Financial and Total Wealth Inequality with Declining Interest Rates*

Daniel L. Greenwald — NYU Stern

Matteo Leombroni — Boston College

Hanno Lustig — Stanford GSB, NBER

Stijn Van Nieuwerburgh — Columbia Business School, NBER, CEPR, ABFER

September 2, 2023

Abstract

Financial wealth inequality and long-term real interest rates track each other closely over the post-war period. We investigate how much of the increase in measured financial wealth inequality can be accounted for by the decline in rates, and study the implications for inequality in total wealth (lifetime consumption). To do so, we measure the exposure of household-level financial portfolios to interest rates. We find enough heterogeneity in household portfolio revaluations to explain the entire rise in financial wealth inequality since the 1980s. A standard incomplete markets model calibrated to these data implies that the low-wealth young lose when rates decline, while the high-wealth old gain.

JEL: E21, E25, E44, G12

Keywords: wealth inequality, interest rates, secular stagnation, human wealth, duration

*First version March 2021. Greenwald: Department of Finance, NYU Stern School of Business, 44 W 4th Street, New York, NY 10012; dlg340@stern.nyu.edu. Leombroni: Department of Finance, Boston College Carroll School of Management, 140 Commonwealth Ave, Chestnut Hill, MA 02467; matteo.leombroni@bc.edu. Lustig: Department of Finance, Stanford Graduate School of Business, Stanford CA 94305 and NBER; hlustig@stanford.edu. Van Nieuwerburgh: Department of Finance, Columbia Business School, Columbia University, 3022 Broadway, New York, NY 10027, NBER, CEPR, and ABFER; svnieuwe@gsb.columbia.edu.

The authors thank Sebastian Hanson and Yuan Tian for excellent research assistance. We thank Andy Atkeson, Adrien Auclert, Ravi Bansal, Max Croce, Atif Mian, Ben Moll, Sylvain Catherine, Peter DeMarzo, Alp Simsek, and seminar participants at Stanford GSB, Columbia GSB (econ and finance), Duke Fuqua Research Triangle seminar, JHU Econ and Finance, MIT Sloan, University of Texas at Dallas, the New York Federal Reserve Bank, the University of Chicago, UC Davis, USC Marshall, New York University, the ECB, and INSEAD.

Over the post-war period, interest rates and financial wealth inequality have displayed a remarkable negative correlation. This time series relationship is displayed in Figure 1, which compares the share of financial wealth held by the top-10% of the financial wealth distribution, against the implied price of a 10-year inflation-adjusted zero coupon bond. As can be seen, the top-10% wealth share falls as the real 10-year bond price falls (yield increases) from the 1960s to the 1980s, and rises as real bond prices rise (yield falls) from the 1980s through the end of our sample in 2019. In terms of magnitudes, real yields rose from an average of 0.17% in the 1950s to an average of 4.94% in the 1980s, before falling to an average of 0.63% in the 2010s. Over the same period, the top-10% share of financial wealth fell from 70.4% in 1947 to 62.4% in 1983, before rising to 70.8% in 2019. We find evidence of broad-based declines in expected real returns across asset classes, not only on bonds but also on stocks and housing.

Figure 1: Top-10% Wealth Share vs. 10-Year Real Bond Prices

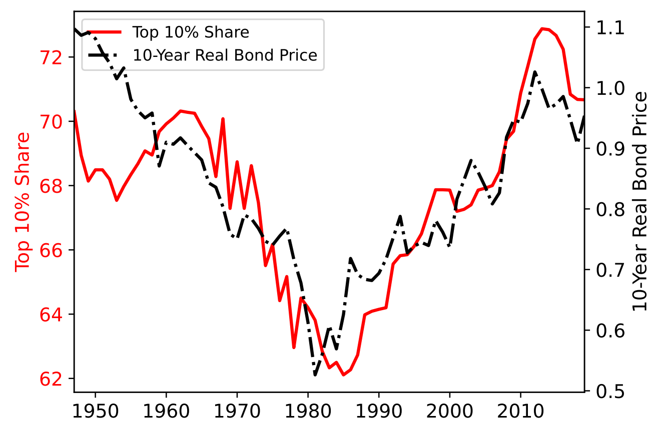

Note: The red solid line displays the top-10% financial wealth share for the United States, obtained annually from 1947 until 2019 from the World Inequality Database. The black dash-dot line displays an estimate of the 10-year real bond price, obtained from a dynamic affine term structure model, estimated on quarterly data from 1947.Q1–2019.Q4 (see Section 2 and Appendix A for details).

Since discount rates have a direct link to the valuation of financial assets, a natural hypothesis is that falling rates cause rising financial wealth inequality. Characterizing this channel, however, requires overcoming two obstacles. First, the ultimate impact on financial wealth inequality depends not on the average effect of discount rates on the value of financial wealth, but on the *heterogeneity* of these revaluations across the population, and how this heterogeneity covaries with initial levels of wealth. Second, to the extent that discount rates influence financial wealth inequality, whether the resulting gains and losses occur only "on paper" (revaluing the same consumption stream) or actually influence future consumption and welfare is far from clear (Moll, 2020).

In this paper, we study the link between real interest rates and wealth inequality to answer two research questions: one positive, one normative. On the positive side, what share of the observed rise in financial wealth inequality can be accounted for by falling interest rates? On the normative side, what are the implications for inequality in the consumption possibilities that actually determine welfare?

We begin by developing intuition using theoretical insights. Our analysis centers on *cash flow duration*, which summarizes the sensitivity of an asset's value to changes in long-term interest rates. A fall in rates will increase financial wealth inequality if household portfolio durations are increasing in wealth, for which a sufficient condition is that aggregate (value-weighted) duration in the economy exceeds average (equal-weighted) duration. We find that this condition is satisfied in U.S. data. We provide a simple formula mapping durations and changes in rates into top wealth shares, and find that it predicts a rise in inequality close to that observed in the data.

Building on work by Auclert (2019), we also show that a household's consumption is unaffected by a small change in interest rates if and only if the duration of its lifetime excess consumption plan is equal to the duration of its financial wealth portfolio. This result allows us to measure in a simple and robust way whether and how much a change in rates influences household lifetime consumption. Households with too little duration in their financial portfolio are worse off when rates decline, while households with sufficiently high financial duration gain.

To develop these insights into quantitative results, we combine a set of novel empirical estimates with a structural model. Our main empirical contribution is to measure households' financial wealth duration. We use microdata from the Survey of Consumer Finances to characterize households' portfolio allocations across asset classes. We then use asset pricing data to assign each asset class a cash flow duration at every point in time. For private business wealth, a key portfolio component for the wealthiest households, we use data on corporate status to separate small and potentially stagnant private businesses from larger and faster-growing ones. By combining observed household portfolio shares with durations for each asset and liability class, we obtain the duration of financial wealth portfolios at each point in time.

Examining the cross-section, we observe substantial heterogeneity in financial wealth durations by wealth level and age. Low-wealth households have low financial durations, due to higher portfolio shares in deposit-like assets, vehicles, and housing. High-wealth households have high financial durations, driven by higher portfolio shares in public equity and private business wealth. Conditional on wealth, financial durations are declining in age. Our paper is the first to document this heterogeneity in financial duration, a crucial ingredient in the transmission of long-term real rate shocks to wealth inequality.

We pair these empirical estimates with a calibrated life-cycle model, which serves two key functions. First, the model accounts for savings behavior, which influences how a change in rates at one point in time dynamically affects the wealth distribution in the future. Second, the model maps observed changes in wealth and interest rates into changes in the household's consumption plan, allowing us to speak to inequality in total wealth (the present value of lifetime consumption) and consumption opportunities.

The model features a bequest motive, rich idiosyncratic income risk calibrated to the Panel Survey of Income Dynamics, and a superstar income state that enables it to exactly match the top-10% financial wealth share in 1983. To capture our key empirical findings, we calibrate heterogeneity in the duration of financial wealth to match our empirical estimates by wealth bin and age.

We first answer the positive question: how has the secular decline in real rates between 1983 and 2019 affected financial wealth inequality? We initialize the model at a long-term interest rate of 4.94%, which we estimate to have prevailed in 1983, then feed in a sequence of unexpected permanent interest rate changes matching the actual 1983–2019 time path. We revalue all assets after each year's rate change using the actual distribution of household portfolios and estimates of asset durations obtained from the data.

The resulting path of the wealth distribution, which we denote the *repriced* distribution, displays a rise in the top-10% financial wealth share of 7.9pp between 1983 and 2019, explaining 95% of the observed 8.3pp rise in the data. The model similarly explains 57% of the rise in the top-1% share, and 113% of the rise in the Gini coefficient. These results are the net effect of an even larger rise in inequality caused by revaluations around interest rate changes, partially offset by household consumption and savings responses between rate changes that gradually pull inequality toward a new steady state level *lower* than observed in the 1980s.

Beyond financial wealth, we also quantify how these interest rate changes affect total wealth, equal to the sum of financial and human wealth. We find a lower initial level and a smaller rise in total wealth inequality than in financial wealth inequality. This occurs because young households, while typically lacking in financial wealth, are rich in long-duration human wealth, which increases sharply in value when rates fall. However, the same qualitative patterns for the repriced distribution of financial wealth also hold for total wealth.

We next evaluate the robustness of our results for the repriced distribution. Alternative assumptions regarding the duration of housing and private business wealth deliver a the rise in the top-10% financial wealth share that ranges from 6.8pp to 12.3pp. Excluding the primary home from the duration calculation results in a rise in the top-10% wealth share of 8.7pp, while including the observed rise in income inequality results in an increase of 12.5pp. Time variation in asset durations contributes a meaningful 1.4pp of the overall increase, while time variation in portfolio shares contributes only 0.3pp.

Turning to our normative question, we measure how the observed path of rates has affected household total wealth and consumption. To do this, we compute the *compensated* financial wealth distribution: the level of financial wealth that each household would require to keep its consumption plan unchanged following each movement in rates. Our theoretical results imply that household consumption is unaffected by rates if and only if the compensated and repriced distributions are identical. Instead, we find large deviations between the two, implying major effects on consumption and total wealth inequality. In fact, to ensure that all households could afford their pre-shock consumption plans, we would have needed to see financial wealth inequality *decline* over our sample, with the top-10% share falling by 1.5pp.

This large mismatch between the compensated and repriced distribution implies that falling interest rates have not merely adjusted financial wealth "on paper." Instead, there are winners and losers from falling rates, with life-cycle dynamics playing a key role. The young plan to save in middle age and dissave in retirement, giving them a high duration of excess consumption. However, young agents on average have little financial wealth accumulated and not enough financial wealth duration to fully hedge this excess consumption plan. As a result, these households see consumption possibilities contract when rates fall. Intuitively, they find it more challenging to accumulate wealth for retirement in the absence of high returns. In contrast, older households who have already accumulated wealth benefit when rates fall, earning large capital gains on their assets that more than offset lost investment income going forward.

This heterogeneity in the effect of rate changes across the age distribution implies that the total impact of falling rates on an individual household's consumption and welfare depends not only on how much rates fall during its lifetime, but also on exactly when in its life cycle these falls occur. We use the model to study which cohorts gained and lost from the fall in rates between the 1980s and 2010s. We find that households born in the 1920s through the 1940s gained substantially from falling rates, as these households had largely accumulated their peak financial wealth by the time rates fell in the 1980s. In contrast, households born in the 1960s or later generally lost from declining rates. This drop in consumption is severe for recent cohorts, with households born in the 2000s losing more than 8% of lifetime consumption at birth due to the decline in rates.

Related Literature. Our paper joins a large body of work on the evolution of financial wealth inequality. Empirical evidence suggests that financial wealth inequality has increased in many countries over the past several decades.[1] To account for these facts, the literature has often adopted a backward-looking approach that explores the connection between past returns and current wealth, identifying heterogeneity in past rates of return as a key driver of the increase in financial wealth inequality.[2] However, financial wealth is also a forward-looking valuation metric (the present discounted value of future consumption minus income) so that discount rates could matter quantitatively for wealth inequality. Our paper explores the relation between discount rates, wealth inequality and welfare. In doing so, we bring a novel asset pricing perspective to the discussion on wealth inequality, as well as a new mechanism — differences in duration across households — explaining variation in past realized returns across households.

[1] See e.g., Piketty and Saez (2003); Piketty (2015); Alvaredo, Chancel, Piketty, Saez, and Zucman (2018).

[2] See e.g., Piketty and Zucman (2015); Benhabib, Bisin, and Luo (2017); Cox (2020); Fagereng, Guiso, Malacrino, and Pistaferri (2020); Bach, Calvet, and Sodini (2020); Hubmer, Krusell, and Smith (2020).

Our paper is closely related to Catherine, Miller, and Sarin (2020), who show that accounting for the revaluation of Social Security benefits — an implicitly held asset with very long duration — can dramatically influence measured changes in financial wealth inequality over time. In this paper, we take a broader view, measuring both financial and total wealth inequality, the latter of which incorporates all forms of income, including Social Security benefits. Further, our results on consumption and the compensated distribution allow us to speak to what measured inequality *should be* to keep consumption opportunities constant, beyond this measurement.

We also build on Auclert (2019), who derives the effect of transitory shocks to interest rates on consumer welfare and consumption, and similarly finds that it depends on the structure of household portfolios. For this transitory shock, the sufficient statistic is unhedged interest rate exposure (URE) — the net difference between assets and liabilities that pay in the future rather than today — which does not distinguish between future cash flows arriving at different times (e.g., two-year vs. five-year zero coupon bonds). In contrast, we consider the impact of a *permanent* shock to interest rates on consumption and wealth *inequality*, for which the exact timing of the cash flows is critical. In our context, the sufficient statistics are the duration of financial wealth and excess consumption, which for a permanent shock can drive variation among households with the same URE.[3] We see these works, studying distinct objects following interest rate shocks of different persistence, as highly complementary.

For an alternative to our duration measure, recent work by Fagereng, Gomez, Gouin-Bonenfant, Holm, Moll, and Natvik (2022) proposes using future net asset purchases to gauge the welfare effects of real rate declines. They apply this measure to Norwegian household-level data, where housing is the dominant form of wealth, and show that their flow-based approach is equivalent to our duration approach following a permanent shock to interest rates. While this flow-based approach is ideal to study the effects of total asset price changes, our work considers the contribution of a *single factor* driving asset prices — falling interest rates — on inequality over time. This allows us to use duration as a sufficient statistic, making our measurement exercise possible in a US context that lacks the detailed transaction data used by Fagereng et al. (2022). Last, our model allows us to study the dynamic effects of changes in valuations over time, as they propagate through consumption and savings behavior, rather than on impact. We again view these differing research questions and methodologies as highly complementary.

In additional work, Gomez and Gouin-Bonenfant (2020) also study the effects of lower interest rates on inequality, through their impact on the cost of raising new capital for entrepreneurs. This investment-based channel operates through growth in real assets and cash flows, complementing our duration-based mechanism operating through financial revaluation of cash flows.

[3]The working paper version of Auclert (2019) features numerical exercises showing that the impact of persistent monetary policy shocks on aggregate consumption depends on asset durations. We complement this work by formalizing and estimating sufficient statistics for the response of consumption and wealth inequality to permanent shocks.

Kuhn, Schularick, and Steins (2020) use novel microdata to document that heterogeneity in the shares of equity and housing in household portfolios drove much of the rise in inequality since the 1970s. We build on this work by constructing new and detailed statistics on the duration of financial wealth for US households, which summarize how declining interest rates over this period translated into changes in the prices of these assets. Our approach is also closely linked to that of Doepke and Schneider (2006), who focus on the distributional consequences of inflation. We follow these authors in using household-level portfolio data to measure the exposure to the shock of interest, but focus on the effects of changes in long-term real rates rather than inflation.

Additional recent work analyzes the measurement challenges for private business income and wealth.[4] In our empirical work, we advance this agenda by producing several new measures of private business wealth that combine microdata from the Survey of Consumer Finances with asset pricing data from a range of data sources. We perform various exercises to establish the robustness of the results to the specifics of the private business wealth duration estimation.

Last, our paper links to the rich literature on the mechanisms behind the decline in interest rates over our sample.[5] Our conclusions regarding the *consequences* of low interest rates for wealth inequality stemming from asset revaluations should not depend on this fundamental source, as a change in discount rates will affect the valuation of a fixed stream of cash flows in the same way regardless of its ultimate cause. We thus view our work as complementary to, but distinct from, this important literature.

Overview. The rest of the paper is organized as follows. Section 1 derives the link between cash flow duration and financial wealth inequality. Section 2 presents our key empirical facts on expected returns and household portfolio durations. Sections 3 constructs the model, while Section 4 calibrates it and analyzes its stationary distribution. Sections 5 and 6 discuss the main quantitative results for the repriced and compensated wealth distributions, respectively. Section 7 computes the impact of rates on total wealth by birth cohort. Section 8 concludes.

The paper also includes a comprehensive appendix. Appendix A provides an auxiliary asset pricing model used to infer real interest rates and expected returns on the components of financial wealth. Appendices B, C, and D contain details on the data, construction of the financial wealth portfolio, and the duration of household assets and liabilities. Appendix E discusses income data and estimation. Appendix F contains proofs of our propositions in the main text. Appendix G presents a general equilibrium model with aggregate risk and endogenous interest rates that generalizes our baseline model. Appendix H displays supplementary model results.

[4] See e.g., Kopczuk (2017); Saez and Zucman (2016); Piketty, Saez, and Zucman (2018); Smith, Yagan, Zidar, and Zwick (2022.); Kopczuk and Zwick (2020).

[5] See e.g., Bernanke (2005); Caballero, Farhi, and Gourinchas (2008); Summers (2014); Eggertsson and Mehrotra (2014); Eichengreen (2015); Gutiérrez and Philippon (2017); Mian, Straub, and Sufi (2020).

1 Duration and Inequality

In this section we define cash flow duration, display its basic properties, and derive our first key theoretical result linking variation in duration across the wealth distribution to the effects of declining interest rates on inequality. We consider the response to an unexpected and permanent change in the interest rate from R to $\tilde{R} = R\exp(\varepsilon)$ for some shock ε, where $\varepsilon < 0$ corresponds to a fall in rates. Throughout the paper we will use tildes (e.g., $\tilde{R}$) to indicate updated values following an interest rate shock, while equivalent variables without tildes indicate pre-shock values. Proofs of all propositions can be found in Appendix F.

Definition 1. Assume that the annualized discount rate is equal to r at all maturities, and let $R = 1 + r$. Given a sequence of cash flows $\{x_t\}$, define its present value at $t = 0$ by

$$P_0 = \sum_{t=0}^{\infty} R^{-t} x_t \tag{1}$$

The *cash flow duration* (or simply, duration) of this stream of cash flows is defined by

$$D \equiv \frac{\sum_{t=0}^{\infty} t \times R^{-t} x_t}{P_0}. \tag{2}$$

In words, cash flow duration is the weighted average of the time remaining to receive each cash flow, weighted by the share of the asset's value attributable to that cash flow. Given this definition, we now present the key properties of this statistic.

Proposition 1. Consider a sequence of cash flows $\{x_t\}$ with present value P_0. Then

$$\frac{\partial \log P_0}{\partial \log R} = -D$$

where D satisfies (2). For a small shock $\varepsilon \to 0$, this implies the approximate revaluation

$$\tilde{P}_0 \simeq P_0 \exp(-D \times \varepsilon) \simeq P_0(1 - D \times \varepsilon). \tag{3}$$

For a portfolio of assets indexed by k, equation (3) holds using the value-weighted duration

$$D^{VW} \equiv \sum_k \omega(k) D(k) \tag{4}$$

where $\omega(k)$ is the share of the portfolio's value in asset k, and $D(k)$ is the duration of asset k.

This proposition derives the well-known property that for a small, permanent shock to interest rates, duration is a sufficient statistic for the resulting change in present value. We now use this property to derive our main theoretical result linking interest rates and financial wealth inequality.

Proposition 2. For a small negative change in rates ($\varepsilon < 0$):

(a) Household wealth growth due to revaluation ($\tilde{\theta}/\theta$) has a positive covariance with wealth if and only if financial wealth duration (D^{θ}) has a positive covariance with financial wealth (θ). A sufficient condition for this positive covariance is that aggregate (value-weighted) financial wealth duration in the pre-shock economy exceeds average (equal-weighted) duration.

(b) The top-α share of wealth S^{α} increases if and only if D^{top}, the value-weighted duration for the top-α wealthiest share of households exceeds D^{bottom}, the value-weighted duration for the bottom $1-\alpha$ share of the wealth distribution (or equivalently, exceeds the overall value-weighted duration D). The change in the top-α wealth share is approximated by

$$dS^{\alpha} \simeq -S^{\alpha}(1-S^{\alpha})(D^{top} - D^{bottom})\varepsilon. \quad (5)$$

This proposition conveys an important intuition: while a decline in discount rates pushes up the value of all assets, whether or not it increases wealth *inequality* depends on whether the financial wealth portfolios of the rich grow by more than those of the poor. This is in turn determined by the relative exposures of these portfolios to interest rates, summarized by cash flow duration.

The proposition also allows us to use simple summary statistics to approximate the effects of a decline in interest rates on inequality. To do so, however, requires that we are able to measure the distribution of duration in the data, motivating our empirical exercise in the next section.

2 Wealth Inequality and Real Rates: Empirical Evidence

In this section we document three key empirical facts. First, we show evidence for a large decline in long-term real rates and expected returns on risky assets more broadly. Second, we provide additional support for a strong time-series correlation between long-term real interest rates and wealth inequality, not only in the U.S. but also in the United Kingdom and France. Third, we present our main empirical finding of the paper: household financial wealth portfolios have highly heterogeneous durations that correlate positively with the level of financial wealth.

2.1 Decline in Real Rates

We begin by documenting a broad-based decline in expected returns across all major asset classes. To do so, we estimate an auxiliary no-arbitrage asset pricing model, detailed in Appendix A. This statistical model prices bonds of various maturities, both nominal and real, the aggregate stock market, several cross-sectional stock market factors including small, growth, value, and infrastructure stocks, and households' housing wealth. According to this model, the ten-year real bond

Table 1: Expected Real Returns Decade Averages

Asset	1980s	2010s	Decline
Ten-year real bond yield	4.94%	0.63%	4.31%
Aggregate stock market	7.98%	2.00%	5.98%
Housing wealth	8.24%	4.89%	3.35%
Growth stocks	5.21%	3.53%	1.68%
Value stocks	18.50%	7.19%	11.31%
Infrastructure stocks	11.75%	2.35%	9.40%
Small stocks	3.57%	3.18%	0.39%

Note: The table reports model-implied real expected real returns and average them over the 40 quarters in the 1980s and the 40 quarters of the 2010s. The model that generates these statistics is detailed in Appendix A.

yield averaged 4.94% over the 1980s decade and 0.63% over the 2010s decade.[6] Table 1 displays the average expected real returns over these decades for a broader set of assets, showing similarly large declines for the aggregate stock market and housing wealth. This pattern displays some variation, with returns on value and infrastructure stocks showing larger declines, and returns on growth and small stocks showing smaller declines. Overall, however, the decline in expected returns is substantial and broad-based.

2.2 Increased Wealth Inequality

We next supplement Figure 1 with additional evidence showing that the strong negative comovement between financial wealth inequality and long-term real interest rates is robust across wealth measures, interest rate measures, and countries. To this end, Figure 2 compares top wealth shares and the prices of 30-year real annuities across the U.S., U.K., and France over the period 1947–2019.[7] For the U.S., we augment our series with wealth shares constructed from the Survey of Consumer Finances (SCF), augmented by the SCF+ data of Kuhn, Schularick, and Steins (2020). We construct the annuity price using either real discount rates from our auxiliary asset pricing model, or by combining nominal yields with inflation forecasts (see Appendix B for details). Since annuity prices move inversely with long-term real rates, we predict a *positive* comovement between annuity prices and inequality.

For both inequality measures, we observe a strong positive correlation between top wealth shares and the price of a long-term real annuity (equivalently, a negative correlation with long-term real rates). Between 1947 and 1983, the top-10% (top-1%) wealth share falls by 8.0pp (5.2pp) in the U.S. as the annuity becomes cheaper. From 1983 until 2019, the top-10% (top-1%) wealth

[6] The asset pricing model matches the available data on Treasury Inflation-Indexed Securities over the period for which they are available. The model-implied yield changes are similar for real bonds of different maturities.

[7] For France we start our sample in 1950 since inflation was very high coming out of WW-II, resulting in implausible real bond yield estimates.

Figure 2: Top Financial Wealth Inequality and Cost of Real Annuity

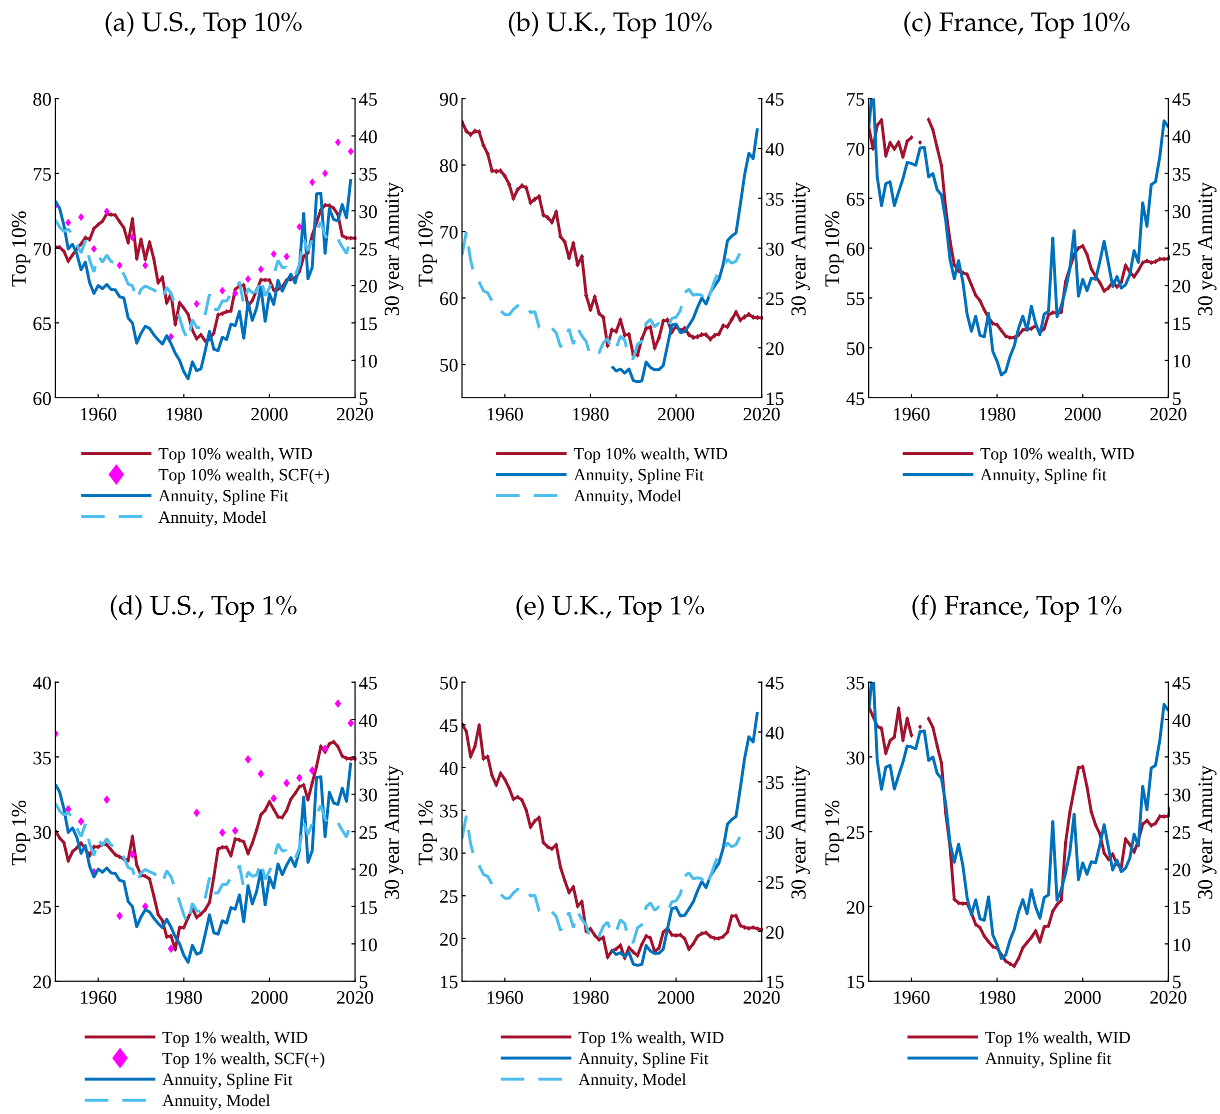

Note: Each panel plots a financial wealth inequality measure against a measure of the cost of a 30-year real annuity. The inequality measure in the left panels is the share of financial wealth going to the top-10% of the population. The right panels plot the share of the top-1% of the population. The wealth shares are from the World Inequality Database (and the SCF+ for the U.S.). Details on annuities and wealth shares in Appendix B.

share rises by 8.3pp (11.3pp) as the cost of the annuity more than doubles. Computing U.S. top wealth shares using the SCF and SCF+ (displayed as pink diamonds) yields a nearly identical pattern. In the U.K. and in France we observe the same general pattern, although the top-10% share falls by more than in the U.S. in the period before 1983, and rises by less in the period after 1983. The top-10% wealth share increases by 4.3pp in the U.K. and by 6.5pp in France from 1984 until 2021.[8] These results show that the strong negative correlation between inequality and real rates displayed in Figure 1 holds across countries and across definitions of top wealth shares.

2.3 Household Heterogeneity in Financial Duration

While declines in real rates increase asset prices, this only matters for wealth *inequality* if household exposures to real rates vary systematically with wealth. In this section, we measure portfolio exposures at the household level, and document how they vary across the wealth distribution.

Proposition 1 above shows that a household's exposure to exposure to a permanent change in real discount rates is summarized by the cash flow duration of its portfolio. From equation (4), we can compute the duration of household i's financial wealth portfolio at time t, denoted $D^{\theta}_{i,t}$, as

$$D^{\theta}_{i,t} = \sum_{k} \omega_{i,t}(k) D_t(k). \tag{6}$$

where $\omega_{i,t}(k)$ is the share of household i's portfolio in asset k at time t (negative for liabilities), and $D_t(k)$ is the duration of asset k at time t. Thus, to measure a household's duration, we combine measures of that household's portfolio shares with measures of duration for each asset.

Measuring Portfolio Shares. For data on portfolio shares, we use the Survey of Consumer Finances (SCF), a detailed survey of household wealth conducted every three years. Since these data do not distinguish holdings at the individual asset level, we consider portfolios constructed across asset *classes*, for which the SCF provides detailed holdings, allowing us to directly measure $\omega_{i,t}(k)$ for each household in each SCF wave (see Appendix C for details). These asset classes include public equities (held both directly and indirectly in mutual funds and pension accounts), real estate, corporate and non-corporate private business wealth (PBW), vehicles, fixed income assets (held both directly and indirectly), and cash, deposits, and money market instruments. On the liabilities side, we also include mortgages, vehicle debt, student debt, and other debt. Since the SCF begins in 1989, after interest rates have already been falling for some time, we supplement our data with the 1983 wave of the SCF+.

The resulting portfolio shares, averaged across all survey waves, can be seen in Table 2. The four right columns display wealth-weighted average portfolio shares for the full sample, bottom-90%, middle 90%-99%, and top-1% of the wealth distribution. The table shows substantial heterogeneity in portfolio composition across the wealth distribution. In particular, the share of financial wealth in real estate is sharply decreasing in wealth, as are the shares of vehicles, and cash and deposits. In contrast, the shares held in equities and private business wealth, particularly the

[8] Rachel and Smith (2017) show that the decline in the real rate occurred across a broad set of developed and emerging market countries. While many other factors certainly differ across countries, this shared trend suggests a global link between falling interest rates and rising financial wealth inequality.

Table 2: Duration of the Household Financial Wealth Portfolio

	Duration	Portfolio Shares			
		All	Bottom 90	P90-P99	Top 1
Assets					
Equities	49.78	0.21	0.15	0.24	0.23
Real Estate	12.28	0.47	0.79	0.40	0.23
Corporate PBW	55.93	0.09	0.01	0.07	0.19
Non-Corporate PBW	16.33	0.12	0.05	0.11	0.20
Vehicles	3.41	0.04	0.09	0.02	0.01
Fixed Income	5.28	0.15	0.15	0.16	0.14
Cash and Deposits	0.25	0.08	0.10	0.08	0.05
Liabilities					
Mortgage Debt	4.81	0.13	0.29	0.08	0.02
Vehicle Debt	1.45	0.01	0.03	0.00	0.00
Student Debt	4.50	0.00	0.01	0.00	0.00
Other Debt	1.00	0.01	0.02	0.01	0.01
Average (EW) Duration	**19.50**				
Aggregate (VW) Duration	**25.42**				
Bottom 90 (VW) Duration	**19.88**				
P90-P99 (VW) Duration	**25.61**				
Top 10 (VW) Duration	**28.01**				
Top 1 (VW) Duration	**30.74**				

Note: The column "Duration" reports the duration of each asset, averaged across all years from 1983 to 2019. The columns "Portfolio Shares" report the value-weighted (aggregate) portfolio weights from the Survey of Consumer Finances (SCF) and SCF+, averaged across all survey waves from 1983 through 2019. Liabilities receive negative portfolio weights. We report portfolio shares for the aggregate household sector (All), for households in the bottom 90 percentiles (Bottom 90), in the 90th to 99th percentile (P90-P99), and the top-1 percentile (Top 1). Portfolio shares for each group are computed as value-weighted averages, using the product of wealth and the SCF sampling weights as weights. The bottom panel reports all households' average or equally weighted (EW) duration and the aggregate or value-weighted (VW) duration. We also report the value-weighted duration for the Bottom 90, P90-P99, Top-10%, and Top-1%. Durations are first computed for each survey wave and then averaged across all survey waves.

corporate type (Corporate PBW), are strongly increasing in wealth.

Measuring Asset Durations. With each household's portfolio shares in hand, we now measure the durations of each asset class at each point in time. For most risky assets, we compute durations using a simple yet robust method: the Gordon Growth Model (GGM). Specifically, we assume that the asset's cash flows follow a growing perpetuity structure, so that cash flows or dividends, denoted Div_t, satisfy $Div_{t+1} = (1+g)Div_t$ for all t, and that these cash flows are discounted using a constant discount rate r. These assumptions generate a simple closed-form formula for duration

$$D_t^{GGM} = 1 + \frac{P_t}{Div_t}. \tag{7}$$

This allows us to compute an asset's duration at each point in time using only its contemporaneous price-dividend ratio (P_t/Div_t).[9]

With this general approach in mind, we now describe details of our duration measurement for each asset class, with full details available in Appendix D. The resulting durations, averaged over our 1980 - 2019 sample period, can be found in the first column of Table 2.

We measure the duration of equities using (7), obtaining the price-dividend ratio from the value-weighted CRSP stock market index. This yields an average duration of 50 years for equities.[10] We similarly measure the duration of real estate from (7), using rents as the asset's cash flows. For the price-rent ratio we use Zillow data from 2015 onwards. Prior to 2015, we compute the price-rent ratio by scaling its 2015 value using the proportional changes in the Federal Housing Financing Agency's price index and the CPI shelter index. This yields an average duration of 12 years for real estate.[11]

The most challenging duration measurement is that of private business wealth, since it is not publicly traded. Moreover, this category contains both fast-growing companies with high duration, as well as more stagnant businesses whose existence is tied to the human capital of the owner (Smith, Yagan, Zidar, and Zwick, 2019). To address this, we use detailed data in the SCF on the legal structure of the private business to divide private business into those with a corporate structure (Corporate PBW), and those with a non-corporate structure (Non-Corporate PBW), which proxy for the fast-growing and slow-growing types described above (see Section C.1.1 for details).

For non-corporate private businesses (Non-Corporate PBW), we again use (7), obtaining price-payout ratios from the Non-Corporate Business sector in the FAUS.[12] This yields an average du-

[9] Appendix A.2.4 shows that this approach delivers durations similar to those implied by our more sophisticated auxiliary asset pricing model.

[10] We consider a set of alternative measures for the duration of equity, detailed in Appendix D.2. The first one uses the S&P 500 cyclically-adjusted price-earnings ratio, combined with a earnings-dividend ratio of 0.5, yielding a duration of 45 years. The second one uses the price/dividend ratio computed from the JST macro-database (Jordà, Schularick, and Taylor, 2017), yielding a duration of 48 years. The third one uses the valuation for the entire non-financial corporate sector from the Financial Accounts of the United States (FAUS), yielding a lower duration of 34 years. The FAUS includes privately-held companies and imputes a valuation to private companies, which is adjusted downward by 25% to account for lower liquidity. Thus, we obtain similar duration numbers for all the measures except for the FAUS measure, which is not a pure equity duration measure.

[11] Appendix D.3 shows a higher housing duration of 21 years using data from Jordà, Schularick, and Taylor (2017), and an intermediate value of 16 using FAUS data and the VAR model.

[12] The corresponding cash-flow series includes both remuneration for labor and capital. To isolate the capital remuneration, we follow the approach used by PSID and split business income equally into labor and capital remuneration. This 50% labor income share is conservative compared to the literature. For example, Quadrini and Rıos-Rull (1997) and Krueger and Perri (2006) use a 86% labor share, or a 14% capital share, which would lead to a much higher valuation ratio and duration estimate of 57 years, as shown in Table D2 of Appendix D.4. A second alternative measurement based on the SCF, and detailed in Appendix D.6, also results in a higher duration of 70 years. Our baseline estimate is

ration of 16 years. For the corporate type of private business wealth (Corporate PBW), we adapt this approach. These businesses include fast-growing start-ups early in their corporate life cycle, which often have cash flows growing at a much faster rate than they will when the firm reaches its mature state. As a result, our assumption that cash flows grow at a constant rate, underlying the Gordon Growth Model, is likely inappropriate. To address this, we use a two-stage Gordon Growth Model that allows for an initial 20-year phase of higher cash flow growth, followed by a second absorbing stage with a lower growth rate. This two-stage model also has a closed-form solution for duration, presented in Appendix D.1.1.

To fit this model, we use data on publicly-listed stocks from CRSP. We assume that these private businesses begin similar to small public companies, then transition to behave like average public companies. Correspondingly, we obtain the initial price-dividend ratio from the smallest decile of public companies. We obtain cash flow growth rates for the initial high-growth stage and second lower-growth stage using 10-year real realized dividend growth for the smallest decile and aggregate (value-weighted) portfolios, respectively. This procedure yields an average duration for Corporate PBW of 56 years — slightly higher than our measure for public equities due to the initial high-growth phase. This measure is our most conservative (lowest duration) compared to the six alternative methods we consider, detailed in Appendix D.5.[13] We perform sensitivity analysis of our main results to alternative values for the duration of PBW in Section 5.3.

Computing the durations of our remaining asset classes is more straightforward, as they have simpler or fixed cash flow structures. For vehicles, we compute an asset duration based on a depreciation rate from the Fixed Asset Tables of the Bureau of Economic Analysis and the average age of vehicles in use from the Bureau of Transportation Statistics (see Appendix D.8 for details). For fixed income assets, we use the effective duration of the ICE government and corporate bond index. This series starts in 1996, prior to which we use the duration of the Treasury portfolio to impute the duration of Treasuries and corporate bonds (see Appendix D.9 for details). For cash and deposits, we assume a constant duration of 0.25 years.

We next compute the durations of household liabilities. While 30-year fixed-rate mortgages are dominant in the U.S., the average outstanding mortgage has duration far below 30 years due to the time elapsed since origination, amortization, coupon payments, and prepayment. Instead, we obtain durations for mortgage debt from a representative portfolio of all outstanding U.S.

the most conservative in that it results in the lowest Non-Corporate PBW duration.

[13] The first two alternatives use the same conceptual approach for Corporate PBW as the benchmark measure, but use an alternative to small stocks for the first, high-growth stage in the GGM. The first alternative uses high-growth private businesses that receive venture capital funding using data from Pitchbook, yielding a duration of 59 years. The second alternative uses stocks that just went through an initial public offering (IPO) using data from (see Ritter, 2022), yielding a duration of 56 years. The third alternative uses small stocks for the measurement of long-duration PBW, but uses a single-stage GGM, which generates a duration of 96 years. A fourth alternative measure, based on SCF data as detailed in Appendix D.6, generates a Corporate PBW duration of 64 years. All four alternatives result in a higher duration for Corporate PBW than our benchmark measure. Our approach is also robust to using micro data on the transitions between size deciles for publicly-listed firms, as explained in Appendix D.7.

pass-through mortgage-backed securities from the Bloomberg-Barclays Aggregate MBS Index (see Appendix D.10 for details). These durations are estimated by the data provider using a model to account for prepayment. For vehicle debt, we observe the periodic loan payment, the remaining number of periods, and the loan rate in the SCF, allowing us to directly compute the duration of each loan (see Appendix D.11 for details). For student debt, we assume a constant duration of 4.5 years.[14] For other debt, we assume a constant duration of 1 year as a compromise between its two main subtypes: revolving debt and amortizing 24-month personal loans.

