The Inequality Multiplier:

Market Inelasticity and the Persistence of Wealth Inequality

C. Christopher Hyland [*] Aditya Khemka [†]

July 20, 2026

Abstract

We study how frictions in financial intermediation transmit income inequality into wealth inequality. When mandate-constrained intermediaries maintain stable allocations between equity and debt—keeping equity demand inelastic—household saving and borrowing flows affect equity prices rather than being absorbed entirely through quantities. Greater saving by wealthy households and greater borrowing elsewhere in the distribution both expand intermediary balance sheets and raise equity valuations. Because equity ownership is concentrated, these revaluations redistribute wealth toward the top—an amplification mechanism we call the *inequality multiplier*. We develop a general equilibrium model with heterogeneous households and calibrate it to the United States from 1989 to 2023. Inelastic pricing roughly doubles the response of wealth inequality relative to an elastic market benchmark, and the multiplier strengthens as wealth becomes more concentrated. Because the income inequality generating these flows is persistent, the resulting wealth concentration is durable. Consistent with the mechanism, the component of equity returns predicted by fund flows generates persistent shifts in wealth shares, whereas the response to fundamental components mean reverts. Our results show that wealth inequality depends not only on household saving and portfolio exposure, but also on how financial institutions transmit distributional flows into asset prices.

Keywords: wealth inequality, market inelasticity, household debt.

JEL Classification: E44, D31, G12, G23, G51.

A supplemental appendix to the paper can be accessed at this link.

We thank Andrea Ferrero and Dimitrios Tsomocos for helpful comments and continued support. We also thank Joseph Abadi, Pedro Bordalo, Abhinash Borah, Markus Brunnermeier, Riccardo Cioffi, David Domeij, Goutham Gopalakrishna, Sebastian Hanson, Niels Johannesen, Jiri Knesl, Ralph Koijen, Seung Joo Lee, Mingze Ma, Zach Mazlish, Max Miller, Benjamin Moll, Ken Okamura, Shumiao Ouyang, Udara Peiris, Ricardo Reis, Joel Shapiro, Amir Sufi, Paolo Surico, Oren Sussman, Purnoor Tak (Discussant), Christoph Trebesch, Florian Trouvain, Gianluca Violante, Mungo Wilson, Kieran Walsh, Tim Willems and other colleagues and conference participants for helpful comments and discussions.

[*] Christopher Hyland: Saïd Business School and Exeter College, University of Oxford and Oxford-Man Institute of Quantitative Finance, christopher.hyland.dphil@said.oxford.edu

[†] Aditya Khemka (Corresponding Author): Saïd Business School and New College, University of Oxford aditya.khemka.dphil@said.oxford.edu

1. Introduction

Over recent decades, wealth inequality in the United States has risen much faster than income inequality.[1] Existing explanations emphasise heterogeneity in saving rates, portfolio composition, and asset returns across the wealth distribution. Wealthy households—the top decile—save more and hold a disproportionate share of equity, while households outside the top decile hold less equity and are more indebted (Figure A1). These differences shape not only who benefits from asset price appreciation, but also how changes in saving and borrowing shift asset demands and, through them, equilibrium prices, credit, and interest rates. Understanding the dynamics of wealth inequality therefore requires studying the saving decisions of wealthy households and the borrowing decisions of the rest of the distribution jointly.[2]

How do these changing asset demands translate into asset prices and quantities? A growing literature in finance shows that the answer depends on the *market macrostructure*[3]: the institutional organisation of asset markets and the structure of financial intermediation. This macrostructure has changed markedly in recent decades, most visibly through the growth of passive ownership to roughly one-third of the equity market. Haddad et al. (2025) show that the rise of passive investment has contributed to making asset demand in equity markets more *inelastic*. We study how this inelasticity shapes the joint dynamics of equity prices, debt, interest rates, and wealth inequality—trends that have moved together in the United States since 1989 (Appendix A, Figure A2).

A common implication of market inelasticity, irrespective of its precise microfoundation,[4] is that quantities of assets in fixed or inelastic supply respond weakly to flows entering or leaving the market. Instead, flows generate large price movements and adjustments in more elastically supplied assets, even when fundamentals are unchanged. If households differ in their exposure to these assets, the resulting price and quantity adjustments redistribute wealth across the distribution.

We first show that the distinction between flow-driven and fundamentals-driven returns is empirically relevant. Gomez (2025) documents that equity returns generate large and persistent movements in top wealth shares, but those returns may reflect either changes in fundamentals or investor flows (Gabaix and

[1] For example, Kuhn et al. (2020) show that the evolution of asset prices introduced a wedge between the evolution of the wealth distribution and the income distribution. Using different methodologies, Piketty and Zucman (2014) and Smith et al. (2023) show that wealth to income ratios have increased substantially.

[2] Theories based on return heterogeneity, such as Gomez (2025), do not by themselves account for the simultaneous rise in aggregate debt. Conversely, theories that link wealth inequality to the indebtedness of households at the bottom of the wealth distribution, such as Kumhof et al. (2015), do not explain the large increase in equity values. Benhabib et al. (2017) further show that standard models of income inequality cannot generate the much larger and faster rise in wealth inequality.

[3] For an overview of the literature, please refer to Haddad and Muir (2025).

[4] Inelasticity may arise from intermediary-level frictions and constraints including fixed mandates, capital constraints, preferred habitat, transaction costs, or rational inattention. Supplemental Appendix 2 provides a more complete, though non-exhaustive, discussion.

Koijen 2021). We decompose excess equity returns into a component predicted by residualised mutual fund and ETF flows, constructed in the spirit of Coval and Stafford (2007), and a residual component capturing discount-rate and cash-flow news. We then estimate the response of top 10% and bottom 90% wealth shares to each component using local projections. Appendix I describes the construction and empirical strategy.

Figure 1 reports the results. The flow-driven component persistently raises the top 10% wealth share and lowers the bottom 90% share, with no meaningful reversal over the following three years. By contrast, the response to the residual fundamentals-driven component mean reverts and eventually crosses zero. Thus, in this exercise, the persistent redistribution associated with equity returns is concentrated in the component predicted by investor flows. This pattern is consistent with the interaction of inelastic markets and heterogeneous household portfolios.

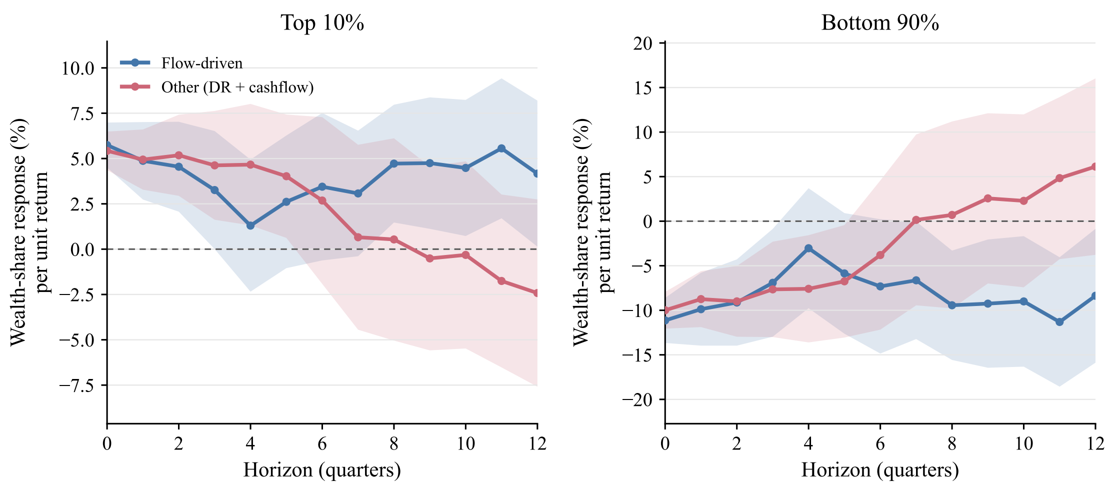

FIGURE 1. Response of the wealth share to the flow-driven and fundamentals-driven components of equity returns, top 10% (left) and bottom 90% (right).

Note: Local projections of the log net-worth share on the flow-driven component (blue) and the residual 'other' component (red) of Gomez (2025)'s excess equity return, per unit of log return, at horizons of 0–12 quarters. The flow-driven component is the part of the excess return predicted by a residualised mutual fund and ETF flow innovation; the other component collects discount-rate and cash-flow news. Shaded bands are 90% Newey–West confidence intervals. Net-worth shares are from the Distributional Financial Accounts, quarterly, 1989Q3 onward. Construction and estimation are detailed in Appendix I.

Motivated by this evidence, we build and calibrate a quantitative consumption–saving model with three household groups: wealthy households in the top 10%, middle class households in the middle 40%, and non-wealthy households in the bottom 50%. Wealthy households invest in equity and debt through an intermediary whose portfolio allocation is costly to adjust, while middle class and non-wealthy households hold little or no equity, hold owner-occupied housing, and borrow in debt markets, in part against housing collateral. Because the intermediary maintains a relatively stable allocation between equity and bonds, its balance sheet creates a tight equilibrium link between the market value of equity and the quantity of outstanding debt. Under constrained intermediation, an expansion of the intermediary's bond position requires a corresponding increase in the market value of its equity position, making equity values and debt first-order complements, rather than substitutes as in a canonical elastic markets model.

This balance sheet link has direct distributional consequences. The wealthy own the intermediary's assets, while middle class and non-wealthy households issue much of the debt that it holds. An increase in equity prices is therefore accompanied by an expansion of liabilities elsewhere in the wealth distribution, redistributing balance sheet value from borrowers toward asset holders. We call the amplification of an equity price movement into a wealth share movement the *inequality multiplier*. Relative to an elastic benchmark, the intermediary mandate raises the wealthy household's effective equity exposure from its explicit equity share to nearly one: the bond component of its intermediated portfolio co-moves with the equity price rather than hedging it. The multiplier is also self-reinforcing: as wealth becomes more concentrated, a given shock produces a larger movement in top wealth shares. At the calibrated mandate and wealth distribution, the inelastic-to-elastic amplification is about 1.9 in 1989 and rises to roughly 2.2 by 2023.

A simple thought experiment illustrates the mechanism. Consider a wealthy household that invests through an intermediary and a non-wealthy household that can borrow but cannot hold equity. Suppose the intermediary has net worth of \$100 and is required to hold 80% in equity and 20% in bonds. Its initial balance sheet therefore contains \$80 of equity and \$20 of bonds. Now suppose the wealthy household receives a \$1 income shock and invests it in the intermediary. Because equity supply is fixed, the intermediary directs the dollar into bonds, increasing its bond position to \$21.[5] To preserve its 80–20 allocation, the equity side of the balance sheet must rise to \$84. Since the number of shares is fixed, this adjustment occurs through a higher equity price. The additional demand for bonds also increases the supply of credit and lowers the interest rate, inducing further borrowing by non-wealthy households. Higher equity prices and higher debt therefore arise together: the former increase the wealth of equity holders, while the latter reduce the net worth of borrowers. In an elastic market model, by contrast, equity values and debt are not tied by this balance sheet restriction. This is the *equity investment* channel; the *borrowing* channel runs in reverse, when an exogenous rise in household or government debt forces the intermediary to bid up equity to preserve its mandate.

The paper makes three contributions. First, we derive the equilibrium pricing relation generated by constrained intermediation and an analytical expression for the inequality multiplier. The expression decomposes the wealth share response into a portfolio heterogeneity term and a price response term, and shows

[5] This is a simplification. We do not claim that all household debt is held directly by mutual funds or similar intermediaries. As shown by Mian et al. (2025), household borrowing can be traced through layers of intermediation to its ultimate funding sources. Their analysis suggests that more than 25% of household debt is intermediated through relatively inelastic institutions, and that mutual funds account for a substantial share of the increase in demand for household-debt instruments.

that amplification increases with existing wealth concentration.

Second, we calibrate the model to the U.S. economy beginning in 1989 and study its response to a reduced-form portfolio demand shock corresponding to the flow innovation in Figure 1. Both the elastic and inelastic models raise the equity price and top wealth share on impact. Only the inelastic model, however, generates the empirically relevant combination of higher equity prices, greater aggregate borrowing, a lower safe rate, and a hump-shaped, persistent increase in the top wealth share. The elastic model instead reduces credit supply, raises the safe rate, and produces a wealth share response that mean reverts relatively quickly.

Third, we assess the model's ability to account for the historical dynamics of wealth shares, household debt, and interest rates. Under perfect foresight, inelasticity improves the joint fit of wealth shares at the top and bottom of the distribution, while also generating more household borrowing and a larger decline in the safe rate than the elastic benchmark. Introducing unanticipated (MIT) shocks—so that households do not foresee future innovations and are surprised anew each period—reactivates endogenous borrowing and helps account for the post-crisis rise in top wealth and the deterioration in bottom wealth. Under unanticipated shocks, aggregate historical paths are less informative about the pricing regime, making independent evidence on price sensitivity to flows important for identification. The equity investment channel accounts for the level of wealth inequality across the distribution and for the path of the top share throughout the sample, including its post-2008 acceleration, while the borrowing channel's distinctive contribution is concentrated at the bottom of the distribution and around the financial crisis. Throughout, the model matches the comovement of equity prices with borrowing and the associated balance sheet transfer between the top and the bottom, rather than the level of the equity price run-up, which it underpredicts. These main results survive the introduction of a realistically calibrated unconstrained arbitrageur, which attenuates but does not eliminate the mechanism (Appendix H).

The framework reconciles two prominent explanations of rising wealth inequality. Investment-based accounts emphasise the higher saving rates, capital accumulation, and asset returns of wealthy households (Hubmer et al. 2021; Kuhn et al. 2020). Debt-based accounts emphasise the increasing leverage of lower wealth households and the role of wealthy households in financing that debt (Rajan 2011; Mian et al. 2025). These two channels map onto the capital-based and debt-based explanations of inequality, which our model shows operate in tandem. In doing so, we endogenise the capital flows that Gabaix and Koijen (2021) take as given: the income distribution determines household saving and borrowing, so the persistence of flows inherits the persistence of income shocks. Because persistent income inequality continually feeds the multiplier, the resulting wealth concentration need not mean revert in the manner suggested by Cioffi (2021).

Why does the inequality multiplier matter? First, it provides a mechanism through which modest and persistent changes in the income distribution can generate larger and faster changes in wealth concentration. By amplifying the effect of asset price movements on wealthy households' balance sheets—and by becoming stronger as wealth becomes more concentrated—the multiplier helps explain both the widening gap between income and wealth inequality and the rapid transition dynamics of top wealth shares. Second, it identifies a potential distributional consequence of changes in financial market structure. The growth of passive and mandate-constrained intermediation may reduce costs and improve diversification, but by lowering market elasticity it can also cause flows to have larger and more persistent effects on asset prices and wealth shares. Third, the mechanism shows that, in inelastic markets, an equity boom is accompanied by debt creation and a balance sheet redistribution from borrowers toward asset holders. This distinction identifies different policy margins. The unequal incidence of asset price appreciation points toward broadening equity participation, while the concentration of borrowing effects at the bottom and around periods of financial stress points toward household leverage, public debt supply, and intermediary portfolio rigidity. Because we do not solve an optimal policy problem, we interpret these results as identifying relevant mechanisms rather than prescribing particular instruments.

Outline of the Paper. We first discuss the paper's positioning and contribution to the literature. Then, Section 2 sets out a quantitative DSGE model and discusses the main qualitative results. Section 3 calibrates and provides the empirical matching of moments. Section 4 studies the impulse response of wealth shares to a reduced-form flow shock. Section 5 analyses wealth inequality dynamics and conducts the main simulation exercises. Section 6 concludes.

1.1. Related Literature and Contribution

This paper connects the market macrostructure literature, reviewed by Haddad and Muir (2025), to the study of wealth inequality. Demand-system asset pricing establishes that institutional demand affects equilibrium valuations (Koijen and Yogo 2019; Koijen et al. 2024), while Gabaix and Koijen (2021) show that flows have large price effects when aggregate equity demand is inelastic. Haddad et al. (2025) provide related evidence that institutional changes, including the rise of passive investing, have made equity demand less elastic.[6] We endogenise the flows that move prices: the distribution of income determines household saving and borrowing, while portfolio mandates govern how intermediaries allocate these flows between debt and equity. This links income inequality to equity valuations, aggregate debt, and the distribution of wealth.

[6]Supplemental Appendix 2 discusses the evidence and microfoundations for inelastic equity demand.

The closest mechanisms in household finance are those of Chien et al. (2011) and Melcangi and Sterk (2025). The former studies equilibrium risk sharing between passive and active investors, while the latter links heterogeneous saving and stock market participation to mutual fund flows, investment, and monetary transmission. In our model, mandate-constrained intermediaries transmit both wealthy household saving and household borrowing into secondary market equity prices. Related work embeds institutional asset demand in general equilibrium (Abadi 2026; Gopalakrishna et al. 2026); our focus is the distributional origin of flows and the joint determination of equity values and household debt.

The paper also contributes to the literature on asset prices and wealth inequality. Portfolio heterogeneity makes asset returns an important source of wealth redistribution (Kuhn et al. 2020; Hubmer et al. 2021; Cioffi 2021; Fagereng et al. 2025). Toda and Walsh (2020) link income inequality to the equity premium through heterogeneity in risk aversion or beliefs; we instead obtain asset-price effects of inequality through market structure, without requiring preference heterogeneity. Greenwald et al. (2021) and Catherine et al. (2025) show that falling interest rates disproportionately revalue the long-duration financial assets held by wealthier households. Our mechanism reinforces this duration channel: lower required returns revalue equity, while inflows into constrained intermediaries generate additional equity demand. It also operates when rates do not fall, since greater borrowing expands intermediary balance sheets and, through portfolio mandates, raises equity valuations. Unlike Gomez (2025), who study the interaction between asset prices and the Pareto tail, we show how saving and borrowing throughout the distribution feed back into prices and wealth concentration. Because the multiplier amplifies each shock, wealth shares respond quickly, in contrast to the slow transition dynamics of standard random growth models of wealth inequality (Gabaix et al. 2016).

More broadly, our paper relates to work on financial frictions and the wealth distribution. Fernández-Villaverde et al. (2023) show how borrowing and lending constraints create high leverage states in which shocks have larger distributional effects, while Ding and Jiang (2026) and Hui (2025) connect the wealth distribution to safe asset demand and asset valuations. Our friction instead restricts the portfolio allocation of financial intermediaries. As a result, changes in household or government debt affect inequality not only through interest rates and borrower–creditor exposures, but also through the revaluation of equity. This cross-asset transmission is central to the inequality multiplier.

Finally, our paper connects debt- and capital-based accounts of rising inequality. In Mian et al. (2021), heterogeneous saving propensities generate debt accumulation, lower aggregate demand, and a lower natural rate. We retain this household-side force but add an asset pricing channel: debt absorbed by mandate-constrained intermediaries raises their equity demand and redistributes wealth toward equity owners. Our elastic market benchmark preserves the indebted demand mechanism while removing this balance sheet linkage, isolating the contribution of market inelasticity. Relatedly, Mian et al. (2025) trace saving by the wealthy to middle class and government borrowing, while Auclert et al. (2025) study how inequality shifts aggregate asset demand. We show how these flows also affect equity valuations and, because equity ownership is concentrated, wealth inequality.

2. Quantitative Model

Our aim is to isolate a single feedback loop—between the wealth distribution and inelastic asset markets—so we reduce both sides to their minimal representation. The mechanism turns on two objects: the between-group wedge in saving and portfolio composition, and the aggregate price-elasticity of equity demand. We therefore model neither the full cross-section of the wealth distribution nor a structural asset-demand system, since richer versions of each refine magnitudes without altering the sign or logic of the amplification. To study the effect of inelastic markets on the dynamics of wealth inequality, we develop a dynamic stochastic general equilibrium model. In our setting, wealthy households must hold equity through a financial intermediary, whose (constrained) portfolio-allocation decisions are central to determining asset price and wealth dynamics.

2.1. Setup

The model is discrete time with three representative household groups, indexed by $i \in \{R, M, P\}$, where R denotes wealthy households (top 10% of the wealth distribution), M denotes middle class households (middle 40%, i.e. the 50th–90th wealth percentiles), and P denotes non-wealthy households (bottom 50%). To keep the analysis simple, we assume that households are fixed within their wealth groups and cannot transition across groups.[7]

Markets. There are two long-lived assets.[8] The first is a Lucas-tree equity claim in fixed supply $\bar{Q}$. It pays a risky dividend D_t and trades at ex-dividend price p_t. The second is owner-occupied housing in fixed

[7] Kuhn et al. (2020) show using the Panel Study of Income Dynamics (PSID) that a high share of households (over 80%) remain in the same wealth group across time. Gomez (2023) shows that the importance of inequality driven by movements between wealth groups has decayed over time. Thus, the 'synthetic' assumption of keeping households fixed within wealth groups provides a good approximation of aggregate wealth dynamics.

[8] We keep the asset side deliberately spare: a single aggregate equity claim rather than a cross-section of stocks, no endogenous risk premia (the model is solved by log-linearisation, so prices are certainty-equivalent), and a reduced-form mandate-deviation cost in place of a structural demand system. Each is a modelling choice rather than a limitation for our question. The inequality multiplier depends on the *aggregate* price-elasticity of equity demand and on the concentration of equity ownership, not on cross-asset substitution or time-varying risk compensation. The mandate cost is isomorphic to the downward-sloping aggregate demand curve a multi-asset non-linear system would imply, and we discipline its slope directly on the empirical micro-elasticity of Gabaix and Koijen (2021) and Haddad et al. (2025) rather than on any single institutional friction; Supplemental Appendix 2 discusses the microfoundations.

aggregate supply $\bar{H}$. Housing pays a risky housing-service dividend D_t^H and trades at price p_t^H. In addition, one-period risk-free bonds are issued by household borrowers and the government. We write B_t^M, B_t^P, and B_t^G as positive quantities of debt issued by middle class households, non-wealthy households, and the government, respectively.

Positions chosen at the end of period t pay gross returns at the beginning of period $t+1$. Define the gross risk-free return, the gross equity return and the gross housing return as:

$$(1) \qquad R_t \equiv 1 + r_t, \qquad R_t^E \equiv \frac{D_{t+1} + p_{t+1}}{p_t}, \qquad R_t^H \equiv \frac{D_{t+1}^H + p_{t+1}^H}{p_t^H}.$$

Segmentation. Motivated by the evidence on heterogeneous portfolio holdings (see Appendix A.1), middle class and non-wealthy households cannot participate in equity markets,[9] but they may hold housing and borrow in debt markets. By contrast, wealthy households may invest in both equity and bonds indirectly through the financial intermediary.[10]

Intermediary. The financial intermediary ('mutual fund') receives flows from the wealthy household and holds a portfolio of equity and risk-free bonds. The intermediary is subject to portfolio frictions in the form of a mandate-deviation cost (detailed below). Therefore, it cannot fully arbitrage between the three asset classes. In Appendix H, we demonstrate that incorporating a realistically-calibrated arbitrageur into the model does not materially alter the results.

Define the gross return on the mutual fund as:

$$(2) \qquad R_t^{MF} \equiv \theta_t R_t^E + (1 - \theta_t) R_t,$$

where θ_t is the equity value share in the mutual fund portfolio.

Households. Households face perpetual-youth mortality with death probability δ each period (Yaari (1965); Blanchard (1985); Farhi and Werning (2019)). The effective discount factor of household group i is $\beta^i = \frac{1}{1+\rho^i+\delta}$, where ρ^i is the rate at which household i discount future utility. The calibration makes wealthy households the most patient group and non-wealthy households the least patient group. To match empirical evidence on

[9]This can be microfounded by differences in tastes and preferences, margin requirements, higher risk aversion, or other institutional constraints, as documented in Campbell (2006) and Vissing-Jørgensen (2002).

[10]For tractability—and because housing is a minor component of wealthy household portfolios—we assume that wealthy households do not invest in housing. Additionally, our model does not feature private equity. Although private equity now constitutes a quantitatively significant share of wealthy portfolios, and the rise of private capital markets has amplified inequality, the results of Göcmen et al. (2025) (Figure 10) show that accounting for private capital is most relevant for studying inequalities *within* the top 10 per cent (especially the top 1 or 0.1 per cent), rather than between the top 10 and bottom 50, which is our focus.

heterogeneous marginal propensities to save, we incorporate non-homothetic preferences for the wealthy household (Straub (2019), Mian et al. (2021)).[11]

Shocks. In addition to equity and housing dividends, shocks in the model are wealthy, middle class and non-wealthy endowments, government debt, the middle class loan-to-value ratio, and wealthy demand for equity. These are introduced in their relevant subsections below and their log-linear laws of motion are reported in Appendix Section C.2

2.2. Financial Intermediary

The mutual fund purchases bonds (or provides loans) to the middle class, non-wealthy households and the government, and holds the entire supply of equity. Each period, it pays dividend and interest income back to the wealthy household and receives flows (savings) from the wealthy household. The mutual fund balance sheet is

$$X_t^{MF} = p_t Q_t + B_t^{MF}, \tag{3}$$

where X_t^{MF} denotes mutual fund assets under management and B_t^{MF} denotes the bond position held by the fund. The portfolio share of equity is

$$\theta_t \equiv \frac{p_t Q_t}{X_t^{MF}}. \tag{4}$$

Equivalently, equity and bond holding values can be written as

$$p_t Q_t = \theta_t X_t^{MF}, \qquad B_t^{MF} = (1 - \theta_t) X_t^{MF}. \tag{5}$$

The intermediary's objective is to choose its equity share to maximise the return on its assets under management, subject to a quadratic cost of deviating from its target equity share θ^*.[12] A reduced-form representation

[11] We keep the household side spare relative to the quantitative-inequality literature: three groups rather than a continuous distribution, exogenous group endowments rather than a calibrated persistent–transitory labour-income process, and non-homothetic preferences rather than a full cross-section of marginal propensities to consume. These devices suffice for our question. The multiplier operates through the *between-group* wedge in saving rates and portfolio shares—which we calibrate to the data—and transmits through balance-sheet revaluation, not through the consumption/aggregate-demand channel in which within-group MPC dispersion is first-order. A richer income process and a continuous distribution are what one needs to *generate* the wealth distribution endogenously; here its forcing process is taken as the input to the asset-market question, so that machinery is upstream of, and separable from, the mechanism we isolate.

[12] We introduce inelasticity through a quadratic cost of deviating from the target mandate share θ^*, motivated by the evidence on infrequent portfolio adjustment in Bacchetta et al. (2023) and Giglio et al. (2021). This makes the portfolio share sticky, in the same spirit as the price-setting frictions of the monetary economics literature (Rotemberg 1982); up to a first-order approximation, the same results would arise if the intermediary instead adjusted its portfolio at a Poisson arrival rate, following Calvo (1983).

of the fund's problem is

$$\max_{\theta_t} \mathbb{E}_t \left[\beta^{MF} \{\theta_t R_t^E + (1-\theta_t) R_t\} - \frac{\chi^{MF}}{2} \left(\frac{\theta_t - \theta^*}{\theta^*} \right)^2 \theta^* \right]. \tag{6}$$

The parameter β^{MF} is the mutual fund's time discount rate, while χ^{MF} governs the degree of portfolio rigidity. As χ^{MF} rises, the mutual fund becomes less willing to move away from its target equity share, holding the market closer to its inelastic limit. In the limiting case $\chi^{MF} \to \infty$, the equity share is fixed at θ^*.

The mutual fund's demand for equity is given by its first-order condition

$$\theta_t = \theta^* \left[1 + \frac{\beta^{MF}}{\chi^{MF}} \underbrace{(R_t^E - R_t)}_{\text{Equity Premium}} \right]. \tag{7}$$

The intermediary's equity allocation is thus proportional to its target θ^* and positively related to the equity premium. When the equity premium is high, the intermediary may choose to hold more equity, albeit at the cost of adjustment. This mechanism allows for time-varying market inelasticity.

It is worth separating the two parameters that govern market inelasticity, since they play distinct roles. The target share θ^* sets the *structural* level of inelasticity—the slope of intermediary demand, $\xi \approx 1 - \theta^*$, and the price impact $\theta^*/(1-\theta^*)$—and, together with the wealth distribution, the strength of the inequality multiplier. The rigidity χ^{MF} instead governs how tightly the intermediary holds to θ^*: as $\chi^{MF} \to \infty$ the share is fixed and the market sits exactly at its inelastic limit $\xi = 1 - \theta^*$, whereas a finite χ^{MF} lets θ_t chase the equity premium, adding elasticity of order $1/\chi^{MF}$ (Appendix E.1).

2.3. Wealthy Households

Wealthy households (top 10% of the wealth distribution) hold their entire wealth invested in the mutual fund. This wealth generates bequest utility in addition to consumption utility (non-homothetic preferences). Bequests play a quantitatively significant role in the aggregate economy and are central to explaining wealth inequality (De Nardi and Fella 2017), while non-homothetic preferences are crucial for generating higher (gross) savings rates towards the right tail of the wealth distribution (Straub 2019; Mian et al. 2025; Fagereng et al. 2025).[13]

The wealthy household receives risky income denoted in the numeraire consumption good and flow returns

[13] We model bequests in line with the seminal frameworks of De Nardi (2004) and Straub (2019), treating them as a luxury good. Non-homothetic preferences thus capture the observation that wealthier households save a larger fraction of income, either for bequests or for high-cost future expenditures such as education, medical care, or charitable giving (Benhabib and Bisin 2018).

from the mutual fund, and chooses a stream of consumption and mutual fund holdings to maximise lifetime utility (consumption utility plus the expected utility of bequests conditional on death), subject to its budget constraint:

$$\max_{\{c_t^R, X_t^{MF}\}_{t\geq 0}} \mathbb{E}_0 \sum_{t=0}^{\infty} (\beta^R)^t \left[u_R(c_t^R) + \delta \nu_R(a_t^R) \right] \tag{8}$$

subject to

$$a_t^R = X_t^{MF}, \tag{9}$$

$$c_t^R + X_t^{MF} + \Gamma^{MF}(X_t^{MF}) = e_t^R + R_{t-1}^{MF} X_{t-1}^{MF} + S_t^R. \tag{10}$$

Here, $\Gamma^{MF}(\cdot)$ is wealthy portfolio holding costs, e_t^R is risky income and S_t^R is the wealthy household's share of net government transfers.

