

Heterogeneous Risk Exposure and the Dynamics of Wealth Inequality

Riccardo A. Cioffi *

Paris School of Economics[†]

First version: November 2021

This version: October 2026

Abstract

In this paper I propose a quantitative theory for the dynamics of wealth inequality based on heterogeneous exposure to aggregate risk in asset returns. I develop an optimal portfolio choice model in which housing is a necessary good and households endogenously shift their portfolios towards riskier assets as they get wealthier. Since households are exposed to different sources of aggregate return risk, wealth inequality naturally responds to movements in asset returns. I show that return shocks have large and persistent effects on wealth inequality and that the model replicates the U-shape observed in top U.S. wealth shares since the 1960s.

JEL: D14, D31, E21, E44, E32, G11, G51

*Contact: riccardo.cioffi@psemail.eu

[†]I am indebted to Gianluca Violante, Richard Rogerson, Moritz Lenel, and Benjamin Moll for their continued guidance and support. I would also like to thank Andrea Ajello, Adrien Auclert, Stefano Baratucho, Christian Bayer, Florin Bilbiie, Alberto Bisin, Alex Clymo, Wei Cui, Andreas Fagereng, Ester Faia, Axelle Ferriere, Dan Greenwald, Martin Holm, Jonathan Parker, Jonathan Payne, Lukasz Rachel, Elisa Rubbo, George Sorg-Langhans, Francisco Vazquez-Grande, Maximilian Vogler, Christian Wolf, and Yucheng Yang for their thoughtful comments and discussions, as well as seminar participants at various institutions and conferences. Leonardo Endrizzi and Gustavo Garcia-Bernal provided excellent research assistance.

1. Introduction

One of the key issues in macroeconomics and finance is understanding what shapes wealth inequality and its dynamics. While most countries have extremely high degrees of wealth concentration and have experienced large increases in it since the second half of the twentieth century, we are still missing a theory that can explain both.

In this paper I develop a framework that accounts for the observed dynamics of wealth inequality and that is based on heterogeneous exposure to aggregate risk in asset returns: in the presence of systematic differences in portfolio composition along the wealth distribution, changes in returns naturally induce changes in wealth inequality. I therefore build a heterogeneous-agent, partial equilibrium model of optimal portfolio choice with aggregate risk featuring a rich menu of assets which captures the main properties of households' portfolio shares along the wealth distribution. By matching portfolio heterogeneity, the model is capable of replicating the response of wealth inequality to movements in asset returns found in the data. When fed the realized sequence of asset returns, I show that the model can also replicate the overall U-shape observed in top wealth shares in the U.S. over the post-war period.

When asset returns are subject to aggregate risk, the wealth distribution is stochastic and its law of motion is defined by a *stochastic* Kolmogorov forward equation (KFE). I therefore use the model solution to analyze the response of inequality along its entire ergodic set and to understand how such response depends on the initial state of the economy, as well as on the size and sign of shocks. I also propose a novel decomposition of the KFE that makes it possible to break up the total inequality response in its different parts, separating the direct revaluation effect from the endogenous effects of savings and portfolio choice as well as along intensive and extensive margins and across asset classes. Finally, I use the above machinery to understand how the nature of households' beliefs about the return process affects inequality dynamics.

The model has several implications for the dynamics of wealth inequality: First, I find that shocks to asset returns have large and persistent effects on wealth inequality. In particular, positive equity shocks imply large and very persistent increases in inequality; positive shocks to safe asset returns also raise inequality, albeit to a smaller extent, while a rise in housing returns reduces it. Second, by feeding the realized sequence of U.S. asset returns into the model I find that, even absent any other structural change, the model can explain nearly three quarters of the fall in inequality and about half of the subsequent rise. Third, I decompose the model response and show that most of the inequality response to equity shocks is driven by the endogenous portfolio-choice reaction rather than direct revaluation; when applied to the historical simulation, behavioral channels account for about a third of the model's fall in inequality between 1963 and 1979 and more than half of its subsequent rise between 1985 and 2000. Finally, I find that the nature of the return process has key implications for the evolution of inequality: temporary return shocks generate changes in inequality around eleven times larger than identical but permanent changes in returns; for similar reasons, whether households believe realized return movements to be unexpected windfalls or the result of a repricing of future income flows also affects the overall inequality response.

Existing models of return heterogeneity usually assume return risk to be idiosyncratic, which means that inequality can only move when the economy’s fundamentals change — for instance, when the dispersion of returns shifts (e.g., Gabaix et al. 2016). I instead let return heterogeneity arise endogenously from portfolio choice over a common menu of assets subject to aggregate return risk. As a consequence, realized aggregate returns can move wealth inequality even absent any change in fundamentals. For example, since in the data wealthier households are more exposed to the stock market, top wealth shares rise during stock-market booms and fall during crashes — exactly what has been shown to be the case for the post-war U.S. economy.

In this paper I highlight this specific mechanism: in section 2, I build a partial equilibrium model of wealth inequality that accounts for the role of portfolio choice and asset returns. The model extends the continuous-time heterogeneous-agent framework of Achdou et al. (2022) to a setting in which households can save in three different assets: bonds, housing, and equity. Consistent with data on expenditure shares, I assume that intra-temporal preferences over housing services and non-housing consumption are non-homothetic, and therefore that the composition of households’ consumption baskets varies with the level of total expenditure. This implies that risk tolerance also varies with expenditure, which drives portfolio shares. Since the model is agnostic about where changes in returns are coming from, what matters most is matching households’ observed behavior in response to changes in returns: De La O and Myers (2021) show that survey measures of return expectations barely move when prices rise, and Fagereng et al. (2019) that households hold on to capital gains rather than consuming them. Consistent with these findings, I therefore model asset returns as described by a stochastic process with constant investment opportunities but subject to aggregate return risk. The assumption of constant expected returns means that every movement in realized returns is a one-time windfall gain/loss in wealth: it changes the value of current wealth but carries no information about future returns. I then assess how much this assumption matters for inequality dynamics.

The fact that, in the data, the expenditure share of housing is declining in total consumption indicates that housing is a necessary good and, consequently, that intratemporal preferences are non-homothetic. In section 3 I show that such non-homotheticity causes both the elasticity of intertemporal substitution (EIS) and relative risk aversion (RRA) to change with wealth. By effectively introducing decreasing relative risk aversion, the model generates an optimal share of equity that is increasing in wealth.¹

In section 4 I then calibrate the model and show that it is in fact capable of replicating the observed heterogeneity in households’ portfolios along the wealth distribution, which generates the sort of scale dependence in returns to total wealth observed in the data and contributes to generating a high level of wealth inequality.

In section 5 I show that movements in returns — especially in the equity market — have large and persistent effects on top wealth shares: a one standard deviation increase in equity returns

¹The basic intuition is that, if preferences are non-homothetic, the share of luxury goods in the consumption basket is increasing in wealth and wealthier households care less about consumption fluctuations; see Atkeson and Ogaki (1996); Browning and Crossley (2000) for early formal analyses of this argument.

raises the top 10% share by 0.5 p.p. (0.8%) at its peak, an equally sized shock to safe returns raises it by 0.2 p.p., and a housing shock of the same magnitude instead reduces the top share by -0.4 p.p.; an asymmetry consistent with recent empirical evidence (Bricker et al. 2025). The effect of movements in equity returns in particular is extremely persistent with a half-life of 48.5 years and generates large, slow swings in wealth inequality. I then use the stochastic KFE to compute generalized impulse responses along the model’s stochastic path and show that the response is state-, size-, and sign-dependent: the same equity shock moves inequality more when inequality is already high than when it is low, larger shocks produce smaller responses, and positive shocks imply a slightly smaller response than negative ones.

Since the KFE fully determines the evolution of the distribution, I then provide a novel decomposition that splits the KFE operator into clearly interpretable channels — the direct revaluation from the shock, portfolio choice, and savings — and use it to attribute the inequality response to its various components. When applied to a positive equity shock, such decomposition shows that the model response is mainly driven by the endogenous portfolio-choice margin rather than the direct revaluation itself: as gains accumulate, wealthier households further tilt their portfolios towards equity, and the positive equity premium compounds these larger positions into additional inequality.

Because the model generates fluctuations in wealth inequality even absent any changes in fundamentals, I ask how much of the observed U-shape in U.S. top wealth shares can be explained by fluctuations in realized returns alone. I feed the post-war sequence of U.S. asset returns into the model and find that it delivers 73.7% of the observed decline and 50.8% of the subsequent rise, where each share is measured between the series’ own turning points. I then run a series of counterfactual exercises to understand where these changes are coming from and find that movements in equity and safe returns drive most of the long 1963–79 decline and the 1985–00 rise, where the model shows a remarkable match with the data, while housing returns play a larger role in the two remaining windows (1979–85 and 2000–08), which is broadly consistent with the empirical literature (Kuhn et al. 2020). I also apply the KFE decomposition to the same simulation and find that about a third of the fall in inequality in 1963–79 and more than half of the rise in 1985–00 come from endogenous behavioral responses in savings and portfolio choice.

Importantly, I also show that the assumptions one makes about returns have extremely different implications for inequality: when the model is fed with the exact same sequence of realized returns, either as a long sequence of temporary changes or as a one-time change in expected returns, temporary shocks raise the top 10% share by 5.3%, while a permanent change lowers the top share slightly at first (-0.1%) and in the long run it raises it by only 0.46%. This is because a permanent rise in expected returns also shifts households’ expenditure and saving schedules so that equity holders consume more of their gains while households at the bottom get a further incentive to save and invest into equity; eventually, almost nothing accumulates into additional inequality in the long run. Conversely, temporary shocks leave those schedules unchanged and wealth gains compound as households keep investing them into equity. A similar mechanism would be at play if one were to assume that households accounted for return predictability (e.g., as in Catherine et al. 2026; Greenwald et al. 2025). In that scenario a drop in future expected returns in response to a stock market boom means

wealthy equity holders have an additional incentive to save to smooth consumption while poorer households who do not already hold equity have fewer reasons to enter the equity market, this pushes inequality up further and produces a larger response in inequality to an equivalent equity shock that is fully temporary. I find that the peak response of inequality would be about 9% larger when future expected returns are assumed to negatively covary with realized returns.

Relative to the rest of the literature, in this paper I therefore provide a different interpretation for the observed changes in U.S. wealth inequality over the last 60 years and show that it might have been just the result of chance rather than of structural changes. Such a sharp rise in inequality is in fact consistent with an otherwise unchanged economy in which, among the many possible realizations of asset returns, the observed one happened to be especially favorable to the portfolios of the wealthy. This is not to say that the U.S. economy is the same today as it was 60 years ago, but simply that if we really want to understand the dynamics of inequality we should more closely turn our attention to households' portfolios and asset returns, which is the main objective of this paper.²

RELATION TO LITERATURE. This paper is connected to several strands of literature: First, it relates to models of optimal portfolio choice with incomplete markets, non-homothetic preferences, and housing. In the classic Merton (1971) model, CRRA preferences imply the optimal equity share is independent of wealth; much of the portfolio-choice literature therefore mainly focuses on life-cycle effects. I instead emphasize the direct role of wealth when intra-temporal utility is non-homothetic. In this respect, this paper is closely related to papers with non-homothetic preferences (Atkeson and Ogaki 1996; Blundell et al. 1994; Meeuwis 2020; Pakoš 2011; Wachter and Yogo 2010, among others) and, more specifically, decreasing RRA: Meeuwis (2020) finds substantial evidence for DRRA and non-homotheticity in risk preferences, while Wachter and Yogo (2010) features equity shares increasing in wealth by incorporating non-homothetic utility over basic and luxury composite bundles which also generates decreasing RRA (see also Carroll 2000). Two key differences relative to Wachter and Yogo's framework are that their basic and luxury goods are aggregate bundles with no specific goods attached, and that there is no role for housing as an investment asset.

My paper also naturally relates to papers that match households' portfolio shares in the presence of income risk (e.g. Benzoni et al. 2007; Catherine 2022; Chang et al. 2018; Fagereng et al. 2018). Relative to this literature, I show that once non-homotheticity is accounted for, matching equity shares along the wealth distribution does not necessarily require income risk and equity returns to be correlated.

Because non-homotheticity implies variation in both risk aversion and intertemporal elasticity, my paper is also related to models of preference heterogeneity (Calvet et al. 2021; Gomes and Michaelides 2008; Guvenen 2009). These models typically feature different types of investors with

²When it comes to prices, in this paper I take a very reduced-form approach and simply assume that asset returns follow an exogenous process consistent with households' behavior and calibrated to fit the data. This is because, solving for prices in equilibrium would require the model to simultaneously deliver a realistic characterization of three objects: the joint distribution of wealth and portfolio shares, the equity premium (a well-known challenge in its own right), as well as return fluctuations. Of these three, I deliver the first; solving for the other two jointly is at least a few orders of magnitude more difficult, but something that the literature has very recently started to tackle (albeit in considerably simpler settings, e.g. Cioffi et al. 2026; Yang et al. 2025).

varying degrees of RRA, EIS, or patience, which are usually fixed for a given household and constant in wealth. While my model features heterogeneity in discount rates to help match the concentration in wealth, because agents are otherwise *ex-ante* identical, heterogeneity in RRA and EIS is instead endogenous and my model provides a microfoundation for DRRA.³ Azzalini et al. (2023) also use types with different degrees of patience and risk aversion to generate an empirically realistic gradient of portfolio shares, whereas here the same pattern arises endogenously as wealth grows.⁴

Since the focus is on portfolio shares and wealth inequality rather than on the housing market, the housing problem as specified in my model is deliberately simple; hence, the most related papers in that literature are early contributions as in Cocco (2005); Flavin and Nakagawa (2008); Flavin and Yamashita (2002); Piazzesi et al. (2007).⁵ Of these the closest is Cocco (2005) in which, however, housing preferences are homothetic and housing crowds out equity only through indivisibility and participation costs.⁶

Second, my paper connects to the recent literature on return heterogeneity as a key driver of wealth concentration. Relative to existing contributions (see the empirical work in Bach et al. 2020; Daminato and Pistaferri 2024; Fagereng et al. 2020; and model-based analyses in Benhabib et al. 2019; 2011; Gabaix et al. 2016; Hubmer et al. 2021; Xavier 2020; see also Atkeson and Irie 2022; Kacperczyk et al. 2019; Kasa and Lei 2018; Kaymak et al. 2026), my focus is specifically on households' *optimal* portfolio choice as the driver of return heterogeneity.⁷ Relatedly, Gomez (2025) shows that when one group of agents is constrained to hold levered equity positions, aggregate shocks generate long-lasting movements in top wealth shares, though his exposure model "cannot fully capture" the U-shape in wealth inequality observed in the data. Relative to his paper, the heterogeneity in exposure to aggregate shocks in my model is instead the outcome of an optimal portfolio choice decision. To the best of my knowledge, this paper is the first to characterize and solve the full *stochastic* Kolmogorov forward equation for the cross-sectional distribution in a continuous-time setting with optimal portfolio choice.⁸

Finally, my paper also speaks to the literature linking the evolution of the wealth distribution to

³Bach et al. (2020) suggest DRRA as one of two candidate explanations for the scale dependence in returns.

⁴Belda et al. (2025) offer an alternative microfoundation for the same pattern, based instead on heterogeneous beliefs measured using survey data on households' stock-market expectations. Relative to that framework, the specific advantage of non-homothetic preferences is that these can be directly disciplined using households' expenditure data.

⁵See Piazzesi and Schneider (2016) for a survey of this literature.

⁶While housing has been shown to also shape stock-market participation through mortgage debt and home equity (Chetty et al. 2017; Vestman 2019), I abstract from mortgages to simplify the numerical solution of my model. In a more recent contribution calibrated to Swedish administrative data, Rácz (2025) also shows that the interaction of housing and portfolio choice helps generate an increasing risky share because housing crowds out risky financial assets for households with a lower ratio of financial wealth to human wealth.

⁷Gabaix et al. (2016) offer an early theoretical treatment of how return heterogeneity shapes the *dynamics* of wealth inequality. While their main focus is on the dynamics of *income* inequality, their analysis applies identically to the case of wealth, as they also show in appendix E. Gabaix et al. specifically express the hope that "other models will be developed that generate scale dependence, calibrated on dynamic micro data" (Gabaix et al. 2016, p. 2098); this paper provides such a model: one in which scale dependence in returns to wealth arises endogenously from households' optimal portfolio choice.

⁸A number of very recent papers now tackle the general-equilibrium version of this problem to solve for asset prices and endogenous portfolios with the wealth distribution as a state variable when the latter is stochastic (Cioffi et al. 2026; Gopalakrishna et al. 2026; Yang et al. 2025). As argued above this is an extremely complicated problem to solve, and these other papers therefore operate in considerably simpler settings.

movements in asset returns. My findings are consistent with empirical work relating changes in wealth inequality to booms and busts in equity and housing markets (Kuhn et al. 2020; Martínez-Toledano 2020; see also Fagereng et al. 2019; Hyland and Khemka 2025; Mian et al. 2020). Recent evidence at business-cycle frequency also finds that equity-price shocks raise measured wealth inequality while house-price shocks reduce it (Bricker et al. 2025), an asymmetry that the model specifically replicates. More tightly connected are papers that explain the observed rise in U.S. wealth inequality through changes in asset returns: in a model without housing Favilukis (2013) shows that the stock market boom helps explain the observed rise in inequality; Hubmer et al. (2021) run a similar exercise to mine and feed the realized sequence of asset returns into a model of portfolio heterogeneity and find that asset-return changes are key to explaining the U-shape of top wealth shares but “have also dampened the increase in wealth concentration on net, in particular explaining the initial dip” (Hubmer et al. 2021, p. 430). My paper departs from Hubmer et al. in two key respects: First, their portfolio weights are fixed functions of total wealth taken from the data. Second, they assume all return movements reflect changes in expected returns and feed the model an anticipated, smooth path for returns. In section 5 I show that these distinctions matter: optimal portfolios generate significant amplification in response to return changes, and temporary changes in returns have much larger effects on inequality than permanent ones.

The distinction between temporary and permanent changes in returns also differentiates this paper from Gomez and Gouin-Bonenfant (2024) and Greenwald et al. (2025). Gomez and Gouin-Bonenfant argue that a decline in interest rates sustains the growth of new fortunes by making it cheaper for entrepreneurs to raise capital, while Greenwald et al. argue that it raises measured wealth inequality through portfolio revaluation, as the wealthy tilt toward longer-duration, equity-like assets whose value rises when rates fall (a mechanism similar to the one at play in my model).

Finally, since this paper was first circulated in 2021, several papers have turned to the joint dynamics of household portfolios, asset returns, and wealth concentration (e.g., Blanchet 2022; Catherine et al. 2026; Cioffi et al. 2026; Fernández-Villaverde and Levintal 2026; Mijakovic 2025). Of these, the closest is Catherine et al., which features a similar model but, using the same history of returns, reaches a different conclusion regarding the main drivers of inequality: because in their model a fall in interest rates also raises stock and house prices, their decomposition credits the rate-driven part of those capital gains to movements in discount rates, and most of the post-war fall and rise in top wealth shares is then explained by the path of real rates. In this paper, instead, I assume that expected returns are constant and that realized returns vary only through temporary innovations; hence changes in inequality are attributed to movements in the specific asset markets themselves. This assumption is consistent with survey measures of investors’ expectations (De La O and Myers 2021) and with how households respond to capital gains (Fagereng et al. 2019); I will later return to why such assumption is appropriate for my model. Finally, none of the above papers characterizes the full stochastic evolution of the wealth distribution under aggregate return risk, the stochastic Kolmogorov forward equation dynamics at the center of this paper.

2. Model

The model is set in continuous-time and is a partial-equilibrium version of the standard heterogeneous-agent model in macroeconomics in which households face idiosyncratic income risk and smooth consumption over time by saving in a menu of assets (e.g. Achdou et al. 2022). When compared to the rest of the literature on wealth inequality, the model directly innovates in the richness of households' portfolio choice.

The model's purpose is to propose a parsimonious yet realistic characterization of households' optimal portfolio choice along the wealth distribution: to generate an increasing equity share along the wealth distribution, the model features both participation costs in the equity market and non-homotheticity in preferences; to capture the main features of the housing market, I account for both homeownership decisions and the illiquid nature of housing.

2.1. Main Environment

DEMOGRAPHICS. There is a continuum of households facing a constant death rate ζ . Population is constant and normalized to 1 so every time an agent dies she is immediately replaced by a newborn agent according to distribution Ψ , to be specified later. I assume that households differ in their patience: each newborn draws a permanent discount type $m \in \{1, \dots, N_\rho\}$ from a distribution π_ρ and keeps it for life, so that ρ_m is fixed at birth.⁹

PREFERENCES. Households have recursive inter-temporal preferences as in Duffie and Epstein (1992):

$$v_t = \mathbb{E}_t \left[\int_t^\infty f(u_s, v_s) ds \right] \quad (1)$$

where $f(u, v)$ aggregates intra-temporal utility u and continuation utility v and takes the form:

$$f(u, v) = \frac{\rho_m + \zeta}{1 - \psi^{-1}} (1 - \gamma) v \left\{ \left[\frac{u}{((1 - \gamma)v)^{\frac{1}{1-\gamma}}} \right]^{1-\psi^{-1}} - 1 \right\} \quad (2)$$

where ρ_m is the rate of time preference for a household of type m , and γ and ψ govern risk aversion and elasticity of intertemporal substitution, respectively.

Intra-temporal utility u is an aggregator of housing services n and non-housing consumption c . The aggregator is of the addi-log form as in Wachter and Yogo (2010):

$$u(c, n) = \left((1 - \omega) \frac{1 - \varepsilon_h^{-1}}{1 - \varepsilon_c^{-1}} c^{\frac{\varepsilon_c - 1}{\varepsilon_c}} + \omega n^{\frac{\varepsilon_h - 1}{\varepsilon_h}} \right)^{\frac{\varepsilon_h}{\varepsilon_h - 1}} \quad (3)$$

where ω captures the relative preference of housing services relative to non-housing consumption.

⁹Discount-rate heterogeneity is a standard mechanism in the literature to help match wealth concentration (Krusell and Smith 1998). In section 4 I discuss in more detail how this feature of the model helps alleviate some of the tensions in matching key data moments while remaining mostly orthogonal to the main model mechanism.

When $\varepsilon_h = \varepsilon_c$, equation (3) nests the constant elasticity of substitution (CES) case; however, whenever $\varepsilon_h \neq \varepsilon_c$, preferences are non-homothetic. In particular, under the assumption that $\varepsilon_h < \varepsilon_c$, housing is a necessary good and non-housing consumption is a luxury good. As we will see later, the assumption that housing is a necessary good is well supported by data on consumption expenditure and has significant implications for portfolio choice. In fact, the non-homotheticity assumption implies that the composition of the consumption basket changes with total expenditure. Because households are less willing to tolerate fluctuations in consumption when the share of housing consumption is higher, both relative risk aversion and elasticity of intertemporal substitution vary along the expenditure distribution, which in turn affects optimal allocations along the expenditure and wealth distributions. I return to this in section 3.

ASSET STRUCTURE. Households can smooth consumption over time by trading in financial assets (bonds and stocks), and housing. Denote by $dr_{j,t}$ the asset j 's instantaneous return over a time interval of length dt , with $j \in \{B, S, H\}$ and collect them in a vector $d\mathbf{r}_t = (dr_{B,t}, dr_{S,t}, dr_{H,t})$.¹⁰ Assuming return innovations are correlated across assets and expected returns are constant, we can write instantaneous returns as:

$$d\mathbf{r}_t = \mathbf{r} dt + \boldsymbol{\sigma} d\mathbf{W}_t \quad (4)$$

where $\mathbf{r} = (r_B, r_S, r_H)$ is the vector of expected returns, $\mathbf{W}_t = (W_{1,t}, W_{2,t}, W_{3,t})$ is a vector of three mutually independent standard Brownian motions, and $\boldsymbol{\sigma}$ is the 3×3 matrix of sensitivities of returns to the three independent shocks. The return variances and covariances are the entries of $\boldsymbol{\sigma}\boldsymbol{\sigma}'$: I write $\sigma_j^2 \equiv (\boldsymbol{\sigma}\boldsymbol{\sigma}')_{jj}$ for the variance of return j and $\sigma_{jk} \equiv (\boldsymbol{\sigma}\boldsymbol{\sigma}')_{jk}$ for the covariance between returns j and k , with $j, k \in \{B, S, H\}$ so that, for example, σ_S^2 is the equity-return variance and σ_{SH} the equity-housing return covariance.

The assumption of constant expected returns requires some additional discussion. The key observation is that this is a model of households' behavior and its main objective is therefore matching their consumption-saving and portfolio-choice dynamics; this is *not* a model of asset pricing. Returns enter the problem exogenously and, as a consequence, the main criteria for choosing a return process are to have a process that is consistent with observed data and, most importantly, that leads households in the model to behave as they do in the data. In this regard, administrative data from Norway show that households hold on to capital gains rather than consuming them (Fagereng et al. 2019), and survey measures of investors' expectations move little when prices rise (De La O and Myers 2021). Hence, under rational expectations, a process in which future expected returns negatively comove with asset prices would attribute to households a response not shown in the data, thereby potentially also getting wrong the dynamics of inequality. In my model, the assumption of constant expected returns means that every realized gain or loss is effectively a one-time windfall: it changes the value of current wealth but carries no news about the returns that wealth will earn going forward.¹¹ I return to this assumption in section 5, where I measure how the model's conclusions

¹⁰Consistent with the fact that households directly get utility from how much housing they own, we ignore the rental yield component in the returns to housing. $dr_{H,t}$ therefore only captures the capital gain component of housing returns.

¹¹If asset prices were instead equilibrium objects, their data-generating process would matter directly to also get asset pricing behavior right; in this partial-equilibrium setting, however, it matters only through the behavior it implies for

would change if the same gains were instead followed by lower expected returns going forward (e.g., driven by changes in discount rates).

EARNINGS. At each instant, households receive an endowment of earnings subject to idiosyncratic risk. Following Kaplan et al. (2018), I model log-earnings as the sum of two independent components — a persistent component $z_{1,i,t}$ and a transitory component $z_{2,i,t}$, with $z_{i,t} = z_{1,i,t} + z_{2,i,t}$ — each of which follows a jump-drift process:

$$dz_{j,i,t} = -\beta_j z_{j,i,t} dt + dJ_{j,i,t}, \quad j \in \{1, 2\} \quad (5)$$

where $\beta_j > 0$ is the mean-reversion rate, $J_{j,i,t}$ is a Poisson process with arrival rate λ_j , and conditional on a jump the process resets to a new level z'_j from $\mathcal{N}(0, \sigma_j^2)$. Such process matches the persistent-transitory specification that is commonly used in the literature together with the evidence provided by Guvenen et al. (2021) on non-Gaussian earnings dynamics.¹²

FRICTIONS. Households face a participation decision in both housing and equity markets. When it comes to housing, they can be owners or renters. Owners derive utility directly from the amount of housing they own with multiplier χ , while renters are free to choose every period how much housing services to consume. With respect to the equity market, participants can freely adjust their portfolio of financial assets between bonds and stocks at all times, while non-participants will have 100% of their financial assets invested in bonds.

The entry and exit structure for both equity and housing markets is similar and has two components: a stochastic component that governs movements across participation states, and a transaction-cost component. In particular, with respect to the housing market, households get a chance to buy or sell at constant rate λ^h ; then, those who take that opportunity choose their optimal housing size h subject to a transaction cost $\kappa^h(h) = \kappa_0^h + \kappa_1^h h$. Infrequent trading, as captured by the λ^h arrival rate, has the dual role of capturing search and matching frictions in the housing market as well as simplifying the numerical computation of the model.

Similarly, in the equity markets, non-participants get an opportunity to enter the equity market at constant rate λ_0^p , at which point they can choose to pay an entry cost κ_0^p and become a participant. To stay in the market, participants also have to pay a (yearly) participation fee κ_1^p at a constant rate λ_1^p ; if they choose not to pay the participation fee, they exit the equity market and become non-participants. One can think of κ_0^p as the cost of setting up a trading account, and of κ_1^p as the per-period maintenance fees of such account.¹³

Finally, households are only allowed to borrow in bonds up to an exogenous borrowing limit ϕ and, when doing so, face a wedge κ^b that increases the effective interest rate on borrowing.¹⁴

households.

¹²Since expected returns are constant and all movements in realized returns are entirely driven by temporary shocks, I abstract from any correlation between asset returns and the earnings distribution.

¹³The stochastic component of entry and exit decision in the equity market here serves mainly the purpose of not having to solve a much harder optimal stopping time problem. As I will show in section 4, the values of λ_0^p and λ_1^p will be chosen to have minimal economic implications.

¹⁴The model's main focus is on the interaction between households' portfolio choice and wealth inequality, and not

HOUSEHOLD PROBLEM. At each instant, all households choose the optimal level of non-housing consumption c , and renters get to choose how much housing services to consume each period, n , at rental rate r_n ; participants additionally choose the optimal equity-share θ .

I denote renters as households with zero housing, $h = 0$, the participation state in the equity market as $p \in \{0, 1\}$, and collapse financial assets into a single state variable a . Hence financial wealth, a , and housing, h , evolve according to:

$$da_{i,t} = \left(e^{z_{i,t}} + (r_B + \zeta)a_{i,t} + (r_S - r_B)\theta_{i,t}p_{i,t}a_{i,t} - x_{i,t} - \kappa^b(a_{i,t}) \right) dt + a_{i,t}\sigma_a d\mathbf{W}_t \quad (6)$$

$$dh_{i,t} = h_{i,t}r_H dt + h_{i,t}\sigma_3 d\mathbf{W}_t \quad (7)$$

where $x_{i,t} = c_{i,t} + r_n n_{i,t} \mathbf{1}\{h_{i,t} = 0\}$ is total consumption expenditure in non-housing consumption and rental services, $\kappa^b(a) = \kappa^b \max\{(\theta p - 1)a, 0\}$ is the additional cost of borrowing, proportional to the amount borrowed in bonds, $(\theta p - 1)a$, and zero otherwise, σ_i is the i -th row of the σ matrix in equation (4), and $\sigma_a = (1 - \theta_{i,t}p_{i,t})\sigma_1 + \theta_{i,t}p_{i,t}\sigma_2$ is the row vector of sensitivities of financial wealth to the three independent shocks, which directly depends on the optimal equity share $\theta_{i,t}$ and the participation state $p_{i,t}$. Finally, the drift on a carries an annuity premium ζ , which is a rebate of the deceased households' liquid wealth to survivors in proportion to their own holdings.

It will be useful to collect the household's individual state in the vector $\mathbf{x} \equiv (a, h, z_1, z_2, p, \rho_m)$.¹⁵ Exposure to aggregate shocks $d\mathbf{W}_t$ is collected in the $\sigma_{\mathbf{x}}(\mathbf{x})$ matrix, whose rows are the loadings of the continuous states (a, h, z_1, z_2) on the three independent shocks:

$$\sigma_{\mathbf{x}}(\mathbf{x}_{i,t}) = \begin{bmatrix} a_{i,t} \sigma_a \\ h_{i,t} \sigma_3 \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \quad (8)$$

where the two zero rows reflect that the earnings components carry no diffusion. If we denote by $v(\mathbf{x})$ the household's value function, solving the household's problem boils down to finding the solution to the following HJB equation:

$$0 = \max \{ f(u, v) + \mathcal{A}v \} \quad (9)$$

where $\mathcal{A}$ is the infinitesimal generator of $\mathbf{x}$ and the choice set for the max operator depends on their current participation and homeownership state. The exact definition of $\mathcal{A}$ is left to appendix A.

CROSS-SECTIONAL DISTRIBUTION. Given the assumption of exogenous returns and aggregate return risk, characterizing our model economy amounts to describing two things: households' behav-

on providing a comprehensive treatment of the housing market. I therefore take some necessary shortcuts to keep the model tractable, in particular with respect to borrowing: notice in fact that households cannot use housing as collateral and can only borrow in short-term bonds. While mortgages are obviously a crucial feature of the housing market, I leave them out because they would make the household problem considerably harder to solve.

¹⁵Note that, while only total earnings $z = z_1 + z_2$ matter for households' decisions, in order to write the problem in recursive form we need to keep track of its individual components. This is because, even though both z_1 and z_2 are Markov, their sum is not itself a Markov process as they have different degrees of mean reversion.

ior — which is characterized by the HJB equation (9) — and the cross-sectional distribution $g_t(\mathbf{x})$, whose evolution is determined by the following proposition:¹⁶

Proposition 1. *The distribution of households over individual states, $g_t(\mathbf{x})$, solves the following Kolmogorov Forward Equation (KFE):*

$$dg_t(\mathbf{x}) = \left\{ \mathcal{A}^* g_t(\mathbf{x}) + \zeta \left(\Psi_t(\mathbf{x}) - g_t(\mathbf{x}) \right) \right\} dt - \partial_{\mathbf{x}} \left\{ \left[\boldsymbol{\sigma}_{\mathbf{x}}(\mathbf{x}) d\mathbf{W}_t \right] g_t(\mathbf{x}) \right\} \quad (10)$$

where $\mathcal{A}^*$ is the adjoint of the differential operator in the HJB equation (9), evaluated at households' optimal policies.

