

Social Security and Trends in Wealth Inequality*

Sylvain Catherine[†]

Max Miller[‡]

Natasha Sarin[§]

December 28, 2020

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 revisits this conclusion by incorporating Social Security retirement benefits into measures of wealth inequality. We find that top wealth shares have not increased in the last three decades when Social Security is properly accounted for. This finding is robust to assumptions about how taxes and benefits may change in response to system financing concerns. When discounted at the risk-free rate, real Social Security wealth increased substantially from \$4.8 trillion in 1989 to \$41.3 trillion in 2016. When we adjust the discount rate for long-run macroeconomic risk, this increase remains sizable, growing from over \$3.9 trillion in 1989 to \$33.9 trillion in 2016. Consequently, by 2016, Social Security wealth represents 57% of the wealth of the bottom 90% of the wealth distribution.

Keywords: Social Security, Inequality, Top Wealth Shares

JEL codes: D31, E21, G51, H55, N32

*We thank Alberto Bisin, Steve Cecchetti, Jason Furman, Wojtek Kopczuk, Francisco Gomes, Joao Gomes, Fatih Karahan, Antoine Levy, Olivia Mitchell, Greg Mankiw, Adair Morse, Kim Schoenholtz, Andrei Shleifer, Lawrence Summers, Constantine Yannelis, Owen Zidar, Gabriel Zucman, and Eric Zwick, and participants at the Wharton Micro Brownbag Seminar, the Chicago Booth Household Finance Conference, Northern Finance Association, NBER Summer Institute, the Red Rock Finance Conference, CEPR Household Finance Conference, NBER Public Economics, NYU Stern Finance Seminar, Stanford GSB Finance Seminar, Federal Reserve Board Seminar, and George Washington University Seminar for helpful comments. 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, Finance department. Email: scath@wharton.upenn.edu

[‡]University of Pennsylvania, Wharton, Finance department. Email: maxmil@wharton.upenn.edu

[§]University of Pennsylvania, Law School and Wharton, Finance department. Email: nsarin@law.upenn.edu

1 Introduction

It is widely believed that wealth inequality in the United States is on the rise. This belief is supported by several studies which, though they differ in their methodology, all use [Piketty \(2013\)](#)'s definition of wealth: the market value of all assets owned by households, net of debt. This paper builds on past work to broaden the definition of wealth to include the value of Social Security retirement benefits, the primary savings vehicle for most households. We find that top wealth shares have not increased since 1989 when Social Security is accounted for.

Why should Social Security wealth count when we estimate wealth concentration? Middle-class households mostly accumulate wealth for retirement ([Modigliani and Brumberg, 1954](#)), which means that a generous public pension system can reduce their private savings rate. This has no impact on lifetime consumption, but exacerbates marketable wealth inequality. If researchers are focused on how living standards diverge and whether the distribution of resources is fair, then consideration of Social Security is imperative.

If instead researchers focus on private wealth accumulation because it is a tool for the exercise of political power ([Saez and Zucman, 2019](#)), then our exercise is still important. Social Security's existence is revelatory on the counterfactual: private wealth—and thus, political power—would have evolved differently if individuals had to fund their later-life consumption through private savings.¹

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 accrues 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

¹Moreover, households can afford to make higher political contributions if they can rely on Social Security during retirement.

old-age retirement program, and we exclude disability insurance, which would lead to an even larger reduction in top wealth shares.

For retirees, we can calculate Social Security wealth from the SCF directly, since benefits are reported. 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, Karahan, Ozkan and Song, 2019; Güvenen, Kaplan, Song and Weidner, 2018). 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. We first offer a risk-free valuation of Social Security wealth using the treasury market yield curve. We find that the top 10% and top 1% “marketable wealth” share grew by 10 percentage points between 1989 and 2016, in line with existing estimates (Piketty, Saez and Zucman, 2018; Smith, Zidar and Zwick, 2020). Once Social Security wealth is included, these trends disappear: the share of the top 10% actually dropped by 3.3 percentage points. The top 1% share is basically flat, rising by only 0.6 percentage points.

However, discounting should reflect the risks associated with the Social Security program (Geanakoplos, Mitchell and Zeldes, 1999). 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 (Catherine, 2019; Geanakoplos and Zeldes, 2010). Our risk-adjustment decreases the stock of Social Security wealth by nearly 20 percent. This has a disproportionate impact on young workers who are most exposed to long-run macroeconomic risk. These workers are nearly always in the bottom 90% of the wealth distribution, and so adjusting for labor market risk decreases the Social Security wealth of this group.

Even after this correction, we find that inequality trends are substantially attenuated relative to past estimates that exclude Social Security. From 1989 to 2016, the top 10% wealth share decreases

by 1.3 percentage points. The top 1% share increases, but only by 1.6 percentage points. The conclusion that Social Security significantly attenuates the recent growth in marketable wealth inequality is also robust to alternative assumptions that incorporate the risk that benefits will be cut (or taxes will rise) in the future, differences in life expectancy among the rich and the poor, and the possibility of persistently low economic growth.

Social Security dramatically impacts inequality trends because the growth in Social Security wealth outpaces 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 that is 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 percent. Moreover, life expectancy increased by nearly 4 years.

Finally, and most importantly, real interest rates have fallen. This means that to fund consumption during retirement with private savings, an investor has to save more (since less interest will accrue) or buy a higher-priced annuity in 2016 than 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 (Cochrane, 2020). In fact, Saez and Zucman (2016) find that the growth in household debts explain most of the decline in the bottom 90% share of marketable wealth. This is at least in part because identical mortgage payments correspond to a larger liability when rates are low than when rates are high. It is errant to consider the role of macroeconomic trends when valuing some kinds of assets and liabilities, but not others.

The low rate environment has an outsized impact on Social Security wealth because rate declines redistribute wealth away from holders of short-duration assets, favoring those with long-term investments (Auclert, 2019), like future retirement benefits. Further, of recent, long-duration assets have outperformed their short-duration counterparts (Binsbergen, 2020). This has increased

the value of Social Security relative to other asset classes: An investor who replicated Social Security cashflows by buying long-term bonds (representing future benefits) and selling short and medium-term bonds (representing payroll taxes) would have become much wealthier over the last three decades.

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.

Second, we quantify how the value of Social Security has increased over the last thirty years, such that it attenuates the upward trend in marketable wealth inequality. Early work dating back to [Feldstein \(1976\)](#) shows how marketable wealth, excluding Social Security, is more unevenly distributed than total wealth. But we are the first to document how sensitive recent trends in top wealth shares are to its inclusion, and why. Our findings 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.

Third, 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 public and private sources of wealth is misleading ([Attanasio and Brugiavini, 2003](#); [Attanasio and Rohwedder, 2003](#); [Scholz, Seshadri and Khitatrakun, 2006](#)).

Fourth, we contribute by building a measure of wealth inequality that is better suited to policy evaluation and cross-country comparisons than prior work. A focus on narrow marketable wealth inequality means that investments in 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.

Finally, we provide a framework for the valuation of long-lived benefit streams, whose value is closely tied to shifts in the interest rate environment. Our exercise is about Social Security, but these insights are similarly important to the valuation of private annuities like defined benefit pension plans. Contemporaneous work values Social Security at a fixed interest rate over the last three decades, despite significant declines in interest rates ([Sabelhaus and Volz, 2020](#)). But there is no reason to believe that increases in the Social Security risk premium exactly offset the decline in interest rates, as a fixed-discounting approach implies. And ignoring the role that rate changes play in Social Security's valuation makes it impossible to incorporate Social Security into existing estimates of private wealth concentration, where recent trends are closely related to the low rate environment ([Cochrane, 2020](#)).

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 ([Saez and Zucman, 2016](#)) 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](#)).

Based on the insights of this past literature, we augment our definition of wealth to include Social Security benefits that workers accrue. In essence, we extend [Feldstein \(1974\)](#), who relied on survey data 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 show 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.

Finally, our work is related to an extensive literature on the magnitude and beneficiaries of redistribution through Social Security. Because the Social Security benefit formula replaces a greater fraction of the lifetime earnings of lower earners than higher earners, it is generally thought of as progressive. Past work documents how much of the intracohort redistribution in the United States is related to factors beyond income: for example, benefits are transferred from those with low life expectancies to those with higher, and from single workers to non-working spouses ([Feldstein and Liebman, 2002](#); [Gustman and Steinmeier, 2000, 2001](#); [Liebman, 2002](#)).

The remainder of our paper is organized as follows. Section 2 presents stylized facts regarding the size, growth and distribution of Social Security wealth. Section 3 describes our data sources. Section 4 lays out our approach to estimating Social Security wealth and its distribution and presents our baseline results. Section 5 adjusts our valuation for macroeconomic risk. Section 6 explains the relationship between changes in the interest rate environment and Social Security’s valuation. Section 7 provides additional discussion of our results, highlighting how the composition of wealth has changed for those outside of the top of the distribution. Section 8 shows that our conclusions are robust to alternative assumptions; as well as concerns that benefits will not be fully paid out to beneficiaries, or taxes will be raised, to deal with the Social Security trust fund’s imminent insolvency. Section 9 concludes.

2 Stylized facts

2.1 Social Security benefits are evenly distributed

Today, Social Security provides the majority of income to most elderly Americans: nearly 90 percent of individuals above the age of 65 receive Social Security benefits, and for over half of beneficiaries, these benefits represent 50 percent or more of their total income ([Dushi, Iams and Brad, 2017](#)).


Figure 1 shows the distribution of these retirement benefits by decile of net worth in the SCF. They are larger for wealthier retirees, who receive greater benefits because they paid more into Social Security over their lifetimes. But, compared to the distribution of capital income, these are minor differences. Among recent retirees, the top decile receives less than 15 percent of Social Security benefits, and more than 60 percent of income from capital.

2.1.1 Benefits formula

Social Security is evenly distributed because the benefit formula is progressive, meaning that the program replaces a larger share of earnings for lower-income workers. While in the labor force, workers pay a 12.4% payroll tax to finance Social Security. Only earnings below the Social Security cap (\$132,900 in 2019) are taxed and count towards future benefits.

Figure 1: Distribution of Social Security Benefits and Capital Income around retirement

This figure shows the distribution of Social Security retirement benefits and capital income in the SCF for respondents between 62 and 75 years of age in 2016. Capital income is total income reported in the SCF less wage income, Social Security income, and government transfers. The horizontal axis is the non-Social Security wealth decile computed based on the SCF population aged 62 to 75.



Once they retire, benefits are determined based on individuals’ historical earnings. Actual earnings are indexed to average wages in the year they were earned. Practically, wage indexing in this manner adjusts workers’ benefits for both inflation and real wage growth. After the earnings in each year have been indexed, the best 35 years are kept and averaged to determine an individual’s average indexed yearly earnings (AIYE).²

The Social Security formula determines benefits in a progressive manner as a function of the AIYE. In 2019, an individual who retires at the full retirement age will receive first year benefits as the sum of:

1. 90% of the share of the AIYE below the first bend point (\$11,112);
2. 32% of the AIYE between the first and second bend point (\$66,996);

²This is the procedure for earnings before age 60. Afterward, earnings can increase benefits, but they enter in nominal terms.

3. and 15% of the AIYE above the second bend point

These “bend points” are determined when a cohort reaches age 62. Early retirement (before age 66) reduces benefits and delayed retirement increases benefits up to age 70.

2.2 Social Security wealth has increased

According to the SSA, Social Security promises rose in value by over 200 percent in real terms between 1989–2016. Given its progressivity, this growth is likely to offset part of the rise in marketable wealth concentration. We document this rise and discuss three factors that contribute to it: changes in the scope of the program; population aging, which boosts the share of U.S. citizens who receive Social Security benefits; and the decline in interest rates.


2.2.1 SSA actuarial estimates

Figure 2 illustrates how the stock of Social Security wealth has changed over time according to SSA’s annual reports. We graph data reported annually by the SSA’s Office of the Chief Actuary, which estimates the theoretical “transition cost” for the program, intended to provide a rough estimate of the cost of meeting obligations to current beneficiaries and terminating the program for all new entrants. The closed-group transition cost estimates reported reflect the present value of expected benefits that will accrue to those currently contributing to Social Security, net of their expected payroll tax payments and current Social Security reserves. SSA transition costs include a 100-year projection period that is intended to capture the lifetime of all current participants included in the valuation exercise.

The total value of Social Security benefits owed to workers net of the taxes they will pay into the Social Security program in their lifetimes has more than doubled in real terms over the last three decades. By 2019, this total was over \$33 trillion or 42% of estimates of household net worth, excluding Social Security ([Piketty, Saez and Zucman, 2018](#)).

Figure 2: Aggregate Social Security wealth implied by SSA reports

This figure represents the net present value of the Old-Age, Survivor and Disability Insurance programs, inferred from actuarial estimates by the Office of the Chief Actuary. We define this net present value as the sum of the “closed-group transition cost” and the value of the Social Security Trust Fund. The “closed-group transition cost” refers to the present value of expected future benefits to current Social Security participants net of their future payroll tax payments, minus the value of the Trust fund. SSA actuaries estimate the “closed-group transition cost” every year using a 100-year projection. Details are available in SSA’s *Actuarial Note #2019.1*.



2.2.2 Expansion of Social Security in the 1970s


Over time there have been several concerns about the financial stability of the U.S. old-age retirement program. In response, policymakers have repeatedly increased the maximum salary subject to the payroll tax. In the short-run, these changes provide additional funding for Social Security beneficiaries through progressive tax increases on those with highest earnings. Nevertheless, raising the tax ceiling also increases what Social Security must pay these beneficiaries in the future.

Over the last 70 years, maximum taxable earnings increased by three times as much as wages grew. This means that the scope of Social Security, and its relative importance for retirement savings, grew substantially. These policy choices increased the benefits owed to middle and upper-income individuals later in life, with important implications for patterns of wealth accumulation during labor market years ([Gustman, Steinmeier and Tabatabai, 2012](#)).

Figure 3 reports the evolution of the relationship between the maximum taxable earnings base and average earnings in the United States from 1961 onward. The share of the working-age population whose earnings fall below the maximum base increased from 73 to 94 percent over this period;³ with the result being that Social Security benefits replace a greater share of lifetime benefits today than they did previously for the upper middle-class. The ratio of maximum taxable earnings to average annual wages was less than 1.5 until 1970; by the 1980s, when the tax cap was automatically indexed to changes to wage growth, this ratio stabilized at around 2.5. The oldest workers in our sample entered the labor force after World War II, and the top earners in this group saw their contributions to Social Security double during their time in the labor force.

Figure 3: Social Security Benefits and Wages: 1961-2018

This figure shows the differential growth of the base of revenue supporting Social Security benefits compared to the Wage Index used by the SSA and median personal wage coming from the U.S. Census Bureau. Panel A shows the levels for each series in 2018 dollars. Panel B shows the benefits base series divided by both the wage index and the median wage.