The Euler equation for mutual fund holdings is

$$1 + \Gamma^{MF\prime}(X_t^{MF}) = \delta \zeta_t^R + \beta^R \mathbb{E}_t \left[R_t^{MF} \frac{u_R'(c_{t+1}^R)}{u_R'(c_t^R)} \right], \qquad \zeta_t^R \equiv \frac{\nu_R'(a_t^R)}{u_R'(c_t^R)}. \tag{11}$$

The quantitative specification introduces a mean-zero risky portfolio demand wedge, ζ_t^d, directly into this Euler equation:

$$1 + \Gamma^{MF\prime}(X_t^{MF}) - \phi_{MF} \zeta_t^d = \delta \zeta_t^R + \beta^R \mathbb{E}_t \left[R_t^{MF} \frac{u_R'(c_{t+1}^R)}{u_R'(c_t^R)} \right]. \tag{12}$$

We interpret ζ_t^d as a marginal-valuation or latent-demand shock for mutual fund shares. A positive ζ_t^d increases desired mutual fund holdings at a given path of consumption, wealth, and returns. Because the mutual fund must deploy its assets under management into both equity and bonds according to its mandate, this demand shift affects not only the equity price but also the risk-free bond market. This provides the closest model counterpart for flow shocks from Gabaix and Koijen (2021).

2.4. Non-Wealthy Households

The representative non-wealthy household (bottom 50% of the wealth distribution) receives risky income and housing service returns, and chooses a stream of consumption, housing purchases and sales, and borrowing to maximise lifetime utility:

$$\max_{\{c_t^P, X_t^{H,P}, B_t^P\}_{t\geq 0}} \mathbb{E}_0 \sum_{t=0}^{\infty} (\beta^P)^t u_P(c_t^P) \tag{13}$$

subject to

$$c_t^P + X_t^{H,P} + \Gamma_H^P(X_t^{H,P}) + \Gamma_B^P(B_t^P) + R_{t-1}B_{t-1}^P = e_t^P + (D_t^H + p_t^H)H_{t-1}^P + B_t^P + S_t^P. \tag{14}$$

where $X_t^{H,P} \equiv p_t^H H_t^P$ is the value of its housing position, $\Gamma_B^P(\cdot)$ and $\Gamma_H^P(\cdot)$ are debt and housing holding costs respectively, e_t^P is risky income and S_t^P denotes its share of government transfers. Unlike for wealthy households, we do not assume non-homotheticity in the preferences of non-wealthy and middle class households, since saving behaviour at the lower end of the wealth distribution is well approximated by homothetic preferences.

Non-wealthy borrowing is determined by

$$1 - \Gamma_B^{P\prime}(B_t^P) = \beta^P \mathbb{E}_t \left[\frac{u_P'(c_{t+1}^P)}{u_P'(c_t^P)} R_t \right], \tag{15}$$

while housing demand is determined by

$$1 + \Gamma_H^{P\prime}(X_t^{H,P}) = \beta^P \mathbb{E}_t \left[\frac{u_P'(c_{t+1}^P)}{u_P'(c_t^P)} \frac{D_{t+1}^H + p_{t+1}^H}{p_t^H} \right]. \tag{16}$$

2.5. Middle Class Households

The representative middle class household (the middle 40%) receives risky income and housing service returns, and chooses consumption, housing purchases, and borrowing to maximise lifetime utility, subject to its budget constraint and a borrowing constraint limiting debt against housing collateral, in the spirit of Justiniano et al. (2019):

$$\max_{\{c_t^M, X_t^{H,M}, B_t^M\}_{t\geq 0}} \mathbb{E}_0 \sum_{t=0}^{\infty} (\beta^M)^t u_M(c_t^M) \tag{17}$$

subject to the budget constraint

$$c_t^M + X_t^{H,M} + \Gamma_H^M(X_t^{H,M}) + R_{t-1}B_{t-1}^M = e_t^M + (D_t^H + p_t^H)H_{t-1}^M + B_t^M + S_t^M, \tag{18}$$

where $X_t^{H,M} \equiv p_t^H H_t^M$ is the value of its housing position, $\Gamma_H^M(\cdot)$ is the housing holding cost, e_t^M is risky income and S_t^M denotes its share of government transfers. The collateral constraint is

$$B_t^M \leq \tau_t^M X_t^{H,M}. \tag{19}$$

where $\tau_t^M \in \mathbb{R}_+$ is the loan-to-value ratio. In equilibrium, housing is supplied inelastically at $\overline{H}$, and we solve the model locally in the regime where the borrowing constraint binds with equality in every period.

Let μ_t^M be the multiplier on the collateral constraint and let $\psi_t^M \equiv \mu_t^M / u'_M(c_t^M)$ be the consumption-normalised multiplier. The bond Euler equation is

$$1 = \beta^M \mathbb{E}_t \left[\frac{u'_M(c_{t+1}^M)}{u'_M(c_t^M)} R_t \right] + \psi_t^M, \tag{20}$$

and the housing Euler equation is

$$1 + \Gamma_H^{M\prime}(X_t^{H,M}) = \beta^M \mathbb{E}_t \left[\frac{u'_M(c_{t+1}^M)}{u'_M(c_t^M)} \frac{D_{t+1}^H + p_{t+1}^H}{p_t^H} \right] + \tau_t^M \psi_t^M. \tag{21}$$

The maintained local-regime assumption is $\mu_t^M > 0$ along the simulated paths.

2.6. Government

The government issues (exogenously determined) bonds to fund transfers to all households based on a predetermined transfer scheme. Each period it maintains a balanced budget, such that new borrowing covers the cost of repaying interest on the previous period's borrowing plus transfers:

$$T_t = B_t^G - R_{t-1} B_{t-1}^G. \tag{22}$$

In the absence of household credit risk, bonds issued by the government and by households are qualitatively equivalent. Accordingly, government bonds are also purchased by the financial intermediary.

Transfers are allocated across the three household groups according to fixed exogenous shares s^i:

$$S_t^i = s^i T_t, \qquad i \in \{R, M, P\}, \qquad \sum_i s^i = 1. \tag{23}$$

2.7. Market Clearing and Equilibrium

DEFINITION 1. *An equilibrium consists of allocations* $\{c_t^R, c_t^P, c_t^M, X_t^{MF}, X_t^{H,M}, X_t^{H,P}, B_t^P, B_t^M\}_{t=0}^{\infty}$*, intermediary positions* $\{Q_t, B_t^{MF}, \theta_t\}_{t=0}^{\infty}$*, prices* $\{p_t, p_t^H, r_t\}_{t=0}^{\infty}$*, endowments and exogenous processes* $\{\overline{Q}, \overline{H}, e_t^P, e_t^R, e_t^M, D_t, D_t^H, B_t^G, \tau_t^M, \zeta_t^d\}_{t=0}^{\infty}$*, government policies and multipliers, such that households optimise, the mutual fund satisfies its mandate and balance sheet identities, the government budget constraint holds, all markets clear, and the sequence converges to the terminal steady state implied by the terminal values of the exogenous processes.*

For a full description of market clearing conditions, refer to Appendix C. The model is log-linearised and solved using Dynare. The deterministic steady state is derived in Appendix C.1, while the full log-linearised version of the model is presented in Appendix C.2.

2.8. Elastic Direct-Holding Benchmark

In order to establish the contribution of inelasticity as a mechanism, we first define an appropriate benchmark with elastic markets. Crucially, this benchmark continues to assume non-homothetic preferences for the wealthy, and uses the same specification as before for the middle class and the non-wealthy. To remove inelasticity, we remove the mutual fund, and allow wealthy households to directly allocate between equity and bonds, with no portfolio allocation restrictions.[14] In spirit, the model is similar to the model from Mian et al. (2021).

Specifically, wealthy households hold the equity claim, with value $E_t \equiv p_t Q_t$, and bonds B_t^R directly, subject to quadratic holding costs $\Gamma^E(E_t)$ and $\Gamma^B(B_t^R)$. Their assets are $a_t^R = E_t + B_t^R$, and the budget constraint is

$$c_t^R + E_t + B_t^R + \Gamma^E(E_t) + \Gamma^B(B_t^R) = e_t^R + (D_t + p_t)\bar{Q} + R_{t-1}B_{t-1}^R + S_t^R. \tag{24}$$

The equity and bond Euler equations are

$$1 + \Gamma^{E\prime}(E_t) - \phi_E \zeta_t^d = \delta\zeta_t^R + \beta^R \mathbb{E}_t\left[R_t^E \frac{u_R'(c_{t+1}^R)}{u_R'(c_t^R)}\right], \tag{25}$$

$$1 + \Gamma^{B\prime}(B_t^R) = \delta\zeta_t^R + \beta^R \mathbb{E}_t\left[R_t \frac{u_R'(c_{t+1}^R)}{u_R'(c_t^R)}\right]. \tag{26}$$

Equity is therefore priced by the wealthy household's Euler equation rather than by the mandate identity of the mutual fund, and bond-market clearing becomes $B_t^R = B_t^P + B_t^M + B_t^G$. The portfolio demand wedge enters

[14] This is equivalent to allowing the mutual fund to choose its portfolio allocation freely, i.e., converting the intermediary into a 'veil'.

the equity Euler equation, so it shifts the equity price directly but has no bond-market counterpart. Because there is no mandate linking the bond leg to the equity leg, the borrowing channel is inactive: changes in household and government debt are absorbed by the wealthy household's bond position without any required revaluation of equity.

2.9. Equity Price

Imposing equity-market clearing $Q_t = \bar{Q}$ into the equity and bond holding equations of the mutual fund (Equation 5) yields the following proposition for the price of equity

PROPOSITION 1 (Price of Equity). *The price of equity in the inelastic model is tightly linked to the bond position of the mutual fund:*

$$p_t = \frac{\theta_t}{1-\theta_t}\frac{B_t^{MF}}{\bar{Q}}. \tag{27}$$

Because the mutual fund absorbs all bond supply, $B_t^{MF} = B_t^P + B_t^M + B_t^G$, and equilibrium equity prices satisfy

$$p_t = \frac{\theta_t}{1-\theta_t}\frac{B_t^P + B_t^M + B_t^G}{\bar{Q}}. \tag{28}$$

In the inelastic model, any fundamental force which changes borrowing (holding equity share constant) directly affects equity prices. As mandate deviation costs $\chi^{MF} \to \infty$ and $\theta_t \to \theta^*$, any increase in borrowing increases equity prices. Therefore, equity and bonds are *complements* in the inelastic model. Moreover, the sensitivity of equity market capitalisation to aggregate debt is $\frac{\theta_t}{1-\theta_t}$. This is the model counterpart of Gabaix and Koijen (2021)'s *price multiplier*. This form for the price of equity admits both backward-looking and forward-looking expressions,[15] as well as leads to a time-varying elasticity of markets,[16] which relates to the time-varying inequality multiplier discussed below.

2.10. The Inequality Multiplier

We now derive the inequality multiplier and identify the mechanisms behind it. The supporting derivations are collected in Appendix E.5. The multiplier is the organising concept of the paper: the amplification we trace in impulse responses, the balance sheet transfer between the top and bottom of the distribution, and

[15]The equity price can be represented as both backward-looking — accounting for cumulative bond issuance over time — and forward-looking — accounting for future changes in dividends and discount rates. A full derivation of these two alternate pricing equations is given in Appendix E.4.

[16]The elasticity of demand for equities is no longer time-invariant: intermediary equity demand responds endogenously to the evolving equity premium. In Appendix E.1 we derive the expression for time-varying Hicksian elasticity.

the widening wedge between income and wealth inequality are all facets of this single object, and we return to it at each step of the results.

The wealthy household's wealth share deviation decomposes into an own-wealth term and a revaluation of the aggregate portfolio:

$$\eta_R^{dev} = \bar{\eta}_R\Big[\widehat{a}^R - \big(\omega_E\,\widehat{p} + \omega_H\,\widehat{p}_h + \omega_G\,\widehat{b}_g\big)\Big], \tag{29}$$

where $\bar{\eta}_R \equiv \bar{a}^R/\bar{\mathcal{W}}$ is the wealthy wealth share, $\bar{a}^R$ is wealthy net worth, and $\omega_E, \omega_H, \omega_G$ are the shares of equity, housing, and government debt in aggregate net worth $\bar{\mathcal{W}} = \bar{p}\bar{Q} + \bar{X}^H + \bar{B}^G$.[17]

The *inequality multiplier* is the response of η_R^{dev} to a fundamental shock ε, which factors into a portfolio heterogeneity term and a price-elasticity term:

$$\frac{\partial\eta_R^{\text{dev}}}{\partial\varepsilon} = \underbrace{\frac{\partial\eta_R^{\text{dev}}}{\partial\widehat{p}}}_{\text{portfolio heterogeneity}} \times \underbrace{\frac{\partial\widehat{p}}{\partial\varepsilon}}_{\text{price elasticity}}. \tag{30}$$

We examine how each term differs across the inelastic and elastic regimes.

2.10.1. Portfolio Heterogeneity Multiplier

Differentiating (29) with respect to $\widehat{p}$, holding the other revaluation terms fixed, yields the portfolio heterogeneity multiplier

$$M \equiv \frac{\partial\eta_R^{\text{dev}}}{\partial\widehat{p}} = \eta_R\big(\beta_R^E - \omega_E\big), \tag{31}$$

where $\beta_R^E \equiv \partial\widehat{a}^R/\partial\widehat{p}$ is the wealthy group's *effective equity beta* and ω_E is the aggregate's equity beta.[18]

PROPOSITION 2. *$M = \eta_R(\beta_R^E - \omega_E)$ is the wealthy's wealth weight times the gap between their equity exposure and the economy's. A common equity price move raises the top share if and only if the wealthy are more equity-exposed than average.*

This portfolio heterogeneity identity is emphasised most directly by Gomez (2025) and is closely related to the heterogeneous-returns literature of Fagereng et al. (2020) and Cioffi (2021)'s heterogeneous aggregate-risk

[17] Hats denote proportional deviations from steady state, $\widehat{z}_t \equiv (z_t - \bar{z})/\bar{z}$; see Appendix E.5.

[18] This is the partial-equilibrium multiplier. In general equilibrium, if house prices co-move with equity ($\partial\widehat{p}_h/\partial\widehat{p} > 0$, the usual boom case), aggregate net worth $\bar{\mathcal{W}}$ appreciates through a component the wealthy do not own, *diluting* their share; M then *overstates* the equilibrium equity-to-inequality pass-through. We discuss this more in Appendix E.7.

exposure model. Nothing yet is model-specific: the two regimes differ *only* in the value of β_R^E. In the elastic model, the effective beta is exactly the wealthy's equity portfolio weight, $\beta_{R,\text{el}}^E = w_R^E$; under the mandate it rises to $\beta_{R,\text{inel}}^E \approx 1$, up to a linear additive correction that vanishes as χ^{MF} increases. Intuitively, to hold the equity share fixed while equity supply $\bar{Q}$ is fixed, the fund's bond leg must move one-for-one (in logs) with the equity price; the bond leg therefore stops being a hedge, and the wealthy's *entire* balance sheet co-moves with $\widehat{p}$ (Appendix E.5).

These betas deliver $M^{\text{el}} = \omega_E^{\text{el}}(1 - \eta_R^{\text{el}})$ and $M^{\text{inel}} = \eta_R^{\text{inel}}(1 - \omega_E^{\text{inel}})$. Assuming a common steady state across regimes, their ratio depends only on $(\bar{\theta}, \eta_R)$:

$$\frac{M^{\text{inel}}}{M^{\text{el}}} = \frac{1 - \bar{\theta}\,\eta_R}{\bar{\theta}\,(1 - \eta_R)}. \tag{32}$$

At the calibrated $\bar{\theta} = 0.78$ and a top 10% wealth share $\eta_R = 0.682$ (Blanchet et al. 2022), inelasticity amplifies the response of wealth shares to equity price changes by 1.9x: the mandate makes the wealthy a levered equity claim, delivering $M^{\text{inel}} > M^{\text{el}}$ for any shock that raises $\widehat{p}$. This ratio depends only on the mandate share $\bar{\theta}$ and the wealth distribution η_R, not on the rigidity χ^{MF}: the rigidity governs whether the market is inelastic, whereas $\bar{\theta}$ and η_R govern how strongly a given price move is levered into the top share. The ratio is strictly decreasing in $\bar{\theta}$ (Appendix E.5), so amplification is largest when the wealthy's nominal equity weight is smallest—there, the mandate converts a large bond position into equity exposure and β_R^E jumps from $\bar{\theta}$ to 1.

The same ratio also depends on the wealth share, in a telling direction.

PROPOSITION 3 (Self-reinforcing amplification). *For $\bar{\theta} < 1$, the ratio (32) is strictly increasing in η_R, with $\partial(M^{\text{inel}}/M^{\text{el}})/\partial\eta_R = (1-\bar{\theta})/[\bar{\theta}(1-\eta_R)^2] > 0$: inelasticity amplifies a given equity price move by more, relative to the elastic benchmark, the more concentrated wealth already is.*

The two regimes redistribute through different margins. Because the wealthy alone hold equity—through the fund, at weight $\bar{\theta}$—aggregate equity exposure is $\omega_E = \bar{\theta}\,\eta_R$, giving $M^{\text{el}} = \bar{\theta}\,\eta_R(1-\eta_R)$ and $M^{\text{inel}} = \eta_R(1-\omega_E)$. The elastic multiplier vanishes as $\eta_R \to 1$: once the wealthy own the market, an equity revaluation lifts everyone in proportion and cannot move shares. The inelastic multiplier does not, because it runs through the wealthy's role as lenders—the mandate turns the bonds that fund others' borrowing into equity exposure—so as inequality rises the elastic channel decays while the inelastic one persists, and the multiplier climbs with the very top share it produces.

Price elasticity. The second factor of (30), the price response $\partial\widehat{p}/\partial\varepsilon$, cannot be signed in general: it depends on which shocks move the price, and that differs across regimes (full treatment in Appendix E.6). In the elastic model the price is a forward-looking present value of fundamentals à la Lucas (1978), so dividend and discount-rate shocks move it first-order; in the inelastic model it tracks aggregate borrowing nearly contemporaneously, so debt and flow shocks do. Inelasticity thus works through a second margin beyond the larger M—it makes borrowing a first-order driver of the equity price—which is what activates the borrowing channel absent from the elastic benchmark.

2.10.2. What does inelasticity add?

Relative to the elastic benchmark, inelasticity adds four elements. *First,* a higher portfolio heterogeneity multiplier: it raises the wealthy's effective equity exposure from their equity share to (nearly) one, converting their bond holdings into a levered equity position. *Second,* a higher price elasticity to debt shocks, so equity prices respond far more to changes in borrowing—connecting to the literature that frames rising wealth inequality as debt-driven and funded by the wealthy. *Third,* and relatedly, a debt–equity cross-equation restriction: factors other than fundamentals (declining r^*, factor shares, risk premia) can move equity prices and account for residual movements in wealth inequality. *Fourth,* a joint implication for both ends of the wealth distribution—the equity boom becomes a balance sheet redistribution, as every equity revaluation is matched by additional borrowing by other households, whose liabilities rise as the wealthy's assets appreciate.[19] In the elastic model, an equity boom revalues holdings but carries no such bundled balance sheet transfer.

3. Calibration

Having set up the model and characterised the equilibrium, we are now in a position to specify functional forms and calibrate parameter values according to data.

3.1. Functional Forms

Following Straub (2019), each household type has CRRA felicity over consumption, with a type-specific elasticity parameter, and the wealthy additionally value bequests.

[19] This is a *balance sheet* transfer—a revaluation of stocks of claims and liabilities—not a cash-flow or welfare transfer at the instant of the shock.

ASSUMPTION 1 (Household-specific utility). *For $i \in \{R, M, P\}$, felicity over consumption is*

$$u_i(c) = \frac{c^{1-\sigma_i}}{1-\sigma_i}, \tag{33}$$

and the wealthy household alone derives utility from bequests,

$$\nu_R(a) = \frac{a^{1-\Sigma_R}}{1-\Sigma_R}, \qquad \Sigma_R < \sigma_R. \tag{34}$$

The wealthy have the highest elasticity parameter σ_R and an operative bequest motive; the middle class and the non-wealthy share a common lower value $\sigma_M = \sigma_P < \sigma_R$ and do not value bequests.

Because elasticities are household-specific, so are the utility functions: the model is not a representative-agent economy with a single curvature. The restriction $\Sigma_R < \sigma_R$ makes bequests a luxury good for the wealthy—the marginal utility of wealth decays more slowly than that of consumption, so their saving rate rises with permanent income and they accumulate disproportionately. The middle class and non-wealthy, having pure CRRA preferences with no bequest term, behave homothetically. Appendix D formalises this non-homotheticity. Labour endowments of all three groups, together with equity and housing dividends, follow AR(1) processes whose steady-state levels, persistences, and innovation variances are specified in Appendix D.

3.2. Holding Costs

The household and intermediary problems feature convex holding costs, written generically as $\Gamma(\cdot)$ in the model and the elastic benchmark of Section 2.8. We give their explicit forms here; all are quadratic. The mutual fund's *mandate* cost χ^{MF} is specified and calibrated with the intermediary in the model and is not repeated.

Wealthy household (inelastic model). The wealthy pay a convex cost of holding mutual fund shares, quadratic in the position X_t^{MF} around its steady-state level $\bar{X}^{MF}$:

$$\Gamma^{MF}\left(X_t^{MF}\right) = \frac{\phi_{MF}}{2\,\bar{X}^{MF}}\left(X_t^{MF}\right)^2. \tag{35}$$

This cost enters the wealthy Euler equation. As discussed below, ϕ_{MF} is not a free parameter: it is pinned residually by that equation given the equity premium.

Housing and non-wealthy borrowing. Both the middle class and the non-wealthy pay a convex cost of holding housing, quadratic in the value of the position $X_t^{H,i} = p_t^H H_t^i$ around its steady-state level $\bar{X}^{H,i}$:

$$\Gamma_H^i\left(X_t^{H,i}\right) = \frac{\phi_H^i}{2\,\bar{X}^{H,i}}\left(X_t^{H,i}\right)^2, \qquad i \in \{M, P\}. \tag{36}$$

These costs make each group's housing demand smoothly downward-sloping and, through the steady-state housing Euler equations, help pin the house price. We set them to modest values, $\phi_H^M = 0.25$ and $\phi_H^P = 0.20$. They are distinct from the mortgage loan-to-value ratio τ_t^M, which governs collateralised borrowing and is calibrated separately. The non-wealthy, who are not collateral constrained, instead face a convex borrowing cost, quadratic in the level of debt B_t^P:

$$\Gamma_B^P\left(B_t^P\right) = \frac{\phi_B^P}{2}\left(B_t^P\right)^2. \tag{37}$$

The parameter ϕ_B^P has no external counterpart and is calibrated internally, jointly with τ_t^M and the equity premium, to match the bottom-50% debt-to-income ratio of 0.5 from Bartscher et al. (2025).

Wealthy household (elastic benchmark). In the elastic direct-holding benchmark of Section 2.8 the mutual fund is removed and the wealthy hold equity $E_t = p_t\bar{Q}$ and bonds B_t^R directly, each subject to a convex holding cost:

$$\Gamma^E(E_t) = \frac{\phi_E}{2\,\bar{E}}(E_t)^2, \qquad \Gamma^B\left(B_t^R\right) = \frac{\phi_B}{2\,\bar{B}^R}\left(B_t^R\right)^2, \tag{38}$$

where ϕ_B is the wealthy household's bond-holding cost in the benchmark, distinct from the non-wealthy borrowing cost ϕ_B^P. To keep the benchmark comparable to the baseline, the two costs are set to reproduce the same steady-state equity premium, satisfying $\phi_E - \phi_B = \beta^R(R_t^E - R_t)$.

These holding costs are not incidental to the mechanism. Following Abadi (2026) and a broader literature on slow-moving capital (Duffie 2010) and gradual portfolio adjustment (Bacchetta et al. 2023), they serve two purposes. First, they prevent the inelasticity of equity demand from being fully arbitraged away: absent a cost of moving the portfolio, the intermediary and the wealthy household would trade away any deviation from the mandate, and the market would collapse back to elastic, present-value pricing. Second, they bound the price-elasticity of household and intermediary positions, keeping asset demands interior and downward-sloping so that positions respond smoothly to the 1989–2023 shock sequence rather than jumping discontinuously to their terminal values. The magnitudes we adopt are chosen with both requirements in mind.

3.3. Parameters

Table 1 collects the parameter values and their sources. It is helpful to group the parameters by how each is disciplined. A first group is fixed externally—either normalised, or taken directly from the microeconomic and asset-pricing literature. A second group is matched to observed data levels: the income, dividend, rental, and government-debt shares. A third group has no direct external counterpart; we solve for these jointly, so that the 1989 reference steady state reproduces a small set of target moments (Table 2).

The model's frequency reflects two distinct tasks, and we match each to its own. The structural relationships we discipline—patience, mortality, the safe rate, and the distributional shares—are inherently annual or lower-frequency objects: a wealth share or a discount rate has no natural monthly counterpart, and manufacturing one would require an arbitrary compounding convention. We therefore calibrate the parameters and the steady-state targets of Table 2 at annual frequency. The historical transition, by contrast, is a high-frequency object: the driver series—income shares, household and government debt, and rents—are available monthly, and the unanticipated-shock exercise of Section 5.2 feeds a per-period sequence of surprises. We therefore solve the 1989–2023 transition on a monthly grid,[20] which uses the data at native resolution and resolves the surprise sequence finely, rather than aggregating the sample to roughly thirty annual points. Calibrating at one frequency while simulating at another is legitimate here because the inequality multiplier prices contemporaneously: through the mandate identity of Proposition 1, $p_t = \frac{\theta_t}{1-\theta_t} B_t^{MF}/\bar{Q}$, the equity price is pinned by the *current* outstanding debt stock rather than by a forward-looking present value, so the map from drivers to wealth shares is period-by-period and does not couple the calibration frequency to the simulation grid. Re-solving on an annual grid confirms this: the inelastic wealth-share and market-moment paths are materially unchanged.[21]

Externally fixed parameters. The mortality rate δ matches the 2022 U.S. crude death rate of 984.1 per 100,000 (CDC National Vital Statistics System). We fix the wealthy and non-wealthy discount rates and recover the middle class rate. The wealthy rate is pinned to the 1989 yield on ten-year Treasuries through a simple benchmark: in a model *without* a bequest motive the steady-state interest rate is exactly $\bar{r} = \rho + \delta$, so we choose $\rho^R = 0.080$ such that $\rho^R + \delta$ equals the 1989 Treasury yield.[22] The non-wealthy are the least patient

[20] The quantitative exercise treats 1989 and 2023 as steady states and computes the perfect-foresight transition between them. Parameters that we allow to vary across the two dates (the income shares, the dividend and rental levels, the government-borrowing share, and the mortgage loan-to-value ratio) are reported at both endpoints; all others are held fixed.

[21] The residual sensitivity to the grid is second-order in the inelastic model, where the equity price is a contemporaneous balance-sheet identity, but first-order in the elastic benchmark, where it is a forward-looking present value of dividends and discount rates. Consistent with this, the annual re-solution shifts the elastic wealth-share paths by noticeably more than the inelastic ones, while leaving the qualitative comparison between the two regimes intact.

[22] The wealthy household in the full model does have a bequest motive, so this is a benchmark used to anchor the *level* of ρ^R rather than an equilibrium restriction. It takes the simplifying assumption of setting the steady-state rate to the Treasury yield via $\bar{r} = \rho + \delta$.

TABLE 1. Calibration of the quantitative model.

Parameter	Description	Value	Source / target
Fixed externally			
δ	Mortality rate	0.0098	CDC NVSS (984 per 100,000, 2022)
ρ^R	Wealthy discount rate	0.080	$\rho^R + \delta$ = 1989 Treasury (no bequest)
ρ^P	Non-wealthy discount rate	0.140	Least-patient household
σ_R	Wealthy income elasticity	4.10	Straub (2019)
$\sigma_M = \sigma_P$	Middle/non-wealthy income elasticity	1.43	Straub (2019) × EIS gradient, Calvet et al. (2021)
Σ_R	Wealthy bequest curvature	2.87	0.699 σ_R, Straub (2019)
θ^*	Fund mandate target	0.78	Avg. of Gabaix and Koijen (2021) and Haddad et al. (2025)
χ^{MF}	Fund mandate-deviation cost	171	Empirical regression (see text)
β^{MF}	Fund discount factor	0.98	Standard
$\bar{Q}, \bar{H}$	Equity / housing supply	10, 10	Normalisation
ϕ_H^M, ϕ_H^P	Housing holding cost (middle, non-wealthy)	0.25, 0.20	Set (see text)
g_D	Long-run dividend growth	0.05	Shiller data (Shiller 2023)
ρ_x	Persistence, all processes	0.90	Standard
$\{s_T^R, s_T^M, s_T^P\}$	Transfer shares	0.01, 0.31, 0.68	Assumption (see text)
Matched to data levels			
$\bar{D}$	Steady-state equity dividend	0.34	Div. = 7.9% of top 10% income, SCF+ (Kuhn et al. 2020)
$\bar{D}^H$	Steady-state housing services	1.69	bottom 90% housing wealth, SCF+ (Kuhn et al. 2020)
$\{\bar{e}^P, \bar{e}^M, \bar{e}^R\}^{1989}$	1989 labour-income shares	0.101, 0.463, 0.435	Blanchet et al. (2022)
$\{\bar{e}^P, \bar{e}^M, \bar{e}^R\}^{2023}$	2023 labour-income shares	0.079, 0.390, 0.530	Blanchet et al. (2022)
$\{\sigma^P, \sigma^M, \sigma^R\}$	Income-process innovations	Δ-share	Blanchet et al. (2022) (1989→2023 change)
$\bar{\alpha}^{1989}, \bar{\alpha}^{2023}$	Government-borrowing share	0.512, 0.609	Fed Z.1 Nonfinancial Debt
Solved to match target moments			
ρ^M	middle class discount rate	0.134	Steady-state rate $\bar{R}$ = 1.095
r_x	Equity premium ($\bar{R}^E - \bar{R}$)	0.012	top 10% wealth share η^R
τ_m	middle class LTV	0.088	middle 40% wealth share η^M; SCF (+57%)
ϕ_B^P	Non-wealthy borrowing cost	0.010	bottom 50% DTI
ϕ_{MF}	Wealthy fund-holding cost	0.146	Wealthy Euler (residual)

group, $\rho^P = 0.140$. The middle class rate ρ^M is then recovered so that the model's steady-state rate matches its 1989 value. That rate is not a free object: in the steady-state construction the middle class block pins it—the house price follows from the non-wealthy housing Euler equation, and the middle class household's housing and bond first-order conditions, with its collateral constraint binding, give

$$(39) \qquad \bar{R} = \frac{1 - \bar{\psi}^M}{\beta^M}, \qquad \bar{\psi}^M = \frac{(1 + \phi_H^M) - \beta^M \bar{G}_h}{\bar{\tau}^M},$$

where $\bar{\psi}^M \geq 0$ is the consumption-normalised collateral multiplier and $\bar{G}_h$ the gross housing return.[23] Since $\bar{R}$ depends on ρ^M through $\beta^M = 1/(1 + \rho^M + \delta)$, we solve for ρ^M to hit $\bar{R} = 1.095$.

