EPFL

Thèse n° 9409

Essays in macro-finance and deep learning

Présentée le 28 juillet 2023

Collège du management de la technologie
Chaire du Prof. ordinaire Collin-Dufresne
Programme doctoral en finance

pour l'obtention du grade de Docteur ès Sciences

par

Goutham GOPALAKRISHNA

Acceptée sur proposition du jury

Prof. S. Malamud, président du jury
Prof. P. Collin Dufresne, directeur de thèse
Prof. M. Brunnermeier, rapporteur
Prof. F. Celentano, rapporteur
Prof. J. Hugonnier, rapporteur
E

École
polytechnique
fédérale
de Lausanne

2023

To Divya and my parents.

Acknowledgements

First of all, I would like to thank my advisor Pierre Collin-Dufresne for his continuous support and guidance over the years. My Ph.D. thesis is mainly shaped by his knack for picking out the essentials from all that I put on his table. He is more of a role model than an advisor to me. His sense of fair play and chasing one's passion without necessarily expecting instant reward are some of the many attributes that I admire.

I am forever indebted to Markus Brunnermeier for taking me under his wing and offering generous support at Princeton University. The saying in Tamizh goes '*Mata, pita, guru, deivam (Mother, father, teacher, and the Divine)*'. *Guru* comes before the Divine, because he/she is an instrument and a doorway into enlightenment. It is through this doorway that one sees the world beyond. Markus has been nothing short of a *Guru*. His immense support on the job market has opened up the doorway to what lies ahead for me in finance academia.

I also thank my jury members Julien Hugonnier, Semyon Malamud, and Francesco Celentano for their valuable time. Special thanks to Erwan Morellec for his support in the early days of the Ph.D. program, and the University of Geneva for accepting me at the Swiss Finance Institute doctoral program. I would also like to thank Lorraine Dupart and Sophie Cadena for their help in navigating through the administrative maze.

Luciano and Benoit made my Ph.D. life in Lausanne idyllic. The much-needed distraction at the basketball court with Benoit kept me going, along with the stimulating discussions with Mads about asset pricing. Coffee/dinner chats with Tammaro, Oliver, Martina, and Federico were always fun-filled. Debating with Narasimhan on topics related to financial markets and politics was a constant reminder of differing perspectives outside academia. Don and Philippe made my life at Princeton memorable, and I am thankful for that.

Divya (my wife) has been my biggest support in every inch of the way. It is hard for one hand to thank the other for successfully playing the piano in unison, so I will not try. This thesis is the product of her hard work, as much as it is mine. Last but not the least, I thank my family (my parents and in-laws) for their unconditional love, support and patience. I especially want to thank my mother for instilling a love for books in me since childhood. My father would have been proud of my accomplishment.

Lausanne, April, 2023 G. Goutham

Abstract

This thesis studies the origins and consequences of financial crises, and computational techniques to solve continuous-time economic models that explain such crises.

The first chapter shows that financial recessions are typically characterized by a large risk premium and a slow recovery. However, macro-finance models have trouble matching these empirical features, especially when they are calibrated to match both the observed unconditional and conditional macroeconomic and asset-pricing moments simultaneously. In this chapter, I build a macro-finance model that quantitatively explains the salient features of a financial crisis, such as a large drop in output, a spike in the risk premium, reduced financial intermediation, and a long duration of economic distress. The model has leveraged intermediaries with stochastic productivity and a state-dependent exit rate that governs the transition into and out of a crisis. A model without these two features suffers from a trade-off between the amplification and persistence of a crisis. I show that my model resolves this tension and generates realistic crisis dynamics.

In the second chapter, I develop a new computational framework called Actively Learned and Informed Equilibrium Nets (ALIENs) to solve continuous time economic models with endogenous state variables and highly non-linear policy functions. I employ neural networks that are trained to solve supervised learning problems that respect the laws governed by the economic system in the form of general parabolic partial differential equations. The subdomain of the high dimensional state space that carries the most economic information is learned actively in an iterative loop, enforcing the random training points to be sampled from areas that matter the most to ensure convergence. The method is applied to successfully solve a model of macro-finance that is notoriously difficult to handle using traditional finite difference schemes.

