

Social Security and Trends in Wealth Inequality

SYLVAIN CATHERINE, MAX MILLER, and NATASHA SARIN*

Forthcoming at *The Journal of Finance*

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 changed much over the last three decades when Social Security is properly accounted for. This is because Social Security wealth increased substantially from \$7.2 trillion in 1989 to \$40.6 trillion in 2019 and now represents nearly 50% 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.

As work dating back to Feldstein (1976) shows, wealth is more equally distributed when it includes the value of accrued Social Security benefits. This paper highlights that including Social Security does not just change the level of wealth inequality, but its evolution as well. Specifically, although recent studies show that wealth inequality in the United States has risen

*Sylvain Catherine is with University of Pennsylvania. Max Miller is with Harvard Business School. Natasha Sarin is with Yale University. 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 and Desmond O'Shaughnessy for their 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." We have read the *Journal of Finance's* disclosure policy and have no conflicts of interest to disclose.

Correspondence: Natasha Sarin, Yale Law School. Email: natasha.sarin@yale.edu

in the last several decades (Saez and Zucman, 2016), we find that top wealth shares have not changed much between 1989 and 2019 when Social Security is properly taken into account. Wealth remains highly concentrated at the top, but not substantially more than in 1989. In 2019, those in the top 1% of the distribution hold just under 24% of wealth, inclusive of Social Security, up from 22% in 1989.

Why does the inclusion of Social Security change trends in wealth inequality? First, Social Security is large: in 2019, it represented almost half of the total wealth of US households in the bottom 90% of the wealth distribution. Second, because they are long-duration cash flows, the market value of Social Security benefits increased a lot with the decline in real interest rates. Existing measures of wealth inequality focused on marketable assets register the rise in the price of long-term assets held by the richest households, such as shares of corporations, but overlook the parallel increase in the market value of Social Security, another long duration asset representing a disproportionate part of the balance sheet of low and middle-class households.

To incorporate Social Security into top wealth estimates, 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 calculate Social Security wealth from the SCF directly by valuing reported benefits as a lifetime annuity. 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 (Güvenen et al., 2021). We then assign these simulated earnings trajectories to individuals in the SCF and apply the Social Security formulas to construct estimates of future retirement benefits that these households have accrued. We validate our simulation model by comparing our results to aggregate wealth estimates reported by the SSA and to benefits reported for retirees in the SCF.

The present value of Social Security benefits depends on the choice of an appropriate discount rate. 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 9 and 6 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 are significantly attenuated: the shares of the top 10% and top 1% only increased by 1.0 and 1.5 pp, respectively.

However, the market value of retirement benefits should reflect the systematic risk of the Social Security program. Importantly, Social Security is wage-indexed, which means that the benefits of future retirees are directly tied to economic growth. Because the labor and stock markets are cointegrated over long horizons (Benzoni, Collin-Dufresne, and Goldstein, 2007), wage-indexed cash flows have a positive market beta, which increases the rate at which they should be discounted (Geanakoplos and Zeldes, 2010). We find that adjusting discount rates to account for systematic risk decreases the stock of Social Security wealth by approximately 11%, with a disproportionate effect on the bottom 90%. This is because of an age effect: Younger workers, who are disproportionately in the bottom 90%, are the furthest from retirement, so they are the most exposed to long-run macroeconomic risk. Based on our risk-adjusted valuation of benefits, the share of the top 10% and top 1% increased by 1.7 and 1.9 pp, respectively.

Our estimates of top wealth shares depend on several assumptions regarding the future of Social Security and the discount rate applied to benefits. For instance, households may value Social Security less than its fair market value because of its illiquidity. We address this in two ways: First, we apply a significant (up to 3%) liquidity premium to future benefits, which reduces aggregate Social Security wealth in half. To be conservative, we do not similarly adjust the valuation of highly illiquid marketable assets, such as private businesses, often owned by households at the top of the wealth distribution. Even under this extreme assumption, Social Security's inclusion substantially attenuates the rise in top wealth shares. Second, we allow for heterogeneity in individual discount rates, essentially using households' own cost of capital as the discount rate to apply to Social Security wealth. We infer this cost of capital from the interest rate paid on debt by borrowers at different points in the income distribution and over the life-cycle. As expected, this adjustment reduces the value of Social Security for constrained households at the bottom of the distribution relative to unconstrained households at the top. So,

relative to the market discounting specification, there is a larger rise in top 10% and 1% wealth share under this specification, but the rise in top wealth is still substantially attenuated relative to trends in marketable wealth, which exclude Social Security all together.

Future Social Security benefits are also exposed to policy risk: absent entitlement reform, the US is now less than a decade away from not being able to pay out Social Security benefits in full (Congressional Budget Office, 2023). We reflect this risk in our valuation of Social Security in a variety of ways. Our most conservative adjustment assumes the worst macroeconomic scenario from SSA actuarial projections and an indiscriminate across-the-board cut of benefits by 40%, which decreases the stock of Social Security wealth by 28.3%. But our headline fact, that Social Security's inclusion substantially attenuates the growth in top wealth shares, is unchanged: the share of wealth held by the top 10% and the top 1% would only have increased by 4.3 and 3.4 percentage points, relative to marketable wealth alone, which rose by 9.5 and 6.4 percentage points, respectively.

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, since less interests will accrue, in order to fund the same level of consumption during retirement, an investor in 2019 has to save considerably more 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. Valuing both marketable assets and Social Security claims using contemporaneous interest rates is the only consistent way to compare their incidence on households' future

consumption.

Falling interest rates affect wealth inequality by redistributing wealth away from holders of short-duration assets, favoring those with long-term investments (Auclert, 2019). Greenwald et al. (2021) document that, 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. However, they also note that the distributional effect on welfare depends on other sources of wealth and the changes in households' intertemporal budget constraint. Our paper shows that the strong link between households' marketable wealth and the average duration of their assets is largely reversed when Social Security wealth is accounted for, since it is a long-term investment representing a disproportionate share of the total wealth of the bottom 90%. As Catherine et al. (2023) show, the large implicit Social Security wealth of poorer households reduces the optimal share of long-term assets in their market wealth, so they experience 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 long-duration assets disproportionately owned by the rich but not the increased value of those owned by the rest of the population, mainly their Social Security benefits. 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.

For many questions related to inequality trends, 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. Failure to do so distorts our understanding of inequality trends: For example, proposals to let Americans invest part of their Social Security contributions in personal retirement accounts would mechanically reduce

measures of private wealth inequality, regardless of their effect on total wealth inequality.

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.¹ A more inclusive wealth concept, in contrast, helps policy-makers evaluate the role redistributive public programs play in curbing inequality.

That said, there are real limitations to the exercise that we undertake in this paper. A total wealth concept that considers marketable wealth and Social Security wealth is far from the last word on the inequality debate. We do not consider the impact of other government-provided benefits (e.g., health benefits), nor how changes in the size of those programs have impacted inequality trends. Further, another important source of wealth is one's human capital, which we do not include in this analysis.

Our findings show that Social Security represents a disproportionate share of the balance sheet of most households, which, given Social Security's unique features, raises questions regarding its private valuation and optimal program design. On the one hand, Social Security benefits cannot be used to finance consumption today or be bequeathed to heirs. On the other hand, Social Security may solve market failures by providing insurance against longevity and income risk. This is an important area for future work. Overall, we hope this paper serves as a useful interim step in illustrating the significant role of Social Security in the household balance sheet, and how Social Security's import has grown in recent history. It represents a shift toward broader wealth concepts 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

¹That is not to say there is a 1-for-1 substitution between liquid wealth and social security wealth: There are a range of estimates with respect to the elasticity of private savings rates to public savings like Social Security, for example (Attanasio and Brugiavini, 2003; Attanasio and Rohwedder, 2004; Chetty et al., 2014; Feldstein, 1974; Lachowska and Myck, 2018; Scholz, Seshadri, and Khitatrakun, 2006).

the level of and changes in measured wage inequality. Gustman et al. (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). However, recent work points out that generous welfare states that mitigate income inequality may not ameliorate intergenerational gaps in income trajectories, educational attainment, or family dynamics (Heckman and Landersø, 2022; Landersø and Heckman, 2017).

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. This fact matters for understanding trends in racial wealth inequality as well, since those at the bottom of the distribution—where Social Security is most impactful to households—are disproportionately Black and Hispanic Americans (Catherine and Sarin, 2023). 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 (2023) who point out that their measure of remaining lifetime spending is much more equally distributed than net wealth or current income.

I. Wealth Concept

This paper introduces a new time series of wealth inequality that includes the present value of Social Security benefits. A natural question to ask is when this broader wealth concept is appropriate.²

To evaluate this question, consider the various motives for wealth accumulation during one's lifetime. One is life-cycle smoothing: individuals accumulate wealth to defer consumption within a lifetime, so that they are able to continue to consume in retirement. Another is precautionary savings: wealth helps households meet liquidity needs as they arrive, for example to self-insure against unemployment or illness. Third, wealth accumulation is a means of spreading consumption across generations, as bequests build dynastic wealth that will serve as a source of consumption beyond one's lifetime.