The KF equation (10) is a *stochastic* partial differential equation and it highlights that, unlike in models of idiosyncratic return risk, the cross-sectional distribution g_t is itself stochastic. The KFE (10) is composed of a drift term — which is present in all continuous-time models *à la* Achdou et al. (2022) — that essentially determines the average level of inequality, and a diffusion term — which is only present in models with aggregate risk — that specifies how the evolution of the cross-sectional distribution g_t depends both on the specific path of shocks $d\mathbf{W}_t$ and, crucially, on the exposure of households' portfolios to aggregate shocks through $\boldsymbol{\sigma}_{\mathbf{x}}(\mathbf{x})$. Specifically, the second term only appears because the aggregate shocks *directly* affect households' state variables in equations (6) and (7).

BIRTH DISTRIBUTION. Finally, to close the model, we are only left to characterize the birth distribution $\Psi_t(\mathbf{x}) = \Psi(\mathbf{x}, g_t)$. Consistent with the existence of annuity markets for liquid wealth, I assume that all households are born as non-participants with no liquid wealth ($a = 0$). Housing on the other hand is fully passed down from parents to children, while earnings and the discount-rate type are drawn from their stationary distributions. The asymmetry between the treatment of liquid assets and housing wealth for inheritance follows from the assumption of incomplete annuity markets. Unlike liquid financial wealth, housing cannot realistically be annuitized (Pashchenko 2013) and is therefore passed on directly to the next generation.¹⁷

2.2. Numerical Solution

The model is solved by discretizing the state vector $\mathbf{x}$ and approximating both the HJB (9) and KFE (10) using finite differences. With respect to the HJB, solving the household problem in equation (9) essentially boils down to writing a discretization of the infinitesimal operator $\mathcal{A}$; then, with that in hand, the drift term of the Kolmogorov forward equation (10) comes for free, as discretizing $\mathcal{A}^*$ only involves taking the transpose of the discretization of $\mathcal{A}$. With respect to the KFE, I also discretize the diffusion term in (10) using finite differences; then, solving the KFE simply requires stepping forward in time starting from an initial distribution g_0 . A detailed description of the numerical procedure is available in online appendix F.

¹⁶The proof of proposition 1 is fairly standard in the applied math literature and I therefore leave it to online appendix D. For a comprehensive treatment and derivation of the KF equation in the context of mean-field games with common noise, see instead Carmona and Delarue (2018); Carmona et al. (2016).

¹⁷Additionally, under complete annuity markets (Blanchard 1985) the entire wealth of the deceased, including their house, would be rebated as cash to surviving households, which would artificially distort the wealth distribution.

Unless stated otherwise, the model is solved by simulating a random sequence of returns from equation (4) at a quarterly frequency. The simulation runs for 21,000 years, of which the first 1,000 years are discarded to remove dependence on the initial condition. The initial distribution g_0 is the risky steady state of the model, which I define below.

3. The Role of Non-homothetic Preferences

We saw in proposition 1 that the evolution of the wealth distribution g_t depends on both the specific path of aggregate shocks $d\mathbf{W}_t$ and on households' exposure to such shocks (through $\sigma_{\mathbf{x}}$). In this section I show how non-homothetic preferences make the elasticity of intertemporal substitution (EIS) and relative risk aversion (RRA) vary along the wealth distribution, thereby shaping households' portfolio composition and, through it, their exposure to aggregate risk.

Early arguments for why non-homothetic preferences generate variation in RRA and EIS can be found in Atkeson and Ogaki (1996); Browning and Crossley (2000) and the intuition is straightforward: because wealthier households allocate a larger share of their consumption basket to luxury goods (here, non-housing consumption), and because households care less about fluctuations in the consumption of luxuries, the elasticity of intertemporal substitution rises with wealth. Additionally, since this is true both across states and over time, the same logic extends to risk aversion and Wachter and Yogo (2010) in fact exploit exactly this channel to generate RRA that varies with wealth. Under recursive preferences RRA and EIS are naturally decoupled, so the two must be characterized independently. The results below extend the above logic to a model with recursive preferences.

To do so, consider a simplified version of the model: a modified Merton (1971) economy with recursive intertemporal preferences and a non-homothetic intratemporal aggregator. Specifically, assume that households live forever ($\zeta = 0$), have zero earnings ($e^{z_{i,t}} = 0$), and can invest only in two assets: riskless bonds ($\sigma_1 = \mathbf{0}$) and a single risky asset. Further assume that households are not allowed to own housing but continue to rent housing services n at rate r_n , so that the two-good intratemporal problem of equation (3) is unaffected.¹⁸ Finally, assume households face no other frictions: all households participate in the stock market and there are no restrictions on borrowing.

Let $V(a)$ denote the value function over wealth a in this simplified economy, and let $x^*(a)$ denote the associated optimal policy for total expenditure. I first characterize how expenditure shares change along the expenditure distribution, then turn to the elasticity of intertemporal substitution and, finally, to relative risk aversion.¹⁹

The following proposition shows that, as long as $\varepsilon_h < \varepsilon_c$, the expenditure share in housing declines in total expenditure and that housing is a necessary good.²⁰

¹⁸While this section formally focuses only to the renters' problem, since in the full model homeowners have the opportunity to infrequently adjust their housing consumption, it will nonetheless provide a useful benchmark to understand the mechanism behind non-homothetic preferences in the quantitative model.

¹⁹The proofs for this section are mostly algebra but can be quite involved, so I leave them to online appendix D.

²⁰In the calibration section I will be more specific about how to pin down ε_h and ε_c using data on expenditure and show that $\varepsilon_h < \varepsilon_c$ is indeed the empirically relevant assumption.

Proposition 2. Define $\alpha(x) \equiv c(x)/x$ the expenditure share in non-housing consumption, where $x = c + r_n n$. Then, if $0 < \varepsilon_h < \varepsilon_c < 1$:

$$\frac{\partial \alpha(x)}{\partial x} > 0 \quad \forall x.$$

Furthermore:

$$\alpha_0 \equiv \lim_{x \rightarrow 0} \alpha(x) = 0; \quad \alpha_\infty \equiv \lim_{x \rightarrow \infty} \alpha(x) = 1.$$

Before moving to RRA and the EIS it is useful to first define a few objects we will be using in what follows. For a twice-differentiable function F of a scalar argument s , let $\eta_F(s) \equiv sF'(s)/F(s)$ denote its elasticity and $\mathcal{R}_F(s) \equiv -sF''(s)/F'(s)$ its Arrow–Pratt curvature. Let $U(x) \equiv u(c(x), n(x))$ denote indirect intratemporal utility along the optimal intratemporal choice. Finally, define $A \equiv \frac{1-\varepsilon_c^{-1}}{1-\varepsilon_h^{-1}}$ as an (inverse) measure of the degree of non-homotheticity in preferences: under our assumption that $0 < \varepsilon_h < \varepsilon_c < 1$, we have that $A \in (0, 1)$ and that the closer it is to one, the closer preferences are to the homothetic benchmark (CES). As we are about to see, A also determines the limiting values of both the EIS and RRA below.

The following proposition characterizes the behavior of the EIS and shows that it increases with total expenditure and has well-defined limits that depend on the distance between ε_h and ε_c :

Proposition 3. The elasticity of intertemporal substitution is given by:

$$EIS(x) \equiv -\frac{\partial \left(\frac{d \ln x_t}{dt} \right)}{\partial \left(\frac{d \ln \Lambda_t}{dt} \right)} = \frac{\alpha(x)\varepsilon_c + (1 - \alpha(x))\varepsilon_h}{1 - \eta_U(x) (\alpha(x)\varepsilon_c + (1 - \alpha(x))\varepsilon_h) \left(\frac{1}{\varepsilon_h} - \frac{1}{\psi} \right)}$$

where $\Lambda_t = e^{\int_0^t f_v(u_s, v_s) ds} f_u(u_t, v_t) U'(x_t)$ is the utility gradient. Then, if $\psi < \varepsilon_h < \varepsilon_c < 1$, we have that $\frac{\partial EIS}{\partial x} > 0$ and

$$\psi_0 \equiv \lim_{x \rightarrow 0} EIS(x) = \psi; \quad \psi_\infty \equiv \lim_{x \rightarrow \infty} EIS(x) = \frac{1}{1 + A \left(\frac{1}{\psi} - 1 \right)} > \psi_0$$

Obtaining a similar result for risk aversion is slightly more involved and requires one more object. In particular, it is useful to denote as $H(x) \equiv U(x)^{1-\gamma}/(1-\gamma)$ the “risk transform” of U , which is what actually enters the recursive certainty equivalent term, so that $\mathcal{R}_H(x)$ is its curvature, which I call the “Hanoch curvature” after Hanoch (1977). The following lemma derives an analytical expression for the risk aversion coefficient:

Lemma 1. Let $\mathcal{R}_V(a) \equiv -aV''(a)/V'(a)$ denote the Arrow–Pratt coefficient of relative risk aversion of the value function $V(a)$. Then, if $\gamma > 1$, $\psi > 0$, and $\psi \neq 1$:

$$\mathcal{R}_V(a) = \eta_x(a) \mathcal{R}_H(x^*(a)) + \left(\gamma - \frac{1}{\psi} \right) \left[\frac{\eta_V(a)}{1-\gamma} - \eta_x(a) \eta_U(x^*(a)) \right]. \quad (11)$$

The first term in equation (11), $\eta_x(a) \mathcal{R}_H(x^*(a))$, is a *static* risk-aversion channel which arises simply from differentiating the consumption first-order condition with respect to wealth: $\mathcal{R}_H(x^*)$

is the risk aversion that an agent with utility H would display when facing gambles over total expenditure in a static problem, while η_x simply translates the curvature in expenditure $\mathcal{R}_H$ into curvature in wealth $\mathcal{R}_V$. In fact, $\mathcal{R}_H$ can also be decomposed as $\mathcal{R}_H(x) = \gamma \eta_U(x) + \mathcal{R}_U(x)$, which makes clear how $\mathcal{R}_H$ includes the curvature of the intratemporal indirect utility, $\mathcal{R}_U$, together with the curvature contributed by the γ power in the expression for H . As we will see momentarily, $\mathcal{R}_H$ and $\mathcal{R}_V$ coincide in the limits as expenditure goes to 0 or ∞ , which makes it worthwhile to characterize the Hanoch curvature on its own.

The second term is instead a *dynamic* risk-aversion channel which arises because, with recursive preferences, the first-order condition also picks up the dependence of the marginal utility of expenditure on the continuation value. Note that the term in square brackets can be written as:

$$\frac{\eta_V(a)}{1-\gamma} - \eta_x(a) \eta_U(x^*(a)) = \frac{1}{1-\gamma} \left[\frac{d \ln((1-\gamma)V(a))}{d \ln a} - \frac{d \ln((1-\gamma)H(x^*(a)))}{d \ln a} \right],$$

which captures the gap between the wealth-elasticity of lifetime utility V and the wealth-elasticity of flow utility H . When lifetime utility is proportional to flow utility, as is the case under homothetic preferences, the two scale together and the bracket equals zero. The dynamic channel is therefore zero in two special cases: under time-separable preferences — which implies households do not care about the timing of the resolution of uncertainty and risk aversion collapses to $\mathcal{R}_V = \eta_x \mathcal{R}_H$ — and under homothetic preferences — which imply risk aversion is independent of expenditure, $\eta_x = \eta_U = 1$ and $\eta_V = 1 - \gamma$, so that $\mathcal{R}_V = \mathcal{R}_H = \gamma$.

Figuring out the sign of the derivative of $\mathcal{R}_V$ would therefore require characterizing both $\eta_x(a)$ and the dynamic channel at every wealth level: under recursive non-homothetic preferences this is considerably involved and I therefore take a less direct approach. In fact, as the corollary to the next proposition shows, characterizing $\mathcal{R}_H$ still provides relevant information about RRA.

Proposition 4. *The Hanoch curvature $\mathcal{R}_H(x)$ is given by:*

$$\mathcal{R}_H(x) = \frac{\alpha(x)\varepsilon_c \left(1 - \varepsilon_h^{-1}\right) \left(1 + \frac{1-\varepsilon_c^{-1}}{1-\varepsilon_h^{-1}}(\gamma - 1)\right) + (1 - \alpha(x))\varepsilon_h \left(1 - \varepsilon_c^{-1}\right) \gamma}{\left[\alpha(x)\varepsilon_c + (1 - \alpha(x))\varepsilon_h\right] \left[\alpha(x) \left(1 - \varepsilon_h^{-1}\right) + (1 - \alpha(x)) \left(1 - \varepsilon_c^{-1}\right)\right]}$$

Then, if $0 < \varepsilon_h < \varepsilon_c < 1 < \gamma$ and $\gamma > \varepsilon_c/\varepsilon_h$, we have that $\frac{\partial \mathcal{R}_H}{\partial x} < 0$ and

$$\gamma_0 \equiv \lim_{x \rightarrow 0} \mathcal{R}_H(x) = \gamma; \quad \gamma_\infty \equiv \lim_{x \rightarrow \infty} \mathcal{R}_H(x) = 1 + A(\gamma - 1) < \gamma_0$$

Corollary 4.1. *Under the hypotheses of lemma 1 and proposition 4, $\psi < \varepsilon_h$, a strictly positive equity premium, and regularity conditions stated in online appendix D, true relative risk aversion converges, at each end of the wealth distribution, to the corresponding limit of the Hanoch curvature:*

$$\lim_{a \rightarrow 0} \mathcal{R}_V(a) = \lim_{x \rightarrow 0} \mathcal{R}_H(x) = \gamma_0 \quad \text{and} \quad \lim_{a \rightarrow \infty} \mathcal{R}_V(a) = \lim_{x \rightarrow \infty} \mathcal{R}_H(x) = \gamma_\infty.$$

In particular, $\mathcal{R}_V(a) - \mathcal{R}_H(x^*(a)) \rightarrow 0$ at both ends.

Hence, relative risk aversion exactly coincides with the Hanoch curvature at the two limits, γ_0 and γ_∞ . Intuitively, the asymptotic limits of $\mathcal{R}_H$ coincide with the Arrow–Pratt curvature of the recursive-utility value function in wealth because at these limits the intratemporal aggregator reduces to a single-good power utility. Conversely, away from those limits the two measures will generally differ because the dynamic channel is nonzero and because the wealth-elasticity η_x might differ from one, so that curvature in expenditure does not translate one-for-one into curvature in wealth.²¹

In the quantitative model of section 2, poorer households devote a larger share of their consumption basket to the necessary good (housing services) and therefore have higher RRA and lower EIS. Hence, poorer households have a high preference for a smooth consumption profile and, since equity is risky and housing illiquid, they prefer to invest in safer assets. On the other hand, wealthier households consume primarily luxury goods and care less about consumption fluctuations; they therefore prefer to invest in equity to capture the higher return.

4. Model Calibration

The model is calibrated to match key features of the U.S. cross-sectional distribution, with special attention given to matching the cross-sectional distribution of portfolio shares. Some parameters are calibrated outside the model, in particular the processes for asset prices and earnings, while others are internally calibrated by directly comparing moments from the model to the data.

Table 1 collects the externally calibrated parameters, with the exception of asset prices and the earnings process which are reported in tables 2 and 3, while the internally calibrated parameters are reported in table 5. Median earnings are normalized to one and the unit of time is one year, all parameter values are expressed accordingly.

4.1. External

ASSET PRICES. Asset prices are calibrated using data from 1964 to 2020 from the MacroHistory Database of Jordà et al. (2019) and adjusted for inflation. I take the safe return to be the average of the returns on Treasury bills and government bonds. For housing, since the model separates between renters and owners, I only use the capital gain component of total housing returns; these are then adjusted for leverage and costs of mortgages by using the average loan-to-value ratio and the average cost for 30-year fixed rate mortgages — consistent with housing portfolio shares being calibrated on households’ net home equity position; online appendix C provides further details on such adjustment. Since in the data most of the observed variation in the aggregate rent-to-price ratio comes from movements in prices, I set the rental rate r_n to the average rental returns over the period from the same data (5.16%).

²¹Relative to the results in Wachter and Yogo (2010), the expression for $\mathcal{R}_H$ and the limiting values of γ_0 and γ_∞ exactly coincide with the static risk-aversion formula therein, while the monotonicity condition and the connection to risk aversion in the presence of recursive preferences as in lemma 1 and corollary 4.1 are new results.

Table 1: Externally Calibrated Parameters

PARAMETER	VALUE	TARGET / SOURCE
<i>Demographics</i>		
ζ	1/40	avg. working life
<i>Taxes</i>		
τ	0.181	tax progressivity (Heathcote et al. 2017)
<i>Borrowing frictions</i>		
ϕ	0.75	median credit limit
κ^b	0.06	Kaplan et al. (2018)
<i>Housing frictions</i>		
λ^h	5.2	10 wks. avg. housing search time (Garriga and Hedlund 2020)
κ_0^h	0.05	housing transaction costs (Favilukis et al. 2016)
κ_1^h	0.055	"
<i>Equity market frictions</i>		
λ_0^p	12	see text
λ_1^p	365	"
κ_0^p	0	"

Table 2: Asset Returns (Yearly, Real)

	MEAN	STD. DEV.	CORRELATIONS	
			Safe	Equity
Safe	0.025	0.058	.	.
Equity	0.073	0.157	0.12	.
Housing	0.007	0.062	-0.14	0.14

Table 2 reports the estimated average returns, implied volatilities, and pairwise correlations of the three assets; the resulting matrix of sensitivities σ is reported in online appendix C.

EARNINGS. I calibrate the jump-drift earnings process in equation (5) following Kaplan et al. (2018), by targeting moments of annual male pre-tax earnings from the administrative records of Guvenen et al. (2021) and adjusting for household earnings as the unit of observation using the Survey of Consumer Finances (SCF). Specifically, I calibrate the process to match the standard deviation and kurtosis of one- and five-year earnings changes, the fraction of one-year changes at several thresholds as measured in Guvenen et al., and adjust to match the cross-sectional dispersion of household earnings in level as measured in the SCF.

To obtain post-tax earnings, I then apply the progressive tax schedule of Heathcote et al. (2017) to pre-tax earnings in levels. I set the progressivity parameter τ to 0.181 and the level parameter λ to match an implied average tax rate of 28% (Heathcote et al. 2017; Trabandt and Uhlig 2011). Table 3 reports the calibrated parameter values and table 4 the targeted moments and the model's fit, where the model-implied moments are computed from the pre-tax Monte Carlo simulation of the

continuous process. Online Appendix E.1 reports further details on the calibration.

The calibrated earnings process has similar properties to the one in Kaplan et al. (2018): shocks to the first component arrive infrequently, on average once every 20 years, and are extremely persistent, while shocks to the second component arrive on average once a year and mean-revert much faster.

Table 3: Earnings Process: Calibrated Parameters (Yearly)

PARAMETER	PERSISTENT	TRANSITORY
Arrival rate λ_j	0.0518	0.9350
Mean-reversion β_j	0.0152	11.0711
Jump size σ_j	1.246	1.495

Table 4: Earnings Process: Moment Fit

MOMENT	DATA	MODEL	SOURCE
<i>Cross-section</i>			
Std. dev.	1.027	1.016	SCF
<i>One-year changes</i>			
Std. dev.	0.510	0.490	Gucenen et al. (2021)
Kurtosis	14.930	14.626	"
Frac. $ \Delta z < 5\%$	0.306	0.313	"
Frac. $ \Delta z < 10\%$	0.488	0.485	"
Frac. $ \Delta z < 20\%$	0.665	0.646	"
<i>Five-year changes</i>			
Std. dev.	0.780	0.817	"
Kurtosis	9.510	9.713	"

BORROWING AND HOUSING FRICTIONS. The exogenous borrowing limit ϕ is set to 75% of median earnings, approximately corresponding to the median credit limit in the U.S. I take the value of κ^b directly from Kaplan et al. (2018).²² Transaction costs in the housing markets are taken directly from Favilukis et al. (2016), who also study a quantitative model with illiquid housing: variable costs κ_1^h are set to 5.5% of the housing value, while the fixed cost is set to 0.05 (which approximately corresponds to their 3.2% of average consumption). The buying and selling frequency λ^h is set to match an average search time in the housing market of 10 weeks (Garriga and Hedlund 2020).

EQUITY-MARKET FRICTIONS. Since there is no clear data moment that can help setting the stochastic components governing entry and exit in the equity market, I set them so as to minimize their economic implications: the frequency of the entry shock λ_0^p is set such that no more than 1%

²²Kaplan et al. (2018) calibrate the corresponding parameter internally to match the share of households with negative net liquid assets. However, in our model the share of households with negative net worth — which is a targeted moment in the internal calibration described below — is mainly driven by the EIS and by ρ_3 (the discount rate of the most impatient households), which makes κ^b harder to identify independently.

of households would ever want to be a stockholder but is forbidden to do so because they did not get an opportunity to enter, while the frequency for the per-period participation cost λ_1^p is simply set so as to arrive on average once a day (therefore mimicking a continuous-time cost structure).

With respect to the entry and participation costs, κ_0^p and κ_1^p , Vissing-Jørgensen (2002) finds evidence for both an entry cost and a per-period participation cost and shows that the per-period cost rationalizes the nonparticipation of half of stock-market nonparticipants. As we will see below, the calibrated value for the degree of risk aversion is already fairly high, which means that a recurring cost alone can account for the degree of nonparticipation observed in the data. I therefore set the entry cost κ_0^p to zero and calibrate only the per-period cost κ_1^p , keeping free parameters to a minimum.

4.2. Internal

I choose the remaining parameters by minimizing a weighted sum of squared relative deviations between model and data moments across the 18 targeted moments listed in table 6.

On the model side, the most obvious candidate distribution for computing model moments is the long-run cross-sectional distribution under aggregate shocks $\bar{g} = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T g_t dt$. Because the KF equation (10) is linear in g_t , this can be found by solving the *mean* KFE, which retains the second-order Itô correction in the drift term but sets the realized shocks to zero. I define this distribution as the *risky steady state* of the model and compute all calibration moments on it; because it can be found directly as the solution to a stationary problem, this has the advantage of not requiring any long simulation and keeps the calibration routine manageable.²³ Other objects of interest, e.g. generalized impulse response functions, will be computed along the *stochastic steady state*, which I define as the ergodic distribution of the random density g_t along the simulated path. Online Appendix F.3 details how I solve for the risky and stochastic steady states.

Correspondingly, on the data side, I also build moments from long-run averages. All data targets are therefore computed as averages across all SCF waves from 1989 to 2022, making the data moments the empirical counterpart of the stationary objects the calibration is matching.

Weights are chosen to achieve a better match for the moments that are central to the paper’s mechanism: portfolio-composition, especially at the top of the wealth distribution, and the participation margins in both equity and housing markets. I use a parallelized multistart algorithm which first evaluates a large set of Sobol points over the relevant parameter domain, and then starts local searches from the most promising initial points using the DFO-LS algorithm of Cartis et al. (2019), a derivative-free algorithm tailored to least-squared objectives.²⁴ I describe in detail the calibration algorithm and the construction of moments in online appendix E.2 and online appendix C.

Table 5 collects the full vector of 10 internally calibrated parameters, and table 6 the model’s fit relative to the 18 targeted moments. In what follows I discuss each group of moments in turn.

²³Notice that, because inequality moments are nonlinear (e.g., top shares and the Gini coefficient) these will differ when computed at $\bar{g}$ rather than averaged along the simulated path. In section 5 I show that the differences are small.

²⁴The algorithm resembles the TikTak routine of Arnoud et al. (2022) while exploiting parallelization as much as possible to reduce total running time.

Table 5: Internally Calibrated Parameters

PARAMETER	VALUE	TARGET	SOURCE
ρ	0.037, 0.161, 2.295	wealth-to-earnings ratios	SCF
γ	10.11	portfolio shares	"
ψ	0.53	share with negative net worth	"
χ	2.44	avg. homeownership rate	"
ω	0.29	avg. expenditure share	CEX
ε_h	0.856	expenditure shares	"
ε_c	0.897	avg. expenditure elasticity	Aguiar and Bils (2015)
κ_1^p	0.004	avg. participation	SCF

Note: The “Target” column is only suggestive of which moment is the most directly related to each parameter, as in the actual calibration routine all parameters affect all moments simultaneously. The implied values for the non-homotheticity parameters are $A = 0.681$, $\gamma_\infty = 7.21$, and $\psi_\infty = 0.62$.

Table 6: Matched Moments

MOMENT	DATA	MODEL	WEIGHT
<i>Participation</i>			
Homeownership rate	0.681	0.583	4
Equity participation	0.518	0.509	3
<i>Portfolio shares</i>			
Safe, top 10%	0.254	0.271	3
Equity, top 10%	0.356	0.388	3
Safe, middle 40%	0.180	0.142	2
Housing, middle 40%	0.614	0.532	2
Safe, bottom 50%	-0.236	-0.220	2
Housing, bottom 50%	1.021	0.978	2
<i>Housing expenditure</i>			
Avg. housing expenditure share	0.290	0.224	1
Avg. expenditure elasticity	0.920	0.964	1
Q5/Q1 ratio	0.750	0.549	1
<i>Wealth-to-earnings</i>			
Aggregate	5.401	4.601	1
Middle 40%	3.647	3.948	1
Top 10%	13.544	15.329	1
<i>Wealth inequality</i>			
Share with negative net worth	0.229	0.261	1
Top 10% wealth share	0.703	0.584	1
Top 1% wealth share	0.302	0.160	1
Gini coefficient	0.813	0.745	1

Note: Data moments are from the SCF (1989–2022) except the expenditure moments, which are from the CEX, and the average expenditure elasticity, which is from Aguiar and Bils (2015). Weights are the relative weights in the calibration loss.

WEALTH-TO-EARNINGS RATIOS. To calibrate the ρ vector note that, on top of generating additional wealth inequality, discount-rate heterogeneity decouples wealth accumulation at the top from the aggregate: the most patient type is disproportionately concentrated at the top and keeps the wealth-to-earnings ratio at the top of the wealth distribution in line with the data, while the more impatient types keep aggregate wealth holdings in check. In this respect, group-level ratios are therefore the natural moments to target to discipline the ρ values: in the SCF, the wealth-to-earnings ratio increases sharply with wealth (from roughly 3.6 for the middle 40% to 13.5 for the top 10%, against an aggregate of about 5.4). Hence, I calibrate the vector $\rho = (\rho_1, \rho_2, \rho_3)$ to match the wealth-to-earnings ratio in the overall economy as well as within the top 10% and middle 40% groups, as measured in the SCF. I correspondingly set π_ρ to (0.1, 0.4, 0.5) to match these three groups.²⁵ This also helps alleviate another tension in the model between matching portfolio shares at the top (and therefore wealth inequality) and wealth-to-earnings ratio. Specifically, a single discount-rate model could either match the aggregate wealth-to-earnings ratio or top equity shares, but not both: this is because matching top equity shares in a single discount-rate model would cause too much accumulation at the top, driving up the aggregate wealth-to-earnings ratio.²⁶

WEALTH INEQUALITY. I target four moments related to wealth inequality: the fraction of households with negative net worth, the share of wealth held by the top 10% and top 1%, and the Gini coefficient. The fraction of households with negative net worth is directly tied to the EIS parameter ψ ; while the others mainly discipline the dispersion of discount rates. The model generates a fraction of households with negative net worth of 0.26, which matches the data fraction of 0.23 fairly well.

The inequality moments come in somewhat below their data counterparts. There are two main reasons that help explain this discrepancy: First, target moments are averages over the pooled 1989–2022 SCF waves, a period in which U.S. wealth concentration was especially high, while calibration moments are taken from the cross-sectional distribution in the risky steady state; in the model inequality fluctuates over time and actually spans the targeted values along the ergodic distribution, with the top 10% wealth share ranging from 50.0% to 77.9% and the Gini coefficient from 0.686 to 0.852. Second, beyond return heterogeneity — which is pinned down by design by the portfolio-share targets — discount rate heterogeneity is the main mechanism generating inequality. However, as discussed above, these rates are mainly pinned down by the wealth-to-earnings targets; leaving little room for the model to exactly match inequality. The model would plausibly deliver a higher level of inequality if we were to give larger weight to the inequality targets, but at the cost of a poorer fit to the wealth-to-earnings ratios.²⁷

²⁵While there is obviously going to be a positive correlation between households’ discount-rate types and their position in the wealth distribution, these are not fixed groups and households in the top 10% of the wealth distribution do not necessarily belong to the lowest discount-rate group (e.g. due to their specific idiosyncratic earnings history).

²⁶While the calibrated discount rate for the most impatient type is quite high, it is indeed the natural outcome given the calibration targets: the aggregate wealth-to-earnings ratio, together with the ratios for the top 10% and middle 40% jointly pin down the implied ratio for the bottom 50% too, which in the SCF is very close to zero (0.04). As reported in table B.2, the most impatient type holds 4.5% of total wealth and only 0.2% of wealth in the top decile; hence, a high discount rate for the most impatient households is key to hit the wealth inequality moments but determines very little of behavior at the top of the distribution, which is what this paper focuses on.

²⁷Additionally, the model does not include standard ingredients used to match wealth inequality (e.g., bequest motives

HOUSING EXPENDITURE. The theoretical results derived in section 3 provide useful guidance for the calibration of the intra-temporal preference parameters $(\omega, \varepsilon_h, \varepsilon_c)$. In fact, while ω can naturally be used to pin down the average share of consumption in housing services, whenever $\varepsilon_h < \varepsilon_c$, proposition 2 implies that the expenditure share of housing is declining in total expenditure. Furthermore, the expenditure elasticity of housing consumption is given by:

$$\frac{\partial \log n}{\partial \log x} = \frac{\varepsilon_h}{\alpha(x)\varepsilon_c + (1 - \alpha(x))\varepsilon_h}. \quad (12)$$

which, as households get wealthier, gradually decreases from 1 to $\varepsilon_h/\varepsilon_c < 1$. This suggests that two moments informative about the non-homotheticity parameters are the average expenditure elasticity of housing and the distribution of expenditure shares along the expenditure distribution.

Hence, I ask the model to match three targets related to housing expenditure: the average expenditure elasticity, the average expenditure share, and the ratio between average expenditure shares at the top and bottom quintiles of the expenditure distribution, all estimated from the Consumer Expenditure Survey (CEX) using the same methodology and sample as in Aguiar and Bils (2015).²⁸ As reported in table 6, the model generates an average elasticity that is a bit too high, and expenditure shares that are slightly too low and decline more than in the data. These gaps arise because a single pair of parameters $(\varepsilon_c, \varepsilon_h)$ simultaneously pins down the housing Engel curve derived in proposition 2 and equation (12) as well as the wealth profile of both risk aversion and EIS characterized in propositions 3 and 4. Because housing is the model's only source of non-homotheticity, this single channel has to generate two distinct cross-sectional patterns at once: the decline of the housing expenditure share along the expenditure distribution, and the rise of risky portfolio shares along the wealth distribution. This generates a tension between fitting the housing Engel curve and fitting the wealth gradient in portfolio shares so that the model cannot match both sets of targets exactly. I resolve this trade-off by overweighting the portfolio-share moments in the loss function which are the key determinant for the evolution of wealth inequality and main calibration objective. This allows the model to closely match portfolio shares along the wealth distribution, at the expense of a less precise fit of expenditure shares. Both sets of cross-sectional moments are shown in figure 1.²⁹

The resulting calibrated values of ε_h and ε_c are 0.856 and 0.897, respectively, with an implied value for the non-homotheticity parameter $A = \frac{1-\varepsilon_c^{-1}}{1-\varepsilon_h^{-1}}$ of approximately 0.681 suggesting that, although the average expenditure elasticity of housing is close to 1, preferences remain meaningfully non-homothetic. Consistent with the pattern of expenditure shares shown in figure 1a, even at $A = 0.681$, I still find a sizable variation in relative risk aversion and EIS across the wealth distribution. The implied range for the EIS from section 3 would be $\psi_0 = 0.53$ to $\psi_\infty = 0.62$, which is well within the range of cross-sectional estimates in Calvet et al. (2021), who report a mean EIS of 0.99 (median 0.42,

or entrepreneurship) and it is therefore not surprising the model falls somewhat short of fully matching the data.

²⁸Housing expenditure is computed using actual rent paid for renters and (self-reported) rental equivalence for home owners (see Aguiar and Bils (2015) for more details about data definitions).

²⁹A richer demand system with several non-homothetic goods, each loading differently on effective risk aversion, would decouple the housing Engel curve from the wealth gradient in risk attitudes and reconcile these moments. However, this would make the model much less tractable for relatively little gain in terms of the model's main results.


standard deviation 0.96). On the other hand the implied Hanoch curvature falls with wealth from $\gamma_0 = 10.11$ to $\gamma_\infty = 7.21$ and is above the mean estimate of risk aversion 5.24 (median 5.30, standard deviation of 0.47) in Calvet et al. (2021). The calibrated value for the risk aversion parameter is at the high-end of values used in the portfolio-choice literature and is required to keep equity holdings at the top of the wealth distribution closer to the data. Lower risk aversion would push all households to hold more equity than in the data, which would therefore mechanically amplify the model’s implied transmission from stock returns to wealth inequality. As we will see in section 5, stock returns turn out to be the main determinant of the evolution of inequality and, by keeping the top equity share low, a high γ calibration is therefore a conservative choice for the paper’s main mechanism.