In the third chapter, I investigate the origins of bank failures, an important feature of financial crises. I analyze a panel of bank holding companies and offer empirical evidence for the franchise value to be associated with a higher probability of failure. The results indicate that the changing scope of the banking industry with declining franchise value compared to the pre-crisis period is worrisome, despite strong capital ratios.

Keywords: Macro-Finance; Financial intermediation; machine-learning; banking.

Résumé

Cette thèse étudie les origines et les conséquences des crises financières, ainsi que les techniques informatiques permettant de résoudre les modèles économiques en temps continu qui expliquent ces crises.

Le premier chapitre montre que les récessions financières sont typiquement caractérisées par une prime de risque importante et une reprise lente. Cependant, les modèles de macro-finance ont du mal à correspondre à ces caractéristiques empiriques, en particulier lorsqu'ils sont calibrés pour correspondre simultanément aux moments macroéconomiques et d'évaluation des actifs, inconditionnels et conditionnels, observés. Dans ce chapitre, je construis un modèle de macro-finance qui explique quantitativement les principales caractéristiques d'une crise financière, telles qu'une forte baisse de la production, un pic de la prime de risque, une réduction de l'intermédiation financière et une longue période de dépression économique. Le modèle comporte des intermédiaires à effet de levier avec une productivité stochastique et un taux de sortie dépendant de l'état qui régit la transition vers et hors d'une crise. Un modèle sans ces deux caractéristiques souffre d'un arbitrage entre l'amplification et la persistance d'une crise. Je montre que mon modèle résout cette tension et génère une dynamique de crise réaliste.

Dans le deuxième chapitre, je développe un nouveau cadre informatique appelé Actively Learned and Informed Equilibrium Nets (ALIENs) pour résoudre des modèles économiques en temps continu avec des variables d'état endogènes, et des fonctions modélisant les politiques économiques non linéaires. J'utilise des réseaux neuronaux entrainés pour résoudre des problèmes d'apprentissage supervisé qui respectent les lois régies par le système économique sous la forme d'équations différentielles partielles paraboliques générales. Les informations économiques sont codées sous forme de régularisateurs qui disciplinent le réseau neuronal profond dans le processus d'apprentissage. Le sous-domaine de l'espace d'état à haute dimension qui contient le plus d'informations économiques est appris activement dans une boucle itérative, obligeant les points d'entraînement aléatoires à être échantillonnés dans les domaines les plus importants pour garantir la convergence. La méthode est appliquée pour résoudre avec succès un modèle de macro-finance qui est notoirement difficile à traiter à l'aide de schémas traditionnels de différences finies.

Dans le troisième chapitre, j'étudie les origines des faillites bancaires, une caractéristique

importante des crises financières. J'analyse un panel de holdings bancaires et j'apporte des preuves empiriques que la valeur de franchise est associée à une plus grande probabilité de faillite. Les résultats indiquent que l'évolution du périmètre du secteur bancaire, avec une valeur de franchise en baisse par rapport à la période précédant la crise, est inquiétante, malgré des ratios de fonds propres solides.

Mots clés : Macro-Finance ; Intermédiation Financiaire ; machine-learning ; Secteur Bancaire.

Contents

1 A Macro-Finance model with Realistic Crisis Dynamics

1.1 Introduction

It is well known that recessions are marked by high equity risk premia, low investment rates, and low output. The great recession of 2007-2008 emphasized the importance that financial intermediaries play in propagating shocks to the real economy. Since then, there has been a growing literature on the leverage of intermediaries as a key factor in moving asset prices and the real economy.[1] Recessions that feature a sharp decrease (increase) in the investment rate and output (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 (2014a) (BS2014 henceforth), He and Krishnamurthy (2013), Di Tella (2017), etc.). In these models, the financial constraints bind only at certain times which leads to non-linearity in the asset prices. In normal times, financial markets facilitate capital allocation to the most productive agents. In such states, intermediaries are sufficiently capitalized and the premium on the risky asset is low. In bad times, financial constraints bind and the capital gets misallocated to less productive agents, who do not value capital as much. This leads to a deterioration of intermediary balance sheets and pushes the economy into a crisis where the premium on the risky asset shoots up. These models explain a high risk premium in the crisis periods but the contribution has largely been qualitative except Maxted (2020) and Krishnamurthy and Li (2020).[2]

