Who hedges interest-rate risk? Implications for wealth inequality*

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

March 1, 2023

Abstract

Falling interest rates increase wealth inequality by raising the market value of long-duration assets held by wealthy households. To understand this phenomenon, we present a life-cycle model in which households can invest in short- or long-term assets to hedge against interest-rate risk. Our model matches important stylized facts. First, the share of long-term assets in households' wealth is hump-shaped over the life-cycle. Within cohorts, it increases with wealth and earnings. Second, wealth inequality grows when interest rates fall, but only when wealth does not include the value of Social Security. Hedging demand against interest-rate risk can explain 40% of long-run changes in wealth inequality since 1960.

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, and Columbia University for helpful comments.

†University of Pennsylvania, Wharton, Finance department. Email: scath@wharton.upenn.edu
‡University of Pennsylvania, Wharton, Finance department. Email: maxmil@wharton.upenn.edu.
§University of Pennsylvania, Wharton, Finance department. Email: jparon@wharton.upenn.edu
¶Yale Law School. Email: natasha.sarin@yale.edu

1 Introduction

When households invest in different assets, and those held by the rich deliver higher returns, portfolio choices contribute to trends in inequality (Fagereng et al., 2020; Hubmer et al., 2021). In particular, long-term assets delivered exceptional returns in the last five decades (Binsbergen, 2021). Because wealthier households invest more in these assets, these exceptional returns increased private wealth inequality (Greenwald et al., 2021). But why do some households invest more in long-term assets than others?

To answer this question, we present a life-cycle model that provides theoretical foundations to the cross-section of long-term asset holdings. Long-term assets are a desirable hedge for those who are reliant on the rate of return on private savings, because they generate large capital gains when rates fall. These gains compensate for the deterioration of investment opportunities. We show that there are predictable differences in the dependence of households on private savings that drive certain types of people to be more heavily invested in long-duration assets.

Households need to hedge against declining interest rates to the extent that their lifetime consumption depends on the rate of return on private savings. The need to hedge is lower for three classes of households. First, older households who have shorter investment horizons, making rate decreases less meaningful in the presence of fewer periods of compounding. Second, younger households who are already implicitly invested in a long-duration asset through their own human capital. Third, low earners whose primary savings vehicle is Social Security, because the rate of return on Social Security contributions is not impacted by changes in interest rates. These predictions are consistent with the cross-sectional relationships we document between wealth, age, earnings, and households' propensity to invest in long-term assets in the Survey of Consumer Finances (SCF).

Our primary contribution is to build a life-cycle model with uninsurable income risk, a Social Security system, bequests, differences in life expectancy, and stochastic interest rates. Households can choose how sensitive their wealth is to changes in interest rates by mixing two assets: a short- and a long-term bond. The finitely lived agents change their portfolio choices as they age; and the presence of human capital (itself a long-term asset) and Social Security give heterogeneous households different incentives to protect against interest-rate risk.

We then use an overlapping generations (OLG) version of our model to study the implications for the dynamics of wealth inequality. Heterogeneity in sensitivity to interest-rate changes drives a significant portion of historical trends in private wealth inequality. Overall, in our model, 40% of the long-run variations in U.S. wealth inequality since the 1960s can be explained by these differences in optimal portfolio choices that lead middle-aged, higher-earning households to be more heavily tilted toward longer-duration assets. In the presence of a long-run rise in interest rates, such as between 1960 and 1985, these differences drive down private wealth inequality; in the presence of a long-run decline in interest rates, such as since 1985, the same phenomenon drives wealth inequality up.

However, we also show that, in our model, differences in households' exposure to interest rates within an age cohort vanish when we extend the concept of wealth to include Social Security and human capital. Said another way, the hedging motive that drives households to invest in long-term assets is less pronounced for households who already implicitly hold long-term assets through human capital and those who prepare retirement primarily through Social Security, which is not exposed to interest-rate risk.

The design of Social Security magnifies these differences. Because low earners have higher Social Security replacement rates than high earners, they need to save less on their own and invest a smaller share of their private wealth in long-term assets. As a consequence, Social Security increases not only the correlations between earnings and private wealth accumulation, but also the correlation of wealth and portfolio returns when interest rates fall. Importantly, because these mechanisms are driven by substitutions effects, they vanish once we consider a wealth concept that includes the present value of Social Security cash flows. That is why the increase in private wealth inequality is significantly attenuated by the inclusion of Social Security in the measure of wealth (Catherine et al., 2020).

For them, interest rates represent the price of transferring resources from periods in which their income exceeds their consumption to times when the opposite is true. The uncertainty regarding the evolution of these prices can be covered by investing their wealth in a portfolio of zero-coupons that replicates the differences between their consumption plan and future income over their life cycle. Once bought, this portfolio will deliver the payoffs they need regardless of the evolution of interest rates.

In reality, households engage in voluntary or mandatory financial transactions that emulate this abstract portfolio. At the bottom of the earnings distribution, workers save primarily through Social Security and the rate of return on their contributions does not depend on the evolution of interest rates. Lower middle-class households 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 their non-residential consumption in retirement. Finally, because they receive low replacement rates from Social Security, higher earners rely more on private savings, such as retirement accounts invested in the stock market. Since stocks are a high-duration asset, they protect retirement consumption because their market value goes up when rates fall.

The differences in hedging behavior relate to observed differences in private wealth inequality, which in the U.S. fell from the 1960s to the 1980s and has steadily risen since, while real short-term rates have followed the opposite trend. Using the historical time series of real interest rates, our OLG model can explain 40% of the decline in the top 10% wealth share between 1960 and 1985 and of its subsequent rise between 1985 and 2019.

Our paper extends past literature by providing a theoretical foundation for 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 et al. (2019), Hubmer et al. (2021) and Fagereng et al. (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 are more likely to invest in long-term assets.

In explaining these facts, we complement a recent strand of the literature that studies why wealthy households invest more in the stock market. This fact can be explained by decreasing relative risk aversion (Meeuwis, 2022), the crowding-out effect of housing (Cioffi, 2021), or the exposure of less wealthy households to counter-cyclical consumption risk (Catherine, 2021; Catherine et al., 2021). Unlike these studies, we focus our attention on the trade-off between short- and long-term assets, because stocks have not outperformed government bonds of similar duration over the last four decades (Binsbergen, 2021).

Auclert (2019) argues that falling rates redistribute wealth towards investors holding long-term assets; Greenwald et al. (2021) show that, empirically, these investors tend to be wealthier. We complement these papers by providing micro-foundations to the cross-section of households' propensity to invest in long-term assets. Kumhof et al. (2015) and Mian et al. (2021) suggest the reverse causal relationship: wealth inequality caused interest rates to decline because the rich have higher savings rates. Even if this is the case, the driver of the decline in interest rates is not relevant to our underlying contribution: as the interest-rate environment shifts, differences in portfolio choice between households mechanically lead to marked changes in private wealth inequality.

We also contribute to a longstanding literature on household portfolio choices. This literature has almost entirely studied how households allocate wealth between the stock market portfolio and a riskfree asset (Benzoni et al., 2007; Catherine, 2021; Cocco et al., 2005; Merton, 1969; Viceira, 2001). Campbell and Viceira (2001) studies a choice between a short- and a long-term zero-coupon bond and show that the desire to hedge rate fluctuations can increase the demand for long-term bonds. However, their model does not account for the effects of the life-cycle or the substitution effects from human capital and Social Security. We show that these effects are essential to explaining the data.

2 Stylized facts

This section discuss our empirical measurement of interest-rate sensitivity and present four 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. the time series of wealth inequality follows the decline in interest rates

2.1 Measuring interest-rate sensitivity

Asset interest-rate sensitivity The interest-rate sensitivity measures how the price of an asset responds to changes in interest rates. Given the difficulty of measuring this elasticity directly in the data, a widely used approximation is an asset's cash-flow duration, the value-weighted timing of its cashflows. This is due to the fact that duration is equal to rate sensitivity when shocks to rates are permanent.

In reality, shocks to interest rates are not permanent and exhibit some level of mean reversion. Under our assumption that the riskfree rate follows an AR(1) process with persistence φ, an asset's actual interest-rate sensitivity is a concave transformation of its cash-flow duration, approximately given by

$$\hat{\varepsilon}(\text{Asset}, r_f) = \frac{1 - \varphi^{\text{dur(Asset)}}}{1 - \varphi}. \tag{1}$$

Note, this result assumes that the interest-rate sensitivity of the asset is the same as the interest-rate sensitivity of a riskfree zero-coupon bond with the same duration.[1]

Portfolio interest-rate sensitivity With the rate sensitivity of each asset adjusted to account for the persistence of interest-rate shocks, we compute the rate sensitivity of the overall wealth portfolio for household i as the value-weighted sum of each component elasticity:

$$\hat{\varepsilon}(\text{Wealth}_i, r_f) = \sum_j \frac{\text{Asset}_{ji}}{|\text{Wealth}_i|} \times \hat{\varepsilon}(\text{Asset}_{ji}, r_f), \tag{2}$$

where Asset_{ji} denotes the value of the asset or debt, Wealth_i denotes the value of all assets the household holds less debts, and $\hat{\varepsilon}(\cdot, r_f)$ is the estimated interest-rate elasticity of that asset or debt.

