Financial and Total Wealth Inequality with Declining Interest Rates*

Daniel L. Greenwald Matteo Leombroni Hanno Lustig
Stijn Van Nieuwerburgh

September 2, 2022

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 inequality can be explained by the decline in rates, and what the implications are for inequality in total wealth (lifetime consumption). We estimate the exposure of financial portfolios to interest rates at the household level to show that there is enough heterogeneity in portfolio revaluations to explain 75% of the rise in financial wealth inequality since the 1980s. A standard incomplete markets model calibrated to these data implies that declining rates are not consumption neutral. Instead, the low-wealth young lose, while the high-wealth old gain.

JEL: E21, E25, E44, G12
Keywords: wealth inequality, interest rates, duration

---

*First Version March 2021. Greenwald: Finance Department, MIT Sloan School of Management; greenw2@gmail.com. Leombroni: Department of Economics, Stanford University, Stanford CA 94305; leombm@stanford.edu; Lustig: Department of Finance, Stanford Graduate School of Business, Stanford CA 94305 and NBER; hlustig@stanford.edu; https://people.stanford.edu/hlustig/. 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, Ben Moll, Sylvain Catherine, Peter DeMarzo, 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 ECB, and INSEAD.

1 Introduction

Over the post-war period, interest rates and financial wealth inequality have displayed a remarkable negative correlation. Figure 1 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, obtained from a statistical asset pricing model (see Section 3 for details). 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 (yields fall) 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 4.82% in the 1980s, before falling to an average of 0.34% 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.

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

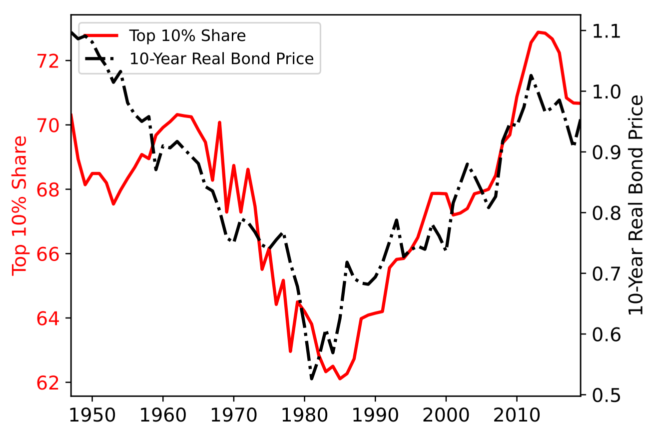

Note: The red solid line displays the top-10% financial wealth share for the United States, obtained annually from 1947 until 2019 from the World Inequality Database. The black dash-dot line displays an estimate of the 10-year real bond price, obtained from a dynamic affine term structure model, estimated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

Since discount rates have a direct link to the values of financial assets, a natural hypothesis is that falling interest rates result in rising financial wealth inequality, and vice versa. 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" or actually influence consumption and hence welfare is far from clear. In this paper, we study the link between real interest rates and wealth inequality to answer two research questions. First, what share of the rise in financial wealth inequality displayed can be quantitatively explained by falling interest rates? Second, what are the implications for inequality in total wealth — defined as the present value of lifetime consumption — that actually determines welfare?

To answer these questions, we combine a set of novel empirical estimates with a quantitative structural model. For our first question, we directly measure the exposure of financial wealth portfolios to changes in real interest rates at the household level, allowing us to directly estimate the effect of asset revaluations on financial wealth inequality. For our second question, we use a calibrated life-cycle consumption-savings model to compute the implied exposure of each household's consumption plan to the change in rates. We show that the extent to which this exposure aligns with or diverges from the exposure of the household's financial wealth portfolio determines the ultimate consequences of the change in rates for welfare inequality.

We begin by summarizing the exposure of agents' portfolios to a change in interest rates using the cash flow duration of those portfolios. A sufficient condition for whether a fall in rates will increase financial wealth inequality is that the aggregate (value-weighted) duration of financial wealth exceeds the average (equal-weighted) duration of financial wealth. To check this condition, and quantify the impact of a fall in rates on financial wealth inequality, we turn to the data. We use microdata from the Survey of Consumer Finances (SCF) to characterize households' portfolio allocations across asset classes. We then use an auxiliary asset pricing model, estimated to fit asset prices and cash flows quarter-by-quarter, to compute the cash flow duration of each asset class. To assign a duration to private business wealth — a key portfolio component for the wealthiest households — we use subcategory-level data on private business types to separate small and potentially stagnant private businesses from larger and faster-growing ones. Combined, we are able to characterize the distribution of financial wealth durations across the population. We find that the aggregate, or wealth-weighted average duration of financial wealth in the 1980s US economy was 18.31, compared to an average, or equally-weighted average duration of 14.83. We show that this gap between the aggregate and average duration implies that a fall in interest rates will raise measured financial wealth inequality.

We observe substantial heterogeneity in financial durations by wealth level and age. High-wealth households have higher financial durations, as their portfolios are more heavily weighted toward high-duration assets like housing, private business, and stock market wealth, and less weighted toward deposit-like assets. Conditional on wealth, financial durations are declining in age. This heterogeneity in financial duration is a new empirical finding, and crucial for the response of financial inequality to the decline in long-term real rates.

To study the quantitative implications of these measured durations for financial wealth inequality, as well as future consumption (total wealth) inequality, we turn to a calibrated life-cycle model of the U.S. economy. The model features a rich idiosyncratic income risk process, calibrated using Panel Survey of Income Dynamics data, paired with 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 flexibly match our empirical estimates by wealth bin and age. For our main experiment, we initialize the model at a long-term real interest rate of 4.94%, the level we estimate to have prevailed in 1983. We then let the model undergo the observed sequence of interest rate changes matching the 1983-2019 time path. Households assume these changes are unexpected and permanent.

First, we answer the positive question: what happens to financial wealth inequality in the calibrated model after rates decline unexpectedly? To do so, we compute the implied financial wealth distribution after revaluing all assets using the new interest rate and our distribution of durations fitted to the data, which we denote the *repriced distribution*. The repriced distribution in 2019 has a substantially higher top-10% financial wealth share, increasing from the 1983 levels by 6.2pp in our preferred model calibration, thereby explaining 75% of the observed 8.3pp rise in the data. The model also predicts substantial increases in the top-1% share and the Gini coefficient of financial wealth on the order of those observed in the data. Alternative measures for the duration of private business wealth yield a range of increases in the top-10% wealth share from 4.3pp to 9.5pp, explaining between 52% and 114% of the observed rise in the data. Adding the rise in (top) income inequality to the decline in rates leads the model to account for the remaining rise in the top-10% financial wealth share. We thus find robust evidence that this repricing mechanism explains the majority of the rise in financial wealth inequality since the 1980s.

Having established the strength of this mechanism, we turn to our second, normative question: what are the implications of this change for total wealth or consumption inequality? Since wealth gains can occur "on paper" without translating into actual consumption gains and losses, establishing a benchmark is essential. We compute this benchmark as the change in financial wealth that would be required under the new, lower interest rate, so that each household would be able to afford their prior consumption plans formed under the old interest rate. We denote this object as the *compensated* financial wealth distribution.[1]

To be completely hedged and keep its consumption unchanged, a household's financial portfolio duration needs to match the duration of its excess consumption plan, defined as its future consumption minus labor income. We find that attaining this compensated distribution requires a rightward shift in the wealth distribution following a decline in rates. Compensation requires a large increase in financial wealth, much of it accruing to top-1% of the financial wealth distribution. These results imply that large increases in the financial wealth of the richest individuals need not necessarily imply any actual change in consumption beyond "paper" gains.

In practice, however, we find that the aggregate duration of U.S. financial wealth is smaller than the duration of a typical household's excess consumption plan. As a result, to ensure that all

[1] Since we change the time discount rate alongside with the interest rate so as to keep their product constant, we show that the old equilibrium consumption allocation is not just feasible but remains optimal.

households could afford their pre-shock consumption plan, we would need to see financial wealth inequality *decline*. In particular, the top-10% financial wealth share of the compensated distribution falls by 3.1pp. The Gini coefficient and top-1% share of the compensated distribution would also need to fall in order to hold agents' consumption plans fixed. In summary, the compensated financial wealth distribution is less unequal than either the original or the repriced distribution.

This large deviation between what is observed under repricing and under compensation implies that household portfolios provide far from perfect hedging. Instead, repricing meaningfully reallocates consumption possibilities across the population and affects welfare. Both wealth and welfare inequality increases with declining rates. This normative result arises as the balance of two opposing forces: luck and the life cycle. A household that experienced a sequence of lucky income shocks has both high financial wealth and a high duration of excess consumption. Its duration of human wealth is low, because of mean reversion in income, while its duration of consumption is higher because of consumption smoothing. Persistent but mean-reverting income shocks create a positive cross-sectional covariance between the level of financial wealth and the excess consumption duration.

The second force is life-cycle dynamics. The young, who plan to save in middle age and dissave in retirement, have a very high duration of excess consumption, because they keep consuming after retirement. This renders their consumption plans much more expensive as rates fall. Young households are forced later to buy financial assets at higher prices. The life-cycle force creates a negative cross-sectional covariance between the level of financial wealth and the excess consumption duration, since young households tend to have low financial wealth.

The first force predicts that a fall in rates should lead to a rise in inequality in the compensated financial wealth distribution, which the life-cycle force prescribes a reduction in wealth inequality to keep all households equally well off after the change in rates. The calibrated model shows that the life-cycle force dominates. Specifically, the young struggle to accumulate enough wealth for retirement under low rates, and face a serious contraction in their consumption possibilities. We find that this degree of under-hedging is decreasing with age. The oldest agents are actually over-hedged, and see their consumption possibilities expand due to large capital gains on their wealth positions. In addition, we find that consumption gains are increasing in households' financial wealth position, conditional on age. We also document how the same historical change in interest rates affected lifetime consumption of different birth cohorts, showing that cohorts born before 1950 largely gained from the historical change in interest rates, while cohorts born after 1950 face substantial losses in the present value of lifetime consumption.

Finally, we study how interest rate changes affect total wealth, the sum of financial and human wealth. We find a lower level and a smaller rise in total wealth inequality than in financial wealth inequality. This is because human wealth is less unequally distributed and lower rates benefit younger, high-human wealth, low financial-wealth households more.

The rest of the paper is organized as follows. Section 2 discusses the related literature. Section 3 shows that top-wealth shares track interest rates in the U.S., U.K., and France and documents the main novel empirical fact on the cross-sectional correlation between financial wealth and the duration of financial wealth. Section 4 sets up the model. Section 5 contains the main theoretical results. Section 6 calibrates the model, while Sections 7 and 8 discuss the main quantitative results for the repriced and compensated wealth distributions, respectively. Section 9 contains the effect of repricing on various birth cohorts. Section 10 concludes. The paper also includes a comprehensive appendix. Appendix A provides an auxiliary asset pricing model used to infer real interest rates and durations of the components of financial wealth. Appendix B contains details on data sources and construction. Appendix C contains the proofs of the propositions in the main text. Appendix D sets up and solves an equilibrium model with aggregate risk and endogenous interest rates that generalizes the results in the main text. Appendix E displays supplementary model results.

2 Related Literature

A large strand of recent literature documents the evolution of income and financial wealth inequality over the past century (Piketty and Saez, 2003; Piketty, 2015; Alvaredo, Chancel, Piketty, Saez, and Zucman, 2018). Most of the evidence suggests that financial wealth inequality has increased in many countries over the past decades. Zucman (2019) reviews the empirical and Benhabib and Bisin (2018) the theoretical literatures on wealth inequality.

Much of the literature on wealth inequality adopts a backward-looking approach and explores the connection between past returns and current wealth. This literature has argued that high past rates of return and heterogeneity therein helps account for the increase in financial wealth inequality (Piketty and Zucman, 2015; Fagereng, Guiso, Malacrino, and Pistaferri, 2020; Bach, Calvet, and Sodini, 2020; Hubmer, Krusell, and Smith, 2020; Cox, 2020).

But wealth is also the current value of the household's future consumption stream. Human wealth is the value of future labor income and financial wealth is the value of future consumption minus income. We bring an asset pricing perspective to the discussion on inequality. We impute a valuation by discounting future cash flows. When discount rates decline, households need more wealth to finance the same consumption stream. Households that have mostly human wealth are better hedged, all else equal. Households with mostly financial wealth need enough duration in their portfolio in order to finance future consumption. To keep consumption shares unchanged, a decline in real rates needs to entail a reallocation of financial wealth towards those households who rely mostly on their (current and future) financial wealth to finance future consumption.

Discount rates matter. In a simple partial equilibrium model, Moll (2020) explains that small discount rate-induced changes in the wealth distribution may have smaller welfare effects than cash flow-induced changes. Recently, Catherine, Miller, and Sarin (2020) show that discounting social security transfers at time-varying discount rates has quantitatively important implications for wealth inequality. The auxiliary asset pricing model in Appendix A shows declines in expected real returns that are broad-based, not only on bonds but also on stocks and housing.

Our paper is related to recent work by Auclert (2019), who explores the effect of cross-sectional variation in the duration of households' financial assets for the effectiveness of monetary policy. We consider a setting with aggregate risk, we develop measures of household duration based on a no-arbitrage dynamic asset pricing model and household financial portfolios, and we assess quantitatively the extent to which households have hedged their consumption plan against interest rate innovations. In earlier work, Doepke and Schneider (2006) focus on the distributional consequences of inflation and Glover, Heathcote, Krueger, and Rios-Rull (2020) on the distributional effects of the Great Recession. Our work instead focuses on the distributional effects of changes in long-term real rates. Gomez and Gouin-Bonenfant (2020) study the effects of lower interest rate on the cost of raising new capital for entrepreneurs, linking the decline in interest rates to the rise in wealth inequality through a different channel.

By emphasizing total wealth (inequality), of which human wealth (inequality) forms a very significant component, our work contributes to the literature on measuring wealth (inequality). Our paper provides new and detailed statistics on the duration of financial wealth for U.S. households. Related, Kuhn, Schularick, and Steins (2020) study how housing and equity portfolio shares differ across the wealth distribution and result in differing financial wealth dynamics for the middle class and the top of the financial wealth distribution. Recent work discusses the measurement of private business income and wealth (Kopczuk, 2017; Saez and Zucman, 2016; Piketty, Saez, and Zucman, 2018; Smith, Yagan, Zidar, and Zwick, 2022.; Kopczuk and Zwick, 2020). In our theoretical work, we sidestep this issue by recognizing that financial wealth is the present discounted value of the future stream of consumption minus labor income. In our empirical work, we infer the duration of private business wealth by combining information from the SCF and the duration of small stocks, and perform various exercises to gauge robustness of the results.

Our positive conclusions regarding the change in wealth inequality are not sensitive to the source of the decline in interest rates. Our normative conclusions are based on a standard Bewley model (Appendix D) where the interest rate is endogenous, and the decline in rates arises from a change in the time discount factor, the expected growth rate of the economy, or growth uncertainty. The theoretical results are invariant to the mix of these three sources of rate decline. When calibrating the model, we attribute the change to the discount factor. A long literature, complementary to our work, studies these and other mechanisms for the decline in interest rates.[2]

[2]One part of the literature proposes candidates for a slowdown in economic growth: demographics (Summers, 2014; Eggertsson and Mehrotra, 2014; Eichengreen, 2015), a productivity slowdown due to a plateau in educational attainment or diminishing technological progress (Gordon, 2017), government spending that leads to depressed future aggregate demand (Mian, Straub, and Sufi, 2020), a decline in competition (Gutiérrez and Philippon, 2017), a decline in desired investment due to lower relative prices of capital goods (Rachel and Smith, 2017), a global saving glut

3 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 detail a strong time-series correlation between the evolution of long-term real interest rates and wealth inequality, not only in the U.S. but also in the United Kingdom and France. Third, we show the main novel empirical fact of the paper: household financial wealth portfolios have highly heterogeneous durations that correlate positively with the level of financial wealth.

3.1 Decline in Real Rates

We start by documenting a broad-based decline in expected returns across all major asset classes. To do so, we develop an auxiliary no-arbitrage asset pricing model, with details provided in Appendix A. The 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 yield averaged 4.82% in the 40 quarters of the 1980s decade and 0.34% in the 2010s decade.[3] The asset pricing model shows similarly large declines in expected real returns on the aggregate stock market and on housing wealth, as shown in Table 1. Other stock indices such as value and infrastructure stocks show larger declines, while growth and small stocks show smaller declines. However, the decline in expected returns is robust across asset classes and broad-based.

3.2 Increased Wealth Inequality

We next provide additional evidence for the strong positive co-movement between financial wealth inequality and long-term real bond prices displayed previously in Figure 1. We show that this relationship is robust across wealth measures, interest rate measures, and countries. Figure 2 shows the wealth share of the top-10% of the population in the left panels and the wealth share of the top-1% of the population in the right panels. The top panel is for the U.S., the middle panel for the U.K., and the bottom panel is for France. Our main source of wealth inequality data is the World

and/or shortage of safe assets (Bernanke et al., 2005; Caballero, Farhi, and Gourinchas, 2008), among others. A rise in income inequality could be the origin of lower interest rates. Hubmer et al. (2020) show that a rise in earnings risk actually lowers wealth inequality as it strengthens precautionary savings motives meaningfully for all but the richest households. A rise in top-income inequality, in contrast, can increase wealth inequality. Mian et al. (2020) argue that the rich have a higher propensity to save than the poor; Fagereng, Blomhoff Holm, Moll, and Natvik (2019) provide empirical evidence consistent with this from Norway. This reduces aggregate demand and the real rate of interest in the wake of an exogenous increase in income inequality, for example, due to skill-biased technological change. Lower tax progressivity could lead to more saving by the rich, more aggregate wealth, and lower rates (Hubmer et al., 2020). However, Heathcothe, Storesletten, and Violante (2020) argue that once transfers and actual tax paid are considered, the U.S. tax system has not become less progressive.

[3]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.

Table 1: Expected Real Returns Decade Averages

Asset	1980s	2010s	Decline
Ten-year real bond yield	4.82%	0.34%	4.48%
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.

Inequality Database. For the U.S., we also plot the wealth shares constructed from the Survey of Consumer Finances (SCF).Each panel also plots the price of a thirty-year real annuity, computed either from nominal yields and inflation or alternatively from our auxiliary asset pricing model. Construction details are in Appendix B.1. The sample is 1947-2019.[4] The 30-year real annuity is an intuitive measure of the long-run cost of a constant consumption bundle in real terms, which is inversely related with long-term real interest rates.

For both inequality measures, there is a strong positive correlation between top wealth shares and the price of a long-term real annuity. 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 share rises by 8.3pp (11.3pp) as the cost of the annuity more than doubles. In the U.K. and in France we observe the same general pattern, although the top-10% share falls by more in France and in the U.K. 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. Rachel and Smith (2017) show that the decline in the real rate has 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.

3.3 Household Heterogeneity in Financial Duration

Variation in real rates only matters for wealth inequality if households' portfolios have heterogeneous exposure to real rates. We next document this heterogeneity in the data.

Every asset has its own duration, which measures the sensitivity of the asset's market price to interest rates. By combining the actual portfolio shares of a household with the duration of the

[4]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

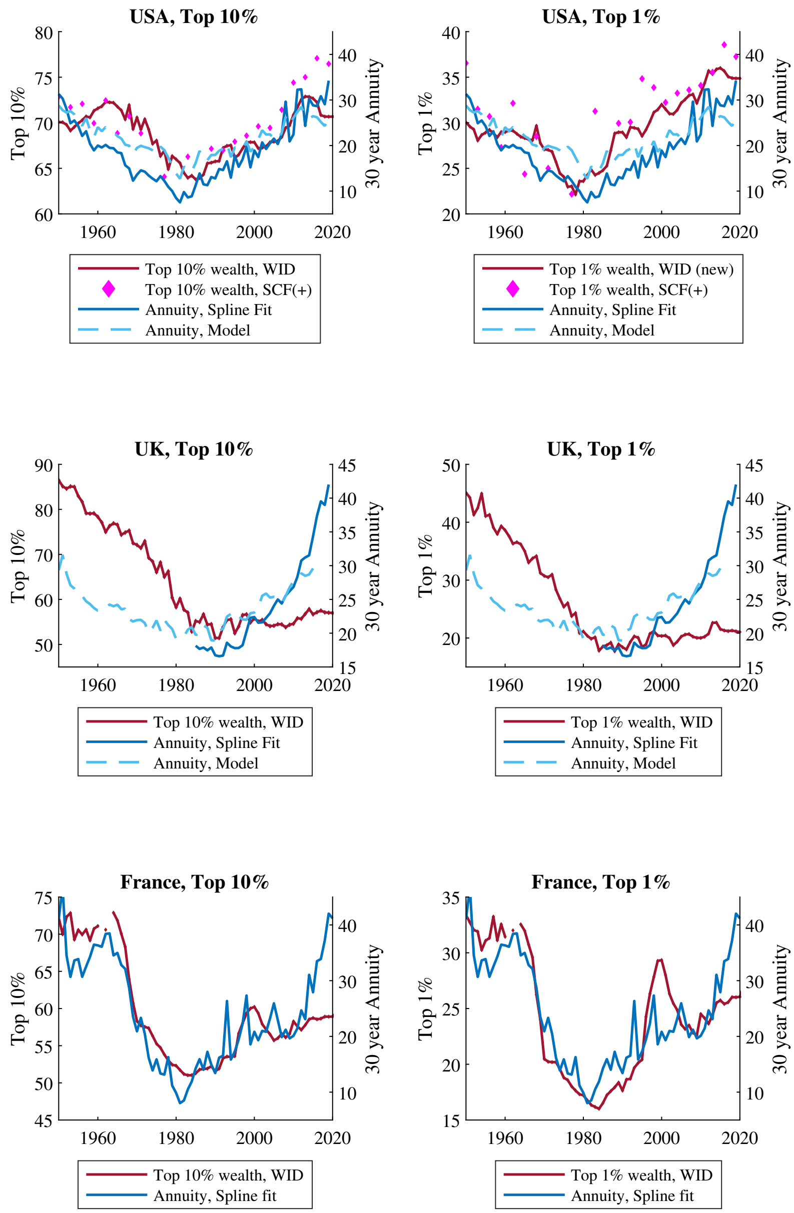

Note: Each panel plots a financial wealth inequality measure against a measure of the cost of a 30-year real annuity. The inequality measure in the left panels is the share of financial wealth going to the top-10% of the population. The right panels plot the share of the top-1% of the population. The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

each asset in the household's portfolio, we can compute the actual duration of that household's portfolio, and determine the change in value of that portfolio when rates change. As Table 2 shows, household assets consist of (i) cash, deposits, and money market instruments, (ii) stocks held directly and indirectly in mutual funds and pension accounts, (iii) real estate, (iv) private business wealth, and (v) fixed income assets (directly and indirectly held). Household liabilities consist of mortgage, student, and consumer debt.

We compute the duration of an asset with risky payoff stream $\{X_t\}$ as

$$D_t(k) = \frac{\sum_{j=0}^{\infty} \mathbb{E}_t \left[M_{t,t+j} X_{t+j} \right] \times j}{\sum_{j=0}^{\infty} \mathbb{E}_t \left[M_{t,t+j} X_{t+j} \right]},$$

where $M_{t,t+j}$ denotes the pricing kernel for payoffs that accrue at $t+k$. To obtain expectations of $M_{t,t+j} X_{t+j}$ for public equities and housing, we use the no-arbitrage asset pricing model detailed in Appendix A. For these assets, the model provides a McCauley duration in each quarter from 1947.Q1 until 2019.Q4, which we average across the entire sample.