Combining our time-varying durations with the portfolio shares for each household, equation (6) delivers a financial wealth duration for each household in each wave of the SCF. The bottom panel of Table 2 displays summary statistics from this distribution, averaged over all waves of the SCF. We compute an aggregate (value-weighted) duration of 25.42, which is substantially higher than the average (equal-weighted) duration of 19.50. Similarly, value-weighted average durations at the top of the wealth distribution (either the top-10% or top-1% groups) exceed those lower in the wealth distribution (the bottom-90% or middle 90%-99% groups). These statistics reflect a positive association between wealth and financial durations, as wealthier households hold larger shares of long-duration assets such as public equity and corporate private business wealth (Corporate PBW), and smaller shares of short-duration assets such as vehicles and cash.

Implied Links Between Interest Rates and Inequality. To interpret implications of these findings for links between interest rates and inequality, we can return to Proposition 2. Since value-weighted durations exceed equal-weighted durations, part (a) of this proposition implies that measured financial wealth inequality should increase when interest rates fall. To quantify these effects, we can turn to part (b) of the theorem. Using equation (5), we can approximate the impact on the top-10% wealth share of our measured fall in long-term real rates of 4.31pp between 1983 and 2019. For S^{α}, we use 67.6%, which is the average of the top-10% wealth share over our sample, while D^{top} and D^{bottom} are obtained as 28.01 and 19.88 from Table 2. Substituting these values into (5), we find that the top-10% financial wealth share should have increased by approximately $-0.676 \times (1.0 - 0.676) \times (28.01 - 19.88) \times (-0.0431) = 7.7$pp, very close to the actual increase of 8.3pp in the data. These results show the simple and robust intuition behind our finding that declining real rates explain most of the rise in financial wealth inequality over this period.

Despite the excellent fit of the data, several caveats are worth noting. First, equation (5) should be applied repeatedly each time rates change, using the S_t^{α}, D_t^{top}, and D_t^{bottom} prevailing at that time, with the resulting effects chained together, instead of all at once as in this example. This more correct computation allows the covariances of the terms to matter beyond their average values. For example, if rates fall more when the gap in durations between the top and bottom

[14] Student loans are typically 10-year annuities. At the average rate on outstanding student loans in 2017 of 5.8%, the duration is 4.56. At the higher interest rates prevailing in the 1980s the duration would be slightly smaller.

Figure 3: Financial Duration by Net Wealth Percentiles and by Age

(a) By Net Wealth Bin (b) By Age

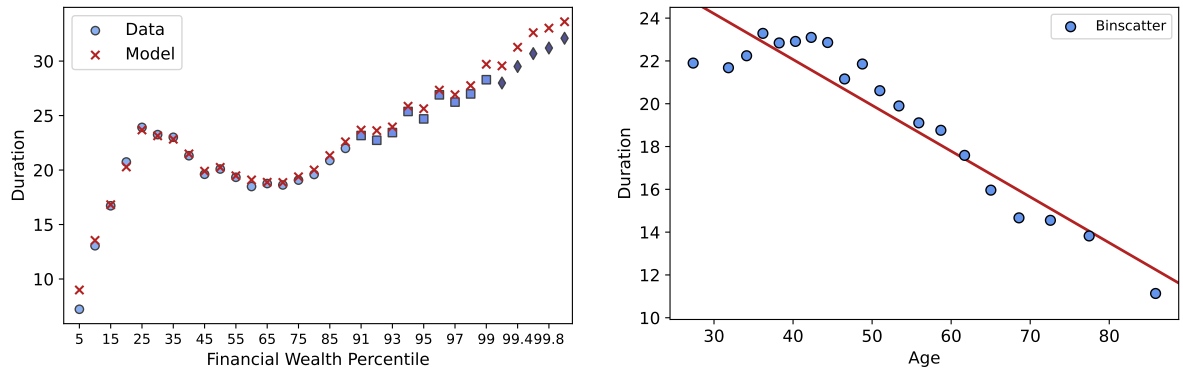

Note: This left panel displays average duration by financial wealth bin in the model and data (source: SCF). The x-axis is measured in percentiles, which each tick representing the right edge of the bin, so that e.g., "5" corresponds to households with financial wealth percentile in the interval $[0, 5]$. Red crosses display model equivalents (see Section 4). This right panel displays a binscatter of average duration by age in the data, after controlling for the financial wealth bins displayed in Figure 3a, while the red line represents the least squares fit (source: SCF), using a regression that pools over households in all SCF waves of our sample (1983–2019) with each wave weighted equally.

of the wealth distribution is particularly high, then this comovement will amplify the impact on inequality. Moreover, equation (5) measures the response to an instantaneous change, while the actual change in inequality we measure in the data occurred over decades. To the extent that households' consumption and savings behavior is influenced by the wealth changes, this formula will fail to capture these impacts on the dynamic evolution of the wealth distribution. We address both of these issues in a more comprehensive quantitative experiment in Section 5.

Cross-Sectional Patterns of Duration. To close our empirical section, we analyze how our measures of household-level duration vary with characteristics in the cross-section.

To begin, Panel (a) of Figure 3 plots the average duration by wealth bin in the SCF (blue dots, squares, and diamonds). Red crosses display equivalents in our structural model, discussed in Section 4. Since higher-wealth households are more influential for aggregate wealth outcomes, the figure displays 5% wealth bins up to the 90th percentile, then 1% bins up to the 99th percentile, and 0.2% bins for the top 1%. The figure shows that duration is generally increasing in wealth. The sole exception is due to a hump shape around the 25th percentile of wealth. This pattern is driven by home ownership, since the most levered homeowners have relatively little wealth but very high portfolio durations. Beyond this segment, however, we find that duration is consistently increasing in wealth up to the very top of the wealth distribution.

The second key characteristic that drives variation in financial duration is age. Panel (b) of Figure 3 displays a binscatter of measured duration in our SCF data by age, after controlling for net wealth using dummies for each of the bins constructed in Panel (a). This figure shows a strong and highly linear negative relationship between age and duration.

Combining these patterns, the duration of household financial wealth is well approximated by

$$D_i^\theta = \alpha + \beta Age_i + \sum_j \gamma_j NetWealthBin_{i,j} + e_i \tag{8}$$

where $NetWealthBin_{i,j}$ is an indicator for whether household i falls in financial wealth bin j. We show in Appendix D.12 that adding other covariates yields little additional power to explain variation in financial duration across households. We estimate this regression, and use the fitted values to impute household durations in our calibrated model in Section 4.

To summarize, our main empirical finding is that financial wealth duration is increasing in financial wealth. Combined with our theoretical results in Section 1, this implies that revaluation effects from a decline in interest rates since the 1980s should have increased financial wealth inequality, with a quantitative increase in top wealth shares on the order observed in the data.

3 Incomplete Markets Model with Household Heterogeneity

To develop theoretical and quantitative insights on how changes in interest rates affect the distribution of wealth, we develop a simple life-cycle model with idiosyncratic labor income risk that connects interest rates, duration, and wealth inequality in a transparent fashion. In order to straightforwardly apply the exact path of rates that occurred in the data, we use a partial equilibrium model (alternatively, a small open economy) where interest rates are taken as given, and abstract from the structural mechanisms or shocks that caused the interest rate to fall. Instead, we analyze the relationship between wealth, interest rates, and consumption as an accounting identity. Regardless of the underlying cause, the duration measures we use accurately describe the change in financial wealth and consumption possibilities due to the observed decline in interest rates relative to a counterfactual world where interest rates had not fallen but all other variables had evolved identically. Our work thus provides a robust quantitative measure that complements work on the underlying forces driving changes in interest rates over this period.

Appendix G generalizes our environment to a general equilibrium setting in which interest rates are determined endogenously, and the aggregate endowment grows at a stochastic rate. We show that changes in the equilibrium interest rates reflect changes in any or all of (i) the subjective time discount rate, (ii) the mean rate of growth, or (iii) the variance of that growth rate. Following Krueger and Lustig (2010), we show how to map the stochastically growing economy into a stationary economy in the style of Bewley (1986) without growth and aggregate risk using a change of measure. Thus, while the exact cause of the decline in interest rates since the 1980s is an important research question, we believe our findings should be robust to the ultimate answer.

3.1 Model Structure

Demographics. The economy is populated by a continuum of households. Households transition through a life cycle, where age j varies from 0 to J. Households survive from age j to age $j+1$ with probability ϕ_j, with $\phi_J = 0$. When a household dies, it is replaced by a newborn household ($j = 0$), which inherits its remaining assets as a bequest.

Endowments. Each household i of age j receives exogenous labor income given by $y_j(z)$, where z is a household-specific (idiosyncratic) stochastic process.

Asset Technology. Households trade a complete set of bonds offering fixed cash flows at future dates.[15] Without loss of generality, we restrict attention to zero coupon bonds, where a zero coupon bond with maturity m promises one unit of the numeraire in m periods. We denote holdings of each bond as x_m, and its price as q_m. Markets are incomplete in that households cannot contract on their idiosyncratic income realizations.

We consider an economy in steady state, so that prices $\{q_m\}$ are expected to hold in all future periods. We further assume that the one-period bond is traded on a global market in which our model economy is a price taker, so that its interest rate takes the exogenous value R. If we normalize $q_0 = 1$, the absence of arbitrage opportunities requires $q_m = R^{-m}, \forall m$ in steady state.

Household Problem. Given start-of-period bond holdings x, labor income y, and bond prices q, a household of age j chooses consumption c and bond holdings x' to solve the recursive problem

$$V_j(x; z) = \max_{c, x'} \underbrace{\frac{c^{1-\gamma}}{1-\gamma}}_{\text{flow utility}} + \underbrace{\phi_j \beta \mathbb{E}\left[V_{j+1}(x'; z') \,\middle|\, z \right]}_{\text{continuation value}} + \underbrace{(1-\phi_j)\chi \frac{\left(\sum_{m=1}^{M} q_{m-1} x'_m \right)^{1-\gamma}}{1-\gamma}}_{\text{bequest utility}} \tag{9}$$

subject to the budget constraint,

$$c \leq y_j(z) - \underbrace{\sum_{m=1}^{M} \left(q_m x'_m - q_{m-1} x_m \right)}_{\text{net saving}}$$

and the borrowing constraint $\sum_m q_m x'_m \geq 0$, which rules out negative bequests. While past work has shown empirical benefits from using non-homothetic bequest functions (De Nardi, 2004; Straub, 2019), our use of a homothetic bequest function offers large gains in tractibility for our theoretical and quantitative analysis. We note that because these non-homothetic bequest functions

[15] The model with aggregate shocks in Appendix G allows for asset payoffs that depend on the aggregate state.

imply greater savings by the rich following wealth gains, and hence greater amplification and persistence of wealth inequality, our results can be seen as, if anything, conservative.

Household Optimality. Each of the m optimality conditions for bond holdings x'_m collapses to the same Euler equation modified to include bequest utility:

$$c^{-\gamma} = R\left\{\phi_j \beta \mathbb{E}\left[(c')^{-\gamma} \,\middle|\, z\right] + (1-\phi_j)\left(\sum_{m=1}^{M} q_{m-1} x'_m\right)^{-\gamma}\right\}. \tag{10}$$

These optimality conditions do not uniquely identify the portfolio holdings, since households expect to receive the same holding period return R on all bond maturities. Instead, only the household's total financial wealth $\theta \equiv \sum_{m=1}^{M} q_{m-1} x_m$ and total savings $s \equiv \sum_{m=1}^{M} q_m x'_m$ matter for the household's problem in a steady state where interest rates do not change. Given this indifference, we assign each household a unique portfolio $\{\hat{x}_m\}$ that matches its empirically predicted duration given its and position in the wealth distribution from (8).

Financial Wealth. Using the results above, we can simplify the household's problem using total financial wealth θ as a single state variable. Since next period financial wealth θ' is given by

$$\theta' = \sum_{m=1}^{M} q_{m-1} x'_m = R^{-1} \sum_{m=1}^{M} q_m x'_m$$

we can rewrite (9) as the optimization problem

$$V_j(\theta; z) = \max_{c,\theta'} \frac{c^{1-\gamma}}{1-\gamma} + \phi_j \beta \mathbb{E}\left[V_{j+1}(\theta'; z') \,\middle|\, z\right] + (1-\phi_j)\chi \frac{(\theta')^{1-\gamma}}{1-\gamma}$$

subject to the budget constraint

$$c \leq y_j(z) + \theta - R^{-1}\theta'. \tag{11}$$

and the no borrowing condition $\theta' \geq 0$, yielding the single optimality condition

$$c^{-\gamma} = R\left\{\phi_j \beta \mathbb{E}\left[(c')^{-\gamma} \,\middle|\, z\right] + (1-\phi_j)(\theta')^{-\gamma}\right\}. \tag{12}$$

Although households believe the economy will remain in steady state forever, we apply unanticipated shocks to interest rates that will revalue financial assets. In this case, although households believed (11) would hold, the actual next period financial wealth is updated according to

$$\tilde{\theta}' = \sum_{m=1}^{M} \tilde{q}_{m-1} x'_m$$

where $\{\tilde{q}_m\}$ is the updated set of bond prices conditional on the new realized interest rate.

3.2 Consumption Effects of Interest Rate Changes.

While the previous exercises clarify the impact of interest rates on financial wealth inequality, it is by no means obvious whether these changes in measured financial wealth inequality reflect changes in consumption possibilities (welfare), or simply represent revaluations of the same consumption plans (paper gains and losses). To distinguish the two, we iterate forward (11) to obtain

$$\theta_0 = \mathbb{E}_0 \left\{ \sum_{t=0}^{T-1} R^{-t}(c_t - y_t) + R^{-T}\theta_T \right\} \tag{13}$$

where T is the (stochastic) death date, measured in number of periods from the present ($t = 0$), and θ_T is the final bequest of this household, which depends on the histories of c_t and y_t, as well as the realization of T. Financial wealth is thus equal to the present value of future excess consumption, defined as consumption minus income $c_t - y_t$, plus the present value of the bequest θ_T.

For this section, we consider the infinite-horizon limit with vanishing mortality risk, so that

$$\theta_0 = \mathbb{E}_0 \left\{ \sum_{t=0}^{\infty} R^{-t}(c_t - y_t) \right\}. \tag{14}$$

With this identity in hand, we turn to our main theoretical insight regarding the link between interest rates and consumption possibilities.

Proposition 3. Assume that the interest rate R and discount factor β both unexpectedly and permanently change to $\tilde{R} = R\exp(\varepsilon)$ and $\tilde{\beta} = \beta\exp(-\varepsilon)$. Define the duration of financial wealth by D^θ, and the duration of future excess consumption D^{c-y} by

$$D^\theta \equiv \frac{\sum_{m=1}^{M} m \times q_m x_m}{\sum_{m=1}^{M} q_m x_m}, \qquad D^{c-y} \equiv \frac{\mathbb{E}_0\left\{\sum_{t=0}^{\infty} t \times R^{-t}(c_t - y_t)\right\}}{\mathbb{E}_0\left\{\sum_{t=0}^{\infty} R^{-t}(c_t - y_t)\right\}}.$$

Let $\tilde{c}_t$ denote the household's consumption plan following the change in rates. Then for a decline in rates ($\varepsilon < 0$):

(a) The set of consumption allocations the household can afford expands if $D^\theta > D^{c-y}$, contracts if $D^\theta < D^{c-y}$, and is unchanged if $D^\theta = D^{c-y}$.

(b) If, following the shock, a household has exactly enough financial wealth $\tilde{\theta}_0$ to afford its pre-shock consumption plan $\{c_t\}$, then this consumption plan remains optimal.

(c) Household consumption is unchanged following a shock to interest rates ($c_t = \tilde{c}_t$, $\forall t$) if and only if $D^\theta = D^{c-y}$.

The proof can be found in Section F.3. Part (a), building on Auclert (2019), shows that whether changes in interest rates expand or contract a household's consumption possibilities depends not on the absolute level of its financial wealth duration, but on how it compares to the duration of that household's lifetime excess consumption.[16] While financial wealth is always equal to present value of future excess consumption by the budget identity (14), the two may be differentially exposed to the same interest rate shock, much like a bank with a maturity mismatch of assets and liabilities. As a result, even if a household gains financial wealth from a decline in rates, it can still see its consumption possibilities contract if the present value of its pre-shock excess consumption plan rises by more than its financial wealth.

Part (b) shows that, as long as a household can exactly afford its pre-shock consumption plan, it will still find this plan optimal, and choose it at equilibrium. This follows directly from (12) in the absence of bequests, and the fact that $\beta R = \tilde{\beta}\tilde{R}$. We will exploit this result to construct a counterfactual wealth distribution that leaves consumption approximately unchanged following a shock in Section 6. Part (c) follows immediately, since the pre-shock consumption plan is optimal if and only if it is exactly affordable after the shock, which occurs if and only if $D^{\theta} = D^{c-y}$.

Proposition 3 assumes zero mortality risk. Although part (a) would hold absent this assumption, parts (b) and (c) would not. This is because we follow the convention of measuring a household's bequest in current value (i.e., dollars). Thus, even if a household's consumption plan is fixed, falling rates that increase the present value of future cash flows will mechanically increase the value of financial wealth the household maintains along its consumption path, thus increasing the "size" of bequests. This issue highlights a potential downside of typical bequest parameterizations, since households receive more utility from a more "valuable" bequest when rates fall, even though the heirs receiving it will not be able to afford larger consumption allocations. Regardless, although parts (b) and (c) hold only approximately, the quantitative discrepancy should be small, particularly for young households whose bequests are far in the future on average.

Proposition 10 in the Appendix generalizes Proposition 3 to a world with stochastic growth and endogenous interest rates. Our proposition goes through irrespective of the combination of discount factor, growth rate, and growth uncertainty changes driving the movement in rates.

[16] Auclert (2019) shows that for perfectly transitory change in the real rate, the consumption response depends on unhedged interest rate exposure (URE). Mapped into the language of this section, this result states that the response depends on whether the present value of future excess consumption or future net cash flows from the household's financial portfolio is greater. This is a sufficient statistic because a perfectly transitory change in the short-term discount rate changes the cumulative discounting of all future cash flows by a constant proportional amount. In our setting with permanent shocks, the exact timing of the cash flows matters, as cash flows further in the future are more affected by a permanent change in rates than cash flows closer to the present. As a result, URE is no longer a sufficient statistic. For example, a household with a portfolio of five-year bonds has the same URE as an otherwise identical household holding two-year bonds, but will experience much larger gains following a decline in rates. Our sufficient statistic of duration accounts for these timing effects, allowing us to extend these results to a setting with permanent shocks.

3.3 Implications: Perfectly Hedged Economy.

Proposition 3 implies that households whose financial wealth and excess consumption durations are perfectly aligned ($D^{\theta} = D^{c-y}$) do not change their consumption plans following a decline in rates. In such a "perfectly hedged" economy, all changes in financial wealth inequality would therefore reflect only "paper" gains, while keeping consumption inequality unchanged. Combined with Proposition 2, this result implies that, in a perfectly-hedged economy, we should see financial wealth inequality rise following a decline in rates if and only if the duration of excess consumption D^{c-y} is higher for wealthy households than poor households. While this condition depends on the parametrization of the model, and there are two offsetting forces at work, we observe here that it is highly dependent on the model's life cycle.

We first consider Bewley models without a life cycle, where households are ex-ante identical. In these models, wealthy households are those who experienced favorable income shocks. These households tend to have high labor income today, which is expected to mean revert over time. Due to consumption smoothing, excess consumption for these households ($c_t - y_t$) tends to be negative in the short run, and positive in the long run, generating a longer duration of excess consumption. Since the wealthy therefore tend to have a higher value for D^{c-y} than the poor, perfect hedging in these economies features a positive covariance between financial wealth and the duration of financial wealth. Such an economy would therefore see financial wealth inequality rise following a decline in interest rates, making the rise in measured financial wealth inequality documented in Section 2 potentially consistent with perfect hedging. In this case, wealthier households would need larger gains in wealth following a fall in rates to finance their future consumption.

This result changes dramatically in the presence of a life cycle. In life cycle economies, the young are typically the least wealthy. At the same time, life-cycle savings motives imply that most households have negative excess consumption in middle age, as they save for retirement, followed by positive excess consumption in retirement. Thus, D^{c-y} is highest for young households, and decreasing in age. This force pushes down the cross-sectional covariance of financial wealth and the duration of excess consumption, potentially making it negative. In a perfectly-hedged life-cycle economy, we may see financial wealth inequality fall following a decline in interest rates, which would make the large rise in financial wealth inequality since the 1980s inconsistent with perfect hedging. Such a discrepancy would imply that there are real consumption effects from interest rate changes, and that these effects are unequally distributed across households.

The two opposing forces highlighted above — the history of stochastic income realizations and the life-cycle effect of saving for retirement — will appear in virtually any model of household inequality. The question of which will dominate, and the ultimate implications for wealth inequality under perfect hedging, is a quantitative one. In the following sections, we answer it (and others) using a standard calibration of a workhorse life-cycle incomplete markets model.

4 Model Quantification

To measure the effect of changing interest rates on wealth inequality, our model must produce quantitatively realistic responses. In this section, we parameterize our model, calibrate it to match data on both household financial wealth durations, income risk, and wealth shares, and present the properties of the model's stationary economy.

4.1 Calibration

Preferences and Mortality Risk. We calibrate the model's mortality risk via the survival probabilities ϕ_j to match Social Security Actuarial tables. Since we model households, we take the average of the male and female mortality rates, weighted by the proportion alive at each age.

We set γ, the risk aversion of the households and the curvature of the bequest function, to a standard value of 2. We set the time discount factor to $\beta = 0.949$ to target a ratio of net wealth to disposable labor income in 1983 of 6.79.[17] We set the bequest utility parameter $\chi = 7.894$ to match a ratio of bequests to GDP of 8.8%, following Auclert, Malmberg, Martenet, and Rognlie (2021).

Interest Rates We set the steady-state interest rate R to match the 10-year real rate in 1983. The interest rate in the data comes from an economy with growth. As shown in Appendix G.3.2, the mapping between the interest rate in our stationary economy and that in the growing economy is:

$$R = \frac{R_g}{G}, \qquad\qquad \beta = G \times \beta_g \tag{15}$$

where R_g and β_g are the gross interest rate and time discount factor in the growing economy, R and β are the corresponding values in our stationary economy, and G is the gross growth rate (adjusted for a Jensen effect and a risk premium). Given average log growth of 1.91%, the observed rate in 1983 of $R_g = 4.94\%$ implies $R = 3.04\%$.

Financial Duration As households in the model go through their life cycle and experience changes in financial wealth, their financial duration is updated according to

$$\widehat{D}_i^{\theta} = \hat{\alpha} + \hat{\beta} Age_i + \sum_j \hat{\gamma}_j NetWealthBin_{i,j} \tag{16}$$

equal to the fitted value from regression (8). The model's equal-weighted and value-weighted durations are 19.87 and 25.96, respectively, close to their empirical counterparts in Table 2.[18] The

[17]Net wealth is measured as the net worth of households and nonprofits in the FAUS (Table B.101). Disposable labor income is household income net of personal taxes, rental, interest, and dividend income, and one half of proprietor's income in the National Income and Product Accounts.

[18]Since the model feeds in the (time-varying) portfolio shares in the SCF data, it matches observed portfolio choices by construction. Endogenizing portfolio choice is an interesting extension, but cannot improve the model's implications

model's durations by wealth bin can be seen as the red crosses in Panel (a) of Figure 3, showing that the model achieves a close fit of the data.[19]

Regular Income Parameters. The labor income process consists of a regular component and a superstar component. The regular income process for household i of age a at time t that is not currently in the superstar state takes the form, standard in the literature, given by:

$$\log\left(y_{t,a}^i\right) = m_t + \chi' X_t^i + z_t^i, \tag{17}$$

$$z_{t+1}^i = \alpha_i + \eta_{t+1}^i + \nu_{t+1}^i, \tag{18}$$

$$\eta_{t+1}^i = \rho \eta_t^i + u_{t+1}^i, \tag{19}$$

where m_t is a year-fixed effect and X_t^i is a vector of household characteristics that includes a cubic function of age.[20] We normalize the mean of the age profile to unity during working life.

The stochastic income component z_t^i contains a household-fixed effect α^i, a persistent component η_{t+1}^i, and an i.i.d. component ν_{t+1}^i. We set $\mathbb{E}[\nu^i] = E[u^i] = \mathbb{E}[\alpha^i] = \mathbb{E}[\eta_0^i] = 0$, while $Var[\nu^i] = \sigma_\nu^2$, $Var[u^i] = \sigma_u^2$, $Var[\alpha^i] = \sigma_\alpha^2$, and $Var[\eta_0^i] = \sigma_{\eta,0}^2$. To allow for lower income risk during retirement, we re-estimate (17) – (19) separately for households above and below age 65, and assume that model households face income risk that switches when they turn 65. The parameters are estimated by GMM using PSID data from 1970 until 2017, as detailed in Appendix E. Figure E1 plots the deterministic life-cycle income profile.

The literature typically estimates (17) – (19) on labor income for males between ages 25 and 55. We deviate from this practice by: (i) considering a broader income concept, (ii) modeling the entire life-cycle from age 18 to 80, and (iii) focusing on households rather than individuals.

First, from the model's perspective, the relevant notion of income used for measuring excess consumption is broad and includes transfers. This approach extends that of Catherine, Miller, and Sarin (2020), who incorporate Social Security benefits as a component of income, to include all sources of labor income and transfers available in our data. To that end, we measure income in the data as income from wages and salaries, the labor income component of proprietor's income, government transfers (unemployment benefits, Social Security, other government transfers), and private defined-benefit pension income. To obtain consistent data series, we reconcile differences in variable definitions among the various waves of the PSID (a non-trivial task, see Appendix E).

Second, we are interested in the entire life-cycle. To this end, we estimate the income distribu-

for portfolio choice, financial duration, and wealth inequality.

[19]Since we are using fitted values, the main reason why the model durations by wealth bin would deviate from their data equivalents is if the joint distribution of age and net wealth bin differs between model and data. This close fit rules out an important discrepancy of this kind.

[20]Our results are similar if we estimate the year fixed effect and the age profile separately for groups of households that depend on education (college completion or not), race (white or non-white), and gender of head of household (8 groups total). Since it makes little difference for our main results, we assume ex-ante identical households for simplicity.

tion for a wide range of ages from 18 to 80.[21] Because our income concept includes transfers such as unemployment benefits and retirement income from public or private defined-benefit pension plans, we do not need to assume that households are working full-time. Instead, our estimation takes into account the full cross-section of sample households, including retirees, receipients of unemployment benefits, part-time employees, and so on.

Third, we focus on households, aggregating income across its adult members. This avoids the obligation to model demographic changes such as becoming married, divorced, or widowed. Instead, we simply follow households identified by the head of household in the data.

Superstar Income Parameters. To help the model match the level of wealth inequality in the high-interest rate regime (1980s), we follow Castaneda, Diaz-Gimenez, and Rios-Rull (2003) in enriching the income process in (17) – (19) with a superstar income state, which has a high income level $\eta_t^i = \eta^{sup}$. Households enter this state with probability p_{12}^{sup} when they are in the regular income state, and return to a regular income state with probability p_{21}^{sup} when they are in the superstar income state. The transition probability parameters $p_{12}^{sup} = 0.0002$ and $p_{21}^{sup} = 0.975$ are taken from Boar and Midrigan (2020), and imply a roughly 1% probability of entering in the superstar income state over one's lifetime. Conditional on entering, the state has an expected duration of 40 years. The income level η^{sup} is then chosen to match the top-10% wealth share in 1983 exactly, which requires a value equal to 32.46 times average income.

Income Process Discretization. For the persistent component of the income process η, we use a nested structure, so that a household first draws whether it transitions into or out of the superstar state using the probabilities stated above. Conditional on ending out of the superstar state, the household then draws a new value of η from the non-superstar distribution. For households beginning and ending in the non-superstar state, we approximate (18) with a discrete Markov chain P_η using the method of Rouwenhorst (1995). Households transitioning from the superstar to non-superstar states draw a random non-superstar η state from the ergodic distribution of P_η. Conditional on η, we draw i.i.d. values for ν using nodes and weights from Gaussian quadrature.

4.2 Stationary Economy: High Interest-Rate Regime

With the model calibrated, we now explore the quantitative properties of its stationary distribution under the initial high interest rate regime. Figure 4 displays the life cycle profiles of several key variables. The first two columns show income, consumption, financial wealth, and human wealth, with the axes are normalized such that 1 represents the median income during working

[21] Since model households can survive to age 100, we assume that the age profile embedded in $\chi' X_t^i$ remains constant from age 80 onward.

Figure 4: Life Cycle Profiles

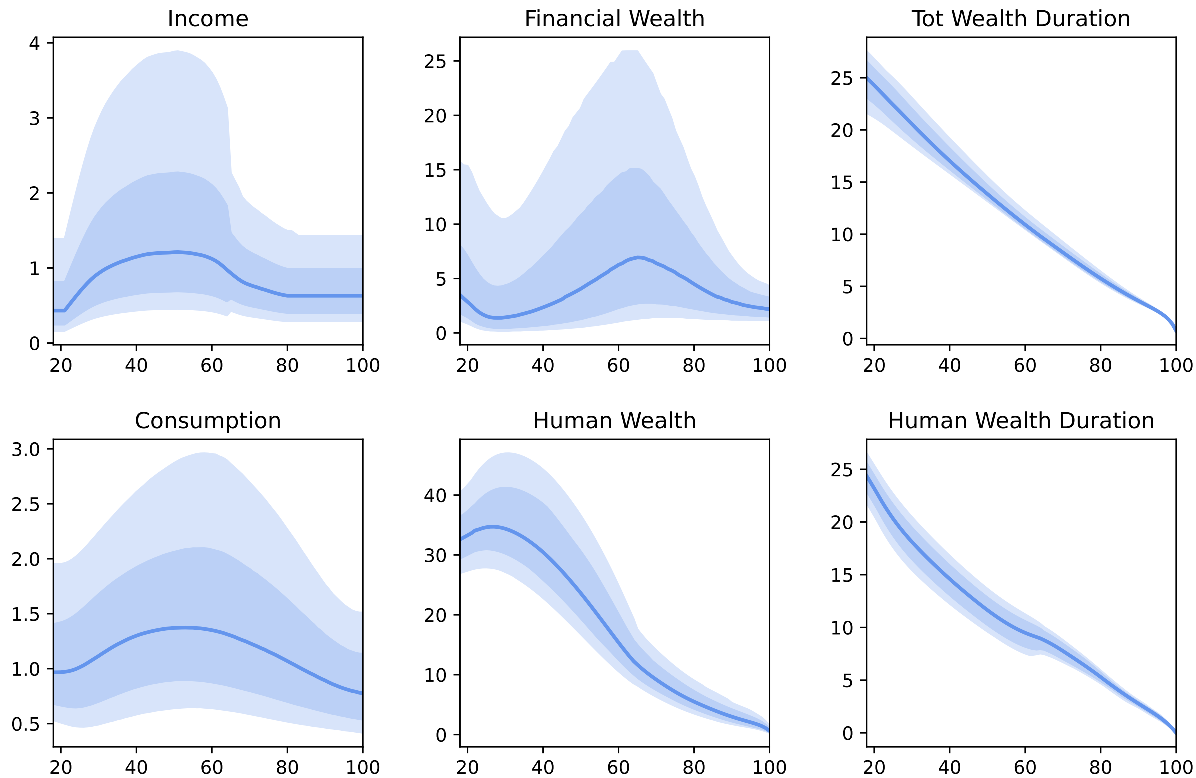

Note: This figure plots the life cycle profiles by age for the all agents of all groups combined. The axes are normalized so that the average income across all agents of all ages is equal to unity. The center line displays the median, while the dark and light bands represent 66.7% and 95% percentile bands.

life. Income displays the traditional hump-shape over the life cycle. Income inequality is increasing over the first half of the life cycle as income shocks accumulate. After retirement, households switch to our estimated over-65 income process which has lower dispersion. While this compresses income inequality beyond this point, we note that it remains non-negligible since agents have heterogeneous retirement income and still face some income risk after age 65.

The bottom right panel shows that both the level and dispersion of consumption are rising over the working part of the life cycle, with dispersion falling in retirement when income risk declines. This is consistent with the data, which show that consumption inherits the hump-shaped profile from income (e.g., Krueger and Perri, 2006).

The top middle panel shows financial wealth, which increases in preparation for retirement, and is subsequently run down during retirement. Financial wealth inequality rises and falls over the life cycle. The first few years are influenced by bequests, which households receive at age 18. On average, households spend these down before beginning to save for retirement closer to age 30, while the large dispersion in the size of bequests generates the initial peak in financial wealth dispersion at the start of the life cycle.[22]

The bottom middle panel shows human wealth, defined as the present value of income. Human wealth is generally declining in age as there are fewer remaining periods of income remaining. However, human wealth rises in the early years of the life cycle as the households' highest-earning periods are brought closer to the present. Total wealth for young households consists almost exclusively of human wealth, except for a small share of households who received large bequests. As households age and prepare for retirement, a larger share of total wealth becomes financial wealth. However, human wealth remains the largest component of total wealth for most households until the typical retirement age.

Appendix Figure H1 displays the Lorenz curves for consumption and wealth for all households (in all groups), and reports the Gini coefficients. The model generates a Gini coefficient for (after-transfer) household income of 0.489. Consistent with the data, consumption inequality is somewhat lower than income inequality, and has a Gini coefficient of 0.406. Financial wealth is much more unequally distributed than human wealth or total wealth, with a Gini coefficient of 0.710 compared to 0.405 and 0.417, respectively. This much lower inequality in total wealth arises from a combination of (i) the importance of human wealth in total wealth, and (ii) the negative cross-sectional correlation between financial wealth, dominated by the middle aged and old, and human wealth, which is highest for the young.

The right panels of Figure 4 display the duration of human and total wealth by age. Human wealth represents a claim on lifetime income whereas total wealth represents a claim on lifetime consumption. Both of these durations are similar because of the importance of human wealth in total wealth. These durations are high when young, around 25, and drop rapidly as age increases, as there are fewer years of income remaining.

5 Results: Repricing Under Falling Interest Rates

In this section, we apply unexpected and permanent interest rate shocks to our economy and study the quantitative effect on financial wealth inequality. We do this in two steps. First, we consider a *one-shot* experiment, in which we apply a single unanticipated and permanent decline in the real interest rate, which captures the entire decline in interest rates between 1983 and 2019. This simple one-shot experiment allows for clear exposition of the mechanism behind our results but lacks realism, since the actual decline in rates occurred gradually over several decades. To generate more quantitatively realistic results, the main analysis considers a *gradual transition* experiment in which we feed in the sequence of annual interest rate changes.

[22] To the extent that actual households receive bequests later in life, this would increase financial inequality between young and old, strengthening our effects.

Figure 5: Histograms, Repriced Financial Wealth Distribution

(a) Hist: Original vs. Repriced

(b) Binscatter: Original vs. Repriced

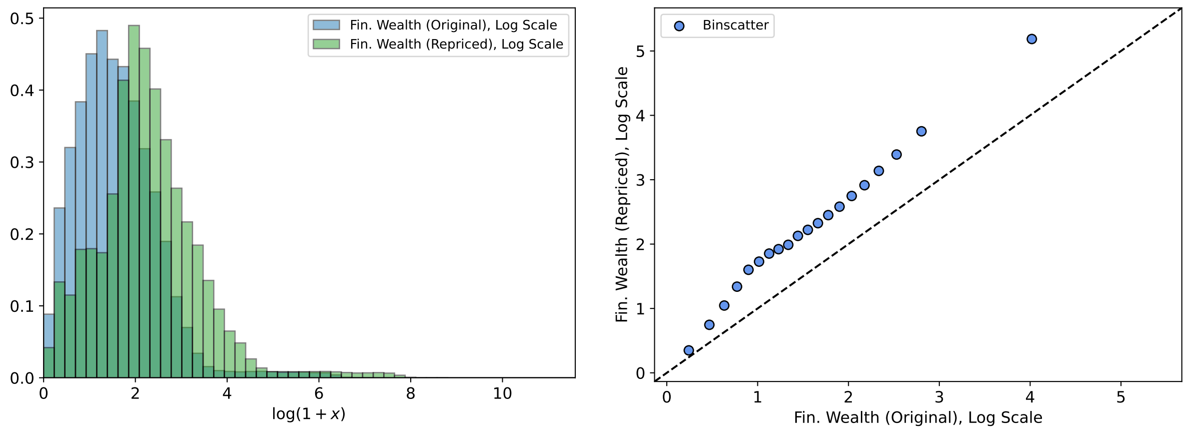

Note: The left panel plots the original wealth distribution in the steady state with high interest rates in blue and the repriced wealth distribution after the decline in interest rates in green. The right panel is a binscatter plot, where each dot represents 5% of the population, that maps the financial wealth in the high interest rate steady state, reported on the x-axis, to the repriced financial wealth after the rate change on the y-axis.

5.1 Repricing: One-Shot Experiment

To build intuition, we begin with the one-shot experiment. The interest rate in the growing economy declines from 4.94% to 0.63%, which corresponds to a decline from 3.04% to -1.19% in our model's stationary economy. Following this change, we update the value of financial wealth for each household in the economy using

$$\tilde{\theta}_{i,t} = \theta_{i,t} \exp(-D^{\theta}_{i,t} \times \Delta \log R). \tag{20}$$

where financial wealth duration for household i, $D^{\theta}_{i,t}$ is obtained as the the fitted value of (8), averaged over all SCF waves, applied to each household in the model. We refer to the resulting wealth distribution $\tilde{\theta}_{i,t}$ as the *repriced wealth distribution*.

Figure 5 Panel (a) shows the repriced distribution in green, alongside the initial wealth distribution in blue. The decline in rates pushes the distribution to the right, as financial wealth increases in value under falling discount rates. Overall, this decline in rates increases the value of financial wealth by 208.4%. Panel (b) presents a binscatter showing how these gains are allocated across the wealth distribution, showing that all but the poorest agents see large asset valuation gains, with larger gains for the wealthier households due to their higher financial wealth durations. Since the model reproduces the fact that the value-weighted average financial duration exceeds the equally-weighted average, falling rates increase financial wealth inequality in our model, consistent with our theoretical results in Section 1.

5.2 Repricing: Transition Experiment

We now turn to our gradual transition experiment, which provides our main quantitative results. For each year between 1983 and 2019, we obtain the real 10-year rate $R_{g,t}$ implied by our auxiliary asset pricing model in Appendix A. We detrend these rates with equation (15) using a constant gross growth rate $G = 1.0193$ to obtain the time series of R_t to use in our stationary economy. We assume that each of these annual changes is permanent and unexpected. We note that, if the real rate declines had instead been anticipated by investors in 1983, the entire effect on valuations would have been front-loaded, and the term structure would have been steeply downward sloping. Both of these predictions are counterfactual.

At the start of each period, we revalue financial wealth with (20) using that year's change in rates $\Delta \log R_t$. To compute household portfolio durations for each year, we re-estimate equation (8) for each SCF wave (typically every three years), then interpolate the coefficients for the years between SCF waves.[23] Applying these coefficients to (16), we assign each model household a financial duration based on its age and position in the wealth distribution.

Importantly, the economy does not jump to the new steady state following each interest rate shock. Instead, following each year's change in rate, financial wealth is revalued, households update their optimal consumption-savings plans, and the economy begins a long transition to the new steady state. However, the following year this transition is interrupted by a new unexpected interest rate shock, and the process repeats. The financial wealth distributions we compute in each year thus reflect the entire history of both rates and household consumption-savings decisions.

Effects on Financial Wealth Inequality. The results of our gradual transition experiment are displayed in Panel (a) of Figure 6. The black line displays the actual top-10% share in the data (WID), while the red line displays its implied path in our model's transition experiment. The figure shows that our model-implied series provides a close fit for the data, ultimately explaining 95% of the increase in the top-10% share observed over the sample. Thus, we find that heterogeneity in financial wealth durations, with wealthier households more exposed to interest rate changes, is sufficient to generate effectively all of the rise in inequality since the 1980s.

This overall rise in inequality is the combination of two forces: (i) revaluations of financial wealth according to measured household financial durations, and (ii) households' optimal response to these revaluations in their consumption-savings plans. We now decompose their separate contributions. Define S_t^{10} to be to the top-10% share of financial wealth at the start of time t, prior to any revaluation. Once the new interest rate R_t is realized, each household's pre-revaluation wealth $\theta_{i,t}$ is updated to $\tilde{\theta}_{i,t}$, leading the top-10% share of wealth to be updated from S_t^{10} to $\tilde{S}_t^{10}$. We define the *instantaneous revaluation effect* in each period, corresponding to force (i)

[23]Since the first wave of the SCF is 1989, we use the fitted values from 1989 for all years between 1983 and 1989.

