Wealth Inequality with Declining Interest Rates*

Daniel L. Greenwald
NYU Stern, NBER

Matteo Leombroni
Boston College

Hanno Lustig
Stanford GSB, NBER

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

September 24, 2026

Abstract

US wealth inequality and long-term real interest rates exhibit a strong negative correlation over the post-war period. We quantify how much of the observed increase in wealth inequality from 1983 to 2023 can be accounted for by the decline in rates. To do so, we combine asset holdings data with asset exposures to interest rates to measure the exposure of households' portfolios to interest rates. The portfolios of the wealthy have higher interest rate exposure due to a tilt toward equity-like assets with long duration. As a result, wealth inequality increases when rates fall. When we feed in the observed path of real interest rates, we find that this revaluation effect explains the majority of the increase in measured wealth inequality over the past forty years.

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

*First draft: March 2021. Greenwald: Department of Finance, NYU Stern School of Business, 44 W 4th Street, New York, NY 10012 and NBER; dlg340@stern.nyu.edu. Leombroni: Department of Finance, Boston College Carroll School of Management, 140 Commonwealth Ave, Chestnut Hill, MA 02467; matteo.leombroni@bc.edu. Lustig: Department of Finance, Stanford Graduate School of Business, Stanford, CA 94305 and NBER; hlustig@stanford.edu. Van Nieuwerburgh: Department of Finance, Columbia Business School, Columbia University, 3022 Broadway, New York, NY 10027, NBER, CEPR, and ABFER; svnieuwe@gsb.columbia.edu. The authors thank Sebastian Hanson and Yuan Tian for excellent research assistance. We thank Andy Atkeson, Adrien Auclert, Ravi Bansal, Max Croce, Atif Mian, Ben Moll, Sylvain Catherine, Peter DeMarzo, Keshav Dogra, Andrej Mijakovic, Indrajit Mitra, Alp Simsek, and seminar participants at Stanford GSB, Columbia GSB (econ and finance), Duke Fuqua Research Triangle seminar, JHU Econ and Finance, MIT Sloan, University of Texas at Dallas, the New York Federal Reserve Bank, the University of Chicago, UC Davis, USC Marshall, New York University, the ECB, and INSEAD. This paper originally circulated under the title: "Financial and Total Wealth Inequality with Declining Interest Rates."

Over the post-war period, US interest rates and wealth inequality have displayed a remarkably strong negative correlation. This time series relationship is displayed in Figure 1, which compares the share of wealth held by the top-10% of the wealth distribution, against the implied price of a 10-year inflation-adjusted zero-coupon bond.[1] Figure 1 shows that the top-10% wealth share declines as the real 10-year bond price falls (yield increases) from the 1960s to the early 1980s, and increases as the real bond price rises (yield falls) from the early 1980s through the end of our sample in 2023. In terms of magnitudes, estimated real yields rose from 1.7% in 1960 to a peak of 7.6% in 1981, before falling to 1.6% at the end of our sample in 2023. At the same time, the top-10% share of wealth fell from 70.5% in 1960 to 62.7% in 1985, before rising to 69.6% by 2023.

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

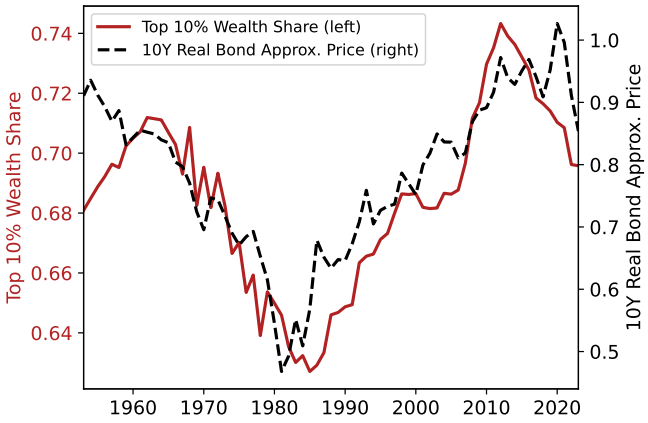

Note: The red solid line displays the top-10% wealth share for the United States from the World Inequality Database. The black dash-dot line displays the 10-year real bond price approximating the real yield using the 10-year nominal yield and contemporaneous inflation. See Appendix B.2 for further details.

Since discount rates determine the valuation of all long-lived assets, a natural hypothesis is that rising asset values due to falling rates have increased measured wealth inequality. However, such a mechanism depends crucially on *heterogeneity* in household portfolio formation. If all households used the same portfolio weights, then a decrease in discount rates would increase all portfolios by the same proportion, leaving top wealth shares unchanged. A fall in rates increases wealth inequality through this revaluation channel only if the wealthy hold a larger share of assets that are more exposed to variation in interest rates.

In this paper, we seek to directly estimate the exposure of household portfolios to interest rate risk, in order to quantify the link between real interest rates and measured wealth inequality over our sample period. In a companion paper, Greenwald et al. (2025), we address the normative

[1] See Appendix B.2 for details on the construction of this real bond series. Wealth is defined as household assets minus liabilities.

question that inevitably follows: what are the implications of falling interest rates for inequality in consumption possibilities and welfare?

Our analysis builds on the concept of *cash flow duration*, which summarizes the sensitivity of an asset's value to changes in long-term interest rates. We generalize this concept to the case of non-permanent shocks to rates, derive adjustments for ad valorem taxes and changes in growth expectations, and provide sufficient statistics to directly estimate durations from asset pricing data. We derive a simple expression that maps moments of the duration distribution and changes in rates into changes in top wealth shares. The key determinant is the gap in duration between the top and the bottom of the wealth distribution. The larger the gap, the larger the increase in the top wealth share when rates decline. Our main empirical exercise seeks to measure the joint distribution of duration and wealth, and to study its implications for the level of wealth inequality.

Characterizing this distribution requires two main inputs. First, we need to measure each household's portfolio allocation across asset classes, which we obtain from the Survey of Consumer Finances (SCF). Second, we need to measure the interest rate exposure (duration) of each asset class. By combining observed portfolio shares with asset class durations, we obtain the duration of household wealth portfolios at each point in time.

We observe substantial heterogeneity in the duration of US household portfolios by wealth level and age. Duration is generally increasing in wealth because the wealthy hold much higher shares of public equities and private business wealth, which have the longest durations among the asset classes we consider. In contrast, the bottom 90% of the wealth distribution is heavily tilted toward housing, vehicles, and deposit-like assets that offer much lower duration. As a consequence, the value-weighted duration of top-10% wealth holders far exceeds that of the bottom 90%, implying that a decline in interest rates raises the top-10% wealth share.

To quantify the strength of this revaluation mechanism, we feed in the observed path of real interest rates to study its implications for the wealth distribution. Using our measures of household-level durations at each point in time, we are able to obtain a new distribution of revalued wealth following each change in interest rates, and hence the implied change in top wealth shares. The revaluations due to changes in the interest rate alone can account for virtually all of the observed increase in the top-10% wealth share from 1983 to 2023, if shocks to the real interest rate are perceived to be permanent. If shocks are perceived to be highly persistent but non-permanent, the mechanism still explains more than half of the increase in wealth inequality.

These results do not depend on the fundamental drivers of the real interest rate decline. However, one driver could be a slowdown in growth, which would have accompanying effects on cash flows. We extend our results to directly account for the effects of shocks to long-run cash-flow

growth on wealth inequality, finding that they do not materially reduce our estimated contributions of interest rates to wealth inequality, and in some cases strengthen them.

To validate our asset duration measures, we check that the implied paths of asset prices are consistent with the data under our estimated exposures. Our duration-based asset prices yield a close match for the evolution of equity values, but overstate the actual rise in house prices. Because the portfolios of the bottom 90% are heavily tilted toward housing, this implies that our model-implied changes in wealth inequality are if anything conservative.

An important concern is that changes in inequality due to revaluations might be short-lived due to offsetting changes in household consumption and savings behavior. To address this, we pair our empirical estimates with a calibrated life-cycle model. The model features a bequest motive and a rich income process calibrated to match data from the Panel Study of Income Dynamics (PSID). We calibrate heterogeneity in the duration of wealth in the model to directly match our empirical estimates by wealth bin and age.

When we feed the observed path of real interest rates into the model, we find that the evolution of top wealth shares is very close to our empirical estimates that simply cumulate the instantaneous revaluation effects, implying that interest rates can have quantitatively important effects on inequality even over multiple decades. This finding is due to two forces: (i) the high degree of persistence of the wealth distribution in workhorse life-cycle models; and (ii) the initial conditions present at the start of our sample in 1983, as the compression of the wealth distribution following the rapid rise in interest rates in the early 1980s returned to steady state. We extend the model to allow for a realistic rise in income inequality over time. The combination of falling rates and rising income inequality implies a moderately larger impact on wealth inequality to date, but much higher long-run inequality if current conditions persist.

Finally, we turn to France and the United Kingdom for international evidence on our mechanism. Using micro data on household portfolios, we find that the duration gap between the top-10% and bottom-90% of the wealth distribution is much smaller in the United Kingdom than in the United States, with France falling in between. This is due to smaller equity and larger real estate portfolio shares throughout the wealth distribution. The UK experienced a decline in long-term real rates similar to the US since the 1980s, while the decline was larger in France. As a result, our methodology correctly predicts that the rise in top wealth inequality was much smaller in the UK than in the US, but larger in France.

To summarize, our results indicate that falling interest rates explain most of the increase in measured wealth inequality since the 1980s due to the higher exposure of wealthy households' portfolios to interest rate risk.

Related Literature. Our paper joins a large body of work on the evolution of wealth inequality. Empirical evidence suggests that wealth inequality has increased in many countries over the past several decades (see, e.g., Piketty and Saez, 2003; Piketty, 2015; Alvaredo et al., 2018). A robust literature has argued for the importance of return heterogeneity in driving wealth inequality (see, e.g., Piketty and Zucman, 2015; Benhabib, Bisin, and Luo, 2017; Cox, 2020; Fagereng et al., 2020; Bach, Calvet, and Sodini, 2020; Stachurski and Toda, 2019; Hubmer, Krusell, and Smith Jr, 2021). We complement these works by arguing that the combination of heterogeneity in interest rate exposure and observed changes in interest rates have been a powerful driver of variation in realized returns, explaining most of the increase in inequality observed over the sample.

Our approach is closely linked to that of Doepke and Schneider (2006), who focus on the distributional consequences of inflation, and Auclert (2019), who focuses on variation in how financial products are indexed to short-term interest rates. Our methodology, which focuses on how changes in exposure to *discount rates* influence inequality, complements these works that measure the cash flow exposure of household portfolios.

Our paper is related to Catherine, Miller, and Sarin (2025), who show that accounting for the revaluation of Social Security benefits—an implicitly held asset with very long duration—can influence measured changes in wealth inequality over time. We take a broader view by showing that valuation effects from the decline in real rates account for most of the observed increase in wealth inequality. Gomez and Gouin-Bonenfant (2024) study the effects of lower interest rates on inequality through their impact on the cost of raising new capital for entrepreneurs. This investment-based channel operates through growth in real assets and cash flows, complementing our duration-based mechanism that operates through the financial revaluation of cash flows.

Kuhn, Schularick, and Steins (2020) use novel microdata to document that heterogeneity in the shares of equity and housing in household portfolios drove much of the rise in inequality since the 1970s. We show that much of this difference can be explained by the large difference in duration between these assets. Extending these insights to cross-country analysis, we show that economies with higher housing wealth shares have smaller differences in duration between the top and bottom of the wealth distribution, and hence smaller increases in wealth inequality when rates fall.

Recent work analyzes the challenges with measuring private business income and wealth (e.g., Kopczuk, 2017; Saez and Zucman, 2016; Piketty, Saez, and Zucman, 2018; Smith et al., 2022; Kopczuk and Zwick, 2020). In our empirical work, we advance this agenda by producing several new measures of private business wealth that combine microdata from the Survey of Consumer Finances with various asset pricing data, and establish robustness of our results to the specifics of

our estimation procedure for the duration of private business wealth.

Last, our paper links to the rich literature on the mechanisms behind the decline in interest rates over our sample (e.g., Bernanke, 2005; Caballero, Farhi, and Gourinchas, 2008; Summers, 2014; Eggertsson, Mehrotra, and Robbins, 2019; Eichengreen, 2015; Gutiérrez and Philippon, 2017; Mian, Straub, and Sufi, 2020). While our estimated effects of interest rate changes do not depend directly on the source of these changes, we extend our analysis to consider interest rate declines driven by a slowdown in growth, as emphasized by Gormsen and Lazarus (2025).[2] We find that our main results on the evolution of inequality are similar after accounting for changes in expected growth rates. Similarly, Heathcote, Storesletten, and Violante (2020) highlight the role of rising income inequality in driving wealth inequality. An extension of our model that captures both time-varying labor income risk and the rising share of earnings that goes to top earners strengthens the rise in top wealth inequality and allows the model to capture the large rise in top-1% wealth shares.

Overview. Section 1 derives the theoretical link between cash flow duration and wealth inequality. Section 2 presents our key empirical facts on household portfolio durations. Section 3 verifies that we match the valuation of equities and real estate over time when we apply our duration-based revaluations to the actual observed path of real interest rates and growth rates. Section 4 contains our main empirical estimates of the effects of real interest rate and growth rate variation on wealth inequality. Section 5 considers the robustness of this analysis. Section 6 uses a calibrated life-cycle model to account for savings dynamics. Section 7 presents international evidence for France and the United Kingdom, furthering the duration mechanism. Section 8 concludes.

1 Duration and Inequality

In this section we define a measure of interest rate sensitivity, which we will refer to for simplicity as duration. We present its basic properties, and derive our key theoretical result linking variation

[2]The decline in growth could be due to demographic change (Summers, 2014; Eggertsson, Mehrotra, and Robbins, 2019; 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 rise in labor income risk which strengthens precautionary savings motives (Hubmer, Krusell, and Smith Jr, 2021), or a global saving glut and/or shortage of safe assets (Bernanke, 2005; Caballero, Farhi, and Gourinchas, 2008), among others. Mian, Straub, and Sufi (2020) argue that the rich have a higher propensity to save than the poor, a hypothesis supported by Norwegian data in Fagereng et al. (2019). As a result, a rise in income inequality, for example due to skill-biased technological change, could reduce aggregate demand, leading to a decline in real rates. Similarly, lower tax progressivity could lead to more saving by the rich, more aggregate wealth, and lower rates (Hubmer, Krusell, and Smith Jr, 2021). However, Heathcote, Storesletten, and Violante (2020) argue that once transfers and actual tax paid are considered, the US tax system has not become less progressive.

in duration across the wealth distribution to the effects of declining interest rates on inequality.

1.1 Duration

The sensitivity of an asset's value to a permanent change in discount rates is summarized by its *cash flow duration*. Formally, this relationship can be written

$$\Delta \log P \simeq -D \times \Delta \log R, \tag{1}$$

where P is the asset price, D is the cash flow duration, and R is the gross discount rate. If all interest rate changes were permanent over our sample, (1) would be sufficient to approximately describe all asset price changes. However, we also consider a more general interest rate process that allows for autoregressive dynamics, which nests the permanent case. We derive sufficient statistic formulas that nest (1), and provide a mapping between the generalized duration statistic and the price-dividend ratio under the special case where cash flows grow at a constant rate. Finally, we extend the model to allow for time-varying cash flow growth rates.

For our implementation, we consider an autoregressive process for an asset's discount rate:

$$\log R_t = \log \bar{R} + z_t, \qquad z_t = \phi z_{t-1} \tag{2}$$

where z_t denotes the time-varying component of the interest rate. We consider a deterministic environment in which cash flows can be interpreted as expected cash flows, and where adjustments for risk occur through the discount rate. Thus, we allow each asset to potentially have its own mean discount rate $\bar{R}$, although we omit an asset-specific index for parsimony. However, we assume all assets share the same value of z_t.

For our derivations, we assume the expected path of interest rates is given by (2), for some initial value z_0. However, at time 0, the interest rate path unexpectedly changes to $\tilde{z}_0 = z_0 + \varepsilon$, so that $\tilde{R}_0 = R_0 \exp(\varepsilon)$, generating a new path of expected interest rates $\tilde{R}_0, \tilde{R}_1, \ldots$ following (2). Throughout the paper we will use tildes (e.g., $\tilde{R}$) to indicate updated values following an interest rate shock, while equivalent variables without tildes indicate pre-shock values. Proofs of all propositions can be found in Appendix A.

Proposition 1. Define

$$P_0 \equiv \sum_{t=1}^{\infty} \left(\prod_{j=0}^{t-1} R_j^{-1} \right) x_t \tag{3}$$

to be the price of an asset with cash flows $\{x_t\}_{t=1}^{\infty}$ at $t = 0$ prior to the shock, and $\tilde{P}_0$ to be the price

at $t = 0$ following the arrival of the shock. Then:

(a) The change in the valuation of the asset is approximately described by

$$\log \tilde{P}_0 - \log P_0 \simeq -\mathcal{D}(\phi) \times \varepsilon \tag{4}$$

where the *generalized duration* for an interest rate shock with persistence ϕ is defined by

$$\mathcal{D}(\phi) = \frac{1}{P_0} \sum_{t=1}^{\infty} \left(\sum_{j=0}^{t-1} \phi^j \right) \left(\prod_{j=0}^{t-1} R_j^{-1} \right) x_t. \tag{5}$$

(b) In the case of permanent shocks ($\phi = 1$) we obtain the traditional duration formula

$$\mathcal{D}(1) = \frac{1}{P_0} \sum_{t=1}^{\infty} t \times \left(\prod_{j=0}^{t-1} R_j^{-1} \right) x_t \tag{6}$$

(c) For a portfolio of assets indexed by k, equation (4) holds using the value-weighted duration

$$\mathcal{D}^{VW}(\phi) \equiv \sum_k \omega(k) \mathcal{D}_k(\phi) \tag{7}$$

where $\omega(k)$ is the share of the portfolio's value in asset k, and $\mathcal{D}_k(\phi)$ is the generalized duration of asset k.

Proposition 1 generalizes the typical duration formula to describe the effect of interest rate changes on asset prices when interest rate shocks may not be permanent. In the traditional case with permanent interest rate changes, duration is equal to the weighted average of the time to receive an asset's cash flows, weighted by the share of that asset's value derived from cash flows at that date. Proposition 1 shows that the more general autoregressive case nests this case, using a weighted average of terms of the form $\sum_{j=0}^{t-1} \phi^j$ instead. These terms, and hence the generalized duration of the asset, become smaller when the interest rate change dissipates more quickly (smaller ϕ).

We next turn to the important special case where an asset's cash flows grow at a constant rate.

Proposition 2. Assume that an asset's cash flows are expected to grow at a constant rate G, and that the discount rate begins in steady state with $z_0 = 0$ and $R_t = R, \forall t$. Then (5) takes the form:

$$\mathcal{D}(\phi) = \frac{(P_0/x_0) + 1}{(1-\phi)(P_0/x_0) + 1}. \tag{8}$$

Under permanent shocks ($\phi = 1$), equation (8) simplifies to the price-dividend ratio plus one:

$$\mathcal{D}(1) = \frac{1+r}{r-g} = \frac{P_0}{x_0} + 1. \tag{9}$$

Proposition 2 shows that we obtain a simple formula for the generalized duration in this setting. More important, this formula can be computed directly using only the current price-dividend ratio of the asset (P_0/x_0) given a value for ϕ. This provides us with a simple method for computing generalized durations from observed asset prices in our empirical exercise below.

Next, we extend this duration concept to deal with property taxes, insurance costs, and maintenance costs, all of which are relevant for real estate. We model these as "ad valorem" taxes that are levied on the price of the property.

Proposition 3. Maintain the assumptions of Proposition 2, and further assume that while the asset has gross cash flow x_t and price P_t, it incurs a tax or cost of τP_t in each period. Then

$$\log \tilde{P}_0 - \log P_0 \simeq -\mathcal{D}(\phi, \tau) \times \varepsilon \tag{10}$$

$$\mathcal{D}(\phi, \tau) \equiv \frac{(1-\tau)(P_0/x_0) + 1}{(1-\phi)(1-\tau)(P_0/x_0) + 1}. \tag{11}$$

Equation (11) nests (8) for $\tau = 0$, but contains two differences when $\tau > 0$. First, it introduces $(1-\tau)$ terms multiplying the price-to-cash-flow ratio, a change that is quantitatively small for typical values of τ. Second and more important, it implies that we should use the *gross* cash flow x_0 in the denominator of the price-to-cash-flow ratio without deducting the ad valorem costs τP_0. This leads to a smaller price-to-cash-flow ratio, and hence a lower generalized duration.[3]

1.2 Duration Heterogeneity and the Wealth Distribution

Having defined our measures of interest rate exposure, we now present our main theoretical proposition on the impact of interest rates on inequality. Define the wealth of household i, W^i, as the sum of all household assets minus its liabilities.

Proposition 4. For a small negative change in rates ($\varepsilon < 0$):

(a) The growth in household wealth due to revaluation, $\tilde{W}^i/W^i$, has a positive cross-sectional covariance with wealth W^i if and only if the generalized duration of wealth, $\mathcal{D}^i$, has a posi-

[3] Intuitively, if the price of housing rises due to falling interest rates, so do the required property tax payments, decreasing the asset's cash flows, and dampening the overall rise in price. This adjustment is economically significant in our analysis, and is important for allowing the model to match the observed path of price-rent ratios in residential real estate in Section 3.3.

tive covariance with the level of wealth W^i. A necessary and sufficient condition for this positive covariance is that aggregate (value-weighted) wealth duration exceeds average (equal-weighted) generalized duration in the pre-shock economy.

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

$$d\tilde{S}^\alpha \simeq -S^\alpha(1-S^\alpha)(\mathcal{D}^{top}-\mathcal{D}^{bottom})\varepsilon. \tag{12}$$

Proposition 4 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 wealth portfolios of the rich grow by more than those of the poor. This is in turn determined by the relative exposures of these portfolios to interest rates, as summarized by generalized duration. The proposition also allows us to use simple summary statistics to approximate the effects of a decline in interest rates on inequality. To do so, however, we need to measure the distribution of generalized duration in the data, motivating our empirical exercise in the next section.

1.3 Shocks to Growth

For our final derivations, we extend our analysis to the case where the change in interest rates is partially caused by a change in the economy's long-term growth rate, as discussed in Gormsen and Lazarus (2025). We update the environment of Proposition 1 to assume that aggregate economic output $\{y_t\}$ grows at rate $\{G_t\}$, so that $y_{t+1} = G_t y_t$. We define $g_t = \log G_t$ and assume that the time-varying component of growth $z_{g,t}$ follows an autoregressive process:

$$g_t = \bar{g} + z_{g,t}, \qquad z_{g,t} = \phi z_{g,t-1}. \tag{13}$$

where $\bar{g}$ is a constant, and ϕ is the same persistence parameter governing the discount rate process (2).[4] We further specify that log discount rates depend on log growth rates with loading θ, so that

$$\log R_t = \log \bar{R} + z_t + \theta z_{g,t}. \tag{14}$$

Proposition 5. Assume that the economy receives unexpected shocks so that $\tilde{z}_0 = z_0 + \varepsilon$, and $\tilde{z}_{g,0} = z_{g,0} + \varepsilon_g$, so that $\tilde{G}_0 = G_0 \exp(\varepsilon_g)$ and $\tilde{R}_0 = R_0 \exp(\varepsilon + \theta\varepsilon_g)$. Then:

(a) Consider an asset whose cash flows will grow at rate G_t at time t (i.e., $x_{t+1} = G_t x_t$, $\forall t$). The change in the value of this asset is approximately

$$\log \tilde{P}_0 - \log P_0 \simeq \mathcal{D}(\phi) \times \varepsilon_g - \mathcal{D}(\phi) \times (\log \tilde{R}_0 - \log R_0) \tag{15}$$

(b) Consider a portfolio of assets, each of which either has cash flows that grow at rate $\{G_t\}$ or has cash flows that do not depend on $\{G_t\}$. This portfolio's value is approximated by

$$\log \tilde{V}_0 - \log V_0 \simeq \mathcal{D}_G^{VW} \times \varepsilon_g - \mathcal{D}^{VW} \times (\log \tilde{R}_0 - \log R_0). \tag{16}$$

where $\mathcal{D}_G^{VW}$ is defined as

$$\mathcal{D}_G^{VW}(\phi) \equiv \sum_k \omega(k) \mathcal{D}_k(\phi) \mathbf{1}_k, \tag{17}$$

with portfolio shares $\omega(k)$ and asset durations $\mathcal{D}_k(\phi)$ defined as in (5), and where

$$\mathbf{1}_k \equiv \begin{cases} 1 & \text{if asset } k\text{'s cash flows grow at rate } \{G_t\} \\ 0 & \text{if asset } k\text{'s cash flows do not depend on } \{G_t\}. \end{cases} \tag{18}$$

(c) The change in the top-α wealth share is approximately

$$dS^\alpha \simeq -S^\alpha(1 - S^\alpha)\Big[(\mathcal{D}^{top} - \mathcal{D}^{bottom})(\log \tilde{R}_0 - \log R_0) - (\mathcal{D}_G^{top} - \mathcal{D}_G^{bottom})\varepsilon_g\Big]. \tag{19}$$

Part (a) of Proposition 5 contains several important results. First, the loadings on these changes are simply $\mathcal{D}(\phi)$ and $-\mathcal{D}(\phi)$, respectively, and hence can be directly computed using only the generalized duration formula. Second, we can summarize the change in the asset's value using

[4] The case where these processes have differing persistence is straightforward to derive, but substantially more complicated to express since $\log R_t$ goes from an AR(1) process to a mixture of such processes, which is not AR(1). Greenwald, Lettau, and Ludvigson (2025) estimate a time series model of the primitive components driving stock prices, and find that the persistence of the low-frequency components of the growth rate and risk-free discount rates are very close, justifying this assumption.

only the change in the growth rate and the change in the discount rate, without needing to isolate the "pure discounting" term z_t. Importantly, this means that the implied change in the value of the asset does not directly depend on the parameter θ, which governs the response of rates to growth shocks. This result bolsters the robustness of our method, since θ typically depends on the elasticity of intertemporal substitution, the value of which is the subject of disagreement in the literature. While θ is very important for decomposing movements in the interest rate into a component driven by growth and a pure discounting term, as in Gormsen and Lazarus (2025), it does not matter for the overall effect on valuations conditional on the change in the interest rate.[5]

Parts (b) and (c) of Proposition 5 extend Proposition 4 to allow for shocks to the growth rate. A decline in growth rates may mitigate the effect of a decline in real rates on wealth inequality if the wealthiest have more exposure to growth shocks in their portfolio.

Combined, these results make it straightforward to extend our analysis to account for changes in valuations driven both by changes in the expected growth rate of the economy as well as changes in the pure rate of time preference.

2 Measuring Household Portfolio Duration

While declines in real interest rates raise asset prices, this only affects wealth *inequality* if households' exposures to real rates vary systematically with wealth. In this section, we quantify portfolio exposures at the household level and document how they differ across the wealth distribution. As shown in Proposition 5 and equation (7) of Section 1, measuring household portfolio exposures requires two sets of estimates: (i) generalized durations for each asset class, and (ii) household portfolio shares. We compute these objects to obtain household-level exposures, then feed in the observed path of real interest rates to measure the implied effects on inequality.

2.1 Measuring Asset Durations

To begin, we compute the durations $\mathcal{D}_{k,t}(\phi)$ of each asset class at each point in time. We report results for $\phi = 1$ (permanent shocks) and for $\phi = 0.99$. For most risky assets, we compute the duration at each point in time using only its contemporaneous price-dividend ratio and the parameter ϕ, using (8). With this general approach in mind, we briefly describe our duration measurement for each asset class, with full details available in Appendix D. The resulting durations, averaged over the 1983–2023 sample period, are reported in the first two columns of Table 1.

[5] For intuition, the total change in the interest rate governs the effect on discount rates, while the size of the growth shock (but not the value of θ) governs the effect on cash flows.

Equities. We measure the duration of equities using Shiller's cyclically-adjusted price-earnings (CAPE) ratio. To convert the price-earnings ratio into a price-dividend ratio, we assume a constant payout ratio of 0.496, corresponding to the average dividend-earnings ratio observed over our 1983–2023 sample, which yields an average duration of 50.5 years in the permanent case. Appendix E.3 shows that our results are robust to alternative sample windows for the payout ratio and to rescaling dividends by total payouts, including net equity repurchases. Using scaled earnings rather than dividends means that our duration estimates are not sensitive to dividend timing. Moreover, our results are highly similar using alternative measures such as the CRSP value-weighted price-dividend ratio.[6]

Real Estate. As discussed in Proposition 3, we take into account that the costs associated with real estate take an ad valorem form. We measure the duration of real estate using (11) using rent as the asset's gross cash flows. The price-rent ratio is constructed using Zillow data from 2015 onward.[7] We use $\tau = 3\%$ to capture property taxes, insurance, and maintenance costs, in line with estimates from Damen, Korevaar, and Van Nieuwerburgh (2025).[8] This yields an average duration of 13.8 years for real estate in the permanent case.[9]

Private Business Wealth. Measuring the duration of private business wealth (PBW) is challenging, since we cannot rely on market valuations. This asset class contains both fast-growing companies with a long duration and more stagnant companies whose existence is tied to the human

[6]We consider a wide set of alternative measures for the duration of equity, detailed in Appendix D.1. The first one uses the price-dividend ratio of the CRSP value-weighted stock market index, yielding a duration of 53.8 years. The second one uses the price-dividend ratio computed from the JST macro-database (Jordà, Schularick, and Taylor, 2017; Jordà et al., 2019), yielding a duration of 49.1 years. The third one uses the S&P 500 price-dividend ratio, yielding a duration of 50.3 years. The fourth one uses the valuation for the entire non-financial corporate sector from the Financial Accounts of the United States (FAUS), yielding a duration of 35.7 years. The FAUS includes privately-held companies and imputes a valuation to private companies, which is adjusted downward by 25% to account for lower liquidity. Thus, we obtain similar duration numbers for all the measures except for the FAUS measure, which is not a pure equity duration measure due to the liquidity adjustment.

[7]Specifically, we employ the Zillow Home Value Index (ZHVI) to measure house prices and the Zillow Observed Rent Index (ZORI) to capture rental rates. For the pre-2015 period, we extend the series backward by scaling the 2015 price-rent ratio according to the proportional changes in the Federal Housing Finance Agency (FHFA) house price index and the Bureau of Labor Statistics' CPI shelter index.

[8]Damen, Korevaar, and Van Nieuwerburgh (2025) calculate the value of these costs for Belgium and the Netherlands, finding total costs for the median rental between 1.6% and 2.3%. We adjust these estimates to account for cross-country differences in property taxation: property taxes amount to approximately 0.3% in Belgium and the Netherlands, compared to an average of 1% in the United States. After this adjustment, we obtain an implied value of τ of approximately 3% to capture property taxes, insurance, and maintenance.

[9]In Appendix D.2 we consider two alternative estimates. The first uses the price-to-rent ratio from Jordà et al. (2019), yielding a similar duration of 13.3. As detailed in the Appendix, the JST database reports the net rental yield. We use our estimate of $\tau = 3\%$ to recover the gross rental yield and then apply our duration formula. The second is the duration measure that matches the increase in the real estate price-to-rent ratio over the sample. Specifically, using (20) and the observed change in real yields, we back out the duration that exactly matches the change in the log price-to-dividend ratio from 1983 to 2023. This measure yields a duration of 6.2 years.

capital of the owner (Smith et al., 2019). We use detailed data from the SCF on the legal structure of private businesses to divide PBW into those with a corporate structure (e.g., S-corporations and all other types of corporations) and those with a non-corporate structure (e.g., partnerships and sole proprietorships). The corporate (non-corporate) segment proxies for the fast-growing (slow-growing) types. Appendix C.1.1 discusses the details of the split. [10]

For non-corporate PBW, we measure duration using (8), obtaining price-payout ratios from the Non-Corporate Business sector in the FAUS.[11] This yields an average duration of 25.7 years.

Corporate PBW includes fast-growing start-ups, which often have cash flows growing at a much faster rate early in their corporate life cycle than later when they reach a mature state of growth. To capture this, we use a two-stage Gordon Growth Model with an initial 20-year phase of higher cash flow growth, followed by a second stage with a lower growth rate.

We use data on publicly listed stocks from CRSP to fit this model. We assume that these private businesses begin similar to small public companies, then transition to behave like average public companies. Correspondingly, we obtain the initial price-dividend ratio from the smallest decile of public companies. We obtain cash flow growth rates for the initial high-growth stage and second lower-growth stage using 10-year real realized dividend growth for the smallest decile and aggregate (value-weighted) portfolios, respectively. This procedure yields an average duration for Corporate PBW of 58.3 years for the permanent case, slightly higher than our measure for public equities due to the initial high-growth phase. This measure is our most conservative (lowest duration) compared to the four alternative methods we consider, as detailed in Appendix D.4.[12]

Persistent Shocks. These duration numbers assume that the rate shocks are permanent. If the interest rate shocks are perceived to be persistent ($\phi = 0.99$) rather than permanent, then the duration of long-lived assets decreases, as transitory changes in discount rates have smaller effects on asset prices. The duration of equities declines from 50.5 to 33.0 years. The duration of real estate

[10] The two groups correspond to corporate businesses and non-corporate businesses in the FAUS.

[11] The corresponding cash-flow series includes both remuneration for labor and capital. To isolate the capital remuneration, we follow the literature and assign 30% of the business income to capital (Alvaredo et al., 2020; Garbinti, Goupille-Lebret, and Piketty, 2021; Martínez-Toledano, 2020). The appendix does robustness with respect to a smaller capital income share of 14%, used by Quadrini and Rıos-Rull (1997) and Krueger and Perri (2006), and a higher share of 50% used by the PSID. We consider a third alternative measurement based on SCF data as detailed in Appendix D.5. Table D.3 in Appendix D.3 compares the results and shows that our baseline duration is at the lower end of the range of alternative estimates.

[12] The first two alternatives use the same conceptual approach for Corporate PBW as the benchmark measure, but use an alternative to small stocks for the first, high-growth stage in the GGM. The first alternative uses high-growth private businesses that receive venture capital funding using data from Pitchbook, yielding a duration of 60.1 years. The second alternative uses stocks that just went through an initial public offering (IPO) using data from Ritter (2022), yielding a duration of 58.5 years. The third alternative uses small stocks for the measurement of long-duration PBW, but uses a single-stage GGM, which generates a duration of 100.3 years. A fourth alternative measure, based on SCF data as detailed in Appendix D.5, generates a Corporate PBW duration of 65.9 years. All four alternative numbers refer to the permanent case of $\phi = 1$.

declines from 13.8 to 12.2 years. Similarly, the duration for Corporate PBW declines from 58.3 to 37.3 years, while the duration for Non-Corporate PBW declines from 25.7 years to 20.5 years.

Remaining Asset Classes. Computing the durations of the remaining asset classes is straightforward, as they have simpler or fixed cash-flow structures. For vehicles, we compute an asset duration based on a depreciation rate from the Fixed Asset Tables of the Bureau of Economic Analysis and the average age of vehicles in use from the Bureau of Transportation Statistics (see Appendix D.6 for details). For fixed income assets, we use the effective duration of the ICE government and corporate bond index. This series starts in 1996, prior to which we use the duration of the Treasury portfolio to impute the duration of Treasuries and corporate bonds (see Appendix D.7 for details). For cash and deposits, we assume a constant duration of 0.25 years.

