Who gains when interest rates fall?*

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

September 7, 2025

working version, please do not circulate

Abstract

We study the interest-rate sensitivity of household wealth in a realistic life-cycle model. The model predicts that middle-aged and wealthier households should hold more long-term assets, as observed in the US data. Consequently, optimal portfolio rules imply that falling interest rates increase wealth inequality, while rising rates reduce inequality, consistent with historical experience. However, these long-run shifts in wealth inequality are largely offset by changes in the valuation of human capital and Social Security benefits, mitigating the passthrough of interest-rate fluctuations to expected lifetime consumption.

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

*This paper previously circulated under the title "Interest-rate risk and household portfolios." We thank seminar participants at the University of Maryland, Dartmouth Tuck, Berkeley Haas, Columbia University, the London School of Economics, the University of Southern California, the University of Texas at Austin, the University of Washington Seattle, and Stanford GSB for helpful comments.

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

Changes in real interest rates have significant but ambiguous effects on households. On the one hand, a decline in rates raises asset prices. These capital gains are substantial and unevenly distributed, as they primarily accrue to owners of long-term assets. Over the past four decades, declining interest rates have been a driving force behind exceptional stock-market returns (Binsbergen, 2021) and the surge in wealth inequality (Greenwald et al., 2023). On the other hand, lower rates make it more difficult to accumulate wealth going forward. In order to finance $1 of consumption at age 70, a 35-year-old must set aside 71 cents at a 1% return, compared to only 36 cents at a 3% return. That capital gains occur precisely when future returns decline complicates our understanding of who gains or loses from long-run fluctuations in interest rates.

Given the ambiguous effects of interest-rate changes on lifetime consumption, it is not obvious which households should hold longer- or shorter-term assets in their portfolios. We therefore use a life-cycle portfolio choice model to study the optimal interest-rate sensitivity of wealth. Households demand long-term assets to protect against future interest-rate declines. The model predicts that the optimal duration of household balance sheets should be highest for middle-aged, high-income, and high-wealth households. These patterns align well with the data, suggesting that household portfolios are close to individually optimal. The model further predicts large and persistent increases (decreases) in wealth concentration when real interest rates fall (rise). However, these trends in inequality are attenuated when measured in terms of lifetime consumption (utility) or when the present value of Social Security is included as wealth.

Understanding *why* households hold portfolios of different durations yields new insights for the debate on wealth inequality. Looking at wealth data alone, one might infer that falling interest rates have disproportionately benefited those at the top. But within the full life-cycle context, households at the bottom are better insulated against rate changes than they appear. Observed portfolios align well with individually optimal allocations, so that changes in wealth inequality driven by interest-rate movements do not imply commensurate changes in inequality of lifetime consumption. Our findings have two important implications. First, broader wealth concepts—incorporating background assets such as Social Security and human capital—offer a more welfare-relevant measure of long-run inequality trends. Second, given current public policies and financial institutions, US households appear capable of constructing portfolios that are broadly consistent with individually optimal rules.

Our first objective is to explain key stylized facts. At first glance, household balance sheets suggest that some Americans are better positioned to gain from falling interest rates. First, the interest-rate sensitivity of wealth is hump-shaped over the life cycle. Second, within an age group, the interest-rate sensitivity of wealth increases with both earnings and wealth. In other words, middle-aged and richer households invest more in long-term assets and should capture a disproportionately larger share of aggregate capital gains than their poorer and younger counterparts when rates decline. Consistent with this mechanism, Greenwald et al. (2023) show that, over the past 60 years, trends in wealth inequality have tracked changes in the valuation of long-term bonds.

To understand these facts, we use our life-cycle model to compute individually optimal portfolio rules and compare them to portfolios in the data. In the model, households decide how much to consume and can invest their wealth in housing, stocks, and bonds (or debts). Their bond portfolio is composed to two assets (or liabilities): a short- and a long-term zero-coupon bond. By mixing these securities, households can reach the desired interest-rate sensitivity of their balance sheets. Taking into account households' stochastic stream of labor income, the progressive Social Security system, and stochastic interest rates, we find that the patterns we document in the data are broadly consistent with optimal portfolio rules. We calibrate and validate our model by matching other dimensions of portfolios, including wealth levels and allocations to equities and housing.

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

Another important background asset is Social Security. Its role can be understood in two ways. First, because it pays benefits from retirement until death, Social Security remains a long-term asset over most of the life cycle. Equivalently, because it represents an alternative savings vehicle with a

rate of return that does not depend on the market rate, Social Security partly insulates workers from interest-rate fluctuations. Either way, it reduces the hedging demand for long-term assets. This substitution effect is larger at the bottom of the earnings distribution because retirement benefits replace a smaller portion of lifetime earnings at the top. Consequently, the demand for long-term assets increases with earnings and wealth within an age group, as observed in the data

How do these abstract portfolio rules translate into households' financial behavior? A household that wants to eliminate interest-rate risk completely should invest in assets with cashflows of different maturities so that, together with expected future earnings and benefits, these payments would match the timing of its desired consumption plan. If held to maturity, these assets would finance a predetermined consumption path regardless of interest-rate fluctuations, because each cashflow would arrive exactly when needed. In practice, households engage in a mix of voluntary and mandatory arrangements that mimic aspects of this abstract portfolio rule.

At the lower end of the earnings distribution, workers save primarily through contributions to Social Security, whose rate of return is unaffected by interest rates. Payroll taxes and benefits redistribute labor earnings over the life cycle to match the timing of a smoothed consumption plan, so these households do not need to actively invest in long-term assets to hedge against interest-rate changes—Social Security already provides that protection. For lower middle–income households, Social Security covers most, but not all, of retirement consumption. They typically supplement it with private saving, most often by buying a house with a fixed-rate mortgage. This arrangement exchanges future housing services for a fixed stream of payments, locking in the cost of residential consumption. It protects households against both interest-rate risk as well as aggregate house-price risks. Social Security then finances most of their other retirement spending. Higher earners, who receive the lowest replacement rates from Social Security, rely more heavily on private retirement saving, often in stock-based accounts. Because stock values tend to rise when interest rates fall, these assets help sustain retirement consumption when rates of return fall.

Our second objective is to study what optimal portfolio rules imply for trends in wealth inequality and their interpretation. We show that, given the historical path of interest rates, these rules generate large fluctuations in the concentration of wealth. These fluctuations account for nearly half of the decline in wealth inequality between 1960 and 1985, and for much of its subsequent increase in the following decades. Moreover, rising stock prices—net of what would be

predicted by the stock market's interest-rate exposure—further amplified these trends. On the other hand, similar changes in national house prices had the opposite effect.

These trends look markedly different under broader wealth concepts, a prediction that echoes the empirical finding that trends in inequality strongly depends on the inclusion or exclusion of accrued Social Security benefits (Catherine et al., 2025). In our model, heterogeneity in portfolios (and thus in wealth returns) are largely driven by substitution effects from off-balance-sheet assets. Therefore, once Social Security and human capital are included in wealth, much of this heterogeneity disappears. We document that this prediction of our model is consistent with the data: including Social Security in our measure of wealth evens out the interest-rate sensitivity of wealth across the earnings and wealth distributions.

A consequence of this insight is that trends in wealth inequality induced by interest-rate fluctuations must be interpreted with caution. Indeed, our model generates large changes in the concentration of wealth—roughly half of historical variations since 1960—without any increase in within-cohort *welfare* inequality. This follows intuitively from the optimal portfolio decision: households target the interest-rate sensitivity of their wealth to offset the implicit exposure of their background assets, so that, ex post, their lifetime consumption is similarly exposed to interest-rate shocks.

There are some important caveats to our main findings. First, not all households can attain their optimal exposure to interest-rate risk. Entrepreneurs, for instance, typically hold most of their wealth in a single high-duration asset—their private business. The composition of their balance sheet is driven less by portfolio-choice considerations than by professional necessity. In our model, entrepreneurs sell their business at retirement, making its value directly sensitive to interest rates at that point. As a result, declines (increases) in interest rates raise (reduce) entrepreneurial wealth, translating into actual, uncompensated gains (losses) in lifetime consumption.

A second example is young adults. Starting their working lives with little or no wealth, they are unable to hedge against the interest-rate environment they face at entry. Moreover, because they do not fully insure against interest-rate risk (except under infinite risk aversion), younger, long-horizon investors remain somewhat more exposed to falling rates than their parents and grandparents. As a result, interest-rate levels have economically significant effects on lifetime consumption.

Related literature Our paper bridges the literatures on portfolio choice and wealth inequality.

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

Previous work justifies this fresh perspective. First, Lustig et al. (2013) find that events in the long-term bond market, not the stock market, are the greater source of risk for households. This motivates our focus on portfolio duration rather than the equity share. Second, for the subset of wealth held in stocks, a key feature is its high interest-rate sensitivity: Binsbergen (2021) shows that, over the last fifty years, duration explains most of the market's exceptional returns. Gomez and Gouin-Bonenfant (2024) show that lower rates increase wealth inequality through entrepreneurship, because raising equity is cheaper and the valuation of private business is higher. Greenwald et al. (2023) document that wealthy households invest more in long-term assets, so their capital gains are larger when rates fall. Our contribution is to derive optimal portfolio rules instead of taking them from the data. This allows us to explain why households choose different durations and to determine how rate fluctuations ultimately pass through to lifetime consumption.

Our findings complement existing work analyzing the welfare implications of interest-rate-driven redistribution. Fagereng et al. (2025) develop a sufficient-statistics approach to measure how historical changes in asset prices redistributed welfare using transaction data. Auclert (2019) studies the implications of unhedged interest-rate exposure for monetary policy transmission. Like Auclert (2019), we emphasize that the interest-rate exposure of wealth inclusive of non-financial income is the welfare-relevant measure. Our life-cycle approach provides a benchmark for evaluating the extent to which the interest-rate exposure of wealth is individually optimal.

In endogenizing portfolio choices, we explain an important driver of portfolio-return heterogeneity. The importance of this work is highlighted by Moll (2021), who argues that explaining the portfolio choices that generate heterogeneous returns is essential. Benhabib et al. (2019), Bach et al. (2020), and Hubmer et al. (2021) have empirically documented that the higher returns of the

wealthy are crucial for explaining wealth inequality and its evolution. We provide an explanation for a large part of this heterogeneity: wealthier households should invest more in long-term assets to be equally hedged against interest-rate risk.

Recent work by Fagereng et al. (2021) shows that capital gains explain most of the rise in wealth inequality because the rich "save by holding" instead of selling assets to consume. This behavior is consistent with interest-rate hedging. As an extreme example, if households hold a perfectly interest-rate-hedged portfolio to maturity, then they will be able to afford the same consumption plan regardless of discount-rate-induced changes in asset prices. As Cochrane (2022) puts it, a family should not pay attention to these fluctuations if it intends to live off of the coupons.

Finally, this paper builds on recent studies that have focused on the ways in which progressive government programs attenuate wealth inequality (Catherine et al., 2025), income inequality (Auten and Splinter, 2024), and their passthrough to lifetime consumption (Auerbach et al., 2023). Our paper shows that including transfer income in wealth yields a measure that is more relevant for the cross-sections of both portfolios and welfare. In particular, we show how the interplay between public programs like Social Security and optimal portfolio choices can generate diverging trends in wealth and consumption inequality. Our mechanism helps explain Meyer and Sullivan (2023)'s finding that, over the past five decades, the rise in overall consumption inequality was small.

1 Data

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

1.1 Definition of interest-rate sensitivity

Consider an asset with current price P_t and future cashflows $\{D_{t+k}\}_{k=1}^{\infty}$. Let r_{ft} denote the log real interest rate. We define the *interest-rate sensitivity* of the asset as the change in the asset's price in response to a decline in the interest rate, holding future cashflows and risk premia fixed:

$$\varepsilon_r(P_t) \equiv -\frac{\partial \log P_t}{\partial r_{ft}}. \tag{1}$$

Appendix A gives a formal definition, along with an explanation of how this concept relates to an asset's *duration*, which is defined as the value-weighted timing of its cashflows.

1.2 Survey of Consumer Finances

Our goal is to measure the interest-rate sensitivity of assets on households' balance sheets. To do this, our primary data source is the triennial Survey of Consumer Finances (SCF), from which we use all survey years from 1989 to 2019. We use three main series: (i) detailed information on household assets and liabilities, (ii) interest rate and maturity data for household liabilities, and (iii) income data for household assets, which we use to compute valuation ratios. We supplement the SCF with data on Social Security wealth and interest-rate sensitivity from Catherine et al. (2025), duration estimates from Greenwald et al. (2023) and Bloomberg, and the real yield curve, computed by subtracting inflation projections from the SSA annual reports from the nominal yield curve from the Federal Reserve.[1]

1.3 Interest-rate process

The interest-rate sensitivity of assets depends on how shocks to current interest rates will affect future interest rates. We assume that log interest rates follow a first-order autoregression, given by:

$$r_{f,t+1} = (1-\varphi)\bar{r}_f + \varphi r_{ft} + \sigma_r \epsilon_{r,t+1}, \tag{2}$$

where ϵ_r is a standard normal shock. When interest rates follow an AR(1) process, the interest-rate sensitivity of a zero-coupon bond paying off 1 in k years is

$$\hat{\varepsilon}_r(P_{kt}) = \frac{1-\varphi^k}{1-\varphi}. \tag{3}$$

Here, the bond's maturity k is also its duration. The interest-rate sensitivity is increasing in the duration; they are identical ($\hat{\varepsilon}_r(P_{kt}) \to k$) if interest-rate shocks are permanent ($\varphi \to 1$).

Since our focus is on obtaining realistic asset-price levels and capital gains for rate-sensitive

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

assets, we calibrate the stationary mean, persistence, and volatility to match moments of the real yield curve. Over our sample period of 1989–2019, we target (1) the slope from a linear 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. Table 1 reports the data moments, their relationship to model parameters, and implied parameter estimates.

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

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

1.4 Interest-rate sensitivity of assets

We obtain the rate sensitivity of households' wealth in three steps. First, we estimate cashflow duration for each asset and liability on households' balance sheets. Second, we apply equation (3) to this cashflow duration, taking the rate sensitivity of the asset to be the same as the rate sensitivity of a riskfree zero-coupon bond with the same duration. Third, we take a value-weighted sum of the assets' rate sensitivities to arrive at the overall portfolio rate sensitivity of each household.[2]

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

Equity, real estate, liquid assets, and fixed income For all equity, real estate, and liquid assets, we apply the annual group-wide duration estimates provided by Greenwald et al. (2023). For fixed-

[2]We formally derive the second and third steps in Appendix A.

income assets, we collect annual average duration estimates from Bloomberg for government debt, municipal bonds, mortgage-backed securities, foreign bonds, and corporate bonds and apply them to each asset within the fixed-income group accordingly.

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

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

$$\text{dur}(\text{Private Business}_t) = \frac{\text{Business Value}_t}{\text{Business Income}_t - \text{Wages}_t}. \tag{4}$$

We use the price-dividend ratio here because it is equal to the duration when discount rates and cashflow growth are constant (Binsbergen, 2021; Gordon and Shapiro, 1956).[3] The business value is the value of the household's ownership share and dividends are the household's business income less wages to business owners, which are either reported in the SCF or estimated using information on their level of education and age. We obtain an annual time series by summing each component within each SCF survey year and applying equation (4).

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

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

$$\text{dur}(\text{Debt}_t) = \sum_{n=1}^{N} \left(\frac{P_{nt}}{\sum_{n'=1}^{N} P_{n't}} \right) n, \tag{5}$$

[3]That is, for constant cashflow growth rate g, riskfree rate r_f, and risk premium μ, the price-dividend ratio is $P_t/D_t = \int_0^\infty e^{-(r_f+\mu-g)n} dn = 1/(r_f+\mu-g)$ and the duration is $\int_0^\infty \frac{e^{-(r_f+\mu)n} D_{t+n}}{P_t} n dn = 1/(r_f+\mu-g)$.

where N is the number of years remaining on the loan given in the SCF and $P_{nt} = e^{-ny_{nt}}$ where y_{nt} is the n-year yield. The number of years remaining on the loan is either given explicitly in the SCF or can be inferred from the interest rate and loan balance outstanding. The exception to this is adjustable rate mortgages, whose duration we assign a maximum of 4—the average number of years in the data until they begin to float.

1.5 Interest-rate sensitivity of wealth

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

$$\hat{\varepsilon}_r(\text{Wealth}_i) = \sum_j \frac{\text{Asset}_{ji}}{\text{Wealth}_i} \times \hat{\varepsilon}_r(\text{Asset}_{ji}), \tag{6}$$

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

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

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

Table 2: Average duration and portfolio share by asset group

	IR Sensitivity	Duration	Portfolio share	
			equal-weighted	wealth-weighted
Assets				
Private Business	16.26	28.12	0.04	0.18
Equity	20.38	51.45	0.09	0.18
Real Estate	9.78	12.53	0.44	0.40
Fixed Income	5.00	5.53	0.08	0.12
Vehicles	3.17	3.34	0.20	0.03
Cash and Deposits	0.25	0.25	0.15	0.08
Liabilities				
Mortgage Debt	8.23	10.01	0.52	0.83
Other Debt	2.59	2.69	0.48	0.17
Networth				
Value-weighted	12.56	17.94		
Equal-weighted	8.36	10.21		

Note: This table reports the average interest-rate sensitivity, duration, and portfolio share of each asset group for households in the SCF. Interest-rate sensitivity is calculated from asset duration as in equation (3). To calculate average duration for each asset group, we take the duration estimate for each household's holdings in the group and then average them across households, weighting each household's contribution by its SCF sample weight and the value of its holding in the asset group. To obtain the equal-weighted portfolio shares, we take the share of each asset group within a household's portfolio and then average the shares across households using SCF sample weights. We repeat this process for the wealth-weighted portfolio shares, but in this case we weight each household by both its SCF sample weight and its networth.

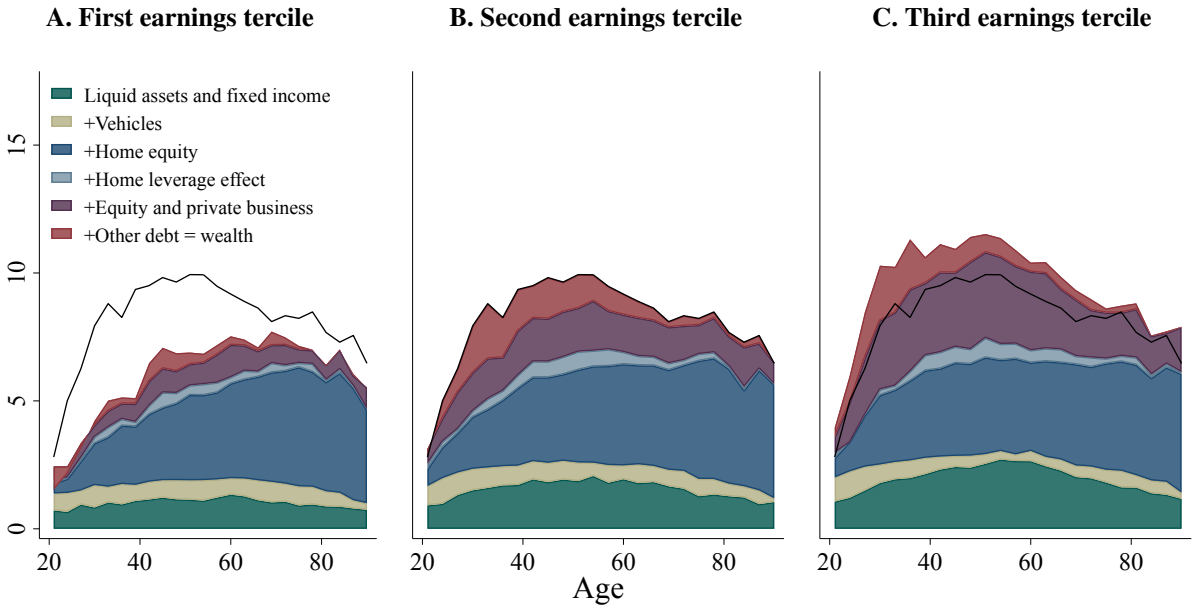
2 Stylized facts

We begin by documenting facts about the directly observable side of household's interest-rate exposure: the interest-rate sensitivity of wealth $\varepsilon_r(W)$. Even though it only paints a partial picture of households' exposures to interest rates, the distribution of $\varepsilon_r(W)$ is interesting because it reflects households' portfolio allocation decisions. Five stylized facts will later guide our structural analysis: (1) interest-rate sensitivity is hump-shaped over the life cycle; (2) high earners hold assets with higher interest-rate sensitivity; (3) interest-rate sensitivity is increasing in wealth; (4) trends in wealth inequality follow trends in interest rates; and (5) Social Security offsets differences in exposure to interest-rate risk.

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

The first stylized fact is that the rate sensitivity of wealth is hump-shaped over the life cycle: it is lowest for 20-year-olds, rises to a high for 40- to 45-year-olds, and steadily declines thereafter. Figure 1 decomposes this pattern clearly, showing the relative contribution of each asset to the total portfolio rate sensitivity. The difference in portfolio interest-rate sensitivities at each age is determined by the assets households choose to hold. For example, 20- to 25-year-old households have relatively low interest-rate sensitivities because the majority of their wealth (70.4%) is invested in liquid accounts (e.g., checking and savings accounts) and vehicles. Their holdings of longer-term assets like stocks and home equity are substantially smaller than later in life.

