Asset Pricing, Participation Constraints, and Inequality

Goutham Gopalakrishna*, Zhouzhou Gu†, Jonathan Payne‡

— Preliminary Draft —

March 27, 2024

Abstract

How do asset returns interact with wealth inequality? Empirical evidence shows that portfolio choices and financial constraints lead to unequal risk exposure across households and financial intermediaries. To understand the dynamic general equilibrium implications, we build a macroeconomic model with heterogeneous households, a financial sector, asset market participation constraints, and endogenous asset price volatility. We develop a new deep learning methodology for characterizing global solutions to this class of macro-finance models. We show that wealth inequality, financial sector recovery, and asset price dynamics depends on which households are able to purchase assets during crisis. This means the government faces a trade-off between tighter leverage constraints and a more equal recovery. In our calibrated model, asset returns and participation constraints account for a large fraction of the change in wealth inequality over the past half-century.

Keywords: Inequality, Participation Constraints, Heterogeneous Agent Macroeconomic Models, Asset Pricing, Deep Learning.

*University of Toronto, Rotman School of Management.
Email: goutham.gopalakrishna@rotman.utoronto.ca

†Princeton, Department of Economics. Email: zg3990@princeton.edu

‡Princeton, Department of Economics. Email: jepayne@princeton.edu

1 Introduction

Many researchers have documented that US wealth inequality has increased during the post-war period (e.g. Piketty (2014), Saez and Zucman (2016)). This has motivated extensive debate about the macroeconomic drivers of wealth inequality. A much discussed channel is that asset portfolios vary across the wealth distribution and so earn different returns. This suggests that asset price dynamics and government interventions in financial markets have an important impact on inequality. However, macroeconomists have had difficulty studying these effects because they have struggled to characterize global solutions to heterogeneous agent models with aggregate risk and complicated portfolio restrictions. We develop a deep learning algorithm that overcomes these technical difficulties and allows researchers to bring insights from the asset pricing literature into the debate over inequality. We use a calibrated general equilibrium macro-finance model to understand how successfully portfolio heterogeneity can explain inequality dynamics and to what extent the government needs to balance financial sector stabilization against inequality.

We study a heterogeneous agent business cycle model with financial intermediaries, participation constraints, aggregate risk, and endogenous asset price volatility. The economy contains short-term risk-free assets and risky capital stock that generates output. It is populated by a large collection of price-taking households who face idiosyncratic death shocks and a penalty for holding capital that is relaxed as agents get wealthier. This leads to a non-trivial household distribution of wealth and portfolio choices. The economy also contains a financial expert that can freely hold capital but cannot raise equity. The financial frictions on the households and financial expert generate endogenous volatility in the capital price process that depends upon the distribution wealth in the economy.

We start by investigating the different mechanisms connecting inequality and asset pricing in our model. The evolution of wealth inequality can be decomposed into the following forces: (i) wealthier agents have more access to capital markets, which allows them to build wealth more quickly (sometimes referred to as the "scaling" force in the literature) but also increases their exposure to aggregate risk, (ii) wealthier agents have a higher propensity to consume out of wealth, and (iii) agent exits increase the wealth share of poorer agents because agents have an imperfect bequest motive. The strength of the first "scaling" force depends on the health of the financial sector. When the financial sector takes losses in a crisis, it becomes less willing to hold capital and so the excess return on capital increases and wealthy households participating in the capital market growth their wealth much quickly than poorer agents. This generates the feature of the data that wealthier households have a more positive exposure to the risk premium during recessions. The magnitude of the crisis excess returns depends upon the severity of the participation constraint and the inequality in the household sector. If household wealth is equally distributed and participation costs are high, then the household sector cannot act as a "backstop" in the capital market during crises and so the scaling force in the model is very strong. If the household wealth is very unequally distributed, then the wealthiest agents participate in the asset market and act as the "back-stop" so the scaling force is weak. In this sense, it is the interaction between inequality and the endogenous general equilibrium price dynamics that determine the extent to which wealthy households can build wealth more quickly.

In Section 5, we calibrate our model to match the average risk free rate, leverage ratio, and portfolio distribution. We perform two non-targeted comparisons to the data to sense-check our model performance. First, the model approximately matches the average equity returns, and the equity risk premium observed in the US data, despite these moments not being explicitly targeted in the calibration. Second, we compare to estimates of heterogeneous business cycle exposure. Using the data from 1976 to 2023 we run local projections that regress the change in wealth shares on the risk premium conditional and unconditional on being in a recession. We find that wealthier households and financial intermediaries have a more positive exposure to the equity risk premium. We also show that this positive exposure is higher during recessions. The projection analysis reveals that relative to the average household, a 10% increase in the equity risk premium *increases* the wealth of affluent agents by 1.0%, and *decreases* the wealth of poor agents by 0.3%. Conditioned on recessions, this gap widens to 1.3% for the wealthy and 1.5% for the poor households, respectively. We interpret this as evidence that poorer agents have less access to higher return assets in general and that the constraints restricting participation bind more during recessions. Local projections using simulated data from our model lead to similar patterns, although poorer agents in our model are more negatively exposed to risk premia than in the data.

Our calibrated model captures a large fraction of the change in wealth shares since 1980. To see this, we simulate the model starting from a household wealth distribution resembling the data and show that it generates wealth distribution evolution consistent with the evolution of empirical distribution. The top 1% wealth share increases from approximately 25% to approximately 35% in the both the data and the model. It is worthy of note that a minimal departure from the literature in the form of participation constraints on households generates a lot of action and matches the empirical moments, indicating the success of our calibrated model.

Finally, we use our model to study the impact of financial macro-prudential policy on inequality. We show that imposing a tighter leverage constraint on the financial sector has two effects: a) it slows down the recovery of financial sector relative to the benchmark with no leverage constraint, and b) increases the risk premium in equilibrium since households have to hold the risky asset and they demand a participation premium due to their constraints. Since wealthier households are more able afford to pay the participation cost than the poorer households, they buy more of the risky assets, earning the high risk premium and widening the inequality gap. Thus, the regulator faces a trade-off between stabilization policy in the form of leverage constraint and wealth inequality.

From a technical point of view, we solve our model by using deep learning tools to train an Economic Model Informed Neural Network (EMINN). General equilibrium for our economy can be characterized by a collection of blocks: (1) a high, but finite dimensional PDE capturing agent *optimization*, (2) a law of motion for the *distribution* of wealth shares and other aggregate state variables, and (3) a set of conditions that ensure the price processes are *consistent* with equilibrium. We develop a new solution approach that can handle complexity in all three blocks. We use neural networks to approximate derivatives of the value function and the price volatility of long-term assets. We then use stochastic gradient descent to train the neural network to minimize the error in the "master" equations that characterize equilibrium for the system. Our approach connects and expands the algorithms developed in Gu, Laurière, Merkel and Payne (2023) and Gopalakrishna (2021). We exploit our continuous time formulation to construct an algorithm that imposes portfolio choice and market clearing explicitly in the master equations. This allows us to circumvent the problems that have occurred in other deep learning papers trying to solve models with portfolio choice. We test our solution approach by solving a collection of canonical macro-finance models that have finite difference solutions.

We believe our algorithm is the first method than can satisfactorily find a global solution to models with non-trivial optimization, distribution evolution, and equilibrium blocks, without having to resort to low-dimensional approximations of the wealth distribution. Other macro-finance models make assumptions to ensure that at least one of these blocks has a closed form solution. To understand this, it is instructive to compare to some canonical models. First, for a representative agent model, the distribution block 2 is not applicable because there is no agent heterogeneity and equilibrium block 3 is less complicated because the goods market condition becomes much simpler. Second, for the continuous time version of Krusell and Smith (1998) discussed in Gu et al. (2023), we have a distribution of agents so distribution block 2 is non-trivial. However, this model only has short-term assets, which leads to closed form expressions for prices in terms of the distribution. So, the equilibrium block 3 is trivial to satisfy. Third, for models such as Basak and Cuoco (1998) and Brunnermeier and Sannikov (2014) discussed in Gopalakrishna (2021), the HJBE can be solved in closed form. This means that agent optimization block 1 can be solved analytically and substituted into the rest of the equations.

Literature Review: We are part of an active literature studying how asset pricing can impact inequality (recent examples include Gomez et al. (2016), Cioffi (2021), Gomez and Gouin-Bonenfant (2024), Fagereng, Gomez, Gouin-Bonenfant, Holm, Moll and Natvik (2022), Basak and Chabakauri (2023), Fernández-Villaverde and Levintal (2024), Irie (2024) amongst many others). Our contribution is to introduce endogenous capital market participation and endogenous price volatility into a heterogeneous agent macroeconomic model.

Our solution approach is part of a growing computational economics literature using deep learning techniques to solve economic models and overcome the limitations of the traditional solution techniques (e.g. Azinovic, Gaegauf and Scheidegger (2022), Han, Yang and E (2021), Maliar, Maliar and Winant (2021), Kahou, Fernández-Villaverde, Perla and Sood (2021), Bretscher, Fernández-Villaverde and Scheidegger (2022), Fernández-Villaverde, Marbet, Nuño and Rachedi (2023), Han, Jentzen and E (2018), Huang (2022), Duarte (2018), Gopalakrishna (2021), Fernandez-Villaverde, Nuno, Sorg-Langhans and Vogler (2020), Sauzet (2021), Gu et al. (2023)). Very few deep learning literature have solved models with long-term asset pricing and complicated portfolio choice, as in our model. Fernández-Villaverde, Hurtado and Nuno (2023) and Huang (2023) solve an extension of Krusell and Smith (1998) with portfolio choice between short-term assets with different risks. Azinovic and Žemlička (2023) solves a general equilibrium model with long-term assets in discrete time by encoding equilibrium conditions and financial constraints into neural network layers. Azinovic, Cole and Kubler (2023) employ low-dimensional approximation of the wealth distribution, following Kubler and Scheidegger (2018), and analyze long-term asset prices in the presence of aggregate and idiosyncratic risk. The difficulty involved with pricing long-term assets with heterogeneous agents is that the equilibrium allocation and individual choices must be determined together, as also pointed out by Guvenen (2009). We demonstrate that in continuous time, on the wealth share space, equilibrium objects can be determined through a unified framework simultaneously. The main contribution of this paper to the deep learning literature is to show how we can globally solve general macro-finance problems without having to resort to low-dimensional approximations of the wealth distribution.

The rest of this paper is structured as follows. Section 2 outlines our economic model. Section 3 introduces our numerical algorithm. Section 4 explores the different mechanisms connecting inequality and asset pricing. Section 5 presents results from our empirical analysis and from our calibrated model.

2 Economic Model

In this section, we outline the economic model we use throughout the paper. We study a continuous time, heterogeneous agent real business cycle macroeconomy where households face asset market participant constraints and the financial sector has equity raising frictions.

2.1 Environment

Setting: The model is in continuous time with infinite horizon. There is a perishable consumption good and a durable capital stock. The economy has the following assets: short-term risk free bonds and capital stock.

Production: The production technology in the economy creates consumption goods according to $Y_t = e^{z_t} K_t$ where K_t is the capital used at time t and z_t is aggregate productivity. Aggregate productivity evolves according to:

$$dz_t = \zeta(\bar{z} - z_t)dt + \sigma_z dW_t,$$

where W_t denotes an aggregate Brownian motion process. We let $\mathcal{F}_t$ denote the filtration generated by W_t. Any agent can use goods to create capital stock, k_t, but all face adjustment costs so that their capital evolves according to:

$$dk_t = (\phi(\iota_t)k_t - \delta k_t)dt$$

where $\Phi(\iota)k := (\iota - \phi(\iota_t))k$ represents the resources used from investment rate ι_t and δ is a depreciation rate.

Agents: The economy is populated by a large, finite collection of infinitely lived price taking agents, indexed by $i \in \mathcal{I} = \{i : 0 \leq i \leq I\}$.[1] We interpret agents $1 \leq i \leq I-1$ as "households" (h) and agent $i = I$ as a financial "expert" (e). Each household i has discount rate ρ and gets flow utility $u(c_{i,t}) = c_{i,t}^{1-\gamma}/(1-\gamma)$. Households receive idiosyncratic death shocks at rate λ_h. We let $\rho_h := \rho + \lambda_h$. Following a death, a new agent enters with a fraction $1-\beta$ of the dying agent's wealth while the remaining β is distributed evenly across the population. The economy also contains a financial expert with the same utility as the household and discount rate $\rho_e > \rho_h$.

Assets, markets, and financial frictions: Each period, there are competitive markets for goods and capital trading. We use goods as the numeraire. We let q_t denote the price of capital and r_t denote the interest rate on bonds. We guess and verify that the capital price process satisfies:

$$dq_t = \mu_{q,t}q_t dt + \sigma_{q,t}q_t dW_{i,t}$$

where $\mu_{q,t}, \sigma_{q,t}$ are the geometric drift and volatility of the q_t respectively. Asset markets are incomplete so households cannot insure their idiosyncratic labor shocks.

Financial frictions: Households face a capital market participation constraint, which we model as a "soft" constraint by imposing the utility penalty function[2]:

$$\Psi_h(k_{i,t}, a_{i,t}, \eta_{i,t}) = \psi_{h,t}\Xi_{i,t}a_{i,t}, \quad \text{where} \quad \psi_{h,t} = \psi_h(k_{i,t}, a_{i,t}, \eta_{i,t}) = \frac{\bar{\psi}\sigma^2}{2\eta_{i,t}}\left(\frac{k_{i,t}}{a_{i,t}}\right)^2$$

where $\bar{\psi}$ is the severity of the constraint, $a_{i,t}$ is household i's wealth, $k_{i,t}$ is household i's

[1] We interpret the economy as an approximation to a competitive equilibrium with a continuum of price-taking agents. In Gu et al. (2023) we compare this to other ways of approximating such equilibria.

[2] We model this as a utility penalty since the cost is small as a function of aggregate output in the economy.

capital holdings, $\eta_{i,t}$ is agent i's wealth share in the economy, and $\Xi_{i,t}$ is the equilibrium SDF in the economy for an agent with wealth $a_{i,t}$. We can observe that participation penalty is softened as agents hold more wealth. The expert cannot raise equity but does not face a capital market participation constraint so $\Psi_e(\cdot) = 0$.

2.2 Equilibrium

Agent problems: Given their belief about the price processes, $(\hat{r}, \hat{q})$, each agent i chooses consumption $c_{i,t}$ and bond holding $b_{i,t}$ to solve problem (2.1) below:

$$\begin{aligned} \max_{c_i, b_i} & \left\{ \mathbb{E}_0 \left[\int_0^\infty e^{-\rho_i t} \left(u(c_{i,t}) - \Psi_i(a_{i,t} - b_{i,t}, a_{i,t}, \eta_{i,t}) \right) dt \right] \right\} \\ s.t. \quad & da_{i,t} = \hat{r}_{i,t} dt + (a_{i,t} - b_{i,t}) d\hat{R}_{k,t} - c_{i,t} dt + \tau_t dt \end{aligned} \tag{2.1}$$

where $\tau = \beta \lambda A_t$ is the transfer from dying agent, Ψ_i is the participation constraint faced by agent i and $d\hat{R}_{k,t}$ is the agent's belief about the return on holding capital:

$$\begin{aligned} d\hat{R}_{k,t} &:= \frac{z - \iota_t k_t}{\hat{q}_t k_t} + \frac{d(\hat{q}_t k_t)}{\hat{q}_t k_t} \\ &= \left(\frac{z}{\hat{q}_t k_t} - \frac{\iota_t}{\hat{q}_t} + (\phi(\iota_t) - \delta) + \hat{\mu}_{q,t} \right) dt + \hat{\sigma}_{q,t} dW_t \\ &=: \hat{r}_{k,t} dt + \hat{\sigma}_{q,t} dW_t \end{aligned}$$

Expanding out the price processes allows the wealth evolution to be written as:

$$da_{i,t} = \mu_{a,t} dt + \sigma_{a,t} dW_t, \quad \text{where} \quad \begin{aligned} \mu_{a,i,t} &:= b_{i,t} \hat{r}_t + (a_{i,t} - b_{i,t}) \hat{r}_{k,t} - c_{i,t} \\ \sigma_{a,i,t} &:= (a_{i,t} - b_{i,t}) \hat{\sigma}_{q,t} \end{aligned} \tag{2.2}$$

Distribution: The uninsurable idiosyncratic shocks and idiosyncratic differences in agent portfolio constraints potentially generate a non-degenerate distribution of agent wealth positions across the economy. We let $g_t = \{a_{i,t} : i \in \mathcal{I}\}$ denote the positions of agents across the economy at time t for a given filtration $\mathcal{F}_t$, where $\mathcal{F}_t$ is generated by aggregate shock process $\{W_t\}_{t \geq 0}$. With some abuse of terminology, we refer to g_t as the distribution across the economy.

(Sequential) Equilibrium Definition: Given an initial distribution g_0, an equilibrium for this economy consists a collection of $\mathcal{F}$-adapted stochastic processes $\{c_t^i, b_t^i, g_t, r_t, q_t, K_t, y_t : t \geq 0, i \in \mathcal{I}\}$ such that:

1. Agent decision processes solve problems (2.1), given their belief about the price process $(\hat{r}, \hat{q})$;

2. At each time t, equilibrium prices (r_t, q_t) solve the market clearing conditions: (i) goods market $\sum_i c_{i,t} + \sum_i \Phi(\iota_{i,t}) k_{i,t} = y$, (ii) bond market $\sum_i b_{i,t} = 0$, and (iii) capital market $\sum_i (a_{i,t} - b_{i,t}) = q_t K_t$;

3. Agent beliefs about the price process are consistent with the optimal behaviour of all agents in the sense that $(\hat{r}, \hat{q}) = (r, q)$.

2.3 Recursive Characterization of Equilibrium

In this section, we characterize the equilibrium recursively. We work through the problem in detail in Appendix A and summarize the key results here. The "natural" state variables for the equilibrium are:

$$(z, K, g = \{a_i\}_{1 \leq i \leq I}).$$

However, it turns out that the recursive characterization in agent wealth levels leads to a complicated fixed point problem that is hard for the Neural Network to train (we discuss in detail in Section 3.3 after we introduce the algorithm.). Instead, it will be convenient to characterize the equilibrium in terms of wealth shares. Let $A := \sum_{j \geq 1} a_j$ denote total wealth in the economy. Let $\eta_i := a_i / A$ denote the share of wealth held by agent i. Then, the aggregate state of the economy can be written in terms of wealth shares (with some abuse of the g notation) as

$$(z, K, g = \{\eta_j\}_{1 \leq i \leq I}).$$

State variables and beliefs: We assume there exists a solution to the equilibrium that is recursive in the aggregate state variables which we denote by $(\cdot)$. This means that the states that appear in the household decision problem are $(a_i, \cdot)$. In this case, beliefs about the price process can be characterized by beliefs about how the distribution and aggregate capital stock evolves since prices are all implicitly functions of the aggregate state variables. Formally, an agent's beliefs about the evolution of the distribution are characterized by their beliefs about the drift and covariance of other agents wealth and the drift of capital stock, $\{\hat{\mu}_{a_j}(\cdot), \hat{\sigma}_{a_j}(\cdot), \hat{\mu}_K(\cdot) : j \neq i\}$, which imply beliefs about prices through the pricing functions $(r(\cdot), q(\cdot))$. We let $V_i(a_i, \cdot)$ denote household i's value function. For notational convenience, we drop the explicit dependence on $(\cdot)$ where possible.

