Asset Pricing with Realistic Crises Dynamics*

Goutham Gopalakrishna†

November 18, 2020

Abstract

What causes deep recessions and slow recovery? I revisit this question and develop a macro-finance asset pricing model that *quantitatively* matches the salient empirical features of financial crises such as a large drop in the output, a high risk premium, reduced financial intermediation, and a long duration of economic distress. The model features leveraged intermediaries who are subjected to both capital and productivity shocks, and face a regime-dependent exit rate. I show that the model without time varying intermediary productivity and exit, which reduces to Brunnermeier and Sannikov [2016], suffers from a tension between the *amplification* and the *persistence* of financial crises. In particular, there is a trade-off between the unconditional risk premium, the conditional risk premium, and the probability and duration of crisis. Features that generate high financial amplification also induce faster recovery, at odds with the data. I show that my model resolves this tension and generates realistic crises dynamics. The model is solved using a novel numerical method with active machine learning that is scalable and alleviates the curse of dimensionality.

*I thank my advisor Pierre Collin-Dufresne for invaluable guidance. I also thank Oliver Krek for helpful comments.

†EPFL and Swiss Finance Institute. Email-address: goutham.gopalakrishna@epfl.ch

1 Introduction

It is well known that recessions are marked by a high equity risk premium, a lower investment rate, and a lower output. The great recession of 2007-2008 emphasized the importance that the financial intermediaries play in propagating shocks to the real economy. Since then, there has been a growing literature with the leverage of intermediaries as a key factor in moving the asset prices and the real economy.[1] Figure (1) shows the evolution of investment rate, the equity risk premium, and the leverage of bank holding companies (BHC) in the United States. Recessions that feature a sharp decrease (increase) in the investment rate (risk premium) also feature a sharp increase in the leverage of BHCs. While the intermediaries take a central role in the recent macro-finance literature, the financial constraints that they face are of particular importance (see, example, Brunnermeier and Sannikov [2014] (BS2014), He and Krishnamurthy [2013], Di Tella [2017]). In these models, the financial constraints bind only in certain times which lead to non-linearity in the asset prices. In normal times, the financial markets facilitate capital allocation to the most productive agents. In such states, the intermediaries are sufficiently capitalized and the premium on the risky asset is low. In bad times, the financial constraints bind and capital gets misallocated to the less productive agents, who do not value capital as much. This leads to a deterioration of intermediary balance sheets and pushes the system into the crisis region where the premium on the risky asset shoots up. While these models explain a high risk premium in the crisis periods, the contribution has largely been qualitative with the exception of He and Krishnamurthy [2019](HK2019), and Krishnamurthy and Li [2020].

The contribution of this paper is two-fold. First, I build an overlapping-generation incomplete-market asset pricing model with stochastic productivity and exit of the intermediaries that occasionally generates capital misallocation and fire-sales. I solve the model using a novel deep learning based numerical method that encodes the economic information as regularizers.[2] This methodology, as shown in the companion paper Gopalakrishna [2020], is scalable and can be applied to similar high dimensional problems. Second, I show that a simpler model with constant productivity and no exit of intermediaries, which reduces to Brunnermeier and Sannikov [2016](BS2016) with recursive preferences, suffers from a tension between the *amplification* and the *persistence* of financial crises. In particular, there is a trade-off between the unconditional risk premium, the conditional risk premium, and the probability and duration of crisis. The model with stochastic productivity and regime-dependent intermediary exit rate resolves this tension and provides reasonable crisis dynamics and a better match to the empirical asset pricing moments. More specifically, my model simultaneously generates a realistic unconditional risk premium, conditional risk premium, probability of crisis, and the duration of crisis without compromising on the other dimensions such as the GDP growth rate, and the intermediary leverage patterns. The model also generates a higher time variation in the risk premium, the risk free rate, and the investment rate compared to the benchmark model.

The literature on incomplete market macro-finance models, following BS2014 and BS2016, assumes a higher productivity rate of experts relative to households, but it is constant throughout the state space. I depart from this assumption and consider a time varying productivity rate of experts. A negative shock that hits the capital reduces the size of

[1] See, for example, Brunnermeier and Sannikov [2014], He and Krishnamurthy [2013], Di Tella [2017], Adrian et al. [2014], Phelan [2016], Moreira and Savov [2017], etc.

[2] Regularizer is a commonly used tool in machine learning to reduce overfitting. See Glorot and Bengio [2010a] for details.

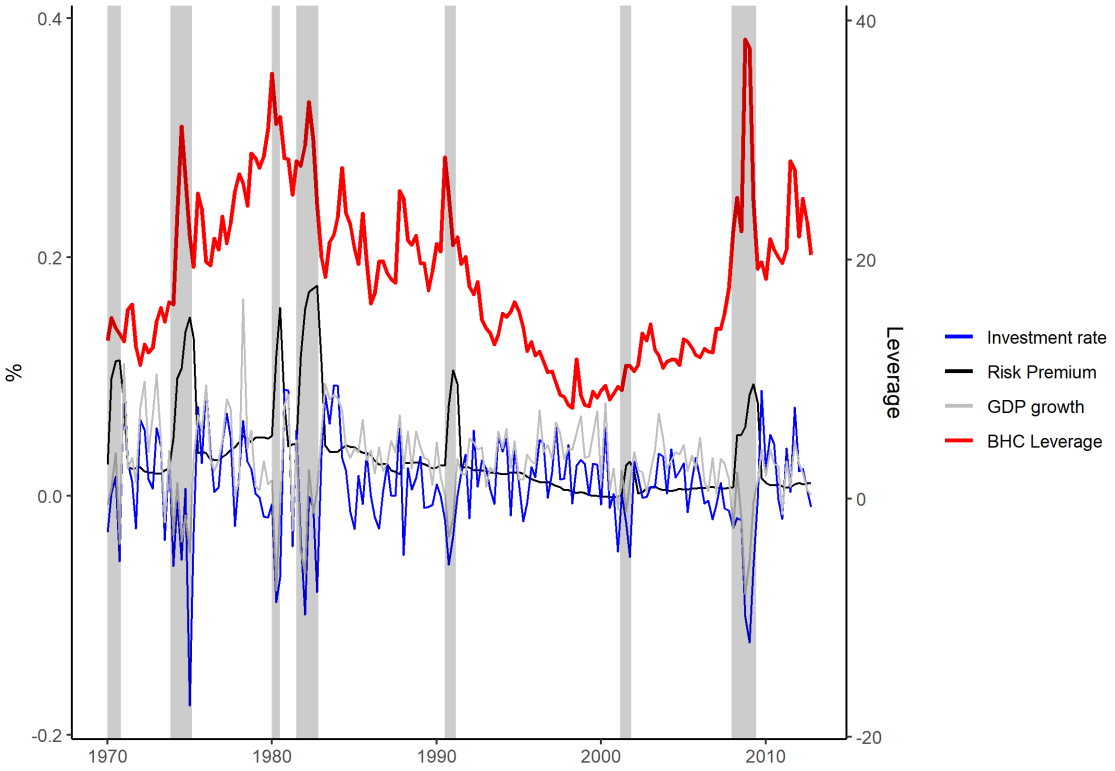

Figure 1: The red line (leverage) corresponds to right axis and the remaining lines correspond to the left axis. Leverage is computed from Federal flow of funds data. The risk premium is computed from the regression $R^e_{t+1} = a + \beta * D_t/P_t + \beta_{rec} * \mathbb{1}_{Rec} * D_t/P_t + \beta_{fin} * \mathbb{1}_{fin} * D_t/P_t + \epsilon_t$ where R^e_{t+1} is the one-year ahead excess return on S&P500, D_t/P_t denotes dividend yield on S&P500, and the dummy variables are flag for recessionary and financial crisis periods. The shaded region in the graph correspond to the NBER recessionary periods.

the experts and pushes down the productivity rate. A shrinking balance sheet of experts causes them to loose the comparative advantage that they hold over the households due to loss in economies of scale. As a result, when the wealth share of the experts becomes sufficiently low such that the crisis region is entered, it gets harder for the experts to regain wealth quickly and revert to the normal region. In addition, I assume an exogenous regime-dependent exit rate of experts which can be thought of as a parsimonious way of capturing bank defaults. The data from Federal Deposit Insurance Corporation (FDIC) shows that a total of 297 banks failed between the period 2009-2010 in the United States, which is a strikingly large number compared to 25 bank failures in the 7 years that preceded the crisis, and 23 bank failures between 2015-2020.[3] Similarly, when measured by default volume, around 80% of the Moody's rated issuers defaults in the year 2008 came from the financial institutions.[4] Figure (2) shows the evolution of bank failures from 2001 till 2020. Both in terms of the count and the default volume, bank failures during the Great recession were

[3]Source: https://www.fdic.gov/bank/historical/bank/.

[4]Source: Moody's Corporate Default and Recovery Rates, 1920-2008. Financial institutions include Bank holding companies, Real estate and insurance companies.

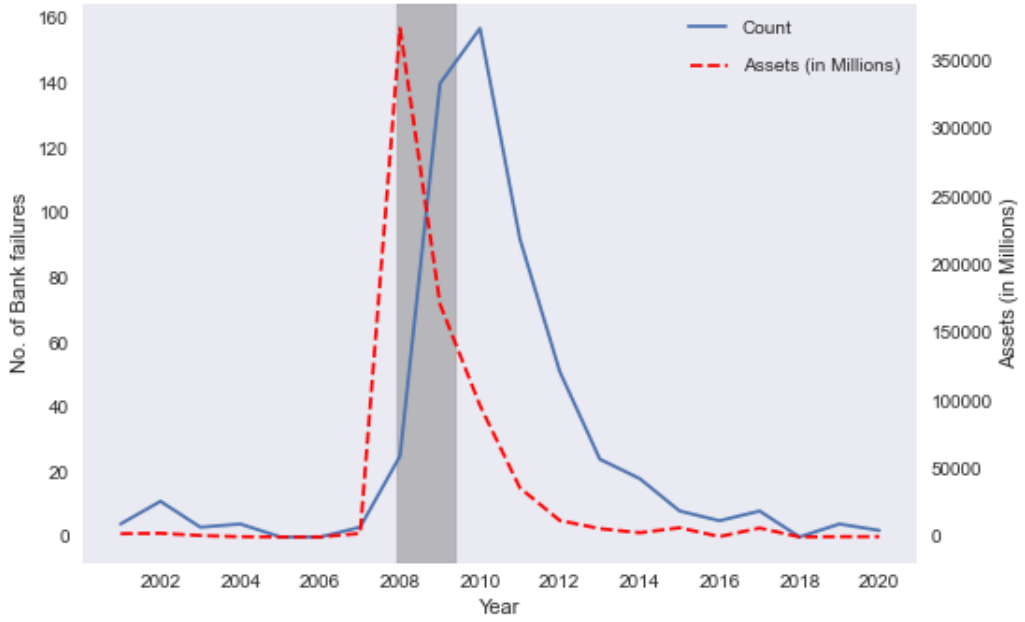

Figure 2: Bank failures from 2001 till 2020. The solid line indicates the number of bank failures and the dashed line indicates the default volume. The shaded region represent the NBER recessionary period. Source: Federal Deposit Insurance Corporation.

far greater than the other years.[5] While a lot of non-financial institutions failed too during the Great recession, the fact that 80% of Moody's issuer default in terms of volume came from financial institutions alone indicates that the intermediaries default to a large extent particularly during financial crises.[6] I capture this empirical phenomenon in reduced form through an exogenous regime-dependent intermediary exit rate. While the crisis point is endogenously determined in my model, once the system enters the crisis state, a higher fraction of the experts exit than when the system is in the normal region.

The model with two state variables- *wealth share* of the experts, and *productivity* of the experts, is solved using a deep learning based numerical algorithm that takes advantage of the universal approximation theorem by Hornik et al. [1989], which states that a neural network with one hidden layer can approximate any Borel measurable function. This method is scalable since it alleviates the curse of dimensionality that plagues the finite-difference schemes in higher dimensions. The main difficulty that arises from the grid-based solutions such as finite-difference schemes is the combination of an explosion in the number of grid points and the need for a reduced time step size as the dimensions grow large. My solution side-steps these limitations since it is mesh-free. I rely on Tensor-flow, a deep learning library developed by Google Brain, that computes the numerical derivatives efficiently. This algorithm dominates the finite-difference method used in BS2016, Hansen et al. [2018], etc., since it has the advantage of being easier to code when scaling to higher dimensions. Appendix A.5.1 shows that the solution obtained from this algorithm matches

[5]The list of banks include only those that are insured by the FDIC. Failure of investment banks such as Lehman Brothers in 2008 are not included. Source: Federal Deposit Insurance Corporation.

[6]Note that I use the terms 'intermediaries' and 'experts' interchangeably.

the solution from the finite difference method when applied to a simpler model with one state variable. I also demonstrate how one can modify few lines of code and jump from a low to a high dimensional state space. The companion paper Gopalakrishna [2020] discusses the algorithm in detail and applies it to similar problems with the number of dimensions as high as five.

In the absence of a stochastic productivity and intermediary exit, the model reduces to BS2016 augmented with OLG and recursive preferences. The assumption of OLG offers a non-degenerate stationary distribution of the state variable[7] (similar to Gârleanu and Panageas [2015]), while recursive preference helps with obtaining realistic asset pricing moments. I quantify this benchmark model, similar in spirit to HK2019 and Krishnamurthy and Li [2020] but with notable differences. The model that I consider has both the households and the experts consuming by solving an infinite horizon optimization problem, whereas, in HK2019 the experts do not consume and solve a myopic optimization problem. Both models feature non-linear asset prices arising due to occasionally binding financial intermediary constraints. However, the transition from the normal to the crisis state is smooth in HK2019. On the contrary, the model that I consider, similar to BS2016, features an endogenous jump in the risk prices that reflects the fact that periods prior to financial crises are typically calm with an exceedingly low risk premium (Baron and Xiong [2017]) and rises dramatically once the crisis period begins. The endogenous jump in the model is caused by the fire-sale effect since the households have a lower valuation of the capital. The effect of fire sales on the asset markets is crucial in times of distress, as is emphasized in Kiyotaki and Moore [1997], Shleifer and Vishny [2011], and Kurlat [2018]. Importantly, due to the endogenous jump, the point in the state space at which the financial crisis occurs is well-defined. In models where the transition is smooth, one has to rely on an exogenously defined threshold at which the system enters the crisis region. Krishnamurthy and Li [2020] considers the model with an endogenous jump similar to this paper but focuses on matching credit spreads across several financial crisis episodes with an emphasis on the pre-crisis froth in credit markets. While the agents in their model have log utility with the capital subject to both Brownian and Poisson shocks, I consider a recursive utility function and focus on matching a broader set of macroeconomic and asset pricing moments such as the intermediary leverage patterns, the risk-free rate, the equity risk premium, the investment rate, the GDP growth rate, the probability and duration of crisis among others. Recursive utility has the advantage of separating the risk aversion from the IES (Bansal and Yaron [2004]) and also helps with obtaining better asset pricing moments.

Models of intermediary asset pricing highlight the *persistence* and the *amplification* of shocks caused by the leveraged agents. A direct measure of persistence is the duration of crisis. A 10% probability of crisis, with an average of 3 years duration is a lot different than the same probability with 1 year duration since in the former case, the capital moves slower to the productive agents. It is common for asset pricing models in the literature to explain a large risk premium with highly risk averse agents.[8] However, whether one can quantitatively generate a large risk premium for a realistic occupation time in the distressed region is an important question that is overlooked in the literature. The quantification of the benchmark model done in this paper addresses this question and reveals a trade-off

[7] The OLG assumption provides a non-degenerate distribution even when there is no discount rate heterogeneity.

[8] For example, the risk aversion of the households is set to 10 in Gârleanu and Panageas [2015]. Such high upper bound for the risk aversion is also used in Mehra and Prescott [1985].

between the probability and duration of crises and the risk aversion of agents. A high risk aversion increases the risk premium that the experts earn in the normal region and does not cause enough deterioration in their net worth to have a realistic occupation time in the crisis state. Moreover, as soon as the system enters the crisis region, the risk premium spikes, enabling the experts to gain wealth quickly and revert to the normal regime. With larger values of risk aversion, the experts build wealth even faster resulting in a higher speed of reversion to the normal state. This poses a direct challenge to the heterogeneous agent models with leveraged agents that are calibrated with high risk aversion to generate a large risk premium. The benchmark model has its strengths in capturing the non-linearity of the asset prices, the output growth, and the leverage patterns of intermediaries. The biggest weakness lies in generating a realistic duration of crisis, non-linearity in the investment rate, and the level of unconditional risk premium. The capital price does not drop enough in the crises for the investment rate to fall. The richer model with stochastic productivity and regime-dependent exit rate of the experts resolves this tension and generates reasonable asset pricing moments. In particular, the time varying productivity of the experts helps to produce realistic volatility of the risk premium, the investment rate, and the risk free rate, while the regime-dependent exit rate generates reasonable amplification and persistence simultaneously. Embedding these two features that has empirical support brings the model closer to the data in important aspects.

Related Literature: This paper relates to several strands of the literature. On the modeling front, it is most closely related to BS2016 who introduce a continuous time macro-finance model based on capital misallocation and fire-sales. It fits within a large body of intermediary based asset pricing models such as BS2014, He and Krishnamurthy [2013], Di Tella [2017], Adrian and Boyarchenko [2012], Moreira and Savov [2017], etc. While BS2014 assume risk neutral agents with an exogenous interest rate, the agents in BS2016 are risk averse with CRRA utility function, and the risk free rate is endogenous. The capital misallocation in BS2016 occurs due to bad shocks and the subsequent fire-sale effect. Moll [2014] analyses a model where the inability of the productive agents to lever up due to collateral constraints causes the capital misallocation.

The empirical evidence for intermediary-based asset pricing highlights the role that the banks and the hedge funds play in pricing assets (He et al. [2017a], and Adrian et al. [2014]). While these papers provide a theory based on the intermediary leverage as a motivation for empirical findings, the literature that tightly tests the ability of general equilibrium asset pricing models with financial frictions to match the data is sparse. Two related papers that attempt to fill the gap are Muir [2017], and HK2019. However, the experts in their model do not consume and solve a myopic optimization problem, whereas, in my model both the households and the experts consume a fraction of the total output by solving an infinite horizon optimization problem. While HK2019 focus on matching the non-linearity of their model with the data and consider an exogenously defined probability of crisis, the goal of this paper goes beyond matching just the non-linearity, and deals with an endogenous crisis boundary- a slightly more daunting task since there is one less degree of freedom. In this regard, this paper comes closer to Krishnamurthy and Li [2020] who attempt to match the pre-crisis froth in the credit market through a Bayesian learning model. Muir [2017] analyses risk premia during downturns for a large panel of countries and finds that financial crises are crucial in understanding the variation in risk premium. Also, the intermediary based asset pricing model is shown to fare better compared to the consumption based representative agent models with long run risk (Bansal and Yaron

[2004]), habit (Campbell and Cochrane [1999]), and rare disaster (Barro [2006]) features.

Hansen et al. [2018] provide a framework that nests several models based on financial frictions. Even though the frictions prevent the economy from achieving a first-best outcome, their model features a dynamically complete market since the households can hedge their risk exposures through the derivative market. Their contribution is largely to provide qualitative insights by comparing different nested models, whereas, this paper is guided by quantitative analysis. While they consider a multi-dimensional problem with auxiliary shocks to the volatility and the long run growth, my model has stochastic productivity and exit rate of experts. More importantly, I conduct extensive simulations to test the model performance in matching a broader set of the macroeconomic and the asset pricing moments. My model assumes that the productivity of experts is a function of its size (wealth share of experts) which holds empirical relevance (Hughes et al. [2001], Feng and Serletis [2010], and Berger and Mester [1997]). I consider a parsimonious way to capture bank defaults through an exogenous exit rate of experts which complements a large literature on the endogenous bank runs and defaults (Gorton and Ordoñez [2014], Gertler et al. [2019], Li [2020]).

Lastly, this paper also relates to the literature on global solution methods for heterogeneous agent models using continuous time machinery (see Achdou et al. [2014b] for an overview). The assumption that the agents can consume and invest continuously in response to their instantaneous change in wealth not only greatly simplifies the computation, it also reflects the reality that people do not take these decisions only at the end of a quarter. Another advantage of the continuous-time method is the analytical tractability of equilibrium prices up-to a coupled or decoupled system of partial differential equations. Achdou et al. [2014a], BS2016, and Fernández-Villaverde et al. [2020] offer a solution technique involving implicit scheme with *up-winding* to solve the PDEs that ensures faster convergence. D'Avernas and Vandeweyer [2019] document that finite difference methods are difficult to implement in higher dimensions not only because of the curse of dimensionality but also due to the difficulty in preserving the monotonicity of the finite difference scheme. They offer a solution method based on Bonnans et al. [2004] that involves rotating the state space and finding the right direction to approximate the cross partial derivatives such that the monotonicity of the scheme is preserved. With the advancements in machine learning, recent papers have turned to neural network to solve equilibrium models. Duarte [2017] considers a method based on deep learning to solve asset pricing problems in high dimensions. Fernández-Villaverde et al. [2020] solves for the high dimensional law of motion of households using a deep neural network.[9] The algorithm proposed in this paper is similar in spirit but also incorporates prior information from the crisis boundary as *regularizers* and is particularly geared towards solving heterogeneous agent incomplete market problems with capital misallocation and *endogenous* jump in prices. It also seeks inspiration from active machine learning where the algorithm learns to sample in an informed manner. To the best of my knowledge, this is the first paper to apply a deep learning based algorithm to solve such type of a model.

The paper is organized as follows. Section 2 introduces the model. Section 3 presents the benchmark model and quantifies it to shed light on the tension between the amplification and the persistence of crises. Section 4 shows that the model with stochastic productivity and exit rate of experts resolves the tension and brings the model closer to the data. Section

[9]There is a substantial literature on the deep-learning techniques to solve PDEs in Applied Mathematics, which I cover in the companion paper Gopalakrishna [2020]. For the application of deep learning techniques to solve discrete time DSGE models, see Azinovic et al. [2019].

5 concludes. The proofs and details on numerical methodology can be found in Appendix *A*.

2 Model

In this section, I present a heterogeneous agent model with stochastic productivity and exit rate of the experts. There is an infinite horizon economy with a continuum of agents, who are of two types: Household ($\mathbb{H}$) and Expert ($\mathbb{E}$). The aggregate capital in the economy is denoted by k_t, where $t \in [0, \infty)$ denotes time. Within each group, the agents are identical and hence we can index the representative household and the expert by $h \in \mathbb{H}$ and $e \in \mathbb{E}$ respectively.[10] The experts can issue risk-free debt, and obtain a higher return to holding capital as they are more productive than the households. The friction is such that the experts have to retain at least some amount of equity on their balance sheet. In the absence of this friction, it is desirable for the experts to hold all capital as they are more productive users. Also, the agents are precluded from shorting the risky capital. The production technology can be written as

$$y_{j,t} = a_{j,t} k_{j,t} \qquad j \in \{e, h\} \tag{1}$$

where the capital evolves as[11]

$$\frac{\partial k_{j,t}}{k_{j,t}} = (\Phi(\iota_{j,t}) - \delta)dt + \sigma dZ_t^k \tag{2}$$

with $\iota_{j,t}$ as the investment rate, and $\{Z_t \in \mathbb{R}; \mathcal{F}_t, \Omega\}$ is the standard Brownian motions representing the aggregate uncertainty in $(\Omega, \mathbb{P}, \mathcal{F})$. The parameter σ denotes the exogenous volatility of capital process. The investment function $\Phi(\cdot)$ is concave and captures the decreasing returns to scale, and δ is the depreciation rate of capital. As in BS2016, $\Phi(\cdot)$ captures the technological illiquidity. The depreciation rate is the same for both the households and the experts. I assume that the investment cost function takes the logarithmic form[12] $\Phi(\iota) = \dfrac{\log(\kappa\iota + 1)}{\kappa}$ where κ is the adjustment cost parameter that controls the elasticity of the investment technology.

I assume that the productivity of the experts is governed by the following stochastic differential equation

$$da_{e,t} = \pi(\hat{a}_e - a_{e,t})dt + \nu \underbrace{(\overline{a}_e - a_{e,t})(a_{e,t} - \underline{a}_e)}_{\sigma_{ae,t}} dZ_t^a \tag{3}$$

where the Brownian shock dZ_t^a has a correlation φdt with the Brownian shock dZ_t^k with $\varphi > 0$. That is, the expert productivity follows an Ornstein–Uhlenbeck process with stochastic volatility such that it moves between a lower level $\underline{a}_e$ and an upper level $\bar{a}_e$ with a persistence parameter π and mean $\hat{a}_e \in (\underline{a}_e, \bar{a}_e)$. Since $a_h < \underline{a}_e < \bar{a}_e$, the productivity of the experts is always higher than that of the households even though it fluctuates between $\underline{a}_e$ and $\bar{a}_e$.[13] The capital prices q_t follows

$$\frac{\partial q_t}{q_t} = \mu_t^q dt + \sigma_t^{q,k} dZ_t^k + \sigma_t^{q,a} dZ_t^a$$

[10] This is also due to the homogeneity of preferences of agents within each group as explained later.

[11] Note that $k_{j,t}$ is the capital held by agent j.

[12] This is a valid investment cost function since $\Phi(0) = 0$, $\Phi' > 0$, and $\Phi'' \leq 0$.

[13] I denote $(a_{j,t}; j \in \{e, h\})$ to have concise notation but it is to be understood that $a_{h,t}$ is just a constant a_h, whereas $a_{e,t}$ follows equation (3).

The return process for each type of agent is given by

$$\partial R_{j,t} = \Bigg(\underbrace{\mu_t^q + \Phi(\iota) - \delta + \sigma\sigma_t^{q,k} + \varphi\sigma\sigma_t^{q,a} + \frac{a_{j,t} - \iota_t}{q_t}}_{\mu_{j,t}^R} \Bigg) dt + (\sigma_t^{q,k} + \sigma) dZ_t^k + \sigma_t^{q,a} dZ_t^a \quad (4)$$

The aggregate output in the economy is given by

$$y_t = A_t k_t$$

where $k_t = \int_{\mathbb{E}\cup\mathbb{H}} k_{j,t}$ and A_t is the aggregate dividend that satisfies

$$A_t = \int_{\mathbb{H}} a_h \frac{k_{h,t}}{k_t} + \int_{\mathbb{E}} a_{e,t} \frac{k_{e,t}}{k_t}$$

Let the capital share held by the expert sector be denoted by

$$\psi_t := \frac{\int_{\mathbb{E}} k_{j,t}}{\int_{\mathbb{H}\cup\mathbb{E}} k_{j,t}}$$

The experts and the households trade capital and the experts face a skin-in-the-game constraint that forces them to retain at least a fraction $\underline{\chi} \in [0, 1]$ of the equity on their balance sheet. The agents can also trade in the risk free security that pays a return r_t that is determined in the equilibrium. The stochastic discount factor (SDF) process for each type of agent is given by

$$\frac{\partial \xi_{j,t}}{\xi_{j,t}} = -r_t dt - \zeta_{j,t}^k dZ_t^k - \zeta_{j,t}^a dZ_t^a \quad (5)$$

where $\zeta_{j,t}^k$ and $\zeta_{j,t}^a$ are the prices of risk for the shocks dZ_t^k and dZ_t^a respectively.