Figure 1: Interest-rate sensitivity of wealth by age



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

As households approach midlife, the composition of assets changes and the interest-rate sensitivities of their portfolios grow. Shorter-term liquid assets and vehicles contribute roughly the same to interest-rate sensitivity as they do for the young, but now, the majority of the portfolio

(48.7% for 40-year-olds) is made up of longer-term assets like equity and real estate. Moreover, leverage—in particular, mortgages and other debts—plays a more important role, increasing the rate sensitivity of the wealth portfolio by nearly 20%. The reason leverage increases the rate exposure of the household's portfolio is because the (equal-weighted) average rate sensitivity of assets is approximately 40% higher than that of debts over our sample. The net position, therefore, has a higher interest-rate sensitivity

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

2.2 Interest-rate sensitivity is increasing in earnings

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

2.3 Interest-rate sensitivity is increasing in wealth

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

For low-wealth households, liquid accounts, vehicles, and non-mortgage debt contribute the most to the interest-rate sensitivity of their portfolios. For middle-wealth households, real estate becomes the dominant asset, with its rate sensitivity amplified by the mortgage taken on to finance the purchase. The large indivisible nature of houses leads lower-middle-wealth homebuyers to take out large mortgages and expose themselves to interest-rate fluctuations, which is reflected in the

small bump in rate sensitivity near the lower-middle portion of the wealth distribution. As wealth increases, portfolio rate exposures increase with larger positions in highly rate-sensitive assets like publicly traded equity and private businesses.

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



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

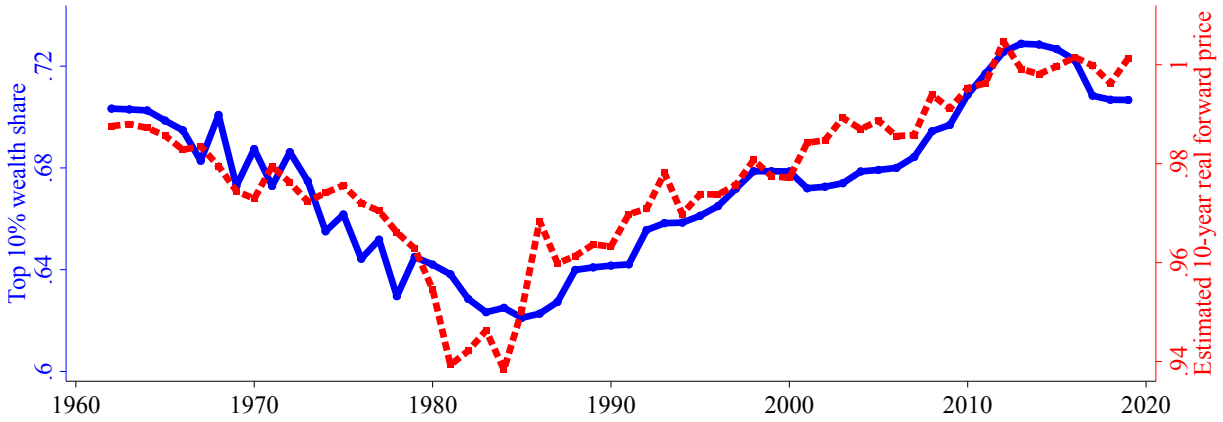
2.4 Wealth inequality follows interest rates

Our fourth stylized fact is that, as document by Greenwald et al. (2023), over the last six decades, the wealth share of the top 10% of the wealth distribution has closely tracked the price of real bonds. Figure 3 replicate their findings, approximating the real bond price by one minus of the estimated 10-year forward rate.

2.5 Social Security offsets differences in rate sensitivity

We now extend our definition of wealth to include the net present value of Social Security payments—that is, of expected benefits minus expected payroll taxes to be paid into the system. Figure 4

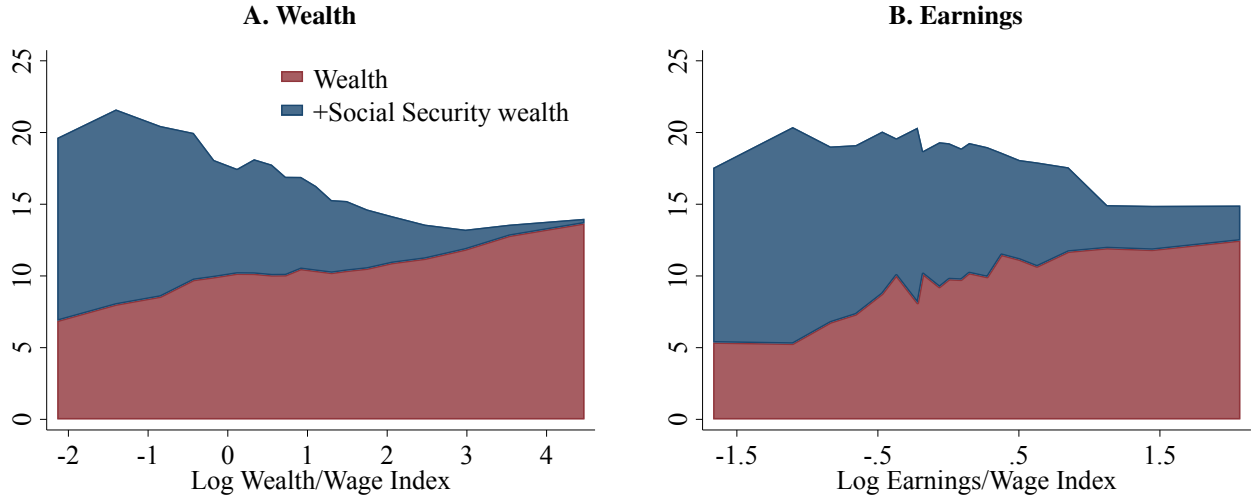
Figure 3: Wealth inequality and estimated 10-year real forward rates



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 (D.1).

displays the fifth stylized fact: the inclusion of Social Security wealth strongly attenuates the relationships between interest-rate exposure and wealth (Panel A) and earnings (Panel B).

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



Note: This figure decomposes the interest-rate sensitivity for households in which the head of the household is between 40 and 45. The methodology is the same as in Figure 1, except that here the x-axis is the log of wealth scaled by the Social Security Wage Index in the survey year in Panel A and scaled earnings in Panel B. Estimates for the net present value of Social Security at the individual level come from the risk-adjusted valuation of Catherine et al. (2025).

3 Model

We model household consumption and investment decisions over a life cycle divided into two stages: working age and retirement. We first describe the economic environment and then define households' consumption-saving problems, including market frictions.

3.1 Interest rates and economic growth

Interest rates and economic growth vary over time in a corrrated way. We thus model them as cointegrated stochastic processes. As in our empirics, we assume that the log riskfree rate, $r_{ft} = \log R_{ft}$, follows a first-order autoregression:

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

To capture the cointegrating relationship between rates and growth, we assume that the log growth rate of aggregate income (defined below) is the sum of an autoregressive process and a loading on the riskfree rate:[5]

$$g_{t+1} = g_{1,t+1} + \lambda_{gr} r_{f,t+1}, \tag{7}$$

where

$$g_{1,t+1} = (1-\varphi_g)\bar{g} + \varphi_g g_{1t} + \epsilon_{g,t+1}. \tag{8}$$

To capture business-cycle dynamics, we assume that systematic shocks to growth follow a normal mixture distribution,

$$\epsilon_{g,t+1} = \begin{cases} \epsilon^-_{g,t+1} \sim \mathcal{N}\left(\mu_g^-, {\sigma_{g_1}^-}^2\right) & \text{with probability } p_g, \\ \epsilon^+_{g,t+1} \sim \mathcal{N}\left(\mu_g^+, {\sigma_{g_1}^+}^2\right) & \text{with probability } 1-p_g, \end{cases} \tag{9}$$

where $p_g < 0.5$ is the probability of a recession, as $\mu_g^+ > \mu_g^-$. The shocks $\epsilon_{r,t+1}$ and $\epsilon_{g,t+1}$ are independent.

[5] In an asset pricing model, one would typically think of interest rates as depending on growth, not the other way around. Econometrically, though, these are equivalent representations and therefore without loss of generality.

3.2 Asset returns

Households can invest their wealth in long-term bonds, stocks, and housing. Henceforth, we continue to denote log returns by lowercase $r = \log R$.

Long-term bonds We model the long-term bond as a riskless claim to one unit of real consumption in n periods. Its price, denoted P_{nt}, satisfies the expectations hypothesis, generalized to include constant term premia. Specifically, 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 C.1, these assumptions imply an explicit relation between the dynamics of long-term bond returns and short-term rate fluctuations: the log bond return equals

$$r_{n,t+1} = r_{ft} + \mu_n - \sigma_n \epsilon_{r,t+1}, \tag{10}$$

where the sensitivity to rate shocks σ_n is given by

$$\sigma_n = \frac{1 - \varphi^{n-1}}{1 - \varphi} \sigma_r. \tag{11}$$

In addition, we set $\mu_n = -\sigma_n^2/2$, so that there is no risk premium.[6]

Recall that the interest-rate sensitivity of the bond (or any asset) is the percentage change in the price caused by an unexpected decline in the interest rate, which equals

$$\varepsilon_r(P_{nt}) \equiv -\frac{\partial \log P_{nt}}{\partial r_{ft}} = \frac{1 - \varphi^n}{1 - \varphi}. \tag{12}$$

This sensitivity is increasing in maturity n. This expression summarizes the effect of unexpected changes in interest rates: if the riskfree rate unexpectedly falls, then the long-term bond has an unexpectedly high return from capital gains next period. The longer is the maturity n, and the higher is the persistence φ of the rate shock, the larger is this response.

[6] This is not inconsistent with an upward-sloping term structure for nominal bonds, as nominal bonds may have a risk premium due to inflation risk. In contrast, the term premium on inflation-indexed bonds (TIPS) is likely close to zero: Ang et al. (2008) find that nearly all the slope of the nominal term structure is explained by inflation risk, whereas the real term structure is on average flat. Indeed, many leading asset pricing models predict a flat or downward-sloping real term structure (Bansal and Yaron, 2004; Wachter, 2013). As such, this assumption accords well with prior work.

Stock market The stock market is a claim to public firms' cashflows. Appendix C.2 shows that, if cashflows are cointegrated with aggregate income, the stock's log return can be written

$$r_{s,t+1} = r_{ft} + \mu_s - \lambda_{sr}\epsilon_{r,t+1} + \lambda_{sg}\epsilon_{g,t+1} + \epsilon_{s,t+1}, \tag{13}$$

where the market-specific shock is

$$\epsilon_{s,t+1} = \begin{cases} \epsilon^{-}_{s,t+1} \sim \mathcal{N}\left(0, \sigma_s^{-2}\right) & \text{if } \epsilon_{g,t+1} = \epsilon^{-}_{g,t+1}, \\ \epsilon^{+}_{s,t+1} \sim \mathcal{N}\left(0, \sigma_s^{+2}\right) & \text{if } \epsilon_{g,t+1} = \epsilon^{+}_{g,t+1}. \end{cases} \tag{14}$$

The market-specific shocks depend on the realization of a recession, capturing kurtosis and heteroskedasticity in returns, specifically the fact that returns are more volatile in in recessions ($\sigma_s^{-2} > \sigma_s^{+2}$). The loading λ_{sr} summarizes the net effect of interest-rate shocks on returns. On the one hand, a fall in rates will increase the value of the market, because it is a long-duration asset; but on the other hand, if growth and interest rates are positively correlated ($\lambda_{gr} > 0$), then this will also tend to come with bad news about cashflows.[7] We will find that the former dominates, so $\lambda_{sr} > 0$.

Housing The price of one unit of housing, denoted P_{it}, is the product of a regional (i.e., systematic) component P_t and a house-specific (i.e., idiosyncratic) component $\tilde{P}_{it}$.[8] The regional component can be decomposed as

$$P_t = P_{ht} P_{n_h t}, \tag{15}$$

where $P_{n_h t}$ is the price of an n_h-period bond, matching the interest-rate sensitivity (and duration) of housing; and where

$$\log \frac{P_{h,t+1}}{P_{ht}} = \mu_h + \sigma_h \epsilon_{h,t+1}, \tag{16}$$

captures the drift of house prices over time with $\epsilon_{h,t+1}$ standard normal. The idiosyncratic component $\tilde{P}_{it}$ evolves as a random walk with Gaussian innovations $\tilde{\sigma}_h \tilde{\epsilon}_{hi,t+1}$.

[7] This implies the imperfect passthrough of interest rates to equities emphasized by Gormsen and Lazarus (2025).
[8] Piazzesi et al. (2007) show that about half of the volatility of housing returns in the data is idiosyncratic.

For an owner-occupier, the return on a house is then[9]

$$r_{hi,t+1} = \underbrace{r_{ft} + \mu_{n_h} - \sigma_{n_h}\epsilon_{r,t+1}}_{r_{n_h,t+1}} + \underbrace{\mu_h + \sigma_h\epsilon_{h,t+1}}_{\Delta \log P_{h,t+1}} + \underbrace{\tilde{\sigma}_h\tilde{\epsilon}_{hi,t+1}}_{\Delta \log \tilde{P}_{i,t+1}}, \tag{17}$$

which includes both the return on the long-term bond component and the sum of the regional and idiosyncratic price appreciation.

3.3 Labor income

We model labor-income dynamics using the empirically realistic process estimated by Catherine (2022) from the data of Guvenen et al. (2022). Each household i earns labor income L_{it}, which is the product of the aggregate wage index $\bar{L}_t$ and an idiosyncratic component $\tilde{L}_{it}$. The aggregate wage grows at the rate

$$\log \frac{\bar{L}_{t+1}}{\bar{L}_t} = g_{t+1}, \tag{18}$$

defined in (7). The idiosyncratic component equals

$$\tilde{L}_{it} = \exp\left\{\ell(a_{it}) + z_{it} + \eta_{it}\right\}. \tag{19}$$

The deterministic component $\ell(a_{it})$ is a polynomial of age a_{it}; it captures the common life-cycle profile of income. The persistent component of earnings, denoted by z_{it}, follows a first-order autoregression

$$z_{it} = \rho z_{i,t-1} + \zeta_{it}, \tag{20}$$

with innovations ζ_{it} drawn from a mixture of normal distributions

$$\zeta_{it} = \begin{cases} \zeta_{it}^{-} \sim \mathcal{N}(\mu_{zt}^{-}, \sigma_z^{-2}) & \text{with probability } p_z, \\ \zeta_{it}^{+} \sim \mathcal{N}(\mu_{zt}^{+}, \sigma_z^{+2}) & \text{with probability } 1 - p_z. \end{cases} \tag{21}$$

Idiosyncratic labor-income growth tends to be negatively skewed, so the state $\zeta_{it} = \zeta_{it}^{-}$ is rare ($p_z < 0.5$) and tends to be negative. This skewness is also cyclical (more negative in recessions),

[9] In Appendix C.3, we derive this return and discuss how it accounts for maintenance and transaction costs.

so the mean outcome co-moves with aggregate income growth[10]

$$\mu_{zt}^{-} = \bar{\mu}_{z}^{-} + \lambda_{zg} g_t, \tag{22}$$

with $\lambda_{zg} > 0$. This state is also more volatile ($\sigma_z^- \gg \sigma_z^+$), meaning that extreme income events happen more frequently in this rare state. The initial cross-sectional distribution of the persistent component of earnings is given by $z_{i0} \sim \mathcal{N}(0, \sigma_{z0}^2)$.

The transitory component of idiosyncratic earnings η_{it} is also drawn from a mixture of normal distributions, with outcomes contingent on the permanent-income shock:

$$\eta_{it} \sim \begin{cases} \mathcal{N}(\mu_\eta^-, \sigma_\eta^{-2}) & \text{if } \zeta_{it} = \zeta_{it}^-, \\ \mathcal{N}(\mu_\eta^+, \sigma_\eta^{+2}) & \text{if } \zeta_{it} = \zeta_{it}^+. \end{cases} \tag{23}$$

3.4 Social Security and income taxes

Social Security Agents pay Social Security payroll taxes T_{it} on their labor income during working life, then receive benefits B_{it} in retirement. We assume all workers retire at the full-retirement age $a_{t_{\text{ret}}}$, 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:

$$T_{it}^{\text{SS}} = 0.106 \min\{L_{it}, 2.5\bar{L}_t\}. \tag{24}$$

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

$$L_{it}^{\text{indexed}} = \min\{L_{it}, 2.5\bar{L}_t\} \frac{\bar{L}_{t_{60}}}{\bar{L}_t} \tag{25}$$

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

[10]The mean of the other state, μ_{zt}^+, is then such that average shocks to z_{it} are zero: $0 = p_z \mu_{zt}^- + (1 - p_z)\mu_{zt}^+$.

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:

$$B_{it} = \begin{cases} 0.9\text{AIYE}_{it_{\text{ret}}} & \text{if AIYE}_{it_{\text{ret}}} < b_{1t}, \\ 0.9b_{1t} + 0.32(\text{AIYE}_{it_{\text{ret}}} - b_{1t}) & \text{if } b_{1t} \leq \text{AIYE}_{it_{\text{ret}}} < b_{2t}, \\ 0.9b_{1t} + 0.32(b_{2t} - b_{1t}) + 0.15(\text{AIYE}_{it_{\text{ret}}} - b_{2t}) & \text{if } b_{2t} \leq \text{AIYE}_{it_{\text{ret}}}. \end{cases} \tag{26}$$

The kinks in this benefit formula are determined by the "bend points" b_{1t} and b_{2t}, 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 indexed earnings as:

$$\text{AIYE}_{it} = \sum_{s=t_0}^{t} \min\{L_{is}, 2.5\bar{L}_s\} \frac{\bar{L}_t}{\bar{L}_s} = \bar{L}_t \sum_{s=t_0}^{t} \min\{\tilde{L}_{is}, 2.5\}. \tag{27}$$

Safety net Households can receive SNAP benefits ("food stamps") if their income falls sufficiently low. Under SNAP, households receive a transfer equal to 6% of the wage index minus 30% of their pre-tax earnings, if this difference is positive.

Income tax 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. Marginal tax rates are progressively increasing in idiosyncratic income $\tilde{L}_i$; we report the formula in Appendix D.3.

3.5 Households

Objective and preferences Household i chooses goods consumption C_{it}, housing consumption H_{it}, and portfolio shares $\pi_{it} = \{\pi_{n,it}, \pi_{s,it}, \pi_{h,it}\}$ to maximize lifetime utility

$$V_{it} = \max_{\{C_{is}, H_{is}, \pi_{is}\}} \mathbb{E}_t \sum_{s=t}^{t_{\max}} \beta^{s-t} p_{it,s-1} \left[(1 - m_{i,s-1}) u(C_{is}, H_{is}) + m_{i,s-1} b_s(W_{is}) \right], \tag{28}$$

where β is the rate of time preference, $t_{\max}$ corresponds to the maximum lifespan, m_{it} is the age- and income-dependent mortality probability from t to $t+1$, and $p_{its} = \prod_{u=t}^{s-1}(1-m_{iu})$ is the probability of surviving from t to s.

The household's flow utility is given by

$$u(C_{it}, H_{it}) = \frac{1}{1-\gamma}(C_{it}^{1-\nu} H_{it}^{\nu})^{1-\gamma}. \tag{29}$$

The parameter ν governs the preference for housing relative to goods; γ is both the coefficient of relative risk aversion and the inverse of the elasticity of intertemporal substitution.

Households also 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. Fluctuating interest rates therefore affect the value of bequests by changing the price at which future consumption can be purchased by the bequeathed wealth. Hence, we model the bequest motive as the expected utility from distributing the household's terminal wealth W_{it} over the next $\bar{b}$ years. In particular, the bequest utility is

$$b_t(W_{it}) = b(\bar{L}_t, P_{ht}, W_{it}, r_{ft}, g_{1t}) = \max_{\{C_{is}, H_{is}, \pi_{is}\}} \mathbb{E}_t \sum_{s=t}^{t+\bar{b}} \beta^{s-t} u(C_{is}, H_{is}) \tag{30}$$

subject to the budget constraints (below) and assuming that no labor income or Social Security benefits are earned over this period.

Budget constraints and frictions Households cannot short-sell stocks or housing, so portfolio shares must be non-negative. If the household decides to invest a positive amount in the stock market, then it must pay both a 1% management fee and a fixed, per-period participation cost of

$$\Phi_{s,it} = \begin{cases} c_s \bar{L}_t & \text{if } \pi_{s,it} > 0, \\ 0 & \text{otherwise.} \end{cases} \tag{31}$$

Households also cannot borrow to invest in the stock market, meaning the equity share $\pi_{s,it} \leq 1$.

The household must also decide whether to rent or own a house. If the household chooses to

rent ($\pi_{h,it} = 0$), then the rental price per unit of housing equals

$$X_t = \begin{cases} \chi P_t & \text{if } \pi_{h,it} = 0 \text{ (renter)}, \\ 0 & \text{if } \pi_{h,it} > 0 \text{ (homeowner)}, \end{cases} \tag{32}$$

for rent-to-price ratio χ. A renter does not have any collateral with which to borrow, and therefore must have $W_{it} \geq 0$.

If instead the household chooses to own ($\pi_{h,it} > 0$), then, as in Cocco (2005), the size of the house must exceed a minimum size, given by

$$P_t H_{it} \geq \kappa_{\min} \bar{L}_t. \tag{33}$$

Homeowners are also constrained to buy a house that is not too large:

$$P_t H_{it} \leq \kappa_{\max} W_{it} \tag{34}$$

This constraint is equivalent to a downpayment requirement: the household must have enough wealth W_{it} to cover a fraction $1/\kappa_{\max}$ of the house's value. Homeowners pay proportionate property taxes. Finally, homeowners can use their housing and equity as collateral for a mortgage. That is, if their investments in housing and equity exceed their networth, $\pi_{h,it} + \pi_{s,it} > 1$, then they must finance this with a mortgage that charges interest $R_{Mt} = e^{\theta} R_{ft}$ at a premium $\theta > 0$.

Putting these constraints together, the utility maximization (28) is solved subject to the dynamic budget constraint

$$W_{i,t+1} = (W_{it} + L_{it} + B_{it} - T_{it} - C_{it} - X_t H_{it} - \Phi_{s,it}) R_{Wi,t+1}, \tag{35}$$

where the return on savings equals

$$\begin{aligned} R_{Wi,t+1} = R_{ft} &+ \pi_{n,it}(R_{n,t+1} - R_{ft}) + \pi_{s,it}(R_{s,t+1} - R_{ft}) \\ &+ \pi_{h,it}(R_{h,t+1} - R_{ft}) - (\pi_{h,it} + \pi_{s,it} - 1)^{+}(R_{Mt} - R_{ft}) \end{aligned} \tag{36}$$

for $(\pi_{h,it} + \pi_{s,it} - 1)^+ \equiv \max(\pi_{h,it} + \pi_{s,it} - 1, 0)$.

