

Asset Pricing, Participation Constraints, and Inequality*

Goutham Gopalakrishna[†], Zhouzhou Gu[‡], Jonathan Payne[§]

November 8, 2024

Abstract

How does the regulation of the financial sector impact household welfare and inequality? To answer this question, we build a macroeconomic model with banks, insurance/pension funds, heterogeneous households, 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 how asset price dynamics, financial stability, and wealth inequality all depend upon which investors are able to purchase assets in bad states of the world. In particular, allowing pension/insurance funds broad access to asset markets leads to greater stability at the business cycle frequency but exposes the household to other risks. Ultimately, the government faces complicated trade-offs between ensuring stability, lowering borrowing costs, and maintaining household equality.

Keywords: Inequality, Market Segmentation, Asset Pricing, Heterogeneous Agent Macroeconomic Models, Deep Learning.

JEL: C63, C68, E27, G12, G21, G22, G23

*We are very grateful to the comments and discussion from Jaroslav Borovička, Anmol Bhandari, Markus Brunnermeier, Paul Ehling, Lars Hansen, Felix Kubler, Benjamin Moll, Michael Reiter, Martin Schneider, Yucheng Yang, Moto Yogo, Jan Žemlička, and Adam Zhang. We are also thankful for the comments from participants in the Blue Collar Working Group at the University of Chicago, the Workshop on the Frontier of Quantitative Macroeconomics: Methods and Models, the 2024 SED, and NBER 2024 Summer Institute.

[†]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

US policy makers have segmented the financial sector very differently throughout US history. In the second half of the nineteenth century, regulation allowed insurance companies to hold a wide range of risky assets but restricted banks to holding short-term low risk loans. This was reversed in the 1930s when New Deal regulation started to encourage banks to hold long-term assets and restrict insurance companies to holding long-term government or corporate bonds that hedge their duration risk. A recent finance literature has shown that these institutional constraints are important for explaining asset pricing because they bind in complex ways for different financial intermediaries and so generate price “over reaction” or “under reaction” to shocks that generate additional risk premia in the economy (e.g. [Kojien and Yogo \(2019, 2023\)](#)). Understanding how these price reactions impact different households is a priority for policy makers. However, macroeconomists have struggled to address this question because they have typically been unable to characterize global solutions to heterogeneous agent models with aggregate risk, endogenous price volatility, and complicated portfolio restrictions. In this paper, we develop a deep learning algorithm to overcome the technical difficulties and allow researchers to bring insights from the asset pricing literature into the heterogeneous agent macroeconomic literature. We use our model to examine how the regulation of the financial sector impacts household inequality.

We study a heterogeneous agent business cycle model with financial intermediaries and endogenous asset price volatility. There are two aggregate shock processes in the economy, capital productivity and the volatility of capital productivity, both of which follow mean reverting processes. The economy is populated by price-taking households who have inelastic or restricted demand for different assets. Households face idiosyncratic death shocks, which leads to demand for long duration insurance or pension products, similar to the “preferred habitat” demand in [Vayanos and Vila \(2021\)](#). Households also face a penalty for holding capital that is relaxed as they get wealthier. This leads to heterogeneous household portfolio choices, with wealthier households having greater exposure to high return assets, and ultimately, a non-degenerate household wealth distribution. The economy also contains two types of financiers who provide financial services to households: bankers who issue deposits and fund managers who issue insurance/pension products. Neither type of financier

can raise equity.¹ Finally, there is a government that issues long term bonds and raises wealth taxes. The financial friction on the households and financial intermediaries mean that the distribution of wealth across the agents impacts the capital, insurance, and government bond price processes.

In Section 4, we illustrate the different economic mechanisms at play in our model. We start by analyzing the economy at the sector level to understand how intermediary asset pricing is working. As in many macro-finance models, the fundamental friction in our environment is that the banking sector does not internalize their price impact and so takes on high leverage, which leads to endogenous asset price volatility (and ultimately household consumption volatility). We show that, without financial regulation, the insurance/pension fund sector partially mitigates this behaviour by acting as a “backstop” for the economy during recessions by purchasing capital from the banking sector. This helps to stabilize the economy at the business cycle frequency but does not necessarily stabilize household wealth and consumption with respect to volatility shocks. To understand this, consider the different risk exposure of the bank and fund liabilities. The bank issues short-term risk-free deposits while the fund issues long-term, tradable, insurance products that depreciate in recessions (because the marginal utility of consumption is high) but can depreciate or appreciate during times of high volatility (because high high volatility also encourages households to rebalance towards insurance products). This means that the market value of fund liabilities decreases during recessions and increases during volatility spikes, giving the funds a natural hedge against business cycle shocks but greater exposure to volatility shocks. Contemporary regulation limits funds’ ability to play this role by restricting their participation in capital markets and incentivizing them to hold government and corporate bonds. In our economy, this destabilizes the banking sector at the business cycle frequency but increases both the level and stability of government bond prices by creating a captive market for government debt. In this sense, by choosing how they restrict the fund sector, the government can choose where the endogenous volatility appears in the economy.

We next investigate the different mechanisms connecting asset pricing and inequality. The evolution of wealth inequality in our model can be decomposed into the following forces: (i) wealthier agents have more access to capital markets, which allows them

¹The model can be extended naturally to an environment where financial intermediaries can raise a limited quantity of outside equity.

to access higher returns and build wealth more quickly (sometimes referred to as the “scaling” force) but also increases their exposure to aggregate risk, (ii) poorer agents rely more on the pension/insurance products, which earn a lower return (an amplification of the “scaling” force in the literature), (iii) wealthier agents have a higher propensity to consume out of wealth, and (iv) taxes and transfers that redistribute wealth towards lower wealth agents. The strength of the first two “scaling” forces depends upon the health and organization of the financial sector.

To decompose these drivers of inequality, we first consider an economy where the funds are excluded from participating in capital markets and so capital prices are highly exposed to the business cycle. In this case, when the banks take losses during a recession, they become less willing to hold capital so the excess return on capital increases and wealthy households participating in the capital market grow their wealth more quickly than poorer agents. This generates the feature in 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 is much less able to act as a “backstop” in the capital market during crises and so the scaling force in the model is very strong. If 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 the unregulated economy where funds are allowed to participate in all markets, the capital risk premium increases less in recessions because the banker can sell capital to the fund sector. This means the change in household inequality is less driven by the capital risk premium and more driven by the insurance premium in the economy. Poorer households can only access the capital markets through the pension/insurance products issued by the fund. The spread between the capital return and the pension return is governed by the inelasticity of household demand for pension products and the financial health of the fund. This means that the it now volatility spikes in the economy that open increase insurance risk premium and lead to large increases in equality.

In Section 5, we calibrate our model to match the average capital risk premium,

average term premium, and portfolio distribution. To gauge the non-targeted performance of the calibration, 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.

In our the counterfactual exercises discussed so far, we have focused on models with an ergodic wealth distribution. However, it is not clear from the data that wealth inequality is a stationary process, particularly over the post-war period. An advantage of our global solution compared to perturbation is that we can study non-stationary economies, which we do in Subsection 5.4. We take out the tax and transfer system that ensures there is a stationary wealth distribution and instead work with a model where the consumption-wealth ratio is the only potential stabilizer. Our calibrated, non-stationary version of the 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.

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 collection of high, but finite dimensional PDEs 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 that can satisfactorily find a global solution to continuous-time models with non-trivial optimization, distribution evolution, and equilibrium blocks, without having to resort to low-dimensional approximations of the wealth distribution. (In discrete time, [Azinovic, Cole and Kubler \(2023\)](#) take on a similar class of models using a homotopy approach). 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 \(2017\)](#), [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 et al. (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 Baseline Economic Model

In this section, we outline the baseline economic model used throughout the paper. We study a continuous time, real business cycle macroeconomy with heterogeneous households who face retirement shocks and asset market participation constraints. The households are serviced by two types of financial intermediaries: a banker that accepts deposits and makes loans, and a fund manager that offers pensions to insure the retirement shocks.

2.1 Environment

Setting: The model is in continuous time with infinite horizon. There is a perishable consumption good and two types of durable capital stock. The economy is populated by a unit continuum of price taking households (h), indexed by $i \in [0, 1]$, a unit continuum of price taking bankers (b), indexed by $j \in [0, 1]$, and a unit continuum of price taking funds (f), indexed by $l \in [0, 1]$. Ultimately there will be a representative agent for the banking sector and the fund sector but not for the household sector. The economy has the following assets: short-term bank deposits, pension fund shares, bank loans, capital stock, and government bonds.

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:

$$\begin{aligned} dz_t &= \beta_z(\bar{z} - z_t)dt + \sqrt{\zeta_t}dW_{z,t}, \\ d\zeta_t &= \beta_\zeta(\bar{\zeta} - \zeta_t)dt + \sigma_\zeta\sqrt{\zeta_t}dW_{\zeta,t} \end{aligned}$$

where $W_{z,t}$ and $W_{\zeta,t}$ denotes an aggregate Brownian motion process. 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 $\Phi(\iota)k := (\iota - \phi(\iota))k$ represents the resources used from investment rate ι_t and δ is a depreciation rate.

Household: There is a unit measure of households. Each household has discount rate ρ_h and gets flow utility $u(c_{h,t}) = \beta c_{h,t}^{1-\gamma} / (1-\gamma)$ from consumption $c_{h,t}$. Households receive idiosyncratic death shocks at rate λ_h . When agents die, we assume they get a positive measure of utility $\mathcal{U}(\mathcal{C}_{h,t}) = (1-\beta)\mathcal{C}_{h,t}^{1-\Gamma} / (1-\Gamma)$ from consuming a positive measure of goods $\mathcal{C}_{h,t}$, which generates a “preferred-habitat” style preference for pension/insurance products with duration λ_h . After an agent dies, they are immediately replaced by a new agent who receives initial wealth $\underline{a}_{h,t} = \phi_h A_t$. When a household dies, it gets a payoff from its pension holdings and can consume a fraction of capital stock holdings (after paying estate taxes). For numerical convenience, we model this using a CES aggregator:

$$\mathcal{C}_{h,T} = \mathcal{C}(k_{h,T}, n_{h,T}) = \begin{cases} (\alpha(q_T k_{h,t}(1-\tau))^\nu + (1-\alpha)(n_{h,T})^\nu)^{1/\nu}, & \text{if } \nu \neq 0 \\ (q_T k_{h,t}(1-\tau))^\alpha (n_{h,T})^{1-\alpha}, & \text{if } \nu = 0 \end{cases}$$

For sufficient complementarity (e.g. Cobb-Douglas), this implies that $k_{h,t}, n_{h,t} > 0$ and eliminates the need for a short selling constraint. We will often write this using the notation:

$$\mathcal{C}_{h,T} = a\mathcal{W}(\theta_{h,T}^k, \theta_{h,T}^n), \quad \text{where:}$$

$$\mathcal{W}(\theta_{h,T}^k, \theta_{h,T}^n) = \begin{cases} (\alpha(\theta_{h,T}^k(1-\tau))^\nu + (1-\alpha)(\theta_{h,T}^n/q_T^n)^\nu)^{1/\nu}, & \text{if } \nu \neq 0 \\ (\theta_{h,T}^k(1-\tau))^\alpha (\theta_{h,T}^n/q_T^n)^{1-\alpha}, & \text{if } \nu = 0 \end{cases}$$

The households face two financial constraints:

$$\Psi_{h,k}(k_{h,t}, a_{h,t}) = \psi_{h,k,t} \bar{\Xi}_{h,t} a_{h,t}, \quad \text{where } \psi_{h,k,t} = \psi_h(\theta_{h,t}^k, \eta_{h,t}) = \frac{\bar{\psi}_k}{2\eta_{h,t}} (\theta_{h,t}^k)^2$$

$$\Psi_{h,n}(n_{h,t}, a_{h,t}) = \psi_{h,n,t} \bar{\Xi}_{h,t} a_{h,t}, \quad \text{where } \psi_{h,n,t} = \psi_h(\theta_{h,t}^n) = \frac{\bar{\psi}_n}{2} (\theta_{h,t}^n - \chi)^2$$

where $\bar{\psi}_k$ and $\bar{\psi}_n$ are the severity of the constraints, $a_{h,t}$ is the household’s wealth, $k_{h,t}$ is the household’s capital holdings, $n_{h,t}$ is the household’s pension shares, $\theta_{h,t}^k := q_t^k k_{h,t} / a_{h,t}$ is the household’s wealth share in capital, $\theta_{h,t}^n := q_t^n n_{h,t} / a_{h,t}$ is the household’s wealth share in pension shares, and $\eta_{h,t}$ is the household’s share of wealth in the economy.

Financial intermediaries: There are two types of financial intermediaries: bankers (b) and fund managers (f). Each type of intermediary $j \in (b, f)$ has discount rate ρ_j and gets log utility $u_j(c_{j,t}) = \log(c_{j,t})$ from consuming $c_{j,t}$ flow goods. The bankers issue risk-free short-term deposits that pay deposit rate r_t^d and invest in risky capital and government bonds. The fund managers sell shares to households and invest in risky capital and government bonds. Each fund share pays one good to the holder of the share when they die so it can be interpreted as a combination of a life-insurance and a pension product. Financial intermediaries of type $j \in (b, f)$ die at rate λ_j and are replaced by new financial intermediaries with initial wealth $\underline{a}_{f,t} = \phi_f A_t$. Ultimately, both banker and fund manager policies will be independent of wealth so we can replace the continuum of bankers and funds by a representative banker and fund.

Government: We treat government fiscal policy as exogenous. The government issues zero coupon bonds that mature at rate λ_m and pay 1 good at maturity. We start with the assumption that bonds are in fixed supply: M . We let q_t^m denote the price of a government bond at time t . Thus, the flow rate of bond maturity is $\lambda_m M$ and the proceeds from the re-issuance of the bonds is $q_t^m \lambda_m M$. The government raises taxes to redistribute wealth to the new entrants in the economy, which could equivalently be decentralized as an inheritance system. The government also raises two types of taxes: a flow wealth tax, τ_j , on agent of type $j \in \{h, b, f\}$ and an inheritance tax on dying households, τ_d . This implies that the government budget constraint is given by:

$$\sum_j \tau_j A_j + \lambda_h (\tau_d K_h + D_h) + \lambda_b A_b + \lambda_f A_f + q_t^m \lambda_m M = \lambda_h \underline{a}_h + \lambda_b \underline{a}_b + \lambda_f \underline{a}_f$$

where A_j is the wealth of agents in sector $j \in \{h, b, f\}$, K_h is capital owned by the household, and D_h is deposits owned by the household. If all agents have the same flow wealth tax, $\tau_a = \tau_j$ for all $j \in \{h, b, f\}$, then:

$$\tau_a = \frac{\sum_j \lambda_j \phi_j A - \lambda_h (\tau_d K_h + D_h) - \lambda_b A_b - \lambda_f A_f + (1 - q_t^m) \lambda_m M}{A}$$

Assets, markets, and financial frictions: Each period, there are competitive markets for goods and capital trading. We use goods as the numeraire. Let r_t^d denote the interest rate on deposits. We let $q_t := (q_t^k, q_t^n, q_t^m)$ denote a vector with the price of

capital, pension shares, and government bonds respectively. We guess and verify that the long-term price processes satisfy:

$$\frac{dq_t^j}{q_t^j} = \mu_{q^j,t} dt + \boldsymbol{\sigma}_{q^j,t}^T d\mathbf{W}_t, \quad j \in \{k, n, m\}$$

where $\mu_{q^j,t}$ and $\boldsymbol{\sigma}_{q^j,t} := [\sigma_{q^j,z,t}, \sigma_{q^j,\zeta,t}]^T$ are the drift and volatility for asset $j \in \{k, n, m\}$ and T indicates the transpose. We also define the return processes by:

$$\begin{aligned} dR_t^k &= r_t^k dt + \boldsymbol{\sigma}_{q^k,t}^T d\mathbf{W}_t, & r_t^k &:= \mu_{q^k,t} + \Phi(\iota) - \delta + \frac{e^{z_t} - \iota}{q_t^k}, \\ dR_{h,t}^n &= r_{h,t}^n dt + \boldsymbol{\sigma}_{q^n,t}^T d\mathbf{W}_t + \frac{1}{q_t^m} dN_{h,t}, & r_{h,t}^n &:= \mu_{q^n,t} \\ dR_{f,t}^n &= r_{f,t}^n dt + \boldsymbol{\sigma}_{q^n,t}^T d\mathbf{W}_t, & r_{f,t}^n &:= \mu_{q^n,t} + \left(\frac{1}{q_t^n} - 1 \right) \lambda \\ dR_t^m &= r_t^m dt + \boldsymbol{\sigma}_{q^m,t}^T d\mathbf{W}_t & r_t^m &:= \mu_{q^m,t} + \left(\frac{1}{q_t^m} - 1 \right) \lambda_m \end{aligned}$$

where fund shares have different flow returns for the household, $R_{h,t}^n$, and fund, $R_{f,t}^n$, because the fund aggregates across a continuum of households.

2.1.1 Discussion of Key Environmental Features

This environment is set up to nest a collection of models and forces commonly used in the literature. We discuss these connections below:

- (i) *Preferred habitat literature* (e.g. [Vayanos and Vila \(2021\)](#)): In our model, the household need for consumption at death generates demand for fund liabilities with an average maturity of $1/\lambda$. If we choose $\mathcal{U}(\cdot)$ to be either the Type I or Type II agents from the Appendix in [Vayanos and Vila \(2021\)](#) and introduce variance in λ across the household population, then we nest their preferred habitat model of the yield curve. However, our model has an important extension compared to the preferred habitat literature—we integrate the preferred habitat demand into a standard portfolio choice problem so that overall household demand is a combination of the “preferred-habitat” component and a standard portfolio choice problem that balances risk and return. This allows us to understand how risk and inelastic demand interact in a general equilibrium model.

- (ii) *Perpetual youth literature (Blanchard (1985))*: Our model nests the asset demand from both the perpetual youth and preferred habitat literatures, which can be viewed as opposite sides of the same death shock. If we set $\beta = 1$, then the households only care about consumption while alive and we recover the perpetual youth model, in which households demand annuities that pay until they die (and which could be synthetically created by shorting the life insurance products offered by the funds and purchasing bonds). If we set $\beta = 0$, then the households only care about consumption at death and we recover the pure preferred habitat model of Vayanos and Vila (2021), in which household only take a long position in fund life insurance products. Our model can be viewed as an intermediate model that nests these two forces. Throughout the paper we focus on parametrizations where the households take long positions in the fund shares. However, the model could just as easily be solved for the case where the CES is linear and some households end up shorting the pension products.
- (iii) *Participation constraint models (e.g. Basak and Cuoco (1998))*: We have set up the household participation penalty so that households increase their fraction of wealth in capital as get older. In this sense, as household wealth goes to zero, the model becomes the Basak and Cuoco (1998) environment in which households cannot participate in the capital market. However, as household wealth becomes large, the agents become unconstrained like in the Brunnermeier and Sannikov (2014) environment where households can freely participate in the capital market. At either extreme, household portfolio choices become homogeneous and the household sector aggregates. The technical difficulties come from having the household move between the extremes as they accumulate wealth.
- (iv) *Different type models (e.g. Chan and Kogan (2002), Gomez (2017))*: There is a collection models in which households have different types ex-ante (e.g. because they have heterogeneous risk aversion) but all agents within a particular type make the same portfolio decisions. These models can generate heterogeneous portfolio choices across the population and so can generate the aggregate asset portfolio for the household sector. However, they cannot match any of the portfolio data at the micro level which shows that portfolio decisions vary with household wealth.

2.2 Equilibrium

We now setup the agent problems and the equilibrium. We use the notation $\mathbf{x} = (x_t)_{t \geq 0}$ to denote the stochastic process for variable x_t . We let $a_{h,t} := q_t^k k_{h,t} + q_t^n n_{h,t} + d_{h,t}$ denote the household wealth, $a_{b,t} := q_t^k k_{b,t} + q_t^m m_{b,t} + d_{h,t}$ denote the banker wealth, and $a_{f,t} := q_t^k k_{f,t} + q_t^m m_{f,t} + q_t^n n_{f,t}$ denote the fund wealth. We let $\mu_{a_j,t}$ and $\sigma_{a_j,t}$ denote the geometric drift and volatility for the wealth of type j with wealth $a_{j,t}$. We let $\theta_{j,t}^l = q_t^l l_{j,t}/a_{j,t}$ denote the share of wealth that an agent of type j with wealth $a_{j,t}$ has in asset l . We let $\boldsymbol{\theta}_{j,t}$ denote the vector of wealth shares chosen by an agent of type j with wealth $a_{j,t}$ at time t .