Social Security's existence bears on these motives differently. It is obviously relevant to life-cycle smoothing: the program was designed to ensure that retirees would have sufficient resources to consume after leaving the labor force.

Social Security wealth is somewhat less relevant for precautionary savings, because, by design, Social Security is illiquid, so it cannot be used to buffer against shocks for the subset of households that are liquidity constrained.³ But so too are forms of private marketable wealth

²It is also perhaps worth noting that the progressive design of the old-age retirement program is not a substantial contributor to Social Security's impact on top wealth shares. The main driver of our results is that Social Security is a form of mandatory retirement saving that makes up a large share (and a larger share over time) of the bottom 99%'s total wealth. Accounting for it therefore substantially attenuates the rise in private wealth inequality. That is not to say Social Security's progressive design is not relevant. Indeed, part of why Social Security attenuates racial wealth inequality is that replacement rates are higher for the lowest earners, who are disproportionately minorities (Catherine and Sarin, 2023). One of the limitations of the measurement of inequality using top wealth shares is that it obscures many important dynamics within the bottom 99%.

³As Catherine, Miller, and Sarin (2020) note, this is a policy choice which guarantees that participants will not be destitute in retirement and allows the program to provide longevity insurance to retirees. In principle, policy makers can make Social Security wealth more liquid without changing the intertemporal budget constraint of the government.

that are part of baseline private wealth inequality calculations, such as shares of private corporations.⁴ Moreover, other public programs, such as Medicaid and Unemployment Insurance, complement Social Security by supporting households in times of needs. If anything, then, understanding precautionary savings calls for an even broader wealth concept, that incorporates these other sources of public wealth.

Social Security wealth also cannot be bequeathed and is therefore an imperfect substitute for market wealth in that regard. However, Social Security's existence reduces the need to consume other sources of bequestable wealth during retirement, and as such bears indirectly on the ability to transfer wealth across generations. Further, benefits that are paid out to beneficiaries and not consumed can of course be passed down. In practice, a substantial share of Social Security benefits are transferred to children (Lee and Tan, 2023; Mukherjee, 2018, 2022). The ability to substitute market and Social Security wealth during retirement is limited for workers who die prematurely, but, in this case, their children and spouse can benefit from the survivor insurance program, which paid \$141 billions to widows and children in 2022.⁵

More broadly, not all types of wealth are created equal: different types of wealth are differentially situated to meet individuals' motives for wealth accumulation. Social Security and private wealth appear close, albeit imperfect, substitutes—and better substitutes for some drivers of wealth accumulation than others. A wealth concept that includes marketable wealth and wealth from public programs, like Social Security, is therefore especially (though not exclusively) relevant to our understanding of consumption inequality during one's lifetime. That is an inquiry that is incomplete without the inclusion of public sources of wealth, because the ex-

⁴The lack of tradability also generates uncertainty regarding the value of Social Security claims, but this is also true of private business wealth. Bhandari et al. (2020) document very large discrepancies in both aggregate national business income and implied valuation ratios across survey and administrative datasets. Because measures of business wealth from administrative data depends on the capitalization factor, which is itself validated using survey data, they conclude the measurement of business wealth makes estimate of wealth inequality unreliable. Even using transaction multiples to value VC-backed companies leads to large and systematic valuation errors due to the issuance of different classes of shares (Gornall and Strebulaev, 2020).

⁵SSA Annual Report, 2023, page 38

istence of programs to provide income during retirement (like Social Security) or offset health costs (like Medicaid) results in households accumulating less private wealth to support those needs later in life.

This paper is far from the final word, particularly since it excludes other important sources of wealth, like human capital, that bear on lifetime consumption inequality. Given its narrow focus on only Social Security and private wealth, it is an interim step that illustrates the important role public wealth plays, but more work must be done to arrive at a more comprehensive wealth inequality measure.

II. 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 top of the distribution. The SCF is well suited to measure marketable wealth shares as it covers most asset classes and provides detailed information on liabilities. It is also well suited to compute the Social Security wealth of retirees as it provides detailed data on retirement and survivor benefits. 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%. Furthermore, the SCF lacks data on the present value of defined benefit pension plans. To ensure these are included, we incorporate the aggregate value of defined benefit pensions into our calculations using data from the Distributional Financial Accounts (DFA) provided by the Federal Reserve Board of Governors (Batty et al., 2019). The DFA provide data on the dollar value of the stock of defined benefit pension obligations going to the top 1%, the rest of the top 10%, and the bottom 90%. We include these data in all results where we examine wealth inequality.

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.

III. 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 for individuals in the SCF. We proceed differently for retirees and workers.

A. Retirees

For retirees, calculating Social Security wealth is relatively straightforward, since we observe their Social Security benefits in the SCF. Specifically, the Social Security wealth S of individual i in year t is:

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

where nominal benefits B_{it} are indexed to the consumer price index P_t , and m_{itk} is an individual's per-year mortality rate in k years. 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 D).

B. Workers

For workers at or below age 66, we simulate income paths for each survey year-gender-age combination. We take earnings to be the product of the Social Security wage index $L_{1,t}$ and an idiosyncratic component $L_{2,i}$:

$$L_{it} = L_{1,t} \cdot L_{2,it}. \quad (2)$$

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

$$\text{Idiosyncratic earnings:} \quad L_{2,it} = (1 - \nu_t^i) e^{(g(t) + \alpha^i + \beta^i t + z_t^i + \varepsilon_t^i)} \quad (3.1)$$

$$\text{Persistent component:} \quad z_t^i = \rho z_{t-1}^i + \eta_t^i \quad (3.2)$$

$$\text{Innovations to AR(1):} \quad \eta_t^i \sim \begin{cases} \mathcal{N}(\mu_{\eta,1}, \sigma_{\eta,1}^2) & \text{with prob. } p_z \\ \mathcal{N}(\mu_{\eta,2}, \sigma_{\eta,2}^2) & \text{with prob. } 1 - p_z \end{cases} \quad (3.3)$$

$$\text{Initial condition of } z_t^i: \quad z_0^i \sim \mathcal{N}(0, \sigma_{z,0}^2) \quad (3.4)$$

$$\text{Transitory shock:} \quad \varepsilon_t^i \sim \begin{cases} \mathcal{N}(\mu_{\varepsilon,1}, \sigma_{\varepsilon,1}^2) & \text{with prob. } p_\varepsilon \\ \mathcal{N}(\mu_{\varepsilon,2}, \sigma_{\varepsilon,2}^2) & \text{with prob. } 1 - p_\varepsilon \end{cases} \quad (3.5)$$

$$\text{Nonemployment duration:} \quad \nu_t^i \sim \begin{cases} 0 & \text{with prob. } 1 - p_\nu(t, z_t^i) \\ \min\{1, \text{Exp}\{\lambda\}\} & \text{with prob. } p_\nu(t, z_t^i) \end{cases} \quad (3.6)$$

$$\text{Prob. of Nonemp. shock:} \quad p_\nu^i(t, z_t) = \frac{e^{a+bt+cz_t^i+dz_t^i t}}{1 + e^{a+bt+cz_t^i+dz_t^i t}} \quad (3.7)$$

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 (3.1), $g(t)$ 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 . The parameters values are shown in Table IAI.

We then simulate 10,000 earnings paths for each survey year-gender-age combination detailed in equations (3.1)-(3.7). We then match these paths to individuals in the SCF based on

their current wage income, gender, and age.⁶ We then apply the Social Security benefit and tax formulas to compute the average Social Security wealth by cohort, gender, and year. 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. A more comprehensive description of the individual assignment procedure is provided in Appendix C.

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

$$T_{it} = .106 \cdot \min \{L_{it}, SSWB_t\}. \quad (4)$$

Benefits. Benefits depend on each individuals' average indexed yearly earnings (AIYE). A worker's indexed taxable earnings in year t are:

$$L_{it}^{\text{indexed}} = \min \{L_{it}, SSWB_t\} \frac{L_{1,60}}{L_{1,t}}, \quad (5)$$

where $L_{1,60}$ and $L_{1,t}$ denotes the value of the national wage index the year of his 60th birthday.

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 L_{it}^{indexed} , as defined in equation (5). Annual 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 piecewise linear function of AIYE:

$$B_{it} = \begin{cases} \frac{P_t}{P_{c_i+60}} \cdot .9 \cdot \text{AIYE}_i & \text{if } \text{AIYE}_i < b_{1,c_i} \\ \frac{P_t}{P_{c_i+60}} [.9 \cdot b_{1,c_i} + .32(\text{AIYE}_i - b_{1,c_i})] & \text{if } b_{1,c_i} \leq \text{AIYE}_i < b_{2,c_i} \\ \frac{P_t}{P_{c_i+60}} [.9 \cdot b_{1,c_i} + .32(b_{2,c_i} - b_{1,c_i}) + .15(\text{AIYE}_i - b_{2,c_i})] & \text{if } b_{2,c_i} \leq \text{AIYE}_i. \end{cases} \quad (6)$$

⁶The ideal matching procedure would consider both past and current income to match future income paths to individuals. Starting in 1995, the SCF inquires if the current year's income is similar to the income received in past years. We also match the simulation using this past income data and find minimal differences compared to the method that uses only current wage income.

where $\frac{P_t}{P_{c_i+60}}$ is an adjustment for the increase in the consumer price index since the retiree turned 60.⁷ After retirement, benefits grow at the rate of inflation.