Figure 6: Top-10% Share, Gradual Transition Exercise

(a) Top-10% Share: Repriced (b) Top-10% Share: Compensated

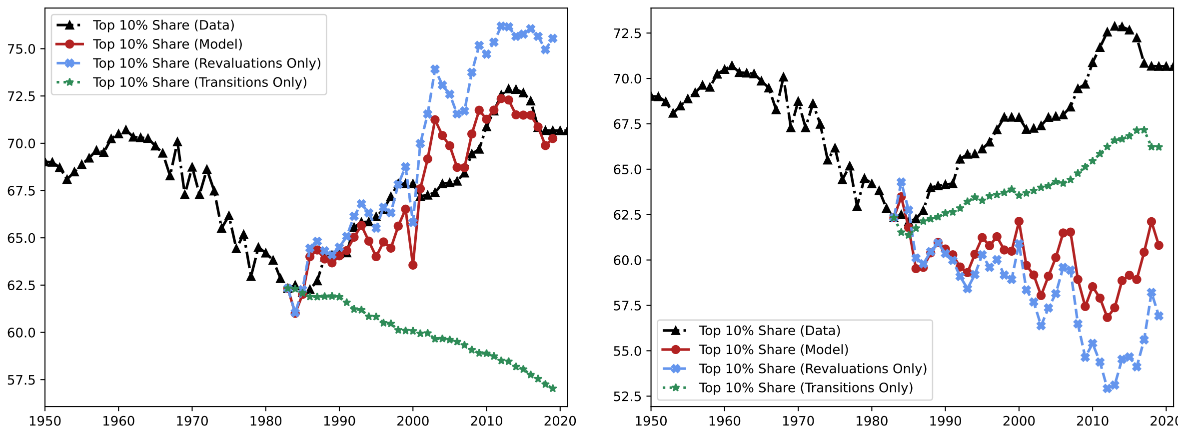

Note: The left panel plots the top-10% financial wealth share in the data (black dash-dotted line) and in the model with repricing (red solid line). It decomposes the evolution in the top-10% financial wealth share into a component solely due to instantaneous repricing (blue line, $\hat{S}_t^{10,REV}$) and component due to optimal consumption-savings decisions (green line, $\hat{S}_t^{10,MR}$). The right panel provides the same information as the left panel except for the compensated wealth distribution.

above, to be $d\hat{S}_t^{10,REV} = \tilde{S}_t^{10} - S_t^{10}$. We cumulate these revaluation effects to obtain the series

$$\hat{S}_t^{10,REV} = S_0^{10} + \sum_{\tau=1}^{t} d\hat{S}_\tau^{10,REV} \tag{21}$$

where $t = 0$ represents the base period 1983. To obtain an additive decomposition, we can define the *mean reversion effect*, corresponding to force (ii) above, to be equal to $\hat{S}_t^{10,MR} = S_t^{10} - \hat{S}_t^{10,REV}$.

The resulting series for $\hat{S}_t^{10,REV}$ and $\hat{S}_t^{10,MR}$ are displayed as the blue and green lines in Panel (a), respectively. The cumulative effect of instantaneous revaluations along the transition path (blue line) exceeds the overall path (red line), and explains more than 100% of the increase in inequality. The effect of revaluations is offset by households' endogenous consumption-savings responses (the mean reversion effect), which push wealth inequality down.

To understand this pattern we note that, despite the large increase in financial wealth inequality along the *transition* to low interest rates, the level of financial inequality in the low interest rate steady state is not higher than in the high interest rate steady state. In fact, it is somewhat lower. Intuitively, this lower steady state inequality results from fewer opportunities for households to build large wealth positions by compounding returns when rates are low. The large rise in inequality observed in Figure 6 is a temporary—albeit highly persistent—effect of a sequence of capital gains stemming from the fall in interest rates, rather than the low level of rates themselves. Thus, in the absence of future interest rate shocks, the model's endogenous transitions would see inequality decline after 2019 as it gradually transitions to the low-rate steady state. The

Table 3: Change in Inequality, Gradual Transition Experiment

	Data (WID)	Repriced	Compensated
Top-10% FW	+8.3pp	+7.9pp	-1.5pp
Top-1% FW	+11.3pp	+6.4pp	+0.6pp
Gini FW	+0.054	+0.061	-0.020
Top-10% HW	–	+1.1pp	+1.1pp
Top-1% HW	–	-2.2pp	-2.2pp
Gini HW	–	+0.066	+0.066
Top-10% TW	–	+0.6pp	-1.6pp
Top-1% TW	–	-1.9pp	-2.7pp
Gini TW	–	+0.059	+0.039

Note: The table reports the change in the Top-10% share, Top-1% share, and Gini Coefficient of financial wealth (FW, top panel), human wealth (HW, middle panel), and total wealth (TW, bottom panel). The change is measured between 1983 and 2019 in the model (Repriced and Compensated columns) as well as in the Data (WID) column.

very strong effects of revaluation are partially offset by this mean reversion effect, yielding our overall result that the decline in rates explains most of the rise in financial wealth inequality.

The resulting changes in financial wealth (FW) inequality are summarized in the top panel of Table 3. Table H1 in the appendix reports the levels rather than the changes for these same inequality moments. The column "Data (WID)" reports the change between 1983 and 2019 measured in the World Inequality Database, while the column "Repriced" displays the change over the same period in our model economy's gradual transition experiment (red line in Panel (a) of Figure 6). The increase in the top-10% share of financial wealth is virtually identical to that of the data. The increases in the FW top-1% share and FW Gini coefficient are also large and of the order as the actual increase in the data, although model understates the former and overstates the latter. In summary, our results indicate that the revaluation of assets following a decline in interest rates has been a powerful driver of inequality, and accounts for most, if not all, of the increase in financial wealth inequality since the 1980s.

Effects on Human and Total Wealth Inequality. Beyond financial wealth, we can also consider the effects of repricing on human wealth (the present value of lifetime income) and total wealth (the present value of lifetime consumption, the sum of financial and human wealth). In this section we continue our positive exercise measuring changes in the valuations of these streams. However, we note that a change in valuation does not necessarily reflect a change in welfare. For instance, a household with a fixed consumption stream will see its total wealth rise following a decline in rates even though its consumption is unchanged. We return to evaluate the *normative* implications for consumption opportunities in Section 6.

The Repriced column of the bottom two panels of Table 3 display the corresponding impacts of falling rates on human and total wealth inequality. The results are generally similar for human and total wealth, but vary substantially by the exact statistic used, with the Ginis rising by around as much as for financial wealth, the top-10% share rising, but by less, and the top-1% share actually falling. We explain these in turn.

To understand the change in the Gini and top-10% share, recall from Figure 4 that both human wealth and human wealth duration peak early in the life cycle and steadily decrease thereafter. As a result, the younger households with the most human wealth also experience the largest gains in human wealth, increasing human wealth inequality. At the same time, the distribution of human wealth duration is not as skewed as that of financial wealth, since human wealth durations are strictly limited by finite working lives and lifespans. As a result, we observe a broad-based increase in inequality that increases the Gini more than the top-10% share.

Turning to the total wealth response, recall that total wealth is the sum of financial and human wealth. Because human wealth makes up the larger share, the response of total wealth inequality is closer to that of human wealth inequality. However, because the gains from repricing to financial wealth (increasing in age) and human wealth (decreasing in age) are negatively correlated, the increases in the top-10% share and Gini for total wealth are smaller than for either component.

At the very top of the total wealth distribution, the response to the rate change is more subtle. The top-1% of the total wealth distribution is made up of two types of households. The first group consists of older households who hold most of their wealth in financial assets. These households have typically saved for a long time, and likely entered the superstar state sometime in the past but have since transitioned out of it. The wealth dynamics of this group are governed by the dynamics of the top-1% financial wealth share, which increases sharply when rates fall. The second group are households who currently are in the superstar income state. They are younger on average and have much higher ratios of human to total wealth. Since the superstar labor income state has a lower duration than the regular income state, due to a substantial exit probability of 2.5%, the value-weighted duration for the top-1% wealthiest share of households by human wealth is below the value-weighted duration for bottom-99% (generally, non-superstar) households. A decrease in rates thus leads to a decrease in top-1% human wealth share. Table 3 shows that this second group dominates the total response, causing the top-1% total wealth share to decline as rates fall.

5.3 Repricing: Robustness

Having established our main results on repricing, we now test their robustness to alternative assumptions on the duration of different asset classes and the evolution of income inequality.

Table 4: Transition Experiment, Alternative Specifications

Specification	Top-10% FW	Top-1% FW	Gini FW
Data (WID)	+8.3pp	+11.3pp	+0.054
Baseline	+7.9pp	+6.4pp	+0.061
Panel A. Robustness to Private Business Wealth			
1. Corporate PBW from IPO data	+8.0pp	+6.6pp	+0.061
2. Corporate PBW from Pitchbook	+8.0pp	+6.7pp	+0.062
3. Corporate PBW from SCF	+7.9pp	+6.4pp	+0.060
4. All PBW from SCF	+12.3pp	+10.8pp	+0.093
5. All PBW from equities	+10.1pp	+8.9pp	+0.078
Panel B. Robustness to Housing Wealth			
6. Housing from JST	+6.8pp	+5.4pp	+0.050
7. Excluding primary home	+8.7pp	+7.5pp	+0.069
Panel C. Time-Varying Income Risk			
8. Matching TV income inequality	+12.5pp	+12.7pp	+0.095
Panel D. Sources of Time Variation in Duration			
9. Assets have average duration	+6.5pp	+5.2pp	+0.050
10. Households have average portfolios	+7.6pp	+5.7pp	+0.059
11. Average duration and portfolios	+6.6pp	+4.5pp	+0.051

Private Business Wealth Duration. As noted above, the most challenging component of household wealth to measure is private business wealth. In Appendices D.4 - D.7, we discuss several alternative measures for the durations of Corporate and Non-Corporate PBW. Panel A of Table 4 displays the implied changes in financial wealth inequality under the same gradual transition experiment using one of these alternative measures. In rows 1 and 2, we use IPO and Pitchbook data in place of CRSP data on small stocks to measure the high-growth stage of the two-stage GGM. In row 3, we use SCF data to measure the duration of Corporate PBW, while in row 4 we use SCF data to also measure the duration of Non-Corporate PBW. Row 5 uses the duration of the overall equity market as an alternative proxy for the durations of both Corporate and Non-Corporate PBW. The first three alternatives result in similar increases in FW inequality, while the last two result in larger increases. Hence, we believe our baseline specification represents, if anything, conservative estimates for the impact of falling rates on wealth inequality.

Housing Duration. Panel B of Table 4 shows that our results are robust to the duration measurement of housing. First, row 6 repeats our measurement exercise obtaining prices and rents from the JST data, which results in similar but slightly lower increases in financial wealth inequality.

Next, we recompute our results excluding housing from the household portfolio altogether. Our baseline treatment of the primary home and the associated mortgage debt parallels that of other household assets. An alternative view is that households do not use their primary home as a transactable asset, but instead simply remain in it permanently, consuming the housing services. In this case, a change in the discount rate would increase the value of the house and the value of the lifetime consumption of housing services by the same amount, with no effect on consumption. While this view is a extreme, requiring that households never sell their property or extract equity from their home, we test robustness to it by excluding the primary home from the household's financial wealth portfolio, setting the duration of housing assets and mortgages to zero.[24] We find that our results strengthen relative to the baseline, with a rise in the top-10% financial wealth share of 8.7pp. The reason for this larger rise is that housing wealth is a relatively long duration asset held broadly across the top three quintiles of the wealth distribution, compared to equities and corporate private business wealth which are more concentrated toward the top. As a result, removing housing results in a larger gap between the average durations at the top and bottom of the wealth distribution, increasing the effect of repricing on financial wealth inequality.

Time-Varying Top Income Inequality While our main experiments apply a decline in interest rates holding income risk fixed, income inequality also rose over this period. Figure 7 Panel (a) plots the top-10% share from the WID data, showing an increase from 34% in 1983 to 46% in 2019. This increase could also contribute to rising wealth inequality.[25]

In this section, we extend our main experiment to match observed changes in top-income inequality alongside the decline in interest rates. Specifically, we re-estimate our regular income process (18) each year on rolling samples to generate distinct income risk parameters for each year of our transition experiment, as detailed in Appendix E.2. Given those time-varying income risk parameters, we then calibrate the superstar income state η^{sup} to exactly match the level of top-10% income inequality observed in the WID data in each year.

Figure 7 Panel (b) displays the resulting paths for financial wealth inequality in the gradual transition exercise, normalized to zero in 1983.[26] Combining declining interest rates and rising top income inequality generates an increase in the top-10% wealth share of 12.5pp, as shown in Panel C of Table 4. Our model-implied top-10% share fits the actual share well until 2000, before overstating the rise in inequality over the last 20 years of the sample. Notably, this version of the

[24] Sodini, Van Nieuwerburgh, Vestman, and von Lilienfeld-Toal (2023) provides evidence that households can and do borrow against rising collateral values to increase consumption opportunities.

[25] Hubmer, Krusell, and Smith (2020) show that this is not necessarily the case, using a model in which rising earnings risk actually lowers wealth inequality as it strengthens precautionary savings motives meaningfully for all but the richest households.

[26] Since we are no longer calibrating the superstar income state to match the 1983 top-10% wealth share in model and data, the initial levels of inequality differ in this exercise versus the benchmark model. For easier comparison, we normalize the plot by subtracting the level in 1983, so that each series is equal to zero in that year.

Figure 7: Top-10% Share, Gradual Transition + Time Varying Income Inequality Exercise

(a) Top-10% Income Share (WID)

(b) Transition Experiment: Rates + Income Ineq.

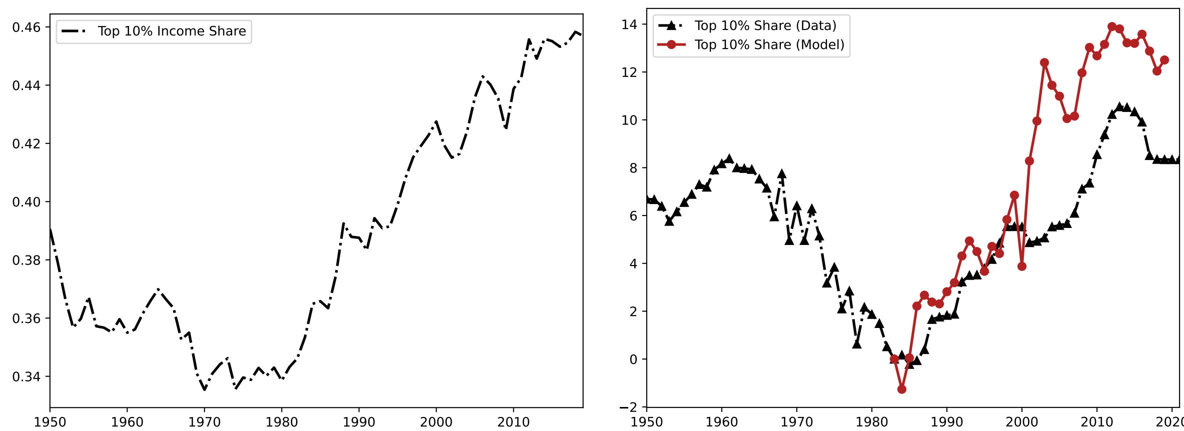

Note: The left panel shows the top-10% income share form the World Inequality database. The right panel plots the top-10% financial wealth share in the data (black line) and in the model that feeds in both gradual interest rate changes and gradual changes in the income process (red line).

model fully explains for the rise in the top-1% share of financial wealth, whereas the benchmark model accounted for less than 60% of it.

Contribution from Time-Varying Asset Durations vs. Portfolio Shares. Our benchmark results allow the duration of each asset and liability class, as well as household portfolio shares across asset classes, to vary over time. Panel D of Table 4 now isolates the importance of each source of time variation. Row 9 keeps the durations of all assets and liabilities constant over time, at their full-sample values, but maintains time variation in portfolio shares. This exercise generates a somewhat smaller increase in the top-10% income share of 6.5pp. This implies that asset durations in our baseline experiment were higher on average in periods when interest rates fell by more, increasing the overall impact. The opposite exercise, which holds households' portfolios constant over time but allows asset durations to fluctuate, also dampens the rise in inequality, implying that households held portfolios tilted toward higher duration assets when rates fell. However, the difference is much smaller, implying that time variation in portfolios is less important for our results. Last, row 11 holds both duration and portfolios constant, finding results close to those of row 9. These results show that the dynamics of asset durations contribute positively to the rise in inequality, while time variation in portfolio shares had much less impact.

6 Results: The Compensated Wealth Distribution

Having answered our first question on the quantitative role of interest rates in driving financial wealth inequality, we now turn to our second question: what are the implications of this change for consumption and total wealth inequality?

To measure the effects of interest rate changes on consumption opportunities, we compare the evolution of the financial wealth distribution in our baseline economy compared to a counterfactual *compensated* wealth distribution in which, following each shock, households receive exactly enough wealth to be able afford their previous consumption plans. We compute this as

$$\tilde{\theta}^{comp} = \mathbb{E}_0 \left\{ \sum_{t=0}^{T-1} \tilde{R}^{-t}(c_t - y_t) + \tilde{R}^{-T}\theta_T \right\} \tag{22}$$

where $\{c_t\}$ and θ_T (bequest) make up the pre-shock consumption plan, and $\tilde{R}$ is the post-shock interest rate. From Proposition 3, we know that if a household ends up with this amount of wealth, it will exactly keep its pre-shock consumption plan in the zero-mortality limit, meaning that under the compensated wealth distribution household consumption is approximately unaffected by interest rate shocks. Thus, we can use deviations between the repriced and compensated wealth distribution to measure the real consumption effects of a change in rates, as distinct from paper gains and losses. Because this measurement depends only on the budget constraint, and not on the household utility function, we view it as robust compared to welfare calculations that could depend heavily on household preferences.

6.1 The Compensated Distribution: One-Shot Experiment

We begin our analysis of the compensated distribution in our one-shot experiment. Figure 8 presents the distribution of financial wealth following our one-time decline in interest rates, alongside the original (pre-shock) distribution, showing two major differences between the two.[27] First, the compensated distribution is shifted substantially to the right from the original wealth distribution as the present values of households' excess consumption plans increase under low rates. As a result, the aggregate amount of financial wealth in the compensated distribution exceeds the pre-shock total by 196.7%. As can be seen from the plot, this rightward shift extends up to the very top, implying that even the wealthiest individuals must be compensated with additional financial assets to attain their old consumption plans. Indeed, nearly one third (29.0%) of new financial wealth accrues to top-1% financial wealth holders under the compensated distribution.

Second, although all households including the wealthiest see financial wealth increase under the compensated distribution, the less wealthy gain proportionally more, reducing financial

[27] To ensure that the full distribution is visible, we display transformed variables $\log(1 + x)$ on the x-axis.

Figure 8: Histogram, Compensated vs. Original Financial Wealth Distribution

(a) Original vs. Compensated

(b) Compensated vs. Repriced

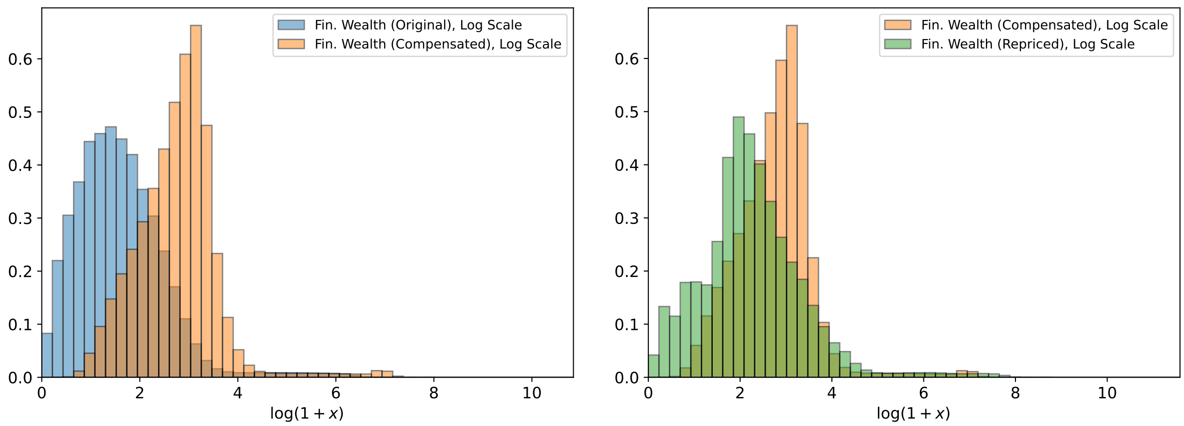

Note: This plot displays the distribution of financial wealth under the stationary distribution and under the compensated distribution drawn from the stationary distribution of the economy. The x-axis displays a transformation $\log(1 + x)$ of the original data. Each distribution is top coded at the top 0.1% of the pre-shock wealth distribution.

wealth inequality. Visually, while the original high interest-rate distribution of financial wealth is heavily right-skewed, the compensated distribution with low rates is actually left-skewed. Quantitatively, the share of financial wealth held by the top-10% decreases from 62.3% in the baseline economy to 52.3% in the compensated economy.

To see why inequality falls in the compensated distribution, we turn to Figure 9. Panel (a) compares medians by age under the original (horizontal axis) and compensated (vertical axis) financial wealth distributions by age. The youngest agents (light/yellow) have a low or intermediate level of wealth in the original distribution, but require the largest increase in financial wealth (vertical distance above the 45-degree line) in the compensated distribution. Households approaching retirement have more initial financial wealth, and require less compensation following the shock. Finally, the oldest households (dark/purple) have low wealth and require the least compensation.

This result may be surprising, since the young have the majority of their portfolio "invested" in human wealth, which has a long duration (bottom right panel of Figure 4), and thus provides a natural hedge against interest rate changes. However, the young plan to save during middle age ($c_t < y_t$), then dissave during retirement ($c_t > y_t$). This gives the young a very long duration of excess consumption D^{c-y}, making their original excess consumption plan much more expensive under low rates. In practical terms, the young will struggle to build retirement wealth and earn income on that wealth in retirement under low rates, making their pre-shock consumption plans unattainable without large infusions of financial wealth today. In contrast, older agents have lower values of D^{c-y} as retirement spending is brought closer to the present, and require less compensation. These households have already benefited from the higher rate of return in accumulating their retirement assets, while the oldest are already dissaving, consuming principal rather than

Figure 9: Scatterplots, Compensated vs. Original Financial Wealth Distribution

(a) Medians by Age

(b) Means by Fin. Wealth

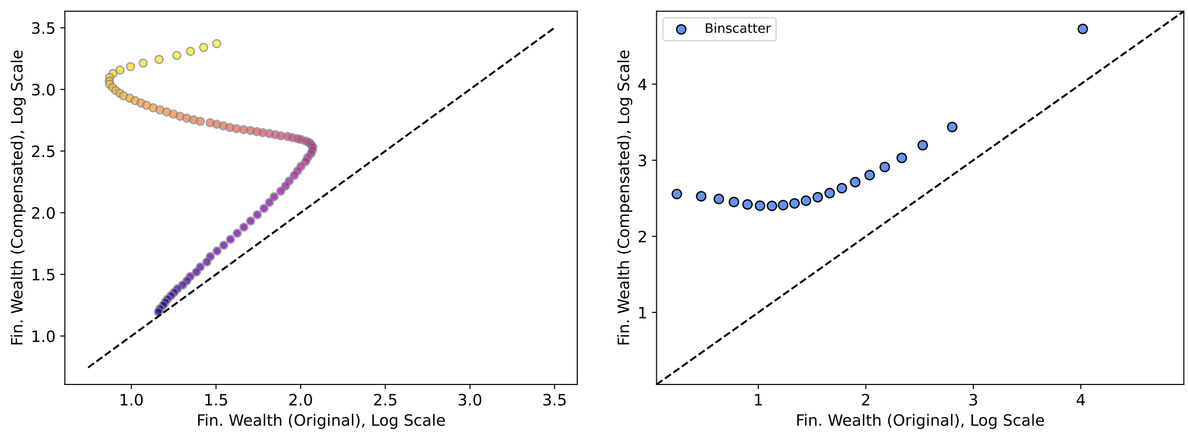

Note: Panel (a) plots the distribution of original financial wealth against the distribution of compensated financial wealth by age. Each dot represents the population mean for one year of age, with the lightest (yellow) dots representing the youngest agents and the darkest (purple) dots representing the oldest agents. Both variables are plotted using the transform $\log(1+x)$. The dashed line represents equality between the original and compensated distributions. Panel (b) plots the same distribution by bins of original financial wealth in place of age.

interest, making them less affected by the loss of high-return investment opportunities.

Panel (b) of Figure 9 aggregates over ages to present the total compensation required for various levels of pre-shock financial wealth. The lowest levels of financial wealth mix young agents who have not begun saving with old agents who are spending down assets late in life. As a result, the low wealth group mixes over agents requiring the largest and smallest amounts of compensation. Quantitatively, the young make up a larger share of this group and dominate the aggregate result, so that the least wealthy agents in this economy require the most compensation. Higher wealth levels contain an increasing share of middle-aged individuals, who require non-zero levels of compensation, but less than those at the bottom of the wealth distribution.

To separate the influence of age and wealth, Panel (a) of Appendix Figure H3 reproduces Panel (b) controlling for age fixed effects. It shows that wealthier households also require less compensation even after controlling for age.

To summarize, we observe a negative relationship between initial wealth and required compensation, due both to the life-cycle pattern of wealth, and to the links between wealth and required compensation independent of age. This negative relationship means that wealthier households need to gain the least under the compensated distribution, leading to a decrease in inequality in this counterfactual environment.

Comparison: Compensated vs. Repriced Distributions Having computed the compensated financial wealth distribution required to keep consumption plans constant, we can compare it to the repriced financial wealth distributions actually observed under low interest rates. Panel (b) of Figure 8 contrasts repriced and compensated distributions from our one-shot experiments. Since lower interest rates increase aggregate financial wealth by 208.4% by slightly more than the 196.7% increase in aggregate financial wealth required under the compensated distribution, there is more than enough wealth gain from repricing in aggregate to compensate all households. However, the compensated and repriced distributions display strikingly different shapes, with many more agents at low wealth levels in the repriced distribution compared to the compensated distribution.

To zoom in on the winners and losers from lower rates, Figure 10 compares changes in the repriced vs. compensated distributions by age in Panel (a) and by wealth in Panel (b). Panel (a) shows that while repricing delivers some gains to the young, it does not satisfy their large need for compensating transfers, leaving them well below the 45-degree line. In contrast, the middle-aged and old are over-hedged, receiving more wealth under repricing than needed to afford their former consumption plan, as shown by their position above the 45-degree line. Thus, the young will see their consumption possibilities contract, while older households will see them expand.

Panel (b) of Figure 10 displays the net gain from repricing, defined as the change in repriced wealth net of the change in compensated wealth. The figure reinforces our previous finding, showing that wealthier households gain on net from repricing, while the least wealthy experience a large net loss from the interest rate change, as repricing fails to appropriately compensate these households. Panel (b) of Appendix Figure H3 shows that the same pattern holds to a lesser degree across the wealth distribution after controlling for age fixed effects.

6.2 The Compensated Distribution: Transition Experiment

Our main quantitative experiment feeds in a gradual sequence of interest rate changes. The red line in Panel (b) of Figure 6 shows how the top-10% wealth share would have evolved under the compensated wealth distribution — the path of inequality that would have been required to keep consumption plans constant. The last column of Table 3 summarizes these changes. The compensated wealth distribution sees a substantial reduction in financial wealth inequality. The top-10% financial wealth share falls by -1.5pp between 1983 and 2019 and the Gini falls by -0.020 , while the top-1% share rises modestly by +0.6pp. The effect is largely captured by the instantaneous compensation following each interest rate change (blue line in Panel (b) of Figure 6), somewhat offset by the effect coming from consumption-savings choices (green line).

In sum, the starkly diverging paths for financial wealth inequality in the repriced and compensated distributions imply that changing interest rates do not merely result in paper gains but result in important changes in consumption possibilities. The young and poor see substantial deterio-

Figure 10: Scatterplots, Repriced Financial Wealth Distribution

(a) Compensated vs. Repriced: Age Medians

(b) Scatter: Repriced Gain

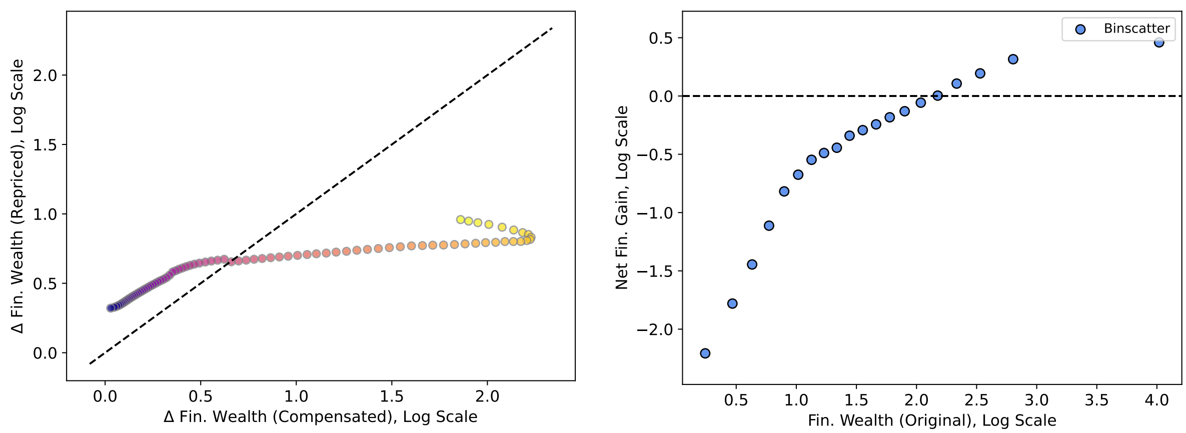

Note: This plot displays the distribution of financial wealth under the repriced distribution, compared to the compensated distribution. Panel (a) displays the change in financial wealth relative to the original distribution for the compensated (x-axis) and repriced (y-axis) distributions. Both axes display a transformation $\log(1+x)$ of the original data. Each dot represents one year of age, with the lightest (yellow) dots representing the youngest agents and the darkest (purple) dots representing the oldest agents. Panel (b) displays original financial wealth on the x-axis and the net financial gain (repriced minus compensated wealth) on the y-axis. The x-axis displays the transform $\log(1+x)$, while the y-axis displays the difference in transformed values. Each dot represents 5% of households from the original wealth distribution. All distributions are drawn from the stationary distribution of the economy.

ration in their consumption because the duration of their financial wealth is below the duration of their excess consumption. The opposite is true for the old and the rich, whose consumption opportunities expand from declining rates.

Last, we can evaluate changes under the compensated distribution in total wealth inequality, which summarizes household lifetime consumption. Table 3 shows that total wealth inequality increases by less (or decreases by more) under the compensated distribution compared to the repriced distribution. Thus, our conclusions regarding the real effects of falling rates on inequality carry over from financial wealth to total wealth. Quantitatively, however, the gaps between repriced and compensated inequality measures are substantially smaller for total than for financial wealth. This reflects the influence of human wealth, which represents the majority of total wealth, and undergoes identical changes in the compensated and repriced distributions.

7 Cohort-Level Analysis

For our final set of results, we use the model to shed light on how different cohorts have gained and lost over our sample period. Recall that a decline in rates generally harms households when they are younger ($D^{\theta} < D^{c-y}$) and benefits them when they are older ($D^{\theta} > D^{c-y}$). Thus, to

Figure 11: Cohort Total Wealth Outcomes, Transition Experiment

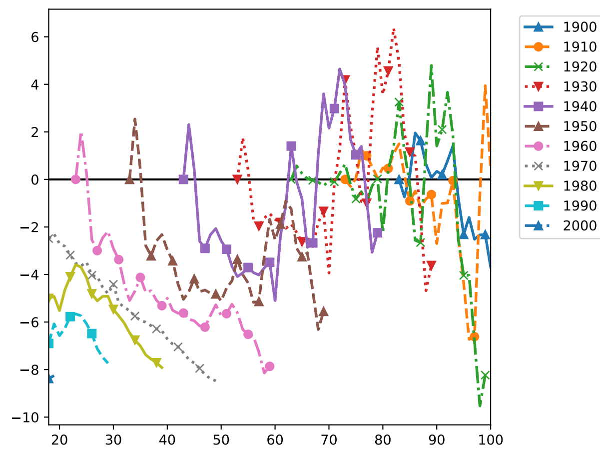

Note: The graph plots total wealth, the present discounted value of consumption, along the observed interest rate path for the median member of a birth cohort, indicated by the various colored lines, in percent deviation of what that wealth path would have been in the steady state with high (1983) interest rates. Each cohort is labeled with the first year of the relevant birth decade, for example "1900" represents households born between 1900 and 1909.

determine the total impact of falling rates on a given household, it matters not only how much rates fell during its lifetime, but also exactly when during its life cycle this occurred.

To address this, Figure 11 plots total wealth (present value of lifetime consumption) for the median member of each birth-decade cohort at each age in their lives. We normalize the series as percent deviations in median total wealth compared to households of the same age in the initial 1983 stationary distribution. These series thus represent the effect of the fall in rates on the lifetime consumption value of each cohort compared to a world in which rates remained unchanged. Each line represents a cohort of ten birth years, with the oldest cohort (born 1900–1909) well into retirement at the start of our sample, and the youngest cohort (born 2000–2009) entering the workforce only in the sample's final year.

The graph shows that the older cohorts gain while the younger cohorts lose. Households born before 1920 observe only modest gains since they have mostly run down their wealth by the time the interest rate declines begin. Households born in the 1920s through 1940s are the biggest winners, experiencing peak total wealth gains in excess of 5%. The youngest cohorts (Gen X, Millennials, and Gen Z) strictly lose, experiencing peak total wealth losses approaching 8% compared to households of the same age in the stationary economy.

8 Conclusion

A persistent decline in real interest rates, like the one experienced in much of the world between the 1980s and the 2010s, leads to a rise in financial wealth inequality when there is a positive covariance between financial wealth levels and the duration of financial wealth across households. Using detailed portfolio data, we show that this condition is met in the U.S. data, and that the duration heterogeneity is large enough to account for the entire rise in the top-10% share of financial wealth. With the help of a standard consumption-savings model, we show that the reduction in interest rates not only leads to "paper valuation gains" but affects consumption possibilities. In particular, young and less wealthy households are forced to save at lower rates for their retirement by purchasing more expensive assets in the future. They see their consumption possibilities contract when rates fall. Older and wealthier households have more than enough duration in their portfolio to allow them to afford the old consumption plan under the new, lower interest rates, thanks to large capital gains. We show how these effects played out in the data by studying how different cohorts' consumption possibilities were affected by the observed path of interest rates.

Recently, long-term real rates have begun to rise after a 40-year decline. Between March 2022 and August 2023, the 10-year real bond yield increased 2.65% points. Our paper predicts that this sharp rise in real rates will lower financial wealth inequality and benefit the consumption opportunities of the young and the poor in the years to come.

References

Alvaredo, Facundo, Lucas Chancel, Thomas Piketty, Emmanuel Saez, and Gabriel Zucman. 2018. *World Inequality Report 2018*. Belknap Press.

Alvarez, Fernando and Urban J Jermann. 2001. "Quantitative Asset Pricing Implications of Endogenous Solvency Constraints." *The Review of Financial Studies* 14 (4):1117–1151.

Auclert, Adrien. 2019. "Monetary Policy and the Redistribution Channel." *American Economic Review* 109 (6):2333–2367.

Auclert, Adrien, Hannes Malmberg, Frederic Martenet, and Matthew Rognlie. 2021. "Demographics, Wealth, and Global Imbalances in the Twenty-First Century." Tech. rep., National Bureau of Economic Research.

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

Benhabib, Jess, Alberto Bisin, and Mi Luo. 2017. "Earnings Inequality and Other Determinants of Wealth Inequality." *American Economic Review* 107 (5):593–597.

Bernanke, Ben. 2005. "The Global Saving Glut and the US Current Account Deficit." Tech. rep., Board of Governors of the Federal Reserve System (US).

Bewley, Truman. 1986. "Stationary Monetary Equilibrium With a Continuum of Independently Fluctuating Consumers." In *Contributions to Mathematical Economics in Honor of Gerard Debreu*, edited by Werner Hildenbrand and Andreu Mas-Collel. Amsterdam: North-Holland.

Boar, Corina and Virgiliu Midrigan. 2020. "Efficient Redistribution." Working Paper.

Caballero, Ricardo J, Emmanuel Farhi, and Pierre-Olivier Gourinchas. 2008. "An Equilibrium Model of 'Global Imbalances' and Low Interest Rates." *American Economic Review* 98 (1):358–93.

Castaneda, Ana, Javier Diaz-Gimenez, and Jose-Victor Rios-Rull. 2003. "Accounting for the US Earnings and Wealth Inequality." *Journal of Political Economy* 111 (4):818–857.

Catherine, Sylvain, Max Miller, James D Paron, and Natasha Sarin. 2022. "Who Hedges Interest-Rate Risk? Implications for Wealth Inequality." *Implications for Wealth Inequality (May 23, 2022). Jacobs Levy Equity Management Center for Quantitative Financial Research Paper* .

Catherine, Sylvain, Max Miller, and Natasha Sarin. 2020. "Social Security and Trends in Inequality." Working Paper.

Cox, Josue. 2020. "Return Heterogeneity, Information Frictions, and Economic Shocks." Working Paper New York University.

De Nardi, Mariacristina. 2004. "Wealth Inequality and Intergenerational Links." *The Review of Economic Studies* 71 (3):743–768.

Deaton, Angus S and Christina Paxson. 1994. "Saving, Growth, and Aging in Taiwan." In *Studies in the Economics of Aging*. University of Chicago Press, 331–362.

Doepke, Matthias and Martin Schneider. 2006. "Inflation and the Redistribution of Nominal Wealth." *Journal of Political Economy* 114 (6):1069–1097.

Eggertsson, Gauti B and Neil R Mehrotra. 2014. "A Model of Secular Stagnation." Tech. rep., National Bureau of Economic Research.

Eichengreen, Barry. 2015. "Secular Stagnation: The Long View." *American Economic Review* 105 (5):66–70.

Fagereng, Andreas, Matthieu Gomez, Emilien Gouin-Bonenfant, Martin Holm, Benjamin Moll, and Gisle Natvik. 2022. "Asset-Price Redistribution." Tech. rep., Working Paper.

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

Gomez, Matthieu and Emilien Gouin-Bonenfant. 2020. "A Q-Theory of Inequality." Working Paper, Columbia University.

Gupta, Arpit and Stijn Van Nieuwerburgh. 2021. "Valuing Private Equity Investments Strip by Strip." *The Journal of Finance* 76 (6):3255–3307.

Gutiérrez, Germán and Thomas Philippon. 2017. "Declining Competition and Investment in the US." Tech. rep., National Bureau of Economic Research.

Heathcote, Jonathan, Fabrizio Perri, and Giovanni L Violante. 2010. "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967–2006." *Review of Economic Dynamics* 13 (1):15–51.

Hubmer, Joachim, Per Krusell, and Anthony A. Smith. 2020. "Sources of U.S. Wealth Inequality: Past, Present, and Future." In *NBER Macroeconomics Annual*, vol. 35.

Jiang, Zhengyang, Hanno Lustig, Stijn Van Nieuwerburgh, and Mindy Xiaolan. 2019. "The U.S. Public Debt Valuation Puzzle." NBER Working Paper No. 26583.

———. 2021. "Public Debt Valuation: The UK Over the Long Run." Working Paper.

Jordà, Òscar, Moritz Schularick, and Alan M Taylor. 2017. "Macrofinancial History and the New Business Cycle Facts." *NBER Macroeconomics Annual* 31:213–263.

Kopczuk, Wojciech. 2017. "U.S. Capital Gains and Estate Taxation: A Status Report and Directions for a Reform." *The Economics of Tax Policy* :265–91.

Kopczuk, Wojciech and Eric Zwick. 2020. "Business Incomes at the Top." *Journal of Economic Perspectives* 34 (4):27–51.

Krueger, Dirk and Hanno Lustig. 2010. "When Is Market Incompleteness Irrelevant for the Price of Aggregate Risk (And When Is It Not)?" *Journal of Economic Theory* 145 (1):1–41.

Krueger, Dirk and Fabrizio Perri. 2006. "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory1." *The Review of Economic Studies* 73 (1):163–193.

Kuhn, Moritz, Moritz Schularick, and Ulrike I Steins. 2020. "Income and Wealth Inequality in America, 1949–2016." *Journal of Political Economy* 128 (9):000–000.

