Asset Pricing with Heterogeneous Agents under Limited Information*

Riccardo A. Cioffi[†] Samuel Hurtado[‡] Galo Nuño[‡]

March 25, 2026

Abstract

In this paper, we study the interaction between wealth inequality and asset prices by developing a framework that features 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 tractable without sacrificing on 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: through heterogeneous portfolio exposure, aggregate shocks reshape the wealth distribution which in turn feeds back into asset demand and prices 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 portfolio heterogeneity, asset price movements have differential effects on households' wealth

*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.

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 heterogeneous agents, endogenous portfolio choice, and aggregate risk – a combination that has proven particularly challenging 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 income-wealth distribution, we restrict our attention to a limited information environment. Specifically, we assume that agents directly forecast asset prices and that they believe that prices 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 – exploiting their universal approximation properties – to capture the complex nonlinear relationship between prices and the distribution without imposing restrictive functional forms. The resulting equilibrium is 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 allows us to endogenously generate 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.

Such time variation in risk premia is a direct consequence of our approach: 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. A simple linear benchmark in the spirit of Krusell and Smith (1998) would capture the *expected* direction of price movements well, but also entirely miss the nonlinearity introduced by the volatility of price movements. Incidentally, this also makes precise where linear methods succeed and where they fail: the model requires estimating two objects from the same low-dimensional state, and a linear approximation would capture well only one of them. Such nonlinearity allows the model to endogenously generate state-dependent risk premia, which – combined with aggregate risk in equity returns and heterogeneous portfolio exposure – also produce significant distributional consequences that feed back into asset prices through market clearing.

Model We study a continuous-time endowment economy populated by households facing idiosyncratic income risk. Households optimally allocate their savings between a risky asset – a Lucas tree whose price is determined endogenously in equilibrium – and a risk-free bond. Uninsurable earnings risk, combined with an exogenous borrowing limit, generates heterogeneity in wealth and, crucially, in portfolio composition. A single aggregate shock drives fluctuations in dividends and – because households hold heterogeneous portfolios – directly reshapes the cross-sectional wealth distribution, making it stochastic.

To better understand why a model with these ingredients is particularly challenging, consider what solving it requires: households making endogenous portfolio choices across the wealth distribution, aggregate risk that directly reshapes the distribution through heterogeneous portfolio holdings, and endogenous asset prices determined by market clearing. Under rational expectations, households recognize that prices are equilibrium objects determined by the cross-sectional wealth distribution through market clearing; to forecast prices, they must therefore forecast how the distribution – an infinite-dimensional state variable – evolves. The household's value function becomes a function of this infinite-dimensional object, and the recursive problem takes the form of a Master equation – an infinite-dimensional equation in the space of distributions.

In a setting like ours, in which aggregate shocks *directly* affect households' wealth, this difficulty is particularly severe: rational households must have a perfect understanding not just of the expected distribution changes, but also of its stochastic dynamics – how aggregate shocks differentially reshape the distribution through heterogeneous portfolio exposure. Although recent computational advances have made progress in finding approximations to the Master equation, they usually maintain the assumption that agents track the entire cross-sectional distribution.[1] An assumption that, as Moll (2024) argues, raises fundamental conceptual concerns above and beyond its computational viability: real-world households and firms solve relatively simple problems; they do not forecast prices by forecasting cross-sectional distributions. We follow Moll's insight and take a different approach.[2]

Equilibrium 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. However, as they also point out, low-dimensional coupling alone does not really provide any simplification: the fact that prices depend on the distribution through market clearing implies that the process for prices would be it-

[1]See Bilal (2023) for an approach using perturbation methods and Gu et al. (2023) for one using neural networks.

[2]Our model is designed to isolate the mechanisms linking wealth inequality and asset prices under limited information, rather than being a fully quantitative exercise meant to match specific moments of the data; we discuss our specific calibration choices and their implications in Section 3.3.

self infinite-dimensional. Nonetheless, it does suggest a natural alternative to the Master equation approach: rather than requiring agents to solve for the infinite-dimensional diffusion process for prices, we model households as forecasting prices directly based on a low-dimensional vector of observable state variables – current dividends and the current price. The fact that the households' information set consists of dividends, prices, and the household's own state variables represents the natural mapping to information available to real-world investors. As such, modeling households as forecasting these variables directly is not a modeling compromise, but rather, a way to formalize what real agents actually observe and act upon.[3] Of course, the question remains of *how* agents forecast prices, given that the true price dynamics – as implied by the rational expectations model – would be infinite-dimensional functions of the entire cross-sectional distribution.

The way we tackle this is by approximating the drift and volatility of the price process with two separate neural networks, each specializing on one component of the dynamics: the drift captures the expected price movements in response to aggregate shocks, while the volatility captures how such shocks propagate through the economy. Additionally, since under limited information the distribution does not enter the households' state space, the households' problem is now a standard finite-dimensional problem. Hence, given a guess for the perceived law of motion for prices, we can use "standard" continuous-time methods to solve the household problem, track the wealth distribution by simulating the stochastic Kolmogorov forward equation, and clear the asset market at each instant to determine realized prices; finally, we train the neural networks on the simulated price series to update the perceived law of motion and iterate until convergence.

We also show that, in the self-confirming equilibrium (Fudenberg and Levine 1993; Cho and Sargent 2018), agents' forecasts are statistically consistent with the price dynamics that actually materialize: given the information they have at their disposal, no agent, upon observing the realized path of prices, would have reason to revise their price forecasts. Quantitatively, forecast errors computed across independent simulations are approximately normally distributed. As such, we never claim that our neural network solution is an approximation of the rational expectations equilibrium; rather, we study how agents with limited information form beliefs about prices using flexible forecasting rules. Indeed, because the Master equation has no known solution in this setting, a direct comparison between the two equilibria is infeasible – the intractability of the rational expectations benchmark is precisely what motivates our limited information approach. In this sense, our approach connects to the adaptive learning tradition, where agents condition on a subset of the available information; the neural network simply generalizes the parametric models used in the traditional learning

[3]Similarly, we maintain an exogenous risk-free rate as a deliberate choice to isolate the feedback between equity prices and the wealth distribution; endogenizing the interest rate – which would require clearing an additional market at every instant – is a computationally demanding extension that we leave for future work.

literature by not imposing a functional form on the perceived law of motion.

The structure of households' optimal portfolio decisions is central to the model's equilibrium mechanism. In the absence of substantive trading frictions, the optimal portfolio is a step function: households want to hold a 100% share in equity when expected excess returns are high – that is, when prices are low relative to dividends – and no equity when expected excess returns are low, with an extremely narrow transition region between the two corner solutions.[4] By construction, market clearing pushes the equilibrium to sit in this narrow transition region: if the typical household were at either corner, excess demand or supply would push prices until portfolio choice becomes interior and the market would clear. At the same time, however, the wealth distribution evolves continuously and stochastically through heterogeneous portfolio exposure to the common aggregate shock, and market clearing must tie the equilibrium price to this moving distribution at every instant. This combination – a step-function portfolio surface, a stochastic distribution, and a narrow equilibrium region – makes the problem fundamentally nonlinear and poorly suited to local approximation methods. The step-function structure also means that the household's HJB equation has a kink at the portfolio boundary, making second-derivative terms highly sensitive to numerical precision, which is one of the key reasons why we solve the household problem with standard finite-difference methods and reserve the neural network approximation for the perceived law of motion alone.