Social Security wealth. To value Social Security wealth, our baseline specification is focused on the value of future benefits that households have accrued based on payments they have already made to the program through payroll taxes, expressed by:

$$S_{it} = \frac{\text{Past Taxes}_{it}}{\text{Past Taxes}_{it} + \text{Future Taxes}_{it}} \sum_{s=t+1}^T \frac{\mathbb{E}[B_{is}]}{(1 + r_{ts})^{s-t}}. \quad (7)$$

where future taxes are defined by

$$\text{Future Taxes}_{it} = \sum_{s=t+1}^T \frac{\mathbb{E}[T_{is}]}{(1 + r_{ts})^{s-t}} \quad (8)$$

and past taxes by

$$\text{Past Taxes}_{it} = \sum_{s=t-a}^t \frac{T_{is}}{(1 + r_{ts})^{s-t}}. \quad (9)$$

Note, that the exponent in the denominator on past taxes is negative, so these taxes are grossed up by the realized nominal return on 30-year treasury bonds from $t - a$ to t .⁸

Although this is not the only approach to calculating the present value of accrued Social Security benefits, we choose it for our baseline because of its conceptual appeal. At the start of their careers, workers owe a flow of taxes in exchange for their future Social Security benefits. To put these tax payments over time into similar units, we discount future taxes at the appropriate rate and gross up past taxes by the realized return on the 30-year nominal treasury bond. This also means that in years when workers make significant tax contributions, they accrue significant benefits. More generally, our approach reflects the structure of the program, which

⁷In practice, workers can claim benefits as early as age 62 or as late as age 70. However, this option is, on average, 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).

⁸The realized nominal return on 30-year treasury bonds is taken from CRSP US Treasury and Inflation Indexes.

is designed for workers to contribute to Social Security during their working years in exchange for benefits in retirement that are derivative of those contributions.

In practice, the results using this method are nearly identical to two other reasonable alternatives: calculating benefits assuming workers stop working and accruing benefits *pro rata* with age. These alternatives are discussed in detail in Section F.

Another alternative approach would be to construct a value of Social Security based on the present value of future benefits less future taxes, expressed by:

$$S_{it} = \sum_{s=t+1}^T \frac{\mathbb{E}[B_{is} - T_{is}]}{(1 + r_{ts})^{s-t}} \quad (10)$$

An NPV-based approach would be consistent with the way that shares of stock, for example, in household portfolios are generally valued in the wealth inequality literature (Saez and Zucman, 2016). Private business stakes are valued based on an expectation of future cashflows, too. But an NPV-based approach to the valuation of Social Security would mean including cashflows derivative of future years in the labor force for the purpose of the old-age retirement program, but not for the purpose of including human capital in our wealth concept in a way that feels conceptually at odds.

As such, our baseline specification values Social Security based on benefits households have already accrued and we explore the NPV-based approach in Section E.

C. Calibration

Lifetime income profiles. We assume $g(t)$ to be cohort and gender-specific. Guvenen et al. (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: $L_{2,cgt}$. Then, we estimate $g_{cg}(t)$ by regressing $\ln(L_{2,cgt})$ 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 ($SSWB_t = 2.5 \cdot L_{1,t}$), and the bend points of the benefits formula will remain 0.21 and 1.25 multiplied by 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 E for details).

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 II. 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 A.

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, Friedman, Leth-Peterson, Nielsen, and Olsen, 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.⁹ Our adjustment effectively makes high income retirees younger and low income retirees older. However, this adjustment has little impact on the final results. Appendix C offers a more detail description of this procedure and its consequences.

⁹We proxy for the permanent income distribution using the Social Security benefits distribution because benefits are, by construction, a proxy for lifetime earnings.

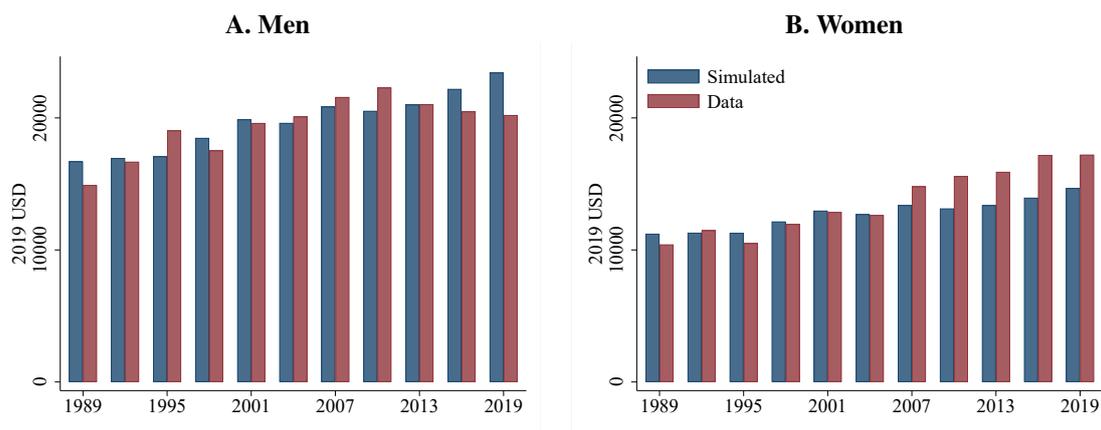
D. Validation

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 1, 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). The simulated data track observed benefits closely.

Matching SSA estimates of aggregate Social Security wealth. Every year since 1996, the SSA estimates the present value of accrued benefits of people currently paying Social Security taxes or receiving benefits. 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 secular decline in interest rates. As reported in Appendix Figure IA5, the SSA discount rates fell by only 2 pp between 1989 and 2019 while the market yield curve fell by three times that amount.

Figure 1. Simulated and actual full retirement age benefits



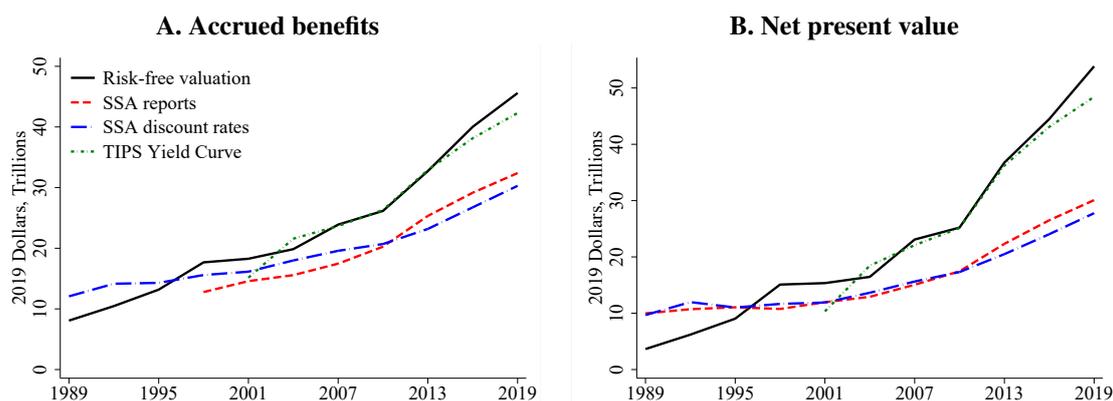
Nevertheless, Panel A of Figure 2 shows that, if we choose to use the SSA's discount rates,

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 just over \$30 trillion, compared to approximately \$45 trillion when using market rates.

Panel B of Figure 2 shows the same series but using the net present value wealth concept. One advantage of using this concept in validation is that the SSA produce the aggregate net present value of Social Security benefits back to 1989. Our net present value series discounted at the SSA’s discount rates also matches their series closely.

Figure 2. 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) for both their accrued benefits and net present value approach. 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.



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

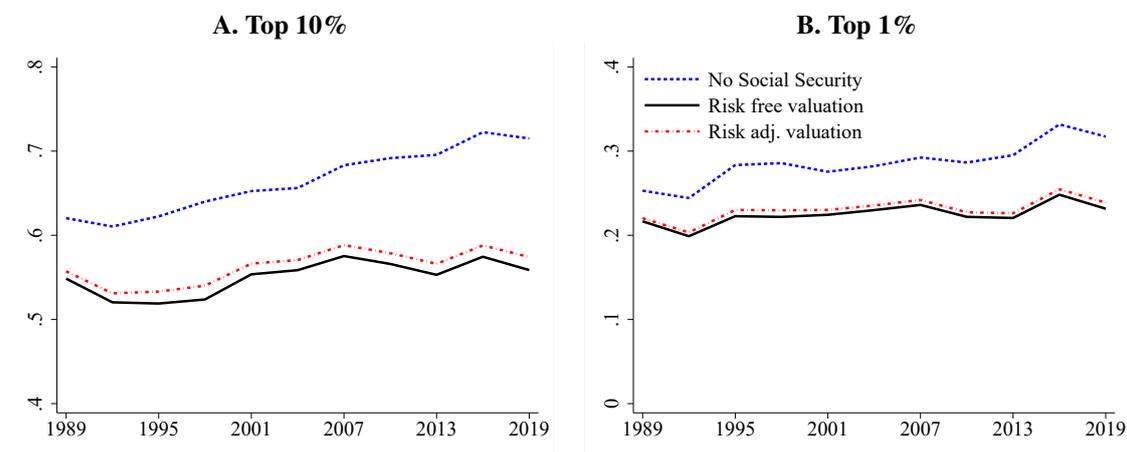
we take from SSA annual reports. To make sure that our results are not driven by these assumptions, we also discount future cash flows using the real yield curve implied by the price Treasury Inflation Protected Securities (TIPS) and assume no inflation. This exercise can only be done for the 1999-2019 period. As reported in Figure 2, this alternative methodology implies a faster increase in aggregate Social Security wealth than ours;¹⁰ as such, our findings are not driven by challenges with forecasting inflation.