Liabilities. Last, we compute the durations of household liabilities. While 30-year fixed-rate mortgages are dominant in the US, the average outstanding mortgage has duration far below 30 years due to the time elapsed since origination, amortization, coupon payments, and prepayment. We obtain the duration of mortgage debt from a representative portfolio of all outstanding US pass-through mortgage-backed securities included in the Bloomberg-Barclays Aggregate MBS Index (see Appendix D.8 for details). These durations are estimated by the data provider accounting for prepayment, and result in a duration just below 5 years. For vehicle debt, we observe the periodic loan payment, the remaining number of periods, and the loan rate in the SCF, allowing us to directly compute the duration of each loan (see Appendix D.9 for details). Similarly, for student debt, we compute duration using SCF information on residual maturity, scheduled cash flows, and interest rates (see Appendix D.10). For other debt, we assume a constant duration of 1 year as a compromise between its two main subtypes: revolving debt and amortizing 24-month personal loans.

2.2 Household Heterogeneity in Duration

With our measures of generalized duration for each asset class in hand, we now measure household portfolio shares, compute each household's portfolio exposure, and analyze how this exposure varies with household characteristics.

Measuring Portfolio Shares. For data on portfolio shares, we use the Survey of Consumer Finances (SCF), a detailed survey of household wealth conducted every three years from 1983 until

Table 1: Duration of the Household Wealth Portfolio

	Duration		Portfolio Share (%)				
Asset Class	$\phi = 1$	$\phi = 0.99$	All	Bottom 90	P90–99	Top 10	Top 1
Equities	50.5	33.0	21.1	15.0	24.2	23.6	23.0
Real Estate	13.8	12.2	46.8	83.7	40.6	31.9	22.4
Corporate PBW	58.3	37.3	9.4	1.3	6.9	12.7	19.1
Non-Corporate PBW	25.7	20.5	12.0	4.6	10.8	14.9	19.4
Vehicles	3.5	3.4	3.9	10.3	2.0	1.4	0.6
Fixed Income	5.6	5.4	15.1	15.3	16.5	15.1	13.5
Cash and Deposits	0.3	0.3	7.7	10.8	7.8	6.5	4.9
Mortgage Debt	4.8	4.7	12.9	32.3	7.8	5.2	2.3
Vehicle Debt	1.4	1.4	1.0	3.1	0.3	0.2	0.1
Student Debt	2.4	2.4	0.7	2.4	0.1	0.0	0.0
Other Debt	1.0	1.0	1.3	3.0	0.6	0.6	0.7
Equal-weighted Duration ($\phi = 0.99$)			14.4	13.9	18.4	18.7	21.7
Equal-weighted Duration ($\phi = 1$)			18.3	17.5	25.5	26.0	31.0
Value-weighted Duration ($\phi = 0.99$)			19.5	16.4	19.0	20.7	22.6
Value-weighted Duration ($\phi = 1$)			**27.1**	**21.2**	**26.6**	**29.4**	**32.5**

Note: The columns "Duration" report the generalized duration of each asset $\mathcal{D}_k$, averaged across all years from 1983 to 2023. We report the durations estimated under $\phi = 1$ (permanent case) and $\phi = 0.99$. The columns "Portfolio Shares" report the value-weighted (aggregate) portfolio weights from the Survey of Consumer Finances (SCF) and SCF+, averaged across all survey waves from 1983 through 2022. We report portfolio shares for the aggregate household sector (All), for households in the bottom 90 percentiles (Bottom 90), in the top-10 percentiles, in the 90th to 99th percentile (P90-P99), and the top-1 percentile (Top 1). Portfolio shares for each group are computed as value-weighted averages, using the product of wealth and the SCF sampling weights as weights. The bottom panel reports the equally weighted (EW) duration and the aggregate or value-weighted (VW) duration. Durations are first computed for each survey wave and then averaged across all survey waves. Liabilities receive negative portfolio weights.

2022.[13] These data contain household holdings by asset class, enabling us to measure $\omega_t^i(k)$ for each asset class k and household i in each SCF wave year t. The resulting portfolio shares, averaged across all survey waves, can be seen in Table 1. The five rightmost columns display wealth-weighted average portfolio shares for the full sample, bottom 90%, 90%-99%, top 10%, and top 1% of the wealth distribution. The bottom 90%, 90%-99%, and top 1% groups hold roughly equal proportions of wealth. The table shows substantial heterogeneity in portfolio composition across the wealth distribution. In particular, the share of wealth in real estate is sharply decreasing with wealth, as are the shares of vehicles, and cash and deposits. In contrast, the shares held in equities and private business wealth, particularly the corporate type, are strongly increasing with wealth.

[13]Since the SCF begins in 1989, after interest rates have already been falling for some time, we supplement our data with the 1983 wave of the SCF+, constructed by Kuhn, Schularick, and Steins (2020). See Appendix C for details and additional graphs. There is no SCF+ in 1986.

Household Duration. Following equation (7), we combine our time-varying asset class durations with our measured portfolio shares to obtain a wealth duration for each household in each wave of the SCF. The bottom panel of Table 1 displays summary statistics from this distribution, averaged over all SCF waves. We compute an aggregate (value-weighted) duration of 27.1 in the permanent case (19.5 for $\phi = 0.99$), which is substantially higher than the average (equal-weighted) duration of 18.3 (14.4 for $\phi = 0.99$). Similarly, value-weighted average durations at the top of the wealth distribution exceed those lower in the wealth distribution. For example, the value-weighted average duration for the top-10% is 29.4 vs. 21.2 for the bottom-90% for $\phi = 1$ (20.7 vs. 16.4 for $\phi = 0.99$). These statistics reflect a positive association between wealth and duration, as wealthier households hold larger shares of long-duration assets such as public equity and corporate private business wealth and smaller shares of short-duration assets such as vehicles and cash.

Given that value-weighted duration exceeds equal-weighted duration, Proposition 4 implies that a decrease in real rates will increase wealth inequality. Similarly, the sizeable gap between the duration of the top decile and other US households implies a large impact of declining real rates on the wealth share of the top decile.

Duration by Wealth. Figure 2 shows how our measure of household-level duration varies in the cross-section of households when we sort households by wealth into groups, using increasingly finer wealth percentile groups higher up in the wealth distribution. The figure shows that duration is generally increasing in wealth. The rise in duration at the top of the wealth distribution is driven by increasing portfolio shares in equities and PBW. There is a non-monotonicity around the 30th percentile of wealth, driven by households transitioning from renting to homeownership. New homeowners tend to have large amounts of housing wealth and a large mortgage, relative to the rest of their assets. Since housing assets have higher duration than mortgage debt, this results in high portfolio durations for new homeowners.[14] The relative patterns observed when we use $\phi = 0.99$ (shown in Figure D.4) are similar, but with all generalized durations shifted downward.

Duration by Age. A second characteristic that drives variation in duration is age. Figure 3a documents the raw age profile. We split households into 10-year age groups (20–29, 30–39, ...) and compute, for each SCF wave, the wealth-weighted average duration within each group. We then average these age-group statistics across all survey years. On average, duration rises from age

[14] At this part of the wealth distribution, we also observe the largest gap between $\mathcal{D}^i(\phi)$ and $\mathcal{D}^i_G(\phi)$. This is because the $\mathcal{D}^i_G(\phi)$ measure retains the duration from the housing asset but zeros out the negative duration from the mortgage liability.

Figure 2: Duration by Wealth Percentiles

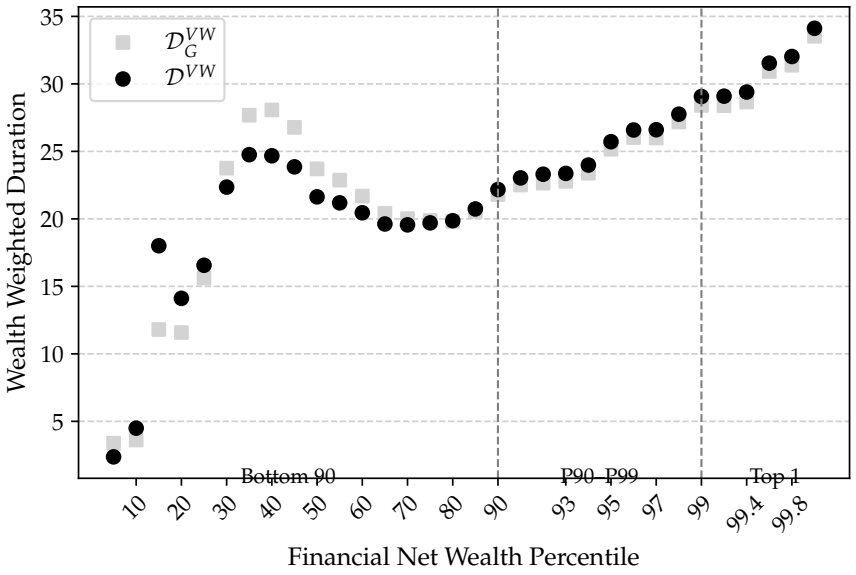

Note: The figure displays average duration by wealth bin in the data. The x-axis is measured in percentiles, with each tick representing the right edge of the bin, so that e.g., "10" corresponds to households with wealth percentile in the interval $[5, 10]$. We plot 1-percent bins between the 90th and 99th percentiles and 0.2 percent bins above the 99th percentile. The data are averages over all SCF waves between 1983 and 2022. The black dots display $\mathcal{D}^i(\phi)$ from (7) while the gray dots display $\mathcal{D}_G^i(\phi)$ from (17), and $\phi = 1$.

30 to age 40, remains relatively flat through age 60, and declines thereafter. Because age and wealth are correlated in the cross-section, part of this pattern reflects the wealth gradient rather than the pure effect of age. Figure 3b shows the age pattern, controlling for wealth.[15] We find durations that decline with age after age 30, as households gradually de-lever their housing position by paying off their mortgage and reduce their share of equities later in life.[16]

To summarize, our main empirical finding is that the duration of households' wealth portfolios is increasing in wealth. The theoretical results in Section 1 imply that revaluation effects from a decline in interest rates since the 1980s should have increased wealth inequality. We quantify the strength of this effect below.

3 Marking Assets to Market

Having measured portfolio durations, we next measure changes in interest rates and growth rate expectations. Combined with our previous estimates, these will allow us to apply Propositions 4

[15] Details on our computation controlling for wealth are available in Appendix D.12. The appendix also reports the results for $\phi = 0.99$.

[16] Appendix Table D.5 provides a further exploration of empirical determinants of household portfolio duration.

Figure 3: Duration by Age

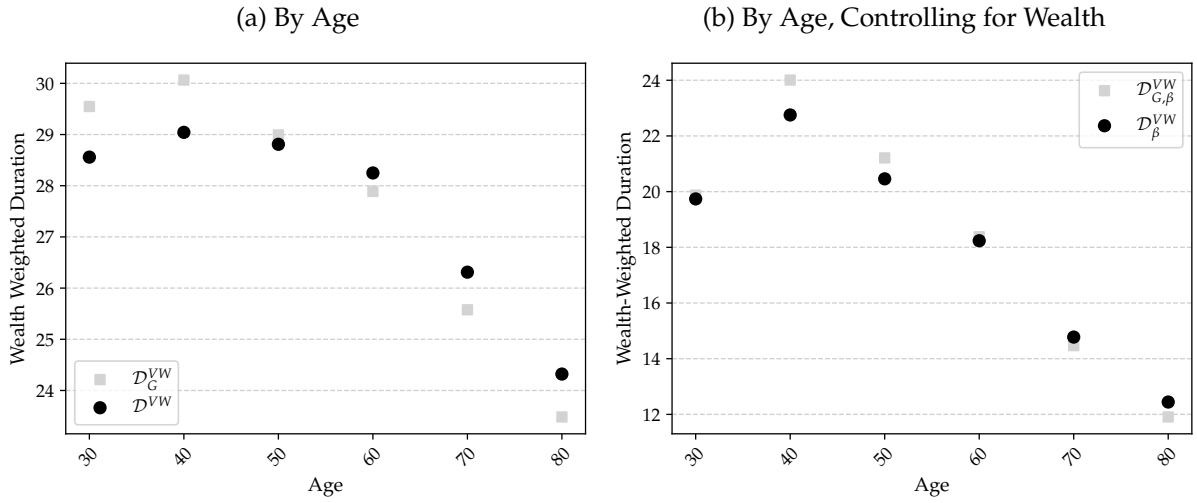

Note: The left panel plots wealth-weighted portfolio duration across 10-year age groups, where each point corresponds to the upper bound of the respective age bin. The right panel reports the analogous age profile after netting out cross-sectional differences in household wealth: for each SCF wave, we regress duration on age-group dummies and 5-percent net-wealth bins, recover the age-group coefficients, re-center them using the wave-specific weighted mean, and then average the resulting fitted values across survey years. The sample covers all SCF surveys from 1983 to 2022. In each panel, the black markers display $\mathcal{D}^i(\phi)$ as defined in (7), while the gray markers report $\mathcal{D}_G^i(\phi)$ from (17).

and 5 to gauge the impact of the decline in real rates on wealth inequality.

3.1 Real Interest Rates

Our empirical strategy recovers the dynamics of the short-term real rate from the evolution of the 10-year real yield. As we show below, this procedure produces an implied short-rate series that is smoother than the observed one-year real rate, and thus provides a more reliable measure of the underlying long-run trend in real rates—the primary object of interest in this paper. To do so, we rely on the generalized expectations hypothesis.

We use 1-year and 10-year real interest rates computed by the Center for Inflation Research at the Cleveland Federal Reserve Bank. According to the generalized expectations hypothesis, the yield on a 10-year zero-coupon bond equals the average of expected future short-term rates plus a constant term premium η. Given the assumed one-year real interest rate process in (14), the implied 10-year yield satisfies:

$$\log R_t^{(10)} = \begin{cases} \log \bar{R} + \eta + \left(\dfrac{1-\phi^{10}}{10(1-\phi)} \right) (z_t + \theta z_{g,t}) & \text{for } \phi < 1 \\ \log \bar{R} + \eta + (z_t + \theta z_{g,t}) & \text{for } \phi = 1. \end{cases}$$

We set $\log \bar{R} = 1.26\%$, the observed sample average of the 1-year real rate. We set $\eta = 1.03\%$ to match the sample average of the 10-year real rate of 2.29% (1.26% + 1.03%). We use the observed $\log R_t^{(10)}$ to recover the unobserved component $z_t + \theta z_{g,t}$. This inferred component is then used to compute the short-term real rate, $\log R_t = \log \bar{R} + z_t + \theta z_{g,t}$ as in (14).

Figure 4a displays the observed one-year and ten-year real interest rates together with the implied one-year rate extracted from the ten-year yield dynamics for $\phi = 1$ (in gray) and $\phi = 0.99$ (in green). The observed 1- and 10-year real rates exhibit a pronounced, similarly-sized decline over most of the sample period. However, the 1-year rate displays considerably greater short-term volatility. By using the 10-year rate to infer the 1-year rate, our methodology filters out high-frequency noise and isolates the persistent component of real rates. The implied real rate dynamics are nearly identical for $\phi = 1$ and $\phi = 0.99$. Using a value of $\phi = 1$ ($\phi = 0.99$) implies a change in the 1-year real rate of -4.08pp (-4.26pp) over the 1983–2023 sample period.[17]

3.2 Long-term Growth Rate Expectations

As discussed in Section 1, long-term growth expectations (g) influence asset prices through both expected cash flows and discount rates per (14). Because g is not directly observable, we employ two alternative proxies in the analysis. Our baseline measure is the average realized GDP growth rate over the preceding six years.[18] The underlying assumption is that investors extrapolate future growth from recently realized outcomes.[19] Figure 4b shows the resulting measure of long-term GDP growth expectations. In this measure, long-term GDP growth expectations change only modestly from 2.30% at the beginning of the sample to 2.54% at the end, a difference of 0.23pp. As an alternative, we proxy g using the 10-year real GDP growth forecast from the Survey of Professional Forecasters (SPF). Because the SPF data begin only in 1992, they cover only a limited portion of our sample. To construct a consistent historical series, we extend SPF growth expectations backward using the model-based extrapolation of Greenwald, Lettau, and Ludvigson (2025).

Depending on the value of ϕ, we recover two corresponding measures of g, plotted in Figure 4b in gray ($\phi = 1$) and green ($\phi = 0.99$). The figure shows these two series are nearly indistinguishable. The full-sample change in this growth rate measure is about 1pp. Additional details

[17] An earlier version of this paper used a VAR for nominal one-year and ten-year interest rates, nominal GDP growth, and inflation along with an affine stochastic discount factor model to infer the ten- and one-year real yields. This alternative approach resulted in a very similar decline in real interest rates.

[18] Since the SCF is conducted at three-year intervals, this window corresponds roughly to the average growth realized over the previous two survey periods (and exactly matches the interval at the beginning of the sample, when the survey was conducted every six years).

[19] Adaptive expectations approximate rational expectations when agents face stable and persistent environments (Muth, 1961), limited information (Sims, 2003), or costly forecasting, such that past observations contain most relevant information for predicting the future.

Figure 4: Evolution of Interest Rates and Growth Rate Expectations

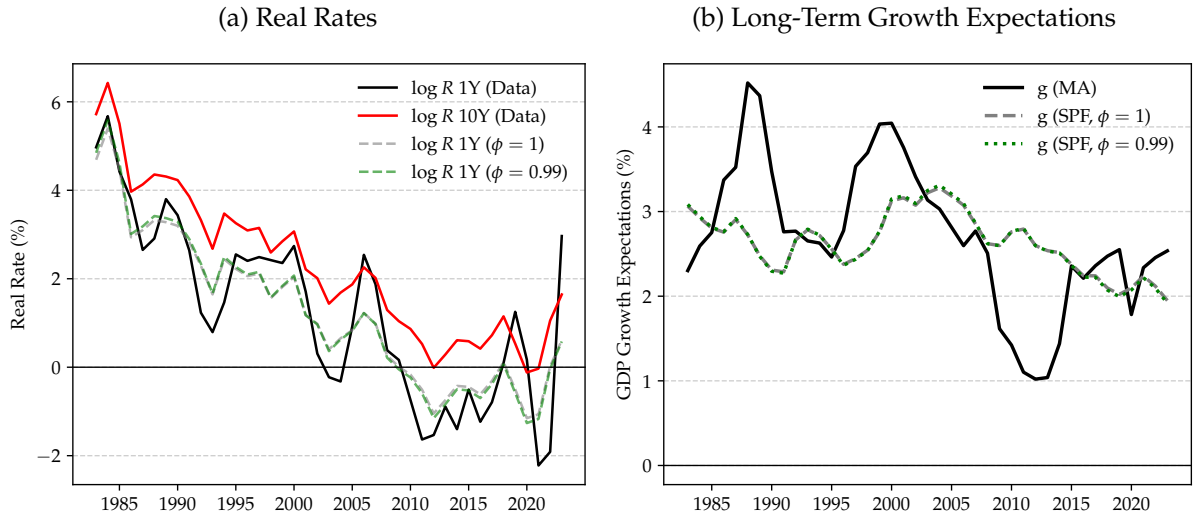

Note: The left panel plots the observed one- and ten-year real yields from the Center for Inflation Research at the Cleveland Fed. It also plots the one-year real yield series implied by the 10-year real yield under the generalized expectations hypothesis for $\phi = 1$ and $\phi = 0.99$. The right panel plots the GDP growth expectations. In black, we plot the growth expectations using the moving average realized growth over the past six years. In gray, long-term real GDP growth expectations from the Survey of Professional Forecasters from 1992 and from Greenwald, Lettau, and Ludvigson (2025) prior to 1992.

are provided in Appendix E.1.

3.3 Validation: Implied Valuations in Equity and Real Estate

To check whether our duration mechanism is consistent with secular changes in valuations, we feed in the observed decline in real interest rates and long-term growth expectations. By using the asset durations we can compute the implied change in the price-to-cash-flow ratio for both equity and real estate using the approximation, recalling (15):

$$\log \tilde{P}_t - \log \tilde{P}_{t-1} \simeq -\mathcal{D}(\phi)\,(\log R_t - \log R_{t-1}) + \mathcal{D}(\phi)(g_t - g_{t-1}). \tag{20}$$

This change in the log price, which is also the change in the log price-dividend ratio given that we hold cash flows fixed, consists of two components: the change in the real interest rate and the change in the long-term growth rate.

Cumulating the left-hand side of equation (20) yields a *duration-implied* measure of the log price-dividend ratio that captures fluctuations due to changes in real interest rates and growth rates. Figure 5 plots the duration-implied measure of the price-dividend ratio alongside the observed series (solid black line). The green dashed line reports the duration-implied measure for

$\phi = 1$, while the red dashed line corresponds to $\phi = 0.99$. The left panel displays the results for the stock market, the right panel for the housing market. Since equation (20) provides the implied changes in the log price-dividend ratio but not its level, we normalize the duration-implied log price-dividend ratio by adding a constant so that its average matches the observed average.

Figure 5a shows the observed and implied price-dividend ratio for equities. The observed log price-dividend ratio rose from 3.0 in 1983 to 4.4 by the end of the sample.[20] The implied price-dividend ratio obtained under $\phi = 1$ is highly correlated with the observed data (0.72) and replicates almost exactly the cumulative increase of 1.4 units over the sample period. However, the $\phi = 1$ series exhibits greater short-term variability. When we instead set $\phi = 0.99$, the implied price-dividend ratio remains strongly correlated with the observed series (0.74) but shows a smaller cumulative increase of 1.1 units and smoother dynamics. Overall, the $\phi = 1$ specification provides a near-perfect match to the long-term increase in the price-dividend ratio—the main focus of this paper—while the $\phi = 0.99$ specification provides a closer fit to short-run fluctuations.

Figure 5: Duration-Implied Stock and Housing Prices

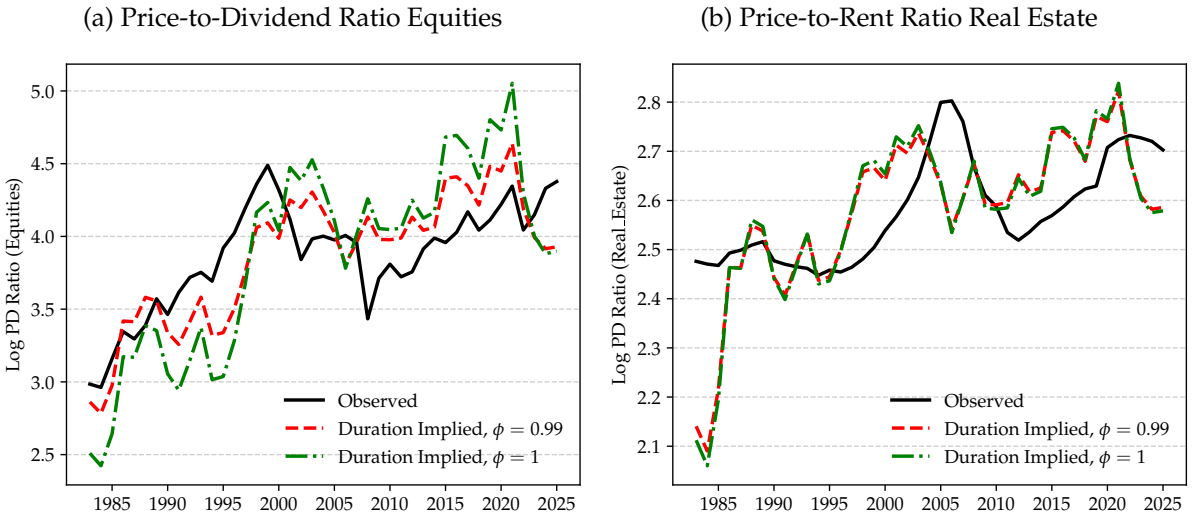

Note: The figure plots the observed price-dividend (or price-rent) ratio together with the duration-implied series computed from Equation 20 for $\phi = 1$ and $\phi = 0.99$. Panel (a) shows the price-dividend ratio for equities, while Panel (b) displays the price-rent ratio for housing.

Figure 5b presents the corresponding results for housing. The observed log price-rent ratio increased from 2.5 in 1983 to 2.7 by the end of the sample, a rise of 0.2 units. The duration-implied series overpredicts this increase, rising by 0.5 under $\phi = 1$ and by 0.4 under $\phi = 0.99$. This pattern

[20] We use data through 2025 in this case, as our duration and interest rate series extend to this date. In our main estimation sample, we instead use data only through 2023, as we are constrained by the availability of WID inequality data, which extend only through this date.

suggests that the effective duration of housing may be lower than what we infer from the data.[21] In particular, at the beginning of the sample, long-term real interest rates decline sharply, yet the observed price-rent ratio remains flat, inconsistent with our measured duration.[22] Outside this early period the duration-implied measure tracks the observed variation more closely.

In summary, our duration-based approximation either matches or slightly underestimates the rise in equity valuations across different values of ϕ, suggesting an appropriate value for the duration of equity. Our approximation overestimates the increase in housing values over our sample period, suggesting that our estimated duration for housing is if anything too high. Consequently, the duration gap between equities and housing used in our analysis is likely an understatement of the true duration gap. Our use of a conservative duration gap estimate will tend to weaken our duration heterogeneity mechanism in accounting for the rise in wealth inequality.

4 Marking the Wealth Distribution to Market

In the preceding sections, we introduced household-level measures of duration and estimated shocks to interest rates and long-term growth expectations. We now combine these elements to assess the paper's central prediction: how heterogeneity in duration shapes the dynamics of wealth inequality. First, in Section 4.1, we approximate the effect of the total change in rates over the sample in a single "one-shot" formula. Second, in Section 4.2, we dynamically mark the wealth distribution to market each period by gradually feeding in real rate shocks. While our baseline exercise focuses on the evolution of inequality driven solely by changes in real interest rates, as in Proposition 4, we also report results that incorporate variation in g, derived in Proposition 5.

4.1 One-Shot Repricing

We first use Equation (12) to compute the predicted change in wealth inequality arising from observed changes to interest rates. In this stylized one-shot repricing experiment, we set $S^{\alpha} = 68.7\%$, the average top-10% share over the sample 1983–2023, yielding $S^{\alpha}(1 - S^{\alpha}) = 21.5\%$. We use the average value-weighted durations across the top 10% and bottom 90% of the wealth distribution. For $\phi = 1$ ($\phi = 0.99$), the top 10% have a value-weighted duration of 29.4 (20.7) years, while the bottom 90% have a value-weighted duration of 21.2 (16.4) years. We compute the implied change using the cumulative shift in the real interest rate, $\log R_T - \log R_0$, where $t = 0$ corresponds to

[21] For instance, Glaeser, Gottlieb, and Gyourko (2012) show that differences in discount rates across households can lead changes in interest rates to have smaller effects on house values compared to a standard user-cost framework.

[22] A similar pattern is visible in the price-rent ratio from the JST database.

1983 and $t = T$ to 2023. Substituting these values into (12), we obtain:

$$d\tilde{S}(\phi = 0.99) \simeq -\underbrace{21.5\%}_{S^{\alpha}(1-S^{\alpha})}\Big[\underbrace{(20.7-16.4)}_{\mathcal{D}^{top}-\mathcal{D}^{bottom}}\underbrace{(-4.3\%)}_{\log R_T-\log R_0}\Big] = 4.0\%,$$

$$d\tilde{S}(\phi = 1) \simeq -\underbrace{21.5\%}_{S^{\alpha}(1-S^{\alpha})}\Big[\underbrace{(29.4-21.2)}_{\mathcal{D}^{top}-\mathcal{D}^{bottom}}\underbrace{(-4.1\%)}_{\log R_T-\log R_0}\Big] = 7.2\%.$$

This exercise shows that differences in financial duration between the top 10% and bottom 90%—combined with the observed decline in real interest rates—represent an important driver of the rise in the top-10% wealth share. In the data, the top-10% wealth share rose by 6.6pp between 1983 and 2023. Our duration-based mechanism accounts for 109.3% of this increase for $\phi = 1$ and 60.3% for $\phi = 0.99$.

Next, we incorporate the effects of changes in long-term growth expectations. We now also use the growth-adjusted durations $\mathcal{D}_G^{top} = 28.8$ (20.1), and $\mathcal{D}_G^{bottom} = 21.6$ (16.8), for $\phi = 1$ ($\phi = 0.99$), together with the cumulative change in expected growth $g_T - g_0$, which is 0.2% for our baseline measure. The overall change in the top 10% share implied by (19) is:

$$d\tilde{S}(\phi = 0.99) \simeq -\underbrace{21.5\%}_{S^{\alpha}(1-S^{\alpha})}\Big[\underbrace{(20.7-16.4)}_{\mathcal{D}^{top}-\mathcal{D}^{bottom}}\underbrace{(-4.3\%)}_{\log R_T-\log R_0} - \underbrace{(20.1-16.8)}_{\mathcal{D}_G^{top}-\mathcal{D}_G^{bottom}}\underbrace{(0.2\%)}_{g_T-g_0}\Big] = 4.1\%,$$

$$d\tilde{S}(\phi = 1) \simeq -\underbrace{21.5\%}_{S^{\alpha}(1-S^{\alpha})}\Big[\underbrace{(29.4-21.2)}_{\mathcal{D}^{top}-\mathcal{D}^{bottom}}\underbrace{(-4.1\%)}_{\log R_T-\log R_0} - \underbrace{(28.8-21.6)}_{\mathcal{D}_G^{top}-\mathcal{D}_G^{bottom}}\underbrace{(0.2\%)}_{g_T-g_0}\Big] = 7.5\%.$$

This simple calculation highlights the core intuition of our results: the long-term decline in real interest rates has been a key contributor to the rise in wealth inequality over the past four decades. The change in expected growth rates acts as a modest wealth inequality amplifier.

4.2 Dynamic Repricing

We now refine this calculation in several ways to arrive at our main results. First, rather than applying the full 40-year change in interest rates all at once using sample-average durations and wealth shares, we now dynamically update the wealth distribution in each year using the wealth shares and household durations prevailing in each period, along with the observed changes in interest and growth rates between each year t and $t+1$. This process generates the duration-implied change in wealth shares for each year, which we then chain together to obtain the full-sample effect. This approach naturally accounts for covariances. For instance, if interest rates decline more

in periods when the duration gap between the top and bottom of the wealth distribution is large, this comovement amplifies the impact of declining rates on wealth inequality.

Given household-level portfolio durations in (7) and (17), computed for all households across SCF survey waves from 1983 to 2022, together with annual changes in interest rates and growth expectations, we revalue each household's wealth W^i between periods t and $t+1$ using

$$\tilde{W}^i_{t+1} = W^i_t \times \exp\left\{-\mathcal{D}^i_t(\phi)\left(\log R_{t+1} - \log R_t\right) + \mathcal{D}^i_{G,t}(\phi)(g_{t+1} - g_t)\right\}. \tag{21}$$

We then compute $d\tilde{S}_{t,t+1}$ as the difference between the top-10% wealth shares under the repriced distribution $\{\tilde{W}^i_{t+1}\}$ and the original distribution $\{W^i_t\}$. This difference measures the change in the top-10% wealth share between periods t and $t+1$ that can be attributed to changes in interest rates, growth expectations, and the distribution of portfolio durations. We initialize $\tilde{S}$ to match the observed top 10% wealth share in the data in 1983 and then cumulate the successive $d\tilde{S}_{t,t+1}$ values to obtain the full time series of duration-implied wealth inequality.[23]

Figure 6 shows the predicted evolution of the top-10% wealth share implied by the chained repricing exercise for the permanent ($\phi = 1$) and persistent ($\phi = 0.99$) cases, alongside the observed top-10% wealth share from the WID data (black line). As discussed earlier the top-10% share increased from 63.0% in 1983 to 69.6% in 2023, a difference of 6.6pp.

Under the permanent case ($\phi = 1$), applying the observed variation in interest rates (red dashed line in Panel 6a) generates an implied path of inequality that is extremely close to its data counterpart. The top-10% share increases by 6.3pp, almost exactly matching the observed full-sample rise in the top-10% wealth share.

Incorporating variation in g (green dash-dot line) increases the implied change in inequality, for a full-sample total of 9.0pp. Even though expected growth g barely changed from the start until the end of the sample, including growth rate shocks in our computation amplifies the rise in duration-implied wealth inequality. To understand this result, it is useful to return to (19). In that expression, the impact of ε_g depends on two components: $S^\alpha(1-S^\alpha)$ and $(\mathcal{D}^{\text{top}}_G - \mathcal{D}^{\text{bottom}}_G)$. The first term captures that effects on inequality are larger when wealth shares are more balanced.[24] This term was larger at the beginning of the sample when the top-10% wealth share was smaller and g was increasing.[25] The second term measures the duration gap between the top 10% and the bottom 90%. This gap was larger in periods when g was rising (which was mostly in the first

[23] The SCF is generally triennial while the repricing is cumulated annually. Since each wave's portfolio is held fixed until the next wave, we assign to each wave the average of each asset's duration over the years that wave spans.

[24] For instance, if the wealthy hold 100% of all wealth, revaluations cannot change wealth shares at all.

[25] The term $S^\alpha(1-S^\alpha)$ reaches its maximum when $S^\alpha = 50\%$. Since the top-10% share has increased over time from a value of 63% in 1983, this term declined over time.

Figure 6: Predicted and Observed Top-10% Wealth Shares

(a) $\phi = 1$ (b) $\phi = 0.99$

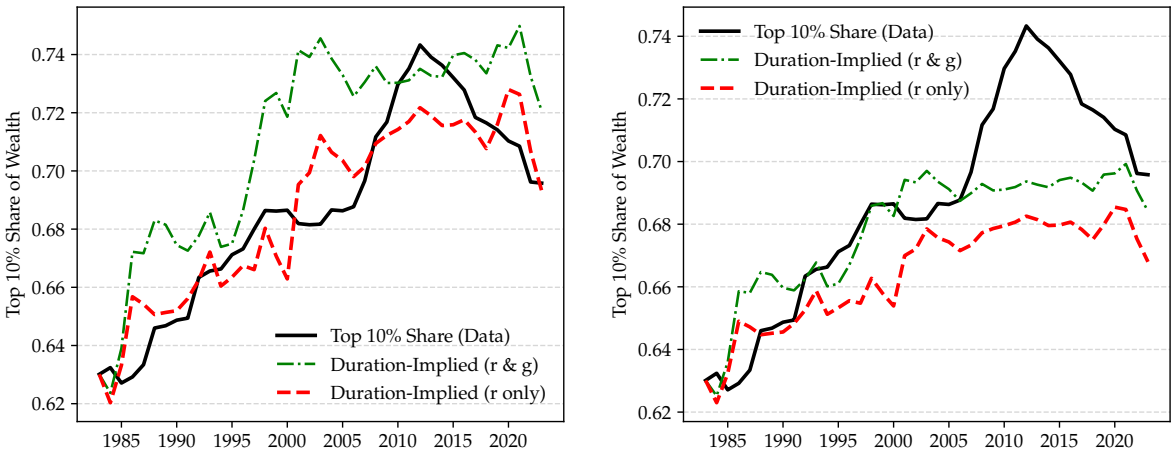

Note: The figure compares the observed top-10% wealth share (black line) with the duration-implied series for the permanent ($\phi = 1$) and persistent ($\phi = 0.99$) cases. The repricing simulation compounds annual changes in interest rates and growth expectations, starting from the observed 1983 level.

half of the sample) than in periods when g was falling (mostly in the last twenty years). These temporal patterns lead to net effects on total inequality even though the growth rate eventually returned to its starting point. Together, these patterns highlight the importance of accounting for the covariance between shocks to g (and r) and the cross-sectional distribution of wealth and durations, rather than relying solely on average effects.

Figure 6b presents the corresponding results under $\phi = 0.99$. The implied paths have a similar shape but lower level, due to the diminished effect of transitory interest rate changes on asset prices. The total effect of interest rate changes on the top-10% share (red dashed line) is 3.8 percentage points, while incorporating changes in g increases this difference to 5.4pp.

