defined marketable wealth understates the wealth of workers and consequently overstates inequality substantially. It also ignores a long literature that documents the importance of Social Security for the distribution of income and wealth. For instance, Wolff (1992, 1996) shows that the inclusion of pension and Social Security wealth impacts both the level of and changes in measured wage inequality. Gustman, Mitchell, Samwick and Steinmeier (1999) investigate the importance of pension and Social Security wealth for those nearing retirement, showing that it accounts for half—or more—of the total wealth of all those below the 95th percentile of the wealth distribution. Poterba (2014) also sheds light on the importance of Social Security to the elderly, documenting that for people over age 65, this stream of cash flows accounts for more than half of total income for the bottom three quartiles of the income distribution. Outside of the US, evidence confirms that ignoring the effects of redistributive pension programs inflates measured wage inequality (Domeij and Klein, 2002).

We build on the insights of past literature to augment our definition of wealth by including the Social Security benefits that workers accrue. Feldstein (1974) does such an exercise to show that in 1962, the ownership of total wealth, inclusive of Social Security, was much less concentrated than the ownership of market wealth. We focus on trends in wealth inequality, showing that this pattern remains true, and the differences between the "market wealth" and "total wealth" series are of growing importance over time. We thus contribute to the literature by documenting the sizable impact of Social Security on trends in wealth inequality. Our exercise confirms Weil (2015) who suggests that the concept of market wealth is incomplete and overstates inequality by ignoring transfer wealth, which is both large and, unlike market wealth, not skewed to the top of the distribution. A related point has been made by Auten and Splinter (2019) in the context of income inequality, who highlight that including government transfer programs decreases top income shares, and by Auerbach, Kotlikoff and Koehler (2019) who point out that their measure of remaining lifetime spending is much more equally distributed than net wealth or current income.

2 Conceptual Framework

Should Social Security wealth be accounted for in estimates of wealth inequality? To answer this question, we calibrate an illustrative model in which households prepare for retirement by accumulating private savings and contributing to Social Security. Because of the design of the program, Social Security savings represent a larger share of the total savings of low-earning households. Further, these contributions are mandatory, which means that when optimal savings rates decline, they are not a lever of adjustment for households. As a result, low-earners adjust private savings proportionately more to arrive at a equivalent reduction in overall retirement resources. As our simple model illustrates, this means that marketable wealth inequality can rise—while consumption inequality is unchanged.

2.1 Model

In the model, agents choose their consumption C to maximize lifetime expected utility:

$$V_{it} = \mathbb{E} \sum_{s=t}^{T} \beta^{s-t} (1 - m_s) \frac{C_{it}^{1-\gamma}}{1-\gamma}, \tag{1}$$

where β is the discount factor, γ the coefficient of relative risk aversion, $1 - m_{is}$ the age-dependent probability of being alive in year s, and T the maximum lifespan.

Agents receive labor earnings L before retirement and Social Security benefits B thereafter. They invest their private wealth W in a risk-free asset with return r. Hence, the dynamic budget constraint is:

$$W_{it+1} = (W_{it} + L_{it} + B_{it} - \mathrm{T}_{it} - C_{it})\,(1+r)\,, \tag{2}$$

where T denotes Social Security payroll taxes.

Social Security Payroll taxes represent 10.6% of earnings, up to a earnings cap, which we call the Social Security wage base SSWB_t:

$$\mathrm{T}_{it} = .106 \cdot \min\left\{L_{it}, \mathrm{SSWB}_t\right\}. \tag{3}$$

For simplicity, we assume that workers retire at full-retirement age R. Benefits depend on each individuals' average indexed yearly earnings (AIYE). A worker's indexed taxable earnings in year t are:

$$\mathrm{L}_{it}^{\text{indexed}} = \min\left\{L_{it}, \mathrm{SSWB}_t\right\}\frac{L_{1,60}}{L_{1,t}}, \tag{4}$$

where $L_{1,60}$ and $L_{1,t}$ denotes the value of the national wage index the year of his 60th birthday. In practice, the AIYE is the average of the best 35 years of indexed earnings up to retirement. To keep the model tractable, we assume, for the moment, that the AIYE is the average of all working years. Social Security benefits are progressive because the replacement rate is a concave function of the AIYE. The marginal replacement rate drops at two "bend points" b_1 and b_2, which, since the 1980, have represented 21% and 125% the wage index. Hence, yearly benefits in the first retirement year are:

$$B_{it} = \begin{cases} .9 \cdot \mathrm{AIYE}_i & \text{if } \mathrm{AIYE}_i < b_1 \\ .9 \cdot b_1 + .32(\mathrm{AIYE}_i - b_1) & \text{if } b_1 \leq \mathrm{AIYE}_i < b_2 \\ .9 \cdot b_1 + .32(b_2 - b_1) + .15(\mathrm{AIYE}_i - b_2) & \text{if } b_2 \leq \mathrm{AIYE}_i. \end{cases} \tag{5}$$

After retirement, benefits grow at the rate of inflation.

Labor Income Earnings are the product of the wage index and an idiosyncratic component $L_{2,i}$:

$$L_{it} = L_{1,t} \cdot L_{2,it}. \tag{6}$$

Throughout the paper, we model idiosyncratic risk using the rich income process estimated in Guvenen, Karahan, Ozkan and Song (2021). Specifically, we assume that the idiosyncratic component of a worker's earnings $L_{2,i}$ evolves as follows:

$$\text{Idiosyncratic earnings:} \qquad L_{2,it} = (1-\nu_t^i)e^{\left(g(t)+\alpha^i+\beta^i t+z_t^i+\varepsilon_t^i\right)} \tag{10.1}$$

$$\text{Persistent component:} \qquad z_t^i = \rho z_{t-1}^i + \eta_t^i \tag{10.2}$$

$$\text{Innovations to AR(1):} \qquad \eta_t^i \sim \begin{cases} \mathcal{N}(\mu_{\eta,1}, \sigma_{\eta,1}^2) & \text{with prob. } p_z \\ \mathcal{N}(\mu_{\eta,2}, \sigma_{\eta,2}^2) & \text{with prob. } 1-p_z \end{cases} \tag{10.3}$$

$$\text{Initial condition of } z_t^i\text{:} \qquad z_0^i \sim \mathcal{N}(0, \sigma_{z,0}^2) \tag{10.4}$$

$$\text{Transitory shock:} \qquad \varepsilon_t^i \sim \begin{cases} \mathcal{N}(\mu_{\varepsilon,1}, \sigma_{\varepsilon,1}^2) & \text{with prob. } p_\varepsilon \\ \mathcal{N}(\mu_{\varepsilon,2}, \sigma_{\varepsilon,2}^2) & \text{with prob. } 1-p_\varepsilon \end{cases} \tag{10.5}$$

$$\text{Nonemployment duration:} \qquad \nu_t^i \sim \begin{cases} 0 & \text{with prob. } 1-p_\nu(t, z_t^i) \\ \min\{1, \text{Exp}\{\lambda\}\} & \text{with prob. } p_\nu(t, z_t^i) \end{cases} \tag{10.6}$$

$$\text{Prob. of Nonemp. shock:} \qquad p_\nu^i(t, z_t) = \frac{e^{a+bt+cz_t^i+dz_t^i t}}{1+e^{a+bt+cz_t^i+dz_t^i t}} \tag{10.7}$$

The persistent component of earnings z_i follows an AR(1) process with innovations drawn from a mixture of normal distributions. Transitory shocks ε_i are also drawn from a normal mixture. These normal mixtures capture high-order moments of the distribution of income shocks. Workers can experience a non-employment shock with some probability p_ν that depends on age, income, and gender, and exponentially distributed duration. In Equation (10.1), $g(t)$ captures the life-cycle profile of earnings common to all workers. The vector (α_i, β_i) determines heterogeneity in the level and growth rate of earnings and is drawn from a multivariate normal distribution with zero mean and correlation coefficient $\text{corr}_{\alpha\beta}$. Heterogeneity in initial conditions of the persistent process is captured by z_0.

Calibration We assume a risk aversion $\gamma = 1$ and a discount factor $\beta = 0.97$. The wage process parameters are calibrated as shown in Table E.1. Survival probabilities are from the Human Mortality Database for the year 2018.

2.2 The role of Social Security in household portfolios

Lifetime consumption is bounded by accrued financial resources and future labor income. The intertemporal budget constraint is:

$$\sum_{s=t}^{T} \frac{C_{is}}{(1+r)^{s-t}} \leq W_{it} + \sum_{s=t}^{T} \frac{B_{is} - \mathrm{T}_{is}}{(1+r)^{s-t}} + \sum_{s=t}^{T} \frac{L_{is}}{(1+r)^{s-t}}. \tag{11}$$

Thus, in addition to future earnings, agents have access to two resources: marketable wealth, and the net present value of Social Security cash flows. We define total wealth as the sum of these two resources:

$$\overline{W}_{it} = W_{it} + S_{it}, \tag{12}$$

where S_{it} is the net present value of Social Security cash flows:

$$S_{it} = \sum_{s=t+1}^{T} \frac{\mathbb{E}\left[\mathrm{B}_{is} - \mathrm{T}_{is}\right]}{\left(1 + r_{ts}\right)^{s-t}} \tag{13}$$

Conceptually, we show that the net present value of Social Security cash flows should be treated as wealth for two reasons.

The first reason is that Social Security and private wealth are substitutes: the more workers accumulate Social Security wealth, the less they need to save privately. By contrast, wealth and earnings tend to be complements: high earners accumulate more wealth. To illustrate the substitution between private and Social Security wealth, let's decompose the total saving rate into a private and a Social Security component:

$$\frac{\overline{W}_{it+1} - \overline{W}_{it}}{Y_{it}} = \frac{W_{it+1} - W_{it}}{Y_{it}} + \frac{S_{it+1} - S_{it}}{Y_{it}} \tag{14}$$

Figure 1: Saving rates by earnings decile

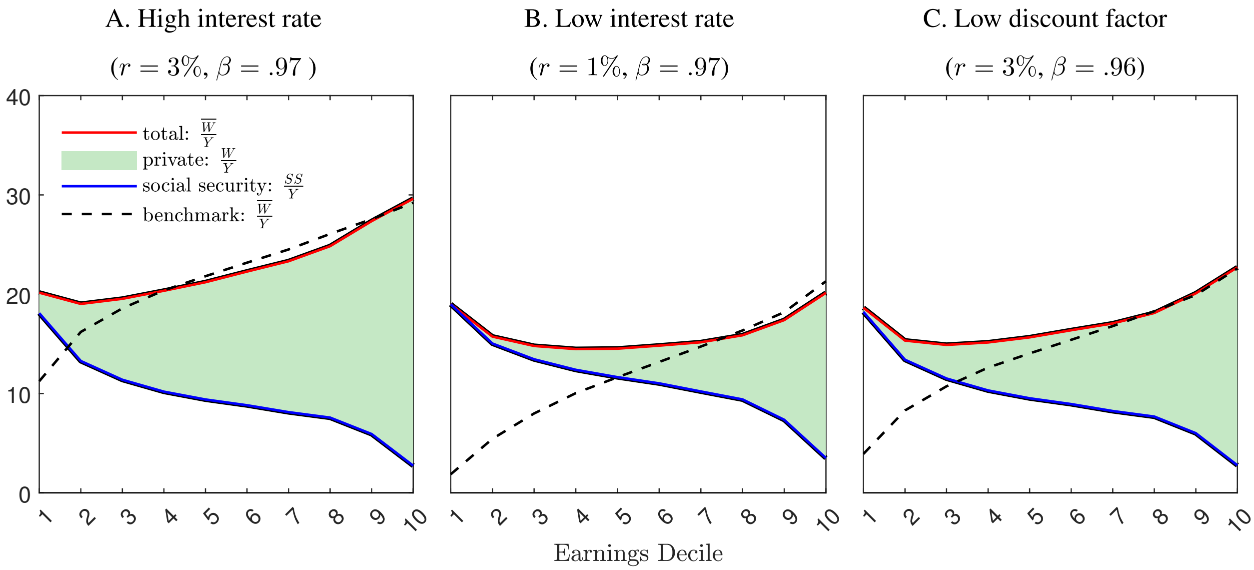

where $Y_{it} = L_{it} + \overline{W}r$ is yearly revenues. When interest rates are low, economic theory suggests that households should favor consumption, rather than saving, today. Reflecting this insight, Figure 1 decomposes the saving rates of 40-year old workers by earnings decile for different levels of interest rates r and discount factors β. As a benchmark, we also report savings rates in the absence of mandatory public savings through Social Security.

Figure 1 illustrates an attractive property of a holistic wealth concept: for unconstrained households, total saving rates are nearly identical with and without Social Security. Each dollar's worth of Social Security savings reduces private savings by the same amount as long as workers have enough liquid precautionary savings.[1]

Figure 1 shows that high earners have slightly higher total saving rates but lower Social Security rates. In our model, the relationship between total saving rates and earnings is a consequence of income uncertainty. Workers in the top deciles earn more than they expected, and thus their optimal

[1]Social Security wealth would not necessarily be more valuable if it were liquid as the program could no longer offer insurance against longevity and adverse income shocks. Importantly, many forms of marketable wealth such as retirement accounts, real-estate and private businesses are not perfectly liquid either.

consumption path for retirement has risen faster than their current level of savings—so their savings rate must rise. The opposite is true for workers in the bottom deciles.

On the other hand, Social Security savings rates are decreasing with earnings. Because payroll taxes are capped, low earners contribute a larger share of their earnings to Social Security. Moreover, because the benefit formula is progressive, low earners buy a larger annuity for each dollar of contribution. Overall, Social Security reduces the private savings rate of low earners much more than that of high earners and thus contributes to market wealth inequality.

The second reason Social Security should be treated as wealth is that, like private wealth, and unlike future labor earnings, it allows households to consume without working. Ideally, we would like to disentangle the share of Social Security wealth that households have already accrued through past work from benefits that they will earn from years in the labor force going forward. However, the nature of the Social Security formula means that the value of benefits that workers have already accrued is derivative of their future earnings trajectory: For example, note that workers early in their careers will be mislabeled as low earners (with high Social Security replacement rates) if we estimate their Social Security benefits only based on years already worked.

But in practice, the difficulty of disentangling accrued from future Social Security wealth is not problematic. As we show in Section 7.3, quantifying the value of Social Security based on the present value of expected benefits and doing so based on benefits that workers have already accrued result in substantially similar estimates. Essentially, to a first approximation, going forward, Social Security is a zero NPV investment for workers.[2]

2.3 Wealth inequality and interest rates

Saez and Zucman (2016) argue that falling saving rates explain most of the drop in the bottom 90% share, whereas Greenwald et al. (2021) show that the evolution of top wealth shares tracks

[2]If anything, the value of accrued benefits *exceeds* the net present value concept by roughly 10% (Nickerson and Burkhalter, 2019) since contributions made in the later years of a worker's career are actually negative NPV investments because the benefit formula is concave and only considers the best 35 years of earnings.

Figure 2: Top 10% shares at age 66

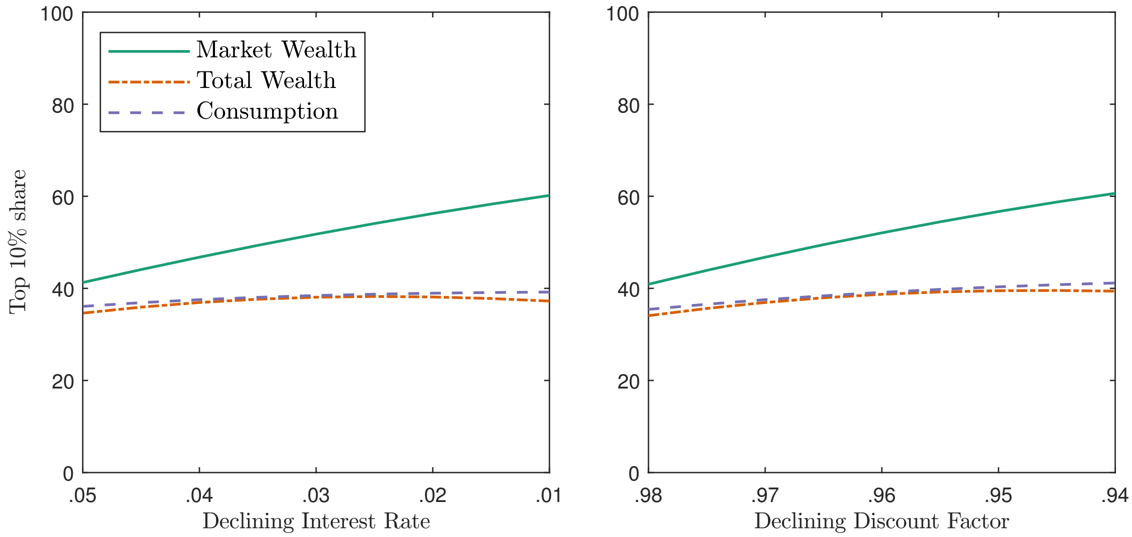

that of interest rates.

As Figure 1 shows, in the absence of Social Security, a lower interest rate or discount factor depresses wealth accumulation across the entire income distribution. However, in the presence of Social Security, the private saving rate of low and middle-class earners falls proportionally more because contributions to Social Security do not change. If anything, workers accumulate *more* Social Security wealth in the low rate environment because their contribution buy annuities with higher market value. Overall, in the low interest rate environment, private savings are much more concentrated in the upper part of the income distribution. Yet, if anything, total savings are more equally distributed.