Preferences and equilibrium: I assume that the agents have recursive utility with IES=1. That is, the utility is given by

$$U_{j,t} = E_t \Bigg[\int_t^\infty f(c_{j,s}, U_{j,s}) ds \Bigg]$$

with

$$f(c_{j,t}, U_{j,t}) = (1 - \gamma_j)\rho_j U_{j,t} \Bigg(\log(c_{j,t}) - \frac{1}{1 - \gamma_j} \log\Big((1 - \gamma_j) U_{j,t} \Big) \Bigg) \quad (6)$$

where γ_j and ρ_j are the risk aversion and the discount rate of agent j respectively. Following Gârleanu and Panageas [2015], I assume that some agents are born and die at each time instant with a probability λ_d. Let $\bar{z}$ and $1 - \bar{z}$ denote the proportion of the experts and the households that are born at each instant respectively. The death risk is not measurable under the filtration generated by the Brownian process $\mathcal{F}_t$ and the agents do not have bequest motives. Hence, once the agents die, the wealth is pooled and distributed on a pro-rata basis. As a result of the death risk, the rate of time preference parameter ρ_j can be thought of as inclusive of the death rate λ_d. I abstract away from the insurance markets to hedge the death risk, similar to Hansen et al. [2018] for simplicity. I assume that at

each time instant dt, a fraction $\tau_t dt$ of the experts become households. I allow the exit rate to be regime-dependent such that τ_t is larger in the crisis region.[14] This assumption is a parsimonious way to capture bank failures, which are particularly high during financial crises as seen in Figure (2). The agents optimize by maximising their respective utility functions, subject to the wealth constraints[15] starting from some initial wealth $w_{j,0}$. They solve

$$\begin{aligned}\sup_{c_{j,t},\chi_t,\psi_t,\iota_{j,t}} \quad & E_t\left[\int_t^\infty f(c_{j,s},U_{j,s})ds\right] \\ \text{s.t.}\frac{dw_{j,t}}{w_{j,t}} &= (r_t - \frac{c_{j,t}}{w_{j,t}} + \frac{\chi_t\psi_t}{z_t}\epsilon_{j,t})dt \\ &\quad + \frac{\chi_t\psi_t}{z_t}(\sigma+\sigma_t^{q,k})dZ_t^k + \frac{\chi_t\psi_t}{z_t}(\sigma_t^{q,a})dZ_t^a; \quad j\in\{e,h\}\end{aligned} \tag{7}$$

where $\iota_{j,t}, \chi_t, \psi_t$ denote the investment rate, the experts' inside equity share, and the experts' capital share respectively, and

$$\epsilon_{e,t} := \zeta_{e,t}^k(\sigma+\sigma_t^{q,k}) + \zeta_{e,t}^a\sigma_t^{q,a} + \varphi(\zeta_{e,t}^a(\sigma+\sigma_t^{q,k}) + \zeta_{e,t}^k\sigma_t^{q,a}) \tag{8}$$

$$\epsilon_{h,t} := \zeta_{h,t}^k(\sigma+\sigma_t^{q,k}) + \zeta_{h,t}^a\sigma_t^{q,a} + \varphi(\zeta_{h,t}^a(\sigma+\sigma_t^{q,k}) + \zeta_{h,t}^k\sigma_t^{q,a}) \tag{9}$$

There are two prices of risk for each type of the agent: $\zeta_{j,t}^k$ and $\zeta_{j,t}^a$, corresponding to the capital shock and the productivity shock respectively. By borrowing in the risk free market at a rate r_t and investing in the risky capital, they obtain the prices of risk $\zeta_{j,t}^k$ and $\zeta_{j,t}^a$. There are in fact an infinite number of agents in the economy but each individual in type $\mathbb{E}$ and $\mathbb{H}$ are identical, hence they have the same preferences. Therefore, one can seek an equilibrium in which all agents in the same group take the same policy decisions. For completeness, I present the full version of the equilibrium first.

Definition 2.1. A competitive equilibrium is a set of aggregate stochastic processes adapted to the filtration generated by the Brownian motions Z_t^k and Z_t^a. Given an initial distribution of wealth between the experts and households, the processes are prices (q_t, r_t), policy functions $(c_{j,t}, \iota_{j,t}, \chi_t, \psi_t; j\in\{e,h\})$ and net worth $(w_{j,t}; j\in\{e,h\})$, such that

- Capital market clears: $\int_{\mathbb{H}}(1-\psi_t)k_{j,t}dj + \int_{\mathbb{E}}\psi_t k_{j,t}dj = \int_{\mathbb{H}\cup\mathbb{E}} k_{j,t}dj \quad \forall t$
- Goods market clear: $\int_{\mathbb{H}\cup\mathbb{E}} c_{j,t}dj = \int_{\mathbb{H}\cup\mathbb{E}}(a_{j,t}-\iota_{j,t})k_{j,t}dj \quad \forall t$
- $\int_{\mathbb{H}\cup\mathbb{E}} w_{j,t}dj = \int_{\mathbb{H}\cup\mathbb{E}} q_t k_{j,t} \quad \forall t$

Asset pricing conditions: The equilibrium conditions map the optimal consumption, the investment, the capital share, and the capital price to the history of Brownian shocks

[14] These are unexpected changes and hence they don't affect the optimization problem, although it will have an impact on the wealth share. Gomez [2019] uses a similar assumption that applies to the leveraged wealthy households.

[15] Note that since all agents within the same group are identical, the wealth equation is presented for the aggregated agents. For wealth dynamics of individual agent within the group, see Appendix A.1.2.

Z_t^k and Z_t^a through the state variables $(z_t, a_{e,t})$. The agents choose the optimal investment rate by maximizing their return to holding the capital. That is, $\iota_{j,t}$ solves

$$\max_{\iota_{j,t}} \Phi(\iota_{j,t}) - \frac{\iota_{j,t}}{q_t}$$

The optimal investment rate is obtained as

$$\iota_{j,t}^* = \frac{q_t - 1}{\kappa} \tag{10}$$

The investment rate is the same for both types of the agents since it depends only on q_t. This is a standard 'q-theory' result which implies a tight relation between the price of capital and the investment rate. Thus, the growth rate of the economy is endogenously determined by the investment rate through the capital price. A higher price increases the investment rate, and causes a hike in the growth rate of output (since $\Phi'(\cdot) > 0$). The asset pricing relationship for the experts is given by[16]

$$\frac{a_{e,t} - \iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_t^{q,k} + \varphi\sigma\sigma_t^{q,a} - r_t = \chi_t\epsilon_{e,t} + (1 - \chi_t)\epsilon_{h,t} \tag{11}$$

where $\epsilon_{j,t}$ is defined in (8). The experts will issue maximum allowed equity $\underline{\chi}$ if the premium demanded by them is higher than that demanded by the households. The pricing condition of the households is given by

$$\frac{a_h - \iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_t^{q,k} + \varphi\sigma\sigma_t^{q,a} - r_t \leq \epsilon_{h,t} \tag{12}$$

where the equality holds if $\psi_t < 1$.

I solve for the decentralized Markov equilibrium by summarizing the system in terms of two state variables: wealth share of the experts denoted by z_t, and the productivity of the experts $a_{e,t}$.[17] The wealth share is defined as

$$z_t = \frac{w_{e,t}}{q_t k_t} \qquad \in (0, 1)$$

where $w_{e,t} = \int_{\mathbb{E}} w_{j,t}$ and $k_t = \int_{\mathbb{E}} k_{j,t} + \int_{\mathbb{H}} k_{j,t}$. Moving forward, I write $x_{h,t}$ and $x_{e,t}$ to denote the aggregated quantity $\int_{\mathbb{H}} x_{j,t}$ and $\int_{\mathbb{E}} x_{j,t}$ respectively.[18]

Proposition 1. *The law of motion of the wealth share of experts is given by*

$$\frac{\partial z_t}{z_t} = \mu_t^z dt + \sigma_t^{z,k} dZ_t^k + \sigma_t^{z,a} dZ_t^a \tag{13}$$

[16] This can be shown using a Martingale argument. See Appendix A.1.1 for the proof.

[17] All relevant objects scale with the capital k_t and hence we can summarize the economy in just two state variables.

[18] This is a slight abuse of notation. Since all agents within the same group are identical, we can think of $x_{e,t}$ as the quantity pertaining to a representative expert.

where

$$\mu_t^z = \frac{a_{e,t} - \iota_t}{q_t} - \frac{c_{e,t}}{w_{e,t}} + \left(\frac{\chi_t \psi_t}{z_t} - 1\right)\left((\sigma + \sigma_t^{q,k})(\hat{\zeta}_{e,t}^1 - (\sigma + \sigma_t^{q,k})) + \sigma_t^{q,a}(\hat{\zeta}_{e,t}^2 - \sigma_t^{q,a}) - 2\varphi(\sigma + \sigma_t^{q,k})\sigma_t^{q,a}\right)$$
$$+ (1 - \chi_t)\left((\sigma + \sigma_t^{q,k})(\hat{\zeta}_{e,t}^1 - \hat{\zeta}_{h,t}^1) + \sigma_t^{q,a}(\hat{\zeta}_{e,t}^2 - \hat{\zeta}_{h,t}^2)\right) + \frac{\lambda_d}{z_t}(\bar{z} - z_t) - \tau_t$$

$$\hat{\zeta}_{j,t}^1 = \zeta_{j,t}^k + \varphi\zeta_{j,t}^a; \quad j \in \{e, h\}$$

$$\hat{\zeta}_{j,t}^2 = \zeta_{j,t}^a + \varphi\zeta_{j,t}^k; \quad j \in \{e, h\}$$

$$\sigma_t^{z,k} = \left(\frac{\chi_t \psi_t}{z_t} - 1\right)(\sigma + \sigma_t^{q,k})$$

$$\sigma_t^{z,a} = \left(\frac{\chi_t \psi_t}{z_t} - 1\right)\sigma_t^{q,a}$$

Proof: See Appendix A.1.2.

The parameters λ_d and $\bar{z}$ denote the death rate and mean proportion of experts in the economy respectively at each time instant. The exit rate τ_t enters the drift of the wealth share.

2.1 Model solution

The solution method is reminiscent of the value function iteration with an inner static loop to solve for the equilibrium quantities $(\chi_t, \psi_t, q_t, \sigma_t^{q,k}, \sigma_t^{q,a})$ using a Newton-Raphson method, and an outer static loop to solve for the value functions $J_{j,t}$ using a deep neural network architecture. The first step starts from a time T and solves for the equilibrium policies from the value function that is set to take an arbitrary value. This is analogous to 'policy improvement' in the reinforcement learning literature. In the second step, the neural network solves for the value function taking the policies computed in first step as given, which is then used to update the policies in the subsequent step. This corresponds to the 'policy evaluation' in the language of reinforcement learning.[19] The two-step procedure is performed repeatedly until the value function converges. I first present and discuss the equilibrium policies in the static loop and then discuss the deep learning methodology used in the outer loop. Further details on equilibrium quantities and algorithm are relegated to Appendix A.1.5.

Static decisions and HJB equations: The value function is given by $U_{j,t}$ and the HJB for optimization problem (7) can be written as

$$\sup_{c_{j,t}, \iota_t, \theta_{j,t}} f(c_{j,t}, U_{j,t}) + E[dU_{j,t}] = 0 \tag{14}$$

Homothetic preferences imply that the value function is of the form

$$U_{j,t} = J_{j,t}(z_t, a_{e,t})\frac{k_{j,t}^{1-\gamma_j}}{1 - \gamma_j}$$

[19] While there are similarities between the value function iteration and reinforcement learning, the state space in my model is known ahead. A large part of the reinforcement learning deals with exploring new state space which is not relevant for the setup considered in this paper.

with the process for the stochastic opportunity set defined as

$$\frac{dJ_{j,t}}{J_{j,t}} = \mu^J_{j,t}dt + \sigma^{J,k}_{j,t}dZ^k_t + \sigma^{J,a}_{j,t}dZ^a_t$$

The HJB equation is computed as

$$\begin{aligned}\sup_{c_{j,t},\iota_t} \quad & \frac{1}{1-\gamma_j}\mu^J_{j,t} + \Phi(\iota_t) - \delta - \frac{1}{1-\gamma_j}\sigma^2 + (\sigma\sigma^{J,k}_{j,t} + \varphi\sigma\sigma^{J,a}_{j,t}) \\ & + \rho_j\left(\log\frac{c_{j,t}}{w_{j,t}} - \frac{1}{1-\gamma_j}\log J_{j,t} + \log(q_t z_t)\right) \end{aligned} \tag{15}$$

Proposition 2. *The optimal consumption policy, investment policy, and prices of risk are given by*

$$\hat{c}_{j,t} = \rho_j \tag{16}$$

$$\iota_{j,t} = \frac{q_t - 1}{\kappa} \tag{17}$$

$$\zeta^k_{e,t} = -\sigma^{J,k}_{e,t} + \sigma^{z,k}_t + \sigma^{q,k}_t + \gamma_e\sigma \tag{18}$$

$$\zeta^a_{e,t} = -\sigma^{J,a}_{e,t} + \sigma^{z,a}_t + \sigma^{q,a}_t \tag{19}$$

$$\zeta^k_{h,t} = -\sigma^{J,k}_{h,t} - \frac{1}{1-z_t}\sigma^{z,k}_t + \sigma^{q,k}_t + \gamma_h\sigma \tag{20}$$

$$\zeta^a_{h,t} = -\sigma^{J,a}_{h,t} - \frac{1}{1-z_t}\sigma^{z,a}_t + \sigma^{q,a}_t \tag{21}$$

Proof: See Appendix A.1.3.

The consumption-wealth ratio $\hat{c}_{j,t}$ is constant and is equal to the discount rate because IES=1. The optimal policies are given in terms of the other equilibrium quantities $(J_{j,t}, \chi_t, \psi_t, q_t, \xi_t)$ which are found by solving for a Markov equilibrium in the state space $(\mathbf{z_t} \in (\mathbf{0},\mathbf{1}), \mathbf{a_{e,t}} \in (\underline{\mathbf{a}}_\mathbf{e}, \bar{\mathbf{a}}_\mathbf{e}))$.

Definition 2.2. A Markov equilibrium in $(\mathbf{z_t} \in (\mathbf{0},\mathbf{1}), \mathbf{a_{e,t}} \in (\underline{\mathbf{a}}_\mathbf{e}, \bar{\mathbf{a}}_\mathbf{e}))$ is a set of adapted processes $q(z_t, a_{e,t}), r(z_t, a_{e,t}), J_e(z_t, a_{e,t}), J_h(z_t, a_{e,t})$, policy functions $\hat{c}_e(z_t, a_{e,t}), \hat{c}_h(z_t, a_{e,t}), \psi(z_t, a_{e,t})$, and state variables $\{z_t, a_{e,t}\}$ such that

- $J_{j,t}$ solves the HJB equation and corresponding policy functions $\hat{c}_{j,t}, \psi_t$
- Markets clear

$$(\hat{c}_{e,t}z_t + \hat{c}_{h,t}(1-z_t))q_t = \psi_t(a_{e,t} - \iota_t) + (1-\psi_t)(a_h - \iota_t)$$

$$w_{e,t}z_t + w_{h,t}(1-z_t) = 1$$

- z_t and $a_{e,t}$ satisfy (13) and (3) respectively

Proposition 3. *The total return variance is given by*

$$||\sigma^R_t||^2 := (\sigma + \sigma^{q,k}_t)^2 + (\sigma^{q,a}_t)^2 = \frac{\sigma^2 + \left(\frac{\sigma^2_{ae,t}}{q_t}\frac{\partial q_t}{\partial a_{e,t}}\right)^2}{\left(1 - \frac{1}{q_t}\frac{\partial q_t}{\partial z_t}z_t\left(\frac{\psi_t\chi_t}{z_t} - 1\right)\right)^2} \tag{22}$$

Proof. See Appendix A.1.4.

The first term in the numerator on the R.H.S of equation (22) reflects the fundamental volatility while the second term captures the contribution of productivity shocks. There are two effects that drive the volatility: (a) Since $\frac{\partial q_t}{\partial z_t} > 0$, and $\frac{\psi_t \chi_t}{z_t} > 1$ in equilibrium, the denominator contributes towards a higher return volatility than the fundamental volatility σ (b) Since $\frac{\partial q_t}{\partial a_{e,t}} > 0$, the second part in the numerator adds to the amplification caused by (a).

2.1.1 Neural network solution method

The outer loop involves solving for a de-coupled system of quasi-linear PDEs- one for the households and one for the experts, taking as given the equilibrium quantities that are determined from the static loop. The PDE obtained by applying Ito's lemma to $J_{j,t}(z_t, a_{e,t})$ and using the HJB equation (15) is[20]

$$\begin{aligned}\mu^J J = \frac{\partial J}{\partial t} + \frac{\partial J}{\partial z}\mu^z + \frac{\partial J}{\partial a}\mu^a + \frac{1}{2}\frac{\partial^2 J}{\partial z^2}\Big((\sigma^{z,k})^2 + (\sigma^{z,a})^2 + 2\varphi\sigma^{z,k}\sigma^{z,a}\Big) + \frac{1}{2}\frac{\partial^2 J}{\partial a^2}\sigma_a^2 \\ + \frac{\partial^2 J}{\partial z_t \partial a}\big(z\sigma^{z,k}\sigma_a\varphi + \sigma_a\sigma^{z,a}\big)\end{aligned} \tag{23}$$

with the boundary conditions

$$\begin{aligned} J(z,a,T) &= J_0 \\ \frac{\partial J(0,a,t)}{\partial z_t} &= \frac{\partial J(1,a,t)}{\partial z_t} = 0 \\ \frac{\partial J(z,\underline{a}_e,t)}{\partial a_{e,t}} &= \frac{\partial J(z,\bar{a}_e,t)}{\partial a_{e,t}} = 0 \end{aligned} \tag{24}$$

I take advantage of the universal approximation theorem that states that a neural network with at least one hidden layer can approximate any Borel measurable function, and solve for the function $J(z,a,t)$ that is governed by the PDE (23). Starting from an arbitrary terminal value at time T, the task is to solve for $J(z,a,T-\Delta t)$. The PDE coefficients and the terminal value are in the form of a grid but not all grid points are required in the algorithm as will be explained. While the space of admissible solutions to the function given the sample data from terminal value and other boundary conditions is potentially large, I use the residuals from PDE and the boundary conditions as *regularizers* that constrain the space to a manageable size. This encoding of prior information into the learning algorithm amplifies the information content from the economic problem and makes it possible for the deep neural network to head towards the correct solution even with the limited training sample.

[20] I ignore the time and agent indices in order to avoid cluttering of notations. The productivity of the expert $a_{e,t}$, and the volatility $\sigma_{ae,t}$ are denoted as a and σ_a for simplicity in the PDEs.

Consider the PDE residual from (23)

$$f := \frac{\partial J}{\partial t} + \frac{\partial J}{\partial z}\mu^z + \frac{\partial J}{\partial a}\mu^a + \frac{1}{2}\frac{\partial^2 J}{\partial z^2}\left((\sigma^{z,k})^2 + (\sigma^{z,a})^2 + 2\varphi\sigma^{z,k}\sigma^{z,a}\right) + \frac{1}{2}\frac{\partial^2 J}{\partial a^2}\sigma_a^2 + \frac{\partial^2 J}{\partial z \partial a}\left(z\sigma^{z,k}\sigma_a\varphi + \sigma_a\sigma^{z,a}\right) - \mu^J J \tag{25}$$

Starting from a neural network $\hat{J}(z, a, t; \Theta)$ parameterized by an arbitrary Θ, the optimal parameter Θ^* that ensures that $\hat{J}(z, a, t; \Theta)$ is close to J is obtained by mimimizing the following loss function

$$\mathcal{L} = \lambda_f \mathcal{L}_f + \lambda_j \mathcal{L}_j + \lambda_b \mathcal{L}_b + \lambda_c \mathcal{L}_c \tag{26}$$

where[21]

$$\text{PDE loss} \qquad \mathcal{L}_f = \frac{1}{N_f}\sum_{i=1}^{N_f} |f(z_f^i, a_f^i, t_f^i)|^2 \tag{27}$$

$$\text{Bounding loss-1} \qquad \mathcal{L}_j = \frac{1}{N_j}\sum_{i=1}^{N_j} |\hat{J}(z_j^i, a_j^i, t_j^i) - J_0^i|^2 \tag{28}$$

$$\text{Bounding loss-2} \qquad \mathcal{L}_b = \frac{1}{N_b}\sum_{i=1}^{N_b} |\nabla\hat{J}(z_b^i, a_b^i, t_b^i)|^2 \tag{29}$$

$$\text{Crisis loss} \qquad \mathcal{L}_c = \frac{1}{N_c}\sum_{i=1}^{N_c} |\hat{J}(z_c^i, a_c^i, t_c^i) - J_0^i|^2 \tag{30}$$

The parameters $(\lambda_f, \lambda_j, \lambda_b, \lambda_c)$ are weights attached to the corresponding losses, $(z_j^i, a_j^i, t_j^i, J_0^i)_{i=1}^{N_j}$ and $(z_b^i, a_b^i, t_b^i)_{i=1}^{N_b}$ denote the boundary training data, and $(z_f^i, a_f^i, t_f^i)_{i=1}^{N_f}$ denote the collocation points for the PDE residual $f(z, a, t)$. The crisis boundary collocation points $(z_c^i, a_c^i, t_c^i)_{i=1}^{N_c}$ are sampled from the neighborhood of state space where fire-sale gets initiated, that is endogenously determined in the static inner loop. The quantities (N_f, N_j, N_b, N_c) denote the number of points to minimize the PDE loss, the two bounding losses, and the crisis boundary loss respectively. By encoding the crisis boundary loss, the neural network is forced to learn better around the crisis threshold which is where the policy functions are highly non-linear. The sampling is done uniformly with replacement in each domains. The construction of crisis loss is inspired from *active machine learning* (Settles [2012]), a budding area in the artificial intelligence literature. Active learning algorithms work by providing better training samples at each iteration to ensure quick convergence. At every iteration, the points in the state space where crisis occurs might change, and sampling more points from around this region dynamically provides better training samples. I consider artificial collocation points for time such that $\{t^i\} \in [t^i - \Delta t^i, t^i]$ are sampled uniformly so as to reduce errors in numerical derivatives with respect to the time dimension. The number of collocation points $(N_j, N_b, N_f, N_c, N_t)$ in total need not be large and is taken to be 10% of the total grid size. This makes the algorithm mesh-free and scalable to higher dimensions.

[21] I write $\nabla\hat{J}$ to denote $\begin{bmatrix}\frac{\partial\hat{J}}{\partial z} & \frac{\partial\hat{J}}{\partial a}\end{bmatrix}^T$.

	Value
No. of hidden layers	4
Hidden units	[30,30,30,30]
Activation function	Tanh
Optimizer	ADAM + L-BFGS-B
Learning rate	0.01
Loss function weights$(\lambda_f, \lambda_j, \lambda_b, \lambda_c)$	{1,1,0.001,1}
Batch size	Full batch

Table 1: Network architecture

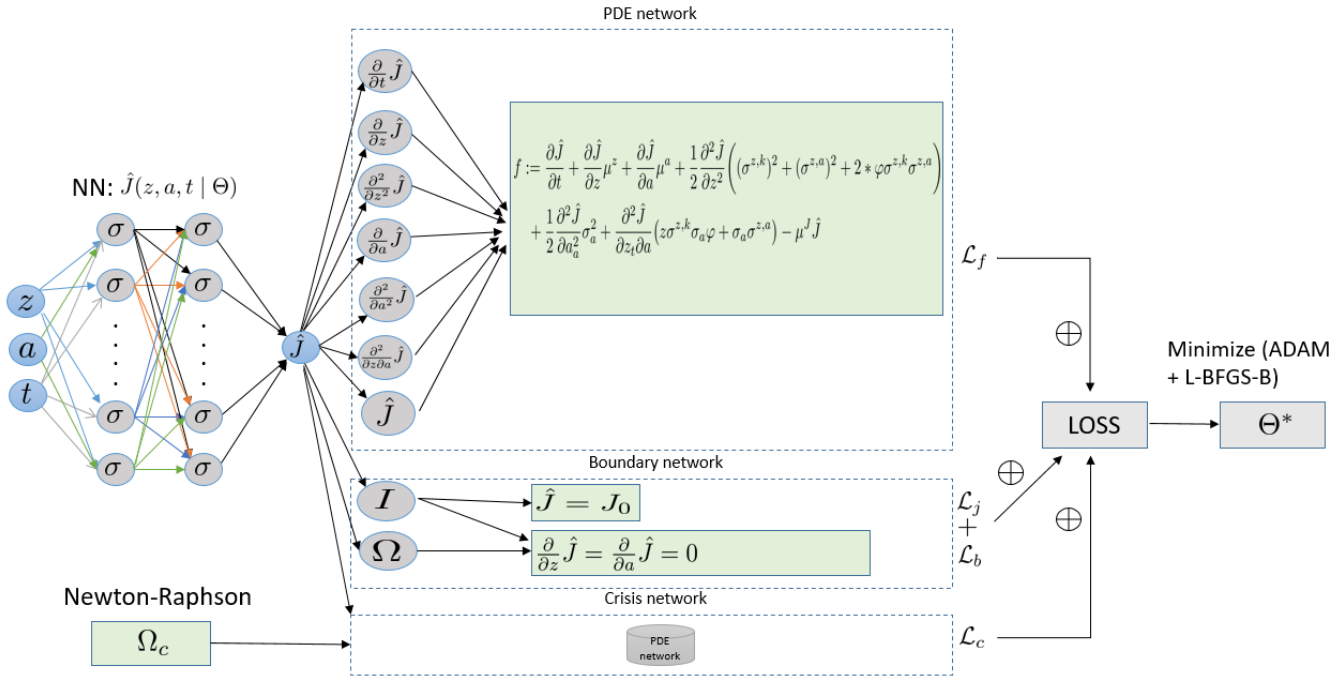

Figure 3: Neural network architecture. The quantities I and Ω denote the domain of the state space pertaining to the initial and boundary conditions respectively. The domain Ω_c refers to the crisis neighborhood and is endogenously determined in the inner static loop.

Opening the black box: The success of a deep neural network model often relies on the network architecture and the hyperparameters. The machine learning models in finance literature use extensive hyperparameter search in the tuning process to select the 'right' model (see Gu et al. [2020], Chen et al. [2019], etc.). The deep learning model used in this paper does not suffer from this problem for two reasons. First, there is no training/test/validation set really which means that one does not have to worry about the classical overfitting problem.[22] Second, and more importantly, the proposed regularization mechanism encodes the economic problem into the learning algorithm by building a meaningful loss function which enables a simple feed-forward network to arrive at the right solution. Using complex architectures such as Convolution neural network, LSTM, etc. create a 'black-box' problem which limits the ability to understand what makes the

[22]The boundary conditions provide us with data points which can be thought of as training sample, but it does not carry the same meaning as it does in supervised machine learning.

algorithm succeed. On the contrary, using a simple feed-forward network and encoding the economic information as regularizers provides a lot more visibility on how the model steers towards the right solution.

