Asset Pricing with Heterogeneous Agents under Limited Information*

Riccardo A. Cioffi† Samuel Hurtado‡ Galo Nuño‡

July 2026

Abstract

In this paper, we study the interaction between wealth inequality and asset prices by developing a framework featuring heterogeneous agents, endogenous portfolio choice, and aggregate risk. We solve the model globally under the assumption of limited information: agents forecast prices directly using observable aggregate variables, which makes the problem computationally tractable without sacrificing the economic mechanism. We approximate the drift and volatility of the perceived price process with neural networks and find that, while the drift component of the evolution of prices is close to linear, the volatility is highly nonlinear. This asymmetry endogenously generates time variation in risk premia that local approximation methods cannot produce and that is key to explaining the interaction between inequality and asset returns in our model: through heterogeneous portfolio exposure, aggregate shocks reshape the wealth distribution, which in turn feeds back into asset prices and asset demand through market clearing.

1 Introduction

Over recent decades, developed economies have witnessed exceptionally high returns to equity, often coupled with unprecedented increases in wealth concentration. Understanding the dynamic relationship between asset prices and wealth inequality is one of the central challenges in modern macroeconomics and finance: on the one hand, in the presence of heterogeneous portfolio exposure, asset price movements have differential effects on households'

*We thank Benjamin Moll, Jonathan Payne, Gianluca Violante, and Yucheng Yang for helpful comments and discussions. We also thank seminar participants at the Paris Macro Group, SED meetings, UCL, the University of Lausanne, and the University of Surrey.

†Paris School of Economics.

‡Bank of Spain, CEMFI and CEPR.

wealth accumulation; on the other, changes in the distribution of wealth alter aggregate asset demand and, consequently, equilibrium prices. Together, these two channels create a feedback loop in which movements in asset returns affect portfolio allocation and the wealth distribution, which in turn feeds back into asset demand and, through market clearing, asset prices. Existing frameworks have struggled to provide a convincing approach to study these interactions in a unified environment, because doing so requires heterogeneity, endogenous portfolio choice, and aggregate risk – a combination that has proven particularly complicated to handle jointly.

In this paper, we directly tackle this challenge: we develop a global-solution framework for studying the joint dynamics of wealth inequality and asset prices in an economy with heterogeneous agents, aggregate risk, and endogenous portfolio choice. Rather than requiring agents to track the infinite-dimensional wealth-income distribution, we restrict our attention to a limited information environment in which we posit that agents forecast asset prices *directly* by assuming they follow a diffusion process where both the drift and volatility functions only depend on observable aggregate variables. We then approximate these functions using neural networks and exploit their universal approximation properties to capture the complex nonlinear relationship between prices and the distribution without imposing restrictive functional forms on the resulting price process. The model delivers a self-confirming equilibrium which completely sidesteps the Master equation, making the problem computationally tractable without sacrificing the economic mechanism: the interaction between heterogeneous portfolio exposure, asset prices, and the wealth distribution remains at the core of our model and endogenously generates time-varying risk premia – as asset prices move, they reshape not only the distribution of wealth but also the compensation investors require for holding risk.[1]

Model. We study a continuous-time endowment economy in which households facing idiosyncratic income risk optimally allocate savings between a risk-free bond and a risky asset – a Lucas tree whose price is determined endogenously in equilibrium. Uninsurable earnings risk and an exogenous borrowing limit generate heterogeneity in wealth and, importantly, in households' exposure to equity risk: a single aggregate shock to dividends then directly reshapes the cross-sectional wealth distribution, making it stochastic.

To better understand why such a model is particularly challenging, consider what is required for solving it: households making endogenous portfolio choices across the wealth distribution, aggregate risk that directly reshapes the distribution through heterogeneous

[1]Such time variation in risk premia is a direct consequence of explicitly modeling return risk: because we model aggregate risk as *directly* affecting household wealth through portfolio returns, the process that governs the evolution of the cross-sectional distribution features a diffusion component that introduces nonlinearities otherwise absent from standard approaches.

portfolio holdings, and endogenous asset prices determined by market clearing. Under rational expectations, households recognize that prices are equilibrium objects, determined by the distribution through market clearing: forecasting prices means forecasting the distribution, an infinite-dimensional state variable which enters the value function and turns the recursive problem into a Master equation in the space of distributions. In a setting like ours, where aggregate shocks *directly* affect households' wealth, this requirement is particularly severe: households would need to understand not just the distribution's expected changes but also its *stochastic dynamics*, how shocks reshape it through heterogeneous portfolio exposure. Recent computational advances have made progress in approximating the Master equation, but usually keep agents tracking the entire cross-sectional distribution;[2] an assumption that, as Moll (2024) argues, raises fundamental conceptual concerns above and beyond its computational viability: "real-world households and firms do not forecast prices by forecasting cross-sectional distributions and instead solve simpler problems." We follow this insight and instead model households as forecasting prices directly.[3]

Our model, like most heterogeneous-agent models, exhibits what Moll and Ryzhik (2025) call "low-dimensional coupling": prices are the only channel through which the distribution enters individual decisions. Even though – as they also point out – low-dimensional coupling alone provides no simplification (the price process inherits the distribution's infinite dimensionality), it does suggest a natural alternative to the Master equation approach: modeling households as forecasting prices directly from *observable* state variables – in our model, the current dividend and the current price – a formalization of what real-world investors actually observe and act upon.[4] The question, of course, is *how* agents then forecast prices, which we tackle by approximating the drift and volatility of the perceived price process with two separate neural networks. With the distribution out of the state space, the household problem is now finite-dimensional: given a guess for the perceived law of motion, we can therefore solve it using "standard" continuous-time methods, simulate the stochastic Kolmogorov forward equation to track the distribution, clear the asset market at each instant to obtain realized prices, retrain the networks on the simulated price series, and iterate until convergence.

Equilibrium. In the resulting self-confirming equilibrium (Fudenberg and Levine 1993; Cho and Sargent 2018), we show that agents' forecasts are statistically consistent with the price dynamics that actually materialize and that forecast errors across independent simu-

[2] See Bilal (2023) for a perturbation approach and Gu et al. (2023) for one using neural networks.

[3] Our model is specifically designed to isolate the mechanisms linking wealth inequality and asset prices under limited information, not to match specific moments of the data; we discuss this more in detail in section 3.2.

[4] In our model we deliberately keep the risk-free rate exogenous to isolate the feedback between equity prices and the wealth distribution; endogenizing it means clearing an additional market at every instant – a relatively straightforward but computationally demanding extension we leave for future work.

lations are approximately normally distributed: given their available information, no agent observing the realized path of prices would have reason to revise their forecasts. Importantly, we never claim that this solution approximates the rational expectations equilibrium; rather, we study how agents with limited information form beliefs about prices. Indeed, because the Master equation has no known solution in this setting, a direct comparison is infeasible at the scale of our model and it is this intractability that motivates our limited information approach.

Additionally, we also find that in the ergodic region each dividend level maps almost one-to-one into a price. If any other variable (e.g. the wealth distribution or past prices) mattered for prices the ergodic set would map a given level of dividends to multiple market-clearing prices; the fact that they do not means that future prices depend exclusively on current dividends and prices, which supports our assumption that asset prices are indeed Markov.

Having households that make optimal portfolio decisions is central to the equilibrium mechanism. Absent substantive trading frictions, the optimal portfolio is a step function: households hold a 100% equity share when expected excess returns are high (i.e., when prices are low relative to dividends), and no equity when they are low, with an extremely narrow transition region separating the two corner solutions.[5] By construction, market clearing pushes the equilibrium into precisely this narrow region – at either corner, excess demand or supply would move prices until portfolio choice turns interior – while the wealth distribution keeps evolving continuously and stochastically, tying the price to the moving distribution at every instant. The problem is therefore fundamentally nonlinear and poorly suited to local approximation methods. Additionally, the step function makes second-derivative terms in the household's Hamilton-Jacobi-Bellman (HJB) equation highly sensitive to numerical precision. This is one key reason we solve the household problem with standard finite-difference methods and reserve the neural networks only for the perceived law of motion.

Price dynamics. We find that the learned price dynamics display a clear asymmetry between drift and volatility: the drift is approximately linear – prices above their long-run level tend to fall, those below tend to rise – precisely the mean-reversion that a linear approximation in the spirit of Krusell and Smith (1998) would capture well; while the volatility is instead highly nonlinear, near zero when prices are low and rising as prices increase, with substantial curvature even within the ergodic region.[6] The reason is that the volatility cap-

[5] We do not allow for short-selling of bonds or stocks.

[6] Notice that the predictability implied by a mean-reverting drift in prices is not an inefficiency that households can exploit: the drift is already part of households' perceived law of motion, so portfolio decisions already condition on it, and the associated expected excess returns are compensation for bearing state-dependent risk, which is the source of the time-varying risk premia documented in section 4.

tures how dividend shocks propagate through the economy as a function of its current state: high asset prices require high equity demand, which means more households are exposed to aggregate fluctuations, so that a given shock hits more portfolios at once, amplifying its transmission to prices and raising volatility. In fact, risk premia are time-varying precisely because price volatility moves with the aggregate state, and so does the compensation required by equity holders. Sorting ergodic states by their price-dividend ratio, we find that the premium is largest in the middle of the range. Specifically, it is about 0.5 p.p. in the highest-ratio states, 2.5 p.p. in the middle of the range, and 0.8 p.p. in the lowest-ratio states.

These premia shape how the economy absorbs a shock: a negative dividend shock moves the average portfolio share through two channels. First, a *valuation channel*, which captures the shock's direct effect on equity values: in response to a fall in dividends asset prices fall, which mechanically make households wealth and portfolio shares fall; second, a *portfolio rebalancing channel* which kicks in because the fall in prices raises expected excess returns and reduces perceived risk: households rebalance their portfolio toward equity, which pushes equity prices back up and makes portfolio shares climb back. The average share eventually overshoots its pre-shock level before falling back, because prices recover toward their pre-shock value relatively quickly while aggregate wealth remains depressed for longer.

Repeating the same experiment from different points along the ergodic path, we show that the price response is strongly state-dependent, with a cumulative price decline that is roughly two and a half times larger when the shock hits high-price, high-dividend states – where equity-market participation is broadest – than when it hits low-price states. Importantly, this state dependence arises naturally from our global, nonlinear solution, a pattern that local or linear approximation methods cannot generate.

Distributional feedback. Another key finding is that, in equilibrium, households' portfolio choice depends almost entirely on the aggregate state. The optimal portfolio step function described above is common to every household and the transition region does not depend on wealth or income. This mainly follows from the simplicity of our model economy: CRRA preferences, limited heterogeneity in income, and a portfolio problem that is frictionless apart from the no-borrowing and short-sale constraints. Such uniformity has two main consequences: first, the asset price is a (state-dependent) fraction of aggregate wealth. Second, portfolio heterogeneity is therefore driven by exposure levels, $\theta^* a$, rather than shares: households hold the same portfolio share of different amounts of wealth.

Nonetheless, the wealth distribution differs systematically across the different states that the economy visits along its ergodic path: aggregate wealth runs about 20% higher in high-price, high-dividend states than in low-price, low-dividend states. A relationship that can

also be read in the other direction: states with high prices are states that carry a larger mass of wealthy households, and the distribution feeds back into the price through the market clearing condition – higher aggregate wealth means higher equity demand, a higher market value for the risky asset, and higher exposure to aggregate risk which makes asset prices more volatile.

Then, in response to the same negative aggregate dividend shock, each household's wealth falls in proportion to its exposure level. As a result, mass accumulates near the borrowing constraint, and the right tail thins as wealthier households – who carry the largest equity exposure in levels – absorb most of the decline in equity values. This two-way feedback between the distribution and prices is precisely the mechanism that motivates the paper.

The rest of the paper is organized as follows: Section 2 presents the model and the limited information framework, section 3 the neural-network methodology and calibration, section 4 the results, and section 5 concludes.

Related Literature. Our paper is connected to three main strands of literature: wealth inequality and asset prices; limited information, learning, and alternatives to rational expectations; and neural networks and machine learning in macroeconomics.

A large empirical literature documents that asset returns vary systematically with wealth, driven in part by portfolio composition (Fagereng et al. 2020; L. Bach et al. 2020; Hubmer, Halvorsen, et al. 2023).[7] Kuhn et al. (2020) show that the resulting portfolio asymmetry – the top 10% holds more than 90% of stocks while middle-class wealth is concentrated in housing – has shaped the U.S. wealth distribution through a race between housing and stock market returns; recent work quantifies this channel in heterogeneous-agent economies (Cioffi 2021; Greenwald et al. 2023; Fernández-Villaverde and Levintal 2024).

Time-varying expected returns are a foundational empirical fact in asset pricing (Campbell and Shiller 1988; Fama and French 1988; Cochrane 2011), which representative-agent models usually generate through non-standard preferences or long-run consumption risks under Epstein-Zin utility (Bansal and Yaron 2004; Campbell and Cochrane 1999, among others). Instead, our model generates them through distributional dynamics: the wealth distribution is a state variable for equilibrium prices, as in the heterogeneous-agent tradition (e.g. Chan and Kogan 2002; Chien and Lustig 2010; Chien, Cole, et al. 2012; Dumas 1989; Gârleanu and Panageas 2015; Guvenen 2009; Storesletten et al. 2007; Toda and Walsh 2019).

Hubmer, Krusell, et al. (2021) find that return heterogeneity is the single largest driver of the level of the U.S. top-10% wealth share, but in their model portfolio shares and asset

[7] Evidence also suggests that risk tolerance increases with wealth (Calvet and Sodini 2014; Meeuwis 2020). See Benhabib and Bisin (2018) for a survey of the mechanisms generating wealth inequality.

prices are both exogenous, so myopia and perfect foresight produce nearly identical wealth distributions (roughly 2 percentage points apart). Endogenizing both is, in their words, "one or two orders of magnitude more challenging"; while our model is quantitatively less rich, we take precisely those steps: portfolio choice derives from household optimization under limited information, so that beliefs drive portfolios, portfolios drive prices, and prices feed back into beliefs. Gomez (2025) embeds portfolio heterogeneity in a two-agent continuous-time general equilibrium model, generating a feedback loop between wealth concentration and asset prices, while our model derives portfolio choice endogenously across the entire distribution.

Relatedly, Krueger and Lustig (2010) formalize conditions under which incomplete markets are irrelevant for the equity premium: irrelevance breaks down under countercyclical cross-sectional income variance or under limited participation – the latter only when participation constraints are imposed exogenously (Gomes and Michaelides 2008). To break these irrelevance conditions, a strand of literature pairs wealth heterogeneity with belief dispersion: Martin and Papadimitriou (2022) show analytically that wealth redistribution among heterogeneous-belief agents generates time-varying sentiment, while Belda et al. (2025) estimate a 400 basis point gap in expected stock returns between the bottom 50% and top 1% of the wealth distribution, with belief heterogeneity substantially amplifying the top-1% share. Instead, we break the irrelevance through a third channel, without any belief dispersion: even under common beliefs, agents' perceived price dynamics generate wealth-dependent portfolio distortions.

Greenwood and Shleifer (2014) show that survey expectations of stock returns are negatively correlated with model-based expected returns, with a correlation of -0.57; De La O and Myers (2021) reverse the Campbell-Shiller decomposition under subjective expectations and find that the cash-flow component dominates while discount-rate variation nearly vanishes – the opposite of the rational expectations benchmark. Beyond the empirical evidence, Moll (2024) argues that rational expectations is computationally prohibitive in heterogeneous-agent settings – agents must forecast the infinite-dimensional wealth distribution because prices are non-Markovian – so the departure is not merely a modeling choice but a response to the conceptual impossibility of requiring agents to track objects that even the econometrician cannot observe. Moll and Ryzhik (2025) formalize how, in mean field games with low-dimensional coupling, departures from rational expectations make it possible to replace the infinite-dimensional Master equation with a standard finite-dimensional Hamilton-Jacobi-Bellman (HJB) equation in which prices are the relevant state variable.

Our limited information framework builds on internal rationality (Adam and Marcet 2011): agents optimize given dynamically consistent subjective beliefs, without knowledge of objective probabilities. Adam, Marcet, and Nicolini (2016) and Adam, Marcet, and Beutel

(2017) show that even a representative-agent model with constant-gain learning matches price-dividend ratio volatility and return predictability under standard CRRA preferences. Our equilibrium concept is a self-confirming equilibrium (Fudenberg and Levine 1993; Cho and Sargent 2018), also interpretable as a restricted perceptions equilibrium (RPE) (Branch 2006): agents condition on a subset of available information, forming a perceived law of motion whose fixed point follows the T-mapping of Marcet and Sargent (1989).[8]

On the computational side, the perceived-law-of-motion approach of Krusell and Smith (1997) and Krusell and Smith (1998) relies on approximate aggregation, which fails when aggregate risk *directly* influences wealth evolution and distributional effects require substantially richer representations of the cross-sectional density. Local methods – sequence-space Jacobians (Auclert, Bardóczy, et al. 2021; Auclert, Rognlie, et al. 2024), discrete-time perturbation (Bayer and Luetticke 2020; Bhandari et al. 2023), and other linearization techniques (Bilal 2023; Reiter 2009) – are by construction insufficient for capturing the nonlinearities that generate time-varying risk premia; Lee (2025) develops a global nonlinear solution in sequence space that handles aggregate uncertainty and portfolio choice without linearization, but maintains rational expectations and the requirement that agents track the aggregate state. Building on the continuous-time framework of Achdou et al. (2022), we provide a global solution method for a heterogeneous-agent economy in which aggregate risk loads *directly* on household wealth through endogenous portfolio exposure, which makes the cross-sectional density itself stochastic: the Kolmogorov forward equation acquires a diffusion term, and the associated Master equation a second-order distributional derivative – features absent from the settings of Bilal (2023) and Gu et al. (2023). To our knowledge, ours is the first quantitative implementation of the Moll (2024) and Moll and Ryzhik (2025) program in this class of models: rather than tracking the distribution globally under rational expectations, as Lee (2025) does, agents forecast prices directly and learn both the drift and the diffusion of their perceived price process. In contemporaneous work, Yang et al. (2025) implement the same program through structural reinforcement learning – learning policies directly from simulated price paths without estimating a perceived law of motion – in economies without a portfolio margin.

Fernández-Villaverde, Hurtado, et al. (2023) show that neural-network perceived laws of motion reveal equilibrium features invisible to log-linear methods in a continuous-time heterogeneous-agent model with financial frictions. We extend their framework from single-asset intermediation to two-asset household portfolio choice with endogenous prices; additionally, while they approximate the equilibrium law of motion with the first moment of the wealth distribution, we model agents' actual beliefs about price dynamics. The math-

[8] Guarda (2025) develops a nonparametric generalization of the RPE concept (NRPE), applied to income perception in a production economy without risky assets.

ematical foundations rest on mean field game theory (Lasry and Lions 2007) and universal approximation (Hornik et al. 1989; Barron 1993; F. Bach 2017); several other papers develop deep learning solutions for heterogeneous-agent models (Azinovic et al. 2022; Azinovic-Yang and Žemlička 2023; Gopalakrishna et al. 2026; Maliar et al. 2021; Gu et al. 2023; Han et al. 2022).

The combination of these three strands allows us to study a channel that requires all three ingredients: portfolio heterogeneity across the wealth distribution, amplified by self-referential price dynamics under limited information, generates time-varying risk premia under standard CRRA preferences.

2 Model

2.1 Environment

2.1.1 Households

We consider a continuous-time endowment economy populated by a unit mass of *ex ante* identical households who face both idiosyncratic earnings risk and aggregate uncertainty through fluctuations in dividends and asset prices.

Demographics. To ensure a stationary cross-sectional distribution in the presence of wealth accumulation, we follow the approach of Blanchard (1985) and Yaari (1965) by assuming that households face a constant probability of death $\phi > 0$, implying an expected lifetime of $1/\phi$ years. Upon death, households are immediately replaced by newborn agents with zero wealth and labor productivity drawn from the ergodic distribution, which we describe in more detail below. We assume the existence of a perfect annuity market so that when agents die their asset holdings get redistributed to living households in proportion to their wealth.[9]

Preferences. Households have time-separable preferences defined over consumption streams, with utility function:

$$\mathbb{E}_0 \left[\int_0^\infty e^{-(\rho+\phi)t} u(c_t)\, dt \right] \tag{1}$$

where $\rho > 0$ is the rate of time preference. We assume that the period utility function $u(c)$ has the usual CRRA specification $u(c) = c^{1-\gamma}/(1-\gamma)$ with $\gamma > 0$ representing the coefficient of relative risk aversion.

[9]Specifically, households can purchase actuarially fair life insurance that pays a return ϕ on their wealth holdings in exchange for transferring all assets to the insurance company upon death. To simplify notation, we include the annuity adjustment directly into asset returns.

Idiosyncratic Earnings Process. Households receive stochastic labor income z_t that follows a two-state Markov chain $z_t \in \{z_1, z_2\}$ as in Achdou et al. (2022). Specifically, we assume $z_1 < z_2$, which can therefore be interpreted as unemployment and employment states, respectively. The transition rates are given by λ_1 (unemployment to employment) and λ_2 (employment to unemployment). The household's earnings state therefore evolves according to:

$$\mathbb{P}(z_{t+dt} = z_j | z_t = z_i) = \begin{cases} \lambda_i dt & \text{if } i \neq j \\ 1 - \lambda_i dt & \text{if } i = j \end{cases} \tag{2}$$

for $i, j \in \{1, 2\}$. The ergodic distribution of this process, which we denote $\pi(z)$, assigns measure $\pi(z_1) = \lambda_2/(\lambda_1 + \lambda_2)$ to the unemployed state and $\pi(z_2) = \lambda_1/(\lambda_1 + \lambda_2)$ to the employed state. We normalize the process such that total income (including dividends) equals 1.

This parsimonious specification of income risk captures the essential features of labor market dynamics – employment transitions and income heterogeneity – while maintaining analytical and computational tractability. The two-state structure allows us to introduce idiosyncratic earnings risk and to match key labor market statistics (unemployment rates and transition probabilities) while avoiding the additional computational burden associated with discretizing more complicated continuous income processes.

2.1.2 Assets and Dividends

Households have access to two assets that they can use to smooth consumption and accumulate wealth. The first is a Lucas tree – representing the aggregate capital stock of the economy – which generates a stochastic dividend stream $\{Y_t\}_{t \geq 0}$ and trades at the endogenously determined price Q_t. We assume that there is exactly one unit of this tree. The second asset is a risk-free bond available in infinitely elastic supply at the exogenously given world interest rate $r^* > 0$.[10] Importantly, we assume that households face a no-borrowing constraint in both assets: they cannot hold negative positions in either the tree or the risk-free asset. Additionally, we impose a short-sale constraint so that the portfolio share θ in the risky asset must satisfy $\theta \in [0, 1]$ – aside from such constraints, households are otherwise free to choose their portfolio share at any point in time.

The dividend process Y_t follows a mean-reverting diffusion in logs. Specifically, we assume that $y_t = \log(Y_t)$ evolves according to an Ornstein-Uhlenbeck process:

$$dy_t = \eta \left(\bar{y} - y_t \right) dt + \sigma dW_t \tag{3}$$

[10]The assumption that the safe interest rate is exogenous allows us to solve for a single price, which keeps the relation between inequality and prices more direct. Extending our framework to solve for two asset prices is conceptually straightforward, but it does complicate the computational side.

where W_t is a standard Brownian motion representing aggregate uncertainty. The process in equation (3) is the continuous-time analogue of an AR(1) process with long-run mean $\bar{y} = \log \bar{Y}$, mean-reversion coefficient $\eta > 0$, and constant volatility $\sigma > 0$. We will often also make use of the same process in levels, which we will simply rewrite as

$$dY_t = \mu_Y(Y_t)dt + \sigma_Y(Y_t)dW_t \tag{4}$$

where μ_Y and σ_Y can be found by a straightforward application of Itô's lemma:

$$\mu_Y(Y) = Y\left[\eta\left(\log \bar{Y} - \log Y\right) + \frac{1}{2}\sigma^2\right] \tag{5}$$

$$\sigma_Y(Y) = \sigma Y. \tag{6}$$

2.2 Limited Information

Before introducing our limited information framework, it is useful to first formulate the household problem under the full-information rational-expectations (FIRE) assumption to illustrate the challenges that motivate our approach.

2.2.1 Rational Expectations and the Master Equation Problem

In a FIRE equilibrium, households know that prices are equilibrium objects and that, as such, they depend on the cross-sectional distribution. Hence, in order for their decisions to be rational, households need to have expectations about the evolution of prices that are consistent with their actual evolution – which means that, in order to forecast future prices, households have to correctly forecast the evolution of the cross-sectional distribution. This requires the latter, an infinite-dimensional object, to enter as a state variable in the individual optimization problem. Under the rational expectations assumption, one can therefore write a so-called "Master equation" – an infinite-dimensional equation in the space of distributions – which encapsulates the evolution of the distribution directly into the households' problem.