Interest-rate sensitivity in the SCF To estimate interest-rate sensitivity in the cross-section of households in the U.S., we apply a three-step process based on the framework above to data on household portfolios, income, and wealth from the

[1]The derivation of expression (1) is presented in Section 3.2.

triennial Survey of Consumer Finances. First, we estimate cash-flow duration for each asset and liability on household balance sheets. Second, we apply our concave transformation to the component duration estimates to obtain the interest-rate sensitivity of each asset and liability. Third, we sum their value-weighted sensitivities to arrive at the overall portfolio interest-rate sensitivity of each household.

We adopt different methods to compute the cash-flow durations of assets and liabilities for the first step. We start by grouping assets into equity, real estate, liquid assets, fixed income, vehicles, and private business wealth, and grouping debts into mortgage debt and other debt. For all equity, real estate, and liquid assets, we apply average group-wide duration estimates provided by Greenwald et al. (2021). For fixed-income assets, we collect yearly 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.

The duration of vehicles is determined using the age of each vehicle provided in the SCF, an estimate of its maximum lifetime, and an assumed constant depreciation rate. The duration of private business wealth is estimated from an aggregate price-dividend ratio of businesses owned by households in the SCF, where price is the value of the household's ownership share and dividends are the household's business income less wages. In addition, to account for the possibility that households of different ages and wealth may hold equity and private business assets of differing duration, we adjust the aggregate duration estimates by age and wealth for these asset classes using valuation ratios implied by the SCF. For more detailed information on each component duration calculation, see Appendix A.2.

For liabilities, we assume a fixed repayment schedule and estimate duration as

$$\text{dur}(\text{Debt}) = \sum_{n=1}^{N} \left(\frac{e^{-ny_{nt}}}{\sum_{n'=1}^{N} e^{-n'y_{n't}}} \right) n, \tag{3}$$

where N is the number of years remaining on the loan, provided in the SCF, and y_{nt} is the riskfree spot rate at horizon n in year t which we obtain from the nominal yield curve from the Fed less SSA inflation projections. More information about our construction of the debt duration can be found in Appendix A.2. Table 1 presents the averages of our duration estimates by asset group for households in the SCF.

Table 1: Average duration and portfolio share by asset group

	Duration	Portfolio share (equal-weighted)	Portfolio share (value-weighted)
Assets			
Private Business	45.5	0.04	0.18
Equity	36.3	0.09	0.19
Real Estate	16.5	0.44	0.40
Fixed Income	5.40	0.08	0.12
Vehicles	2.61	0.20	0.03
Cash and Deposits	0.25	0.15	0.08
Liabilities			
Mortgage Debt	10.3	0.53	0.83
Other Debt	2.70	0.47	0.17

Note: This table reports the average duration and portfolio share of each asset group for households in the SCF. 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 value-weighted portfolio shares, but in this case we weight each household by both its SCF sample weight and its networth.

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

The first stylized fact is that the rate sensitivity of household wealth portfolios is hump-shaped over the life-cycle: it is lowest for 20-year-olds, rises to a high for

Figure 1: Interest-rate sensitivity of wealth by age

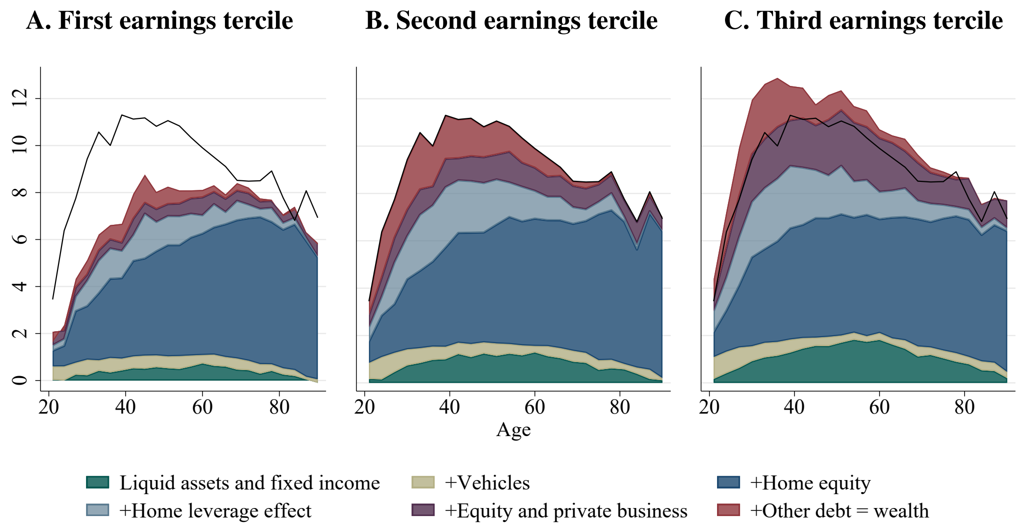

Note: This figure reports the interest-rate sensitivity of wealth by age and tercile of earnings. The rate sensitivity is decomposed into the contribution of six components of wealth. From bottom to top, we calculate the sensitivity of partial portfolios, adding components step-by-step. First, we report 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 $\hat{\varepsilon}(\text{Portfolio}_k, r_f) = \frac{\text{Portfolio}_{k-1}}{\text{Portfolio}_k}\hat{\varepsilon}(\text{Portfolio}_{k-1}, r_f) + \frac{\text{Component}_k}{\text{Portfolio}_k}\hat{\varepsilon}(\text{Component}_k, r_f)$.

40- to 45-year-olds, and steadily declines thereafter. Figure 1 decomposes this pattern clearly, showing 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.

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 average rate sensitivity of assets is nearly double 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.

2.3 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% larger capital gains than those of the bottom earnings tercile. In the model presented below, earnings are the primary source of heterogeneity within a cohort, with higher-income households holding portfolios with higher interest-rate sensitivities. In the data, the main difference between the bottom tercile (Panel A) and the middle tercile (Panel B) is in home equity, while the main difference in rate sensitivity between the top (Panel C) and middle tercile is due to differences in equity holdings.

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

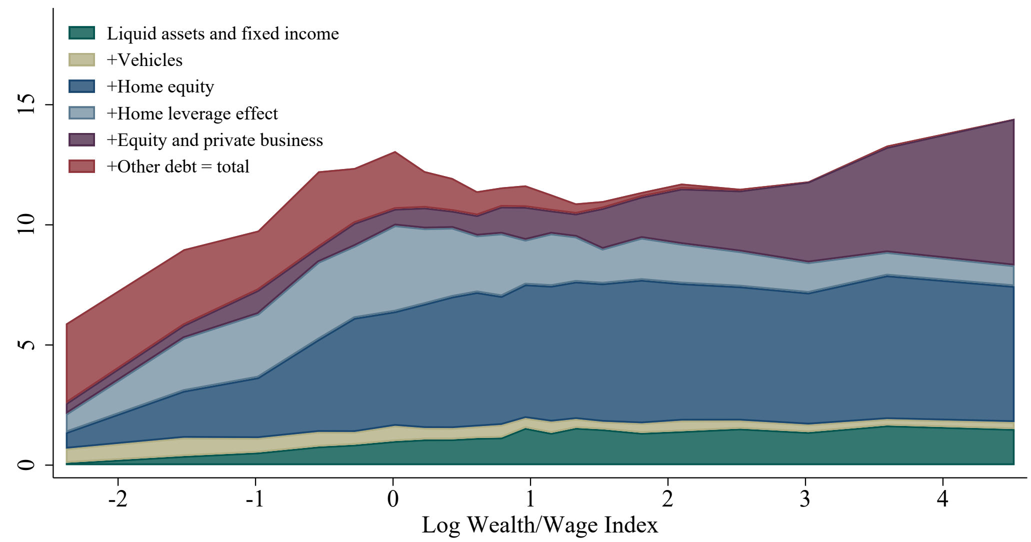

Note: This figure decomposes the interest-rate sensitivity for households in which the head of the household is between 40 and 45. The methodology is the same as in Figure 1, except that here the x-axis is the log of wealth scaled by the Social Security Wage Index in the survey year.

2.4 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, the assets that contribute most to the interest-rate sensitivities of their portfolios are liquid accounts, vehicles, and non-mortgage debt. 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 requires lower-middle-wealth homebuyers to take out large mortgages and excessively 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.

2.5 Wealth inequality follows interest rates

The fourth and final stylized fact is that, over the last six decades, the wealth share of the top 10% of the wealth distribution has closely tracked the price of a one-year bond. Figure 3 plots one minus our estimated 10-year forward rate.[2] We use 10-year real forward rates here because they represent expectations of where interest rates will be, thus containing information about the persistent component of the interest rate.