3.6 Entrepreneurs

Some households own and manage a private business instead of working in the labor market. In particular, with some probability $p_E(a_{it})$, a worker of age a_{it} will start a private business and operate it through working life. If this occurs, then the entrepreneur will no longer earn labor income according to (19); rather, he or she will earn idiosyncratic business income according to

$$\tilde{L}_{it} = \exp\left\{\ell_E(a_{it}) + z_{it}\right\}, \tag{37}$$

which is the sum of a polynomial of age $\ell_E(a_{it})$ and a random walk

$$z_{i,t+1} = z_{it} + \sigma_{Ez}\epsilon_{z,t+1}. \tag{38}$$

At the beginning of the life cycle, the initial value of z_{it} is drawn $N(0, \sigma^2_{Ez0})$.

The entrepreneur operates the business until retirement or death—whichever is first—at which point the business is sold on the market and liquidated into market wealth at a value E_{it}. In Appendix C.4, we show that the value of the business can be expressed as

$$E_{it} = L_{it} \exp\left\{\upsilon + \frac{1}{1 - \kappa_E(\upsilon)\varphi_g}(g_{1t} - \bar{g}_1) - \frac{1 - \lambda_{gr}}{1 - \kappa_E(\upsilon)\varphi}(r_{ft} - \bar{r}_f)\right\}, \tag{39}$$

where υ is the mean log valuation ratio (a function of mean discount rates and cashflow growth) and $\kappa_E(\upsilon) \equiv 1/(1 + \exp\{-\upsilon\})$. All business valuation ratios fluctuate over time with predictable variation in interest rates r_{ft} and growth g_{1t}. Note that the value (39) also corresponds to the market value of the business at any time before retirement.[11]

[11] In our data, business owners are asked to report estimates of the value of their businesses if they were to sell them, so these values indeed represent estimates of public market values.

4 Economic intuition

To communicate the key intuitions of our model, we present an analytical solution to a linearized version with no income risk or bequests, which we derive in Appendix G. We proceed in three steps. First, we discuss consumption and portfolio rules for an agent without human capital or Social Security, trading in only a long- and a short-term bond. Second, we show that the same portfolio rule applies to *total* wealth in the presence of background assets such as human capital and Social Security. And third, we explain why these results continue to hold even when we introduce other tradable assets (e.g. equity, real estate) and risk factors (e.g., growth, house prices).

4.1 Optimal choices without labor income

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

$$\frac{C_{it}^*}{W_{it}} = \underbrace{(1-\beta(1-m_{it}))}_{\text{time discounting}} \times \underbrace{\exp\left\{\left(1-\frac{1}{\gamma}\right)(\varrho_{0t}+\varrho_{rt}r_{ft})\right\}}_{\text{income and substitution effects}}. \tag{40}$$

The first term represents the positive effect of impatience and mortality on consumption. The second term represents the net of income and substitution effects from interest rates. Higher rates mean higher interest income, so that households can consume more today (the income effect). At the same time, higher rates mean agents get more consumption tomorrow in exchange for their savings (the substitution effect). The income effect dominates the substitution effect when the elasticity of intertemporal substitution (the EIS, $1/\gamma$) is less than one ($\gamma > 1$). The sensitivity of consumption to interest rates depends on the coefficient[12]

$$\varrho_{rt} = \sum_{j=1}^{\infty} \varphi^{j-1}\beta^j p_{it,t+j}, \tag{41}$$

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

[12] We omit the subscript on ϱ_{rit} to simplify notation and to emphasize that it is homogeneous within a cohort.

Portfolio rule The optimal allocation to the n-period bond is:[13]

$$\pi_{it}^* = \underbrace{\frac{1}{\gamma}\frac{\mu_n + \frac{1}{2}\sigma_n^2}{\sigma_n^2}}_{\text{myopic demand}} + \underbrace{\left(1 - \frac{1}{\gamma}\right)\varrho_{rt}\left(\frac{1-\varphi^{n-1}}{1-\varphi}\right)^{-1}}_{\text{interest-rate risk hedging demand}}. \tag{42}$$

The first term represents the traditional risk-return tradeoff of Merton (1969). Our assumption of $\mu_n = -\sigma_n^2/2$ sets this term to zero. The second term is the demand from intertemporal hedging of interest-rate fluctuations, the focus of our paper. Because its value increases when rates unexpectedly decline, the long-term bond offers protection against the deterioration of investment opportunities. This portfolio rule implies that the optimal interest-rate sensitivity of wealth equals[14]

$$\varepsilon_r^*(W_{i,t+1}) = \pi_{it}^*\frac{\sigma_n}{\sigma_r} = \frac{1}{\gamma}\frac{\mu_n + \frac{1}{2}\sigma_n^2}{\sigma_n\sigma_r} + \left(1 - \frac{1}{\gamma}\right)\varrho_{rt}. \tag{43}$$

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

The sensitivity of consumption to rate shocks ϱ_{rt} declines with the investor's horizon, so the hedging demand decreases in age toward zero. From Equation (41), we see that ϱ_{rt} represents the cumulative effect of an interest-rate shock over time, weighted by the importance of each future period in the agent's lifetime utility, taking into account mortality and impatience. In other words, it results from the interplay between the agent's horizon and the persistence of interest-rate shocks. Agents with shorter horizons are less sensitive to future rate changes; therefore, for $\gamma > 1$, the optimal interest-rate sensitivity of wealth declines over the life cycle.

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

[13] As we verify in Appendix G.2, (42) holds even if we separate the coefficient of relative risk aversion from the elasticity of intertemporal substitution. Thus, the portfolio share is indeed governed by risk aversion, not the EIS.

[14] Notice that the agent targets the interest-rate sensitivity of next period's wealth, $W_{i,t+1}$, because that is what will respond to the interest-rate shock $\epsilon_{r,t+1}$ next period. This distinction disappears as the time interval becomes small.

4.2 Substitution effects of labor income and Social Security

Now let us consider the effect of labor income and Social Security. Suppose that labor income L_{it}, taxes T_{it}, and benefits B_{it} are deterministic. The value of human capital, H_{it}, is the present value of future earnings discounted using the prevailing yield curve and Social Security wealth, S_{it}, is the present value of future benefits net of future payroll taxes. Define total wealth $\overline{W}_{it} = W_{it}+H_{it}+S_{it}$ as the sum of wealth and these present values of these background assets.

Implementing the same linearization implies that 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 equation (40). Similarly, the optimal allocation to bonds out of total wealth is $\bar{\pi}_i = \pi_i^*$ from equation (42). The optimal allocation out of wealth W therefore takes the form

$$\pi_{it} = \pi_{it}^* - \underbrace{(\pi_{it}^H - \pi_{it}^*)\frac{H_{it}}{W_{it}}}_{\substack{\text{human capital}\\ \text{substitution effect}}} - \underbrace{(\pi_{it}^S - \pi_{it}^*)\frac{S_{it}}{W_{it}}}_{\substack{\text{Social Security}\\ \text{substitution effect}}}, \tag{44}$$

or, in terms of interest-rate sensitivities,

$$\varepsilon_r(W_{i,t+1}) = \varepsilon_{r,it}^* - (\varepsilon_r(H_{i,t+1}) - \varepsilon_{r,it}^*)\frac{H_{it}}{W_{it}} - (\varepsilon_r(S_{i,t+1}) - \varepsilon_{r,it}^*)\frac{S_{it}}{W_{it}}, \tag{45}$$

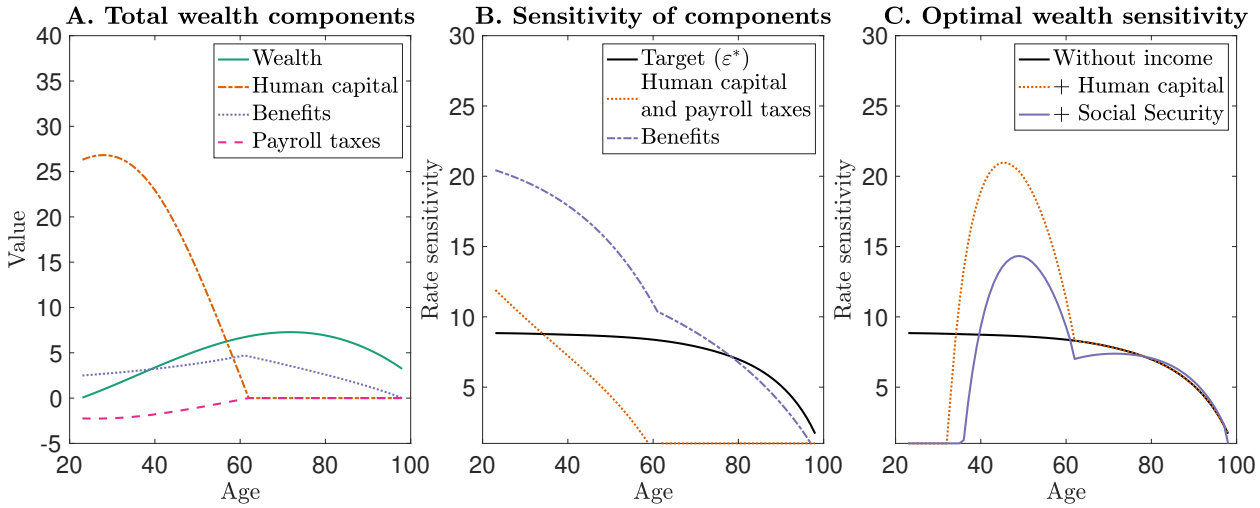
such that the interest-rate sensitivity of total wealth $\varepsilon_r(\overline{W}_{i,t+1})$ remains equal to $\varepsilon_{r,it}^*$.

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

Figure 5 illustrates the life-cycle pattern generated by this model. Early in life, most agents have little financial wealth and a large endowment of high-duration human capital. To match their ideal total-wealth rate exposure, they mostly hold assets with low rate sensitivity. As households get closer to retirement, they increase holdings of the long-term asset to offset short-term labor income and taxes, net of long-term benefits. As they progress through retirement, households

reduce long-term bond holdings, in line with the declining target allocation implied by the policies above. In sum, substitution and aging effects explain the hump-shaped pattern in the data.

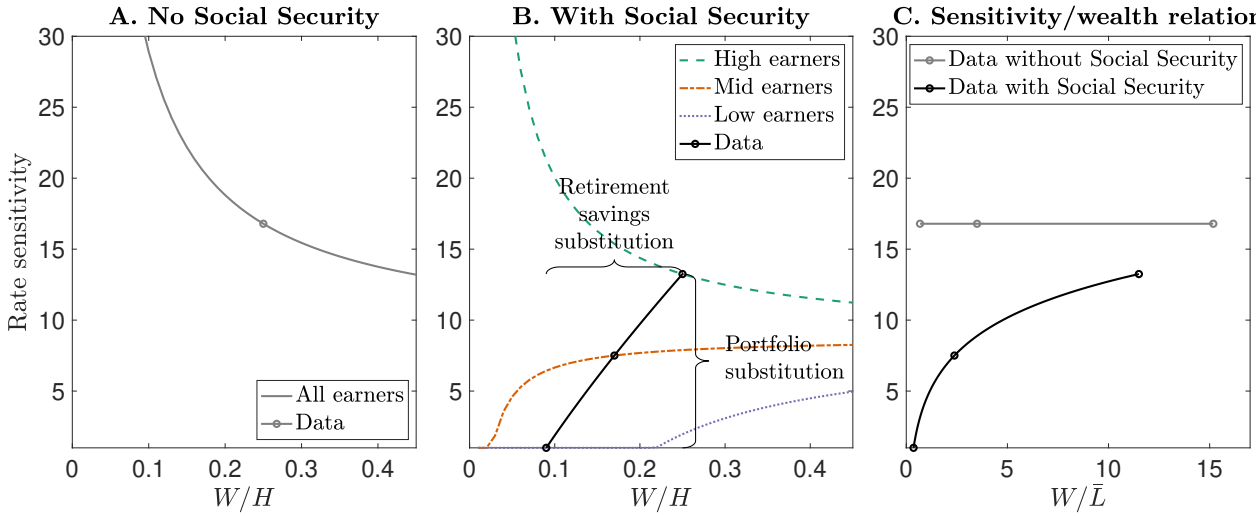
Figure 5: Effect of labor income and Social Security on long-term asset share



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

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

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



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

4.3 Adding other assets and sources of systematic risk

All of the above intuition generalizes to multiple risk factors and multiple assets with different exposures to interest-rate risk. To see this, suppose that there are J risky assets and J risk factors—for example, the long-term bond, stock market, and housing. Letting $\varepsilon \sim N(0, I)$ denote a $J \times J$ vector of independent shocks, we can write the interest-rate process as

$$r_{f,t+1} = (1 - \varphi)\bar{r}_f + \varphi r_{ft} + \sigma_r^\top \varepsilon_{t+1} \tag{46}$$

and write the J-dimensional vector of risky-asset returns as

$$r_{t+1} = r_{ft}\iota + \mu + \Sigma^\top \varepsilon_{t+1}. \tag{47}$$

The only assumption we require is that these assets are non-redundant (the $J \times J$ matrix Σ is invertible). Appendix G.5 shows that, under this assumption, the vector of optimal portfolio weights

on the risky assets (in total wealth) equals

$$\pi_{it}^{*} = \frac{1}{\gamma}(\Sigma^{\top}\Sigma)^{-1}\left(\mu + \frac{1}{2}\text{diag}(\Sigma^{\top}\Sigma)\right) - \left(1 - \frac{1}{\gamma}\right)\varrho_{rt}\Sigma^{-1}\sigma_r, \tag{48}$$

which, like before, is the sum of the myopic demand and the rate-hedging demand for each asset. If asset j is the long-term bond, then $\pi_t^*(j)$ still equals (42). The vector of interest-rate sensitivities across assets is then given by

$$\varepsilon_r = -\frac{\Sigma^{\top}\sigma_r}{\sigma_r^{\top}\sigma_r}, \tag{49}$$

so the interest-rate sensitivity of wealth is

$$\varepsilon_r^*(W_{i,t+1}) = \pi_{it}^{*\top}\varepsilon_r = \underbrace{-\frac{1}{\gamma}\frac{\Lambda^{\top}\sigma_r}{\sigma_r^{\top}\sigma_r}}_{\text{risk premia}} + \underbrace{\left(1 - \frac{1}{\gamma}\right)\varrho_{rt}}_{\text{hedging}}, \tag{50}$$

where

$$\Lambda = (\Sigma^{\top})^{-1}\left(\mu + \frac{1}{2}\text{diag}(\Sigma^{\top}\Sigma)\right) \tag{51}$$

is a vector of factor loadings ("Sharpe ratios"): the prices of risk for exposure to each risk factor. If $\Lambda^{\top}\sigma_r = 0$, then there is no additional compensation for holding interest-rate risk, and (50) takes the exact same value as in the two-asset model above.[15]

Intuitively, adding additional systematic risk factors has no consequence on the overall optimal interest-rate sensitivity of wealth. This is essentially an application of the portfolio separation theorem, which states that, no matter the specific menu of assets, households should choose the same overall exposure to risk factors—exposure to interest-rate risk, house-price risk, etc.—which can be formed by a set of "separating portfolios" (often called "funds" in the finance literature).

4.4 Real-life interpretation

Thus far, we have built intuition by representing the interest-rate exposure of household balance sheets with two zero-coupon securities of different maturities. We now turn to a more concrete explanation of how households can construct real-world portfolios that achieve an interest-rate

[15]There still may be risk premia on every individual asset when $\Lambda^{\top}\sigma_r = 0$, but these premia come from exposure to the other $J - 1$ risk factors. For a long-term nominal bond, for example, this would include an inflation risk premium.

sensitivity close to the optimal portfolio rule.

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

We can illustrate this intuition most clearly in the limiting case of an investor with infinite risk aversion and a zero EIS.[16] In this case, the investor's desire to smooth consumption over states and time yields a constant, deterministic policy $C_{it} = \bar{C}_i$. Let Y_i denote the agent's deterministic stream of income. Wealth is the present value of the excess consumption plan:

$$W_{it} = \sum_{k=1}^{t_{\max}} P_{kt}(\bar{C}_i - Y_{i,t+k}). \tag{52}$$

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 G.6, the optimal allocation π_i replicates exactly this buy-and-hold strategy as $\gamma \to \infty$.

Real-world implementation In reality, households do not hold portfolios of zero-coupon bonds. Instead, they implement this strategy using the assets and contracts available to them. We illustrate this mechanism with three hypothetical households, each represented in our full model.

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

Second, consider a middle-class worker. Because replacement rates fall with lifetime earnings, the worker needs to save privately as well. However, he can execute large intertemporal transfers at current prices by buying a house with a fixed-rate mortgage. By doing so, he effectively trades a flow of coupon payments later in life, when $C > Y$, in the form of rent-free housing, in ex-

[16] See Appendix G.6 for a derivation and more detailed technical discussion of this case. In that appendix, we prove that this case requires both infinite risk aversion and a zero EIS. Intuitively, in order for an agent to desire perfect hedging, he or she must be unwilling to substitute consumption across both states (risk aversion) and time (EIS).

change for a stream of mortgage payments earlier in life, when $C < Y$. This strategy eliminates interest-rate risk for workers whose Social Security benefits cover non-residential consumption in retirement.

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

4.5 Distributional effects of interest-rate fluctuations

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

As we showed above, the rate elasticity of total wealth is $\varepsilon_r(\overline{W}_{i,t+1}) = \varepsilon^*_{r,it}$, which depends only on risk aversion and the agent's investment horizon ϱ_{rt}. In other words, for any two agents with the same risk aversion and horizon, a rate shock will have an identical effect on total wealth, even if those agents have different long-term asset shares. Likewise, the optimal consumption rule (40) implies that the response of consumption to a rate shock $\varepsilon_r(C^*_{it})$ is also homogeneous within a cohort.

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

expected utility:

$$U_{it} = ((1-\gamma)V_{it})^{1/(1-\gamma)}, \tag{53}$$

where V_{it} is the expected utility maximand (28). 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.[17] In our linearized model, lifetime utility has the closed-form solution

$$U_{it} = \overline{W}_{it}(1-\beta(1-m_{it}))\exp\left\{\varrho_{0t} + \varrho_{rt} r_{ft}\right\}. \tag{54}$$

This expression illustrates the complementarity between total wealth and interest rates: agents have high expected utility when they enjoy high wealth with high rates. The rate sensitivity of lifetime utility is therefore[18]

$$\varepsilon_r(U_{i,t+1}) = \underbrace{-\frac{\partial \log U_{i,t+1}}{\partial r_{f,t+1}}}_{-\varrho_{r,t+1}} + \underbrace{\frac{\partial \log U_{i,t+1}}{\partial \log \overline{W}_{i,t+1}} \varepsilon_r^*(\overline{W}_{it})}_{(1-1/\gamma)\varrho_{rt}} \approx -\frac{1}{\gamma}\varrho_{rt}. \tag{55}$$

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

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

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

[18] This expression remains unchanged when separating risk aversion from the EIS. The approximate equality becomes exact as the time interval becomes small.

5 Model calibration

We calibrate the model to macroeconomic time series and cross-sectional household data from the SCF. In this section, we exclude entrepreneurs from both the model and data. We report the calibration and model fit including entrepreneurs in Appendices D.4 and H.1, respectively.

5.1 Asset markets and aggregate growth

Table 1 reports estimates of the interest-rate process. We estimate the processes for stock returns and growth jointly via simulated method of moments. Appendix D.1 explains the estimation procedure and Tables D.1 and D.2 report the parameter estimates and moments. Historically, economic growth has been positively exposed to interest-rate shocks ($\lambda_{gr} > 0$), while market returns have been negatively exposed ($\lambda_{sr} > 0$). Not surprisingly, the market crashes in recessions ($\lambda_{sg} > 0$).

For housing returns (17), the drift rate is taken to be $\mu_h = -0.0105$, which matches the historical average housing return (net of maintenance costs) of 6.1% estimated by Jordà et al. (2019). The regional house price volatility is $\sigma_h = 0.07$ and the idiosyncratic volatility is $\tilde{\sigma}_h = 0.12$, as implied by estimates from Flavin and Yamashita (2002) and Piazzesi and Schneider (2016).[19] For a renter, the rent-to-price ratio χ is 5.33% (Jordà et al., 2019).

Households investing in these assets face the following costs and constraints. Equity investors pay a 1 percentage-point management fee and a 1.2 percentage-point capital gains tax.[20] The equity participation cost is $c_s = 0.02$ times the average wage. Homeowners pay a property tax rate of 1.4% on the value of their house. The maximum house size is $\kappa_{\max} = 5$ times wealth, which represents a down-payment constraint of 20%. The minimum house size is $\kappa_{\min} = 1.5$ times the average wage. The mortgage premium θ is 0.7%.

5.2 Preferences

We calibrate households' preferences to match the evolution of wealth over the life cycle and the average equity share observed in the SCF. We find that a discount factor of $\beta = 0.97$ and a bequest

[19] Flavin and Yamashita (2002) report a total house-price volatility of 0.14, which implies an idiosyncratic volatility $\tilde{\sigma}_h$ such that $\tilde{\sigma}_h^2 + \sigma_h^2 = 0.14^2$.

[20] This is implied by a 15% capital gains tax rate on a 8 percentage-point expected return.

motive equivalent to $\bar{b} = 10$ years of consumption match the growth of wealth until the retirement age and its evolution thereafter. See Panel A of Figure 7.

A coefficient of relative risk aversion of $\gamma = 6$ matches the average allocation of wealth to stocks. Our calibration of γ is consistent with other studies matching the life-cycle profile of the share of wealth invested in stocks, which typically use values between 5 and 6 (Benzoni et al., 2007; Catherine, 2022; Huggett and Kaplan, 2016; Lynch and Tan, 2011; Meeuwis, 2022). Based on portfolios in Swedish administrative data, Calvet et al. (2021) estimate an average γ of 5.2.

