Interest-rate risk and household portfolios*

Sylvain Catherine† Max Miller‡ James D. Paron§
Natasha Sarin¶

October 14, 2023

Abstract

How are households exposed to interest-rate risk? When rates fall, households face lower future expected returns but those holding long-term assets—disproportionately the wealthy and middle-aged—experience capital gains. We study the hedging demand for long-term assets in a portfolio choice model. The optimal interest-rate sensitivity of wealth is hump-shaped over the life cycle. Within cohorts, it increases with wealth and earnings. These predictions fit observed patterns in the United States, suggesting a relatively efficient distribution of interest-rate risk. By protecting workers from rate fluctuations, Social Security limits the welfare consequences of rising wealth inequality when rates fall.

Keywords: Interest rates, Portfolio choices, Inequality, Social Security
JEL codes: D31, E21, G51, H55

---

*We thank seminar participants at the University of Maryland, Dartmouth Tuck, Berkeley Haas, Columbia University, the London School of Economics, the University of Southern California, the University of Texas at Austin, and the University of Washington Seattle for helpful comments.

†University of Pennsylvania, Wharton. Email: scath@wharton.upenn.edu.
‡Harvard Business School. Email: mamiller@hbs.edu.
§University of Pennsylvania, Wharton. Email: jparon@wharton.upenn.edu.
¶Yale Law School. Email: natasha.sarin@yale.edu.

Changes in real interest rates have important consequences for households. First, asset prices rise when interest rates fall, generating capital gains for the owners of long-term assets. These gains are large and unequally distributed. Over the past four decades, the fall in interest rates explained the exceptional stock market returns (Binsbergen, 2021) and greatly contributed to the rise in wealth inequality (Greenwald, Leombroni, Lustig and Nieuwerburgh, 2023). Second, lower rates help households borrow against future income, but also make it more difficult to save for retirement. To fund \$1 of spending at 70 years old, a 35-year-old worker must invest 71 cents at a rate of return of 1%, but only 36 cents if the rate of return is 3%. The fact that capital gains occur when future rates of return deteriorate complicates our understanding of who benefits or loses from long-run interest-rate fluctuations.

The goal of our paper is twofold. First, we seek to understand why certain households choose to invest more in long-term assets. Second, we evaluate the extent to which the distribution of interest-rate risk among households is efficient.

We start by analyzing US households' balance sheets to identify who benefits from changes in the market value of assets and liabilities induced by interest-rate fluctuations. At first glance, the data suggest that interest-rate risk is unequally distributed, as three stylized facts illustrate. First, the interest-rate sensitivity of wealth is hump-shaped over the life cycle. Second, within an age group, the interest-rate sensitivity of wealth increases with both earnings and wealth. In other words, middle-aged and richer households invest more in long-term assets that rapidly appreciate in value when rates fall. Third, over the last 60 years, trends in wealth inequality have tracked changes in the valuation of long-term bonds.

The interest-rate sensitivity of wealth must be interpreted in the broader context of a household's life-cycle problem. Therefore, we set up a parsimonious life-cycle model to assess these stylized facts against a normative benchmark. In our model, households decide how much to consume and choose the interest-rate sensitivity of their wealth by mixing two assets: a short- and a long-term zero-coupon bond. The allocation between these two abstract securities defines the optimal interest-rate sensitivity of their balance sheets. Taking into account households' stochastic stream of labor income, the progressive Social Security system, and stochastic interest rates, we find that the patterns we document in the data are broadly consistent with the model's predictions, reflecting an efficient distribution of interest-rate risk across households.

Why can the model match the US data well? Because investors with longer horizons benefit from more years of compounding, they are more sensitive to the rate of return on wealth. Consequently, absent any background assets, the need to hedge against falling rates by holding long-term assets declines with age, and so does the optimal interest-rate sensitivity of wealth. In the presence of background assets, the same portfolio rule applies to *total* wealth. Long-term background assets reduce the optimal rate sensitivity of financial wealth, while short-term assets do the opposite. Because its duration falls over the life cycle, human capital acts as a substitute for long-term assets during the first decade of working life and the opposite afterwards. The combination of the investment horizon and the changing substitution effects from human capital generates a hump-shaped demand for long-term assets over the life cycle, as observed in the data.

Another important background asset is Social Security. Its role can be understood in two ways. First, because it pays benefits from retirement until death, Social Security remains a long-term asset over most of the life cycle. Equivalently, because it represents an alternative savings vehicle with a rate of return that does not depend on the market rate, Social Security partly insulates workers from interest-rate fluctuations. Either way, it reduces the hedging demand for long-term assets. This substitution effect is larger at the bottom of the earnings distribution because retirement benefits replace a smaller portion of lifetime earnings at the top. Consequently, the demand for long-term assets increases with earnings and wealth within an age group, as observed in the data.

How do these abstract portfolio rules translate to the hedging dynamics that we actually observe? To eliminate interest-rate risk entirely, households could buy a portfolio of zero-coupon bonds of various maturities that, combined with expected future earnings, exactly matches their desired consumption plan. Holding this portfolio to maturity would allow them to finance a preset consumption plan regardless of interest-rate fluctuations because payments would be made exactly when needed. This strategy is only optimal for an infinitely risk-averse agent, but it helps convey economic intuition because, in reality, households engage in both voluntary and mandatory financial transactions that emulate this portfolio.

At the bottom of the earnings distribution, workers save primarily through contributions to Social Security, the rate of return on which does not depend on the evolution of interest rates. Payroll taxes and benefits essentially redistribute labor earnings over the life cycle to match the timing of a smoothed consumption plan. These workers do not need to invest in long-term assets to hedge interest-rate risk, because Social Security automatically does the hedging for them. Lower middle-class households, who have most but not all of their earnings replaced by Social Security, need to complement Social Security with private savings and do so by buying a house with a fixed-rate mortgage. Through this arrangement, they trade a flow of future residential consumption for a stream of future payments at current spot prices. Social Security benefits cover most of their non-residential consumption in retirement. Finally, because they receive the lowest replacement rates from Social Security, higher earners rely more on private savings, such as retirement accounts invested in the stock market. Since stocks are high-duration assets, their value goes up when rates fall, protecting retirement consumption when saving becomes harder.

In complete markets, risk-averse agents share all risks optimally (Arrow, 1964). A direct implication is that the realization of uncertainty should not redistribute welfare across investors with identical preferences, as they would hold the same portfolio and any risk would be equally allocated among them. At first glance, differences in the duration of household balance sheets suggest that interest-rate risk is far from equally distributed and that the rich accrue larger capital gains when interest rates fall. However, our model shows that, to truly bear identical interest-rate risk, households with the same investment horizon and risk preference need to hold different portfolios. These differences align well with those observed in the data.

A consequence of this insight is that trends in wealth inequality induced by interest-rate fluctuations must be interpreted with caution. Indeed, our model generates large changes in the concentration of wealth—roughly half of historical variations since 1960—without any increase in within-cohort *welfare* inequality. As two key empirical facts demonstrate, Social Security plays an important role in explaining this paradox. First, the interest-rate sensitivity of wealth no longer increases with wealth or earnings once Social Security wealth is accounted for. Second, wealth inequality, inclusive of Social Security, did not rise much over the past four decades (Catherine et al., 2023).

Our study shows that portfolio-choice theory is a useful tool for understanding why wealth inequality rises and falls over long periods of time, why the inclusion of Social Security wealth significantly changes these trends, and interpreting their implications for welfare.

Related literature Our paper contributes to the literature on portfolio choices and the dynamics of inequality.

First, we contribute to the study of households' portfolio choices. Following seminal studies by Samuelson (1969) and Merton (1969), this literature has focused on the allocation of wealth between different assets classes, and in particular the stock market portfolio and a riskfree bond (Benzoni, Collin-Dufresne and Goldstein, 2007; Catherine, 2022; Cocco, Gomes and Maenhout, 2005; Lynch and Tan, 2011; Viceira, 2001). Our approach differs in its focus on a risk factor—the real interest rate—rather than a specific asset class. Campbell and Viceira (2001) show that investors can hedge interest-rate risk by holding long-term bonds. We build on this intuition to quantitatively explain household's demand for long-term assets and demonstrate how mortality, human capital, and Social Security shape the optimal hedging demand.

Previous work justifies this fresh angle. First, Lustig, Van Nieuwerburgh and Verdelhan (2013)'s finding that events in the long-term bond market, not the stock market, are the greater source of risk for households motivates our focus on the optimal duration of their portfolio rather than their equity share. Second, for the subset of wealth held in stocks, a key feature is its high interest-rate sensitivity: Binsbergen (2021) shows that, over the last fifty years, duration explains most of the market's exceptional returns. Greenwald, Leombroni, Lustig and Nieuwerburgh (2023) document that wealthy households invest more in long-term assets, which mechanically increases wealth inequality when interest rates fall. We offer a theoretical foundation to these portfolio differences not only across the wealth distribution, but also over the life cycle and across the earnings distribution.

In doing so, we explain an important driver of portfolio-return heterogeneity. The importance of this work is highlighted by Moll (2021), who argues that explaining the portfolio choices that generate heterogeneous returns is first-order. Benhabib, Bisin and Luo (2019), Bach, Calvet and Sodini (2020) and Hubmer, Krusell and Smith (2021) have empirically documented that the higher returns of the wealthy are essential for explaining wealth inequality and its evolution. We provide an explanation for this heterogeneity: wealthier households should invest more in long-term assets to be equally hedged against interest-rate risk.

Recent work by Fagereng, Holm, Moll and Natvik (2021) shows that capital gains explain most of the rise in wealth inequality because the rich "save by holding" instead of selling assets to consume. This behavior is consistent with interest-rate hedging. At the limit, if households hold the perfect interest-rate hedging portfolio, holding it to maturity will let them afford the same consumption plan regardless of discount-rate-induced changes in asset prices. As Cochrane (2022) notes, a family should not pay attention to these changes if it intends to live off of the coupons.

Finally, this paper builds on recent studies that have focused on the ways in which progressive government programs attenuate the incidence of wealth inequality on lifetime consumption (Auerbach, Kotlikoff and Koehler, 2023; Catherine, Miller and Sarin, 2023). Our model shows how the interplay between public programs like Social Security and optimal portfolio choices can generate diverging trends in wealth and consumption inequality. Our mechanism helps explain Meyer and Sullivan (2023)'s finding that, over the past five decades, the rise in overall consumption inequality was small.

1 Data

This section describes our data sources and the methodology we adopt to measure the interest-rate sensitivity of wealth. We provide more details in Appendix A.

1.1 Survey of Consumer Finances

Our goal is to measure the interest-rate sensitivity of assets on households' balance sheets. To do this, our primary data source is the triennial Survey of Consumer Finances (SCF), from which we use all survey years from 1989 to 2019. We use three main series: (i) detailed information on household assets and liabilities, (ii) interest rate and maturity data for household liabilities, and (iii) income data for household assets, which we use to compute valuation ratios. We supplement the SCF with data on Social Security wealth and interest-rate sensitivity from Catherine et al. (2023), duration estimates from Greenwald et al. (2023) and Bloomberg, and the real yield curve, computed by subtracting inflation projections from the SSA annual reports from the nominal yield curve from the Federal Reserve.[1]

1.2 Interest-rate process

Measuring the interest-rate sensitivity of assets requires some structure over how shocks to current interest rates will affect future interest rates. To this end, we assume that log interest rates follow a first-order autoregression, given by:

$$r_{f,t+1} = (1-\varphi)\bar{r}_f + \varphi r_{ft} + \sigma_r \epsilon_{r,t+1}, \tag{1}$$

where ϵ_r is a standard normal shock.

Since our focus is on obtaining realistic asset price levels and capital gains for rate-sensitive assets, we calibrate the stationary mean, persistence, and volatility to match moments of the real yield curve. In particular, over our sample period of 1989–2019, we target (1) the slope from a linear regression of the 30-year real

[1] Data on the nominal yield curve can be found here. The zero-coupon yield curve is estimated using off-the-run Treasury coupon securities for horizons up to 30 years.

forward rate (f_{30}) on the current one-year real yield, (2) the average 30-year real forward rate, and (3) the unconditional volatility of the one-year real yield. These three moments provide an exactly identified system that defines each parameter in terms of data moments. Table 1 reports the data moments, their relationship to model parameters, and the implied estimates.

When interest rates follow an AR(1) process, the interest-rate sensitivity of the price of a zero-coupon bond paying off 1 in k years (P_{kt}) is given by:

$$\hat{\varepsilon}_r(P_{kt}) = \frac{1-\varphi^k}{1-\varphi}. \tag{2}$$

We derive this expression in the model below.

Table 1: Parameter values and moments for the riskfree rate process

Moment condition			Estimate	
Data moment	Model equiv.	Data value	Parameter	Value
$\text{cov}(f_{30,t}, r_{ft})/\text{var}(r_{ft})$	φ^{30}	0.2569	φ	0.9557
$\bar{f}_{30,t}$	$\bar{r}_f$	0.0193	$\bar{r}_f$	0.0193
$\text{var}(r_{ft})$	$\sigma_r^2/(1-\varphi^2)$	0.0167	σ_r	0.0049

1.3 Interest-rate sensitivity of assets

We obtain the rate sensitivity of households' wealth in three steps. First, we estimate cashflow duration for each asset and liability on households' balance sheets. Second, we apply equation (2) to this cashflow duration, taking the rate sensitivity of the asset to be the same as the rate sensitivity of a riskfree zero-coupon bond with the same duration. Third, we take a value-weighted sum of the assets' rate sensitivities to arrive at the overall portfolio rate sensitivity of each household.

We adopt different methods to compute the cashflow durations of assets and liabilities for the first step. For equity, real estate, liquid assets, and fixed income, we use cashflow duration measures from outside sources. For vehicles, private businesses and liabilities, we use information provided in the SCF. More information about our construction of all of the different components of interest-rate sensitivity can be found in Appendix A.2.

Equity, real estate, liquid assets, and fixed income For all equity, real estate, and liquid assets, we apply the annual group-wide duration estimates provided by Greenwald et al. (2023). For fixed-income assets, we collect annual average duration estimates from Bloomberg for government debt, municipal bonds, mortgage-backed securities, foreign bonds, and corporate bonds and apply them to each asset within the fixed-income group accordingly.

Vehicles To determine the duration of vehicles, we compute the time left on a vehicle's life using the age of the vehicle provided in the SCF and an estimate of its maximum lifetime. We then assume a constant depreciation rate to determine its cashflow duration.

Private business wealth The time series of private business wealth duration is taken from the annual, aggregate price-dividend ratio of businesses owned by households in the SCF, given by

$$\text{dur}(\text{Private Business}_t) = \frac{\text{Business Value}_t}{\text{Business Income}_t - \text{Wages}_t}. \tag{3}$$

We use the price-dividend ratio here because it is equal to the duration when discount rates and cashflow growth are constant (Binsbergen, 2021). The business value is the value of the household's ownership share and dividends are the household's business income less wages to business owners, which are either reported in the SCF or estimated using information on their level of education and age. We obtain an annual time series by summing each component within each SCF survey year and applying equation (3).

The problem with this method is that valuation ratios are higher at the top of the earnings distribution, since wealthy entrepreneurs tend to run different businesses than other business owners (Schoar, 2010). To address this, we allow for cashflow duration to vary across different wealth categories using the gradient of valuation ratios across the wealth distribution. More detail on the estimation of private business duration is outlined in Appendix A.2.4

Liabilities For each household's liabilities in the SCF, we assume a fixed repayment schedule and estimate duration as

$$\text{dur}(\text{Debt}_t) = \sum_{n=1}^{N} \left(\frac{P_{nt}}{\sum_{n'=1}^{N} P_{n't}} \right) n, \tag{4}$$

where N is the number of years remaining on the loan given in the SCF and $P_{nt} = e^{-ny_{nt}}$ where y_{nt} is the n-year yield. The number of years remaining on the loan is either given explicitly in the SCF or can be inferred from the interest rate and loan balance outstanding.

1.4 Interest-rate sensitivity of wealth

Finally, we compute the rate sensitivity of the overall wealth portfolio for household i as the value-weighted sum of each component elasticity:

$$\hat{\varepsilon}_r(\text{Wealth}_i) = \sum_j \frac{\text{Asset}_{ji}}{|\text{Wealth}_i|} \times \hat{\varepsilon}_r(\text{Asset}_{ji}), \tag{5}$$

where Asset_{ji} denotes the value of the asset or debt j, $\hat{\varepsilon}_r(\cdot)$ its interest-rate elasticity, and Wealth_i the household's net worth.

Table 2: Average duration and portfolio share by asset group

	IR Sensitivity	Portfolio share	
		equal-weighted	wealth-weighted
Assets			
Private Business	16.26	0.04	0.18
Equity	20.38	0.09	0.18
Real Estate	9.78	0.44	0.40
Fixed Income	5.00	0.08	0.12
Vehicles	3.17	0.20	0.03
Cash and Deposits	0.25	0.15	0.08
Liabilities			
Mortgage Debt	8.23	0.52	0.83
Other Debt	2.59	0.48	0.17
Networth			
Value-weighted	12.56		
Equal-weighted	8.36		

Note: This table reports the average interest-rate sensitivity and portfolio share of each asset group for households in the SCF. Interest-rate sensitivity is calculated from asset duration as in equation (2). To calculate average duration for each asset group, we take the duration estimate for each household's holdings in the group and then average them across households, weighting each household's contribution by its SCF sample weight and the value of its holding in the asset group. To obtain the equal-weighted portfolio shares, we take the share of each asset group within a household's portfolio and then average the shares across households using SCF sample weights. We repeat this process for the wealth-weighted portfolio shares, but in this case we weight each household by both its SCF sample weight and its networth.