3 Model

We model household consumption and investment decisions over a life-cycle which can be 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

$$V_{it} = \max_{\{C_{is},\pi_{is}\}} \mathbb{E}_t \sum_{s=t}^{t_{\max}} \beta^{s-t} \left[(1-m_{is}) \frac{C_{is}^{1-\gamma}}{1-\gamma} + m_{is} b(W_{is}, r_{fs}) \right], \tag{4}$$

[2]For details on the construction of our time series of the short-term real interest rate, see Appendix A.3

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

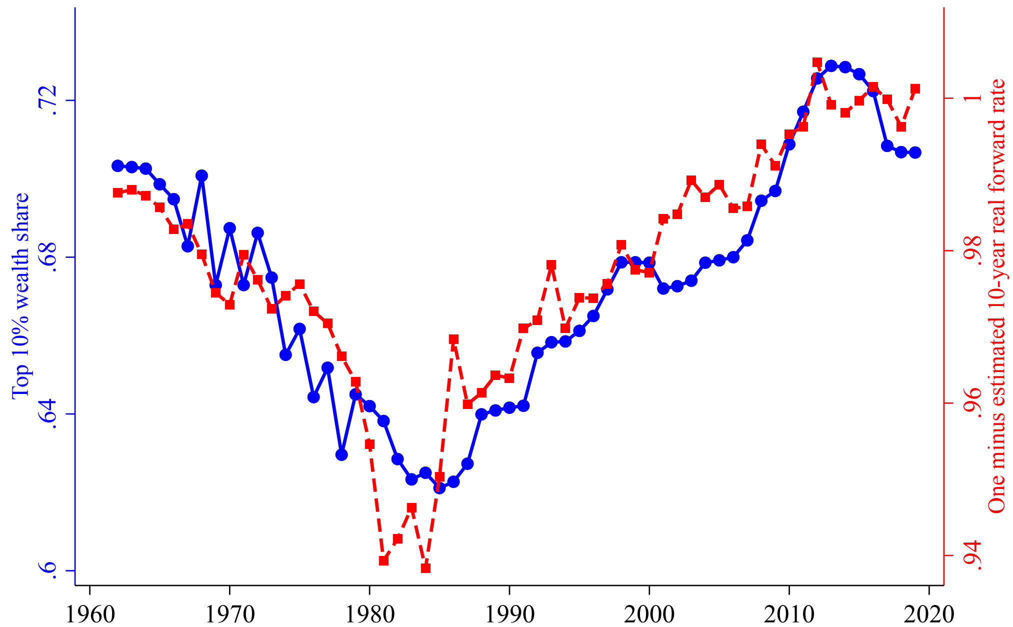

Note: This figure presents the time series of the top 10% wealth share from the World Inequality Database and $1 - \hat{f}_{10,t}$, one minus our estimated 10-year real forward rate from Equation (A.3).

where β is the rate of time preference, γ is the coefficient of relative risk aversion, $t_{\max}$ is the maximum lifespan, m_i is the age- and income-dependent mortality probability, 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

$$W_{i,t+1} = (W_{it} + L_{it} + B_{it} - T_{it} - C_{it}) R_{Wi,t+1}, \tag{5}$$

with gross return on savings

$$R_{Wi,t+1} = R_{ft} + \pi_{it}(R_{n,t+1} - R_{ft}). \tag{6}$$

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 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 $r = \log R$. We assume that the riskfree rate follows a first-order autoregression given by

$$r_{f,t+1} = (1 - \varphi)\bar{r}_f + \varphi r_{ft} + \sigma_r \epsilon_{r,t+1}. \tag{7}$$

We model the long-term asset as a risk-free claim to one unit of real consumption in n periods. Its payoff is riskless in real terms. 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

$$r_{n,t+1} = r_{ft} + \mu_n - \sigma_n \epsilon_{r,t+1}, \tag{8}$$

where sensitivity to rate shocks σ_n is given by

$$\sigma_n = \frac{1 - \varphi^{n-1}}{1 - \varphi} \sigma_r. \tag{9}$$

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.

It is straightforward to show that the rate elasticity of the long-term bond price is

$$\varepsilon(P_{nt}, r_{ft}) = -\frac{\partial \log P_{nt}}{\partial r_{ft}} = \frac{1 - \varphi^n}{1 - \varphi}. \tag{10}$$

This sensitivity is increasing in maturity n. We see immediately the effect of unexpected changes in interest rates: if the riskfree rate unexpectedly falls, then 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 (6).

We purposely choose not to model the actual assets–housing, bonds, stocks– and liabilities that actually constitute most of households' balance sheets for two reasons. First, to make the economic intuition of the model as transparent as we can. 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 has 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

$$\tilde{L}_{it} = \exp\left\{g(t) + \zeta_{i0} + z_{it} + \epsilon_{it}\right\}. \tag{11}$$

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

$$z_{it} = \rho z_{i,t-1} + \eta_{it}, \tag{12}$$

with innovations η_i drawn from a mixture of normal distributions

$$\eta_{it} \sim \begin{cases} \mathcal{N}(\mu_{\eta 1}, \sigma^2_{\eta 1}) & \text{with probability } p_z, \\ \mathcal{N}(\mu_{\eta 2}, \sigma^2_{\eta 2}) & \text{with probability } 1 - p_z. \end{cases} \tag{13}$$

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

$$\epsilon_{it} \sim \begin{cases} \mathcal{N}(\mu_{\epsilon 1}, \sigma^2_{\epsilon 1}) & \text{with probability } p_\epsilon, \\ \mathcal{N}(\mu_{\epsilon 2}, \sigma^2_{\epsilon 2}) & \text{with probability } 1 - p_\epsilon. \end{cases} \tag{14}$$

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:

$$T_{it} = 0.106 \min\{L_{it}, 2.5\bar{L}_t\}. \tag{15}$$

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

$$L_{it}^{\text{indexed}} = \min\{L_{it}, 2.5\bar{L}_t\} \frac{\bar{L}_{t_{60}}}{\bar{L}_t} \tag{16}$$

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:

$$B_{it} = \begin{cases} 0.9\text{AIYE}_{it_{\text{ret}}} & \text{if AIYE}_{it_{\text{ret}}} < b_1, \\ 0.9b_1 + 0.32(\text{AIYE}_{it_{\text{ret}}} - b_1) & \text{if } b_1 \leq \text{AIYE}_{it_{\text{ret}}} < b_2, \\ 0.9b_1 + 0.32(b_2 - b_1) + 0.15(\text{AIYE}_{it_{\text{ret}}} - b_2) & \text{if } b_2 \leq \text{AIYE}_{it_{\text{ret}}}. \end{cases} \tag{17}$$

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:

$$\text{AIYE}_{it} = \sum_{s=t_0}^{t} \min\{L_{is}, 2.5\bar{L}_s\}\frac{\bar{L}_t}{\bar{L}_s} = \bar{L}_t \sum_{s=t_0}^{t} \min\{\tilde{L}_{is}, 2.5\}. \tag{18}$$

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

$$b(W_{it}, r_{ft}) = \bar{b}\frac{\bar{C}_{it}^{1-\gamma}}{1-\gamma}, \tag{19}$$

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:

$$W_{it} = \bar{C}_{it} \sum_{k=0}^{\bar{b}} P_{kt}. \tag{20}$$

4 Economic intuition

To communicate the first-order intuition of our model, we first present an analytical solution to a linearized version with no idiosyncratic income risk or bequests.[3]

4.1 Optimal choices without labor income

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

$$\frac{C_{it}^{*}}{W_{it}} = \underbrace{(1-\beta(1-m_{it}))}_{\text{time discounting}} \times \underbrace{\exp\left\{\left(1-\frac{1}{\gamma}\right)(\varrho_{0t}+\varrho_{rt}r_{ft})\right\}}_{\text{income and substitution effects}}. \tag{21}$$

The first term represents the positive effect of impatience and mortality on consumption. The second term represents the net effect of the 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 ϱ_{rt}, which is positive and declining in age.[4] Because they depend less on future rates of return, older households' consumption reacts less to changes in interest rates.

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

[4] We use the shorthand notation ϱ_{0t} and ϱ_{rt} for $\varrho_0(\{m_{is}\}_{s\geq t})$ and $\varrho_r(\{m_{is}\}_{s\geq t})$, respectively. Both quantities approach zero as $m_{it} \to 1$: agents approaching the end of life consume everything.

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

$$\pi_{it}^{*}=\underbrace{\frac{1}{\gamma}\frac{\mu_n+\frac{1}{2}\sigma_n^2}{\sigma_n^2}}_{\text{myopic demand}}+\underbrace{\left(1-\frac{1}{\gamma}\right)\varrho_{rt}\left(\frac{1-\varphi^{n-1}}{1-\varphi}\right)^{-1}}_{\text{duration/hedging demand}}. \tag{22}$$

The first term represents the traditional risk-return tradeoff of Merton (1969). 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. The sensitivity of consumption to rate shocks declines with the investor's horizon, so the hedging demand decreases in age toward zero with the coefficient ϱ_{rt}. Therefore, absent labor income and Social Security, the rate exposure of households' portfolios should decline over the lifecycle.