Price dynamics The asymmetry between drift and volatility functions we discussed above emerges endogenously from the learned price dynamics. The drift is easy to understand: prices above their long-run level tend to fall, and prices below it tend to rise – precisely the pattern of mean-reversion one would obtain from a standard linear approximation. The volatility, however, is different and highly nonlinear: near zero for low asset prices and rising as prices increase, with substantial curvature even within the ergodic region. This difference is driven by the fact that, unlike the drift, the volatility term captures how dividend shocks propagate through the economy depending on its current state, which involves the full nonlinear interaction between heterogeneous portfolios, market clearing, and the wealth distribution. The intuition is that high asset prices require high equity demand, which means more households are exposed to aggregate fluctuations; a given dividend shock therefore affects more portfolios simultaneously, amplifying its transmission to prices and raising volatility.

In fact, because price volatility varies with the aggregate state, the risk borne by equity holders is itself time-varying, and so is the compensation they require. When prices are low, the dividend yield is high and the mean-reverting drift points upward, while volatility is

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

low – producing high expected excess returns at low perceived risk. Conversely, when prices are high, the yield falls and the drift turns negative, while volatility rises – compressing the premium. The result is a negative relationship between the price level and risk premia that arises naturally from our global nonlinear solution – a relationship which local methods up to the second order cannot generate.

Finally, this time variation in risk premia, coupled with aggregate risk in asset returns and heterogeneous portfolio exposure, also produces significant distributional consequences that reflect back into the evolution of prices. Households' exposure to asset-price fluctuations depends on their equity holdings and on the share of their total income coming from labor: households that are near the constraint receive most of their income in the form of labor earnings, which keeps them relatively isolated from fluctuations in equity returns; households that are farther away from the constraint have significant wealth invested in equity and most of their income is tied to the risky asset.[5] As a consequence, when a negative aggregate shock pushes prices down, households with positive equity exposure suffer capital losses, while those near the constraint are largely shielded; the wealth distribution shifts, with mass accumulating near the borrowing constraint and the right tail thinning. This redistribution feeds back into aggregate asset demand through market clearing – the reshaped distribution changes who is the marginal buyer of equity, affecting prices. The fall in prices, however, temporarily creates a regime of high expected returns through the time-varying premium; households respond by substituting toward equity, and the market-clearing condition translates this higher demand into a price recovery. This two-way feedback between the distribution and prices – aggregate shocks reshaping the distribution through heterogeneous portfolios, and the reshaped distribution feeding back into asset demand and prices through market clearing – is precisely the mechanism that motivates the paper, which is made quantitatively tractable via our limited information framework and operational via our neural network solution.

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

1.1 Relation to Literature

Our paper is connected to at least three main strands of the literature: first, it relates to the literature examining the interplay between wealth inequality and asset prices; second,

[5]Notice that this heterogeneous exposure arises endogenously from a common perceived law of motion: all households share the same beliefs about price dynamics, but the portfolio first-order condition maps these common beliefs into different optimal equity shares depending on wealth – the portfolio distortion is wealth-dependent, not belief-dependent.

it relates to the literature on limited information, learning, and alternatives to rational expectations; and third, it relates to the literature applying neural networks and machine learning methods in macroeconomics.

A large empirical literature documents that asset returns vary systematically with wealth, driven by portfolio composition (Fagereng et al. 2020; L. Bach et al. 2020; Hubmer, Halvorsen, et al. 2023).[6,7] 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. Cioffi (2021) quantifies this channel computationally, finding that asset price movements disproportionately affect the wealthy and can account for a substantial fraction of the rise in top wealth shares. Greenwald et al. (2023) show that declining real interest rates have concentrated financial wealth through portfolio duration heterogeneity, and Fernández-Villaverde and Levintal (2024) quantify the distributional welfare effects of asset returns in a heterogeneous-agent economy with endogenous portfolio choices – both confirming that the composition of household portfolios is a key channel through which returns shape inequality. If asset returns affect wealth inequality through portfolios, how does wealth inequality feed back into asset prices?

Time-varying expected returns are a foundational empirical fact in asset pricing (Campbell and Shiller 1988; Fama and French 1988; Cochrane 2011). Representative-agent models usually generate time-varying returns either through non-standard preferences – e.g., external habit formation (Campbell and Cochrane 1999) – or long-run consumption risks with Epstein-Zin preferences (Bansal and Yaron 2004). Our approach instead generates time-varying risk premia through distributional dynamics: the wealth distribution is a state variable for equilibrium prices, as in the heterogeneous-agent tradition (see also 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, among others).[8]

Hubmer, Krusell, et al. (2021) quantify the importance of portfolio heterogeneity for wealth inequality: in their decomposition of the U.S. top-10% wealth share, return heterogeneity contributes +29.5 percentage points – the single largest channel. However, their model features exogenous portfolio shares and exogenous asset prices; they describe endogenizing both as "one or two orders of magnitude more challenging" than the problem they solve. As a consequence, myopia and perfect foresight produce nearly identical wealth distributions in their framework – approximately 2 percentage points difference at the top-10%

[6] Evidence also suggests that risk tolerance increases with wealth (Calvet and Sodini 2014; Meeuwis 2020), reinforcing the portfolio heterogeneity channel.

[7] For a survey of the mechanisms generating wealth inequality, including return heterogeneity, see Benhabib and Bisin (2018).

[8] See Panageas (2019) for a survey of heterogeneous-agent asset pricing models.

share – because households cannot act on foreseen return changes. While our model is quantitatively less rich, we take precisely those steps: we endogenize both portfolio choice and asset prices, so that beliefs drive portfolios, portfolios drive prices, and prices feed back into beliefs – a loop that requires both margins to operate.

Gomez (2025) embeds portfolio heterogeneity in a two-agent continuous-time general equilibrium model, generating a distributional feedback loop between wealth concentration and asset prices. Our paper instead derives portfolio choice endogenously from optimal household decisions under limited information, so the feedback between the wealth distribution and prices arises from agents' perceived price dynamics. Relatedly, Krueger and Lustig (2010) formalize conditions under which incomplete markets are irrelevant for the equity premium; countercyclical cross-sectional income variance and limited participation are two channels through which this irrelevance breaks down – though Gomes and Michaelides (2008) show that the latter requires participation constraints to be imposed exogenously rather than arising endogenously. Our model breaks these irrelevance conditions through a third channel: agents' perceived price dynamics generate wealth-dependent portfolio distortions even under common beliefs, because households at different points in the wealth distribution respond differently to the same perceived return moments.

Our limited information assumption is motivated both empirically and conceptually. 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.[9] Beyond the empirical evidence, Moll (2024) argues that rational expectations is fundamentally intractable in heterogeneous-agent settings: agents must forecast the infinite-dimensional wealth distribution because prices are non-Markovian (even though the distribution itself might follow a Markov process). Departing from rational expectations is therefore not merely a modeling choice but a response to the conceptual impossibility of requiring agents to track objects that even the econometrician cannot observe.

Our limited information framework builds on Adam and Marcet (2011), who formalize internal rationality: agents optimize given dynamically consistent subjective beliefs without requiring 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 – establishing learning as an alternative to preference engineering for

[9]See also Coibion et al. (2018) for evidence that 60% of New Zealand firms do not track inflation and Angeletos et al. (2021) for a survey of systematic departures from rational expectations.

generating time-varying risk premia. 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 – where perceived and actual dynamics coincide – follows the T-mapping framework of Marcet and Sargent (1989); Bhandari, Borovička, et al. (2025) develop tools for validating the resulting belief wedges against survey data.[10] The mathematical foundations come from Moll and Ryzhik (2025): in mean field games with low-dimensional coupling, our departure from rational expectations allows us to replace the infinite-dimensional Master equation with a standard finite-dimensional HJB where prices are the relevant state variable.