The contribution of this chapter is two-fold. First, I build an overlapping-generation incomplete-market asset pricing model with stochastic productivity and state-dependent exit of the experts 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 regular-

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

[2] Gertler et al. (2020) incorporates bank run into a standard New Keynesian model that explains financial crisis quantitatively. However, they focus on matching a specific crisis episode- the great recession of 2008.

izers.[3] This methodology, as shown in the second chapter, is scalable and can be applied to similar high-dimensional problems. The fluctuating productivity of experts is a source of dividend risk and is a crucial driver of systemic instability along with capital shocks. In addition, I introduce state-dependent exit of experts as a parsimonious way of capturing bank defaults. The data from Federal Deposit Insurance Corporation (FDIC) shows that a total of 297 banks failed in 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. 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 (1.1) 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 far greater than the other years. While a lot of non-financial institutions failed too during the Great recession, the fact that 80% of Moody's issuer defaults in terms of volume came from financial institutions alone indicates that the intermediaries default to a large extent, particularly during financial crises. I capture this empirical phenomenon through a an expert exit rate that is calibrated to observed bank default rates.

The second contribution is the quantification of my model to dissect the mechanisms of the financial crisis. To this end, I show that a simpler model with constant productivity and no exit of experts, which reduces to Brunnermeier and Sannikov (2016a) (BS2016 henceforth), suffers from a tension between the *amplification* and the *persistence* of financial crises. In particular, there is a trade-off between the conditional risk premium and the duration of crisis.[5] During bad times, the premium on the risky asset shoots up due to capital misallocation and fire-sale. The leveraged experts earn the higher conditional risk premium allowing them to rebuild sufficient wealth and recover quickly from the crisis. Such a fast rebound is at odds with the data since recessions are empirically long-lasting. Auxiliary features of the model that generate longer crises necessarily attenuate the conditional risk premium (i.e., amplification gets dampened). This is because crises tend to be long when the experts recapitalize slowly, which can only happen when the risk premium that the experts earn is low in the model. To give a concrete example, when the simpler model is calibrated to generate a realistic 18-month duration of the crisis, the model implied conditional risk premium is 2%, which is much lower than the empirically observed premium of 25%.[6] On the other hand, when the model is calibrated to generate a realistic conditional risk premium of 25%, the model-implied average

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

[4] Source for bank failures: https://www.fdic.gov/bank/historical/bank/, and Moody's Corporate Default and Recovery Rates, 1920-2008. Financial institutions include Bank holding companies, Real estate, and insurance companies. The list of banks includes only those that are insured by the FDIC. Failure of investment banks such as Lehman Brothers in 2008 is not included.

[5] In this chapter, *conditional* risk premium refers to the premium on the risky asset in the crisis state. Another interesting trade-off that emerges from this simpler model is between the unconditional risk premium and the probability of a crisis. This is explored in detail in Section 1.3.

[6] See Table (1.6) in Section 1.3 for the estimated conditional risk premium. The average contraction period from the NBER website is around 18 months. Source: https://www.nber.org/cycles.html. This is a conservative measure compared to around 3 years peak to trough period reported in Muir (2017).