Interestingly, for a $\gamma = 0$, 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 short-sell the long-term asset if he can. 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 Adding 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

[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.

capital H_{it} and Social Security wealth S_{it} are

$$H_{it} = \sum_{k=1}^{t_{\text{ret}}-t} \left[\prod_{s=1}^{k} (1 - m_{i,t+s}) \right] P_{kt} L_{i,t+k}, \tag{23}$$

and

$$S_{it} = \sum_{k=1}^{t_{\max}} \left[\prod_{s=1}^{k} (1 - m_{i,t+s}) \right] P_{kt} (B_{i,t+k} - T_{i,t+k}), \tag{24}$$

where $\prod_{s=1}^{k}(1 - m_{i,t+s})$ is the cumulative probability of surviving from t to $t + k$ and 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 (21). Similarly, the optimal allocation to bonds out of total wealth is $\bar{\pi}_i = \pi_i^*$ from (22). The optimal allocation out of financial wealth W then takes the form

$$\pi_{it} = \pi_{it}^* + (\pi_{it}^* - \pi_{it}^H) \frac{H_{it}}{W_{it}} + (\pi_{it}^* - \pi_{it}^S) \frac{S_{it}}{W_{it}}. \tag{25}$$

The endowments of human capital and Social Security wealth are implicit holdings of long-term assets, and thus substitutes 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 4: Effect of labor income and Social Security on long-term asset share

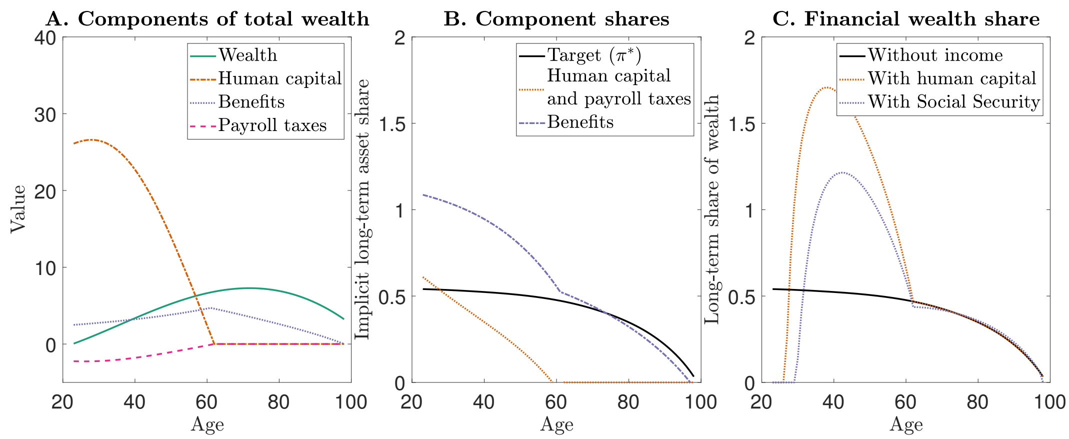

Note: This figure shows a representative path of total wealth components, their duration, and their effect on wealth allocations over the life-cycle. Panel A plots the average values of each component of total wealth, defined as the sum of wealth and the present values of labor income (human capital) and Social Security taxes and benefits. Panel B shows the implicit share of each component in the n-period bond. Panel C illustrates the incremental effect of each component on the financial-wealth allocation to the long-term bond.

Figure 4 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 hold mostly short-term bonds. 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 5: Wealth-duration relation with income and Social Security

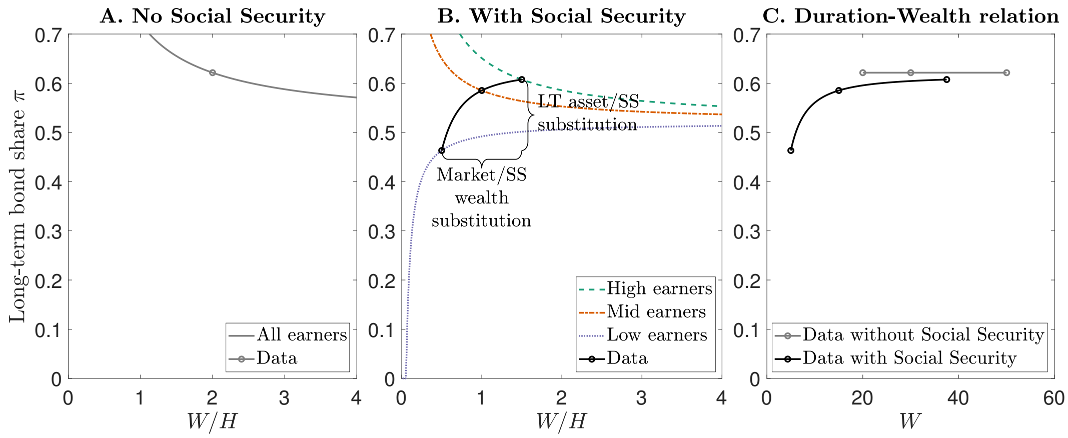

Note: This figure illustrates the effect of Social Security on intra-cohort allocations to the long-term asset. Panel A plots the optimal long-term bond share as a function of the ratio of wealth W to human capital H when there is no Social Security. The round marker represents the ratio W/H observed in the data. 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.

In addition to these effects, the progressivity of Social Security implies that households with lower earnings will hold less rate-sensitive portfolios, even after controlling for wealth and income. Figure 5 illustrates this prediction. Without Social Security, wealth-income ratios and portfolio allocations within a cohort show little variation. But because Social Security yields higher replacement rates for low-earning households, it has a larger effect on their portfolios. This arises out of two compounding effects. First, fixing wealth-income ratios, low earners have more Social Security per dollar of income, which they offset by decreasing financial-wealth duration. Second, the comparatively large endowment of Social Security reduces the savings rates of low earners, reducing their wealth-income ratios. As the third panel of the figure shows, these effects combine to generate a steep positive relation between wealth and duration, just as in the data.

4.3 Implications for wealth inequality

Recall that an individual's financial wealth evolves according to

$$\frac{W_{i,t+1}}{W_{it}} = \underbrace{\left(1 - \frac{C_{it} - Y_{it}}{W_{it}}\right)}_{\text{savings}} \underbrace{R_{Wi,t+1}}_{\text{portfolio}}, \tag{26}$$

where Y_i is the sum of labor income, taxes, and benefits. There are thus two channels through which the interest rate can affect inequality: consumption-wealth ratios (savings) and portfolio allocations (returns).

We can decompose changes in inequality over time by taking logs of (26) and then computing cross-sectional variances. Doing so yields the change in wealth inequality from one period to the next,[6]

$$\begin{aligned} \mathrm{var}_I(w_{i,t+1}) - \mathrm{var}_I(w_{it}) &= \mathrm{var}_I(s_{it}) + \mathrm{var}_I(r_{wi,t+1}) \\ &\quad + 2\mathrm{cov}_I(w_{it}, s_{it}) + 2\mathrm{cov}_I(w_{it}, r_{wi,t+1}) + 2\mathrm{cov}_I(s_{it}, r_{wi,t+1}). \end{aligned} \tag{27}$$

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 the wealthy tend to save more. It also increases if the wealthy experience higher returns. Finally, it increases if households with higher savings rates experience higher returns. Our model reveals why, in the presence of income and Social Security, these covariance channels are all positive when interest rates fall.

First, consider intra-cohort wealth inequality. Absent income and Social Security, there is no variation in savings rates or portfolio choices within a cohort, so inequality is fixed at its initial condition. That is, all variance and covariance

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

terms on the right-hand side of (27) are zero. The presence of income and Social Security, in contrast, generates changing inequality through 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 illustrated in Figure 5, the substitution effects of Social Security create a positive correlation between a household's wealth and its long-term asset share. 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.

The decomposition (27) applies similarly to between-cohort wealth inequality. In particular, differences in human capital endowments between age groups drive the return channel of inequality. Figure 4 illustrates this fact. Following a rate decline, middle-aged households' wealth returns are exceptionally high, so that these cohorts accumulate more wealth than both the young and the old. These households have the highest levels of wealth, so the wealth-return covariance in this instance is positive, driving wealth inequality upward. Unlike in the within-cohort case, the savings-return covariance will remain close to zero, since savings rates are positive for the young, near zero for the middle-aged, and negative for the old.

4.4 Targeting consumption duration

These optimal policies ultimately point to the fact that individuals choose assets to hedge interest-rate risk and finance a smooth consumption plan. When we add intertemporal income, agents simply adjust this trading strategy to target a consumption plan of the same shape — that is, with the same duration. In our framework, agents achieve this by buying and selling zero-coupon bonds with payoffs arriving at approximately the same time as the desired consumption flows.

We can illustrate this intuition most clearly in the limiting case of an infinitely risk-averse investor.[7] 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. Financial wealth is the present value of the excess consumption plan:

$$W_{it} = \sum_{k=1}^{t_{\max}} P_{kt}(\bar{C}_i - Y_{i,t+k}). \tag{28}$$

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.4, the optimal allocation π_i replicates exactly this buy-and-hold strategy in the limit as $\gamma \to \infty$.

In reality, households do not invest their wealth in portfolio of zero-coupon bonds. We illustrate how they replicate this strategy, with the set of assets and contracts at their disposal, with three hypothetical households.

First, consider workers 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

[7] See Appendix C.4 for a derivation and more detailed technical discussion of this case.

so independently of the rate of return on private savings.