The proliferation of deep learning application in the past decade can be largely attributed to the *automatic differentiation* which has enabled reduced computation time of the derivatives of functions. In the deep learning literature, the parameters of a network are optimized through backpropogation by taking the derivative of a loss function with respect to the parameters. The approach presented in this paper explicitly uses automatic differentiation to take derivatives with respect to the space and the time dimensions. In Figure (3), the left most part of the neural network ($NN : \hat{J}(z, a, t \mid \Theta)$) is the familiar simple feed-forward architecture. The output from this network ($\hat{J}$) is fed into the PDE, boundary, and crisis network respectively that utilizes automatic differentiation in the customized loss functions. The separation of fundamental neural network with a simple architecture and the informed PDE network allows us to peek into the black-box and witness the automatic differentiation fully in action, which is the key driver of the algorithm's learning in both low and high dimensions.

Hyperparameter choices: Table (1) presents the chosen hyperparameters of the model. I use 4 hidden layers with 30 neurons each since a deep layer is empirically observed to be better than a wide layer. While a rectified linear unit is the common choice for activation function, I use a hyperbolic tangent function based on its superior performance for the problem at hand. The optimizers are chosen based on empirical observation. I use an adaptive momentum (ADAM) optimizer with a learning rate of 0.01 until error is minimized to the order of 1e-4 and then use a quasi-newton method called L-BFGS-B until convergence is ensured. The network weights and biases are initialized using Xavier initialization in order to avoid the ubiquitous vanishing/exploding gradient problem in deep learning (see Glorot and Bengio [2010a]). The weights of loss functions ($\lambda_f, \lambda_j, \lambda_c$) are uniform to give equal importance for each of these components. I use a smaller weight for the second bounding loss $\mathcal{L}_b$. Since the training sample size is much smaller than the full grid size, full batch is used in optimizer as opposed to mini-batches which is common in deep learning algorithms.

2.2 Calibration

RBC parameters: The main parameter that governs the evolution of capital is the volatility. While BS2016 uses a value of 10%, the exogenous volatility of stock market dividends is empirically observed to be lower. In fact, the consumption volatility from 1975 till 2015 is found to be just 1.24% (HK2019). I choose a value of 6% so as to obtain a non-negligible time variation in the prices. The productivity of the experts and the households, and the investment cost parameter are chosen to match the average output growth rate of 1-3% and the investment rate of 5-8%. A lower investment cost parameter increases the investment rate but also pushes up the output.

Preferences and demographics: The discount rates are chosen to match a low average risk free rate to reflect the current environment.[23] Although the discount rates are the same as in BS2016, they are inclusive of the death rate which is chosen to be 3% meaning that

[23]The 3-month T-bill rate is 0.1% in October 2020, for exmaple.

experts live on average for 37 years.[24] The fraction of newborns designated to be the experts is taken from Hansen et al. [2018]. I assume that the risk aversion parameter for the households and experts are same and equal to 5. The assumption of unitary IES greatly simplifies the numerical computation since the consumption-wealth ratio becomes constant.

Intermediation parameters: Finally, the equity retention threshold is set to be 0.65. This is comparable to the value of 0.5 used in BS2016 and Hansen et al. [2018].

	Description	Symbol	Value
Technology/Preferences	Volatility of output	σ	0.06
	Discount rate (experts)	ρ_e	0.05
	Discount rate (households)	ρ_h	0.05
	Depreciation rate of capital	δ	0.05
	Investment cost	κ	5
	Productivity (experts)	$\{\underline{a}_e, \hat{a}_e, \bar{a}_e\, \pi, \nu\}$	{0.1,0.15,0.2,0.01,4.16}
	Productivity (households)	a_h	0.02
	Correlation of shocks	φ	0.9
Utility	Recursive utility	γ_e, γ_h, IES	$[5, 5, 1]$
	Mean proportion of experts	$\bar{z}$	0.10
Demographics	Turnover	λ_d	0.03
	Experts exit rate	$\{\tau_{normal}, \tau_{crisis}\}$	{0.06,0.4}
Friction	Equity retention	$\underline{\chi}$	0.65

Table 2: Calibrated parameters in the model. All values are annualized.

The exit rate of the experts is chosen to be 6% in the normal regime and 40% in the crisis regime reflecting the empirical evidence of numerous bank failures during financial crises. Figure (4) presents the equilibrium quantities obtained from the numerical solution. The productivity level has a large effect on the capital price. A lower level of expert productivity implies a lower capital price throughout the state space. The presence of productivity shocks allow the return volatility to be higher than the fundamental volatility even in the normal regime.

When the wealth share of the more productive experts is higher, capital is fully held by them. They always operate with leverage in equilibrium and therefore, when a negative shock hits the capital, their net worth decreases disproportionately more than that of the households resulting in a deterioration of their wealth share. When it falls below a threshold $\{z^*, a_e^*\}$, the system endogenously enters into the crisis region featuring depressed asset prices, and higher asset volatility. The jump in prices occurs due to the fire sale effect. In the crisis region, experts start selling capital to the households who always value it less. Hence, the capital price has to fall enough for households to purchase it and clear the market. The fall in capital price is an inefficiency caused by failure to internalize the pecuniary externality by the agents. This is because each individual in the economy takes

[24] Gârleanu and Panageas [2015], and Hansen et al. [2018] use a value of 2% which is comparable to the value of 3% used in this paper.

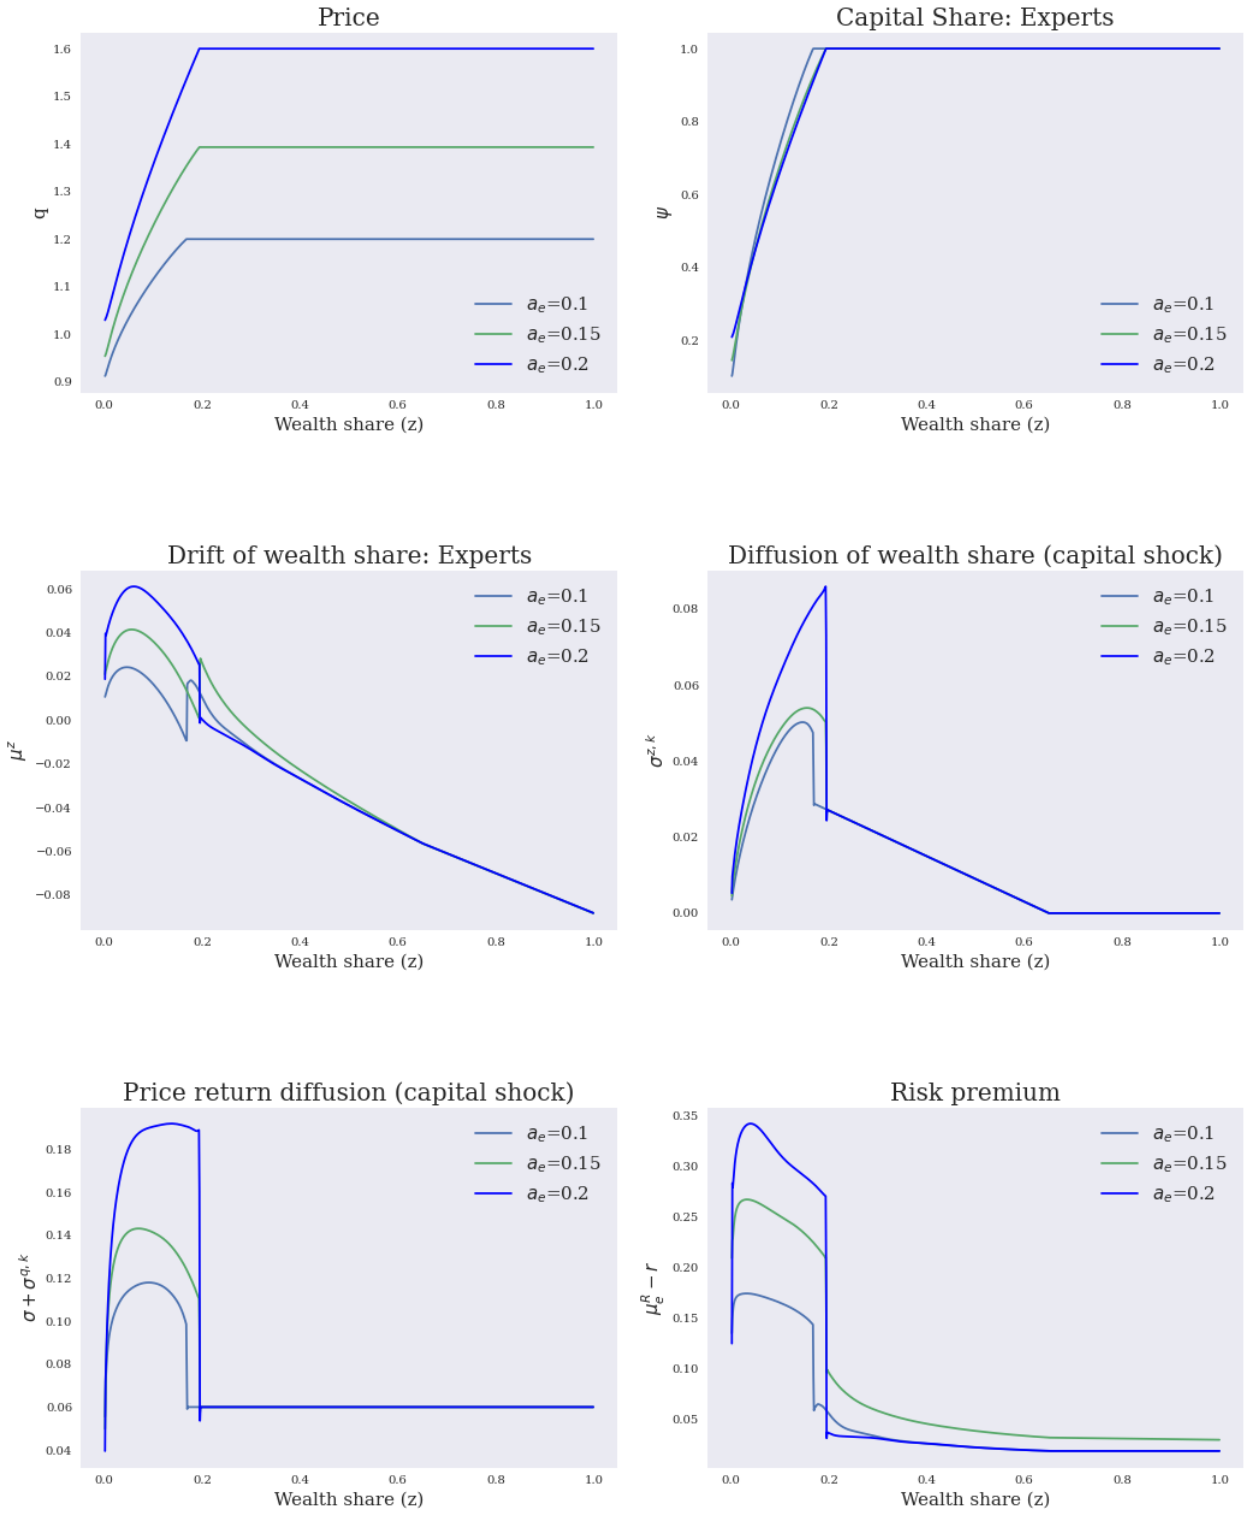

Figure 4: Equilibrium values as functions of state variable z_t for different values of $a_{e,t}$.

prices as given in their respective decision making process. To be more concrete, whenever experts choose not to hold capital, they fail to take into account the fact that the households will be forced to hold it by market clearing. Since the households value capital less, they

will demand a higher premium resulting in a fall in the capital price. This feeds-back into the experts balance sheet since they are leveraged, and causes further inefficiency and misallocation of resources. There is a second externality that the experts do not take into consideration, which is the increase in exit rate when the system enters the crisis region. The pricing dynamics is different from the heterogeneous risk aversion literature in complete markets (see Gârleanu and Panageas [2015], for example). With homogeneous productivity and heterogeneous risk aversion, experts will sell capital to household during periods of distress who will demand a higher premium (and lower price) due to their *higher risk aversion*. Although both models feature a drop in prices during the crisis, the change will be gradual in the latter. The jump in prices due to fire-sale effect can only be explained from the differences in productivity rates in an incomplete market setting and no-shorting constraint. There will be a state space where the experts hold all capital since the risk premium of households is lower than that of the experts. In such states, the households would desire to hold a negative quantity of capital but since shorting is disallowed, they will hold no capital at all. In contrast, if the productivity of households is the same as experts, they will face the same risk premium as experts. Therefore, even if their risk aversion is smaller, they would desire to hold some positive quantity of capital. This makes the transition from the normal to crisis regime smoother.[25]

3 Quantitative analysis

In this section, I present a simpler model without stochastic productivity and exit rate of the experts that will serve as a benchmark model for the quantitative analysis. Through simulation studies, I show that there is a trade-off between the amplification and the persistence of financial crises in this simpler model. While there are many channels that generate this tension, I focus on the risk aversion channel.[26]

3.1 Benchmark model

I assume that the productivity rate of both the experts and the households is constant such that $a_e > a_h$ holds. I also set the exit rate of the experts to zero in both the normal and the crisis regime. With these two simplifications, the model reduces to BS2016 augmented with recursive preference and OLG elements. While the agents have CRRA utility function in BS2016, I assume that they have recursive preference so as to disentangle the risk aversion and the inter-temporal elasticity of substitution. The rest of the assumptions carry over from the stochastic productivity model in Section 2. That is, the output is given by AK technology as in (1), with a_e and a_h as the productivity rates of the experts and the households respectively. The evolution of capital is governed by (2) as before. The capital price per unit q_t follows the process

$$\frac{dq_t}{q_t} = \mu_t^q dt + \sigma_t^q dZ_t^k$$

The terms μ_t^q, and σ_t^q are endogenously determined in the equilibrium. Note that the productivity shocks are absent in the benchmark model. Using this dynamics for the price,

[25] This dynamics is present in Gârleanu and Panageas [2015]. Hansen et al. [2018] offer additional insights for the case of heterogeneous productivity vs heterogeneous risk aversion.

[26] See Appendix A.4 for details on the skin-in-the-game constraint generating a similar trade-off.

the return process can be written as

$$dR_{j,t} = \underbrace{\left(\frac{a_j - \iota_{j,t}}{q_t} + \Phi(\iota_{j,t}) - \delta + \mu_t^q + \sigma\sigma_t^q\right)}_{\mu_{j,t}^R} dt + (\sigma + \sigma_t^q) dZ_t^k \tag{31}$$

Let $\xi_{e,t}$ and $\xi_{h,t}$ denote the SDF of the experts and the households respectively that follows

$$\frac{d\xi_{j,t}}{\xi_{j,t}} = -r_t dt - \zeta_{j,t} dZ_t^k \tag{32}$$

where, $\zeta_{j,t}$ is the market price of risk for agent j. Similar to the stochastic productivity model, both agents invest in the risk-free asset, and hence the drift of the SDF process is the same for all agents. The asset pricing conditions for the experts and the households respectively take the simpler form[27]

$$\frac{\frac{a_e - \iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_{q,t} - r_t}{\sigma + \sigma_{q,t}} = \chi_t \zeta_{e,t} + (1 - \chi_t)\zeta_{h,t} \tag{33}$$

$$\frac{\frac{a_h - \iota_t}{q_t} + \Phi(\iota_t) - \delta + + \mu_t^q + \sigma\sigma_{q,t} - r_t}{\sigma + \sigma_{q,t}} \leq \zeta_{h,t} \tag{34}$$

The equality holds in (34) if the households own some amount of capital ($\psi_t < 1$). The optimal investment rate is the same as before and is given in (10). The agents solve

$$\begin{aligned} &\sup_{c_{j,t}, \chi_t, \psi_t, \iota_{j,t}} E_t\left[\int_t^\infty f(c_{j,s}, U_{j,s}) ds\right] \\ &\text{s.t.} \frac{dw_{j,t}}{w_{j,t}} = (r_t - \frac{c_{j,t}}{w_{j,t}} + \frac{\chi_t \psi_t}{z_t}(\sigma + \sigma_t^q)\zeta_{j,t})dt + \frac{\chi_t \psi_t}{z_t}(\sigma + \sigma_t^q) dZ_t^k \qquad j \in \{e, h\} \end{aligned} \tag{35}$$

where the aggregator $f(c_{j,s}, U_{j,s})$ is given in (6). Since all agents within the group j are identical as before, I solve for the decentralized economy with wealth share of the experts z_t as the sole state variable.

Proposition 4. *The law of motion of the wealth share of experts is given by*

$$\frac{dz_t}{z_t} = \mu_t^z dt + \sigma_t^z dZ_t^k \tag{36}$$

where

$$\mu_t^z = \frac{a_e - \iota_t}{q_t} - \frac{c_{e,t}}{w_{e,t}} + \left(\frac{\chi_t \psi_t}{z_t} - 1\right)(\sigma + \sigma_{q,t})(\zeta_{e,t} - (\sigma + \sigma_t^q)) + (1 - \chi_t)(\sigma + \sigma_t^q)(\zeta_{e,t} - \zeta_{h,t}) + \frac{\lambda_d}{z_t}(\bar{z} - z_t)$$

$$\sigma_t^z = \left(\frac{\chi_t \psi_t}{z_t} - 1\right)(\sigma + \sigma_t^q)$$

Proof: See Appendix A.2.2.

The expression for the wealth share dynamics is similar to the model with stochastic productivity except that only the price of risk for capital shock matters, and the exit rate

[27] This can be proved using the Martingale argument similar to the model with stochastic productivity. See Appendix A.2.1 for the proof.

τ_t disappears from the drift. The solution methodology is also the same as before where equilibrium policies are determined in the static inner step and the value function is solved in the outer time step by solving a couple of PDEs. I use an implicit finite difference method with up-winding to solve the PDEs. The up-winding preserves the monotonicity of the PDEs and helps achieve convergence. In Appendix A.5.1, I show that the solution to the PDEs obtained using the finite difference method is the same as the solution obtained form the neural network method.

3.2 Comparative statics:

Figure (5) plots the key equilibrium quantities for the parameters used in Table (3). The static comparison from Figure (5) reveals that as the risk aversion increases, the price of capital decreases and the price of risk increases. Even though the price volatility is lower for higher risk aversion, there is a region in the parameter space where it is much higher than the case of lower risk aversion. This is because the crisis boundary z^* moves to the right with increasing risk aversion. The differences in the market price of risk translate to vast differences in the drift of the wealth share. This has a direct impact on how the system transitions in and out of the crisis region.

	Description	Symbol	Value
Technology/Preferences	Volatility of output	σ	0.06
	Discount rate (experts)	ρ_e	0.06
	Discount rate (households)	ρ_h	0.04
	Depreciation rate of capital	δ	0.02
	Investment cost	κ	10
	Productivity (experts)	a_e	0.11
	Productivity (households)	a_h	0.03
Utility	CRRA utility	γ_e, γ_h	$[1, 15]$
	Recursive utility	γ_e, γ_h	$[1, 15]$
Demographics	Mean proportion of experts	$\bar{z}$	0.10
	Turnover	λ_d	0.03
Friction	Equity retention	$\underline{\chi}$	0.5

Table 3: Calibrated parameters for the benchmark model. All values are annualized.

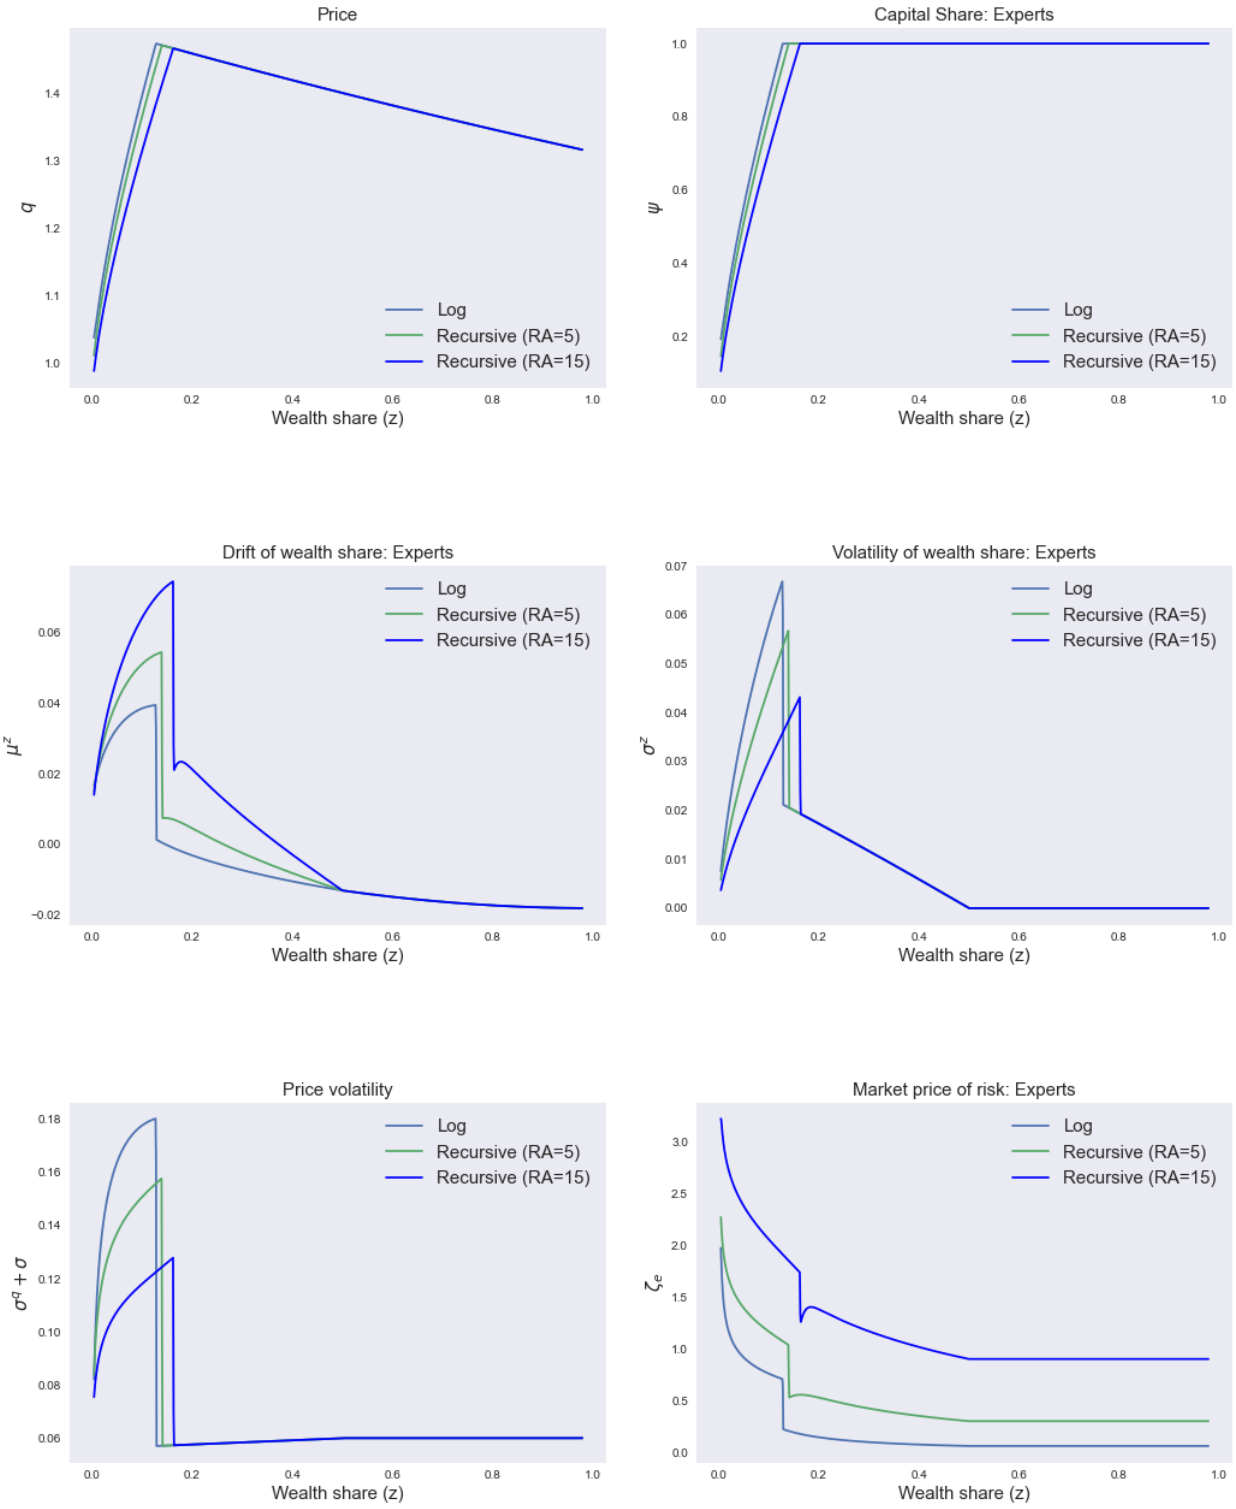

Figure 5: Equilibrium values as functions of state variable z_t. The recursive utility plots have IES equal to 1. Log utility has RA=1 by construction.

Stationary distribution: While Figure (5) gives us a qualitative description of the economy, the stationary distribution of the wealth share is required to confront the model

with the data. The stationary distribution represents the average location of the state variable z_t in the interval $[0, 1]$ as $t \to \infty$ for any given starting point z_0. I obtain this distribution by numerically simulating the model for 5000 years at monthly frequency. The simulation maps the Brownian shocks Z_t^k to state variable z_t which is governed by the law of motion given by (36). I repeat the procedure 1000 times and ignore the first 1000 years so that the distribution is not sensitive to the arbitrarily chosen initial value z_0. I annualize the result and repeat the procedure for different initial values to ensure that the economy has converged. I explain the numerical procedure in detail in Appendix A.3. Figure (6) plots the stationary distribution of the wealth share for different risk aversion levels. As the risk aversion increases, the mass of wealth share that lies in the crisis zone diminishes. In fact, it shrinks rather quickly and this result also holds if I allow for heterogeneous risk aversion with the experts being less risk averse. The stationary distribution gives us additional insights that one cannot obtain from studying the comparative static plots. In Figure (5), it appears as if increasing risk aversion will not have a drastic impact on the frequency of a crisis since the boundary z^* moves only slightly to the right. However, the drift of wealth share moves up a lot as the risk aversion is higher and this pushes the stationary distribution away from the crisis region to a greater extent. Since the experts operate with leverage, a higher price of risk will have a positive effect on their wealth share. Note that the crisis boundary z^* is far from the steady state[28] $\hat{z}$ for higher levels of risk aversion. This means that a much longer sequence of negative shocks are required to push the system into the crisis region.

[28] The steady state can be defined as $\{z_t : \mu_t^z(z_t) = 0\}$.