We take the wealthy income-elasticity parameter $\sigma_R = 4.1$ and the bequest curvature $\Sigma_R = 0.699\,\sigma_R = 2.87$ directly from Straub (2019), so that bequests are a luxury good ($\Sigma_R < \sigma_R$) and saving rises with permanent income. For the middle and non-wealthy we set a common, lower elasticity $\sigma_M = \sigma_P = 1.43$. We discipline the gap between the wealthy and the rest using the cross-sectional evidence in Calvet et al. (2021), who estimate that the elasticity of intertemporal substitution rises with the wealth–income ratio; the ratio $\sigma_M/\sigma_R \approx 0.35$ corresponds to the decline in estimated substitutability from the top of the wealth distribution to its middle.

[23]The alternative could be pinning a common ρ to the 10-year Treasury yield. The wealthy hold equity through the fund and the non-wealthy borrow subject to a convex cost, so neither pins the risk-free rate; it is the middle class's binding collateral margin, through β^M in (39), that does. Given the other parameters, the non-wealthy debt level and the wealthy holding cost then adjust so that their own Euler equations hold at this $\bar{R}$.

Imposing $\sigma_M = \sigma_P$ is a simplification that leaves the results substantially unchanged.[24]

The fund-mandate target $\theta^* = 0.78$ is the average of the equity-share mandates implied by the aggregate price multipliers of Gabaix and Koijen (2021) and Haddad et al. (2025). The multiplier of 5.3 in Gabaix and Koijen (2021) maps to a target equity share of 0.84, while the more competitive multiplier of 2.5 estimated by Haddad et al. (2025) maps to 0.72; we take their average, $\theta^* = 0.78$. No external estimate exists for the fund's mandate cost χ^{MF}, so we recover it from the mandate rule (Equation 7): taking $\beta^{MF} = 0.98$ and $\theta^* = 0.78$, we regress the realised equity share relative to target, θ_t/θ^*, on the equity premium $R^E_t - R_t$ using quarterly U.S. data, and the slope implies $\chi^{MF} = 171$. The dynamics are, however, insensitive to the precise value of χ^{MF}: because the equity share is nearly fixed at this level of rigidity, the simulated 1989–2023 wealth-share transition is essentially unchanged across its inelastic range (Supplemental Appendix 4)—what matters for our results is that the market is inelastic, not the exact degree of rigidity. Equity and housing supply are normalised to $\bar{Q} = \bar{H} = 10$. The household holding costs $\phi^M_H, \phi^P_H, \phi^P_B$ are specified and calibrated in Section 3.2. All exogenous processes share a persistence of $\rho_x = 0.9$. Finally, we allocate lump-sum transfers across groups in the shares $\{s^R_T, s^M_T, s^P_T\} = \{0.01, 0.31, 0.68\}$, directing the bulk to the non-wealthy; this is an assumption, and setting the shares equal or fully concentrating them on the non-wealthy does not materially affect the dynamics.

Data-matched levels. The steady-state equity dividend $\bar{D}$ is set so that equity dividends equal 7.9% of top-10% income, a share we compute from equity holdings in the SCF+ (Kuhn et al. 2020), imputing dividends at an aggregate dividend yield and expressing the result relative to top-10% income. The housing-service flow $\bar{D}^H$ is the gross imputed rent on owner-occupied housing. Symmetrically with the equity dividend, we normalise it so that bottom-90% housing enters the balance sheet at its empirical size, disciplining the implied housing-wealth-to-income ratio against the SCF+ (Kuhn et al. 2020), in which bottom-90% housing wealth is roughly 1.9 times bottom-90% income in 1989. The model's price–rent multiple is pinned by the non-wealthy's discount rate and holding cost and lies below its empirical counterpart, so it is the housing-*wealth* stock, rather than the rent flow, that we match; the resulting normalisation is conservative, leaving the model's bottom-90% housing wealth below its data value (Table 2). Steady-state equity price assumes an expected long-run dividend growth rate $g_D = 5\%$ (Shiller 2023). Labour-income endowments $\bar{e}^P, \bar{e}^M, \bar{e}^R$ match the 1989 and 2023 income shares of the bottom 50%, middle 40%, and top 10% from Blanchet et al. (2022), net of these asset-income flows; the innovation variances $\sigma^P, \sigma^M, \sigma^R$ are scaled to the observed 1989→2023 change in each group's income share. The government-borrowing share $\bar{\alpha}_t$, defined by $B^G_t = \alpha_t(B^P_t + B^M_t + B^G_t)$,

[24] Calvet et al. (2021) report a size-weighted cross-sectional distribution of the EIS with mean 0.99 and median 0.42, positively correlated with the wealth–income ratio (correlation +0.32). We use this gradient only to discipline the relative magnitude of σ across groups, not its level, which follows Straub (2019).

matches the Federal Reserve Z.1 Nonfinancial Debt series at both endpoints.

Internally calibrated parameters and target moments. The remaining parameters have no direct external counterpart and are solved jointly so that the 1989 reference steady state reproduces the moments in Table 2. The middle class discount rate ρ^M is pinned down by the steady-state gross interest rate $\bar{R} = 1.095$ (the 1989 Treasury yield), subject to the requirement that the middle class collateral constraint bind: the consumption-normalised multiplier is positive, $\bar{\psi}^M = 1 - \beta^M \bar{R} > 0$ (Equation A18), which holds whenever $\bar{r} < \rho^M + \delta$;[25] the steady-state equity premium $r_x \equiv \bar{R}^E - \bar{R}$ (equivalently the price–dividend ratio) is disciplined by the top-10% wealth share $\eta^R = 0.682$; the middle class loan-to-value ratio τ_t^M by the middle-40% wealth share $\eta^M = 0.310$, with its +57% rise over the transition taken from the SCF; and the non-wealthy borrowing cost ϕ_B^P by the bottom-50% debt-to-income ratio of 0.5 from Bartscher et al. (2025) (wealth shares from Blanchet et al. (2022)). Because ρ^M enters the wealth shares through the interest-rate block, the four parameters $\{\rho^M, r_x, \tau_t^M, \phi_B^P\}$ are solved simultaneously for the four targets $\{\bar{R}, \eta^R, \eta^M, \mathrm{DTI}^P\}$; the wealthy fund-holding cost ϕ_{MF} is then pinned residually by the wealthy Euler equation. The price–dividend ratio is not itself a target: it floats, and emerges somewhat above its pre-1989 average of 15 (Shiller 2023).

Table 2 reports the fit. The four moments in the upper panel are targeted by construction. The lower panel reports residually targeted or untargeted moments—the bottom-50% wealth share, the price–dividend ratio, and the bottom-90% housing-wealth-to-income ratio—which the model is not calibrated to match directly but which provide an out-of-sample check. The bottom-90% housing-wealth ratio comes out below its empirical value, reflecting the compressed structural price–rent multiple: the model understates, rather than overstates, the housing wealth of the bottom of the distribution.

TABLE 2. Targeted and untargeted moments, 1989 reference steady state.

Moment	Description	Data	Model	Source
Targeted				
$\bar{R}$	Gross risk-free rate	1.095	1.095	1989 10Y Treasury
η^R	top 10% wealth share	0.682	0.682	Blanchet et al. (2022)
η^M	middle 40% wealth share	0.310	0.310	Blanchet et al. (2022)
DTI^P	bottom 50% debt-to-income	0.5	0.49	Bartscher et al. (2025)
Untargeted				
η^P	bottom 50% wealth share	≈ 0.01	0.008	Blanchet et al. (2022)
$p/\bar{D}$	Price–dividend ratio	≈ 15	17.6	Shiller; floats with η^R
$p_h\bar{H}/Y^{b90}$	bottom 90% housing wealth / income	≈ 1.9	0.79	SCF+ (Kuhn et al. 2020)

The elastic benchmark shares the 1989 steady state. We calibrate the elastic direct-holding benchmark to the identical 1989 reference steady state as the inelastic model. Because the mutual fund nets out at the steady

[25] This is the only regularity restriction the implemented model imposes on ρ^M. If we assume middle class households also value bequests, a more general condition involving a bequest curvature will be operative here.

state, the inelastic allocation, prices, and wealth shares are also an equilibrium of the elastic economy: the wealthy hold the same equity and bonds directly, at holding costs that reproduce the same equity premium (Section 3.2). The two economies therefore begin the 1989–2023 transition from the same wealth distribution and the same asset prices, and differ only in how equity is priced along the path. This gives the elastic benchmark the most favourable footing: any divergence in outcomes over the transition reflects the pricing mechanism alone, not a difference in starting points or in calibration.

Shock series for the quantitative exercise. The main quantitative exercise is a deterministic, perfect-foresight transition: the economy begins at the 1989 steady state and is fed observed data paths for the exogenous processes, converging to the 2023 steady state. Seven of the eight exogenous processes are *active* shocks in this exercise—the three group labour endowments e_t^R, e_t^M, e_t^P, the equity dividend D_t, the housing rent D_t^H, the government-borrowing quantity B_t^G, and the middle class loan-to-value ratio τ_t^M. Each empirical series is expressed as a proportional deviation from its 1989 steady-state value and fed through its AR(1) law of motion (persistence $\rho_x = 0.9$) to recover the innovation sequence that reproduces it. Concretely: the endowment paths are the top-10%, middle-40%, and bottom-50% income shares of Blanchet et al. (2022), net of dividend and (imputed) rental income; the dividend path is cumulated from the Shiller dividend-growth series and the housing-rent path from the growth of owner-occupied rents (FRED, the CPI for Owners' Equivalent Rent of Residences); the government-debt path is built from the government share of nonfinancial debt in the Federal Reserve Z.1 accounts; and the loan-to-value path is a smooth ramp reproducing the +57% rise in middle class mortgage leverage in the SCF. The eighth process, the portfolio demand wedge ζ_t^d, is held at zero throughout the quantitative transition; it enters only the wealthy household's mutual fund Euler equation and is used solely to construct the impulse responses discussed below.

4. The Mechanism in One Shock

For our first quantitative validation exercise, we ask whether the inelastic model, relative to the elastic benchmark, can reproduce the response of top wealth shares to a flow shock—the model counterpart to the residualised equity-flow innovations we construct in the introduction. We model the flow as a positive 10% shock to the latent portfolio demand wedge ζ_t^d, the closest model analogue of a Gabaix and Koijen (2021) flow shock, which decays at $\rho_\zeta = 0.9$; we then trace the impulse responses of the key endogenous variables.[26] The exercise also lays bare the mechanisms that distinguish inelastic from elastic markets.

Both regimes raise the equity price and the top wealth share on impact; what separates them is the *comovement*

[26]The full set of model IRFs, for every shock and endogenous variable, is reported in the Supplemental Appendix.

and the *persistence*. Only the inelastic model raises the equity price *together with* higher aggregate borrowing and a lower safe rate, and only the inelastic model delivers a hump-shaped, highly persistent rise in the top wealth share. In the elastic benchmark the same demand shock *lowers* borrowing and *raises* the safe rate—the opposite of the empirical comovement—and the wealth share response, though larger on impact, reverts monotonically. The inelastic model thus reproduces, if only directionally, the joint rise in equity values, debt, and top wealth shares alongside falling rates documented in the introduction; the elastic model cannot match this configuration of signs.

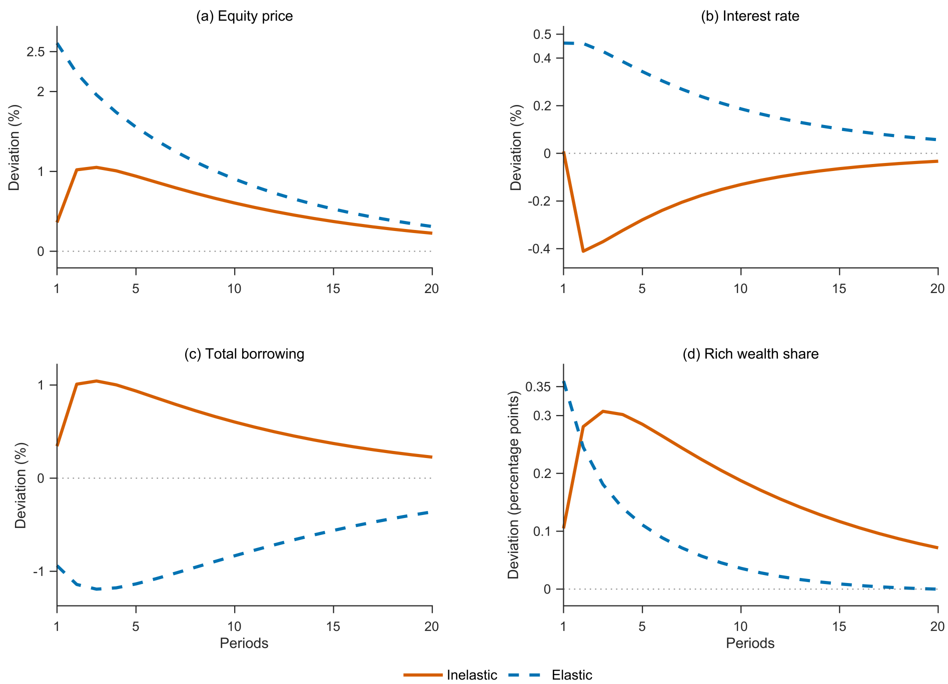

FIGURE 2. Impulse response functions to a positive latent demand shock

Each panel shows the IRF response of one endogenous variable to a positive 10% latent demand shock (ζ_t^d), in the inelastic (solid red) and elastic (dashed blue) models. Each panel, except panel (d), is plotted as a percentage deviation from the steady state. Panel (d), which plots the wealth share of the top 10%, plots the deviation in percentage points from the steady state wealth share. Given the discussion in the calibration section, each period denotes one month, and each panel plots a 20-period response.

In the elastic market, positive latent equity demand induces wealthy households to substitute from bonds into equities. Total credit supply falls, so the safe rate must rise to clear the higher borrowing demanded by non-wealthy households. Because equity supply is fixed, the increased demand raises the equity price, and this revaluation dominates the offsetting effect of the higher discount rate. The top wealth share therefore rises—by more than in the inelastic model on impact—but mean reverts monotonically, returning to steady state by roughly period 17.[27]

[27]This fixed-supply elastic benchmark is deliberately conservative. Our mechanism turns on the borrowing-driven comovement and the persistence of the wealth share response, not on the size of the price impact. With the arbitrageur of Appendix H, part of a demand shock in the elastic market is absorbed on quantities rather than price; fixing equity supply therefore makes the elastic price *more* reactive and stacks the price comparison *against* the inelastic mechanism.

In the inelastic market, the increased allocation to the constrained mutual fund *raises* the supply of credit: a fraction $1 - \theta_t$ of the additional assets under management is routed into bonds by the mandate. Crucially, this response is hump-shaped. Higher mutual fund valuations raise wealthy wealth positions, and non-homothetic bequest utility converts these into further *endogenous* flows into the fund, which are again split between equity and additional credit supply. Every endogenous variable inherits this hump—the signature of a *borrowing-driven* mechanism.

To clear the market at a higher level of debt, the safe rate falls. Through the tight borrowing–price link of Proposition 1, the equity price rises in step with total borrowing.[28] Most important, the top wealth share rises with the same hump-shaped, highly persistent dynamics, remaining clearly above steady state at period 20 (about a fifth of its peak), long after the elastic response has reverted. The gap between the two wealth share paths is the inequality multiplier of Equation (32) in a single shock: the common equity price move is levered into the top share at M^{inel} under the mandate against M^{el} in the elastic benchmark—an amplification of roughly 1.9 at the calibrated 1989 distribution—while its persistence is the multiplier's dynamic signature, the hump the elastic response lacks. The inelastic model therefore generates the persistent response of top wealth shares to flow shocks—together with the empirically relevant rise in borrowing and fall in rates—that motivates the paper.

This exercise qualitatively validates the model's central mechanism: inelastic markets convert a flow into a persistent, borrowing-driven, amplified wealth share response. Whether the same mechanism can *quantitatively* account for the long-run dynamics of U.S. wealth inequality is the question we turn to next.

5. Dynamics of Wealth Inequality

One of the key contributions of Straub (2019), Benhabib et al. (2017), and Mian et al. (2025), among others, is to highlight the importance of incorporating non-homothetic, wealth-dependent preferences in order to match the dynamics of top wealth inequality. Our model extends this approach by adding inelasticity in asset markets. We ask whether inelastic markets improve the model's ability to match the dynamics of wealth inequality relative to a comparable elastic markets economy, and, if so, through which channels. We proceed in four steps, each addressing a limitation of the last: a perfect-foresight simulation, its extension to unanticipated shocks, a decomposition into the equity-investment and borrowing channels, and their unification in the inequality multiplier.

[28] In deviations the mandate price is $\hat{p} = \hat{b}^{MF} + \hat{\theta}/(1 - \bar{\theta})$, which for $\hat{\theta} \approx 0$ (given $\chi^{MF} = 171$) collapses to $\hat{p} \approx \hat{b}^{MF}$.

5.1. Counterfactual Simulations under Perfect Foresight

We begin with deterministic simulations under perfect foresight. We initialise the economy in January 1989 and feed in the shock series constructed in Section 3.3. Agents in January 1989 perfectly anticipate the entire path of exogenous shocks up to March 2023, at which point the economy returns to steady state. We trace the evolution of wealth inequality—measured by the top 10%, middle 40% and bottom 50% wealth shares, all untargeted moments—under both the inelastic and the elastic model, isolating the contribution of inelastic markets to the dynamics of wealth inequality.

We solve the resulting perfect-foresight transition numerically, imposing the March 2023 steady state as the terminal condition.[29] Figure 3 plots empirical wealth shares against simulated wealth shares under two specifications: (i) elastic markets; (ii) inelastic markets.

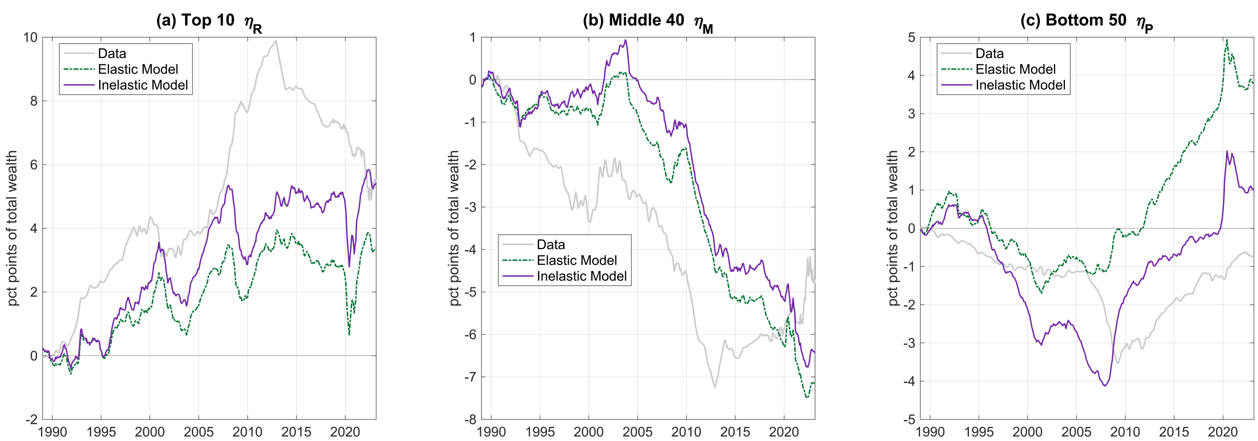

FIGURE 3. Change in wealth shares from 1989: perfect foresight

Note: The perfect foresight simulation initialises the economy at the January 1989 levels of labour income shares and dividends. Then, agents are provided with the full realisation of exogenous processes from January 1989 to March 2023, and the model is solved as a deterministic transition. The figure plots the percentage-point change in wealth shares relative to their 1989 level. The grey line plots the empirical wealth shares, the purple line the inelastic model, and the green line the elastic model. Empirical wealth share data is taken from Realtime Inequality at the household level (Blanchet et al. 2022).

Under perfect foresight, the inelastic model improves the fit of simulated wealth shares to the data, especially for the top 10% and the bottom 50%. The improvement is most pronounced for the bottom 50%, where the inelastic model replicates the negative net wealth shares seen in the data over much of the sample. For the top 10%, inelastic markets deliver a level improvement, though both models fail to capture the run-up in wealth inequality until 2012, which is driven by rapid equity price appreciation that perfect foresight cannot generate. Appendix H.4 shows that introducing an arbitrageur attenuates but does not remove the fit improvement from inelasticity.

[29]The transition is solved under rational expectations with a standard Newton-type perfect-foresight solver: we guess the paths of the endogenous variables, solve forward, and iterate to convergence.

Both models generate the uptrend in equity prices, driven by rising dividends and the increasing income share of the wealthy. The two simulations differ in their borrowing response. Under elastic markets, perfect anticipation of falling labour income eliminates the incentive of the middle 40% and bottom 50% to borrow, which is also why the elastic model cannot reproduce the negative net worth of the bottom 50% in the data. Under inelastic markets, household borrowing demand is similarly muted, but the intermediary's mandate *forces* additional borrowing into the economy, leading bottom 50% households to lever up further.

To make the underlying mechanism transparent, Figure 4 plots the simulated series for household borrowing, interest rates and equity prices against their empirical counterparts over 1989–2023.

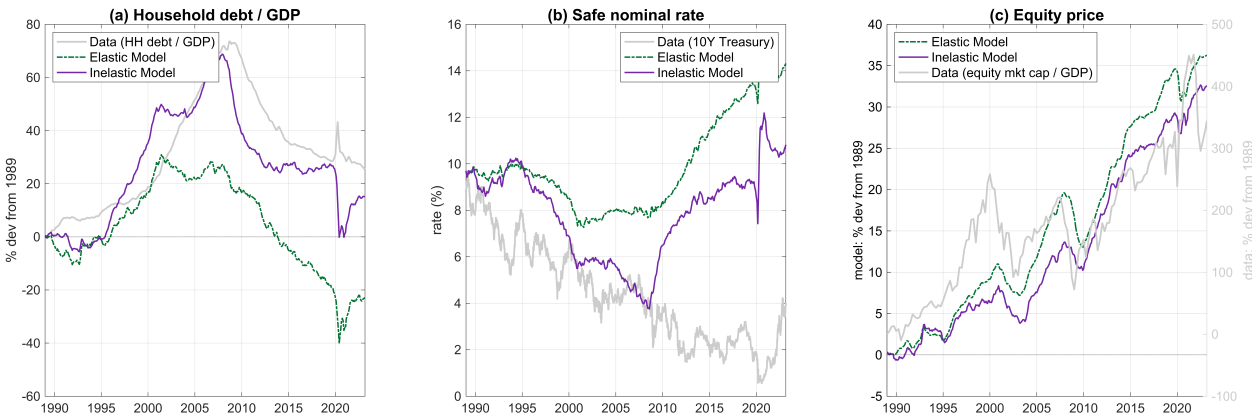

FIGURE 4. Data vs. simulation: borrowing, interest rate and equity price under perfect foresight

Note: Simulated series from the perfect-foresight exercise against their empirical counterparts, 1989–2023. Panel (a) household debt to GDP; panel (b) the safe interest rate; panel (c) the equity price, plotted on a secondary axis because the model underpredicts the level of the price run-up. Data lines in light grey.

Inelasticity improves not only the *fit* to wealth-inequality data but the underlying *mechanism*. The inelastic model closely replicates the path of household debt to GDP over the simulation horizon, whereas the elastic model fails to generate enough borrowing. This borrowing fit is the first ingredient of the multiplier: under the mandate, debt moves the equity price one-for-one with the fund's bond leg (Proposition 1), so the borrowing the inelastic model reproduces is precisely the flow an elastic market would leave unpriced into wealth inequality.

Under inelasticity the safe nominal rate falls more, tracking the data closely until around 2010 but missing the post-crisis 'low-for-long' regime. The turn is driven almost entirely by the post-2008 surge in public debt, which we feed in exogenously and which an inelastic bond market absorbs only at a higher rate: holding government debt at its 1989 level removes the reversal, leaving the safe rate several percentage points lower and still falling as the wealthy's saving continues to bid it down.[30] We therefore read the reversal as an

[30] In the channel decomposition (Figure 6), the equity-investment (wealthy-saving) block lowers the safe rate monotonically over the

artefact of the imposed fiscal path rather than as a prediction of the inelasticity mechanism.

On equity prices, both models fail to capture the fourfold increase in equity values over this period but match the dynamics reasonably well, as expected given the model's stylised structure and its limited sources of aggregate risk. We read the equity price panel as evidence on the co-movement between wealth inequality, borrowing and equity prices, and on the equity-to-borrowing balance sheet transfer, rather than on the level of the price.[31]

Three lessons emerge from the perfect-foresight exercise. Inelasticity improves the joint fit to wealth shares, household borrowing and interest rates; it matches the borrowing *mechanism* rather than the wealth share moments alone; and the amplification it generates—the inequality multiplier—is itself endogenous, rising with market concentration over the sample (we quantify this self-reinforcement in Section 5.4). The model nonetheless understates the sharp rise in the top 10% share between 2008 and 2015 and its partial reversal after 2016. Because perfect foresight suppresses precisely the borrowing that a crisis provokes, we turn next to a simulation that relaxes it.

5.2. Unanticipated Shocks and the Post-Crisis Surge

The perfect-foresight simulation fits the data well through 2008 but understates inequality thereafter because it suppresses households' own borrowing demand. When middle class and non-wealthy households perfectly observe their falling income shares, they meet the decline by lowering consumption rather than by borrowing, so the leverage channel is shut down. Perfect foresight therefore captures the rise in equity prices, and some of the forced leveraging under inelastic markets, but not the endogenous borrowing that amplifies wealth inequality around a crisis.

To address this limitation, we simulate an economy under perfect foresight *with unanticipated shocks*. Under this setup, agents act under perfect foresight conditional on their current information, but may later receive unexpected 'surprise' shocks which force them to recompute their optimal plans going forward based on the new information. This is close to the extended path algorithm proposed by Fair and Taylor (1983) and utilised in Afrouzi et al. (2024), among others. Intuitively, the economy is hit with a sequence of MIT-style surprise shocks, in the sense of Boppart et al. (2018), where the shocks are the per-period innovation in the shock series constructed for the perfect foresight exercise.

sample, while the post-2008 reversal is contributed entirely by the government-debt path.

[31] This division of labour is broadly consistent with Abadi (2026), who shows that in a general-equilibrium demand-based model the demand friction shapes the dynamics and cross-sectional incidence of asset prices while the average level of risk premia is governed by household preferences rather than by the friction itself. Our inelastic and elastic economies share the same calibrated 1989 equity premium, so inelasticity here operates on the dynamics and the balance sheet transfer rather than on the level of the price run-up.

Figure 5 plots the empirical wealth shares and simulated wealth shares. The corresponding market aggregates—borrowing, equity price, and rate— and comparison to the elastic model under unanticipated shocks are reported in Figure A7 and Figure A6.

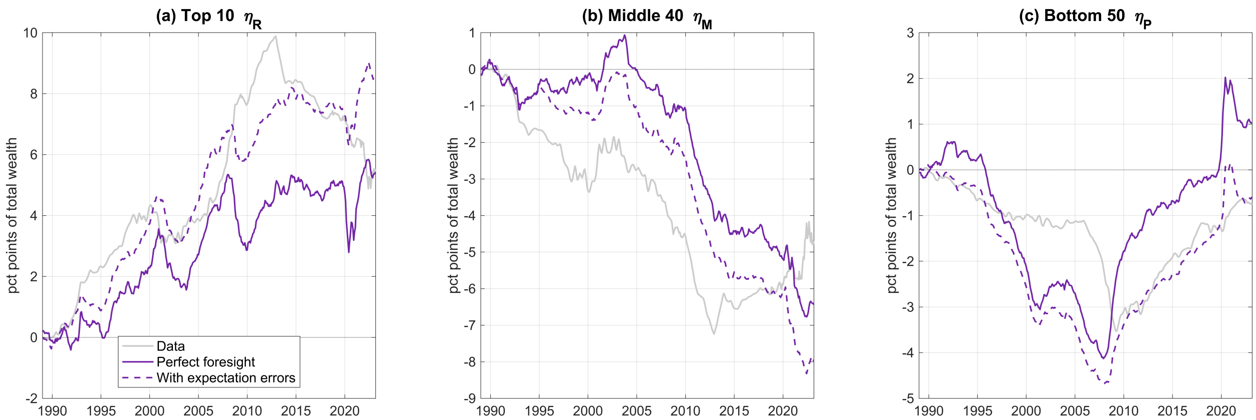

FIGURE 5. Change in wealth shares from 1989: perfect foresight with unanticipated shocks

Note: The perfect foresight with unanticipated shocks simulation initialises the economy at the January 1989 levels of labour income shares and dividends. Before each innovation is realised, agents forecast that the exogenous state will evolve according to its estimated autoregressive law, with future innovations set to their conditional mean of zero. Then, each period from January 1989 to March 2023, the 'surprise' realises in the form of the period-*t* innovation, post which agents update the current state and correctly anticipate the propagation of that innovation through the AR(1) process, but they do not anticipate subsequent innovations. The figure plots percentage-point change in wealth shares relative to their 1989 level. The grey line plots the empirical wealth shares, the solid purple line plots wealth shares from the inelastic model simulated under perfect foresight, the dashed purple line plots wealth shares from the inelastic model simulated under perfect foresight with unanticipated shocks. Empirical wealth share data is taken from Realtime Inequality at the household level (Blanchet et al. 2022).