To further illustrate the role of Social Security, Figure 2 reports three measures of inequality at retirement age when the model is simulated with interest rates ranging, from left to right, between $r = .05$ and $r = .01$ and discount factors between $\beta = .98$ and $.94$. These measures are the top 10% share of wealth (W), total wealth ($\overline{W} = W + S$), and consumption. Figure 2 shows that a decline in interest rates or patience leads to an increase in market wealth inequality: a 4 pp drop in rates or the discount factor translates into a 20 point increase in the top 10% share of market wealth among new retirees. Lower incentives to save have a disproportionate effect on the *private* savings rates of low-earners, because they contribute a larger share of their income towards Social Security. On the other hand, the top 10% share of total wealth or consumption is barely changed. Moreover, measures of wealth inequality are more distorted by the omission of Social Security in a low rate environment or when households save less.

Hence, in a simple consumption model, when interest rates fall, measures of wealth concentration that overlook Social Security suggest an increase in wealth inequality that is not representative of the allocation of resources across the population.

3 Data

We use the triennial SCF for two main purposes: (i) measuring marketable wealth shares, and (ii) estimating aggregate Social Security wealth, and determining the share of Social Security wealth going to the wealthy. The SCF is well suited to measure marketable wealth shares as it covers most asset classes and provides detailed information on liabilities. We measure marketable wealth using the net worth variable: the sum of assets net of liabilities. One caveat is that the SCF does not survey extremely wealthy households. To fill this gap, we follow Saez and Zucman (2016) by adding the Forbes 400 list to the richest 0.01%.

We compute the Social Security wealth of retirees using detailed data on retirement and survivor benefits. The ability to observe benefits at the household level allows us to observe the joint distribution of Social Security and marketable wealth among retirees. We rely on this joint distribution to assign Social Security wealth between the top and bottom of the distribution.

Yield curve data come from the Federal Reserve. These data provide estimates of the zero-coupon yield curve using off-the-run Treasury coupon securities for horizons up to 30 years. To obtain interest rates at horizons greater than 30 years, we extend this series by repeatedly applying the 29-to-30 year forward rate to the annualized spot rate at 30 years. Hence, the annualized spot rate at $30 + h$ is $r_{t,t+30+h} = \left((r_{t+29,t+30})^h (r_{t,t+30})^{30}\right)^{\frac{1}{30+h}}$. Our assumption is that this forward rate represents the long-run interest rate on nominal government claims.

We use historical inflation, wage growth, and discount rate projections from past SSA Annual Reports to calibrate and validate our valuation model. We also collect Social Security parameters such as the time series of the Social Security bend points, national wage index, maximum taxable earnings, and cost-of-living index from the SSA website.

4 Valuing Social Security

In this paper, we trace out how accounting for Social Security impacts trends in wealth concentration. To do so, we estimate the evolution of Social Security wealth by cohort. Then, we distribute this wealth between the top 10% or 1% and the rest of the population. We proceed differently for retirees and workers.

4.1 Retirees

For retirees, calculating Social Security wealth is relatively straightforward, since we observe their marketable wealth and Social Security benefits in the SCF. As there are no more taxes to be paid, Social Security wealth is

$$S_{it} = \sum_{s=t}^{T} \left(\prod_{k=t}^{s-1} (1 - m_{itk}) \right) \frac{B_{it}}{(1 + r_{t,s})^{s-t}} \frac{\mathbb{E}[P_s]}{P_t} \tag{15}$$

where nominal benefits are indexed to the consumer price index P_t. We also include survivor benefits in this calculation. Survivor benefits are paid to the surviving spouse and can represent up to 100% of the benefits of the deceased husband or wife (see details in Appendix B.4).

4.2 Workers

For working-age cohorts, we simulate the distribution of income paths of their members using the income process laid out in Equations (10.1)-(10.7). Then, we apply the Social Security benefit and tax formulas to compute average Social Security wealth by cohort, gender, and year.

Taxes As in Equation (3), we assume that workers will keep paying 10.6% of their earnings below the Social Security wage base.

Benefits For simplicity, we assume that workers retire at the cohort-specific full retirement age. We compute the AIYE as the average of the best 35 years of indexed earnings $\mathrm{L}_{it}^{\text{indexed}}$, as defined in equation (4). Yearly benefits depend on year of birth c, and the marginal replacement rate drops at two cohort-specific bend points, $b_{1,c}$ and $b_{2,c}$. Hence, benefits are concave and piece-wise linear function of AIYE:

$$B_{it} = \begin{cases} \frac{P_t}{P_{c_i+60}} \cdot .9 \cdot \mathrm{AIYE}_i & \text{if } \mathrm{AIYE}_i < b_{1,c_i} \\ \frac{P_t}{P_{c_i+60}} \left[.9 \cdot b_{1,c_i} + .32(\mathrm{AIYE}_i - b_{1,c_i}) \right] & \text{if } b_{1,c_i} \leq \mathrm{AIYE}_i < b_{2,c_i} \\ \frac{P_t}{P_{c_i+60}} \left[.9 \cdot b_{1,c_i} + .32(b_{2,c_i} - b_{1,c_i}) + .15(\mathrm{AIYE}_i - b_{2,c_i}) \right] & \text{if } b_{2,c_i} \leq \mathrm{AIYE}_i. \end{cases} \tag{16}$$

where $\frac{P_t}{P_{c_i+60}}$ is an adjustment for the increase in the consumer price index since the retiree turned 60.[3]

4.3 Aggregate Social Security wealth by cohort

For households below the minimum retirement age, we aggregate the Social Security wealth of workers at the cohort level using the SCF demographics weights ω_{it}. Specifically, we have:

$$S_{ct}^{\text{total}} = \sum_{i \in c} \omega_{it} \cdot S_{gct}^{\text{mean}}, \tag{17}$$

where S_{gct}^{mean} is the mean social security wealth by gender and cohort. We account for the estimated 20% of women and 10% of men do not contribute to Social Security, as detailed in Appendix B.5.

For respondents from 62 to 66, the simulated data and SCF overlap. For those whose benefits are reported in the SCF, we rely on these estimates. For individuals without benefits, we fill in the average simulated Social Security wealth, adjusting for the non-contributing share of the population. For individuals aged 66-69 who have not yet claimed their benefits, we backfill average benefits and wealth from the succeeding survey for respondents from 70 to 73 years of age (see

[3]In practice, workers can claim benefits as early as age 62 or as late as age 70. However, this option is relatively fairly priced as retiring earlier (later) reduces (increases) benefits in a proportion consistent with life expectancy at retirement, such that overall the total present value of benefits remains the same (Auerbach et al., 2017).

details in Appendix C). Finally, for cohorts above age 69, we simply aggregate individual-level estimates of Social Security wealth using SCF survey weights.

4.4 Calibration

Lifetime income profiles We assume $g(t)$ to be cohort and gender-specific. Guvenen, Kaplan, Song and Weidner (2018) report the average earnings of each cohort c and gender g by year from 1957 to 2013. First, we divide these time series by the wage index $L_{1,t}$ to get the average realization of L_2 of each cohort-gender group: $\mathrm{L}_{2,cgt}$. Then, we estimate $g_{cg}(t)$ by regressing $\ln(\mathrm{L}_{2,cgt})$ on a cubic polynomial of age. The data includes workers who enter the labor force from 1949-2016. For cohorts where there is insufficient labor market data to estimate $g(t)$ directly, we rely on estimates for nearby cohorts, whose earnings trajectories follow similar paths.

Social Security parameters To obtain Social Security wealth for a given year, we use actual Social Security parameters up to that year as they are stated on the SSA website. We then assume that future Social Security parameters will scale up with the wage index, which has been the case over our sample period. Hence, we assume that the Social Security wage base will remain 2.5 times the wage index ($\mathrm{SSWB}_t = 2.5 \cdot L_{1,t}$), and the bend points of the benefits formula will remain .21 and 1.25 the wage index ($b_{1,t} = 0.21 \cdot L_{1,t}$ and $b_{2,t} = 1.25 \cdot L_{1,t}$). We assume that Social Security respectively covers 90% and 80% of the male and female populations (see Appendix Figure E.4).

Macroeconomic assumptions Because they are inflation-indexed, Social Security cash flows should be discounted using the real yield curve. In our baseline specification, we use the nominal yield curve for Treasury notes with data coming from the Federal Reserve, described in Section 3. Therefore, we let cash flows grow with the consumer price index. We use inflation projection from SSA reports, as we are not aware of another source for long-term inflation projections since 1989. Wage growth projections also come from the SSA reports. We discuss alternative growth scenarios in Section 8.1.

Mortality and differences in life expectancy Survival probabilities are calibrated to the historical mortality rates by gender from 1989–2017 coming from the Human Mortality Database (HMD), which we adjust for differences in life expectancy by income. Individuals with higher earnings live longer: life expectancy for men in the top 1% by income is nearly 15 years longer than average life expectancy for the bottom 1% (Chetty et al., 2014). We adjust for these differences using data from the Health Inequality Project (HIP) by allowing survival probabilities of SCF respondents receiving Social Security retirement benefits to differ by income.[4] Our adjustment effectively makes high income retirees younger and low income retirees older, a procedure discussed in Appendix B.3.

4.5 Validation

To validate our methodology, we check (i) that the benefits predicted by our simulation match the data, (ii) that, when using the same discount rates as the SSA, we obtain similar estimates of the evolution of aggregate Social Security wealth, and (iii) that the use, due to data availability, of a nominal rather than real yield curve is not driving our results.

Matching observed benefits at retirement age In Figure 3, we compare simulated and observed benefits for retirees between ages 62 and 67. For those who did not retire at full retirement age, we use Social Security rules to determine what their full retirement age benefits would be if they had (see Appendix B.2). The simulated data track observed benefits closely.

Matching SSA estimates of aggregate Social Security wealth Every year, the SSA estimates the present value of expected benefits to current participants, net of their expected payroll taxes. Our goal is not to replicate the SSA estimates, as the SSA actuaries' assumptions regarding the level and slope of the yield curve are inaccurate. Rather than using a market-implied spot rate to discount future cash flows, the SSA projects rates based on interest rate movements in prior business cycles, which drastically understates the decline in interest rates. As reported in Appendix

[4]We proxy for the permanent income distribution using the Social Security benefits distribution because benefits are, by construction, a proxy for lifetime earnings.

Figure 3: Simulated and actual full retirement age benefits

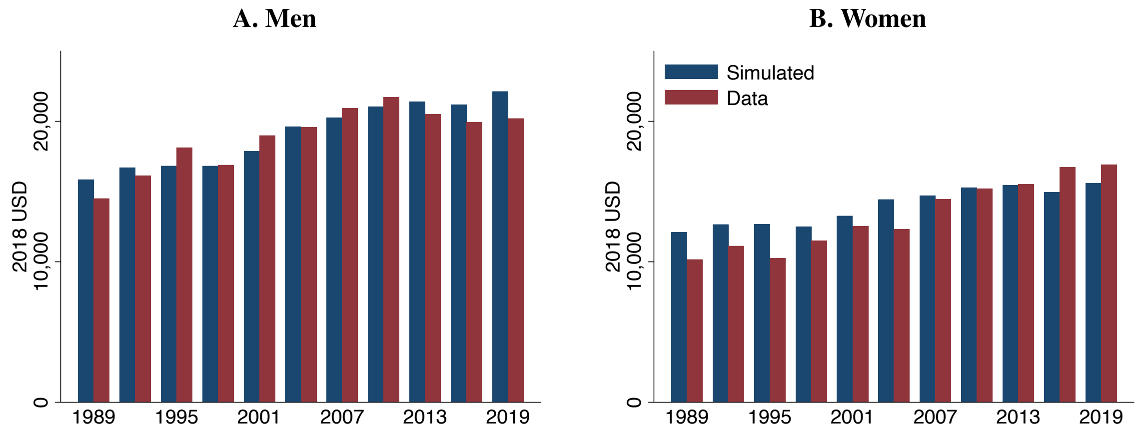

Figure E.3, the SSA discount rates fell by only 2 pp between 1989 and 2019 while the market yield curve fell by three times that amount.

Nevertheless, Figure 4 shows that, if we choose to use the SSA's discount rates, the evolution of aggregate Social Security wealth reported by the SSA tracks our estimates. This gives us confidence in our simulated estimate of workers' lifetime earnings histories, from which we derive their Social Security wealth. For comparison, we also include our estimate of aggregate Social Security wealth discounting based on the market-implied yield curve. The deviations between discounting based on SSA projections and Treasury reported rates is fairly small in the first decade of our sample, but it grew substantially in the last 15 years. In 2019, SSA-implied aggregate Social Security wealth was nearly $32 trillion, compared to over $50 trillion when using market rates.

Using the real yield curve to validate inflation forecasts Finally, because we discount future cash flows using the nominal yield curve, our findings are sensitive to inflation forecasts, which we take from SSA annual reports. To make sure that our results are not driven by these assumptions, we also discount future cash flows using the real yield curve implied by the price Treasury Inflation Protected Securities (TIPS) and assume no inflation. This exercise can only be done for the 1999-2016 period. As reported in Figure 4, this alternative methodology implies a faster increase in

Figure 4: Aggregate Social Security wealth under alternative discount rates

This figure reports estimates of the aggregate present value of Social Security. The "SSA Reports" line reports estimates by the Office of the Chief Actuary (OACT). We subtract the value of the Disability Insurance program by assuming that it represents 1.8/12.4 of the total, which is consistent with the allocation of payroll tax revenues. The "SSA Discount Rates" line reports our estimates using OACT discount rates. The "TIPS Yield curve" line reports our estimates when we assume no inflation and use the real yield curve implied by treasury inflation-protected securities. The "Risk-free valuation" line reports our estimates using the nominal market yield curve.

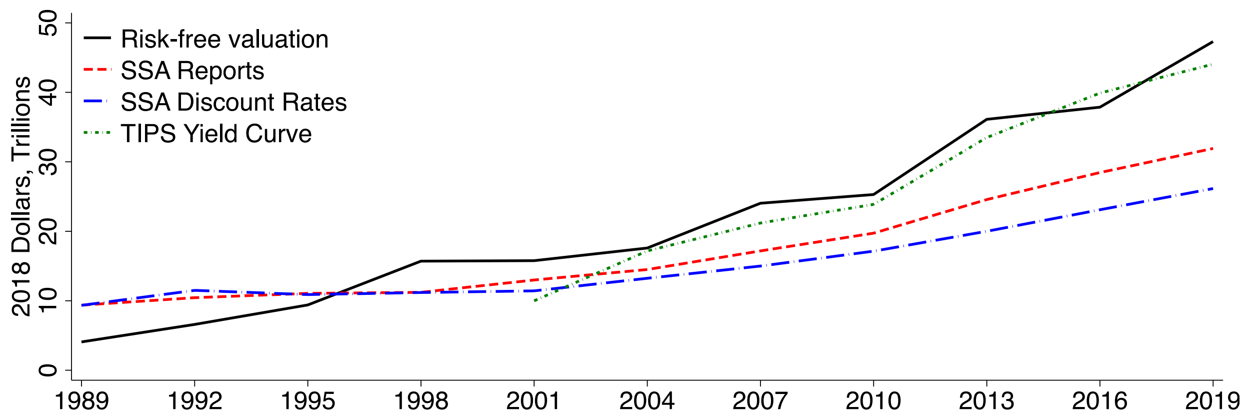

aggregate Social Security wealth than ours;[5] as such, our findings are not driven by challenges with forecasting inflation.

4.6 Assigning Social Security wealth to the top

With an aggregate value of Social Security wealth, we now turn to allocating this wealth across the distribution. Our strategy to assign each cohort's Social Security wealth to the top and bottom of the distribution depends on whether households have already claimed their benefits.

Retirees For retirees, we compute Social Security wealth at the individual level. Hence, we can precisely estimate the share of Social Security wealth that is captured by each centile of the overall marketable wealth distribution.

[5]There is an economically significant deviation between the nominal and TIPS discounted valuations in 2001. However, TIPS rates were not representative of the real risk-free rate in the early part of the sample from 1999-2003 (Fleming and Krishnan, 2004).

Figure 5: Assignment of Social Security wealth by working-age cohort in 2019

This figure illustrates how we allocate the Social Security wealth of working-age cohorts between the top 10% and bottom 90% of the wealth distribution. Following Equation (18), the Social Security wealth of cohort c going to the top 10% is:

$$S^{10\%}_{c,2019} = \phi_{2019}(\text{Share in Top }10\%_{c,2019}) \cdot S^{\text{total}}_{c,2019}$$

In Panel A, we report the share of households falling in the top 10% of the overall wealth distribution, that is Share in Top $10\%_{c,2019}$. In Panel B, we estimate the share of the Social Security wealth of young retirees (65–75) that goes to the richest $x\%$ of that group, that is the function $\phi_t(x)$. Panel C reports total social security wealth by cohort, split between the top 10% and the rest of the population.