Household problem: Given their belief about price processes, (\tilde{r}, \tilde{q}) , and initial wealth, $a_{h,0}$, a household chooses processes $(\mathbf{c}_h, \boldsymbol{\theta}_h, \iota_h)$ to solve the Problem (2.1) below:

$$\begin{aligned} & \max_{\mathbf{c}_h, \boldsymbol{\theta}_h} \mathbb{E} \left[\int_0^T e^{-\rho_h t} \left(u(c_{h,t}) + \psi_h(\boldsymbol{\theta}_{h,t}^k) \Xi_{h,t} a_{h,t} \right) dt + e^{-\rho T} \mathcal{U}(\mathcal{C}_{h,t}) \right] \quad s.t. \\ & \frac{da_{h,t}}{a_{h,t}} = \theta_{h,t}^n d\tilde{R}_{h,t}^n + \theta_{h,t}^k d\tilde{R}_{h,t}^k + \left((1 - \theta_{h,t}^k - \theta_{h,t}^n) \tilde{r}_{d,t} - c_{h,t}/a_{h,t} - \tau_{h,t} \right) dt \quad (2.1) \\ & \mathcal{C}_{h,t} \leq a_i \mathcal{W}(\boldsymbol{\theta}_{h,t}^k, \boldsymbol{\theta}_{h,t}^n) \\ & \theta_{h,t}^k, \theta_{h,t}^n, \theta_{h,t}^d \geq 0 \end{aligned}$$

where $\tau_{h,t}$ is the net tax or transfer (per unit of wealth) while agents are alive.

Banker problem: Given their belief about price processes, (\tilde{r}, \tilde{q}) , and initial wealth, $a_{b,0}$, a banker chooses processes $(\mathbf{c}_b, \boldsymbol{\theta}_b, \iota_b)$ to solve the Problem (2.2) below:

$$\begin{aligned} & \max_{\mathbf{c}_b, \boldsymbol{\theta}_b, \iota_b} \left\{ \int_0^\infty e^{-\rho_b t} u(c_{b,t}) dt \right\} \quad s.t. \\ & \frac{da_{b,t}}{a_{b,t}} = \theta_{b,t}^k d\tilde{R}_t^k + \theta_{b,t}^m d\tilde{R}_t^m + \left((1 - \theta_{b,t}^k - \theta_{b,t}^m) \tilde{r}_t^d - c_{b,t}/a_{b,t} - \tau_{b,t} \right) dt \quad (2.2) \end{aligned}$$

Fund problem: Given their belief about price processes, (\tilde{r}, \tilde{q}) , and initial wealth, $a_{f,0}$,

a fund manager chooses processes $(\mathbf{c}_f, \boldsymbol{\theta}_f, \iota_f)$ to solve the Problem (2.3) below:

$$\begin{aligned} & \max_{\mathbf{c}_f, \boldsymbol{\theta}_f, \iota_f} \left\{ \int_0^\infty e^{-\rho_f t} u(c_{f,t}) dt \right\} \quad s.t. \\ & \frac{da_{f,t}}{a_{f,t}} = \theta_{f,t}^k d\tilde{R}_t^k + \theta_{f,t}^m d\tilde{R}_t^m + (1 - \theta_{f,t}^k - \theta_{f,t}^m) d\tilde{R}_t^n + (-c_{f,t}/a_{f,t} - \tau_{f,t}) dt \end{aligned} \quad (2.3)$$

Distribution: Throughout this paper, we work with distribution of wealth shares rather than wealth levels. The bank and fund sectors aggregate so we will only need to track of the aggregate states for each sector. We let $\eta_{b,t} := a_{b,t}/A_t$ and $\eta_{f,t} := a_{f,t}/A_t$ denote the share of aggregate wealth held by the banking and fund sectors. The uninsurable idiosyncratic shocks and wealth dependent differences in household portfolio constraints generate a non-degenerate cross-section distribution of household wealth across the economy. We let $g_{h,t} = \{\eta_{i,t} = a_{i,t}/A_t : i \in \mathcal{I}\}$ denote measure of household wealth shares across the economy at time t for a given filtration \mathcal{F}_t , where \mathcal{F}_t is generated the by aggregate shock processes $\{\mathbf{W}_t\}_{t \geq 0}$. With some abuse of notation, we let $G = (\eta_{b,t}, \eta_{f,t}, g_{h,t})$ denote the collection of distribution states in the economy.

Definition 1 (Equilibrium). Let $\mathbf{r} = (r_t^d)_{t \geq 0}$ and $\mathbf{q} = (q_t^k, q_t^n, q_t^m)_{t \geq 0}$ denote the stochastic processes for interest rates and long term asset prices. For a given set of government taxation policies, an equilibrium is a collection of \mathcal{F}_t -adapted processes $(\mathbf{K}, \mathbf{r}, \mathbf{q}, G)$ and household decision processes $(\mathbf{c}_i, \boldsymbol{\nu}_i, \boldsymbol{\theta}_i)$ for $i \in I$ and financial intermediary decision process $(\mathbf{c}_j, \boldsymbol{\nu}_j, \boldsymbol{\theta}_j)$ for $j \in \{b, f\}$ such that:

1. Given beliefs $(\tilde{\mathbf{r}}, \tilde{\mathbf{q}})$, households ($i \in I$) solve Problem 2.1, bankers solve Problem 2.2, and fund managers solve Problem 2.3.
2. The price processes (\mathbf{r}, \mathbf{q}) satisfies market clearing conditions at each t (where the capital letter $Y_{j,t}$ refers to the aggregate volume of variable Y held by an agent of type j):
 - (a) Goods market clears: $C_{h,t} + C_{b,t} + C_{f,t} + \lambda C_{h,t} = e^{z_t} K_t - \iota_t K_t$ where K_t is the aggregate capital stock.
 - (b) Capital market clears: $\sum_{j \in \{b,n\}} K_{j,t} = K_t$
 - (c) Annuity market clears: $N_{h,t} + N_{f,t} = 0$

- (d) Deposit market clears: $D_{h,t} + D_{b,t} = 0$
 - (e) Bond market clears: $M_{b,t} + M_{f,t} = M$
3. Agent beliefs are consistent with equilibrium $(\tilde{\mathbf{r}}, \tilde{\mathbf{q}}) = (\mathbf{r}, \mathbf{q})$.

2.3 Recursive Characterization of Equilibrium

In this section, we characterize the equilibrium recursively. The finite dimensional aggregate state vector is: $\mathbf{s} := (z, \zeta, K, \eta_b, \eta_f)$, where η_j is the fraction of wealth held by sector $j \in \{h, b, f\}$. The other aggregate state is the density of household wealth shares, $g_h(\eta)$. Let $\mathbf{S} := (\mathbf{s}, g_h)$ denote the full aggregate state space. Under the recursive formulation of the problem, agent beliefs about the price process become agent beliefs about the evolution of sector wealth, $(\tilde{\mu}_{\eta,j}, \tilde{\sigma}_{\eta,j})_{j \in \{h, f\}}$, and the evolution of the density of household wealth shares:

$$dg_{h,t}(a) = \tilde{\mu}_g(a, \mathbf{S})dt + \tilde{\sigma}_g(a, \mathbf{S})^T d\mathbf{W}_t$$

For convenience, we also define the state spaces for an individual agent by $\mathbf{x} := (a, z, \zeta, K, \eta_b, \eta_f)$ and $\mathbf{X} := (\mathbf{x}, g_h)$.

Let $V_j(\mathbf{X})$ denote the value function for an agent of type $j \in \{h, b, f\}$ with state variable \mathbf{X} . Let $\xi_j(\mathbf{X}) := \partial_a V_j(\mathbf{X})$ denote the partial derivative of the value function for an agent of type $j \in \{h, b, f\}$. We let $\mu_{\xi_h}(a, \mathbf{s})$ and $\sigma_{\xi_h}(a, \mathbf{s})$ denote drift and volatility of the process for ξ_j . Theorem 1 below summarizes the recursive characterization of equilibrium. For convenience we group the characterization into three blocks: (i) the optimization problems of the agents, (ii) the evolution of the distribution, and (iii) market clearing.

Theorem 1. *Block 1: Agent Optimization: Given prices and price processes $(r^d, (q^l, r^l, \sigma_{q^l})_{l \in \{k, n, m\}})$, for all agents the investment rate satisfies $\Phi'(\iota) = (q^k)^{-1}$ and:*

- Household optimization implies that $(\xi_h, c_h, \boldsymbol{\theta}_h)$ satisfy the Euler and FOCs:

$$\begin{aligned}
0 &= -(\rho + \lambda) + \mu_{\xi_h} + r^d - \tau_h + \psi_{h,k}(\theta_h^k) - \partial_{\theta_h^k} \psi_{h,k}(\theta_h^k) \theta_h^k \\
&\quad + \psi_{h,n}(\theta_h^n) - \partial_{\theta_h^n} \psi_{h,n}(\theta_h^n) \theta_h^n \\
0 &= u'(c_h) - \xi_h \\
0 &= r^k - r^d + \lambda \partial_{\theta_h^k} \mathcal{W}(\theta_h^k, \theta_h^n) \frac{\mathcal{U}'(\mathcal{C})}{a \xi_h} + \partial_{\theta_h^k} \psi_{h,k} + \boldsymbol{\sigma}_{\xi_h}^T \boldsymbol{\sigma}_{q^k} \\
0 &= r^n - r^d + \lambda \partial_{\theta_h^n} \mathcal{W}(\theta_h^k, \theta_h^n) \frac{\mathcal{U}'(\mathcal{C})}{a \xi_h} + \partial_{\theta_h^n} \psi_{h,n} + \boldsymbol{\sigma}_{\xi_h}^T \boldsymbol{\sigma}_{q^n}
\end{aligned}$$

- Bank optimization implies that $(\xi_b, c_b, \boldsymbol{\theta}_b)$ satisfy the Euler and FOCs:

$$\xi_b = \frac{1}{(\rho_b + \lambda_b) a_b}, \quad c_b = (\rho_b + \lambda_b) a_b, \quad 0 = r_t^k - r^d + \boldsymbol{\sigma}_{\xi_b}^T \boldsymbol{\sigma}_{q^k}$$

- Fund optimization implies that $(\xi_f, c_f, \boldsymbol{\theta}_f)$ satisfy the Euler and FOCs:

$$\xi_f = \frac{1}{(\rho_f + \lambda_f) a_f}, \quad c_f = (\rho_f + \lambda_f) a_f, \quad 0 = r^k - r^d + \boldsymbol{\sigma}_{\xi_f}^T \boldsymbol{\sigma}_{q^k}$$

where in all cases, the drift and volatility of the ξ_j for $j \in (h, b, f)$ are given by ITO's lemma:

$$\begin{aligned}
\mu_{\xi_j} \xi_j(\mathbf{X}) &= (D_x \xi_j(\mathbf{X}))^T \boldsymbol{\mu}_x + \frac{1}{2} \text{tr} \left\{ (\boldsymbol{\sigma}_x(\mathbf{X}, \boldsymbol{\theta}_j) \odot \mathbf{X})^T (\boldsymbol{\sigma}_x(\mathbf{X}, \boldsymbol{\theta}_j) \odot \mathbf{x}) D_x^2 \xi_j(\mathbf{X}) \right\} \\
&\quad + \mathcal{L}_g \xi_j(\mathbf{X}) \\
\boldsymbol{\sigma}_{\xi_j} \xi_j &= (\boldsymbol{\sigma}_x \odot \mathbf{x})^T (D_x \xi_j)
\end{aligned}$$

and $\mathcal{L}_g \xi_j(\mathbf{X})$ denotes the collection of terms with Frechet derivatives of ξ_j with respect to ξ_j .

Block 2: Distribution evolution. Given prices and price processes $(r^d, (q^l, r^l, \boldsymbol{\sigma}_{q^l})_{l \in \{k, n, m\}})$ and agent decisions $(\xi_j, c_j, \boldsymbol{\theta}_j, \iota)_{j \in \{h, b, f\}}$, we can track the distribution evolution. At

the sector level, the bank and fund wealth shares evolve according to:

$$\begin{aligned}\frac{d\eta_{b,t}}{\eta_{b,t}} &= \left(\mu_{A_b,t} - \mu_{A,t} + (\boldsymbol{\sigma}_{A,t} - \boldsymbol{\sigma}_{A_b,t})^T \boldsymbol{\sigma}_{A,t} \right) dt + (\boldsymbol{\sigma}_{A_b,t} - \boldsymbol{\sigma}_{A,t})^T d\mathbf{W}_t \\ \frac{d\eta_{f,t}}{\eta_{f,t}} &= \left(\mu_{A_f,t} - \mu_{A,t} + (\boldsymbol{\sigma}_{A,t} - \boldsymbol{\sigma}_{A_f,t})^T \boldsymbol{\sigma}_{A,t} \right) dt + (\boldsymbol{\sigma}_{A_f,t} - \boldsymbol{\sigma}_{A,t})^T d\mathbf{W}_t\end{aligned}$$

where:

$$\begin{aligned}\mu_{A_b,t} &= r^d + \theta_h^k (r^k - r^d) - (\rho_b + \lambda_b) - \tau_b + \lambda_b (\phi_b \eta_{b,t}^{-1} - 1) \\ \mu_{A_f,t} &= r^b + \theta_f^k (r^k - r_f^b) + \theta_f^m (r^m - r_f^m) - (\rho_f + \lambda_f) - \tau_f + \lambda_f (\phi_f \eta_{f,t}^{-1} - 1) \\ \mu_{A,t} &= \vartheta_t (\mu_{q^k} + \Phi(\iota) - \delta) + (1 - \vartheta_t) \mu_{q^m}\end{aligned}$$

and where aggregate wealth is given by $A_t = q_t^k K_t + q_t^m M$ and $\vartheta_t = q_t^k K_t / (q_t^k K_t + q_t^m M)$. Within the household sector, density of households wealth shares evolves according to:

$$\begin{aligned}dg_{h,t}(\eta) &= + \lambda_h \phi(\eta) - \lambda_h g_{h,t}(\eta) - \partial_\eta [\mu_\eta(\eta, \mathbf{s}_t, g_{h,t}) g_{h,t}(\eta)] \\ &+ \frac{1}{2} \partial_\eta \left[(\sigma_{\eta,z}^2(\eta, \mathbf{s}_t, g_{h,t}) + \sigma_{\eta,\zeta}^2(\eta, \mathbf{s}_t, g_{h,t})) g_{h,t}(\eta) \right] dt \\ &- \partial_\eta [\sigma_{\eta,z}(\eta, \mathbf{s}_t, g_{h,t}) g_{h,t}(\eta)] dW_{z,t} - \partial_\eta [\sigma_{\eta,\zeta}(\eta, \mathbf{s}_t, g_{h,t}) g_{h,t}(\eta)] dW_{\zeta,t}\end{aligned}$$

where:

$$\begin{aligned}\mu_{\eta_i,t} &= \mu_{a_i,t} - \mu_{A,t} + (\boldsymbol{\sigma}_{A,t} - \boldsymbol{\sigma}_{a_i,t})^T \boldsymbol{\sigma}_{A,t} \\ \boldsymbol{\sigma}_{\eta_i,t} &= \boldsymbol{\sigma}_{a_i,t} - \boldsymbol{\sigma}_{A,t}\end{aligned}$$

Block 3: Market clearing and belief consistency: The equilibrium prices satisfy:

$$\begin{aligned}\int c_h g_h(\eta) d\eta + \frac{C_b}{A_b} \eta_b + \frac{C_f}{A_f} \eta_f + \lambda \int C_h g_h(\eta) d\eta &= \frac{(e^z - \iota) K}{q^k K + q^m M} \\ \int \theta_h g_h(\eta) d\eta + \theta_f \eta_f + \theta_b \eta_b &= \vartheta \\ \theta_f^n \eta_f + \int \theta_h^n g_h(\eta) d\eta &= 0 \\ \int \theta_h^d g_h(\eta) d\eta + \theta_b^d a_b &= 0 \\ \theta_f^m \eta_f &= 1 - \vartheta\end{aligned}$$

where $\vartheta = q^k K / (q^k K + q^m M)$. The long term assets prices are only implicitly defined by the asset pricing equations and so must satisfy consistency with Itô's Lemma for $l \in (k, n, m)$:

$$\begin{aligned}\mu_{q^l} q^l(\mathbf{X}) &= (D_x q^l(\mathbf{X}))^T \boldsymbol{\mu}_x + \frac{1}{2} \text{tr} \left\{ (\boldsymbol{\sigma}_x(\mathbf{X}, \boldsymbol{\theta}_h) \odot \mathbf{X})^T (\boldsymbol{\sigma}_x(\mathbf{X}, \boldsymbol{\theta}_h) \odot \mathbf{x}) D_x^2 q^l(\mathbf{X}) \right\} \\ &\quad + \mathcal{L}_g \xi_h(\mathbf{X}) \\ \boldsymbol{\sigma}_{q^l} q^l &= (\boldsymbol{\sigma}_x \odot \mathbf{x})^T (D_x q^l)\end{aligned}$$

Proof. See Appendix A. □

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 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. Conceptually, we can view our approach as a type of “projection” onto a neural network. At a high level, this involves:

- (a) Replacing the agent continuum by a finite dimensional distribution approximation,
- (b) Representing the equilibrium functions by neural networks with the states as inputs, and
- (c) Training the neural network parameters to minimize the loss in the equilibrium conditions on randomly sampled points from the state space.

Although this approach is straightforward to describe at a high level, implementing it successfully involves many non-trivial decisions that we discuss in this section.

3.1 Finite dimensional distribution approximation

In [Gu et al. \(2023\)](#), we compare the three main finite dimensional distribution approximation approaches: working with a finite number of agents, discretizing the wealth space onto a grid, and projecting the distribution onto a collection of basis functions. In this paper, we focus on the finite agent approximation approach because we have found it is convenient for handling competitive markets. Conceptually, this approach replaces the household distribution state by a finite collection of agents while imposing analytically that the idiosyncratic agent noise averages out.

Formally, we impose that there are $I < \infty$ household “dynasties” in the economy. When a household in a dynasty dies, they are replaced by a new household with wealth $a_i = \phi_h A_t$ as before. Let $\eta_i := a_i/A$ denote the wealth share for household dynasty i . We solve for an equilibrium in which agents behave as price takers and forecast prices under the assumption that the idiosyncratic death shocks have averaged out and so perceive the law of motion:

$$\begin{aligned} \frac{d\eta_{i,t}}{\eta_{i,t}} &= \lambda_h \left(\frac{\phi_h}{\eta_i} - 1 \right) + \left(\mu_{a_i,t} - \mu_{A,t} + (\boldsymbol{\sigma}_{A,t} - \boldsymbol{\sigma}_{a_i,t})^T \boldsymbol{\sigma}_{A,t} \right) dt \\ &\quad + (\boldsymbol{\sigma}_{a_i,t} - \boldsymbol{\sigma}_{A,t})^T d\mathbf{W}_t \end{aligned}$$

That is, we solve the model with agents behaving as if they are in a model with continuum of agents.

A natural and frequently raised concern with this approach (and finite agent approaches more generally) is that the simulation of the I agent economy would contain idiosyncratic noise and so be inconsistent with perceived law of motion. However, there are number of reasons why this is not a problem. First, we do not use simulation to solve the model because we are working with a continuous time analytical formulation of the problem. So, concerns about simulation are unrelated to the accuracy of neural network solution. Second, when we need to simulate the solved model (e.g. to generate time paths or impulse responses), we use our neural network solution to approximate a finite difference approximation to the KFE and so are able to simulate the limiting economy with a continuum of agents. We do this following the approach we developed in [Gu et al. \(2023\)](#) and summarize in Appendix B to this paper. In this sense, we exploit the continuous time analytical formulation to be able to maintain the convenience of a finite agent economy without have to deal with the finite sample noise problem that appears in discrete time simulation based training methods.

3.2 Neural network representation and loss function

A “direct” implementation of our deep learning approach would be to parameterize the equilibrium objects and then train the neural networks to minimize a large loss function that combines all the general equilibrium equations described in Theorem 1. Although this should work in principle, many researchers have found it very difficult to implement in practice. Instead, we rewrite the equilibrium characterization to “help” the deep learning algorithm to train the neural networks.