The model now matches the large post-2008 increase in top wealth inequality and improves the fit of bottom 50% wealth shares over the same period. The reactivated leverage channel is responsible: middle class and non-wealthy households are repeatedly shocked on the downside but expect reversion to trend, and borrow to smooth consumption. This raises borrowing, interest rates and, through the mandate, equity prices, widening wealth inequality. Higher rates on a larger stock of household debt transfer resources from debtors (the non-wealthy and middle class) to creditors (the wealthy), who as net lenders and equity holders capture the gain.

Under unanticipated shocks the elastic markets model reproduces the wealth share dynamics almost as well as the inelastic one, and the market aggregates—borrowing, the safe rate, and the equity price—are close across the two regimes (Figure A7). This reflects observational equivalence under unanticipated shocks rather than a decisive case for elastic markets: the common, sizeable borrowing surge that surprises introduce drives both economies, so the aggregate series no longer separate the pricing regimes. The separation comes instead from the perfect-foresight exercise, where borrowing is muted and the inelastic model alone reproduces the wealth divergence (Figure 3), and from independent evidence on the elasticity of prices to flows.[32]

The post-crisis surge in top wealth is, in this model, a borrowing phenomenon: reintroducing the household leverage that perfect foresight suppresses is what allows the model to match it. This raises a further question—how much of the wealth-inequality path reflects that borrowing, and how much the wealthy's direct accumulation of equity—which we address by decomposing the two channels.

5.3. Equity Investment vs. Borrowing

The unanticipated-shock simulation showed that household leverage is central to the off-trend rise in wealth inequality after the crisis. We now separate the two channels through which inelastic pricing raises wealth inequality—the wealthy's direct accumulation of equity, and the borrowing of other households—and ask how much each contributes, and to which part of the distribution.

To answer this question, we consider the following two counterfactual exercises, which decompose the total increase in wealth inequality into components related to drivers of increased wealthy investment into equity (equity investment) and drivers of increased leveraging (borrowing). To isolate the equity investment channel, we simulate a version of the model (with unanticipated shocks) in which the top 10% income share and equity dividends follow their empirical paths, while all other exogenous processes are fixed at their 1989 steady-state level. To isolate the borrowing channel, we instead let the income shares of the middle 40% and bottom 50%, housing dividends, and government borrowing follow their empirical paths, holding the top 10% income share and equity dividends at their 1989 levels. The secular trend in middle class borrowing capacity, τ_t^M, is held at its 1989 level in *both* counterfactuals; the two channels therefore isolate the income and debt drivers of wealth inequality and need not sum exactly to the total.

Figure 6 plots the full unanticipated-shock simulation from Figure 5 against the equity investment and borrowing counterfactuals.

[32] This near-equivalence is the form the excess volatility puzzle takes in our setting. Stock prices move by more than the present value of subsequent dividends can account for (Shiller 1981; LeRoy and Porter 1981; Grossman and Shiller 1981), and the standard resolution attributes the excess movement to time-varying discount rates (Campbell and Shiller 1988). Inelastic pricing is an observationally close alternative: because the demand curve for equities slopes down, flows move the price directly, with no accompanying news about discount rates. The two accounts can generate the same price path—as the elastic and inelastic simulations here nearly do—and the same safe rate and wealth shares along with it; they differ in the *source* of the price movements, discount-rate news versus flows, which the level series do not reveal. What separates them is the elasticity of the equity price to flows, an object that can be estimated directly and outside the model. The demand-system and identified-flow estimates (Gabaix and Koijen 2021; Haddad et al. 2025) place this elasticity well above the near-zero value that frictionless pricing implies; that is the over-identifying restriction on which our inelastic reading rests. In general equilibrium, Abadi (2026) shows that a demand-based model of this kind reproduces excess volatility through precisely this redistributive amplification—a fundamental shock moves prices, which redistributes wealth toward the intermediaries that hold risky assets and amplifies the initial move, the asset-pricing counterpart of our inequality multiplier—and that the macro price-to-flow elasticity of Gabaix and Koijen (2021) captures only the direct demand response, omitting this redistributive term and so understating the total sensitivity of prices to flows.

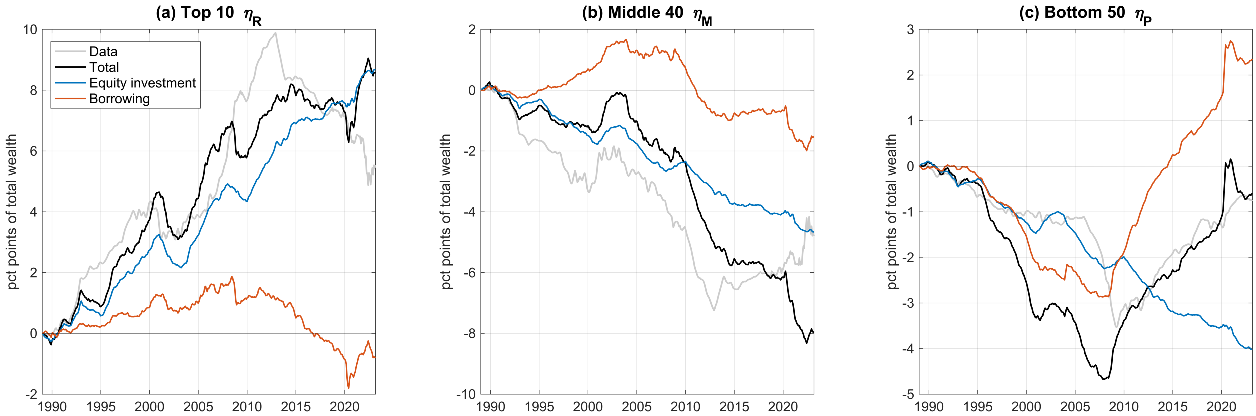

FIGURE 6. Dynamics of wealth inequality: decomposition into channels

Note: The figure plots percentage-point change in wealth shares relative to their 1989 level. The grey line plots the empirical wealth shares, the black line plots the total wealth share response from Figure 5, the blue line plots the total wealth share response in a model with only e_t^R, D_t shocks, and the red line plots the total wealth share response in a model with only $e_t^M, e_t^P, D_t^H, B_t^G$ shocks. The secular middle class borrowing-capacity trend τ_t^M is held at its 1989 level in both counterfactuals, so the blue and red responses need not sum to the black total.

Two features stand out. The first is that the equity investment channel is the dominant force behind the level of wealth inequality, throughout the sample and across the distribution. For the top 10% (panel a), it accounts for almost the entire path of the wealth share—both its secular rise and its acceleration after 2008—so that the blue line and the total move together. For the middle 40% and bottom 50% (panels b and c) the same channel traces a smooth, low-frequency decline: holding little equity, these groups are progressively left behind as the appreciation of wealthy portfolios raises the top share. This mirrors the perfect-foresight simulation of Figure 3: when agents fully anticipate their income paths they have little reason to borrow, and the rise in equity prices comes almost entirely from higher saving by the wealthy, intermediated into an inelastic market. On its own the equity investment channel replicates the low-frequency component of wealth inequality with considerable precision, but misses the large off-trend movements around and after the crisis.

The second feature is that the borrowing channel's distinctive contribution is concentrated at the *bottom* of the distribution and in the post-crisis window. Absent unanticipated shocks it would leave wealth inequality broadly flat; its role is to generate the off-trend swings the investment channel cannot. This is clearest for the bottom 50% (panel c), where borrowing drives the deleveraging through the 2000s and the sharp reversal after 2008, and it holds in milder form for the middle 40% (panel b). For the top 10% the post-2008 acceleration runs through the investment channel; the borrowing channel adds only a modest positive contribution before the crisis and turns slightly negative thereafter, as higher rates on a larger public and household debt stock redistribute towards the wealthy while also raising the wealth of the rest through their bond holdings. Equity investment is the through-line for the top of the distribution, while borrowing leaves its clearest mark at the bottom and in the years around the crisis. Together the two channels reproduce the dynamics of wealth inequality at both low and high frequencies.

The two counterfactuals group shocks by their dominant channel rather than isolate a single transmission mechanism: in general equilibrium each shock moves both prices and borrowing. So construed, they map onto the two families of explanation in the literature and indicate that these operate in tandem: debt-based accounts, in which a build-up of household leverage exacerbates inequality,[33] and investment-based accounts, which attribute inequality to greater capital accumulation by, or higher returns to, the wealthy.[34] From a policymaker's perspective, the appropriate lever depends on the objective—the long-run rise in inequality, its shorter-run implications for financial stability, or both[35].

In sum, the equity investment and borrowing channels are complementary drivers of wealth inequality—one secular and distribution-wide, the other cyclical and concentrated at the bottom—and inelastic pricing is the common margin that translates both sets of flows into the equity price movements that widen the wealth distribution. It is the inequality multiplier that governs the *size* of that translation, and—as we show next—strengthens it as inequality itself rises.

5.4. The Self-Reinforcing Multiplier and the Income–Wealth Wedge

Both channels feed a single object: the inequality multiplier of Section 2. Because the wealthy hold equity through the mandate, every flow—their own saving or the borrowing of others—is levered onto their balance sheet at an effective equity beta near one rather than at their portfolio weight, so a given equity price move raises the top share by M^{inel} rather than M^{el}. At the calibrated mandate share and the 1989 wealth distribution, $M^{\text{inel}}/M^{\text{el}} \approx 1.9$: inelastic markets convert the same fundamental shock into a wealth-inequality response nearly twice as large as an elastic market would.

This amplification is not constant. Figure 7 plots the multiplier along the simulated transition. Because $M^{\text{inel}}/M^{\text{el}} = (1 - \theta_t \eta_{R,t})/\big(\theta_t(1 - \eta_{R,t})\big)$ depends only on the (roughly stable) mandate share and on wealth concentration, it rises with the very top share it produces—from about 1.9 in 1989 to roughly 2.2 by 2023. The multiplier is thus self-reinforcing: a more concentrated economy converts each additional shock into a larger movement in wealth inequality, so inequality begets inequality.

[33] Examples include Rajan (2011), Kumhof et al. (2015), Mian and Sufi (2015), Mian et al. (2025).

[34] Examples include Piketty (2014), Saez and Zucman (2016), De Nardi and Fella (2017), Kuhn et al. (2020).

[35] If the long-run increase in inequality is the concern, it is more crucial for policy to target the investment-based channel, such as through higher asset market participation or taxes on unrealised capital gains. If the short- to medium-run implications of inequality on financial stability is the concern, policy should target borrowing-based channels, such as through leverage restrictions on households, higher loan-to-value ratios, or reduced bond demand from intermediaries. If both effects are a concern, policymakers should look to regulate inelastic financial intermediaries better to alleviate the effects of higher inelasticity, through, for example, ensuring that financial regulation does not impose a high degree of portfolio rigidity.

FIGURE 7. The time-varying inequality multiplier

Note: The inequality multiplier $M_{\text{inel}}/M_{\text{el}} = (1 - \theta_t \eta_{R,t})/(\theta_t(1 - \eta_{R,t}))$, evaluated along the inelastic transition using the simulated paths of the mandate share θ_t and the top 10% wealth share $\eta_{R,t}$. The solid purple line (left axis) plots the multiplier; the dashed grey line (right axis) plots the top 10% wealth share. The multiplier rises with wealth concentration, so inelastic amplification is self-reinforcing.

This self-reinforcement is the mechanism behind the growing wedge between income and wealth inequality that motivates the paper. A rise in the top income share lifts the wealthy's saving and, through the equity investment channel, the equity price; the multiplier converts that price move into a *larger* rise in the top wealth share—and by more as concentration grows. A steadily but only mildly rising income concentration is thereby amplified into a steeper, self-perpetuating rise in wealth concentration, driving a wedge between the two. This is the general-equilibrium *redistributive* effect of Abadi (2026); whereas it dissipates in his representative-household setting as intermediaries' wealth shares revert, here it is durable, because the income inequality that drives concentration is persistent.

6. Conclusion

This paper studies how frictions in financial intermediation transmit income inequality into wealth inequality. When intermediaries maintain stable allocations between equity and debt, household saving and borrowing flows affect equity prices rather than being absorbed entirely through quantities. Greater saving by wealthy households and greater borrowing elsewhere in the distribution both expand intermediary balance sheets and raise equity valuations. Because equity ownership is concentrated, these adjustments redistribute wealth toward the top—an amplification mechanism we call the *inequality multiplier*.

In the data, the component of equity returns predicted by mutual fund and ETF flows generates persistent changes in wealth shares, whereas the response to the remaining component mean reverts. In the model, inelastic intermediation roughly doubles the response of wealth inequality relative to an elastic market benchmark. This amplification increases with wealth concentration, so the multiplier is self-reinforcing: a more concentrated distribution converts each subsequent shock into a larger movement in inequality. Because the income inequality that drives concentration is itself persistent, wealth concentration need not mean revert over empirically relevant horizons. When fed the observed 1989–2023 shocks, the model captures the broad comovement among equity prices, borrowing, interest rates, and wealth shares, although it underpredicts the level of the equity-price run-up and misses the post-2010 decline in the safe rate. The equity investment channel accounts primarily for the secular evolution of wealth inequality, while the borrowing channel is most important at the bottom of the distribution and around the financial crisis. Together, they connect the long-run divergence between income and wealth inequality to shorter-run movements in leverage and asset prices. These results are robust to an unconstrained arbitrageur sector of realistically calibrated size.

The results also identify a policy trade-off. Broadening equity participation may reduce the unequal incidence of asset-price appreciation, but participation through passive or mandate-constrained institutions may further reduce market elasticity. Likewise, policies affecting household or government debt can influence wealth inequality not only through interest rates and borrower balance sheets, but also through equity revaluation. Policy should therefore consider both who owns financial assets and how the institutions holding them transmit flows into prices, which requires carefully understanding heterogeneous institutional constraints.

Several questions remain open. Future work will focus on solving for optimal policy in the presence of intermediary-driven increases in wealth inequality. Further work should also identify flow-driven price movements more sharply, trace debt through the intermediation network, and allow for heterogeneous mandates, endogenous arbitrage entry, heterogeneous arbitrage capacity, and elastic equity supply. Endogenising fiscal and monetary policy would permit a joint analysis of asset supply, interest rates, and redistribution, while cross-country evidence could reveal how the multiplier varies with financial structure. The broader implication is that wealth inequality depends not only on household saving and portfolio exposure, but also on how the financial system converts distributional flows into asset prices.

References

Abadi, Joseph. 2026. "Demand-Based Asset Pricing in General Equilibrium." Working Paper 26-12, Federal Reserve Bank of Philadelphia. https://www.philadelphiafed.org/-/media/FRBP/Assets/working-papers/2026/wp26-12.pdf.

Afrouzi, Hassan, Andres Blanco, Andres Drenik, and Erik Hurst. 2024. "A Theory of How Workers Keep Up with Inflation.", National Bureau of Economic Research.

Auclert, Adrien, Hannes Malmberg, Matthew Rognlie, and Ludwig Straub. 2025. "The Race Between Asset Supply and Asset Demand." Working Paper 34470, National Bureau of Economic Research. https://www.nber.org/papers/w34470.

Bacchetta, Philippe, Eric Van Wincoop, and Eric R Young. 2023. "Infrequent random portfolio decisions in an open economy model." *The Review of Economic Studies* 90 (3): 1125–1154.

Bartscher, Alina K., Moritz Kuhn, Moritz Schularick, and Ulrike I. Steins. 2025. "The distribution of household debt in the United States, 1950-2022." *Review of Economic Dynamics* 57: 101288. https://doi.org/10.1016/j.red.2025.101288.

Benhabib, Jess, and Alberto Bisin. 2018. "Skewed Wealth Distributions: Theory and Empirics." *Journal of Economic Literature* 56 (4): 1261–1291. https://ideas.repec.org/a/aea/jeclit/v56y2018i4p1261-91.html.

Benhabib, Jess, Alberto Bisin, and Mi Luo. 2017. "Earnings inequality and other determinants of wealth inequality." *American Economic Review* 107 (5): 593–597.

Blanchard, Olivier J. 1985. "Debt, deficits, and finite horizons." *Journal of Political Economy* 93 (2): 223–247.

Blanchet, Thomas, Emmanuel Saez, and Gabriel Zucman. 2022. "Real-time inequality.", National Bureau of Economic Research.

Boppart, Timo, Per Krusell, and Kurt Mitman. 2018. "Exploiting MIT shocks in heterogeneous-agent economies: The impulse response as a numerical derivative." *Journal of Economic Dynamics and Control* 89: 68–92.

Calvet, Laurent E., John Y. Campbell, Francisco Gomes, and Paolo Sodini. 2021. "The Cross-Section of Household Preferences." Working Paper 28788, National Bureau of Economic Research. http://www.nber.org/papers/w28788.

Calvo, Guillermo A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12 (3): 383–398.

Campbell, John Y. 2006. "Household Finance." *The Journal of Finance* 61 (4): 1553–1604. https://EconPapers.repec.org/RePEc:hrv:faseco:3157877.

Campbell, John Y., and Robert J. Shiller. 1988. "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors." *The Review of Financial Studies* 1 (3): 195–228.

Catherine, Sylvain, Max Miller, James D Paron, and Natasha Sarin. 2025. "Who gains when interest rates fall.", unpublished working paper.

Chien, YiLi, Harold Cole, and Hanno Lustig. 2011. "A multiplier approach to understanding the macro implications of household finance." *The Review of Economic Studies* 78 (1): 199–234.

Chinco, Alex, and Marco Sammon. 2024. "The passive ownership share is double what you think it is." *Journal of Financial Economics* 157 (C). https://ideas.repec.org/a/eee/jfinec/v157y2024ics0304405x24000837.html.

Cioffi, Riccardo A. 2021. "Heterogeneous Risk Exposure and the Dynamics of Wealth Inequality." *Working paper*.

Coval, Joshua, and Erik Stafford. 2007. "Asset fire sales (and purchases) in equity markets." *Journal of Financial Economics* 86 (2): 479–512.

De Nardi, Mariacristina. 2004. "Wealth Inequality and Intergenerational Links." *The Review of Economic Studies* 71 (3): 743–768. https://EconPapers.repec.org/RePEc:oup:restud:v:71:y:2004:i:3:p:743-768.

De Nardi, Mariacristina, and Giulio Fella. 2017. "Saving and Wealth Inequality." *Review of Economic Dynamics* 26: 280–300. https://ideas.repec.org/a/red/issued/16-340.html.

Ding, Xuning, and Zhengyang Jiang. 2026. "Wealth Inequality and Safe Asset Demand." Working Paper 35393, National Bureau of Economic Research. https://www.nber.org/papers/w35393.

Duffie, Darrell. 2010. "Presidential address: Asset price dynamics with slow-moving capital." *The Journal of Finance* 65 (4): 1237–1267.

Fagereng, Andreas, Matthieu Gomez, Émilien Gouin-Bonenfant, Martin Holm, Benjamin Moll, and Gisle Natvik. 2025. "Asset-Price Redistribution." *Journal of Political Economy* 133 (11): 3494–3549. https://www.journals.uchicago.edu/doi/10.1086/736769.

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

Fair, Ray, and John Taylor. 1983. "Solution and Maximum Likelihood Estimation of Dynamic Nonlinear Rational Expectations Models." *Econometrica* 51 (4): 1169–85. https://EconPapers.repec.org/RePEc:ecm:emetrp:v:51:y:1983:i:4:p:1169-85.

Farhi, Emmanuel, and Iván Werning. 2019. "Monetary Policy, Bounded Rationality, and Incomplete Markets." *American Economic Review* 109 (11): 3887–3928.

Fernández-Villaverde, Jesús, Samuel Hurtado, and Galo Nuno. 2023. "Financial frictions and the wealth distribution." *Econometrica* 91 (3): 869–901.

Gabaix, Xavier, and Ralph S. J. Koijen. 2021. "In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis." *Working paper*. https://doi.org/10.3386/w28967.

Gabaix, Xavier, Ralph S. J. Koijen, Federico Mainardi, Sangmin Oh, and Motohiro Yogo. 2025. "Limited Risk Transfer Between Investors: A New Benchmark for Macro-Finance Models." *NBER Working Paper No. 33336*. https://doi.org/10.3386/w33336.

Gabaix, Xavier, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. 2016. "The Dynamics of Inequality." *Econometrica* 84 (6). https://doi.org/10.3982/ECTA13569.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus. 2021. "Five facts about beliefs and portfolios." *American Economic Review* 111 (5): 1481–1522.

Göcmen, Ararat, Clara Martínez-Toledano, and Vrinda Mittal. 2025. "Private Capital Markets and Inequality." https://ssrn.com/abstract=5166981. Working paper; available at SSRN.

Gomez, Matthieu. 2023. "Decomposing the growth of top wealth shares." *Econometrica* 91 (3): 979–1024.

Gomez, Matthieu. 2025. "Wealth Inequality and Asset Prices." *The Review of Economic Studies* 92 (6): 3924–3967.

Gopalakrishna, Goutham, Zhouzhou Gu, and Jonathan Payne. 2026. "Institutional Asset Pricing with Segmentation and Household Heterogeneity." working paper, Princeton University. https://jepayne.github.io/files/GGP_AssetPricingInequality.pdf. Previously circulated as "Asset Pricing, Participation Constraints, and Inequality".

Greenwald, Daniel L, Matteo Leombroni, Hanno Lustig, and Stijn Van Nieuwerburgh. 2021. "Financial and total wealth inequality with declining interest rates.", National Bureau of Economic Research.

Grossman, Sanford, and Robert Shiller. 1981. "The Determinants of the Variability of Stock Market Prices." *American Economic Review* 71 (2): 222–27. https://EconPapers.repec.org/RePEc:aea:aecrev:v:71:y:1981:i:2:p:222-27.

Haddad, Valentin, Paul Huebner, and Erik Loualiche. 2025. "How Competitive Is the Stock Market? Theory, Evidence from Portfolios, and Implications for the Rise of Passive Investing." *American Economic Review* 115 (3): 975–1018. https://www.aeaweb.org/articles?id=10.1257/aer.20230505.

Haddad, Valentin, and Tyler Muir. 2025. "Market Macrostructure: Institutions and Asset Prices.", National Bureau of Economic Research.

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

Hui, Xitong. 2025. "Asset Prices, Wealth Inequality, and Welfare: Safe Assets as a Solution." ECB Working Paper 2025/3162, European Central Bank. https://www.ecb.europa.eu/pub/pdf/scpwps/ecb.wp3162~339eaf8c4a.en.pdf.

Jordà, Òscar. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 95 (1): 161–182.

Justiniano, Alejandro, Giorgio E Primiceri, and Andrea Tambalotti. 2019. "Credit supply and the housing boom." *Journal of Political Economy* 127 (3): 1317–1350.

Kiyotaki, Nobuhiro, and John Moore. 1997. "Credit cycles." *Journal of Political Economy* 105 (2): 211–248.

Koijen, Ralph S J, Robert J Richmond, and Motohiro Yogo. 2024. "Which Investors Matter for Equity Valuations and Expected Returns?." *The Review of Economic Studies* 91 (4): 2387–2424. https://ideas.repec.org/a/oup/restud/v91y2024i4p2387-2424..html.

Koijen, Ralph S. J., and Motohiro Yogo. 2019. "A Demand System Approach to Asset Pricing." *Journal of Political Economy* 127 (4): 1475–1515.

Kuhn, Moritz, Moritz Schularick, and Ulrike I. Steins. 2020. "Income and Wealth Inequality in America, 1949–2016." *Journal of Political Economy* 128: 3469–3519. https://doi.org/10.1086/708815.

Kumhof, Michael, Romain Rancière, and Pablo Winant. 2015. "Inequality, leverage, and crises." *American Economic Review* 105 (3): 1217–1245.

LeRoy, Stephen F., and Richard D. Porter. 1981. "The Present-Value Relation: Tests Based on Implied Variance Bounds." *Econometrica* 49 (3): 555–574.

Lucas, Robert. 1978. "Asset Prices in an Exchange Economy." *Econometrica* 46 (6): 1429–45. https://EconPapers.repec.org/RePEc:ecm:emetrp:v:46:y:1978:i:6:p:1429-45.

Melcangi, Davide, and Vincent Sterk. 2025. "Stock market participation, inequality, and monetary policy." *Review of Economic Studies* 92 (4): 2656–2690.

Mian, Atif R, Ludwig Straub, and Amir Sufi. 2025. "The Saving Glut of the Rich." Working Paper 26941, National Bureau of Economic Research. http://www.nber.org/papers/w26941.

Mian, Atif, Ludwig Straub, and Amir Sufi. 2021. "Indebted demand." *The Quarterly Journal of Economics* 136 (4): 2243–2307.

Mian, Atif, and Amir Sufi. 2015. *House of debt: How they (and you) caused the Great Recession, and how we can prevent it from happening again.*: University of Chicago Press.

Piketty, Thomas. 2014. *Capital in the twenty-first century.*: Harvard University Press.

Piketty, Thomas, and Gabriel Zucman. 2014. "Capital is back: Wealth-income ratios in rich countries 1700–2010." *The Quarterly journal of economics* 129 (3): 1255–1310.

Rajan, Raghuram G. 2011. *Fault Lines: How Hidden Fractures Still Threaten the World Economy.*: Princeton University Press.

Rotemberg, Julio J. 1982. "Sticky Prices in the United States." *Journal of Political Economy* 90 (6): 1187–1211.

Saez, Emmanuel, and Gabriel Zucman. 2016. "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data." *The Quarterly Journal of Economics* 131 (2): 519–578. https://doi.org/10.1093/qje/qjw004.

Sammon, Marco, and John J Shim. 2026. "Who Clears the Market When Passive Investors Trade?" *The Review of Financial Studies*. http://dx.doi.org/10.1093/rfs/hhag032.

Shiller, Robert J. 1981. "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?" *American Economic Review* 71 (3): 421–436.

Shiller, Robert J. 2023. "Online Data: Stock Market Data Used in Irrational Exuberance." http://www.econ.yale.edu/~shiller/data.htm. Accessed: 2025-08-26.

Smith, Matthew, Owen Zidar, and Eric Zwick. 2023. "Top wealth in America: New estimates under heterogeneous returns." *The Quarterly Journal of Economics* 138 (1): 515–573.

Straub, Ludwig. 2019. "Consumption, Savings, and the Distribution of Permanent Income." *Working paper*. https://scholar.harvard.edu/files/straub/files/cons_ineq_rates.pdf.

Toda, Alexis Akira, and Kieran James Walsh. 2020. "The equity premium and the one percent." *The Review of Financial Studies* 33 (8): 3583–3623.

Vissing-Jørgensen, Annette. 2002. "Limited asset market participation and the elasticity of intertemporal substitution." *Journal of Political Economy* 110 (4): 825–853.

Welch, Ivo, and Amit Goyal. 2008. "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction." *The Review of Financial Studies* 21 (4): 1455–1508.

Yaari, Menahem E. 1965. "Uncertain lifetime, life insurance, and the theory of the consumer." *The Review of Economic Studies* 32 (2): 137–150.

Appendix

Appendix A. Motivating Facts

Two sets of empirical facts motivate our analysis: the heterogeneity of household portfolios across the wealth distribution, and the coincident aggregate trends in asset values, debt, and inequality since 1989. We describe each in turn.

A.1. Portfolio Heterogeneity

Ex-ante, why should equity market inelasticity matter for wealth inequality? The link is portfolio heterogeneity. The composition of household portfolios (including the amount of borrowing) differs systematically along the wealth distribution, with wealthy (top 10%) household portfolios dominated by equity (Figure A1C), and non-wealthy (bottom 50%) household portfolios by borrowing (Figure A1A). Due to this portfolio heterogeneity, equity price increases (induced by higher inelasticity) cause large increases in the share of wealth held by wealthy households.