Table 2 presents the averages of our interest-rate sensitivity estimates by asset group and for wealth for households in the SCF. The average balance sheet interest-rate sensitivity suggests that interest-rate risk generates substantial volatility in returns to wealth. The combination of a rate volatility of 0.49% and an interest-rate sensitivity of 8.21 translates into a volatility of 4% for the return on wealth. Interestingly, holdings of public and private equity only imply a volatility from stock market returns of 2.4%.[2] This implies that variation in interest rates constitutes a risk to households that is of similar magnitude to the risk imposed from variation in stock markets.

2 Stylized facts

We begin by documenting facts about the directly observable side of household's interest-rate exposure: the interest-rate sensitivity of wealth $\varepsilon_r(W)$. Even though it only paints a partial picture of households' exposures to interest rates, the distribution of $\varepsilon_r(W)$ is interesting because it reflects households' portfolio choice problems. We document five stylized facts:

(1) interest-rate sensitivity is hump-shaped over the life cycle;

(2) high earners hold assets with higher interest-rate sensitivity;

(3) interest-rate sensitivity is increasing in wealth;

(4) trends in wealth inequality follow trends in interest rates;

(5) Social Security offsets differences in exposure to interest-rate risk.

These patterns will later guide our structural analysis.

2.1 Interest-rate sensitivity is hump-shaped over the life cycle

The first stylized fact is that the rate sensitivity of wealth is hump-shaped over the life cycle: it is lowest for 20-year-olds, rises to a high for 40- to 45-year-olds, and steadily declines thereafter. Figure 1 decomposes this pattern clearly, showing

[2] Public and private equity shares represent 13% of networth of the average household, so, assuming a volatility of market returns of 18.5%, we obtain 2.4%. Using the wealth-weighted averages of rate sensitivity and equity share yield similar implied volatilities from stock market and interest-rate exposures of roughly 6%. This assumes that other balance sheet assets are not correlated with the stock market, which is true for housing (Flavin and Yamashita, 2002).

the relative contribution of each asset to the total portfolio rate sensitivity. The difference in portfolio interest-rate sensitivities at each age is determined by the assets households choose to hold. For example, 20- to 25-year-old households have relatively low interest-rate sensitivities because the majority of their wealth (70.4%) is invested in liquid accounts (e.g., checking and savings accounts) and vehicles. Their holdings of longer-term assets like stocks and home equity are substantially smaller than later in life.

Figure 1: Interest-rate sensitivity of wealth by age

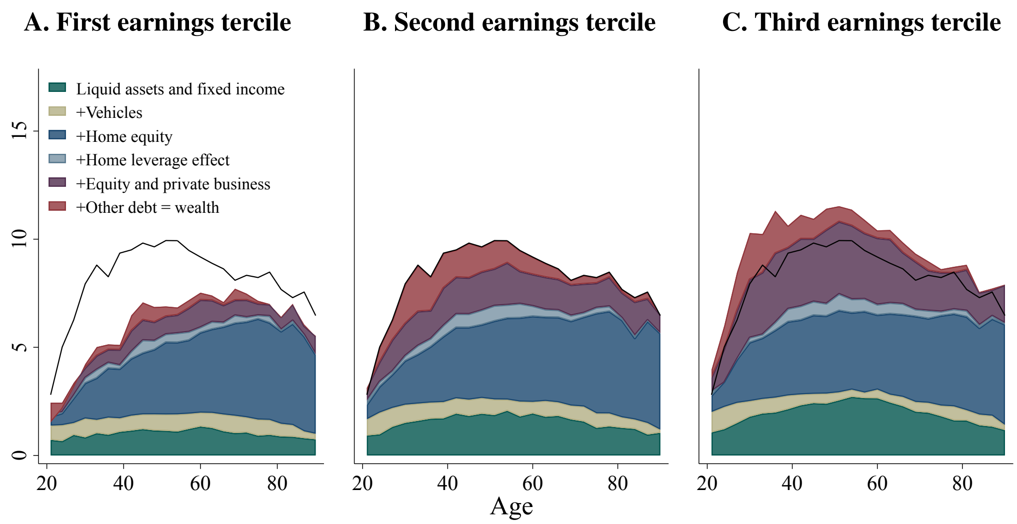

Note: This figure reports the interest-rate sensitivity of wealth by age and tercile of earnings. The rate sensitivity is decomposed into the contribution of six components of wealth. From bottom to top, we calculate the sensitivity of partial portfolios, adding components step-by-step. First, we report the interest-rate sensitivity of liquid assets and fixed-income assets. We then report the rate sensitivity of a larger portfolio that also includes vehicles, and so forth. Thus, the interest-rate sensitivity of the partial portfolio inclusive of the first k components of wealth is $\hat{\varepsilon}_r(\text{Portfolio}_k) = \frac{\text{Portfolio}_{k-1}}{\text{Portfolio}_k}\hat{\varepsilon}_r(\text{Portfolio}_{k-1}) + \frac{\text{Component}_k}{\text{Portfolio}_k}\hat{\varepsilon}_r(\text{Component}_k)$.

As households approach midlife, the composition of assets changes and the interest-rate sensitivities of their portfolios grow. Shorter-term liquid assets and vehicles contribute roughly the same to interest-rate sensitivity as they do for the young, but now, the majority of the portfolio (48.7% for 40-year-olds) is made up of longer-term assets like equity and real estate. Moreover, leverage—in particular, mortgages and other debts—plays a more important role, increasing the rate sensitivity of the wealth portfolio by nearly 20%. The reason leverage increases the rate exposure of the household's portfolio is because the (equal-weighted) average rate sensitivity of assets is approximately 40% higher than that of debts over our sample. The net position, therefore, has a higher interest-rate sensitivity

As midlife turns to retirement, the rate sensitivity of household portfolios begins to fall. The decline in rate exposure is driven not by the asset side of the portfolio, but rather by the disappearance of leverage, which reduces the interest-rate sensitivity of the wealth portfolio. This is consistent with the conventional narrative in saving for retirement: households with a large stock of human capital take on mortgages in early adulthood to guarantee housing consumption flows in old age.

2.2 Interest-rate sensitivity is increasing in earnings

The second stylized fact is that high-earning households hold more rate-sensitive portfolios, on average, as seen by comparing the three panels of Figure 1. The three panels show that, for a 1% decline in interest rates, the top earnings tercile will see approximately 4 percentage points larger capital gains than those of the bottom earnings tercile. High earners investing more in equity drives the relationship between earnings and the rate-sensitivity of wealth.

2.3 Interest-rate sensitivity is increasing in wealth

The third stylized fact is that interest-rate sensitivity is generally increasing in wealth. This fact is shown in Figure 2, which decomposes the average rate sensitivity for households between ages 40 and 45 over the log of their wealth scaled by the Social Security Wage Index in their survey year.

For low-wealth households, liquid accounts, vehicles, and non-mortgage debt contribute the most to the interest-rate sensitivity of their portfolios. For middle-wealth households, real estate becomes the dominant asset, with its rate sensitivity amplified by the mortgage taken on to finance the purchase. The large indivisible nature of houses leads lower-middle-wealth homebuyers to take out large mortgages and expose themselves to interest-rate fluctuations, which is reflected in the small bump in rate sensitivity near the lower-middle portion of the wealth distribution. As wealth increases, portfolio rate exposures increase with larger positions in highly rate-sensitive assets like publicly traded equity and private businesses.

Figure 2: Interest-rate sensitivity of wealth at ages 40–45 by level of wealth

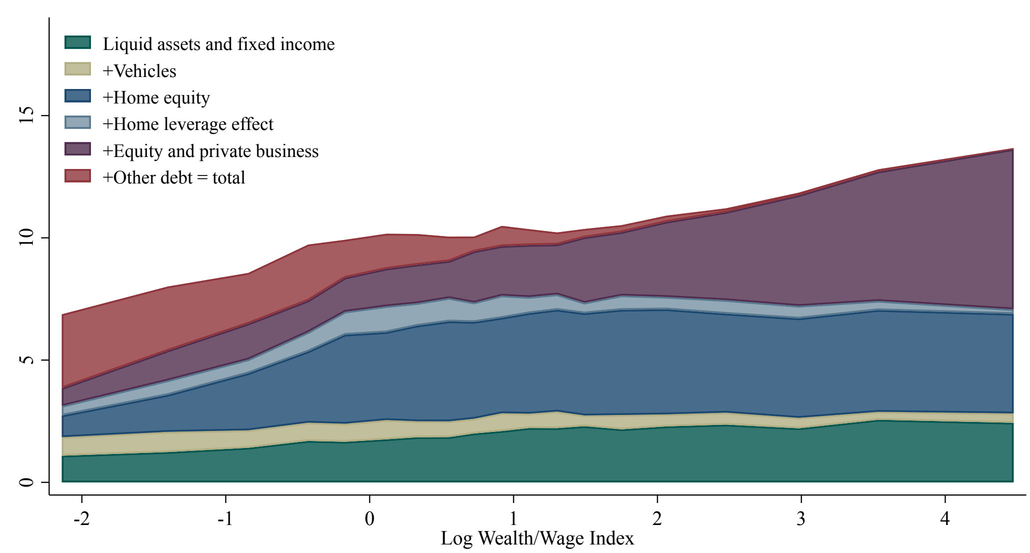

Note: This figure decomposes the interest-rate sensitivity for households in which the head of the household is between 40 and 45. The methodology is the same as in Figure 1, except that here the x-axis is the log of wealth scaled by the Social Security Wage Index in the survey year.

2.4 Wealth inequality follows interest rates

Figure 3 illustrates the fourth stylized fact: Over the last six decades, the wealth share of the top 10% of the wealth distribution has closely tracked the price of real bonds, which we approximate by one minus of the estimated 10-year forward rate.

Figure 3: Wealth inequality and estimated 10-year real forward rates

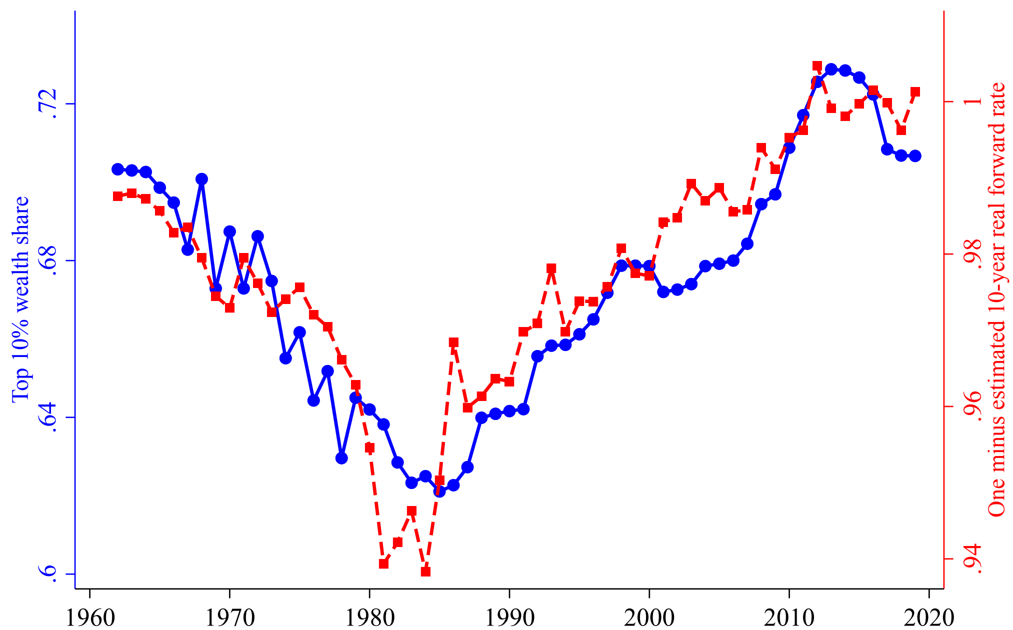

Note: This figure presents the time series of the top 10% wealth share from the World Inequality Database and $1 - \hat{f}_{10,t}$, one minus our estimated 10-year real forward rate from equation (A.2).

2.5 Social Security offsets differences in rate sensitivity

We now extend our definition of wealth to include the net present value of Social Security payments—that is, of expected benefits minus expected payroll taxes to be paid into the system. Figure 4 displays the fifth stylized fact: the inclusion of Social Security wealth strongly attenuates the relationships between interest-rate exposure and wealth (Panel A) and earnings (Panel B).

Figure 4: Interest-rate sensitivity of wealth at ages 40–45: Role of Social Security

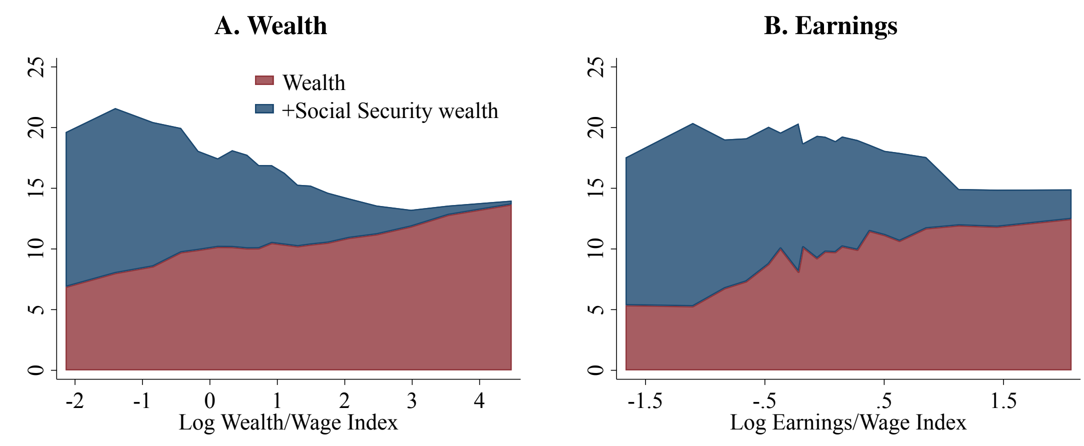

Note: This figure decomposes the interest-rate sensitivity for households in which the head of the household is between 40 and 45. The methodology is the same as in Figure 1, except that here the x-axis is the log of wealth scaled by the Social Security Wage Index in the survey year in Panel A and scaled earnings in Panel B. Estimates for the net present value of Social Security at the individual level come from the risk-adjusted valuation of Catherine et al. (2023).

3 Model

We model household consumption and investment decisions over a life cycle divided into two stages: working age and retirement.

3.1 Agents

Agent i chooses consumption C_i and portfolio allocation π_i to maximize lifetime utility

$$V_{it} = \max_{\{C_{is},\pi_{is}\}} \mathbb{E}_t \sum_{s=t}^{t_{\max}} \beta^{s-t} p_{it,s-1} \left[(1 - m_{i,s-1}) \frac{C_{is}^{1-\gamma}}{1-\gamma} + m_{i,s-1} b(W_{is}, r_{fs}) \right], \quad (6)$$

where β is the rate of time preference, γ is the coefficient of relative risk aversion, $t_{\max}$ is the maximum lifespan, m_{it} is the age- and income-dependent mortality probability from t to $t+1$, $p_{its} = \prod_{u=t}^{s-1}(1-m_{iu})$ is the probability of surviving from t to s, and b is the bequest motive over terminal wealth W and the interest rate r_f.

While working-aged, the agent receives labor income L_i and pays Social Security taxes T_i; in retirement, which begins at a given time t_{ret}, he or she receives benefits B_i. The utility maximization is therefore subject to the budget constraint for wealth

$$W_{i,t+1} = (W_{it} + L_{it} + B_{it} - T_{it} - C_{it})R_{Wi,t+1}, \tag{7}$$

with gross return on savings

$$R_{Wi,t+1} = R_{ft} + \pi_{it}(R_{n,t+1} - R_{ft}). \tag{8}$$

In this expression, R_f is the return on a riskfree bond, R_n is the return on the long-term asset, and π_i is the share of wealth invested in this asset.

3.2 Interest rates and wealth returns

Rates of return on assets vary over time; we thus model stochastic processes for the short- and long-term asset returns and constrain their joint dynamics using equilibrium pricing conditions. Denote log returns by lowercase $r = \log R$. As in our empirics, we assume that the riskfree rate follows a first-order autoregression:

$$r_{f,t+1} = (1-\varphi)\bar{r}_f + \varphi r_{ft} + \sigma_r \epsilon_{r,t+1}.$$

We model the long-term asset as a riskless claim to one unit of real consumption in n periods. Its price, denoted P_n, satisfies the expectations hypothesis, generalized to include constant term premia.

We assume that the term premium on each n-period bond is some constant μ_n (with $\mu_1 = 0$). As we show in Appendix B.2, these assumptions imply an explicit relation between the dynamics of long-term bond returns and short-term rate fluctuations: the log bond return equals

$$r_{n,t+1} = r_{ft} + \mu_n - \sigma_n \epsilon_{r,t+1}, \tag{9}$$

where sensitivity to rate shocks σ_n is given by

$$\sigma_n = \frac{1-\varphi^{n-1}}{1-\varphi}\sigma_r. \tag{10}$$

In addition, we set $\mu_n = -\sigma_n^2/2$, so that there is no risk premium. This is because we are interested in the effect of rate fluctuations, not the additional risk compensation for holding long-term government debt.