For assets with a simple or fixed cash flow structure, we directly calibrate the duration. For cash and deposits, we assume a duration of 0.25 years. For fixed income, we assume a duration of 4 years.[5] For student debt, we assume a duration of 4.5 years.[6] For consumer debt, we assume a duration of 1 year as a compromise between its two main subtypes: revolving debt and amortizing 24-month personal loans.[7] For mortgage debt, we obtain data for the Bloomberg-Barclays Aggregate MBS Index. It is a representative portfolio of all outstanding U.S. pass-through mortgage-backed securities. The average McCauley duration of this representative mortgage portfolio in 1989 and 1990 was 5.2 years. While most mortgage debt in the U.S. has a maturity of 30 years, its duration is much shorter due to amortization, high interest rates (discounting) and prepayment.[8] The resulting durations are reported in the first column of Table 2.

The most challenging duration measurement is for private business wealth, a risky asset class that accounts for 22% of aggregate wealth, 30% of the wealth of the top-10%, and 39% of the wealth of the top-1%. Unlike public equity, however, these assets lack the high-quality price data needed to compute their duration using our auxiliary asset pricing model. Moreover, private business wealth likely contains both fast-growing start-ups, with very high duration, as well as smaller businesses whose existence is heavily tied to the human capital of the owner and hence have

[5]For reference, the maturity of outstanding U.S. Treasury marketable securities averages 62 months between 2000 and 2020. The duration is strictly smaller than the maturity since bonds pay coupons. For example, if the coupon rate is 4.65% and the bond pays semi-annual coupons, then the duration is 4.5 years. Other corporate and international bonds and loans held by U.S. households tends to have somewhat lower duration than U.S. Treasuries because there are fewer long-term bonds and coupons are higher.

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

[7]We exclude both vehicles from assets and auto debt from liabilities.

[8]MBS outstanding in 1989-1990 had an average maturity of 9.8 years and an average coupon rate of 9.35%.

Table 2: Duration of the Household Financial Wealth Portfolio

	Duration	Portfolio Shares			
		All	Bottom 90	P90-P99	Top 1
Assets					
Cash and Deposits	0.25	0.10	0.14	0.11	0.06
Equities	28.78	0.11	0.07	0.12	0.13
Real Estate	14.89	0.54	0.84	0.43	0.29
PBW Short	10	0.12	0.06	0.14	0.17
PBW Long	61.25	0.10	0.01	0.09	0.22
Fixed Income	4.0	0.16	0.15	0.18	0.15
Liabilities					
Mortgage Debt	5.2	0.11	0.24	0.06	0.01
Student Debt	4.5	0.00	0.00	0.00	0.00
Other Debt	1.0	0.02	0.03	0.01	0.01
Average (EW) Duration		**14.83**	**14.57**	**16.48**	**22.69**
Aggregate (VW) Duration		**18.31**	**15.16**	**17.11**	**23.76**

Note: The column "Duration" reports the duration of the asset, averaged over the full sample. For Equities, Private Business Wealth, and Real Estate, the durations are computed form the asset pricing model in Appendix A. The columns "Portfolio Shares" reports the wealth-weighted average portfolio weights from the 1989 Survey of Consumer Finances for i) all households (All), ii) subsample of households with net-wealth in the bottom 90 percentile (Bottom 90), iii) subsample of households with net-wealth in the 90th to 99th percentile (P90-P99), iv) subsample of households with net-wealth in the top 1 percentile (Top 1). Liabilities receive negative portfolio weights.

lower duration (Smith, Yagan, Zidar, and Zwick, 2019). We exploit the detailed categorization of private business wealth in the SCF to split business wealth into high and low duration types, using both the legal structure (limited partnership, S-corporation, other corporation, sole proprietorship, LLC, other) and the size of the business, measured using the number of employees (see Appendix B.3.7 for details). Our main finding is that while the bottom-90% hold almost exclusively the low-duration subcategory, the high-duration share is increasing with wealth, and makes up a majority for the top-1%, boosting their exposure to interest rate changes.

For the high-duration category, we equalize the duration to the duration of small stocks estimated using our auxiliary asset pricing model (61.25), since small public companies may be a good proxy for fast-growing private businesses. For the low duration category, we apply a conservative duration of 10.[9] This benchmark approach results in a private business duration estimate for each individual that depends on the composition of its private business wealth. We perform extensive robustness checks on these assumptions in Section 7.

[9]Under the Gordon Growth Model, this corresponds to a business that has an annual growth rate of cash flows of 1% per year and a constant discount rate of 12% per year. Alternatively, it corresponds to a business that has constant cash flows for the next 23 years before shutting down, using the 1983 real interest rate as the discount rate.

Figure 3: Financial Duration by Net Worth Wealth Percentiles

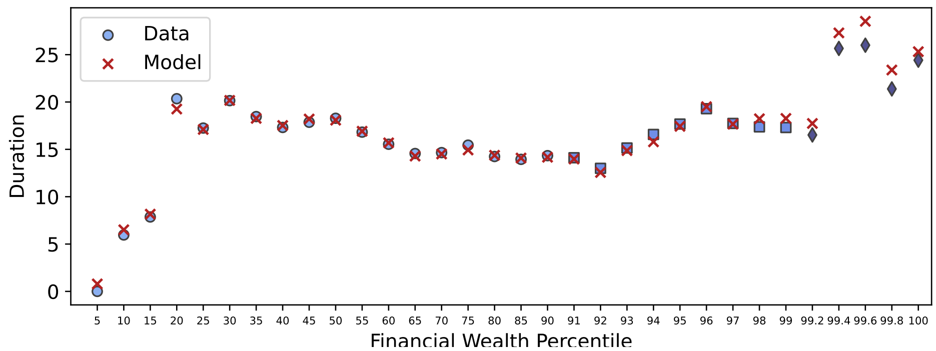

Note: This plot displays average duration by financial wealth bin in the model and data (source: SCF). The x-axis is measured in percentiles, which each tick representing the right edge of the bin, so that e.g., "5" corresponds to households with financial wealth percentile in the interval $[0, 5]$.

Given the duration of each asset class, the duration of a household's portfolio is

$$D_{t,i}^{fin} = \sum_k \omega_t^i(k) D_t(k). \tag{1}$$

where $\omega(k)$ is the weight of asset k in the household's portfolio. We compute these portfolio weights $\omega(k)$ using data from the Survey of Consumer Finances (SCF), and report them in Table 2; see Appendix B.3 for more details.

Having computed household portfolio durations using (1), we compute a number of aggregate statistics that will be useful inputs for the theory in Section 5. We first compute an aggregate (value-weighted) duration of 18.31 that weights households by their wealth, and an average (equal-weighted) duration of 14.83 that weights households equally, and is substantially lower than its value-weighted counterpart. We also compute value-weighted durations of the top-10% and top-1% wealthiest households of 20.02 and 23.76, respectively, both of which are much higher than the overall value-weighted measure.

These statistics reflect a positive association between wealth and financial durations. Richer households tend to hold more equities and (high-duration) private business wealth, which are high-duration assets, and hold fewer short-duration assets such as cash. These greater exposures imply that the portfolios of the wealthy increase in value by proportionally more when rates decline, leading to an increase in financial wealth inequality.

To quantify the empirical correlation between financial duration and the level of financial wealth, Figure 3 plots the average duration by wealth bin in the SCF (dots). Since higher-wealth agents are more important for aggregate wealth outcomes, Figure 3 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 the wealthiest households hold the longest-duration financial portfolios. Further

Figure 4: Financial Duration by Age

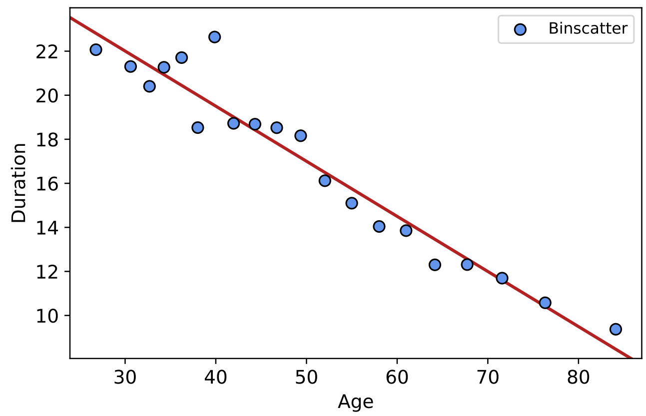

Note: This plot displays a binscatter of average duration by age in the data, after controlling for the financial wealth bins displayed in Figure 3, while the red line represents the least squares fit (source: SCF).

analysis of the duration profile can be found in Appendix B.5.[10]

The second key data pattern is variation in financial duration by age. Figure 4 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 Figure 3. Figure 4 shows that there is a strongly negative relationship between age and duration.

Our empirical results in Appendix B.3 show that adding other covariates yields little additional power to explain variation in financial duration across households. Therefore, we approximate financial duration using the regression:

$$D_i^\theta = \alpha + \beta Age_i + \sum_j \gamma_j NetWealthBin_{i,j} + \varepsilon_i, \tag{2}$$

where $NetWealthBin_{i,j}$ is a dummy for whether household i falls in financial wealth bin j. We use the fitted value from this regression to assign durations to household portfolios in the structural model below. The red crosses in Figure 3 show the equivalent durations in our structural model, to be discussed in Section 6.

To summarize, the key empirical finding is that financial wealth duration increases in the level of financial wealth. We show in the next section that this is a sufficient condition for a decline in real rates to increase wealth inequality.

[10] In particular, this appendix shows that the non-monotonic shape, with duration initially peaking at the 20th - 30th percentile of net worth, then falling to around the 92nd percentile, before rising again, is due to the effects of mortgage leverage. Mortgage leverage amplifies the duration of home equity at the lower end of the net wealth distribution. At higher wealth levels, this effect dissipates due to lower home leverage and higher shares of long-duration assets such as stocks and long-duration private business wealth.

4 Incomplete Markets Model with Household Heterogeneity

4.1 Causality, General Equilibrium, Source of Rate Decline

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 where interest rates are taken as given, and therefore 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 therefore provides a robust and important quantitative measure that complements work on the underlying forces driving changes in interest rates over this period.

Appendix D generalizes our environment to a general equilibrium setting in which interest rates are determined endogenously, and the aggregate endowment grows at a stochastic rate. There, changes in the equilibrium interest rates reflect changes in any or all of (i) in the subjective time discount rate, (ii) in the mean rate of growth, and (iii) in the variance of that growth rate. Following Krueger and Lustig (2009), we show how to map the stochastically growing economy into a stationary economy in the style of Bewley (1986) without growth and aggregate risk. Proposition D.7 in the Appendix generalizes Proposition 5.3 in the main text on the relationship between interest rates and wealth inequality to a world with stochastic growth and endogenous interest rates, showing that the proposition goes through irrespective of the combination of discount factor, growth rate, and growth uncertainty changes driving the movement in rates.

4.2 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 $\phi_J = 0$.

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

Asset Technology. Households trade a complete set of bonds offering fixed cash flows at future dates.[11] 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 also assume that households cannot short bonds, so $x_m \geq 0$ for all households, maturities, and dates.

We 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. We normalize the bond price vector so that $q_0 = 1$, and

$$q_m = R^{-m}, \quad \forall m. \tag{3}$$

Household Problem. Given beginning-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

$$V_j(x; z, q) = \max_{c,x'} \frac{c^{1-\gamma}}{1-\gamma} + \phi_j \beta \mathbb{E}\left[V_{j+1}(x'; z', q') \;\middle|\; z\right] \tag{4}$$

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 $\sum_m q_m x'_m \geq 0$.

Household Optimality. The optimality condition for bond holdings is

$$1 = \left(\frac{q_{m-1}}{q_m}\right) \beta \mathbb{E}\left[\left(\frac{c'}{c}\right)^{-\gamma} \;\middle|\; z\right] + \mu, \quad \forall m.$$

where μ is the multiplier on the borrowing constraint. Since $q_{m-1}/q_m = R$, we obtain the standard Euler equation

$$1 \geq \beta R \mathbb{E}\left[\left(\frac{c'}{c}\right)^{-\gamma} \;\middle|\; z\right] \tag{5}$$

which holds with equality whenever the borrowing constraint is not binding.

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

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. As a result the household's financial wealth $\theta \equiv \sum_{m=1}^{M} q_{m-1} x_m$ is a sufficient state variable in a steady state where interest rates do not change. Given this indifference, we assign each household a unique set of portfolio shares $\hat{\omega}_m = x_m/\theta$, chosen to match its predicted duration in our empirical results given the household's age and position in the wealth distribution (see Section 6 for more details).

Recasting the household problem in terms of θ, households optimize by solving the problem

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

subject to the budget constraint

$$c \leq y_j(z) + \theta - R^{-1}\theta'. \tag{6}$$

and the no borrowing condition $\theta' \geq 0$, where $\theta' = \sum_{m=1}^{M} q_{m-1} x'_m$. However, if there is an unexpected revaluation between steady states, the actual next period financial wealth will be

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

where q' is the updated vector of bond prices conditional on the new realized interest rate.

5 Wealth Inequality and Heterogeneity in Duration

Next, we let the economy undergo a decline in the interest rate and characterize the conditions under which this increases the inequality in financial wealth. To study the implications for consumption and total wealth inequality, we also derive conditions under which household consumption possibilities expand or contract following a shock, as well as the conditions under which household consumption is invariant to the change in rates. We present our main results below and relegate the proofs to Appendix C.

For our main experiment, we allow the economy to undergo an unexpected and permanent change in the interest rate from R to $\tilde{R} = R\exp(-\varepsilon)$ for some shock $\varepsilon > 0$.[12] We simultaneously update the discount factor from β to $\tilde{\beta} = \beta\exp(\varepsilon)$ so that $\tilde{\beta}\tilde{R} = \beta R$. This simultaneous change in β and R can be interpreted as reflecting that the decline in rates was caused by a change in time preference, but is also the correct way to account for a decline in interest rates caused by a

[12] Our experiment assumes that changes in the interest rate are perceived by households as zero probability events before they arrive, and as permanent changes after they arrive. As a result, households at each point in time believe they are solving a stationary problem. As such, we mostly suppress dependence on the calendar time t in our notation for parsimony, despite the fact that transitions over time are an important element of our analysis.

decline in the growth rate in the stationary version of a growing economy (see Appendix D). For the theoretical results of this section we assume for convenience that financial wealth is strictly positive, allowing us to take logs.[13]

5.1 Repricing

We first consider the implications of a decline in interest rates on measured financial wealth inequality. We begin by deriving the standard link between an asset's cash flow duration and its exposure to interest rates.

Lemma 5.1. Given a sequence of cash flows z_t and a valuation

$$Z_0 = \sum_{t=0}^{\infty} R^{-t} z_t \tag{7}$$

we have

$$\frac{\partial \log Z_0}{\partial \log R} = -D \equiv \frac{\sum_{t=0}^{\infty} R^{-t} z_t \times t}{Z_0}$$

where D is the asset's duration. For a small shock $\varepsilon \to 0$, this implies the approximate revaluation

$$\tilde{Z}_0 \simeq Z_0 \exp(-D \times \varepsilon) \simeq Z_0(1 - D \times \varepsilon). \tag{8}$$

Next, we define the duration of a household's financial wealth portfolio, and of its excess consumption claim. For notation, let D^θ to be the duration of the household's portfolio of financial assets, and let D^{c-y} be the duration of the household's excess consumption claim, defined by

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

With these preliminaries in hand, we present our main theoretical result linking interest rates and financial wealth inequality.

Proposition 5.2. Assume that the natural borrowing limits hold. Then for a small negative change in rates ($\varepsilon \to 0$), we have

(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 wealth. A sufficient condition for this positive covariance is that aggregate (value-weighted) financial wealth duration in the pre-shock economy exceeds the average (equal-weighted) duration.

[13] While this assumption would follow naturally from the existence of a zero income state in the presence of Inada conditions, we do not impose this restriction in the calibrated model.

(b) The top-α share of wealth increases if and only if the value-weighted duration for the top-α wealthiest share of households exceeds the value-weighted duration for all households prior to the shock.

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 increase by more than the those of the poor. This is in turn determined by the relative exposures of these portfolios to interest rates, summarized by duration.

This proposition also provides summary statistics that determine whether inequality will rise or fall with interest rates: the value-weighted versus the equal-weighted averages of financial wealth duration in the population of households. If the value-weighted duration is higher, it indicates that the wealthy hold relatively higher-duration portfolios and thus benefit more when interest rates fall. Similarly, determining whether declines in interest rates will increase top wealth shares following a small shock to rates depends solely on whether the financial wealth duration of those in the top group exceed the national aggregate (equivalently, those in the other group).

These summary statistics guide our measurement. Table 2 showed that value-weighted wealth duration (18.31) is substantially larger than equal-weighted duration (14.83), implying that declining interest rates should increase measured financial wealth inequality. Similarly, the value-weighted durations of the top-10% (20.02) and the top-1% (23.76) each exceed the overall value-weighted duration, implying that top financial wealth shares should increase when rates fall.

5.2 Consumption Possibilities

While the results of the previous exercise 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 on equation (6) to obtain

$$\theta = \mathbb{E}_0 \sum_{t=0}^{\infty} R^{-t}(c_t - y_t). \tag{9}$$

Financial wealth is equal to the present value of future excess consumption, defined as consumption minus income. With this identity in hand, we turn to our main theoretical insight regarding the link between interest rates and consumption possibilities.

Proposition 5.3. Let tildes (e.g., $\tilde{c}$) denote allocations following the change in rates, and assume that the natural borrowing limits hold. Then we have:

(a) If $D^{\theta} > D^{c-y}$ then the household's consumption possibilities expand when the interest rate falls, while if $D^{\theta} < D^{c-y}$ the household's consumption possibilities contract.

(b) Household consumption is unchanged ($c_t = \tilde{c}_t, \forall t$) if and only if $D^\theta = D^{c-y}$.

Part (a) 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 that duration compares to the duration of that household's lifetime excess consumption. While financial wealth is always equal to present value of future excess consumption by the budget identity (9), 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. Intuitively, a decline in interest rates increases the cost of a given consumption plan. Thus, if the value of the household's human and financial wealth has not risen sufficiently at the same time, its former consumption plan may no longer be affordable.

For a household that is perfectly hedged ($D^\theta = D^{c-y}$), its pre-shock consumption plan is still budget feasible. Part (b) of the proposition states that this plan remains not only feasible but optimal, so that households will keep their pre-shock consumption plan even following the decline in rates. Intuitively, since $\tilde{\beta}\tilde{R} = \beta R$, any consumption plan that satisfied the Euler equations under (β, R) will continue to do so under $(\tilde{\beta}, \tilde{R})$. As a result, Proposition 5.3 implies that if households arrange their portfolios to perfectly hedge the interest rate risk they face on their excess consumption claim, there should be zero response in household consumption plans and welfare, and that changes in financial wealth should therefore occur purely on paper.

5.3 Implications: Perfectly Hedged Economy

Proposition 5.3 implies that, under the natural borrowing limits, 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 5.2, 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.

In infinite-horizon Bewley models (without a life cycle), households are ex-ante identical. 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. Excess consumption $c_t - y_t$ tends to be negative in the short run and positive in the long run as a result of consumption smoothing, generating a long duration of excess consumption. Since the wealthy therefore tend to have a higher value for D^{c-y} than the poor, an infinite-horizon economy with perfect hedging features a positive cross-sectional covariance between financial wealth and the duration of financial wealth. It sees financial wealth inequality *rise* following a decline in interest rates. This would then make the rise in measured financial wealth inequality documented in Section 3 consistent with perfect hedging. Put differently, rising financial wealth inequality in such an economy is not inconsistent with a welfare-neutral outcome, with every household's consumption allocation unchanged due to perfect hedging.

This result changes dramatically in the presence of a life cycle. Under typical life cycle parameterizations, the young tend to be the least wealthy agents in the economy. 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 (dissaving) in retirement. As a result, life cycle models predict that D^{c-y} is highest for young households, and decreasing in age. The cross-sectional covariance of financial wealth and the duration of excess consumption is therefore negative. This implies that in a perfectly hedged life cycle economy, we should expect to see financial wealth inequality *fall* following a decline in interest rates. Thus, observing a large rise in financial wealth inequality as interest rates have fallen since the 1980s is inconsistent with perfect hedging. This discrepancy implies that there are real consumption effects from interest rate changes, and that these effects are unequally distributed across households.

5.4 Model Extensions

Appendix D shows that the theoretical results from this section generalize to a model where interest rates are determined in equilibrium and where there is aggregate growth risk. In that equilibrium model, the underlying source of the change in rates can be a change in the subjective time discount factor, the average growth rate, the variance of that growth rate, or any combination of these three sources.

The two forces highlighted above, namely (i) the history of stochastic income realizations and (ii) the life-cycle effect of saving for retirement are generic forces that are present in any model and shape the cross-sectional covariance between the level of financial wealth and the duration of financial wealth. A good starting point to evaluate the quantitative strength of these forces is to do so in a standard calibration of a workhorse life-cycle incomplete markets model.

6 Model Quantification

In this section, we quantify this effect of changing interest rates on wealth inequality in a model with realistic heterogeneity among households. We start the model in 1983 when interest rates where high, then trace the impact on inequality as interest rates decline over time.

6.1 Calibration

Preferences and Mortality Risk. We calibrate the model's mortality risk via the survival probabilities s_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.

Households have CRRA preferences with risk aversion γ equal to 2. We set the time discount factor equal to $\beta = 0.963$ to target a ratio of net wealth to disposable labor income in 1983 of 6.79.[14] We do not model an operative bequest motive since accidental bequests turn out to be large enough to capture key features of the observed bequest distribution. The bequest of one household that dies goes to one randomly chosen young household.

Interest Rates We calibrate the interest rate R to match the average value of the 10-year real interest rate in 1983. The interest rate in the data comes from an economy with growth. As explained in Appendix D.3.2, the relationship between the interest rate in our stationary economy and that in the growing economy is:

$$R = \frac{R_g}{G}, \qquad\qquad \beta = G \times \beta_g \tag{10}$$

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 changes according to

$$\widehat{D}_i^{\theta} = \hat{\alpha} + \hat{\beta} Age_i + \sum_j \hat{\gamma}_j NetWealthBin_{i,j}, \tag{11}$$

which is the fitted value estimated from the data. The resulting model durations by wealth, shown by the red crosses in Figure 3, provide a close fit of the data.[15] The model delivers an equal-weighted duration of 14.83, a value-weighted duration of 19.40, and value-weighted durations for the top-10% and top-1% of 21.99 and 24.99, all close to their empirical counterparts in Table 2.

[14] Net wealth is measured as the net worth of households and nonprofits in the Flow of Funds (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.

[15] The small discrepancies are due to differences in the covariance of age and net wealth bins in model and data.

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:

$$\log\left(y^i_{t,a}\right) = \chi' X^i_t + z^i_t, \tag{12}$$
$$z^i_{t+1} = \alpha_i + \varepsilon^i_{t+1} + \nu^i_{t+1}, \tag{13}$$
$$\varepsilon^i_{t+1} = \rho\varepsilon^i_t + u^i_{t+1}, \tag{14}$$

where X^i_t is a vector of household characteristics that includes a cubic function of age.[16] We normalize the mean of the age profile to unity during working life.

The stochastic income component z^i_t contains a household-fixed effect α^i, a persistent component ε^i_{t+1}, and an i.i.d. component ν^i_{t+1}. We have: $\mathbb{E}[z^i] = \mathbb{E}[\alpha^i] = E[\nu^i] = \mathbb{E}[\varepsilon^i] = 0$ and $Var[\nu^i] = \sigma^2_\nu$, $Var[u^i] = \sigma^2_u$, $Var[\alpha^i] = \sigma^2_\alpha$, and $Var[\varepsilon^i_0] = \sigma^2_{\varepsilon,0}$. To allow for lower income risk during retirement, we re-estimate (12) - (14) 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 B.2. Figure B.1 in that appendix plots the deterministic life-cycle income profile.

The literature typically estimates (12)-(14) on labor income for males between ages 25 and 55. We deviate from this practice in three ways, all of which are important for our purposes. First, we consider a broader income concept. Second, we consider the entire life-cycle from age 18 to 80. Third, we focus on households rather than individuals.