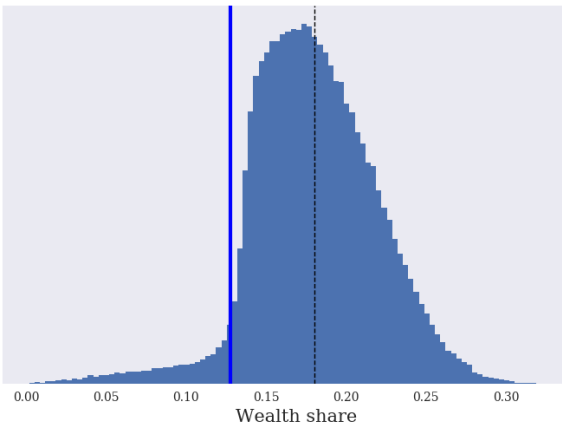

(a) RA=1

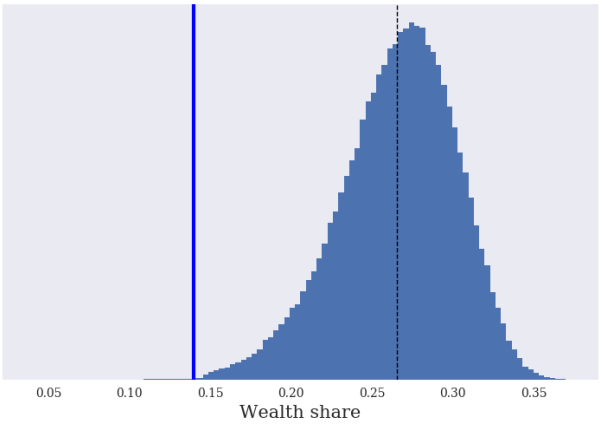

(b) RA=5

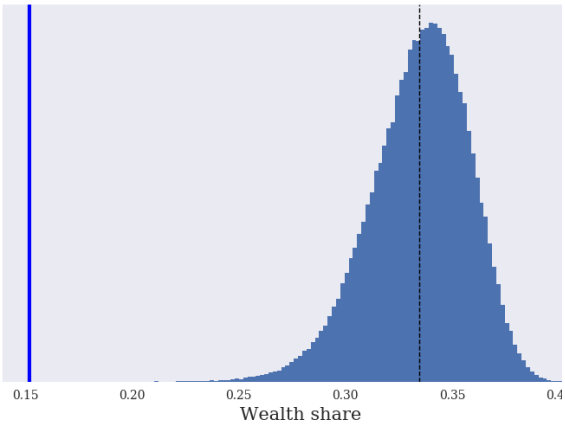

(c) RA=10

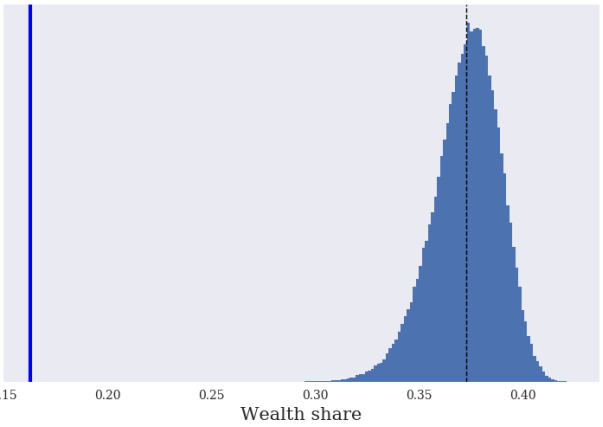

(d) RA=20

Figure 6: Stationary distribution of wealth share. Plots (a), (b), (c), and (d) represent benchmark model with risk aversion parameter set to 1, 5, 10, and 20 respectively. The vertical blue line represents crisis boundary. The vertical dotted line is the mean of the distribution.

Comparison to Data: While the crisis point is well defined and endogenously determined in the model, defining the crisis episodes in the data is a challenge. Reinhart and Rogoff [2009a] determine the frequency of crisis states to be around 7% for the advanced economy. This figure is much lower than the 29% percentage NBER recessionary periods from year 1874 till today.[29] The stark difference in the frequency between Reinhart and Rogoff [2009a] and NBER data is due to the fact that in the former, recessionary periods need to feature severe banking panic to be qualified as financial crises. This relates to the findings by Muir [2017] and Gorton and Ordoñez [2020] that not all recessions are financial crisis episodes. Muir [2017] finds that the risk premium is higher during financial crises than recessions, where a financial crisis occurs when the wealth share of intermediaries deteriorate sufficiently, just like in the model considered in this paper. HK2019 argue that the past decade in the US featured roughly three financial crisis periods. I take the probability of being in the crisis period as 3-10% for the purpose of quantitative calibration.[30] For each z_t simulated from the discretized version of its dynamics governed by (36), the equilibrium quantities are computed using the mapping given the equilibrium functions.[31] Following this, various model-implied moments are computed and compared to the data as will be explained. Since the empirical risk premium is not observed, I estimate its mean and volatility using return forecasting regression (37). I split the NBER recessionary periods into crisis (financial recession) and non-crisis (non-financial recession) periods based on the definition of Reinhart and Rogoff [2009a]. I then run predictive regressions with dividend yield (D_t/P_t) as the regressor and 1-year ahead stock returns as the dependent variable. Regression (I) in Table (4) uses just the dividend yield as regressor and the regressions (II), and (III) include a dummy for non-financial recession and financial crisis respectively. The dividend yield and stock return data are from Robert Shiller's website. I use monthly frequency from years 1945 till 2018. The indicator functions $\mathbb{1}_{Rec}$, and $\mathbb{1}_{fin}$ take a value 1 in months of NBER non-financial recession and financial recession respectively. The dummy variable corresponding to the financial crisis is positive and statistically significant.[32] The R-squared value is also higher controlling for recession and financial crises indicating a better predictive power. This confirms the finding in Muir [2017] that the risk premium is much higher during financial crises and the predictive power is improved by conditioning on the recessionary periods. I take the fitted value from regression (III) in Table (4) and compute the standard deviation to obtain the volatility of the risk premium.

$$R^e_{t+1} = a + \beta * D_t/P_t + \beta_{rec} * \mathbb{1}_{Rec} * D_t/P_t + \beta_{fin} * \mathbb{1}_{fin} * D_t/P_t + \epsilon_t \tag{37}$$

3.3 Tension between amplification and persistence of crises

A trade-off between the amplification and the persistence of financial crises arises in the benchmark model. One such channel that generates this trade-off is the risk aversion of the agents. The level of amplification required to match the empirical asset pricing moments leads to a two fold problem. First, the probability of a crisis implied by the model with high

[29] The percentage of NBER recessionary periods since the beginning of Federal Reserve (1914) is around 20%.

[30] A wide range is considered for the probability of crisis since it is undesirable to reject the success of model based on a metric for which there is substantial disagreement in the literature.

[31] See Appendix A.3.4 for details.

[32] This finding is robust to using different time periods such as 1871-2018 (time since Shiller's data is available), and 1914-2018 (since the start of Federal Reserve).

	(I)	(II)	(III)
const	-0.01	0.00	0.00
	(0.02)	(0.02)	(0.02)
D_t/P_t	2.70***	1.76**	1.75**
	(0.47)	(0.53)	(0.53)
$\mathbb{1}_{Rec}$		2.12***	1.77***
		(0.42)	(0.45)
$\mathbb{1}_{fin}$			2.02***
			(0.69)
N	876	876	876
R2	0.06	0.10	0.11

Table 4: Predictive regression. Model I sets both dummy variables to zero. Model II sets financial crisis dummy to zero. Model III uses both dummy variables.

risk aversion becomes too small to reconcile with the data. Second, and more importantly, higher the amplification, less persistent the crises episodes become. I first explain the trade-off between the amplification and the probability of crisis, and then explain how a higher amplification can be attained only at the cost of a lower persistence.

Figure (7) plots the unconditional risk premium, the conditional risk premium, and the probability of crisis. With a risk aversion equal to 1, the parameters in Table (3) lead to 7.8% probability of crisis. The unconditional mean risk premium is around 1.7%. One way to obtain even higher risk premium is by pumping up the risk aversion. However, increasing the risk aversion leads to the probability of crisis declining rapidly. As the values in Table (5) suggest, to obtain an empirically observed unconditional risk premium of 7.5%, the risk aversion has to be around 20. For this high level of risk aversion, the economy almost never enters into the crisis state. The reason is that a higher risk premium increases the wealth share of experts and therefore, a series of large negative shocks is required for the wealth share to diminish enough and push the system into the crisis zone. The standard deviation of the risk premium is 2.8% (see column 5 of Table (5)) which occurs solely due to the non-linearity in the model between the normal and the crisis regime. Since empirically estimating the risk premium in the crises episodes is a challenge, the calibration is performed to match the unconditional risk premium moments. The point is that while the comparative static plots in Figure (5) show a spike in the risk premium in some regions of the state space, if the dynamics of the model is such that these regions are barely reached, then the model cannot match the high risk premium in the data.

The persistence of financial crises is as much an important empirical phenomenon as the amplification. A direct measure of persistence is the duration. The average length of the crisis that the model can generate is around 5-8 months, which is much shorter than observed in the data. While there is disagreement regarding the empirical length of crises in the literature, the consensus is that it is larger than eight months.[33] Figure (8) plots the frequency distribution of the crisis length observed in the model. Most of the mass lies in periods less than 5 months and a crisis length of more than 10 months is probabilistically very small. The reason for this is that the only shocks in the model are Brownian, whose increments are i.i.d normal. Hence, a negative shock which impairs the intermediary wealth

[33] See He and Krishnamurthy [2013], Muir [2017] for example.

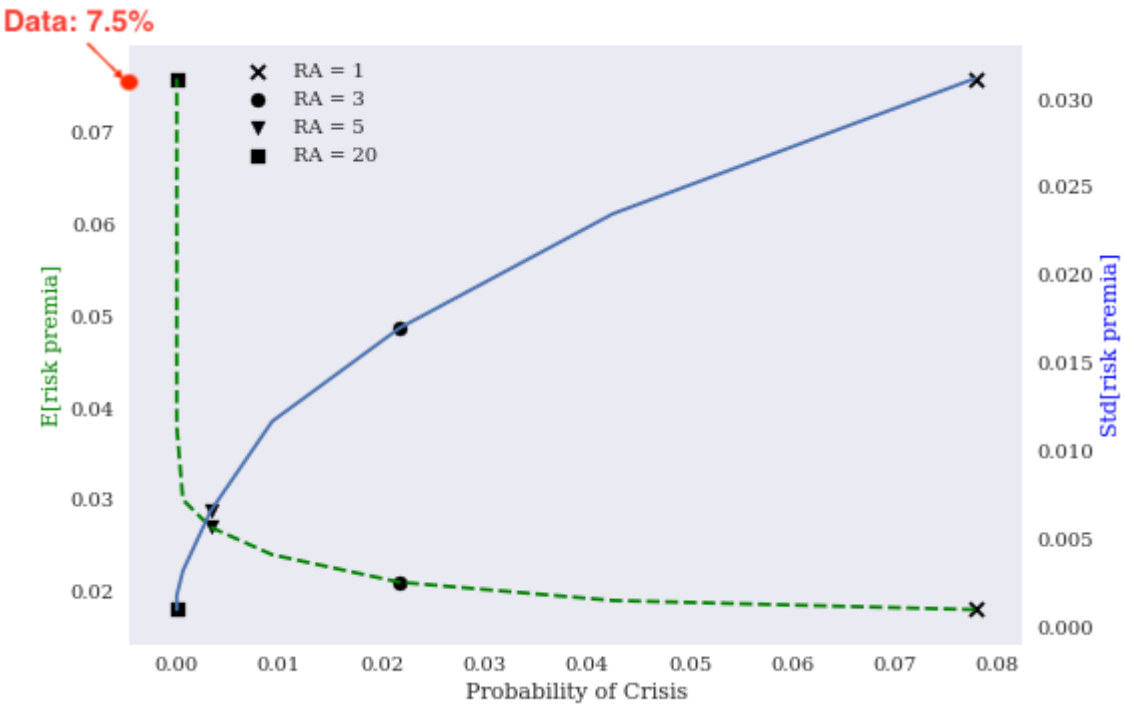

Figure 7: Trade off between the unconditional asset pricing moments and the probability of crisis for different risk aversion parameters (RA). The dashed line represents Expected risk premium (see left axis). The full line represents standard deviation of risk premium (see right axis). Risk aversion decreases from left to right.

	Data			Benchmark Model (RA=1)		Benchmark Model (RA = 20)	
	All	Recession	Crisis	All	Crisis	All	Crisis
E(Risk premium)	7.5	16.6	25.0	1.7	13.4	7.3	-
Std(Risk premium)	5.1	6.5	7.4	2.8	1.3	0	-
Probability of Crisis	7			7.8		0	

Table 5: Risk premium moments: model and data. All values are in annualized percentage.

share is on average followed by a positive shock which restores the lost wealth quickly. This is the case despite the model featuring leveraged experts. To be more concrete, imagine that the system has just entered the crisis period following a series of negative shocks. From Figure (5), the capital price is lower which puts a downward pressure on the net worth of intermediaries. However, the risk premium is higher and as the intermediaries operate with leverage, they earn more since they hold a larger proportion of risky capital. The latter effect is larger than the former and makes the drift of the wealth share high enough to push the system back to the normal regime. When risk aversion is higher, the risk premium effect is even larger resulting in the average length of the crisis to fall even more. In other words, a larger risk aversion creates higher amplification but dampens the persistence. Figure (8) shows the average length of crisis for different values of risk aversion. As the risk aversion increases, the mass of crisis length that lies in the range 1-2 months increases. As for the mass of crisis length that lies above 2 months, the opposite is true. This indicates that crises periods are far too infrequent when the agents are more risk averse. The dynamics explained above corroborates with this observation.

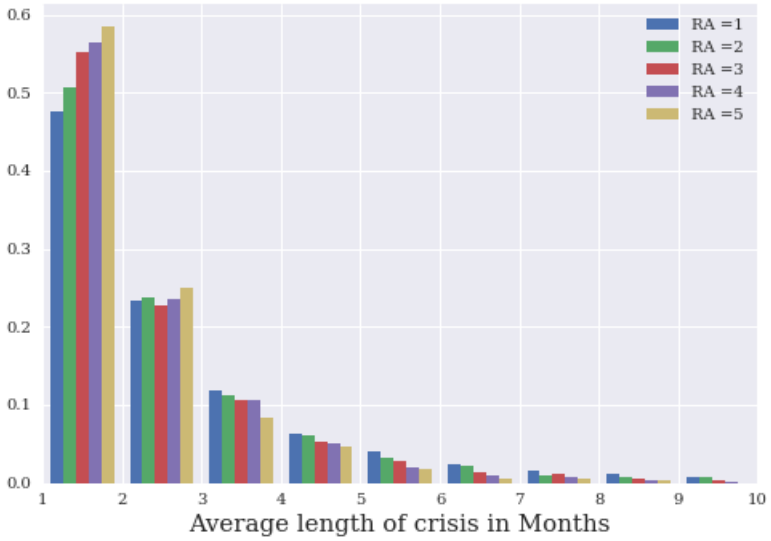

Figure 8: Frequency distribution of average crisis duration for different values of Risk aversion (RA). The graph shows only till months 10 since the frequency for months larger than 10 is negligible.

This tension between the persistence and the amplification is robust to the choices of any parameter values and utility functions. In the case of CRRA utility, and recursive utility with non-unitary IES, the consumption-wealth ratio is time varying and affects the drift of the wealth share in addition to the risk premium and the capital price. However, the effect of the risk premium highly dominates the other two effects and hence this tension is pervasive for more general preferences as well.[34] There are also other channels through which this tension becomes evident. In Appendix A.4, I show that decreasing the skin-in-the-game constraint leads to a more amplified crisis, but reduces the persistence. When the experts are constrained to keep a larger (smaller) fraction of the equity on their balance sheet, the risk premium becomes larger (smaller) in the crisis state, which increases (decreases) the wealth share of the experts leading to a quick (late) recovery. This indicates that the tension observed is not a matter of calibration. Regardless of how one calibrates the model to generate a high amplification to the extent that is observed in the data, the high risk premium in the amplified crisis state causes the experts to repair their balance sheets by quickly building sufficient capital, thereby failing to match the prolonged crisis that we see in the data.

Other moments: The benchmark model delivers an unconditional average GDP growth rate of around 2% and an investment rate of around 5-8%. Recall that the calibration is done to match the unconditional moments. Therefore, one measure of success of the model is to see how well it captures the non-linearity in the data. The GDP growth rate conditional on being in the crisis region is around -8%. The annualized GDP growth rate during the third quarter of 2008 was -8.2%. In this respect, the model captures the non-linearity quite well. However, the drop in investment rate implied by the model during the crisis is not sufficient to reconcile with the data. The private investment rate fell by 8% during the third quarter of 2008 whereas the model implied investment rate conditional on being in the crisis is 4%. Note that even though the output of experts and households individually move in sync with the capital due to the assumption of AK technology, the aggregate output depends on the aggregate dividend, which is a function of the capital share. During the crisis period, less productive households hold capital and hence the aggregate dividend drops to a large extent, and this causes the output to drop a lot as well. On the other hand, the investment rate is determined by the capital prices alone. The drop in capital price during the crisis period is not large enough to generate the observed drop in investment rate.[35] The volatility of investment rate implied by the model is close to zero. Overall, while the model captures the non-linearity in output growth, it misses out on the non-linearity in mean and volatility of investment rate. This result is similar to HK2019 who obtain a realistic consumption volatility but too low an investment volatility due to the assumption of AK technology. This calls for future work to match both output and investment dynamics. The mean leverage of intermediary sector implied by the model with unitary risk aversion is 3.23, comparable to the empirical leverage of 3.77.[36] The model also features a counter-cyclical leverage. Even though the experts fire sell the assets

[34] I experiment with log, CRRA, recursive utility with IES=1, and recursive utility with IES different from 1. Appendix A solves the benchmark model with these utility functions using the finite difference upwinding scheme. The results from simulation studies for the case of all utility functions are not included in the paper but they display the same tension between the persistence and the amplification that is explained in the paper.

[35] The result is not much quantitatively different if one assumes a quadratic functional form instead of logarithmic for the capital adjustment costs $\Phi(\cdot)$.

[36] This number is taken from HK2019.

to the households in periods of distress, the price of capital also drops, which depresses the experts' equity. Since the experts operate with leverage in equilibrium, the drop in expert equity is more than the drop in assets, which results in a rising leverage. Table (9) shows that the correlation between the shock and the leverage ranges from -19% to -22% for different risk aversion levels. This matches the empirical correlation of -18% quite well. However, as risk aversion increases, the level of leverage falls. Table (9) shows that with a risk aversion level as high as 10, the leverage is 1.43, well short of the empirical leverage of 3.77. Overall, for lower risk aversion levels, the model seems to do well in matching the leverage patterns. Lastly, the model does not generate excess asset return volatility (Shiller [1981]). The unconditional return volatility is more or less the same as the exogenous fundamental volatility of 6%, even though it shoots up in the crisis state. This is because the endogenous risk σ_t^q becomes zero in the normal regime. The conditional volatility, albeit high, is not large enough to make the unconditional one match the data.

Table (6) summarizes the ability of the benchmark model to succeed in different aspects. By far, matching the intermediary leverage pattern and the non-linearity in output growth seem to be the strongest suit of the model. The model cannot resolve the tension between unconditional risk premium, conditional risk premium, and persistence of crisis for any reasonable parameters in calibration. The focus of the next section is to provide a resolution to this problem.

	Quantity of interest	Success level	Comments
Macroeconomic	GDP/Output growth	High	✓
	Investment rate	Low	Low variation and not enough drop in crisis
Intermediary	Leverage	High	✓
	Cyclicality of leverage	High	✓
Crises	Probability of crises	Moderate	Matching prob. of crisis attenuates crisis dynamics
	Duration of crises	Low	Not enough persistence
Asset price	Conditional risk premium	High	✓
	Unconditional risk premium	Low	Matching unconditional risk premium attenuates prob. of crisis
	Std. of risk premium	Moderate	-
	Conditional volatility	High	✓
	Unconditional volatility	Low	Shiller puzzle

Table 6: Model success summary.

4 Resolution of the tension between amplification and persistence of crises

In this section, I quantify the model with stochastic productivity and exit rate of experts and show that it resolves the tension between the persistence and the amplification of financial crises, and provides reasonable time variation in the prices. Figure (9) shows a simulated sample path for the expert productivity. I assume a low mean reversion rate ($\pi = 1\%$) to generate paths that resemble a regime switching process. Positive shocks will push the productivity to the upper limit of 20%, whereas a series of negative shocks pushes it to a lower limit of 10%. Due to the low mean reversion rate, it takes a long time for the process to switch towards the other limit. I assume that the system is in the crisis state whenever $a_{e,t}$ is below its mean value of $\hat{a}$ and the wealth share is below the crisis threshold z^*.[37] The right panel of Figure (9) plots the stationary distribution of the wealth share obtained through simulation.[38]

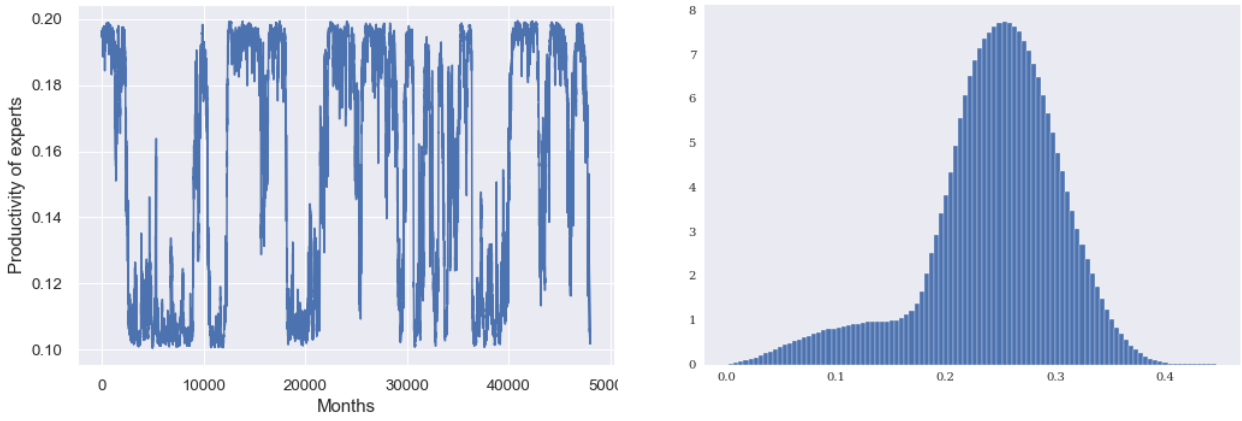

Figure 9: Left panel: Sample path of expert productivity. Right panel: Stationary distribution of wealth share. Both correspond to the model with stochastic productivity.

Table (7) presents the average duration of the crisis in the benchmark model and the stochastic productivity model and compares it against the data. There is substantial controversy in the literature regarding the duration of crises (Reinhart and Rogoff [2009a]). The NBER reports that the Great Recession started at December 2007 and ended at June 2009, indicating an 18 month duration.[39] To facilitate comparisons, I adjust the parameters to generate a comparable probability of the crisis in the range of 7-8% across the the benchmark and the richer model. The numbers in Table (7) can be thought of as the ability of various models to generate the stated duration for a reasonable crisis probability. Both of the benchmark models deliver a duration of crisis that is much lower than in the data. The mean duration from the richer model matches the data quite well although the

[37] The simulation results show that the system does enter the crisis region mostly when the productivity is well below its mean.

[38] The simulation method is similar to the benchmark model except that the equilibrium quantities are two-dimensional.

[39] The average duration of recession in the past 33 cycles from year 1854 to 2009 is 17.5 months. Source: https://www.nber.org/cycles.html.

10th and 50th percentile values are lower. The parameters used for calibrating the richer model are shown in Table (2). Figure (10) plots the frequency distribution of the wealth

	Data (NBER)	Benchmark model (RA=1, IES = 1)	Benchmark model (RA=2, IES = 1.5)	My model (RA=5, IES=1)
10th percentile	8.0	1.0	1.0	1.0
50th percentile	13.5	2.0	2.0	3.0
90th percentile	31.2	13.0	16.0	49.0
Mean	17.5	5.8	6.5	17

Table 7: The empirical data for duration of crisis (months) is from NBER website. The model implied duration (months) is obtained for a 6% probability of crisis fixed across models to facilitate comparison.

share during the time the system spends in the crisis region. In the benchmark model (left panel), a lot of the mass lies near the crisis boundary of 0.11. This indicates that once the system enters the crisis region, the experts gain wealth quickly and revert to the normal region. In contrast, the frequency distribution of the wealth share in the crisis region in my model, as shown in the right panel in Figure (10), features fatter tails. A negative shock that hits the capital also lowers the productivity of the experts which in turn reduces the risk premium that the experts earn. This puts a downward pressure on the drift of the wealth share and helps to achieve a realistic probability of the crisis even for higher risk aversion levels. The fact that the crisis zone features both a lower wealth share and a lower productivity of experts can be seen in the right panel of Figure (11). Moreover, once the system enters the crisis region, exit rate of the experts shoots up and generates persistent crises by depressing the drift of the wealth share. Most of the mass lie in the range of 10-11%. After spending a sufficient amount of time in the crisis zone, the expert productivity reverts to upper level, taking the system out to the normal regime. The variation in the investment rate, and the risk free rate is also higher compared to the benchmark model due to the assumption of time varying experts productivity. The model implied unconditional volatility of the risk premium is 5.3%, well in line with the empirical value of 5.1% reported in Table (5). Overall, my model does a good job of balancing the persistence and the amplification, and delivers a reasonable time variation in the prices.

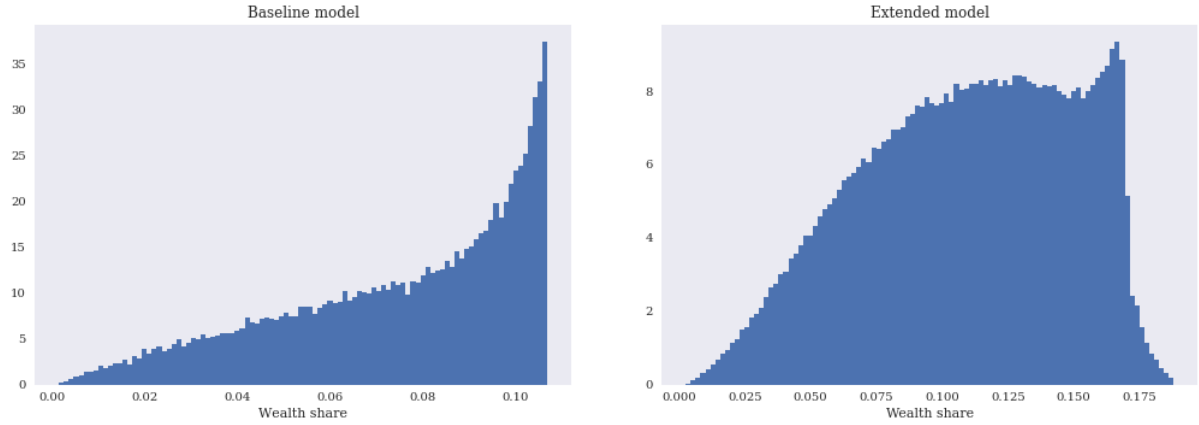

Figure 10: Left panel: Tail of experts wealth share distribution from the benchmark model. Right panel: Tail of experts wealth share distribution from the model with stochastic productivity.

