

Institutional Asset Pricing with Segmentation and Household Heterogeneity*

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

March 12, 2026

Abstract

How are different households affected by the structure and regulation of the financial sector? To answer this question, we build a heterogeneous-agent macro-finance model with households facing asset market participation constraints, banks providing deposits, funds providing insurance/pension products, and endogenous asset price volatility. We develop a novel deep learning approach to globally solve this and other macro-finance models. Counterfactual experiments reveal policy trade-offs: tighter financial sector restrictions increase stability but lead to lower growth and/or higher inequality because richer households are better able to take advantage of the higher spreads created by the regulations.

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

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

*This paper was previously circulated as “Asset Pricing, Participation Constraints, and Inequality”. We are very grateful to the comments and discussion from Jaroslav Borovička, Anmol Bhandari, Saki Bigio, Markus Brunnermeier, Riccardo Cioffi (Discussant), Paul Ehling (Discussant), Lars Hansen, Felix Kübler, Moritz Lenel, Galo Nuño, Michael Reiter, Richard Rogerson, Simon Scheidegger, Martin Schneider, Yucheng Yang, Moto Yogo, Jan Žemlička, and Adam Zhang (Discussant). We are also thankful for the comments from participants in the Bank of Spain 8th Annual Research Conference “Economics of Artificial Intelligence”, the Blue Collar Working Group at the University of Chicago, the CESifo Conference, the EFA, the SED, the NFA, the Sargent Alumni Reading Group, and the Zurich Workshop on the Frontier of Quantitative Macroeconomics: Methods and Models. We thank Rotman FinHub lab for the computational resources.

[†]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

A large literature has established that financial institutions face frictions and regulations that constrain lending, segment markets, and influence asset pricing. What is less understood is how these features of the financial sector affect and are affected by different households. Are poorer households more exposed to financial intermediary risk because they are more dependent on banking, pension, and/or insurance products? Can wealthier households take advantage of the higher asset pricing spreads from financial regulations? How does household demand for different financial services impact the funding costs of different financial intermediaries? Does household inequality affect risk sharing and investment? This paper studies these macro-finance connections and shows that policy makers face novel and subtle tradeoffs between managing financial stability, economic growth, and household inequality.

We make three main contributions. First, we develop a novel quantitative heterogeneous-agent macro-finance (HAMF) model that embeds different types of financial intermediaries (namely banks and pension/insurance funds) and endogenous volatility into a continuous-time heterogeneous agent real business cycle environment. This allows us to bridge the different intermediary asset pricing and macroeconomics literatures. Second, we develop a deep-learning approach to characterize global solutions to our model. This overcomes the major technical challenge for heterogeneous-agent macro-finance that standard numerical techniques cannot handle environments like ours with aggregate risk, constrained portfolio choice, and a non-degenerate distribution of household wealth. Finally, we calibrate the model to the post-2010 period, examine how frictions in the household sector interact with asset pricing and inequality, and conduct counterfactual experiments with alternative regulation on banks and pension/insurance funds.

The details of our model are outlined in Section 2. The economy is populated by overlapping generations of price-taking households who retire at random times, which we model as exit from the productive economy. Households face two frictions: they incur costs when they participate directly in capital markets and they cannot contract among each other to insure retirement shocks. There are two types of financiers who provide services to help households overcome these frictions: bankers and fund managers. Bankers issue risk-free short-term deposits while the fund managers issue contracts that pay out when agents retire, which we refer to as “pensions” and interpret as a combination of pension and life insurance services. Both types of financiers can also issue long-term bonds and face equity raising constraints that prevent them from quickly recapitalizing after negative wealth shocks. Finally, there is a government that issues long-term bonds, raises wealth

taxes, and potentially places regulatory portfolio restrictions on the financial intermediaries.

Our environment is constructed to nest a collection of key macro-finance models. If we restrict household preferences so that households only consume at retirement, then we recover the inelastic “preferred habitat” asset demand functions for long-maturity assets in [Vayanos and Vila \(2021\)](#). Alternatively, if we eliminate retirement and household participation frictions, then we recover the [Brunnermeier and Sannikov \(2014\)](#) model in which bankers borrow from households and all types of agents participate together in the capital market. If we take the household participation friction to infinity, then we recover a segmented market model similar to [Gertler and Kiyotaki \(2010\)](#) where only bankers have access to direct investment in productive capital. The flexibility (and complexity) in our household choice problem allows us to incorporate these different models into one framework. It also allows us to link frictions in the household sector directly to financial intermediary portfolio choices and constraints.

A key feature of our model is that households make portfolio choices by balancing the standard frictionless risk-return tradeoff with the additional distortions from the frictions that penalize capital market participation and generate demand for pension products. One implication is that richer households choose to hold a greater fraction of their wealth in capital while poorer households choose relatively more deposits and all agents allocate a similar fraction of their wealth to pensions. Ultimately, this household portfolio heterogeneity, combined with the incompletely insurable retirement shocks, generates a non-degenerate household wealth distribution. Another implication is that household pension demand becomes highly return inelastic once pension holdings become large whereas household deposit demand does not. This means that pension demand is less sensitive to returns and household wealth changes than deposit demand.

We calibrate our baseline model to match asset pricing spreads and agent portfolios in the post-financial crisis period (2010Q1 to 2024Q4). We draw a collection of lessons about the quantitative impact of the household portfolio choice frictions. First, wealthier households earn a higher return on their wealth than poorer households because they hold more of their wealth in capital and so earn higher premiums. This means that shocks or policy changes that end up increasing Sharpe ratios have subtle and ambiguous effects on inequality. To understand this, consider how a higher Sharpe ratio impacts different households, holding all else constant. On the one hand, it disproportionately benefits the richer households who already hold a large fraction of their wealth directly in capital stock (an “intensive margin” effect). That is, richer households have better access to asset markets and so are better able to take advantage of higher returns. On the other hand, a higher risk adjusted capital return encourages poorer households into the capital market and so reduces the portfolio

differences across the wealth distribution. This is because, for poorer households, the Sharpe ratio offsets the participation cost and so leads them to participate more in the capital market (an “extensive” margin effect). The result is that, in partial equilibrium, a small increase in the Sharpe ratio increases inequality because the intensive margin effect dominates while a sufficiently large increase in the Sharpe ratio decreases inequality because the extensive margin effect dominates. The lower the household participation cost, the more the extensive margin effect dominates. In general equilibrium, the entry of poorer households into the capital market then affects aggregate risk sharing and so the Sharpe ratio.

Second, a higher capital market participation penalty leads to lower social mobility. In our model, social mobility is governed by how quickly a generation of agents is able to accumulate wealth. This, in turn, is governed by how severely the household portfolio constraints bind for poorer agent and the size of the endogenous risk premium. We show that an increase in the participation cost leads to a decrease in social mobility. An important advantage of solving a model with a full wealth distribution is that we can study social mobility because the model captures how households endogenously transition between being rich and poor. By contrast, if we had a lower dimensional model with two types of households (e.g. capital market participants and non-capital market participants), then households would be permanently rich or poor (or potentially transition between rich and poor at an exogenous rate).

Third, relative to bankers, fund managers are both less able to expand their balance sheet and more able to absorb risk. The former occurs because household demand for pensions becomes relatively unresponsive to returns once pension holdings are high and so it is very costly for fund managers to scale up their balance sheet. By contrast, household demand for deposits grows approximately linearly with deposit return spreads and so the bankers can increase their balance sheet at lower cost. With regard to the ability to absorb risk, fund liabilities (pensions) fall in value during recessions, while banks liabilities (deposits) value remain unaffected. As a result, funds have a natural hedge against business cycle risk and help to “insure” the banking sector by buying their long-term bonds.

A final implication of the household portfolio frictions is that general equilibrium investment and output growth are connected to inequality. If the economy dis-intermediates and the households start to hold more capital directly, then capital prices and investment fall leading to lower growth because households face higher costs of holding capital and have a higher risk aversion. The effect is less pronounced if the household distribution is more unequal because then more wealth is held by the households who are least constrained.

We close our paper by using our calibrated model to revisit proposed changes to macroprudential regulations: (i) tightening bank leverage requirements to capture suggestions that current restrictions remain too lax (e.g. [Admati and Hellwig \(2013\)](#)) and (ii) imposing bank-style leverage restrictions on funds to respond to observations that the pension-insurance sector has taken on additional risk under Basel III (e.g. [Kojen and Yogo \(2022\)](#), [Begenau, Liang and Siriwardane \(2024\)](#)). Restricting the banking sector leads to lower volatility but also decreases equality and growth. By contrast, restricting the fund sector leads to lower volatility and higher growth but also much higher inequality. In both case, the restrictions end up increasing the Sharpe ratio. The difference between the two counterfactual experiments appears because the bankers and fund managers have different ability to expand their balance sheets and share risk. When the fund sector is restricted, the bankers expands their capital holdings and household capital market participation stays similar. This means that investment and growth remain high while inequality increases because households stay out of the capital market (i.e. the intensive margin affect dominates). By contrast, when the banking sector is restricted, the fund managers are less able to expand and so households end up holding more capital. This leads to a decrease in investment and growth. Inequality also increases, although not by as much as when funds are restricted because the extensive margin effect is stronger and more households enter the capital market. This suggests a trilemma style policy tradeoff that macroprudential regulation cannot simultaneously decrease volatility, increase growth and decrease inequality.

From a technical point of view, our model is unsolvable using traditional techniques and so we develop a new deep learning approach for characterizing solutions to this and other heterogeneous-agent macro-finance models. Mathematically, general equilibrium for our economy can be characterized by three blocks: (1) a collection of high dimensional partial differential equations (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 market clearing. The technical challenge is that, unlike for existing macro-finance models, we need a solution approach that can handle complexity in all three blocks. We do this by drawing upon and expanding recent advances in the continuous time deep learning macroeconomics literature (e.g. [Duarte, Duarte and Silva \(2024\)](#); [Gopalakrishna \(2021\)](#); [Fernández-Villaverde, Hurtado and Nuno \(2023\)](#); [Gu, Lauriere, Merkel and Payne \(2024\)](#)). This involves using neural networks to approximate derivatives of value functions, the price volatility of long-term assets, portfolio choices, and other equilibrium objects. We then train the neural network to minimize the error in the “master” equations that characterize the equilibrium blocks. The novel challenge is

that the combination of complicated portfolio choice and a non-degenerate household distribution leads to an endogenously stochastic Kolmogorov Forward Equation (KFE), sharp curvature in the policy functions, and market clearing conditions that are complicated to impose. To the best of our knowledge, this is the first paper to solve a model involving a stochastic KFE with endogenous volatility.

Literature Review: We contribute to several areas of research. The first is the macro-finance literature studying how financial institutions generate endogenous risk in dynamic general equilibrium models. Many of these papers focus on contracting frictions and incomplete risk sharing between an aggregate household sector and a representative financial intermediary (e.g. [Gertler and Kiyotaki \(2010\)](#); [Brunnermeier and Sannikov \(2016\)](#); [He and Krishnamurthy \(2013\)](#)). More recent work introduces heterogeneity across intermediary sector. For example, [Coimbra and Rey \(2024\)](#) and [Corbae and D’Erasmus \(2014\)](#) introduce heterogeneous banks, [Gertler, Kiyotaki and Prestipino \(2016\)](#) and [Kargar, \(2021\)](#) introduce wholesale banking and/or broker-dealers, and [Koijen and Yogo \(2023a,b\)](#) and [Coimbra, Gomes, Michaelides and Shen \(2023\)](#) introduce a pension/insurance sector. Our analysis of financial sector regulation connects to the part of the literature on the general equilibrium impact of macro-prudential regulation (e.g. [Van den Heuvel \(2008\)](#), [Begeau \(2016\)](#), [Begenau and Landvoigt \(2017\)](#), [Elenev, Landvoigt and Nieuwerburgh \(2021\)](#), [Gale and Yorulmazer \(2020\)](#)). Our contribution relative to these literatures is to combine intermediary heterogeneity with a non-degenerate household wealth distribution, thereby producing rich cross-sectional outcomes for households and sector-level outcomes between the intermediaries. This allows us to study the impact of macro-prudential policy on inequality and social mobility.

Secondly, we are part of the literature studying how asset pricing can impact the household wealth distribution (e.g. [Gomez \(2017\)](#), [Cioffi \(2021\)](#), [Basak and Chabakauri \(2024\)](#), [Cioffi, Nuño and Hurtado \(2023\)](#)). Many of these papers focus on uninsurable idiosyncratic income and return risk. We complement this work by studying heterogeneity across overlapping generations (similar to [Garleanu and Panageas \(2023\)](#)) in a general equilibrium model with aggregate risk, endogenous volatility, and a frictional financial sector. This allows us to match empirical evidence documenting the heterogeneity in household portfolio choices (e.g. [Bach, Calvet and Sodini \(2020\)](#)) in a general equilibrium model with endogenous asset returns.

Thirdly, we connect to the research on limited financial market participation. A long literature has studied frictions on household and investor asset portfolio choice (e.g. [Merton \(1987\)](#), [Gomes and Michaelides \(2007\)](#), [Guvenen \(2009\)](#), [Favilukis \(2013\)](#), [Khorrami \(2021\)](#)). Of particu-

lar relevance is the recent institutional asset pricing literature studying how regulatory constraints, inelastic demand, and financial market segmentation impact asset prices (e.g. [Chien, Cole and Lustig \(2012\)](#), [Greenwood and Vissing-Jorgensen \(2018\)](#), [Kojen and Yogo \(2019\)](#), [Vayanos and Vila \(2021\)](#)). Our contribution to this literature is to connect the household portfolio restrictions to the institutional financial sector portfolio choices.

Finally, we are part of the computational economics literature employing deep learning and other new tools to solve complex heterogeneous-agent models that challenge 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 \(2021\)](#); [Duarte et al. \(2024\)](#); [Gopalakrishna \(2021\)](#); [Gu et al. \(2024\)](#); [Kogan and Mitra \(2025\)](#)). While several papers apply deep learning to heterogeneous-agent models with portfolio choice between short-term risky assets (e.g. [Fernández-Villaverde et al. \(2023\)](#); [Huang \(2023\)](#)), very few tackle long-term asset pricing in general equilibrium because this has proven to be a difficult challenge for the deep learning techniques. One contemporaneous attempt to resolve the challenge is [Azinovic and Zemlicka \(2026\)](#), which solves a discrete-time model with long-lived assets using an iterative homotopy procedure. Our contribution is to provide a general, stable technique for solving continuous macro-finance models with heterogeneous intermediaries, multiple long-lived assets, and rich wealth distributions without resorting to restrictive distribution and portfolio choice approximations. We do this by exploiting the continuous time formulation of the problem to construct a loss function that helps the neural network avoid the difficulties with general equilibrium portfolio choice problems.

The rest of this paper is structured as follows. Section 2 outlines our economic model. Section 3 introduces our numerical algorithm. Section 4 describes the calibration of the baseline model. Section 5 discusses the key mechanisms. Section 6 presents counterfactual regulatory experiments before concluding.

2 Economic Model

In this section, we outline our economic model. We study a stochastic production economy with heterogeneous households who face retirement shocks and asset market participation constraints. The households are serviced by two types of financiers: bankers who issue deposits and fund managers who issue pensions to provide consumption at retirement.

2.1 Environment

Setting: The model is in continuous time with an infinite horizon. There is a perishable consumption good and a durable capital stock. The economy is populated by unit continua of price-taking households (h), bankers (b), and fund managers (f), the later of which that we interpret as the combining pension and insurance services. The banking and fund sectors will each aggregate to pseudo representative agents but the household sector will not. The economy has the following assets: short-term deposits, pension contracts, capital stock, and long-term bonds.

Production: All agents have access to a production technology that creates consumption goods according to the linear production function $y_t = e^{z_t} k_t$ where k_t and z_t are capital deployed and aggregate productivity at time t respectively. We impose that z_t evolves according to a mean reverting process $dz_t = \beta_z(\bar{z} - z_t)dt + \sigma_z dW_t$ where W_t denotes the aggregate Brownian motion shock process for the economy. All agents can create capital stock using an investment technology that converts $\iota_t k_t$ goods into $\phi(\iota_t) k_t$ capital, where ι_t is referred to as the investment rate. Capital depreciates at rate $\delta > 0$. So, an agent's physical capital stock evolve according to $dk_t = (\phi(\iota_t) k_t - \delta k_t)dt$.

Households: Households are born as “young” agents and then transition to “retired” agents at rate λ_h . While young, each household has discount rate ρ and gets flow utility $u(c_{h,t}) = \beta c_{h,t}^{1-\gamma} / (1-\gamma)$ from flow consumption $c_{h,t}$. When a household retires, they disengage from the productive sector and consume using their accumulated wealth. We follow the macro-finance literature and model retirement by imposing that households receive a lump sum of utility $\mathcal{U}(\mathcal{C}_{h,t}) = (1-\beta)\mathcal{C}_{h,t}^{1-\Gamma} / (1-\Gamma)$ from consuming a lump sum of consumption $\mathcal{C}_{h,t}$ at retirement. Here the parameter $1-\beta$ captures the relative weight that households place of retirement consumption. In Appendix B we micro-found this expression as the present discounted value of consumption during retirement and show that $1-\beta = 1/(\rho + \lambda_d)$ where $1/\lambda_d$ is the average retirement duration. After one household retires, it is immediately replaced by a new young household who receives initial wealth from a density $\phi_h(a)$ with mean $\bar{\phi}_h A_{h,t}$, where $A_{h,t}$ is the total wealth in the household sector. Transfers to new households are financed through a proportional wealth tax on existing households.

Households face two financial frictions. First, they face a “participation restriction” that hold-

ing capital stock k_t incurs the flow penalty:

$$\Psi_{h,k,t}(k_t, a_{h,t}) = \psi_h \left(\frac{q_t^k k_t}{a_{h,t}}, \frac{a_{h,t}}{A_t} \right) \Xi_{h,t} a_{h,t}, \text{ where } \psi_h \left(\frac{q_t^k k_t}{a_{h,t}}, \frac{a_{h,t}}{A_t} \right) = \frac{\bar{\psi}_k}{2(a_{h,t}/A_t)^{1+\alpha}} \left(\frac{q_t^k k_t}{a_{h,t}} \right)^2,$$

where $\bar{\psi}_k > 0$ and $\alpha > 0$ are parameters representing the severity of the constraint, $q_t^k k_t$ is the market value of the household's capital, $a_{h,t}/A_t$ is the household's share of wealth in the economy, A_t is the aggregate wealth, and $\Xi_{h,t}$ is a normalization factor to convert the penalty to marginal utility units. The function $\varphi_h(\cdot)$ is increasing in the fraction of wealth in capital $q_t^k k_t/a_{h,t}$ and decreasing in agent wealth $a_{h,t}$. So, the constraint imposes that wealthier agents face a lower penalty when directly holding capital, which we interpret as capturing the costs, education differences, and/or behavioral constraints associated with direct capital market participation. Second, households cannot write contracts with each other to insure against retirement shocks.

Financial intermediaries: There are two types of financiers servicing households: bankers (b) and fund managers (f). Each type of financier $j \in \{b, f\}$ has discount rate ρ and gets flow utility $u_j(c_{j,t}) = c_{j,t}^{1-\gamma_j}/(1-\gamma_j)$ from consuming $c_{j,t}$ flow goods. The financiers differ in what type of liabilities they issue. Bankers issue risk-free short-term deposits that pay an instantaneous rate r_t^d . Fund managers sell contracts to households that pay one good to the holder when they retire and exit. For convenience, we refer to these contracts as pensions although conceptually they combine both retirement provision and life-insurance. In addition, both bankers and fund managers can issue long-term bonds that mature at rate λ_m and pay 1 unit of the consumption good at maturity. On the asset side of the balance sheet, both banks and fund managers can purchase long-term bonds and capital subject to government regulatory frictions described below. Financiers of type $j \in (b, f)$ exit at rate λ_j and are replaced by new financiers with initial wealth drawn from a density $\phi_f(a)$ with mean $\bar{\phi}_f A_t$. Any remaining wealth from exiting financiers left over after recapitalizing new financiers is rebated to households proportional to their wealth. This means that financier exit can be interpreted as dividend payout and recapitalization in the manner of [Gertler and Kiyotaki \(2010\)](#) or as financial taxation scheme.

Markets: Each period, there are competitive markets for goods and assets. We use goods as the numeraire. Let r_t^d denote the interest rate on deposits and let $\mathbf{q}_t := (q_t^k, q_t^n, q_t^m)$ denote a vector with the price of capital, pensions, and bonds respectively. We guess and verify that each long-term asset $l \in \{k, n, m\}$ has a price process $dq_t^l = q_t^l(\mu_{q^l,t}dt + \sigma_{q^l,t}dW_t)$ where $\mu_{q^l,t}$ and $\sigma_{q^l,t}$ are the

geometric drift and volatility. We express the corresponding return processes by:

$$\begin{aligned}
dR_t^k &= r_t^k dt + \sigma_{q^k,t} dW_t, & r_t^k &:= \mu_{q^k,t} + \Phi(\iota) - \delta + (e^{z_t} - \iota)/q_t^k, \\
dR_{h,t}^n &= r_{h,t}^n dt + \sigma_{q^n,t} dW_t + (1/q_t^n) dN_{h,t}, & r_{h,t}^n &:= \mu_{q^n,t} \\
dR_{f,t}^n &= r_{f,t}^n dt + \sigma_{q^n,t} dW_t, & r_{f,t}^n &:= \mu_{q^n,t} + (1/q_t^n - 1) \lambda_h \\
dR_t^m &= r_t^m dt + \sigma_{q^m,t} dW_t & r_t^m &:= \mu_{q^m,t} + (1/q_t^m - 1) \lambda_m
\end{aligned}$$

where $dN_{h,t}$ denotes a Poisson process that is 1 when the household retires and $R_{h,t}^n$ and $R_{f,t}^n$ are the return on pensions to household and fund manager respectively.¹

Government: The government also issues zero coupon bonds that mature at rate λ_m and pay 1 unit of the consumption good at maturity. We impose that the government follows a fiscal rule that scales bond supply with aggregate capital stock, K_t , so $M_t = \mathcal{M}K_t$. The government raises flow wealth taxes at rate τ on all agents to finance the issuance of debt subject to the budget constraint $\tau A_t + q_t^m \mu_{M_t} M_t = \lambda_m M_t$ where μ_{M_t} is the growth rate of government debt.

In addition to its fiscal policy, the government also imposes regulatory constraints on the financial sector. Financiers of type $j \in \{b, f\}$ face the flow portfolio penalty:

$$\Psi_{j,t}(k_{j,t}, m_{j,t}, a_{j,t}) = \sum_{l \in \{k, m\}} \psi_{j,l,t} \Xi_{j,t} a_{j,t}, \text{ where } \psi_{j,l,t} = \frac{\bar{\psi}_{j,l}}{2} \max \left\{ 0, \frac{q_t^l l_{j,t}}{a_{j,t}} - \bar{\theta}_j^l \right\}^2, \quad (2.1)$$

for holding assets $\{k_{j,t}, m_{j,t}\}$ where $\bar{\theta}_j^l$ is the regulatory portfolio share limit for agent j in asset l , $\bar{\psi}_j^l$ is the tightness or weight of the regulatory requirement for asset j , $a_{j,t}$ is the wealth the financier, and $\Xi_{j,t}$ is again a normalization factor to convert the penalty to marginal utility units. We will internally calibrate $\{\bar{\theta}_j^l\}_{j \in \{b, f\}, l \in \{k, m\}}$ to match equilibrium portfolios and asset prices so we interpret the parameters as the implicit portfolio restrictions (or “asset market segmentation”) resulting from the combined current macro-prudential regulations rather than a direct mapping to one particular item of regulation.

Connection to the existing models: This environment has been constructed to nest the following collection of models and economic forces commonly studied in the macro-finance literature:

¹The fund manager earns a different return to a household because the fund managers diversifies across a continuum of households.

- (i) *Preferred habitat and perpetual youth models:* If we set $\beta = 0$ so the households only care about consumption at retirement and set $\mathcal{U}(\cdot)$ to be either the Type I or Type II agents from the Appendix in [Vayanos and Vila \(2021\)](#), then our households become analogous to the “preferred habitat agents” in [Vayanos and Vila \(2021\)](#) who only demand maturities with average maturity $1/\lambda_h$.² At the other extreme, if we set $\beta = 1$ so the households only care about consumption while young, then we recover the “perpetual youth” model from [Blanchard \(1985\)](#). In this case, households demand annuities that pay until they die, which, in principle, they could recreate synthetically by shorting the pensions and purchasing bonds.³ In this sense, the two extreme parameterizations, $\beta \in \{0, 1\}$, nest two of the most commonly used models of demand for pension/insurance products or long-term bonds. Our environment can be viewed as an intermediate model that nests these two forces. In particular, for the general setup with $\beta \in (0, 1)$, 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) *Segmented market models:* If we take the limit as $\bar{\psi}_k$ goes to infinity and the financial regulatory constraints are relaxed, then we get complete capital market segmentation: households have no access to the capital market while the financiers have full access. This is similar to [Basak and Cuoco \(1998\)](#), [Chien et al. \(2012\)](#), [Gertler and Kiyotaki \(2015\)](#), and other models with market segmentation that amplifies asset price volatility. Alternatively, if we take the limit as $\bar{\psi}_k$ goes to zero, then the households become unconstrained to participate in capital markets, like in [Brunnermeier and Sannikov \(2014\)](#). At either extreme, household portfolio choices become homogeneous and the household sector aggregates. By introducing heterogeneous household portfolio choice, we allow our model to intermediate between these extremes. However, it also what leads to many of the technical difficulties.
- (iii) *Heterogeneous type models:* There is a collection of models in which households have different fixed ex-ante types (e.g. different risk aversion) but all agents within a particular type

²In order to generate the full yield curve model in [Vayanos and Vila \(2021\)](#), we would also need to introduce variation in λ_h across households and bonds with different average maturities.

³Throughout the paper we focus on parametrizations where the households take long positions in the fund contracts and do not directly participate in the bond market. However, the model could just as easily be solved for the case where some households end up shorting the pension products.

make the same portfolio decisions (e.g. [Chan and Kogan \(2002\)](#), [Gomez \(2017\)](#)). These models can generate heterogeneous portfolio choices across the different types in the population and so can generate the aggregate asset portfolio for the household sector. Our model enriches these environments to allow household portfolio decisions to depend on wealth, which allows us to endogenize the transitions between types and match the micro level data on household portfolio choice.

- (iv) *Non-homothetic utility*: Another branch of the literature has introduced non-homothetic preferences that include wealth directly in the utility function so that wealthier agents “enjoy” wealth accumulation (e.g. [Carroll \(1998\)](#)). In our model, the combination of the consumption at retirement and the capital participation penalty introduces a similar force that wealthier agents hold more capital. We extend this literature by expanding the environment to include risky assets, which effectively allows us to use the asset pricing data to help discipline the shape of the non-homothetic preferences.

2.2 Equilibrium

To set up equilibrium, we use the following notation. We let $a_{j,t}$ denote the wealth of a generic individual agent in sector $j \in \{h, b, f\}$, where the indices h , b , and f refer to the household, banking and fund sectors respectively. 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 and let $\theta_{j,t}$ denote the vector of wealth shares chosen by an agent of type j with wealth $a_{j,t}$ at time t . We let $A_{j,t}$ denote the total wealth in sector j and A_t denote the total wealth in economy. We let $\omega_{j,t} := a_{j,t} / A_t$ denote the wealth share of a generic individual agent in sector j and $\Omega_{j,t} := A_{j,t} / A_t$ denote the share of total wealth in sector j . More generally, we use lower case letters to denote variables for an individual agent and upper case letters to denote sector or economy aggregates. We also use the notation $\mathbf{x} = (x_t)_{t \geq 0}$ to denote the stochastic process for variable x_t . Finally, we let $\mathcal{F}_t$ denote the filtration generated by $\{W_t\}_{t \geq 0}$ and use $\mathbb{E}_t[\cdot] := \mathbb{E}[\cdot | \mathcal{F}_t]$ to denote the expectation conditional on $\mathcal{F}_t$.

Household problem: The wealth of a household at time t is $a_{h,t} = q_t^k k_{h,t} + q_t^n n_{h,t} + d_{h,t}$, where $k_{h,t}$, $n_{h,t}$, and $d_{h,t}$ denote the household’s capital, pension, and deposit holdings. While they are

young, $a_{h,t}$ evolves by:

$$\begin{aligned} \frac{da_{h,t}}{a_{h,t}} &= \theta_{h,t}^k dR_t^k(\iota_{h,t}) + \theta_{h,t}^n dR_{h,t}^n + (1 - \theta_{h,t}^k - \theta_{h,t}^n) r_t^d dt - c_{h,t}/a_{h,t} dt - \tau_{h,t} dt \\ &=: \mu_{a_{h,t}} dt + \sigma_{a_{h,t}} dW_t \end{aligned} \quad (2.2)$$

where $(c_{h,t}, \theta_{h,t}, \iota_{h,t})$ are the household's consumption, portfolio shares, and investment rate at time t and $\tau_{h,t}$ is the net tax or transfer (per unit of wealth) incorporating government taxes, transfers, and any dividends from the financial sector. By expanding the returns we get:

$$\mu_{a_{h,t}} = r_t^d + \sum_{l \in \{n,k\}} \theta_{h,t}^l (r_t^l - r_t^d) - c_{h,t}/a_{h,t} - \tau_{h,t}, \quad \sigma_{a_{h,t}} = \sum_{l \in \{n,k\}} \theta_{h,t}^l \sigma_{q^l,t}$$

At retirement, households consume their remaining wealth and exit. So, taking the price process and their initial wealth a_{h,t_0} as given, a household born at t_0 chooses processes (c_h, θ_h, ι_h) to solve Problem (2.3) below:

$$\begin{aligned} \max_{c_h, \theta_h, \iota_h} \mathbb{E}_{t_0} \left[\int_{t_0}^T e^{-\rho_h(t-t_0)} \left(u(c_{h,t}) - \psi_{h,k}(\theta_{h,t}^k, \omega_{h,t}) \Xi_{h,t} a_{h,t} \right) dt + e^{-\rho T} \mathcal{U}(\mathcal{C}_{h,T}) \right] \\ \text{s.t. (2.2) and } \mathcal{C}_{h,T} \leq \left(1 - \theta_{h,t}^n + \theta_{h,t}^n/q_t^n \right) a_{h,t}. \end{aligned} \quad (2.3)$$

Financier problems: The wealth of a financier in sector $j \in \{b, f\}$ is given by $a_{b,t} = q_t^k k_{b,t} + q_t^m m_{b,t} + d_{b,t}$ for bankers and $a_{f,t} := q_t^k k_{f,t} + q_t^m m_{f,t} + q_t^n n_{f,t}$ for funds, where $k_{j,t}$, $m_{j,t}$, $d_{j,t}$, and $n_{j,t}$ are holdings of capital, bonds, deposits, and pensions respectively. It evolves by:

$$\begin{aligned} \frac{da_{j,t}}{a_{j,t}} &= \theta_{j,t}^k dR_t^k(\iota_{j,t}) + \theta_{j,t}^m dR_t^m + (1 - \theta_{j,t}^k - \theta_{j,t}^m) dR_t^j - c_{j,t}/a_{j,t} dt - \tau dt \\ &=: \mu_{a_{j,t}} dt + \sigma_{a_{j,t}} dW_t \end{aligned} \quad (2.4)$$

where $(c_{j,t}, \theta_{j,t}, \iota_{j,t})$ are the financier's consumption, portfolio shares, and investment rate at time t and dR_t^j is the return on the liabilities issued only by a financier of type j (i.e. the return on deposits $r_t^d dt$ for bankers $j = b$ and the return on pensions $dR_{f,t}^n = r_{f,t}^n dt + \sigma_{q^n,t} dW_t$ for fund

managers $j = f$). Once again expanding terms we get:

$$\begin{aligned}\mu_{a_j,t} &= r_t^j + \sum_{l \in \{m,k\}} \theta_{h,t}^l (r_t^l - r_t^j) - c_{h,t}/a_{h,t} - \tau_{j,t} \\ \sigma_{a_j,t} &= \theta_{h,t}^m \sigma_{q^m,t} + \theta_{h,t}^k \sigma_{q^k,t} + (1 - \theta_{h,t}^m - \theta_{h,t}^k) \sigma_{q^j,t}\end{aligned}$$

where r_t^j and $\sigma_{q^j,t}$ are the expected return and volatility of return process for the liabilities only issued by financiers of type j . So, taking as given the price processes and their initial wealth, $a_{j,0}$, a financier born at t_0 in sector $j \in \{b, f\}$ chooses processes (c_j, θ_j, ι_j) to solve Problem (2.5) below:

$$\max_{c_j, \theta_j, \iota_j} \left\{ \mathbb{E}_{t_0} \int_{t_0}^{\infty} e^{-\rho_j(t-t_0)} \left(u(c_{j,t}) - \sum_{l \in \{k,m\}} \psi_{j,l}(\theta_{j,t}^l) \Xi_{j,t} a_{j,t} \right) dt \right\} \text{ s.t. (2.4).} \quad (2.5)$$

Distributions: Throughout this paper, we work with the distribution of wealth shares, rather than wealth levels. We will show in Corollary 1 that bank and fund sectors aggregate because they face constraints that are independent of wealth and so it is sufficient to keep track of the aggregate wealth share in the financial sectors, $\{\Omega_{j,t}\}_{j \in \{b,f\}}$, rather than the distribution of wealth in each sector. We denote the evolution of $\Omega_{j,t}$ by:

$$d\Omega_{j,t} = \mu_{\Omega,j} \Omega_{j,t} dt + \sigma_{\Omega,j} \Omega_{j,t} dW_t$$

where the geometric drift and volatility, $(\mu_{\Omega,j}, \sigma_{\Omega,j})$, are equilibrium objects. The household sector will not aggregate because household portfolio constraints are wealth dependent and the incompletely insurable idiosyncratic retirement shocks generates a non-degenerate cross-section distribution of household wealth across the economy. We let $g_{h,t}$ denote the measure function of household wealth shares across the economy at time t for a filtration $\mathcal{F}_t$. We denote the evolution of $g_{h,t}$ by:

$$dg_{h,t}(\omega) = \mu_{g,t}(\omega) \omega dt + \sigma_{g,t}(\omega) \omega dW_t$$

where the geometric drift and volatility, $(\mu_{g,t}, \sigma_{g,t})$, are equilibrium objects. With some abuse of notation, we let $G_t = (\Omega_{b,t}, \Omega_{f,t}, g_{h,t})$ denote the collection of “distribution” states in the economy.