³ *Social Security Administration, Annual Statistical Supplement, 2018, Table 4.B1*

2.2.3 Demographic changes

There is a direct link between the level of Social Security wealth and the age distribution. Social Security wealth peaks around retirement, when individuals have paid in maximally to the program, yet accrued benefits have yet to be disbursed. The share of the population near retirement age and for whom Social Security wealth is at its peak (those between the ages of 50-70) has increased by over 40 percent.

In the coming years, the age pyramid will shift further: the share of the U.S. population over the age of 65 has risen from 12.5% to 16.9%, and it is projected to grow to 22% by 2050, as the Baby Boomer generation begins to claim Social Security payments. Barring other contemporaneous changes, population aging will decrease the value of Social Security wealth, since a growing share of accrued benefits will be paid out to new retirees. As a result of this demographic shift, Social Security will be the primary source of income for an even larger swath of the population.

2.2.4 Decline in interest rates

We define the expected Social Security wealth of current program participants as the present value of benefits net of the present value of payroll taxes to be paid, discounted using the average nominal yield curve in each survey year.

To illustrate, consider an individual who was 40 years old in 1989. By this point, he had spent 20 years paying Social Security payroll taxes. He will spend 25 more years working before he begins receiving Social Security benefits at age 65, for 20 years (his assumed life expectancy is 85). The present value of his aggregate Social Security wealth is thus:

$$\begin{aligned} PV_{SS \text{ Wealth}} = & \frac{\text{Benefit}}{(1 + r_{25})^{25}} + \frac{\text{Benefit}}{(1 + r_{26})^{26}} + \dots \\ & - \text{Tax}_{1989} - \frac{\text{Tax}_{1990}}{(1 + r_2)^2} - \dots - \frac{\text{Tax}_{2014}}{(1 + r_{25})^{25}}, \end{aligned} \tag{1}$$

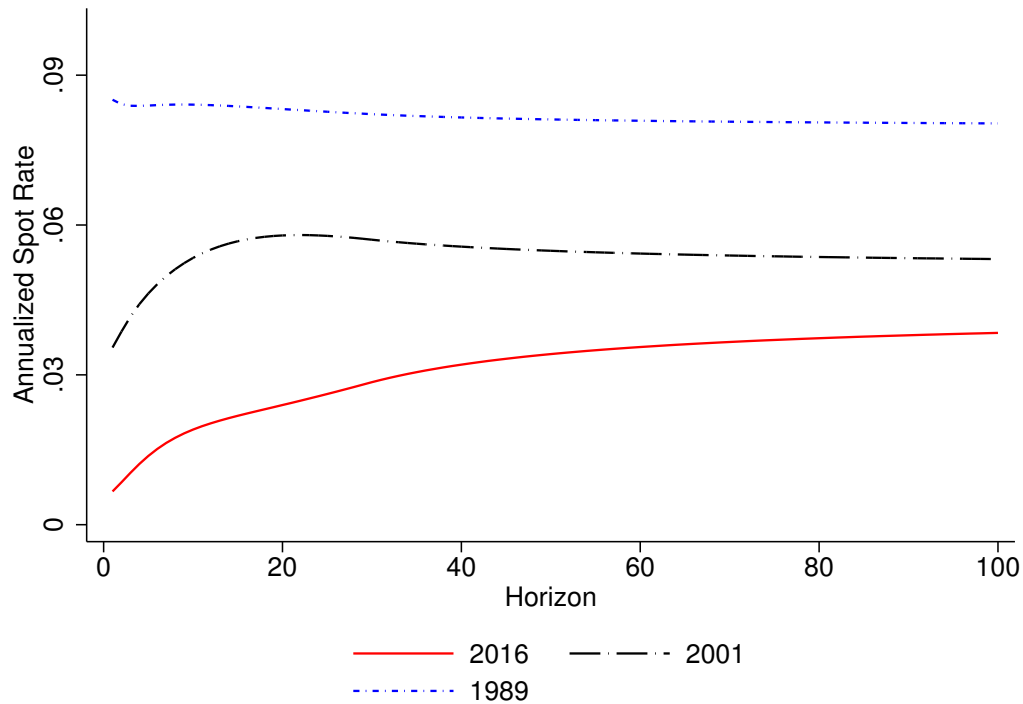
where r_i represents the annualized zero-coupon spot rate i periods into the future; and benefits are benchmarked to economy-wide average annual earnings presently, but to prices post-retirement, thus reflecting both the trajectory of inflation and wage growth. Payroll taxes are a percentage of wages, up to the cap.

Figure 4 traces out the evolution of the market yield curve over our sample period (1989–2016)

following [Gürkaynak, Sack and Wright \(2007\)](#). During this period, the average nominal yield curve fell by 70 percent. This mechanically increases the value of future retirement benefits, as Equation (1) makes clear.

Figure 4: Market Implied Yield Curve

This figure shows the annualized zero coupon rates taken from [Gürkaynak, Sack and Wright \(2007\)](#) from 1 to 100 years for 1989, 2001, and 2016—the beginning, middle, and end of the SCF time series. The data are extended beyond 30 years by applying the 29 to 30 year forward rate to the annualized spot rate at 30 years, under the assumption that this forward rate represents the long-run interest rate on nominal government claims.



3 Data

3.1 Survey of Consumer Finances

We use the triennial SCF for four purposes: (i) measuring marketable wealth shares, (ii) estimating aggregate Social Security wealth, (iii) determining the share of Social Security wealth going to the wealthy, and (iv) validating our simulation by comparing predicted and observed retirement benefits.

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, with the notable omission of defined benefit pension plans. Importantly, the SCF over-samples wealthy households to provide a more accurate view of their assets. One caveat is that the SCF does not survey extremely wealthy households, per agreement with the U.S. Treasury, thus their wealth is excluded. To fill this gap, we follow [Saez and Zucman \(2016\)](#) by adding data from the Forbes 400 list to the richest 0.01%.

To compute the aggregate Social Security wealth time series, we combine our simulated data with the SCF demographic weights. These weights take into account the over-sampling of wealthier households and the probability of non-response in order to create a representative sample of the U.S. population.

To determine the share of Social Security wealth going to the top, we rely on detailed data on retirement and survivor benefits which allow us to compute Social Security wealth for each retiree. 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

Finally, to validate our simulation, we compare predicted and observed benefits for cohorts that retired between 1989 and 2016. For this exercise, we rely on the age at which pensioners retired. Because the SCF reports this information, we can reconstruct how much pensioners would have received if they claimed their benefits at full retirement age, which is essential to making an apples-to-apples comparison between simulated and observed benefits.

3.2 Other sources

Mortality Historical mortality rates come from the Human Mortality Database (HMD) operated by the University of California, Berkeley and the Max Planck Institute for Demographic Research. The HMD provides data on life expectancy and conditional survival probabilities by gender from 1933–2017.

Yield curve Yield curve data come from the Federal Reserve Board of Governors who broadly follow the methods of [Gürkaynak, Sack and Wright \(2007\)](#). These data provides an estimate of the zero-coupon yield curve using off-the-run Treasury coupon securities for horizons up to 30

years. This series indicates the rate of return investors require to hold government debt, and is often thought of as the nominal risk-free rate of return. 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 $r_{t,t+30+h} = ((r_{t+29,t+30})^h (r_{t,t+30})^{30})^{\frac{1}{30+h}}$. Our assumption is that this forward rate represents the long-run interest rate on nominal government claims. In addition, we use data from [Gürkaynak, Sack and Wright \(2008\)](#) on the implied real yield curve from Treasury Inflation Protected Securities (TIPS) to test the sensitivity of our results to alternative discount rate assumptions.

Social Security reports We use inflation, wage growth, and discount rate projections from the SSA Annual Reports. Historical wage growth and inflation forecasts are used in to calibrate the model, and discount rate projections to compare our results to the SSA’s own estimates of aggregate Social Security wealth. We also collect important 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.

For retirees, benefits are already being paid and reported in the SCF. Their present value depends on mortality, expected inflation and the contemporaneous market yield curve. We are thus able to compute the Social Security wealth of each household and whether they belong to the top of the marketable wealth distribution from the SCF directly.

For workers, taxes and benefits are simulated and discounted using the market yield curve. The simulation produces an average Social Security wealth by cohort, gender and year. We combine these averages with the SCF demographic weights to estimate the aggregate Social Security wealth of each cohort. Finally, we determine the share of each cohort’s Social Security wealth going to the top 10% or top 1%.

4.1 Social Security wealth

We define Social Security wealth as the present value of future benefits net of future taxes, including cash flows in the present year.⁴ Assuming a retirement age of 66, the Social Security wealth of a worker from cohort c in year t is:

$$\begin{aligned} \text{Social Security Wealth}_{it} = & \sum_{s=c+66}^T \left(\prod_{k=t}^{s-1} (1 - m_{itk}) \right) \frac{\mathbb{E}[\text{Benefits}_{it}]}{(1 + r_{ts})^{s-t}} \\ & - \sum_{s=t+1}^{c+65} \left(\prod_{k=t}^{s-1} (1 - m_{itk}) \right) \frac{\mathbb{E}[\text{Taxes}_{it}]}{(1 + r_{ts})^{s-t}} \end{aligned} \quad (2)$$

where T is the maximum age and expectation terms take mortality into account. This definition of Social Security wealth is consistent with the valuation of other forms of marketable wealth, for example businesses and real estate, which capture the present value of cashflows generated by existing assets and the net present value of future investments.

4.2 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

$$\text{Social Security Wealth}_{it} = \sum_{s=t}^T \left(\prod_{k=t}^{s-1} (1 - m_{itk}) \right) \frac{\text{Benefits}_{it}}{(1 + r_{t,s})^{s-t}} \frac{\mathbb{E}[\text{CPI}_s]}{\text{CPI}_t} \quad (3)$$

where nominal benefits are indexed to the consumer price index CPI_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. Actual survivor benefits are added to the benefits of the surviving spouse up to a family maximum which depends on the benefits of the deceased.⁵

⁴We discuss an alternative definition of Social Security wealth based on accrued benefits in Section 8.1.

⁵We provide more details on the computation of their present value in Appendix B.4. Although we include survivor benefits, which represent 16% of old-age benefits in our calculations, we exclude spousal benefits that accrue to those who have not worked or worked but earned less than their partners, as these constitute a much smaller share (less than 4%) of old-age benefits (Social Security Administration, 2016).

4.3 Workers

For workers, Social Security wealth is the present value of future benefits net of future taxes Equation (2). Both benefits and taxes depends on the evolution of a worker's earnings relative to the national average. Worker i 's earnings are:

$$L_{it} = L_{1,t} \cdot L_{2,it}, \quad (4)$$

where $L_{1,t}$ denotes the average national wage and $L_{2,it}$ is the idiosyncratic component of earnings.

4.3.1 Taxes

Taxes depend on workers' earnings below the Social Security wage base ($SSWB_t$) and the level of the tax rate. Specifically, future Social Security contributions of worker i will be:

$$\text{Taxes}_{it} = \text{Tax Rate} \times \min \{L_{it}, SSWB_t\}. \quad (5)$$

4.3.2 Benefits

For simplicity, we assume that workers retire at the full retirement age of 66. Yearly benefits depend on a workers' average indexed yearly earnings (AIYE). Only the taxable share of earnings is taken into account. Hence, the indexed earnings of year t are:

$$\text{Indexed Earnings}_{it} = \min \{L_{it}, SSWB_t\} \frac{L_{1,c+60}}{L_{1,t}} \quad (6)$$

where $\frac{L_{1,c+60}}{L_{1,t}}$ is the indexation coefficient. The AIYE is the average of the best 35 years of indexed earnings up to retirement. Yearly benefits depend on year of birth c , and are a concave and piecewise linear function of AIYE. The marginal replacement rate drops at two cohort-specific bend points $\text{Bend}_{1,c}$ and $\text{Bend}_{2,c}$.

- Below the first bend point, the replacement rate is 90%, so if $\text{AIYE}_i < \text{Bend}_{1,c}$, then:

$$\text{Benefits}_{it} = \frac{\text{CPI}_t}{\text{CPI}_{c+60}} \times 0.9 \times \text{AIYE}_i \quad (7)$$

where $\frac{\text{CPI}_t}{\text{CPI}_{c+60}}$ is an adjustment for the increase in the consumer price index since the retiree turned 60.

- Between the first and second bend points, the marginal replacement rate drops to 32%, so if $\text{Bend}_{1,c} \leq \text{AIYE}_i < \text{Bend}_{2,c}$:

$$\text{Benefits}_{it} = \frac{\text{CPI}_t}{\text{CPI}_{c+60}} \left[0.9 \times \text{Bend}_{1,c} + 0.32 \times (\text{AIYE}_i - \text{Bend}_{1,c}) \right] \quad (8)$$

- Then, the marginal replacement rate drops to 15%, so if $\text{AIYE}_i \geq \text{Bend}_{2,c}$:

$$\text{Benefits}_{it} = \frac{\text{CPI}_t}{\text{CPI}_{c+60}} \left[0.9 \times \text{Bend}_{1,c} + 0.32 \times (\text{Bend}_{2,c} - \text{Bend}_{1,c}) + 0.15 \times (\text{AIYE}_i - \text{Bend}_{2,c}) \right] \quad (9)$$

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).

4.3.3 Simulating past and future earnings trajectories

To estimate the expected future taxes and benefits of each cohorts, we need to simulate the distribution of income paths of its members. To do so, we use the flexible labor income process estimated in Guvenen, Karahan, Ozkan and Song (2019). Specifically, we assume that the idiosyncratic component of a worker's earnings $L_{2,it}$ evolves as follows:

$$\text{Level of idiosyncratic earnings: } L_{2,it} = (1 - \nu_t^i) e^{(g(t) + \alpha^i + \beta^i t + z_t^i + \varepsilon_t^i)} \quad (10.1)$$

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

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

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

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

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

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

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 and fully mean revert within the year. Workers can also experience a non-employment shock with some probability p that can vary with any mixture of age, income, and gender, and whose duration is exponentially distributed.

In Equation (10.1), $g(t)$ is a quadratic polynomial of age that 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 final component of the earnings process is the nonemployment shock Equation (10.6), which is realized with probability p_ν in each period. The duration ν_t reflects the duration of full-year nonemployment (zero annual income).

Once we have simulated workers' earnings trajectories, we apply the Social Security benefit and tax formulas detailed above and average Social Security wealth for each cohort, by gender, each year.

4.4 Aggregate Social Security wealth by cohort

We aggregate Social Security wealth at the cohort level using the SCF demographics weights. Specifically, we have:

$$\text{Cohort Social Security Wealth}_{ct} = \sum_{i \in c} \text{weight}_{it} \times \text{Mean Social Security Wealth}_{gct}. \quad (11)$$

For retirees, Social Security wealth is computed at the individual level, so the aggregation is straightforward. For workers, we obtain the mean Social Security wealth by cohort and gender from our simulation. 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 65, the simulated data and SCF overlap. For those whose benefits are reported in the SCF, we rely on these estimates. For individuals without explicit benefits, we fill in the average simulated Social Security wealth, adjusting for the non-contribution share of the population. We then aggregate both the SCF and the simulation estimates to account for SCF respondents not currently receiving benefits that will receive benefits in the future.