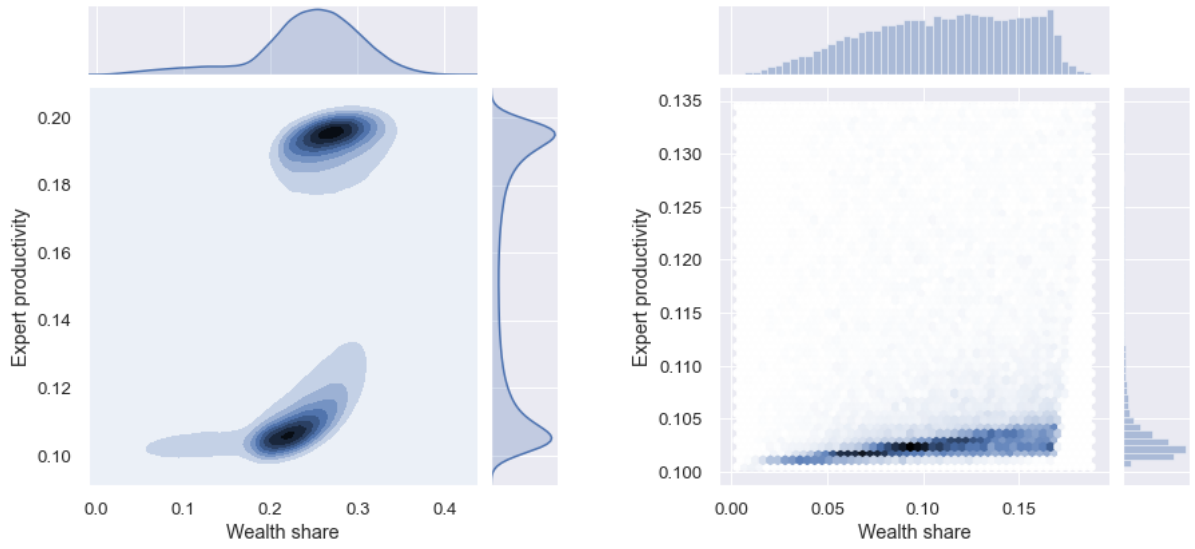

Figure 11: Left panel: Joint density of wealth share and productivity of experts along with respective marginals. Right panel: Joint density of wealth share and productivity of experts along with respective marginals in crisis region.

	All	Crisis	Normal	All	Crisis	Normal
E[leverage]	2.80	4.79	2.62	3.23	5.50	3.10
E[inv. rate]	7.70%	2.80%	8.20%	6.00%	5.00%	6.00%
E[risk free rate]	0.90%	-7.20%	1.70%	4.80%	0.00%	5.00%
E[risk premia]	6.70%	17.50%	5.70%	1.70%	13.40%	1.00%
E[price]	1.39	1.14	1.41	1.42	1.34	1.42
E[return volatility]	7.14%	11.63%	6.83%	6.20%	15.80%	5.70%
E[GDP growth rate]	1.20%	-8.00%	1.90%	2.30%	-7.90%	2.70%
Std[inv. rate]	3.18%	1.31%	2.91%	0.36%	1.09%	0.11%
Std[risk premia]	5.35%	1.57%	4.45%	2.82%	1.31%	0.18%
Std[risk free rate]	3.98%	1.64%	3.21%	1.19%	0.42%	0.28%
Std[GDP growth rate]	11.40%	21.00%	9.59%	7.17%	19.63%	5.15%
Corr(leverage,shock)	-0.25	-0.17	-0.30	-0.28	-0.05	-0.25
Corr(price return, riskf ree rate)	0.16	-0.25	0.18	0.20	0.01	-0.23
Corr(risk premia, volatility)	0.98	0.76	0.96	0.98	-0.34	0.57
Probability of crisis	7.66%			7.80%		
Duration of crisis (months)	17.0			6		

Table 8: Summary of moments from the model with stochastic productivity and the benchmark model with RA=1.

5 Conclusion

Financial recessions are typically characterized with high risk premium and slow revival. I have built a macro-finance asset pricing model with intermediaries facing productivity shocks and regime-dependent exit rate. A sequence of bad shocks reduces the net worth of the leveraged intermediaries, and at the same time diminishes their comparative advantage over the households in terms of the productivity differential. A simpler model with constant intermediary productivity and no exit rate cannot simultaneously generate amplified and persistent financial crises. There is a trade-off between the conditional risk premium, the unconditional risk premium, and the probability and duration of crises. I show that any auxiliary feature that improves the financial amplification channel dampens the persistence of crisis. The model is successful in capturing the non-linearity in output and the intermediary leverage patterns. However, the capital price does not fall enough to match the empirical negative investment rates in times of distress.

The richer model with stochastic productivity and exit rate of the intermediaries can resolve this tension and quantitatively generate a high risk premium, a large drop in output, decreased financial intermediation, pro-cyclical leverage, and prolonged distress periods. The model also generates a large time variation in the investment rate, which is absent in the benchmark model. While my model generates a higher percentage drop of investment than the benchmark model in the crisis region, the dip is still not enough to reconcile with the empirically large disinvestment during financial crises. An interesting avenue for future research is to build a model that can accommodate realistic investment dynamics in addition to the output and the asset pricing dynamics.

I have utilized a novel method of solving the model based on active machine learning that encodes the economic information as regularizers in a deep neural network. The algorithm is scalable and has the potential to solve high dimensional problems with less effort in coding, opening up new avenues to model asset pricing with frictions in potentially large dimensions.

References

Yves Achdou, Francisco J. Buera, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. PDE Models in Macroeconomics. *Proceedings of the Royal Society*, 2014a.

Yves Achdou, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. Heterogeneous agent models in continuous time. *Preprint*, 2014b.

Yves Achdou, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. Income and wealth distribution in macroeconomics: A continuous-time approach. *NBER Working Paper Series*, 2017. doi: 10.3386/w23732.

Tobias Adrian and Nina Boyarchenko. Intermediary Leverage Cycles and Financial Stability. *SSRN Electronic Journal*, (August 2012), 2012. doi: 10.2139/ssrn.2133385.

Tobias Adrian, Erkko Etula, and Tyler Muir. Financial Intermediaries and the Cross-Section of Asset Returns. *Journal of Finance*, 69(6):2557–2596, 2014. ISSN 15406261. doi: 10.1111/jofi.12189.

Marlon Azinovic, Luca Gaegauf, and Simon Scheidegger. Deep equilibrium nets. *SSRN Electronic Journal*, 2019.

Ravi Bansal and Amir Yaron. Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance*, 59(4):1481–1509, 2004. ISSN 00221082. doi: 10.1111/j.1540-6261.2004.00670.x.

Ravi Bansal, Dana Kiku, Ivan Shaliastovich, and Amir Yaron. Volatility, the Macroeconomy, and Asset Prices. *Journal of Finance*, 69(6):2471–2511, 2014. ISSN 15406261. doi: 10.1111/jofi.12110.

Matthew Baron and Wei Xiong. Credit expansion and neglected crash risk. *Quarterly Journal of Economics*, 132(2):713–764, 2017. ISSN 15314650. doi: 10.1093/qje/qjx004.

Robert J. Barro. Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics*, 121(3):823–866, 2006. ISSN 00335533. doi: 10.1162/qjec.121.3.823.

Jason Beeler. The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment. *Critical Finance Review*, 2012. ISSN 21645744. doi: 10.1561/104.00000004.

Allen N. Berger and Loretta J. Mester. Inside the black box: What explains differences in the efficiencies of financial institutions? *Journal of Banking & Finance*, 21(7):895–947, jul 1997. URL https://linkinghub.elsevier.com/retrieve/pii/S0378426697000101.

J. Frédéric Bonnans, Élisabeth Ottenwaelter, and Housnaa Zidani. A fast algorithm for the two dimensional HJB equation of stochastic control. *Mathematical Modelling and Numerical Analysis*, 2004. ISSN 0764583X. doi: 10.1051/m2an:2004034.

M. K. Brunnermeier and Y. Sannikov. Macro, Money, and Finance: A Continuous-Time Approach. In *Handbook of Macroeconomics*. 2016. ISBN 9780444594877. doi: 10.1016/bs.hesmac.2016.06.002.

Markus K. Brunnermeier and Yuliy Sannikov. A macroeconomic model with a financial sector. *American Economic Review*, 104(2):379–421, 2014. ISSN 00028282. doi: 10.1257/aer.104.2.379.

John Y. Campbell and John H. Cochrane. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 1999. ISSN 00223808. doi: 10.1086/250059.

Luyang Chen, Markus Pelger, and Jason Zhu. Deep Learning in Asset Pricing. *SSRN Electronic Journal*, 2019. ISSN 1556-5068. doi: 10.2139/ssrn.3350138.

Adrien D'Avernas and Quentin Vandeweyer. A Solution Method for Continuous-Time Models. pages 1–33, 2019.

Sebastian Di Tella. Uncertainty shocks and balance sheet recessions. *Journal of Political Economy*, 125(6):2038–2081, 2017. ISSN 1537534X. doi: 10.1086/694290.

Victor Duarte. Machine Learning for Continuous-Time Economics. 2017.

Guohua Feng and Apostolos Serletis. Efficiency, technical change, and returns to scale in large US banks: Panel data evidence from an output distance function satisfying theoretical regularity. *Journal of Banking and Finance*, 34(1):127–138, 2010. ISSN 03784266. doi: 10.1016/j.jbankfin.2009.07.009.

Jesús Fernández-Villaverde, Samuel Hurtado, and Galo Nuno. Financial Frictions and the Wealth Distribution. *SSRN Electronic Journal*, 2020. ISSN 1556-5068. doi: 10.2139/ssrn.3615695.

Nicolae Gârleanu and Stavros Panageas. Young, old, conservative, and bold: The implications of heterogeneity and finite lives for asset pricing. *Journal of Political Economy*, 123(3):670–685, 2015. ISSN 1537534X. doi: 10.1086/680996.

Mark Gertler, Nobuhiro Kiyotaki, and Andrea Prestipino. A Macroeconomic Model with Financial Panics. *American Historical Review*, 2019. ISSN 00028762. doi: 10.1093/restud/rdz032.

Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Journal of Machine Learning Research*, 2010a.

Xavier Glorot and Yoshua Bengio. Xavier Initialization. *Journal of Machine Learning Research*, 2010b. ISSN 15324435.

Matthieu Gomez. Asset Prices and Wealth Inequality. *Working Paper*, 2019.

Goutham Gopalakrishna. Informed Neural Nets and Deep Recessions. *Work in progress*, 2020.

Gary Gorton and Guillermo Ordoñez. Collateral crises. *American Economic Review*, 2014. ISSN 00028282. doi: 10.1257/aer.104.2.343.

Gary Gorton and Guillermo Ordoñez. Good Booms, Bad Booms. *Journal of the European Economic Association*, 2020. ISSN 1542-4766. doi: 10.1093/jeea/jvy058.

Shihao Gu, Bryan Kelly, and Dacheng Xiu. Autoencoder asset pricing models. *Journal of Econometrics*, 2020. ISSN 18726895. doi: 10.1016/j.jeconom.2020.07.009.

Fatih Guvenen. A Parsimonious Macroeconomic Model for Asset Pricing: Habit Formation or Cross-sectional Heterogeneity? *SSRN Electronic Journal*, 2005. doi: 10.2139/ssrn.641141.

Lars Peter Hansen, Paymon Khorrami, and Fabrice Tourre. Comparative Valuation Dynamics in Models with Financing Restrictions. (February), 2018. URL http://www.zccfe.uzh.ch/dam/jcr:c4a09093-fc00-4953-9f31-701dc0cd3fab/macro-model-print.pdf.

Zhiguo He and Arvind Krishnamurthy. Intermediary asset pricing. *American Economic Review*, 103(2):732–770, 2013. ISSN 00028282. doi: 10.1257/aer.103.2.732.

Zhiguo He and Arvind Krishnamurthy. A macroeconomic framework for quantifying systemic risk. *American Economic Journal: Macroeconomics*, 11(4):1–37, 2019. ISSN 19457715. doi: 10.1257/mac.20180011.

Zhiguo He, Bryan Kelly, and Asaf Manela. Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics*, 126(1):1–35, 2017a. ISSN 0304405X. doi: 10.1016/j.jfineco.2017.08.002.

Zhiguo He, Bryan Kelly, and Asaf Manela. Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics*, 126(1):1–35, 2017b. ISSN 0304405X. doi: 10.1016/j.jfineco.2017.08.002. URL https://doi.org/10.1016/j.jfineco.2017.08.002.

Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 1989. ISSN 08936080. doi: 10.1016/0893-6080(89)90020-8.

Joseph P. Hughes, Loretta J. Mester, and Choon Geol Moon. Are scale economies in banking elusive or illusive?: Evidence obtained by incorporating capital structure and risk-taking into models of bank production. *Journal of Banking and Finance*, 25(12):2169–2208, 2001. ISSN 03784266. doi: 10.1016/S0378-4266(01)00190-X.

Nobuhiro Kiyotaki and John Moore. Credit cycles. *Journal of Political Economy*, 1997. ISSN 00223808. doi: 10.1086/262072.

Arvind Krishnamurthy and Wenhao Li. Dissecting Mechanisms of Financial Crises: Intermediation and Sentiment. *SSRN Electronic Journal*, 2020. ISSN 1556-5068. doi: 10.2139/ssrn.3554788.

Arvind Krishnamurthy and Tyler Muir. How Credit Cycles across a Financial Crisis. *National Bureau of Economic Research*, 2017. ISSN 0898-2937. doi: 10.3386/w23850.

Pablo Kurlat. How I Learned to Stop Worrying and Love Fire Sales. *National Bureau of Economic Research*, 2018. ISSN 0898-2937. doi: 10.3386/w24752.

Wenhao Li. Public liquidity, bank runs, and financial crises. *Working Paper*, 2020.

Rajnish Mehra and Edward C. Prescott. The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2):145–161, 1985. ISSN 03043932. doi: 10.1016/0304-3932(85)90061-3.

Benjamin Moll. Productivity losses from financial frictions: Can self-financing undo capital misallocation? *American Economic Review*, 104(10):3186–3221, 2014. ISSN 00028282. doi: 10.1257/aer.104.10.3186.

John Moore and Nobuhiro Kiyotaki. Credit Cycles Author (s): Nobuhiro Kiyotaki and John Moore Nobuhiro Kiyotaki. *Journal of Political Economy*, 105(2):211–248, 1997.

Alan Moreira and Alexi Savov. The Macroeconomics of Shadow Banking. *Journal of Finance*, 72(6):2381–2432, 2017. ISSN 15406261. doi: 10.1111/jofi.12540.

Tobias J. Moskowitz, Yao Hua Ooi, and Lasse Heje Pedersen. Time series momentum. *Journal of Financial Economics*, 2012. ISSN 0304405X. doi: 10.1016/j.jfineco.2011.11.003.

Tyler Muir. Financial crises and risk premia. *Quarterly Journal of Economics*, 132(2):765–809, 2017. ISSN 15314650. doi: 10.1093/qje/qjw045.

Gregory Phelan. Financial intermediation, leverage, and macroeconomic instability. *American Economic Journal: Macroeconomics*, 2016. ISSN 19457715. doi: 10.1257/mac.20140233.

Carmen M. Reinhart and Kenneth S. Rogoff. *This time is different: Eight centuries of financial folly*. 2009a. ISBN 9780691152646. doi: 10.1080/09585206.2010.512722.

Carmen M. Reinhart and Kenneth S. Rogoff. The aftermath of financial crises. *American Economic Review*, 99(2):466–472, 2009b. ISSN 00028282. doi: 10.1257/aer.99.2.466.

Burr Settles. Active learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 2012. ISSN 19394608. doi: 10.2200/S00429ED1V01Y201207AIM018.

Robert J Shiller. Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends? *American Economic Review*, 1981.

Andrei Shleifer and Robert Vishny. Fire sales in finance and macroeconomics. *Journal of Economic Perspectives*, 25(1):29–48, 2011. ISSN 08953309. doi: 10.1257/jep.25.1.29.

David C. Wheelock and Paul W. Wilson. Do Large Banks Have Lower Costs? New Estimates of Returns to Scale for U.S. Banks. *Journal of Money, Credit and Banking*, 44(1):171–199, 2012a. ISSN 00222879. doi: 10.1111/j.1538-4616.2011.00472.x.

David C. Wheelock and Paul W. Wilson. Do Large Banks Have Lower Costs? New Estimates of Returns to Scale for U.S. Banks. *Journal of Money, Credit and Banking*, 44(1):171–199, 2012b. ISSN 00222879. doi: 10.1111/j.1538-4616.2011.00472.x.

A Appendix

A.1 Model with stochastic productivity

A.1.1 Proof of the Asset pricing conditions:

The expected return that the experts earn from investing in the capital is given by

$$dr_t^v = (\mu_{e,t}^R - (1-\chi_t)\epsilon_{h,t})dt + \chi_t(\sigma_t^{q,k} + \sigma)dZ_t^k + \chi_t\sigma_t^{q,a}dZ_t^a$$

where $\epsilon_{h,t} = \zeta_{h,t}^k(\sigma_t^{q,k} + \sigma) + \zeta_{h,t}^a\sigma_t^{q,a} + \varphi(\zeta_{h,t}^a(\sigma + \sigma_t^{q,a}) + \sigma_t^{q,a}\zeta_{h,t}^k)$. That is, $(1-\chi_t)\epsilon_{h,t}$ is the part of the expected excess return that is paid by the experts to the outside equity holders, which is netted out. Since the experts hold a fraction χ_t of the inside equity, the volatility terms are multiplied by this quantity. Consider a trading strategy of investing \$1 into the capital at time 0. Let v_t be the value of this investment strategy at time t. Then, we have $\frac{dv_t}{v_t} = dr_t^v$, and

$$\frac{d(\xi_e v_t)}{\xi_e v_t} = (-r_t + \mu_{e,t}^R - (1-\chi_t)\epsilon_{h,t} - \chi_t\epsilon_{e,t})dt + \text{diffusion terms}$$

where $\epsilon_{e,t} = \zeta_{e,t}^k(\sigma + \sigma_t^{q,k}) + \zeta_{e,t}^a\sigma_t^{q,a} + \varphi(\zeta_{e,t}^a(\sigma + \sigma_t^{q,k}) + \zeta_{e,t}^k\sigma_t^{q,a})$, and $\xi_{e,t}$ follows the process in (5). Since $\xi_e v_t$ is a martingale, the drift equals to zero, which implies

$$\mu_{e,t}^R - r_t = \chi_t\epsilon_{e,t} + (1-\chi_t)\epsilon_{h,t}$$

■

A.1.2 Proof of Proposition 1

The law of motion of wealth for the households and the experts are given by equation (35). Using the law of large numbers to aggregate the wealth of individual household and expert, we get

$$\frac{dw_{h,t}}{w_{h,t}} = \left(r_t - \frac{c_{h,t}}{w_{h,t}} - \lambda_d + \theta_{h,t}(\mu_{h,t}^R - r_t) + \frac{(1-\bar{z})\lambda_d}{1-z_t}\right)dt + \theta_{h,t}(\sigma + \sigma_t^q)dZ_t^k + \theta_{h,t}\sigma_t^a dZ_t^a$$
$$\frac{dw_{e,t}}{w_{e,t}} = \left(r_t - \frac{c_{e,t}}{w_{e,t}} - \lambda_d + \theta_{e,t}\epsilon_{e,t} + \frac{\bar{z}\lambda_d}{z_t}\right)dt + \theta_{e,t}(\sigma + \sigma_t^{q,k})dZ_t^k + \theta_{e,t}\sigma_t^{q,a}dZ_t^a$$

where $w_{h,t} = \int_{j\in H} w_{j,t}$ and $w_{e,t} = \int_{j\in E} w_{j,t}$ denotes aggregated wealth among respective group, $z_t = \frac{w_{e,t}}{w_{h,t}+w_{e,t}}$, and $\theta_{e,t} := \frac{\chi_t\psi_t}{z_t}, \theta_{h,t} := \frac{1-\chi_t\psi_t}{z_t}$. By Ito's lemma, the dynamics of the wealth share becomes

$$\frac{dz_t}{z_t} = \frac{dw_{e,t}}{w_{e,t}} - \frac{d(q_tk_t)}{q_tk_t} + \frac{d\langle q_tk_t, q_tk_t\rangle}{(q_tk_t)^2} - \frac{d\langle q_tk_t, w_{e,t}\rangle}{(q_tk_tw_{e,t})}$$

where

$$\frac{d(q_t k_t)}{q_t k_t} = (\epsilon_{e,t}(\sigma + \sigma_t^q) - \frac{(a_{e,t} - \iota_t)}{q_t} + r_t)dt + (\sigma + \sigma_t^{q,k})dZ_t^k + \sigma_t^{q,a} dZ_t^a$$
$$\frac{d\langle q_t k_t, q_t k_t\rangle}{(q_t k_t)^2} = (\sigma_t^{q,k} + \sigma)^2 + (\sigma_t^{q,a})^2 + 2\varphi(\sigma_t^{q,k} + \sigma)\sigma_t^{q,a} dt$$
$$\frac{d\langle q_t k_t, w_{e,t}\rangle}{q_t k_t w_{e,t}} = \big(\theta_{e,t}(\sigma_t^{q,k} + \sigma)^2 + \theta_{e,t}(\sigma_t^{q,a})^2 + 2\varphi(\sigma_t^{q,k} + \sigma)\sigma_t^{q,a}\big)dt$$

and the result follows from here after some algebra. ■

Note that we can write $\theta_{e,t}\epsilon_{e,t} = \theta_{e,t}\chi_t^{-1}(\mu_{e,t}^R - r_t - (1-\chi_t)\epsilon_{h,t})$ from the asset pricing condition in A.1.1, which allows us to write the experts wealth dynamics as

$$\frac{dw_{e,t}}{w_{e,t}} = \big(r_t - \frac{c_{e,t}}{w_{e,t}} - \lambda_d + \frac{\psi_t}{z_t}(\mu_{e,t}^R - r_t) - (1-\chi_t)\frac{\psi_t}{z_t}\epsilon_{h,t} + \frac{\bar{z}\lambda_d}{z_t}\big)dt + \frac{\chi_t\psi_t}{z_t}(\sigma + \sigma_t^{q,k})dZ_t^k + \frac{\chi_t\psi_t}{z_t}\sigma_t^{q,a}dZ_t^a$$

A.1.3 Proof of Proposition 2

The value function conjecture is

$$U_{j,t} = J_{j,t}(z_t, a_t)\frac{k_{j,t}^{1-\gamma_j}}{1-\gamma_j}$$

where $J_{j,t}$ follows the stochastic differential equation $\frac{dJ_{j,t}}{J_{j,t}} = \mu_{j,t}^J dt + \sigma_{j,t}^{J,k} dZ_t^k + \sigma_{j,t}^{J,a} dZ_t^a$ whose drift and volatility needs to be determined in the equilibrium. The HJB equation is derived directly in terms of the capital k_t instead of the wealth share z_t. The value function derivatives are

$$\frac{\partial U_{j,t}}{\partial J_{j,t}} = \frac{k_{j,t}^{1-\gamma_j}}{1-\gamma_j}; \quad \frac{\partial U_{j,t}}{\partial k_{j,t}} = J_{j,t}k_{j,t}^{-\gamma_j} \qquad (38)$$
$$\frac{\partial^2 U_{j,t}}{\partial J_{j,t}^2} = 0; \quad \frac{\partial^2 U_{j,t}}{\partial k_{j,t}^2} = -\gamma_j J_{j,t}k_{j,t}^{-(1+\gamma_j)}; \quad \frac{\partial^2 U_{j,t}}{\partial J_{j,t}\partial k_{j,t}} = k_{j,t}^{-\gamma_j}$$

Applying Ito's lemma to $J_{j,t}$ and using HJB, we get

$$\sup_{c,\theta} \; -\rho\frac{J_{j,t}k_{j,t}^{1-\gamma_j}}{1-\gamma_j} + \frac{c_{j,t}^{1-\gamma_j}}{1-\gamma_j} + \frac{J_{j,t}k_{j,t}^{1-\gamma_j}}{1-\gamma_j}\mu_{j,t}^J + J_{j,t}k_{j,t}^{1-\gamma_j}(\Phi(\iota_t) - \delta) \qquad (39)$$
$$- \sigma^2\frac{\gamma_j}{2}J_{j,t}k_{j,t}^{1-\gamma_j} + J_{j,t}k_{j,t}^{1-\gamma_j}(\sigma\sigma_{j,t}^{J,k} + \varphi\sigma\sigma_{j,t}^{J,a}) = 0$$

At the optimum, the marginal utilities of wealth and consumption become equal. Using the value function expression in terms of the wealth, we have

$$\frac{\partial U_{j,t}}{\partial w_{j,t}} = \frac{\partial f}{\partial c_{j,t}}$$
$$\tilde{J}_{j,t} w_{j,t}^{-\gamma_j} = (1-\gamma_j)\rho_j \frac{U_{j,t}}{c_{j,t}} \implies \frac{c_{j,t}}{w_{j,t}} = \rho_j$$

The stochastic discount factor for recursive utility is given by

$$\xi_{j,t} = \exp\left(\int_0^t \frac{\partial f(c_{j,s}, U_{j,s})}{\partial U} ds \right) \frac{\partial U_{j,t}}{\partial w_{j,t}}$$

Writing the value function conjecture in terms of the wealth instead of the capital, we have

$$U_{j,t} = \tilde{J}_{j,t} \frac{w_{j,t}^{1-\gamma_j}}{1-\gamma_j}; \quad f(c_{j,t}, U_{j,t}) = (1-\gamma_j)\rho_j U_{j,t}\big(\log\rho_j - \frac{1}{1-\gamma_j}\tilde{J}_{j,t}\big)$$

where $\tilde{J}_{j,t} = \dfrac{J_{j,t}}{(q_t z_t)^{1-\gamma_j}}$. The SDF then becomes

$$\xi_{j,t} = (1-\gamma_j)\exp\left(\int_0^t \big[\rho_j\big((1-\gamma_j)\log c_{j,s} - \log\big((1-\gamma_j)U_{j,s}\big) - 1\big] ds \right) \frac{U_{j,t}}{w_{j,t}}$$

This implies that $\sigma(\xi_{j,t}) = \sigma\left(\dfrac{U_{j,t}}{w_{j,t}}\right)$. To compute the R.H.S., we have to compute the dynamics of $\partial\left(\dfrac{U_{j,t}}{w_{j,t}}\right) := \partial v(J_{j,t}, z_t, q_t, k_{j,t})$. Using the derivatives

$$\frac{1}{v}\frac{\partial v}{\partial J_{j,t}} = \frac{1}{J_{j,t}}; \quad \frac{1}{v}\frac{\partial v}{\partial z_t} = \frac{1}{z_t}$$
$$\frac{1}{v}\frac{\partial v}{\partial q} = \frac{-1}{q_t}; \quad \frac{1}{v}\frac{\partial v}{\partial k_{j,t}} = \frac{-\gamma_j}{k_{j,t}}$$

and applying Ito's lemma, we get

$$\frac{\partial v}{v} = \underbrace{[\ldots\ldots]}_{\text{drift term}} dt + (\sigma_{j,t}^{J,k} dZ_t^k + \sigma_{j,t}^{J,a} dZ_t^a) - (\sigma_{j,t}^{Z,k} dZ_t^k + \sigma_{j,t}^{z,k} dZ_t^a) - (\sigma_t^{q,k} dZ_t^k + \sigma_t^{q,a} dZ_t^a) - \gamma_j \sigma dZ_t^k$$