Definition 1 (Equilibrium). For a given set of government taxation policies, an equilibrium is a collection of $\mathcal{F}_t$ -adapted processes (K, r, q, G) and agent decision processes $(c_{i,j}, \iota_{i,j}, \theta_{i,j})$ for $i \in I$ and $j \in \{h, b, f\}$ such that:

1. Given price processes, households solve (2.3) and financial intermediaries solve (2.5),

2. The price processes (r, q) clear each market at each t :⁴ (a) Goods market: $\sum_{j \in \{h, b, f\}} C_{j,t} + \lambda_h \mathcal{C}_{h,t} = e^{z_t} K_t - \iota_t K_t$, where $\lambda_h \mathcal{C}_{h,t}$ is the aggregate household consumption upon retirement, (b) Capital market: $\sum_{j \in \{h, b, f\}} K_{j,t} = K_t$, (c) Pension market: $\sum_{j \in \{h, f\}} N_{j,t} = 0$, (d) Deposit market: $\sum_{j \in \{h, b\}} D_{j,t} = 0$, and (e) Bond market: $\sum_{j \in \{b, f\}} M_{j,t} = M_t$.

2.3 Recursive Characterization of Equilibrium

We characterize the equilibrium recursively. We denote the finite dimensional components of the aggregate state vector by $\mathbf{s} := (z, K, \Omega_b, \Omega_f)$ and the full aggregate state vector incorporating the household wealth share distribution by $\mathbf{S} := (\mathbf{s}, g_h)$.⁵ Individual agents also have their wealth a as an idiosyncratic state. For convenience, we define the state spaces for an individual agent by $\mathbf{x} := (a, z, K, \Omega_b, \Omega_f)$ and $\mathbf{X} := (\mathbf{x}, g_h)$. We let $\mu_{\mathbf{s}}, \sigma_{\mathbf{s}}, \mu_{\mathbf{x}}, \sigma_{\mathbf{x}}$ denote the geometric drift and volatility of $\mathbf{s}$ and $\mathbf{x}$ respectively. We let $V_j(a, \mathbf{S})$ and $\xi_j(a, \mathbf{S}) := \partial_a V_j(a, \mathbf{S})$ denote the value function and the derivative of the value function (which we refer to as the Stochastic Discount Factor (SDF)) for an agent of type $j \in \{h, b, f\}$ with individual state a . We let $\mu_{\xi_h}(a, \mathbf{S})$ and $\sigma_{\xi_h}(a, \mathbf{S})$ denote the geometric drift and volatility of ξ_j .

Theorems 1, 2, and 3 in the following subsections summarize the recursive characterization of equilibrium while the detailed steps are derived in Appendix A. We group the characterization into three blocks: (i) the optimization problems of the agents, (ii) the evolution of the distribution, and (iii) market clearing.

2.3.1 Block 1: Agent Optimization

We solve the agent optimization problems by setting up Hamilton-Jacobi-Bellman (HJB) equations for the households and financiers. The equations are involved so we put the details in Appendices A.1 and A.2. Here, we collect the main results.

Theorem 1 (Agent Optimization). *Given the price functions $(r^d, (q^l, r^l, \sigma_{q^l})_{l \in \{k, n, m\}})$, the agent choices at state $\mathbf{S}$ satisfy the following first-order-conditions (FOCs):*

- (i) *The (c_j, ι_j) choices satisfy $u'(c_j) = \xi_j(a, \mathbf{S})$ and $\Phi'(\iota_j) = (q^k(\mathbf{S}))^{-1}$ for all $j \in \{h, b, f\}$,*

⁴Here we continue with the notation that the capital letter $Y_{j,t}$ refers to the aggregate quantity of variable y in sector $j \in \{h, b, f\}$ and Y_t refers to the aggregate quantity of variable y across the economy.

⁵Technically, including all of Ω_b, Ω_f , and g_h in the state space is redundant because they sum to 1 but we write it this way to clarify what the full distribution looks like.

(ii) The household portfolio choice, $\theta_h = (\theta_h^k, \theta_h^n)$, satisfies:

$$[\theta_h^k] : \quad r^k(\mathbf{S}) - r^d(\mathbf{S}) = -\sigma_{\xi_h}(a, \mathbf{S})\sigma_{q^k}(\mathbf{S}) - \partial_{\theta_h^k}\psi_{h,k}(\theta_h^k, \omega_h) \quad (2.6)$$

$$[\theta_h^n] : \quad r^n(\mathbf{S}) - r^d(\mathbf{S}) = -\sigma_{\xi_h}(a, \mathbf{S})\sigma_{q^n}(\mathbf{S}) - \lambda_h \left(\frac{1}{q^n(\mathbf{S})} - 1 \right) \frac{\mathcal{U}'(\mathcal{C})}{\xi_h(a, \mathbf{S})}, \quad (2.7)$$

(iii) The portfolio choice for financiers in sector j , $\theta_j = (\theta_j^k, \theta_j^m)$, satisfies the FOC:

$$[\theta_j^k] : \quad r^k(\mathbf{S}) - r^j(\mathbf{S}) = -\sigma_{\xi_j}(\mathbf{S}; \theta_j)(\sigma_{q^k}(\mathbf{S}) - \sigma_{q^j}(\mathbf{S})) - \partial_{\theta_j^k}\psi_{j,k}(\theta_j^k) \quad (2.8)$$

$$[\theta_j^m] : \quad r^m(\mathbf{S}) - r^j(\mathbf{S}) = -\sigma_{\xi_j}(\mathbf{S}; \theta_j)(\sigma_{q^m}(\mathbf{S}) - \sigma_{q^j}(\mathbf{S})) - \partial_{\theta_j^m}\psi_{j,m}(\theta_j^m), \quad (2.9)$$

where (r^j, σ_{q^j}) is the return and volatility for the liability issued by sector j financiers⁶. The SDFs ξ_j for $j \in \{h, b, f\}$ satisfy the continuous time Euler equations:

$$\begin{aligned} \rho_h + \lambda_h &= \mu_{\xi_h}(a, \mathbf{S}) + r^d(\mathbf{S}) - \tau_h - \psi_{h,k}(\theta_h^k, \omega_h) - \partial_{\omega_h}\psi_{h,k}(\theta_h^k, \omega_h)\omega_h + \partial_{\theta_h^k}\psi_{h,k}(\theta_h^k, \omega_h)\theta_h^k, \\ \rho_j + \lambda_j &= \mu_{\xi_j}(a, \mathbf{S}) + r^j(\mathbf{S}) - \tau_j + \sigma_{\xi_j}(\mathbf{S})\sigma_{q^j}(\mathbf{S}) - \sum_{l \in \{k, m\}} \left(\psi_{j,l} - \partial_{\theta_j^l}\psi_{j,l}(\theta_j^l) \right), \quad j \in \{b, f\} \end{aligned}$$

where the drift and volatility of ξ_j are characterized by Itô's lemma:

$$\begin{aligned} \mu_{\xi_j}(a, \mathbf{S})\xi_j(a, \mathbf{S}) &= (D_x\xi_j(a, \mathbf{S}))^T \boldsymbol{\mu}_x(a, \mathbf{S}) + 0.5 \text{tr} \left\{ \Sigma^x(a, \mathbf{S}) D_x^2 \xi_j(a, \mathbf{S}) \right\} \\ &\quad + \langle \sigma_a(\mathbf{S}; \theta_j) a \partial_a \partial_g \xi_j(a, \mathbf{S})(\cdot), \tilde{\sigma}_g(\cdot, \mathbf{S}) \rangle + \mathcal{G}_g \xi_j(a, \mathbf{S}) \\ \sigma_{\xi_j}(a, \mathbf{S})\xi_j(a, \mathbf{S}) &= (\boldsymbol{\sigma}_x(a, \mathbf{S}) \odot \mathbf{x})^T D_x \xi_j(a, \mathbf{S}) + \langle \partial_g \xi_j(a, \mathbf{S})(\cdot), \sigma_g(\cdot, \mathbf{S}) \rangle, \end{aligned}$$

and $\Sigma^x(a, \mathbf{S}) := (\boldsymbol{\sigma}_x(a, \mathbf{S}) \odot \mathbf{x})^T (\boldsymbol{\sigma}_x(a, \mathbf{S}) \odot \mathbf{x})$, $\partial_g \xi_j$ is the Frechet derivative of ξ_h w.r.t. g , and $\mathcal{G}_g \xi_h(a, \mathbf{S})$ is a collection of Frechet derivative terms given by equation (A.1) in Appendix A.1.

Proof. See Appendix A.4. □

The FOCs for c_h and ι_h are standard in the continuous time macro-finance literature. The former equates the household's marginal utility of consumption to their marginal value of wealth (i.e. their stochastic discount factor). The later equates the marginal marginal cost of producing capital to the marginal benefit of selling capital.

⁶As before, for $j = b$ we have $r^j = r^d$ and $\sigma_{q^j} = 0$ while for $j = f$ we have $r^j = r_f^n$ and $\sigma_{q^j} = \sigma_{q^n}$.

The FOCs for portfolio choice (2.6)-(2.9) are less standard and so are central to understanding the economics in our model. The portfolio FOCs are partially characterized by the standard trade-off between earning an expected excess return (the LHS of each equation) and facing risk from the co-movement between the agent’s stochastic discount factor and the asset prices (the first term on the RHS of each equation). However, in addition, they are also characterized by distortions from the household financial frictions and government financial regulation (the final term on the RHS of each equation). The equilibrium dynamics will ultimately be governed by the interaction between these distortions that restrict portfolio adjustment and endogenous price volatility.

We can use the portfolio FOCs to study how each constraint influences asset demand. To understand the impact of household retirement consumption, we can rearrange the FOC for pension contracts, equation (2.7), in the form of [Vayanos and Vila \(2021\)](#) to get the following expression for pension demand:

$$\underbrace{\lambda_h \left(\frac{1}{q_t^n(\mathbf{S})} - 1 \right) \frac{\mathcal{U}'(\mathcal{C})}{\xi_h(a, \mathbf{S})}}_{\text{“Preferred habitat” component}} = - \underbrace{(r_{h,t}^n(\mathbf{S}) - r_t^d(\mathbf{S}))}_{\text{Excess return}} - \underbrace{\sigma_{\xi_h}(a, \mathbf{S}) \sigma_{q_t^n}(\mathbf{S})}_{\text{“risk adjustment”}} \quad (2.10)$$

“Demand shifter”

which can be interpreted as a generalization of the preferred habitat demand function from [Vayanos and Vila \(2021\)](#). Like [Vayanos and Vila \(2021\)](#), we have the “inelastic” preferred habitat term coming from the need for assets with duration λ_h (the LHS in equation (2.10)). However, instead of exogenous demand shifters, we instead have that the household’s risk-return tradeoff (the RHS in equation (2.10)) shifts household demand. We expand on this comparison both theoretically and quantitatively in our counterfactual experiments in Subsection 6.4.

To understand the role of the household capital market participation constraint, observe from the FOC for θ_h^k , equation (2.7), that the term $-\partial_{\theta_h^k} \psi_{h,k}(\theta_h^k, \omega_h)$ decreases household demand for capital by penalizing capital holding. The penalty is less pronounced as the household gets wealthier (ω_h increases) so, all else equal, wealthier households hold more capital. This increasing relationship between household wealth and capital holdings is what ultimately allows us to match the empirical data showing that poorer households predominately hold deposits while wealthier households predominately hold capital. It is also what breaks aggregation across the household sector and so necessitates keeping track of the household wealth distribution. By contrast, the portfolio choice problems for the financial intermediaries, equations (2.8) and (2.9), do not depend upon financial intermediary wealth and so the banking and fund management sectors can be aggregated. We formalize this in Proposition 1 below.

Proposition 1. *The firm and banking sectors aggregate but the household sector does not.*

Proof. See Appendix A.2. □

2.3.2 Block 2: Distribution Evolution

Given the individual optimal decisions, we proceed to study how the distributions evolve. From Proposition 1, we only need to keep track of the aggregate wealth share dynamics for the financial sectors because we can aggregate within each of those sectors. However, the household sector does not aggregate and so we need to keep track of the full distribution of household wealth. We summarize the equilibrium laws of motion for the different distributions in Theorem 2 below.

Theorem 2 (Distribution evolution). *Given price functions $(r^d, (q^l, r^l, \sigma_{q^l})_{l \in \{k,n,m\}})$ and agent optimization functions $(\xi_j, c_j, \theta_j)_{j \in \{h,b,f\}}$, ι :*

(i) *At the sector level, the wealth share for financial sector $j \in \{b, f\}$ follows:*

$$\begin{aligned} \frac{d\Omega_{j,t}}{\Omega_{j,t}} = & \left(\mu_{A_j}(\mathbf{S}_t) + \lambda_j (\bar{\phi}_j / \Omega_{j,t} - 1) - \mu_A(\mathbf{S}_t) + (\sigma_A(\mathbf{S}_t) - \sigma_{A_j}(\mathbf{S}_t)) \sigma_A(\mathbf{S}_t) \right) dt \\ & + (\sigma_{A_j}(\mathbf{S}_t) - \sigma_A(\mathbf{S}_t)) dW_t \end{aligned} \quad (2.11)$$

where (μ_A, σ_A) are the geometric drift and volatility of aggregate wealth:

$$\begin{aligned} \mu_A(\mathbf{S}_t) &= \vartheta(\mathbf{S}_t)(\mu_{q^k}(\mathbf{S}_t) + \Phi(\iota) - \delta) + (1 - \vartheta(\mathbf{S}_t))\mu_{q^m}(\mathbf{S}_t) \\ \sigma_A(\mathbf{S}_t) &= \vartheta(\mathbf{S}_t)\sigma_{q^k}(\mathbf{S}_t) + (1 - \vartheta(\mathbf{S}_t))\sigma_{q^m}(\mathbf{S}_t) \end{aligned}$$

where aggregate wealth is given by $A_t = q_t^k K_t + q_t^m M_t$, and the aggregate wealth share in capital is $\vartheta(\mathbf{S}_t) := q^k(\mathbf{S}_t)K_t / (q^k(\mathbf{S}_t)K_t + q^m(\mathbf{S}_t)M_t)$.

(ii) *Within the household sector, the density of household wealth shares follows:*

$$\begin{aligned} dg_{h,t}(\omega) = & \left(\underbrace{\lambda_h \phi_h(\omega)}_{\text{Entry}} - \underbrace{\lambda_h g_{h,t}(\omega)}_{\text{Retirement}} - \underbrace{\partial_\omega [\mu_\omega(\omega, \mathbf{S}_t) \omega g_{h,t}(\omega)]}_{\text{Drift}} \right. \\ & \left. + \underbrace{0.5 \partial_{\omega\omega} [\sigma_\omega^2(\omega, \mathbf{S}_t) \omega^2 g_{h,t}(\omega)]}_{\text{Ave. Dispersion from TFP Shocks}} \right) dt - \underbrace{\partial_\omega [\sigma_\omega(\omega, \mathbf{S}_t) \omega g_{h,t}(\omega)]}_{\text{TFP Shock Exposure}} dW_t, \quad \text{s.t.} \quad (2.12) \\ \mu_\omega(\omega, \mathbf{S}_t) &= \mu_a(\omega A(\mathbf{S}_t), \mathbf{S}_t) - \mu_A(\mathbf{S}_t) + (\sigma_A(\mathbf{S}_t) - \sigma_a(\omega A(\mathbf{S}_t), \mathbf{S}_t)) \sigma_A(\mathbf{S}_t), \\ \sigma_\omega(\omega, \mathbf{S}_t) &= \sigma_a(\omega A(\mathbf{S}_t), \mathbf{S}_t) - \sigma_A(\mathbf{S}_t). \end{aligned}$$

Proof. See Appendix A.3. □

Theorem 2 shows the forces that spread out and condense the wealth distribution. At the sector level, the dynamics are described by equation (2.11). A positive drift spread $\mu_{A_j} - \mu_A > 0$ implies that the sector gains wealth share because it accumulates wealth more quickly than the overall economy while $\Omega_{j,t} > \bar{\phi}_j$ implies that the financial sector loses wealth share through financier exit, dividend payment, and recapitalization because the sector currently has more wealth share than new financiers. Within the household sector, the dynamics are described by the *stochastic Kolmogorov Forward Equation (KFE)* (2.12). The first two terms are the rate of household entry and the rate of exit at position ω in the density. The third term captures the impact of the drift of ω on the evolution of the density. The final two terms capture the impact of aggregate TFP shocks on the distribution. Because households have different portfolios with different risk exposure, TFP shocks potentially spread out the distribution. From all these terms, we can see that, within the household sector, differences in μ_a across households potentially spread out households while retirement and entry terms stabilize the distribution.

2.3.3 Block 3: Market Clearing and Itô Consistency

Finally, we impose equilibrium by ensuring markets clear. The following theorem restates the equilibrium conditions using our recursive formulation.

Theorem 3 (Market clearing and consistency). *At each $\mathbf{S}$, the goods market clearing is:*

$$\int_0^1 (c_h(\omega A, \mathbf{S}_t) + \lambda_h C_h(\omega A, \mathbf{S}_t)) g_{h,t}(\omega) d\omega + \sum_{j \in \{b,n\}} C_j(\mathbf{S}_t) = (e^z - \iota)K$$

and the equilibrium price functions (r, q^n, q^k, q^m) satisfy asset market clearing conditions:

$$\begin{aligned} \int \theta_h^d(\omega A(\mathbf{S}), \mathbf{S}) g_h(\omega) d\omega + \theta_b^d(\mathbf{S}) \Omega_b &= 0 \\ \theta_f^n(\mathbf{S}) \Omega_f + \int \theta_h^n(A(\mathbf{S}), \mathbf{S}) g_h(\omega) d\omega &= 0 \\ \int \theta_h(A(\mathbf{S}), \mathbf{S}) g_h(\omega) d\omega + \theta_f(\mathbf{S}) \Omega_f + \theta_b(\mathbf{S}) \Omega_b &= \vartheta(\mathbf{S}) \\ \theta_b^m(\mathbf{S}) \Omega_b + \theta_f^m(\mathbf{S}) \Omega_f &= 1 - \vartheta(\mathbf{S}) \end{aligned}$$

and the long term assets prices must satisfy consistency with Itô's Lemma for $l \in \{k, n, m\}$:

$$\begin{aligned}\mu_{q^l}(\mathbf{S})q^l(\mathbf{S}) &= (D_x q^l(\mathbf{S}))^T(\boldsymbol{\mu}_s \odot \mathbf{s}) + \frac{1}{2}tr\{\Sigma^s D_s^2 q^l(\mathbf{S})\} + \mathcal{G}_g q^l(\mathbf{S}) \\ \sigma_{q^l}(\mathbf{S})q^l(\mathbf{S}) &= (\sigma_s(\mathbf{S}) \odot \mathbf{S})^T(D_s q^l(\mathbf{S})) + \langle \partial_g q^l(a, \mathbf{S})(\cdot), \sigma_g(\cdot, \mathbf{S}) \rangle\end{aligned}\tag{2.13}$$

where once again $\Sigma^s(\mathbf{S}) := (\sigma_s(\mathbf{S}) \odot \mathbf{S})^T(\sigma_s(\mathbf{S}) \odot \mathbf{s})$, $\mathcal{G}_g$ is the collection of Frechet derivatives defined by equation (A.1) in the Appendix, and $\partial_g q^l$ is the Frechet derivative.

Proof. See Appendix A.4. □

Theorem 3 illustrates one of the main difficulties for macro-finance. The asset market clearing conditions only implicitly pin down the pricing functions for the long-term assets (q^k, q^n, q^m). Consequently, we also need to impose additional cross equation restrictions to ensure that long-term asset price processes are consistent with equilibrium through Itô's Lemma.

2.4 Technical Comparison to Other Models

Our equilibrium characterization is difficult to solve because, unlike in most other models, all three blocks are non-trivial. To understand why this is the case, it is instructive to compare our model to other macro-finance models, as summarized in Table 1 and discussed below:

- (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. So, the set of equations can be simplified to a differential equation for q . We use this as a test case in Appendix G.
- (ii). For the continuous time version of [Krusell and Smith \(1998\)](#) discussed in [Gu et al. \(2024\)](#), there is a distribution of agents so block 2 is non-trivial. However, this model has no long-term assets and so we can derive closed form expressions for all prices in terms of the distribution. This implies that block 3 can be trivially satisfied and we can combine all equilibrium conditions into one “master” partial differential equation. Other models with multiple short term assets, like [Krusell and Smith \(1997\)](#), have similarly trivial market clearing conditions because they do not have long-term asset prices that necessitate Itô consistency conditions of the form of equation (2.13).
- (iii). For models such as [Basak and Cuoco \(1998\)](#) and [Brunnermeier and Sannikov \(2014\)](#) discussed in [Gopalakrishna \(2021\)](#), there are long-term assets and so we need the non-trivial

Models	Non-Trivial Blocks			Method
	1 (Opt.)	2 (Dist.)	3 (Asset q)	
Representative Agent (RA) (à la Lucas (1978))	simple	NA	simple	Finite difference
Heterogeneous Agents (HA) (à la Krusell and Smith (1998))	✓	✓	simple	Gu et al. (2024)
Long-lived assets (à la Brunnermeier and Sannikov (2014))	closed-form	low-dim	✓	Duarte et al. (2024) Gopalakrishna (2021)
HA + Long-lived assets	✓	✓	✓	This paper

Table 1: Key models and solution methods. A tick indicates that the block is non-trivial.

Itô consistency conditions of the form (2.13). However, in these models the HJBE can be solved in closed form so that we can get an analytical expression for the dependence of the value function on idiosyncratic wealth. This means that block 1 can be reduced to a scaled PDE and substituted into the block 3.

3 Computational Methodology and Algorithm

In this section, we outline our algorithm for characterizing solutions to our model. Conceptually, our approach can be viewed as a type of “projection” onto a neural network (building on [Duarte et al. \(2024\)](#), [Gopalakrishna \(2021\)](#), [Gu et al. \(2024\)](#), and other papers that adapt the Neural Galerkin literature for using deep learning to solve differential equations). 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 constructing a loss function that captures the equilibrium conditions, and
- (c) Finding the neural network parameters that minimize the loss function on randomly sampled points from the state space (often referred to as “training” the neural network).

Once we have trained neural network approximations to the equilibrium functions, we use these neural networks to simulate the equilibrium evolution of the state space and compute ergodic outcomes, impulse response functions, and other analyses.

Although this approach may appear straightforward to describe, implementing it successfully for macro-finance models has proven difficult and involves many non-trivial decisions. In particular, we need to decide which distribution approximation is most effective, how to set up the loss function, and how to customize the algorithm. The main reason these become important decisions is because the combination of multiple long-term assets, aggregate risk, and a non-degenerate wealth distribution make the equilibrium conditions difficult to impose. In this section we describe the features of our algorithm and discuss the tradeoffs associated with different approaches.

3.1 Part (a): Finite Dimensional Distribution Approximation

A fundamental challenge for heterogeneous agent macro-finance models is that the state space contains a density, g_h , which is an infinite-dimensional object. To make computational progress, we must adopt a finite dimensional approximation to the density. In this paper, we use a hybrid approach. When training the neural network to approximate the equilibrium functions, we approximate the distribution by a finite collection of $I < \infty$ price taking agents so the approximate aggregate state space becomes $\hat{\mathcal{S}} := (z, K, \omega_1, \dots, \omega_I, \Omega_b, \Omega_f) \in \mathcal{S}$ where ω_i is the wealth share of agent i , (Ω_b, Ω_f) are again the wealth shares of the financial sectors, and $\mathcal{S} = \mathbb{R} \times \mathbb{R}^+ \times [0, 1]^{I+2}$ is space of possible aggregate states.⁷ We rewrite the equilibrium conditions formally on the finite state space in Appendix C.2.

When simulating the economy we use our neural network approximations to the equilibrium functions to construct a finite difference approximation to the equilibrium KFE (2.12). We do this by extending the approach developed in Gu et al. (2024) and summarized in Appendix C.5.

Discussion: Previous work has studied the benefits and costs associated with different distribution approximations (e.g. see Gu et al. (2024)). We highlight key points here:

- (i) A frequently raised concern with the finite-agent approach is that simulations of a finite agent economy contain idiosyncratic noise that could make the training unstable and lead to an equilibrium that is inconsistent with the original model. However, we mitigate these concerns in a number of ways:
 - (a) We solve for an equilibrium in which agents behave as price takers and forecast prices under the assumption that the idiosyncratic exit shocks have averaged out. This means

⁷Once again, there is redundancy because the shares $\omega_1, \dots, \omega_I, \Omega_b, \Omega_f$ but we write it this way for clarity and consistency.

that the perceived state space evolution does not contain idiosyncratic risk and so our equilibrium differential equations contain no idiosyncratic noise.

- (b) We do not simulate the model to draw training points and so concerns about simulation accuracy are unrelated to neural network approximation accuracy.
 - (c) When we do need to simulate the solved model to generate time paths or impulse responses, we use our finite-agent 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 rather than the finite agent economy. Figure 8 in the Appendix shows a histogram of the difference in the impulse responses computed using the finite difference approximation to the KFE and the N-agent simulation to illustrate the importance of working with the KFE for simulations.
- (ii) Another concern that is sometimes raised about finite dimensional distributional approximations is that they lead to master equations with high dimensional Itô terms that are impracticable to solve. Our hybrid approach of training on a finite agent economy and then reconstructing the KFE does not require such a large population to make the Itô terms intractable, even when calculating the Itô terms naively. Moreover, [Duarte et al. \(2024\)](#) shows that the Itô term in high-dimensional differential equations can be computed very efficiently by calculating second order directional derivatives. For all these reasons, we find that dimensionality is not what makes the problem most challenging. Instead, it is the high curvature and sensitivity of the market clearing conditions that make the problem challenging. In the next section, we discuss how we construct the loss function to overcome these issues.

3.2 Part (b): Neural Network Representation and Loss Function

We need to construct a loss function for the deep learning algorithm to minimize. One approach would be to represent all variables by neural networks on the approximate state space $\hat{S}$ and then sum every condition in Theorems 1, 2, and 3 into one large loss function. Although this approach might be feasible in principle, it has proven difficult to implement in practice for macro-finance models (e.g. see discussion in [Azinovic and Zemlicka \(2026\)](#)). We overcome these problems by reorganizing the equilibrium equations to construct a loss function that is “easier” for the deep learning algorithm to train in terms of stability and convergence time. We summarize our approach here, provide a more detailed description in Appendix C, and illustrate how alternative loss functions create difficulties in Appendix E.

First, we express the equilibrium objects as functions of the aggregate states without any explicit dependence on idiosyncratic states. To understand what this means, consider the household stochastic discount factor function ξ_h . This was defined in the household problem in Section 2.3 as a function of individual household wealth and the aggregate state space $\xi_h(a_i, \hat{\mathbf{S}})$. However, once general equilibrium is imposed it also has an equilibrium representation that is only a function of the aggregate states $\Xi_h(\hat{\mathbf{S}}) := \xi_h(\omega_i A(\hat{\mathbf{S}}), \hat{\mathbf{S}})$ where $A(\hat{\mathbf{S}}) = q^k(\hat{\mathbf{S}})K + q^m(\hat{\mathbf{S}})M$ is equilibrium aggregate wealth. Likewise, we can write all agent value and policy functions directly in terms of the aggregate state $\hat{\mathbf{S}}$. Conceptually, this is analogous to imposing equilibrium using the “little-k and big-K” approach described in Chapter 7 of [Sargent and Ljungqvist \(2000\)](#). We solve directly for the equilibrium functions (i.e. Ξ_h) rather than for the partial equilibrium functions (i.e. ξ_h). We discuss this further in Appendix C.3.

Second, we parametrize the functions by Neural Nets:

$$\begin{aligned}
\hat{\eta}_j : \mathcal{S} &\rightarrow \mathbb{R}, (\hat{\mathbf{S}}, \Theta_{\eta_h}) \mapsto \hat{\eta}_j(\hat{\mathbf{S}}; \Theta_{\eta_h}), \quad \forall j \in \{h, f, b\} \\
\hat{\theta}_h^l : \mathcal{S} &\rightarrow \mathbb{R}, (\hat{\mathbf{S}}, \Theta_{\theta_h}) \mapsto \hat{\theta}_h^l(\hat{\mathbf{S}}; \Theta_{\theta_h^l}), \quad \forall l \in \{k, n\}, \\
\hat{\theta}_f^m : \mathcal{S} &\rightarrow \mathbb{R}, (\hat{\mathbf{S}}, \Theta_{\theta_f}) \mapsto \hat{\theta}_f^m(\hat{\mathbf{S}}; \Theta_{\theta_f^m}), \\
\hat{q}^l : \mathcal{S} &\rightarrow \mathbb{R}, (\hat{\mathbf{S}}, \Theta_{q^l}) \mapsto \hat{q}^l(\hat{\mathbf{S}}; \Theta_{q^l}), \quad \forall l \in \{n, m\} \\
\hat{\mu}_{q^k} : \mathcal{S} &\rightarrow \mathbb{R}, (\hat{\mathbf{S}}, \Theta_{\mu, q^k}) \mapsto \hat{\mu}_{q^k}(\hat{\mathbf{S}}; \Theta_{\mu, q^k}), \\
\hat{\sigma}_{q^l} : \mathcal{S} &\rightarrow \mathbb{R}, (\hat{\mathbf{S}}, \Theta_{\sigma, q^l}) \mapsto \hat{\sigma}_{q^l}(\hat{\mathbf{S}}; \Theta_{\sigma, q^l}), \quad \forall l \in \{k, n, m\}
\end{aligned} \tag{3.1}$$

where $\eta_j := c/a$ denotes the flow consumption-to-wealth ratio for an agent of type $j \in \{h, f, b\}$, Θ_ν denotes the parameters for the Neural Net approximation of variable ν , and we use Θ to refer the collection of neural network parameters across all approximations. In Appendix C.4 we show that, given the neural network approximation functions for $(\hat{\eta}_j)_{j \in \{h, f, b\}}$, $(\hat{\theta}_h^l)_{l \in \{k, n\}}$, θ_f^m , $(\hat{q}^l)_{n, m}$, and $(\hat{\mu}_{q^l}, \hat{\sigma}_{q^l})_{l \in \{k, n, m\}}$, at each state $\hat{\mathbf{S}}$ we can solve for the other equilibrium objects explicitly using linear algebra.