Overall, both the permanent ($\phi = 1$) and the persistent ($\phi = 0.99$) cases account for a large share of the observed increase in top wealth shares. The permanent case explains virtually all (96%) of the observed rise in inequality from variation in interest rates alone, while the persistent case explains approximately 58% of the observed increase.[26] These main results show that the cross-sectional distribution of duration, combined with the observed evolution of real interest rates and growth expectations, is a leading driver of the rise in wealth inequality over our sample.

[26]Figure E.2 in the appendix repeats the above exercise for the alternative measure of expected growth g, based on the SPF. For $\phi = 1$, the model-implied increase in the top-10% wealth share is 7.0pp. This is still larger than the observed increase, but closer to the data compared to the baseline measure of g. For $\phi = 0.99$, the corresponding increase is 3.9pp.

4.3 Heterogeneity in Returns

Fagereng et al. (2020) and Bach, Calvet, and Sodini (2020) use Scandinavian data to show that wealthier households have experienced higher portfolio returns. Our paper identifies one mechanism contributing to this return variation: heterogeneity in portfolio durations. To quantify the impact of this mechanism on realized returns over the past forty years, we compute the returns on wealth for each wealth decile. We find that US households in the top-10% experienced an extra annualized real return of 3.66% from the unexpected shocks to r and g, while the bottom-90% experienced an extra return of 2.17%. An average 1.50pp difference in annual returns sustained over 40 years exerts powerful effects on top-wealth inequality.[27]

5 Robustness of the Mechanism

Our main empirical results, presented in Section 4, show that variation in interest rates has been an important driver of the rise in wealth inequality over the past forty years, operating through heterogeneity in portfolio duration. We now re-estimate this implied increase in wealth inequality under a range of alternative assumptions. This exercise not only serves as a robustness check on our core findings, but also sheds light on the relative importance of the different channels contributing to the long-term rise in wealth inequality.

Table 2 reports changes in wealth inequality between 1983 and 2023. Columns 2–3 present results for the permanent case ($\phi = 1$), while Columns 4–5 report results for $\phi = 0.99$. We report outcomes for the top-10% (Columns 2 and 4) and the top-1% (Columns 3 and 5). The first row reports the increase in the top-10% wealth share observed in the WID data (6.57pp). Row 2 reports the duration-implied increase in the top-10% wealth share in the baseline (discussed above) where expected growth g is constant, equal to 6.31pp for $\phi = 1$ (3.79pp for $\phi = 0.99$).

Growth Expectations. Panel A explores alternative assumptions for the expected growth rate of the economy. Row 3 is our baseline case, discussed in Section 3.2, where we use the average realized GDP growth rate over the preceding six years (MA). As we noted earlier, duration-implied increases in wealth inequality are 9.05pp (5.40pp for $\phi = 0.99$). Row 4 replaces the baseline growth rate expectations measure with long-term real GDP growth expectations from the SPF. Under this specification, wealth inequality increases by 7.01pp (3.93pp for $\phi = 0.99$). While this is a smaller effect than in Row 3, including the alternative growth expectations still results in larger increases

[27] Table E.2 reports the arithmetic average of the additional returns on wealth experienced by households in different wealth deciles due to unexpected shocks to interest rates, r, and expected growth, g, over our sample period.

in wealth inequality than in the baseline with constant growth expectations in Row 2. Overall, our conclusions are robust to either measure of g.

Private Business Wealth. As noted above, the most challenging component of household wealth to measure is PBW. Appendices D.3–D.5 discuss several alternative approaches for estimating the durations of Corporate and Non-Corporate PBW. Panel B of Table 2 reports the implied changes in wealth inequality using these alternative measures. Rows 5 and 6 replace CRSP small-stock data with IPO and PitchBook data to characterize the high-growth stage in the two-stage GGM. The implied increases in wealth inequality are close to the baseline. Row 7 uses data on business income and valuation from SCF to estimate the duration of Corporate PBW, resulting in a somewhat larger increase in the top-10% wealth share of 8.17pp (4.61pp under $\phi = 0.99$). Row 8 uses SCF data to compute the duration of both Corporate and Non-Corporate PBW, leading to a substantially larger effect (13.15pp under $\phi = 1$ and 6.82pp under $\phi = 0.99$). Finally, Row 9 applies the duration of the overall equity market as a proxy for both Corporate and Non-Corporate PBW, producing effects that are slightly smaller but broadly comparable to the baseline (5.24pp under $\phi = 1$ and 3.21pp under $\phi = 0.99$). In sum, our conclusions are robust to a wide range of alternative PBW duration measures. Under most alternatives, the implied rise in wealth inequality is either larger than or similar to the baseline.

Housing Duration. Panel C of Table 2 shows that our findings are robust to alternative measures of housing duration. Row 10 replicates our measurement exercise using housing duration constructed from the JST dataset and yields an increase in the top-10% wealth share nearly identical to our baseline estimate. Alternatively, Row 11 considers a robustness check where we set the duration of housing such that the duration-implied rise in the price-rent ratio matches that in the data. This implies a housing duration of 6.2 under $\phi = 1$ and 5.9 under $\phi = 0.99$, which are lower than our baseline estimates of 13.8 and 12.2, respectively. The rise in wealth inequality exceeds that in the baseline substantially. Under $\phi = 1$, the implied top-10% wealth share increases by 8.71pp, while under $\phi = 0.99$ it increases by 6.09pp. Row 12 recomputes our results excluding primary home and mortgage from the household portfolio altogether.[28] We implement this by setting the duration of housing assets and mortgages to zero. Our results strengthen relative to the baseline,

[28]The motivation for this exercise is as follows. Since the house has a higher duration than the mortgage, our baseline model implies that a fall in rates increases household wealth, thereby expanding consumption opportunities. An alternative view is that, when households stay in the same house permanently, a change in the discount rate increases the value of the house and the value of the lifetime consumption of housing services by the same amount, leaving wealth and future non-housing consumption opportunities unchanged. This alternative view implicitly assumes that households do not change house size over the life-cycle nor extract home equity. Sodini et al. (2023) provides evidence from Sweden that households can and do borrow against rising collateral values to increase consumption opportunities.

Table 2: Alternative Specifications

	$\phi = 1$		$\phi = 0.99$	
	Top-10%	Top-1%	Top-10%	Top-1%
1. Data (WID)	+6.57pp	+10.87pp	+6.57pp	+10.87pp
2. Baseline	+6.31pp	+7.92pp	+3.79pp	+4.80pp
Panel A. Including Changes in Growth Expectations				
3. Moving Average (MA) g	+9.05pp	+10.82pp	+5.40pp	+6.34pp
4. SPF g	+7.01pp	+8.86pp	+3.93pp	+4.79pp
Panel B. Robustness to Private Business Wealth				
5. Corporate PBW from IPO data	+6.22pp	+7.74pp	+3.70pp	+4.66pp
6. Corporate PBW from Pitchbook	+6.92pp	+8.67pp	+4.27pp	+5.47pp
7. Corporate PBW from SCF	+8.17pp	+10.33pp	+4.61pp	+5.88pp
8. All PBW from SCF	+13.15pp	+14.80pp	+6.82pp	+7.59pp
9. All PBW from Equities	+5.24pp	+6.45pp	+3.21pp	+4.05pp
Panel C. Robustness to Housing Wealth				
10. Housing from JST	+6.33pp	+7.94pp	+3.79pp	+4.81pp
11. Matching rise in price-rent	+8.71pp	+9.55pp	+6.09pp	+6.50pp
12. Excluding primary home	+10.85pp	+10.84pp	+8.31pp	+7.96pp
Panel D. Sources of Time Variation in Duration				
13. Assets have average duration	+8.53pp	+10.56pp	+4.68pp	+5.93pp
14. Households have average portfolios	+5.10pp	+7.08pp	+3.02pp	+4.40pp
15. Average duration and portfolios	+7.42pp	+9.10pp	+3.90pp	+5.05pp

Note: The table reports changes in wealth inequality between 1983 and 2023, based on the WID data (Row 1) or on our duration-implied revaluations under different sets of assumptions (Rows 2 through 15). Columns 2–3 present results for the permanent case ($\phi = 1$), while columns 4–5 report results for $\phi = 0.99$. Within each case, we report outcomes for the top-10% (columns 2 and 4) and the top-1% (columns 3 and 5).

with a rise in the top-10% wealth share of 10.85pp.

These three exercises share a key insight. Housing is a relatively long-duration asset held broadly by the middle class, compared to public and private equity holdings which are concentrated at the top of the wealth distribution. Removing housing or lowering its duration results in a larger gap between the durations of the top-10% and bottom-90% of the wealth distribution, increasing the effect of repricing on wealth inequality.

Time-Varying Asset Durations vs. Portfolio Shares. Our benchmark results allow both the durations of each asset and liability class and household portfolio shares across asset classes to vary over time. Panel D of Table 2 isolates the importance of each source of time variation. Row 13

keeps the durations of all assets and liabilities constant over time, at their full-sample averages, but allows time variation in portfolio shares. This exercise generates a larger increase in the top-10% wealth share of +8.53pp. Time-varying asset durations dampen the rise in wealth inequality because top-minus-bottom duration gaps were smaller in periods of larger rate declines.

The opposite exercise—holding households' portfolio shares constant over time while allowing asset durations to vary—dampens the rise in wealth inequality, as shown in Row 14. This implies that time-varying portfolio shares contributed positively to wealth inequality. Over our sample, the bottom 90% increased the share of wealth allocated to public and private equities while the top-10% maintained a relatively stable exposure to long-duration assets. In the 1980s, the gap in equity and PBW portfolio shares between the top 10% and the bottom 90% was particularly large. Since this decade also experienced the largest declines in interest rates (Figure 4a), these changes in portfolio composition amplified the rise in wealth inequality.

Row 15 holds both duration and portfolios constant, finding results in between these individual cases, though somewhat above the baseline in Row 2. In sum, the dynamics in households' portfolio choices contributed positively to the rise in inequality, while time variation in durations dampened inequality. These partially offset when combined, explaining why our one-shot and our dynamic repricing ended up with similar results.

Top-1% Inequality. Table 2 also reports changes in the wealth share of the top-1%. In the WID data, the top-1% share increases by 10.87pp, 4.3pp more than the increase in the top-10% share. Under the baseline assumptions, the duration-based mechanism implies an increase in the top-1% wealth share of 7.92pp in the permanent case and 4.80pp for $\phi = 0.99$. The mechanism therefore correctly captures the larger increase for the top-1% relative to the top-10%. As discussed in Section 2.2 and shown in Figure 2, households in the top-1% have higher portfolio duration than those in the 90%–99% group. This difference is driven by higher shares of private business wealth. Allowing for time variation in g brings the model even closer to the data. Using the baseline moving-average measure of g in the permanent case ($\phi = 1$), Row 3 shows that the top-1% share increases by 10.82pp, almost exactly matching the data. Using the SPF-based measure of g, the implied increase is 8.86pp (Row 4). The remaining robustness checks in rows 5–15 deliver results for the top-1% wealth share that are similar to those discussed above for the top-10% share. Overall, the duration-based mechanism captures the dynamics of the top-1% wealth share well, explaining a larger increase for the top-1% than for the top-10% wealth share.

6 Long-Term Dynamics from a Calibrated Life-Cycle Model

Up to this point, our results have used empirical measures of the joint distribution of wealth and interest rate exposure to study how inequality evolves with changes in interest rates, without imposing additional structural assumptions. This model-free approach abstracts from the fact that interest rate and asset price changes unfolded over multiple decades, giving households substantial time to adjust their consumption and savings behavior in response to wealth revaluations. Because the effects of asset revaluations on inequality are driven by changes in interest rates rather than their new level, they may eventually dissipate as the economy transitions back to steady state. If this reversion is sufficiently rapid, initially large inequality changes may not persist over our relatively long sample period. To assess this possibility, we develop a structural model of consumption and savings: a workhorse life-cycle framework with idiosyncratic labor income risk. We use this model to examine the indirect effects of interest rate changes on inequality through household consumption-savings responses.

In this setting, we seek to match as closely as possible the sequence of repricing events estimated in Section 4. To this end, we study a partial-equilibrium model (equivalently, a small open economy) in which interest rates are exogenous, allowing us to feed in the observed path of rates. By assigning each household a portfolio exposure (duration) consistent with our empirical estimates by age and wealth, we ensure that revaluation effects in the model almost exactly replicate those in the data. After each revaluation event, consumption adjusts endogenously as the economy transitions toward the new steady state. This framework enables us to study how the model's implied consumption-savings responses shape the propagation of revaluation-driven changes in the wealth distribution over time.

We find that the total implied change in the model is very similar to the model-free results obtained by simply cumulating the revaluation-implied changes. This finding is due to the high persistence of changes in the wealth distribution in the model, as well as the unusually compressed initial wealth distribution in 1983 following a sequence of earlier interest rate increases.

6.1 Model Setup

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 ψ_j, with $\psi_J = 0$. When a household dies, it is replaced by a newborn household ($j = 0$), which inherits its remaining assets as a bequest.

Each household i of age j is endowed with exogenous labor income given by $y_j(z)$, where

z is a household-specific (idiosyncratic) stochastic process. In addition, households can trade a complete set of risk-free bonds offering fixed cash flows at future dates. However, markets are incomplete in that households cannot contract on their idiosyncratic income realizations. 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 the bond with maturity m as x_m, and its price as q_m.

We consider an economy in which no future aggregate shocks are expected to occur, and bond prices are expected to remain constant forever, corresponding to a permanent interest rate process ($\phi = 1$). The case $\phi = 1$ is advantageous for simplicity, and because it allows us to separate current effects on inequality from effects on steady-state inequality. We provide results for $\phi = 0.99$ in Appendix G.3 and show that the implications for model dynamics are similar. We further assume that the one-period bond is traded on a global market, so that its interest rate takes the exogenous value R. If we normalize $q_0 = 1$, the absence of arbitrage opportunities requires $q_m = R^{-m}, \forall m$.

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

$$V_j(x;z) = \max_{c,x'} \underbrace{\frac{c^{1-\gamma}}{1-\gamma}}_{\text{flow utility}} + \beta \left\{ \underbrace{\psi_j \mathbb{E}\left[V_{j+1}(x';z') \,\middle|\, z \right]}_{\text{continuation value}} + \underbrace{(1-\psi_j)\chi \frac{\left(\sum_{m=1}^{M} q_{m-1} x'_m\right)^{1-\gamma}}{1-\gamma}}_{\text{bequest utility}} \right\} \tag{22}$$

subject to the budget constraint[29]

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

While past work has shown empirical benefits from using non-homothetic bequest functions (De Nardi, 2004; Straub, 2019), our use of a homothetic bequest function offers large gains in tractability for our theoretical and quantitative analysis. Because these non-homothetic bequest functions imply greater savings by the rich following wealth gains, and hence greater amplification and persistence of wealth inequality, these specifications would imply even larger increases in inequality from declining interest rates.

[29] An explicit borrowing constraint is not necessary, as the marginal utility of bequests increases without bound as bequests approach zero, implying that households will always save a nonzero amount.

Each of the m optimality conditions for bond holdings x'_m collapses to the same Euler equation

$$c^{-\gamma} = \beta R \left\{ \psi_j \mathbb{E}\left[(c')^{-\gamma} \,\middle|\, z \right] + (1-\psi_j)\chi \left(\sum_{m=1}^{M} q_{m-1} x'_m \right)^{-\gamma} \right\}. \tag{23}$$

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, and are thus indifferent between them. Instead, only total wealth $W \equiv \sum_{m=1}^{M} q_{m-1} x_m$ and total savings $s \equiv \sum_{m=1}^{M} q_m x'_m$ matter for the household's problem in the absence of future interest rate volatility. Given this indifference, we assign each household a unique portfolio $\{\hat{x}_m\}$ that matches its empirically predicted duration given its age and position in the wealth distribution, as discussed below.

6.2 Calibration

We calibrate our model at annual frequency. Under our baseline assumption of a permanent interest rate process ($\phi = 1$) the model has a different steady state for each value of R. For calibration, we match empirical averages taken over our sample to corresponding moments in the steady state where R is equal to its sample mean. Because our model is stationary, while the true economy features growth, we map parameters and variables from the growing economy to an equivalent stationary economy. Following Krueger and Lustig (2010), the average observed log real interest rate of r^* in the growing economy maps into $r = r^* - g + (\gamma - 1/2)\sigma^2$ in the stationary model economy, where g is the unconditional average of the log real growth rate, and σ is its volatility. For our sample period, $g = 0.0189$ is the average growth in log real per-capita GDP over our sample period, while $\sigma = 0.0213$ is its volatility.[30] Given that $r^* = 0.0126$ is the mean log real 1-year interest rate computed in Section 3.1, the growth-adjusted log real rate has mean $r = -0.0056$, or $R = 0.9944$.

We calibrate the model's survival probabilities $\{\psi_j\}$ to match Social Security Actuarial tables. Since we model households, we take the average of the male and female mortality rates, weighted by the proportion alive at each age. We set γ, the risk aversion of the households and the curvature of the bequest function, to a standard value of 2. We set the real time discount factor in the stationary economy to $\beta = 1.007$, corresponding to $\beta = 0.988$ in the growing economy, to target a ratio of wealth to disposable labor income under the average interest rate of 8.36, equal to the average value in the data over our sample.[31] We set the bequest utility parameter $\chi = 0.330$ to

[30] Real per capita GDP growth data are from the Bureau of Economic Analysis, retrieved from the St. Louis Fed FRED system (code A939RX0Q048SBEA).

[31] For this calculation, wealth is measured as the net worth of households and nonprofits in the FAUS (Table B.101). Disposable labor income is household income net of personal taxes, rental, interest, and dividend income, and one half

match a ratio of bequests to GDP of 8.8%, following Auclert et al. (2021).

In our model exercise we seek to faithfully capture the effects of asset repricing on household wealth estimated in Section 4, so that we can study the resulting consumption-savings response. To this end, we assign each household an empirically estimated duration so that the model reproduces as closely as possible the instantaneous effects of asset repricing on inequality. This is consistent with optimality as our model households are indifferent with respect to the duration of their portfolio because they do not anticipate rate shocks. Our theoretical results in Proposition 4 and equation (12) imply that sufficient statistics for matching revaluations are the average value-weighted durations in each wealth bin. To obtain these, we run the value-weighted regression

$$\mathcal{D}_{i,t}^{VW} = \alpha_t + \beta_t \widetilde{Age}_{i,t} + \sum_j \gamma_{j,t} NetWealthBin_{i,j,t} + e_{i,j,t} \tag{24}$$

for each SCF wave t, where $NetWealthBin_{i,j,t}$ is a dummy for being in wealth bin j, and $\widetilde{Age}_{i,t}$ is the residual from a separate value-weighted regression of age on our wealth bins.[32] To implement (24), we use 20 bins each corresponding to 5% of the value-weighted wealth distribution. We then assign to each model household at each date the duration $\hat{\mathcal{D}}_{i,t}^{VW}$ equal to the fitted value of (24), given that household's age and wealth bin. We re-estimate equation (24) for each SCF wave (typically every three years), then interpolate the coefficients for the years between SCF waves.

The labor income process consists of a regular component and a superstar component, detailed in Appendix F. As is standard in the literature, the regular income process contains a deterministic age profile, a vector of household characteristics, and a stochastic component. The stochastic income component contains a household-fixed effect, a persistent component, and an i.i.d. component. We allow the variances of these components to be different before and after age 65. The parameters are estimated by GMM using PSID data from 1970 until 2017.

The literature typically estimates the labor income process on data for males between ages 25 and 55. We deviate from this practice by: (i) considering a broader income concept, (ii) allowing a wider range of ages including retirement, and (iii) focusing on households rather than individuals.

First, from the model's perspective, the relevant notion of income includes transfers. This approach extends that of Catherine, Miller, and Sarin (2025), who incorporated Social Security benefits as a component of income, to include all sources of labor income and transfers available in the data. To that end, we measure income in the data as income from wages and salaries, the labor

of proprietor's income in the National Income and Product Accounts. For our computation of the discount factor in the growing economy, we follow Krueger and Lustig (2010) in computing $\beta^* = \beta \exp(-g - \gamma\sigma^2)$.

[32] Orthogonalizing age to wealth bins ensures that the model's value-weighted mean duration within each wealth bin matches its empirical counterpart. The two might otherwise differ due to discrepancies in the covariance of age and wealth between the model and data.

income component of proprietor's income, government transfers (unemployment benefits, Social Security, other government transfers), and private defined-benefit pension income. To obtain consistent data series, we reconcile differences in variable definitions among the various waves of the PSID, a non-trivial task detailed in Appendix F.2.

Second, we are interested in the entire life-cycle. To this end, we estimate the income distribution for a wide range of ages from 21 to 80, and back- and forward-fill our estimated parameters for ages outside this age range.[33] Because our income concept includes transfers such as unemployment benefits and retirement income from public or private defined-benefit pension plans, we do not need to assume that households are working full-time. Instead, our estimation takes into account the full cross-section of sample households, including retirees, recipients of unemployment benefits, part-time employees, and so on.

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

We enrich the regular income process with a superstar income state following Castaneda, Diaz-Gimenez, and Rios-Rull (2003). Households enter and exit this state with probabilities taken from Boar and Midrigan (2022). We choose this specification rather than mechanisms based on return heterogeneity for clarity, since realized return heterogeneity is an *output* of our model mechanism. These imply a roughly 1% probability of entering the superstar income state over one's lifetime. We calibrate the superstar income level so that the top-10% wealth share in the model's stationary economy matches the average top-10% wealth share in the WID data over our 1983–2023 sample.

6.3 Transition Path Results

In this section, we use our structural model to provide a full characterization of the contribution of interest rates to wealth inequality over the period 1983–2023.

Transition Experiment. The 1983 wealth distribution does not correspond to a stationary distribution, as long-term real interest rates had risen to very high levels in the 1980s. This temporarily suppressed inequality due to the repricing force we described above. To account for this, we initialize the 1983 wealth distribution by applying a one-time update to the steady-state distribution:

$$W_{i,1983} = K_{init} \times W_{i,steady} \exp\left(-\hat{\mathcal{D}}_{i,steady} \times \varepsilon_{r,init}\right), \tag{25}$$

[33] Since model households begin at age 18, we use the age 21 age profile embedded in $\chi' X_t^i$ for ages 18–20. Similarly, since model households can survive to age 100, we assume that the age profile in $\chi' X_t^i$ remains constant from age 80 onward.

where $W_{i,steady}$ is financial wealth for household i in the stationary distribution for our model computed at the sample average interest rate, $\hat{\mathcal{D}}_{i,steady}$ is the fitted value from (24) using a pooled regression over all SCF waves, $\varepsilon_{r,init}$ is the initial shock to interest rates that determines initial inequality, and K_{init} is a scaling constant set to match the aggregate wealth-to-income ratio in the 1983 economy. We jointly choose the values of $\varepsilon_{r,init}$ and K_{init} so that our initialized 1983 wealth distribution produces the observed top-10% wealth share from the WID in 1983 (63.01%) and 1983 wealth-to-income ratio from the Financial Accounts of the United States (6.79).

Next, for each year between 1983 and 2023, we feed in the annual changes in the detrended real rate, under the assumption that the growth rate is constant and that interest rate changes are permanent ($\phi = 1$) and unexpected.[34] In each period, we apply an offsetting change to the subjective discount rate from β to $\tilde{\beta}$ such that $\beta R = \tilde{\beta}\tilde{R}$, to approximate a change in interest rates associated with a change in effective patience.[35] At the start of each period, we revalue wealth using (4) applying that year's change in rates $\Delta \log R_t$ and a duration for each household equal to its fitted value from (24) based on its age and position in the wealth distribution.

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

Main Results. The implied path for the top-10% wealth share from our transition experiment is displayed in Figure 7, while Table 3 summarizes the full-sample changes for various inequality measures for our main specification with $\phi = 1$, as well as for the case of $\phi = 0.99$.[36] Table 3 shows that the overall effect of interest rates on inequality in our baseline model is close to that observed in the data for the top-10% share and the Gini coefficient, while explaining around half

[34] We note that, if the real rate declines had instead been anticipated by investors in 1983, the entire effect on valuations would have been front-loaded, and the term structure would have been steeply downward sloping. Both of these predictions are counterfactual.

[35] Jointly shocking R and β has two benefits relative to shocking R alone. First, it approximates the effect of an economic shock by which R is changing due to a change in β, which would be exactly true in a representative agent model, but will not hold exactly in our heterogeneous agents setting. Second, as we show in Greenwald et al. (2025), jointly shocking R and β implies that households' preferred slope of expected consumption via the Euler equation remains the same. This means that our consumption-savings responses are driven by wealth effects and precautionary savings behavior rather than changes in households' optimal rate of consumption growth, which is highly dependent on the calibration of the elasticity of intertemporal substitution.

[36] The model implies a modestly smaller rise in the top-10% share for the $\phi = 0.99$ case relative to the $\phi = 1$ case, but a slightly larger rise in the top-1% share. Appendix G.3 explains that the latter results from the larger initial rate shock required to match the 1983 wealth data.

of the total implied rise in the top-1% share. Thus, changes in inequality due to revaluations are highly persistent in our model, and not undone by endogenous consumption-savings responses, even over our long sample period.

Table 3: Model-Implied Change in Wealth Inequality, Gradual Transition Experiment

		Baseline		Matching Income Inequality	
	Data (WID)	$\phi = 1$	$\phi = 0.99$	$\phi = 1$	$\phi = 0.99$
Top-10%	+6.57pp	+5.6pp	+4.9pp	+8.3pp	+8.0pp
Top-1%	+10.87pp	+6.6pp	+7.5pp	+10.4pp	+12.3pp
Gini	+0.039	+0.033	+0.029	+0.049	+0.047

Note: The table reports the change in the Top-10% share, Top-1% share, and Gini Coefficient of wealth from 1983 to 2023. The first column displays this change for the WID data. The second and third columns reflect the implied change in the baseline model with constant income risk for the specifications $\phi = 1$ and $\phi = 0.99$, respectively. The fourth and fifth columns reflect the same model outputs in our specification with time-varying income risk calibrated to match the change in income inequality.

To better visualize this result, we decompose the total change in the top-10% wealth share into a component due to revaluations and a component due to transitions to the steady state. Let S_t^{10} be the top-10% share of wealth in the model at the start of time t, prior to any revaluation. Once the new interest rate R_t is realized, each household's pre-revaluation wealth $W_{i,t}$ is updated to $\tilde{W}_{i,t}$, 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 as $dS_t^{10,rev} = \tilde{S}_t^{10} - S_t^{10}$. We cumulate these revaluation effects to obtain the series $S_t^{10,rev} = S_0^{10} + \sum_{\tau=1}^{t} dS_\tau^{10,rev}$, where $t = 0$ represents the base period 1983. The resulting "Revaluations Only" series corresponds closely to our model-free contribution of interest rates to inequality in Section 4. To obtain an additive decomposition, we can also define a "Transitions Only" series, representing the effect of endogenous transitions back to steady state, to be equal to $S_t^{10,trans} = S_0^{10} + S_t^{10} - S_t^{10,rev}$.

The resulting series are displayed in Figure 7a, alongside the top-10% share in the WID data (black dash-dot line). The red solid line displays the total change in inequality, while the blue dashed line displays the Revaluations Only series, and the green dotted line displays the Transitions Only series. The main result of this analysis is that the model's total implied change in inequality hews closely to the Revaluations Only series. In other words, in a workhorse life-cycle model, the effects of portfolio revaluations on inequality are highly persistent, and are not undone by endogenous consumption-savings responses even over our 40-year sample. Correspondingly, the Transitions Only series contributes little to the overall response.

Closer inspection shows a more nuanced temporal pattern, where the model-implied inequal-

Figure 7: Model-Implied Top-10% Share

(a) Model vs. Revaluations, Baseline

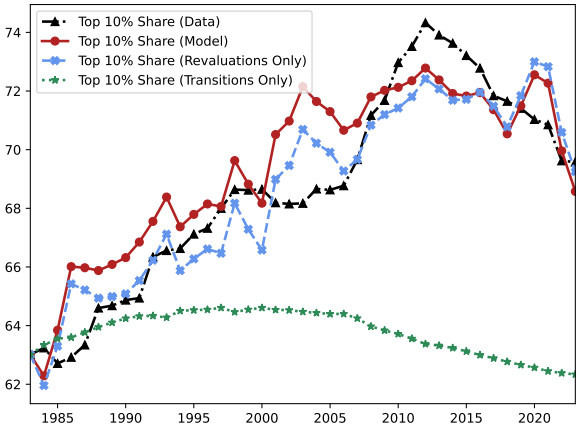

(b) Model vs. Steady State, Baseline

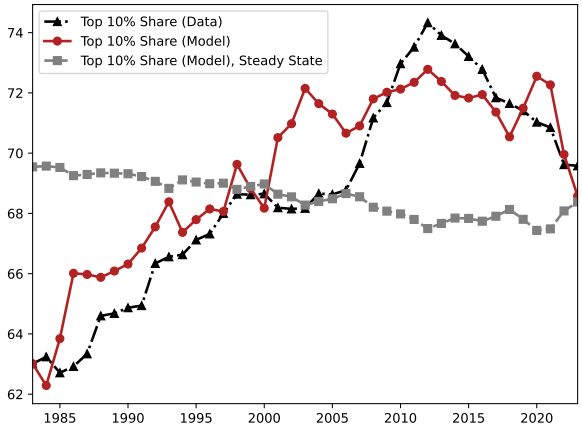

(c) Model vs. Revaluations, TV Inc. Risk

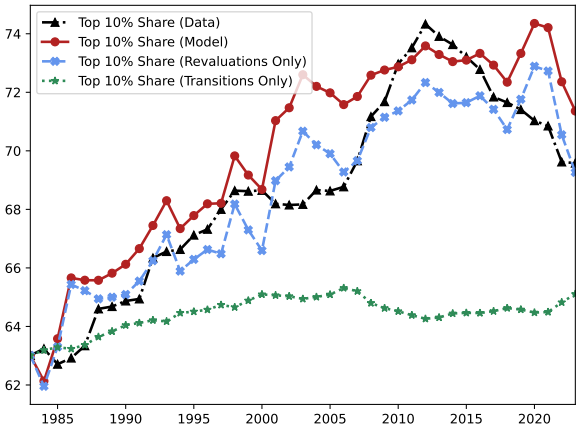

(d) Model vs. Steady State, TV Inc. Risk

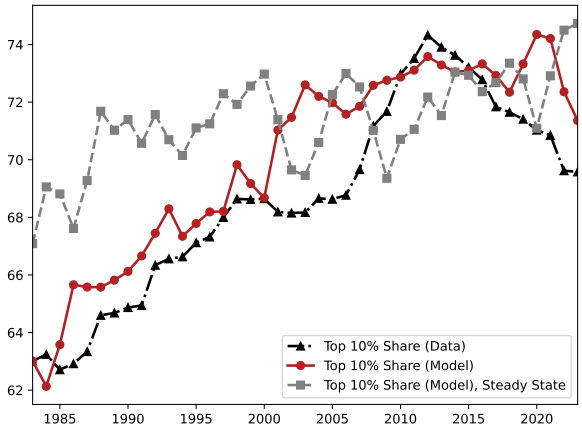

Note: In all panels, the series are defined as follows. "Top 10% Share (Data)" is the top-10% wealth share from the World Inequality Database. "Top 10% Share (Model)" is the model's implied path of the top-10% wealth share from our gradual transition experiment. "Top 10% Share (Revaluations Only)" is the implied contribution of revaluations $S_t^{10,rev}$. "Top 10% Share (Transitions Only)" is the implied contribution of endogenous transitions $S_t^{10,trans}$. "Top 10% Share (Model), Steady State" is the model's steady state top-10% share assuming the current interest rate holds in perpetuity. Panels (a) and (b) plot these series for the baseline model and experiment. Panels (c) and (d) plot these series for the alternative experiment in which we match the relative change in income inequality in each year.

ity series at first outpaces the Revaluations Only series until around the 2010s, after which overall model-implied inequality falls below the Revaluations Only series. To provide intuition for this result, Panel (b) displays the model's steady state top-10% wealth share (gray circles) that would hold in the absence of any future shocks, as the model households believe to be the case.

At the start of the sample, the top-10% wealth share is far below its steady state level. This is the result of the large rise in interest rates just prior to the start of our sample. Following our main mechanism, this devalued the portfolios of the wealthy by relatively more, compressing

inequality—a pattern we match through our initialization procedure described above. Thus, even though interest rates begin falling rapidly after 1983, inequality initially remains below steady state, and the model's transition back to steady state pulls top wealth shares *upwards* until the late 1990s, leading total model inequality to rise above that implied by the Revaluations Only series. In the second half of the sample, however, further falls in the real interest rate push inequality above its steady state, setting off an opposing transition force which pulls inequality down. This force leads the model's top-10% share to fall below the Revaluations Only series after the 2010s.

Panel (b) conveys an important additional intuition: while falling real interest rates increase inequality over our sample, they actually *decrease* inequality in the long-run steady state. Intuitively, under lower interest rates, there are fewer opportunities to accumulate wealth by compounding high returns. Thus, it is not the low level of interest rates that leads to rising inequality, but rather the *fall* in interest rates. While falling interest rates can lead to a highly persistent rise in inequality, this is ultimately a transitory phenomenon, with long-run inequality returning to a similar or even lower level under the 2023 level of interest rates compared to 1983. Panel (b) implies that inequality at the end of the sample is close to steady state, implying that we would expect inequality to remain stable going forward, absent further interest rate shocks.

Time-Varying Income Inequality. While our analysis focuses on the role of revaluations in driving wealth inequality, it is well known that income inequality has also increased over this period (e.g., Hubmer, Krusell, and Smith Jr, 2021). We now extend our model in two ways to study the effect of rising income inequality alongside our duration mechanism. First, we re-estimate our regular income process to allow for annually-varying income risk, using rolling 15-year samples of the PSID (Appendix F.3). Second, we update the superstar income state each period so that we exactly match the normalized top-10% *income* share in each year.[37] Both changes are assumed by households to be permanent, with no further updates expected at each point in time.

The increase in income inequality over the past forty years increases wealth inequality substantially. However, the combined effect of income inequality and revaluations produces a plausible series for the dynamics of top wealth shares that remains broadly consistent with the observed path. Table 3 shows that this model somewhat overstates the rise in the top-10% share and Gini coefficient, but provides a good match for the total increase in the top-1% wealth share. Appendix Figure G.2 shows that the model with $\phi = 0.99$ and time-varying income inequality provides a

[37] To implement each year's change in the income process in the model, we assume that each household maintains its position or index in the discrete grid of incomes, and simply update the value associated with that position. Because we calibrate our superstar income state to match wealth inequality rather than income inequality, our model's top-10% income share does not match that of the data on average. However, our experiment ensures that the ratio of the top-10% income share across any two points in time is the same in the model as in the data.

relatively close fit of the path of the top-10% wealth share.

Figure 7 panels (c) and (d) are the counterparts of panels (a) and (b) for the case with time-varying income risk. The model-implied top-10% share remains consistently above the Revaluations Only series throughout the sample. The reason for this can be seen in Panel (d), which shows the evolution of steady-state wealth inequality under rising income inequality. In contrast to the baseline experiment, in which rising wealth inequality surpassed steady-state wealth inequality in the late 1990s, this experiment features steady-state wealth inequality levels that rise throughout the sample. As a result, in the model with rising income inequality, transitions to the new steady state do not meaningfully decrease wealth inequality.

Finally, while the realized time series for wealth inequality are broadly similar in our baseline and time-varying income risk models, the long-run implications are meaningfully different. In our baseline experiment in Panel (b), observed top wealth shares are close to the final steady state at the end of our sample, implying that inequality should be expected to remain stable going forward. In the model with increasing income inequality in Panel (d), steady-state wealth inequality at the end of the sample substantially exceeds both the final empirical and the contemporaneous model-implied wealth inequality observations. Under this model, in the absence of further rate shocks, we should expect inequality to gradually rise to the highest levels yet observed in the sample. Higher income inequality takes a long time to pass through to higher wealth inequality in the model. Hence, while rising income inequality has a limited effect on wealth inequality over our sample, over which the persistent yet temporary effects of asset revaluations have dominated, it is expected to play a larger role in shifting inequality over the long term.