Second, because replacement rates fall with lifetime earnings, middle-class workers need to save privately as well. However, they can execute large inter-temporal transfers at current prices by buying a house with a fixed-rate mortgage. By doing so, they effectively trade 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.

Because this is not true for workers in the upper half of the earnings distribution, they complement 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, households would need to increase their contribution rate to maintain the same consumption level in retirement. However, if their account were mostly invested in long-term assets, capital gains would offset the decline of rates of 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 perfect sense, as it invests retirement contributions in stocks early in the life-cycle and moves towards safer assets when workers get older.

5 Matching the stylized facts

5.1 Calibration

Interest-rate dynamics To estimate our riskfree interest-rate process presented in equation (7), we calibrate its stationary mean, persistence, and volatility to match moments of the real yield curve, computed by subtracting inflation projections from the SSA annual reports from the nominal yield curve from the Federal Reserve.[8] In particular, over our sample period of 1989–2019, we target 1) the slope on a regression of the 30-year real 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 that can be estimated using a method of moments counterpart. The data moments we use for this procedure, their model counterparts, and the parameter values we obtain are shown in Table 2.

Table 2: 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

The reason we calibrate to the yield curve, as opposed to the time series of short-term real returns, is that the yield curve captures the expectations investors have over future interest rates. As such, matching the yield curve is more important for obtaining realistic asset price levels and capital gains for rate-sensitive assets.

To obtain a time series of short-term real interest rates, we use a methodology similar to that of Beeler and Campbell (2012) on the 10-year nominal Treasury bond yield and annual inflation, also described in detail in Appendix 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 variation in measured short-

[8] 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.

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.

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 $\beta = .95$ and a bequest motive equivalent to $\bar{b} = 10$ years of consumption matches the growth of wealth until retirement age and its evolution afterwards. Moreover, a coefficient of relative risk aversion of $\gamma = 6$ 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, 2021; 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 D.1.

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

$$m_{it} = \min \left\{ \chi \left(\frac{\text{AIYE}_{it}}{\bar{L}_t} \right) \times m(\text{age}_{it}), 1 \right\} \tag{29}$$

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 his age, which we calibrate as the average across gender 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 $\chi\left(\mathrm{AIYE}_{it}/\bar{L}_t\right)$ at each point of the numerical grid of the $\mathrm{AIYE}_{it}/\bar{L}_t$ state variable such that, given our labor-income process, we obtain the same life expectancy differential across percentiles of $\chi\left(\mathrm{AIYE}_{it}/\bar{L}_t\right)$ at age 40 as those reported by percentiles of earnings in Chetty et al. (2016).

5.2 Cross-section of interest-rate sensitivity

Figure 6 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.

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 own life expectancy.

Figure 6: Life-cycle profiles of wealth and its interest-rate sensitivity

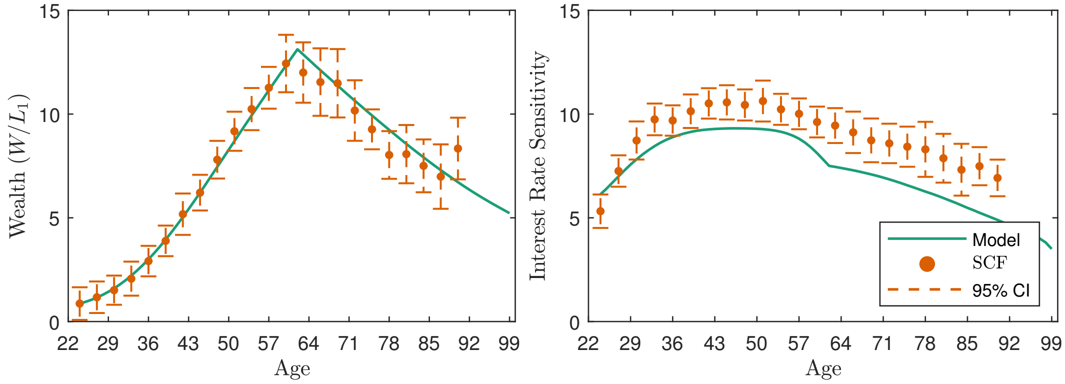

Note: This figure reports the evolution of market wealth and its sensitivity to interest rates over the life-cycle in our benchmark calibration and in the SCF. In the data, wealth is computed per adult, including deceased spouses, and scaled by the Social Security wage index. 95% confidence intervals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

Figure 7: Interest-rate sensitivity at age 40–45

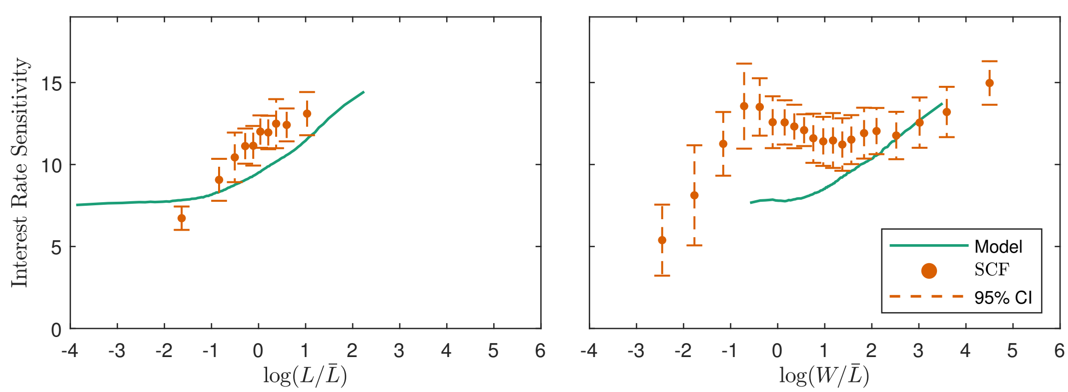

Note: This figure reports the relationship between the interest-rate sensitivity of wealth and wealth (left panel) and earnings (right panel). In the data, wealth and earnings are computed per adult and scaled by the Social Security wage index. In the left panel, each bin represents a decile of earnings. 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 ± 1.96 standard errors, clustered by year.

The left panel of Figure 7 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.

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 notably difficult for life-cycle models to match the allocation of household portfolios between stocks and short-term bonds. By 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.

5.3 Trends in wealth inequality

How much of the evolution of wealth inequality can our model and the historical path of interest rates explain? To answer this question, 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. We provide more details on the construction of this time series 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 focus is on understanding the effect of changing interest rates on trends in wealth, but our model is not designed 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 in our baseline model. Therefore, we focus our attention on the empirical evolution of the top 10% share within the bottom 99%. Second, within the bottom 99%, wealth inequality is not just generated by differences in earnings but also by heterogeneity in preferences and idiosyncratic returns on wealth.

For these reasons, we increase the σ_α from .472 to 1, such that the average wealth share of the top 10% over our sample period matches the top 10% share (within the bottom 99%) in the WID.

Figure 8 illustrates our results. In our historical simulation, the top 10% share falls from 53.9% in 1956 to 50% in 1984, then rises back to 54.2% in 2019. According to the WID, the top 10% (within the bottom 99%)[9] share fell from 58.6% to 49% in 1984, then rose back to 55% in 2019. Consequently, our model can explain roughly half of the evolution in the top 10% share (within the bottom 99%) since the 1960. 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.

Figure 8 also shows that, in the simulation, the top 10% share inclusive of Social Security has not increased since 1989, consistent with the empirical findings of Catherine et al. (2020). From the point of view of workers, Social Security is a leveraged position in the long-term asset. First, they are required to pay contributions, which is equivalent to a short position on a medium-term bond; in exchange, they will receive pension benefits far into the future, equivalent to a long position on a very long-term bond. The consequence of this leverage is that the net present value of Social Security cash flows is highly sensitive to the yield curve.

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

Figure 8: Evolution of top 10% share

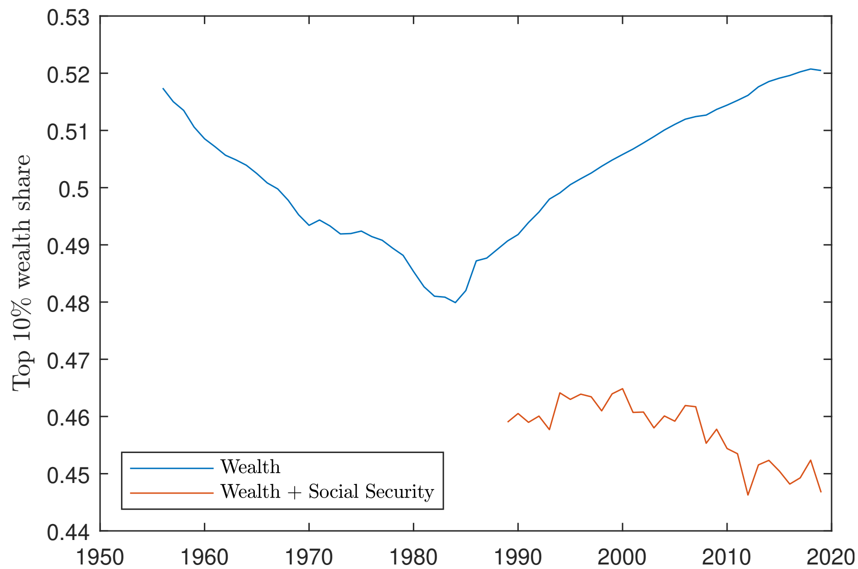

6 Discussion

In this section, we discuss the quantitative implications of the model for the cross-section. We first verify that the full model captures the same economic intuition as does the linearized benchmark. The results reveal that mortality differences are not of first-order importance to the cross-section, but Social Security is. We then study the sensitivity of financial wealth, total wealth, and lifetime utility (welfare) to interest rates. While there is a great deal of cross-sectional heterogeneity in the rate sensitivity of financial wealth, there is little heterogeneity in that of total wealth or lifetime utility. Redistribution of wealth from interest-rate shocks is inconsequential for consumption and welfare.