Leombroni, Matteo, Monika Piazzesi, Martin Schneider, and Ciaran Rogers. 2020. "Inflation and the Price of Real Assets." Working Paper 26740, National Bureau of Economic Research.

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

Mian, Atif, Ludwig Straub, and Amir Sufi. 2020. "The Saving Glut of the Rich and the Rise in Household Debt." NBER Working Paper No. 26941.

Moll, Ben. 2020. "Comment on Hubmer, Krusell & Smith "Sources of US Wealth Inequality: Past, Present, and Future"." 469–479.

Piketty, Thomas. 2015. "About Capital in the Twenty-First Century." *American Economic Review* 105 (5):48–53.

Piketty, Thomas and Emmanuel Saez. 2003. "Income Inequality in the United States, 1913–1998." *The Quarterly Journal of Economics* 118 (1):1–41.

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

Piketty, Thomas and Gabriel Zucman. 2015. "Wealth and Inheritance in the Long Run." In *Handbook of income distribution*, vol. 2. Elsevier, 1303–1368.

Quadrini, V and José-Vıctor Rıos-Rull. 1997. "Dimensions of Inequality: Facts on the US Distribution of Earnings, Income and Wealth." *Federal Reserve Bank of Minneapolis Quarterly Review* 21 (2):3–21.

Rachel, Lukasz and Thomas Smith. 2017. "Are Low Real Interest Rates Here to Stay?" *International Journal of Central Banking* 13 (3).

Ritter, Jay R. 2022. "IPO Statistics for 2021." https://site.warrington.ufl.edu/ritter/ipo-data/. Accessed: 2022-12-7.

Rouwenhorst, K Geert. 1995. "Asset Pricing Implications of Equilibrium Business Cycle Models." *Frontiers of Business Cycle Research* 1:294–330.

Saez, Emmanuel and Gabriel Zucman. 2016. "Wealth Inequality in the United States Since 1913: Evidence From Capitalized Income Tax Data." *The Quarterly Journal of Economics* 131 (2):519–578.

Smith, Matthew, Danny Yagan, Owen Zidar, and Eric Zwick. 2019. "Capitalists in the Twenty-First Century." *The Quarterly Journal of Economics* 134 (4):1675–1745.

———. 2022. "The Rise of Pass-Throughs and the Decline of the Labor Share." *American Economic Review: Insights* Forthcoming.

Sodini, Paolo, Stijn Van Nieuwerburgh, Roine Vestman, and Ulf von Lilienfeld-Toal. 2023. "Identifying the Benefits From Homeownership:A Swedish Experiment." *American Economic Review* Forthcoming.

Straub, Ludwig. 2019. "Consumption, Savings, and the Distribution of Permanent Income." Unpublished Manuscript, Harvard University.

Summers, Lawrence H. 2014. "US Economic Prospects: Secular Stagnation, Hysteresis, and the Zero Lower Bound." *Business Economics* 49 (2):65–73.

Online Appendix

A Affine Asset Pricing Model

This appendix develops a reduced-form asset pricing model. The asset pricing model is used for two main purposes. First, to compute long-term real bonds yields, the cost of a 30-year real annuity, and expected returns on stocks and housing wealth. Second, to compute the McCauley duration of the aggregate stock market, small stocks, and housing wealth in a manner that is consistent with the history of bond and stock prices.

The asset pricing model in the class of exponentially-affine SDF models. A virtue of the reduced-form model is that it can accommodate a substantial number of aggregate risk factors. We argue that it is important to go beyond the aggregate stock and bond markets to capture the risk embedded in households' financial asset portfolios as well as the aggregate risk in consumption and labor income claims. Similar models are estimated in Lustig, Van Nieuwerburgh, and Verdelhan (2013); Jiang, Lustig, Van Nieuwerburgh, and Xiaolan (2019); Gupta and Van Nieuwerburgh (2021).

A.1 Setup

A.1.1 State Variable Dynamics

Time is denoted in quarters. We assume that the $N \times 1$ vector of state variables follows a Gaussian first-order VAR:

$$\boldsymbol{z}_t = \boldsymbol{\Psi} \boldsymbol{z}_{t-1} + \boldsymbol{\Sigma}^{\frac{1}{2}} \boldsymbol{\varepsilon}_t, \tag{23}$$

with shocks $\boldsymbol{\varepsilon}_t \sim i.i.d. \mathcal{N}(0, I)$ whose variance is the identity matrix. The companion matrix Ψ is a $N \times N$ matrix. The vector z is demeaned. The covariance matrix of the innovations to the state variables is Σ; the model is homoscedastic. We use a Cholesky decomposition of the covariance matrix, $\Sigma = \Sigma^{\frac{1}{2}} \Sigma^{\frac{1}{2}\prime}$, which has non-zero elements only on and below the diagonal. The Cholesky decomposition of the residual covariance matrix allows us to interpret the shock to each state variable as the shock that is orthogonal to the shocks of all state variables that precede it in the VAR. We discuss the elements of the state vector and their ordering below. The (demeaned) one-quarter bond nominal yield is one of the elements of the state vector: $y_{t,1}^{\$} = y_{0,1}^{\$} + e_{yn}' z_t$, where $y_{0,1}^{\$}$ is the unconditional average 1-quarter nominal bond yield and e_{yn} is a vector that selects the element of the state vector corresponding to the one-quarter yield. Similarly, the (demeaned) inflation rate is part of the state vector: $\pi_t = \pi_0 + e_{\pi}' z_t$ is the (log) inflation rate between $t-1$ and t. Lowercase letters denote logs.

A.1.2 Stochastic Discount Factor

The nominal SDF $M^{\$}_{t+1} = \exp(m^{\$}_{t+1})$ is conditionally log-normal:

$$m^{\$}_{t+1} = -y^{\$}_{t,1} - \frac{1}{2}\mathbf{\Lambda}'_t\mathbf{\Lambda}_t - \mathbf{\Lambda}'_t\varepsilon_{t+1}. \tag{24}$$

Note that $y^{\$}_{t,1} = -\mathbb{E}_t[m^{\$}_{t+1}] - 0.5\text{Var}_t[m^{\$}_{t+1}]$. The real log SDF $m_{t+1} = m^{\$}_{t+1} + \pi_{t+1}$ is also conditionally Gaussian. The innovations in the vector ε_{t+1} are associated with a $N \times 1$ market price of risk vector $\mathbf{\Lambda}_t$ of the affine form:

$$\mathbf{\Lambda}_t = \mathbf{\Lambda}_0 + \mathbf{\Lambda}_1 z_t. \tag{25}$$

The $N \times 1$ vector $\mathbf{\Lambda}_0$ collects the average prices of risk while the $N \times N$ matrix $\mathbf{\Lambda}_1$ governs the time variation in risk premia. Asset pricing amounts to estimating the market prices of risk $(\mathbf{\Lambda}_0, \mathbf{\Lambda}_1)$. We specify the moment conditions used to identify the market prices of risk below.

A.1.3 State Vector Elements

The state vector contains the following $N = 22$ variables, in order of appearance: (1) real GDP growth, (2) GDP price inflation, (3) the nominal short rate (3-month nominal Treasury bill rate), (4) the spread between the yield on a five-year Treasury note and a three-month Treasury bill, (5) the log price-dividend ratio on the CRSP value-weighted stock market, (6) the log real dividend growth rate on the CRSP stock market. Elements 7, 9, 11, and 13 are the log price-dividend ratios on the first size quintile of stocks (small), the first book-to-market quintile of stocks (growth), the fifth book-to-market quintile of stocks (value), and a listed infrastructure index (infra). Elements 8, 10, 12, and 14 are the corresponding log real dividend growth rates. Element 15 is the log price-dividend ratio on housing wealth, element 16 is log real dividend growth on housing wealth. Finally, the state vector contains the log change in the consumption/GDP ratio Δcx in 17th, the log change in the log labor income/GDP ratio Δlx in 18th, the log level of the consumption/GDP ratio cx in 19th, and the log level of the labor income/GDP ratio lx in 20th position.

$$\begin{aligned} z_t = \quad & \Big[\pi_t, x_t, y^{\$}_{t,1}, y^{\$}_{t,20} - y^{\$}_{t,1}, pd^m_t, \Delta d^m_t, pd^{small}_t, \Delta d^{small}_t, \\ & pd^{growth}_t, \Delta d^{growth}_t, pd^{value}_t, \Delta d^{value}_t, pd^{infra}_t, \Delta d^{infra}_t \\ & pd^{hw}_t, \Delta d^{hw}_t, \Delta cx_{t+1}, \Delta lx_{t+1}, cx_{t+1}, lx_{t+1}\Big]'. \end{aligned} \tag{26}$$

This state vector is observed at quarterly frequency from 1947.Q1 until 2019.Q4 (292 observations). This is the longest available time series for which all variables are available. Inflation is the log change in the GDP price deflator. For the yields, we use the average of daily Constant

Maturity Treasury yields within the quarter. All dividend series are deseasonalized by summing dividends across the current month and past 11 months. Small stocks are the bottom 20% of the market capitalization distribution, growth stocks the bottom 20% of the book-to-market distribution, and value stocks the top 20% of the book-to-market distribution. The infrastructure stock index is measured as the value-weighted average of the eight relevant Fama-French industries (Aero, Ships, Mines, Coal, Oil, Util, Telcm, Trans). We subtract inflation from all nominal dividend growth rates to obtain real dividend growth rates.

Dividend growth on housing wealth is measured as housing services consumption growth from the Bureau of Economic analysis Table 2.3.5. The price-dividend ratio is the ratio of owner-occupied housing wealth from the Financial Accounts of the United States Table B.101.h divided by housing services consumption. The resulting price-dividend ratio on housing wealth averages 16.1 (for annualized dividends) between 1947 and 2019. We subtract inflation from dividend growth on housing wealth and we also subtract 0.6% per quarter to reflect the fact that the size of the housing stock is growing and we are only interested in the rental price change, not the change in the quantity of housing. The resulting real rental growth rate is 1.82% per year, which is in line with (and still on the higher end of the numbers reported in) the literature.

Aggregate consumption is measured as non-durables plus services plus durable services consumption. Durable services consumption is constructed as the depreciation rate (20%) multiplied by the stock of durables. The stock of durables itself is computed using the perpetual inventory method. This series is divided by nominal GDP and logs are taken.

Aggregate labor income is measured as wages and salaries plus business income (proprietors' income with inventory valuation and capital consumption adjustments) plus transfer income (personal current transfer receipts) minus taxes (Personal current taxes and Contributions for government social insurance, domestic). This series is divided by nominal GDP and logs are taken. Real consumption growth can then be written as the sum of real GDP growth plus the change in the consumption/GDP ratio:

$$\Delta c^a_{t+1} = x_{t+1} + \Delta cx_{t+1}$$

and similar for labor income growth.

All state variables are demeaned with the observed full-sample mean. The first 18 equations of the VAR are estimated by OLS equation by equation. We recursively zero out all elements of the companion matrix Ψ whose t-statistic is below 2.2. The resulting point estimates for Ψ and $\Sigma^{\frac{1}{2}}$ are reported below.

The dynamics of cx are pinned down by the dynamics of Δcx:

$$cx_{t+1} = cx_t + \Delta cx_{t+1} = \left(e_{cx} + e_{cxgr}\Psi\right)' z_t + e_{cxgr}\gamma^{\frac{1}{2}}\varepsilon_{t+1}$$

Therefore the 19st row of Ψ is identical to the 17th row, except that $\Psi(19,19) = \Psi(17,19) + 1$.

Similarly, the 20th row of Ψ is identical to the 18th row, except that $\Psi(20,20) = \Psi(18,20) + 1$. The innovations to the 19th and 20th row are not independent innovations but determined by the innovations that precede it. The level variables cx and lx are only added to the VAR to enforce cointegration between consumption and GDP and between labor income and GDP. As a result of this cointegration, the aggregate consumption and labor income claims will have the same aggregate risk as the GDP claim.

A.2 Estimation

A.2.1 Bond Pricing

In this setting, nominal bond yields of maturity τ are affine in the state variables:

$$y_{t,\tau}^{\$} = -\frac{1}{\tau}A_{\tau}^{\$} - \frac{1}{\tau}\left(B_{\tau}^{\$}\right)' z_t.$$

The scalar $A^{\$}(\tau)$ and the vector $B_{\tau}^{\$}$ follow ordinary difference equations (ODE) that depend on the properties of the state vector and on the market prices of risk. Real bond yield are also exponentially affine with coefficients that follow their own ODEs. We will price the cross-section of nominal and real bond yields (price levels), putting more weight on matching the time series of one- and twenty-quarter nominal bond yields since those yields are part of the state vector $\mathbf{z}_t$. We also fit the dynamics of 20-quarter nominal bond risk premia (price changes).

Figure A1 plots the nominal bond yields on bonds of maturities 1 quarter, 1-, 2-, 3-, 5-, 7-, 10-, 20-, and 30-years. These are all available bond yields in the data. The 20-, and 3-year bond yields are not available in parts of the sample, but the estimation minimizes the distance between observed and model-implied yields for every period where data is available. The model matches the time series of bond yields in the data closely. It matches nearly perfectly the 1-quarter and 5-year bond yield which are part of the state space.

Figure A2 shows that the model also does a good job matching real bond yields. These yields are available over a much shorter sample in the data, and we only plot the relevant subsample for the model-implied yields as well.

The top panels of Figure A3 show the model's implications for the average nominal (left panel) and real (right panel) yield curves at longer maturities. These long-term yields are well behaved. The bottom left panel shows that the model matches the dynamics of the nominal bond risk premium, defined as the expected excess return on five-year nominal bonds. The compensation for interest rate risk varies substantially over time, both in data and in the model. The bottom right panel shows a decomposition of the yield on a five-year nominal bond into the five-year real bond yield, annual expected inflation over the next five years, and the five-year inflation risk premium. The importance of these components fluctuates over time. This graph shows the secular rise and fall of real bond yields, with a peak in the early 1980s.

Figure A1: Dynamics of the Nominal Term Structure of Interest Rates

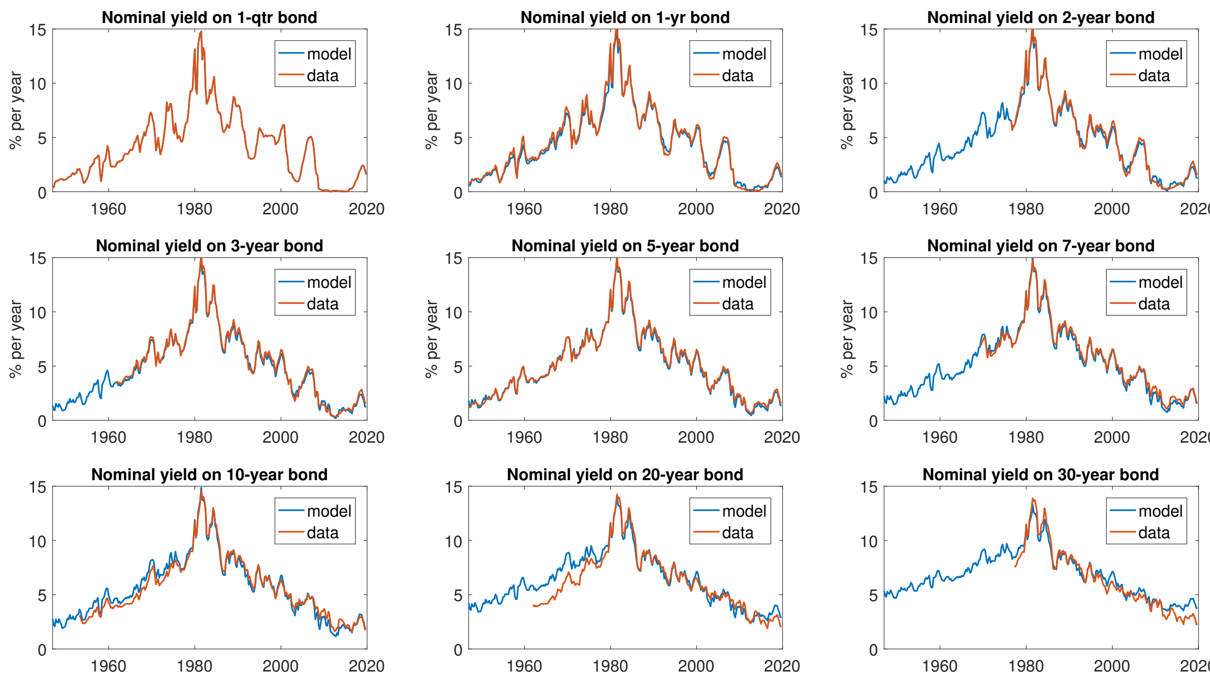

Note: The figure plots the observed and model-implied nominal bond yields. Data are from FRED: constant-maturity Treasury yields, daily averages within the quarter.

A.2.2 Equity Factors and Housing Wealth Pricing

The VAR contains both the log price-dividend ratio and log dividend growth for five equity risk factors (the aggregate stock market, small stocks, growth stocks, value stocks, and infrastructure stocks), and residential real estate wealth. Together these two time-series imply a time-series for log returns through the definition of a log stock return. Hence, the VAR implies linear dynamics for the expected excess stock return, or equity risk premium, for each of these seven assets. We estimate market prices of risk to match the VAR-implied risk premium levels and dynamics.

The price of a stock equals the present-discounted value of its future cash-flows. By value-additivity, the price of the aggregate stock index, P_t^m, is the sum of the prices to each of its future cash-flows D_t^m. These future cash-flow claims are the so-called market dividend strips or zero-coupon equity. Dividing by the current dividend D_t^m:

$$\frac{P_t^m}{D_t^m} = \sum_{\tau=1}^{\infty} P_{t,\tau}^d \tag{27}$$

$$\exp\left(\overline{pd} + e_{pd^m}' z_t\right) = \sum_{\tau=0}^{\infty} \exp\left(A_\tau^m + B_\tau^{m\prime} z_t\right), \tag{28}$$

Figure A2: Dynamics of the Real Term Structure of Interest Rates

Note: The figure plots the observed and model-implied real bond yields. Data are from FRED: constant-maturity Treasury inflation-indexed bond yields, daily averages within the quarter.

where $P^d_{t,\tau}$ denotes the price of a τ-period dividend strip divided by the current dividend. The log price-dividend ratio on each dividend strip, $p^d_{t,\tau} = \log\left(P^d_{t,\tau}\right)$, is affine in the state vector and the coefficients $\left(A^m_\tau, \boldsymbol{B}^m_\tau\right)$ follow an ODE. Since the log price-dividend ratio on the stock market is an element of the state vector, it is affine in the state vector by assumption. Equation (28) restates the present-value relationship from equation (27). It articulates a non-linear restriction on the coefficients $\{(A^m_\tau, B^m_\tau)\}_{\tau=1}^{\infty}$ at each date (for each state z_t), which we impose in the estimation. Analogous present value restrictions are imposed for each of the other four equity factors, and for housing wealth.

If dividend growth were unpredictable and its innovations carried a zero risk price, then dividend strips would be priced like real zero-coupon bonds. The strips' dividend-price ratios would equal yields on real bonds with the coupon adjusted for deterministic dividend growth. All variation in the price-dividend ratio would reflect variation in the real yield curve. In reality, the dynamics of real bond yields only account for a small fraction of the variation in the price-dividend ratio, implying large prices of risk associated with shocks to dividend growth that are orthogonal to shocks to bond yields. Hence, matching price-dividend ratios (price levels) and expected returns (price changes) allow us to pin down the market prices of risk associated with orthogonal dividend growth shocks (shocks to the state variables in rows 6, 8, 10, 12, 14, 16, and 18 of the

Figure A3: Long-term Yields and Bond Risk Premia

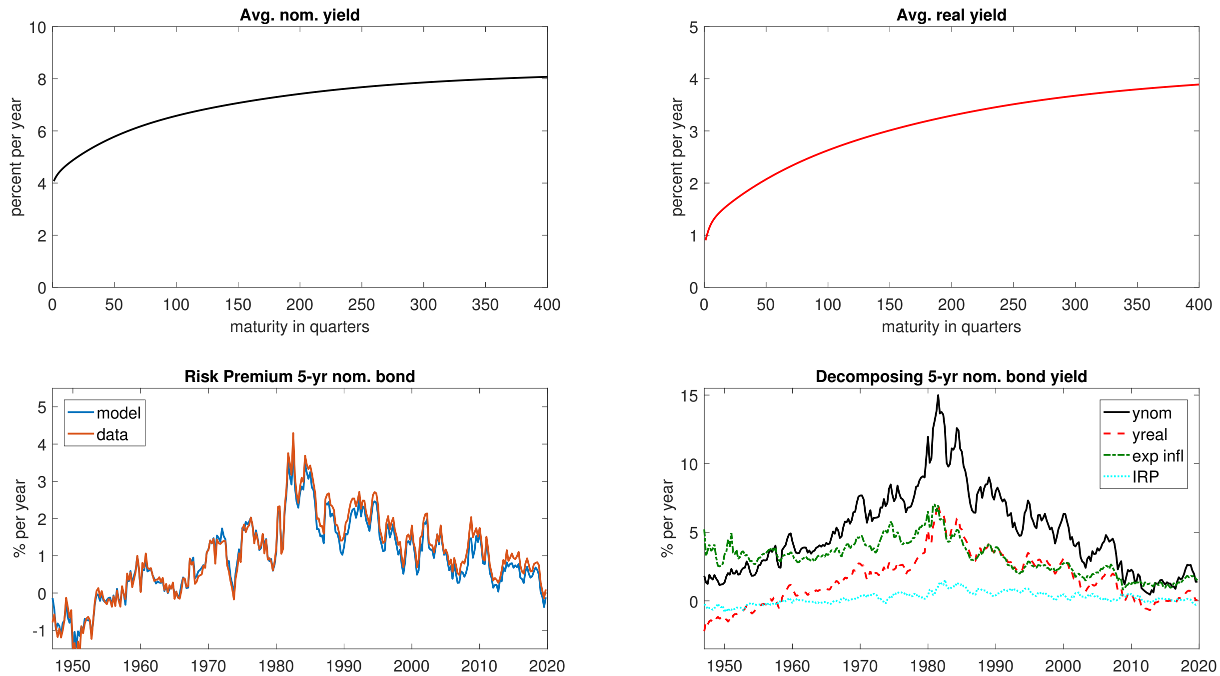

Note: The top panels plot the average bond yield on nominal (left panel) and real (right panel) bonds for maturities ranging from 1 quarter to 400 quarters. The bottom left panel plots the nominal bond risk premium in model and data. The bottom right panel decomposes the model's five-year nominal bond yield into the five-year real bond yield, the five-year inflation risk premium and the five-year real risk premium.

VAR).

Figures A4 and A5 show the equity risk premium, the expected excess return, in the left panels and the price-dividend ratio in the right panels. The various rows cover the five equity indices and the housing wealth series we price. The dynamics of the risk premia in the data are dictated by the VAR. The model chooses the market prices of risk to fit these risk premium dynamics as closely as possible alongside with the price-dividend ratio levels. The price-dividend ratios in the model are formed from the price-dividend ratios on the strips of maturities ranging from 1 to 3600 quarters, as explained above. The figure shows an excellent fit for price-dividend levels and a good fit for risk premium dynamics. Some of the VAR-implied risk premia have outliers which the model does not fully capture. This is in part because the good deal bounds restrict the SDF from becoming too volatile and extreme. We note large level differences in valuation ratios across the various stock factors, as well as big differences in the dynamics of both risk premia and price levels, which the model is able to capture well.

Figure A4: Equity Risk Premia and Price-Dividend Ratios (1/2)

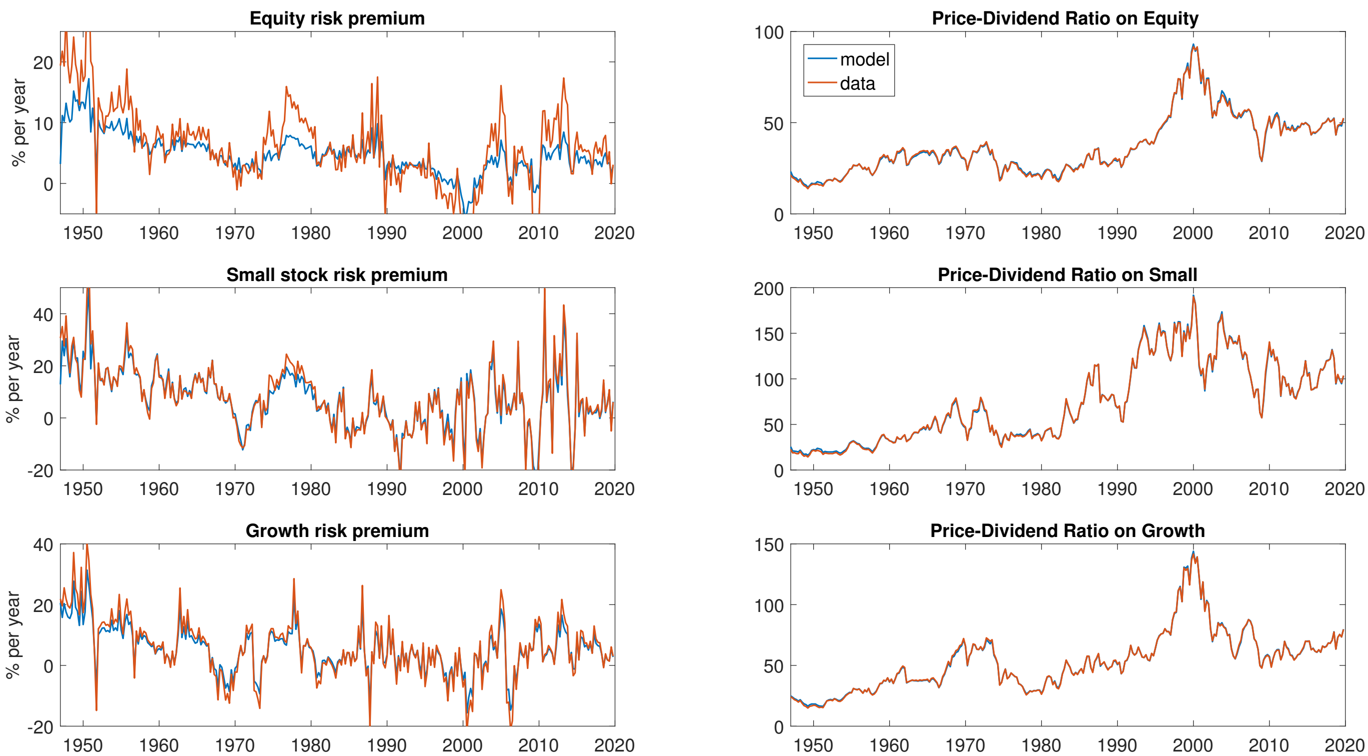

Note: The figure plots the observed and model-implied equity risk premium on the overall stock market, small stocks, and growth stocks in the left panels, as well as the corresponding price-dividend ratio in the right panels. The model is the blue line, the data are the red line.

A.2.3 Pricing Claims to Aggregate Consumption and Labor Income

Shocks to the growth rate in consumption/GDP (labor income/GDP) ratio are priced only to the extent that they are correlated with other priced sources of risk. The innovation to the change in the consumption/GDP (labor income/GDP) ratio that is orthogonal to all prior shocks is not priced. Since consumption/GDP growth and labor income/GDP growth appear last in the VAR and the model includes many sources of priced aggregate risk, those innovations are as small as possible.

Figure A6 plots the annual price-dividend ratios on the claims to GDP, aggregate consumption, and aggregate labor income. It contrasts these valuation ratios to those for the aggregate stock market, and housing wealth. The valuation ratios of GDP, aggregate consumption, and aggregate labor income claims are all highly correlated. They are high at the start of the sample, low in the early 1980s, and high at the end of the sample. Since total wealth is a claim to aggregate consumption, this suggests that expected returns on total wealth were highest in the early 1980s and have been falling ever since.

Figure A5: Equity Risk Premia and Price-Dividend Ratios (2/2)

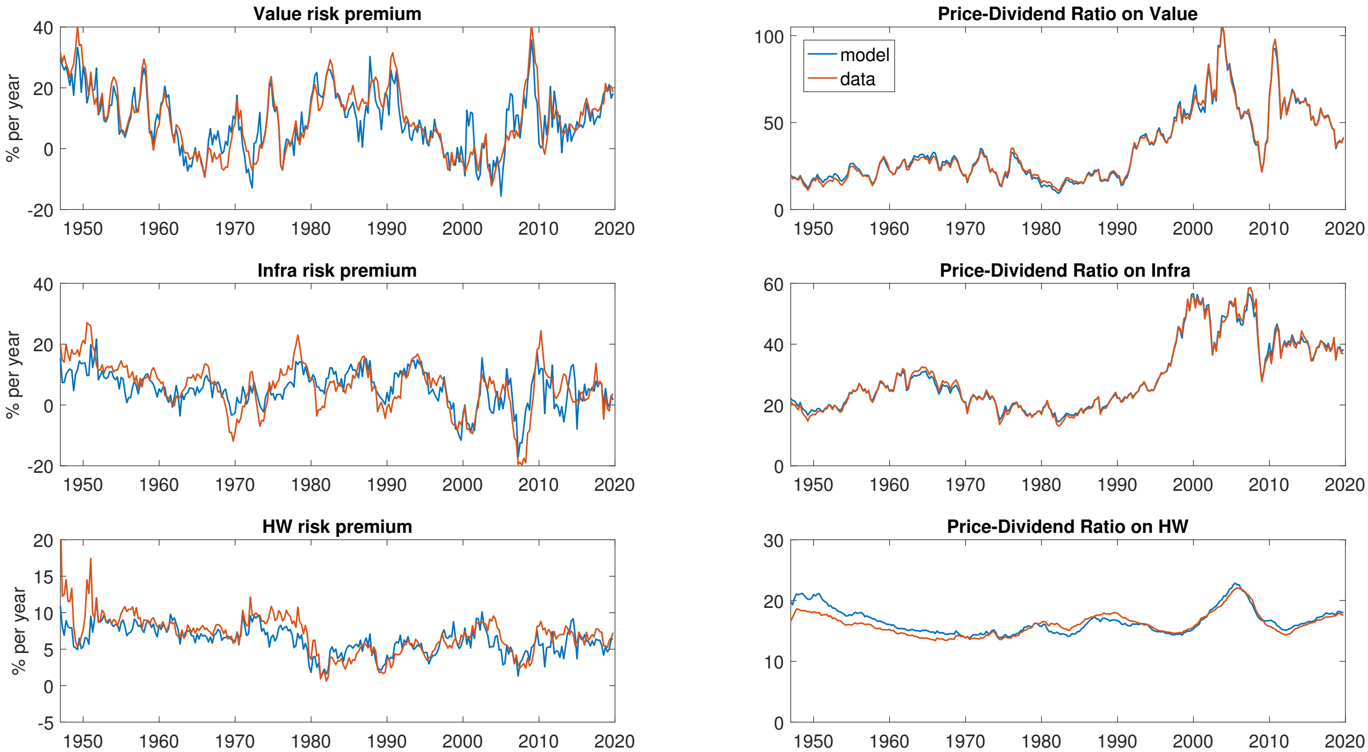

Note: The figure plots the observed and model-implied equity risk premium on value stocks, infrastructure stocks, and housing wealth in the left panels, as well as the corresponding price-dividend ratio in the right panels. The model is the blue line, the data are the red line.

A.2.4 Cash-flow Duration

The (McCauley) duration is the weighted average time for an investor to receive cash flows. For the aggregate stock market, this measure is computed as follows:

$$D_t^{CF,m} = \sum_{\tau=1}^{\infty} w_{t,\tau}\tau, \qquad w_{t,h} = \frac{P_{t,\tau}^{d}}{\frac{P_t^m}{D_t^m}} = \frac{\exp\left(A_\tau^m + B_\tau^{m\prime} z_t\right)}{\exp\left(\overline{pd} + e_{pd^m}' z_t\right)}$$

where $P_{t,\tau}^{d}$ is the price-dividend ratio of a τ-period dividend strip. Since durations are usually expressed in years while time runs in quarters in our model, we divide by 4. Duration is defined analogously for the other four equity indices, housing wealth, and for the GDP, consumption, and labor income claims. Note that for a nominal or real zero-coupon bond of maturity τ, $D_t^{CF} = \tau$.

Figure A7 The figure plots the model-implied time series of cash-flow durations on the overall stock market, small stocks, growth stocks, value stocks, infrastructure stocks, housing wealth, the GDP claim, the aggregate consumption claim, and the aggregate labor income claim. Durations tend to be positively correlated with the price-dividend ratios: high at the start of the sample,

Figure A6: Valuation Ratios

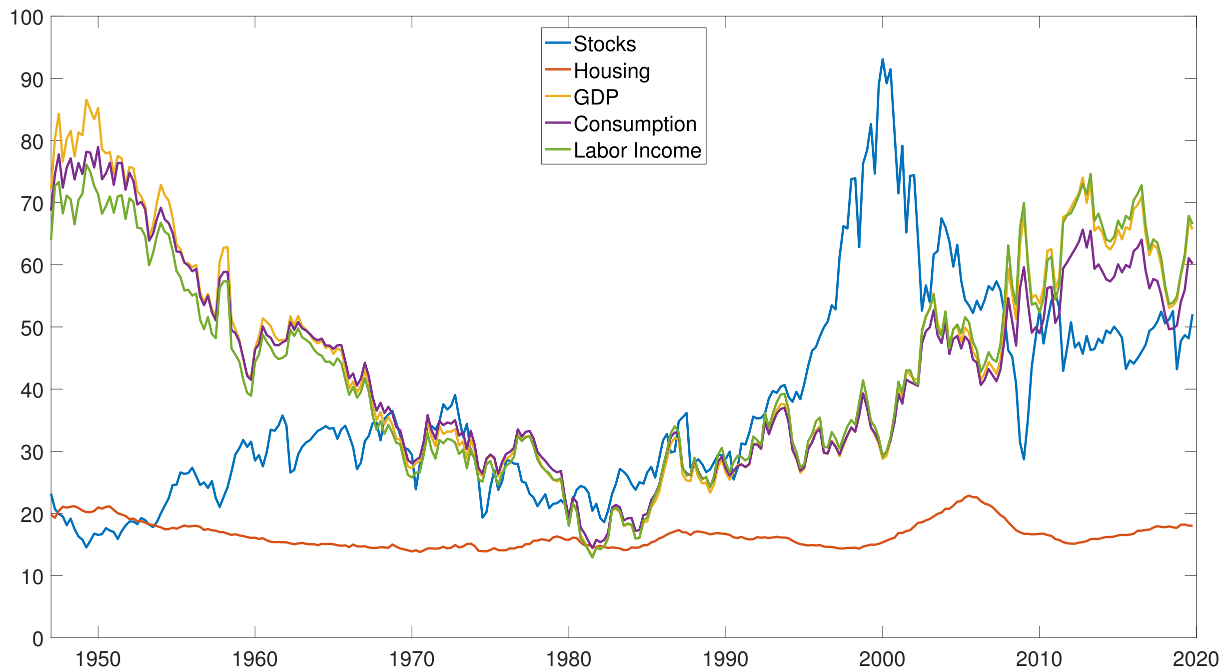

Note: The figure plots the annual price-dividend ratios on the aggregate stock market, housing wealth, and on claims to GDP, aggregate consumption, and aggregate labor income.

lowest in the early 1980s, and high at the end of the sample. The duration of housing wealth is highest during the housing boom in 2003–2007 when the valuation ratio of housing peaks. It then falls sharply in the housing bust before rising again in the housing boom that starts in 2013.

A.2.5 Market Price of Risk Estimates

The market prices of risk are pinned down by the moments discussed in the main text. Here we report and discuss the point estimates. Note that the prices of risk are associated with the orthogonal VAR innovations $\varepsilon \sim \mathcal{N}(0, I)$. Therefore, their magnitudes can be interpreted as (quarterly) Sharpe ratios. The constant in the market price of risk estimate $\widehat{\Lambda_0}$ is:

0.11	0.00	-0.36	0.06	0.00	0.43	0.00	-0.01	0.00	0.12	0.00	0.25	0.00	0.26	0.00	2.76	0.00	0.00	0.00	0.00

The matrix that governs the time variation in the market price of risk is estimated to be $\widehat{\Lambda_1} =$:

Figure A7: Cash-Flow Duration

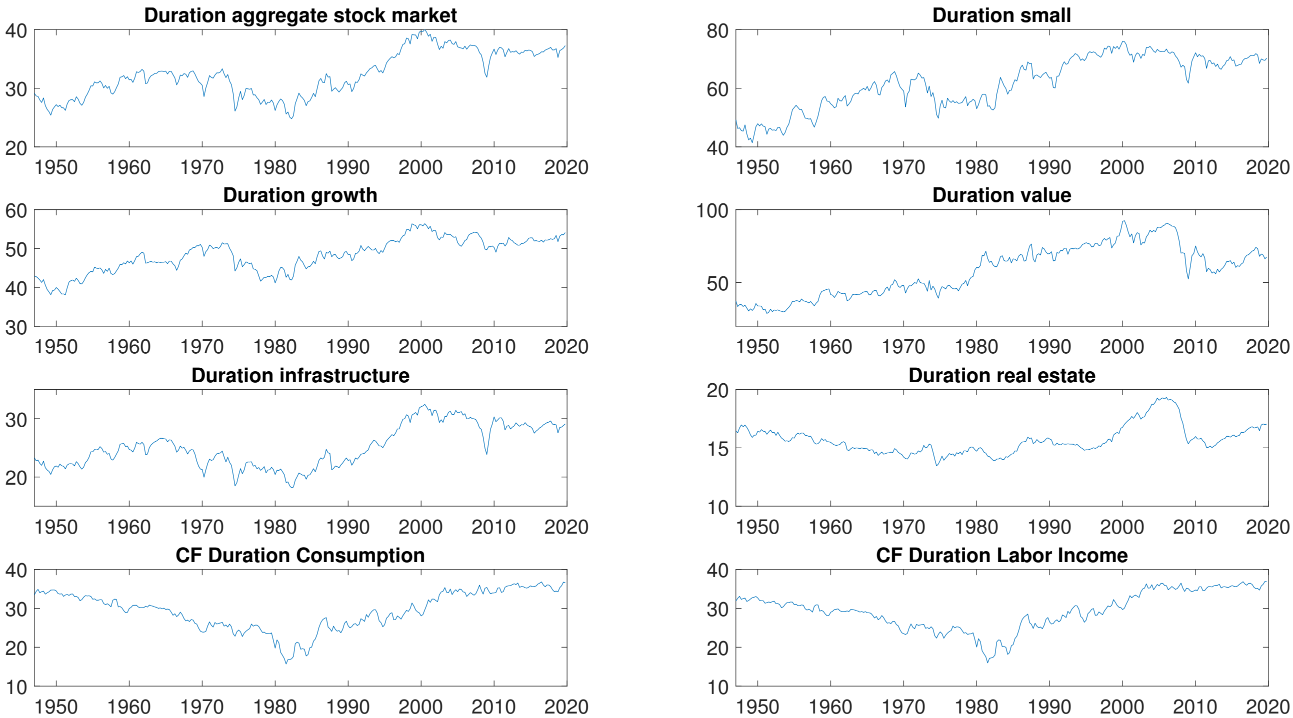

Note: The figure plots the model-implied time series of cash-flow durations on the overall stock market, small stocks, growth stocks, value stocks, infrastructure stocks, housing wealth, the GDP claim, the aggregate consumption claim, and the aggregate labor income claim. The duration is expressed in years.

0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
19.3	16.5	-31.7	-250.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
37.5	15.1	0.0	3.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	2.2	19.5	0.9	0.9
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
23.3	4.0	-29.7	-160.0	-0.9	12.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	3.6	0.1	1.1
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
-5.0	-1.6	1.5	-21.1	0.6	2.0	-0.5	8.8	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	2.9	-4.0	0.2	0.7
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
1.1	34.3	-23.7	14.0	0.9	-11.2	-0.1	1.8	-1.1	16.8	0.0	0.0	0.0	0.0	0.0	0.0	0.8	1.2	-0.1	-0.7
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	34.7	-11.1	34.1	0.8	-15.6	0.1	-1.1	-0.3	0.5	-2.0	0.6	0.0	0.0	0.0	0.0	-1.5	1.3	17.6	0.1
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
10.1	46.0	-40.9	-20.8	7.6	-10.9	1.0	-2.4	-5.9	-2.3	0.5	1.6	-4.4	0.1	0.0	0.0	3.3	3.7	-5.2	-2.3
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
8.6	27.2	-113.3	66.6	5.7	1.3	-2.2	-5.7	-0.4	-3.0	0.8	0.2	-3.3	0.0	1.3	0.4	0.1	-0.2	-15.0	-6.1
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

B Inequality and Interest Rate Data

B.1 Wealth Inequality

Our primary source of data for the top wealth shares presented in Figure 2 is the World Inequality Database maintained by WID team. For the US, we also report survey estimates of top wealth shares from the Survey of Consumer Finances (SCF) and the SCF+, the database developed by Kuhn, Schularick, and Steins (2020). We use data from the SCF+ from 1950 to 1983. We use SCF data from 1989 to 2019. Appendix C provides detailed information on SCF and SCF+. Figure B1 plots the net wealth shares held by the top-10% and top-1% as well as the gini coefficient for financial wealth. Our definition of net wealth calculated from the SCF and SCF+, used in Figure B1 and detailed in Appendix C, excludes the variable "Other non-financial assets".