We define the marginal value of wealth and the partial derivatives of the marginal value of wealth (the so called "stochastic discount factors") by:

$$\xi_i := \frac{\partial V_i}{\partial a_i}, \qquad \partial_a \xi_i := \frac{\partial \xi_i}{\partial a_i} = \frac{\partial^2 V_i}{\partial a_i^2}, \qquad \partial_{a_j} \xi_i := \frac{\partial \xi_i}{\partial a_j} = \frac{\partial^2 V_i}{\partial a_i a_j}$$

Once equilibrium is imposed, all the endogenous objects in the economy must be functions of $(z, K, \{\eta_j\}_{j\geq 1})$. Thus, we can use Ito's Lemma to express the drift and volatility of ξ_i in terms of derivatives of ξ_i with respect to $(z, K, \{\eta_j\}_{j\geq 1})$ in equilibrium:

$$\begin{aligned}
\xi_i \mu_{\xi_i} = & \frac{\partial \xi_i}{\partial z}\mu_z + \frac{\partial \xi_i}{\partial K}\mu_K + \sum_j \frac{\partial \xi_i}{\partial \eta_j}\eta_j \mu_{\eta_j,t} + \sum_j \frac{\partial^2 \xi_i}{\partial z \partial \eta_j}\eta_j \sigma_{\eta_j,t}\sigma_z \\
& + \frac{1}{2}\frac{\partial^2 \xi_i}{\partial z^2}\sigma_z^2 + \frac{1}{2}\sum_{j,j'} \frac{\partial^2 \xi_i^2}{\partial \eta_j \partial \eta_{j'}}\eta_j \eta_{j'} \sigma_{\eta_j,t}\sigma_{\eta_{j'},t} && (2.3) \\
\xi_i \sigma_{\xi_i} = & \frac{\partial \xi_i}{\partial z}\sigma_z + \sum_j \frac{\partial \xi_i}{\partial \eta_j}\eta_j \sigma_{\eta_j,t}. && (2.4)
\end{aligned}$$

It is helpful to characterize the equilibrium in terms of three blocks that impose belief consistency and market clearing conditions, where possible.

1. Agent optimization block: Applying the Envelope Theorem to the HJBE (equation (A.1) in the Appendix), imposing belief consistency, and using Ito's Lemma to collect terms leads to the continuous time Euler equation (the so called "master equation" for the economy) for ξ_i. That is, given prices $(r, r_k, q, \mu_q, \sigma_q)$, agent optimization implies that $(\xi_i, c_i, b_i, \iota_i)$ satisfy:

$$\begin{aligned}
0 = & -\rho_i + r + \mu_{\xi_i,t} && (2.5) \\
u'(c_i) = & \xi_i \\
r - r_k = & \sigma_{\xi_i}\sigma_q + \frac{\partial \psi_i}{\partial b_i}\Big|_{a_i = \eta_i q} && (2.6) \\
\iota_i = & (\phi')^{-1}\left(q^{-1}\right)
\end{aligned}$$

where μ_{ξ_i} satisfies (2.3) and σ_{ξ_i} satisfies (2.4).

2. State evolution block: Given prices $(r_t, r_k, q, \mu_q, \sigma_q)$ and agent optimization (ξ, c, b, ι), we can use Ito's Lemma to get the law of motion for each wealth share $\eta_{j,t} = a_{j,t}/(q_t K_t)$:

$$\begin{aligned} \mu_{\eta_j,t} = {} & r_{k,t} - \mu_{q,t} - \mu_{K,t} + \frac{b_{j,t}}{\eta_{j,t} q_t K_t}(r_t - r_{k,t}) - \frac{(u')^{-1}(\xi_{j,t})}{\eta_{j,t} q_t K_t} \\ & + \frac{b_{j,t}}{\eta_{j,t} q_t K_t}\sigma_{q,t}^2 + \frac{\beta\lambda(1-\eta_{j,t})}{\eta_{j,t}} \end{aligned} \tag{2.7}$$

$$\sigma_{\eta_j,t} = -\frac{b_{j,t}}{\eta_{j,t} q_t K_t}\sigma_{q,t} \tag{2.8}$$

The evolution of K_t satisfies:

$$dK_t = (\phi(\iota_t)K_t - \delta K_t)dt.$$

3. Equilibrium block: The market clearing conditions now become:

$$\sum_i c_i + \Phi(\iota)K = y \qquad \sum_i b_i = 0 \qquad \sum_i (\eta_i A - b_i) = K$$

where the aggregate household wealth satisfies $A := \sum_{j\geq 1} a_j = qK$ and so the capital market clearing condition simply becomes $\sum_i \eta_i = 1$. The risk free rate can only be implicitly expressed in terms of the state variables through its dependence on the stochastic processes for ξ and q (using the agent first order conditions):

$$r = r_k + \sigma_{\xi_i}\sigma_q + \frac{\partial \psi_i}{\partial b_i}$$

The price of capital is even more difficult to handle because capital is a long-lived asset for which the price can only be implicitly expressed in terms of the state variables using Itô's Lemma:

$$\begin{aligned} q\mu_{q,t} = {} & \sum_j \frac{\partial q}{\partial \eta_j}\eta_j \mu_{\eta_j,t} + \frac{\partial q}{\partial z}\mu_{z,t} + \frac{\partial q}{\partial K}\mu_{K,t} + \sum_j \frac{\partial^2 \xi_i}{\partial z \partial \eta_j}\eta_j \sigma_{\eta_j,t}\sigma_z \\ & + \frac{1}{2}\sum_{j,j'} \frac{\partial^2 q}{\partial \eta_j \partial \eta_{j'}}\eta_j \eta_{j'}\sigma_{\eta_j,t}\sigma_{\eta_{j'},t} + \frac{1}{2}\frac{\partial^2 q}{\partial z^2}\sigma_z^2 \\ q\sigma_{q,t} = {} & \sum_j \frac{\partial q}{\partial \eta_j}\eta_j \sigma_{\eta_j,t} + \frac{\partial q}{\partial z}\sigma_{z,t}. \end{aligned}$$

These expressions for $\mu_{q,t}$ and $\sigma_{q,t}$ are what makes the law of motion for capital "consistent" with the process that we posited in the environment and so are often referred to as the price consistency differential equations.

2.3.1 Comparison to Other Models

Why is this system of equations difficult to solve in our model? Because, unlike in most models, all three blocks are non-trivial. To our knowledge, no other paper is able to satisfactorily solve this system globally without imposing assumptions to make one of the blocks trivial. To understand why this is the case, it is instructive to compare the model to other macro-finance models.

(i). For a representative agent model, block 2 is not applicable because there is no distribution and block 3 is less complicated because the goods market condition simply becomes $c + (\iota - \phi(\iota))K = y$, which can be substituted into equations in block 1. In this case, the model can be simplified to a differential equation for q. For heterogeneous agent models, following Krusell and Smith (1998), other papers approximate the distribution by a low dimensional collection of moments and do not need to work the agent distribution.

(ii). For the continuous time version of Krusell and Smith (1998) discussed in Gu et al. (2023), we have a distribution of agents so block 2 is non-trivial. However, this model has no long-term assets and closed form expressions for all prices in term of the distribution. So, block 3 is can be trivially satisfied and we can combine all equilibrium conditions into one master equation.

(iii). For models such as Basak and Cuoco (1998) and Brunnermeier and Sannikov (2014) discussed in Gopalakrishna (2021), the HJBE can be solved in closed form. This means that block 1 can be solved analytically and substituted into the block 3.

3 Algorithm

In this section, we outline our algorithm for solving the model. A "direct" application of deep learning would be to parameterize the equilibrium objects and then train the neural networks to minimize a loss function that combines condensed set of the general equilibrium equations described in subsubsection A.2. Although this approach should work in principle, many researchers have found it very difficult to implement in practice. Instead, we simplify the equations, choose a parsimonious parametrization and break the problem up into "linear" blocks.

3.1 Neural network parametrization and loss function

Let $\boldsymbol{X} := (z, K, (\eta_i)_{i\leq I}) \in \mathcal{X}$ denote the state vector in the economy and let $\mathcal{X}$ denote the state space. We use neural networks to approximate sufficiently many variables to allow us to calculate the remaining variables using matrix algebra. For our general model, this requires approximating: the equilibrium consumption-to-wealth ratio policy for the first agent in the economy with each type of financial constraint, $\{\omega_h(\boldsymbol{X})\}_{h\in\mathcal{H}}$, and the price volatility, $\sigma_q(\boldsymbol{X})$.We denote the approximations by:

$$
\begin{aligned}
\hat{\omega}_h &: \mathcal{X} \to \mathbb{R},\ (\boldsymbol{X}, \Theta_{\omega_h}) \mapsto \hat{\omega}_h(\boldsymbol{X}; \Theta_{\omega_h}), \quad \forall h \in \mathcal{H} \\
\hat{\sigma}_q &: \mathcal{X} \to \mathbb{R},\ (\boldsymbol{X}, \Theta_q) \mapsto \hat{\sigma}_q(\boldsymbol{X}; \Theta_q)
\end{aligned}
$$

where $\{\Theta_{\omega_h}\}_{h\in\mathcal{H}}$, Θ_q are the parameters in the neural network approximations of $\hat{\omega}_h$ and $\hat{\sigma}_q$respectively.

We can recover the approximate consumption policy function for each agent i with constraint h from $\hat{\omega}_h$ because policies for all agents with a particular financial constraint are symmetric. That is, let $n(h)$ denote the position of the first agent economy in the economy with constraint h. Then $\hat{\omega}_i(\boldsymbol{X})$ for any i with constraint h can be recovered by swapping the positions of the states for $n(h)$ and i:

$$
\begin{aligned}
\hat{\omega}_i(\boldsymbol{X}) = \hat{\omega}_{H(i)}\Big(& z, K, (\ldots, (\eta_{n(h)}) = (\eta_i), \\
& \ldots, (\eta_i) = (\eta_{n(h)}), \ldots)\Big).
\end{aligned} \tag{3.1}
$$

At state $\boldsymbol{X}$, the error (or "loss") in the Neural network approximations is given by the following equations for $h \in H$:

$$
\begin{aligned}
\mathcal{L}_{\omega_j}(\boldsymbol{X}) = (r-\rho)\hat{\xi}_i &+ \frac{\partial \hat{\xi}_i}{\partial z}\mu_z + \frac{\partial \hat{\xi}_i}{\partial K}(\phi((\phi')^{-1}(q^{-1}))K_t - \delta K_t) \\
&+ \sum_j \frac{\partial \hat{\xi}_i}{\partial \eta_j}\eta_j \mu_{\eta_j,t} + \sum_j \frac{\partial^2 \hat{\xi}_i}{\partial z \partial \eta_j}\eta_j \sigma_{\eta_j,t}\sigma_z + \frac{1}{2}\frac{\partial^2 \hat{\xi}_i}{\partial z^2}\sigma_z^2 \\
&+ \frac{1}{2}\sum_{j,j'} \frac{\partial^2 \hat{\xi}_i^2}{\partial \eta_j \partial \eta_{j'}}\eta_j \eta_{j'}\sigma_{\eta_j,t}\sigma_{\eta_{j'},t}
\end{aligned} \tag{3.2}
$$

$$
\mathcal{L}_{\sigma}(\boldsymbol{X}) = \ -q\hat{\sigma}_q + \sum_j \frac{\partial q}{\partial \eta_j}\eta_j \sigma_{\eta_j} + \frac{\partial q}{\partial z}\sigma_z \tag{3.3}
$$

where $\hat{\xi}_j = \hat{\xi}(\hat{\omega}_j(\boldsymbol{X}))$ for all $j \in J$, $\hat{\psi}_j = \hat{\psi}_j(\boldsymbol{X})$ for all $j \in J$, $\hat{\sigma}_q = \hat{\sigma}_q(\boldsymbol{X})$, and the other variables are evaluated by solving the relevant equations in section B.

Discussion: which neural network objects need to be approximated. We approximate variables to ensure that the equations are linear given the neural network approximated variables. This means that we always need to approximate ω_j (or ξ_j) because the Euler equation is non-linear. If there are no financial constraints, $\psi_j = 0$ for all $j \in J$, then we do not need to make any additional approximations because the risk allocation equation can be solved using matrix inversion.

3.2 Algorithm

We outline the algorithm in Algorithm 1 below. Given the current guesses of the neural networks, we solve for equilibrium using the matrix algebra. The exact set of condensed equations that we solve are provided Appendix B. We then update our guesses for the neural network approximations.

3.3 Imposing Market Clearing in the Sampling

The major difficulty faced by the deep learning macroeconomics literature is that it is necessary to impose market clearing in the sampling. This is partly because trying to impose market clearing in the loss function generates instability. It is also because for asset pricing problems, in particular, sampling schemes that don't impose market clearing often lead the neural network learn a trivial mapping "$q = q$" due to the summation of individual wealth equating to q. To overcome these problems, we restrict the sample space to enforce market clearing.

If we sample in the a space, then we end up needing to restrict a to a subspace that depends upon equilibrium prices. To make this concrete, consider the goods market clearing condition, the capital market clearing condition, and the borrowing constraint:

$$\sum_i c(a_i) + \sum_i \Phi(\iota_{i,t}) k_{i,t} = e^z K, \qquad \sum_i a_i = qK, \qquad a_i \geq \bar{a}$$

If we sample in a space, then we need to draw a values in a way that respects these conditions. This restricts a to an $I-1$ dimensional hyperplane $\mathcal{A}(z, K, q)$ that depends upon z, K, and the equilibrium q.

Restricting a to the equilibrium hyperplane causes a number of problems when we

Algorithm 1 Pseudo Code

1: Initialize neural network objects $\{\hat{\omega}_h\}_{h\leq H}$, and $\hat{\sigma}_q$ with parameters $\{\Theta_{\omega_h}\}_{h\leq H}$, and $\{\Theta_q\}$ respectively.

2: Initialize optimizer.

3: **while** Loss $>$ tolerance **do**

4: Sample N new training points: $\left(\boldsymbol{X}^n = \left(z^n, K^n, (\eta_i)^n_{i\leq N-1}\right)\right)^N_{n=1}$.

5: Calculate equilibrium at each training point $\boldsymbol{X}^n$:

- a. Compute $(\hat{\omega}^n_i)_{i\leq I}$ using equation (3.1) and the current approximation $\{\hat{\omega}_h\}_{h\leq H}$ evaluated at $\boldsymbol{X}^n$.
- b. Compute q^n and $(\xi^n_i)_{i\leq I}$ using $(\hat{\omega}^n_i)_{i\leq I}$.
- c. Solve for $\boldsymbol{\sigma}^n_{\boldsymbol{\eta}}$ and s^n the current approximations for $\{\hat{\omega}_h\}_{h\leq H}$, $\{\partial_b\hat{\psi}_h\}_{h\leq H}$, and $\hat{\sigma}_q$ (and their automatic derivatives).
- d. Solve for portfolio choice $\boldsymbol{\theta}^n$.
- e. Compute μ_η, μ_q, r.

5: Construct loss as:

$$\hat{\mathcal{L}}(\boldsymbol{X}) = \sum_h \frac{1}{N}\sum_n |\hat{\mathcal{L}}_{\omega_h}(\boldsymbol{X}^n)| + \frac{1}{N}\sum_n |\hat{\mathcal{L}}_\sigma(\boldsymbol{X}^n)|$$

where $\hat{\mathcal{L}}_{\omega_h}$ and $\hat{\mathcal{L}}_\sigma$ are defined by (3.2)and (3.3) with ω_h and σ_q replaced by their neural network approximation.

6: Update $\{\Theta_{\omega_h}\}_{1\leq i\leq H}$and $\{\Theta_q\}$ using ADAM optimizer.

7: **end while**

don't have a closed form expression for q and so need neural network approximations for both $\hat{V}$ and $\hat{q}$.[3] First, it is hard to control how frequently the sampled agents hit the borrowing constraint. Second, numerical instability arises because the $\hat{V}$ has another neural network, $\hat{q}$, as an input. This second problem is particularly acute for deep learning based algorithms because there is no easy way to retain the computational graph for q when calculating auto-derivatives for V. To understand this, recall that the loss function depends upon $\hat{V}(\boldsymbol{X};\Theta_V)$ and $\hat{q}(\boldsymbol{X};\Theta_q)$:

$$Loss(\boldsymbol{X}) = \mathcal{F}(\boldsymbol{X}, \hat{V}(\boldsymbol{X};\Theta_V), \hat{q}(\boldsymbol{X};\Theta_q))$$

This means that, in principle, the parameter update step in the stochastic gradient descent algorithm should look like the following:

$$\begin{aligned}\theta_{V,n+1} &= \theta_{V,n} - \alpha_{V,n}\frac{\partial Loss}{\partial \theta_V},\\ \theta_{q,n+1} &= \theta_{q,n} - \alpha_{q,n}\frac{\partial Loss}{\partial \theta_q},\end{aligned}$$

where $\alpha_{V,n}$ and $\alpha_{q,n}$ denote the rate of updating. However, when we impose equilibrium sampling and so need to express $\hat{V}$ as an implicit function of $\hat{q}$, then we need to detach $\hat{q}$ in the θ_V update step and detach $\hat{V}$ in the θ_q update step. So, in practice, the algorithm looks like:

$$\begin{aligned}\theta_{V,n+1} &= \theta_{V,n} - \alpha_{V,n}\frac{\partial Loss}{\partial \hat{V}}\frac{\partial \hat{V}}{\partial \theta_V},\\ \theta_{q,n+1} &= \theta_{q,n} - \alpha_{q,n}\frac{\partial Loss}{\partial \hat{q}}\frac{\partial \hat{q}}{\partial \theta_q},\end{aligned}$$

This "diverted" gradient based updating is very likely to get stuck at local minima, particularly when there is high curvature in the problem. An additional problem is that computing the evolution of the distribution requires the asset returns r_q, r, μ_q, σ_q but at the same time μ_q, σ_q are pinned down by the the consistency conditions, which in turn depend upon the distribution evolution. This creates a fixed point problem that does not have a simple closed form solution for μ_q, σ_q in most heterogeneous agent economic models and so suggests that we need to introduce auxiliary neural networks for μ_q, σ_q. Resolving these issues requires a staggered updating approach similar to

[3] For example, in Gu et al. (2023) we had a closed form expression for the prices and so we did not face these difficulties.

that proposed by Guvenen (2009).

Working in the wealth share space rather than the wealth space resolves these issues. This is because, in the wealth share characterization, the capital market clearing condition is automatically satisfied because of the accounting relation: $\sum_i \eta_i = 1$. This, in turn, means that we are able impose market clearing in the sampling without needing to allow the neural network approximation $\hat{V}$ to take $\hat{q}$ as an input.

3.4 Three Testable Models

We "test" our approach by using our algorithm to characterize the solution to three macro-finance models that can be solved using conventional methods: a complete markets model, Basak and Cuoco (1998), and Brunnermeier and Sannikov (2014). Appendix C studies the comparison in detail. Here we summarize the key results. For all models, we use simple feed-forward neural networks and an ADAM optimizer. The details of the neural network parameters for each model are shown in Table 1.

Model	Num of Layers	Num of Neurons	Learning Rate
"As-if" Complete Model	4	64	0.001
Limited Participation Model	5	64	0.001
BruSan Model	5	32	0.001

Table 1: Neural network parameters for the three testible models

Table 2 summarizes the mean squared error between the conventional solution and the neural network solution. Evidently, the neural network and conventional methods converge to very similar characterizations of equilibrium. We compares plots from the models visually in Appendix C.