First, we approximate the equilibrium functions rather than the partial equilibrium functions. Under the finite agent approximation, the aggregate state space is $\hat{\mathbf{S}} := (z, \zeta, K, \eta_1, \dots, \eta_I, \eta_b, \eta_f) \in \mathcal{S}$. In equilibrium, all the functions can be expressed directly in terms of the aggregate state $\hat{\mathbf{S}}$. For example, the function $\xi_h(a_i, \hat{\mathbf{S}})$ has an equilibrium representation given by:

$$\Xi_h(\hat{\mathbf{S}}) := \xi_h\left(\eta_i A(\hat{\mathbf{S}}), \hat{\mathbf{S}}\right)$$

where $A(\hat{\mathbf{S}}) = q^k(\hat{\mathbf{S}})K + q^m(\hat{\mathbf{S}})M$ is equilibrium aggregate wealth. We solve directly

for the equilibrium functions (i.e. Ξ_h) rather than for the partial equilibrium functions (i.e. ξ_h) in order to avoid having to nest the neural network approximation for prices (i.e. q^k, q^m) inside the neural network approximation for other variables when we impose equilibrium.

Second, we use Neural Nets to represent variables that are relatively easy to train. This leads us to parameterize the following variables by Neural Nets:

$$\begin{aligned}
\hat{\omega}_h &: \mathcal{S} \rightarrow \mathbb{R}, (\hat{\mathbf{S}}, \Theta_{\omega_h}) \mapsto \hat{\omega}_h(\hat{\mathbf{S}}; \Theta_{\omega_h}), \\
\hat{\Omega}_h &: \mathcal{S} \rightarrow \mathbb{R}, (\hat{\mathbf{S}}, \Theta_{\Omega_h}) \mapsto \hat{\Omega}_h(\hat{\mathbf{S}}; \Theta_{\Omega_h}), \\
\hat{\theta}_h^l &: \mathcal{S} \rightarrow \mathbb{R}, (\hat{\mathbf{S}}, \Theta_{\theta_j}) \mapsto \hat{\theta}_j(\hat{\mathbf{S}}; \Theta_{\theta_j}), \quad \forall l \in \{k, n\} \\
\hat{q}^l &: \mathcal{S} \rightarrow \mathbb{R}, (\hat{\mathbf{S}}, \Theta_{q^l}) \mapsto \hat{q}^l(\hat{\mathbf{S}}; \Theta_{q^l}), \quad \forall l \in \{n, m\} \\
\hat{\mu}_{q^k} &: \mathcal{S} \rightarrow \mathbb{R}, (\hat{\mathbf{S}}, \Theta_{\mu, q^k}) \mapsto \hat{\mu}_{q^k}(\hat{\mathbf{S}}; \Theta_{\mu, q^k}), \quad \forall l \in \{n, m\} \\
\hat{\sigma}_{q^l} &: \mathcal{S} \rightarrow \mathbb{R}, (\hat{\mathbf{S}}, \Theta_{\sigma, q^l}) \mapsto \hat{\sigma}_{q^l}(\hat{\mathbf{S}}; \Theta_{\sigma, q^l}), \quad l \in \{k, n, m\}
\end{aligned}$$

where $\omega := c/a$ denotes the household consumption-to-wealth ratio during their lifetime, $\Omega := C/a$ denotes the household consumption-to-wealth ratio at death, and Θ_ν denotes the parameters for the Neural Net approximation of variable ν . Why do we approximate these variables? In general, it is easier for the Neural Net to approximate well behaved, bounded functions. This guides our choices about how to parametrize the household optimization variables. It is easier to approximate $\xi_h = \partial_a V_h$ than V_h because it is easier to impose concavity. It is even easier to approximate the consumption to wealth ratio $\omega_h(\eta_i)$ and then reconstruct $\xi_h(\eta_i) = (\omega_h(\eta_i)\eta_i A)^{-\gamma}$ because then the explosive curvature is encoded analytically.

Third, we impose market clearing explicitly rather than including the market clearing conditions as part of the loss function. Given the neural network approximations $(\hat{\omega}_h, \hat{\Omega}_h, \hat{\theta}_j, \hat{q}^l, \hat{\mu}_{q^k}, \hat{\sigma}_{q^l})$, we can solve for the other equilibrium variables explicitly using linear algebra. The neural network approximations then need to satisfy the following equations (after imposing market clearing and with $\hat{\Xi} = u'(\hat{\omega}(\hat{\mathbf{S}})\eta\hat{q}^k(\hat{\mathbf{S}})K)$):

$$\begin{aligned}
\mathcal{L}_\omega(\hat{\mathbf{S}}) &= (r^d - \tau_h - \rho_h - \lambda_h)\hat{\Xi} + \mu_\Xi(\hat{\mathbf{S}}) + \psi_{h,k}(\theta_h^k(\hat{\mathbf{S}})) \\
&\quad - \partial_{\theta_h^k} \psi_{h,k}(\theta_h^k(\hat{\mathbf{S}}))\theta_h^k(\hat{\mathbf{S}}) + \psi_{h,n}(\theta_h^n(\hat{\mathbf{S}})) - \partial_{\theta_h^n} \psi_{n,k}(\theta_h^n(\hat{\mathbf{S}}))\theta_h^n(\hat{\mathbf{S}}) && \dots \text{Euler eq.} \\
\mathcal{L}_\Omega(\hat{\mathbf{S}}) &= \hat{\Omega} - \eta A(\hat{\mathbf{S}})\mathcal{W}(\theta_h^k(\hat{\mathbf{S}}), \theta_h^n(\hat{\mathbf{S}})) && \dots \text{Death cons.} \\
\mathcal{L}_{\theta_h^l}(\hat{\mathbf{S}}) &= \theta_h^{l*} - \theta_h^l(\hat{\mathbf{S}}), \quad l \in (k, n) && \dots \text{Portfolio opt.} \\
\mathcal{L}_{q^n}(\hat{\mathbf{S}}) &= \lambda - q^n(\hat{\mathbf{S}})(r_f^n - r^k + \lambda - (\mu_{q^n} - \mu_{q^k} - z/q^k)) && \dots \text{FOC diff + MC} \\
\mathcal{L}_{q^m}(\hat{\mathbf{S}}) &= r^m - f^k - (\mu_{q^m} - \mu_{q^k} - z/q^k + \lambda_m(1/q^m(\hat{\mathbf{S}}) - 1)) && \dots \text{FOC diff + MC} \\
\mathcal{L}_{\mu_{q^k}}(\hat{\mathbf{S}}) &= (D_{\hat{\mathbf{S}}}q^l)^T \boldsymbol{\mu}_{\hat{\mathbf{S}}} + \frac{1}{2} \text{tr} \left\{ (\boldsymbol{\sigma}_{\hat{\mathbf{S}}}(\hat{\mathbf{S}}, \boldsymbol{\theta}_h) \odot \hat{\mathbf{S}})^T (\boldsymbol{\sigma}_x(\hat{\mathbf{S}}, \boldsymbol{\theta}_h) \odot \hat{\mathbf{S}}) D_{\hat{\mathbf{S}}}^2 q^l \right\} && \dots \text{Consistency} \\
\mathcal{L}_\sigma(\hat{\mathbf{S}}) &= \hat{\boldsymbol{\sigma}}_q(\hat{\mathbf{S}}) - (\boldsymbol{\sigma}_{\hat{\mathbf{S}}} \odot \hat{\mathbf{S}})^T (D_{\hat{\mathbf{S}}}q^l), \quad \forall l \in (k, n, m) && \dots \text{Consistency}
\end{aligned} \tag{3.1}$$

3.3 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. We then update our guesses for the neural network approximations.

Algorithm 1: Pseudo Code

- 1: Initialize neural network objects $(\hat{\omega}_h, \hat{\Omega}_h, \hat{\boldsymbol{\theta}}_j, \hat{q}^l, \hat{\mu}_{q^k}, \boldsymbol{\sigma}_{q^l})$ with parameters $\boldsymbol{\Theta}$,
- 2: Initialize optimizer.
- 3: **while** Loss > tolerance **do**
- 4: Sample N new training points: $(\hat{\mathbf{S}}^n = (z^n, \zeta^n, K^n, (\eta_i)_{i \leq I}^n, \eta_b^n, \eta_f^n))_{n=1}^N$.
- 5: Calculate equilibrium at each training point $\hat{\mathbf{S}}^n$:
- 5: Construct loss as:

$$\hat{\mathcal{L}}(\hat{\mathbf{S}}^n) = (\mathcal{L}_\omega + \mathcal{L}_\Omega + \mathcal{L}_{\theta_h^l} + \mathcal{L}_{q^l} + \mathcal{L}_{\mu_{q^k}} + \mathcal{L}_\sigma)(\hat{\mathbf{S}}^n; \boldsymbol{\Theta})$$

where $\hat{\mathcal{L}}_v$ is defined by equation (3.1) for each variable v .

- 6: Update $\boldsymbol{\Theta}$ using stochastic gradient descent.
 - 7: **end while**
-

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 $(\{\hat{\omega}_h\}_{h=1}^I, \hat{\sigma}_q)$. 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} .²




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 the Economic Mechanisms

Our model contains a rich and novel set of economic forces. In this section, we introduce the different forces gradually in order to isolate the different mechanisms at play. Subsection 4.1 studies the model without household heterogeneity. This allows us to focus on asset pricing dynamics and the sector level implications of restricting fund involvement in the capital market. Subsection 4.2 introduces household heterogeneity into a model. This allows us to study how financial distress in the financial sector and participation constraints in the household sector impact inequality.

²Figure 1 only shows for 300 epochs since we ignore epochs whenever the loss is larger than the running minimum.

4.1 Sector Level Asset Pricing

In this section, we study a version of our model in which households have homogeneous wealth and so there are three representative agents in the economy: a household, a banker, and a fund manager. This would be nested in our general environment by a parametrization in which $\underline{a}_h = A_h$ and all households start with the same wealth. We collect our analysis into a number of lessons.

Lesson 1: Regulation determines which agent acts as the “back-stop” during bad times.

Figure 2 plots key equilibrium variables as a function of aggregate TFP (z) and stochastic volatility (ζ). (Additional plots are shown in Appendix F.) The blue lines correspond to an economy where the fund portfolio is unrestricted in its asset holdings (similar to the second half of the nineteenth century in the US) and the dashed orange line corresponds to an economy where the fund is restricted to holding long duration bonds (similar to contemporary regulation). The left column depicts portfolio choices across the different sectors. From this, we can see the bank leverage channel that is present in many macro-finance models: in good times the bank takes on additional leverage which means that in bad times they take additional losses and end up deleveraging. When the fund is able to participate in the capital market (the blue line), then they respond by purchasing capital from banks during recessions. However, if funds are restricted from capital markets (the orange line), then either the bankers have to continue to hold the capital themselves or sell the capital to the household sector. In other words, the regulatory restrictions on the fund determine whether the funds or households end up acting as the backstop to the banking sector in bad times.

What drives the trading behaviour between the bankers and the fund managers? The key force is that the market value of bank and fund liabilities respond very differently to the aggregate shocks in the economy. The bankers issue short term risk-free liabilities. By contrast, the funds issue long-term insurance products that decrease in price during recessions (when goods are scarce and the marginal value of consumption is high) and could increase or decrease in price during periods of high volatility (when goods are scarce but households also want to rebalance their portfolios into safer assets). This means that when a recession hits, the networth the bankers falls while the net-worth of the funds increases and so the funds have a

comparative advantage in hedging business cycle risk. As a result, the bankers trade their capital to the funds during recessions. However, having the funds play this role is not necessarily costless for the economy because the funds are not necessarily a good hedge against the volatility shocks.

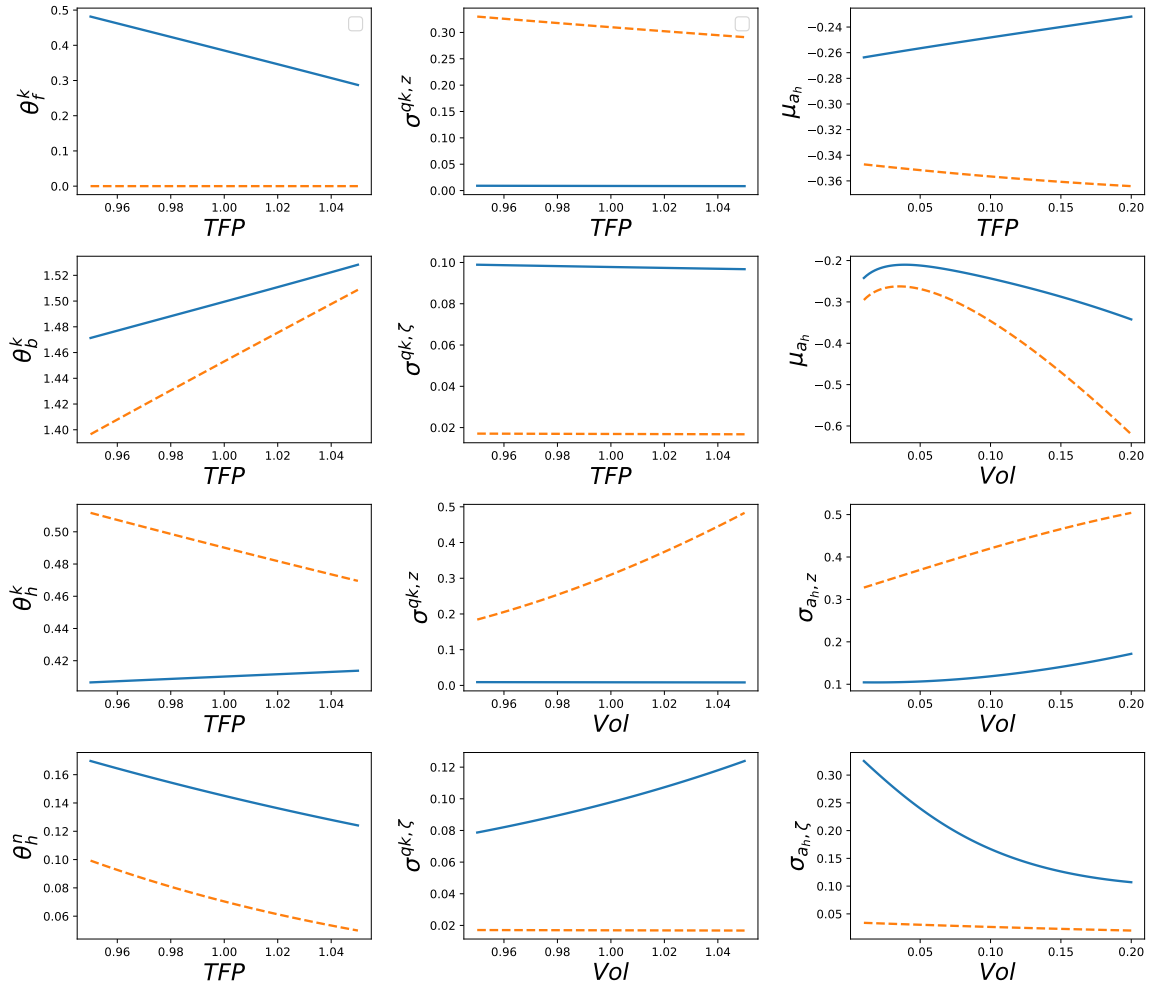


Figure 2: The figure plots equilibrium variables as a function of the wealth share of banks and funds. The solid blue line refers to the unregulated economy in which the fund can hold any assets while the dashed orange line refers to the regulated economy in which the fund can only hold long duration bonds.

Lesson 2: Regulation determines which shocks generate high endogenous volatility and which variables experience high volatility.

The second column in Figure 2 shows how the bank and fund trading behavior ends up affecting the volatility of capital prices. Evidently, the capital prices in the unregulated economy (the blue lines) are less volatile in response to business cycle shocks (i.e. lower $\sigma_{q^k,z}$) but more volatile in response to uncertainty shocks (i.e. lower $\sigma_{q^k,\zeta}$). This occurs because the funds purchase capital during recessions (and so stabilize the price of capital) but not necessarily willing to purchase capital during high volatility events.

The final column of Figure 2 shows the impact on the household wealth evolution. Again, we see the trade-off that giving the fund broad access to asset markets stabilizes household wealth in response to business cycle shocks but destabilizes household wealth in response to volatility shocks. In this sense, allowing the fund to act as the backstop during recessions does not come free: funds do not internalize how their capital purchases increase household wealth exposure to uncertainty shocks.

Lesson 3: Restricting fund participation increases capital risk premium and can also increase the convenience yield on government debt.

Finally, Figure 3 shows the portfolios and excess returns as the wealth shares of the fund sector and the banking sector change. Evidently, the impact of decreasing banker networth varies considerably when the fund is allowed to participate in capital markets. In an unregulated economy, low net-worth banks short capital while funds go long capital. In the restricted economy, the banks are forced to hold capital even as they take losses. This ends up leading to a higher risk premium in the regulated economy. The impact on government borrowing costs is ambiguous. On the one hand, forcing funds to hold government debt creates a captive debt market. When banks and funds are well capitalized, this leads to a higher convenience yield on government debt. However, forcing funds to hold government debt also increases volatility elsewhere in the economy and so changes household demand for fund products. Ultimately, this can lead to lower demand for government debt (and in-fact a negative convenience yield).

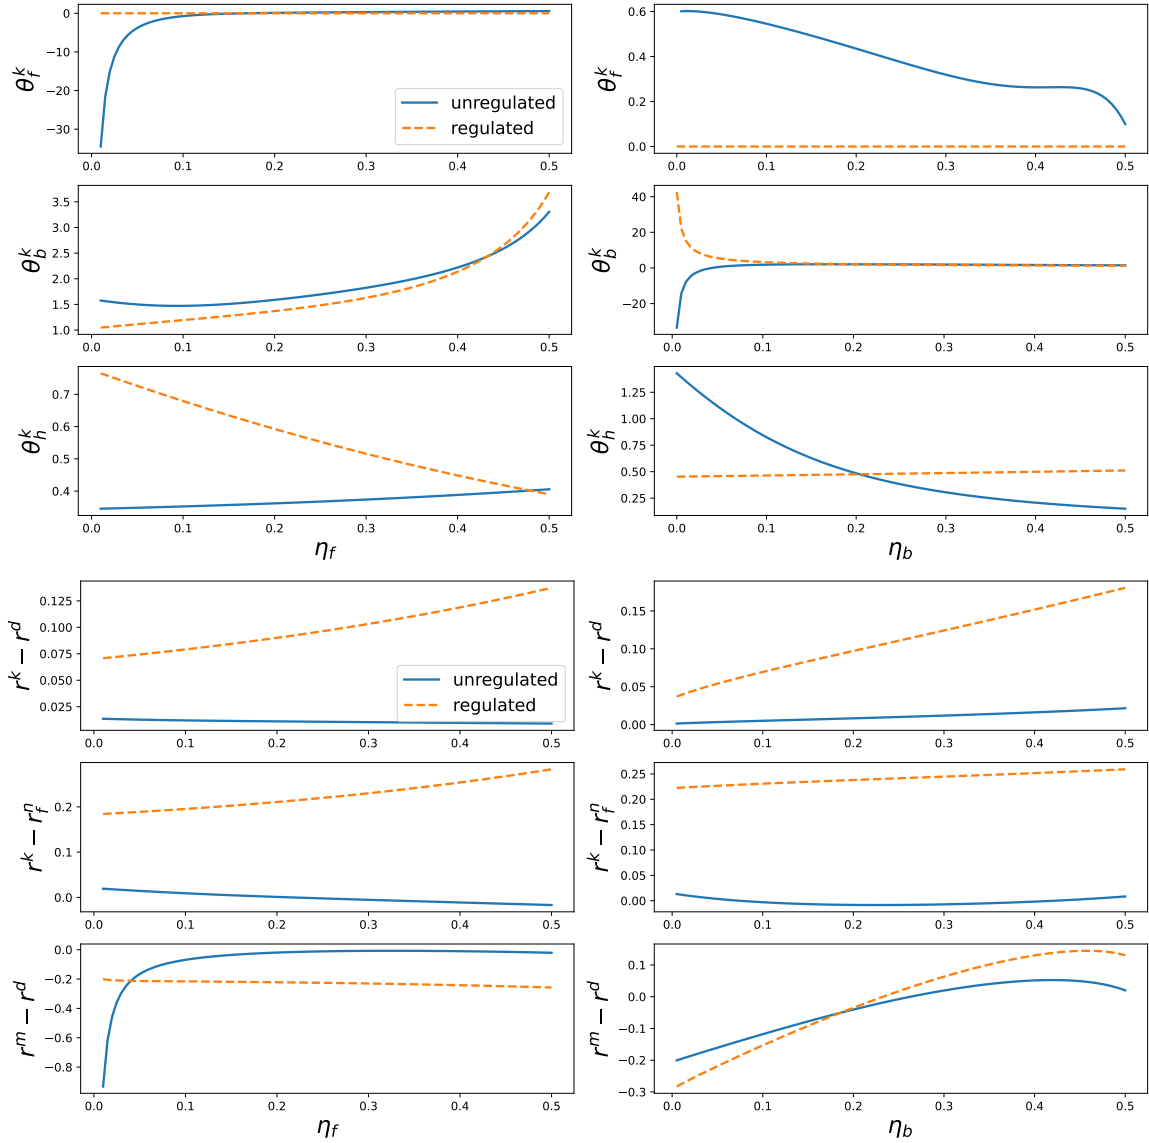


Figure 3: The figure plots equilibrium variables as a function of TFP (and volatility). The solid blue line refers to the unregulated economy in which the fund can hold any assets while the dashed orange line refers to the regulated economy in which the fund can only hold long duration bonds.