We define the interest-rate sensitivity of an asset with current price P_t as the percentage change in the price caused by an unexpected decline in the interest rate:[3]

$$\varepsilon_r(P_t) \equiv -\frac{\log P_{t+1} - \mathbb{E}_t \log P_{t+1}}{r_{f,t+1} - \mathbb{E}_t r_{f,t+1}}. \tag{11}$$

It follows that the rate sensitivity of the long-term bond is

$$\varepsilon_r(P_{nt}) = \frac{1-\varphi^{n-1}}{1-\varphi}, \tag{12}$$

This sensitivity is increasing in maturity n. This expression summarizes the effect of unexpected changes in interest rates: if the riskfree rate unexpectedly falls, then

[3] A related but conceptually distinct concept is the instantaneous interest-rate elasticity $-\partial \log P_t/\partial r_{ft}$. This elasticity tells us how the price depends on the level of the rate instead of how it will respond to an unexpected future change. In certain cases, as the length of the time interval between periods becomes infinitesimally small, these objects converge to the same concept.

the long-term bond has an unexpectedly high return from capital gains. The longer is the maturity n, the larger is this response. These return processes, together with the agent's portfolio allocation π_i, give us the return on wealth from (8).

We purposely choose not to model the full menu of assets—housing, bonds, stocks—and liabilities that actually constitute most of a household's balance sheets for two reasons. First, this choice makes the economic intuition of the model as transparent as possible. Second, by the portfolio separation theorem, two bonds of different maturities are sufficient to target the optimal interest-rate sensitivity of wealth. Adding other assets with their own risk factors will have no effect on this dimension of portfolio choice.

3.3 Labor income

We model labor-income dynamics using the empirically realistic process estimated by Guvenen, Karahan, Ozkan and Song (2022). Each agent's income L_i is the product of the aggregate wage index $\bar{L}_t$ and an idiosyncratic component

$$\tilde{L}_{it} = \exp\left\{g(t) + \zeta_{i0} + z_{it} + \epsilon_{it}\right\}. \tag{13}$$

The deterministic component $g(t)$ is a quadratic polynomial of age; it captures common life-cycle patterns in income. The parameter ζ_{i0} governs heterogeneous levels of earnings. The persistent component of earnings, denoted by z_i, follows a first-order autoregression

$$z_{it} = \rho z_{i,t-1} + \eta_{it}, \tag{14}$$

with innovations η_i drawn from a mixture of normal distributions

$$\eta_{it} \sim \begin{cases} \mathcal{N}(\mu_{\eta 1}, \sigma^2_{\eta 1}) & \text{with probability } p_z, \\ \mathcal{N}(\mu_{\eta 2}, \sigma^2_{\eta 2}) & \text{with probability } 1 - p_z. \end{cases} \tag{15}$$

The initial cross-sectional distribution of the persistent component of earnings is given by $z_{i0} \sim \mathcal{N}(0, \sigma^2_{z0})$. The transitory component of idiosyncratic earnings ϵ_{it} is also drawn from a mixture of normal distributions

$$\epsilon_{it} \sim \begin{cases} \mathcal{N}(\mu_{\epsilon 1}, \sigma^2_{\epsilon 1}) & \text{with probability } p_\epsilon, \\ \mathcal{N}(\mu_{\epsilon 2}, \sigma^2_{\epsilon 2}) & \text{with probability } 1 - p_\epsilon. \end{cases} \tag{16}$$

These mixture processes serve to match higher-order moments of income growth.

3.4 Social Security

Agents pay Social Security payroll taxes T_i on their labor income during working life, then receive benefits B_i in retirement. We assume all workers retire at the full-retirement age t_{ret}, which is the age at which they receive 100% of their scheduled benefits. The tax payments are 10.6% of all income below the Social Security wage base, which is 2.5 times the average wage:

$$T_{it} = 0.106 \min\{L_{it}, 2.5\bar{L}_t\}. \tag{17}$$

Social Security retirement benefits depend on the agent's average indexed yearly earnings (AIYE), which is an average of the highest 35 years of indexed earnings

$$L_{it}^{\text{indexed}} = \min\{L_{it}, 2.5\bar{L}_t\} \frac{\bar{L}_{t_{60}}}{\bar{L}_t} \tag{18}$$

up to retirement, where $\bar{L}_{t_{60}}$ is the wage index during the period in which the worker is 60. In words, indexed earnings are the income below the wage base at a given age, adjusted for growth in the aggregate wage index $\bar{L}_t$ up to age 60. Income earned after age 60 but before retirement at t_{ret} can still contribute to the worker's AIYE, but it is indexed to t_{60}. Total benefits are then a piecewise-linear function of the AIYE when the worker retires:

$$B_{it} = \begin{cases} 0.9\text{AIYE}_{it_{\text{ret}}} & \text{if AIYE}_{it_{\text{ret}}} < b_1, \\ 0.9b_1 + 0.32(\text{AIYE}_{it_{\text{ret}}} - b_1) & \text{if } b_1 \leq \text{AIYE}_{it_{\text{ret}}} < b_2, \\ 0.9b_1 + 0.32(b_2 - b_1) + 0.15(\text{AIYE}_{it_{\text{ret}}} - b_2) & \text{if } b_2 \leq \text{AIYE}_{it_{\text{ret}}}. \end{cases} \tag{19}$$

The kinks in this benefit formula are determined by the "bend points" b_1 and b_2, which historically are about 21% and 125% of the wage index, respectively. The formula is progressive: as AIYE (lifetime income) increases, the marginal benefit declines. Note that AIYE is itself bounded above due to the wage base, so benefits have an upper bound. Benefits after the retirement year are held constant in real terms — that is, they are adjusted in nominal terms to account for CPI inflation. Before retirement, we keep track of average index earnings as:

$$\text{AIYE}_{it} = \sum_{s=t_0}^{t} \min\{L_{is}, 2.5\bar{L}_s\} \frac{\bar{L}_t}{\bar{L}_s} = \bar{L}_t \sum_{s=t_0}^{t} \min\{\tilde{L}_{is}, 2.5\}. \tag{20}$$

3.5 Income taxes

Households pay taxes on income and benefits according to the income tax brackets faced by U.S. households in 2020, adjusted for changes in the aggregate wage index. Marginal tax rates are progressively increasing in idiosyncratic income $\tilde{L}_i$; we report the formula for these rates in Appendix B.1.

3.6 Bequests

Individuals bequeath to their children an inheritance from their terminal financial wealth. In modeling utility over bequests, one must consider the fact that inheritance does not necessarily constitute a one-time transfer of liquid wealth; it might instead be a long-lived flow of consumption, such as from real estate. Hence, we model the bequest motive as a function of an annuity flow $\bar{C}_i$ which takes into account both the value of financial wealth and the time value of money. Specifically, we assume

$$b(W_{it}, r_{ft}) = \bar{b}\frac{\bar{C}_{it}^{1-\gamma}}{1-\gamma}, \tag{21}$$

where $\bar{b}$ can be interpreted as the number of years of consumption that the agent wants to bequeath, and $\bar{C}_i$ is the coupon implicit in the annuity of $\bar{b}$ years:

$$W_{it} = \bar{C}_{it}\sum_{k=1}^{\bar{b}} P_{kt}. \tag{22}$$

4 Economic intuition

To communicate the first-order intuitions of our model, we present an analytical solution to a linearized version with no idiosyncratic income risk or bequests.[4] We proceed in two steps. First, we discuss consumption and portfolio rules for an agent without human capital or Social Security. Second, we show that the same portfolio rule applies to *total* wealth in the presence of background assets such as human capital and Social Security.

[4]See Appendix C for derivations and further discussion.

4.1 Optimal choices without labor income

Consumption rule Without labor income, the linearized model implies the optimal consumption policy

$$\frac{C_{it}^*}{W_{it}} = \underbrace{(1-\beta(1-m_{it}))}_{\text{time discounting}} \times \underbrace{\exp\left\{\left(1-\frac{1}{\gamma}\right)(\varrho_{0t}+\varrho_{rt}r_{ft})\right\}}_{\text{income and substitution effects}}. \tag{23}$$

The first term represents the positive effect of impatience and mortality on consumption. The second term represents the net of income and substitution effects from interest rates. Higher rates mean higher interest income, so that households can consume more today (the income effect). At the same time, higher rates mean agents get more consumption tomorrow in exchange for their savings (the substitution effect). The income effect dominates the substitution effect when the elasticity of intertemporal substitution (the EIS, $1/\gamma$) is less than one ($\gamma > 1$). The sensitivity of consumption to interest rates depends on the coefficient

$$\varrho_{rt} = \sum_{j=1}^{\infty} \varphi^{j-1}\beta^j p_{it,t+j}, \tag{24}$$

where $p_{it,t+j}$ is the survival probability from year t to $t+j$. It is declining in age because agents with shorter horizons are less affected by persistent rate changes.

Portfolio rule The optimal allocation to the n-period bond is:[5]

$$\pi_{it}^* = \underbrace{\frac{1}{\gamma}\frac{\mu_n+\frac{1}{2}\sigma_n^2}{\sigma_n^2}}_{\text{myopic demand}} + \underbrace{\left(1-\frac{1}{\gamma}\right)\varrho_{rt}\left(\frac{1-\varphi^{n-1}}{1-\varphi}\right)^{-1}}_{\text{interest-rate risk hedging demand}}. \tag{25}$$

[5]As we verify in Appendix C.2, this solution holds true even if we separate the coefficient of relative risk aversion from the elasticity of intertemporal substitution (EIS). Thus, the portfolio share is indeed governed by risk aversion, and not the EIS.

The first term represents the traditional risk-return tradeoff of Merton (1969). Our assumption of $\mu_n = -\sigma_n^2/2$ sets this term to zero. The second term is the demand from intertemporal hedging of interest-rate fluctuations, the focus of our paper. Because its value increases when rates unexpectedly decline, the long-term bond offers protection against the deterioration of investment opportunities.

This portfolio rule implies that the optimal interest-rate sensitivity of wealth, which equals $\varepsilon_r^*(W_{it}) = \pi_{it}^* \sigma_n / \sigma_r$ from (11), is

$$\varepsilon_r^*(W_{it}) = \frac{1}{\gamma} \frac{\mu_n + \frac{1}{2}\sigma_n^2}{\sigma_n \sigma_r} + \left(1 - \frac{1}{\gamma}\right) \varrho_{rt}. \tag{26}$$

Note that this optimal sensitivity is independent of the long-term bond maturity n. This illustrates the fact that households can mix any combination of short- and long-term assets to target their optimal total exposure to rate risk.

The sensitivity of consumption to rate shocks ϱ_{rt} declines with the investor's horizon, so the hedging demand decreases in age toward zero. From (24), we see that ϱ_{rt} represents the cumulative effect of an interest-rate shock over time, weighted by the importance of each future period in the agent's lifetime utility, taking into account mortality and the psychological preference for the present. In other words, it results from the interplay between the agent's horizon and the persistence of interest-rate shocks. Agents with shorter horizons are less sensitive to future rate changes; therefore, absent labor income and Social Security, the exposure of a household's portfolio to rate shocks should decline over the life cycle.

Interestingly, for a risk-neutral agent ($\gamma = 0$), the hedging demand is infinitely negative. A risk-neutral agent prefers to receive capital gains when they can be reinvested at a higher rate of return and would therefore like to short-sell the long-term asset. In the log-utility case ($\gamma = 1$), this force is perfectly offset by the will to insure against the deterioration of investment opportunities and the portfolio rule is reduced to the myopic demand.

4.2 Effects of labor income and Social Security

Now let us consider the effect of labor income and Social Security. Suppose that labor income L_i, taxes T_i, and benefits B_i are deterministic. The values of human capital H_{it} and Social Security wealth S_{it} are

$$H_{it} = \sum_{k=1}^{t_{\text{ret}}-t} p_{t,t+k} P_{kt} L_{i,t+k}, \tag{27}$$

and

$$S_{it} = \sum_{k=1}^{t_{\max}} p_{t,t+k} P_{kt} (B_{i,t+k} - T_{i,t+k}), \tag{28}$$

where P_{kt} is the price of a k-maturity zero-coupon bond. Define total wealth $\overline{W}_i$ as the sum of wealth W_i and these present values.

Implementing the same linearization implies the consumption rule relative to total wealth is the same as in the no-income solution: $C_i/\overline{W}_i$ equals the right-hand side of (23). Similarly, the optimal allocation to bonds out of total wealth is $\bar{\pi}_i = \pi_i^*$ from (25). The optimal allocation out of wealth W then takes the form

$$\pi_{it} = \pi_{it}^* - \underbrace{(\pi_{it}^H - \pi_{it}^*)\frac{H_{it}}{W_{it}}}_{\substack{\text{human capital} \\ \text{substitution effect}}} - \underbrace{(\pi_{it}^S - \pi_{it}^*)\frac{S_{it}}{W_{it}}}_{\substack{\text{Social Security} \\ \text{substitution effect}}}, \tag{29}$$

or, in terms of interest-rate sensitivities,

$$\varepsilon_r(W_{it}) = \varepsilon_{r,it}^* - (\varepsilon_r(H_{it}) - \varepsilon_{r,it}^*)\frac{H_{it}}{W_{it}} - (\varepsilon_r(S_{it}) - \varepsilon_{r,it}^*)\frac{S_{it}}{W_{it}}, \tag{30}$$

such that the interest-rate sensitivity of total wealth $\varepsilon_r(\overline{W}_{it})$ remains equal to $\varepsilon_{r,it}^*$.

The endowments of human capital and Social Security wealth are implicit holdings of long-term assets, and thus substitute for the traded n-period bond. The values π_i^H and π_i^S represent the implicit percentage of each asset invested in the n-period bond. The agent adjusts the allocation to wealth π_i such that the duration of total wealth matches π_i^*. If, for instance, agents are endowed with a large stock of high-duration Social Security (i.e., π_i^S and S_i are large), they adjust their allocations to long-term bonds π_i downward to offset this high rate exposure.

Figure 5: Effect of labor income and Social Security on long-term asset share

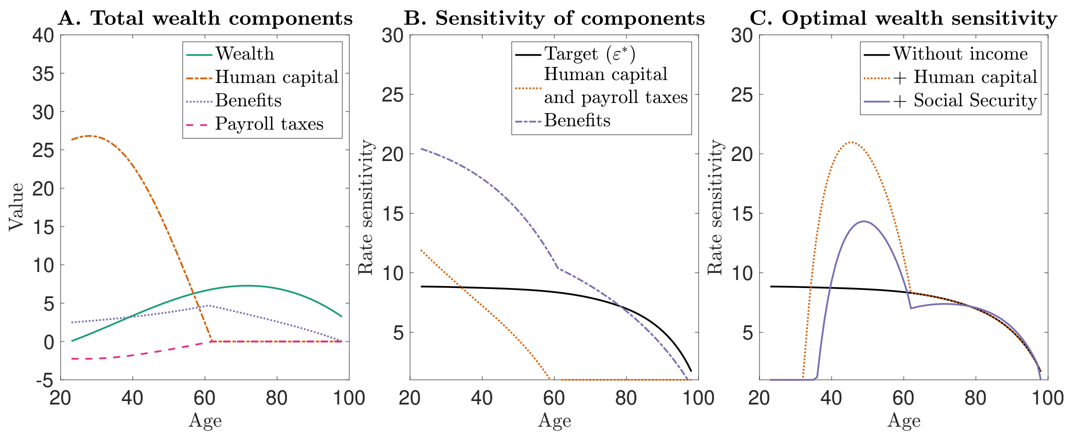

Note: This figure shows a representative path of total wealth components, their interest-rate sensitivities, and their effect on wealth allocations over the life cycle. Panel A plots the average values of each component of total wealth, defined as the sum of wealth and the present values of labor income (human capital) and Social Security taxes and benefits. Panel B shows the interest-rate sensitivity of each component. Panel C illustrates the incremental effect of each component on the optimal interest-rate sensitivity. We assume $\gamma = 5$ and $\beta = 0.95$. The life-cycle profile of wealth is approximated from the data, the present values of human capital and Social Security wealth are simulated.

Figure 5 illustrates the life-cycle pattern generated by this model. Early in life, most agents have little financial wealth and a large endowment of high-duration human capital. To match their ideal total-wealth rate exposure, they mostly hold assets with low rate sensitivity. As households get closer to retirement, they increase holdings of the long-term asset to offset short-term labor income and taxes, net of long-term benefits. As they progress through retirement, households reduce long-term bond holdings, in line with the declining target allocation implied by the policies above. In sum, substitution and aging effects explain the hump-shaped pattern in the data.

Figure 6: Substitution effects of Social Security within an age group

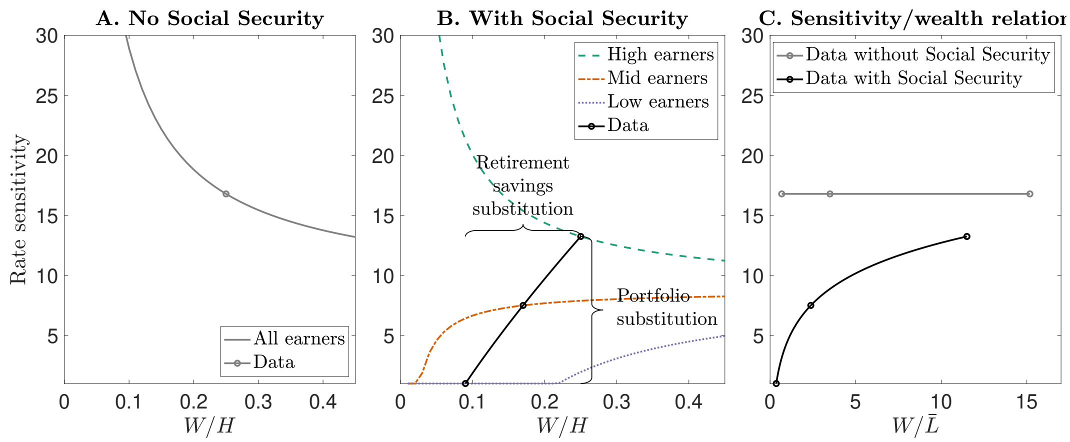

Note: This figure illustrates the effect of Social Security on intra-cohort allocations to the long-term asset. Panel A plots the optimal long-term bond share as a function of the ratio of wealth W to human capital H when there is no Social Security. The round markers represent hypothetical wealth-to-human capital ratios W/H. Panel B shows the same relation but in the presence of Social Security. In Panel C, we re-plot the points in Panels A and B in terms of wealth only. Policy functions are drawn for age 42, $\gamma = 5$, and $\beta = 0.95$.