Collecting the diffusion terms, using $\sigma_{e,t}^{z,i} = \sigma_t^{z,i}, \sigma_{h,t}^{z,i} = -\dfrac{1}{1-z_t}\sigma_t^{z,i}; i \in \{k, a\}$, and comparing it to the SDF equation

$$\frac{d\xi_{j,t}}{\xi_{j,t}} = -r_t dt - \zeta_{j,t}^k dZ_t^k - \zeta_{j,t}^a dZ_t^a$$

we get the desired result. ■

Plugging in the consumption-wealth ratio into the HJB equation (65), we obtain the

expressions for $\mu^J_{j,t}$

$$\mu^J_{e,t} = (\gamma_e - 1)\big(\rho_e \log \rho_e + \log(q_t z_t)\big) - \rho_e \log J_{e,t} - (1-\gamma_e)(\Phi(\iota_t) - \delta - \frac{\gamma_e}{2}\sigma^2 + \sigma\sigma^{J,k}_{e,t} + \varphi\sigma\sigma^{J,a}_{h,t})$$
$$\mu^J_{h,t} = (\gamma_h - 1)\big(\rho_h \log \rho_h + \log(q_t(1-z_t))\big) - \rho_h \log J_{h,t} - (1-\gamma_h)(\Phi(\iota_t) - \delta - \frac{\gamma_h}{2}\sigma^2 + \sigma\sigma^{J,k}_{h,t} + \varphi\sigma\sigma^{J,a}_{h,t}) \tag{40}$$

A.1.4 Proof of Proposition 3

Applying Ito's lemma to $q(z_t, a_t)$, we have

$$\partial q_t = \frac{\partial q_t}{\partial z_t} dz_t + \frac{\partial q_t}{\partial a_t} da_t + \frac{1}{2}\frac{\partial^2 q_t}{\partial z_t^2} d\langle z_t, z_t\rangle + \frac{1}{2}\frac{\partial^2 q_t}{\partial a_{e,t}^2} d\langle a_{e,t}, a_{e,t}\rangle + \frac{\partial^2 q_t}{\partial z_t \partial a_{e,t}} d\langle z_t, a_{e,t}\rangle$$

Matching the drift and the volatility terms, we get

$$\begin{aligned}
\mu_{q,t} &= \frac{\partial q_t}{\partial z_t}\frac{1}{q_t}\mu^z_t + \frac{\partial q_t}{\partial a_{e,t}}\mu_{ae,t} + \frac{1}{2}\frac{\partial^2 q_t}{\partial z_t^2}\big((\sigma^{z,k}_t)^2 + (\sigma^{z,a}_t)^2 + 2\varphi\sigma^{z,k}_t\sigma^{z,a}_t\big) \\
&\quad + \frac{1}{2}\frac{\partial^2 q_t}{\partial a_t^2}\sigma^2_{ae,t} + \frac{\partial^2 q_t}{\partial z_t \partial a_t}\big(\varphi\sigma^{z,k}_t\sigma_{ae,t} + \sigma^{z,a}_t\sigma_{ae,t}\big) \\
\sigma^{q,k}_t &= \frac{\partial q_t}{\partial z_t}\frac{1}{q_t}\sigma^{z,k}_t \\
\sigma^{q,a}_t &= \frac{\partial q_t}{\partial z_t}\frac{1}{q_t}\sigma^{z,a}_t + \frac{\partial q_t}{\partial a_{e,t}}\frac{1}{q_t}\sigma_{ae,t}
\end{aligned}$$

where $\sigma_{a,t} = \nu(\bar{a}_e - a_t)(a_t - \underline{a}_t)$ and $\mu_{ae,t} = \pi(\hat{a}_e - a_{e,t})$ Plugging in the expression for $\sigma^{z,k}_t$ and $\sigma^{z,a}_t$ from the dynamics of wealth share (13) in the above equation and rearranging, we get the result. ■

A.1.5 Numerical solution

Static step: We need to solve for the equilibrium quantities $\{\psi_t, (\sigma + \sigma^{q,k}_t), \sigma^{q,a}_t, q_t\}$. The other equilibrium quantities $\theta_{e,t}, \theta_{h,t}, \zeta^k_{e,t}, \zeta^a_{e,t}, \zeta^k_{h,t}, \zeta^a_{h,t}, r_t, \mu^R_{e,t}, \mu^R_{h,t}, \iota_t$ can be derived from the goods market clearing and the HJB first order conditions. To solve for these four quantities, four equations are required. The first equation is given by subtracting the expected return of each type of the agent. That is, we have

$$\chi_t(\bar{\zeta}_{e,t} - \bar{\zeta}_{h,t}) = \mu^R_{e,t} - \mu^R_{h,t}$$

Plugging in the expression for the return processes from (4), and from Proposition 2, we get

$$\frac{a_t - a_h}{q_t} = \underline{\chi}\Bigg(\underline{\chi}(\psi_t - z_t)\big((\sigma^{q,k}_t + \sigma)^2 + (\sigma^{q,a}_t)^2 + 2\varphi(\sigma + \sigma^{q,k}_t)\big)\left(\frac{\partial J_{h,t}}{\partial z_t}\frac{1}{J_{h,t}} - \frac{\partial J_{e,t}}{\partial z_t}\frac{1}{J_{e,t}} + \frac{1}{z_t(1-z_t)}\right) \tag{41}$$

$$+ \left(\frac{\partial J_{h,t}}{\partial a_{h,t}}\frac{1}{J_{h,t}} - \frac{\partial J_{e,t}}{\partial a_{e,t}}\frac{1}{J_{e,t}}\right)\sigma_{ae,t}\sigma^{q,a}_t + \sigma\big((\sigma^{q,k}_t + \sigma) + \varphi\sigma^a_t\big)(\gamma_e - \gamma_h)\big) \tag{42}$$

The second condition comes from the goods market clearing

$$(z_t\rho_e + (1-z_t)\rho_h)q_t = \psi_t(a_{e,t} - \iota_t) + (1-\psi_t)(a_h - \iota_t) \tag{43}$$

The third and fourth conditions are the return variance components

$$\sigma_t^{q,k} + \sigma = \frac{\sigma}{1 - \frac{1}{q_t}\frac{\partial q_t}{\partial z_t}(\underline{\chi}\psi_t - z_t)} \tag{44}$$

$$\sigma_t^{q,a} = \frac{\frac{1}{q_t}\frac{\partial q_t}{\partial a_{e,t}}\sigma_{ae,t}}{1 - \frac{1}{q_t}\frac{\partial q_t}{\partial z_t}(\underline{\chi}\psi_t - z_t)} \tag{45}$$

which are partial differential equations solved using a Newton-Raphson scheme. The algorithm is given below.

Consider tensor grids of size N_z and N_a with step size Δ_i, and Δ_j where $\{i\}_1^{N_z}, \{j\}_1^{N_a}$ denote the dimensions for the wealth share and the expert productivity respectively. There are two regions in the state space: one where the capital share held by experts $\psi_t < 1$, and one where $\psi_t = 1$. In the first region, the households also hold capital and hence equation (12) holds with equality. In this case, the equations (41), (43), and 44 are used to solve for ψ_t, q_t, $(\sigma + \sigma_t^{q,k})$, and $\sigma_t^{q,a}$. In the second region, the households do not hold capital and hence the equation (12) holds with an inequality. In this case, set $\psi_t = 1$, and use (44), and (43) to solve for q_t, $(\sigma + \sigma_t^{q,k})$, and $\sigma_t^{q,a}$.

- For the first iteration on the wealth share $\{i = 1, \forall j\}$, set $\psi_t = 0$, and take the limiting case of the goods market clearing condition to get q_t. That is

$$\inf_{z \to 0^+} q_t = \frac{a_h\kappa + 1}{\rho_h\kappa + 1} \tag{46}$$

- For iterations $i > 1, \forall j$, use the discretized versions of the equations (44)

$$(\sigma^{q,k} + \sigma)_{i,j} = \sigma\left(1 - \frac{1}{q_{i,j}}\left(\frac{q_{i,j} - q_{i-1,j}}{\Delta_i} z_i\left(\frac{\psi_{i,j}}{z_i} - 1\right)\right)\right)^{-1} \tag{47}$$

$$(\sigma^{q,a})_{i,j} = \left(\frac{q_{i,j} - q_{i,j-1}}{\Delta_j}\sigma_{ae,j}\right)\left(1 - \frac{1}{q_{i,j}}\left(\frac{q_{i,j} - q_{i-1,j}}{\Delta_i} z_i\left(\frac{\psi_{i,j}}{z_i} - 1\right)\right)\right)^{-1} \tag{48}$$

  along with the equations (41), and (43) to solve for $q_{i,j}, \psi_{i,j}, (\sigma+\sigma^q)_{i,j}, (\sigma^{q,a})_{i,j}$.[40] This set of non-linear equations is solved using the Newton-Raphson method. Repeat this procedure until $\psi_t = 1$, in which case the system enters the second region. Then, use (43), (47), and (48) to solve for $q_{i,j}$, $(\sigma + \sigma^{q,k})_{i,j}$ and $(\sigma^{q,a})_{i,j}$.

[40] For $j = 1$, set $\frac{\partial q_t}{\partial a_{e,t}} = 0$ since $a_{e,t} \in [\underline{a}_e, \bar{a}_e]$. That is, the lower and the upper boundaries $\underline{a}_e$ and $\bar{a}_e$ respectively act as reflecting barriers forcing the derivative of the price to be zero.

Time step: Applying Ito's lemma to $J_{j,t}(z_t, a_{e,t})$, matching the drift terms, and augmenting the resulting coupled PDEs with a time step (falst-transient method), we get

$$\mu^J_{j,t} J_{j,t} = \frac{\partial J_{j,t}}{\partial t} + \frac{\partial J_{j,t}}{\partial z_t}\mu^z_t + \frac{\partial J_{j,t}}{\partial a_{e,t}}\mu^a_t + \frac{1}{2}\frac{\partial^2 J_{j,t}}{\partial z_t^2}\left((\sigma^{z,k}_{j,t})^2 + (\sigma^{z,a}_{j,t})^2 + 2\varphi\sigma^{z,k}_{j,t}\sigma^{z,a}_{j,t}\right) + \frac{1}{2}\frac{\partial^2 J_{j,t}}{\partial a_{e,t}^2}\sigma^2_{ae,t} + \frac{\partial^2 J_{j,t}}{\partial z_t \partial a_{a,e}}\left(z_t\sigma^{z,k}_{j,t}\sigma_{ae,t}\varphi + \sigma_a\sigma^{z,a}_{j,t}\right) \tag{49}$$

The coefficients μ^z_t and σ^z_t can be computed from the equilibrium quantities in the static step and $\mu^J_{j,t}$ is computed from the equations in (40). The PDEs are solved using the neural network method explained in Section 2.1.1. Using the updated function $J_{j,t}$, the static step is performed again. The procedure is repeated until the function $J_{j,t}$ convergence until a pre-specified tolerance level is reached.

A.1.6 Three-dimensional plots

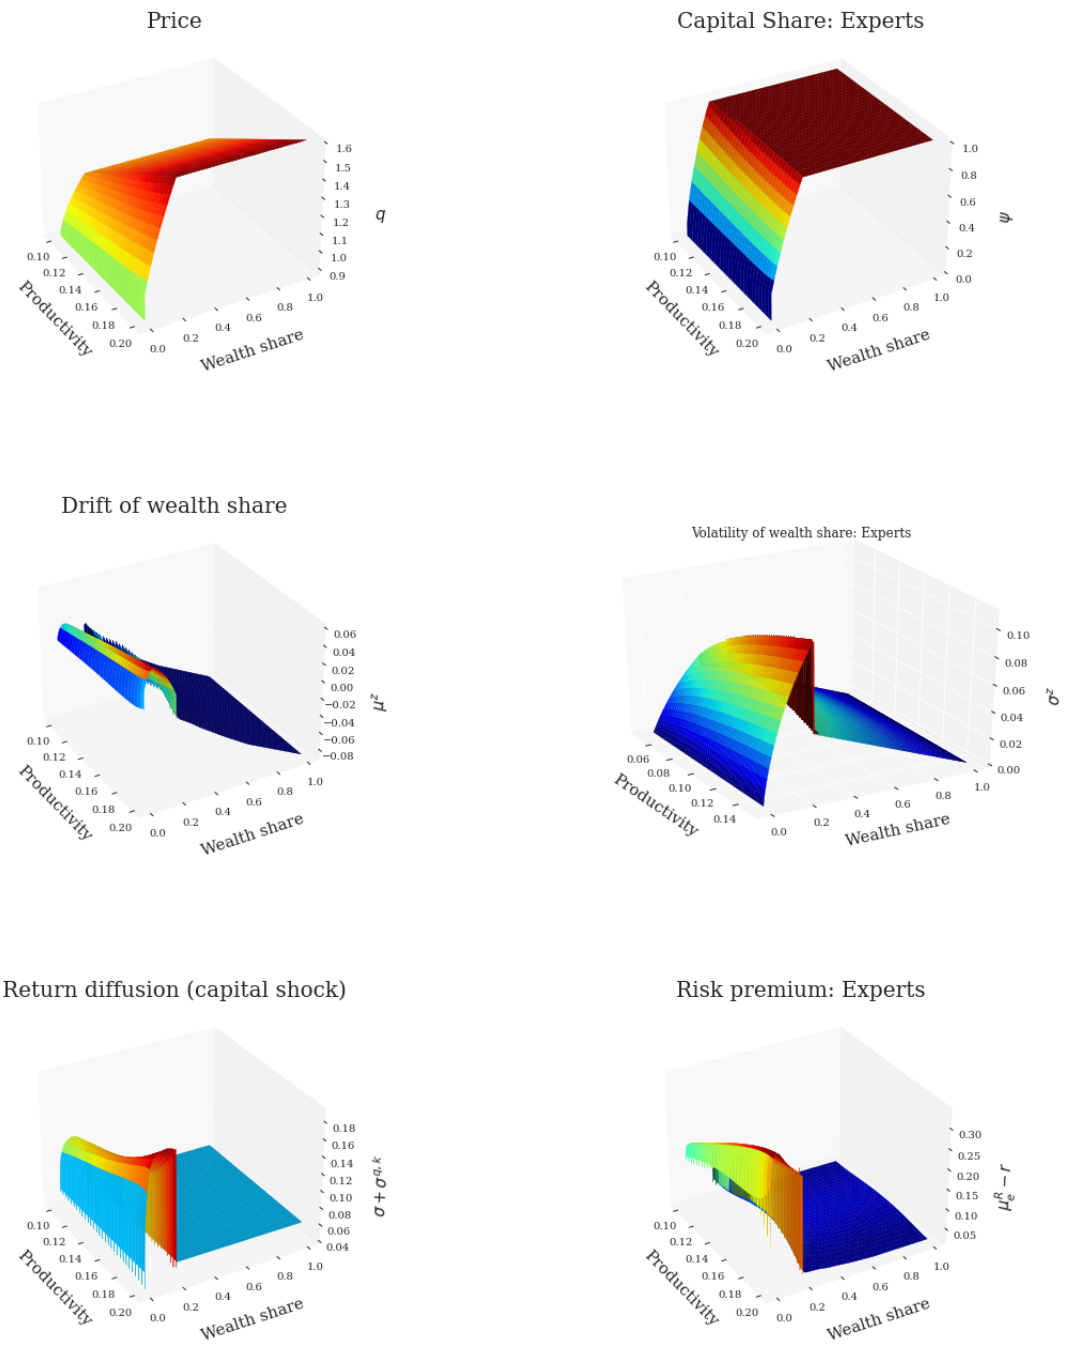

Figure 12: Equilibrium values as functions of state variables z_t and a_t for the stochastic productivity model.

A.2 Benchmark model

A.2.1 Asset pricing conditions

The expected return that the experts earn from investing in the capital is given by

$$dr_t^v = (\mu_{e,t}^R - (1-\chi_t)\epsilon_{h,t})dt + \chi_t(\sigma_t^{q,k} + \sigma)dZ_t^k$$

where $\epsilon_{h,t} = \zeta_{h,t}(\sigma_t^q + \sigma)$. That is, $(1-\chi_t)\epsilon_{h,t}$ is the part of the expected excess return that is paid by the experts to the outside equity holders, which is netted out. Consider a trading strategy of investing \$1 into the capital at time 0. Denoting v_t as the value of this investment strategy at time t, we have $\dfrac{dv_t}{v_t} = dr_t^v$, and

$$\frac{d(\xi_e v_t)}{\xi_e v_t} = (-r_t + \mu_{e,t}^R - (1-\chi_t)\epsilon_{h,t} - \chi_t\epsilon_{e,t})dt + \text{diffusion terms}$$

where $\epsilon_{e,t} = \zeta_{e,t}(\sigma + \sigma_t^q)$, and $\xi_{e,t}$ follows the process in (32). Since $\xi_e v_t$ is a martingale, the drift equals to zero, which implies $\mu_{e,t}^R - r_t = \chi_t\epsilon_{e,t} + (1-\chi_t)\epsilon_{h,t}$ ■

A.2.2 Proof of Proposition 4:

The law of motion of wealth for the households and the experts are given by equation (35). Using the law of large numbers to aggregate the wealth of individual household and expert, we get

$$\frac{dw_{h,t}}{w_{h,t}} = \left(r_t - \frac{c_{h,t}}{w_{h,t}} - \lambda_d + \frac{\chi_t\psi_t}{z_t}(\mu_{h,t}^R - r_t) + \frac{(1-\bar{z})\lambda_d}{1-z_t}\right)dt + \frac{\chi_t\psi_t}{z_t}(\sigma + \sigma_t^q)dZ_t$$
$$\frac{dw_{e,t}}{w_{e,t}} = \left(r_t - \frac{c_{e,t}}{w_{e,t}} - \lambda_d + \frac{\chi_t\psi_t}{z_t}\zeta_{e,t}(\sigma + \sigma_t^q) + \frac{\bar{z}\lambda_d}{z_t}\right)dt + \frac{\chi_t\psi_t}{z_t}(\sigma + \sigma_t^q)dZ_t$$

where $w_{h,t} = \displaystyle\int_{j\in H} w_{j,t}$ and $w_{e,t} = \displaystyle\int_{j\in E} w_{j,t}$ denotes the aggregated wealth among respective group[41], and $z_t = \dfrac{w_{e,t}}{w_{h,t} + w_{e,t}}$. By Ito's lemma, the dynamics of the wealth share becomes

$$\frac{dz_t}{z_t} = \frac{dw_{e,t}}{w_{e,t}} - \frac{d(q_t k_t)}{q_t k_t} + \frac{d\langle q_t k_t, q_t k_t\rangle}{(q_t k_t)^2} - \frac{d\langle q_t k_t, w_{e,t}\rangle}{(q_t k_t w_{e,t})}$$

where

$$\frac{d(q_t k_t)}{q_t k_t} = ((\chi_t\zeta_{e,t} + (1-\chi_t)\zeta_{h,t})(\sigma + \sigma_t^q) - \frac{(a_e - \iota_t)}{q_t} + r_t)dt + (\sigma + \sigma_t^q)dZ_t$$

and the result follows from here after some algebra. ■

While the main text presents and analyzes the benchmark model with recursive utility and IES=1, I present and solve the model for a broader range of preference specifications.

[41] There is a slight abuse of notation here. The quantities $w_{h,t}$ and $w_{e,t}$ represent individual households and experts wealth, as well as the aggregated households and experts wealth.

I consider four different types of utility functions. Let

$$f(c_{j,s},U_{j,s})=\begin{cases}\rho_j\log(c_{j,t})-\rho_jU_{j,t} & \text{if}\quad \gamma_j=1,\varrho_j=1\\ \dfrac{c_{j,t}^{1-\gamma_j}}{1-\gamma_j}-\rho_jU_{j,t} & \text{if}\quad \gamma_j=\varrho_j^{-1}\neq 1\\ (1-\gamma_j)\rho_jU_{j,t}\left(\log(c_{j,t})-\dfrac{1}{1-\gamma_j}\log\Big((1-\gamma_j)U_{j,t}\Big)\right) & \text{if}\quad \gamma_j\neq 1,\varrho_j=1\\ \dfrac{1-\gamma_j}{1-\frac{1}{\varrho_j}}U_{j,t}\left[\left(\dfrac{c_{j,t}}{\big((1-\gamma_j)U_{j,t}\big)^{1/(1-\gamma_j)}}\right)^{1-\frac{1}{\varrho_j}}-\rho_j\right] & \text{if}\quad \gamma_j\neq 1,\varrho_j\neq 1\end{cases} \tag{50}$$

Proposition 5. *The optimal consumption policy and price of risk are given by*

$$\hat{c}_{e,t}=\begin{cases}\rho_e & \textit{if}\quad (\log\quad \textit{or}\quad \textit{Recursive (IES=1)})\\ J_{e,t}^{-1/\gamma_e}(z_tq_t)^{\frac{1-\gamma_e}{\gamma_e}} & \textit{if}\quad \textit{CRRA}\\ \dfrac{J_{e,t}^{\frac{1-\varrho_j}{1-\gamma_e}}}{(z_tq_t)^{1-\varrho_j}} & \textit{if}\quad \textit{Recursive (IES}\neq 1\textit{)}\end{cases} \tag{51}$$

$$\hat{c}_{h,t}=\begin{cases}\rho_h & \textit{if}\quad (\textit{log}\quad \textit{or}\quad \textit{Recursive (IES=1)})\\ J_{h,t}^{-1/\gamma_h}((1-z_t)q_t)^{\frac{1-\gamma_h}{\gamma_h}} & \textit{if}\quad \textit{CRRA}\\ \dfrac{J_{h,t}^{\frac{1-\varrho_j}{1-\gamma_h}}}{((1-z_t)q_t)^{1-\varrho_j}} & \textit{if}\quad \textit{Recursive (IES}\neq 1\textit{)}\end{cases} \tag{52}$$

$$\zeta_{e,t}=\begin{cases}\dfrac{\chi_t\psi_t}{z_t}(\sigma+\sigma_t^q) & \textit{if}\quad \textit{log}\\ -\sigma_{e,t}^J+\sigma_t^z+\sigma_t^q+\gamma_e\sigma & \textit{if}\quad (\textit{CRRA}\quad \textit{or}\quad \textit{Recursive})\end{cases} \tag{53}$$

$$\zeta_{h,t}=\begin{cases}\dfrac{(1-\chi_t\psi_t)}{1-z_t}(\sigma+\sigma_t^q) & \textit{if}\quad \textit{log}\\ -\sigma_{h,t}^J-\dfrac{z_t}{1-z_t}\sigma_t^z+\sigma_t^q+\gamma_h\sigma & \textit{if}\quad (\textit{CRRA}\quad \textit{or}\quad \textit{Recursive})\end{cases} \tag{54}$$

A.2.3 Proof of Proposition 5 :

The HJB equation is given by

$$\sup_{c,\theta}f(c_{j,t},U_{j,t})+E[dU_{j,t}]=0 \tag{55}$$

(a) Log utility: The value function conjecture takes a logarithmic form

$$U_{j,t}=\log k_{j,t}+J_{j,t}(z_t)=\log w_{j,t}+\tilde{J}_{j,t}$$

and where the second equality follows from $z_t=\dfrac{w_{e,t}}{q_tk_t}=1-\dfrac{w_{h,t}}{q_tk_t}$. Also, $f(c_{j,t},U_{j,t})=\rho_j\log(c_{j,t})-\rho_jU_{j,t}$. The value function derivatives are

$$\frac{\partial U_{j,t}}{\partial w_{j,t}}=\frac{dw_{j,t}}{w_{j,t}};\quad \frac{\partial^2U_{j,t}}{\partial w_{j,t}^2}=-\frac{d\langle w_{j,t},w_{j,t}\rangle}{w_{j,t}^2};\quad \frac{\partial U_{j,t}}{\partial \tilde{J}_{h,t}}=1;\quad \frac{\partial^2U_{j,t}}{\partial \tilde{J}_{j,t}^2}=\frac{\partial^2\tilde{J}_{j,t}}{\partial\tilde{J}_{j,t}\partial w_{j,t}}=0$$

Applying Ito's lemma and using the HJB, we get

$$\sup_{c,\theta} \quad \rho_j \log c_{j,t} - \rho(\log w_{j,t} + \tilde{J}_{j,t}) + r_t - \frac{c_{j,t}}{w_{j,t}} + \theta_{j,t}(\sigma + \sigma_t^q)\zeta_{j,t} - \frac{1}{2}\theta_{j,t}^2(\sigma + \sigma_t^q)^2 + \mu_t^{\tilde{J}} = 0$$

Taking the first order conditions and recognizing the fact that $\theta_{e,t} = \frac{\chi_t \psi_t}{z_t}$ and $\theta_{h,t} = \frac{(1-\chi_t\psi_t)}{1-z_t}$, we get the following result for log utility.

$$\hat{c}_{j,t} = \rho_j \tag{56}$$

$$\zeta_{e,t} = \frac{\chi_t \psi_t}{z_t}(\sigma + \sigma_t^q) \tag{57}$$

$$\zeta_{h,t} = \frac{1-\chi_t \psi_t}{1-z_t}(\sigma + \sigma_t^q) \tag{58}$$

(b) CRRA Utility: The value function conjecture is

$$U_{j,t} = J_{j,t}(z_t)\frac{k_{j,t}^{1-\gamma_j}}{1-\gamma_j}$$

where $J_{j,t}$ follows the stochastic differential equation $\frac{dJ_{j,t}}{J_{j,t}} = \mu_{j,t}^J dt + \sigma_{j,t}^J dZ_t$ whose drift and volatility needs to be determined in the equilibrium. The HJB equation is derived directly in terms of the capital k_t instead of the wealth share z_t. The value function derivatives are

$$\begin{aligned} &\frac{\partial U_{j,t}}{\partial J_{j,t}} = \frac{k_{j,t}^{1-\gamma_j}}{1-\gamma_j}; \quad \frac{\partial U_{j,t}}{\partial k_{j,t}} = J_{j,t}k_{j,t}^{-\gamma_j} \\ &\frac{\partial^2 U_{j,t}}{\partial J_{j,t}^2} = 0; \quad \frac{\partial^2 U_{j,t}}{\partial k_{j,t}^2} = -\gamma_j J_{j,t} k_{j,t}^{-(1+\gamma_j)}; \quad \frac{\partial^2 U_{j,t}}{\partial J_{j,t}\partial k_{j,t}} = k_{j,t}^{-\gamma_j} \end{aligned} \tag{59}$$

Applying Ito's lemma and using HJB, we get

$$\begin{aligned} \sup_{c,\theta} \quad &-\rho\frac{J_{j,t}k_{j,t}^{1-\gamma_j}}{1-\gamma_j} + \frac{c_{j,t}^{1-\gamma_j}}{1-\gamma_j} + \frac{J_{j,t}k_{j,t}^{1-\gamma_j}}{1-\gamma_j}\mu_{j,t}^J + J_{j,t}k_{j,t}^{1-\gamma_j}(\Phi(\iota_t) - \delta) \\ &- \sigma^2\frac{\gamma_j}{2}J_{j,t}k_{j,t}^{1-\gamma_j} + J_{j,t}k_{j,t}^{1-\gamma_j}\sigma\sigma_{j,t}^J = 0 \end{aligned} \tag{60}$$

At the optimum, the marginal utilities of consumption and wealth become equal. Rewriting the value function in terms of the wealth and using the mapping $q_t k_t = \frac{w_{e,t}}{z_t} = \frac{w_{h,t}}{1-z_t}$, we get the equilibrium consumption-wealth ratio

$$\frac{c_{e,t}}{w_{e,t}} = \frac{(z_t q_t)^{\frac{1-\gamma_e}{\gamma_e}}}{J_{e,t}^{\frac{1}{\gamma_e}}}; \qquad \frac{c_{h,t}}{w_{h,t}} = \frac{((1-z_t)q_t)^{\frac{1-\gamma_h}{\gamma_h}}}{J_{h,t}^{\frac{1}{\gamma_h}}} \tag{61}$$