B.2 Interest Rates

We construct the price of a real 30 year annuity by estimating the historical real yield curve for each country. Letting $y^r_t(h)$ denote the real yield at maturity h at time t the cost of the annuity is calculated as:

$$\sum_{h=1}^{30} \frac{1}{(1+y^r_t(h))^h} \ .$$

Due to varying availability of data and for robustness, we use three different approaches to estimate the real yield curve that lead to broadly consistent estimates.

First, for the UK post 1985 we use historical time series of real yields of various maturities available from the Bank of England. We fit a spline through these points and construct the real yield curve directly.

Second, for the U.S. and France we use the time series of historical nominal yields and inflation provided by Global Financial Data, augmented with data from the Macrohistory database constructed by Jordà, Schularick, and Taylor (2017), to estimate real yields at different maturities and then fit a spline through the estimated real yields to construct the real yield curve. We construct real yields for each year by estimating an AR(1) process for inflation on a rolling sample of 50 years of past data, and then subtracting forecasted inflation from nominal yields at all available maturities. Those are 3-month treasury yields and 10-year government bond yields for all periods, as well as 30-year government bond yields for later years.

Third, for the U.K. and U.S. we also use model estimates of the real yield curve. The U.S. estimates are from the model in Section A. The U.K. estimates are from a similar model estimated for the U.K. in Jiang, Lustig, Van Nieuwerburgh, and Xiaolan (2021).

Figure B1: US Wealth Inequality

(a) Top 10% Share

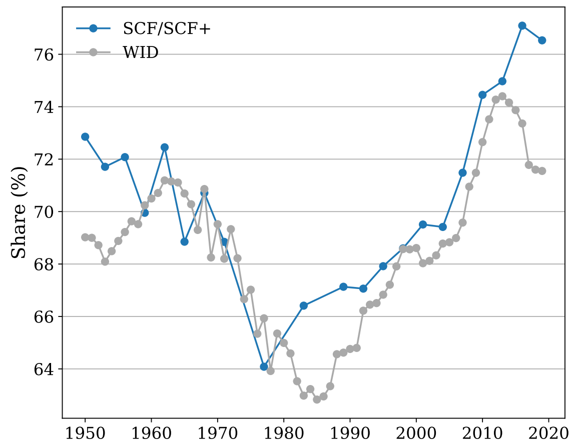

(b) Top 1% Share

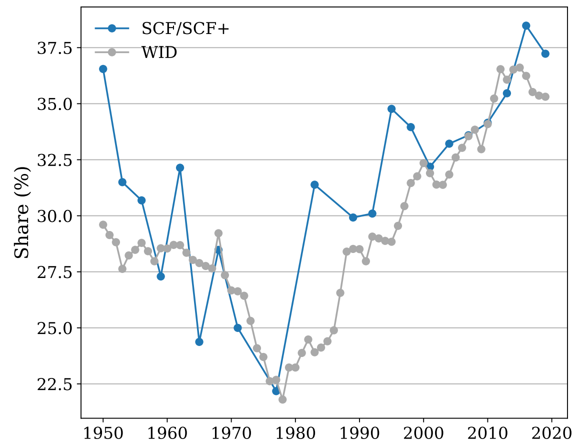

(c) Gini Coefficients

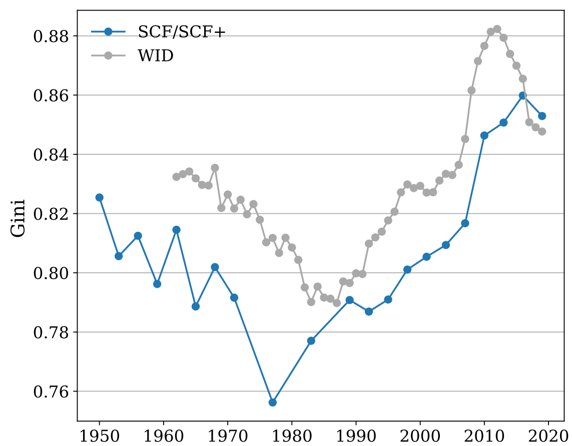

Note: Data are based on WID, SCF as well as the SCF+ database developed by Kuhn, Schularick, and Steins (2020). Figure B1a plots the shares of net wealth held by households in the top 10% net wealth percentile. Figure B1b plots the shares of net wealth held by households in the top 1% net wealth percentile. Figure B1c plots the Gini coefficients for net wealth.

C Household Wealth and Portfolio Shares

In this section, we detail the variables used in calculating household net wealth and portfolio shares. Our data encompasses the 1950-2019 time range, based on Survey of Consumer Finance (SCF) waves from 1989 through 2019 and data from the SCF+ database, compiled by Kuhn, Schularick, and Steins (2020), for the period from 1950 through 1983. In sections C.1 and C.2, we discuss the specific variables used in our analysis for these two sources.

C.1 Survey of Consumer Finances (SCF)

The SCF is a statistical survey of the balance sheet, pension, income and other demographic characteristics of households in the United States. We use data from the Summary Extract Data—that is, the extract data set of summary variables used in the Federal Reserve Bulletin[28]—as well as granular information from the SCF original surveys. We construct the following variables for our analysis.

Total Financial Assets. SCF variable *FIN*. This includes: All types of transaction accounts (liquid assets), Certificates of deposit, Directly held pooled investment funds (excl. money market funds), Savings bonds, Directly held stocks, Directly held bonds (excl. bond funds savings bonds), Cash value of whole life insurance, Other managed assets, Quasi-liquid retirement accounts, Other misc. financial assets.

Equities (direct and indirect). SCF variables *EQUITY*. Total value of financial assets held by household that are invested in stock. That includes: directly-held stock, Stock mutual funds, RAs/Keoghs invested in stock, Other managed assets with equity interest, Thrift-type retirement accounts invested in stock.

Real Estate. SCF variables *HOUSES* + *ORESRE* + *NNRESRE*. The real estate variable includes: Primary residence; Residential property excluding primary residence (e.g., vacation homes); Net equity in non-residential real estate.

Private Business Wealth. SCF variable *BUS*. Businesses with either an active or passive interest.

Vehicles. SCF variable *VEHIC*. Value of all vehicles.

Cash & Deposits. SCF variables *LIQ* + *CDS*. This includes all types of transaction account (Money market accounts, Checking accounts, Savings accounts, Call accounts, Prepaid cards) and certificated of deposits.

Fixed Income. SCF variable *FIN* minus Cash & Deposits and Equities. Fixed income is calculated as the residual of total financial assets minus Cash & Deposits and Equity (direct and indirect).

Mortgage Debt. SCF variables *MRTHEL* + *RESDBT*. This includes: Debt secured by primary residence (mortgages, home equity loans, HELOCs); Debt secured by other residential property.

Student Debt. SCF variable *EDN_INST*. Total value of education loans held by household. This includes education loans that are currently in deferment and loans in scheduled repayment period.

Vehicles Debt. SCF variable *VEH_INST*. Total value of vehicle loans held by household.

Other Debt. SCF variable *DEBT* minus Mortgage Debt, Student Debt, and Vehicles Debt. This includes: Other lines of credit (not secured by resid. real estate); Credit card balances after last payment; Other installments other than vehicles debt and student debt.

[28]The SCF Flow Chart provides information on how variables are constructed https://www.federalreserve.gov/econres/files/networth%20flowchart.pdf. The code on how different variables in the Summary Extract Data are constructed can be found here: https://www.federalreserve.gov/econres/files/bulletin.macro.txt

C.1.1 Private Business Wealth Corporate and Non-Corporate

The SCF presents a wealth of granular information on individual businesses that collectively constitute Private Business Wealth (PBW). We use this information to split the total value of Private Business Wealth into the Corporate component and the Non-Corporate component.

Private Business Wealth Corporate. It includes the value of all businesses reported as Subchapter-S corporations or all other types of corporations (both actively and passively managed). These categories are reported as Corporate Business in the Financial Accounts of the United States. Figure C1a plots the shares of Corporate PBW as a percentage of total Private Business Wealth in 1989, the first SCF survey wave we use in our sample. Figure C1b plots the shares in 2019, the last available SCF survey wave.

Private Business Wealth Non-Corporate. It includes all businesses which are not corporations (both actively and passively-managed), namely limited partnership, other partnership, LLCs, Co-operative, sole-proprietorships , and other types. These categories are reported as Non-Corporate Business in the Financial Accounts of the United States.

We use the following methodology to construct the value of Private Business Wealth Corporate and Non-Corporate. For each actively-managed business, we follow the methodology used by the SCF to measure the overall total value of private business wealth. From 1989 to 2010, the SCF provides information on the three largest active businesses, while from 2010 onward it provides information on the two largest active businesses. We now provide further details on the data construction. We label the counter of the SCF variable used for each of the business where the information is available: #1 for the first business, #2 for the second business, #3 for the third business. For each active business, we measure its value as: the net worth of the share of the business (#1 X3229, #2 X3229, #3 X3329), plus the amount the business owes the households (#1 X3124, #2 X3224, #3 X3324), minus the amount the household owes the business (# 1 X3126, #2 X3226, #3 X3326), plus the amount of loans that are collateralized/guaranteed (#1 X3121, #2 X3221, #3 X3321). After having computed the value of each business, we then use information on the types of business entities and legal structure (#1 X3119, #2 X3219, #3 X3319) to split business wealth into a Corporate and a Non-Corporate component.

For passively-held businesses, we directly observe the value of: limited partnership (X3408), other partnership (X3412), S-Corporations (X3416), other corporations (X3420), sole-proprietorships (X3424), LLCs (X3452), other types (X3428). We use this information on the value of the business and the legal structure to split passively-held businesses into the Corporate and Non-Corporate.

C.2 SCF+

We use the SCF+ database developed by Kuhn, Schularick, and Steins (2020) from 1950 through 1983. In this section, we outline the list of variables utilized from the SCF+ database to calculate

Figure C1: Ratio of Corporate PBW to Private Business Wealth by Net Wealth percentile groups

(a) 1989 (b) 2019

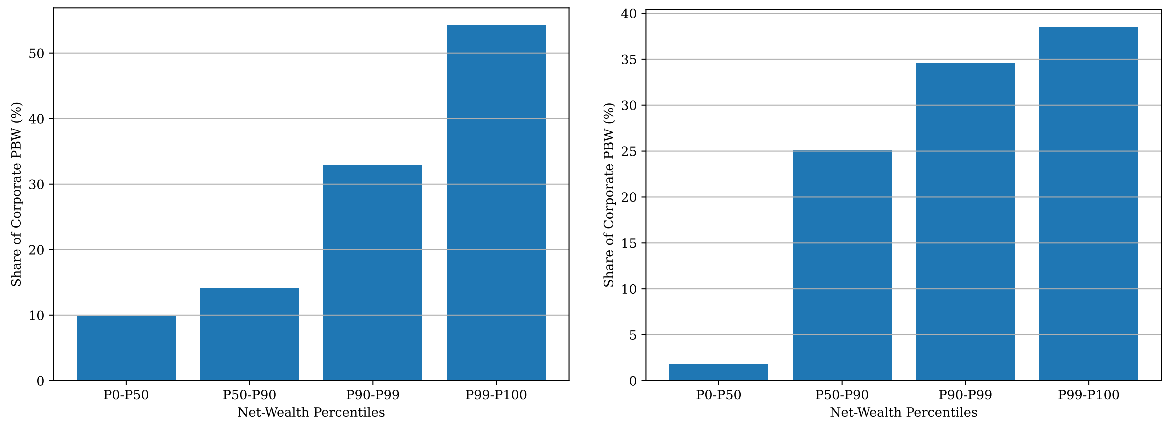

Note: The Figure plots the ratio of Corporate PBW to total Private Business Wealth for different net wealth percentile groups in 1989 and 2019. Initially, we determine the division between Corporate PBW and Non-Corporate PBW for each individual household. Subsequently, we categorize households into four groups based on their net wealth ranges (P0-P50, P50-P90, P90-P99, P99-P100). Finally, we calculate the total dollar amount of Corporate PBW and Private Business Wealth for all households within each net wealth percentile group and derive the corresponding ratios. Figure C1a plots the shares in 1989, the first SCF survey wave we use in our sample. Figure C1b plots the shares in 2019, the last available SCF survey wave.

households' net wealth and portfolio shares.

Equities. SCF+ variable *ffaequ*. This includes equity and other managed assets. We also add indirect holdings of equities through mutual funds and pension funds (see Appendix C.2.1).
Real Estate. This includes: SCF+ variable *house*, asset value of house; SCF+ variable *oest*, other real estate (net position); SCF+ variable *hoestdebt*, other real estate debt (note: we add back the debt to the other real estate net position).
Private Business Wealth. SCF+ variable *ffabus*, business wealth.
Vehicles. SCF+ variable *vehi*.
Cash and Deposits. SCF+ variable *liqcer*, liquid assets and certificates of deposit.
Fixed Income. SCF+ variable *ffafin*, financial assets minus Cash and Deposits and Equities.
Mortgage Debt. This includes: SCF+ variable *hdebt*, housing debt on owner-occupied real estate; SCF+ variable *oestdebt*, other real estate debt.
Vehicle Debt. Information on vehicle debt is not available from SCF+. We infer vehicle debt using the procedure described in Section C.2.1.
Student Debt. Information on student debt is not available from SCF+. We infer student debt using the procedure described in Section C.2.1.
Total Other Debt. *pdebt*, personal debt.

C.2.1 Adjustments to SCF+ data

Indirect Equity holdings The SCF+ database does not provide information on indirect holdings of equities. We follow the methodology of Leombroni, Piazzesi, Schneider, and Rogers (2020) to compute indirect holdings exploiting aggregate data from the US Financial Accounts of the United States. Households hold indirect exposure to equities through mutual funds shares and pension accounts.

We assume that the indirect exposure to equities through mutual fund shares reflects the allocation to equities within the aggregate mutual fund sector. To calculate this equity exposure, we use data from the Financial Accounts of the united States. We compute the ratio of mutual funds' holdings of Corporate Equities (FAUS LM653064100) divided by the total amount of mutual fund assets (FAUS LM654090000). This ratio allows us to allocate a portion of households' holdings of mutual fund shares to equities. For instance, Corporate Equities accounted for 70% of mutual funds' assets in 1983 and hence we attribute 70% of households' mutual fund holdings to equities.

We perform a similar procedure to determine the indirect exposures to equities through households' pension accounts. Initially, we calculate the portion of assets allocated to equities by the DC pension funds sector using data from the Financial Accounts of the United States. A significant part of DC pension funds' assets is allocated to mutual fund shares (FAUS LM573064255). Hence, for the pension funds sector, we consider both this indirect exposure to equities (using the aforementioned procedure for the mutual fund sector) and the direct exposure of DC pension funds to Corporate Equities (FAUS LM573064133).

Subsequently, we compute the percentage of equities holdings relative to the total assets of DC pension funds (FAUS FL574090055). Utilizing this ratio, we allocate a portion of households' pension accounts to equities. For instance, in 1983, equities holdings represented 40% of DC pension funds' assets. Accordingly, we allocate 40% of households' pension accounts to equities.

Private Business Wealth Corporate and Non-Corporate. The SCF+ database provides only the total households' holdings of private business wealth, without distinguishing between the Corporate and Non-Corporate. To estimate this split, we rely on data from the 1989 SCF survey. We use the calculated ratio of Corporate PBW to total Private Business Wealth for different net wealth percentile groups in the 1989 SCF survey (see Figure C1a) to allocate Private Business Wealth into a Corporate and Non-Corporate components in the SCF+ data.

Student Debt and Vehicle Debt. We encounter a lack of data for the variables Student Debt and Vehicle Debt within the SCF+ database. To overcome this limitation, we apply a methodology similar to the one used for inferring Private Business Wealth. By leveraging data from the 1989 SCF survey, we calculate the ratio of Student Debt and Vehicle Debt to the Total Other Debt, represented as the sum of Student Debt, Vehicle Debt, and all other types of debt. This ratio is computed for different percentile groups (P0-P50, P50-P90, P90-P99, P99-P100) in 1989. As we only have access to a composite variable called Total Other Debt in the SCF+ dataset, which includes Student Debt, Vehicle Debt, and other forms of debt, we utilize the shares obtained from the 1989 survey to allocate the components of Student Debt and Vehicle Debt from the overall category of Other Debt.

C.3 Portfolio Shares by Net Wealth Percentile

To compute the household's portfolio share in each asset, we divide the dollar holdings in the asset (or liability) by the household's net wealth. Households hold significantly different portfolios, depending on their net wealth. Figure C2 plots wealth-weighted portfolio shares for households of different net wealth percentiles. We compute the wealth-weighted shares for each survey and the average across all years from 1983 to 2019. Finally, we re-scale the weights to sum to one. Figure C2a plots the portfolio shares in Equities, Non-Corporate PBW, and Corporate PBW. Figure C2b plots the portfolio shares in Real Estate and Vehicles. Figure C2c plots the portfolio shares in Fixed Income and Cash and Deposits. Figure C2d plots the portfolio shares in Mortgage Debt, Vehicle Debt, Student Debt and Other Debt. Liabilities are reported as a negative number.

Figure C3 shows the median portfolio shares within each net wealth percentile bin. It then averages across all waves and re-scales to make portfolio weights sum to one.

Figure C2: Wealth-Weighted Portfolio Shares by Net Wealth Percentiles

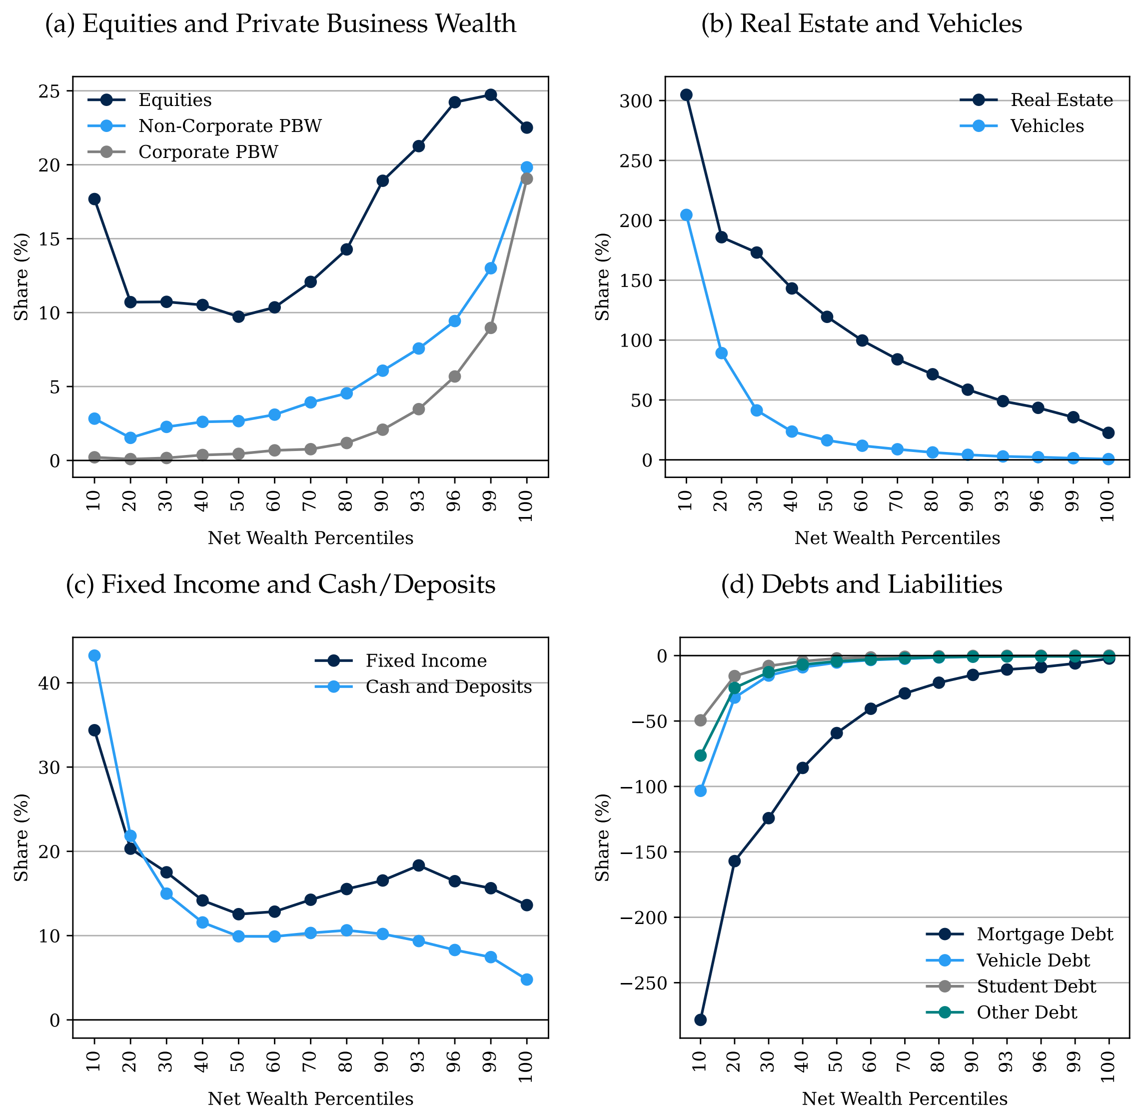

Note: Figure C2 plots wealth-weighted portfolio shares for households of different net wealth percentiles. We compute the wealth-weighted shares for each survey and the average across all years from 1983 to 2019. Finally, we re-scale the weights to sum to one. The sample runs from 1983 to 2019.

Figure C3: Median Portfolio Shares by Net Wealth Percentiles

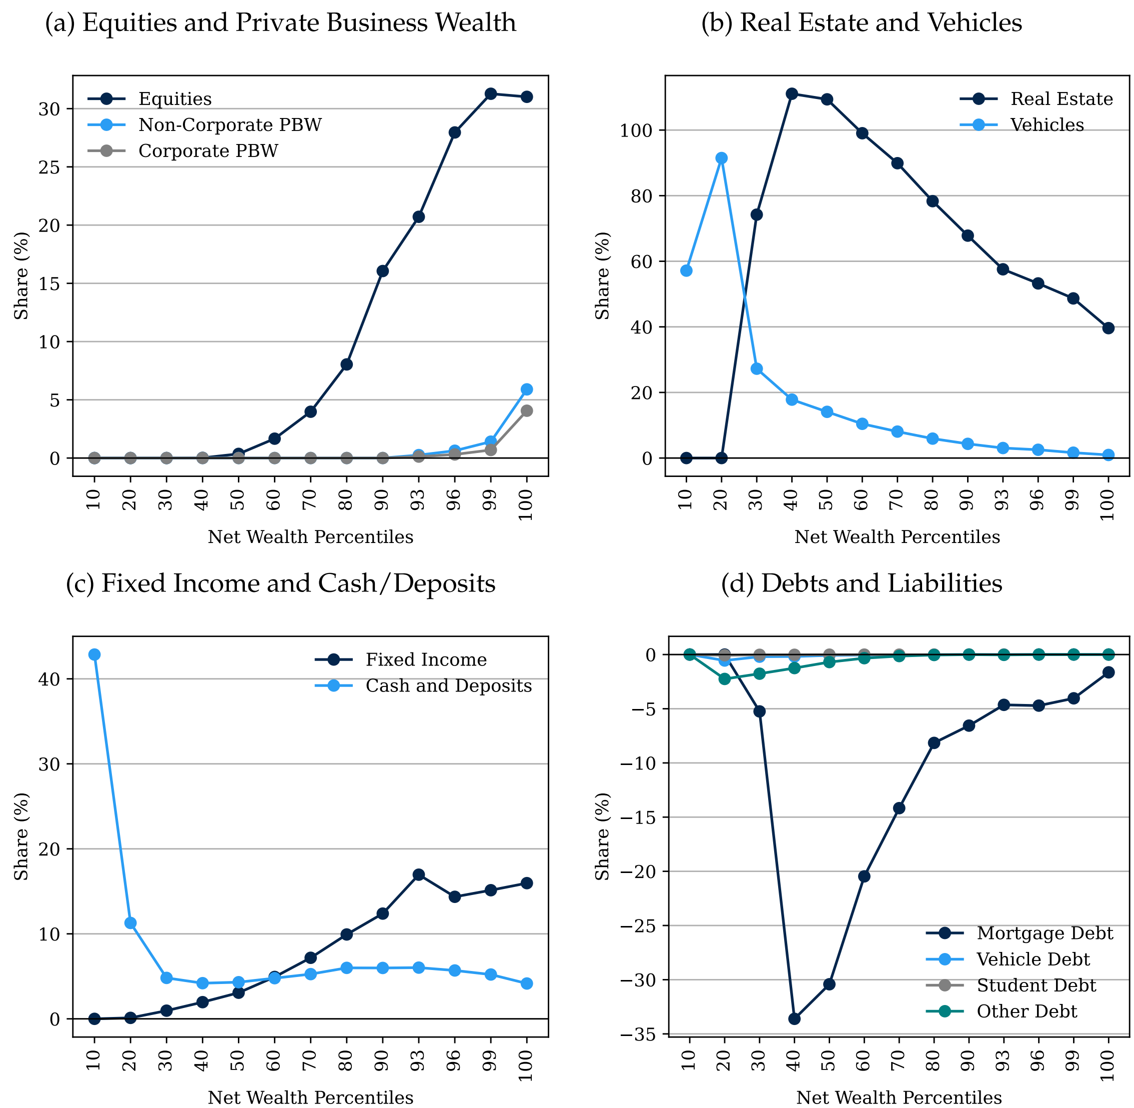

Note: Figure C3 shows the median portfolio shares within each net wealth percentile bin. It then averages across all years and re-scale to make portfolio weights summing to one. The sample runs from 1983 to 2019.

C.4 Portfolio Shares by Age

Figure C4 displays portfolio shares by age, using survey data from 1983 until 2019. Households are categorized into different age groups: 25-35, 35-45, 45-55, 55-65, 65-75, and 75-85 years old. Figure C4a plots value-weighted portfolio shares, while Figure C4b plots median portfolio shares. The shares are computed for each survey wave and demographic group, averaged across years, and subsequently rescaled to ensure they sum to 100%.

Figure C4: Portfolio Shares by Cohorts

(a) Wealth-Weighted Portfolio Shares

(b) Median Portfolio Shares

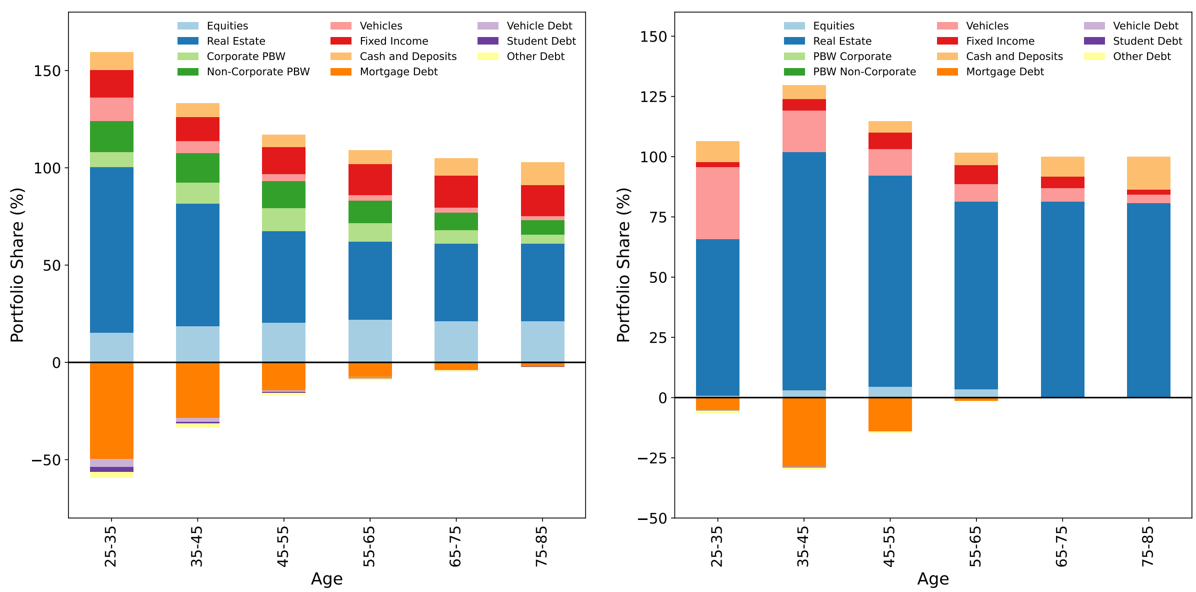

Note: Portfolio shares by age group, averaged across years. Households are categorized into different cohort groups: 25-35, 35-45, 45-55, 55-65, 65-75, and 75-85. Figure C4a utilizes value-weighted portfolio shares. Figure C4b plots the median portfolio share in each asset category. The shares are computed for each survey wave and demographic group, averaged across years, and subsequently rescaled to ensure they sum to 100%.

D Duration Measurement

D.1 Duration Formula

For risky assets, we use Gordon's growth model to estimate the duration. The formula for the duration is defined in equation (7). The derivation of this formula is as follows. Define $k = (1+g)/(1+r)$ where r denotes the expected return and g the expected cash-flows growth rate. The present value of cash flows is:

$$\begin{aligned} P_0 =& Div_0 \frac{(1+g)}{(1+r)} + Div_0 \frac{(1+g)^2}{(1+r)^2} + ... = Div_0\, k + Div_0\, k^2 + ... = Div_0 \left(\sum_{t=1}^{\infty} k^t \right) \\ &= Div_0 \frac{k}{1-k} = Div_0 \frac{(1+g)}{r-g}, \end{aligned} \tag{29}$$

while time-weighted present value of cash flows is:

$$\begin{aligned} PVT_0 &= Div_0(k + 2k^2 + 3k^3 + ...) = Div_0\, k(1 + 2k + 3k^2 ...) = Div_0\, k \left(\sum_{j=0}^{\infty} \sum_{t=j}^{\infty} k^t \right) \\ &= Div_0\, k \left(\sum_{j=0}^{\infty} k^j \sum_{t=j}^{\infty} k^{t-j} \right) = Div_0\, k \left(\sum_{j=0}^{\infty} k^j \sum_{t=0}^{\infty} k^t \right) = Div_0 \frac{k}{1-k} \left(\sum_{j=0}^{\infty} k^j \right) \\ &= Div_0 \frac{k}{(1-k)^2}. \end{aligned}$$

The duration is therefore:

$$D^{GGM} = \frac{PVT_0}{P_0} = \frac{1}{1-k} = \frac{1+r}{r-g}\,. \tag{30}$$

Using (29), we also find that:

$$\frac{1}{r-g} = \frac{P_0}{Div_0(1+g)} \tag{31}$$

which implies

$$\frac{P_0}{Div_0} = \frac{1+g}{r-g}. \tag{32}$$

Manipulating (32), we obtain

$$1 + \frac{P_0}{Div_0} = \frac{1+r}{r-g} = D^{GGM}. \tag{33}$$

We use the observed price-dividend ratio at each point in time t, P_t/Div_t to compute a time series for duration.

D.1.1 Two-Stage GGM

We use a two-stage Gordon's growth model for the duration of Corporate Private Business Wealth. Denote the first stage's dividend growth rate and required return by g and r, respectively. In the second stage, the firm's cash flows grow at a rate of $\tilde{g}$ and the required return is $\tilde{r}$. Let n be the length of the first stage (in years). Define $k = (1+g)/(1+r)$ and $\tilde{k} = (1+\tilde{g})/(1+\tilde{k})$. The duration of the firm's cash flows in the two-stage GGM is given by the formula:

$$D^{2SGGM} = \frac{\frac{k}{(1-k)^2}\left(1-(n+1)k^n + nk^{n+1}\right) + k^n \frac{\tilde{k}}{(1-\tilde{k})^2}\left(n(1-\tilde{k})+1\right)}{k\frac{1-k^n}{1-k} + k^n \frac{\tilde{k}}{1-\tilde{k}}}. \tag{34}$$

The derivation of equation (34) is as follows. The present value of the firm's cash flows is:

$$\begin{aligned} P_0 &= \frac{Div_0(1+g)}{(1+r)} + \frac{Div_0(1+g)^2}{(1+r)^2} + ... + \frac{Div_0(1+g)^{(n)}}{(1+r)^n} \\ &+ \frac{Div_0(1+g)^n}{(1+r)^n}\frac{(1+\tilde{g})}{(1+\tilde{r})} + \frac{Div_0(1+g)^n}{(1+r)^n}\frac{(1+\tilde{g})^2}{(1+\tilde{r})^2} + ... \\ &= Div_0\left[(k+k^2+...+k^n) + k^n(\tilde{k}+\tilde{k}^2+...)\right] \\ &= Div_0\left[k\frac{1-k^n}{1-k} + k^n\frac{\tilde{k}}{1-\tilde{k}}\right]. \end{aligned} \tag{35}$$

while the time-weighted present value of cash flow is:

$$\begin{aligned} PVT = Div_0\, k + 2Div_0\, k^2 + 3Div_0\, k^3 + ... + nDiv_0\, k^n + \\ + (n+1)Div_0\, k^n\tilde{k} + (n+2)Div_0\, k^n\tilde{k}^2 + ... \end{aligned}$$

We split PVT into the first stage PVT^{1st} and second stage PVT^{2nd}:

$$\begin{aligned} PVT^{1st} &= Div_0\, k + 2Div_0\, k^2 + 3Div_0\, k^3 + ... + nDiv_0\, k^n \\ &= Div_0\, k\left(\sum_{j=0}^{n-1}\sum_{t=j}^{n-1} k^t\right) \\ &= Div_0\, k\left(\sum_{j=0}^{n-1} k^j \sum_{t=j}^{n-1} k^{t-j}\right) \\ &= Div_0\, k\left(\sum_{j=0}^{n-1} k^j \sum_{t=0}^{n-j-1} k^t\right) \\ &= Div_0\, k\left(\sum_{j=0}^{n-1} k^j \frac{1-k^{n-j}}{1-k}\right) \end{aligned}$$

$$= \frac{Div_0\, k}{1-k}\left(\frac{1-k^n}{1-k} - nk^n\right)$$
$$= \frac{Div_0\, k}{(1-k)^2}\left(1-(n+1)k^n + nk^{n+1}\right)$$

$$\begin{aligned} PVT^{2nd} &= (n+1)Div_0\, k^n\tilde{k} + (n+2)Div_0\, k^n\tilde{k}^2 + ... \\ &= nDiv_0\, k^n\tilde{k}\left(1+\tilde{k}+\tilde{k}^2+...\right) + Div_0\, k^n\tilde{k}\left(1+2\tilde{k}+3\tilde{k}^2...\right) \\ &= nDiv_0\, k^n\frac{\tilde{k}}{1-\tilde{k}} + Div_0\, k^n\frac{\tilde{k}}{(1-\tilde{k})^2} \\ &= Div_0\, k^n\frac{\tilde{k}}{(1-\tilde{k})^2}\left(n(1-\tilde{k})+1\right). \end{aligned}$$

Finally, the duration is calculated as:

$$D^{2SGGM} = \frac{PVT_0}{P_0} = \frac{\frac{k}{(1-k)^2}\left(1-(n+1)k^n + nk^{n+1}\right) + k^n\frac{\tilde{k}}{(1-\tilde{k})^2}\left(n(1-\tilde{k})+1\right)}{k\frac{1-k^n}{1-k} + k^n\frac{\tilde{k}}{1-\tilde{k}}}.$$

D.2 Duration of Equities

We use data on the value-weighted market portfolio of publicly listed companies from CRSP for the duration of the equity market. Using annual dividends (Div_t^{eq}) and end-of-year prices (P_t^{eq}) we determine the duration based on equation (7). Figure D2a plots the time series of these duration estimates.

For robustness, we consider four alternative ways of measuring equity duration. The first three alternatives use the same approach, but different data. As a first alternative, we estimate the duration using Shiller's S&P500 historical data. We obtain the price-dividend ratio from the CAPE price-earnings ratio by assuming a dividend-earnings ratio of 0.5, which equals the historical average. The result is shown in Column (2) of Table D1. As a second alternative, we use, the equity price and dividend series from the Jordà-Schularick-Taylor Macrohistory Database (JST). The results are again similar and displayed in Column (3) of Table D1. As a third alternative, we compute the duration using Financial Accounts of the United States data. For the price, we sum up Corporate Equities of Non-Financial Corporate Business (LM103164103) and Foreign Direct Investment in U.S. Non-financial Corporate Business (LM103192105). To measure the dividend series, we utilize Dividends Paid by Non-Financial Corporate Business (FA106121001). The results are displayed in Column (4) of Table D1. These numbers are somewhat lower than the other three series. As a fourth alternative, we can compute the duration based on the affine model from

Appendix A. That model does not use the GGM but rather has growth rates and valuation ratios that depend on the state variables in the VAR. The VAR-SDF model generates an average equity duration of 34.2, and follows a similar pattern as the baseline series over time.

Table D1: Duration of Equities and Housing

	Equities				Real Estate	
	Baseline	S&P 500	JST	FAUS Corporate	Baseline	JST
	(1)	(2)	(3)	(4)	(5)	(6)
1980-2019	49.78	45.25	48.08	33.64	12.28	21.46
1980-1989	26.59	24.03	26.55	22.37	11.51	19.97
1990-1999	52.20	43.42	50.18	36.47	11.31	19.93
2000-2009	65.90	62.98	61.50	37.23	13.68	23.77
2010-2019	54.43	50.57	54.08	38.51	12.63	22.17

Note: The table reports the duration for equities and real estate. Column (1) to (4) show the duration for equities calculated using data from (1) CRSP, (2) S&P 500, (3) Jordà-Schularick-Taylor Macrohistory Database (JST) and (4) FAUS. Column (5) to (6) report the duration for real estate using data from (5) Zillow and (6) Jordà-Schularick-Taylor Macrohistory Database (JST).

D.3 Duration of Real Estate

Our baseline measure of real estate duration is based on Zillow data. The main advantage of using Zillow data to compute the price-to-rent ratio is that Zillow compares similar homes when it constructs rental rates and home prices, and adjusts transactions data for variation in characteristics (hedonics). As a measure of house prices P_t^{re}, we use the *Zillow Home Value Index (ZHVI)* for all houses in the U.S. This index measures the typical home value and market changes across regions and housing types. It reflects the typical value for homes in the 35th to 65th percentile range. As a measure of dividends (i.e., rents), Div_t^{re}, we use the *Zillow Observed Rent Index (ZORI)*. This index is a smoothed measure of the typical observed market rate rent across a region. ZORI is a repeat-rent index that is weighted to the rental housing stock to ensure representativeness across the entire market, not just those homes currently listed for rent. We use Zillow data starting in March 2015 since the Zillow ZORI index is only available since then. Prior to March 2015, we iterate the price/rent ratio backward using alternative data sources. We calculate the growth rate in prices P_t^{re}/P_{t-1}^{re} from January 1980 to March 2015 using the growth rate in the House Price Index from the U.S. Federal Housing Finance Agency. We compute the growth rate in dividends (rents), $Div_t^{re}/Div_{t-1}^{re}$ using the growth rate in the Shelter Consumer Price Index for All Urban Consumers from the U.S. Bureau of Labor Statistics. Using the level of prices and dividends in March 2015, we use the growth rate from 1980 to 2015 to reconstruct the full time series of P_t^{re} and Div_t^{re}. We then compute the duration of Real Estate assets using equation (30). The results are in Column (5) of Table D1. Figure D2b plots the time series of these duration estimates.

As an alternative measure, we also consider the prices and rents (dividends) from the Jordà-Schularick-Taylor Macrohistory Database (JST), with results reported in Table D1, column (6). As a second alternative, we compute real estate duration from the affine SDF-VAR model, which does not rely on the GGM assumptions. It delivers an average duration for 1980–2019 of 16.0, in between the baseline and the JST measure.

D.4 Duration of Private Business Wealth Non-Corporate

In Appendix C.1.1, we elaborated on our approach to splitting Private Business Wealth (PBW) into two components: Corporate and Non-Corporate PBW.

To estimate the duration of Non-Corporate PBW, we utilize data on Non-Corporate Business wealth from the Financial Accounts of the United States. Our measure of P^{pbwnc} is derived by summing Proprietors' Equity in Non-corporate Business (LM112090205) with Foreign Direct Investment in U.S. Non-Financial Non-Corporate Business (LM115114103). For the dividend series Div_t^{pbwnc}, we use Withdrawals from Income of Non-Financial Non-Corporate Business (FA116122001). This cash flow measure includes both labor and capital remuneration. Since we want to capture only the capital remuneration component, we split business income equally into labor and capital remuneration. Consequently, we multiply the series Withdrawals from Income of Non-Financial Non-Corporate Business by 0.5. Figure D2c illustrates the duration time series for the Non-Corporate component of Private Business Wealth.