In addition to these effects, the progressivity of Social Security implies that households with lower earnings will hold less rate-sensitive portfolios. Figure 6 illustrates the economic intuition behind this prediction. Panel A describes the case without Social Security: all households have the same savings rate, hence the wealth-income ratios and portfolio allocations show little variation within an age group. Panel B shows the two effects of Social Security. Social Security offers a higher replacement rate to low-earners, which means that they need to save less for retirement, leading to a lower wealth-to-human capital ratio at a given age: the retirement savings substitution effect. Second, because Social Security represents a greater share of their total endowment, it reduces the demand for long-term asset of low-earners disproportionately: the portfolio substitution effect. As Panel C shows, these effects combine to generate a steep positive relation between wealth and rate exposure, as in the data.

4.3 Real-life interpretation

Fully hedged consumption plan In principle, agents can eliminate interest-rate risk by buying, at current prices, a portfolio of zero-coupon bonds that perfectly matches the difference between their desired consumption plan and their expected earnings in all periods, then holding this portfolio to maturity.

We can illustrate this intuition most clearly in the limiting case of an infinitely risk-averse investor.[6] In this case, the investor's desire to smooth consumption over time yields a constant, deterministic policy $C_{it} = \bar{C}_i$. Let Y_i denote the agent's deterministic stream of income. Wealth is the present value of the excess consumption plan:

$$W_{it} = \sum_{k=1}^{t_{\max}} P_{kt}(\bar{C}_i - Y_{i,t+k}). \tag{31}$$

The agent can secure the optimal consumption plan by buying $\bar{C}_i - Y_{i,t+k}$ of each k-period zero-coupon bond and consuming the coupons and income at maturity.

The strategy is unaffected by capital gains and losses from interest-rate changes. As we prove in Appendix C.5, the optimal allocation π_i replicates exactly this buy-and-hold strategy in the limit as $\gamma \to \infty$.

Real-world implementation In reality, households do not invest their wealth in portfolios of zero-coupon bonds. We illustrate how they implement this strategy,

[6] See Appendix C.5 for a derivation and more detailed technical discussion of this case.

with the set of assets and contracts at their disposal, with three hypothetical households.

First, consider a worker at the bottom of the earnings distribution. Because of its high replacement rate, Social Security taxes and benefits execute all inter-temporal transfers of income required to smooth consumption over the life cycle, and does so independently of the rate of return on private savings. The worker need only hold short-term assets.

Second, consider a middle-class worker. Because replacement rates fall with lifetime earnings, the worker needs to save privately as well. However, he can execute large inter-temporal transfers at current prices by buying a house with a fixed-rate mortgage. By doing so, he effectively trades a flow of coupon payments later in life, when $C > Y$, in the form of rent-free housing, in exchange for a stream of mortgage payments earlier in life, when $C < Y$. This operation is priced using the spot yield curve. This strategy eliminates interest-rate risk for households whose Social Security benefits cover their non-residential consumption in retirement.

Finally, consider a high-income worker. Because Social Security benefits are relatively small in the upper half of the earnings distribution, she complements this strategy with additional investments. In the United States, this complement typically takes the form of a retirement account. If these savings were invested in short-term assets, she would need to increase her contribution rate to maintain the same consumption level in retirement. However, if her account were mostly invested in long-term assets, capital gains would offset potential declines in future rates of return. As we see in Figure 1, high earners follow the long-term asset strategy by investing their "extra" wealth in stocks. From this point of view, the glide path strategy of pension funds also makes sense, as it invests retirement contributions in stocks early in the life cycle and moves towards safer assets when workers get older.

4.4 Distributional effects of interest-rate fluctuations

Because of substitution effects, the interest-rate exposure of wealth is heterogeneous and correlated with key household characteristics, both within and between cohorts. At the same time, the model suggests substantially less heterogeneity in *total* rate exposure. In fact, all agents within a given age group are identically exposed to interest-rate risk, regardless of differences in wealth, income, or benefits. The only source of heterogeneous total rate exposure is different investment horizons. To see this, let us consider three measures of total rate exposure: the interest-rate sensitivity of total wealth, of consumption, and of lifetime utility.

As we showed above, the rate elasticity of total wealth is $\varepsilon_r(\overline{W}_{it}) = \varepsilon^*_{it}$, which depends only on risk aversion and the agent's investment horizon ϱ_{rt}. In other words, for any two agents with the same risk aversion and horizon, a rate shock will have an identical effect on total wealth, even if those agents have different long-term asset shares. Likewise, the optimal consumption rule (23) implies that the response of consumption to a rate shock $\varepsilon_r(C^*_{it})$ is also homogeneous within a cohort.

The third—and arguably most relevant—measure of total rate exposure is the interest-rate sensitivity of lifetime utility. Specifically, we calculate the rate sensitivity of a transformation of expected utility:

$$U_{it} = ((1-\gamma)V_{it})^{1/(1-\gamma)}, \tag{32}$$

where V_{it} is the expected utility maximand (6). This transformation backs out a total-wealth certainty equivalent—it is the value of total wealth implied by the value function V taking a power form.[7] In our linearized model, lifetime utility has the

[7]The other, more mathematical reason for the transformation is that V is negative, so it does not have a well-defined rate sensitivity.

closed-form solution

$$U_{it} = \overline{W}_{it}(1 - \beta(1 - m_{it})) \exp\left\{\varrho_{0t} + \varrho_{rt} r_{ft}\right\}. \tag{33}$$

This expression illustrates the complementarity between total wealth and interest rates: agents have high expected utility when they enjoy high wealth with high rates. The rate sensitivity of lifetime utility is therefore[8]

$$\varepsilon_r(U_{it}) = \left(1 - \frac{1}{\gamma}\right) \varrho_{rt} - \varrho_{r,t+1} \approx -\frac{1}{\gamma} \varrho_{rt}. \tag{34}$$

For any finite value of risk aversion, this elasticity is negative, meaning that a decline in rates is bad news for lifetime utility, while an increase is good news. In welfare terms, this means that, when rates fall, the deterioration of future investment opportunities outweighs the capital gains on long-term cash-flow claims. Agents are willing to accept this unhedged interest-rate exposure because of the option value of compounding at a higher rate, as explained above.

Most notably, the amount of wealth, income, and benefits a household possesses is irrelevant to the rate sensitivity of utility, because the household always trades in long-term asset markets to rebalance back to this optimal exposure. Unexpected rate changes may redistribute wealth, but they do not redistribute welfare.

[8] This expression remains unchanged with Epstein-Zin preferences. The approximate equality is exact as the length of the time interval shrinks to zero.

5 Matching the stylized facts

5.1 Calibration

Preferences We calibrate households' preferences to match the evolution of wealth over the life cycle and the average interest-rate sensitivity of wealth observed in the SCF. We find that a discount factor of $\beta = 0.95$ and a bequest motive equivalent to $\bar{b} = 10$ years of consumption match the growth of wealth until the retirement age and its evolution thereafter. Moreover, a coefficient of relative risk aversion of $\gamma = 5$ matches the average rate sensitivity of wealth.

Our calibration of γ is consistent with studies matching the life-cycle profile of the share of wealth invested in stocks, which typically use values between 5 and 6 (Benzoni, Collin-Dufresne and Goldstein, 2007; Catherine, 2022; Huggett and Kaplan, 2016; Lynch and Tan, 2011; Meeuwis, 2022). Based on portfolios observed in Swedish administrative data, Calvet, Campbell, Gomes and Sodini (2021) estimate an average γ of 5.2.

Income process We calibrate the stochastic parameters of the labor process using estimates from Guvenen et al. (2022), which we report in Appendix E.1.

Initial wealth Households enter working life with $0.1\times$ the national wage index in networth, the equivalent of \$5,400 in 2019.

Mortality We model mortality as a function of age and past lifetime earnings:

$$m_{it} = \min\left\{\chi\left(\frac{\text{AIYE}_{it}}{\bar{L}_t}\right) \times m(\text{age}_{it}), 1\right\}, \tag{35}$$

where χ is an adjustment coefficient which only depends on the average indexed earnings of the agent up to time t and $m(\text{age}_{it})$ is the average mortality rate at that age, which we calibrate as the average across genders from the 2017 Social Security actuarial life tables. While $\chi\left(\text{AIYE}_{it}/\bar{L}_t\right)$ does not depend on age, the agent sees his life expectancy change as he moves up or down the wage ladder. An advantage of our method is that the agent's life expectancy is less volatile than if it were a function of persistent income z_{it}. We calibrate the value of $\chi\left(\text{AIYE}_{it}/\bar{L}_t\right)$ at each point of the numerical grid of the $\text{AIYE}_{it}/\bar{L}_t$ state variable such that, given our labor-income process, we obtain the same life expectancy differential across percentiles of $\chi\left(\text{AIYE}_{it}/\bar{L}_t\right)$ at age 40 as those reported by percentiles of earnings in Chetty et al. (2016).

5.2 Cross-section of interest-rate sensitivity

Figure 7 reports the evolution of wealth and its sensitivity to interest rates over the life cycle, in the data and in the model. The left panel shows that the model matches the evolution of wealth very well. The right panel shows that, like in the data, the interest-rate sensitivity of wealth increases over the first twenty years and declines afterwards. The increase is explained by the substitution effect of human capital and Social Security early in life. Both of these assets have higher rate sensitivity than the agent's target, thus reducing the optimal long-term asset share. Over the life cycle, the duration of human capital declines and drops below the agent's target, reversing the sign of the hedging demand and increasing the long-term asset share. As the weight of human capital declines with age, the magnitude of the hedging begins to fall at retirement.

The introduction of labor income risk in the model reduces the agent's valuation of his human capital, which explains why the predicted hump-shaped relationship with age is less pronounced than in Figure 5.

During retirement, the decline in the agent's investment horizon becomes the dominant force and reduces the need to hedge against falling interest rates. As a result, the long-term asset share falls. This decline is moderated by the bequest motive, which effectively increases the investment horizon of the agent beyond his or her own life expectancy.

Figure 7: Life-cycle profiles of wealth and its interest-rate sensitivity

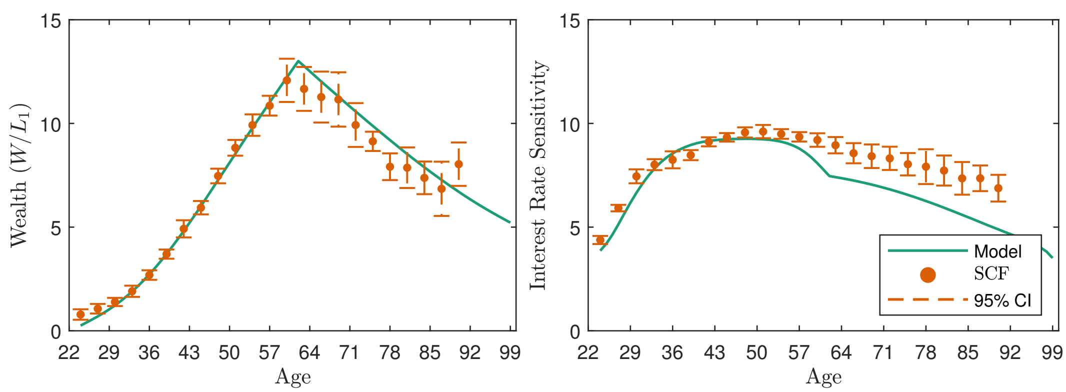

Note: This figure reports the evolution of market wealth and its sensitivity to interest rates over the life cycle in our benchmark calibration and in the SCF. In the data, wealth is computed per adult, including deceased spouses, and scaled by the Social Security wage index. 95% confidence intervals represent $\pm$ 1.96 standard errors, clustered by cohort.

Figure 8: Interest-rate sensitivity of wealth at age 40–45

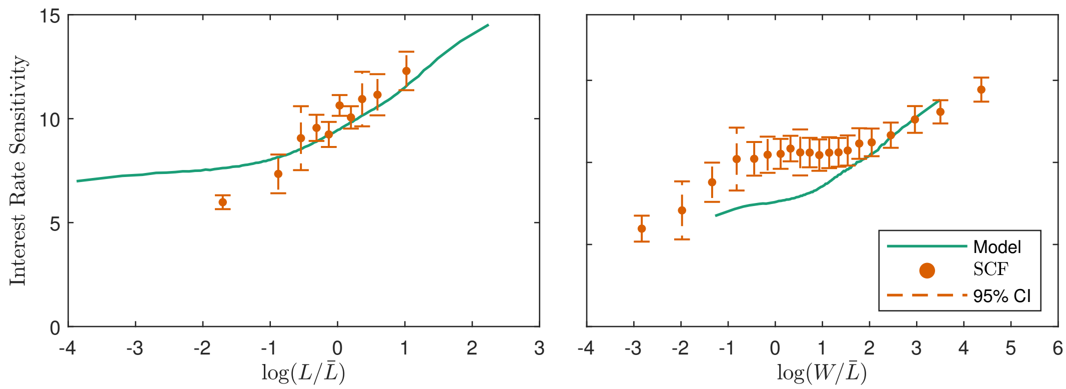

Note: This figure reports the relationship between the interest-rate sensitivity of wealth and wealth (left panel) and earnings (right panel). In the data, wealth and earnings are computed per adult and scaled by the Social Security wage index. In the left panel, each bin represents a decile of earnings. In the right panel, each bin represents 5% of observations, except for the four wealthiest bins which represent 2.5% each. Simulated data report the average interest-rate sensitivity per centile of wealth and earnings, respectively. 95% confidence intervals represent $\pm$ 1.96 standard errors, clustered by cohort.

The left panel of Figure 8 reports the relationship between the interest-rate sensitivity of wealth and income between age 40 and 45. In the model, high earners invest more in the long-term asset because Social Security covers a smaller share of their retirement consumption and, to a lesser extent, because they have higher life expectancy. Appendix D.2 decomposes these two effects quantitatively.

The right panel shows that the model also produces a positive relationship between the long-term asset share and wealth within an age group. This is partly explained by the fact that wealthier households tend to be high earners and that human capital and Social Security represent smaller fractions of their total wealth, and thus have weaker substitution effects. In the data, we observe a hump around the third decile of the wealth distribution, which could reflect the need for these households to borrow to buy houses and vehicles.

It is notoriously difficult for life-cycle models to match household portfolio allocations between stocks and short-term bonds. In contrast, our findings show that a relatively simple model can match the key cross-sectional features of the allocation of wealth between short- and long-term assets.

Robustness Importantly, if labor earnings and Social Security are modeled properly, our cross-sectional predictions should be robust to model misspecifications. This is because cross-sectional predictions directly follow from the substitution effects of Social Security and human capital described in equation (30) and would not change as long as the relative weights and interest-rate sensitivities of background assets remain the same.

We chose to estimate φ based on moments from the yield curve, which implies a persistence of rate shocks of $\varphi = .956$. It seems, however, that investors and economic forecasters have been repeatedly surprised by the long-run decline in interest rates (Council of Economic Advisers, 2015), which suggests that the true persistence may be higher than the yield curve implied. In Appendix D.1, we consider how our findings would change if φ was .02 higher, meaning equal to $.976$. We find that this change in calibration raises empirical and simulated estimates of the interest-rate sensitivity of wealth in a similar fashion, such that the model still matches well the data with the same preference parameters.

5.3 Household interest-rate exposure

We now study the rate sensitivities of two measures that are more relevant for welfare: wealth inclusive of Social Security and expected lifetime utility. As in the linearized model, we find that there is less heterogeneity in these measures (especially in utility), suggesting that the recent rise in wealth inequality has not necessarily come with a rise in welfare inequality.

To calculate wealth inclusive of Social Security, we capitalize the expected benefits and taxes into a present value. To measure welfare, we calculate the transformed expected utility U defined in (32). Because U is a function of both wealth W and rates r_f, this elasticity can be approximated, to a first order, as

$$\varepsilon_r(U) \approx \underbrace{-\frac{\partial \log U}{\partial r_f}}_{\substack{\text{change in investment}\\ \text{opportunities}}} + \underbrace{\frac{\partial \log U}{\partial \log W}\varepsilon_r(W)}_{\text{capital gains}} \tag{36}$$

When rates decline, expected utility decreases because investment opportunities are worse, but also increases because of capital gains in financial wealth. If $\varepsilon_r(U)$ is negative, as we find, then a decline in rates decreases welfare.