A. Share of cohort in the Top 10%

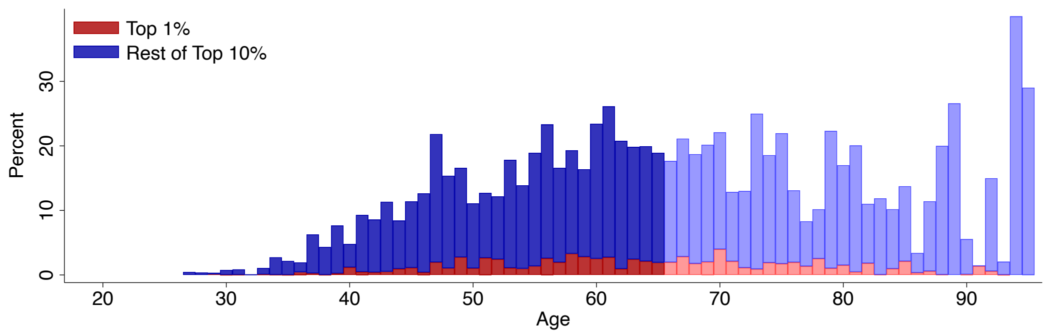

B. Wealthiest retirees' share of Social Security wealth $\phi_t(x)$

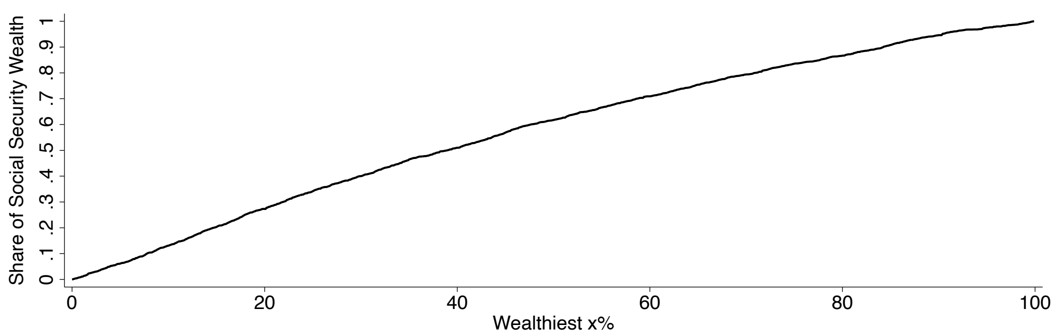

C. Social Security Wealth by Age — Risk-free Valuation

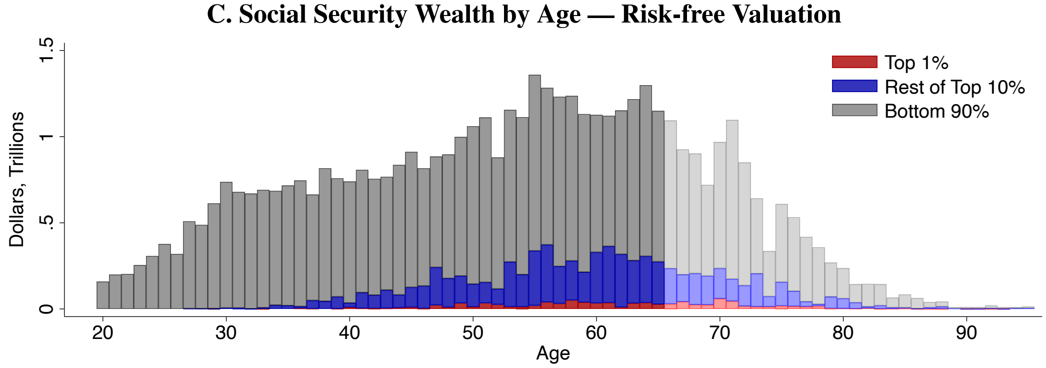

Workers For workers, we need to divide our cohort-level estimates of aggregate Social Security wealth between the top 10% or 1% and the rest of the population.We assume that the share of Social Security Wealth going to the top 10% is:

$$S_{ct}^{10\%} = \phi_t(\text{Share in Top } 10\%_{ct}) \cdot S_{ct}^{\text{total}} \tag{18}$$

where Share in Top $10\%_{ct}$ is the share of cohort c in the overall top 10% of the market wealth distribution in year t. The function $\phi_t(x)$ represents the share of Social Security wealth held by the wealthiest $x\%$ of young retirees in the same year. Both Share in Top $10\%_{ct}$ and $\phi_t(.)$ are readily observable in the SCF.

A numerical example illustrates our approach.

1. In 2019, 60 year-olds had \$1.2 trillion in Social Security wealth.

2. As Panel A of Figure 5 shows, in 2019, 23% of 60-year old households were in the top 10% of the overall marketable wealth distribution.

3. Panel B shows that, within the population of young retirees (65-75), the wealthiest 23% held 30% of the Social Security wealth of this age group.

4. Hence, we allocate \$360 billion (\$1.2 trillion x .30) of 60 year-olds' Social Security wealth to the top 10% in 2019.

By repeating this exercise for all working-age cohorts, we determine the overall amount of simulated Social Security wealth owned by the top 10% and bottom 90% in 2019. We use the same procedure for other survey years and the top 1%.

In this exercise, our key assumption is that the share of Social Security wealth that accrues to different centiles of the market wealth distribution is constant across ages. There are several reasons why this assumption is reasonable for our exercise. First, this assumption is most tenuous for the youngest workers. However, as illustrated by Panel A of Figure 5, the implications of any potential mis-allocation of Social Security wealth for these cohorts are quantitatively irrelevant to our exercise, because their chances of being in the top 10% of the overall population are negligible. Relatedly, the Social Security wealth of current workers is concentrated among those approaching retirement, who are nearly finished paying into Social Security and have yet to claim their benefits. As illustrated by Panel C of Figure 5, 79% of the Social Security wealth of the top 10% goes to households above age 55 and the share going to those below 45 is close to zero. For workers above 55, relying on the relationship between marketable wealth and Social Security wealth observed for retirees is sensible.

If anything, our assumption overstates the share of Social Security wealth that accrues to the top 10% because the value of Social Security is low and perhaps even negative for the wealthiest individuals in younger cohorts. Social Security is progressive, and so it offers higher replacement rates to low earners. Though high earners who recently retired have more Social Security wealth than low earners, each dollar has been bought at a higher price. At retirement, this price is sunk and does not change their Social Security wealth. However, for younger cohorts, a large fraction of this cost remains to be paid, which reduces the net present value of Social Security disproportionately for high earners.

4.7 Baseline top wealth shares

Figure 6 reports the levels and trends of top wealth shares with and without Social Security wealth.[6] We define top wealth shares based on the top 10% and top 1% of the population by measures of marketable wealth. This allows for comparison of how previously documented inequality trends are impacted by the inclusion of Social Security.

Panel A focuses on the top 10%. The top 10% share of market wealth grew by 10 pp between 1989–2019. Once Social Security wealth is included, this trend is reversed. Rather than rising, the top 10% wealth share falls by 5.4 pp over this period. Panel B shows the impact of Social Security wealth on top 1% wealth share. When Social Security wealth is excluded, the top 1% share has

[6] We find similar declines in top wealth shares when other data on marketable wealth are used, as discussed in Appendix Section 8.5.

Figure 6: Top 10% and Top 1% Wealth Shares with and without Social Security

This figure reports the evolution of the top 10% and 1% wealth shares with and without Social Security wealth. In the risk-free valuation, cash flows are discounted using the yield curves implied by the price of government bonds. In the risk-adjusted valuation, we adjust discount rates to account for the long-run cointegration between the labor and stock markets, as detailed in Section 5.1.

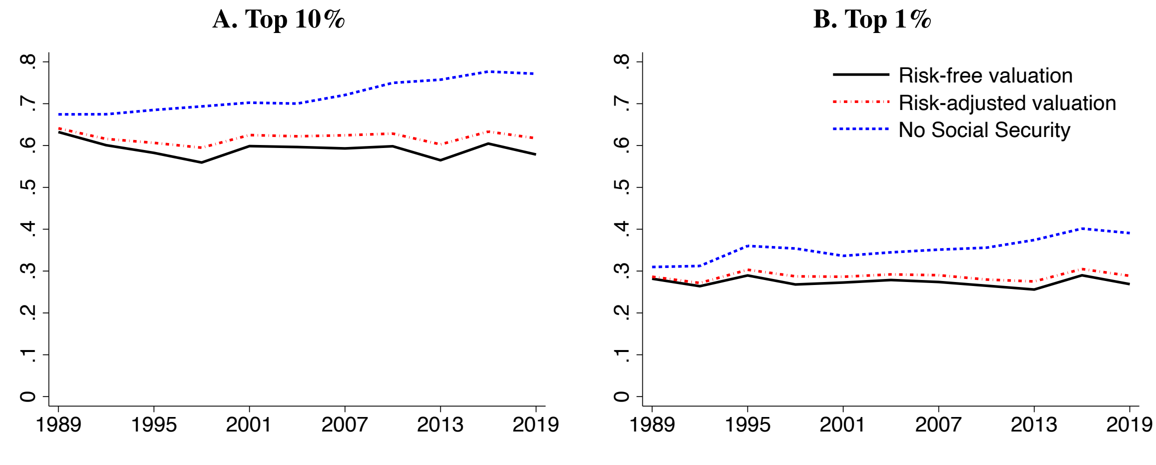

grown by 10 pp over our sample period. Once it is included, the top 1% share has risen by 1.3 pp.

5 Accounting for macroeconomic risk

The rate of return of pay-as-you go systems is the sum of the growth rates of the population and per capita earnings (Samuelson, 1958). For U.S. Social Security, the relationship between returns on contributions and the long-run growth in earnings is explicitly achieved through wage-indexation. Wage-indexation exposes Social Security participants to long-run macroeconomic risk, and discount rates should reflect this systematic risk.

Social Security cash flows perfectly scale up with the national wage index. Since 1980, the Social Security wage base and bend points have been growing at the same rate as earnings. In Section 4.2, we show that tax payments are proportional to the wage index, whereas benefits are proportional to the wage index the year a worker turns 60. Therefore, a diversified investor would discount these cash flow using the expected return on an asset delivering a single wage-indexed coupon with the same years of indexation and payment. In this Section, we determine the expected return for such a security.

5.1 Market beta of Social Security cash flows

At what rate should we discount a cash flow that is proportional to the average level of earnings $L_{1,t+n}$ in n years? To answer this question, we assume that the stock and labor markets are cointegrated as documented in Benzoni et al. (2007) and as would be expected if the shares of labor and profits are stable over long periods. Specifically, we assume that the log of L_1 evolves as follows:

$$dl_{1,t} = \left((\phi - \kappa) y_t + \mu - \delta - \frac{\sigma_l^2}{2} \right) dt + \sigma_l dz_{1,t}, \tag{19}$$

where $\mu - \delta$ determines the unconditional log aggregate growth rate of earnings and σ_l its volatility. Log stock market gains follow:

$$ds_t = \left(\mu + \phi y_t - \frac{\sigma_s^2}{2} \right) dt + \sigma_s dz_{2,t}, \tag{20}$$

where μ and σ_s represent expected stock market log returns and their volatility. The state variable y_t keeps track of whether the labor market performed better or worse than the stock market relative to expectations. Specifically, y_t evolves as follows:

$$dy_t = -\kappa y_t + \sigma_l dz_{1,t} - \sigma_s dz_{2,t}, \tag{21}$$

where κ determines the strength of the cointegration. If the two markets are cointegrated, y_t should mean revert to zero. Mean reversion takes two forms. If stock markets gains are caused by higher long-run economic growth, wages will catch up. If stock market returns have nothing to do with future economic growth, we should expect them to mean revert. The parameter ϕ controls the fraction of the mean reversion in y_t caused by mean reversion in stock market returns.

In Appendix D, we show that the market beta of a security delivering a single coupon proportional to $L_{1,t+n}$ is:

$$\beta_t^{L_1,n} = \left(1 - \frac{\phi}{\kappa} \right) \left(1 - e^{-\kappa n} \right) \ , \tag{22}$$

and we demonstrate that, under the no-arbitrage condition, the expected return on this security is:

$$\mathrm{E}_t \left[r_t^{L_1,n} \right] = \beta_t^{L_1,n} (\mu - r) + r \tag{23}$$

where r is the risk-free rate. Note that, assuming policy risk away, any Social Security payment proportional to L_{t+n} would deliver the same expected return if it were publicly traded, as all other sources of risk are purely idiosyncratic.

Our empirical exercise is in discrete time, so we approximate our results by assuming that the discount factor for a cash flow proportional to $L_{1,n}$ paid in year k is:

$$\chi_{t,n,k} \approx \left[\prod_{s=t}^{n} \left(1 + \beta_s^{L_{1},n} (\mu - r) + r_{ts}\right) \prod_{s=n+1}^{k} (1 + r_{ts}) \right]^{-1}, \tag{24}$$

and the risk-adjusted present value of Social Security is:

$$\text{Adj. S}_{it} = \sum_{s=t+1}^{T} \left(\prod_{k=t}^{s-1} (1 - m_{itk}) \right) \left(\mathbb{E}\left[\text{B}_{it}\right] \cdot \chi_{t,c_i+60,s} - \mathbb{E}\left[\text{T}_{it}\right] \cdot \chi_{t,s,s} \right) \tag{25}$$

where real benefits are indexed on L_1 in the year in which the worker turns 60.

We calibrate the model as in Benzoni et al. (2007) who estimate $\kappa = .16$ and $\phi = .08$ using U.S. data from 1929 to 2004. This implies a market beta of 0.5 for for the most distant indexed cash flows. The equity premium is $\mu - r = .06$.

5.2 Risk-adjusted results

Panel A of Figure 7 reports aggregate Social Security wealth with and without adjusting for systematic labor market risk. Panel B shows that the adjustment is larger for young workers: it cuts the Social Security wealth of a 25-year old's benefits by 60%.

Once macroeconomic risk associated with Social Security cashflows is factored in, Figure 6 shows that the share of the top 10% decreased by 2.4 pp and that the top 1% share has increased by 0.2 pp. This finding differs from our baseline risk-free specification because Social Security wealth is smaller, and therefore plays a lesser role in the evolution of wealth inequality. The risk-adjusted results primarily decrease Social Security wealth for younger workers, who are rarely in the top 10%. Consequently, our risk adjustment decreases the wealth of the bottom 90%, with only a small impact on the Social Security wealth of the top 10%. Regardless, top wealth shares remain substantially attenuated relative to prior work.

Figure 7: Risk-adjusted valuation

Panel A presents aggregate Social Security wealth in 2018 dollars. Panel B presents average Social Security wealth by age in 2019.

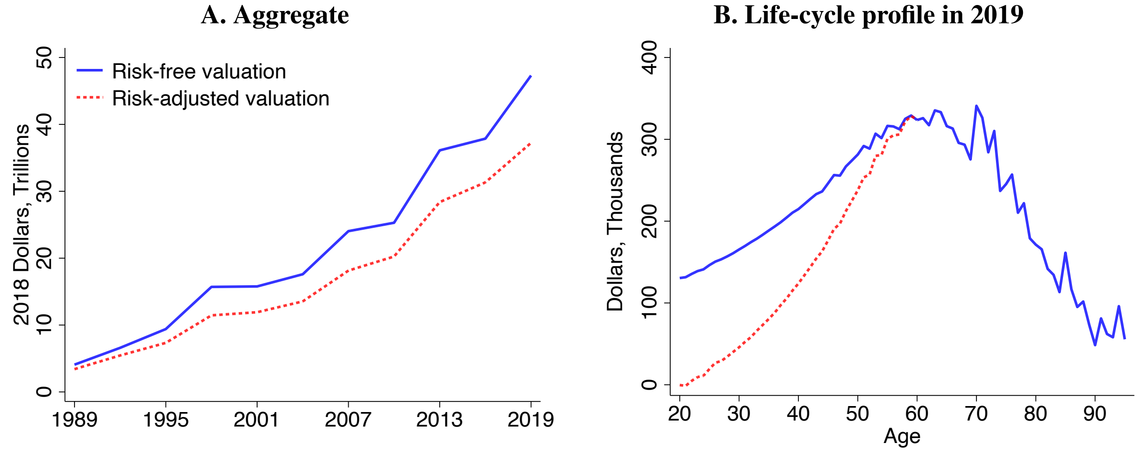

6 Role of interest rate environment

The decline in interest rates has important implications for wealth inequality. Over the last 30 years, long-duration assets have dramatically outperformed short-duration assets (Binsbergen, 2020). Because rich households invest in longer-duration assets such as stocks and private businesses, the decline in interest rates can explain most of the increase in market wealth inequality (Greenwald et al., 2021). However, the focus on market wealth overlooks the largest long-duration investment of most households: their Social Security contributions.

6.1 Social Security as a leveraged exposure to duration

For working-age households, Social Security benefits are disbursed years into the future, while taxes are paid into the program today. Essentially, the exposure to rates through future tax payments can be replicated by selling short- and medium-term bonds, and the exposure through benefits can be replicated by buying long-term bonds. In Table 1, we report the weights of benefits and taxes in the overall net present value of Social Security and the change in the value of each component since 1989. Because benefits (the long position) have a longer duration, when rates fall, their present

Table 1: Impact of leverage and interest rates on Social Security wealth

This tables decomposes the increases in Social Security wealth between 1989 and 2019. Columns (a) and (b) report the weights of future benefits and taxes in the net present value of Social Security in 1989. Columns (c) and (d) report the increase in the present values of benefits and taxes. The last column reports the percentage change in Social Security wealth.

	Share of Social Security wealth in 1989		Change since 1989		
	Benefits (a)	Taxes (b)	Benefits (c)	Taxes (d)	NPV (a)·(c)+(b)·(d)
Bottom 99%	248%	-148%	491%	136%	1015%
Top 1%	123%	-23%	436%	142%	504%
Entire population	243%	-143%	490%	136%	994%

value rises faster than that of taxes (the short position). The result is a rapid increase in the net present value of Social Security.