We set the Cobb-Douglas housing preference parameter $\nu = 0.25$, which is close to the aggregate housing expenditure share (Piazzesi and Schneider, 2016).

5.3 Other parameters

Income process We calibrate the stochastic parameters of the labor process .... in Appendix D.2.

Table 3: Estimated parameters: Idiosyncratic labor-income shocks

	Persistent shocks					Transitory shocks	
p_z	ρ_z	$\overline{\mu_z^-}$	λ_{zl}	σ_z^-	σ_z^+	σ_η^-	σ_η^+
.224	.993	-.026	5.219	.352	.059	.414	.124

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

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

$$m_{it} = \min\left\{\chi\left(\text{AIYE}_{it}/\bar{L}_t\right) \times \overline{m}_{it}, 1\right\}, \tag{56}$$

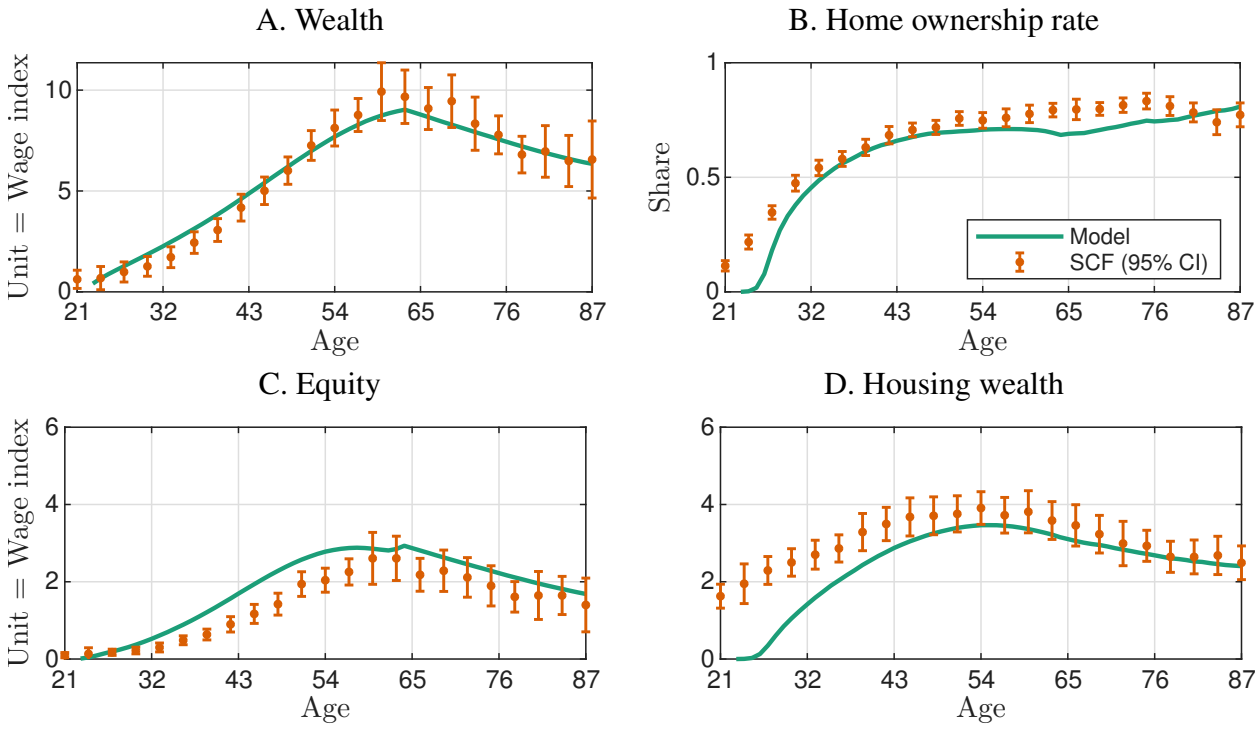
where χ is an adjustment coefficient which only depends on the average indexed earnings of the agent up to time t and $\overline{m}_{it}$ is the average mortality rate at his age, which we calibrate as the average across genders from the 2019 Social Security actuarial life tables. While $\chi\left(\text{AIYE}_{it}/\bar{L}_t\right)$ does not depend on age, the agent sees his life expectancy change as he moves up or down the

wage ladder. An advantage of our method is that the agent's life expectancy is less volatile than if it were a function of persistent income z_{it}. We calibrate the value of $\chi\left(\text{AIYE}_{it}/\bar{L}_t\right)$ at each point of the numerical grid of the $\text{AIYE}_{it}/\bar{L}_t$ state variable such that, given our labor-income process, we obtain the same life expectancy differential across percentiles of $\text{AIYE}_{it}/\bar{L}_t$ at age 40 as those reported by percentiles of earnings in Chetty et al. (2016).

5.4 Validation

We calibrate preference parameters (γ, β and b) to fit the trajectories of wealth and portfolio allocations over the life cycle. Figure 7 compares the model-implied wealth allocations to the data over the life cycle. The model captures the hump-shaped profiles of wealth accumulation, equity holdings, and conditional house value. The homeownership rate increases to a stable level over the life cycle.

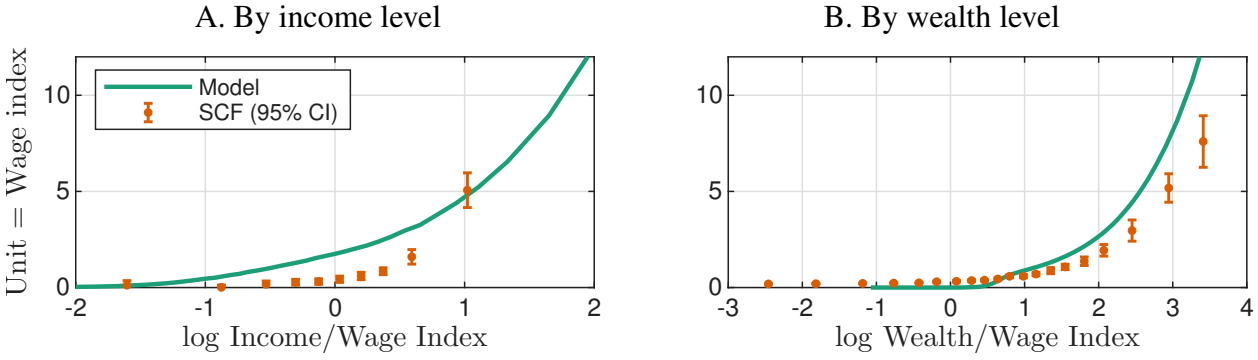
Figure 7: Wealth and portfolios over the life cycle



Note: This Figure reports wealth (Panel A), home ownership rates (Panel B), wealth invested in stocks (Panel C) and housing (Panel D) by age, in the model and in the data. Equity wealth, income and wealth are measured in units of the national wage index. Entrepreneurs are excluded. Appendix X presents the same figure including entrepreneurs.

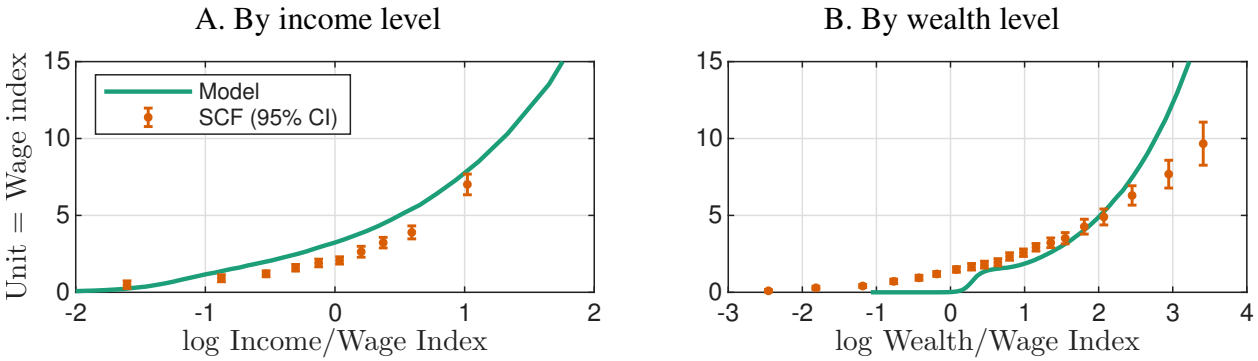
Figure 8 and 9 reports average holdings of equity and housing wealth respectively, as a function of income and wealth, for households between age 40 and 45. The model matches the fact that high-income and high-wealth households tend to invest more in equities and housing, not only in dollars but also in shares of wealth. Life-cycle models often struggle to explain why young and low-income investors invest so little in stocks. Our model is able to explain this primarily due to the presence of counter-cyclical labor income risk (Catherine, 2022).

Figure 8: Equity holdings at age 40-45



Note: This figure reports the amount of wealth invested in stocks in the model and in the data, for households between age 40 and 45, by level of income and wealth. Equity wealth, income and wealth are measured in units of the national wage index. Entrepreneurs are excluded. Appendix X presents the same figure including entrepreneurs.

Figure 9: Housing wealth at age 40-45



Note: This figure reports the amount of wealth invested in housing in the model and in the data, for households between age 40 and 45, by level of income and wealth. Housing wealth, income and wealth are measured in units of the national wage index. Entrepreneurs are excluded. Appendix X presents the same figure including entrepreneurs.

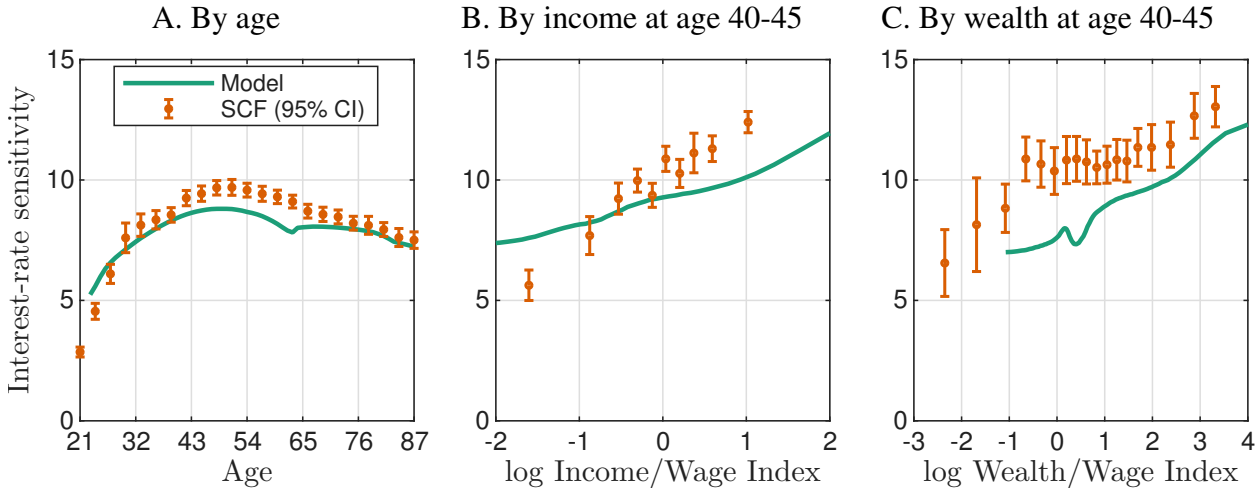
6 Matching the stylized facts

Although the model is calibrated to match other dimensions of household balance sheets, our primary interest remains the interest-rate sensitivity of wealth. We therefore now examine how well the model reproduces the stylized facts documented in Section 2.

6.1 Cross-section of interest-rate sensitivity

Panel A of Figure 10 reports the evolution of the interest-rate sensitivity of wealth over the life cycle, in the data and in the model. It 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, as explained in Section 4.2. The introduction of labor income risk in the model reduces the agent's valuation of his human capital, which explains why the predicted hump-shaped relationship with age is less pronounced than in Figure 5.

Figure 10: Interest-rate sensitivity in the cross-section



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.

Panel B reports the relationship between the interest-rate sensitivity of wealth and income between age 40–45. In the model, high earners invest more in the long-term asset because Social Security covers a smaller share of their retirement consumption and, to a lesser extent, because they have higher life expectancy. Appendix H.3 decomposes these two effects quantitatively.

Panel C 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.

6.2 Household interest-rate exposure

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

To calculate wealth inclusive of Social Security, we capitalize the expected benefits and taxes into a present value. To measure welfare, we calculate the transformed expected utility U defined in (53). Because U is a function of both wealth W and rates r_f, this elasticity can be approximated, to a first order, as

$$\varepsilon_r(U) \approx \underbrace{-\frac{\partial \log U}{\partial r_f}}_{\substack{\text{change in investment} \\ \text{opportunities}}} + \underbrace{\frac{\partial \log U}{\partial \log W}\varepsilon_r(W)}_{\text{capital gains}} \tag{57}$$

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

Figure 11: Interest-rate sensitivities for different measures of wealth



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

Panel A of Figure 11 shows the average paths of these sensitivities over the life cycle. Adding Social Security wealth increases the average sensitivity for the young, consistent with the fact that it is a very long-term asset.[21] The rate sensitivity of expected utility, on the other hand, is flatter over the life cycle. At all ages, households are negatively affected by rate declines on net. The magnitude of this effect is slowly declining over the life cycle as the investment horizon declines with age. For instance, $\varepsilon_r(U)$ is -2.7 at age 25 and -2.2 at age 65. The closeness of these numbers means that even different cohorts have relatively similar total exposure to rate risk.

Panels B and C report the distribution of these sensitivities within a middle-aged cohort, as functions of income and wealth. 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 why, when Social Security is accounted for and discounted using the market yield curve, wealth inequality has not increased since 1989 (Catherine et al., 2025).

Panels B and C are the model's counterparts to Figure 3. The data and the model convey the same message: once Social Security is accounted for, rich households no longer hold more interest-rate-sensitive assets. Like in the data, the interest-rate sensitivity of wealth, inclusive of Social Security, is around 15, though a bit higher for low-wealth households in the data.

[21] We only report the interest-rate sensitivity of Wealth+Social Security starting at age 25. Because Social Security wealth can be negative in high rate environments, we find that Wealth+Social Security can be negative or extremely close to zero in the first two years, leading to diverging values of ϵ_r on very small dollar amounts.

Second, within a cohort, expected utility is uniformly elastic to interest rates across the earnings and wealth distributions, save for a very minor effect from income-driven mortality differences. As equation (44) predicts, once human capital and Social Security wealth are accounted for, all households within a cohort are equally exposed to rate fluctuations.

7 Trends in wealth inequality and welfare

So far, we have demonstrated that, under a reasonable calibration, the interest-rate risk hedging demand for long-term assets can explain differences in the interest-rate sensitivity of wealth over the life cycle, and across the earnings and wealth distribution. Consequently, in the model, interest-rate fluctuations will, through our mechanism, redistribute wealth across the population without significant implications for the distribution of welfare or lifetime consumption.

In this section, we use an overlapping-generations version of our life-cycle model to study the historical distributions of wealth and welfare. First, we assess the extent to which interest-rate fluctuations explain the long-run fluctuations in wealth inequality. Second, we quantify the implications of these fluctuations for welfare.

7.1 Intuition: From rate shocks to inequality

Wealth evolves according to:

$$\frac{W_{i,t+1}}{W_{it}} = \underbrace{\left(1 - \frac{C_{it} - Y_{it}}{W_{it}}\right)}_{\text{savings}} \underbrace{R_{Wi,t+1}}_{\text{portfolio}}, \tag{58}$$

where Y_{it} summarizes disposable income. We can decompose changes in inequality over time by taking logs of (58) and then computing cross-sectional variances. Doing so yields the change in

wealth dispersion from one period to the next:[22]

$$\begin{aligned}\text{var}_I(w_{i,t+1}) - \text{var}_I(w_{it}) = \text{var}_I(s_{it}) + \text{var}_I(r_{wi,t+1}) \\ + 2\text{cov}_I(w_{it}, s_{it}) + 2\text{cov}_I(w_{it}, r_{wi,t+1}) + 2\text{cov}_I(s_{it}, r_{wi,t+1}). \quad (59)\end{aligned}$$

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

Consider how wealth inequality within a cohort responds to a negative interest-rate shock. 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 ($\text{var}_I(s_{it}) > 0$) and a positive wealth-savings correlation ($\text{cov}_I(w_{it}, s_{it}) > 0$). Second, Social Security gives rise to changes in inequality via its impact on portfolio choices. As Figure 6 illustrates, the substitution effects of Social Security create a positive correlation between a household's wealth and its interest-rate exposure. Thus, households within a cohort may experience different wealth returns, and the direction of reallocation will depend on the direction of the interest-rate shock. All unexpected rate changes result in heterogeneous returns ($\text{var}_I(r_{wi,t+1}) > 0$). A negative rate shock will result in disproportionately high returns for the wealthy ($\text{cov}_I(w_{it}, r_{wi,t+1}) > 0$), increasing inequality. Finally, since the wealthy save more, the savings-return covariance $\text{cov}_I(s_{it}, r_{wi,t+1})$ is also positive given a rate decline, amplifying the increase in inequality.

The same redistribution will not, in general, occur for welfare (lifetime utility) or for broader concepts of wealth that include human capital and Social Security.[23] This is because background assets offset interest-rate shocks; indeed, they are precisely the reason that households choose differential savings rates and portfolios to begin with.

[22] Lowercase letters denote logs; $s \equiv \log(1 - (C - Y)/W)$ denotes the log savings rate. We abstract from entry and exit, which keep the wealth distribution stationary on a balanced-growth path (i.e., absent aggregate shocks).

[23] We say "in general" because this will only be precisely true insofar as households are not constrained in making consumption and portfolio choices. We account for the presence of such constraints in our model.

7.2 Overlapping-generations simulation

To quantify the role of the interest-rate risk hedging channel on wealth inequality trends, we set up an overlapping-generations (OLG) version of our life-cycle model. Specifically, we simulate the lives of cohorts born since 1880, feeding in the history of macroeconomic shocks. The time series of interest rates is described in Appendix D.1.1; time series for growth, stocks, and house prices are as in the model calibration. We describe the OLG simulation in more detail in Appendix F

The only substantive assumption we must make for the OLG model is how to distribute bequests across generations. In the data, the size of inheritance is correlated with the income and wealth of the recipient. To capture this, we assume that the persistent component of income z_{it} is correlated across generations, via its initial value, and calibrate this correlation to the persistence documented in Chetty et al. (2014). Inheritance flows are unanticipated.

7.3 Implications for wealth inequality

We use our model to quantify the effect of each shock on wealth inequality, measured by the top 10% wealth share. To measure the effect of each shock, we simulate the model multiple times, starting from a counterfactual with no aggregate shocks and then sequentially add one shock from the data at a time. We start with a specification that includes data on interest rates and growth. The next specifications add stock-market shocks, followed by national house-price shocks. The only exception is that we do not add the growth shock to stock returns in the first specification, as to make clear the impact of stock returns in the data on wealth inequality in the second specification.

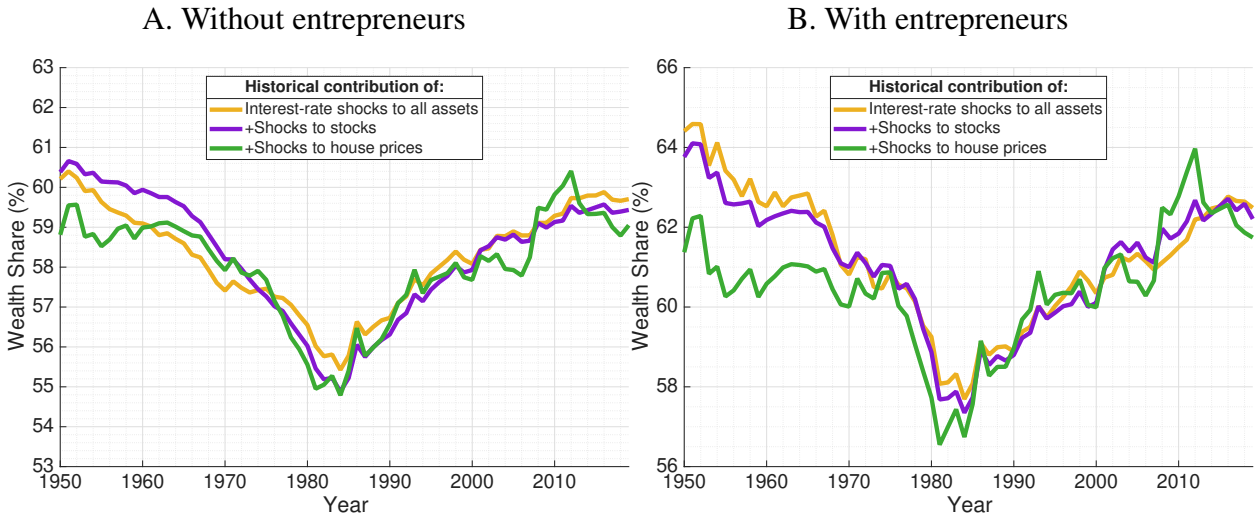
Figure 12 illustrates the model-implied inequality series. We consider the model both without entrepreneurs (Panel A) and with entrepreneurs (Panel B). The latter case is essential for explaining the level of top wealth shares in the data. The results of the simulation suggest that the majority of post-war variation in top wealth shares is explained by declining interest rates. From 1950–1980, the top 10% share declines by over 4 percentage points in the specification where only interest-rate shocks are included. The result is nearly identical when adding data on stock returns and house prices. Similarly, the subsequent rise in top wealth shares is approximately 4 percentage points and almost entirely driven by declining interest rates.

The specification with entrepreneurs looks quite similar, but explains a greater share of the

overall fluctuations in the top wealth shares since 1950. Here, interest-rate fluctuations implied 6 percentage points (pp) of the decline in top wealth shares from 1950–1980 and 4.5 pp of the subsequent rise. Adding housing puts a damper on this effect, but the model still generates a 5 pp decline and nearly 6 pp rise in top wealth shares pre- and post-1980.