Method	Error
Complete markets	1.0×10^{-5}
Basak and Cuoco (1998)	4.9×10^{-4}
Brunnermeier and Sannikov (2014)	7.0×10^{-5}

Table 2: Summary of the algorithm performance and computational speed. Error calculates the difference between solution by neural network and finite difference. All errors are in absolute value (L1).

3.5 Convergence For Our Full Quantitative Model

We solve the quantitative model by training the the deep neural networks $\left(\{\hat{\omega}_h\}_{h=1}^{I}, \hat{\sigma}_q\right)$. Each neural network is fully-connected feed-forward type, and has 4 hidden layers and 32 neurons in each layer. We train using an ADAM optimizer with a learning rate of 0.0005 for 1400 iterations. Figure 1 presents the L-1 loss from the quantitative model over iterations. The loss decreases over time although not monotonically due to the stochastic nature of learning process. The HJB loss is higher than the consistency loss due since the HJB equations involve Euler equations which are complicated since they embed the market clearing conditions. After 10,000 iterations, the total L-1 loss is 0.018. The corresponding L-2 loss is 1.4×10^{-4}.[4]

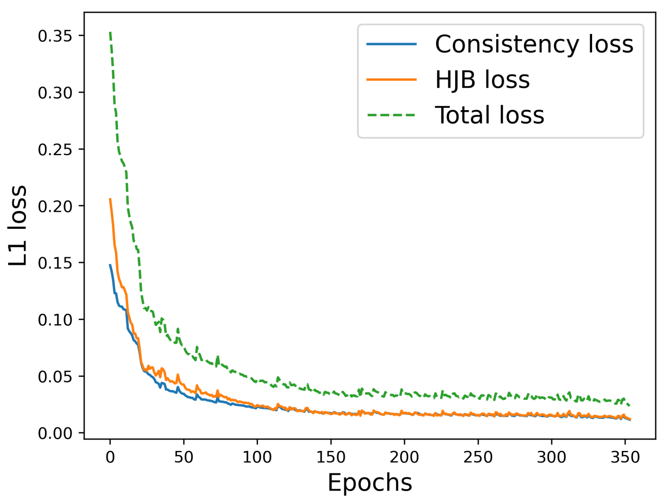

Figure 1: The L-1 loss from the quantitative model over iterations. The neural network architecture is 4 hidden layers with 32 neurons in each layer trained using an ADAM optimizer.

4 Understanding Inequality and Asset Price Dynamics

An important feature of the model is the ability to characterize the general equilibrium relationship between participation constraints, inequality, and asset price dynamics. In this section, we explore these connections.

In our model, the difference between the drift of the wealth share of any two

[4]Figure 1 only shows for 300 epochs since we ignore epochs whenever the loss is larger than the running mimimum.

households i and j is given by:

$$\begin{aligned}\mu_{\eta_j,t} - \mu_{\eta_i,t} = (\theta_{j,t} - \theta_{i,t})(r_{k,t} - r_t) - (\theta_{j,t} - \theta_{i,t})\sigma_{q,t}^2 \\ + \beta\lambda\left(\frac{1}{\eta_{j,t}} - \frac{1}{\eta_{i,t}}\right) - (\omega_j - \omega_i)\end{aligned} \tag{4.1}$$

where $\theta_{i,t} := k_{i,t}/a_{i,t}$ is the fraction of wealth household i allocates to capital and the portfolio choice must satisfy the FOC (with the function form for ψ substituted in):

$$\sigma_{\xi,i}(\theta_i) - \frac{\sigma}{\sigma_q}\frac{\bar{\psi}}{\eta_{i,t}}\theta_i = \frac{r_k - r_f}{\sigma_q}$$

The first term in (4.1) captures how participation constraints and risk aversion impact the excess return that different agents can earn. When $\eta_{j,t} > \eta_{i,t}$ is higher, then agent j holds more wealth in capital and so gains wealth share compared to the poorer agents who are unwilling to pay the cost to participate in the capital market. This has sometimes been referred to as the “scaling” effect in the literature—wealthier agents have access to better investment opportunities and so gain wealth more quickly. The second term in (4.1) captures the impact of risk exposure on the average wealth drift. Agents holding more capital are also more exposed to aggregate risk in the economy. This is additional impact of scaling up into risky investment opportunities that is not present in macroeconomic inequality models without aggregate risk. The third term in (4.1) captures the impact of the death rate in the economy. This is the main force that stabilizes the wealth distribution in economy. Other model have attributed this many possible features (e.g. new entrants with better skills, idiosyncratic risk, ...). We have little to say about it in our model and so simply allocate it to a death rate. The final term in (4.1) captures how a lower marginal propensity to consume out of wealth, $\omega_j < \omega_i$, leads to greater wealth accumulation.

How does the financial sector impact these dynamics? The expert does not have a capital market participation constraint and so

$$\sigma_{\xi,e}(\theta_e) = \frac{r_k - r_f}{\sigma_q}$$

Imposing this in equation (4.1) gives the alternative expression:

$$\mu_{\eta_j,t} - \mu_{\eta_i,t} = (\theta_{j,t} - \theta_{i,t})\sigma_{\xi,e}(\theta_e)\sigma_{q,t} - (\theta_{j,t} - \theta_{i,t})\sigma_{q,t}^2 - (\omega_j - \omega_i) \tag{4.2}$$

So, the portfolio choice of the expert and resulting capital price volatility end up dictating to what extent capital market participation constraints impact inequality.

We use our model to characterize how these different forces play out. We start by studying the approximate solution to the log-utility version of the model, which has some analytical tractability. We then consider the more general CRRA preferences.

4.1 Log Utility and Fixed Consumption-to-Wealth Ratios

We start by considering the following log-case of the environment, which is analytically tractable. We impose that all agents have log utility $u(c) = \log(c)$, and set $\beta = 0$ for transfers. In Appendix E, we show that in this environment the consumption-to-wealth ratios are given by:

$$\frac{c_{e,t}}{a_{e,t}} = \rho_e, \qquad \qquad \frac{c_{i,t}}{a_{i,t}} = \rho_i + \mathcal{O}(\sigma^2)$$

and so, for analytical tractability, we work with the approximation that the consumption wealth ratio constant for all agents. Under this approximation, the portfolio weights on capital are:

$$\theta_{i,t} = \frac{r_{q,t} - r_{f,t}}{\sigma_{q,t}^2 + \bar{\psi}\sigma^2/\eta_{i,t}}, \quad i \in \{1, \dots, I-1\}, \qquad \qquad \theta_{e,t} = \frac{r_{q,t} - r_{f,t}}{\sigma_{q,t}^2}.$$

This means that $\sigma_{\xi,e} = -\theta_{e,t}\sigma_{q,t}$ and so the relative drift in wealth shares equation (4.2) is simply given by:

$$\mu_{\eta_j,t} - \mu_{\eta_i,t} = (\theta_{j,t} - \theta_{i,t})(\theta_{e,t} - 1)\sigma_{q,t}^2$$

We show in Appendix E that the experts are levered in equilibrium so $\theta_e \geq 1$ and so the experts widen the wealth share gaps between agents by amplifying the risk premium.

How does asset pricing impact household inequality? We start by investigating how the severity of the participation constraint, $\bar{\psi}$ and expert decision making impact asset pricing and ultimately inequality dynamics.

Theorem 1 characterizes the equilibrium in the limiting cases as $\bar{\psi} \to 0$ and $\bar{\psi} \to \infty$. The first case looks similar to Brunnermeier and Sannikov (2014) in the sense that households can frictionlessly purchase capital from the experts and so as act as the "back-stop" during financial crises. The second case looks similar to Basak and Cuoco

(1998) in the sense that households can never purchase capital and so the expert sector has to be its own "back-stop" during the crisis. In both cases, the household sector aggregates because all households make the same portfolio choice.

Theorem 1. *Given the wealth distribution $\{\eta_i\}, i \in \{1, ..., N\}$, we have that:*

(i) In the limit as $\bar{\psi} \to 0$, all households have the same portfolio weight on capital $\theta_i = (r_{k,t} - r_{f,t})/\sigma_q^2$.

(ii) In the limit as $\bar{\psi} \to \infty$, all households hold no capital $\theta_i = 0$.

In both limits, inequality is unchanging, $\mu_{\eta_j,t} - \mu_{\eta_i,t} = 0$ for all $i, j < I$, and only the aggregate wealth share of the household sector, $\sum_{i=1}^{I-1} \eta_{i,t}$ matters for asset pricing.

Proof. See Appendix E. □

A strength of our solution approach is that we can characterize the equilibrium for $0 < \bar{\psi} < \infty$, where household have heterogeneous portfolio decisions. In Figure 2, we plot the numerical solution for three intermediate participation constraints $\bar{\psi} \in \{0.2, 1.0, 2.0\}$. Unsurprising, we see that a higher participation cost means that households hold less capital and experts hold more capital, even when expert wealth drops low. This means that the risk premium in the economy must be higher so that experts are compensated for the additional risk they bear. In particular, we see that higher participation constraints lead to much higher risk premia during crises.

Figure 3 plots the resulting relationship between the drift in inequality and the expert wealth share for different levels of the participation constraint. Evidently, the impact of a higher participation constraint is non-monotonic. When the expert wealth share is small, having a higher participation constraint leads to faster growth in household inequality because the risk premium is high but only the wealthy agents can access the capital market. When the expert wealth share is large, a higher participation constraint can lead to slower growth in household inequality because even the wealthy households are unwilling to compete with expert to hold capital.

How does inequality impact asset prices? We now consider the feedback from household inequality back into asset prices. Aggregate capital demand is given by:

$$\sum_{i=1}^{I-1} \theta_{i,t}\eta_{i,t}A_t + \theta_{e,t}\eta_{e,t}A_t = \left(\sum_{i=1}^{I-1} \frac{\eta_{i,t}^2}{\bar{\psi}\sigma^2 + \sigma_{q,t}^2\eta_i} + \frac{\eta_{I,t}}{\sigma_{q,t}^2}\right)(r_{k,t} - r_{f,t})q_t K_t$$

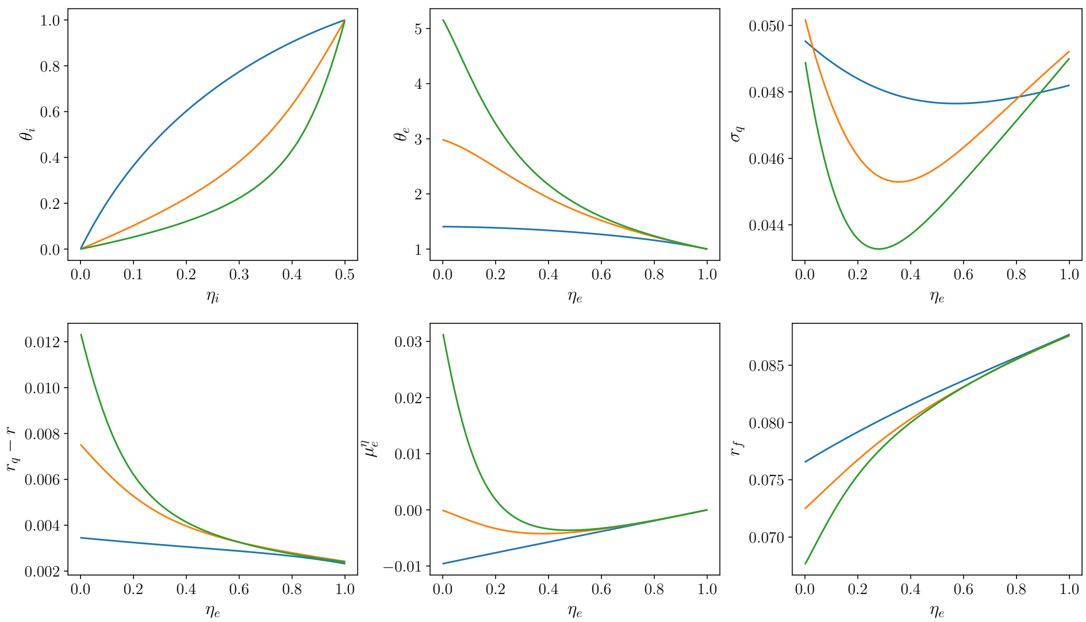

Figure 2: Equilibrium functions for different participation constraints for geometric TFP process. The blue plot has $\bar{\psi} = 0.2$. The orange plot has $\bar{\psi} = 1.0$. The green plot has $\bar{\psi} = 2.0$. Wealth distribution within the household sector is set to be equal. $\rho_e = 0.04, \rho_h = 0.03, \mu = 0.02, \sigma = 0.05$ (the same below for all figures in this section).

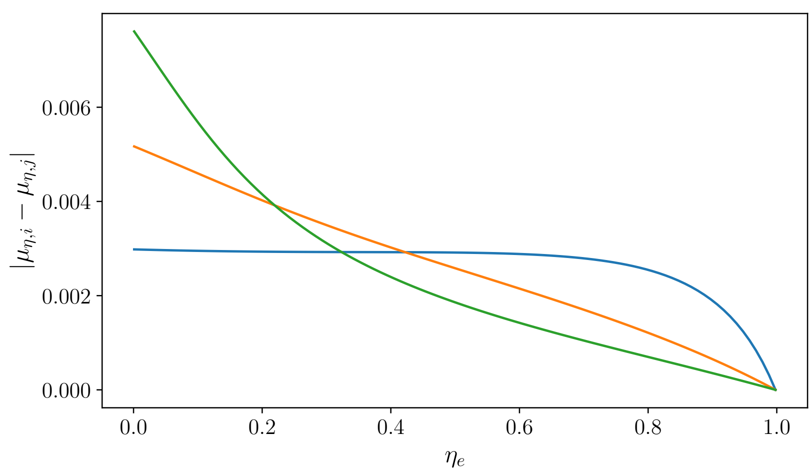

Figure 3: Equilibrium inequality drift for different participation constraints. The blue plot has $\bar{\psi} = 0.2$. The orange plot has $\bar{\psi} = 1.0$. The green plot has $\bar{\psi} = 2.0$. Wealth distribution within the household sector is set to be equal.

For $\bar{\psi} \in (0, \infty)$, the agent portfolio choices $\{\theta_{i,t}\}_{i \le I}$ are heterogeneous across the population and so the wealth distribution impacts the aggregate capital price. Holding $\sigma_{q,t}$ constant, we can see that a more unequal distribution leads to a lower excess return on capital because most of the household wealth in held by an agent facing a small participation constraint. Theorem 2 shows that as $\eta_{e,t} \to 0$, the $\sigma_{q,t}$ becomes constant and intuition above is precisely true. Figure 4 plots the equilibrium functions when the distribution is equal (the blue line) and when one household owns all the wealth (the orange line). This also shows numerically that high household inequality pushes up the risk premium and pushes down the risk free rate.

Theorem 2. *As $\eta_e \to 0$, the $\sigma_q \to \sigma$ and greater household inequality leads to a lower excess return on capital.*

Proof. Available in the appendix E. □

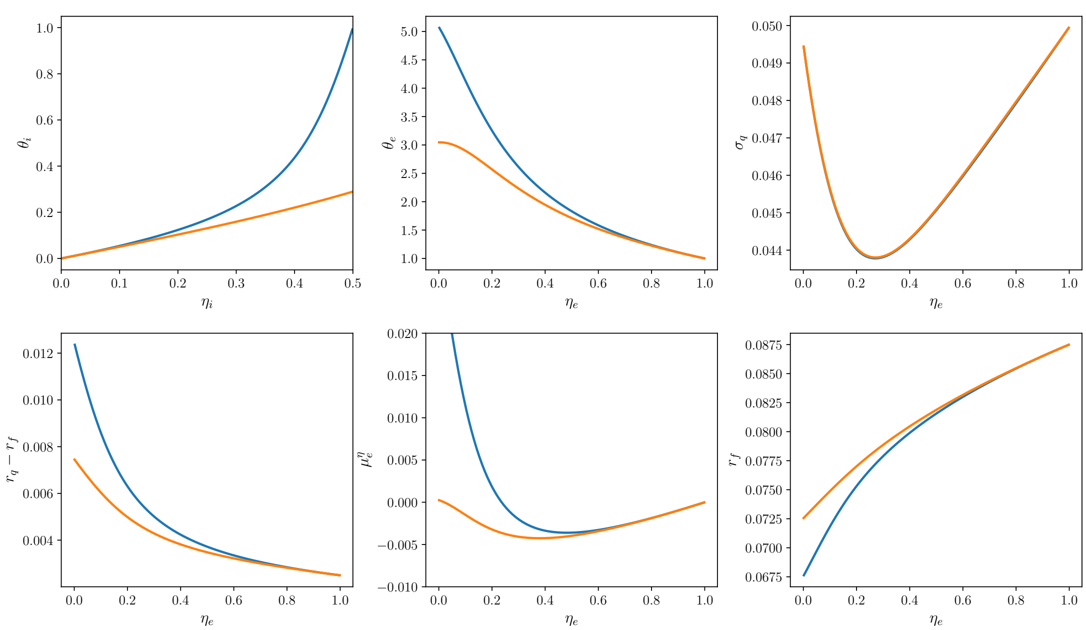

Figure 4: Equilibrium functions in a two households economy with log utility and varying inequality. The orange line is when one household has all the wealth. The blue line is when each household takes half of the wealth in household sector. Participation constraint is $\bar{\psi} = 2.0$.

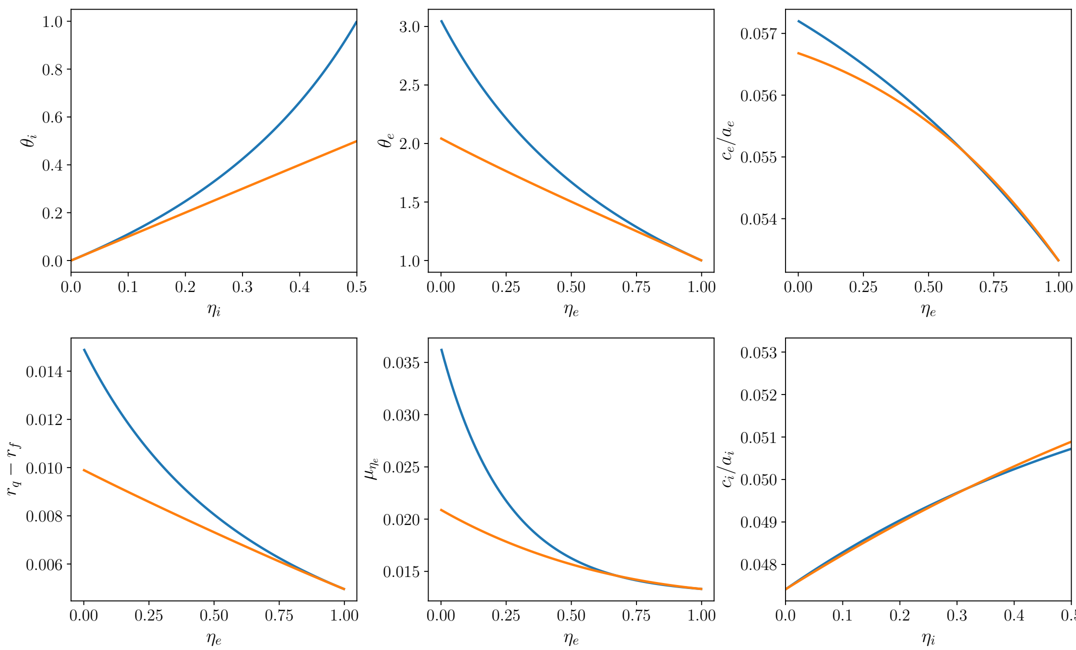

Figure 5: Equilibrium functions in a two households economy with CRRA utility ($\gamma = 2$) and varying inequality. The orange line is when one household has all the wealth. The blue line is when each household takes half of the wealth in household sector. Participation constraint is $\bar{\psi} = 2.0$.