First, from the model's perspective, the relevant notion of income is broad and includes transfers. It is the risk in this income that the household is hedging by trading in financial markets (borrowing and saving). 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. Obtaining consistent data on the various components of transfers is involved because successive waves of the PSID use different variable codes for the same concepts. Appendix B.2 provides the details. This data definition is similar to that of Catherine et al. (2020), who share with us a focus on after-transfer income.

Second, we are interested in the entire life-cycle. To this end, we estimate the income distribution for a wide range of ages from 18 to 80.[17] Because our income concept includes transfers such

[16]We have verified that 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), gender (male or non-male), giving rise to 8 groups in total. Since it makes little difference for our main results, we only consider one group here.

[17]Since agents in the model can survive to age 100, we assume that the age profile embedded in $\chi' X^i_t$ remains constant from age 80 onward.

as unemployment benefits and retirement income from public or private defined-benefit pension plans, we do not need to separately model labor force participation or retirement decisions. Our approach instead captures the average decisions made in the data. For example, we do not need to make the assumption that retirement starts at age 65, that income in retirement is some constant fraction of pre-retirement income, or that income risk disappears in retirement. We can let the data speak on these issues. Since our income concept includes income from part-time work, it captures a broader set of earners, such as students.

Third, we focus on households, aggregating income across its adult members. This absolves us from having to model demographic changes such as getting married, divorced, or widowed. We simply follow households identified by the head of household as designated in the data.

Superstar Income Parameters. To help the model match the level of wealth inequality in the high-interest rate regime (1980s), we enrich the income process in (12)-(14) with a superstar income state. This state has a high income level $\varepsilon_t^i = \varepsilon^{sup}$. Households enter in this state with probability p_{12}^{sup} when they are in the regular income state, and return to the regular 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 1.5% probability of entering the superstar income state over one's lifetime. Conditional on entering, the state has an expected duration of 40 years. The income level Y^{sup} is then chosen to match the top-10% wealth share in 1983 exactly, which requires a value equal to 32.71 times average income.

Income Process Discretization. We discretize the regular labor income process (13) using a discrete Markov chain using the method of Rouwenhorst (1995) for ε and supplement it with Gaussian quadrature over ν. For the ε transition probabilities, we assume a nested form so that households first transition between the superstar and non-superstar states. Households that transition from superstar to non-superstar states draw a random non-superstar income state from the ergodic distribution of the non-superstar distribution.

6.2 Stationary Economy: High Interest-Rate Regime

We begin by describing the properties of the model in its stationary distribution under the high interest rate regime. Figure 5 displays the life cycle profiles of income, consumption, financial wealth, and human wealth. The axes are normalized such that 1 represents the median income during working life. Income displays the traditional hump-shape over the life cycle. Income inequality is increasing over the first half of the life cycle because of the accumulation of income shocks and because of the increase in average income over the life cycle profile. Income inequality drops after retirement but is still non-negligible since agents have heterogeneous retirement income and still face some income risk after age 65.

Figure 5: Life Cycle Profiles

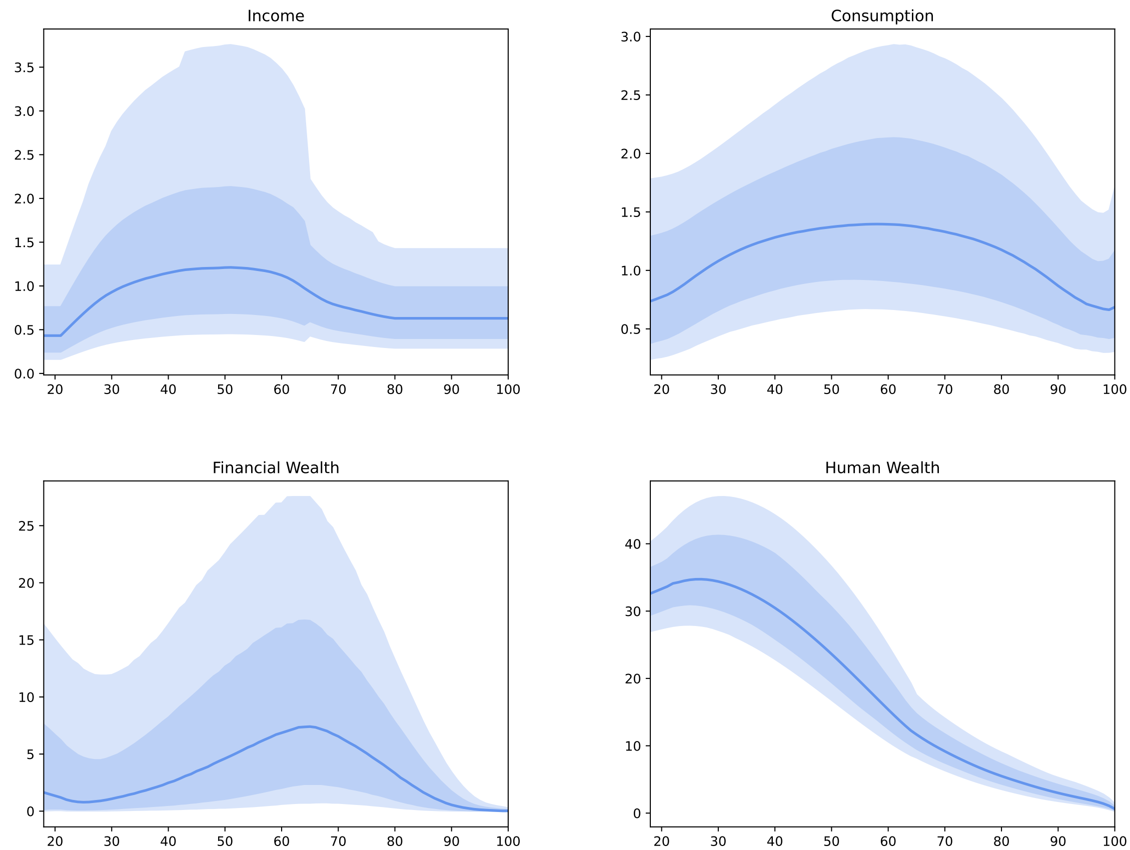

Note: This figure plots the life cycle profiles by age for the all agents of all groups combined. The axes are normalized so that the average income across all agents of all ages is equal to unity. The center line displays the median, while the dark and light bands represent 66.7% and 95% percentile bands. Although agents in the model have a maximum age of 100, we truncate the plot at age 90 due the relatively small sample of agents surviving past this age.

The top-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 reduces. This is consistent with the data which show that consumption inherits the hump-shaped profile from income (e.g., Krueger and Perri, 2006).

Financial wealth in the bottom left panel increases in preparation for retirement, and is subsequently run down during retirement. Financial wealth inequality rises and falls over the life cycle. The initial financial wealth inequality at age 18 reflects the unequal distribution of bequests.

Human wealth in the bottom right panel is generally decreasing in age, except for an initial hump early in the life cycle. There are two effects at play. Human wealth rises in the early parts of the life cycle as the households' highest-earning periods are brought closer to the present. Human wealth falls due to the overall decrease in the remaining periods of work. The latter effect

Figure 6: Wealth Durations

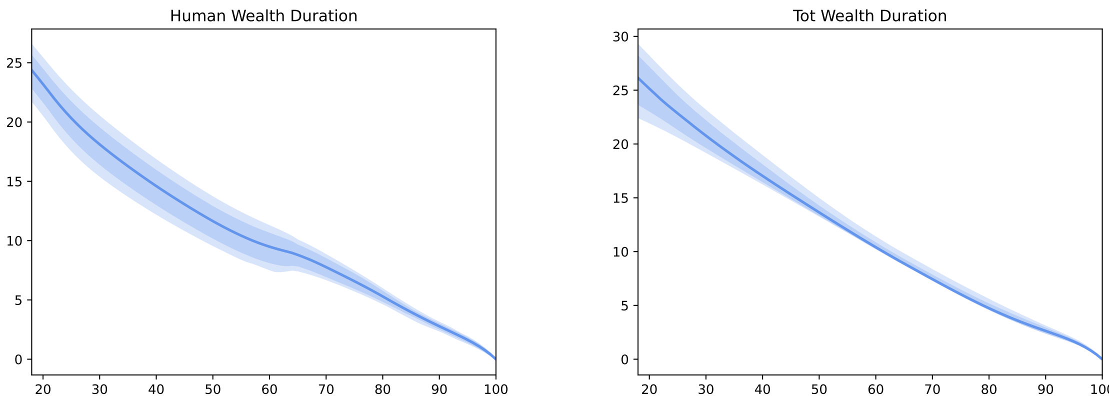

Note: This figure plots the durations of labor income (human) wealth (left panel) and consumption (right panel). The plots display durations computed for many agents simulated from the stationary equilibrium of the model. The economy is normalized so that the average income is equal to unity. The center line displays the median, while the dark and light bands represent 66.7% and 95% percentile bands.

dominates past age 30. Total wealth consists almost exclusively of human wealth when young, except for the few households who receive a bequest. As households age and accumulate financial wealth, a larger share of total wealth becomes financial wealth. However, human wealth remains a large component of total wealth throughout the life-cycle.

Appendix Figure E.1 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 0.468. Consistent with the data, consumption inequality is somewhat lower than income inequality, and has a Gini coefficient of 0.395. Financial wealth is much more unequally distributed than human wealth or total wealth. The Gini coefficients of human and total wealth are 0.384 and 0.405, compared to the Gini of financial wealth of 0.740. The low total wealth inequality arises from (i) the importance of human wealth in total wealth, and (ii) the negative cross-sectional correlation between financial and human wealth.

Figure 6 displays 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, since there are fewer years of life remaining to earn labor/pension income.

7 Results: Repricing Under Falling Interest Rates

In this section, we apply unexpected negative 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 that captures the entire decline in interest rates between 1983 and 2019. This 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 actual sequence of annual interest rate changes.

7.1 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. We update the value of financial wealth for each household in the economy. Given our assumption that households hold zero-coupon bonds with maturity D_i^{θ}, we can mark each household's wealth to market following the rate change using

$$\tilde{\theta}_{i,t} = \theta_{i,t} \exp(-D_{i,t}^{\theta} \times \Delta \log R). \tag{15}$$

Household financial wealth duration $D_{i,t}^{\theta}$ comes from applying the fitted value of (2) to each household in the model. We refer to the resulting wealth distribution $\tilde{\theta}_{i,t}$ as the *repriced wealth distribution*.

Figure 7 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 132.9%. Panel (b) presents a binscatter showing how these gains are allocated across the wealth distribution. The figure shows 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 5.

7.2 Repricing: Transition Experiment

We now turn to our transition experiment, which provides our main quantitative results. For each year between 1983 and 2019, we use the actual real 10-year rate $R_{g,t}$ implied by our auxiliary asset pricing model in Appendix A. We detrend using equation (10) with a constant growth rate G to obtain the time series of R_t to use in our stationary economy. We assume that each of these

Figure 7: Histograms, Repriced Financial Wealth Distribution

(a) Hist: Original vs. Repriced

(b) Binscatter: Original vs. Repriced

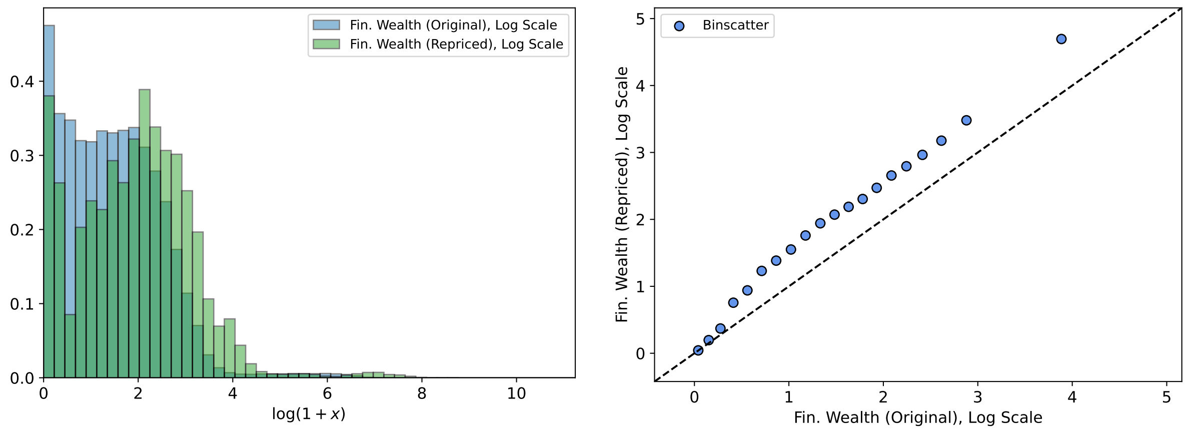

Note: The left panel plots the original wealth distribution in the steady state with high interest rates in blue and the repriced wealth distribution after the decline in interest rates in green. The right panel is a binscatter plot, where each dot represents 5% of the population, that maps the financial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

annual changes is permanent and unexpected. At the start of each period, we update the value of financial wealth according to (15), using $\Delta \log R_t$ as our change in rates. To compute household portfolio durations, we re-estimate equation (2) for each SCF wave (three years), then interpolate the coefficients for the years between SCF waves.[18] These time-varying regressions capture both changes in portfolio shares over time, as well as changes in the financial duration of each asset over time according to our auxiliary asset pricing model. Each model household receives a financial duration based on its age and position in the wealth distribution, using the fitted value of equation (11). Between periods, we update the financial wealth of each household based on its optimal consumption-savings condition (5).

The results are displayed in Panel (a) of Figure 8. The black line displays the actual top-10% share in the data (WID), while the red line displays the implied path of the top-10% share in our model's transition experiment. The figure shows that our model-implied series explains 75% 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 the vast majority 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. To separate these, 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

[18] Since the first wave of the SCF is 1989, we use the fitted values of equation (2) for all years between 1983 and 1989.

Figure 8: Top-10% Share, Gradual Transition Exercise

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

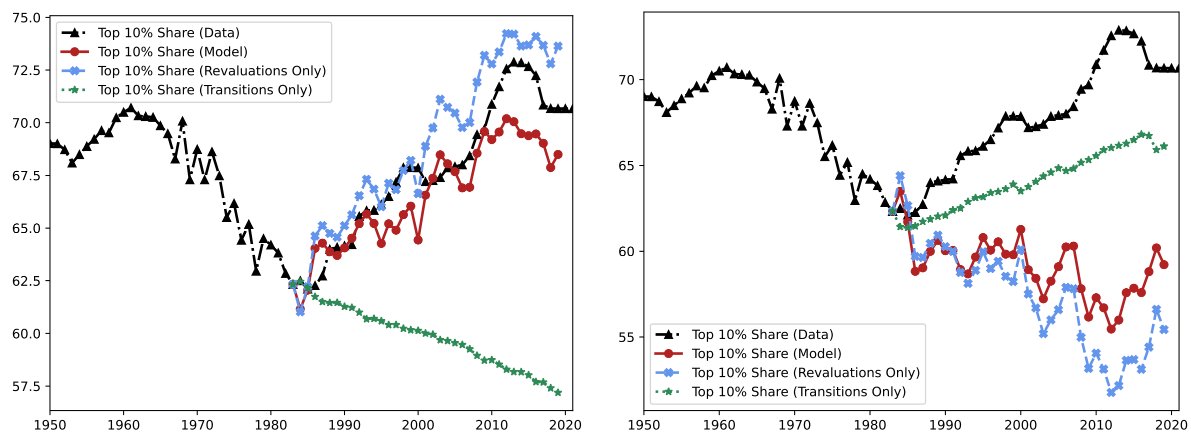

Note: The left panel plots the top-10% financial wealth share in the data (black dash-dotted line) and in the model with repricing (red solid line). It decomposes the evolution in the top-10% financial wealth share into a component solely due to instantaneous repricing (blue line, $\hat{S}_t^{10,REV}$) 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.

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) above, to be

$$d\hat{S}_t^{10,REV} = \tilde{S}_t^{10} - S_t^{10}.$$

and cumulate these revaluation effects to obtain the series

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

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

$$\hat{S}_t^{10,MR} = S_0^{10} + \left(S_t^{10} - \hat{S}_t^{10,REV}\right).$$

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. This effect of revaluations is offset by households' endogenous consumption-savings responses, which push wealth inequality down.

To understand this pattern, we note that, despite the large increase in financial wealth inequal-

Table 3: Change in Inequality: Gradual Transition Experiment

	Data WID	Repriced	Compensated
Top-10% FW	+8.3pp	+6.2pp	-3.1pp
Top-1% FW	+11.3pp	+8.1pp	-0.9pp
Gini FW	+0.054	+0.039	-0.040
Top-10% HW	–	+1.4pp	+1.4pp
Top-1% HW	–	-1.9pp	-1.9pp
Gini HW	–	+0.070	+0.070
Top-10% TW	–	+0.2pp	-1.9pp
Top-1% TW	–	-1.5pp	-3.2pp
Gini TW	–	+0.054	+0.036

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 (World Inequality Database) column.

ity 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 under low rates for wealthy households and households saving for retirement to build their wealth by compounding returns. The large rise in inequality observed in Figure 8 is a temporary—albeit highly persistent—effect of the large capital gains stemming from the fall in interest rates, rather than the low level of rates themselves. Thus, in the absence of shocks, the model's endogenous transitions would see inequality decline as it transitions back to this steady state. The very strong effects of revaluation are partially offset by this mean reversion effect, yielding our overall result that the decline in rates explains most, but not all, of the rise in inequality.

The magnitudes of our results are summarized in the top panel of Table 3, which displays changes in financial wealth (FW) inequality. Table E.1 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. The "Repriced" column displays the change between 1983–2019 in our benchmark model economy in the gradual repricing experiment (red line in Figure 8a). The increase in the top-10% share of financial wealth is 75% of that in the data. The changes in the financial wealth top-1% share and Gini coefficient are also on the order of their data counterparts. 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 of the secular increase in financial wealth inequality since the 1980s.

7.3 Robustness to Private Business Duration Measurement

Because we do not directly measure the duration of private business wealth (PBW) in the data, we approximate it, assuming that short-duration PBW has a duration of 10 and long-duration PBW has the same duration as a portfolio of small public equities, which our empirical asset pricing model estimates to be 61.25 in the 1980s.[19] The decomposition of PBW for each household is detailed in Appendix B.3.7. It shows that the share of PBW that is of the long-duration type is much higher for the richest households (top-10% or top-1%) than for the bottom-90%.

To explore robustness to our baseline assumptions on the duration of PBW, Table 4 reproduces our main results using four alternatives for the duration of PBW. In the top panel, we keep our assumption that the long-duration PBW has the same duration as that of small stocks (61.25) but change the duration of short-duration PBW from 10 (baseline) to 28.78 (overall equity market, middle column) or to 61.25 (last column). The latter assumes that all of PBW has the same duration as small stocks. It results in the largest rise in financial wealth inequality among the four alternatives, overshooting the data for the top-10% share (114% explained) and closely matching it for the top-1% share (103% explained).

The bottom panel lowers the duration of long-duration PBW from 61.25 to 52. The latter value comes from an analysis, detailed in Appendix B.4, which, instead of assuming that PBW behaves like a portfolio of perpetually small stocks that is continually resorted, uses CRSP microdata on market value and payouts (dividends or dividends and repurchases) to directly estimate the cash flow duration of the smallest public companies. This procedure yields values between 52 and 62, depending on whether we look at the smallest quintile or decile of listed firms. In the interest of robustness, we choose the lower value of 52. It then combines this assumption with the same three choices for the duration of the short-duration PBW. They deliver a rise in the top-10% share between 4.3pp and 7.5pp.

To sum up, while our findings vary with the estimated duration of PBW, with the share of observed variation in the top-10% financial wealth share explained ranging from 52% to 114%, the ultimate conclusion that our duration mechanism explains the majority of the rise in financial wealth inequality observed since the 1980s is robust.

[19] The asset pricing model in Appendix A matches both the observed time-series of price-dividend ratios and dividend growth rates on a small stock portfolio, the smallest size decile of publicly traded firms. This model implies a time-varying expected return and McCauley duration. To understand the high value for the duration of small stocks intuitively, consider a back-of-the-envelope calculation based on the Gordon Growth Model where the McCauley duration is $(1+r)/(r-g)$. Under perfect foresight, we can use the average realized real return and average realized real dividend growth rate from 1985–2020 to proxy for the expected real return and expected real dividend growth rate in the 1980s. For the smallest market capitalization decile, we find $r = 8.01\%$ and $g = 6.55\%$. This delivers a duration of 75.8, close to the number we obtain using our more sophisticated SDF model and the bottom quintile of market caps. For comparison, for stocks in the largest decile of market capitalization, we obtain $r = 7.94\%$ and $g = 0.81\%$, resulting in a duration of 15.2. The expected returns on large and small stocks are nearly identical. Thus, the high duration for small stocks arises from its high cash flow growth rate.

Table 4: Inequality by Private Business Duration (Transition)

Baseline: $D_H^{PBW} = D^{smallstocks}$				
		Repriced, Alternative D_L^{PBW} Specifications		
Fin. Wealth	Data WID	Baseline	$D_L^{PBW} = D^{eq}$	$D_L^{PBW} = D_H^{PBW}$
Top-10%	+8.3pp	+6.2pp	+7.4pp	+9.5pp
Top-1%	+11.3pp	+8.1pp	+9.5pp	+11.6pp
Gini	+0.054	+0.039	+0.048	+0.062
Robustness: $D_H^{PBW} = 52$				
		Repriced, Alternative D_L^{PBW} Specifications		
Fin. Wealth	Data WID	Baseline	$D_L^{PBW} = D^{eq}$	$D_L^{PBW} = D_H^{PBW}$
Top-10%	+8.3pp	+4.3pp	+5.4pp	+7.5pp
Top-1%	+11.3pp	+5.6pp	+6.9pp	+8.9pp
Gini	+0.054	+0.027	+0.035	+0.048

Note: The table reports the change in the Top-10% share, Top-1% share, and Gini Coefficient of financial wealth under various assumptions on the duration of private business wealth. The change is measured between 1983 and 2019 in the model columns (Repricing) as well as in the Data WID column. The top panel uses the baseline value of 61.25 for the duration of long-duration private business wealth, while the bottom panel uses a lower value of 52. Each panel reports on three values for the duration of short-duration private business wealth: the baseline value of 10, the duration of the overall stock market of 28.78, and a case where the duration of the two types of private business wealth is the same.

7.4 Robustness: Time-Varying Income Inequality

While our main experiments isolate the response of wealth inequality to declining rates, movements in interest rates were by no means the only force influencing wealth inequality over this period. Notably, (top) income inequality also increased over this period. Figure 9 Panel (a) displays the evolution of the top-1% income share in the U.S. from the same WID data source. The top-1% share of income increases from 11.5% in 1983 to 19.1% in 2019.

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 (13) each year on rolling samples to generate distinct income risk parameters for each year of our transition experiment, as detailed in Appendix B.2.3. Given those time-varying income risk parameters, we then calibrate the ratio of income in the superstar income state relative to the regular income state to exactly match the level of top-1% income inequality observed in the WID data in each year.

Figure 9 Panel (b) displays the resulting paths for financial wealth inequality in the gradual transition exercise.[20] The figure shows that, when augmented by the actual observed rise in in-

[20]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. For easier comparison, we thus normalize the plot by subtracting the level in 1983, so that each series is equal to zero in that year.

Figure 9: Top-10% Share, Gradual Transition + Time Varying Income Inequality Exercise

(a) Top-1% Income Share (WID) | (b) Transition Experiment: Rates + Income Ineq.