The risk premium of the experts and the households can be derived from the stochastic discount factor which is given by

$$\xi_{j,t} = \xi_{j,0} e^{-\rho_j t} \left(\frac{c_{j,t}}{c_{j,0}} \right)^{-\gamma_j}$$

This gives a relationship between the volatility of SDF and consumption: $\sigma^{\xi}_{j,t} = -\gamma_j \sigma^{c}_{j,t}$. The consumption-capital ratio for the households and the experts is given by $\frac{c_{h,t}}{k_t} = \frac{((1-z_t)q_t)^{1/\gamma_h}}{J_{h,t}^{1/\gamma_h}}$ and $\frac{c_{e,t}}{k_t} = \frac{(z_t q_t)^{1/\gamma_e}}{J_{e,t}^{1/\gamma_e}}$. Combining this with the differential equation for SDF

$$\frac{d\xi_{j,t}}{\xi_{j,t}} = -r_t dt - \zeta_{j,t} dZ_t$$

we get

$$\zeta_{e,t} = \gamma_e \sigma^c_{e,t} = -\sigma^J_{e,t} + \sigma^z_t + \sigma^q_t + \gamma_e \sigma; \quad \zeta_{h,t} = \gamma_h \sigma^c_{h,t} = -\sigma^J_{h,t} - \frac{z_t}{1-z_t}\sigma^z_t + \sigma^q_t + \gamma_h \sigma \quad (62)$$

Plugging in the optimal consumption-wealth ratio from (61) into HJB equation (60), we get the expressions for $\mu^J_{j,t}$

$$\mu^J_{e,t} = \rho_e - \frac{(z_t q_t)^{\frac{1-\gamma_e}{\gamma_e}}}{J_{e,t}^{1/\gamma_e}} - (1-\gamma_e)\big(\Phi(\iota_t) - \delta - \gamma_e \sigma^2\big) + \sigma^J_{e,t}\sigma) \quad (63)$$

$$\mu^J_{h,t} = \rho_e - \frac{((1-z_t)q_t)^{\frac{1-\gamma_h}{\gamma_h}}}{J_{h,t}^{1/\gamma_h}} - (1-\gamma_h)\big(\Phi(\iota_t) - \delta - \gamma_h \sigma^2\big) + \sigma^J_{h,t}\sigma) \quad (64)$$

(c) Recursive Utility (IES=1): The value function conjecture is the same as that of CRRA utility, and $f(c_{j,t}U_{j,t}) = (1-\gamma_j)\rho_j U_{j,t}\left(\log c_{j,t} - \frac{1}{1-\gamma_j}\log\big((1-\gamma_j)U_{j,t}\big) \right)$. Plugging in the conjecture for value function in HJB equation (55) and applying Ito's lemma[42], we get

$$\begin{aligned} \sup_{c,\theta} \quad & \rho J_{j,t} k_t^{1-\gamma_j} [\log \frac{c_{j,t}}{w_{j,t}} - \frac{1}{1-\gamma_j}\log J_{j,t} + \log(q_t z_t)] + J_{j,t}\frac{k_t^{1-\gamma_j}}{1-\gamma_j}\mu^J_{j,t} \\ & + J_{j,t} k_t^{1-\gamma_j}(\Phi(\iota_t - \delta)) - J_{j,t} k_t^{1-\gamma_j}\frac{1}{2}\gamma_j \sigma^2 + J_{j,t} k_t^{1-\gamma_j}\sigma\sigma^J_{j,t} = 0 \end{aligned} \quad (65)$$

As before, at the optimum, the marginal utilities of the wealth and the consumption become equal. Using the value function expression in terms of wealth, we have

$$\frac{\partial U_{j,t}}{\partial w_{j,t}} = \frac{\partial f}{\partial c_{j,t}}$$

$$\tilde{J}_{j,t} w_{j,t}^{-\gamma_j} = (1-\gamma_j)\rho_j \frac{U_{j,t}}{c_{j,t}} \implies \frac{c_{j,t}}{w_{j,t}} = \rho_j$$

[42] The value function derivatives are the same as in the CRRA case given by (59).

The stochastic discount factor for recursive utility is given by

$$\xi_{j,t} = \exp\left(\int_0^t \frac{\partial f(c_{j,s}, U_{j,s})}{\partial U} ds \right) \frac{\partial U_{j,t}}{\partial w_{j,t}}$$

Writing the value function conjecture in terms of the wealth instead of the capital, we have

$$U_{j,t} = \tilde{J}_{j,t} \frac{w_{j,t}^{1-\gamma_j}}{1-\gamma_j}; \quad f(c_{j,t}, U_{j,t}) = (1-\gamma_j)\rho_j U_{j,t}\big(\log\rho_j - \frac{1}{1-\gamma_j}\tilde{J}_{j,t}\big)$$

where $\tilde{J}_{j,t} = \dfrac{J_{j,t}}{(q_t z_t)^{1-\gamma_j}}$. The SDF then becomes

$$\xi_{j,t} = (1-\gamma_j)\exp\left(\int_0^t \big[\rho_j\big((1-\gamma_j)\log c_{j,s} - \log\big((1-\gamma_j)U_{j,s}\big) - 1\big] ds \right) \frac{U_{j,t}}{w_{j,t}}$$

This implies that $\sigma(\xi_{j,t}) = \sigma\left(\dfrac{U_{j,t}}{w_{j,t}}\right)$. Computing the R.H.S and using

$$\frac{d\xi_{j,t}}{\xi_{j,t}} = -r_t dt - \zeta_{j,t} dZ_t$$

we get the desired result. Plugging in the consumption-wealth ratio and the market price of risk into the HJB equation (65), we obtain the expressions for $\mu_{j,t}^J$

$$\begin{aligned}
\mu_{e,t}^J &= (\gamma_e - 1)\big(\rho_e \log\rho_e + \log(q_t z_t)\big) + \rho_e \log J_{e,t} - (1-\gamma_e)(\Phi(\iota_t) - \delta - \frac{\gamma_e}{2}\sigma^2 + \sigma\sigma_{e,t}^J) \\
\mu_{h,t}^J &= (\gamma_h - 1)\big(\rho_h \log\rho_h + \log(q_t(1-z_t))\big) + \rho_h \log J_{h,t} - (1-\gamma_h)(\Phi(\iota_t) - \delta - \frac{\gamma_h}{2}\sigma^2 + \sigma\sigma_{h,t}^J)
\end{aligned} \tag{66}$$

(d) Recursive Utility (IES different from unity): The optimization problem is

$$\sup_{c_{j,t}.\theta_{j,t},\iota_t} f(c_{j,t}, U_{j,t}) + E[dU_{j,t}] = 0$$

where

$$f(c_{j,t}, U_{j,t}) = \frac{1-\gamma_j}{1-\frac{1}{\varrho_j}} U_{j,t}\left[\left(\frac{c_{j,t}}{\big((1-\gamma_j)U_{j,t}\big)^{1/(1-\gamma_j)}} \right)^{1-\frac{1}{\varrho_j}} - \rho_j \right]$$

where ϱ_j denotes the IES parameter. The conjecture for the value function is

$$U_{j,t} = J_{j,t}(z_t)\frac{k_{j,t}^{1-\gamma_j}}{1-\gamma_j}$$

where $J_{j,t}$ follows the stochastic differential equation $\dfrac{dJ_{j,t}}{J_{j,t}} = \mu_{j,t}^J dt + \sigma_{j,t}^J dZ_t$ whose drift and volatility needs to be determined in the equilibrium.[43]

[43] Since the value function conjecture is the same as in CRRA case, the value function derivatives are given by (59).

The HJB equation is derived directly in terms of the capital k_t instead of the wealth share z_t. Applying Ito's lemma and using the HJB, we get

$$\sup_{c,\theta} \quad \frac{1}{1-\frac{1}{\varrho_j}}\left(\frac{c_{j,t}^{1-\frac{1}{\varrho_j}}}{J_{j,t}^{\frac{1-\frac{1}{\varrho_j}}{1-\gamma_j}} k_{j,t}^{1-\frac{1}{\varrho_j}}} - \rho_j\right) J_{j,t}k_{j,t}^{1-\gamma_j} + \frac{J_{j,t}k_{j,t}^{1-\gamma_j}}{1-\gamma}\mu_{j,t}^J + J_{j,t}k_{j,t}^{1-\gamma_j}(\Phi(\iota_t)-\delta) \tag{67}$$

$$- \sigma^2\frac{\gamma_j}{2}J_{j,t}k_{j,t}^{1-\gamma_j} + J_{j,t}k_{j,t}^{1-\gamma_j}\sigma\sigma_{j,t}^J = 0$$

At the optimum, the marginal utilities of the consumption and the wealth become equal. Rewriting the value function in terms of the wealth and using the mapping $q_t k_t = \frac{w_{e,t}}{z_t} = \frac{w_{h,t}}{1-z_t}$, we have

$$\begin{aligned}
\frac{\partial f_{e,t}}{\partial c_{e,t}} &= c_{e,t}^{-\frac{1}{\varrho_e}} J_{e,t}^{\frac{\frac{1}{\varrho_e}-\gamma_e}{1-\gamma_e}} (z_t q_t)^{\gamma_j - \frac{1}{\varrho_e}} \\
\frac{\partial f_{h,t}}{\partial c_{h,t}} &= c_{h,t}^{-\frac{1}{\varrho_h}} J_{h,t}^{\frac{\frac{1}{\varrho_h}-\gamma_h}{1-\gamma_h}} ((1-z_t) q_t)^{\gamma_j - \frac{1}{\varrho_h}} \\
\frac{\partial U_{e,t}}{\partial w_{e,t}} &= \frac{J_{e,t}}{(z_t q_t)^{1-\gamma_e}} w_{e,t}^{1-\gamma_e} \\
\frac{\partial U_{h,t}}{\partial w_{h,t}} &= \frac{J_{h,t}}{((1-z_t) q_t)^{1-\gamma_h}} w_{h,t}^{1-\gamma_h}
\end{aligned}$$

Equating the marginal values, we get the respective optimal consumption-wealth ratios

$$\frac{c_{e,t}}{w_{e,t}} = \frac{J_{e,t}^{\frac{1-\varrho_e}{1-\gamma_e}}}{(z_t q_t)^{1-\varrho_e}}; \qquad \frac{c_{h,t}}{w_{h,t}} = \frac{J_{h,t}^{\frac{1-\varrho_h}{1-\gamma_h}}}{((1-z_t) q_t)^{1-\varrho_h}} \tag{68}$$

The stochastic discount factor for recursive utility is given by

$$\xi_{j,t} = \exp\left(\int_0^t \frac{\partial f(c_{j,s}, U_{j,s})ds}{\partial U}\right)\frac{\partial U_{j,t}}{\partial w_{j,t}}$$

Writing the value function conjecture in terms of the wealth instead of the capital, we have

$$U_{j,t} = \tilde{J}_{j,t}\frac{w_{j,t}^{1-\gamma_j}}{1-\gamma_j}; \quad f(c_{j,t}, U_{j,t}) = \frac{\tilde{J}_{j,t} w_{j,t}^{1-\gamma_j}}{1-\frac{1}{\varrho_j}}\left[\left(\frac{c_{j,t}}{w_{j,t}}\right)^{1-\frac{1}{\varrho_j}} \tilde{J}_{j,t}^{\frac{1-\frac{1}{\varrho_j}}{\gamma_j - 1}} - \rho_j\right]$$

where $\tilde{J}_{j,t} = \frac{J_{j,t}}{(q_t z_t)^{1-\gamma_j}}$. Plugging in the above expression in the stochastic discount factor, we notice that $\sigma(\xi_{j,t}) = \sigma\big(\frac{U_{j,t}}{w_{j,t}}\big)$. Computing the R.H.S and using

$$\frac{d\xi_{j,t}}{\xi_{j,t}} = -r_f dt - \zeta_{j,t} dZ_t$$

we get the following result.

$$\zeta_{e,t} = -\sigma^J_{e,t} + \sigma^z_t + \sigma^q_t + \gamma_e \sigma \tag{69}$$

$$\zeta_{h,t} = -\sigma^J_{h,t} - \frac{z_t}{1-z_t}\sigma^z_t + \sigma^q_t + \gamma_h \sigma \tag{70}$$

Substituting the consumption-wealth ratio into the HJB equation (67), we the expression for $\mu^J_{j,t}$

$$\mu^J_{e,t} = \frac{(\gamma_e - 1)}{1 - \frac{1}{\varrho_e}}\left((q_t z_t)^{\varrho_e - 1} J_{e,t}^{\frac{1-\varrho_e}{1-\gamma_e}} - \rho_e\right) - (1-\gamma_e)(\Phi(\iota_t) - \delta - \frac{\gamma_e}{2}\sigma^2 + \sigma\sigma^J_{e,t}) \tag{71}$$

$$\mu^J_{h,t} = \frac{(\gamma_h - 1)}{1 - \frac{1}{\varrho_h}}\left((q_t(1-z_t))^{\varrho_h - 1} J_{h,t}^{\frac{1-\varrho_h}{1-\gamma_h}} - \rho_h\right) - (1-\gamma_h)(\Phi(\iota_t) - \delta - \frac{\gamma_h}{2}\sigma^2 + \sigma\sigma^J_{h,t}) \tag{72}$$

This proves the proposition. ■

A.3 Numerical solution method

A.3.1 Model solution (Log utility):

I rely on the solution technique from BS2016 and Hansen et al (2018) that solves the partial differential equations using an up-winding finite difference scheme. The method involves a static inner loop that solves for the equilibrium quantities $\{\psi_t, (\sigma_{q,t} + \sigma), q_t\}$, and an outer loop that updates the value function from $J_{j,t}$ to $J_{j,t-\Delta t}$ using a finite difference method, similar to the model with stochastic productivity.

Static step: To solve for the quantities in inner loop, three equations are required. The first equation is given by subtracting the portfolio choices of the households and the experts. That is, we have

$$(\theta_{e,t} - \theta_{h,t})(\sigma^q_t + \sigma)^2 = \mu^R_{e,t} - (\mu^R_{h,t})$$

Plugging in the expressions for $\mu^R_{e,t}, \mu^R_{h,t}$ from the return process (31), and using $\theta_{e,t} = \frac{\chi_t \psi_t}{z_t}$ as well as from the capital market clearing condition $\theta_{h,t} = \frac{1 - \chi_t \psi_t}{1 - z_t}$, we get

$$\frac{\chi_t \psi_t - z_t}{z_t(1-z_t)}(\sigma^q_t + \sigma)^2 = \frac{a_e - a_h}{q_t} \tag{73}$$

The second equation comes from the goods market clearing condition

$$(z_t \hat{c}_{e,t} + (1-z_t)\hat{c}_{h,t})q_t = \psi_t(a_e - \iota_t) + (1-\psi_t)(a_h - \iota_t) \tag{74}$$

where $\iota_t = \frac{q_t - 1}{\kappa}$. For the third equation, apply Ito's lemma to $q(z_t)$ and match the drift and the volatility terms to get $\sigma^q_t = \frac{\partial q_t}{\partial z_t}\frac{1}{q}\sigma^z_t$. Combining this with the volatility of wealth

share, we get

$$\sigma_t^q + \sigma = \frac{\sigma}{1 - \frac{\partial q_t}{\partial z_t}\frac{1}{q}(\frac{\chi_t \psi_t}{z_t} - 1)} \tag{75}$$

Equations (73), (74), and (75) are solved using the Newton-Raphson method[44] yielding $\{\psi_t, (\sigma_{q,t}+\sigma), q_t\}$. In the case of log utility, the static step is enough since the consumption-wealth share is equal to the discount rate and the risk premium is not dependent on $J_{j,t}$.

A.3.2 Model solution (CRRA and Recursive utility):

The portfolio choice in the case of CRRA and recursive utility includes the hedging demand that needs to be taken into account. From equations (33) and (34), we get

$$\frac{a_e - a_h}{q_t(\sigma + \sigma_t^q)} \geq \underline{\chi}(\zeta_{e,t} - \zeta_{h,t})$$

with equality if $\psi_t = 1$. Plugging in the expressions for $\zeta_{e,t}$ and $\zeta_{h,t}$ from proposition (5), we have

$$\begin{aligned}
\frac{a_e - a_h}{q_t} &= \underline{\chi}\left(\frac{1}{J_{h,t}}\frac{\partial J_{h,t}}{\partial z_t} - \frac{1}{J_{e,t}}\frac{\partial J_{e,t}}{\partial z_t} + \frac{1}{z_t(1-z_t)}\right)(\underline{\chi}\psi_t - z_t)(\sigma + \sigma_t^q)^2 \\
\frac{a_e - a_h}{q_t} &= \underline{\chi}\left(\sigma_{h,t}^J - \sigma_{e,t}^J + \frac{\sigma_t^z}{1-z_t}\right)(\sigma + \sigma_t^q)
\end{aligned}$$

where the second expression comes from using the dynamics of the wealth share (36).[45] The goods market clearing condition (74) and return volatility (75) remain the same. Similar to the case of log utility, the Newton-Raphson method is used to solve for the $\{\psi_t, q_t, (\sigma+\sigma_t^q)\}$. Given these equilibrium functions, $J_{j,t}$ needs to be solved for, which is done in the dynamic time step.

Time step: Applying Ito's lemma to $J_{j,t}(z_t)$, matching the drift terms, and augmenting the resulting coupled PDEs with a time step (falst-transient method), we get

$$\mu_{h,t}^J J_{h,t} = \frac{\partial J_{h,t}}{\partial z_t}\mu_t^z + \frac{1}{2}\frac{\partial^2 J_{h,t}}{\partial z_t^2}(\sigma_t^z)^2 \tag{76}$$

$$\mu_{e,t}^J J_{e,t} = \frac{\partial J_{e,t}}{\partial z_t}\mu_t^z + \frac{1}{2}\frac{\partial^2 J_{e,t}}{\partial z_t^2}(\sigma_t^z)^2 \tag{77}$$

The coefficients μ_t^z and σ_t^z can be computed from the equilibrium quantities in the static step and $\mu_{j,t}^J$ is computed from the equations in (66). The PDEs are solved using an implicit method with an up-winding scheme explained in the next part.

[44] BS2016 and Hansen et al. [2018] provide details of the algorithm. The state space is segmented into the crisis region and the normal region. The static step is solved for iteratively until the system enters the crisis region in which case the capital share ψ is set to 1 and the remaining quantities (q_t, σ_t^q) are solved for using equations (74) and (75).

[45] Note that by Ito's lemma, we have $\sigma_{j,t} = \frac{1}{J_{j,t}}\frac{\partial J_{j,t}}{\partial z_t}\sigma_t^z = \frac{1}{J_{j,t}}\frac{\partial J_{j,t}}{\partial z_t}(\theta_{e,t} - 1)(\sigma + \sigma_t^q)^2$

A.3.3 Up-winding scheme

The PDEs (76) are solved by considering artificial time-derivatives. To be specific, the modified system

$$0 = \frac{\partial J_{h,t}}{\partial t} - \mu^J_{h,t} J_{h,t} + \frac{\partial J_{h,t}}{\partial z_t}\mu^z_t + \frac{1}{2}\frac{\partial^2 J_{h,t}}{\partial z_t^2}(\sigma^z_t)^2 \tag{78}$$

$$0 = \frac{\partial J_{e,t}}{\partial t} - \mu^J_{e,t} J_{e,t} + \frac{\partial J_{e,t}}{\partial z_t}\mu^z_t + \frac{1}{2}\frac{\partial^2 J_{e,t}}{\partial z_t^2}(\sigma^z_t)^2 \tag{79}$$

is solved backwards in time with the corresponding terminal conditions $(J_{h,T}, J_{e,T})$. Consider a general quasi-linear PDE of the form

$$A\left(z, g, \frac{\partial g}{\partial z}\right) + tr\left[B\left(z, g, \frac{\partial g}{\partial z}\right)\frac{\partial^2 g}{\partial z^2}B\left(z, g, \frac{\partial g}{\partial z}\right)'\right] + \frac{\partial g}{\partial t} = 0$$

Consider a two-dimensional grid of size N_z and N_t with step sizes Δ_i and Δ_j respectively where $\{i\}_1^{N_z}$,$\{j\}_1^{N_t}$ denote the dimensions for space and time respectively. The function $g(z_t, t)$ evaluated at (i, j) is denoted as $g_{i,j}$. The derivatives of the function are discretized as

$$\frac{\hat{\partial} g_{i,j}}{\hat{\partial} z} = (\mu^z_j)^+ \frac{g_{i+1,j} - g_{i,j}}{\Delta_i} + (\mu^z_j)^- \frac{g_{i,j} - g_{i-1,j}}{\Delta_i}$$

$$\frac{\hat{\partial}^2 g_{i,j}}{\hat{\partial} z^2} = \frac{g_{i+1,j} - 2g_{i,j} + g_{i-1,j}}{\Delta_i^2}$$

$$\frac{\hat{\partial} g_{i,j}}{\hat{\partial} t} = \frac{g_{i,j+1} - g_{i,j}}{\Delta_j}$$

where $(\mu^z_j)^+ = \begin{cases} \mu^z_t & \text{if} \quad \mu^z_t > 0 \\ 0 & \text{if} \quad \text{otherwise} \end{cases}$ $(\mu^z_j)^- = \begin{cases} \mu^z_t & \text{if} \quad \mu^z_t < 0 \\ 0 & \text{if} \quad \text{otherwise} \end{cases}$

Discretizing the derivatives at $j+1$ and applying it to the PDE, we get

$$g_{i,j+1} = g_{i,j} + \Delta_j \left\{ A\left(z, g_{i,j+1}, \frac{\hat{\partial} g_{i,j+1}}{\hat{\partial} z}\right) + tr\left[B\left(z, g_{i,j+1}, \frac{\hat{\partial} g_{i,j+1}}{\hat{\partial} z}\right)\frac{\hat{\partial}^2 g_{i,j+1}}{\hat{\partial} z^2}B\left(z, g_{i,j+1}, \frac{\hat{\partial} g_{i,j+1}}{\hat{\partial} z}\right)'\right]\right\}$$

Solving for $g_{i,j+1}$ requires solving a linear system of equations which can be done using a standard procedure such as the Richardson method. The up-winding scheme ensures monotonicity of the numerical scheme (see d'Avernas and Vandeweyer (2019)). Since the method is implicit, a large time step can be set which considerably reduces the computation time.

A.3.4 Numerical simulation:

The state variable in the model is z_t whose law of motion is governed by the equation (36). Once the mapping between z_t and (μ^z_t, σ^z_t) are determined numerically from the previous section, we can simulate z_t using an Euler-Maruyama scheme. Specifically, the task is to simulate

$$dz_t = \mu^z_t dt + \sigma^z_t dZ_t$$

where the shock dZ_t is the standard Brownian motion. The law of motion is discretized as

$$z_{t+\Delta t} = z_t + \mu_z(z_t)\Delta t + \sigma_t^z(z_t) * \sqrt{\Delta t} Z$$

where $Z \sim N(0,1)$. The steps are as follows

1. Set z_0 to an arbitrary initial value, say 0.5.
2. Simulate Z from the standard normal distribution and compute $z_{t+\Delta t}$ using the discretized equation for $\Delta = 1/12$. The mapping between z_t and (μ_t^z, σ_t^z) is in a grid since it is solved for numerically and hence I use a spline interpolation to obtain the intermediate values.
3. Repeat the procedure for $z_1, z_2, ...$ and obtain $\{z_t\}_1^{60,000}$. That is, the simulation is done for 5000 years at monthly frequency.

The first 1000 years are eliminated so as to reduce the dependency on the initial condition. I experiment with different initial values to make sure that the obtained distribution is indeed stationary. The procedure is repeated for 1000 times and Figure (6) plots the resulting distributions.

Comparison with Fokker-Planck equation: The density of the wealth share $g(z_t, t)$ can be expressed in the form of Fokker-Planck (or Kolmogorov Forward Equation) equation

$$\frac{\partial g(z_t,t)}{\partial t} = -\frac{\partial}{\partial z_t}(\mu_t^z g(z_t,t)) + \frac{1}{2}\frac{\partial^2}{\partial z_t^2}((\sigma_t^z)^2 g(z_t,t))$$

We have $\lim_{z_t \to 0^+, z_t \to 1^-} \sigma_t^z = 0$ by construction and ($\lim_{z_t \to 0^+} \mu_t^z > 0$, $\lim_{z_t \to 1^-} \mu_t^z < 0$) due to the overlapping generations assumption. This forces the distribution to be non-degenerate. Also, a stationary density implies that $\frac{\partial g}{\partial t} = 0$. Thus, we can integrate the Fokker-Planck equation to obtain

$$0 = \text{constant} - (\mu_t^z g(z_t)) + \frac{1}{2}\frac{\partial}{\partial z_t}((\sigma_t^z)^2) g(z_t)$$

I solve this equation numerically using an explicit finite difference scheme and compare it with the stationary distribution obtained through the simulation. Figure (13) shows that the density obtained from the simulation is a good approximation for the theoretical density dictated by the Focker-Planck equation. The simulated wealth share is annualized so as to make the comparison with the empirical data. The proportion of annualized wealth share that fall below the theoretically obtained crisis boundary z^* is taken to be the probability of crisis implied by the model. Table (9) presents the moments of equilibrium quantites obtained using the annualized wealth share

A.4 Other trade-offs in the benhmark model

One key quantity that governs the time spent in the crisis region is the drift of the wealth share. The parameter λ_d controls the death rate of experts which is necessary to ensure model stationarity. As the death rate increases, all else equal, the system stays in the crisis region longer. A similar effect is observed when the mean proportion of experts $\overline{z}$ is

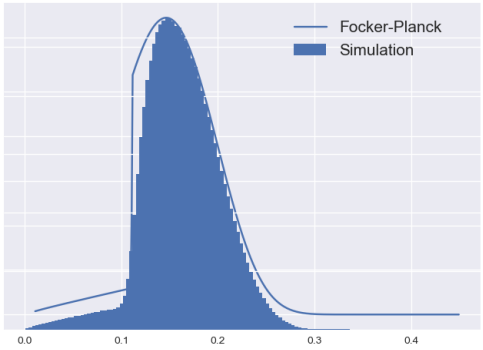

Figure 13: Comparison of the stationary density obtained from Focker-Planck equation and the simulation for the benchmark model with RA=1.

decreased. Figure (14) presents the static comparison of the drift of the wealth share for different values of λ_d and $\overline{z}$. A higher death rate pushes the system into the crisis region by making the drift of wealth share more negative in the normal regime. However, there is only a minimal effect on the drift once the system enters the crisis region. The second panel varies the mean population share of experts by keeping the death rate fixed. As the population share decreases, the drift becomes more negative making the crisis more likely. Once the system enters the crisis region, the drift becomes less positive pushing the system back into the normal regime at a slower rate. Both of these effects work towards increasing the frequency of crisis. Figure (15) shows the probability of crisis for several values of λ_d

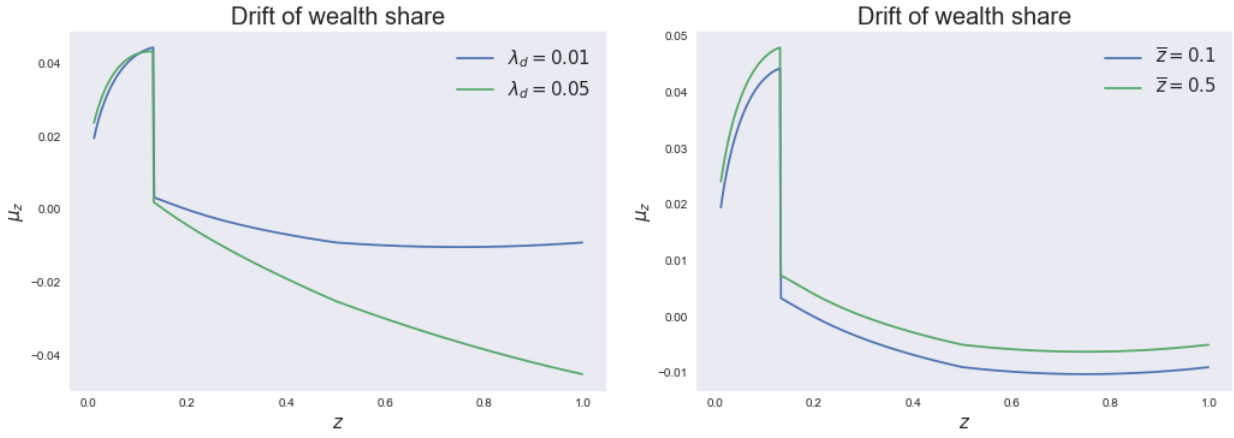

Figure 14: Left panel shows the drift of wealth share for two different values of death rate λ_d for $\overline{z}$ fixed at 0.1. The second panel shows the drift of wealth share for two different values of mean expert population for λ_d fixed at 0.02. The risk aversion is set to 2 for both the plots.

and $\overline{z}$ for the recursive utility model with IES=1 and risk aversion equal to 2. To obtain a 7% probability of crisis, the population share of experts have to be less than 10%, with a death rate of 3%. Since the discount rate assumed in the model is inclusive of death rate, a 3% death rate means that the households do not discount at all. The second panel of Figure

(15) reveals that changing the OLG parameters doesn't affect unconditional risk premium much. While it is possible to achieve a realistic probability of crisis and unconditional risk premium simultaneously, this comes at the cost of extremely high death risk, and more importantly, it still does not generate persistent recessions. This is because the duration of the crisis is unaffected by a high death risk and thus leads to a quick recovery.