It is particularly instructive to derive the Master equation in our context; not because we are interested in solving it but because it will allow us to illustrate what challenges we face when solving our model under limited information.[11]

Market Clearing. Let G_t denote the cross-sectional probability measure over wealth and income, with $G_t \in \mathcal{P}(\mathbb{R}_+ \times \{z_1, z_2\})$, and let $g_t(a, z)$ denote the corresponding density, where

[11] Since we will not actually solve the Master equation, our derivation is somewhat informal. For a formal derivation of the Master equation as well as conditions for the existence and uniqueness of its solution, see Cardaliaguet et al. (2019) and Carmona and Delarue (2018).

a denotes individual wealth. The asset price is determined by market clearing:

$$Q_t = \int_{\mathbb{R}_+ \times \{z_1, z_2\}} \hat{\theta}^*(a, z; Y_t, g_t) a \, g_t(a, z) \, da \, dz \equiv \mathcal{Q}(Y_t, g_t) \tag{7}$$

where $\hat{\theta}^*(\cdot)$ is the optimal portfolio choice decision of households under rational expectations.[12] This equation directly establishes that Q_t is a function of the dividend level Y_t and the distribution g_t.

Stochastic Distribution Evolution. Since our economy features a single aggregate shock W_t that drives the dividend process in equation (4), and this same shock directly affects individual wealth through households' portfolio exposure, the cross-sectional distribution itself becomes stochastic.

Specifically, the evolution of the distribution is governed by the following stochastic Kolmogorov forward equation (KFE):[13]

$$dg_t(a, z) = \hat{\mathcal{A}}_t^*[g_t](a, z) \, dt + \hat{\mathcal{B}}_t^*[g_t](a, z) \, dW_t \tag{8}$$

where the drift component $\hat{\mathcal{A}}_t^*[g]$ captures the deterministic flow of households across wealth and income states due to optimal consumption, portfolio choice, income transitions, and demographic turnover (with $\psi(a, z) = \delta(a) \, \pi(z)$ the birth distribution: a point mass at $a = 0$ scaled by the ergodic income weights $\pi(z)$); and the diffusion component $\hat{\mathcal{B}}_t^*[g]$ represents how the aggregate noise term dW_t induces stochastic changes in the distribution through heterogeneous portfolio exposure.[14]

Specifically, the operator $\hat{\mathcal{A}}_t^*$ acts on the density as:

$$\begin{aligned} \hat{\mathcal{A}}_t^*[g](a, z) = &- \frac{\partial}{\partial a} \left[\hat{\mu}_a(a, z; Y_t, g_t) g(a, z) \right] + \frac{1}{2} \frac{\partial^2}{\partial a^2} \left[\hat{\sigma}_a^2(a, z; Y_t, g_t) g(a, z) \right] \\ &+ \lambda_{z'} g(a, z') - \lambda_z g(a, z) + \phi \left(\psi(a, z) - g(a, z) \right) \end{aligned} \tag{9}$$

[12] We will henceforth use hats to indicate variables and functions in the rational expectations equilibrium.

[13] See Appendix A for a derivation of the KFE (28) below.

[14] With some abuse of notation we use the subscript t in $\hat{\mathcal{A}}_t^*$ and $\hat{\mathcal{B}}_t^*$ to denote the dependence of these operators on the current state of the economy (here given by (Y_t, g_t)).

where $\hat{\mu}_a(\cdot)$ and $\hat{\sigma}_a(\cdot)$ are the drift and diffusion coefficients of individual wealth:

$$
\begin{aligned}
da_{i,t} &= \hat{\mu}_a(a_{i,t}, z_{i,t}; Y_t, g_t)dt + \hat{\sigma}_a(a_{i,t}, z_{i,t}; Y_t, g_t)dW_t \\
\hat{\mu}_a(a, z; Y, g) &= \underbrace{z}_{\text{labor income}} + \underbrace{r^* a + (\hat{R}(Y, g) - r^*)\hat{\theta}^*(a, z; Y, g)a}_{\text{capital income}} - \underbrace{\hat{c}^*(a, z; Y, g)}_{\text{consumption}} \\
\hat{\sigma}_a(a, z; Y, g) &= \frac{\hat{\theta}^*(a, z; Y, g)a}{\mathcal{Q}(Y, g)}\hat{\sigma}_Q(Y, g),
\end{aligned}
$$

and $\hat{c}^*(\cdot)$ and $\hat{\theta}^*(\cdot)$ are the optimal consumption and portfolio policies. The key components of the operator $\hat{\mathcal{A}}^*$ are asset returns $\hat{R}(Y, g)$ and volatility $\hat{\sigma}_Q(Y, g)$, which we define below.

Similarly, the diffusion operator $\hat{\mathcal{B}}^*$ acts on the density as:

$$
\hat{\mathcal{B}}_t^*[g](a, z) = -\frac{\partial}{\partial a}\left[\hat{\theta}^*(a, z; Y_t, g_t)a\frac{\hat{\sigma}_Q(Y_t, g_t)}{\mathcal{Q}(Y_t, g_t)}g(a, z)\right]. \tag{10}
$$

This term essentially captures that when a (common) shock occurs, households who are exposed to the shock (through a non-zero $\hat{\theta}^*(\cdot)$) experience wealth changes, causing the entire distribution to shift stochastically.

Price Dynamics. The instantaneous expected return on the risky asset is given by the usual formula of dividends plus (expected) capital gains:

$$
\hat{R}(Y_t, g_t) = \frac{Y_t + \mathbb{E}_t[d\mathcal{Q}(Y_t, g_t)]/dt}{\mathcal{Q}(Y_t, g_t)}. \tag{11}
$$

To characterize the evolution of the asset price $Q_t = \mathcal{Q}(Y_t, g_t)$, we apply Itô's lemma and substitute the dynamics of dY_t and dg_t from equations (4) and (8), which gives us:

$$
\begin{aligned}
d\mathcal{Q}(Y_t, g_t) = \Bigg[&\frac{\partial \mathcal{Q}(Y_t, g_t)}{\partial Y}\mu_Y(Y_t) + \left\langle \frac{\delta \mathcal{Q}(Y_t, g_t)}{\delta g}, \hat{\mathcal{A}}_t^*[g_t] \right\rangle \\
&+ \frac{1}{2}\frac{\partial^2 \mathcal{Q}(Y_t, g_t)}{\partial Y^2}\sigma_Y^2(Y_t) + \frac{1}{2}\left\langle \left\langle \frac{\delta^2 \mathcal{Q}(Y_t, g_t)}{\delta g^2}, \hat{\mathcal{B}}_t^*[g_t] \otimes \hat{\mathcal{B}}_t^*[g_t] \right\rangle \right\rangle \\
&+ \sigma_Y(Y_t)\left\langle \frac{\partial}{\partial Y}\left[\frac{\delta \mathcal{Q}(Y_t, g_t)}{\delta g}\right], \hat{\mathcal{B}}_t^*[g_t] \right\rangle \Bigg] dt \\
&+ \left[\frac{\partial \mathcal{Q}(Y_t, g_t)}{\partial Y}\sigma_Y(Y_t) + \left\langle \frac{\delta \mathcal{Q}(Y_t, g_t)}{\delta g}, \hat{\mathcal{B}}_t^*[g_t] \right\rangle \right] dW_t
\end{aligned} \tag{12}
$$

where, importantly, in the last term the coefficient in front of dW_t in square brackets gives

us asset price volatility:

$$\hat{\sigma}_Q(Y,g) = \frac{\partial \mathcal{Q}(Y,g)}{\partial Y}\sigma_Y(Y) + \left\langle \frac{\delta \mathcal{Q}(Y,g)}{\delta g}, \hat{\mathcal{B}}^*[g] \right\rangle \tag{13}$$

which shows that asset price volatility arises from both direct dividend volatility and the stochastic evolution of the wealth distribution.[15]

Equation (12) is one of the key equations of the rational expectations model: it fully characterizes the evolution of asset prices as a function of the current aggregate state variables (the dividend Y_t and the distribution g_t) and of the optimal households' choices (through the operators $\hat{\mathcal{A}}_t^*$ and $\hat{\mathcal{B}}_t^*$).

Distributional Derivatives. When taking the derivative of the $\mathcal{Q}$ functional with respect to the infinite-dimensional distribution g the correct notion of derivative is that of the Fréchet derivative, which we denote as $\frac{\delta \mathcal{Q}}{\delta g}$. Specifically, consider how the asset price $Q_t = \mathcal{Q}(Y_t, g_t)$ responds to small changes in the distribution g_t. If the household state space were discrete with only finitely many possible states $\{a_1, a_2, \ldots, a_N\}$ (ignoring z for simplicity), then the distribution would be characterized by a vector of masses $(g_1, g_2, \ldots, g_N)$ at each state, and the price sensitivity would be given by the vector of standard partial derivatives $\left(\frac{\partial \mathcal{Q}}{\partial g_i}\right)_{i=1}^N$ measuring how price changes when the mass at state i increases or decreases.

The Fréchet derivative simply extends this concept to continuous state spaces. For the price functional $\mathcal{Q}(Y,g)$, the Fréchet derivative $\frac{\delta \mathcal{Q}(Y,g)}{\delta g}[a',z']$ measures how the asset price responds to adding an infinitesimal mass of households at the specific location (a',z') in the wealth-income space. Since different household types (a',z') have different portfolio demands $\hat{\theta}^*(a',z';Y,g)$, adding mass at different locations (a',z') will have different effects on the equilibrium price, which is precisely what $\frac{\delta \mathcal{Q}(Y,g)}{\delta g}[a',z']$ captures.[16]

Then, the angle brackets $\langle \cdot, \cdot \rangle$ represent the L^2 inner product on the space of functions. For any two functions f and g defined on $\mathbb{R}_+ \times \{z_1, z_2\}$ we write:

$$\langle f, g \rangle = \int_{\mathbb{R}_+ \times \{z_1,z_2\}} f(a,z)g(a,z)\, da\, dz. \tag{14}$$

[15] It might be worth noting at this point that both $\hat{\mathcal{A}}^*$ and $\hat{\mathcal{B}}^*$ in equations (9) and (10) are defined recursively. Specifically, $\hat{\mathcal{A}}^*$ depends both on $\hat{R}(Y,g)$ – which itself depends on both $\hat{\mathcal{A}}^*$ and $\hat{\mathcal{B}}^*$ – and on $\hat{\sigma}_Q(Y,g)$ – which itself depends on $\hat{\mathcal{B}}^*$ – while $\hat{\mathcal{B}}^*$ depends on $\hat{\sigma}_Q(Y,g)$ (on top of their "indirect" dependency through optimal consumption and portfolio choice). While this recursive formulation does not pose a problem in and of itself, as it will be clear later, it makes the problem highly nonlinear, which is part of what motivates our neural-network approach.

[16] The notation $\frac{\delta \mathcal{Q}(Y,g)}{\delta g}[a',z']$ emphasizes that this derivative depends on both the current state (Y,g) where we evaluate the functional and the specific location (a',z') where we consider the perturbation – just as a finite-dimensional gradient depends on both the evaluation point and the direction of the perturbation.

For example, the term $\left\langle \frac{\delta \mathcal{Q}}{\delta g}, \hat{\mathcal{B}}^*[g] \right\rangle$ which shows up in equation (13) represents how the stochastic evolution of the distribution affects asset price volatility:

$$\left\langle \frac{\delta \mathcal{Q}(Y,g)}{\delta g}, \hat{\mathcal{B}}^*[g] \right\rangle = \int_{\mathbb{R}_+ \times \{z_1, z_2\}} \frac{\delta \mathcal{Q}(Y,g)}{\delta g}[a', z'] \cdot \hat{\mathcal{B}}^*[g](a', z')\, da'\, dz'$$

where $\frac{\delta \mathcal{Q}(Y,g)}{\delta g}[a', z']$ captures how sensitive the asset price is to changes in the distribution at point (a', z'), and $\hat{\mathcal{B}}^*[g](a', z')$ represents how the common noise term dW_t affects the distribution at that same point.

Similarly, double angle brackets $\langle\langle \cdot, \cdot \rangle\rangle$ represent a bilinear form for second-order Fréchet derivatives:[17]

$$\left\langle\left\langle \frac{\delta^2 \mathcal{Q}}{\delta g^2}, \hat{\mathcal{B}}^*[g] \otimes \hat{\mathcal{B}}^*[g] \right\rangle\right\rangle = \int\int \frac{\delta^2 \mathcal{Q}(Y,g)}{\delta g^2}[a', z'][a'', z''] \cdot \hat{\mathcal{B}}^*[g](a', z') \cdot \hat{\mathcal{B}}^*[g](a'', z'')\, da' dz' da'' dz''$$

The Master Equation. Under rational expectations, each household correctly anticipates the stochastic evolution of both dividends Y_t and the distribution g_t. The household's value function $V(\cdot)$ then becomes a function of the infinite-dimensional distribution g, and the

[17] The second-order derivative $\frac{\delta^2 \mathcal{Q}}{\delta g^2}$ is a bilinear functional that takes two "directions" in the distribution space; and the tensor product $\hat{\mathcal{B}}^*[g] \otimes \hat{\mathcal{B}}^*[g]$ creates the appropriate bilinear object. That is, for any two points (a', z') and (a'', z''), the tensor product equals $\hat{\mathcal{B}}^*[g](a', z') \cdot \hat{\mathcal{B}}^*[g](a'', z'')$.

recursive problem can be written in its so-called "Master equation" form as:

$$
\begin{aligned}
(\rho+\phi)\hat{V}(a,z;Y,g) = \max_{c,\theta}\Bigg\{ & u(c) + \hat{\mu}_a(a,z,c,\theta;Y,g)\frac{\partial \hat{V}(a,z;Y,g)}{\partial a} \\
& + \lambda_z\left[\hat{V}(a,z';Y,g) - \hat{V}(a,z;Y,g)\right] \\
& + \mu_Y(Y)\frac{\partial \hat{V}(a,z;Y,g)}{\partial Y} \\
& + \left\langle \frac{\delta \hat{V}(a,z;Y,g)}{\delta g}, \hat{\mathcal{A}}^*[g] \right\rangle \\
& + \frac{1}{2}\hat{\sigma}_a^2(a,z,\theta;Y,g)\frac{\partial^2 \hat{V}(a,z;Y,g)}{\partial a^2} \\
& + \frac{1}{2}\sigma_Y^2(Y)\frac{\partial^2 \hat{V}(a,z;Y,g)}{\partial Y^2} \\
& + \frac{1}{2}\left\langle\left\langle \frac{\delta^2 \hat{V}(a,z;Y,g)}{\delta g^2}, \hat{\mathcal{B}}^*[g]\otimes\hat{\mathcal{B}}^*[g] \right\rangle\right\rangle \\
& + \hat{\sigma}_a(a,z,\theta;Y,g)\sigma_Y(Y)\frac{\partial^2 \hat{V}(a,z;Y,g)}{\partial a \partial Y} \\
& + \sigma_Y(Y)\left\langle \frac{\partial}{\partial Y}\left[\frac{\delta \hat{V}(a,z;Y,g)}{\delta g}\right], \hat{\mathcal{B}}^*[g] \right\rangle \\
& + \hat{\sigma}_a(a,z,\theta;Y,g)\left\langle \frac{\partial}{\partial a}\left[\frac{\delta \hat{V}(a,z;Y,g)}{\delta g}\right], \hat{\mathcal{B}}^*[g] \right\rangle \Bigg\}
\end{aligned}
\tag{15}
$$

which is an infinite-dimensional equation in the space of distributions. This is similar to the Master equation derived in both Bilal (2023) and Gu et al. (2023), the main difference being the presence of a second-order term due to the presence of a common noise term directly affecting the evolution of households' states.[18]

[18]Notice that even in the applications with aggregate shocks considered by Bilal (2023) and Gu et al. (2023), their Master equations still do not feature a second-order term since aggregate shocks only affect households' states *indirectly* through aggregate equations. In our model, on the other hand, because aggregate shocks *directly* affect households' wealth, common noise results in second-order terms in the Master equation. Note that the "second-order" in the method of Bilal (2023) refers instead to the perturbation order of the solution technique, not to a second-order derivative term of the kind discussed here. Notice, however, that this does *not* mean that our Master equation is now a stochastic equation. In fact, as pointed out by Carmona and Delarue (2018, p. 286) "It is tempting to believe that the PDE ought to be stochastic because of the presence of the common noise in the dynamics of the players. However [...] the coefficients of the infinite dimensional forward-backward system comprising the Fokker-Planck and the Hamilton-Jacobi-Bellman equations are deterministic functions of the unknowns, namely the equilibrium measure (which is random) and the value function (which is random as well). For that reason [...] the presence of the common noise is merely the reason for the existence of a nontrivial *second-order term* in the PDE."

The Problem with Rational Expectations. The extremely complicated structure of the Master equation (15) is a direct consequence of the extraordinary informational requirements that are imposed on households by the rational expectations assumption in the context of heterogeneous-agent models. In our context, due to the presence of aggregate shocks directly affecting households, these informational burdens are particularly severe.

Beyond the standard terms capturing the household's own state dynamics, the Master equation (15) features both first-order functional derivatives with respect to the distribution – terms like $\langle \frac{\delta V}{\delta g}, \hat{\mathcal{A}}^*[g] \rangle$ – capturing the distribution's "average" evolution and second-order functional derivatives – terms like $\langle\langle \frac{\delta^2 V}{\delta g^2}, \hat{\mathcal{B}}^*[g] \otimes \hat{\mathcal{B}}^*[g] \rangle\rangle$ – capturing its stochastic dynamics. As argued by Moll (2024) and Moll and Ryzhik (2025), the assumption that households can perfectly solve (or even understand) these problems merits reconsideration, particularly if we recognize that even the most sophisticated market participants – e.g., institutional investors, hedge funds, or policy institutions – do not actually forecast prices by forecasting the evolution of the entire cross-sectional distribution.

As if this were not enough, relative to a "standard" model like Krusell and Smith (1998), the *direct* effect of aggregate risk in our framework further amplifies these difficulties. The fact that in our model rational households would need to solve a second-order Master equation means that they must also account for the distribution's *stochastic* evolution – they must understand not just where the distribution is heading "on average", but the entire probability distribution over possible future wealth distributions. This includes understanding how aggregate dividend shocks differently affect different parts of the wealth distribution through heterogeneous portfolio exposures, and how this in turn translates into changes in households' wealth holdings and therefore asset prices. Rational households in our model would need to have not only a perfect understanding of the complex feedback between households' portfolio choice, asset prices, and the distribution but also how this feedback changes in response to aggregate shocks.

The fundamental issue with rational expectations is that we are attributing to households – and indeed to *all* market participants – a level of sophistication that goes beyond what is observed in practice. Moll (2024) and Moll and Ryzhik (2025) make the same argument: households and firms rely on simple forecasting heuristics rather than tracking the full cross-sectional distribution. Rather than assuming all agents solve these highly complex equations, we should model the simpler price-forecasting procedures they actually employ. This motivates our limited information approach, where households forecast prices directly without needing to track the infinite-dimensional state of the economy.

2.2.2 Low-Dimensional Coupling

A critical observation is that our model, like most heterogeneous-agent macroeconomic models, exhibits what Moll and Ryzhik (2025) term "low-dimensional coupling": households' optimization problem depends on the infinite-dimensional distribution only through equilibrium prices.[19]

To see this clearly, let us go back to the household's budget constraint. In the presence of a single source of uncertainty, just as in equation (12), prices will follow a (possibly infinite-dimensional) diffusion process:

$$d\mathcal{Q}(\Omega_t) = \mu_Q(\Omega_t)dt + \sigma_Q(\Omega_t)dW_t \tag{16}$$

where Ω_t is the set of relevant aggregate states and, by construction

$$\mu_Q(\Omega_t) = \frac{\mathbb{E}_t[d\mathcal{Q}(\Omega_t)]}{dt} \tag{17}$$

$$\sigma_Q(\Omega_t) = \sqrt{\frac{\mathbb{V}_t(d\mathcal{Q}(\Omega_t))}{dt}}. \tag{18}$$

As a consequence, household wealth evolves according to:

$$da_{i,t} = \left[z_{i,t} + r^* a_{i,t} + \theta_{i,t} a_{i,t} \left(\frac{Y_t + \mu_Q(\Omega_t)}{Q_t} - r^*\right) - c_{i,t}\right] dt + \theta_{i,t} a_{i,t} \frac{\sigma_Q(\Omega_t)}{Q_t} dW_t. \tag{19}$$

Equation (19) makes clear that, setting aside for a moment what Ω_t is, the only things that households need to know to make optimal decisions are μ_Q and σ_Q.

Somewhat unfortunately, as we saw in section 2.2.1, under rational expectations the set of relevant state variables, Ω_t, includes the distribution g_t, and this low-dimensional coupling provides no simplification whatsoever: the Master equation remains infinite-dimensional because rational forward-looking households must forecast the distribution's evolution to determine future price dynamics. However, the fact that households only care about the distribution because they care about μ_Q and σ_Q suggests a natural alternative: to model them as forecasting prices directly.

In our limited information framework below we will specifically assume that $\Omega_t = (Y_t, Q_t)$. By modeling how households forecast the evolution of this low-dimensional state vector directly – without reference to the underlying cross-sectional distribution of asset holdings

[19] As emphasized by Moll and Ryzhik (2025), this special structure arises for two main reasons: first, the distribution g_t does not *directly* enter either in the households' objective function or in their constraints; second, the competitive equilibrium assumption implies that market-clearing conditions compress the infinite-dimensional distribution into finite-dimensional prices.

– we will therefore be able to sidestep the Master equation entirely and to work with finite-dimensional state spaces that are both computationally tractable and conceptually plausible. In fact, as we will argue momentarily, the asset price Q_t itself is *the* natural endogenous aggregate state to use to forecast future prices.

Importantly, such a formulation also has a simple information-theoretic interpretation: households' decision problem depends only on easily-observable aggregate variables (dividends and prices) and on their own states – precisely the information available to real-world investors. The filtration generated by $(Y_t, Q_t, a_{i,t}, z_{i,t})$ represents the natural information set that agents actually possess, rather than the more complex distributional information assumed either in the standard Krusell-Smith approach or, worse, the complete distributional information assumed under rational expectations. In this sense, the limited information set is not a modeling compromise but rather a realistic representation of how agents interact with financial markets. Notice that by assuming $\Omega_t = (Y_t, Q_t)$ we are directly assuming that prices follow a Markovian process, whereas in reality the true price process – being a function of the infinite-dimensional distribution – is not Markovian in (Y, Q) alone (see Moll and Ryzhik 2025).[20]

There remains, of course, the question of *how* agents might forecast these prices in this framework. Equation (12) makes clear that the drift and diffusion terms can be quite complicated objects. Ideally, one would have a perceived law of motion that is sufficiently simple so as to not complicate the model solution, but also sufficiently flexible to allow for potentially highly nonlinear dynamics. We turn to this question next.

2.2.3 Limited Information and Neural Networks

Our limited information framework is composed of two key elements: the first, as we just argued, is that households forecast prices directly rather than tracking the infinite-dimensional distribution. Specifically, we assume that households believe that the asset price follows equation (16) and that $\Omega_t = (Y_t, Q_t)$ – that is, the set of relevant aggregate state variables consists only of current dividends and the current price level. The second is that we allow both the drift and volatility functions $\mu_Q(\Omega)$ and $\sigma_Q(\Omega)$ to be completely arbitrary and possibly highly nonlinear. Specifically, we approximate the perceived law of motion with two neural networks: $\mu_Q(\Omega; \alpha)$ and $\sigma_Q(\Omega; \beta)$, where α and β are network weights to be found. The neural networks act as universal function approximators, allowing us to capture complex relationships between states and price dynamics without imposing restrictive functional forms.

This formulation generalizes the Krusell and Smith (1998) approach, where agents use

[20] Yang et al. (2025) show that solutions assuming Markovian prices perform very similarly to more complex approaches that account for the non-Markovian structure, lending support to our specification.

only the first moment of the wealth distribution to forecast prices. In their economy, this works well because there is a one-to-one mapping between interest rates and aggregate capital through the firm's first-order condition. Our framework maintains their insight that agents can forecast future economic conditions using a much lower-dimensional state space, but allows for richer dynamics.[21]

The use of neural networks to learn μ_Q and σ_Q also extends the approach in Fernández-Villaverde, Hurtado, et al. (2023) to a context in which aggregate shocks directly affect households' states and the volatility term σ_Q must also be learned. While the two approaches are computationally similar, we stress that we do not employ neural networks to approximate the rational expectations equilibrium but simply to model how agents might form beliefs about price dynamics in a context in which they have access to limited information. This distinction is crucial: we do not claim that neural networks provide a good approximation to the rational expectations equilibrium, but rather that they can help us understand how agents might forecast prices in complex environments using simple rules.