Several papers study the interaction between wealth heterogeneity and beliefs, including wealth-weighted belief aggregation (Caballero and Simsek 2020) and experience-based learning (Ehling et al. 2018; Collin-Dufresne et al. 2017). Martin and Papadimitriou (2022) show analytically that wealth redistribution among heterogeneous-belief agents generates time-varying sentiment; however, their aggregation to a representative agent requires log utility. Our model addresses a related question – whether distributional dynamics can generate time-varying risk premia under incomplete markets – however, our model does not feature heterogeneous beliefs. 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 and find that belief heterogeneity amplifies the top-1% wealth share substantially. Our approach differs in two respects: all agents share a common perceived law of motion for prices, and beliefs satisfy self-confirming equilibrium discipline, so that forecasts are not systematically wrong along the equilibrium path. The effect of distributional dynamics on risk premia in our model therefore operates through portfolio heterogeneity under common beliefs – endogenous portfolio choice maps homogeneous beliefs into heterogeneous risk exposures across the wealth distribution – rather than through cross-sectional belief dispersion.

On the computational side, Krusell and Smith (1998) introduced the perceived-law-of-motion approach, showing that approximate aggregation holds ($R^2 > 0.9999$) in a single-asset economy. Ahn et al. (2017) show that this fails in two-asset models, where distributional effects require substantially richer representations of the cross-sectional density. Existing methods provide local approximations – sequence-space Jacobians (Auclert, Bardóczy, et al. 2021; Auclert, Rognlie, et al. 2024), discrete-time perturbation (Bayer and Luetticke 2020; Bhandari, Bourany, et al. 2023), and other linearization techniques (Bilal 2023; Reiter 2009) – but are by construction insufficient for capturing the nonlinearities that generate time-varying risk premia. Lee (2025) develops a global nonlinear solution in sequence space

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

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 with multiple assets and aggregate risk directly affecting household wealth through a stochastic Kolmogorov Forward equation – to our knowledge, the first such method in this class of models.

Fernández-Villaverde, Hurtado, et al. (2023) show that neural-network perceived laws of motion ($R^2 = 0.99$ versus 0.83 for linear) reveal equilibrium features invisible to log-linear methods in a continuous-time heterogeneous-agent model with financial frictions. We extend their infrastructure from a single-asset intermediary model to two-asset household portfolio choice with endogenous Lucas-tree pricing; the main conceptual distinction is that, while they use neural networks to approximate the equilibrium law of motion using the first moment of the wealth distribution, we instead use them to model agents' actual beliefs about price dynamics. The mathematical 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; Yang et al. 2025).

Our paper brings together three strands: (i) the wealth distribution as a state variable for asset pricing, building on the heterogeneous-agent literature from Dumas (1989) through Hubmer, Krusell, et al. (2021) and Gomez (2025); (ii) limited information about price dynamics, building on Adam and Marcet (2011) and Moll and Ryzhik (2025); and (iii) neural-network global solution methods, extending Fernández-Villaverde, Hurtado, et al. (2023) to two-asset household portfolio choice. The combination 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. The next section presents the model.

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.[11]

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.

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 assigns measure $\lambda_2/(\lambda_1 + \lambda_2)$ to the unemployed state and $\lambda_1/(\lambda_1 + \lambda_2)$ 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.

[11] 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.

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 $r^* > 0$.[12] 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:

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

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.

[12] 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.

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.[13]

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 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.[14] 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 their 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):

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

where ψ is the birth distribution; the drift component $\hat{\mathcal{A}}_t^*[g]$ captures the deterministic flow

[13] 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).

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

of households across wealth and income states due to optimal consumption, portfolio choice, and income transitions; 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.[15]

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

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

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

[15] 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)).

substitute the dynamics of dY_t and dg_t from Equations (4) and (8), which gives us:

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

where, importantly, in the last term the coefficient in front of dW_t in square brackets gives us asset price volatility:

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

which shows that asset price volatility arises from both direct dividend volatility and the stochastic evolution of the wealth distribution.[16]

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 the standard partial derivative $\left(\frac{\partial \mathcal{Q}}{\partial g_i}\right)_{i=1}^N$ 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

[16]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.

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.[17]

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:

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

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

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

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 $\hat{\mathcal{B}}_t^*[g](a', z')$ 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:[18]

$$\left\langle \left\langle \frac{\delta^2 \mathcal{Q}}{\delta g^2}, \hat{\mathcal{B}}_t^*[g] \otimes \hat{\mathcal{B}}_t^*[g] \right\rangle \right\rangle = \int \int \frac{\delta^2 \mathcal{Q}(Y, g)}{\delta g^2}[a_1, z_1][a_2, z_2] \cdot \hat{\mathcal{B}}_t^*[g](a_1, z_1) \cdot \hat{\mathcal{B}}_t^*[g](a_2, z_2)\, da_1 dz_1 da_2 dz_2$$

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 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.

[18] 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 $\hat{\mathcal{B}}_t^*[g] \otimes \hat{\mathcal{B}}_t^*[g]$ creates the appropriate bilinear object. That is, for any two points (a_1, z_1) and (a_2, z_2), the tensor product equals $\hat{\mathcal{B}}^*[g](a_1, z_1) \cdot \hat{\mathcal{B}}^*[g](a_2, z_2)$.

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 very 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.[19]

The Problem with Rational Expectations The extremely complicated structure of the Master equation (15) is a direct consequence of the extraordinary informational requirements

[19]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. Notice that, however, 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."

that are imposed on households by the rational expectations assumption in the context of heterogeneous-agent models – an extreme form of the curse of dimensionality. 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 as well as 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. As Moll (2024) and Moll and Ryzhik (2025) argue, real-world households and firms do not actually forecast prices by forecasting distributions and instead solve simpler problems. 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.[20]

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:

$$dQ(\Omega_t) = \mu_Q(\Omega_t)dt + \sigma_Q(\Omega_t)dW_t \tag{16}$$

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

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

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

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 – 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.

[20] 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.

Importantly, such 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).[21]

There remains of course to be answered 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 price level. The second is that we allow both the drift and diffusion 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 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

[21] 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.

allows for richer dynamics.[22]

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.

[22]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):[23]

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

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.

[23]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:[24]

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

where the cross-derivative terms arise because the aggregate shock dW_t simultaneously affects individual wealth (through portfolio returns), dividends, and asset prices.[25]

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. The first term is the standard

[24] 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

[25] 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).

myopic demand: the household holds the risky asset in proportion to the excess return $R(Y,Q) - r^*$, scaled by the ratio $-v_a/v_{aa}$ – which equals a/γ under CRRA preferences – 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 propagate 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:

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

where $\psi(a, z)$ is the birth distribution and the operators $\mathcal{A}^*$ and $\mathcal{B}^*$ are defined as:

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

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

where μ_a and σ_a are defined as in Equations (23) and (24).

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

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

(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 functional

[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.

forms *ex ante*. Second, they provide a global solution method that can handle nonlinearities throughout the state space, avoiding the local approximation errors arising 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

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

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.