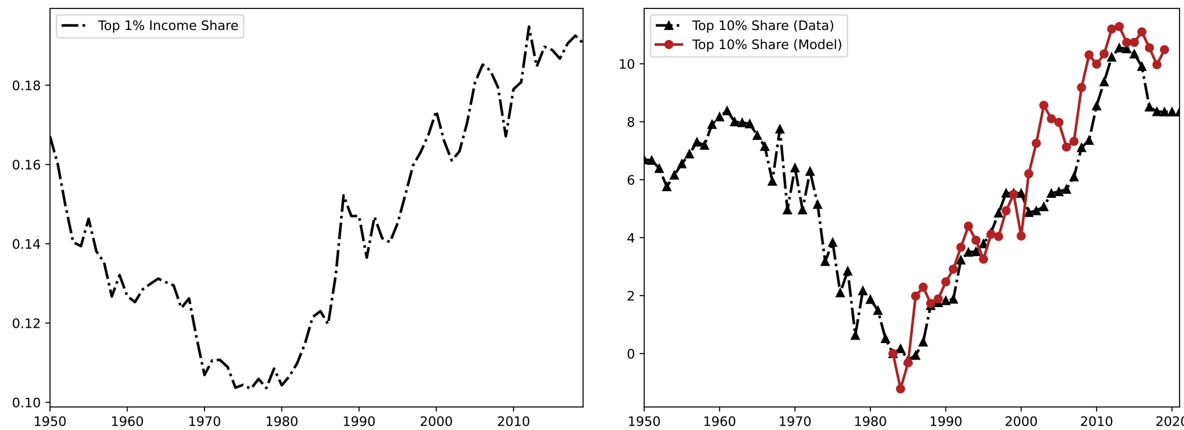

Note: The left panel shows the top-1% income share form the World Inequality database. The right panel plots the top-10% financial wealth share in the data (black line) and in the model that feeds in both gradual interest rate changes and gradual changes in the income process (red line). Both series in the right panel are normalized to zero in 1983 by subtracting the 1983 value from the entire series.

come inequality, our transition path reproduces the actual dynamics of wealth inequality very well. Combining declining interest rates and rising top income inequality generates an increase in the top-10% wealth share of 10.5pp, slightly more than the actual increase of 8.3pp. In sum, our results are robust to allowing for rising income inequality over this period.

8 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 establish an intuitive baseline that is consistent with the theoretical analysis in Proposition 5.3 of Section 5, we compute the change in financial wealth that would be required to maintain each household's pre-shock consumption allocation from the high interest-rate economy after the rate change. We refer to the counterfactual wealth allocation in which "fully hedged" households receive this required financial wealth as the *compensated* financial wealth distribution. As in Section 7, we begin with a one-shot experiment in which we apply the full decline in rates using a single shock, before we present our more realistic gradual transition experiment for our main quantitative results.

Before presenting the results, we note that evaluation of the compensated distribution depends only on pre-shock consumption plans and the interest rate. While the model's functional forms and parameters are used to generate the consumption plans, they do not independently influence

Figure 10: Histogram, Compensated vs. Original Financial Wealth Distribution

(a) Original vs. Compensated

(b) Compensated vs. Repriced

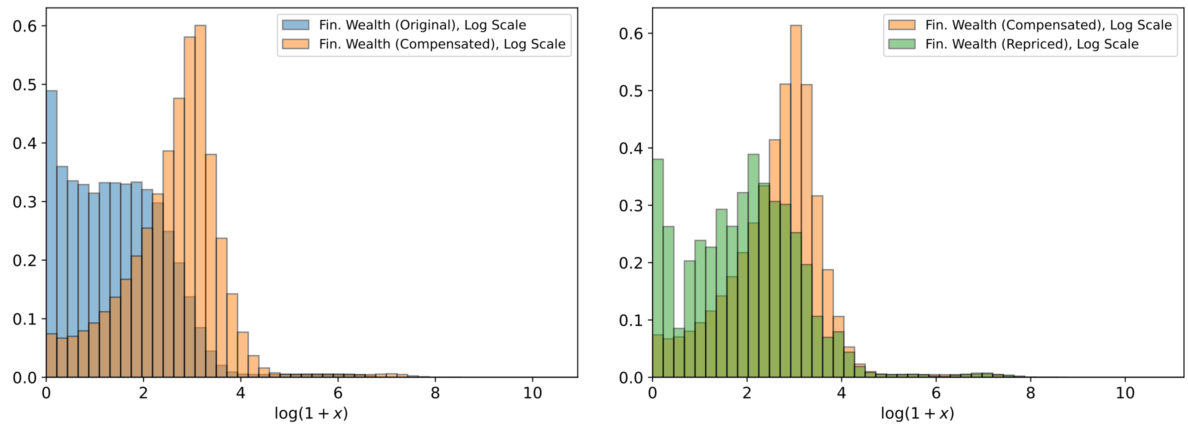

Note: This plot displays the distribution of financial wealth under the stationary distribution and under the compensated distribution drawn from the stationary distribution of the economy. The x-axis displays a transformation $\log(1 + x)$ of the original data. Each distribution is top coded at the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

the compensated distribution conditional on those consumption plans. Thus, to the extent that our model's consumption plans are reasonable, our normative results regarding the compensated distribution should be relatively robust to our particular parametric and modeling choices.

8.1 The Compensated Distribution: One-Shot Experiment

Figure 10 presents the distribution of financial wealth following our one-shot decline in interest rates, alongside the original (pre-shock) distribution.[21] This comparison shows two major differences between the pre-shock and compensated distribution. First, the compensated distribution is shifted substantially to the right from the original wealth distribution. Households in this economy mostly save ($c_t < y_t$) earlier in life before dissaving ($c_t > y_t$) in old age. When rates are much lower, households lose much of the effect of compound interest on their retirement savings and require substantial compensation to be able to afford the old consumption plan. As a result, the aggregate amount of financial wealth in the compensated distribution exceeds the pre-shock total by 193.8%. 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 (30.6%) of new financial wealth accrues to top-1% financial wealth holders under the compensated distribution.

Second, although the wealthiest gain under this compensated distribution, the less wealthy gain proportionally more. The financial wealth Gini is substantially lower in the compensated distribution than in the original one. Visually, while the original high interest-rate distribution of

[21] To ensure that the full distribution is visible, we display transformed variables $\log(1 + x)$ on the x-axis. Because many agents have zero financial wealth, a standard log transform would be inappropriate in this context.

Figure 11: Scatterplots, Compensated vs. Original Financial Wealth Distribution

(a) Medians by Age

(b) Means by Fin. Wealth

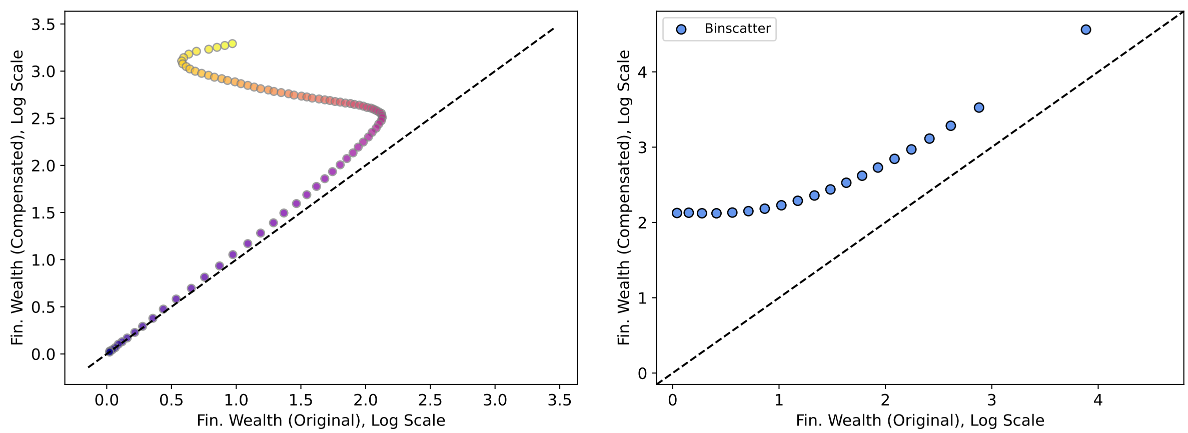

Note: Panel (a) plots the distribution of original financial wealth against the distribution of compensated financial wealth by age. Each dot represents the population mean for one year of age, with the lightest (yellow) dots representing the youngest agents and the darkest (purple) dots representing 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.

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-1% decreases from 62.3% in the baseline economy to 52.0% in the compensated economy.

To see why inequality falls in the compensated distribution, we turn to Figure 11. Panel (a) compares the original (horizontal axis) and compensated financial wealth distributions (vertical axis) by age. The youngest agents (light/yellow) have a relatively low level of wealth in the original distribution, but require the most financial wealth in the compensated distribution. As households age, their actual wealth increases, but their compensated wealth falls. Late in life (dark/purple), both actual and compensated wealth fall toward zero. Actual and compensated distributions are close to coinciding for these older households.

This result is perhaps surprising, since the young have the majority of their portfolio "invested" in human wealth, which has a very long duration (recall Figure 6), and thus provides a natural hedge against interest rate changes. The key challenge the young face in a low interest rate environment, however, is not from their current portfolio, but their future portfolios. Due to the life-cycle profile of income, the young plan to save during middle age, then dissave during retirement. Under a low interest rate, the young will be unable to accumulate enough interest on their future savings, making their original consumption plans unattainable without large infusions of financial wealth today. In contrast, older agents have already benefited from the higher rate of return in accumulating their retirement assets, while the oldest are dissaving, consuming principal rather than interest. These households are less affected by the loss of high-return investment opportunities, and hence require less compensation.

Panel (b) of Figure 11 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. Although the average financial wealth of the young appears higher than that of the old in Panel (a), the financial wealth of the young is largely attained through bequests, which are highly skewed. Many young hold close to zero financial wealth, yet require large transfers under the compensated distribution. As a result, the low wealth group mixes over agents requiring the largest and smallest amounts of compensation. Quantitatively, the young make up a disproportionate share of this group and dominate the aggregate result, so that the least wealthy agents in this economy require the most compensation, measured as the vertical distance from the dot to the dashed 45-degree line. As wealth increases, we move toward the middle-aged individuals in the economy, who require a non-zero level of compensation, but less than those at the bottom of the wealth distribution. Finally, the wealthiest agents in the top bin, whose wealth is more driven by their income realizations than by demographics, also require a strictly positive level of compensation, but less than that of the least wealthy.

To separate the influence of age and wealth, Appendix Figure E.2a reproduces Panel (b) controlling for age fixed effects. It shows that the less wealthy still require more compensation even after controlling for age. Hence, both age and wealth play independent roles in understanding the effects of interest rates on consumption opportunities.

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 that actually result under low interest rates. Figure 10b contrasts histograms of the repriced and compensated distributions. The compensated and repriced distributions display strikingly different shapes, with the repriced distribution leaving many more agents at lower wealth levels than in the compensated distribution.

Figure 12 provides more direct evidence on the winners and losers from lower rates. Panel (a) compares by age, contrasting the change in financial wealth required under the compensated distribution to that actually obtained under the repriced distribution. It 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, older households 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.

Panel (b) of Figure 12 reproduces this comparison by initial wealth, displaying the net gain from repricing, defined as the change in repriced wealth net of the change in compensated wealth.

Figure 12: Scatterplots, Repriced Financial Wealth Distribution

(a) Compensated vs. Repriced: Age Medians

(b) Scatter: Repriced Gain

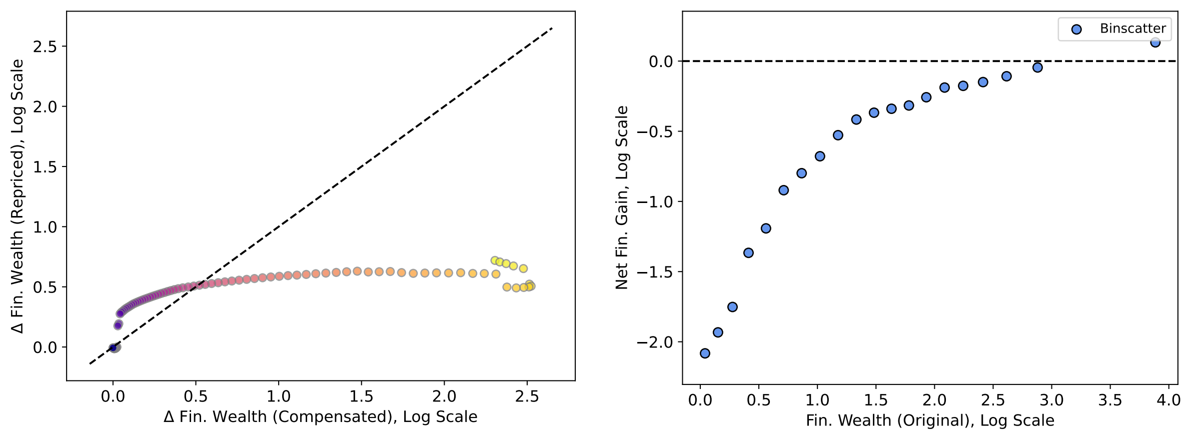

Note: This plot displays the distribution of financial wealth under the repriced distribution, compared to the compensated distribution. Panel (a) displays the change in financial wealth relative to the original distribution for the compensated (x-axis) and repriced (y-axis) distributions. Both axes 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.

The figure reinforces the previous finding, showing that wealthier households (top-5%) 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. Appendix Figure E.2b shows that the same pattern across the wealth distribution holds (to a lesser degree) after controlling for age fixed effects.

8.2 The Compensated Distribution: Transition Experiment

As before, our main quantitative experiment feeds in the observed sequence of interest rate changes gradually. Figure 8b shows the evolution of the top-10% share of the compensated wealth distribution in the red line, along with the same decomposition into instantaneous revaluation and mean reversion effects. The last column of Table 3 summarizes these changes. The compensated wealth distribution sees a substantial reduction in financial wealth inequality, with the top-10% financial wealth share falling by -3.1pp between 1983 and 2019, the top-1% by -0.9pp, and the Gini by -0.040 . The effect is more than explained by the instantaneous compensation following each interest rate change (blue line in Figure 8b), again offset to a lesser degree by consumption-savings choices (green line). In sum, the starkly diverging paths for wealth inequality in the repriced and compensated distributions imply that changing interest rates do not merely have paper gains but result in important changes in consumption possibilities. The young and poor see substantial deterioration in their consumption; the duration of their financial wealth is below the duration of their excess consumption. The opposite is true for the old and the rich; their consumption opportunities expand from declining rates.

8.3 Total Wealth Inequality

From a welfare perspective, what matters are consumption opportunities and how those are distributed across households. We now turn to the implications of changing rates for total wealth inequality, where total wealth is defined as the present discounted value of consumption, equal to the sum of human and financial wealth.

First, we observe that all three total wealth inequality measures are much lower than their financial wealth inequality counterparts in the initial distribution. For example for the Gini coefficient, the total wealth Gini (0.405) is close to the human wealth Gini (0.384) and much lower than the financial wealth Gini (0.740). This occurs because (i) human wealth is a large component of total wealth for most households, (ii) human wealth is more equally distributed than financial wealth, and (iii) human and financial wealth tend to be negatively correlated in the cross-section of households.

Second, the bottom two panels of Table 3 show that lower rates increase human and total wealth inequality. The effects are stronger for the Gini coefficient than for the top-10% and top-1% shares. Lower interest rates increase the human wealth Gini by 0.070 . Younger households own most of the human wealth and have a high duration of human wealth. Given the positive cross-sectional covariance between human wealth levels and human wealth durations, a decline in the interest rate generates an increase in human wealth inequality, consistent with part (a) of Proposition 5.2. In the bottom-90% of the total wealth distribution, human and financial wealth tend to be negatively correlated. The young tend to have low financial wealth but high human wealth, whose high duration provides a natural interest rate hedge that buffers the blow from declining rates on their consumption possibility set. Due to this negative correlation between human and financial wealth, the total wealth Gini only increases by slightly more than the human wealth Gini, and by less than the sum of the increase in the financial and human wealth Ginis.

At the top of the total wealth distribution, the response to the rate change is more subtle. The top-1% of total wealth 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 state. They are younger on average and have much lower ratios of financial to total wealth. The wealth dynamics of this group are governed by the dynamics of the top-1% human wealth share, which decreases by 1.9pp when rates fall. Because superstar labor income has a lower duration than regular income, 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 all households. A decrease in rates thus leads to a decrease in top-1% human wealth share, consistent with Proposition 5.2 part (b).

Combined, the top-10% total wealth share shows a very modest response to lower rates, increasing by only 0.2pp, lower than the top-10% human wealth share growth of 1.4pp, and much lower than the 6.2pp growth in the top-10% financial wealth share.

Last, we note that the repriced distribution for total wealth features more inequality than the compensated distribution. Our conclusion regarding the real effects of falling rates on inequality carries over from financial wealth to total wealth. Quantitatively, the gap between repriced and compensated inequality measures is smaller for total than for financial wealth.

9 Cohort-Level Analysis

Our model also helps to shed light on how different cohorts have gained and lost over our sample period. Our results in Section 8 showed that the duration structure of older households ($D^{\theta} > D^{c-y}$) expands their consumption possibilities when interest rates fall, while younger households see their own consumption possibilities contract ($D^{\theta} < D^{c-y}$). Since falling interest rates arrive at different times in the life cycle of different birth cohorts, we should expect differential net benefits from the same path of realized interest rates over our sample.

Figure 13 illustrates such cohort effects implied by the model. Each line displays the present value of remaining lifetime consumption $\sum_{k=0}^{J-j} R^{-k} c_{t+k}$ over the 1983–2019 sample. To make the series comparable, we use the 1983 steady state interest rate as R when computing the present value at all dates. Each line represents a different birth cohort, with the oldest cohort (born in 1900) well into retirement at the start of our sample and the youngest cohort (born in 2000) entering the workforce only in the sample's final year. To show the impact of falling interest rates while controlling for typical life cycle effects, we present these results as percent deviations of cohort medians from the corresponding medians for households of the same age in the 1983 steady state.

The graph shows that the older cohorts gain while the younger cohorts lose. The very oldest (1900 and 1910 cohorts) see modest gains since they have mostly run down their wealth by the time the interest rate declines begin. The 1920 and 1930 cohorts are the biggest winners experiencing peak total wealth gains in excess of 10%. In contrast, the youngest cohorts (Gen X, Millennials, and Gen Z) experience large total wealth losses approaching 10% in some years.

Figure 13: Cohort Total Wealth Outcomes, Transition Experiment

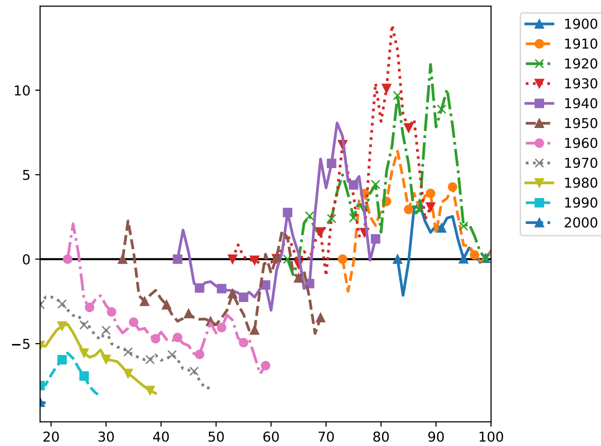

Note: The graph plots total lifetime wealth, the present discounted value of consumption, along the observed interest rate path for the median member of a birth cohort, indicated by the various colored lines, in percent deviation of what that wealth path would have been in the steady state with high (1983) interest rates.

10 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 most of the 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. 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.

References

ALVAREDO, F., L. CHANCEL, T. PIKETTY, E. SAEZ, AND G. ZUCMAN (2018): *World inequality report 2018*, Belknap Press.

ALVAREZ, F. AND U. J. JERMANN (2001): “Quantitative asset pricing implications of endogenous solvency constraints,” *Rev. Financ. Stud.*, 14, 1117–1151.

AUCLERT, A. (2019): “Monetary Policy and the Redistribution Channel,” *Am. Econ. Rev.*, 109, 2333–2367.

BACH, L., L. E. CALVET, AND P. SODINI (2020): “Rich Pickings? Risk, Return, and Skill in Household Wealth,” *American Economic Review*, 110.

BENHABIB, J. AND A. BISIN (2018): “Skewed wealth distributions: Theory and empirics,” *Journal of Economic Literature*, 56, 1261–91.

BERNANKE, B. S. ET AL. (2005): “The global saving glut and the US current account deficit,” Tech. rep.

BEWLEY, T. (1986): “Stationary Monetary Equilibrium with a Continuum of Independently Fluctuating Consumers,” in *Contributions to Mathematical Economics in Honor of Gerard Debreu*, ed. by W. Hildenbrand and A. Mas-Collel, Amsterdam: North-Holland.

BOAR, C. AND V. MIDRIGAN (2020): “Efficient Redistribution,” Working Paper.

CABALLERO, R. J., E. FARHI, AND P.-O. GOURINCHAS (2008): “An equilibrium model of” global imbalances” and low interest rates,” *American economic review*, 98, 358–93.

CATHERINE, S., M. MILLER, AND N. SARIN (2020): “Social Security and Trends in Inequality,” .

COX, J. (2020): “Return Heterogeneity, Information Frictions, and Economic Shocks,” Working Paper New York University.

DEATON, A. S. AND C. PAXSON (1994): “Saving, growth, and aging in Taiwan,” in *Studies in the Economics of Aging*, University of Chicago Press, 331–362.

DOEPKE, M. AND M. SCHNEIDER (2006): “Inflation and the Redistribution of Nominal Wealth,” *J. Polit. Econ.*, 114, 1069–1097.

EGGERTSSON, G. B. AND N. R. MEHROTRA (2014): “A model of secular stagnation,” Tech. rep., National Bureau of Economic Research.

EICHENGREEN, B. (2015): “Secular stagnation: the long view,” *American Economic Review*, 105, 66–70.

FAGERENG, A., M. BLOMHOFF HOLM, B. MOLL, AND G. NATVIK (2019): "Saving Behavior Across the Wealth Distribution: The Importance of Capital Gains," NBER Working Paper No. 26588.

FAGERENG, A., L. GUISO, D. MALACRINO, AND L. PISTAFERRI (2020): "Heterogeneity and Persistence in Returns to Wealth," *Econometrica*, 88, 115–170.

GLOVER, A., J. HEATHCOTE, D. KRUEGER, AND J.-V. RIOS-RULL (2020): "Intergenerational Redistribution in the Great Recession," *Journal of Political Economy*, 128.

GOMEZ, M. AND E. GOUIN-BONENFANT (2020): "A Q-Theory of Inequality," Working Paper Columbia University.

GORDON, R. J. (2017): *The rise and fall of American growth: The US standard of living since the civil war*, vol. 70, Princeton University Press.

GUPTA, A. AND S. VAN NIEUWERBURGH (2021): "Valuing Private Equity Strip by Strip," *The Journal of Finance*.

GUTIÉRREZ, G. AND T. PHILIPPON (2017): "Declining Competition and Investment in the US," Tech. rep., National Bureau of Economic Research.

HEATHCOTE, J., F. PERRI, AND G. L. VIOLANTE (2010): "Unequal we stand: An empirical analysis of economic inequality in the United States, 1967–2006," *Review of Economic dynamics*, 13, 15–51.

HEATHCOTHE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2020): "How Should Tax Progressivity Respond to Rising Income Inequality," NBER Working Paper No. 28006.

HUBMER, J., P. KRUSELL, AND A. A. SMITH (2020): "Sources of U.S. Wealth Inequality: Past, Present, and Future," in *NBER Macroeconomics Annual*, vol. 35.

JIANG, Z., H. LUSTIG, S. VAN NIEUWERBURGH, AND M. 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, Ò., M. SCHULARICK, AND A. M. TAYLOR (2017): "Macrofinancial history and the new business cycle facts," *NBER macroeconomics annual*, 31, 213–263.