The flexibility provided by neural networks in approximating the law of motion of asset prices is particularly important given the complexity of the true price dynamics. As equation (12) from the rational expectations model demonstrates, the actual law of motion involves intricate interactions between the distribution's evolution and dividend shocks. In the presence of aggregate shocks that directly affect both distributions and prices, there is no reason to expect simple linear or log-linear relationships. Neural networks can therefore capture these nonlinearities while maintaining computational tractability – transforming an infinite-dimensional problem into a two-dimensional one that remains rich enough to generate interesting dynamics.

In fact, neural networks can naturally capture the nonlinear aggregation of micro-level heterogeneity into macro-level price dynamics. While individual portfolio decisions $\theta^*(a, z; Y, Q)$ may vary smoothly across the wealth distribution, their aggregate effect on prices through market clearing can be highly nonlinear. Neural networks learn these nonlinear mappings without imposing restrictions on their functional form, allowing micro features – such as wealth concentration – to have disproportionate effects on asset price dynamics.

[21] Indeed note that if portfolio shares were constant $\theta^*(\cdot) = \bar{\theta}$, from market clearing we would find that the asset price is proportional to aggregate wealth $Q_t = \bar{\theta} \int a dG_t \propto A_t$ and our approach would reduce to a version of their first-moment approximation in which we replace the log-linear regression with a neural-network approximation.

2.3 Household Problem

Under limited information, we assume that households believe that asset prices evolve according to the following perceived law of motion (PLM):[22]

$$dQ_t = \mu_Q(Y_t, Q_t)dt + \sigma_Q(Y_t, Q_t)dW_t \tag{20}$$

where $\mu_Q(\cdot)$ and $\sigma_Q(\cdot)$ are functions to be determined.

Households choose consumption c and portfolio share θ to maximize expected discounted utility:

$$\max_{c,\theta} \mathbb{E}_0 \left[\int_0^\infty e^{-(\rho+\phi)t} u(c_t)\, dt \right] \tag{21}$$

subject to the borrowing and short-sale constraints $a_t \geq 0$, $0 \leq \theta \leq 1$, and the law of motion for wealth:

$$da_t = \mu_a(a_t, z_t; Y_t, Q_t)dt + \sigma_a(a_t, z_t; Y_t, Q_t)dW_t \tag{22}$$

where

$$\mu_a(a, z; Y, Q) = [z + r^* a + \theta a \left(R(Y, Q) - r^*\right) - c] \tag{23}$$

$$\sigma_a(a, z; Y, Q) = \theta a \frac{\sigma_Q(Y, Q)}{Q} \tag{24}$$

and $R(Y, Q) = (Y + \mu_Q(Y, Q))/Q$ is the instantaneous total expected return on the risky asset.

[22] Notice that, in equilibrium, the *actual* evolution of prices will naturally result from the evolution of the distribution through the market clearing condition (31). It is only the *perceived* law of motion of prices that satisfies equation (20). In fact, we could in principle write a differential equation for the actual evolution of prices similar to (12) in which the operators $\mathcal{A}$ and $\mathcal{B}$ are consistent with the household problem under limited information. However, because "realized" prices depend on the evolution of the distribution, such differential equation would still need to be written in the space of distribution. Hence, we would not be able to actually use it in the model solution, so we can safely ignore it.

2.3.1 Recursive Formulation

Imposing the Bellman principle and applying Itô's lemma to the value function, we obtain the Hamilton-Jacobi-Bellman equation:[23]

$$\begin{aligned}(\rho+\phi)v(a,z;Y,Q) = \max_{c,\theta}\Bigg\{ & u(c) + v_a\left[z + r^*a + \theta a(R(Y,Q)-r^*) - c\right] \\ & + \frac{1}{2}v_{aa}\left(\theta a\frac{\sigma_Q(Y,Q)}{Q}\right)^2 + \lambda_z\left[v(a,z';Y,Q) - v(a,z;Y,Q)\right] \\ & + v_Y\mu_Y(Y) + \frac{1}{2}v_{YY}\sigma_Y(Y)^2 \\ & + v_Q\mu_Q(Y,Q) + \frac{1}{2}v_{QQ}\sigma_Q(Y,Q)^2 \\ & + v_{aY}\theta a\frac{\sigma_Q(Y,Q)}{Q}\sigma_Y(Y) + v_{aQ}\theta a\frac{\sigma_Q(Y,Q)^2}{Q} \\ & + v_{YQ}\sigma_Y(Y)\sigma_Q(Y,Q)\Bigg\}\end{aligned} \tag{25}$$

where the cross-derivative terms arise because the aggregate shock dW_t simultaneously affects individual wealth (through portfolio returns), dividends, and asset prices.[24]

2.3.2 First-Order Conditions

The first-order condition for consumption yields the familiar expression:

$$c^*(a,z;Y,Q) = \left[\partial_a v(a,z;Y,Q)\right]^{-\frac{1}{\gamma}} \tag{26}$$

while that for portfolio choice gives:

$$\theta^*(a,z;Y,Q) = \underbrace{-\frac{v_a}{v_{aa}}\frac{R(Y,Q)-r^*}{a\left(\frac{\sigma_Q}{Q}\right)^2}}_{\text{risk-return tradeoff}} \underbrace{-\frac{v_{aY}}{v_{aa}}\frac{\sigma_Y}{\sigma_Q}\frac{Q}{a}}_{\substack{\text{dividend}\\ \text{hedging}}} \underbrace{-\frac{v_{aQ}}{v_{aa}}\frac{Q}{a}}_{\substack{\text{price}\\ \text{hedging}}} \tag{27}$$

which decomposes optimal portfolio choice into three components: a myopic risk-return tradeoff and two hedging demands that arise because the household's value function is exposed to aggregate shocks through dividends and prices.

[23] To ease notation we write v_a for $\frac{\partial v(a,z;Y,Q)}{\partial a}$ whenever it is unambiguous and do the same for second-order derivatives.

[24] Notice that the value function v does not depend on time because, in the limited-information equilibrium, households solve a stationary problem in which the evolution of prices does not depend on the distribution (even though the latter is time-varying).

The first term is the standard myopic demand: the household holds the risky asset in proportion to the excess return $R(Y,Q) - r^*$, scaled by the ratio $-v_a/v_{aa}$ and inversely proportional to the variance of portfolio returns $(\sigma_Q/Q)^2$. This is the only component that would appear in a portfolio problem without additional aggregate state variables.

The second term captures a hedging demand driven by the cross-derivative v_{aY}, which measures how the marginal value of wealth responds to dividend movements. When $v_{aY} < 0$ – that is, when higher dividends reduce the marginal value of wealth, as one would expect if rising dividends raise asset returns and hence wealth – the household has an incentive to hold the risky asset as a hedge, since it pays off precisely in states where the marginal value of wealth is low. The ratio σ_Y/σ_Q governs the hedging effectiveness by measuring the correlation between dividend and price risk.

The third term is a hedging demand arising from the household's exposure to price risk through the cross-derivative v_{aQ}. Because the price Q enters the household problem as a state variable – through the perceived law of motion – households internalize that price movements affect their future investment opportunities and hedge accordingly. Notice that this channel is specific to our limited-information setup: under rational expectations, the household's state would be (Y, g) rather than (Y, Q), and the hedging demand would take a different form. In fact, because μ_Q and σ_Q from the perceived law of motion enter the portfolio condition both directly – through the risk-return term – and indirectly – through the price-hedging term – households' beliefs about price dynamics are a first-order determinant of portfolio choice, which then feeds back into equilibrium prices through market clearing.

The presence of second-derivative terms, and especially of v_{aa}, in the portfolio choice formula underscores why high-precision numerical methods are essential: small errors in approximating these derivatives can lead to substantial misspecification of portfolio demand, which then propagates to the wealth distribution and – through the market clearing condition – to equilibrium prices.

2.4 Equilibrium Definition

We now formally define the recursive competitive equilibrium under limited information.

A recursive competitive equilibrium consists of: value function $v(a, z; Y, Q)$ and associated policy functions $c^*(a, z; Y, Q)$ and $\theta^*(a, z; Y, Q)$, perceived law of motion functions $\mu_Q(Y, Q)$ and $\sigma_Q(Y, Q)$, a distribution $g_t(a, z)$ over wealth and income states, and the asset price Q, such that:

(i) **Households optimize:** Given the perceived law of motion (20), the value function $v(a, z; Y, Q)$ and policy functions $c^*(a, z; Y, Q)$, $\theta^*(a, z; Y, Q)$ solve the Hamilton-Jacobi-Bellman equation (25).

(ii) **Distribution evolution:** The distribution $g_t(a, z)$ evolves according to the Kolmogorov forward equation:

$$dg_t(a, z) = \mathcal{A}_t^*[g_t](a, z)\, dt + \mathcal{B}_t^*[g_t](a, z)\, dW_t \tag{28}$$

where the operators $\mathcal{A}^*$ and $\mathcal{B}^*$ are defined as

$$\begin{aligned}\mathcal{A}_t^*[g](a, z) = &-\frac{\partial}{\partial a}[\mu_a(a, z; Y_t, Q_t)g(a, z)] + \frac{1}{2}\frac{\partial^2}{\partial a^2}[\sigma_a^2(a, z; Y_t, Q_t)g(a, z)] \\ &+ \lambda_{z'}g(a, z') - \lambda_z g(a, z) + \phi(\psi(a, z) - g(a, z))\end{aligned} \tag{29}$$

$$\mathcal{B}_t^*[g](a, z) = -\frac{\partial}{\partial a}\left[\theta^*(a, z; Y_t, Q_t)a\frac{\sigma_Q(Y_t, Q_t)}{Q_t}g(a, z)\right], \tag{30}$$

and μ_a and σ_a are defined as in equations (23) and (24).[25]

(iii) **Market clearing:** The asset market clears:

$$Q_t = \int_{\mathbb{R}_+ \times \{z_1, z_2\}} \theta^*(a, z; Y_t, Q_t)a\, g_t(a, z)\, da\, dz \quad \forall\ t. \tag{31}$$

(iv) **Consistency of beliefs:** Households' perceived law of motion (20) is consistent with the evolution of Q_t implied by equations (28) and (31).[26]

3 Methodology

3.1 Neural Network Approach

While low-dimensional coupling and limited information allow us to sidestep the curse of dimensionality inherent in the Master equation, we still face the challenge of approximating potentially highly nonlinear mappings from the two-dimensional state space to price dynamics. Specifically, the limited information framework discussed in section 2 requires approximating the perceived law of motion for asset prices – i.e., the drift and diffusion functions $\mu_Q(\Omega)$ and $\sigma_Q(\Omega)$ where $\Omega_t = (Y_t, Q_t)$. Neural networks provide a natural solution to this challenge and offer two key advantages: First, they serve as universal function approximators, allowing us to capture complex relationships without imposing restrictive

[25] The derivation of equation (28) can be found in the Appendix A.

[26] As condition (iv) makes clear, our equilibrium concept is that of a self-confirming equilibrium in the spirit of Cho and Sargent (2018): agents' subjective beliefs about price dynamics generate actual price dynamics consistent with those beliefs.

functional forms *ex ante*. Second, they provide a global solution method that can handle nonlinearities throughout the state space, avoiding the local approximation errors that arise from perturbation methods and that are particularly problematic in models with asset pricing.

We approximate the perceived law of motion using two separate neural networks: $\mu_Q(\Omega;\alpha)$ for the drift and $\sigma_Q(\Omega;\beta)$ for the diffusion, where α and β denote the network weights to be determined. By construction, if the evolution of prices is consistent with households' PLM equation (16), it must be the case that

$$\begin{aligned}\mathbb{E}_t[dQ_t] &= \mu_Q(\Omega_t)dt \\ \mathbb{V}_t[dQ_t] &= \sigma_Q^2(\Omega_t)dt,\end{aligned}$$

which we can use to train $\mu_Q(\Omega;\alpha)$ and $\sigma_Q(\Omega;\beta)$ on the corresponding simulated objects.

The algorithm proceeds iteratively. Starting from an initial guess for the perceived law of motion, denoted $\mu_{Q,n}(\Omega;\alpha^n)$ and $\sigma_{Q,n}(\Omega;\beta^n)$ at iteration n, we:

(i) **Solve the household problem.** Given the PLM functions $\mu_{Q,n}(\cdot), \sigma_{Q,n}(\cdot)$, we solve the Hamilton-Jacobi-Bellman equation (25) using a finite-difference scheme; we detail the discretization of this step – and of step (ii) – in section D.

(ii) **Simulate the economy.** Using the optimal policies from step (i), we simulate the economy forward using the Kolmogorov Forward equation (28) to track the distribution's evolution. For each period t in our simulation, we compute the market-clearing price Q_t from the asset market clearing condition (31) and record the realized state $\Omega_t = (Y_t, Q_t)$.

(iii) **Train the neural networks.** From the simulated time series $\{\Omega_t\}_{t=1}^T$, we construct empirical estimates of the drift and diffusion:

$$\hat{\mu}_{Q,t} = \frac{\Delta Q_{t+\Delta t}}{\Delta t} \tag{32}$$

$$\hat{\sigma}^2_{Q,t} = \frac{(\Delta Q_{t+\Delta t})^2 - (\mu_{Q,n}(Y_t, Q_t)\,\Delta t)^2}{\Delta t} \tag{33}$$

We then update the network weights by minimizing the mean squared error:

$$\alpha^{n+1} = \arg\min_{\alpha} \frac{1}{2}\sum_{t=1}^{T} \|\mu_Q(\Omega_t;\alpha) - \hat{\mu}_{Q,t}\|^2 \tag{34}$$

$$\beta^{n+1} = \arg\min_{\beta} \frac{1}{2}\sum_{t=1}^{T} \|\sigma^2_Q(\Omega_t;\beta) - \hat{\sigma}^2_{Q,t}\|^2 \tag{35}$$

(iv) **Check convergence.** If $\|\mu_{Q,n+1} - \mu_{Q,n}\| < \epsilon$ and $\|\sigma_{Q,n+1} - \sigma_{Q,n}\| < \epsilon$ for some arbitrarily small tolerance ϵ, we have found a fixed point of the perceived law of motion. Otherwise, return to step (i).

This iterative procedure searches for a self-confirming equilibrium in the spirit of Cho and Sargent (2018) and Adam, Marcet, and Nicolini (2016), where agents' subjective beliefs about price dynamics are not disappointed on the equilibrium path.[27] Like Adam, Marcet, and Nicolini (2016), we depart from rational expectations in modeling agents' beliefs about asset price dynamics, but employ neural networks rather than adaptive learning to approximate the perceived law of motion. The neural networks approximate the equilibrium mapping from states to price changes; the resulting beliefs are not time-varying but instead correspond to the fixed point of the self-confirming equilibrium. Importantly, because agents forecast prices directly rather than tracking distributions, the computational burden remains manageable even as the underlying heterogeneity becomes arbitrarily complex.

The training process itself employs standard backpropagation, though the economic structure of our problem introduces several implementation details worth noting. First, we find that separate networks for drift and volatility perform better than a single network with multiple outputs, likely because the two functions exhibit different degrees of nonlinearity (as should be expected, volatility $\sigma_Q(\cdot)$ is much more nonlinear than the drift $\mu_Q(\cdot)$). Second, the simulation length T must be chosen carefully: too short and the networks lack sufficient training data, particularly for rarely visited regions of the state space; too long and computational costs become prohibitive. In practice, we find that simulating 800 years at a monthly frequency provides adequate coverage of the ergodic distribution while maintaining tractability.

We relegate the remaining implementation details to section C, which describes the network architecture and training procedure – a single hidden layer with $K = 8$ softplus units per network, trained by full-batch gradient descent with exact line search – and discusses the approximation properties that make neural networks well-suited to this problem.

3.2 Calibration

For our calibration approach we directly assign values to a subset of parameters based on standard choices in the literature, while the remaining parameters are chosen to match key moments of the data.

The discount rate ρ is set to 0.06, which is in the range of typical values used in the literature and allows us to match the wealth-to-income ratio given our simplified asset structure.

[27] While, at least in theory, such beliefs may be incorrect off-equilibrium, neural networks have particularly good extrapolation properties – as shown also by Fernández-Villaverde, Hurtado, et al. (2023).

The coefficient of relative risk aversion γ is set to 2.5, within the conventional range used in macroeconomic studies, which ensures that households' portfolio decisions generate realistic exposure to risky assets. The death rate $\phi = 0.02$ implies an average work-life duration of $1/\phi = 50$ years, consistent with the perpetual-youth framework of Blanchard (1985).

For the earnings process, we adopt a two-state specification with employed and unemployed states and follow a similar strategy as in Fernández-Villaverde, Hurtado, et al. (2023): the transition intensities $(\lambda_1, \lambda_2) = (0.986, 0.052)$ are chosen to generate a 5% unemployment rate in steady state, as well as a job finding rate of 0.3 at monthly frequency. We normalize average labor earnings so that total income (including dividends) equals 1 in steady state, $\frac{\lambda_2}{\lambda_1+\lambda_2} z_1 + \frac{\lambda_1}{\lambda_1+\lambda_2} z_2 = 1 - \bar{Y}$, and set z_1 to be 71% of z_2 as in Hall and Milgrom (2008).

The dividend process parameters are chosen to be illustrative rather than to match specific empirical moments, with values within the range commonly used in the asset pricing literature. We set the mean dividend level $\bar{Y} = 0.1$ so that dividends represent roughly 10% of total income – which yields a plausible dividend-price ratio in equilibrium. The mean reversion parameter $\eta = 0.1$,[28] implies moderately persistent dividend shocks, while the volatility parameter $\sigma = 0.1$ controls the amplitude of dividend fluctuations and, through general equilibrium, the volatility of equity returns. The goal of this calibration is to study the qualitative mechanisms of the model rather than to replicate the exact quantitative properties of U.S. equity markets.

Table 1 presents the calibrated parameter values alongside their targets. The resulting economy captures the essential features needed to study the interaction between wealth inequality and asset prices.

4 Results

We organize the discussion of our findings in five parts. First, we justify our neural-network approach by characterizing the converged perceived law of motion and showing that its drift component is nearly linear, while the volatility component is highly nonlinear and state-dependent. Second, we provide evidence that our mechanism works by verifying that the converged PLM constitutes a self-confirming equilibrium – i.e., that agents' forecasts are statistically consistent with the price dynamics that actually materialize in the simulated economy. Third, we trace the aggregate dynamics of the model by computing impulse responses to a negative dividend shock, documenting the response of prices, dividend yields, and portfolio shares. Fourth, we examine the portfolio surface, the endogenous equity premium, and the mechanism through which portfolio choice and time-varying risk premia

[28] The Ornstein-Uhlenbeck process for dividends (3) is equivalent to an AR(1) in discrete time. A mean reversion parameter $\eta = 0.1$ is equivalent to an autoregressive parameter of 0.9 at annual frequency.

Table 1: Calibrated Parameters

Parameter	Value	Target
Preferences		
ρ	0.06	Wealth-to-income ratio
γ	2.5	Average portfolio share
ϕ	0.02	Average work-life duration
Assets		
r^*	0.04	World interest rate
Earnings Process		
λ_1	0.986	Job-finding rate (30%)
λ_2	0.052	Share of unemployed (5%)
z_1	0.648	Unemployed income
z_2	0.913	Normalization
Dividend Process		
$\bar{Y}$	0.1	Dividend yield
η	0.1	Dividend persistence
σ	0.1	Return volatility

generate the aggregate dynamics documented in the impulse responses. Fifth, we illustrate the distributional implications of aggregate shocks, tracing the mechanism from households' heterogeneous equity exposure to state-dependent wealth distributions.

4.1 The Perceived Law of Motion

Figures 1 and 2 display the converged perceived law of motion for asset prices – that is, the learned functions $\mu_Q(Y, Q)$ and $\sigma_Q(Y, Q)$ that characterize agents' beliefs about price dynamics at the fixed point of the algorithm. The darker band in each figure indicates the ergodic region of the state space – the (Y, Q) pairs that the economy visits most frequently in the long run.

The learned drift function $\mu_Q(Y, Q)$, shown in figure 1, is close to linear in the ergodic region. The drift is positive when prices are low relative to dividends and negative when prices are high, reflecting mean-reversion: prices that are above their long-run level given current dividends tend to fall, and prices below it tend to rise. This pattern is similar to what one would obtain from a standard log-linear approximation as in Krusell and Smith (1998), and it suggests that the drift alone – being approximately linear – would probably capture the *average* direction of price movements well. The near-linearity of μ_Q indicates that, for the drift component alone, the nonlinear flexibility offered by neural networks is not essential. At the same time, however, as we will see momentarily, such flexibility is crucial for capturing the volatility – and it is the volatility, not the drift, that drives the key economic mechanisms of the model.

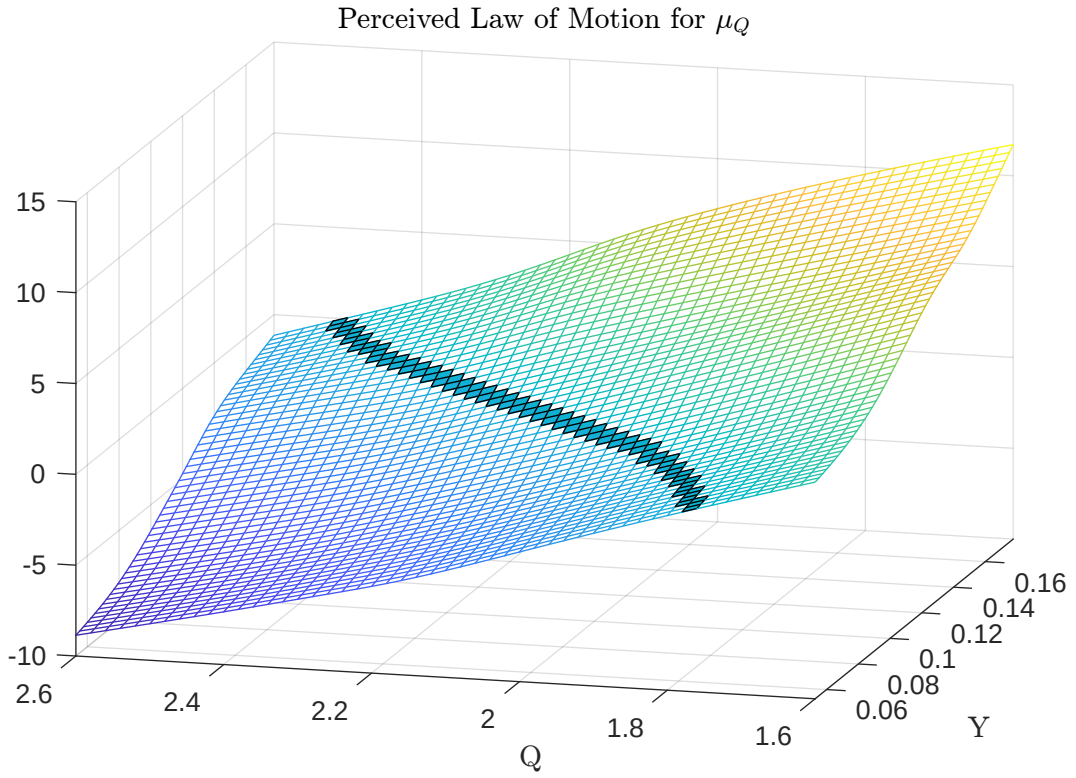


Figure 1: Perceived law of motion: drift $\mu_Q(Y, Q)$.

Notes: The surface shows the learned drift function $\mu_Q(Y, Q)$ as a function of the dividend level Y and the price level Q. The darker band indicates the ergodic region. The near-linearity of the surface indicates that the drift alone captures the average direction of price movements well.

In fact, the learned diffusion function $\sigma_Q(Y, Q)$ shown in figure 2 reveals a fundamentally different pattern. Price volatility is highly nonlinear and state-dependent: it is near zero for low asset prices and rises sharply as prices increase, with substantial curvature even within the ergodic region. This asymmetry between the drift and volatility is one of the key findings of our analysis. While the *expected* direction of price movements is well approximated by a linear function – consistent with the success of the Krusell and Smith (1998) approach in simpler settings – the *volatility* of price movements is not.