PORTFOLIO SHARES. The primary objective of the calibration is to match households’ portfolio shares along the wealth distribution. I therefore directly target the composition of household portfolios across the same three main groups used to calibrate discount-rate heterogeneity: the bottom 50%, the middle 40%, and the top 10%. To compute the target shares I use data from the SCF to aggregate assets into three classes that map directly onto the model’s three assets. Risky assets capture public equity: stocks, stock mutual funds, IRAs/Keoghs invested in stocks, and other managed assets with an equity interest, held either directly or indirectly. I specifically exclude private business equity from the underlying wealth concept altogether: the model has no private businesses and private equity — which is highly concentrated at the top, very illiquid, and partly idiosyncratic (Kartashova 2014; Moskowitz and Vissing-Jørgensen 2002; Quadrini 1999) — is very different from public equity and cannot simply be folded into the model’s risky asset. Private business assets and liabilities are therefore completely removed from the definition of wealth, which should therefore be interpreted as total wealth ex-business. After removing asset categories with no model counterpart, listed in table C.1, safe assets are defined as the residual of total financial assets net of equity and net of consumer debt (non-housing debt). Households’ housing position is computed as the value of residential and nonresidential real estate net of housing debt; this is consistent with the model only allowing for short-term borrowing and no mortgages. Online Appendix C provides more details on the construction of each data series and all its components.

Figure 1b plots portfolio holdings in the SCF pooling all waves from 1989 to 2022 and shows a clear pattern along the wealth distribution: the bottom 50% hold negative safe assets (−0.24); housing is *the* asset of the bottom 90% (with a share of 1.02 for the bottom 50% and 0.61 for the middle 40%); while equity becomes a major portfolio component only at the top (the top 10% hold an average equity share of 0.36). The model directly targets safe shares across all three groups, the average housing share for the bottom two groups, and the equity share for the top group. As shown in figure 1b the model closely reproduces this pattern across the three groups.

PARTICIPATION. The last two targeted moments, the average equity-market participation rate and the homeownership rate, are computed consistently with the definitions of the corresponding shares in the SCF pooling all waves from 1989 to 2022. I calibrate the yearly per-period participation cost κ_1^p to match the average participation rate and the multiplier χ to match the homeownership

Figure 1: Housing Expenditure and Portfolio Shares



rate.³⁰ As reported in table 6, the model’s participation rate is in line with the 0.52 observed in the data, while the homeownership rate is somewhat below its data counterpart. This is mainly due to the existence of a minimum house size and the absence of mortgages: in order to become homeowners households must immediately hold a substantial net housing equity position at the time of purchase, which requires them to accumulate that equity first, keeping lower-wealth households renting longer than in the data and pushing down the homeownership rate.

5. Results

Having shown that the model matches households’ portfolio shares along the wealth distribution, I now move to the paper’s main objective: to quantify the role of heterogeneous exposure for the dynamics of wealth inequality. I focus on the top 10% wealth share as the main inequality statistic, and proceed in two steps: First, I show that the model replicates the observed U-shape in inequality. Second, I look at the model IRF to return shocks and propose a novel decomposition of the KF equation (10) to map the model response into its underlying economic channels.


5.1. The Dynamics of Wealth Inequality

THE REALIZED SEQUENCE OF RETURNS. Since the model generates large and persistent changes in inequality along the ergodic distribution, can it also replicate the observed increase in inequality given the realized sequence of asset returns? Figure 2 feeds the realized sequence of U.S. asset returns into the model and plots the resulting path of the top 10% wealth share against the same series from the World Inequality Database (WID) from 1963 to 2020. Given the two series slightly differ in levels, in panel 2a I normalize both to equal one in 1963.³¹

³⁰The calibrated κ_1^p value of 0.004 is expressed in yearly terms: with $\lambda_1^p = 365$, each arrival charges κ_1^p/λ_1^p , so that on average participants pay the yearly fee in daily installments.

³¹To feed in the realized sequence of returns I first extract a sequence of shocks for $\mathbf{W}_t$ using the calibrated process in equation (4) and then simulate the economy forward at a yearly frequency using the resulting realizations of $d\mathbf{W}_t$, starting from the risky steady state. Consistent with the model assumption of constant investment opportunities, all movements in returns are therefore treated as temporary, an assumption to which I return below.

Figure 2: Top 10% Wealth Share Dynamics – Model vs. Data




The model closely tracks the overall U-shape pattern in wealth inequality and, with the exception of two specific windows between 1979–85 and 2000–08 — the areas shaded in red in both panels — the model delivers a remarkable fit to the data: over the entire 50+ years sample, it delivers 73.7% of the total decline and 50.8% of the subsequent rise, where each share is measured from 1963 to the series’ trough and then to the series’ peak.³² Panel 2b instead renormalizes the two series at the end of each “mismatch window”, in 1985 and 2008, which further highlights how such mismatch is mostly concentrated in the two shaded areas: outside of these two periods, the model delivers 102.4% of the decline and 62.8% of the rise in the top 10% share. As we will see momentarily these two mismatch windows share similar characteristics, which help us understand along which dimensions the model does not fit the data as well; both windows include periods in which the U.S. economy experienced major recessions and both involved large movements in returns — especially in the housing market — so it is only natural to turn to the return process to understand what is going on.

COUNTERFACTUALS. To attribute the movements in inequality to individual asset returns, in figure 3 I construct a set of counterfactuals in which I shut down variation in each asset’s return series to assess how much that asset’s return movements contributed to the evolution of inequality. I do so by replacing the realized path with each asset’s calibrated mean return while keeping the other two assets’ return paths unchanged. Panel 3a plots, for each asset, the cumulative contribution of its realized return variation to the change in the top 10% wealth share since 1963, while panel 3b breaks each contribution into the total rise and the total decline within each period.

I find that if on the one hand movements in returns to equity and safe assets are the main drivers of the broad U-shape pattern from 1963 to 2000, on the other, housing played a significantly

³²The comparison with the WID data series is not exactly apples-to-apples. In fact, unlike the WID series and consistent with its SCF calibration targets, the model excludes private business wealth; since the latter grew substantially over precisely this 1980–2012 window (Smith et al. 2023), the mismatch is not entirely innocuous. Different data sources present different trade-offs and I take the WID as the most commonly used data in the literature, especially for long-run comparisons. The overall match between data and model turned out to be extremely robust across different iterations of the model, different data sources and targets, as well as various calibrations — e.g. both with and without discount-rate heterogeneity, with or without private equity, etc. — albeit of course to slightly different extents.

Figure 3: Top 10% Wealth Share Dynamics – Counterfactual Attribution



Note: Each curve in panel (a) gives one asset’s contribution: the baseline change minus the change in the counterfactual that replaces that asset’s realized returns with its calibrated mean, cumulated since 1963. Panel (b) sums each contribution’s annual increases and decreases separately within each period, with the yellow marker identifying the net change; the “Total” block reads on the left axis, while the period blocks on the right axis. Note that, because the model is not linear, the three assets’ net changes need not sum exactly to the baseline’s net change.

larger role from then onwards. In the two 1963–79 and 1985–00 periods movements in inequality are mainly driven by equity and safe return changes: in 1963–79 they account for about 58% and 35% of the decline, respectively; in 1985–00 about 60% and 44%, while housing returns’ contribution is smaller and negative on net. As section 5.2 below will show, the model response to high equity returns peaks late and decays very slowly (figure 4b), so a decade of abnormal equity returns in either direction translate in a persistent shift in the level of inequality. Additionally, because that persistence is primarily driven by the portfolio-choice response (figure 6) and the response is state-dependent (figure 5a) a run of consistently high or consistently low equity return years compounds into a larger total response. Together, these three forces allow the model to generate the sort of slow and cumulative changes in inequality that we observe from 1963 to 1979 and then from 1985 to 2000; even a sequence of small temporary return changes can significantly drive inequality.

The reason why housing plays a larger role for the remaining two 1979–85 and 2000–08 periods is correspondingly simple: in the model housing booms (respectively, busts) imply a reduction (rise) in top shares because housing is disproportionately the middle class’s asset; this is in line with findings from Kuhn et al. (2020), who estimate a house-price elasticity of the top 10% wealth share of about -0.17 . In 1980 housing (real) returns turn negative and remain so every year through 1984; to understand why, recall that the return series for housing is the leverage-adjusted real return on home equity that I construct from the Macrohistory Database’s house price series (online appendix C). In those years house-price growth fell short of inflation while mortgage rates remained high during the Volcker disinflation period, which leads to negative real housing returns and a rise in the top 10% share.³³ Between 2000 and 2008 the same mechanism is at work but in the opposite direction: the model falls by 4.5 log points while the data rise by 3.6. Specifically, there are two parts to this

³³Additionally, on the data side, changes in the role of business wealth account for roughly half of the 1979–85 gap. Specifically, business wealth fell from 21.4% of top 10% wealth in 1979 to 10.1% in 1986. When recomputing the same decomposition against an ex-business series, the 1979–1985 gap roughly halves, from 7.9 to 3.2–3.5 log points.


explanation. First, between 2000 and 2006 the model’s decline in inequality reflects two forces which both push the top share down: the housing boom, which increases wealth for the middle class more than for the top, and the dot-com bubble bust, which hits the top 10% directly and brings down inequality in 2001 and 2002. Second, between 2007 and 2009 the effects of strong negative returns to both equity and housing largely offset each other: the equity crash alone would push the top share down by 1.8 p.p. and the housing bust alone would push it up by 1.0 p.p.; together with the movement in safe returns, inequality in the model moves by only -0.49 p.p. over the two years.

In the data, instead, much of the measured rise over this window was driven by middle-class households borrowing heavily against their housing gains (Mian and Sufi 2011; Saez and Zucman 2016) and by returns to private equity, which remained high through the dot-com collapse and earned a large premium over public equity between 1999 and 2007 (Kartashova 2014). Both are factors that the model simply cannot match since it lacks private equity and levered borrowing on housing only enters through the aggregate loan-to-value ratio embedded in a common housing return.

While this exercise is similar in spirit to the one performed by Hubmer et al. (2021), I reach different conclusions on the role of asset returns after 1985 who find that return premia explain the initial dip and dampened the increase in concentration. Two differences help explain the gap: first, my model features an optimal portfolio choice decision that depends on the household’s entire state vector, whereas their portfolio weights are fixed functions of wealth; as section 5.2 below shows, this is a key component of the model response. Second, I assume that movements in returns are unanticipated and temporary, whereas they feed the model an anticipated, smooth path of returns under perfect foresight. Instead, my results are much more similar and consistent with those found using SCF+ data by Kuhn et al. (2020), who find that the evolution of post-war U.S. inequality was mainly driven by a race between equity and housing markets. Relative to them, I find that movements in safe asset returns also played a major role, bringing this closer to the findings in Greenwald et al. (2025), though they mostly focus on the role of duration in an exogenous portfolio-choice model.

My results are also broadly consistent with the findings of Catherine et al. (2026) who, however, find that most of the swing in top wealth shares has been driven by changes in real interest rates. In their model, both stock and house prices rise when rates fall, because both are claims to future cash flows; so a large part of the gains that my counterfactuals assign to each asset market is attributed, in their decomposition, to rate movements. They interpret return movements as revaluations caused by persistent, mean-reverting declines in discount rates which lower expected returns going forward, while I treat return movements as windfalls which leave them unchanged. Because the return process is exogenous in both models, our disagreement amounts to a choice between different descriptions of the same returns. In the context of such partial-equilibrium models, the natural criterion for choosing between the two is households’ behavior, which is what both models actually target. Existing empirical evidence finds that investors’ stated expectations of future returns move little when prices rise (De La O and Myers 2021), and households hold on to capital gains rather than consume them (Fagereng et al. 2019). The repricing interpretation would require investors to revise their expected returns downward whenever asset prices rise, which survey measures of return expectations do not show. Below, I measure how much the choice between the two interpretations matters.

Figure 4: Response of the Top 10% Wealth Share to Return Shocks



5.2. Understanding the Dynamics of Inequality

To better understand the model's response to shocks, I look at both generalized and standard impulse response functions, which I refer to as gIRF and sIRF, respectively.

Let g_0 denote the cross-sectional distribution at the time the shock hits, $S_{0.1}(g)$ the top 10% wealth share implied by a distribution g , and ϵ_t the vector of return innovations at t . Following Koop et al. (1996), I define the gIRF of inequality at horizon t to an innovation ϵ hitting at $t = 0$ as


$$\text{gIRF}(t; \epsilon, g_0) = \mathbb{E} [S_{0.1}(g_t) \mid \epsilon_0 = \epsilon, g_0] - \mathbb{E} [S_{0.1}(g_t) \mid \epsilon_0 = \mathbf{0}, g_0], \quad (13)$$

where both expectations integrate over all subsequent innovations $\{\epsilon_s\}_{s>0}$. The shock ϵ is scaled so that the shocked asset's return moves by one standard deviation on impact, and orthogonalized so that the other two asset returns are unaffected. I evaluate both expectations by simulating the model with the same shocks: for each initial state, I draw 100 future innovation paths; along each path, I then solve the KF equation (10) forward twice from the same g_0 , once with the impact innovation set to ϵ and once with it set to $\mathbf{0}$ (and with identical innovations from then on), so that the difference isolates the effect of the shock and that the sampling noise from future realizations cancels. For the initial states, I pick fifty initial distributions from a 20,000-year simulation of the model; these are taken to be at least a century apart and chosen to span the realized range of top 10% wealth shares. The mean gIRF then averages the gIRF across the fifty initial states.

The sIRF and gIRF have complementary roles: On the one hand the sIRF traces the response of the *average* cross-sectional distribution: it starts at the risky steady state, is hit by a one-time shock, and then evolves under the mean forward equation with all later shocks averaged out. It therefore isolates the effect of a shock on the average distribution and, because the KFE is linear in g_t , makes equation (14) below an exact decomposition of the density's law of motion. On the other, because the model never converges to a constant distribution, the gIRF allows us to show the dependence of the model response on the state of the economy and the nonlinearity of the inequality statistics.

Figure 4 plots the sIRF to all three shocks in the left panel and both gIRF and sIRF to an equity shock in the right panel. The first things we can observe from panel 4a are that both equity and

Figure 5: Equity gIRF Across Initial States, Shock Size, and Shock Sign



Note: The color bars encode the initial top 10% share of each starting state in panel (a) and the absolute shock size, in standard deviations, in panel (b).

bond shocks increase inequality but equity has by far the largest and most persistent effect, while housing reduces inequality and its effect declines much faster. Peak responses to equity, housing, and bond shocks are 0.8% (0.5 p.p.), -0.6% (-0.4 p.p.), and 0.4% (0.2 p.p.), with half-lives of 48.5, 3.2, and 35.8 years, respectively.

The gIRF in panel 4b plots the mean response to the same equity shock as well as interquartile and 90-10 bands across initial states. Relative to the sIRF, the mean gIRF response peaks roughly 10 years later and is about one quarter smaller at the peak, which is consistent with the average top share being slightly below the top share in the average distribution (due to Jensen's inequality); see figure 5a below.³⁴ The gIRF also shows that the spread around the mean response is quite large; as we see next, this spread is almost entirely determined by the level of initial inequality.

STATE, SIZE, AND SIGN DEPENDENCE. Figure 5 plots the model response as a function of the initial state of the economy when the shock hits, the size of the shock, and its sign.³⁵

The state-dependence plot in panel 5a displays the average gIRF as well as the per-state response across all fifty initial states separately and sorted by color based on the initial level of inequality.³⁶ The model response to the same equity shock varies widely depending on initial inequality, and top shares increase much more when inequality is already high; this is because initial inequality is high when the top holds more equity, which naturally makes top wealth shares further increase in response to shocks. That is, positive equity shocks make the economy more unequal, which makes future positive equity shocks have larger effects. This is a key aspect to understand the evolution of inequality in the U.S.: it only takes a few years of above-average equity returns for inequality to increase significantly and, combined with the persistence of the model's response, to remain high.

Panel 5b then turns to the size of the shock: holding the same fifty starting states and the same future draws fixed, it plots the mean response for shock sizes corresponding to $\frac{1}{2}$, 1, 2, and 2.5 standard

³⁴The top 10% wealth share in the stochastic steady state, $E[S_{0.1}(g_t)]$, is 2.9% smaller than the one in the risky steady state, $S_{0.1}(E[g_t])$, where both expectations are taken with respect to the ergodic distribution.

³⁵Online Appendix B.1 reports the gIRF for housing and bond shocks.

³⁶Each state's shaded band spans one standard error of its mean response across the 100 simulated future paths.

deviations in both directions *normalized* by the size of the shock itself (the responses to negative shocks are flipped in sign to make the comparison more immediate). We see that the response varies with the size (and sign) of the shock: larger positive shocks produce smaller responses, while larger negative shocks imply bigger changes. At peak the per-standard-deviation response to a 2.5σ boom falls 13.4% short of the 1σ response, while the response to a bust of the same size runs 9.6% higher; overall, a 2.5σ bust moves the top 10% share about 1.5 times as a corresponding boom. However, note that while the size dependence is apparently large, not all of it is due to behavioral differences in response to booms and busts. Rather, some is a strictly mechanical consequence of the top 10% share being given by the ratio between wealth at the top and overall wealth.³⁷ Also, while not shown here, I find no interaction between state and size dependence: the asymmetry between larger and smaller shocks is almost completely independent of the initial level of inequality.

DECOMPOSING THE MODEL RESPONSE. The KF equation (10) uniquely determines the evolution of the cross-sectional distribution and can therefore be decomposed into its different parts to understand what is driving the model response to shocks. Specifically, since the KFE operator $\mathcal{A}^*$ is linear, it can be expanded into its separate components, which we can then collect into economically interpretable channels. The decomposition of the KF equation (10) below follows directly from the definition of $\mathcal{A}^*$ by rearranging terms:

$$\begin{aligned}
dg_t = & \left\{ \underbrace{-\partial_a[\mu_{\text{sav}} g_t]}_{\text{consumption-saving}} + \overbrace{\sum_{j \in \{1,2\}} \partial_{z_j} [\beta_j z_j g_t] + \mathcal{Z}_J^* g_t}^{\text{saving}} \right. \\
& + \overbrace{\frac{1}{2} \sum_{i,j \in \{a,h\}} \partial_{ij} [\Sigma_{ij} g_t]}^{\text{portfolio choice}} - \underbrace{\partial_a[\mu_{\text{eq}} g_t] + \mathcal{P}^* g_t}_{\text{equity}} - \underbrace{\partial_h[\mu_h g_t] + \mathcal{H}^* g_t}_{\text{housing}} \\
& \left. + \underbrace{\zeta(\Psi_t - g_t)}_{\text{life-cycle}} \right\} dt + \underbrace{\mathcal{B}^* g_t}_{\text{aggregate revaluation}} d\mathbf{W}_t
\end{aligned} \tag{14}$$

³⁷Consider the following simple example: two households starting with almost identical wealth, but one holding 100% of its wealth in equity and the other 0%, and consider the share of total wealth held by the equity-holder, which at time 0 begins at approximately one half. Let us also assume that a 1σ return shock increases the value of equity by 50%. Such a shock would raise the equity-holders' wealth share to $1.5/2.5 = 0.6$, a gain of 10 p.p.; a 2σ return shock would instead double the value of equity and raise the wealth share to $2/3$, which, when normalized by the shock size, implies a gain of just 8.3 p.p.; conversely, a -1σ shock would halve the value of equity and lower the share to $0.5/1.5 = 1/3$, a loss of 16.7 p.p. That is, positive shocks inflate total wealth thereby diluting increases in top shares, while negative shocks shrink total wealth thereby amplifying the per-unit reductions in top shares (I have here tracked the equity holder's share rather than the top share, since a bust would also re-rank the two households, which introduces a distinct margin from the pure revaluation mechanism at work in this example).

where $\mu_h \equiv r_H h$ is the return on the housing stock and the only other change has been to split the wealth drift of equation (6) as $\mu_a = \mu_{\text{sav}} + \mu_{\text{eq}}$ by defining $\mu_{\text{eq}} \equiv (r_S - r_B) \theta p a$ as the excess return on the household's equity position and $\mu_{\text{sav}} \equiv e^z + (r_B + \zeta) a - x - \kappa^b(a)$ as the “saving” component that is standard in all incomplete-market models. Because the operator is linear in g_t , any subset of these channels can be “deactivated” (i.e. set to zero) and the remaining terms still define a well-posed forward equation (provided the resulting operator remains the adjoint of a valid generator).³⁸

The saving channel includes the consumption-saving drift μ_{sav} and the earnings process, therefore collecting all terms that would still move wealth around even if households held a single safe asset. The portfolio-choice channel, by contrast, is composed only of terms driven by households holding multiple assets: the return-risk correction coming from Σ_{ij} , and the two asset margins, equity and housing. Within each asset margin, the drift term contains the intensive margin of adjustment, while the jump operator the extensive one. Along the intensive margins, conditional on being a stock-market participant ($p = 1$), households earn μ_{eq} on their equity position; similarly, conditional on being a homeowner ($h > 0$), households earn μ_h on their house. Along the extensive margins, $\mathcal{P}^*$ gathers the terms that move households into and out of the equity market and $\mathcal{H}^*$ into and out of homeownership. The life-cycle term replaces dying households with newborns drawn from Ψ_t . Finally, the last term — the direct effect from aggregate revaluations — shifts households' wealth in proportion to their risk exposure when an aggregate shock realizes.³⁹

Figure 6 plots the contribution of the various channels for the sIRF to a positive equity shock and panel 6a immediately shows that the total response is much larger and much more persistent than what the direct revaluation effect alone would imply and that most of the total response is actually driven by the portfolio choice component.⁴⁰


The direct effect, the green solid line, is straightforward to interpret: since equity is held primarily by households at the top of the distribution, a positive equity return increases the value of their equity portfolios and translates into additional inequality; at its peak the direct effect contributes 47.0% of the peak total response and its total contribution (the area under the curve) amounts to 28.5% of the total. The saving channel, the black dotted line, contributes to generating some additional inequality early on (at its peak, in year 5, it accounts for 29.2% of the total response) but, by year 28, its net effect on inequality turns negative and its total contribution ends up at -8.2% of the total. To understand the saving channel, observe that households that gain from the shock will spend only a small share of the windfall in any given year, so at first the positive equity shock translates into

³⁸Since inequality statistics are nonlinear in g_t , in order for the various channels to add up to the total response I solve equation (14) by adding each channel sequentially and compute its contribution as the corresponding change in inequality.

³⁹Because expected returns are constant, the policy functions are the same before and after each shock and the operators in equation (14) are therefore built from these constant policy functions. Nonetheless, return shocks do indeed affect household behavior — movements in returns change households' states and correspondingly affect their optimal policies — and the decomposition in equation (14) specifically measures the behavioral channels: the direct effect captures the counterfactual in which households *do not* react, and every other channel is the cumulative effect of households re-optimizing their behavior along their optimal policy as their state vector changes.

⁴⁰In what follows we pair the direct effect with a life-cycle component simply because, by definition, the direct effect alone has a *permanent* effect on wealth and the shock mean-reverts only because of life-cycle effects. Notice also that the effect of the earnings channel is zero since return shocks do not affect earnings inequality.

Figure 6: Equity sIRF Decomposition




additional inequality. Over time however, as the revaluation has carried these households towards the top of the distribution — where savings are generally negative — their spending gradually erodes the gains and the contribution eventually turns negative. The portfolio response, the purple dash-dotted line, is by far the largest of all channels: at its peak, in year 32, it accounts for 76.1% of the total response, and its total contribution is 79.8% of the total response.

To better understand where the portfolio choice response is coming from, panels 6b and 6c further decompose it into its sub-components. We can immediately see that the equity margin accounts for the bulk of the total portfolio response (around 85%), and that this is driven mainly by the intensive margin, which alone accounts for 57.6% of the *total* response in the top 10% share: as households grow wealthier they tilt their spending towards luxury goods, which makes them less risk averse; a larger and larger share of the population therefore holds more of their wealth in equity, which the equity risk premium compounds into additional inequality. Similarly, the risk component — the yellow dashed line — increases as households' equity positions grow and accounts for 23.3% of the portfolio effect (or, equivalently, 18.6% of the total inequality response). On the other hand, housing actually contributes to dampening the total response, mostly through its *extensive* margin (−12.7% of the total): households who already hold equity but are still renting spend the extra funds to become homeowners, while existing owners sell (the wealthier ones trading up to a larger house, while the constrained ones trading down or leaving homeownership altogether) and this churn in and out of homeownership gradually redistributes the gains back across the distribution.

As the above decomposition shows, the response to shocks is strongly driven by the portfolio choice response: modeling households' optimal portfolio choice decision is of first order importance to capture how the economy responds to return shocks. To more directly answer how much this mattered for the evolution of inequality in the U.S., we can apply the same channel decomposition to the historical simulation of section 5.1. Figure 7 presents the results: panel 7a plots the top 10% share's deviation from its 1963 level together with the contribution of each channel, while panel 7b breaks each channel's contribution into its total rise and its total decline within each period.

In the two long periods driven by equity and safe returns, namely 1963–79 and 1985–00, the behavioral channels amplify the direct effect: between 1963 and 1979 the top share falls by 5.1 percentage points, with direct revaluation accounting for 3.5 of those, the portfolio response adding

Figure 7: Top 10% Wealth Share Dynamics – Channel Decomposition



Note: Panel (a) plots the top 10% wealth share’s deviation from its 1963 level in the historical simulation of section 5.1, together with the contribution of each channel of equation (14). Panel (b) sums each channel’s annual increases and decreases separately within each period, with the yellow marker identifying the net change; the “Total” block reads on the left axis, while the period blocks read on the right axis.


0.9, and the saving response 0.7 — so about a third of the total decline in inequality comes from behavioral responses. Between 1985 and 2000 inequality rises by 3.0 percentage points, with households’ behavioral response doing most of the work: direct revaluation contributes 1.3 percentage points, while portfolio choice and saving contribute 0.8 and 0.9, respectively — more than half of the model’s increase in inequality. In the two 1979–85 and 2000–08 windows, instead, the response is almost entirely given by the direct revaluation channel, while the behavioral margins together push mildly in the opposite direction (the offset in 2000–08 comes from savings alone).

THE ROLE OF THE RETURN PROCESS. As mentioned above, another difference relative to Greenwald et al. (2025); Hubmer et al. (2021) and Catherine et al. (2026) is the return process itself. Specifically, how would the model response change if households believed that a rise in current returns is followed by a fall in future expected returns? And how would it change if shocks were permanent rather than temporary? To answer these questions, in figure 8 I run two sets of experiments: in the first I compare different interpretations of the same equity shock: a one-time windfall gain vs. a repricing of future dividend flows; in the second I compare two simulations that generate exactly the same sequence of realized equity returns under different assumptions about the nature of realized returns.

For the first experiment, I take a one-time equity revaluation of one standard deviation and trace its response under the following assumptions on households’ beliefs about future expected returns: in the first case, households believe the equity shock is a pure windfall gain — future expected returns are unchanged — so the model response is the baseline sIRF to an equity shock; in the second case, households believe that the increase in equity returns is due to a repricing of future dividends (e.g. due to a fall in discount rates) and the same revaluation is paired with a permanent decline in the expected equity return of 0.92 p.p.;⁴¹ in the third case, households believe that this repricing

⁴¹For the fall in future expected returns I use the revaluation divided by the loading of stock returns on interest-rate

Figure 8: Equity IRF Under Different Return Beliefs



is persistent but not fully permanent — future expected returns fall on impact but slowly return back to pre-shock levels — which allows for a direct comparison with Catherine et al. (2026).⁴² Panel 8a plots the model responses under the three assumptions on beliefs. The two repricing paths show a response in inequality that is almost identical for the first 25 years and that, at its peak, is approximately 9% larger than the windfall baseline. Then, after about 50 years, inequality falls below the baseline response; and under the assumption of permanently lower expected returns, inequality in the final steady state is about 0.3% lower than where it started from.

For the second experiment, the first simulation feeds the model with a continuous sequence of temporary innovations that holds the realized equity return one percentage point above its long-run mean in every period, leaving expected returns unchanged; the second simulation is a single, one-time, permanent increase in the expected equity return of one percentage point.⁴³

Panel 8b shows that the responses to these two identical realized return sequences are extremely different. The sequence of temporary shocks raises the top 10% share by 5.3% in the long run, while the permanent revision to expected returns leaves inequality almost untouched: an extremely small decline (-0.1%) over the first two decades, followed by a long-run increase of only 0.46%. This second experiment helps explain the difference relative to Hubmer et al. (2021), who find that asset-price fluctuations mainly explain the initial dip in inequality.

To understand where these differences are coming from we can run again the same decomposition of equation (14) on the temporary sequence of shocks, which is shown in panel 8c. The portfolio choice and direct revaluation components play by far the largest role and account for about three quarters and 45% of the total long-run change in inequality; savings, on the other hand, play a much smaller role and eventually contribute to a small reduction in the total change of about a fifth. The

movements estimated by Catherine et al. (2026), about 17. The revaluation arrives as a one-time shock under the pre-shock decision rules. After the shock, the household problem is then re-solved with the lower expected return and the KFE solved forward from the post-shock distribution.

⁴²The main difference is that I here solve for the transition to a one-time shock under perfect foresight. This is the cleanest possible comparison in the context of my model since a full repricing implementation of the historical simulation would essentially require writing a completely different model in which expected returns comove with realized returns and households have proper expectations about such comovement. To keep the comparison as close as possible, I run it by using their estimate for the annual persistence of expected returns which implies a half-life of roughly fifteen years.

⁴³To compute the response to a one-time permanent change I re-solve the household problem under the higher expected return and solve the KF equation forward starting from the pre-shock wealth distribution.

sustained excess return to equity mechanically revalues households' equity position upwards, which increases inequality; as wealthy households get wealthier, they keep investing more and more into the equity market, which further compounds the direct effect; on the other hand, as wealth grows, households move up their declining saving schedule, which eats part of the equity gains.

Conversely, permanent shocks imply very different behavior: on the one hand, equity holders increase their expenditure and reduce their savings because they anticipate higher equity returns in the future, which does not happen in the face of temporary shocks since in that case the expenditure schedule does not move; on the other hand, the bottom of the distribution (i.e., non-participants) now has an additional incentive to save and invest in equity, which further compresses the distribution due to a reduction in heterogeneity in both wealth and portfolios. Hence, even though the larger equity premium implies a larger effect on steady state inequality through return heterogeneity, the total effect is much smaller than if those changes in returns were temporary.

This also explains why the response is stronger when realized and expected returns are negatively correlated: in the face of lower future expected returns, wealthy equity holders increase their savings without investing the wealth gain into equity — a pattern that is different from the “saving by holding” behavior documented by Fagereng et al. (2019) — while, at the other end of the distribution, poorer non-participants have fewer incentive to invest into equity now that future returns are lower. This contributes to the even larger effect on top wealth shares that can be observed in panel 8a.

6. Conclusions

In this paper I propose a theory for the dynamics of wealth inequality based on heterogeneous exposure to aggregate risk. I show that large movements in wealth inequality are the natural outcome of a model in which households are differently exposed to common shocks in asset returns. One key conclusion is that much of the observed U-shape in top wealth shares can be explained by realized returns happening to be first unfavorable and then favorable to the portfolios of the wealthy.

I use the model to show that temporary shocks to returns have large and persistent effects on top wealth shares. When fed the realized sequence of U.S. asset returns, the model can explain nearly three quarters of the long post-war decline in top wealth shares and about half of the subsequent rise. I find that most of the response to an equity shock comes from households' endogenous portfolio choice. Finally, the nature of the return process — whether return changes are temporary or permanent, and whether households believe realized returns comove with expected future returns — also matters.

One of this paper's main implications is that, to better understand the evolution of wealth inequality, we should turn our attention to *both* households' portfolios and asset prices. When it comes to prices, in this paper I take a very reduced-form approach and simply assume that asset returns follow an exogenous process. However, asset prices are obviously equilibrium objects and they are themselves likely to be influenced by changes in the wealth distribution too. To fully characterize the interaction between asset prices and wealth inequality would mean to be able to provide a model that can match the cross-sectional distribution of both wealth and portfolio shares, as well as the equity premium and its variation; a challenging task that future work should definitely tackle.

Appendix A. HJB

The generator $\mathcal{A}$ in the HJB equation (9) collects all the information about how the household state (a, h, z_1, z_2, p, m) evolves: its drift, its diffusion, and its idiosyncratic jumps.⁴⁴

$$\mathcal{A} = \mathcal{L}_c + \mathcal{Z} + \mathbf{1}_{\{p=0\}} \mathcal{P}^0 + \mathbf{1}_{\{p=1\}} \mathcal{P}^1 + \mathbf{1}_{\{h=0\}} \mathcal{H}^0 + \mathbf{1}_{\{h>0\}} \mathcal{H}^+.$$

The continuous operator $\mathcal{L}_c$ carries the drift and diffusion of wealth and housing, common to both participation states: participation enters only through the θp factor in the wealth drift and in the portfolio loading σ_a .

$$\begin{aligned} \mathcal{L}_c v = & \left[e^z + (r_B + \zeta)a + (r_S - r_B)\theta p a - x - \kappa^b(a) \right] \frac{\partial}{\partial a} v + \frac{1}{2} a^2 (\sigma_a \sigma'_a) \frac{\partial^2}{\partial a^2} v + \\ & + r_H h \frac{\partial}{\partial h} v + \frac{1}{2} h^2 (\sigma_3 \sigma'_3) \frac{\partial^2}{\partial h^2} v + a h (\sigma_a \sigma'_3) \frac{\partial^2}{\partial a \partial h} v, \end{aligned}$$

where $\sigma_a = (1 - \theta p)\sigma_1 + \theta p\sigma_2$ as in equation (6), so that $\sigma_a = \sigma_1$ at $p = 0$, $z = z_1 + z_2$, $x = c + r_n n \mathbf{1}_{\{h=0\}}$, and $\kappa^b(a) = \kappa^b \max\{(\theta p - 1)a, 0\}$, as in equations (6) and (7).