Table D2: Duration of Non-Corporate PBW

	Non-Corporate PBW		
	Baseline	0.86 Labor Share	SCF
	(1)	(2)	(3)
1980-2019	16.33	57.36	69.51
1980-1989	22.08	78.48	75.17
1990-1999	15.43	54.04	65.55
2000-2009	14.71	51.40	71.26
2010-2019	13.11	45.51	66.06

Note: In this table we provide the duration estimates for PBW Short using different approaches. Column (1) displays the baseline estimates for PBW Short. In Column (2), we apply the 0.86 adjustment of labor share, as estimated by Quadrini and Rıos-Rull (1997). Additionally, Column (3) presents the duration estimated using the SCF, as described in Appendix D.6.

Quadrini and Rıos-Rull (1997) and Krueger and Perri (2006) adopt a labor income to business income ratio of approximately 0.86, implying that only 14% of business income is capital income. This naturally leads to a lower dividend on Non-Corporate PBW and a higher price-dividend ratio. Column (2) of Table D2 displays the duration estimates using this alternative labor income-to-business income ratio of 0.86. As a second robustness check, column (3) of Table D2 shows the duration calculated using data series from the SCF. Appendix D.6 explains the methodology based on SCF data in more detail. Both alternative measures result in higher estimates for Non-Corporate PBW, so that our benchmark measure in column (1) is conservative.

D.5 Duration of Private Business Wealth Corporate

The estimation of the duration for the Corporate component of PBW is conducted using the 2-stage Gordon Growth model in Appendix D.1.1. Private businesses in this category go through an initial stage of high growth, followed by a second, mature stage. Working backwards, in the second stage of the GGM, we assume that the growth rate ($\tilde{g}$) and the required rate of return ($\tilde{r}$) equal those of the value-weighted public equity market. We assume high-growth private businesses transition into this second stage after a period of twenty years in the first stage (i.e., $n = 20$).[29]

To establish the values of g and r for the first stage, our benchmark approach computes g_t based on dividends from small stocks, defined as the subset of companies within the first market capitalization decile of publicly-traded companies on NYSE-Amex-NASDAQ. The data are obtained from CRSP. Subsequently, we derive the implied required rate of return (r_t) by using the price-dividend ratio of small stocks. Specifically, we identify the r_t value at each point in time as the value that minimizes the difference between the implied price-dividend ratio from equation (35) and the observed price-dividend ratio of small stocks. Finally, with the set of parameters ($g_t, r_t, \tilde{g}_t, \tilde{r}_t, n$) in hand, we use equation (34) to calculate the duration of Corporate PBW. Figure D2d shows our baseline duration time series for Corporate PBW. Column (1) of Table D3 shows the full-sample and decade-by-decade averages.

We also compute a set of six alternative duration measures for Corporate PBW. The first two alternatives use the same two-stage GGM approach but with alternative empirical price-dividend targets. The first alternative uses Pitchbook data on price-revenue ratios for private equity portfolio companies with positive revenues. The Pitchbook data only start in 2003. We average this ratio in several ways: equally-weighted, revenue-weighted, or valuation-weighted, with either 1-5% trimming or winsorization when calculating these averages. The revenue-weighted average with 5% winsorization delivers the most conservative valuation ratios. Averaged from 2003-2019, we get a price/revenue ratio of 8.3. As with the IPO data, we use the same average revenue-earnings ratio of 7.7 for the first decile of publicly listed companies, and the earnings-dividend ratio of 2.0 for all public companies to translate the price-revenue ratio into a price-dividend ratio

[29] Using $n = 10$ or $n = 30$ for the length of the first stage does not materially change the final duration estimate.

($\frac{P}{Div} = \frac{P}{S} \times \frac{S}{E} \times \frac{E}{Div}$). Column (3) of Table D3 shows the duration computed using Pitchbook data. This value is much higher than the benchmark measure, making the latter conservative.

The second alternative uses the median price-revenue ratio reported by Ritter for companies that do an IPO (1980-2021). We use the average revenue-earnings ratio of 7.7 for the first decile of publicly listed companies, and the earnings-dividend ratio of 2.0 for all public companies to translate the price-revenue ratio into a price-dividend ratio ($\frac{P}{Div} = \frac{P}{S} \times \frac{S}{E} \times \frac{E}{Div}$). These calculations are conservative since they use sales-earnings and earnings-dividend ratios from more mature companies, which are lower than for companies that IPO. Column (3) of Table D3 shows the duration using IPO data, which is higher than our benchmark measure.

The third alternative uses a different method, namely a one-stage GGM applied to small stocks. It results in a much higher duration estimate. Average durations obtained from the one-stage Gordon Growth model for small stocks are listed in Column (4) of Table D3.

The fourth alternative calculates valuation ratios and durations using detailed information from the SCF dataset. Appendix D.6 explains the methodology. Column (5) of Table D3 shows a measure that is substantially higher than our baseline measure.

Table D3: Duration of Corporate PBW

	Corporate PBW				
	Baseline	Pitchbook	IPO	1-stage GGM	SCF
	(1)	(2)	(3)	(4)	(5)
1980-2019	55.93	58.92	56.08	95.67	64.13
1980-1989	33.12	41.82	32.42	73.43	64.24
1990-1999	59.00	61.23	58.59	116.19	65.60
2000-2009	72.17	70.38	72.35	107.54	65.33
2010-2019	59.42	62.27	60.94	85.51	61.37

Note: In this table we present the duration for Corporate PBW across various estimation methods. Column (1) displays the baseline estimates for Corporate PBW. In Column (2), we utilize the two-stage Gordon Growth model with Pitchbook data to estimate the duration. In Column (3), we use the two-stage Gordon Growth model with IPO data for duration estimation. Moving to Column (4), we present the duration estimated through a one-stage Gordon Growth model on CRSP Small stocks data. Additionally, in Column (5), the duration is estimated using the SCF, as elaborated in Appendix D.6.

A fifth alternative is to compute the Corporate PBW duration based on the affine model from Appendix A, using small stocks as our proxy for high-growth private businesses. That model does not use the GGM but rather has growth rates and valuation ratios that depend on the state variables in the VAR. The VAR-SDF model generates an average Corporate PBW duration of 68.0, again higher than our baseline estimate.

The final alternative is to compute the Corporate PBW duration based on small stocks, like in our benchmark approach, but instead of using the two-stage GGM, we use the observed evolution along the size distribution and the associated payout ratios, computed from the CRSP-Compustat micro data. This approach is detailed in Section D.7, and delivers a similar average duration of Corporate PBW between 52 and 62 depending on whether we use the smallest quintile or decile of stocks.

In conclusion, while the measurement of Corporate PBW duration is certainly not easy, the various approaches we have pursued result in similar estimates. Our benchmark estimate is at the low end of the empirical estimates, and hence conservative.

D.6 Alternative Calculations of Duration for Non-Corporate PBW and Corporate PBW Using SCF

As noted in Appendix Sections D.4 and D.5, one alternative measure of the duration of Non-Corporate PBW and Corporate PBW was based on SCF data. This appendix details those calculations.

D.6.1 Payout

We utilize data from the SCF to calculate the payout of each private business. For passively-held businesses, the reported net income (series X3410, X3414, X3418, X3422, X3426, X3430, X3454) serves as the measure of payout. For actively-managed businesses, we calculate the payout for each business (denoted by b) at a given time t using the following formula:

$$\text{Payout}^{active}_{b,t} = \text{Adjustment} * (1 - \text{NIPA Tax Rate}_t) * (\text{Net Income}_{b,t} - \text{Labor Income}_{b,t}). \tag{36}$$

We calculate the Net Income$_{b,t}$ by considering the reported income from the business (X3132, X3232, X3332) and multiplying it by the shares held by the household (X3128, X3228, X3328). Since the SCF provides net income before taxes, we apply a tax rate calculated using aggregate information from NIPA. We use an *Adjustment* coefficient of 0.5, the ratio of dividends (or payouts) to earnings.

Labor Income. To account for labor compensation, we subtract the Labor Income$_{b,t}$. This value is determined by summing the wages of the head and spouse (if they work in the business). We extract information from the SCF on whether the head is self-employed (X4106) and whether their second job is in their own business (X4504). The same information is collected for the spouse (X4706, X5104, respectively). If the head/spouse reports their wage, we use their reported wage; otherwise, we use an estimated shadow labor income.

To estimate the shadow labor income for actively managed businesses, we follow an approach similar tos Catherine, Miller, Paron, and Sarin (2022). First, we determine the shadow labor income for households. This involves calculating the wage income for both the head and spouse from their first and second jobs. We use data on the wage (head: X4112, X4509; spouse: X4712, X5108) and weeks of work (head: X4111, X4508, spouse: X4711, X5108) to compute the annual wage income for all households in the survey. The wage income is adjusted from weekly wages to annual wages using the frequency variables (head: X4113, X4508, spouse: X4713, X5108). Next, we perform regression analysis to estimate the shadow labor income. The wage income is regressed on a year fixed-effect, a fixed-effect of race, education, and sex, as well as a cubic function of age. Separate regression models are estimated for the head's first job, head's second job, spouse's first job, and spouse's second job. If households have multiple businesses, we imputed a pro-quota labor income, where the quota is based on the share of the business as percentage of total private business wealth.

D.6.2 Valuation Ratios and Duration

By employing this methodology, we can determine the payout for each private business. We also extract information on the market value of each actively-managed business (X3129, X3229, X3329) and the passively-held businesses (X3408, X3412, X3416, X3420, X3424, X3428, X3452). Combining information on the market value and the payout, we compute valuation ratios for each individual business. To ensure accurate calculations, we exclude private businesses with payouts equal to or less than 0, as this would result in negative or undefined ratios. We then compute dollar-weighted averages for each year. The weighting mechanism involves multiplying survey-weights with the corresponding market value of each business. This approach guarantees that businesses with higher market values carry more influence in the final valuation ratios, appropriately reflecting their significance in the overall assessment. We separately estimate valuation ratios for Corporate Businesses and Non-Corporate Businesses. To derive the duration from valuation ratios, we employ equation (7). To ensure completeness, we assume that the duration for the 1983 survey is equal to the average of 1989 and 1992, and we perform linear interpolations for the missing years. The summary statistics for the Non-Corporate Businesses are reported in Table D2 while those for Corporate Businesses are shown in Table D3.

D.7 Alternative Calculations of Duration of Corporate PBW Using CRSP-Computstat

Our benchmark Corporate PBW duration measurement uses the two-stage GGM assuming that private businesses behave like small stocks in the first phase of live and transition to become like the average publicly-listed company in the second stage. This approach may be understating the duration of Corporate PBW, to the extent that firms in the first decile of publicly-listed firms already experienced a lot of (cash flow) growth leading up to their inclusion in the publicly-listed universe. Including the cash-flow growth leading up to the IPO would result in a higher duration.

The approach may also be overstating duration in that it measures the duration of small public firms, holding fixed inclusion in this group for the duration of the first stage, set to 20 years. In reality, firms in the bottom decile of publicly-listed firms may grow faster and transition more quickly into higher deciles of the market capitalization distribution. Since larger firms may have lower cash flow pay-out ratios, assuming slower transitions than observed may lead us to overstate the duration of Corporate PBW. In this appendix, we address this potential overstatement issue, by computing the duration of firms that are *currently* in the bottom decile (or quintile) of the market cap distribution , but may not remain there in the future.

D.7.1 Measuring Duration with Firm Life-Cycles

The duration of a firm is the weighted average time to its cash flows:

$$D = \sum_{t=1}^{\infty} t \frac{PV_t}{\sum_{t=1}^{\infty} PV_t}$$

Let s indicate the current-year size group of a firm, where size is measured by market capitalization. Let there be S groups. We assume that $PV_t = CF_t(1+R)^{-t}$ for some constant discount rate R, calibrated as discussed below. We model the cash flow of the median firm in size group s_t, which came from size group s_{t-1} in the previous period, as the product of the payout-asset ratio of the median firm in that size group and the assets of the median firm in that size group:

$$CF_t(s_t|s_{t-1}) = (CF_t/A_t)\,(s_t|s_{t-1}) \cdot A_t(s_t|s_{t-1}) \tag{37}$$

The state (market capitalization group) transition matrix is denoted by $\mathcal{P}(s_t|s_{t-1})$. Conditional on starting out in the smallest decile at time zero, the cash flow of a typical firm t periods later is:

$$Div_t|s_0 = \sum_{s_t=1}^{S} \mathcal{P}^{t-1} \cdot (\mathcal{P} \cdot (Div_t/A_t) \cdot A_t) \tag{38}$$

D.7.2 Implementation

We use CRSP-Compustat data on the universe of publicly-listed firms for the standard sample from 1967–2020. Market capitalization is measured as price per share times shares outstanding, properly adjusted for stock splits. We also make an adjustment for mergers & acquisitions. As is commonly done, we delete stocks whose price is below $1 per share and whose market capitalization is less than $10 million at the first time of observation (and only then).

Cash flow CF is either computed as cash dividends or as cash dividends plus net share repurchases, with the latter bounded from below at zero. Cash flows and assets are deflated by the consumer price index. To compute assets and the cash flow-to-asset ratio in each size group, we first compute book assets and CF/asset ratios for each firm, then winsorize at the 1% level, then compute the median across the firms that are in size group s_t in the current year and were in size group s_{t-1} in the prior year. This delivers a time series for the $S \times S$ matrices $(CF_t/A_t)\,(s_t|s_{t-1})$ and $A_t(s_t|s_{t-1})$. We then average these objects across years.

Our groups are either market capitalization deciles ($S = 10$) or quintiles ($S = 5$). When computing the size transition probability matrix $\mathcal{P}$, we collapse set all transition probabilities that are more than three notches up (down) to zero and add the empirical weight of those transitions to the state that is exactly three notches up (down). We take the time-series average of the state transition probability matrices in each year.

Finally, we calibrate the discount rate R, needed in the duration calculation, in order to obtain a duration of 28 for the value-weighted market portfolio of all stocks. This is the duration of the aggregate stock market we estimate in the auxiliary asset pricing model. This enables comparability across approaches.

D.7.3 Results

Size Deciles. Using deciles for size groups, the transition probability matrix is $\mathcal{P}(s'|s) =$

$$\begin{bmatrix}
75.1\% & 19.5\% & 3.6\% & 1.7\% & 0.0\% & 0.0\% & 0.0\% & 0.0\% & 0.0\% & 0.0\% \\
20.3\% & 49.8\% & 22.2\% & 5.8\% & 2.0\% & 0.0\% & 0.0\% & 0.0\% & 0.0\% & 0.0\% \\
3.7\% & 20.9\% & 43.9\% & 22.8\% & 6.9\% & 1.8\% & 0.0\% & 0.0\% & 0.0\% & 0.0\% \\
0.8\% & 5.3\% & 20.7\% & 42.2\% & 23.8\% & 5.9\% & 1.2\% & 0.0\% & 0.0\% & 0.0\% \\
0.0\% & 1.6\% & 5.2\% & 19.8\% & 43.3\% & 24.5\% & 4.9\% & 0.6\% & 0.0\% & 0.0\% \\
0.0\% & 0.0\% & 1.7\% & 4.2\% & 18.8\% & 47.2\% & 24.3\% & 3.6\% & 0.2\% & 0.0\% \\
0.0\% & 0.0\% & 0.0\% & 1.4\% & 3.5\% & 17.6\% & 52.2\% & 23.7\% & 1.5\% & 0.0\% \\
0.0\% & 0.0\% & 0.0\% & 0.0\% & 1.4\% & 2.5\% & 15.4\% & 61.0\% & 19.5\% & 0.2\% \\
0.0\% & 0.0\% & 0.0\% & 0.0\% & 0.0\% & 1.2\% & 1.3\% & 12.1\% & 72.9\% & 12.5\% \\
0.0\% & 0.0\% & 0.0\% & 0.0\% & 0.0\% & 0.0\% & 1.8\% & 0.9\% & 8.6\% & 88.6\%
\end{bmatrix}$$

Table D4 shows, for each of the size groups, the dividend/asset ratio, the payout/asset ratio (which includes net share repurchases in the numerator), log assets, and the duration using either dividends or payouts. For the smallest decile of listed firms, which is our proxy for private businesses, we obtain a duration of 62.5 using cash dividends and 62.3 using the broader payout measure. We conclude that this number is quite similar to the 61.25 number we use in our benchmark results.

Table D4: Duration by Size Decile

Deciles	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10
Log asset	4.09	4.82	5.18	5.53	6.04	6.43	6.94	7.48	8.32	9.58
CF / asset (div, %)	0.10	0.21	0.30	0.37	0.51	0.56	0.77	1.03	1.36	1.93
CF / asset (payout, %)	0.13	0.34	0.42	0.49	0.68	0.81	1.06	1.36	1.70	2.35
Duration (div)	**62.5**	59.8	56.7	53.4	49.6	45.4	40.8	35.6	29.7	23.2
Duration (payout)	**62.3**	59.6	56.5	53.3	49.5	45.3	40.7	35.6	29.7	23.3

Note: The first row reports the log of book assets of the median firm in each decile of market capitalization. The second and third rows report the ratio of cash flows to book assets for the median firm in each decile of market capitalization, where cash flows are measured as cash dividends (div) in the first instance and dividends plus the max of net share repurchases and zero in the second instance. Assets and CF/assets depend on both the current size decile and the prior year's size decile, but are integrated across the prior year's size deciles for presentation purposes. The last two rows report the durations, using either dividends or dividends plus net share repurchases as the measure of cash flow.

Size Quintiles. As a further robustness check, we also compute durations for quintiles, assuming that private businesses resemble firms in the bottom-20% of the size distribution of listed firms. Using quintiles for size groups, the transition probability matrix is $\mathcal{P}(s'|s) =$

$$\begin{bmatrix} 81.8\% & 17.0\% & 1.2\% & 0.1\% & 0.0\% \\ 15.2\% & 64.9\% & 19.1\% & 0.7\% & 0.0\% \\ 1.0\% & 15.1\% & 66.9\% & 16.8\% & 0.1\% \\ 0.2\% & 1.0\% & 12.1\% & 76.1\% & 10.7\% \\ 0.0\% & 0.5\% & 0.7\% & 7.7\% & 91.0\% \end{bmatrix}$$

Table D5 shows, for each of the size groups, the dividend/asset ratio, the payout/asset ratio (which includes net share repurchases in the numerator), log assets, and the duration using either dividends or payouts. For the smallest decile of listed firms, which is our proxy for private businesses, we obtain a duration of 52.0 using cash dividends and 51.9 using the broader payout measure.

Combining the results for deciles and quintiles suggests a value between 51.9–62.5 for the duration of Corporate PBW.

D.8 Duration of Vehicles

To determine the average duration of vehicles in the US for each year, we adopt a model that assumes a constant depreciation rate (δ) for the car's value. If $t = 0$ is when the car is new, then we calculate the future cash flow at time t using the formula:

$$CF_t = (1 - \delta)^t \delta, \tag{39}$$

Table D5: Duration by Size Quintile

Quintiles	Q1	Q2	Q3	Q4	Q5
Log asset	4.44	5.34	6.19	7.20	8.88
CF / asset (div, %)	0.12	0.32	0.50	0.89	1.63
CF / asset (payout, %)	0.20	0.42	0.70	1.18	1.99
Duration (div)	**52.0**	47.7	41.8	34.3	25.2
Duration (payout)	**51.9**	47.6	41.7	34.2	25.3

Note: The first row reports the log of book assets of the median firm in each quintile of market capitalization. The second and third rows report the ratio of cash flows to book assets for the median firm in each quintile of market capitalization, where cash flows are measured as cash dividends (div) in the first instance and dividends plus the max of net share repurchases and zero in the second instance. Assets and CF/assets depend on both the current-year and the prior year's size quintiles, but are integrated across the prior-year's size quintiles for presentation purposes. The last two rows report the durations, using either dividends or dividends plus net share repurchases as the measure of cash flow.

where $(1-\delta)^t$ represents the value after depreciation. The duration for a car that is j years old is then computed as follows:

$$D_j^v = \frac{\sum_{t=j}^{T}(t-j)\frac{CF_t}{(1+r)^{t-j}}}{\sum_{t=j}^{T}\frac{CF_t}{(1+r)^{t-j}}}, \tag{40}$$

where r denotes the 1-year real Treasury rate in each year.

We determine the value for j as the average age of vehicles in use in the U.S., based on data from the Bureau of Transportation Statistics.[30] Additionally, we set the maximum vehicle age, T, equal to double the average age of vehicles: $T = 2j$. To calculate the depreciation rate of vehicles, we utilize data from the U.S. Bureau of Economic Analysis (BEA). The depreciation rate is computed as the current-cost depreciation of autos (BEA Table 2.4) divided by the current-cost net stock of autos in the previous year (BEA Table 2.1). For the years 1980 to 2019, the average depreciation rate (δ) is found to be 22%. Figure D2e plots the duration time series for Vehicles.

D.9 Duration of Fixed Income Assets

For fixed-income assets, we utilize the duration of ICE-BofA US Corporate & Government bonds, which is available starting in 1996. To estimate the duration for the period 1980 through 1995, we employ the following imputation approach. First, we establish a linear regression model using the 1996-2020 sample of ICE-BofA duration data and the duration from CRSP government bond data (available since 1946), along with a constant. Second, we apply the parameters obtained from this regression model to estimate the ICE-BofA duration values for the sample period of 1980-1995. Figure D2f plots the time-varying duration series.

[30] The data is available at: `https://www.bts.gov/content/average-age-automobiles-and-trucks-operation-united-states#:~:text=2018%2D19%3A%20IHS%20Markit%20Co,%2Dmarkit%2D%20as%20of%20Sep`. Data has been available since 1995, and we iterate it backward using a constant growth rate estimated from the period 1995 to 2022.

D.10 Mortgage Duration

For mortgage debt analysis, we obtain data from the Bloomberg-Barclays Aggregate MBS Index, which provides a comprehensive representation of all outstanding U.S. pass-through mortgage-backed securities. In the United States, the most prevalent mortgage product is the 30-year fixed-rate mortgage. However, it is essential to note that the average outstanding mortgage has a significantly lower duration due to factors such as loan aging, amortization, coupon payments, and prepayment. These factors contribute to variations in mortgage duration over time and are critical considerations when studying mortgage debt dynamics. See Figure D2g.

D.11 Vehicle Debt Duration

We compute the vehicle debt duration using the SCF. The SCF provide detailed information on up to four vehicles loans (the information is limited to three vehicles loans for the survey 1989 and 1992). We use information on all the available vehicles loans. First, we calculate the number of residual monthly payments. We then subtract the number of payments already made from the total number of payments to determine the residual number of payments T. Second, we calculate the monthly payment C_t of each loan. Third, we calculate the monthly interest rate r_t charged on the loan. With this information, we calculate the monthly duration of the vehicle debt j of household i:

$$D_{i,j}^{vd} = \left[\frac{\sum_{t=1}^{T_{i,j}} \frac{t * C_{i,j}}{(1+r_{i,j})^t}}{\sum_{t=1}^{T_{i,j}} \frac{C_{i,j}}{(1+r_{i,j})^t}} \right] /12.$$

We calculate the duration of vehicle debt for each household i by taking the weighted average of the duration of its individual vehicle debts. To determine the weights, we use the outstanding balance of each debt j. The vehicle debt duration measure D^{vd} is computed as the median duration of households' vehicle debts in each survey year. Our first estimate is available in 1989, based on the initial SCF survey. For the years from 1980 until 1988, we perform extrapolation using the trend growth rate observed in the sample from 1989 to 2019. Additionally, we linearly interpolate for the years between surveys. Finally, for the 1992 survey we use the interpolated value between 1989 and 1995 as the duration estimate is anomalously low compared to all other observations.

Figure D1 shows the distribution of information on each individual vehicle debt in 2019. The figure provides information on the cross-section of: (i) the length of vehicle debts measured in number of years at inception or residual, (ii) the duration estimates ($D_{i,j}^{vd}$), (iii) the annual interest rates of the loan, and (iv) the monthly payment. Figure D2h shows the time series of Vehicle Debt duration.

Figure D1: Vehicle Debt

(a) Years to Maturity

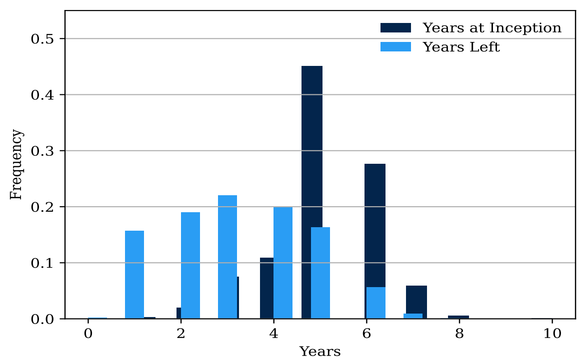

(b) Duration

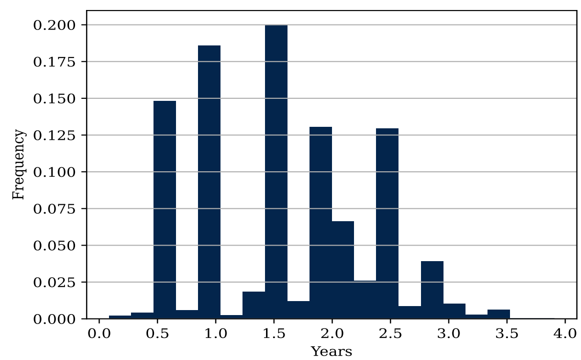

(c) Annual Rate

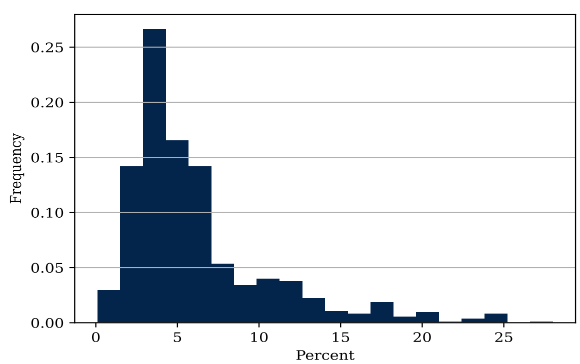

(d) Monthly Payment

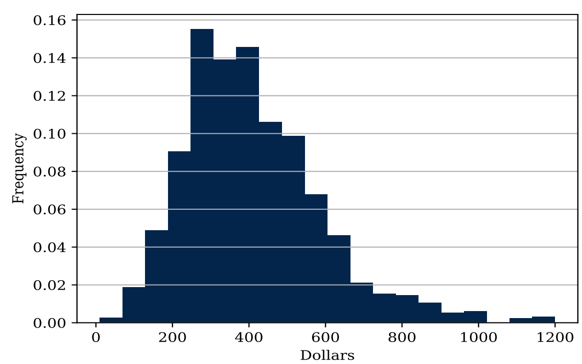

Note: The figures plot the distribution of vehicle debts in 2019. Figure D1a plots the distribution of number of years left and number of years at inception. Figure D1b plots the distribution of duration estimates. Figure D1c plots the distribution of annualized interest rates on loans. Figure D1d plots the distribution of monthly payments. Source: SCF 2019.

Figure D2: Duration Estimates

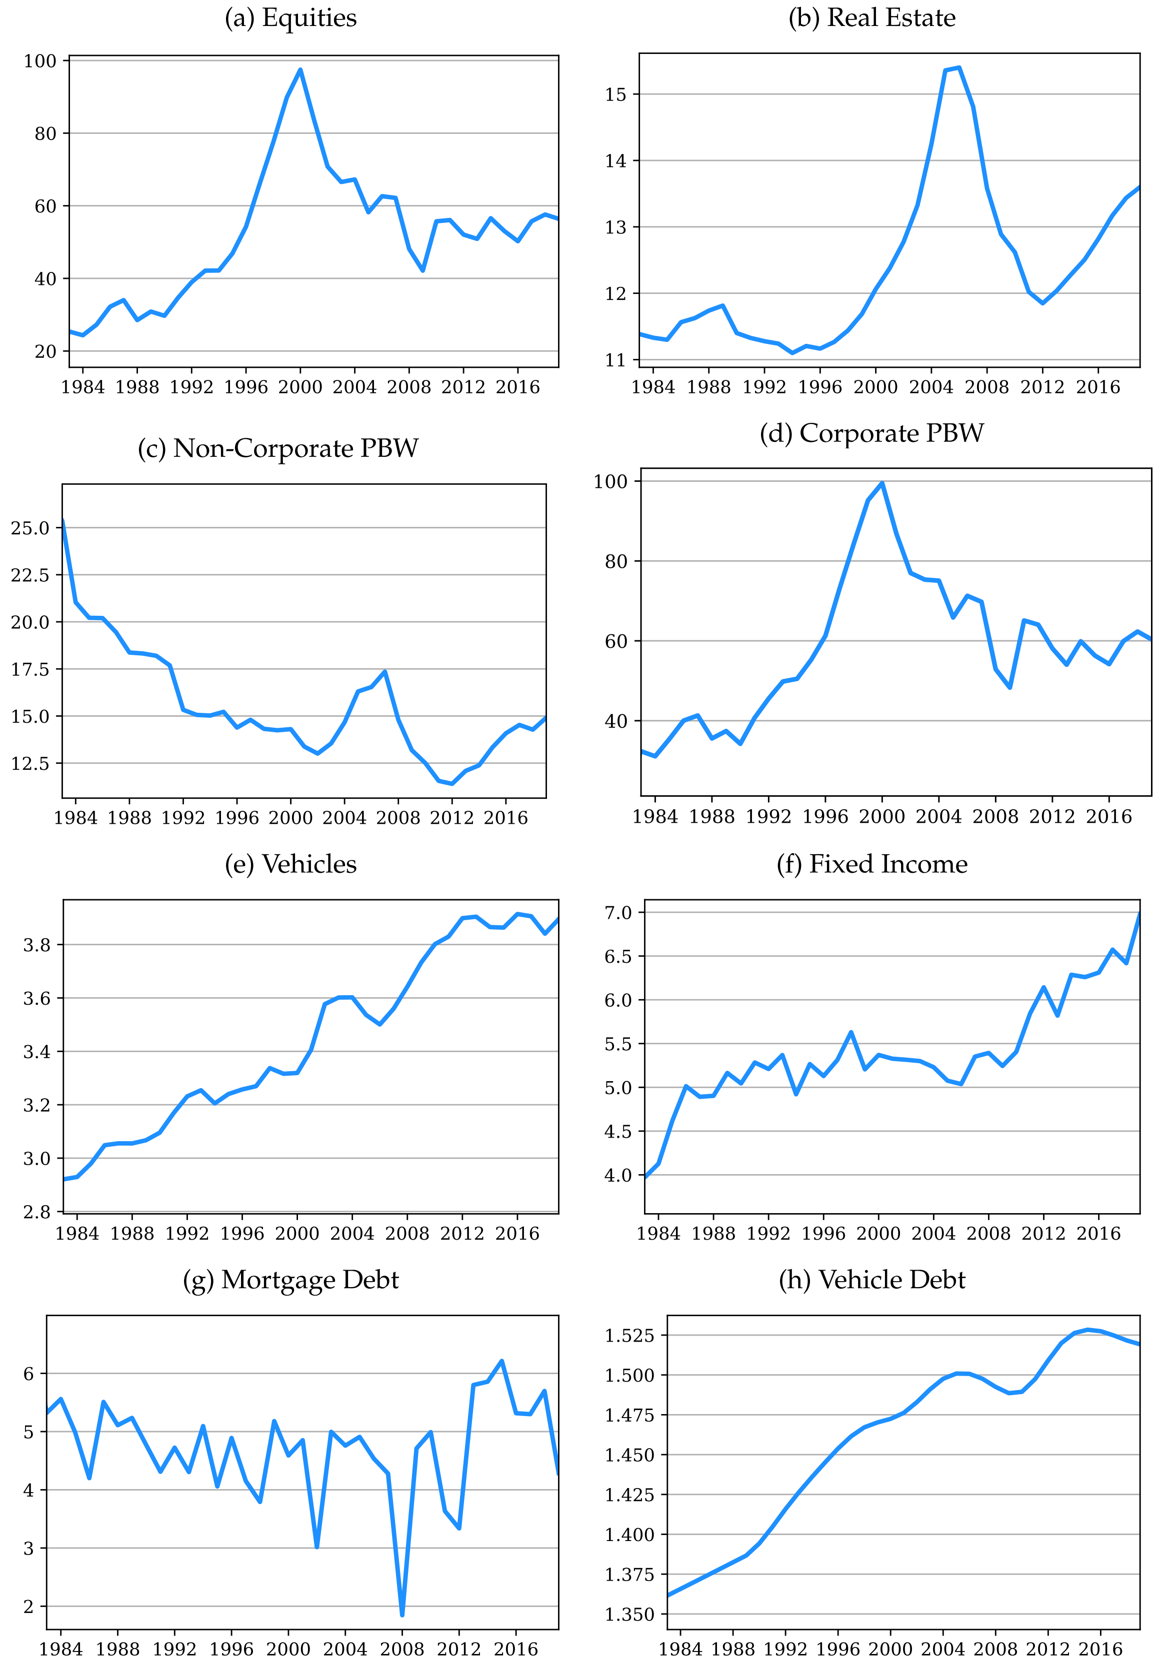

Note: The figures plot the time varying duration estimates for the set of assets and liabilities.

D.12 Portfolio Duration

We calculate the portfolio duration of each household by multiplying the duration of each asset (or liability) by its corresponding portfolio share and summing over all assets (liabilities) in the portfolio. In each year, we winsorize the top/bottom 2.5% of households ranked by the duration of their portfolio.

In Section 2.3 we discussed the heterogeneity in portfolio duration by net wealth and age. Figure D3a shows the heterogeneity in portfolio duration by income percentile. We rank households by their income and compute the median duration in for each income percentile bucket, year by year. We then average across all surveys from 1983 to 2019. In Figure D3b we use the same procedure but instead use wealth-weighted net wealth percentile. Each wealth-weighted percentile bin is designed such that the share of total wealth held by the households in each bin is the same across different bins.

Figure D3: Distribution of Durations

(a) Duration by Income

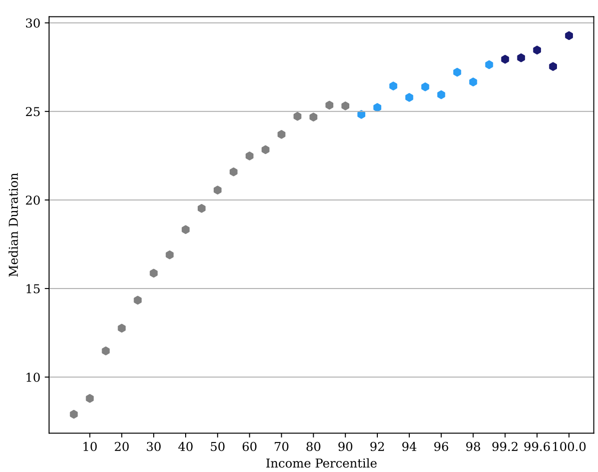

(b) Duration by Wealth-Weighted Net Wealth

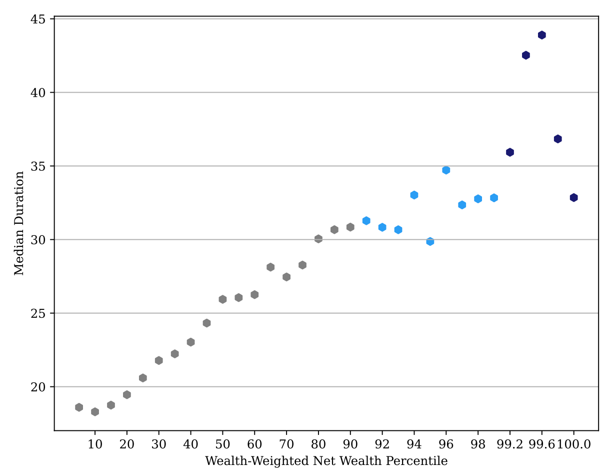

Note: Figure D3a shows the heterogeneity by income percentile. We rank households by their income and compute the median duration for each income percentile bucket, year by year. We then average across all surveys from 1983 to 2019. In Figure D3b we use the same procedure but instead use wealth-weighted net wealth percentile. Each wealth-weighted percentile bin is designed such that the share of total wealth held by the households in each bin is the same across different bins.

Households' portfolio duration notably differs by age. However, there is also dispersion within age group. Figure D4 provides further information on the within age group dispersion of duration. For each age group, we rank households by the duration of their portfolio. The figure includes the 5%, 25%, 50%, 75% and 95% percentile, as well as the median (blue line) and mean (orange triangle). We pool households across all surveys from 1983 to 2019.

We also evaluate more formally the correlation between financial duration and some covariates

Figure D4: Distribution of Durations

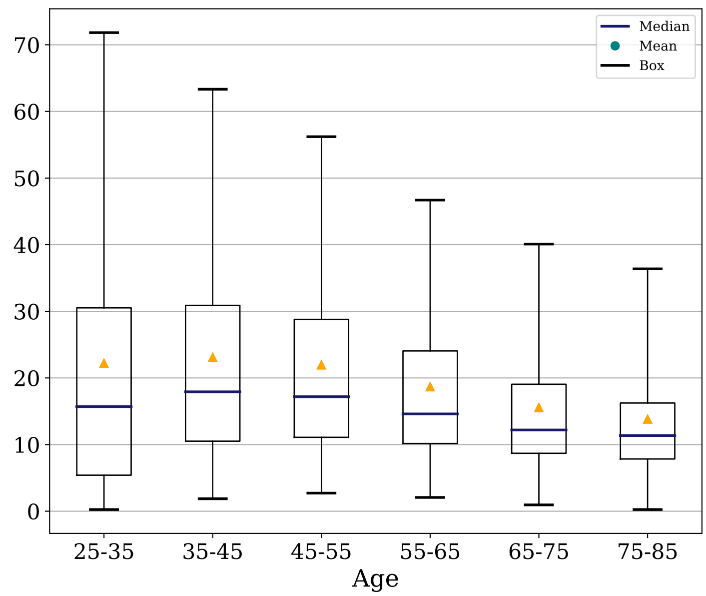

Note: Figure D4 provides further information on the within age group dispersion of duration. For each age group, we rank households by the duration of their portfolio. The figure includes the 5%, 25%, 50%, 75% and 95% percentile, as well as the median (blue line) and mean (orange triangle). We pool households across all surveys from 1983 to 2019.

of interest. Table D6 reports the estimation results. Data are based on all SCF survey from 1983 to 2019. All regression models include year fixed effects. In column (1), we regress household financial duration on household age. In column (2), we regress financial durations household position in the Lorenz Curve. To calculate households' positions, we rank households by their net wealth, then calculate the cumulative sum of net wealth and divide by the aggregate net wealth. In column (3), we regress financial durations on both age and Lorenz Curve position. In column (4), we add a quadratic function of age. In column (5), we add the log of household income. In column (6), we add the logarithm of households net wealth.

D.12.1 Financial Duration Over Time

Figure D2a-D2h highlighted the degree of time variation in the duration estimates of each component of the household portfolio. In this section, we study the time-varying properties of the duration of the overall household portfolios. Figure D5a displays the duration estimates for all households included in the sample. The graph shows both wealth-weighted and equally-weighted averages; they show similar dynamics. The population of households is further broken down into the bottom-90% in Figure D5c, the top-10% in Figure D5e, and the top-1% in Figure D5g. All series

Table D6: Determinants of Household-level Financial Duration

	(1)	(2)	(3)	(4)	(5)	(6)
Age	-0.20***		-0.24***	0.022	-0.10***	-0.14***
	(0.0024)		(0.0023)	(0.016)	(0.015)	(0.015)
Lorenz Curve		0.14***	0.20***	0.19***	-0.045***	-0.10***
		(0.0021)	(0.0019)	(0.0020)	(0.0031)	(0.0029)
Age Squared				-0.0023***	-0.00028**	-0.00020
				(0.00014)	(0.00013)	(0.00013)
Log-Income					5.72***	4.89***
					(0.050)	(0.055)
Log-Net-Wealth						0.82***
						(0.019)
Constant	20.7***	9.31***	20.3***	14.0***	-44.3***	-42.2***
	(0.13)	(0.060)	(0.13)	(0.43)	(0.66)	(0.66)
Year effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	256361	256361	256361	256361	256361	256361
R^2	0.109	0.090	0.137	0.138	0.197	0.204

Note: The table reports the estimation results of regressing households' portfolio duration on a set of covariates. Data are based on all SCF survey from 1983 to 2019. All regression models include year fixed effects. In column (1), we regress household financial duration on household age. In column (2), we regress financial durations household position in the Lorenz Curve. To calculate households' positions, we rank households by their net wealth, then calculate the cumulative sum of net wealth and divide by the aggregate net wealth. In column (3), we regress financial durations on both age and Lorenz Curve position. In column (4), we add a quadratic function of age. In column (5), we add the log of household income. In column (6), we add the logarithm of households net wealth. Standard Errors in parentheses (* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$).

peak around the year 2000 when the duration of equity and PBW peak, and the portfolio shares of these wealth categories are large, especially for the wealthiest households.