E. Baseline Top Wealth Shares

Figure 3 reports the levels and trends of top wealth shares with and without Social Security wealth. We define top wealth shares based on the top 10% and top 1% of the population by measures of wealth concept being examined. This means that we rank according to marketable wealth when examining marketable wealth inequality and total wealth when examining total wealth inequality.

Figure 3. 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 A.



¹⁰There 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).

Panel A focuses on the top 10%. The top 10% share of market wealth grew by 9.5 pp between 1989–2019. Once Social Security wealth is included, this trend is substantially attenuated. Panel A shows the top 10% wealth share, which now only rises by 1.0 pp over this period. Similarly, 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 grown by 6.4 pp over our sample period. Once it is included, the top 1% share has risen by 1.5 pp.

IV. Accounting for macroeconomic risk

The rate of return of pay-as-you go systems is tied to 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 B, 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.

A. 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, Collin-Dufresne, and Goldstein (2007). This 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:

$$dl_{1,t} = \left((\phi - \kappa)y_t + \mu - \delta - \frac{\sigma_l^2}{2} \right) dt + \sigma_l dz_{1,t}, \quad (11)$$

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

$$ds_t = \left(\mu + \phi y_t - \frac{\sigma_s^2}{2} \right) dt + \sigma_s dz_{2,t}, \quad (12)$$

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 according to:

$$dy_t = -\kappa y_t + \sigma_l dz_{1,t} - \sigma_s dz_{2,t}, \quad (13)$$

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:

$$\beta_t^{L_{1,n}} = \left(1 - \frac{\phi}{\kappa} \right) (1 - e^{-\kappa n}), \quad (14)$$

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

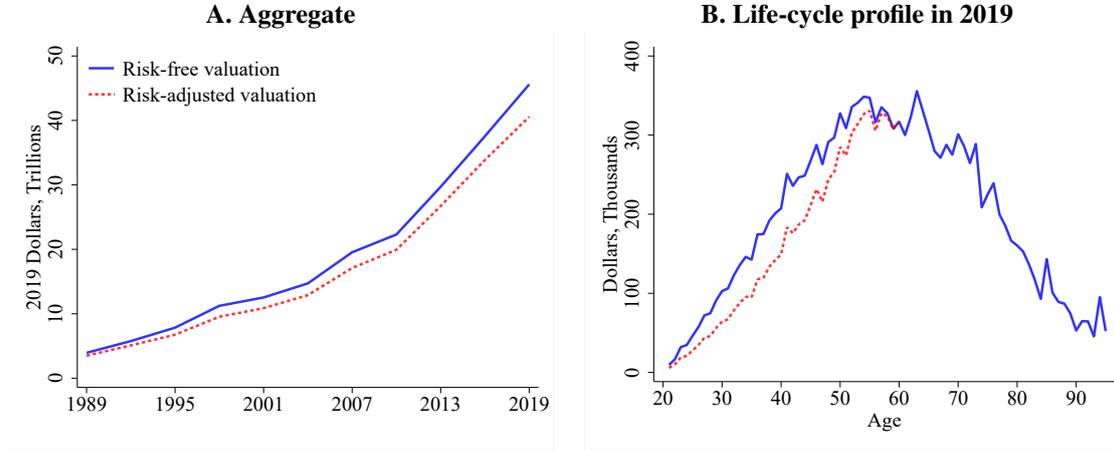
$$E_t \left[r_t^{L_{1,n}} \right] = \beta_t^{L_{1,n}} (\mu - r) + r \quad (15)$$

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

Figure 4. Risk-adjusted valuation

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



the discount factor for a cash flow proportional to $L_{1,n}$ paid in year k is:

$$\chi_{t,n,k} \approx \left[\prod_{s=t}^n (1 + \beta_s^{L_{1,n}} (\mu - r) + r_{ts}) \prod_{s=n+1}^k (1 + r_{ts}) \right]^{-1}. \quad (16)$$

The risk-adjusted present value of Social Security is then given by equation (7) with equation (16) used as the discount rate on benefits and future taxes. The parameter n here is equal to 60 less the current age, because real benefits are indexed on L_1 in the year in which the worker turns 60.

We calibrate the model as in Benzoni, Collin-Dufresne, and Goldstein (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 the most distant indexed cash flows. The equity premium is assumed to be $\mu - r = .06$.

B. Risk-adjusted Results

Panel A of Figure 4 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 over 40%.

Once macroeconomic risk associated with Social Security cashflows is factored in, Figure 3 shows that the share of the top 10% and top 1% increased by 1.7 and 1.9 pp, respectively.

This finding differs from our baseline risk-free specification because Social Security wealth is smaller, and therefore plays less of a 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.

V. Discussion

A. *Factors Contributing to Social Security's Growth*

Table I lays out the contributors to the growth in Social Security wealth. 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 which drives 45.8% of Social Security's growth (48.3% with risk-free valuation).

A.1. *Shifts in the Interest Rate Environment*

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.

For working-age households, Social Security benefits are disbursed years into the future and so can be replicated by a portfolio of long-term bonds. In Table II, we show how the present value of benefits has increased over our sample for the top 1% and the bottom 99%. 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 third column

Table I. 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	0.843	0.802
Shift in age distribution	0.157	0.184
Life expectancy	0.120	0.121
Social Security expansion & other	0.302	0.319
Log total per capita	1.423	1.426
Population growth	.323	.323
Log total	1.746	1.749

of Table II shows, the value of Social Security has increased approximately 120 pp more for the bottom 99% than for those in the top 1%. This is an age effect; the present value Social Security grows more with declines in interest rates for younger workers who are disproportionately in the bottom 99%. Offsetting this, to some degree, is the change in the amount of benefits individuals have accrued. When interest rates decline, the present value of future taxes grow, which, all else equal, reduces the portion of benefits accrued. We can see the effect in the first row where the bottom 99% have accrued only 73% of their benefits, on average, relative to just over 74% in 1989. At the same time, the US population has aged substantially since 1989, which should increase the portion of benefits accrued. We can see that this effect dominates for the top 1% in the second row of Table II.

Even if it is primarily driven by interest rate changes, the evolution of Social Security wealth matters for our understanding of inequality trends. For one, it is important when comparing the rise in Social Security wealth to the rise marketable wealth inequality of which the decline

Table II. Impact of interest rates on Social Security wealth – Risk-adjusted valuation

This tables decomposes the increases in Social Security wealth between 1989 and 2019. The first three columns present the rise in the present value of all benefits, both accrued and not yet accrued. The next three columns present the change in the portion of benefits accrued.

	Present value of all benefits (2019 dollars)			Percent of benefits accrued		
	1989	2019	Change	1989	2019	Change
Bottom 99%	58,103	268,614	362%	74.4	73.0	-1.4
Top 1%	93,942	327,995	249%	88.9	91.3	2.3
Full Population	58,460	269,207	360%	74.6	73.2	-1.3

in interest rates is also a primary driver (Greenwald et al., 2021). To understand why this is necessary for an apples-to-apples comparison, 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.

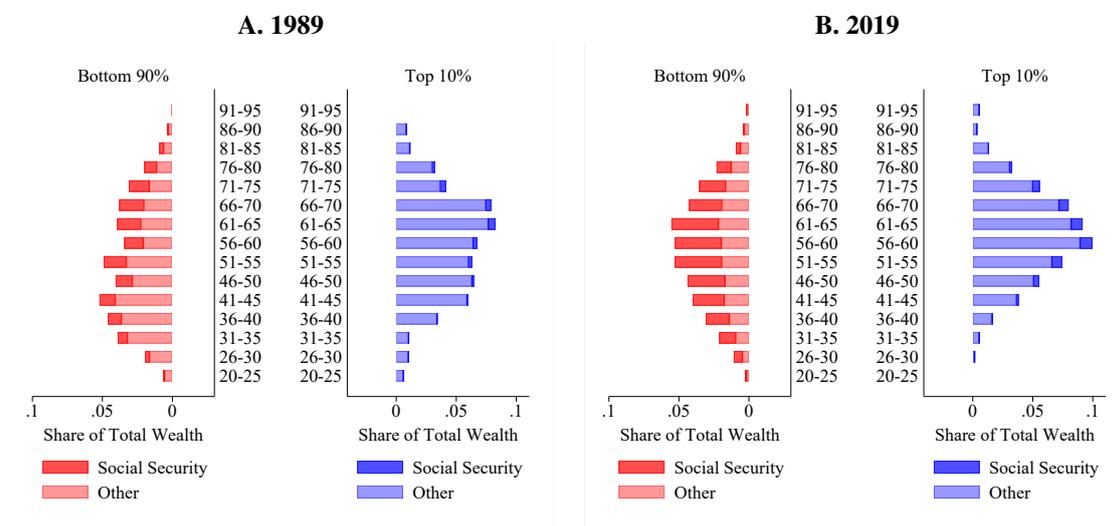
If the interest rate is 0.8%, as in 2019, this household needs to save \$0.85 annually. Now, if the rate of return on Social Security contributions has been approximately constant at 2.5%, which it has since 1989, the same can be achieved by contributing approximately \$0.61 every year to Social Security. 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.¹¹¹²

¹¹The 0.8% rate in 2019 and 5% rate in 1989 correspond to the 20-year real forward rate less projected inflation from the SSA. The internal real rate of return (IRR) on Social Security comes from the SSA Actuarial Note 2019.5 (Clingman, Burkhalter, and Chaplain, 2020). The IRR for a middle income couple born in 1949 (the 40-year olds in 1989) is 2.61%. The same couple born in 1973 (data for 1976 is not provided) has an IRR of 2.79%. For simplicity, we round to 2.5%.