There are also individuals aged 66-69 who have yet to receive benefits, but for this population there is no overlap of simulated data and survey data. For these individuals, we backfill⁶ average benefits and wealth from the succeeding survey for respondents from 70 to 73 years of age, adjusting for inflation and the number of respondents that will not receive benefits. This aggregation process is discussed in greater detail in Appendix C.

4.5 Calibration

Income process We calibrate the idiosyncratic income process described in Equations (10.1) to (10.7) using estimates from [Guiso, Karahan, Ozkan and Song \(2019\)](#) listed in Appendix Table E.1. These parameters are estimated using the Social Security Master Earnings File and match many moments from the cross-section and dynamics of individual of earnings.⁷

We assume $g(t)$ to be a cohort and gender-specific quadratic polynomial of age. To estimate these polynomials, we rely on data in [Guiso, Kaplan, Song and Weidner \(2018\)](#). In their appendix, these authors report the average earnings of each cohort c and gender g by year from 1957 to 2013, which we denote Cohort Earnings $_{cgt}$. We estimate life cycle income profile of each cohort and gender by running the following OLS regression:

$$\hat{g}_{cg}(t) = \ln \left(\frac{\text{Cohort Earnings}_{cgt}}{L_{1,t}} \right) = \alpha_{cg} + \beta_{cg,1} \times \text{Age}_{ct} + \beta_{cg,2} \times \text{Age}_{ct}^2 + \beta_{cg,3} \times \text{Age}_{ct}^3 + \epsilon_{gct} \quad (12)$$

where $L_{1,t}$ is the Social Security wage index. Our data includes workers who enter the labor force from 1949-2016, so it is broader than [Guiso, Kaplan, Song and Weidner \(2018\)](#)'s sample. 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. In our simulation, we use the

⁶For 2016, we cannot backfill, so we fill in directly average benefits and wealth for 70–73 year olds within the same survey.

⁷One important caveat is that this process is estimated only for males, though [Guiso, Karahan, Ozkan and Song \(2019\)](#) note that their “analysis for women found qualitatively similar patterns.”

predicted values derived from Equation (12) for each gender and cohort as our calibration for $g(t)$, from which we subtract half the variance generated by idiosyncratic shocks and heterogeneity in income profiles to adjust for Jensen’s inequality.

Social Security parameters Over our sample period, Social Security parameters have scaled up nearly perfectly with the level of the wage index. We assume that this pattern will persist in the future. Hence, we assume that the Social Security wage base will remain 2.5 times the wage index ($SSWB_t = 2.5 \times L_{1,t}$), and the bend points of the benefits formula will remain .21 and 1.25 the wage index ($Bend_{1,t} = 0.21 \times L_{1,t}$ and $Bend_{2,t} = 1.25 \times L_{1,t}$). Part of our simulation covers historical years before the 1980’s when the Social Security wage base was lower, so we use actual historical values of the Social Security wage base to simulate covered earnings preceding any SCF survey year. We assume that Social Security respectively covers 90% and 80% of the male and female populations. We estimate these coverage ratios by looking at the share of male and female above 70 years old that do not receive any benefit (Appendix Figure E.2).

Macroeconomic assumptions Since Social Security cash flows are inflation-indexed, they should be discounted using the real yield curve. Blocker, Kotlikoff, Ross and Vallenas (2019) argue that the Treasury Inflation Protected Securities (TIPS) term structure should be used to price Social Security. Unfortunately, TIPS data are only available from 1999 onward, so we elect against using the TIPS yield curve to preserve consistency in our methodology. But we use TIPS data to validate our approach in later years.

In our baseline specification, we use the nominal yield curve for Treasury notes. Therefore, we let cash flow grow with the consumer price index. We use inflation projection from SSA reports as we are not aware a single dataset that provides long-term inflation projections throughout our entire sample. We also take real wage growth forecasts from the SSA reports “intermediate cost” scenario. We discuss the more pessimistic “high cost” and the more optimistic “low cost” scenarios in Section 8.2.⁸

⁸No other data (e.g., inflation expectations, breakeven rates, Michigan Survey of Consumers inflation expectations) provide information required by our methodology for our entire sample period. In alternate specifications, we apply TIPS rates to the post-1999 period and find similar results.

4.6 Validation

We validate our methodology in two ways. First, we check that the Social Security benefits predicted by our simulation match those observed in the data for cohorts that retired over our sample period. Second, when using the same discount rate assumption as the SSA, we check that we obtain similar estimates of the evolution of aggregate Social Security wealth.

Matching observed benefits at retirement age In Figure 5, we compare simulated and observed Social Security benefits for pensioners between age 62 and 67, for each SCF survey year. For those who retire before or after full retirement age, we use Social Security rules to determine what their full retirement age benefits would have been. The simulated data track observed benefits closely.

Matching SSA estimates of aggregate Social Security wealth As reported in Figure 2, the SSA estimates the aggregate stock of Social Security wealth each year. It calculates the present value of benefits to current participants, net of the present value of payroll taxes. Our goal in this paper is not to replicate the SSA estimates of Social Security wealth, 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.⁹ This is why in our risk-free valuation, we instead discount cash flows using Treasury estimates of the off-the-run yield curve based on a large set of outstanding Treasury notes and bonds, reported daily [Gürkaynak, Sack and Wright \(2007\)](#). However, if we chose to use the SSA's discount rates, we should be able to match its estimates.

⁹Appendix Figure E.1 reports the evolution of the yield curve and SSA discount rates between 1989 and 2016. The SSA discount rates fell by roughly 2 percentage points. The market yield curve fell by three times that amount.

Figure 5: Simulated and actual full retirement age benefits

This figure reports mean Social Security benefits at full retirement age predicted by the model and observed in the SCF data, by gender and survey, and conditional on receiving benefits. Individuals not receiving benefits are not included. Panel A represents mean benefits for men and Panel B represents mean benefits for women. Because pensioners retire at different ages, we use Social Security rules to compute the benefits they would receive had they claimed their benefits at full retirement age, a process described in Appendix B.2. Benefits are reported in 2018 dollars.





Figure 6 shows these results. The evolution of aggregate Social Security wealth reported by the SSA tracks our estimates, giving us confidence in our simulated estimate of workers' lifetime earnings histories, from which we derive their Social Security wealth.¹⁰

¹⁰Note that in this exercise, we only report 85% of the SSA's estimates because 15% of Social Security cash flows and revenues can be attributed to the disability insurance program, which is outside of the scope of our study. In principle and in practice, our estimates should not and do not match the SSA's aggregate valuation perfectly. One discrepancy is that the latter includes future workers as young as 15 whereas we restrict our analysis to adults 20 and older.

Figure 6: Aggregate Social Security Wealth under Alternative Yield Curve Assumptions

This figure reports different estimates of the aggregate present value of Social Security. The “SSA Reports” line reports estimates by the Office of the Chief Actuary (OACT) as in Figure 2. We adjust these estimates to remove the value of Disability Insurance program by assuming that the Old Age and Survivor program represents 10.6/12.4 of the total, which is consistent with the allocation of payroll tax revenues. The “SSA Discount Rates” line reports our estimates when we use the same discounting assumptions as the OACT. The “TIPS Yield curve” line reports our estimates when we assume no inflation and discount future cash flows using the real yield curve implied by treasury inflation-protected securities. Finally, the ”risk-free valuation” line reports our estimates based on the methodology outlined in Section 4, which uses the nominal market yield curve.



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. SSA-implied aggregate Social Security wealth was nearly \$30 trillion in 2016 and nearly \$40 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 our projections for the consumer price index, 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 TIPS prices and assuming no inflation. This exercise can only be done for the 1999-

2016 period. As reported in Figure 6, this alternative methodology implies a faster increase in aggregate Social Security wealth than ours, and a greater value for 2016.¹¹ Hence, we feel confident that our findings are not driven by incorrect inflation forecasts.

4.7 Assigning Social Security wealth to the top

Our goal is to understand how trends in inequality documented in prior work are impacted by the large and growing stock of Social Security wealth. To do so, we must determine the distribution of Social Security wealth. The appropriate assignment strategy depends on whether households have already claimed their retirement benefits, or are still in the labor force.

Retirees In the first case, we observe benefits and compute the Social Security wealth of retirees at the individual level. Hence, for this part of our sample, we can precisely estimate the share of Social Security wealth that is captured by each centile of the overall marketable wealth distribution.

Workers For households which are not retired, our simulation exercise only produces an estimate of aggregate Social Security wealth by cohort. This wealth needs to be divided 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:

$$\text{Top 10\% Social Security Wealth}_{ct} = \phi_t(\text{Share in Top 10\%}_{ct}) \times \text{Cohort Social Security Wealth}_{ct} \quad (13)$$

where $\text{Share in Top 10\%}_{ct}$ represents the percentage of individuals in cohort c who fell in the overall top 10% of the marketable 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 $\text{Share in Top 10\%}_{ct}$ and $\phi_t(\cdot)$ are readily observable in the SCF.

A numerical example illustrates our approach.

1. According to our simulation, in 2016, 60 year-olds had \$1.2 trillion in Social Security wealth.

¹¹However, 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).

2. As Panel A of Figure 7 shows, in 2016, 21% of 60-year old households were in the top 10% of the overall marketable wealth distribution.¹²
3. Panel B shows that, within the population of young retirees (65-75), the wealthiest 21% held 26% of the Social Security wealth of this age group.
4. Hence, we allocate \$312 billion ($\$1.2 \text{ trillion} \times .26$) of 60 year-olds' Social Security wealth to the top 10% in 2016.

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 2016. 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 marketable wealth distribution is constant across ages. To be sure, this relationship is likely not constant across ages. However, there are several reasons why assuming the reverse is reasonable for our exercise. First, our assumption is most tenuous for the youngest workers, whose earnings trajectories will evolve substantially prior to retirement. However, as illustrated by Panel A of Figure 7, 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. Moreover, 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 7, 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 is conservative and overstates the share of Social Security wealth that accrues to the top 10%. This is 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.

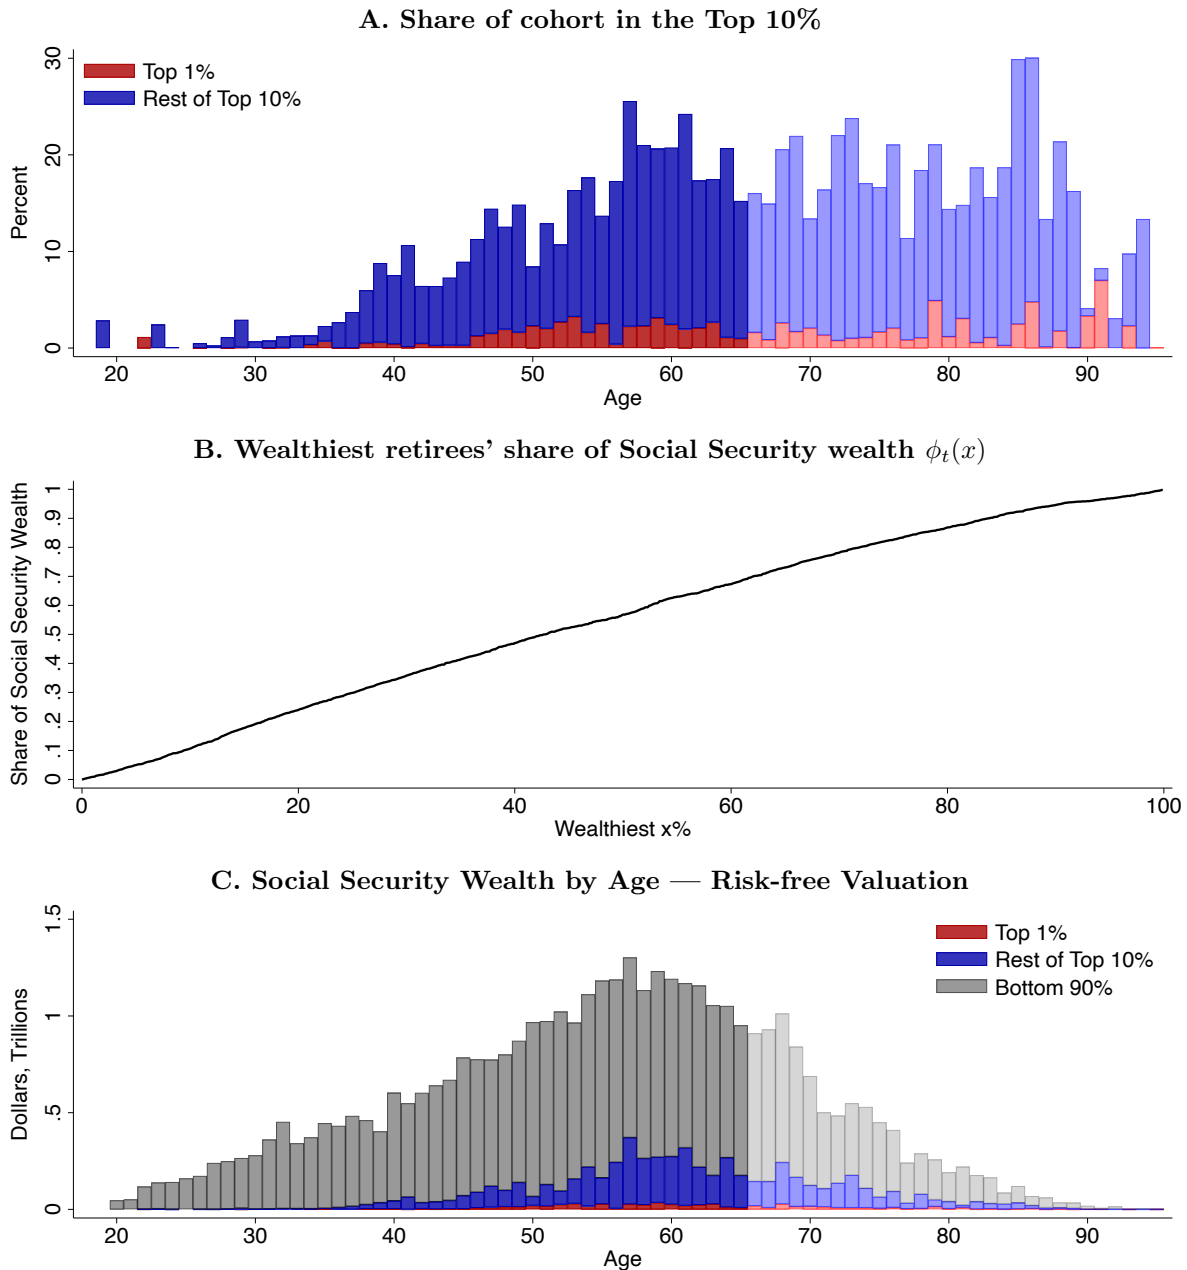
¹²By way of contrast, there are no 20-year olds in the top 10% of the overall distribution in the SCF in 2016. The mechanical relationship between age and wealth accumulation suggests the importance of intra-cohort estimates of inequality.