To better understand the impact of portfolio weights, we conduct a separate analysis that fixes the duration of assets and liabilities at the sample average (measured over the period from 1980 to 2019). This approach isolates the role of time-varying portfolio shares. The results are presented in the right panels: Figures D5b, D5d, D5f, and D5h are for all households, bottom-90%, top-10%, and top-1%, respectively. They show much more muted dynamics, suggesting that most of the time variation in household portfolio duration arises from time variation in the duration of the underlying asset durations.

Figure D5: Financial Duration Over Time

(a) All Households

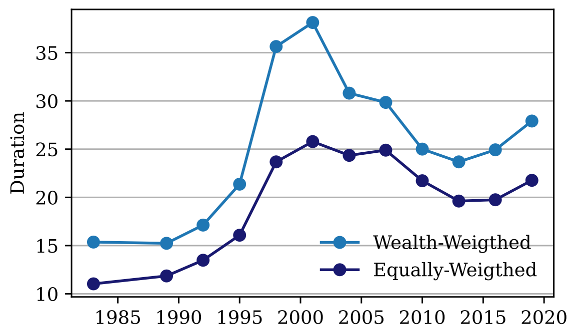

(b) All Households with Fixed Duration

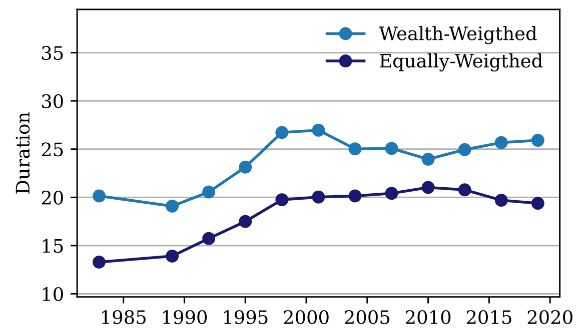

(c) Bottom 90

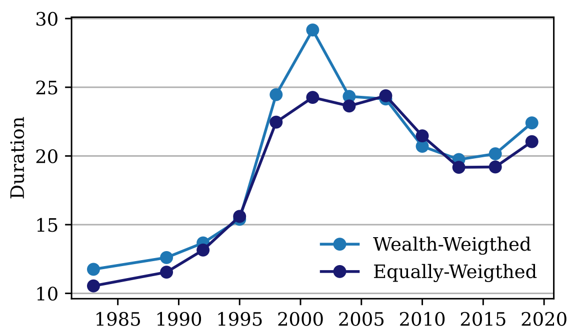

(d) Bottom 90 with Fixed Duration

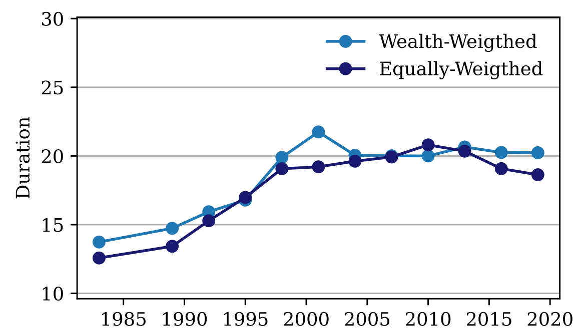

(e) Top 10

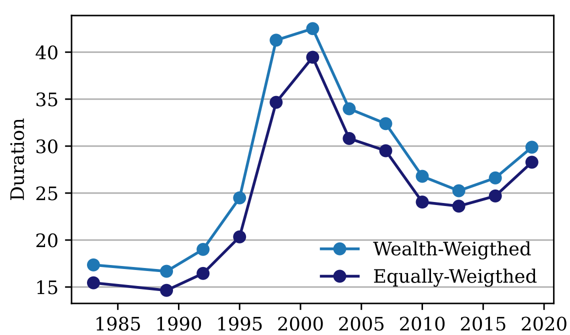

(f) Top 10 with Fixed Duration

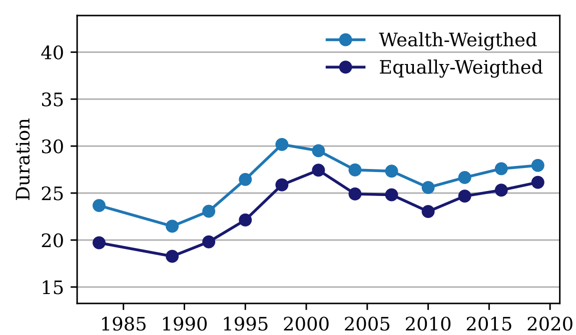

(g) Top 1

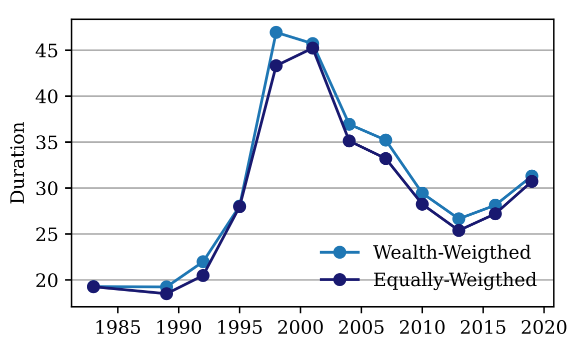

(h) Top 1 with Fixed Duration

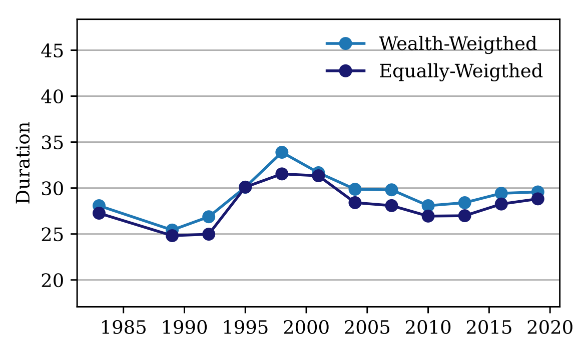

Note: Figure D5 shows the time varying duration of households' portfolios. We use information from each SCF survey from 1983 till 2019 to measure the equally-weighted and wealth-weighted duration over time. Figure D5a plots the duration for all households in our sample. Figure D5c, Figure D5e and Figure D5g plot the duration for the Bottom 90%, the top 10% and the top 1%. In Figure D5b, D5d, D5f and D5h we estimate households' portfolios duration when the duration of assets and liabilities is fixed to their sample averages.

E Income Process

E.1 Data Source: PSID

The Panel Study of Income Dynamics (PSID) is a household panel survey that began in 1968. The PSID was originally designed to study the dynamics of income and poverty. Thus, the original 1968 PSID sample was drawn from two independent samples: a sample of 1,872 low income households from the Survey of Economic Opportunity (the "SEO sample") and a nationally representative sample of 2,930 households designed by the Survey Research Center at the University of Michigan (the "SRC sample"). In this paper, we use the "SRC sample" for the time period from 1970 until 2017.

E.1.1 PSID Income variables

We now describe the construction of the relevant income variables used in the paper. We construct the following variables: *labinc2f* is labor income excluding transfers but including the labor part of business and farm income for both head and eventual spouse; *transf* which are total households transfer (including Social Security Income and other transfers); *labinc3f*, which is our measure of total household income for both head and eventual spouse, is the sum of *labinc2f* and *transf*.

We provide further details on how we build these three variables. As the variables included in the PSID are subject to change, the variable construction vary with different sample period. For this reason, below we provide details on the variables used in different time periods. Moreover, the ticker for each variable changed in each survey. We therefore define the ticker used in a specific year as (YYYY:Ticker).[31]

labinc2f In the 1970 - 1993 sample, this variable is defined as the sum of Total labor income of head, including wages and salaries, labor part of business income and farm income (1993:V23323), and Spouse's total labor income, including labor part of business income and farm income (1993:V23324). In the 1993 - 2017 sample, this variable is defined as the sum of Reference Person's total labor (including wages and other labor) excluding Farm and Unincorporated Business Income, (2017:ER71293), Labor Part of Business Income from Unincorporated Businesses (2017:ER71274), Reference Person's and Spouse's/Partner's Income from Farming (2017:ER71272), Wife's Labor Income, Excluding Farm and Unincorporated Business Income (2017:ER71321), Wife's Labor Part of Business Income from Unincorporated Businesses (2017:ER71302). Note that farm's income includes both labor and asset portions of income.

[31] The PSID website provides information on how to harmonize tickers across different surveys.

transf In the 1970-1993 sample, this variable is defined as Total Transfer Income of Head and Wife/"Wife" (1993:V22366) and Total Transfer Income of Others (1993:V22397). In the 1994-2003 sample, this variable is defined as Head's and Wife's Total Transfer Income, Except Social Security (2017:ER71391), Other Total Transfer Income, Except Social Security (2017:ER71419), Total Family Income from Social Security (1994:ER4152). In the 2004-2017 sample, this variable is defined as: Head's and Wife's Total Transfer Income, Except Social Security (2017:ER71391), Other Total Transfer Income, Except Social Security (2017:ER71419), Reference Person's Income from Social Security (2017:ER71420), Spouse's/Partner's Income from Social Security (2017:ER71422), Others Income from Social Security (2017:ER71424).

labinc3f We then construct *labinc3f* by summing total family labor income (*labinc2f*) and total family transfers (*transf*).[32]

E.2 Estimating the Income Process

Age Profile We estimate the age profile of income following Deaton and Paxson (1994). First, we estimate the average income for each cohort in each year, using PSID data. $y_{c,t}$ is the average income of cohort c at time t, based on our *labinc3f* definition of income. Then, we estimate the following regression model:

$$\log y_{c,t} = \beta + \gamma_a + \gamma_c + \gamma_t + \varepsilon_{c,t}, \tag{41}$$

where the subscript c, t, a refers to cohort, time and age, respectively. We define as c the age at time $t = 0$ (i.e. 1970). Due to the linear relationship between age, cohort and time, we cannot separately identify the different fixed effects. We hence resort to the method used by Deaton and Paxson (1994): we attribute growth to age and cohort effects, while we use the year effects to capture cyclical fluctuations or business-cycle effects that average to zero over the long run. We hence constraint the year fixed effects to be orthogonal to a time trend and to sum to zero. We then estimate Equation 41 using constrained OLS.

Figure E1 plots the estimates for the age dummies. The dots are based on the estimated dummies. The dashed lines apply a Savitzky–Golay filter to smooth the estimates and characterize our deterministic age-profile.

Income Risk In the second stage, we estimate income risk. We estimate (17) year by year and include cubic function of age as well as a set of fixed-effects: education, race, gender, state. We then extract the residuals z_{it}. Finally we estimate the risk parameters by GMM as detailed below.

[32] We have verified that aggregating the PSID results in a series that is close to the NIPA series.

Figure E1: Income Profile

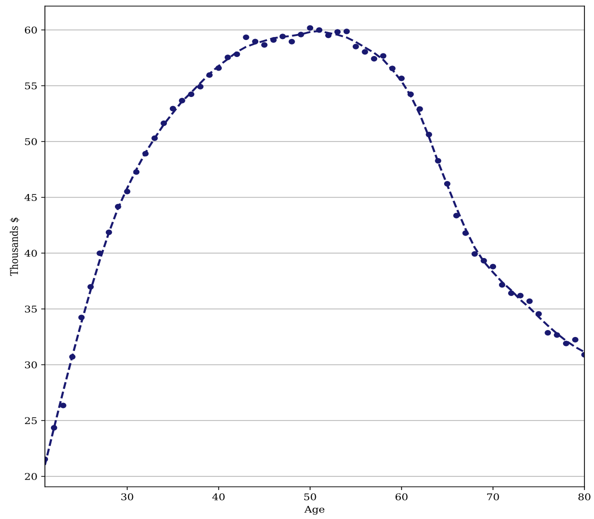

Note: This figure displays the expected income profile evaluated at the 2016 year-fixed effects. The graph plots the expected income profile for the average person who is 21 years old in 2016, expressed in thousands of 2016 dollars. The model is estimated according to Equation (41) on PSID data from 1970 to 2017.

Using Equation (17)-(19), and define j as equal to the age of the households minus the minimum age (21), we find that:

$$\begin{aligned} E[z_j^i, z_{j+h}^i] &= \sigma_\alpha^2 + E[{\varepsilon_j^i}^2] + \sigma_\nu^2 \qquad \text{if } h = 0, \\ E[z_j^i, z_{j+h}^i] &= \sigma_\alpha^2 + \rho^h E[{\eta_j^i}^2] \qquad \text{if } h > 0, \\ E[{\eta_j^i}^2] &= \rho^{2j}\sigma_{\eta_0}^2 + \sum_{k=1}^{j} \rho^{2(j-k)}\sigma_u^2. \end{aligned}$$

We allow the variance to differ in working age (w) and retirement age (r), where the retirement age starts at 65. We fixed the variance of initial persistent shocks $\sigma_{\eta 0}^2 = 0$, then use a GMM estimation to estimate $\theta = (\rho, \sigma_{\nu,w}, \sigma_{u,r}, \sigma_{\nu,r}, \sigma_{u,r}, \sigma_\alpha)$. We use a Minimum Distance Estimator, where the weighting matrix is the identity matrix. We only include sample moments estimated on 100 or more observations.

Sample Selection. We use PSID data from 1970 to 2017. As discussed in Heathcote, Perri, and Violante (2010), after survey year 1997, the data frequency goes from annual to biannual. To make the estimation consistent, in the first part of the sample 1970-1997 we also sample data at biannual frequency. We only include households whose head is 21 to 80 years old. We only include households which were in the survey for three or more periods. We exclude households with zero or negative income. In each year, we trim the top 2.5% of households by their income.

The point estimates are displayed in Table E1. These are the parameters used in the main text.

Table E1: Idiosyncratic Risk Parameter Estimates

	σ^2_α	$\sigma^2_{\nu,w}$	$\sigma^2_{u,w}$	$\sigma^2_{\nu,r}$	$\sigma^2_{u,r}$	ρ
Estimated Parameters	0.0762	0.1605	0.0413	0.0906	0.0255	0.9152

Note: $\theta = (\rho, \sigma_{\nu,w}, \sigma_{u,r}, \sigma_{\nu,r}, \sigma_{u,r}, \sigma_\alpha)$, are estimated using Equation (17)-(19); $\sigma^2_{\varepsilon_0}$ is fixed equal to 0. Data are based on PSID and runs from 1970 to 2017.

E.2.1 Time Varying Income Risk

We also estimate the risk parameters using rolling sample of the PSID from 1983 till 2016. We use sample of 15 years (apart from 1983 and 1984 where we include data from 1970 to 1983 and 1984, respectively). We fix the autocorrelation ρ of persistent shocks to the full sample value estimated in Table E1. Figure E2 plots time varying estimates of $\sigma^2_{u,w}$, $\sigma^2_{\nu,w}$, $\sigma^2_{u,r}$, $\sigma^2_{\nu,r}$ and σ^2_α.

Figure E2: Time Varying Income Risk

(a) Transitory Shocks - Working Age

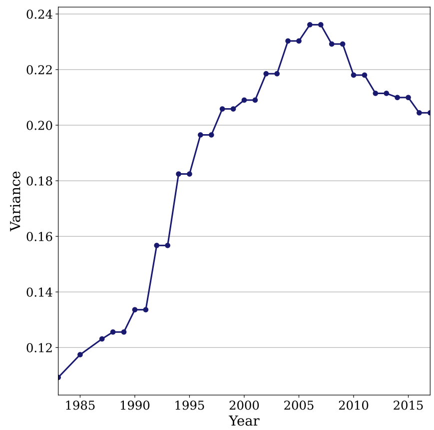

(b) Persistent Shocks - Working Age

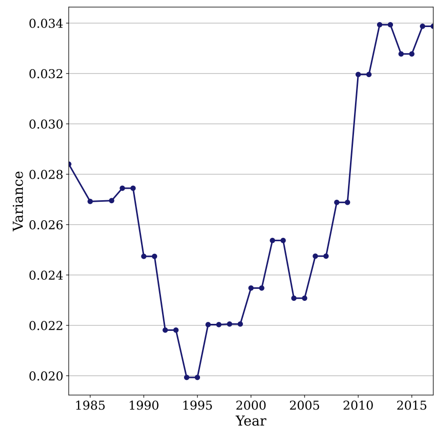

(c) Transitory Shocks - Retirement

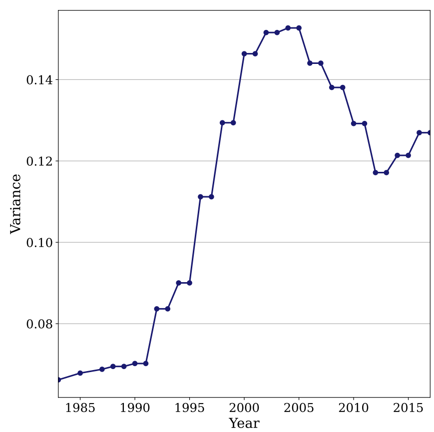

(d) Persistent Shocks - Retirement

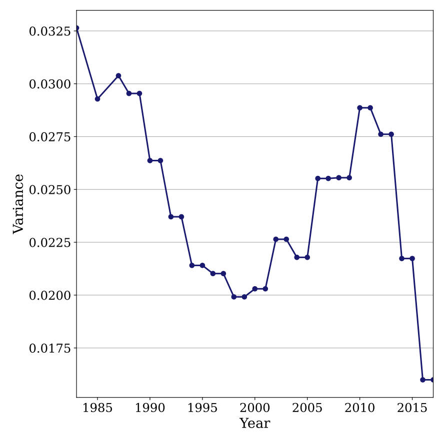

(e) Fixed-Effect Shocks

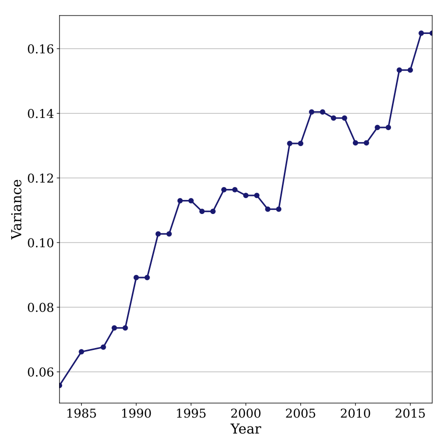

Note: This figure displays the the risk parameters using rolling sample of the PSID. We use sample of 15 years (apart from 1983 and 1984 where we include data from 1970 to 1983 and 1984, respectively). We fix the autocorrelation ρ of persistent shocks to the full sample value estimated in Table E1. Panel E2a plots time varying estimates of $\sigma^2_{u,w}$, Panel E2b plots time varying estimates of $\sigma^2_{v,w}$, Panel E2c plots time varying estimates of $\sigma^2_{u,r}$, E2d plots time varying estimates of $\sigma^2_{v,r}$ and E2e plots time varying estimates of σ^2_{α}. The model is estimated according to Equation (17)-(19) on PSID data from 1970 to 2017.

F Proofs

F.1 Proof of Proposition 1

Rewriting (1) yields

$$\log P_0 = \log \left\{ \sum_{t=0}^{\infty} \exp(-t \times \log R) z_t \right\}.$$

Taking the derivative, we now obtain

$$\begin{aligned} \frac{\partial \log P_0}{\partial \log R} &= \left\{ \sum_{t=0}^{\infty} \exp(-t \times \log R) z_t \right\}^{-1} \sum_{t=0}^{\infty} (-t) \exp(-t \times \log R) z_t \\ &= P_0^{-1} \sum_{t=0}^{\infty} (-t) R^{-t} z_t = -\frac{\sum_{t=0}^{\infty} R^{-t} z_t \times t}{P_0} = -D. \end{aligned}$$

A first order approximation now immediately yields

$$\log \tilde{P}_0 \simeq \log P_0 - D \times \varepsilon$$

which implies

$$\tilde{P}_0 \simeq P_0 \exp(-D \times \varepsilon).$$

Applying the approximation $x \simeq \log(1 + x)$ as $x \to 0$ yields the final approximation

$$\tilde{P}_0 \simeq P_0 \times (1 - D \times \varepsilon).$$

We now relate these results to a portfolio of assets, indexed by k. Let $v(k)$ be the value of asset k in the portfolio, and let $V = \sum_k v(k)$ be the total value of the portfolio. Then following a small shock ε we have

$$\begin{aligned} \tilde{V} &= \sum_k \tilde{v}(k) = \sum_k v(k)(1 - D(k)\varepsilon) = V \times \sum_k \omega(k)(1 - D(k)\varepsilon) \\ &= V \times \left[\sum_k \omega(k) - \left(\sum_k \omega(k) D(k) \right) \varepsilon \right] \\ &= V(1 - D^{VW} \varepsilon) \end{aligned}$$

where $D^{VW} = \sum_k \omega(k) D(k)$, and where $\omega(k) = v(k)/V$, which implies $\sum_k \omega(k) = 1$.

F.2 Proof of Proposition 2

From Proposition 1, we have

$$\tilde{\theta}_i = \theta_i(1 - D_i^{\theta}\varepsilon)$$

where θ_i and D_i^{θ} are the pre-shock financial wealth for household i and its duration, and $\tilde{\theta}_i$ is post-shock wealth after revaluation. Rearranging, we obtain

$$\frac{\tilde{\theta}_i}{\theta_i} = 1 - D_i^{\theta}\varepsilon.$$

Taking the covariance with respect to θ_i, we obtain

$$\mathrm{Cov}\left(\frac{\tilde{\theta}_i}{\theta_i}\right) = -\mathrm{Cov}\left(D_i^{\theta}, \theta_i\right)\varepsilon.$$

For $\varepsilon < 0$ this is positive if and only if $\mathrm{Cov}\left(D_i^{\theta}, \theta_i\right) > 0$.

For the sufficient condition, note that

$$D^{\theta,VW} = \int \frac{\theta_i}{\Theta} D_i^{\theta}\, di$$

where Θ is aggregate wealth. Manipulating this expression, we obtain

$$\begin{aligned}
D^{\theta,VW} &= \Theta^{-1}\int \theta_i D_i^{\theta}\, di \\
&= \Theta^{-1}\left\{\int \theta_i\, di \int D_i^{\theta}\, di + \int \left(\theta_i - \int_j \theta_j\, dj\right)\left(D_i^{\theta} - \int_j D_j^{\theta}\, dj\right) di\right\} \\
&= \Theta^{-1}\left\{\Theta D^{\theta,EW} + \int (\theta_i - \Theta)\left(D_i^{\theta} - D^{\theta,EW}\, di\right)\right\} \\
&= D^{\theta,EW} + \Theta^{-1}\mathrm{Cov}(\theta_i, D_i^{\theta}).
\end{aligned}$$

It follows immediately that $D^{\theta,VW} > D^{\theta,EW}$ if and only if $\mathrm{Cov}(D_i^{\theta}, \theta_i) > 0$.

Turning to part (b), let Θ^A denote the wealth of the top-α share of the wealth distribution, and let Θ^B denote the wealth of the bottom $1-\alpha$ share of the wealth distribution. Then the top-α wealth share is given by

$$S^{\alpha} = \frac{\Theta^A}{\Theta^A + \Theta^B} = \frac{1}{1 + \Theta^B/\Theta^A} = \frac{1}{1 + \exp(\log\Theta^B - \log\Theta^A)}. \tag{42}$$

Taking the derivative with respect to ε, we obtain

$$\frac{\partial S^{\alpha}}{\partial \varepsilon} = -\frac{1}{(1+\exp(\log\Theta^{B} - \log\Theta^{A}))^{2}} \times \exp(\log\Theta^{B} - \log\Theta^{A}) \times \left(\frac{\partial \log\Theta^{B}}{\partial \varepsilon} - \frac{\partial \log\Theta^{A}}{\partial \varepsilon}\right). \quad (43)$$

The first term is equal to $-(S^{\alpha})^{2}$, while is straightforward to verify from (42) that the second term satisfies

$$\exp(\log\Theta^{B} - \log\Theta^{A}) = \frac{1-S^{\alpha}}{S^{\alpha}}.$$

Substituting, we obtain

$$\frac{\partial S^{\alpha}}{\partial \varepsilon} = -S^{\alpha}(1-S^{\alpha})\left(\frac{\partial \log\Theta^{B}}{\partial \varepsilon} - \frac{\partial \log\Theta^{A}}{\partial \varepsilon}\right).$$

Applying a first-order approximation, and substituting the first approximation in (3), yields (5). It follows immediately that S^{α} increases for $\varepsilon < 0$ if and only if $D^{A} > D^{B}$.

F.3 Proof of Proposition 3

We begin by proving (a). First, consider any consumption plan $\{c_t\}$ that exactly satisfies the lifetime budget constraint (14), which is without loss of generality due to local non-satiation. Similarly, let $\{\tilde{c}_t\}$ be a consumption plan that exactly satisfies the post-shock lifetime budget constraint

$$\theta_0 = \mathbb{E}_0\left\{\sum_{t=0}^{\infty} \tilde{R}^{-t}(\tilde{c}_t - y_t)\right\} \quad (44)$$

Approximating the left hand side of (14) following a change in rates using (3) delivers

$$\log\tilde{\theta}_0 \simeq \log\theta_0 - D^{\theta} \times \varepsilon$$

which can be rearranged to yield

$$\log\theta_0 \simeq \log\tilde{\theta}_0 + D^{\theta} \times \varepsilon. \quad (45)$$

Similarly, we can approximate the right hand side of (14) to obtain

$$\log\mathbb{E}_0\left\{\sum_{t=0}^{\infty} \tilde{R}^{-t}(c_t - y_t)\right\} \simeq \log\mathbb{E}_0\left\{\sum_{t=0}^{\infty} R^{-t}(c_t - y_t)\right\} - D^{c-y} \times \varepsilon$$

which can be rearranged to yield

$$\log \mathbb{E}_0 \left\{ \sum_{t=0}^{\infty} R^{-t}(c_t - y_t) \right\} \simeq \log \mathbb{E}_0 \left\{ \sum_{t=0}^{\infty} \tilde{R}^{-t}(c_t - y_t) \right\} + D^{c-y} \times \varepsilon. \tag{46}$$

Substituting (45) and (46) into (14) and rearranging, we obtain

$$\log \tilde{\theta}_0 = \log \mathbb{E}_0 \left\{ \sum_{t=0}^{\infty} \tilde{R}^{-t}(c_t - y_t) \right\} + (D^{c-y} - D^{\theta}) \times \varepsilon.$$

Finally, applying (44) to the left hand side delivers

$$\log \mathbb{E}_0 \left\{ \sum_{t=0}^{\infty} \tilde{R}^{-t}(\tilde{c}_t - y_t) \right\} \simeq \log \mathbb{E}_0 \left\{ \sum_{t=0}^{\infty} \tilde{R}^{-t}(c_t - y_t) \right\} + (D^{c-y} - D^{\theta}) \times \varepsilon \tag{47}$$

Assume that $\varepsilon < 0$. Examining (44) shows that the pre-shock consumption plan $\{c_t\}$ is affordable under the post-shock budget constraint if and only if

$$\mathbb{E}_0 \left\{ \sum_{t=0}^{\infty} \tilde{R}^{-t}(c_t - y_t) \right\} \leq \mathbb{E}_0 \left\{ \sum_{t=0}^{\infty} \tilde{R}^{-t}(\tilde{c}_t - y_t) \right\}. \tag{48}$$

From equation (47), we know this is true if and only if $D^{\theta} \geq D^{c-y}$. If $D^{\theta} > D^{c-y}$, then we have that (48) holds strictly. In this case, the household is able to afford a new consumption plan $\{\tilde{c}_t\}$ that is at least as large as $(\{c_t\}, b_T)$ in histories, and is strictly larger after some histories. As a result, this household's consumption possibilities expand. A symmetric argument shows that if $D^{\theta} < D^{c-y}$, then the household can no longer afford its pre-shock consumption plan following the fall in rates, implying that its consumption opportunities contract. Last, if $D^{\theta} = D^{c-y}$ the household's consumption opportunities are unchanged, meaning that a post-shock consumption plan is affordable under (44) if and only if it was also affordable prior to the shock under (14). We note that a symmetric proof using (13) in place of (14) would prove part (a) of the proposition even in the case where mortality risk is nonzero.

Part (b) follows directly from the optimality condition (12), which as mortality risk goes to zero takes the familiar form

$$c^{-\gamma} = \beta R \mathbb{E}\left[(c')^{-\gamma} \,\middle|\, z \right]. \tag{49}$$

Since $\beta R = \tilde{\beta}\tilde{R}$, any consumption plan that satisfied (12) will still satisfy (49) under $\tilde{R}$ and $\tilde{\beta}$. Since the budget constraint (44) is satisfied by assumption, this completes the proof of part (b).

Part (c) follows directly from parts (a) and (b). From part (a) we know that if $D^{\theta} = D^{c-y}$ then the household will have exactly enough financial wealth post-shock to afford its pre-shock consumption plan, while part (b) implies that it will still find this plan optimal.

Alternatively, assume that $D^\theta \neq D^{c-y}$. If $D^\theta < D^{c-y}$ we know from part (a) that the original consumption plan is no longer affordable, and cannot be an equilibrium selection of the household. If $D^\theta > D^{c-y}$ then the household's consumption opportunity set has strictly expanded. In this case, the household should pick a sequence with $\tilde{c}_t > c_t$ in some states due to local non-satiation. Thus, if $D^\theta \neq D^{c-y}$ the pre-shock consumption plan is either infeasible or suboptimal.

Combined, we have shown that a household's consumption plan is unchanged following the shock if and only if $D^\theta = D^{c-y}$.

G Incomplete Markets Model with Aggregate Risk

This appendix sets up an infinite-horizon model where ex-ante identical households face both idiosyncratic and aggregate income risk. Interest rates are determined in equilibrium. It first shows how to map the model with stochastic growth into a stationary model without aggregate risk in the spirit of Bewley (1986). We define an equilibrium under high interest rates. We show how to compute the value of human wealth in a matter consistent with the aggregate resource constraint. The main results are in Section G.5. They characterize how a decline in rates affects wealth inequality.

G.1 Model Setup

G.1.1 Endowments

Time is discrete, infinite, and indexed by $t \in [0, 1, 2, ...)$. The aggregate endowment e follows the stochastic process:

$$e_t(z^t) = e_{t-1}(z^{t-1})\lambda_t(z_t)$$

where $\lambda(z_t)$ denotes the stochastic growth rate of the aggregate endowment and z_t the aggregate state. The history of aggregate shocks is denoted by $z^t = \{z_t, z_{t-1}, \cdots\}$. A share $\alpha_t(z_t)$ of the aggregate endowment is financial income (dividends), the remaining $1 - \alpha_t(z_t)$ share represents aggregate labor income.

Households are subject to idiosyncratic income shocks, whose history is denoted by $\eta^h = \{\eta_h, \eta_{h-1}, \cdots\}$. The η_h shocks are i.i.d. across households and persistent over time. The idiosyncratic shock process is assumed to be independent from the aggregate shock process. Labor income y follows the following stochastic process:

$$y_t(z^t, \eta^h) = \widehat{y}_t(z^t, \eta^h)(1 - \alpha_t(z_t))e_t(z^t),$$

The ratio of individual to aggregate labor income, which we refer to as the labor income share, is given by $\widehat{y}_t(z^t, \eta^h)$. We use (z^t, η^h) to summarize the history of aggregate and idiosyncratic shocks, and $\pi(z^t, \eta^h)$ to denote the unconditional probability that state s^t will be realized. If the aggregate and idiosyncratic states are independently distributed, then we can decompose state transition probabilities into an aggregate and idiosyncratic component:

$$\pi(z_{t+1}, \eta_{h+1}|z^t, \eta^h) = \phi(z_{t+1}|z^t)\varphi(\eta_{h+1}|\eta^h),$$

We make this assumption of independence between aggregate and idiosyncratic labor income risk in what follows.

G.1.2 Preferences

A household maximizes discounted expected utility:

$$U(c) = \sum_{j=1}^{\infty} \beta^j \sum_{(z^{t+j}, \eta^j)} \phi(z^{t+j}) \varphi(\eta^j) \frac{c(z^{t+j}, \eta^j)^{1-\gamma}}{1-\gamma},$$

where the coefficient of relative risk aversion $\gamma > 1$, and the subjective time discount factor $0 < \beta < 1$.

G.1.3 Technology

Households choose a portfolio of state-contingent bonds $a_t(z^t, \eta^h; z_{t+1})$ for each state z_{t+1}, which trade at prices $q_t(z^t, z_{t+1})$, and shares in the Lucas tree (stocks) $\sigma_t(z^t, \eta^h)$, which trade at price $\nu_t(z^t)$ satisfying the budget constraint:

$$c_t(z^t, \eta^h) + \sum_{z_{t+1}} a_t(z^t \eta^h; z_{t+1}) q_t(z^t, z_{t+1}) + \sigma_t(z^t, \eta^h) \nu_t(z^t) \leq W_t(z^t, \eta^h).$$

Household cash on hand W evolves according to:

$$\begin{aligned} W_{t+1}(z^{t+1}, \eta^{h+1}) &= a_t(z^t \eta^h; z_{t+1}) + \hat{y}_{t+1}(z^{t+1}, \eta^{h+1})(1 - \alpha(z_{t+1})) e_{t+1}(z^{t+1}) \\ &+ \left(\alpha(z_{t+1}) e_{t+1}(z^{t+1}) + \nu_{t+1}(z^{t+1}) \right) \sigma_t(z^t, \eta^h). \end{aligned}$$

Households are subject to state-uncontingent and state-contingent borrowing constraints:

$$\sum_{z_{t+1}} a_t(z^t \eta^h; z_{t+1}) q_t(z^t, z_{t+1}) + \sigma_t(s^t) \nu_t(z^t) \geq K_t(s^t)$$

$$a_t(z^t \eta^h; z_{t+1}) + \left(\alpha(z_{t+1}) e_{t+1}(z^{t+1}) + \nu_{t+1}(z^{t+1}) \right) \sigma_t(s^t) \geq M_t(s^t, z_{t+1})$$

where K and M denote generic borrowing limits. Incomplete risk sharing arises from two sources: the lack of an asset whose payoff depends on the idiosyncratic income shock η^t and the borrowing constraints.

G.2 Transformation into Stationary Economy

We can transform the stochastically growing economy into a stationary economy with a constant aggregate endowment following Alvarez and Jermann (2001); Krueger and Lustig (2010). To that end, define the stationary consumption allocations:

$$\hat{c}_t(z^t, \eta^h) = \frac{c_t(z^t, \eta^h)}{e_t(z^t, \eta^h)}, \forall (z^t, \eta^h),$$

the stationary transition probabilities and the stationary subjective time discount factor:

$$\begin{aligned}\widehat{\phi}(z_{t+1}|z^t) &= \frac{\phi(z_{t+1}|z^t)\lambda_{t+1}(z_{t+1})^{1-\gamma}}{\sum_{z_{t+1}}\phi(z_{t+1}|z^t)\lambda_{t+1}(z_{t+1})^{1-\gamma}},\\ \widehat{\beta}(z^t) &= \beta\sum_{z_{t+1}}\phi(z_{t+1}|z^t)\lambda_{t+1}(z_{t+1})^{1-\gamma}.\end{aligned}$$

Agents in the stationary economy with these preferences:

$$U(\widehat{c})(z^t,\eta^h) = \frac{\widehat{c}(z^t,\eta^h)^{1-\gamma}}{1-\gamma} + \sum_{z_{t+1}}\widehat{\beta}(z_{t+1},z^t)\widehat{\phi}(z_{t+1}|z^t)\sum_{\eta_{h+1}}\varphi(\eta_{h+1}|\eta^h)U(\widehat{c})(z^{t+1},\eta^{h+1}) \tag{50}$$

rank consumption plans identically as in the original economy.

When there is predictability in aggregate consumption growth, shocks to expected growth manifest themselves as time discount rate shocks in the stationary economy. If aggregate growth shocks are i.i.d. over time, then the stationary time discount factor is constant and given by:

$$\widehat{\beta} = \beta\sum_{z_{t+1}}\phi(z_{t+1})\lambda_{t+1}(z_{t+1})^{1-\gamma}. \tag{51}$$

This i.i.d. assumption on aggregate growth shocks is the assumption we will make, noting that it can easily be relaxed. In what follows, we also assume that aggregate factor shares are constant: $\alpha_t(z_t) = \alpha,\ \forall t$. By definition, labor income shares average to one across households:

$$\sum_{t_0\geq 1}\sum_{\eta^h}\varphi(\eta^h|\eta_0)\widehat{y}_t(\eta^h) = 1,\ \forall t.$$

G.3 Equilibrium in the Stationary Economy

In the stationary economy, agents trade a single risk-free bond and a stock. Both securities have the same returns in the absence of aggregate risk. The stock yields a dividend α in each period. Given initial financial wealth θ_0, interest rates $\widehat{R}_t$ and stock prices $\widehat{v}_t$, households choose consumption $\{\widehat{c}_t(\theta_0,\eta^h)\}$, bond positions $\{\widehat{a}_t(\theta_0,\eta^h)\}$, and stock positions $\{\widehat{\sigma}_t(\theta_0,\eta^h)\}$ to maximize expected utility (50) subject to the budget constraint:

$$\widehat{c}_t(\eta^h) + \frac{\widehat{a}_t(\theta_0,\eta^h)}{\widehat{R}_t} + \widehat{\sigma}_t(\theta_0,\eta^h)\widehat{v}_t = (1-\alpha)\widehat{y}_t(\eta^h) + \widehat{a}_{t-1}(\theta_0,\eta^{h-1}) + \widehat{\sigma}_{t-1}(\theta_0,\eta^{h-1})(\widehat{v}_t+\alpha),$$

and subject to borrowing constraints:

$$\frac{\widehat{a}_t(\theta_0,\eta^h)}{\widehat{R}_t} + \widehat{\sigma}_t(\theta_0,\eta^h)\widehat{v}_t \geq \widehat{K}_t(\eta^h), \quad \forall\eta^h$$

$$\widehat{a}_t(\theta_0,\eta^h)+\widehat{\sigma}_t(\theta_0,\eta^h)(\widehat{v}_{t+1}+\alpha)\geq \widehat{M}_t(\eta^h),\quad \forall\eta^h.$$

Definition 2. For a given initial distribution of wealth Θ_0, a Bewley equilibrium is a list of consumption choices $\{\widehat{c}_t(\theta_0,\eta^h)\}$, bond positions $\{\widehat{a}_t(\theta_0,\eta^h)\}$, and stock positions $\{\widehat{\sigma}_t(\theta_0,\eta^h)\}$ as well as stock prices $\widehat{v}_t$, and interest rates $\widehat{R}_t$ such that each household maximizes its expected utility, and asset markets and goods markets clear.

$$\sum_{t_0\geq 1}\int\sum_{\eta^h}\varphi(\eta^h|\eta_{t_0})\widehat{a}_t(\theta_0,\eta^h)d\Theta_0=0,$$

$$\sum_{t_0\geq 1}\int\sum_{\eta^h}\varphi(\eta^h|\eta_{t_0})\widehat{\sigma}_t(\theta_0,\eta^h)d\Theta_0=1.$$

$$\sum_{t_0\geq 1}\int\sum_{\eta^h}\varphi(\eta^h|\eta_{t_0})\widehat{c}_t(\theta_0,\eta^h)d\Theta_0=1.$$

In the stationary economy, the return on the aggregate stock equals the risk-free rate:

$$\widehat{R}_t=\frac{\widehat{v}_{t+1}+\alpha}{\widehat{v}_t}. \tag{52}$$

The equilibrium stock price equals the present discounted value of the dividends:

$$\widehat{v}_t=\sum_{\tau=0}^{\infty}\widehat{R}^{-1}_{t\to t+\tau}\alpha,$$

discounted at the cumulative gross risk-free rate, defined as: $\widehat{R}_{t\to t+T}=\Pi_{k=0}^{T}\widehat{R}_{t+k}$. Note that $\widehat{R}_{t\to t}=\widehat{R}_t$ and define $\widehat{R}_{t\to t-1}=1$. Since both assets, the stock and the risk-free bond, earn the same risk-free rate of return in the stationary economy, households are indifferent between them. This indifference extends to any other assets with different durations since interest rates are deterministic in the stationary economy.

G.3.1 Connection with the Equilibrium in the Growing Economy

We can map the equilibrium in the stationary economy into an equilibrium in the stochastically growing economy.