7 Wealth Inequality and Real Rates: International Evidence

To close, we look beyond the United States for further evidence on the relationship between long-term real interest rates and wealth inequality, and study the United Kingdom (UK) and France over the same period from 1983 until 2023. While the wealth inequality changes were noticeably different from those in the US, the same duration mechanism goes a long way towards accounting for these differences. As we did for the US in Section 2, we quantify the increase in wealth inequality by combining survey-based measures of household portfolios with asset class duration estimates. Appendix H provides a detailed discussion of measurement and data sources.

For the UK, we use the Wealth and Assets Survey (WAS), a biennial longitudinal survey conducted by the Office for National Statistics that collects detailed information on households' assets, savings, debts, and retirement planning. We use the micro-level data to construct household

portfolio shares using the same asset class definitions as in Table 1.

The WAS does not separately report private business wealth. We therefore group all equity-like components into a single *Equities* category. Although the survey reports holdings in mutual funds and defined contribution pension plans, the reported equity and fixed income categories exclude indirect equity exposures through investment funds. We therefore assign 60% of mutual fund and defined contribution pension holdings to equities and the remaining 40% to fixed income, consistent with a standard, balanced portfolio.[38] Finally, because private business wealth is not separately reported, our equity measure may understate total equity holdings. To address this concern, we rescale household equity holdings so that average equity wealth in the survey matches aggregate equity holdings in the UK flow-of-funds data.[39]

For France, we use the Household Finance and Consumption Survey (HFCS), conducted by the ECB for all Euro area countries. Unlike in the UK, the HFCS reports private business wealth. To make the results comparable between the UK and France, we again combine public equity and private business wealth into a single Equities category and apportion investment funds to equity and fixed income using the 60/40 approach.

Unlike for the US, household wealth surveys for the UK and France are available only in recent years. As a result, we cannot construct portfolio shares over the full 1983–2023 period. For the UK, the WAS is available from 2006 onward, while for France the HFCS is available from 2010. In addition, the earliest survey waves provide less detailed asset information than later waves. We therefore rely on the most recent survey wave available for each country, which offers the most comprehensive and consistent asset coverage.[40]

Table 4 reports the portfolio shares for the UK in Panel A and for France in Panel B. A number of interesting facts stand out. First, the portfolio share in real estate for all households is notably higher in both UK (82.1%) and France (76.9%), compared to the US (46.8%). Correspondingly, equity holdings are much lower in the UK (12.3%) and France (14.4%) than in the US (42.5%).[41] The comparatively low allocation to equity in the UK and France persists even at the top of the wealth distribution, leading to more modest differences between the portfolios of the bottom 90% and the top 10% in these countries relative to the United States.

We then combine portfolio shares with the duration of assets and liabilities. Durations are computed under the assumption that interest rate changes are permanent ($\phi = 1$). For equities,

[38] This assumption is supported by UK flow-of-funds data, which indicate that the average equity share of pension funds over the period 1987–2022 (the time period when these data are available) is 62%.

[39] This adjustment follows the approach used in Kuhn, Schularick, and Steins (2020).

[40] For the UK, we use the 2020 WAS. For France, we use the 2021 HFCS.

[41] The value of 42.5% is obtained by summing the portfolio shares of equities, corporate PBW, and non-corporate PBW in the United States.

we use duration estimates based on the MSCI World Index.[42] The resulting equity duration over the full sample is 41.3, which is lower than the corresponding US measure.

Table 4: Portfolio Composition and Asset Duration: UK and France

Panel A. United Kingdom

	Duration	Portfolio Share (%)				
Asset Class		All	Bottom 90	P90-99	Top 10	Top 1
Equities	41.31	12.33	6.52	14.60	19.07	31.19
Real Estate	15.06	82.11	95.87	70.77	66.16	53.67
Vehicles	3.46	3.36	4.79	1.91	1.70	1.14
Fixed Income	5.58	8.27	6.78	9.94	9.99	10.12
Cash and Deposits	0.25	11.17	12.31	10.44	9.84	8.22
Mortgage Debt	0.83	-15.46	-23.22	-7.28	-6.46	-4.21
Other Debt	1.00	-1.79	-3.05	-0.38	-0.32	-0.13
Equal-Weighted (EW) Duration		14.90	14.63	16.98	17.33	20.38
Value-Weighted (VW) Duration		**17.92**	**17.48**	**17.27**	**18.43**	**21.55**

Panel B. France

	Duration	Portfolio Share (%)				
Asset Class		All	Bottom 90	P90-99	Top 10	Top 1
Equities	41.31	14.40	4.89	13.67	23.59	42.87
Real Estate	16.82	76.92	92.55	71.20	61.81	43.58
Vehicles	3.46	3.34	5.25	1.83	1.49	0.81
Fixed Income	5.58	9.90	7.33	12.63	12.39	11.91
Cash and Deposits	0.25	10.14	13.83	7.90	6.58	4.01
Mortgage Debt	5.89	-12.75	-20.70	-6.27	-5.05	-2.69
Other Debt	1.00	-1.95	-3.14	-0.97	-0.81	-0.50
Equal-Weighted (EW) Duration		13.25	12.71	17.60	18.17	23.30
Value-Weighted (VW) Duration		**18.81**	**16.96**	**18.03**	**20.60**	**25.58**

Note: The table reports, for each country, the duration of each asset class in column (1) and portfolio shares by wealth group in columns (2) through (6): all households (All), the bottom 90%, households in the 90th to 99th percentiles (P90–P99), the top 10%, and the top 1%. Durations are computed as sample averages over the period 1983–2023 under the assumption that interest rate changes are permanent ($\phi = 1$). The bottom rows of each panel report the equal-weighted (EW) and value-weighted (VW) portfolio durations of each group. Data for the United Kingdom (Panel A) are based on the WAS, Wave 8; data for France (Panel B) are based on the HFCS, 2021 wave.

For real estate in the UK, we use the price-to-rent ratio from Giglio et al. (2021), who estimate an average rent-to-price of 6.9% over the 1988–2016 period — a sample that broadly overlaps with our own. The duration estimate of 15.1 is above the US duration measure of 13.8. For France, we calculate a similar rental yield of 6.6% over the sample 1983–2023.[43] We maintain the same cost

[42] Our choice to use an international equity measure, rather than France- or UK-specific stock market index, is motivated by the fact that equity exposure in both countries is not confined to domestic markets. A substantial share of equity holdings is held through mutual funds and pension funds who hold internationally-diversified portfolios.

[43] We start from the average rental yield for France of 4.6% from the June 2025 edition of Global Property Guide. We

Figure 8: Predicted and Observed Top-10% Wealth Share for France and the UK

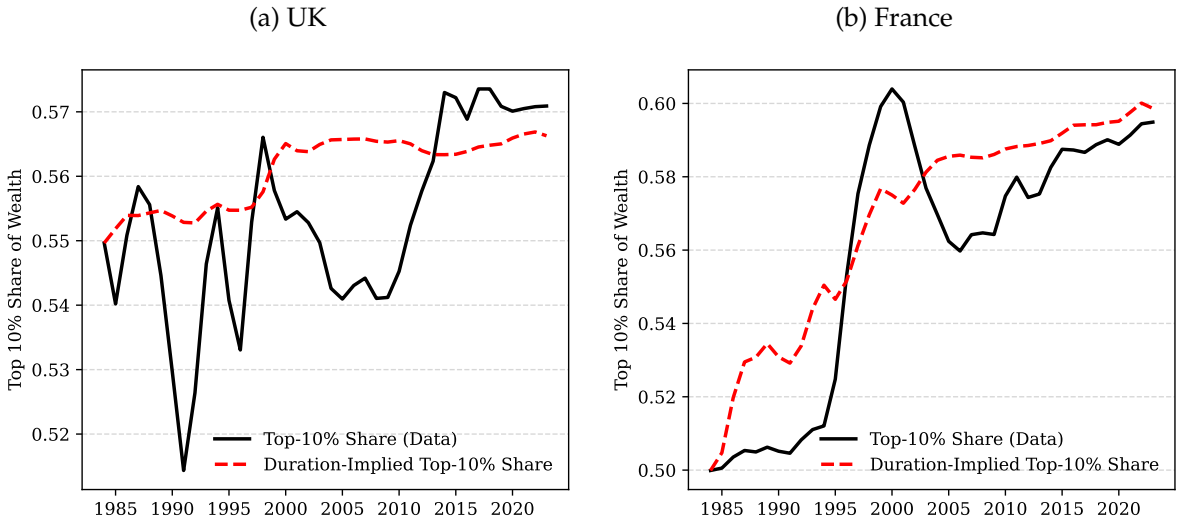

Note: The figure compares the observed top-10% wealth share (black line) with the duration-implied series. The repricing simulation compounds annual changes in interest rates starting from the observed 1983 level to 2023. We use duration and interest rate changes under the permanent case ($\phi = 1$). Panel 8a shows the results for the UK while Panel 8b shows the results for France.

estimate $\tau = 3\%$ used previously for the US for France and the UK. For mortgage debt, we use a duration of 0.83 years for the UK. The UK is a predominantly adjustable-rate mortgage (ARM) country. In contrast, in France, the typical mortgage is a fixed-rate mortgage that fully amortizes over 20 years. Given the history of mortgage rates, an average mortgage duration of 5.9 years is appropriate for France. For the other asset and liability classes, which are quantitatively much smaller, we use the same duration estimates as for the US.

A key observation from Table 4 is that the value- and equally-weighted durations are much closer together in the UK and France than in the US, with the value-weighted average duration of the top 10% only modestly higher than that of the bottom 90%. This is due to the dominance of real estate in the portfolio of all households, including the very wealthy. As a result, there is a smaller gap in equity allocations between the bottom 90% and the top 10%, as well as a smaller difference in duration between equities and real estate. A decline in real interest rates similar in magnitude to that in the US, will therefore have a more modest impact on wealth inequality than it did in the US.

Using the portfolio durations of the bottom 90% and the top 10%, we apply (12) to compute the change in the top-10% wealth share explained by the decline in interest rates over the period 1983–2023. For both the UK and France, we use the 10-year real log-interest rate (see Appendix H for

then use the rental price index and the property price index, to extend the rent-price ratio backwards in time to 1983.

details). In the UK, the real yield fell from 5.2% in 1983 to 0.5% in 2023, a decline of 4.7pp. In France, the real interest rate declined from 7.9% to 0.5%, a larger reduction of 7.4pp. For comparison, the change in real rates in the US was -4.08pp, similar to the UK but smaller than in France.

Figure 8 shows the top-10% wealth share implied by our revaluation mechanism alongside the observed top wealth share.[44] Panel 8a displays results for the UK, where wealth inequality only rises by 2.1pp, similar to what our revaluation predicts. This modest increase occurs despite the sizable decline in UK real rates because the duration gap between the top and bottom of the UK wealth distribution is small. Panel 8b displays results for France, where the top-10% wealth share increases by 9.5pp. This rise is also broadly consistent with our revaluation-based estimate, which predicts an increase of 9.9pp. The large rise in the top-10% wealth share results from the combination of a large duration gap and a large decline in real rates.

We conclude that declines in long-term real interest rates, together with heterogeneity in household portfolio durations, play an important role in explaining changes in top wealth inequality in both the UK and France. Our results also show that the strength of this mechanism varies substantially across countries. The duration gap between the top 10% and the bottom 90% is largest in the US, intermediate in France, and smaller in the UK. This cross-country variation is a key determinant of differences in the evolution of wealth inequality over the past four decades. For France, the large decline in real rates amplified the modest duration gap to deliver a large rise in top-wealth inequality.

8 Conclusion

A persistent decline in real interest rates, like the one experienced in much of the world between the 1980s and the 2010s, leads to a rise in wealth inequality when there is a positive covariance between wealth levels and the duration of wealth across households. Using detailed portfolio data, we show that this condition is met in the US data, and that the duration heterogeneity is large enough to account for most of the rise in the top-10% share of wealth since 1983. The same mechanism also helps account for the even larger rise in top-wealth inequality in France and the modest rise in the United Kingdom.

Recently, long-term real rates have begun to rise after a 40-year decline. The 10-year real bond yield increased by 2.0pp between March 2021 and March 2023 and is 2.2pp higher as of September 2026. Our revaluation mechanism predicts that this sharp rise in real rates lowers top-wealth inequality, a prediction borne out in recent inequality data. In the longer term, our model points to

[44] Since the WID data on the top-10% wealth share are more choppy for France and the UK than for the US, we compute a two-year average in both data and model.

even higher levels of top-wealth inequality, driven mostly by rising income inequality rather than expected asset revaluations.

References

Alvaredo, Facundo, Anthony B. Atkinson, Thomas Blanchet, Lucas Chancel, Luis Bauluz, Max Fisher-Post, Isabel Flores, Bertrand Garbinti, Jonathan Goupille-Lebret, Clara Martinez-Toledano, Marc Morgan, Tim Neef, Thomas Piketty, Anne-Sophie Robilliard, Emmanuel Saez, Li Yang, and Gabriel Zucman. 2020. "Distributional National Accounts Guidelines 2020: Concepts and Methods Used in the World Inequality Database." WID.world Working Paper.

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

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

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

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

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

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

Boar, Corina and Virgiliu Midrigan. 2022. "Efficient redistribution." *Journal of Monetary Economics* 131:78–91.

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

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

Catherine, Sylvain, Max Miller, James D. Paron, and Natasha Sarin. 2023. "Interest-rate Risk and Household Portfolios." *Jacobs Levy Equity Management Center for Quantitative Financial Research Paper* .

Catherine, Sylvain, Max Miller, and Natasha Sarin. 2025. "Social security and trends in wealth inequality." *The Journal of Finance* 80 (3):1497–1531.

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

Damen, Sven, Matthijs Korevaar, and Stijn Van Nieuwerburgh. 2025. "An Alpha in Affordable Housing?" Tech. rep., National Bureau of Economic Research.

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

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

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

Eggertsson, Gauti B, Neil R Mehrotra, and Jacob A Robbins. 2019. "A model of secular stagnation: Theory and quantitative evaluation." *American Economic Journal: Macroeconomics* 11 (1):1–48.

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

Fagereng, Andreas, Martin Blomhoff Holm, Benjamin Moll, and Gisle Natvik. 2019. "Saving Behavior Across the Wealth Distribution: The Importance of Capital Gains." NBER Working Paper No. 26588.

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

Garbinti, Bertrand, Jonathan Goupille-Lebret, and Thomas Piketty. 2021. "Accounting for Wealth-Inequality Dynamics: Methods, Estimates and Simulations for France." *Journal of the European Economic Association* 19 (1):620–663.

Giglio, Stefano, Matteo Maggiori, Krishna Rao, Johannes Stroebel, and Andreas Weber. 2021. "Climate Change and Long-Run Discount Rates: Evidence from Real Estate." *The Review of Financial Studies* 34 (8):3527–3571. URL https://doi.org/10.1093/rfs/hhab032.

Glaeser, Edward L, Joshua D Gottlieb, and Joseph Gyourko. 2012. "Can cheap credit explain the housing boom?" In *Housing and the financial crisis*. University of Chicago Press, 301–359.

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

Gordon, Robert J. 2017. *The Rise and Fall of American Growth: The US Standard of Living Since the Civil War*, vol. 70. Princeton University Press.

Gormsen, Niels Joachim and Eben Lazarus. 2025. "Interest rates and equity valuations." Working paper, University of Chicago.

Greenwald, Daniel, Matteo Leombroni, Hanno Lustig, and Stijn Van Nieuwerburgh. 2025. "Winners and Losers When Interest Rates Change." Tech. rep., Working Paper.

Greenwald, Daniel L, Martin Lettau, and Sydney C Ludvigson. 2025. "How the wealth was won: Factor shares as market fundamentals." *Journal of Political Economy* 133 (4):1083–1132.

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

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

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2020. "How Should Tax Progressivity Respond to Rising Income Inequality." NBER Working Paper No. 28006.

Hubmer, Joachim, Per Krusell, and Anthony A Smith Jr. 2021. "Sources of US wealth inequality: Past, present, and future." *NBER Macroeconomics Annual* 35 (1):391–455.

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

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

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

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

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

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

Kuhn, Moritz, Moritz Schularick, and Ulrike I Steins. 2020. "Income and Wealth Inequality in America, 1949–2016." *Journal of Political Economy* 128 (9):3469–3519.

Martínez-Toledano, Clara. 2020. "House Price Cycles, Wealth Inequality and Portfolio Reshuffling." WID.world Working Paper.

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

Muth, John F. 1961. "Rational Expectations and the Theory of Price Movements." *Econometrica* 29 (3):315–335.

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

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

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

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

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

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

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

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

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

Sims, Christopher. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50(3):665–90.

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

———. 2022. "The Rise of Pass-Throughs and the Decline of the Labor Share." *American Economic Review: Insights* 4 (3):323–340.

Sodini, Paolo, Stijn Van Nieuwerburgh, Roine Vestman, and Ulf von Lilienfeld-Toal. 2023. "Identifying the Benefits From Homeownership: A Swedish Experiment." *American Economic Review* 113 (12):3173–3212.

Stachurski, John and Alexis Akira Toda. 2019. "An impossibility theorem for wealth in heterogeneous-agent models with limited heterogeneity." *Journal of Economic Theory* 182:1–24.

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

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

Online Appendix

A Proofs for Section 1

In this section we provide the proofs for the main theoretical insights discussed in Section 1.

A.1 Proof of Proposition 1

Proof. **Part (a).** Following the shock ε, it is straightforward to verify that $\tilde{z}_t = z_t + \phi^t \varepsilon$, and so the expected path of interest rates becomes

$$\log \tilde{R}_t = \log R_t + \phi^t \varepsilon. \tag{A.1}$$

The post-shock asset price can now be written as

$$\begin{aligned} \tilde{P}_0 &= \sum_{t=1}^{\infty} \left(\prod_{j=0}^{t-1} \tilde{R}_j^{-1} \right) x_t \\ &= \sum_{t=1}^{\infty} \exp \left\{ - \sum_{j=0}^{t-1} \log \tilde{R}_j \right\} x_t \\ &= \sum_{t=1}^{\infty} \exp \left\{ - \sum_{j=0}^{t-1} \left(\log R_j + \phi^j \varepsilon \right) \right\} x_t \end{aligned} \tag{A.2}$$

Since $\tilde{P}_0$ is a differentiable function of ε, we can apply a first-order Taylor expansion around $\varepsilon = 0$

$$\log \tilde{P}_0 \simeq \log P_0 + \frac{1}{P_0} \times \left. \frac{d\tilde{P}_0}{d\varepsilon} \right|_{\varepsilon=0} \times \varepsilon. \tag{A.3}$$

Taking the derivative of (A.2), we obtain

$$\frac{d\tilde{P}_0}{d\varepsilon} = - \sum_{t=1}^{\infty} \exp \left\{ - \sum_{j=0}^{t-1} \left(\log R_j + \phi^j \varepsilon \right) \right\} \left(\sum_{j=0}^{t-1} \phi^j \right) x_t. \tag{A.4}$$

Evaluating this expression at $\varepsilon = 0$ yields

$$\begin{aligned} \left. \frac{d\tilde{P}_0}{d\varepsilon} \right|_{\varepsilon=0} &= - \sum_{t=1}^{\infty} \exp \left\{ - \sum_{j=0}^{t-1} \log R_j \right\} \left(\sum_{j=0}^{t-1} \phi^j \right) x_t \\ &= - \sum_{t=1}^{\infty} \left(\prod_{j=0}^{t-1} R_j^{-1} \right) \left(\sum_{j=0}^{t-1} \phi^j \right) x_t. \end{aligned} \tag{A.5}$$

Substituting into (A.3) now delivers

$$\log \tilde{P}_0 - \log P_0 \simeq -\frac{1}{P_0}\left\{\sum_{t=1}^{\infty}\left(\prod_{j=0}^{t-1} R_j^{-1}\right)\left(\sum_{j=0}^{t-1}\phi^j\right)x_t\right\} \times \varepsilon = -\mathcal{D}(\phi) \times \varepsilon \tag{A.6}$$

where generalized duration $\mathcal{D}(\phi)$ is defined as

$$\mathcal{D}(\phi) = \frac{1}{P_0}\sum_{t=1}^{\infty}\left(\sum_{j=0}^{t-1}\phi^j\right)\left(\prod_{j=0}^{t-1} R_j^{-1}\right)x_t. \tag{A.7}$$

Part (b). In the special case $\phi = 1$, we have $\sum_{j=0}^{t-1}\phi^j = \sum_{j=0}^{t-1} 1 = t$, and so (A.7) becomes

$$\mathcal{D}(1) = \frac{1}{P_0}\sum_{t=1}^{\infty} t \times \left(\prod_{j=0}^{t-1} R_j^{-1}\right)x_t \tag{A.8}$$

recovering the traditional formula.

Part (c). Let the portfolio consist of assets indexed by k, each with price $P_0^{(k)}$ and generalized duration $\mathcal{D}_k$. The total portfolio value is:

$$P_0^{\text{port}} = \sum_k P_0^{(k)}, \tag{A.9}$$

and after the shock:

$$\tilde{P}_0^{\text{port}} = \sum_k \tilde{P}_0^{(k)}. \tag{A.10}$$

From part (a), we know that for each asset k, the log change in its price is approximately:

$$\log \tilde{P}_0^{(k)} - \log P_0^{(k)} \simeq -\mathcal{D}_k(\phi) \cdot \varepsilon. \tag{A.11}$$

To approximate the change in the log value of the portfolio, we use a first-order log-linearization:

$$\log \tilde{P}_0^{\text{port}} \simeq \log P_0^{\text{port}} + \frac{1}{P_0^{\text{port}}}\sum_k \frac{d\tilde{P}_0^{(k)}}{d\varepsilon} \cdot \varepsilon. \tag{A.12}$$

Now, from part (a), we also have:

$$\frac{d\tilde{P}_0^{(k)}}{d\varepsilon} \simeq -\mathcal{D}_k(\phi) \cdot P_0^{(k)}. \tag{A.13}$$

Substituting in:

$$\log \tilde{P}_0^{\text{port}} \simeq \log P_0^{\text{port}} - \frac{1}{P_0^{\text{port}}} \sum_k \mathcal{D}_k(\phi) \cdot P_0^{(k)} \cdot \varepsilon. \tag{A.14}$$

Rewriting the weighted sum using portfolio value weights:

$$\omega(k) \equiv \frac{P_0^{(k)}}{P_0^{\text{port}}}, \tag{A.15}$$

we have:

$$\log \tilde{P}_0^{\text{port}} - \log P_0^{\text{port}} \simeq -\sum_k \omega(k) \mathcal{D}_k(\phi) \cdot \varepsilon. \tag{A.16}$$

Defining the value-weighted generalized duration as:

$$\mathcal{D}^{VW}(\phi) \equiv \sum_k \omega(k) \mathcal{D}_k(\phi), \tag{A.17}$$

we obtain the final result:

$$\log \tilde{P}_0^{\text{port}} - \log P_0^{\text{port}} \simeq -\mathcal{D}^{VW}(\phi) \cdot \varepsilon. \tag{A.18}$$

□

A.2 Proof of Proposition 2

Proof. Under the assumptions of the proposition, the interest rate is constant at $R_j = R, \ \forall j$, and cash flows grow deterministically at rate G, so $x_t = G^t x_0$. The generalized duration from Proposition 1 becomes

$$\mathcal{D}(\phi) = \frac{1}{P_0} \sum_{t=1}^{\infty} \left(\sum_{j=0}^{t-1} \phi^j \right) R^{-t} G^t x_0. \tag{A.19}$$

Define $\rho \equiv R/G$, so that $R^{-t} G^t = \rho^{-t}$. Then

$$\mathcal{D}(\phi) = \frac{x_0}{P_0} \sum_{t=1}^{\infty} \left(\sum_{j=0}^{t-1} \phi^j \right) \rho^{-t}. \tag{A.20}$$

Assuming $\phi < 1$, we can apply the partial geometric sum formula $\sum_{j=0}^{t-1} \phi^j = \frac{1-\phi^t}{1-\phi}$ to obtain

$$\mathcal{D}(\phi) = \frac{x_0}{P_0(1-\phi)} \sum_{t=1}^{\infty} (1 - \phi^t) \rho^{-t}. \tag{A.21}$$

We can split this sum to yield

$$\sum_{t=1}^{\infty}(1-\phi^t)\rho^{-t} = \sum_{t=1}^{\infty}\rho^{-t} - \sum_{t=1}^{\infty}(\phi/\rho)^t = \frac{1}{\rho-1} - \frac{\phi}{\rho-\phi}. \tag{A.22}$$

Substituting this into the generalized duration formula now delivers

$$\mathcal{D}(\phi) = \frac{x_0}{P_0(1-\phi)}\left(\frac{1}{\rho-1} - \frac{\phi}{\rho-\phi}\right) = \left(\frac{x_0}{P_0}\right)\frac{\rho}{(\rho-1)(\rho-\phi)}. \tag{A.23}$$

From the asset pricing formula for a growing perpetuity, we have

$$P_0 = \frac{Gx_0}{R-G} = \frac{x_0}{\rho-1} \tag{A.24}$$

which implies

$$\frac{x_0}{P_0} = \rho - 1. \tag{A.25}$$

Substituting from (A.25) into (A.23), we obtain

$$\mathcal{D}(\phi) = (\rho-1)\cdot\frac{\rho}{(\rho-1)(\rho-\phi)} = \frac{\rho}{\rho-\phi}, \tag{A.26}$$

which derives equation (8). Further, we can rearrange (A.25) to yield

$$\rho = \frac{(P_0/x_0)+1}{(P_0/x_0)}. \tag{A.27}$$

Substituting from (A.27) into (A.26), we obtain

$$\mathcal{D}(\phi) = \frac{(P_0/x_0)+1}{(1-\phi)(P_0/x_0)+1}, \tag{A.28}$$

deriving (8).

For the case $\phi = 1$, we must replace our partial geometric sum formula with $\sum_{j=0}^{t-1} 1 = t$. In this case, (A.20) becomes

$$\mathcal{D}(\phi) = \frac{x_0}{P_0}\sum_{t=1}^{\infty} t \times \rho^{-t}.$$

To compute the sum term, note that if we define

$$G(a) = \sum_{t=0}^{\infty} a^t = \frac{1}{1-a} \tag{A.29}$$

then

$$G'(a) = \sum_{t=0}^{\infty} t \times a^{t-1} = \sum_{t=1}^{\infty} t \times a^{t-1} = \frac{1}{(1-a)^2} \tag{A.30}$$

where the second equality follows from the fact that $t \times a^{t-1} = 0$ for $t = 0$. Multiplying by a we obtain

$$aG'(a) = \sum_{t=1}^{\infty} t \times a^{t} = \frac{a}{(1-a)^2}. \tag{A.31}$$

Applying $a = \rho^{-1}$ and substituting now yields

$$\mathcal{D}(\phi) = \left(\frac{x_0}{P_0}\right) \frac{\rho^{-1}}{(1-\rho^{-1})^2} = \left(\frac{x_0}{P_0}\right) \frac{\rho}{(\rho-1)^2}.$$

Applying (A.25), we can substitute for (x_0/P_0) to obtain

$$\mathcal{D}(\phi) = \frac{\rho}{\rho - 1}, \tag{A.32}$$

We can similarly derive (8) for the case $\phi = 1$, exactly as above. Last, we can substitute $\rho = R/G$ into (A.32) to obtain

$$\mathcal{D}(\phi) = \frac{R/G}{R/G - 1} = \frac{R}{R - G} = \frac{1+r}{r-g},$$

which derives (9). □

A.3 Proof of Proposition 3

For this proof, we first derive a formula for the general case with arbitrary cash flows, and then specialize to the Gordon Growth environment considered in the proposition. To begin, consider the general case

$$\begin{aligned}
\tilde{P}_0 &= \sum_{t=1}^{\infty} \left(\prod_{j=0}^{t-1} \tilde{R}_j^{-1} \right) (x_t - \tau \tilde{P}_t) \\
&= \sum_{t=1}^{\infty} \exp\left\{ - \sum_{j=0}^{t-1} \log \tilde{R}_j \right\} (x_t - \tau \tilde{P}_t) \\
&= \sum_{t=1}^{\infty} \exp\left\{ - \sum_{j=0}^{t-1} \left(\log R_j + \phi^j \varepsilon \right) \right\} (x_t - \tau \tilde{P}_t) .
\end{aligned}$$

Then

$$\frac{d\tilde{P}_0}{d\varepsilon} = -\sum_{t=1}^{\infty} \exp\left\{-\sum_{j=0}^{t-1}\left(\log R_j + \phi^j \varepsilon\right)\right\}\left(\sum_{j=0}^{t-1}\phi^j\right)\left(x_t - \tau\tilde{P}_t\right) - \sum_{t=1}^{\infty} \exp\left\{-\sum_{j=0}^{t-1}\left(\log R_j + \phi^j \varepsilon\right)\right\} \times \tau \frac{d\tilde{P}_t}{d\varepsilon}.$$

Evaluating this expression at $\varepsilon = 0$ yields

$$\begin{aligned}\left.\frac{d\tilde{P}_0}{d\varepsilon}\right|_{\varepsilon=0} &= -\sum_{t=1}^{\infty} \exp\left\{-\sum_{j=0}^{t-1}\log R_j\right\}\left(\sum_{j=0}^{t-1}\phi^j\right)\left(x_t - \tau P_t\right) - \sum_{t=1}^{\infty} \exp\left\{-\sum_{j=0}^{t-1}\log R_j\right\} \times \tau \left.\frac{d\tilde{P}_t}{d\varepsilon}\right|_{\varepsilon=0} \\ &= -\sum_{t=1}^{\infty}\left(\prod_{j=0}^{t-1} R_j^{-1}\right)\left(\sum_{j=0}^{t-1}\phi^j\right)\left(x_t - \tau P_t\right) - \sum_{t=1}^{\infty}\left(\prod_{j=0}^{t-1} R_j^{-1}\right) \times \tau \left.\frac{d\tilde{P}_t}{d\varepsilon}\right|_{\varepsilon=0}.\end{aligned}$$

If we define the generalized duration of the asset to be

$$\mathcal{D}(\phi, \tau) = -\left.\frac{d\log\tilde{P}_0}{d\varepsilon}\right|_{\varepsilon=0} = -\left(\frac{1}{P_0}\right)\left.\frac{d\tilde{P}_0}{d\varepsilon}\right|_{\varepsilon=0} \tag{A.33}$$

then we obtain

$$\mathcal{D}(\phi, \tau) \simeq \frac{1}{P_0}\sum_{t=1}^{\infty}\left(\prod_{j=0}^{t-1} R_j^{-1}\right)\left(\sum_{j=0}^{t-1}\phi^j\right)\left(x_t - \tau P_t\right) + \frac{1}{P_0}\sum_{t=1}^{\infty}\left(\prod_{j=0}^{t-1} R_j^{-1}\right) \times \tau \left.\frac{d\tilde{P}_t}{d\varepsilon}\right|_{\varepsilon=0}. \tag{A.34}$$

To unpack this expression, let us define some of its subcomponents. From the first term on the right hand side, we can define

$$\hat{\mathcal{D}}(\phi, \tau) \equiv \frac{1}{P_0}\sum_{t=1}^{\infty}\left(\prod_{j=0}^{t-1} R_j^{-1}\right)\left(\sum_{j=0}^{t-1}\phi^j\right)\left(x_t - \tau P_t\right) \tag{A.35}$$

to be the generalized duration of a hypothetical asset that paid *fixed* cash flows equal to the initial values of $x_t - \tau P_t$. The hat notation indicates that this duration measure ignores the fact that P_t and hence the tax portion of the cash flow will in practice vary with discount rates, and therefore does not reflect the true total interest rate sensitivity of the asset. If we define

$$T_t = \left(\prod_{j=0}^{t-1} R_j^{-1}\right)\tau P_t$$

to be the present value of the tax payments at date t, then we can substitute into (A.34) to obtain

$$\mathcal{D}(\phi,\tau) \simeq \hat{\mathcal{D}}(\phi,\tau) + \frac{1}{P_0}\sum_{t=1}^{\infty}\left\{T_t \times \frac{1}{P_t}\left.\frac{d\tilde{P}_t}{d\varepsilon}\right|_{\varepsilon=0}\right\} = \hat{\mathcal{D}}(\phi,\tau) + \frac{1}{P_0}\sum_{t=1}^{\infty}\left\{T_t \times \left.\frac{d\log\tilde{P}_t}{d\varepsilon}\right|_{\varepsilon=0}\right\}.$$

Assume we are in a stationary environment so that

$$\left.\frac{d\log\tilde{P}_t}{d\varepsilon}\right|_{\varepsilon=0} = \phi^t \left.\frac{d\log\tilde{P}_0}{d\varepsilon}\right|_{\varepsilon=0} = -\phi^t\mathcal{D}(\phi,\tau)$$

or alternatively

$$\left.\frac{d\log\tilde{P}_t}{d\varepsilon_t}\right|_{\varepsilon_t=0} = \left.\frac{d\log\tilde{P}_0}{d\varepsilon}\right|_{\varepsilon=0} = -\mathcal{D}(\phi,\tau)$$

for $\varepsilon_t = \phi^t\varepsilon$. In words, this assumption means that the price of the asset is expected to be equally sensitive to a contemporaneous change in the interest rate of the same persistence at all times. In this case, we obtain

$$\mathcal{D}(\phi,\tau) \simeq \hat{\mathcal{D}}(\phi,\tau) - \frac{1}{P_0}\sum_{t=1}^{\infty}\phi^t T_t\mathcal{D}(\phi,\tau).$$

Solving for $\mathcal{D}(\phi,\tau)$ now yields

$$\mathcal{D}(\phi,\tau) = \left[1 + \sum_{t=1}^{\infty}\phi^t\left(\frac{T_t}{P_0}\right)\right]^{-1}\hat{\mathcal{D}}(\phi,\tau). \tag{A.36}$$

Special Case: Gordon Growth. In the special case that interest rates are initially expected to be constant ($R_t = R_0, \forall t$) and cash flow growth is expected to be constant ($x_t = G^t x_0$), we can guess and verify that $P_t = G^t P_0$. In this case we obtain

$$\begin{aligned} P_0 &= \sum_{t=1}^{\infty}(G/R)^t(x_0 - \tau P_0) = \sum_{t=1}^{\infty}\rho^{-t}(x_0 - \tau P_0) = \left(\frac{\rho^{-1}}{1-\rho^{-1}}\right)(x_0 - \tau P_0) \\ &= \left(\frac{1}{\rho - 1}\right)(x_0 - \tau P_0). \end{aligned}$$

Solving for P_0 now yields

$$P_0 = \frac{x_0}{\rho - (1-\tau)}.$$

A symmetric argument would show that $P_t = x_t/(\rho - (1-\tau))$, and since $x_t = G^t x_0$, this verifies the guess that $P_t = G^t P_0$. With our guess verified, we can compute the present value of tax

payments as

$$T_t = R^{-t}\tau P_t = \tau R^{-t} G^t P_0 = \tau \rho^{-t} P_0$$

and so

$$\frac{T_t}{P_0} = \tau \rho^{-t}.$$

This in turn implies

$$\sum_{t=1}^{\infty} \phi^t \left(\frac{T_t}{P_0}\right) = \tau \sum_{t=1}^{\infty} (\phi/\rho)^t = \tau \left(\frac{(\phi/\rho)}{1-(\phi/\rho)}\right) = \left(\frac{\tau\phi}{\rho-\phi}\right).$$

Substituting into (A.36) now yields

$$\begin{aligned}
\mathcal{D}(\phi,\tau) &= \left[1 + \left(\frac{\tau\phi}{\rho-\phi}\right)\right]^{-1} \hat{\mathcal{D}}(\phi,\tau) \\
&= \left[\left(\frac{\rho-(1-\tau)\phi}{\rho-\phi}\right)\right]^{-1} \hat{\mathcal{D}}(\phi,\tau) \\
&= \left(\frac{\rho-\phi}{\rho-(1-\tau)\phi}\right) \hat{\mathcal{D}}(\phi,\tau) \\
&= \left(\frac{\rho-\phi}{\rho-(1-\tau)\phi}\right)\left(\frac{\rho}{\rho-\phi}\right) \\
&= \frac{\rho}{\rho-(1-\tau)\phi},
\end{aligned}$$

which implies that, as a growing perpetuity, we have $\hat{\mathcal{D}}(\phi,\tau) = \rho/(\rho-\phi)$.