(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:

$$\hat{\mu}_{Q,t} = \frac{\Delta Q_{t+\Delta t}}{\Delta t} \tag{32}$$
$$\hat{\sigma}_{Q,t}^2 = \frac{(\Delta Q_{t+\Delta t})^2 - (\mu_{Q,n}(Y_t, Q_t))^2}{\Delta t} \tag{33}$$

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

$$\alpha^{n+1} = \arg\min_{\alpha} \frac{1}{2}\sum_{t=1}^{T} \|\mu_Q(\Omega_t;\alpha) - \hat{\mu}_{Q,t}\|^2 \tag{34}$$
$$\beta^{n+1} = \arg\min_{\beta} \frac{1}{2}\sum_{t=1}^{T} \|\sigma_Q(\Omega_t;\beta) - \hat{\sigma}_{Q,t}\|^2 \tag{35}$$

(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.

3.1.1 Network Architecture and Training

The functional approximation problem we face requires capturing potentially complex nonlinear 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:

$$\mu_Q(Y, Q; \alpha) = \alpha_0^{(2)} + \sum_{k=1}^{K} \alpha_k^{(2)} \varphi\left(\alpha_{k,0}^{(1)} + \alpha_{k,1}^{(1)} Y + \alpha_{k,2}^{(1)} Q\right) \tag{36}$$

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

[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).

analogous structure with its own parameter vector β.[28]

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) provides 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), 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.[29]

[28]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.

[29]The algorithm runs for several thousand outer iterations, and we declare the model converged when

3.2 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 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, bounded 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

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 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.

than would be possible with a single, compromise architecture. This modular approach also facilitates debugging and interpretation – 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.

3.3 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. 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 72% 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 $\eta = 0.1$ implies moderately persistent dividend shocks, while the volatility parameter $\sigma = 0.09$ 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.

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
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	0.918	Normalization
Dividend Process		
$\bar{Y}$	0.1	Dividend yield
η	0.1	Dividend persistence
σ	0.09	Return volatility

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 heterogeneity and time-varying risk premia generate the aggregate dynamics documented in the impulse responses. Fifth, we illustrate the distributional implications of aggregate shocks, tracing the mechanism from heterogeneous portfolios 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 a simple first-moment summary of the distribution – here, the price level itself – 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.

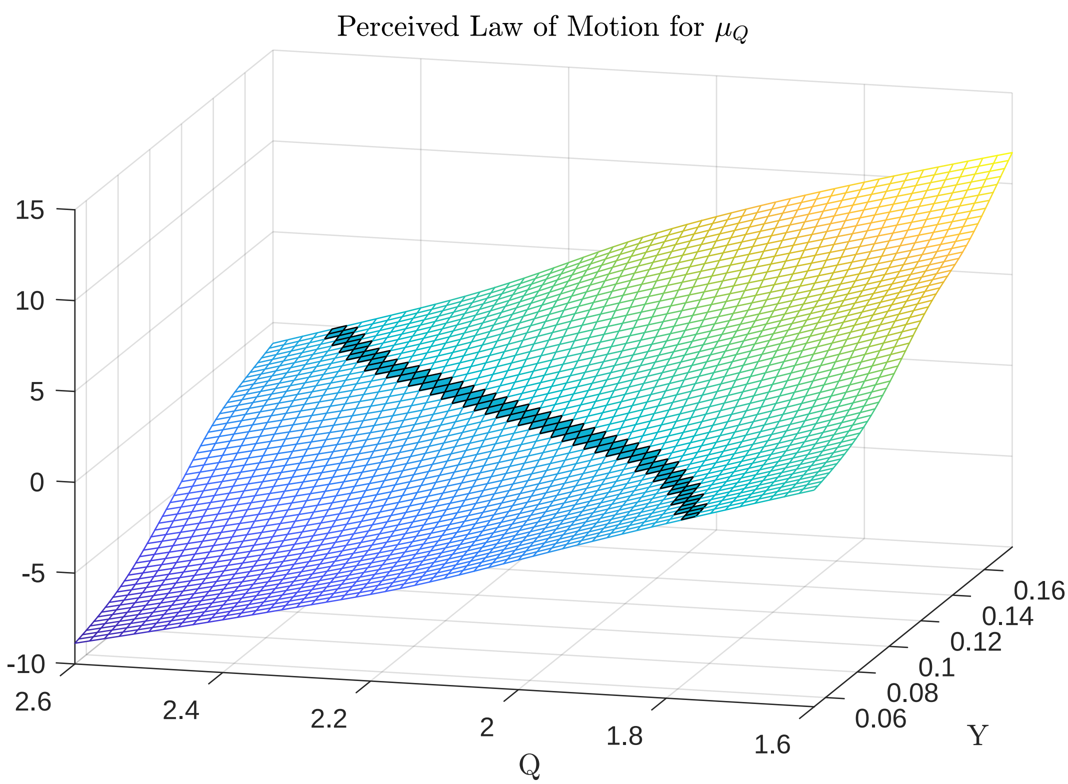

Figure 1: Perceived law of motion: drift $\mu_Q(Y, Q)$.

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 a first-moment summary 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.

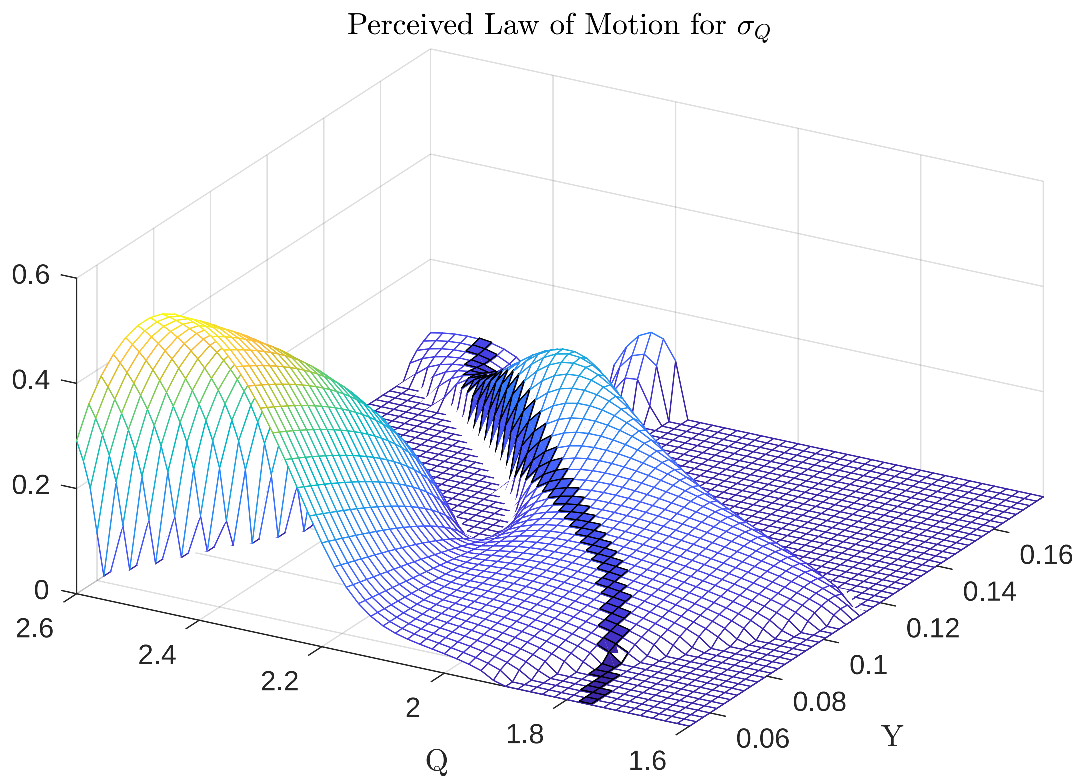

Figure 2: Perceived law of motion: diffusion $\sigma_Q(Y, Q)$.