4.2 Inequality and Asset Pricing 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.

The difference between the drift of the wealth share of any two households i and

j can be expressed as:

$$\begin{aligned} \mu_{\eta_j,t} - \mu_{\eta_i,t} = & \underbrace{(\theta_{j,t}^k - \theta_{i,t}^k)(r_t^k - r_t^l - \boldsymbol{\sigma}_{q,t}^k \cdot \boldsymbol{\sigma}_{q,t}^k)}_{=:(\mu_{\eta_j,t} - \mu_{\eta_i,t})^K} + \underbrace{(\theta_{j,t}^n - \theta_{i,t}^n)(r_t^n - r_t^l - \boldsymbol{\sigma}_{q,t}^n \cdot \boldsymbol{\sigma}_{q,t}^n)}_{=:(\mu_{\eta_j,t} - \mu_{\eta_i,t})^N} \\ & - (\omega_j - \omega_i) + \lambda_h \phi_h \left(\frac{1}{\eta_{j,t}} - \frac{1}{\eta_{i,t}} \right) \end{aligned} \quad (4.1)$$

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 models 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.

Figures 4 and 5 plot the decomposition of household inequality evolution for economies with different participation constraints (with respect to changing η_b and η_f respectively) and no fund participation in the capital market. This is the most restricted version of our economy. First, we can see households respond to an increase in wealth by increasing their portfolio share in capital and decreasing their share in pension/insurance products (as is seen stylistically in the data). This is because the capital market participation constraint is relaxed as the household accumulates wealth. We can also see that the household capital share increases much more quickly when the capital market participation constraint is lower ($\bar{\psi}$ is lower). How does this impact the evolution of inequality in general equilibrium? The effect is ambiguous. On the one hand, the economy with tighter capital market participation constraints

generates a higher capital risk premium so wealthier agents can earn a higher return. However, on the other hand, the economy with tighter capital market participation constraints also has less variation in household portfolios, which means the difference in exposure to the risk premium is smaller.

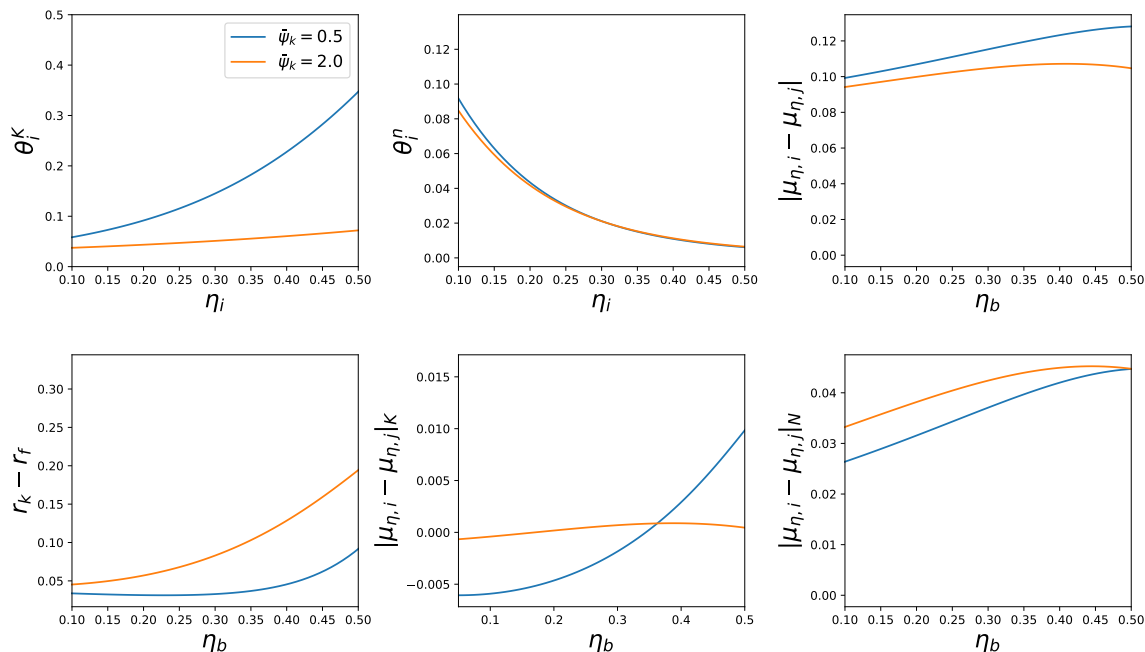


Figure 4: Inequality evolution decomposition for different participation constraints. The blue plot has $\bar{\psi} = 0.5$. The orange plot has $\bar{\psi} = 2.0$. The wealth distribution within the household sector is set to be equal. $\rho_e = 0.04, \rho_h = 0.03, \mu = 0.02, \sigma = 0.05$.

Figures 6 and 7 show the decomposition of inequality evolution for the unregulated economy (in orange) and the regulated economy where funds are restricted from accessing capital markets. We can see that the regulated economy leads to less drift in inequality when household participation constraints are high.

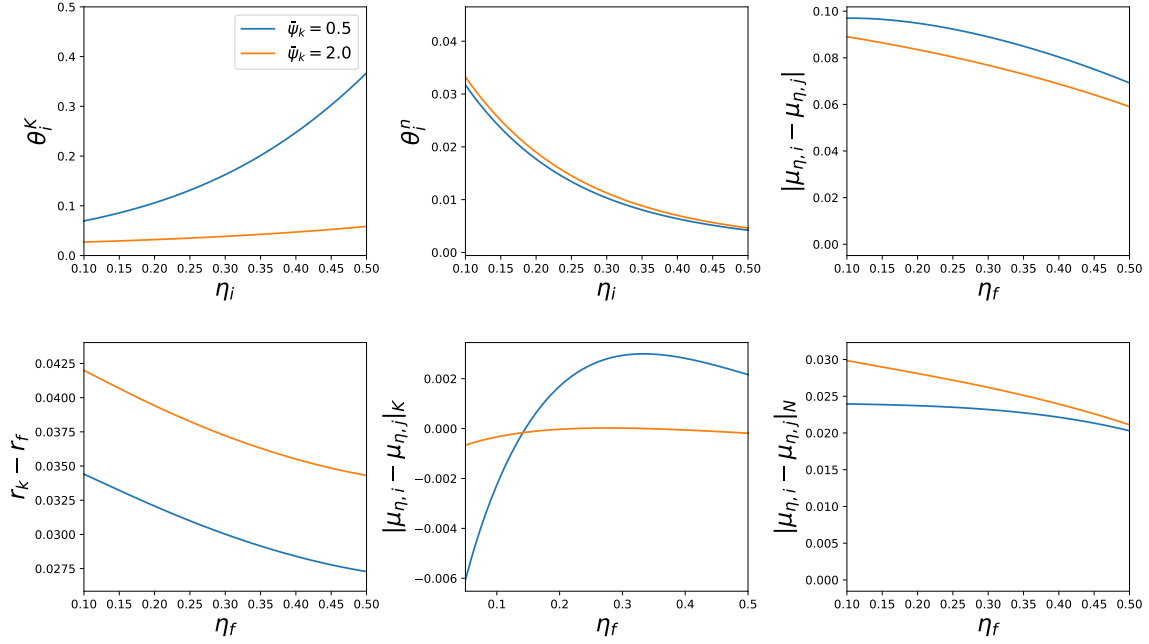


Figure 5: Inequality evolution decomposition for different participation constraints. The blue plot has $\bar{\psi} = 0.5$. The orange plot has $\bar{\psi} = 2.0$. The wealth distribution within the household sector is set to be equal. $\rho_e = 0.04, \rho_h = 0.03, \mu = 0.02, \sigma = 0.05$.

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.

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 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, which we attempt to match in our 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*.

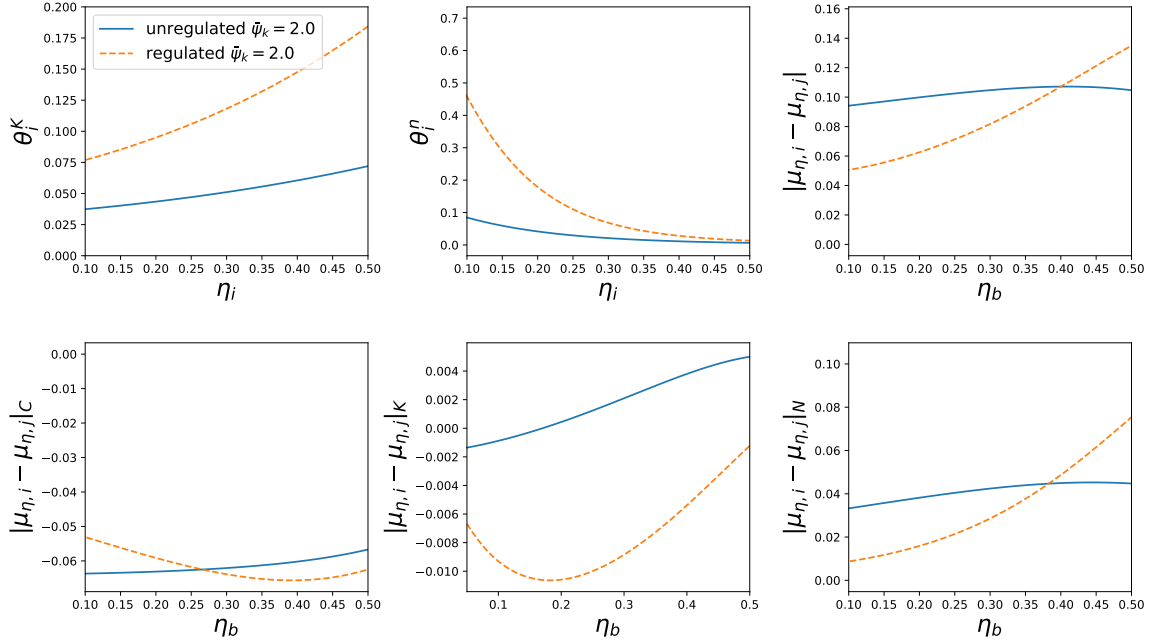


Figure 6: Decomposition of inequality for the regulated and unregulated economy. The blue line shows the The blue plot has $\bar{\psi} = 0.2$. The orange plot has $\bar{\psi} = 1.0$. The green plot has $\bar{\psi} = 2.0$. Wealth distribution within the household sector is set to be equal. $\rho_e = 0.04, \rho_h = 0.03, \mu = 0.02, \sigma = 0.05$ (the same below for all figures in this section).

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:

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

where $R_{t \rightarrow t+k}$ is the cumulative k-period future returns, r_f is the risk free rate, and dp_t is the dividend yield.³ For the baseline specification, we use $k = 1$ but the results do not materially change for other values.

³For robustness, we estimate the risk premium using the Fama-French three factor model instead of dividend yield and get similar results.

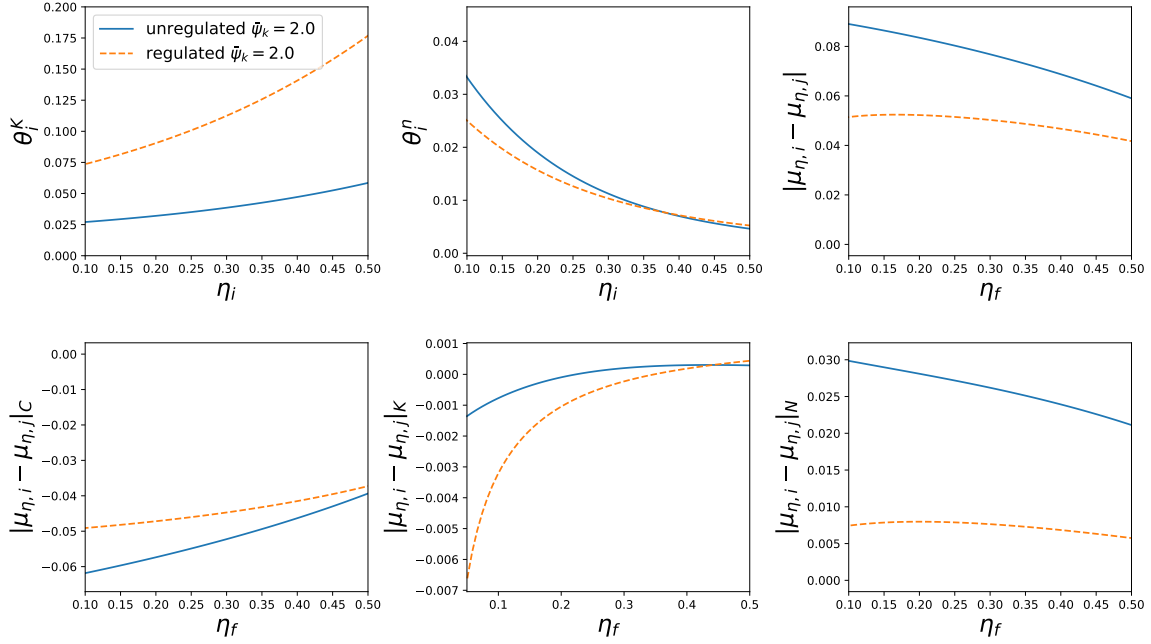


Figure 7: Decomposition of inequality for the regulated and unregulated economy. The blue line shows the The blue plot has $\bar{\psi} = 0.2$. The orange plot has $\bar{\psi} = 1.0$. The green plot has $\bar{\psi} = 2.0$. Wealth distribution within the household sector is set to be equal. $\rho_e = 0.04, \rho_h = 0.03, \mu = 0.02, \sigma = 0.05$ (the same below for all figures in this section).

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

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

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 8 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 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 8 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 8. Evidently, the conditional effect of risk premium on wealth is larger for top wealth percentiles.

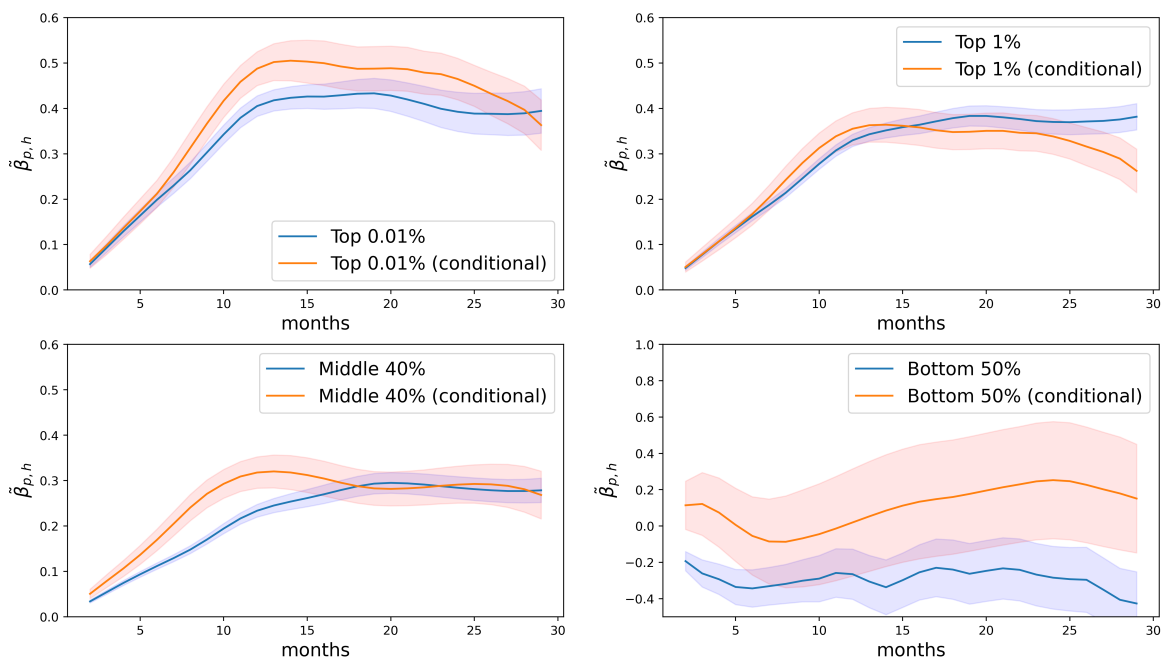


Figure 8: The figure plots the impulse response of wealth distribution to risk premium ($\beta_{p,h}$) obtained from the regression $\log(W_{p,t+h}/W_{p,t}) = \alpha_{p,h} + \beta_{p,h}rp_t^K + \epsilon_{p,t+h}$. The red lines are the conditional impulse response of wealth distribution to risk premium ($\tilde{\beta}_{p,h}$) obtained from the regression $\log(W_{p,t+h}/W_{p,t}) = \alpha_{p,h} + \tilde{\beta}_{p,h}rp_t^K \times 1_{Rec} + \epsilon_{p,t+h}$. The data for wealth percentiles come from [Saez and Zucman \(2016\)](#), and risk premium is estimated using a factor model.




Figure 9: The left panel plots plots the impulse response of wealth distribution to risk premium ($\beta_{BHC,h}$) obtained from the regression $\log(W_{BHC,t+h}/W_{BHC,t}) = \alpha_{BHC,h} + \beta_{BHC,h}r_t^K + \epsilon_{BHC,t+h}$. The right panel plots ($\beta_{BHC,h}$) obtained from the regression $\log(W_{BHC,t+h}/W_{BHC,t}) = \alpha_{BHC,h} + \hat{\beta}_{BHC,h}r_t^K \times 1_{Rec} + \epsilon_{BHC,t+h}$. The right panel plots the ratio of conditional exposure to unconditional exposure $\hat{\beta}_{p,h}/\beta_{p,h}$ for top three wealth distribution percentiles $p \in \{0.01, 0.1, 10.0\}$.

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.). Financial intermediaries 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 0.2.⁴ The portfolio constraint parameter is calibrated to generate a 32% portfolio share from the middle income households.

⁴While 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	Financial sector 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 [Güvenen \(2009\)](#), [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%	<i>Amit Goyal's</i> website
Std[Equity returns]	19.3%	13%	<i>Amit Goyal's</i> website
E[Risk-free rate]	4.3%	4.5%	<i>Amit Goyal's</i> 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

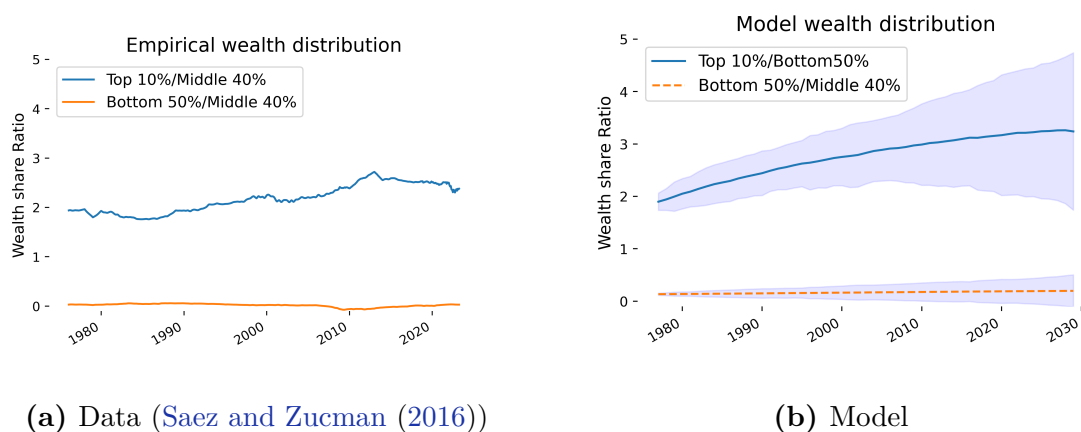


Figure 10: The left panel presents the share of total wealth for the households in top 1%, middle 40%, and bottom 50%, respectively. The time period is from 1976 till 2023 at monthly frequency. The data is taken from Saez and Zucman (2016). The right panel presents the share of total wealth produced 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 10 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 10 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 in 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 11: 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 11 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.