Since

$$\rho = \frac{x_0}{P_0} + (1-\tau),$$

we can substitute to obtain

$$\mathcal{D}(\phi,\tau) = \frac{(1-\tau)(P_0/x_0)+1}{(1-\phi)(1-\tau)(P_0/x_0)+1} \tag{A.37}$$

completing the derivation of (11).

A.4 Proof of Proposition 4

Proof. **Part (a).** From Proposition 1, household i's financial wealth W^i changes with a shock ε to the expected path of interest rates according to:

$$\log \tilde{W}^i - \log W^i \simeq -\mathcal{D}^i \cdot \varepsilon,$$

where $\mathcal{D}^i = \mathcal{D}^i(\phi)$ denotes the generalized duration of household i's financial wealth. Exponentiating both sides and applying a first-order Taylor approximation of the exponential around $\varepsilon = 0$ yields:

$$\frac{\tilde{W}^i}{W^i} \simeq \exp(-\mathcal{D}^i \cdot \varepsilon) \simeq 1 - \mathcal{D}^i \cdot \varepsilon.$$

Taking the covariance of both sides with respect to W^i, we obtain:

$$\text{Cov}\left(\frac{\tilde{W}^i}{W^i}, W^i\right) \simeq -\varepsilon \cdot \text{Cov}(\mathcal{D}^i, W^i).$$

When $\varepsilon < 0$, the right-hand side is positive if and only if $\text{Cov}(\mathcal{D}^i, W^i) > 0$. That is, revaluation effects from declining interest rates amplify wealth differences if and only if the generalized duration is positively correlated with the level of financial wealth.

To derive this condition, let $\bar{W} = \int W_i \, di$ denote total (aggregate) financial wealth. Then the value-weighted average duration is:

$$\mathcal{D}^{VW} = \int \frac{W^i}{\bar{W}} \mathcal{D}^i \, di.$$

We decompose this as:

$$\begin{aligned}
\mathcal{D}^{VW} &= \bar{W}^{-1} \int W^i \mathcal{D}^i \, di \\
&= \bar{W}^{-1} \left\{ \int W^i \, di \cdot \int \mathcal{D}^i \, di + \int \left(W^i - \int W^j \, dj \right) \left(\mathcal{D}^i - \int \mathcal{D}^j \, dj \right) \, di \right\} \\
&= \mathcal{D}^{EW} + \bar{W}^{-1} \text{Cov}(W^i, \mathcal{D}^i),
\end{aligned}$$

where $\mathcal{D}^{EW} = \int \mathcal{D}^i \, di$ is the equal-weighted average generalized duration. This implies:

$$\mathcal{D}^{VW} > \mathcal{D}^{EW} \quad \text{if and only if} \quad \text{Cov}(W^i, \mathcal{D}^i) > 0.$$

Hence, a necessary and sufficient condition for declining interest rates to lead to proportional wealth gains that covary positively with wealth is that value-weighted generalized duration ex-

ceeds equal-weighted generalized duration.

Part (b). Let $\bar{W}^{top}$ denote the total financial wealth of the top-α share of households, and $\bar{W}^{bottom}$ denote the total financial wealth of the bottom $1-\alpha$ share. Then the top share of financial wealth is:

$$S^{\alpha} = \frac{\bar{W}^{top}}{\bar{W}^{top} + \bar{W}^{bottom}} = \frac{1}{1 + \bar{W}^{bottom}/\bar{W}^{top}} = \frac{1}{1 + \exp(\log \bar{W}^{bottom} - \log \bar{W}^{top})}. \tag{A.38}$$

To analyze how S^{α} responds to an interest rate shock ε, we differentiate with respect to ε:

$$\frac{\partial S^{\alpha}}{\partial \varepsilon} = -\frac{\exp(\log \bar{W}^{bottom} - \log \bar{W}^{top})}{(1 + \exp(\log \bar{W}^{bottom} - \log \bar{W}^{top}))^2} \cdot \left(\frac{\partial \log \bar{W}^{bottom}}{\partial \varepsilon} - \frac{\partial \log \bar{W}^{top}}{\partial \varepsilon} \right). \tag{A.39}$$

From equation (A.38), it is straightforward to check that

$$\exp(\log \bar{W}^{bottom} - \log \bar{W}^{top}) = \frac{1 - S^{\alpha}}{S^{\alpha}},$$

which implies

$$-\frac{\exp(\log \bar{W}^{bottom} - \log \bar{W}^{top})}{(1 + \exp(\log \bar{W}^{bottom} - \log \bar{W}^{top}))^2} = -\frac{(1 - S^{\alpha})/S^{\alpha}}{(1 + (1 - S^{\alpha})/S^{\alpha})^2} = -S^{\alpha}(1 - S^{\alpha}).$$

Substituting back into (A.39), we obtain

$$\frac{\partial S^{\alpha}}{\partial \varepsilon} = -S^{\alpha}(1 - S^{\alpha}) \left(\frac{\partial \log \bar{W}^{bottom}}{\partial \varepsilon} - \frac{\partial \log \bar{W}^{top}}{\partial \varepsilon} \right).$$

Applying a first-order approximation, and using the fact that:

$$\frac{\partial \log \bar{W}^{G}}{\partial \varepsilon} \simeq -\mathcal{D}^{G}, \quad \text{for } G \in \{top, bottom\},$$

where $\mathcal{D}^{top}$ and $\mathcal{D}^{bottom}$ are the value-weighted generalized durations of the top-α and bottom $(1-\alpha)$ groups respectively, we obtain:

$$\frac{\partial S^{\alpha}}{\partial \varepsilon} \simeq -S^{\alpha}(1 - S^{\alpha})(\mathcal{D}^{top} - \mathcal{D}^{bottom}).$$

Thus, for $\varepsilon < 0$, the top share increases if and only if $\mathcal{D}^{top} > \mathcal{D}^{bottom}$. Since overall aggregate duration $\mathcal{D}^{VW}$ is a weighted average of $\mathcal{D}^{top}$ and $\mathcal{D}^{bottom}$, the same condition holds if and only if $\mathcal{D}^{top} > \mathcal{D}^{VW}$. □

A.5 Proof of Proposition 5

Proof. **Part (a).** The asset price after the joint shock to the discount and growth rates is:

$$\begin{aligned}\tilde{P}_0 &= \sum_{t=1}^{\infty}\left(\prod_{j=0}^{t-1}\tilde{R}_j^{-1}\right)\left(\prod_{j=0}^{t-1}\tilde{G}_j\right)x_0 \\ &= \sum_{t=1}^{\infty}\left(\prod_{j=0}^{t-1}R_j^{-1}G_j\right)\exp\left(-\sum_{j=0}^{t-1}\phi^j\varepsilon+(1-\theta)\sum_{j=0}^{t-1}\phi^j\varepsilon_g\right)x_0, \end{aligned} \tag{A.40}$$

where

$$\tilde{R}_j = R_j\exp\left(\phi^j\varepsilon+\theta\phi^j\varepsilon_g\right), \qquad \tilde{G}_j = G_j\exp\left(\phi^j\varepsilon_g\right).$$

Taking a first-order Taylor expansion of $\log\tilde{P}_0$ around $\varepsilon=0,\varepsilon_g=0$, we obtain:

$$\log\tilde{P}_0-\log P_0 \simeq \frac{1}{P_0}\frac{d\tilde{P}_0}{d\varepsilon}\bigg|_{\varepsilon=0,\varepsilon_g=0}\cdot\varepsilon+\frac{1}{P_0}\frac{d\tilde{P}_0}{d\varepsilon_g}\bigg|_{\varepsilon=0,\varepsilon_g=0}\cdot\varepsilon_g.$$

We already computed the sensitivity to ε in Proposition 1, which yields:

$$\frac{1}{P_0}\frac{d\tilde{P}_0}{d\varepsilon}\bigg|_{\varepsilon=0,\varepsilon_g=0} = -\frac{1}{P_0}\sum_{t=1}^{\infty}\left(\sum_{j=0}^{t-1}\phi^j\right)\left(\prod_{j=0}^{t-1}R_j^{-1}\right)x_t = -\mathcal{D}(\phi),$$

where $x_t = x_0\prod_{j=0}^{t-1}G_j$.

We now compute the effect of the growth shock ε_g. From (A.40), it is clear that the effect of ε_g is identical to the effect of ε, except for the scaling by $1-\theta$. It is straightforward to verify that

$$\frac{1}{P_0}\frac{d\tilde{P}_0}{d\varepsilon_g}\bigg|_{\varepsilon=0,\varepsilon_g=0} = (1-\theta)\cdot\frac{1}{P_0}\sum_{t=1}^{\infty}\left(\sum_{j=0}^{t-1}\phi^j\right)\left(\prod_{j=0}^{t-1}R_j^{-1}\right)x_t = (1-\theta)\mathcal{D}(\phi).$$

Combining both effects, the first-order change in the asset price is:

$$\log\tilde{P}_0-\log P_0 \simeq (1-\theta)\mathcal{D}(\phi)\cdot\varepsilon_g-\mathcal{D}(\phi)\cdot\varepsilon.$$

Since the change in the initial discount rate is given by:

$$\log\tilde{R}_0-\log R_0 = \varepsilon+\theta\varepsilon_g,$$

we can substitute to obtain the desired result

$$\log\tilde{P}_0-\log P_0 \simeq \mathcal{D}(\phi)\cdot\varepsilon_g-\mathcal{D}(\phi)\cdot(\log\tilde{R}_0-\log R_0).$$

Part (b). Consider a portfolio of assets indexed by k, with total value $V_0 = \sum_k P_0^{(k)}$. Each asset has a generalized duration $\mathcal{D}_k$, and some assets have cash flows that grow with the process $\{G_t\}$, while others do not. The change in the log value of asset k is:

$$\log \tilde{P}_0^{(k)} - \log P_0^{(k)} \simeq \mathcal{D}_k \cdot \varepsilon_g \cdot \mathbf{1}_k - \mathcal{D}_k \cdot (\log \tilde{R}_0 - \log R_0),$$

where $\mathbf{1}_k = 1$ if asset k's cash flows grow with G_t, and $\mathbf{1}_k = 0$ otherwise.

Taking a value-weighted average across assets, with weights $\omega(k) = \frac{P_0^{(k)}}{V_0}$, the change in the value of the portfolio is:

$$\log \tilde{V}_0 - \log V_0 \simeq \left(\sum_k \omega(k) \mathcal{D}_k \mathbf{1}_k \right) \cdot \varepsilon_g - \left(\sum_k \omega(k) \mathcal{D}_k \right) \cdot (\log \tilde{R}_0 - \log R_0).$$

Define the value-weighted growth-sensitive duration:

$$\mathcal{D}_G^{VW} \equiv \sum_k \omega(k) \mathcal{D}_k \mathbf{1}_k, \quad \mathcal{D}^{VW} \equiv \sum_k \omega(k) \mathcal{D}_k.$$

This gives the desired portfolio-level result:

$$\log \tilde{V}_0 - \log V_0 \simeq \mathcal{D}_G^{VW} \cdot \varepsilon_g - \mathcal{D}^{VW} \cdot (\log \tilde{R}_0 - \log R_0).$$

Part (c). Let $\bar{W}^{top}$ and $\bar{W}^{bottom}$ denote the total financial wealth of the top-α and bottom-$(1-\alpha)$ groups of households, respectively. The top-share of financial wealth is

$$S^{\alpha} = \frac{\bar{W}^{top}}{\bar{W}^{top} + \bar{W}^{bottom}} = \frac{1}{1 + e^{z}}, \qquad \text{where } z := \log \bar{W}^{bottom} - \log \bar{W}^{top}.$$

We first approximate the total wealth of each group using

$$d \log \bar{W}^{k} \approx \mathcal{D}_G^{k} \, \varepsilon_g - \mathcal{D}^{k} \, (\log \tilde{R}_0 - \log R_0), \qquad k \in \{top, bottom\}.$$

Then:

$$dz = (\mathcal{D}_G^{bottom} - \mathcal{D}_G^{top}) \, \varepsilon_g - (\mathcal{D}^{bottom} - \mathcal{D}^{top}) (\log \tilde{R}_0 - \log R_0).$$

Since

$$dS^{\alpha} = -S^{\alpha}(1 - S^{\alpha}) \, dz,$$

it follows that

$$dS^{\alpha} \approx -S^{\alpha}(1-S^{\alpha})\Big[(\mathcal{D}_G^{bottom} - \mathcal{D}_G^{top})\,\varepsilon_g - (\mathcal{D}^{bottom} - \mathcal{D}^{top})(\log \tilde{R}_0 - \log R_0)\Big].$$

Rewriting in terms of top-bottom differences, we obtain

$$dS^{\alpha} \approx -S^{\alpha}(1-S^{\alpha})\Big[(\mathcal{D}^{top} - \mathcal{D}^{bottom})(\log \tilde{R}_0 - \log R_0) - (\mathcal{D}_G^{top} - \mathcal{D}_G^{bottom})\varepsilon_g\Big].$$

□

B Inequality and Interest Rate Data

B.1 Wealth Inequality

Our main source for top wealth shares in Figure 1 is the World Inequality Database (WID), maintained by the WID team. The data are annual and span 1950–2023. Figure B.1 reports the wealth shares held by the top 10% and the top 1%, together with the Gini coefficient for financial wealth. Each panel presents estimates based on both the WID and the Survey of Consumer Finances (SCF), which we complement with SCF+, the extended dataset developed by Kuhn, Schularick, and Steins (2020). Our wealth measure, constructed from SCF and SCF+, is described in Appendix C and excludes the category "Other non-financial assets."

Figure B.1: US Wealth Inequality

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

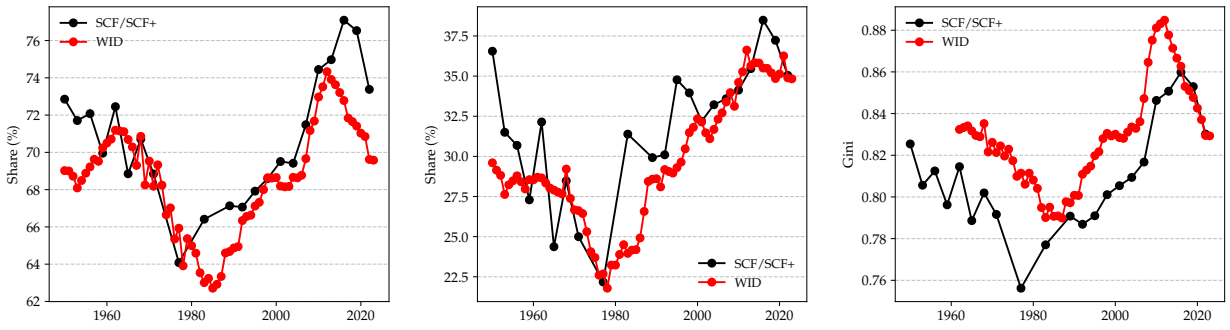

Note: Data are drawn from WID, SCF, and SCF+. Figure B.1a reports the wealth share held by the top 10%, Figure B.1b reports the share held by the top 1%, and Figure B.1c displays the Gini coefficient for wealth. The sample spans 1950–2022 for SCF/SCF+ and 1950–2023 for WID.

B.2 Interest Rates

Our primary series for the real interest rate is sourced from the Federal Reserve Bank of Cleveland, which provides estimates of the 10-year real rate beginning in 1982. For our motivating Figure 1, we extend the series backward. To do so, we regress the Cleveland Fed 10-year real yield series on the 10-year nominal yield (Federal Reserve Board of Governors, FRED code GS10) and the current year's inflation rate $\Delta \log CPI_t$, where CPI_t is the average of the CPI price index over the year (US Bureau of Labor Statistics, FRED code CPIAUCSL). Specifically, we run the regression

$$\log R^{10}_{real,t} = \beta_0 + \beta_1 \log R^{10}_{nom,t} + \beta_2 \Delta \log CPI_t + \varepsilon_t \tag{B.1}$$

where e.g., $R^{10}_{real,t}$ is the gross real interest rate. The regression achieves an R^2 of 99.3%, meaning a near-perfect fit. We then obtain the fitted log yield

$$\hat{r}^{10}_{real,t} = \hat{\beta}_0 + \hat{\beta}_1 \log R^{10}_{nom,t} + \hat{\beta}_2 \Delta \log CPI_t \tag{B.2}$$

Figure B.2 shows that this fitted value achieves a near perfect fit with the original series over the sample for which they are available. For Figure 1, we compute the implied 10-year real bond price using

$$q^{10}_{real,t} = \exp(-10 \times \hat{r}^{10}_{real,t}).$$

Figure B.2: 10-Year Real Yield, Actual vs. Proxy

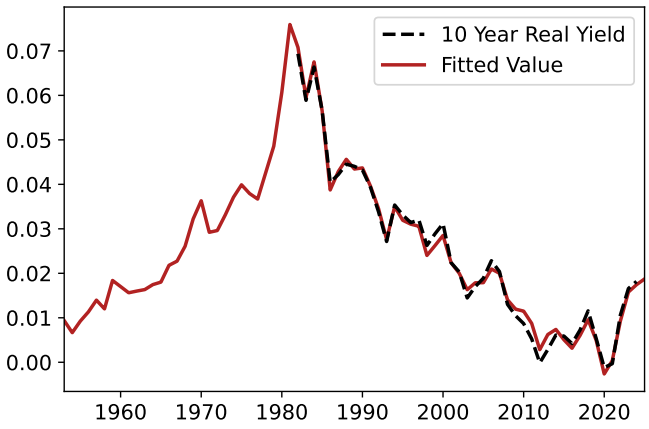

Note: This figure displays the Cleveland Fed 10-year real yield series alongside the fitted value $\hat{r}^{10}_{real,t}$ from (B.2), obtained by regressing the Cleveland Fed 10-year real yield against the 10-year nominal yield (Federal Reserve Board of Governors, FRED code GS10) and the current year's inflation rate $\Delta \log CPI_t$, where CPI_t is the average of the CPI price index over the year (US Bureau of Labor Statistics, FRED code CPIAUCSL).

B.3 Income Inequality

For our model extension with time-varying income inequality, we recalibrate the model to match the normalized level of income inequality in each year. We target the top-10% income share up to a constant of proportionality, so that the ratio of the share between any two dates in the model matches the data, while the level itself is not targeted (the superstar income state is calibrated to wealth inequality). For our measure of inequality, we use the top-10% income share in the WID, plotted in Figure B.3.

Figure B.3: Top-10% Income Share (WID)

Note: This figure displays the top-10% income share for the US. Source: World Inequality Database.

C Household Wealth and Portfolio Shares

In this section, we detail the variables used in calculating household wealth and portfolio shares. Our data encompasses the 1950–2022 time range, based on Survey of Consumer Finances (SCF) waves from 1989 through 2022 and data from the SCF+ database, compiled by Kuhn, Schularick, and Steins (2020), for the period from 1950 through 1983. In sections C.1 and C.2, we discuss the specific variables used in our analysis for these two sources.

C.1 Survey of Consumer Finances (SCF)

The SCF is a statistical survey of the balance sheet, pension, income and other demographic characteristics of households in the United States. We use data from the Summary Extract Data—that is, the extract data set of summary variables used in the Federal Reserve Bulletin[45]—as well as granular information from the SCF original surveys. We construct the following variables for our analysis. We use the SCF from 1989 to 2022 (the latest available survey).

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

[45]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

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

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

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

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

Cash & Deposits. SCF variables *LIQ* + *CDS*. This includes all types of transaction accounts (money market accounts, checking accounts, savings accounts, call accounts, prepaid cards) and certificates of deposit.

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

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

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

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

Other Debt. SCF variable *DEBT* minus Mortgage Debt, Student Debt, and Vehicle 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.

C.1.1 Private Business Wealth: Corporate vs. Non-Corporate

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

Corporate Private Business Wealth. It includes the value of all businesses reported as Subchapter-S corporations or all other types of corporations (both actively and passively managed). These categories are reported as Corporate Business in the Financial Accounts of the United States. Figure C.1a plots the shares of Corporate PBW as a percentage of total Private Business Wealth in 1989, the first SCF survey wave we use in our sample, along the wealth distribution. Figure C.1b plots the shares in 2022, the last available SCF survey wave.

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

We use the following methodology to construct the value of Corporate and Non-Corporate Private Business Wealth. From 1989 to 2010, the SCF provides information on the three largest active businesses, while from 2013 onward it provides information on the two largest active businesses. For each actively-managed business, we follow the methodology used by the SCF to measure the overall total value of private business wealth. We label the counter of the SCF variable used for each of the business where the information is available: #1 for the first business, #2 for the second business, #3 for the third business. For each active business, we measure its value as: the net worth of the share of the business (#1 X3129, #2 X3229, #3 X3329), plus the amount the business owes the household (#1 X3124, #2 X3224, #3 X3324), minus the amount the household owes the business (# 1 X3126, #2 X3226, #3 X3326), plus the amount of loans that are collateralized/guaranteed (#1 X3121, #2 X3221, #3 X3321). After having computed the value of each business, we then use information on the types of business entities and legal structure (#1 X3119, #2 X3219, #3 X3319) to split business wealth into a Corporate and a Non-Corporate component.

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

Figure C.1: Ratio of Corporate PBW to Private Business Wealth by Wealth Percentile Groups

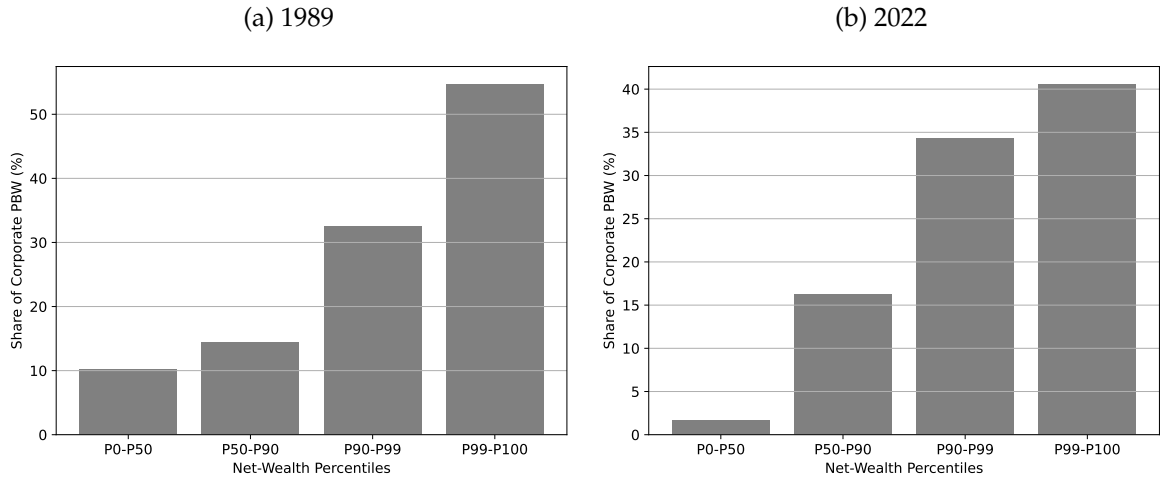

Note: The Figure plots the ratio of Corporate PBW to total Private Business Wealth for different wealth percentile groups in 1989 and 2022. Initially, we determine the division between Corporate PBW and Non-Corporate PBW for each individual household. Subsequently, we categorize households into four groups based on their wealth ranges (P0-P50, P50-P90, P90-P99, P99-P100). Finally, we calculate the total dollar amount of Corporate PBW and Private Business Wealth for all households within each wealth percentile group and derive the corresponding ratios. Figure C.1a plots the shares in 1989, the first SCF survey wave we use in our sample. Figure C.1b plots the shares in 2022, the last available SCF survey wave.

C.2 SCF+

We use SCF+ (Kuhn, Schularick, and Steins, 2020) to construct household portfolios in 1983. In this section, we outline the list of variables utilized from the SCF+ database to calculate households' wealth and portfolio shares.

Equities. SCF+ variable *ffaequ*. This includes equity and other managed assets. We also add indirect holdings of equities through mutual funds and pension funds (see Appendix C.2.1).

Real Estate. This includes: SCF+ variable *house*, asset value of house; SCF+ variable *oest*, other real estate (net position); SCF+ variable *oestdebt*, other real estate debt (note: we add back the debt to the other real estate net position).

Private Business Wealth. SCF+ variable *ffabus*, business wealth.

Vehicles. SCF+ variable *vehi*.

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

Fixed Income. SCF+ variable *ffafin*, financial assets minus Cash and Deposits and Equities.

Mortgage Debt. This includes: SCF+ variable *hdebt*, housing debt on owner-occupied real estate; SCF+ variable *oestdebt*, other real estate debt.

Vehicle Debt. Information on vehicle debt is not available from SCF+. We infer vehicle debt using the procedure described in Section C.2.1.

Student Debt. Information on student debt is not available from SCF+. We infer student debt using the procedure described in Section C.2.1.

Total Other Debt. *pdebt*, personal debt.

C.2.1 Adjustments to SCF+ data

Indirect Equity Holdings. SCF+ does not report indirect equity holdings. We therefore infer them using aggregate data from the US Financial Accounts. Indirect exposure arises through mutual fund shares and pension accounts.

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

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

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

Private Business Wealth Corporate and Non-Corporate. The SCF+ database provides only the total households' holdings of private business wealth, without distinguishing between the Corporate and Non-Corporate. To estimate this split, we rely on data from the 1989 SCF survey. We use the calculated ratio of Corporate PBW to total Private Business Wealth for different wealth percentile groups in the 1989 SCF survey (see Figure C.1a) to allocate Private Business Wealth into corporate and non-corporate components in the SCF+ data.

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

C.3 Portfolio Shares by Wealth Percentile

To compute the household's portfolio share in each asset, we divide the dollar holdings in the asset (or liability) by the household's wealth. Households hold significantly different portfolios, depending on their wealth. Figure C.2 plots wealth-weighted portfolio shares for households of different wealth percentiles. We compute the wealth-weighted shares for each survey and the average across all years from 1983 to 2022. Finally, we re-scale the weights to sum to one. Figure C.2a plots the portfolio shares in Equities, Non-Corporate PBW, and Corporate PBW. Figure C.2b plots the portfolio shares in Real Estate and Vehicles. Figure C.2c plots the portfolio shares in Fixed Income and Cash and Deposits. Figure C.2d plots the portfolio shares in Mortgage Debt, Vehicle Debt, Student Debt and Other Debt. Liabilities are reported as a negative number.

Figure C.3 shows the median portfolio shares within each wealth percentile bin. It then averages across all waves and re-scales to make portfolio weights sum to one.

C.4 Portfolio Shares by Age

Figure C.4 displays portfolio shares by age, using SCF data from 1983 until 2022. Households are categorized into different age groups: 25-35, 35-45, 45-55, 55-65, 65-75, and 75-85 years old. Figure

Figure C.2: Wealth-Weighted Portfolio Shares by Wealth Percentile

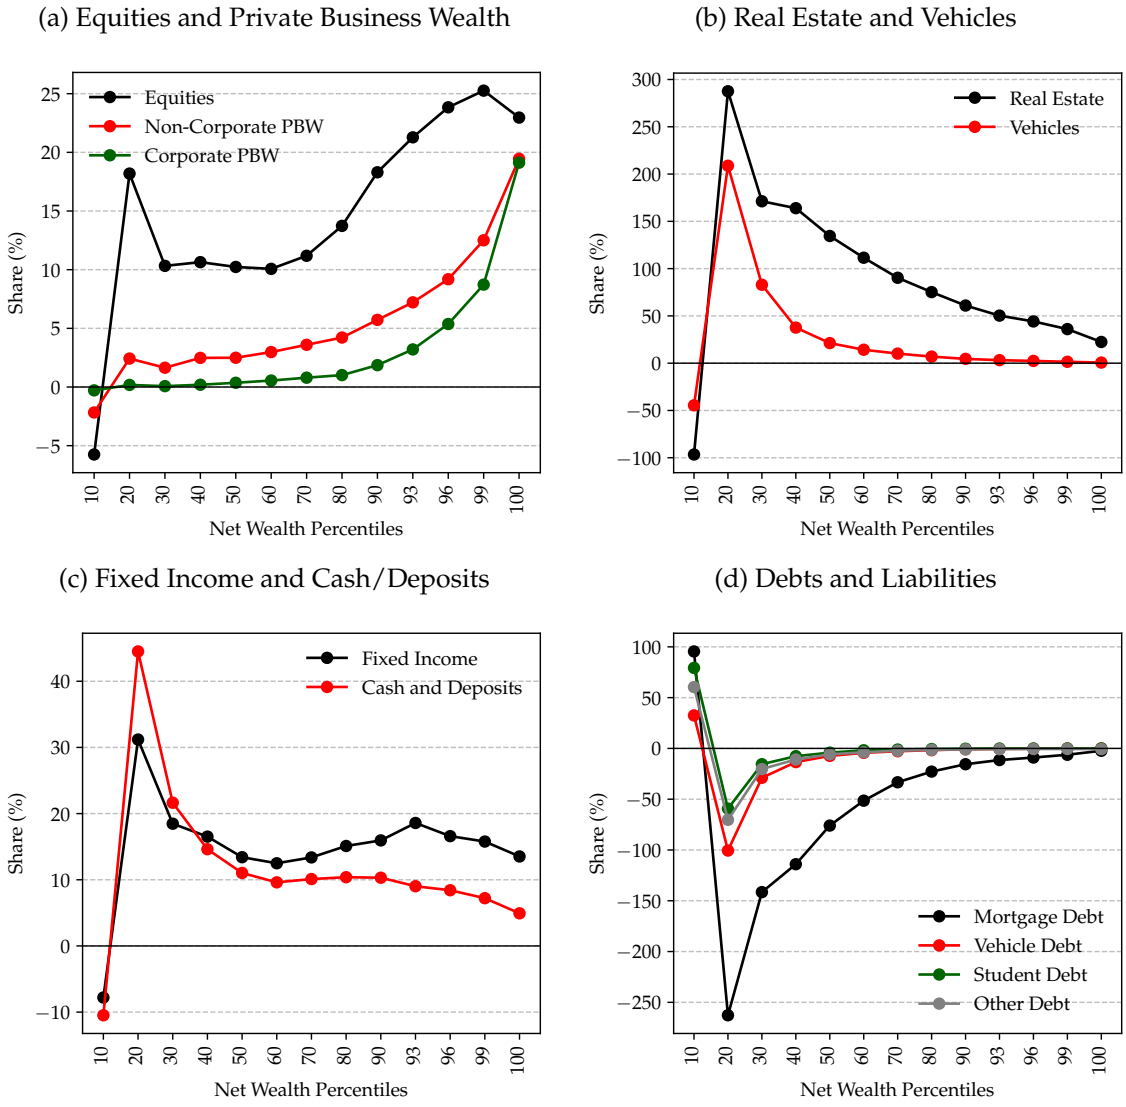

Note: Figure C.2 plots wealth-weighted portfolio shares for households of different wealth percentiles. We compute the wealth-weighted shares for each survey and the average across all years from 1983 to 2022. Finally, we re-scale the weights to sum to one. The sample runs from 1983 to 2022.

C.4a plots value-weighted portfolio shares, while Figure C.4b plots median portfolio shares. The shares are computed for each survey wave and demographic group, averaged across years, and subsequently rescaled to ensure they sum to 100%.

Figure C.3: Median Portfolio Shares by Wealth Percentile

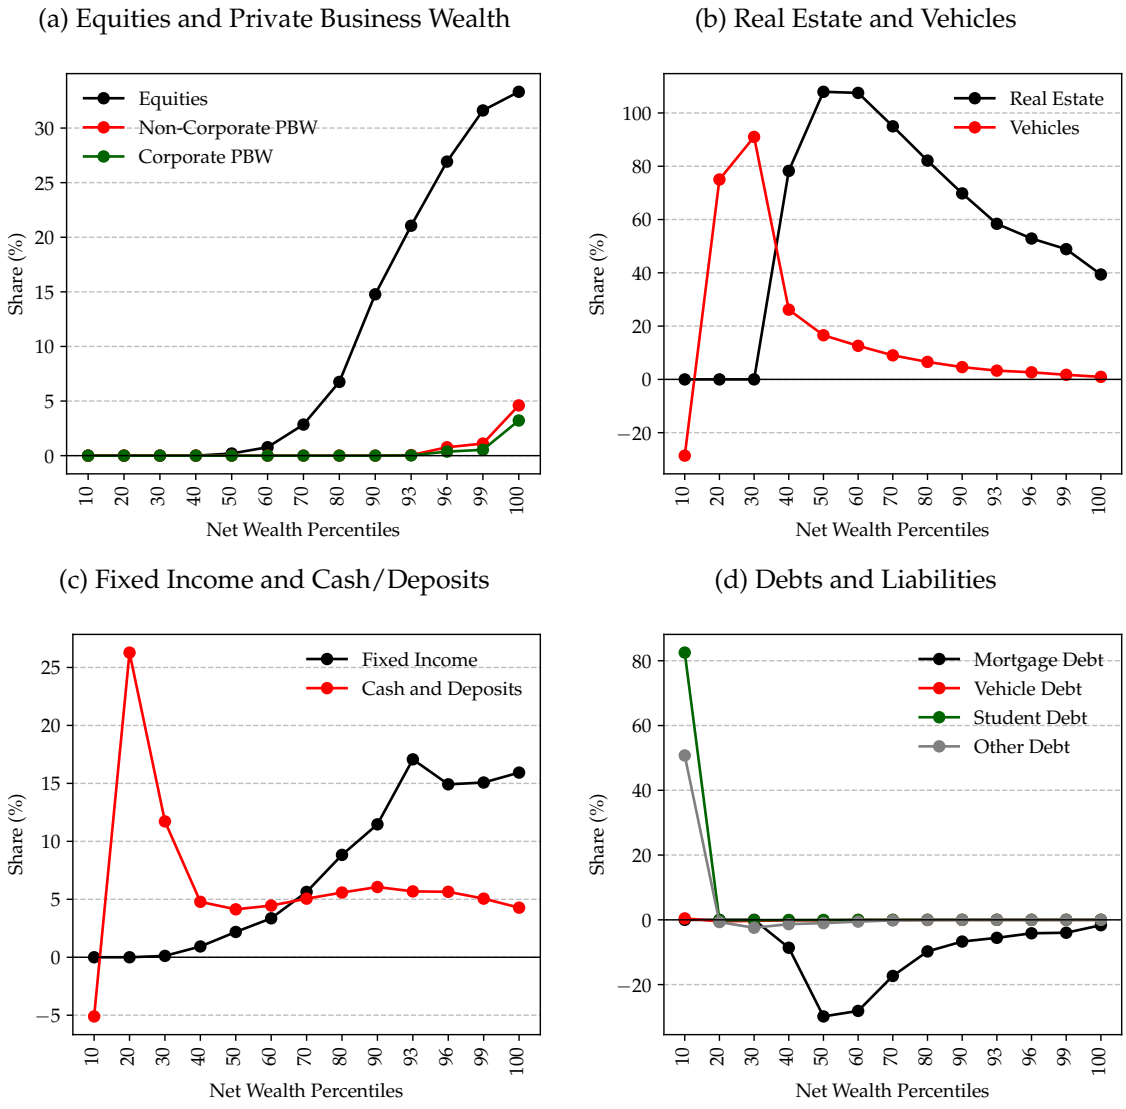

Note: Figure C.3 shows the median portfolio shares within each wealth percentile bin. It then averages across all years and re-scale to make portfolio weights sum to one. The sample runs from 1983 to 2022.