4.2 CRRA With Varying Consumption-to-Wealth Ratios

Finally, we can consider the CRRA case where the consumption-to-wealth ratio is no longer approximately constant. Figure 5 shows the analogous figure for CRRA preferences. Evidently, the participation constraint changes the curvature in the consumption-to-wealth ratio of the households.

5 Quantitative Model

In this section, we show that a calibrated version of our model can match cross-section local projections and long term trends in inequality. We then use our calibrated version of our model to study how macroprudential policy impacts inequality.

5.1 Evidence on Asset Pricing and Wealth Inequality

Before moving to our calibrated model, we estimate how the US equity risk premium impacts the wealth distribution in recessions and expansions. We show that there is wealthier households and financial intermediaries have a more positive exposure to the equity risk premium, particularly in recessions. We interpret this as evidence that poorer agents are less able to take advantage of business cycle frequency asset return risk. We match this data in our quantitative model.

5.1.1 Data Sources

We use data from the following sources. Stock market returns are from Welch and Goyal (2008). Dividend and risk free rate data are from the *Shiller Online Database*. Wealth distribution data is from the updated version of Saez and Zucman (2016). Financial institution data is constructed from the *CRSP Database*. For all empirical analysis, we use times series from 1976 until 2023 at a monthly frequency.

We estimate the equity risk premium since it is not directly observed. We proxy the risk premium by the fitted value of the following regression:

$$\sum_{k=1}^{K} R_{t\to t+k} - r_{f,t} = \beta_0 + \beta_1 dp_t + \epsilon_t$$

where $R_{t\to t+k}$ is the cumulative k-period future returns, r_f is the risk free rate, and dp_t is the dividend yield.[5] For the baseline specification, we use $k = 1$ but the results do not materially change for other values.

5.1.2 Household Risk Premium Exposure

To measure the impact of risk premium on the household wealth distribution, we perform a Jordà (2005) style local projections and run the following regression

$$\log\left(\frac{W_{p,t+h}}{W_{p,t}}\right) = \alpha_{p,h} + \beta_{p,h} rp_t^K + \epsilon_{p,t+h}$$

for horizon $h = 1$ to 30 months, where $w_{p,t+h}$ is the real wage growth of households in $p - th$ percentile at horizon h. We repeat this regression for 4 different wealth percentiles $p \in \{0.01, 0.1, 40.0, 50.0\}$ denoting the top 0.01%, top 0.1%, middle 40%, and bottom 50% of the household wealth distribution, respectively. The top panel of Figure 6 displays the coefficient $\beta_{p,h}$ for different percentile levels. First, risk premium tends to affect wealth positively over the longer horizon. Second, this effect is larger

[5] For robustness, we estimate the risk premium using the Fama-French three factor model instead of dividend yield and get similar results.

among the top wealth percentiles compared to the bottom percentiles. The results do not change if we add lagged risk premium as controls to account for the possibility that wealth share moves because risk premium is correlated.

Next, we study the response of wealth distribution to risk premium conditional on the economy being in a recessionary state. Recessionary periods correspond to the NBER recessionary dates. We run the following regression:

$$\log\left(\frac{W_{p,t+h}}{W_{p,t}}\right) = \alpha_{p,h} + \tilde{\beta}_{p,h} rp_t^K \times 1_{Rec} + \epsilon_{p,t+h}$$

where 1_{Rec} is a dummy variable taking a value 1 during NBER recessionary periods, and 0 otherwise. The coefficient $\tilde{\beta}_{p,h}$ measures the response of wealth distribution to *conditional* risk premium. The bottom panel of Figure 6 presents the coefficients, where the unconditional patterns also hold conditional on recessionary periods. We plot the ratio of the conditional exposure in a recession to the unconditional exposure in the bottom panel of Figure 6. Evidently, the conditional effect of risk premium on wealth is larger for top wealth percentiles.

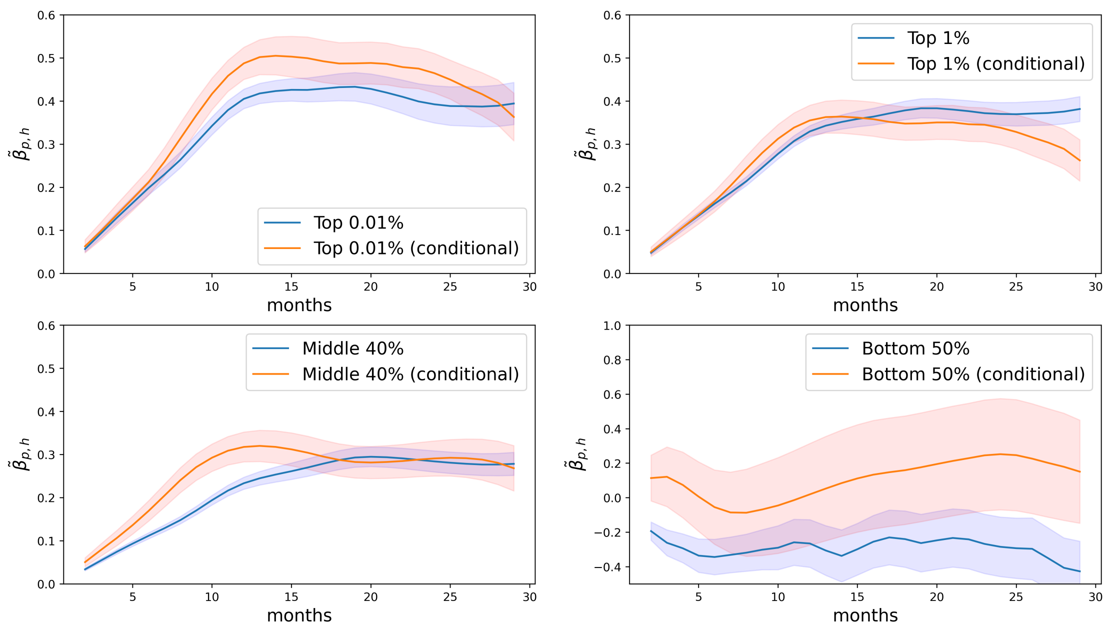

Figure 6: The figure plots the impulse response of wealth distribution to risk premium ($\beta_{p,h}$) obtained from the regression $\log\left(W_{p,t+h}/W_{p,t}\right) = \alpha_{p,h} + \beta_{p,h} rp_t^K + \epsilon_{p,t+h}$. The red lines are the conditional impulse response of wealth distribution to risk premium ($\beta_{p,h}$) obtained from the regression $\log\left(W_{p,t+h} W_{p,t}\right) = \alpha_{p,h} + \tilde{\beta}_{p,h} rp_t^K \times 1_{Rec} + \epsilon_{p,t+h}$. The data for wealth percentiles come from Saez and Zucman (2016), and risk premium is estimated using a factor model.

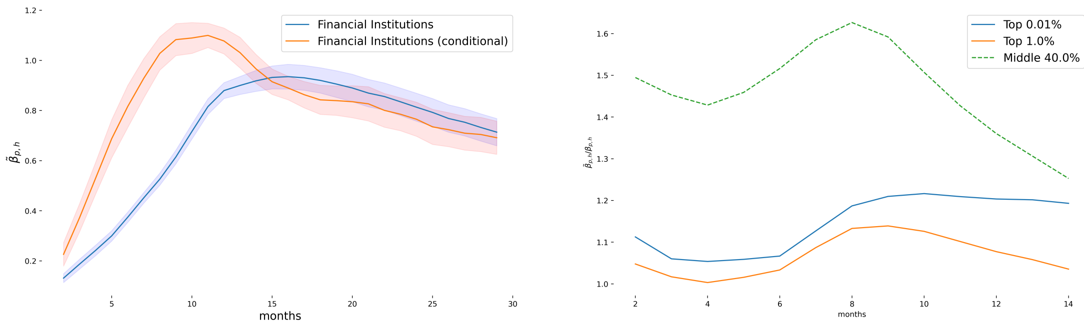

Figure 7: The left panel plots plots the impulse response of wealth distribution to risk premium $(\beta_{BHC,,h})$ obtained from the regression $\log\left(W_{BHC,t+h}/W_{BHC,t}\right) = \alpha_{BHC,h} + \beta_{BHC,h} rp_t^K + \epsilon_{BHC,,t+h}$. The right panel plots $(\beta_{BHC,h})$ obtained from the regression $\log\left(W_{BHC,t+h}/W_{BHC,t}\right) = \alpha_{BHC,h} + \tilde{\beta}_{BHC,h} rp_t^K \times 1_{Rec} + \epsilon_{BHC,t+h}$. The right panel plots the ratio of conditional exposure to unconditional exposure $\tilde{\beta}_{p,h}/\beta_{p,h}$ for top three wealth distribution percentiles $p \in \{0.01, 0.1, 10.0\}$.

5.1.3 Financial Institution Risk Premium Exposure

We repeat the same empirical exercise for the financial institutions in the US. The wealth of financial institutions is computed as the sum of market capitalization of financial companies in the CRSP universe. Figure 7 confirms that the response of financial institution wealth to risk premium are similar to the top wealth percentile households, both unconditionally and conditional on recessionary periods.

5.2 Calibration

We calibrate the model with a strategy that combines targeting model moments and data moments. Remaining parameters are taken from the literature. Table 3 displays the calibrated parameters. The discount rate is set to 5% based on the literature (Krishnamurthy and Li (2020), Gertler and Kiyotaki (2010) etc.). Expert's discount rate is 7% that includes a death rate of 2% in line with Gârleanu and Panageas (2015). The risk aversion parameter is calibrated to match expert sector leverage ratio of 6.6. This number is closer to the value of 6 used in Krishnamurthy and Li (2020). The volatility parameter is set to 0.2.[6] The portfolio constraint parameter is calibrated to generate a 32% portfolio share from the middle income households.

[6]While this is higher than the historical volatility of 4% of real GDP growth (Bohn's historical data), we set it to a higher value since the only shock in the model is a Brownian TFP shock with which we aim to match the entire evolution of wealth distribution in the past century. A lower value of σ does not materially change the asset pricing moments since participation constraints remain to be the major driver.

Parameter	Symbol	Value	Target
Risk aversion	γ	3.0	Expert leverage
Households' Discount rate	ρ_h	0.05	Literature
Experts' Discount rate	ρ_e	0.07	Literature
Reversion rate	β	0.5	Data
Volatility	σ	0.2	Long-run Volatility of TFP
Portfolio constraint	$\bar{\psi}$	10	Middle-40 pctl. portfolio share

Table 3: Calibrated parameters.

5.3 Comparison to Asset Pricing Data

Table 4 reports the asset pricing moments in the data and from the model. None of the asset pricing moments are specifically targeted and hence a measure of success of our model is to see how well it matches these moments. The table shows that the model generates a sizable equity returns and risk premium, and also generates endogenous volatility close to the data. Having the expert sector in the model helps generate amplification. The agents in the model have CRRA utility with a risk aversion parameter calibrated to $\gamma = 3$. Unlike Guvenen (2009), Gârleanu and Panageas (2015), Gomez (2017) and Basak and Chabakauri (2024), we do not have preference heterogeneity between the agents in the economy. More generally, the asset pricing literature typically generates a high risk premium using either Epstein-Zin utility and/or calibrating with a high risk aversion parameter. We require neither of these features to match the equity premium since participation constraints of households generate all the intended effects.

	Data	Model	Source
E[Risk premium]	5.5%	3.8%	Predictive regression
Std[Risk premium]	4.7%	1.1%	Predictive regression
E[Equity returns]	6.4%	8.5%	*Amit Goyal's* website
Std[Equity returns]	19.3%	13%	*Amit Goyal's* website
E[Risk-free rate]	4.3%	4.5%	*Amit Goyal's* website

Table 4: The table reports the asset pricing moments in the data and the model. The time period is from 1950Q1 till 2021Q1. All values are in annualized terms.

5.4 Comparison to wealth distribution

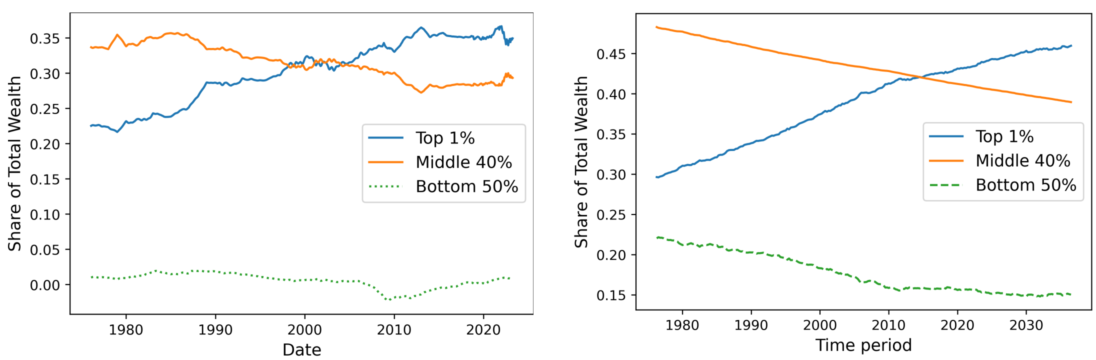

(a) Data (Saez and Zucman (2016)) (b) Model

Figure 8: The left panel presents the share of total wealth for the households in top 1%, middle 40%, and bottom 50%, respectively. The time period is from 1976 till 2023 at monthly frequency. The data is taken from Saez and Zucman (2016). The right panel presents the share of total wealth produced by the model for the same percentiles of household wealth.

Instead of matching specific moments of wealth distribution as in Gomez (2017), we feed-in an initial wealth distribution resembling the data and track the model implied evolution of wealth distribution over time. The left panel of Figure 8 displays the empirical wealth distribution from Saez and Zucman (2016) between the time periods 1976 and 2023. The top 1 pctl. households start out at a lower share of total wealth compared to the bottom 40 pctl., but gradually take over the latter. The bottom 50 pctl. households instead start with a much lower share of wealth, and remain there for the rest of the time period. The right panel of Figure 8 displays the evolution of wealth distribution implied by the model. It is important to note that the wealth distribution is not particularly targeted in the calibration. The participation constraints on the households alone generates and matches the empirical evolution of wealth share over a comparable time period. While the share of wealth held by bottom 50 pctl. households is close to zero in the data, we feed in a larger value because the wealth of the agents do not go below zero like in the data due to their risk aversion. Nevertheless, the model captures the declining trend of these households pretty well. Apart from such minor differences in the way we feed in the initial distribution, the model successfully captures the long-term trend of the “hollowing-out" of the wealth distribution. Notably, in addition to the widening gap between the top 1 pctl. and bottom 50 pctl. households that is much talked about the literature, the model also captures the declining wealth share of middle 40 pctl. households that resonates with the disappearance of middle-class in the US.

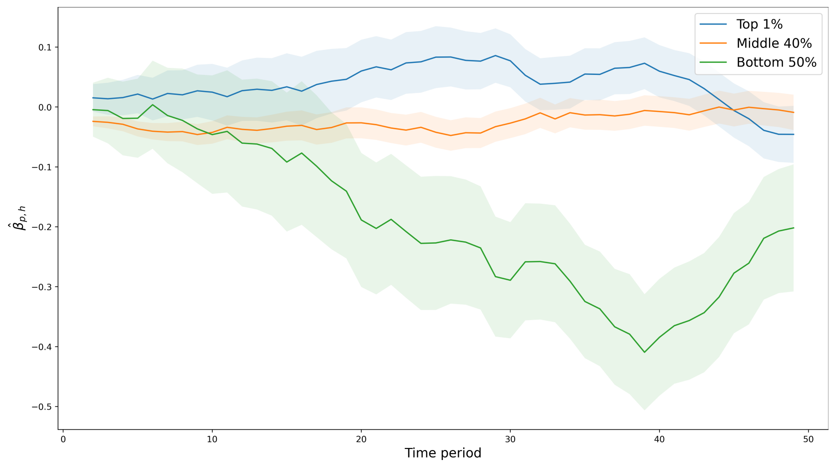

Figure 9: The figure plots the model implied impulse response of wealth distribution to risk premium ($\beta_{p,h}$) obtained from the regression $\log\left(\frac{w_{p,t+h}}{w_{p,t}}\right) = \alpha_{p,h} + \hat{\beta}_{p,h} rp_t \times 1_{REC} + \epsilon_{p,t+h}$. The wealth levels are proxied by the wealth-shares η_p from the model for different percentiles. The indicator function 1_{REC} takes a value of 1 if productivity level is below its mean. The risk premium rp_t used in the regression is the model implied risk premium.

Lastly, we perform local projection using the model implied equilibrium quantities to show the hollowing-out effect. Figure 9 displays the result where we regress change in wealth shares of households in different wealth percentiles on the risk premium implied by the model. Consistent with the empirical observation, the top 1-pct. households have a higher exposure to risk premium compared to agents in the other wealth percentiles. Admittedly, the effects on the middle 40 pctl. and bottom 50 pctl. are much stronger than what we see in the data. This could be because in the data, households have access to other assets such as housing, private equity, which affect wealth distribution in complicated ways. Nevertheless, the local projections capture the spirit of empirical observation which is that the equity markets have played a dominant role in hollowing out the wealth distribution in the US.

5.5 Macroprudential Policy and Inequality

We close the paper by considering how macroprudential policy impacts inequality. We do this by performing the following counterfactual exercise. We introduce an exogenous leverage constraint $\theta_{e,t} = \bar{\ell}$ for the intermediary sector. We then simulate a collection of productivity paths, track the distribution evolution starting from the same initial wealth distribution as in section 5.4, and plot the fan chart for relative differences between wealth share dynamics (in percentage) with and without the leverage constraint in

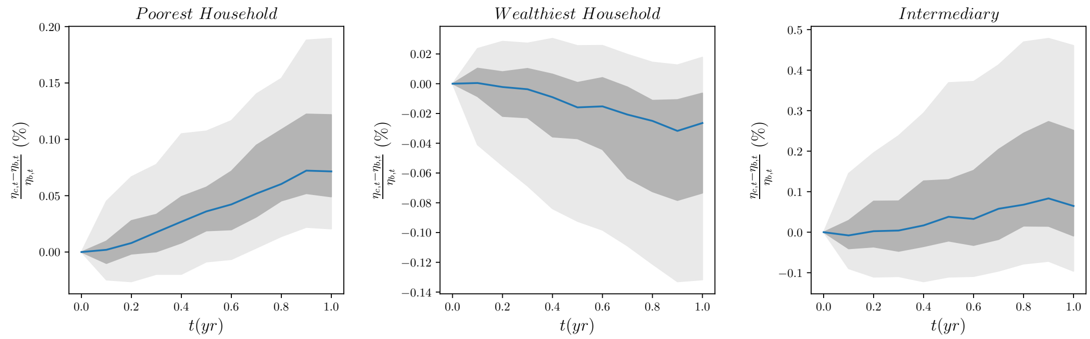

Figure 10: Forecasted distributional dynamics with $\bar{\ell} = 2.0$. The left, middle and right panel are fan charts of relative wealth share responses' difference), at quantile 10%, 30 %, 50%, 70% and 90%. Subscript c stands for counter factual and b stands for baseline model without leverage constraint (same as below).