This increase is especially important for the bottom of the wealth distribution for two reasons. First, Social Security represents a larger share of their total wealth. Second, as the last column of Table 1 shows, the value of Social Security has increased 511 pp more for the bottom 99% than for those in the top 1%.

6.2 Why low rates make Social Security more valuable

Even if it is primarily driven by interest rate changes, the evolution of Social Security wealth matters to our understanding of inequality trends. Consider a household that is 20 years away from retirement and seeks to save enough to finance one dollar of consumption for 20 retirement years. Assuming an interest rate of 5%, as in 1989, this household needs to save $0.38 per year over the next two decades. If the interest rate is 0.8%, as in 2019, this household needs to save $0.85 annually. Now, if the rate of returns on Social Security contributions has been constant at 2.5%, which it has since 1989, the same can be achieved by contributing approximately $0.61 every year to Social Security. So, in effect, from this household's point of view, one dollar of Social Security contribution was equivalent to $0.62 ($0.38/$0.61) of private saving in 1989, but is equivalent to \$1.39 (\$0.85/\$0.61) in 2019. Said another way, the future purchasing power of \$1 of Social Security contributions corresponds to more private savings when rates are low. Discounting based on market rates is important as it captures the fact that the attractiveness of Social Security depends on its rate of returns relative to comparable investment opportunities on the private market.[7]

Incidentally, whether the increase in Social Security wealth corresponds to a real improvement in well-being is not key to our choice of an appropriate discount rate. This same question can be asked of marketable assets whose value has increased faster than the cashflows they generate (Greenwald et al., 2021). To compute shares correctly, one dollar of Social Security wealth must afford the same consumption as one dollar of marketable wealth in any given year. To do this, we need to discount cash flows using the prevailing cost of exchanging current and future dollars—the contemporary yield curve.

To grasp the importance of consistency, consider a household who just retired in 2019 and bought an annuity that exactly matches their Social Security benefits. It is difficult to understand why we would value this annuity using 1989 market prices. It is equally difficult to see why Social Security benefits would not have the same value as this annuity. If the annuity—and other marketable assets—are valued using contemporary prices, the same should be done for Social Security benefits.

Similarly, if a household has mortgage payments representing 10.6% of their earnings, then from a consumption standpoint, these cashflows have the same implications as Social Security taxes. But they correspond to a smaller mortgage in 1989, when rates were high, than in 2019. Therefore, they register as a greater liability in the computation of marketable net wealth in 2019 than in 1989. This mechanical upsurge in debt contributes to the rise in marketable wealth inequality documented by Saez and Zucman (2016). It is inconsistent that an increase in the present value of the mortgage payments would be accounted for in wealth inequality estimates, but not an

[7] In work following ours, Sabelhaus and Volz (2020) instead apply a constant discount rate to Social Security cashflows. This is a mistake because it ignores the effect of interest rates on asset prices, one of the main causes of rising marketable wealth inequality.

equivalent increase in the value of Social Security taxes.

7 Discussion

7.1 Factors contributing to Social Security's growth

Table 2 lays out the several contributors to Social Security's growth. These include changes in demographics (Social Security wealth is highest for those nearing retirement, who are a larger share of the population today), increasing life expectancy (average life expectancy increased by 3.5 years since 1989), and the expansion of the program (the share of earnings subject to Social Security taxes increased from 1.25 times average earnings to 2.5 times), as well as the interest rate environment. But by far the largest contributor is changes in the yield curve previously discussed, which drives 46% of Social Security's growth (51% with risk-free valuation).

7.2 Shifts in the composition of wealth

Figure 8 reports how total wealth is distributed by age and between the top 10% and the rest of the population. The overall share of the top 10% has not changed much between 1989 and 2019, nor has its composition. On the other hand, the composition of the wealth of the bottom 90% has changed dramatically. In 1989, Social Security only represented 17.2% of the total wealth of the bottom 90%. Because the rate of return on Social Security contributions was lower than risk-adjusted discount rates, Social Security wealth was even negative for households below age 40. In 2019, Social Security represents 56.9% of the wealth of the bottom 90%. The constituents of wealth held by the bottom and top of the distribution have diverged, making clear why a focus on marketable wealth inequality alone is misleading.

7.3 Comparing Social Security and private wealth

Some suggest Social Security should be excluded from wealth concentration estimates based on a few arguments: first, that Social Security wealth is uncertain, without a readily available market value; second, that Social Security benefits cannot be passed down to heirs like private wealth; and third, that Social Security wealth is illiquid and cannot be used to absorb shocks today

Table 2: Decomposing the increase in Social Security wealth

This table lays out the relative importance of changes in interest rates, the aging of the population, life expectancy, the scope of the Social Security retirement program, and the size of the population to the growth of aggregate Social Security wealth. The first row is calculated as the difference between log per capita Social Security wealth in 2019 and log per capita Social Security wealth in 2019 under the 1989 yield curve. The second row is computed by subtracting log per capita Social Security wealth in 2019 under the 1989 yield curve from Social Security wealth in 2019 under the 1989 age distribution and yield curve. The third row is computed by subtracting log per capita Social Security wealth in 2019 under the 1989 age distribution and yield curve from log per capita Social Security wealth in 2019 under the 1989 age distribution, yield curve, and survival probabilities. The fourth row is computed by subtracting log per capita Social Security wealth in 2019 under the 1989 age distribution, yield curve, and survival probabilities from log per capita Social Security wealth in 1989. The total log per capita wealth change is given by $\log(SSW^{2019}) - \log(SSW^{1989})$ where both terms are calculated under the 2019 and 1989 populations, life expectancies, benefit policies, and yield curves, respectively.

	Valuation method	
	Risk-free	Risk-adjusted
Change in yield curve	1.239	1.099
Shift in age distribution	0.290	0.348
Life expectancy	0.262	0.254
Social Security expansion & other	0.338	0.369
Log total per capita	2.128	2.069
Population growth	0.323	0.323
Log total	2.451	2.392

(Zucman, 2019).

None of these arguments are compelling. First, many sources of wealth included in existing estimates, for example pension wealth, are also illiquid. It is true that, relative to retirement accounts, there is more uncertainty in Social Security's value given policy risk (which we address in Section 8.1). But a significant contributor to rising top wealth shares—private business wealth—is similarly illiquid, and of much more uncertain value than Social Security (Bhandari et al., 2020). Further, unless beneficiaries die prematurely, retirement benefits not used to finance consumption in retirement are bequestable. Finally, the illiquidity of Social Security is in and of itself a policy choice so that the program can provide longevity insurance to retirees and guarantee a minimum level of wealth to those who may otherwise save too little. This means it can be relaxed (Catherine

Figure 8: Total Wealth Distribution by Age — Risk-adjusted valuation

This figure plots the shares of total wealth by age group for Social Security and non-Social Security wealth for 2019 and 1989 using the risk-adjusted valuation method.

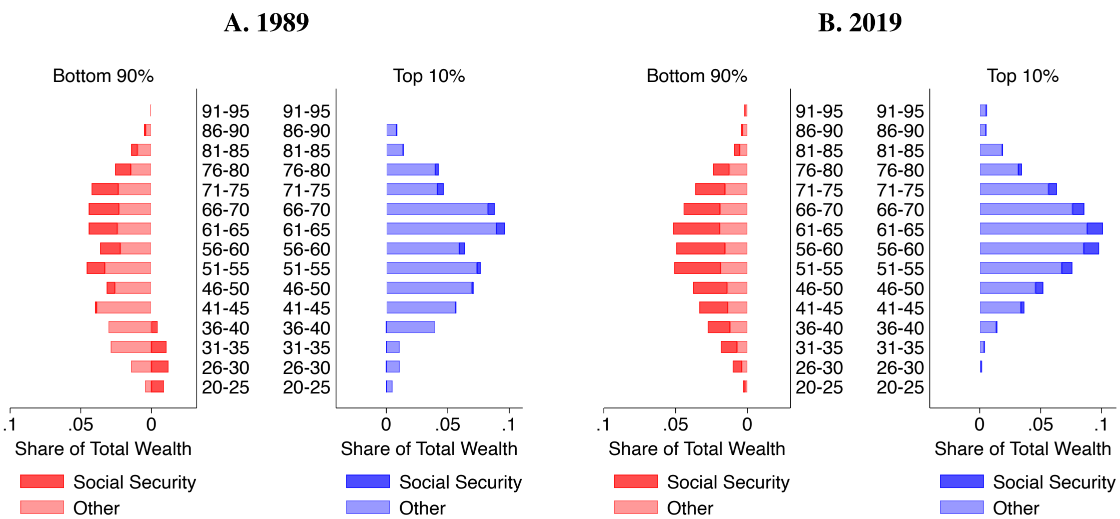

et al., 2020). But this choice does not detract from the fact that Social Security is the primary source of income for all but the very wealthiest retirees, and so is relevant to our understanding of inequality.

A significant share of Social Security wealth accrues to households who are not liquidity constrained. Households above age 50 own 79% of Social Security wealth. As Figure E.6 shows, of those 50 and older, 64% have more than $10,000 in accessible wealth and 44% have more than $100,000.

8 Robustness

We next consider the extent to which our baseline results are sensitive to alternative assumptions that impact our estimates of aggregate Social Security wealth, including policy risk that beneficiaries will not receive all promised benefits or that taxes will rise to replenish a depleted trust fund; weak economic growth; and differences in mortality between the rich and the poor. Table 3 presents results using alternative assumptions, which we discuss in turn below.

Table 3: Robustness checks

Panel A reports our baseline results. First, we report top shares of marketable wealth in the SCF+Forbes 400. We then report top wealth shares including our risk-free and risk-adjusted valuations of Social Security. In Panel B, we address the projected funding gap by cutting Social Security benefits or increasing taxes. We calibrate our wage growth assumptions and benefits cuts/tax increases based on the baseline ("Intermediate cost") and pessimistic scenarios ("High cost") used by the SSA. Under the high cost assumptions, the trust fund is depleted earlier and wages grow less than under the intermediate cost assumptions. Panel C shows additional robustness tests. First, we assume that expected wage growth declined linearly from 1% in 1989 to 0% in 2019. Second, we report three specifications which adds 1, 2 or 3 percentage points to our discount rates to reflect a hypothetical liquidity premium. All specifications in Panels B and C use the risk-adjusted valuation method.

	Share of Top 10%			**Share of Top 1%**		
	1989	2019	Change	1989	2019	Change
Panel A: Baseline results						
Marketable wealth	67.5%	77.2%	9.7%	31.0%	39.1%	8.1%
Risk-free valuation	63.2	57.9	-5.4	28.2	26.9	-1.3
Risk-adjusted valuation	64.1	61.7	-2.4	28.6	28.9	0.2
Panel B: Funding gap						
Benefit cut (Intermediate Cost)	64.1	64.0	-0.1	28.7	30.2	1.6
Benefit cut (High Cost)	64.8	67.0	2.3	29.0	32.0	3.0
Tax hike (Intermediate Cost)	64.1	62.1	-2.0	28.6	29.1	0.4
Tax hike (High Cost)	64.3	63.7	-0.6	28.8	29.9	1.1
Panel C: Robustness						
Declining wage growth	64.3	63.6	-0.7	28.8	29.9	1.1
Liquidity Premium (1%)	64.8	64.1	-0.6	29.0	30.3	1.3
Liquidity Premium (2%)	65.6	66.9	1.3	29.5	31.9	2.4
Liquidity Premium (3%)	66.2	69.0	2.7	29.9	33.3	3.3

Our overall conclusion—that the inclusion of Social Security substantially attenuates the growth in top wealth shares—is not sensitive to the specification chosen. The top 10% and 1% shares of marketable wealth (excluding Social Security) rose by 9.7 and 8.1 pp respectively between 1989–2016. Once Social Security is included, using our most conservative set of assumptions, the top 10% and 1% shares grow by only a small fraction of that over this horizon.

8.1 Accounting for Social Security policy risk

An important caveat to our baseline calculations is the imminent depletion of the Social Security trust fund: within the next 15 years, absent entitlement reform, the SSA will not be able to meet their full obligations to beneficiaries. To ascertain the impact of policy risk on our results, we modify our estimates of Social Security wealth by directly adjusting the cashflows that beneficiaries will receive or the taxes they will pay. Even under the most conservative assumptions—that beneficiaries will receive only benefits that are payable at current tax rates (eventually cutting benefits by up to 40%), or that taxes will rise for all but the top of the wealth distribution— our conclusion regarding the substantial impact Social Security has on estimates of wealth inequality is unchanged.

Balancing the budget by cutting benefits The SSA provides benchmark estimates of the extent to which the trust fund's bankruptcy will impair its obligations under three scenarios: low cost, intermediate, and high cost. Appendix Figure E.5 reports the proportion of payable benefits under each of the SSA's 1989 and 2019 cost scenarios. We assume that benefits will decrease across the board to the payable amounts reported by the SSA in each scenario, despite potential political pressure for more progressive entitlement reform.

To understand the impact of insolvency risk on our estimates, we collect annual data from the SSA on the year that the trust fund is projected to run out, the total revenue generated from Social Security payroll taxes, and the total obligations to beneficiaries. Once the Social Security fund is extinguished (estimated to be between 2030-2035), benefits paid in a year must be less than or equal to total tax revenue going forward.

Assuming maximal cuts to expected Social Security benefits decreases the bottom 99% wealth share by 3.2 pp, wiping out a quarter of Social Security's impact. But as Figure 9 shows, top wealth shares are still significantly attenuated. This is for two reasons. First, for people close to retirement, the impact of the fund's depletion is small, since benefits will pay out as normal for the first 10-15 years. Second, even for cohorts impacted, 60% of expected Social Security benefits

Figure 9: Top 10% and Top 1% wealth shares — Funding gap adjustment

This figure presents top 10% and 1% wealth shares under four, risk-adjusted specifications. The "Low cost" specification refers to the SSA's high economic growth scenario in which benefits are fully paid. In the "Intermediate cost" and "High cost" specifications, benefits are cut to match expected tax revenues under the baseline and worst-case economic growth scenarios. All specifications is the risk-adjusted valuation method.

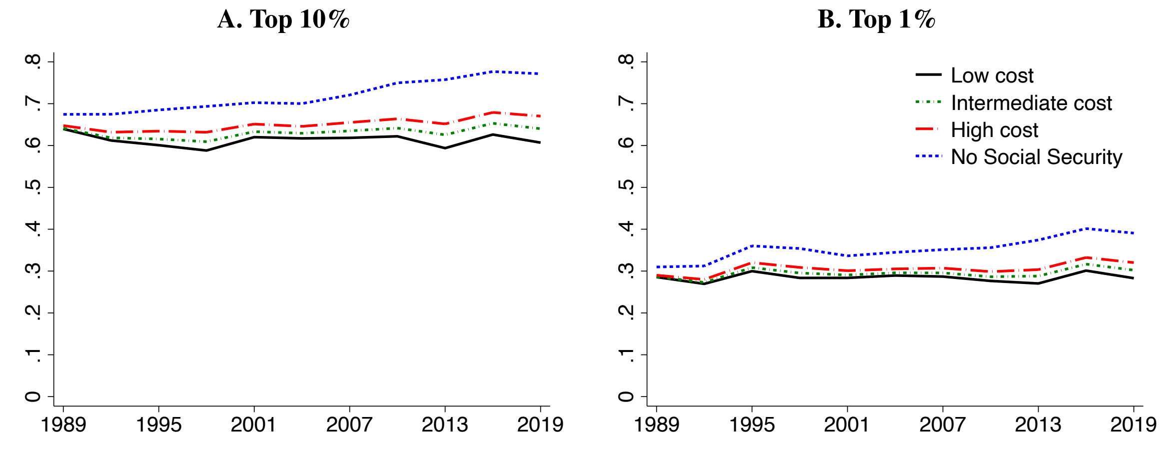

represents a sizable sum relative to their marketable wealth.

Balancing the budget by raising taxes Alternatively, taxes could be raised to avoid cutting benefits. To assess this possibility, we adopt the most conservative assumption from the perspective of our baseline results: that the additional tax burden will be borne entirely by the bottom 90%, or bottom 99%. Nonetheless, the top 10% share still declines by 0.6 pp; the top 1% share rises slightly, by 1.1 pp. Interestingly, raising taxes has less of an impact on aggregate Social Security wealth than cutting benefits. This is because raising taxes pushes a greater portion of the funding gap to future generations.

8.2 Adjusting for the illiquidity of Social Security

While excluding Social Security wealth based on its illiquidity is misguided for the reasons described in Section 7.3, one may be interested in how the results would change if discount rates account for an illiquidity premium. Who should be covered by this illiquidity premium is not clear since a significant share of Social Security wealth goes to households with ample liquid wealth.

Nevertheless, we apply liquidity discounts of a 1%, 2%, and 3% to Social Security wealth for all households and report the effect on our main results in Figure 10.

Figure 10: Liquidity premium adjustment

This figure reports the evolution of the top 10% and 1% wealth shares, average Social Security wealth by age in 2019, and the aggregate value of Social Security wealth when we add 1, 2 or 3 percentage points to our discount rates to reflect a hypothetical liquidity premium. All specifications is the risk-adjusted valuation method.