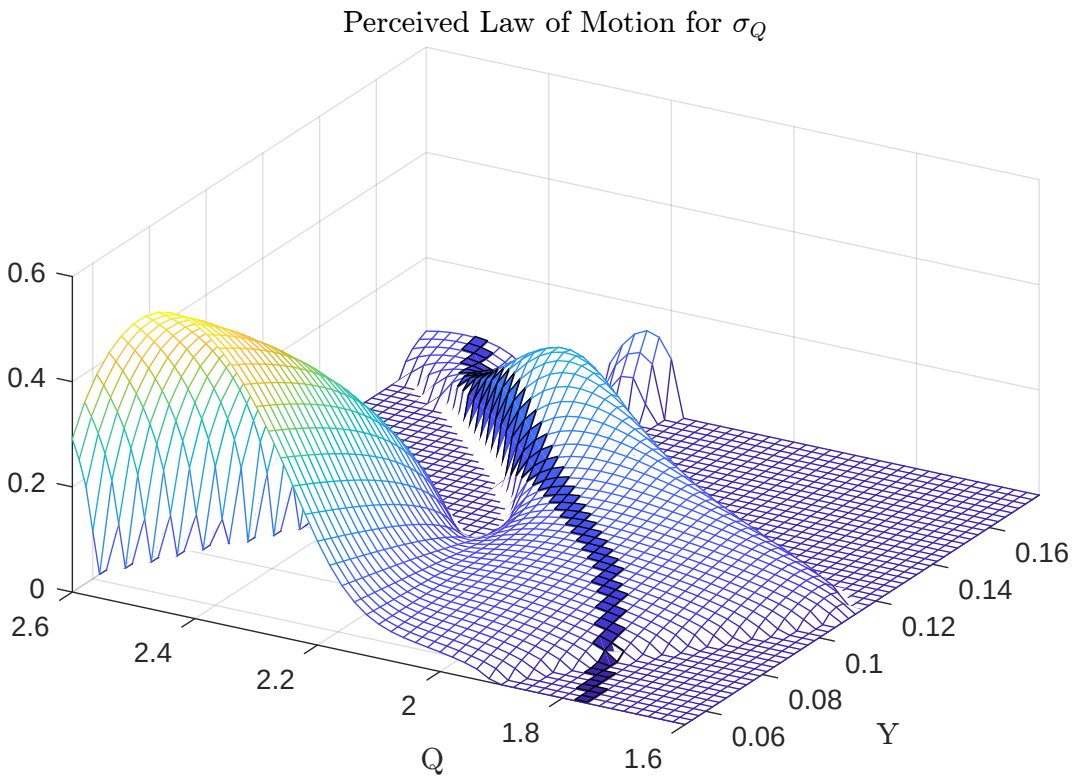


Figure 2: Perceived law of motion: diffusion $\sigma_Q(Y, Q)$.

Notes: The surface shows the learned diffusion function $\sigma_Q(Y, Q)$. The nonlinear shape indicates that price volatility varies substantially with the aggregate state, requiring a flexible approximation that a linear specification cannot provide.

The economic forces behind this state-dependent volatility – the interaction between individual portfolio decisions and market clearing – are examined in section 4.4, after we document the aggregate dynamics of the model.

4.2 Self-Confirming Equilibrium

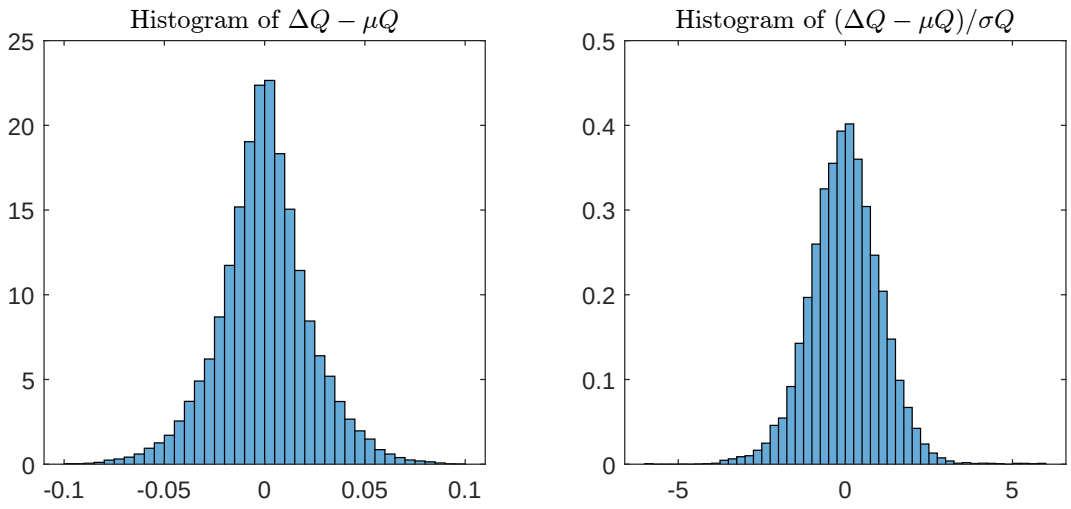


Figure 3: Out-of-sample validation of the perceived law of motion.

Notes: Left panel: histogram of raw forecast errors $\Delta Q - \mu_Q(Y_t, Q_t)\,\Delta t$. Right panel: histogram of standardized residuals $\frac{\Delta Q - \mu_Q(Y_t,Q_t)\Delta t}{\sigma_Q(Y_t,Q_t)\sqrt{\Delta t}}$. The standardized residuals are approximately standard normal, confirming that the PLM correctly captures both drift and diffusion.

We now verify that the converged PLM constitutes a self-confirming equilibrium. Figure 3 displays two panels: the histogram of raw forecast errors (left) and the histogram of standardized residuals (right). While the raw forecast errors clearly confirm that agents are not systematically over- or under-predicting price movements (mean of -0.0002), the standardized residuals also allow for a stronger test: if the PLM correctly captures both the drift and the diffusion, the normalized residuals should be normally distributed.[29] We test this by simulating the economy 1,000 times for 50 years each, discarding the initial 25 years to eliminate transient dynamics. Within the ergodic region of the state space – excluding the tails of the dividend distribution where the learned price volatility is near zero and the normalization is ill-conditioned – a Kolmogorov-Smirnov test rejects the null of normally distributed innovations in approximately 5% of simulations at the 5% significance level, consistent with the nominal size of the test. In the terminology of Fudenberg and Levine (1993), agents' beliefs are confirmed by their experience along the equilibrium path (Cho and Sargent 2018). This normality property ensures that no agent, upon observing the realized path of prices, would have reason to revise the perceived law of motion.

Finally, Figure 4 displays the joint ergodic distribution of dividends Y and asset prices Q. The distribution is well-defined and concentrated in a compact region of the state space

[29] From Equation (20) we know that $\frac{dQ - \mu_Q(Y,Q)dt}{\sigma_Q(Y,Q)} \sim dW$ which implies $\frac{\Delta Q - \mu_Q(Y_t,Q_t)\Delta t}{\sigma_Q(Y_t,Q_t)\sqrt{\Delta t}} \underset{\Delta t \to 0}{\longrightarrow} \mathcal{N}(0,1)$.

– around $Q \in [1.8, 2.2]$ and $Y \in [0.08, 0.14]$ – confirming the long-run stability of the equilibrium. This region corresponds to the darker band in the PLM surfaces of figures 1 and 2.

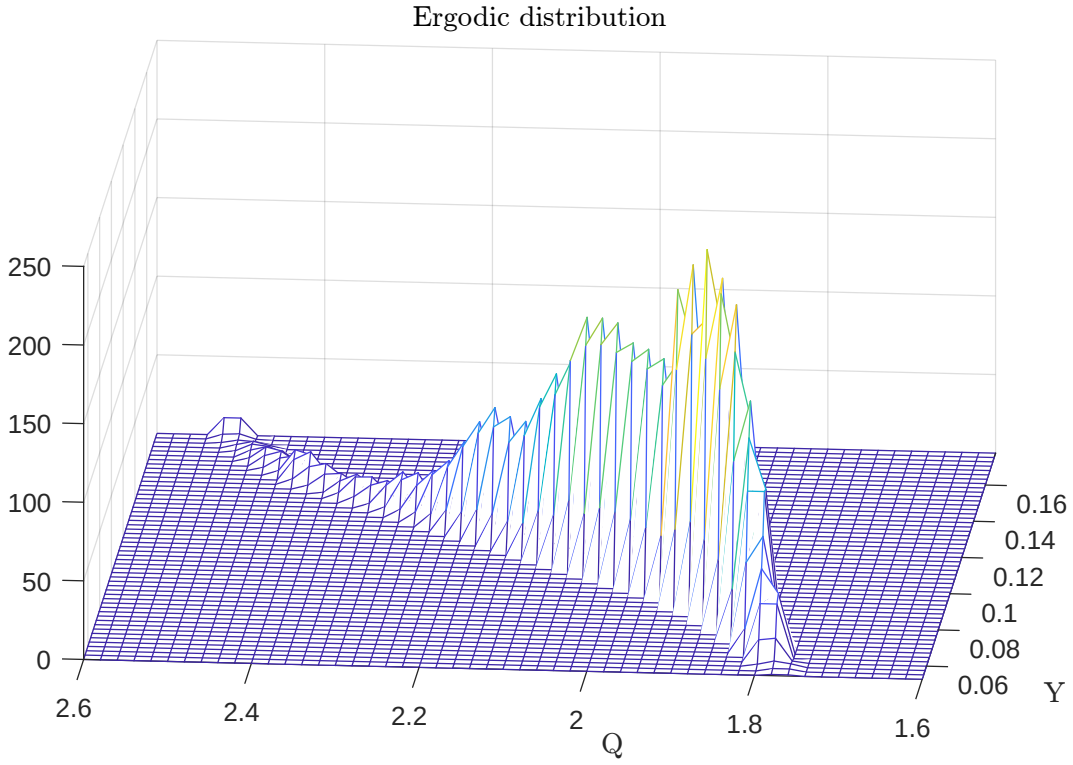


Figure 4: Joint ergodic distribution of dividends Y and asset prices Q.

Notes: The distribution is concentrated around $Q \in [1.8, 2.2]$ and $Y \in [0.08, 0.14]$, confirming the long-run stability of the equilibrium. Prices and dividends co-move positively, but the distribution of Q is more concentrated than that of Y, reflecting the smoothing of dividend shocks through equilibrium prices.

The ergodic distribution has several notable features: First, dividends and prices are strongly positively correlated, as one would expect from market clearing: higher dividends raise the return on equity, which increases asset demand and pushes up the equilibrium price. Second, the distribution of prices is more concentrated than the distribution of dividends – that is, prices are less volatile than the underlying fundamentals. This smoothing of dividend shocks through prices is consistent with the mean-reverting drift shown in figure 1: when a positive dividend shock pushes prices up, the negative drift pulls them back toward their long-run level, partially absorbing the shock. In this sense, the equilibrium price acts as a buffer between dividend fluctuations and household wealth.

Third, the ergodic distribution is very thin in the Y-Q space, almost a line: each dividend maps almost one-to-one to a price.[30] If different wealth distributions, corresponding to

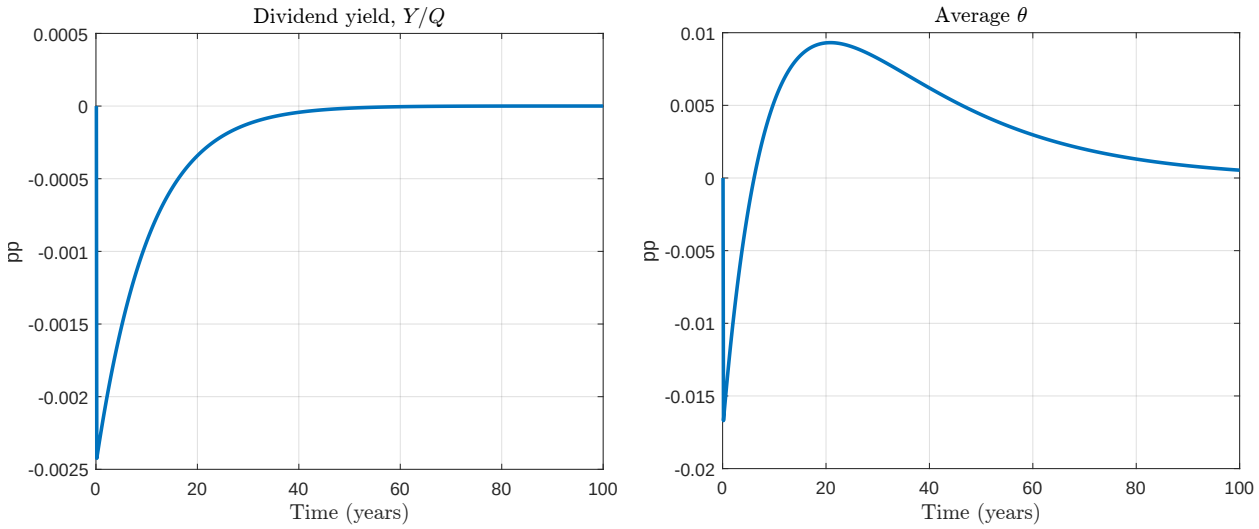
[30] See figure 16 in appendix B.

different prices, were visited in equilibrium for a given dividend, the ergodic region would display a thicker area. The fact that this is not the case should be taken as evidence that each dividend maps to a certain wealth distribution with an associated price. Combining this evidence with the PLM in Figures Figures 1 and 2, future prices depend (almost) exclusively on current prices and dividends, which supports our assumption of Markovian prices.

Finally: because prices move less than one-for-one with dividends, the price-dividend ratio Q/Y *falls* along the ergodic path as dividends (and prices) rise. Figure 14 in appendix B documents this pattern: Q/Y ranges between 30 in low-dividend, low-price states to around 14 in high-dividend, high-price states. When compared to the subjective-beliefs learning model of Adam, Marcet, and Nicolini (2016), in which belief momentum in expected capital gains feeds back into prices, amplifying swings in the price-dividend ratio and generating sustained booms and busts, in our limited-information economy mean-reversion dampens price movements rather than amplifying them, which keeps our equilibrium closer to a rational-expectations benchmark in which prices smooth fundamentals. We return to the role of the inverse mapping between Q/Y and the level of prices in section 4.4 when interpreting the state-dependence in the model's impulse responses, where the low-Q/Y states are precisely the high-price, high-dividend states.

4.3 Aggregate Dynamics and Impulse Responses

We now trace the aggregate dynamics of the model by computing the impulse response to a negative aggregate dividend shock of 0.1%, starting from the stochastic steady state. Because the shock is deliberately small, the exercise confines the economy to a small neighborhood of the stochastic steady state and the responses reported here trace the model's dynamics locally, while the state dependence of the transmission is documented separately in section 4.4.



(a) Dividend yield Y/Q.

(b) Average portfolio share θ^*.

Figure 5: Impulse response of the dividend yield and average portfolio share.
Notes: Left panel: response of the dividend yield Y/Q following the negative aggregate shock. Right panel: response of the average portfolio share θ^*, which drops on impact and subsequently overshoots, reflecting a portfolio rebalancing cycle.

Figure 5 presents the impulse response of two aggregate quantities that summarize the shock's transmission: the dividend yield Y/Q (left panel) and the average portfolio share θ^* (right panel). The two series display strikingly different dynamics. The dividend yield drops on impact and then recovers monotonically toward its steady-state level; the average portfolio share, by contrast, drops on impact and then *overshoots* – rising above its pre-shock level before gradually reverting.

A first observation is that the price decline is smaller than the dividend decline and that prices recover faster (Section B reports the responses of Y and Q separately). This is consistent with the mean-reverting drift and low volatility at low prices documented in figures 1 and 2: equilibrium prices partially absorb fundamental shocks rather than amplifying them. The asymmetry in speed reflects two distinct forces: the slower dividend recovery is governed by the exogenous mean-reversion of the Y process (through the parameter η), while the faster price adjustment arises endogenously through portfolio rebalancing and the market-clearing condition.

The second, and more interesting, observation is that the non-monotone response in θ^* is precisely the mechanism that absorbs the shock. Explaining it requires two objects not yet developed: households' optimal portfolio choice over the aggregate state and the time-varying risk premia to which it responds. section 4.4 analyzes these two elements before returning to this impulse response to show that the response reflects two distinct channels: a valuation channel and a portfolio-rebalancing channel.

4.4 Portfolio Allocation and Time-Varying Risk Premia

The impulse responses of section 4.3 left open the question of what drives the non-monotonicity in the portfolio response; this subsection looks at households' optimal portfolio choice and the time-varying risk premia to which it responds, and then returns to the impulse response with both in hand.

Optimal portfolio share. Figure 6 displays both the optimal portfolio share θ^* for an employed household at average wealth, plotted as a surface over (Y, Q), and the average portfolio share realized across the simulation. As we show below, the exact level of income or wealth at which we show this surface turns out not to matter.[31] As is often the case in "frictionless" portfolio choice models, the figure shows that households' optimal share is a step function: when prices are low relative to dividends households want to hold all of their wealth in equity ($\theta^* = 1$); conversely, when prices are high, households want to sell their risky-asset holdings ($\theta^* = 0$). The transition region between these two "regimes" is extremely narrow, and its location is common across households, a uniformity we explain below through the portfolio first-order condition. The blue scatter represents the average θ^* across the simulation, which ranges approximately between 0.53 and 0.65 and lies almost exactly on the transition region.

[31] To improve readability, the Q axis in figure 6 is inverted, with higher prices on the left and lower prices on the right.

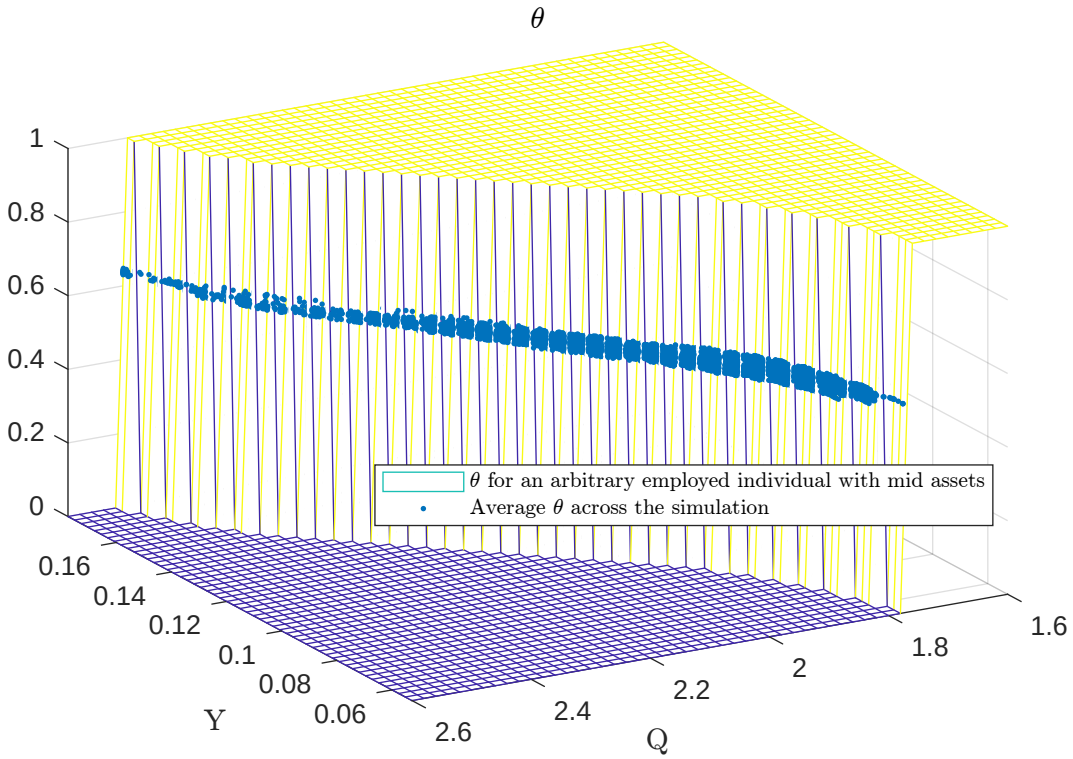


Figure 6: Optimal portfolio share θ^*.

Notes: The yellow surface shows θ^* for an employed household at average wealth as a function of dividends Y and prices Q (the Q axis is inverted: low prices on the right, high prices on the left). The surface is a step function: $\theta^* = 0$ at high Q and $\theta^* = 1$ at low Q, with a sharp transition region in between. The blue scatter shows the average θ^* across the simulation (approximately 0.53–0.65), which lies on the transition region.

The fact that optimal portfolio choice is a step function and that the simulation closely tracks the transition region is not incidental; rather, it reflects an equilibrium condition enforced through market clearing. If the aggregate state (Y, Q) were in the region where $\theta^* = 1$, that would mean the average household would be allocating its entire portfolio to equity, and the resulting excess demand would drive prices upward. Conversely, if (Y, Q) were in the region where $\theta^* = 0$, it would mean the average household would exit the equity market entirely, and the resulting excess supply would drive prices downward. The market-clearing condition therefore restricts the equilibrium to lie *on* the transition region, where prices adjust so that portfolio choice (when aggregated) is at an interior solution and the market clears. In fact, it turns out that in our economy, households portfolio choice depends exclusively on the aggregate states, i.e. $\theta^*(a, z; Y, Q) = \theta(Y_t, Q_t)$.[32] This implies

[32]Figure 15 in appendix B plot the optimal portfolio choice profile across income and wealth for different values of (Y, Q) and shows that it is completely flat across wealth and the same whether households are

both that the asset price is effectively proportional to aggregate wealth and that the average portfolio share plotted in figure 6 (and in figure 5) equals Q_t/A_t. To see this note that when $\theta^*(a, z; Y, Q) = \theta(Y_t, Q_t)$, market clearing (31) collapses to $Q_t = \theta(Y_t, Q_t)\, A_t$, where $A_t \equiv \int a\, g_t(a, z)\, da\, dz$ denotes aggregate wealth. Note also that this proportionality is not something we impose but a property of the solved equilibrium: as we discussed in section 2, if portfolio shares were constant the asset price would be a constant fraction of aggregate wealth, our model delivers a state-dependent version of that result.

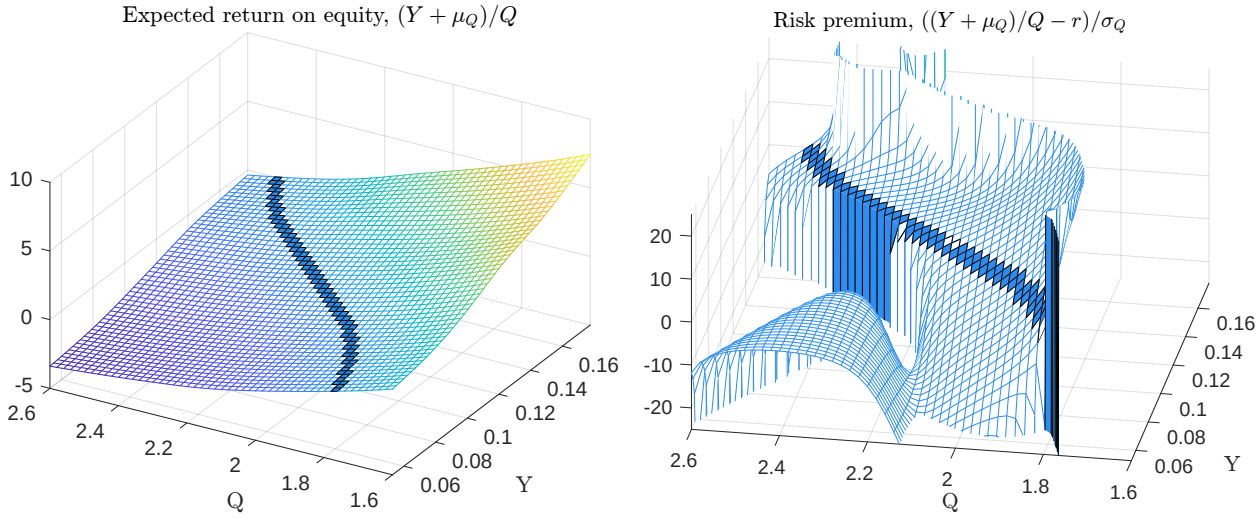
The first-order condition equation (27) derived in section 2 explains why: quantitatively, the two hedging demand terms turn out to be negligible in most of the visited region, while most of the variation in optimal portfolio choice is coming from the risk-return tradeoff term. That term is itself the product of a household-specific risk tolerance and a risk-return ratio $(R(Y, Q) - r^*)/(\sigma_Q/Q)^2$ – where $R(Y, Q) = (Y + \mu_Q)/Q$ is the expected equity return defined in section 2 – which is common to all households. While risk tolerance does vary across the wealth distribution, its effect is dwarfed by variation in the risk-return ratio as (Q, Y) varies. This uniformity is partially also a result of the deliberate simplicity of the model: CRRA preferences, limited income heterogeneity, and a portfolio problem that is frictionless apart from the no-borrowing and short-sale constraints (with no participation costs or similar frictions) all contribute to limited cross-sectional heterogeneity in portfolio shares. In this model, heterogeneity in equity exposure is mainly a level effect, $\theta \cdot a$ – the dimension along which aggregate shocks reshape the distribution, as we show in section 4.5 – rather than a share effect.