Figure 9: Interest-rate sensitivities over the life cycle

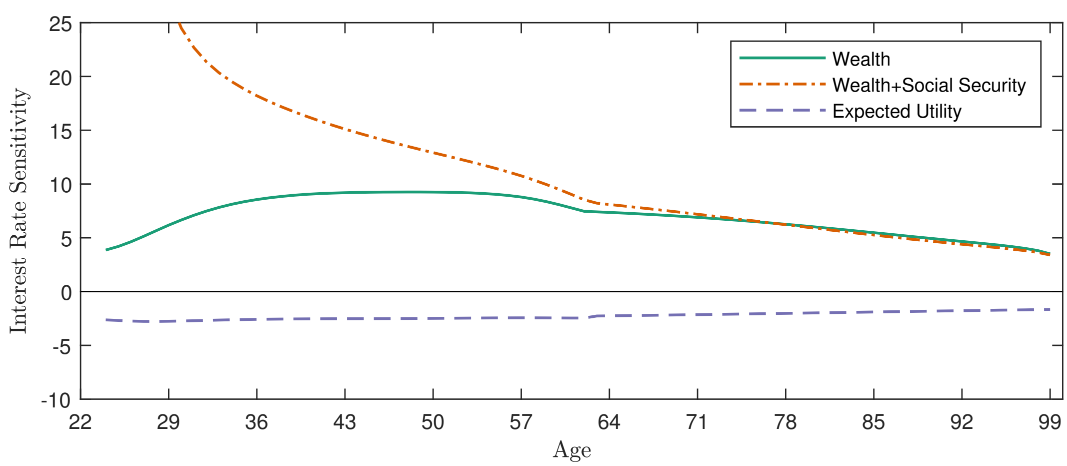

Note: This figure reports the interest-rate sensitivity of wealth, wealth inclusive of Social Security, and expected utility over the life cycle.

Figure 9 shows the average paths of these sensitivities over the life cycle. Adding Social Security wealth increases the average sensitivity for the young, consistent with the fact that it is a very long-term asset.[9] The rate sensitivity of expected

[9]We only report the interest-rate sensitivity of Wealth+Social Security starting at age 25. Be-

utility, on the other hand, is flatter over the life cycle. At all ages, households are negatively affected by rate declines on net. The magnitude of this effect is slowly declining over the life cycle as the investment horizon declines with age. For instance, $\varepsilon_r(U)$ is -2.7 at age 25 and -2.2 at age 65. The closeness of these numbers means that even different cohorts have relatively similar total exposure to rate risk.

Figure 10 reports the distribution of these sensitivities within a middle-aged cohort. First, when Social Security is taken into account, the wealth of the rich and of high earners is no longer more sensitive to interest rates. This explains the findings of Catherine et al. (2023) that, when Social Security is accounted for and discounted using the market yield curve, wealth inequality has not increased since 1989. Figure 10 is the model's counterpart to Figure 3. The data and the model convey the same message: once Social Security is accounted for, rich households no longer hold more interest-rate-sensitive assets. Like in the data, the interest-rate sensitivity of wealth, inclusive of Social Security, is around 15, though a bit higher for low-wealth households in the data.

Second, within a cohort, expected utility is uniformly elastic to interest rates across the earnings and wealth distributions, save for a very minor effect from income-driven mortality differences. As equation (29) predicts, once human capital and Social Security wealth are accounted for, all households within a cohort are equally exposed to rate fluctuations.

cause Social Security wealth can be negative in high rate environments, we find that Wealth+Social Security can be negative or extremely close to zero in the first two years, leading to diverging values of ϵ_r on very small dollar amounts.

Figure 10: Interest-rate sensitivities at age 42

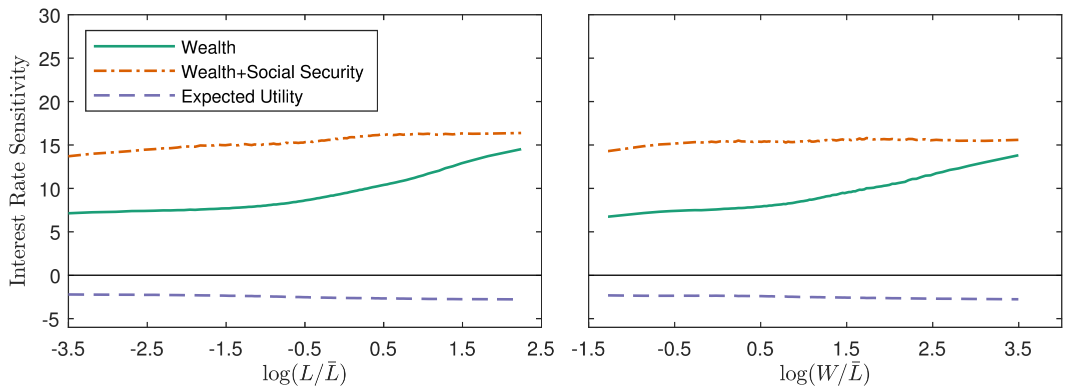

Note: This figure reports the sensitivity of wealth, excluding and including Social Security, and of expected utility to an interest-rate decline at age 42.

6 Trends in wealth inequality

So far, we have demonstrated that, under a reasonable calibration, the interest-rate risk hedging demand for long-term assets can explain differences in the interest-rate sensitivity of wealth over the life cycle, and across the earnings and wealth distribution. Consequently, in the model, interest-rate fluctuations will, through our mechanism, redistribute wealth across the population without significant implications for the distribution of welfare or lifetime consumption. The goal of this section is to assess how much of the long-run fluctuations in wealth inequality can be explained by this channel.

6.1 Economic intuition

Wealth evolves according to:

$$\frac{W_{i,t+1}}{W_{it}} = \underbrace{\left(1 - \frac{C_{it} - Y_{it}}{W_{it}}\right)}_{\text{savings}} \underbrace{R_{Wi,t+1}}_{\text{portfolio}}, \tag{37}$$

where Y_i is the sum of labor income, taxes, and benefits. Interest rates can shape trends in wealth inequality through two channels: differences in consumption-wealth ratios (savings) and differences in portfolio allocations (returns).

We can decompose changes in inequality over time by taking logs of (37) and then computing cross-sectional variances. Doing so yields the change in wealth inequality from one period to the next,[10]

$$\begin{aligned} \text{var}_I(w_{i,t+1}) - \text{var}_I(w_{it}) &= \text{var}_I(s_{it}) + \text{var}_I(r_{wi,t+1}) \\ &\quad + 2\text{cov}_I(w_{it}, s_{it}) + 2\text{cov}_I(w_{it}, r_{wi,t+1}) + 2\text{cov}_I(s_{it}, r_{wi,t+1}). \end{aligned} \tag{38}$$

The first two channels through which wealth inequality may increase are the direct effects of heterogeneous savings rates s and realized portfolio returns r_w. The remaining three channels are captured by the covariance terms. Inequality increases if (i) the wealthy tend to save more, (ii) the wealthy experience higher returns and (iii) households with higher savings rates experience higher returns. Our model reveals why these covariance channels are all positive when interest rates fall.

First, consider intra-cohort wealth inequality. Absent Social Security, there is no variation in savings rates or portfolio choices within a cohort, so wealth inequality will just mirror post-tax income inequality.

The presence of Social Security, in contrast, generates changing inequality through

[10] Lowercase letters denote logs; $s \equiv \log(1 - (C - Y)/W)$ denotes the log savings rate.

both the savings and portfolio channels. First, inequality increases because Social Security induces differential savings rates: low-income, low-wealth households with higher replacement rates will save less into financial wealth. This savings substitution effect of Social Security results in a dispersion in savings rates ($\mathrm{var}_I(s_{it}) > 0$) and a positive wealth-savings correlation ($\mathrm{cov}_I(w_{it}, s_{it}) > 0$). Second, Social Security gives rise to changes in inequality via its impact on portfolio choices. As Figure 6 illustrates, the substitution effects of Social Security create a positive correlation between a household's wealth and its rate exposure. Thus, households within a cohort may experience different wealth returns, and the direction of reallocation will depend on the direction of the interest-rate shock. All unexpected rate changes result in heterogeneous returns ($\mathrm{var}_I(r_{wi,t+1}) > 0$). A negative rate shock will result in disproportionately high returns for the wealthy ($\mathrm{cov}_I(w_{it}, r_{wi,t+1}) > 0$), increasing inequality. A positive rate shock will do the opposite. Finally, since the wealthy save more, the savings-return covariance $\mathrm{cov}_I(s_{it}, r_{wi,t+1})$ is also positive given a rate decline, amplifying the increase in inequality.

6.2 Overlapping-generations simulation

To quantify the role of the interest-rate risk hedging channel on wealth inequality trends, we set up an overlapping-generations version of our life-cycle model. Specifically, we simulate the lives of cohorts born between 1880 and 1986 and feed the model with the historical time series of interest rates and interest-rate shocks. To obtain this time series, we use a methodology similar to that of Beeler and Campbell (2012) described in detail in Appendix A.3. We assume that, when a household dies, its wealth is transferred to a household from a cohort that is thirty years younger. For simplicity, we assume this transfer of wealth to be unexpected.

Our model is not ideally situated to match the *level* of wealth inequality. First, the wealth concentration in the top 1% of the distribution comes primarily from business income, which is omitted from the model. Therefore, we calibrate our model to match the level of inequality within the bottom 99%,[11] by increasing the standard deviation of earnings fixed effects $\sigma_{z,0}$ from .652 to 1.1.[12]

Figure 11 illustrates our results. In our historical simulation, the top 10% share falls from 56% in 1956 to 52% in 1984, then rises back to 56% in 2019. According to the World Inequality Database, the top 10% (within the bottom 99%) share fell from 58.6% to 49% in 1984, then rose back to 55% in 2019. The top 10%, inclusive of the top 1%, fell from 70.3% in 1962 to 62.1% in 1985, then rose back to 70.7% in 2019. Consequently, the interest-rate risk hedging channel can explain roughly half of the long-run fluctuations in wealth inequality since the mid-1950s.

As predicted in Section 4.4, these trends in wealth inequality do not translate into increased welfare inequality within a cohort. To illustrate this point, Figure 11 also reports the evolution of welfare inequality within cohorts, using the certainty equivalent measure from equation (32). As predicted by theory, welfare inequality remains largely unchanged.[13]

[11] We approximate this measure as (Top 10% share − Top 1% share)/(100% − Top 1% share).

[12] There are several reasons why earnings inequality in administrative data, measured at the worker level, would be lower than at the household level across the entire population. For example, part of the population does not participate in the labor force and the number of earners varies across households.

[13] If anything, welfare inequality trends go in the opposite direction. This is because, in the full model, high earners have higher life expectancy and are therefore slightly more exposed to interest-rate risk, as predicted in equation (34).

Figure 11: Evolution of top 10% share

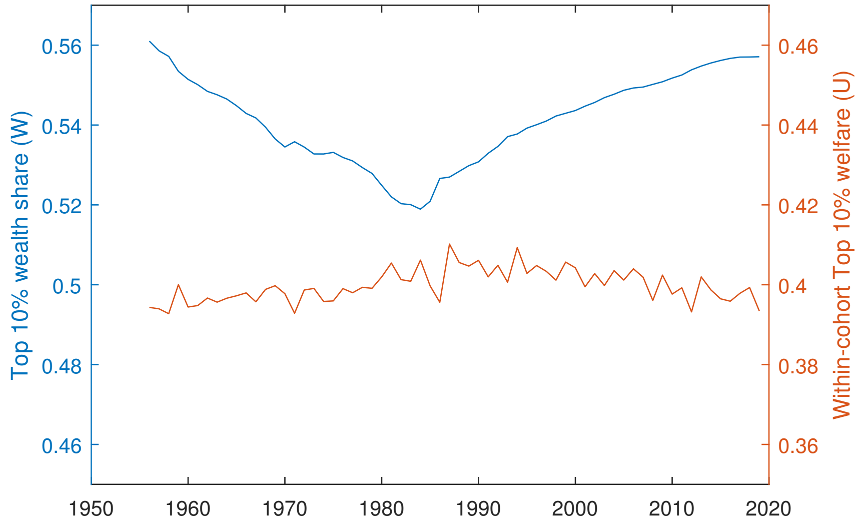

Note: This figure reports, on the left axis, the evolution of the top 10% wealth share in simulated data using the historical path of real interest rates. On the right axis, it reports the evolution of the share of *total wealth*, or welfare (U), as defined in equation (32), going to the top 10% of each cohort, and averaged across cohorts in any given year. In each time series, the demographic weights of living cohorts are weighted using the historical age pyramid.

Greenwald et al. (2023) conduct a similar quantitative exercise and find that the fall in real interest rate explains 75% of the rise in wealth inequality because the rich invest more in long-term assets. Our model only explains 50% but the gap is explained by the different natures of the two exercises. Greenwald et al. (2023) estimate the role played by the *empirical* relationship between wealth and duration, including the rapid growth of private business wealth owned by the very rich. In contrast, in our study, the long-term asset share is endogenous and limited to the interest-rate risk hedging demand alone. Hence, our model provides a micro-foundation for two-thirds of the effect documented in their paper.

7 Conclusion

We study household exposure to interest-rate risk, beginning from the empirical observation that middle-aged and richer households invest more in long-term assets that appreciate in value when rates fall. This interest-rate sensitivity must be interpreted in light of households' life-cycle problem, so we develop a life-cycle model that incorporates the roles that human capital and Social Security play in households' portfolio choices.

The data and the normative model produce strikingly close patterns: the optimal long-term asset share is hump-shaped over the life cycle and, within cohorts, increases with wealth and earnings. This is driven in part by the heterogeneous role our background assets play: human capital displaces the need for long-term holdings for the young, and Social Security has the same effect for the low- and middle-income households on which its impact is most pronounced.

So, in practice, households are targeting the right interest-rate exposure over their life cycles. Still, we cannot conclude from our findings that households would continue to do this if the economic environment were to change. For example, we cannot say whether households adjust their portfolios optimally in response to Social Security or whether Social Security is well-designed to correct investment mistakes that would arise in its absence. Determining the direction of this causality is necessary for evaluating policy counterfactuals and is therefore an essential step for future research.

Our paper focuses on interest-rate risk, motivated by the fact that several empirical studies have shown it is a significant driver of past returns and inequality dynamics (Binsbergen, 2021; Greenwald et al., 2023). Still, the core logic of the paper could, in principle, apply to any risk factor: differences in observed wealth exposure to systematic risks may reflect differences in exposure to household-specific background risks. Thus, more research is also needed to determine the extent to which other major risk factors—like economic growth, house prices, and inflation—shape household portfolio choices. In general, our paper advances a view of portfolio choice in which households decide asset allocations across common risk factors, as opposed to across asset types.

References

Arrow, K. J., "The Role of Securities in the Optimal Allocation of Risk-bearing," *The Review of Economic Studies*, 1964, *31* (2), 91–96.

Auerbach, Alan J., Laurence J. Kotlikoff, and Darryl Koehler, "US Inequality and Fiscal Progressivity: An Intragenerational Accounting," *Journal of Political Economy*, 2023, *131* (5), 1249–1293.

Bach, Laurent, Laurent E. Calvet, and Paolo Sodini, "Rich Pickings? Risk, Return, and Skill in Household Wealth," *American Economic Review*, September 2020, *110* (9), 2703–47.

Beeler, Jason and John Y. Campbell, "The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment," *Critical Finance Review*, 2012, *1* (1).

Benhabib, Jess, Alberto Bisin, and Mi Luo, "Wealth Distribution and Social Mobility in the US: A Quantitative Approach," *American Economic Review*, May 2019, *109* (5), 1623–47.

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.

Binsbergen, Jules van, "Duration-Based Stock Valuation," 2021. NBER Working Paper No. 27367.

Calvet, Laurent E, John Y Campbell, Francisco Gomes, and Paolo Sodini, "The Cross-Section of Household Preferences," May 2021, (28788).

Campbell, John Y. and Luis M. Viceira, "Who Should Buy Long-Term Bonds?," *American Economic Review*, 2001, *91* (1), 99–127.

Catherine, Sylvain, "Countercyclical Labor Income Risk and Portfolio Choices over the Life Cycle," *The Review of Financial Studies*, 2022. hhab136.

Catherine, Sylvain and Constantine Yannelis, "The Distributional Effects of Student Loan Forgiveness," 2021. BFI Working Paper No. 2020-169.

Catherine, Sylvain, Max Miller, and Natasha Sarin, "Social security and trends in inequality," 2023.

Chetty, Raj, 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.

Cocco, Joao F., Francisco J. Gomes, and Pascal J. Maenhout, "Consumption and Portfolio Choice over the Life Cycle," *Review of Financial Studies*, 2005, *18* (2), 491–533.

Cochrane, John H, "Portfolios for Long-Term Investors*," *Review of Finance*, 2 2022, *26* (1), 1–42.

Council of Economic Advisers, "Long-tern Interest Rates: A Survey," 2015.

Fagereng, Andreas, Martin Blomhoff Holm, Benjamin Moll, and Gisle Natvik, "Saving Behavior Across the Wealth Distribution: The Importance of Capital Gains," 2021. NBER Working Paper No. 26588.

Flavin, Marjorie and Takashi Yamashita, "Owner-Occupied Housing and the Composition of the Household Portfolio," *American Economic Review*, March 2002, *92* (1), 345–362.

Greenwald, Daniel, Matteo Leombroni, Hanno N. Lustig, and Stijn van Nieuwerburgh, "Financial and Total Wealth Inequality with Declining Interest Rates," 2023. Stanford University Graduate School of Business Research Paper.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song, "What Do Data on Millions of U.S. Workers Reveal About Life-Cycle Earnings Risk?," *Econometrica*, 2022, *89* (5), 2303–2339.

Hubmer, Joachim, Per Krusell, and Anthony A. Smith, "Sources of US Wealth Inequality: Past, Present, and Future," *NBER Macroeconomics Annual*, 2021, *35*, 391–455.

Huggett, Mark and Greg Kaplan, "How large is the stock component of human capital?," *Review of Economic Dynamics*, 2016, *22*, 21–51.

Lustig, Hanno, Stijn Van Nieuwerburgh, and Adrien Verdelhan, "The Wealth-Consumption Ratio," *The Review of Asset Pricing Studies*, 04 2013, *3* (1), 38–94.