6 Conclusion

In this paper, we have studied the feedback between intermediary 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.

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.
- Blanchard, Olivier J**, “Debt, deficits, and finite horizons,” *Journal of political economy*, 1985, *93* (2), 223–247.
- 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.
- Chan, Yeung Lewis and Leonid Kogan**, “Catching up with the Joneses: Heterogeneous preferences and the dynamics of asset prices,” *Journal of Political Economy*, 2002, *110* (6), 1255–1285.
- 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 Guin-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.
- 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.
- Koijen, Ralph S. J. and Motohiro Yogo**, “Understanding the Ownership Structure of Corporate Bonds,” *American Economic Review: Insights*, 3 2023, 5, 73–92.
- Koijen, Ralph SJ and Motohiro Yogo**, “A demand system approach to asset pricing,” *Journal of Political Economy*, 2019, 127 (4), 1475–1515.
- 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.
- 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.
- Vayanos, Dimitri and Jean-Luc Vila**, “A Preferred-Habitat Model of the Term Structure of Interest Rates,” *Econometrica*, 1 2021, 89, 77–112.
- 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

The finite dimensional aggregate state variables are: $\mathbf{s} := (z, \zeta, K, \eta_b, \eta_f)$, where η_j is the fraction of wealth held by sector $j \in \{h, b, f\}$. We also have that the infinite dimensional state variable: g_h . With some abuse of notation, we let $G = \{\eta_b, \eta_f, g_h\}$ denote the collection of distribution state variables in the economy. The belief about prices becomes the belief about the evolution of sector wealth, $(\tilde{\mu}_{\eta,j}, \tilde{\sigma}_{\eta,j})_{j \in \{h, f\}}$, and the evolution of the household measure function g_h . For convenience, we define the following notation (all summarized here):

$$\begin{aligned}\mathbf{s} &:= (z, \zeta, K, \eta_b, \eta_f) \\ \mathbf{S} &:= (\mathbf{s}, g_h) \\ \mathbf{x} &:= (a, z, \zeta, K, \eta_b, \eta_f) \\ \mathbf{X} &:= (\mathbf{x}, g_h)\end{aligned}$$

Let $V_i(\mathbf{X})$ denote the value function for agent of type $i \in \{h, b, f\}$ with state variable \mathbf{X} .

In matrix form, the evolution of \mathbf{x}_t is (under household beliefs):

$$d\mathbf{x}_t = (\boldsymbol{\mu}_x(\mathbf{x}_t, c_h, \boldsymbol{\theta}_h) \odot \mathbf{x}_t)dt + (\boldsymbol{\sigma}_x(\mathbf{x}_t, \boldsymbol{\theta}_h) \odot \mathbf{x}_t)^T d\mathbf{W}_t$$

where

$$\boldsymbol{\mu}_x(\mathbf{x}, c_h, \boldsymbol{\theta}_h) = \begin{bmatrix} \mu_a(x, c_h, \boldsymbol{\theta}_h) \\ \mu_z \\ \mu_\zeta \\ \mu_K \\ \tilde{\mu}_{\eta_h} \\ \tilde{\mu}_{\eta_f} \end{bmatrix}, \quad \boldsymbol{\sigma}_x(\mathbf{x}, \boldsymbol{\theta}_h)^T = \begin{bmatrix} \sigma_{a,z}(\mathbf{x}, \boldsymbol{\theta}_h) & \sigma_{a,\zeta}(\mathbf{x}, \boldsymbol{\theta}_h) \\ \sigma_z & 0 \\ 0 & \sigma_\zeta \\ 0 & 0 \\ \tilde{\sigma}_{\eta_h,z} & \tilde{\sigma}_{\eta_h,\zeta} \\ \tilde{\sigma}_{\eta_f,z} & \tilde{\sigma}_{\eta_f,\zeta} \end{bmatrix}$$

where $\boldsymbol{\sigma}_y = [\sigma_{y,z}, \sigma_{y,\zeta}]^T$ is the vector of volatilities for variable y . And (dropping

explicit dependence on \mathbf{x} to save space) we have:

$$\begin{aligned} \Sigma &= \sigma_x \sigma_x^T(\mathbf{x}, \boldsymbol{\theta}_h) \\ &= \begin{bmatrix} \boldsymbol{\sigma}_a^T \boldsymbol{\sigma}_a(\boldsymbol{\theta}_h) & \sigma_{a,z}(\boldsymbol{\theta}_h) \sigma_z & \sigma_{a,\zeta}(\boldsymbol{\theta}_h) \sigma_\zeta & 0 & \boldsymbol{\sigma}_a^T(\boldsymbol{\theta}_h) \tilde{\boldsymbol{\sigma}}_{\eta_h} & \boldsymbol{\sigma}_a^T(\boldsymbol{\theta}_h) \tilde{\boldsymbol{\sigma}}_{\eta_f} \\ \sigma_z \sigma_{a,z}(\boldsymbol{\theta}_h) & \sigma_z^2 & 0 & 0 & \sigma_z \tilde{\boldsymbol{\sigma}}_{\eta_h,z} & \sigma_z \tilde{\boldsymbol{\sigma}}_{\eta_f,z} \\ \sigma_\zeta \sigma_{a,\zeta}(x, \boldsymbol{\theta}_h) & 0 & \sigma_\zeta^2 & 0 & \sigma_\zeta \tilde{\boldsymbol{\sigma}}_{\eta_h,\zeta} & \sigma_\zeta \tilde{\boldsymbol{\sigma}}_{\eta_f,\zeta} \\ 0 & 0 & 0 & 0 & 0 & 0 \\ \tilde{\boldsymbol{\sigma}}_{\eta_h}^T \boldsymbol{\sigma}_a & \tilde{\boldsymbol{\sigma}}_{\eta_h,z} \sigma_z & \tilde{\boldsymbol{\sigma}}_{\eta_h,\zeta} \sigma_\zeta & 0 & \tilde{\boldsymbol{\sigma}}_{\eta_h}^T \tilde{\boldsymbol{\sigma}}_{\eta_h} & \tilde{\boldsymbol{\sigma}}_{\eta_h}^T \tilde{\boldsymbol{\sigma}}_{\eta_f} \\ \tilde{\boldsymbol{\sigma}}_{\eta_f}^T \boldsymbol{\sigma}_a & \tilde{\boldsymbol{\sigma}}_{\eta_f,z} \sigma_z & \tilde{\boldsymbol{\sigma}}_{\eta_f,\zeta} \sigma_\zeta & 0 & \tilde{\boldsymbol{\sigma}}_{\eta_f}^T \tilde{\boldsymbol{\sigma}}_{\eta_h} & \tilde{\boldsymbol{\sigma}}_{\eta_f}^T \tilde{\boldsymbol{\sigma}}_{\eta_f} \end{bmatrix} \end{aligned}$$

We use analogous notation for the law of motion for \mathbf{s}_t . Agent's belief about the law of motion of $g_{h,t}$ is denoted by:

$$dg_{h,t}(a) = \tilde{\mu}_g(a, \mathbf{S}) dt + \tilde{\boldsymbol{\sigma}}_g(a, \mathbf{S})^T d\mathbf{W}_t$$

A.1 Household Optimization

HBJE: Given beliefs about the evolution of the wealth shares, the household value function $V_h(a, \cdot)$ solves the HJBE (A.1) below (written in matrix form):

$$\begin{aligned} \rho_h V_h(\mathbf{X}) &= \max_{c_h, \boldsymbol{\theta}_h, \iota_h} \left\{ u(c_h) + (\psi_{h,k}(\theta_h^k) + \psi_{h,n}(\theta_h^n)) \Xi_h a \right. \\ &\quad + \lambda \left(\mathcal{U} \left(a \mathcal{W}(\theta_{h,t}^k, \theta_h^n) \right) - V_h(\mathbf{X}) \right) \\ &\quad + (\boldsymbol{\mu}_x(\mathbf{x}, c_h, \boldsymbol{\theta}_h, \iota_h) \odot \mathbf{x})^T D_x V_h(\mathbf{X}) \\ &\quad + \frac{1}{2} \text{tr} \left\{ (\sigma_x(\mathbf{x}, \boldsymbol{\theta}_h) \odot \mathbf{x})^T (\sigma_x(\mathbf{x}, \boldsymbol{\theta}_h) \odot \mathbf{x}) D_x^2 V_h(\mathbf{X}) \right\} \\ &\quad \left. + \mathcal{L}_g V_h(\mathbf{X}) \right\} \end{aligned} \tag{A.1}$$

where $\mathcal{L}_g V_h(\mathbf{X})$ is the collection of Frechet derivative terms and:

$$D_x V_h(\mathbf{x}) = \begin{bmatrix} \partial_a V_h(\mathbf{x}) \\ \partial_z V_h(\mathbf{x}) \\ \partial_\zeta V_h(\mathbf{x}) \\ \partial_K V_h(\mathbf{x}) \\ \partial_{A_h} V_h(\mathbf{x}) \\ \partial_{A_b} V_h(\mathbf{x}) \\ \partial_{A_f} V_h(\mathbf{x}) \end{bmatrix}, \quad D_x^2 V_h(\mathbf{x}) = \begin{bmatrix} \partial_{aa}^2 V_h & \dots & \partial_{aA_f}^2 V_h \\ \partial_{za}^2 V_h & \dots & \partial_{zA_f}^2 V_h \\ \partial_{\zeta a}^2 V_h & \dots & \partial_{\zeta A_f}^2 V_h \\ \partial_{Ka}^2 V_h & \dots & \partial_{KA_f}^2 V_h \\ \partial_{A_h a}^2 V_h & \dots & \partial_{A_h A_f}^2 V_h \\ \partial_{A_b a}^2 V_h & \dots & \partial_{A_b A_f}^2 V_h \\ \partial_{A_f a}^2 V_h & \dots & \partial_{A_f A_f}^2 V_h \end{bmatrix}$$

HBJE (partial matrix): We can rewrite the HJBE with the controlled variables taken out of the matrices. Then the HJBE is given by:

$$\begin{aligned} \rho_h V_h(a, \mathbf{s}, g_h) &= \max_{c_h, \boldsymbol{\theta}_h, \iota_h} \left\{ u(c_h) + (\psi_{h,k}(\theta_h^k) + \psi_{h,n}(\theta_h^n)) \Xi_h a \right. \\ &+ \lambda \left(\mathcal{U}(a \mathcal{W}(\theta_{h,t}^k, \theta_h^n)) - V_h(a, \mathbf{s}, g_h) \right) \\ &+ \mu_a(a, \mathbf{s}, c_h, \boldsymbol{\theta}_h, \iota_h) a + (\boldsymbol{\mu}_s(\mathbf{s}, g_h) \odot \mathbf{s})^T D_s V_h(a, \mathbf{s}, g_h) \\ &+ \frac{1}{2} \partial_{aa}^2 V_h(a, \mathbf{s}, g_h) \boldsymbol{\sigma}_a^T(\boldsymbol{\theta}_h, \mathbf{s}, g_h) \boldsymbol{\sigma}_a(\boldsymbol{\theta}_h, \mathbf{s}, g_h) a^2 + \sum_j \partial_{as_j} V_h(a, \mathbf{s}, g_h) \boldsymbol{\sigma}_a^T(\boldsymbol{\theta}_h, \mathbf{s}, g_h) \boldsymbol{\sigma}_{s_j} a s_j \\ &\left. + \frac{1}{2} \text{tr} \left\{ (\boldsymbol{\sigma}_s(\mathbf{s}, g_h) \odot \mathbf{s})^T (\boldsymbol{\sigma}_s(\mathbf{s}, g_h) \odot \mathbf{s}) D_s^2 V_h(a, \mathbf{s}, g_h) \right\} + \mathcal{L}_g V_h(a, \mathbf{s}, g_h) \right\} \end{aligned}$$

where:

$$\begin{aligned} \psi_h(\theta_{h,t}^k) &= \frac{\bar{\psi}_k}{2} (\theta_{h,t}^k)^2 \\ \psi_h(\theta_{h,t}^n) &= \frac{\bar{\psi}_n}{2} (\theta_{h,t}^n - \chi)^2 \\ \mathcal{W}(\theta_{h,t}^k, \theta_h^n) &= (\theta_h^k (1 - \tau))^\alpha ((q^n)^{-1} \theta_h^n)^{1-\alpha} \\ \mu_a(a, c_h, \boldsymbol{\theta}_h, \cdot) &= \left(\tilde{r}_t^d + \theta_{h,t}^n (\tilde{r}_t^n - \tilde{r}_t^d) + \theta_{h,t}^k (\tilde{r}_t^k - \tilde{r}_t^d) - c_{h,t} / a_{h,t} - \tau_{h,t} \right) \\ \boldsymbol{\sigma}_a^T(\boldsymbol{\theta}_h, \cdot) &= \boldsymbol{\theta}_h^T \tilde{\sigma}_q = \left[\theta_h^n \tilde{\sigma}_{q^n, z} + \theta_h^k \tilde{\sigma}_{q^k, z}, \theta_h^n \tilde{\sigma}_{q^n, \zeta} + \theta_h^k \tilde{\sigma}_{q^k, \zeta} \right] \end{aligned}$$

and so:

$$\begin{aligned}
(\boldsymbol{\sigma}_a^T \boldsymbol{\sigma}_a)(a, \boldsymbol{\theta}_h, \cdot) &= \boldsymbol{\theta}_h^T \tilde{\boldsymbol{\sigma}}_q \tilde{\boldsymbol{\sigma}}_q^T \boldsymbol{\theta}_h \\
&= (\theta_h^n \tilde{\sigma}_{q^n, z} + \theta_{h, t}^k \tilde{\sigma}_{q^k, z})^2 + (\theta_{h, t}^n \tilde{\sigma}_{q^n, \zeta} + \theta_{h, t}^k \tilde{\sigma}_{q^k, \zeta})^2 \\
\boldsymbol{\sigma}_a^T(a, \boldsymbol{\theta}_h, \cdot) \tilde{\boldsymbol{\sigma}}_{\eta_j} &= \boldsymbol{\theta}_h^T \tilde{\boldsymbol{\sigma}}_q \tilde{\boldsymbol{\sigma}}_{\eta_j}^T \\
&= (\theta_{h, t}^n \tilde{\sigma}_{q^n, z} + \theta_{h, t}^k \tilde{\sigma}_{q^k, z}) \tilde{\sigma}_{\eta_j, z} + (\theta_{h, t}^n \tilde{\sigma}_{q^n, \zeta} + \theta_{h, t}^k \tilde{\sigma}_{q^k, \zeta}) \tilde{\sigma}_{\eta_j, \zeta} \\
\sum_j \partial_{as_j} V_h(a, \mathbf{s}) \boldsymbol{\sigma}_a^T(\boldsymbol{\theta}_h, \mathbf{s}) \boldsymbol{\sigma}_{s_j} &= \partial_{az}^2 V_h(a, \mathbf{s}) \sigma_{a, z}(\boldsymbol{\theta}_h, \mathbf{s}) \sigma_z a z \\
&\quad + \partial_{a\zeta}^2 V_h(a, \mathbf{s}) \sigma_{a, \zeta}(\boldsymbol{\theta}_h, \mathbf{s}) \sigma_\zeta a \zeta + \sum_j \partial_{aA_j}^2 V_h(a, \mathbf{s}) \boldsymbol{\sigma}_a^T(\boldsymbol{\theta}_h, \mathbf{s}) \tilde{\boldsymbol{\sigma}}_{A_j}(\mathbf{s}) a \eta_j
\end{aligned}$$

The HJBE becomes:

$$\begin{aligned}
\rho_h V_h(a, \mathbf{s}) &= \max_{c_h, \boldsymbol{\theta}_h, \iota_h} \left\{ u(c_h) + (\psi_{h, k}(\theta_h^k) + \psi_{h, n}(\theta_h^n)) \Xi_h a \right. \\
&\quad + \lambda \left(\mathcal{U} \left(a \mathcal{W} \left(\theta_h^k, \theta_h^n \right) \right) - V_h(a, \mathbf{s}) \right) \\
&\quad + \mu_a(a, \mathbf{s}, c_h, \boldsymbol{\theta}_h, \iota_h) a + (\boldsymbol{\mu}_s(\mathbf{s}) \odot \mathbf{s})^T D_s V_h(a, \mathbf{s}) \\
&\quad + \frac{1}{2} \partial_{aa}^2 V_h(a, \mathbf{s}) \boldsymbol{\theta}_h^T \tilde{\boldsymbol{\sigma}}_q \tilde{\boldsymbol{\sigma}}_q^T \boldsymbol{\theta}_h a^2 + \sum_j \partial_{as_j} V_h(a, \mathbf{s}) \boldsymbol{\theta}_h^T \tilde{\boldsymbol{\sigma}}_q \boldsymbol{\sigma}_{s_j} a s_j \\
&\quad \left. + \frac{1}{2} \text{tr} \left\{ (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s})^T (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s}) D_s^2 V_h(a, \mathbf{s}) \right\} + \mathcal{L}_g V_h(a, \mathbf{s}, g_h) + \mathcal{L}_g V_h(a, \mathbf{s}, g_h) \right\}
\end{aligned}$$

FOCs: The first order conditions are given by (as as the problem with out the

household distribution):

$$\begin{aligned}
[c_h] : \quad & 0 = u'(c_h) - \partial_a V_h \\
[t_h] : \quad & 0 = \Phi'(t) - \frac{1}{q_t^k} \\
[\theta_h^k] : \quad & 0 = \partial_a V_h(\tilde{r}^k - \tilde{r}^d)a + \lambda \partial_{\theta_h^k} \mathcal{W}(\theta_h^k, \theta_h^n) \mathcal{U}'(\mathcal{C}) \\
& \quad + \partial_{\theta_h^k} \psi_{h,k} \Xi_h a + \partial_{aa} V_h \tilde{\sigma}_{q^k}^T \sigma_a a^2 + \sum_j \partial_{as_j} V_h \tilde{\sigma}_{q^k}^T \sigma_{s_j} a s_j \\
& \quad = \partial_a V_h(\tilde{r}^k - \tilde{r}^d)a + \lambda \partial_{\theta_h^k} \mathcal{W}(\theta_h^k, \theta_h^n) \mathcal{U}'(\mathcal{C}) \\
& \quad \quad + \partial_{\theta_h^k} \psi_{h,k} \Xi_h a + (D_x(\partial_a V_h))^T (\sigma_x \odot \mathbf{x})^T \tilde{\sigma}_{q^k} a \\
[\theta_h^n] : \quad & 0 = \partial_a V_h(\tilde{r}^n - \tilde{r}^d)a + \lambda \partial_{\theta_h^n} \mathcal{W}(\theta_h^k, \theta_h^n) \mathcal{U}'(\mathcal{C}) \\
& \quad + \partial_{\theta_h^n} \psi_{h,n} \Xi_h a + \partial_{aa} V_h \tilde{\sigma}_{q^n}^T \sigma_a a^2 + \sum_j \partial_{as_j} V_h \tilde{\sigma}_{q^n}^T \sigma_{s_j} a s_j \\
& \quad = \partial_a V_h(\tilde{r}^n - \tilde{r}^d)a + \lambda \partial_{\theta_h^n} \mathcal{W}(\theta_h^k, \theta_h^n) \mathcal{U}'(\mathcal{C}) \\
& \quad \quad + \partial_{\theta_h^n} \psi_{h,k} \Xi_h a + (D_x(\partial_a V_h))^T (\sigma_x \odot \mathbf{x})^T \tilde{\sigma}_{q^n} a
\end{aligned}$$