Finally, we construct the loss function for the deep learning algorithm to minimize. In doing so, we impose market clearing explicitly rather than including the market clearing conditions as part of the loss function. In particular, we compute the bank portfolio choices $(\theta_b^k, \theta_b^m, \theta_b^d)$ as residuals to impose market clearing in the capital, bond, and deposit markets, θ_f^n as a residual to impose clearing in the pension market, and q_t^k as a residual to impose clearing in the goods market.⁸

⁸We enforce the household and fund manager budget constraints explicitly. The banker’s budget constraint is enforced by Walras’s law. The choice of which θ ’s to treat as residuals and approximate with neural networks is not

This implies that the neural network approximations need to satisfy the following equations for all $j \in \{b, f\}$, $l \in \{k, m\}$, and $i \in \{k, m, n\}$:

$$\mathcal{L}_{\eta_h}(\hat{\mathbf{S}}) = r^d(\hat{\mathbf{S}}) - \rho_h - \lambda_h - \tau_h + \mu_{\hat{\Xi}_h}(\hat{\mathbf{S}}) - \psi_{h,k}(\hat{\mathbf{S}}) - \partial_{\hat{\omega}_h} \psi_{h,k}(\hat{\mathbf{S}}) \hat{\omega}_h(\mathbf{S}) + \partial_{\hat{\theta}_h^k} \psi_{h,k}(\hat{\mathbf{S}}) \hat{\theta}_h^k(\hat{\mathbf{S}}),$$

$$\mathcal{L}_{\eta_j}(\hat{\mathbf{S}}) = r^j(\hat{\mathbf{S}}) - \rho_j - \lambda_j - \tau_j + \mu_{\Xi_j}(\hat{\mathbf{S}}) + \sigma_{\Xi_j}(\hat{\mathbf{S}}) \sigma_{q^j}(\hat{\mathbf{S}}) - \sum_{l \in \{k, m\}} \left(\psi_{j,l}(\hat{\mathbf{S}}) - \partial_{\theta_j^l} \psi_{j,l}(\hat{\mathbf{S}}) \right),$$

$$\mathcal{L}_{\theta_h^k}(\hat{\mathbf{S}}) = r^k(\hat{\mathbf{S}}) - r^d(\hat{\mathbf{S}}) + \sigma_{\hat{\Xi}_h}(\hat{\mathbf{S}}) \hat{\sigma}_{q^k}(\hat{\mathbf{S}}) + \partial_{\theta_h^k} \psi_{h,k}(\hat{\mathbf{S}}),$$

$$\begin{aligned} \mathcal{L}_{\theta_h^n}(\hat{\mathbf{S}}) &= r^n(\hat{\mathbf{S}}) - r^d(\hat{\mathbf{S}}) + \frac{\lambda_h}{\hat{\Xi}_h(\hat{\mathbf{S}})} \left(1/\hat{q}^n(\hat{\mathbf{S}}) - 1 \right) \mathcal{U}' \left((1 + \theta^n(\mathbf{S})(1/\hat{q}^n(\hat{\mathbf{S}}) - 1)) \omega A(\hat{\mathbf{S}}) \right) \\ &\quad + \sigma_{\hat{\Xi}_h}(\hat{\mathbf{S}}) \hat{\sigma}_{q^n}(\hat{\mathbf{S}}), \end{aligned}$$

$$\mathcal{L}_{\theta_f^l}(\hat{\mathbf{S}}) = r^l(\hat{\mathbf{S}}) - r_f^n(\hat{\mathbf{S}}) + \sigma_{\hat{\Xi}_f}(\hat{\mathbf{S}}) (\hat{\sigma}_{q^l}(\hat{\mathbf{S}}) - \hat{\sigma}_{q^n}(\hat{\mathbf{S}})) + \partial_{\theta_f^l} \psi_{f,l}(\theta_f^l(\hat{\mathbf{S}})),$$

$$\mathcal{L}_{\theta_b^m}(\hat{\mathbf{S}}) = r^m(\hat{\mathbf{S}}) - r^d(\hat{\mathbf{S}}) + \sigma_{\hat{\Xi}_b}(\hat{\mathbf{S}}) (\hat{\sigma}_{q^m}(\hat{\mathbf{S}}) - \hat{\sigma}_{q^n}(\hat{\mathbf{S}})) + \partial_{\theta_b^m} \psi_{b,m}(\theta_b^m(\hat{\mathbf{S}})),$$

$$\mathcal{L}_{\mu_{q^k}}(\hat{\mathbf{S}}) = \mu_{q^k}(\hat{\mathbf{S}}) \hat{q}^k(\hat{\mathbf{S}}) - (D_{\hat{\mathbf{S}}} \hat{q}^k(\hat{\mathbf{S}}))^T (\boldsymbol{\mu}_{\hat{\mathbf{S}}}(\hat{\mathbf{S}}) \odot \hat{\mathbf{S}}) - 0.5 \text{tr} \left\{ \Sigma^{\hat{\mathbf{S}}}(\mathbf{S}) D_{\hat{\mathbf{S}}}^2 \hat{q}^k(\hat{\mathbf{S}}) \right\}$$

$$\mathcal{L}_{\sigma_{q^i}}(\hat{\mathbf{S}}) = \hat{\sigma}_{q^i}(\hat{\mathbf{S}}) \hat{q}^i(\hat{\mathbf{S}}) - (\boldsymbol{\sigma}_{\hat{\mathbf{S}}}(\hat{\mathbf{S}}) \odot \hat{\mathbf{S}})^T (D_{\hat{\mathbf{S}}} \hat{q}^i(\hat{\mathbf{S}})),$$

where $\hat{\Xi}_j(\hat{\mathbf{S}}) = u'(\hat{\eta}_j(\hat{\mathbf{S}}) \omega_j \hat{q}^k(\hat{\mathbf{S}}) K)$ and $\Sigma^{\hat{\mathbf{S}}}(\mathbf{S}) := (\boldsymbol{\sigma}_{\hat{\mathbf{S}}}(\hat{\mathbf{S}}) \odot \hat{\mathbf{S}})^T (\boldsymbol{\sigma}_{\hat{\mathbf{S}}}(\hat{\mathbf{S}}) \odot \hat{\mathbf{S}})$. The first two equations are Euler equations, the next four equations are agent first order conditions, and the final two equations are the Itô consistency conditions for the price functions. The mapping between the losses and the neural network approximation is mostly explicit in the naming. The exceptions are $\mathcal{L}_{\theta_f^k}(\hat{\mathbf{S}})$ and $\mathcal{L}_{\theta_b^m}(\hat{\mathbf{S}})$ which correspond to loss functions for $(q^l)_{l \in \{n, m\}}$. The total loss function is:

$$\hat{\mathcal{L}}(\hat{\mathbf{S}}; \boldsymbol{\Theta}) = \left[\sum_{j \in \{h, b, f\}} \mathcal{L}_{\eta_j} + \sum_{l \in \{k, n\}} \mathcal{L}_{\theta_h^l} + \sum_{l \in \{k, m\}} \mathcal{L}_{\theta_f^l} + \mathcal{L}_{\theta_b^m} + \mathcal{L}_{\mu_{q^k}} + \sum_{l \in \{k, n, m\}} \mathcal{L}_{\sigma_{q^l}} \right] (\hat{\mathbf{S}}; \boldsymbol{\Theta}) \quad (3.2)$$

Discussion Our loss function attempts to balance a collection of difficulties that the macro-finance deep learning literature has faced. We describe the key ideas here and provide computational illustrations of these points in Online Appendix E.

- (i) We solve directly for the equilibrium functions (e.g. $\Xi_h(\mathbf{S})$) rather than for the partial equilibrium functions (e.g. $\xi_h(a, \mathbf{S})$) because it is a more stable and computationally efficient way of imposing general equilibrium. When we solve for the partial equilibrium represen-

unique. It is equally efficient to treat the fund portfolio choice as the residual to clear the capital and bond markets.

tation $\xi_h(a, \mathbf{S})$ we end up wasting effort training on off-equilibrium states and then need to impose additional loss functions to impose general equilibrium. By contrast, solving for the general equilibrium representation $\Xi_h(\mathbf{S})$ imposes market clearing directly and only trains on the equilibrium states. We illustrate the difference in Appendix E.1 by comparing training on the different representations.

- (ii) We attempt to approximate variables that are easier for the deep learning algorithm to train, which typically means approximating variables that are smooth, bounded functions of the aggregate states. To understand this, observe that it is easier for the Neural network to approximate $\xi_h = \partial_a V_h$ than V_h because it is easier to impose concavity on the derivative than the function. It is even easier to approximate the consumption-to-wealth ratio $\eta_h(\omega_i)$ and then reconstruct $\xi_h(\omega_i) = (\eta_h(\omega_i)\omega_i A)^{-\gamma}$ because $\eta_h(\omega_i)$ is typically bounded and so the explosive curvature in the SDF is encoded analytically. We illustrate this in Appendix E.3 by comparing training for different functional forms.
- (iii) We impose the market clearing conditions explicitly whenever possible because we find that allowing the deep learning algorithm to violate market clearing during the training process leads to instability. We illustrate this in Appendix E.4 by comparing training with and without asset market clearing explicitly imposed.

3.3 Part (c) Training Algorithm

We outline the key steps for training the neural networks in Algorithm 1 below and provide additional details in Algorithm 3 in the Appendix. Given the current guesses of the neural network equilibrium function approximations, we sample states, calculate the equilibrium at those states, compute the loss at those states with (3.2), and then update the parameters to decrease the loss.

Each neural network is a fully-connected feed-forward type and has 2 hidden layers, with 256 neurons in each layer. We train using an ADAM optimizer with a learning rate scheduler (ReduceLROnPlateau) for a maximum of 200k iterations. Every 50 epochs, the algorithm evaluates the validation loss on a test dataset. If the validation loss improves, it saves a checkpoint and resets the patience counter. Otherwise, it increments the counter. Training terminates either when the patience counter exceeds 15,000 validation checks without improvement (early stopping) or when the variance of validation loss is lower than a threshold. We use Latin hypercube sampling and make sure that each agent's wealth share sample ranges from 0.0 to 0.5.

Algorithm 1: Neural Network Training with Validation and Adaptive Learning

```
1: Initialize neural network objects for the functions in (3.1) with parameters  $\Theta$ 
2: Set patience counter  $p \leftarrow 0$  and validation loss  $\mathcal{L}_{\text{val}}^{\text{best}} \leftarrow \infty$ 
3: for  $n_{\text{epoch}} < N_{\text{max-epoch}}$  do
4:   Sample  $N$  new training points:  $(\hat{\mathbf{S}}^n = (z^n, K^n, (\omega_i)_{i \leq I}^n, \Omega_b^n, \Omega_f^n))_{n=1}^N$ .
5:   Use current NN approximations to calculate equilibrium at each  $\hat{\mathbf{S}}^n$  with Appendix C.4
6:   Construct the loss as  $\hat{\mathcal{L}}(\hat{\mathbf{S}}^n)$  using equation (3.2).
7:   Update parameters using ADAM optimizer to minimize loss.
8:   if  $t \bmod 50 = 0$  then
9:     Evaluate validation loss  $\mathcal{L}_{\text{val}}$  on test dataset; scheduler step with  $\mathcal{L}_{\text{val}}$ 
10:    if  $\mathcal{L}_{\text{val}} < \mathcal{L}_{\text{val}}^{\text{best}}$  then save checkpoint;  $\mathcal{L}_{\text{val}}^{\text{best}} \leftarrow \mathcal{L}_{\text{val}}$ ;  $p \leftarrow 0$  else  $p \leftarrow p + 1$ 
11:    if  $p \geq P_{\text{max}}$  then break {early stopping}
12:  end if
13: end for
```

Figure 9 in the Appendix presents the L-2 loss from the quantitative model over iterations. The loss decreases over time, although not monotonically, due to the stochastic nature of the learning process. After 60,000 iterations, the average Euler equation training loss and validation loss (MSE) are 8.5×10^{-5} and 1.2×10^{-4} , respectively, in marginal utility units.⁹ The right panel figure reveals that different components in the loss function converge at different speeds. The algorithm focuses on the HJB equation first, followed by the consistency and portfolio choice conditions.

Discussion: The algorithm includes a few key features:

- (i) We evaluate the progress of the algorithm on an “out-of-sample” validation dataset and save models that reduce loss on the validation set. We do this to try and limit the possibility of overfitting. This echoes standard practice for machine learning.
- (ii) To test the efficacy of our numerical algorithm in a controlled environment, in the Online Appendix G we use our algorithm to solve three simpler canonical models from the macro-finance literature that can also be solved with traditional methods: a multi-agent Lucas asset pricing model, [Basak and Cuoco \(1998\)](#), and [Brunnermeier and Sannikov \(2014\)](#). In each case, our L2 out-of-sample validation errors are in the order of 10^{-9} and our L2 difference to traditional finite difference techniques are in the order of 10^{-7} . This is a similar accuracy

⁹For demonstration, Figure 9 displays losses until 200k epochs. The early stopping criterion gets triggered after 60k epochs since there is little improvement in the validation loss when trained longer.

to what [Azinovic et al. \(2022\)](#) and [Duarte et al. \(2024\)](#) achieve when they illustrate their techniques on commonly used models from the macroeconomics and finance literatures. We summarize our results in Table 2.

Method	Training Error	Validation Error	Difference
20-Agent Complete markets	1.8×10^{-9}	2.6×10^{-9}	8.0×10^{-9}
Basak and Cuoco (1998)	3.0×10^{-9}	3.2×10^{-9}	1.6×10^{-7}
Brunnermeier and Sannikov (2014)	4.6×10^{-8}	7.0×10^{-8}	8.8×10^{-7}

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 L-2.

4 Calibration

For standard macro-finance parameters, we externally calibrate using accepted values in the literature. For the preference and regulatory parameters governing the household and financial intermediary portfolio problems, we internally calibrate using simulated method of moments to match target moments. Our data sample spans 2010Q1-2025Q2. Table 3 reports the complete set of parameters and the set of targeted moments. We classify the calibration targets into three groups: (i) macroeconomic, (ii) asset pricing, and (iii) cross-sectional moments. For each group, we discuss how successfully we fit the targeted moments and how robust our calibration is to matching “out-of-sample” untargeted moments.

Macroeconomic parameters: The mean-reversion of TFP (β_z) is set to 0.3, following [Gertler, Kiyotaki and Prestipino \(2019\)](#), and the volatility of TFP (σ_z) is set to target a 3.5% output growth volatility. The depreciation rate (δ) is chosen to match an annualized output growth rate of 2.6%, which is consistent with the real long-run output growth rate used in the literature. The investment friction parameter (κ) is calibrated to match the volatility of the private investment-to-capital ratio. The model also successfully generates an investment-to-capital ratio of 12.8%, close to the data, even though this is untargeted in the calibration.

Asset pricing parameters: The risk aversion parameters are internally calibrated to match spreads. We set household risk aversion to target an average equity risk premium ($r^k - r^d$ in the model)

<i>Targeted moments</i>	Parameter	Value	Data Target	Model
TFP mean reversion	β_z	0.30	Literature	-
TFP volatility	σ_z	0.03	Output growth vol = 0.034	0.032
Depreciation	δ	0.01	Output growth =0.025	0.021
Investment friction	κ	80	Investment vol =0.076	0.086
Discount rate (hh, fund, bank)	ρ	0.05	Literature	-
Household risk aversion	γ, Γ	0.3	Risk premium = 0.060	0.057
Bank risk aversion	γ_b	0.10	Capital Sharpe Ratio = 0.40	0.40
Fund risk aversion	γ_f	0.05	Pension return = 0.035	0.033
Capital constraint	ψ_k	1.5e-4	50th pctl. capital share =0.39	0.37
Death shock intensity (hh.)	λ_h	0.10	Wtd. avg. time to retirement =10	10
Death shock intensity (fund)	λ_f	0.20	Equity recapitalization	5yrs
Death shock intensity (bank)	λ_b	0.20	Equity recapitalization	5yrs
Transfer weight (hh.)	$\bar{\phi}_h$	0.03	10th pctl. capital share = 0.02	0.00
Transfer weight (fund and bank)	$\bar{\phi}_f = \bar{\phi}_b$	0.10	Fin. wealth/Total wealth = 0.39	0.37
Bond maturity	λ_m	0.10	Avg. Maturity of LT bonds	10 yrs
Bank capital ceiling	$\bar{\theta}_b^k$	0.9	Bank K/E=1.61	1.73
Bank capital tightening	$\bar{\psi}_b^k$	0.1	Implied regulatory cost = 0.040	0.032
Fund capital ceiling	$\bar{\theta}_f^k$	1.0	Fund K/E= 1.70	1.81
Bank capital tightening	$\bar{\psi}_f^k$	0.1	Implied regulatory cost = 0.040	0.054
<i>Untargeted moments</i>			Data	Model
Investment/capital rate (%)			14.0	12.8
Capital price volatility (%)			18.0	16.1
Gini coefficient (income)			0.40	0.47

Table 3: List of model parameters and calibration targets. All values are annualized. The time period is from 2010 Q1 to 2025Q2. The final column gives our model value for the data target.

equal to the sample average 6.0%. We set banker risk aversion to target an average capital Sharpe ratio $((r^k - r^d)/\sigma_{q^k})$ in the model) equal to the sample value of 0.40, where we compute the same value by estimating the risk premium using a factor model.¹⁰ Finally, we set fund risk aversion to target an average pension return to the fund $(r^n - r^d)$ in the model) equal to the literature value of 3.5%.¹¹ For all agents, our internally calibrated risk aversion parameters are less than one, which is closer to risk neutrality than is used in the frictionless consumption based asset pricing literature.

¹⁰Specifically, we regress equity market returns on dividend yield and estimate the Sharpe ratio from the fitted values.

¹¹Pension returns are difficult to compute and there is no clear literature consensus. We target a conservative estimate of 3.5% net of fees, which comes from (Mitchell, Poterba, Warshawsky and Brown, 1999).

This reflects that the participation and regulatory constraints generate significant curvature in the agent value function and so we do not need high risk aversion to match spreads. For all agents, the discount rate is set at 0.05, consistent with the literature (e.g., [Gertler et al. \(2019\)](#)).

The transfer weights and exit rates of the funds and banks $(\bar{\phi}_j, \lambda_j)_{j \in \{b, f\}}$ target the equity share of financial institutions (mapped to $(\theta_b^k + \theta_f^k)A/K$ in the model). We calculate the financial sector equity share as the fraction of U.S. corporate assets held by pension funds, insurance companies, and US-chartered banks using the Federal Reserve’s Z.1. Table L.223. Over 2010-2024, this target share averages 39% of total market value. We target the equity share of the total financial sector rather because it is unclear how to allocate equity holdings between pension funds, insurance companies, and banks.

We internally calibrate the regulatory parameters to target market outcomes rather than attempting to take them directly from the relevant regulations (e.g. Basel III, Insurance Capital Standard (ICS), and The Dodd-Frank Act). This is because our model effectively has a combined bank-firm sector and a combined pension-insurance sector so there is no literal mapping between our regulatory parameters and the regulations. First, the intermediary portfolio regulatory limits and recapitalization rates $(\bar{\theta}_j^k, \bar{\phi}_j)_{j \in \{b, f\}}$ are set to target risky asset to equity ratios in the bank and pension/insurance fund sectors (mapped to θ_b^k and θ_f^k in our model).¹² The risky-asset to equity ratio of bank holding companies, broker dealers, and non-financial institutions in the data are 3.77, 5.2, and 1.12, respectively.¹³ Since the “banking” sector in our model combines both financial and non-financial institutions that hold risky capital, we target the asset weighted average ratio of 1.61. For the pension fund sector, the empirical leverage ratio is 1.7.¹⁴ The resulting values of $(\bar{\theta}_j^k)_{j \in \{b, f\}}$ are shown in Table 3 in the baseline model column and will be varied in counterfactual experiments. Second, the soft penalty parameters $(\bar{\psi}_j)_{j \in \{b, f\}}$ capture the cost of violating risk-based capital requirements. They are chosen to target empirical implied regulatory costs. The model implied cost of regulation is given by the penalty functions Ψ_b^k, Ψ_f^k in equation (2.1). They are chosen to target 4% as [Koijen and Yogo \(2015\)](#) reports that insurers are willing to give up \$0.96 of value to obtain \$1 more in statutory capital.

The household participation constraint, transfers, and regulatory parameters are set to target the portfolio holdings. The household transfer weight is calibrated to match the risky asset share

¹²We set $\bar{\psi}_b^m = \bar{\psi}_f^m = 0$ and calibrate the capital penalty parameters. The model includes penalty on all assets for generality.

¹³Source: Federal Flow of Funds.

¹⁴Source: FRED database. We use ‘BOGZ1FL594090005Q’ for total pension fund assets, and ‘BOGZ1FL574190005Q’ for liabilities.

of the bottom 10th percentile household in the SCF. When working with the SCF, our empirical counterpart to our ‘capital stock’ is the aggregate of directly held mutual funds, publicly traded equity, bonds, and weighted owner-occupied housing. Because housing is only partially exposed to aggregate market risk, converting housing wealth into ‘risky capital’ requires an assumption on the housing market beta β^H . Empirical estimates of β^H vary significantly and lie in the range $[0.1, 0.9]$ in the post-2010 sample period (Piazzesi (2025)). We therefore set $\beta^H = 0.5$ as a benchmark midpoint for the calibration target. We set the household capital constraint parameter $\bar{\psi}_k$ to match the risky asset holdings of the median household in the SCF data, and then evaluate how well the model replicates moments across the distribution.

The exit rates are externally calibrated. The household retirement intensity λ_h reflects the wealth-weighted average time to retirement in the population and is set to 0.1. We choose this so we can match the fraction of aggregate wealth in pensions. We set $\beta = 0.5$ so that flow and terminal utility carry the same weight.

As an out-of-sample test, we study the risk exposure of the financial sector. The model directionally matches the exposure of fund’s net worth to capital returns, even though this is not directly targeted. Formally, in Table 4 we report factor coefficients from Koijen and Yogo (2022), who regress publicly traded annuity insurer equity returns against aggregate stock market and 10-year Treasury bond returns, and our model counterpart which regresses the fund wealth share against capital market and bond returns. Our model coefficients are broadly consistent with Koijen and Yogo (2022) estimate, which show that insurers’ stock returns have a positive beta with respect to stock and a negative beta with respect to bond returns. While the interest rate is endogenous in our model, the quantitative exercise indicates that an unexpected increase in the rate does decrease the fund’s wealth, in line with Koijen and Yogo (2022). The period where the long term bond return coefficient does not align is the peak of the Quantitative Easing (QE) during 2010-17, which suggests QE brought some additional forces not currently in our model.

Cross-sectional moments: Our model generates rich cross-sectional moments that are broadly consistent with the data, even though we only explicitly targeted the portfolio shares of the bottom 10th percentile and the median household. Figure 1 shows households’ risky capital, annuity, and risk-free asset shares as a fraction of total wealth in both the data and the model. Under our calibration, wealthy households choose relatively more capital, poorer households choose relatively more deposits, and all households allocate a similar fraction of their wealth to pensions. Given the difficulty of finding a single point estimate for the risky portfolio share, we assess the fit across the distribution by comparing to an empirically plausible interval computed under two extreme spec-

Factor	Data (1999-2017)	Data (2010-2017)	Model
Stock market return	1.36 (0.19)	1.11 (0.08)	0.82 (0.00)
Long term bond return	-0.01 (0.32)	-1.28 (0.43)	-0.03 (0.66)
Observations	228	96	299

Table 4: Risk exposure of fund sector. The table reports betas from a factor regression of fund equity returns on stock returns and long-term bond returns. Data values are taken from [Kojien and Yogo \(2022\)](#), and corresponds to the period 2010-2017. Heteroskedasticity adjusted standard errors are given in parenthesis.




Figure 1: The top panel figures present the portfolio holdings of households across 10 percentile buckets from the SCF data and model, respectively. The bottom left figure plots the impulse response of wealth distribution to risk premium ($\beta_{p,h}$) obtained from the regression $\log(W_{p,t+h}/W_{p,t}) = \alpha_{p,h} + \beta_{p,h}\hat{f}_t + \epsilon_{p,t+h}$, where p denotes top 10 percent. The data for wealth percentiles come from [Saez and Zucman \(2016\)](#), and risk premium is estimated using a factor model. The bottom right figure plots the model counterpart.

ifications: $\beta^H = 0$ (housing as a safe-asset), and $\beta^H = 1$ (housing as a risky asset). Our model implied risky capital share broadly lies inside this interval. In particular, our model says the bottom decile holds almost no wealth in risky assets, which exactly fits the data, and the top decile holds approximately half their wealth in risky assets, which is close to the upper bound with $\beta^H = 1$.

We also assess the model's performance in matching the exposure of household wealth to the financial sector. Empirically, we measure changes in the wealth distribution following shocks to the net worth of U.S. insurance funds. Shocks are estimated as innovations in the autoregression $\hat{a}_t = \beta_0 + \beta_1 \hat{a}_{t-1} + u_t$, normalized by lagged equity value $\hat{f} := \frac{u_t}{\hat{a}_{t-1}}$. We then apply local projections in the spirit of [Jordà \(2005\)](#), regressing shocks to fund net worth on the wealth share of households at different percentiles. Specifically, we run the regression $\log \left(\frac{W_{p,t+h}}{W_{p,t}} \right) = \alpha_{p,h} + \beta_{p,h} \hat{f}_t + \epsilon_{p,t+h}$ for horizons $h = 1$ to 25 quarters where $w_{p,t+h}$ denotes household wealth at percentile p at horizon h , where p denotes the top 10 percent wealth. The household wealth distribution is taken from [Saez and Zucman \(2016\)](#), and the fund wealth share series is constructed from the Financial Accounts of the United States (FRB) following [Kojien and Yogo \(2021\)](#). Figure 1 compares the empirical $\beta_{p,h}$ with the model-implied ones. We can see that our model produces a similar hump shaped response of household wealth to the shocks, even though it is untargeted.

5 Household Participation, Asset Prices, & Inequality

A key feature of our model is that households face frictions that restrict portfolio adjustment: a penalty on capital market participation and an inability to directly provide pensions to each other. Section 2 showed that these frictions appear as distortions in the household portfolio choice FOCs (equations (2.6) and (2.7)) that disrupt the standard risk-return tradeoff. Our calibration in Section 4 showed that these portfolio choice distortions lead to wealthy households choosing relatively more capital, poorer households choosing relatively more deposits, and all households allocating a similar fraction of their wealth to pensions (Figure 1). In this section, we study the impact of these household portfolio choice frictions on the broader economy. We start by exploring how the cross sectional variation in household capital market participation impacts inequality and social mobility. We then study how the relative elasticity of household demand for bank deposits and fund pensions impacts financial intermediary borrowing costs and risk sharing. Finally, we study how aggregate portfolio allocations impact investment and output.

5.1 Inequality

In this subsection, we study how the household frictions influence the equilibrium household wealth distribution. To understand this, we start by considering how asset returns influence relative agent wealth shares. Let the share of household wealth owned by household i be denoted by $\omega_{i,h,t} := a_{i,t}/A_{h,t} = \omega_{i,t}/\Omega_{h,t}$. From Theorem 2, the difference between the geometric wealth share growth of any two young households j and i can be decomposed using equation (5.1) below:

$$d(\omega_{j,h,t} - \omega_{i,h,t}) = \left[(\theta_{j,t}^k - \theta_{i,t}^k)(r_t^k - r_t^d - \sigma_{A_{h,t}} \sigma_{q,t}^k) + (\theta_{j,t}^n - \theta_{i,t}^n)(r_t^n - r_t^d - \sigma_{A_{h,t}} \sigma_{q,t}^n) - (\eta_{j,t} - \eta_{i,t}) \right] dt + \left[(\theta_{j,t}^k - \theta_{i,t}^k) \sigma_{q,t}^k + (\theta_{j,t}^n - \theta_{i,t}^n) \sigma_{q,t}^n \right] dW_t \quad (5.1)$$

where we have used time subscripts rather than explicit dependence on the states and $\sigma_{A_{h,t}}$ is the risk exposure of aggregate household wealth $\sigma_{A_{h,t}} = \frac{K_{h,t}}{A_{h,t}} \sigma_{q,t}^k + \frac{N_{h,t}}{A_{h,t}} \sigma_{q,t}^n$.

The first terms in the drift and volatility in equation (5.1) capture how capital market participation constraints and risk aversion impact the total return that different agents can earn. If agent j is richer, $a_{j,t} > a_{i,t}$, then under our calibration they hold more wealth in capital, $\theta_{j,t}^k > \theta_{i,t}^k$, and so have more exposure to the risk adjusted premium in the economy, $r_t^k - r_t^d - \sigma_{A_{h,t}} \sigma_{q,t}^k$, and to the risk in the economy, $\sigma_{q,t}^k$. This means that, on average, richer households gain wealth share compared to the poorer agents who are unwilling to pay the cost to participate in the capital market. The literature has sometimes referred to this as a “scaling” effect: wealthier agents have access to better investment opportunities and so gain wealth more quickly. Our model endogenizes the strength of scaling effect by determining the spread $r_t^k - r_t^d$ in general equilibrium. The second terms in the drift and volatility of equation (5.1) capture how pension costs impact inequality. In our model, agents pay an insurance fee to the pension system since $r_{f,t}^n - r_t^d < 0$ so households with relatively more pension exposure end up losing wealth. The third and final term in the drift of (5.1) capture how differences in marginal propensity to consume out of wealth impact wealth accumulation.

The left subplot in Figure 2 shows the equilibrium decomposition of the drift in equation (5.1) between agents at the 25th and 75th percentiles at the ergodic distribution for our baseline economy (the black bars) and for an economy with higher participation costs (the dotted bars). Evidently, the first term dominates the decomposition because pension holdings are relatively similar across the wealth distribution. This suggests that in equilibrium the scaling effect is both strong and largely driven by the capital market. We can also see that an increase in participation costs amplifies these effects because it makes poorer household even less able to participate in capital markets. The net

change in the left plot is positive, which indicates that, on average, asset returns lead to rich and poor agents drifting apart. Ultimately, around the ergodic mean density this scaling effect is offset by the redistribution from the exit and entry process.¹⁵




Figure 2: The left panel presents the inequality decomposition between the 75th and 25th percentile household. Here $\Delta \text{capital} = (\theta_{j,t}^k - \theta_{i,t}^k)(r_t^k - r_t^d - \sigma_{A_h,t}\sigma_{q,t}^k)$ and $\Delta \text{pension} = (\theta_{j,t}^n - \theta_{i,t}^n)(r_t^n - r_t^d - \sigma_{A_h,t}\sigma_{q,t}^n)$. The center and right panel plots present capital portfolio choice difference between 75th and 25th percentile household for varying Sharpe ratios. We set $\bar{\psi}_k = 0.15$ and $\bar{\psi}_k = 0.05$ in the high cost and low cost economies, respectively.