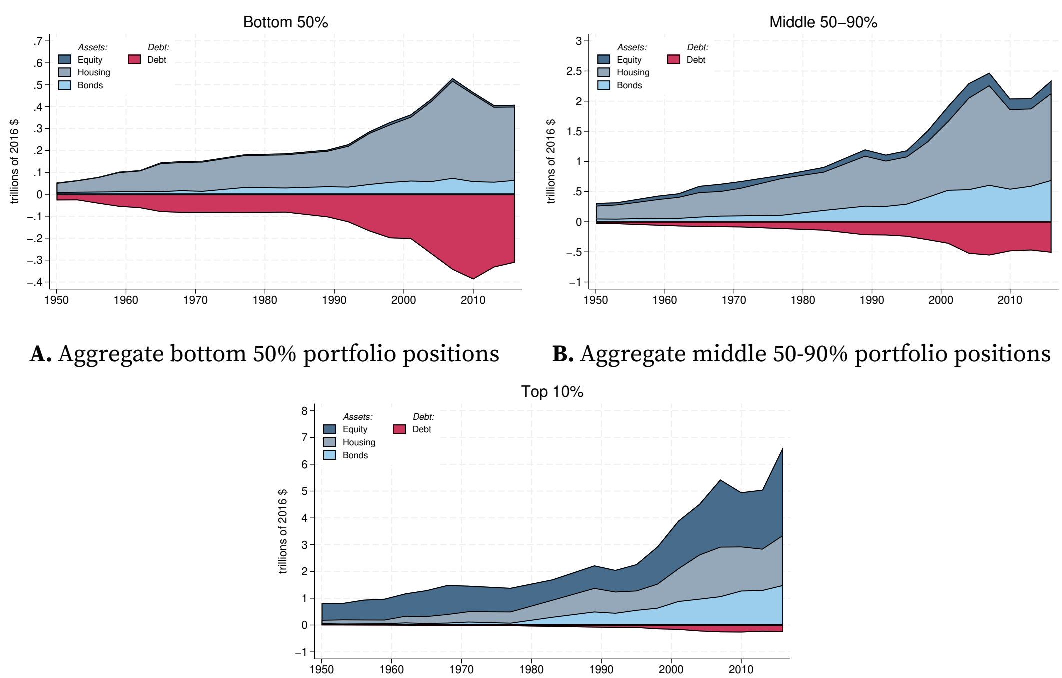

A. Aggregate bottom 50% portfolio positions

B. Aggregate middle 50-90% portfolio positions

C. Aggregate top 10% portfolio positions

FIGURE A1. Panel A shows bottom 50% households' aggregate portfolio holdings, with large increases in indebtedness. Panel B shows aggregate portfolio holdings of the middle 50-90% — housing is *the* asset of the middle class. Panel C shows aggregate portfolio holdings of the top 10%, dominated by equity holdings.

Source: Authors' calculation using Kuhn, Schularick and Steins, JPE 2020 and FRED. All portfolio components and wealth levels are shown in trillions of dollars (2016 dollars).

A.2. Aggregate Trends

A second set of facts motivates our focus on market inelasticity. Four aggregate trends have unfolded together in the United States since 1989, summarised in Figure A2: equity market capitalisation and aggregate debt have both risen sharply relative to GDP; the wealth gap between the top 10% and the bottom 50% has widened far faster than the corresponding income gap; and passive investors' share of the equity market has grown to roughly one-third. No single existing mechanism jointly accounts for these facts; the inequality multiplier links them.

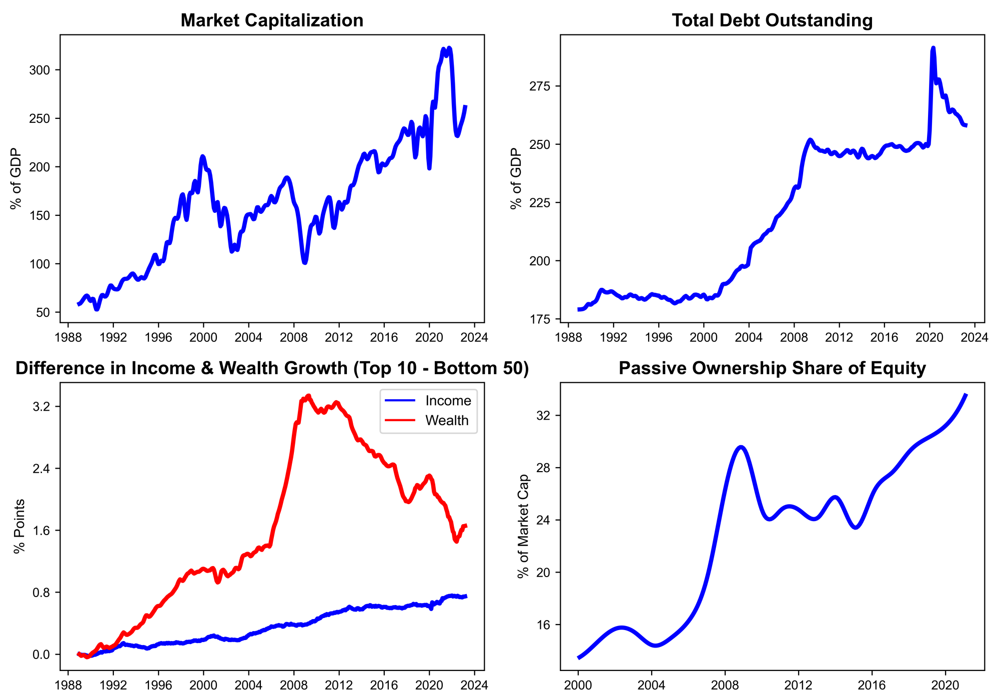

FIGURE A2. Four coincident trends since 1989: the rise in equity market capitalisation and in aggregate debt (both relative to GDP), the widening gap between wealth and income inequality, and the growth of passive ownership of U.S. equities.

Note: *Market Capitalization* and *Total Debt Outstanding* are expressed as a percentage of GDP; total debt includes federal, household, and business debt. In the lower-left panel, the increase in wealth (income) inequality is the difference between the growth rate of average wealth (average factor income) of the top 10th percentile and that of the bottom 50th percentile, with January 1989 as the base year. *Passive Ownership Share* is the passive share of equity-market capitalisation.

Source: Corporate-equity market capitalisation and debt from the Financial Accounts of the United States (Z.1 Flow of Funds); GDP from FRED. Top-10% and bottom-50% income and wealth growth from Blanchet et al. (2022). Passive ownership share from Chinco and Sammon (2024).

Appendix B. A Stylised Three-Period Model

This appendix presents the simplest version of the mechanism in Section 2: two household types, two assets, and a mandate-constrained intermediary, over three dates. It strips out housing, the middle class, non-homothetic preferences, and the borrowing frictions of the quantitative model, retaining only the portfolio mandate that welds the equity price to the outstanding stock of debt. The resulting pricing identity, $p_0 = p_{-1} + \frac{\theta}{1-\theta}\frac{B_0^P}{\bar{Q}}$, is the exact three-date analogue of the quantitative mandate identity $p_t\bar{Q} = \frac{\theta_t}{1-\theta_t}B_t^{MF}$.

The wealthy household (R) saves only through the intermediary, which allocates its wealth between equity and debt; the non-wealthy household (P) receives labour income and can only borrow, being institutionally excluded from equity. The intermediary (a mutual fund) is required to hold a fixed proportion of equity and bonds, which is the sole friction. Figure A3 illustrates the setup.

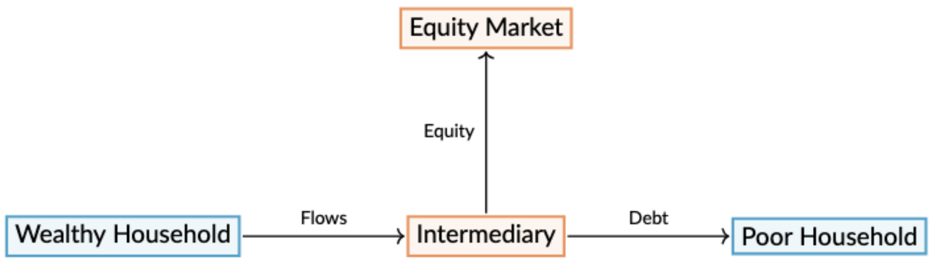

FIGURE A3. Environment

Timing. At $t = -1$ the intermediary allocates its initial wealth between equity and bonds to maximise $t = 1$ returns, unconstrained by any mandate; prices are set in a no-trade equilibrium, and the household positions are inherited from it—the wealthy own the intermediary's portfolio, the non-wealthy hold legacy debt equal to its bond leg. At $t = 0$ both households solve a two-period consumption–saving problem: the wealthy choose flows F_0 into the intermediary, the non-wealthy choose borrowing B_0^P, and—now bound by its mandate—the intermediary cannot rebalance, so the fixed equity supply $\bar{Q}$ clears through the price p_0. A negatively correlated endowment shock at $t = 0$ generates the initial wealth disparity, and the resulting price movements revalue both portfolios. At $t = 1$ the intermediary unwinds and the non-wealthy repay. Figure A4 summarises.

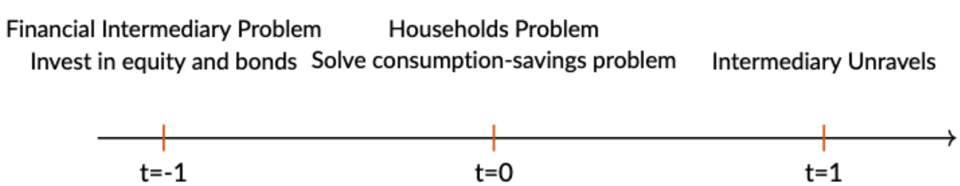

FIGURE A4. Timeline

B.1. Financial Intermediary

At $t = -1$, given initial wealth W_{-1}, dividend D_1, and the two-period annualised rate $r_{f,-1}$, the intermediary chooses its equity share θ to maximise returns on final-period wealth:

$$\max_{\theta} \frac{1}{(1+r_{f,-1})^2}\left[\theta \frac{D_1}{p_{-1}} + (1-\theta)(1+r_{f,-1})^2\right]. \tag{A1}$$

The first-order condition yields the fundamental price at $t = -1$:

$$p_{-1} = \frac{D_1}{(1+r_{f,-1})^2}. \tag{A2}$$

Here θ is the aggregate, equity-weighted share of a continuum of intermediaries $i = 1, \ldots, I$ with heterogeneous shares θ_i; with $S_i \equiv pQ_i / \sum_j pQ_j$, following Gabaix and Koijen (2021), $\theta \equiv \sum_i S_i \theta_i$. At $t = -1$ the aggregate accounting identities are

$$\theta W_{-1} = p_{-1} Q_{-1}, \qquad (1-\theta) W_{-1} = B_{-1}. \tag{A3}$$

At $t = 0$ the intermediary receives inflow F_0 and the price adjusts to p_0, but it cannot readjust θ. Its holdings satisfy

$$\theta\left(p_0 Q_{-1} + B_{-1} + F_0\right) = p_0 Q_0(F_0), \qquad (1-\theta)\left(p_0 Q_{-1} + B_{-1} + F_0\right) = B_0(F_0), \tag{A4}$$

where $Q_0(F_0)$ and $B_0(F_0)$ are the new equity and bond quantities. Intermediary wealth is

$$W_0 = p_0 Q_{-1} + B_{-1} + F_0 = p_0 Q_0 + B_0, \tag{A5}$$

so flows are absorbed as new equity plus new bonds:

$$F_0 = p_0(Q_0 - Q_{-1}) + (B_0 - B_{-1}). \tag{A6}$$

The $t = 0$ portfolio rigidity captures institutional mandates or costly rebalancing; potential microfoundations are reviewed in Supplemental Appendix 2, and Appendix H shows that adding a realistically-calibrated arbitrageur does not materially change the results.

B.2. Households

Non-Wealthy Household. The non-wealthy household chooses $\{c_0^P, c_1^P, B_0^P\}$ to solve max $u[c_0^P] + \beta u[c_1^P]$, given endowments $\{e_0^P, e_1^P\}$ and legacy debt $(1 + r_{f,-1})^2 B_{-1}$,[36] subject to

$$c_0^P \leq e_0^P + B_0^P, \qquad c_1^P + (1 + r_{f,-1})^2 B_{-1} + (1 + r_{f,0}) B_0^P \leq e_1^P. \tag{A7}$$

Its Euler equation prices the risk-free bond.

Wealthy Household.. The wealthy household chooses $\{c_0^R, c_1^R, F_0\}$ to solve max $u[c_0^R] + \beta u[c_1^R]$, given $\{e_0^R, e_1^R\}$ and end-of-period income $D_1 Q_0 + (1 + \tilde{r}_{f,0}) B_0$, subject to

$$c_0^R + F_0 \leq e_0^R, \qquad c_1^R \leq e_1^R + D_1 Q_0(F_0) + (1 + \tilde{r}_{f,0}) B_0(F_0), \tag{A8}$$

where $1 + \tilde{r}_{f,0} = (1 + r_{f,-1})^2 B_{-1}/B_0 + (1 + r_{f,0}) B_0^P / B_0$. Its Euler equation prices returns on the intermediary, not equity directly. Equity clears at fixed supply, $Q_{-1} = Q_0 = \bar{Q}$.

B.3. Equilibrium

DEFINITION A1. *An equilibrium consists of choices* $\{c_0^R, c_1^R, F_0, c_0^P, c_1^P, B_0^P, \theta\}$, *quantities* $\{Q_0, B_0\}$, *prices* $\{p_{-1}, p_0, r_{f,-1}, r_{f,0}\}$, *and endowments* $\{Q_{-1}, B_{-1}, W_{-1}, e_0^R, e_1^R, e_0^P, e_1^P, D_1\}$ *such that households and the intermediary optimise and all markets clear:* $Q_{-1} = Q_0 = \bar{Q}$; $B_0 = B_0^P + B_{-1}$; *and* $c_0^R + c_0^P = e_0^R + e_0^P$, $c_1^R + c_1^P = e_1^R + e_1^P + D_1 Q_0$.

B.4. Equilibrium Pricing

Imposing equity-market clearing yields the central pricing result. Because flows from the wealthy and borrowing by the non-wealthy are two sides of the same transaction, it is stated directly in terms of the bond position.

PROPOSITION A1. *The price of equity at* $t = 0$ *is*

$$p_0 = p_{-1} + \frac{\theta}{1 - \theta} \frac{B_0^P}{\bar{Q}}, \tag{A9}$$

where p_{-1} *is the fundamental price and* $\frac{\theta}{1-\theta} \frac{B_0^P}{\bar{Q}}$ *is the non-fundamental component, tied one-for-one to the outstanding debt.*

[36] Without loss of generality, this is two-period debt with no intermediate coupon.

PROOF. Equity-market clearing requires $(Q_0 - Q_{-1})/Q_{-1} = 0$. Substituting $Q_0 = \frac{\theta}{p_0}(p_0 Q_{-1} + B_{-1} + F_0)$ from (A4),

$$\theta + \frac{\theta B_{-1}}{p_0 Q_{-1}} + \frac{\theta F_0}{p_0 Q_{-1}} = 1.$$

Using $B_{-1} = (1-\theta)W_{-1}$ and $Q_{-1} = \theta W_{-1}/p_{-1}$ from (A3),

$$\theta p_0 + (1-\theta)p_{-1} + \frac{\theta F_0}{Q_{-1}} = p_0 \implies p_0 = p_{-1} + \frac{\theta}{1-\theta}\frac{F_0}{Q_{-1}}.$$

From (A6) with $Q_0 = Q_{-1} = \bar{Q}$, $F_0 = B_0 - B_{-1}$; since the bond leg grows only through new lending to the non-wealthy, $B_0 - B_{-1} = B_0^P$, so $F_0 = B_0^P$. Substituting and imposing $Q_{-1} = \bar{Q}$ gives (A9). □

For every additional dollar of borrowing absorbed by the intermediary, equity market value must rise by $\frac{\theta}{1-\theta} = 3.55$ at the calibrated $\theta = 0.78$; we call $\frac{\theta}{1-\theta}$ the *price impact* of borrowing. The mechanism is a single one: saving by the wealthy is intermediated as new lending to the non-wealthy, so $F_0 = B_0^P$; this dollar cannot buy more equity (supply is fixed at $\bar{Q}$), so it is absorbed as a bond, but the mandate forces the equity leg $p_0\bar{Q}$ to rise in step. With fixed supply, the equity price is pinned to the outstanding bond position rather than to fundamentals alone.

What determines borrowing? With log utility, borrowing is pinned down by endowments:

(A10) $$F_0 = B_0^P = \frac{e_1^P - \beta(1+r_{f,0})e_0^P - (1+r_{f,-1})^2 B_{-1}}{(1+\beta)(1+r_{f,0})},$$

so borrowing is positive whenever future non-wealthy endowments exceed legacy repayment and desired consumption. Although the non-wealthy's income *share* has fallen, their absolute income *levels* have risen, sustaining more debt.

B.5. Price Decomposition and Downward-Sloping Demand

Since prices depend on borrowing, which depends on endowments, the fundamental drivers can be made explicit.

PROPOSITION A2. *Under log utility, the price of equity depends on dividends, the income distribution, equity supply, and the intermediary constraint:*

(A11) $$p_0 = \frac{\beta}{1+\beta}\left[\frac{D_1}{(1+r_{f,-1})^2}\right] + \frac{1}{1+\beta}\left[\frac{1}{\bar{Q}}\frac{\theta}{1-\theta}\left\{\frac{e_1^P - \frac{c_1^P}{c_0^P}e_0^P}{1+r_{f,0}}\right\}\right].$$

PROOF. Start from (A9) and use (A7) to write $B_0^P = c_0^P - e_0^P$. The non-wealthy intertemporal budget constraint is

$$c_0^P + \frac{c_1^P}{1+r_{f,0}} + \frac{(1+r_{f,-1})^2 B_{-1}}{1+r_{f,0}} = e_0^P + \frac{e_1^P}{1+r_{f,0}}.$$

Under log utility $c_1^P = \beta(1+r_{f,0})c_0^P$, so

$$c_0^P = \frac{1}{1+\beta}\left(e_0^P + \frac{e_1^P - (1+r_{f,-1})^2 B_{-1}}{1+r_{f,0}}\right).$$

Substituting into the pricing equation and using $p_{-1} = D_1/(1+r_{f,-1})^2$ and $B_{-1}/(1-\theta) = W_{-1} = p_{-1}Q_{-1}/\theta$ from (A2)–(A3),

$$p_0 = \frac{D_1}{(1+r_{f,-1})^2}\left(1 - \frac{(1+r_{f,-1})^2}{(1+\beta)(1+r_{f,0})}\right) + \frac{\theta}{1-\theta}\frac{1}{\bar{\bar{Q}}}\frac{1}{1+\beta}\left(\frac{e_1^P}{1+r_{f,0}} - \beta e_0^P\right).$$

By the expectations hypothesis $(1+r_{f,-1})^2 = (1+r_{f,-1\to 0})(1+r_{f,0})$, and $r_{f,-1\to 0} = 0$ in equilibrium (no risk-free trade from $t=-1$ to $t=0$), so $(1+r_{f,-1})^2 = (1+r_{f,0})$. Substituting gives

$$p_0 = \frac{\beta}{1+\beta}\frac{D_1}{(1+r_{f,-1})^2} + \frac{1}{1+\beta}\frac{\theta}{1-\theta}\frac{1}{\bar{\bar{Q}}}\left(\frac{e_1^P}{1+r_{f,0}} - \beta e_0^P\right). \qquad \square$$

Income distribution therefore matters for equity prices, contrary to Grossman and Shiller (1981). Unlike Kiyotaki and Moore (1997), where borrowing is driven by heterogeneous patience or preferences, both households here are equally patient; the non-wealthy borrow a positive amount whenever

$$e_1^P - \frac{c_1^P}{c_0^P}e_0^P > 0 \Leftrightarrow \frac{e_1^P}{e_0^P} > \beta(1+r_{f,0}),$$

i.e. whenever their absolute income growth exceeds their consumption growth. Lower equity supply and higher θ both amplify the price impact of a given flow. Figure A5 contrasts the model's downward-sloping demand with the flat Lucas (1978) benchmark: price adjustments are required to absorb flows even absent fundamental changes.

B.6. Elasticity of Equity Demand

The price is decreasing in quantity, $\partial p_0/\partial Q_0 < 0$: inelastic markets generate downward-sloping demand. The elasticity is pinned down by θ.

PROPOSITION A3. *The signed price elasticity of equity demand is* $\xi = 1-\theta$.

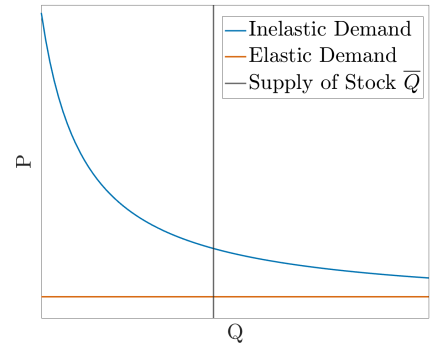

FIGURE A5. Inelastic demand curve ($\xi = 0.2$) versus elastic demand curve.

PROOF. From (A4) and (A5), $Q_0 = \theta W_0/p_0$. Taking logs and first-differencing, $\Delta \ln Q = \Delta \ln \theta + \Delta \ln W - \Delta \ln p$. A first-order approximation gives $\Delta \ln W \approx \theta \Delta \ln p + F_0/W_{-1}$; with $\Delta \ln \theta = 0$,

$$\Delta \ln Q = (\theta - 1)\Delta \ln p + \frac{F_0}{W_{-1}} \implies \xi = -\frac{\partial \Delta \ln Q}{\partial \Delta \ln p} = 1 - \theta. \qquad \square$$

As θ rises the market becomes more inelastic and the price impact $\frac{\theta}{1-\theta}$ grows. When $\theta = 0$, $p_0 = p_{-1}$; this is *not* the Lucas benchmark, since the price does not evolve despite changing discount rates and $\xi = 1 \neq \infty$. Appendix B.10 derives the Lucas model as the special case in which the intermediary re-optimises θ at $t = 0$.

B.7. Wealth Inequality

DEFINITION A2. *For $t \in \{-1, 0\}$, wealth is $W_t^R \equiv p_t Q_t + B_t$ (wealthy) and $W_t^P \equiv \mathcal{O}^P - B_t$ with $\mathcal{O}^P > B_t$ (non-wealthy), where $\mathcal{O}^P$ is non-marketable assets (e.g. primary residence) ensuring a non-negative net position. Wealth shares are $\eta_t^R \equiv W_t^R/(W_t^R + W_t^P)$ and $\eta_t^P \equiv 1 - \eta_t^R$.*

We consider $t \in \{-1, 0\}$ because the wealth evolution from $t = -1$ to $t = 0$ is what inelastic markets shape (at $t = 1$ wealth trivially equals income).

PROPOSITION A4. *The wealthy wealth share is increasing in equity prices:*

$$\frac{\partial(\eta_0^R - \eta_{-1}^R)}{\partial p_0} > 0. \tag{A12}$$

PROOF. Substituting the wealth share definitions,

$$\frac{\partial(\eta_0^R - \eta_{-1}^R)}{\partial p_0} = \frac{Q_0(\mathcal{O}^P - B_0)}{(p_0 Q_0 + \mathcal{O}^P)^2} > 0,$$

since $\mathcal{O}^P > B_0$ by assumption. $\square$

By the pricing equation, the wealthy share rises whenever borrowing is higher, θ is higher, $\bar{Q}$ is lower, or p_{-1} is higher—linking wealth inequality directly to market inelasticity.

B.8. Comparative Statics on Income

We study how endowments map into prices and inequality. Closed forms for prices in income shares are unavailable, so we simulate with log utility, $D_1 = 2$, $\theta = 0.8$, $\beta = 0.99$, $W_{-1} = 5$, $\bar{Q} = 10$, and endowments $\{e_0^P, e_0^R\} = \{2, 8\}$, $\{e_1^P, e_1^R\} = \{8, 12\}$ (top 10% income share 80%).

DEFINITION A3. *For $t \in \{0,1\}$, income is $I_t^R \equiv e_t^R + \mathbb{1}_{t=1} D_1 Q_t$ and $I_t^P \equiv e_t^P$, with shares $\mathcal{J}_t^R \equiv I_t^R/(I_t^R + I_t^P)$ and $\mathcal{J}_t^P \equiv 1 - \mathcal{J}_t^R$.*

PROPOSITION A5. *The partial derivatives of the equity price with respect to endowments are*

$$\frac{\partial p_0}{\partial e_0^P} < 0, \quad \frac{\partial p_0}{\partial e_1^P} > 0, \quad \frac{\partial p_0}{\partial e_0^R} > 0, \quad \frac{\partial p_0}{\partial e_1^R} < 0.$$

A higher current non-wealthy endowment lowers borrowing (prices fall); a higher future non-wealthy endowment raises borrowing capacity (prices rise); a higher current wealthy endowment raises flows (prices rise); a higher future wealthy endowment lowers saving (prices fall). The wealthy effect dominates:

PROPOSITION A6. *Prices increase when the current-period wealthy income share increases: $\partial p_0 / \partial \mathcal{J}_0^R > 0$.*

Holding $\mathcal{J}_0^R = \mathcal{J}_1^R = 0.8$, the top 10% wealth share rises from 0.625 to 0.644 (+3%). When the top 10% income share instead *falls* from 0.8 to 0.75, the top 10% wealth share rises from 0.625 to 0.7 (+12%): a fall in the wealthy's relative income raises their saving while the non-wealthy's higher absolute income raises borrowing, and under the mandate both feed the same equity price rise, widening wealth inequality even as income inequality narrows.

B.9. Exogenous Increase in Equity Supply

The fixed-supply assumption is not essential. Let supply change exogenously by δ_q, so $Q_0 = Q_{-1}(1 + \delta_q)$.

PROPOSITION A7. *The price of equity at $t = 0$ is*

$$p_0 = \left(\frac{1-\theta}{1-\theta+\delta_q}\right)\left[p_{-1} + \frac{\theta}{1-\theta}\frac{F_0}{Q_{-1}}\right]. \tag{A13}$$

PROOF. Clearing now requires $(Q_0 - Q_{-1})/Q_{-1} = \delta_q$. Repeating the proof of Proposition A1 with this right-hand side gives

$$(1-\theta)p_{-1} + \frac{\theta F_0}{Q_{-1}} = p_0(1-\theta+\delta_q),$$

which rearranges to (A13). □

Since $\partial p_0/\partial\delta_q < 0$, the multiplier $\frac{\theta}{1-\theta+\delta_q}$ is dampened for $\delta_q > 0$, but short of $\theta = 0$ prices remain flow-sensitive. The percentage price change is $\frac{p_0 - p_{-1}}{p_{-1}} = \frac{\theta F_0/(p_{-1}Q_{-1}) - \delta_q}{1-\theta+\delta_q}$, so $\Delta p > 0 \Leftarrow \delta_q < \theta F_0/(p_{-1}Q_{-1})$. Writing flows as a share f of intermediary AuM and using $p_{-1}Q_{-1} = \theta W_{-1}$, this is simply $\delta_q < f$. Empirically $f \approx 2\%$ per quarter (Flow of Funds) while $\delta_q \approx 0.2\%$ per quarter (SIFMA issuance), so issuance does not negate the flow effect. Finally, with variable supply the tight B_0^P–F_0 link loosens, since part of each flow finances new shares:

COROLLARY A1. *When supply changes by δ_q, the price–debt relationship is*

$$p_0 = \left(\frac{1}{1+\delta_q}\right)\left[p_{-1} + \frac{\theta}{1-\theta}\frac{B_0^P}{Q_{-1}}\right]. \tag{A14}$$

PROOF. Substitute $F_0 = p_0(Q_0 - Q_{-1}) + (B_0 - B_{-1})$ into (A13) and solve for p_0. □

As δ_q rises, the equity price impact of higher indebtedness falls.

B.10. The Lucas Model as a Special Case

If instead the intermediary re-optimises θ at $t = 0$, it solves $\max_\theta \frac{1}{1+r_{f,0}}[\theta\frac{D_1}{p_0} + (1-\theta)(1+r_{f,0})]$. The first-order condition with market clearing (no trade) gives $p_0 = D_1/(1+r_{f,0})$. The Lucas benchmark is thus nested as the case in which the intermediary can rebalance at $t = 0$ in response to flows.

Appendix C. The Inelastic Quantitative Model

This section lays out the deterministic steady state and the log-linearised representation for the model presented in Section 2.

C.1. Deterministic Steady State

The deterministic steady state is the point around which the linear system is constructed. The following restrictions are useful for interpreting the log-linear equations. Gross returns satisfy

$$\text{(A15)} \qquad \bar{R}^E = \frac{\bar{D} + \bar{p}}{\bar{p}}, \qquad \bar{R}^{MF} = \bar{\theta}\bar{R}^E + (1 - \bar{\theta})\bar{R}.$$

The mutual fund balance sheet implies

$$\text{(A16)} \qquad \bar{X}^{MF} = \bar{p}\bar{Q} + \bar{B}^{MF}, \qquad \bar{\theta} = \frac{\bar{p}\bar{Q}}{\bar{X}^{MF}}, \qquad \bar{B}^{MF} = \bar{B}^P + \bar{B}^M + \bar{B}^G.$$

The wealthy household's Euler equation, evaluated at $\widehat{\zeta}^d = 0$, implies

$$\text{(A17)} \qquad 1 + \phi_{MF} = \delta\bar{\zeta}^R + \beta^R\bar{R}^{MF}, \qquad \bar{\zeta}^R = (\bar{a}^R)^{-\Sigma_R}(\bar{c}^R)^{\sigma_R}.$$

For the middle class,

$$\text{(A18)} \qquad \bar{\psi}^M = 1 - \beta^M\bar{R}, \qquad 1 + \phi_H^M = \beta^M\frac{\bar{D}^H + \bar{p}^H}{\bar{p}^H} + \bar{\tau}^M\bar{\psi}^M,$$

and the binding collateral constraint gives

$$\text{(A19)} \qquad \bar{B}^M = \bar{\tau}^M\bar{X}^{H,M}.$$

For non-wealthy households,

$$\text{(A20)} \qquad 1 - \phi_B^P\bar{B}^P = \beta^P\bar{R}, \qquad 1 + \phi_H^P = \beta^P\frac{\bar{D}^H + \bar{p}^H}{\bar{p}^H}.$$

Housing-market clearing in steady state is

$$\text{(A21)} \qquad \bar{X}^{H,M} + \bar{X}^{H,P} = \bar{p}^H\bar{H} \equiv \bar{X}^H.$$

C.2. Log-Linearised Equilibrium System

This subsection gives the log-linearised system corresponding to the inelastic model. All variables are proportional deviations from the 1989 deterministic steady state unless otherwise noted. Transfer deviations $\widetilde{T}_t$ are level deviations. The notation $\widehat{\mu}_t^M$ denotes the proportional deviation of the level collateral multiplier μ_t^M; its steady-state consumption-normalised value is $\bar{\psi}^M \equiv \bar{\mu}^M / u'_M(\bar{c}^M)$. In the stochastic version, all variables dated $t+1$ inside Euler equations are conditional expectations; in the perfect-foresight implementation, the expectations operator is suppressed. The equation tags match the corresponding blocks of the implementation.[37]

A. Mutual Fund and Wealthy Household

The mutual fund mandate rule is

$$\bar{\theta}\widehat{\theta}_t = \theta^* \frac{\beta^{MF}}{\chi^{MF}} \left(\bar{R}^E \widehat{R}_t^E - \bar{R}\widehat{R}_t\right). \tag{A1}$$

The equity-return identity is

$$(\bar{D} + \bar{p})\widehat{R}_t^E = \bar{D}\widehat{D}_{t+1} + \bar{p}\widehat{p}_{t+1} - (\bar{D} + \bar{p})\widehat{p}_t. \tag{A2}$$

The mutual fund portfolio return is

$$\bar{R}^{MF}\widehat{R}_t^{MF} = \bar{\theta}\bar{R}^E\widehat{R}_t^E + (1-\bar{\theta})\bar{R}\widehat{R}_t + \bar{\theta}(\bar{R}^E - \bar{R})\widehat{\theta}_t. \tag{A3}$$

The value of the mutual fund position satisfies

$$\widehat{X}_t^{MF} = \bar{\theta}\widehat{p}_t + (1-\bar{\theta})\widehat{B}_t^{MF}. \tag{A4}$$

Wealthy household assets are

$$\bar{a}^R\widehat{a}_t^R = \bar{X}^{MF}\widehat{X}_t^{MF}. \tag{A5}$$

The wealthy household budget constraint is

$$\bar{c}^R\widehat{c}_t^R + (1+\phi_{MF})\bar{X}^{MF}\widehat{X}_t^{MF} = \bar{e}^R\widehat{e}_t^R + \bar{R}^{MF}\bar{X}^{MF}\left(\widehat{R}_{t-1}^{MF} + \widehat{X}_{t-1}^{MF}\right) + s_T^R\widetilde{T}_t. \tag{A6}$$

[37] The tag A8 is reserved in the implementation for an equation that is absent from this specification, so the block runs A1–A7 and A9.