How does this compare to the data? According to the World Inequality Database, the top 10% share fell from 70.3% in 1962 to 62.1% in 1985, then rose back to 70.7% in 2019. Consequently, the interest-rate risk hedging channel explains roughly two-thirds of the long-run fluctuations in wealth inequality over this period.

Figure 12: Evolution of the top 10% wealth share



Notes: This figure shows the evolution of the top 10% wealth share in our model using historical asset returns. Panel A reports the top 10% share excluding business owners; Panel B includes them. The "interest-rates only" specification includes the historical time series for economic growth and interest rates, along with shocks to interest rates ϵ_r. The other stochastic components of returns (ϵ_g, ϵ_s, ϵ_h) are drawn from random distributions. Relative to this specification, the "stock market returns" specification adds historical data on the remaining shocks that affect equity returns (ϵ_s, ϵ_g). The "house prices" specification instead combines aggregate national house price shocks with idiosyncratic shocks drawn from a random distribution. Appendix F details how we infer the historical values of these shocks.

Clearly, most of the long-run fluctuations in wealth inequality are explained by interest-rate fluctuations. After accounting for the effect of interest rates on stock and house prices, these asset classes prove to be relatively less important for explaining inequality. This is consistent with the finding of Greenwald et al. (2023), who conduct a similar exercise, taking portfolio choices as

given from the data, and find that the fall in real rates since the 1980s explains 75% of the rise in wealth inequality. In our study, portfolios are endogenous, meaning that our model provides an explanation for the effect documented in their paper.

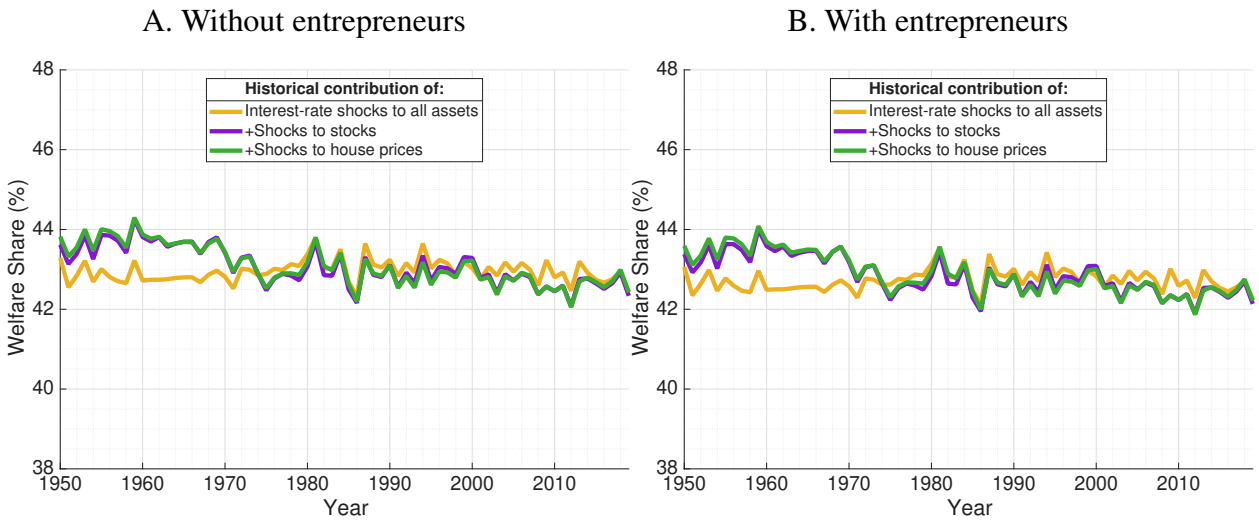
The patterns we document would also allow us to explain a substantial fraction of the rise in income inequality documented in prior work (Auten and Splinter, 2024; Piketty et al., 2018; Smith et al., 2020). A substantial fraction of rising top incomes comes from capital gains on the investments of the wealthy. As such, the large increase in capital income from realized gains—driven by falling interest rates—also helps explain rising income inequality over the last 4 decades.

7.4 Implications for welfare

Are the wealthy actually better off from the interest-rate-induced increase in wealth inequality? As predicted in Section 4.5, increases in top wealth shares stemming from falling rates do not translate into increased welfare inequality. To be precise, falling rates have no distributional effect on welfare within cohorts and a small amount of welfare redistribution to older cohorts with shorter investment horizons. To illustrate this point, Figure 13 reports the evolution of welfare inequality in the model, using the certainty equivalent measure from equation (53). As predicted by theory, welfare inequality remains largely unchanged.[24] This remains true even accounting for multiple assets, risk factors, portfolio constraints, and frictions.

[24] If anything, welfare inequality trends go in the opposite direction. This is because, in the full model, high earners have higher life expectancy and are therefore slightly more exposed to interest-rate risk, as predicted in equation (55).

Figure 13: Evolution of the top 10% welfare



8 Conclusion

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

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

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

for future research.

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

References

Ang, Andrew, Geert Bekaert, and Min Wei, "The Term Structure of Real Rates and Expected Inflation," *The Journal of Finance*, 2008, *63* (2), 797–849.

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

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

Auten, Gerald and David Splinter, "Income Inequality in the United States: Using Tax Data to Measure Long-term Trends," *Journal of Political Economy*, 2024, *132* (7), 2179–2530.

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

Bansal, Ravi and Amir Yaron, "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *The Journal of Finance*, 2004, *59* (4), 1481–1509.

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

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

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

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

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

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

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

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

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

Catherine, Sylvain, Max Miller, and Natasha Sarin, "Social Security and Trends in Wealth Inequality," *The Journal of Finance*, 2025, *80* (3), 1497–1531.

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.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez, "Where is the land of Opportunity? The Geography of Intergenerational Mobility in the United States *," *The Quarterly Journal of Economics*, 09 2014, *129* (4), 1553–1623.

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.

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

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.

Fagereng, Andreas, Matthieu Gomez, Émilien Gouin-Bonenfant, Martin Blomhoff Holm, Benjamin Moll, and Gisle Natvik, "Asset-Price Redistribution," *Journal of Political Economy*, 2025. Forthcoming.

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

Gomez, Matthieu and Émilien Gouin-Bonenfant, "Wealth Inequality in a Low Rate Environment," *Econometrica*, 2024, *92* (1), 201–246. Working Paper.

Gordon, Myron J. and Eli Shapiro, "Capital Equipment Analysis: The Required Rate of Profit," *Management Science*, 1956, *3* (1), 102–110.

Gormsen, Niels Joachim and Eben Lazarus, "Interest Rates and Equity Valuations," 2025. Working Paper.

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

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

Guvenen, Fatih, Serdar Ozkan, and Jae Song, "The Nature of Countercyclical Income Risk," *Journal of Political Economy*, 2014, *122* (3), 621 – 660.

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

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

Jordà, Òscar, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M Taylor, "The Rate of Return on Everything, 1870–2015*," *The Quarterly Journal of Economics*, 04 2019, *134* (3), 1225–1298.

Longstaff, Francis A. and Monika Piazzesi, "Corporate earnings and the equity premium," *Journal of Financial Economics*, 2004, *74* (3), 401–421.

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

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

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

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

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

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

Piazzesi, Monika and Martin Schneider, "Housing and Macroeconomics," in John B. Taylor and Harald Uhlig, eds., *John B. Taylor and Harald Uhlig, eds.*, Vol. 2 of *Handbook of Macroeconomics*, Elsevier, 2016, chapter 19, pp. 1547–1640.

Piazzesi, Monika, Martin Schneider, and Selale Tuzel, "Housing, consumption and asset pricing," *Journal of Financial Economics*, 2007, *83* (3), 531–569.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman, "Distributional National Accounts: Methods and Estimates for the United States," *The Quarterly Journal of Economics*, 2018, *133* (2), 553–609.

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

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

Schwert, G. William, "Stock Returns and Real Activity: A Century of Evidence," *The Journal of Finance*, 1990, *45* (4), 1237–1257.

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

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

Wachter, Jessica A., "Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?," *The Journal of Finance*, 2013, *68* (3), 987–1035.

INTERNET APPENDIX

A Interest-rate sensitivity: Definitions and derivations

A.1 Interest-rate sensitivity of an asset

Consider a claim to a stream of cashflows $\{D_{t+k} : k \in \{1, 2, \dots\}\}$. Let P_t denote the price of the asset and $\{R_{t+k} : k \in \{1, 2, \dots\}\}$ its future one-period returns (i.e., $R_{t+k} = (P_{t+k} + D_{t+k})/P_{t+k-1}$).

The *interest-rate sensitivity* of an asset is defined as the elasticity of its price with respect to the interest rate:

$$\varepsilon_r(P_t) \equiv -\frac{\partial \log P_t}{\partial r_{ft}}, \tag{A.1}$$

holding fixed the distributions of future cashflows $\{D_{t+1+k}\}$ and future excess returns $\{R_{t+1+k}/R_{f,t+k}\}$.

Following this definition, let us derive an expression for interest-rate sensitivity under the assumption that interest rates follow an AR(1). It will be easiest to do this if we separate the asset into separate claims to the cashflows at each horizon k (the "strips"). Letting P_{kt} denote the price of the claim to the single cashflow D_{t+k}, we have

$$P_t = \sum_{k=1}^{\infty} P_{kt}. \tag{A.2}$$

The interest-rate sensitivity of the asset is then the value-weighted average of the sensitivities of each strip:

$$\varepsilon_r(P_t) = \sum_{k=1}^{\infty} \frac{P_{kt}}{P_t} \varepsilon_r(P_{kt}). \tag{A.3}$$

To derive these elasticities, note that, from the definition of a return,

$$P_{kt} = \frac{P_{k-1,t+1}}{R_{t+1}} = \left(\prod_{k'=1}^{k} R_{t+k'}^{-1} \right) D_{t+k}. \tag{A.4}$$

Thus, defining the excess return $\hat{r}_{t+k} \equiv \log(R_{t+k}/R_{f,t-1+k})$ and taking expectations,

$$\log P_{kt} = -\sum_{k'=1}^{k} (\mathbb{E}_t[r_{f,t-1+k'}] + \mathbb{E}_t[\hat{r}_{t+k'}]) + \mathbb{E}_t[\log D_{t+k}]. \tag{A.5}$$

Now, using the fact that

$$\mathbb{E}_t[r_{f,t-1+k'}] = \bar{r}_f + \varphi^{k'-1}(r_{ft} - \bar{r}_f), \tag{A.6}$$

we have that

$$\varepsilon_r(P_{kt}) = -\frac{\partial \log P_{kt}}{\partial r_{ft}} = -\sum_{k'=1}^{k}\left(\varphi^{k'-1} + \frac{\partial \mathbb{E}_t[\hat{r}_{t+k'}]}{\partial r_{ft}}\right) + \frac{\partial \mathbb{E}_t[\log D_{t+k}]}{\partial r_{ft}} = \frac{1-\varphi^k}{1-\varphi}. \tag{A.7}$$

The last equality follows from the assumption that expected excess returns and cashflow growth are held constant, so $\partial \mathbb{E}_t[\hat{r}_{t+k'}]/\partial r_{ft} = \partial \mathbb{E}_t[D_{t+k}]/\partial r_{ft} = 0$. Substituting this into the expression above, we then have that the asset's interest-rate sensitivity is

$$\varepsilon_r(P_t) = \sum_{k=1}^{\infty} \frac{P_{kt}}{P_t}\frac{1-\varphi^k}{1-\varphi}. \tag{A.8}$$

A.2 Relating interest-rate sensitivity to duration

The (Macaulay) *duration* of an asset is defined as the value-weighted timing of its cashflows:

$$\text{dur}(P_t) \equiv \sum_{k=1}^{\infty} \frac{P_{kt}}{P_t} k, \tag{A.9}$$

where the strip prices P_{kt} are defined above.

One can see from this expression that

$$\lim_{\varphi \to 1} \varepsilon_r(P_t) = \text{dur}(P_t), \tag{A.10}$$

meaning that the duration is the same as the response to a permanent change in interest rates. Our empirical measure of interest-rate sensitivity is derived from a first-order approximation of ε_r around the duration. To see this, note that, to a first-order,

$$\varphi^k \approx \varphi^{\text{dur}(P_t)} + \varphi^{\text{dur}(P_t)} \log \varphi (k - \text{dur}(P_t)). \tag{A.11}$$

Substituting this approximation back into the expression for ε_r gives

$$\varepsilon_r(P_t) = \frac{1 - \sum_{k=1}^{\infty} \frac{P_{kt}}{P_t}\varphi^k}{1-\varphi} \approx \frac{1-\varphi^{\text{dur}(P_t)}}{1-\varphi}. \tag{A.12}$$

Thus, high duration implies high interest-rate sensitivity.

A.3 Interest-rate sensitivity of wealth

Just as the rate sensitivity of an asset is the value-weighted average of the sensitivities of its strips, the rate sensitivity is the value-weighted average of the sensitivities of its assets. Suppose a household has wealth

$$W_t = \sum_j A_{jt} - \sum_k D_{kt}. \tag{A.13}$$

In the case of positive net worth $W_t > 0$, it is clear that

$$\varepsilon_r(W_t) = -\sum_j \frac{A_{jt}}{W_t}\frac{\partial \log A_{jt}}{\partial r_{ft}} + \sum_k \frac{D_{kt}}{W_t}\frac{\partial \log D_{kt}}{\partial r_{ft}} \tag{A.14}$$

$$= \sum_j \frac{A_{jt}}{W_t}\varepsilon_r(A_{jt}) - \sum_k \frac{D_{kt}}{W_t}\varepsilon_r(D_{kt}). \tag{A.15}$$

The definition of rate sensitivity above applies only to positive-wealth households. This is not an issue in our empirics, as we only include these households in our computations. Still, it is worth discussing how one would measure rate sensitivity for a household with negative wealth. A simple generalization that handles this case is

$$\varepsilon_r(W_t) \equiv -\frac{\partial W_t}{\partial r_{ft}}\frac{1}{|W_t|}. \tag{A.16}$$

For $W_t > 0$, one can see that this reverts to the original definition (by the chain rule). For a household with negative networth,

$$\varepsilon_r(W_t) = \sum_j \frac{A_{jt}}{|W_t|}\varepsilon_r(A_{jt}) - \sum_k \frac{D_{kt}}{|W_t|}\varepsilon_r(D_{kt}). \tag{A.17}$$

The absolute value in the denominator is necessary for getting the sign of the elasticity right. For example, suppose a household has \$100 of assets and \$200 of debt, both with the same rate sensitivity (say, 2). If rates fall by 1%, networth should fall by roughly \$2, from $-$\$100 to about $-$\$102.[25] In this case, (A.17) implies a rate sensitivity of wealth of -2, reflecting the fact that a decline in rates has a negative effect on wealth.

B Data appendix

B.1 Survey of Consumer Finances

Data on household portfolios come from the Survey of Consumer Finances (SCF). We construct networth as:[26]

$$\texttt{networth_d} = \texttt{cash_dep} + \texttt{equity} + \texttt{fixed_inc} + \texttt{real_estate} + \texttt{bus} + \texttt{vehic} - \texttt{mortgage_dbt} - \texttt{vehic_dbt} - \texttt{other_dbt},$$

[25] Under a rate sensitivity of 2, assets will rise from \$100 to \$102 and debt will rise from \$200 to \$204.

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

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 on the number of people in the household and the Social Security wage index to create a per capita wage measure that is comparable over time.

B.2 Duration component calculations

B.2.1 Duration of equity

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

B.2.2 Duration of fixed income

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

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

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

B.2.3 Duration of real estate

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

B.2.4 Duration of private business wealth

The duration of private business wealth is computed for each household as the value of household businesses, bus, divided by the annual cashflows from those equity holdings. However, the annual cashflows from those equity holdings are not reported in the SCF, the major issue being that 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. We do this in four ways, depending on the household's role in the business and what is reported.

1. For households whose main respondent has an active management role in either of the household's potential actively managed businesses, reports being self-employed, and reports not receiving a salary, we estimate their predicted wage.
    - The predicted wage is estimated via ordinary least squares on all SCF respondents j where the household's wage income is the dependent variable, and the independent variables are a third-degree polynomial in age interacted with dummies for each Race $\times$ Education $\times$ Gender group.

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

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

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

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

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

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

B.2.5 Duration of vehicles

The `vehic` category in the SCF contains detailed information on up to 4 automobiles, up to 2 non-automobile vehicles, and an aggregation of additional automobiles and non-automobile vehicles owned by the household. For the primary automobiles of the house, we attribute an expected lifetime of 8 years for 1989 and 12 years for 2019, linearly interpolating in intermediate years. We calculate the time left on an automobile's life as the model year plus the expected age minus the survey year. We assume a fixed depreciation rate to 0 over the car's remaining years, and calculate the duration using (5). 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.

B.2.6 Duration of debts

For the debt categories, `mortgage_dbt`, `vehic_dbt`, and `other_dbt`, we break each up into their component loans as described in the SCF extract and calculate the duration of each loan separately. For each loan, we assume a fixed payment schedule, and thus its duration can be calculated using equation (5), where N is the maturity of the loan and y_{nt} is the riskfree spot rate at horizon n in year t.

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

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

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

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

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

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

Finally, since adjustable rate mortgages (ARMs) have no interest-rate exposure after they float, we treat them differently. In particular, we cap their duration at 4—the average number of years until the interest rate begins to float in the SCF. Approximately 12% of mortgages and 19% of investment property loans are adjustable rate in the data.

B.2.7 Interest-rate sensitivity of Social Security wealth

The interest-rate sensitivity of Social Security wealth comes from the methodology used in Catherine et al. (2025). We generate their baseline risk-adjusted Social Security wealth under the net present value wealth concept and the Treasury yield curve. We then generate Social Security wealth under an identical specification where the log forward rate at horizon h in survey year t is given by $\tilde{f}_{h,t} = f_{h,t} + \varphi^h 0.01$ where $f_{h,t}$ is the unshocked log forward rate and φ is our calibrated persistence for the interest rate process shown in Table 1. This is the same thing as applying a one-period shock of 0.01 to the log riskfree rate under the process in equation (2).

C Model derivations

C.1 Long-term bond returns

This section explains how the riskfree rate dynamics (2) imply the n-period bond returns (10). Since it has no intermediate cashflows, the bond's return from t to $t+1$ is

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

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

$$y_{nt} \equiv \frac{1}{n}\log\left(\frac{1}{P_{nt}}\right) = \frac{1}{n}\sum_{j=1}^{n}(\mathbb{E}_t r_{f,t+j-1} + \mu_j). \tag{C.2}$$

Moreover, note that (2) iterates backward to the expression

$$r_{f,t+j} = (1-\varphi^j)\bar{r}_f + \varphi^j r_{ft} + \sum_{k=1}^{j}\varphi^{j-k}\sigma_r\epsilon_{r,t+k}. \tag{C.3}$$

Substituting the riskfree rates (C.3) into the yield expression (C.2) and evaluating expectations implies

$$y_{nt} = \bar{r}_f + \frac{1}{n}\frac{1-\varphi^n}{1-\varphi}(r_{ft} - \bar{r}_f) + \frac{1}{n}\sum_{j=1}^{n}\mu_j. \tag{C.4}$$

Taking logs of (C.1) and substituting (C.4) into the yield then implies the log return

$$r_{n,t+1} = ny_{nt} - (n-1)y_{n-1,t+1} \tag{C.5}$$

$$= \bar{r}_f + \frac{1-\varphi^n}{1-\varphi}(r_{ft} - \bar{r}_f) - \frac{1-\varphi^{n-1}}{1-\varphi}(r_{f,t+1} - \bar{r}_f) + \mu_n \tag{C.6}$$

$$= \bar{r}_f + \frac{1-\varphi^n}{1-\varphi}(r_{ft} - \bar{r}_f) - \frac{\varphi-\varphi^n}{1-\varphi}(r_{ft} - \bar{r}_f) - \frac{1-\varphi^{n-1}}{1-\varphi}\sigma_r\epsilon_{r,t+1} + \mu_n \tag{C.7}$$

$$= r_{ft} + \mu_n - \frac{1-\varphi^{n-1}}{1-\varphi}\sigma_r\epsilon_{r,t+1}, \tag{C.8}$$

the stated expression (10).

C.2 Stock returns

Let D_t denote the dividends of the market and V_t its value. We make two assumptions. First, expected log returns equal the riskfree rate plus a constant risk premium:

$$\mathbb{E}_t[r_{s,t+1}] = r_{ft} + \mu_s. \tag{C.9}$$

Second, dividends are cointegrated with aggregate income. Second, following Longstaff and Piazzesi (2004), the log dividend-income ratio $dl_t \equiv \log(D_t/L_t)$ follows the autoregression

$$dl_{t+1} = (1-\varphi_d)\overline{dl} + \varphi_d dl_t + \epsilon_{d,t+1}, \tag{C.10}$$

for some shock $\epsilon_{d,t+1}$ with $\mathbb{E}_t[\epsilon_{d,t+j}] = 0$. This implies the log dividend growth rate

$$g_{D,t+1} = g_{t+1} + \Delta dl_{t+1} = \lambda_{gr} r_{f,t+1} + g_{1,t+1} - (1-\varphi_d)(dl_t - \overline{dl}) + \epsilon_{d,t+1}. \tag{C.11}$$

Next, we linearize the stock return under these assumptions. By definition, the return on the stock market is

$$R_{s,t+1} = \frac{D_{t+1} + V_{t+1}}{V_t} = \frac{D_{t+1}}{D_t}\left(\frac{1 + V_{t+1}/D_{t+1}}{V_t/D_t}\right). \tag{C.12}$$

Now take logs and linearize to get the Campbell-Shiller decomposition. Letting $g_{D,t+1} = \log(D_{t+1}/D_t)$ and $vd_t = \log(V_t/D_t)$ represent the log of dividend growth and of the price-dividend ratio, respectively, we have

$$r_{s,t+1} \approx g_{D,t+1} + \kappa_0 + \kappa_1 vd_{t+1} - vd_t, \tag{C.13}$$

where the linearizing constants are

$$\kappa_1 \equiv \frac{1}{1 + \exp\{-\overline{vd}\}} \quad \text{and} \quad \kappa_0 \equiv -\log\kappa_1 - \kappa_1\overline{vd}, \tag{C.14}$$

for $\overline{vd} = \mathbb{E}[vd_t]$. Solving the linearized equation forward for the price-dividend ratio implies

$$vd_t = g_{D,t+1} - r_{s,t+1} + \kappa_0 + \kappa_1 vd_{t+1} = \kappa_0 + \sum_{j=0}^{\infty} \kappa_1^j (g_{D,t+1+j} - r_{s,t+1+j}). \tag{C.15}$$

which holds both in realization and in expectation.