To understand the different mechanisms behind the decomposition (5.1), we study the household's capital choice in their optimization problem. Rearranging the capital choice FOC gives:

$$\frac{r_t^k - r_t^d}{\sigma_{q,t}^k} = -\sigma_{\xi_{h,t}} - \frac{\bar{\psi}_k}{\omega_{h,t}^{1+\alpha} \sigma_{q,t}^k} \theta_{h,t}^k$$

for a household with wealth share $\omega_{h,t}$ where again we have dropped the explicit dependence on the states. As the term $\bar{\psi}_k \theta_{h,t}^k / (\omega_{h,t}^\alpha \sigma_{q,t}^k)$ becomes insignificant relative to the other terms, this portfolio choice becomes the standard trade-off between risk and return. So, as we saw in Figure 1, for wealthy agents with high $\omega_{h,t}$ this is the standard Merton portfolio choice FOC while for poorer agents with low $\omega_{h,t}$ this is a distorted FOC in which the participation constraint dominates and households choose low θ_h^k . This means that a decrease in the participation cost, $\downarrow \bar{\psi}_k$, closes the difference in household portfolios and so, holding all else equal, reduces the first term in (5.1). It also means that an increase in the Sharpe ratio $(r^k - r^d)/\sigma_{q,t}^k$ has a heterogeneous impact on portfolio choice across the wealth distribution. For richer households who are already unconstrained,

¹⁵On average, redistribution from the exit and entry process adjusts the wealth share gap $\omega_j - \omega_i$ at rate $\lambda_h \bar{\phi}_h \omega_h (\omega_{j,t}^{-1} - \omega_{i,t}^{-1})$ so the more $\omega_j > \omega_{i,t}$, the more aggressively the death and rebirth brings agents back together.

the risk-return tradeoff becomes more attractive and so they allocate more wealth to capital. For poorer households, there is an additional effect because an increase in the Sharpe ratio offsets the participation cost and so increases participation in the capital market. That is, there is an “extensive margin” effect where high Sharpe ratios encourage more households to overcome the participation barriers and enter the capital market. In this sense, across the wealth distribution, the portfolio responsiveness to spread changes is governed by the participation constraint.

The second two plots in Figure 2 illustrate how these forces in the capital market impact inequality in partial equilibrium for different spreads. The middle subplot shows the difference between household capital choices, $\theta_{j,t}^k - \theta_{i,t}^k$, for increasing values of the Sharpe ratio and right subplot shows the first term in the drift in equation (5.1), $\Delta\mu_\eta^k = (\theta_{j,t}^k - \theta_{i,t}^k)(r_t^k - r_t^d - \sigma_{A_h}\sigma_{q,t}^k)$, for increasing risk premia, where j and i refer to the 25% and 75% percentiles of household wealth distribution. The solid line on each plot depicts an economy with low participation costs while the dashed line depicts an economy with high participation costs. Evidently, when the Sharpe ratio is low, the richer households hold relatively more capital. For small increases in the Sharpe ratio, this leads to a large increase in inequality because it is the richer households who benefit the most from higher capital profitability. However, if the Sharpe ratio becomes sufficiently high, then low wealth households enter the capital market and the difference in portfolio choices erodes. Ultimately, this means that that inequality stops increasing. The size of the Sharpe ratio increase that is required in order to close the portfolio choice gap depends on the tightness of the household participation constraint. For the solid line depicting the economy with low participation constraints, $\Delta\mu_\eta^k$ starts decreasing around a risk premium of 5% and Sharpe ratio of 0.45 while, for the dashed line depicting the economy with a higher participation constraint, $\Delta\mu_\eta^k$ is still increasing at an 8% risk premium and Sharpe ratio of 0.65. We summarize our observations in the following takeaway.

Takeaway 1: *Shocks or parameter changes that lead to higher Sharpe ratios both allow richer households to earn higher risk adjusted returns and also encourage poorer households to enter the market. The impact on inequality is ambiguous.*

5.2 Social Mobility

So far, we have analyzed how participation costs and household portfolio choices impact the difference in household wealth shares at different quantiles. Now, we focus on how different households within a quantile move around the distribution, which we refer to as “social mobility”. Formally, we let $s_t(q, q')$ denote the transition probability that a household in quartile q moves to quartile

q' over a time horizon t years. Households move up the wealth distribution on average as they live longer because they accumulate wealth while the richer households gradually exit. The rate at which they do this depends upon their access to the capital market and the risk premium.

Figure 3 shows the average equilibrium social mobility that emerges in our calibrated model. The first two panels show $s_5(q, q')$ and $s_{20}(q, q')$ between the four wealth quartiles, bottom 25%, 25-50%, 50-75%, and above 75%, while the second two panels show the change in $s_5(q, q')$ and $s_{20}(q, q')$ when the participation cost is increased. Evidently, a household's position in the wealth distribution is persistent with the probability of households staying in the wealth bucket they start from ranging from 0.57 to 0.88. In particular, the transition probability is highest for the top wealth bucket, followed by the bottom wealth bucket. High persistence for the bottom wealth households is due to the capital participation cost that becomes binding, preventing them from growing their wealth by earning risk premium relative to the other groups. While middle income households (i.e., 25-50% and 50-75%) have less persistence, the probability of moving to the upper wealth quartile is lower than that of moving to the lower wealth quartile, indicating that middle wealth households have a higher chance of drifting downward in the wealth distribution. This is broadly consistent with the data even though we don't explicitly target any mobility moments in our calibration.¹⁶ We can also see that persistence is higher for all wealth groups in the high participation cost economy compared to the baseline economy, particularly for the middle two wealth buckets (i.e., 25-50% and 50-75%). This leads to the takeaway:

Takeaway 2: *Higher participation costs reduce social mobility.*




Figure 3: The first two panels present the transition probabilities representing the social mobility across four wealth share household quartiles after $t = 5$ years and $t = 20$ years from the stochastic steady state. The right two panels present the difference in transition probabilities between the baseline ($\bar{\psi}_k = 1.50e^{-4}$) and high-cost economies ($\psi_k = 1.55e^{-4}$).

¹⁶Table 15 in the Internet Appendix H provides the comparison between the model and data moments.

5.3 Financier Liability Costs

In this subsection, we explore how the household frictions impact financier liability issuance costs. Absent regulation, the difference between the bankers and fund managers is that they have different liability structures that provide different services to the households: bankers issue short-term deposits while fund managers issue tradable long-term pensions. A stylized version of their balance sheets is shown in Table 5 (for the banker) and Table 6 (for the fund manager). Here, we depict bonds on the asset side of the balance sheet but, in principle, both banks and funds can take a positive or a negative position in the bond market.

Assets		Liabilities	
Capital	$q_t^k k_{b,t}$	Deposits	$-d_{b,t}$
Bonds	$q_t^m m_{b,t}$	Equity	$-a_{b,t}$

Table 5: Banker Balance Sheet

Assets		Liabilities	
Capital	$q_t^k k_{f,t}$	Pensions	$-q_t^n n_{f,t}$
Bonds	$q_t^m m_{f,t}$	Equity	$-a_{f,t}$

Table 6: Fund Manager Balance Sheet

There are two key differences between the deposits issued by the bankers and the pensions issued by the fund managers: (i) household deposit demand is more sensitive to return changes than pension demand and (ii) pensions provide a natural hedge against business cycle risk. To understand the first difference, the left subplots of Figure 4 shows household deposit and pension demand for different spreads. We can see that the mean household's choice of θ_h^d is approximately linear in the spread $r^d - r^k$ so the banker can expand its liabilities at linear cost. By contrast, the mean household's choice of θ_h^n becomes unresponsive to the spread $r_f^n - r^k$ as the spread becomes large and so it become very costly for the fund manager to expand its liabilities. That is, pension demand hits a "satiation point" and becomes relatively unresponsive to return and risk, similar to the preferred habitat demand.

To understand the second difference, consider the risk exposure of banker and fund manager liabilities. Bankers owe short-term deposits so the market value of their liabilities is unaffected by productivity decreases: they have to repay the full amount next period regardless of the state of the economy. By contrast, fund managers owe long-term tradable pensions that decrease in price during recessions (as do the prices of all long-term assets because goods are scarce and the marginal value of consumption is high) so the market value of their liabilities decreases in recessions. This means that, all else equal, funds gain wealth relative to banks during recessions. We can also see the different financial intermediary risk exposure mathematically by examining the difference in

the evolution of fund and banking wealth share growth rates:

$$d(\Omega_{f,t} - \Omega_{b,t}) = \left[\sum_{l \in \{k,m\}} \left(\theta_f^l (r_t^l - r_{f,t}^n - \sigma_{A,t}(\sigma_{q,t}^l - \sigma_{q,t}^n)) - \theta_b^l (r_t^l - r_t^d - \sigma_{A,t}\sigma_{q,t}^l) \right) - (\eta_{f,t} - \eta_{b,t}) \right] dt + \sum_{l \in \{k,m\}} \left[\theta_f^l (\sigma_{q,t}^l - \sigma_{q,t}^n) - \theta_b^l \sigma_{q,t}^l \right] dW_t$$

Here we can see that the fund manager's risk exposure is $\sum_{l \in \{k,m\}} \theta_f^l (\sigma_{q,t}^l - \sigma_{q,t}^n)$ whereas the bankers' risk exposure is $\sum_{l \in \{k,m\}} \theta_b^l \sigma_{q,t}^l$, which implies that the fund's liabilities partially offset their asset risk exposure whereas the bank's liabilities do not.

For these reasons, the funds are potentially a natural “backstop” for the banking sector in an economy so long as the households have little participation in the financial sector and banks and funds are insuring each other. This is illustrated in the right panel of Figure 4, which depicts banker and fund manager portfolio choices for different values of the household participation cost. In our baseline model, the fund manager takes a positive position in the bond market whereas the banker takes a negative position. In other words, the fund managers lends to the bankers to help the bankers offset their capital risk exposure. That is, they indirectly “insure” the banking sector through the bond market.

Takeaway 3: *Fund managers are less able to scale up their balance sheets but more able to bear risk because pension demand is more stable.*




Figure 4: Sector level: The left panel presents the deposit-capital and annuity-capital spread against household deposit and annuity holdings, respectively. The right panel presents the bank and fund capital holdings across the baseline and high participation cost economy, respectively.

5.4 Aggregate Real Outcomes

We close this section by studying how the household and financier choices from the previous two subsections translate to real investment. The real economy is connected to the capital price through the investment equation $\iota = (\Phi')^{-1}(1/q^k(s))$. The capital price is, in turn, pinned down through the capital market clearing condition $A \left(\int \theta_h^k \omega(\omega) g_h(\omega) d\omega + \theta_f^k \Omega_f + \theta_b^k \Omega_b \right) = q^k K$, which illustrates how shifts in average wealth shares and portfolio choices impact investment in the economy. If wealth moves from agents with high θ^k to agents with low θ^k , then the price of capital has to fall (or equivalently the return on capital has to rise) in order for the capital market to clear. This leads to lower investment and growth.

A first implication is that reallocation of capital to the household sector decreases growth in the economy because θ_f^k and θ_b^k are greater than $\theta_h(\omega)$ for households at all wealth share ω . This is illustrated in Figure 5, which shows output growth, investment, and spreads as a function of the total financial sector wealth share. Evidently, growth is higher when the financial sector holds more capital directly.

A second implication is that inequality and investment are interconnected. Consider two economies with the same overall wealth in the household sector Ω_h but with different dispersion across the households. If wealth in the household sector is evenly dispersed, then average θ_h is low and capital prices fall. By contrast, if the wealth in the household sector is very unevenly distributed, then average θ_h is high and capital prices fall by less. In this sense, wealth equality and high growth are inversely related, which plays key role in the regulatory counterfactual experiments in the next section. We can see this in Figure 5 by comparing the solid line (for a more equal economy) and the dashed line (for a less equal economy). Evidently, for the same overall household wealth, the more unequal economy delivers higher investment and growth.

Takeaway 4: *Output growth decreases when the household sector holds more capital directly. The decrease is smaller when the household distribution is more unequal.*

6 Revisiting Risk Based Capital Requirements

In Section 4, we calibrated our baseline economy using data from 2010Q1-2024Q4 and interpreted the estimated regulatory parameters as implicitly characterizing the restrictions financial intermediaries face under the current regulations (e.g. Basel III, Insurance Capital Standard (ICS), and The Dodd-Frank Act). In this section, we study the macroeconomic impact of alternative regulatory pa-




Figure 5: The top panel presents the output growth rate and investment rate vs. intermediary wealth shares. The bottom panel presents the pension premiums. In all plots, the solid (dashed) line corresponds to the economy with equal (unequal) household wealth distribution. All households have the same wealth share in the ‘equal’ case. Households have a dispersed wealth distribution in the ‘unequal’ case.

parameters that place different restrictions on bank and fund capital holding. Previous macro-finance work has focused on how bank capital requirements can lead to a high-level regulatory tradeoff between aggregate growth and aggregate stability. Our model allows us to extend this analysis to consider the heterogeneous impact of macro-prudential regulation across different financial intermediaries and households. We show that policy makers now potentially face an additional tradeoff with inequality.

6.1 Counterfactual Macroprudential Policies

Our baseline calibration finds that $\bar{\theta}_b^k$ is smaller than $\bar{\theta}_f^k$, which implies that current regulations restrict bank participation in capital markets more than fund participation. This essentially means that the current regulation allows funds to hold a greater share of high volatility assets than banks. This is consistent with a common interpretation that bank regulation (e.g. Basel III) penalizes banks holding volatile assets while non-bank regulation (e.g. the Insurance Capital Standard and other regulations on the pension/insurance sector) penalize long-term insolvency risk rather than exposure to short-term market volatility.

In this section we study the macroeconomic implications of changing the regulatory parameters $\bar{\theta}_f^k$ and $\bar{\theta}_b^k$. In our first counterfactual experiment, we tighten the maximum costless capital exposure

on the banking sector by reducing the parameter $\bar{\theta}_b^k$ from 0.9 to 0.5 so the bank risky exposure goes from 1.73 to 1.25. We refer to this as the “bank restricted economy” and interpret the decrease in $\bar{\theta}_b^k$ as reflecting a higher risk weight on volatile assets (i.e. a higher risk on non-government debt assets). This policy experiment is a stylized representation of proposals from critics concerned that banks are currently circumventing key areas of the Basel III regulation by holding too many risky assets (e.g. [Admati and Hellwig \(2013\)](#), [Acharya, Engle and Pierret \(2014\)](#)) and is similar to counterfactual policy experiments in the literature.

In our second counterfactual experiment, we tighten restrictions on the fund sector by decreasing the parameter $\bar{\theta}_f^k$ to the same level as in the banking sector. We refer to this as the “fund restricted economy” and interpret the decrease in $\bar{\theta}_f^k$ as reversing the implicit regulatory advantage funds have when holding assets with price volatility. This reflects policy maker concerns that bank risk taking has been tightly restricted since the Global Financial Crisis and so non-bank financial intermediaries—including pension and insurance funds—have absorbed a larger share of long-duration risk assets (as documented by [Koijen and Yogo \(2022\)](#), [Begenau et al. \(2024\)](#)).

6.2 Long Run Impact: A Growth-Stability-Inequality Tradeoff

We start by studying the long run impact of macro-prudential regulation. Table 7 reports the long-run average outcomes under the different counterfactual policy regimes. The first group shows two classic macro-finance summary statistics: the ergodic mean output growth rate in each regime and the ergodic mean volatility of household wealth share. The second group shows a collection of different measures of inequality. The third group shows ergodic mean equilibrium spreads and prices. The final group shows ergodic mean sector wealth shares and portfolios.

To understand the table, we need to consider how the tighter portfolio restrictions impact the macroeconomy. Decreasing $\bar{\theta}^j$ for financial intermediary j forces that intermediary to decrease capital demand. In general equilibrium, in a frictionless economy, the other financial intermediary sector and the households can respond by increasing demand for capital. However, in our environment the other agents face frictions that limit their capacity to offset the impact of the regulation. The other intermediary sector is restricted by their own regulatory constraints and their cost of issuing liabilities to expand their balance sheet. The households are restricted by their participation constraints. We make a collection of observations about how these forces play out in general equilibrium for our counterfactual experiments.

Observation 1: Tighter banking sector restrictions lead to a decrease in banker and fund manager capital holdings and an increase in household capital holdings. This leads to higher stability

but lower growth. By contrast, tighter fund sector restrictions lead to a decrease in fund manager capital holdings and an increase in both banker and household capital holdings. This leads to higher stability and higher growth. These asymmetries arise from the different liability structure of the bankers and the fund managers discussed in Section 5.3. When the bankers are restricted from holding capital, the fund managers are less willing and able to issue additional pensions and purchase capital because household demand becomes inelastic to spreads and so it is costly for the fund managers to expand when the banking sector contracts (see the second subplot in Figure 4). Instead, it is households who primarily respond and increase their capital holdings. By contrast, when the fund managers are restricted from holding capital, the bankers respond by increasing capital holdings relative to the households. This is because household demand for deposits is relatively elastic and so the bank can expand at low cost (see the first subplot of Figure 4). In this case, household participation in the capital market changes little.

In both counterfactual experiments, volatility decreases because the overall financial sector decreases leverage, as has been documented in many macro-finance papers. The differential impacts on investment and output growth come from the different portfolio responses. As we saw in Section 5.4, when bankers are more restricted and households end up holding additional capital, then capital prices must fall to compensate households for entering the capital market and so investment and output fall. By contrast, when fund managers are restricted, it is the bankers who increase capital holdings, which drives up the price of capital leading to higher investment and output growth.

Observation 2: Tighter financial sector restrictions lead to higher Sharpe ratios and potentially higher inequality, especially among poorer agents. In both counterfactual experiments, the tighter regulations end up increasing the Sharpe ratio in the economy. We can understand this mathematically through the portfolio FOCs (equations (2.6-2.9) in Theorem 1). Tightening a particular financial sector's regulatory constraint increases the wedge in that intermediary's FOC. In addition, the wedges in the other FOC equations also end up binding more tightly if those agents respond by increasing their capital holdings. The result is that spreads increase while volatility decreases leading to a higher Sharpe ratio.

As discussed in Section 5.1, how this higher Sharpe ratio impacts inequality depends upon the relative size of the intensive margin effect (i.e. how much higher the return is for households with large capital holdings) and the extensive margin effect (i.e. to what extent households increase capital holdings). How these effects play out across the household distribution are subtle and so we provide a collection of different measures. We start in Table 7 by reporting the rate at which the

top 80-th percentile gain wealth share relative to the 50-th percentile before any taxes, transfers, and deaths (equation (5.1) at the ergodic mean). We then report the analogous number for the 20th-50th percentile gap. These differential wealth share growth rates reflect the “primary” increase in inequality coming from the economics in the model. By this metric, we can see that inequality increases more quickly in both the fund and bank restricted economies, with a larger increase in the fund restricted economy.

To better understand this, Figure 6 plots the individual Sharpe ratios faced by agents across the wealth distribution (the LHS subplot) and the difference between the change in wealth shares for different percentiles relative to the median (the RHS subplot). By both metrics, we can see that, even after risk adjustments, inequality increases more quickly among poorer households and less quickly among richer households (with the the wealth shares essentially not changing above the median in the baseline economy). This is because for richer households, all agents become unconstrained and so their portfolios converge. By contrast, poorer households are stuck out of the capital market and earn lower returns. The basic intuition is that increased regulation make financial intermediation difficult, which opens up spreads. Richer households can better take advantage of these spreads because they can better circumvent the regulations and invest directly into capital markets to earn high returns. Poorer household loose out because it is too costly for them to invest directly.

To illustrate how these forces play out dynamically, we also report the 80th-50th and 20th-50th percentile changes in wealth shares along the transition paths, measured as the difference in ω_h 10 years after the regulation change for cohorts that are alive at the time of the reform. By this metric, we also see an increase in inequality under the bank and fund restricted economies.

Observation 3: No counterfactual experiments deliver higher output growth, less volatility, and lower inequality. Bringing the observations together, our analysis show a complicated third dimension to the standard macro-prudential trade-offs in the literature. The bank restricted economy delivers higher stability but lower growth and higher inequality. The fund restricted economy delivers higher stability and higher growth but also the highest inequality, particularly at the bottom end. This suggests that governments face a difficult trilemma style tradeoff between prioritizing financial stability and reducing inequality.

6.3 Business Cycle Dynamics

The previous subsection largely studied long-run ergodic outcomes. We now examine how the economy responds to recessions under different regulatory regimes. Figure 7 shows the impulse re-

	Baseline	Counterfactual: Bank-restricted	Counterfactual: Fund-restricted
<i>Growth-Stability</i>			
Output Growth (%)	1.9332	1.7423	1.9801
Wealth share Risk (hh.) (%)	0.1924	0.1211	0.1546
<i>Inequality</i>			
Wealth Share Growth(80th-50th pct)	0.0799	0.5901	0.4688
Wealth Share Growth(20th-50th pct)	-0.2733	-0.9954	-1.0270
Transition Wealth Share Growth(80th-50th pct)	0.2108	0.2644	3.0820
Transition Wealth Share Growth(20th-50th pct)	-0.2330	-0.2699	-2.5560
<i>Prices</i>			
Sharpe Ratio $(r^k - r^d)/\sigma_{q^k}$	0.3987	0.6042	0.5352
Govt. bond price q^m	0.4884	0.4440	0.4833
Pension price q^n	0.8813	0.8287	0.8727
<i>Sector level results</i>			
Fund K/E	1.8058	1.7721	1.4946
Bank K/E	1.7321	1.2521	1.7464
Household K/E (Ave)	0.3763	0.5240	0.4133
Household K/E (Top 1%)	0.5279	0.7445	0.5895
Wealth Share (hh.)	0.6442	0.6856	0.6287
Wealth Share (fund.)	0.1725	0.1585	0.1676
Wealth Share (bank.)	0.1833	0.1559	0.2037
Investment/Capital (%)	12.1610	11.1715	12.4147

Table 7: Equilibrium across different regulatory regimes. Wealth share risk (hh.) represents the average wealth share volatility ($\sigma_{\omega_h, z}$) across households. Fund K/E and Bank K/E represent the respective risky capital to equity ratios. Wealth share risk denotes $\sigma_{\omega_h, z}$, the loading on the TFP shock to the household wealth-share.

sponse functions to flow consumption, retirement consumption, output growth, investment, wealth shares, and asset prices to a deep recession under the different regulatory regimes. Evidently, restricting the financial intermediaries leads to a less deep decrease in output during the recessions. This reconciles with second line of Table 7, which shows that both regulated economies have greater stability across the business cycle. We can also see that the recession mitigation is most pronounced for the bank restricted economy because capital demand stays most stable and so capital prices fall the least. Finally, the impulse responses show that restricting the fund sector leads to a greater decrease in flow consumption but a more muted initial decrease in retirement consumption. This is because the risk premium increases more in the fund restricted economy,




Figure 6: The left panel presents the sharpe ratio on household wealth share at different percentiles across the three economies. The right panel presents the expected growth rate of household wealth share across the three economies.

households conduct more precautionary saving, and the cost of providing pension products does not increase as much.

6.4 Policy Variant Preferred Habitat Demand

Following [Vayanos and Vila \(2021\)](#) (henceforth VV), many researchers have studied and estimated environments with preferred habitat preferences. Our paper extends this tradition by integrating preferred habitat preferences into a standard macroeconomic environment. We close this section by making this precise. Specifically, we map our model into the VV environment and show that our FOCs can be interpreted as endogenizing the policy dependence of the inelastic demand “shifters” in VV. We then study how these shifters change across our counterfactual experiments.

To understand these connections formally, we first express our pension demand in an analogous form to the literature. VV has a long-term bond demand function:

$$\mathcal{U}'(C_t) = q_t^n \beta_t, \quad \beta_t = \theta_0 + \sum_{k=1}^K \theta_k \beta_{k,t} \quad (6.1)$$

where the $\{\beta_{k,t}\}_{k=1}^K$ are exogenous, mean reverting factors and the $\{\theta_k\}_{k=1}^K$ are exogenous factor loadings. Our asset pricing condition for pensions can be rearranged to get:

$$\mathcal{U}'(C_t) = \frac{1}{\lambda_h} \frac{q_t^n}{1 - q_t^n} \xi_h \left(r_t^d - r_t^n - \sigma_{\xi_h,t} \sigma_{q_t^n} \right) = \frac{1}{\lambda_h} \frac{q_t^n}{1 - q_t^n} \mathbb{E} \left[\int_0^\infty e^{-\rho t} \xi_{h,t} d c_{h,t} dt \right],$$

which implicitly characterizes our equilibrium demand function for pensions. Our demand function shares the same basic trade-off as the micro-foundation in the appendix of VV: households




Figure 7: The figure presents the impulse response of equilibrium quantities to a large negative TFP shock. The first row presents the shocked path, flow, and terminal consumption. The second row presents the output growth, investment rate, and capital price. The third row presents the wealth shares and capital return spread.

can invest an extra unit in the pension for marginal benefit $U'(C)$ (the LHS) but face the cost of shifting down their consumption path (the RHS). However, our formula differs in three ways: (i) we model a random horizon while VV has fixed horizon¹⁷; (ii) we have one good so the gain from purchasing one share of a pension right before retirement shock hits is $1/q_n - 1$, while in VV, there are two goods; (iii) in VV, they assume risk-neutral flow utility function $\xi_{h,t} = 1$; and (iv) in VV the RHS “shifter” $\beta_h = \xi_h (r^d - r^n - \sigma_{\xi_h} \sigma_{q^n}) / \lambda_h$ is treated as exogenous whereas in our model it comes from the portfolio choice problem of a risk averse agent.

To make this explicit, we take the linear approximation for our endogenous RHS shifter $\beta_h = \xi_h (r^d - r^n - \sigma_{\xi_h} \sigma_{q^n}) / \lambda_h$ and compare it to the exogenous linear factor model for β_h in VV. We do this by taking a first-order Taylor expansion of $\beta_h(\omega, \Omega_f, z)$ around the stochastic steady state

¹⁷To see this more clearly, we could further rewrite $1/\lambda_h$ as: $\int_0^\infty e^{-\lambda_h s} ds$.

Comparison	Test Statistic	p -value
Bank restricted vs Baseline	0.0660	3.28×10^{-4}
Fund restricted vs Baseline	0.0555	4.22×10^{-3}

Table 8: Kolmogorov–Smirnov test for equality of the β_h distribution under three counterfactor regimes. The Statistic is the maximum absolute difference between the empirical CDFs of the two samples. Reported p -values are two-sided.

(SSS) point $(\omega_{h,ss}, \dots, \omega_{h,ss}, \Omega_{f,ss}, z_{ss})$:

$$\beta_h(\omega, \Omega_f, z) \approx \beta_{h,0} + \sum_{i=1}^I \frac{\partial \beta_h}{\partial \omega_i} \Big|_{\text{SSS}} (\omega_i - \omega_{h,ss}) + \frac{\partial \beta_h}{\partial \Omega_f} \Big|_{\text{SSS}} (\Omega_f - \Omega_{f,ss}) + \frac{\partial \beta_h}{\partial z} \Big|_{\text{SSS}} (z - z_{ss}).$$

Define the centered factors $\beta_{k,t} = \omega_{k,t} - \omega_{h,ss}$ for $k = 1, \dots, I$, $\beta_{I+1,t} = \Omega_{f,t} - \Omega_{f,ss}$, $\beta_{I+2,t} = z_t - z_{ss}$. Then $K = I + 2$, and the loadings in VV formulation in equation (6.1) map to:

$$\theta_0 = \beta_{h,0}, \quad \theta_k = \frac{\partial \beta_h}{\partial \omega_k} \Big|_{\text{SSS}} \quad (k = 1, \dots, I), \quad \theta_{I+1} = \frac{\partial \beta_h}{\partial \Omega_f} \Big|_{\text{SSS}}, \quad \theta_{I+2} = \frac{\partial \beta_h}{\partial z} \Big|_{\text{SSS}}.$$

In Table 8, we compare the CDFs of the β_h distribution in our counterfactual regimes to the CDF in our baseline case and perform a statistical test for whether to reject the baseline β_h process when the financial regulatory regime changes. For both counterfactual experiments, we reject the baseline process for β_h , indicating we need to be able to predict the change in the β_h process in order to understand the general equilibrium impact of regulatory changes.

7 Conclusion

This paper has examined how household frictions interact with the constraints on the financial sector and how this ultimately affects household inequality. To do this, we constructed a novel heterogeneous-agent macro-finance model with households facing asset market participation constraints, banks providing deposits, funds providing insurance/pension products, and endogenous asset price volatility. We solved the model by developing a new deep learning methodology that enables global solutions in environments with long-term assets, endogenous aggregate volatility, and binding portfolio constraints. Quantitatively, we calibrated the model to the post financial crisis United States economy from 2010 to 2024 and showed that it matches key macroeconomic, asset pricing, and cross sectional portfolio moments. Counterfactual experiments highlight the

policy tradeoffs associated with alternative regulatory regimes. Tighter financial sector restrictions increase stability but at the expense of lower growth and/or higher inequality. This is because richer households are better able to take advantage of the higher spreads created by the regulations.

We believe the technical innovations in our paper open up new pathways for the macro-finance literature. Our model combines advances in intermediary asset pricing with macroeconomics and household finance. Our solution approach brings aggregate risk and endogenous volatility into the macroeconomic inequality literature. Our policy counterfactuals introduce household inequality into the macroprudential policy tradeoffs. This points the way to a more nuanced and sophisticated integration of finance and macroeconomics.

References

- Acharya, Viral V., Robert F. Engle, and Diane Pierret**, “Testing Macroprudential Stress Tests: The Risk of Regulatory Risk Weights,” *Journal of Monetary Economics*, 2014, 65, 36–53.
- Admati, Anat R. and Martin F. Hellwig**, *The Bankers’ New Clothes: What’s Wrong with Banking and What to Do about It*, Princeton, NJ: Princeton University Press, 2013.
- Azinovic, Marlon and Jan Zemlicka**, “Intergenerational consequences of rare disasters,” 2026.
- , **Luca Gaegauf, and Simon Scheidegger**, “Deep equilibrium nets,” *International Economic Review*, 2022, 63 (4), 1471–1525.
- Bach, Laurent, Laurent E. Calvet, and Paolo Sodini**, “Rich Pickings? Risk, Return, and Skill in Household Wealth,” *American Economic Review*, September 2020, 110 (9), 2703–2747.
- 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,” *SSRN Electronic Journal*, 1 2024.
- Begenau, Juliane**, “Capital Requirements, Risk Choice, and Liquidity Provision in a Business Cycle Model,” *Working Paper, Harvard*, 2016, (November 2013), 1–73.
- **and Tim Landvoigt**, “Financial Regulation in a Quantitative Model of The Modern Banking System,” *SSRN Electronic Journal*, 2017.
- , **Pauline Liang, and Emil Siriwardane**, “The Rise of Alternatives,” Technical Report Working Paper 25-016, Harvard Business School August 2024.
- Blanchard, Olivier J**, “Debt, deficits, and finite horizons,” *Journal of political economy*, 1985, 93 (2), 223–247.
- Brunnermeier, Markus and Yuliy Sannikov**, “The I Theory of Money,” *Working paper*, 2016.

- Brunnermeier, Markus K and Yuliy Sannikov**, “A macroeconomic model with a financial sector,” *American Economic Review*, 2014, 104 (2), 379–421.
- Carroll, Christopher D**, “Why do the rich save so much?,” 1998.
- 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.
- Chien, YiLi, Harold Cole, and Hanno Lustig**, “Is the volatility of the market price of risk due to intermittent portfolio rebalancing?,” *American Economic Review*, 2012, 102 (6), 2859–2896.
- Cioffi, Riccardo A**, “Heterogeneous risk exposure and the dynamics of wealth inequality,” *Working Paper*, 2021.
- Cioffi, Riccardo, Galo Nuño, and Samuel Hurtado**, “When Does Wealth Inequality Matter for Asset Pricing?,” 2023.
- Coimbra, Nuno and Hélène Rey**, “Financial Cycles with Heterogeneous Intermediaries,” *The Review of Economic Studies*, mar 2024, 91 (2), 817–857.
- , **Francisco Gomes, Alexander Michaelides, and Jialu Shen**, “Asset Pricing and Risk Sharing Implications of Alternative Pension Plan Systems,” *SSRN Electronic Journal*, 3 2023.
- Corbae, Dean and Pablo D’Erasmus**, “Capital Requirements in a Quantitative Model of Banking Industry Dynamics,” 2014, (14).
- den Heuvel, Skander J Van**, “The welfare cost of bank capital requirements,” *Journal of Monetary Economics*, 2008, 55 (2), 298–320.
- Duarte, Victor, Diogo Duarte, and Dejanir H Silva**, “Machine Learning for Continuous-Time Finance,” *The Review of Financial Studies*, 09 2024, 37 (11), 3217–3271.
- Elenev, Vadim, Tim Lando, and Stijn Van Nieuwerburgh**, “A Macroeconomic Model with Financially Constrained Producers and Intermediaries,” *Econometrica*, 2021, 89 (4), 1361–1418.
- Favilukis, Jack**, “Inequality, stock market participation, and the equity premium,” *Journal of Financial Economics*, 2013, 107 (3), 740–759.
- Fernández-Villaverde, Jesús, Samuel Hurtado, and Galo Nuno**, “Financial frictions and the wealth distribution,” *Econometrica*, 2023, 91 (3), 869–901.
- Gale, Douglas and Tanju Yorulmazer**, “Bank capital, fire sales, and the social value of deposits,” *Economic Theory*, 2020, 69 (4), 919–963.
- Garleanu, Nicolae and Stavros Panageas**, “Heterogeneity and Asset Prices: An Intergenerational Approach,” *Journal of Political Economy*, None 2023, 131 (4), 839–876.
- 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.