The earnings operator $\mathcal{Z}$ includes each component's mean reversion and its Poisson redraw. Writing $\mathbf{x}_{j'}$ for the state $\mathbf{x}$ with its j -th earnings component reset to z'_j , so $\mathbf{x}_{1'} = (a, h, z'_1, z_2, p, m)$ and $\mathbf{x}_{2'} = (a, h, z_1, z'_2, p, m)$, we have:

$$\mathcal{Z}v = \sum_{j \in \{1,2\}} \left[-\beta_j z_j \frac{\partial}{\partial z_j} v + \underbrace{\lambda_j \int [v(\mathbf{x}_{j'}) - v(\mathbf{x})] \phi_j(z'_j) dz'_j}_{\mathcal{Z}_j v} \right],$$

where ϕ_j is the $\mathcal{N}(0, \sigma_j^2)$ density from which the new component is drawn, as in equation (5).

When a Poisson participation opportunity arrives, the household relocates to the policy-chosen destination encoded in the kernel $\mathcal{K}_l^p$, so that each operator is simply $\mathcal{P}^l v = \lambda_l^p (\mathcal{K}_l^p - \mathcal{I})v$ with $l \in \{0, 1\}$:

$$\mathcal{P}^0 v = \lambda_0^p \left[\max \{v(a - \kappa_0^p, h, z_1, z_2, 1, m), v(a, h, z_1, z_2, 0, m)\} - v(a, h, z_1, z_2, 0, m) \right]$$

$$\mathcal{P}^1 v = \lambda_1^p \left[\max \{v(a - \kappa_1^p, h, z_1, z_2, 1, m), v(a, h, z_1, z_2, 0, m)\} - v(a, h, z_1, z_2, 1, m) \right]$$

Housing adjustment works in exactly the same way: a renter draws a chance to buy ($\mathcal{K}_B^h$, also choosing the new size h'), an owner a chance to sell ($\mathcal{K}_S^h$), each at transaction cost $\kappa^h(h) = \kappa_0^h + \kappa_1^h h$

⁴⁴Because the discount type m enters the household's problem only through the time-preference rate ρ_m , which appears in the HJB equation (9) separately from the generator, $\mathcal{A}$ itself does not depend on m .

so that $\mathcal{H}^0 v = \lambda^h (\mathcal{K}_B^h - \mathcal{I}) v$ and $\mathcal{H}^+ v = \lambda^h (\mathcal{K}_S^h - \mathcal{I}) v$, where:

$$\mathcal{H}^0 v = \lambda^h \left[\max \left\{ \max_{h'} v(a - h' - \kappa^h(h'), h', z_1, z_2, p, m), v(a, 0, z_1, z_2, p, m) \right\} - v(a, 0, z_1, z_2, p, m) \right]$$

$$\mathcal{H}^+ v = \lambda^h \left[\max \left\{ v(a + h - \kappa^h(h), 0, z_1, z_2, p, m), v(a, h, z_1, z_2, p, m) \right\} - v(a, h, z_1, z_2, p, m) \right]$$

For deriving the KFE in online appendix D, it will also be useful to collect the continuous terms and the Poisson jumps together and write $\mathcal{A} = \mathcal{A}_c + \mathcal{A}_J$, where $\mathcal{A}_c = \mathcal{L}_c - \sum_j \beta_j z_j \partial_{z_j}$ and $\mathcal{A}_J = \sum_j \mathcal{Z}_j + \mathbf{1}_{\{p=0\}} \mathcal{P}^0 + \mathbf{1}_{\{p=1\}} \mathcal{P}^1 + \mathbf{1}_{\{h=0\}} \mathcal{H}^0 + \mathbf{1}_{\{h>0\}} \mathcal{H}^+$.

References

- ACHDOU, Y., J. HAN, J.-M. LASRY, P.-L. LIONS, and B. MOLL (2022). “Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach”. *The Review of Economic Studies*.
- AGUIAR, M. and M. BILS (2015). “Has Consumption Inequality Mirrored Income Inequality?” *American Economic Review*.
- ARNOUD, A., F. GUVENEN, and T. KLEINEBERG (2022). “Benchmarking Global Optimizers”. *Working Paper*.
- ATKESON, A. G. and M. IRIE (2022). “Rapid Dynamics of Top Wealth Shares and Self-Made Fortunes: What Is the Role of Family Firms?” *American Economic Review: Insights*.
- ATKESON, A. G. and M. OGAKI (1996). “Wealth-Varying Intertemporal Elasticities of Substitution: Evidence from Panel and Aggregate Data”. *Journal of Monetary Economics*.
- AZZALINI, G., M. KONDZIELLA, and Z. RÁCZ (2023). “Preference Heterogeneity and Portfolio Choices over the Wealth Distribution”. *Working Paper*.
- BACH, L., L. E. CALVET, and P. SODINI (2020). “Rich Pickings? Risk, Return, and Skill in Household Wealth”. *American Economic Review*.
- BELDA, P., J. HEINEKEN, and A. IFRIM (2025). “How Expectations Amplify Wealth Inequality”. *Working Paper*.
- BENHABIB, J., A. BISIN, and M. LUO (2019). “Wealth Distribution and Social Mobility in the US: A Quantitative Approach”. *American Economic Review*.
- BENHABIB, J., A. BISIN, and S. ZHU (2011). “The Distribution of Wealth and Fiscal Policy in Economies With Finitely Lived Agents”. *Econometrica*.
- BENZONI, L., P. COLLIN-DUFRESNE, and R. S. GOLDSTEIN (2007). “Portfolio Choice over the Life-Cycle When the Stock and Labor Markets Are Cointegrated”. *The Journal of Finance*.
- BLANCHARD, O. J. (1985). “Debt, Deficits, and Finite Horizons”. *Journal of Political Economy*.
- BLANCHET, T. (2022). “Uncovering the Dynamics of the Wealth Distribution”. *Working Paper*.
- BLUNDELL, R., M. BROWNING, and C. MEGHIR (1994). “Consumer Demand and the Life-Cycle Allocation of Household Expenditures”. *The Review of Economic Studies*.
- BRICKER, J., J. BRIGGS, S. FRIEDMAN, and K. SOMMER (2025). “The Business Cycle Dynamics of the Wealth Distribution”. *Journal of Political Economy Macroeconomics*.
- BROWNING, M. and T. F. CROSSLEY (2000). “Luxuries Are Easier to Postpone: A Proof”. *Journal of Political Economy*.

- CALVET, L. E., J. Y. CAMPBELL, F. GOMES, and P. SODINI (2021). “The Cross-Section of Household Preferences”. *Working Paper*.
- CARMONA, R. and F. DELARUE (2018). *Probabilistic Theory of Mean Field Games with Applications II: Mean Field Games with Common Noise and Master Equations*. Springer.
- CARMONA, R., F. DELARUE, and D. LACKER (2016). “Mean Field Games with Common Noise”. *The Annals of Probability*.
- CARROLL, C. D. (2000). “Why Do the Rich Save So Much?” In: *Does Atlas Shrug? The Economic Consequences of Taxing the Rich*.
- CARTIS, C., J. FIALA, B. MARTEAU, and L. ROBERTS (2019). “Improving the Flexibility and Robustness of Model-Based Derivative-Free Optimization Solvers”. *ACM Transactions on Mathematical Software*.
- CATHERINE, S. (2022). “Countercyclical Labor Income Risk and Portfolio Choices over the Life Cycle”. *The Review of Financial Studies*.
- CATHERINE, S., M. MILLER, J. D. PARON, and N. SARIN (2026). “Who Gains When Interest Rates Fall?” *Working Paper*.
- CHANG, Y., J. H. HONG, and M. KARABARBOUNIS (2018). “Labor Market Uncertainty and Portfolio Choice Puzzles”. *American Economic Journal: Macroeconomics*.
- CHETTY, R., L. SÁNDOR, and A. SZEIDL (2017). “The Effect of Housing on Portfolio Choice”. *The Journal of Finance*.
- CIOFFI, R. A., G. NUÑO, and S. HURTADO (2026). “Asset Pricing with Heterogeneous Agents under Limited Information”. *Working Paper*.
- COCCO, J. F. (2005). “Portfolio Choice in the Presence of Housing”. *The Review of Financial Studies*.
- DAMINATO, C. and L. PISTAFERRI (2024). “Returns Heterogeneity and Consumption Inequality Over the Life Cycle”. *Working Paper*.
- DE LA O, R. and S. MYERS (2021). “Subjective Cash Flow and Discount Rate Expectations”. *The Journal of Finance*.
- DUFFIE, D. and L. G. EPSTEIN (1992). “Stochastic Differential Utility”. *Econometrica*.
- FAGERENG, A., L. GUISO, D. MALACRINO, and L. PISTAFERRI (2020). “Heterogeneity and Persistence in Returns to Wealth”. *Econometrica*.
- FAGERENG, A., L. GUISO, and L. PISTAFERRI (2018). “Portfolio Choices, Firm Shocks, and Uninsurable Wage Risk”. *The Review of Economic Studies*.
- FAGERENG, A., M. B. HOLM, B. MOLL, and G. NATVIK (2019). “Saving Behavior Across the Wealth Distribution: The Importance of Capital Gains”. *Working Paper*.
- FAVILUKIS, J. (2013). “Inequality, Stock Market Participation, and the Equity Premium”. *Journal of Financial Economics*.
- FAVILUKIS, J., S. C. LUDVIGSON, and S. VAN NIEUWERBURGH (2016). “The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium”. *Journal of Political Economy*.
- FERNÁNDEZ-VILLAVARDE, J. and O. LEVINTAL (2026). “The Distributional Effects of Asset Returns”. *Working Paper*.
- FLAVIN, M. and S. NAKAGAWA (2008). “A Model of Housing in the Presence of Adjustment Costs: A Structural Interpretation of Habit Persistence”. *American Economic Review*.
- FLAVIN, M. and T. YAMASHITA (2002). “Owner-Occupied Housing and the Composition of the Household Portfolio”. *American Economic Review*.
- GABAIX, X., J.-M. LASRY, P.-L. LIONS, and B. MOLL (2016). “The Dynamics of Inequality”. *Econometrica*.
- GARRIGA, C. and A. HEDLUND (2020). “Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession”. *American Economic Review*.

- GOMES, F. and A. MICHAELIDES (2008). “Asset Pricing with Limited Risk Sharing and Heterogeneous Agents”. *The Review of Financial Studies*.
- GOMEZ, M. (2025). “Wealth Inequality and Asset Prices”. *The Review of Economic Studies*.
- GOMEZ, M. and É. GOUIN-BONENFANT (2024). “Wealth Inequality in a Low Rate Environment”. *Econometrica*.
- GOPALAKRISHNA, G., Z. GU, and J. PAYNE (2026). “Institutional Asset Pricing with Segmentation and Household Heterogeneity”. *Working Paper*.
- GREENWALD, D. L., M. LEOMBRONI, H. LUSTIG, and S. VAN NIEUWERBURGH (2025). “Wealth Inequality with Declining Interest Rates”. *Working Paper*.
- GUVENEN, F. (2009). “A Parsimonious Macroeconomic Model for Asset Pricing”. *Econometrica*.
- GUVENEN, F., F. KARAHAN, S. OZKAN, and J. SONG (2021). “What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics?” *Econometrica*.
- HANOCH, G. (1977). “Risk Aversion and Consumer Preferences”. *Econometrica*.
- HEATHCOTE, J., K. STORESLETTEN, and G. L. VIOLANTE (2017). “Optimal Tax Progressivity: An Analytical Framework”. *The Quarterly Journal of Economics*.
- HUBMER, J., P. KRUSELL, and A. A. SMITH (2021). “Sources of U.S. Wealth Inequality: Past, Present, and Future”. *NBER Macroeconomics Annual*.
- HYLAND, C. and A. KHEMKA (2025). “The Inequality Multiplier: Market Inelasticity and the Persistence of Wealth Inequality”. *Working Paper*.
- JORDÀ, Ò., K. KNOLL, D. KUVSHINOV, M. SCHULARICK, and A. M. TAYLOR (2019). “The Rate of Return on Everything, 1870–2015”. *The Quarterly Journal of Economics*.
- KACPERCZYK, M., J. NOSAL, and L. STEVENS (2019). “Investor Sophistication and Capital Income Inequality”. *Journal of Monetary Economics*.
- KAPLAN, G., B. MOLL, and G. L. VIOLANTE (2018). “Monetary Policy According to HANK”. *American Economic Review*.
- KARTASHOVA, K. (2014). “Private Equity Premium Puzzle Revisited”. *American Economic Review*.
- KASA, K. and X. LEI (2018). “Risk, Uncertainty, and the Dynamics of Inequality”. *Journal of Monetary Economics*.
- KAYMAK, B., D. LEUNG, and M. POSCHKE (2026). “Accounting for Wealth Concentration in the United States”. *American Economic Journal: Macroeconomics*.
- KOOP, G., M. H. PESARAN, and S. M. POTTER (1996). “Impulse Response Analysis in Nonlinear Multivariate Models”. *Journal of Econometrics*.
- KRUSELL, P. and A. SMITH (1998). “Income and Wealth Heterogeneity in the Macroeconomy”. *Journal of Political Economy*.
- KUHN, M., M. SCHULARICK, and U. I. STEINS (2020). “Income and Wealth Inequality in America, 1949–2016”. *Journal of Political Economy*.
- MARTÍNEZ-TOLEDANO, C. (2020). “House Price Cycles, Wealth Inequality and Portfolio Reshuffling”. *Working Paper*.
- MEEUWIS, M. (2020). “Wealth Fluctuations and Risk Preferences: Evidence from U.S. Investor Portfolios”. *Working Paper*.
- MERTON, R. C. (1971). “Optimum Consumption and Portfolio Rules in a Continuous-Time Model”. *Journal of Economic Theory*.
- MIAN, A., L. STRAUB, and A. SUFI (2020). “The Saving Glut of the Rich”. *Working Paper*.
- MIAN, A. and A. SUFI (2011). “House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis”. *American Economic Review*.
- MIJAKOVIC, A. (2025). “Income Inequality and the Rise of Risky Capital”. *Working Paper*.

- MOSKOWITZ, T. J. and A. VISSING-JØRGENSEN (2002). “The Returns to Entrepreneurial Investment: A Private Equity Premium Puzzle?” *American Economic Review*.
- PAKOŠ, M. (2011). “Estimating Intertemporal and Intratemporal Substitutions When Both Income and Substitution Effects Are Present: The Role of Durable Goods”. *Journal of Business & Economic Statistics*.
- PASHCHENKO, S. (2013). “Accounting for Non-Annuitization”. *Journal of Public Economics*.
- PIAZZESI, M. and M. SCHNEIDER (2016). “Housing and Macroeconomics”. In: *Handbook of Macroeconomics*.
- PIAZZESI, M., M. SCHNEIDER, and S. TUZEL (2007). “Housing, Consumption and Asset Pricing”. *Journal of Financial Economics*.
- QUADRINI, V. (1999). “The Importance of Entrepreneurship for Wealth Concentration and Mobility”. *Review of Income and Wealth*.
- RÁCZ, Z. (2025). “Housing and Portfolio Choice over the Wealth Distribution”. *Working Paper*.
- SAEZ, E. and G. ZUCMAN (2016). “Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data”. *The Quarterly Journal of Economics*.
- SMITH, M., O. ZIDAR, and E. ZWICK (2023). “Top Wealth in America: New Estimates under Heterogeneous Returns”. *The Quarterly Journal of Economics*.
- TRABANDT, M. and H. UHLIG (2011). “The Laffer Curve Revisited”. *Journal of Monetary Economics*.
- VESTMAN, R. (2019). “Limited Stock Market Participation Among Renters and Homeowners”. *The Review of Financial Studies*.
- VISSING-JØRGENSEN, A. (2002). “Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures”. *Working Paper*.
- WACHTER, J. A. and M. YOGO (2010). “Why Do Household Portfolio Shares Rise in Wealth?” *The Review of Financial Studies*.
- XAVIER, I. (2020). “Wealth Inequality in the US: The Role of Heterogeneous Returns”. *Working Paper*.
- YANG, Y., C. WANG, A. SCHAAB, and B. MOLL (2025). “Structural Reinforcement Learning for Heterogeneous Agent Macroeconomics”. *Working Paper*.

Supplemental Appendix for “Heterogeneous Risk Exposure and the Dynamics of Wealth Inequality”

Riccardo A. Cioffi

October 2026

B. Additional Results

B.1. Generalized Impulse Responses: Bonds and Housing

Figure B.1 plots the gIRF to bond and housing return shocks, drawn from the same set of ergodic starting points used for the equity gIRF in figure 5a. Unlike for equity, the mean bond and housing gIRFs essentially coincide with their standard impulse responses.

B.2. Non-homotheticity and ρ -Heterogeneity

The baseline model has two features other than return heterogeneity that affect inequality: discount-rate heterogeneity and non-homotheticity. Hence, I hereby ask how much the additional features contribute to the average level of inequality. Because neither channel enters the decomposition equation (14) as a separate term, I isolate each by re-solving the model with each feature switched off while holding every other parameter at its calibrated value, and I compare the resulting stationary distribution to that of the full model.

Table B.1 shows that removing non-homotheticity leaves wealth concentration essentially untouched with a reduction of less than 1 p.p. in both the top 10% share and Gini coefficient. Hence, while non-homotheticity is the main mechanism to generate wealth-dependence in risk aversion and intertemporal substitution — and therefore for the rising equity share — it does not, by itself, add to wealth inequality.

On the other hand, discount-rate heterogeneity accounts for most of wealth concentration in excess of earnings inequality. Setting a common discount rate at the middle type’s value lowers the top 10% wealth share from 58.4% to 36.5% and the Gini coefficient from 0.745 to 0.466, i.e. a wealth distribution no more unequal than the underlying earnings distribution. This is not at all surprising: the vector of discount rates ρ is calibrated to match group-level wealth-to-earnings ratios precisely because that spread decouples wealth accumulation at the top from the aggregate (see section 4.2).

At the same time, the single discount-rate economy also displays much higher total wealth accumulation, with a wealth-to-earnings ratio of 7.9 against the benchmark economy value of 4.60 (and well above the SCF aggregate of 5.4). This is the same tension highlighted in section 4.2: a model with a single discount rate cannot match both the aggregate wealth-to-earnings ratio and the concentration of wealth at the same time.¹

¹Even recalibrating the single discount-rate model to hit the wealth to income ratio exactly delivers a Gini coefficient of at most 0.5 to 0.55.

Figure B.1: Generalized Impulse Responses: Distribution of Individual Bond- and Housing-Shock Responses




Table B.1: Wealth Inequality – Policy Channels

CHANNEL REMOVED	Top 10%	Gini	Δ Top 10% (pp)	Δ Gini (pts)	W/Y
<i>Full model</i>	0.584	0.745	–	–	4.60
Discount-rate heterogeneity	0.365	0.466	+21.9	+27.9	7.90
Non-homotheticity ($A = 1$)	0.583	0.743	+0.07	+0.12	5.45

Note: Each channel row re-solves the model with the named feature removed and reports the resulting counterfactual top 10% share, Gini coefficient, and wealth-to-earnings ratio, alongside the channel’s contribution in points — the full-model value minus the counterfactual — so a positive entry means the feature raises the statistic. The non-homotheticity row imposes $A = 1$ by setting $\varepsilon_c = \varepsilon_h = 0.856$; ω and all other parameters stay at their calibrated values.

Since the level of inequality comes mostly from discount-rate heterogeneity, I next ask how much of it is driven by *between*- and *within*-type dispersion; table B.2 reports results for all three types. As argued in section 4.2 discount-rate type and wealth position are positively correlated but do not map one to one; as a matter of fact most of the households in the top decile belong to the intermediate discount-rate type. Nonetheless, as expected, patient households are the wealthiest: they hold 45.9% of wealth in the top wealth decile while only accounting for 37.5% of households in that decile. The middle 40% ($\rho = 0.161$) holds 61% of aggregate wealth and 53.9% of top-decile wealth, while the impatient half ($\rho = 2.295$) holds 4.5% of total wealth and almost none of the top decile’s share. Using the Heikkuri and Schief (2026) decomposition I find that between-type differences account for 45.2% of the aggregate Gini coefficient; within-type dispersion accounts for the remaining 54.8%. Hence, even though discount-rate heterogeneity accounts for most of wealth inequality, more than half of measured inequality is dispersion among households that share a single discount rate, driven by earnings histories and by the return differences that portfolio composition generates, which is particularly relevant at the bottom of the wealth distribution: the impatient type’s own Gini coefficient reaches 0.835 against roughly 0.55 for the other two types.

Table B.2: Wealth Inequality – Within- and Between-Type Inequality

TYPE	ρ	Aggregate		Top Wealth Decile		Within-type Gini
		pop. share	wealth share	pop. share	wealth share	
Patient	0.037	10%	34.5%	37.5%	45.9%	0.549
Middle	0.161	40%	61.0%	62.3%	53.9%	0.561
Impatient	2.295	50%	4.5%	0.3%	0.2%	0.835

Note: The top-decile share columns report each type’s share of *households* in the wealthiest decile and the share of the *wealth* they hold (each column sums to 100%).

C. Data Construction and Moment Computation

This appendix documents in detail the data construction and moment computation underlying the calibration in section 4. The primary data source is the *Survey of Consumer Finances* (SCF) and it is used to compute the wealth distribution and the portfolio composition of households in the United States. Other data sources used are the *MacroHistory DataBase* (MHDB) by Jordà et al. (2019), which I use for prices and returns data, and the *Consumer Expenditure Survey* (CEX) which I use to target consumer housing expenditures in the same fashion as in Aguiar and Bils (2015).

SURVEY OF CONSUMER FINANCES. To compute cross-sectional moments, I use all 12 currently available SCF waves covering 1989–2022. The sample is restricted to household heads aged 25–65, which excludes students and retirees whose wealth dynamics the model does not target. All dollar values are expressed in constant 2020 U.S. dollars using the CPI index.

Consistently with the model definition, I group household assets in three main categories $k \in \{B, S, H\}$: safe assets, equity, and (net) housing. Importantly, consistent with the model definition of equity, private businesses are excluded from all variable definitions. I therefore define the wealth measure as net worth excluding business equity. With respect to liabilities, I bundle mortgages into the value of housing — which is therefore expressed in net terms — and consumer debt in safe assets.² Table C.1 maps SCF survey variables to the three model asset classes. Business equity, vehicles, life insurance cash value, and other miscellaneous assets are excluded because they have no model counterpart.

Within each SCF wave, I rank households into groups d of the weighted distribution of w_i , using SCF sample weights ϖ_i — the survey weight assigned to household i . The portfolio share of asset class k in group d is then the ratio of weighted sums:

$$\frac{\sum_{i \in d} \varpi_i w_i^k}{\sum_{i \in d} \varpi_i w_i},$$

where w_i^k is household i ’s net holding of asset class k . Households with $|w_i| < \$500$ are excluded from the ratio computations to avoid dividing by near-zero wealth; this affects fewer than 0.3% of

²Since around 20-25% of households have negative net wealth, for the computation of the Gini coefficient I use the adjustment of Raffinetti et al. (2015) which extends the standard Gini to distributions containing negative values by normalizing with the sum of the absolute positive and negative totals rather than the mean.

Table C.1: Asset Decomposition: SCF Variables to Model Asset Classes

MODEL ASSET CLASS	SCF COMPONENTS
<i>Equity</i>	Directly held stocks, stock mutual funds, equity in IRAs/Keoghs and thrift plans, other managed assets with equity interest
<i>Safe (net)</i>	Transaction accounts, CDs, bonds, savings bonds, non-equity retirement allocations, minus consumer debt ^a
<i>Housing (net)</i>	Primary residence, nonresidential and other residential real estate, minus mortgage debt
<i>Excluded</i>	Business equity, vehicles, life insurance cash value, other financial assets

^a Consumer debt = total debt minus housing debt (mortgages and home equity loans). Non-equity retirement allocations include bond and balanced mutual funds held in IRA, Keogh, and defined-contribution accounts.

weighted observations. Equity market participation is defined as any positive stock-market exposure, including holdings through retirement accounts, while homeownership is defined as positive gross housing wealth. For the wealth-to-earnings ratio I convert pre-tax earnings by using the Heathcote et al. (2017) tax function $y_{\text{post}} = \lambda(y/\bar{y})^{1-\tau}\bar{y}$ with $\tau = 0.181$ and setting λ to match the average tax rate as in the model. All calibration targets are equal-weight cross-year averages across the 12 SCF waves.

CONSUMER EXPENDITURE SURVEY. Housing expenditure shares are constructed from the Consumer Expenditure Survey following Aguiar and Bils (2015). Housing expenditure is defined as actual rent for renters and self-reported rental equivalence for owners; the model targets the expenditure share schedule by decile and the average expenditure elasticity of housing, estimated at 0.92 by Aguiar and Bils.

RETURNS DATA. Asset return data come from the MacroHistory Database (MHDB) of Jordà et al. (2019). The safe rate is taken to be the average of the returns on Treasury bills and government bonds. The equity return is the total return on the aggregate equity index including dividends. Both series are deflated to real returns using the corresponding CPI series from the same database.

Since the model calibrates on net home equity — gross housing minus mortgage debt — while the MHDB housing series measures capital gains on the unlevered asset, I convert the raw capital-gain series to a leverage-adjusted return on equity using the following formula:³

$$r_t^H = \frac{r_t^{\text{cg}} - i_t(1 - \tau_t^m) \cdot \text{LTV}_t}{1 - \text{LTV}_t}, \quad (\text{C.1})$$

where r_t^{cg} is the nominal housing capital gain from the MHDB, τ_t^m is the federal marginal income tax rate (Barro and Redlick 2011), and i_t and LTV_t are the nominal mortgage rate and the aggregate loan-to-value ratio from FRED. The numerator captures the net return on the gross asset after subtracting the tax-deductible cost of mortgage financing; dividing by $1 - \text{LTV}_t$ scales this net return by the household's equity stake. The mortgage rate i_t is blended from two different data series to cover a longer sample period: I use FHA30 only for 1964–1970, an equal-weight average of

³Since the housing capital-gains series is in nominal terms, the resulting levered return is then deflated by the CPI as the other return series.

FHA30 and MORTGAGE30US for 1971–2000, and MORTGAGE30US only for 2001–2020.⁴

The effective sample period for returns is 57 years, from 1964 to 2020. The main constraint on sample size is the housing leverage adjustment in equation (C.1), which requires the FHA30 mortgage rate series available from FRED only from 1964 onward. The implied average returns, volatilities, and correlations are reported in table 2 in the main text. The corresponding matrix of sensitivities σ in equation (4) — the Cholesky factor of the sample covariance matrix, with assets ordered as (safe, equity, housing) — is:

$$\sigma = \begin{bmatrix} 0.058 & . & . \\ 0.019 & 0.156 & . \\ -0.009 & 0.010 & 0.060 \end{bmatrix}$$

D. Proofs

D.1. Proof of Proposition 1

We want to show that $g_t(\mathbf{x})$, the cross-sectional density over the household state vector $\mathbf{x} = (a, h, z_1, z_2, p, m)$, satisfies equation (10) which we restate here for convenience:

$$dg_t(\mathbf{x}) = \{ \mathcal{A}^* g_t(\mathbf{x}) + \zeta (\Psi_t(\mathbf{x}) - g_t(\mathbf{x})) \} dt - \partial_{\mathbf{x}} \{ [\sigma_{\mathbf{x}} d\mathbf{W}_t] g_t(\mathbf{x}) \}.$$

The proof is rather standard in the math literature, the idea behind it is to study the dynamics of a population with a finite number of agents and then take the limit as the number of agents goes to infinity. We start by setting up the problem with a finite number of agents N . Then, we apply Itô’s lemma to describe the evolution of a “test” density, take the limit as $N \rightarrow \infty$, apply integration by parts and get a differential equation for g_t .⁵

In the N agents case, we can define the empirical density function as:

$$g_t^N(\mathbf{x}) \equiv \frac{1}{N} \sum_{i=1}^N \delta(\mathbf{x} - \mathbf{x}_{i,t})$$

where δ is the Dirac-delta measure. Our object of interest, the population density for a continuum of agents is going to be given by:

$$g_t(\mathbf{x}) = \lim_{N \rightarrow \infty} g_t^N(\mathbf{x}).$$

Defining $\langle \mu, f \rangle \equiv \int f d\mu$ between a measure μ and a test function f we note that, when evalu-

⁴FHA30 is the “30-Year FHA Mortgage Rate: Secondary Market” series, available at monthly frequency for 1964–2000 and discontinued thereafter. MORTGAGE30US is the series “30-Year Fixed Rate Mortgage Average in the United States,” available at weekly frequency from 1971. Both series come from the FRED database; I convert each to annual frequency by averaging all observations within a calendar year. The loan-to-value ratio LTV_t equals one minus owners’ equity as a share of household real estate: FRED series HOEREPHRE.

⁵Notice that, because asset returns follow an exogenous process, each household’s drift and diffusion are functions only of its own state and of the common price path — they do not depend on g_t — so g_t evolves according to a linear stochastic partial differential equation driven by the *finite-dimensional* common shocks, with no price or return feedback from the distribution and no master equation to solve.

ated against the empirical measure g_t^N this corresponds to the empirical average (i.e. $\langle g_t^N, f \rangle = N^{-1} \sum_i f(\mathbf{x}_{i,t})$), whereas once the measure carries a density g this collapses to the form $\langle g, f \rangle = \int g f d\mathbf{x}$.⁶

Throughout the proof, I make use of two assumptions: First, conditional on the common shock path, the cross-sectional empirical measure admits a well-defined density limit g_t . Second, $\boldsymbol{\mu}$, $\boldsymbol{\sigma}_{\mathbf{x}}$, and g_t decay fast enough at the boundaries of the continuous state space that the boundary terms produced by integration by parts vanish.

ITÔ'S LEMMA. Let f be twice continuously differentiable with compact support in the continuous coordinates (a, h, z_1, z_2) on the region above the borrowing limit, $a > a_{\min}$, and bounded in the discrete ones. Applying Itô's formula for jump-diffusions to $f(\mathbf{x}_{i,t})$ under equations (5) to (7) together with the participation and housing jump kernels of appendix A of the main text yields

$$df(\mathbf{x}_{i,t}) = (\mathcal{A}_c f + \mathcal{A}_J f)(\mathbf{x}_{i,t}) dt + (\nabla_{\mathbf{x}} f)' \boldsymbol{\sigma}_{\mathbf{x}}(\mathbf{x}_{i,t}) d\mathbf{W}_t + dM_{i,t}, \quad (\text{D.1})$$

The first term, the average rate at which f changes, has both a continuous and a jump piece. The continuous piece is given by $\mathcal{A}_c f = \boldsymbol{\mu} \cdot \nabla_{\mathbf{x}} f + \frac{1}{2} \text{Tr}[\boldsymbol{\sigma}_{\mathbf{x}} \boldsymbol{\sigma}_{\mathbf{x}}' \nabla_{\mathbf{x}}^2 f]$ and it includes the drift $\boldsymbol{\mu}$ of equations (5) to (7) as well as the second-order term from the common return shocks, which loads only on wealth and housing (since z_1 and z_2 have no diffusion). Equivalently, $\mathcal{A}_c = \mathcal{L}_c - \sum_j \beta_j z_j \partial_{z_j}$, where $\mathcal{L}_c$ is the continuous generator in appendix A of the main text. The jump piece $\mathcal{A}_J f$ collects the household's idiosyncratic Poisson events and is given by:

$$\mathcal{A}_J f = \underbrace{\sum_{j \in \{1,2\}} \mathcal{Z}_j f}_{\text{earnings redraws } \mathcal{Z}_J f} + \underbrace{\mathbf{1}_{\{p=0\}} \mathcal{P}^0 f + \mathbf{1}_{\{p=1\}} \mathcal{P}^1 f}_{\text{participation } \mathcal{P} f} + \underbrace{\mathbf{1}_{\{h=0\}} \mathcal{H}^0 f + \mathbf{1}_{\{h>0\}} \mathcal{H}^+ f}_{\text{housing } \mathcal{H} f}.$$

where each term is defined as in appendix A of the main text: a participation opportunity arrives at rate λ_0^p for non-participants or λ_1^p for participants, and a housing opportunity at rate λ^h for renters and owners; on arrival the household relocates to the policy-chosen destination encoded in the matching kernel ($\mathcal{K}_0^p$ or $\mathcal{K}_1^p$ for participation, $\mathcal{K}_B^h$ or $\mathcal{K}_S^h$ for housing) exactly as in equation (9). The last two terms in equation (D.1) are the exposure of the households' state vector to the common return shocks and the idiosyncratic random jumps, respectively. Specifically, the last term $dM_{i,t}$ is the difference between the household's realized path and its average (already accounted for in $\mathcal{A}_J$). In short, $\mathcal{A}_J$ and $dM_{i,t}$ are the average and the realization of the same jumps.