KOPCZUK, W. (2017): "US CAPITAL GAINS AND ESTATE TAXATION," *The Economics of Tax Policy*.

KOPCZUK, W. AND E. ZWICK (2020): "Business incomes at the top," *Journal of Economic Perspectives*, 34, 27–51.

KRUEGER, D. AND H. LUSTIG (2009): "When is Market Incompleteness Irrelevant for the Price of Aggregate Risk?" *Journal of Economic Theory*, 145, 1–41.

——— (2010): "When is market incompleteness irrelevant for the price of aggregate risk (and when is it not)?" *J. Econ. Theory*, 145, 1–41.

KRUEGER, D. AND F. PERRI (2006): "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory," *The Review of Economic Studies*, 73, 163–193.

KUHN, M., M. SCHULARICK, AND U. I. STEINS (2020): "Income and wealth inequality in America, 1949–2016," *Journal of Political Economy*, 128, 000–000.

LEOMBRONI, M., M. PIAZZESI, M. SCHNEIDER, AND C. ROGERS (2020): "Inflation and the Price of Real Assets," Working Paper 26740, National Bureau of Economic Research.

LUSTIG, H., S. VAN NIEUWERBURGH, AND A. VERDELHAN (2013): "The Wealth-Consumption Ratio," *Review of Asset Pricing Studies*, 3(1), 38–94., review of Asset Pricing Studies.

MIAN, A., L. STRAUB, AND A. SUFI (2020): "The Saving Glut of the Rich and the Rise in Household Debt," NBER Working Paper No. 26941.

MOLL, B. (2020): "Comment on "Sources of U.S. Wealth Inequality: Past, Present, and Future"," in *NBER Macro Annual*, ed. by Martin Eichenbaum And, 25.

PIKETTY, T. (2015): "About capital in the twenty-first century," *American Economic Review*, 105, 48–53.

PIKETTY, T. AND E. SAEZ (2003): "Income inequality in the United States, 1913–1998," *The Quarterly journal of economics*, 118, 1–41.

PIKETTY, T., E. SAEZ, AND G. ZUCMAN (2018): "Distributional national accounts: methods and estimates for the United States," *The Quarterly Journal of Economics*, 133, 553–609.

PIKETTY, T. AND G. ZUCMAN (2015): "Chapter 15 - Wealth and Inheritance in the Long Run," in *Handbook of Income Distribution*, ed. by A. B. Atkinson and F. Bourguignon, Elsevier, vol. 2, 1303–1368.

RACHEL, L. AND T. SMITH (2017): "Are Low Real Interest Rates Here to Stay?" *International Journal of Central Banking*, 13.

ROUWENHORST, K. G. (1995): "Asset pricing implications of equilibrium business cycle models," *Frontiers of business cycle research*, 1, 294–330.

SAEZ, E. AND G. ZUCMAN (2016): "Wealth inequality in the United States since 1913: Evidence from capitalized income tax data," *The Quarterly Journal of Economics*, 131, 519–578.

SMITH, M., D. YAGAN, O. ZIDAR, AND E. ZWICK (2019): "Capitalists in the Twenty-First Century," *The Quarterly Journal of Economics*, 134, 1675–1745.

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

SUMMERS, L. H. (2014): "US economic prospects: Secular stagnation, hysteresis, and the zero lower bound," *Business economics*, 49, 65–73.

WACHTER, J. (2005): "Solving Models with External Habit," *Finance Research Letters*, 2, 210–226.

ZUCMAN, G. (2019): "Global wealth inequality," *Annual Review of Economics*, 11, 109–138.

A Affine Asset Pricing Model

This appendix develops a reduced-form asset pricing model. The asset pricing model is used for three 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 real estate wealth in a manner that is consistent with the history of bond and stock prices. Third, the model delivers the price and duration of a claim to aggregate consumption and to aggregate labor income.

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:

$$\boldsymbol{z}_t = \boldsymbol{\Psi} \boldsymbol{z}_{t-1} + \boldsymbol{\Sigma}^{\frac{1}{2}} \boldsymbol{\varepsilon}_t, \tag{16}$$

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:

$$m^{\$}_{t+1} \quad = \quad -y^{\$}_{t,1} - \frac{1}{2}\mathbf{\Lambda}'_t\mathbf{\Lambda}_t - \mathbf{\Lambda}'_t\boldsymbol{\varepsilon}_{t+1}. \tag{17}$$

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:

$$\mathbf{\Lambda}_t = \mathbf{\Lambda}_0 + \mathbf{\Lambda}_1 z_t. \tag{18}$$

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.

$$\begin{aligned} z_t = \quad & \Big[\pi_t, x_t, y^{\$}_{t,1}, y^{\$}_{t,20} - y^{\$}_{t,1}, pd^m_t, \Delta d^m_t, pd^{small}_t, \Delta d^{small}_t, \\ & pd^{growth}_t, \Delta d^{growth}_t, pd^{value}_t, \Delta d^{value}_t, pd^{infra}_t, \Delta d^{infra}_t \\ & pd^{hw}_t, \Delta d^{hw}_t, \Delta cx_{t+1}, \Delta lx_{t+1}, cx_{t+1}, lx_{t+1}\Big]'. \end{aligned} \tag{19}$$

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 $\boldsymbol{z}_t$. We also fit the dynamics of 20-quarter nominal bond risk premia (price changes).

Figure A.1 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 A.2 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 A.3 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

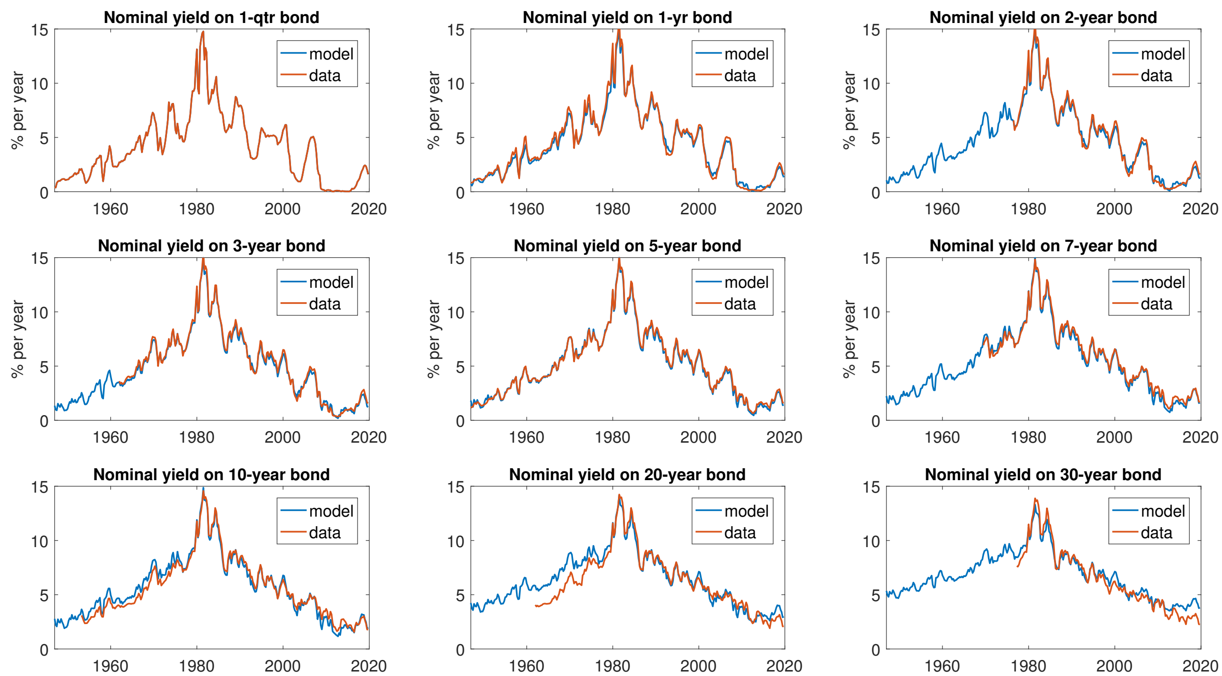

Note: The figure plots the observed and model-implied nominal bond yields. Data are from FRED: constant-maturity Treasury yields, daily averages within the quarter.

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^m_t, is the sum of the prices to each of its future cash-flows D^m_t. These future cash-flow claims are the so-called market dividend strips or zero-coupon equity (Wachter, 2005). Dividing by the current dividend D^m_t:

$$\frac{P^m_t}{D^m_t} = \sum_{\tau=1}^{\infty} P^d_{t,\tau} \tag{20}$$

$$\exp\left(\overline{pd} + e'_{pd^m} z_t\right) = \sum_{\tau=0}^{\infty} \exp\left(A^m_\tau + B^{m\prime}_\tau z_t\right), \tag{21}$$

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

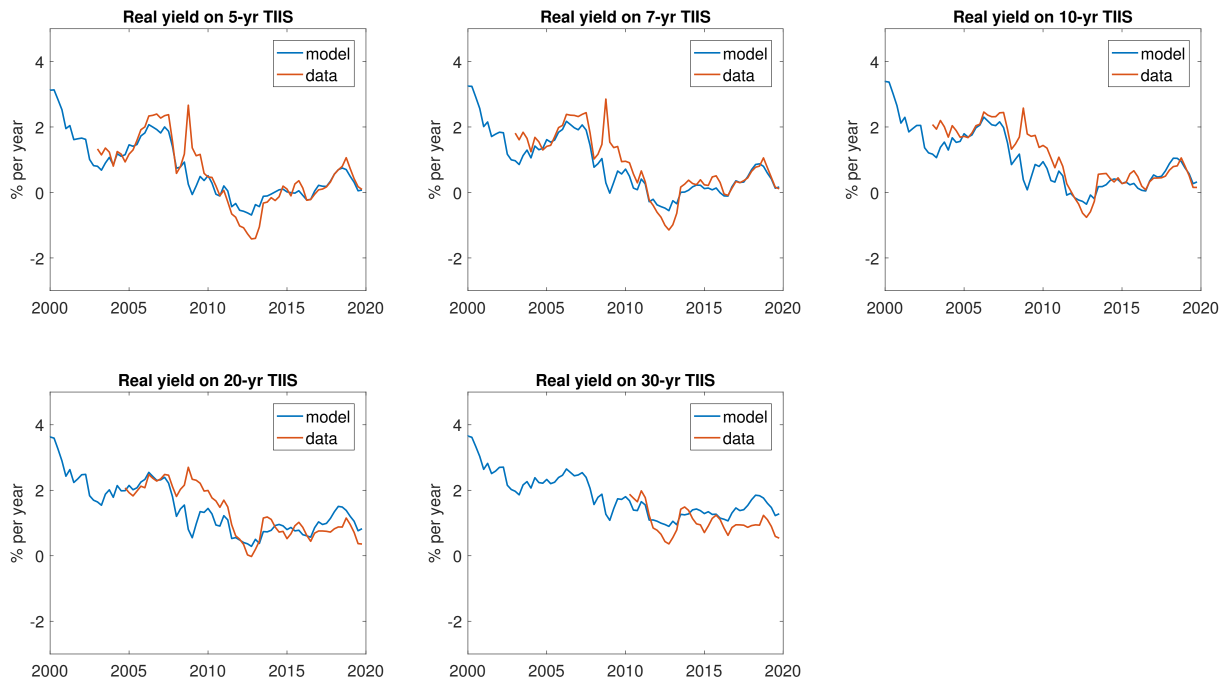

Note: The figure plots the observed and model-implied real bond yields. Data are from FRED: constant-maturity Treasury inflation-indexed bond yields, daily averages within the quarter.

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 (21) restates the present-value relationship from equation (20). 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

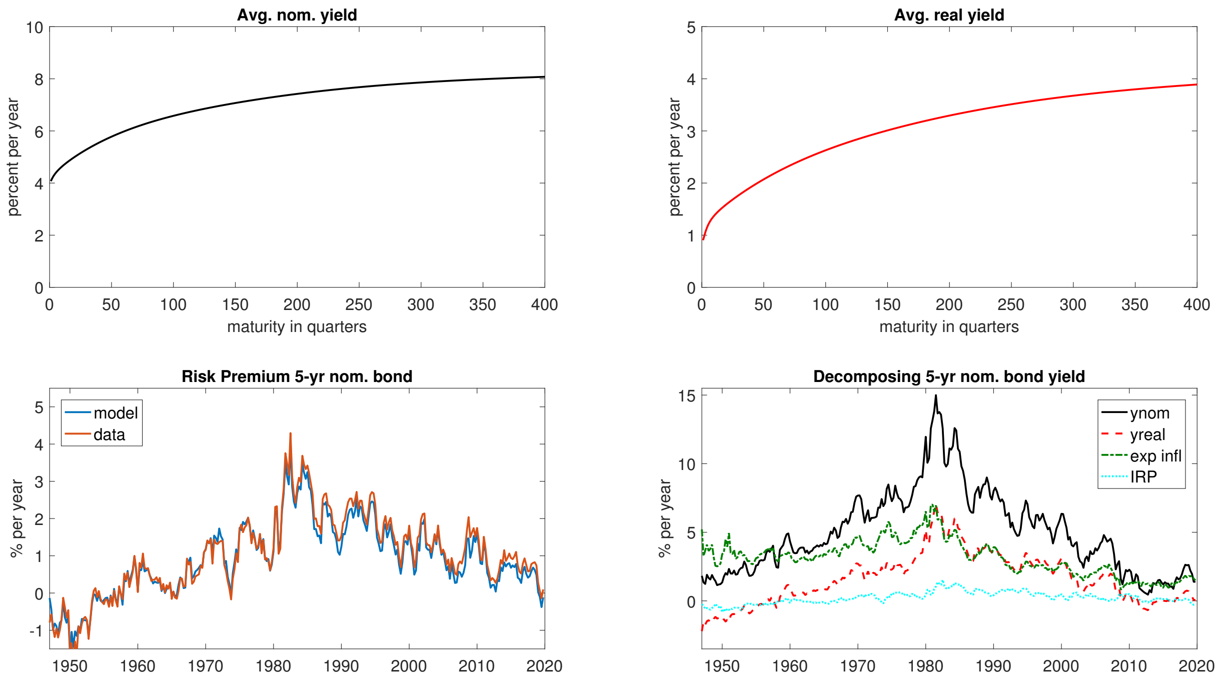

Note: The top panels plot the average bond yield on nominal (left panel) and real (right panel) bonds for maturities ranging from 1 quarter to 400 quarters. The bottom left panel plots the nominal bond risk premium in model and data. The bottom right panel decomposes the model's five-year nominal bond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

VAR).

Figures A.4 and A.5 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)

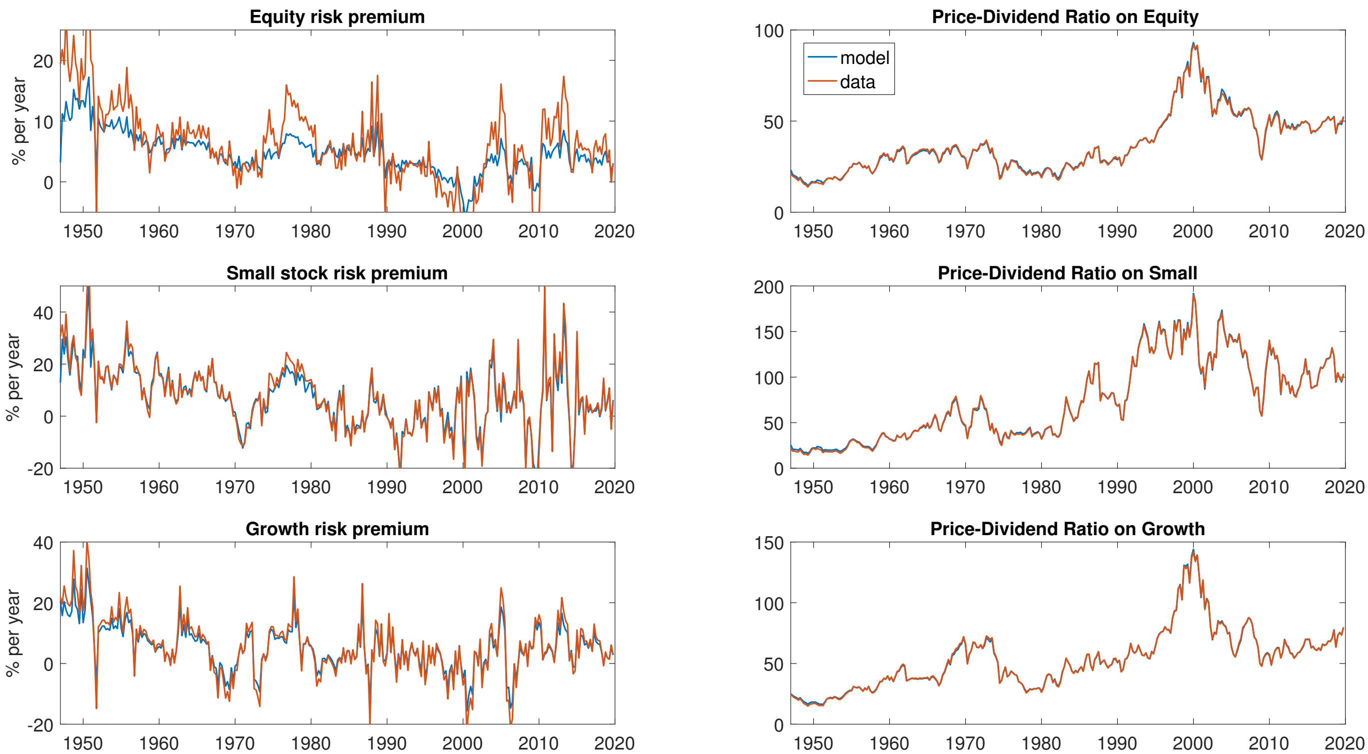

Note: The figure plots the observed and model-implied equity risk premium on the overall stock market, small stocks, and growth stocks in the left panels, as well as the corresponding price-dividend ratio in the right panels. The model is the blue line, the data are the red line.

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 A.6 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)

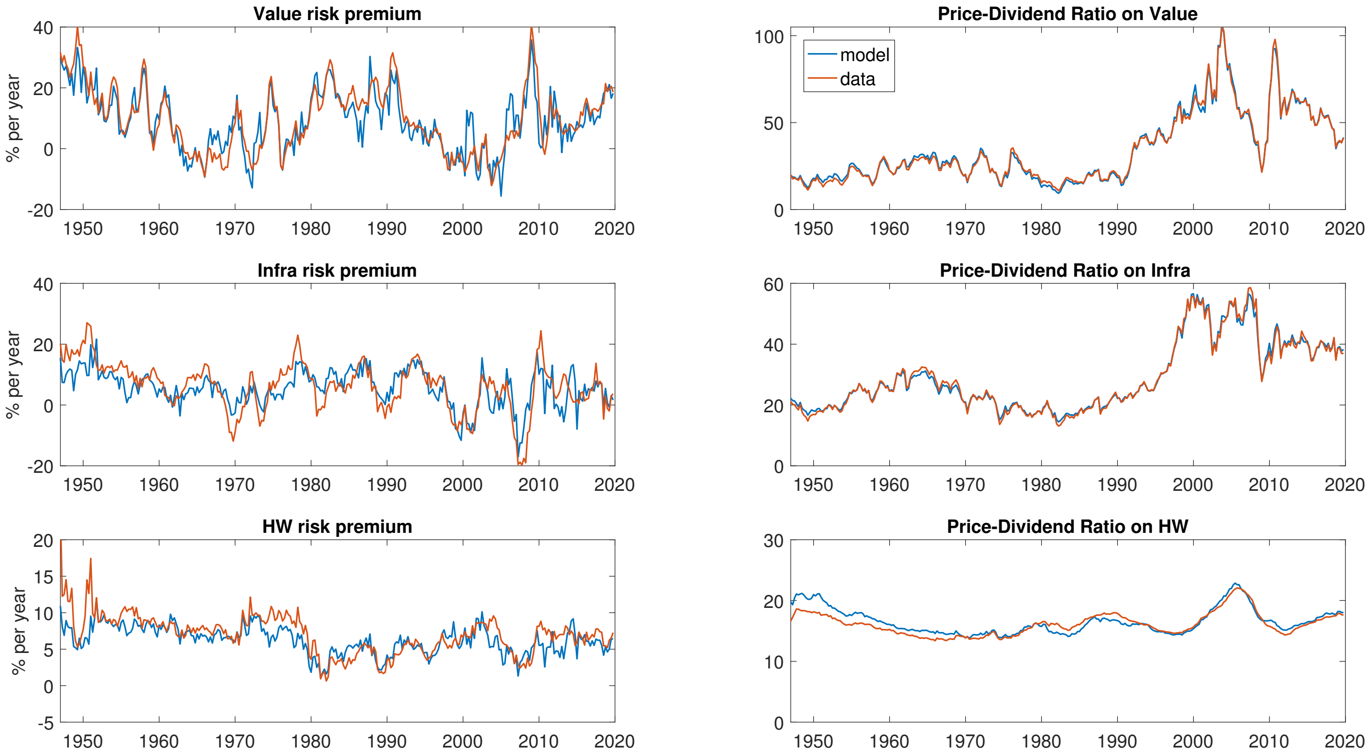

Note: The figure plots the observed and model-implied equity risk premium on value stocks, infrastructure stocks, and housing wealth in the left panels, as well as the corresponding price-dividend ratio in the right panels. The model is the blue line, the data are the red line.

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:

$$D_t^{CF,m} = \sum_{\tau=1}^{\infty} w_{t,\tau}\tau, \qquad w_{t,h} = \frac{P^d_{t,\tau}}{\frac{P^m_t}{D^m_t}} = \frac{\exp\left(A^m_\tau + B^{m\prime}_\tau z_t\right)}{\exp\left(\overline{pd} + e'_{pd^m} z_t\right)}$$

where $P^d_{t,\tau}$ 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 A.7 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

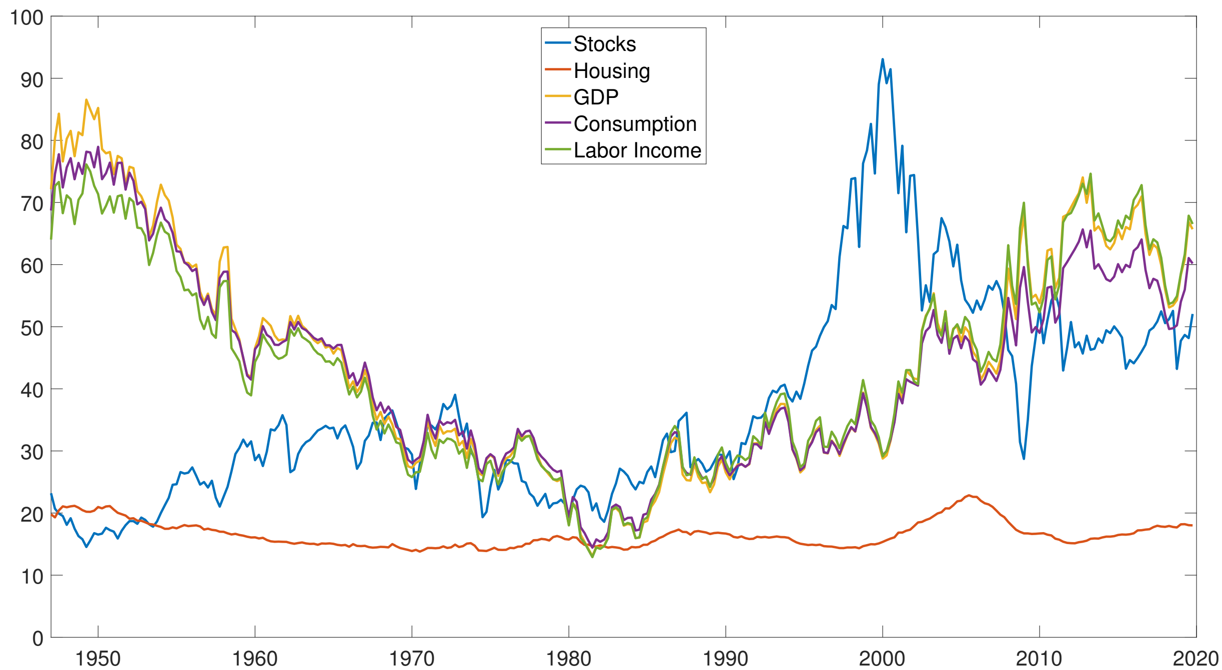

Note: The figure plots the annual price-dividend ratios on the aggregate stock market, housing wealth, and on claims to GDP, aggregate consumption, and aggregate labor income.