Lynch, Anthony W. and Sinan Tan, "Labor income dynamics at business-cycle frequencies: Implications for portfolio choice," *Journal of Financial Economics*, 2011, *101* (2), 333–359.

Meeuwis, Maarten, "Wealth Fluctuations and Risk Preferences: Evidence from U.S. Investor Portfolios," *Working Paper*, 2022.

Merton, Robert C., "Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case," *The Review of Economics and Statistics*, August 1969, *51* (3), 247–57.

Meyer, Bruce D. and James X. Sullivan, "Consumption and Income Inequality in the United States since the 1960s," *Journal of Political Economy*, 2023, *131* (2), 247–284.

Moll, Benjamin, "Comment on "Sources of US Wealth Inequality: Past, Present, and Future"," *NBER Macroeconomics Annual*, 2021, *35*, 468–479.

Samuelson, Paul A., "Lifetime Portfolio Selection By Dynamic Stochastic Programming," *The Review of Economics and Statistics*, 1969, *51* (3), 239–246.

Schoar, Antoinette, "The Divide between Subsistence and Transformational Entrepreneurship," *Innovation Policy and the Economy*, 2010, *10* (1), 57–81.

Viceira, Luis M., "Optimal portfolio choice for long-horizon investors with nontradable labor income," *The Journal of Finance*, 2001, *56* (2), 433–470.

INTERNET APPENDIX

A Data appendix

A.1 Survey of Consumer Finances

Data on household portfolios come from the Survey of Consumer Finances (SCF). We construct networth as:[14]

$$\begin{aligned}\texttt{networth\_d} = \texttt{cash\_dep} + \texttt{equity} + \texttt{fixed\_inc} + \texttt{real\_estate}\\ + \texttt{bus} + \texttt{vehic} - \texttt{mortgage\_dbt} - \texttt{vehic\_dbt} - \texttt{other\_dbt},\end{aligned}$$

where each of the constituent variables are defined as:

- `cash_dep`: value of cash deposits defined as liquid accounts (`liq`) which are the sum of all checking, savings, and money market accounts, call accounts at brokerages, and prepaid cards, added to certificates of deposit (`cds`).
- `equity`: value of all financial assets invested in stock, which include directly held stock, stock mutual funds, and the portion of any combination mutual funds, annuities, trusts, IRA/Keogh accounts, and other retirement accounts invested in stock.

[14]Note that we do not include student debt in our analysis for several reasons. In the US, student debt is largely repaid through income-driven repayment (IDR) programs. In IDR, borrowers' monthly payments are determined as a fraction (10–15%) of their earnings, above a certain family-size-dependent threshold. After 10 to 25 years of payments, their remaining balance is forgiven. This system has both practical and conceptual implications. Practically, it means that the information collected by the SCF regarding balance and scheduled payments do not accurately reflect the actual present values and cashflow duration of the debt (Catherine and Yannelis, 2021). Conceptually, student debt payments are more akin to a progressive tax on earnings and therefore would be better treated as a deduction to human capital. For context, Catherine and Yannelis (2021) estimate that in 2019, the average working-age American had $4,922 in student debt (in present value terms) and that student debt had an average interest-rate-sensitivity of 6.

– `fixed_inc`: value of all other remaining financial assets (`fixed_inc` = `fin` − `cash_dep` − `equity`). The largest component of this asset category is bonds held outright, in mutual funds, and in retirement accounts.

– `real_estate`: value of the primary residence (`houses`) plus the value of other residential real estate (`oresre`) and net equity in nonresidential real estate (`nnresre`).

– `bus`: reported market value of private business interest.

– `vehic`: prevailing retail value for all vehicles owned by household.

– `mortgage_dbt`: housing debt from mortgages, home equity loans, and home equity lines of credit (`mrthel`) plus debt for other residential property (`resdbt`).

– `vehic_dbt`: debt from vehicle loans (`veh_inst`)

– `other_dbt`: other debt, including other lines of credit plus credit card balance (`ccbal`) plus installment loans less education loans and vehicle loans (`other_dbt` = `othloc` + `ccbal` + `install` − `edn_inst` − `veh_inst`).

In addition to portfolio data, we also use data on household wage income (`wageinc`) which we combine with data on the number of people in the household and the Social Security wage index to create a per capita wage measure that is comparable over time.

A.2 Duration component calculations

A.2.1 Duration of equity

The duration of `equity` is obtained using annual estimates for the duration of the aggregate stock market from Greenwald et al. (2023), Figure D2 of the September 2023 working paper version. These estimates are applied uniformly to all individuals in the SCF by survey year.

A.2.2 Duration of fixed income

Data on the Macaulay duration of government bonds, municipal bonds, corporate bonds, and mortgage backed securities come from Bloomberg where the series used are:

– U.S. gov/credit: `LUGCTRUU`

- U.S. Treasury: `LUATTRUU`
- Government-related: `LD08TRUU`
- U.S. aggregate: `LBUSTRUU`
- Municipal bond: `LMBTTR`
- Corporate: `LUACTRUU`
- U.S. MBS: `LUMSTRUU`
- Global aggregate: `LEGATRUU`

For holdings of U.S. government bonds (`govtbnd` + `gbmutf` + `savbnd`), we use the market-value weighted average Macaulay duration of the U.S. gov/credit, U.S. Treasury, and government-related bond categories. For holdings of tax-free and municipal bonds (`notxbnd` + `tfbmutf`), mortgage-backed securities (`mortbnd`), corporate bonds (`corpbnd`), and foreign bonds (`forbnd`), we use the Macaulay duration of municipal bonds, corporate bonds, U.S. MBS, and the global aggregate, respectively. For all other fixed income assets that we do not have duration measures for, we assign 5.64, which is the average fixed income wealth-weighted duration in the SCF assets which we have data.

A.2.3 Duration of real estate

The duration of `real_estate` is obtained using the annual estimates of the duration of aggregate real estate from Greenwald et al. (2023), Figure D2 of the September 2023 working paper version. These estimates are applied uniformly to all individuals in the SCF by survey year.

A.2.4 Duration of private business wealth

The duration of private business wealth is computed for each household as the value of household businesses, `bus`, divided by the annual cashflows from those equity holdings. However, the annual cashflows from those equity holdings are not reported in the SCF, the major issue being that cash-flows from private businesses partially contain implicit or explicit labor income for the entrepreneur. As such, we must estimate or difference out this labor income. We do this in four ways, depending on the household's role in the business and what is reported.

1. For households whose main respondent has an active management role in either of the household's potential actively managed businesses, reports being self-employed, and reports not receiving a salary, we estimate their predicted wage.

   – The predicted wage is estimated via ordinary least squares on all SCF respondents j where the household's wage income is the dependent variable, and the independent variables are a third-degree polynomial in age interacted with dummies for each Race $\times$ Education $\times$ Gender group.

2. For households whose main respondent has an active management role in either of the household's potential actively managed businesses and reports being self-employed and receiving a salary or reports being employed by someone else, we subtract the maximum of their predicted wage and reported wage from `busefarminc`.

3. We repeat steps 1) and 2) for spouses who have an active management role in either of the household's potential actively managed businesses.

4. All other households with positive private business wealth who do not meet the criteria for a wage subtraction are given cashflows equal to `busefarminc`.

We then aggregate `bus` and the estimated annual cashflows within each survey year and divide them to obtain an annual time series of valuation ratios.

Next, to allow our aggregate estimates of private business duration to vary over the wealth distribution, we perform a mean-preserving adjustment to these aggregate duration estimates. First, we split the population into different networth groups defined by whether they are in the bottom 50%, 50–90%, 90–99%, 99–99.9%, 99.9–99.99% or the top 0.01% of the wealth distribution. We then take the business wealth (`bus`) divided by total income from businesses (`busefarminc`) for each household, and take the business wealth-weighted average for each networth group. Provided that cashflows from equity are proportional to labor income, this provides a proxy for duration for each group. These price-total income ratios are then divided by the business wealth-weighted average for the aggregate population to obtain a mean-preserving adjustment which is applied to the annual aggregate private business duration estimates. This is given by

$$\text{dur}(\text{Private business}_{ct}) = \frac{\text{Price-total income ratio}_c}{\text{Price-total income ratio}} \times \text{dur}(\text{Private business}_t). \tag{A.1}$$

A.2.5 Duration of vehicles

The `vehic` category in the SCF contains detailed information on up to 4 automobiles, up to 2 non-automobile vehicles, and an aggregation of additional automobiles and non-automobile vehicles owned by the household. For the primary automobiles of the house, we attribute an expected lifetime of 8 years for 1989 and 12 years for 2019, linearly interpolating in intermediate years. We calculate the time left on an automobile's life as the model year plus the expected age minus the survey year. We assume a fixed depreciation rate to 0 over the car's remaining years, and calculate the duration using (4). We attribute a duration of one to vehicles whose age exceeds their expected lifetime.

For the aggregation of additional automobiles owned, we attribute a duration equal to the average of the duration of the first four automobiles owned by the household. For all non-automobile vehicles owned by the household, we ascribe a duration of 6 years.

A.2.6 Duration of debts

For the debt categories, `mortgage_dbt`, `vehic_dbt`, and `other_dbt`, we break each up into their component loans as described in the SCF extract and calculate the duration of each loan separately. For each loan, we assume a fixed payment schedule, and thus its duration can be calculated using equation (4), where N is the maturity of the loan and y_{nt} is the riskfree spot rate at horizon n in year t.

Under our fixed payment assumption, the only metric we need for each loan is its time remaining. Since different loan component variables contain different amounts of information in the raw SCF, we calculate the time remaining differently depending on the available information for each component loan group: primary component loans, aggregated additional loans, and lines of credit. The primary component loans of each debt category contain information on loan origination, balance, payments, and interest rates. For these loans, we calculate the number of years remaining on the loan payments using the reported origination year, length of loan at origination, and survey year. For respondents with a positive loan balance who have missing responses for loan length or a negative calculated time remaining, we impute time remaining with balance (B), initial amount (L), interest rate (R), and year of origination (p) using the equation

$$T = \frac{\log(R^p - B/L) - \log(1 - B/L)}{\log R} - p.$$

The aggregated additional loans group contains loan variables that capture an aggregation of loans that the respondents hold in addition to the primary ones in each debt category. These loans include data on only loan balance and payments (X). Using the average interest rates for primary loans in the same debt category, we calculate time remaining as

$$T = -\frac{\log(1 - B(R-1)/X)}{\log R}.$$

The third group of component loans is the lines of credit. The line of credit variables contain information on loan balance, typical payments, and interest rates. With these data points, we calculate time remaining according to the same formula used for the aggregated additional loans group. Finally, there is an aggregated additional lines of credit variable, which we assign a duration equal to the average of the duration of the other lines of credit.

We replace the duration of loans with a predicted time remaining under one year with a duration of one and give the median duration to respondents with a positive loan amount but insufficient information to calculate time remaining on the loan.

A.2.7 Interest-rate sensitivity of Social Security wealth

The interest-rate sensitivity of Social Security wealth comes from the methodology used in Catherine et al. (2023). We generate their baseline risk-adjusted Social Security wealth under the net present value wealth concept and the Treasury yield curve. We then generate Social Security wealth under an identical specification where the log forward rate at horizon h in survey year t is given by $\tilde{f}_{h,t} = f_{h,t} + \varphi^h 0.01$ where $f_{h,t}$ is the unshocked log forward rate and φ is our calibrated persistence for the interest rate process shown in Table 1. This is the same thing as applying a one-period shock of 0.01 to the log riskfree rate under the process in equation (1).

A.3 Time series of riskfree rates

Figure A.1: Time series of riskfree rates, Post-war sample

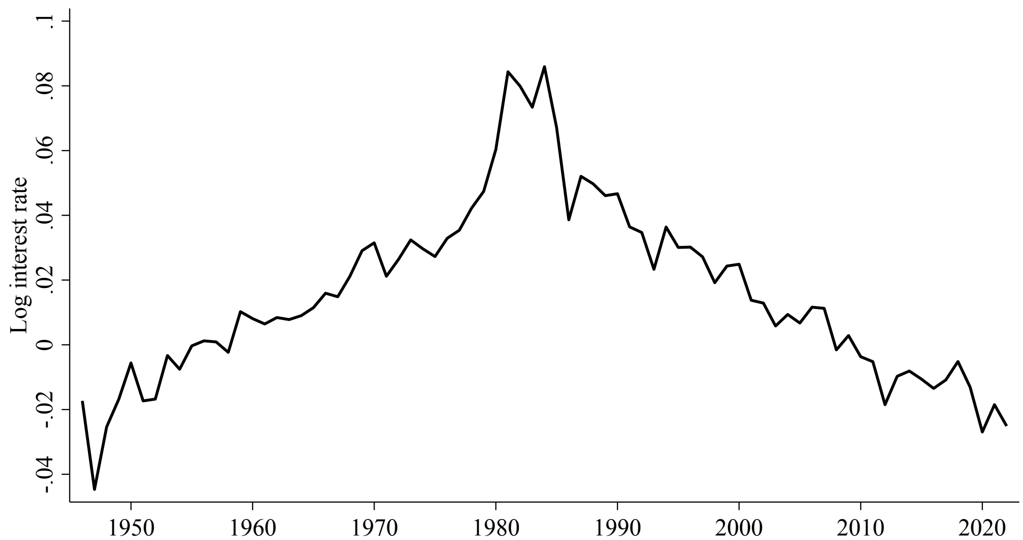

Note: This figure presents the time series of short-term riskfree rates as estimated by equation (A.2) and transformed by equation (A.3).

To obtain a time series of the short-term real interest rate, we use a methodology similar to that of Beeler and Campbell (2012). Using the yield on the 10-year nominal Treasury bond y_{10} and annual inflation rate π from Global Financial Data, we estimate the annual regression

$$y_{10,t} - \pi_{t,t+1} = \beta_0 + \beta_1 y_{10,t} + \beta_2 \pi_{t-1,t} + \epsilon_{t+1} \tag{A.2}$$

on the post-war period. The fitted values are then taken as our estimate of the expected riskfree rate 10-years from time t, $\hat{f}_{10,t}$. From this, equation (1) yields the time-t riskfree rate:

$$r_{ft} = \varphi^{-10}(\hat{f}_{10,t} - (1 - \varphi^{10})\bar{r}_f). \tag{A.3}$$

We use this methodology for two main reasons. First, by using long-term rates to back out short-term rates, we smooth much of the short-term variation in measured short-term real rates that are potentially outside of our model. Second, this methodology allows us to extend our real rate series further into the past, allowing for a longer simulation prior to our period of interest. This procedure yields a time series of annual realizations of real rates $\{r_{ft}\}$ and shocks $\{\epsilon_{rt}\}$ from 1789 to 2020. The post-war time series of these rates are shown in Figure A.1.

B Model appendix

B.1 Details on income tax rates

Section 3.5 discusses the taxes paid on labor income and Social Security benefits. In the model, households face the following marginal tax rates:

$$\text{Marginal Tax Rate}_{it} = \begin{cases} 0.10 & \text{if } \tilde{L}_{it} < 0.18, \\ 0.12 & \text{if } 0.18\tilde{L}_{it} < 0.72, \\ 0.22 & \text{if } 0.72\tilde{L}_{it} < 1.54, \\ 0.24 & \text{if } 1.54\tilde{L}_{it} < 2.94, \\ 0.32 & \text{if } 2.94\tilde{L}_{it} < 3.73, \\ 0.35 & \text{if } 3.73\tilde{L}_{it} < 9.32, \\ 0.37 & \text{if } \tilde{L}_{it} > 9.32. \end{cases} \tag{B.1}$$

The bendpoints in this formula are the limits of the 2020 tax brackets divided by the wage index.

B.2 Derivation of long-term bond returns

This section explains how the riskfree rate dynamics (1) imply the n-period bond returns (9). Since it has no intermediate cash flows, the bond's return from t to $t+1$ is

$$R_{n,t+1} = \frac{P_{n-1,t+1}}{P_{nt}} \equiv \frac{e^{-(n-1)y_{n-1,t+1}}}{e^{-ny_{nt}}}, \tag{B.2}$$

where the yield y_{nt}, under the expectations hypothesis, is given by

$$y_{nt} \equiv \frac{1}{n}\log\left(\frac{1}{P_{nt}}\right) = \frac{1}{n}\sum_{j=1}^{n}(\mathbb{E}_t r_{f,t+j-1} + \mu_j). \tag{B.3}$$

Moreover, note that (1) iterates backward to the expression

$$r_{f,t+j} = (1-\varphi^j)\bar{r}_f + \varphi^j r_{ft} + \sum_{k=1}^{j} \varphi^{j-k}\sigma_r \epsilon_{r,t+k}. \tag{B.4}$$

Substituting the riskfree rates (B.4) into the yield expression (B.3) and evaluating expectations implies

$$y_{nt} = \bar{r}_f + \frac{1}{n}\frac{1-\varphi^n}{1-\varphi}(r_{ft} - \bar{r}_f) + \frac{1}{n}\sum_{j=1}^{n}\mu_j. \tag{B.5}$$

Taking logs of (B.2) and substituting (B.5) into the yield then implies the log return

$$\begin{aligned}
r_{n,t+1} &= ny_{nt} - (n-1)y_{n-1,t+1} \\
&= \bar{r}_f + \frac{1-\varphi^n}{1-\varphi}(r_{ft} - \bar{r}_f) - \frac{1-\varphi^{n-1}}{1-\varphi}(r_{f,t+1} - \bar{r}_f) + \mu_n \\
&= \bar{r}_f + \frac{1-\varphi^n}{1-\varphi}(r_{ft} - \bar{r}_f) - \frac{\varphi-\varphi^n}{1-\varphi}(r_{ft} - \bar{r}_f) - \frac{1-\varphi^{n-1}}{1-\varphi}\sigma_r\epsilon_{r,t+1} + \mu_n \\
&= r_{ft} + \mu_n - \frac{1-\varphi^{n-1}}{1-\varphi}\sigma_r\epsilon_{r,t+1},
\end{aligned}$$

the stated expression (9).