The wealthy household mutual fund Euler equation, including the demand wedge, is

$$\phi_{MF}\left(\widehat{X}_t^{MF} - \widehat{\zeta}_t^d\right) = \delta\bar{\zeta}^R\left(\sigma_R\widehat{c}_t^R - \Sigma_R\widehat{a}_t^R\right) + \beta^R\bar{R}^{MF}\left[\widehat{R}_t^{MF} - \sigma_R\left(\widehat{c}_{t+1}^R - \widehat{c}_t^R\right)\right]. \tag{A7}$$

Finally, the equity price follows from the mutual fund balance sheet:

$$\widehat{p}_t = \widehat{B}_t^{MF} + \frac{1}{1-\bar{\theta}}\widehat{\theta}_t. \tag{A9}$$

Equation (A9) is the log-linear version of the inelastic pricing identity (27). It shows that equity prices move one-for-one with the mutual fund bond leg and are amplified by changes in the mandate share.

B. Middle Class Households

The middle class budget constraint is

$$\begin{aligned} &\bar{c}^M\widehat{c}_t^M + (1+\phi_H^M)\bar{X}^{H,M}\widehat{X}_t^{H,M} + \bar{R}\bar{B}^M\left(\widehat{R}_{t-1} + \widehat{B}_{t-1}^M\right) \\ &\quad = \bar{e}^M\widehat{e}_t^M + \bar{D}^H\bar{H}^M\widehat{D}_t^H + \bar{X}^{H,M}\widehat{p}_t^H \\ &\qquad + (\bar{D}^H + \bar{p}^H)\bar{H}^M\left(\widehat{X}_{t-1}^{H,M} - \widehat{p}_{t-1}^H\right) + \bar{B}^M\widehat{B}_t^M + s_T^M\widetilde{T}_t. \end{aligned} \tag{B1}$$

The middle class bond Euler equation is

$$\bar{\psi}^M\widehat{\mu}_t^M = -\beta^M\bar{R}\left[\widehat{R}_t - \sigma_M\left(\widehat{c}_{t+1}^M - \widehat{c}_t^M\right)\right] - \bar{\psi}^M\sigma_M\widehat{c}_t^M. \tag{B2}$$

The middle class housing Euler equation is

$$\begin{aligned} \phi_H^M\bar{p}^H\widehat{X}_t^{H,M} = \beta^M\Big[&\bar{D}^H\widehat{D}_{t+1}^H + \bar{p}^H\widehat{p}_{t+1}^H - (\bar{D}^H + \bar{p}^H)\widehat{p}_t^H \\ &- \sigma_M(\bar{D}^H + \bar{p}^H)\left(\widehat{c}_{t+1}^M - \widehat{c}_t^M\right)\Big] \\ &+ \bar{\tau}^M\bar{p}^H\bar{\psi}^M\left(\widehat{\mu}_t^M + \sigma_M\widehat{c}_t^M + \widehat{\tau}_t^M\right). \end{aligned} \tag{B3}$$

The binding collateral constraint is

$$\widehat{B}_t^M = \widehat{\tau}_t^M + \widehat{X}_t^{H,M}. \tag{B4}$$

C. Non-Wealthy Households

The non-wealthy budget constraint is

$$
\begin{aligned}
&\bar{c}^P \widehat{c}_t^P + (1+\phi_H^P)\bar{X}^{H,P}\widehat{X}_t^{H,P} + \bar{R}\bar{B}^P\left(\widehat{R}_{t-1} + \widehat{B}_{t-1}^P\right) + \phi_B^P(\bar{B}^P)^2\widehat{B}_t^P \\
&\quad = \bar{e}^P\widehat{e}_t^P + \bar{D}^H\bar{H}^P\widehat{D}_t^H + \bar{X}^{H,P}\widehat{p}_t^H \\
&\qquad + (\bar{D}^H + \bar{p}^H)\bar{H}^P\left(\widehat{X}_{t-1}^{H,P} - \widehat{p}_{t-1}^H\right) + \bar{B}^P\widehat{B}_t^P + s_T^P\widetilde{T}_t.
\end{aligned}
\tag{C1}
$$

The unconstrained debt Euler equation is

$$
0 = -\beta^P\bar{R}\left[\widehat{R}_t - \sigma_P\left(\widehat{c}_{t+1}^P - \widehat{c}_t^P\right)\right] - \phi_B^P\bar{B}^P\widehat{B}_t^P. \tag{C2}
$$

The non-wealthy housing Euler equation is

$$
\phi_H^P\bar{p}^H\widehat{X}_t^{H,P} = \beta^P\Big[\bar{D}^H\widehat{D}_{t+1}^H + \bar{p}^H\widehat{p}_{t+1}^H - (\bar{D}^H + \bar{p}^H)\widehat{p}_t^H - \sigma_P(\bar{D}^H + \bar{p}^H)\left(\widehat{c}_{t+1}^P - \widehat{c}_t^P\right)\Big]. \tag{C3}
$$

D. Government and Market Clearing

The government budget constraint is

$$
\widetilde{T}_t = \bar{B}^G\widehat{B}_t^G - \bar{R}\bar{B}^G\widehat{B}_{t-1}^G - \bar{R}\bar{B}^G\widehat{R}_{t-1}. \tag{D1}
$$

Bond-market clearing is

$$
\bar{B}^{MF}\widehat{B}_t^{MF} = \bar{B}^P\widehat{B}_t^P + \bar{B}^M\widehat{B}_t^M + \bar{B}^G\widehat{B}_t^G. \tag{D2}
$$

Housing-market clearing is

$$
\bar{X}^{H,M}\widehat{X}_t^{H,M} + \bar{X}^{H,P}\widehat{X}_t^{H,P} = \bar{X}^H\widehat{p}_t^H. \tag{D3}
$$

E. Exogenous Processes

The exogenous processes follow AR(1) laws of motion:

$$
\widehat{e}_t^R = \rho_{e^R}\widehat{e}_{t-1}^R + \varepsilon_t^{e^R}, \tag{E1}
$$

$$\widehat{e}^M_t = \rho_{e^M}\widehat{e}^M_{t-1} + \varepsilon^{e^M}_t, \tag{E2}$$

$$\widehat{e}^P_t = \rho_{e^P}\widehat{e}^P_{t-1} + \varepsilon^{e^P}_t, \tag{E3}$$

$$\widehat{D}_t = \rho_D\widehat{D}_{t-1} + \varepsilon^D_t, \tag{E4}$$

$$\widehat{D}^H_t = \rho_{D^H}\widehat{D}^H_{t-1} + \varepsilon^{D^H}_t, \tag{E5}$$

$$\widehat{B}^G_t = \rho_{B^G}\widehat{B}^G_{t-1} + \varepsilon^{B^G}_t, \tag{E6}$$

$$\widehat{\tau}^M_t = \rho_{\tau^M}\widehat{\tau}^M_{t-1} + \varepsilon^{\tau^M}_t, \tag{E7}$$

$$\widehat{\zeta}^d_t = \rho_\zeta\widehat{\zeta}^d_{t-1} + \varepsilon^\zeta_t. \tag{E8}$$

C.3. Auxiliary Reporting Identities

The following identities are used to report total debt, wealth shares, and income shares. They do not add independent equilibrium restrictions.

Total bonds are reported as

$$\bar{B}^{tot}\widehat{B}^{tot}_t = \bar{B}^{MF}\widehat{B}^{MF}_t. \tag{AUX1}$$

Define aggregate measured wealth as

$$\bar{W} = \bar{p}\bar{Q} + \bar{X}^H + \bar{B}^G, \qquad \bar{X}^H \equiv \bar{X}^{H,M} + \bar{X}^{H,P}. \tag{A22}$$

The wealthy household's wealth share deviation is

$$\eta^{dev}_{R,t} = \frac{\bar{a}^R}{\bar{W}}\widehat{a}^R_t - \bar{\eta}_R\frac{\bar{p}\bar{Q}\widehat{p}_t + \bar{X}^H\widehat{p}^H_t + \bar{B}^G\widehat{B}^G_t}{\bar{W}}. \tag{AUX2}$$

The middle class wealth share deviation is

$$\eta^{dev}_{M,t} = \frac{\bar{X}^{H,M}\widehat{X}^{H,M}_t - \bar{B}^M\widehat{B}^M_t}{\bar{W}} - \bar{\eta}_M\frac{\bar{p}\bar{Q}\widehat{p}_t + \bar{X}^H\widehat{p}^H_t + \bar{B}^G\widehat{B}^G_t}{\bar{W}}. \tag{AUX3}$$

The non-wealthy wealth share deviation is

$$\eta^{dev}_{P,t} = \frac{\bar{X}^{H,P}\widehat{X}^{H,P}_t - \bar{B}^P\widehat{B}^P_t}{\bar{W}} - \bar{\eta}_P\frac{\bar{p}\bar{Q}\widehat{p}_t + \bar{X}^H\widehat{p}^H_t + \bar{B}^G\widehat{B}^G_t}{\bar{W}}. \tag{AUX4}$$

Let the aggregate income deviation be

$$\Delta Y_t = \bar{e}^R \widehat{e}_t^R + \bar{D}\bar{Q}\widehat{D}_t + \bar{e}^M \widehat{e}_t^M + \bar{D}^H \bar{H}^M \left(\widehat{D}_t^H + \widehat{X}_{t-1}^{H,M} - \widehat{p}_{t-1}^H\right) + \bar{e}^P \widehat{e}_t^P + \bar{D}^H \bar{H}^P \left(\widehat{D}_t^H + \widehat{X}_{t-1}^{H,P} - \widehat{p}_{t-1}^H\right). \tag{A23}$$

Then the wealthy household's income-share deviation is

$$\widehat{inc}_{R,t} = \frac{\bar{e}^R}{\bar{Y}^R}\widehat{e}_t^R + \frac{\bar{D}\bar{Q}}{\bar{Y}^R}\widehat{D}_t - \frac{\Delta Y_t}{\bar{Y}}. \tag{AUX5}$$

The middle class income-share deviation is

$$\widehat{inc}_{M,t} = \frac{\bar{e}^M}{\bar{Y}^M}\widehat{e}_t^M + \frac{\bar{D}^H \bar{H}^M}{\bar{Y}^M}\left(\widehat{D}_t^H + \widehat{X}_{t-1}^{H,M} - \widehat{p}_{t-1}^H\right) - \frac{\Delta Y_t}{\bar{Y}}. \tag{AUX6}$$

The non-wealthy income-share deviation is

$$\widehat{inc}_{P,t} = \frac{\bar{e}^P}{\bar{Y}^P}\widehat{e}_t^P + \frac{\bar{D}^H \bar{H}^P}{\bar{Y}^P}\left(\widehat{D}_t^H + \widehat{X}_{t-1}^{H,P} - \widehat{p}_{t-1}^H\right) - \frac{\Delta Y_t}{\bar{Y}}. \tag{AUX7}$$

The equilibrium system used in the inelastic simulations consists of equations (A1)–(A7), (A9), (B1)–(B4), (C1)–(C3), (D1)–(D3), and the exogenous processes (E1)–(E8). The auxiliary identities are used only for reporting simulated debt, wealth share, and income-share series. Because all aggregate borrowing is absorbed by the mutual fund, changes in household and government debt directly affect the mutual fund's balance sheet and, through the mandate identity, the equilibrium equity price.

Appendix D. Non-Homothetic Preferences and Calibration of Exogenous Processes

D.1. Non-Homothetic Preferences

Section 3.1 specifies household-specific felicity u_i and, for the wealthy alone, a bequest function ν_R with $\Sigma_R < \sigma_R$. Here we show why this single restriction generates non-homothetic saving among the wealthy, while the middle class and non-wealthy—who have no bequest motive—remain homothetic.

For the wealthy household, define the ratio of the marginal utility of bequeathing to the marginal utility of consuming at wealth level a,

$$\zeta(a) = \frac{\nu_R'(a)}{u_R'(a)} = a^{\sigma_R - \Sigma_R}. \tag{A24}$$

As in Mian et al. (2025), when $\Sigma_R = \sigma_R$ this ratio is constant: the two marginal utilities are proportional at every wealth level and preferences are homothetic. When $\Sigma_R < \sigma_R$, $\zeta(a)$ is strictly increasing—the marginal utility of wealth decays more slowly than that of consumption, so wealthier households have a stronger *relative* incentive to save. Bequests are thus a luxury good, the saving rate rises with permanent income, and wealth accumulates disproportionately at the top: this is the mechanism that drives the divergence between wealth and income shares over the transition. In the quantitative model of Appendix C, the corresponding equilibrium object is $\zeta_t^R = \nu_R'(a_t^R)/u_R'(c_t^R)$.

The middle class and the non-wealthy have pure CRRA felicity (u_M, u_P) with no bequest term, so no analogous margin exists: their intertemporal marginal rate of substitution is independent of the level of wealth, and their consumption–saving behaviour is homothetic—saving rates do not vary with permanent income. Non-homotheticity in the model therefore operates entirely through the wealthy household's bequest luxury. The type-specific elasticities $\sigma_R > \sigma_M = \sigma_P$ additionally shape the *level* of each group's intertemporal substitution but, absent a bequest term for M and P, do not by themselves break homotheticity for those groups.

D.2. Exogenous Processes

The exogenous drivers of the quantitative model are the labour endowments of the three household groups, equity dividends, housing rents, government debt, the middle class loan-to-value ratio, and the portfolio demand wedge, $\{e_t^R, e_t^M, e_t^P, D_t, D_t^H, B_t^G, \tau_t^M, \zeta_t^d\}$. Each follows an AR(1) process in proportional deviations from its 1989 steady-state value,

$$\widehat{z}_t = \rho_z \widehat{z}_{t-1} + \varepsilon_t^z, \qquad z \in \{e^R, e^M, e^P, D, D^H, B^G, \tau^M, \zeta^d\},$$

with i.i.d. normal innovations. The complete set of laws of motion is reported as Block E of the log-linearised system in Appendix C.2. Steady-state levels $\bar{e}^P$, $\bar{e}^M$, $\bar{e}^R$, $\bar{D}$, and $\bar{D}^H$ are calibrated in Table 1, and the innovation standard deviations σ^P, σ^M, σ^R, σ^D, and σ^{D^H} are chosen as described in Section 3.

Appendix E. Additional Qualitative Results

E.1. Elasticity in Infinite-Period Model

In the quantitative model, elasticity of demand for the stock is time-varying. We can derive the aggregate (signed) time-varying Hicksian elasticity of demand for stock as follows:

(A25)
$$\xi_t = (1-\theta^*) - \frac{\beta^{MF}}{\chi^{MF}}\left[\theta^* \underbrace{(r^e_{t-1} - r_{t-1})}_{\text{Equity Premium}} - \left\{\frac{D}{p} - 1\right\}\right].$$

where $\frac{D}{p}$ is the average dividend-price ratio.

PROOF. Start with an inverted version of the mutual fund's mandate rule, Equation 7:

$$\frac{\theta_t - \theta^*}{\theta^*} = \frac{\beta^{MF}}{\chi^{MF}}\left(\frac{D_{t+1} + p_{t+1} - p_t}{p_t} - r_t\right) = \Delta \ln \theta_t$$

This is equivalent to the first-order Taylor expansion of θ_t around θ^*. Consider the term on the right-hand side within brackets, i.e. the equity premium π_t, which can be rewritten as:

$$\pi_t = \frac{D_{t+1}}{p_t} + \Delta \ln p_t - r_t$$

Consider the percentage change in the dividend-price ratio, i.e. $\Delta \ln \frac{D_{t+1}}{p_t}$. Express the percentage log deviation of dividends as $\tilde{d}$ and the percentage log deviation of prices as $\tilde{p}$. Then, $\Delta \ln \frac{D_{t+1}}{p_t} = \tilde{d} - \tilde{p}$.
Next, consider a first-order Taylor approximation of $\Delta \ln \frac{D_{t+1}}{p_t}$ around some baseline dividend-price ratio $\frac{D}{p}$:

$$\Delta \ln \frac{D_{t+1}}{p_t} = \frac{\frac{D_{t+1}}{p_t} - \frac{D}{p}}{\frac{D}{p}}$$

$$\tilde{d} - \tilde{p} = \frac{\frac{D_{t+1}}{p_t}}{\frac{D}{p}} - 1$$

$$\frac{D_{t+1}}{p_t} = (\tilde{d} - \tilde{p} + 1)\frac{D}{p}$$

Substituting into π_t:

$$\pi_t = (1 + \tilde{d})\frac{D}{p} + \tilde{p}(1 - \frac{D}{p}) - r_t$$

Go back to the mandate equation and invert to find quantity demanded:

$$Q_t = \frac{\theta_t X_t^{MF}}{p_t}$$

Taking logs and first differencing, and writing the wealthy household's net inflow into the fund as F_t, so that $f_t \equiv F_t/X_{t-1}^{MF}$ is the inflow rate relative to assets under management:

$$\begin{aligned}
\Delta \ln Q_t &= \Delta \ln \theta_t + \Delta \ln X_t^{MF} - \Delta \ln p_t \\
q_t &= \Delta \ln \theta_t + \left(\frac{p_{t-1} Q_{t-1}}{X_{t-1}^{MF}} \frac{\Delta p_t}{p_{t-1}} + \underbrace{\frac{F_t}{X_{t-1}^{MF}}}_{f_t} \right) - \tilde{p} \\
q_t &= \Delta \ln \theta_t + \theta_{t-1} \tilde{p} + f_t - \tilde{p}
\end{aligned}$$

Substituting from π_t and $\Delta \ln \theta_t$:

$$q_t = \frac{\beta^{MF}}{\chi^{MF}} \left((1 + \tilde{d}) \frac{D}{p} + \tilde{p}(1 - \frac{D}{p}) - r_t \right) + \tilde{p}(\theta_{t-1} - 1) + f_t$$

Finally, to find the elasticity, take the derivative of q_t with respect to $\tilde{p}$:

$$\xi_t \equiv \frac{\partial q_t}{\partial \tilde{p}} = \frac{\beta^{MF}}{\chi^{MF}} (1 - \frac{D}{p}) + (\theta_{t-1} - 1)$$

Substituting in the value of θ_{t-1} yields the desired proposition for unsigned elasticity:

$$\xi_t = (\theta^* - 1) + \frac{\beta^{MF}}{\chi^{MF}} \left[\theta^* (r_{t-1}^e - r_{t-1}) - \left\{ \frac{D}{p} - 1 \right\} \right].$$

□

As the equity premium increases, the intermediary demands more equity, driving unsigned elasticity up in the subsequent period. Similarly, as the cost of adjustment χ^{MF} increases, the unsigned elasticity decreases with the intermediary finding it more costly to deviate from target. The limiting case of $\chi^{MF} \to \infty$ (or $\beta^{MF} \to 0$) gets us back to the tight mandate condition and corresponding elasticity of the three-period model.

E.2. An Alternate Elastic Benchmark in the Quantitative Model

The elastic markets benchmark implemented in the quantitative simulations replaces the mutual fund altogether: wealthy households hold equity and bonds directly, and equity is priced by their Euler equation.

That benchmark is described in Appendix 2.8. An instructive alternative retains the mutual fund but removes the mandate friction. Setting $\chi^{MF} = 0$ in the fund's problem (Equation 6), the intermediary solves

$$\max_{\theta_t} \beta^{MF} \mathbb{E}_t \left[\theta_t \left\{ \frac{D_{t+1} + p_{t+1}}{p_t} \right\} + (1 - \theta_t)(1 + r_t) \right].$$

The first-order condition with respect to θ_t yields the standard pricing function

$$p_t = \mathbb{E}_t \left[\frac{D_{t+1} + p_{t+1}}{1 + r_t} \right].$$

This is the well-established Lucas (1978) result, where the price of an asset depends upon its discounted stream of future dividends. The Lucas model is therefore nested in our framework as the frictionless limit of the mandate problem.

E.3. Exogenous Increase in Equity Supply

The baseline quantitative model fixes equity supply at $\bar{Q}$. We now relax this and let supply differ from its baseline level by an exogenous δ_q, so that $Q_t = \bar{Q}(1 + \delta_q)$, and ask how the equilibrium price and its link to debt respond. In the quantitative model, supply enters the mandate pricing identity only through the denominator.

PROPOSITION A8. *With equity supply $Q_t = \bar{Q}(1 + \delta_q)$, the equilibrium equity price is*

$$p_t = \frac{1}{1 + \delta_q} \frac{\theta_t}{1 - \theta_t} \frac{B_t^{MF}}{\bar{Q}}. \tag{A26}$$

PROOF. The mandate identities $p_t Q_t = \theta_t X_t^{MF}$ and $B_t^{MF} = (1 - \theta_t) X_t^{MF}$ from (5) give $p_t Q_t = \frac{\theta_t}{1-\theta_t} B_t^{MF}$. Substituting $Q_t = \bar{Q}(1 + \delta_q)$ and solving for p_t yields (A26). □

Since $\partial p_t / \partial \delta_q < 0$, exogenous issuance dampens the price–debt multiplier $\frac{\theta_t}{1-\theta_t}$ by the factor $\frac{1}{1+\delta_q}$ relative to the fixed-supply identity of Proposition 1, which is recovered at $\delta_q = 0$. Short of $\theta_t \to 0$, prices remain sensitive to the fund's bond position, and hence to borrowing and flows.

To see that realistic issuance does not overturn the flow effect, hold the mandate share at its target and let the fund's assets under management grow with net inflows, $X_t^{MF} = (1 + f) X_{t-1}^{MF}$, where f is the flow into the fund as a share of assets under management. From $p_t = \theta_t X_t^{MF}/Q_t$ with $Q_t = \bar{Q}(1+\delta_q)$ and $Q_{t-1} = \bar{Q}$,

$$\frac{p_t - p_{t-1}}{p_{t-1}} = \frac{1+f}{1+\delta_q} - 1 = \frac{f-\delta_q}{1+\delta_q},$$

so prices rise, $\Delta p > 0$, if and only if $\delta_q < f$: as long as equity issuance grows more slowly than the fund's inflows, positive flows still raise prices. Empirically, $f \approx 2\%$ per quarter (Flow of Funds) while $\delta_q \approx 0.2\%$ per quarter (SIFMA issuance), so issuance is an order of magnitude too small to negate the flow effect, and the fixed-supply assumption is without loss of generality for our qualitative results.

Finally, with variable supply the tight link between the fund's bond leg and household borrowing loosens, since part of each inflow finances newly issued shares rather than bonds. By bond-market clearing, $B_t^{MF} = B_t^P + B_t^M + B_t^G$, so the price–debt relationship under issuance is

$$p_t = \frac{1}{1+\delta_q}\,\frac{\theta_t}{1-\theta_t}\,\frac{B_t^P + B_t^M + B_t^G}{\bar{Q}}. \tag{A27}$$

As δ_q rises, the equity price impact of higher household and government indebtedness falls, but—for realistic issuance—remains first-order.

Endogenous issuance.. The dampening is weaker still once issuance responds to prices. Sammon and Shim (2026) show that firms are the marginal sellers who clear the market for passive inflows, issuing shares (largely through non-primary channels such as stock-based compensation) in response to the higher prices those inflows generate, with prices acting as the coordinating mechanism. We capture this with a reduced-form supply rule in which net issuance is increasing in the proportional price change,

$$\delta_q = \varphi\,\frac{p_t - p_{t-1}}{p_{t-1}}, \qquad \varphi \geq 0, \tag{A28}$$

where φ is the elasticity of net share supply to price. Substituting (A28) into $\frac{p_t - p_{t-1}}{p_{t-1}} = \frac{f-\delta_q}{1+\delta_q}$ and writing $g \equiv (p_t - p_{t-1})/p_{t-1}$ gives the fixed point

$$\varphi g^2 + (1+\varphi)g - f = 0 \qquad \Longrightarrow \qquad g = \frac{-(1+\varphi) + \sqrt{(1+\varphi)^2 + 4\varphi f}}{2\varphi}.$$

To first order in f, $g \approx f/(1+\varphi)$: endogenous issuance dampens the price impact of inflows by the factor $1/(1+\varphi)$ but leaves its sign unchanged. Prices rise if and only if $f > 0$ for any finite φ; only in the perfectly elastic limit $\varphi \to \infty$ does the price impact vanish, recovering the Lucas-type benchmark in which quantities, not prices, absorb flows. This is consistent with Sammon and Shim (2026): issuers accommodate part of passive demand—so realized issuance co-moves with price—yet passive-driven flows continue to move prices because supply is far from perfectly elastic.

E.4. Other Representations of Equity Price

Backward-looking representation. Writing the fund's bond leg as its initial level plus the cumulated history of net issuance, $B_t^{MF} = B_{-1}^{MF} + \sum_{s=0}^{t} \Delta B_s^{MF}$ with $\Delta B_s^{MF} \equiv B_s^{MF} - B_{s-1}^{MF}$, the price becomes

$$p_t = \frac{\theta_t}{1-\theta_t} \frac{1}{\bar{Q}} \left(B_{-1}^{MF} + \sum_{s=0}^{t} \Delta B_s^{MF} \right). \tag{A29}$$

The current price is the mandate-scaled accumulated history of borrowing by middle class, non-wealthy, and government borrowers.

The backward-looking expression highlights how the accumulated stock of borrowing maps into the present equity price. Because equity supply is fixed at $\bar{Q}$, the mandate scales the fund's entire bond leg into equity market value through the multiplier $\frac{\theta_t}{1-\theta_t}$: a debt stock B_t^{MF} carried into period t has price impact $\frac{\theta_t}{1-\theta_t}\frac{1}{\bar{Q}}$, and the same stock carried into $t+1$ has impact $\frac{\theta_{t+1}}{1-\theta_{t+1}}\frac{1}{\bar{Q}}$. The impact of a given quantity of outstanding debt therefore evolves with $\theta_{t+1} - \theta_t$, itself determined by changes in the equity premium through the mandate rule (A31). If the premium at $t+1$ exceeds that at t, the fund tilts toward equity, θ rises, and the price impact of the outstanding debt stock strengthens. This is why a disturbance that raises borrowing at $t = 0$ can become more salient when later positive dividend shocks occur: the dividend shock raises the equity premium, lifting θ and amplifying the price impact of debt already on the fund's balance sheet.[38]

Forward-looking representation. Iterating the equity-return identity $1 + r_t^E = (D_{t+1} + p_{t+1})/p_t$ forward and imposing the no-bubble condition yields

$$p_t = \mathbb{E}_t \left[\sum_{s=1}^{\infty} \frac{D_{t+s}}{\prod_{j=0}^{s-1} \left(1 + r_{t+j}^E\right)} \right], \tag{A30}$$

where the required gross return on equity embeds a time-varying, mandate-driven equity premium,

$$1 + r_t^E = 1 + r_t + \frac{\chi^{MF}}{\beta^{MF}} \frac{\theta_t - \theta^*}{\theta^*}. \tag{A31}$$

[38] It is straightforward to extend the model to allow the debt stock to decay at a 'half-life' rate, but we omit this to preserve tractability.

The difference between equity prices in the inelastic benchmark compared to the elastic benchmark is that in the former, equity premia are pinned down by quantities, while simultaneously satisfying both the backward-looking and the (more typical) forward-looking specifications. In the elastic model, prices are set by future dividends and discount rates, and equity premia is directly tied to the wealthy household's stochastic discount factor.

E.5. The Inequality Multiplier

This subsection derives the wealth share decomposition (29), the regime-specific effective equity betas, and the multiplier ratio (32) used in the discussion of the inequality multiplier in Section 2. Throughout, for any variable z_t with strictly positive deterministic steady state $\bar{z}$, the proportional deviation is $\widehat{z}_t \equiv (z_t - \bar{z})/\bar{z}$, so that $z_t \simeq \bar{z}(1+\widehat{z}_t)$.