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

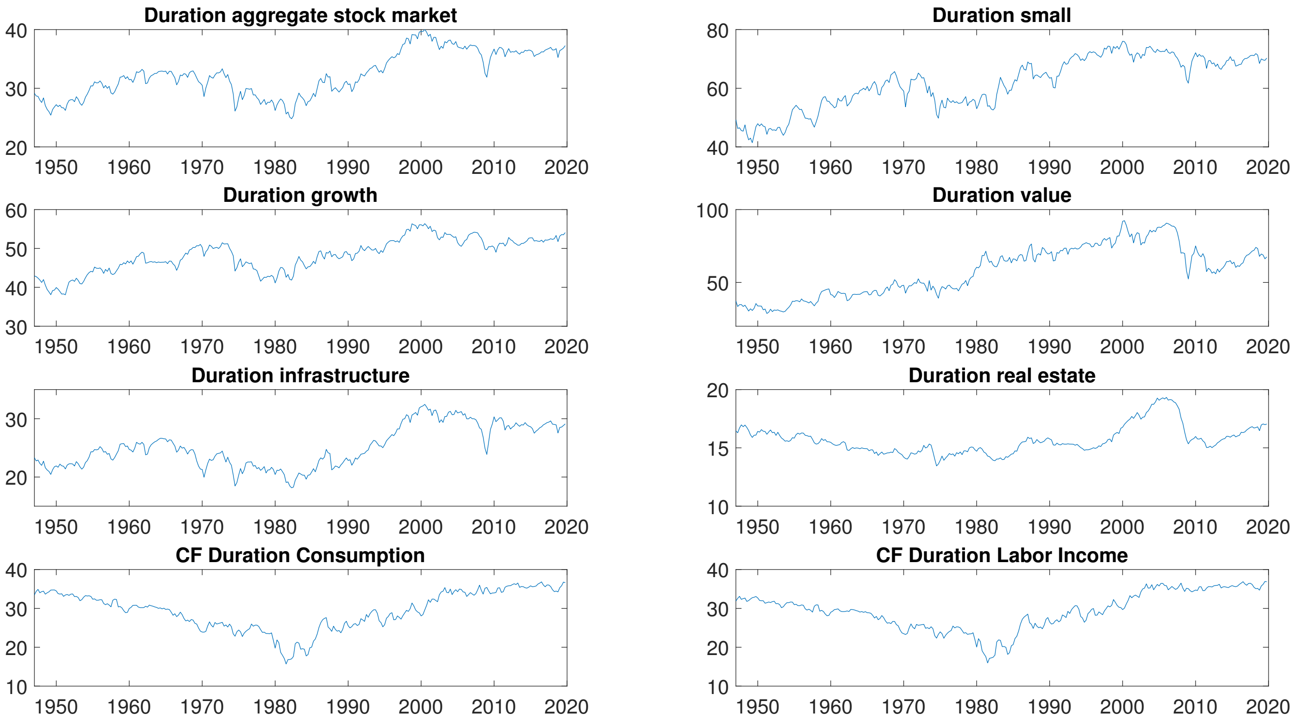

Note: The figure plots the model-implied time series of cash-flow durations on the overall stock market, small stocks, growth stocks, value stocks, infrastructure stocks, housing wealth, the GDP claim, the aggregate consumption claim, and the aggregate labor income claim. The duration is expressed in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

B.1 Inequality Data

Our primary source of data for the top wealth shares presented in 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+. We use data from the SCF+, the database developed by Kuhn et al. (2020), from 1950 to 1983. We use SCF data from 1989 to 2019.

We construct the price of a real 30 year annuity by estimating the historical real yield curve for each country. Letting $y^r_{t,m}$ denote the real yield at maturity m at time t the cost of the annuity is calculated as

$$a_t = \sum_{m=1}^{30} \frac{1}{(1 + y^r_{t,m})^m}$$

Due to varying availability of data we use three different approaches to estimate the real yield curve that lead to broadly consistent estimates. Firstly, for the UK post 1985 we use historical time series of real yields from to fit a spline through these points and construct the real yield curve directly. Secondly, 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 annually 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 sample of 50 years prior and then subtracting forecasted inflation from nominal yields at all available maturities (3-month treasury yields and 10-year government bond yields for all periods, as well as 30-year government bond yields for later periods). Thirdly, for the U.K. and U.S. we also use model estimates of the real yield curve from .

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 et al. (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).

B.2 Income Data

B.2.1 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: an over-sample of 1,872 low income families from the Survey of Economic Opportunity (the "SEO sample") and a nationally representative sample of 2,930 families 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: 1970 to 2017.

B.2.2 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).[22]

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

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.

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*).[23]

B.2.3 Estimating 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:

$$\log y_{c,t} = \beta + \gamma_a + \gamma_c + \gamma_t + \varepsilon_{c,t}, \tag{22}$$

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

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

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 22 using constrained OLS.

Figure B.1 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.

Figure B1: Income Profile

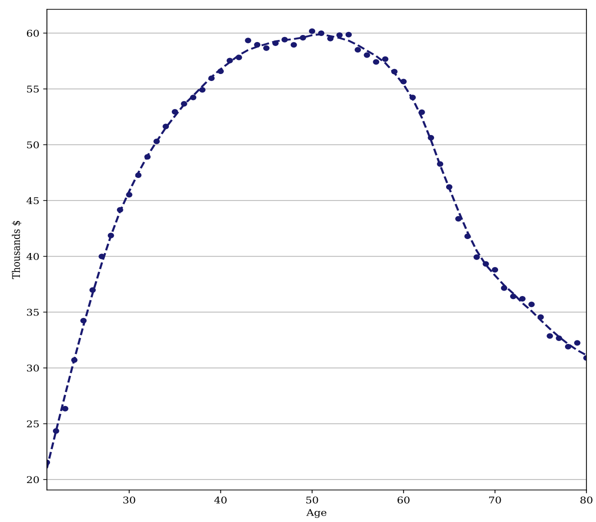

Note: This figure displays the expected income profile evaluated at the 2016 year-fixed effects. The graph plots the expected income profile for the average person who is 21 years old in 2016, expressed in thousands of 2016 dollars. The model is estimated according to Equation (22) on PSID data from 1970 to 2017.

In the second stage, we estimate income risk. We estimate (12) 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.

Using Equation (12)-(14), and define j as equal to the age of the households minus the minimum age (21), we find that:

$$
\begin{aligned}
E[\eta^i_j, \eta^i_{j+h}] &= \sigma^2_\alpha + E[{\varepsilon^i_j}^2] + \sigma^2_\nu \qquad \text{if } h = 0, \\
E[\eta^i_j, \eta^i_{j+h}] &= \sigma^2_\alpha + \rho^h E[{\varepsilon^i_j}^2] \qquad \text{if } h > 0, \\
E[{\varepsilon^i_j}^2] &= \rho^{2j}\sigma^2_{\varepsilon_0} + \sum_{k=1}^{j} \rho^{2(j-k)}\sigma^2_u.
\end{aligned}
$$

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 $\sigma^2_{\varepsilon_0} = 0$, 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 B.1. These are the parameters used in the main text.

Table B1: Idiosyncratic Risk Parameter Estimates

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

Note: $\theta = (\rho, \sigma_{v,w}, \sigma_{u,r}, \sigma_{v,r}, \sigma_{u,r}, \sigma_\alpha)$, are estimated using Equation (12)-(14); $\sigma^2_{\varepsilon_0}$ is fixed equal to 0. Data are based on PSID and runs from 1970 to 2017.

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 B.1. Figure B.2 plots time varying estimates of $\sigma^2_{u,w}$, $\sigma^2_{v,w}$, $\sigma^2_{u,r}$, $\sigma^2_{v,r}$ and σ^2_α.

Figure B2: Time Varying Income Risk

(a) Transitory Shocks - Working Age

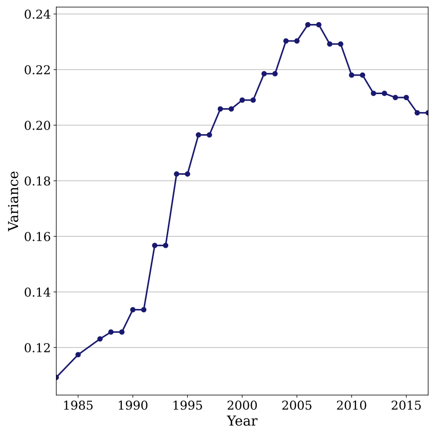

(b) Persistent Shocks - Working Age

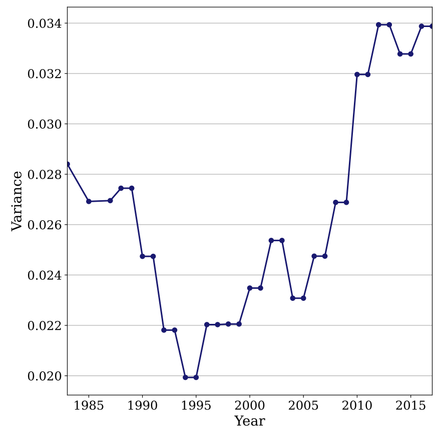

(c) Transitory Shocks - Retirement

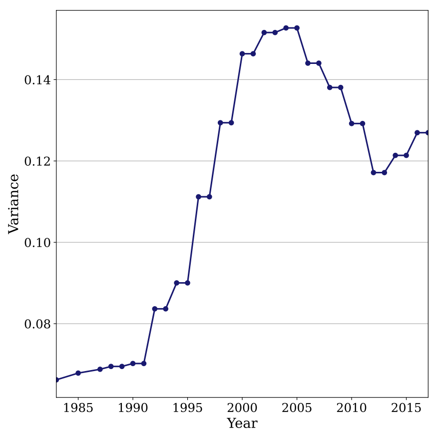

(d) Persistent Shocks - Retirement

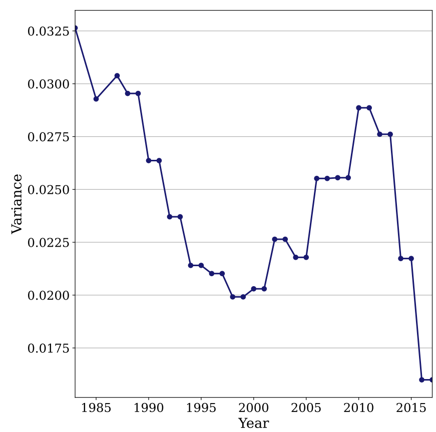

(e) Fixed-Effect Shocks

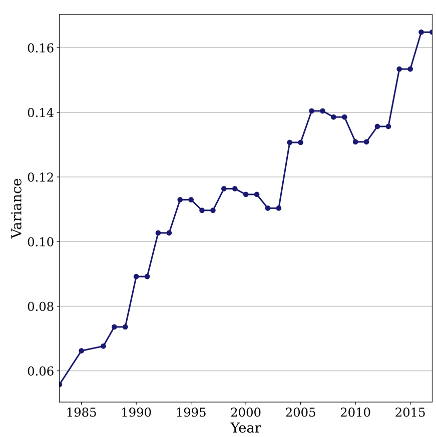

Note: This figure displays the the risk parameters using rolling sample of the PSID. We use sample of 15 years (apart from 1983 and 1984 where we include data from 1970 to 1983 and 1984, respectively). We fix the autocorrelation ρ of persistent shocks to the full sample value estimated in Table B.1. Panel B.2a plots time varying estimates of $\sigma^2_{u,w}$, Panel B.2b plots time varying estimates of $\sigma^2_{v,w}$, Panel B.2c plots time varying estimates of $\sigma^2_{u,r}$, B.2d plots time varying estimates of $\sigma^2_{v,r}$ and B.2e plots time varying estimates of σ^2_{α}. The model is estimated according to Equation (12)-(14) on PSID data from 1970 to 2017.

B.3 Portfolio Shares

B.3.1 Data Source: SCF

The Survey of Consumer Finances (SCF) is a statistical survey of the balance sheet, pension, income and other demographic characteristics of families 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. It includes data from the triennial surveys beginning in 1989.[24] We collect the following variables.

Total Financial Assets. This includes: All types of transaction account (liquid assets), Certificates of deposit, Directly held pooled investment funds (exc. money mkt 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.

Cash & Deposits This includes all types of transaction account (liquid assets) and certificated of deposits. The list of variables are: Money market accounts; Checking accounts (excl. money market); Savings accounts; Call accounts; Prepaid cards; Certificates of deposit.

Equities (direct & indirect). Total value of financial assets held by household that are invested in stock. That includes: directly-held stock, Stock mutual funds: full value if described as stock mutual fund, 1/2 value of combination mutual funds; RAs/Keoghs invested in stock: full value if mostly invested in stock, 1/2 value if split between stocks/bonds or stocks/money market, 1/3 value if split between stocks/bonds/money market; Other managed assets with equity interest (annuities, trusts, MIAs): full value if mostly invested in stock, 1/2 value if split between stocks/MFs & bonds/CDs, or ”mixed/diversified”, 1/3 value if ”other”; Thrift-type retirement accounts invested in stock: full value if mostly invested in stock,1/2 value if split between stocks and interest earning assets.

Real Estate. 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. Businesses (with either an active or nonactive interest). Businesses include both actively and nonactively-managed business(es). Value of active business(es) calculated as net equity if business(es) were sold today, plus loans from the household to the business(es), minus loans from the business(es) to the household not previously reported, plus value of personal assets used as collateral for business(es) loans that were reported earlier. Value of nonactive business(es) is calculated as the market value of the business(es).

Fixed Income. Fixed income is calculated as the residual of Total financial assets minus Cash

[24] 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

& Deposits and Equity (direct & indirect).

Mortgage Debt. This includes: Debt secured by prim. resid. (mortgages, home equity loans, HELOCs); Debt secured by other residential property.

Student Debt. Total value of education loans held by household. This includes education loans that are currently in deferment and loans in scheduled repayment period. We exclude installment loans: these are mostly student loans (which we accounts for separately), vehicle loans (which we do not account as debt as vehicles are part of consumption).

Consumer and Other 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

Net Wealth. We calculate net wealth for each household as the difference between total assets (Cash & Deposits, Equities (direct & indirect), Real Estate, Private Business Wealth and Fixed Income) and total liabilities (Mortgage Debt, Student Debt and Consumer and Other Debt).

B.3.2 Data Source: SCF+

We use the SCF+ database developed by Kuhn et al. (2020) in order to extend our sample backward. Here we detail how we construct income and financial variables to be consistent with the data in the SCF described above.

Income: *tinc*, total household income, excluding capital gains **Real Estate.** This incldues: *house*, asset value of house; *oest*, other real estate (net position); *hoestdebt*, other real estate debt (note: we add back the debt to the other real estate net position).

Cash and Deposits. *liqcer*, liquid assets and certificates of deposit.

Equities. This includes: *ffaequ*, equity and other managed assets; Indirect holdings through mutual funds and pension funds (as described below)

Fixed Income: *ffafin*, financial assets; excluding Equities.

Private Business Wealth: *ffabus*, business wealth.

Indirect Equity holdings

As in the SCF+ the indirect holdings of equity are not available, we follow the methodology of Leombroni, Piazzesi, Schneider, and Rogers (2020) to compute indirect holdings exploiting aggregate information from the US financial account. In order to look-through the mutual funds and pension holdings we use data from the US financial account. We compute the mutual funds equity holdings using Corporate Equities (LM653064100)[25]. We compute the shares dividing equity holdings by total mutual funds assets (LM654090000).

We compute the DC pension total equity holdings as the sum of Corporate equities (LM573064133) and indirect holdings of equity through mutual funds. We compute the pension fund indirect

[25] The code refers to the ones used in the US financial account.

holdings as mutual fund shares (LM573064255) times mutual funds equity shares (as computed above). We then divide the total equity holdings by total DC pension assets (FL574090055) to estimate the pension equity shares. For each household we calculate the indirect equity holdings by multiplying the holdings (in dollar) of mutual funds and pension by the calculated indirect portfolio shares.

Mortgage Debt: *hdebt*, housing debt on owner-occupied real estate; *oestdebt*, other real estate debt.

Student Debt: Student debt is not available in the SCF+. We assume it is 0.

Other Debt: *pdebt*, personal debt.

B.3.3 Wealth Shares and Holdings

We estimate the net-wealth shares held by the top-10% and top-1%. We also estimate gini coefficients. We use the SCF+ database developed by Kuhn et al. (2020) in order to have a longer time series of wealth and income. We slightly modify their definition of total financial net-wealth by subtracting vehicles and other non-financial wealth. Note that our definition of net wealth slightly differ from the SCF: we exclude vehicles, other miscellaneous nonfinancial assets as well as vehicles loans. Figure B.3 plots the top shares and the gini coefficient for our definition of financial (net) wealth.

Figure B3: Financial Wealth Inequality in the SCF+

(a) Top 10% Share (b) Top 1% Share (c) Gini

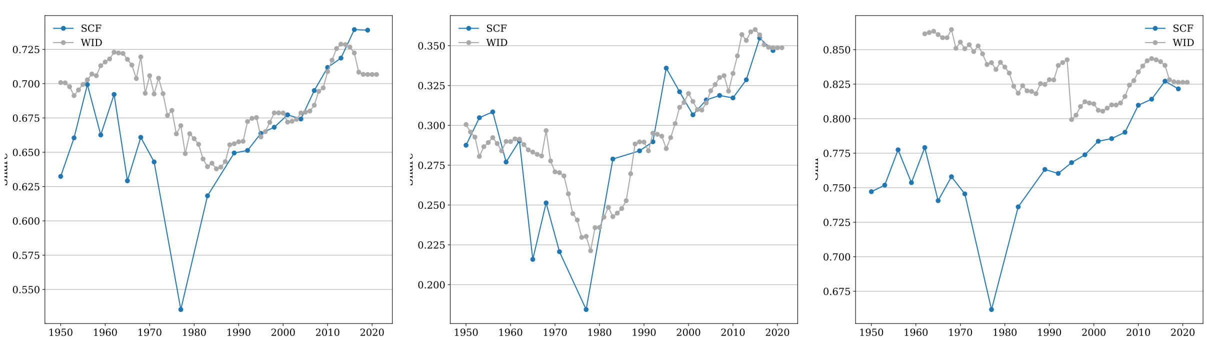

Note: Data are based on SCF/+ database developed by Kuhn et al. (2020), SCF and WID database. We exclude household with negative wealth/income from SCF/SCF/+.

We compute the holdings for each of the assets and liabilities for each household. Table B.2 shows summary statistics for the distribution of asset holdings in 1989, the first available SCF survey.[26]

[26]Note that Private Business Wealth is a measure net of loans from the business to the households and hence may also be a negative number.

Table B2: Portfolio Holdings

	Mean	std	Min	25%	50%	75%	90%	95%	Max
Cash and Deposits	38.08	235.73	0.00	0.99	5.60	24.27	89.25	161.03	67791
Equities	38.67	328.57	0.00	0.00	0.00	7.47	56.02	153.12	103866
Real Estate	206.37	849.42	0.00	9.34	102.70	231.54	429.47	653.54	201100
Private Business Wealth	84.75	1214.02	-13.07	0.00	0.00	0.00	20.54	233.41	258186
Fixed Income	59.26	459.66	-0.00	0.00	4.06	31.74	112.41	235.27	171205
Mortgage Debt	42.36	93.72	0.00	0.00	0.00	56.02	130.71	199.80	30006.82
Student Debt	0.54	3.95	0.00	0.00	0.00	0.00	0.00	1.79	166.19
Other Debt	5.94	51.86	-0.00	0.00	0.07	2.76	9.34	20.54	4201.33
Net Wealth	378.30	1974.83	0.00	26.14	112.41	306.88	711.43	1363.10	290295.01

Note: Data are reported in 2016 thousands dollars. Note that Private Business Wealth is a measure net of loans from the business to the households. For this reason some observations are negative. Data are based on SCF 1989. We exclude household with negative wealth as well as households with zero or negative income.

B.3.4 Financial Duration

We compute household's portfolio share in each asset by dividing the dollar holdings in the asset by the households net wealth. Using the portfolio shares, we compute the durations of the household's financial portfolio by multiplying the asset duration of an asset (assets durations are reported in the first column of Table 2) by the portfolio share of that asset, and summing over all assets in the portfolio. In the one but last row of Table 2, we report the average duration, by averaging over all households (using the SCF sampling weights) for the 1989 SCF survey.

Table 2 also reports value-weighted portfolio shares for each asset. They are obtained by summing dollar holdings of an asset among all households (households in a group) by the total dollar holdings of all assets among all households (households in a group). Aggregate durations are then obtained by multiplying the value-weighted portfolio weights for each asset by the duration of that asset, and summing over assets. The last row of Table 2 reports the aggregate duration for the 1989 SCF survey.

Figure B.4 shows portfolio shares by age in the 1989 SCF. We bundle households into different cohort groups: 25-35, 35-45, 45-55, 55-65, 65-75, 75-85. Figure B.4a uses the value-weighted portfolio shares. Figure B.4b plots the median portfolio share in each asset category, and then rescales the resulting shares so that they sum to 100%.

Figure B.5 provides further information on the distribution of durations across households. Figure B.5a plots the average duration by cohort. We bundle households into cohort groups and estimate the average duration. Figure B.5b bundles households in wealth-weighted percentile and estimate the average duration of households in each bin.[27] Panel B.5c and Panel B.5d rank households according to their income and estimate the average duration of each group against the

[27] Households are ranked according to their net-wealth and allocated to different bins. Each bin is designed such that the share of total wealth held by the households in each bin is the same across different bins.

Figure B4: Portfolio Shares by Cohorts

(a) Wealth-Weighted Portfolio Shares (b) Median Portfolio Shares

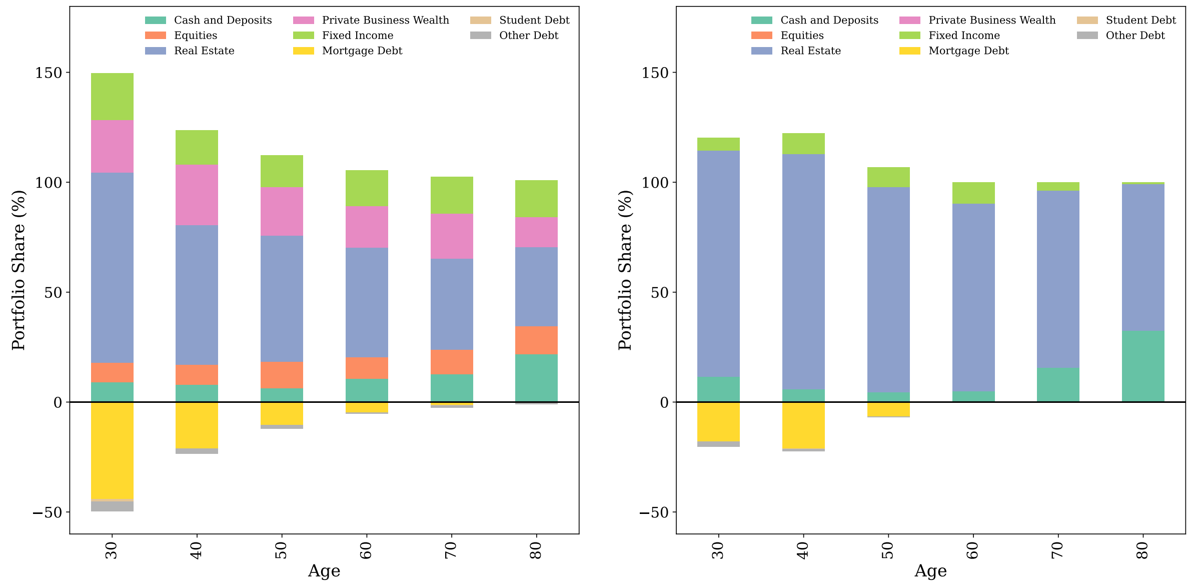

Note: Portfolio shares by age in the 1989 SCF. We bundle households into different cohort groups: 25-35, 35-45, 45-55, 55-65, 65-75, 75-85. The top panel uses the value-weighted portfolio shares. The bottom panel uses the median portfolio share in each asset category, and then rescales the resulting shares so that they sum to 100%. Data are based on SCF 1989. We exclude household with negative wealth as well as households with zero or negative income. We winsorize the top/bottom 2.5% of households ranked by the duration of their portfolio.

average in income (B.5c) or the income percentile (B.5d).