We know the expected stock return $\mathbb{E}_t[r_{s,t+1}]$ by assumption. Using the linearization, the unexpected component of the stock return equals

$$r_{s,t+1} - \mathbb{E}_t[r_{s,t+1}] = g_{D,t+1} - \mathbb{E}_t[g_{D,t+1}] + \kappa_1(vd_{t+1} - \mathbb{E}_t[vd_{t+1}]). \tag{C.16}$$

Because risk premia are constant, the unexpected change in the price-dividend ratio equals:

$$\kappa_1(vd_{t+1} - \mathbb{E}_t[vd_{t+1}]) = \sum_{j=1}^{\infty} \kappa_1^j (\mathbb{E}_{t+1}[g_{D,t+1+j}] - \mathbb{E}_t[g_{D,t+1+j}]) - \sum_{j=1}^{\infty} \kappa_1^j (\mathbb{E}_{t+1}[r_{f,t+j}] - \mathbb{E}_t[r_{f,t+j}]), \tag{C.17}$$

and hence

$$\kappa_1(vd_{t+1} - \mathbb{E}_t[vd_{t+1}]) = \sum_{j=0}^{\infty} \kappa_1^j (\mathbb{E}_{t+1}[g_{D,t+1+j}] - \mathbb{E}_t[g_{D,t+1+j}]) - \sum_{j=1}^{\infty} \kappa_1^j (\mathbb{E}_{t+1}[r_{f,t+j}] - \mathbb{E}_t[r_{f,t+j}]). \tag{C.18}$$

Note first that

$$\mathbb{E}_{t+1}[r_{f,t+j}] - \mathbb{E}_t[r_{f,t+j}] = \varphi^{j-1}\sigma_r \epsilon_{r,t+1}, \tag{C.19}$$

and therefore that

$$\sum_{j=1}^{\infty} \kappa_1^j (\mathbb{E}_{t+1}[r_{f,t+1+j}] - \mathbb{E}_t[r_{f,t+1+j}]) = \frac{\kappa_1}{1-\kappa_1\varphi}\sigma_r \epsilon_{r,t+1}. \tag{C.20}$$

Next, note that

$$\begin{aligned}\mathbb{E}_{t+1}[g_{D,t+1+j}] - \mathbb{E}_t[g_{D,t+1+j}] = \lambda_{gr}(\mathbb{E}_{t+1}[r_{f,t+1+j}] - \mathbb{E}_t[r_{f,t+1+j}]) + (\mathbb{E}_{t+1}[g_{1,t+1+j}] - \mathbb{E}_t[g_{1,t+1+j}]) \\ - (1-\varphi_d)(\mathbb{E}_{t+1}[dl_{t+j}] - \mathbb{E}_t[dl_{t+j}]) + \mathbb{I}_{j=0}\epsilon_{d,t+1}.\end{aligned} \tag{C.21}$$

The riskfree rate component equals

$$\lambda_{gr}(\mathbb{E}_{t+1}[r_{f,t+1+j}] - \mathbb{E}_t[r_{f,t+1+j}]) = \lambda_{gr}\varphi^j(r_{f,t+1} - \mathbb{E}_t[r_{f,t+1}]) = \lambda_{gr}\varphi^j\sigma_r\epsilon_{r,t+1}. \tag{C.22}$$

Similarly,

$$\mathbb{E}_{t+1}[g_{1,t+1+j}] - \mathbb{E}_t[g_{1,t+1+j}] = \varphi_g^j(g_{1,t+1} - \mathbb{E}_t[g_{1,t+1}]) = \varphi_g^j\epsilon_{g,t+1}, \tag{C.23}$$

and

$$(1-\varphi_d)(\mathbb{E}_{t+1}[dl_{t+j}] - \mathbb{E}_t[dl_{t+j}]) = (1-\varphi_d)\varphi_d^{j-1}\epsilon_{d,t+1}. \tag{C.24}$$

Consequently,

$$\sum_{j=0}^{\infty} \kappa_1^j (\mathbb{E}_{t+1}[g_{D,t+1+j}] - \mathbb{E}_t[g_{D,t+1+j}]) = \frac{\lambda_{gr}\sigma_r}{1-\kappa_1\varphi}\epsilon_{r,t+1} + \frac{1}{1-\kappa_1\varphi_g}\epsilon_{g,t+1} + \frac{1-\kappa_1}{1-\kappa_1\varphi_d}\epsilon_{d,t+1}. \tag{C.25}$$

Putting all of this together, the unexpected stock return equals

$$r_{s,t+1} - \mathbb{E}_t[r_{s,t+1}] = \frac{\lambda_{gr}-\kappa_1}{1-\kappa_1\varphi}\sigma_r\epsilon_{r,t+1} + \frac{1}{1-\kappa_1\varphi_g}\epsilon_{g,t+1} + \frac{1-\kappa_1}{1-\kappa_1\varphi_d}\epsilon_{d,t+1}. \tag{C.26}$$

Substituting in the expected return, we have an expression for return.

Finally, to get to the expression (13) in the main text, we assume that the shock ϵ_d can be decomposed such that

$$\frac{1-\kappa_1}{1-\kappa_1\varphi_d}\epsilon_{d,t+1} = \lambda_{dg}\epsilon_{g,t+1} + \epsilon_{s,t+1} \tag{C.27}$$

where ϵ_g is the growth shock and ϵ_s is the market-specific shock. In the data, dividends fall by more than consumption and output in recessions, so $\lambda_{dg} > 0$ (Longstaff and Piazzesi, 2004). For the loading λ_{sr} on interest-rate shocks in (13), we use the fact that, under this linearization, the interest-rate sensitivity of the stock equals

$$\frac{\kappa_1}{1-\kappa_1\varphi} = \frac{\sigma_{n_s}}{\sigma_r} = \frac{1-\varphi^{n_s-1}}{1-\varphi}, \tag{C.28}$$

where n_s is the duration of the stock market. This implies

$$\kappa_1 = \frac{\sigma_{n_s}/\sigma_r}{1+\varphi\sigma_{n_s}/\sigma_r} = \frac{1-\varphi^{n_s-1}}{1-\varphi^{n_s}}. \tag{C.29}$$

We then use this value to compute the loading on the interest rate,

$$\lambda_{sr} = \frac{\kappa_1 - \lambda_{gr}}{1-\kappa_1\varphi}. \tag{C.30}$$

Together, this gives the stock return (13).

C.3 Housing returns

To derive housing returns, we need to make assumptions about depreciation and maintenance of the housing stock, as well as transaction costs from buying and selling housing. First, we assume that the housing stock depreciates each period and owners offset this by investing a proportionate maintenance cost. Second, we assume that homeowners incur proportionate transaction costs, representing the cost of moving. In the model, it is costless for agents to change house size, and thus households move houses every period and pay this transaction cost each period.

These assumptions manifest themselves in the housing return, which equals

$$R_{hi,t+1} = \frac{P_{n_h-1,t+1}}{P_{n_h t}} \frac{P_{h,t+1}}{P_{ht}} \frac{\tilde{P}_{i,t+1}}{\tilde{P}_{it}}, \tag{C.31}$$

as follows. The bond component of the return accounts for the fact that the house depreciates over the period, so its duration falls from n_h to $n_h - 1$, as would a long-term bond. The drift μ_h in (16) includes not only common price appreciation, but also nets out the costs to the homeowner of transacting and of maintaining the property. The investment in maintenance offsets the depreciation, and in doing so returns the bond portion of the house back to its original duration n_h for the new set of homebuyers. And finally, because the house is sold after this return is realized, the idiosyncratic house price shock only accumulates over the one period, at which point the homeowner sells the house and buys a new one at the regional price P_t.

C.4 Private business valuation

Here we derive the valuation of private business at exit (39). Both in our model and in empirical measures of the value of the business, the relevant valuation is under the discount rate and cashflow dynamics after the company is sold. For discount rates, we assume that the business is sold to diversified investors—this will be the case, for instance, if the company goes public. Because the new owners will be diversified, they will not demand compensation for idiosyncratic risk, so the relevant risk premium will be the same as that of the market. Future cashflows L_i can follow a general process on which we impose only one fairly minor assumption: growth in the persistent part of the idiosyncratic component of income is not predictable from current information. Formally, for all $j \geq 0$, $\mathbb{E}_t[\tilde{L}_{i,t+1+j}/\tilde{L}_{i,t+j}] = \mathbb{E}[\tilde{L}_{i,t+1+j}/\tilde{L}_{i,t+j}]$. Note that this assumption allows for the possibility that the cashflow growth dynamics could change after exit. The only purpose of this assumption is to simplify the model; it affects fluctuations in valuation over time, but not the average level of valuations, as we will see from what follows.

Let E_{it} denote the present value of future earnings $\{L_{i,t+j}\}$. Let e_{it} and $l_{i,t+j}$ denote logs of these variables and $el_{it} \equiv e_{it} - l_{it}$ the current valuation ratio. In what follows, we will suppress i subscripts. Applying the same Campbell-Shiller linearization as we did for the stock market in Appendix C.2, we have that the one-period log return on the business is

$$r_{e,t+1} \approx g_{L_i,t+1} + \kappa_{E0} + \kappa_E el_{t+1} - el_t, \tag{C.32}$$

where

$$\kappa_E \equiv \frac{1}{1 + \exp\{-\overline{el}\}} \quad \text{and} \quad \kappa_{E0} \equiv -\log \kappa_E - \kappa_E \overline{el}, \tag{C.33}$$

and $\overline{el} \equiv \mathbb{E}[el_t]$ is the unconditional mean valuation ratio. Iterating forward and taking expectations, we have

$$el_t = \kappa_{E0} + \sum_{j=0}^{\infty} \kappa_E^j (\mathbb{E}_t[g_{l,t+1+j}] - \mathbb{E}_t[r_{e,t+1+j}]), \tag{C.34}$$

or, in deviations from the mean,

$$el_t - \overline{el} = \sum_{j=0}^{\infty} \kappa_E^j (\mathbb{E}_t[g_{l,t+1+j}] - \bar{g}_l) - \sum_{j=0}^{\infty} \kappa_E^j (\mathbb{E}_t[r_{e,t+1+j}] - \bar{r}_e). \tag{C.35}$$

Because the firm is sold to diversified investors, it has the same expected return as the stock market, so

$$\mathbb{E}_t[r_{e,t+1+j}] = \mathbb{E}_t[r_{s,t+1+j}] = \mathbb{E}_t[r_{f,t+j}] + \mu_s = (1-\varphi^j)\bar{r}_f + \varphi^j r_{ft} + \mu_s. \tag{C.36}$$

This implies

$$\sum_{j=0}^{\infty} \kappa_E^j (\mathbb{E}_t[r_{e,t+1+j}] - \bar{r}_e) = \frac{1}{1-\kappa_E \varphi} (r_{ft} - \bar{r}_f). \tag{C.37}$$

The cashflow growth component can be further divided into aggregate wage index growth and idiosyncratic growth. The aggregate growth component is

$$\sum_{j=0}^{\infty} \kappa_E^j (\mathbb{E}_t[g_{1,t+1+j}] - \bar{g}_1) + \frac{\lambda_{gr}}{1-\kappa_E \varphi} (r_{ft} - \bar{r}_f) = \frac{1}{1-\kappa_E \varphi_g} (g_{1t} - \bar{g}_1) + \frac{\lambda_{gr}}{1-\kappa_E \varphi} (r_{ft} - \bar{r}_f). \tag{C.38}$$

And we have assumed that the idiosyncratic component is not predictable, so

$$\sum_{j=0}^{\infty} \kappa_E^j (\mathbb{E}_t[g_{\tilde{l},t+1+j}] - \mathbb{E}[g_{\tilde{l},t+1+j}]) = 0. \tag{C.39}$$

Consequently,

$$el_t = \overline{el} + \frac{1}{1-\kappa_E \varphi_g} (g_{1t} - \bar{g}_1) - \frac{1-\lambda_{gr}}{1-\kappa_E \varphi} (r_{ft} - \bar{r}_f). \tag{C.40}$$

Taking the exponential of this and relabelling $\overline{el}$ as the parameter υ, we get the expression (39).

D Model calibration

D.1 Macroeconomic variables

D.1.1 Time series of riskfree interest 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

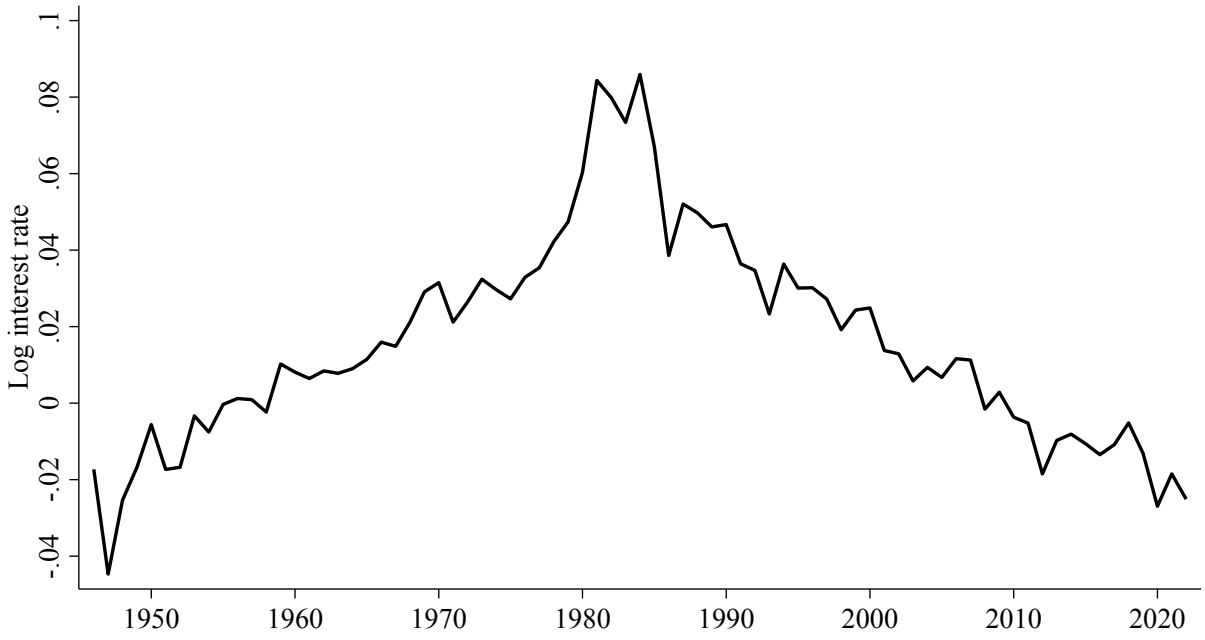
$$y_{10,t} - \pi_{t,t+1} = \beta_0 + \beta_1 y_{10,t} + \beta_2 \pi_{t-1,t} + \epsilon_{t+1} \tag{D.1}$$

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 (2) yields the time-t riskfree rate:

$$r_{ft} = \varphi^{-10}(\hat{f}_{10,t} - (1 - \varphi^{10})\bar{r}_f). \tag{D.2}$$

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 D.1.

Figure D.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 (D.1) and transformed by equation (D.2).

D.1.2 Economic growth and stock-market returns

We jointly estimate the dynamics of aggregate growth (7) and stock returns (13) via the simulated method of moments (SMM). The parameter vector to be estimated is[27]

$$\Theta \equiv \begin{bmatrix} \bar{g} & \varphi_g & \mu_g^- & \sigma_{g_1}^- & \sigma_{g_1}^+ & p_g & \lambda_{gr} & \mu_s & \lambda_{sg} & \sigma_s^- & \sigma_s^+ \end{bmatrix}. \tag{D.3}$$

[27]Recall that λ_{sr} is given by (C.30) and that μ_g^+ sets $p_g\mu_g^- + (1-p_g)\mu_g^+ = 0$, so these are implicitly estimated here.

We estimate the probability of a recession p_g directly from the data as the frequency of NBER recessions. To identify the remaining 11 parameters, we use the first four moments (i.e., mean, variance, skewness, and kurtosis) of log stock returns and of log GDP growth, the autocorrelation of GDP growth, the correlation of stock returns and GDP growth, and the correlation of riskfree rates with five-year moving average growth. We use the correlation between lagged returns and GDP growth as a moment because stock markets are forward looking. This shows up in the data—as noted in Schwert (1990); **?**); **?**—by stock markets returns predicting future real activity more strongly than they predict current real activity. The higher-order moments (skewness and kurtosis) are particularly useful for identifying parameters of the normal mixtures. The correlation of rates with a moving average of the growth rate, as opposed to the growth rate itself, provides identification for λ_{gr} by isolating the low-frequency covariance.

The SMM procedure then estimates Θ as follows. Letting $m(\Theta)$ denote the vector of simulated moments implied by the model and $\hat{m}$ the corresponding vector of data moments, our estimated parameters minimize the squared (percentage) errors:

$$\hat{\Theta} = \arg\min_{\Theta} \left\{ \left(\frac{\hat{m} - m(\theta)}{\hat{m}} \right)^{\top} \left(\frac{\hat{m} - m(\theta)}{\hat{m}} \right) \right\}, \tag{D.4}$$

where division of vectors is an element-wise operation. For the data moments, we use historical data from a number of sources. We construct the growth series using per capita real GDP growth as a proxy for aggregate wage dynamics. The primary data source is the Global Macro Database (GMD), from which we extract U.S. real GDP per capita from 1900 to 2021. Stock-market return data cover 1871–2021 and come from Shiller. We use the log total return on the S&P 500.

Table D.1 reports the estimated parameters, and Table D.2 reports the moments in the data and model.

Table D.1: Estimated parameters for growth and stock returns

Parameter	Notation	Estimate
Mean growth	$\bar{g}$	0.0166
Persistence of growth shocks	φ_g	0.1541
Mean growth (recessions)	μ_g^-	-0.0354
Growth volatility (recession)	$\sigma_{g_1}^-$	0.0119
Growth volatility (normal times)	$\sigma_{g_1}^+$	0.0135
Recession probability	p_g	0.164
Riskfree-rate growth loading	λ_{gr}	0.1568
Stock return drift	μ_s	0.0514
Stock volatility (recession)	σ_s^-	0.3000
Stock volatility (normal times)	σ_s^+	0.1255
Stock-growth loading	λ_{sg}	5.1385

D.2 Idiosyncratic income process

Our calibration of the idiosyncratic income component of earnings follows Catherine (2022) closely, with two differences. In this study, the parameters of the income process are estimated by targeting moments of the distribution of log earnings growth from 1978 to 2011, feeding the historical mean log earnings growth of each of these years into

Table D.2: Moments for estimation of growth and stock returns

Moment	Description	Data	Model
$\mathbb{E}[r_{st}]$	Mean stock return (1900-2021)		
$\sigma(r_{st})$	Standard deviation of stock returns		
Skew(r_{st})	Skewness of stock returns		
Kurt(r_{st})	Kurtosis of stock returns		
$\mathbb{E}[g_t]$	Mean wage growth (1960-2021)		
$\sigma(g_t)$	Standard deviation of wage growth		
Skew(g_t)	Skewness of wage growth		
Kurt(g_t)	Kurtosis of wage growth		
Corr(g_t, g_{t-1})	Autocorrelation of wage growth		
Corr(g_t, r_{st})	Contemporaneous correlation		
Corr$(r_{f,t}, \bar{g}_{5,t})$	Risk-free rate and 5-year average growth		
p_g	Recession probability	0.164	0.164

the simulation. The targeted moments include the standard deviation of log earnings growth at the 1– and 5–year horizons, and the skewness at the 1–, 3–, and 5–year horizons. These moments are computed using Social Security administrative data and reported in Guvenen et al. (2014).

We estimate the process using the same procedure as Catherine (2022), described in detail in Appendix B.1 of his study, with two differences. First, for years prior to 1978, we feed into the simulation the historical time series of log changes in real per capita GDP growth, which we use to calibrate the process for economic growth as in Section D.1. For the period 1978–2011, we feed into the simulation the historical mean log earnings growth rate reported in Guvenen et al. (2014), but scale this variable such that it exhibits the same volatility over that period as per capita GDP growth, which is roughly 1.5 times lower. Our goal is to obtain comparable variability in idiosyncratic income risk over the business cycle in the model and in the data, while leveraging the longer time span of the per capita GDP growth series relative to the moments reported in Guvenen et al. (2014).

The second difference is that we account for the diversification arising from two earners in the typical American household. As Guvenen et al. (2014) only report moments for individual earners, we proceed in two steps to estimate the income process for a two-earner household. First, we estimate all parameters for a single-earner household by targeting the moments reported in Guvenen et al. (2014). The estimated parameters are reported in the 'Single' row of Table D.3 and are very close to those in Catherine (2022), except for the parameter λ_{zl}, which governs the relationship between GDP growth and the skewness of idiosyncratic income shocks. This parameter is roughly 1.5 times larger than in Catherine (2022), reflecting the fact that the GDP growth time series is itself about 1.5 times less volatile than the mean log earnings growth series used in this study.

Second, we use the estimated parameters to generate new moments for two-earner households. To do so, we simulate individuals using our estimated income process and randomly match them to create two-earner households. We then compute new log income growth moments from these households, assuming that the highest earner in each household receives 70% of the combined income. Finally, we rerun the estimation procedure targeting these new moments. The estimated parameters are reported in the 'Two earners' row of Table D.3.