SDF Evolution: Let $\xi_h(\mathbf{X}) := \partial_a V_h(a, \mathbf{s}, g_h)$. From Ito's Lemma, we have that the drift and volatility of ξ_h are given by:

$$\begin{aligned}
\mu_{\xi_h} \xi_h(\mathbf{X}) &= (D_x \xi_h(\mathbf{X}))^T \boldsymbol{\mu}_x \\
& \quad + \frac{1}{2} \text{tr} \left\{ (\sigma_x(\mathbf{X}, \boldsymbol{\theta}_h) \odot \mathbf{X})^T (\sigma_x(\mathbf{X}, \boldsymbol{\theta}_h) \odot \mathbf{x}) D_x^2 \xi_h(\mathbf{X}) \right\} \\
& \quad + \mathcal{L}_g \xi_h(\mathbf{X}) \\
&= \partial_a \xi_h(a, \mathbf{s}, g_h) \mu_a(a, c_h, \boldsymbol{\theta}_h, \mathbf{s}, g_h) a \\
& \quad + (D_s \xi_h(a, \mathbf{s}, g_h))^T (\boldsymbol{\mu}_s(\mathbf{s}) \odot \mathbf{s}) \\
& \quad + \frac{1}{2} \partial_{aa}^2 \xi_h(a, \mathbf{s}, g_h) \sigma_a^T(\boldsymbol{\theta}_h, \mathbf{s}, g_h) \sigma_a(\boldsymbol{\theta}_h, \mathbf{s}, g_h) a^2 \\
& \quad + \sum_j \partial_{as_j} \xi_h(a, \mathbf{s}, g_h) \sigma_a^T(\boldsymbol{\theta}_h, \mathbf{s}, g_h) \sigma_{s_j} a s_j \\
& \quad + \frac{1}{2} \text{tr} \left\{ (\sigma_s(\mathbf{s}) \odot \mathbf{s})^T (\sigma_s(\mathbf{s}) \odot \mathbf{s}) D_s^2 \xi_h(a, \mathbf{s}, g_h) \right\} + \mathcal{L}_g \xi_h(\mathbf{X}) \\
\sigma_{\xi_h} \xi_h &= (\sigma_x \odot \mathbf{x})^T (D_x \xi_h) \\
&= \begin{bmatrix} \partial_a \xi_h \sigma_{a,z} a + \partial_z \xi_h \sigma_z z + \sum_j \partial_{a_j} \xi_h \sigma_{a_j, z} \eta_j \\ \partial_a \xi_h \sigma_{a,\zeta} a + \partial_\zeta \xi_h \sigma_\zeta \zeta + \sum_j \partial_{a_j} \xi_h \sigma_{a_j, \zeta} \eta_j \end{bmatrix}
\end{aligned}$$

Thus, we can rewrite the FOCs as:

$$\begin{aligned}
[\theta_h^k] : \quad 0 &= \xi_h(\tilde{r}^n - \tilde{r}^d) + \lambda \partial_{\theta_h^k} \mathcal{W}(\theta_h^k, \theta_h^n) \frac{\mathcal{U}'(\mathcal{C})}{a} \\
&\quad + \partial_{\theta_h^k} \psi_{h,k} \Xi_h + (\boldsymbol{\sigma}_{\xi_h} \xi_h)^T \boldsymbol{\sigma}_{q^k} \\
[\theta_h^n] : \quad 0 &= \xi_h(\tilde{r}^n - \tilde{r}^d) + \lambda \partial_{\theta_h^n} \mathcal{W}(\theta_h^k, \theta_h^n) \frac{\mathcal{U}'(\mathcal{C})}{a} + \partial_{\theta_h^n} \psi_{h,n} \Xi_h + (\boldsymbol{\sigma}_{\xi_h} \xi_h)^T \boldsymbol{\sigma}_{q^n}
\end{aligned}$$

Imposing belief consistency and using the equilibrium result that $\Xi_h = \xi_h$, we get the simplified FOCs:

$$\begin{aligned}
[\theta_h^k] : \quad r^k - r^d &= -\lambda \partial_{\theta_h^k} \mathcal{W}(\theta_h^k, \theta_h^n) \frac{\mathcal{U}'(\mathcal{C})}{a \xi_h} - \partial_{\theta_h^k} \psi_{h,k} - \boldsymbol{\sigma}_{\xi_h}^T \boldsymbol{\sigma}_{q^k} \\
[\theta_h^n] : \quad r^n - r^d &= -\lambda \partial_{\theta_h^n} \mathcal{W}(\theta_h^k, \theta_h^n) \frac{\mathcal{U}'(\mathcal{C})}{a \xi_h} - \partial_{\theta_h^n} \psi_{h,n} - \boldsymbol{\sigma}_{\xi_h}^T \boldsymbol{\sigma}_{q^n}
\end{aligned}$$

Euler equation: We close this section by using the Envelope theorem to get the Euler equation (using the partial matrix representation). To do this, we treat θa as the

choice rather than θ when taking the envelope theorem. This gives:

$$\begin{aligned}
\rho_h \xi_h(a, \mathbf{s}) &= (\psi_{h,k}(\theta_h^k a/a) + \psi_{h,n}(\theta_h^n a/a)) \Xi_h \\
&\quad - \left(\partial_{\theta_h^k} \psi_{h,k}(\theta_h^k a/a) \frac{\theta_h^k}{a} + \partial_{\theta_h^n} \psi_{h,n}(\theta_h^n a/a) \frac{\theta_h^n}{a} \right) \Xi_h a \\
&\quad - \lambda \xi_h(a, \mathbf{s}) + \partial_a \xi_h(a, \mathbf{s}) \mu_a(a, c_h, \boldsymbol{\theta}_h, \cdot) a + \xi_h(a, \mathbf{s}) (r^d + \tau_h) \\
&\quad + (D_s \xi_h(a, \mathbf{s}))^T (\boldsymbol{\mu}_s(\mathbf{s}) \odot \mathbf{s}) \\
&\quad + \frac{1}{2} \partial_{aa}^2 \xi_h(a, \mathbf{s}) \boldsymbol{\sigma}_a^T(\boldsymbol{\theta}_h, \mathbf{s}) \boldsymbol{\sigma}_a(\boldsymbol{\theta}_h, \mathbf{s}) a^2 \\
&\quad + \sum_j \partial_{as_j} \xi_h(a, \mathbf{s}) \boldsymbol{\sigma}_a^T(\boldsymbol{\theta}_h, \mathbf{s}) \boldsymbol{\sigma}_{s_j} a s_j \\
&\quad + \frac{1}{2} \text{tr} \left\{ (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s})^T (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s}) D_s^2 \xi_h(a, \mathbf{s}) \right\} + \mathcal{L}_g \xi_h \\
&= (\psi_{h,k}(\theta_h^k) - \partial_{\theta_h^k} \psi_{h,k}(\theta_h^k a/a) \theta_h^k + \psi_{h,n}(\theta_h^n) - \partial_{\theta_h^n} \psi_{h,n}(\theta_h^n a/a) \theta_h^n) \Xi_h \\
&\quad - \lambda \xi_h(a, \mathbf{s}) + (D_s \xi_h(\mathbf{x}))^T (\boldsymbol{\mu}_x(\mathbf{x}) \odot \mathbf{x}) + \xi_h(a, \mathbf{s}) (r^d + \tau_h) \\
&\quad + \frac{1}{2} \text{tr} \left\{ (\boldsymbol{\sigma}_x(\mathbf{x}, \boldsymbol{\theta}_h) \odot \mathbf{x})^T (\boldsymbol{\sigma}_x(\mathbf{x}, \boldsymbol{\theta}_h) \odot \mathbf{x}) D_x^2 \xi_h(\mathbf{x}) \right\} + \mathcal{L}_g \xi_h \\
&= (\psi_{h,k}(\theta_h^k) - \partial_{\theta_h^k} \psi_{h,k}(\theta_h^k) \theta_h^k + \psi_{h,n}(\theta_h^n) - \partial_{\theta_h^n} \psi_{h,n}(\theta_h^n) \theta_h^n) \Xi_h \\
&\quad - \lambda \xi_h(a, \mathbf{s}) + \mu_{\xi_h} \xi_h(a, \mathbf{s}) + \xi_h(a, \mathbf{s}) (r^d - \tau_h)
\end{aligned}$$

So, imposing belief consistency and we get that:

$$\rho + \lambda = \mu_{\xi_h} + r^d - \tau_h + \psi_{h,k}(\theta_h^k) - \partial_{\theta_h^k} \psi_{h,k}(\theta_h^k) \theta_h^k + \psi_{h,n}(\theta_h^n) - \partial_{\theta_h^n} \psi_{h,n}(\theta_h^n) \theta_h^n$$

A.1.1 Banker Optimization

The banker HJBE is given by:

$$\begin{aligned}
\rho_b V_b(a, \mathbf{s}, g_h) &= \max_{c_h, \boldsymbol{\theta}_h, \iota_h} \left\{ u(c_b) + \mu_a(a, \mathbf{s}, c_b, \boldsymbol{\theta}_b, \iota_b) a + (\boldsymbol{\mu}_s(\mathbf{s}) \odot \mathbf{s})^T D_s V_b(a, \mathbf{s}, g_h) \right. \\
&\quad + \frac{1}{2} \partial_{aa}^2 V_b(a, \mathbf{s}, g_h) \boldsymbol{\sigma}_a^T(\boldsymbol{\theta}_b, \mathbf{s}) \boldsymbol{\sigma}_a(\boldsymbol{\theta}_b, \mathbf{s}) a^2 + \sum_j \partial_{as_j} V_b(a, \mathbf{s}, g_h) \boldsymbol{\sigma}_a^T(\boldsymbol{\theta}_b, \mathbf{s}) \boldsymbol{\sigma}_{s_j} a s_j \\
&\quad \left. + \frac{1}{2} \text{tr} \left\{ (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s})^T (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s}) D_s^2 V_b(a, \mathbf{s}) \right\} + \mathcal{L}_g V_b(a, \mathbf{s}, g_h) \right\}
\end{aligned}$$

where:

$$\begin{aligned}\mu_a(a, c_b, \theta_b, \cdot) &= \left(\tilde{r}^d + \theta_h^k (\tilde{r}^k - \tilde{r}^d) - c_{b,t}/a_b - \tau_b \right) \\ \boldsymbol{\sigma}_a^T(\theta_b, \cdot) &= \left[\theta_b^k \tilde{\sigma}_{q^k, z}, \theta_b^k \tilde{\sigma}_{q^k, \zeta} \right] = \theta_b^k \tilde{\boldsymbol{\sigma}}_{q^k}\end{aligned}$$

Following the same steps, the equilibrium FOCs for the banker are given by:

$$\begin{aligned}[c_b] : & \quad 0 = u'(c_b) - \partial_a V_b(a, \mathbf{s}, g_h) \\ [\iota_b] : & \quad 0 = \Phi'(\iota) - \frac{1}{q_t^k} \\ [\theta_b^k] : & \quad 0 = r_t^k - r^d + \boldsymbol{\sigma}_{\xi_b}^T \boldsymbol{\sigma}_{q^k}\end{aligned}$$

and the Euler equation is:

$$\rho = \mu_{\xi_b} + r^d - \tau_b$$

We impose that $u(c_b) = \log(c_b)$ and so we can solve the Euler equation analytically to get:

$$\begin{aligned}c_b &= (\rho_b + \lambda_b) a_b \\ \iota &= (\phi')^{-1}(1/q^k) \\ \frac{da_{b,t}}{a_{b,t}} &= \left(r^d + \theta_b^k (r_t^k - r_t^d) - (\rho_b + \lambda_b) + \tau_{b,t} \right) dt + \boldsymbol{\theta}_b^T \tilde{\sigma}_{q,t} d\mathbf{W}_t\end{aligned}$$

where $\boldsymbol{\theta}_b^T$ is independent of a .

A.1.2 Fund Manager Optimization

The banker HJBE is given by:

$$\begin{aligned}\rho_f V_f(a, \mathbf{s}, g_h) &= \max_{c_h, \boldsymbol{\theta}_h, \iota_h} \left\{ u(c_f) + \mu_a(a, \mathbf{s}, c_f, \boldsymbol{\theta}_f, \iota_f) a + (\boldsymbol{\mu}_s(\mathbf{s}) \odot \mathbf{s})^T D_s V_f(a, \mathbf{s}, g_h) \right. \\ &+ \frac{1}{2} \partial_{aa}^2 V_f(a, \mathbf{s}) \boldsymbol{\sigma}_a^T(\boldsymbol{\theta}_f, \mathbf{s}, g_h) \boldsymbol{\sigma}_a(\boldsymbol{\theta}_f, \mathbf{s}, g_h) a^2 + \sum_j \partial_{as_j} V_f(a, \mathbf{s}, g_h) \boldsymbol{\sigma}_a^T(\boldsymbol{\theta}_f, \mathbf{s}) \boldsymbol{\sigma}_{s_j} a s_j \\ &\left. + \frac{1}{2} \text{tr} \left\{ (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s})^T (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s}) D_s^2 V_f(a, \mathbf{s}, g_h) \right\} + \mathcal{L}_g V_f(a, \mathbf{s}, g_h) \right\}\end{aligned}$$

where:

$$\begin{aligned}\mu_a(a, c_f, \boldsymbol{\theta}_f, \mathbf{s}, g_h) &= \tilde{r}_f^n + \theta_f^k(\tilde{r}^k - \tilde{r}_f^n) + \theta_f^m(\tilde{r}_t^m - \tilde{r}_f^n) - c_{f,t}/a_f - \tau_f \\ \boldsymbol{\sigma}_a(\boldsymbol{\theta}_f, \mathbf{s}, g_h) &= \begin{bmatrix} \theta_f^k \tilde{\sigma}_{q^k, z} + \theta_f^m \tilde{\sigma}_{q^m, z} + (1 - \theta_f^k - \theta_f^m) \tilde{\sigma}_{q^n, z}, \\ \theta_f^k \tilde{\sigma}_{q^k, \zeta} + \theta_f^m \tilde{\sigma}_{q^m, \zeta} + (1 - \theta_f^k - \theta_f^m) \tilde{\sigma}_{q^n, \zeta} \end{bmatrix} = \boldsymbol{\theta}_f^T \tilde{\boldsymbol{\sigma}}_{q,t}\end{aligned}$$

Following the same steps, the equilibrium FOCs for the fund are given by:

$$\begin{aligned}[c_f] : & \quad 0 = u'(c_f) - \partial_a V_f(a, \cdot) \\ [\iota_f] : & \quad 0 = \Phi'(\iota) - \frac{1}{q^k} \\ [\theta_f^k] : & \quad 0 = r^k - r_f^n + \boldsymbol{\sigma}_{\xi_f}^T (\boldsymbol{\sigma}_{q^k} - \boldsymbol{\sigma}_{q^n}) \\ [\theta_f^m] : & \quad 0 = r^m - r_f^n + \boldsymbol{\sigma}_{\xi_f}^T (\boldsymbol{\sigma}_{q^m} - \boldsymbol{\sigma}_{q^n})\end{aligned}$$

and the Euler equation is:

$$\rho = \mu_{\xi_f} + r_f^n - \tau_h$$

We impose that $u(c_f) = \log(c_f)$ and so we can solve the Euler equation analytically to get:

$$\begin{aligned}c_f &= (\rho_f + \lambda_f) a_f \\ \iota &= (\phi')^{-1}(1/q^k) \\ \frac{da_{f,t}}{a_{f,t}} &= \left(r^d + \theta_f^k (r_t^k - r_t^d) - (\rho_f + \lambda_f) + \tau_{f,t} \right) dt + \boldsymbol{\theta}_f^T \tilde{\boldsymbol{\sigma}}_{q,t} d\mathbf{W}_t\end{aligned}$$

where $\boldsymbol{\theta}_f^T$ is independent of a .

A.1.3 Equilibrium Functions

The agent optimization problem has the terms:

$$\xi_h(a, z, \zeta, K, g), \quad D_x \xi_h(a, z, \zeta, K, g) = \begin{bmatrix} \partial_a \xi_h(a, z, \zeta, K, g) \\ \partial_z \xi_h(a, z, \zeta, K, g) \\ \partial_\zeta \xi_h(a, z, \zeta, K, g) \\ \partial_K \xi_h(a, z, \zeta, K, g) \\ \partial_{\eta_h} \xi_h(a, z, \zeta, K, g) \\ \partial_{\eta_f} \xi_h(a, z, \zeta, K, g) \end{bmatrix}$$

In equilibrium we have that that $a = \eta_h A(\mathbf{s})$ where $A(\mathbf{s}) = q^k(\mathbf{s})K + q^m(\mathbf{s})M$:

$$\begin{aligned} \Xi_h(z, \zeta, g, K) &= \xi_h(a, z, \zeta, K, g)|_{a=\eta_h A(\mathbf{s})} \\ \Xi'_h(z, \zeta, g, K) &= D_x \xi_h(a, z, \zeta, K, g)|_{a=\eta_h A(\mathbf{s})} \end{aligned}$$

The term $\Xi'_i(z, \zeta, g, K)$ appears throughout the FOCs equations so we need approximate it. However, we have:

$$\Xi'_h(z, \zeta, g, K) \neq D_s \Xi_h(z, \zeta, g, K)$$

for the obvious reason that the dimension is different. Instead, we have that:

$$D_s \Xi_h(z, \zeta, g, K) = \begin{bmatrix} \partial_z \xi_h(\eta_h A(\mathbf{s}), z, \zeta, K, g) \\ \partial_\zeta \xi_h(\eta_h A(\mathbf{s}), z, \zeta, K, g) \\ \partial_K \xi_h(\eta_h A(\mathbf{s}), z, \zeta, K, g) \\ \partial_a \xi_h(a, z, \zeta, K, g)|_{a=\eta_h A(\mathbf{s})} A(\mathbf{s}) + \partial_{\eta_h} \xi_h(a, z, \zeta, K, g)|_{a=\eta_h A(\mathbf{s})} \\ \partial_{\eta_f} \xi_h(\eta_h A(\mathbf{s}), z, \zeta, K, g) \end{bmatrix}$$

Proposition 1. *In equilibrium, we have that for $j \in \{h, b, f\}$:*

$$\begin{aligned} \mu_{\xi_j} \xi_j(a, \mathbf{s})|_{a=\eta_j A(\mathbf{s})} &= \mu_{\Xi_j} \Xi_j(\mathbf{s}) \\ \sigma_{\xi_j} \xi_j(a, \mathbf{s})|_{a=\eta_j A(\mathbf{s})} &= \sigma_{\Xi_j} \Xi_j(\mathbf{s}) \end{aligned}$$

Proof. For clarity, I show this in the non-matrix form for the household rather than

using the matrix chain rule. For the volatility, we have that:

$$\begin{aligned}\boldsymbol{\sigma}_{\xi_h} \xi_h &= (\boldsymbol{\sigma}_x \odot \boldsymbol{x})^T (D_x \xi_h) \\ &= \begin{bmatrix} \partial_a \xi_h \sigma_{a,z} a + \partial_z \xi_h \sigma_z z + \sum_j \partial_{\eta_j} \xi_h \sigma_{\eta_j, z} \eta_j \\ \partial_a \xi_h \sigma_{a,\zeta} a + \partial_\zeta \xi_h \sigma_\zeta \zeta + \sum_j \partial_{\eta_j} \xi_h \sigma_{\eta_j, \zeta} \eta_j \end{bmatrix}\end{aligned}$$