Taking the cross-sectional limit of $\langle g_t^N, f \rangle$ we therefore recover:⁷

$$d\langle g_t, f \rangle = \langle g_t, (\mathcal{A}_c + \mathcal{A}_J) f \rangle dt + \langle \zeta(\Psi_t - g_t), f \rangle dt + \langle g_t, (\nabla_{\mathbf{x}} f)' \boldsymbol{\sigma}_{\mathbf{x}} d\mathbf{W}_t \rangle \quad (\text{D.2})$$

⁶In this $L^2/\text{density}$ form $\int g f d\mathbf{x}$, the integral runs over the continuous coordinates (a, h, z_1, z_2) while the sum — under the counting measure — runs over the discrete coordinates p and m .

⁷The key observation in the passage to the limit is that, conditional on the common shock path $\mathbf{W}_t$, the idiosyncratic jump martingales $dM_{i,t}$ are independent across households (their coefficients do not depend on g_t) so the jump term $N^{-1} \sum_i dM_{i,t}$ washes out in the aggregate, whereas the common shock $d\mathbf{W}_t$ is shared across all households and therefore survives — which is precisely why g_t is itself stochastic. The same argument applies to the death and birth events, and so to the demographic term in equation (D.2).

where the middle term in equation (D.2) is the demographic flow coming from death and birth: over an interval dt a fraction ζdt of households die and are replaced by a newborn household drawn from the birth distribution Ψ_t as described in section 2.1. Each replacement therefore swaps f evaluated at the state at the moment of death for f evaluated at the state at the moment of birth. In the cross-sectional limit, death and birth therefore add

$$\zeta(\langle \Psi_t, f \rangle - \langle g_t, f \rangle) dt = \langle \zeta(\Psi_t - g_t), f \rangle dt$$

to $d\langle g_t, f \rangle$. Notice that, for this term, no operator acts on f so — unlike the terms above — it needs no integration by parts and will pass into equation (10) directly.

THE ADJOINT OPERATOR. Equation (D.2) above pairs g_t against a test function f . To obtain an equation for g_t itself, we need to rewrite each term so that the operators act on g_t rather than on f , which amounts to finding the adjoint of $\mathcal{A}$ where, for each generator $\mathcal{A}$, its *adjoint* $\mathcal{A}^*$ is the operator satisfying $\langle g, \mathcal{A}f \rangle = \langle \mathcal{A}^*g, f \rangle$ for every test function f . To do that, we can start by performing integration by parts on the continuous states.

The continuous generator collects the first-order drifts and the second-order common-noise diffusion,

$$\mathcal{A}_c f = \sum_{k \in \{a, h, z_1, z_2\}} \mu_k \partial_k f + \frac{1}{2} \sum_{k, l \in \{a, h, z_1, z_2\}} \Sigma_{kl} \partial_{kl} f, \quad \Sigma \equiv \sigma_{\mathbf{x}} \sigma'_{\mathbf{x}},$$

where $\partial_k \equiv \partial_{x_k}$ and Σ is the instantaneous covariance of the continuous states.⁸

Pairing against g_t and integrating by parts — once for each drift, twice for the second-order term — we get:

$$\begin{aligned} \langle g_t, \mu_k \partial_k f \rangle &= \left[\mu_k g_t f \right]_{\Gamma} - \langle \partial_k [\mu_k g_t], f \rangle, \\ \langle g_t, \frac{1}{2} \Sigma_{kl} \partial_{kl} f \rangle &= \left[\frac{1}{2} \Sigma_{kl} g_t \partial_l f - \frac{1}{2} \partial_k [\Sigma_{kl} g_t] f \right]_{\Gamma} + \langle \frac{1}{2} \partial_{kl} [\Sigma_{kl} g_t], f \rangle, \end{aligned}$$

where $[\cdot]_{\Gamma}$ denotes the terms evaluated on the boundary Γ of the continuous domain Ω . These boundary terms all vanish: as a , h , and z_j go to infinity, the test function f and its first derivatives are all zero because f has compact support. For homeowners, at $h = 0$ the drift and diffusion terms all vanish because they're multiplicative in h . At the borrowing constraint, liquid wealth still carries return risk ($\Sigma_{aa} = a^2 \sigma_a \sigma'_a$ is positive since $a_{\min}$ can be negative), so the drift and diffusion coefficients do not vanish there; the boundary terms are nonetheless zero because f is zero in a neighborhood of $a_{\min}$. At $a_{\min}$ itself, I impose a reflecting boundary condition: a household whose wealth would otherwise fall below the borrowing limit stays at the limit, so that no mass leaves the state space; this is also the boundary condition used in the numerical solution (online

⁸Since earnings have zero loadings on the common shocks, Σ has a single nonzero block,

$$\Sigma_{aa} = a^2 \sigma_a \sigma'_a, \quad \Sigma_{hh} = h^2 \sigma_3 \sigma'_3, \quad \Sigma_{ah} = \Sigma_{ha} = ah \sigma_a \sigma'_3,$$

and $\Sigma_{kl} = 0 \ \forall k, l \in \{z_1, z_2\}$; while the off-diagonal Σ_{ah} is nonzero because wealth and housing load on the same correlated returns (equation (4)).

appendix F.2.4).⁹

Collecting the integral terms gives

$$\mathcal{A}_c^* g = - \sum_k \partial_k [\mu_k g] + \frac{1}{2} \sum_{k,l} \partial_{kl} [\Sigma_{kl} g],$$

so that we can rewrite the first term of equation (D.2) as:

$$\langle g_t, \mathcal{A}_c f \rangle dt = \langle \mathcal{A}_c^* g_t, f \rangle dt.$$

To derive the adjoint of the jump component, $\mathcal{A}_j^*$, let us first start by taking the adjoint of the earnings generators. At rate λ_j , each component z_j is reset to an independent draw from $\mathcal{N}(0, \sigma_j^2)$, whose density we write ϕ_j . The generator $\mathcal{Z}_j$ is as in appendix A of the main text:

$$\mathcal{Z}_j f = \lambda_j \int [f(\mathbf{x}_{j'}) - f(\mathbf{x})] \phi_j(z'_j) dz'_j$$

where $\mathbf{x}_{j'}$ denotes the state $\mathbf{x}$ with its j -th earnings component set to z'_j , as in appendix A of the main text. Additionally, let us denote by $\mathbf{x}_{-j}$ the vector of all $\mathbf{x}$ coordinates minus coordinate z_j (e.g. $\mathbf{x}_{-1} = (a, h, z_2, p, m)$), and by $\bar{g}_j$ the marginal of g over coordinate z_j , that is $\bar{g}_j \equiv \int g dz_j$.

Then, to find $\mathcal{Z}_j^*$:

$$\begin{aligned} \langle g, \mathcal{Z}_j f \rangle &= \lambda_j \int g(\mathbf{x}) [f(\mathbf{x}_{j'}) - f(\mathbf{x})] \phi_j(z'_j) dz'_j d\mathbf{x} \\ &\stackrel{(1)}{=} \lambda_j \int \bar{g}_j \phi_j(z'_j) f(\mathbf{x}_{j'}) dz'_j d\mathbf{x}_{-j} - \lambda_j \langle g, f \rangle \\ &\stackrel{(2)}{=} \lambda_j \langle \phi_j(z_j) \bar{g}_j, f \rangle - \lambda_j \langle g, f \rangle \\ &= \lambda_j \langle \phi_j(z_j) \bar{g}_j - g, f \rangle \end{aligned}$$

where in (1) we have split the “redraw” and “stay” terms and integrated over z_j , while in (2) we have simply renamed $z'_j \rightarrow z_j$ so that $\mathbf{x}_{j'} \rightarrow \mathbf{x}_j$ and $dz'_j d\mathbf{x}_{-j} \rightarrow d\mathbf{x}$. Hence:

$$\mathcal{Z}_j^* g = \lambda_j (\phi_j(z_j) \bar{g}_j - g).$$

To take the adjoint of the participation and housing generators, let us start by observing that each adjustment opportunity is a jump process: at a given Poisson rate the household is offered a chance to change his homeownership or participation state, and on acceptance the household jumps to the optimally-chosen destination. Since the participation structure in the two markets is exactly symmetric and the only thing that changes is the cost of entering/exiting each market, their treatment is the same and we here use participation as the running example. Recall from appendix A of the main text that each generator was written as $\mathcal{J} = \lambda(\mathcal{K} - \mathcal{I})$ where the kernel $\mathcal{K}$ returns the value of

⁹Because the reflection acts only at $a_{\min}$, where f is zero, it does not enter Itô’s formula in equation (D.1), and the derivation therefore establishes equation (10) for $a > a_{\min}$.

responding optimally.

For a non-participant $\mathbf{x} = (a, h, z_1, z_2, 0, m)$, the kernel of appendix A of the main text is

$$(\mathcal{K}v)(\mathbf{x}) = \max \{ v(a - \kappa_0^p, h, z_1, z_2, 1, m), v(\mathbf{x}) \}.$$

Let T be the “relocation” map that the optimal policy of equation (9) selects:

$$T(\mathbf{x}) = \begin{cases} (a - \kappa_0^p, h, z_1, z_2, 1, m) & \text{if } v(a - \kappa_0^p, h, z_1, z_2, 1, m) \geq v(\mathbf{x}), \\ \mathbf{x} & \text{otherwise,} \end{cases}$$

so that, by definition, $(\mathcal{K}v)(\mathbf{x}) = v(T(\mathbf{x}))$. Since in the KFE policy functions are taken as given, T is a fixed deterministic mapping and the kernel reduces to composition with T :

$$\mathcal{J}f = \lambda(\mathcal{K}f - f), \quad (\mathcal{K}f)(\mathbf{x}) = f(T(\mathbf{x})).$$

To find the adjoint, change variables along T :

$$\langle g, \mathcal{K}f \rangle = \int g(\mathbf{x}) (\mathcal{K}f)(\mathbf{x}) \, d\mathbf{x} = \int g(\mathbf{x}) f(T(\mathbf{x})) \, d\mathbf{x} = \int (T^*g)(\mathbf{y}) f(\mathbf{y}) \, d\mathbf{y} = \langle T^*g, f \rangle,$$

where T^*g — the density g *relocated* along T — collects at each destination $\mathbf{y}$ the mass of every source the policy sends there,

$$(T^*g)(\mathbf{y}) = \sum_{\mathbf{x}: T(\mathbf{x})=\mathbf{y}} g(\mathbf{x}).$$

For entry into participation the only source mapped to a participant state $(a, h, z_1, z_2, 1, m)$ is the non-participant household who pays the entry cost, which is a simple change in wealth

$$(T^*g)(a, h, z_1, z_2, 1, m) = g(a + \kappa_0^p, h, z_1, z_2, 0, m)$$

if entering is optimal at $(a + \kappa_0^p, h, z_1, z_2, 0, m)$, and zero otherwise. Hence $\mathcal{K}^*g = T^*g$ and $\mathcal{J}^*g = \lambda(T^*g - g)$: the adjoint strips mass λg from each state and moves it where assigned by the policy.

We can similarly build the adjoint operators for the other participation state and for housing.¹⁰ I write the assembled adjoints as $\mathcal{P}^*$ and $\mathcal{H}^*$, matching the sub-operators $\mathcal{P}$ and $\mathcal{H}$ of appendix A of the main text.¹¹ Hence, defining $\mathcal{A}_J^*g = \mathcal{Z}_J^*g + \mathcal{P}^*g + \mathcal{H}^*g$, we can write the second term in equation (D.2) as

$$\langle g_t, \mathcal{A}_J f \rangle \, dt = \langle \mathcal{A}_J^* g_t, f \rangle \, dt.$$

And, assembling the continuous adjoint with the earnings, participation, and housing jump adjoints

¹⁰Specifically, for housing, T^*g is defined as $\langle T^*g, f \rangle = \langle g, f \circ T \rangle$ because all renters have $h = 0$.

¹¹Note that the permanent discount-rate type m enters differently: because there are no transitions in m during a household’s life, the generator is block-diagonal across types and carries no jump adjoint; the birth operator Ψ is the only channel that mixes discount-rate types, assigning newborns type m with probability $\pi_\rho(m)$.

derived above, the adjoint of the full generator is:¹²

$$\begin{aligned} \mathcal{A}^* g = & \underbrace{-\partial_a[\mu_a g] - \partial_h[\mu_h g] + \partial_{z_1}[\beta_1 z_1 g] + \partial_{z_2}[\beta_2 z_2 g] + \frac{1}{2} \sum_{i,j \in \{a,h\}} \partial_{ij}[\Sigma_{ij} g]}_{\mathcal{A}_c^* g} \\ & + \underbrace{\mathcal{Z}_J^* g + \mathcal{P}^* g + \mathcal{H}^* g}_{\mathcal{A}_J^* g}. \end{aligned} \quad (\text{D.3})$$

It now only remains to take the adjoint of the surviving common-noise martingale in the evolution of $\langle g_t, f \rangle$ above. Because $\sigma_{\mathbf{x}}$ has zero rows for the earnings components (only liquid wealth and housing carry common-return loadings $a\sigma_a$ and $h\sigma_3$) only the a and h rows of $(\nabla_{\mathbf{x}} f)' \sigma_{\mathbf{x}} d\mathbf{W}_t$ appear below. Integrating by parts in a and h , gives:

$$\begin{aligned} \langle g_t, (\nabla_{\mathbf{x}} f)' \sigma_{\mathbf{x}} d\mathbf{W}_t \rangle &= \int g_t \left[(a\sigma_a \cdot d\mathbf{W}_t) \partial_a f + (h\sigma_3 \cdot d\mathbf{W}_t) \partial_h f \right] d\mathbf{x} \\ &= \left[(a\sigma_a \cdot d\mathbf{W}_t) g_t f + (h\sigma_3 \cdot d\mathbf{W}_t) g_t f \right]_{\Gamma} \\ &\quad - \int \left\{ \partial_a[(a\sigma_a \cdot d\mathbf{W}_t) g_t] + \partial_h[(h\sigma_3 \cdot d\mathbf{W}_t) g_t] \right\} f d\mathbf{x} \\ &= - \left\langle \partial_a[(a\sigma_a \cdot d\mathbf{W}_t) g_t] + \partial_h[(h\sigma_3 \cdot d\mathbf{W}_t) g_t], f \right\rangle \\ &= \left\langle -\partial_{\mathbf{x}}\{[\sigma_{\mathbf{x}} d\mathbf{W}_t] g_t\}, f \right\rangle, \end{aligned}$$

where the boundary terms $[\cdot]_{\Gamma}$ vanish for the same reasons as the drift and diffusion terms above. Collecting this last term, it is also going to be convenient to name the diffusion adjoint $\mathcal{B}^*$ — the common-noise counterpart of $\mathcal{A}^*$, which we reuse for the numerical scheme in online appendix F.3.2 — through its action against the common shock,

$$\mathcal{B}^* g d\mathbf{W}_t \equiv -\partial_{\mathbf{x}}\{[\sigma_{\mathbf{x}} d\mathbf{W}_t] g\} = -\partial_a[(a\sigma_a \cdot d\mathbf{W}_t) g] - \partial_h[(h\sigma_3 \cdot d\mathbf{W}_t) g].$$

CONCLUSION. Substituting the continuous, jump, and diffusion adjoints back into equation (D.2), every operator now acts on g_t rather than on f :

$$d\langle g_t, f \rangle = \langle \mathcal{A}^* g_t + \zeta(\Psi_t - g_t), f \rangle dt + \langle \mathcal{B}^* g_t d\mathbf{W}_t, f \rangle.$$

¹²Because the discount type m never changes within a household's life, $\mathcal{A}^*$ is block-diagonal in m : writing $g_m \equiv g(\cdot, m)$,

$$(\mathcal{A}^* g)_m = \mathcal{A}_m^* g_m, \quad m = 1, \dots, N_{\rho},$$

where $\mathcal{A}_m^*$ is equation (D.3) with the policies — and hence the drifts, diffusion, and relocation maps — evaluated at the type- m value function, which differs across types only through the discount rate ρ_m . The types are coupled only at birth, through Ψ .

Since the test function f is fixed (it is neither random nor time-dependent) we also have $d\langle g_t, f \rangle = \langle dg_t, f \rangle$.¹³ Hence:

$$\left\langle dg_t - [\mathcal{A}^* g_t + \zeta(\Psi_t - g_t)] dt - \mathcal{B}^* g_t d\mathbf{W}_t, f \right\rangle = 0,$$

which must hold for every admissible test function f and it must therefore be the case that¹⁴

$$dg_t(\mathbf{x}) = [\mathcal{A}^* g_t(\mathbf{x}) + \zeta(\Psi_t(\mathbf{x}) - g_t(\mathbf{x}))] dt + \mathcal{B}^* g_t(\mathbf{x}) d\mathbf{W}_t,$$

which is exactly our KFE (10). □

D.2. Proof of Proposition 2

We first collect notation used throughout the proofs of section 3. Recall the non-housing expenditure share $\alpha(x) = c(x)/x$, with $x = c + r_n n$ total expenditure, and let

$$\beta_h \equiv 1 - \varepsilon_h^{-1}, \quad \beta_c \equiv 1 - \varepsilon_c^{-1}, \quad P(\alpha) \equiv \alpha \varepsilon_c + (1 - \alpha) \varepsilon_h, \quad Q(\alpha) \equiv \alpha \beta_h + (1 - \alpha) \beta_c.$$

Under the assumption $\varepsilon_h < \varepsilon_c < 1$ we have $\beta_h < \beta_c < 0$, while $P(\alpha) > 0$ and $Q(\alpha) < 0$ for all $\alpha \in [0, 1]$. From the addi-log intratemporal utility $u(c, n)$ of equation (3), its derivatives are

$$u_c = (1 - \omega) u^{1/\varepsilon_h} c^{-1/\varepsilon_c}, \quad u_n = \omega u^{1/\varepsilon_h} n^{-1/\varepsilon_h},$$

and the intratemporal Euler equation reads $u_n = r_n u_c$.

We need to prove that $\partial\alpha(x)/\partial x > 0$ and that

$$\lim_{x \rightarrow 0} \alpha(x) = 0 \quad \text{and} \quad \lim_{x \rightarrow \infty} \alpha(x) = 1.$$

From the definition of $\alpha(x)$, we immediately get:

$$\alpha'(x) = \frac{1}{x} (c'(x) - \alpha(x))$$

and, given x is always positive, proving that $\alpha'(x) > 0$ boils down to showing $c'(x) > \alpha(x)$.

Using the derivatives above, the intratemporal Euler equation reads:

$$r_n = \frac{u_n}{u_c} = \frac{\omega}{1 - \omega} \frac{n^{-1/\varepsilon_h}}{c^{-1/\varepsilon_c}}.$$

We can then use the intratemporal Euler equation and the definition of x to find:

$$\alpha(x) = \frac{c}{x} = \frac{1}{1 + r_n \left(\frac{1-\omega}{\omega} r_n \right)^{-\varepsilon_h} c(x)^{\frac{\varepsilon_h}{\varepsilon_c} - 1}}. \quad (\text{D.4})$$

¹³For the dt term this is a direct consequence of linearity of the integral; for the $d\mathbf{W}_t$ term, this formally requires to invoke the stochastic Fubini theorem, which lets the integral pass through the Itô integral against the common noise.

¹⁴This is a straightforward application of the fundamental lemma of the calculus of variations: a locally integrable function that integrates to zero against every compactly supported f is zero almost everywhere.

Then, to find $c'(x)$, we use the fact that $c(x)$ solves:

$$\begin{aligned} x &= c(x) + r_n n(c(x)) \\ &= c(x) + r_n \left(\frac{1-\omega}{\omega} r_n \right)^{-\varepsilon_h} c(x)^{\varepsilon_h/\varepsilon_c} \end{aligned} \quad (\text{D.5})$$

Hence:

$$1 = c'(x) + r_n \left(\frac{1-\omega}{\omega} r_n \right)^{-\varepsilon_h} \frac{\varepsilon_h}{\varepsilon_c} c(x)^{\frac{\varepsilon_h}{\varepsilon_c}-1} c'(x)$$

which immediately gets us $c'(x)$:

$$c'(x) = \frac{1}{1 + r_n \left(\frac{1-\omega}{\omega} r_n \right)^{-\varepsilon_h} \frac{\varepsilon_h}{\varepsilon_c} c(x)^{\frac{\varepsilon_h}{\varepsilon_c}-1}}. \quad (\text{D.6})$$

Showing that $c'(x) > \alpha(x)$ therefore requires that

$$1 + r_n \left(\frac{1-\omega}{\omega} r_n \right)^{-\varepsilon_h} \frac{\varepsilon_h}{\varepsilon_c} c(x)^{\frac{\varepsilon_h}{\varepsilon_c}-1} < 1 + r_n \left(\frac{1-\omega}{\omega} r_n \right)^{-\varepsilon_h} c(x)^{\frac{\varepsilon_h}{\varepsilon_c}-1}$$

which is satisfied as long as $\varepsilon_h < \varepsilon_c$.

Now we have left to prove that $\alpha_0 = 0$ and $\alpha_\infty = 1$. From equation (D.5) above, total expenditure is a strictly increasing function of c so that $x \rightarrow 0$ if and only if $c \rightarrow 0$ and $x \rightarrow \infty$ if and only if $c \rightarrow \infty$. Because $\varepsilon_h < \varepsilon_c$, the exponent of $c(x)$ in equation (D.4) is negative: $c^{\frac{\varepsilon_h}{\varepsilon_c}-1} \rightarrow \infty$ as $c \rightarrow 0$, and $c^{\frac{\varepsilon_h}{\varepsilon_c}-1} \rightarrow 0$ as $c \rightarrow \infty$. It immediately follows that

$$\alpha_0 \equiv \lim_{x \rightarrow 0} \alpha(x) = 0 \quad \text{and} \quad \alpha_\infty \equiv \lim_{x \rightarrow \infty} \alpha(x) = 1.$$

□

D.3. Proof of Proposition 3

INDIRECT UTILITY AND ENVELOPE IDENTITY. Let $U(x) \equiv u(c(x), n(x))$ denote the indirect utility of total expenditure $x = c + r_n n$ along the optimal intratemporal expansion path. The envelope theorem gives

$$U'(x) = u_c c'(x) + u_n n'(x) = u_c [c'(x) + r_n n'(x)] = u_c,$$

where we have used the intratemporal Euler equation $u_n = r_n u_c$ and the budget derivative $c'(x) + r_n n'(x) = 1$.

The elasticity of intertemporal substitution is defined with respect to total expenditure growth $d \ln x_t / dt$:

$$\text{EIS}(x) \equiv - \frac{\partial (d \ln x_t / dt)}{\partial (d \ln \Lambda_t / dt)},$$

where $\Lambda(t)$ is the utility gradient, defined following Schroder and Skiadas (1999) as:

$$\Lambda(t) = \exp \left[\int_0^t f_v(u_s, v_s) ds \right] f_u(u_t, v_t) U'(x_t).$$

Here f_u and f_v denote partial derivatives of the intertemporal aggregator f defined in equation (2), and we write $U'(x_t)$ in place of u_c to make explicit that the gradient depends on t only through $x(t)$.

From the definition of Λ :

$$\frac{\dot{\Lambda}}{\Lambda} = f_v + \frac{\dot{f}_u}{f_u} + \frac{\dot{U}'(x)}{U'(x)}.$$

We expand each term in turn. For the second term, we use the chain rule on $f_u(u, v)$ along the optimal path. Since $u = U(x)$ along the optimal path, $\dot{u} = U'(x) \dot{x}$; therefore:

$$\frac{\dot{f}_u}{f_u} = \frac{f_{uu}}{f_u} \dot{u} + \frac{f_{uv}}{f_u} \dot{v} = \frac{f_{uu}}{f_u} U'(x) \dot{x} + \frac{f_{uv}}{f_u} \dot{v}.$$

For the third term, since U' depends on t only through $x(t)$:

$$\frac{\dot{U}'(x)}{U'(x)} = \frac{U''(x)}{U'(x)} \dot{x}.$$

Collecting the terms proportional to $\dot{x}$:

$$\frac{\dot{\Lambda}}{\Lambda} = f_v + \frac{f_{uv}}{f_u} \dot{v} + \left[\frac{f_{uu}}{f_u} U'(x) x + \frac{x U''(x)}{U'(x)} \right] \frac{\dot{x}}{x}.$$

Notice that our object of interest is precisely the coefficient of $\dot{x}/x$, whose negative is the inverse of the EIS, that is:

$$-\frac{1}{\text{EIS}} = \frac{f_{uu}}{f_u} U'(x) x + \frac{x U''(x)}{U'(x)}.$$

Using the definition of f in equation (2) and $u = U(x)$ along the optimal path, direct calculation gives:

$$\frac{f_{uu}}{f_u} = -\frac{1}{\psi U(x)}, \quad \frac{f_{uv}}{f_u} = \frac{\psi^{-1} - \gamma}{1 - \gamma} \cdot \frac{1}{v}.$$

Therefore:

$$\frac{f_{uu}}{f_u} \cdot U'(x) x = -\frac{1}{\psi} \cdot \frac{U'(x) x}{U(x)} \equiv -\frac{\eta_U(x)}{\psi}, \quad (\text{D.7})$$

where we define the expenditure-elasticity of the indirect utility:

$$\eta_U(x) \equiv \frac{U'(x) x}{U(x)}.$$

To obtain a closed form, write the addi-log aggregator as $u = (C + N)^{1/\beta_h}$, where $C \equiv (1 - \omega) \frac{\beta_h}{\beta_c} c^{\beta_c}$ and $N \equiv \omega n^{\beta_h}$, so that $u_c c/u = \frac{\beta_c}{\beta_h} \frac{C}{C+N}$ and $u_n n/u = \frac{N}{C+N}$. The intratemporal Euler equation and the budget constraint together imply $N/C = (\beta_c/\beta_h)(1 - \alpha)/\alpha$; substituting, and using the envelope

identity $U'(x) = u_c$, yields

$$\eta_U(\alpha) = \frac{\beta_c}{\alpha \beta_h + (1 - \alpha) \beta_c} = \frac{\beta_c}{Q(\alpha)}. \quad (\text{D.8})$$

In particular, $\eta_U(0) = 1$ and $\eta_U(1) = \beta_c/\beta_h = A$.

Differentiating $U'(x) = u_c(c(x), n(x))$ along the optimal expansion path:

$$U''(x) = u_{cc} c'(x) + u_{cn} n'(x).$$

Differentiating u_c and u_n , using equation (3) and the intratemporal Euler equation $u_n = r_n u_c$, gives the second-order partial derivatives

$$u_{cc} = \frac{1}{\varepsilon_h} \frac{u_c^2}{u} - \frac{1}{\varepsilon_c} \frac{u_c}{c}, \quad u_{cn} = \frac{1}{\varepsilon_h} \frac{u_c u_n}{u}, \quad u_{nn} = \frac{1}{\varepsilon_h} \frac{u_n^2}{u} - \frac{1}{\varepsilon_h} \frac{u_n}{n}. \quad (\text{D.9})$$

Substituting into the expression for $U''(x)$:

$$U''(x) = \frac{1}{\varepsilon_h} \frac{u_c}{u} [u_c c'(x) + u_n n'(x)] - \frac{1}{\varepsilon_c} \frac{u_c}{c} c'(x).$$

The bracket $u_c c'(x) + u_n n'(x)$ equals $U'(x) = u_c$ by the envelope identity. Therefore:

$$U''(x) = \frac{1}{\varepsilon_h} \frac{u_c^2}{u} - \frac{u_c}{\varepsilon_c c} c'(x).$$

Dividing both sides by $U'(x) = u_c$ and multiplying by $-x$, the curvature $\mathcal{R}_U(x) \equiv -x U''(x)/U'(x)$ of the indirect utility is:

$$\mathcal{R}_U = \frac{\eta_c(x)}{\varepsilon_c} - \frac{\eta_U(x)}{\varepsilon_h},$$

where $\eta_c(x) \equiv \frac{c'(x)x}{c(x)}$ is the expenditure elasticity of non-housing consumption which, from equations (D.4) and (D.6), can be written as

$$\eta_c = \frac{\varepsilon_c}{\alpha \varepsilon_c + (1 - \alpha) \varepsilon_h} = \frac{\varepsilon_c}{P(\alpha)}.$$

Hence we obtain:

$$\mathcal{R}_U(\alpha) = \frac{1}{P(\alpha)} - \frac{\eta_U(\alpha)}{\varepsilon_h}. \quad (\text{D.10})$$

Substituting equation (D.7) and equation (D.10) in the coefficient of $\dot{x}/x$:

$$\frac{1}{\text{EIS}(\alpha)} = \frac{\eta_U(\alpha)}{\psi} + \mathcal{R}_U(\alpha) = \eta_U(\alpha) \left(\frac{1}{\psi} - \frac{1}{\varepsilon_h} \right) + \frac{1}{P(\alpha)},$$

which yields the expression for $\text{EIS}(x)$ stated in proposition 3. Note that $\text{EIS}(x) > 0$ for any $\psi > 0$: $\eta_U = \beta_c/Q > 0$ and, substituting equation (D.8) into equation (D.10),

$$\mathcal{R}_U(\alpha) = \frac{\alpha (\varepsilon_h - \varepsilon_c)}{\varepsilon_h P(\alpha) Q(\alpha)} \geq 0.$$

Further substituting $\eta_U(\alpha) = \beta_c/Q(\alpha)$ and multiplying numerator and denominator by $Q(\alpha)$ gives the closed form we work with below:

$$\text{EIS}(\alpha) = \frac{P(\alpha) Q(\alpha)}{Q(\alpha) + \beta_c \left(\frac{1}{\psi} - \frac{1}{\varepsilon_h} \right) P(\alpha)}. \quad (\text{D.11})$$

MONOTONICITY. Write $D(\alpha) \equiv Q(\alpha) + \beta_c \left(\frac{1}{\psi} - \frac{1}{\varepsilon_h} \right) P(\alpha)$ for the denominator of equation (D.11), so that $\text{EIS} = PQ/D$ and the quotient rule gives:

$$\frac{d \text{EIS}}{d\alpha} = \frac{(PQ)' D - (PQ) D'}{D^2}.$$

P , Q , and D are all affine in α , with slopes $P' = \varepsilon_c - \varepsilon_h$, $Q' = \beta_h - \beta_c$, and $D' = Q' + \beta_c \left(\frac{1}{\psi} - \frac{1}{\varepsilon_h} \right) P'$. Expanding the numerator:

$$\begin{aligned} (PQ)' D - (PQ) D' &= (P'Q + PQ') \left[Q + \beta_c \left(\frac{1}{\psi} - \frac{1}{\varepsilon_h} \right) P \right] - PQ \left[Q' + \beta_c \left(\frac{1}{\psi} - \frac{1}{\varepsilon_h} \right) P' \right] \\ &= P' Q^2 + \beta_c \left(\frac{1}{\psi} - \frac{1}{\varepsilon_h} \right) Q' P^2. \end{aligned}$$

Substituting P' and Q' , the sign of each element follows from the parameter restriction $\psi < \varepsilon_h < \varepsilon_c < 1$:

$$(PQ)' D - (PQ) D' = \underbrace{(\varepsilon_c - \varepsilon_h)}_{(+)} \underbrace{Q^2}_{(+)} + \underbrace{\beta_c}_{(-)} \underbrace{\left(\frac{1}{\psi} - \frac{1}{\varepsilon_h} \right)}_{(+)} \underbrace{(\beta_h - \beta_c)}_{(-)} \underbrace{P^2}_{(+)} > 0.$$

Since $D^2 > 0$, it follows that $d \text{EIS}/d\alpha > 0$ for all $\alpha \in [0, 1]$. Combined with $\alpha'(x) > 0$ from proposition 2, we obtain

$$\frac{\partial \text{EIS}}{\partial x} = \frac{d \text{EIS}}{d\alpha} \alpha'(x) > 0.$$

LIMITS. The limits follow from the closed form equation (D.11) and the boundary values $\alpha_0 = 0$ and $\alpha_\infty = 1$ established in proposition 2.

Setting $\alpha = 0$: we have $P(0) = \varepsilon_h$, $Q(0) = \beta_c$, and $\eta_U(0) = 1$. Substituting these into equation (D.11) immediately gives us:

$$\psi_0 \equiv \lim_{x \rightarrow 0} \text{EIS}(x) = \psi.$$

As $\alpha \rightarrow 0$ expenditure is allocated entirely to housing, and the model reduces to a single-good Duffie-Epstein economy with primitive parameters (γ, ψ) ; the EIS therefore recovers its one-good value ψ .

Setting $\alpha = 1$ we have $P(1) = \varepsilon_c$, $Q(1) = \beta_h$, and $\eta_U(1) = \beta_c/\beta_h$. Hence, some trivial algebra gets us:

$$\psi_\infty \equiv \lim_{x \rightarrow \infty} \text{EIS}(x) = \frac{1}{1 + A \left(\frac{1}{\psi} - 1 \right)},$$

where $A = \beta_c/\beta_h = \frac{1-\varepsilon_c^{-1}}{1-\varepsilon_h^{-1}}$. □

D.4. Proof of Lemma 1

In the simplified economy of section 3, the household's first-order condition for total expenditure equalizes the marginal flow value of expenditure to the marginal value of wealth:

$$f_u(U(x^*(a)), V(a)) U'(x^*(a)) = V'(a).$$

Since $\gamma > 1$ we have $V(a) < 0$ and hence $(1 - \gamma)V(a) > 0$, while $f_u > 0$ and $U' > 0$, so both sides of the first-order condition are strictly positive.