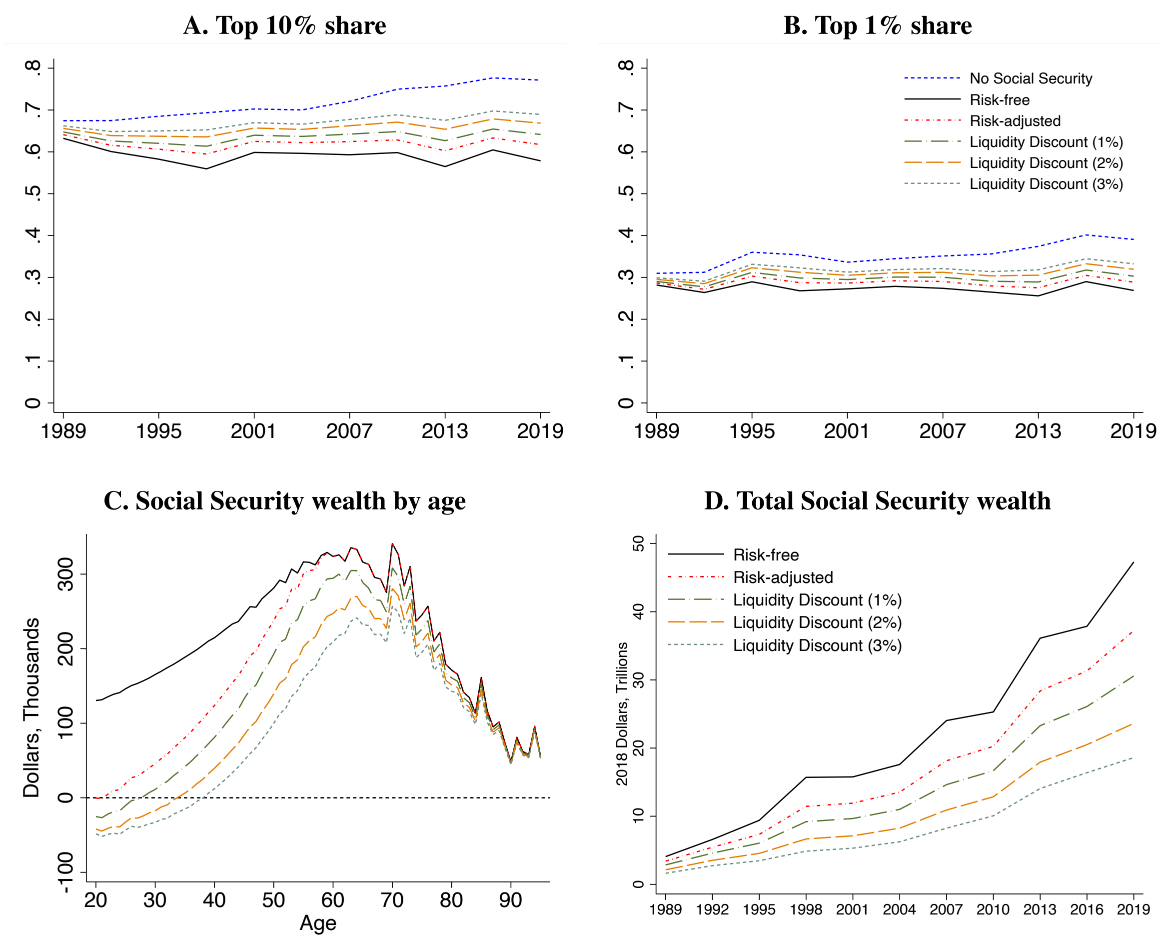

Panels C and D show that applying a 3% illiquidity premium reduces aggregate Social Security wealth by half, and Social Security wealth even becomes negative for young households. Yet, even under such a drastic adjustment, the rise in top wealth shares remains substantially attenuated when we include Social Security, as shown in Panels A and B. The top 10% and top 1% wealth shares only rise by 2.7 and 3.3 pp respectively, instead of 9.7 and 8.1 pp when Social Security is not included. These computations do not take into account that, for such an exercise to make sense, some important components of market wealth, such as the private business wealth of wealthy entrepreneurs, would also need to be adjusted for their illiquidity.

8.3 Decline in productivity growth

The decline in interest rates could be symptomatic of lower future long-run economic growth, which reduces the value of wage-indexed Social Security benefits. Our baseline estimates already assume a decline in the growth rate of wages: we rely on assumptions from SSA reports, which, as of 2019, assumed a 1.2% long-term annual wage growth rate, down from 1.7% in 1989. When we consider a more pessimistic scenario in which the real growth rate of wages declines linearly from 1% to 0% between 1989 and 2019, our main result is qualitatively unchanged: the top 10% decreased by 0.7 pp whereas the top 1% share increased by 1.1 pp.

Figure 11: Top 10% and Top 1% wealth shares — Accrued benefits

This figure shows the top 10% and top 1% wealth shares with and without the risk-adjusted value of accrued Social Security benefits.

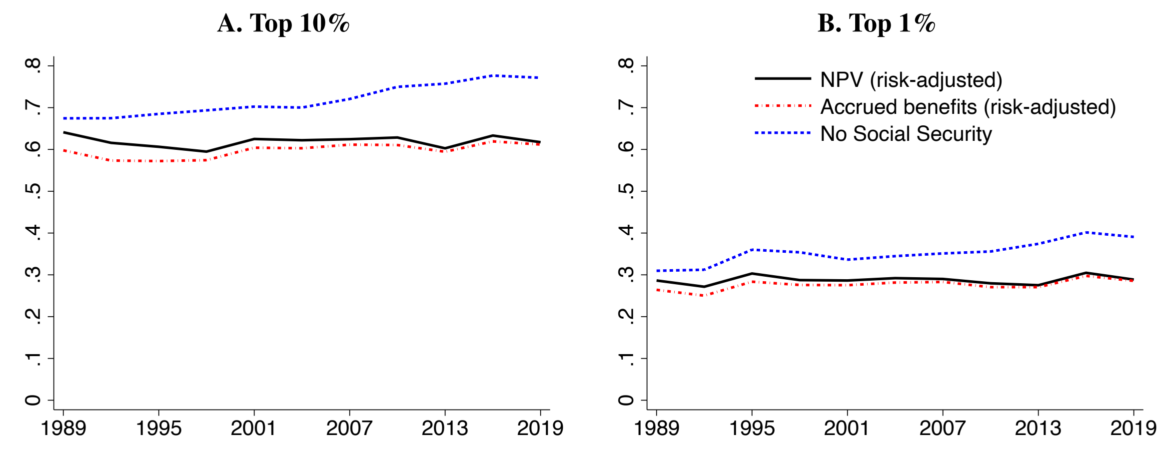

8.4 Restricting Social Security wealth to accrued benefits

An alternative definition of Social Security wealth is the value of expected benefits based on past contributions only. As shown in Figure 11, our findings are similar when we use this definition. This should be expected if the present value of future taxes are close to the corresponding increase in benefits.

8.5 Adjusting previous studies on wealth inequality

Previous studies compute top wealth shares using other datasets than the SCF . In Figure 12, we adjust these studies to include our estimates of the Social Security wealth of the top 1% and bottom 99%. Our main results remain unchanged.

Figure 12: Top 1% from previous studies

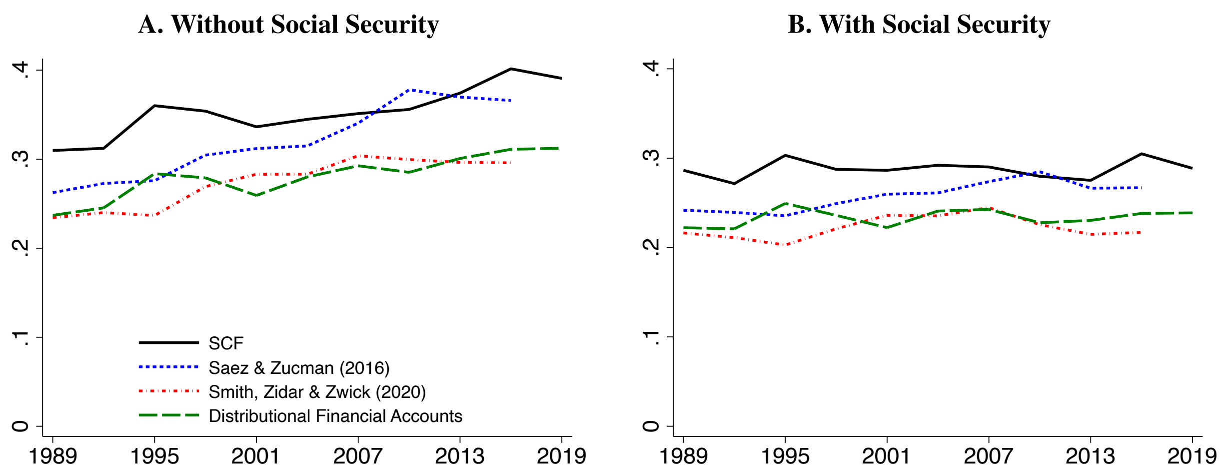

9 Conclusion

Prior studies find large increases in U.S. wealth inequality over the last three decades based on measures of wealth concentration that exclude Social Security. We find that, when Social Security is incorporated into inequality estimates, top wealth shares have not increased since 1989. Our top wealth estimates may still be overstated because we exclude programs like disability insurance and Medicare, which accrue disproportionately to the bottom of the wealth distribution. Overall, our paper shows that public transfer programs like Social Security make the U.S. economy more progressive, and it is important for inequality estimates to reflect this. Much more work is needed to arrive at a fuller understanding of wealth concentration in America.

References

Attanasio, Orazio P. and Agar Brugiavini, "Social Security and Households' Saving," *The Quarterly Journal of Economics*, 2003, *118* (3), 1075–1119.

— **and Susann Rohwedder**, "Pension Wealth and Household Saving: Evidence from Pension Reforms in the United Kingdom," *American Economic Review*, 2003, *93* (5), 1499–1521.

Auclert, Adrien, "Monetary Policy and the Redistribution Channel," *American Economic Review*, 2019, *109* (6).

Auerbach, Alan J., Kerwin K. Charles, Courtney C. Coile, William Gale, Dana Goldman, Los Angeles, Ronald Lee, Charles M. Lucas, Peter R. Orszag, Louise M. Sheiner, Bryan Tysinger, David N. Weil, Justin Wolfers, and Rebeca Wong, "How the Growing Gap in Life Expectancy May Affect Retirement Benefits and Reforms," *The Geneva Papers on Risk and Insurance — Issues and Practice*, 2017, *42* (3), 475–499.

— **, Laurence J. Kotlikoff, and Darryl Koehler**, "U.S. Inequality and Fiscal Progressivity — An Intragenerational Accounting," 2019. Working Paper.

Auten, Gerald and David Splinter, "Income Inequality in the United States: Using Tax Data to Measure Long-term Trends," 2019. Working Paper.

Benzoni, Luca, Pierre Collin-Dufresne, and Robert S. Goldstein, "Portfolio Choice over the Life-Cycle when the Stock and Labor Markets Are Cointegrated," *Journal of Finance*, 2007, *62* (5), 2123–2167.

Bhandari, Anmol, Serdar Birinci, Ellen R. McGrattan, and Kurt See, "What Do Survey Data Tell Us about US Businesses?," *American Economic Review: Insights*, 2020, *Forthcoming.*

Binsbergen, Jules van, "Duration-Based Stock Valuation," 2020. NBER Working Paper No. 27367.

Catherine, Sylvain, Max Miller, and Natasha Sarin, "Relaxing Household Liquidity Constraints through Social Security," *Journal of Public Economics*, 2020, *Forthcoming.*

— **, — , James D. Paron, and Natasha Sarin**, "Who Hedges Interest-rate Risk? Implications for Wealth Inequality," 2022. Working Paper.

Chetty, Raj, John N Friedman, Soren Leth-Peterson, Torben Heien Nielsen, and Tore Olsen, "Active vs Passive Decisions and Crowd-out in Retirement Savings Accounts: Evidence from Denmark Faculty Research," *The Quarterly Journal of Economics*, 2014, *121* (4), 1311–1346.

Domeij, David and Paul Klein, "Public pensions: To what extent do they account for Swedish wealth inequality?," *Review of Economic Dynamics*, 2002, *5* (3), 503–534.

Feldstein, Martin, "Social Security, Induced Retirement, and Aggregate Capital Accumulation," *Journal of Political Economy*, 1974, *82* (5), 905–926.

— , "Social Security and the Distribution of Wealth," *Journal of the American Statistical Association*, 1976, *71* (356), 800–807.

Fleming, Michael J. and Neel Krishnan, "The Microstructure of the TIPS Market," 2004. Federal Reserve Bank of New York Economic Policy Review.

Geanakoplos, John and Stephen P. Zeldes, "Market Valuation of Accrued Social Security Benefits," in Deborah Lucas, ed., *Measuring and Managing Federal Financial Risk*, National Bureau of Economic Research, 2010, pp. 213–233.

Greenwald, Daniel L, Matteo Leombroni, Hanno Lustig, and Stijn van Nieuwerburgh, "Financial and Total Wealth Inequality with Declining Interest Rates," 2021. NBER Working Paper Series No. 28613.

Gustman, Alan L., Olivia S. Mitchell, Andrew A. Samwick, and Thomas L. Steinmeier, "Pension and Social Security Wealth in the Health and Retirement Study," *Wealth, Work, and Health: Innovations in Survey Measurement in the Social Sciences*, 1999, pp. 150–208.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song, "What Do Data on Millions of U.S. Workers Reveal About Life-Cycle Earnings Risk?," *Econometrica (forthcoming)*, 2021.

— , Greg Kaplan, Jae Song, and Justin Weidner, "Lifetime Incomes in the United States Over Six Decades," 2018. NBER Working Paper Series No. 23371.

Hubmer, Joachim, Per Krusell, and Anthony A. Smith, "Sources of US Wealth Inequality: Past, Present, and Future," in Martin Eichenbaum and Erik Hurst, eds., *NBER Macroeconomics Annual*, National Bureau of Economic Research, 2020, pp. 391–455.

Miron, Jeffrey A. and David N. Weil, "The Genesis and Evolution of Social Security," in Michael D. Bordo, Claudia Goldin, and Eugene N. White, eds., *The Defining Moment: The Great Depression and the American Economy in the Twentieth Century*, University of Chicago Press, 1998, pp. 297–322.

Moll, Benjamin, "Comment on Hubmer, Krusell and Smith (2020) "Sources of U.S. Wealth Inequality: Past, Present, and Future"," in Martin Eichenbaum and Erik Hurst, eds., *NBER Macroeconomics Annual*, National Bureau of Economic Research, 2020, pp. 468–479.

Nickerson, Daniel and Kyle Burkhalter, "Unfunded Obligation and Transition Costs for the OASDI Program," Actuarial Note, Social Security Administration 2019.

Novy-Marx, Robert and Joshua Rauh, "Public Pension Promises: How Big Are They and What Are They Worth?," *The Journal of Finance*, 2011, *66* (4), 1211–1249.

Poterba, James M., "Retirement Security in an Aging Population," *American Economic Review*, 2014, *104* (5), 1–30.

Saez, Emmanuel and Gabriel Zucman, "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data," *The Quarterly Journal of Economics*, 2016, *131* (2), 519–578.

Samuelson, Paul A., "An Exact Consumption-Loan Model of Interest with or without the Social Contrivance of Money," *Journal of Political Economy*, 1958, *66* (6), 467–482.

Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun, "Are Americans Saving "Optimally" for Retirement?," *Journal of Political Economy*, 2006, *114* (4), 607–643.

Smith, Matthew, Owen Zidar, and Eric Zwick, "Top Wealth in the United States: New Estimates and Implications for Taxing the Rich," 2020. Working Paper.

Weil, David N, "Capital and Wealth in the 21st Century," *The American Economic Review: Papers and Proceedings*, 2015, *105* (5).

Wolff, Edward N., "Changing Inequality of Wealth," *American Economic Review*, 1992, *82* (2), 552–558.

— , "International comparisons of wealth inequality," *Review of Income and Wealth*, 1996, *42* (4), 433–451.

Zucman, Gabriel, "Global Wealth Inequality," *Annual Review of Economics*, 2019, *11* (1), 109–138.

For Online Publication

INTERNET APPENDIX

In this section, we give a detailed account of the methodology described in Section 4. We explain the construction of our dataset to allow for replication and explain our discount rate assumptions. We then describe the adjustments we make to reflect life expectancy differences, early/late retirement choices, and benefit adjustments for those who receive survivor benefits, or do not receive benefits at all. Finally, we provide a lengthy discussion of the steps followed to assign simulated Social Security wealth to the top and bottom of the marketable wealth distribution.

A Data Appendix

A.1 Survey of Consumer Finances variables

We use the triennial Survey of Consumer Finances for two main purposes: (i) measuring marketable wealth shares, and (ii) estimating aggregate Social Security wealth, and determining the share of Social Security wealth going to the wealthy. The SCF is well suited to measure marketable wealth shares as it covers most asset classes and provides detailed information on households' liabilities. We measure marketable wealth using the net worth variable: the sum of assets net of liabilities. One caveat is that the SCF does not survey extremely wealthy households. To fill this gap, we follow Saez and Zucman (2016) by adding data from the Forbes 400 list to the richest 0.01%.

Raw SCF To study Social Security in the SCF, we collect several variables from the raw SCF data which are listed below. We report the variable name for the second person in the household (typically the spouse) in parentheses.