After imposing equilibrium $a = \eta_1 A(\boldsymbol{s})$, where $A(\boldsymbol{s}) = q^k K + q^m M$, we have that the RHS is:

$$\begin{aligned}RHS &= \begin{bmatrix} \partial_a \xi_h \sigma_{a,z} \eta_1 A(\boldsymbol{s}) + \partial_z \xi_h \sigma_z z + \sum_j \partial_{\eta_j} \xi_h \sigma_{\eta_j, z} \eta_j \\ \partial_a \xi_h \sigma_{a,\zeta} \eta_1 A(\boldsymbol{s}) + \partial_\zeta \xi_h \sigma_\zeta \zeta + \sum_j \partial_{\eta_j} \xi_h \sigma_{\eta_j, \zeta} \eta_j \end{bmatrix} \\ &= \begin{bmatrix} \sigma_{a,z} & \sigma_{a,\zeta} \\ \sigma_z & 0 \\ 0 & \sigma_\zeta \\ 0 & 0 \\ \sigma_{\eta_h, z} & \sigma_{\eta_h, \zeta} \\ \sigma_{\eta_f, z} & \sigma_{\eta_f, \zeta} \end{bmatrix}^T \begin{bmatrix} \partial_z \xi_h(\eta_h A(\boldsymbol{s}), z, \zeta, K, g) \\ \partial_\zeta \xi_h(\eta_h A(\boldsymbol{s}), z, \zeta, K, g) \\ \partial_K \xi_h(\eta_h A(\boldsymbol{s}), z, \zeta, K, g) \\ \partial_a \xi_h(\eta_h A(\boldsymbol{s}), z, \zeta, K, g) A(\boldsymbol{s}) + \partial_{\eta_h} \xi_h(\eta_h A(\boldsymbol{s}), z, \zeta, K, g) \\ \partial_{\eta_f} \xi_h(\eta_h A(\boldsymbol{s}), z, \zeta, K, g) \end{bmatrix} \\ &= (\boldsymbol{\sigma}_s \odot \boldsymbol{s})^T (D_s \Xi_h) \\ &= \boldsymbol{\sigma}_{\Xi_h} \Xi_h\end{aligned}$$

□

Equilibrium Portfolio Choice: Imposing Proposition 1 we have that the portfolio choices satisfy:

$$\begin{aligned}[\theta_h^k] : \quad & r^k - r^d = -\lambda \partial_{\theta_h^k} \mathcal{W}(\theta_h^k, \theta_h^n) \frac{\mathcal{U}'(\mathcal{C})}{a \Xi_h} - \partial_{\theta_h^k} \psi_{h,k} - \boldsymbol{\sigma}_{\Xi_h}^T \boldsymbol{\sigma}_{q^k} \\ [\theta_h^n] : \quad & r^n - r^d = -\lambda \partial_{\theta_h^n} \mathcal{W}(\theta_h^k, \theta_h^n) \frac{\mathcal{U}'(\mathcal{C})}{a \Xi_h} - \partial_{\theta_h^n} \psi_{h,n} - \boldsymbol{\sigma}_{\Xi_h}^T \boldsymbol{\sigma}_{q^n} \\ [\theta_b^k] : \quad & r^k - r^d = -\boldsymbol{\sigma}_{\Xi_b}^T \boldsymbol{\sigma}_{q^k} \\ [\theta_f^k] : \quad & r^k - r_f^n = -\boldsymbol{\sigma}_{\Xi_f}^T (\boldsymbol{\sigma}_{q^k} - \boldsymbol{\sigma}_{q^n}) \\ [\theta_f^m] : \quad & r^m - r_f^n = -\boldsymbol{\sigma}_{\Xi_f}^T (\boldsymbol{\sigma}_{q^m} - \boldsymbol{\sigma}_{q^n})\end{aligned}$$

where:

$$r_f^n = r_h^n + \left(\frac{1}{q_t^n} - 1 \right) \lambda$$

Equilibrium Euler Equations: Imposing Proposition 1 we have that the Euler equations satisfy:

$$\begin{aligned} \rho + \lambda &= \mu_{\Xi_h} + r^d - \tau_h + \psi_{h,k}(\theta_h^k) - \partial_{\theta_h^k} \psi_{h,k}(\theta_h^k) \theta_h^k + \psi_{h,n}(\theta_h^n) - \partial_{\theta_h^n} \psi_{h,k}(\theta_h^n) \theta_h^n \\ \rho &= \mu_{\Xi_b} + r^d - \tau_b \\ \rho &= \mu_{\Xi_f} + r_f^n - \tau_h \end{aligned}$$

A.1.4 Equilibrium Block 1: Summary of Optimization

Given equilibrium prices and price processes:

$$(r^d, q^k, r^k, \boldsymbol{\sigma}_{q^k}, q^n, r^n, \boldsymbol{\sigma}_{q^n}, q^m, r^m, \boldsymbol{\sigma}_{q^m})$$

the household, banker, and fund optimization variables (13 variables):

$$(\Xi_h, \Xi_b, \Xi_f, c_h, C_h, c_b, c_f, \theta_h^k, \theta_h^n, \theta_b^k, \theta_f^k, \theta_f^m, \iota)$$

satisfy the optimization equations (13 equations):

$$\begin{aligned}
0 &= -(\rho + \lambda) + \mu_{\Xi_h} + r^d - \tau_h + \psi_{h,k}(\theta_h^k) - \partial_{\theta_h^k} \psi_{h,k}(\theta_h^k) \theta_h^k \\
&\quad + \psi_{h,n}(\theta_h^n) - \partial_{\theta_h^n} \psi_{h,n}(\theta_h^n) \theta_h^n \\
0 &= -\rho_b + \mu_{\Xi_b} + r^d - \tau_b \\
0 &= -\rho_f + \mu_{\Xi_f} + r_f^n - \tau_h \\
0 &= u'(c_h) - \Xi_h \\
0 &= u'(c_b) - \Xi_b \\
0 &= u'(c_f) - \Xi_f \\
0 &= -\mathcal{C}_h + a(\theta_h^k(1 - \tau))^\alpha ((q^n)^{-1} \theta_h^n)^{1-\alpha} \\
0 &= r^k - r^d + \lambda \partial_{\theta_h^k} \mathcal{W}(\theta_h^k, \theta_h^n) \frac{\mathcal{U}'(\mathcal{C})}{a \Xi_h} + \partial_{\theta_h^k} \psi_{h,k} + \boldsymbol{\sigma}_{\Xi_h}^T \boldsymbol{\sigma}_{q^k} \\
0 &= r^n - r^d + \lambda \partial_{\theta_h^n} \mathcal{W}(\theta_h^k, \theta_h^n) \frac{\mathcal{U}'(\mathcal{C})}{a \Xi_h} + \partial_{\theta_h^n} \psi_{h,n} + \boldsymbol{\sigma}_{\Xi_h}^T \boldsymbol{\sigma}_{q^n} \\
0 &= r_t^k - r^d + \boldsymbol{\sigma}_{\Xi_b}^T \boldsymbol{\sigma}_{q^k} \\
0 &= r^k - r_f^n + \boldsymbol{\sigma}_{\Xi_f}^T (\boldsymbol{\sigma}_{q^k} - \boldsymbol{\sigma}_{q^n}) \\
0 &= r^m - r_f^n + \boldsymbol{\sigma}_{\Xi_f}^T (\boldsymbol{\sigma}_{q^m} - \boldsymbol{\sigma}_{q^n}) \\
0 &= \Phi'(l) - \frac{1}{q^k}
\end{aligned}$$

A.1.5 Equilibrium Block 2: Distribution Evolution

Kolmogorov Forward Equation (KFE): Financial Sector: We consider two levels of the distribution evolution. Let $A_{h,t}$, $A_{b,t}$, and $A_{f,t}$ denote the aggregate wealth in the household, banking, and fund sectors. Let $g_{j,t}$ denote the measure function of wealth for type $j \in \{h, b, f\}$. We start with the evolution of aggregate wealth in the banking

sector:

$$\begin{aligned}
\frac{dA_{b,t}}{A_{b,t}} &= \frac{1}{A_{b,t}} \int_0^\infty \mu_{a_b}(c_h(a), \boldsymbol{\theta}_b, \mathbf{S}) a g_{b,t}(a) da dt + \frac{1}{A_{b,t}} \lambda_b (\phi_b A_t - A_{b,t}) dt \\
&\quad + \frac{1}{A_{b,t}} \int_0^\infty \boldsymbol{\sigma}_{b,t}^T(\boldsymbol{\theta}_b, \mathbf{S}) a d\mathbf{W}_t a g_{b,t}(a) da \\
&= \left(\mu_{a_b}(c_h(a), \boldsymbol{\theta}_b, \mathbf{S}) + \lambda_b \left(\frac{\phi_b}{\eta_{b,t}} - 1 \right) \right) dt + \boldsymbol{\sigma}_{b,t}^T(\boldsymbol{\theta}_b, \mathbf{S}) d\mathbf{W}_t \\
&= \left(r^d + \theta_h^k (r^k - r^d) - (\rho_b + \lambda_b) - \tau_b + \lambda_b \left(\frac{\phi_b}{\eta_{b,t}} - 1 \right) \right) dt \\
&\quad + \boldsymbol{\sigma}_{b,t}^T(\boldsymbol{\theta}_b, \mathbf{S}) d\mathbf{W}_t
\end{aligned}$$

Likewise, the evolution of aggregate wealth in the fund section is:

$$\begin{aligned}
\frac{dA_{f,t}}{A_{f,t}} &= \left(r^b + \theta_f^k (r^k - r_f^b) + \theta_f^m (r^m - r_f^m) - (\rho_f + \lambda_f) - \tau_f + \lambda_f \left(\frac{\phi_f}{\eta_{f,t}} - 1 \right) \right) dt \\
&\quad + \boldsymbol{\sigma}_{f,t}^T(\boldsymbol{\theta}_f, \mathbf{S}) d\mathbf{W}_t
\end{aligned}$$

Aggregate wealth is given by $A_t = q_t^k K_t + q_t^m M$. Let $\vartheta_t = q_t^k K_t / (q_t^k K_t + q_t^m M)$. The evolution of aggregate wealth follows:

$$\begin{aligned}
\frac{dA_t}{A_t} &= \vartheta \left(\frac{dq^k}{q^k} + \frac{dK}{K} \right) + (1 - \vartheta) \frac{dq^m}{q^m} \\
&= \underbrace{\vartheta (\mu_{q^k} + \Phi(\iota) - \delta) + (1 - \vartheta) \mu_{q^m}}_{\mu_A} + \vartheta (\sigma_{q^k, z} dW_z + \sigma_{q^k, \zeta} dW_\zeta) \\
&\quad + (1 - \vartheta) (\sigma_{q^m, z} dW_z + \sigma_{q^m, \zeta} dW_\zeta)
\end{aligned}$$

So, the evolution of $\eta_{b,t} = A_{b,t}/A_t$ is given by:

$$\begin{aligned}
\frac{d\eta_{b,t}}{\eta_{b,t}} &= \frac{dA_{b,t}}{A_{b,t}} - \frac{dA_t}{A_t} - \frac{dA_{b,t}}{A_{b,t}} \frac{dA_t}{A_t} + \left(\frac{dA_t}{A_t} \right)^2 \\
&= \left(\mu_{A_{b,t}} - \mu_{A,t} - \boldsymbol{\sigma}_{A_{b,t}}^T \boldsymbol{\sigma}_{A,t} + \boldsymbol{\sigma}_{A,t}^T \boldsymbol{\sigma}_{A_{b,t}} \right) dt + (\boldsymbol{\sigma}_{A_{b,t}} - \boldsymbol{\sigma}_{A,t})^T d\mathbf{W}_t \\
&= \left(\mu_{A_{b,t}} - \mu_{A,t} + (\boldsymbol{\sigma}_{A,t} - \boldsymbol{\sigma}_{A_{b,t}})^T \boldsymbol{\sigma}_{A,t} \right) dt + (\boldsymbol{\sigma}_{A_{b,t}} - \boldsymbol{\sigma}_{A,t})^T d\mathbf{W}_t
\end{aligned}$$

and the evolution of $\eta_{f,t} = A_{f,t}/A_t$ is given by:

$$\frac{d\eta_{f,t}}{\eta_{f,t}} = \left(\mu_{A_{f,t}} - \mu_{A,t} + (\boldsymbol{\sigma}_{A,t} - \boldsymbol{\sigma}_{A_{f,t}})^T \boldsymbol{\sigma}_{A,t} \right) dt + (\boldsymbol{\sigma}_{A_{f,t}} - \boldsymbol{\sigma}_{A,t})^T d\mathbf{W}_t$$

KFE: Within Households: The KFE for the household distribution in levels, a , is given by:

$$\begin{aligned} dg_{h,t}(a) &= + \lambda_h \phi(a) A_t - \lambda_h g_{h,t}(a) - \partial_a [\mu_a(a, \mathbf{s}_t, g_{h,t}) g_{h,t}(a)] \\ &\quad - \partial_a [\boldsymbol{\sigma}_a^T(a, \mathbf{s}_t, g_{h,t}) d\mathbf{W}_t g_{h,t}(a)] + \frac{1}{2} \partial_a \left[\boldsymbol{\sigma}_a^T \boldsymbol{\sigma}_a(a, \mathbf{s}_t, g_{h,t}) g_{h,t}(a) \right] dt \\ &= + \lambda_h \phi(a) A_t - \lambda_h g_{h,t}(a) - \partial_a [\mu_a(a, \mathbf{s}_t, g_{h,t}) g_{h,t}(a)] \\ &\quad - \partial_a [\sigma_{a,z}(a, \mathbf{s}_t, g_{h,t}) g_{h,t}(a)] dW_{z,t} - \partial_a [\sigma_{a,\zeta}(a, \mathbf{s}_t, g_{h,t}) g_{h,t}(a)] dW_{\zeta,t} \\ &\quad + \frac{1}{2} \partial_a \left[(\sigma_{a,z}^2(a, \mathbf{s}_t, g_{h,t}) + \sigma_{a,\zeta}^2(a, \mathbf{s}_t, g_{h,t})) g_{h,t}(a) \right] dt \end{aligned}$$

The evolution $\eta_{i,t} := a_{i,t}/A_t$ is given by:

$$\begin{aligned} \frac{d\eta_{i,t}}{\eta_{i,t}} &= \left(\mu_{a_{i,t}} - \mu_{A,t} + (\boldsymbol{\sigma}_{A,t} - \boldsymbol{\sigma}_{a_{i,t}})^T \boldsymbol{\sigma}_{A,t} \right) dt + (\boldsymbol{\sigma}_{a_{i,t}} - \boldsymbol{\sigma}_{A,t})^T d\mathbf{W}_t \\ &=: \mu_{\eta_{i,t}} dt + \boldsymbol{\sigma}_{\eta_{i,t}}^T d\mathbf{W}_t \end{aligned}$$

where we have softened the entry function from ϕ_h to $\phi(a)$, where $\phi(a)$ is a function with mean ϕ_h . For a , a natural candidate would be $\phi(a) = \text{LogNormal}(\phi_h, \sigma)$. Likewise, the KFE for the distribution in shares is:

$$\begin{aligned} dg_{h,t}(\eta) &= + \lambda_h \phi(\eta) - \lambda_h g_{h,t}(\eta) - \partial_\eta [\mu_\eta(\eta, \mathbf{s}_t, g_{h,t}) g_{h,t}(\eta)] \\ &\quad - \partial_\eta [\boldsymbol{\sigma}_\eta^T(\eta, \mathbf{s}_t, g_{h,t}) d\mathbf{W}_t g_{h,t}(\eta)] + \frac{1}{2} \partial_\eta \left[\boldsymbol{\sigma}_\eta^T \boldsymbol{\sigma}_\eta(\eta, \mathbf{s}_t, g_{h,t}) g_{h,t}(\eta) \right] dt \\ &= + \lambda_h \phi(\eta) - \lambda_h g_{h,t}(\eta) - \partial_\eta [\mu_\eta(\eta, \mathbf{s}_t, g_{h,t}) g_{h,t}(\eta)] \\ &\quad - \partial_\eta [\sigma_{\eta,z}(\eta, \mathbf{s}_t, g_{h,t}) g_{h,t}(\eta)] dW_{z,t} - \partial_\eta [\sigma_{\eta,\zeta}(\eta, \mathbf{s}_t, g_{h,t}) g_{h,t}(\eta)] dW_{\zeta,t} \\ &\quad + \frac{1}{2} \partial_\eta \left[(\sigma_{\eta,z}^2(\eta, \mathbf{s}_t, g_{h,t}) + \sigma_{\eta,\zeta}^2(\eta, \mathbf{s}_t, g_{h,t})) g_{h,t}(\eta) \right] dt \end{aligned}$$

where again we have softened the entry function ϕ_h to $\phi(\eta)$, where $\phi(\eta)$ is a function with mean ϕ_h . For η , a natural candidate would be $\phi(\eta) \sim \text{Beta}$ with mean ϕ_h .

B Additional Details on Simulating

In order to simulate the economy we need to compute the evolution of the household wealth distribution. This is complicated for the finite agent approximation method because the neural network policy rules are functions of the positions of the N other agents rather than a continuous density. To overcome this difficulty, we deploy the “hybrid” approach described in Algorithm 2 that uses the neural network solution to approximate a finite difference approximation to the KFE. Let $\mathbf{a} = (a_m : m \leq M)$ denote the grid in the a -dimension. Let $\mathbf{g}_t = (g_{m,t} : m \leq M)$ denote the marginal density on the a -grid. At each time step, our method draws N_{sim} different samples of N agents from the current density g_t . For each draw $k \leq N_{sim}$, denoted by $\hat{\varphi}_t^k = (a_i : 1 \leq i \leq N)$, the KFE is replaced by the following finite difference equation:

$$dg_{m,t} = \mu_{g,m}(\hat{\varphi}_t^k)dt + \boldsymbol{\sigma}_{g,m}^T(\hat{\varphi}_t^k)d\mathbf{W}_t, \quad m \leq M \quad (\text{B.1})$$

where the drift at point (m) is defined by the finite difference approximation for the KFE using the policy rules from our finite population neural network solution. From this approximation we can calculate the transition matrix $\mathcal{A}_{t,k}$ for the finite difference approximation at the draw φ^k . We repeat this procedure many times then compute an average transition matrix, which we use for simulation. We summarize the steps in Algorithm 2.

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 $y_t = e^{z_t}$.

Algorithm 2: Finding Transition Paths In Finite Agent Approximation

Input : Initial distribution, neural network approximations to the policy and price functions, number of agents N , time step size Δt , number of time steps N_T , number of simulations N_{sim} , grid $\underline{a} = \{a_m : m \leq M\}$ for the finite difference approximation.

Output: A transition path $g = \{g_t : t = 0, \Delta t, \dots, N_T \Delta t\}$

```

for  $n = 0, \dots, N_T - 1$  do
  for  $k = 1, \dots, N_{sim}$  do
    Sample  $\Delta W_{t,z}$  from the normal distribution  $N(0, \Delta t)$ , construct TFP shock paths by:  $z_{t+\Delta t} = z_t + \eta(\bar{z} - z_t) + \sigma \Delta B_t^0$ . Do likewise to construct the volatility shock path.
    Draw states for  $N$  agents  $\{\varphi_i^k : i = 1, \dots, N\}$  from density  $g_t$  at  $t = n\Delta t$ .
    Given state  $(z_{t+\Delta t}, \varphi_t^k)$ , compute equilibrium prices and returns.
    At each grid point  $a_m \in \underline{a}$ , calculate the consumption and portfolio choices.
    Construct the transition matrix  $\mathcal{A}_{t,k}$  using finite difference on the grid  $\underline{a}$ , as described by (B.1).
  end
  Take the average:  $\bar{\mathcal{A}}_t = \frac{1}{N_{sim}} \sum_{k=1}^{N_{sim}} \mathcal{A}_{t,k}$ 
  Update  $g_t$  by implicit method:  $g_{t+\Delta t} = (I - \bar{\mathcal{A}}_t^\top \Delta t)^{-1} g_t + \sigma^T dW_{t,z}$ 
end

```

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:

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

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

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

Therefore, asset prices are given by:

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

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 (??) 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).