Differentiating $f(U(x), V(a))$ twice in x , the interior second-order condition is¹⁵

$$f_{uu} (U')^2 + f_u U'' < 0.$$

Because it is nonzero, the implicit function theorem guarantees that $x^*(a)$ is continuously differentiable: $\eta_x(a)$ is therefore well defined, and the first-order condition can be differentiated with respect to a .

Differentiating the first-order condition with respect to a :

$$V''(a) = [f_{uu} (U')^2 + f_u U''] \frac{dx^*}{da} + f_{uv} U' V'(a). \quad (\text{D.12})$$

As computed in the proof of proposition 3, we also have that, along the optimal path

$$\frac{f_{uu}}{f_u} = -\frac{1}{\psi U}, \quad \frac{f_{uv}}{f_u} = \frac{\psi^{-1} - \gamma}{(1 - \gamma)V}.$$

Using these together with the first-order condition $f_u U' = V'$ to eliminate f_u , the three terms become

$$f_{uu} (U')^2 \frac{dx^*}{da} = -\frac{U'}{\psi U} V' \frac{dx^*}{da}, \quad f_u U'' \frac{dx^*}{da} = \frac{U''}{U'} V' \frac{dx^*}{da}, \quad f_{uv} U' V' = \frac{\psi^{-1} - \gamma}{(1 - \gamma)V} (V')^2.$$

Then, multiplying equation (D.12) by $-a/V'(a)$ and rewriting by using the elasticities $\eta_x = \frac{a}{x^*} \frac{dx^*}{da}$, $\eta_U = x^* U'/U$, $\eta_V = a V'/V$ and the curvature $\mathcal{R}_U = -x^* U''/U'$, we get:

$$\mathcal{R}_V(a) = \eta_x(a) \left[\frac{\eta_U(x^*)}{\psi} + \mathcal{R}_U(x^*) \right] + \left(\gamma - \frac{1}{\psi} \right) \frac{\eta_V(a)}{1 - \gamma}.$$

Finally, it only remains to rewrite the static term as a function of the Hanoch curvature. From $H(x) = U(x)^{1-\gamma}/(1-\gamma)$ we have $H' = U^{-\gamma} U'$, so that

$$\mathcal{R}_H(x) \equiv -\frac{x H''(x)}{H'(x)} = \gamma \eta_U(x) + \mathcal{R}_U(x), \quad (\text{D.13})$$

and, substituting for $\mathcal{R}_U$ into the collected expression and using $\psi^{-1} - \gamma = -(\gamma - \psi^{-1})$ we finally

¹⁵The inequality is strict because $f_{uu} = -f_u/(\psi U) < 0$ and, as shown in online appendix D.3, $\mathcal{R}_U \geq 0$ which implies $U'' \leq 0$.

get an expression for relative risk aversion:

$$\mathcal{R}_V(a) = \eta_x(a) \mathcal{R}_H(x^*(a)) + \left(\gamma - \frac{1}{\psi} \right) \left(\frac{\eta_V(a)}{1 - \gamma} - \eta_x(a) \eta_U(x^*(a)) \right).$$

□

D.5. Proof of Proposition 4

CLOSED FORM. Recall from the proof of lemma 1 the decomposition equation (D.13), $\mathcal{R}_H(x) = \gamma \eta_U(x) + \mathcal{R}_U(x)$. In the proof of proposition 3 we found $\eta_U(\alpha) = \beta_c/Q(\alpha)$ and, from equation (D.10), $\mathcal{R}_U(\alpha) = \frac{1}{P(\alpha)} - \frac{\eta_U(\alpha)}{\varepsilon_h}$. Substituting,

$$\mathcal{R}_H = \gamma \frac{\beta_c}{Q} + \frac{1}{P} - \frac{\beta_c}{\varepsilon_h Q} = \frac{\gamma \beta_c P + Q - \beta_c P/\varepsilon_h}{P(\alpha) Q(\alpha)} = \frac{N(\alpha)}{P(\alpha) Q(\alpha)},$$

where

$$N(\alpha) \equiv \alpha \varepsilon_c (\beta_h + \beta_c (\gamma - 1)) + (1 - \alpha) \varepsilon_h \beta_c \gamma.$$

Expanding $N/(PQ)$ in α and $1 - \alpha$ reproduces the expression stated in proposition 4. Under $\varepsilon_h < \varepsilon_c < 1 < \gamma$ we have $N < 0$, $P > 0$, $Q < 0$, so $\mathcal{R}_H = N/(PQ) > 0$.

MONOTONICITY. To find the sign of $\partial \mathcal{R}_H / \partial x$ under $\varepsilon_h < \varepsilon_c < 1 < \gamma$ and $\gamma > \varepsilon_c / \varepsilon_h$, let's start by defining the intermediate ratio $S(\alpha) \equiv N(\alpha)/P(\alpha)$, so that $\mathcal{R}_H = S/Q$. Since N and P are affine in α ,

$$S'(\alpha) = \frac{N'(\alpha)P(\alpha) - N(\alpha)P'(\alpha)}{P(\alpha)^2} = \frac{\varepsilon_c \varepsilon_h (\beta_h - \beta_c)}{P(\alpha)^2} < 0,$$

since $\beta_h - \beta_c < 0$. The quotient rule applied to $\mathcal{R}_H = S/Q$ then gives

$$\frac{\partial \mathcal{R}_H}{\partial \alpha} = \frac{\beta_h - \beta_c}{P(\alpha)^2 Q(\alpha)^2} [\varepsilon_c \varepsilon_h Q(\alpha) - N(\alpha)P(\alpha)]. \quad (\text{D.14})$$

Since $\beta_h - \beta_c < 0$ and $P^2 Q^2 > 0$, the sign of $\partial \mathcal{R}_H / \partial \alpha$ is opposite to that of the term in square brackets, so it remains to show $\varepsilon_c \varepsilon_h Q - NP > 0$. Since N , P , and Q are affine in α , the left-hand side is a quadratic in α :

$$\varepsilon_c \varepsilon_h Q - NP = a_2 \alpha^2 + a_1 \alpha + a_0,$$

with

$$\begin{aligned} a_0 &= -\beta_c \varepsilon_h (\varepsilon_h \gamma - \varepsilon_c), \\ a_1 &= -2\beta_c \varepsilon_h \gamma (\varepsilon_c - \varepsilon_h), \\ a_2 &= (\varepsilon_c - \varepsilon_h) [\varepsilon_c (\beta_c - \beta_h) - \beta_c \gamma (\varepsilon_c - \varepsilon_h)]. \end{aligned}$$

$a_0 > 0$ because $-\beta_c > 0$ and $\varepsilon_h \gamma > \varepsilon_c$; $a_1 > 0$ because $-\beta_c > 0$ and $\varepsilon_c > \varepsilon_h$; and $a_2 > 0$ because $\varepsilon_c > \varepsilon_h$, $\beta_c - \beta_h > 0$, and $-\beta_c > 0$. All three coefficients are strictly positive, so $\varepsilon_c \varepsilon_h Q - NP > 0$

for all $\alpha \in [0, 1]$, hence $\partial \mathcal{R}_H / \partial \alpha < 0$. Since $\alpha'(x) > 0$ by proposition 2, we have

$$\frac{\partial \mathcal{R}_H}{\partial x} = \frac{\partial \mathcal{R}_H}{\partial \alpha} \alpha'(x) < 0. \quad (\text{D.15})$$

LIMITS. Finally, the limits follow from the boundary shares $\alpha_0 = 0$ and $\alpha_\infty = 1$ of proposition 2.

$$\gamma_0 \equiv \lim_{x \rightarrow 0} \mathcal{R}_H(x) = \frac{\varepsilon_h \beta_c \gamma}{\varepsilon_h \beta_c} = \gamma, \quad \gamma_\infty \equiv \lim_{x \rightarrow \infty} \mathcal{R}_H(x) = \frac{\varepsilon_c (\beta_h + \beta_c (\gamma - 1))}{\varepsilon_c \beta_h} = 1 + A(\gamma - 1).$$

Since $A \in (0, 1)$, $\gamma_\infty = 1 + A(\gamma - 1) < \gamma = \gamma_0$. □

D.6. Proof of Corollary 4.1

Proposition 4 characterizes the Hanoch curvature in closed form over the entire domain of expenditure and, in particular, gives its two limiting values: $\gamma_0 = \gamma$ and $\gamma_\infty = 1 + A(\gamma - 1)$. In this corollary I prove that those values apply to true risk aversion as well: at each end of the wealth distribution, $\mathcal{R}_V$ converges to the corresponding limit of $\mathcal{R}_H$.

I proceed by showing that, at each limit, the simplified economy of section 3 reduces to a one-good Merton economy whose curvature is defined as the corresponding limit of the Hanoch curvature and whose elasticity of intertemporal substitution is the corresponding limit from proposition 3. Then using the household's HJB equation I show that, in the presence of regularity conditions, $\mathcal{R}_V$ converges to the curvature of the one-good benchmark, i.e., to the limit of $\mathcal{R}_H$. More specifically, I start by setting out the two one-good benchmark economies and the required regularity conditions; I then solve the one-good HJB equation for $\mathcal{R}_V$, leaving as unknowns only the expenditure-to-wealth ratio $\kappa(a) = x^*(a)/a$ and the flow-utility term $f/(aV')$; then, I show that each unknown converges to its one-good benchmark counterpart and finally conclude the proof.

Because the same argument applies at both limits, the proof is written just once, treating only one of the two limits. Throughout, the hatted benchmark objects $(\hat{\gamma}, \hat{\psi}, \hat{\kappa}, \hat{\eta}_U, \hat{\eta}_V)$ denote limits of the corresponding variable at the limit under consideration; their exact values at the two limits are collected in the summary array below.

ONE-GOOD BENCHMARK ECONOMY. By proposition 2, as total expenditure x tends to 0 or ∞ the expenditure share $\alpha(x)$ converges to 0 or 1, and the intratemporal problem collapses to a single good: housing at $x \rightarrow 0$, non-housing consumption at $x \rightarrow \infty$. The benchmark economy at each limit is therefore a one-good economy, and the only preference inputs it requires are a curvature and an elasticity of intertemporal substitution.

The curvature of the one-good economy, $\hat{\gamma}$, is defined as the limit of the Hanoch curvature $\mathcal{R}_H(x)$ at the limit under consideration ($x \rightarrow 0$ at the bottom, $x \rightarrow \infty$ at the top); proposition 4 gives this limit in closed form. The elasticity of intertemporal substitution, $\hat{\psi}$, is defined analogously, as the corresponding limit of $\text{EIS}(x)$ from proposition 3.

It is also useful to define the remaining objects we need from the one-good benchmark economy, $(\hat{\eta}_U, \hat{\eta}_V, \hat{\kappa})$. First, the expenditure-elasticity of indirect utility, $\eta_U(x) = xU'(x)/U(x)$, has a well-defined limit: from equation (D.8) we have its closed-form expression and evaluating it at $\alpha_0 = 0$

and $\alpha_\infty = 1$ gives $\hat{\eta}_U = 1$ and $\hat{\eta}_U = A$, respectively. Second, define $\hat{\eta}_V \equiv 1 - \hat{\gamma}$ as the wealth-elasticity $a\hat{V}'/\hat{V}$ of the value function in the benchmark's one-good economy (which is constant, by the power form), the benchmark's counterpart of the wealth-elasticity $\eta_V(a) = aV'(a)/V(a)$ from lemma 1; the proof below shows that $\eta_V(a)$ converges to it. Finally, to define $\hat{\kappa}$, note that the relevant benchmark is then a one-good Merton economy with Duffie–Epstein preferences with risk aversion $\hat{\gamma}$, elasticity of intertemporal substitution $\hat{\psi}$, discount rate ρ , and return process characterized by the vector of expected returns and risky-asset volatility (r_B, r_S, σ_S) . Hence, optimal expenditure is proportional to wealth and we can define $\hat{\kappa}$ as the ratio of optimal expenditure to wealth: $\hat{\kappa} = \hat{x}(a)/a$. Additionally, note that the value function of the benchmark economy is the power function $\hat{V}(a) = (\hat{K}a)^{1-\hat{\gamma}}/(1-\hat{\gamma})$.¹⁶ Then, substituting the power form $\hat{V}$ into the HJB equation of the one-good Merton model and solving for $\hat{\kappa}$ we can get the expenditure-to-wealth ratio in closed form as:

$$\hat{\kappa} = r_B + \hat{\psi}(\rho - r_B) + (1 - \hat{\psi}) \frac{(r_S - r_B)^2}{2\hat{\gamma}\sigma_S^2} \quad (\text{D.16})$$

which, when evaluated at the two limits will yield the values κ_0 and κ_∞ below.

For reference, we collect below the endpoint values of the benchmark objects:

	α	$\hat{\gamma}$	$\hat{\psi}$	$\hat{\kappa}$	$\hat{\eta}_U$	$\hat{\eta}_V$
$a \rightarrow 0$:	0	$\gamma_0 = \gamma$	$\psi_0 = \psi$	κ_0	1	$1 - \gamma$
$a \rightarrow \infty$:	1	$\gamma_\infty = 1 + A(\gamma - 1)$	$\psi_\infty = \frac{1}{1+A(\psi^{-1}-1)}$	κ_∞	A	$A(1 - \gamma)$

Formally, the proof requires three regularity conditions:

- (R1) The model's primitives are such that the benchmark expenditure-to-wealth ratio $\hat{\kappa}$ is strictly positive at both limits: $\kappa_0 > 0$ and $\kappa_\infty > 0$.
- (R2) At each limit, the value function is asymptotically equivalent to that of the corresponding benchmark economy, that is, their ratio converges to one: $V(a)/\hat{V}(a) \rightarrow 1$ as $a \rightarrow 0$ and as $a \rightarrow \infty$.
- (R3) At each limit, indirect utility U is asymptotically a power function of expenditure with exponent $\hat{\eta}_U$, that is, $U(x)/(Bx^{\hat{\eta}_U}) \rightarrow 1$ for some constant $B > 0$.

Condition (R1) can be trivially checked using equation (D.16), while condition (R2) asks that, at each limit, the value function of the simplified two-good economy converges to that of the corresponding one-good Merton economy. Condition (R2) is relatively mild: since the intratemporal problem converges to a single good problem in the limit, the scale invariance of the recursive problem makes the value function a power function of wealth (Skiadas 2007).¹⁷

Condition (R3) is the intratemporal counterpart of (R2) and strengthens proposition 2: expenditure shares at the limit imply that the elasticity converges, $\eta_U(x) \rightarrow \hat{\eta}_U$, whereas (R3) requires

¹⁶The constant $\hat{K} > 0$ depends on the units in which the benchmark's flow utility is measured, but its actual value plays no role in the rest of the proof.

¹⁷I take (R2) as an assumption rather than proving it formally.

convergence of the level $U(x)/x^{\hat{\eta}_U}$. It follows directly from the functional form of intratemporal utility; in order to see this, recall from online appendix D.3 that

$$u = (C + N)^{1/\beta_h}, \quad C = (1 - \omega) \frac{\beta_h}{\beta_c} c^{\beta_c}, \quad N = \omega n^{\beta_h}, \quad \frac{C}{N} = \frac{\alpha}{A(1 - \alpha)}$$

We can then write indirect utility in two equivalent ways; first, factor out N

$$U(x) = \frac{\omega^{1/\beta_h}}{r_n} x \cdot (1 - \alpha) \left(1 + \frac{C}{N}\right)^{1/\beta_h}.$$

Hence, as $x \rightarrow 0$ we have $\alpha \rightarrow 0$ and $C/N \rightarrow 0$ so $U(x) \rightarrow Bx$ with $B = \omega^{1/\beta_h}/r_n$. Second, factor out C

$$U(x) = \left(\frac{1 - \omega}{A}\right)^{1/\beta_h} x^A \cdot \alpha^A \left(1 + \frac{N}{C}\right)^{1/\beta_h}.$$

Hence, as $x \rightarrow \infty$ we have $\alpha \rightarrow 1$ and $N/C \rightarrow 0$ so $U(x) \rightarrow Bx^A$ with $B = ((1 - \omega)/A)^{1/\beta_h}$. Condition (R3) therefore holds at both limits.

Note that, because (R2) compares the household's value function with that of the benchmark, the two economies must measure flow utility in the same units which requires a normalization. Throughout the proof, $\hat{V}$ denotes the benchmark's value function in the household's units: the constant $\hat{K}$ therefore absorbs the normalization factor.¹⁸

SOLVING THE HJB. First, I turn the household's HJB equation into an expression for $\mathcal{R}_V$. In the simplified economy the household HJB equation is:¹⁹

$$0 = \max_{x, \theta} \left\{ f(U(x), V) + V' (r_B a + (r_S - r_B) \theta a - x) + \frac{1}{2} \sigma_S^2 \theta^2 a^2 V'' \right\},$$

and the optimal portfolio share satisfies the first-order condition

$$\theta(a)a = -\frac{(r_S - r_B)V'(a)}{\sigma_S^2 V''(a)}.$$

Substituting this first-order condition into the household's HJB equation gives

$$0 = f(U(x^*), V) + V' (r_B a - x^*) - \frac{(r_S - r_B)^2 (V')^2}{2\sigma_S^2 V''}.$$

¹⁸Specifically, at each limit, the household's indirect utility behaves like $Bx^{\hat{\eta}_U}$ by (R3), whereas the benchmark's flow utility is x itself. Rewriting the benchmark's value function in the household's units multiplies its value function by the constant $B^{1-\gamma}\hat{\eta}_U$ but leaves its optimal policies unchanged. With $c \equiv B^{1-\gamma}\hat{\eta}_U$, the household's aggregator f and the benchmark's aggregator $\hat{f}$, which has parameters $(\hat{\gamma}, \hat{\psi})$, satisfy $f(Bx^{\hat{\eta}_U}, cv) = c\hat{f}(x, v)$, so the two HJB equations coincide up to the factor c .

¹⁹Note that the HJB is quadratic in θ , and the household is free to choose any value of θ . A maximum therefore exists only if the coefficient on θ^2 is negative, that is, only if $V'' < 0$.

Since $\mathcal{R}_V = -aV''/V'$, the last term is $\frac{(r_S - r_B)^2 aV'}{2\sigma_S^2 \mathcal{R}_V}$. Dividing by $aV'(a) > 0$, writing $\kappa(a) = x^*(a)/a$, and solving for $\mathcal{R}_V$, gives:

$$\mathcal{R}_V(a) = \frac{(r_S - r_B)^2}{2\sigma_S^2} \left[\kappa(a) - r_B - \frac{f(U(x^*(a)), V(a))}{a V'(a)} \right]^{-1}. \quad (\text{D.17})$$

DECOMPOSING THE FLOW-UTILITY TERM. To further decompose $f/(aV')$ notice that, as a function of u , the intertemporal aggregator f defined in equation (2) (with $\zeta = 0$) is the sum of a term proportional to $u^{1-1/\psi}$ and a term that does not involve u . Differentiating with respect to u we can therefore write

$$f(u, v) = \frac{f_u(u, v) u - \rho(1 - \gamma)v}{1 - \psi^{-1}}.$$

Using the expenditure first-order condition $f_u U' = V'$ from the proof of lemma 1 and $\eta_U = x^* U'/U$, this identity evaluated at $u = U(x^*(a))$ and $v = V(a)$ gives

$$\frac{f(U(x^*), V)}{a V'} = \frac{1}{1 - \psi^{-1}} \left[\frac{\kappa(a)}{\eta_U(x^*(a))} - \rho \frac{(1 - \gamma)V(a)}{a V'(a)} \right]. \quad (\text{D.18})$$

The factor $(1 - \gamma)V/(aV')$ is the reciprocal of $\eta_V(a)/(1 - \gamma)$, the wealth-elasticity ratio appearing in equation (11).

CONVERGENCE OF THE MARGINAL VALUE OF WEALTH. Condition (R2) gives convergence of the value function itself. However, equation (D.17) and equation (D.18) involve its derivative V' so it is also necessary to show that $V'(a)/\hat{V}'(a) \rightarrow 1$, which we can obtain by the concavity of V . Because V is concave, for $0 < h < 1$:

$$\frac{V(a(1 + h)) - V(a)}{ah} \leq V'(a) \leq \frac{V(a) - V(a(1 - h))}{ah}.$$

Dividing all terms by $\hat{V}'(a) > 0$, we would like to show that both the left and right side of the inequality converge, at each of the limits under consideration, to bounds that tend to 1 as $h \rightarrow 0$.

Start by multiplying and dividing each term by $\hat{V}$ which, for the left-most term, gives us:

$$\frac{V(a(1 + h)) - V(a)}{ah \hat{V}'(a)} = \frac{V(a(1 + h))}{\hat{V}(a(1 + h))} \cdot \frac{\hat{V}(a(1 + h))}{ah \hat{V}'(a)} - \frac{V(a)}{\hat{V}(a)} \cdot \frac{\hat{V}(a)}{ah \hat{V}'(a)},$$

where both $V/\hat{V}$ ratios converge to 1 by (R2). We are then left to evaluate the two remaining terms; specifically, note that by the power form of $\hat{V}$ we have $\hat{V}'(a) = \hat{K}^{1-\hat{\gamma}} a^{-\hat{\gamma}}$ so that:

$$\begin{aligned} \frac{\hat{V}(a(1 + h))}{ah \hat{V}'(a)} &= \frac{(1 + h)^{1-\hat{\gamma}}}{(1 - \hat{\gamma})h} \\ \frac{\hat{V}(a)}{ah \hat{V}'(a)} &= \frac{1}{(1 - \hat{\gamma})h}. \end{aligned}$$

Hence, holding h fixed and taking the limit as a tends to the limit under consideration (i.e., 0 or ∞), we have:

$$\frac{V(a(1+h)) - V(a)}{ah \hat{V}'(a)} \rightarrow \frac{(1+h)^{1-\hat{\gamma}} - 1}{(1-\hat{\gamma})h}$$

and similarly for the right-most term:

$$\frac{V(a) - V(a(1-h))}{ah \hat{V}'(a)} \rightarrow \frac{1 - (1-h)^{1-\hat{\gamma}}}{(1-\hat{\gamma})h}.$$

We therefore get:

$$\frac{(1+h)^{1-\hat{\gamma}} - 1}{(1-\hat{\gamma})h} \leq \liminf_a \frac{V'(a)}{\hat{V}'(a)} \leq \limsup_a \frac{V'(a)}{\hat{V}'(a)} \leq \frac{1 - (1-h)^{1-\hat{\gamma}}}{(1-\hat{\gamma})h},$$

where $\liminf$ and $\limsup$ are taken as a tends to the limit under consideration; both bounds tend to 1 as $h \rightarrow 0$, and therefore $V'(a)/\hat{V}'(a) \rightarrow 1$.

CONVERGENCE OF THE EXPENDITURE POLICY. The expenditure policy enters equations (D.17) and (D.18) through two terms: the expenditure-to-wealth ratio $\kappa(a)$ and the elasticity $\eta_U(x^*(a))$. In this and the next steps of the proof I establish that $\eta_U(x^*(a)) \rightarrow \hat{\eta}_U$ and $\kappa(a) \rightarrow \hat{\kappa}$.

We start by separating the expenditure first-order condition into a side that depends only on expenditure and a side that depends only on wealth; at each of the two limits for a , the wealth side behaves like a power of wealth with exponent $-1/\hat{\psi}$, which gives $\eta_U(x^*(a)) \rightarrow \hat{\eta}_U$. Then, in the following step, I use this result to show that $\kappa(a) \rightarrow \hat{\kappa}$.

Optimal expenditure $x^*(a)$ is determined by the expenditure first-order condition

$$f_u(U(x^*), V) U'(x^*) = V'(a).$$

To compute f_u , collect the u -dependence of the aggregator: with $\zeta = 0$, equation (2) reads

$$f(u, v) = \frac{\rho}{1 - \psi^{-1}} \left[u^{1-\psi^{-1}} ((1-\gamma)v)^\nu - (1-\gamma)v \right], \quad \nu \equiv \frac{\gamma - \psi^{-1}}{\gamma - 1},$$

so that differentiating with respect to u gives

$$f_u(u, v) = \rho u^{-1/\psi} ((1-\gamma)v)^\nu.$$

Substituting f_u into the first-order condition, evaluated at $u = U(x^*)$ and $v = V(a)$, and dividing both sides by $((1-\gamma)V(a))^\nu$ separates the first-order condition into a side that depends only on expenditure and a side that depends only on wealth:

$$\begin{aligned} \rho U(x^*)^{-1/\psi} U'(x^*) &= V'(a) ((1-\gamma)V(a))^{-\nu} \\ \varphi(x^*(a)) &= \Phi(a) \end{aligned}$$

To characterize the limits of the expenditure policy function, we observe two things: First, the expenditure side φ is strictly positive, finite at every $x > 0$, and strictly decreasing: its elasticity is

$$\frac{d \ln \varphi}{d \ln x} = -\frac{1}{\psi} \frac{x U'(x)}{U(x)} + \frac{x U''(x)}{U'(x)} = -\left(\frac{\eta_U(x)}{\psi} + \mathcal{R}_U(x) \right) = -\frac{1}{\text{EIS}(x)},$$

where $\text{EIS}(x)$ is strictly positive, as shown in the proof of proposition 3.

Second, at the two limits the wealth side converges to its benchmark counterpart $\hat{\Phi}(a)$ defined as

$$\hat{\Phi}(a) \equiv \hat{V}'(a) \left((1 - \gamma) \hat{V}(a) \right)^{-\nu},$$

which has the following two limits: $\lim_{a \rightarrow 0} \hat{\Phi}(a) = \infty$ and $\lim_{a \rightarrow \infty} \hat{\Phi}(a) = 0$. To see this, note that by condition (R2) and the convergence of the marginal value of wealth established above,

$$\frac{\Phi(a)}{\hat{\Phi}(a)} = \frac{V'}{\hat{V}'} \cdot \left(\frac{V}{\hat{V}} \right)^{-\nu} \rightarrow 1.$$

Substituting the power form of $\hat{V}$, as in the previous step, allows us to evaluate $\hat{\Phi}$ explicitly:

$$\hat{\Phi}(a) = \hat{K}^{(1-\hat{\gamma})(1-\nu)} \left(\frac{1-\gamma}{1-\hat{\gamma}} \right)^{-\nu} a^{-\hat{\gamma}-(1-\hat{\gamma})\nu},$$

where at the two limits of interest, the wealth exponent equals $-1/\hat{\psi}$. In fact, at the bottom we have $\hat{\gamma} = \gamma$ and

$$-\gamma - (1 - \gamma) \nu = -\gamma + \left(\gamma - \psi^{-1} \right) = -\frac{1}{\psi},$$

while at the top we have $1 - \hat{\gamma} = A(1 - \gamma)$,

$$-\hat{\gamma} - (1 - \hat{\gamma}) \nu = -1 + A(1 - \gamma) + A \left(\gamma - \psi^{-1} \right) = -(1 + A(\psi^{-1} - 1)) = -\frac{1}{\psi_{\infty}}.$$

Since, by proposition 3, $1/\psi > 0$, we have that $\hat{\Phi} \propto a^{-1/\hat{\psi}}$ goes to ∞ at the bottom limit and 0 at the top one, and so does Φ : $\Phi(a) \rightarrow \infty$ as $a \rightarrow 0$ and $\Phi(a) \rightarrow 0$ as $a \rightarrow \infty$. The limits of optimal expenditure then follow from monotonicity of the expenditure side, $\varphi(x)$: $x^*(a) \rightarrow 0$ at the bottom and $x^*(a) \rightarrow \infty$ at the top. Using equation (D.8) and proposition 2 this also implies $\eta_U(x^*(a)) \rightarrow \hat{\eta}_U$ at each of the two limits.

CONVERGENCE OF THE EXPENDITURE-TO-WEALTH RATIO. The limits we just derived for the optimal expenditure function do not yet determine the limit of the ratio $\kappa(a) = x^*(a)/a$, because both numerator and denominator have the same limit. To pin down the limit for $\kappa(a)$, we compare the household's optimal expenditure $x^*(a)$ with that of the one-good benchmark economy, $\hat{x}(a) = \hat{\kappa}a$: showing that $x^*(a)/\hat{x}(a) \rightarrow 1$ is the same as showing that $\kappa(a) \rightarrow \hat{\kappa}$. To do this comparison, we look at the FOC in both economies.

Let's start with the household's FOC which, once we substitute $\Phi(a)/\hat{\Phi}(a) \rightarrow 1$ reduces to

$$\frac{\varphi(x^*(a))}{\hat{\Phi}(a)} \rightarrow 1. \quad (\text{D.19})$$

On the other hand, in the one-good benchmark indirect utility takes the asymptotic power form $U(x) \sim B x^{\hat{\eta}_U}$, by (R3), and its derivative is $U'(x) \sim B \hat{\eta}_U x^{\hat{\eta}_U - 1}$. The expenditure side in the one-good benchmark is therefore defined as:

$$\begin{aligned} \hat{\varphi}(x) &\equiv \rho \left(B x^{\hat{\eta}_U} \right)^{-1/\psi} B \hat{\eta}_U x^{\hat{\eta}_U - 1} \\ &= \rho \hat{\eta}_U B^{1-\psi^{-1}} x^{\hat{\eta}_U(1-\psi^{-1})-1} \\ &= \rho \hat{\eta}_U B^{1-\psi^{-1}} x^{-1/\hat{\psi}} \end{aligned}$$

where for the last equality we have used the fact that $\hat{\eta}_U(1 - \psi^{-1}) - 1$ equals $-1/\psi$ at the bottom ($\hat{\eta}_U = 1$) and $-1/\psi_\infty$ at the top ($\hat{\eta}_U = A$). Separating the benchmark's expenditure first-order condition exactly as we did for the household problem gives $\hat{\varphi}(x) = \hat{\Phi}(a)$ which, since a change of units leaves policies unchanged, also holds at the benchmark's optimal expenditure:

$$\hat{\varphi}(\hat{x}(a)) = \hat{\Phi}(a). \quad (\text{D.20})$$

Finally, to connect the two economies, we need to show that φ and $\hat{\varphi}$ coincide asymptotically. First, condition (R3) states that indirect utility approaches the benchmark's flow utility:

$$\frac{U(x)}{B x^{\hat{\eta}_U}} \rightarrow 1.$$

Second, the same holds for the derivatives which can be seen by using its elasticity, $U'(x) = \eta_U(x) U(x)/x$:

$$\frac{U'(x)}{B \hat{\eta}_U x^{\hat{\eta}_U - 1}} = \frac{\eta_U(x)}{\hat{\eta}_U} \cdot \frac{U(x)}{B x^{\hat{\eta}_U}} \rightarrow 1,$$

where the first tends to one because $\eta_U(x) \rightarrow \hat{\eta}_U$ and the second by (R3). Third, φ and $\hat{\varphi}$ are built from these two objects in the same way, so their ratio also tends to one:

$$\frac{\varphi(x)}{\hat{\varphi}(x)} = \frac{\rho U(x)^{-1/\psi} U'(x)}{\rho (B x^{\hat{\eta}_U})^{-1/\psi} B \hat{\eta}_U x^{\hat{\eta}_U - 1}} = \left(\frac{U(x)}{B x^{\hat{\eta}_U}} \right)^{-1/\psi} \cdot \frac{U'(x)}{B \hat{\eta}_U x^{\hat{\eta}_U - 1}} \rightarrow 1. \quad (\text{D.21})$$

Putting the three pieces together,

$$\frac{\hat{\varphi}(x^*(a))}{\hat{\varphi}(\hat{x}(a))} = \frac{\hat{\varphi}(x^*(a))}{\varphi(x^*(a))} \cdot \frac{\varphi(x^*(a))}{\hat{\Phi}(a)} \cdot \frac{\hat{\Phi}(a)}{\hat{\varphi}(\hat{x}(a))} \rightarrow 1$$

by equations (D.19) to (D.21). Finally, by the definition of $\hat{\varphi}$, the left-hand side equals $(x^*(a)/\hat{x}(a))^{-1/\hat{\psi}}$, so $x^*(a)/\hat{x}(a) \rightarrow 1$ and therefore $\kappa(a) \rightarrow \hat{\kappa}$.

CONCLUSION. Now that we have a well-defined limit for every object in the flow-utility equation (D.18) and in the HJB equation (D.17) we can collect them and evaluate the resulting expressions. As a approaches each of the two limits, $\kappa \rightarrow \hat{\kappa}$ and $\eta_U(x^*) \rightarrow \hat{\eta}_U$ by the previous two steps. Moreover, by (R2) and the convergence of V' we have

$$\frac{\eta_V(a)}{\hat{\eta}_V} = \frac{V'/V}{\hat{V}'/\hat{V}} \rightarrow 1 \quad \implies \quad \frac{(1-\gamma)V}{aV'} = \frac{1-\gamma}{\eta_V(a)} \rightarrow \frac{1-\gamma}{\hat{\eta}_V}.$$

Hence, taking the limit of equation (D.18) gives

$$\begin{aligned} \frac{f(U(x^*), V)}{aV'} &\longrightarrow \frac{1}{1-1/\psi} \left[\frac{\hat{\kappa}}{\hat{\eta}_U} - \frac{\rho(1-\gamma)}{\hat{\eta}_V} \right] \\ &= \frac{1}{1-1/\psi} \left[\frac{\hat{\kappa} - \rho}{\hat{\eta}_U} \right] \\ &= \frac{1}{1-1/\hat{\psi}} (\hat{\kappa} - \rho) \end{aligned}$$

where the first equality follows from $\hat{\eta}_V = 1 - \hat{\gamma} = \hat{\eta}_U(1 - \gamma)$ and the second from $1 - 1/\hat{\psi} = \hat{\eta}_U(1 - 1/\psi)$. Hence the bracket of equation (D.17) converges to $\hat{\kappa} - r_B - (\hat{\kappa} - \rho)/(1 - \hat{\psi}^{-1})$.