6.1 Mechanisms

The quantitative model validates our economic intuition and allows us to study counterfactuals. Before discussing welfare implications, we analyze the importance of two novel mechanisms in our model: income-based differences in mortality rates and the presence of Social Security. Figure 9 plots quantities of interest with and without these features.

Figure 9: Effect of Social Security and differences in life expectancy

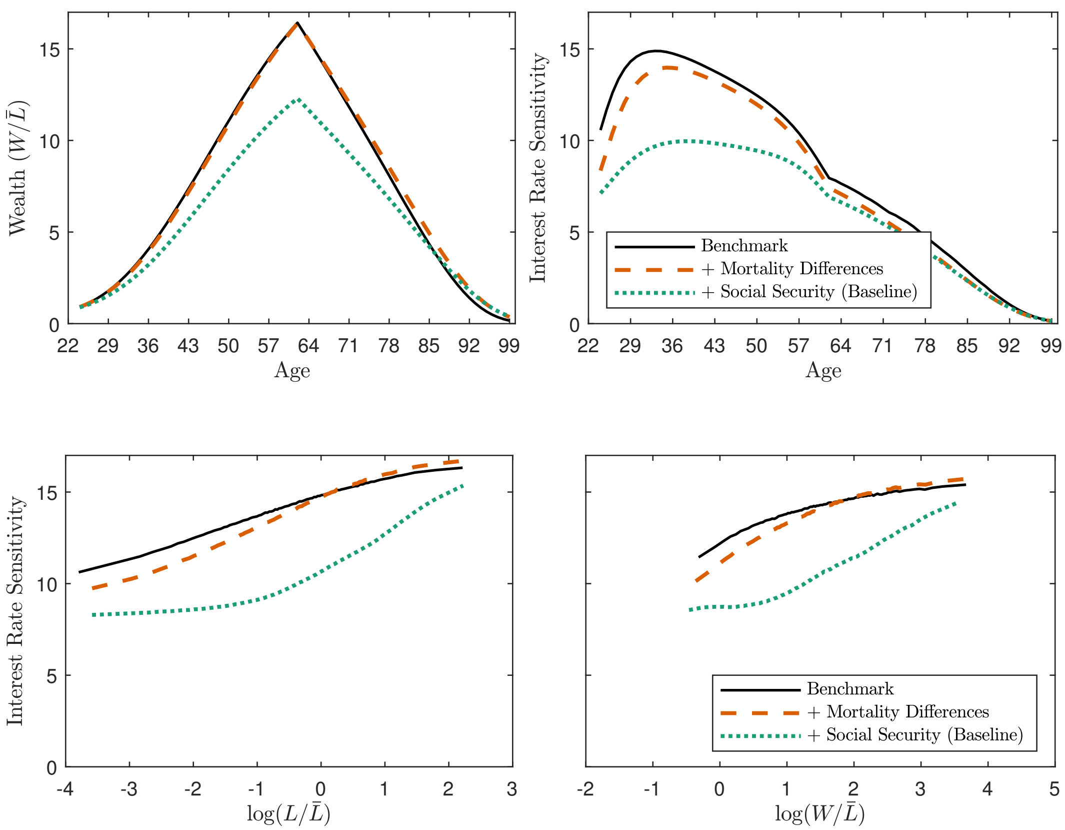

Note: This figure shows the effects of mortality differences and Social Security on life-cycle wealth accumulation and the interest-rate sensitivity of wealth in the model. In the benchmark case, mortality probabilities are constant within an age cohort and there are no Social Security taxes or benefits. 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 9. 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 9, 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.

6.2 Exposure to interest-rate shocks

So far, we have focused on explaining the interest-rate sensitivity of households' financial wealth. We now study the rate sensitivities of two measures that are more relevant for welfare: wealth inclusive of Social Security and expected lifetime utility. We find that there is less heterogeneity in these measures (especially expected utility), suggesting that the recent rise in financial wealth inequality has not necessarily come with a rise in welfare inequality.

Recall that the rate sensitivity of financial wealth W is the elasticity

$$\varepsilon(W, r_f) = -\frac{\partial \log W}{\partial r_f}. \tag{30}$$

We calculate the analogous elasticities for our other two measures. To calculate wealth inclusive of Social Security, we capitalize the expected benefits and taxes into a present value. To measure welfare, we calculate the sensitivity $\varepsilon(U, r_f)$ of a transformation of expected utility:

$$U_{it} = ((1-\gamma)V_{it})^{1/(1-\gamma)}, \tag{31}$$

where V_{it} is the expected utility maximand (4). 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.[10] Because V is a function of both wealth W and rates r_f, this elasticity can be decomposed as

$$\varepsilon(U, r_f) = \frac{1}{1-\gamma}\Bigg(\underbrace{-\frac{\partial \log((1-\gamma)V)}{\partial r_f}}_{\substack{\text{change in investment}\\ \text{opportunities}}} + \underbrace{\frac{\partial \log((1-\gamma)V)}{\partial W}\varepsilon(W, r_f)}_{\text{capital gains}}\Bigg). \tag{32}$$

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

Figure 10 shows the average paths of these elasticities over the life-cycle. Adding Social Security wealth does not change the average elasticity very much. This is consistent with the fact that most of the hump-shaped pattern is driven by human

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

capital and a diminishing investment horizon. The rate sensitivity of expected utility, on the other hand, is virtually constant over the life-cycle. At all ages, households are negatively affected by lower rates as it reduces their lifetime consumption. The magnitude of this effect is slowly declining over the life-cycle as the investment horizon declines with age.

Figure 10: Interest sensitivities over the life-cycle

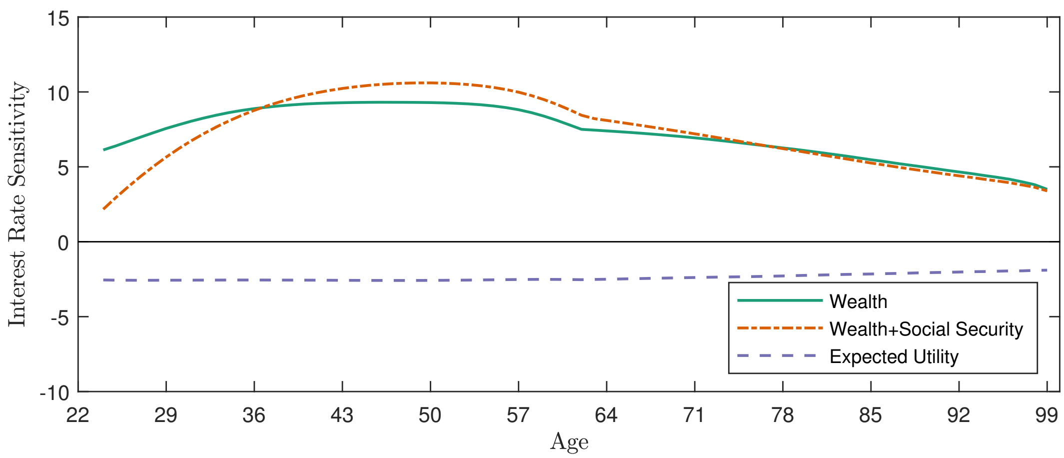

Note: This figure reports the interest rate sensitivity of wealth, wealth inclusive of Social Security, and expected utility over the life-cycle.

Figure 11 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. (2020) that, when Social Security is accounted for and discounted using the market yield curve, wealth inequality has not increased since 1989.

Second, within a cohort, expected utility is uniformly elastic to interest rates across the earnings and wealth distributions. As Equation (25) predicts, once human capital and Social Security wealth are accounted for, all households within a cohort are equally exposed to rate fluctuations.