C Derivation of the linearized model

This section lays out the details of the linearization and analytical solutions presented in Section 4. The approach follows that of Campbell and Viceira (2001), except that we add finite lives and, ultimately, intertemporal income. To fully understand the economics, we first solve for policies in the general case of recursive utility (i.e., disentangling risk aversion and the EIS), then reduce to the time-additive case in the main text. For the remainder of this appendix section, we will suppress i indices and state approximate (i.e., linearized) equalities as exact.

C.1 Linearized conditions

Suppose that there is no intertemporal income, so the budget constraint (7) simplifies to

$$W_{t+1} = (W_t - C_t)R_{W,t+1}. \tag{C.1}$$

The first-order condition for a recursive-utility agent takes the familiar form

$$1 = \mathbb{E}_t \left[(\beta(1-m_t))^\theta \left(\frac{C_{t+1}}{C_t} \right)^{-\theta/\psi} R_{W,t+1}^{\theta-1} R_{j,t+1} \right], \tag{C.2}$$

where $\beta(1-m_t)$ is mortality-adjusted patience, ψ is the EIS, $\theta = (1-\gamma)/(1-1/\psi)$, and $R_j \in \{R_f, R_n, R_W\}$. The analytical solution follows from linearizing this budget constraint and first-order condition.

Let lowercase letters denote logs and the Δ operator denote first-differences. Scaling the budget constraint (C.1) by financial wealth W_t, taking logs, and linearizing $\log(1-e^{c_t-w_t})$ around $c_t - w_t = \log(1-\beta(1-m_t))$ implies

$$\Delta w_{t+1} = \kappa_w(m_t) + \left(1 - \frac{1}{\rho_c(m_t)}\right)(c_t - w_t) + r_{w,t+1}, \tag{C.3}$$

where $\rho_c(m_t) = \beta(1-m_t)$ and $\kappa_w(m_t) = \log \rho_c(m_t) + (1-\rho_c(m_t)) \log(1-\rho_c(m_t))/\rho_c(m_t)$.[15] (Notice that, as $m_t \to 1$, $c_t \to w_t$; agents who will die almost surely consume everything.) We can also get the linearized approximation to the log wealth return

$$r_{w,t+1} = r_{ft} + \pi_t(r_{n,t+1} - r_{ft}) + \frac{1}{2}\pi_t(1-\pi_t)\text{var}_t(r_{n,t+1}). \tag{C.4}$$

This expression is a discretization of the exact continuous-time law of motion. Finally, log-linearize the Euler equation (C.2) up to a second order:

$$\begin{aligned} 0 = \theta \log(\beta(1-m_t)) + \mathbb{E}_t \left[-\frac{\theta}{\psi}\Delta c_{t+1} + (\theta-1) r_{w,t+1} + r_{j,t+1} \right] \\ + \frac{1}{2}\text{var}_t \left(-\frac{\theta}{\psi}\Delta c_{t+1} + (\theta-1) r_{w,t+1} + r_{j,t+1} \right). \end{aligned} \tag{C.5}$$

Substituting in $r_j = r_n$ and then $r_j = r_f$ and subtracting the two equations implies the risk premium on the long-term bond

$$\mathbb{E}_t[r_{n,t+1} - r_{f,t+1}] + \frac{1}{2}\text{var}_t(r_{n,t+1}) = \frac{\theta}{\psi}\text{cov}_t(r_{n,t+1}, \Delta c_{t+1}) + (1-\theta)\text{cov}_t(r_{n,t+1}, r_{w,t+1}). \tag{C.6}$$

[15] In infinite-horizon models like that of Campbell and Viceira (2001), one typically chooses $\rho_c = 1 - \exp\{\mathbb{E}[c_t - w_t]\}$, which reduces to $\rho_c = \beta$ for EIS of 1. Here, to capture the effect of aging, we linearize instead around the unit-EIS solution, which is exact in our model.

Using the decomposition

$$\Delta c_{t+1} = (c_{t+1} - w_{t+1}) - (c_t - w_t) + \Delta w_{t+1} \tag{C.7}$$

and the expression for Δw_{t+1} from the linearized budget constraint (C.3), we can rewrite

$$\mathrm{cov}_t(r_{n,t+1}, \Delta c_{t+1}) = \mathrm{cov}_t(r_{n,t+1}, c_{t+1} - w_{t+1}) + \mathrm{cov}_t(r_{n,t+1}, r_{w,t+1}).$$

Substituting this and the fact that

$$\mathrm{cov}_t(r_{n,t+1}, r_{w,t+1}) = \pi_t \mathrm{var}_t(r_{n,t+1}) \tag{C.8}$$

into (C.6) and using $\theta/\psi + 1 - \theta = \gamma$ implies the solution

$$\pi_t = \frac{1}{\gamma} \frac{\mathbb{E}_t[r_{n,t+1} - r_{f,t+1}] + \frac{1}{2}\mathrm{var}_t(r_{n,t+1})}{\mathrm{var}_t(r_{n,t+1})} - \frac{1 - 1/\gamma}{1 - \psi} \frac{\mathrm{cov}_t(r_{n,t+1}, c_{t+1} - w_{t+1})}{\mathrm{var}_t(r_{n,t+1})}. \tag{C.9}$$

As explained in the main text, the first term is the myopic risk-return portfolio; the second is intertemporal hedging of rate risk.

Another fact that will become useful is that the first-order condition (C.5) for wealth returns $(r_j = r_w)$ simplifies to

$$\mathbb{E}_t[\Delta c_{t+1}] = \psi \log(\beta(1 - m_t)) + \psi \mathbb{E}_t\left[r_{w,t+1}\right] + \frac{1}{2}\frac{\theta}{\psi}\mathrm{var}_t\left(\Delta c_{t+1} - \psi r_{w,t+1}\right). \tag{C.10}$$

Using fact (C.8) and the decomposition of Δc from (C.7), the variance term can be rewritten as

$$\begin{aligned} \mathrm{var}_t\left(\Delta c_{t+1} - \psi r_{w,t+1}\right) &= \mathrm{var}_t\left(c_{t+1} - w_{t+1} + (1 - \psi) r_{w,t+1}\right) \\ &= \mathrm{var}_t\left(c_{t+1} - w_{t+1}\right) + (1 - \psi)^2 \pi_t^2 \mathrm{var}_t\left(r_{n,t+1}\right) \\ &\quad + (1 - \psi)\pi_t \mathrm{cov}_t\left(r_{w,t+1}, c_{t+1} - w_{t+1}\right). \end{aligned} \tag{C.11}$$

We will use these expressions to solve for the equilibrium consumption-wealth ratio.

C.2 Optimal policies in the linearized model

We will now solve for the optimal consumption and portfolio choices using the conditions derived above. Conjecture that the optimal consumption-wealth ratio takes the form

$$c_t - w_t = \log(1-\beta(1-m_t)) + (1-\psi)(\varrho_{0t} + \varrho_{rt} r_{ft}), \tag{C.12}$$

for some functions $\varrho_{0t} = \varrho_0(\{m_s\}_{s\geq t})$ and $\varrho_{rt} = \varrho_r(\{m_s\}_{s\geq t})$ of the future mortality probabilities. Increasing utility implies the boundary conditions $\lim_{m\to 1}(1-\psi)\varrho_0(m) = 0$ and $\lim_{m\to 1}\varrho_r(m) = 0$. This conjecture implies that

$$\begin{aligned}(1-\psi)^{-1}\mathrm{cov}_t(r_{n,t+1}, c_{t+1} - w_{t+1}) &= \varrho_{rt}\mathrm{cov}_t(r_{n,t+1}, r_{f,t+1})\\ &= -\varrho_{rt}\sigma_n\sigma_r.\end{aligned}$$

Substituting this expression into (C.9), we obtain

$$\begin{aligned}\pi_t &= \frac{1}{\gamma}\frac{\mu_n + \frac{1}{2}\sigma_n^2}{\sigma_n^2} + \left(1-\frac{1}{\gamma}\right)\varrho_{rt}\frac{\sigma_r}{\sigma_n}\\ &= a_0 + a_r\varrho_{rt},\end{aligned}$$

which, combined with (9), is our expression for the optimal share in the n period bond given by (25).

To solve for ϱ_0 and ϱ_r, notice that substituting this solution for π into the expectation of our log-linearized wealth return (C.4) implies

$$\begin{aligned}\mathbb{E}_t[r_{w,t+1}] &= r_{ft} + \pi_t\mu_n + \pi_t(1-\pi_t)\sigma_n^2\\ &= r_{ft} + \left(a_0\mu_n + (a_0 - a_0^2)\sigma_n^2\right) + \left(a_r\mu_n + (a_r - 2a_0a_r)\sigma_n^2\right)\varrho_{rt} - a_r^2\sigma_n^2\varrho_{rt}^2 \qquad \text{(C.13)}\\ &= r_{ft} + d_0 + d_1\varrho_{rt} - d_2\varrho_{rt}^2.\end{aligned}$$

It also implies

$$\begin{aligned}\mathrm{var}_t\left(c_{t+1} - w_{t+1}\right) &= (1-\psi)^2\varrho_{rt}^2\sigma_r^2,\\ (1-\psi)^2\pi_t^2\mathrm{var}_t\left(r_{n,t+1}\right) &= (1-\psi)^2(a_0^2 + 2a_0a_r\varrho_{rt} + a_r^2\varrho_{rt}^2)\sigma_n^2,\\ (1-\psi)\pi_t\mathrm{cov}_t\left(r_{n,t+1}, c_{t+1} - w_{t+1}\right) &= (1-\psi)^2(a_0 + a_r\varrho_{rt})(-\varrho_{rt}\sigma_n\sigma_r).\end{aligned}$$

Substituting these three facts into (C.11), we have

$$\operatorname{var}_t\left(\Delta c_{t+1} - \psi r_{w,t+1}\right) = (1-\psi)^2(g_0 + g_1\varrho_{rt} + g_2\varrho_{rt}^2) \tag{C.14}$$

for constants g_j. Finally, substituting our log-linearized budget constraint (C.3) into our decomposition (C.7) and applying our conjecture (C.12), we have

$$\begin{aligned}\mathbb{E}_t[\Delta c_{t+1}] &= \mathbb{E}_t[c_{t+1} - w_{t+1}] - \rho_c(m_t)^{-1}(c_t - w_t) + \kappa_w(m_t) + \mathbb{E}_t[r_{w,t+1}] \\ &= (1-\psi)(\varrho_{0,t+1} + \varrho_{r,t+1}((1-\varphi)\bar{r}_f + \varphi r_{ft})) \\ &\quad - \rho_c(m_t)^{-1}(1-\psi)(\varrho_{0t} + \varrho_{rt}r_{ft}) + \kappa_w(m_t) + \mathbb{E}_t[r_{w,t+1}].\end{aligned} \tag{C.15}$$

Substituting (C.14), (C.15), and (C.11) into the Euler equation for wealth returns (C.10), then collecting coefficients on r_{ft}, implies the difference equation

$$\varphi\varrho_{r,t+1} = \rho_c(m_t)^{-1}\varrho_{rt} - 1.$$

Now iterate forward and use the boundary condition $\lim_{t\to\infty}\varrho_{rt} = 0$:

$$\begin{aligned}\varrho_{rt} &= \rho_c(m_t)(\varphi\varrho_{r,t+1} + 1) \\ &= \rho_c(m_t) + \varphi\rho_c(m_t)\rho_c(m_{t+1}) + \varphi^2\rho_c(m_t)\rho_c(m_{t+1})\rho_c(m_{t+2}) + \dots \\ &= \beta(1-m_t)\left(1 + \sum_{j=1}^{\infty}\varphi^j\beta^j\prod_{k=1}^{j}(1-m_{t+k})\right) \\ &= \sum_{j=1}^{\infty}\varphi^{j-1}\beta^j p_{t,t+j}\end{aligned}$$

The lower are the probabilities of future survival $p_{t,t+j}$, the less relevant are fluctuations in the interest rate to consumption and portfolio choices. For reference, note that, for infinitely lived agents ($m_t = 0$ for all t), this converges to $\varrho_r = \rho_c/(1-\varphi\rho_c)$, the result from Campbell and Viceira (2001).

Collecting the remaining constant terms implies a difference equation for ϱ_{0t}:

$$\varrho_{0,t+1} = \rho_c(m_t)^{-1}\varrho_{0t} - q_{0t}$$

for the deterministic constant

$$
\begin{aligned}
q_{0t} \equiv \varrho_{r,t+1}(1-\varphi)\bar{r}_f + \log\rho_c(m_t) + (1-\psi)^{-1}(\rho_c(m_t)^{-1}-1)\log(1-\rho_c(m_t)) \\
+ d_0 + d_1\varrho_{rt} - d_2\varrho_{rt}^2 + \frac{1}{2}(\gamma-1)(g_0+g_1\varrho_{rt}+g_2\varrho_{rt}^2).
\end{aligned}
$$

Note that q_{0t} converges to a finite constant: $\lim_{m\to 1} q_{0t} = d_0 + (\gamma-1)g_0/2$. We can similarly iterate this expression forward with terminal condition $(1-\psi)\varrho_0 \to 0$ to arrive at a solution:

$$
\begin{aligned}
\varrho_{0t} &= \rho_c(m_t)(q_{0t} + \varrho_{0,t+1}) \\
&= \rho_c(m_t)q_{0t} + \rho_c(m_t)\rho_c(m_{t+1})q_{0,t+1} + \rho_c(m_t)\rho_c(m_{t+1})\rho_c(m_{t+2})q_{0,t+2} + \ldots \\
&= \sum_{j=1}^{\infty}\beta^j p_{t,t+j} q_{0,t+j-1}.
\end{aligned}
$$

This verifies the conjecture.

C.3 Adding labor income and Social Security

We now introduce a deterministic stream of labor income L and, in turn, Social Security taxes T and benefits B. The present values of labor income (human capital) H and Social Security wealth S are as stated in the main text.

As we did with the wealth return above, let us linearize the returns on human capital and Social Security wealth using a continuous-time approximation. For human capital, the log return is

$$
r_{H,t+1} = r_{ft} + \mu_{Ht} + \underbrace{\left(\sum_{j=1}^{t_{\text{ret}}-t}\omega_{jt}^H\left(\frac{\sigma_j}{\sigma_n}\right)\right)}_{\pi_t^H}(r_{n,t+1} - r_{f,t+1})
$$

where

$$
\omega_{jt}^H = \frac{p_{t,t+j}P_{jt}L_{t+j}}{\sum_{j'=1}^{t_{\text{ret}}-t}p_{t,t+j'}P_{j't}L_{t+j'}} = \frac{p_{t,t+j}P_{jt}L_{t+j}}{H_t}
$$

is the value weight of the jth labor-payment, and therefore π^H is a value-weighted rate-sensitivity adjustment. More specifically, π^H represents the percent holdings of n-period bonds implicit in the human capital asset. To see this, note that the rate sensitivity implied by π^H is[16]

$$\pi_t^H \frac{\sigma_n}{\sigma_r} = \sum_{j=1}^{t_{\text{ret}}-t} \omega_{jt}^H \left(\frac{\sigma_j}{\sigma_r}\right) = \sum_{j=1}^{t_{\text{ret}}-t} \omega_{jt}^H \varepsilon_r(P_{jt}) = \varepsilon_r(H_t).$$

In words, the interest-rate sensitivity of a portfolio with π^H percent allocated to the n-period bond has the exact same interest-rate elasticity as the human capital asset.

Identical logic leads us to conclude that the log return on Social Security is

$$r_{S,t+1} = r_{ft} + \mu_{St} + \underbrace{\left(\sum_{j=1}^{\infty} \omega_{jt}^S \left(\frac{\sigma_j}{\sigma_n}\right)\right)}_{\pi_t^S} (r_{n,t+1} - r_{f,t+1}),$$

where the value weights take the form

$$\omega_{jt}^S = \omega_{jt}^B - \omega_{jt}^T = \frac{p_{t,t+j} P_{jt}(B_{t+j} - T_{t+j})}{S_t},$$

the difference between the benefits claim and the tax liability.

Now, as in the main text, define total wealth as

$$\overline{W}_t = W_t + (L_t + H_t) + (B_t - T_t + S_t). \tag{C.16}$$

(Recall that H and S do not include their contemporaneous "dividends," so we must add them back in this expression.) Grossing up at the rates of return on these assets implies

$$\overline{W}_{t+1} = (W_t + L_t + B_t - T_t - C_t)R_{W,t+1} + H_t R_{H,t+1} + S_t R_{S,t+1}. \tag{C.17}$$

Multiplying and dividing by $\overline{W}_t - C_t$, we have the dynamic budget constraint

$$\overline{W}_{t+1} = (\overline{W}_t - C_t)R_{\overline{W},t+1}.$$

[16]The last equality is exact in continuous time and true to a first order in discrete time.

where the return on total wealth is

$$
\begin{aligned}
R_{\overline{W},t+1} &= \left(\frac{W_t + L_t + B_t - T_t - C_t}{\overline{W}_t - C_t}\right) R_{W,t+1} + \left(\frac{H_t}{\overline{W}_t - C_t}\right) R_{H,t+1} + \left(\frac{S_t}{\overline{W}_t - C_t}\right) R_{S,t+1} \\
&= \alpha_{Wt} R_{W,t+1} + \alpha_{Ht} R_{H,t+1} + \alpha_{St} R_{S,t+1},
\end{aligned}
$$

and the return on financial wealth R_W is as it was in the original problem.