Figure C.4: Portfolio Shares by Cohorts

(a) Wealth-Weighted Portfolio Shares

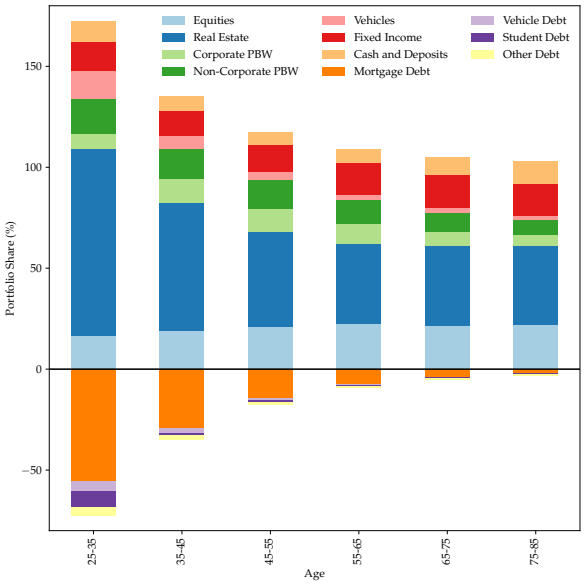

(b) Median Portfolio Shares

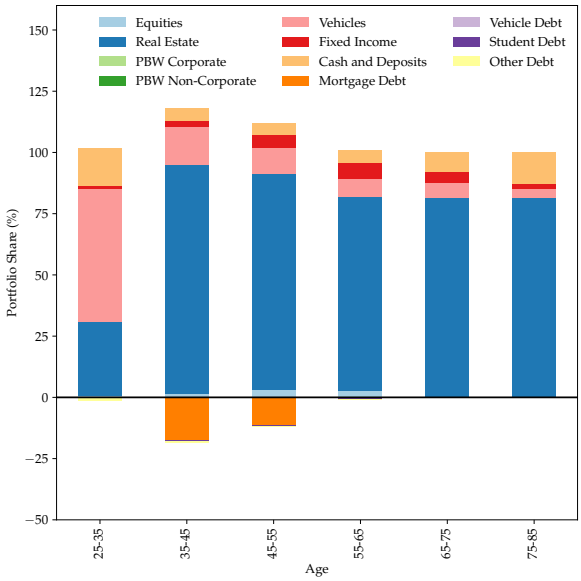

Note: Portfolio shares by age group, averaged across years. Households are categorized into different cohort groups: 25-35, 35-45, 45-55, 55-65, 65-75, and 75-85. Figure C.4a utilizes value-weighted portfolio shares. Figure C.4b plots the median portfolio share in each asset category. The shares are computed for each survey wave and demographic group, averaged across years, and subsequently rescaled to ensure they sum to 100%.

D Duration Measurement

D.1 Duration of Equities

Our baseline duration measure uses Shiller's cyclically adjusted price-earnings (CAPE) ratio. We convert the price-earnings ratio into a price-dividend ratio using a constant payout ratio of 0.496, equal to the average dividend-earnings ratio over our 1983–2023 sample (see Appendix E.3 for alternative payout scales). We then compute duration using equation (8). Figure D.3a shows the resulting time series.

We next consider several alternative duration measures, reported in Table D.1 for $\phi = 1$ and $\phi = 0.99$. First, we compute duration using the price-dividend ratio of the value-weighted market portfolio from CRSP (column 2). Second, we use equity price and dividend data from the Jordà–Schularick–Taylor Macrohistory Database (JST; column 3).[46] Third, we compute duration from the historical S&P 500 price-dividend ratio available on Shiller's website (column 4). All three approaches yield estimates consistent with the baseline measure.

As a fourth approach, we utilize data from the Financial Accounts of the United States, summing the market value of Corporate Equities of Non-Financial Corporate Business (LM103164103) and Foreign Direct Investment in U.S. Non-Financial Corporate Business (LM103192105) to measure the price, and using Dividends Paid by Non-Financial Corporate Business (FA106121001) for the dividend series. These results, shown in column (5) of Table D.1, are somewhat lower than those of the other three methods, but still display a duration far longer than that of alternative assets like real estate.

Table D.1: Duration of Equities

	Baseline	CRSP	JST	S&P 500	FAUS Corporate
	(1)	(2)	(3)	(4)	(5)
$\phi = 1$	50.5	53.8	49.1	50.3	35.7
$\phi = 0.99$	33.0	34.3	32.4	32.9	26.0

Note: The table reports duration estimates for equities. Columns (1)–(5) present equity duration measures based on price-dividend ratios constructed from: (1) Shiller's CAPE, (2) CRSP, (3) the Jordà–Schularick–Taylor Macrohistory Database (JST), (4) the S&P 500, and (5) FAUS. The sample spans 1983–2023.

[46] JST data are available through 2020. For 2021–2023, we carry forward the 2020 values.

D.2 Duration of Real Estate

Our baseline measure of real estate duration is based on Zillow data. The main advantage of using Zillow data to compute the price-to-rent ratio is that Zillow compares similar homes when it constructs rental rates and home prices, and adjusts transactions data for variation in characteristics (hedonics). As a measure of house prices, we use the *Zillow Home Value Index (ZHVI)* for all houses in the U.S. This index measures the typical home value and market changes across regions and housing types. It reflects the typical value for homes in the 35th to 65th percentile range. As a measure of dividends (i.e., rents), we use the *Zillow Observed Rent Index (ZORI)*. This index is a smoothed measure of the typical observed market rate rent across a region. ZORI is a repeat-rent index that is weighted to the rental housing stock to ensure representativeness across the entire market, not just those homes currently listed for rent. We use the Zillow ZORI index for the available sample, which begins in March 2015.

Prior to March 2015, we impute the price/rent ratio backward using alternative data sources. We calculate the growth rate in prices from January 1980 to March 2015 using the growth rate in the House Price Index from the U.S. Federal Housing Finance Agency. We compute the growth rate in dividends (rents), using the growth rate in the Shelter Consumer Price Index for All Urban Consumers from the U.S. Bureau of Labor Statistics. Using the level of prices and dividends in March 2015, we use the growth rate from 1980 to 2015 to reconstruct the full time series of prices and rents. We then compute the duration of Real Estate assets using equation (11). The results are displayed in Column (1) of Table D.2. Figure D.3b plots the time series of these duration estimates.

Table D.2: Duration of Real Estate

	Baseline	JST	Match P-D increase
	(1)	(2)	(3)
$\phi = 1$	13.8	13.3	6.2
$\phi = 0.99$	12.2	11.9	5.9

Note: The table reports the duration for real estate. Column (1) reports the duration using data from Zillow, Column (2) uses data from the Jordà–Schularick–Taylor Macrohistory Database (JST), and Column (3) reports the duration implied by matching the rise in the price-to-dividend (i.e., price-to-rent) ratio. The sample spans 1983–2023.

As an alternative measure, we consider prices and rents (dividends) from the Jordà–Schularick–Taylor Macrohistory Database Jordà, Schularick, and Taylor (2017); Jordà et al. (2019), which we abbreviate as JST. The JST database reports the rental yield net of the ad valorem costs discussed in Section 2. As discussed in Proposition 3, the correct duration formula should be computed using

the gross rental yield that does not deduct these costs. To recover the gross rental yield, we assume the same proportion of ad valorem costs $\tau = 3\%$ to capture property taxes, insurance, and maintenance costs, as discussed in Section 2. We then impute the gross rental yield as the net rental yield plus τ and apply the standard duration formula in (11). Our results using this measure are reported in column (2) of Table D.2.

As a third measure, we compute the duration of real estate that would be required to match the increase in the log price-to-dividend (i.e., price-to-rent) ratio over our sample. Specifically, using (20) and the change in real yields, we back out the duration that exactly matches the change in the log price-to-dividend ratio from 1983 to 2023. The results are displayed in column (3).

D.3 Duration of Non-Corporate Private Business Wealth

Appendix C.1.1 explains how we split Private Business Wealth (PBW) into two components: Corporate and Non-Corporate PBW.

To estimate the duration of Non-Corporate PBW, we utilize data on Non-Corporate Business wealth from the Financial Accounts of the United States. The value (or price) is derived by summing Proprietors' Equity in Non-corporate Business (LM112090205) with Foreign Direct Investment in U.S. Non-Financial Non-Corporate Business (LM115114103). For the dividend series, we use Withdrawals from Income of Non-Financial Non-Corporate Business (FA116122001). This cash flow measure includes both labor and capital remuneration. Since we want to capture only the capital remuneration component, we follow the literature and assign 30% of the business income to capital (Alvaredo et al., 2020; Garbinti, Goupille-Lebret, and Piketty, 2021; Martínez-Toledano, 2020). Thus, we multiply the series Withdrawals from Income of Non-Financial Non-Corporate Business by 0.3. Figure D.3c illustrates the duration time series for the Non-Corporate component of Private Business Wealth. Table D.3 reports the average baseline duration over the sample, as well as a number of alternative measures.

Quadrini and Rıos-Rull (1997) and Krueger and Perri (2006) adopt a labor income to business income ratio of 0.86, implying that only 14% of business income is capital income. This naturally leads to a lower dividend on Non-Corporate PBW and a higher price-dividend ratio. Column (2) of Table D.3 displays the duration estimates using a capital-income share of 14%, yielding durations of 55.6 when $\phi = 1$ and 35.7 when $\phi = 0.99$, notably higher than our baseline measures.

As a second robustness check, we use a 50% capital-income share, yielding durations of 15.8 when $\phi = 1$ and 13.8 when $\phi = 0.99$ (column (3)). This is the assumption used by the PSID data to separate labor and capital income.

As a third robustness check, column (4) of Table D.3 shows the SCF-based estimate, yielding

Table D.3: Duration of Non-Corporate PBW

Capital income share	Baseline (30%)	Lower (14%)	Higher (50%)	SCF
	(1)	(2)	(3)	(4)
$\phi = 1$	25.7	55.6	15.8	71.1
$\phi = 0.99$	20.5	35.7	13.8	41.7

Note: In this table we provide the duration estimates for Non-Corporate PBW using different approaches. Column (1) displays the baseline estimates for Non-Corporate PBW. In Column (2), we apply the 14% capital income share estimated by Quadrini and Rıos-Rull (1997). Column (3) uses the higher 50% capital income share of the PSID. Column (4) presents the duration estimated using the SCF, as described in Appendix D.5. The sample spans 1983–2023.

durations of 71.1 when $\phi = 1$ and 41.7 when $\phi = 0.99$. Appendix D.5 explains the SCF-based methodology in more detail. This third alternative produces a duration that is markedly higher than the baseline.

D.4 Duration of Corporate Private Business Wealth

The estimation of the duration for the Corporate component of PBW is conducted using the 2-stage Gordon Growth model. Private businesses in this category go through an initial stage of high growth, followed by a second, mature stage. Working backwards, in the second stage of the GGM, we assume that the growth rate ($\tilde{g}$) and the required rate of return ($\tilde{r}$) equal those of the value-weighted public equity market. We assume high-growth private businesses transition into this second stage after a period of twenty years in the first stage (i.e., $n = 20$).[47]

To establish the values of g and r for the first stage, our benchmark approach computes g based on dividends from small stocks, defined as the subset of companies within the first market capitalization decile of publicly-traded companies on NYSE-Amex-NASDAQ. The data are obtained from CRSP. Subsequently, we derive the implied required rate of return (r) by using the price-dividend ratio of small stocks. Specifically, we identify the r value at each point in time as the value that minimizes the difference between the implied price-dividend ratio (given $g, r, \tilde{g}, \tilde{r}, n$) and the observed price-dividend ratio of small stocks. We assume at each point in time the value of $g, r, \tilde{g}, \tilde{r}$ to be fixed, meaning there are no shocks to these underlying parameters.

Finally, given the parameter set $(g, r, \tilde{g}, \tilde{r}, n)$, we compute the duration of Corporate PBW. Since both the required rate of return and the growth rate shift after period n, it is not feasible to apply the standard closed-form expression in equation (8). Instead, we evaluate the duration numerically using equation (5), employing g, r for the first n periods and $\tilde{g}, \tilde{r}$ from period n onward.

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

Figure D.3d shows our baseline duration time series for Corporate PBW. Column (1) of Table D.4 shows the full-sample averages.

We also compute a number of alternative duration measures for Corporate PBW. The first two alternatives use the same two-stage GGM approach but with alternative empirical price-dividend targets. The first alternative uses Pitchbook data on price-revenue ratios for private equity portfolio companies with positive revenues. The Pitchbook data only start in 2001. We average this ratio in several ways: equally-weighted, revenue-weighted, or valuation-weighted, with either 1-5% trimming or winsorization when calculating these averages. The revenue-weighted average with 5% winsorization delivers the most conservative valuation ratios. Averaged from 2001–2019, we get a price/revenue ratio of 7.9. We use a revenue-earnings ratio of 7.9 for the first decile of publicly listed companies, and the earnings-dividend ratio of 2.0 for all public companies to translate the price-revenue ratio into a price-dividend ratio ($\frac{P}{Div} = \frac{P}{S} \times \frac{S}{E} \times \frac{E}{Div}$). Column (2) of Table D.4 shows the duration computed using Pitchbook data. This value is higher than the benchmark measure, making the latter conservative, yielding 60.1 years when $\phi = 1$ and 38.7 years when $\phi = 0.99$.

The second alternative uses the median price-revenue ratio reported by Ritter for companies that do an IPO (1980–2022). We use the average revenue-earnings ratio of 8.4 for the first decile of publicly listed companies, and the earnings-dividend ratio of 2.0 for all public companies to translate the price-revenue ratio into a price-dividend ratio ($\frac{P}{Div} = \frac{P}{S} \times \frac{S}{E} \times \frac{E}{Div}$). These calculations are conservative since they use sales-earnings and earnings-dividend ratios from more mature companies, which are lower than for companies that IPO. Column (3) of Table D.4 shows the duration using IPO data, which is higher than our benchmark measure, yielding 58.5 years when $\phi = 1$ and 37.3 years when $\phi = 0.99$.

The third alternative uses a different method, namely a one-stage GGM applied to small stocks. It results in a much higher duration estimate. Average durations obtained from the one-stage Gordon Growth model for small stocks are listed in Column (4) of Table D.4, yielding 100.3 years when $\phi = 1$ and 49.7 years when $\phi = 0.99$.

The fourth alternative calculates valuation ratios and durations using detailed information from the SCF dataset. Appendix D.5 explains the methodology. Column (5) of Table D.4 shows a measure that is substantially higher than our baseline measure, yielding 65.9 years when $\phi = 1$ and 39.8 years when $\phi = 0.99$.

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

Table D.4: Duration of Corporate PBW

	Baseline	Pitchbook	IPO	1-stage GGM	SCF
	(1)	(2)	(3)	(4)	(5)
$\phi = 1$	58.3	60.1	58.5	100.3	65.9
$\phi = 0.99$	37.3	38.7	37.3	49.7	39.8

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

D.5 Alternative Duration for Non-Corporate PBW and Corporate PBW Using SCF

As noted in Appendix Sections D.3 and D.4, one alternative measure of the duration of Non-Corporate PBW and Corporate PBW is based on SCF data, an approach detailed here.

D.5.1 Payout

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

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

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

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

To estimate the shadow labor income for actively managed businesses, we follow an approach similar to Catherine et al. (2023). First, we determine the shadow labor income for households. This involves calculating the wage income for both the head and spouse from their first and second jobs. We use data on the wage (head: X4112, X4509; spouse: X4712, X5109), hours worked per week (head: X4110, X4507; spouse: X4710, X5107) and weeks of work (head: X4111, X4508; spouse: X4711, X5108) to compute the annual wage income for all households in the survey, where in each pair the first code refers to the respondent's main job and the second to their second job. Because respondents report pay at different frequencies, we annualize using the corresponding frequency variables. Next, we perform regression analysis to estimate the shadow labor income. The wage income is regressed on a year fixed-effect, a fixed-effect of race, education, and sex, as well as a cubic function of age. Separate regression models are estimated for the head's first job, head's second job, spouse's first job, and spouse's second job. If households have multiple businesses, we impute a pro-quota labor income, where the quota is based on the share of the business as a percentage of total private business wealth.

D.5.2 Valuation Ratios and Duration

By employing this methodology, we can determine the payout for each private business. We also extract information on the market value of each actively-managed business (X3129, X3229, X3329) and the passively-held businesses (X3408, X3412, X3416, X3420, X3424, X3428, X3452). Combining information on the market value and the payout, we compute valuation ratios for each individual business. To ensure accurate calculations, we exclude private businesses with payouts equal to or less than 0, as this would result in negative or undefined ratios. When computing the dollar-weighted average valuation ratio we further exclude businesses with a price-payout ratio of 500 or above, whose payouts are negligible relative to their value. We then compute dollar-weighted averages for each year. The weighting mechanism involves multiplying survey-weights with the corresponding market value of each business. This approach guarantees that businesses with higher market values carry more influence in the final valuation ratios, appropriately reflecting their significance in the overall assessment. We separately estimate valuation ratios for Corporate Businesses and Non-Corporate Businesses. To derive the duration from valuation ratios, we employ equation (8). To ensure completeness, we assume that the duration for the 1983 survey is equal to the average of 1989 and 1992, and we perform linear interpolations for the missing years. The summary statistics for the Non-Corporate Businesses are reported in Table D.3 while those for Corporate Businesses are shown in Table D.4.

D.6 Duration of Vehicles

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

$$CF_t = (1 - \delta)^t \delta, \tag{D.2}$$

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

$$D_j^v = \frac{\sum_{t=j+1}^{T} (t - j) \frac{CF_t}{(1+r)^{t-j}}}{\sum_{t=j+1}^{T} \frac{CF_t}{(1+r)^{t-j}}}, \tag{D.3}$$

where r denotes the 10-year real rate in each year.

We determine the value for j as the average age of vehicles in use in the U.S., based on data from the Bureau of Transportation Statistics.[48] Additionally, we set the maximum vehicle age, T, equal to double the average age of vehicles: $T = 2j$. To calculate the depreciation rate of vehicles, we utilize data from the U.S. Bureau of Economic Analysis (BEA). The depreciation rate is computed as the current-cost depreciation of autos (BEA Table 2.4) divided by the current-cost net stock of autos in the previous year (BEA Table 2.1). For the years 1980 to 2022, the average depreciation rate (δ) is found to be 22%. Figure D.3e plots the duration time series for Vehicles.

D.7 Duration of Fixed Income Assets

For fixed-income assets, we utilize the duration of ICE-BofA US Corporate & Government bonds, which is available starting in 1996. To estimate the duration for the period 1983 through 1995, we employ the following imputation approach. First, we establish a linear regression model using the 1996–2020 sample of ICE-BofA duration data and the duration from CRSP government bond data (available since 1946), along with a constant. Second, we apply the parameters obtained from this regression model to estimate the ICE-BofA duration values for the sample period of 1983–1995. Figure D.3f plots the time-varying duration series.

[48] The data are available at: https://www.bts.gov/content/average-age-automobiles-and-trucks-operation-united-states. Since the sample for these data begins in 1995, we extend it backward using a constant growth rate estimated from the period 1995 to 2022.

D.8 Mortgage Duration

For mortgage debt analysis, we obtain data from the Bloomberg-Barclays Aggregate MBS Index, which provides a comprehensive representation of all outstanding U.S. pass-through mortgage-backed securities. In the United States, the most prevalent mortgage product is the 30-year fixed-rate mortgage. However, it is essential to note that the average outstanding mortgage has a significantly lower duration due to factors such as loan aging, amortization, coupon payments, and prepayment. These factors contribute to variations in mortgage duration over time and are critical considerations when studying mortgage debt dynamics. See Figure D.3g.

D.9 Vehicle Debt Duration

We compute the vehicle debt duration using the SCF. The SCF provides detailed information on up to four vehicle loans, limited to three vehicle loans for the 1989 and 1992 waves. We use information on all the available vehicle loans. First, we calculate the number of residual monthly payments. We then subtract the number of payments already made from the total number of payments to determine the residual number of payments $T_{i,j}$. Second, we calculate the monthly payment $C_{i,j}$ of each loan. Third, we calculate the monthly interest rate $r_{i,j}$ charged on the loan.

To apply the same annual timing convention used for the other asset classes, we aggregate the monthly payments into annual cash flows: year s collects the payments due in months $12(s-1)+1$ through $12s$, so that $x_{i,j,s}$ equals $C_{i,j}$ times the number of monthly payments falling in year s (twelve, except possibly in the final year). Because these payments arrive spread out over the year, we date each annual cash flow at mid-year, $t_s = s - \frac{1}{2}$, with the final partial year dated at the midpoint of the months it covers. We discount at the annual rate $r^a_{i,j} = (1 + r_{i,j})^{12} - 1$. The duration of vehicle debt j of household i under permanent shocks then follows (6):

$$D^{vd}_{i,j} = \frac{\sum_{s=1}^{\lceil T_{i,j}/12 \rceil} t_s \, x_{i,j,s} \, (1 + r^a_{i,j})^{-t_s}}{\sum_{s=1}^{\lceil T_{i,j}/12 \rceil} x_{i,j,s} \, (1 + r^a_{i,j})^{-t_s}}.$$

For the case where $\phi < 1$, we replace the weight t_s with $\left(1 - \phi^{t_s}\right) / \left(1 - \phi\right)$, the closed form of the geometric weights in (5) evaluated at the mid-year dates t_s, which converges back to t_s as $\phi \to 1$.

We calculate the duration of vehicle debt for each household i by taking the weighted average of the duration of its individual vehicle debts. To determine the weights, we use the outstanding balance of each debt j. The vehicle debt duration measure D^{vd} is computed as the median duration of households' vehicle debts in each survey year. Our first estimate is available in 1989, based on the initial SCF survey. For the years from 1983 until 1988, we perform extrapolation using the

trend growth rate observed in the sample from 1989 to 2022. Additionally, we interpolate for the years between surveys with a cubic spline. Finally, for the 1992 survey we use the interpolated value between 1989 and 1995 as the duration estimate is anomalously low compared to all other observations.

Figure D.1 shows the distribution of information on each individual vehicle debt in 2022, the latest available SCF survey. The figure provides information on the cross-section of: (i) the length of vehicle debts measured in number of years at inception or residual, (ii) the duration estimates ($D^{vd}_{i,j}$), (iii) the annual interest rates of the loan, and (iv) the monthly payment. Figure D.3h shows the time series of Vehicle Debt duration.

Figure D.1: Vehicle Debt (2022)

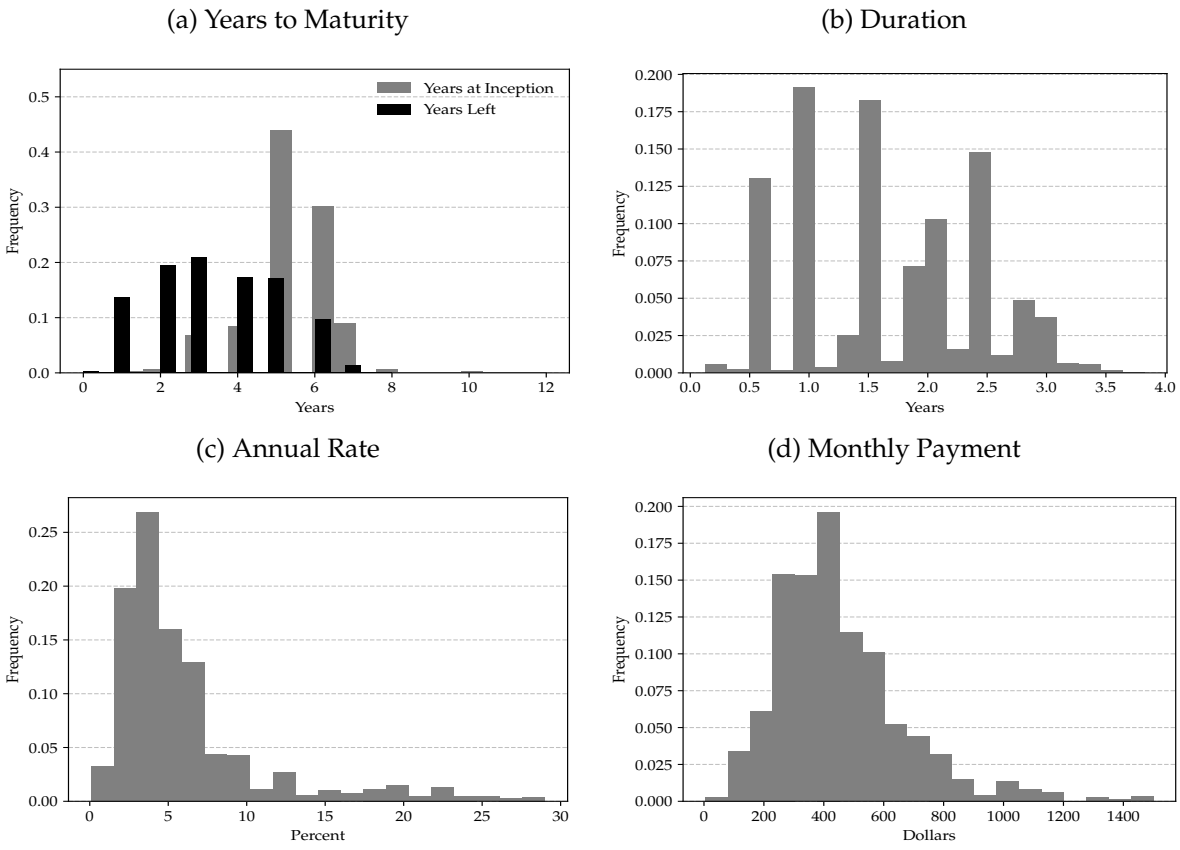

Note: The figures plot the distribution of vehicle debt in 2022. Figure D.1a plots the distribution of number of years left and number of years at inception. Figure D.1b plots the distribution of duration estimates. Figure D.1c plots the distribution of annualized interest rates on loans. Figure D.1d plots the distribution of monthly payments. Source: SCF 2022.

D.10 Student Debt Duration

For student debt, we use an approach similar to the vehicle debt duration (Appendix D.9). The SCF provides information on up to six education loans per household. For each of them, it provides information on (i) the year in which the loan is expected to be repaid, (ii) the year in which the borrower (most recently) started making payments, (iii) the agreed number of years at inception, (iv) the agreed number of payments at inception, (v) the payment amount, (vi) the frequency of payments, (vii) the interest rate, and (viii) the outstanding balance. Based on this information, we start by estimating a residual maturity for each loan. If the expected payoff year is available, we use that information and subtract the year of the survey to measure the remaining maturity. If it is not available, we use the agreed number of payments net of the payments already made (converted to months using the payment frequency), or the agreed number of years net of the years elapsed since payments started; these two measures encode the same underlying answer in different units, as the SCF records the agreed term either in years or in number of payments. We restrict the sample to loans in repayment at the time of the interview, since no payment stream is observed for loans in deferment, grace periods, or forbearance.

Figure D.2a plots the distribution of the residual maturity implied by the expected payoff year for student debt in 2022. The median is 5.0 years and the mean is 7.7 years. As we can see from the figure, this is lower than the maturity computed based on the initially agreed schedule. This second measure, plotted in Figure D.2b, has a median of 9.0 years and a mean of 10.7 years. The main reason is that—according to the SCF in 2022—about 60 percent of paying loans are reported as being paid off ahead of schedule, so borrowers expect to repay before the agreed schedule runs out. In our computation we use the first measure (expected) and only supplement with the second measure when the first is not available.[49] We then use the cash flows and the interest rates (plotted in Figure D.2c) to compute the duration of each loan, applying the same annual timing convention and duration formula used for vehicle debt in Appendix D.9: the monthly payments are aggregated into annual mid-year cash flows and the duration follows (6), with the geometric weights of (5) when $\phi < 1$.

As for vehicle debt, we aggregate the loan-level durations within each household using the outstanding balances as weights, and compute the student debt duration as the median duration across households in each survey year. The education-loan grid is available starting with the 1992 survey; for the years 1983 through 1991 we extrapolate using the trend growth rate estimated over the 1992–2022 sample, and we interpolate between surveys. The duration histogram for 2022 is

[49] The expected payoff year is asked of essentially all loans only from the 2016 wave onward; in earlier waves it is asked only of loans that are not being repaid on schedule, so it covers between roughly 10 and 55 percent of paying balances, with the agreed-schedule measure supplying the remainder.

plotted in Figure D.2d. The loan-level median for 2022 is 2.9 years, which is slightly lower than the mean of 3.6 years.

Figure D.3i shows the time series of Student Debt duration.

Figure D.2: Student Debt (2022)

(a) Residual Years to Maturity Expected

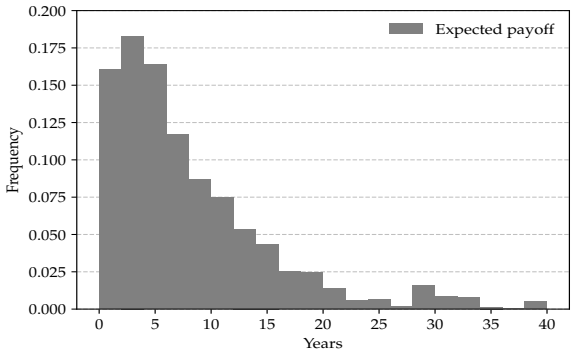

(b) Residual Years to Maturity Agreed

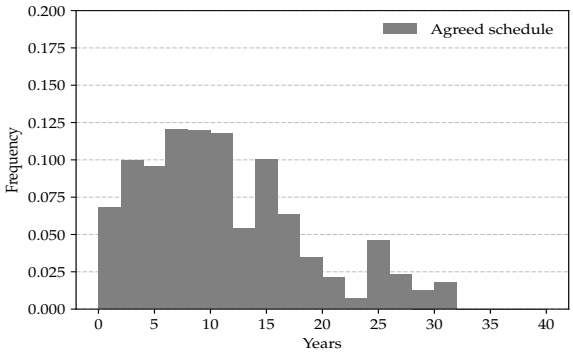

(c) Interest Rates

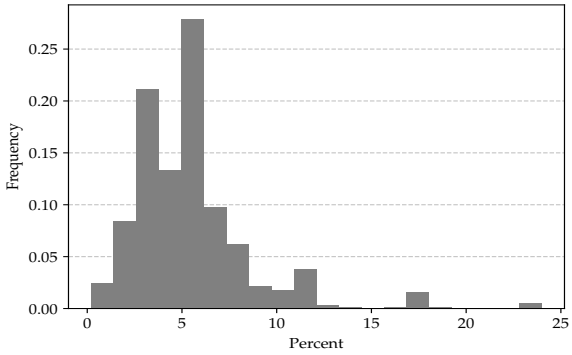

(d) Duration

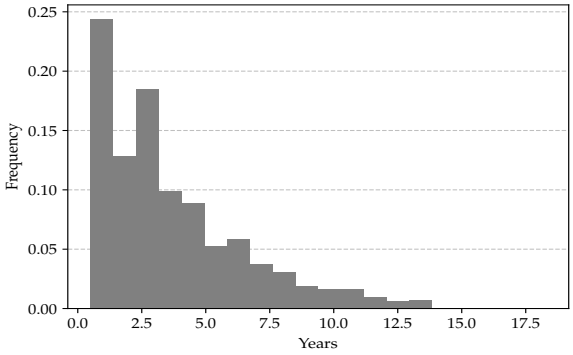

Note: The figures plot cross-sectional distributions for education loans in repayment in 2022, at the loan level and using survey weights. Figure D.2a plots the residual maturity implied by the year in which the borrower expects the loan to be repaid. Figure D.2b plots the residual maturity implied by the agreed repayment schedule. Figure D.2c plots the distribution of annual interest rates on the loans. Figure D.2d plots the distribution of the duration estimates. Panels (c) and (d) refer to the loans entering the duration computation. Source: SCF 2022.

Figure D.3: Duration Estimates Over Time

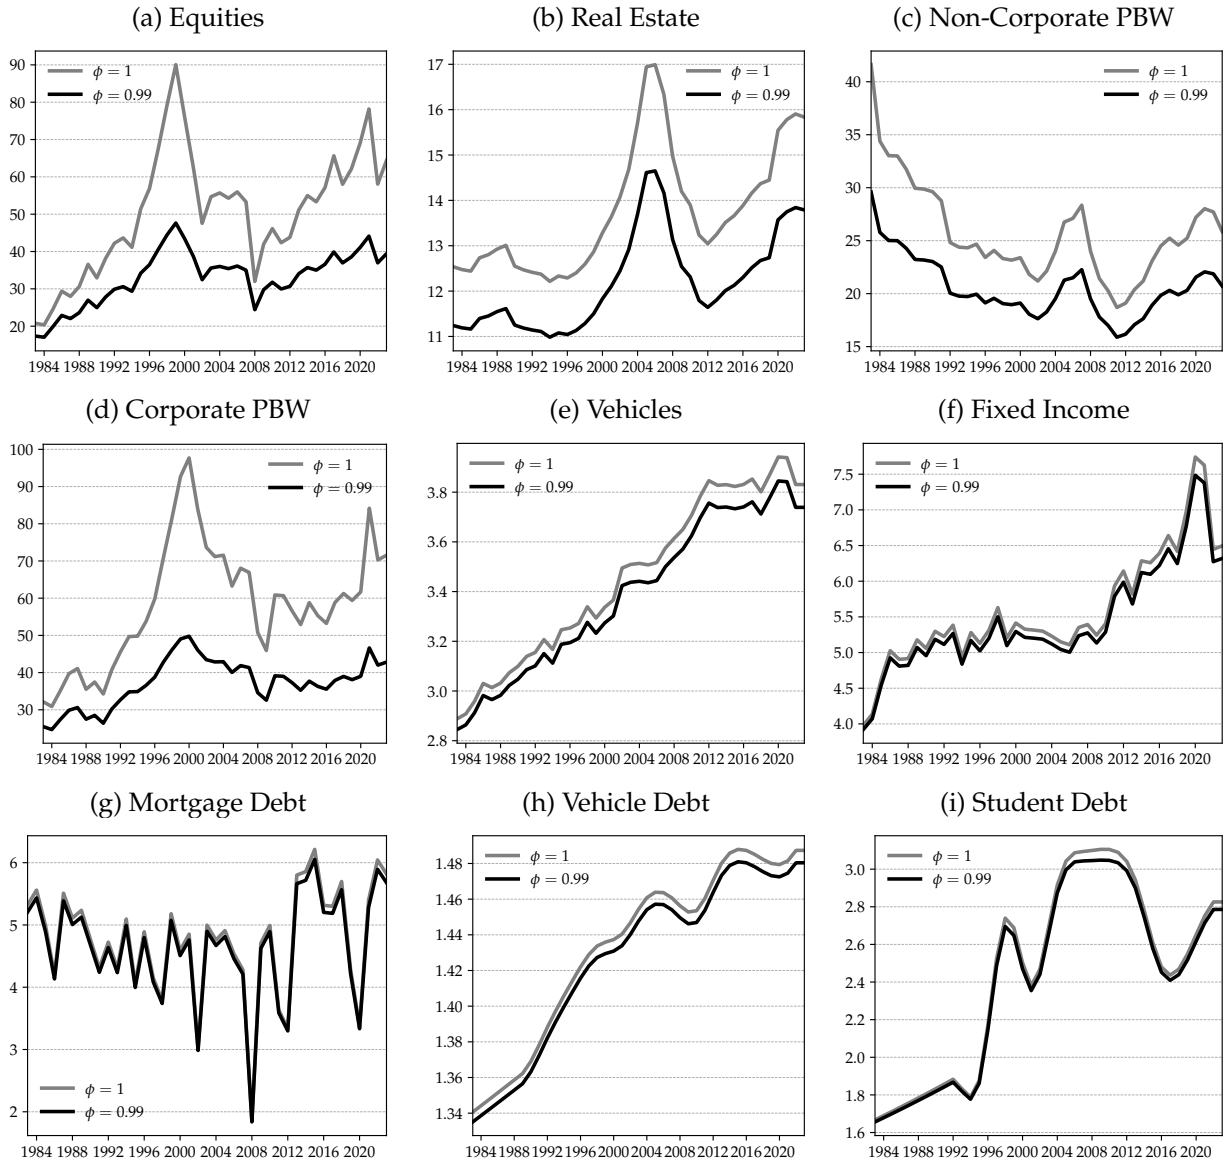

Note: The figures plot the time varying duration estimates for the full set of assets and liabilities. The sample runs from 1983 to 2023.