- **and** — , “Banking, Liquidity, and Bank Runs in an Infinite Horizon Economy,” *American Economic Review*, 2015, 105 (7), 2011–2043.
- , — , **and Andrea Prestipino**, “Wholesale banking and bank runs in macroeconomic modeling of financial crises,” in “Handbook of macroeconomics,” Vol. 2, Elsevier, 2016, pp. 1345–1425.
- , — , **and** — , “A Macroeconomic Model with Financial Panics,” *The Review of Economic Studies*, 05 2019, 87 (1), 240–288.
- Gomes, Francisco J. and Alexander Michaelides**, “Asset Pricing with Limited Risk Sharing and Heterogeneous Agents,” *Review of Financial Studies*, 2007, 21 (1), 415–448.
- Gomez, Matthieu**, “Asset Prices and Wealth Inequality,” 2017 Meeting Papers 1155, Society for Economic Dynamics 2017.
- Gopalakrishna, Goutham**, “Aliens and continuous time economies,” *Swiss Finance Institute Research Paper*, 2021, 21 (34).
- Greenwood, Robin and Annette Vissing-Jorgensen**, “The Impact of Pensions and Insurance on Global Yield Curves,” Finance Working Paper 19-059, Harvard Business School 2018. HBS Finance Working Paper No. 19-059.
- Gu, Zhouzhou, Mathieu Lauriere, Sebastian Merkel, and Jonathan Payne**, “Global Solutions to Master Equations for Continuous Time Heterogeneous Agent Macroeconomic Models,” *Available at SSRN 4871228*, 2024.
- Güvenen, Fatih**, “A parsimonious macroeconomic model for asset pricing,” *Econometrica*, 2009, 77 (6), 1711–1750.
- Han, Jiequn, Yucheng Yang, and Weinan E**, “Deephm: A global solution method for heterogeneous agent models with aggregate shocks,” *arXiv preprint arXiv:2112.14377*, 2021.
- He, Zhiguo and Arvind Krishnamurthy**, “Intermediary asset pricing,” *American Economic Review*, 2013, 103 (2), 732–770.
- Huang, Ji**, “Breaking the Curse of Dimensionality in Heterogeneous-Agent Models: A Deep Learning-Based Probabilistic Approach,” *SSRN Working Paper*, 2023.
- 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.
- Kargar, Mahyar**, “Heterogeneous intermediary asset pricing,” *Journal of Financial Economics*, 8 2021, 141, 505–532.
- Kennickell, Arthur B. and Martha Starr-Mccluer**, “Household saving and portfolio change: evidence from 1983-89 SCF panel,” *Review of Income and Wealth*, 1997, 43 (4), 381–399.
- Khorrami, Paymon**, “Entry and Slow-Moving Capital: Using Asset Markets to Infer the Costs of Risk Concentration,” Technical Report, SSRN Working Paper 2021. Available at SSRN 2777747.

- Kogan, Leonid and Indrajit Mitra**, “Near-Rational Equilibria in Heterogeneous-Agent Models: A Verification Method,” *The Review of Financial Studies*, 05 2025, 38 (8), 2227–2274.
- Koijen, Ralph S. J. and Motohiro Yogo**, “The Cost of Financial Frictions for Life Insurers,” *American Economic Review*, 2015, 105 (1), 445–475.
- **and** —, “A demand system approach to asset pricing,” *Journal of Political Economy*, 2019, 127 (4), 1475–1515.
- **and** —, “The Evolution from Life Insurance to Financial Engineering,” Working Paper 29030, National Bureau of Economic Research July 2021.
- **and** —, “The Fragility of Market Risk Insurance,” *The Journal of Finance*, 4 2022, 77, 815–862.
- **and** —, “Understanding the Ownership Structure of Corporate Bonds,” *American Economic Review: Insights*, 3 2023, 5, 73–92.
- Koijen, Ralph SJ and Motohiro Yogo**, *Financial economics of insurance*, Princeton University Press, 2023.
- Krusell, Per and Anthony A Smith**, “Income and wealth heterogeneity, portfolio choice, and equilibrium asset returns,” *Macroeconomic dynamics*, 1997, 1 (2), 387–422.
- **and** —, “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 1998, 106 (5), 867–896.
- Lucas, Robert E.**, “Asset Prices in an Exchange Economy,” *Econometrica*, 1978, 46 (6), 1429–1445.
- Maliar, Lilia, Serguei Maliar, and Pablo Winant**, “Deep learning for solving dynamic economic models,” *Journal of Monetary Economics*, 2021, 122, 76–101.
- Merton, Robert C.**, “A Simple Model of Capital Market Equilibrium with Incomplete Information,” *The Journal of Finance*, July 1987, 42 (3), 483–510. Presidential Address to the American Finance Association.
- Mitchell, Olivia S., James M. Poterba, Mark J. Warshawsky, and Jeffrey R. Brown**, “New Evidence on the Money’s Worth of Individual Annuities,” *Brookings Papers on Economic Activity*, 1999, 1999 (2), 131–192.
- Piazzesi, Monika**, “Presidential Address: Housing Betas,” *The Journal of Finance*, 2025, 80 (6), 3103–3136.
- 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.
- Sargent, Thomas J and Lars Ljungqvist**, “Recursive macroeconomic theory,” *Massachusetts Institute of Technology*, 2000, p. 4.
- Vayanos, Dimitri and Jean-Luc Vila**, “A Preferred-Habitat Model of the Term Structure of Interest Rates,” *Econometrica*, 1 2021, 89, 77–112.

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

A Recursive Characterization of Equilibrium

In this section, we derive the recursive representation of equilibrium described in Theorems 1, 2, and 3 in the main text. Recall the notation for the finite dimensional aggregate state vector at time t , $\mathbf{s}_t := (z_t, K_t, \Omega_{b,t}, \Omega_{f,t})$, and the total collection of aggregate states at time t , $\mathbf{S}_t := (\mathbf{s}_t, g_{h,t})$. We let N_s denote the dimension of $\mathbf{s}_t$ and denote the law of motion for $\mathbf{s}_t$ as:

$$d\mathbf{s}_t = (\boldsymbol{\mu}_s(\mathbf{s}_t) \odot \mathbf{s}_t)dt + (\boldsymbol{\sigma}_s(\mathbf{s}_t) \odot \mathbf{s}_t)^T dW_t$$

where $\boldsymbol{\mu}_s$ and $\boldsymbol{\sigma}_s$ are functions for the geometric drift and volatility for $\mathbf{s}_t$ and where $\odot$ denotes element-wise multiplication.¹⁸ For convenience, we also define the companion arithmetic drift as $\boldsymbol{\mu}^s(\mathbf{s}_t) := (\boldsymbol{\mu}_s(\mathbf{s}_t) \odot \mathbf{s}_t)$ and the arithmetic volatility as $\boldsymbol{\sigma}^s(\mathbf{s}_t) := (\boldsymbol{\sigma}_s(\mathbf{s}_t) \odot \mathbf{s}_t)$. Moreover, throughout this Appendix we use subscripts for geometric drifts and superscripts for arithmetic drifts. We denote the law of motion for the household wealth share density by:

$$dg_{h,t}(\omega) = \mu^g(\omega, \mathbf{S})dt + \sigma^g(\omega, \mathbf{S})^T dW_t$$

where μ^g and σ^g are the arithmetic drift and volatility of the density g respectively.

To be clear about what households internalize in the recursive formulation, we model rational expectations belief consistency explicitly. We set up the agent optimization problems with agents who believe that the endogenous aggregate states evolve by the functions $(\tilde{\mu}_K, \tilde{\boldsymbol{\mu}}_s, \tilde{\boldsymbol{\sigma}}_s, \tilde{\mu}_g, \tilde{\sigma}_g)(\cdot)$, understand that their individual actions do not influence the aggregate states (i.e. are price takers), and understand the equilibrium mapping between the aggregate states and price functions $(r^d, (q^l, r^l, \sigma_{q^l})_{l \in \{k, m, n\}})(\cdot)$. We then impose that, in equilibrium, agent beliefs are consistent with the actual laws of motion in the sense that: $(\tilde{\mu}_K, \tilde{\boldsymbol{\mu}}_s, \tilde{\boldsymbol{\sigma}}_s, \tilde{\mu}_g, \tilde{\sigma}_g)(\cdot) = (\mu_K, \boldsymbol{\mu}_s, \boldsymbol{\sigma}_s, \mu_g, \sigma_g)(\cdot)$.

A.1 Household Optimization

In this subsection, we solve the household optimization problem. We start by setting up the household's state evolution. We then construct the Hamilton-Jacobi-Bellman-Equation (HJBE) equation for the households and take the first order conditions (FOCs). Finally, we derive the law of motion

¹⁸We write the diffusion term that's applicable for a more general form with multiple shocks. In the case of a single shock, as is the case in our model, the diffusion should be read as $(\boldsymbol{\sigma}_s(\mathbf{s}_t) \odot \mathbf{s}_t) dW_t$.

for the household Stochastic Discount Factor (SDF) and get the Euler equation.

Perceived idiosyncratic and aggregate state evolution: For each household, their idiosyncratic wealth, a_t , is their individual state. We let $\mathbf{x}_t := (a_t, z_t, K_t, \Omega_{b,t}, \Omega_{f,t})$ and $\mathbf{X}_t := (\mathbf{x}_t, g_{h,t})$ denote the total states for a particular household with individual wealth a_t . Under their beliefs about the endogenous aggregate state variables $(\tilde{\mu}_K, \tilde{\mu}_s, \tilde{\sigma}_s, \tilde{\mu}_g, \tilde{\sigma}_g)$, the geometric drift and volatility of the household's finite state vector $\mathbf{x}_t$ are:

$$\tilde{\mu}_x(a, \mathbf{S}; c_h, \theta_h, \iota_h) = \begin{bmatrix} \mu_a(a, \mathbf{S}; c_h, \theta_h, \iota_h) \\ \mu_z(z) \\ \tilde{\mu}_K(\mathbf{S}) \\ \tilde{\mu}_{\Omega_b}(\mathbf{S}) \\ \tilde{\mu}_{\Omega_f}(\mathbf{S}) \end{bmatrix}, \quad \tilde{\sigma}_x(a, \mathbf{S}; \theta_h) = \begin{bmatrix} \sigma_a(a, \mathbf{S}; \theta_h) \\ \sigma_z \\ 0 \\ \tilde{\sigma}_{\Omega_b}(\mathbf{S}) \\ \tilde{\sigma}_{\Omega_f}(\mathbf{S}) \end{bmatrix},$$

the arithmetic drift and volatility are $\tilde{\mu}^x(\mathbf{x}_t) := (\tilde{\mu}_x(\mathbf{x}_t) \odot \mathbf{x}_t)$ and $\tilde{\sigma}^x(\mathbf{x}_t) := (\tilde{\sigma}_x(\mathbf{x}_t) \odot \mathbf{x}_t)$, and their belief about the law of motion of $g_{h,t}$ satisfies the equation:

$$dg_{h,t}(\omega) = \tilde{\mu}^g(\omega, \mathbf{S})dt + \tilde{\sigma}^g(\omega, \mathbf{S})^T dW_t$$

Here, households understand and internalize the law of motion for their individual state a but take as given the evolution of the exogenous aggregate state, z_t , and their belief about the law of motion for the endogenous aggregate state variables, $(K_t, \Omega_{b,t}, \Omega_{f,t}, g_{h,t})$. We use the tilde notation $(\tilde{\mu}_x, \tilde{\sigma}_x)$ to emphasize the evolution of $\mathbf{x}_t$ depends on agent beliefs about the endogenous aggregate states but do not use the tilde notation for functions that implicitly depend upon beliefs.

HBJE: Let $V_h(a, \mathbf{S})$ denote the value function for a household with state vector $(a, \mathbf{S})$. Given their beliefs about the evolution of the endogenous aggregate states, the $V_h(a, \mathbf{S})$ solves the HJBE:

$$\begin{aligned} \rho_h V_h(a, \mathbf{S}) = & \max_{c_h, \theta_h, \iota_h} \left\{ u(c_h) - \psi_{h,k}(\theta_h^k, \omega_h) \Xi_h(\mathbf{S}) a + \lambda (\mathcal{U}(1 - \theta_h^n + \theta_h^n/q^n) - V_h(a, \mathbf{S})) \right. \\ & + (\tilde{\mu}^x(a, \mathbf{S}; c_h, \theta_h, \iota_h))^T D_x V_h(a, \mathbf{S}) + \langle \partial_g V_h(a, \mathbf{S})(\cdot), \tilde{\mu}^g(\cdot, \mathbf{S}) \rangle \\ & + 0.5 \text{tr} \left\{ (\tilde{\sigma}^x(\mathbf{S}; \theta_h))^T \tilde{\sigma}^x(\mathbf{S}; \theta_h) D_x^2 V_h(a, \mathbf{S}) \right\} \\ & \left. + \langle ((\tilde{\sigma}^x(\mathbf{S}; \theta_h))^T D_{x,g} V_h)(a, \mathbf{S})(\cdot), \tilde{\sigma}^g(\cdot, \mathbf{S}) \rangle + 0.5 \langle D_{gg} V(a, \mathbf{S})(\cdot), \tilde{\sigma}^g(\cdot, \mathbf{S}) \otimes \tilde{\sigma}^g(\cdot, \mathbf{S}) \rangle \right\}, \end{aligned}$$

where $D_x V_h(a, \mathbf{S})$ and $D_x^2 V_h(a, \mathbf{S})$ are the gradient vector and Hessian with respect to the finite dimensional state vector $\mathbf{x}$ at $(a, \mathbf{S})$, $\partial_g V_h(a, \mathbf{S})(\omega')$ is the kernel of the Riesz representation

of the Frechet derivative of V with respect to the distribution g at $(a, \mathbf{S})$, $D_{xg}V(a, \mathbf{S})(\omega') := D_x \partial_g V_h(a, \mathbf{S})(\omega')$ is the Riesz representation of the cross partial of the kernel of the Frechet derivative at $(a, \mathbf{S})$, $D_{gg}(a, \mathbf{S})V(\omega', \omega'') := \partial_g \partial_g V_h(a, \mathbf{S})(\omega', \omega'')$ is the kernel of the Riez representation of the second order Frechet derivative at $(a, \mathbf{S})$, and the inner products are $\langle f, h \rangle := \int f(\omega')h(\omega')d\omega'$ and $\langle\langle K, u \otimes v \rangle\rangle := \int \int K(\omega', \omega'')u(\omega')v(\omega'')d\omega'd\omega''$. We collect the terms:

$$\begin{aligned} \mathcal{G}_g V_h(a, \mathbf{S}) &:= \langle \partial V_h / \partial g(a, \mathbf{S})(\cdot), \tilde{\mu}_g(\cdot, \mathbf{S}) \rangle + \langle ((\tilde{\sigma}^s(\mathbf{S}; \theta_h))^T D_{x,g} V_h)(a, \mathbf{S})(\cdot), \tilde{\sigma}^g(\cdot, \mathbf{S}) \rangle \\ &+ \frac{1}{2} \langle\langle D_{gg} V(a, \mathbf{S})(\cdot), \tilde{\sigma}^g(\cdot, \mathbf{S}) \otimes \tilde{\sigma}^g(\cdot, \mathbf{S}) \rangle\rangle \end{aligned} \quad (\text{A.1})$$

For clarity, we can rewrite the HJBE with the controlled variables taken out of the matrices:

$$\begin{aligned} \rho_h V_h(a, \mathbf{S}) &= \max_{c_h, \theta_h, \iota_h} \left\{ u(c_h) - \psi_{h,k}(\theta_h^k, \omega_h) \Xi_h(\mathbf{S})a + \lambda (\mathcal{U}((1 - \theta_h^n + \theta_h^n/q^n)a) - V_h(a, \mathbf{S})) \right. \\ &+ \mu_a(a, \mathbf{S}; c_h, \theta_h, \iota_h) a \partial_a V_h(a, \mathbf{S}) + (\boldsymbol{\mu}^s(\mathbf{S}))^T D_s V_h(a, \mathbf{S}) + 0.5 \partial_{aa}^2 V_h(a, \mathbf{S}) \sigma_a^2(\mathbf{S}; \theta_h) a^2 \\ &+ \sum_{l \leq N_s} \partial_{as_l} V_h(a, \mathbf{S}) \sigma_a(\mathbf{S}; \theta_h) \sigma_{s_l}(\mathbf{S}) a s_l + 0.5 \text{tr} \left\{ (\boldsymbol{\sigma}^s(\mathbf{S}))^T \boldsymbol{\sigma}^s(\mathbf{S}) D_s^2 V_h(a, \mathbf{S}) \right\} \\ &\left. + \langle \sigma_a(\mathbf{S}; \theta_h) a \partial_a \partial_g V_h(a, \mathbf{S})(\cdot), \tilde{\sigma}^g(\cdot, \mathbf{S}) \rangle + \mathcal{G}_g V_h(a, \mathbf{S}) \right\} \end{aligned}$$

where $\psi_{h,k}(\theta_h^k, \omega_h) = 0.5 \bar{\psi}_k \omega_h^{-(1+\alpha)} (\theta_h^k)^2$ and:

$$\mu_a(a, \mathbf{S}; c_h, \theta_h, \iota_h) = r^d(\mathbf{S}) + \sum_{l \in \{n,k\}} \theta_h^l (r^l(\mathbf{S}) - r^d(\mathbf{S})) - \frac{c_h}{a_h} - \tau_h, \quad \sigma_a(\mathbf{S}; \theta_h) = \sum_{l \in \{n,k\}} \theta_h^l \sigma_{q^l}(\mathbf{S})$$

Then the HJBE becomes:

$$\begin{aligned} \rho_h V_h(a, \mathbf{S}) &= \max_{c_h, \theta_h, \iota_h} \left\{ u(c_h) - \psi_{h,k}(\theta_h^k, \omega_h) \Xi_h a + \lambda (\mathcal{U}((1 - \theta_h^n + \theta_h^n/q^n)a) - V_h(a, \mathbf{S})) \right. \\ &+ \mu_a(a, \mathbf{S}; c_h, \theta_h, \iota_h) a \partial_a V_h(a, \mathbf{S}) + (\tilde{\boldsymbol{\mu}}^s(\mathbf{S}))^T D_s V_h(a, \mathbf{S}) + 0.5 \partial_{aa}^2 V_h(a, \mathbf{S}) (\theta_h^T \sigma_q(\mathbf{S}))^2 a^2 \\ &+ \sum_l \partial_{as_l} V_h(a, \mathbf{S}) \theta_h^T \sigma_q(\mathbf{S}) \tilde{\sigma}_{s_l}(\mathbf{S}) a s_l + 0.5 \text{tr} \left\{ (\tilde{\boldsymbol{\sigma}}^s(\mathbf{S}))^T \tilde{\boldsymbol{\sigma}}^s(\mathbf{S}) D_s^2 V_h(a, \mathbf{S}) \right\} \\ &\left. + \langle \theta_h^T \sigma_q(\mathbf{S}) \partial_a \partial_g V_h(a, \mathbf{S})(\cdot), \tilde{\sigma}^g(\cdot, \mathbf{S}) \rangle a + \mathcal{G}_g V_h(a, \mathbf{S}) \right\} \end{aligned} \quad (\text{A.2})$$

and the FOCs are given by the following equations:

$$\begin{aligned} [c_h] : \quad & 0 = u'(c_h) - \partial_a V_h(a, \mathbf{S}) \\ [\iota_h] : \quad & 0 = \Phi'(\iota) - 1/q^k(\mathbf{S}) \end{aligned}$$

$$\begin{aligned}
[\theta_h^k] : \quad 0 &= (r^k(\mathbf{S}) - r^d(\mathbf{S}))a\partial_a V_h(a, \mathbf{S}) - \partial_{\theta_h^k} \psi_{h,k}(\theta_h^k, \omega_h) \Xi_h a \\
&\quad + (D_x(\partial_a V_h(a, \mathbf{S}))^T (\tilde{\sigma}^x(\mathbf{S}))^T \sigma_{q^k}(\mathbf{S}))a + \langle \partial_a \partial_g V_h(a, \mathbf{S})(\cdot), \tilde{\sigma}^g(\cdot, \mathbf{S}) \rangle \sigma_{q^k}(\mathbf{S})a \\
[\theta_h^n] : \quad 0 &= (r^n(\mathbf{S}) - r^d(\mathbf{S}))a\partial_a V_h(a, \mathbf{S}) + \lambda(1/q^n(\mathbf{S}) - 1)a\mathcal{U}'(\mathcal{C}) \\
&\quad + (D_x(\partial_a V_h(a, \mathbf{S}))^T (\tilde{\sigma}^x(\mathbf{S}))^T \sigma_{q^n}(\mathbf{S}))a + \langle \partial_a \partial_g V_h(a, \mathbf{S})(\cdot), \tilde{\sigma}^g(\cdot, \mathbf{S}) \rangle \sigma_{q^n}(\mathbf{S})a
\end{aligned}$$

SDF evolution: Let $\xi_h(a, \mathbf{S}) := \partial_a V_h(a, \mathbf{S})$. From Itô's Lemma, we have that the drift and volatility of ξ_h (under the household's beliefs) are given by:

$$\begin{aligned}
\mu_{\xi_h}(a, \mathbf{S})\xi_h(a, \mathbf{S}) &= (D_x \xi_h(a, \mathbf{S}))^T \tilde{\mu}^x(a, \mathbf{S}) + \frac{1}{2} \text{tr} \left\{ (\tilde{\sigma}^x(a, \mathbf{S}))^T \tilde{\sigma}^x(a, \mathbf{S}) D_x^2 \xi_h(a, \mathbf{S}) \right\} \\
&\quad + \langle \sigma_a(\mathbf{S}; \theta_h) a \partial_a \partial_g \xi_h(a, \mathbf{S})(\cdot), \tilde{\sigma}^g(\cdot, \mathbf{S}) \rangle + \mathcal{G}_g \xi_h(a, \mathbf{S}) \\
\sigma_{\xi_h}(a, \mathbf{S})\xi_h(a, \mathbf{S}) &= (\tilde{\sigma}^x(a, \mathbf{S}))^T D_x \xi_h(a, \mathbf{S}) + \langle \partial_g \xi_h(a, \mathbf{S})(\cdot), \tilde{\sigma}^g(\cdot, \mathbf{S}) \rangle
\end{aligned}$$

Thus, we can rewrite the FOCs as:

$$\begin{aligned}
[\theta_h^k] : \quad 0 &= \xi_h(a, \mathbf{S})(r^k(\mathbf{S}) - r^d(\mathbf{S})) - \partial_{\theta_h^k} \psi_{h,k}(\theta_h^k, \omega_h) \Xi_h(a, \mathbf{S}) + \sigma_{\xi_h}(a, \mathbf{S})\xi_h(a, \mathbf{S})\sigma_{q^k}(\mathbf{S}) \\
[\theta_h^n] : \quad 0 &= \xi_h(a, \mathbf{S})(r^n(\mathbf{S}) - r^d(\mathbf{S})) + \lambda(1/q^n(\mathbf{S}) - 1)\mathcal{U}'(\mathcal{C}(a, \mathbf{S})) \\
&\quad + \sigma_{\xi_h}(\mathbf{S}; \theta_h)\xi_h(a, \mathbf{S})\sigma_{q^n}(\mathbf{S})
\end{aligned}$$

Finally, by differentiating the HJBE (A.2) with respect to a we get the Euler equation:

$$\begin{aligned}
\rho_h \xi_h(a, \mathbf{S}) &= (-\psi_{h,k}(\theta_h^k, \omega_h) - \partial_{\omega_h} \psi_{h,k}(\theta_h^k, \omega_h) \omega_h + \partial_{\theta_h^k} \psi_{h,k}(\theta_h^k, \omega_h) \theta_h^k) \Xi_h - \lambda \xi_h(a, \mathbf{S}) \\
&\quad + \xi_h(a, \mathbf{S}) (r^d(\mathbf{S}) + \tau_h) + \frac{1}{2} \text{tr} \left\{ (\tilde{\sigma}^x(\mathbf{S}, \theta_h))^T \tilde{\sigma}^x(\mathbf{S}, \theta_h) D_x^2 \xi_h(a, \mathbf{S}) \right\} \\
&\quad + (D_s \xi_h(\mathbf{x}))^T \tilde{\mu}^x(a, \mathbf{S}) + \langle \sigma_a(\mathbf{S}; \theta_h) \partial_a \partial_g \xi_h(a, \mathbf{S})(\cdot), \tilde{\sigma}^g(\cdot, \mathbf{S}) \rangle a + \mathcal{G}_g \xi_h(a, \mathbf{S}) \\
&= (\mu_{\xi_h} - \lambda_h) \xi_h(a, \mathbf{S}) + \xi_h(a, \mathbf{S}) (r^d(\mathbf{S}) - \tau_h) \\
&\quad - (\psi_{h,k}(\theta_h^k, \omega_h) + \partial_{\omega_h} \psi_{h,k}(\theta_h^k, \omega_h) \omega_h - \partial_{\theta_h^k} \psi_{h,k}(\theta_h^k, \omega_h) \theta_h^k) \Xi_h.
\end{aligned}$$

A.2 Financial Intermediary Optimization

Let $V_j(a, \mathbf{S})$ denote the value function for a financier of type $j \in \{b, f\}$ with state variable $(a, \mathbf{S})$. Given their beliefs about the evolution of the aggregate states, the value function $V_j(a, \mathbf{S})$ for a

financier solves the HJBE below:

$$\begin{aligned} \rho_j V_j(a, \mathbf{S}) = & \max_{c_j, \theta_j, \iota_j} \left\{ u(c_j) - \sum_{l \in \{k, m\}} \psi_{j,l}(\theta_j^l) \Xi_j(\mathbf{S}) a + \mu_a(a, \mathbf{S}; c_j, \theta_j, \iota_j) a \partial_a V_j(a, \mathbf{S}) \right. \\ & + (\tilde{\boldsymbol{\mu}}^s(\mathbf{S}))^T D_s V_j(a, \mathbf{S}) + \frac{1}{2} \partial_{aa}^2 V_j(a, \mathbf{S}) \sigma_a^2(\mathbf{S}; \theta_j) a^2 + \sum_{l \leq N_s} \partial_{as_l} V_j(a, \mathbf{S}) \sigma_a(\mathbf{S}; \theta_j) \tilde{\sigma}_{s_l}(\mathbf{S}) \\ & \left. + \langle \sigma_a(\mathbf{S}; \theta_j) a \partial_{ag} V_j(a, \mathbf{S})(\cdot), \tilde{\sigma}^g(\cdot, \mathbf{S}) \rangle \right\} + \frac{1}{2} \text{tr} \left\{ (\tilde{\boldsymbol{\sigma}}^s(\mathbf{S}))^T \tilde{\boldsymbol{\sigma}}^s(\mathbf{S}) D_s^2 V_j(a, \mathbf{S}) \right\} + \mathcal{G}_g V_j(a, \mathbf{S}) \end{aligned} \quad (\text{A.3})$$

where:

$$\begin{aligned} \mu_a(a, \mathbf{S}; c_j, \theta_j, \iota_j) &= r^j(\mathbf{S}) + \sum_{l \in \{m, k\}} \theta_j^l (r^l(\mathbf{S}) - r^j(\mathbf{S})) - c_{j,t}/a_j - \tau \\ \sigma_a(\mathbf{S}; \theta_j) &= \theta_j^k \sigma_{q^k}(\mathbf{S}) + \theta_j^m \sigma_{q^m}(\mathbf{S}) + (1 - \theta_j^k - \theta_j^m) \sigma_{q^j}(\mathbf{S}) \end{aligned}$$

where $r^j(\cdot)$ and $\sigma_{q^j}(\cdot)$ are the average return and volatility of return process for the liabilities issued by a financial intermediary of type j (i.e. deposits for bankers and pensions for fund managers). Following the same steps as for the household, the FOCs for the financial intermediary are:

$$\begin{aligned} [c_j] : \quad & 0 = u'(c_j) - \partial_a V_j(a, \mathbf{S}) \\ [\iota_j] : \quad & 0 = \Phi'(\iota_j) - 1/q^k(\mathbf{S}) \\ [\theta_j^k] : \quad & 0 = r^k(\mathbf{S}) - r^j(\mathbf{S}) + \sigma_{\xi_j}(\mathbf{S}; \theta_j) (\sigma_{q^k}(\mathbf{S}) - \sigma_{q^j}(\mathbf{S})) - \partial_{\theta_j^k} \psi_{j,k}(\theta_j^k) \Xi_j(\mathbf{S}) \\ [\theta_j^m] : \quad & 0 = r^m(\mathbf{S}) - r^j(\mathbf{S}) + \sigma_{\xi_j}(\mathbf{S}; \theta_j) (\sigma_{q^m}(\mathbf{S}) - \sigma_{q^j}(\mathbf{S})) - \partial_{\theta_j^m} \psi_{j,m}(\theta_j^m) \Xi_j(\mathbf{S}) \end{aligned}$$

and the Euler equation is:

$$\rho_j + \lambda_j = \mu_{\xi_j}(a, \mathbf{S}) + r^j(\mathbf{S}) - \tau + \sigma_{\xi_j}(\mathbf{S}) \sigma_{q^j}(\mathbf{S}) - \sum_{l \in \{k, m\}} \left(\psi_{j,l} - \partial_{\theta_j^l} \psi_{j,l}(\theta_j^l) \right) \Xi_j(\mathbf{S})$$

Proof of Proposition 1 (Aggregation Within the Financial Sector). We guess and verify that the value function for financier $j \in \{b, f\}$ satisfies $V_j(a, \mathbf{S}) = \eta_j(\mathbf{S})^{-\gamma_j} a^{1-\gamma_j} / (1 - \gamma_j)$ and $\xi_j(a, \mathbf{S}) = \eta_j(\mathbf{S})^{-\gamma_j} a^{-\gamma_j}$. This implies a consumption policy function $c_j = \eta_j(\mathbf{S}) a$ and so the geometric drift and volatility of a for a financier of type j are:

$$\begin{aligned} \mu_a(a, \mathbf{S}; c_j, \theta_j, \iota_j) &= r^j(\mathbf{S}) + \sum_{l \in \{m, k\}} \theta_j^l (r^l(\mathbf{S}) - r^j(\mathbf{S})) - \eta_j(\mathbf{S}) - \tau_b =: \mu_{a_j}(\mathbf{S}; \eta_j, \theta_j, \iota) \\ \sigma_a(a, \mathbf{S}; \theta_j) &= \theta_j^k \sigma_{q^k}(\mathbf{S}) + \theta_j^m \sigma_{q^m}(\mathbf{S}) + (1 - \theta_j^k - \theta_j^m) \sigma_{q^j}(\mathbf{S}) =: \sigma_{a_j}(\mathbf{S}, \theta_j). \end{aligned}$$

This implies that the portfolio choice θ_j^l is independent of a and characterized by:

$$0 = r^l(\mathbf{S}) - r^j(\mathbf{S}) + \sigma_{\xi_j}(\mathbf{S}; \theta_j)(\sigma_{q^l}(\mathbf{S}) - \sigma_{q^j}(\mathbf{S})) - \partial_{\theta_j^l} \psi_{j,l}(\theta_j^l) \Xi_j(\mathbf{S})$$

Recall that we set the normalizing constant $\Xi_j = \xi_j$. Finally, we can verify the the guess by substituting the guesses into the Euler equation. It follows that financier choices are either independent of or linear in wealth and so we get aggregation. $\square$

A.3 Distribution Evolution

In this subsection, we characterize the evolution of the wealth distributions implied by the agent choice functions and price functions. We start by studying the evolution of aggregate wealth in the different sectors. For the banking and fund sectors, from Proposition 1, all agents choose the same portfolio allocations and so we do not need financier distributions to understand the aggregate dynamics. By contrast, for the household sector, different households choose different portfolios and so we do need the distribution of household wealth to understand aggregate dynamics.