¹²Sabelhaus 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.

Figure 5. 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.



B. Shifts in the Composition of Wealth

Figure 5 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 26.0% of the total wealth of the bottom 90%. In 2019, Social Security represents 49.8% 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.

VI. 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; potential illiquidity discounts; using heterogeneous discount rates that account for individuals own cost of capital; and weak economic growth. Table III presents results using alternative assumptions, which we discuss in turn below.

Table III. Robustness checks

Panel A reports our baseline results. First, we report top shares of marketable wealth in the SCF. 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. Third, we present a specification that assigns heterogeneous discount rates to individuals according to their own cost of capital. Fourth, we provide the change in top wealth shares under the net present value wealth concept, calculated as in equation (10). 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	62.0	71.5	9.5	25.3	31.7	6.4
Risk-free valuation	54.9	55.9	1.0	21.7	23.2	1.5
Risk-adjusted valuation	55.7	57.4	1.7	22.0	23.9	1.9
Panel B: Funding Gap						
Benefit cut (Intermediate Cost)	55.7	58.8	3.1	22.0	24.6	2.6
Benefit cut (High Cost)	56.4	60.7	4.3	22.3	25.7	3.4
Tax hike (Intermediate Cost)	55.7	57.6	1.9	22.0	23.9	1.9
Tax hike (High Cost)	56.2	58.2	2.0	22.2	24.3	2.1
Panel C: Robustness						
Declining wage growth	56.0	58.5	2.5	22.2	24.4	2.2
Liquidity premium (1%)	56.5	59.4	2.9	22.4	25.0	2.6
Liquidity premium (2%)	57.2	61.1	3.9	22.7	25.9	3.2
Liquidity premium (3%)	57.7	62.4	4.7	23.0	26.6	3.6
Heterogenous discount rates	56.4	60.6	4.2	22.3	25.5	3.2
Heterogenous discount rates and benefit cut	56.4	61.5	5.1	22.3	26.0	3.7
Net present value	60.2	57.2	-3.0	24.0	23.8	-0.2

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 9.5 and 6.4 pp respectively between 1989–2019. Once Social Security is included, using our most conservative set of assumptions, the top 10% and 1% shares grow by only a fraction of that over this horizon.

A. 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 cash flows 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.¹³

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. Figure 6 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 1.8 pp, wiping out a quarter of Social Security’s impact. But as Figure 7 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 represents a sizable sum relative to their marketable wealth.

¹³Note that here we only cut the benefits that households receive or increase the taxes they pay and do not increase the discount rates that they may face. The reason for this is that in each scenario we assume that the policy change is deterministic such that there is no increase in uncertainty or risk. This is a more conservative assumption than placing some distribution over potential outcomes and increasing the discount rate to account for this increased policy risk.

Figure 6. 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.

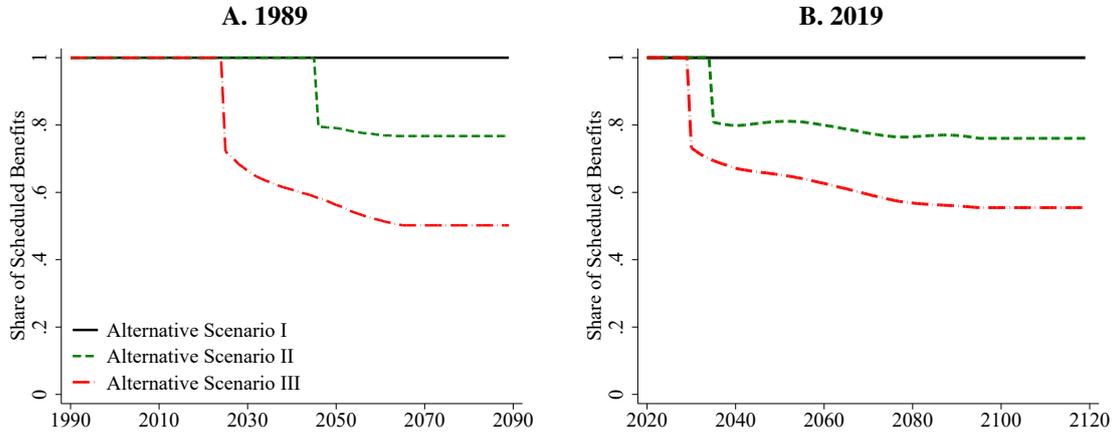
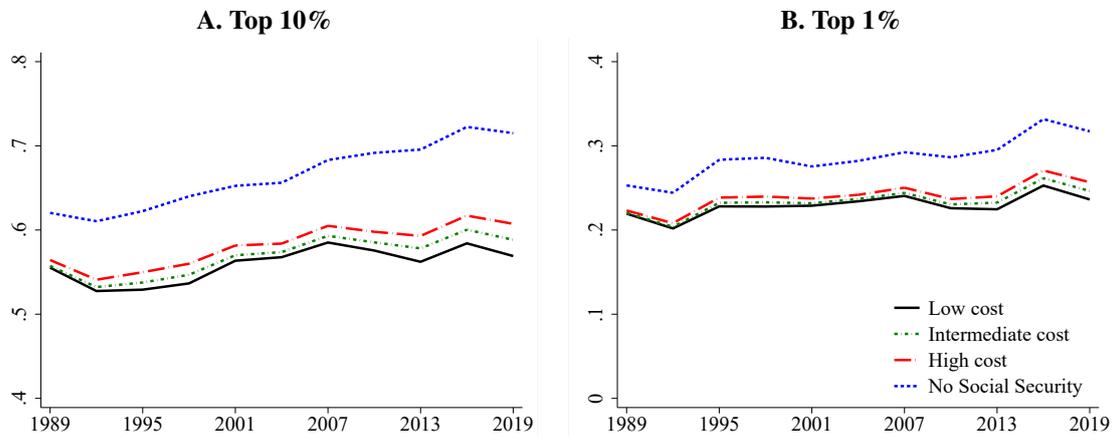


Figure 7. Top 10% and Top 1% wealth shares — Funding gap adjustment

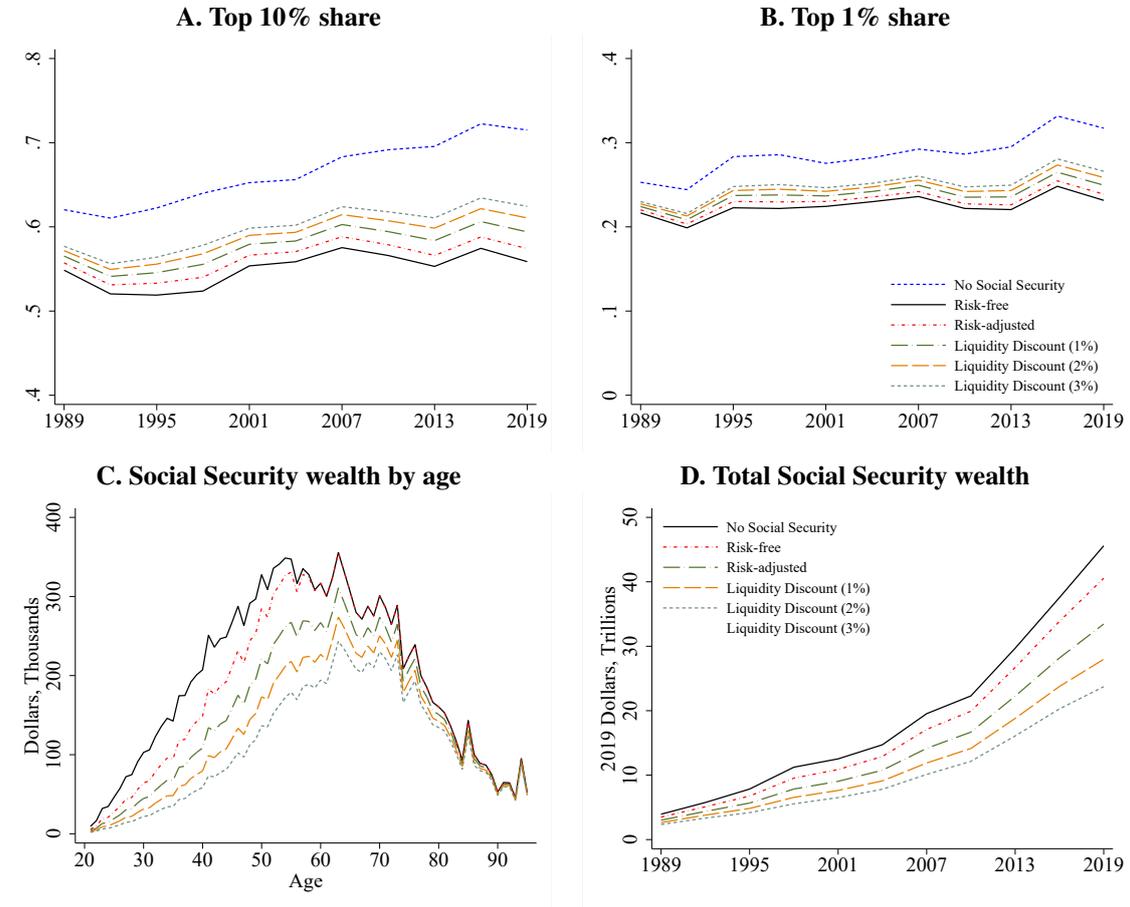
This figure presents top 10% and 1% wealth shares under four, risk-adjusted specifications. The “Low cost” specification refers to the SSA’s high economic growth scenario in which benefits are fully paid. In the “Intermediate cost” and “High cost” specifications, benefits are cut to match expected tax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