– X5306 (X5311): Social Security benefit amount. Note that these are reported at different frequencies.

– X5307 (X5312): Social Security benefit frequency. The variable values and their corresponding frequencies are as follows: 4) monthly, 5) quarterly, 6) annually, 12) every two months, -7) other, 0) no benefits.

– X5304 (X5309): Social Security benefit type. This variable takes three values, which represent three benefit categories: 1) retirement, 2) disability, and 3) survivor.

– X5305 (X5310): Number of years receiving Social Security benefits.

– X19: Age of second person.

– X103: Gender of second person.

From these we create a series of variables. First, we create a payment frequency variable, given by

$$\texttt{pay\_freq} = \begin{cases} 12 & \text{if } \texttt{X5307}\ (\texttt{X5312}) = 4 \\ 4 & \text{if } \texttt{X5307}\ (\texttt{X5312}) = 5 \\ 1 & \text{if } \texttt{X5307}\ (\texttt{X5312}) = 6 \\ 2 & \text{if } \texttt{X5307}\ (\texttt{X5312}) = 12 \\ 0 & \text{otherwise} \end{cases}$$

which allows us to calculate annual benefits, given by

$$\texttt{ssinc} = \begin{cases} \texttt{X5306} * \texttt{pay\_freq} & \text{if Head of Household} \\ \texttt{X5311} * \texttt{pay\_freq} & \text{if Second Person in Household.} \end{cases}$$

We further subdivide this income by benefit type, with retirement income given by

$$\texttt{ssinc\_ret} = \begin{cases} \texttt{ssinc} & \text{if } \texttt{X5304}\ (\texttt{X5309}) = 1 \\ \texttt{ssinc} & \text{if } \texttt{X5304}\ (\texttt{X5309}) = 2\ \&\ \texttt{age}\ (\texttt{X19}) \geq 62 \end{cases}$$

and observed survivor benefits given by

$$\texttt{ssinc\_ben} = \texttt{ssinc} \quad \text{if } \texttt{X5304}\ (\texttt{X5309}) = 3.$$

Note that the second condition for retirement benefits assigns disability benefits going to people of retirement age as retirement benefits, consistent with the SSA. Finally, we calculate the age at retirement, which is given by

$$\texttt{ret\_age} = \begin{cases} \texttt{age} - \texttt{X5305} & \text{if Head of Household} \\ \texttt{X19} - \texttt{X5310} & \text{if Second Person in Household} \end{cases}$$

and is used to calculate full retirement age benefits in Section B.2.

Cleaned SCF Extract All wealth variables come from the cleaned SCF extract data. In particular, we use the `networth` variable to calculate the wealth distribution in each survey. This variable includes all assets less debt given in the SCF. We add to this the wealth held by the Forbes 400 as listed in the replication code of Saez and Zucman (2016). The SCF does not survey people beyond a certain wealth threshold, so people in the Forbes 400 are excluded from the sample. To fill this gap, we add aggregate Forbes 400 to the aggregate wealth of the Top 0.01%.

We also calculate a liquid wealth variable which is used to construct Appendix Figure E.6, Panel A. The component pieces of this variable are as follows:

- `liq`: liquid accounts, which is the sum of all checking, savings, and money market accounts, call accounts at brokerages, and prepaid cards.
- `cds`: certificates of deposit.
- `nmmf`: directly held mutual funds.
- `stocks`: wealth held in stocks.
- `bond`: wealth held in bonds of any type excluding savings bonds.

– `retqliq`: quasi-liquid retirement accounts, which are the sum of IRAs, thrift-type accounts, current pensions, and future pensions.

– `savbnd`: savings bonds.

– `homeeq`: home equity, which is the value of the home less the outstanding mortgage principal.

From these, liquid wealth is given by

$$\texttt{liquid\_wealth} = \texttt{liq} + \texttt{cds} + \texttt{nmmf} + \texttt{stocks} + \texttt{bond} + \texttt{retqliq} + \texttt{savbnd} + \texttt{homeeq}.$$

Finally, it is important to note that the Raw SCF values are in nominal terms (e.g. the 1995 Raw SCF is in 1995 dollars) while the Cleaned SCF Extract are in the dollars of the most recent survey year (e.g. 2019 dollars at the time of this writing). The SCF adjusts the Cleaned SCF Extract using the Consumer Price Index for all urban consumers (CPI-U-RS) from the Bureau of Labor Statistics. To make the two datasets consistent, we adjust the Cleaned SCF Extract to nominal dollars.

B Assumptions and adjustments

B.1 Market implied vs. SSA yield curve assumptions

Appendix Figure E.3 shows the differences in the yield curve assumptions implied from Treasury notes and the assumptions used by the SSA to compute the present value of Social Security obligations. The SSA discount rates are based on historical business cycles rather than market-implied rates, which is erroneous given the persistence of the current low interest rate environment.[8] An additional piece of evidence of the issues with the SSA's approach comes from the

[8] Summers, Lawrence, "U.S. Economic Prospects: Secular Stagnation, Hysteresis, and the Zero Lower Bound," *Business Economics*, 2014, *49* (2).

Federal Reserve, which reported in December, 2019 FOMC meeting projections that median long-run nominal rates are expected to be around 2.4-2.8%, with an upper bound of 3.3%, significantly below the 5+% suggested by the SSA.

B.2 Full retirement benefits

To validate the simulation methodology, we compare benefits in the simulated and SCF data. In reality, individuals can choose to retire early or delay retirement, meaning we must adjust their benefits in the data to compare them with benefits implied by the simulation. Beneficiaries retiring before the full retirement age receive reduced benefits, while beneficiaries retiring after the full retirement age receive increased benefits. Therefore, we define individual i's *full retirement benefit* as

$$\text{Full Retirement Benefit}_i = \frac{\text{Benefit}_i}{\text{Adjustment}}$$

where the adjustment term depends on the number of years that the beneficiary retires early or late.

For beneficiaries retiring early, the discount is 5/9% for each month before the full retirement age, up to 36 months, and 5/12% for each additional month. For beneficiaries retiring late, the amount of the credit depends of the beneficiary's birth year and can be found here. Further, the full retirement age is different for each cohort and can be found here. From these data, we create the `full_retirement_age` variable allowing us to determine the number of years of early or late retirement as

$$\texttt{ret\_discount\_years} = \texttt{full\_retirement\_age} - \texttt{ret\_age}.$$

This variable allows us to compute the appropriate benefit adjustment.

Here is an example to help clarify the procedure: Take a 62 year retiring in 2019. This person was born in 1957, meaning that the full retirement age for her cohort is 66 years and 6 months old. For this person, we have $\text{Adjustment} = (1 - \frac{5}{9} \cdot 36 - \frac{5}{12} \cdot 18)$, meaning that the full retirement benefit is given by

$$\text{Full Retirement Benefit}_i = \frac{\text{Benefit}_i}{(1 - \frac{5}{9} \cdot 36 - \frac{5}{12} \cdot 18)}.$$

In this case, the observed benefit is adjusted upward to account for the early retirement discount. Conversely, if the individual retires late, her observed benefit will be greater than the calculated full retirement benefit.

B.3 Adjusting life expectancy by income

We adjust for differential life expectancy across income centiles using data from Chetty, Friedman, Leth-Peterson, Nielsen and Olsen (2014) as reported by the Health Inequality Project (HIP). These data provide life expectancy at age 40 for each lifetime income centile from 2001 to 2014. Since our sample starts in 1989 and goes until 2014, we apply the 2001 data for all years between 1989–2001 and the 2014 data for 2014–2019. Assigning the 2001 values to previous years seems to be a reasonable assumption, as the life expectancy differential between high and low income individuals is flat from 2001–2007, then expands after the 2008 Financial Crisis, as shown in Figure B.1.

Using these data, we compute the number of years fewer (more) that a retired SCF respondent will live given their lifetime income centile. We then adjust the respondents age to reflect the shorter (longer) longevity implied by the data. To do this, the compute the *life expectancy spread* for each lifetime income centile in the HIP data, which is given by

$$\text{life expectancy spread}_{centile,t} = \frac{\text{life expectancy}_{centile,t}}{\frac{1}{100}\sum_{centile=1}^{100} \text{life expectancy}_{centile,t}}.$$

We then take these life expectancy spreads and merge them with our primary mortality dataset coming from the Human Mortality Database (HMD). We then calculate the number of years fewer (more) people in the lower (higher) centiles of the income distribution live based on the unconditional life expectancy (i.e. at age 0). We define this as the *year difference* which is given by

$$\text{year difference}_{centile,t} = (\text{life expectancy spread}_{centile,t} - 1) \times \text{unconditional life expectancy}_t.$$

Note, that this will be negative for people in the bottom half of the lifetime income distribution and positive for people in the top half. From this, we calculate the *effective mortality age* for each SCF

Figure B.1: Life expectancy differential, 2001–2014

This figure plots the difference in life expectancy for people in the top half and bottom half of the lifetime earnings distribution. The differences for men and women are plotted separately. The vertical line in the middle of the graph denotes the period before and after 2007.

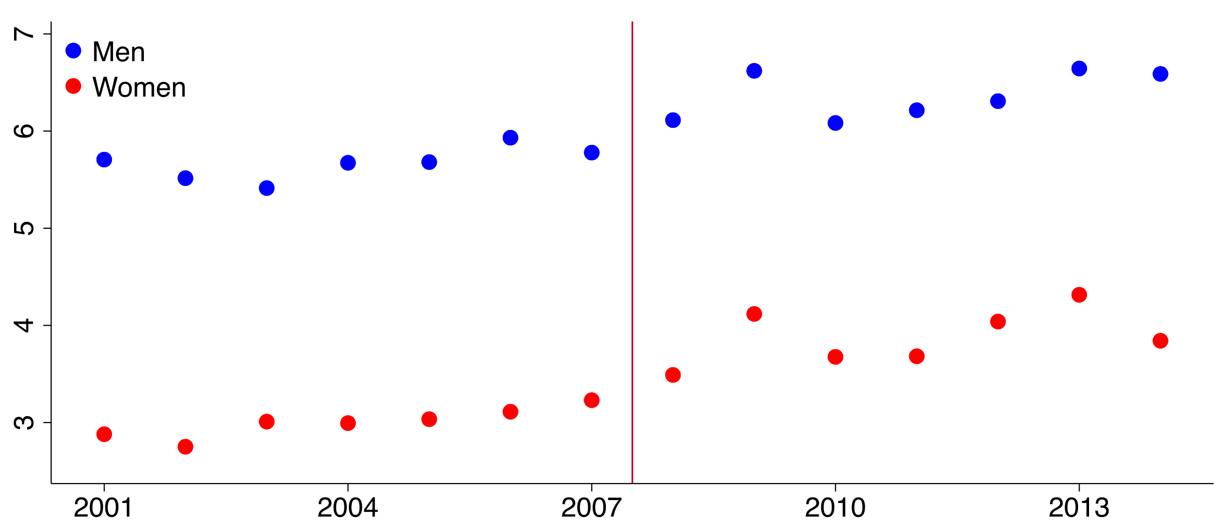

respondent, which is given by

$$\text{effective mortality age}_{i,centile,t} = \text{current age}_i - \text{year difference}_{centile,t}.$$

We then assign survival probabilities to that individual based on their effective mortality age.

Completing the life expectancy adjustment requires a valid proxy for lifetime income. Unfortunately, the SCF does not provide income histories. However, we can extrapolate based on the Social Security retirement benefits centile. Since Social Security benefits are a monotonically increasing function of lifetime income, this proxy allows us to preserve the order of individuals within the lifetime income distribution, which we then apply to the life expectancy adjustment.

An example is illustrative on this procedure: the life expectancy for men in 2019 in the HMD data is 76 years, and in that year, a person in the 1st lifetime income centile lives approximately 9 years less than the average person. Therefore, a 40 year old man in the 1st lifetime income centile

Figure B.2: Adjusting for differential in life expectancy

This figure shows per capita Social Security wealth for each person in the SCF, applying population weights, for people in the top 10% (Panel A) and bottom 90% (Panel B) of the non-Social Security wealth distribution. Life expectancy adjusted values incorporate differential life expectancy across income centiles using data from the Health Inequality Project (HIP), as outlined in Appendix B.3.

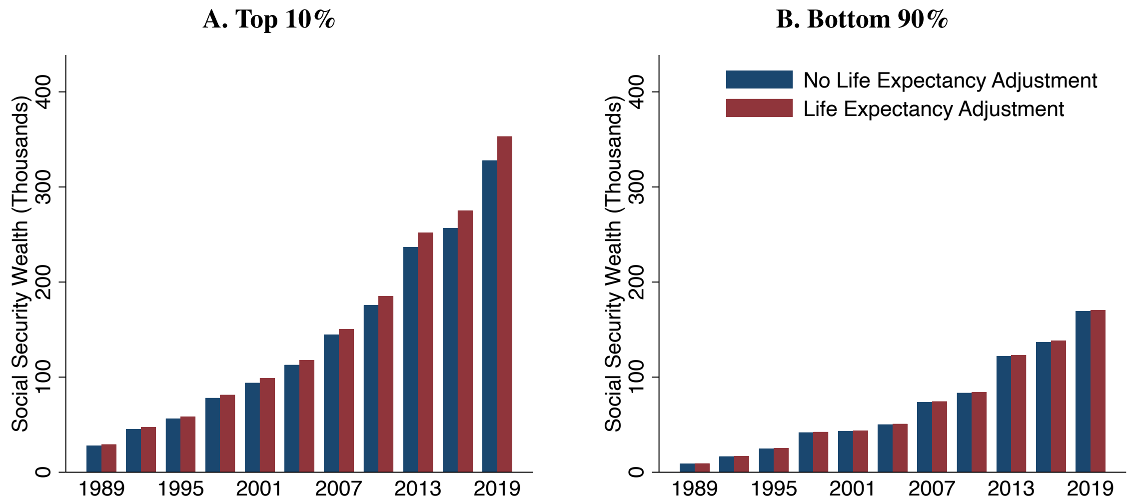

has an effective mortality age of 49 years old, and he would be assigned the survival probabilities of a 49 year old man in 2019. We apply this life expectancy correction both to retired workers and to those still in the workforce, whose earnings histories we simulate.

Figure B.2 shows how a life expectancy adjustment impacts Social Security wealth across the Social Security benefits distribution among current retirees. When differences in mortality rates are accounted for, per capital Social Security wealth that accrues to the bottom decile falls by nearly 25%, and per capita Social Security wealth falls for all but the top three deciles. We modify our estimates of cohort Social Security wealth to reflect these differences.

However, this adjustment does not have a large impact on top wealth shares. This is due to an increase in the benefit-weighted average life expectancy of beneficiaries in the bottom 90%. Specifically, those in upper deciles of the marketable wealth distribution live for longer (more years of benefits) than those in lower deciles. Within the bottom 90%, the effect of this adjustment is to decrease benefit-years for individuals with lower benefits, and increase benefit-years for individuals with higher benefits.

As such, adjusting for the relationship between income level and mortality rates increases Social Security wealth for both the top and bottom of the overall wealth distribution. Though the increase in aggregate Social Security wealth goes disproportionately to the wealthy, it remains, nonetheless, much more equally distributed than marketable wealth.[9]

B.4 Capitalizing implied survivor benefits

Widows can receive a share of the Social Security benefits of their deceased spouses. We account for this when capitalizing benefits by computing how likely it is that a respondent's spouse is alive given that the respondent is deceased, under the assumption that the survival probabilities of the couple are uncorrelated. In particular, widows can receive the maximum of their benefit and their deceased spouse's benefit. The implied present value of survivor benefits is therefore given by

$$\text{Implied Survivor Benefits}_{i,t} = \max\left\{\text{Spouse Benefits}_{i,t} - \text{Benefits}_{i,t}, 0\right\} \times \sum_{s=0}^{\infty} \frac{\prod_{k=t}^{s-1} m_{i,t+k}(1 - m_{i,t+k}^{spouse})}{1 + r_{t,t+s}}$$

where m represents the survival probability and r the real discount rate.

B.5 Proportion of people with no benefits

The vast majority of retirees receive some form of Social Security benefits. However, a fraction of retirees have insufficient work history to receive benefits. When aggregating Social Security benefits, we must take this into account. This requires a reasonable estimation of the proportion of people in each cohort that do not receive benefits.

[9] It is worth noting that this exercise illustrates the issue with a singular focus on top shares as a measure of wealth inequality. Differences in life expectancy disproportionately impact those at the bottom of the wealth distribution, but standard measures of wealth concentration focus on the share of aggregate wealth accruing to those at the top, thus missing out on such dynamics.

We estimate this using Deaton-Paxson regressions for each gender, which is a constrained regression of the following form

$$\log(Pr(\text{No Retirement Benefits}))_{t,a,b} = \gamma_t + \eta_a + \delta_b + \varepsilon_{t,a,b} \tag{B.1}$$

subject to

$$\sum_{1989}^{2016} \gamma_t = 0 \tag{B.2}$$

$$\sum_{1989}^{2016} \gamma_t (t - 2002.5) = 0 \tag{B.3}$$

$$\eta_{72} = 0. \tag{B.4}$$

where a represents each age, t each survey year, and b each birth year.[10][11] The coefficients of interest are the birth year fixed effects, where this empirical set-up allows us to adjust for survey specific sampling error and age specific effects. The fitted values by birth year are shown in Appendix Figure E.4, where the average number of zero Social Security income respondents is shown to be 10% for men and 20% for women. In the simulation, these estimates are used to determining average Social Security wealth.