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 – which are rooted in the interaction between heterogeneous portfolios and market clearing – are examined in Section 4.4, after we document the aggregate dynamics of the model. For now, we note that the asymmetric pattern in σ_Q reflects the fact that the transmission of dividend shocks to prices depends on how many households are actively participating in the equity market at a given aggregate state – a mechanism that becomes visible once we examine portfolio choice and impulse responses below.

To quantify more precisely where the linear benchmark succeeds and where it fails, we fit linear regressions of the form $\mu_Q = a_0 + a_1 Y + a_2 Q$ and $\sigma_Q = b_0 + b_1 Y + b_2 Q$ on simulated data and compare in-sample fit. The linear drift regression achieves an R^2 of 0.6 – consistent with the near-linearity of the drift surface visible in Figure 1. By contrast, the linear diffusion regression achieves an R^2 of only 0.016, confirming that the nonlinear neural network specification is essential for capturing the diffusion. [30] In the Krusell and Smith (1998) framework, the first moment of the wealth distribution is a sufficient statistic for prices because the model is sufficiently linear that higher-order variation is economically negligible. Our model lacks that linearity: the drift and volatility carry distinct economic content – the drift captures the direction of price movements, while the diffusion captures state-dependent risk – and the latter requires nonlinear approximation.

4.2 Self-Confirming Equilibrium

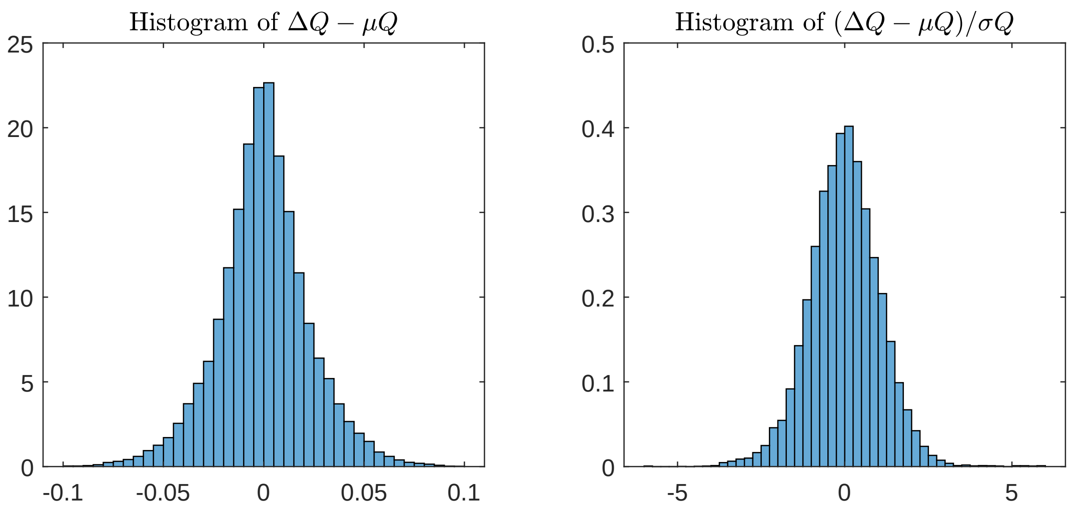

Figure 3: Out-of-sample validation of the perceived law of motion.

Notes: Left panel: histogram of raw forecast errors $\Delta Q - \mu_Q(Y_t, Q_t)$. 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 0.0003), 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.[31] We test this by simulating the economy 100 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

[30] For comparison, the neural networks achieve R^2 coefficients of 0.79 and 0.03 for the drift and volatility terms, respectively.

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

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 95% confidence 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 – 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.

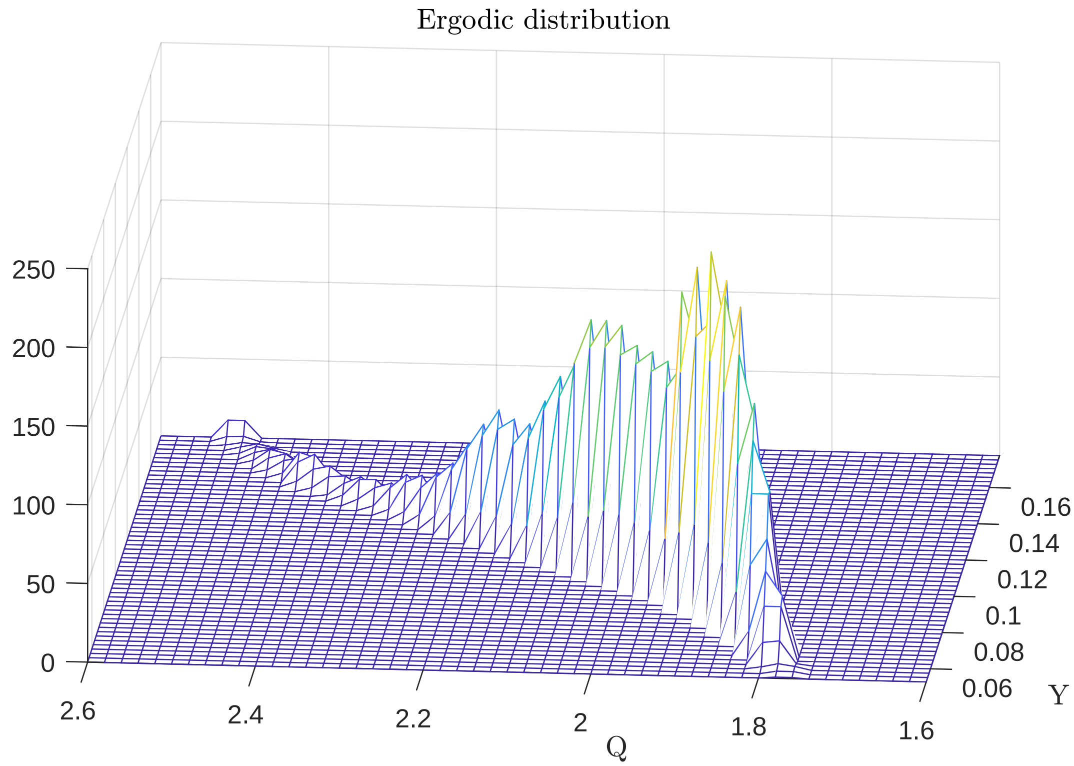

Figure 4: Joint ergodic distribution of dividends Y and asset prices Q.
Notes: The distribution is concentrated around $Q \in [1.8, 2.2]$ and $Y \in [0.08, 0.14]$, confirming the long-run stability of the equilibrium. Prices and dividends co-move positively, but the distribution of Q is more concentrated than that of Y, reflecting the smoothing of dividend shocks through equilibrium prices.

The ergodic distribution has two 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.

4.3 Aggregate Dynamics and Impulse Responses

We now trace the aggregate dynamics of the model by computing an impulse response to a negative aggregate dividend shock of 0.1%, starting from the stochastic steady state.