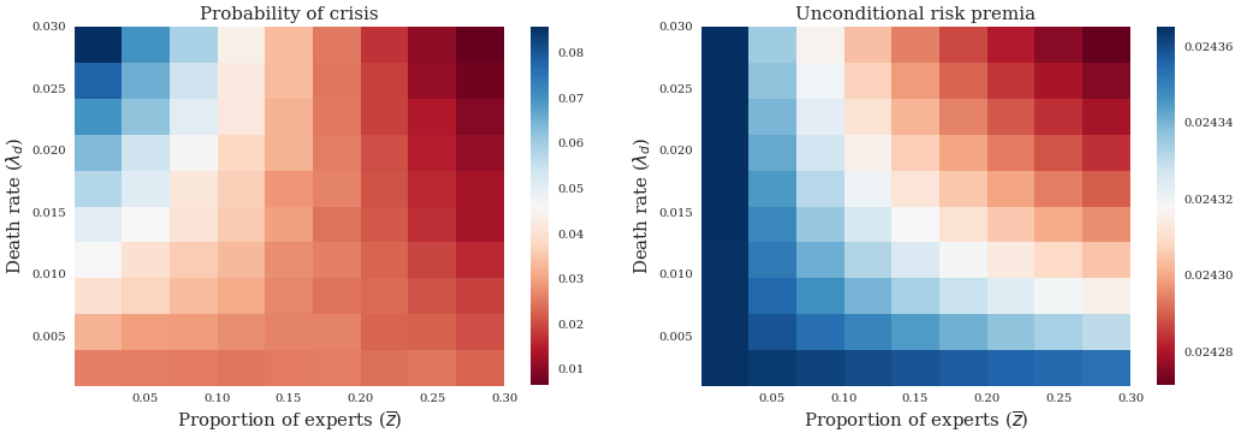

Figure 15: Left panel shows the drift of wealth share for two different values of exit rate λ_d for $\overline{z}$ fixed at 0.1. The right panel shows the drift of wealth share for two different values of mean expert population for λ_d fixed at 0.02. Both plots are from recursive utility model with risk aversion equal 2 and IES=1.

Tightening financial constraint: One of the key assumptions of the model is the inability of experts to fully issue outside equity. The parameter $\underline{\chi}$ governs how much equity the experts are forced to retain and hence it is of interest to study the model by varying this parameter. As the financial constraint tightens, the probability of crisis increases. The left panel of Figure (16) plots the risk premium of experts for three different values of the skin-in-the-game constraint. As the constraint increases, the crisis boundary shifts to the right but the unconditional risk premium is lower. This effect can be seen in the simulation result on the right panel of Figure (16). While a higher value of $\underline{\chi}$ leads to a higher probability of the crisis, the conditional risk premium drops drastically leading to only a marginal increase in the unconditional premium.

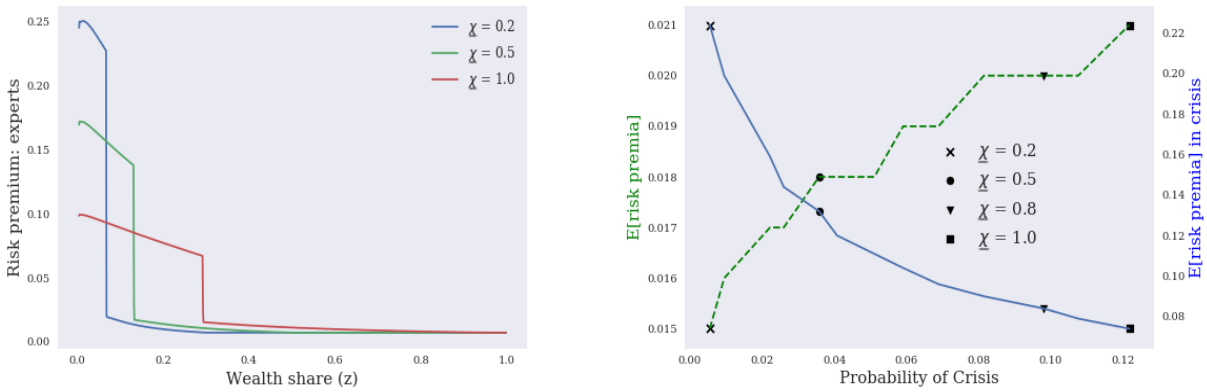

Figure 16: Left panel: Static comparison of the risk premium by changing the skin-in-the game constraint for the baseline model with RA=1 and IES=1. Right panel: Trade off between the conditional risk premium and the probability of crisis by varying the skin-in-the-game constraint. The parameter $\underline{\chi}$ increases from left to right. Left (dashed line) and right (blue line) axes correspond to the unconditional and the conditional risk premium respectively.

A.5 Deep learning methodology

A.5.1 One-dimensional model

I first present the solution to the benchmark model using deep learning method and then demonstrate how and why it is easy to scale to higher dimensions by presenting the solution to richer model with two state variables. I consider the case of recursive utility with IES=1 and RA=2 for demonstration.[46] The PDE that needs to be solved is given in (78). Construct a neural network $\hat{J}(z,t \mid \Theta)$ and define the PDE residual to be

$$f := \frac{\partial \hat{J}}{\partial t} + \frac{\partial \hat{J}}{\partial z}\mu^z + \frac{1}{2}\frac{\partial^2 \hat{J}}{\partial z^2}(\sigma^z)^2 - \mu^J \hat{J}$$

The network architecture is given in Figure (18) with the hyperparameters in Table (1).[47] Figure (17) plots the full grid and the training sample. The inner static loop uses a grid size of 1000 points in space dimension while the neural network only uses 300 points for training. In the case of a single space dimensional model, sampling one-third of the grid points is enough to find the right solution. In higher dimensions, the proportion of grid points required as training sample can be set much lower than one-third.

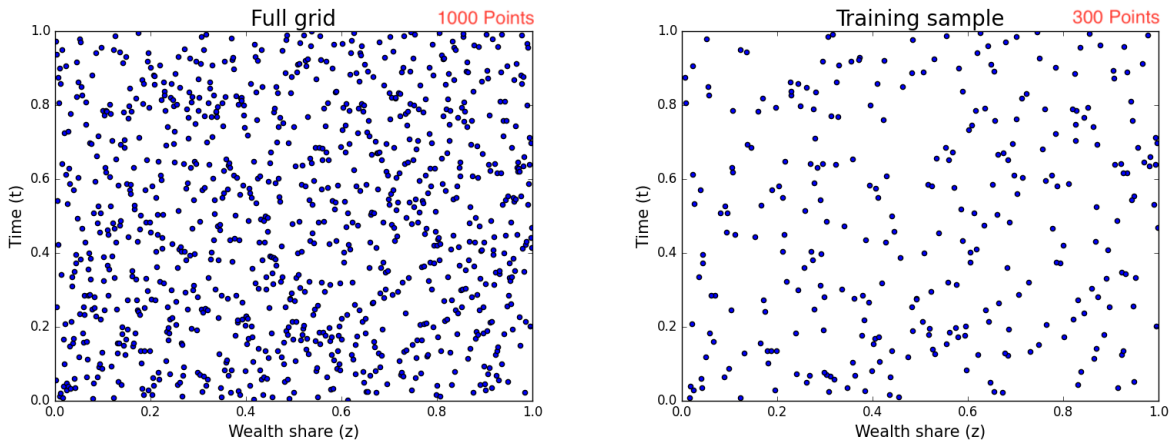

Figure 17: Grid used in numerical procedure: 1D model.

I illustrate the simplicity of coding the neural network solution using code snippets that uses Tensorflow library. The first step is to construct a neural network $\hat{J}$ using the space and time dimensions as training data, and weights and biases as parameters initialized arbitrarily.[48] This is illustrated in the code snippet (1) and it corresponds to the left most feed-forward neural network ($NN : \hat{J}(z,t \mid \Theta)$) in Figure (18). The next step is to construct the regularizers using PDE residual as given in code snippet (2). This forms the PDE network in Figure (18). The PDE coefficients (advection, diffusion, and linear terms) are taken as given and form part of the training sample. The automatic differentiation

[46] The deep learning algorithm works for any type of utility function. For larger risk aversion values, it takes longer to achieve convergence due to the highly non-linear value function near the boundaries.

[47] The algorithm works even for 2 hidden layers with 30 neurons each instead of 4 hidden layers but may be prone to instabilities for some extreme parameter values such as setting $\underline{\chi} = 0.1$. It is recommended to have four layers to capture the non-linearity well.

[48] I use Xavier initialization to avoid the vanishing gradient problem.

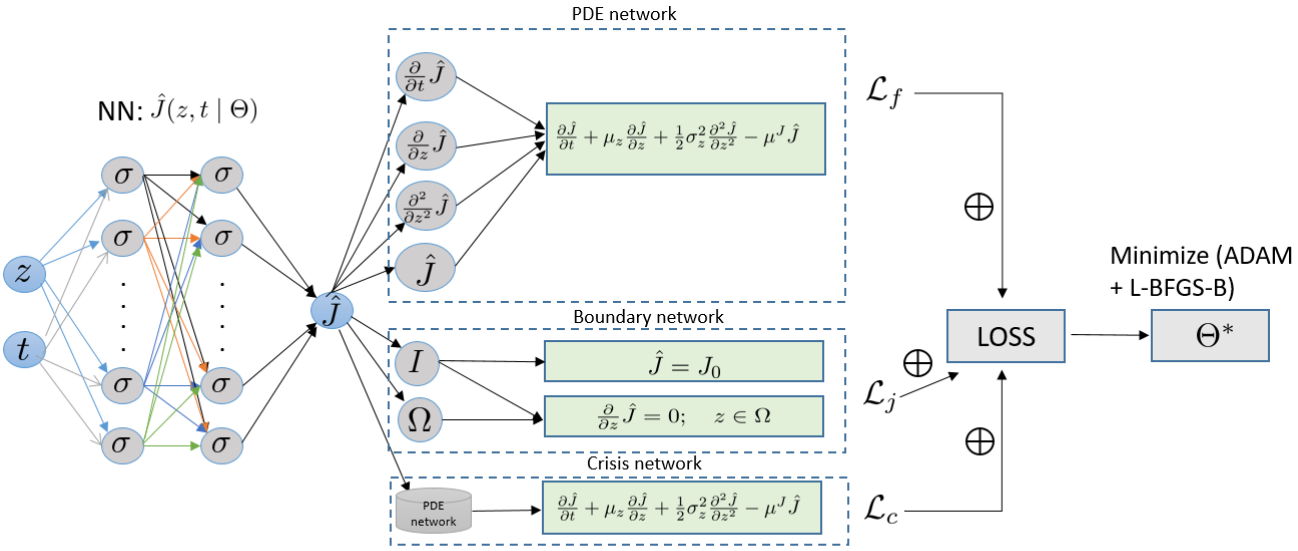

Figure 18: Network architecture: benchmark model.

in Tensorflow ($tf.gradients$) enables fast computation of derivatives in the regularizers which guides the parameterized neural network $\hat{J}$ towards the right solution even when the training sample is small. In addition to the PDE bounding loss, one can easily set up the boundary loss and crisis region loss in a similar fashion.

```
def J(z,t):
  J = neural_net(tf.concat([z,t],1),weights,biases)
  return J

```

Listing 1: Approximating J using a neural network: 1D model

```
def f(z,t):
  J = J(z,t)
  J_t = tf.gradients(J,t)[0]
  J_z = tf.gradients(J,z)[0]
  J_zz = tf.gradients(J_z,z)[0]
  f = J_t + advection * J_z + diffusion * J_zz - linearTerm
  return f
```

Listing 2: Constructing regularizer: 1D model

Since the analytical solution to the benchmark model is not available, I compare the neural network solution with the those obtained from the finite difference method explained in Appendix (A.3.3). Figure (19) shows the comparison. They are not only qualitatively similar, they are quantitatively the same up to the order of 1e-4.

A.5.2 Two-dimensional model

The PDE that needs to be solved in the two-dimensional model is given in (23). As in the case of one-dimensional model, construct the neural network $\hat{J}(z, a, t \mid \Theta)$ with the PDE

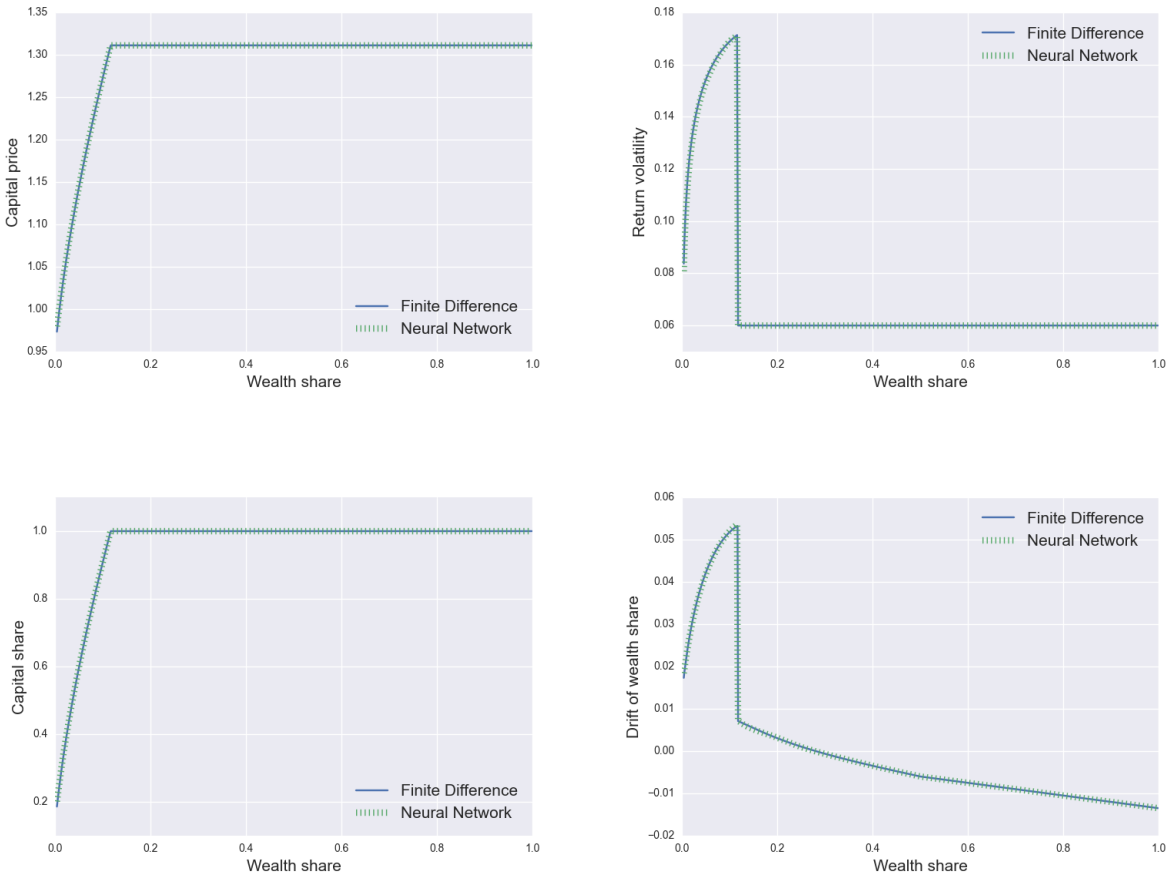

Figure 19: Comparison of equilibrium quantities using finite difference and neural network in one-dimensional benchmark model.

residual taking the form

$$f := \frac{\partial \hat{J}}{\partial t} + \frac{\partial \hat{J}}{\partial z}\mu^z + \frac{\partial \hat{J}}{\partial z}\mu^a + \frac{1}{2}\frac{\partial^2 \hat{J}}{\partial z^2}\left((\sigma^{z,k})^2 + (\sigma^{z,a})^2 + 2\varphi\sigma^{z,k}\sigma^{z,a}\right) + \frac{1}{2}\frac{\partial^2 \hat{J}}{\partial a_a^2}\sigma_a^2 + \frac{\partial^2 \hat{J}}{\partial z_t \partial a}\left(z\sigma^{z,k}\sigma_a\varphi + \sigma_a\sigma^{z,a}\right) - \mu^J \hat{J}$$

The network architecture and hyperparameters are given in Figure (3) and (1) respectively. The grid size becomes larger compared to the one-dimensional model but the chosen training sample size is 3000 which is much smaller than the full grid size of 30,000 as is illustrated in Figure (20).

To appreciate the simplicity involved in scaling to higher dimensions, I present the code snippets for the 2D model in (3) and (4). Similar to the 1D model, the neural network J is parameterized the same way except that the network takes three inputs-two space dimensions (z, a) and one time dimension (t). This corresponds to the leftmost feed-forward neural network in Figure (3) where three neurons enter the network instead of two as in Figure (18). The construction of regularizer as shown in code snippet (4) simply adds new derivative terms to the PDE network taking as given the coefficients (advection, diffusion, linear, and cross terms). Moving from one to two dimensions in

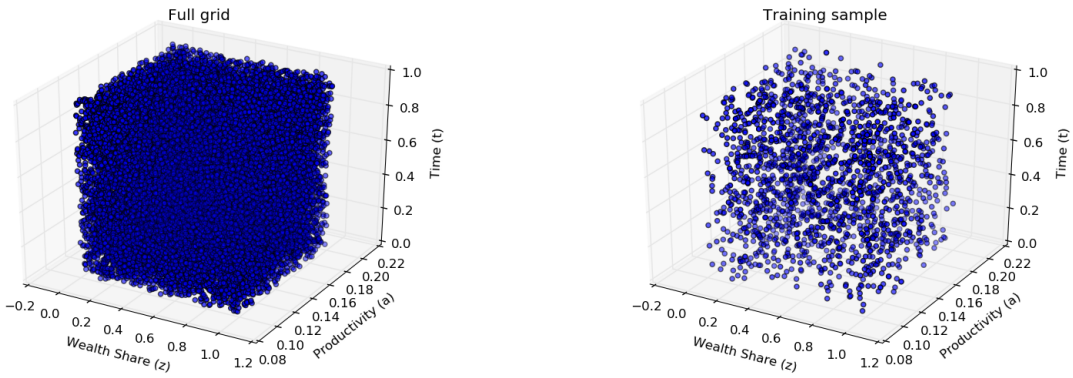

Figure 20: Grid used in numerical procedure: 2D model. The full grid contains 30,000 points and the training sample contains 3000 points.

an implicit finite difference method is not trivial since one has to set up the system of linear equations to be solved numerically. In even higher dimensions, as demonstrated in Gopalakrishna (2020), the PDE network simply adds further derivative terms. This is easier to do in comparison with setting up the system of equations. In dimensions more than two with correlated state variables, preserving monotonicity of the numerical schemes adds further complications, which the neural network method sidesteps. The literature has used advanced C++ tools like Paradiso (see Hansen et al. [2018]) which requires much more effort than simply augmenting the PDE network. Since most of the heavy lifting is done by the automatic differentiation in the regularizers, learning in high dimensions is accomplished effectively through a few lines of coding.

```
def J(z,a,t):
  J = neural_net(tf.concat([z,a,t],1),weights,biases)
  return J

```

Listing 3: Approximating J using a neural network: 2D model

```
def f(z,a,t):
  J = J(z,a,t)
  J_t = tf.gradients(J,t)[0]
  J_z = tf.gradients(J,z)[0]
  J_a = tf.gradients(J,a)[0]
  J_zz = tf.gradients(J_z,z)[0]
  J_aa = tf.gradients(J_a,a)[0]
  J_az = tf.gradients(J_a,z)[0]
  f = J_t + advection_z * J_z + advection_a * J_a + diffusion_z * J_zz +
        diffusion_a * J_aa  + crossTerm * J_az - linearTerm
  return f
```

Listing 4: Constructing regularizer: 2D model

	All	Crisis	Normal	All	Crisis	Normal	All	Crisis	Normal	All	Crisis	Normal
E[leverage]	3.22	5.50	3.09	2.08	4.36	2.08	1.57		1.57	1.29		1.29
E[inv. rate]	6.00%	4.90%	6.00%	5.80%	5.60%	5.80%	5.50%		5.50%	5.30%		5.30%
E[risk free rate]	4.80%	-0.10%	5.00%	3.30%	-1.30%	3.30%	1.60%		1.60%	-1.80%		-1.80%
E[risk premia]	1.70%	13.50%	1.00%	2.70%	16.50%	2.70%	4.50%		4.50%	8.00%		8.00%
E[price]	1.42	1.34	1.42	1.40	1.39	1.40	1.39		1.39	1.37		1.37
E[return volatility]	6.20%	15.90%	5.70%	5.80%	14.30%	5.80%	5.80%		5.80%	5.90%		5.90%
E[price of risk]	0.21	0.86	0.18	0.47	1.15	0.47	0.77		0.77	1.35		1.35
E[GDP growth rate]	2.30%	-7.90%	2.90%	2.10%	-10.70%	2.10%	1.90%		1.90%	1.80%		1.80%
Std[inv. rate]	0.38%	1.12%	0.11%	0.13%	0.23%	0.12%	0.08%		0.08%	0.04%		0.04%
Std[risk premia]	2.84%	0.91%	0.23%	0.43%	0.56%	0.21%	0.17%		0.17%	0.10%		0.10%
Std[risk free rate]	1.19%	0.42%	0.28%	0.32%	0.15%	0.28%	0.22%		0.22%	0.14%		0.14%
Std[GDP growth rate]	7.17%	19.36%	5.15%	4.93%	5.04%	4.90%	4.89%		4.89%	4.86%		4.86%
Corr(leverage,shock)	-0.27	-0.04	-0.24	-0.19	0.29	-0.19	-0.21		-0.21	-0.26		-0.26
Corr(price return, riskf ree rate)	0.16	-0.05	-0.23	-0.09	0.48	-0.16	-0.17		-0.17	-0.20		-0.20
Corr(risk premia, volatility)	0.98	-0.60	0.43	0.40	-0.09	-0.92	-1.00		-1.00	-1.00		-1.00
Prof. of crisis	7.80%			0.10%			0.01%			0.00%		
Risk aversion	1			5			10			20		

Table 9: Benchmark model implied moments for different risk aversion levels with parameters from Table (3).

B Online Appendix

B.1 Benchmark model

Solving the incomplete market capital misallocation model with fire-sales and endogenous regimes involves numerical techniques that are non-standard from the asset pricing literature viewpoint. In addition to the complexity involving in solving the PDEs, the coefficients of the PDEs change with respect to the form of utility function. Thus, comparing model solutions across different utility specifications require manual intervention to modify the equations in static step, and the PDE coefficients. Part of the contribution of this paper is to offer a simpler way to perform comparative valuation dynamics through numerical libraries made available[49] at `https://github.com/goutham-epfl/MacroFinance`. The simplicity of using the library is that model can be solved and simulated in a few lines facilitating comparative valuation. Code snippet (5) presents an example of solving the model with different utility specifications. Code snippet (6) shows examples of simulating different models from the general framework.

[49] Advanced users can also choose among implicit and explicit finite difference schemes to solve the model, use different interpolation methods, and modify the frequency of time used in the simulation.

```
from model_recursive_class import model_recursive
from model_class import model
from model_general_class import model_recursive_general
import matplotlib.pyplot as plt

#Input parameters
params={'rhoE': 0.06, 'rhoH': 0.03, 'aE': 0.11, 'aH': 0.03,
            'alpha':0.5, 'kappa':7, 'delta':0.025, 'zbar':0.1,
            'lambda_d':0, 'sigma':0.06, 'gammaE':2, 'gammaH':2, 'IES=1.5'
    }

#solve model1
model1 = model_recursive_general(params)
model1.solve()

#solve model2
#switch to model with unitary IES
params['IES'] =1.0
#solve model
model2 = model_recursive(params)
model2.solve()

#plot capital price (Q) from the model1 and model2
plt.plot(model1.Q), plt.plot(model2.Q)
```

Listing 5: Solving the model using Python library

```
from model_recursive_class import model_recursive
from simulation_model_class import simulation_benchmark


#Input parameters
params={'rhoE': 0.06, 'rhoH': 0.03, 'aE': 0.11, 'aH': 0.03,
            'alpha':0.5, 'kappa':7, 'delta':0.025, 'zbar':0.1,
            'lambda_d':0, 'sigma':0.06, 'gammaE':2, 'gammaH':2, 'IES=1.0'
    }
#set number of simulations
params['nsim'] = 500
params['utility'] = 'recursive'
#simulate model1
simulate_model1 = simulation_benchmark(params)
simulate_model1.compute_statistics()
print(simulate_model1.stats) #print key statistics
simulate_model1.write_files() #store key statistics for later use

#simulate model2
#change volatility
params['sigma'] =0.10
simulate_model2 = simulation_benchmark(params)
simulate_model2.compute_statistics()

#compare stationary distribution from two models
plt.plot(simulate_model1.z_sim.reshape(-1))
plt.hist(simulate_model2.z_sim.reshape(-1))
```

Listing 6: Simulating the model using Python library