In fact, note that the transition region's location is pinned down by the drift and diffusion surfaces (figures 1 and 2) independently of the cross-sectional wealth distribution. The wealth distribution only enters through aggregate wealth A_t in the market clearing equation $Q_t = \theta(Y_t, Q_t)\, A_t$, by "selecting" the point on the transition region at which households' optimal portfolio choice and market clearing jointly pin down the equilibrium price. It is this interaction between households' optimal portfolio choice, market-clearing conditions, and the perceived law of motion that makes the model both interesting and computationally demanding: the market-clearing condition requires sufficient grid resolution in the transition region, and the nonlinear PLM must correctly position this transition in the (Y, Q) space at each outer iteration of the algorithm.

Risk premium. Within the transition region, small movements in the aggregate state (Y, Q) translate into large changes in aggregate equity demand. The risk-return tradeoff driving these portfolio decisions is itself state-dependent; Figure 7 plots the expected equity return $R(Y, Q) = (Y + \mu_Q)/Q$ and the risk premium $((Y + \mu_Q)/Q - r^*)/\sigma_Q$ as surfaces over

employed or unemployed.

the aggregate state. Away from the ergodic region, expected returns climb as the state approaches the low-price, high-dividend corner of the state space. Along the simulated ergodic path, the mean expected return comes to 5.11% with a standard deviation of 2.15 p.p.[33] Sorting ergodic states by the price-dividend ratio Q/Y reveals a hump-shaped premium: about 0.5 p.p. in the highest-Q/Y quintile, peaking around 2.5 p.p. in the middle of the range, and about 0.8 p.p. in the lowest-Q/Y quintile. The excess return itself decomposes into two components: the dividend yield net of r^* and the expected capital gain μ_Q/Q. At the bottom of the ergodic region the dividend yield sits almost exactly at r^*, about 4.0%, so the premium there is small and mostly driven by the positive expected capital gains.[34] Moving up the ergodic region toward high-price, high-dividend states, the dividend yield rises steadily, reaching about 5.9%, but the expected capital gain μ_Q/Q moves against it: positive on average in the middle, where it is largest, and negative at the top, where prices are high and the perceived mean-reverting drift pulls them back down. The premium therefore peaks in the middle of the range, where a large dividend yield combines with the largest capital-gain component, and stays small at both ends of the ergodic region.



(a) Expected equity return $(Y+\mu_Q)/Q$.

(b) Risk premium $((Y+\mu_Q)/Q-r^*)/\sigma_Q$.

Figure 7: Expected equity return and risk premium over the aggregate state (Y, Q).

Notes: Left panel: expected return on equity $(Y+\mu_Q)/Q$ as a function of dividends Y and asset prices Q. Right panel: risk premium, defined as excess return per unit of perceived volatility $((Y+\mu_Q)/Q-r^*)/\sigma_Q$. The darker band indicates the ergodic region.

Across the same quintiles, perceived volatility $\sigma_Q(Y,Q)/Q$ declines monotonically from

[33] This is against a risk-free rate of 4%, so a mean equity premium of 1.11 p.p.

[34] Recall that, as discussed in section 4.2, the low-price, low-dividend states correspond to the highest values of Q/Y.

about 4.6% to about 1.1%. In fact, because $\sigma_Q(Y, Q)$ varies with the aggregate state, so does the risk borne by equity holders (and therefore the compensation they require). Holding dividends fixed, at lower prices the dividend yield Y/Q is higher and the mean-reverting drift points upward, while volatility is near zero – producing a large expected excess return at low risk. At higher prices the dividend yield is lower and the drift turns negative, while volatility is elevated – compressing the premium.

Hence, as a direct consequence of state-dependent volatility, our model endogenously generates a negative conditional relationship between prices and expected excess returns: for a given level of dividends, higher prices are associated with lower expected excess returns.[35] This conditional relationship runs in the same direction as the well-documented empirical pattern in the finance literature that a high dividend-price ratio predicts higher subsequent excess equity returns (Campbell and Shiller 1988; Fama and French 1988; Cochrane 2011). We stress that the empirical pattern is an unconditional predictive relationship whose strength builds over multi-year horizons, while ours is a property of the model's conditional-expectation surfaces; we do not compute the corresponding unconditional regressions here.

These portfolio and risk-premium properties explain the impulse response dynamics documented in section 4.3. The non-monotone response of θ^* to the negative dividend shock visible in figure 5b reflects two channels that operate at different horizons.

Consider first the *valuation channel*: the direct effect of the shock on equity values and dividend income. By market clearing, the average share θ^* equals the ratio of the market value of equity to aggregate wealth, Q_t/A_t; its impulse response therefore decomposes exactly into a price component Q_t and a wealth component A_t at every horizon. On impact, prices fall by roughly 1.7 times as much as aggregate wealth (-0.054 versus -0.032; figure 13 in appendix B plots both paths). Specifically, only the equity slice of household wealth is revalued on impact, which is why wealth falls by less than prices. The on-impact decline visible in figure 5b is this valuation effect: prices repricing by more than aggregate wealth.

The *portfolio rebalancing channel* instead relies on changes in the risk-return tradeoff. On impact, the excess return $R(Y, Q) - r^*$ rises by 13.3 basis points. The rise is almost entirely due to the strengthened expected capital gain – a higher mean-reverting drift at lower prices – which contributes 13.5 basis points, while the dividend yield itself edges down slightly, since dividends fall by more than prices (figure 5a), while the perceived variance falls. The rebalancing incentive is strongest at impact and dominates the recovery: households substitute toward equity, and the average share recovers, eventually overshooting its pre-shock level. This is time-varying risk premia in action: the negative shock temporarily creates a regime of high expected returns at low perceived risk, which the portfolio surface

[35] Notice that this is not at odds with the hump-shaped premium across Q/Y quintiles documented above: sorting by Q/Y moves along the ergodic region, where dividends and prices rise together, whereas here we hold dividends fixed and vary the price.

reflects through an upward shift in θ^*.

The overshoot in the average share reflects, through market clearing, prices recovering toward their pre-shock level while aggregate wealth remains depressed for longer. Q_t/A_t therefore rises above its pre-shock value, peaking roughly two decades after the shock. Households are willing to hold the higher share because expected excess returns stay elevated throughout the recovery. The entire impulse response traces the interaction between the step-function portfolio surface and the time-varying risk premia documented above: the transition region is where small movements in (Y, Q) translate into meaningful changes in aggregate equity demand, and the market-clearing condition converts these demand shifts into the price dynamics that close the loop.

Figure 8 confirms how the time-variation visible in the static surfaces of figure 7 plays out dynamically following an aggregate shock. When the negative dividend shock hits, prices fall and the economy moves to a region of higher expected returns and lower volatility – precisely the pattern documented in figures 1 and 2 – so both the expected equity return and the risk premium spike on impact. As prices recover, the economy gradually returns to the ergodic region and the premium decays; the magnitudes are small because the shock itself is small (a 0.1% decline in dividends), but the pattern is sharp and confirms that even modest shocks produce meaningful temporary variation in the equity premium through the nonlinear interaction between prices, volatility, and portfolio choice.

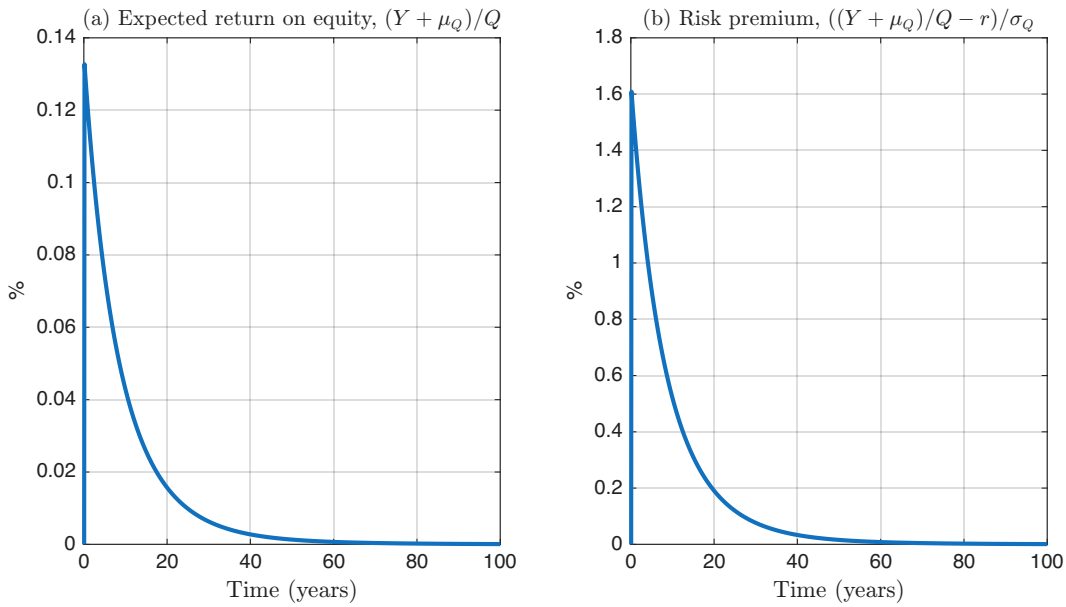


Figure 8: Impulse response of the expected equity return and risk premium to a negative aggregate shock.

Notes: Left panel: expected return on equity $(Y+\mu_Q)/Q$. Right panel: risk premium $((Y + \mu_Q)/Q - r^*)/\sigma_Q$. Both respond to the same negative dividend shock used in section 4.3. The expected return and risk premium spike on impact and decay monotonically.

The state dependence of σ_Q documented in figure 2 is driven by how dividend shocks propagate differently through the economy depending on its current state. Consider a positive dividend shock: the direct effect is to raise returns for equity holders; a distributional effect then shifts wealth toward households with high equity exposure, increasing aggregate asset demand; higher prices then reduce expected future returns and dampen demand. These channels interact nonlinearly, and the market-clearing condition makes the mechanism precise: high prices require high aggregate demand for the risky asset, which in turn means that more households are actively participating in the equity market and are away from their borrowing constraints. In such states, a given dividend shock affects more portfolios simultaneously, amplifying its impact on prices through the redistribution channel. When prices are low, by contrast, many households will be near their constraints with low equity exposure, dampening the transmission of aggregate shocks to prices.

Figure 9 makes this state dependence explicit by repeating the impulse-response experiment from different points of the ergodic distribution. Each dot reports the cumulative response of the asset price to the same negative dividend shock, as a function of the price-dividend ratio Q/Y at the moment the shock arrives. The cumulative decline is roughly two and a half times larger when the shock hits states with low Q/Y – which, as docu-

mented in section 4.2, are along the ergodic path the high-price, high-dividend states where equity-market participation is broadest – than when it hits high-Q/Y states, where more households sit near their constraints and transmission is dampened. The same shock thus produces systematically different price dynamics depending on the state of the economy – a form of state dependence that linear solution methods rule out by construction.

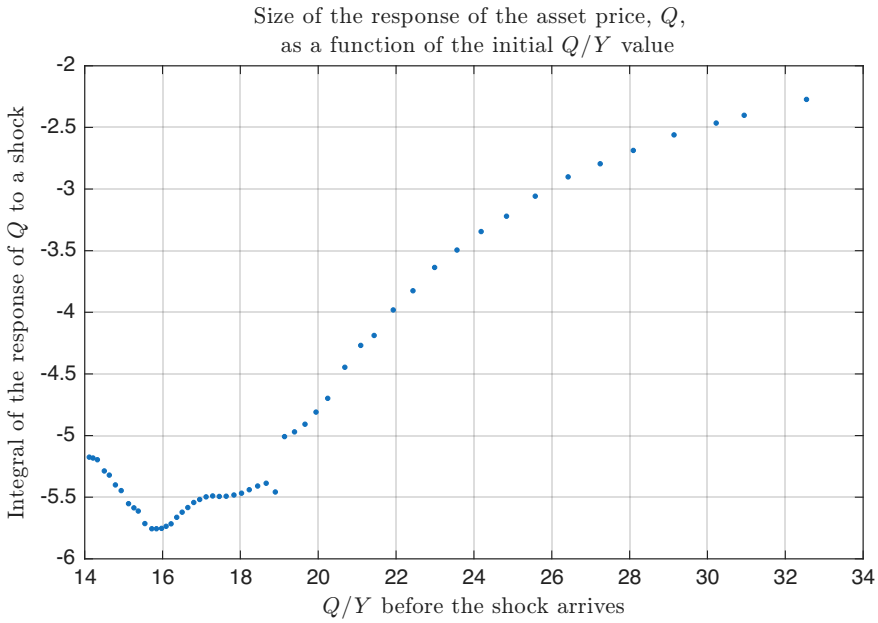


Figure 9: State-dependent impulse responses.

Notes: Cumulative response of the asset price Q to a negative dividend shock, as a function of the price-dividend ratio Q/Y in the period before the shock arrives. Each dot corresponds to a bin of initial states drawn from the simulated ergodic distribution.

This whole mechanism – from the step-function portfolio surface to time-varying risk premia to the joint recovery of prices and portfolio shares – arises naturally from our global nonlinear solution under limited information with optimal portfolio choice, and it is a pattern that would *not* emerge from a linear model, nor from a local approximation to a nonlinear model.

4.5 Distributional Dynamics

The mechanisms documented in section 4.4 also reshapes the cross-sectional wealth distribution.[36] Figure 10 presents the $g_t(a, z)$ at two representative aggregate states: a high-dividend

[36]Note that these distributional changes are present under common beliefs and, as we showed in section 4.4, a common portfolio share: what differs across households is equity exposure *in levels*, $\theta^* \cdot a$, which scales with wealth, while labor income is completely insulated from the aggregate shock and therefore cushions

state with correspondingly high prices ($Q/Y \approx 15$) and a low-dividend state with low prices ($Q/Y \approx 26$). In the low-dividend state more mass sits near the borrowing constraint and the right tail is thinner; in the high-dividend state, instead, the distribution is more dispersed, with a thicker right tail. Aggregate wealth is about 18% higher in the high-dividend state relative to the low-dividend state with corresponding average portfolio shares of 0.63 and 0.55, near the two ends of the simulated scatter in figure 6.

Higher dividends sustain higher saving which, together with the capital gains accumulated along the path into such states, shifts the wealth distribution to the right. The logic can also be read in reverse: states with high dividends carry a larger mass of wealthy households and, with the portfolio share pinned down by the aggregate state, market clearing $Q_t = \theta(Y_t, Q_t)\, A_t$ translates the larger aggregate wealth into higher equity demand, and therefore into a higher market value of the risky asset.

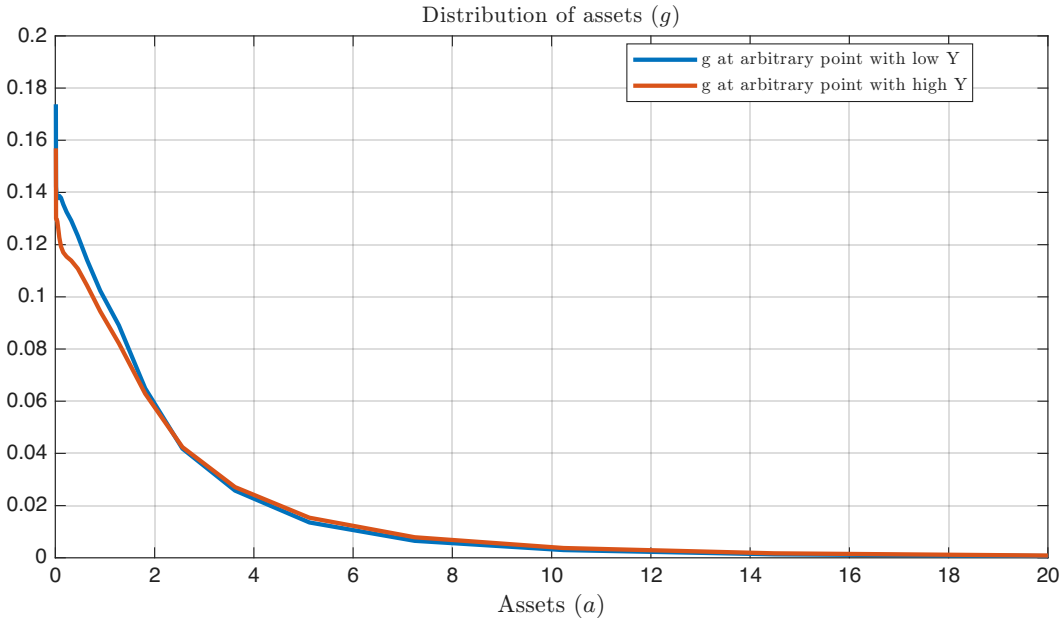


Figure 10: Cross-sectional wealth distribution across states.

Notes: The figure shows the cross-sectional wealth distribution $g_t(a, z)$ at two states: a high-Y state and a low-Y state. When dividends are low, the distribution carries more mass near the borrowing constraint and a thinner right tail; when dividends are high, it is more dispersed, with a thicker right tail, reflecting state-dependent wealth inequality driven by households' equity exposure.

The distribution is therefore also time-varying: aggregate shocks reshape the wealth distribution through households' heterogeneous equity exposure. This mechanism has a

poorer households proportionally more. Hence, the same proportional revaluation of equity holdings moves total resources very differently along the wealth distribution.

precise footprint in the distribution's response to the IRF shock, which we examine next.

Figure 11 completes the picture by showing the change in the aggregate wealth distribution Δg relative to the pre-shock stochastic steady state, 82 months after the negative aggregate shock – which corresponds to the moment the asset price Q is halfway back to its steady-state level. The distributional footprint of the shock is concentrated at two locations: near the borrowing constraint, where there is a positive spike reflecting a mass of households whose wealth is compressed toward the lower bound; and farther out in the right tail, where a shallow trough reflects capital losses among wealthier households with higher equity exposure. The directional pattern is precisely what differential equity exposure predicts: the shock redistributes wealth toward the lower end of the distribution, because households near the borrowing constraint have little equity exposure and are therefore partially shielded, while wealthier households bear the brunt of the decline in equity values.

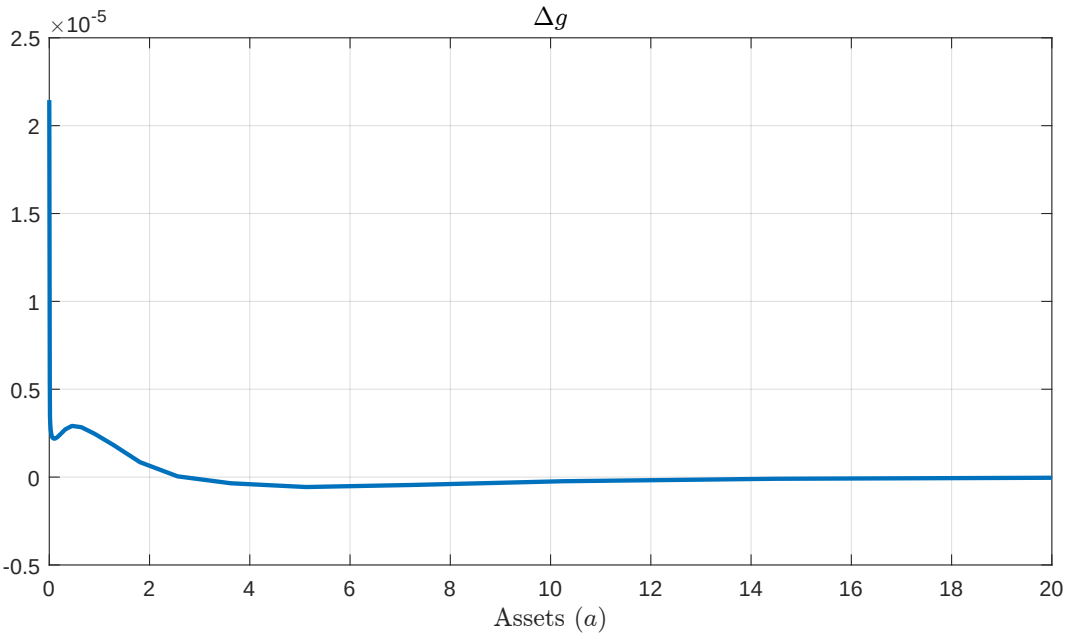


Figure 11: Change in the cross-sectional wealth distribution following a negative aggregate shock.

Notes: The figure shows the change $\Delta g(a)$ in the aggregate wealth distribution relative to the pre-shock stochastic steady state, 82 months after the negative aggregate shock – which corresponds to the moment the asset price Q is halfway back to its steady-state level. The positive spike near $a = 0$ and the shallow trough in the right tail reflect the heterogeneous impact of the shock through differential equity exposure.

This two-way feedback between the distribution and prices – aggregate shocks reshape the distribution through heterogeneous portfolios, and the reshaped distribution feeds back

into asset demand and prices through market clearing – is precisely the main mechanism that motivates the paper. Our limited information framework makes it tractable to study this feedback quantitatively: agents need not track the infinite-dimensional distribution that they are collectively reshaping, yet the equilibrium they sustain exhibits rich distributional dynamics via the interaction of aggregate risk and household heterogeneity, time-varying wealth inequality and risk premia, and portfolio rebalancing.

5 Conclusion

In this paper, we develop a framework for studying the joint dynamics of wealth inequality and asset prices in an economy with heterogeneous agents, endogenous portfolio choice, and aggregate risk. By restricting attention to a limited information environment – in which agents forecast prices directly using observable aggregate variables – we sidestep the Master equation and make the problem not only computationally tractable but also economically more plausible: agents in our model solve the kind of simple forecasting problems that real-world investors actually face, rather than tracking the infinite-dimensional distributions that rational expectations would require. Neural networks approximate the drift and diffusion of the perceived price process, capturing nonlinear relationships without imposing restrictive functional forms. We show that this combination uncovers a key asymmetry: the drift of prices is close to linear, capturing mean-reversion, but the diffusion is highly nonlinear and state-dependent, which endogenously generates time variation in risk premia that local approximation methods cannot produce. This time variation, combined with heterogeneous portfolio exposure, produces a two-way feedback between the wealth distribution and asset prices: aggregate shocks reshape the distribution through differential capital gains and losses; the reshaped distribution feeds back into aggregate asset demand and, through market clearing, into prices.

Because our primary objective is methodological – to show that the limited information framework can solve a class of models that has so far resisted global solution methods – we keep the economic environment deliberately simple so that the mechanism remains transparent. An exogenous risk-free rate isolates the feedback between equity prices and the wealth distribution by reducing the problem to a single endogenous price; a two-state income process captures the essential features of employment transitions and income heterogeneity without adding complexity that is orthogonal to the mechanism. We also do not attempt to match asset-pricing moments quantitatively: doing so would require taking a stance on the sources of realistic risk premia – and the finance literature itself does not have a consensus on which mechanisms generate high equity premia, high return volatility, and time-varying risk premia simultaneously.

The tractability of our framework makes it possible to study questions that were previously computationally out of reach. One question concerns welfare: in an economy where portfolios are heterogeneous and risk premia are time-varying, who gains and who loses from equity booms and busts, and by how much? Another concerns policy: fiscal redistribution, macroprudential regulation, or monetary policy interventions can be evaluated in an environment where the distribution and asset prices co-evolve – an environment in which policy effects propagate through both the direct channel and the distributional feedback channel. A third concerns information itself: how do different information structures – what agents observe, how they form beliefs about prices – shape distributional dynamics and equilibrium outcomes? These questions, while beyond the scope of this paper, illustrate how the combination of limited information and neural-network approximation can be applied to richer environments where the interplay between heterogeneity and aggregate risk is central.

References

Achdou, Y., J. Han, J.-M. Lasry, P.-L. Lions, and B. Moll (2022). "Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach". In: *The Review of Economic Studies* 89.1, pp. 45–86.

Adam, K. and A. Marcet (2011). "Internal Rationality, Imperfect Market Knowledge and Asset Prices". In: *Journal of Economic Theory* 146.3, pp. 1224–1252.

Adam, K., A. Marcet, and J. Beutel (2017). "Stock Price Booms and Expected Capital Gains". In: *American Economic Review* 107.8, pp. 2352–2408.

Adam, K., A. Marcet, and J. P. Nicolini (2016). "Stock Market Volatility and Learning". In: *The Journal of Finance* 71.1, pp. 33–82.

Auclert, A., B. Bardóczy, M. Rognlie, and L. Straub (2021). "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models". In: *Econometrica* 89.5, pp. 2375–2408.

Auclert, A., M. Rognlie, L. Straub, and T. Tapak (2024). "When Do Endogenous Portfolios Matter for Hank". In: *Working Paper*.

Azinovic, M., L. Gaegauf, and S. Scheidegger (2022). "Deep Equilibrium Nets". In: *International Economic Review* 63.4, pp. 1471–1525.

Azinovic-Yang, M. and J. Žemlička (2023). "Intergenerational Consequences of Rare Disasters". In: *Working Paper*.

Bach, F. (2017). "Breaking the Curse of Dimensionality with Convex Neural Networks". In: *Journal of Machine Learning Research* 18.19, pp. 1–53.

Bach, L., L. E. Calvet, and P. Sodini (2020). "Rich Pickings? Risk, Return, and Skill in Household Wealth". In: *American Economic Review* 110.9, pp. 2703–2747.