Sector level wealth 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 $j \in \{h, b, f\}$. Then the aggregate wealth in financial sector $j \in \{b, f\}$ evolves according to:

$$\begin{aligned} \frac{dA_{j,t}}{A_{j,t}} &= \frac{1}{A_{j,t}} \int_0^\infty \mu_{a_j}(\mathbf{S}_t; c_j(a, \mathbf{S}_t), \theta_j(a, \mathbf{S}_t), \iota) a g_{b,t}(a) da dt + \frac{1}{A_{j,t}} \lambda_j(\bar{\phi}_j A_t - A_{j,t}) dt \\ &\quad + \frac{1}{A_{j,t}} \int_0^\infty \sigma_{a_j}(\theta_j(a, \mathbf{S}_t), \mathbf{S}_t) dW_t a g_{b,t}(a) da \end{aligned}$$

For the financial sectors, all agents choose the same consumption function and portfolio so $c_j(a, \mathbf{S}_t) = \eta_j(\mathbf{S}_t)a$ and $\theta_j(a, \mathbf{S}_t) = \theta_j(\mathbf{S}_t)$ for all a and $j \in \{b, f\}$. This implies that for financial sector $j \in \{b, f\}$ we have:

$$\begin{aligned} \frac{dA_{j,t}}{A_{j,t}} &= \left(r^j(\mathbf{S}_t) + \sum_{l \in \{m, k\}} \theta_j^l(\mathbf{S}_t)(r^l(\mathbf{S}_t) - r^j(\mathbf{S}_t)) - \eta(\mathbf{S}_t)^{-1/\gamma_j} - \tau + \lambda_j \left(\frac{\bar{\phi}_j}{\Omega_{j,t}} - 1 \right) \right) dt \\ &\quad + \left(\theta_j^k(\mathbf{S}_t) \sigma_{q^k}(\mathbf{S}_t) + \theta_j^m(\mathbf{S}_t) \sigma_{q^m}(\mathbf{S}_t) + (1 - \theta_j^k(\mathbf{S}_t) - \theta_j^m(\mathbf{S}_t)) \sigma_{q^j}(\mathbf{S}_t) \right) dW_t \end{aligned}$$

We now consider aggregate wealth $A(\mathbf{S}_t) = q^k(\mathbf{S}_t)K_t + q^m(\mathbf{S}_t)M_t$. Denote the share of wealth

in capital stock by $\vartheta_t(\mathbf{S}) = q^k(\mathbf{S})K_t/(q^k(\mathbf{S})K_t + q^m(\mathbf{S})M_t)$. The evolution of A_t follows:

$$\begin{aligned} \frac{dA_t}{A_t} = & \left(\vartheta(\mathbf{S}_t)(\mu_{q^k}(\mathbf{S}) + \Phi(\iota) - \delta) + (1 - \vartheta(\mathbf{S}_t))(\mu_{q^m}(\mathbf{S}_t) + \Phi(\iota) - \delta) \right) dt \\ & + (\vartheta(\mathbf{S}_t)\sigma_{q^k}(\mathbf{S}_t) + (1 - \vartheta(\mathbf{S}_t))\sigma_{q^m}(\mathbf{S}_t))dW_t \end{aligned}$$

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

$$\frac{d\Omega_{j,t}}{\Omega_{j,t}} = \left(\mu_{A_j}(\mathbf{S}_t) - \mu_A(\mathbf{S}_t) + (\sigma_A(\mathbf{S}_t) - \sigma_{A_j}(\mathbf{S}_t))\sigma_A(\mathbf{S}_t) \right) dt + (\sigma_{A_j}(\mathbf{S}_t) - \sigma_A(\mathbf{S}_t))dW_t \quad (\text{A.4})$$

Household wealth evolution: Let $\check{g}_{h,t}(a)$ and $g_{h,t}(\omega) := \check{g}_{h,t}(\omega A)$ denote the measure functions over household wealth and household wealth shares respectively. So, we have $\int_0^\infty \check{g}_{h,t}(a)ada = A_{h,t}$ and $\int g_{h,t}(\omega)\omega d\omega = \Omega_h$. The Kolmogorov Forward Equation (KFE) for $\check{g}_{h,t}(a)$ is given by:

$$\begin{aligned} d\check{g}_{h,t}(a) = & \left(\lambda_h \bar{\phi}_h \phi_h(a/A_{h,t}) - \lambda_h \check{g}_{h,t}(a) - \partial_a [\mu_a(a, \mathbf{S}_t) a \check{g}_{h,t}(a)] \right. \\ & \left. + \frac{1}{2} \partial_{aa} [\sigma_a^2(a, \mathbf{S}_t) a^2 \check{g}_{h,t}(a)] \right) dt - \partial_a [\sigma_a(a, \mathbf{S}_t) a \check{g}_{h,t}(a)] dW_t \end{aligned}$$

The evolution of the household wealth share, $\omega_{i,t} := a_{i,t}/A_t$, is:

$$\begin{aligned} \frac{d\omega_{i,t}}{\omega_{i,t}} = & (\mu_{a_i}(\omega_{i,t}A(\mathbf{S}), \mathbf{S}_t) - \mu_A(\mathbf{S}_t) + (\sigma_A(\mathbf{S}_t) - \sigma_{a_i}(\omega_{i,t}A(\mathbf{S}), \mathbf{S}_t))\sigma_A(\mathbf{S}_t)) dt \\ & + (\sigma_{a_i}(\omega_{i,t}A(\mathbf{S}_t), \mathbf{S}_t) - \sigma_{A,t}(\mathbf{S}_t))dW_t =: \mu_{\omega_i}(\omega_{i,t}, \mathbf{S}_t)dt + \sigma_{\omega_i}(\omega_{i,t}, \mathbf{S}_t)dW_t \end{aligned}$$

so the KFE for the measure function over wealth shares, $g_{h,t}(\omega) := \check{g}_{h,t}(\omega A)$, is:

$$\begin{aligned} dg_{h,t}(\omega) = & \left(\lambda_h \bar{\phi}_h \phi(\omega/\Omega_h) - \lambda_h g_{h,t}(\omega) - \partial_\omega [\mu_\omega(\omega, \mathbf{S}_t) \omega g_{h,t}(\omega)] \right. \\ & \left. + \frac{1}{2} \partial_{\omega\omega} [\sigma_\omega^2(\omega, \mathbf{S}_t) \omega^2 g_{h,t}(\omega)] \right) dt - \partial_\omega [\sigma_\omega(\omega, \mathbf{S}_t) \omega g_{h,t}(\omega)] dW_t. \end{aligned} \quad (\text{A.5})$$

A.4 Recursive Rational Expectations Equilibrium

Definition 2 (Recursive Rational Expectations Equilibrium). Is a collection of endogenous state evolution functions $(\mu_K, \boldsymbol{\mu}_s, \boldsymbol{\sigma}_s, \mu_g, \sigma_g)(\cdot)$, price functions $(r^d, (q^l, r^l, \sigma_{q^l})_{l \in \{k,m,n\}})(\cdot)$, and agent functions $((\xi_j, c_j, \theta_j)_{j \in \{h,b,f\}}, \mathcal{C}_h, \iota)(\cdot)$ such that:

1. Given their beliefs about the functions $(\tilde{\mu}_K, \tilde{\boldsymbol{\mu}}_s, \tilde{\boldsymbol{\sigma}}_s, \tilde{\mu}_g, \tilde{\sigma}_g)(\cdot)$, the agent functions $((\xi_j, c_j, \theta_j)_{j \in \{h,b,f\}}, \mathcal{C}_h, \iota)(\cdot)$ solve the HJBs (A.2) and (A.3),

2. The price functions $(r^d, (q^l, r^l, \sigma_{q^l})_{l \in \{k, m, n\}})(\cdot)$ satisfy market clearing for: (a) Goods: $\int_{\omega} c_h(\omega A(\mathbf{S}), \mathbf{S}) g(\omega) d\omega + \sum_{j \in \{b, f\}} C_j(\mathbf{S}) + \lambda_h \int \mathcal{C}_h(\omega A(\mathbf{S}), \mathbf{S}) g(\omega) d\omega = e^{z_t} K_t - \iota_t K_t$, (b) Capital: $\int_{\omega} \theta_h^k(\omega A(\mathbf{S}), \mathbf{S}) A_h(\mathbf{S}) g(\omega) d\omega + \sum_{j \in \{b, f\}} \theta_j^k(\mathbf{S}) A_j(\mathbf{S}) = q_t^k K_t$, (c) Pensions: $\int_{\omega} \theta_h^n(\omega A(\mathbf{S}), \mathbf{S}) A_h(\mathbf{S}) g(\omega) d\omega + \theta_f^n(\mathbf{S}) A_f(\mathbf{S}) = 0$, (d) Deposits: $\int_{\omega} \theta_h^d(\omega A(\mathbf{S}), \mathbf{S}) A_h(\mathbf{S}) g(\omega) d\omega + \theta_b(\mathbf{S}) A_b(\mathbf{S}) = 0$, and (e) Bonds: $\sum_{j \in \{b, f\}} \theta_j(\mathbf{S}) A_j(\mathbf{S}) = q_t^m \mathcal{M} K_t$.
3. Given the agent and price functions, the evolution of the endogenous states $(K_t, \{\Omega_{j,t}\}_{j \in \{f, b\}}, g_{h,t})$ satisfy $dK_t = (\phi(\iota_t) K_t - \delta K_t) dt$, (A.4), and the KFE (A.5) respectively,
4. Agent beliefs are consistent: $(\tilde{\mu}_K, \tilde{\mu}_s, \tilde{\sigma}_s, \tilde{\mu}_g, \tilde{\sigma}_g) = (\mu_K, \mu_s, \sigma_s, \mu_g, \sigma_g)$.

We can now prove Theorems 1, 2, and 3, which characterize the recursive equilibrium.

Proof of Theorem 1 (Block 1: Agent Optimization). Impose belief consistency $(\tilde{\mu}_K, \tilde{\mu}_s, \tilde{\sigma}_s, \tilde{\mu}_g, \tilde{\sigma}_g) = (\mu_K, \mu_s, \sigma_s, \mu_g, \sigma_g)$ and $\Xi_j(\cdot) = \xi_j(\cdot)$ on the results from Subsections A.1 and A.2. Then, given equilibrium prices and price functions $((\xi_j, c_j, \theta_j)_{j \in \{h, b, f\}}, \mathcal{C}_h, \iota)(\cdot)$ the household, banker, and fund choices (11 variables), $(c_h, \mathcal{C}_h, c_b, c_f, \theta_h^k, \theta_h^n, \theta_b^k, \theta_b^m, \theta_f^k, \theta_f^m, \iota)$, satisfy the optimization equations (11 equations) for any $\mathbf{S}$:

$$\begin{aligned}
0 &= u'(c_j) - \xi_j(a, \mathbf{S}), \quad j \in \{h, b, f\} \\
0 &= -\mathcal{C}_h + (1 - \theta_h^n + \theta_h^n/q^n(\mathbf{S})) a \\
0 &= r^k(\mathbf{S}) - r^d(\mathbf{S}) + \sigma_{\xi_h}(a, \mathbf{S}) \sigma_{q^k}(\mathbf{S}) - \partial_{\theta_h^k} \psi_{h,k}(\theta_h^k, \omega_h) \\
0 &= r^n(\mathbf{S}) - r^d(\mathbf{S}) + \sigma_{\xi_h}(a, \mathbf{S}) \sigma_{q^n}(\mathbf{S}) + \lambda_h (1/q^n(\mathbf{S}) - 1) \mathcal{U}'(\mathcal{C}_h) \\
0 &= r^l(\mathbf{S}) - r^d(\mathbf{S}) + \sigma_{\xi_b}(a, \mathbf{S}) \sigma_{q^l}(\mathbf{S}) - \partial_{\theta_b^l} \psi_{b,l}(\theta_b^l), \quad l \in \{m, k\} \\
0 &= r^l(\mathbf{S}) - r^n(\mathbf{S}) + \sigma_{\xi_f}(a, \mathbf{S}) (\sigma_{q^l}(\mathbf{S}) - \sigma_{q^n}(\mathbf{S})) - \partial_{\theta_f^l} \psi_{f,l}(\theta_f^l), \quad l \in \{m, k\} \\
0 &= \Phi'(\iota) - 1/q^k(\mathbf{S})
\end{aligned}$$

and the 3 agent marginal value functions $(\xi_h, \xi_b, \xi_f)(a, \mathbf{S})$ satisfy the 3 Euler equations:

$$\begin{aligned}
\rho_h + \lambda_h &= \mu_{\xi_h}(a, \mathbf{S}) + r^d(\mathbf{S}) - \tau_h - \psi_{h,k}(\theta_h^k, \omega_h) - \partial_{\omega_h} \psi_{h,k}(\theta_h^k, \omega_h) \omega_h + \partial_{\theta_h^k} \psi_{h,k}(\theta_h^k, \omega_h) \theta_h^k \\
\rho_b + \lambda_b &= \mu_{\xi_b}(a, \mathbf{S}) + r^d(\mathbf{S}) - \tau - \sum_{l \in \{k, m\}} \left(\psi_{b,l}(\theta_b^l) - \partial_{\theta_b^l} \psi_{b,l}(\theta_b^l) \right) \\
\rho_f + \lambda_f &= \mu_{\xi_f}(a, \mathbf{S}) + r_f^n(\mathbf{S}) - \tau + \sigma_{\xi_f}(\mathbf{S}) \sigma_{q^n}(\mathbf{S}) - \sum_{l \in \{k, m\}} \left(\psi_{f,l} - \partial_{\theta_f^l} \psi_{f,l}(\theta_f^l) \right)
\end{aligned}$$

where the equilibrium household net tax rate is $\tau_h = \tau - \lambda_h A_{h,t} - (\lambda_b A_{b,t} + \lambda_f A_{f,t} - \sum_{j \in \{b, f\}} \lambda_j \phi_j A_t) / A_{h,t}$

and the drift and volatility of ξ_j are characterized by Itô's lemma:

$$\begin{aligned}\mu_{\xi_j}(a, \mathbf{S})\xi_j(a, \mathbf{S}) &= (D_x \xi_j(a, \mathbf{S}))^T \boldsymbol{\mu}^x(a, \mathbf{S}) + \frac{1}{2} \text{tr} \left\{ (\boldsymbol{\sigma}^x(a, \mathbf{S}))^T \boldsymbol{\sigma}^x(a, \mathbf{S}) D_x^2 \xi_j(a, \mathbf{S}) \right\} \\ &\quad + \langle \sigma_a(\mathbf{S}; \theta_j) a \partial_a \partial_g \xi_j(a, \mathbf{S})(\cdot), \tilde{\sigma}^g(\cdot, \mathbf{S}) \rangle + \mathcal{G}_g \xi_j(a, \mathbf{S}) \\ \sigma_{\xi_j}(a, \mathbf{S})\xi_j(a, \mathbf{S}) &= (\boldsymbol{\sigma}^x(a, \mathbf{S}))^T D_x \xi_j(a, \mathbf{S}) + \langle \partial_g \xi_j(a, \mathbf{S})(\cdot), \tilde{\sigma}^g(\cdot, \mathbf{S}) \rangle\end{aligned}$$

Proof of Theorem 2 (Block 2: Distribution Evolution). This follows directly from the KFEs (A.4) and (A.5) in Subsection A.3 with the equilibrium policy and price functions substituted in. $\square$

Proof of Theorem 3 (Block 3: Equilibrium and Consistency). The statement of asset market clearing conditions in the main text comes from dividing each equation by $A(\mathbf{S})$ and observing that $\vartheta = q^k K/A(\mathbf{S})$. The Itô's lemma expression is standard. $\square$

Internet Appendix

B Modeling Retirement (Internet Appendix)

In this section, we offer a microfoundation for the terminal utility $\mathcal{U}(\mathcal{C})$. Consider the setup from Section 2 but with explicit agent retirement. Formally, agents are now born as active participants in the economy, then transition to retirement at rate λ_h , and then finally die at rate λ_d . When agents retire they stop participating in economic production and exchange their assets for resources to consume during their retirement. They receive flow utility $u(c_{h,t}) = c_{h,t}^{1-\gamma}/(1-\gamma)$ from consuming during retirement. One way to model this is for the households to purchase a stock of goods at retirement, which they progressively consume. We focus on this case in this Appendix since it leads to the exact market clearing conditions as the main text. An alternative model would be for households to purchase an annuity that pays goods throughout their retirement, which would lead to a similar household problem but minor differences in the market clearing conditions.

As in the main text, let $V_h(a, \mathbf{S})$ denote the value function of household with wealth a when the aggregate state is $\mathbf{S}$. Now, let $W_h(A, \mathbf{S})$ denote the value function of a retired household with remaining wealth A . The value of W_h is derived in Proposition 2 below.

Proposition 2. *The household's value function at retirement is:*

$$W_h(A, \mathbf{S}) = \varrho A^{1-\gamma} / ((1-\gamma)(\rho + \lambda_d)) \quad (\text{B.1})$$

where ϱ satisfies:

$$\frac{\varrho(\rho + \lambda_d)}{(\rho + \lambda_d)(1-\gamma)} = \frac{1}{1-\gamma} \left(\frac{\rho + \lambda_d}{\varrho} \right)^{(1-\gamma)/\gamma} + \frac{\varrho}{\rho + \lambda_d} \left(1 - \left(\frac{\rho + \lambda_d}{\varrho} \right)^{1/\gamma} \right) \quad (\text{B.2})$$

Proof. The optimization problem of a household that retires at time T_r with a_{T_r} is:

$$\begin{aligned} W_h(A_{T_r}, \mathbf{S}_{T_r}) &= \max_c \mathbb{E}_{T_r} \left[\int_{T_r}^{T_d} e^{-\rho(t-T_r)} u(c_{T_r+t}) dt \right] \text{ s.t.} \\ dA_t &= (A_t - c_t) dt, \quad A_{T_r} = (1 - \theta_{h,t}^n + \theta_{h,t}^n / q_t^n) a_{T_r} \end{aligned}$$

where A_{T_r} is the total household wealth after redeeming pension contracts with the fund managers. From the setup we can see that W_h does not depend upon $\mathbf{S}_{T_r}$ and so we can drop it from the

function. It follows that $W_h(A)$ satisfies the HJBE:

$$(\rho + \lambda_d)W_h(A) = \max_c \{u(c) + \partial W_h(A)(A - c)\}$$

for which the FOC is $u'(c) = \partial_A W(A)$. We guess (and verify) that W_h takes the form (B.1) and impose the functional form $u(c_{h,t}) = c_{h,t}^{1-\gamma}/(1-\gamma)$. Then the FOC becomes $c = ((\rho + \lambda_d)/\varrho)^{1/\gamma} A$ and the HJBE becomes:

$$\frac{\varrho A^{1-\gamma}(\rho + \lambda_d)}{(\rho + \lambda_d)(1-\gamma)} = \frac{1}{1-\gamma} \left(\left(\frac{\rho + \lambda_d}{\varrho} \right)^{1/\gamma} A \right)^{1-\gamma} + \frac{\varrho}{\rho + \lambda_d} A^{-\gamma} \left(A - \left(\frac{\rho + \lambda_d}{\varrho} \right)^{1/\gamma} A \right)$$

so that ϱ implicitly satisfies (B.2) (after simplification) and our guess is verified. $\square$

Extension with annuities: If the household invested in an annuity with coupon rate ζ instead of storing consumption goods, then the household problem is the same except that now $da_t = (\zeta a_t - c_t)dt$ and so the implicit expression for ϱ is adjusted by ζ to become:

$$\frac{\varrho(\rho + \lambda_d)}{(\rho + \lambda_d)(1-\gamma)} = \frac{1}{1-\gamma} \left(\frac{\rho + \lambda_d}{\varrho} \right)^{(1-\gamma)/\gamma} + \frac{\varrho}{\rho + \lambda_d} \left(\zeta - \left(\frac{\rho + \lambda_d}{\varrho} \right)^{1/\gamma} \right).$$

Micro-foundation for main text: Returning to the formulation in the main text, we can micro-found consumption at death with this retirement model by setting $\mathcal{U}(\cdot) = W(\cdot)$ and noting that at retirement $C_{T_r} = (1 - \theta_{h,t}^n + \theta_{h,t}^n/q_t^n)A_{T_r} = A_{T_r}$ and so we have $U(C_{T_r}) = V(A_{T_r})$ as required. Under this formulation, we have the parameter mapping: $1 - \beta = 1/(\rho + \lambda_d)$.

C Construction of The Loss Function (Internet Appendix)

In this section, we provide the details for the construction of the loss function in Subsection 3.2. We first re-characterize the equilibrium blocks on the finite agent state space. We then re-express the system in terms of equilibrium functions. Finally, we show how to characterize equilibrium given the approximated functions.

C.1 Finite Agent State Space

Formally, we approximate the household wealth-share measure function g_h by a collection of I agents. The aggregate state space for the finite agent approximation becomes $\hat{\mathcal{S}} := (z, K, \omega_1, \dots,$

$\omega_I, \Omega_b, \Omega_f$), where $\hat{\omega} := \omega_1, \dots, \omega_I$ are the wealth shares of each agent.¹⁹ The KFE for the density $g(\omega)$ is approximated by the evolution of $\{\omega_{i,t} = a_{i,t}/A_t\}_{i \leq I}$, which we impose follows

$$\frac{d\omega_{i,t}}{\omega_{i,t}} = \left(\mu_{a_i}(\hat{\mathbf{S}}_t) - \mu_A(\hat{\mathbf{S}}_t) + (\sigma_A(\hat{\mathbf{S}}_t) - \sigma_{a_i}(\hat{\mathbf{S}}_t))\sigma_A(\hat{\mathbf{S}}_t) \right) dt + (\sigma_{a_i}(\hat{\mathbf{S}}_t) - \sigma_A(\hat{\mathbf{S}}_t))dW_t \quad (\text{C.1})$$

for all $i \leq I$ where:

$$\begin{aligned} \mu_{a_i}(\hat{\mathbf{S}}_t) &= r^d(\hat{\mathbf{S}}_t) + \sum_{l \in \{n,k\}} \theta_h^l(\hat{\mathbf{S}}_t)(r^l(\hat{\mathbf{S}}_t) - r^d(\hat{\mathbf{S}}_t)) - \eta_h(\hat{\mathbf{S}}_t) - \tau_h + \lambda_h \left(\frac{\bar{\phi}_h \Omega_{h,t}}{\omega_{i,t}} - 1 \right) \\ \sigma_{A_i}(\hat{\mathbf{S}}_t) &= \theta_i^k(\hat{\mathbf{S}}_t)\sigma_{q^k}(\hat{\mathbf{S}}_t) + \theta_i^m(\hat{\mathbf{S}}_t)\sigma_{q^m}(\hat{\mathbf{S}}_t) \end{aligned}$$

This follows the approach developed in [Gu et al. \(2024\)](#) where the law of motion averages across idiosyncratic death shocks so that our finite agent economy approximates the density in which idiosyncratic noise is averaged out. The literal interpretation is that the law of motion for $\omega_{i,t}$ is the expected law of motion for a family of agents whose members die at rate λ_h and are reborn with wealth share drawn from $\phi_h(\omega)$ after they die.

Define the corresponding agent state space by $\hat{\mathbf{X}} = (a, \hat{\mathbf{S}})$ for an agent with wealth a . Under the finite agent approximation, the application of Itô's lemma to a variable $\nu_j(a, \hat{\mathbf{S}})$ leads to the following expressions that no longer involve any functional derivatives:

$$\begin{aligned} \mu_{\nu_j}(a, \hat{\mathbf{S}})\nu_j(a, \hat{\mathbf{S}}) &= (D_{\hat{\mathbf{X}}} \nu_j(a, \hat{\mathbf{S}}))^T \boldsymbol{\mu}_{\hat{\mathbf{X}}}(a, \hat{\mathbf{S}}) \\ &\quad + \frac{1}{2} \text{tr} \left\{ (\boldsymbol{\sigma}_{\hat{\mathbf{X}}}(a, \hat{\mathbf{S}}) \odot \hat{\mathbf{X}})^T (\boldsymbol{\sigma}_{\hat{\mathbf{X}}}(a, \hat{\mathbf{S}}) \odot \hat{\mathbf{X}}) D_{\hat{\mathbf{X}}}^2 \nu_j(a, \hat{\mathbf{S}}) \right\} \\ \sigma_{\nu_j}(a, \hat{\mathbf{S}})\nu_j(a, \hat{\mathbf{S}}) &= (\boldsymbol{\sigma}_{\hat{\mathbf{X}}}(a, \hat{\mathbf{S}}) \odot \hat{\mathbf{X}})^T D_{\hat{\mathbf{X}}} \nu_j(a, \hat{\mathbf{S}}) \end{aligned}$$

C.2 Equilibrium Characterization on Finite Agent State Space

We now bring together the three blocks and rewrite the equations with the finite agent state space. Let $\eta_j := c_j/a$ denote the consumption-to-wealth function for an agent in sector $j \in \{h, b, f\}$ with wealth a . The agent functions, distribution evolution functions, and equilibrium price functions:

$$\begin{aligned} &\left(\{\eta_j, \xi_j, \mu_{\xi_j}, \sigma_{\xi_j}\}_{j \in \{h,b,f\}}, \{\theta_h^l\}_{l \in \{k,n\}}, \{\theta_j^l\}_{j \in \{b,f\}, l \in \{k,n,m\}}, \ell \right) (\hat{\mathbf{S}}), \\ &\left(\{\mu_{\omega_i}, \sigma_{\omega_i}\}_{i \leq I}, \{\mu_{\Omega_j}, \sigma_{\Omega_j}\}_{j \in \{b,f\}} \right) (\hat{\mathbf{S}}), \\ &\left(r^d, \{q^l, r^l, \sigma_{q^l}\}_{l \in \{k,n,m\}} \right) (\hat{\mathbf{S}}) \end{aligned}$$

¹⁹Once again, for clarity, we include all the wealth shares in the state space although technically there is redundancy because by definition they sum to 1.

satisfy the following conditions:

(i) *Optimization:* The first order conditions:

$$0 = u'(\eta_j a) - \xi_j(a, \hat{\mathbf{S}}), \quad \forall j \in \{h, b, f\} \quad (\text{C.2})$$

$$0 = r^k(\hat{\mathbf{S}}) - r^d(\hat{\mathbf{S}}) + \sigma_{\xi_h}(a, \hat{\mathbf{S}})\sigma_{q^k}(\hat{\mathbf{S}}) - \partial_{\theta_h^k} \psi_{h,k}(\theta_h^k, \omega_h) \quad (\text{C.3})$$

$$0 = r_h^n(\hat{\mathbf{S}}) - r^d(\hat{\mathbf{S}}) + \sigma_{\xi_h}(a, \hat{\mathbf{S}})\sigma_{q^n}(\hat{\mathbf{S}}) + \lambda(1/q^n(\hat{\mathbf{S}}) - 1)\mathcal{U}'\left(\left(1 + (1/q^n(\hat{\mathbf{S}}) - 1)\theta_h^n\right)a\right) / \xi_h \quad (\text{C.4})$$

$$0 = r^k(\hat{\mathbf{S}}) - r^d(\hat{\mathbf{S}}) + \sigma_{\xi_b}(a, \hat{\mathbf{S}})\sigma_{q^k}(\hat{\mathbf{S}}) - \partial_{\theta_b^k} \psi_{b,k}(\theta_b^k) \quad (\text{C.5})$$

$$0 = r^m(\hat{\mathbf{S}}) - r^d(\hat{\mathbf{S}}) + \sigma_{\xi_b}(a, \hat{\mathbf{S}})\sigma_{q^m}(\hat{\mathbf{S}}) - \partial_{\theta_b^m} \psi_{b,m}(\theta_b^m) \quad (\text{C.6})$$

$$0 = r^k(\hat{\mathbf{S}}) - r_f^n(\hat{\mathbf{S}}) + \sigma_{\xi_f}(a, \hat{\mathbf{S}})(\sigma_{q^k}(\hat{\mathbf{S}}) - \sigma_{q^n}(\hat{\mathbf{S}})) - \partial_{\theta_f^k} \psi_{f,k}(\theta_f^k) \quad (\text{C.7})$$

$$0 = r^m(\hat{\mathbf{S}}) - r_f^n(\hat{\mathbf{S}}) + \sigma_{\xi_f}(a, \hat{\mathbf{S}})(\sigma_{q^m}(\hat{\mathbf{S}}) - \sigma_{q^n}(\hat{\mathbf{S}})) - \partial_{\theta_f^m} \psi_{f,m}(\theta_f^m) \quad (\text{C.8})$$

$$0 = \Phi'(\iota) - 1/q^k(\hat{\mathbf{S}}),$$

and the Euler equations:

$$\rho_h + \lambda_h = \mu_{\xi_h}(a, \hat{\mathbf{S}}) + r^d(\hat{\mathbf{S}}) - \tau_h - \psi_{h,k}(\theta_h^k, \omega_h) - \partial_{\omega_h} \psi_{h,k}(\omega_h^k, \omega_h)\omega_h + \partial_{\theta_h^k} \psi_{h,k}(\theta_h^k, \omega_h)\theta_h^k \quad (\text{C.9})$$

$$\rho_b + \lambda_b = \mu_{\xi_b}(a, \hat{\mathbf{S}}) + r^d(\hat{\mathbf{S}}) - \tau_b - \sum_{l \in \{k, m\}} \left(\psi_{b,l}(\theta_b^l) - \partial_{\theta_b^l} \psi_{b,l}(\theta_b^l) \right) \quad (\text{C.10})$$

$$\rho_f + \lambda_f = \mu_{\xi_f}(a, \hat{\mathbf{S}}) + r_f^n(\hat{\mathbf{S}}) - \tau_h + \sigma_{\xi_f}(\mathbf{S})\sigma_{q^n}(\hat{\mathbf{S}}) - \sum_{l \in \{k, m\}} \left(\psi_{f,l} - \partial_{\theta_f^l} \psi_{f,l}(\theta_f^l) \right) \quad (\text{C.11})$$

where the drift and volatility of ξ_j are characterized by Itô's lemma:

$$\begin{aligned} \mu_{\xi_j}(a, \hat{\mathbf{S}})\xi_j(a, \hat{\mathbf{S}}) &= (D_{\hat{\mathbf{X}}}\xi_j(a, \hat{\mathbf{S}}))^T \boldsymbol{\mu}_{\hat{\mathbf{X}}}(a, \hat{\mathbf{S}}) \\ &\quad + \frac{1}{2} \text{tr} \left\{ (\boldsymbol{\sigma}_{\hat{\mathbf{X}}}(a, \hat{\mathbf{S}}) \odot \hat{\mathbf{X}})^T (\boldsymbol{\sigma}_{\hat{\mathbf{X}}}(a, \hat{\mathbf{S}}) \odot \hat{\mathbf{X}}) D_{\hat{\mathbf{X}}}^2 \xi_j(a, \hat{\mathbf{S}}) \right\} \end{aligned} \quad (\text{C.12})$$

$$\sigma_{\xi_j}(a, \hat{\mathbf{S}})\xi_j(a, \hat{\mathbf{S}}) = (\boldsymbol{\sigma}_{\hat{\mathbf{X}}}(a, \hat{\mathbf{S}}) \odot \hat{\mathbf{X}})^T D_{\hat{\mathbf{X}}}\xi_j(a, \hat{\mathbf{S}}), \quad (\text{C.13})$$

(ii) *State evolution:* The evolution of K_t is given by: $dK_t = (\phi(\iota_t)K_t - \delta K_t)dt$. The evolution of $\Omega_{j,t} = A_{j,t}/A_t$ for any financial intermediary $j \in \{b, f\}$ is given by:

$$\frac{d\Omega_{j,t}}{\Omega_{j,t}} = \left(\mu_{A_j}(\hat{\mathbf{S}}_t) - \mu_A(\hat{\mathbf{S}}_t) + (\sigma_A(\hat{\mathbf{S}}_t) - \sigma_{A_j}(\hat{\mathbf{S}}_t))\sigma_A(\hat{\mathbf{S}}_t) \right) dt + (\sigma_{A_j}(\hat{\mathbf{S}}_t) - \sigma_A(\hat{\mathbf{S}}_t))dW_t \quad (\text{C.14})$$

where:

$$\begin{aligned}\mu_{A_j}(\hat{\mathbf{S}}_t) &= r^j(\hat{\mathbf{S}}_t) + \sum_{l \in \{m, k\}} \theta_j^l(\hat{\mathbf{S}}_t)(r^l(\hat{\mathbf{S}}_t) - r^j(\hat{\mathbf{S}}_t)) - \eta_j(\hat{\mathbf{S}}_t) - \tau + \lambda_j \left(\frac{\bar{\phi}_j}{\Omega_{j,t}} - 1 \right) \\ \sigma_{A_j}(\hat{\mathbf{S}}_t) &= \theta_j^k(\hat{\mathbf{S}}_t)\sigma_{q^k}(\hat{\mathbf{S}}_t) + \theta_j^m(\hat{\mathbf{S}}_t)\sigma_{q^m}(\hat{\mathbf{S}}_t) + (1 - \theta_j^k(\hat{\mathbf{S}}_t) - \theta_j^m(\hat{\mathbf{S}}_t))\sigma_{q^j}(\hat{\mathbf{S}}_t)\end{aligned}$$

The evolution of $\omega_{i,t}$ for $i \leq I$ is given by equation (C.14).