Wealth share decomposition.. Aggregate steady-state net worth across both regimes is

$$\bar{W} = \bar{p}\bar{Q} + \bar{X}^H + \bar{B}^G, \qquad \bar{X}^H \equiv \bar{X}^{H,M} + \bar{X}^{H,P}. \tag{A32}$$

The wealthy household's wealth share deviation is

$$\begin{aligned} \eta_R^{dev} &= \frac{\bar{a}^R}{\bar{W}}\widehat{a}^R - \bar{\eta}_R\Bigg[\underbrace{\frac{\bar{p}\bar{Q}}{\bar{W}}}_{\omega_E}\widehat{p} + \underbrace{\frac{\bar{X}^H}{\bar{W}}}_{\omega_H}\widehat{p}^H + \underbrace{\frac{\bar{B}^G}{\bar{W}}}_{\omega_G}\widehat{B}^G\Bigg] \\ &= \bar{\eta}_R\Big[\widehat{a}^R - \big(\omega_E\,\widehat{p} + \omega_H\,\widehat{p}_h + \omega_G\,\widehat{b}_g\big)\Big], \end{aligned} \tag{A33}$$

which is (29).

Effective equity betas.. Differentiating the decomposition with respect to $\widehat{p}$, holding the other revaluation terms fixed, gives $\partial\eta_R^{\text{dev}}/\partial\widehat{p} = \eta_R(\beta_R^E - \omega_E) \equiv M$, with $\beta_R^E \equiv \partial\widehat{a}^R/\partial\widehat{p}$. In the elastic model, wealth is $a_t^R = E_t + B_t^R$, so holding the bond revaluation fixed,

$$\beta^E_{R,\text{el}} = \frac{\partial\widehat{a}_R}{\partial\widehat{p}} = \frac{\bar{p}\bar{Q}}{\bar{a}_R} \equiv w_R^{E,\text{el}}; \tag{A34}$$

the effective beta is exactly the wealthy's equity portfolio weight. In the inelastic model, wealth is $a_t^R = X_t^{MF}$ and the mandate ties equity and bond values together through (27), so

$$\beta^E_{R,\text{inel}} = \frac{\partial\widehat{a}_R}{\partial\widehat{p}} = 1 - \frac{\partial\widehat{\theta}}{\partial\widehat{p}} \approx 1, \tag{A35}$$

where the correction is $O(1/\chi^{MF})$ and vanishes as the mandate tightens. To hold the equity share fixed at $\bar{\theta}$ while the equity quantity $\bar{Q}$ is fixed, the fund's bond leg must move one-for-one (in logs) with the equity price. The bond leg therefore stops being a hedge: the wealthy's *entire* balance sheet—equity leg and bond leg—co-moves with $\widehat{p}$. In the elastic model the bond holding is free to move independently, which is what makes $\beta^E_{R,\text{el}} = w^E_R < \beta^E_{R,\text{inel}}$.

Multiplier ratio and comparative static.. Combining these betas with $M = \eta_R(\beta^E_R - \omega_E)$,

$$M^{\text{el}} = \eta^{\text{el}}_R\big(\beta^E_{R,\text{el}} - \omega^{\text{el}}_E\big) = \eta^{\text{el}}_R\Big(\frac{\omega^{\text{el}}_E}{\eta^{\text{el}}_R} - \omega^{\text{el}}_E\Big) = \omega^{\text{el}}_E\big(1-\eta^{\text{el}}_R\big), \tag{A36}$$

$$M^{\text{inel}} = \eta^{\text{inel}}_R\big(1-\omega^{\text{inel}}_E\big). \tag{A37}$$

Assuming a common steady state across regimes and using $\bar{\eta}_R = \bar{a}^R/\bar{W} = \omega^{\text{inel}}_E/\bar{\theta}$, the ratio of multipliers can be expressed in $(\bar{\theta}, \eta_R)$ only:

$$\frac{M^{\text{inel}}}{M^{\text{el}}} = \frac{1-\bar{\theta}\,\eta_R}{\bar{\theta}\,(1-\eta_R)},$$

which is (32). The comparative static in the target share is

$$\frac{\partial}{\partial\bar{\theta}}\left(\frac{M^{\text{inel}}}{M^{\text{el}}}\right) = -\,\frac{1}{\bar{\theta}^2(1-\eta_R)} \;<\; 0, \qquad \frac{M^{\text{inel}}}{M^{\text{el}}} \to 1 \text{ as } \bar{\theta}\to 1, \qquad \frac{M^{\text{inel}}}{M^{\text{el}}} \to \infty \text{ as } \bar{\theta}\to 0. \tag{A38}$$

Amplification is largest when the wealthy's nominal equity weight is smallest. When $\bar{\theta}\to 1$ the wealthy are nearly all-equity even in the frictionless benchmark, so the additional bond leg adds nothing and the regimes converge. When $\bar{\theta}$ is low—wealthy nominally bond-heavy—the mandate converts a large bond position into equity exposure, β^E_R jumps from $\bar{\theta}$ to 1, and the ratio blows up.

E.6. Price Elasticity between Inelastic and Elastic Models

M converts a price move into a share move. The full impulse response also needs the price move itself, $\partial\widehat{p}/\partial\varepsilon$. Both regimes share the log-linear Campbell–Shiller identity (A2). With $g \equiv \bar{p}/(\bar{D}+\bar{p})$ and $d \equiv \bar{D}/(\bar{D}+\bar{p}) = 1-g$,

$$\widehat{p}_t = g\,\widehat{p}_{t+1} + d\,\hat{D}_{t+1} - \hat{R}_{e,t}. \tag{$\star$}$$

The regimes differ only in what pins the required return $\hat{R}_e$.

E.6.1. Elastic price recursion

Solve the wealthy equity Euler (A1′), $\phi_e\widehat{p} = \delta_\zeta(\sigma_r\hat{c}_r - \Sigma_r\hat{a}_r) + \beta_R\bar{R}_e(\hat{R}_e - \sigma_r\Delta\hat{c}_{r,t+1})$, for $\hat{R}_e$ and insert into ($\star$). Collecting the $\widehat{p}$ terms gives, with $B \equiv \phi_e/(\beta_R\bar{R}_e)$,

$$(1+B)\,\widehat{p}_t = g\,\widehat{p}_{t+1} + d\,\hat{D}_{t+1} - \sigma_r\Delta\hat{c}_{r,t+1} + \frac{\delta_\zeta}{\beta_R\bar{R}_e}\big(\sigma_r\hat{c}_{r,t} - \Sigma_r\hat{a}_{r,t}\big), \tag{A39}$$

with forward solution

$$\frac{\partial\widehat{p}_t^{\,\mathrm{el}}}{\partial\varepsilon} = \frac{1}{1+B}\sum_{j\geq 0}\Big(\tfrac{g}{1+B}\Big)^j\frac{\partial}{\partial\varepsilon}\Big[\,d\,\hat{D}_{t+1+j} - \sigma_r\Delta\hat{c}_{r,t+1+j} + \tfrac{\delta_\zeta}{\beta_R\bar{R}_e}(\sigma_r\hat{c}_r - \Sigma_r\hat{a}_r)_{t+j}\Big]. \tag{A40}$$

The discount factor $g/(1+B)$ is near $g\approx 1$: the elastic price is a *forward-looking present value of fundamentals and the wealthy Euler*. There is no borrowing term.

E.6.2. Inelastic price recursion

Use A9 and the mandate rule (A1), $\bar{\theta}\hat{\theta} = \theta_{tgt}\kappa(\bar{R}_e\hat{R}_e - \bar{R}\hat{R})$ with $\kappa \equiv \beta_{MF}/\chi^{MF}$, and substitute $\hat{R}_e$ from ($\star$). With

$$A \equiv \frac{\theta_{tgt}\,\kappa\,\bar{R}_e}{\bar{\theta}(1-\bar{\theta})} = O(1/\chi^{MF}), \tag{A41}$$

one obtains

$$(1+A)\,\widehat{p}_t = \hat{b}_{mf,t} + A\,d\,\hat{D}_{t+1} + A\,g\,\widehat{p}_{t+1} - A\frac{\bar{R}}{\bar{R}_e}\hat{R}_t, \tag{A42}$$

with forward solution

$$\frac{\partial\widehat{p}_t^{\,\mathrm{inel}}}{\partial\varepsilon} = \frac{1}{1+A}\sum_{j\geq 0}\Big(\tfrac{Ag}{1+A}\Big)^j\frac{\partial}{\partial\varepsilon}\Big[\,\hat{b}_{mf,t+j} + A\,d\,\hat{D}_{t+1+j} - A\tfrac{\bar{R}}{\bar{R}_e}\hat{R}_{t+j}\Big]. \tag{A43}$$

Because $A = O(1/\chi^{MF})$ (at the baseline $\chi^{MF} = 171$, $\kappa \approx 0.006$, A small), the discount factor $Ag/(1+A)\approx 0$ and

$$\frac{\partial\widehat{p}_t^{\,\mathrm{inel}}}{\partial\varepsilon} \;\approx\; \frac{\partial\hat{b}_{mf,t}}{\partial\varepsilon} \;+\; \underbrace{O(1/\chi^{MF})\ \text{fundamental terms}}_{\text{via the }\hat{\theta}\text{ channel}}. \tag{A44}$$

By bond-market clearing (D2), $\hat{b}_{mf} = \sum_i s_i\,\hat{b}_i$ with $s_i = \bar{B}_i/\bar{B}_{mf}$ ($i = m, p, g$). The inelastic price is driven, nearly contemporaneously, by aggregate borrowing, with only a small forward fundamental leakage through $\hat{\theta}$.

E.7. The Multiplier in General Equilibrium

M in (31) is a *partial* derivative: it holds $\widehat{p}_h$ and $\hat{b}_g$ fixed and isolates the equity channel. This is the portfolio heterogeneity object, but it is not the impulse-response coefficient. In general equilibrium the house price and (in the elastic model) the wealthy bond holding respond to the same shock that moves $\widehat{p}$, so the *total* derivative is

$$\text{(A45)} \qquad \frac{d\eta_R^{\text{dev}}}{d\widehat{p}} = \underbrace{\eta_R(\beta_R^E - \omega_E)}_{M\ \text{(partial)}} - \underbrace{\eta_R\,\omega_H\,\frac{\partial \widehat{p}_h}{\partial \widehat{p}}}_{\text{housing dilution}} - \underbrace{\eta_R\omega_G\,\frac{\partial \hat{b}_g}{\partial \widehat{p}}}_{\text{govt. debt adjustment}} + \underbrace{\omega_B\,\frac{\partial \hat{b}_r}{\partial \widehat{p}}}_{\substack{\text{elastic only}\\ (=0\ \text{inelastic})}} .$$

The middle term is unambiguous in sign: the wealthy hold no housing, so if house prices co-move with equity ($\partial\widehat{p}_h/\partial\widehat{p} > 0$, the usual boom case) the denominator $\bar{W}$ appreciates through a component the wealthy do not own, diluting their share. The partial M therefore overstates the equilibrium equity-to-inequality pass-through whenever housing co-moves with equity. The government debt term matters for a $\hat{b}_g$ shock which makes both $\hat{b}_g$ and $\widehat{p}$ move along the impulse response.

Next, what is the gap between wealth shares across the two models? Taking the common steady-state assumption ($\eta_R, \omega_E, \omega_H, \omega_G$ are taken equal across regimes), write (29) out in each regime, restoring the elastic bond term.

Inelastic ($\hat{a}_r = \widehat{p} - \hat{\theta}$):

$$\text{(A46)} \qquad \eta_R^{\text{dev,inel}} = \eta_R(1-\omega_E)\,\widehat{p}^{\text{inel}} - \eta_R\hat{\theta} - \eta_R\omega_H\,\widehat{p}_h^{\text{inel}} - \eta_R\omega_G\,\hat{b}_g.$$

Elastic ($\hat{a}_r = (\omega_E\widehat{p} + \omega_B\hat{b}_r)/\eta_R$):

$$\text{(A47)} \qquad \eta_R^{\text{dev,el}} = \omega_E(1-\eta_R)\,\widehat{p}^{\text{el}} + \omega_B\,\hat{b}_r - \eta_R\omega_H\,\widehat{p}_h^{\text{el}} - \eta_R\omega_G\,\hat{b}_g.$$

Subtracting (the exogenous $\hat{b}_g$ cancels):

$$\text{(A48)} \qquad \eta_R^{\text{dev,inel}} - \eta_R^{\text{dev,el}} = \underbrace{\eta_R(1-\omega_E)\widehat{p}^{\text{inel}} - \omega_E(1-\eta_R)\widehat{p}^{\text{el}}}_{\text{(I) equity revaluation}} \underbrace{-\,\omega_B\,\hat{b}_r}_{\text{(II) elastic bond}} \underbrace{-\,\eta_R\omega_H(\widehat{p}_h^{\text{inel}} - \widehat{p}_h^{\text{el}})}_{\text{(III) housing}} - \eta_R\hat{\theta}.$$

This reveals three structural sources of the gap:

a. Concentration / loss of hedge. Even at equal prices $\widehat{p}^{\text{inel}} = \widehat{p}^{\text{el}}$, the coefficient is larger in the inelastic regime:

$\eta_R(1-\omega_E) > \omega_E(1-\eta_R)$ as long as $\eta_R > \omega_E$, i.e. the wealthy share of wealth is higher than the equity share of total wealth. This is purely β_R^E: the mandate makes the wealthy a levered equity claim. It alone delivers $M^{\text{inel}} > M^{\text{el}}$ for any shock that raises $\widehat{p}$ in both regimes (dividends, fundamentals).

b. Elastic bond revaluation. $-\omega_B \hat{b}_r$ has no inelastic analogue: the elastic-model wealthy, as creditors, take a direct mark-to-market on their bonds when aggregate borrowing moves. This makes the regime diverge through shocks on the credit side (e.g. a non-wealthy endowment shock that makes them de-lever, $\hat{b}_r < 0$).

c. price response difference. $\widehat{p}^{\text{inel}}$ vs $\widehat{p}^{\text{el}}$ from (A40) and (A43): borrowing-driven and contemporaneous vs fundamentals-driven and forward-looking, as well as the differential house price response across both models. Both depend upon the source of the fundamental shock — for example, a positive dividend shock raises $\widehat{p}^{\text{el}}$ more than $\widehat{p}^{\text{inel}}$, but a positive government borrowing shock has the opposite effect.

Appendix F. Dynamics of Wealth Inequality: Additional Results

This appendix collects results referenced from Section 5: the elastic-versus-inelastic wealth shares under unanticipated shocks, and the market aggregates under the unanticipated-shock simulation. The simulations with the arbitrageur are collected in Appendix H.

F.1. Simulations under Unanticipated Shocks: Comparison to Elastic Model

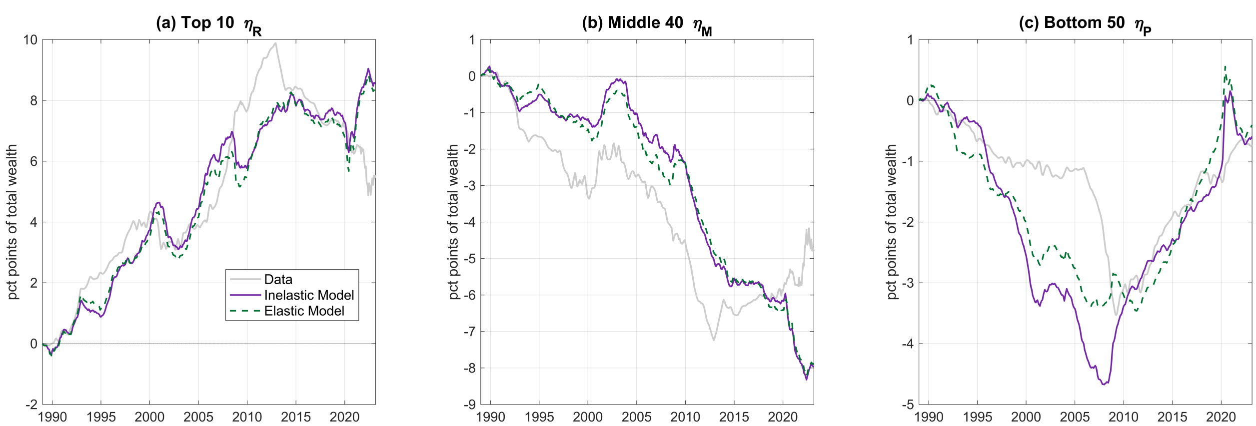

FIGURE A6. Perfect foresight with unanticipated shocks: elastic vs. inelastic wealth shares

Note: Wealth shares under the unanticipated-shock simulation for the inelastic (solid purple) and elastic (dashed green) models, against the data (grey), for the top 10%, middle 40% and bottom 50%. The figure plots the percentage-point change in wealth shares relative to their 1989 level.

Under unanticipated shocks the inelastic and elastic wealth share paths nearly coincide: both reach a top 10% share about 8.5 percentage points above its 1989 level, against the data's peak of roughly ten. This is a marked change from perfect foresight, where the inelastic model separated clearly from the elastic one (a top 10% rise of 5.4 against 3.4 percentage points). The common, sizeable borrowing surge that unanticipated shocks introduce drives both economies, so the wealth shares alone no longer discriminate between the two pricing regimes. The discriminating evidence is the perfect-foresight exercise of Figure 3, where borrowing is muted and only the pricing mechanism differs.

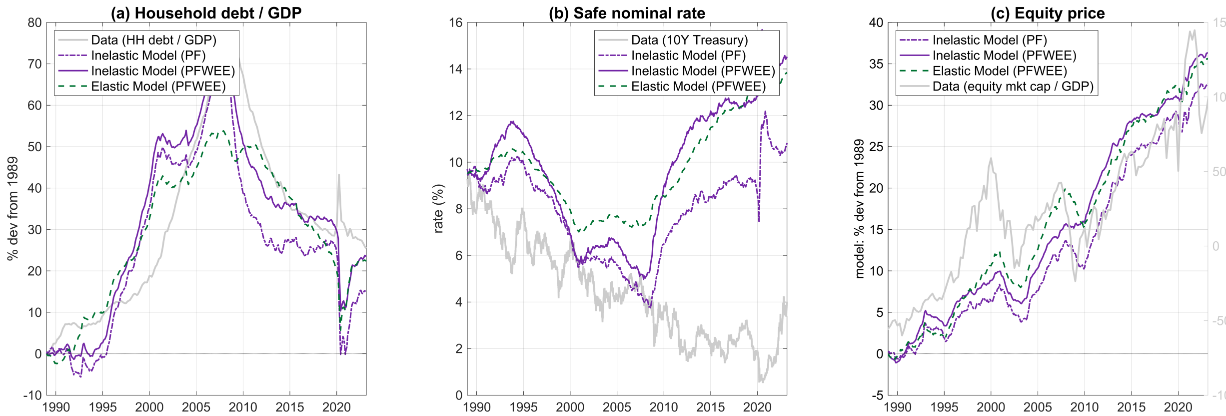

FIGURE A7. Perfect foresight with unanticipated shocks: borrowing, interest rate, and equity price

Note: As Figure 4, but for the unanticipated-shock simulation. Panel (a) household debt to GDP; panel (b) the safe nominal rate; panel (c) the equity price, with the model on the left axis and the data on the right, both in percentage deviations from their 1989 level. Each panel plots the inelastic model under perfect foresight (dash-dot purple) and under unanticipated shocks (solid purple), the elastic model under unanticipated shocks (dashed green), and the data (grey). The elastic model reproduces the post-2008 borrowing dynamics only with a substantially larger swing in household debt and the safe rate.

The market aggregates tell the same story. Relative to perfect foresight, unanticipated shocks spare non-wealthy households the large anticipatory deleveraging that otherwise suppresses debt—bottom 50% borrowing ends near its 1989 level rather than some twenty percent below it—lifting household debt to GDP onto the hump observed in the data (panel a). Both pricing regimes reproduce this path, and both track the dynamics of the equity price while understating its level (panel c). Neither captures the post-2010 'low-for-long' rate: with public debt fed in exogenously, the post-crisis supply of bonds pushes the modelled safe rate up rather than down (panel b), as discussed in Section 5. Across all three aggregates the inelastic and elastic paths remain close under unanticipated shocks, reinforcing that the two regimes are separated not by these series but by the elasticity of prices to flows.

Appendix G. Additional Model Impulse Responses

The flow shock of Section 4 is one of eight structural shocks in the quantitative model. Figures A8 and A9 report the two whose transmission is most diagnostic of the borrowing channel: a government-borrowing shock and a middle-class loan-to-value shock. In both, the inelastic and elastic models diverge in *sign*. A rise in borrowing—whether from public debt or looser household collateral—*raises* the equity price in the inelastic model, because the mandate routes the additional credit supply into equity; in the elastic model the same shock *lowers* the equity price through the discount-rate channel. This sign reversal is the borrowing-channel counterpart of the inequality multiplier. The remaining shock (the equity dividend) and a full by-variable decomposition across all shocks are collected in Supplemental Appendix 3.

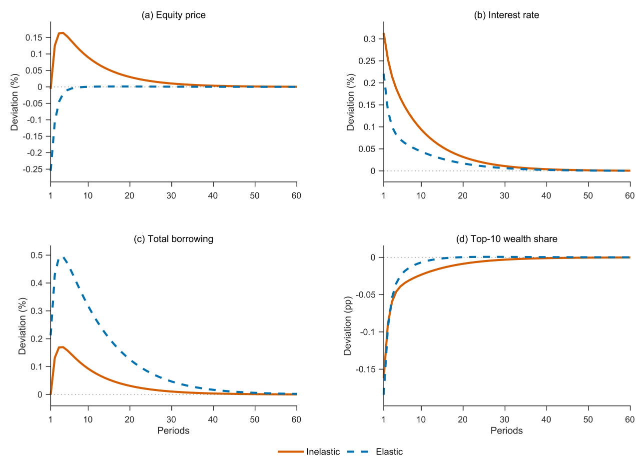

FIGURE A8. Impulse responses to a government-borrowing shock

Response to a one-standard-deviation innovation to government borrowing (B_t^G), inelastic (solid red) versus elastic (dashed blue), over 60 months. Panel (a) equity price; (b) interest rate; (c) total borrowing; (d) top 10% wealth share. Panels (a)–(c) are percentage deviations from steady state; panel (d) is in percentage points.

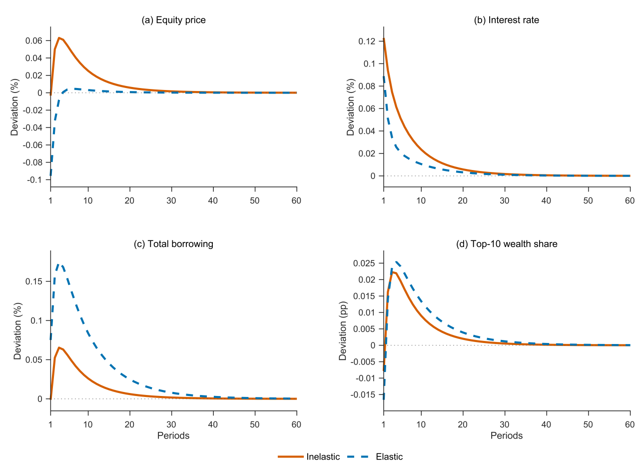

FIGURE A9. Impulse responses to a middle class loan-to-value shock

Response to a one-standard-deviation innovation to the middle class loan-to-value ratio (τ_t^M), inelastic (solid red) versus elastic (dashed blue), over 60 months. Panels as in Figure A8.

Appendix H. Model with Arbitrageur

In the baseline model, we have assumed that there is no arbitrageur to drive prices back towards fundamental values. In this section, we introduce one and trace how it modifies the results. As we shall see, a finite-capacity arbitrageur attenuates but does not remove the inelasticity channel unless its risk-bearing capacity is implausibly large.

H.1. Arbitrageur in the Inelastic Model

At time period t, the financial intermediary is restricted in its portfolio choice. We introduce an arbitrageur that is ultimately owned by the wealthy household but operates as a separate vehicle with predetermined risk capital W_t^A.[39] It finances its equity position in an external funding market at the domestic risk-free rate, so that its safe-asset leg does not enter domestic bond-market clearing.[40]

The arbitrageur chooses equity shares q_t^A to maximise CARA utility over next-period wealth,

$$\max_{q_t^A} \mathbb{E}_t\big[-\exp(-\gamma W_{t+1}^A)\big], \qquad W_{t+1}^A = (1+r_t)W_t^A + q_t^A\big[D_{t+1} + p_{t+1} - (1+r_t)p_t\big],$$

where γ is the arbitrageur's risk aversion. Since $(1+r_t)W_t^A$ is predetermined and CARA utility is translation-invariant, the choice depends only on the trading profit and loss,

$$\Pi_{t+1}^A \equiv q_t^A\big[D_{t+1} + p_{t+1} - (1+r_t)p_t\big],$$

so the objective is equivalently $\max_{q_t^A} \mathbb{E}_t[-\exp(-\gamma\Pi_{t+1}^A)]$.

Write the equity premium as:

$$\pi_{t+1} \equiv r_t^e - r_t = \frac{D_{t+1} + p_{t+1}}{p_t} - (1+r_t).$$

The premium is a function of the model's state variables $S_t = (e_t^R, e_t^M, e_t^P, D_t, D_t^H, B_t^G, \tau_t^M, \zeta_t^d)$. Because the model is linearised around the steady state, the premium is approximately affine in the (approximately Gaussian) state vector and is therefore conditionally normal to first order,

$$(\pi_{t+1} \,|\, \mathcal{F}_t) \sim \mathcal{N}\big(\mu_t(S_t),\ \sigma_t^2(S_t)\big),$$

where $\mathcal{F}_t$ is the σ-algebra generated by the history of states, and we make explicit the dependence of the

[39] This segmentation is essential. Because the arbitrageur is not bound by the fund mandate, unrestricted contemporaneous injections of capital by the wealthy would let them circumvent the mandate through the vehicle and undo the inelasticity; this amounts to the assumption that the wealthy cannot alter its risk-bearing capacity within the period, i.e. the arbitrageur is 'ring-fenced'. The vehicle is externally financed, and we do not consolidate its trading profit into the wealthy household's flow budget or Euler equation. Ultimately for the wealth distribution, however, the arbitrageur is owned by the wealthy: the reported top share credits them with the residual float $\bar{Q}(1-\kappa_t)$ they and their fund hold directly *and* the $\bar{Q}\kappa_t$ shares the vehicle holds on their behalf, valued at $\bar{p}\bar{Q}\kappa_t$. This is an ownership attribution rather than the output of a fully consolidated balance sheet, and with a zero steady-state position ($\bar{\kappa} = 0$) it leaves the deterministic steady state undisturbed.

[40] The alternative is domestic financing, in which the arbitrageur's bond position $B_t^A = W_t^A - p_t q_t^A$ enters bond-market clearing and hence feeds back into the safe rate, household borrowing, and the mandate's bond channel. This is more internally consistent for a domestic vehicle but changes the model's interest-rate and borrowing predictions: leaning against the mandate's mispricing would then push the safe rate and household borrowing toward their elastic market values as well as the price, so the arbitrageur would damp the debt-driven divergence too and compress the wealth share gap by more than under external financing. We adopt the external-financing benchmark here to keep the pricing relation transparent.

conditional mean and variance on S_t.

Given CARA utility and normality of returns, the arbitrageur's maximisation function can be re-written as:

$$\max_{q_t^A} q_t^A p_t \mu_t - \frac{1}{2}\gamma (q_t^A)^2 p_t^2 \sigma_t^2,$$

resulting in the familiar first-order condition:

$$q_t^A = \frac{\mu_t}{\gamma p_t \sigma_t^2}.$$

The expected premium $\mu_t = \mathbb{E}_t(\pi_{t+1})$ deviates from its normal level when the price departs from its correctly-priced benchmark. That benchmark is the price that would prevail in a frictionless market—one free of the mandate and portfolio demand distortions but bearing the same equity risk. This is precisely the price of the elastic direct-holding benchmark of Appendix 2.8, in which the wealthy hold equity directly and price it through their stochastic discount factor at the required return, the risk-free rate plus the equity premium r^x. We therefore take the fundamental to be this elastic market price,

$$p_t^f \equiv p_t^{\mathrm{el}},$$

which shares the common deterministic steady state, $\bar{p}^f = \bar{p}$, so that the normal excess return is exactly the premium, $\bar{\mu} = r^x$, and the arbitrageur corrects deviations of the price from its frictionless value—the inelasticity and demand distortions.[41] In the quantitative model we feed the simulated elastic-benchmark price directly as p_t^f. Deviations of the expected premium from its normal level then decompose as

$$\mu_t - \bar{\mu} = (1 + r_t + r^x)\left(\frac{p_t^f - p_t}{p_t}\right) + \underbrace{\mathbb{E}_t\left[\frac{p_{t+1} - p_{t+1}^f}{p_t}\right]}_{\lambda_t \text{ (expected future mis-pricing)}},$$

where λ_t is an endogenous expected-return term, the discounted expected future mis-pricing. We set $\lambda_t = 0$, so that the arbitrageur is a myopic one-period investor trading on current mis-pricing alone. This is an approximation: with persistent demand shocks ($\rho_\zeta = 0.9$), some mis-pricing is expected to persist into the next period. The arbitrageur's active demand—its position relative to the zero it holds when correctly

[41] A risk-*neutral* reference (discounting at $1 + r_t$) would instead place $\bar{p}^f > \bar{p}$ and lead the arbitrageur to trade against the entire equity premium: it would read equity as permanently cheap and lean *into* rather than against the cycle. Anchoring on the elastic price, which embeds the required return, isolates the distortion the arbitrageur is meant to correct.

priced—is then

$$q_t^A = \frac{(1 + r_t + r^x)(p_t^f - p_t) + \lambda_t p_t}{\gamma p_t^2 \sigma_t^2}.$$

Meanwhile, the mandate-constrained mutual fund's demand is (as in Equation 4):

$$q_t^{MF} = \frac{\theta_t X_t^{MF}}{p_t}.$$

Imposing market clearing:

$$\bar{Q} = q_t^{MF} + q_t^A = \frac{\theta_t X_t^{MF}}{p_t} + q_t^A \implies p_t = \frac{\theta_t X_t^{MF}}{\bar{Q} - q_t^A}.$$

Equivalently, the price can be written as

$$p_t = \underbrace{\frac{\theta_t}{1-\theta_t} \frac{B_t^{MF}}{\bar{Q}}}_{\text{Price with no arbitrageur}} \frac{1}{1-\kappa_t},$$

where we have used that $X_t^{MF} = \frac{B_t^{MF}}{1-\theta_t}$ and bond-market clearing $B_t^{MF} = B_t^P + B_t^M + B_t^G$ (which holds because the arbitrageur funds its position externally), and define κ_t to be the arbitrageur's *position share* of the equity float:

$$\kappa_t \equiv \frac{q_t^A}{\bar{Q}} = \frac{(1 + r_t + r^x)(p_t^f - p_t) + \lambda_t p_t}{\gamma \bar{Q} p_t^2 \sigma_t^2}.$$

The position κ_t is an equilibrium object, and its sign is the opposite of current mis-pricing (with $\lambda_t \approx 0$). When inflows into the mutual fund push the price above fundamentals the arbitrageur sells, $\kappa_t < 0$, and the factor $1/(1 - \kappa_t) < 1$ dampens the price impact of flows; the reverse holds for outflows. At correct pricing ($p_t = p_t^f$, $\lambda_t \approx 0$) the arbitrageur holds nothing, $\kappa_t = 0$, and the restricted-intermediary identity (28) is recovered exactly. The strength of the attenuation is set by the risk-bearing *capacity* $1/(\gamma p_t \sigma_t^2)$, which falls with γ and, under CARA, is independent of the arbitrageur's wealth W_t^A: balance sheet size alone does not constrain the position without a leverage or margin constraint.