Bansal, R. and A. Yaron (2004). "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles". In: *The Journal of Finance* 59.4, pp. 1481–1509.

Barron, A. (1993). "Universal Approximation Bounds for Superpositions of a Sigmoidal Function". In: *IEEE Transactions on Information Theory* 39.3, pp. 930–945.

Bayer, C. and R. Luetticke (2020). "Solving Discrete Time Heterogeneous Agent Models with Aggregate Risk and Many Idiosyncratic States by Perturbation". In: *Quantitative Economics* 11.4, pp. 1253–1288.

Belda, P., J. Heineken, and A. Ifrim (2025). "How Expectations Amplify Wealth Inequality". In: *Working Paper*.

Benhabib, J. and A. Bisin (2018). "Skewed Wealth Distributions: Theory and Empirics". In: *Journal of Economic Literature* 56.4, pp. 1261–1291.

Bhandari, A., T. Bourany, D. Evans, and M. Golosov (2023). "A Perturbational Approach for Approximating Heterogeneous Agent Models". In: *Working Paper*.

Bilal, A. (2023). "Solving Heterogeneous Agent Models with the Master Equation". In: *Working Paper*.

Blanchard, O. J. (1985). "Debt, Deficits, and Finite Horizons". In: *Journal of Political Economy* 93.2, pp. 223–247.

Branch, W. A. (2006). "Restricted Perceptions Equilibria and Learning in Macroeconomics". In: *Post Walrasian Macroeconomics: Beyond the Dynamic Stochastic General Equilibrium Model. Cambridge University Press, New York*, pp. 135–160.

Calvet, L. E. and P. Sodini (2014). "Twin Picks: Disentangling the Determinants of Risk-Taking in Household Portfolios". In: *The Journal of Finance* 69.2, pp. 867–906.

Campbell, J. Y. and J. H. Cochrane (1999). "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior". In: *Journal of Political Economy* 107.2, pp. 205–251.

Campbell, J. Y. and R. J. Shiller (1988). "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors". In: *The Review of Financial Studies* 1.3, pp. 195–228.

Cardaliaguet, P., F. Delarue, J.-M. Lasry, and P.-L. Lions (2019). *The Master Equation and the Convergence Problem in Mean Field Games*. Princeton: Princeton University Press.

Carmona, R. and F. Delarue (2018). *Probabilistic Theory of Mean Field Games with Applications I-II*. Springer.

Chan, Y. L. and L. Kogan (2002). "Catching Up with the Joneses: Heterogeneous Preferences and the Dynamics of Asset Prices". In: *Journal of Political Economy* 110.6, pp. 1255–1285.

Chien, Y., H. Cole, and H. Lustig (2012). "Is the Volatility of the Market Price of Risk Due to Intermittent Portfolio Rebalancing?" In: *American Economic Review* 102.6, pp. 2859–2896.

Chien, Y. and H. Lustig (2010). "The Market Price of Aggregate Risk and the Wealth Distribution". In: *The Review of Financial Studies* 23.4, pp. 1596–1650.

Cho, I.-K. and T. J. Sargent (2018). "Self-Confirming Equilibria". In: *The New Palgrave Dictionary of Economics*. Springer, pp. 12147–12151.

Cioffi, R. A. (2021). "Heterogeneous Risk Exposure and the Dynamics of Wealth Inequality". In: *Working Paper*.

Cochrane, J. H. (2011). "Presidential Address: Discount Rates". In: *The Journal of Finance* 66.4, pp. 1047–1108.

De La O, R. and S. Myers (2021). "Subjective Cash Flow and Discount Rate Expectations". In: *The Journal of Finance* 76.3, pp. 1339–1387.

Dumas, B. (1989). "Two-Person Dynamic Equilibrium in the Capital Market". In: *Review of Financial Studies* 2.2, pp. 157–188.

Fagereng, A., L. Guiso, D. Malacrino, and L. Pistaferri (2020). "Heterogeneity and Persistence in Returns to Wealth". In: *Econometrica* 88.1, pp. 115–170.

Fama, E. F. and K. R. French (1988). "Dividend Yields and Expected Stock Returns". In: *Journal of Financial Economics* 22.1, pp. 3–25.

Fernández-Villaverde, J., S. Hurtado, and G. Nuño (2023). "Financial Frictions and the Wealth Distribution". In: *Econometrica* 91.3, pp. 869–901.

Fernández-Villaverde, J. and O. Levintal (2024). "The Distributional Effects of Asset Returns". In: *Working Paper*.

Fudenberg, D. and D. K. Levine (1993). "Self-Confirming Equilibrium". In: *Econometrica* 61.3, pp. 523–545.

Gârleanu, N. and S. Panageas (2015). "Young, Old, Conservative, and Bold: The Implications of Heterogeneity and Finite Lives for Asset Pricing". In: *Journal of Political Economy* 123.3, pp. 670–685.

Gomes, F. and A. Michaelides (2008). "Asset Pricing with Limited Risk Sharing and Heterogeneous Agents". In: *The Review of Financial Studies* 21.1, pp. 415–449.

Gomez, M. (2025). "Wealth Inequality and Asset Prices". In: *The Review of Economic Studies* 92.6, pp. 3924–3967.

Gopalakrishna, G., Z. Gu, and J. Payne (2026). "Institutional Asset Pricing with Segmentation and Household Heterogeneity". In: *Working Paper*.

Greenwald, D. L., M. Leombroni, H. Lustig, and S. Van Nieuwerburgh (2023). "Financial and Total Wealth Inequality with Declining Interest Rates". In: *Working Paper*.

Greenwood, R. and A. Shleifer (2014). "Expectations of Returns and Expected Returns". In: *Review of Financial Studies* 27.3, pp. 714–746.

Gu, Z., M. Laurière, S. Merkel, and J. Payne (2023). "Deep Learning Solutions to Master Equations for Continuous Time Heterogeneous Agent Macroeconomic Models". In: *Working Paper*.

Guarda, S. (2025). "Narrow and Short Beliefs in Macroeconomics with Heterogeneous Agents". In: *Working Paper*.

Guvenen, F. (2009). "A Parsimonious Macroeconomic Model for Asset Pricing". In: *Econometrica* 77.6, pp. 1711–1750.

Hall, R. E. and P. R. Milgrom (2008). "The Limited Influence of Unemployment on the Wage Bargain". In: *American economic review* 98.4, pp. 1653–1674.

Han, J., Y. Yang, and W. E (2022). "DeepHAM: A Global Solution Method for Heterogeneous Agent Models with Aggregate Shocks". In: *Working Paper*.

Hornik, K., M. Stinchcombe, and H. White (1989). "Multilayer Feedforward Networks Are Universal Approximators". In: *Neural Networks* 2.5, pp. 359–366.

Hubmer, J., E. Halvorsen, S. Salgado, and S. Ozkan (2023). "Why Are the Wealthiest So Wealthy? New Longitudinal Empirical Evidence and Implications for Theories of Wealth Inequality". In: *Working Paper*.

Hubmer, J., P. Krusell, and A. A. Smith (2021). "Sources of U.S. Wealth Inequality: Past, Present, and Future". In: *NBER Macroeconomics Annual* 35.1, pp. 391–455.

Kloeden, P. E. and E. Platen (1995). *Numerical Solution of Stochastic Differential Equations*. Springer.

Krueger, D. and H. Lustig (2010). "When Is Market Incompleteness Irrelevant for the Price of Aggregate Risk (and When Is It Not)?" In: *Journal of Economic Theory* 145.1, pp. 1–41.

Krusell, P. and A. Smith (1997). "Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns". In: *Macroeconomic Dynamics* 1.2, pp. 387–422.

— (1998). "Income and Wealth Heterogeneity in the Macroeconomy". In: *Journal of Political Economy* 106.5, pp. 867–896.

Kuhn, M., M. Schularick, and U. I. Steins (2020). "Income and Wealth Inequality in America, 1949–2016". In: *Journal of Political Economy* 128.9, pp. 3469–3519.

Lasry, J.-M. and P.-L. Lions (2007). "Mean Field Games". In: *Japanese Journal of Mathematics* 2, pp. 229–260.

Lee, H. (2025). "Global Nonlinear Solutions in Sequence Space and the Generalized Transition Function". In: *Working Paper*.

Maliar, L., S. Maliar, and P. Winant (2021). "Deep Learning for Solving Dynamic Economic Models." In: *Journal of Monetary Economics* 122, pp. 76–101.

Marcet, A. and T. J. Sargent (1989). "Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models". In: *Journal of Economic Theory* 48.2, pp. 337–368.

Martin, I. W. R. and D. Papadimitriou (2022). "Sentiment and Speculation in a Market with Heterogeneous Beliefs". In: *American Economic Review* 112.8, pp. 2465–2517.

Meeuwis, M. (2020). "Wealth Fluctuations and Risk Preferences: Evidence from U.S. Investor Portfolios". In: *Working Paper*, p. 100.

Moll, B. (2024). "The Trouble with Rational Expectations in Heterogeneous Agent Models: A Challenge for Macroeconomics". In: *Working Paper*.

Moll, B. and L. Ryzhik (2025). "Mean Field Games without Rational Expectations". In: *Working Paper*.

Øksendal, B. (2014). *Stochastic Differential Equations: An Introduction with Applications*. 6th Ed. Springer.

Reiter, M. (2009). "Solving Heterogeneous-Agent Models by Projection and Perturbation". In: *Journal of Economic Dynamics and Control* 33.3, pp. 649–665.

Storesletten, K., C. I. Telmer, and A. Yaron (2007). "Asset Pricing with Idiosyncratic Risk and Overlapping Generations". In: *Review of Economic Dynamics* 10.4, pp. 519–548.

Toda, A. A. and K. J. Walsh (2019). "The Equity Premium and the One Percent". In: *The Review of Financial Studies*.

Yaari, M. E. (1965). "Uncertain Lifetime, Life Insurance, and the Theory of the Consumer1". In: *The Review of Economic Studies* 32.2, pp. 137–150.

Yang, Y., C. Wang, A. Schaab, and B. Moll (2025). "Structural Reinforcement Learning for Heterogeneous Agent Macroeconomics". In: *Working Paper*.

A KFE Derivation

In this appendix we show that the cross-sectional density $g_t(a,z)$ satisfies the equilibrium Kolmogorov forward equation (28), which we restate here for convenience:

$$dg_t(a,z) = \mathcal{A}_t^*[g_t](a,z)\,dt + \mathcal{B}_t^*[g_t](a,z)\,dW_t,$$

$$\begin{aligned}
\mathcal{A}_t^*[g](a,z) &= -\partial_a[\mu_a\, g] + \tfrac{1}{2}\,\partial_{aa}[\sigma_a^2\, g] + \lambda_{z'}\, g(a,z') - \lambda_z\, g(a,z) + \phi(\psi(a,z) - g(a,z)),\\
\mathcal{B}_t^*[g](a,z) &= -\partial_a\left[\theta^*(a,z;Y_t,Q_t)\, a\, \frac{\sigma_Q(Y_t,Q_t)}{Q_t}\, g(a,z)\right],
\end{aligned}$$

with μ_a and σ_a as in equations (23) and (24). The idea behind the derivation is to study the dynamics of a population with a finite number of households and then take the limit

as the number of households goes to infinity. We start with a population of N households; then, we apply Itô's lemma to a test function evaluated at each household's state, take the cross-sectional limit $N \to \infty$, and integrate by parts to get a differential equation for g_t directly.[37]

Setup. Let $g_t^N(a,z) = \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}_{\{z_{i,t}=z\}}\,\delta(a - a_{i,t})$ denote the empirical density of a population of N households – with δ the Dirac measure – and let $g_t = \lim_{N\to\infty} g_t^N$, conditional on the common-noise path, in the sense of distributions. For a density g and a test function f recall that

$$\langle g, f\rangle \equiv \sum_{z\in\{z_1,z_2\}} \int_0^\infty f(a,z)\, g(a,z)\, da.$$

Against the empirical measure, this boils down to the cross-sectional average: $\langle g_t^N, f\rangle = \frac{1}{N}\sum_i f(a_{i,t}, z_{i,t})$. Throughout the derivation, we make use of two assumptions: First, conditional on the common-noise path, the empirical measure admits a well-defined density limit g_t. Second, μ_a, σ_a, and g_t decay fast enough at the boundaries of the wealth domain that the boundary terms produced by the integrations by parts below all vanish.

Itô's lemma. Let $f(\cdot, z)$ be twice continuously differentiable with compact support in a, defined on both income states. Applying Itô's lemma for jump-diffusions to $f(a_{i,t}, z_{i,t})$ yields

$$df(a_{i,t}, z_{i,t}) = \big(\mathcal{A}_{a,t}f + \mathcal{A}_{z,t}f\big)(a_{i,t}, z_{i,t})\, dt + \sigma_a(a_{i,t}, z_{i,t}; Y_t, Q_t)\, \partial_a f(a_{i,t}, z_{i,t})\, dW_t + dM_{i,t},$$

with wealth generator $\mathcal{A}_{a,t}f = \mu_a\, \partial_a f + \frac{1}{2}\,\sigma_a^2\, \partial_{aa} f$ and income generator $\mathcal{A}_{z,t}f = \lambda_z\,[f(a,z') - f(a,z)]$, both with coefficients evaluated at $(a, z; Y_t, Q_t)$, and z' the other income state. The dt term, the average rate at which f changes, has both a continuous and a jump piece: the continuous piece comes from the wealth drift and from the second-order correction induced by the common return shock; the jump piece from the income switch that occurs at rate λ_z. The dW_t term is the household's exposure to the common shock through its portfolio; $dM_{i,t}$ is the compensated idiosyncratic income-switch martingale – the difference between the realized income jumps and their average, which $\mathcal{A}_z$ already accounts for. In short, $\mathcal{A}_z$ and $dM_{i,t}$ are the average and the realization of the same jumps.

[37] Under limited information each household takes the aggregate path (Y_t, Q_t) as given, so its own drift and diffusion coefficients are functions of its own state (a, z) and of the realized path (Y_t, Q_t) only – they do not depend on g_t. Conditional on that path, g_t therefore evolves according to a linear, finite-dimensional stochastic partial differential equation driven by the common shock W_t, and there is no Master equation to solve; the distribution feeds back on prices only through the market-clearing condition (31), condition (iii) of the equilibrium definition. For a formal treatment of the propagation-of-chaos limit, see Carmona and Delarue (2018).

Cross-sectional limit. Taking the cross-sectional limit of $\langle g_t^N, f\rangle$ we therefore recover:[38]

$$d\langle g_t, f\rangle = \big\langle g_t,\, (\mathcal{A}_c + \mathcal{A}_z) f\big\rangle\, dt + \big\langle \phi(\psi - g_t),\, f\big\rangle\, dt + \big\langle g_t,\, \sigma_a\, \partial_a f\big\rangle\, dW_t \tag{36}$$

where the middle term is the flow coming from death and birth: over an interval dt, a fraction ϕdt of households die and are replaced byt a newborn agent drawn from ψ. Each replacement therefore swaps f evaluated at the state at the moment of death for f evaluated at the state at the moment of birth. In the cross-sectional limit death and birth therefore add

$$\phi(\langle \psi, f\rangle - \langle g_t, f\rangle)\, dt = \langle \phi(\psi - g_t),\, f\rangle\, dt$$

to $d\langle g_t, f\rangle$. Notice that no operator acts on f in this term, so – unlike the others – it needs no integration by parts and will pass into equation (28) directly.

The adjoint operator. Equation (36) pairs g_t against the test function f; to obtain an equation for g_t itself, we need to rewrite each term so that the operators act on g_t rather than on f, which amounts to finding the adjoint $\mathcal{A}^*$ defined by $\langle g,\, \mathcal{A} f\rangle = \langle \mathcal{A}^* g,\, f\rangle$ for every test function f. Integrating the continuous piece by parts gives

$$\begin{aligned}
\big\langle g_t,\, \mu_a\, \partial_a f\big\rangle &= \big[\, \mu_a\, g_t\, f\,\big]_\Gamma - \big\langle \partial_a[\mu_a\, g_t],\, f\big\rangle, \\
\big\langle g_t,\, \tfrac{1}{2}\sigma_a^2\, \partial_{aa} f\big\rangle &= \big[\, \tfrac{1}{2}\sigma_a^2\, g_t\, \partial_a f - \tfrac{1}{2}\partial_a[\sigma_a^2\, g_t]\, f\,\big]_\Gamma + \big\langle \tfrac{1}{2}\, \partial_{aa}[\sigma_a^2\, g_t],\, f\big\rangle,
\end{aligned}$$

where $[\,\cdot\,]_\Gamma$ denotes evaluation on the boundary of the wealth domain. These boundary terms all vanish. As $a \to \infty$, f and $\partial_a f$ are zero because f has compact support. At $a = 0$, the diffusion coefficient vanishes because $\sigma_a = \theta^*\, a\, \sigma_Q/Q$ is multiplicative in a, and the borrowing constraint implies non-negative saving at the constraint (from equation (23), $\mu_a(0,z) = z - c^*(0,z) \geq 0$ whenever $c^*(0,z) \leq z$), so no probability flux exits the domain and we impose the corresponding state-constraint (zero-flux) boundary condition. The income adjoint follows simply by relabeling the two-state sum:

$$\sum_z \lambda_z\, [f(a,z') - f(a,z)]\, g(a,z) = \sum_z f(a,z)\, [\lambda_{z'}\, g(a,z') - \lambda_z\, g(a,z)].$$

Collecting the surviving terms, the adjoint acts as $-\partial_a[\mu_a\, g] + \frac{1}{2}\partial_{aa}[\sigma_a^2\, g] + \lambda_{z'}\, g(a,z') - \lambda_z\, g(a,z)$; together with the demographic flow $\phi(\psi(a,z) - g(a,z))$ from the birth-death term,

[38] Conditional on the common path, the idiosyncratic jump martingales $dM_{i,t}$ are mean-zero and independent across households, since their coefficients do not depend on g_t, so their cross-sectional average has conditional quadratic variation of order $1/N$ and vanishes in the limit; the same holds for the finite-N fluctuations of the birth-death events. The common shock dW_t, by contrast, multiplies every household's term and therefore survives the average, which is precisely why g_t is itself stochastic.

this is exactly the drift operator $\mathcal{A}_t^*[g]$ of equation (28).

The common-noise adjoint. It now only remains to take the adjoint of the surviving common-noise term in equation (36). One integration by parts – with boundary terms vanishing for the same reasons as above – gives

$$\big\langle g_t,\, \sigma_a\, \partial_a f \big\rangle\, dW_t = \big\langle -\partial_a[\sigma_a\, g_t],\, f \big\rangle\, dW_t.$$

This defines the diffusion adjoint through its action against the common shock, $\mathcal{B}_t^*[g]\, dW_t \equiv -\partial_a[\sigma_a\, g]\, dW_t$, which matches $\mathcal{B}_t^*$ in equation (28) after substituting $\sigma_a = \theta^*\, a\, \sigma_Q/Q$ from equation (24).[39]

Conclusion. Substituting the two adjoints back into equation (36), every operator now acts on g_t:

$$d\langle g_t, f\rangle = \big\langle \mathcal{A}_t^*[g_t],\, f \big\rangle\, dt + \big\langle \mathcal{B}_t^*[g_t]\, dW_t,\, f \big\rangle,$$

with $\mathcal{A}_t^*$ the full operator, demographic term included. Since the test function f is fixed (it is neither random nor time-dependent), we also have $d\langle g_t, f\rangle = \langle dg_t, f\rangle$.[40] Hence

$$\big\langle dg_t - \mathcal{A}_t^*[g_t]\, dt - \mathcal{B}_t^*[g_t]\, dW_t,\; f \big\rangle = 0,$$

which must hold for every admissible test function f. It must therefore be the case that

$$dg_t = \mathcal{A}_t^*[g_t]\, dt + \mathcal{B}_t^*[g_t]\, dW_t,$$

which is exactly the equilibrium KFE equation (28).[41]

The subscript t on $\mathcal{A}_t^*, \mathcal{B}_t^*$ simply records their dependence on the current aggregate state (Y_t, Q_t). □

B Additional Results

Figure 12 plots the impulse responses of dividends Y and the asset price Q separately, complementing the dividend-yield response in figure 5a. Dividends fall by the full size of the

[39] This is a first-order term in dW_t and must be kept distinct from the second-order dt Itô correction inside $\mathcal{A}_t^*$: both come from the same loading σ_a, yet one enters as σ_a^2 in the dt term and the other as σ_a at first power in the dW_t term.

[40] For the dt term this is a direct consequence of linearity of the integral; for the dW_t term, this formally requires to invoke the stochastic Fubini theorem, which lets the integral pass through the Itô integral against the common noise – see Øksendal (2014).

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

shock and recover slowly, governed by the exogenous mean reversion of the Y process. The asset price falls by roughly half as much on impact and recovers faster: equilibrium prices partially absorb the fundamental shock rather than amplifying it, as discussed in section 4.3.

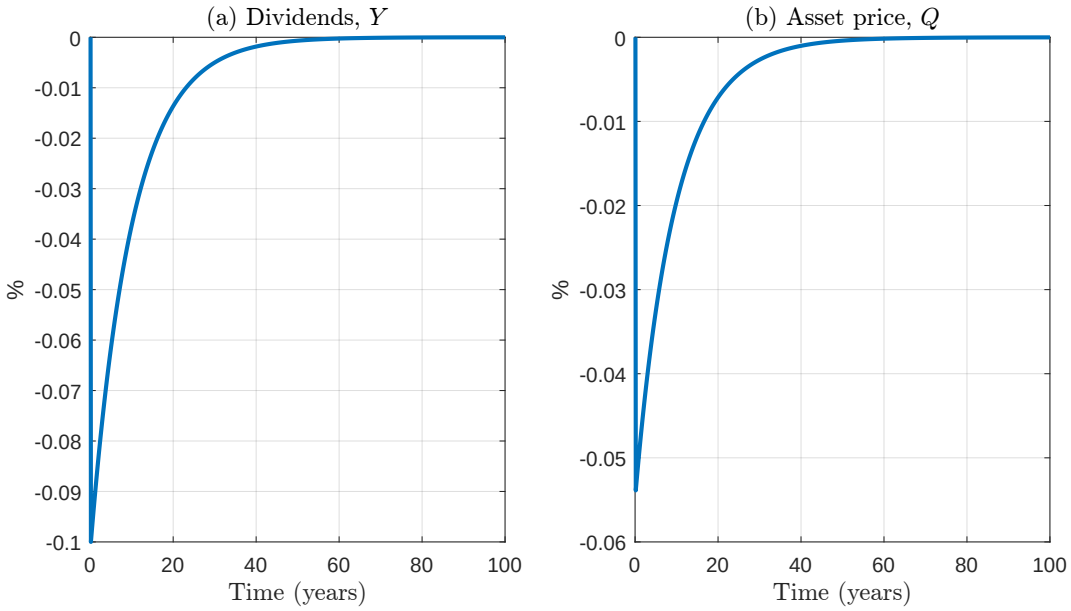


Figure 12: Impulse response of dividends and the asset price.

Notes: Response of dividends Y (left panel) and the asset price Q (right panel) to the negative aggregate dividend shock of section 4.3, in percent deviations from the stochastic steady state.

Figure 13 plots the responses of the asset price Q and of aggregate wealth A on the same axes. Since the average portfolio share equals Q_t/A_t, the vertical gap between the two paths reproduces the portfolio-share response of figure 5b at every horizon: the price declines by more than wealth on impact and recovers faster, and the two paths cross about six years after the shock – precisely where the share overshoot begins.

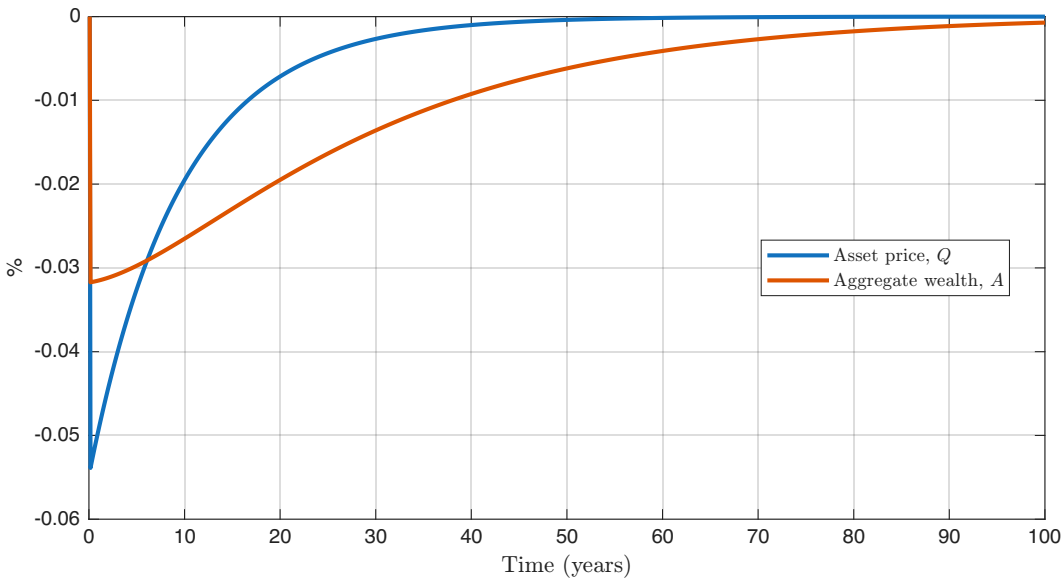


Figure 13: Impulse response of the asset price and aggregate wealth.