By equation (D.16) we have that in the one-good benchmark economy risk aversion is determined by

$$\frac{(r_S - r_B)^2}{2\sigma_S^2 \hat{\gamma}} = \hat{\kappa} - r_B - \frac{\hat{\kappa} - \rho}{1 - 1/\hat{\psi}}.$$

The bracket of equation (D.17) converges to the right side of the above expression, which is strictly positive, so

$$\mathcal{R}_V(a) \longrightarrow \frac{(r_S - r_B)^2 / (2\sigma_S^2)}{(r_S - r_B)^2 / (2\sigma_S^2 \hat{\gamma})} = \hat{\gamma}.$$

Concluding, at each of the two limits true risk aversion converges to the benchmark's curvature and that curvature is, by definition, the corresponding limit of the Hanoch curvature established in proposition 4. Additionally, since $x^*(a) \rightarrow 0$ as $a \rightarrow 0$ and $x^*(a) \rightarrow \infty$ as $a \rightarrow \infty$, the same limit holds along the optimal expenditure path.

The limits of $\mathcal{R}_V$ and $\mathcal{R}_H$ therefore exist and coincide:

$$\lim_{a \rightarrow 0} \mathcal{R}_V(a) - \mathcal{R}_H(x^*(a)) = \lim_{a \rightarrow \infty} \mathcal{R}_V(a) - \mathcal{R}_H(x^*(a)) = 0$$

or, as stated in the corollary:

$$\lim_{a \rightarrow 0} \mathcal{R}_V(a) = \lim_{x \rightarrow 0} \mathcal{R}_H(x) = \gamma_0, \quad \lim_{a \rightarrow \infty} \mathcal{R}_V(a) = \lim_{x \rightarrow \infty} \mathcal{R}_H(x) = \gamma_\infty.$$

The values of $\gamma_0 = \gamma$ and $\gamma_\infty = 1 + A(\gamma - 1)$ come from proposition 4. □

E. Calibration Details

E.1. Earnings Process Calibration

DATA AND STRATEGY. I calibrate the process for household earnings described in equation (5) using two sources: moments of changes in male pre-tax labor earnings from Guvenen et al. (2021) (henceforth GKOS), reported in Figures 1 and 2 and Table I of that paper; and the cross-sectional dispersion of household earnings from the SCF. The calibration proceeds in two stages and follows the same general strategy as Kaplan et al. (2018):²⁰ in the first stage, I search for the parameters of the continuous-time process — the jump arrival rates λ_j , mean-reversion speeds β_j , and jump sizes σ_j for both the persistent and transitory components — that minimize the distance between simulated and data moments via Monte Carlo simulation; in the second stage, after applying a tax adjustment, I discretize each component on a grid and search for the width and curvature of the discrete grids so that the resulting Markov chain for after-tax log-earnings replicates the moments of the corresponding calibrated continuous process. I describe these two stages in more details below.

MOMENT TARGETS AND FIT. Since the earnings process is the same as in Kaplan et al. (2018) I follow the same approach and target a similar set of moments: the cross-sectional standard deviation of log-earnings and seven additional moments of the distribution of log-earnings changes — the standard deviation and kurtosis of one-year and five-year changes as well as the fraction of workers whose absolute one-year change falls below 5%, 10%, and 20% — which allow me to jointly discipline the tails, the persistence, and the high-frequency volatility of the earnings process.²¹ Since GKOS do not report the cross-sectional standard deviation of male earnings and their data follow individuals rather than households, I take this moment from household earnings in the SCF and compute the standard deviation of log labor earnings among households with positive earnings (using the same sample selection as for the other cross-sectional moments) under the labor-earnings definition of Kuhn and Ríos-Rull (2016).²² The loss function is:

$$L(\theta) = \sqrt{\mathbf{d}(\theta)' \mathbf{W} \mathbf{d}(\theta)}, \quad d_k(\theta) = \frac{m_k(\theta) - m_k^{\text{data}}}{m_k^{\text{data}}}.$$

The weighting matrix $\mathbf{W}$ places equal weight on every moment and is normalized so that its entries sum to one, which makes $L(\theta)$ the root mean square of the relative deviations $d_k(\theta)$ of the model moments from their targets. In both stages I minimize $L(\theta)$ using a parallel multistart algorithm which resembles the TikTak global optimization algorithm in Arnoud et al. (2022) and combines a quasi-random Sobol search over the parameter space with local refinements using DFO-LS (Cartis et al. 2019) as the local optimizer — a derivative-free algorithm tailored to least-squares objective

²⁰Notice that, relative to them, since my model abstracts from labor supply I do not need to apply a labor-supply grid correction. I do, however, adjust the process to be consistent with household-level earnings data.

²¹Skewness is zero by construction and is therefore not targeted.

²²Following Kuhn and Ríos-Rull, labor earnings are wages and salaries plus a fraction 0.878 of business and farm income, where this labor share is the economy-wide ratio of wage income to the sum of wage, interest, dividend, and capital-gains income; the implied variance of log labor earnings is 1.0537.

functions. More details on the algorithm are presented below.

Table E.1 reports the targeted moments and the model’s fit, while table E.2 reports the calibrated parameters for both components of the process. The persistent component, with an average inter-arrival time of roughly 19 years, accounts for the bulk of cross-sectional earnings dispersion; the transitory component, by contrast, is largely absorbed within a single year, contributing to high-frequency volatility. The large jump sizes are the main mechanism behind the high kurtosis in the data: because jumps are large but rare, the distribution of annual earnings changes has very fat tails, with most workers experiencing small changes and a small fraction experiencing very large ones, consistent with the pattern documented in Guvenen et al. (2021).

Table E.1: Earnings Process: Moment Fit

		MODEL		
MOMENT	DATA (GKOS)	Pre-tax	After-tax	Discretized
<i>Cross-section</i>				
Std. dev. of z (Kuhn and Ríos-Rull 2016)	1.027	1.016	0.832	0.899
<i>One-year changes</i>				
Std. dev.	0.510	0.490	0.401	0.368
Skewness	−1.07	—	—	—
Kurtosis	14.930	14.626	14.626	14.604
Frac. $ \Delta z < 5\%$	0.306	0.313	0.354	0.362
Frac. $ \Delta z < 10\%$	0.488	0.485	0.536	0.497
Frac. $ \Delta z < 20\%$	0.665	0.646	0.688	0.693
<i>Five-year changes</i>				
Std. dev.	0.780	0.817	0.669	0.696
Skewness	−1.25	—	—	—
Kurtosis	9.510	9.713	9.713	8.613

Table E.2: Earnings Process: Calibrated Parameters (Yearly)

PARAMETER	PERSISTENT	TRANSITORY
Arrival rate λ	0.0518	0.9350
Mean-reversion β	0.0152	11.0711
Jump size σ_j	1.246	1.495
Grid states N	7	5
Grid width	3.137	4.265
Grid curvature	0.392	0.961

IDENTIFICATION. Figure E.1 shows how the targeted moments identify the parameters of the earnings process. Panel E.1a reports the elasticity of each moment with respect to each parameter, evaluated at the calibrated values. Panel E.1b reports the implied sensitivity of each estimated parameter to each target, following Andrews et al. (2017). As expected, the parameters of the transitory component are mainly determined by moments of the distribution of one-year changes:

Figure E.1: Earnings Process: Identification

(a) Elasticities of Targeted Moments							(b) Sensitivity of Estimates to Targets						
$\sigma(z)$	0.17	-0.17	0.93	0.04	-0.01	0.13	$\sigma(z)$	0.12	-9.85	-0.80	-0.07	-0.22	-0.08
$\sigma(\Delta z)$	0.20	-0.03	0.38	0.29	-0.11	1.15	$\sigma(\Delta z)$	-0.42	5.10	0.90	-0.63	-0.26	0.60
$K(\Delta z)$	0.36	-0.07	1.08	-0.83	0.32	-1.81	$K(\Delta z)$	0.52	2.27	0.34	-0.09	-0.23	-0.29
$P(\Delta z < 5\%)$	-0.08	0.00	-0.02	-0.91	0.58	-0.50	$P(\Delta z < 5\%)$	-0.31	-1.76	-0.28	-2.78	-2.23	0.70
$P(\Delta z < 10\%)$	-0.10	0.00	-0.01	-0.64	0.66	-0.48	$P(\Delta z < 10\%)$	0.26	1.09	0.16	2.61	3.78	-0.42
$P(\Delta z < 20\%)$	-0.07	0.00	-0.02	-0.37	0.22	-0.45	$P(\Delta z < 20\%)$	-0.05	-1.09	-0.18	-0.57	-0.47	0.07
$\sigma(\Delta_5 z)$	0.40	-0.07	0.77	0.10	-0.04	0.42	$\sigma(\Delta_5 z)$	0.24	3.93	0.76	0.47	0.77	0.15
$K(\Delta_5 z)$	-0.49	-0.04	0.44	-0.25	0.11	-0.79	$K(\Delta_5 z)$	-1.58	3.74	1.00	0.03	-0.07	0.10
	λ_1	β_1	σ_1	λ_2	β_2	σ_2		λ_1	β_1	σ_1	λ_2	β_2	σ_2
	Parameter							Parameter					


Note: Rows are the targeted moments: $\sigma(x)$ and $K(x)$ indicate the cross-sectional standard deviation and kurtosis of log earnings ($x = z$), one-year changes ($x = \Delta z$), and five-year changes ($x = \Delta_5 z$), while $P(|\Delta z| < x\%)$ is the share of one-year changes smaller than $x\%$ in absolute value, for $x \in \{5, 10, 20\}$. Columns are the parameters of the persistent ($j = 1$) and transitory ($j = 2$) components. Panel (a) reports the elasticity $\partial \log m_k / \partial \log \theta_\ell$ of moment k with respect to parameter ℓ , evaluated at the calibrated parameters by central finite differences ($\pm 1\%$ perturbations of each parameter) on the same Monte Carlo panel ($N_i = 250,000$ agents) used in the final polishing search of Stage 1. Panel (b) reports $\partial \log \hat{\theta}_\ell / \partial \log m_k^{\text{data}}$, the elasticity form of the sensitivity measure of Andrews et al. (2017) for the calibration loss with equal weights on all moments, computed from panel (a) and the ratios of model to data moments.

the transitory arrival rate and mean reversion, λ_2 and β_2 , are pinned down primarily by the shares of small one-year changes; the transitory jump size σ_2 is pinned down by the same two shares and by the dispersion of one-year changes; finally, the persistent arrival rate λ_1 is pinned down mainly by the kurtosis of five-year changes. On the other hand, the parameters of the persistent component are primarily identified by the cross-sectional dispersion and moments of the distribution of five-year changes: the persistent jump size σ_1 depends about equally on the dispersion of one-year and five-year changes, the kurtosis of five-year changes, and the cross-sectional dispersion target; the mean reversion coefficient for the persistent component, β_1 , is only weakly identified. At the optimum, the value of β_1 is fairly small relative to the jump arrival rate λ_1 , so most of the movements in the persistent component are accounted for by jumps and all targeted moments therefore respond only weakly to β_1 . As a result, β_1 is pinned down mainly by the cross-sectional standard deviation of log earnings.²³ Figure E.2 reports the profile of the loss around the calibrated parameters: each panel moves one parameter away from its calibrated value in 5% steps and re-optimizes the remaining five parameters. The loss rises on both sides for every parameter which shows the algorithm has indeed found the optimum, but the profile is especially flat for β_1 which is consistent with it being weakly identified.

STAGE 1 – MONTE CARLO SIMULATION. Stage 1 simulates the two components using the Euler-Maruyama method at a monthly frequency for 56 years starting from the stationary distribution of the process (computed using the process generator); I then discard the first 50 years to remove dependence on initial conditions, and only use the last 6 years of the simulation to compute

²³The moments of earnings changes respond the least: the elasticities of all seven moments with respect to β_1 are at most 0.07 in absolute value, compared with 0.17 for the cross-sectional dispersion target.

Figure E.2: Earnings Process: Profile Loss



Note: Each panel fixes one parameter at its calibrated value and at $\pm 5\%$ and $\pm 10\%$ perturbations of it, and plots the loss $L(\theta)$ obtained by re-optimizing the other five parameters at each of these values. The first row shows the persistent component ($j = 1$) and the second row the transitory component ($j = 2$). The dashed line marks the loss at the calibrated parameters.

moments.²⁴ The two components are simulated independently and summed to obtain $z_t = z_{1,t} + z_{2,t}$.

AFTER-TAX ADJUSTMENT AND FREQUENCY CONVERSION. After simulation I then apply the Heathcote et al. (2017) tax function to convert pre-tax annual earnings to after-tax annual earnings before computing the Stage 2 targets. Specifically, for pre-tax earnings level y , after-tax earnings are given by:

$$y_{\text{net}} = \lambda_{\text{HSV}} \cdot y^{1-\tau_{\text{HSV}}},$$

where $\tau_{\text{HSV}} = 0.181$ is the progressivity parameter estimated by Heathcote et al. (2017) and λ_{HSV} is the level parameter, which I choose so that the implied aggregate average tax rate equals 28%, a standard target in the literature (Trabandt and Uhlig 2011).

STAGE 2 – GRID DISCRETIZATION AND MARKOV CHAIN SIMULATION. Each component is first discretized on its own grid, and then the grids and corresponding discretized generators are combined to a unique earnings grid with its associated discretized generator. The persistent and transitory components use 7 and 5 grid points, respectively, yielding a joint grid of 35 points. The total width of the grids, as well as their spacings are optimized to minimize the distance between the moments produced by a Markov-chain simulation of the combined process and the after-tax moments

²⁴I use a panel of $N_i = 100,000$ agents for both stages, and then run a final local polishing search from the global optimum with $N_i = 250,000$. Computed moments are reported from the larger panel and were found to be essentially unchanged when using a longer burn-in period.

implied by Stage 1, using the same loss function and optimizer as Stage 1.²⁵

For each component, the drift term is discretized using an upwind finite-difference scheme while the jump generator is constructed by evaluating the CDF of the jump distribution at the grid mid-points and assigning the resulting mass to the appropriate grid points. Because the two components are independent, the discretized generator of the joint process is going to be the Kronecker sum of the discretized generators of the individual components:

$$\mathbf{Q} = \mathbf{Q}_1 \oplus \mathbf{Q}_2 = \mathbf{Q}_1 \otimes \mathbf{I}_{N_2} + \mathbf{I}_{N_1} \otimes \mathbf{Q}_2,$$

where $\mathbf{I}_{N_j}$ is the $N_j \times N_j$ identity matrix and $\mathbf{Q}_j$ is the generator of component j . The joint grid is likewise formed as the Cartesian product of the two component grids.

E.2. Calibration Algorithm

The internally-calibrated parameters are chosen to minimize the distance $L(\theta)$ between model and data moments. To find the optimum I use a custom asynchronous parallel multistart derivative-free algorithm. The algorithm follows the same broad logic as TikTak (Arnoud et al. 2022): first, it begins with a Sobol pre-test over the bounded parameter space, retains the best resulting points as candidate seeds, and then launches local refinements from starts formed by mixing promising seeds with the current best guess. At the local stage these refinements are run using the DFO-LS algorithm, a derivative-free, bounded algorithm tailored to least-squares objectives rather than to a generic scalar objective (Cartis et al. 2019). This is particularly useful when calibrating the full model because the loss is expensive to evaluate and naturally organized around residual-based moment matching. The local DFO-LS step is applied to the entire vector of residuals $r(\theta)$ chosen so that $|r(\theta)|_2^2 = d(\theta)'Wd(\theta)$.

Like in TikTak, the algorithm first evaluates $L(\theta)$ on a Sobol sequence spanning the bounded parameter space. These evaluations provide a coarse ranking of the objective surface and are then ordered by their objective values to identify a set of candidate points to start the local refinement stage from. Rather than launching every refinement from the raw Sobol candidates, the algorithm forms each start by mixing a promising seed with the current best point, so later starts are gradually tilted toward regions that have already performed well.

Relative to TikTak, the algorithm implementation also changes how local chains are dispatched. New local starts are in fact generated adaptively during the run, using the current best point and sent out asynchronously as workers become available, so as to maximize parallelism and in order for improvements in the loss function to immediately affect newly launched local searches. At the same time, rather than shifting all work towards local optimization, the algorithm also preserves an exploration stage after the initial Sobol search.²⁶

²⁵The number of points for each grid was not formally included in the set of parameters to be optimized, but experimenting with different grids of different sizes and re-running the optimization step always led to a higher loss (conditional on keeping approximately the same total number of grid points).

²⁶Since with high values of RRA the model solution can sometimes be numerically unstable to large changes in γ , when calibrating the main model the mixing rule is also more conservative along that coordinate, which keeps risk

The full model calibration also uses an optional single-linkage filter to reduce repeated searches of the same basin of attraction. After the optimum has been found, the algorithm performs one final local polish from the best point found using tighter convergence parameters. The same algorithm is useful for both calibrations: for calibrating the earnings process it mainly provides a thorough global search and a reliable way to find a good local minimum, whereas for calibrating the main model in the paper the other features are more important, in particular the asynchronous parallelism, the conservative movement in γ , and the optional filter.

F. Numerical Appendix

F.1. Preferences

Relative to the standard heterogeneous-agent continuous time models of Achdou et al. (2022), preferences in this model feature two main differences that are relevant for its numerical solution: recursive intertemporal utility and a non-homothetic intratemporal aggregator. The former implies that the standard implicit method cannot be applied one-for-one, the latter that we potentially need to solve a non-linear problem at each step. In the following two sections, I describe how I tackle each of these in details.

F.1.1. Intertemporal Recursive Utility

Recall from equation (2) that the aggregator $f(u, v)$ takes the following explicit form, as in Duffie and Epstein (1992):

$$f(u, v) = \frac{\rho + \zeta}{1 - \psi^{-1}} \left[((1 - \gamma)v)^{\frac{\psi^{-1} - \gamma}{1 - \gamma}} u^{1 - \psi^{-1}} - (1 - \gamma)v \right]. \quad (\text{F.1})$$

With time-separable utility, equation (F.1) takes the standard form $f(u, v) = (\rho + \zeta) \left(\frac{u^{1 - \gamma}}{1 - \gamma} - v \right)$. In that case, since f is linear in v , for a given $\mathbb{A}^n$ the system of equations

$$0 = f(u^n, v^{n+1}) + \mathbb{A}^n v^{n+1} \quad (\text{F.2})$$

is *linear* in v^{n+1} and can be solved very quickly using the standard implicit method in Achdou et al. (2022).²⁷ With recursive preferences, however, f is nonlinear in v making solving equation (F.2) as is computationally prohibitive since it would require solving a very large (albeit sparse) nonlinear system at each HJB iteration.

To solve this, I instead split f into a linear and a non-linear component by writing:

$$f(u^n, v^n, v^{n+1}) = \underbrace{f(u^n, v^n)}_{\text{explicit}} + \lambda_{\text{rec}} v^n - \lambda_{\text{rec}} v^{n+1} \quad (\text{F.3})$$

aversion closer to its best-guess while still allowing other parameters to move more freely across the admissible set.

²⁷As they also acknowledge, this is formally a *semi*-implicit method since the fully implicit method would have $\mathbb{A}^{n+1}$ instead of $\mathbb{A}^n$ on the right-hand side and would imply solving a *non-linear* system of equations.

where

$$\lambda_{\text{rec}} \equiv (\rho + \zeta) \frac{1 - \gamma}{1 - \psi^{-1}}.$$

Equation (F.3) will obviously collapse to $f(u^n, v^{n+1})$ at the converged solution but has the advantage of being linear in v^{n+1} which means we only have to solve a linear system of equations while still applying a semi-implicit scheme.

Substituting into the discretized HJB in fact yields the linear system $\mathbb{B}^n v^{n+1} = d^n$:

$$\begin{aligned} \mathbb{B}^n &= \left(\frac{1}{\Delta} + \lambda_{\text{rec}} \right) \mathcal{I} - \mathbb{A}^n \\ d^n &= f(u^n, v^n) + \left(\frac{1}{\Delta} + \lambda_{\text{rec}} \right) v^n \end{aligned} \tag{F.4}$$

where Δ is the time step, $\mathcal{I}$ is the identity matrix, and $\mathbb{A}^n$ is the finite-difference discretization of $\mathcal{A}$. When $\psi = 1/\gamma$, $\lambda_{\text{rec}} = \rho + \zeta$ and the system reduces to the standard power-utility implicit scheme of Achdou et al. (2022).²⁸

Importantly, while the standard implicit scheme for power utility converges with a fixed large time step Δ , the explicit splitting of $f(u^n, v^n, v^{n+1})$ in the semi-implicit scheme limits the step size at which the iteration remains stable, so a constant large step is unfortunately insufficient for recursive preferences. Hence, I therefore apply an adaptive scheme that changes Δ based on the size of the residual.

F.1.2. Intratemporal Non-homothetic Aggregator

The intratemporal aggregator is in the addi-log form equation (3):

$$u(c, n) = \left((1 - \omega) \frac{1 - \varepsilon_h^{-1}}{1 - \varepsilon_c^{-1}} c^{\frac{\varepsilon_c - 1}{\varepsilon_c}} + \omega n^{\frac{\varepsilon_h - 1}{\varepsilon_h}} \right)^{\frac{\varepsilon_h}{\varepsilon_h - 1}}.$$

This implies that the household's first-order condition for consumption is given by:

$$f_u(u(c, n), v) \cdot u_c(c, n) = v_a \tag{F.5}$$

where $f_u \equiv \partial f / \partial u$ is the gradient of the Duffie-Epstein aggregator equation (F.1), and $v_a \equiv \partial v / \partial a$ is the marginal value of wealth. Equation (F.5) can also be rearranged as:

$$u(c, n)^{\varepsilon_h^{-1} - \psi^{-1}} c^{-\varepsilon_c^{-1}} = \frac{v_a [(1 - \gamma)v]^{\frac{\gamma - \psi^{-1}}{1 - \gamma}}}{(\rho_m + \zeta)(1 - \omega)} \tag{F.6}$$

²⁸When $\gamma > 1$, the continuation value satisfies $v < 0$, so $(1 - \gamma)v > 0$; however, particularly for larger values of γ , the base $(1 - \gamma)v$ approaches zero near the boundary of the state space, making direct evaluation of $[(1 - \gamma)v]^{\frac{\psi^{-1} - \gamma}{1 - \gamma}}$ numerically fragile. I therefore evaluate (the non-linear part of) f in log-space instead:

$$f(u, v) = \frac{\rho + \zeta}{1 - \psi^{-1}} \left[\exp \left(\frac{\psi^{-1} - \gamma}{1 - \gamma} \log((1 - \gamma)v) + (1 - \psi^{-1}) \log u \right) - (1 - \gamma)v \right].$$

which, except in the separable case $\psi = \varepsilon_h$, is a non-linear equation in c . For homeowners ($h > 0$), housing services are fixed at $n = \chi h$, where χ is the utility multiplier for owner-occupied housing units. Renters ($h = 0$) on the other hand get to choose how much to consume of housing services every period. For them the intratemporal Euler equation gives:

$$\frac{u_n}{u_c} = r_n \quad \implies \quad n(c) = \left(\frac{(1-\omega)}{\omega} r_n \right)^{-\varepsilon_h} c^{\varepsilon_h/\varepsilon_c} \quad (\text{F.7})$$

where r_n is the rental rate of housing. For the main calibration we have that $\psi \neq \varepsilon_h$, which means utility is non-separable and we must solve the non-linear equation at each grid point. To solve for c I use a variant of the secant method with bisection fallback, with an initial guess carried from the adjacent grid point. When provided with a good initial guess, the secant method generally converges very quickly to the solution; while bisection ensures robustness in the rare cases the secant step diverges.²⁹

F.2. Discretized HJB

This section collects the finite difference building blocks used throughout the numerical appendix.

I discretize the household's state (a, h, z_1, z_2, p, m) . Wealth a and housing h are discretized using non-uniform, geometrically spaced grids $\{a_i\}_{i=1}^{N_a}$ and $\{h_j\}_{j=1}^{N_h}$. The two earnings components z_1 and z_2 are combined into a single continuous-time Markov chain on a grid $\{z_k\}_{k=1}^{N_z}$, as described in online appendix F.2.3. The participation state p takes only two values in $\{0, 1\}$. The permanent discount type $m \in \{1, \dots, N_\rho\}$, by contrast, is fixed at birth and affects the problem only through the time-preference rate ρ_m , so the HJB can be solved as a separate block for each type. We define the forward and backward differences in the a -dimension as

$$\Delta_{i,+} \equiv a_{i+1} - a_i, \quad \Delta_{i,-} \equiv a_i - a_{i-1}.$$

All derivative formulas below are written for the a -dimension; the h -dimension is analogous.

The forward and backward first-derivative approximations at node i are

$$v_i^{a,F} \equiv \frac{\partial v}{\partial a} \Big|_{i,F} = \frac{v_{i+1} - v_i}{\Delta_{i,+}}, \quad v_i^{a,B} \equiv \frac{\partial v}{\partial a} \Big|_{i,B} = \frac{v_i - v_{i-1}}{\Delta_{i,-}}.$$

The second derivative also requires a three-point formula that accounts for non-uniform grids:

$$v_{aa,i} \equiv \frac{\partial^2 v}{\partial a^2} \Big|_i = \frac{\Delta_{i,-} v_{i+1} - (\Delta_{i,-} + \Delta_{i,+}) v_i + \Delta_{i,+} v_{i-1}}{d_i} \quad (\text{F.8})$$

where $d_i = \Delta_{i,-} \Delta_{i,+} (\Delta_{i,-} + \Delta_{i,+})/2$ and which reduces to the standard central difference $(v_{i+1} - 2v_i + v_{i-1})/(\Delta a)^2$ with a uniform grid.

The upwind scheme selects the finite difference direction according to the sign of the savings drift.

²⁹Also in this case, when γ is much larger than 1, direct evaluation is numerically fragile so the FOC is instead evaluated in log-space.

Write s for the drift of wealth in equation (6); because the policies entering s depend on the marginal value v_a , evaluating them with the forward and backward first-derivative approximations gives two candidate drifts s_F and s_B (both also depend on the remaining states h, z, p at the node, not on a_i alone). Define the upwind indicators $I_F = \mathbf{1}[s_F > 0]$, $I_B = \mathbf{1}[s_B < 0]$, and $I_0 = \mathbf{1}[s_F \leq 0 \leq s_B]$; the upwinded derivative is then

$$v_{a,i} = I_F v_i^{a,F} + I_B v_i^{a,B} + I_0 \bar{v}_i' \quad (\text{F.9})$$

where $\bar{v}_i'$ is the derivative consistent with zero savings at a_i . When v is concave, $v_i^{a,F} < v_i^{a,B}$ and hence $s_F < s_B$, so the three indicators are mutually exclusive. Where v is locally convex, however, the forward and backward conditions can hold simultaneously and the indicator rule in equation (F.9) no longer applies. In that case the scheme evaluates the discretized Hamiltonian under the forward, backward, and zero-drift candidates — each combining the flow payoff with the corresponding drift and diffusion terms — and selects the direction that gives the largest value.

F.2.1. Participation Costs

Entry costs into both equity and housing markets introduce additional complications to the finite-difference representation A. To simplify exposition I first look at a simplified version of the problem that has the same entry cost structure as the main model; in online appendix F.2.4 I then present the full matrix representation of the infinitesimal generator for the complete model of section 2. I start by looking at the entry problem into the equity market first and then, separately, analyze the entry problem into the housing market.

STOCK-MARKET PARTICIPATION. Consider the simplest 2-asset model without any labor income risk and time-separable preferences. Financial wealth a is therefore the only continuous state variable of the household problem. There are two participation states denoted by 0 and 1: households in state 0 only have access to bonds b , while households in state 1 can freely move assets between bonds and equities s . Bonds earn a risk-free rate of return r , while equities earn a risky return R with volatility σ . Households receive an opportunity to switch from state 0 to state 1 (entering the stock market) with intensity λ_0 and the opportunity to switch from state 1 to state 0 (exiting the stock market) with intensity λ_1 . Households who get hit by the participation shock can only enter the stock market by paying an entry cost κ_0^p ; households who get hit by the exit shock can only stay in the stock market by paying a staying cost κ_1^p .

The household problem therefore reads:

$$\begin{aligned} \rho v_0(a) &= \max_c u(c) + v_0'(a)(z + ra - c) + \lambda_0 \left[\max \{v_1(a - \kappa_0^p), v_0(a)\} - v_0(a) \right] \\ \rho v_1(a) &= \max_{c,s} u(c) + v_1'(a)(z + ra + (R - r)s - c) + v_1''(a) \frac{(s\sigma)^2}{2} + \\ &\quad + \lambda_1 \left[\max \{v_1(a - \kappa_1^p), v_0(a)\} - v_1(a) \right] \end{aligned}$$

The model nests both a model with a fixed cost of participation (with $\lambda_1 = 0$) and a model with a flow cost of participation (when $\lambda_1 \rightarrow \infty$). Making the opportunity to enter/exit the stock market

stochastic has the merit of making the numerical solution significantly simpler while also being able to arbitrarily approximate the continuous entry/exit decision: in fact, using larger and larger λ s allows us to get closer and closer to the continuous entry/exit problem without having to use the (more complicated) methods needed to solve variational inequalities.

The finite-differences approximation of the above HJB is then:

$$\begin{aligned} \frac{v_{i,p}^{n+1} - v_{i,p}^n}{\Delta} + \rho v_{i,p}^{n+1} = & u(c_{i,p}^n) + v_{i,p}^{a,F,n+1} \cdot \left(\dot{a}_{i,p}^{F,n}\right)^+ + v_{i,p}^{a,B,n+1} \cdot \left(\dot{a}_{i,p}^{B,n}\right)^- + \\ & + v_{i,p}^{aa,n+1} \cdot \frac{\left(s_{i,p}^n \sigma\right)^2}{2} + \lambda_p \left[\max \left\{ v_{i',p}^{n+1}, v_{i,0}^{n+1} \right\} - v_{i,p}^{n+1} \right] \end{aligned}$$

where $v_{i',p} = v_1(a_i - \kappa_p^s)$.

When specifying the above problem numerically, there are three things that one needs to be especially careful with: first, the a -state changes after the shock hits; second, $a_i - \kappa_p^s$ does not necessarily fall on a grid point (this is especially true in the presence of non-uniform grids); third, for any two different $\hat{v}, \tilde{v}$, to compute $\max \{ \hat{v}^{n+1}, \tilde{v}^{n+1} \}$ we would need to know their values at iteration $n + 1$.

The first point simply implies that there will be mass on some off-diagonal elements of $\mathbb{A}$, which means that, in the presence of non-uniform grids, the $\mathbb{A}$ matrix will not have a band-matrix structure. This is not a problem if one uses general sparse solvers, but it does exclude the possibility of exploiting the band structure to simplify the construction of $\mathbb{A}$ and/or the solution of the linear system.

To address the second point I assign mass to both the two closest points to $a_i - \kappa_p^s$ in proportion to their relative distance: for each grid point a_i , denote by $a_{i',p}$ and $a_{i'+1,p}$ the two grid points closest to $a_i - \kappa_p^s$. That is,

$$a_{i',p} \equiv \sup \left\{ a \in \{a_1, a_2, \dots, a_{N_a}\} : a < a_i - \kappa_p^s \right\}.$$

We can then compute the mass assigned to grid-point i' as $\omega_{i',p}^s = \frac{a_{i'+1,p} - (a_i - \kappa_p^s)}{a_{i'+1,p} - a_{i',p}}$ and the mass assigned to grid-point $i' + 1$ as $1 - \omega_{i',p}^s$. Notice also that, although in most cases the household would choose to pay (or not) the cost for both $a_{i',p}$ and $a_{i'+1,p}$, there might be edge cases when the household would only choose to pay the cost for $a_{i'+1,p}$ but not for $a_{i',p}$. Because the rows of $\mathbb{A}$ need to always sum to 0, I interpolate the value function across the two points and let the household decide based on the average value. The household is then forced to accept the “bet” and move to both points.