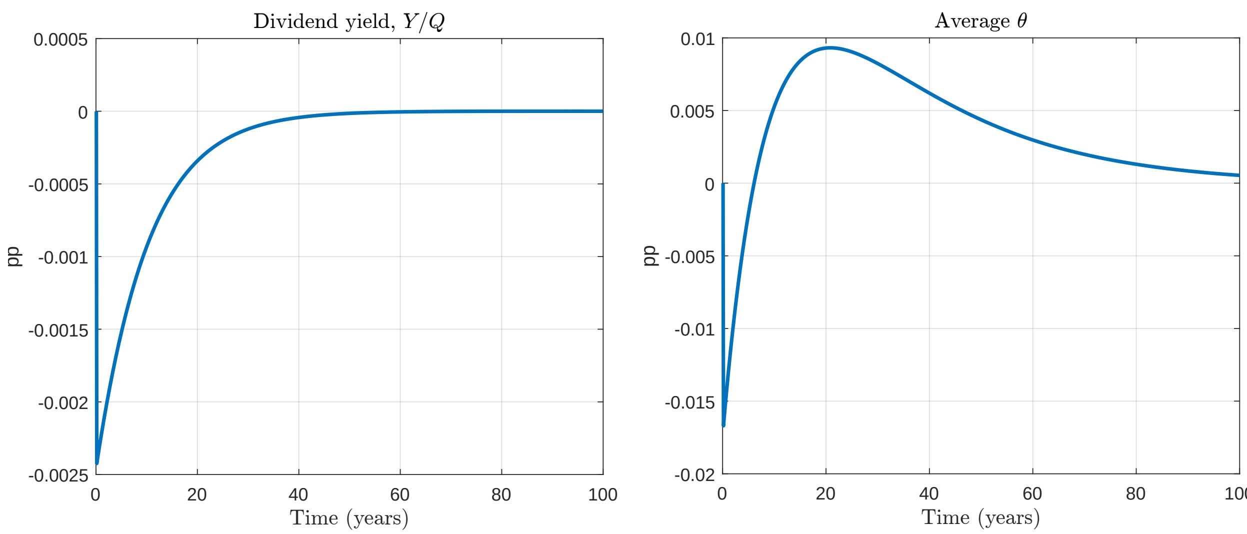

(a) Dividend yield Y/Q.

(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. 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, more interesting observation, is that the non-monotone response in θ^* is precisely the mechanism through which the shock is absorbed, and understanding it requires examining how portfolio heterogeneity and time-varying risk premia interact – which happens through two main channels. First, a *valuation channel*: on impact, the decline in Q mechanically compresses the equity share in household portfolios, driving θ^* down. Second, a *portfolio rebalancing channel*: lower prices raise the dividend yield Y/Q and – via equilibrium volatility – reduce perceived risk, improving the risk-return tradeoff facing each household; this induces a substitution toward equity that pushes θ^* back up and, ultimately, above its pre-shock level. The resulting overshoot in portfolio shares is also what eventually stabilizes prices, with higher equity demand pulling prices back towards the steady state.

The mechanism behind these dynamics – in particular the overshoot in portfolio shares and the asymmetric price recovery – is rooted in portfolio heterogeneity and time-varying risk premia, which we turn to next.

4.4 Portfolio Heterogeneity and Time-Varying Risk Premia

The valuation and rebalancing channels identified in the impulse responses operate through the interaction of households' optimal portfolio choice with time-varying risk premia; we now examine both objects in turn, beginning with the portfolio surface.

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.[32] 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 exact location and shape naturally vary depending on the household's individual state variables (income and wealth). The blue scatter represents the average θ^* across the simulation, it ranges approximately between 0.53 and 0.65 and, since we are showing the surface for a household with average wealth, lies almost exactly on the transition region.

[32] To improve its readability, notice that the Q axis has been inverted in Figure 6, with higher prices on the left and lower prices on the right.

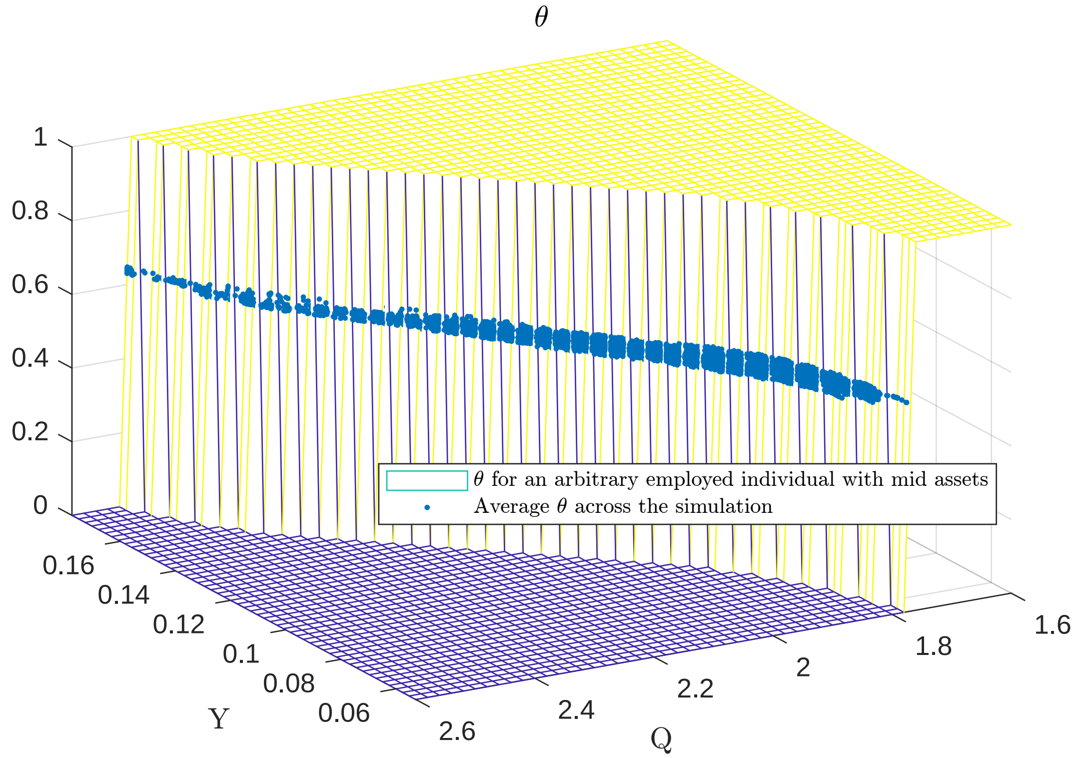

Figure 6: Optimal portfolio share θ^*.

Notes: The yellow wireframe 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 their 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.

The location of the transition region is in turn determined by the interaction of the drift and diffusion surfaces (Figures 1 and 2) with the cross-sectional wealth distribution. 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.

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 $(Y + \mu_Q)/Q$ and the risk premium $((Y + \mu_Q)/Q - r^*)/\sigma_Q$ as surfaces over the aggregate state. Expected returns rise sharply at low prices and high dividends; the risk premium exhibits a pronounced nonlinearity in the interior of the state space, confirming that the state dependence is quantitatively large.