Figure 7: Assignment of Social Security wealth by working-age cohort in 2016

This figure illustrates how we allocate the Social Security wealth of working-age cohort between the top 10% and bottom 90% of the wealth distribution. Following Equation (13), the Social Security wealth of cohort c going to the top 10% is:

$$\text{Top 10\% Social Security Wealth}_{c,2016} = \phi_t(\text{Share in Top 10\%}_{c,2016}) \times \text{Cohort Social Security Wealth}_{c,2016}$$

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% $_{c,2016}$. 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.



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.

4.8 Baseline top wealth shares

This section compares the levels and trends of top wealth shares under alternative specifications, both including and excluding Social Security wealth. 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.

Figure 8: 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, future Social Security 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.

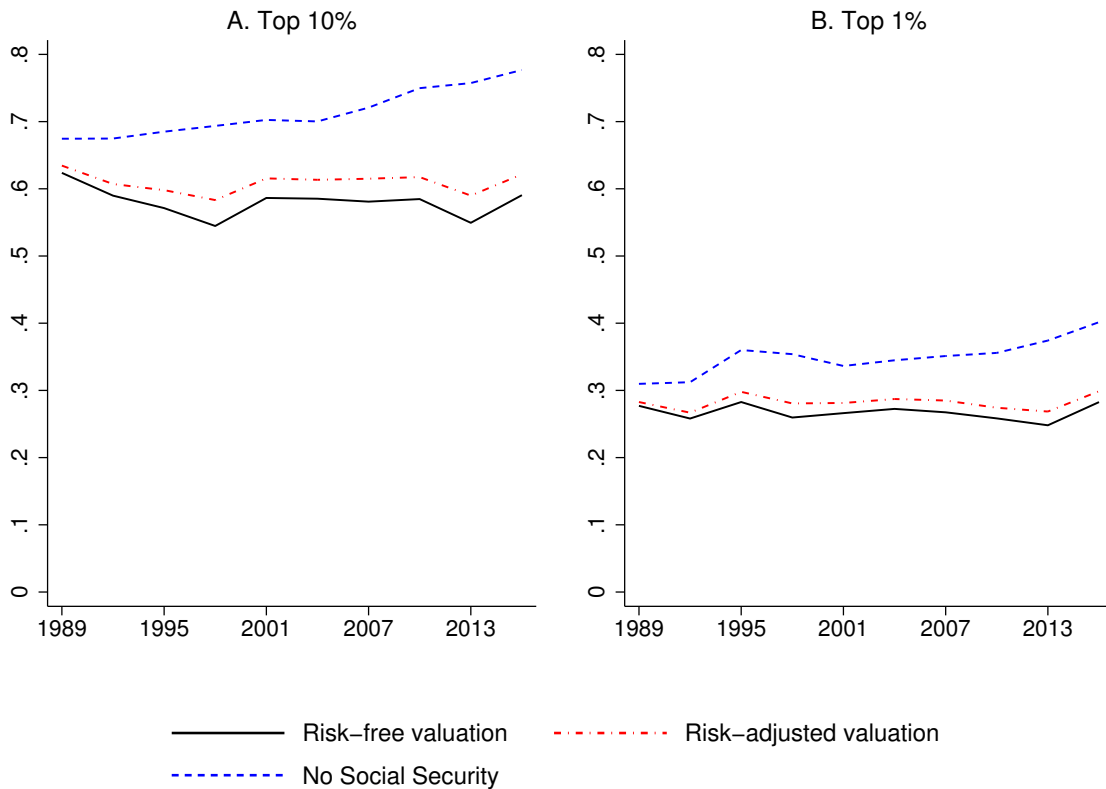


Figure 8 reflects our baseline specification. Panel A focuses on the top 10%. The top 10% wealth share (excluding Social Security) grew by 10 percentage points between 1989–2016. This

is in line with top wealth estimates from others: for example, [Piketty, Saez and Zucman \(2018\)](#) report a 9 percentage point increase in top wealth over this period. Once Social Security wealth is included, this trend is reversed. Rather than rising, the top 10% wealth share falls by 3.3 percentage points over this period.

Panel B of [Figure 8](#) 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 10 percentage points over our sample period. Once it is included, the top 1% share has risen by 0.6 percentage points.

5 Accounting for macroeconomic risk

Overlapping-generation models tell us that 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 per capita earnings is explicitly achieved through wage-indexation. Therefore, Social Security participants are exposed to long-run macroeconomic risk. For this reason, [Geanakoplos and Zeldes \(2010\)](#) and [Catherine \(2019\)](#) argue that Social Security cash flows should not be discounted at the risk-free rate. These studies respectively find that, after adjusting for systematic risk, the market and private values of Social Security obligations are 19% and 37% lower than the sum of future cash flows discounted at the risk-free rate.

In this section, we estimate what the market value of Social Security claims would be if they could be sold to diversified investors. To take systematic risk into account, we assume that future Social Security cash flows perfectly scale up with the level of per capita earnings in the economy. Since 1980, the values of the Social Security wage base and bend points have been growing at the same rate as earnings. In [Section 4.3](#), we show that tax payments are proportional to the level of the wage index ($L_{1,t}$) whereas benefits are proportional to the value of the wage index the year a worker turns 60 ($L_{1,c+60}$). Because she does not care about idiosyncratic risk, a fully diversified investor would discount each tax payment and each benefit he will receive like a security paying a single coupon in the year that the cash flow is realized, which is also indexed on the value of L_1 in the same year. Therefore, we want to determine the expected return for such a security and use it to discount future Social Security cash flows.

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 follow [Geanakoplos and Zeldes \(2010\)](#) and [Catherine \(2019\)](#) by assuming that the stock and labor markets are cointegrated. Cointegration between dividends and earnings is documented in [Benzoni, Collin-Dufresne and Goldstein \(2007\)](#) and would be expected in an economy where the shares of labor and profits are stable over long periods. Specifically, we assume that the log of L_1 evolves as follows:

$$dl_{1,t} = \left((\phi - \kappa)y_t + \mu - \delta - \frac{\sigma_l^2}{2} \right) dt + \sigma_l dz_{1,t}, \quad (14)$$

where $\mu - \delta$ determines the unconditional log aggregate growth rate of earnings and σ_l its volatility.

Log stock market gains follow:

$$ds_t = \left(\mu + \phi y_t - \frac{\sigma_s^2}{2} \right) dt + \sigma_s dz_{2,t}, \quad (15)$$

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:

$$dy_t = -\kappa y_t + \sigma_l dz_{1,t} - \sigma_s dz_{2,t}, \quad (16)$$

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 $L_{1,t+n} = e^{l_{1,t+n}}$ is:

$$\beta_t^{L_1,n} = \left(1 - \frac{\phi}{\kappa} \right) (1 - e^{-\kappa n}) \quad (17)$$

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

$$\mathbb{E}_t \left[r_t^{L_1,n} \right] = \beta_t^{L_1,n} (\mu - r) + r \quad (18)$$

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.

Given the discrete nature of our exercise, we approximate our continuous time results in discrete time by assuming that the discount factor for a Social Security cash flow proportional to the wage index in year n and paid in year $k > n$ is:

$$\text{Discount Factor}_{t,n,k} \approx \left[\prod_{s=t}^n (1 + \beta_s^{L_1, n} (\mu - r) + r_{ts}) \prod_{s=n+1}^k (1 + r_{ts}) \right]^{-1}, \quad (19)$$

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

$$\begin{aligned} \text{Adj. Social Security Wealth}_{it} = & \sum_{s=c_i+66}^T \left(\prod_{k=t}^{s-1} (1 - m_{itk}) \right) \mathbb{E}[\text{Benefits}_{it}] \times \text{Discount Factor}_{t,c_i+60,s} \\ & - \sum_{s=t+1}^{c+65} \left(\prod_{k=t}^{s-1} (1 - m_{itk}) \right) \mathbb{E}[\text{Taxes}_{it}] \times \text{Discount Factor}_{t,s,s} \end{aligned} \quad (20)$$

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


We calibrate our model as in [Benzoni, Collin-Dufresne and Goldstein \(2007\)](#) who estimate κ and ϕ using U.S. macroeconomic data from 1929 to 2004. Specifically, we set $\kappa = .16$ and $\phi = .08$ which, at the limit ($n = \infty$) implies a market beta of 0.5 for distant Social Security cash flows. We assume a constant equity premium of $\mu - r = 0.06$.

5.2 Aggregate Social Security wealth

Figure 9 reports aggregate Social Security wealth with and without adjusting for systematic labor market risk. In line with previous studies, we find that adjusting for systematic risk leads to a large reduction in the net present value of Social Security. The impact of this adjustment varies across the age distribution (Appendix Figure E.7). Most impacted are young workers whose Social Security benefits will be disbursed many years into the future: our risk adjustment cuts the present value of a 25-year old's benefits by 60%.

Figure 9: Total Social Security wealth — Risk-adjusted valuation

This figure shows the present value of Social Security in 2018 dollars under two different discounting specifications. The “risk-free valuation” refers to the present value of Social Security when cash flows are discounted using the market implied yield curve. The “risk-adjusted valuation” also uses the market implied yield curve but is adjusted to take systematic risk coming from wage-indexation into account. As outlined in Section 5, systematic risk only affects discount rates for workers below age 60 and declines over the life-cycle.



5.3 Risk-adjusted top wealth shares

Once macroeconomic risk associated with Social Security cashflows is factored in, Figure 8 shows that the share of top 10% decreased by 1.3 percentage points and that top 1% share has increased by 1.6 percentage points.

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%. This is because the cointegration of labor market income and stock returns is a long-run relationship. By the time older workers are retired or nearing retirement, they are no longer exposed to systematic risk. Consequently, this adjustment decreases the wealth of the bottom 90%, with only a small impact on Social Security wealth of the top of the distribution. Regardless,

top wealth shares remain substantially attenuated relative to prior work.

6 Role of interest rate environment

This section discusses the impact of interest rate changes on our results. First, we show that accounting for Social Security offset trends in marketable wealth inequality because Social Security wealth grew faster than marketable wealth, and the bottom 90% share of aggregate Social Security wealth increased. Then, we explain this phenomenon by showing that Social Security participants have a leveraged exposure to long term bonds, which have performed well over the last couple of decades. Finally, we explain why ignoring the evolution of rates in our valuation of Social Security leads to inconsistent measures of wealth inequality.

6.1 Decomposing the impact of Social Security on top wealth shares

The top 1% share in total wealth, inclusive of Social Security, ($\alpha_{W+SS,t}$) can be decomposed as¹³

$$\alpha_{W+SS,t} = \frac{\alpha_{W,t}W_t + \alpha_{SS,t}SS_t}{W_t + SS_t} \quad (21)$$

where α_{W+SS} is the top 1% share in total wealth (including Social Security), α_W is the top 1% share in marketable wealth, α_{SS} is the top 1% share in Social Security wealth, and $\frac{SS}{W+SS}$ is the share of Social Security in total wealth. Rearranging this equation tells us how taking Social Security into account changes the top 1% wealth share:

$$\alpha_{W+SS,t} - \alpha_{W,t} = (\alpha_{SS,t} - \alpha_{W,t}) \left(\frac{SS_t}{W_t + SS_t} \right) \quad (22)$$

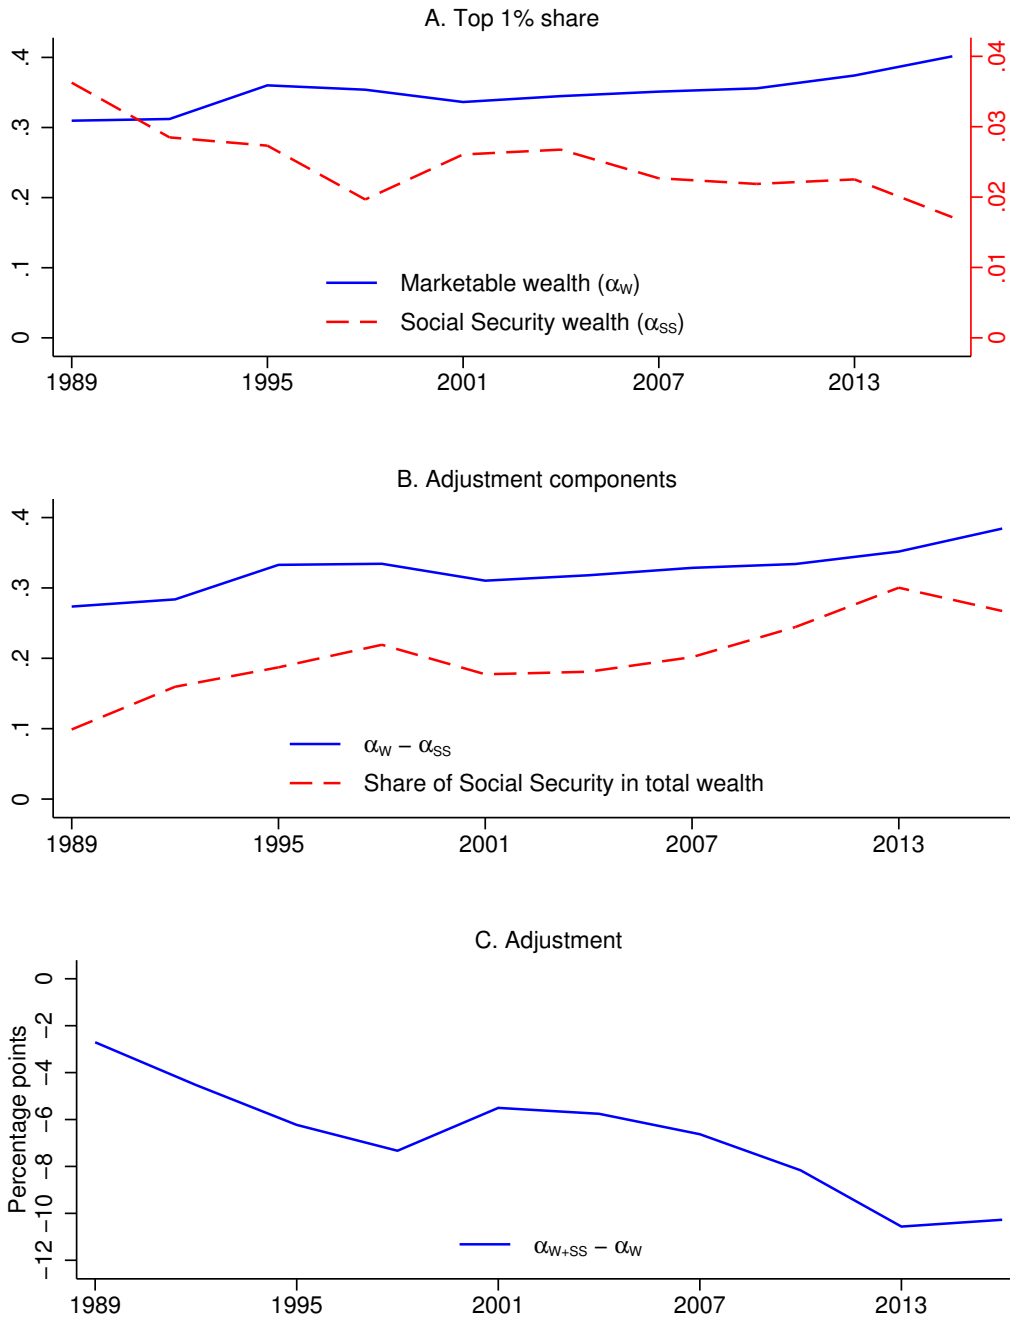
A rise in the top 1% share of marketable wealth can be attenuated by Social Security in two ways. First, there could be a commensurate decline in the top 1% share of Social Security wealth. Second, the relative importance of Social Security wealth, as a share of total wealth, could rise.