Figure B.6 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 also evaluate more formally the correlation between financial duration and some covariates of interest. Table B.3 reports the estimation results. 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. Data are based on SCF 1989. We exclude household with negative wealth as well as households with zero or negative income. We winsorize the top/bottom 2.5% of households ranked by the duration of their portfolio. The regression estimate take into account survey weights.

Figure B5: Distribution of Durations

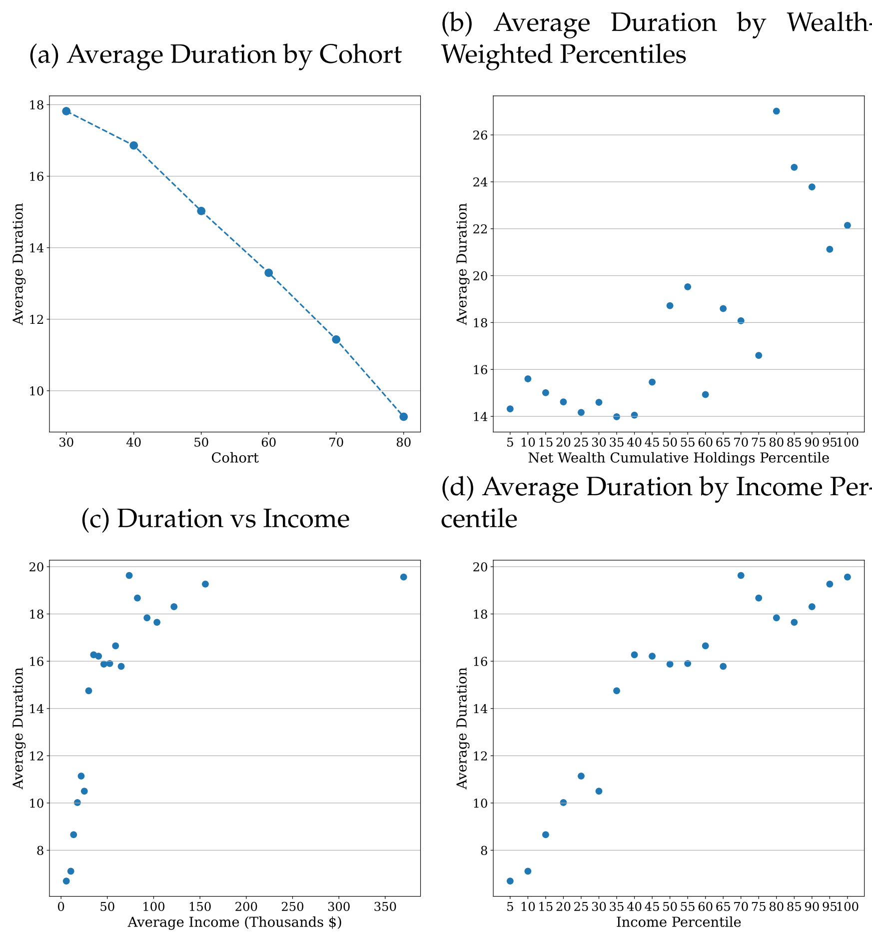

(a) Average Duration by Cohort

(b) Average Duration by Wealth-Weighted Percentiles

(c) Duration vs Income

(d) Average Duration by Income Percentile

Note: Panel B.5a plots the average duration by cohort. We bundle households into cohort groups and estimate the average duration. Panel B.5b bundles households in wealth-weighted percentile and estimate the average duration of households in each bin. Panel B.5c and Panel B.5d rank households according to their income and estimate the average duration of each group against the average in income (B.5c) or the income percentile (B.5d). Data are based on SCF 1989. We exclude household with negative wealth as well as households with zero or negative income. We winsorize the top/bottom 2.5% of households ranked by the duration of their portfolio.

B.3.5 Financial Duration Over Time

Figure B.7 uses information from each SCF survey from 1989 till 2019 to measure average and aggregate duration over time. Figure B.7a compute the aggregate (wealth-weighted) while B.7b computes the average (equally-weighted) duration over time. We use two different specifications for the duration of assets. Full sample computes the duration of the asset using the information over the whole sample; the duration of each asset is kept constant over time. The time varying specification computes time varying duration measures for equity, private business wealth and real estate. We then use these time varying measures to compute the portfolio duration. Table B.4 compute the numbers reported in Table 2 for the full sample.

Figure B6: Distribution of Durations

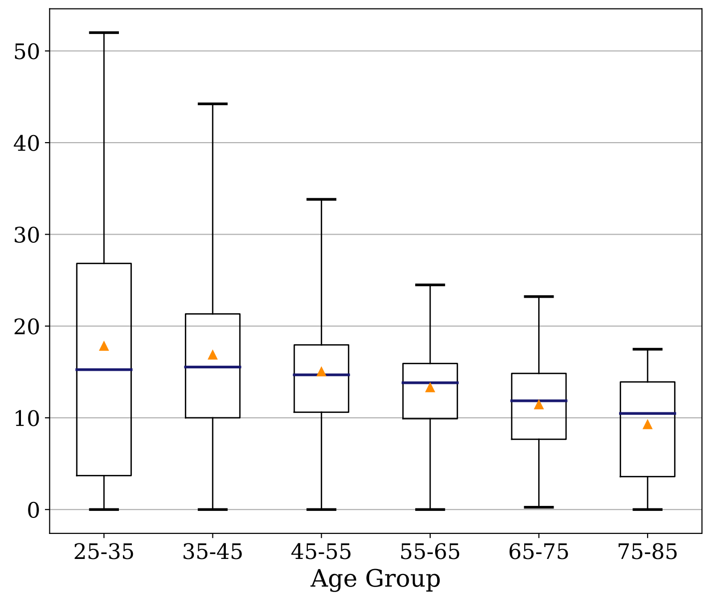

Note: For each age group, we rank households by the duration of their portfolio. The figure includes the 5%, 25%, 50%, 75% and 95% percentile, as well as the median (blue line) and mean (green triangle). Data are based on SCF 1989. We exclude household with negative wealth as well as households with zero or negative income. We winsorize the top/bottom 2.5% of households ranked by the duration of their portfolio.

Table B3: Determinants of Household-level Financial Duration

	(1)	(2)	(3)	(4)	(5)	(6)
Age	-0.16***		-0.18***	-0.0026	-0.14**	-0.23***
	(0.0086)		(0.0082)	(0.061)	(0.058)	(0.054)
Lorenz Curve		0.056***	0.089***	0.085***	0.0018	-0.11***
		(0.0072)	(0.0067)	(0.0067)	(0.0082)	(0.0079)
Age Squared				-0.0016***	0.00036	0.00061
				(0.00052)	(0.00050)	(0.00046)
Log-Income					3.04***	1.33***
					(0.15)	(0.13)
Log-Net-Wealth						1.43***
						(0.032)
Constant	23.0***	14.2***	22.8***	18.6***	-11.6***	-3.64*
	(0.53)	(0.19)	(0.54)	(1.68)	(2.24)	(1.98)
Observations	14161	14161	14161	14161	14161	14161
R^2	0.054	0.006	0.069	0.070	0.118	0.208

Note: Data are based on SCF 1989. We exclude household with negative wealth as well as households with zero or negative income. We winsorize the top/bottom 2.5% of households ranked by the duration of their portfolio. Standard Errors in parentheses (* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$)

Figure B7: Financial Duration Over Time

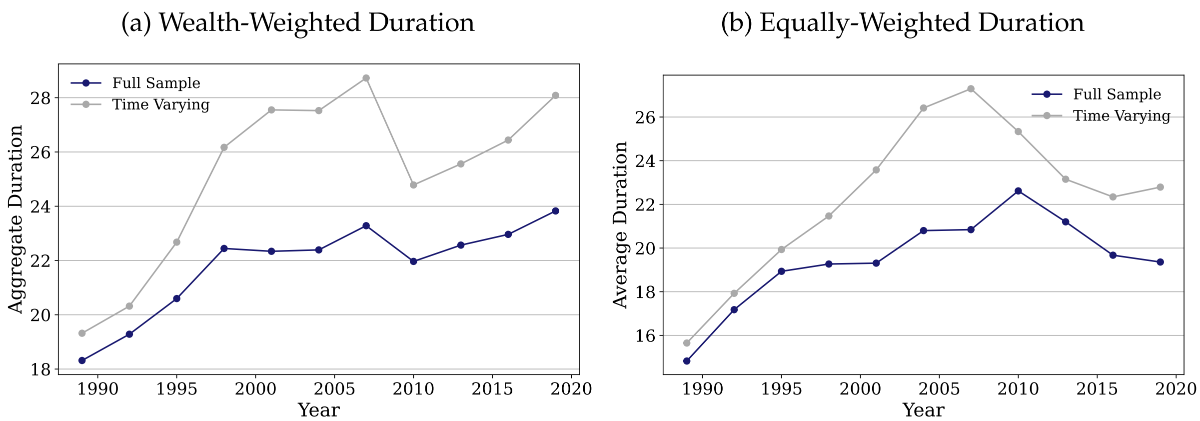

Note: Figure B.7 uses information from each SCF survey from 1989 till 2019. Figure B.7a compute the aggregate (wealth-weighted) while B.7b computes the average average (equally-weighted) duration over time. We use two different specifications for the duration of assets. Full sample computes the duration of the asset using the information over the whole sample; the duration of each asset is kept constant over time. The time varying specification computes time varying duration measures for equity, private business wealth and housing.

B.3.6 Financial Duration - Robustness

To make sure our results are robust we replicate Figure 3 using different duration estimates for Private Business Wealth (PBW). Panel B.8a uses our baseline measure of duration: the duration of PBW Short is equal to 10 and the duration of PBW Long is equal to the duration of small stocks estimated over the full sample, 61.25. Panel B.8b uses the time varying measure of duration for PBW long. Panel B.8c uses a measure of duration for PBW Long equal to 52 (as estimated in Appendix B.4), Panel B.8d uses a measure of duration for both PBW Short and PBW Long equal to the duration for equity (28.78).

B.3.7 Private Business Wealth Measurement in the SCF

The SCF provides more granular information on private business wealth. We collect data on i) the type of business, ii) the value of the business and iii) the number of employees for active business. We define as large a business employing more than 10 people. We split the sample in net-wealth percentile 90 to 99 (p90-p99) and top 1% (Top1).

For each active business, we follow the methodology used by the SCF to measure the overall total value of private business wealth. From 1989 to 2010, SCF provides information on the three largest active business while from 2010 onward it provides information on the two largest active business. We now provide further details on the data construction: in parenthesis we provide information on the ticker 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

Table B4: Pooled Duration of the Household Financial Wealth Portfolio

	Duration	Portfolio Shares			
		All	Bottom 90	P90-P99	Top 1
Assets					
Cash and Deposits	0.25	0.08	0.09	0.08	0.06
Equities	35.50	0.22	0.21	0.26	0.26
Real Estate	16.36	0.48	0.64	0.49	0.34
PBW Short	10.00	0.09	0.07	0.07	0.09
PBW Long	69.82	0.12	0.07	0.09	0.20
Fixed Income	4.00	0.16	0.17	0.17	0.14
Liabilities					
Mortgage Debt	5.2	0.13	0.22	0.14	0.08
Student Debt	4.5	0.00	0.01	0.01	0.00
Other Debt	1.0	0.01	0.01	0.01	0.01
Average (EW) Duration		**19.46**	**19.43**	**19.12**	**24.84**
Aggregate (VW) Duration		**25.20**	**20.93**	**23.04**	**31.16**

Note: The column "Duration" reports the duration of the asset, averaged over all quarters from 1983.Q1 to 2019.Q4. For Equities, Private Business Wealth, and Real Estate, the durations are computed form the asset pricing model in Appendix A. The columns "Portfolio Shares" reports the wealth-weighted average portfolio weights from the Survey of Consumer Finances for i) all households (All), ii) subsample of households with net-wealth in the bottom 90 percentile (Bottom 90), iii) subsample of households with net-wealth in the 90th to 99th percentile (P90-P99), iv) subsample of households with net-wealth in the top 1 percentile (Top 1). Liabilities receive negative portfolio weights. The numbers are average from 1983.Q1 to 2019.Q4. Averages are computed using the SCF weights, rescaled so that weights sum to unity in all years. Top-10% and Top-1% for the aggregate statistics are computed year-by-year before averaging.

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 owe the households (#1 X3124, #2 X3224, #3 X3324), minus the amount the household owe the business (# 1 X3126, #2 X3226, #3 X3326) if it was not reported earlier, plus the amount of this loan that is collateralized/guaranteed (#1 X3121, #2 X3221, #3 X3321). We then divide in the different categories based on variable (#1 X3119, #2 X3219, #3 X3319) that defines the type of business. We further collect the number of employees (#1 X3111, #2 X3211, #3 X3311) and based on that we split into small and large businesses.

For the passive business, we compute directly the value of: limited partnership (X3408), other partnership (X3412), S-Corporations (X3416), other corporations (X3420), sole-proprietorships (X3424), LLCs (X3452), other types (X3428).

Table B.5 reports the Wealth-weighted shares in each type of business wealth. For our baseline measure of private business wealth duration, we divide private businesses in two types: long-duration type (with a duration equal to the baseline measure of duration of small stocks estimated to be 61.25) or low-duration type (with a low duration of 10). The shaded rows indicate our

Figure B8: Duration by Net Worth Percentile, Robustness

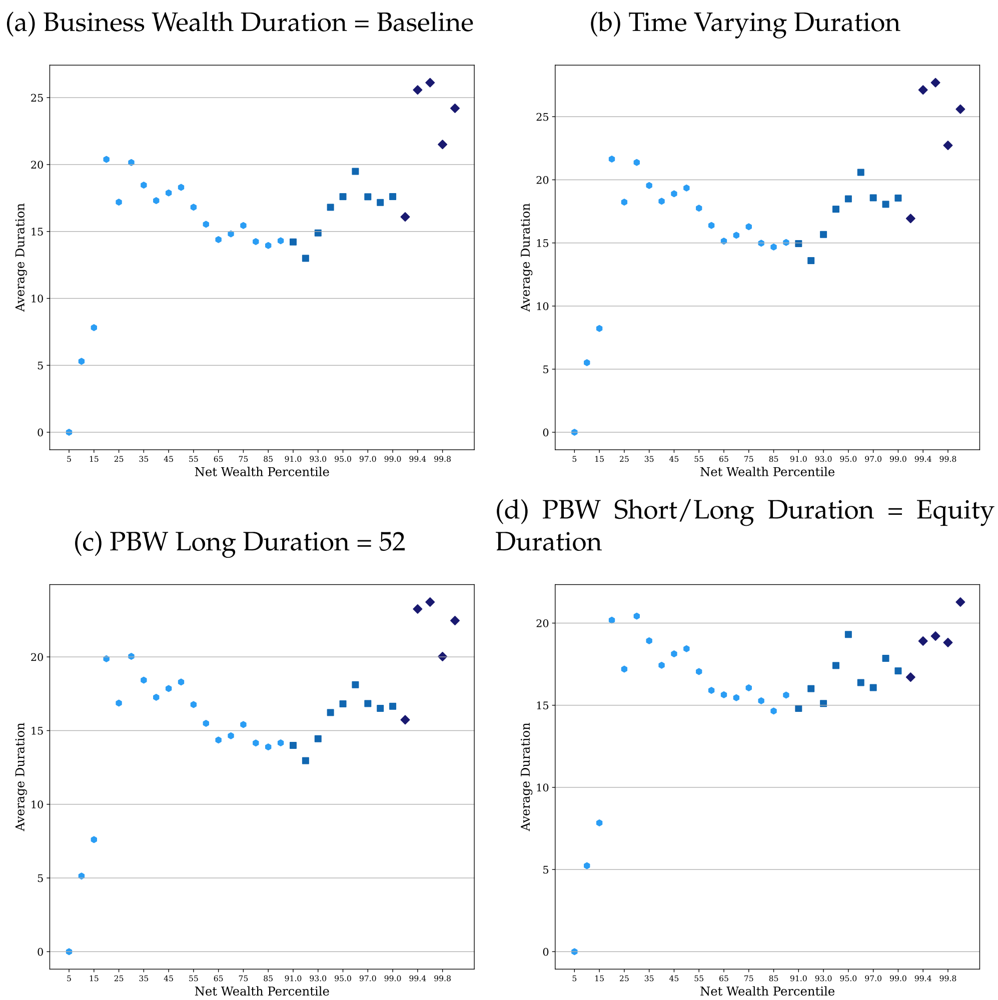

(a) Business Wealth Duration = Baseline

(b) Time Varying Duration

(c) PBW Long Duration = 52

(d) PBW Short/Long Duration = Equity Duration

Note: Data are based on SCF 1989. We exclude household with negative wealth as well as households with zero or negative income. We winsorize the top/bottom 2.5% (if not otherwise specified) of households ranked by the duration of their portfolio. Panel B.8a uses our baseline measure of duration: the duration of PBW Short is equal to 10 and the duration of PBW Long is equal to the duration of small stocks estimated over the full sample, 61.25. Panel B.8b uses the time varying measure of duration for PBW long. Panel B.8c uses a duration for PBW Long equal to 52 (as estimated in Appendix B.4), Panel B.8d uses a measure of duration for both PBW Short and PBW Long equal to the duration for equity (28.78).

classification into the long-duration type. The non-shaded rows are the short-duration private businesses. This classification is based on legal form as well as size (number of employees). The bottom of the table reports the total share of business wealth in long-duration types. For some of the private business wealth, we do not have information on the type of business (e.g. Active Other Large, Active Other or Passive Other); we assigned a duration of 10 to those business. That means we are computing a conservative measure of the overall duration.

	1989		2001		2010		2019	
	P90-P99	Top1	P90-P99	Top1	P90-P99	Top1	P90-P99	Top1
Active Partnership Small	16.87	8.70	5.93	5.18	3.51	4.48	1.97	0.64
Active Sole Partnership Small	18.38	12.05	12.76	4.76	12.61	1.16	6.93	1.49
Active S Corporation Small	7.01	2.69	13.22	8.95	10.35	3.07	8.15	3.12
Active Other Corporation Small	8.36	7.99	7.90	6.69	5.63	0.58	3.96	3.47
Active Foreign Small	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Active LP Small	0.00	0.00	2.33	2.63	20.05	13.26	21.37	10.56
Active Informal Small	0.00	0.00	0.00	0.00	0.00	0.07	0.00	0.00
Active Other Small	1.61	0.10	0.00	0.00	0.00	0.00	0.00	0.00
Active Partnership Large	5.50	7.24	7.60	12.23	1.80	3.08	2.03	2.80
Active Sole Partnership Large	4.93	6.24	7.30	3.54	4.23	1.58	1.23	1.95
Active S Corporation Large	9.68	13.73	10.93	21.86	7.87	25.67	10.78	24.28
Active Other Corporation Large	8.31	25.49	16.54	13.31	6.23	11.74	9.16	5.16
Active Foreign Large	0.00	0.00	0.00	0.00	0.00	0.04	0.00	0.00
Active LP Large	0.00	0.00	1.23	3.28	9.37	16.19	17.39	22.90
Active Informal Large	0.00	0.00	0.00	0.00	0.00	0.03	0.00	0.00
Active Other Large	0.00	1.56	0.00	0.00	0.00	0.00	0.00	0.00
Active Non Categorized	0.11	2.29	0.45	4.19	3.41	3.58	3.51	6.31
Farm	10.68	2.73	4.58	0.50	4.11	0.63	3.49	0.89
Passive LP	3.68	3.99	2.40	5.79	2.93	3.81	4.13	8.73
Passive Other Partnership	0.90	1.73	0.95	2.53	0.63	0.42	0.39	0.47
Passive S Corporation	0.83	1.13	3.72	2.22	3.54	2.96	0.86	1.58
Passive Other Corporation	1.08	1.01	0.34	1.23	2.01	3.86	1.64	2.92
Passive Sole Partnership	1.90	0.45	0.98	0.17	0.00	0.00	0.00	0.00
Passive Other	0.17	0.87	0.86	0.94	0.58	0.73	0.23	0.41
Passive LLC	0.00	0.00	0.00	0.00	1.13	3.05	2.76	2.31
Long-Duration Share	38.95	56.03	56.27	63.32	49.06	70.93	58.83	74.48

Table B5: Private Business Wealth Decomposition

Note:
Data are based on SCF. We split the sample in net-wealth percentile 90 to 99 (p90-p99) and top 1% (Top1). The table reports the Wealth-weighted shares in each type of business wealth. The shaded rows are the long-duration types. The bottom of the table reports the total share of business wealth in long-duration types.

B.4 Measurement of Private Business Duration in CRSP-Computstat

Our benchmark model uses the duration of small stocks—those in the bottom decile of the market capitalization distribution of publicly listed firms—as a proxy for the duration of private business wealth. Since the price-dividend ratio and dividend growth rates of small stocks are included in the state vector of the auxiliary asset pricing model, the latter model fits the post-war quarterly time series of small stock returns and cash flow growth rates exactly and implies a quarterly time series for the duration of small stocks. The resulting private business duration, averaged over the 40 quarters of the 1980s, is 61.25.

This approach may be understating the duration of private business wealth, to the extent that firms in the bottom 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 point of 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. In reality, firms in the bottom decile of publicly-listed firms may grow further and transition into higher deciles of the market capitalization distribution. Since larger firms may have lower cash flow pay-out ratios, taking into account that small firms do not remain small may lead us to overstate the duration of small public firms. In this appendix, we address this potential overstatement issue, and explain how to arrive at 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.

B.4.1 Measuring Duration with Firm Life-Cycles

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

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

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:

$$CF_t(s_t|s_{t-1}) = (CF_t/A_t)(s_t|s_{t-1}) \cdot A_t(s_t|s_{t-1}) \tag{23}$$

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:

$$D_t|s_0 = \sum_{s_t=1}^{S} \mathcal{P}^{t-1} \cdot (\mathcal{P} \cdot (D_t/A_t) \cdot A_t) \tag{24}$$

B.4.2 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.

B.4.3 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 B.6 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.

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.

Table B6: 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.

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 B.7 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.

Table B7: 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.

Combining the results for deciles and quintiles suggests that a value around 50 for the duration of private business wealth is conservative. This is particularly true, given the concern of understatement of private business durations mentioned at the beginning of this appendix.

B.5 Household Duration Profile

This section provides additional context and analysis of the pattern of duration by household wealth displayed in Figure 3. In particular, the durations display a non-monotonic pattern, increasing sharply up to around the 20th percentile or so, then decreasing until around the 92nd percentile, then increasing again at the very top of the distribution. Figure B.9 Panel (a) shows that this pattern is even stronger when evaluated over the full sample, rather than the 1980s starting point used for the one-shot experiment.