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% and top 1% shares still both rise by

Figure 8. 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 use the risk-adjusted valuation method.



only approximately 2.0 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.

B. Adjusting for the Illiquidity of Social Security

Because it is illiquid before retirement, financially constrained households may value accrued Social Security benefits less than an unconstrained and diversified investor. Although the private valuation of non-tradable assets cannot be observed, we evaluate the robustness of

our findings through two exercises. First, through comparative statics, we show how trends in inequality vary when we apply various liquidity premia to our risk-adjusted discount rates. Second, we try to infer which households are financially constrained in the SCF and estimate their opportunity cost of capital to build alternative discount rates.

In Figure 8, we show how applying liquidity discounts of a 1%, 2%, and 3% to Social Security wealth for all households changes our main results. Panels C and D show that applying a 3% illiquidity premium reduces aggregate Social Security wealth by half. 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 4.7 and 3.6 pp respectively, instead of 9.5 and 6.4 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.

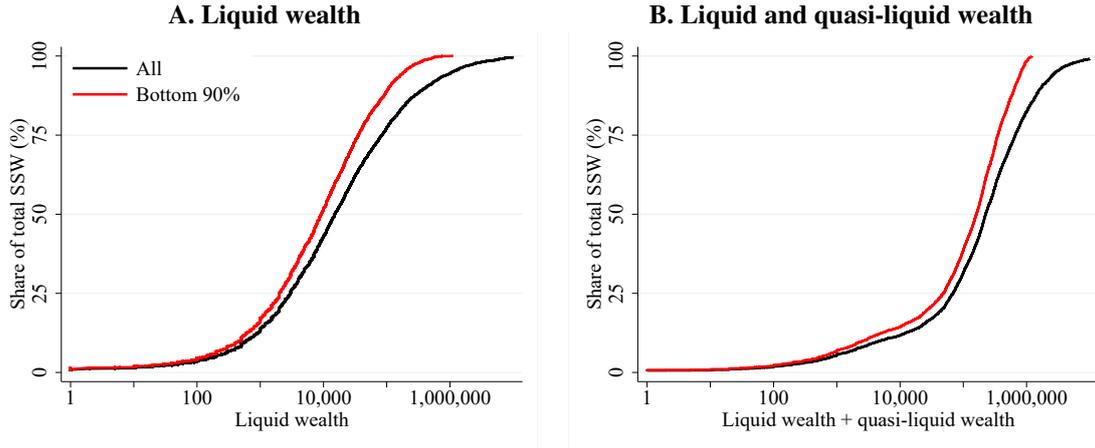
Whether we should apply a uniform liquidity premium is unclear. First, Social Security should be considered as a part of a broader welfare systems since other public programs, such as unemployment insurance, Medicare, Medicaid and the disability insurance program, provide liquidity to households facing unexpected hardship. Second, a significant share of Social Security wealth goes to households with ample liquid wealth, even within the bottom 90%. Panel A of Figure 9 shows that approximately 50% of Social Security wealth is held by households with greater than \$10,000 in liquid wealth. This is even more stark when we consider less liquid forms of wealth like quasi-liquid retirement accounts and home equity: Panel B of Figure 9 shows that 50% of Social Security wealth is held by households with greater than \$100,000 of quasi-liquid wealth.

C. Adjusting for Heterogeneity in Individual Discount Rates

So far, we have valued Social Security wealth from the perspective of a well-diversified investor using market prices or assumed a uniform liquidity premium for all households. However, as Figure 9 illustrates, such a premium would not equally apply to all households. First, because many unconstrained households should value Social Security like diversified investors.

Figure 9. Portion of Social Security wealth accruing to households with liquid wealth

This figure shows the share of Social Security wealth in 2019 held by households with less than $\$X$ dollars of liquid wealth (Panel A) and liquid plus quasi-liquid wealth (Panel B) for the bottom 90% and all individuals, where X is varied across the X-axis. Liquid wealth includes checking, savings or money market account, certificates of deposit, directly held mutual funds, stocks or bonds. Quasi-liquid wealth also includes quasi-liquid retirement accounts and home equity. Appendix B provides more details on the construction of these variables. Individual liquid and quasi-liquid wealth is obtained in the SCF by splitting wealth equally between each member of two person households.



Second, because among those constrained, the opportunity cost of capital is likely to vary across the income and wealth distribution.

In this section, we approximate the private discount rate of households using their own opportunity cost of capital. In this exercise, we distinguish two groups. First, unconstrained households, which we define as having no debt, and, either more than $\$10,000$ in liquid assets or more than $\$50,000$ in illiquid assets. We assume their opportunity cost of capital to be equal to that of diversified investors:

$$f_{h,a,q,t}^{\text{unconstrained}} = f_{t,h}^{\text{risk-adj}} \quad (17)$$

where f represents the risk-adjusted forward rate in time t at horizon h .

Second, constrained households, for which we assume the opportunity cost of capital to be the interest rate on debt. In computing the opportunity cost of capital, we take into account that the cost of debt varies over time, with income and age, and the probability of exiting the constrained state. Specifically, the forward rate of constrained investors at any horizon h , age

a , in within-cohort income quintile q , and in survey year t is:

$$f_{h,a,q,t}^{\text{constrained}} = \begin{cases} f_{t,h}^{\text{risk-adj}} & \text{with probability } p_{a+h-1,q,t} \\ f_{t,h}^{\text{risk-free}} + s_{a+h-1,q,t} & \text{with probability } 1 - p_{a+h-1,q,t} \end{cases}. \quad (18)$$

where $f_{t,h}^{\text{risk-free}}$ is the real forward rate on government bonds, $f_{t,h}^{\text{risk-adj}}$ is the forward rate from our risk-adjusted specification, and $s_{a+h-1,q,t}$ is the average interest rate spread over the safe rate paid by borrowers in income quintile q . Finally, $p_{a+h-1,q,t}$ is the conditional probability of being unconstrained in h years. Hence, we apply different discount rates for SCF observations depending on their age, income, year and whether they are currently categorized as constrained or not.

Empirically, we use the SCF to measure households' interest rate spread $s_{a,q,t}$ and their probability of remaining borrowers over time. First, for each earnings quintile, we model s as a cubic polynomial function of age and a survey year fixed effect, allowing the cost of debt faced by each income quintile to vary differently over the life-cycle and over time. We estimate these functions using Tobit regressions and using each quintile's average balance-weighted interest rate on debt, in excess of the yield on 5-year treasury bonds. Finally, we set the transition probabilities out of the constrained state to match the decline in the share of constrained households over the life-cycle and in each income quintile. Appendix F offers a more detailed step-by-step description of our methodology.