duration of crisis is 5 months, well short of 18-month crisis duration observed in the data. The model with stochastic productivity and state-dependent exit rate resolves this tension and provides reasonable crisis dynamics along three key dimensions: a) *crisis likelihood*, which represents the occupation time of the economy in a crisis state, b) *amplification*, that represents a large conditional risk premium and low output, and c) *persistence*, that represents slow recovery from the crisis. When the economy is in a stochastic steady state, all capital is held by the experts, and the risk premium is low. A negative shock to the level of capital also decreases the productivity of experts, increasing the frequency of crisis since the experts are more likely to fire-sell assets to the households and trigger the financial amplification channel. The crisis state is characterized by a low output, depressed investment, and a large risk premium. The model implies an 8% probability of a crisis, matching the empirical value of 7% from Reinhart and Rogoff (2009). In a crisis state, the rate at which the experts exit and become households is high, reflecting large empirical bank bankruptcies, reducing the proportion of agents who manage capital more productively. This force has a dominating impact on the experts' wealth compared to the effect coming from increased risk premium and pushes the economy deeper into crisis. The productivity eventually mean reverts, and the economy reaches a point where the increased productivity dominates the exit effect, helping the economy climb out of the crisis. The speed of mean reversion in productivity is low, forcing the economy to spend a long amount of time in distress before the increase in productivity ends the gloomy phase. The model implies a crisis duration of 17 months, close to the empirical value of 18 months from the NBER recessionary cycle data. At normalcy, all capital in the economy is held by experts again, and the financial amplification channel is shut down, where the exit rate is small. Thus, the twin forces of stochastic productivity and exit match the empirical moments in all three categories, bringing the model closer to data.

The model 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 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.[7] This algorithm dominates the finite-difference method used in BS2016, Hansen et al. (2018), etc., since the presence of correlated state variables makes it difficult to maintain the monotonicity of finite-difference schemes which is required for convergence.[8] The second chapter discusses the algorithm in detail and applies it to similar problems with several dimensions as high as five.

The simpler benchmark model with constant productivity and no exit is similar in spirit to

[7] I rely on Tensor-flow, a deep learning library developed by Google Brain, that computes the numerical derivatives efficiently.

[8] See D'Avernas and Vandeweyer (2019) and Phelan and Eslami (2022) for issues of monotonicity in finite-difference schemes.

Figure 1.1 – Bank failures

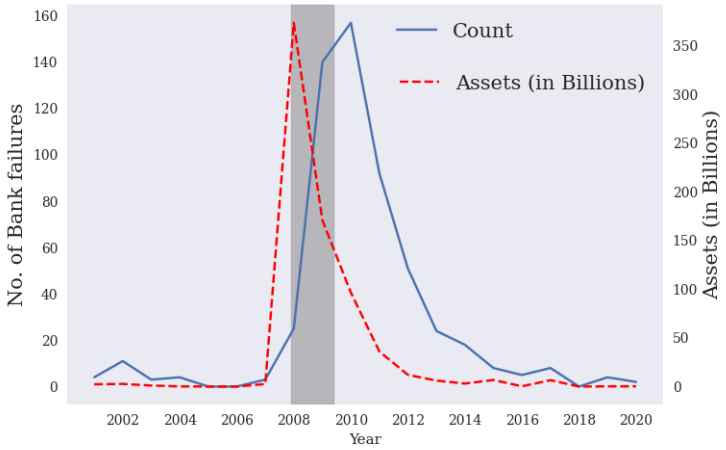

Note: The solid line indicates the number of bank failures and the dashed line indicates the default volume. The shaded region represent the NBER recessionary period. Source: Federal Deposit Insurance Corporation.

BS2016 but there is an overlapping generation of agents (OLG) with recursive preference. The assumption of OLG offers a non-degenerate stationary distribution of the state variable (Gârleanu and Panageas (2015)), while recursive preference helps with obtaining realistic asset pricing moments.[9] I quantify this benchmark model, similar in spirit to He and Krishnamurthy (2019) (HK2019 henceforth) 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 where the experts sell capital to households that have a lower valuation of the capital due to their lower productivity rate. 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 chapter but focuses on matching credit

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

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 subjected to Brownian and Poisson shocks, I consider a recursive utility function and focus on matching a broader set of both unconditional and conditional macroeconomic and asset pricing moments such as the intermediary leverage patterns, the equity risk premium, the investment rate, the GDP growth rate, and the probability and duration of crisis. The 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. Maxted (2020) analyzes a quantitative model of financial intermediation and sentiment, similar to Krishnamurthy and Li (2020) where intermediaries do not consume and have mean-variance preferences over their reputation.