Notes: Response of the asset price Q and aggregate wealth A to the negative aggregate dividend shock of section 4.3, in percent deviations from the stochastic steady state; the vertical gap between the two paths equals the response of the average portfolio share $\theta^* = Q/A$.

Figure 14 plots the price-dividend ratio Q/Y along the ergodic path, against the dividend level Y (left panel) and the asset price Q (right panel). Because equilibrium prices are less volatile than dividends – the smoothing property discussed in section 4.2 – the ratio declines along the ergodic locus, from above 30 in low-dividend, low-price states to around 14 in high-dividend, high-price states.

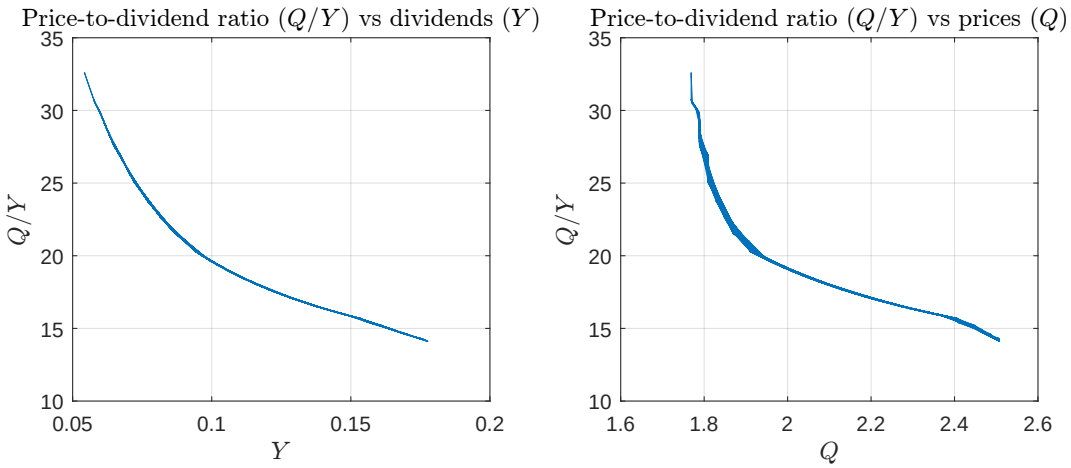
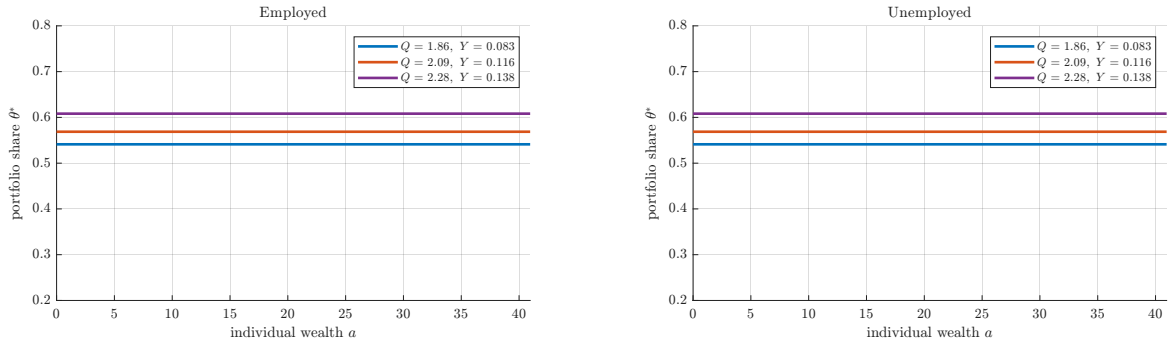


Figure 14: The price-dividend ratio along the ergodic path.

Notes: The left panel plots the price-dividend ratio Q/Y against the dividend level Y; the right panel plots it against the asset price Q. Both panels trace the thin, downward-sloping ergodic locus, from $Q/Y \approx 33$ at the low end ($Y \approx 0.055$, $Q \approx 1.8$) to $Q/Y \approx 14$ at the high end ($Y \approx 0.18$, $Q \approx 2.5$).

Figure 15 shows that optimal portfolio choice is invariant to individual wealth and income, the result established in section 4.4. Each panel plots $\theta^*(a, z; Y, Q)$ against individual wealth a for each employment state, at three representative points of the ergodic distribution – the 5th, 50th, and 95th percentiles of the realized average portfolio share along the simulated path.



(a) Employed.

(b) Unemployed.

Figure 15: Optimal portfolio share as a function of individual wealth at three representative ergodic states.

Notes: Each panel plots $\theta^*(a, z; Y, Q)$ as a function of individual wealth a for each employment state, evaluated at three aggregate states $(Q, Y) = (1.86, 0.083)$, $(2.09, 0.116)$, and $(2.28, 0.138)$ – the 5th, 50th, and 95th percentiles of the realized average portfolio share along the simulated path. Within each state the profile is exactly flat across wealth and the employed and unemployed panels coincide: along the simulated ergodic path, the cross-sectional range of θ^* is zero and the only variation is across aggregate states.

Figure 16 shows a scatter plot of all simulated (Y_t, Q_t) points against the full solution grid. As the scatter shows, for a given level of dividends variability in the equilibrium price is (almost) zero and the solution never crosses grid points.

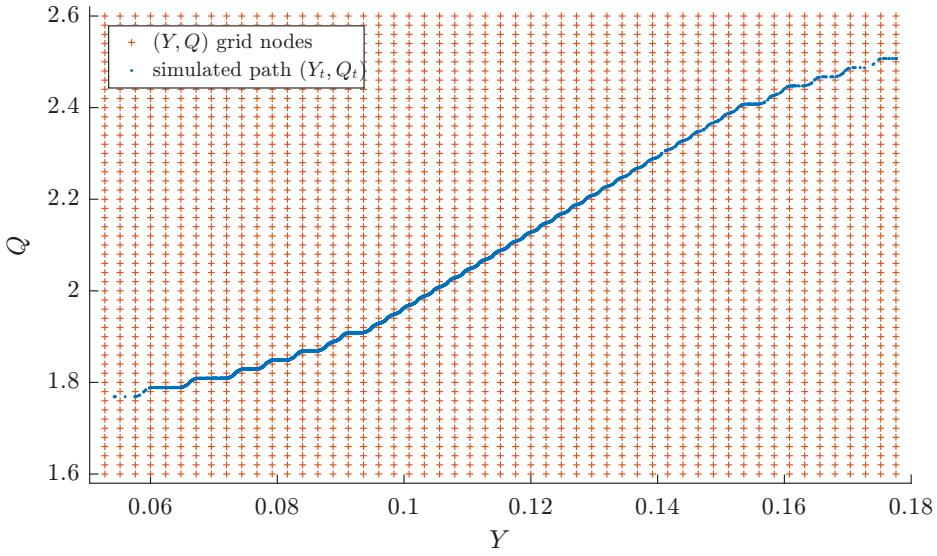


Figure 16: The (Q, Y) solution grid and the simulated ergodic path.

Notes: Markers denote the 51×53 solution-grid nodes, with Q uniform on $[1.6, 2.6]$ (51 nodes) and Y uniform on $[0.0528, 0.1776]$ (53 nodes); the dotted band traces the simulated (Y_t, Q_t) path over the 9,600 ergodic periods. None of the 9,600 simulated periods falls on a grid node (minimum distance 0.21 grid steps, median 0.43), so the market-clearing step described in section D always interpolates between the four surrounding nodes rather than evaluating at one.

C Neural Network Architecture and Training

The functional approximation problem we face requires capturing potentially complex non-linear relationships between the two-dimensional state vector $\Omega = (Y, Q)$ and the price dynamics characterized by drift μ_Q and diffusion σ_Q. Each neural network takes the form of a parametric function mapping from $\mathbb{R}^2$ to $\mathbb{R}$. For the drift network, we use a single hidden-layer structure:

$$\mu_Q(Y, Q; \alpha) = \alpha_0^{(2)} + \sum_{k=1}^{K} \alpha_k^{(2)} \varphi\left(\alpha_{k,0}^{(1)} + \alpha_{k,1}^{(1)} Y + \alpha_{k,2}^{(1)} Q\right) \tag{37}$$

where $\varphi(\cdot)$ denotes the activation function, K is the number of hidden units, and $\alpha = \{\alpha^{(1)}, \alpha^{(2)}\}$ collects all network parameters. The diffusion network $\sigma_Q(Y, Q; \beta)$ follows an analogous structure with its own parameter vector β.[42]

[42]Because $\sigma_Q(\cdot)$ enters in the denominator of the FOC for portfolio shares, we also impose $\sigma_Q \geq 10^{-4}$ to avoid running into numerical problems.

For the activation function, we employ the softplus function, $\varphi(x) = \log(1 + e^x)$, which provides a smooth approximation to the rectified linear unit and maintains differentiability everywhere. This choice proves particularly advantageous in our economic context where the underlying value and policy functions exhibit continuous derivatives. The fact that the softplus activation function is everywhere differentiable avoids the dead neuron problem that can plague ReLU-based architectures when gradients vanish and that is particularly severe when the number of neurons per layer is relatively small. Through experimentation, we find that $K = 8$ hidden units (for each network) provide sufficient expressiveness without overfitting – a balance that reflects the relatively low-dimensional nature of our state space compared to applications in computer vision or natural language processing where hundreds or thousands of units are common.

To minimize the training objectives defined in step (iii) of the algorithm in section 3, we employ gradient descent with exact line search, computing the full gradient over all training data at each iteration. Rather than using a fixed or adaptive learning rate, our algorithm determines the optimal step size through a two-phase procedure: first, a backtracking line search ensures feasibility by finding a step length that produces a valid decrease in the loss function; second, a bracketing and bisection method refines this step size to minimize the loss along the descent direction. This approach, while computationally more intensive per iteration than stochastic methods, offers crucial advantages for our economic application. The guarantee of monotonic loss decrease at each iteration ensures stable convergence of the perceived law of motion – a critical requirement when the neural-network approximation feeds back into the household optimization problem. Moreover, since our training data consists of simulated paths from a single economy rather than independent samples, the full-batch approach naturally aligns with the problem structure.

An important implementation detail concerns the construction of training data. Unlike supervised learning applications where data is given exogenously, our training samples $\{(Y_t, Q_t)\}_{t=1}^T$ emerge endogenously from the model simulation. This creates a potential feedback loop: the quality of our approximation determines the simulated paths, which in turn determine the training data. To ensure adequate coverage of the state space, we employ two strategies. First, in addition to the main simulation, we add a number of shorter simulations starting from multiple off-equilibrium locations drawn from a wide distribution over the (Y, Q) space. Second, in the first iterations of the full algorithm, we add small perturbations to the dividend process, effectively implementing a form of domain randomization that is later phased out with exponential decay.[43]

[43]We run the algorithm for as many outer iterations as needed – 101 in our baseline – and we declare the model converged when the maximum absolute difference between successive PLM evaluations in the visited region of the support is below 10^{-5}. The update step employs a standard relaxation algorithm, with only part of the computed difference incorporated into the next iteration. If the initial values and the support

C.1 Properties of Neural Networks

The choice of neural networks over alternative approximation methods deserves explicit justification. Traditional approaches such as polynomial approximations suffer from several limitations in our context. Chebyshev polynomials, while possessing excellent approximation properties on bounded intervals, exhibit problematic behavior outside their fitting region – a serious concern given that the ergodic distribution may shift substantially. Spline-based methods offer local flexibility but require pre-specifying knot locations, introducing an additional layer of complexity when the relevant regions of state space are not known *ex ante*.

Neural networks address these limitations through their capacity for adaptive, data-driven approximation. The universal approximation theorem guarantees that sufficiently large networks can approximate any continuous function to arbitrary precision, providing a theoretical foundation for our approach. More importantly for practical implementation, neural networks exhibit favorable extrapolation properties – with our choice of softplus activation function, they tend to produce smooth, well-behaved predictions even in regions of state space not visited during training. This "reasonable" behavior outside the training domain proves crucial for maintaining numerical stability during the iterative solution procedure.

We find that the neural-network approach is particularly well-suited to capturing how micro-level heterogeneity translates into aggregate price dynamics. The nonlinear activation functions naturally accommodate threshold effects (e.g., participation constraints), composition effects (e.g., wealth changes), and interaction effects between the wealth distribution and dividend shocks. These nonlinearities are not artifacts of the solution method but fundamental features of how heterogeneous portfolios aggregate into market prices.

The computational demands of training neural networks, while non-trivial, are manageable given modern hardware and software infrastructure. Each gradient update requires a forward pass through the network (to compute predictions) and a backward pass (to compute gradients), both of which parallelize efficiently.

Finally, we note that our use of separate networks for drift and diffusion, rather than a single multi-output network, reflects the distinct nature of these two functions. The drift μ_Q captures the expected price change and typically exhibits smooth variation across the state space. The volatility σ_Q, by contrast, must capture potentially sharp changes in uncertainty. By allowing each network to specialize, we achieve better overall approximation quality than would be possible with a single, compromise architecture. This modular approach

are set appropriately, the PLM parameters stabilize well before the declared tolerance, with the diffusion weights converging somewhat more slowly than the drift weights – consistent with the greater difficulty of learning the nonlinear diffusion surface.

also makes debugging and interpretation easier – we can examine each network's predictions independently to understand how agents' beliefs about different aspects of price dynamics evolve with the state of the economy.

D Numerical Solution

We solve the household's problem with an upwind finite-difference scheme in the style of Achdou et al. (2022). Two features of our environment make the implementation non-standard. First, under limited information the aggregate states (Y_t, Q_t) enter the household's problem directly through the perceived law of motion (20): the Hamilton-Jacobi-Bellman equation (25) is therefore a four-dimensional equation, carrying second-order terms in the aggregates and three cross-derivative terms. Second, the cross-sectional density evolves stochastically under the common shock while the price must clear the market at every instant: the density and the price therefore have to be found jointly at each simulated date. We proceed in four steps – the discretized Hamilton-Jacobi-Bellman equation, the discretized Kolmogorov forward equation, the market-clearing fixed point, and the simulation, stochastic steady state and impulse responses.

D.1 Discretized HJB

We approximate the value function on a finite grid in each of the four state variables. The wealth grid is $a \in \{a_1, \ldots, a_I\}$ with $I = 26$ nodes, geometrically spaced. The income grid is $z \in \{z_1, z_2\}$, with $J = 2$ states. The price grid is $Q \in \{Q_1, \ldots, Q_L\}$, with $L = 51$ nodes uniformly spaced on $[1.6, 2.6]$. The dividend grid is $Y \in \{Y_1, \ldots, Y_M\}$, with $M = 53$ nodes uniformly spaced on $[0.0528, 0.1776]$. The full grid therefore has $I \times J \times L \times M = 140{,}556$ nodes. We write $v_{i,j,l,m} \equiv v(a_i, z_j; Q_l, Y_m)$, and when only a subset of subscripts is reported the remaining ones take their default value, so that $v_{i+1,l+1} \equiv v_{i+1,j,l+1,m}$. We choose the bounds of the Q and Y grids to contain the ergodic region of the equilibrium.

We first define the forward and backward wealth spacings:

$$\Delta^F a_i = a_{i+1} - a_i, \qquad \Delta^B a_i = a_i - a_{i-1}$$

with phantom spacings at the two ends copied from their neighbours, $\Delta^F a_I = \Delta^F a_{I-1}$ and $\Delta^B a_1 = \Delta^B a_2$. We approximate the forward, backward, second, and cross-partial derivatives

of the value function as:

$$\partial_a^F v \equiv \frac{v_{i+1} - v_i}{\Delta^F a_i}, \qquad \partial_a^B v \equiv \frac{v_i - v_{i-1}}{\Delta^B a_i}$$
$$\partial_{aa} v \equiv \frac{v_{i+1}\Delta^B a_i - v_i(\Delta^B a_i + \Delta^F a_i) + v_{i-1}\Delta^F a_i}{\frac{1}{2}\Delta^B a_i \Delta^F a_i (\Delta^B a_i + \Delta^F a_i)}$$
$$\partial_{aQ} v \equiv \frac{v_{i-1,l-1} - v_{i-1,l+1} - v_{i+1,l-1} + v_{i+1,l+1}}{(\Delta^B a_i + \Delta^F a_i)(\Delta^B Q_l + \Delta^F Q_l)}$$

and similarly for the other pairs of states, $\partial_{QQ} v$, $\partial_{YY} v$, $\partial_{aY} v$ and $\partial_{QY} v$. It is convenient to define the weights that the second-difference approximation places on the forward and the backward neighbour:

$$w_a^B = \frac{\Delta^B a_i}{\frac{1}{2}\Delta^B a_i \Delta^F a_i (\Delta^B a_i + \Delta^F a_i)}, \qquad w_a^F = \frac{\Delta^F a_i}{\frac{1}{2}\Delta^B a_i \Delta^F a_i (\Delta^B a_i + \Delta^F a_i)}$$

these are the weights that $\partial_{aa} v$ places on the forward and the backward neighbour respectively; note the crossing – w_a^B multiplies v_{i+1} and w_a^F multiplies v_{i-1} – which is exactly what the coefficient definitions below encode. Because Q and Y are uniform grids, $w_Q^B = w_Q^F = 1/(\Delta Q)^2$ and $w_Y^B = w_Y^F = 1/(\Delta Y)^2$. The cross-derivative stencils for the remaining pairs of states are the centred four-corner differences shown above for $\partial_{aQ} v$.

In an upwind scheme the direction of the difference follows the sign of the drift. The candidate wealth drifts are:

$$\mu_a^F = z_j + r^* a_i + \big(R(Q_l, Y_m) - r^*\big)\theta_{i,j,l,m}\, a_i - c^F$$
$$\mu_a^B = z_j + r^* a_i + \big(R(Q_l, Y_m) - r^*\big)\theta_{i,j,l,m}\, a_i - c^B$$

with $c^F = (\partial_a^F v)^{-1/\gamma}$ and $c^B = (\partial_a^B v)^{-1/\gamma}$ the consumption levels implied by the forward and backward marginal value of wealth, and $R(Y, Q) = (Y + \mu_Q(Y, Q))/Q$ the expected return on the tree as in section 2.[44] We define the convenience functions $(x)^+ \equiv x\,\mathbf{1}\{x > 0\}$ and $(x)^- \equiv x\,\mathbf{1}\{x < 0\}$. Consumption c^n takes the forward value c^F where the forward drift is positive; the backward value c^B where the backward drift is negative and the forward drift is not positive; and otherwise the zero-drift consumption level $\bar{c}^n$ that sets the drift of wealth to zero:

$$\bar{c}^n = z_j + r^* a_i + \big(R(Q_l, Y_m) - r^*\big)\theta_{i,j,l,m}\, a_i$$

We upwind the drifts of Q and Y in exactly the same way, using $(\mu_Q^F)^+$, $(\mu_Q^B)^-$, $(\mu_Y^F)^+$ and $(\mu_Y^B)^-$, with μ_Q read off the perceived law of motion (20) and μ_Y as in equation (4).

[44] The displayed return is the world rate r^*, following the convention of section 2; because annuity markets are perfect, the return on wealth actually used in the solution is annuity-inclusive, so the drift the solver iterates on carries $(r^* + \phi)a$ rather than $r^* a$, and the effective discount rate is $\rho + \phi$.

Let superscript n denote the iteration counter and Δ the implicit step; the upwind approximation of the Hamilton-Jacobi-Bellman equation (25) is:

$$\begin{aligned}\frac{v^{n+1}-v^n}{\Delta}+(\rho+\phi)v^{n+1} = u(c^n)&+\partial_a^F v^{n+1}(\mu_a^F)^+ + \partial_a^B v^{n+1}(\mu_a^B)^- + \frac{1}{2}\partial_{aa}v^{n+1}\sigma_a^2\\ &+\lambda_j\big(v_{-j}^{n+1}-v^{n+1}\big)+\partial_Q^F v^{n+1}(\mu_Q^F)^+ + \partial_Q^B v^{n+1}(\mu_Q^B)^-\\ &+\frac{1}{2}\partial_{QQ}v^{n+1}\sigma_Q^2+\partial_Y^F v^{n+1}(\mu_Y^F)^+ + \partial_Y^B v^{n+1}(\mu_Y^B)^- + \frac{1}{2}\partial_{YY}v^{n+1}\sigma_Y^2\\ &+\partial_{aQ}v^{n+1}\sigma_a\sigma_Q+\partial_{aY}v^{n+1}\sigma_a\sigma_Y+\partial_{QY}v^{n+1}\sigma_Q\sigma_Y\end{aligned}$$

Rearranging and collecting terms node by node gives:

$$\begin{aligned}\frac{v^{n+1}-v^n}{\Delta}+(\rho+\phi)v^{n+1} = u(c^n)&+\alpha v_{i-1}^{n+1}+\beta v^{n+1}+\xi v_{i+1}^{n+1}+\lambda_j v_{-j}^{n+1}\\ &+\phi v_{l-1}^{n+1}+\theta v_{l+1}^{n+1}+\rho v_{m-1}^{n+1}+\chi v_{m+1}^{n+1}\\ &+\nu\big(v_{i-1,l-1}^{n+1}-v_{i-1,l+1}^{n+1}-v_{i+1,l-1}^{n+1}+v_{i+1,l+1}^{n+1}\big)\\ &+\eta\big(v_{i-1,m-1}^{n+1}-v_{i-1,m+1}^{n+1}-v_{i+1,m-1}^{n+1}+v_{i+1,m+1}^{n+1}\big)\\ &+\tau\big(v_{l-1,m-1}^{n+1}-v_{l-1,m+1}^{n+1}-v_{l+1,m-1}^{n+1}+v_{l+1,m+1}^{n+1}\big)\end{aligned}$$

with coefficients:[45]

$$\begin{aligned}\alpha &= -\frac{(\mu_a^B)^-}{\Delta^B a_i}+\frac{1}{2}w_a^F\sigma_a^2, & \xi &= \frac{(\mu_a^F)^+}{\Delta^F a_i}+\frac{1}{2}w_a^B\sigma_a^2\\ \phi &= -\frac{(\mu_Q^B)^-}{\Delta^B Q_l}+\frac{1}{2}w_Q^F\sigma_Q^2, & \theta &= \frac{(\mu_Q^F)^+}{\Delta^F Q_l}+\frac{1}{2}w_Q^B\sigma_Q^2\\ \rho &= -\frac{(\mu_Y^B)^-}{\Delta^B Y_m}+\frac{1}{2}w_Y^F\sigma_Y^2, & \chi &= \frac{(\mu_Y^F)^+}{\Delta^F Y_m}+\frac{1}{2}w_Y^B\sigma_Y^2\\ \nu &= \frac{\sigma_a\sigma_Q}{(\Delta^B a_i+\Delta^F a_i)(\Delta^B Q_l+\Delta^F Q_l)}, & \eta &= \frac{\sigma_a\sigma_Y}{(\Delta^B a_i+\Delta^F a_i)(\Delta^B Y_m+\Delta^F Y_m)}\\ \tau &= \frac{\sigma_Q\sigma_Y}{(\Delta^B Q_l+\Delta^F Q_l)(\Delta^B Y_m+\Delta^F Y_m)}, & \beta &= -\alpha-\xi-\lambda_j-\phi-\theta-\rho-\chi\end{aligned}$$

We build the diagonal $\beta_{i,j,l,m}$ as minus the sum of the off-diagonal coefficients, so each row of the generator sums to zero by construction.

For boundary conditions, at $a_1 = 0$ we set $\partial_a^B v_{1,j,l,m} = u'(z_j)$, which makes $c_1^B = z_j$ and hence $\mu_a^B(a_1) = 0$, so no mass leaves the grid below the borrowing constraint. At a_I we set $\partial_a^F v_{I,j,l,m} = u'(\bar{c}_I)$ with $\bar{c}_I$ the zero-drift consumption level, so that $\mu_a^F(a_I) = 0$, and we set $\xi_{I,j,l,m} = 0$ so that the diffusion cannot push wealth above the top node. Along Q and Y we

[45] Some of the matrix coefficients used here re-use symbols used in the main text. For example $\theta_{i,j,l,m}$ need not be confused with the portfolio share.

impose reflecting barriers, implemented by folding the block that would point outside the grid back onto the diagonal block – a zero-flux condition that preserves the zero row sums exactly. For the cross-derivative stencils, the ghost nodes are folded onto the diagonal.