D.11 Portfolio Duration

We calculate the portfolio duration of each household by multiplying the duration of each asset (or liability) by its corresponding portfolio share and summing over all assets (liabilities) in the portfolio. In each year, we winsorize the top/bottom 2.5% of households ranked by the duration of their portfolio. In Section 2.2 we discussed the heterogeneity in portfolio duration by wealth and age. Figure 2 displayed household duration by wealth for the case of $\phi = 1$. Figure D.4 plots the corresponding figure for the case of $\phi = 0.99$.

In Figure D.5a we use the same procedure but instead use the wealth-weighted wealth per-

Figure D.4: Financial Duration by Wealth Percentiles, $\phi = 0.99$

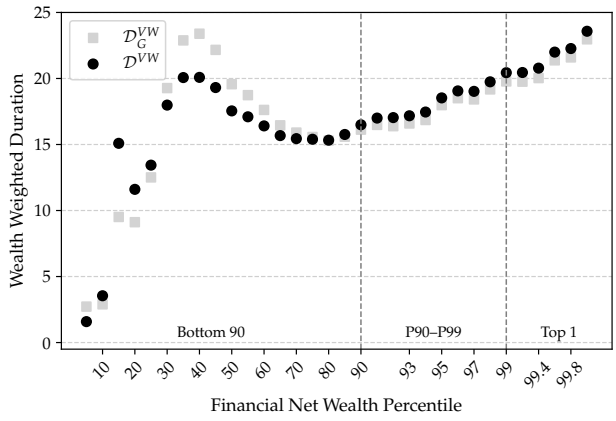

Note: The figure displays average duration by financial wealth bin in the data. The x-axis is measured in percentiles, with each tick representing the right edge of the bin, so that e.g., "10" corresponds to households with financial wealth percentile in the interval $[5, 10]$. We plot 1-percent bins between the 90th and 99th percentiles and 0.2 percent bins above the 99th percentile. The data are averages over all SCF waves between 1983 and 2022. The black dots display $\mathcal{D}^i(\phi)$ from (7), while the gray dots display $\mathcal{D}_G^i(\phi)$ from (17) and $\phi = 0.99$.

centile. Each wealth-weighted percentile bin is designed such that the share of total wealth held by the households in each bin is the same across different bins.

Households' portfolio duration notably differs by age. However, there is also dispersion within age group. Figure D.5b provides further information on the within age group dispersion of duration. For each age group, we rank households by the duration of their portfolio. The figure includes the 5%, 25%, 50%, 75% and 95% percentile, as well as the median (blue line) and mean (orange triangle). We pool households across all surveys from 1983 to 2022.

We also evaluate more formally the correlation between financial duration and some covariates of interest. In each regression we use the baseline duration measure under $\phi = 1$. Table D.5 reports the estimation results. Data are based on all SCF surveys from 1983 to 2022. All regression models include year fixed effects. In column (1), we regress household financial duration on household age. In column (2), we regress financial durations on the household position in the Lorenz Curve. To calculate households' positions, we rank households by their wealth, then calculate the cumulative sum of wealth and divide by the aggregate 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 household wealth.

The results in Table D.5 show a non-linear relationship between age and duration. In columns (4) and (5), duration increases with age up to roughly 45 years and declines thereafter once income

Figure D.5: Distribution of Durations

(a) Duration by Wealth-Weighted Wealth Percentile

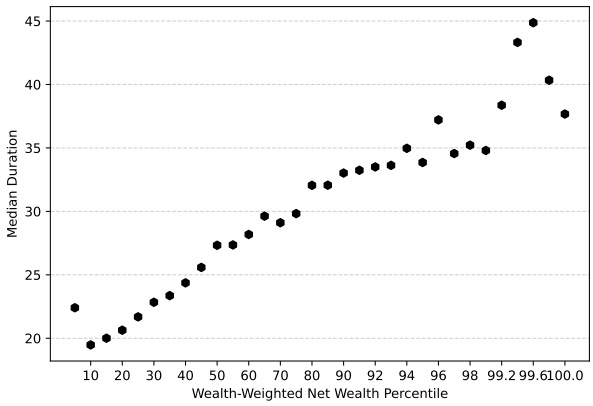

(b) Distribution of Durations by Age

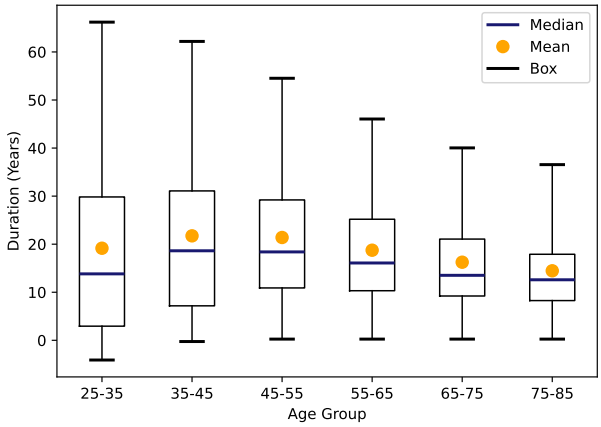

Note: In Figure D.5a we plot the duration by wealth-weighted wealth percentile. Each wealth-weighted percentile bin is designed such that the share of total wealth held by the households in each bin is the same across different bins. Figure D.5b provides further information on the within-age group dispersion 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% percentiles, as well as the median (blue line) and mean (orange triangle). We pool households across all surveys from 1983 to 2022.

and year effects are controlled for. In column (6), where wealth is added, the turning point shifts to the mid-twenties, indicating that the age profile becomes flatter once wealth differences later in life are accounted for.

Column (6) indicates a strong positive association between duration and net wealth, conditional on age and income.

D.11.1 Financial Duration Over Time

Figure D.6 summarizes the time variation in the duration of household portfolios. We plot the value-weighted duration for the bottom 90%, top 10%, and top 1% under the assumption of $\phi = 1$ (Panel D.6a) and $\phi = 0.99$ (Panel D.6b). Portfolio duration rises through the 1990s and peaks in the early 2000s, particularly for the wealthiest households. This pattern reflects both the high duration of equity and private business wealth and their elevated portfolio shares during this period. The figure also shows that the duration gap between the top 10% and the bottom 90% is especially large at the beginning of the sample and widens further toward the end of the 1990s. It subsequently narrows, reaching its minimum around 2010, before diverging again in more recent years.

To better understand the role of portfolio composition, we conduct a separate analysis that fixes the duration of assets and liabilities at their sample averages (measured over the period

Table D.5: Determinants of Household-level Financial Duration

	(1)	(2)	(3)	(4)	(5)	(6)
Age	-0.047***		-0.11***	0.61***	0.43***	0.15***
	(0.0022)		(0.0021)	(0.013)	(0.013)	(0.013)
Lorenz Curve		0.24***	0.28***	0.25***	0.12***	-0.063***
		(0.0019)	(0.0018)	(0.0019)	(0.0030)	(0.0027)
Age Squared				-0.0068***	-0.0047***	-0.0030***
				(0.00011)	(0.00011)	(0.00012)
Log-Income					3.35***	2.69***
					(0.052)	(0.051)
Log-Net-Wealth						1.32***
						(0.018)
Constant	14.0***	9.14***	13.9***	-2.54***	-33.5***	-29.1***
	(0.13)	(0.071)	(0.12)	(0.33)	(0.56)	(0.54)
Year effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	311215	311215	311215	311215	311200	288228
R^2	0.039	0.072	0.082	0.096	0.134	0.176

Note: The table reports the estimation results of regressing households' portfolio duration on a set of covariates. Data are based on all SCF surveys from 1983 to 2022. All regression models include year fixed effects. In column (1), we regress household financial duration on household age. In column (2), we regress financial durations household position in the Lorenz Curve. To calculate households' positions, we rank households by their wealth, then calculate the cumulative sum of wealth and divide by aggregate 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 household wealth. Standard Errors in parentheses (* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$).

1983–2023), thereby isolating the contribution of time-varying portfolio shares. Figures D.6c–D.6d show that the duration gap between the bottom 90% and the top 10% is particularly large at the beginning of the sample and narrows steadily until around 2015, after which it starts to widen again. This pattern reflects two main sources of portfolio heterogeneity. Early in the sample, the bottom 90% held substantially lower shares of equity than the top 10%, which generated a large initial duration gap. Over time, this gap narrowed as the equity share of the bottom 90% increased relative to that of the top 10%. A second important factor is leverage (i.e., borrowing), which raises the portfolio duration for leveraged households. Leverage among the bottom 90% rose until the late 2000s and then declined after the global financial crisis, contributing to the renewed divergence in portfolio duration.

Figure D.6: Value-Weighted Household Portfolio Duration Over Time

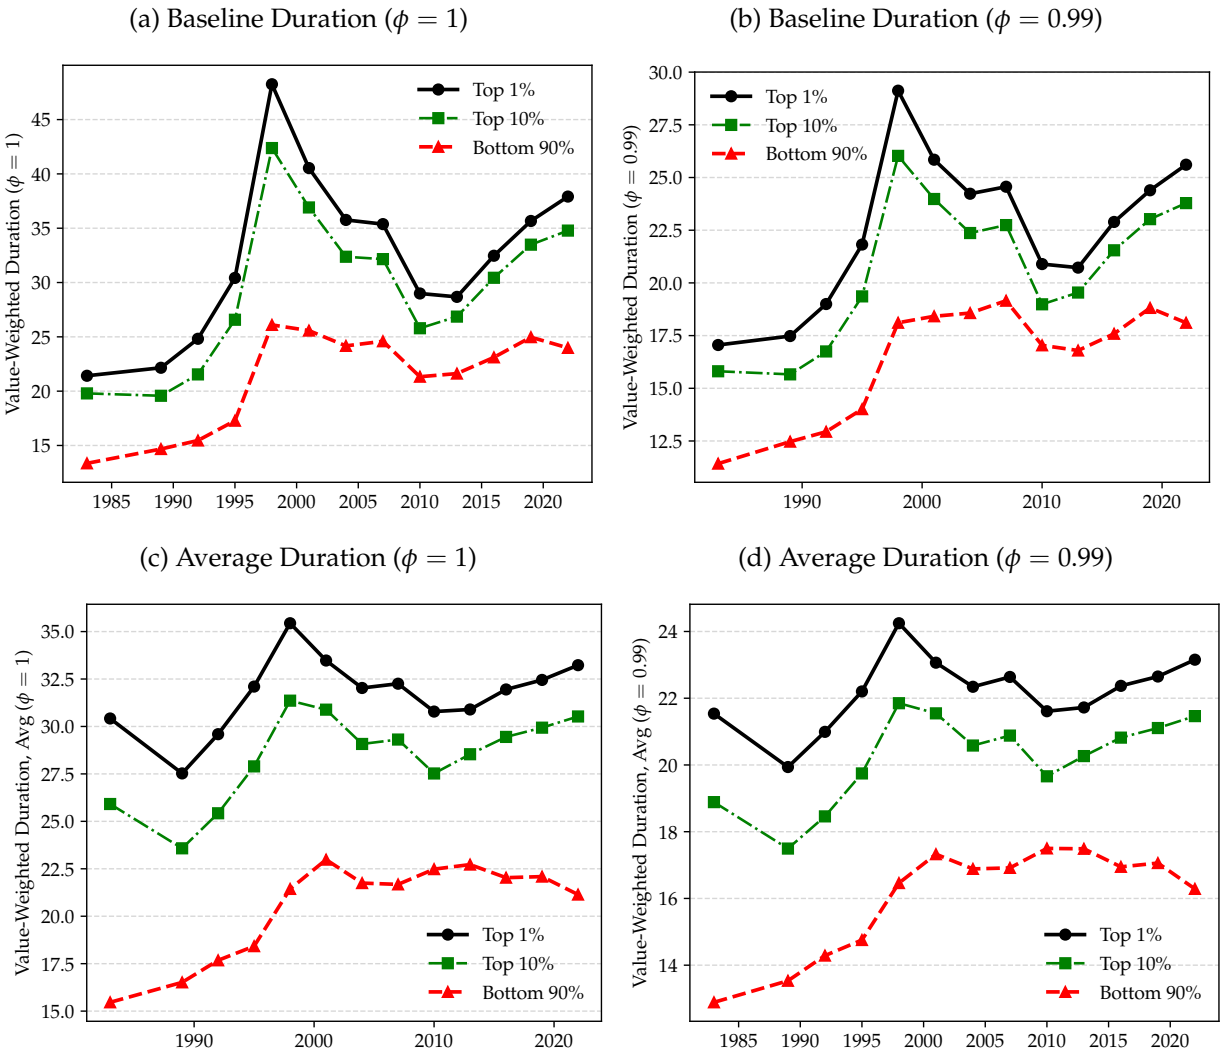

Note: Figure D.6 reports the evolution of the value-weighted duration of household portfolios across survey waves from 1983 to 2022. The top panels display baseline duration estimates, in which both asset and liability durations vary over time, for $\phi = 1$ (left) and $\phi = 0.99$ (right). The bottom panels hold asset and liability durations fixed at their sample averages while allowing portfolio shares to vary across survey years, again for $\phi = 1$ (left) and $\phi = 0.99$ (right). Each panel plots duration separately for the Top 1%, Top 10%, and Bottom 90% of the wealth distribution, where households are assigned to wealth groups based on the net-wealth percentile ranking at each survey date. The sample spans 1983–2022.

D.12 Duration By Age

To compute durations by age, controlling for household wealth, we estimate the following regression for each survey wave y using survey sampling weights:

$$D_y^i = \alpha_y + \sum_k \beta_{y,k} \mathbb{1}\{AgeGroup_i = k\} + \sum_j \gamma_{y,j} \mathbb{1}\{NetWealthBin_i = j\} + e_y^i, \tag{D.4}$$

where the $\beta_{y,k}$ coefficients capture the partial contribution of each age group holding wealth constant. For each wave y, we construct the fitted duration for age group k evaluated at the average wealth-bin effect,

$$\widehat{D}_{y,k} = \alpha_y + \beta_{y,k} + \sum_j \gamma_{y,j} w_{y,j},$$

where $w_{y,j}$ denotes the survey-weighted share of households in wealth bin j in wave y. We then average $\widehat{D}_{y,k}$ across survey waves to obtain the age profile reported in the figure. This procedure yields the wealth-adjusted duration-by-age curve plotted in Figure 3b. Figure D.7 plots the estimates under $\phi = 0.99$.

Figure D.7: Duration by Age, $\phi = 0.99$

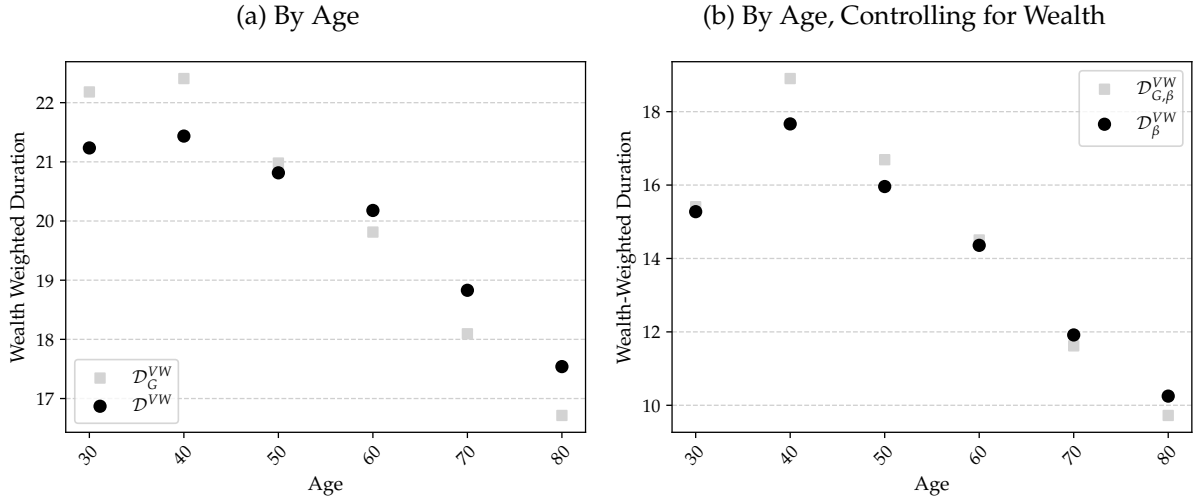

Note: The left panel plots wealth-weighted portfolio duration across 10-year age groups, where each point corresponds to the upper bound of the respective age bin. The right panel reports the analogous age profile after netting out cross-sectional differences in household wealth: for each SCF wave, we regress duration on age-group dummies and 5-percent net-wealth bins, recover the age-group coefficients, re-center them using the wave-specific weighted mean, and then average the resulting fitted values across survey years. The sample covers all SCF surveys from 1983 to 2022. In each panel, the black markers display $\mathcal{D}^i(\phi)$ as defined in (7), while the gray markers report $\mathcal{D}_G^i(\phi)$ from (17).

E Additional Empirical Results

E.1 Alternative Measure of GDP Growth Rate

We use an alternative proxy for g using the 10-year real GDP growth forecast from the Survey of Professional Forecasters (SPF). Since our sample starts in 1983 but this series is only available starting in 1992, we supplement it with a series constructed by Greenwald, Lettau, and Ludvigson (2025). This series perfectly tracks the observed SPF forecasts in the post-1992 sample when available, and is filled in for remaining periods based on the estimated low-frequency trend in the realized growth of real corporate net value added. Since Greenwald, Lettau, and Ludvigson (2025) compute a distribution for the latent state values for this variable accounting for both parameter and latent state uncertainty, we use the median. Figure E.1 displays the long-term GDP growth expectations derived from the model (in dark grey) and those inferred from the SPF (in black). Between 1983 and 2023, long-term GDP growth expectations decreased from 3.06% to 1.95%, a change of -1.12 percentage points. We assume that g_t follows the autoregressive process defined in (13). Using the observed g_t series and setting $\bar{g}$ to the sample average, we recover the sequence of $z_{g,t}$, and, given ϕ, the associated shocks $\varepsilon_{g,t}$.

Figure E.1: Long-Term Growth Expectations using SPF

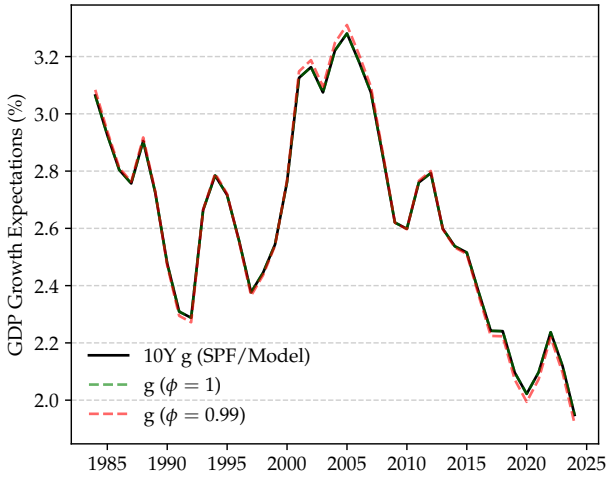

Note: The figure plots long-term real GDP growth expectations from the Survey of Professional Forecasters from 1992 onward and from Greenwald, Lettau, and Ludvigson (2025) prior to 1992.

E.2 Quantifying the Link between Interest Rates and Wealth Inequality, Additional Results

Figure E.2 replicates Figure 6 in the main text, but uses SPF GDP growth forecasts for g. Under this alternative measure, g has slightly declined over the sample. Using this SPF g generates an increase in the top–10% share of wealth that is slightly below the increase implied by our baseline proxy, the moving–average measure of g. However, the rise in the top–10% share remains larger than under the measure that only uses variation in r (i.e., assumes a constant g over the sample).

Figure E.2: Predicted and Observed Top-10% Wealth Shares with SPF g

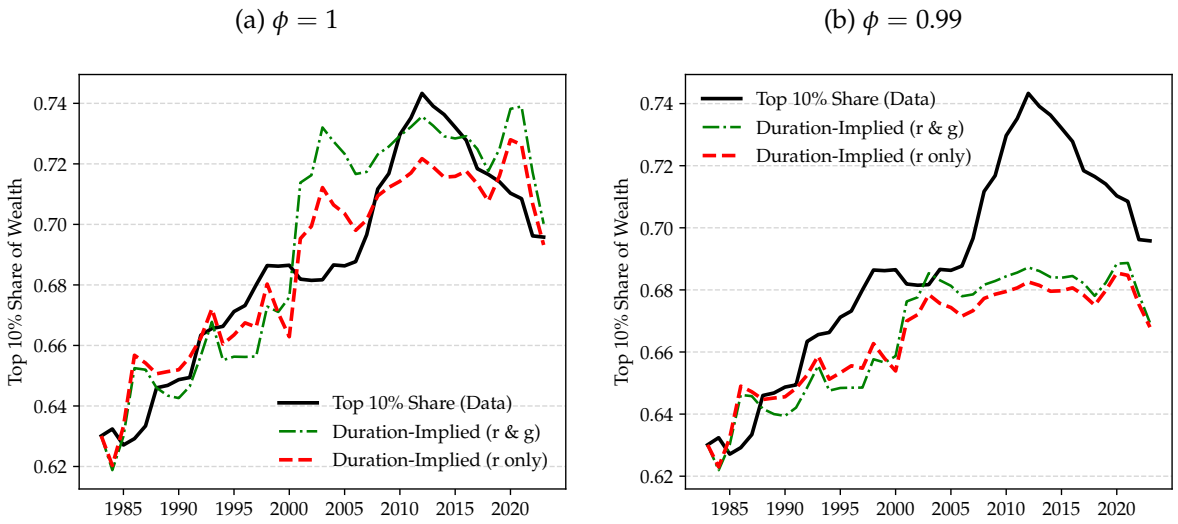

Note: The figure replicates the chained repricing exercise of Figure 6 using the Survey of Professional Forecasters (SPF) measure of long-term growth expectations in place of the baseline g_t.

E.3 Adjusting Equity Duration for the Payout Ratio

Our baseline measures of durations for equity are based on the price to dividend ratio. However, one potential concern is that using dividends as proxy for cash flows may be misspecified if the actual net payout includes capital injections and share buybacks. This may be especially relevant starting in 1982 due to SEC Rule 10b-18, which deregulated share buybacks in 1982.

To mitigate this concern, we rescale the dividend-to-earnings ratio by the ratio of total payouts (dividends plus net equity repurchases) to dividends from the Financial Accounts of the United States (FAUS) over our 1983–2023 sample.[50] The duration of Corporate PBW is constructed from dividend yields as well —the price-dividend ratio of the smallest decile of public firms for the high-growth stage, and of the market for the mature stage—so the same redefinition applies to it. We therefore report two rescaled cases: one in which only the equity payout is rescaled, and one in which the payout of both equities and Corporate PBW is rescaled. Table E.1 reports the scale, the FAUS rescaling factor, and the durations each case implies. The rescaled cases imply larger payouts, hence lower price-payout ratios and shorter durations. Finally, because the dividend-to-earnings ratio itself depends on the window over which it is averaged, we also report the longer 1970–2023 window, which spans both the pre- and post-deregulation payout regimes: once using that ratio directly, and once rescaled by the FAUS payout-to-dividend factor for the same window, again applied to equities and Corporate PBW together. Note that the choice of window affects only the conversion of the cyclically-adjusted price-earnings ratio for equities; it leaves the CRSP dividend yields that underlie Corporate PBW unchanged, so Corporate PBW takes its baseline value whenever no payout rescaling is applied.

The durations are only moderately sensitive to the payout definition. The baseline scale of 0.496 implies an average equity duration of 50.5 years under permanent shocks and a Corporate PBW duration of 58.3 years. Applying the FAUS payout-to-dividend factor of 1.139 yields a scale of 0.565 and an equity duration of 44.5 years, about 12% below the baseline. Rescaling the payout of Corporate PBW as well lowers its duration from 58.3 to 51.8 years, a proportional change close to that of equities, as both are priced as perpetuities whose price-payout ratio is scaled by the same factor. Under persistent shocks ($\phi = 0.99$) the differences compress further. Even under the payout adjustment, equities and Corporate PBW remain the longest-duration assets in the household portfolio, so the ordering of asset classes that drives our results is unaffected. Measuring the dividend-to-earnings ratio over the longer 1970–2023 window instead raises the equity duration

[50] Because net equity repurchases are volatile and total payouts can be close to zero in some periods, a period-by-period payout-to-dividend ratio is poorly behaved. We therefore compute the average payout-to-profit ratio and the average dividend-to-profit ratio in the FAUS over the relevant window, and use the ratio of these two averages as the payout-to-dividend rescaling factor: $\frac{\text{Payout/Profit}}{\text{Dividends/Profit}} = \frac{\text{Payout}}{\text{Dividends}}$. This yields a factor of 1.139 for 1983–2023.

Table E.1: Equity Duration under Alternative Payout Scales

	Scale sample	FoF rescale factor	Scale (D/E)	Equity duration $\phi = 1$	Equity duration $\phi = 0.99$	Corp. PBW duration $\phi = 1$	Corp. PBW duration $\phi = 0.99$
Baseline (dividends)	1983–2023	–	0.496	50.5	33.0	58.3	37.3
Rescaled: equities	1983–2023	1.139	0.565	44.5	30.3	58.3	37.3
Rescaled: equities and Corporate PBW	1983–2023	1.139	0.565	44.5	30.3	51.8	34.6
1970–2023 scale	1970–2023	–	0.487	51.5	33.4	58.3	37.3
1970–2023, rescaled: equities and Corporate PBW	1970–2023	1.053	0.512	49.0	32.3	55.6	36.2

Note: The table reports the dividend-to-earnings scale used to convert the cyclically-adjusted price-earnings ratio into a price-payout ratio, and the Financial Accounts payout-to-dividend rescaling factor. The baseline row uses dividends over the 1983–2023 sample. The second row rescales the equity payout to total payout; the third additionally rescales the payout of Corporate PBW, whose duration is likewise built from dividend yields. The last two rows measure the dividend-to-earnings ratio over the longer 1970–2023 window, first directly and then rescaled, in the latter case for equities and Corporate PBW together. Because the window affects only the equity conversion, Corporate PBW takes its baseline value in the rows without a rescaling factor. The last four columns report the average implied duration of equities and of Corporate PBW over 1983–2023, for permanent ($\phi = 1$) and persistent ($\phi = 0.99$) interest rate shocks.

slightly, to 51.5 years. Applying its FAUS factor of 1.053 to equities and Corporate PBW together gives 49.0 and 55.6 years respectively. The window and the payout rescaling therefore push in opposite directions, and this case falls between the baseline and the most conservative one.

Figures E.3 and E.4 replicate the chained repricing exercise of Figure 6 under the baseline and the four alternative cases, for $\phi = 1$ and $\phi = 0.99$ respectively. The implied paths are similar across all five cases, indicating that our results are robust both to the treatment of share buybacks in the measurement of equity and Corporate PBW duration, and to the window over which the payout ratio is measured.

Figure E.3: Implied Top-10% Share under Alternative Payout Scales, $\phi = 1$

(a) Baseline

(b) Rescaled: equities

(c) Rescaled: equities and Corporate PBW

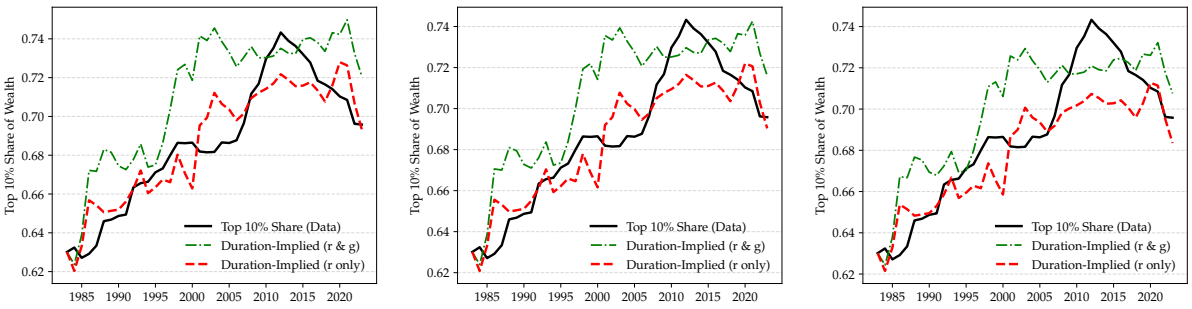

(d) 1970–2023 scale

(e) 1970–2023, rescaled: equities and Corporate PBW

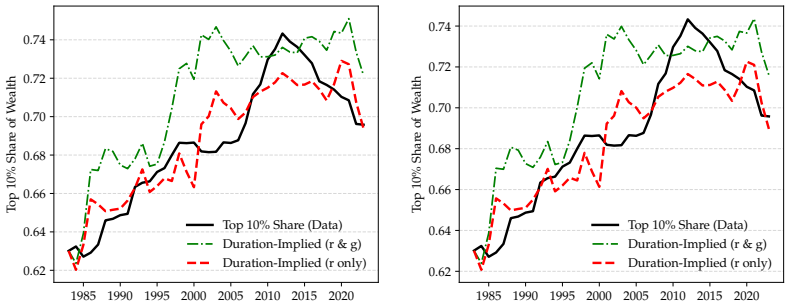

Note: Each panel replicates the chained repricing exercise of Figure 6 under the corresponding case of Table E.1: the baseline dividend-based durations; the durations obtained when the equity payout is rescaled to total payout; those obtained when the payout of Corporate PBW is rescaled as well; and the same two exercises with the dividend-to-earnings ratio measured over the longer 1970–2023 window, used directly and rescaled. Permanent interest rate shocks ($\phi = 1$).

Figure E.4: Implied Top-10% Share under Alternative Payout Scales, $\phi = 0.99$

(a) Baseline

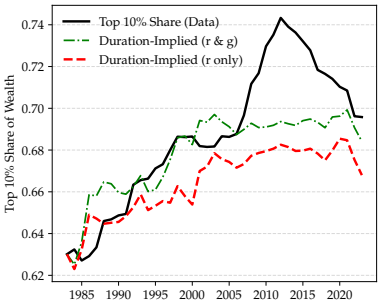

(b) Rescaled: equities

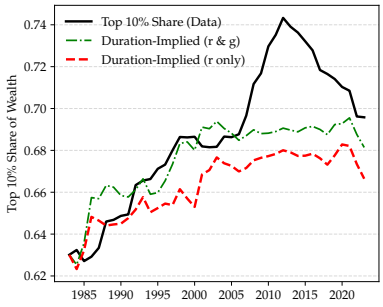

(c) Rescaled: equities and Corporate PBW

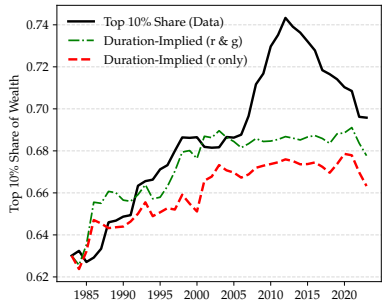

(d) 1970–2023 scale

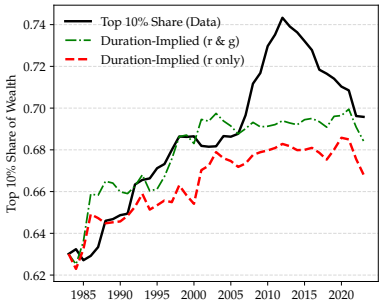

(e) 1970–2023, rescaled: equities and Corporate PBW

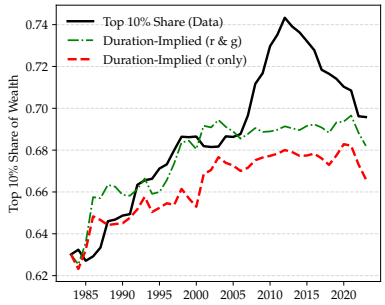

Note: Each panel replicates the chained repricing exercise of Figure 6 under the corresponding case of Table E.1: the baseline dividend-based durations; the durations obtained when the equity payout is rescaled to total payout; those obtained when the payout of Corporate PBW is rescaled as well; and the same two exercises with the dividend-to-earnings ratio measured over the longer 1970–2023 window, used directly and rescaled. Persistent interest rate shocks ($\phi = 0.99$).

E.4 Realized Returns

In Table E.2, we compute the additional returns on wealth experienced by households due to unexpected shocks to interest rates, r, and expected growth, g. We report these returns as arithmetic averages over the full 1983–2023 sample.

Table E.2: Additional Average Realized Returns From 40 Years of Interest Rate and Growth Shocks

	$\phi = 1$		$\phi = 0.99$	
	Baseline	With g	Baseline	With g
P0–10	0.14	0.38	0.11	0.30
P10–20	1.36	1.63	1.20	1.41
P20–30	2.18	1.89	1.85	1.58
P30–40	2.71	2.72	2.30	2.28
P40–50	2.45	2.57	2.08	2.18
P50–60	2.21	2.23	1.87	1.87
P60–70	2.00	2.12	1.68	1.77
P70–80	2.00	2.00	1.65	1.64
P80–90	2.14	2.09	1.70	1.66
P0–P90	2.15	2.17	1.76	1.77
P90–100	3.38	3.66	2.41	2.61

Note: The table displays the additional arithmetic average returns on wealth that households have experienced because of unexpected shocks to interest rates r and expected growth g. The table computes these additional returns for different wealth percentiles under different assumptions. Columns (1) and (2) assume $\phi = 1$: Column (1) uses only changes in r while holding g constant, and Column (2) incorporates changes in both r and g. Columns (3) and (4) replicate the exercise for $\phi = 0.99$.

F Labor Income Process

F.1 Main Specification

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:

$$\log\left(y^i_{t,a}\right) = m_t + \chi' X^i_t + z^i_t, \tag{F.1}$$

$$z^i_{t+1} = \alpha_i + \eta^i_{t+1} + \nu^i_{t+1}, \tag{F.2}$$

$$\eta^i_{t+1} = \rho\eta^i_t + u^i_{t+1}, \tag{F.3}$$

where m_t is a year-fixed effect and X^i_t is a vector of household characteristics that includes a cubic function of age.[51] 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 η^i_{t+1}, and an i.i.d. component ν^i_{t+1}. We set $\mathbb{E}[\nu^i] = E[u^i] = \mathbb{E}[\alpha^i] = \mathbb{E}[\eta^i_0] = 0$, while $Var[\nu^i] = \sigma^2_\nu$, $Var[u^i] = \sigma^2_u$, $Var[\alpha^i] = \sigma^2_\alpha$, and $Var[\eta^i_0] = \sigma^2_{\eta,0}$.

To allow for lower income risk during retirement, we re-estimate (F.1)–(F.3) separately for households above and below age 65, and assume that 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 F.3. Figure F.1 below plots the deterministic life-cycle income profile.