Figure 11: Interest sensitivities at age 42

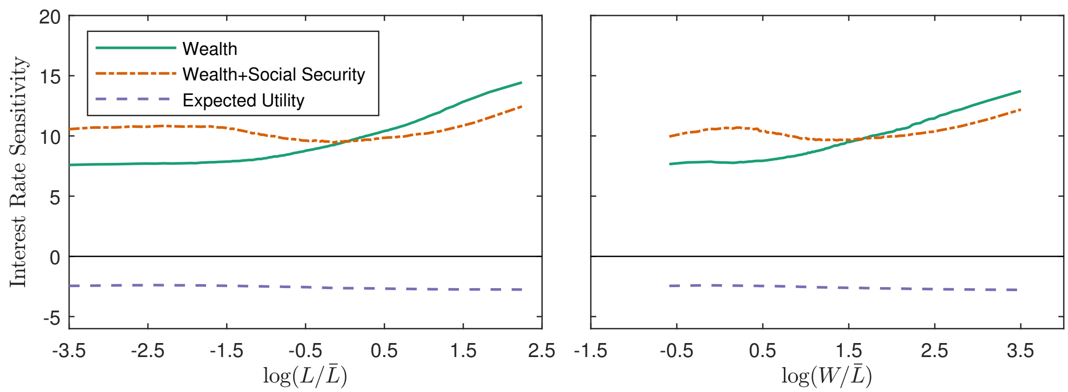

Note: This figure reports the interest-rate sensitivity of wealth, excluding and including Social Security to an interest fall and the interest rate sensitivity of expected utility at age 42.

7 Conclusion

Prior work notes that differences in returns are a key determinant in the rise in wealth inequality over the last forty years. Of particular importance to this explanation is the greater holding of long-term, highly interest-rate-sensitive assets by wealthy and high-income households, which saw greater capital gains due to the large decline in interest rates seen over the same period. This paper shows that a parsimonious life-cycle model with uninsurable income risk, a realistic Social Security system, and stochastic interest rates can generate the patterns of portfolio interest rate sensitivity observed in the cross section and makes concrete the relationship between the rise in wealth inequality and the decline in interest rates.

These results also pave the way for future research. The model above misses along some key dimensions, namely underestimating the interest rate sensitivity of the middle class, much of which is due to the long duration of home equity.

Alternative models may be able to capture this by focusing on the indivisibility of housing or explicitly modeling the consumption of housing services, both of which are beyond the scope of this paper. Moreover, future work may seek to embed the trade-off households face between investing in long- and short-term assets and how this interacts with the policy environment, to better understand the implications of certain monetary and fiscal policies for wealth inequality.

References

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

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*, 2021. hhab136.

Catherine, Sylvain, Max Miller, and Natasha Sarin, "Social security and trends in inequality," 2020.

Catherine, Sylvain, Paolo Sodini, and Yapei Zhang, "Countercyclical Income Risk and Portfolio Choices: Evidence from Sweden," 2021.

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.

Cioffi, Riccardo A., "Heterogeneous Risk Exposure and the Dynamics of Wealth Inequality," 2021.

Cocco, João F., "Portfolio Choice in the Presence of Housing," *The Review of Financial Studies*, 2005, *18* (2), 535–567.

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.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri, "Heterogeneity and Persistence in Returns to Wealth," *Econometrica*, 2020, *88* (1), 115–170.

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.

Greenwald, Daniel, Matteo Leombroni, Hanno N. Lustig, and Stijn van Nieuwerburgh, "Financial and Total Wealth Inequality with Declining Interest Rates," 2021. 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.

Kumhof, Michael, Romain Rancière, and Pablo Winant, "Inequality, Leverage, and Crises," *American Economic Review*, March 2015, *105* (3), 1217–45.

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.

Madrian, Brigitte C. and Dennis F. Shea, "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior," *Quarterly Journal of Economics*, 2001, *116* (4).

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.

Mian, Atif, Ludwig Straub, and Amir Sufi, "Indebted Demand," *The Quarterly Journal of Economics*, 2021, *136* (4), 2243–2307.

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

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. We construct networth as.

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

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.
- `fixed_inc`: value of all other remaining financial assets ($\texttt{fixed\_inc} = \texttt{fin} - \texttt{cash\_dep} - \texttt{equity}$). 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 ($\texttt{other\_dbt} = \texttt{othloc} + \texttt{ccbal} + \texttt{install} - \texttt{edn\_inst} - \texttt{veh\_inst}$).

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

A.2 Duration component calculations

A.2.1 Duration of equity

The duration of `equity` is obtained using yearly estimates for the duration of the aggregate stock market from Greenwald et al. (2021). We them apply a mean preserving adjustment to the aggregate duration value by networth decile and age group using price-dividend ratios from asset holding data in the SCF. The dividends used in the PD ratios are taxable "Ordinary Dividends" reported in IRS form 1040 line 3b, which includes all dividend income from individually held stocks and mutual funds and is given in the raw SCF as X5710. We call the assets corresponding to these dividends `tf_equity`, which we construct from the sum of SCF extract variables as

$$\texttt{tf\_equity} = \texttt{stock} + \texttt{stmutf} + \texttt{comutf} + \texttt{omutf} + \texttt{gbmutf} + \texttt{tfbmutf} + \texttt{obmutf}.$$

The major difference between the `tequity` variable and the `tf_equity` variable is the inclusion of bond mutual funds, in particular, government bond mutual funds (`gbmutf`), tax-free bond mutual funds (`tfbmutf`), other bond mutual funds (`obmutf`), and a portion of combined mutual funds (`comutf`).[11] This is because, for the purpose of Form 1040, income from bond mutual funds are taxed as dividends.

[11] The presence of bond mutual funds in the variables used to construct our adjustment could bias our estimates if bond holdings make up a large portion of `tf_equity` and differ systematically by age group and decile. However, this is not the case in the data, as the stock portion accounts for the vast majority of `tf_equity` and remains stable across age groups and networth deciles.

To understand how duration varies in the cross-section of equity holders, we split the respondents into the age groups of 20-40, 40-60, and 60+ and compute networth deciles within each group. We then sum the value of `tf_equity` and dividends within each networth-age group and divide the total asset value by the total dividend value to obtain each group's price-dividend ratio. We then create a mean-preserving adjustment multiplier by dividing these price-dividend ratios by the aggregate price dividend ratio in the SCF. This implies that the Macaulay duration of equity for each household is given by

$$\text{dur}(\text{Equity}_{dat}) = \frac{\text{PD Ratio}_{da}}{\text{PD Ratio}} \times \text{dur}(\text{Equity}_t) \tag{A.1}$$

where d represents the within age group networth decile, a represents the age group, and t represents the survey year.

Note that this adjustment is only applied to the portion of equity reported on Form 1040, namely, directly held stocks (`stock`), stock mutual funds (`stmutf`), other mutual funds (`omutf`), and a portion of combined mutual funds (`comutf`). To the other elements of `equity` whose income is not reported on Form 1040, such as portions of retirement accounts allocated to stock, we apply the duration of the aggregate stock market in that survey year.[12]

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`

[12]This follows from a literature in behavioral economics that suggests people opt in the default option for their defined contribution pension plans, usually the market portfolio (Madrian and Shea, 2001).

– 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, we assign a cash flow duration of 4.

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. (2021) Appendix E.2.4. 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 cashflows from those equity holdings. However, the annual cashflows cashflows from those equity holdings are not reported in the SCF, the major issue being that that cashflows from private businesses partially contain implicit or explicit labor income for the entrepreneur. As such, we must estimate or difference out this labor income, which we do in four ways depending in 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 households wage income is the dependent variable, and the independent variables are a third degree polynomial in age interacted with dummies for each Race$\times$ Eduction $\times$ 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 don't meet the criteria for a wage subtraction are given cashflows equal to `busefarminc`.

We then aggregate `bus` and the estimated annual cashflows and divided them to obtain our proxy for duration.

Next, to allow our aggregate estimates of private business duration to vary over the wealth distribution, we follow a similar procedure as we did with publicly traded equity. First, we split the population into age groups of 20-40, 40-60, and 60+ and compute networth centiles within each age group. We then sum the business wealth (`bus`) and total income from businesses (`busefarminc`) within each centile-age group to obtain a price-total income ratio. Provided that cashflows from equity are proportional to labor income, this provides a proxy for duration within each networth centile-age group. These price-total income ratios are then divided by the aggregate price-total income ratio ratio to obtain a mean preserving adjustment which is applied to the annual aggregate private business duration estimates. This is given by

$$\text{dur}(\text{Private business}_{cat}) = \frac{\text{Price-total income ratio}_{ca}}{\text{Price-total income ratio}} \times \text{dur}(\text{Private business}_{t}) \tag{A.2}$$

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 cars remaining years, and calculate the duration using (3). 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 (3), where N is the number of years remaining of the vehicles expected lifetime 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

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

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

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

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.3 Time series of riskfree rates

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

$$y_{10,t} - \pi_{t,t+1} = \beta_0 + \beta_1 y_{10,t} + \beta_2 \pi_{t-1,t} + \epsilon_{t+1} \tag{A.3}$$

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 (7) yields the time-t riskfree rate:

$$r_{ft} = \varphi^{-10}(\hat{f}_{10,t} - (1 - \varphi^{10})\bar{r}_f). \tag{A.4}$$