Table D.3: Estimated parameters: Idiosyncratic labor-income shocks

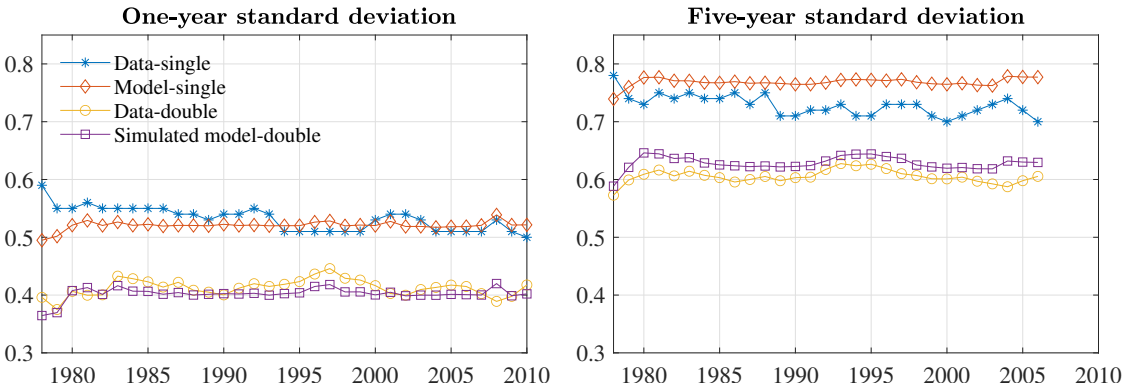
		Persistent shocks					Transitory shocks	
	p_z	ρ_z	$\overline{\mu_z^-}$	λ_{zl}	σ_z^-	σ_z^+	σ_η^-	σ_η^+
One earner	.145	.971	-.085	6.542	.644	.012	.758	.097
Two earners	.224	.993	-.026	5.219	.352	.059	.414	.124

This table reports parameter estimates for the idiosyncratic income process, conditional on the group-level aggregate shock. The single-earner process is estimated as in Catherine (2022) by targeting the time-series of cross-sectional moments of individual income log growth rates in the SSA administrative data from Guvenen et al. (2014). To estimate the double-earner process, we first simulate data for couples using the single-earner process and recompute the target moments using the sum of two incomes. We assume a 30% gap between the highest and lowest earner and no assortative matching.

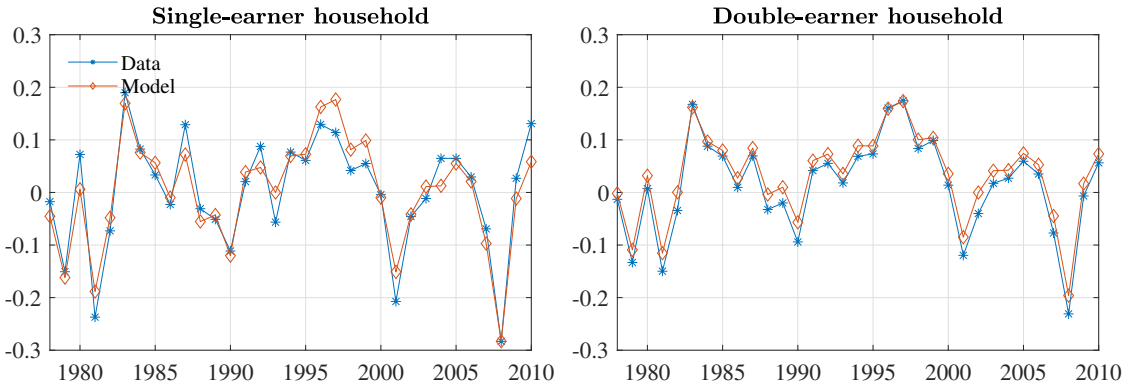
Figure D.2 presents the values and fit of targeted and simulated moments from the estimation of the idiosyncratic labor income shocks. Panels on the left report targeted and simulated moments for single-earners households, where targeted moments are from Guvenen et al. (2014). Panels on the right report targeted and simulated moments for households with two earners.

Figure D.2: Fitness of idiosyncratic labor income risk model

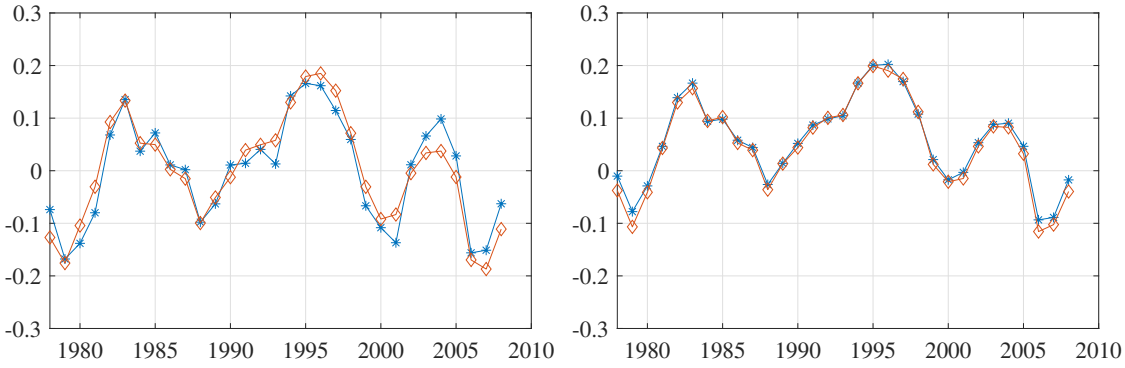
A. Standard deviation



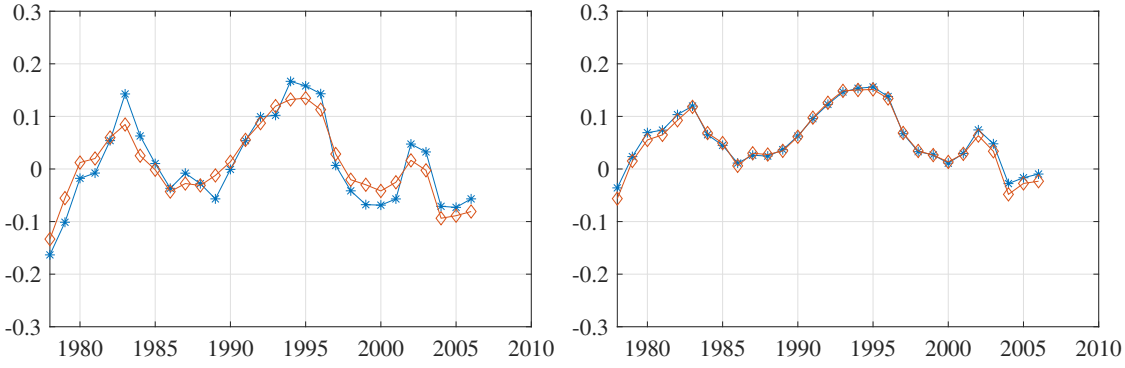
B. 1-Year Skewness



C. 3-Year Skewness



D. 5-Year Skewness



D.3 Income tax rates

Households pay income tax on labor income and Social Security benefits. In the model, households face the following marginal tax rates:

$$\text{Marginal Tax Rate}_{it} = \begin{cases} 0.10 & \text{if } \tilde{L}_{it} < 0.18, \\ 0.12 & \text{if } 0.18\tilde{L}_{it} < 0.72, \\ 0.22 & \text{if } 0.72\tilde{L}_{it} < 1.54, \\ 0.24 & \text{if } 1.54\tilde{L}_{it} < 2.94, \\ 0.32 & \text{if } 2.94\tilde{L}_{it} < 3.73, \\ 0.35 & \text{if } 3.73\tilde{L}_{it} < 9.32, \\ 0.37 & \text{if } \tilde{L}_{it} > 9.32. \end{cases} \tag{D.5}$$

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

D.4 Entrepreneurs

We calibrate the private business income (37) and wealth (39) processes to data from the SCF. We need to calibrate the deterministic component of income $\ell_E(a_{it})$, the initial volatility σ_{Ez0} and time-series volatility σ_{Ez} of the permanent component z_{it}, the mean valuation ratio υ, and the probability of becoming an entrepreneur by age $p_E(a_{it})$. In the SCF, we consider a household to be an entrepreneur if it reports being an active manager of its first or second private business. We restrict our sample to households before retirement (i.e., aged 18 to 60).

We choose the mean log valuation ratio υ to match the average log aggregate valuation ratio:

$$\upsilon = \mathbb{E}\left[\log \frac{\sum_{i=1}^{N_t} E_{it}}{\sum_{i=1}^{N_t}(L_{it} - L_{it}^{\text{wage}})}\right] = 3.124941, \tag{D.6}$$

where L_{it}^{wage} is the predicted wage paid to the entrepreneur, as computed in the empirics. Recall that this aggregate ratio was how we computed the duration and interest-rate sensitivity of private business in Section B, so this calibration is consistent with those moments as well. To estimate the deterministic life-cycle profile of business income, we fit a cubic polynomial in age to the moments

$$\ell_E(a_{it}) = \mathbb{E}\left[\log \frac{L_{it}}{\bar{L}_t}\middle| a_{it}\right], \tag{D.7}$$

which implies an estimated profile

$$\ell_E(a_{it}) = -1.580268 - 0.0822865 \times a_{it} + 0.042289 \times \frac{a_{it}^2}{10} - 0.0044304 \times \frac{a_{it}^3}{100}. \tag{D.8}$$

To estimate the variances of z_{it}, we use the fact that

$$\text{var}(z_{it}|a_{it}) = \text{var}\left(\log \frac{L_{it}}{\bar{L}_t}\middle| a_{it}\right). \tag{D.9}$$

Consequently, the time-series volatility can be computed from the change in variance of business income over the life cycle,

$$\sigma^2_{Ez} = \text{var}\left(\log \frac{L_{it}}{\bar{L}_t}\middle| a_{it}\right) - \text{var}\left(\log \frac{L_{it}}{\bar{L}_t}\middle| a_{it}-1\right), \tag{D.10}$$

and the initial volatility can be computed as the level at the initial age, 20,

$$\sigma^2_{Ez0} = \text{var}\left(\log \frac{L_{it}}{\bar{L}_t}\middle| a_{it}=20\right). \tag{D.11}$$

We estimate these by running a linear regression of the variances $\text{var}(\log(E_{it}/\bar{L}_t)|a_{it})$ on age a_{it}, and find a constant $\sigma^2_{Ez0} = 1.692465$ and a slope $\sigma^2_{Ez} = 0.0430068$. Finally, we compute the entry probability $p_E(a_{it})$ as the change in the share of households who are entrepreneurs from one age to the next.[28] Specifically, we first compute the share of households at each age a_{it} who are entrepreneurs (call this $s_E(a_{it})$), then fit a polynomial to this share, and then take the change in this fitted line as the entry probability, from the fact that

$$p_E(a_{it}) = \frac{s_E(a_{it}) - s_E(a_{it}-1)}{1 - s_E(a_{it}-1)}. \tag{D.12}$$

The change is approximately linear, and so we estimate

$$p_E(a_{it}) = \max\{0.0104395 - 0.0001949 \times a_{it}, 0\}. \tag{D.13}$$

E Numerical appendix

This appendix section details how we solve the model numerically. To render the model solvable, we exploit two key properties of the model that allow us to rescale the problem, eliminating state variables; and to transform the wealth return, simplifying the portfolio choice problem.

E.1 Rescaling the model

We eliminate two state variables, the wage index $\bar{L}_t$ and the house price P_{ht}, by rescaling the problem as follows. Henceforth, let $\hat{Z}_t$ denote wage-indexed quantities $Z_t/\bar{L}_t$ for any variable Z_t. The utility function equals

$$u(C_{it}, H_{it}) = \bar{L}_t^{1-\gamma} P_{ht}^{-\nu(1-\gamma)} u(\hat{C}_{it}, \hat{P}_{ht} H_{it}). \tag{E.1}$$

[28] We do not model exit, so the change in the share of entrepreneurs must all be attributed to entry in the model.

From this, it follows (and we will prove) that there exists a function $\hat{V}$ such that the value function equals[29]

$$V_{it} = V(\bar{L}_t, P_{ht}, W_{it}, r_{ft}, g_{1t}, z_{it}, a_{it}, \text{AIYE}_{it}) = \bar{L}_t^{1-\gamma} P_{ht}^{-\nu(1-\gamma)} \hat{V}(\hat{W}_{it}, r_{ft}, g_{1t}, z_{it}, a_{it}, \text{AIYE}_{it}), \tag{E.2}$$

provided that the agent chooses the rescaled policies $\{\hat{C}_{it}, \hat{P}_{ht} H_{it}\}$ instead of $\{C_{it}, H_{it}\}$. Henceforth, we will use the notation $\hat{H}_{it} \equiv \hat{P}_{ht} H_{it}$ to emphasize that this is a single choice variable.

Before proving the claim, note that all constraints can be rescaled accordingly. Wealth accumulates according to

$$\hat{W}_{i,t+1} = \frac{\bar{L}_t}{\bar{L}_{t+1}} (\hat{W}_{it} + \tilde{L}_{it} + \hat{B}_{it} - \hat{T}_{it} - \hat{C}_{it} - \chi P_{n_h t} \hat{H}_{it} \mathbb{I}_{\{\pi_{h,it}=0\}} - c_s) R_{W i,t+1}, \tag{E.3}$$

which is independent of $\bar{L}_t$ and P_{ht}. The minimum and maximum house size constraints are

$$P_{n_h t} \hat{H}_{it} \geq \kappa_{\min} \quad \text{and} \quad \pi_{h,it} \leq \kappa_{\max}. \tag{E.4}$$

The scaled value of a private business $\hat{E}_{it}$ is also independent of $\bar{L}_t$. Thus, the constraints can all be written in terms of controls $\{\hat{C}_{it}, \hat{H}_{it}\}$ and without any dependence on $\bar{L}_t$ and P_{ht}.

We can now prove the claim by (backward) induction. First, consider the bequest (30). In the terminal period (i.e., $\bar{b}$ periods after death), the claim is true, as

$$V_{it} = \bar{L}_t^{1-\gamma} P_{ht}^{-\nu(1-\gamma)} \max_{\{\hat{C}_{it}, \hat{P}_{ht} H_{it}\}} u(\hat{C}_{it}, \hat{P}_{ht} H_{it}) = \bar{L}_t^{1-\gamma} P_{ht}^{-\nu(1-\gamma)} \hat{V}_{it}. \tag{E.5}$$

Now suppose that it is true at any time $t+1$ within the time span after death. This implies

$$V_{it} = \max_{\{C_{it}, H_{it}, \pi_{it}\}} \left\{ u(C_{it}, H_{it}) + \beta \mathbb{E}_t \left[\bar{L}_{t+1}^{1-\gamma} P_{h,t+1}^{-\nu(1-\gamma)} \hat{V}_{i,t+1} \right] \right\} \tag{E.6}$$

$$= \bar{L}_t^{1-\gamma} P_{ht}^{-\nu(1-\gamma)} \max_{\{\hat{C}_{it}, \hat{H}_{it}, \pi_{it}\}} \left\{ u(\hat{C}_{it}, \hat{H}_{it}) + \beta \mathbb{E}_t \left[\left(\frac{\bar{L}_{t+1}}{\bar{L}_t} \right)^{1-\gamma} \left(\frac{P_{h,t+1}}{P_{ht}} \right)^{-\nu(1-\gamma)} \hat{V}_{i,t+1} \right] \right\} \tag{E.7}$$

$$= \bar{L}_t^{1-\gamma} P_{ht}^{-\nu(1-\gamma)} \hat{V}_{it}. \tag{E.8}$$

The second equality changes the choice variables to $\{\hat{C}_{it}, \hat{H}_{it}, \pi_{it}\}$, which is possible because the rescaled utility function and budget constraints only feature these choice variables. The third equality follows from the fact that that the maximand has no dependence on $\bar{L}_t$ and P_{ht}. It follows that the initial bequest utility equals

$$b(\bar{L}_t, P_{ht}, W_{it}, r_{ft}, g_{1t}) = \bar{L}_t^{1-\gamma} P_{ht}^{-\nu(1-\gamma)} \hat{b}(\hat{W}_{it}, r_{ft}, g_{1t}). \tag{E.9}$$

This proves the claim for the bequest.

The rest of the proof follows the same inductive reasoning. We know, upon death, the value function rescales

[29] For simplicity, we will consider the case of a worker who never becomes an entrepreneur, and the same logic will apply when we add entrepreneurship.

as claimed (since the value function is the bequest utility). Now, for the induction step, suppose the value function conditional on living to $t+1$ rescales as claimed. Then

$$V_{it} = \max_{\{C_{it},H_{it},\pi_{it}\}} \left\{ u(C_{it},H_{it}) + \beta \mathbb{E}_t \left[\bar{L}_{t+1}^{1-\gamma} P_{h,t+1}^{-\nu(1-\gamma)} \left((1-m_{it})\hat{V}_{i,t+1} + m_{it}\hat{b}_{i,t+1} \right) \right] \right\} \tag{E.10}$$

$$= \bar{L}_t^{1-\gamma} P_{ht}^{-\nu(1-\gamma)} \max_{\{\hat{C}_{it},\hat{H}_{it},\pi_{it}\}} \left\{ u(\hat{C}_{it},\hat{H}_{it}) \right. \tag{E.11}$$

$$\left. + \beta \mathbb{E}_t \left[\left(\frac{\bar{L}_{t+1}}{\bar{L}_t} \right)^{1-\gamma} \left(\frac{P_{h,t+1}}{P_{ht}} \right)^{-\nu(1-\gamma)} \left((1-m_{it})\hat{V}_{i,t+1} + m_{it}\hat{b}_{i,t+1} \right) \right] \right\} \tag{E.12}$$

$$= \bar{L}_t^{1-\gamma} P_{ht}^{-\nu(1-\gamma)} \hat{V}_{it}. \tag{E.13}$$

The equalities follow identical reasoning as above. Most importantly, $\hat{V}_{it}$ in the last equation is independent of $\bar{L}_t$ and P_{ht}, proving the conjecture.

E.2 Transforming the portfolio problem

In choosing a portfolio, the agent must select optimal weights $\{\pi_{n,it}, \pi_{s,it}, \pi_{h,it}\}$ on the risky assets. Numerically, this problem becomes intractable because all three assets are exposed to interest-rate shocks. We circumvent this tractability problem by re-expressing the problem in terms of exposures to separate asset-market shocks: the agent instead chooses $\{\varepsilon_{r,it}, \pi_{s,it}, \pi_{h,it}\}$, where $\varepsilon_{r,it}$ is the interest-rate sensitivity of the total portfolio. This transformation is also ideal because the long-term bond share π_n is itself not that meaningful, as it arbitrarily depends on our choice of the bond duration n; in contrast, ε_r is invariant to n.

To implement this change of choice variables, we need to rewrite the wealth return (36) in a way that replaces the long-term bond share $\pi_{n,it}$ with the interest-rate sensitivity of the total portfolio $\varepsilon_{r,it}$. To do this, let us start from the result of Appendix G.5 that, to a second order, the log wealth return (suppressing i subscripts) equals

$$r_{w,t+1} = r_{ft} + \pi_t^\top (r_{t+1} - r_{ft}\iota) + \frac{1}{2}\left(\pi_t^\top \text{diag}(\Sigma\Sigma^\top) - \pi_t^\top \Sigma\Sigma^\top \pi_t \right) - (\pi_{h,t} + \pi_{s,t} - 1)^+ \theta, \tag{E.14}$$

where π_t is the vector of portfolio weights, r_{t+1} is the vector of risky returns, ι is a vector of ones, and $\Sigma\Sigma^\top$ is the covariance matrix of returns with respect to the vector of shocks.