We stack the value function with a varying first, then z, then Q, then Y: the scheme becomes a sparse linear system, whose matrix we build from the bottom up. The $(I\cdot 2)\times(I\cdot 2)$ block at a fixed aggregate node (Q_l, Y_m) stacks two a-tridiagonals, one per income state, with the income-switch intensities λ_1 and λ_2 linking the two blocks:

$$\mathbf{A}^n_{l,m} = \begin{bmatrix} \beta^n_{1,1,l,m} & \xi^n_{1,1,l,m} & 0 & \cdots & 0 & \lambda_1 & 0 & \cdots & 0 \\ \alpha^n_{2,1,l,m} & \beta^n_{2,1,l,m} & \xi^n_{2,1,l,m} & \cdots & 0 & 0 & \lambda_1 & \ddots & 0 \\ 0 & \alpha^n_{3,1,l,m} & \beta^n_{3,1,l,m} & \xi^n_{3,1,l,m} & \cdots & 0 & 0 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & \alpha^n_{I,1,l,m} & \beta^n_{I,1,l,m} & 0 & 0 & \cdots & \lambda_1 \\ \lambda_2 & 0 & \cdots & 0 & 0 & \beta^n_{1,2,l,m} & \xi^n_{1,2,l,m} & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & \lambda_2 & 0 & \cdots & \alpha^n_{I,2,l,m} & \beta^n_{I,2,l,m} \end{bmatrix}$$

The tridiagonal entries $\alpha_{i,j,l,m}$, $\beta_{i,j,l,m}$, $\xi_{i,j,l,m}$ are the wealth-dimension transition rates defined above, and the off-diagonal λ_1, λ_2 entries switch a household between income states without moving it in wealth.

We stack the right-hand side and the value function itself in the same order:

$$\mathbf{u}^n = \begin{bmatrix} u(c^n_{1,1,1,1}) \\ u(c^n_{2,1,1,1}) \\ \vdots \\ \vdots \\ u(c^n_{I,2,L,M}) \end{bmatrix}, \quad \mathbf{v}^{n+1}_{l,m} = \begin{bmatrix} v^{n+1}_{1,1,l,m} \\ v^{n+1}_{2,1,l,m} \\ \vdots \\ v^{n+1}_{I,1,l,m} \\ v^{n+1}_{1,2,l,m} \\ \vdots \\ v^{n+1}_{I,2,l,m} \end{bmatrix}, \quad \mathbf{v}^{n+1}_m = \begin{bmatrix} \mathbf{v}^{n+1}_{1,m} \\ \mathbf{v}^{n+1}_{2,m} \\ \vdots \\ \mathbf{v}^{n+1}_{L,m} \end{bmatrix}, \quad \mathbf{v}^{n+1} = \begin{bmatrix} \mathbf{v}^{n+1}_1 \\ \mathbf{v}^{n+1}_2 \\ \vdots \\ \mathbf{v}^{n+1}_M \end{bmatrix}$$

The vector $\mathbf{u}^n$ stacks the flow utility $u(c^n_{i,j,l,m})$ over the same (i, j, l, m) ordering as the value function.

Coupling across price nodes at a fixed dividend node Y_m produces a block-tridiagonal

matrix:

$$\mathbf{A}_m^n = \begin{bmatrix} \left(\mathbf{A}_{1,m}^n + \mathbf{F}_{1,m}^n\right) & \mathbf{Z}_{1,m}^n & \mathbf{0}_{I\cdot 2} & \cdots & \mathbf{0}_{I\cdot 2} & \mathbf{0}_{I\cdot 2} \\ \mathbf{F}_{2,m}^n & \mathbf{A}_{2,m}^n & \mathbf{Z}_{2,m}^n & \cdots & \mathbf{0}_{I\cdot 2} & \mathbf{0}_{I\cdot 2} \\ \mathbf{0}_{I\cdot 2} & \mathbf{F}_{3,m}^n & \mathbf{A}_{3,m}^n & \cdots & \mathbf{0}_{I\cdot 2} & \mathbf{0}_{I\cdot 2} \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \mathbf{0}_{I\cdot 2} & \mathbf{0}_{I\cdot 2} & \cdots & \mathbf{F}_{L-1,m}^n & \mathbf{A}_{L-1,m}^n & \mathbf{Z}_{L-1,m}^n \\ \mathbf{0}_{I\cdot 2} & \mathbf{0}_{I\cdot 2} & \cdots & \mathbf{0}_{I\cdot 2} & \mathbf{F}_{L,m}^n & \left(\mathbf{A}_{L,m}^n + \mathbf{Z}_{L,m}^n\right) \end{bmatrix}$$

The diagonal blocks $\mathbf{A}_{l,m}^n$ are the wealth-income generators defined above; $\mathbf{F}_{l,m}^n$ and $\mathbf{Z}_{l,m}^n$ carry the coupling to the price nodes immediately below and above l.

The $\mathbf{F}$ and $\mathbf{Z}$ blocks are themselves block-diagonal in the income state:

$$\mathbf{F}_{l,m}^n = \begin{bmatrix} \mathbf{F}_{1,l,m}^n & \mathbf{0}_I \\ \mathbf{0}_I & \mathbf{F}_{2,l,m}^n \end{bmatrix}, \qquad \mathbf{Z}_{l,m}^n = \begin{bmatrix} \mathbf{Z}_{1,l,m}^n & \mathbf{0}_I \\ \mathbf{0}_I & \mathbf{Z}_{2,l,m}^n \end{bmatrix}$$

with the inner $I \times I$ blocks:

$$\mathbf{F}_{j,l,m}^n = \begin{bmatrix} (\phi_{1,j,l,m} + \nu_{1,j,l,m}) & -\nu_{1,j,l,m} & 0 & 0 & \cdots & 0 \\ \nu_{2,j,l,m} & \phi_{2,j,l,m} & -\nu_{2,j,l,m} & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & \nu_{I-1,j,l,m} & \phi_{I-1,j,l,m} & -\nu_{I-1,j,l,m} \\ 0 & 0 & 0 & \cdots & \nu_{I,j,l,m} & (\phi_{I,j,l,m} - \nu_{I,j,l,m}) \end{bmatrix}$$

$$\mathbf{Z}_{j,l,m}^n = \begin{bmatrix} (\theta_{1,j,l,m} - \nu_{1,j,l,m}) & \nu_{1,j,l,m} & 0 & 0 & \cdots & 0 \\ -\nu_{2,j,l,m} & \theta_{2,j,l,m} & \nu_{2,j,l,m} & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & -\nu_{I-1,j,l,m} & \theta_{I-1,j,l,m} & \nu_{I-1,j,l,m} \\ 0 & 0 & 0 & \cdots & -\nu_{I,j,l,m} & (\theta_{I,j,l,m} + \nu_{I,j,l,m}) \end{bmatrix}$$

where $\phi_{i,j,l,m}$ and $\theta_{i,j,l,m}$ are the Q-drift and Q-diffusion coefficients defined above, evaluated at each wealth node; $\nu_{i,j,l,m}$ is the a–Q cross-derivative coefficient, entering with opposite sign on the two off-diagonals of each block.

Finally, adding the dividend state produces the full matrix:

$$\mathbf{A}^n = \begin{bmatrix} (\mathbf{A}_1^n + \mathbf{P}_1^n) & \mathbf{X}_1^n & \mathbf{0}_{I\cdot 2\cdot L} & \cdots & \mathbf{0}_{I\cdot 2\cdot L} & \mathbf{0}_{I\cdot 2\cdot L} \\ \mathbf{P}_2^n & \mathbf{A}_2^n & \mathbf{X}_2^n & \cdots & \mathbf{0}_{I\cdot 2\cdot L} & \mathbf{0}_{I\cdot 2\cdot L} \\ \mathbf{0}_{I\cdot 2\cdot L} & \mathbf{P}_3^n & \mathbf{A}_3^n & \cdots & \mathbf{0}_{I\cdot 2\cdot L} & \mathbf{0}_{I\cdot 2\cdot L} \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \mathbf{0}_{I\cdot 2\cdot L} & \mathbf{0}_{I\cdot 2\cdot L} & \cdots & \mathbf{P}_{M-1}^n & \mathbf{A}_{M-1}^n & \mathbf{X}_{M-1}^n \\ \mathbf{0}_{I\cdot 2\cdot L} & \mathbf{0}_{I\cdot 2\cdot L} & \cdots & \mathbf{0}_{I\cdot 2\cdot L} & \mathbf{P}_M^n & (\mathbf{A}_M^n + \mathbf{X}_M^n) \end{bmatrix}$$

The rest follows a similar structure as above:

$$\mathbf{P}_m^n = \begin{bmatrix} \left(\mathbf{P}_{1,m}^n + \tau_{1,m}\cdot\mathbf{I}_{I\cdot 2}\right) & -\tau_{1,m}\cdot\mathbf{I}_{I\cdot 2} & \mathbf{0}_{I\cdot 2} & \cdots & \mathbf{0}_{I\cdot 2} & \mathbf{0}_{I\cdot 2} \\ \tau_{2,m}\cdot\mathbf{I}_{I\cdot 2} & \mathbf{P}_{2,m}^n & -\tau_{2,m}\cdot\mathbf{I}_{I\cdot 2} & \cdots & \mathbf{0}_{I\cdot 2} & \mathbf{0}_{I\cdot 2} \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \mathbf{0}_{I\cdot 2} & \mathbf{0}_{I\cdot 2} & \cdots & \tau_{L-1,m}\cdot\mathbf{I}_{I\cdot 2} & \mathbf{P}_{L-1,m}^n & -\tau_{L-1,m}\cdot\mathbf{I}_{I\cdot 2} \\ \mathbf{0}_{I\cdot 2} & \mathbf{0}_{I\cdot 2} & \cdots & \mathbf{0}_{I\cdot 2} & \tau_{L,m}\cdot\mathbf{I}_{I\cdot 2} & \left(\mathbf{P}_{L,m}^n - \tau_{L,m}\cdot\mathbf{I}_{I\cdot 2}\right) \end{bmatrix}$$

$$\mathbf{X}_m^n = \begin{bmatrix} \left(\mathbf{X}_{1,m}^n - \tau_{1,m}\cdot\mathbf{I}_{I\cdot 2}\right) & \tau_{1,m}\cdot\mathbf{I}_{I\cdot 2} & \mathbf{0}_{I\cdot 2} & \cdots & \mathbf{0}_{I\cdot 2} & \mathbf{0}_{I\cdot 2} \\ -\tau_{2,m}\cdot\mathbf{I}_{I\cdot 2} & \mathbf{X}_{2,m}^n & \tau_{2,m}\cdot\mathbf{I}_{I\cdot 2} & \cdots & \mathbf{0}_{I\cdot 2} & \mathbf{0}_{I\cdot 2} \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \mathbf{0}_{I\cdot 2} & \mathbf{0}_{I\cdot 2} & \cdots & -\tau_{L-1,m}\cdot\mathbf{I}_{I\cdot 2} & \mathbf{X}_{L-1,m}^n & \tau_{L-1,m}\cdot\mathbf{I}_{I\cdot 2} \\ \mathbf{0}_{I\cdot 2} & \mathbf{0}_{I\cdot 2} & \cdots & \mathbf{0}_{I\cdot 2} & -\tau_{L,m}\cdot\mathbf{I}_{I\cdot 2} & \left(\mathbf{X}_{L,m}^n + \tau_{L,m}\cdot\mathbf{I}_{I\cdot 2}\right) \end{bmatrix}$$

$$\mathbf{P}_{l,m}^n = \begin{bmatrix} \mathbf{P}_{1,l,m}^n & \mathbf{0}_I \\ \mathbf{0}_I & \mathbf{P}_{2,l,m}^n \end{bmatrix}, \qquad \mathbf{X}_{l,m}^n = \begin{bmatrix} \mathbf{X}_{1,l,m}^n & \mathbf{0}_I \\ \mathbf{0}_I & \mathbf{X}_{2,l,m}^n \end{bmatrix}$$

$$\mathbf{P}_{j,l,m}^n = \begin{bmatrix} (\rho_{1,j,l,m} + \eta_{1,j,l,m}) & -\eta_{1,j,l,m} & 0 & 0 & \cdots & 0 \\ \eta_{2,j,l,m} & \rho_{2,j,l,m} & -\eta_{2,j,l,m} & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & \eta_{I-1,j,l,m} & \rho_{I-1,j,l,m} & -\eta_{I-1,j,l,m} \\ 0 & 0 & 0 & \cdots & \eta_{I,j,l,m} & (\rho_{I,j,l,m} - \eta_{I,j,l,m}) \end{bmatrix}$$

$$\mathbf{X}_{j,l,m}^n = \begin{bmatrix} (\chi_{1,j,l,m} - \eta_{1,j,l,m}) & \eta_{1,j,l,m} & 0 & 0 & \cdots & 0 \\ -\eta_{2,j,l,m} & \chi_{2,j,l,m} & \eta_{2,j,l,m} & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & -\eta_{I-1,j,l,m} & \chi_{I-1,j,l,m} & \eta_{I-1,j,l,m} \\ 0 & 0 & 0 & \cdots & -\eta_{I,j,l,m} & (\chi_{I,j,l,m} + \eta_{I,j,l,m}) \end{bmatrix}$$

where $\rho_{i,j,l,m}$ and $\chi_{i,j,l,m}$ are the Y-drift and Y-diffusion coefficients defined above and $\eta_{i,j,l,m}$ is the a–Y cross-derivative coefficient, entering with opposite sign on the two off-diagonals

of each block, in parallel with $\mathbf{F}^n_{j,l,m}$ and $\mathbf{Z}^n_{j,l,m}$ above. The $\tau_{l,m}$ blocks in $\mathbf{P}^n_m$ and $\mathbf{X}^n_m$ are scalar multiples of the identity matrix because the Q–Y cross-derivative coefficient does not depend on the individual states a and z. The resulting matrix $\mathbf{A}^n$ is a 140,556 $\times$ 140,556 sparse matrix.

The full system, incorporating all five matrices above, is:

$$\frac{1}{\Delta}\left(\mathbf{v}^{n+1}-\mathbf{v}^n\right)+(\rho+\phi)\,\mathbf{v}^{n+1}=\mathbf{u}^n+\mathbf{A}^n\mathbf{v}^{n+1}$$

which rearranges to

$$\mathbf{B}^n\mathbf{v}^{n+1}=\mathbf{d}^n,\qquad \mathbf{B}^n=\left(\frac{1}{\Delta}+\rho+\phi\right)\mathbf{I}-\mathbf{A}^n,\qquad \mathbf{d}^n=\mathbf{u}^n+\frac{1}{\Delta}\mathbf{v}^n \tag{38}$$

We solve the system jointly over all four dimensions by a sparse direct factorization at every iteration. The implicit step is $\Delta = 1000$ – large enough that the update behaves like a Newton step, which is why so few iterations are needed. The algorithm to solve the discretized HJB runs as follows.

1. Given $\mathbf{v}^n$, compute c^n from the upwinded marginal value of wealth.
2. Build $\mathbf{A}^n$ and solve equation (38) for $\mathbf{v}^{n+1}$.
3. Update the value function and stop if $\|\mathbf{v}^{n+1}-\mathbf{v}^n\|_\infty$ is below tolerance; otherwise set $n \leftarrow n+1$ and return to step 1.

The value update in step 3 is relaxed,

$$\mathbf{v}^{n+1} \leftarrow (1-\omega)\,\mathbf{v}^n+\omega\,\tilde{\mathbf{v}}^{n+1}$$

with ω adapted to the behaviour of the residual: whenever the sup-norm residual improves, ω grows; whenever it worsens, the step is rejected, the previous iterate restored, and ω halved down. With a warm start from the previous outer iteration the residual improves monotonically and the HJB converges in five to nine iterations.

D.1.1 Portfolio Choice

The portfolio share solves the first-order condition of equation (25), evaluated on the grid with the same finite-difference stencils used for the value function:

$$\theta^\star=-\frac{(R-r^*)\,\partial_a v+\frac{\sigma_Q}{Q}\left(\sigma_Q\,\partial_{aQ}v+\sigma_Y\,\partial_{aY}v\right)}{a\left(\frac{\sigma_Q}{Q}\right)^2\partial_{aa}v}$$

where $\partial_a v$, $\partial_{aa} v$, $\partial_{aQ} v$ and $\partial_{aY} v$ are the finite-difference derivatives defined above and $R = R(Q_l, Y_m)$ is the expected return on the tree. The numerator's three additive pieces correspond, in order, to the risk-return, price-hedging and dividend-hedging terms of the decomposition in section 2. We evaluate the condition twice – once with $\partial_a^F v$ and once with $\partial_a^B v$ – and select between the two candidates by the same upwind indicator used for consumption, taking their average where the drift of wealth is zero.

We impose the short-sale and no-leverage constraints by clamping the candidate to $[0, 1]$: households at a corner solution hold either no equity or their entire portfolio in equity. The portfolio share is not solved jointly with the value function. The HJB solve of equation (38) takes θ as given; we evaluate the first-order condition above once per outer iteration, at the value function that solve returns, and then update the portfolio share as a damped Gauss–Seidel step,

$$\theta^{(n+1)} = (1 - \omega_\theta)\, \theta^{(n)} + \omega_\theta\, \theta^\star\big(v^{(n)}\big)$$

with the relaxation weight ω_θ converging toward one as the outer loop proceeds. We switch the update on only after the perceived law of motion has begun to settle: the value function and the portfolio share therefore converge together as a damped fixed point, rather than being solved simultaneously at each step.

D.2 Discretized KFE

The density evolves according to equation (28), whose drift and diffusion operators $\mathcal{A}_t^*$ and $\mathcal{B}_t^*$ we now discretize; the derivation of the equation itself is in section A. The discretized drift operator is not the transpose of the full four-dimensional Hamilton-Jacobi-Bellman matrix $\mathbf{A}^n$. At each aggregate node (Q_l, Y_m) we build a separate $(I \cdot 2) \times (I \cdot 2)$ generator $\mathbf{A}_{l,m}$ in the individual states only – the same upwind drift and $\frac{1}{2}\sigma_a^2$ diffusion coefficients $\alpha_{i,j,l,m}$ and $\xi_{i,j,l,m}$ as above, plus the income-switch intensities, with the diagonal rebuilt as $-\alpha_{i,j,l,m} - \xi_{i,j,l,m} - \lambda_j$. We move the aggregate states outside the operator: Y is advanced by an explicit step, Q is pinned by market clearing, and the dependence of the generator on the aggregate state is recovered by re-selecting and interpolating the four surrounding nodes at every date. The rows of the matrix $\mathbf{A}_{l,m}$ all sum to zero, and $\mathbf{A}_{l,m}$ is therefore a proper generator in (a, z).

The realized aggregate shock enters the density only through the first-order operator $\mathcal{B}^*$, which we discretize as $\mathbf{B}_{l,m}$: the finite-difference approximation of $-\partial_a[\sigma_a g]$, with $\sigma_a = \theta a \sigma_Q / Q$ as in equation (24). Because $\theta \in [0, 1]$, $a \geq 0$ and $\sigma_Q > 0$, the coefficient σ_a is non-negative everywhere, so the operator is a forward difference at every node.

Because the wealth grid is non-uniform, the object we transport is the vector of node masses, not of density values. We define $\tilde{g} = \mathbf{D} g$, where $\mathbf{D}$ is diagonal with entries the

trapezoidal weights:

$$\tilde{\Delta}a_i = \begin{cases} \Delta^F a_1/2 & i = 1 \\ (\Delta^B a_i + \Delta^F a_i)/2 & 1 < i < I \\ \Delta^B a_I/2 & i = I \end{cases}$$

$\sum_i g_i \tilde{\Delta} a_i$ is the trapezoidal approximation of $\int g\, da$, so working in masses is what makes $\mathbf{A}^\top$ the correct discrete adjoint of the generator on a non-uniform grid; moments are computed against the same weights, and the density is recovered as $g = \mathbf{D}^{-1}\tilde{g}$. This generalizes the uniform-grid construction in the online appendix of Achdou et al. (2022).

With the operators in hand, we must choose how implicitly to treat each term. Fully implicit schemes are inappropriate for the strong approximation of stochastic differential equations: they involve reciprocals of Gaussian random variables – which do not have finite absolute moments – so a strong convergence analysis would not make sense; we instead restrict attention to semi-implicit schemes that are implicit in the drift only, following Kloeden and Platen (1995). Writing $a(g) = \mathcal{A}^* g$ and $b(g) = \mathcal{B}^* g$, the family of semi-implicit schemes is:

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

where $\alpha \in [0,1]$ is the implicitness parameter: $\alpha = 0$ is the drift-explicit Euler scheme, $\alpha = 1$ is the drift-implicit scheme, and $\alpha = 1/2$ generalizes the trapezoidal (Crank–Nicolson) method. We implement the fully drift-implicit case $\alpha = 1$, which is unconditionally stable and first-order accurate.

The implemented step is therefore:

$$\Big[(1+\phi\,\Delta t)\,\mathbf{I} - \Delta t\, \mathbf{A}_{l,m}^\top\Big]\, \tilde{g}_{t+\Delta t}^{(l,m)} = \Big[\mathbf{I} + \sqrt{\Delta t}\,\varepsilon_t\, \mathbf{B}_{l,m}^\top\Big]\, \tilde{g}_t \;+\; \phi\,\Delta t\,\mathbf{e}$$

$\varepsilon_t \sim \mathcal{N}(0,1)$ so that $\Delta W_t = \sqrt{\Delta t}\,\varepsilon_t$; $\mathbf{e}$ is the discretized birth distribution ψ, which places mass $\lambda_2/(\lambda_1+\lambda_2)$ at (a_1, z_1) and $\lambda_1/(\lambda_1+\lambda_2)$ at (a_1, z_2) – that is, newborns enter with zero wealth and an income state drawn from the ergodic distribution – so the total injected mass is exactly $\phi\Delta t$. We renormalize the resulting density to unit mass at every step.

D.2.1 Market Clearing

We advance the dividend by an Euler–Maruyama step on the level of the process:

$$Y_{t+\Delta t} = Y_t + Y_t\,\Delta t\Big[\eta \log\frac{\bar{Y}}{Y_t} + \frac{1}{2}\sigma^2\Big] + Y_t\,\sigma\sqrt{\Delta t}\,\varepsilon_t.$$

Given Y_t and the previous period's distribution, the price and the density must be found jointly, because the price enters the market-clearing condition (31) twice: first through the

portfolio policy, which is evaluated at (Q_t, Y_t); second through the density itself, since the generator used to advance $\tilde{g}$ into period t is the one indexed by (Q_t, Y_t). The price is the fixed point:

$$Q_t = \mathcal{F}(Q_t) \equiv \sum_{i,j} \theta^*\big(a_i, z_j; Q_t, Y_t\big)\, a_i\, \tilde{g}_t\big(a_i, z_j; Q_t\big)$$

For a candidate price, we solve the KFE step described above at each of the four grid nodes surrounding (Q_t, Y_t) – four sparse solves – and combine the four resulting densities by interpolation; the portfolio policy is also interpolated at the same point. We then update the candidate with dampening: the step keeps its size while the candidate moves toward the fixed point; whenever the sign of $\mathcal{F}(Q) - Q$ flips, the step is reversed and halved. The search stops when $|\mathcal{F}(Q) - Q|$ falls below 10^{-8}, or when the step size collapses below 10^{-10}.

D.2.2 Simulation, Stochastic Steady State, and Impulse Responses

Each outer iteration simulates a single path of the economy at monthly frequency ($\Delta t = 1/12$) for one thousand years, starting from the deterministic steady state; we discard the first two hundred years as burn-in, and the remaining eight hundred years (9,600 periods) are the sample we use to train the networks and to report moments, as described in section 3. We draw the Brownian increments once, before the outer loop begins, and reuse the same shock path at every iteration – common random numbers, which make the outer loop a deterministic map and prevent simulation noise from being mistaken for failure to converge.

The stochastic steady state is the point at which the economy settles when the aggregate shock is never realized but households continue to price risk: the value function, the portfolio policies and the density's own diffusion are all computed under the perceived law of motion with $\sigma_Q > 0$; only the realized Brownian increment is set to zero. We compute it by simulating the full system – the density, the market-clearing price and the dividend – with $dW \equiv 0$ from the deterministic steady state until the state stops moving.

The baseline for the impulse responses of section 4.3 is the stochastic steady state; the shocked path is identical except for a single innovation in the first simulated period, a negative dividend shock of 0.1%. We report responses as period-by-period percent deviations of the shocked path from the baseline (the density's response is reported as a level difference).

The state dependence documented in figure 9 is a distribution over initial states, not over shock paths. We take each of the 9,600 periods of the ergodic simulation as an initial condition, recompute the baseline and shocked paths from that state with the same innovation, and report the resulting responses as a function of the state in which the shock arrives.