(iii) *Market clearing conditions:*

$$\sum_{i \leq I} \eta(a_i, \hat{\mathbf{S}})\omega_i + \sum_{j \in \{b, f\}} \eta_j(\hat{\mathbf{S}})\Omega_j + \lambda_h \sum_{i \leq I} \left(1 + \left(\frac{1}{q^n(\hat{\mathbf{S}})} - 1 \right) \theta_h^n(a_i, \hat{\mathbf{S}}) \right) \omega_i = \frac{(e^z - \iota)K}{A(\hat{\mathbf{S}})} \quad (\text{C.15})$$

$$\sum_{i \leq I} \theta_h(a_i, \hat{\mathbf{S}}) + \theta_f(\hat{\mathbf{S}})\Omega_f + \theta_b(\hat{\mathbf{S}})\Omega_b = \frac{q^k(\hat{\mathbf{S}})K}{A(\hat{\mathbf{S}})} \quad (\text{C.16})$$

$$\sum_{i \leq I} \theta_h^n(a_i, \hat{\mathbf{S}})\omega_i + \theta_f^n(\hat{\mathbf{S}})\Omega_f = 0 \quad (\text{C.17})$$

$$\sum_{i \leq I} \theta_h^d(a_i, \hat{\mathbf{S}})\omega_i + \theta_b^d(\hat{\mathbf{S}})\Omega_b = 0 \quad (\text{C.18})$$

$$\theta_b^m(\hat{\mathbf{S}})\Omega_b + \theta_f^m(\hat{\mathbf{S}})\Omega_f = 1 - \frac{q^k(\hat{\mathbf{S}})K}{A(\hat{\mathbf{S}})} \quad (\text{C.19})$$

where $A(\hat{\mathbf{S}}) := q^k(\hat{\mathbf{S}})K + q^m(\hat{\mathbf{S}})M$ and the price processes for the long term asset prices must satisfy consistency with Itô's Lemma for $l \in \{k, n, m\}$:

$$\mu_{q^l}(\hat{\mathbf{S}})q^l(\hat{\mathbf{S}}) = (D_{\hat{\mathbf{X}}}q^l(\hat{\mathbf{S}}))^T \boldsymbol{\mu}_{\hat{\mathbf{S}}}(\hat{\mathbf{S}}) + \frac{1}{2} \text{tr} \left\{ (\boldsymbol{\sigma}_{\hat{\mathbf{S}}}(\hat{\mathbf{S}}) \odot \hat{\mathbf{S}})^T (\boldsymbol{\sigma}_{\hat{\mathbf{S}}}(\hat{\mathbf{S}}) \odot \hat{\mathbf{S}}) D_{\hat{\mathbf{S}}}^2 q^l(\hat{\mathbf{S}}) \right\} \quad (\text{C.20})$$

$$\boldsymbol{\sigma}_{q^l}(\hat{\mathbf{S}})q^l(\hat{\mathbf{S}}) = (\boldsymbol{\sigma}_s(\hat{\mathbf{S}}) \odot \hat{\mathbf{S}})^T (D_s q^l(\hat{\mathbf{S}})) \quad (\text{C.21})$$

C.3 Equilibrium Functions

We express the equilibrium objects as functions of the aggregate state vector $\hat{\mathbf{S}}$ without any explicit dependence on the agents' idiosyncratic state a . To illustrate this, we focus on the household stochastic discount factor function, which is defined in the partial equilibrium household problem as a function of individual household wealth and the aggregate states $\xi_h(a, \hat{\mathbf{S}})$. In equilibrium we have that the wealth for agent i is $a_i = \omega_i A(\hat{\mathbf{S}})$ where ω_i is the agent i 's share the economy's wealth and $A(\hat{\mathbf{S}}) = q^k(\hat{\mathbf{S}})K + q^m(\hat{\mathbf{S}})M$. This implies that $\Xi_h(\hat{\mathbf{S}}) = \xi_h(a, \hat{\mathbf{S}})|_{a=\omega A(\hat{\mathbf{S}})}$. The difficulty is that for some variables, namely the agent SDFs, we need to take derivatives. This is

complicated because $D_{\hat{\mathbf{X}}} \xi_h(a, \hat{\mathbf{S}})|_{a=\eta_h A(\hat{\mathbf{S}})} \neq D_{\hat{\mathbf{S}}} \Xi_h(\hat{\mathbf{S}})$ for the obvious reason that the dimension is different. We resolve the formal connection in Proposition 3 below.

Proposition 3. *In equilibrium, we have that for $j \in \{h, b, f\}$:*

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

Proof. We show the result for the volatility. The working for the drift term is analogous. For the volatility, we have that:

$$\begin{aligned} \sigma_{\xi_j}(a, \hat{\mathbf{S}}) \xi_j(a, \hat{\mathbf{S}}) &= (\sigma_{\hat{\mathbf{X}}} \odot \hat{\mathbf{X}})^T (D_{\hat{\mathbf{X}}} \xi_j(a, \hat{\mathbf{S}})) \\ &= \partial_a \xi_j(a, \hat{\mathbf{S}}) \sigma_a(a, \hat{\mathbf{S}}) a + \partial_z \xi_j(a, \hat{\mathbf{S}}) \sigma_z(z) z \\ &\quad + \sum_{i \leq I} \partial_{\omega_i} \xi_h(a, \hat{\mathbf{S}}) \sigma_{\omega_i}(\hat{\mathbf{S}}) \omega_i + \sum_{j \leq \{b, f\}} \partial_{\omega_j} \xi_h(a, \hat{\mathbf{S}}) \sigma_{\Omega_j}(\hat{\mathbf{S}}) \omega_j \end{aligned}$$

After imposing equilibrium $a = \omega_1 A(\hat{\mathbf{S}})$ (with the normalization that a refers to the first agent in the finite population vector), where $A(\hat{\mathbf{S}}) = q^k(\hat{\mathbf{S}})K + q^m(\hat{\mathbf{S}})M$, we have that the RHS is:

$$\begin{aligned} RHS &= \partial_a \xi_j(\omega_1 A(\hat{\mathbf{S}}), \hat{\mathbf{S}}) \sigma_a(\omega_1 A(\hat{\mathbf{S}}), \hat{\mathbf{S}}) \omega_1 A(\hat{\mathbf{S}}) + \partial_z \xi_j(\omega_1 A(\hat{\mathbf{S}}), \hat{\mathbf{S}}) \sigma_z(z) z \\ &\quad + \sum_{i \leq I} \partial_{\omega_i} \xi_h(\omega_1 A(\hat{\mathbf{S}}), \hat{\mathbf{S}}) \sigma_{\omega_i}(\hat{\mathbf{S}}) \omega_i + \sum_{j \leq \{b, f\}} \partial_{\omega_j} \xi_h(\omega_1 A(\hat{\mathbf{S}}), \hat{\mathbf{S}}) \sigma_{\Omega_j}(\hat{\mathbf{S}}) \omega_j \\ &= \begin{bmatrix} \sigma_a(\omega_1 A(\hat{\mathbf{S}}), \hat{\mathbf{S}}) a \\ \sigma_z(z) z \\ 0 \\ \sigma_{\omega_h, z}(\hat{\mathbf{S}}) \omega_{h, z} \\ \vdots \\ \sigma_{\Omega_f, z}(\hat{\mathbf{S}}) \Omega_{f, z} \end{bmatrix}^T \begin{bmatrix} \partial_a \xi_h(\omega_1 A(\hat{\mathbf{S}}), \hat{\mathbf{S}}) \\ \partial_z \xi_h(\omega_1 A(\hat{\mathbf{S}}), \hat{\mathbf{S}}) \\ \partial_K \xi_h(\omega_1 A(\hat{\mathbf{S}}), \hat{\mathbf{S}}) \\ \partial_a \xi_h(\omega_1 A(\hat{\mathbf{S}}), \hat{\mathbf{S}}) A(s) + \partial_{\omega_h} \xi_h(\omega_1 A(\hat{\mathbf{S}}), \hat{\mathbf{S}}) \\ \vdots \\ \partial_{\Omega_f} \xi_h(\Omega_h A(\hat{\mathbf{S}}), \hat{\mathbf{S}}) \end{bmatrix} \\ &= (\sigma_{\hat{\mathbf{S}}} \odot \hat{\mathbf{S}})^T (D_{\hat{\mathbf{S}}} \Xi_h(\hat{\mathbf{S}})) \\ &= \sigma_{\Xi_h}(\hat{\mathbf{S}}) \Xi_h(\hat{\mathbf{S}}) \end{aligned}$$

□

C.4 Loss Function

We can now finally construct our loss function. We approximate the following functions with neural networks:

$$\begin{aligned}
\hat{\eta}_j : \mathcal{S} &\rightarrow \mathbb{R}, (\hat{\mathbf{S}}, \Theta_{\eta_h}) \mapsto \hat{\eta}_j(\hat{\mathbf{S}}; \Theta_{\eta_h}), \quad \forall j \in \{h, f, b\} \\
\hat{\theta}_h^l : \mathcal{S} &\rightarrow \mathbb{R}, (\hat{\mathbf{S}}, \Theta_{\theta_h}) \mapsto \hat{\theta}_h^l(\hat{\mathbf{S}}; \Theta_{\theta_h^l}), \quad \forall l \in \{k, n\}, \\
\hat{\theta}_f^m : \mathcal{S} &\rightarrow \mathbb{R}, (\hat{\mathbf{S}}, \Theta_{\theta_f}) \mapsto \hat{\theta}_f^m(\hat{\mathbf{S}}; \Theta_{\theta_f^m}), \\
\hat{q}^l : \mathcal{S} &\rightarrow \mathbb{R}, (\hat{\mathbf{S}}, \Theta_{q^l}) \mapsto \hat{q}^l(\hat{\mathbf{S}}; \Theta_{q^l}), \quad \forall l \in \{n, m\} \\
\hat{\mu}_{q^k} : \mathcal{S} &\rightarrow \mathbb{R}, (\hat{\mathbf{S}}, \Theta_{\mu, q^k}) \mapsto \hat{\mu}_{q^k}(\hat{\mathbf{S}}; \Theta_{\mu, q^k}), \\
\hat{\sigma}_{q^l} : \mathcal{S} &\rightarrow \mathbb{R}, (\hat{\mathbf{S}}, \Theta_{\sigma, q^l}) \mapsto \hat{\sigma}_{q^l}(\hat{\mathbf{S}}; \Theta_{\sigma, q^l}), \quad \forall l \in \{k, n, m\}
\end{aligned}$$

Given these neural network approximations, at any state $\mathbf{S}$ we can construct the value of the other equilibrium objects in C.2 from the equilibrium characterization in Subsection C.2 using the following steps:

- (i) Impose the goods market clearing condition (C.15) to calculate $q^k(\hat{\mathbf{S}})$,
- (ii) Impose the asset market clearing conditions (C.16), (C.19), (C.18), (C.17), to get the financial intermediary portfolio choices: $\theta_b^k(\hat{\mathbf{S}})$, $\theta_b^d(\hat{\mathbf{S}})$, $\theta_b^m(\hat{\mathbf{S}})$, and $\theta_f^n(\hat{\mathbf{S}})$,
- (iii) Impose the fund and household budget constraints to get the remaining portfolio choices: $\theta_h^d(\hat{\mathbf{S}})$ and $\theta_f^k(\hat{\mathbf{S}})$ (the bank budget constraint is then imposed through Walras's law),
- (iv) Use the consumption first order conditions (C.2) to reconstruct $\{\xi_j(\hat{\mathbf{S}})\}_{j \in \{h, b, f\}}$,
- (v) Compute $\sigma_{\hat{\mathbf{S}}}(\hat{\mathbf{S}})$ from the state evolution equation,
- (vi) Use the first order conditions (C.5), (C.6), (C.7), and automatic derivatives of $\{\xi_j(\hat{\mathbf{S}})\}_{j \in \{h, b, f\}}$ combined with Itô consistency for ξ volatility (C.13) to compute the capital Sharpe ratio $(r^k(\hat{\mathbf{S}}) - r^d(\hat{\mathbf{S}}))/\sigma_{q^k}(\hat{\mathbf{S}})$, the government bond Sharpe ratio $(r^m(\hat{\mathbf{S}}) - r^d(\hat{\mathbf{S}}))/\sigma_{q^m}(\hat{\mathbf{S}})$, and the pension Sharpe ratio $(r^n(\hat{\mathbf{S}}) - r^d(\hat{\mathbf{S}}))/\sigma_{q^n}(\hat{\mathbf{S}})$,
- (vii) Use the neural network approximation to μ_{q^k} and the agent policy functions to compute $r^k(\mathbf{S})$ and then combine with the capital sharp ratio to compute $r^d(\mathbf{S})$, and
- (viii) Calculate the drifts of the states using (C.14), the drift of SDF $\{\xi_{j \in \{h, b, f\}}\}$ using (C.12), and the drift of q^m and q^n using Itô's lemma (C.20).

After making all these substitutions, we are left with the equations from the loss function constructed in section 3.2: the agent Euler equations ((C.9), (C.10), (C.11)), the first order conditions for the portfolio choices that were approximated by neural networks, ((C.3), (C.4), (C.8)), the Itô consistency equation for the drift of capital prices (C.20), and the Itô consistency equations for price volatility (C.21).

C.5 Simulation

In order to simulate the economy we need to compute the evolution of the household wealth distribution. This is complicated for our finite agent approximation method because the neural network policy rules are equilibrium functions of the positions of the N other agents rather than a continuous density. To overcome this difficulty, we use the trained neural network functions to construct a finite difference approximation to the KFE. This involves two key approximation, which we describe below: fitting a partial equilibrium function for $\mu_a(a|z, K, \Omega_f, \Omega_b, g_h)$ and then reconstructing the finite difference matrices. Throughout, we use the following notation. Let $\underline{a} = (a_m : m \leq M)$ denote the grid in the a -dimension. Let $\underline{g}_t = (g_{m,t} : m \leq M)$ denote the marginal household wealth measure function on the a -grid.

Approximation 1: Fitting the functions $\mu_a(a|z, K, \Omega_f, \Omega_b, g_h)$ and $\sigma_a(a|z, K, \Omega_f, \Omega_b, g_h)$ on the grid points at a given $(z, K, \Omega_f, \Omega_b, g_h)$. We illustrate the procedure for $\mu_a(a|z, K, \Omega_f, \Omega_b, g_h)$. We need to use our a finite agent representation for the equilibrium functions on $\hat{S}$ to approximate the partial equilibrium function $\mu_a(\cdot|\cdot)$. To do this, we exploit the law of iterated expectations that:

$$\mu_a(a_i|z, \Omega_f, \Omega_b, g_h) = \int_{\hat{\omega}} \mu(a_i|z, \Omega_b, \Omega_f, \hat{\omega}) g_h(d\hat{\omega})$$

where $\hat{\omega} := (\hat{\omega}_1, \dots, \hat{\omega}_I)$ and $g_h(d\hat{\omega})$ denotes the density for an I agent draw from g_h . We approximate the RHS at each aggregate state $(z, K, \Omega_f, \Omega_b, g_h)$ using the following:

- Draw N_{fit} I -agent vectors, $\{\hat{\omega}\}_{n=1}^{N_{fit}}$, from g_h and calculate $\mu(a|z, K, \Omega_b, \Omega_f, \hat{\omega})$ for each draw. This gives a sample:

$$\{\mu(\hat{\omega}_1^n A(z, K, \Omega_b, \Omega_f, \hat{\omega}^n) | z, K, \Omega_b, \Omega_f, \hat{\omega}^n), \\ \dots, \mu(\hat{\omega}_I^n A(z, K, \Omega_b, \Omega_f, \hat{\omega}^n) | z, K, \Omega_b, \Omega_f, \hat{\omega}^n)\}_{n=1}^{N_{fit}}$$

- Fit a cubic spline fitting for relationship between a and $\mu(a|z, K, \Omega_b, \Omega_f, g_h)$. Denote the fitted function $\hat{f}_{\mu_a}(\cdot|z, K, \Omega_b, \Omega_f, g_h)$.




Figure 8: This figure presents the histogram of difference in the impulse responses computed using the KFE and the N-agent simulation. The left and right panel plots the level differences in the wealth share of fund and the wealth share of bank, respectively. The right panel presents the level differences in net asset holdings of fund (net of bank asset holdings).

- Compute the drifts at grid points: $\mu_{g,m} = \hat{f}_{\mu_a}(a_m|z, K, \Omega_b, \Omega_f, g_h), m \leq M$.

The volatility $\sigma_a(\cdot|\cdot)$ is fitted in an analogous way.

Approximation 2: Constructing an average KFE operator at a given $(z, K, \Omega_f, \Omega_b, g_h)$. Draw N_{sim} I-agent vectors $\{\hat{\mathbf{a}}_t^k = (a_i : 1 \leq i \leq I)\}_{k=1}^{N_{sim}}$. For each draw $k \leq N_{sim}$, the KFE is replaced by the following finite difference equation:

$$dg_{m,t} = \mu_{g,m}(\hat{\mathbf{a}}_t^k)dt + \sigma_{g,m}^T(\hat{\mathbf{a}}_t^k)dW_t, \quad m \leq M \quad (\text{C.22})$$

where the drift at grid 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 $\hat{\mathbf{a}}_t^k$. We then compute an average transition matrix, which we use for simulation.

Simulation algorithm: Algorithm 2 summarizes how approximations 1 and 2 are used to simulate the evolution of the household wealth measure function.

For comparison, we also simulate the baseline economy using the I-agent equilibrium functions (referred to henceforth as ‘I-agent simulation’). Specifically, we simulate using an Euler scheme for 500 years with a monthly frequency and compute the ergodic mean wealth shares of all N-agents. We then compute impulse responses from these ergodic states to a negative TFP shock and compare them with the responses from the ergodic states obtained using the KFE method. Figure 8 presents the difference between the two.

Algorithm 2: Simulating Paths

Input : Initial distribution, neural network approximations to the policy and price functions, number of agents I , 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 simulated path $(z, K, \Omega_f, \Omega_b, g) = \{(z_t, K_t, \Omega_{f,t}, \Omega_{b,t}, g_{h,t})\}_{t=0, \Delta t, \dots, N_T \Delta t}$

for $n = 0, \dots, N_T - 1$ **do**

At iteration n , the time is $t = n\Delta$ and the current state is $(z_t, K_t, \Omega_{f,t}, \Omega_{b,t}, g_{h,t})$.

for $k = 1, \dots, N_{sim}$ **do**

Draw states for I agents $\{\omega_i : i = 1, \dots, I\}$ from density g_t at $t = n\Delta t$.

Given state current state, compute equilibrium prices and returns.

Given the current state, use Approximation 1 to fit the functions approximating the drift and volatility of a : $\hat{f}_{\mu_a}(\cdot | z_t, K_t, \Omega_{f,t}, \Omega_{b,t}, g_{h,t})$ and

$\hat{f}_{\sigma_a}(\cdot | z_t, K_t, \Omega_{f,t}, \Omega_{b,t}, g_{h,t}, h)$,

Using $\hat{f}_{\mu_a}$ and $\hat{f}_{\sigma_a}$, at each grid point $a_m \in \underline{a}$, calculate the consumption and portfolio choices.

Use Approximation 2 to construct the transition matrix $\mathcal{A}_{t,k}$ using a finite difference approximation on the grid $\underline{a}$, as described by (C.22).

end

Take the average: $\bar{\mathcal{A}}_t = \frac{1}{N_{sim}} \sum_{k=1}^{N_{sim}} \mathcal{A}_{t,k}$

Sample ΔW_t from the normal distribution $N(0, \Delta t)$ and construct TFP next period using $z_{t+\Delta t} = z_t + \eta(\bar{z} - z_t)\Delta t + \sigma \Delta W_t$.

Update $(K_t, \Omega_{f,t}, \Omega_{b,t})$ to $(K_{t+\Delta t}, \Omega_{f,t+\Delta t}, \Omega_{b,t+\Delta t})$ using their respective laws of motion.

Update g_t by implicit method: $g_{t+\Delta t} = (I - \bar{\mathcal{A}}_t^T \Delta t)^{-1} g_t + \sigma^T \Delta W_{t,z}$

end

D Training Algorithm and Errors (Internet Appendix)

The full algorithm is provided in Algorithm 3 below. Figure 9 presents the Euler equation training MSE in marginal utility units over epochs in the top left panel, along with the validation loss in the top right panel. The algorithm converges successfully with the validation loss in the order 10^{-4} . The error saturates after 60k epochs, where the gradient norm of all neural network parameters stabilizes.

E Discussion of The Loss Function (Internet Appendix)

In this section, we explain and illustrate the advantages of our loss function outlined in Section 3.2 and derived in Appendix C. We start in Subsection E.1 by demonstrating the benefit of working with the equilibrium functions rather than the partial equilibrium functions in a representative agent Lucas tree economy. This analysis is extended in Subsection E.2 to an economy with heterogeneous agents. Then in Subsection E.3 we show that parameterizing consumption-to-wealth ratios helps the algorithm to converge and avoid undesirable approximate “cheat” solutions. Finally, in Subsection E.4 we illustrate the benefit of imposing market clearing explicitly.

E.1 Working With Equilibrium Functions

We demonstrate the usefulness of imposing general equilibrium explicitly by solving a representative agent Lucas asset pricing economy under different specifications of the loss function. We use this model because we can compare directly to an analytical solution. This can be interpreted as a version of the “big-K, little-k” approach in chapter 7 of [Sargent and Ljungqvist \(2000\)](#).

Economic model. The economy is in continuous time with infinite horizon. There is one type of consumption good. The economy is populated by a representative agent price taking agent who enjoys flow utility $u(c_t)$ from consuming goods flow c_t and has discount rate ρ . There is a unit share of a Lucas tree that generates output flow y_t following a stochastic process $dy_t = \mu(y_t)dt + \sigma(y_t)dW_t$, where W_t is the standard Brownian motion. Shares in the Lucas tree are traded in a competitive market at price q .

Recursive equilibrium characterization: Let a denote the market value of individual wealth held by the representative agent and let A denote the degenerate wealth distribution at point A . Because the representative agent is a price taker, they do not internalize that their individual wealth will ultimately be the aggregate wealth in the economy and so they take the evolution of the degenerate

Algorithm 3: Neural Network Training with Validation and Adaptive Learning

- 1: Initialize neural network objects $(\hat{\eta}_j)_{j \in \{h,f,b\}}$, $(\theta_h^l)_{l \in \{k,n\}}$, θ_f^m , $(\hat{q}^l)_{n,m}$, and $(\hat{\mu}_{q^l}, \hat{\sigma}_{q^l})_{l \in \{k,n,m\}}$ with parameters Θ
- 2: Initialize Adam optimizer and ReduceLROnPlateau scheduler
- 3: Generate validation dataset: $(\hat{S}^{val} = (z^{val}, K^{val}, (\omega_i)_{i \leq I}^{val}, \Omega_b^{val}, \Omega_f^{val}))$
- 4: Set $patience \leftarrow 0$, $max_patience \leftarrow 1000$, $max_epochs \leftarrow 60,000$
- 5: **while** Epochs < max_epochs AND $patience < max_patience$ **do**
- 6: Compute adaptive learning rate lr using inbuilt scheduler
- 7: Sample N new training points: $(\hat{S}^n = (z^n, K^n, (\omega_i)_{i \leq I}^n, \Omega_b^n, \Omega_f^n))_{n=1}^N$
- 8: Calculate equilibrium at each training point $\hat{S}^n$ given current neural network approximations
- 9: Construct the loss as:

$$\hat{\mathcal{L}}(\hat{S}; \Theta) = \left[\sum_{j \in \{h,b,f\}} \mathcal{L}_{\eta_j} + \sum_{l \in \{k,n\}} \mathcal{L}_{\theta_h^l} + \sum_{l \in \{k,m\}} \mathcal{L}_{\theta_f^l} + \mathcal{L}_{\theta_b^m} + \mathcal{L}_{\mu_{q^k}} + \sum_{l \in \{k,n,m\}} \mathcal{L}_{\sigma_{q^l}} \right]$$

where $\hat{\mathcal{L}}_v$ is defined by equations in Subsection 3.2 for each variable v .

- 10: Apply gradient clipping and update parameters: $\Theta \leftarrow \Theta - lr \cdot \nabla_{\Theta} \hat{\mathcal{L}}_{total}$
 - 11: Update learning rate scheduler based on $\hat{\mathcal{L}}_{val}$
 - 12: **if** $epoch \bmod 50 = 0$ **then**
 - 13: Compute validation loss $\hat{\mathcal{L}}_{val}$ on validation set $\hat{S}^{val}$
 - 14: **if** $\hat{\mathcal{L}}_{val} < min_loss$ **then**
 - 15: Update the minimum loss and patience counter $min_loss \leftarrow \hat{\mathcal{L}}_{val}$, $patience \leftarrow 0$
 - 16: Save all model states, optimizer, and scheduler
 - 17: **else**
 - 18: $patience \leftarrow patience + 1$
 - 19: **end if**
 - 20: **end if**
 - 21: IF validation loss variance (normalized by mean) < $1e - 4$ over window, BREAK
 - 22: $epoch \leftarrow epoch + 1$
 - 23: **end while**
 - 24: Save final checkpoint model.
-




Figure 9: The top left panel plots the Euler equation L-2 training loss from the baseline model over iterations on a logarithmic scale. The neural network architecture is 4 hidden layers with 256 neurons in each layer trained using an ADAM optimizer with a learning rate scheduler. The top right panel produces the Euler equation validation loss. The bottom left panel produces other losses and the bottom right panel produces the normalized losses.

wealth distribution A as given. That is, they do not internalize that in equilibrium $a = A$. So, the aggregate state space is (y, A) and individual agent's state space is (a, y, A) . Let the agent's value function be denoted by $V(a, y, A)$ and let the derivative be denoted by $\xi(a, y, A) = \partial_a V(a, y, A)$, which we refer to as the agent's stochastic discount factor.

Following the approach in Section 2 and Appendix A, we can characterize equilibrium by a collection of optimization, distribution, and market clearing blocks. We refer to this as the indirect general equilibrium formulation because we do not impose market clearing in all the equations. Formally, equilibrium consists of a collection of agent optimization functions and price functions:

$$(c, \xi, \mu_\xi, \sigma_\xi)(a, y, A), \quad (q, \mu_q, \sigma_q)(y, A)$$

satisfying the following set of equations:

$$\begin{aligned} [Optimization] \quad & \begin{cases} 0 = -u'(c) + \xi(a, y, A) \\ 0 = -\rho + \mu_\xi(a, y, A) + \frac{y}{q(y, A)} + \mu_q(y, A) + \sigma_\xi(a, y, A)\sigma_q(y, A) \end{cases} \\ [Distribution] \quad & \begin{cases} \mu_a(a, y, A) = \frac{y}{q(y, A)} + \mu_q(y, A) - \frac{c(a, y, A)}{a} \\ \sigma_a(y, A) = \sigma_q(y, A) \\ A = a \end{cases} \\ [Clearing] \quad & \begin{cases} c = y \\ q(y, A) = A \end{cases} \end{aligned}$$

where $\mu_\xi, \mu_q, \sigma_\xi, \sigma_q$ are the geometric drift and volatility of ξ and q , which can be computed by Itô's lemma in space of $\{a, y, A\}$:

$$\begin{aligned} \xi\mu_\xi(a, y, A) &= \partial_a \xi(a, y, A)a\mu_a(a, y, A) + \partial_y \xi(a, y, A)y\mu_y(y) + \partial_A \xi(a, y, A)A\mu_A(y, A) \\ &\quad + 0.5\partial_{aa}\xi(a, y, A)a^2(\sigma_a)^2(a, y, A) + 0.5\partial_{yy}\xi(a, y, A)y^2(\sigma_y)^2 \\ &\quad + 0.5\partial_{AA}\xi(a, y, A)A^2(\sigma_A)^2(y, A) + \partial_{aA}\xi(a, y, A)aA\sigma_a(a, y, A)\sigma_A(y, A) \\ &\quad + \partial_{yA}\xi(a, y, A)yA\sigma_y(y)\sigma_A(y, A) + \partial_{ay}\xi(a, y, A)ay\sigma_a(a, y, A)\sigma_y(y), \\ \xi\sigma_\xi(a, y, A) &= \partial_a \xi(a, y, A)a\sigma_a(a, y, A) + \partial_A \xi(a, y, A)A\sigma_A(y, A) + \partial_y \xi(y\sigma_y(y)), \\ q\mu_q(y, A) &= \partial_A q(y, A)A\mu_A(y, A) + \partial_y q(y, A)y\mu_y(y) + 0.5\partial_{AA}q(y, A)(A\sigma_A)^2 \\ &\quad + 0.5\partial_{yy}q(y, A)(y\sigma_y)^2 + \partial_{Ay}q(y, A)yA\sigma_y\sigma_A \\ q\sigma_q(y, A) &= \partial_A q(y, A)A\sigma_A + \partial_y q(y, A)y\sigma_y \end{aligned}$$

Observe that this characterization essentially writes the conditions for the partial equilibrium optimization problem along with a collection of equilibrium conditions that must be imposed.