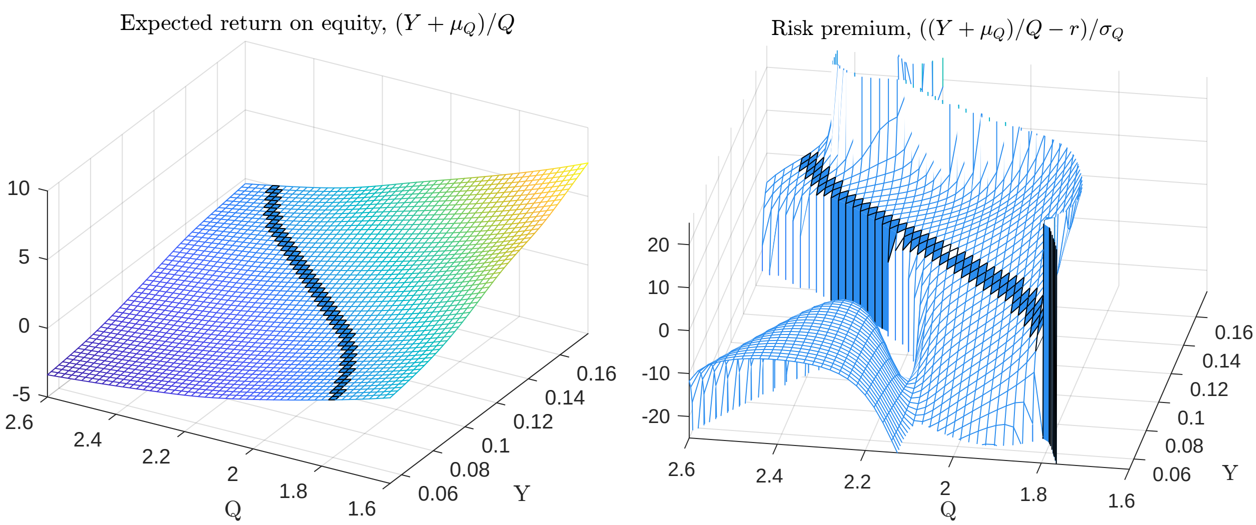

(a) Expected equity return $(Y + \mu_Q)/Q$.

(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 ratio is ill-defined in regions where $\sigma_Q \approx 0$; the economically relevant region is the ergodic band shown in Figure 4.

Figure 8 confirms that 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 sharply 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.

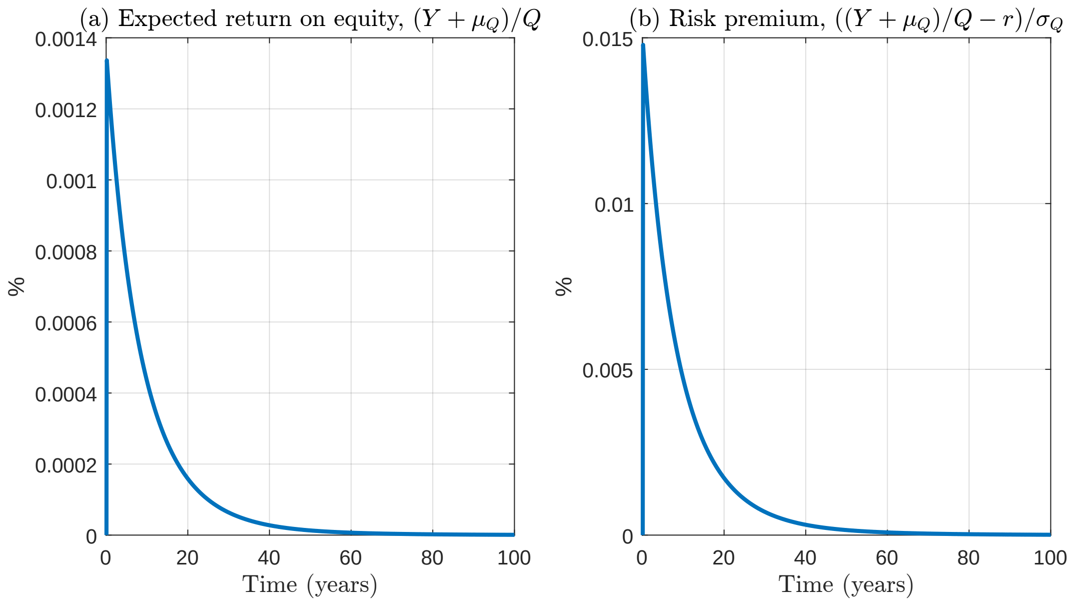

Figure 8: Impulse response of the expected equity return and risk premium to a negative aggregate shock.

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.

Because $\sigma_Q(Y,Q)$ varies with the aggregate state, so does the risk borne by equity holders – and hence the compensation they require. When prices are low, the dividend yield Y/Q is high and the mean-reverting drift points upward, while volatility is near zero – producing a large expected excess return at low risk. When prices are high, the dividend yield falls and the drift turns negative, while volatility is elevated – compressing the premium.

Hence, as a direct consequence of heterogeneity in portfolios and state-dependent volatility, our model endogenously generates a negative relationship between the price level and expected excess returns – a well-documented empirical pattern in the finance literature, which shows that valuation ratios (in particular the dividend-price ratio) positively predict excess equity returns (Campbell and Shiller 1988; Fama and French 1988; Cochrane 2011).

The economic mechanism behind this state dependence is rooted in how dividend shocks propagate differently through the economy depending on its current state, which in turn determines the behavior of σ_Q documented in Figure 2. 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.

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 – the initial drop followed by the overshoot visible in Figure 5b – reflects two distinct channels operating in sequence.

Consider first the *valuation channel* – the direct effect of the shock on equity values and dividend income. The negative dividend shock moves the economy along the transition region described in Figure 6: as dividends fall and prices decline, the economy shifts to a new point on the transition region where the equilibrium portfolio share is lower. The drop in prices mechanically reduces the value of equity holdings and therefore the average θ^*; the on-impact decline visible in Figure 5b is the direct consequence of this movement along the portfolio surface.

As prices fall, however, a second – endogenous – force sets in: the *portfolio rebalancing channel*. The dividend yield Y/Q rises (visible in Figure 5a), and at lower prices the mean-reverting drift points upward while volatility is smaller, precisely the conditions documented in Figures 1 and 2. This combination means that the excess return $R(Y, Q) - r^*$ in the risk-return tradeoff term of the portfolio first-order condition rises while the perceived variance $(\sigma_Q/Q)^2$ falls – both forces pushing θ^* upward. Households respond by substituting toward the riskier asset, and θ^* not only recovers but overshoots 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 reflects through an upward shift in θ^*.

The θ^* overshoot is in turn part of the mechanism of price recovery itself. Higher equity demand – the consequence of temporarily elevated expected returns – pulls prices back toward their steady-state level. In equilibrium, the portfolio rebalancing acts as an endogenous stabilizer: the market-clearing condition translates higher θ^* into higher Q, and the resulting price increase reduces expected returns and brings θ^* back down. The overshoot in portfolios is precisely what prevents an overshoot in prices; if households did not substitute toward equity when expected returns were high, prices would recover more slowly. 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.

This whole mechanism – from the portfolio surface to time-varying risk premia to endogenous stabilization through portfolio rebalancing – 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 interaction between heterogeneous portfolios and time-varying risk premia operates at the aggregate level, but because households differ in their equity exposure at a given (Y, Q), aggregate shocks also reshape the cross-sectional wealth distribution. Notice that this heterogeneity operates entirely under common beliefs: all households perceive the same price dynamics through the PLM, but the portfolio first-order condition maps these into different optimal equity shares depending on each household's position in the wealth distribution.