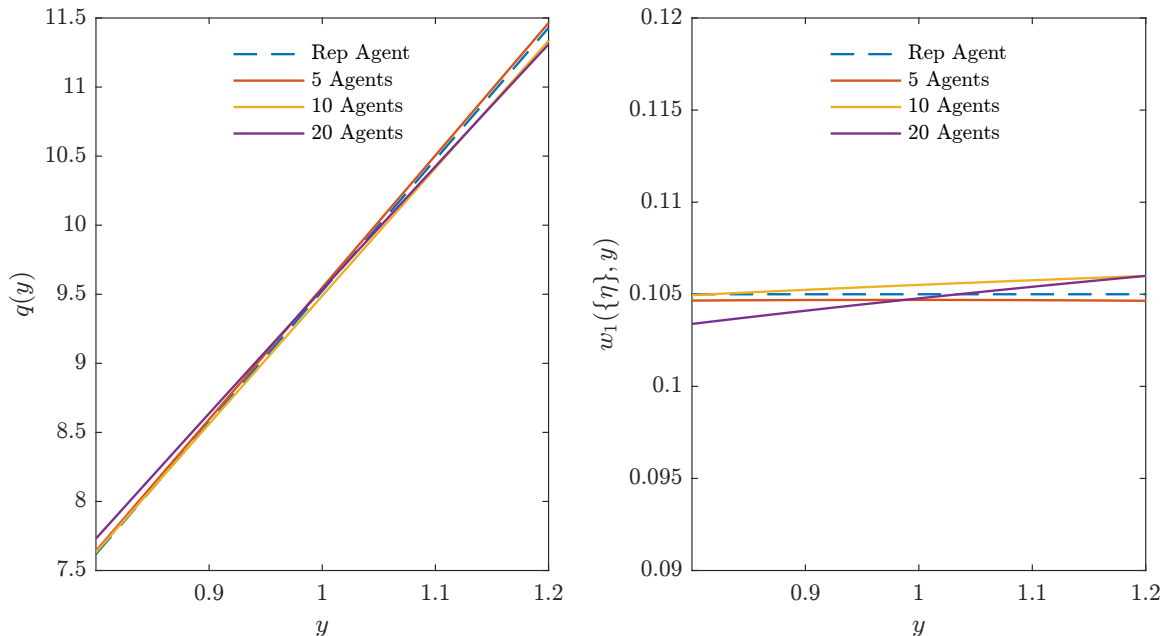


Figure 12: Solution to As-if representative agent model. Right panel: consumption-wealth ratio of agent 1.

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).

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:

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

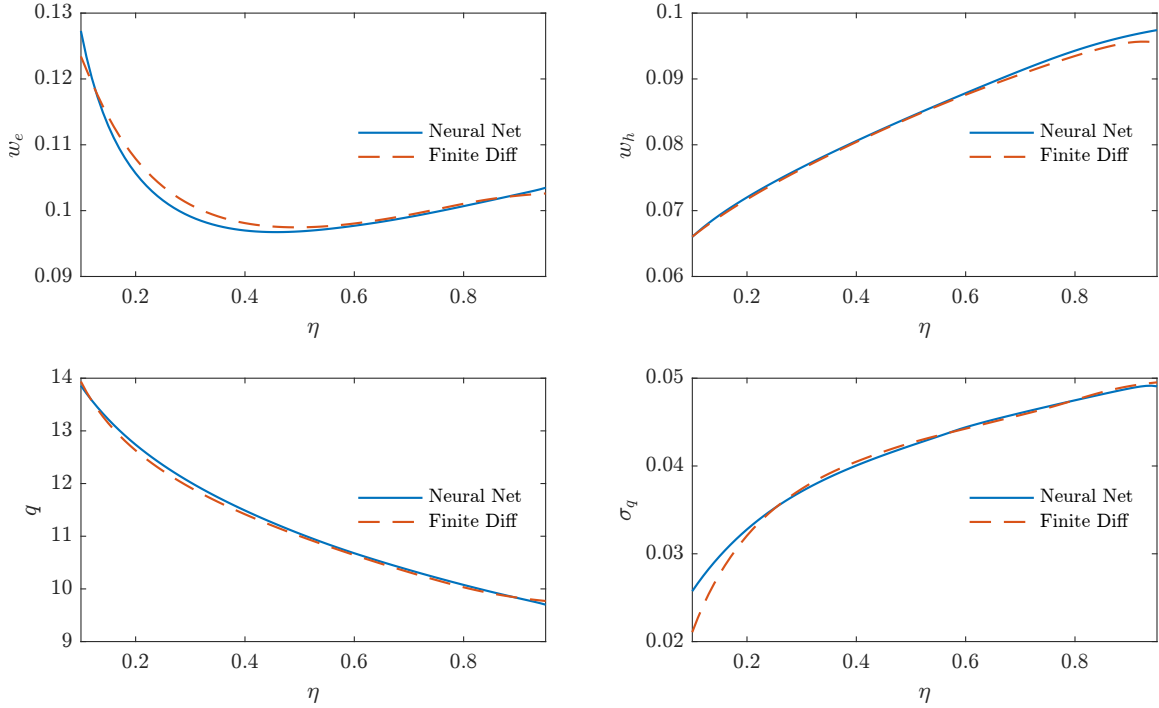


Figure 13: Solution to restricted stock market participation model.

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:

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

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 $\mu_y = y\mu$, volatility $y\sigma$, 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:

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

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:

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

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}_λ is a trainable neural net and β is solved from the asymptotic solution for $\eta \rightarrow 0$. 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 $\eta \rightarrow 0$, 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.

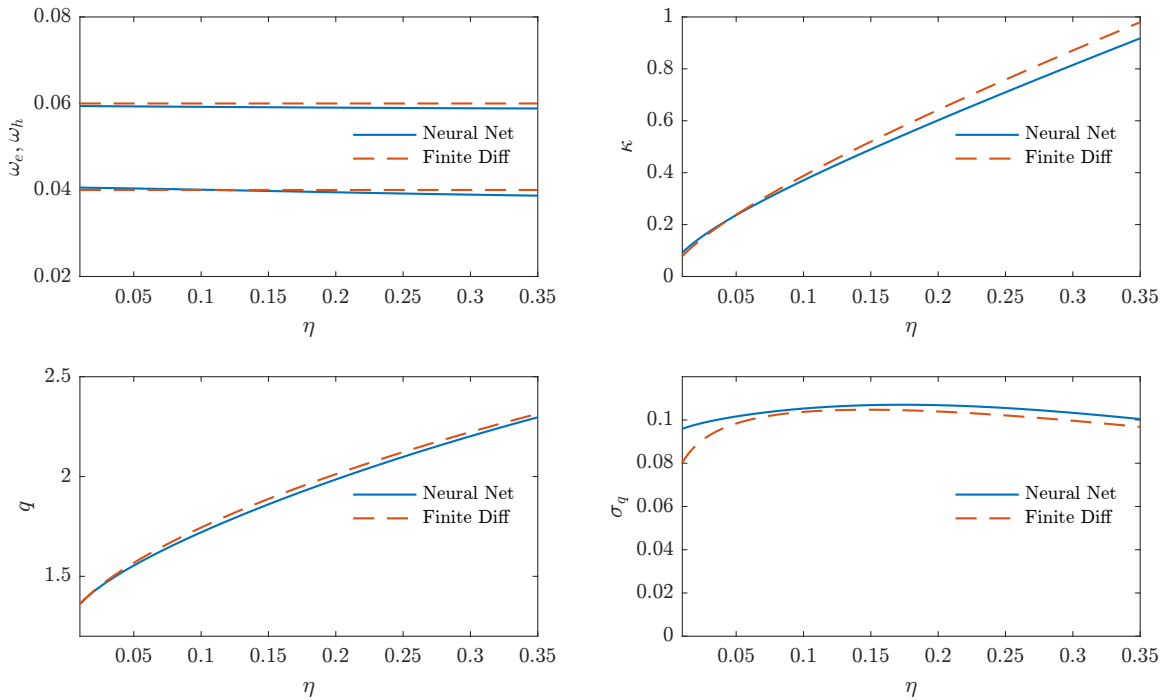


Figure 14: Solution to the model with productivity gap.

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 condition⁵:

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

From the goods market clearing conditon, we have:

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

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

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

where μ_η, σ_η are from (??) and (??). 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):

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

⁵This expression leads to the boundary condition at $\eta = 1$: $\frac{f(1)}{v_e} = 1$

and update value function in an implicit scheme to solve equation

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

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

D.1 Solution to the Limited Participation Model

The distributional dynamics for limited participation model are:

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

By the consistency condition for price volatility, we have:

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

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: $\eta_n = \text{linspace}(\Delta\eta, 1 - \Delta\eta, 1/\Delta\eta - 1)$. 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:

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

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

Algorithm. Set up grids: $\eta_n = \text{linspace}(\Delta\eta, 1 - \Delta\eta, 1/\Delta\eta - 1)$. 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 $\kappa < 1$

$$\left\{ \begin{array}{l} \rho_e \omega_e \eta + \rho_h \omega_h (1 - \eta) = \kappa z_e + (1 - \kappa) z_h \\ \sigma_q = \frac{\sigma}{1 - \frac{f'(\eta)}{f(\eta)}(\kappa - \eta)} \\ \frac{z_e - z_h}{q} = \frac{\kappa - \eta}{\eta(1 - \eta)} \sigma_q^2. \end{array} \right. \quad (\text{D.4})$$

(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 Parameters for Testable Models

E.1 Economic Parameters

E.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%

E.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%

E.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%

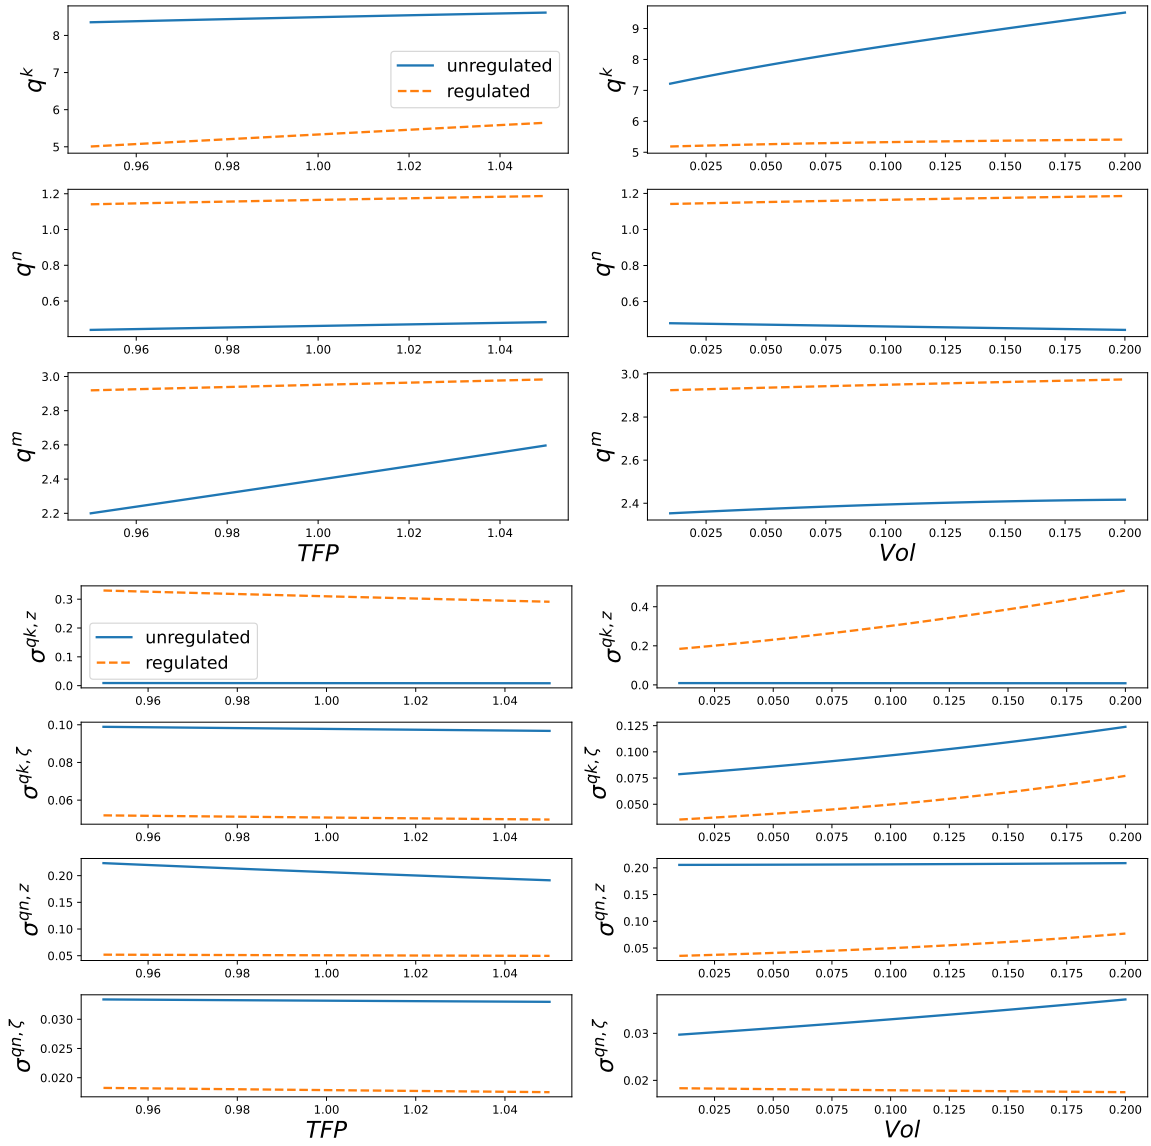


Figure 15: The figure plots the model-implied spreads for the unregulated (regulated) economy in solid (dashed) line. The left (right) panel presents holdings with respect to TFP (volatility). The fund is restricted from participating in the capital market in the regulated economy.

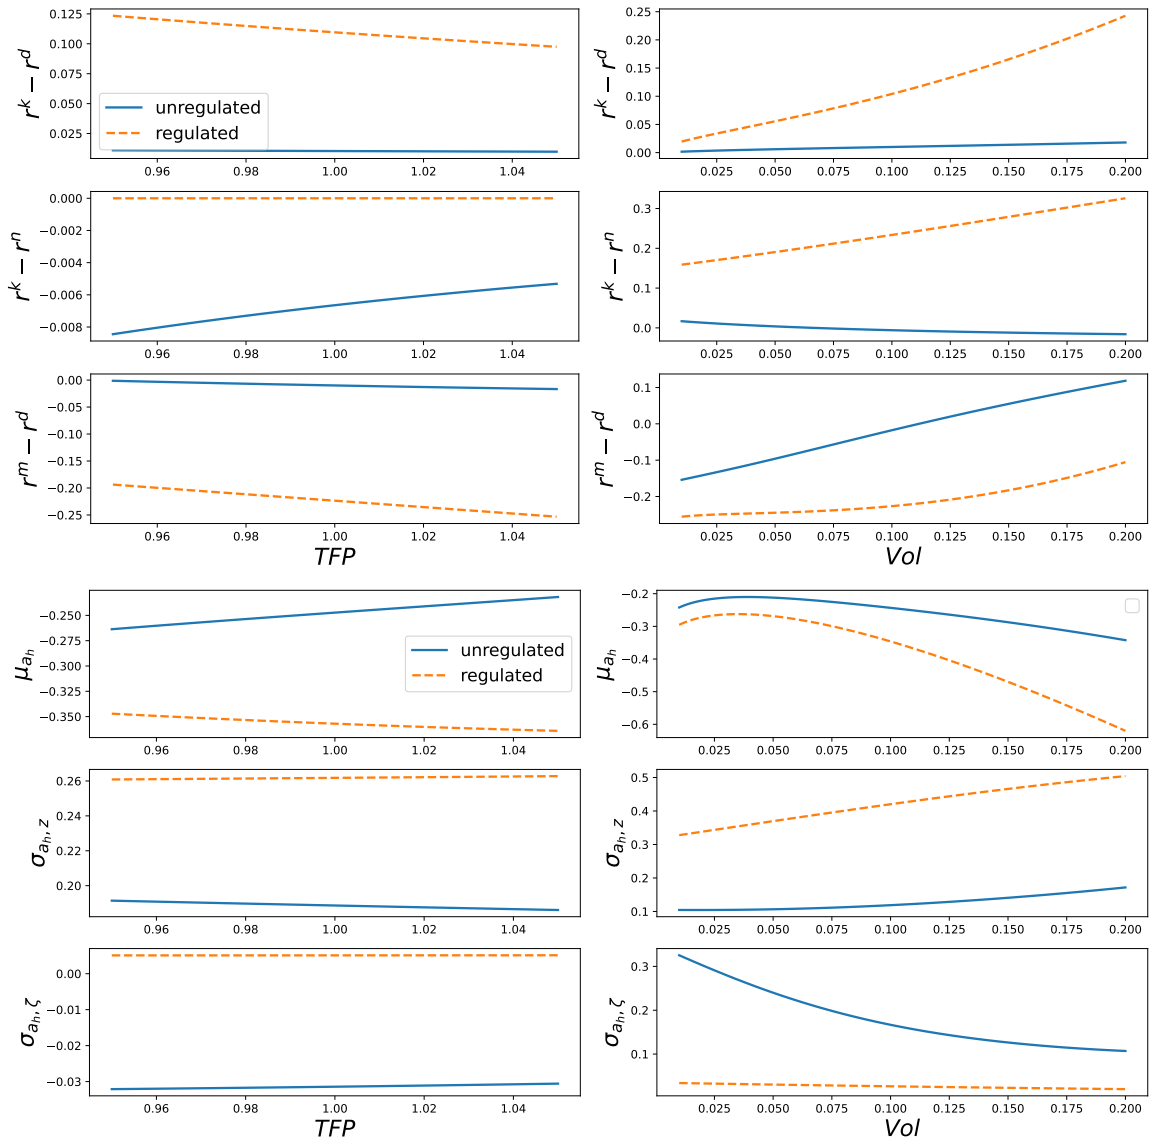


Figure 16: The figure plots the model-implied price volatilities for the unregulated (regulated) economy in solid (dashed) line. The left (right) panel presents holdings with respect to TFP (volatility). The fund is restricted from participating in the capital market in the regulated economy.

F Additional Plots for Section 4.1

G Additional Plots and Working for Section 4.2

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

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

For $\bar{\psi} \in (0, \infty)$, the agent portfolio choices $\{\theta_{i,t}\}_{i \leq I}$ 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 is held by an agent facing a small participation constraint. Theorem 2 shows that as $\eta_{e,t} \rightarrow 0$, the $\sigma_{q,t}$ becomes constant and intuition above is precisely true. Figure 17 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 $\eta_e \rightarrow 0$, the $\sigma_q \rightarrow \sigma$ and greater household inequality leads to a lower excess return on capital.*

Proof. Available in the appendix ??.

□

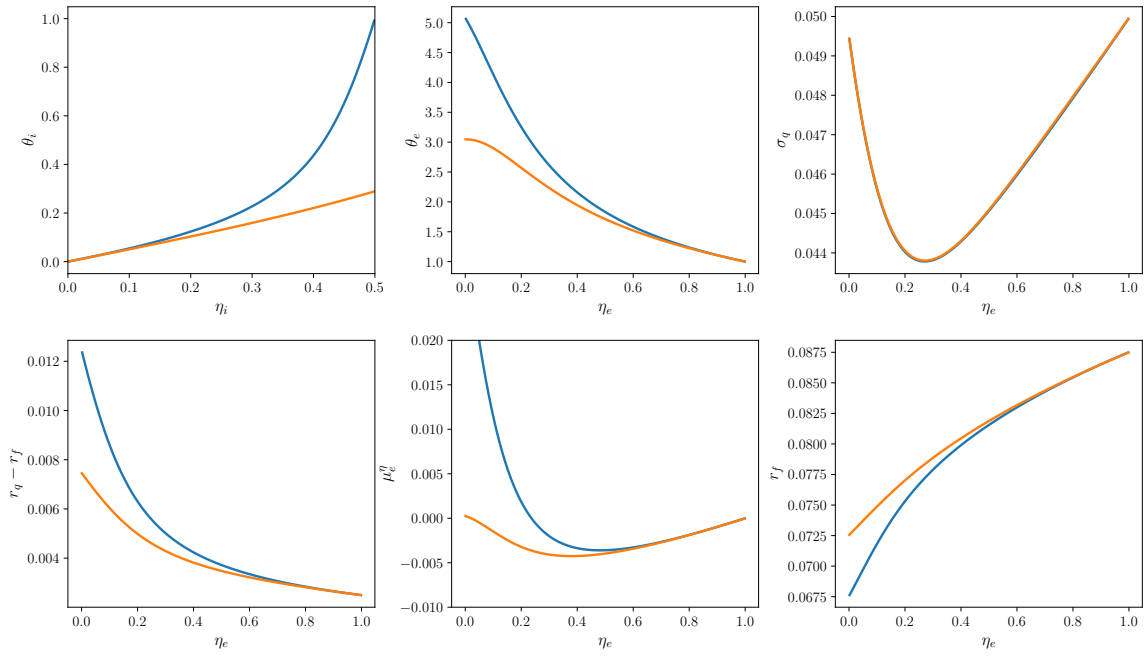


Figure 17: Equilibrium functions in a two households economy with log utility and varying inequality. The orange line is when one household has all the wealth. The blue line is when each household takes half of the wealth in household sector. Participation constraint is $\bar{\psi} = 2.0$.