Alternatively, we can follow the approach in Appendix C.3 and tackle the general equilibrium problem directly. We refer to this as the direct general equilibrium formulation because we impose market clearing and $a = A$ in all the equations. First, the goods market clearing condition pins down the consumption $c = y$ so the SDF can be written entirely as a function of output: $\xi(a, y, A) = u'(y)$. Second, the asset market clearing condition $q(y, A) = A$ implies that aggregate wealth is an implicit function of output: $A(y)$. This means that, in equilibrium, the Lucas tree price is a function y , i.e., $q(y, A(y))$. Imposing these observations and substituting the market clearing conditions into the Euler equation leads to the equilibrium ODE (or “master equation”) for the asset price $q(y)$:

$$0 = -\rho + \mu_\xi(y) + \left(\frac{y}{q(y)} + \mu_q(y) \right) + \sigma_\xi(y)\sigma_q(y)$$

where $\mu_\xi, \mu_q, \sigma_\xi, \sigma_q$ are geometric drift, volatility of SDF and share price, computed by Itô’s lemma in space of $\{y\}$:

$$\begin{aligned} \xi\mu_\xi &= \partial_y \xi(y) y \mu_y(y) & q\mu_q &= \partial_y q(y) y \mu_y(y) \\ \xi\sigma_\xi &= \partial_y \xi(y) y \sigma_y(y) & q\sigma_q &= \partial_y q(y) y \sigma_y(y) \end{aligned}$$

Numerical solution. We compare numerical calculation using the indirect and direct general equilibrium formulations. In each case, we set up the loss function:

- (i) *Indirect general equilibrium formulation:* Denote the aggregate state as $\mathbf{s} = (y, A) \in \mathbb{R}^{+2}$ and denote $\mathbf{S} := (a, \mathbf{s})$. We parameterize a set of neural network $\hat{\xi}(a, \mathbf{s}; \Theta_\xi)$, $\hat{q}(\mathbf{s}; \Theta_q)$, $\hat{\mu}_q(\mathbf{s}; \Theta_{\mu_q})$, $\hat{\sigma}_q(\mathbf{s}; \Theta_{\sigma_q})$, which need to satisfy the optimization, market clearing, and consistency equations:

$$\begin{aligned} \mathcal{L}_\xi(\mathbf{S}) &= -\rho + \mu_\xi(a, \mathbf{s}) + \frac{y}{q(\mathbf{s})} + \mu_q(\mathbf{s}) + \sigma_\xi(a, \mathbf{s})\hat{\sigma}_q(\mathbf{s}) \\ \mathcal{L}_c(\mathbf{S}) &= (u')^{-1}(\hat{\xi}(a, \mathbf{s})) - y \\ \mathcal{L}_q(\mathbf{S}) &= \hat{q}(\mathbf{s}) - A \\ \mathcal{L}_{\sigma_q}(\mathbf{s}) &= \partial_A q(\mathbf{s}) A \sigma_A + \partial_y q(\mathbf{s}) y \sigma_y - \hat{q} \hat{\sigma}_q(\mathbf{s}) \\ \mathcal{L}_{\mu_q}(\mathbf{s}) &= \partial_A q(\mathbf{s}) A \mu_A(\mathbf{s}) + \partial_y q(\mathbf{s}) y \mu_y(y) + 0.5 \partial_{AA} q(\mathbf{s}) (A \sigma_A)^2 \\ &\quad + 0.5 \partial_{yy} q(\mathbf{s}) (y \sigma_y)^2 + \partial_{Ay} q(\mathbf{s}) y A \sigma_y \sigma_A - \hat{q} \hat{\mu}_q(\mathbf{s}), \end{aligned}$$

where we impose the optimal saving condition: $c(a, \mathbf{s}) = (u')^{-1}(\xi(a, \mathbf{s}))$. The loss function for the deep learning algorithm is:

$$\hat{\mathcal{L}}_1(\mathcal{S}; \Theta) = (\mathcal{L}_c + \mathcal{L}_q + \mathcal{L}_\xi + \mathcal{L}_{\sigma_q} + \mathcal{L}_{\mu_q})(\mathcal{S}; \Theta). \quad (\text{E.1})$$

- (ii) *Direct general equilibrium formulation:* Under this formulation there is only one aggregate state variable $s = y \in \mathbb{R}^+$ so we parameterize the asset price $\hat{q}(s; \Theta'_q)$, which needs to satisfy the Euler equation:

$$\mathcal{L}_q(y) = -\rho + \mu_\xi(y) + \frac{y}{q(y)} + \mu_q(y) + \sigma_\xi(y)\sigma_q(y)$$

and the loss function for the deep learning algorithm is:

$$\mathcal{L}_2(y; \Theta') = (\mathcal{L}_q)(y; \Theta') \quad (\text{E.2})$$

For computation, we assume the utility function takes the form $u(c_t) = \log c_t$, and the output follows geometric Brownian motion: $\mu(y_t) = \mu y_t, \sigma(y_t) = \sigma y_t$. We approximate the price-dividend ratio $f(y; \Theta_f) = \hat{q}(y)/y$ by a neural network. We compare training the neural network to minimize the loss function from the indirect formulation (E.1) and the loss function from the direct formulation (E.2). For each case we train neural networks with the same architecture (aside from the number of inputs) and use the same optimization procedure. The only difference is the choice of loss functions constructed.

Figure 10 reports the training loss on the left subplot and compares it with the analytical solution $q(y, A) = y/\rho, A = y/\rho$ on the right subplot. We find that the neural network is less able to solve the problem when minimizing the loss function from the indirect formulation: the L-1 loss is higher and the price function deviates substantially from the analytical solution. In addition, minimizing the implicit problem is less stable. During the epochs between approximately 3000 and 7000, the price curve fits the the analytical benchmark closely while other equilibrium conditions are off. This switches around later in the training and the price curve stops fitting the analytical benchmark closely.

One of reasons why training fails for the indirect formulation of the problem is because the sampling in the partial equilibrium wealth space is inefficient. To understand this, note that in equilibrium, total wealth is equal to total asset supply:

$$q(y, A) = A.$$




Figure 10: Solution to the representative Lucas Tree model with log utility for two parameterizations: the implicit GE View (orange) and the explicit GE View (blue). Left: L-1 training loss; Right: learned $q(y)$ (Explicit GE view) and implied $q(\bar{a}, y)$ (Implicit GE view, $\bar{a}$ clears asset market) compared to the analytical benchmark: $q = y/\rho$.

Given the price function q , the solution to the fixed point problem, if exists, is denoted by $A^*(y)$. Therefore, in equilibrium, the price can be written implicitly as a function of y , i.e., $q(y) \equiv q(y, A^*(y))$. Training the neural network to learn the price as a function of both individual wealth and output y requires understanding a partial equilibrium relationship that is not necessary for computing equilibrium prices.

E.2 Imposing General Equilibrium in Heterogeneous Agent Economies and Working in the Wealth Share Space

In Subsection E.1 we were easily able to impose equilibrium because the economy had a representative agent. We now discuss how to do this in a model with heterogeneous agents and explain why working in the wealth share space makes this easier.

Environment. Consider the same environment as is Section E.1 but now suppose there are I types of price taking agent who enjoys flow utility $u_i(c_{it})$ from consuming goods flow c_{it} and has discount rate ρ_i . The agents can now also trade risk free bonds in zero net supply.

Recursive representation in the wealth level space. Let a_i denote the individual wealth of agent i . The wealth distribution is $\check{g} = \{a_1, \dots, a_N\}$. Let each agent's value function be denoted by

$V_i(a_i, y, \check{g})$, and let the derivative be denoted by $\xi_i(a_i, y, \check{g}) = \partial_{a_i} V_i(a_i, y, \check{g})$. Following the approach in Section 2 and Appendix A, we can again characterize equilibrium by a collection of optimization, distribution, and market clearing blocks. Formally, equilibrium consists of a collection of agent optimization functions and price functions:

$$(c_i, \xi_i, \theta_{i,b}, \mu_{\xi_i}, \sigma_{\xi_i})_{i \leq I}(a, y, \check{g}), \quad (q, \mu_q, \sigma_q)(y, \check{g})$$

satisfying the following set of equations:

$$\begin{aligned} [\text{Optimization}] \quad & \begin{cases} 0 = -u'_i(c) + \xi_i(a, y, \check{g}) \\ 0 = -\rho + \mu_{\xi_i}(a, y, \check{g}) + \frac{y}{q(y, \check{g})} + \mu_q(y, \check{g}) + \sigma_{\xi_i}(a, y, \check{g})\sigma_q(y, \check{g}) \\ 0 = \frac{y}{q(y, \check{g})} + \mu_q(y, \check{g}) + \sigma_{\xi_i}(a, y, \check{g})\sigma_q(y, \check{g}) - r_f(y, \check{g}) \end{cases} \\ [\text{Distribution}] \quad & \begin{cases} \mu_a(a, y, \check{g}) = \theta_b(a, y, \check{g})r_f + (1 - \theta_b(a, y, \check{g})) \left(\frac{y}{q(y, \check{g})} + \mu_q(y, \check{g}) \right) - \frac{c(a, y, \check{g})}{a} \\ \sigma_a(y, \check{g}) = \sigma_q(y, \check{g}) \end{cases} \\ [\text{Clearing}] \quad & \begin{cases} \sum_{i=1}^I c_i(a, y, \check{g}) = y \\ q(y, \check{g}) = \sum_{i=1}^I a_i \end{cases} \end{aligned}$$

where $\mu_{\xi_i}, \mu_q, \sigma_{\xi_i}, \sigma_q$ are once again the geometric drift, volatility of SDF and share price, which can be computed by Itô's lemma in space of $\{a_1, \dots, a_I, y\}$, which is denoted as $\mathbf{s}$:

$$\begin{aligned} \xi_i \mu_{\xi_i}(a_i, \mathbf{s}) &= (\boldsymbol{\mu}_s(\mathbf{s}) \odot \mathbf{s})^T D_s \xi_i + \mu_{a_i}(\mathbf{s}) a_i D_{a_i} \xi_i + \sigma_{a_i}(\mathbf{s}) a_i (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s})^T D_{a_i, s} \xi_i \\ &\quad + 0.5(\sigma_{a_i}(\mathbf{s}) a_i)^2 D_{a_i}^2 \xi_i + 0.5 \text{tr} \left\{ (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s})^T (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s}) D_s^2 \xi_i \right\} \\ \xi_i \sigma_{\xi_i}(a_i, \mathbf{s}) &= (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s})^T D_s \xi_i + \sigma_{a_i}(\mathbf{s}) a_i D_{a_i} \xi_i \\ \xi_i \mu_q(\mathbf{s}) &= (\boldsymbol{\mu}_s(\mathbf{s}) \odot \mathbf{s})^T D_s q + 0.5 \text{tr} \left\{ (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s})^T (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s}) D_s^2 q \right\} \\ \xi_i \sigma_q(\mathbf{s}) &= (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s})^T D_s q \end{aligned}$$

Recursive representation in the wealth share space: Let ω_i denote the individual wealth share of agent i . The wealth distribution is $g = \{\omega_1, \dots, \omega_{I-1}\}$. The direct general equilibrium formulation

consists a collection of optimization functions and price functions:

$$(c_i, \xi_i, \theta_{i,b}, \mu_{\xi_i}, \sigma_{\xi_i})_{i \leq I}(g, y), \quad (q, \mu_q, \sigma_q)(g, y)$$

satisfying the following set of equations:

$$\begin{aligned} [\textit{Optimization}] \quad & \begin{cases} 0 = -u'_i(c) + \xi_i(g, y) \\ 0 = \frac{y}{q(g, y)} + \mu_q(g, y) - r_f + \sigma_{\xi_i}(g, y)\sigma_q(g, y) \\ 0 = -\rho + \mu_{\xi_i}(g, y) + r_f \end{cases} \\ [\textit{Distribution}] \quad & \begin{cases} \mu_{\omega_i}(g, y) = \frac{y}{q(g, y)} + \theta_{i,b} \left[r_f - \left(\frac{y}{q(g, y)} + \mu_q(g, y) \right) \right] - \frac{c_i}{\omega_i q(g, y)} + \theta_{i,b}(\sigma_q)^2 \\ \sigma_{\omega_i}(g, y) = -\sigma_q(g, y)\theta_{i,b} \end{cases} \\ [\textit{Clearing}] \quad & \begin{cases} \sum_{i=1}^I c_i(g, y) = y \\ \sum_{i=1}^I \omega_i = 1 \\ \sum_{i=1}^I \omega_i \theta_{i,b} = 0 \end{cases} \end{aligned}$$

where $\mu_{\xi_i}, \mu_q, \sigma_{\xi_i}, \sigma_q$ are once again the geometric drift, volatility of SDF and share price, which can be computed by Itô's lemma in space of $\{\omega_1, \dots, \omega_{I-1}, y\}$, which is denoted as $\mathbf{s}$:

$$\begin{aligned} \xi_i \mu_{\xi_i}(\mathbf{s}) &= (\boldsymbol{\mu}_s(\mathbf{s}) \odot \mathbf{s})^T D_s \xi_i + 0.5 \text{tr} \left\{ (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s})^T (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s}) D_s^2 \xi_i \right\} \\ \xi_i \sigma_{\xi_i}(\mathbf{s}) &= (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s})^T D_s \xi_i \\ \xi_i \mu_q(\mathbf{s}) &= (\boldsymbol{\mu}_s(\mathbf{s}) \odot \mathbf{s})^T D_s q + 0.5 \text{tr} \left\{ (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s})^T (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s}) D_s^2 q \right\} \\ \xi_i \sigma_q(\mathbf{s}) &= (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s})^T D_s q \end{aligned}$$

Imposing market clearing. To impose asset market clearing explicitly, we treat the wealth of one agent as a residual, e.g.,

$$a_I = q(a_1, \dots, a_I; y) - \sum_{i=1}^{I-1} a_i.$$

Given price function $q(\cdot, y)$, and the wealth level of first $I - 1$ agents, the asset market clearing con-

dition implies: $a_I = a_I^*(a_1, \dots, a_{I-1}, y)$. The existence of a_I^* relies on the implicit function theorem. Let $F(a_1, \dots, a_I, y) = \sum_{i=1}^I a_i - q(a_1, \dots, a_I, y)$. If q is C^1 and at the point of interest $\frac{\partial q(a_1, \dots, a_I, y)}{\partial a_I} \neq 1$, then there exists a (locally unique) C^1 function a_I^* such that $F(a_1, \dots, a_{I-1}, a_I^*(\cdot), y) = 0$. Imposing directly the asset market clearing condition is equivalent to parameterizing price as function of $(a_1, \dots, a_{I-1}, y)$ and enforcing that the last agent takes the residual. In practice, the sampling region should be more restricted in order to make sure the every agent's wealth is greater than 0.

We now show how moving to the wealth share space allows us to use this approach to impose market clearing more easily. Denote the wealth share of agent $i \leq I$ by ω_i . Then, the asset market clearing condition can be written as:

$$q\left(\sum_{i=1}^I \omega_i\right) = q$$

From this, we can see that simple algebra can back out the last agent's wealth share by $\omega_I = 1 - \sum_{i=1}^{I-1} \omega_i$ whereas, in the wealth space, the last agent's wealth depends on the implicit function a_I^* .

Numerical Solution. We solve the problem numerically in the wealth share space. We parameterize a set of neural network $\{\hat{\eta}_i(\mathbf{s}; \Theta_{\eta_i})\}_{i=1}^I$, $\{\hat{\theta}_{i,b}(\mathbf{s}; \Theta_{\theta_{i,b}})\}_{i=1}^{I-1}$, $\hat{\sigma}_q(\mathbf{s}; \Theta_{\sigma_q})$, $\hat{\mu}_q(\mathbf{s}; \Theta_{\mu_q})$ which need to satisfy the optimization, and consistency conditions:²⁰

$$\begin{aligned} \mathcal{L}_{\xi_i}(\mathbf{s}) &= -\rho + \mu_{\xi_i} + r_f \\ \mathcal{L}_{b,i}(\mathbf{s}) &= \frac{y}{q(\mathbf{s})} + \mu_q(\mathbf{s}) - r_f + \sigma_{\xi_i}(\mathbf{s})\sigma_q(\mathbf{s}) \\ \mathcal{L}_{\sigma_q}(\mathbf{s}) &= \frac{(\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s})^T D_{\tilde{s}} \left(\frac{y}{\sum_{i=1}^I \hat{\eta}_i \omega_i} \right)}{\left(\frac{y}{\sum_{i=1}^I \hat{\eta}_i \omega_i} \right)} - \hat{\sigma}_q \\ \mathcal{L}_{\mu_q}(\mathbf{s}) &= \frac{(\boldsymbol{\mu}_s(\mathbf{s}) \odot \mathbf{s})^T D_{\tilde{s}} \left(\frac{y}{\sum_{i=1}^I \hat{\eta}_i \omega_i} \right) + 0.5 \text{tr} \left\{ (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s})^T (\boldsymbol{\sigma}_s(\mathbf{s}) \odot \mathbf{s}) D_s^2 \left(\frac{y}{\sum_{i=1}^I \hat{\eta}_i \omega_i} \right) \right\}}{\left(\frac{y}{\sum_{i=1}^I \hat{\eta}_i \omega_i} \right)} - \hat{\mu}_q \end{aligned}$$

where we impose the optimal saving condition and the goods market clearing condition: $(u'_i)^{-1}(\xi_i(\mathbf{s})) =$

²⁰In simple cases, it is not essential to parameterize $\{\hat{\theta}_{i,b}\}_{i=1}^{I-1}$, $\hat{\sigma}_q$, $\hat{\mu}_q$ as we can solve for the portfolio weight and spreads in closed-form, then express price drift and volatility via Itô's lemma. See more details in subsection G.1's code. In general economic problems, however, we need to parameterize these objects.

$\frac{\hat{\eta}_i \omega_i}{\sum_{i=1}^I \hat{\eta}_i \omega_i} y$. The loss function for the deep learning algorithm is:

$$\hat{\mathcal{L}}(\mathbf{s}; \Theta) = \left(\sum_{i=1}^I (\mathcal{L}_{\xi_i} + \mathcal{L}_{b,i}) + \mathcal{L}_{\sigma_q} \right) (\mathbf{s}; \Theta)$$

In subsection G.1, we solve the asset price and compare with the analytical benchmark.

Connection to Azinovic and Zemlicka (2026). With the asset market clearing condition plugged in $\omega_I = 1 - \sum_{i=1}^{I-1} \omega_i$, asset price q can be expressed in terms of consumption-to-wealth ratio and wealth shares:

$$q = \frac{1}{\sum_{i=1}^I \eta_i \omega_i} y = f(\omega_1, \dots, \omega_{I-1}, y).$$

This expression nests the price-dividend ratio parameterization's idea and extends to multiple-agent setup. Compared to the representative agent's case, we are imposing extra structure about the goods market encoded into the output layer of neural network of $f(\cdot)$, i.e.,

$$f(\mathbf{s}; \Theta_f) = \frac{1}{\sum_{i=1}^I \eta_i(\mathbf{s}; \Theta_{\eta_i}) \omega_i},$$

echoing the same insight from Azinovic and Zemlicka (2026).

E.3 Portfolio Choice and Parameterizing Consumption-to-Wealth Ratios

The previous Subsections E.1 and E.2 focused on imposing equilibrium in an economy with one asset. We now consider the parametrization of economies with multiple assets and portfolio choice. In simple macro-finance economic models, portfolio choice problems simplify because there are closed-form expressions for the stochastic discount factor. However, in more complicated macro-finance models, like in our paper, we need to solve for the stochastic discount factor numerically, which is well known to be a difficult computational problem. In this subsection, we demonstrate that parameterizing the consumption-to-wealth ratio rather than the stochastic discount factor directly helps the Neural Network to learn the equilibrium.

Environment: Consider a standard continuous time Merton portfolio problem. An agent receives flow utility $u(c_t)$ from consuming goods flow c_t and has discount rate ρ . There are two assets. The first is risky asset with price q_t and return $dR_t^q = r^q dt + \sigma_q dW_t$ where W_t is a Brownian motion process. The second is a short-term risk free bond that pays interest rate r .

Recursive equilibrium characterization: Let $a_t := b_t + q_t k_t$ denote the market value of individual wealth held by the agent, where b_t are holdings of the bond and k_t are holdings of capital. Let $\theta_t := q_t k_t / a_t$ denote the fraction of wealth the agent holds in the risky asset. Let $V(a)$ denote the agent's value function and let $\xi(a) := \partial_a V(a)$. The agent choice and value functions $(c, \theta, \xi)(a)$ solves the first order conditions and Euler equation:

$$\begin{aligned} 0 &= -u'(c) + \xi(a) \\ \frac{qk}{a} &= - \left(\frac{r^q - r}{\sigma_q^2} \right) \frac{\xi(a)}{a\xi_a(a)} \\ 0 &= -\rho + r + \mu_\xi(a) \end{aligned}$$

where:

$$\begin{aligned} \mu_\xi(a)\xi(a) &= \partial_a \xi(a) a \mu_a(a) + 0.5 a^2 \sigma_a^2(a) \partial_{aa} \xi(a), \\ \sigma_\xi(a)\xi(a) &= \partial_a \xi(a) a \sigma_a(a) \end{aligned}$$

Numerical solution. We compare numerical calculation using the parameterization of SDF $\xi = \hat{\xi}(a; \Theta_\xi)$ and parameterization of consumption to wealth ratio, where the SDF is constructed as $\hat{\xi} = (\hat{\eta}(a; \Theta_\eta) a)^{-\gamma}$. In each case, we set up the loss function as follows:

$$\mathcal{L}_\xi(a; \Theta) = -\rho + r + \mu_\xi(a).$$

The relative difference compared to the analytical solution is shown for each case is shown in Figure 11. Evidently, the parameterization of the consumption-to-wealth ratio leads to a neural network approximation that is much closer to the analytical solution.

E.4 Discussion: Market Clearing Conditions

The previous subsections focused on simplified models without production. In this final subsection we return to our full model from Section 2 in the main text with investment and production. We use this model to illustrate the advantage of imposing asset market clearing exactly.

Imposing asset market clearing: In the construction of our loss function in Subsection C, we




Figure 11: Solution to the Merton’s portfolio problem. $\hat{\omega}$: neural network implied consumption to wealth ratio. Blue line: ω parameterization; Orange line: ξ parameterization of SDF; Green dashed line: analytical benchmark $\omega^* = \frac{1}{\gamma}[\rho - (1 - \gamma)(r + \frac{(r^q - r)^2}{2\gamma\sigma_q^2})]$ is closed-form benchmark. (Economic parameters: $r = 1\%$, $r^q = 2\%$, $\sigma_q = 2\%$, $\rho = 5\%$, $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$, $\gamma = 2.0$). The analytical benchmark perfectly aligns with the blue line.

cleared the asset market by letting the last agent take the residual asset portfolio share:

$$\theta_I^j = \frac{\vartheta^j - \sum_{i=1}^{I-1} \omega_i \theta_i^j}{\omega_I}, j \in \{m, n, k\}.$$

where ϑ^j is the share of aggregate wealth in asset j . We refer to this as the “exact” imposition of market clearing because it prevents any approximation error in the market clearing equations and instead forces any numerical approximation error into the first order conditions. Alternatively, we could include an additional loss function for market clearing condition:

$$\mathcal{L}_{\text{clr}} = \sum_{j \in \{m, n, k\}} \left(\sum_{i \in \{h, f, b\}} \Omega_i \theta_i^j - \vartheta_j \right)^2.$$

We refer this as the “approximate” imposition of market clearing because we are allowing the Neural network to make small deviations from asset market clearing.

Numerical implementation: Figure 12 presents the household pension holdings computed using two different methods. In the ‘MC Exact’ method, we impose all market clearing conditions




Figure 12: The figure presents household pension holdings under two methods a) MC exact - market clearing is exactly imposed in the solid line, and b) MC Approximate - market clearing is approximately solved by including them in the loss function in the dashed line. For illustration, we compute the pension holdings at randomly sampled household wealth-share points, keeping fixed the other state variables at their respective long-run average values.

explicitly. In the ‘MC Approximate’ method, we include market clearing in the loss function, which means that it is solved for approximately. For illustration, we display the pension holdings at randomly sampled household wealth share values, keeping the other state variables fixed at their respective long-run average values. The ‘MC Exact’ solution shows a downward sloping pension holding as a function of household wealth share, while the ‘MC Approximate’ solution is flat. We found that, for many different parameterizations, allowing for slack in the market clearing condition leads to an inelastic annuity demand curve. This means that allowing for a slack in the market clearing condition leads to an inaccurate risk-shifter in the pension market FOC (equation 2.10).

F Multiple stochastic steady states (Internet Appendix)

Our model admits multiple stochastic steady states. Figure 13 displays the drift of the bank’s wealth share as a function of its own wealth share. We sample 2000 points from the state space and compute the equilibrium drift using the evolution equation in Theorem 2 in two economies: the baseline economy and an alternate economy where households’ demand for deposits dominates other assets. For illustration, we fit a cubic spline curve to the scatter points. In the baseline economy, there is a non-degenerate stationary density around $\eta_b = 0.23$, as evidenced by the drift

curve crossing zero at that point. In the alternate economy, the drift is increasing in the bank’s wealth share, implying that the bank takes over the economy in the long run.




Figure 13: The left panel plots the drift of bank’s wealth share as a function of η_b in the baseline economy. The right panel plot the same in an alternative economy with high household deposit demand. In both panels, scatter points represent the equilibrium drift from a sample of 2000 randomly drawn points from the state space. The red and blue line represent a cubic spline fit to the points.

G Test Models (Internet Appendix)

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 \(2014\)](#). First we summarize the key results, and then present the model details. 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 9.

Model	Num of Layers	Num of Neurons	Learning Rate
Lucas Tree Model	4	64	0.001
Limited Participation Model	5	64	0.001
Brunnermeier Sannikov Model	4	64	0.0005/0.005

Table 9: Neural network parameters for the three testible models. For Brunnermeier Sannikov model, we set the learning rate for consumption-to-wealth ratio as 0.0005, and the learning rate for the portfolio choice of expert as 0.005.

Table 10 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. Each following subsection describes how the model in that section can be nested with the main model along with technical details.

Method	L2-Error
Complete markets	1.8×10^{-9}
Basak and Cuoco (1998)	3.0×10^{-9}
Brunnermeier and Sannikov (2014)	4.6×10^{-8}

Table 10: Summary of the algorithm performance and computational speed. Error calculates the difference between solution by neural network and finite difference.

G.1 Lucas Asset Tree Model

We solve the Lucas Asset Tree model described in subsection E.2. We assume that each type $i \leq I$ of agents takes CRRA utility form with same relative risk aversion γ and the output process y_t of the tree follows a geometric Brownian motion process.

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

Without financial frictions, there is simple aggregation of individuals' Euler equations, which coincides with the representative agent's pricing equation.

Analytical Solution. In a representative agent's world, by the standard Lucas tree pricing formula, the asset price is determined by the discounted flow of dividends:

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

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):

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

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 that the consumption policy is:

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

Neural Network Solution. Though aggregation results hold, we still incorporate the wealth heterogeneity in our solution algorithm, i.e., we have I asset pricing conditions and I Euler equations. We compare the equilibrium asset price $q(\cdot, y)$ and consumption to wealth ratio $\omega_i(\cdot, y)$ with the “as-if” representative agent economy. The estimated time cost for the model with 5, 10, and 20 agents is about 2 mins, 10 mins, and 20 mins, respectively. The difference between the consumption rule solved using our neural network method and the analytical solution is less than 0.1% for 5 and 10 agents, and 0.5% for 20 agents, respectively. The parameters for the complete market model are provided in Table 12.

Num of Agents	Euler Eq Error	Diff	Time Cost
5	<1e-8	<1e-8	2 mins
10	<1e-8	<1e-8	5 mins
20	<1e-8	<1e-8	20 mins

Table 11: Summary of the algorithm performance and computational speed. “Diff” means the difference between representative agent case’s solution and brute-force. All errors and differences are in mean squared error.




Figure 14: Solution to As-if representative agent model. Left panel: asset price; Middle panel: consumption-wealth ratio of agent 1; Right panel: the distribution of losses when training stops.

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

Table 12: Model parameters.

G.2 Asset Pricing with Restricted Participation

We carry forward modifications to the main model from the previous section to mimic the endowment economy. Consider an infinite-horizon economy with two types of price-taking agents: expert (indexed by e) and household (indexed by h). The financial friction is that households cannot participate in the capital market. Experts do not face this constraint. Mathematically, it is stated as

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

As before, the output y_t follows a geometric Brownian motion process.

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

Finite Difference Solution. We exploit the scalability for geometric Brownian motion's case to get a precise solution by focusing on one-dimensional differential equation. For a 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., $\eta = \eta_e$ (and $1 - \eta = \eta_h$). Denote the consumption-to-wealth ratio as $\omega_j, j \in \{h, e\}$, the value function can be written as:

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

where v_i is the scaled value function. From the first-order condition, we get ²¹

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

From the goods market clearing condition, we have:

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

The *HJB equation* for the scaled value function v_i is given by:

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

where μ_η, σ_η 's expressions are as follows.

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

The price of risk which appears in the asset pricing condition is determined by Itô's Lemma as follows.

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

²¹This expression leads to the boundary condition at $\eta = 1$: $\frac{f(1)}{v_e} = 1$

In finite difference solution approach, we introduce a pseudo time-derivative. (G.2):

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

We then update the value function in an implicit scheme to solve the following 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 $\mathbf{I}$ is the identity matrix.

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., there will be nobody left in the economy to hold 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:

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

which implies the boundary condition: $v_e(1) = \frac{1}{\rho_e + (\gamma - 1)\mu - \frac{1}{2}\gamma(\gamma - 1)\sigma^2}$.

The estimated time to solve the limited participation problem by neural network is about 5 minutes. We compare the finite difference solution with the neural network solution on η 's dimension in figure 15 for $y = 1$. We can see that our method well captures the high non-linearity (left-upper panel) and amplification (right-lower panel). The parameters are provided in Table 13.

Parameter	Symbol	Value
Risk aversion	γ	1.5
Households' Discount rate	ρ_h	0.05
Experts' Discount rate	ρ_h	0.05
Output Growth Rate	μ	2%
Volatility of Growth	σ	5%

Table 13: Parameters for the restricted participation model.




Figure 15: Solution to restricted stock market participation model.

G.3 A Macroeconomic Model with Productivity Gap

The setup in this example follows [Brunnermeier and Sannikov \(2014\)](#). There are two types of agents in this infinite horizon economy: experts and households. Both types can hold capital, but experts have a higher productivity rate compared to households. The productivity rates are given by z_h, z_e ($z_h < z_e$), respectively. Their relative risk aversions are the same, denoted by γ . Output grows exogenously by $\mu_y = y\mu$, with volatility $y\sigma$, and experts cannot issue outside equities. In addition, we assume that households cannot short capital, 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 of households and experts follow

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

The expected capital return is

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

We rewrite the financial friction as the difference : $\frac{y_e - y_h}{q\sigma_q}$. For the first two equations, we have:

$$\begin{cases} -\frac{1}{\xi_e} \frac{\partial \xi_e}{\partial y} \sigma_y = \frac{1}{\xi_e} \frac{\partial \xi_e}{\partial \eta} \sigma_\eta - \frac{r_f - r_{q,h}}{\sigma_q} + \frac{y_e - y_h}{q\sigma_q} \\ -\frac{1}{\xi_h} \frac{\partial \xi_h}{\partial y} \sigma_y = \frac{1}{\xi_h} \frac{\partial \xi_h}{\partial \eta} \sigma_\eta - \frac{r_f - r_{q,h}}{\sigma_q} \end{cases}$$

Unlike the fully restricted participation's case where the experts hold all capital, we have to keep track of the capital allocation ratio of experts κ , which is parameterized as $\kappa = \eta + \lambda = \eta + \mathcal{N}_\lambda \eta^\beta$, where $\mathcal{N}_\lambda$ is a trainable neural network, and $\beta = \frac{1}{2}$ captures the power law for $\eta \rightarrow 0$. Given the expert's capital share holding κ , the volatility of wealth share σ_η is $(\kappa - \eta)\sigma_q$. The goods market clearing condition (G.1) is replaced by

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

and the price volatility is revised as

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

At the left boundary, $f(\cdot)$ is determined by $f(0) = \frac{y_h}{\omega_h(0)}$, $f(1) = \frac{y_e}{\omega_e(1)}$.

The estimated time to solve the model by our neural network method is about 5 minutes. Again, we compare the finite difference solution with the neural network solution in Figure 16 for $y = 1$. We restrict the range of η to be the crisis region in [Brunnermeier and Sannikov \(2014\)](#), which is defined by the region where the capital share of experts $\kappa < 1$. This region captures the fire-sale region, where amplification takes place. We can see that the neural network solution well captures most of the amplification in that crisis region. The parameters are provided in Table 14.




Figure 16: Solution to the model with productivity gap.

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%

Table 14: Parameters for the macroeconomic model.

H Social Mobility

From \ To	Data				Model			
	0-25	25-50	50-75	75-100	0-25	25-50	50-75	75-100
0-25	67.2	24.5	6.4	1.9	82.6	17.2	0.1	0.0
25-50	24.6	49.9	19.2	6.3	17.3	68.9	13.8	0.0
50-75	6.7	18.6	46.8	28.0	0.1	13.8	79.5	6.6
75-100	1.4	7.4	27.5	63.7	0.0	0.0	6.6	93.4

Table 15: The left table presents the empirical transition matrix $s_t(q, q')$ that represents the probability that households in quartile q move to q' . The data is taken from [Kennickell and Starr-Mccluer \(1997\)](#) that computes the transition matrix using the SCF data from 1983 to 1989. The right table presents the model counterpart.