We enrich the regular income process in (F.1)–(F.3) with a superstar income state following Castaneda, Diaz-Gimenez, and Rios-Rull (2003). This state has a high income level $\eta^i_t = \eta^{sup}$; households enter the state with probability p^{sup}_{12}, and exit with probability p^{sup}_{21}. The values $p^{sup}_{12} = 0.0002$ and $p^{sup}_{21} = 0.025$ are taken from Boar and Midrigan (2022), and imply a roughly 1% probability of entering in the superstar income state over one's lifetime. Conditional on entering, the state has an expected duration of 40 years (abstracting from mortality risk). We calibrate the superstar income level η^{sup} so that the top-10% wealth share in the model's stationary economy exactly matches the average top-10% wealth share in the WID data over our 1983–2023 sample. This requires a value equal to 46.64 times average income.

A household first draws whether it transitions into or out of the superstar state using the probabilities stated above. Conditional on being in the regular income state, the household then draws a new value of η. For households beginning and ending the period in the non-superstar

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

state, we approximate (F.2) with a discrete Markov chain P_η using the method of Rouwenhorst (1995). Households exiting the superstar state draw a random η from the ergodic distribution of P_η. Conditional on η, we draw i.i.d. values for ν using Gaussian quadrature.

F.2 Data Source: PSID

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

F.2.1 PSID Income variables

We now describe the construction of the relevant income variables used in the paper. We construct the following variables: *labinc2f* is labor income excluding transfers but including the labor part of business and farm income for both the head and eventual spouse; *transf* is total household transfers (including Social Security Income and other transfers); *labinc3f* is total household income for both the head and eventual spouse, computed as 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 varies with the 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).[52]

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

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

of Business Income from Unincorporated Businesses (2017:ER71302). Note that farm income includes both labor and capital 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*).[53]

F.3 Estimating the Income Process

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

$$\log y_{c,t} = \beta + \gamma_a + \gamma_c + \gamma_t + \varepsilon_{c,t}, \tag{F.4}$$

where the subscript c, t, a refers to cohort, time and age, respectively. We define as c the age at time $t = 0$ (i.e. 1970). Due to the linear relationship between age, cohort and time, we cannot separately identify the different fixed effects. We hence resort to the method used by Deaton and Paxson (1994): we attribute growth to age and cohort effects, while we use the year effects to capture cyclical fluctuations or business-cycle effects that average to zero over the long run. We hence constrain the year fixed effects to be orthogonal to a time trend and to sum to zero. We then estimate Equation F.4 using constrained OLS.

Figure F.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. We use data for households aged 21 to 80. To cover ages outside of our

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

estimation sample, we then back-fill the age profile for age 21 for ages 18–20, and forward-fill the age profile for age 80 to ages 81–100.

Figure F.1: Income Profile

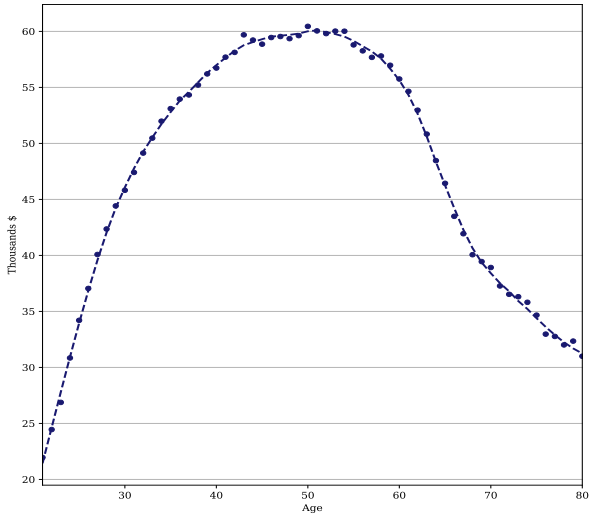

Note: This figure displays the expected income profile evaluated at the 2016 year-fixed effects. The graph plots the expected income profile for the average person who is 21 years old in 2016, expressed in thousands of 2016 dollars. The model is estimated according to Equation (F.4) on PSID data from 1970 to 2017.

Income Risk In the second stage, we estimate income risk. We estimate (F.1) year-by-year, including a cubic function of age as well as fixed effects for education, race, gender, and state. We then extract the residuals z_{it}. Finally we estimate the risk parameters by GMM as detailed below.

Using Equation (F.1)-(F.3), and defining j to be the age of the household minus 20, so that $j = 1$ for the youngest households in our sample, we find that:

$$E[z_j^i z_{j+h}^i] = \begin{cases} \sigma_\alpha^2 + \rho^h E\left[(\eta_j^i)^2\right] & \text{if } h > 0 \\ \sigma_\alpha^2 + E\left[(\eta_j^i)^2\right] + \sigma_\nu^2 & \text{if } h = 0. \end{cases}$$

$$E[\eta_j^{i2}] = \rho^{2j}\sigma_{\eta_0}^2 + \sum_{k=1}^{j} \rho^{2(j-k)}\sigma_u^2.$$

We allow the variance to differ in working age (w) and retirement age (r), where the retirement age starts at 65. We fix the variance of initial persistent shocks $\sigma_{\eta_0}^2 = 0$, then use a GMM estimation to estimate $\theta = (\rho, \sigma_{\nu,w}, \sigma_{u,w}, \sigma_{\nu,r}, \sigma_{u,r}, \sigma_\alpha)$. 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 biennial. To make the estimation consistent, in the first part of the sample 1970–1997 we also sample data at biennial frequency. We only include households whose head is 21 to 80 years old. We only include households that were in the survey for five 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 F.1. These are the parameters used in the main text.

Table F.1: Idiosyncratic Risk Parameter Estimates

	σ_α^2	$\sigma_{\nu,w}^2$	$\sigma_{u,w}^2$	$\sigma_{\nu,r}^2$	$\sigma_{u,r}^2$	ρ
Estimated Parameters	0.0762	0.1605	0.0413	0.0906	0.0255	0.9152

Note: $\theta = (\rho, \sigma_{\nu,w}, \sigma_{u,w}, \sigma_{\nu,r}, \sigma_{u,r}, \sigma_\alpha)$, are estimated using Equation (F.1)-(F.3); $\sigma_{\eta_0}^2$ is fixed equal to 0. Parameters are estimated using data from the PSID over the sample 1970–2017.

F.3.1 Time Varying Income Risk

We also estimate the risk parameters using rolling samples of the PSID from 1983 to 2017. We use samples of 15 years (apart from 1983 where we include data from 1970 to 1983), and forward-fill our final 2017 estimates for years 2018–2023 of our sample. We fix the autocorrelation ρ of persistent shocks to the full sample value estimated in Table F.1. Figure F.2 plots time varying estimates of $\sigma_{\nu,w}^2$, $\sigma_{u,w}^2$, $\sigma_{\nu,r}^2$, $\sigma_{u,r}^2$ and σ_α^2.

Figure F.2: Time Varying Income Risk

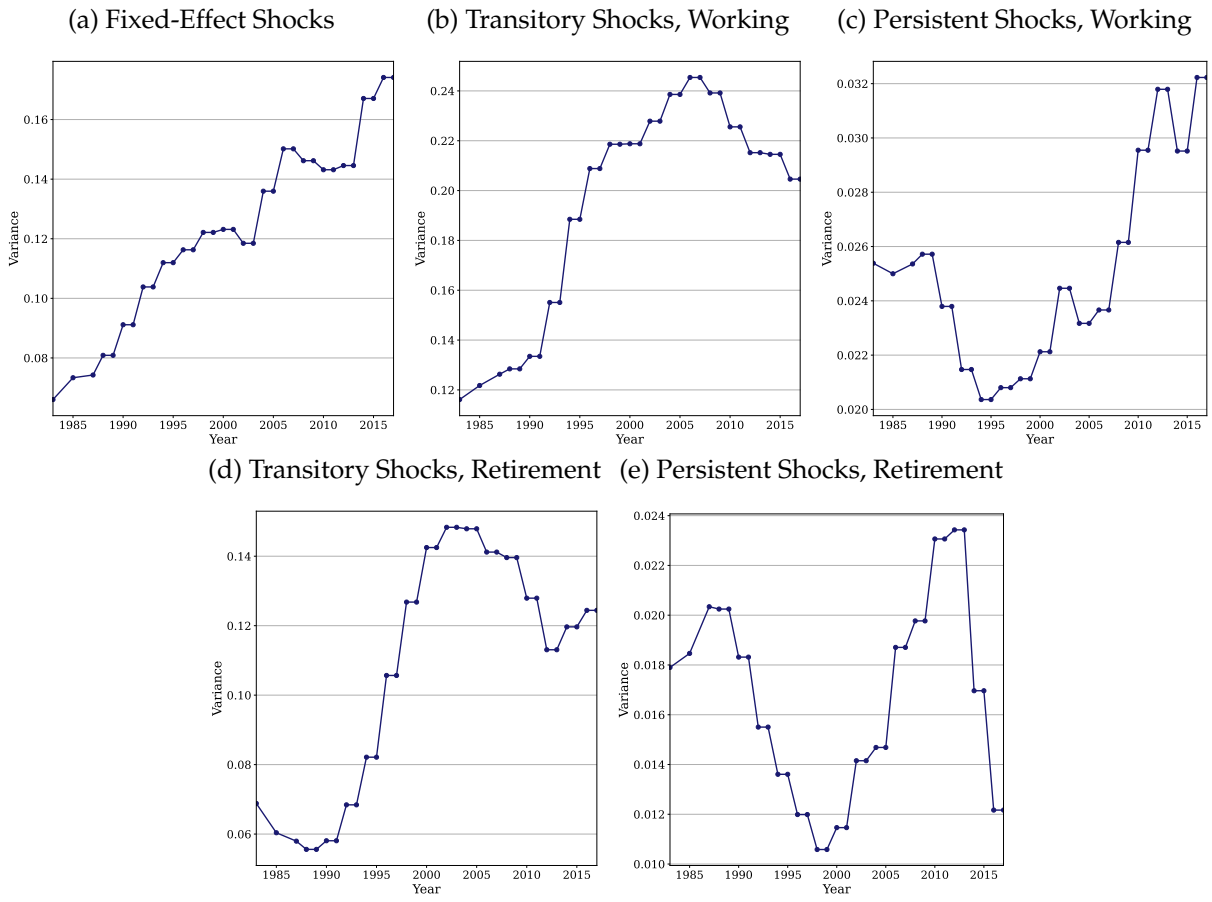

Note: This figure displays the risk parameters estimated from a rolling sample of the PSID. We use samples of 15 years apart from 1983, where we include data from 1970 to 1983. We fix the autocorrelation ρ of persistent shocks to the full sample value estimated in Table F.1. Panel F.2a plots time varying estimates of σ^2_α, F.2b plots time varying estimates of $\sigma^2_{v,w}$, Panel F.2c plots time varying estimates of $\sigma^2_{u,w}$, Panel F.2d plots time varying estimates of $\sigma^2_{v,r}$, and F.2e plots time varying estimates of $\sigma^2_{u,r}$. The model is estimated according to Equation (F.1)-(F.3) on PSID data from 1970 to 2017.

G Additional Model Results

G.1 Additional Derivations

In this section, we demonstrate how to simplify the household's problem using total financial wealth W as a single state variable. Since next period financial wealth W' is given by

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

we can rewrite (22) as the optimization problem

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

subject to the budget constraint

$$c \leq y_j(z) + W - R^{-1} W', \tag{G.1}$$

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

$$c^{-\gamma} = \beta R \left\{ \psi_j \mathbb{E}\left[(c')^{-\gamma} \,\middle|\, z \right] + (1-\psi_j)\chi (W')^{-\gamma} \right\}. \tag{G.2}$$

Although households believe the economy will remain in steady state forever, we apply unanticipated shocks to interest rates that will revalue financial assets. In this case, although households believed (G.1) would hold, the actual next period financial wealth is updated according to

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

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

G.2 Characterizing the Stationary Economy

In this section, we provide more detail on the quantitative properties of the model's stationary distribution under the steady state (average) interest rate. Figure G.1 displays the life cycle profiles of several key variables. The plots are normalized such that 1 represents the unconditional average income. Income displays the traditional hump-shape over the life cycle. Income inequality is increasing over the first half of the life cycle as income shocks accumulate. After retirement, households switch to our estimated over-65 income process which has lower dispersion. While

this compresses income inequality beyond this point, we note that it remains non-negligible since agents have heterogeneous retirement income and still face some income risk after age 65.

Figure G.1: Life Cycle Profiles

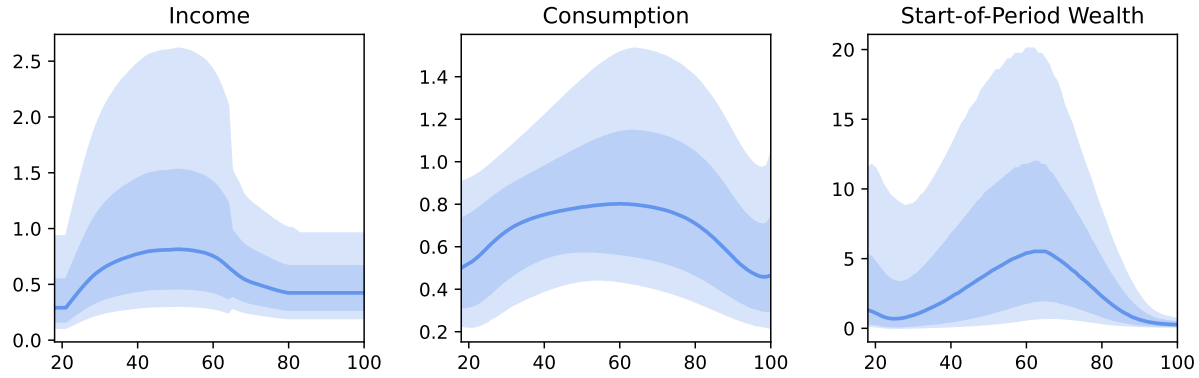

Note: This figure plots life cycle profiles by age. The axes are normalized so that the average income across all agents of all ages is equal to unity. The center line displays the median, while the dark and light bands represent 66.7% and 90% percentile bands.

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

The right panel shows financial wealth, which increases in preparation for retirement, and is subsequently run down during retirement. Financial wealth inequality rises and falls over the life cycle. The first few years are influenced by bequests, which households receive at age 18. On average, households spend these down before beginning to save for retirement closer to age 30, while the large dispersion in the size of bequests generates the initial peak in financial wealth dispersion at the start of the life cycle.[54]

G.3 Model with Non-Permanent Interest Rate Changes

While our benchmark model uses permanent interest rate changes ($\phi = 1$), we can also solve a version of the model that features transitory interest rate changes. This model is largely identical, with the key difference that households now expect future interest rates to differ from the current

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

interest rate due to reversion back to the steady state. In this case, (G.2) becomes

$$c(W,z,R)^{-\gamma} = \beta R \left\{ \psi_j \mathbb{E}\left[c'(W',z',R')^{-\gamma} \,\middle|\, z \right] + (1-\psi_j)\chi(W')^{-\gamma} \right\}. \tag{G.3}$$

This makes explicit that c is a function of (W,z,R), and that next period's consumption c' is a function of (W',z',R'), where

$$\log R' = (1-\phi)\log \bar{R} + \phi \log R. \tag{G.4}$$

Once we solve for the solution $c(W,z,R)$, we apply specific values of R_t at each date in our gradual transition experiment to obtain time-specific policy functions $c_t(W,z) = c(W,z,R_t)$, which anticipate that the interest rate will gradually revert toward steady state going forward.

Our results for this version of the model can be found in Figure G.2, which corresponds to Figure 7 in the main text. As expected given the results in Section 4.2, the cumulated effect of instantaneous revaluations contained in $S_t^{10,rev}$ is smaller under $\phi = 0.99$ in both Panels (a) and (c) compared to our baseline results with $\phi = 1$, reflecting the smaller effects on valuations under transitory rates. However, Panels (a) and (c) show that the total effect of endogenous consumption-savings responses is more positive for the transitory interest rate model compared to our baseline results with $\phi = 1$, leading the total model response to substantially surpass the cumulated revaluation effects in both panels.

To understand this result, we can turn to Panels (b) and (d), which compare the evolution of inequality to the "static" steady state. While our transitory version of the model only features a single true steady state, since the interest rates will revert over time, we can gain intuition by computing an alternative steady state that would hold if the time-t policy functions $c_t(W,z)$ were left in place permanently, which we refer to as the static steady state. As with $\phi = 1$, the initial conditions of the model imply that the top-10% wealth share begins far below this static steady state, implying that the model's mean reversion forces will be pulling inequality upward over time. However, because the smaller revaluations do not push inequality far above the steady state, unlike what we saw for the $\phi = 1$ case, the mean-reversion forces pulling inequality downwards in the second half of the sample are much weaker for $\phi = 0.99$. This leads the total model top-10% share to remain above the Revaluations Only series even at the end of the sample. As a result, the difference in the model's increase in inequality between the $\phi = 1$ and $\phi = 0.99$ top-10% wealth shares is smaller in Table 3 compared to the direct revaluation effects in Table 2.

Finally, we note that for both the models with and without time variation in income inequality, the effect on the top-1% share is larger for $\phi = 0.99$ than for $\phi = 1$. This finding relates to the

Figure G.2: Model-Implied Top-10% Share, $\phi = 0.99$

(a) Model vs. Revaluations, Baseline

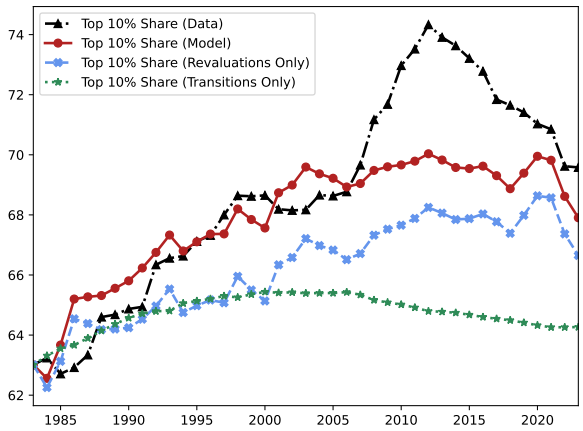

(b) Model vs. "Static" Steady State, Baseline

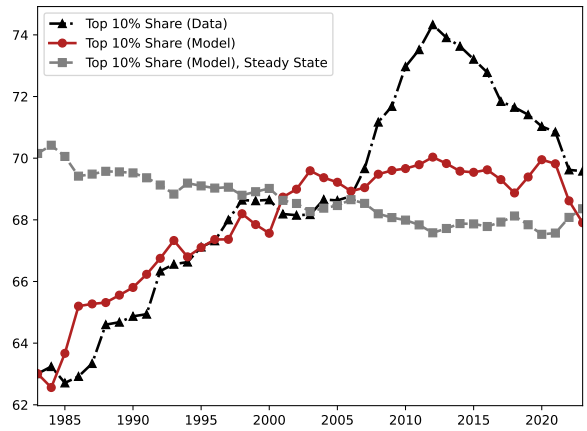

(c) Model vs. Revaluations, TV Inc. Risk

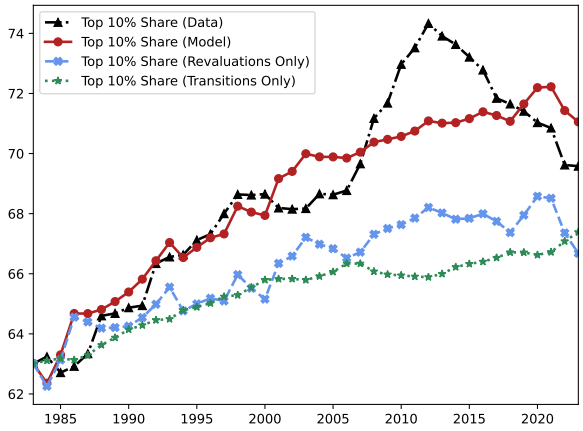

(d) Model vs. "Static" Steady State, TV Inc. Risk

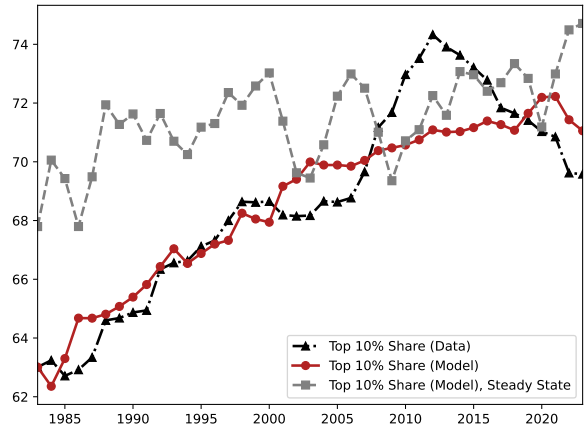

Note: In all panels, the series are defined as follows. "Top 10% Share (Data)" is the top-10% wealth share from the World Inequality Database. "Top 10% Share (Model)" is the model's implied path of the top-10% wealth share from our gradual transition experiment. "Top 10% Share (Revaluations Only)" is the implied contribution of revaluations $S_t^{10,rev}$. "Top 10% Share (Transitions Only)" is the implied contribution of endogenous transitions $S_t^{10,trans}$. "Top 10% Share (Model), Steady State" is the model's steady state top-10% share assuming the current policy functions hold in perpetuity. Panels (a) and (b) plot these series for the baseline model and experiment. Panels (c) and (d) plot these series for the alternative experiment in which we match the relative change in income inequality in each year.

initialization of the model in 1983. Recall that we are choosing an initial increase in the interest rate from steady state in order to match the lower level of inequality in 1983. Because a change in the interest rate of a given size has a smaller effect on inequality due to revaluations in the transitory model, a larger initial shock $\varepsilon_{r,init}$ is required to achieve this target. However, the $\varepsilon_{r,init}$ shock has a larger effect on the top-1% share compared to the top-10% share. As a result, while this larger shock produces the same initial top-10% share in the transitory model compared to the initialization of the permanent model, it generates a larger initial gap between the top-1% share

and its static steady state. As this gap naturally shrinks over time, it leads to a larger total increase in the top-1% share in the transitory model with $\phi = 0.99$ compared to its permanent counterpart.

H International Evidence

Section 7 studies heterogeneity in household portfolio durations in the UK and France and shows that, when combined with the decline in real interest rates, it is an important factor in explaining the rise in the top-10% wealth share in both countries. The section uses household survey data for the UK and France, together with estimated real yields over the period 1983–2023. In this appendix, we provide additional details on the construction of portfolio shares for the UK (Appendix H.1) and France (Appendix H.2). Further information on the estimation of asset and liability durations is given in Appendix H.3. We also describe the data sources and the estimation of real yields for both countries in Appendix H.4.

H.1 UK Wealth and Assets Survey: Data and Variable Construction

For the United Kingdom, we use the Wealth and Assets Survey (WAS), Wave 8 (2020). We harmonize UK survey variables with the asset classes constructed for the United States. All variables are defined at the household level and expressed in nominal pounds. Survey weights are given by `r8xshhwgt`.

H.1.1 Asset Class Definitions

Cash and Deposits. Cash and deposits are defined as the sum of balances held in transaction and savings accounts: current accounts in credit (`dvcacrvalr8_aggr`), savings accounts (`dvsavalr8_aggr`), and cash ISAs (`dvcisavr8_aggr`). This category captures highly liquid household assets.

Equities. Equity holdings include directly held shares, the equity component of mixed investment vehicles, and a fraction of defined contribution pension wealth. Specifically, equities are defined as the sum of UK listed shares (`dvfshukvr8_aggr`), employee shares and stock options (`dvfesharesr8_aggr`), overseas shares (`dvfshosvr8_aggr`), 60% of investment ISAs (`dviisavr8_aggr`), 60% of ISAs of unknown type (`dvkisavr8_aggr`), 60% of unit trusts and mutual funds (`dvfcollvr8_aggr`), 60% of other financial investments (`dvfinvotvr8_aggr`), and 60% of defined contribution pension wealth, constructed as the sum of accumulated DC pension assets (`dvvaldcosr8_aggr`) and additional voluntary contributions (`dvpavcuvr8_aggr`). Unlike the HFCS, the WAS household file does not report private business wealth, so our UK equity measure excludes it. The survey does record the net value of business holdings at the individual level (`BWorthR8`, the value of an individual's share of a business after deducting outstanding debts), but these amounts are small: they are reported by around 1% of individuals and account

for about 1% of aggregate UK household net wealth. Including them would raise the equity share for all households by roughly one percentage point, leaving the comparison with France and the United States unchanged. This contrasts sharply with France, where private business wealth accounts for close to 12% of aggregate net wealth.

Fixed Income. Fixed-income assets include savings products, bonds, and the fixed-income component of mixed investment vehicles, as well as the remaining fraction of defined contribution pension wealth. Specifically, fixed-income assets are defined as the sum of National Savings products (`dvfnsvalr8_aggr`), insurance and life assurance products (`dvflfenvr8_aggr`, `dvflfsivr8_aggr`, `dvflffsvr8_aggr`, `dvflftevr8_aggr`), fixed-term investment bonds (`dvfbondvr8_aggr`), UK government bonds (`dvfgltukvr8_aggr`), foreign government bonds (`dvfgltfovr8_aggr`), 40% of investment ISAs (`dviisavr8_aggr`), 40% of ISAs of unknown type (`dvkisavr8_aggr`), 40% of unit trusts and mutual funds (`dvfcollvr8_aggr`), 40% of other financial investments (`dvfinvotvr8_aggr`), and 40% of defined contribution pension wealth, constructed as the sum of accumulated DC pension assets (`dvvaldcosr8_aggr`) and additional voluntary contributions (`dvpavcuvr8_aggr`).

Real Estate. Real estate wealth is measured as the reported market value of residential property owned by the household (`dvpropertyr8`).

Vehicles. Vehicle wealth is defined as the total reported value of vehicles owned by the household (`dvtotvehvalr8`).

Mortgage Debt. Mortgage debt is measured as the outstanding balance on residential mortgages (`hmortgr8`).

Other Debt. Other debt includes non-mortgage financial liabilities (`hfinlr8_aggr`).

Wealth. Wealth is defined as the sum of all asset classes minus total liabilities.

H.1.2 Rescaling Equities

We find that average equity holdings reported in the survey are substantially lower than aggregate benchmarks. To address this discrepancy, we rescale individual equity holdings to align them with aggregate data. A similar adjustment procedure is employed by Kuhn, Schularick, and Steins (2020) in the construction of the SCF+.

Specifically, we use data from the UK Flow of Funds[55] to compute average household equity holdings. Aggregate equity holdings are constructed as the sum of direct equity holdings and 60% of investment fund shares held by the household sector. We then divide this aggregate value by

[55] Available at https://www.ons.gov.uk/economy/nationalaccounts/uksectoraccounts.

the total number of households in the United Kingdom.[56] This procedure yields average equity holdings of £30,188 per household, compared with £20,258 reported in the survey.[57] We therefore use the ratio of these two values as a scaling factor to rescale individual equity holdings.

This adjustment preserves the cross-sectional distribution of equity wealth within the survey while ensuring consistency with aggregate Flow of Funds benchmarks.

H.2 France Household Finance and Consumption Survey: Data and Variable Construction

For France, we use the Household Finance and Consumption Survey (HFCS), 2021 wave. We harmonize HFCS variables with the asset classes constructed for the UK. All variables are defined at the household level and expressed in nominal euros. Survey weights are given by `HW0010`.

H.2.1 Asset Class Definitions

Cash and Deposits. Cash and deposits are measured by total household deposits (`DA2101`), which include balances held in transaction and savings accounts.

Equities. Equity holdings include directly held publicly traded shares, the equity component of mutual funds and managed accounts, and private business wealth. Specifically, equities are defined as the sum of publicly traded shares (`DA2105`), 60% of mutual funds (`DA2102`) and managed accounts (`DA2106`), and private business wealth, constructed as the sum of non-self-employment private business wealth (`DA2104`) and self-employment business wealth (`DA1140`).

Fixed Income. Fixed-income assets include bonds, voluntary pension and whole life insurance products, and the fixed-income component of mutual funds and managed accounts. Specifically, fixed income is defined as the sum of bonds (`DA2103`), voluntary pension and whole life insurance products (`DA2109`), and 40% of mutual funds (`DA2102`) and managed accounts (`DA2106`).

Real Estate. Real estate wealth is defined as the sum of the value of the household's main residence (`DA1110`) and the value of other real estate property (`DA1120`).

Vehicles. Vehicle wealth is measured as the total reported value of vehicles owned by the household (`DA1130`).

Mortgage Debt. Mortgage debt is defined as the sum of outstanding balances on mortgages for the household's main residence (`DL1110`) and mortgages on other properties (`DL1120`).

[56] Data on the number of households are available at https://www.statista.com/statistics/295545/number-of-households-in-the-uk/.

[57] The survey value of £20,258 includes direct equity holdings and holdings through mutual funds, but excludes indirect holdings through defined contribution pension plans, to ensure consistency with the Flow of Funds definition.

Other Debt. Other debt includes outstanding balances of non-mortgage liabilities (DL1200).
Wealth. Wealth is defined as the sum of all asset classes minus total liabilities.

H.3 Duration Estimates

For the UK and France, we use duration estimates based on the *MSCI World Index*. We download the index from Bloomberg, and obtain a series beginning in 1995. We use its price-earnings ratio and divide it by the average dividend-earnings ratio over the 1995–2023 period, also obtained from Bloomberg. We assume that the price-dividend ratio was at its 1995 level between 1983 and 1995.

For real estate in the UK, we use the price-to-rent ratio from Giglio et al. (2021), who estimate an average rent-to-price of 6.9% over the 1988–2016 period. We estimate a duration of 15.1 years. For France, we calculate a similar rental yield of 6.6% over the sample 1983–2023. We start from the average rental yield for France of 4.6% from the June 2025 edition of Global Property Guide. We then use the rental price index (Fred, ticker FRACP040100IXOBM) and the property price index (Fred, ticker QFRN628BIS) to extend the rent-price ratio backwards in time to 1983. We use the same cost estimate $\tau = 3\%$ as for the US for France and the UK.

For mortgage debt, we use a duration of 0.83 years for the UK. The UK is a predominantly adjustable-rate mortgage (ARM) country. Until 1989, all UK mortgages were ARMs. Fixed-rate mortgages were introduced in 1989, and initially only obtained a small market share of new loans (19% in 1996). By 2010, the FRM share of new loans became 50%, and grew further to 90% in the last decade. Importantly, FRMs in the UK only have a short fixation period, after which they turn into ARMs. The dominant fixation period is 2 years (about 2/3 of FRMs), followed by five years (30% share). Long fixation periods of 7 or 10 years are rare; a 3-year fixation period product also has a small market share. ARMs reset every month, giving them a duration of 0.083. Assuming that 15% of outstanding ARMs refinance and obtain new mortgages according to the new loan product distribution, we find that the duration of the stock of mortgage debt increases from 0.083 in 1983 to 2.15 in 2023. We use the full-sample average of 0.83.

In contrast, in France, the typical mortgage is a fixed-rate mortgage that fully amortizes over 20 years. Given the history of mortgage rates, an average mortgage duration of 5.9 years is appropriate for France. Unlike the US, there is no free mortgage prepayment in France. In each year, we compute the duration of the outstanding mortgage stock as the weighted average durations of each vintage of 20-year FRMs, with weights equal to the share of each vintage in the mortgage stock. To compute these shares, we determine the principal amortization scheme for each vintage, which depends on the mortgage interest rate. We initialize the stock in 1965 and assume that new

mortgages to the tune of 6% of the initial stock are issued each year. The duration of each vintage, and its evolution over time, also depend on the mortgage rate. We use mortgage rates from 1965–2023 from INSEE, Banque de France, and the European Mortgage Federation.

For the other asset and liability classes, which are quantitatively much smaller, we use the same duration estimates as for the US.

H.4 Real Yields for the UK and France

We construct a time series of real interest rates for the UK and France. For the UK, we use the measure of the real yield provided by the Bank of England (BoE). We use the 10-year real yield, which is available since 1985. We extend the series backward using a similar methodology to the US. We regress the 10-year real yield on the 10-year nominal yield (Source: OECD, FRED Code `IRLTLT01GBM156N`) and the inflation rate $\Delta \log CPI_t^{UK}$, where CPI_t^{UK} is the average of the UK CPI price index over the year (Source: OECD, FRED code `GBRCPIALLMINMEI`). For France, we use the same exact approach. We use inflation-protected bond yield available from Bloomberg (`GFRGEN10 Index`) since 2004. We regress this 10-year real yield on the 10-year nominal yield (Source: OECD, FRED Code `IRLTLT01FRM156N`) and the inflation rate $\Delta \log CPI_t^{FR}$, where CPI_t^{FR} is the average of the French CPI price index over the year (Source: OECD, FRED code `FRACPIALLMINMEI`). Specifically, we run the regression

$$r^{10}_{real,j,t} = \beta_{j,0} + \beta_{j,1} r^{10}_{nom,j,t} + \beta_{j,2} \Delta \log CPI_{j,t} + \varepsilon_{j,t}, \tag{H.1}$$

for each country $j \in \{\text{UK}, \text{France}\}$. Based on the estimated coefficients, we construct the time series for the real yield from 1983 to 2023:

$$\hat{r}^{10}_{real,j,t} = \hat{\beta}_{j,0} + \hat{\beta}_{j,1} r^{10}_{nom,t} + \hat{\beta}_{j,2} \Delta \log CPI_t. \tag{H.2}$$

The actual and fitted values are displayed in Figure H.1.

Figure H.1: 10-Year Real Yield, Actual vs. Proxy, UK and France

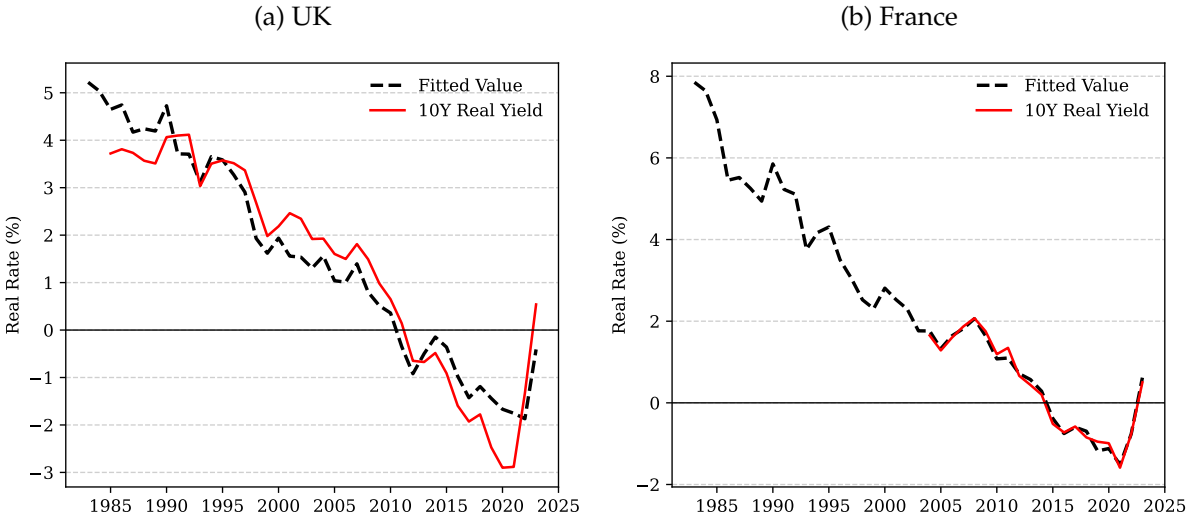

Note: Panel (a) plots the observed 10-year real yield for the UK and its fitted value, which is also used to extrapolate the series backward. Panel (b) presents the corresponding results for France. The fitted value is based on Equation (H.2).