C Assignment and aggregation

After generating the age-year-gender averages from the simulated panel, we merge the simulated data with the SCF. The steps we use to assign Social Security wealth to the top 10% and top 1% of the marketable wealth distribution are as follows:

1. Determine how Social Security wealth is distributed across the marketable wealth distribution among retirees aged 65–75.

[10] Note that respondents are grouped into three-year age and birth year cohorts in this estimation.

[11] Deaton, Angus S. and Christina Paxson, "Saving, Growth, and Aging in Taiwan," *Studies in the Economics of Aging*, National Bureau of Economic Research, 1994, pp. 331–362.

(a) Find the share of full retirement age Social Security wealth accruing to each wealth centile in each survey year.

(b) Define the share going to each wealth centile w as $\alpha_{w,t} = \frac{SSW_{w,t}}{\sum_{w=1}^{100} SSW_{w,t}}$. Then define $\phi_t(x) = \sum_{w=x}^{100} \alpha_{w,t}$ as the cumulative share of benefits going to people above centile x, where the subscript t denotes different survey years. For 2019, this function is shown in Panel B of Figure 5.

2. Determine proportion of the population of each cohort in the top 10% and top 1% of wealth distribution, which we denote by $k^g_{c,t}$, where the subscript c denotes different cohorts, and the superscript g denotes different populations (i.e. in this case either the top 10% or top 1%).

   (a) This means that the top 10% share of population is given by $k^{\text{Top 10\%}}_{c,t} \equiv \frac{N^{\text{Top 10\%}}_{c,t}}{N^{\text{Full}}_{c,t}}$ and the top 1% share by $k^{\text{Top 1\%}}_{c,t} \equiv \frac{N^{\text{Top 1\%}}_{c,t}}{N^{\text{Full}}_{c,t}}$, where $N^g_{c,t}$ is the total size of cohort c in survey t in population g. Mathematically, this is given by $N^g_{c,t} \equiv \sum_i \text{wgt}_i \mathbb{1}_i(c,t,g)$ where wgt_i is the weight in the SCF for observation i and $\mathbb{1}$ is equal to 1 if observation i is in year t, of cohort c, and in population g.

   (b) For example, for respondents aged 40 in 2019, 4.8% are in the top 10% and 1.2% are in the top 1% of the aggregate 2019 wealth distribution.

3. Assign average Social Security wealth by cohort-year-gender from the simulated panel to each person in the SCF. This is denoted by $\overline{\text{SSW}}_{c,t,s}$ where the subscript s denotes each sex.

4. For all cohorts less than 66 years of age, calculate average Social Security wealth from the simulation for the top 1%, the rest of the top 10% and bottom 90%. For people in the top 1%, this is given by

$$\overline{\text{SSW}}^{\text{Top 1\%}}_{c,t} = \frac{\phi_t(k^{\text{Top 1\%}}_{c,t})}{N^{\text{Top 1\%}}_{c,t}} \cdot \text{Cohort Social Security Wealth}_{c,t},$$

for people in the rest of the top 10% by

$$\overline{\text{SSW}}_{c,t}^{\text{Rest of Top 10\%}} = \frac{(\phi_t(k_{c,t}^{\text{Top 10\%}}) - \phi_t(k_{c,t}^{\text{Top 1\%}}))}{(N_{c,t}^{\text{Top 10\%}} - N_{c,t}^{\text{Top 1\%}})} \times \text{Cohort Social Security Wealth}_{c,t},$$

and for people in the bottom 90% by

$$\overline{\text{SSW}}_{c,t}^{\text{Bottom 90\%}} = \frac{1 - \phi_t(k_{c,t}^{\text{Top 10\%}})}{N_{c,t}^{\text{Bottom 90\%}}} \times \text{Cohort Social Security Wealth}_{c,t},$$

where Cohort Social Security Wealth$_{c,t} \equiv \sum_s \left(N_{c,t,s}^{\text{Full}} \times \overline{\text{SSW}}_{c,t,s} \right)$, and $\phi_t(x)$ is the function from Step 1.

5. For respondents less than 62 years of age, nothing else needs to be done. We calculate their aggregate Social Security wealth, which is given by the sum of the SCF weights multiplied by the assigned averages from (4).

6. For respondents aged 62–69, the simulation meets the data, meaning that we have respondents in the data with observed Social Security benefits, as well as respondents not receiving benefits that will receive them in the future. For respondents currently receiving benefits, we calculate the present value of those benefits to determine their Social Security wealth. For respondents not currently receiving benefits, we fill in their benefits using either the average benefits from (4) or a backfilling methodology.

   (a) For respondents aged 62–65, we fill in the average benefits calculated in (4) for all non-recipients.

   (b) For respondents aged 66–69, we fill in the average observed retirement-adjusted Social Security wealth in the data for recipients aged 70–73 from the succeeding survey adjusted for inflation. This, of course, only works for 1989–2016, so for 2019, we fill in the average observed retirement-adjusted Social Security wealth in the data for recipients aged 70–73 from the 2019 survey.

(c) However, we must adjust these filled benefits downward, as these respondents have a higher probability of being a non-recipient. This is because we assume that 10% of males and 20% of females will not receive retirement benefits (this is verified in the data). For people ineligible for benefits (i.e. less than 62 years old), no additional adjustment must be made. But for people above 62, we must adjust. An example will clarify why. Assume that 50% of men will claim benefits at age 62. This means that 50% of male beneficiaries receive no benefits in that year, and 10% of those men will never receive benefits, meaning that the proportion of people never receiving benefits in that subsample is 20%. In this case, the average benefit will be given by $\left(\frac{0.8}{0.9}\right)\overline{\text{SSW}}^{g}_{a,t}$ to account for the increased probability of never receiving benefits in the subpopulation. Formally, this adjustment is given by

$$\text{adj}^{g}_{a,t,s} = \frac{\sum \mathbb{1}\{\text{No Benefits}\} - .1(1 + \mathbb{1}\{\text{Female}\})}{\sum \mathbb{1}\{\text{No Benefits}\}(1 - .1(1 + \mathbb{1}\{\text{Female}\}))}$$

where $\mathbb{1}\{x\}$ is an indicator variable equal to 1 when conditions x are met. This adjustment is calculated for each year-age-sex-population combination.

7. For all cohorts older than 70, we aggregate all values from the data. Nothing needs to be filled in for these observations, as there is no benefit from claiming after 70. In reality, some people may claim later, but we assume that these individuals will not receive benefits for the remainder of their lives.

D Market beta of aggregate labor income

Consider the following exogenous system of stochastic processes

$$dy_t = -\kappa y_t dt + \begin{bmatrix} \sigma_l \\ -\sigma_s \end{bmatrix}^T dz_t, \tag{D.1}$$

$$ds_t = \left(\mu - \frac{\sigma^2}{2} + \phi y_t\right) dt + \begin{bmatrix} 0 \\ \sigma \end{bmatrix}^T dz_t, \tag{D.2}$$

$$l_{1,t} = y_t + s_t - \delta t, \tag{D.3}$$

$$d\pi_t = -r\pi_t dt - \begin{bmatrix} 0 \\ \lambda \end{bmatrix}^T \pi_t dz_t, \tag{D.4}$$

where y_t is log output, $s_t \equiv \log S_t$ is log stock price, $l_{1,t} \equiv \log L_{1,t}$ is log wage, π_t is the state-price density, $\lambda \equiv \frac{\mu - r}{\sigma}$, and $z_t = \begin{bmatrix} z_{1,t} \; z_{2,t} \end{bmatrix}^T$ is a standard Brownian motion. Note that, for now, we allow the $\sigma \neq \sigma_s$, which is different than in Equation (20) and gives us a more general solution.

We want to find the beta at time t on a "wage strip", which is a security that pays out $L_{1,t+n}$ at $t + n$ and is denoted by

$$\beta_t^{L_1,n} = \frac{\text{Cov}_t\left(r_t^m dt, r_t^{L_1,n} dt\right)}{\text{Var}_t\left[r_t^m dt\right]}.$$

In this economy, the instantaneous return on the market r_t^m is defined by

$$r_t^m dt = \frac{dS_t}{S_t} = ds_t + \frac{1}{2}\left(ds_t\right)^2 = \left(\mu + \phi y_t\right) dt + \begin{bmatrix} 0 \\ \sigma \end{bmatrix}^T dz_t,$$

and the instantaneous return on the wage strip $r_t^{L_1,n}$ by

$$r_t^{L_1,n} dt = \frac{dP_t^{L_1,n}}{P_t^{L_1,n}},$$

where $P_t^{L_1,n}$ is the price of the wage strip. By no-arbitrage, the price of the wage strip is given by

$$P_t^{L_1,n} = \mathbb{E}_t\left[\frac{\pi_{t+n}}{\pi_t} L_{1,t+n}\right] = \mathbb{E}_t\left[\exp\left\{\tilde{\pi}_{t+n} - \tilde{\pi}_t + l_{1,t+n}\right\}\right], \tag{D.5}$$

where $\tilde{\pi}_t \equiv \log \pi_t$. The process $\tilde{\pi}_t$ is given by

$$d\tilde{\pi}_t = \frac{d\pi_t}{\pi_t} - \frac{1}{2}\left(\frac{d\pi_t}{\pi_t}\right)^2 = \left(-r - \frac{1}{2}\lambda^2\right) dt - \begin{bmatrix} 0 \\ \lambda \end{bmatrix}^T dz_t$$

$$\Rightarrow \tilde{\pi}_t = \left(-r - \frac{1}{2}\lambda^2\right) t - \begin{bmatrix} 0 \\ \lambda \end{bmatrix}^T z_t$$

which comes from a straightforward application of Ito's lemma.

To solve Equation (D.5), we are left with finding $l_{1,t+n}$, which is equivalent to solving for y_t and s_t. Using Ito's lemma, we find that

$$y_t = e^{-\kappa t}\left(y_0 + \begin{bmatrix} \sigma_l \\ -\sigma_s \end{bmatrix}^T \int_0^t e^{\kappa s} dz_s\right).$$

Now, to find s_t, we introduce a new variable $\tilde{s}_t$ defined as

$$\tilde{s}_t = s_t + \frac{\phi}{\kappa} y_t,$$

which is given by

$$d\tilde{s}_t = ds_t + \frac{\phi}{\kappa} dy_t = \left(\mu - \frac{\sigma^2}{2}\right) dt + \begin{bmatrix} \frac{\phi}{\kappa}\sigma_l \\ \sigma - \frac{\phi}{\kappa}\sigma_s \end{bmatrix}^T dz_t$$

$$\Rightarrow \tilde{s}_t = \left(\mu - \frac{\sigma^2}{2}\right) t + \begin{bmatrix} \frac{\phi}{\kappa}\sigma_l \\ \sigma - \frac{\phi}{\kappa}\sigma_s \end{bmatrix}^T z_t$$

Using this expression, we solve for s_t, yielding

$$s_t = \tilde{s}_t - \frac{\phi}{\kappa} y_t = \left(\mu - \frac{\sigma^2}{2}\right) t + \begin{bmatrix} \frac{\phi}{\kappa}\sigma_l \\ \sigma - \frac{\phi}{\kappa}\sigma_s \end{bmatrix}^T z_t - \frac{\phi}{\kappa} e^{-\kappa t}\left(y_0 + \begin{bmatrix} \sigma_l \\ -\sigma_s \end{bmatrix}^T \int_0^t e^{\kappa s} dz_s\right)$$

which implies that $l_{1,t}$ equals

$$l_{1,t} = y_t + s_t - \delta t = \left(\mu - \frac{\sigma^2}{2} - \delta\right) t + \begin{bmatrix} \frac{\phi}{\kappa}\sigma_l \\ \sigma - \frac{\phi}{\kappa}\sigma_s \end{bmatrix}^T z_t + \left(1 - \frac{\phi}{\kappa}\right) y_t.$$

Plugging everything back into the exponential expression of Equation (D.5), we obtain

$$\begin{aligned}
\tilde{\pi}_{t+n} - \tilde{\pi}_t + l_{1,t+n} &= \left(-r - \frac{1}{2}\lambda^2\right)(t+n) - \begin{bmatrix} 0 \\ \lambda \end{bmatrix}^T z_{t+n} - \left(-r - \frac{1}{2}\lambda^2\right) t + \begin{bmatrix} 0 \\ \lambda \end{bmatrix}^T z_t \\
&\quad + \left(\mu - \frac{\sigma^2}{2} - \delta\right)(t+n) + \begin{bmatrix} \frac{\phi}{\kappa}\sigma_l \\ \sigma - \frac{\phi}{\kappa}\sigma_s \end{bmatrix}^T z_{t+n} + \left(1 - \frac{\phi}{\kappa}\right) y_{t+n} \\
&= \left(-r - \frac{1}{2}\lambda^2\right) n + \left(\mu - \frac{\sigma^2}{2} - \delta\right)(t+n) + \begin{bmatrix} 0 \\ \lambda \end{bmatrix}^T z_t + \\
&\quad \begin{bmatrix} \frac{\phi}{\kappa}\sigma_l \\ \sigma - \frac{\phi}{\kappa}\sigma_s - \lambda \end{bmatrix}^T z_{t+n} + \left(1 - \frac{\phi}{\kappa}\right) y_{t+n}
\end{aligned}$$

Note that all components inside the exponent in Equation (D.5) are normal variables. Hence, we can rewrite the equation as

$$P_t^{L_1,n} = \exp\left\{\mathbb{E}_t\left[\tilde{\pi}_{t+n} - \tilde{\pi}_t + l_{1,t+n}\right] + \frac{1}{2}\text{Var}_t\left[\tilde{\pi}_{t+n} - \tilde{\pi}_t + l_{1,t+n}\right]\right\}, \tag{D.6}$$

which leaves us with finding the two components in the exponent. Also note how we can express y_{t+n} via y_t:

$$\begin{aligned}
y_{t+n} &= e^{-\kappa(t+n)}\left(y_0 + \begin{bmatrix} \sigma_l \\ -\sigma_s \end{bmatrix}^T \int_0^{t+n} e^{\kappa s} dz_s\right) \\
&= e^{-\kappa n}\left(y_t + \begin{bmatrix} \sigma_l \\ -\sigma_s \end{bmatrix}^T \int_t^{t+n} e^{\kappa(s-t)} dz_s\right)
\end{aligned}$$

The first expression, $\mathbb{E}_t\left[\tilde{\pi}_{t+n} - \tilde{\pi}_t + l_{1,t+n}\right]$, is given by

$$\mathbb{E}_t\left[\tilde{\pi}_{t+n} - \tilde{\pi}_t + l_{1,t+n}\right] = \left(\mu - \frac{\sigma^2}{2} - \delta\right) t - \left(\frac{1}{2}\left(\lambda - \sigma\right)^2 + \delta\right) n + \left(1 - \frac{\phi}{\kappa}\right) e^{-\kappa n} y_t + \begin{bmatrix} \frac{\phi}{\kappa}\sigma_l \\ \sigma - \frac{\phi}{\kappa}\sigma_s \end{bmatrix}^T z_t$$

and the second expression, $\mathrm{Var}_t\left[\tilde{\pi}_{t+n} - \tilde{\pi}_t + l_{1,t+n}\right]$, by

$$\begin{aligned} \mathrm{Var}_t\left[\tilde{\pi}_{t+n} - \tilde{\pi}_t + l_{1,t+n}\right] = {} & \left(\left(\frac{\phi}{\kappa}\sigma_l\right)^2 + \left(\sigma - \frac{\phi}{\kappa}\sigma_s - \lambda\right)^2\right) n \\ & + \left(1 - \frac{\phi}{\kappa}\right)^2 \left(\sigma_l^2 + \sigma_s^2\right) \frac{1}{2\kappa}\left(1 - e^{-2\kappa n}\right) \\ & + 2\left(1 - \frac{\phi}{\kappa}\right)\left(\frac{\phi}{\kappa}\sigma_l^2 + \frac{\phi}{\kappa}\sigma_s^2 - \sigma\sigma_s + \lambda\sigma_s\right)\frac{1}{\kappa}\left(1 - e^{-\kappa n}\right). \end{aligned}$$

From this, we obtain the solution for $P_t^{L_1,n}$,

$$P_t^{L_1,n} = \exp\left\{at + b + cy_t + d^T z_t\right\}, \tag{D.7}$$

where

$$\begin{aligned} a &\equiv \mu - \frac{\sigma^2}{2} - \delta \\ b(n) &\equiv -\left(\delta - \frac{1}{2}\frac{\phi^2}{\kappa^2}\left(\sigma_l^2 + \sigma_s^2\right) + \frac{\phi}{\kappa}\sigma_s\left(\sigma - \lambda\right)\right) n + \left(1 - \frac{\phi}{\kappa}\right)^2 \left(\sigma_l^2 + \sigma_s^2\right)\frac{1}{4\kappa}\left(1 - e^{-2\kappa n}\right) \\ &\quad + \left(1 - \frac{\phi}{\kappa}\right)\left(\frac{\phi}{\kappa}\left(\sigma_l^2 + \sigma_s^2\right) - \sigma_s\left(\sigma - \lambda\right)\right)\frac{1}{\kappa}\left(1 - e^{-\kappa n}\right) \\ c(n) &\equiv \left(1 - \frac{\phi}{\kappa}\right) e^{-\kappa n} \\ d &= \begin{bmatrix} \frac{\phi}{\kappa}\sigma_l \\ \sigma - \frac{\phi}{\kappa}\sigma_s \end{bmatrix}. \end{aligned}$$