Figure 9 presents the cross-sectional wealth distribution at two representative aggregate states: a high-dividend state (with correspondingly high prices) and a low-dividend state (with low prices). Both distributions are concentrated at low wealth levels, but they differ systematically. When dividends are high, the distribution is less dispersed with a thinner right tail and a higher peak at $a = 0$; when dividends are low, the distribution is more dispersed with a lower peak near the borrowing constraint.

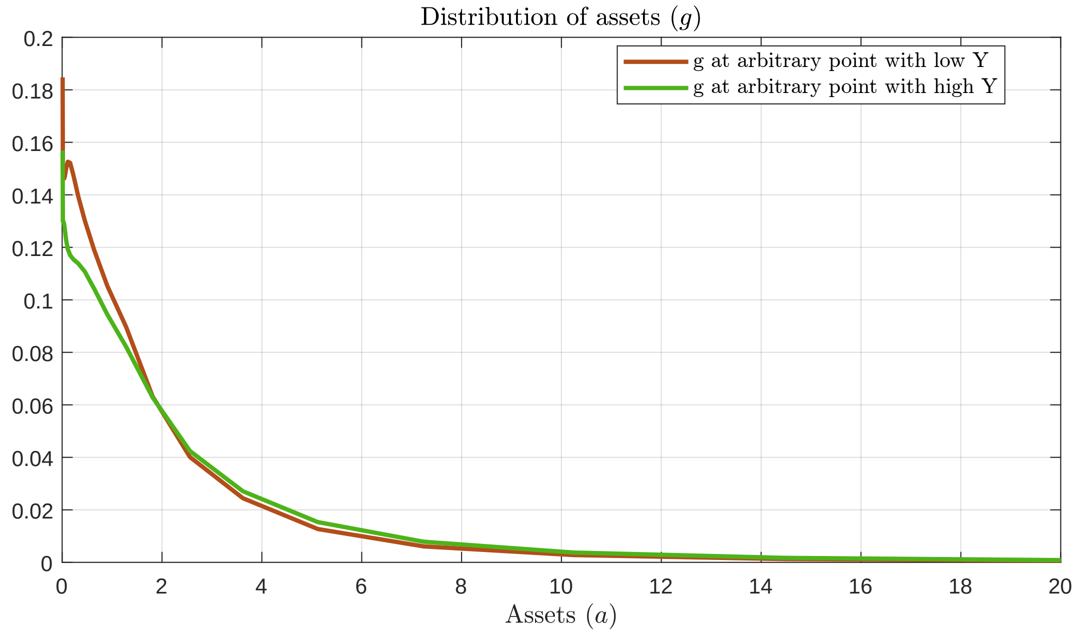

Figure 9: Cross-sectional wealth distribution across states.

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. The distribution is less dispersed with a thinner right tail when dividends are high, reflecting state-dependent wealth inequality driven by heterogeneous portfolio exposure.

The distribution is therefore also time-varying: aggregate shocks reshape it through heterogeneous portfolio exposure. This mechanism has a precise footprint in the distribution's response to the IRF shock, which we examine next.

Figure 10 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 the portfolio heterogeneity channel 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.

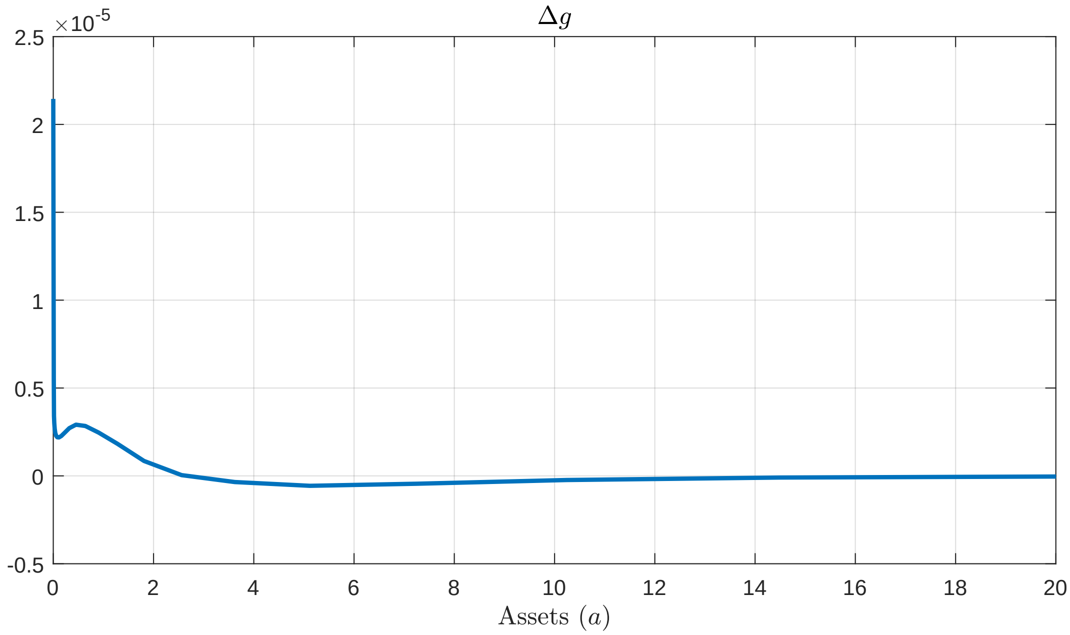

Figure 10: Change in the cross-sectional wealth distribution following a negative aggregate shock.

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, heterogeneous portfolio responses, and endogenous stabilization through rebalancing.

5 Conclusion

In this paper, we developed 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 showed 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 keeping 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. Neither do we 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.

Ahn, S., G. Kaplan, B. Moll, T. Winberry, and C. Wolf (2017). "When Inequality Matters for Macro and Macro Matters for Inequality". In: *NBER Macroeconomics Annual 2017* 32, pp. 1–75.

Angeletos, G.-M., Z. Huo, and K. A. Sastry (2021). "Imperfect Macroeconomic Expectations: Evidence and Theory". In: *NBER Macroeconomics Annual*.

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., J. Borovička, and P. Ho (2025). "Survey Data and Subjective Beliefs in Business Cycle Models". In: *Review of Economic Studies* 92.3, pp. 1375–1437.

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.

Caballero, R. J. and A. Simsek (2020). "A Risk-Centric Model of Demand Recessions and Speculation*". In: *The Quarterly Journal of Economics* 135.3, pp. 1493–1566.

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.

Coibion, O., Y. Gorodnichenko, and S. Kumar (2018). “How Do Firms Form Their Expectations? New Survey Evidence”. In: *The American Economic Review* 108.9, pp. 2671–2713.

Collin-Dufresne, P., M. Johannes, and L. A. Lochstoer (2017). “Asset Pricing When 'This Time Is Different'”. In: *The Review of Financial Studies* 30.2, pp. 505–538.

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.

Ehling, P., A. Graniero, and C. Heyerdahl-Larsen (2018). “Asset Prices and Portfolio Choice with Learning from Experience”. In: *The Review of Economic Studies* 85 (3 (304)), pp. 1752–1780.

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.

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 (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*.

Panageas, S. (2019). “The Implications of Heterogeneity and Inequality for Asset Pricing”. In: *Working Paper*.

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*.