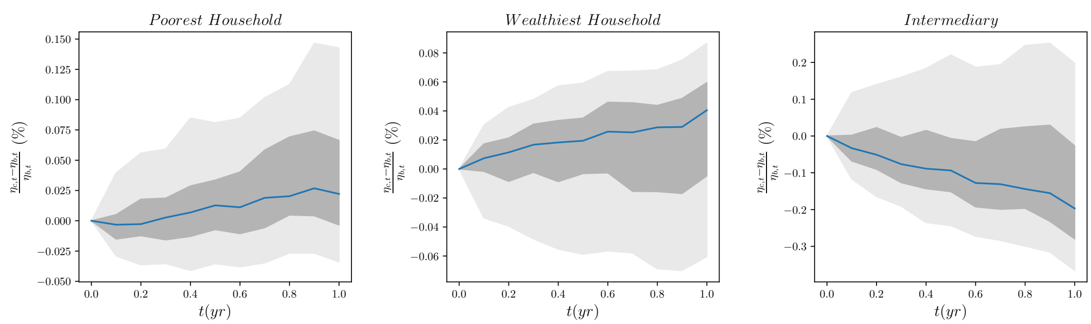

Figure 11: Forecasted distributional dynamics with $\bar{\ell} = 1.5$. The left, middle and right panel are fan charts of relative wealth share responses' difference), at quantile 10%, 30 %, 50%, 70% and 90%.

Figure 10 and 11 show the results from our counterfactual exercise. We can make the following observations. First, the distributional responses are skewed. Intuitively, when intermediary's net worth level is low, the wealthiest household steps in and earns the high return, which explains the skewness of responses. Second, intermediary leverage is closely connected with the rate of recovery and overall inequality dynamics. Comparing Figure 10 and Figure 11, we find that a tighter leverage constraint exacerbates inequality and slows down the intermediary's rate of recover. By contrast, a tighter leverage constraint induces the household to hold more capital and thus furthers push up the premium, which is disproportionately earned by households at different wealth levels.

6 Conclusion

In this paper, we have studied the feedback between asset pricing and inequality when there are participation constraints. This required us to develop a new methodology that uses deep learning to characterize global solutions to macroeconomic models with long-term assets, agent heterogeneity, and non-trivial household portfolio choice. We believe this technique provides a general approach for exploring how asset pricing relates to inequality across investors and institutions. We used a calibrated version of our model to explore how limited participation in asset markets leads to amplification of the capital price process. We find that our relatively simple model does a good job of capturing the trend movements in inequality over the past half-century. We also find that the financial sector stabilization increases inequality.

References

Azinovic, Marlon and Jan Žemlička, "Economics-Inspired Neural Networks with Stabilizing Homotopies," *arXiv preprint arXiv:2303.14802*, 2023.

_ **, Harold Cole, and Felix Kubler**, "Asset Pricing in a Low Rate Environment," NBER Working Papers 31832, National Bureau of Economic Research, Inc 2023.

_ **, Luca Gaegauf, and Simon Scheidegger**, "Deep equilibrium nets," *International Economic Review*, 2022, *63* (4), 1471–1525.

Basak, Suleyman and Domenico Cuoco, "An equilibrium model with restricted stock market participation," *The Review of Financial Studies*, 1998, *11* (2), 309–341.

_ **and Georgy Chabakauri**, "Asset Prices, Wealth Inequality, and Taxation," *Wealth Inequality, and Taxation (June 1, 2023)*, 2023.

_ **and** _ , "Asset Prices, Wealth Inequality, and Taxation," *SSRN Electronic Journal*, 1 2024.

Bretscher, Lorenzo, Jesús Fernández-Villaverde, and Simon Scheidegger, "Ricardian Business Cycles," *Available at SSRN*, 2022.

Brunnermeier, Markus K and Yuliy Sannikov, "A macroeconomic model with a financial sector," *American Economic Review*, 2014, *104* (2), 379–421.

_ **and** _ , "Macro, money, and finance: A continuous-time approach," in "Handbook of Macroeconomics," Vol. 2, Elsevier, 2016, pp. 1497–1545.

Campbell, John Y and Luis M Viceira, *Strategic asset allocation: portfolio choice for long-term investors*, Clarendon Lectures in Economic, 2002.

Cioffi, Riccardo A, "Heterogeneous risk exposure and the dynamics of wealth inequality," *URL: https://rcioffi. com/files/jmp/cioffi_jmp2021_princeton. pdf (cit. on p. 7)*, 2021.

Duarte, Victor, "Machine learning for continuous-time economics," *Available at SSRN 3012602*, 2018.

Fagereng, Andreas, Matthieu Gomez, Emilien Gouin-Bonenfant, Martin Holm, Benjamin Moll, and Gisle Natvik, "Asset-price redistribution," Technical Report, Working Paper 2022.

Fernández-Villaverde, Jesús and Oren Levintal, "The Distributional Effects of Asset Returns," 2024.

Fernandez-Villaverde, Jesus, Galo Nuno, George Sorg-Langhans, and Maximilian Vogler, "Solving high-dimensional dynamic programming problems using deep learning," *Unpublished working paper*, 2020.

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

Fernández-Villaverde, Jesús, Joël Marbet, Galo Nuño, and Omar Rachedi, "Inequality and the Zero Lower Bound," CESifo Working Paper Series 10471, CESifo 5 2023.

Gertler, Mark and Nobuhiro Kiyotaki, "Chapter 11 - Financial Intermediation and Credit Policy in Business Cycle Analysis," in Benjamin M. Friedman and Michael Woodford, eds., *NBER Handbook*, Vol. 3 of *Handbook of Monetary Economics*, Elsevier, 2010, pp. 547–599.

Gomez, Matthieu, "Asset Prices and Wealth Inequality," 2017 Meeting Papers 1155, Society for Economic Dynamics 2017.

— **and Émilien Gouin-Bonenfant**, "Wealth inequality in a low rate environment," *Econometrica*, 2024, *92* (1), 201–246.

— **et al.**, "Asset prices and wealth inequality," *Unpublished paper: Princeton*, 2016.

Gopalakrishna, Goutham, "Aliens and continuous time economies," *Swiss Finance Institute Research Paper*, 2021, *21* (34).

Gu, Zhouzhou, Methieu Laurière, Sebastian Merkel, and Jonathan Payne, "Deep Learning Solutions to Master Equations for Continuous Time Heterogeneous Agent Macroeconomic Models," *Princeton Working Paper*, 2023.

Guvenen, Fatih, "A parsimonious macroeconomic model for asset pricing," *Econometrica*, 2009, *77* (6), 1711–1750.

Gârleanu, Nicolae and Stavros Panageas, "Young, Old, Conservative, and Bold: The Implications of Heterogeneity and Finite Lives for Asset Pricing," *Journal of Political Economy*, 2015, *123* (3), 670–685.

Han, Jiequn, Arnulf Jentzen, and Weinan E, "Solving high-dimensional partial differential equations using deep learning," *Proceedings of the National Academy of Sciences*, 2018, *115* (34), 8505–8510.

— **, Yucheng Yang, and Weinan E**, "Deepham: A global solution method for heterogeneous agent models with aggregate shocks," *arXiv preprint arXiv:2112.14377*, 2021.

Huang, Ji, "A Probabilistic Solution to High-Dimensional Continuous-Time Macro-Finance Models," *Available at SSRN 4122454*, 2022.

— , "Breaking the Curse of Dimensionality in Heterogeneous-Agent Models: A Deep Learning-Based Probabilistic Approach," *SSRN Working Paper*, 2023.

Irie, Magnus, "Innovations in Entrepreneurial Finance," 2024.

Kahou, Mahdi Ebrahimi, Jesús Fernández-Villaverde, Jesse Perla, and Arnav Sood, "Exploiting symmetry in high-dimensional dynamic programming," Technical Report, National Bureau of Economic Research 2021.

Krishnamurthy, Arvind and Wenhao Li, "Dissecting Mechanisms of Financial Crises: Intermediation and Sentiment," NBER Working Papers 27088, National Bureau of Economic Research, Inc May 2020.

Krusell, Per and Anthony A Smith, "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 1998, *106* (5), 867–896.

Kubler, Felix and Simon Scheidegger, "Self-justi ed equilibria: Existence and computation," 2018 Meeting Papers 694, Society for Economic Dynamics 11 2018.

Maliar, Lilia, Serguei Maliar, and Pablo Winant, "Deep learning for solving dynamic economic models.," *Journal of Monetary Economics*, 2021, *122*, 76–101.

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

Saez, Emmanuel and Gabriel Zucman, " Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data," *The Quarterly Journal of Economics*, 02 2016, *131* (2), 519–578.

Sauzet, Maxime, "Projection methods via neural networks for continuous-time models," *Available at SSRN 3981838*, 2021.

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

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

A Recursive Characterization of Equilibrium

In this Appendix, we work through recursive characterization of equilibrium. We start by setting up the optimization problem of the agents recursively in the "natural" state variables:

$$(z, K, g = \{a_i\}_{1 \leq i \leq I}).$$

This characterization is convenient for understanding the agent optimization problem but turns out to be hard for the neural network to solve. We then characterize equilibrium in space of wealth shares, which turns out be more convenient for training the neural network.

A.1 Characterization in Natural State Variables

State variables and beliefs: We assume there exists a solution to the equilibrium that is recursive in the aggregate state variables, (y, K, g), which we denote by $(\cdot)$. This means that the states that appear in the household decision problem are $(a_i, \cdot)$. In this case, beliefs about the price process can be characterized by beliefs about how the distribution and aggregate capital stock evolves since prices are all implicitly functions of the aggregate state variables. Formally, an agent's beliefs about the evolution of the distribution are characterized by their beliefs about the drift and covariance of other agents wealth and the drift of capital stock, $\{\hat{\mu}_{a_j}(\cdot), \hat{\sigma}_{a_j}(\cdot), \hat{\mu}_K(\cdot) : j \neq i\}$, which imply beliefs about prices through the pricing functions $(r(\cdot), q(\cdot))$. Technically, agents also have beliefs about the evolution of other agent's labor status but we leave that implicit since it is unrelated to agent decisions. We let $V_i(a_i, \cdot)$ denote household i's value function. It is helpful to characterize the equilibrim in terms of three blocks.

1. Agent optimization block: Given their beliefs, agent i chooses (c_i, b_i) to solve the Hamilton-Jacobi-Bellman Equation (HJBE) equation (A.1) below:

$$\begin{aligned}
\rho V_i(a_i, \cdot) = \max_{c_i, b_i, \iota_i} \Big\{ & u(c_i) + \psi_i(a_i, b_i, \cdot) a_i \Xi_i + \frac{\partial V_i}{\partial a_i} \mu_{a_i}(a_i, c_i, b_i, \iota, \cdot) + \frac{\partial V_i}{\partial z} \mu_z \\
& + \frac{\partial V_i}{\partial K} \hat{\mu}_K(\cdot) + \frac{1}{2} \frac{\partial^2 V_i}{\partial a_i^2} \sigma_{a_i}^2(b_i, \cdot) + \frac{1}{2} \frac{\partial^2 V_i}{\partial z^2} \sigma_z^2 + \frac{\partial^2 V_i}{\partial a_i \partial z} \sigma_{a_i}(b_i, \cdot) \sigma_z \\
& + \sum_{j \neq i} \frac{\partial^2 V_i}{\partial a_j \partial z} \hat{\sigma}_{a_j}(\cdot) \sigma_z + \frac{1}{2} \sum_{j \neq i, j' \neq i} \frac{\partial^2 V_i}{\partial a_j \partial a_{j'}} \hat{\sigma}_{a_j}(\cdot) \hat{\sigma}_{a'_j}(\cdot) \Big\}
\end{aligned} \tag{A.1}$$

where the first two lines are the standard terms that would appear in the individual agent optimization problem, and the last lines capture the impact of the distribution on the agent's value function. The first order conditions with respect to (c_i, b_i) are given by the following respectively:

$$
\begin{aligned}
[c_i]: \quad & 0 = u'(c_i) - \partial_a V_i(a_i) \\
[b_i]: \quad & 0 = -\frac{\partial V_i}{\partial a_i}(r(\cdot) - r_k(\cdot)) + \frac{\partial^2 V_i}{\partial a_i^2}(a_i - b_i)\sigma_q^2(\cdot) \\
& \quad + \frac{\partial^2 V_i}{\partial a_i \partial z}\sigma_q(\cdot)\sigma_z(\cdot) + \sum_{j \neq i} \frac{\partial^2 V_i}{\partial a_i \partial a_j}\sigma_q(\cdot)\hat{\sigma}_{a_j}(\cdot) + \frac{\partial \psi_i}{\partial b_i}\Xi_i \\
[\iota_i]: \quad & 0 = -\frac{1}{\hat{q}(\cdot)} + \phi'(\iota_i)
\end{aligned}
$$

where $\overline{R}_k(\cdot) - r(\cdot)$ is the "risk-premium" in the economy. From these equations, we can immediately see that $\iota_i = (\phi')^{-1}(1/q) =: \iota$ is the same for agents.

2. State evolution block: The law of motion for each agent satisfies (2.2) and aggregate capital stock satisfies:

$$
dK_t = (\phi(\iota_t)K_t - \delta K_t)dt. \tag{A.2}
$$

3. Market clearing and belief consistency block: The equilibrium pricing functions $(r(\cdot),\ q(\cdot))$ are pinned down implicitly by the market clearing conditions in part 2 of the equilibrium definition. Under this recursive characterization, the belief consistency condition becomes that each agent has correct beliefs about the evolution of wealth for the other agents and aggregate capital stock:

$$
\left(\hat{\mu}_{a_j}(\cdot), \hat{\sigma}_{a_j}(\cdot), \hat{\mu}_K(\cdot)\right) = \left(\mu_{a_j}(\cdot), \sigma_{a_j}(\cdot), \mu_K(\cdot)\right)
$$

A.2 Characterization as Master Equations in Wealth Shares

We now re-characterize the equilibrium as a collection of "master" differential equations for the neural network to train. The first change is to the characterization of the distribution. It turns out that the recursive characterization in agent wealth levels leads to a complicated fixed point problem that is hard for the Neural Network to train (we discuss in detail in Section 3.3 after we introduce the algorithm.). Instead, it will be convenient to characterize the equilibrium in terms of wealth shares. Let $A := \sum_{j \geq 1} a_j$

denote total wealth in the economy. Let $\eta_i := a_i/A$ denote the share of wealth held by agent i. Then, the aggregate state of the economy can be written in terms of wealth shares as $(z, K, \{\eta_j, l_j\}_{j\geq 1})$. We can now restate the equilibrium conditions using the wealth shares as the state. For notational convenience, we drop the explicit dependence on $(z, K, \{\eta_j, l_j\}_{j\geq 1})$ and, where possible.

The second change is to work with the derivative of the value function. We define the marginal value of wealth and the partial derivatives of the marginal value of wealth (the so called "stochastic discount factors") by:

$$\xi_i := \frac{\partial V_i}{\partial a_i}, \qquad \partial_a \xi_i := \frac{\partial \xi_i}{\partial a_i} = \frac{\partial^2 V_i}{\partial a_i^2}, \qquad \partial_{a_j} \xi_i := \frac{\partial \xi_i}{\partial a_j} = \frac{\partial^2 V_i}{\partial a_i a_j}$$

Once equilibrium is imposed, all the endogenous objects in the economy must be functions of $(z, K, \{\eta_j, l_j\}_{j\geq 1})$. We can use Ito's Lemma to express the drift and volatility of ξ_i in terms of derivatives of ξ_i with respect to $(z, K, \{\eta_j, l_j\}_{j\geq 1})$ in equilibrium:

$$\begin{aligned} \xi_i \mu_{\xi_i} &= \frac{\partial \xi_i}{\partial z}\mu_z + \frac{\partial \xi_i}{\partial K}\mu_K + \sum_j \frac{\partial \xi_i}{\partial \eta_j}\eta_j \mu_{\eta_j,t} + \sum_j \frac{\partial^2 \xi_i}{\partial z \partial \eta_j}\eta_j \sigma_{\eta_j,t}\sigma_z \\ &\quad + \frac{1}{2}\frac{\partial^2 \xi_i}{\partial z^2}\sigma_z^2 + \frac{1}{2}\sum_{j,j'} \frac{\partial^2 \xi_i^2}{\partial \eta_j \partial \eta_{j'}}\eta_j \eta_{j'} \sigma_{\eta_j,t}\sigma_{\eta_{j'},t} \qquad \text{(A.3)} \\ \xi_i \sigma_{\xi_i} &= \frac{\partial \xi_i}{\partial z}\sigma_z + \sum_j \frac{\partial \xi_i}{\partial \eta_j}\eta_j \sigma_{\eta_j,t}. \qquad \text{(A.4)} \end{aligned}$$

The third change is that we impose belief consistency and market clearing conditions, where possible.

We now rewrite the general equilibrium blocks with these changes imposed.

1. Agent optimization block: Applying the Envelope Theorem to the HJBE (A.1), imposing belief consistency, and using Ito's Lemma to collect terms leads to the continuous time Euler equation (the so called "master equation" for the economy) for ξ_i. We work through the details heuristically and formally in Appendix A.3. Given prices

$(r, r_k, q, \mu_q, \sigma_q)$, agent optimization implies that $(\xi_i, c_i, b_i, \iota_i)$ satisfy:

$$\begin{aligned} 0 &= -\rho + r + \mu_{\xi_i,t} && \text{(A.5)} \\ u'(c_i) &= \xi_i \\ r - r_k &= \sigma_{\xi_i}\sigma_q + \frac{\partial \psi_i}{\partial b_i}\Big|_{a_i=\eta_i q} && \text{(A.6)} \\ \iota_i &= (\phi')^{-1}\left(q^{-1}\right) \end{aligned}$$

where μ_{ξ_i} satisfies (A.3) and σ_{ξ_i} satisfies (A.4).

2. State evolution block: Given prices $(r_t, r_k, q, \mu_q, \sigma_q)$ and agent optimization (ξ, c, b, ι), we can use Ito's Lemma to get the law of motion for each wealth share $\eta_{j,t} = a_{j,t}/(q_t K_t)$:

$$\begin{aligned} \eta_j \mu_{\eta_j,t} &= \frac{1}{a_{j,t}}\left[r_{k,t}(a_{j,t} - b_{j,t}) + b_{j,t} r_t - (u')^{-1}(\xi_{j,t})\right] \\ &\quad - \mu_{q,t} - \mu_{K,t} + \sigma_{q,t}\left(\sigma_{q,t} - \frac{1}{a_{j,t}}(a_{j,t} - b_{j,t})\sigma_{q,t}\right) \\ &= r_{k,t} - \mu_{q,t} - \mu_{K,t} + \frac{b_{j,t}}{\eta_{j,t} q_t K_t}(r_t - r_{k,t}) - \frac{(u')^{-1}(\xi_{j,t})}{\eta_{j,t} q_t K_t} + \frac{b_{j,t}}{\eta_{j,t} q_t K_t}\sigma_{q,t}^2 && \text{(A.7)} \\ \eta_j \sigma_{\eta_j,t} &= \frac{1}{a_{j,t}}\left(a_{j,t} - b_{j,t}\right)\sigma_{q,t} - \sigma_{q,t} = -\frac{b_{j,t}}{\eta_{j,t} q_t K_t}\sigma_{q,t} && \text{(A.8)} \end{aligned}$$

The evolution of K_t once again satisfies (A.2).