Proposition 4. If $\{\widehat{c}_t(\theta_0,\eta^h),\widehat{a}_t(\theta_0,\eta^h),\widehat{\sigma}_t(\theta_0,\eta^h)\}$ and $\{\widehat{v}_t,\widehat{R}_t\}$ are a Bewley equilibrium, then $\{c_t(\theta_0,z^t,\eta^h),a_t(\theta_0,z^t,\eta^h,z_{t+1}),\sigma_t(\theta_0,z^t,\eta^h)\}$ as well as asset prices $\{v_t(z^t),q_t(z^t,z_{t+1})\}$ are an equilibrium of the stochastically growing economy with:

$$\begin{aligned}
c_t(\theta_0, z^t, \eta^h) &= \widehat{c}_t(\theta_0, \eta^h) e_t(z^t) \\
a_t(\theta_0, z^t, \eta^h; z_{t+1}) &= \widehat{a}_t(\theta_0, \eta^h; z_{t+1}) e_t(z^t) \\
\sigma_t(\theta_0, z^t, \eta^h) &= \widehat{\sigma}_t(\theta_0, \eta^h) \\
\nu_t(z^t) &= \widehat{\nu}_t e_t(z^t) \\
q_t(z^t, z_{t+1}) &= \frac{\widehat{\phi}(z_{t+1})}{\lambda(z_{t+1})} \frac{1}{\widehat{R}_t}.
\end{aligned}$$

The proof is provided in Krueger and Lustig (2010).

The last equation in the proposition above implies the following relationship between the interest rate in the growing economy (R_t) and the stationary economy ($\widehat{R}_t$):

$$R_t = \left(\sum_{z_{t+1}} q_t(z^t, z_{t+1}) \right)^{-1} = \left(\sum_{z_{t+1}} \frac{\widehat{\phi}(z_{t+1})}{\lambda(z_{t+1})} \right)^{-1} \widehat{R}_t. \tag{53}$$

or, plugging in for $\widehat{\phi}(z_{t+1}|z^t)$:

$$\widehat{R}_t = \frac{E_t \left[\lambda_{t+1}^{-\gamma} \right]}{E_t \left[\lambda_{t+1}^{1-\gamma} \right]} R_t$$

G.3.2 Log-normal Growth

Consider a special case where the aggregate endowment growth rate λ_t is i.i.d. log-normally distributed:

$$log(\lambda_t) \sim \mathcal{N}(g, \sigma_\lambda^2)$$

Then:

$$E_t[\lambda_{t+1}^{-\gamma}] = E_t[\exp(-\gamma \log(\lambda_{t+1}))] = \exp(-\gamma g + 0.5\gamma^2 \sigma_\lambda^2)$$

and

$$E_t[\lambda_{t+1}^{1-\gamma}] = E_t[\exp((1-\gamma) \log(\lambda_{t+1}))] = \exp((1-\gamma) g + 0.5(1-\gamma)^2 \sigma_\lambda^2)$$

We obtain

$$\widehat{R} = \frac{\exp(-\gamma g + 0.5\gamma^2 \sigma_\lambda^2)}{\exp((1-\gamma) g + 0.5(1-\gamma)^2 \sigma_\lambda^2)} R_t = \frac{R}{G},$$

where

$$G = \exp(g + 0.5\sigma_\lambda^2 - \gamma \sigma_\lambda^2)$$

which recovers equation (15) in the main text. (Recall that the main text refers to the interest rate in the growing economy as R_g and to the interest rate in the stationary economy as R.)

Using lowercase letters to denote logs:

$$\hat{r} = r - g - 0.5\sigma_\lambda^2 + \gamma\sigma_\lambda^2$$

Changes in Interest Rates Now consider the relationship between the time-series change in the interest rate in the growing economy and the time-series change in the interest rate in the stationary economy. Denote the initial and new steady states by the subscripts 0 and T. Assume that the growth rate uncertainty does not change between steady states, but only the subjective time discount factor and/or the expected growth rate of the economy:

$$\hat{r}_T - \hat{r}_0 = (r_T - r_0) - (g_T - g_0)$$

The interest rate in the growing economy can be written, from the first-order condition, as:

$$r_t = -\log(\beta) + \gamma g - 0.5\gamma^2\sigma_\lambda^2.$$

Under the maintained assumption of no change in growth uncertainty, the change in interest rates in the growing economy is:

$$r_T - r_0 = -\log(\beta_T) + \log(\beta_0) + \gamma(g_T - g_0)$$

.

The change in rates in the stationary economy is lower by the change in the growth rate in the actual economy. We can also write this as:

$$\hat{r}_T - \hat{r}_0 = -(\log(\beta_T) - \log(\beta_0)) + (\gamma - 1)(g_T - g_0)$$

The change in the equilibrium interest rate in the stationary economy reflects either a change in the subjective time discount factor in the growing economy or a change in the expected growth rate of the economy or a combination of the two. The effect of a change in the expected growth rate on the interest rate depends on the inter-temporal elasticity of substitution (IES) γ^{-1}. If the IES is smaller than 1 ($\gamma > 1$), then a decrease in the expected growth rate results in a decrease in the interest rate; the income effect dominates the substitution effect.

Last, we can compute the impact on $\hat{\beta}$. Since

$$\hat{\beta} \quad = \quad \beta E_t[\lambda_{t+1}^{1-\gamma}] = \beta \exp\left\{(1-\gamma)g + 0.5(1-\gamma)^2\sigma_\lambda^2\right\}$$

In logs:

$$\log \hat{\beta} \quad = \quad \log \beta + (1-\gamma)g + \frac{1}{2}(1-\gamma)^2\sigma_\lambda^2$$

Under the maintained assumption that σ_λ^2 does not change between 0 and T, we have that:

$$\log \hat{\beta}_T - \log \hat{\beta}_0 = \log \beta_T - \log \beta_0 + (1-\gamma)(g_T - g_0)$$

implying that

$$\log(\hat{R}_T \hat{\beta}_T) - \log(\hat{R}_0 \hat{\beta}_0) = 0.$$

The change in $\log \hat{\beta}$ is of the same magnitude and opposite sign as the change in $\hat{r}$.

In the calibrated model, we envision the decline in interest rates in the data is driven by an increase in β so that:

$$\hat{r}_T - \hat{r}_0 = r_T - r_0 = -(\log(\beta_T) - \log(\beta_0)) = -(\log \hat{\beta}_T - \log \hat{\beta}_0) = 4.48\%.$$

G.4 Wealth Accounting

What is the right discount rate to use when measuring household wealth? If we want a wealth measure that can be aggregated, we have to use the same discount rate $\widehat{R}$ for all claims.

Proposition 5. At time 0, the financial wealth of each household equals the present discounted value of future consumption minus future labor income.

$$\theta_0 \quad = \quad \sum_{\tau=0}^{\infty} \sum_{\eta^\tau} \frac{\varphi(\eta^\tau)}{\widehat{R}_{0\to\tau-1}} \left(\widehat{c}_\tau(\eta^\tau) - (1-\alpha)\widehat{y}_\tau(\eta^\tau)\right)$$

The proposition follows directly from iterating forward on the one-period budget constraint. In this iteration, we take expectations over financial wealth in all future states using the objective probabilities of the idiosyncratic events $\varphi(\eta^\tau)$, and discount by the cumulative risk-free rate $\widehat{R}_{0\to\tau-1}$. Aggregate financial wealth in the economy in period 0 is given by:

$$\int \theta_0 d\Theta_0 = \int (\widehat{a}_{-1}(\theta_0) + \widehat{\sigma}_{-1}(\theta_0)\widehat{\nu}_0) d\Theta_0 = 0 + 1\widehat{\nu}_0,$$

where we have used market clearing in the bond and stock markets at time 0.

Aggregating the cost of the excess consumption plan across all households, using the fact that labor income shares average to 1, and imposing goods market clearing at time 0, we get:

$$\int \sum_{\tau=0}^{\infty} \widehat{R}^{-1}_{0\to\tau-1} \sum_{\eta^\tau} \varphi(\eta^\tau) \left(\widehat{c}_\tau(\eta^\tau) - (1-\alpha)\widehat{y}_\tau(\eta^\tau)\right) d\Theta_0 = \sum_{\tau=0}^{\infty} \widehat{R}^{-1}_{0\to\tau-1}\alpha = \widehat{v}_0.$$

The aggregate cost of households' excess consumption plan, or households' aggregate financial wealth, exactly equals the stock market value $\widehat{v}_0$, the only source of net financial wealth in the economy. This result relies on market clearing:

$$\int \sum_{\eta^\tau} \varphi(\eta^\tau) \left(\widehat{c}_\tau(\eta^\tau) - (1-\alpha)\widehat{y}_\tau(\eta^\tau)\right) d\Theta_0 = \alpha,$$

at each time t, because $\int \sum_{\eta^\tau} \varphi(\eta^\tau)\widehat{c}_\tau(\eta^\tau) d\Theta_0 = 1$ from market clearing, and the labor income shares sum to one as well.

The choice of the actual probability measure $\varphi(\cdot)$ and rate $\widehat{R}$ to compute an individual's human capital, the expected present discounted value of her labor income stream, may seem arbitrary. After all, claims to labor income are not traded in this model and markets are incomplete. The key insight is that, using any other pricing kernel to discount individual labor income and consumption streams may result in a value of aggregate financial wealth different from the value of the Lucas tree. To see this, consider using a distorted measure $\psi(\eta^\tau)\varphi(\eta^\tau)$ different from the actual measure $\varphi(\eta^\tau)$, where the household-specific wedges satisfy $\mathbb{E}_0[\psi_t] = 1, \forall t$. Under this different measure, the goods markets do not clear and the labor shares do not sum to one, unless the household-specific wedges do not covary with consumption and income shares:

Proposition 6. Wealth measures aggregate if and only if the following orthogonality conditions holds for the househehold-specific wedges and household consumption and income:

$$Cov_0\left(\psi_t, \widehat{c}_t\right) = 0, \quad Cov_0\left(\psi_t, \widehat{y}_t\right) = 0.$$

For all other wedge processes $\psi_t(\eta^\tau)$, the resource constraint is violated:

$$\int \sum_{\eta^\tau} \psi(\eta^\tau)\varphi(\eta^\tau) \left(\widehat{c}_\tau(\eta^\tau) - (1-\alpha)\widehat{y}_\tau(\eta^\tau)\right) d\Theta_0 \neq \alpha,$$

It is common in the literature to use the household's own IMRS to compute human capital. The household's IMRS is a natural choice because it ties the valuation of human wealth directly to welfare. However, this approach does not lend itself to aggregation. The wedges

$$\psi(\eta^{t+1}) = \frac{u'(\widehat{c}(\eta_{t+1}, \eta^t))}{u'(\widehat{c}_t(\eta_0))},$$

do not satisfy the zero covariance restrictions of the proposition. Imperfect consumption insurance implies that:

$$Cov_0(\psi_t, \widehat{c}_t) \leq 0, \quad Cov_0(\psi_t, \widehat{y}_t) \leq 0.$$

Proposition 7. If the cross-sectional covariance between the household-specific wedges and consumption is negative ($Cov_0(\psi_t, \widehat{c}_t) \leq 0$), then the aggregate valuation of individual wealth is less than the market's valuation of total wealth.

When aggregating, this pricing functional undervalues human wealth and therefore also total wealth.[33] In sum, while pricing claims to consumption and labor income using the household's IMRS is sensible from a welfare perspective, this approach does not lend itself to wealth accounting and aggregation.

G.5 Interest Rate Decline

We now analyze the main exercise of the paper, which is to let the economy undergo an unexpected and permanent decrease in the interest rate ("MIT shock"). We study the implications for inequality in financial wealth.

Since interest rates are endogenously determined, we generate the decline in the equilibrium real rate in the stationary model, $\widehat{R}$, through increase in the deflated subjective time discount factor, $\widehat{\beta}$. As discussed in Section G.3.2, the latter arises either from an increase in the subjective time discount factor in the economy with growth, β, a decline in the expected rate of growth of the aggregate endowment, $E[\lambda]$ (or equivalently G), or some combination of the two. We focus on the case of an increase in the subjective time discount factor, but the theoretical results go through if all or some of the change in interest rates comes from a decline in expected growth. We denote the equilibrium of the stationary economy under high interest rates with a hat ($\widehat{x}$) and the equilibrium of the stationary economy under low interest rates with a tilde ($\widetilde{x}$).

It is natural to ask whether the equilibrium consumption allocation $\{\widehat{c}_t(\theta_0, \eta^t)\}$ that prevailed in the economy with high rates is still an equilibrium after the change in interest rates. Given that the time discount factor of all agents increased by the same amount, there should be no motive to trade away from these allocations: $\widetilde{\beta}\widetilde{R} = \widehat{\beta}\widehat{R} = 1$. The following proposition shows that the old consumption allocation is indeed still an equilibrium in the low interest rate economy, provided that initial financial wealth is scaled up for every household.

Proposition 8. If the allocations and asset market positions $\{\widehat{c}_t(\theta_0, \eta^t), \widehat{a}_t(\theta_0, \eta^t), \widehat{\sigma}_t(\theta_0, \eta^t)\}$ and asset prices $\{\widehat{v}_t, \widehat{R}_t\}$ are a Bewley equilibrium in the economy with $\widehat{\beta}$ and natural borrowing limits

[33] Since the factor shares are constant, the consumption claim is in the span of traded assets. Financial wealth is the value of the Lucas tree, which equals α times the value of a claim to total consumption.

$\{\widehat{K}_t(\eta^t)\}$,

$$\widehat{K}_t(\eta^t) = \sum_{\tau=t}^{\infty} \widehat{R}_{t\to\tau-1}^{-1} \sum_{\eta^\tau|\eta^t} \varphi(\eta^\tau|\eta^t)(1-\alpha)\widehat{y}_\tau(\eta^\tau),$$

then the allocations and asset market positions $\{\widehat{c}_t(\widetilde{\theta}_0,\eta^t), \widehat{a}_t(\widetilde{\theta}_0,\eta^t), \widehat{\sigma}_t(\widetilde{\theta}_0,\eta^t)\}$ and asset prices $\{\widetilde{v}_t, \widetilde{R}_t\}$ will be an equilibrium of the economy with $\widetilde{\beta}$ and natural borrowing limits $\{\widetilde{K}_t(\eta^t)\}$,

$$\widetilde{K}_t(\eta^t) = \sum_{\tau=t}^{\infty} \widetilde{R}_{t\to\tau-1}^{-1} \sum_{\eta^\tau|\eta^t} \varphi(\eta^\tau|\eta^t)(1-\alpha)\widehat{y}_\tau(\eta^\tau),$$

asset prices are given by

$$\widetilde{\beta}\widetilde{R}_t \quad = \quad \widehat{\beta}\widehat{R}_t, \text{ and } \widetilde{v}_t = \sum_{\tau=0}^{\infty} \widetilde{R}_{t\to t+\tau}^{-1}\alpha,$$

and every household's initial wealth is adjusted as follows:

$$\widetilde{\theta}_0 = \theta_0 \frac{\sum_{\tau=0}^{\infty} \widetilde{R}_{0\to\tau}^{-1} \sum_{\eta^\tau} \varphi(\eta^\tau)\left(\widehat{c}_\tau(\eta^\tau) - (1-\alpha)\widehat{y}_\tau(\eta^\tau)\right)}{\sum_{\tau=0}^{\infty} \widehat{R}_{0\to\tau}^{-1} \sum_{\eta^\tau} \varphi(\eta^\tau)\left(\widehat{c}_\tau(\eta^\tau) - (1-\alpha)\widehat{y}_\tau(\eta^\tau)\right)}.$$

The proof is below.

Aggregate financial wealth undergoes an adjustment equal to the ratio of the price of two perpetuities:

$$\frac{\sum_{\tau=0}^{\infty} \widetilde{R}_{0\to\tau}^{-1}}{\sum_{\tau=0}^{\infty} \widehat{R}_{0\to\tau}^{-1}} = \frac{\widetilde{v}_0}{\widehat{v}_0}.$$

Intuitively, with lower interest rates, all asset prices are higher than in the high-rate economy. The Lucas tree becomes more valuable. A fraction $1-\alpha$ of this tree reflects aggregate human wealth, the remaining fraction is aggregate financial wealth.

Each individual's financial wealth adjustment differs, and depends on the expected discounted value of the same future excess consumption plan discounted at different rates. The higher one's expected future excess consumption, the larger the initial financial wealth adjustment needed to implement the old equilibrium allocation.

Characterizing Interest Rate Sensitivity Using Duration of Excess Consumption To a first-order approximation, i.e., for a small change in the interest rate, the adjustment in initial financial wealth needed for agents to keep their initial consumption plan is given by the duration of their planned consumption in excess of labor income. This is the duration households will need in their net financial assets in order to be fully hedged against interest rate risk.

Define the duration of a household's excess consumption plan at time 0, following the realization of the idiosyncratic labor income shock η_0, as follows:

$$D^{c-y}(\theta_0,\eta_0) \quad = \quad \frac{\sum_{\tau=0}^{\infty}\sum_{\eta^\tau|\eta_0}\tau\widehat{R}_{0\to\tau}^{-1}\varphi(\eta^t|\eta_0)\left(\widehat{c}_\tau(\eta^\tau|\eta_0)-(1-\alpha)\widehat{y}(\eta^\tau|\eta_0)\right)}{\sum_{\tau=0}^{\infty}\sum_{\eta^\tau|\eta_0}\varphi(\eta^t|\eta_0)\widehat{R}_{0\to\tau}^{-1}\left(\widehat{c}_\tau(\eta^\tau|\eta_0)-(1-\alpha)\widehat{y}(\eta^\tau|\eta_0)\right)}$$

The duration measures the sensitivity of the cost of its excess consumption plan to a change in the interest rate. In our endowment economy, aggregate consumption is fixed. We are interested in the valuation effects of interest rate changes.

The duration of the excess consumption claim equals the value-weighted difference of the duration of the consumption claim and that of the labor income claim:

$$D^{c-y} \quad = \quad \frac{P_0^c}{P_0^{c-y}}D^c - \frac{P_0^y}{P_0^{c-y}}D^y.$$

where $P_0^{c-y} = \theta_0$ is household financial wealth, P_0^y is human wealth, and P_0^c is total household wealth, the sum of financial and human wealth. Households with a high positive duration of excess consumption face a large increase in the cost of their consumption plan when interest rates go down, insofar that this increased cost is not offset fully by the increase in their human wealth.

The duration of the aggregate excess consumption claim, the aggregate duration for short, equals:

$$D^a = \frac{\sum_{\tau=0}^{\infty}\tau\widehat{R}_{0\to\tau}^{-1}}{\sum_{\tau=0}^{\infty}\widehat{R}_{0\to\tau}^{-1}}$$

This is the duration of a claim to aggregate consumption minus aggregate labor income, or equivalently to aggregate financial income. It is the duration of a perpetuity in the stationary economy. Recall that $\widehat{\nu}_0 = \nu_0 = \alpha\sum_{\tau=0}^{\infty}\widehat{R}_{0\to\tau}^{-1}$ denotes aggregate financial wealth.

Proposition 9. The aggregate duration equals the wealth-weighted average duration of households' excess consumption claims:

$$D^a = \int D^{c-y}(\theta_0,\eta_0)\frac{\theta_0}{\nu_0}d\Theta_0.$$

The proof follows directly from the definition of the household specific duration measure and market clearing.

The next proposition is the main result. It shows that, when households that are richer than average tend to have excess consumption plans of higher duration, then the (equally-weighted) average household's excess consumption plan duration is smaller than the aggregate duration.

Proposition 10. If $cov(\theta_0, D^{c-y}(\theta_0)) > 0$ then $\int D^{c-y}(\theta_0,\eta_0)d\Theta_0 \leq D^a$ and lower interest rates increase financial wealth inequality when households are fully hedged.

The proof follows from recognizing the following relationship between (cross-sectional) expectations and covariances:

$$D^a = \mathbb{E}\left[\frac{\theta_0}{\nu_0^a} D^{c-y}(\theta_0, \eta_0)\right] = \mathbb{E}\left[D^{c-y}(\theta_0, \eta_0)\right] + cov\left[\frac{\theta_0}{\nu_0}, D^{c-y}(\theta_0, \eta_0)\right].$$

The proposition says that under the covariance condition, if all households are perfectly hedged in their portfolio, then wealth inequality should increase when rates decline.

Ex-Ante Identical Households In this class of Bewley models, if agents are ex-ante identical, agents with low financial wealth have encountered a bad history of labor income shocks. If labor income is highly persistent, their labor income is low today and in the near future relative to labor income in the distant future (because of mean-reversion). This pattern makes the duration of their labor income stream high. But since the household is smoothing consumption inter-temporally, $D^c < D^y$. As a result, low-wealth agents tend to have low duration of their excess consumption plan. Conversely, rich agents have high labor income and high excess consumption duration. Consumption smoothing is the force that makes the assumption of a positive covariance between the level of financial wealth and the duration of excess consumption satisfied in a Bewley model where the only source of heterogeneity is income shock realizations. It follows immediately from Proposition 10 that the decline in rates (i) increases the cost of the excess consumption plan for the aggregate (per capita) value-weighted household by more than the cost for the equally-weighted average household, and (ii) increases financial wealth inequality. Put differently, in a model where all households are exactly equally well off after the change in rates by construction, i.e., they are perfectly hedged, financial wealth inequality should increase when rates go down.

Low-financial wealth households in a Bewley model have high-duration human wealth, which provides a natural interest rate hedge. High financial-wealth households have low-duration human wealth and need to increase financial wealth by more when rates decline to be able to afford the old consumption plan.

Ex-Ante Heterogeneous Households The insights of this normative proposition apply more broadly to a richer model with ex-ante heterogeneity across households, for example because agents go through a life cycle and differ by age.

Proposition 11. If $cov(\theta_t, D_t^{c-y}(\theta_{t_0})) > 0$ then the average duration is lower than the aggregate duration, $\sum_{t_0} \int D_t^{c-y}(\theta_0, \eta_0) d\Theta_{t_0} \leq D_t^a$ and lower interest rates increase financial wealth inequality when households are fully hedged.

We check this condition in the calibrated version of the model.

Real-world households may not be fully hedged, unlike the households in the Bewley model. The actual duration of the household's financial assets in the data, denoted D^θ, can differ from the duration of the excess consumption claim D^{c-y} in the model where households are fully hedged. Section 4 of the paper considers a calibrated life-cycle version of the Bewley model with overlapping generations to assess how well households are hedged against interest rate risk.

G.6 Proofs of Propositions in this Appendix

G.6.1 Proof of proposition 5

Proof. The one-period budget constraint:

$$\widehat{c}_t(\eta^t) + \frac{\widehat{a}_t(\eta^t)}{\widehat{R}_t} + \widehat{\sigma}_t(\eta^t)\widehat{v}_t = (1-\alpha)\widehat{y}_t(\eta^t) + \widehat{a}_{t-1}(\eta^{t-1}) + \widehat{\sigma}_{t-1}(\eta^{t-1})(\widehat{v}_t + \alpha),$$

can be restated, using equation (52), as:

$$\widehat{c}_t(\eta^t) - (1-\alpha)\widehat{y}_t(\eta^t) + \frac{\widehat{a}_t(\eta^t) + \widehat{\sigma}_t(\eta^t)(\widehat{v}_{t+1} + \alpha)}{\widehat{R}_t} = \widehat{a}_{t-1}(\eta^{t-1}) + \widehat{\sigma}_{t-1}(\eta^{t-1})(\widehat{v}_t + \alpha). \tag{54}$$

Rewriting (54) one period later:

$$\widehat{c}_{t+1}(\eta^{t+1}) - (1-\alpha)\widehat{y}_{t+1}(\eta^{t+1}) + \frac{\widehat{a}_{t+1}(\eta^{t+1}) + \widehat{\sigma}_t(\eta^{t+1})(\widehat{v}_{t+2} + \alpha)}{\widehat{R}_{t+1}} = \widehat{a}_t(\eta^t) + \widehat{\sigma}_t(\eta^t)(\widehat{v}_{t+1} + \alpha).$$

Multiply this equation by $\varphi(\eta_{t+1}|\eta^t)$ and sum across all states η_{t+1} to obtain:

$$\begin{aligned} & \sum_{\eta_{t+1}} \varphi(\eta_{t+1}|\eta^t) \left(\widehat{c}_{t+1}(\eta^{t+1}) - (1-\alpha)\widehat{y}_{t+1}(\eta^{t+1}) + \frac{\widehat{a}_{t+1}(\eta^{t+1}) + \widehat{\sigma}_t(\eta^{t+1})(\widehat{v}_{t+2} + \alpha)}{\widehat{R}_{t+1}} \right) \\ = \; & \widehat{a}_t(\eta^t) + \widehat{\sigma}_t(\eta^t)(\widehat{v}_{t+1} + \alpha), \end{aligned}$$

where we used the fact that $\sum_{\eta_{t+1}} \varphi(\eta_{t+1}|\eta^t) = 1$ on the right-hand side. Next, substitute this expression back into (54) to obtain:

$$\begin{aligned} & \widehat{c}_t(\eta^t) - (1-\alpha)\widehat{y}_t(\eta^t) + \widehat{R}_t^{-1} \sum_{\eta_{t+1}} \varphi(\eta_{t+1}|\eta^t) \left(\widehat{c}_{t+1}(\eta^{t+1}) - (1-\alpha)\widehat{y}_{t+1}(\eta^{t+1}) \right) \\ & + \widehat{R}_{t\to t+1}^{-1} \sum_{\eta_{t+1}} \varphi(\eta_{t+1}|\eta^t) \left(\widehat{a}_{t+1}(\eta^{t+1}) + \widehat{\sigma}_t(\eta^{t+1})(\widehat{v}_{t+2} + \alpha) \right) = \widehat{a}_{t-1}(\eta^{t-1}) + \widehat{\sigma}_{t-1}(\eta^{t-1})(\widehat{v}_t + \alpha). \end{aligned}$$

Define financial wealth, scaled by the aggregate endowment, as:

$$\widehat{\theta}_t = \widehat{a}_{t-1}(\eta^{t-1}) + \widehat{\sigma}_{t-1}(\eta^{t-1})(\widehat{v}_t + \alpha).$$

Continuing the forward substitution, we end up with the following expression:

$$\widehat{\theta}_t = \sum_{\tau=t}^{\infty} \widehat{R}_{t\to\tau-1}^{-1} \sum_{\eta^\tau|\eta^t} \varphi(\eta^\tau|\eta^t)\left(\widehat{c}_\tau(\eta^\tau) - (1-\alpha)\widehat{y}_\tau(\eta^\tau)\right).$$

where $\varphi(\eta^t|\eta^t) = 1$. Financial wealth must equal the cost of the household's excess consumption plan, where excess refers to the part not paid for with labor income. Noting that $e_0 = 1$ so that $\widehat{\theta}_0 = \theta_0$, writing this expression at time zero:

$$\theta_0 = \sum_{\tau=0}^{\infty} \widehat{R}_{0\to\tau-1}^{-1} \sum_{\eta^\tau} \varphi(\eta^\tau)\left(\widehat{c}_\tau(\eta^\tau) - (1-\alpha)\widehat{y}_\tau(\eta^\tau)\right)$$

recovers the statement of the proposition. □

G.6.2 Proof of Proposition 6

Proof. We note that the cross-sectional expectation of the product can be decomposed in the standard way:

$$\int \sum_{\eta^\tau} \varphi(\eta^\tau)\psi(\eta^\tau)\left(\widehat{c}_\tau(\eta^\tau)\right) d\Theta_0 = \mathbb{E}_0[\psi_\tau c_\tau] = \mathbb{C}ov_0[\psi_\tau, c_\tau] + \mathbb{E}_0\left[\psi_\tau\right]\mathbb{E}_0\left[c_\tau\right].$$

If the orthogonality condition is satisfied, then the following result obtains:

$$\int \sum_{\eta^\tau} \varphi(\eta^\tau)\psi(\eta^\tau)\left(\widehat{c}_\tau(\eta^\tau)\right) d\Theta_0 = \mathbb{E}_0[\psi_\tau c_\tau] = \mathbb{E}_0[\psi_\tau]\mathbb{E}_0[c_\tau] = \mathbb{E}_0[c_\tau] = 1,$$

because $\mathbb{E}_0[\psi_t] = 1$. □

G.6.3 Proof of Proposition 7

Proof. This inequality $0 \geq Cov(\psi_t, \widehat{c}_t)$ directly implies that the following inequalities obtain:

$$\int \sum_{\tau=0}^{\infty} \widehat{R}_{0\to\tau-1}^{-1} \sum_{\eta^\tau} \varphi(\eta^\tau)\psi(\eta^\tau)\widehat{c}_\tau(\eta^\tau) d\Theta_0 \leq \int \sum_{\tau=0}^{\infty} \widehat{R}_{0\to\tau-1}^{-1} \sum_{\eta^\tau} \varphi(\eta^\tau)\widehat{c}_\tau(\eta^\tau) d\Theta_0 = \sum_{\tau=0}^{\infty} \widehat{R}_{0\to\tau-1}^{-1},$$
$$\int \sum_{\tau=0}^{\infty} \widehat{R}_{0\to\tau-1}^{-1} \sum_{\eta^\tau} \varphi(\eta^\tau)\psi(\eta^\tau)\widehat{y}_\tau(\eta^\tau) d\Theta_0 \leq \int \sum_{\tau=0}^{\infty} \widehat{R}_{0\to\tau-1}^{-1} \sum_{\eta^\tau} \varphi(\eta^\tau)\widehat{y}_\tau(\eta^\tau) d\Theta_0 = \sum_{\tau=0}^{\infty} \widehat{R}_{0\to\tau-1}^{-1}.$$

As a result, this new measure implies an aggregate value of individual wealth that falls short of total wealth, $\sum_{\tau=0}^{\infty} \widehat{R}_{0\to\tau-1}^{-1}$. Note that even though this claim to total consumption is itself not traded, the Lucas tree is a claim to α of the same cash flow stream. The market value of the Lucas tree is $\alpha\sum_{\tau=0}^{\infty} \widehat{R}_{0\to\tau-1}^{-1}$, and hence the value of total wealth has to be $\sum_{\tau=0}^{\infty} \widehat{R}_{0\to\tau-1}^{-1}$. □

G.6.4 Proof of proposition 8

Proof. An unconstrained household's Euler equation in the high-growth economy is given by:

$$1 = \widehat{\beta}\widehat{R}_t \sum_{\eta_{t+1}} \varphi(\eta_{t+1}|\eta^t) \frac{u'(\widehat{c}(\eta_{t+1}, \eta^t))}{u'(\widehat{c}_t(\eta^t))}.$$

This Euler equation is satisfied because the allocations and prices constitute a Bewley equilibrium in the high-growth economy. This household's Euler equation in the new economy with lower interest rates is still satisfied at the old consumption allocation. This can be seen by plugging in the new equilibrium interest rates:

$$\widetilde{R}_t\widetilde{\beta} = \widehat{\beta}\widehat{R}_t,$$

to recover the unconstrained household's Euler equation in the low-growth economy:

$$1 = \widetilde{\beta}\widetilde{R}_t \sum_{\eta_{t+1}} \phi(\eta_{t+1}|\eta_t) \frac{u'(\widehat{c}(\eta^t, \eta_{t+1})}{u'(\widehat{c}_t(\eta^t))}.$$

We allocate the following amount of financial wealth at time 0 to ensure the household can afford the same consumption plan:

$$\widetilde{\theta}_0(\theta_0, \eta_0) \quad = \quad \sum_{\tau=0}^{\infty} \widetilde{R}^{-1}_{0\to\tau-1} \sum_{\eta^\tau} \varphi(\eta^t) \left(\widehat{c}_\tau(\eta^\tau) - (1-\alpha)\widehat{y}_\tau(\eta^\tau)\right).$$

Aggregating this initial financial wealth across households:

$$\int \widetilde{\theta}_0 d\Theta_0 \quad = \quad \alpha \sum_{\tau=0}^{\infty} \widetilde{R}^{-1}_{0\to\tau} = \tilde{v}_0,$$

where we have used the goods market clearing condition and the definition of labor income shares. The last equation shows that the new allocation of initial financial wealth uses up all aggregate financial wealth in the economy. Finally, note that the natural borrowing constraints are not binding in the high-growth economy. They remain non-binding in the low-growth economy because consumption is nonnegative. Hence, the allocations are feasible, and they satisfy the sufficient conditions for optimality. □

H Additional Model Results

Table H1: Change in Inequality, Transition Experiment (Levels)

	Data		Model		
	Initial	After	Initial	Repriced	Comp.
Top-10% share FW	62.4%	70.8%	62.3%	70.2%	60.8%
Top-1% share FW	23.8%	35.1%	35.3%	41.7%	35.9%
Gini FW	0.772	0.826	0.710	0.770	0.690
Top-10% share HW	–	–	30.0%	31.1%	31.1%
Top-1% share HW	–	–	15.5%	13.3%	13.3%
Gini HW	–	–	0.405	0.470	0.470
Top-10% share TW	–	–	35.5%	36.0%	33.9%
Top-1% share TW	–	–	19.6%	17.7%	16.9%
Gini TW	–	–	0.405	0.476	0.456

Note: Top-10% share, Top-1% share and Gini coefficient of financial wealth are measured in the WID data. For the Initial period we use the value in 1983. For the After period we use the value in 2019. More details on the data computations are provided in Appendix B.1. For model results, the columns represent the pre-shock wealth distribution ("Initial"), the repriced distribution ("Repriced"), and the compensated distribution ("Comp").

Figure H1: Lorenz Curves

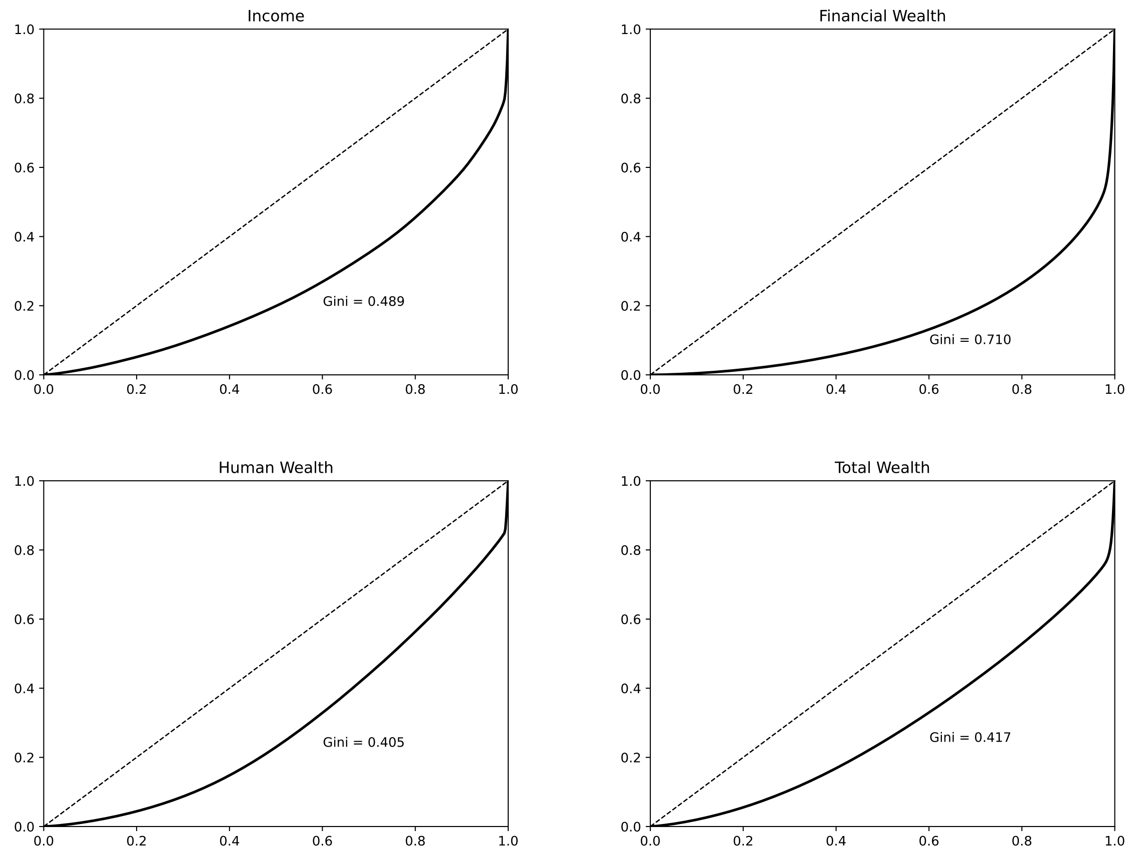

Note: This figure plots the Lorenz curve for each variable, obtained from a long simulation of the model.

Figure H2: Scatterplots by Age: Medians

(a) Compensated vs. Original (b) Compensated vs. Repriced

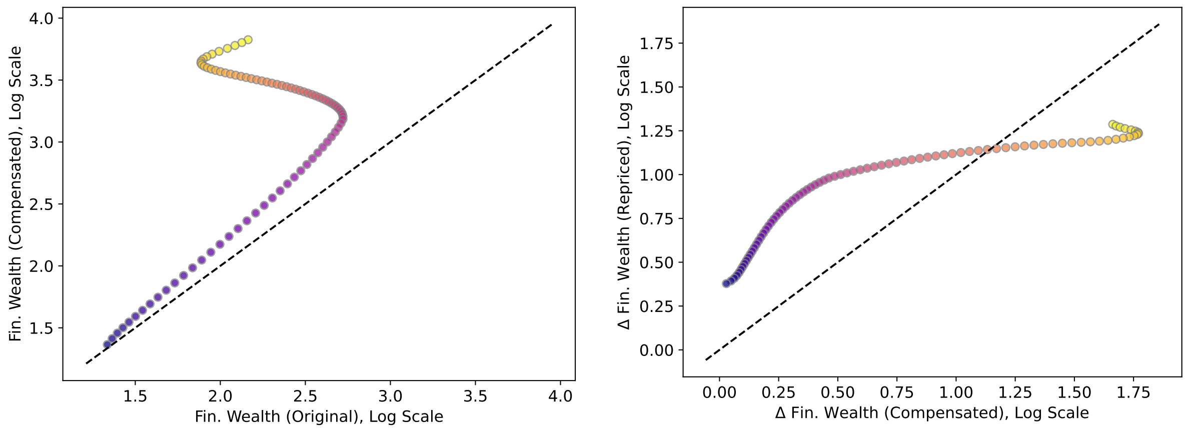

Note: Panel (a) plots the distribution of original financial wealth against the distribution of compensated financial wealth by age. Each dot represents the median value for one year of age, with the lightest (yellow) dots representing the youngest agents and the darkest (purple) dots representing the oldest agents. Both variables are plotted using the transform $\log(1 + x)$. Panel (b) similarly plots medians for one-year bins of the change in financial wealth under repricing, compared to the change in financial wealth under the compensated distribution, both using the transform $\log(1 + x)$ before differencing. The dashed line represents equality between the x and y axes.

Figure H3: Binscatters by Wealth, Controlling for Age

(a) Compensated vs. Original (b) Compensated vs. Repriced

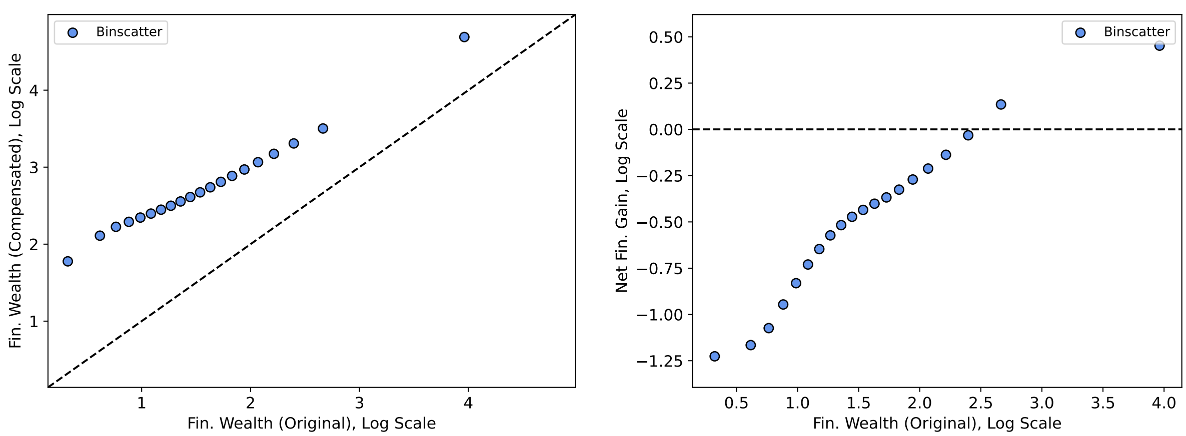

Note: Panel (a) plots the distribution of original financial wealth against the distribution of compensated financial wealth by age. Each dot represents 5% of the original financial wealth distribution. Both variables are plotted using the transform $\log(1 + x)$. Panel (b) similarly plots medians for 5% bins of the change in financial wealth under repricing, compared to the change in financial wealth under the compensated distribution, both using the transform $\log(1 + x)$ before differencing. The dashed line represents equality between the x and y axes.