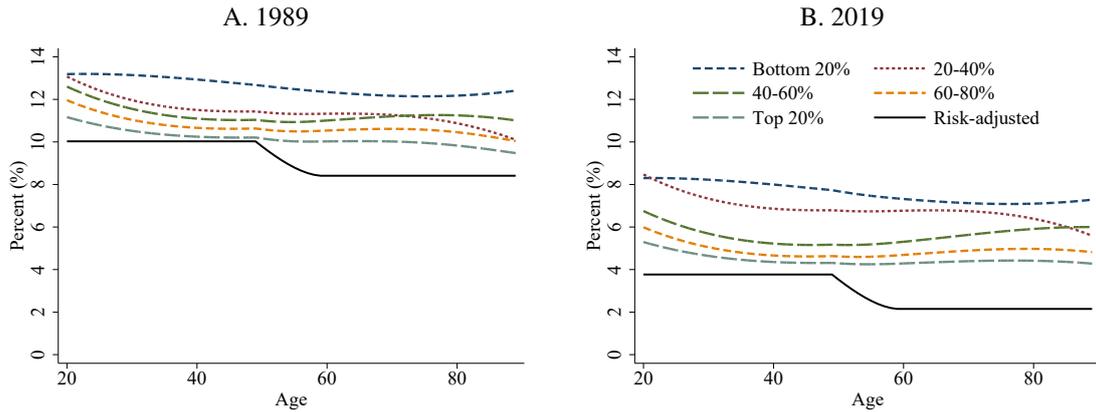
Panel 1 of Figure 10 plots the resulting 10-year yields at different ages and for different income quintiles for constrained households in 1989 and 2019. As expected, our methodology implies significantly higher discount rates for constrained households than the risk-adjusted specification. Moreover, constrained households in lower income quintiles face higher interest rates on debt, and therefore have higher discount rates. Interestingly, our methodology yields a similar distribution of discount rates as Samwick (1998), early work that takes a structural approach to estimate discount rate heterogeneity to study Social Security reform.

Panel 2 plots the portion of Social Security wealth remaining in each wealth decile for constrained households after the heterogeneous discounting is applied relative to the riskfree

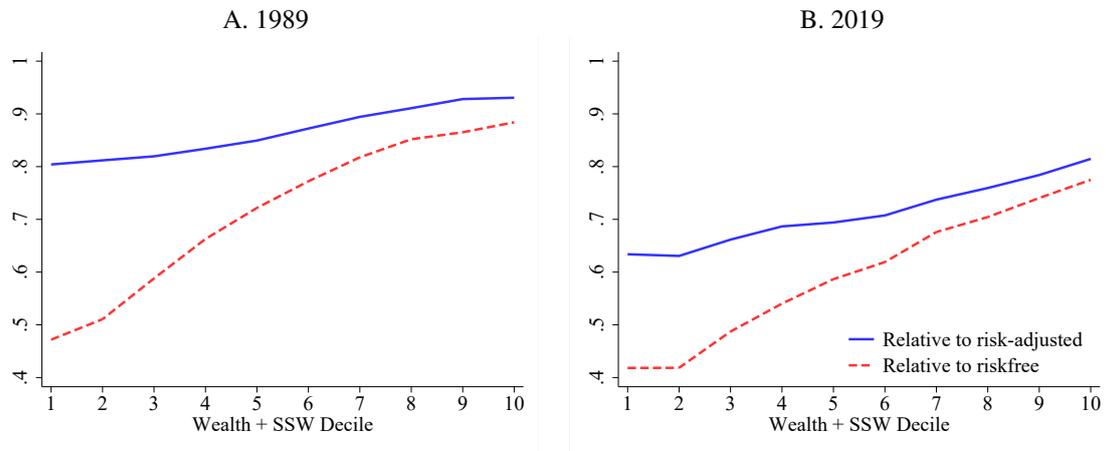
Figure 10. 10-year yield under heterogeneous discounting

This figure shows the 10-year yield used to discount Social Security benefits for each within-cohort earnings quintile and age in 1989 and 2019 calculated from equation (IA9). Panel 1 displays yields for constrained households only. Panel 2 displays the weighted-average yields across constrained and unconstrained households calculated from Equation (IA10). For comparison, the yield used in the risk-adjusted specification is plotted as well.

1. Constrained households



2. Social Security wealth remaining after heterogeneous discounting

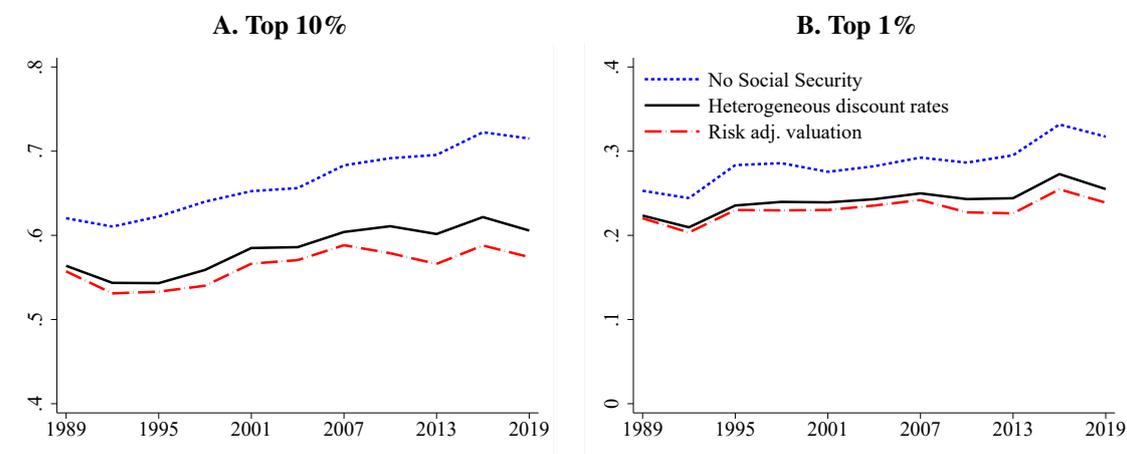


and risk-adjusted specification. The effect of the discounting is stark. Relative to the risk-adjusted specification, households in the first decile lose 20% of their Social Security wealth in 1989. This increases to over 35% as we move to 2019. The increase in the portion of benefits lost from 1989 to 2019, is one of the main reasons heterogeneous discounting affects trends in wealth inequality.

Figure 11 presents how allowing for household-specific discount rates affects the levels and trends in wealth inequality. Relative to the risk-adjusted valuation, there is a larger rise in top

Figure 11. Top 10% and Top 1% wealth shares — Heterogeneous discount rates

This figure shows the top 10% and top 1% wealth shares under individuals private valuations of Social Security wealth.



10% and 1% wealth shares, 4.2 and 3.2 percentage points. However, the rise in top wealth shares is still approximately half of that observed in the specifications without Social Security wealth included. There are three reasons why using heterogeneous discount rates affect our results. First, because discount rates are higher for all groups, aggregate Social Security wealth is lower, mechanically reducing its impact on wealth inequality. Second, spreads were larger in 2019 than 1989, which counteracts the effect of declining interest rates. Third, the difference between the opportunity cost of low and high-earners increased, slightly reducing the redistributive nature of Social Security wealth.

We explore how this result is impacted by 50% higher and lower the discount rate spreads for all constrained households in Appendix Section F. Even with the higher spreads, the rise in wealth inequality is still substantially attenuated, rising by 5.2 and 3.7 percentage points for the top 10% and 1% respectively. Under these higher spreads, the rise in wealth inequality is approximately 60% as large as is it without Social Security wealth.

While this exercise illustrates the robustness of our findings, we want to be careful when considering what it means from a policy perspective. Our analysis misses some of the aspects of the Social Security system that would tend to make private valuations higher, for example, by providing longevity insurance and insurance against idiosyncratic income shocks (Catherine, 2019). A more complete private valuation exercise that accounts for all of these positive

channels is beyond the scope of this paper.

D. 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% and 1% shares now increase by 2.5 and 2.2 pp.

E. Social Security Wealth as the Net Present Value of Benefits and Taxes

To this point, we have defined Social Security wealth as the present value of the portion of benefits that have been accrued from past taxes paid. However, an alternative definition is to take Social Security wealth as the net present value of expected benefits less taxes. As discussed above, there are many good reasons to use this valuation concept for Social Security, chief among them that this is how, in theory, other assets are valued. The net present value approach, therefore, allows for a more apples-to-apples comparison with other sources of wealth.

The results using this approach relative to the accrued benefits concept are presented in Figure 12. Under the net present value concept, the top 10% and top 1% wealth share slightly declined by 3.0 and 0.2 pp. The net present value valuation approach also has a larger effect on top wealth shares than the accrued benefits approach. One reason for this is that under the net present value approach Social Security wealth can be negative. Indeed, this was the case in the high interest rate environment of the late 1980s for individuals just entering the labor force. Because of this, the NPV-based approach starts from a lower aggregate value of Social Security; it rises to a similar level as the accrued benefits approach between 1989-2019, and as such, its impact on inequality trends is slightly larger.

From a slightly more technical point of view, this is because the net present value wealth concept is, in essence, levered 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. Because benefits (the long position) have a longer duration, when rates fall, their present value rises faster than that of taxes (the short position). The result is a rapid increase in the net present value of Social Security.

F. Alternative Measures of Accrued Benefits

This section presents the results for two additional accrued benefits specifications. The first approach assumes that benefits are accrued *pro rata* with each year of work. Since we assume individuals begin working at 20, this means that Social Security wealth is given by

$$S_{it} = \frac{a - 19}{43} \sum_{s=t+1}^T \frac{\mathbb{E}[B_{is}]}{(1 + r_{ts})^{s-t}}. \quad (19)$$

where a is their current age.

The second approach assumes individuals stop working such that they owe no future taxes,

Figure 12. Top 10% and Top 1% wealth shares — Net present value

This figure shows the top 10% and top 1% wealth shares with and without the risk-adjusted value of Social Security wealth under the net present value wealth concept.