First, we show why the value-weighted duration is so much higher than the equal-weighted duration, despite little perceivable trend in duration between the 20th and 99th percentiles. Figure B.9 Panel (b) displays the weights used in the two calculations. Note that because the bins are not equally sized, the equal weights are in fact smaller for the higher bins. As a result, the equal-weighted calculation puts substantial weight not only on the lowest bins, which have very low duration, but also on the bins in the 60th - 90th percentiles, which also have relatively low duration. In contrast, the value weights in Figure (b) show that the value-weighted calculation is very heavily weighted toward the top 1% of the distribution, which is where durations are their highest, driven by increasing portfolio weights on equities and long-duration private business wealth.

Next, we analyze the non-monotonic duration pattern itself, in particular why households with net worth in the 20th - 30th percentile have such high durations. This turns out to be the result of mortgage leverage. Because housing wealth has a higher duration than prepayable mortgage debt, the home equity in a house financed by a mortgage has a much higher duration than the house itself. This area of the net wealth distribution is heavily weighted toward households with high mortgage leverage, leading to these high durations.

To see this in more detail, Figure B.10 displays durations by net wealth, split by whether households have high or low mortgage leverage. Specifically, we define loan-to-value (LTV) to be the ratio of mortgage debt to real estate assets (set equal to zero for households without real estate) and split the sample by whether LTV is above or below 50%. Panel (a) shows that, once high-lTV households are excluded, the duration profile becomes uniformly increasing in wealth. However, Panel (b) shows a sharply different picture for high-LTV households. In particular, low net wealth households with mortgages tend to have very high LTVs, leading potentially very high durations via leverage effects. Panel (b) shows that durations for this subsample are in fact *decreasing* with wealth as leverage is decreasing in wealth, even in this more levered subsample.

The overall sample is a weighted combination of these two subsamples, where the weights depend on the share of high-LTV households. This share is displayed in Figure B.10 Panel (c), which shows that the high-LTV share is very low at low levels of wealth, as the poorest households mostly rent, peaks at around the 30th percentile of net wealth, and then decreases beyond that

Figure B9: Net Wealth vs. Duration, Detail

(a) Net Wealth vs. Duration, Full Sample

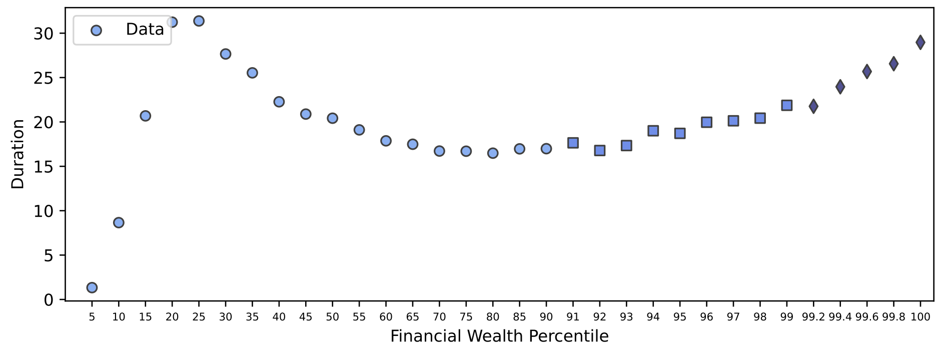

(b) Equal Weights vs. Value Weights by Net Worth Bin

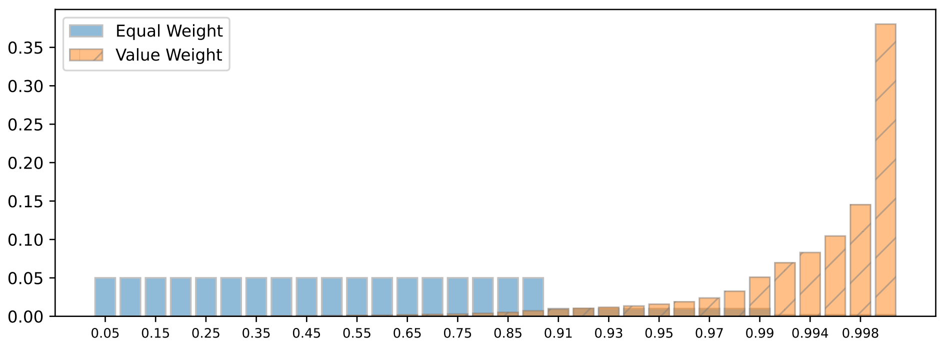

Note: These plots display durations and weights for the full 1983 - 2019 sample. The portfolio shares and net wealth values are obtained from the SCF, while the durations for each asset class are obtained as in Table 2. Value weights are weighted by net wealth.

point. Combining these shares with the durations conditional on share leads to the overall two-peaked profile observed in Figures 3 and B.9.

In summary, our overall duration profile is a combination of two patterns: an increasing profile for households without high mortgage leverage, and a decreasing profile for households with high mortgage leverage. The first effect dominates, leading to our main result that inequality increases with wealth. However, these results imply that there exists a nontrivial subset of households with high mortgage leverage who may have benefited greatly from falling interest rates, a topic deserving of further research in the future.

Figure B10: Equal Weights vs. Value Weights by Net Worth Bin

(a) Durations: LTV ≤ 50

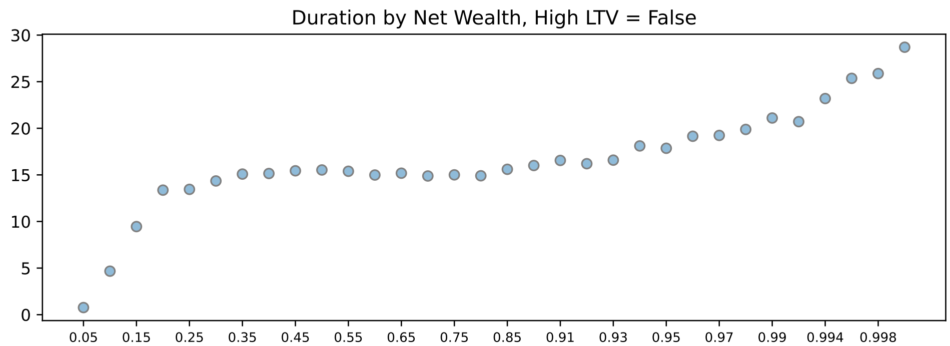

(b) Durations: LTV > 50

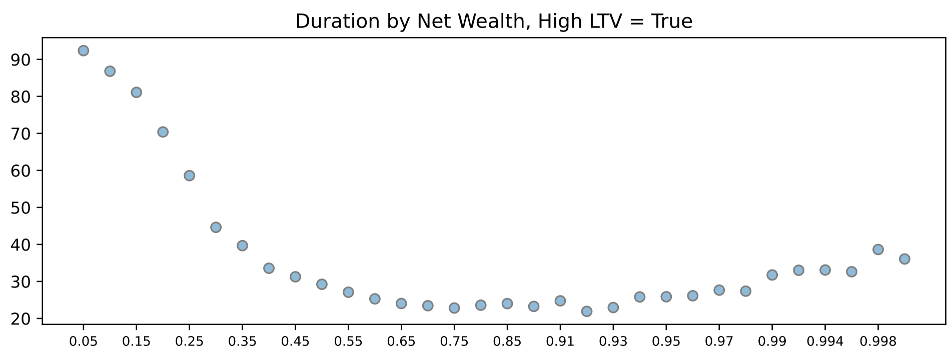

(c) Share with LTV > 50

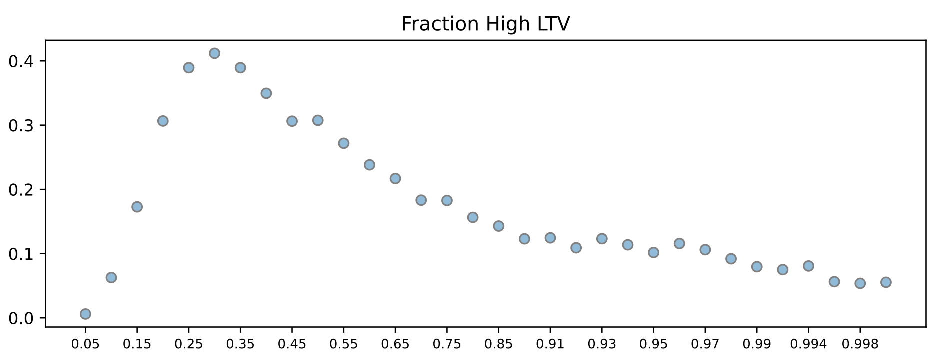

Note: These plots display durations and population shares split by LTV, computed as mortgage debt divided by real estate assets (source: SCF). For households without real estate assets, LTV is defined to be zero. Duration computations otherwise follow those of Figures 3 and B.9.

C Proofs

C.1 Proof of Lemma 5.1

Rewriting (7) yields

$$\log Z_0 = \log \left\{ \sum_{t=0}^{\infty} \exp(-t \times \log R) z_t \right\}.$$

Taking the derivative, we now obtain

$$\begin{aligned}
\frac{\partial \log Z_0}{\partial \log R} &= \left\{ \sum_{t=0}^{\infty} \exp(-t \times \log R) z_t \right\}^{-1} \sum_{t=0}^{\infty} (-t) \exp(-t \times \log R) z_t \\
&= Z_0^{-1} \sum_{t=0}^{\infty} (-t) R^{-t} z_t \\
&= -\frac{\sum_{t=0}^{\infty} R^{-t} z_t \times t}{Z_0}.
\end{aligned}$$

The approximations at the end of the lemma follow immediately, except possibly for the last line

$$\tilde{Z}_0 \simeq Z_0(1 - D \times d \log R)$$

which follows from

$$\tilde{Z}_0 \simeq Z_0 \exp(-D \times d \log R)$$

and the approximation $h \simeq \log(1+h)$ as $h \to 0$.

C.2 Proof of Proposition 5.2

Beginning with part (a), note that household wealth growth is equal to

$$\frac{\tilde{\theta}_i}{\theta_i} = (1 - D_i^{\theta} d \log R)$$

which implies

$$\text{Cov}\left(\frac{\tilde{\theta}_i}{\theta_i}, \theta_i\right) = -d \log R \times \text{Cov}(D_i^{\theta}, \theta_i).$$

Since $d \log R < 0$, this implies that the covariance of wealth growth and wealth is positive if and only if the covariance of duration and wealth is positive. Since wealth-weighted duration is defined by

$$\begin{aligned}
D^{\theta,VW} &= \int D_i^{\theta} \frac{\theta_i}{\bar{\theta}}\, di \\
&= \bar{\theta}^{-1} \int D_i^{\theta}\, di \int \theta_i\, di + \bar{\theta}^{-1} \mathrm{Cov}(D_i^{\theta}, \theta_i) \\
&= \int D_i^{\theta}\, di + \bar{\theta}^{-1} \mathrm{Cov}(D_i^{\theta}, \theta_i) \\
&= D^{\theta,EW} + \bar{\theta}^{-1} \mathrm{Cov}(D_i^{\theta}, \theta_i)
\end{aligned}$$

where

$$\bar{\theta} \equiv \int_i \theta_i\, di.$$

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

Turning to part (b), let A represent the top-α share of households ranked by wealth. Then

$$\begin{aligned}
\tilde{\Theta}_A &\equiv \int_A \tilde{\theta}_i\, di \\
&= \int_A \theta_i \exp(-D_i^{\theta} d \log R) \\
&= \int_A \theta_i (1 - D_i^{\theta} d \log R) \\
&= \Theta_A - \int_A \theta_i D^{\theta} d \log R \\
&= \Theta_A \left(1 - D_A^{\theta,VW} d \log R\right).
\end{aligned}$$

Similarly

$$\tilde{\Theta} = \Theta \left(1 - D^{\theta,VW} d \log R\right).$$

So the top-α share is given by

$$\tilde{S}_A = \frac{\tilde{\Theta}_A}{\tilde{\Theta}} = \frac{\Theta_A \left(1 - D_A^{\theta,VW} d \log R\right)}{\Theta \left(1 - D^{\theta,VW} d \log R\right)} = S_A \left(\frac{1 - D_A^{\theta,VW} d \log R}{1 - D^{\theta,VW} d \log R} \right).$$

For a negative change in interest rates ($d \log R < 0$), the fraction on the right hand side is greater (less) than unity if the value-weighted duration of the top-α wealth households is greater (less) than that of all households, completing the proof.

C.3 Proof of Proposition 5.3

We begin by proving (a). First, consider any consumption plan $\{c_t\}$ that is budget-feasible under the original interest rate. Then, applying Equation 8, we find:

$$\log \tilde{\theta}_0 = \log \theta_0 - D^{\theta} \times d \log R \tag{25}$$

Equivalently, for the valuation of excess consumption claim, under the new interest rate $\log \tilde{R}$, we apply Equation 8 and find:

$$\log \left\{ \mathbb{E}_0 \sum_{t=0}^{\infty} \tilde{R}^{-t}(c_t - y_t) \right\} = \log \left\{ \mathbb{E}_0 \sum_{t=0}^{\infty} R^{-t}(c_t - y_t) \right\} - D^{c-y} \times d \log R. \tag{26}$$

That also means, inverting Equation 26:

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

and using:

$$\log \theta_0 = \log \left\{ \mathbb{E}_0 \sum_{t=0}^{\infty} R^{-t}(c_t - y_t) \right\} = \log \left\{ \mathbb{E}_0 \sum_{t=0}^{\infty} \tilde{R}^{-t}(c_t - y_t) \right\} + D^{c-y} \times d \log R$$

we rewrite equation 25:

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

Combining this with the updated budget constraint:

$$\log \tilde{\theta}_0 = \log \left\{ \mathbb{E}_0 \sum_{t=0}^{\infty} \tilde{R}^{-t}(\tilde{c}_t - y_t) \right\} \tag{27}$$

we obtain

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

If $D^{\theta} > D^{c-y}$, with $d \log R < 0$, we have:

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

and hence, as $\{y\}$ is unchanged:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \tilde{R}^{-t} \tilde{c}_t > \mathbb{E}_0 \sum_{t=0}^{\infty} \tilde{R}^{-t} c_t.$$

that implies that new consumption allocation $\{\tilde{c}_t\}$ with $\tilde{c}_t \geq c_t$ in all periods and $\tilde{c}_t > c_t$ in some periods are affordable under the updated budget constraint (27).

If $D^{\theta} < D^{c-y}$, with $d \log R < 0$, we have:

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

and hence, as $\{y\}$ is unchanged:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \tilde{R}^{-t} \tilde{c}_t < \mathbb{E}_0 \sum_{t=0}^{\infty} \tilde{R}^{-t} c_t.$$

that implies that the original allocation $\{c_t\}$ exceeds the budget constraint.

Turning to part (b), it is clear that for $D^{\theta} < D^{c-y}$ the original consumption plan $\{c_t\}$ is no longer affordable following the change in rates, while for $D^{\theta} > D^{c-y}$ the household should pick a sequence with $\tilde{c}_t > c_t$ in some states due to local non-satiation. Thus, a necessary condition that $\tilde{c}_t = c_t$ in all states is that $D^{\theta} = D^{c-y}$. We next show that this condition is sufficient. This follows from the budget constraint (27) and Euler equation being exactly satisfied, due to our assumption that $\tilde{\beta}\tilde{R} = \beta R$.

D 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 D.5. They characterize how a decline in rates affects wealth inequality.

D.1 Model Setup

D.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.

D.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 $0 < \beta < 1$.

D.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:

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

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.

D.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:

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

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:

$$U(\widehat{c})(z^t,\eta^h) = \frac{\widehat{c}(z^t,\eta^h)^{1-\gamma}}{1-\gamma} + \sum_{z_{t+1}}\widehat{\beta}(z_{t+1},z^t)\widehat{\phi}(z_{t+1}|z^t)\sum_{\eta_{h+1}}\varphi(\eta_{h+1}|\eta^h)U(\widehat{c})(z^{t+1},\eta^{h+1}) \tag{28}$$

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:

$$\widehat{\beta} = \beta\sum_{z_{t+1}}\phi(z_{t+1})\lambda_{t+1}(z_{t+1})^{1-\gamma}. \tag{29}$$

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:

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

D.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 (28) 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 1. 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:

$$\widehat{R}_t = \frac{\widehat{v}_{t+1} + \alpha}{\widehat{v}_t}. \tag{30}$$

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.

D.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 D.1. 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 ($\widehat{R}_t$):

$$R_t = \left(\sum_{z_{t+1}} q_t(z^t, z_{t+1}) \right)^{-1} = \left(\sum_{z_{t+1}} \frac{\widehat{\phi}(z_{t+1})}{\lambda(z_{t+1})} \right)^{-1} \widehat{R}_t. \tag{31}$$

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

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

D.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 (10) 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:

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

D.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 D.2. 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:

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

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 D.3. 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 D.4. 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.[28] 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.

D.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 D.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 D.5. 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

[28] 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

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

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:

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

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:

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

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 D.6. 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 D.7. 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 D.7 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 D.8. 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 6 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.

D.6 Proofs of Propositions in this Appendix

D.6.1 Proof of proposition D.2

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 (30), as:

$$\widehat{c}_t(\eta^t) - (1-\alpha)\widehat{y}_t(\eta^t) + \frac{\widehat{a}_t(\eta^t) + \widehat{\sigma}_t(\eta^t)(\widehat{v}_{t+1} + \alpha)}{\widehat{R}_t} = \widehat{a}_{t-1}(\eta^{t-1}) + \widehat{\sigma}_{t-1}(\eta^{t-1})(\widehat{v}_t + \alpha). \tag{32}$$

Rewriting (32) 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 (32) 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. □

D.6.2 Proof of Proposition D.3

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

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

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 $\mathbb{E}_0[\psi_t] = 1$. □

D.6.3 Proof of Proposition D.4

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 $\sum_{\tau=0}^{\infty} \widehat{R}_{0\to\tau-1}^{-1}$. □

D.6.4 Proof of proposition D.5

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. □

E Additional Model Results

Table E1: Change in Inequality, Transition Experiment (Levels)

	Data		Model		
	Initial	After	Initial	Repriced	Comp.
Top-10% share FW	62.4%	70.8%	62.3%	68.5%	59.2%
Top-1% share FW	23.8%	35.1%	38.7%	46.7%	37.8%
Gini FW	0.772	0.826	0.740	0.779	0.700
Gini HW	–	–	0.384	0.454	0.454
Top-10% share HW	–	–	27.5%	28.9%	28.9%
Top-1% share HW	–	–	12.6%	10.7%	10.7%
Gini TW	–	–	0.384	0.459	0.441
Top-10% share TW	–	–	33.5%	33.6%	31.6%
Top-1% share TW	–	–	18.7%	17.2%	15.5%

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. For model results, the columns represent the pre-shock wealth distribution ("Initial"), the repriced distribution ("Repriced"), and the compensated distribution ("Comp").

Figure E1: Lorenz Curves

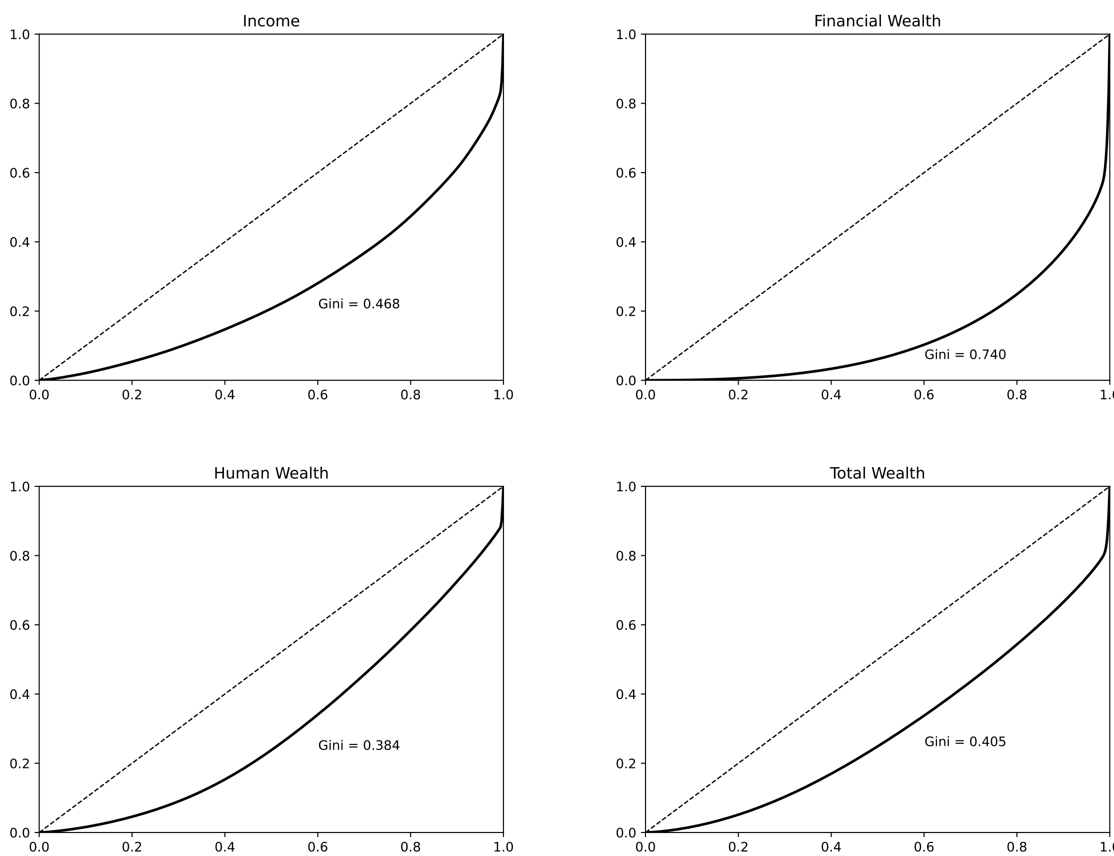

Note: This figure plots the Lorenz curve for each variable, obtained from a long simulation of the model.

Figure E2: Binscatters by Wealth, Controlling for Age

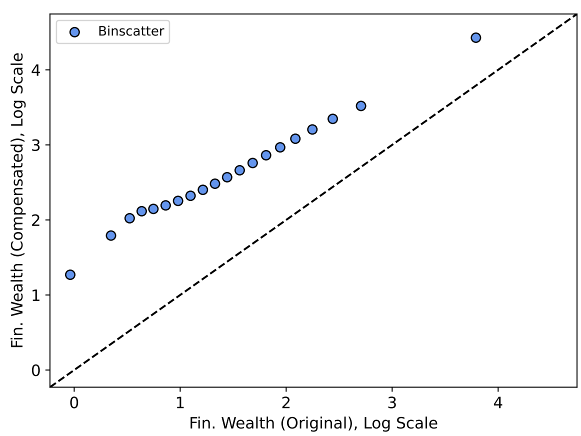

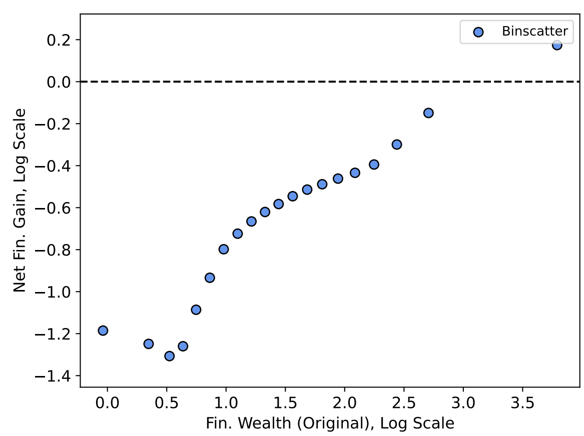

Note: Panel (a) plots the distribution of original financial wealth against the distribution of compensated financial wealth by age. Each dot represents 5% of the original financial wealth distribution. Both variables are plotted using the transform $\log(1+x)$. Panel (b) similarly plots medians for 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.