Figure 10 shows that both of these effects contribute to our results. In Panel A, we observe that the top 1% share of marketable wealth has trended up, increasing between 1989–2016 by around 10 percentage points. But contemporaneously, the top 1% share of Social Security wealth decreased, falling by nearly two-thirds. Compounding this, in Panel B, we observe that the share of Social Security in total wealth has more than doubled over this period.

¹³We are thankful to Stephen Cecchetti and Kermit Schoenholtz for suggesting this exercise.

Figure 10: Decomposition the effect of Social Security on the Top 1% share

This figure decomposes the effect of including Social Security in top wealth share estimates. Panel A reports the share of the top 1% in marketable wealth ($\alpha_{W,t}$) and Social Security wealth ($\alpha_{SS,t}$). As Equation (22) shows, the effect of Social Security is the product of the difference between these two shares ($\alpha_{W,t} - \alpha_{SS,t}$) and the share of Social Security in total wealth. Panel B reports the evolution of these two components. The product of these components is reported in Panel C and represents the difference between the top 1% share of total wealth (including Social Security) and marketable wealth: $\alpha_{W+SS,t} - \alpha_{W,t}$.



6.2 Leverage effect drives Social Security’s impact on top wealth shares

Panel B of Figure 10 illustrates that the impact of Social Security on inequality trends is the consequence of two facts: (i) Social Security wealth grew faster than marketable wealth and (ii) the top 1% share in Social Security wealth has declined.

Both are primarily the consequence of leverage: for working-age individuals, 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. Recent research points out that long-duration assets have outperformed short-duration assets over the last 30 years (Binsbergen, 2020). Consequently, the change in the present value of Social Security benefits (the long position) outpaces the present value of taxes to be paid into the program (the short position), leading to a rapid increase in Social Security wealth relative to other asset classes (Table 1, Row 3).

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

This table decomposes the increases in Social Security wealth between 1989 and 2016. Columns (a) and (b) respectively report the present values of future benefits and taxes in the net present value of Social Security in 1989. Columns (c) and (d) report the percentage increase in the present values of benefits and taxes between 1989 and 2016. 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%	229%	-129%	407%	118%	780%
Top 1%	123%	-23%	303%	282%	308%
Entire population	225%	-125%	405%	119%	763%

This exemplifies the “subtle redistribution” from interest rate declines highlighted by Auclert (2019), who notes that asset holders do not universally benefit when rates fall: instead, wealth is transferred “away from those who are invested primarily in short-term assets, in favor of those with large long-dated investments.” This insight explains why the redistributive effect of Social Security has been magnified by a declining yield curve.

The increase in Social Security wealth is especially pronounced at the bottom of the wealth distribution. As the last column of Table 1 shows, the value of Social Security has increased

472 percentage points more for the bottom 99% than for those in the top 1%. This is an age effect. Being in the top of the wealth distribution is a life-cycle phenomenon. For older workers, who are in the top 1%, Social Security has a shorter duration compared to younger workers, who disproportionately comprise the bottom 99%. This is because older workers are about to receive benefits; whereas younger workers, with many years until retirement and more taxes in front of them, have a more highly levered position.

6.3 Why lower interest rates make Social Security more valuable

The growth in Social Security wealth since 1989 is largely a consequence of declining interest rates. This increase accrues disproportionately to those outside of the top of the wealth distribution, and thus interest rate changes are the primary driver of Social Security's attenuation of top wealth shares. A natural question is: should wealth increases driven by interest rate changes matter for our understanding of inequality trends? If benefits are not substantially larger than in 1989, is Social Security truly more valuable to households than it was in 1989?

In short, yes. To see why, 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 approximately \$0.38 per year over the next two decades.¹⁴ If the interest rate is closer to 1%, as in 2016, this household needs to save \$0.82 annually. Now, if the internal 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 ($\$0.38/\0.61) of private saving in 1989, but to \$1.34 ($\$0.82/\0.61) in 2016. Said another way, the future purchasing power of \$1 of Social Security contribution corresponds to more private saving when rates are low.¹⁵ Discounting based on market rates also captures the fact that the attractiveness of Social Security depends on its

¹⁴The annual contribution to the private annuity in this case is given by $(1 + r)^{-20}$ since the horizons for contribution and retirement are identical.

¹⁵The 1% rate in 2016 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. 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%.

rate of returns relative to comparable investment opportunities on the private market.

Incidentally, whether the increase in the present value of Social Security cash flows 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. To compute shares correctly, it is important that one dollar of Social Security wealth affords 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 2016 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.

Considering the liability side of the household balance sheet is also instructive. 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 2016. Therefore, they register as a greater liability in the computation of marketable net wealth in 2016 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 equivalent increase in Social Security taxes.

It is worth noting that in contemporaneous work, [Sabelhaus and Volz \(2020\)](#) ignore the impact of declines in the interest rate environment on Social Security cashflows, and instead assume a fixed 2.8% discount rate over the last three decades. But this is errant, and thus unsurprisingly paints an inaccurate picture of the evolution of top wealth shares. In economic terms, cashflows generated by Social Security and marketable assets must enter the intertemporal budget constraint of households in a consistent way.

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 percent of Social Security’s growth (50 percent with risk-free valuation).


Table 2: Decomposing the increase in Social Security wealth

This table shows the relative contribution of different effects on Social Security wealth. The first row is calculated as the difference between log per capita Social Security wealth in 2016 and log per capita Social Security wealth in 2016 under the 1989 yield curve. The second row is computed by subtracting log per capita Social Security wealth in 2016 under the 1989 yield curve from Social Security wealth in 2016 under the 1989 age distribution and yield curve. The third row is computed by subtracting log per capita Social Security wealth in 2016 under the 1989 age distribution and yield curve from log per capita Social Security wealth per capita in 2016 under the 1989 age distribution, yield curve, and survival probabilities. The fourth row is computed by subtracting log per capita Social Security wealth in 2016 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 $\log(SSW^{2016}) - \log(SSW^{1989})$ where both terms are calculated under the 2016 and 1989 populations, life expectancies, benefit policies, and yield curves, respectively.

	Valuation method	
	Risk-free	Risk-adjusted
Change in yield curve	1.079	0.990
Shift in age distribution	0.239	0.291
Life expectancy	0.211	0.206
Social Security expansion & other	0.328	0.365
Log total per capita	1.857	1.852
Population growth	0.303	0.303
Log total	2.160	2.155

Figure 11: 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 2016 and 1989 using the risk-adjusted valuation method. The risk-adjusted valuation includes a risk adjustment for the long-run cointegration between the labor and stock markets, as explained in Section 5.



7.2 Shifts in the composition of wealth

In Figure 11, we report 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 2016, nor has its composition. On the other hand, while the share of the bottom 90% remains the same as in 1989, its composition has changed dramatically. In 1989, Social Security only represented 19.8% of the total wealth of the bottom 90%. Because the rate of return on Social Security contributions was substantially lower than interest rates, Social Security wealth was even negative for households below age 40. In 2016, Social Security represents 56.8% 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.¹⁶

7.3 Comparing Social Security and private wealth

We document significant growth in Social Security wealth in recent decades, which has outpaced the growth in private wealth. A long line of theoretical and empirical literature makes clear why ignoring this increase in studies of inequality paints an incomplete picture, since individuals substitute away from private wealth accumulation as social insurance programs become more generous (Attanasio and Brugiavini, 2003; Attanasio and Rohwedder, 2003; Feldstein, 1974).

Still, 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 (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 8.2). 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, Birinci, McGrattan and See, 2020). Further, unless beneficiaries die prematurely, retirement benefits

¹⁶This figure also illustrates the importance of considering within-cohort inequality measures, since aggregate trends are driven by life-cycle dynamics.

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, Miller and Sarin, 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.

7.3.1 Illiquidity adjustment

It is interesting to note that a significant share of Social Security wealth accrues to those for whom liquidity constraints are not binding: As Appendix Figure E.5 shows, of those 50 and older (to whom 79% of Social Security wealth accrues), 64% have more than \$10,000 in accessible wealth—and 44% have more than \$100,000. This is consistent with existing empirical evidence, for example Goda, Ramnath, Shoven and Slavov (2018). Still, in Appendix Figure E.6 we consider the relevance of illiquidity to our valuation of Social Security by applying a premium of 1%, 2%, or 3% to Social Security cashflows.¹⁷ Even with a large liquidity premium, top wealth shares remain substantially attenuated by Social Security’s inclusion.


7.4 Adjusting previous studies on wealth inequality

Figure 12 illustrates the impact of Social Security wealth on existing wealth inequality estimates. We begin in Panel A by presenting four prior estimates of the evolution of the top 1% share: from the Survey of Consumer Finances; Saez and Zucman (2016) which capitalizes income tax returns using a constant rate of return within an asset class; Smith, Zidar and Zwick (2020) which adjusts the Saez and Zucman (2016) estimates to account for heterogeneous returns; and Batty et al. (2019) which relies on data from the Distributional Financial Accounts series. There are significant discrepancies between existing top wealth share estimates (Kopczuk, 2015); however, they all show an upward trend for the top 1% share over our sample period.

¹⁷This is naïve, since the appropriate discount would vary by individual. Since our approach is based on arriving at a market value for Social Security claims if held by a diversified investor, this exercise is outside our scope, as there is no reason to believe that the private and market value of Social Security claims coincide.

Figure 12: Social Security wealth and other measures

This figure shows the wealth shares for the top 1% from other studies before (Panel A) and after (Panel B) adding aggregate Social Security wealth from the risk-adjusted specification. Wealth shares come from the SCF, Saez and Zucman (2016), Smith, Zidar and Zwick (2020), and Batty et al. (2019).

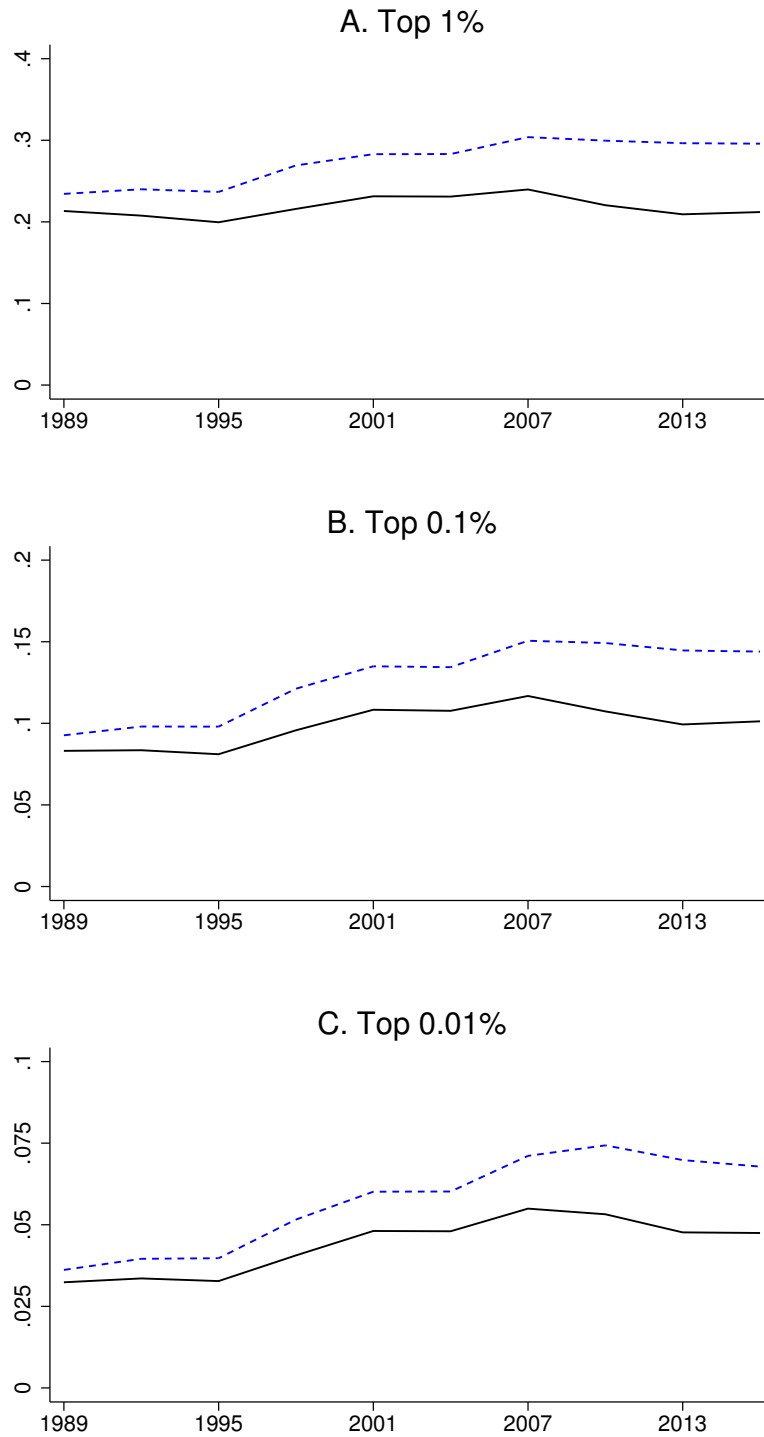


Once Social Security wealth is included, this trend disappears. Given Social Security’s progressivity, it is unsurprising that its inclusion has the effect of scaling down the top 1% estimates. Without Social Security, the top 1% wealth share ranges from 25% to 40% of total wealth. With Social Security included, this share drops to 20% to 30%. What is even more striking than the level effect is the impact on inequality trends: depending on the series, there remains either only a minimal increase in the top 1% share over our sample period, or even a decrease.