3. Equilibrium block: The market clearing conditions now become:

$$\sum_i c_i + \Phi(\iota)K = y \qquad \sum_i b_i = 0 \qquad \sum_i (\eta_i A - b_i) = K$$

where the aggregate household wealth satisfies $A := \sum_{j\geq 1} a_j = qK$ and so the capital market clearing condition simply becomes $\sum_i \eta_i = 1$. The risk free rate can only be implicitly expressed in terms of the state variables through its dependence on the stochastic processes for ξ and q (using the agent first order conditions):

$$r = r_k + \sigma_{\xi_i}\sigma_q + \frac{\partial \psi_i}{\partial b_i}$$

The price of capital is even more difficult to handle because capital is a long-lived asset for which the price can only be implicitly expressed in terms of the state variables using

Itô's Lemma:

$$
\begin{aligned}
q\mu_{q,t} &= \sum_j \frac{\partial q}{\partial \eta_j}\eta_j\mu_{\eta_j,t} + \frac{\partial q}{\partial z}\mu_{z,t} + \frac{\partial q}{\partial K}\mu_{K,t} + \sum_j \frac{\partial^2 \xi_i}{\partial z \partial \eta_j}\eta_j\sigma_{\eta_j,t}\sigma_z \\
&\quad + \frac{1}{2}\sum_{j,j'} \frac{\partial^2 q}{\partial \eta_j \partial \eta_{j'}}\eta_j\eta_{j'}\sigma_{\eta_j,t}\sigma_{\eta_{j'},t} + \frac{1}{2}\frac{\partial^2 q}{\partial z^2}\sigma_z^2 + \sum_j \lambda(l_j)(\tilde{q}_{i,j} - q_i) \\
q\sigma_{q,t} &= \sum_j \frac{\partial q}{\partial \eta_j}\eta_j\sigma_{\eta_j,t} + \frac{\partial q}{\partial z}\sigma_{z,t}.
\end{aligned}
$$

These expressions for $\mu_{q,t}$ and $\sigma_{q,t}$ are what makes the law of motion for capital "consistent" with the process that we posited in the environment and so are often referred to as the price consistency differential equations.

A.3 Derivations of Analytical Results

In this section, we provide the derivations for the Euler equation we used in Section A.1. The first subsection introduces a heuristic derivation from a continuous-time approximation of the discrete time Euler equation without financial frictions and jumps. The second subsection derives from HJB equation and envelop theorem in the most generic setup.

A.3.1 A Heuristic Derivation

We consider the discrete time version of Euler equation without jumps and financial frictions:

$$
\mathbb{E}\left[\beta \frac{u'(c_{t+1})}{u'(c_t)}\left(\frac{q_{t+1} + d_{t+1}}{q_t}\right)\right] = 1,
$$

where q_t, q_{t+1} are asset price at time t and $t+1$, d_{t+1} is the dividend at time $t+1$. Note that marginal value of wealth is connected with marginal utility by optimal consumption decision:

$$
u'(c_t) = \xi_t,
$$

we could essentially rewrite Euler equation as:

$$
\mathbb{E}\left[\beta \frac{\xi_{t+1}}{\xi_t}\left(\frac{q_t + (d_{t+1} + q_{t+1} - q_t)}{q_t}\right)\right] = 1
$$

Now, consider the case that time step is sufficiently small, i.e., replace $t+1$ as $t+\Delta t$:

$$\begin{aligned}\beta &= e^{-\rho\Delta t},\\ \frac{\xi_{t+1}}{\xi_t} &= 1+\mu_{\xi,t}\Delta t+\sigma_{\xi,t}\Delta\epsilon,\\ \frac{q_{t+1}}{q_t} &= 1+\mu_{q,t}\Delta t+\sigma_{q,t}\Delta\epsilon,\\ \frac{d_{t+1}}{q_t} &= \frac{\pi_t\Delta t}{q_t}+\mathcal{O}(\frac{\Delta\pi_t\Delta t}{q_t}),\end{aligned}$$

where π_t is the net profit process from production, $\Delta\epsilon$ is normally distributed random variable with mean zero and variance Δt ($\mathbb{E}[\Delta\epsilon\Delta\epsilon]=\Delta t$). As we only keep up to order of Δt, the dividend price ration can be essentially simplified as $\frac{d_{t+1}}{q_t}=\frac{\pi_t\Delta t}{q_t}$. Plug in all above equations, Euler Equation can be expressed as:

$$\mathbb{E}\left[(1-\rho\Delta t)(1+\mu_{\xi,t}\Delta t+\sigma_{\xi,t}\Delta\epsilon)\left(1+(\frac{\pi_t}{q_t}+\mu_{q,t})\Delta t+\sigma_{q,t}\Delta\epsilon\right)\right]=1$$

Drop all higher order terms $\Delta\epsilon\Delta t, \Delta t\Delta t$ again, and we have:

$$-\rho+\mu_{\xi,t}+\underbrace{\left(\frac{\pi_t}{q_t}+\mu_{q,t}\right)}_{r_{q,t}}+\sigma_{\xi,t}\sigma_{q,t}=0 \tag{A.9}$$

Similarly, we could also derive the Euler equation by considering the return on risk-free assets, which is:

$$-\rho+r_t+\mu_{\xi,t}=0 \tag{A.10}$$

Taking the difference between (A.9) and (A.10), we get the first order condition for portfolio choice again:

$$r_{q,t}-r_t=\sigma_{\xi,t}\sigma_{q,t}.$$

With financial constrains, however, Euler equation becomes inequality and the above derivations no longer apply. The next subsection explores the full problem in a recursive way.

A.3.2 Full Derivation

Taking the envelope condition and imposing belief consistency to get the continuous time "Euler" equation. To apply all first order conditions, including the portfolio choice and consumption decision, we first work on the wealth space, then convert to the wealth share space and impose all equilibrium conditions.

Lets take the first order derivative w.r.t a_i for the HJB equation (A.1):

$$\begin{aligned}
\rho\frac{\partial V_i(a_i,\cdot)}{\partial a_i} =& u'(c_i)\frac{\partial c_i(\cdot)}{\partial a_i} + \frac{\partial \psi_{H(i)}(a_i,b_i)}{\partial a_i} + \frac{\partial \psi_{H(i)}(a_i,b_i)}{\partial b_i}\frac{\partial b_i}{\partial a_i} + \frac{\partial^2 V_i}{\partial a_i^2}\mu_{a_i}(a_i,c_i,b_i,\iota,\cdot) \\
&+ \frac{\partial V_i}{\partial a_i}\Big(\frac{\partial \mu_{a_i}(a_i,c_i,b_i,\iota,\cdot)}{\partial a_i} + \frac{\partial \mu_{a_i}(a_i,c_i,b_i,\iota,\cdot)}{\partial c_i}\frac{\partial c_i}{\partial a_i} + \frac{\partial \mu_{a_i}(a_i,c_i,b_i,\iota,\cdot)}{\partial b_i}\frac{\partial b_i}{\partial a_i}\Big) \\
&+ \frac{1}{2}\Big[\frac{\partial^3 V_i}{\partial a_i^3}\sigma_{a_i}^2(b_i,\cdot) + 2\frac{\partial^2 V_i}{\partial a_i^2}\sigma_{a_i}(b_i,\cdot)\Big(\frac{\partial \sigma_{a_i}(b_i,\cdot)}{\partial a_i} + \frac{\partial \sigma_{a_i}(b_i,\cdot)}{\partial b_i}\frac{\partial b_i}{\partial a_i}\Big)\Big] + \frac{1}{2}\frac{\partial^3 V_i}{\partial a_i \partial z^2}\sigma_z^2 \\
&+ \frac{\partial^3 V_i}{\partial a_i^2 \partial z}\sigma_{a_i}(b_i,\cdot)\sigma_z + \frac{\partial^2 V_i}{\partial a_i \partial z}\Big(\frac{\partial \sigma_{a_i}(b_i,\cdot)}{\partial a_i} + \frac{\partial \sigma_{a_i}(b_i,\cdot)}{\partial b_i}\frac{\partial b_i}{\partial a_i}\Big)\sigma_z \\
&+ \sum_{j\neq i}\frac{\partial^3 V_i}{\partial a_i^2 \partial a_j}\sigma_{a_i}(b_i,\cdot)\hat{\sigma}_{a_j}(\cdot) + \frac{\partial^2 V_i}{\partial a_i \partial a_j}\Big(\frac{\partial \sigma_{a_i}(b_i,\cdot)}{\partial a_i} + \frac{\partial \sigma_{a_i}(b_i,\cdot)}{\partial b_i}\frac{\partial b_i}{\partial a_i}\Big)\hat{\sigma}_{a_j}(\cdot) \\
&+ \sum_{j\neq i}\frac{\partial^2 V_i}{\partial a_i \partial a_j}\hat{\mu}_{a_j}(\cdot) + \sum_{j\neq i}\frac{\partial^3 V_i}{\partial a_i \partial a_j \partial z}\hat{\sigma}_{a_j}(\cdot)\sigma_z + \frac{1}{2}\sum_{j\neq i, j'\neq i}\frac{\partial^3 V_i}{\partial a_i \partial a_j \partial a_{j'}}\hat{\sigma}_{a_j}(\cdot)\hat{\sigma}_{a'_j}(\cdot) \\
&+ \frac{\partial^2 V_i}{\partial a_i \partial z}\mu_z + \frac{\partial^2 V_i}{\partial a_i \partial K}\hat{\mu}_K(\cdot)
\end{aligned}$$

Note that all agents are not internalizing the price effect in the competitive equilibrium, which means there is no need to further differentiate assets' returns with respect to a_i. By plugging in all first order conditions, terms related to $\frac{\partial c_i(\cdot)}{\partial a_i}$, $\frac{\partial b_i}{\partial a_i}$ and $\frac{\partial \sigma_{a_i}(b_i,\cdot)}{\partial a_i}$ are canceled out. Rewrite the above equation in terms of marginal life-time utility $\xi_i = \frac{\partial V_i(a_i,\cdot)}{\partial a_i}$, the simplified expression of HJB after we take the first order derivative w.r.t a_i is:

$$\begin{aligned}
\rho\xi_i(a_i,\cdot) =& \frac{\partial\psi_{H(i)}(a_i,b_i)}{\partial a_i} + \frac{\partial\xi_i}{\partial a_i}\mu_{a_i}(a_i,c_i,b_i,\iota,\cdot) + \hat{r}\xi_i(a_i,\cdot) \\
&+ \frac{\partial\xi_i}{\partial z}\mu_z + \frac{\partial\xi_i}{\partial z}\hat{\mu}_K(\cdot) + \frac{1}{2}\frac{\partial^2\xi_i}{\partial a_i^2}\sigma_{a_i}^2(b_i,\cdot) \\
&+ \frac{1}{2}\frac{\partial^2\xi_i}{\partial z^2}\sigma_z^2 + \frac{\partial^2\xi_i}{\partial a_i\partial z}\sigma_{a_i}(b_i,\cdot)\sigma_z + \sum_{j\neq i}\frac{\partial^2\xi_i}{\partial a_i\partial a_j}\sigma_{a_i}(b_i,\cdot)\hat{\sigma}_{a_j}(\cdot) \\
&+ \sum_{j\neq i}\frac{\partial\xi_i}{\partial a_j}\hat{\mu}_{a_j}(\cdot) + \sum_{j\neq i}\frac{\partial^2\xi_i}{\partial a_j\partial z}\hat{\sigma}_{a_j}(\cdot)\sigma_z + \frac{1}{2}\sum_{j\neq i,j'\neq i}\frac{\partial^2\xi_i}{\partial a_j\partial a_{j'}}\hat{\sigma}_{a_j}(\cdot)\hat{\sigma}_{a'_j}(\cdot).
\end{aligned}$$

To further simplify the expression and make the connection to dynamics on the wealth share space. We consider the generalized Itô's lemma with jump for ξ_i. The expected drift part contains $\mu_{\xi_i}\xi_i$ which summarizes the drift by continuous process.

$$\begin{aligned}
\mu_{\xi_i}\xi_i =& \frac{\partial\xi_i}{\partial a_i}\mu_{a_i}(a_i,c_i,b_i,\iota,\cdot) + \sum_{j\neq i}\frac{\partial\xi_i}{\partial a_j}\hat{\mu}_{a_j}(\cdot) + \frac{\partial\xi_i}{\partial z}\mu_z + \frac{\partial\xi_i}{\partial z}\hat{\mu}_K(\cdot) \\
&+ \frac{1}{2}\frac{\partial^2\xi_i}{\partial a_i^2}\sigma_{a_i}^2(b_i,\cdot) + \frac{1}{2}\frac{\partial^2\xi_i}{\partial z^2}\sigma_z^2 + \frac{1}{2}\sum_{j\neq i,j'\neq i}\frac{\partial^2\xi_i}{\partial a_j\partial a_{j'}}\hat{\sigma}_{a_j}(\cdot)\hat{\sigma}_{a'_j}(\cdot) \\
&+ \frac{\partial^2\xi_i}{\partial a_i\partial z}\sigma_{a_i}(b_i,\cdot)\sigma_z + \sum_{j\neq i}\frac{\partial^2\xi_i}{\partial a_i\partial a_j}\sigma_{a_i}(b_i,\cdot)\hat{\sigma}_{a_j}(\cdot) + \sum_{j\neq i}\frac{\partial^2\xi_i}{\partial a_j\partial z}\hat{\sigma}_{a_j}(\cdot)\sigma_z
\end{aligned}$$

Still, given all the states at time t, the value of μ_{ξ_i} won't change if we switch to the wealth share space. Plug in the expression for μ_{ξ_i} , then we can get Euler equation in continuous-time as in equation (A.5).

To see the expression for risk-premium, we consider the volatility term, still as a scalar which does not vary over different state spaces, loading on the aggregate shock dW_t in Itô's lemma:

$$\sigma_{\xi_i}\xi_i = \frac{\partial\xi_i}{\partial a_i}(a_i - b_i)\sigma_q + \sum_{j\neq i}\hat{\sigma}_{a_j}(\cdot) + \frac{\partial\xi_i}{\partial z}\sigma_z(\cdot)$$

Plug it into the portfolio choice and we can get (A.6).

B Additional Details on The Algorithm

We start by reorganizing the set of equilibrium conditions to prepare the model for neural network training. We start with consumption and goods market clearing. Let

$\omega_i := c_i/a_i = c_i/(\eta_i q)$ and $\theta_i := b_i/a_i = b_i/(\eta_i q)$ denote the equilibrium consumption-to-wealth ratio and bond-to-wealth ratio for agent i. From the goods market clearing condition, q satisfies:

$$q = \frac{e^z K^{1-\alpha} L^{1-\alpha} + \Phi((\phi')^{-1} q^{-1})}{\sum_{i=1}^{I} \omega_i \eta_i}.$$

Individual SDFs can then be expressed as:

$$\xi_i = u'(\omega_i \eta_i q K), \text{ for } i \in \{1, 2, ..., I\},$$

We now combine the first order conditions for portfolio choice. Substituting equation (A.8) (the Ito's Lemma expansion of $\eta_j \sigma_{\eta_j}$) and equation (A.4) (the Ito's Lemma expansion of σ_ξ) into equation (A.6) (the agent portfolio choice first order condition) gives the equations:

$$\xi_i \left(\frac{r - r_k}{\sigma_q} \right) = \sum_{j<I-1} \frac{\partial \xi_i}{\partial \eta_j} \eta_j \sigma_{\eta_j} + \frac{\partial \xi_i}{\partial z} \sigma_z + \frac{1}{\sigma_q} \frac{\partial \psi_i(\eta_i q, -\eta_i^2 \sigma_{\eta_i} q/\sigma_q)}{\partial b_i}, \quad i = 1, \dots I$$

Rearranging and stacking the equations for $i = 1, \dots I$ gives:

$$-\sigma_z \begin{bmatrix} \frac{\partial \xi_1}{\partial z} \\ \vdots \\ \vdots \\ \frac{\partial \xi_I}{\partial z} \end{bmatrix} = \begin{bmatrix} \frac{\partial \xi_1}{\partial \eta_1} \eta_1 & \cdots & \frac{\partial \xi_1}{\partial \eta_{I-1}} \eta_{I-1} & \xi_1 \\ \frac{\partial \xi_2}{\partial \eta_1} \eta_1 & \cdots & \cdots & \xi_2 \\ \vdots & \cdots & \cdots & \vdots \\ \frac{\partial \xi_I}{\partial \eta_1} \eta_1 & \cdots & \frac{\partial \xi_I}{\partial \eta_{I-1}} \eta_{I-1} & \xi_I \end{bmatrix} \begin{bmatrix} \sigma_{\eta_1} \\ \vdots \\ \sigma_{\eta_{I-1}} \\ \frac{r_k - r}{\sigma_q} \end{bmatrix} + \frac{1}{\sigma_q} \begin{bmatrix} \frac{\partial \psi_1}{\partial b_1} \\ \frac{\partial \psi_2}{\partial b_2} \\ \vdots \\ \frac{\partial \psi_I}{\partial b_I} \end{bmatrix} \tag{B.1}$$

where the explicit dependence of ψ_i on σ_{η_i} has been suppressed. This can be written in matrix form in the following way:

$$-\sigma_z \frac{\partial \boldsymbol{\xi}}{\partial z} = M \begin{bmatrix} \boldsymbol{\sigma_\eta} \\ s \end{bmatrix} + \frac{1}{\sigma^q} \text{diag} \left(\frac{\partial \boldsymbol{\psi}}{\partial \boldsymbol{b}} \right) \tag{B.2}$$

where the vectors are $\boldsymbol{\xi} := [\xi_1, \dots, \xi_N]^T$, $\boldsymbol{\eta} := [\eta_1, \dots, \eta_{N-1}]^T$, $\boldsymbol{\sigma_\eta} := [\sigma_{\eta_1}, \dots, \sigma_{\eta_{N-1}}]^T$, $\boldsymbol{\psi} := [\psi_1, \dots, \psi_N]^T$, and $\boldsymbol{b} := [b_1, \dots, b_N]^T$, $s := \frac{r_k - r}{\sigma_q}$ is the Sharpe ratio, and M denotes the matrix:

$$M := \begin{bmatrix} \frac{\partial \boldsymbol{\xi}}{\partial \boldsymbol{\eta}} \odot \begin{bmatrix} \boldsymbol{\eta} \\ | \end{bmatrix} & \boldsymbol{\xi} \end{bmatrix}$$

Equation (B.2) shows the endogenous connection between agent wealth shares and the stochastic price process: agent portfolio decisions react to the price process in the economy and amplify the movement in the distribution. If ψ_i is linear in b_i, then equation (B.2) is a linear equation that can be solved explicitly for $[\boldsymbol{\sigma_\eta}, s]^T$.