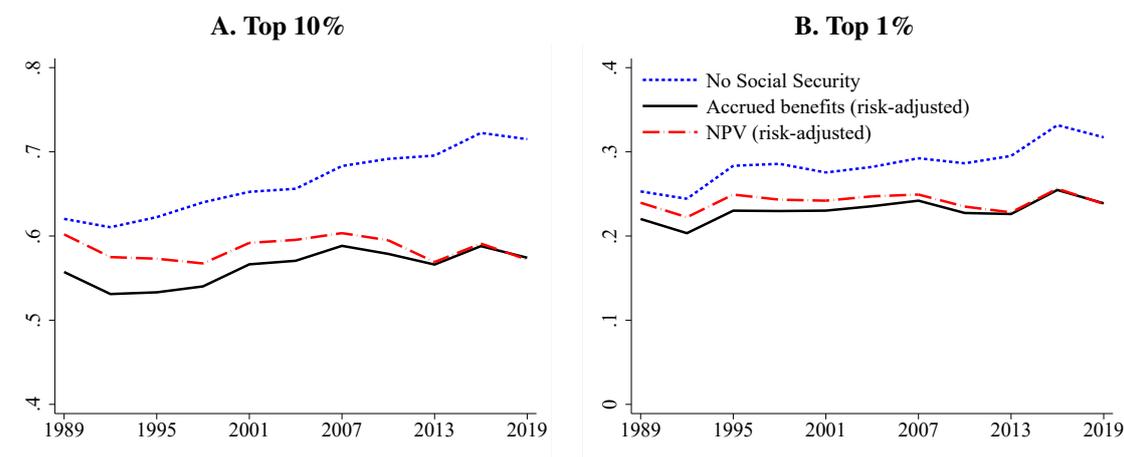
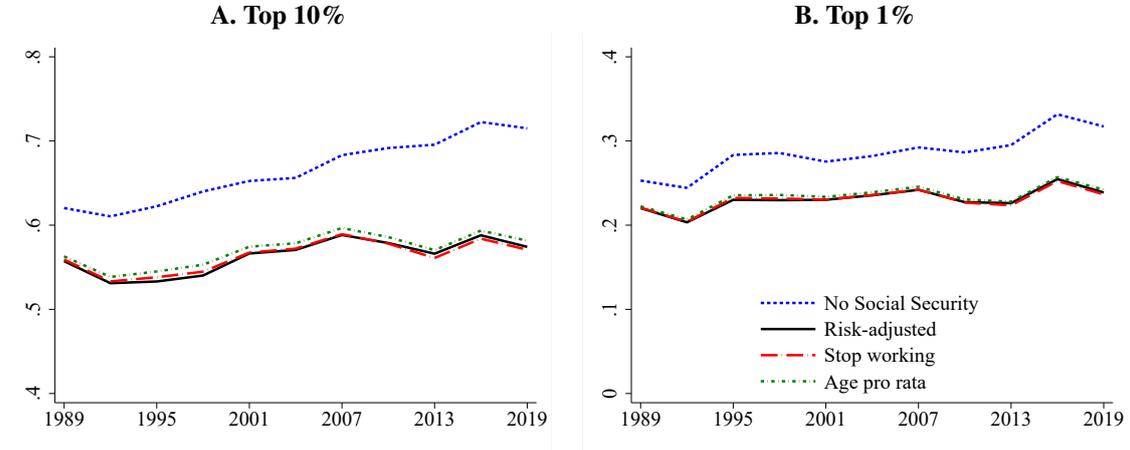


Figure 13. Top 10% and Top 1% wealth shares — Alternative measures of accrued benefits

This figure shows the top 10% and top 1% wealth shares for three different measures of accrued benefits: accrued benefits as a portion of future taxes, assuming individuals stop working tomorrow, and assigned pro rata based on working years.



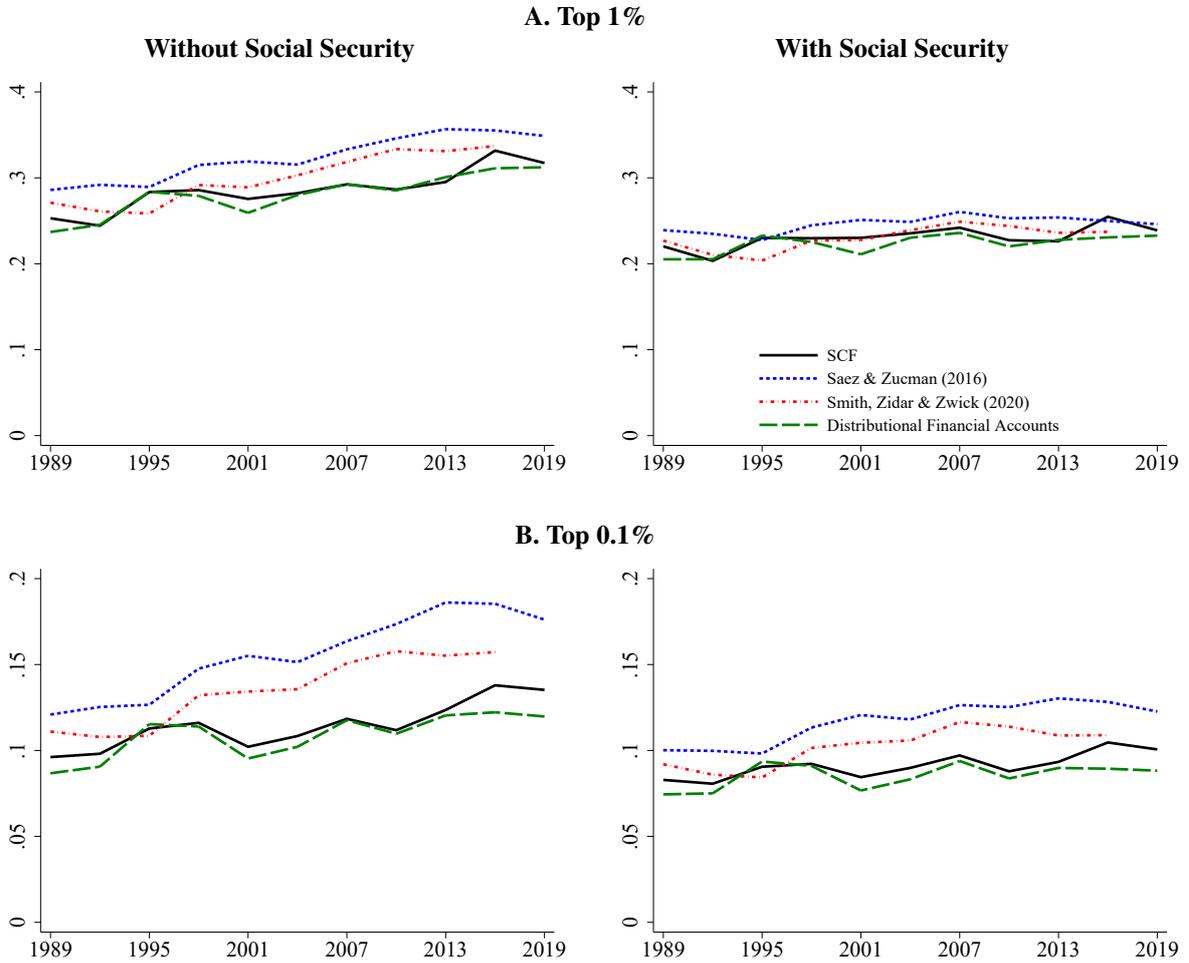
meaning Social Security wealth is given by

$$S_{it} = \sum_{s=t+1}^T \frac{\mathbb{E} [B_{is}]}{(1 + r_{ts})^{s-t}}. \quad (20)$$

This approach has two problems. The first issue arises from the curvature of the benefits function. Because social security is progressive, a larger fraction of benefits are replaced for poorer workers. This means replacement rates are higher under this concept, lead us to overstate Social Security wealth. The second issue is that workers need at least 10 years of work history to receive Social Security benefits. This would create a discontinuity whereby young workers have no expected benefits. In the results, we do not consider this discontinuity to maintain consistency with the other specifications.

Figure 13 presents the results with these different approaches. All three approaches give very similar results, with the stop working specification producing the largest decline. This shows that the results are robust to whichever definition of accrued benefits is employed.

Figure 14. Previous studies Top 1% and Top 0.1% with Social Security



G. *Adjusting Previous Studies on Wealth Inequality*

Previous studies compute top wealth shares using other datasets than the SCF. In Figure 14, we adjust these studies to include our estimates of the Social Security wealth of the top 1% and bottom 99% and top 0.1% and bottom 99.9%.¹⁴ Our main results remain unchanged.

¹⁴For Saez and Zucman (2016), we report updated time series available from Gabriel Zucman’s website (February 2022 version). Saez and Zucman (2016) report an increase of the top 1% share from 27.8% to 41.8% between 1989 and 2013 whereas the updated time series shows an increase from 28.6% to 35.7%.

VII. 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.

Initial submission: 07-30-2022; Accepted: 05-24-2024

Editors: Antoinette Schoar, Urban Jermann, Leonid Kogan, Jonathan Lewellen, and Thomas Philippon