Using the same linearization technique as before, the log total-wealth return can be approximated as

$$
r_{\overline{w},t+1} = r_{ft} + \bar{\mu}_t + \bar{\pi}_t (r_{n,t+1} - r_{ft}) + \frac{1}{2} \bar{\pi}_t (1 - \bar{\pi}_t) \sigma_n^2,
$$

where

$$
\bar{\mu}_t = \alpha_{Ht} \mu_{Ht} + \alpha_{St} \mu_{St}
$$

is a value-weighted drift term from the intertemporal endowments, and

$$
\bar{\pi}_t = \alpha_{Wt} \pi_t + \alpha_{Ht} \pi_t^H + \alpha_{St} \pi_t^S \tag{C.18}
$$

is the value-weighted average of positions in the long-term bond—that is, the percentage of total wealth invested in the bond. Other than the presence of $\bar{\mu}$, this budget constraint is identical in form to that of the problem with no labor income or Social Security. Following the same steps from before, we conclude that

$$
\bar{\pi}_t = \pi_t^*,
$$

where π_t^* is the optimal solution without intertemporal income. Substituting this into (C.18) and rearranging, we see that the optimal allocation to the asset from financial wealth is

$$
\pi_t = \pi_t^* + \left(\frac{H_t}{W_t + L_t + B_t - T_t - C_t}\right) (\pi_t^* - \pi_t^H) + \left(\frac{S_t}{W_t + L_t + B_t - T_t - C_t}\right) (\pi_t^* - \pi_t^S).
$$

In the main text, we slightly simplify notation by redefining wealth to include the contemporaneous income and consumption flows (thus far, we have assumed that it *excludes* these components). Doing this gives us the final expression (29).

C.4 Solution to the value function

This section solves for the value function under optimal policies in closed form. Under recursive preferences, the transformed value function (32) is implicitly defined by the aggregator

$$U_t = \left[(1-\beta)C_t^{1-1/\psi} + \beta(1-m_t)\mathbb{E}_t\left[U_{t+1}^{1-\gamma}\right]^{\frac{1-1/\psi}{1-\gamma}}\right]^{\frac{1}{1-1/\psi}}.$$

By Euler's Theorem,

$$U_t = \frac{\partial U_t}{\partial C_t}C_t + \mathbb{E}_t\left[\frac{\partial U_t}{\partial U_{t+1}}U_{t+1}\right],$$

where

$$\frac{\partial U_t}{\partial C_t} = (1-\beta)\left(\frac{U_t}{C_t}\right)^{1/\psi}$$

and

$$\frac{\partial U_t}{\partial U_{t+1}} = \beta(1-m_t)U_t^{1/\psi}\mathbb{E}_t\left[U_{t+1}^{1-\gamma}\right]^{\frac{\gamma-1/\psi}{1-\gamma}}U_{t+1}^{-\gamma}.$$

Noting that the stochastic discount factor

$$M_{t+1} = \beta(1-m_t)\left(\frac{C_{t+1}}{C_t}\right)^{-1/\psi}\left[\frac{U_{t+1}}{\mathbb{E}_t\left[U_{t+1}^{1-\gamma}\right]^{1/(1-\gamma)}}\right]^{-(\gamma-1/\psi)},$$

and that

$$\frac{(\partial U_t/\partial U_{t+1})(\partial U_{t+1}/\partial C_{t+1})}{\partial U_t/\partial C_t} = M_{t+1},$$

we have that

$$\frac{U_t}{\partial U_t/\partial C_t} = C_t + \mathbb{E}_t\left[M_{t+1}\frac{U_{t+1}}{\partial U_{t+1}/\partial C_{t+1}}\right].$$

Iterating this recursion forward yields

$$\frac{U_t}{\partial U_t/\partial C_t} = \sum_{j=0}^{\infty}\mathbb{E}_t\left[M_{t+j}C_{t+j}\right] = \overline{W}_t,$$

since total wealth is the present value of consumption. Substituting the expression for $\partial U_t/\partial C_t$ and noting that consumption is at an optimum ($C_t = C_t^*$), we get the solution (33).

C.5 Optimal consumption plan in the limit

This section derives the optimal consumption-investment strategy in the limit as risk aversion approaches infinity and the EIS approaches zero. To do so, it is easiest to begin with the first-order condition of a power-utility investor:

$$1 = \mathbb{E}_t \left[\beta(1 - m_t) \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} R_{j,t+1} \right]. \tag{C.19}$$

Conjecture that the optimal consumption policy is a deterministic constant $C_t = \bar{C}_t$. Substituting this conjecture into the first-order condition implies the recursion

$$\bar{C}_t = \left(\beta(1 - m_t)\mathbb{E}_t[R_{j,t+1}] \right)^{-1/\gamma} \bar{C}_{t+1}. \tag{C.20}$$

Now taking the limit as $\gamma \to \infty$ implies that $\bar{C}_t = \bar{C}_{t+1} = \bar{C}$, meaning that consumption is indeed deterministic and in fact time-invariant.

The present value of optimal consumption must equal total wealth, so we have

$$\overline{W}_t = \bar{C} \sum_{j=0}^{t_{\max}-t} P_{jt}, \tag{C.21}$$

where $t_{\max}$ is the first year in which $m_t = 1$.[17] This expression pins down the value of $\bar{C}$. Because the optimal consumption plan is deterministic and constant, the agent finances it by purchasing $\bar{C}$ of each zero-coupon bond and consuming the coupons.

Finally, we wish to relate the optimal portfolio strategy financing this consumption plan to the optimal policy $\bar{\pi}$ derived above. First, using the same linearization technique as above, notice that the wealth return under this consumption policy equals

$$r_{w,t+1} = r_{ft} + \underbrace{\left(\sum_{j=1}^{t_{\max}-t} \frac{P_{jt}}{\sum_{j'=1}^{t_{\max}-t} P_{j't}} \left(\frac{\sigma_j}{\sigma_n} \right) \right)}_{\tilde{\pi}_t} (r_{n,t+1} - r_{ft}). \tag{C.22}$$

As with human capital and Social Security wealth, $\tilde{\pi}$ represents an implicit holding of n-period bonds from the annuity financing consumption. Now let us compare this implicit holding $\tilde{\pi}$ to the

[17] Note that this satisfies the terminal condition $\overline{W}_{t_{\max}} = \bar{C}$, since $P_0 = 1$.

optimal holding π^*. In the limit, the general expression for optimal consumption (23) implies the (negative) elasticity

$$\frac{\partial \log(\bar{C}/W_t)}{\partial r_{ft}} = \varrho_{rt}.$$

Calculating this same left-hand-side derivative from (C.21) and equating these, we have

$$\varrho_{rt} = \sum_{j=0}^{t_{\max}-t} \frac{P_{jt}}{\sum_{j'=0}^{t_{\max}-t} P_{j't}} \left(\frac{\sigma_j}{\sigma_r}\right).$$

Substituting this into the expression for the optimal portfolio $\bar{\pi} = \pi^*$ in (25), then taking $\gamma \to \infty$, we have

$$\bar{\pi}_t = \varrho_{rt}\frac{\sigma_r}{\sigma_n} = \sum_{j=0}^{t_{\max}-t} \frac{P_{jt}}{\sum_{j'=0}^{t_{\max}-t} P_{j't}} \left(\frac{\sigma_j}{\sigma_n}\right).$$

This optimal policy is exactly identical to the expression $\tilde{\pi}$ from (C.22), as claimed.

D Additional simulation results

D.1 Higher persistence of interest-rate shocks

Figure D.2: Life-cycle profiles of wealth and its interest-rate sensitivity for $\varphi = .9756$

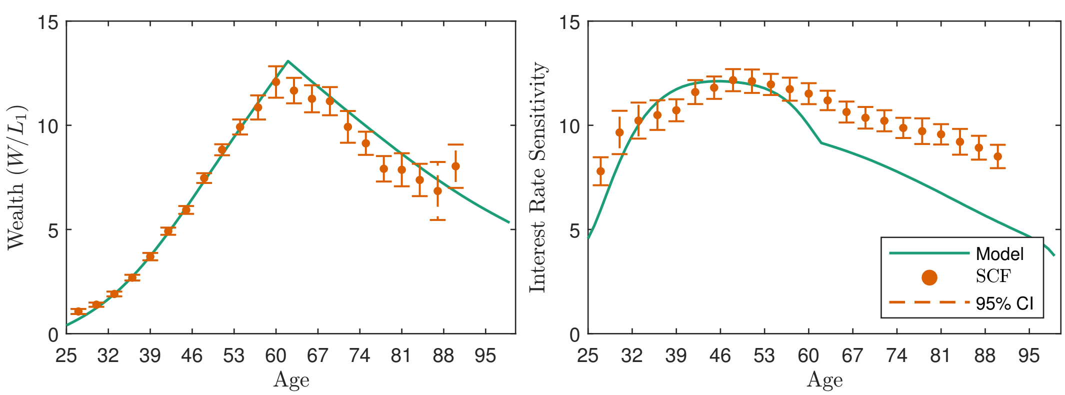

Note: This figure reports the evolution of market wealth and its sensitivity to interest rates over the life cycle in our benchmark calibration and in the SCF. In the data, wealth is computed per adult, including deceased spouses, and scaled by the Social Security wage index. 95% confidence intervals represent $\pm$ 1.96 standard errors, clustered by cohort.

Figure D.3: Interest-rate sensitivity of wealth at age 40–45 for $\varphi = .9756$

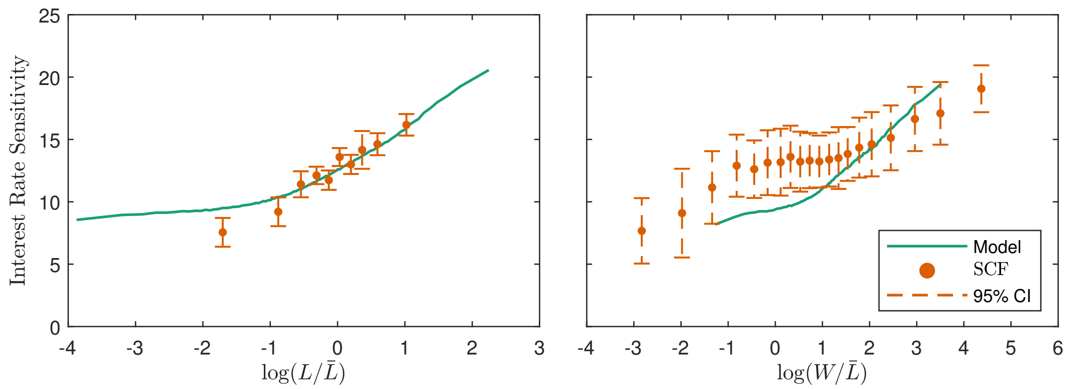

Note: This figure reports the relationship between the interest-rate sensitivity of wealth and wealth (left panel) and earnings (right panel). In the data, wealth and earnings are computed per adult and scaled by the Social Security wage index. In the left panel, each bin represents a decile of earnings. In the right panel, each bin represents 5% of observations, except for the four wealthiest bins which represent 2.5% each. Simulated data report the average interest-rate sensitivity per centile of wealth and earnings, respectively. 95% confidence intervals represent $\pm$ 1.96 standard errors, clustered by cohort.

D.2 Mechanism

The quantitative model validates our economic intuition and allows us to study counterfactuals. To shed more light on the full numerical model, this section analyzes the importance of two novel mechanisms in our model: income-based differences in mortality rates and the presence of Social Security. Figure D.4 plots quantities of interest with and without these features.

Figure D.4: Effect of Social Security and differences in life expectancy

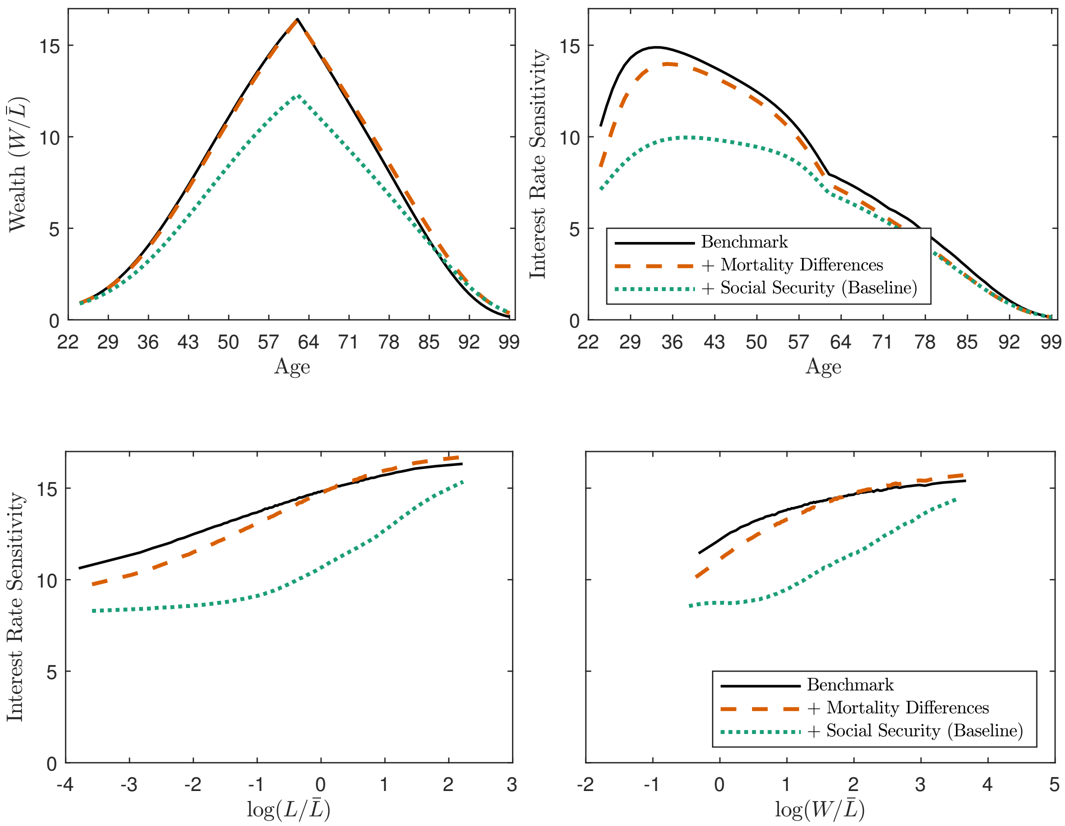

Note: This figure shows the effects of mortality differences and Social Security on life-cycle wealth accumulation and the interest-rate sensitivity of wealth in the model. In the benchmark case, mortality probabilities are constant within an age cohort and there are no Social Security taxes or benefits. Mortality differences are based on lifetime earnings (AIYE). Where relevant, wealth W and income L are scaled by the Social Security wage index $\bar{L}$.

Mortality affects the optimal interest-rate sensitivity through two channels. First, higher mortality rates reduce the value of human capital relative to financial wealth, diminishing its substitution effect. Second, higher morality reduces rate exposure because agents discount the future more. The distributional consequences of this effect are revealed by the bottom two panels of Figure D.4. The income-based adjustment to mortality rates applies mostly to low-income households; the adjustment is small for households with average and high income. As a result, the optimal rate exposure falls noticeably for low earners but does not change much for other households. This means that the average life-cycle path of rate exposure, shown in the top right panel of Figure D.4, tends to be lower in levels than in the benchmark without intracohort mortality differences. Perhaps surprisingly, the overall quantitative effect of mortality differences on most of the cross-section is minimal.

The effect of Social Security is more substantial. The existence of Social Security taxes and benefits leads to less accumulation of financial wealth over the life cycle, because taxes reduce disposable income and benefits crowd out the need to save. Social Security also flattens the "hump" in rate exposure during working life but has little effect in retirement, consistent with the economic intuition discussed in Section 4. Finally, Social Security steepens the relation of rate sensitivity with wealth and income. This, too, is exactly as predicted by the analysis in Section 4.

E Numerical appendix

Table E.1: Calibration of labor income process

Parameter estimates for Section 3.3 come from Specifications (5) in Guvenen et al. (2022). Parameters can be found in Table IV of the published version and Table D.III of the Online Appendix. We also combine the z and α processes, which results in the $\sigma_{z,0}$ parameter listed below. We do this to avoid adding an additional state variable to the model, a decision that has little effect on the results as the z process is extremely persistent. Finally, note the deterministic life-cycle component is given by $g(\text{Age}) = b_{0,g} + b_{1,g}\text{Age} + b_{2,g}\text{Age}^2/10 + b_{3,g}\text{Age}^3/100$ where $b_{0,g}$ is specified to make mean earnings equal to Social Security Wage Index.

Parameter	Calibration
ρ	0.991
p_z	17.6%
$\mu_{\eta,1}$	-0.524
$\sigma_{\eta,1}$	0.113
$\sigma_{\eta,2}$	0.046
$\sigma_{z,0}$	0.652
λ	0.016
p_ε	4.4%
$\mu_{\varepsilon,1}$	0.134
$\sigma_{\varepsilon,1}$	0.762
$\sigma_{\varepsilon,2}$	0.055
$a_\nu \cdot 1$	-2.495
$b_\nu \cdot t$	-1.037
$c_\nu \cdot z_t$	-5.051
$d_\nu \cdot t \cdot z_t$	-1.087
$b_{0,g}$	-6.142
$b_{1,g}$	0.3040
$b_{2,g}$	-0.051
$b_{3,g}$	0.002586