We can also examine the impact of the inclusion of Social Security wealth on the evolution of alternative definitions of top wealth. Figure 13 shows how top 1%, 0.1%, and top 0.01% wealth shares evolve under Smith, Zidar and Zwick (2020)’s assumption of heterogeneous returns with asset classes. Again, the impact on inequality trends is significant: there is little difference in top wealth estimates with and without Social Security in 1989, but these have expanded significantly over the last three decades, such that the gap between the series is large and growing.

Figure 13: Adding Social Security to Smith, Zidar and Zwick (2020)

In this figure, we add Social Security wealth to the top wealth shares estimates in Smith, Zidar and Zwick (2020). Shares with (black) and without (blue) Social Security wealth are shown in each panel.



Without Social Security wealth shares from 1989 to 2016 for the top 1%, 0.1%, and 0.01% have increased by 6.2, 5.1, and 3.2 percentage points, respectively; with Social Security the top 1% share has stayed flat and the top 0.1%, and 0.01% shares have increased by only 1.8 and 1.5 percentage points, respectively.

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.

Our overall conclusion—that the inclusion of Social Security substantially attenuates the growth in top wealth shares—is not sensitive to our assumptions. The top 10% and 1% shares of marketable wealth (excluding Social Security) rose by 10.2 and 9.2 percentage points 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 Alternative definition of Social Security wealth as accrued benefits

Our definition of Social Security wealth is the present value of future benefits, net of taxes that workers will pay into the program over their remaining years in the labor force. It may seem appropriate to instead count as wealth only the portion of Social Security wealth that current workers have already accrued. This is problematic for several reasons. First, it does not enable an apples-to-apples comparison with other forms of wealth (e.g., private business wealth, which drives much of the recent increase in marketable wealth inequality), where the market value captures the present value of disbursements as well as the net present value of future benefits. Second, this approach would fail to consider the entire earnings history of workers in ways that make applying the Social Security benefits formula complex. Workers just starting in the labor force will appear to have low average indexed yearly earnings, and thus higher replacement rates on their past contributions than they will eventually receive.

Table 3: Robustness checks

This table reports the evolution of top wealth shares under different assumptions. 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 take into account that high (low) earners live longer (less long) by assigning them lower (higher) mortality rates, based on [Chetty et al. \(2014\)](#). Second, we assume that expected wage growth declined linearly from 1% in 1989 to 0% in 2016. Finally, we use the value of accrued benefit as an alternative measure of Social Security wealth.

	Share of Top 10%			Share of Top 1%		
	1989	2016	Change	1989	2016	Change
Panel A: Baseline results						
Marketable wealth	67.5%	77.7%	10.2%	31.0%	40.2%	9.2%
Risk-free valuation	62.4	59.1	-3.3	27.7	28.3	0.6
Risk-adjusted valuation	63.4	62.2	-1.3	28.3	29.9	1.6
Panel B: Funding gap						
Benefit cut (Intermediate Cost)	63.5	64.1	0.6	28.3	31.0	2.7
Benefit cut (High Cost)	64.2	66.7	2.5	28.7	32.5	3.9
Tax hike (Intermediate Cost)	63.4	62.4	-1.0	28.3	30.0	1.8
Tax hike (High Cost)	63.8	63.7	-0.1	28.5	30.7	2.2
Panel C: Robustness						
Life expectancy differential	63.4	62.2	-1.2	28.2	29.7	1.5
Declining wage growth	63.7	63.8	0.1	28.4	30.8	2.4

It is thus incorrect to single out Social Security wealth as based on accrued wealth rather than the net present value of future cashflows. But we can test to see whether this alternative definition impacts our headline result.

Because returns on past contributions depend on future earnings in non-trivial ways, there is no obvious way to determine how much benefits current workers have already accrued. We say that the fraction of expected benefits already accrued is equal to the share of a worker’s lifetime contributions that have already been paid. Specifically, we define the value of accrued benefits as:

$$\text{Accrued Benefits}_{it} = \frac{\text{Value of Past Taxes}_{it}}{\text{Value of Past and Future Taxes}_{it}} \times \sum_{s=c+66}^T \left(\prod_{k=t}^{s-1} (1 - m_{itk}) \right) \frac{\mathbb{E} [\text{Benefits}_{it}]}{(1 + r_{ts})^{s-t}} \quad (23)$$


where past taxes are compounded using historical realized returns on 30-year government bonds

whereas future taxes are discounted using the risk-adjusted yield curve as in Equation (20).

Even if we adopt a definition of Social Security wealth based on accrued benefits, our baseline results are qualitatively unchanged, as Figure 14 shows.

Figure 14: Top 10% and Top 1% wealth shares — Accrued benefits

This figure shows the top 10% and top 1% wealth shares with and without the present value of accrued Social Security benefits, which is computed as the value expected benefits multiplied by the ratio of past to expected lifetime contributions. Past contributions are capitalized using historical returns on 30-year Treasury bonds. Benefits and future contributions are discounted as in the risk-adjusted valuation.



Under this definition, the inclusion of Social Security wealth still significantly attenuates the growth in top wealth shares: the increase in the top 10% (1%) share falls from 10.2 (9.2) percentage points to 0.8 (2.6) percentage points.

8.2 Accounting for Social Security policy risk

Our baseline calculations value Social Security wealth based on the net present value of promised benefits to individuals (although, as described above, an accrued benefits concept produces similar results). One complexity for either definition is that the trust fund's depletion is imminent, such

that within the next 15 years, absent entitlement reform, the SSA will not be able to meet their full obligations to beneficiaries. It is strange to count as wealth Social Security benefits that program participants may never receive. Thus, it is imperative to account for this policy risk and ascertain what consequences it has for the evolution of top wealth shares.

Related work by [Sabelhaus and Volz \(2020\)](#) accounts for policy risk by applying a constant 2.8% discount rate to Social Security cashflows, suggesting that risk associated with the expected insolvency of the program would “likely offset” the decrease in interest rates. Existing evidence suggests it is unlikely that assuming policy risk offsets the recent changes in the yield curve. For example, [Luttmer and Samwick \(2018\)](#) conclude individuals are willing to forgo 6 percent of benefits to remove policy risk associated with future Social Security benefits. For a 45-year old who will first receive benefits in 20 years, this suggests a discount rate effect on the order of 40 basis points, much lower than the 180 assumed by [Sabelhaus and Volz](#).

The approach we adopt—arriving at a market value of Social Security based on market rates of return—is consistent with the prior literature in this area, including [Geanakoplos and Zeldes \(2010\)](#) and [Blocker, Kotlikoff, Ross and Vallenas \(2019\)](#). It is also the approach taken for all categories of marketable wealth, which are valued at their market price, not their private value.

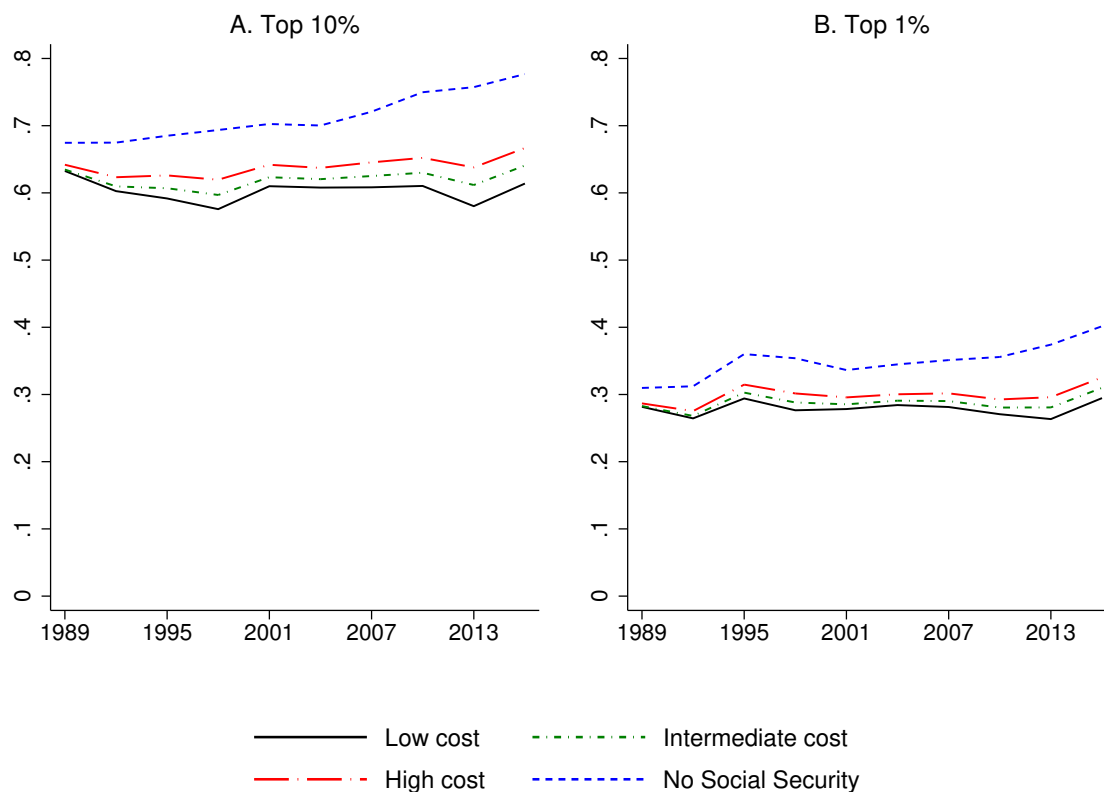
We instead account for policy risk by directly adjusting the cashflows that beneficiaries will receive or the taxes they will pay, as described below. Even under the most conservative assumptions—that beneficiaries will receive only benefits that are payable at current tax rates (eventually cutting benefits by up to 40%), or that taxes will rise for all but the top of the wealth distribution—the substantial impact Social Security has on estimates of wealth inequality is unchanged.

8.2.1 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 (alternative I), intermediate (alternative II), and high cost (alternative III). Appendix Figure [E.4](#) reports the proportion of payable benefits under each of the SSA’s 1989 and 2016 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 (e.g., benefits cuts borne disproportionately by the wealthy).

Figure 15: Top 10% and Top 1% wealth shares — Funding gap adjustment

This figure presents top 10% (Panel A) and 1% (Panel B) wealth shares under four, risk-adjusted specifications. The “No Social Security” specification shows the wealth shares of only non-Social Security assets in SCF. The “Low cost” specification assumes that Social Security benefits are not reduced (Alternative I) and that aggregate wage growth will be higher than in our standard risk-adjusted specification. The “Intermediate cost” specification assumes that Social Security benefits are reduced under the SSA’s intermediate assumptions (Alternative II). The “High cost” specification assumes that Social Security benefits are reduced under the SSA’s high cost assumptions (Alternative III) and that aggregate wage growth will be lower than in our standard risk-adjusted specification.



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, as well as on 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.

The appropriate discount rate would reflect the risk-neutral probability that accounts for the risk inherent to each of the SSA’s proposed scenarios. It is hard to know the right weight to put on each of these outcomes; however, to arrive at a lower bound for the value of Social Security, we can assume that the worst-case scenario (the high cost Alternative III from the SSA’s assumptions) will be realized with probability 1.

Assuming maximal cuts to expected Social Security benefits decreases the bottom 99% wealth share by 2.7 percentage points, wiping out a quarter of Social Security's impact. But as Figure 8.2.1 shows, top wealth shares are still significantly attenuated. This is for two reasons. First, for people close to retirement, the impact of the funds depletion is small, since benefits will pay out as normal for the first 10-15 years after the fund is extinguished. Second, even for cohorts impacted, 60% of expected Social Security benefits represents a sizable sum relative to their marketable wealth.

8.2.2 Balancing the budget by raising taxes

An alternative to benefit cuts would be tax changes to ensure that promised benefits can be paid in full. Along these lines, recent proposals to replenish the Social Security trust fund call for raising payroll taxes at the top of the distribution.¹⁸ The incidence of any tax changes will impact our estimates of Social Security wealth and its distribution.

In theory, an increase in Social Security taxes borne by the *bottom* of the wealth distribution could attenuate our results. The idea would be that higher taxes will decrease the value of Social Security for the bottom of the wealth distribution. To assess this possibility, we adopt the most conservative assumption from the perspective of our baseline results: the possibility that a tax hike to replenish the trust fund will be borne entirely by the bottom 90%, or bottom 99%. Although the top 10% and top 1% shares rise, this increase is slight: by -0.1 and 2.2 percentage points, respectively, from 1989–2016.¹⁹ Interestingly, assuming taxes rise on those at the bottom of the wealth distribution to cover the trust fund's imminent shortfall has less of an impact on their Social Security wealth than assuming benefit cuts. This is because tax hikes, unlike decreases in benefits, push a portion of the consequences of the gap between promised and payable benefits to future generations not yet in the labor force.

¹⁸For example, the Biden campaign has called for an increase in payroll taxes for those making more than \$400,000 annually, who will be required to pay in more to the program during their lifetimes, with no corresponding increase in benefits. <https://budgetmodel.wharton.upenn.edu/issues/2020/3/6/biden-social-security>

¹⁹Table 3, using the SSA's high-cost assumption. The intermediate cost assumption is even less consequential for top wealth shares because the fund shortfall is lower in this case.

8.3 Decline in productivity growth

Another potential concern is that the recent decline in interest rates is symptomatic of lower long-run economic growth, as the secular stagnation hypothesis suggests (Summers, 2014). Since Social Security is wage-indexed, this low growth could decrease our estimate of Social Security wealth and thus attenuate its impact on inequality trends. Our baseline estimates already assume a decline in the growth rate of wages: we rely on assumptions from SSA reports, which, as of 2016, assumed a 1.2% long-term annual wage growth rate, down from 1.7% in 1989.

Table 3 considers the possibility that wage growth declines even more significantly than these SSA projections reflect. Specifically, we assume that the real growth rate of wages declines linearly from 1% to 0% between 1989 and 2016. Our main result is qualitatively unchanged: including Social Security offsets the growth in top wealth shares.

8.4 Differences in life expectancy

Chetty et al. (2016) document that individuals with higher earnings live longer. These differences are large: average 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). Because those at the top of the lifetime income distribution live longer on average, their stream of benefits is longer, and their total Social Security wealth will be understated by using cohort average life expectancy. Similarly, we will overstate the Social Security wealth of those at the bottom of the distribution.

Therefore, we adjust for these differences in life expectancy using data from the Health Inequality Project (HIP), by allowing survival probabilities of SCF respondents receiving Social Security retirement benefits to differ based on their place in the income distribution.²⁰ Kreiner, Nielsen and Serena (2018) suggest life expectancy differentials are less than half as large when income mobility is accounted for as the HIP data suggest. Our adjustment can thus be understood as an imperfect upper bound on the role these differences play on the stock and distribution of Social Security wealth.