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

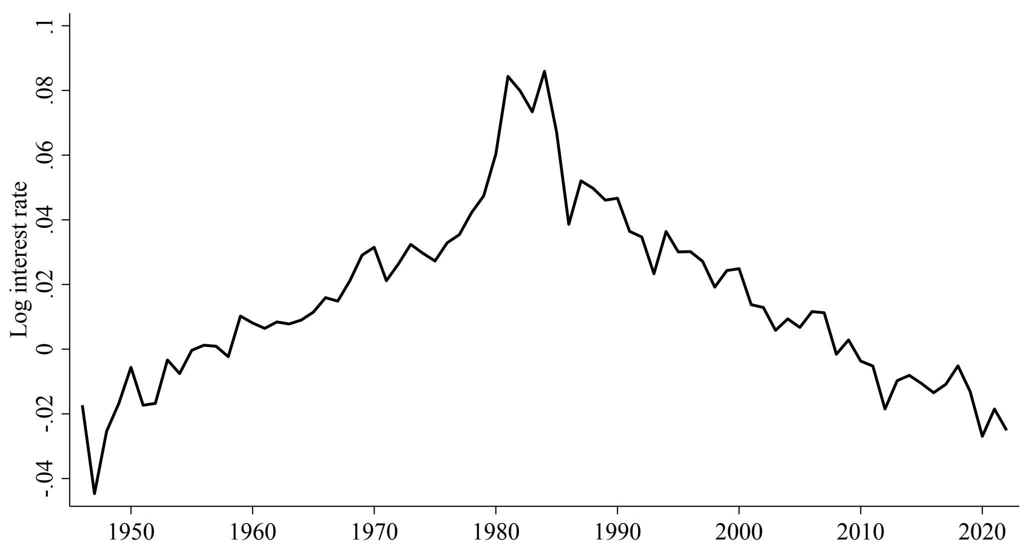

Note: This figure presents the time series of short-term riskfree rates as estimated by Equation (A.3) and transformed by Equation (A.4).

As discussed above, 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 (7) imply the n-period bond returns (8). 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 (7) 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

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

the stated expression (8).

B.3 Derivation of private-business valuation and duration

Let E_n represent the value of business income n periods into the future (the dividend strip with maturity n). By analogy to the zero-coupon bonds, assume that the dividend yield on this strip equals

$$y_{nt}^{(D)} = -\frac{1}{n} \log \frac{E_{nt}}{D_t} = y_{nt} + \frac{1}{n} \sum_{j=1}^{n} \tilde{\mu}_{jD},$$

where we will set $\tilde{\mu}_{jD}$ to get a constant risk premium. Note the boundary condition $E_{0t} = D_t$ (we verify below that this holds in our solution). By definition, the return on this claim is

$$R_{n,t+1}^{(D)} = \frac{E_{n-1,t+1}}{E_{nt}},$$

so the log return

$$r^{(D)}_{n,t+1} = ny^{(D)}_{nt} - (n-1)y^{(D)}_{n-1,t+1} + \log \frac{D_{t+1}}{D_t}$$
$$= r_{ft} + \underbrace{\mu_n - \sigma_n \epsilon_{r,t+1}}_{r_{n,t+1}-r_{ft}} + \underbrace{\tilde{\mu}_{nD} + g_D + \sigma_D \epsilon_{D,t+1}}_{r^{(D)}_{n,t+1}-r_{n,t+1}}.$$

To target a maturity-invariant risk premium μ_E, including adjustments for Jensen's inequality, we need to set

$$\tilde{\mu}_{nD} = \mu_D - g_D - \frac{1}{2}\sigma_D^2 - \underbrace{\left(\mu_n + \frac{1}{2}\sigma_n^2\right)}_{=0}.$$

Combining these results, we have the strip value

$$E_{nt} = E(n, D_t, r_{ft}) = P_{nt} D_t \exp\left\{-n\left(\mu_D - g_D - \frac{1}{2}\sigma_D^2\right)\right\}.$$

Then the total value of private business equity is $E_t = \sum_{n=1}^{\infty} E_{nt}$.

The duration of this claim is simply

$$\mathrm{dur}(E_t) = \sum_{n=1}^{\infty} n \frac{E_{nt}}{\sum_{n'=1}^{\infty} E_{n't}} = \sum_{n=1}^{\infty} n \frac{P_{nt}\exp\left\{-n\left(\mu_D - g_D - \frac{1}{2}\sigma_D^2\right)\right\}}{\sum_{n'=1}^{\infty} P_{n't}\exp\left\{-n'\left(\mu_D - g_D - \frac{1}{2}\sigma_D^2\right)\right\}},$$

a function of r_{ft}. To get a sense of this duration value, let r_{ft} equal $\bar{r}_f$ and ignore the Jensen's inequality term, so that $P_n \approx \exp\{-n\bar{r}_f\}$. Then we have

$$\overline{E}_{nt} = E_t \exp\left\{-n\left(\bar{r}_f + \mu_D - g_D - \frac{1}{2}\sigma_D^2\right)\right\},$$

implying

$$\overline{E}_t = D_t \frac{\exp\left\{-\left(\bar{r}_f + \mu_D - g_D - \frac{1}{2}\sigma_D^2\right)\right\}}{1 - \exp\left\{-\left(\bar{r}_f + \mu_D - g_D - \frac{1}{2}\sigma_D^2\right)\right\}}.$$

It follows that the duration is approximately

$$\mathrm{dur}(\overline{E}_t) = \frac{1}{1 - \exp\left\{-\left(\bar{r}_f + \mu_D - g_D - \frac{1}{2}\sigma_D^2\right)\right\}} \approx \frac{1}{\bar{r}_f + \mu_D - g_D - \frac{1}{2}\sigma_D^2}.$$

The denominator is a risk-adjusted "r minus g" term.

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 (5) 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 $\log\left(1 - e^{c_t - w_t}\right)$ 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 $\kappa_w(m_t) = \log \rho_c(m_t) + (1-\rho_c(m_t)) \log\left(1 - \rho_c(m_t)\right)/\rho_c(m_t)$.[13] (Notice that, as $m_t \to 1$, $c_t \to w_t$; agents who will die almost surely consume everything.) We can

[13] 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.

also get the linearized approximation to the log wealth return

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

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:

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

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

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

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

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

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

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

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

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

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

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

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 (8), is our expression for the optimal share in the n period bond given by (22).

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

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

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

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

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.13), (C.14), 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 $\lim_{t\to\infty} \varrho_{rt} = 0$:

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

The higher is the mortality probability, 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}, which we can similarly iterate forward with terminal condition $(1 - \psi)\varrho_0 \to 0$ to arrive at a solution. 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 value 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. Let $p_{\text{surv}}(t, j) = \prod_{s=1}^{j}(1 - m_{i,t+s})$ denote the cumulative probability of surviving from t to $t + j$. For human capital, the log return is approximately

$$
r_{H,t+1} = r_{ft} + \mu_{Ht} + \underbrace{\left(\sum_{j=0}^{t_{\text{ret}}-t} \omega_{jt}^H \left(\frac{\sigma_j}{\sigma_n}\right)\right)}_{\pi_t^H} (r_{n,t+1} - r_{f,t+1})
$$

where

$$
\omega_{jt}^H = \frac{p_{\text{surv}}(t, j) P_{jt} L_{t+j}}{\sum_{j'=0}^{t_{\text{ret}}-t} p_{\text{surv}}(t, j') P_{j't} L_{t+j'}} = \frac{p_{\text{surv}}(t, j) P_{jt} L_{t+j}}{H_t}
$$

is the value weight of the jth labor-payment, and therefore π^H is a value-weighted rate-sensitivity adjustment. More specifically, a share π^H of the total value of human capital is an implicit holding of n-period bonds; this share is a value-weighted average of the cross-price elasticities

$$\frac{\sigma_j}{\sigma_n} = \frac{-\partial \log P_{jt}/\partial r_{ft}}{-\partial \log P_{nt}/\partial r_{ft}} = \frac{\partial \log P_{jt}}{\partial \log P_{nt}},$$

and hence constitutes an adjustment for the duration of the income stream relative to the traded n-period bond. Identical logic leads us to conclude that the log return on Social Security is

$$r_{S,t+1} = r_{ft} + \mu_{St} + \underbrace{\left(\sum_{j=0}^{\infty} \omega_{jt}^{S} \left(\frac{\sigma_j}{\sigma_n}\right)\right)}_{\pi_t^S} (r_{n,t+1} - r_{f,t+1}),$$

where the value weights take the form

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

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

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

where the return on total wealth

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

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 from 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 (25).

C.4 Optimal consumption plan in the limit

This section derives the 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 some a deterministic constant $C_t = \bar{C}$. 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 $\bar{C}_t = \bar{C}_{t+1} = \bar{C}$ — that is, 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 $m_t = 1$.[14] 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.

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

$$r_{w,t+1} = r_{ft} + \underbrace{\left(\sum_{j=0}^{t_{\max}-t} \frac{P_{jt}}{\sum_{j'=0}^{t_{\max}-t} P_{j't}} \left(\frac{\sigma_j}{\sigma_n}\right)\right)}_{\tilde{\pi}_t} (r_{n,t+1} - r_{ft}) \tag{C.22}$$

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 optimal holding π^*. In the limit, the general expression for optimal consumption (21) 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

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

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

Substituting this into the expression for the optimal portfolio $\bar{\pi} = \pi^*$ in (22), 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 Numerical appendix

Table D.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