Finally, we eliminate $(\iota, \boldsymbol{c}, \boldsymbol{b}, \mu_\xi, \sigma_\xi, \mu_K)$ from the equations in section A.2 by making the appropriate substitutions. This leaves the following system of equations. At state $\boldsymbol{X} = (z, K, (\eta_i)_{i\leq I})$, the equilibrium objects $(\boldsymbol{\xi}, q, \boldsymbol{\omega}, \boldsymbol{\sigma_\eta}, s, \sigma_q, \boldsymbol{\theta}, \mu_\eta, \mu_q, r)$ must satisfy the collection of equations:

$$
\begin{aligned}
0 &= (r-\rho_i)\xi_i + \frac{\partial \xi_i}{\partial z}\mu_z + \frac{\partial \xi_i}{\partial K}(\phi((\phi')^{-1}(q^{-1}))K_t - \delta K_t) \\
&\quad + \sum_j \frac{\partial \xi_i}{\partial \eta_j}\eta_j \mu_{\eta_j,t} + \sum_j \frac{\partial^2 \xi_i}{\partial z \partial \eta_j}\eta_j \sigma_{\eta_j,t}\sigma_z + \frac{1}{2}\frac{\partial^2 \xi_i}{\partial z^2}\sigma_z^2 \\
&\quad + \frac{1}{2}\sum_{j,j'} \frac{\partial^2 \xi_i^2}{\partial \eta_j \partial \eta_{j'}}\eta_j\eta_{j'}\sigma_{\eta_j,t}\sigma_{\eta_{j'},t} \\
q &= \frac{e^z K + \Phi((\phi')^{-1}(q^{-1}))}{\sum_{i=1}^I \omega_i \eta_i}, && \text{(B.3)} \\
\xi_i &= u'\left(\omega_i \eta_i q K\right), \ \text{for } i \in \{1, ..., I\}, \\
0 &= -\left[\frac{\partial \boldsymbol{\xi}}{\partial \boldsymbol{\eta}} \odot \begin{bmatrix} \boldsymbol{\eta} \\ | \end{bmatrix} \quad \boldsymbol{\xi}\right]\begin{bmatrix} \boldsymbol{\sigma_\eta} \\ s \end{bmatrix} - \sigma_z \frac{\partial \boldsymbol{\xi}}{\partial z} - \frac{1}{\sigma_q}\text{diag}\left(\frac{\partial \boldsymbol{\psi}}{\partial \boldsymbol{b}}\right) \\
q\sigma_q &= \sum_j \frac{\partial q}{\partial \eta_j}\eta_j \sigma_{\eta_j} + \frac{\partial q}{\partial z}\sigma_z \\
1-\theta_i &= -\frac{\eta_j \sigma_{\eta_j}}{\sigma_q}, \ \text{for } i \in \{1, ..., I\}, \\
K\mu_K &= \left(\phi\left((\phi')^{-1}\left(q^{-1}\right)\right)K - \delta K_t\right) \\
r_k - \mu_q &= \frac{e^z}{q} - \frac{(\phi')^{-1}\left(q^{-1}\right)}{q} + (\phi(\iota_t) - \delta) && \text{(B.4)} \\
\eta_i \mu_{\eta_i} &= r_k - \mu_q + \theta_i \sigma_q s - \mu_K - \omega_i + \theta_i \sigma_q^2, \ \text{for } i \in \{1, ..., I\} && \text{(B.5)} \\
q\mu_q &= \sum_j \frac{\partial q}{\partial \eta_j}\eta_j \mu_{\eta_j} + \frac{\partial q}{\partial z}\mu_z + \frac{\partial q}{\partial K}\mu_K + \sum_j \frac{\partial^2 \xi_i}{\partial z \partial \eta_j}\eta_j \sigma_{\eta_j}\sigma_z \\
r &= \sigma_q s + \frac{e^z}{q} - \frac{(\phi')^{-1}(q^{-1})}{q} + (\phi((\phi')^{-1}(q^{-1})) - \delta) + \mu_q && \text{(B.6)}
\end{aligned}
$$

Discussion: In general, working in the wealth space $\{a_j\}$ features an additional non-trivial fixed point problem as the wealth dynamics contain μ_q and price dynamics requires μ_{a_j}. Thus, jointly pinning down μ_{a_j} and μ_q requires a iterative scheme, as proposed in the computation part of Guvenen (2009). However, in the wealth share space the state dynamics do not depend on μ_q directly as implied by (B.5) and (B.4), due to the price effect does not affect shares' dynamics. Actually, the price's geometric drift only helps determine the risk free rate in (B.6) which enters into the Euler equation as part of the final loss. Following the execution order from (B.3) to (B.6) turns out to be critical.

C Three Testable Models

We compare neural network solution to analytical results (for complete market model) and finite difference solutions (for incomplete market models) solved by HJB equations.

C.1 Complete Market Model

We make the following modifications to map the model mentioned in section 2 to a Lucas Tree model. We set the capital share α to be one. We set both the capital depreciation rate δ and the capital conversion function to be zero. We fix the capital level K_t to be one and remove all penalty functions. To further simplify our notations, we introduce the output level $y_t = e^{z_t}$.

Without financial frictions, there is simple aggregation of individual's Euler equations as stated in main text, which coincides with the representative agent's pricing equation. Let us consider y's process follows the geometric Brownian motion's case:

$$dy_t = \mu y_t dt + \sigma y_t dW_t^0.$$

In representative agent's world, by standard Lucas tree pricing formula, asset price is determined by discounted flow of dividend:

$$q(y_0) = \mathbb{E}\left[\int_0^\infty e^{-\rho t}\frac{u'(c_t)}{u'(c_0)}y_t dt\right] = y_0\mathbb{E}\left[\int_0^\infty e^{-\rho t}(y_t/y_0)^{1-\gamma}dt\right]$$

Note that for geometric Brownian motion, the distribution of output is given by:

$$\ln(y_t/y_0) \sim \mathcal{N}\left((\mu - \frac{1}{2}\sigma^2)t, \sigma^2 t\right)$$

which means (the integral and expectation operator are interchangeable):

$$\begin{aligned}\mathbb{E}(y_t/y_0)^{1-\gamma} &= (1-\gamma)(\mu - \frac{1}{2}\sigma^2)t + \frac{1}{2}(1-\gamma)^2\sigma^2 t \\ &= (1-\gamma)\mu t + \frac{1}{2}(\gamma-1)\gamma\sigma^2 t \\ &\equiv -\breve{g}t\end{aligned}$$

Therefore, asset prices are given by:

$$q(y_0) = y_0 \int_0^\infty e^{-\rho t} e^{-\breve{g}t} dt = \frac{y_0}{\rho + \breve{g}} = \frac{y_0}{\rho + (\gamma-1)\mu - \frac{1}{2}\gamma(\gamma-1)\sigma^2}$$

By goods market clearing condition, we know that $c_t = y_t$, which means the consumption policy is:

$$c = \left[\rho + (\gamma-1)\mu - \frac{1}{2}\gamma(\gamma-1)\sigma^2\right] q$$

For $\gamma = 5, \mu = 0.02, \sigma = 0.05, \rho = 0.05$ in the numerical example, $c/q = 10.5\%$, which means: $q(1) = 1/10.5\% \approx 9.5$.

Though aggregation results hold, we still incorporate the wealth heterogeneity and solve by our algorithm. Note that the instant risk allocation is determined by simple matrix inversion from (B.1) and there's no other unknowns for price's risk consistency, it is unnecessary to parameterize σ_q. We find that our solution aligns with the "as-if" representative agent's solution quite well. The estimated time cost for model with 5 agents is about 2 mins, 10 agents is about 10 mins and 20 agents is about 20 mins. The difference between consumption rule solved neural network and analytical solution is less than 0.1% (for 5, 20 agents)/ 0.5% (for 20 agents).

Num of Agents	Euler Eq Error	Diff	Time Cost
5	<1e-4	<0.1%	2 mins
10	<1e-4	<0.5%	10 mins
20	<1e-3	<0.5%	20 mins

Table 5: Summary of the algorithm performance and computational speed. "Diff" means the difference between representative agent case's solution and brute-force. All errors are in absolute value (L1 loss).

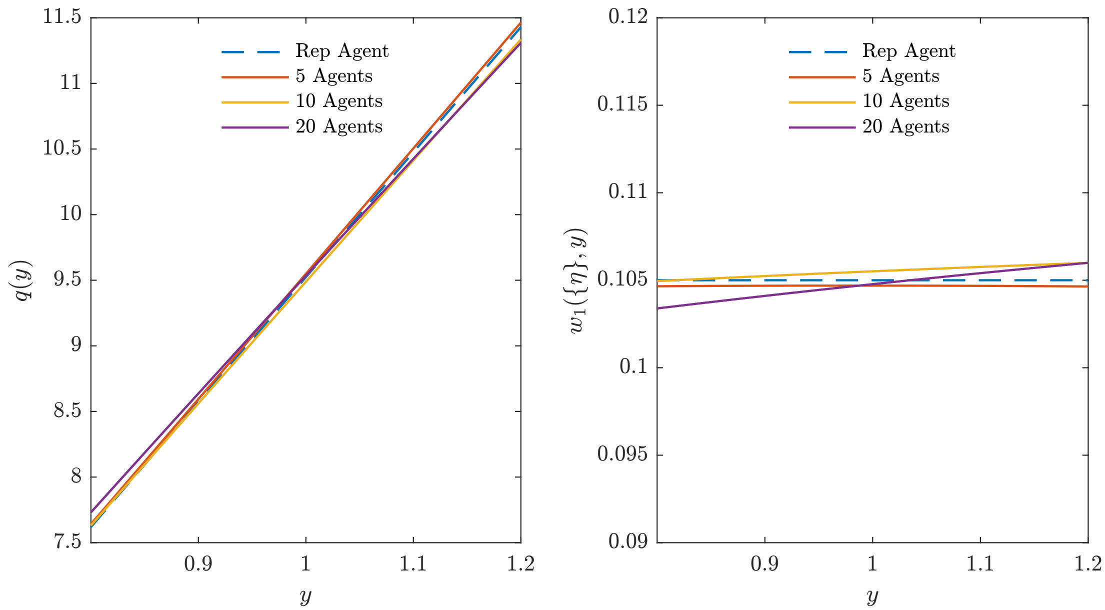

Figure 12: Solution to As-if representative agent model. Right panel: consumption-wealth ratio of agent 1.

C.2 Asset Pricing with Restricted Participation

We still adopt the modifications that are done in the first subsection to mimic the endowment economy. There are two price taking agents in this infinite horizon economy: expert and household. The financial friction we use is that household cannot participate the stock market. Mathematically, it is stated as:

$$\Psi_i(a_i, b_i) = -\frac{\bar{\psi}_i}{2}(a_i - b_i)^2, \bar{\psi}_h = \infty, \bar{\psi}_e = 0.$$

Again, the output y_t follows a geometric Brownian motion:

$$dy_t = \mu y_t dt + \sigma dZ_t.$$

Boundary Conditions. We focus on the case that $\eta \in (0, 1]$, as the economy is ill-defined when experts are wiped out from the economy, i.e., nobody holds the tree in equilibrium. To get the right boundary, we use the asset prices and consumption policy ω^e from the representative agent's solution:

$$\omega_e(1, y) = \rho_e + (\gamma - 1)\mu - \frac{1}{2}\gamma(\gamma - 1)\sigma^2, q(1, y) = \frac{y}{\omega_e(1, y)}.$$

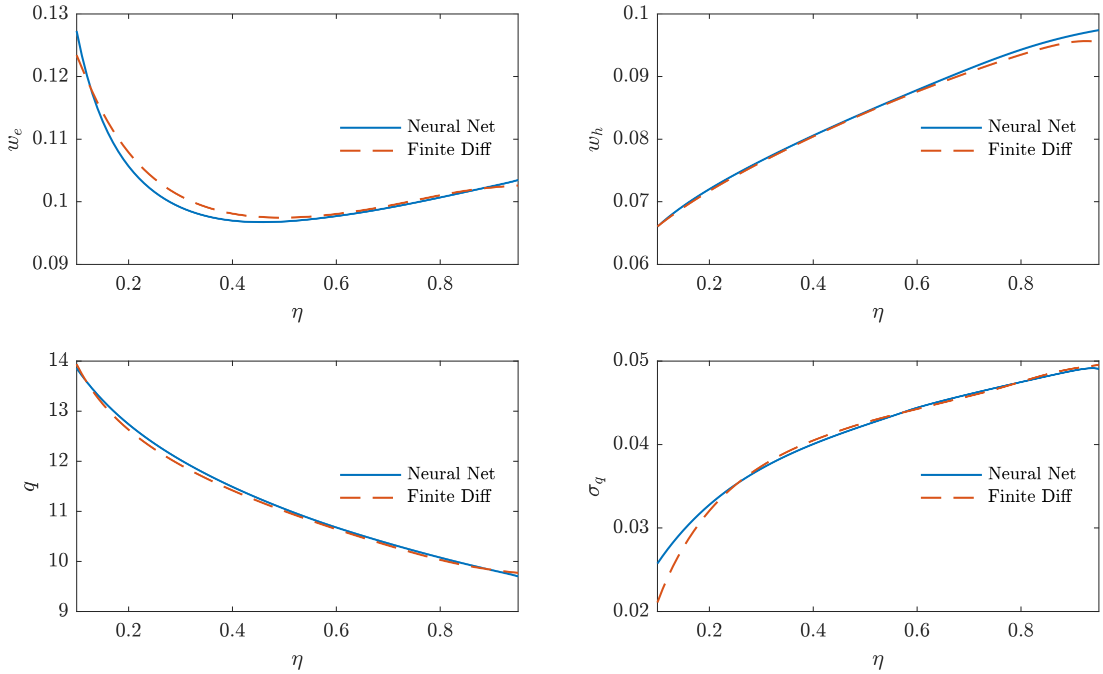

Figure 13: Solution to restricted stock market participation model.

Model Solution. The estimated time to solve the limited participation problem by neural network is about 5 minutes. We compare the finite difference solution (technical details can be found from the appendix) with the neural network solution on η's dimension in figure 13 for $y = 1$. We can see that neural network well captures the high non-linearity (left-upper panel) and amplification (right-lower panel) by high risk-aversion.

C.3 A Macroeconomic Model with Productivity Gap

The setup follows Brunnermeier and Sannikov (2016). There are two types of agents in this infinite horizon economy: experts and households. We allow households to hold capitals but in a less productive way. The productivity of experts and households is z_h, z_e ($z_h < z_e$) respectively. Their relative risk-aversion are both γ. Output grows at exogenous drift $\mu_y = y\mu$, volatilty $y\sigma$, and experts cannot issue outside equities. In addition, we assume there's a constraint for no short-selling from households' side, which can be formally written as:

$$\begin{cases} \Psi_h(a_h, b_h) = -\dfrac{\bar{\psi}_h}{2}(\min\{a_h - b_h, 0\})^2, & \bar{\psi}_h = \infty \\ \Psi_e(a_e, b_e) = -\dfrac{\bar{\psi}_e}{2}(a_e - b_e)^2, & \bar{\psi}_e = 0. \end{cases}$$

The output flow on households' side and experts' side can be written as:

$$d_{e,t} = z_e y_t, d_{h,t} = z_h y_t, dy_t = y_t \mu dt + y_t \sigma dZ_t$$

The capital return from households' side and experts' side:

$$r_{q,e,t} = \frac{d_{e,t}}{q_t} + \mu_{q,t}, r_{q,h,t} = \frac{d_{h,t}}{q_t} + \mu_{q,t}.$$

We could rewrite the financial friction as return's gap: $\frac{a_e - a_h}{q\sigma^q}$. For the first two equations, we have:

$$\begin{cases} -\dfrac{1}{\xi_e}\dfrac{\partial \xi_e}{\partial y}\sigma_y = \dfrac{1}{\xi_e}\dfrac{\partial \xi_e}{\partial \eta}\sigma_\eta - \dfrac{r_f - r_{q,h}}{\sigma_q} + \dfrac{y_e - y_h}{q\sigma_q} \\ -\dfrac{1}{\xi_h}\dfrac{\partial \xi_h}{\partial y}\sigma_y = \dfrac{1}{\xi_h}\dfrac{\partial \xi_h}{\partial \eta}\sigma_\eta - \dfrac{r_f - r_{q,h}}{\sigma_q} + 0 \end{cases} \Leftrightarrow \boldsymbol{n} = \mathbf{M}\begin{bmatrix} \sigma_\eta \\ \frac{r_f - r_{q,h}}{\sigma^q} \end{bmatrix} + \underbrace{\begin{bmatrix} \frac{y_e - y_h}{q\sigma_q} \\ 0 \end{bmatrix}}_{\partial_2 \boldsymbol{\psi}}$$

The main difficulty for Brunnermeier and Sannikov (2016)'s model is that we need to **preserve** computational graph when output is a function of risk allocation, which means resorting to non linear solver, as in Gopalakrishna (2021), is not applicable here. The algorithm in section 3 still applies here, however. Compared to the previous two examples, we have to parameterize only one more equilibrium object, because of the closed form relationships between the equilibrium objects. In practice, we introduce the auxiliary neural network for the capital allocation (or say, the output function), which turned to be most efficient, $\kappa = \eta + \lambda = \eta + \mathcal{N}_\lambda \eta^\beta$, where $\mathcal{N}_\lambda$ is a trainable neural net and β is solved from the asymptotic solution for $\eta \to 0$. Such parameterization effectively captures the high non-linearity as η goes to zero.

Model Solution. The estimated time to solve the model by neural network is about 5 minutes. Again, we compare the finite difference solution with neural network solution in figure 14 for $y = 1$. We set up the range of η to be the crisis region in Brunnermeier and Sannikov (2016), which is defined by inefficient capital allocation as $\kappa < 1$. We can see that the neural network solution well captures most of the amplification in that crisis region, despite the volatility gap between finite difference solution and neural network's when $\eta \to 0$, which is not quantitatively relevant because of the negligible amount of time the economy spends in this deep crisis region. Matching such extremely high non-linearity as η goes to be very close to zero has already been studied well in Gopalakrishna (2021) and is beyond the scope of our paper.

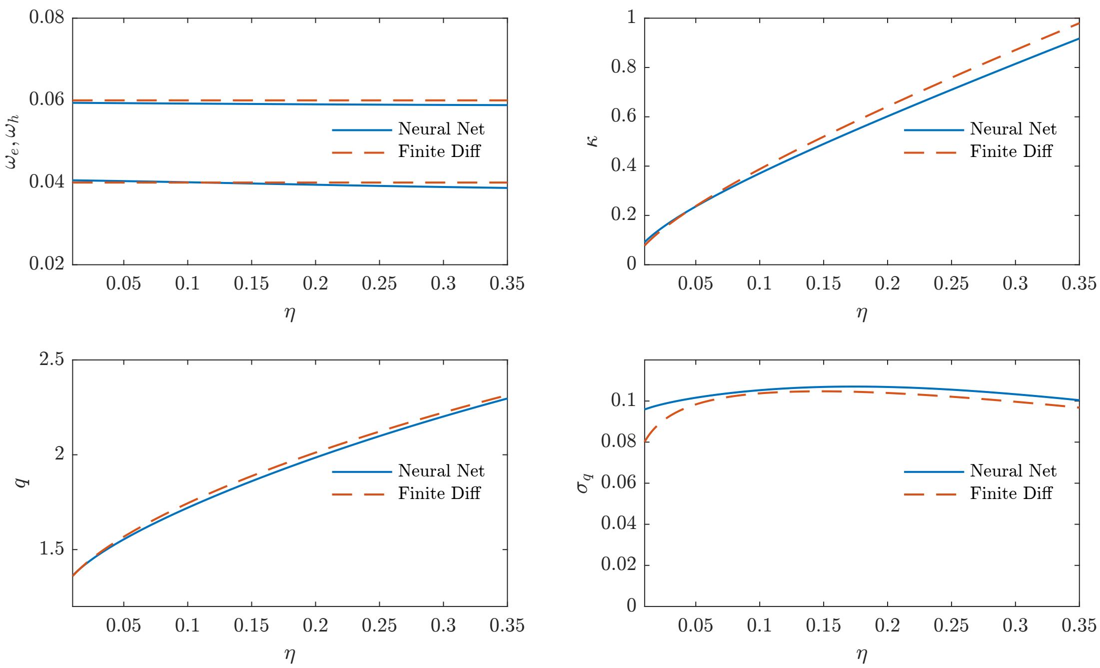

Figure 14: Solution to the model with productivity gap.