To transform the wealth return, first note that we can rewrite the log return on any asset $j \in \{n, s, h\}$ as

$$r_{j,t+1} = \log R_{j,t+1} = r_{ft} + \mu_j - \sigma_{jr}\epsilon_{r,t+1} + \sigma_{j,-r}^\top \epsilon_{-r,t+1}, \tag{E.15}$$

where ϵ_{-r} is defined as the vector of shocks orthogonal to ϵ_r,

$$\epsilon_{-r,t+1} \equiv [\epsilon_{g,t+1}, \epsilon_{s,t+1}, \epsilon_{h,t+1}, \tilde{\epsilon}_{h,t+1}]^\top, \tag{E.16}$$

and $\sigma_{j,-r}$ the corresponding vector of loadings. In this notation, the loadings on the rate shock are $\{\sigma_{nr}, \sigma_{sr}, \sigma_{hr}\} =$

$\{\sigma_n, \lambda_{sr}, \sigma_{n_h}\}$. Stacking these returns in vector form, we have

$$r_{t+1} = r_{ft}\iota + \mu - \Sigma_r \epsilon_{r,t+1} + \Sigma_{-r}\epsilon_{-r,t+1}, \tag{E.17}$$

meaning that the wealth return is

$$\begin{aligned} r_{w,t+1} = r_{ft} + \pi_t^\top(\mu - \Sigma_r\epsilon_{r,t+1} + \Sigma_{-r}\epsilon_{-r,t+1}) + \frac{1}{2}\big(\pi_t^\top \text{diag}(\Sigma_r\Sigma_r^\top) + \pi_t^\top \text{diag}(\Sigma_{-r}\Sigma_{-r}^\top) \\ - \pi_t^\top\Sigma_r\Sigma_r^\top\pi_t - \pi_t^\top\Sigma_{-r}\Sigma_{-r}^\top\pi_t\big) - (\pi_{h,t} + \pi_{s,t} - 1)^+\theta. \end{aligned} \tag{E.18}$$

Next, let us substitute out $\pi_{n,t}$ using the fact that the interest-rate sensitivity of the portfolio equals

$$\varepsilon_{r,t} = \pi_t^\top \Sigma_r. \tag{E.19}$$

Substituting this into the wealth return, we have

$$\begin{aligned} r_{w,t+1} = r_{ft} + \pi_t^\top\mu - \varepsilon_{r,t}\epsilon_{r,t+1} + \pi_t^\top\Sigma_{-r}\epsilon_{-r,t+1} + \frac{1}{2}\big(\pi_t^\top \text{diag}(\Sigma_r\Sigma_r^\top) + \pi_t^\top \text{diag}(\Sigma_{-r}\Sigma_{-r}^\top) \\ - \varepsilon_{r,t}^2 - \pi_t^\top\Sigma_{-r}\Sigma_{-r}^\top\pi_t\big) - (\pi_{h,t} + \pi_{s,t} - 1)^+\theta. \end{aligned} \tag{E.20}$$

Because the first row of Σ_{-r} is zero (i.e., $r_{n,t+1}$ has no loadings on $\epsilon_{-r,t+1}$), and because the first element of $\mu + \text{diag}(\Sigma_r\Sigma_r^\top)/2$ equals zero (i.e., $\mu_n = -\sigma_n^2/2$), this rewritten wealth return has no dependence on $\pi_{n,t}$. We have completely replaced it with $\varepsilon_{r,t}$. Indeed, (E.20) can be written as

$$r_{w,t+1} = r_{ft} + \pi_t^{*\top}(\mu_t^* - \Sigma_t^*\epsilon_{t+1}) + \frac{1}{2}\big(\pi_t^{*\top}\text{diag}(\Sigma_t^*\Sigma_t^{*\top}) - \pi_t^{*\top}\Sigma_t^*\Sigma_t^{*\top}\pi_t^*\big) - (\pi_{h,t} + \pi_{s,t} - 1)^+\theta \tag{E.21}$$

where[30]

$$\pi_t^* \equiv \begin{bmatrix} 1 \\ \pi_s \\ \pi_h \end{bmatrix}, \quad \mu_t^* \equiv \begin{bmatrix} -\varepsilon_{r,t}^2/2 \\ \mu_s + \lambda_{sr}^2/2 \\ \mu_h \end{bmatrix}, \quad \text{and} \quad \Sigma_t^* \equiv \begin{bmatrix} -\varepsilon_{r,t} & 0 & 0 & 0 & 0 \\ 0 & \lambda_{sg} & 1 & 0 & 0 \\ 0 & 0 & 0 & \sigma_h & \tilde{\sigma}_h \end{bmatrix}. \tag{E.22}$$

This transformed expression the implies an equivalent expression for the wealth return:

$$\begin{aligned} R_{Wi,t+1} = R_{ft} + (e^{-\varepsilon_{r,it}^2/2 - \varepsilon_{r,it}\epsilon_{r,t+t}} - 1)R_{ft} + \pi_{s,it}(e^{r_{ft} + \mu_s + \lambda_{sr}^2/2 + \lambda_{sg}\epsilon_{g,t+1} + \epsilon_{s,t+1}} - R_{ft}) \\ + \pi_{h,it}(e^{r_{ft} + \mu_h + \sigma_h\epsilon_{h,t+1} + \tilde{\sigma}_h\tilde{\epsilon}_{hi,t+1}} - R_{ft}) - (\pi_{h,it} + \pi_{s,it} - 1)^+(R_{Mt} - R_{ft}) \end{aligned} \tag{E.23}$$

(To see this, simply apply the same second-order approximation to (E.23) as we did to get (E.14), and the result will be (E.20)). Thus, the optimal allocation can equivalently be stated as a choice over $\{\varepsilon_{r,t}, \pi_{s,t}, \pi_{h,t}\}$, and these choices

[30] The third element of μ_t^* is $\mu_h + \mu_{n_h} + \sigma_{n_h}^2/2$ and we have substituted in $\mu_{n_h} = -\sigma_{n_h}^2/2$.

map to an optimal long-term bond allocation equal to

$$\pi_{n,t} = \frac{\varepsilon_{r,t} - \pi_{s,t}\lambda_{sr} - \pi_{h,t}\sigma_{n_h}}{\sigma_n}. \tag{E.24}$$

This allocation is unique given the maturity n of the bond.

F Simulation of the overlapping-generations model

F.1 Aggregate shocks

Aggregate shocks The overlapping-generations (OLG) has 4 aggregate shocks that must be mapped to data: these are shocks to (1) riskfree rates ϵ_r, wage growth ϵ_g, stock returns ϵ_s, and house prices ϵ_h. Shocks to the riskfree rate come from the procedure outlined in Appendix D.1.1. This procedure also provides us with the level of real interest rates r_f. The series of wage income growth rates g are proxied for by the rate of growth of per capita real GDP, just as in Appendix D.1, . We obtain g_1 using the relation

$$g_{1t} = g_t - \lambda_{rg} r_{ft}. \tag{F.1}$$

Stock returns come from the the S&P500 and are added directly to the simulation after subtracting out the portion driven by shocks to interest rates—which is akin to to subtracting $\sigma_{n_s}\epsilon_r$ from the log stock return. This means that we add both ϵ_g and ϵ_s into the simulation simultaneously. Similarly, aggregate house prices come from the Case-Shiller index is added after differencing out the portion of the capital gain attributable to interest rate shocks $\sigma_{n_h}\epsilon_r$. For the idiosyncratic portion of house price shocks, we assign each simulated individual a shock drawn from normal distribution with variance $\tilde{\sigma}_h^2$.

When not using data in the simulation, we randomly assign agents aggregate shocks from 50 randomly generated time series. When computing top wealth shares, we do so within a group of shocks and then take the average top 10% share among the 50 groups.

F.2 Inheritance and intergenerational income persistence

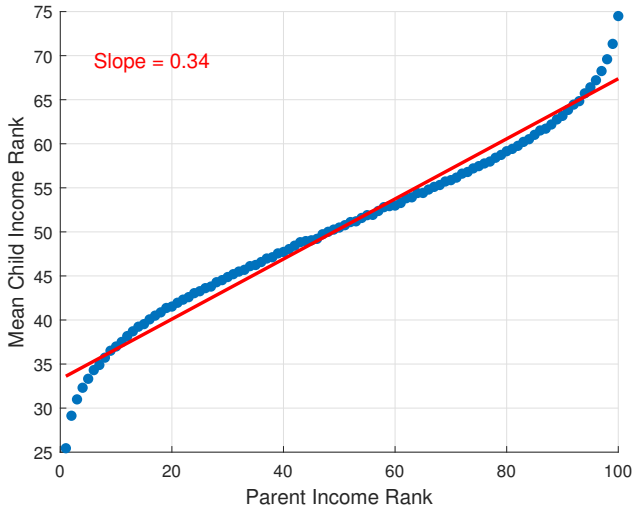
To model intergenerational income persistence, we assume a correlation between the levels of persistent income of parents and their children at labor market entry. Specifically, we model initial persistent income as an AR(1) process:

$$z_{i,t_0} = \rho_{z_0} z_{t_0-30}^{\text{parent}_i} + \sigma_{z_0}\epsilon_{z_0}, \tag{F.2}$$

where $z^{\text{child}}t_0$ denotes the child's initial value of z upon entering the labor force, and $z^{\text{parent}i}t_0 - 30$ is the corresponding value for the parent 30 years earlier. We calibrate the parameters ρz_0 and σ_{z_0} using two moment conditions. First, the

variance of z_{i,t_0} must equal σ_α^2, a parameter chosen to match the observed level of income inequality. Second, we set ρ_{z_0} to reproduce the empirical slope of the relationship between the child's and the parent's income ranks at age 30, as documented in Chetty et al. (2014). Given the rest of our labor income process, these conditions imply $\rho_{z_0} = 0.4564$ and $\sigma_{z_0} = 0.9787$. Figure F.3 plots the child's income rank against the parent's income rank, providing the model counterpart to Figure II in Chetty et al. (2014).

Figure F.3: Child income rank vs parent income rank in the model



G 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.

G.1 Linearized conditions

Suppose that there is no intertemporal income or housing, so the budget constraint (35) simplifies to[31]

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

[31] It would be fine to assume rented housing, in which case consumption here represents total consumption of both goods and housing. We abstract from this for simplicity.

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

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

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 (G.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

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

where $\rho_c(m_t) = \beta(1-m_t)$ and $\kappa_w(m_t) = \log \rho_c(m_t) + (1-\rho_c(m_t)) \log\left(1-\rho_c(m_t)\right)/\rho_c(m_t)$.[32] (Notice that, as $m_t \to 1$, $c_t \to w_t$; agents who will die almost surely consume everything.) We can also get the linearized approximation to the log wealth return

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

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

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

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

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

Using the decomposition

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

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

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

[32] 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.

Substituting this and the fact that

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

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

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

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

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

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

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

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

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

G.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

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

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 (G.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 (10), is our expression for the optimal share in the n-period bond (42).

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

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

It also implies

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

Substituting these three equations into (G.11), we have

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

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

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

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

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

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

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

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

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

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

for the deterministic constant

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

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

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

This verifies the conjecture.

G.3 Adding labor income and Social Security

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

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

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

where

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

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

that the rate sensitivity implied by π^H is

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

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

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

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

where the value weights take the form

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

the difference between the benefits claim and the tax liability.

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

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

(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

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

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 is

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

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

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

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

where

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

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

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

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

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

where π_t^* is the optimal solution without intertemporal income. Substituting this into (G.18) and solving for π_t, 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 W_t 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 (44).

G.4 Solution to the value function

This section solves for the value function under optimal policies in closed form. Under recursive preferences, the transformed value function (53) is implicitly defined by the aggregator

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

By Euler's Theorem,

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

where

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

and

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

Noting that the stochastic discount factor

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

and that

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

we have that

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

Iterating this recursion forward yields

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

since total wealth is the present value of consumption. Substituting the expression for $\partial U_t/\partial C_t$ and noting that consumption is at an optimum ($C_t = C_t^*$), we get the solution (54).

G.5 Adding multiple assets

Consider the setting with J non-redundant risky assets in Section 4.3, summarized by (46) and (47). In this case, the linearized wealth return becomes

$$r_{w,t+1} = r_{ft} + \pi_t^\top(r_{t+1} - r_{ft}\iota) + \frac{1}{2}\left(\pi_t^\top \text{diag}(\Sigma^\top\Sigma) - \pi_t^\top\Sigma^\top\Sigma\pi_t\right).$$

The log-linearized Euler equations are the same for each asset j, so by the exact same logic, we have the vector of optimal portfolio weights

$$\pi_t^* = \frac{1}{\gamma}(\Sigma^\top\Sigma)^{-1}\left(\mathbb{E}_t[r_{t+1} - r_{f,t+1}\iota] + \frac{1}{2}\text{diag}(\Sigma^\top\Sigma)\right) - \left(1-\frac{1}{\gamma}\right)(\Sigma^\top\Sigma)^{-1}\frac{\text{cov}_t(r_{t+1}, c_{t+1} - w_{t+1})}{1-\psi}.$$

Also by the same logic, the consumption-wealth ratio $c_t - w_t$ continues to take the same form, namely affine in r_{ft}. Consequently, substituting in the expected log returns and the fact that

$$\text{cov}_t(r_{t+1}, r_{f,t+1}) = \Sigma^\top\sigma_r,$$

we get the optimal portfolio rule π_t^* and the resulting interest-rate sensitivity of wealth $\varepsilon_r^*(W_{t+1})$ stated in the main text.

G.6 Optimal consumption plan in the limit

This section derives the optimal consumption-investment strategy in the limit as risk aversion γ approaches infinity and the EIS ψ approaches zero. We first show that, in the limit as $\gamma \to \infty$ and $\psi \to 0$, the optimal consumption plan is constant and financed by a perpetuity. We then prove that both of these conditions (i.e., both $\gamma \to \infty$ and $\psi \to 0$) are necessary for this to be true—that is, it is insufficient to just have $\gamma \to \infty$ and not $\psi \to 0$ or vice versa.[33]

First, let us consider what happens when $\gamma = 1/\psi \to \infty$. The first-order condition when $\gamma = 1/\psi$ is

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

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

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

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

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

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

where $t_{\max}$ is the first year in which $m_t = 1$.[34] This expression pins down the value of $\bar{C}$. Because the optimal consumption plan is deterministic and constant, the agent finances it by purchasing $\bar{C}$ of each zero-coupon bond and consuming the coupons.

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

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

[33] This result was first shown by Campbell and Viceira (2001). An alternative (and closely related) proof can be found in their appendix.

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

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 (40) implies the (negative) elasticity

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

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

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

Substituting this into the expression for the optimal portfolio $\bar{\pi} = \pi^*$ in (42), 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 (G.22), as claimed.

Using similar arguments, we can show that $\gamma \to \infty$ and $\psi \to 0$ are both necessary to obtain this result. Suppose, for the sake of contradiction, that $\gamma \neq 1/\psi$ and the consumption plan is constant. Before taking any limits, the optimal consumption rule implies

$$\frac{\partial \log(\bar{C}/W_t)}{\partial r_{ft}} = (1-\psi)\varrho_{rt}.$$

Because the consumption plan is constant, this condition and (G.21) then imply

$$(1-\psi)\varrho_{rt} = \sum_{j=0}^{t_{\max}-t} \frac{P_{jt}}{\sum_{j'=0}^{t_{\max}-t} P_{j't}} \left(\frac{\sigma_j}{\sigma_r}\right),$$

which in turn implies

$$\bar{\pi}_t = \left(1-\frac{1}{\gamma}\right)\varrho_{rt}\frac{\sigma_r}{\sigma_n} = \frac{1-1/\gamma}{1-\psi}\sum_{j=0}^{t_{\max}-t} \frac{P_{jt}}{\sum_{j'=0}^{t_{\max}-t} P_{j't}} \left(\frac{\sigma_j}{\sigma_n}\right).$$

But then this can only be equal to $\tilde{\pi}$ from (G.22) if $\gamma = 1/\psi$.[35] Thus, the only way for a constant consumption plan to be consistent with the optimal policies is to have both $\gamma \to \infty$ and $\psi \to 0$ simultaneously.

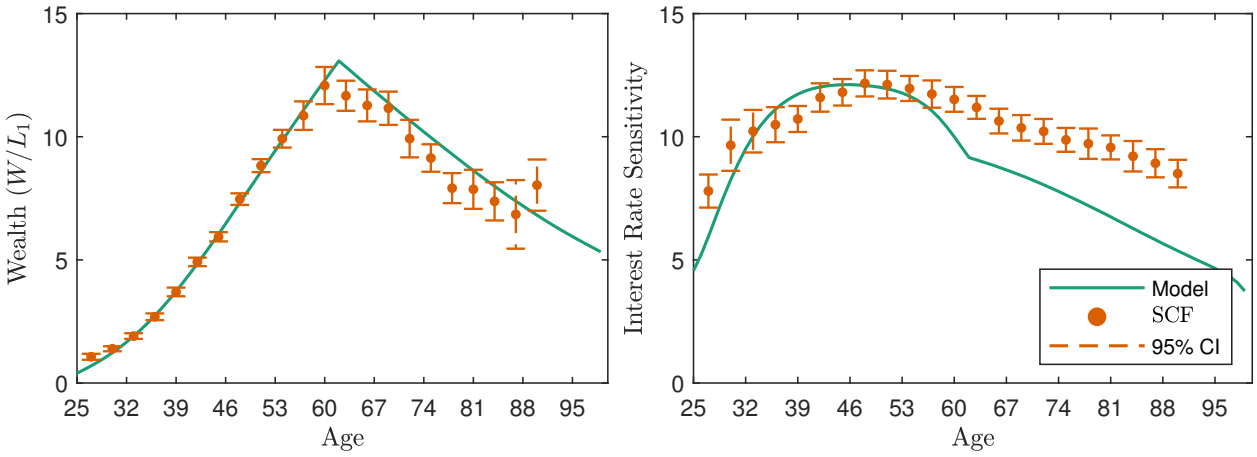
[35] This implies that $\gamma = 1/\psi$ is necessary, but not sufficient, for a constant consumption plan.

H Additional results

H.1 Model fit with entrepreneurs

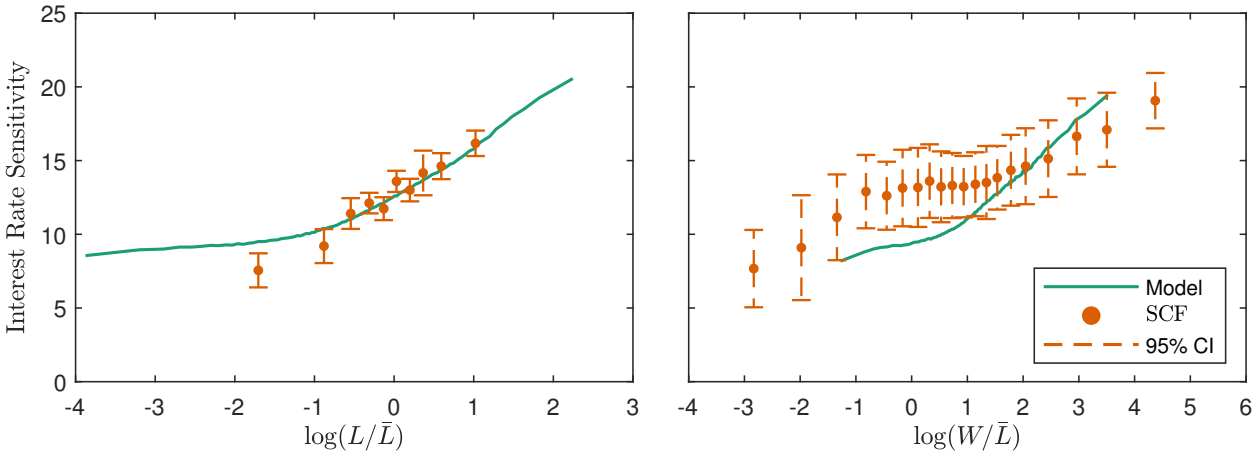
H.2 Higher persistence of interest-rate shocks

Figure H.4: Life-cycle profiles of wealth and its interest-rate sensitivity for $\varphi = .9756$



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

Figure H.5: Interest-rate sensitivity of wealth at age 40–45 for $\varphi = .9756$

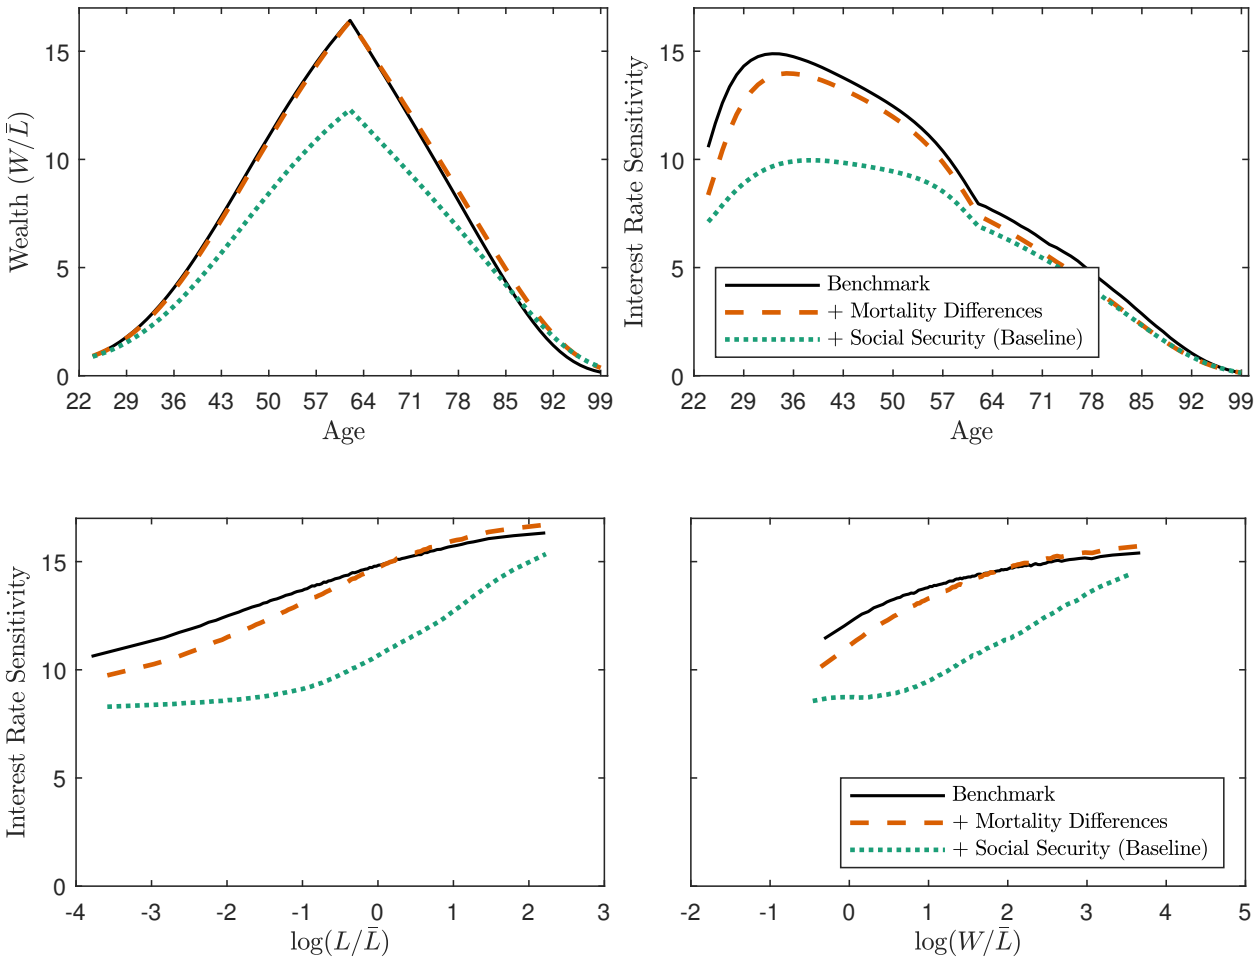


Note: This figure reports the relationship between the interest-rate sensitivity of wealth and wealth (left panel) and earnings (right panel). In the data, wealth and earnings are computed per adult and scaled by the Social Security wage index. In the left panel, each bin represents a decile of earnings.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

H.3 Mechanism

The quantitative model validates our economic intuition and allows us to study counterfactuals. To shed more light on the full numerical model, this section analyzes the importance of two novel mechanisms in our model: income-based differences in mortality rates and the presence of Social Security. Figure H.6 plots quantities of interest with and without these features.

Figure H.6: 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 H.6. 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 H.6, 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.