A closed-form price makes the dependence on risk-bearing capacity explicit. Market clearing requires

$$\bar{Q} = \frac{\theta_t X_t^{MF}}{p_t} + \frac{(1 + r_t + r^x)(p_t^f - p_t) + \lambda_t p_t}{\gamma p_t^2 \sigma_t^2}.$$

Cross-multiplying and taking the positive root of the resulting quadratic yields the pricing function

$$p_t = \frac{\gamma\,\sigma_t^2\,\theta_t X_t^{MF} - \left(1 + r_t + r^x - \lambda_t\right) + \sqrt{\left[\gamma\,\sigma_t^2\,\theta_t X_t^{MF} - (1 + r_t + r^x - \lambda_t)\right]^2 + 4\gamma\,\sigma_t^2 \bar{Q}\,(1 + r_t + r^x)\,p_t^{\mathrm{f}}}}{2\gamma\,\sigma_t^2 \bar{Q}}.$$

When there is no arbitrageur, $\gamma \to \infty$, and $p_t = \frac{\theta_t X_t^{MF}}{\bar{Q}}$, exactly as in the baseline model. When there is unlimited arbitrage, $\gamma \to 0$, and $p_t = \frac{(1+r_t+r^x)\,p_t^f}{1+r_t+r^x-\lambda_t}$, which, for $\lambda_t \approx 0$, gives $p_t = p_t^f$; arbitrageurs drive prices to fundamentals. However, it requires implausibly high risk-bearing capacity in arbitrageurs to fully attenuate the effects of flows on price. With a realistically low calibrated arbitrage capacity, the model and our qualitative findings continue to hold, while our quantitative findings are dampened very slightly. Crucially, this price convergence is not convergence of the economy: even as $a \to \infty$ and the equity price is driven exactly to its elastic value p_t^f, household borrowing and the safe rate are unchanged, because the externally-financed vehicle takes no position in the domestic bond market and so leaves the mandate's bond channel intact. The wealth share distribution therefore need not collapse to its elastic market counterpart—the debt-driven component of the divergence survives even when the price gap is fully closed.

What disciplines the channel is capacity, not ownership. The arbitrageur is a marginal stabilising trader: it holds nothing in steady state and takes a temporary long or short position only when a shock opens a valuation gap. We impose a zero steady-state net position, $\bar{q}^A = 0$ (equivalently $\bar{\kappa} = 0$), so that the constrained fund continues to hold the entire equity supply and the deterministic steady states of both regimes are left unchanged. The arbitrageur then responds to deviations of the expected excess return from its normal level,

$$q_t^A = \frac{\mu_t - \bar{\mu}}{\gamma\,\bar{p}\,\bar{\sigma}^2}, \qquad \bar{q}^A = 0,$$

which linearises to the reduced-form schedule $\kappa_t = a(\hat{p}_t^f - \hat{p}_t)$ used in the next subsection, with $\hat{p}_t^f = \hat{p}_t^{\mathrm{el}}$ the deviation of the elastic-benchmark price and $\bar{p}^f = \bar{p}$ the common steady-state reference.[42] The capacity a is common to both regimes, so the comparison isolates how the same residual risk-bearing capacity propagates shocks under the two market structures.

The empirical scale of natural arbitrageurs enters only as motivation for a finite—and plausibly low—capacity. Gabaix and Koijen (2021) find that hedge funds and similar players hold under 5% of the market and sell into rather than lean against crises, and Gabaix et al. (2025) measure the reallocation of equity risk across

[42] The normal excess return is the constant premium, $\bar{\mu} = r^x > 0$, so an exact CARA investor facing the *level* premium μ_t would hold a positive position even in steady state. Trading on the deviation $\mu_t - \bar{\mu}$—the change in expected excess return relative to the normal premium—nets this out, leaving the steady state undisturbed. Because the reference is the frictionless elastic price ($\bar{p}^f = \bar{p}$), the deviation form $\kappa_t = a(\hat{p}_t^f - \hat{p}_t)$ references the common steady-state price and isolates the arbitrageur's marginal stabilising role.

investors to be an order of magnitude smaller than heterogeneous-investor models imply. We do not use this ownership share to pin a steady-state position—which remains zero—because average ownership does not identify the elasticity of active arbitrage demand: an arbitrageur can hold nothing on average yet trade large temporary positions, and a large passive owner may do no stabilising trading at all. We instead read the 5% figure as a ceiling on *active* capacity and calibrate $a = 0.075$ so that the arbitrageur's mean gross position along the transition, $\frac{1}{T}\sum_t |\kappa_t|$ over 1989–2023, equals 5% in the inelastic model. This disciplines how hard the arbitrageur leans against mis-pricing without tying its balance sheet to a passive ownership share. At this capacity $|\kappa_t|$ stays small and the effect of flows on prices is attenuated but far from removed; only implausibly large capacity (very small γ) makes $|\kappa_t|$ large enough to eliminate the price impact of flows.

H.2. Arbitrageur in the Elastic Model

The elastic direct-holding benchmark of Appendix 2.8 already prices equity at the fundamental: the wealthy absorb the whole float, $q_t^R = \bar{Q}$, and price it through their stochastic discount factor, so the equilibrium price is exactly p_t^f. The arbitrageur therefore faces no mispricing and holds nothing, $\kappa_t = a(\hat{p}_t^f - \hat{p}_t) = 0$, at any capacity. Its role is confined to the inelastic regime, where the mandate holds the price $\hat{p}_t^{\text{noarb}}$ of (28) away from the fundamental and the schedule $\kappa_t = a(\hat{p}_t^f - \hat{p}_t)$ closes a fraction $a/(1+a)$ of the gap.[43] Were the elastic market instead to admit its own departure from fundamentals—a residual demand shock the wealthy did not fully absorb on price—the identical schedule would correct it; in our benchmark no such wedge arises, so the arbitrageur is dormant there by construction.

H.3. Quantitative Effect of the Arbitrageur

Figure A10 traces the economy's response to the flow demand shock, comparing the inelastic model with and without the arbitrageur against the elastic benchmark. In the inelastic regime the arbitrageur's position share κ_t stays small—a few percent of the float—so the price gap $\hat{p}_t - \hat{p}_t^f$ it closes is minor (panels a–c), and the with-arbitrageur paths for the equity price, safe rate, total borrowing, and top 10 wealth share barely depart from the no-arbitrageur inelastic paths plotted alongside (panels a, d–f). In the elastic regime the arbitrageur is inert, $\kappa_t = 0$, because the price already equals the fundamental. A realistically-calibrated arbitrageur therefore attenuates, but does not overturn, the inelasticity channel.

[43] Both regimes share the reduced form $\hat{p}_t = (\hat{p}_t^0 + a\,\hat{p}_t^f)/(\varepsilon + a)$, a capacity-weighted average of the no-arbitrageur price $\hat{p}_t^0$ and the fundamental: under direct holding $\hat{p}_t^0 = \hat{p}_t^f$ with the wealthy's elasticity ε_R, so the arbitrageur is inert; under the mandate $\hat{p}_t^0 = \hat{p}_t^{\text{noarb}}$ with unit-elastic value demand $\varepsilon = 1$. Because $\varepsilon = 1 < \varepsilon_R$ (Appendix E.1) the same capacity bites harder under the mandate, leaving a larger residual mispricing there.

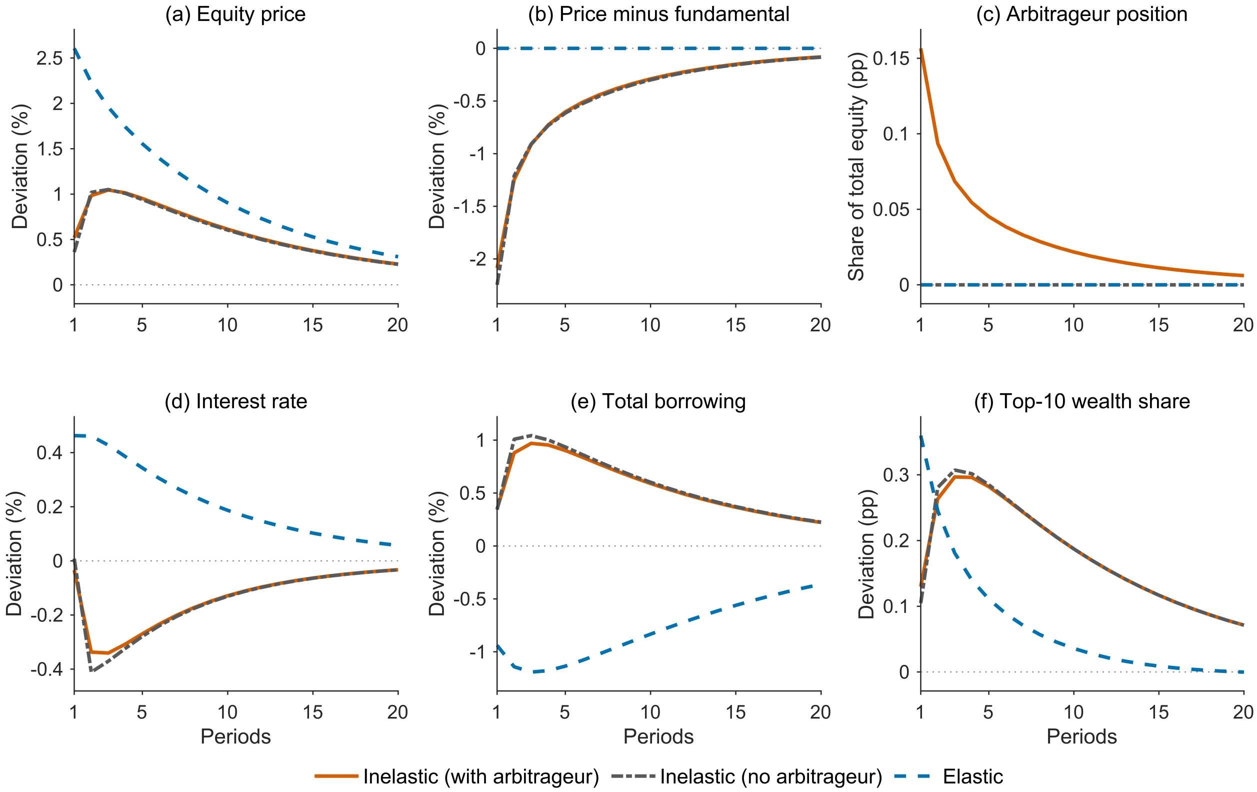

FIGURE A10. Flow demand shock with the arbitrageur

Response to a positive one-standard-deviation (10%) latent demand shock (ζ_t^d): the inelastic model with the arbitrageur at its calibrated capacity (solid), the inelastic model without the arbitrageur (grey dash-dot), and the elastic model (dashed), on which the arbitrageur is inert. Top row: (a) equity price; (b) price minus the fundamental, $\hat{p}_t - \hat{p}_t^f$; (c) arbitrageur position share κ_t. Bottom row: (d) safe interest rate; (e) total household borrowing; (f) top 10 wealth share. Panels (a), (b), (d), (e) are percentage deviations; (c) is in percentage points of the equity float; (f) in percentage points. In the elastic regime the price equals the fundamental, so panel (b) and the position in panel (c) are zero; the with- and no-arbitrageur inelastic paths nearly coincide, the small gap between them being the arbitrageur's attenuation.

H.4. Simulations with Arbitrageur

We repeat the perfect-foresight transition of Section 5 with the arbitrageur active at its calibrated capacity, feeding the elastic-benchmark price in as the fundamental so that the vehicle trades only the inelasticity wedge. Figure A11 plots the resulting wealth shares and Figure A12 the corresponding market aggregates.

The arbitrageur attenuates the inelasticity channel without overturning it. Leaning against the mandate-driven mispricing pulls the inelastic equity price partway toward the elastic price, so the top 10% share rises by slightly less than in the no-arbitrageur transition of Figure 3: the 2023 gap between the inelastic and elastic top shares narrows only modestly—from about two percentage points to just under two—and the bottom 50% improvement is likewise dampened but preserved. The wealth share advantage of inelastic markets therefore survives a realistically-calibrated residual trader.

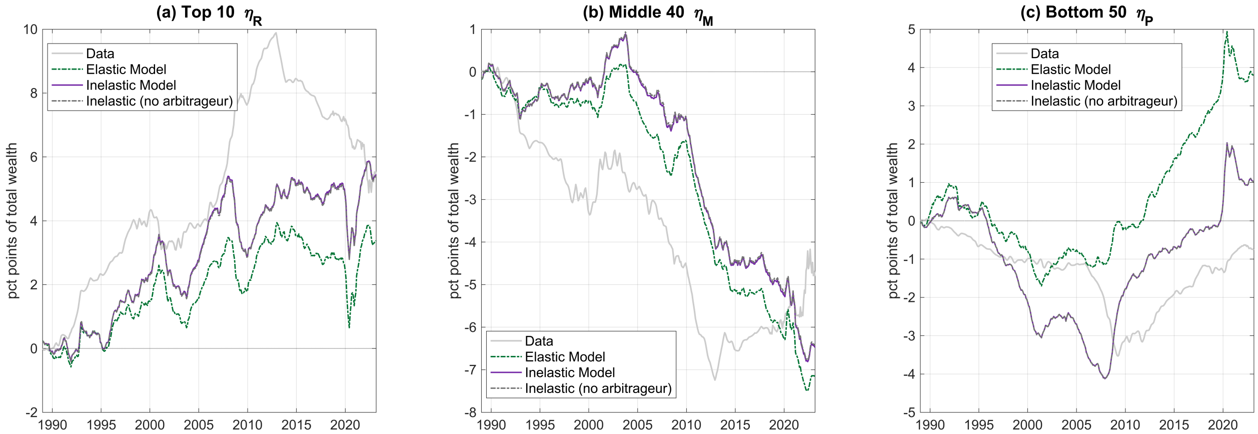

FIGURE A11. Change in wealth shares from 1989: perfect foresight with arbitrageur

Note: As Figure 3, but with the arbitrageur active at its calibrated capacity. The figure plots the percentage-point change in wealth shares relative to their 1989 level. The light-grey line plots the empirical wealth shares, the purple line the inelastic model with the arbitrageur, the grey dash-dot line the inelastic model without the arbitrageur, and the green line the elastic model (on which the arbitrageur is inert). Empirical wealth share data is taken from Realtime Inequality at the household level (Blanchet et al. 2022).

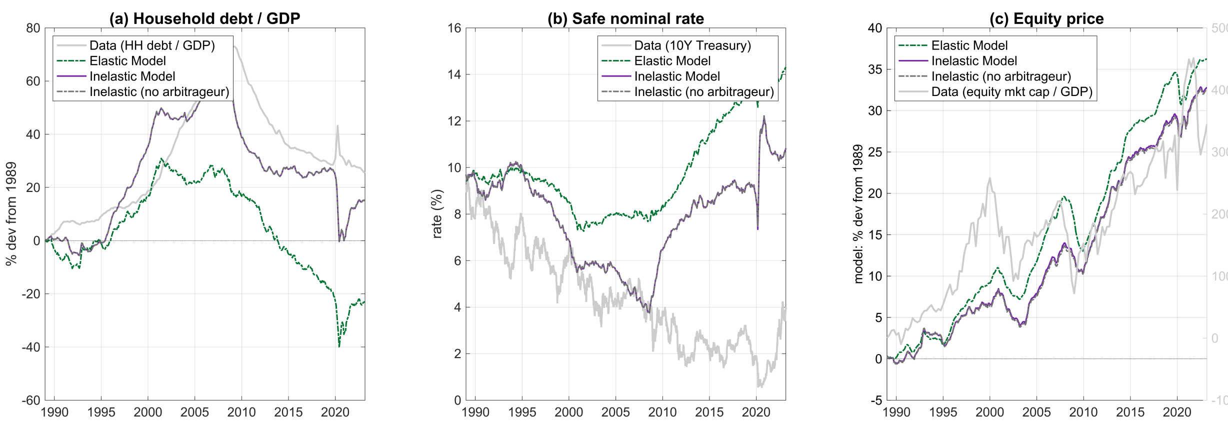

FIGURE A12. Data vs. simulation with arbitrageur: borrowing, interest rate and equity price under perfect foresight

Note: As Figure 4, but with the arbitrageur active at its calibrated capacity, and adding the inelastic model without the arbitrageur (grey dash-dot). Panel (a) household debt to GDP; panel (b) the safe interest rate; panel (c) the equity price, plotted on a secondary axis because the model underpredicts the level of the price run-up. Data lines in light grey.

Figure A12 shows why the channel is so resilient. The arbitrageur funds its equity position externally, so its trades never enter domestic bond-market clearing: the household debt-to-GDP and safe-rate paths are essentially unchanged from the no-arbitrageur transition and remain sharply different across the two regimes. The vehicle compresses only the equity price gap, and only partially. Because the wealth share divergence between inelastic and elastic markets runs primarily through the borrowing channel—the mandate forcing additional debt onto non-wealthy households—rather than through the equity price level, an arbitrageur that touches only the price leaves the mechanism substantially intact.

Appendix I. Construction of Residualised Equity Flows

We construct a residualised measure of aggregate fund flows from mutual fund and ETF data, use it to instrument aggregate equity returns, and estimate the resulting dynamic responses of household wealth and wealth shares using local projections.

I.1. Aggregate fund flows

Let F_t denote quarterly, not-seasonally-adjusted transactions in total financial assets of the consolidated mutual fund and exchange-traded-fund sector, as reported in the Federal Reserve's *Financial Accounts of the United States* (series `FU484090005`). Let V^E_{t-1} denote the lagged market value of U.S. corporate equities held across all sectors (series `LM893064105`). Following Gabaix and Koijen (2021), we define aggregate fund flows as

$$f_t \equiv 100\,\frac{F_t}{V^E_{t-1}}, \tag{A49}$$

expressed as a percentage of lagged equity-market value. The sample runs from 1989Q1 to 2023Q4.

Raw fund flows are persistent and may respond endogenously to past returns, cash-flow news, discount rates, and changes in aggregate risk. To isolate variation not accounted for by these observable factors, we estimate

$$f_t = \alpha_f + \mathbf{b}'_f X_t + u_t, \tag{A50}$$

where

$$X_t \equiv \left(f_{t-1},\ldots,f_{t-4},\ \Delta d_t, \Delta d_{t-1}, \Delta d_{t-2},\ pd_{t-1},\ y^{10}_t, \Delta y^{10}_t, \Delta y^{10}_{t-1},\ \sigma^m_{t-1}, \Delta\sigma^m_t, \Delta\sigma^m_{t-1},\ r^m_{t-1},\ \widetilde{t}\right)'. \tag{A51}$$

Here, Δd_t denotes quarterly log dividend growth, pd_{t-1} the lagged log price–dividend ratio, y^{10}_t the ten-year Treasury yield, σ^m_t realised market volatility over the preceding four quarters, r^m_{t-1} the lagged market return, and $\widetilde{t}$ a standardised linear trend. Price, dividend, and yield data are obtained from Shiller (2023).

The four lags of f_t absorb the persistence of delegated fund flows. Dividend growth and the price–dividend ratio control for cash-flow and valuation news; the Treasury-yield terms capture shifts in the relative attractiveness of bonds and equities; the volatility terms proxy for changes in aggregate risk; and the lagged market return accounts for return-chasing. The linear trend removes any secular trends, so that the residual variation is identified from fluctuations around this trend. We deliberately exclude the distributional outcomes and their lags from (A50), allowing the same residualised innovation to be used for every wealth group. We define this innovation as

$$z_t \equiv \widehat{u}_t. \tag{A52}$$

When reporting responses directly to z_t, we standardise the series to have unit variance.

Our procedure draws on the identification logic of Gabaix and Koijen (2021), but does not reproduce their granular instrument. We residualise an aggregate flow series with respect to a broad set of observed common factors. The residualisation removes variation associated with the included observables, but does not by itself establish the exclusion restriction required for a causal interpretation.

I.2. Returns and the first stage

We construct quarterly equity returns following Gomez (2025), using the Goyal–Welch data set (Welch and Goyal 2008). Let P_t denote the end-of-quarter market index, $D_t^{(4)}$ dividends paid over the preceding four quarters, and $R_t^{f,\text{simple}}$ the quarterly simple risk-free return. The log market return, log risk-free return, and log excess return are

$$r_t^m = \log\left(\frac{P_t + D_t^{(4)}/4}{P_{t-1}}\right), \tag{A53}$$

$$r_t^f = \log\left(1 + R_t^{f,\text{simple}}\right), \tag{A54}$$

$$r_t^e = r_t^m - r_t^f. \tag{A55}$$

Dividing $D_t^{(4)}$ by four assigns one quarter of trailing annual dividends to the quarterly holding period. Our baseline specification uses the excess return r_t^e. As robustness checks, we use the total market return r_t^m and the Shiller price-only return, $\Delta \log P_t$, while holding fixed the flow innovation, outcomes, controls, and local-projection specification (Shiller 2023). If anything, alternative definitions of returns serve to push results further in our direction.

We estimate the first-stage relationship

$$r_t^e = \alpha_r + b z_t + \Lambda_r' C_t + \nu_t, \qquad C_t \equiv \left(\tilde{t}, y_t^{10}, \Delta y_t^{10}\right)'. \tag{A56}$$

The vector C_t contains the controls retained in the return decomposition and the subsequent local projections.

Although z_t was constructed by residualising fund flows with respect to the richer conditioning set X_t, we include C_t again so that the return decomposition and local projections are estimated using the same sample and conditioning variables.

Let $\widetilde{r}_t^e$ and $\widetilde{z}_t$ denote the residuals from separate OLS regressions of r_t^e and z_t on a constant and C_t. By the Frisch–Waugh–Lovell theorem, the first-stage coefficient can be written as

$$\widehat{b} = \frac{\sum_t \widetilde{z}_t \widetilde{r}_t^e}{\sum_t \widetilde{z}_t^2}. \tag{A57}$$

We decompose the residualised excess return into the component predicted by the flow innovation,

$$r_t^F \equiv \widehat{b}\widetilde{z}_t, \tag{A58}$$

and the remaining component,

$$r_t^O \equiv \widetilde{r}_t^e - r_t^F. \tag{A59}$$

It follows that

$$\widetilde{r}_t^e = r_t^F + r_t^O, \qquad \sum_t r_t^F r_t^O = 0. \tag{A60}$$

We refer to r_t^F as the flow-induced return: it is the component of aggregate excess returns predicted by the residualised flow innovation. The component r_t^O contains the remaining return variation, including variation associated with cash-flow news, discount-rate news, aggregate risk, and measurement error.

I.3. Distributional outcomes

Household balance sheet data are from the Federal Reserve's *Distributional Financial Accounts* (DFA). For the two-group analysis, we aggregate the Top 0.1 per cent, Next 0.9 per cent, and Next 9 per cent bins to construct top-decile net worth. We combine the Next 40 per cent and Bottom 50 per cent bins to construct net worth for the bottom 90 per cent. For either group, the wealth share is

$$\eta_{g,t} \equiv \frac{W_{g,t}}{W_{\text{Top}\,10,t} + W_{\text{Bottom}\,90,t}}, \qquad g \in \{\text{Top}\,10,\ \text{Bottom}\,90\}. \tag{A61}$$

The two shares sum to one by construction.

Following Gomez (2025), our baseline net-worth outcome measures cumulative wealth growth in excess of the risk-free return. For group g and horizon h, it is

$$Y^W_{g,t}(h) \equiv 100\left[\log\left(\frac{W_{g,t+h}}{W_{g,t-1}}\right) - \sum_{j=0}^{h} r^f_{t+j}\right]. \tag{A62}$$

As a robustness check, we use cumulative net-worth growth without subtracting the risk-free return — this does not qualitatively change our results.

We measure changes in wealth shares in log points:

$$Y^S_{g,t}(h) \equiv 100\left[\log\eta_{g,t+h} - \log\eta_{g,t-1}\right]. \tag{A63}$$

I.4. Local projections

For each group, outcome, and horizon $h = 0, \dots, 12$ quarters, we estimate

$$Y_{g,t}(h) = \alpha_{g,h} + \beta^F_{g,h} r^F_t + \beta^O_{g,h} r^O_t + \Lambda'_{g,h} C_t + \varepsilon_{g,t+h}, \tag{A64}$$

where $Y_{g,t}(h)$ denotes one of the net-worth, wealth share, or balance sheet outcomes defined above. We estimate (A64) separately for each outcome and horizon, following Jordà (2005). The projections are estimated on the common quarterly sample, which begins in 1989Q3—the first quarter for which the Distributional Financial Accounts are available—so that, while the flow series spans 1989Q1–2023Q4, the local projections use regressor dates from 1989Q3 onward. Standard errors are Newey–West with a Bartlett kernel and bandwidth $h + 1$. The bandwidth increases with the horizon because the cumulative outcomes induce serial correlation in the regression residuals.

The coefficient $\beta^F_{g,h}$ measures the response to the component of equity returns predicted by the residualised flow innovation. We report coefficients per unit of log return. First-stage relevance is assessed using the estimated effect of z_t on excess returns, the corresponding first-stage F statistic, and the share of residualised return variation explained by the flow innovation.

Figure A13 reports the net-worth analogue of the wealth share decomposition shown in the introduction. Here the outcome is excess net-worth growth, $Y^W_{g,t}(h)$ from (A62). For both groups the fundamentals-driven component produces the larger *level* of net-worth growth, while the flow-driven component fades. This is consistent with the share result: fundamentals move both groups' net worth broadly together, so they redistribute little, whereas the flow-driven component—though smaller in levels—acts asymmetrically across groups and is therefore what persistently moves the wealth *share*.

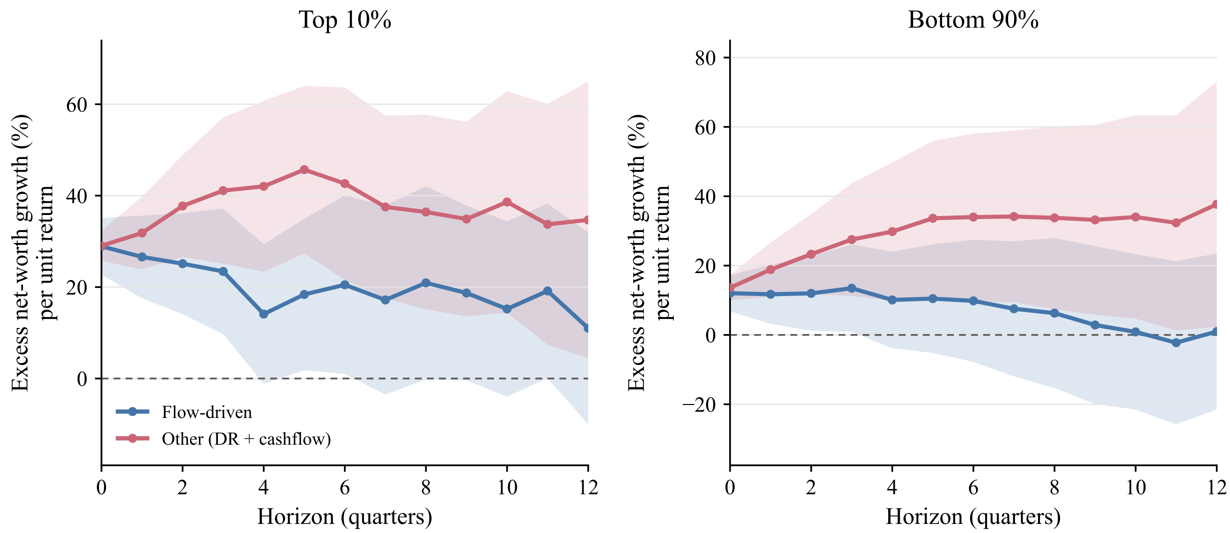

FIGURE A13. Response of excess net-worth growth to the flow-driven and fundamentals-driven components of equity returns, top 10% (left) and bottom 90% (right).

Note: Local projections of excess net-worth growth $Y^{W}_{g,t}(h)$ from (A62) on the flow-driven component (blue) and the residual 'other' component (red) of Gomez (2025)'s excess equity return, per unit of log return, at horizons of 0–12 quarters. Shaded bands are 90% Newey–West confidence intervals. Net worth is from the Distributional Financial Accounts, quarterly, 1989Q3 onward.