From Equation (D.7), we can find the return on the wage strip by differentiating its price. To do that, we can rewrite its price as

$$P_t^{L_1,n} = \exp\left\{P_t^{L_1,n}\right\},$$

where

$$P_t^{L_1,n} = at + b(n) + c(n)y_t + d^T z_t.$$

By Ito's lemma we have (note that $dn = -dt$)

$$dP_t^{L_1,n} = \left(a - b'(n) - c'(n)y_t - \kappa c(n)y_t\right)dt + \left(c(n)\begin{bmatrix}\sigma_l \\ -\sigma_s\end{bmatrix} + d\right)^T dz_t, \tag{D.8}$$

where

$$\begin{aligned}
b'(n) &= \frac{1}{2}\left(\sigma_l^2 + \sigma_s^2\right)\left(\frac{\phi}{\kappa} + c\right)^2 - \sigma_s\left(\sigma - \lambda\right)\left(\frac{\phi}{\kappa} + c\right) - \delta \\
c'(n) &= -\kappa\left(1 - \frac{\phi}{\kappa}\right)e^{-\kappa n} = -\kappa c(n).
\end{aligned}$$

Then, the return on the wage strip equals

$$\begin{aligned}
r_t^{L_1,n}dt &= \frac{dP_t^{L_1,n}}{P_t^{L_1,n}} = dP_t^{L_1,n} + \frac{1}{2}\left(dP_t^{L_1,n}\right)^2 \\
&= \left(a - b'(n) + \frac{1}{2}\left(c\sigma_l + \frac{\phi}{\kappa}\sigma_l\right)^2 + \frac{1}{2}\left(\sigma - \frac{\phi}{\kappa}\sigma_s - c\sigma_s\right)^2\right)dt + \begin{bmatrix}c\sigma_l + \frac{\phi}{\kappa}\sigma_l \\ \sigma - \frac{\phi}{\kappa}\sigma_s - c\sigma_s\end{bmatrix}^T dz_t
\end{aligned}$$

meaning that the expected return is

$$\mathbb{E}_t\left[r_t^{L_1,n}\right] = \mu - \left(\mu - r\right)\frac{\sigma_s}{\sigma}\left(\frac{\phi}{\kappa} + c\right).$$

This gives the beta on the wage strip as

$$\beta^{L_1,n} = \frac{\text{Cov}_t\left(r_t^m dt, r_t^{L_1,n}dt\right)}{\text{Var}_t\left[r_t^m dt\right]} = 1 - \frac{\sigma_s}{\sigma}\left(\frac{\phi}{\kappa} + c\right)$$

Further, we can test if the CAPM holds in this economy. To do this, we assess if $\mathbb{E}_t\left[r_t^{L_1,n} - r\right] = \beta^{L_1,n}\mathbb{E}_t\left[r_t^m - r\right]$ holds. The RHS of the expression is given by

$$\beta^{L_1,n}\mathbb{E}_t\left[r_t^m - r\right] = \left(1 - \frac{\sigma_s}{\sigma}\left(\frac{\phi}{\kappa} + c\right)\right)\left(\mu - r + \phi y_t\right)$$

and the LHS by

$$\mathbb{E}_t\left[r_t^{L_1,n}-r\right]=\left(1-\frac{\sigma_s}{\sigma}\left(\frac{\phi}{\kappa}+c\right)\right)(\mu-r)\,.$$

Therefore, the CAPM only holds when y_t is zero in this economy.

Finally, note that if we assume no contemporaneous correlation between the labor and stock market ($\sigma_s=\sigma$), the results reduce to

$$\begin{aligned}\beta_t^{L_1,n}&=\left(1-\frac{\phi}{\kappa}\right)\left(1-e^{-\kappa n}\right)\\ \mathbb{E}_t\left[r_t^{L_1,n}\right]&=\left(1-\frac{\phi}{\kappa}\right)\left(1-e^{-\kappa n}\right)(\mu-r)+r\end{aligned}$$

while the discount rate remains unchanged as it does not depend on σ_s. So, when $n\to\infty$, the beta converges to $1-\frac{\phi}{\kappa}=1-\frac{0.08}{0.16}=0.5$.

E Additional figures

Figure E.3: Market Implied and Social Security Administration Yield Curve Estimates

This figure presents the differences between the yield curves implied by treasury markets and those used in SSA reports. The market series is extended by extrapolating the 29-to-30 year forward rate into the future, as described in Section 3.

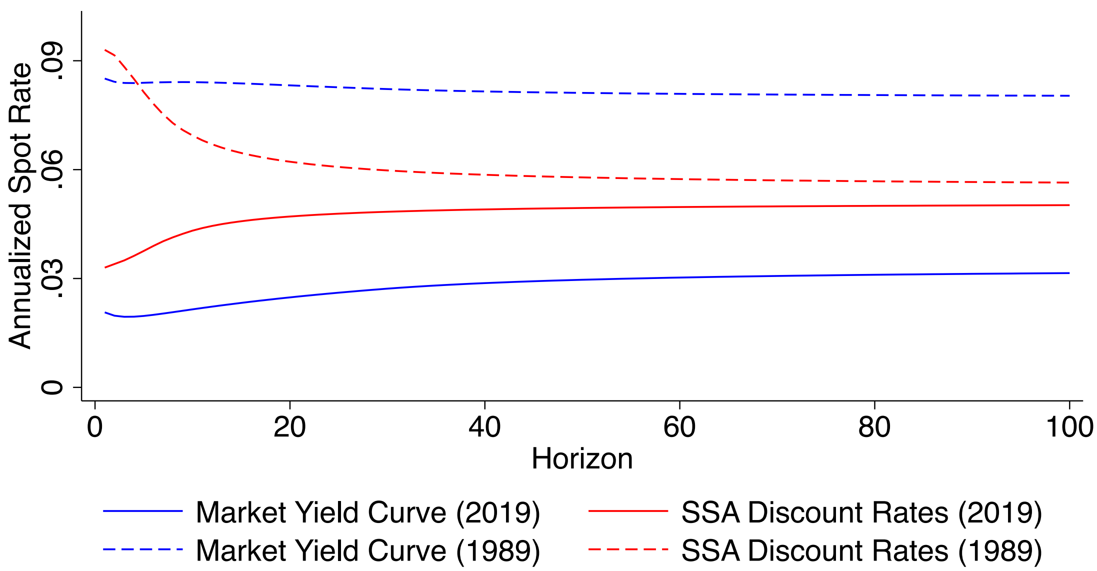

Figure E.4: Zero-Social Security Income Estimates: Deaton-Paxson Regressions

This figure shows the results for the Deaton-Paxson regressions outlined in Appendix B.5. The solid lines represent the estimated proportion of male and female respondents not receiving benefits after adjusting for survey-year and age specific fixed effects in a constrained. The dashed lines represent the mean proportion not receiving benefits for the 1929-1953 birth cohorts.

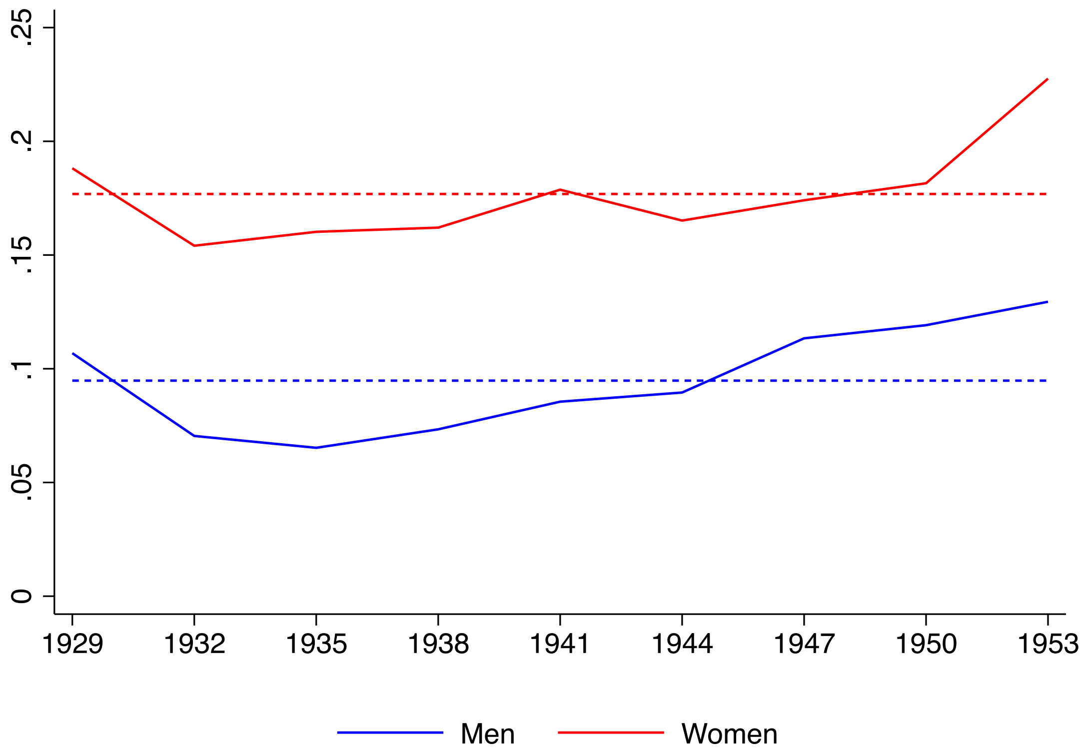

Figure E.5: Funding Gap: Payable Benefits under 1989 and 2019 SSA projections

This figure shows the proportion of payable benefits under the SSA's different funding gap assumptions. Benefits cuts for horizons greater than 75 years are assumed to be the same as the 75th year benefits cuts.

A. 1989

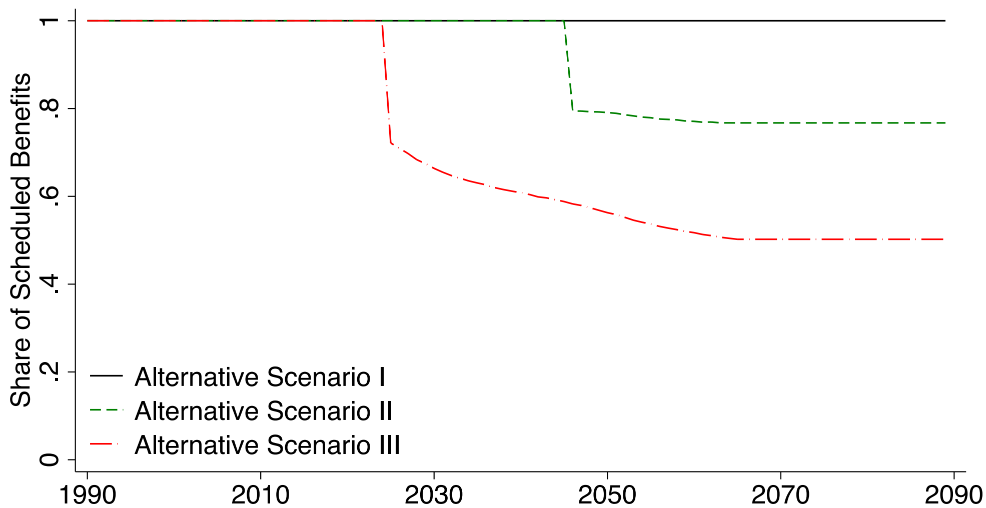

B. 2019

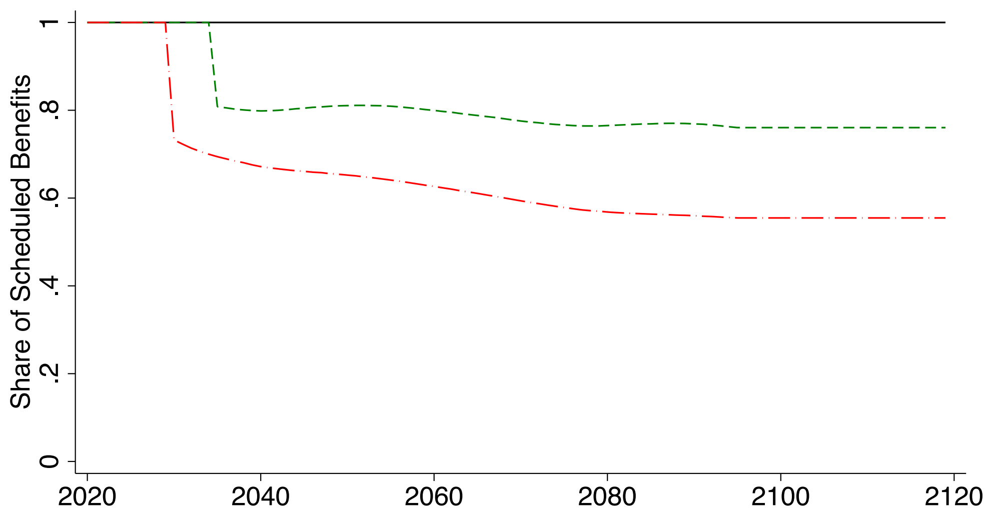

Figure E.6: Accessible and Social Security Wealth over the Lifecycle

Panel A shows the weighted proportion of SCF respondents with more than $10,000, $50,000, and $100,000 of accessible wealth by three-year age group in the 2019 SCF. The measure of accessible wealth we employ sums all wealth from liquid savings, stocks, bonds, mutual funds, quasi-liquid retirement accounts, and home equity and subtracts the total value of all non-mortgage debt. Panel B shows the cumulative share of Social Security wealth by age.

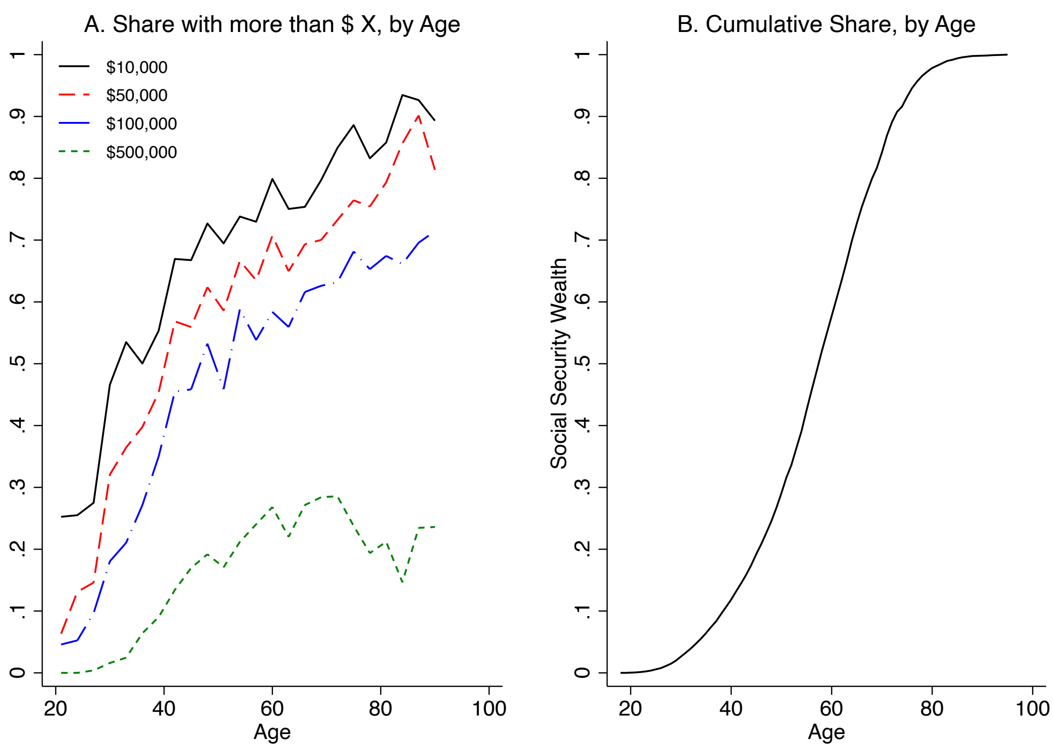

Table E.1: Calibration of labor income process

Parameter estimates for Sections 2.1 and 4 come from Specifications (5) and (6) in Guvenen et al. (2021). Parameters can be found in Table IV and Table D.3 of the 2021 working paper version.

Parameter	Section 2.1 Calibration	Section 4 Calibration
ρ	0.991	0.959
p_z	17.6%	40.7%
$\mu_{\eta,1}$	-0.524	-0.085
$\sigma_{\eta,1}$	0.113	0.364
$\sigma_{\eta,2}$	0.046	0.069
$\sigma_{z_1,0}$	0.450	0.714
λ	0.016	0.0001
p_ε	4.4%	13.0%
$\mu_{\varepsilon,1}$	0.134	0.271
$\sigma_{\varepsilon,1}$	0.762	0.285
$\sigma_{\varepsilon,2}$	0.055	0.037
σ_α	0.472	0.300
$\sigma_\beta \cdot 10$	N/A	0.196
$\text{corr}_{\alpha\beta}$	N/A	0.768
$a_\nu \cdot 1$	-2.495	-3.353
$b_\nu \cdot t$	-1.037	-0.859
$c_\nu \cdot z_t$	-5.051	-5.034
$d_\nu \cdot t \cdot z_t$	-1.087	-2.895
$a_{z_1} \cdot 1$	0.176	0.407