Our adjustment effectively makes high income retirees younger and low income retirees older, a procedure which we outline in detail in Appendix B.3. Figure 16 shows how a life expectancy

²⁰We proxy for the permanent income distribution using the Social Security benefits distribution because benefits are, by construction, a proxy for lifetime earnings.

adjustment impacts Social Security wealth across the Social Security benefits distribution among current retirees. When differences in mortality rates are accounted for, per capita Social Security wealth that accrues to the bottom decile falls by nearly 25% percent, 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.

Figure 16: 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.

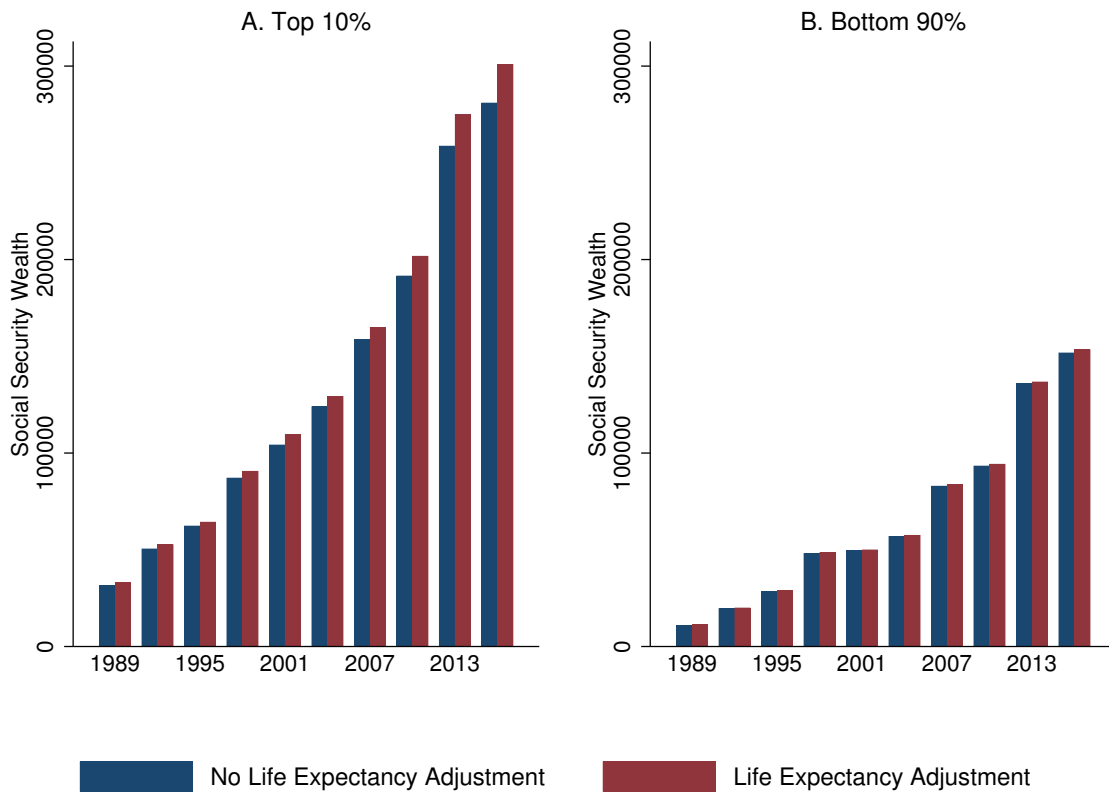


Table 3 shows average Social Security wealth for the top 10% and bottom 90%, both with, and without, adjusting for differences in life expectancy. These differences increase the average stock of Social Security wealth that accrues to the top wealth decile by approximately 7.1% percent in 2016. The Social Security wealth of the bottom 90% increases as well, by roughly 1.1% 2016. 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.²¹

8.5 Most conservative adjustment

Our headline result is that the inclusion of Social Security in estimates of top wealth shares attenuates the increase documented by prior work. A reasonable question is whether this result is sensitive to the alternative assumptions discussed with respect to the risk that promised benefits will not be fully paid or that taxes will increase, declining productivity growth, or differences in life expectancy between the rich or the poor. Table 3 makes clear that any individual adjustment has no impact on our qualitative finding, and relatively little impact quantitatively: the only scenario that results in an increase in, for example, top 10% wealth shares once Social Security is included is when policy risk is resolved by a 40% decrease in benefits for current beneficiaries. Even then, the increase in the top 10% (1%) share is only 2.5 (3.9) percentage points, relative to the 10.2 (9.2) percentage point increase in top wealth shares when Social Security is excluded.

Collectively, moving from the lowest estimate of top wealth shares in 1989 to the highest estimate in 2016 (inclusive of Social Security) would raise the top 10% share only slightly from 63.4% to 66.7% (top 1% from 28.3% to 32.5%). Any mix of assumptions about policy risk, the macroeconomic environment, or life expectancy lead to the same conclusion: Social Security's inclusion substantially attenuates the growth in marketable wealth inequality.

²¹It 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.

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. This paper builds on past work by incorporating Social Security into inequality estimates. We find that top wealth shares have not increased once the old age retirement program is accounted for.

This is because Social Security wealth has risen: In 1989, Social Security represents 23.7% of the wealth held by the bottom 90% of the wealth distribution. By 2016, this share had grown to 62.2%. Even after adjusting for systematic risk, Social Security rose from only 19.8% of the total wealth of the bottom 90% to 56.8%.

Since Social Security and private wealth are substitutes ([Feldstein, 1974](#)), a narrow definition of wealth paints an incomplete picture of inequality trends. Our risk-adjusted estimates suggest that between 1989 and 2016 the top 10% share declined by 1.3 percentage points and the top 1% share increased only slightly by 1.6 percentage points. This differs drastically from recent work that excludes Social Security and finds the top 10% and 1% shares rose by around 10 percentage points over this period.

The top wealth estimates in this paper are still overstated, because we exclude programs like disability insurance and Medicare, which accrue disproportionately to the bottom of the wealth distribution. Overall, this paper makes the point 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.
- Batty, Michael, Jesse Bricker, Joseph Briggs, Elizabeth Holmquist, Susan McIntosh, Kevin Moore, Eric Nielsen, Sarah Reber, Molly Shatto, Kamila Sommer, Tom Sweeney, and Alice Henriques Volz**, “Introducing the Distributional Financial Accounts of the United States,” 2019. Finance and Economics Discussion Series 2019-017.
- 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.
- Blocker, Alexander W., Laurence J. Kotlikoff, Stephen A. Ross, and Sergio Villar Vallenias**, “The True Cost of Social Security,” *Tax Policy and the Economy*, 2019, 33.
- Catherine, Sylvain**, “Labor Market Risk and the Private Value of Social Security,” 2019. SSRN Electronic Journal.
- **, Max Miller, and Natasha Sarin**, “Relaxing Household Liquidity Constraints through Social Security,” *Journal of Public Economics*, 2020, *Forthcoming*.
- 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.
- **, Michael Stepner, Sarah Abraham, Shelby Lin, Benjamin Scuderi, Nicholas Turner, Augustin Bergeron, and David Cutler**, “The association between income and life expectancy in the United States, 2001-2014,” *Journal of the American Medical Association*, 2016, 315 (16), 1750–1766.
- Cochrane, John**, “Wealth and Taxes,” Technical Report, Cato Institute 2020. Tax and Budget Bulletin No. 86.
- Deaton, Angus S. and Christina Paxson**, “Saving, Growth, and Aging in Taiwan,” in David Wise, ed., *Studies in the Economics of Aging*, National Bureau of Economic Research, 1994, pp. 331–362.
- 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.
- Dushi, Irena, Howard M. Iams, and Trenkamp Brad**, “The Importance of Social Security Benefits to the Income of the Aged Population,” *Social Security Bulletin*, 2017, 77 (2).
- 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.
- and **Jeffrey B. Liebman**, “The Distributional Aspects of an Investment-Based Social Security System,” in “The Distributional Aspects of Social Security and Social Security Reform” 2002.
- 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.
- , **Olivia S. Mitchell**, and **Stephen P. Zeldes**, “Social Security Money’s Worth,” in Olivia S. Mitchell, Robert J. Meyers, and Howard Young, eds., *Prospects for Social Security Reform*, University of Pennsylvania Press, 1999, pp. 80–151.
- Goda, Gopi Shah, Shanthi Ramnath, John B. Shoven, and Sita Nataraj Slavov**, “The financial feasibility of delaying Social Security: evidence from administrative tax data,” *Journal of Pension Economics & Finance*, 2018, 17 (4).
- Gürkaynak, Refet, Brian Sack, and Jonathan Wright**, “The U.S. Treasury Yield Curve: 1961 to the Present,” *Journal of Monetary Economics*, 2007, 54 (8), 2291–2304.
- , —, and —, “The TIPS Yield Curve and Inflation Compensation,” *Finance and Economics Discussion Series*, 2008, 2008 (5).
- Gustman, Alan L. and Thomas L. Steinmeier**, “Retirement in Dual-Career Families: A Structural Model,” *Journal of Labor Economics*, 2000, 18 (3), 503–545.
- and —, “How Effective is Redistribution Under the Social Security Benefit Formula?,” *Journal of Public Economics*, 2001, 82 (1), 1–28.
- , **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.
- , **Thomas L. Steinmeier**, and **Nahid Tabatabai**, “The Growth in Social Security Benefits Among the Retirement-Age Population from Increases in the Cap on Covered Earnings,” *Social Security Bulletin*, 2012, 72 (2), 49–61.
- Guvonen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song**, “What Do Data on Millions of U.S. Workers Reveal About Life-Cycle Earnings Risk?,” 2019. Federal Reserve Bank of New York Staff Report No. 710.
- , **Greg Kaplan**, **Jae Song**, and **Justin Weidner**, “Lifetime Incomes in the United States Over Six Decades,” 2018. NBER Working Paper Series No. 23371.
- Kopczuk, Wojciech**, “What do we know about the evolution of top wealth shares in the United States?,” *Journal of Economic Perspectives*, 2015, 29 (1), 47–66.
- Kreiner, Claus Thustrup, Torben Heien Nielsen, and Benjamin Ly Serena**, “Role of income mobility for the measurement of inequality in life expectancy,” *Proceedings of the National Academy of Sciences*, 2018, 115 (46), 11754–11759.
- Liebman, Jeffrey B.**, “Redistribution in the Current U.S. Social Security System,” in Martin Feldstein and Jeffrey B. Liebman, eds., *The Distributional Aspects of Social Security and Social Security Reform*, National Bureau of Economic Research, 2002, pp. 11–41.
- Luttmer, Erzo F. P. and Andrew A. Samwick**, “The Welfare Cost of Perceived Policy Uncertainty: Evidence from Social Security,” *American Economic Review*, 2018, 108 (2).
- Modigliani, Franco and Richard Brumberg**, “Utility Analysis and the Consumption Function: An Interpreta-

- tion of Cross-section Data,” in Kenneth K. Kurihara, ed., *Post-Keynesian Economics*, Rutgers University Press, 1954.
- Piketty, Thomas**, *Capital in the 21st Century*, Cambridge: Harvard University Press, 2013.
- , **Emmanuel Saez, and Gabriel Zucman**, “Distributional National Accounts: Methods and Estimates for the United States,” *The Quarterly Journal of Economics*, 2018, *133* (2), 553–609.
- Poterba, James M.**, “Retirement Security in an Aging Population,” *American Economic Review*, 2014, *104* (5), 1–30.
- Sabelhaus, John and Alice Henriques Volz**, “Social Security Wealth, Inequality, and Lifecycle Saving,” 2020. NBER Working Paper No. 27110.
- 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.
- and —, *The Triumph of Injustice: How the Rich Dodge Taxes and How to Make Them Pay*, New York: W.W. Norton, 2019.
- 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.
- Social Security Administration**, “Monthly Statistical Snapshot, October 2016,” 2016. Research, Statistics & Policy Analysis.
- Summers, Lawrence H.**, “U.S. Economic Prospects: Secular Stagnation, Hysteresis, and the Zero Lower Bound,” *Business Economics*, 2014, *49* (2).
- 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.

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 SCF variables

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

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

which allows us to calculate annual benefits, given by

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

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

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

and observed survivor benefits given by

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

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

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

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. 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.5](#), 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

$$\text{liquid_wealth} = \text{liq} + \text{cds} + \text{nmmf} + \text{stocks} + \text{bond} + \text{retqliq} + \text{savbnd} + \text{homeeq}.$$

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. 2016 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.1 shows the differences in the yield curve assumptions implied from Treasuries notes from [Gürkaynak et al. \(2007\)](#) 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 ([Summers, 2014](#)). An additional piece of evidence of the issues with the SSA’s approach comes from the 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

$$\text{ret.discount_years} = \text{full_retirement_age} - \text{ret.age}.$$

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 2016. This person was born in 1954, meaning that the full retirement age for his/her cohort is 66 years old. For this person, we have $\text{Adjustment} =$

$(1 - \frac{5}{9} \times 36 - \frac{5}{12} \times 12)$, meaning that the full retirement benefit is given by

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

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 et al. \(2014\)](#) 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 2014-2016.

Using these data, we compute the number of years less (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, we 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 less (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 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 2016 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 has an effective mortality age of 49 years old, and he

would be assigned the survival probabilities of a 49 year old man in 2016. We apply this life expectancy correction both to retired workers and to those still in the workforce, whose earnings histories we simulate.

Our baseline exercise requires ascribing each cohort’s Social Security wealth to centiles of the marketable wealth distribution. Interestingly, there is not a one-to-one relationship between where workers fall on the income distribution and the marketable wealth distribution. This attenuates the impact that adjustments for differences in life expectancy—based on where individuals fall in the income distribution—have on our results, since those in the bottom of the income distribution fall in different deciles of the marketable wealth distribution, as Appendix Figure E.3 makes clear.

B.4 Capitalizing implied survivor benefits

Widows 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. We then adjust survivor benefits to reflect the maximal benefits that a surviving spouse can receive, as detailed [here](#). We adjust our survival benefit calculations such that the received benefits do not exceed the maximal family threshold. Once the maximum benefit is calculated, the wealth coming from the implied survivor benefits is given by

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

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.

We estimate this using regressions in the style of [Deaton and Paxson \(1994\)](#) for each gender, which is a constrained regression of the following form

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

subject to

$$\sum_{1989}^{2016} \gamma_t = 0 \quad (\text{B.2})$$

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

$$\eta_{72} = 0. \quad (\text{B.4})$$

where a represents each age, t each survey year, and b each birth year.²² 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.2, 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 Step-by-step guide to 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.
 - (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 2016, this function is shown in Panel B of Figure 7.
2. Determine proportion of the population of each cohort in the top 10% and top 1% of wealth distribution, which we denote by $k_{c,t}^g$ 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 $k_{c,t}^{\text{Top 10\%}} \equiv \frac{N_{c,t}^{\text{Top 10\%}}}{N_{c,t}^{\text{Full}}}$ and the top 1% share by $k_{c,t}^{\text{Top 1\%}} \equiv \frac{N_{c,t}^{\text{Top 1\%}}}{N_{c,t}^{\text{Full}}}$, where $N_{c,t}^g$ is the total size of cohort c in survey t in population g . Mathematically, this is given by $N_{c,t}^g \equiv \sum_i \text{wgt}_i \mathbb{1}_i(c, t, g)$ 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 2016, 7.5% are in the top 10% and 0.4% are in the top 1% of the aggregate 2016 wealth distribution.