Finally, relative to the third point, I use a semi-implicit scheme consistent with how the other choice variables are determined, that is the choice of whether to enter or exit the stock market is based on the value function at iteration n , but the change in value is based on the value function at iteration $n + 1$. Hence, if we define by $\mathcal{I}_{i,p}^n$ the household decision to participate or not, that is:

$$\mathcal{I}_{i,p}^n \equiv \mathbf{1} \left\{ v_0^n(a_i) < \omega_{i',p}^p v_1^n(a_{i',p}) + (1 - \omega_{i',p}^p) v_1^n(a_{i'+1,p}) \right\}$$

we can rewrite the finite difference approximation as

$$\begin{aligned} \frac{v_{i,p}^{n+1} - v_{i,p}^n}{\Delta} + \rho v_{i,p}^{n+1} &= u(c_{i,p}^n) + v_{i,p}^{a,F,n+1} \cdot \left(\dot{a}_{i,p}^{c,F,n}\right)^+ + v_{i,p}^{a,B,n+1} \cdot \left(\dot{a}_{i,p}^{c,B,n}\right)^- \\ &\quad + v_{i,p}^{aa,n+1} \cdot \frac{\left(s_{i,p}^n \sigma\right)^2}{2} - \lambda_p (v_{i,p}^{n+1} - v_{i,0}^{n+1}) \\ &\quad + \lambda_p \left[\left(\omega_{i',p}^p v_1^{n+1}(a_{i',p}) + (1 - \omega_{i',p}^p) v_1^{n+1}(a_{i'+1,p}) \right) - v_{i,0}^{n+1} \right] \cdot \mathcal{I}_{i,p}^n \end{aligned}$$

HOUSE SWITCHING. The treatment above generalizes directly to the housing adjustment decision, which has the same Poisson structure but introduces two additional complications. First, a renter who receives a buying opportunity chooses the optimal housing size h' from the housing grid, so the post-purchase wealth $a' = a - h' - \kappa^h(h')$ generally falls between points on the a -grid; second, the housing grid has an asymmetry between $h = 0$ (renter, faces a buy opportunity at rate λ^h) and $h > 0$ (owner, faces a sell opportunity at rate λ^h). These two cases correspond to the $\mathcal{H}^0$ and $\mathcal{H}^+$ operators defined in appendix A of the main text and the same interpolation and semi-implicit logic as for stock-market participation applies, with the added step that the buyer's search must first resolve the discrete choice over h' before the a -interpolation weight can be computed.

A renter at grid point $(a_i, 0, z_k, p_l)$ receives a buying opportunity at rate λ^h . She searches over all affordable candidate housing levels $h_{j'} \in \{h_1, \dots, h_{N_h}\}$; for each candidate, wealth after buying a house is

$$a'(h_{j'}) = a_i - h_{j'} - \kappa^h(h_{j'}).$$

The optimal housing choice is $h^* = \arg \max_{h_{j'}} \tilde{v}^n(a'(h_{j'}), h_{j'}, z_k, p_l)$, where $\tilde{v}^n$ denotes the value function interpolated in a at iteration n . The renter buys only if the interpolated value at the optimal choice exceeds her current value $v^n(a_i, 0, z_k, p_l)$, in which case she transitions from $h = 0$ to $h = h^*$ and from a_i to $a'(h^*)$; otherwise she does not buy and remains a renter. The semi-implicit decision indicator is

$$\mathcal{I}_{i,B}^{h,n} \equiv \mathbf{1} \{ \tilde{v}^n(a'(h^*), h^*, z_k, p_l) > v^n(a_i, 0, z_k, p_l) \},$$

which equals 1 when buying is optimal at iteration n and 0 otherwise; the value update itself is carried out at iteration $n + 1$.

For each candidate $h_{j'}$, the post-purchase wealth $a'(h_{j'})$ generically falls between two adjacent grid points $a_{i'}$ and $a_{i'+1}$. We define the interpolation weight

$$\omega_{i',j'}^h = \frac{a_{i'+1} - a'(h_{j'})}{a_{i'+1} - a_{i'}}$$

so that the interpolated value for candidate $h_{j'}$ is

$$\omega_{i',j'}^h v^n(a_{i'}, h_{j'}, z_k, p_l) + (1 - \omega_{i',j'}^h) v^n(a_{i'+1}, h_{j'}, z_k, p_l).$$

Since the search is over grid points in h' only, one a -interpolation is needed per candidate; the

switch matrix Λ^h , however, spans both the a and h dimensions of $\mathbb{A}$ because the buyer transitions simultaneously in both state variables.

An owner at (a_i, h_j, z_k, p_l) with $h_j > 0$ receives a selling opportunity at rate λ^h . Post-sale wealth is

$$a' = a_i + h_j - \kappa^h(h_j),$$

and the owner's housing stock deterministically becomes $h = 0$. The owner sells — defined by the indicator $\mathcal{I}_{i,j,S}^{h,n} = 1$ — if $\tilde{v}^n(a', 0, z_k, p_l) > v^n(a_i, h_j, z_k, p_l)$; otherwise she remains an owner.

Between Poisson adjustment shocks, owners' housing stock evolves continuously via the drift μ_h . This drift is discretized with the same upwind finite differences applied to all other drift terms in $\mathbb{A}$. Renters, who have $h = 0$, are not subject to this drift and the corresponding rows of $\mathbb{A}$ carry no h -direction terms.

The Poisson jump terms for both equity participation and housing adjustment enter $\mathbb{A}$ as sparse off-diagonal blocks through switch matrices Λ^p and Λ^h . For equity, Λ^p connects the $p = 0$ and $p = 1$ blocks of $\mathbb{A}$. For a household at (a_i, h_j, z_k, p) that decides to switch participation state, Λ^p assigns rate $\lambda_p^p \cdot \mathcal{I}_{i,p}^n$ to the rows corresponding to $(a_{i'}, h_j, z_k, 1 - p)$ and $(a_{i'+1}, h_j, z_k, 1 - p)$, weighted by $\omega_{i',p}^p$ and $1 - \omega_{i',p}^p$ respectively. The h and z coordinates are unchanged by the equity participation transition, so Λ^p is block-sparse with a regular structure along the h and z dimensions.

For housing, Λ^h is asymmetric because buyers and sellers start from different h -states and arrive at different h -states. The buying rows assign rate $\lambda^h \cdot \mathcal{I}_{i,B}^{h,n}$ to grid points $(a_{i'}, h_{j^*}, z_k, p_l)$ and $(a_{i'+1}, h_{j^*}, z_k, p_l)$, weighted by ω_{i',j^*}^h and $1 - \omega_{i',j^*}^h$, where h_{j^*} is the optimal housing choice. The selling rows assign rate $\lambda^h \cdot \mathcal{I}_{i,j,S}^{h,n}$ to grid points $(a_{i'}, 0, z_k, p_l)$ and $(a_{i'+1}, 0, z_k, p_l)$, weighted by the analogous interpolation weight for post-sale wealth.

F.2.2. Portfolio Choice

The first-order condition for the equity share θ follows from differentiating the HJB with respect to θ and setting it to zero. For a participant ($p = 1$) the unconstrained optimum is the generalized Merton fraction

$$\theta^* = -\frac{v_a}{v_{aa} \cdot a} \cdot \frac{r_S - r_B}{\sigma_S^2 + \sigma_B^2 - 2\sigma_{BS}} - \frac{v_{ah}}{v_{aa} \cdot a} \cdot \frac{h(\sigma_{SH} - \sigma_{BH})}{\sigma_S^2 + \sigma_B^2 - 2\sigma_{BS}} + \frac{\sigma_B^2 - \sigma_{BS}}{\sigma_S^2 + \sigma_B^2 - 2\sigma_{BS}}$$

where the variance and covariance terms are defined as in equation (4). The first term is the standard Merton fraction; the second term is the hedging demand arising from the covariance of portfolio returns with housing wealth; the third term is the equity share of the global minimum-variance portfolio, which does not depend on risk preferences and vanishes when the bond is riskless ($\sigma_B^2 = \sigma_{BS} = 0$).

When the agent is a net borrower in the bond market that is, when $(1 - \theta)a < 0$ borrowed funds carry an additional wedge κ^b , shifting the optimal share to

$$\theta^{\text{cost}} = \theta^* + \frac{v_a}{v_{aa} \cdot a} \cdot \frac{\kappa^b}{\sigma_S^2 + \sigma_B^2 - 2\sigma_{BS}}.$$

I therefore apply two regimes: if the agent is a net borrower, I use θ^{cost} ; otherwise, θ^* . Non-participants ($p = 0$) are forced to $\theta = 0$.

The presence of risky assets generates a second-order term in the HJB operator through the portfolio diffusion coefficient

$$\sigma_{aa} = a^2 \left[(1 - \theta)^2 \sigma_B^2 + \theta^2 \sigma_S^2 + 2\theta(1 - \theta) \sigma_{BS} \right].$$

The savings drift and this diffusion together set the weights α_i and ξ_i that the finite difference approximation of the infinitesimal generator $\mathcal{A}$ places on the backward and forward wealth neighbors of node a_i (at fixed h_j, z_k, p, ρ_m):

$$\alpha_i = -I_B \frac{s_i^B}{\Delta_{i,-}} + \frac{1}{2} \sigma_{aa} w_i^-, \quad \xi_i = I_F \frac{s_i^F}{\Delta_{i,+}} + \frac{1}{2} \sigma_{aa} w_i^+$$

where s^F, s^B and the upwind indicators I_F, I_B are defined in equation (F.9), and $w_i^+ = \Delta_{i,-}/d_i$ and $w_i^- = \Delta_{i,+}/d_i$ are the weights that equation (F.8) places on the forward and backward neighbors.

At the upper boundary a_{max} , the value function approaches the power-law form with exponent $1 - \gamma_\infty$ which implies

$$-\frac{v_{aa}a}{v_a} = \gamma_\infty$$

and, when substituted into the FOC yields the asymptotic Merton fraction so we set v_{aa} analytically rather than via finite differences.

F.2.3. Earnings Process

The earnings process is a finite-state continuous-time Markov chain calibrated as described in online appendix E.1. Its generator is a $N_z \times N_z$ matrix Q . The chain has $N_z = 35$ states, formed as the Cartesian product of a persistent component on 7 grid points and a transitory component on 5 grid points; because the two components evolve independently, Q is the Kronecker sum $Q = Q_1 \oplus Q_2 = Q_1 \otimes \mathcal{I}_5 + \mathcal{I}_7 \otimes Q_2$ of the two component generators (online appendix E.1). Because z is now discrete and orthogonal to all other household states, all z -dynamics enter the infinitesimal generator $\mathbb{A}$ through:

$$\Lambda^z = Q \otimes \mathcal{I}_{N_a \cdot N_h}$$

where $\mathcal{I}_{N_a \cdot N_h}$ is the identity block on wealth $\times$ housing, and Λ^z carries all earnings dynamics and is added as a separate sparse block to $\mathbb{A}$.

F.2.4. Complete System

I discretize the state space on grids $\{a_i\}_{i=1}^{N_a}, \{h_j\}_{j=1}^{N_h}, \{z_k\}_{k=1}^{N_z}$, with participation states $p \in \{0, 1\}$ and discount-rate types indexed by $m \in \{1, \dots, N_\rho\}$. The sparse intensity matrix $\mathbb{A}$ is the finite-difference approximation of the continuous operator $\mathcal{A}$ defined in appendix A of the main text; I now describe its construction from the bottom up.

For a given vector of indices (j, k, p, m) , the a -drift and portfolio diffusion at these coordinates

generate an $N_a \times N_a$ tridiagonal matrix

$$\mathbb{A}_{j,k,p,m} = \begin{bmatrix} \beta_1 & \xi_1 & & & \\ \alpha_2 & \beta_2 & \xi_2 & & \\ & \ddots & \ddots & \ddots & \\ & & \alpha_{N_a-1} & \beta_{N_a-1} & \xi_{N_a-1} \\ & & & \alpha_{N_a} & \beta_{N_a} \end{bmatrix} \quad (\text{F.10})$$

where α_i and ξ_i are the backward and forward off-diagonal weights defined in equations (F.8) and (F.9), and $\beta_i = -\alpha_i - \xi_i$. Additionally, β_1 and β_{N_a} also fold in the reflecting boundary conditions at a_1 and a_{N_a} , respectively.

We now stack the a -tridiagonal blocks across the housing dimension to form the $N_a N_h \times N_a N_h$ matrix for the (a, h) state at fixed indices (k, p, m) . Four distinct contributions enter this block: the tridiagonal block from the savings drift and portfolio diffusion in equation (F.10); the nearest-neighbor terms from the housing-return drift and diffusion (which apply only to owners with $h_j > 0$); the mixed-derivative term generated by the covariance between portfolio and housing returns; and the Poisson jump terms from the housing switch matrix $\Lambda^{h,\text{jump}}$ described in online appendix F.2.1. The housing drift and diffusion generate off-diagonal entries that connect adjacent housing nodes h_j and $h_{j\pm 1}$ where the drift is upwinded using the usual scheme. The Poisson jump terms in $\Lambda^{h,\text{jump}}$ — which correspond to the $\mathcal{H}^0$ and $\mathcal{H}^+$ operators in appendix A of the main text — connect non-adjacent nodes in both the a and h dimensions simultaneously, because buying and selling a house changes both the housing stock and liquid wealth. The resulting block matrix is:

$$\mathbb{A}_{k,p,m} = \underbrace{\begin{bmatrix} \mathbb{A}_{1,k,p,m} & & \\ & \ddots & \\ & & \mathbb{A}_{N_h,k,p,m} \end{bmatrix}}_{a\text{-tridiagonal}} + \underbrace{\Lambda_{k,p,m}^h}_{\text{housing drift and diffusion}} + \underbrace{\Lambda_{k,p,m}^{ah}}_{\text{cross term}} + \underbrace{\Lambda_{k,p,m}^{h,\text{jump}}}_{\text{housing jumps}} \quad (\text{F.11})$$

The housing operator $\Lambda_{k,p,m}^h$ is the h -direction counterpart of equation (F.10): it collects the housing drift $r_H h$ and the housing diffusion $\frac{1}{2} h^2 (\sigma_3 \sigma_3')$ of $\mathcal{L}_c$ in appendix A of the main text, discretized along the housing grid with the same upwind and second-derivative formulas. Because both terms move (a_i, h_j) to $(a_i, h_{j\pm 1})$ at the same a_i , the operator factors as a tridiagonal matrix over the N_h housing nodes times the $N_a \times N_a$ identity $\mathcal{I}_{N_a}$:

$$\Lambda_{k,p,m}^h = \begin{bmatrix} \beta_1^h & \xi_1^h & & & \\ \alpha_2^h & \beta_2^h & \xi_2^h & & \\ & \ddots & \ddots & \ddots & \\ & & \alpha_{N_h-1}^h & \beta_{N_h-1}^h & \xi_{N_h-1}^h \\ & & & \alpha_{N_h}^h & \beta_{N_h}^h \end{bmatrix} \otimes \mathcal{I}_{N_a}$$

where α_j^h and ξ_j^h are the backward and forward weights of equations (F.8) and (F.9) applied along the housing grid, with the housing drift and diffusion in place of the savings drift and of σ_{aa} in online appendix F.2.2, and $\beta_j^h = -\alpha_j^h - \xi_j^h$. The renter node $h_1 = 0$ carries neither drift nor diffusion ($\beta_1^h = \xi_1^h = 0$). The cross operator $\Lambda_{k,p,m}^{ah}$ discretizes the remaining term of $\mathcal{L}_c$, the mixed derivative $ah(\sigma_a \sigma'_3) \frac{\partial^2}{\partial a \partial h} v$. The Poisson jump operator $\Lambda_{k,p,m}^{h,\text{jump}}$ is sparse but generally non-banded, with interpolation along the liquid wealth dimension a as described in online appendix F.2.1.

At the next level we stack $\mathbb{A}_{k,p,m}$ across the N_z earnings states and add the earnings block from online appendix F.2.3:

$$\mathbb{A}_{p,m} = \begin{bmatrix} \mathbb{A}_{1,p,m} & & \\ & \ddots & \\ & & \mathbb{A}_{N_z,p,m} \end{bmatrix} + \Lambda^z.$$

The per-type intensity matrix $\mathbb{A}_m$ stacks the two participation states and adds the equity-participation switch matrix from online appendix F.2.1:

$$\mathbb{A}_m = \begin{bmatrix} \mathbb{A}_{0,m} + \Lambda_{m,0 \rightarrow 0}^p & \Lambda_{m,0 \rightarrow 1}^p \\ \Lambda_{m,1 \rightarrow 0}^p & \mathbb{A}_{1,m} + \Lambda_{m,1 \rightarrow 1}^p \end{bmatrix} \quad (\text{F.12})$$

where $\Lambda_{m,p \rightarrow p'}^p$ encodes the Poisson rate λ_p^p , the semi-implicit decision indicator $\mathcal{I}_{i,p}^n$, and the interpolation weights from the participation appendix.

The model admits N_ρ discount-rate types, each with its own discount rate ρ_m . Because there are no cross-type transitions during the lifetime of an agent, the full intensity matrix $\mathbb{A}$ is block-diagonal across types and cross-type mixing happens only through the birth operator Ψ in the KFE:

$$\mathbb{A} = \begin{bmatrix} \mathbb{A}_1 & & \\ & \ddots & \\ & & \mathbb{A}_{N_\rho} \end{bmatrix}.$$

With recursive preferences, the nonlinear aggregator $f(u, v)$ from equation (F.1) is split into explicit and implicit components as described in online appendix F.1.1, yielding the linear system $\mathbf{B}_m v_m^{n+1} = d_m^n$ with

$$\mathbf{B}_m = \left(\frac{1}{\Delta} + \lambda_{\text{rec}} \right) \mathcal{I} - \mathbb{A}_m, \quad d_m^n = f(u_m^n, v_m^n) + \left(\frac{1}{\Delta} + \lambda_{\text{rec}} \right) v_m^n$$

as in equation (F.4), where $\lambda_{\text{rec}} = (\rho_m + \zeta)(1 - \gamma)/(1 - \psi^{-1})$. Across iterations the time step Δ is adapted using a switched-evolution-relaxation (SER) rule, which rescales the step by the ratio of successive residual norms so that it is small far from the solution and grows as the residual falls.

The solution of the discretized Kolmogorov forward equation is described in online appendix F.3.

F.3. Discretized KFE

This appendix covers how to approximate the distribution with non-uniform grids and the simulation of the KFE of the distribution under aggregate shocks.

F.3.1. KFE Under Non-uniform Grids

This is a generalization of online appendix 7.2 in Achdou et al. (2022) to the case of non-uniform grids along multiple state-variables. I begin by deriving the trapezoidal weights used to approximate the integral.

To simplify notation we here work with a simplified problem with only two states a and z , and no aggregate risk. The results immediately generalize to more states; aggregate risk obviously changes the equation that needs to be solved, but has no impact on the numerical approximation of the cross-sectional distribution in the presence of non-uniform grids.

Following Achdou et al. (2022) to solve the KFE with non-uniform grids we can work with the rescaled density $\tilde{g} = Dg$ where D is a diagonal matrix with elements $\tilde{\Delta}_\iota$ that appropriately adjusts for the presence of non-uniform grids. In particular, if $\mathbb{A}$ is the intensity matrix associated with the HJB problem above we will be solving

$$\frac{\tilde{g}^{n+1} - \tilde{g}^n}{\Delta t} = \mathbb{A}^T \tilde{g}^{n+1}$$

To understand why the above works notice that $\sum_{\iota=1}^N g_\iota \tilde{\Delta}_\iota$ is an approximation of the integral of g using a trapezoidal rule, and we denote by ι the "composite" index of a and z , i.e. $\iota = (i, k)$.

$$\begin{aligned} \int g(a, z) da dz &\approx \sum_{i=1}^{N_a-1} \sum_{k=1}^{N_z-1} \frac{(a_{i+1} - a_i)}{2} \frac{(z_{k+1} - z_k)}{2} (g_{i,k} + g_{i,k+1} + g_{i+1,k} + g_{i+1,k+1}) \\ &= \sum_{i=1}^{N_a} \sum_{k=1}^{N_z} g_{i,k} \tilde{\Delta} a_i \tilde{\Delta} z_k \\ &= \sum_{\iota=1}^N g_\iota \cdot \tilde{\Delta}_\iota \end{aligned}$$

where $N = N_a \cdot N_z$, $\tilde{\Delta}_\iota = \tilde{\Delta} a_i \cdot \tilde{\Delta} z_k$, and the $\tilde{\Delta}$ operator for a generic variable x indexed by j with $j \in \{1, 2, \dots, J\}$ is

$$\tilde{\Delta} x_j = \begin{cases} \frac{\Delta^F}{2} x_j & \text{if } j = 1 \\ \frac{(\Delta^B + \Delta^F)}{2} x_j & \text{if } 1 < j < J \\ \frac{\Delta^B}{2} x_j & \text{if } j = J \end{cases}$$

where $\Delta^F x_j = \Delta^B x_{j+1} = x_{j+1} - x_j$.

After having solved for $\tilde{g}$ we can then back out the true distribution using $g = D^{-1} \tilde{g}$ and, to compute moments, we can use the fact that $\sum_{\iota=1}^N g_\iota \tilde{\Delta}_\iota$ approximates the integral of g . For instance,

to find average asset holdings we would use:

$$\begin{aligned}\mathbb{E}(a) &= \int ag(a, z)dadz \approx \sum_{t=1}^N a_t g_t \tilde{\Delta}_t \\ &= \sum_{i=1}^{N_a} \sum_{k=1}^{N_z} a_i g_{i,k} \tilde{\Delta}_i \tilde{\Delta}_k\end{aligned}$$

F.3.2. KFE Under Aggregate Shocks

Recall the KF equation (10) can be written as

$$dg_t(\mathbf{x}) = \left\{ \mathcal{A}^* g_t(\mathbf{x}) + \zeta [(\Psi g_t)(\mathbf{x}) - g_t(\mathbf{x})] \right\} dt + \mathcal{B}^* g_t(\mathbf{x}) d\mathbf{W}_t$$

where Ψ is the operator associated to the birth distribution, which I describe more in detail below.

The goal is to derive a stable numerical scheme to solve for g_t in the presence of aggregate shocks. In order to do so, it is useful to define the infinitesimal operators for the drift and diffusion terms of the KFE. Just as in the case without aggregate risk, when it comes to solving the boundary-value-problem associated with the KFE, we also have a choice of using an explicit or implicit method. It turns out that, in the presence of a common noise term, the best choice is to use a so-called α -scheme on the drift term only (i.e. a semi-implicit scheme on the drift term and an explicit scheme on the diffusion term). In fact, as pointed out by Kloeden and Platen: “Difficulties can arise in applying fully implicit schemes to obtain strong approximations of solutions of stochastic differential equations, because they usually involve reciprocals of Gaussian random variables which do not have finite absolute moments. Consequently, finite absolute moments generally will not exist for fully implicit strong approximations and a strong convergence analysis would not make sense. For mainly this reason we shall restrict our attention here to ‘semi-implicit’ strong approximations, which we shall call implicit.” (Kloeden and Platen 1995, p. 395)

The drift-explicit and drift-implicit Euler schemes for the KFE have the form:

$$\begin{aligned}g_{t+\Delta}^E &= g_t^E + a(g_t^E, t)\Delta + b(g_t^E, t)\Delta \mathbf{W}_t \\ g_{t+\Delta}^I &= g_t^I + a(g_{t+\Delta}^I, t + \Delta)\Delta + b(g_t^I, t)\Delta \mathbf{W}_t\end{aligned}$$

where Δ is the time step, $\Delta \mathbf{W}_t \equiv \mathbf{W}_{t+\Delta} - \mathbf{W}_t$, $a(g, t) = \{\mathcal{A}_t^* + \zeta(\Psi - \mathcal{I})\}g \equiv \hat{\mathcal{A}}_t^* g$, and $b(g, t) = \mathcal{B}_t^* g$ (and, for the sake of completeness, we allow $\mathcal{A}_t^*$ and $\mathcal{B}_t^*$ to also depend on time). Following Kloeden and Platen (1995), from these two schemes we can construct a whole family of drift-implicit Euler schemes — called α -schemes — which are given by:

$$g_{t+\Delta} = g_t + [\alpha a(g_{t+\Delta}, t + \Delta) + (1 - \alpha)a(g_t, t)] \Delta + b(g_t, t)\Delta \mathbf{W}_t$$

where $\alpha \in [0, 1]$ characterizes the “implicitness-degree”. In particular, note that for $\alpha = 0$ and $\alpha = 1$ we have the drift-explicit and drift-implicit schemes above, whereas for $\alpha = 1/2$ we get

“a generalization of the deterministic trapezoidal method” (sometimes also referred to as Crank-Nicolson).

In our terminology, we therefore have:

$$g_{t+\Delta} - g_t = \left[\alpha \hat{\mathcal{A}}_{t+\Delta}^* g_{t+\Delta} + (1 - \alpha) \hat{\mathcal{A}}_t^* g_t \right] \Delta + \mathcal{B}_t^* g_t \Delta \mathbf{W}_t$$

which, at each time-step, solves the following:

$$\left(\mathcal{I} - \alpha \Delta \hat{\mathcal{A}}_{t+\Delta}^* \right) g_{t+\Delta} = \left(\mathcal{I} + (1 - \alpha) \Delta \hat{\mathcal{A}}_t^* \right) g_t + \mathcal{B}_t^* g_t \Delta \mathbf{W}_t.$$

In terms of the numerical implementation, if we define a time step Δ we can therefore solve the KFE using a finite-difference scheme by solving the corresponding initial value problem. We denote by $\mathbf{g}_t$ the vector which discretizes $g_t(\mathbf{x})$ on a grid of points for $\mathbf{x}$. We assume that $\mathbf{g}_0$ is the density in the risky steady state (i.e. the long-run mean density); the update in the next time period is then given by the KF equation:

$$\mathbf{g}_{t+\Delta} = \left(\mathbf{I} - \alpha \Delta \left[\mathbf{A}_{t+\Delta}^T + \zeta(\Psi - \mathbf{I}) \right] \right)^{-1} \left(\mathbf{I} + (1 - \alpha) \Delta \left[\mathbf{A}_t^T + \zeta(\Psi - \mathbf{I}) \right] + \sqrt{\Delta} \mathbf{B}_t^T \epsilon_t \right) \mathbf{g}_t$$

where $\mathbf{A}_t^T$ is the discretization of the infinitesimal operator $\mathcal{A}_t^*$, $\mathbf{B}_t^T$ is the discretization of $\mathcal{B}_t^*$, and $\epsilon_t \equiv \Delta \mathbf{W}_t / \sqrt{\Delta}$ is the vector of standardized innovations to $\mathbf{W}_t$.

Combining the above α -scheme with the non-uniform adjustment for the density described in online appendix F.3.1, and making use of the fact that the infinitesimal generator $\mathcal{A}^*$ does not depend on t , we can therefore solve the KFE in equation (10) by solving the following sequence of systems of linear equations at each time step n :

$$\mathbf{D}_\alpha \tilde{g}^{n+1} = \mathbf{E}_\alpha \tilde{g}^n \tag{F.13}$$

where $\mathbf{D}_\alpha$ and $\mathbf{E}_\alpha$ are the implicit and explicit matrices:

$$\begin{aligned} \mathbf{D}_\alpha &= (1 + \Delta \alpha \zeta) \mathbf{I} - \Delta \alpha \left(\mathbf{A}^T + \zeta \Psi \right) \\ \mathbf{E}_\alpha &= (1 - \Delta(1 - \alpha) \zeta) \mathbf{I} + \Delta(1 - \alpha) \left(\mathbf{A}^T + \zeta \Psi \right) + \sqrt{\Delta} \mathbf{B}^T \epsilon_n. \end{aligned}$$

Given the above equation (F.13), solving for the risky steady state is trivial as it suffices setting $\epsilon_n = \mathbf{0} \forall n$. To solve for the stochastic steady state I instead run a long simulation starting from the risky steady state with $\epsilon_n \sim \mathcal{N}(\mathbf{0}, \mathcal{I}_3)$ and discard the first N_{burn} periods to remove dependency from the starting distribution.

Finally, to compute the IRF to a shock in asset class j , I set the innovation to ϵ_0 so that the shocked return moves by one standard deviation on impact while the other two returns are unaffected, and set all subsequent innovations to zero; orthogonalized shocks for correlated asset returns are obtained via the Cholesky decomposition of the return volatility matrix $\boldsymbol{\sigma}$.

BIRTH OPERATOR. I conclude this appendix by describing the birth-death structure that enters the KFE through the operator Ψ . Agents die at rate ζ , which raises the effective discount rate to $\rho + \zeta$ in the HJB (see online appendix F.1.1). Because annuity markets are incomplete, liquid wealth is annuitized at $\bar{r}_{B,\text{ann}} = \bar{r}_B + \zeta$, but housing is not — it is instead inherited at death. Newborns enter with no liquid wealth, $a = 0$, in the non-participation state $p = 0$, and draw earnings and housing from stationary marginal distributions:

$$\psi(a, h, z, p) = \mathbf{1}[a = 0] \cdot \mathbf{1}[p = 0] \cdot g_h(h) \cdot \pi_z(z)$$

where $g_h(h) = \sum_{a,z,p} \tilde{g}(a, h, z, p)$ is the marginal housing distribution of the current cross-section and π_z is the stationary distribution of the continuous-time Markov chain for earnings. I write this as a sparse $N \times N$ matrix Ψ such that $\Psi \tilde{g}$ equals the birth distribution: for each housing level h_j , the operator marginalizes $\tilde{g}$ over (a, z, p) and places the resulting mass at $(a = 0, h_j, \pi_z, p = 0)$. Because the marginalization is linear in $\tilde{g}$, the operator Ψ is a fixed, precomputed matrix. The model admits N_ρ discount-rate types; the intensity matrix $\mathbb{A}$ is block-diagonal across types (no cross-type transitions during life), and Ψ is the sole source of cross-type mixing: newborns are assigned type m with probability $\pi_\rho(m)$.

F.4. Computational Details

The model is solved in Julia using sparse direct methods. All linear systems arising from the HJB and KFE iterations are solved with the Panua Pardiso library (Schenk and Gärtner 2004; Schenk et al. 2000).

The state space is discretized on $N_a = 51$ wealth grid points (geometric spacing), $N_h = 11$ housing grid points, $N_z = 35$ earnings states, $N_p = 2$ participation states, and $N_\rho = 3$ discount-rate types, for a total of $N = 117,810$ grid points.

All model simulations are run using a fully drift-implicit backward Euler method (i.e., $\alpha = 1$) at a quarterly frequency $\Delta = 1/4$, with the exception of the historical simulation and its counterfactuals/decompositions which are run at a yearly frequency.

References

- ACHDOU, Y., J. HAN, J.-M. LASRY, P.-L. LIONS, and B. MOLL (2022). “Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach”. *The Review of Economic Studies*.
- AGUIAR, M. and M. BILS (2015). “Has Consumption Inequality Mirrored Income Inequality?” *American Economic Review*.
- ANDREWS, I., M. GENTZKOW, and J. M. SHAPIRO (2017). “Measuring the Sensitivity of Parameter Estimates to Estimation Moments”. *The Quarterly Journal of Economics*.
- ARNOUD, A., F. GUVENEN, and T. KLEINEBERG (2022). “Benchmarking Global Optimizers”. *Working Paper*.
- BARRO, R. J. and C. J. REDLICK (2011). “Macroeconomic Effects From Government Purchases and Taxes”. *The Quarterly Journal of Economics*.

- CARTIS, C., J. FIALA, B. MARTEAU, and L. ROBERTS (2019). “Improving the Flexibility and Robustness of Model-Based Derivative-Free Optimization Solvers”. *ACM Transactions on Mathematical Software*.
- DUFFIE, D. and L. G. EPSTEIN (1992). “Stochastic Differential Utility”. *Econometrica*.
- GUVENEN, F., F. KARAHAN, S. OZKAN, and J. SONG (2021). “What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics?” *Econometrica*.
- HEATHCOTE, J., K. STORESLETTEN, and G. L. VIOLANTE (2017). “Optimal Tax Progressivity: An Analytical Framework”. *The Quarterly Journal of Economics*.
- HEIKKURI, V.-M. and M. SCHIEF (2026). “Subgroup Decomposition of the Gini Coefficient: A New Solution to an Old Problem”. *Econometrica*.
- JORDÀ, Ò., K. KNOLL, D. KUVSHINOV, M. SCHULARICK, and A. M. TAYLOR (2019). “The Rate of Return on Everything, 1870–2015”. *The Quarterly Journal of Economics*.
- KAPLAN, G., B. MOLL, and G. L. VIOLANTE (2018). “Monetary Policy According to HANK”. *American Economic Review*.
- KLOEDEN, P. E. and E. PLATEN (1995). *Numerical Solution of Stochastic Differential Equations*. Springer.
- KUHN, M. and J.-V. RÍOS-RULL (2016). “2013 Update on the U.S. Earnings, Income, and Wealth Distributional Facts: A View from Macroeconomics”. *Federal Reserve Bank of Minneapolis Quarterly Review*.
- RAFFINETTI, E., E. SILETTI, and A. VERNIZZI (2015). “On the Gini Coefficient Normalization When Attributes with Negative Values Are Considered”. *Statistical Methods & Applications*.
- SCHENK, O. and K. GÄRTNER (2004). “Solving Unsymmetric Sparse Systems of Linear Equations with PARDISO”. *Future Generation Computer Systems*.
- SCHENK, O., K. GÄRTNER, and W. FICHTNER (2000). “Efficient Sparse LU Factorization with Left-Right Looking Strategy on Shared Memory Multiprocessors”. *BIT Numerical Mathematics*.
- SCHRODER, M. and C. SKIADAS (1999). “Optimal Consumption and Portfolio Selection with Stochastic Differential Utility”. *Journal of Economic Theory*.
- SKIADAS, C. (2007). “Dynamic Portfolio Choice and Risk Aversion”. In: *Handbooks in Operations Research and Management Science*.
- TRABANDT, M. and H. UHLIG (2011). “The Laffer Curve Revisited”. *Journal of Monetary Economics*.