D Finite Difference Solutions

We exploit the scalability, as in textbook Campbell and Viceira (2002), for geometric Brownian motion's case to get a preciser solution by focusing only on one dimensional differential equation. For scalable income process, we postulate the price function as: $q = f(\eta)y$, where η is the expert's wealth share with no loss of generality, i.e., $\eta = \eta_1$. The value function can be written as:

$$V_i = \frac{1}{\rho_i} \frac{(\omega_i \eta_i q)^{1-\gamma}}{1-\gamma} = \frac{(\omega_i \eta_i f(\eta))^{1-\gamma}}{\rho_i} \frac{y^{1-\gamma}}{1-\gamma} \equiv v_i \frac{y^{1-\gamma}}{1-\gamma},$$

where v_i can be viewed as the value function on η's space only. From the first order condition[7]:

$$c_i^{-\gamma} = \frac{1}{\rho_i}\frac{(\omega_i\eta_i q)^{1-\gamma}}{\eta_i q} \Rightarrow \left(\frac{c_i}{y}\right)^{\gamma} = \frac{\eta_i f(\eta)}{v_i}, \omega_i = [\eta_i f(\eta)]^{\frac{1}{\gamma}-1} v_i^{-\frac{1}{\gamma}} \tag{D.1}$$

From the goods market clearing conditon, we have:

$$1 = \frac{\sum_i c_i}{y} = \sum_i \left(\frac{\eta_i f(\eta)}{v_i}\right)^{\frac{1}{\gamma}} = y \Rightarrow f(\eta) = \frac{1}{\left[\sum_i \left(\frac{\eta_i}{v_i}\right)^{\frac{1}{\gamma}}\right]^{\gamma}} \tag{D.2}$$

The HJB for scaled value functionv_i (note: for $y^{1-\gamma}$ which appears in V, we still need to take the Itô's lemma on it)

$$[\rho_i - (1-\gamma)\mu + \frac{\gamma}{2}(1-\gamma)\sigma^2 - \omega_i]v_i = [\mu_\eta + (1-\gamma)\sigma\sigma_\eta]\eta\frac{\partial v_i}{\partial \eta} + \frac{1}{2}\frac{\partial^2 v_i}{\partial \eta^2}\eta^2\sigma_\eta^2 \tag{D.3}$$

where μ_η, σ_η are from (A.7) and (A.8). The price of risk which appears in the asset pricing condition is determined by Itô's Lemma:

$$\xi_i = \frac{v_i}{\eta_i f(\eta)} y^{-\gamma} \Rightarrow \sigma_\xi = \sigma_v - \sigma_f - \sigma_\eta - \gamma\sigma = \frac{v_i'(\eta)\eta\sigma_\eta}{v_i} - \frac{f'(\eta)\eta\sigma_\eta}{f} - \sigma_\eta - \gamma\sigma.$$

In finite difference, we introduce the pseudo time-steps (D.3):

$$[\rho_i - (1-\gamma)\mu + \frac{\gamma}{2}(1-\gamma)\sigma^2 - \omega_i]v_i = [\mu_\eta + (1-\gamma)\sigma\sigma_\eta]\eta\frac{\partial v_i}{\partial \eta} + \frac{1}{2}\frac{\partial^2 v_i}{\partial \eta^2}\eta^2\sigma_\eta^2 + \frac{\partial v_i}{\partial t},$$

and update value function in an implicit scheme to solve equation

$$\check{\boldsymbol{\rho}}\mathbf{I}\mathbf{v}_{t+dt} = \mathbf{M}\mathbf{v}_{t+dt} + \frac{\mathbf{v}_{t+dt} - \mathbf{v}_t}{dt},$$

where $\mathbf{M}$ is the differential matrix by upwind scheme, and I is the identity matrix.

[7] This expression leads to the boundary condition at $\eta = 1$: $\frac{f(1)}{v_e} = 1$

D.1 Solution to the Limited Participation Model

The distributional dynamics for limited participation model are:

$$\mu_\eta = (1-\eta)(\omega_h - \omega_e) + \left(-\frac{1-\eta}{\eta}\right)(r_f - r_q + (\sigma_q)^2)$$
$$\sigma_\eta = \frac{1-\eta}{\eta}\sigma_q, \text{where } r_f - r_q = \sigma_\xi \sigma_q.$$

By the consistency condition for price volatility, we have:

$$f(\eta)y\sigma_q = f'(\eta)y\sigma_\eta + f(\eta)\sigma y \to \sigma^q = \frac{\sigma}{1 - \frac{f'(\eta)}{f(\eta)}(1-\eta)}.$$

The boundary conditions: $f(1) = \frac{1}{\rho_e + (\gamma-1)\mu - \frac{1}{2}\gamma(\gamma-1)\sigma^2}$, $v_e(1) = f(1)$.

Algorithm. Set up grids: $\eta_n = linspace(\Delta\eta, 1-\Delta\eta, 1/\Delta\eta - 1)$. Initialize the value function as $v_{i,0}(\cdot) = \rho_i + (\gamma-1)\mu - \frac{1}{2}\gamma(\gamma-1)\sigma^2$.
While $Error > \epsilon$:

1. Compute $\omega_e, \omega_h, f(\eta)$ by equation (D.1), (D.2).
2. Compute $\frac{dq}{d\eta}, \frac{dv_e}{d\eta}, \frac{dv_h}{d\eta}$ by upwind scheme, use the boundary condition if $\mu_{1-\Delta\eta} > 0$ required.
3. Construct the terms in HJB. Then update $v_{i,t+dt}$ by implicit scheme.
4. Compute $Error = |v_{e,t+dt} - v_{e,t}| + |v_{h,t+dt} - v_{h,t}|$

D.2 Solution to the Macroeconomic Model with a Financial Sector

Given the expert's capital share holding κ, the wealth share η's risk σ_η is $(\kappa - \eta)\sigma_q$. The goods market clearing condition (D.2) is replaced by:

$$f(\eta) = \frac{\kappa\eta + (1-\kappa)(1-\eta)}{\left[\sum_i \left(\frac{\eta_i}{v_i}\right)^{\frac{1}{\gamma}}\right]^\gamma}$$

By the consistency condition for price volatility, we have:

$$f(\eta)y\sigma_q = f'(\eta)y\sigma_\eta + f(\eta)\sigma y \to \sigma_q = \frac{\sigma}{1 - \frac{f'(\eta)}{f(\eta)}(\kappa-\eta)}$$

The boundary conditions are $f(0) = \frac{a_h}{\omega_h(0)}, f(1) = \frac{a_e}{\omega_e(1)}$.

Algorithm. Set up grids: $\eta_n = linspace(\Delta\eta, 1-\Delta\eta, 1/\Delta\eta - 1)$. Initialize the value function as $v_{i,0}(\cdot) = \rho_i + (\gamma-1)\mu - \frac{1}{2}\gamma(\gamma-1)\sigma^2$.
While $Error > \epsilon$:

1. Compute ω_e, ω_h by equation (D.1).
2. Approximate $f'(\eta)$ by finite difference. For $\eta = \Delta\eta : \Delta\eta : 1-\Delta\eta$, solve $(f(\eta), \kappa, \sigma_q)$ from the following set of equations: (1) if $\kappa < 1$

$$\begin{cases} \rho_e\omega_e\eta + \rho_h\omega_h(1-\eta) = \kappa z_e + (1-\kappa)z_h \\ \sigma_q = \dfrac{\sigma}{1 - \frac{f'(\eta)}{f(\eta)}(\kappa-\eta)} \\ \dfrac{z_e - z_h}{q} = \dfrac{\kappa-\eta}{\eta(1-\eta)}\sigma_q^2. \end{cases} \tag{D.4}$$

   (2) if $\kappa > 1$, set κ to be 1, then only solve q, σ_q from the first two equations in (D.4).

3. Compute $\frac{dv_e}{d\eta}, \frac{dv_h}{d\eta}$ by upwind scheme.
4. Construct the terms in HJB. Then update $v_{i,t+dt}$ by implicit scheme.
5. Compute $Error = |v_{e,t+dt} - v_{e,t}| + |v_{h,t+dt} - v_{h,t}|$

E Proofs for Section 4

Approximate consumption-wealth ratio under log-utility. We now investigate if the constant consumption-to-wealth ratio satisfies the Euler equation. Due to the participation frictions, net worth share dynamics affect the risk premium and in turn affects the risk free rate. We assume constant consumption wealth ratio $c_i/a_i = \rho_i$, then by Itô's Lemma

$$\begin{aligned} \mathbb{E}\frac{d}{dt}[1/(\rho_i\eta_i q)] = & -\frac{1}{\rho_i\eta_i q^2}q\mu_q - \frac{1}{\rho_i\eta_i^2 q}\eta_i\mu_{\eta,i} + \frac{1}{\rho_i\eta_i^3 q}\sigma_{\eta,i}^2\eta_i^2 \\ & + \frac{1}{\rho_i\eta_i q^3}\sigma_q^2 q^2 + \frac{1}{\rho_i\eta_i^2 q^2}\sigma_{\eta,i}\eta_i\sigma_q q, \end{aligned}$$

which means the geometric drift of SDF can thus be expressed as:

$$\mu_{\xi,i} = -\mu_{\eta,i} - \mu_q + \sigma_{\eta,i}\sigma_q + \sigma_{\eta,i}^2 + \sigma_q^2.$$

The left hand side of the Euler equation then can be expressed as:

$$\begin{aligned}
-\rho_i + r_f + \mu_{\xi,i} &= -\rho_i + r_f - \mu_{\eta,i} - \mu_q - \sigma_{\eta,i}\sigma_q + \sigma_{\eta,i}^2 + \sigma_q^2 \\
&= -\rho_i + r_f + \underbrace{\left(-\frac{y}{q} - \theta_i(r_f - r_q + \sigma_q^2) + \rho_i\right)}_{=\mu_{\eta,i}} - \mu_q + \sigma_{\eta,i}\sigma_q + \sigma_{\eta,i}^2 + \sigma_q^2 \\
&= (r_f - r_q + \sigma_q^2) - \theta_i(r_f - r_q + \sigma_q^2) + \theta_i^2\sigma_q^2 - \theta_i\sigma_q^2 \\
&= (1-\theta_i)(r_f - r_q + \sigma_q^2(1-\theta_i))
\end{aligned}$$

For experts, their asset pricing equation implies $r_f - r_q = (\theta_N - 1)\sigma_q^2$, which means the Euler equation is satisfied with the ansatz that consumption wealth ratio is ρ_e.

For households, we plug in the households' asset pricing conditions $r_f - r_q = (\theta_i - 1)\sigma_q^2 + \frac{k}{\eta_i}(\theta_i - 1)$, and we can see that the LHS of Euler equation equals:

$$-(1-\theta_i)^2\frac{k}{\eta_i}.$$

Though mathematically the assumption for households $c_i/a_i = \rho_i$ is not self-contained, the loss for Euler equation is extremely small, because of the fact that k is same order of magnitude of σ_q^2. Therefore, L2 Loss of Euler equation is:

$$L_{Euler,i}^2 = \mathcal{O}(\sigma_q^4) < \mathcal{O}(\sigma^4) \sim 10^{-6}.$$

$\square$

Experts are levered. We first use capital market clearing condition to solve the risk premium:

$$\left(\sum_{i=1}^{I-1}\frac{\eta_{i,t}^2}{\bar{\psi} + \sigma_{q,t}^2\eta_i} + \frac{\eta_{I,t}}{\sigma_{q,t}^2}\right)(r_q - r_f) = 1.$$

Therefore, expert's portfolio share on capital is solved as:

$$\theta_I = \frac{1}{\sigma_q^2} \frac{1}{\frac{\eta_I}{\sigma_q^2} + \sum_{i=1}^{I-1} \frac{\eta_i^2}{\bar{\psi}+\eta_i\sigma_q^2}},$$

which is greater than one since:

$$\frac{1}{\sigma_q^2} \frac{1}{\frac{\eta_I}{\sigma_q^2} + \sum_{i=1}^{I-1} \frac{\eta_i^2}{\bar{\psi}+\eta_i\sigma_q^2}} > \frac{1}{\sigma_q^2} \frac{1}{\frac{\eta_I}{\sigma_q^2} + \sum_{i=1}^{I-1} \frac{\eta_i^2}{\eta_i\sigma_q^2}} = \frac{1}{\sigma_q^2} \frac{\sigma_q^2}{\eta_I + \sum_{i=1}^{I-1} \eta_i} = 1$$

□

Proof of Theorem 1. As experts are levered, we know $1/\sigma_q$ is uniformly bounded. We first consider the case that $\bar{\psi} \to 0$. It can be shown that:

$$\lim_{\bar{\psi}\to 0} \frac{\frac{\eta_i}{\bar{\psi}+\eta_i\sigma_q^2}}{\frac{\eta_I}{\sigma_q^2} + \sum_{i=1}^{I-1} \frac{\eta_i^2}{\bar{\psi}+\eta_i\sigma_q^2}} = \frac{\frac{1}{\sigma_q^2}}{\frac{1}{\sigma_q^2}} = 1,$$

as household's portfolio choice is continuous function of $\bar{\psi}$. On the other extreme $\bar{\psi} \to \infty$, recall that we know $\sigma_q \leq \sigma$ by the expression for σ_q. Taking the limit we know $\theta_i = 0$, for $i \in \{1, ..., I\}$. In these two extreme cases, portfolio differences $(\theta_{j,t} - \theta_{i,t})$ are both zero.

To understand the price volatility dynamics, we rewrite the derivatives $\partial q/\partial \eta_i$ by chain rule:

$$\frac{\partial q}{\partial \eta_i} = \frac{\partial q}{\partial \eta_h} \frac{\partial \eta_h}{\partial \eta_i},$$

where η_h is $\sum_{i=1}^{I-1} \eta_i$, total household wealth share. The consistency condition for capital volatility process can be simplified as

$$q\sigma_q = q\sigma + \frac{\partial q}{\partial \eta_h}\sigma_q \sum_{i=1}^{I-1}(-\theta_i)\eta_i = q\sigma + \frac{\partial q}{\partial \eta_I}(-\theta_I\eta_I)\sigma_q$$

where we applied the bonds market clearing condition. Combine with the intermediary's asset pricing condition: $r_f - r_q = -\sigma_q^2\theta_I$, we know that aggregate households' wealth share is sufficient for asset pricing. □

Proof of Theorem 2. We characterize the asymptotic solution for price volatility and experts' leverage given a low η_N, with different level of participation friction. We first show that, given participation friction k, as $\eta \to 0$, the price volatility will converge to fundamental volatility. Recall that the expert's capital-wealth ratio is given by:

$$\theta_I = \frac{1}{\sigma_q^2} \frac{1}{\frac{\eta_I}{\sigma_q^2} + \sum_{i=1}^{I-1} \frac{\eta_i^2}{\bar{\psi}+\eta_i\sigma_q^2}}.$$

Since $\frac{\eta_i^2}{\bar{\psi}+\eta_i\sigma_q^2}$ is a convex function, given the expert's sector wealth level $\eta_I = 1 - \sum_{i=1}^{I-1} \eta_i$, we know that:

$$(I-1)\frac{\bar{\eta}_h^2}{\bar{\psi}+\bar{\eta}_h\sigma_q^2} \leq \sum_{i=1}^{I-1} \frac{\eta_i^2}{\bar{\psi}+\eta_i\sigma_q^2} \leq \frac{(1-\eta_I)^2}{\bar{\psi}+(1-\eta_I)\sigma_q^2},$$

where $\bar{\eta}_h$ is defined as $\frac{1-\eta_I}{I-1}$, and σ_q is still determined by households' wealth distribution $\{\eta_1, ..., \eta_{I-1}\}$. Using the inequality above, we know that:

$$\frac{1}{\sigma_q^2} \frac{1}{\frac{\eta_I}{\sigma_q^2} + (1-\eta_I)\frac{(1-\eta_I)}{\bar{\psi}+(1-\eta_I)\sigma_q^2}} \leq \theta_I \leq \frac{1}{\sigma_q^2} \frac{1}{\frac{\eta_I}{\sigma_q^2} + (1-\eta_I)\frac{\bar{\eta}_h}{k+\bar{\eta}_h\sigma_q^2}},$$

which can be simplified as:

$$\frac{\frac{\bar{\psi}}{\sigma_q^2} + (1-\eta_I)}{\frac{\bar{\psi}}{\sigma_q^2}\eta_I + (1-\eta_I)} \leq \theta_I \leq \frac{\frac{(I-1)\bar{\psi}}{\sigma_q^2} + (1-\eta_I)}{\frac{(I-1)\bar{\psi}}{\sigma_q^2}\eta_I + (1-\eta_I)}.$$

To show that the endogenous price volatility is exactly the fundamental volatility, we consider the range of $-\theta_N\eta_I$. We deduct both side by 1 and multiply both side by η_I:

$$\frac{\frac{\bar{\psi}}{\sigma_q^2}(1-\eta_I)}{\frac{\bar{\psi}}{\sigma_q^2}\eta_I + (1-\eta_I)}\eta_I \leq \eta_I(\theta_I - 1) \leq \frac{\frac{(I-1)\bar{\psi}}{\sigma_q^2}(1-\eta_I)}{\frac{(I-1)\bar{\psi}}{\sigma_q^2}\eta_I + (1-\eta_I)}\eta_I.$$

Note that $1/\sigma_q$ is uniformly bounded, when taking the limit $\eta_I \to 0$, we know $\lim_{\eta_I \to 0} \theta_N\eta_I = 0$ by the squeeze theorem. Thus, the term $\frac{\partial q}{\partial \eta_I}(-\theta_I\eta_I)\sigma_q$ is zero at the limit $\eta_I = 0$.

Next, for a given small expert wealth share η_I the range for debt wealth ratio $1-\theta_N$ is:

$$-\frac{\bar{\psi}}{\sigma^2} \geq 1 - \theta_I \geq -\frac{(I-1)\bar{\psi}}{\sigma^2}$$

On one extreme, when the inequality within households is high, the expert's leverage ratio is $\frac{\bar{\psi}}{\sigma^2}$. On the other extreme, inequality within households is low, the expert's leverage ratio is $\frac{(I-1)\bar{\psi}}{\sigma^2}$. These two extreme cases shed lights on how household's side inequality determines total household sector's capital share which in turn shapes the intermediary risk taking and affects negatively on its recovery condition on recessions. The proof for generic household distribution is straight forward by the convexity of household's capital taking function. $\square$

F Parameters for Testable Models

F.1 Economic Parameters

F.1.1 Parameters for the "as-if" Complete Market Model

Parameter	Symbol	Value
Risk aversion	γ	5.0
Agents' Discount rate	ρ	0.05
Output Growth Rate	μ	2%
Volatility of Growth	σ	5%

F.1.2 Parameters for the Limited Participation Model

Parameter	Symbol	Value
Risk aversion	γ	5.0
Households' Discount rate	ρ_h	0.05
Experts' Discount rate	ρ_h	0.05
Output Growth Rate	μ	2%
Volatility of Growth	σ	5%

F.1.3 Parameters for the Macroeconomic Model with a Financial Sector

Parameter	Symbol	Value
Risk aversion	γ	1.0
Households' Discount rate	ρ_h	0.04
Experts' Discount rate	ρ_e	0.06
Households' Productivity	z_e	0.11
Experts' Productivity	z_h	0.05
Output Growth Rate	μ	2%
Volatility of Growth	σ	5%