²²Note that respondents are grouped into three-year age and birth year cohorts in this estimation.

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

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

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 1989–2013, so for 2016, we fill in the average observed retirement-adjusted Social Security wealth in the data for recipients aged 70–73 from the 2016 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 $\left(\frac{0.8}{0.9} \right) \overline{\text{SSW}}_{a,t}^g$ to account for the increased probability of never

receiving benefits in the subpopulation. Formally, this adjustment is given by

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

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

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

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

$$l_{1,t} = y_t + s_t - \delta t, \quad (\text{D.3})$$

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

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 (15) 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 $t+n$, which is denoted by

$$\beta_t^{L_{1,n}} = \frac{\text{Cov}_t \left(r_t^m dt, r_t^{L_{1,n}} dt \right)}{\text{Var}_t [r_t^m dt]}.$$

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

$$P_t^{L_1, n} = \exp \{ at + b + cy_t + d^T z_t \}, \quad (\text{D.7})$$

where

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

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

$$P_t^{L1,n} = \exp \left\{ P_t^{L1,n} \right\},$$

where

$$P_t^{L1,n} = at + b(n) + c(n)y_t + d^T z_t.$$

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

$$dP_t^{L1,n} = (a - b'(n) - c'(n)y_t - \kappa c(n)y_t) dt + \left(c(n) \begin{bmatrix} \sigma_l \\ -\sigma_s \end{bmatrix} + d \right)^T dz_t, \quad (\text{D.8})$$

where

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

Then, the return on the wage strip equals

$$\begin{aligned}
r_t^{L1,n} dt &= \frac{dP_t^{L1,n}}{P_t^{L1,n}} = dP_t^{L1,n} + \frac{1}{2} (dP_t^{L1,n})^2 \\
&= \left(a - b'(n) + \frac{1}{2} \left(c\sigma_l + \frac{\phi}{\kappa} \sigma_l \right)^2 + \frac{1}{2} \left(\sigma - \frac{\phi}{\kappa} \sigma_s - c\sigma_s \right)^2 \right) dt + \begin{bmatrix} c\sigma_l + \frac{\phi}{\kappa} \sigma_l \\ \sigma - \frac{\phi}{\kappa} \sigma_s - c\sigma_s \end{bmatrix}^T dz_t
\end{aligned}$$

meaning that the expected return is

$$\mathbb{E}_t \left[r_t^{L1,n} \right] = \mu - (\mu - r) \frac{\sigma_s}{\sigma} \left(\frac{\phi}{\kappa} + c \right).$$

This gives the beta on the wage strip as

$$\beta^{L1,n} = \frac{\text{Cov}_t \left(r_t^m dt, r_t^{L1,n} dt \right)}{\text{Var}_t \left[r_t^m dt \right]} = 1 - \frac{\sigma_s}{\sigma} \left(\frac{\phi}{\kappa} + c \right)$$

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

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

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 ($\sigma_s = \sigma$), then the results reduce to

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

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

E Additional figures

Figure E.1: Market Implied and Social Security Administration Yield Curve Estimates

The figure presents the differences between the yield curves implied by treasury markets (Gürkaynak, Sack and Wright, 2007) 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 3.2.




Figure E.2: 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.3: Distribution of Wealth in the Bottom Decile of Social Security Benefits

This figure shows how marketable wealth is distributed among the bottom decile of Social Security beneficiaries. Each bar represents the share of people in each marketable wealth decile. This exercise is done for current Social Security beneficiaries with deciles of marketable wealth computed for individuals between 62 and 76 year of age.




Figure E.4: Funding Gap: Payable Benefits under 1989 and 2016 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 E.5: 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. 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.




Figure E.6: Top 10% and Top 1% wealth shares – Liquidity premium

This figure reports the evolution of the top 10% and 1% wealth shares with and without Social Security wealth when we add 1, 2 or 3 percentage points to our discount rates to reflect a hypothetical liquidity premium.

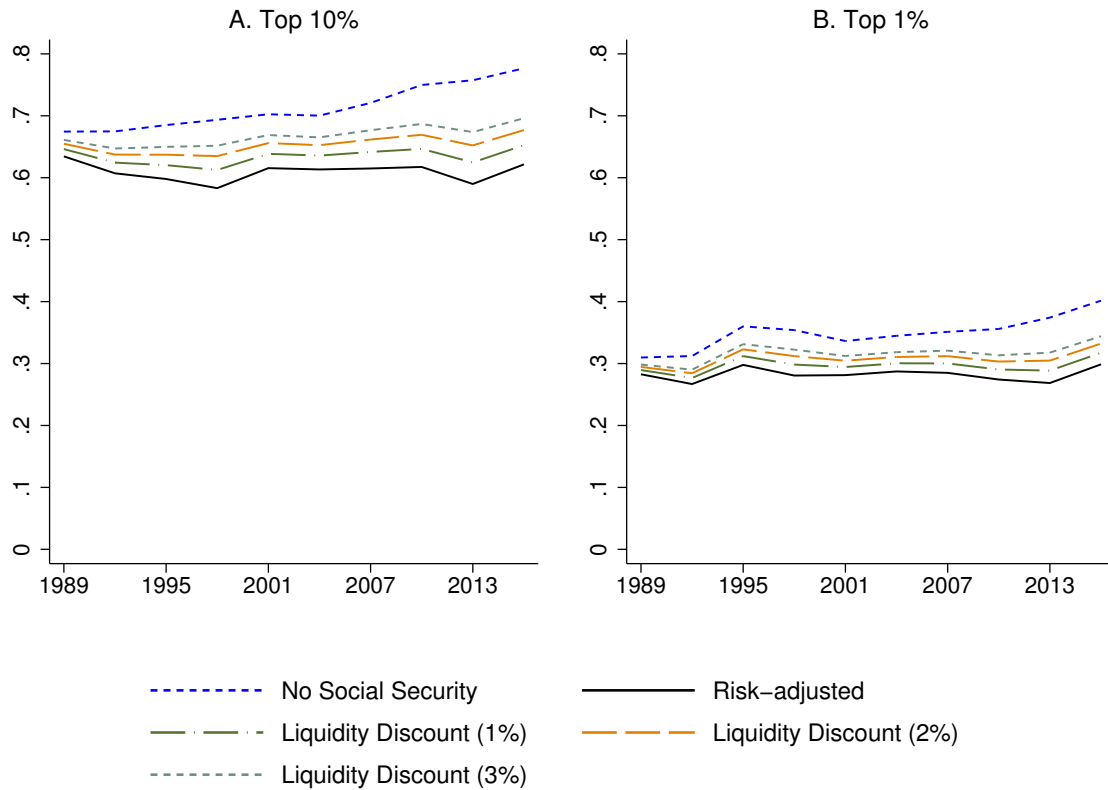


Figure E.7: Social Security wealth by age — Risk-free vs. risk-adjusted

This figure shows average Social Security wealth in 2016 by age for the risk-free and risk-adjusted specifications. The risk-adjusted specification adjusts for the cointegration of stock and labor markets, as detailed in Section 5.1.

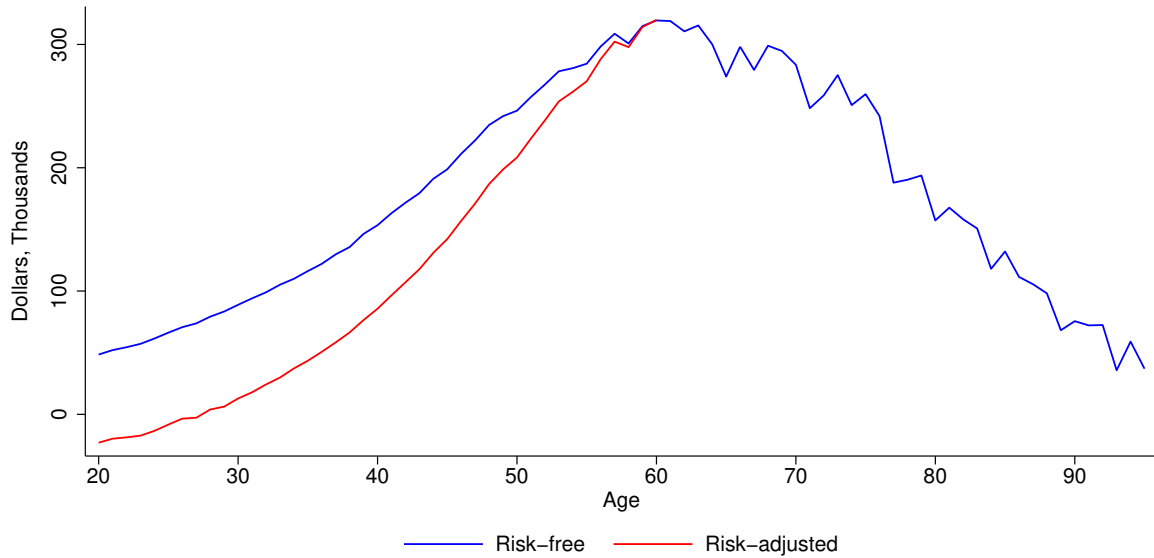


Figure E.8: Total Social Security Wealth: Fixed Yield Curve

This figure shows the present value of Social Security under a fixed yield curve specification. The “Risk-free valuation” shows the present value of Social Security using the market implied yield curve to discount cashflows coming from Social Security benefits. The “Mean yield curve” specification uses the average yield curve from 1989-2016 to discount the cashflows, respectively. The “Sabelhaus and Volz (2020) assumptions” specification uses a fixed 2.8% fixed real discount rate in all years. Finally, the “Sabelhaus and Volz (2020)” specification displays the aggregate Social Security wealth reported in Sabelhaus and Volz (2020). All series are adjusted for inflation by placing them in 2018 dollars.




Figure E.9: Top 10% and Top 1% Wealth Shares–Fixed Yield Curve

This figure shows the top 10% (Panel A) and top 1% (Panel B) wealth shares with and without Social Security included under the market implied yield curve and the average yield curve from 1989-2016. The “Risk-free valuation” shows the present value of Social Security using the market implied yield curve to discount cashflows coming from Social Security benefits. The “Mean yield curve” specification uses the average yield curve from 1989-2016 to discount the cashflows, respectively. The “Sabelhaus and Volz (2020) assumptions” specification uses a fixed 2.8% fixed real discount rate in all years. The “Sabelhaus and Volz (2020)” specification displays the aggregate Social Security wealth reported in Sabelhaus and Volz (2020). Finally, the “No Social Security” specification reports the wealth shares reported in the SCF with the Forbes 400 included. All series are adjusted for inflation by placing them in 2018 dollars.

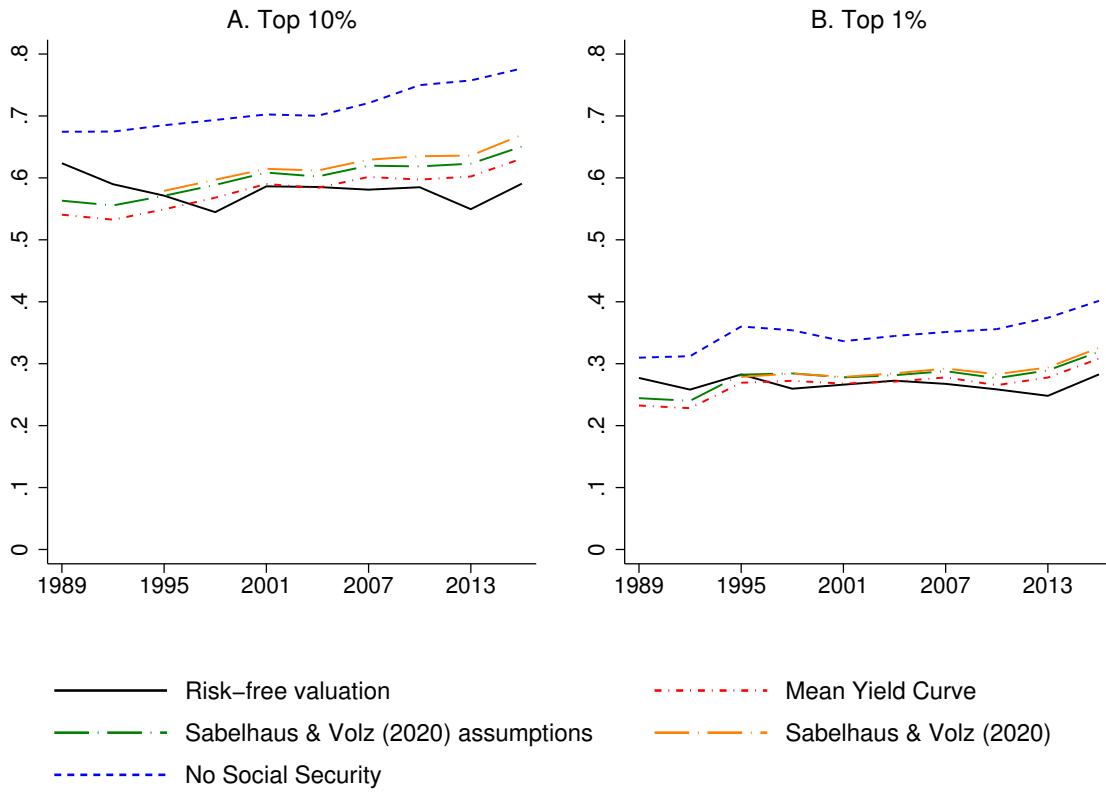


Table E.1: Calibration of labor income process

Parameter	Value	Parameter	Value
ρ	0.958	$\sigma_{\varepsilon,2}$	0.063
p_z	21.9%	σ_α	0.298
$\mu_{\eta,1}$	-0.147	$\sigma_\beta \times 10$	0.185
$\sigma_{\eta,1}$	0.457	$\text{corr}_{\alpha\beta}$	0.976
$\sigma_{\eta,2}$	0.139	$a_\nu \times 1$	-3.2740
$\sigma_{z_1,0}$	0.667	$b_\nu \times t$	-0.8935
λ	0.001	$c_\nu \times z_t$	-4.5692
p_ε	12.6%	$d_\nu \times t \times z_t$	-2.9203
$\mu_{\varepsilon,1}$	0.236	$a_{z_1} \times 1$	0.2191
$\sigma_{\varepsilon,1}$	0.343		