Models of intermediary asset pricing highlight the *persistence* and the *amplification* of shocks caused by the leveraged agents. A measure of persistence and amplification is the duration of the crisis and conditional risk premium, respectively. The quantification of the benchmark model reveals two key trade-offs. First, there is tension between the unconditional risk premium and the probability of a crisis. A high level of risk aversion is required to match the large observed unconditional risk premium. When the experts earn a large premium in the stochastic steady state, small negative shocks to the capital do not cause enough deterioration in their net worth to hit the crisis boundary, thereby diminishing the probability of a crisis. Second, conditional on being in crisis, there is a tension between the risk premium and the duration of the crisis. This is because risk premium spikes as soon as the economy enters a crisis state, enabling the experts to gain wealth quickly and revert to the normal regime leading to fast recovery. With larger values of risk aversion, the experts build wealth even faster through a higher risk premium, resulting in a quicker 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 since larger risk aversion levels mechanically imply a lower probability and duration of crisis.[10] 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 weaknesses are the inability to jointly generate a realistic duration of crisis and risk premium, and sufficient variation in the risk prices.[11] The richer model with stochastic productivity and state-dependent exit rate of the experts generates reasonable asset pricing and crisis moments. Embedding these two features that have empirical support brings the model closer to the data in important aspects.

Related Literature This chapter relates to several strands of the literature. On the modeling front, it is most closely related to BS2016 which introduces 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),

[10] It is common in asset pricing literature to assume a high risk aversion. See, for example Gârleanu and Panageas (2015), who set risk aversion of leveraged agents equal to 10.

[11] Since the q-theory result tightly ties the investment rate to the capital price, a low model implied volatility of price translates to a low variation in the investment rate too.

Adrian and Boyarchenko (2012), Moreira and Savov (2017), etc. While BS2014 assumes 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 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. (2017), and Adrian et al. (2014)). While these papers provide a theory based on 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 a crisis, the goal of this chapter 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 chapter comes closer to Krishnamurthy and Li (2020) which attempts 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. This chapter also relates to Khorrami (2016), who shows that the implied cost of entry to participate in the stock market is as large as 90% of the wealth of the agents. Another interpretation of this result is that the costs of risk concentration are unreasonably large to match the empirically observed level of risk premium. While he focuses on a limited asset market participation model with costly entry, my model features capital misallocation with stochastic productivity that is calibrated to match both the amplification as well as the duration of crisis in the data. Bigio and D'Avernas (2021) build a risk capacity-based model with information asymmetries to explain slow recovery from the financial crisis. The state-dependent exit of experts in this chapter relates to Eisfeldt et al. (2017) who introduce endogenous entry and exit of participants in complex asset markets.[12] Ikeda and Kurozumi (2019) analyze a DSGE model with financial shocks to generate slow recovery from financial crisis. However, they study deviations from a steady-state, whereas, my model allows for studying the global dynamics.

Hansen et al. (2018) provide a framework that nests several models based on financial frictions. Even though the frictions prevent the economy from achieving the first-best outcome, their

[12] In Eisfeldt et al. (2017), the decision to enter and exit is endogenous and hence the agents solve an optimal stopping time problem. In this chapter, the exit rate is assumed to be state-dependent.

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 chapter is guided by quantitative analysis. While they consider a multi-dimensional problem with auxiliary shocks to the volatility and 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 both unconditional and conditional macroeconomic and asset pricing moments. My model assumes that the productivity of experts fluctuates with business cycle, which holds empirical relevance (Hughes et al. (2001), Feng and Serletis (2010)). Dindo et al. (2022) shows that the average intermediation cost for banks in the US is counter-cyclical, and builds a general equilibrium model associating the cost with the business cycle. I consider a parsimonious way to capture bank defaults through an exogenous exit rate of experts which complements a large literature on endogenous bank runs and defaults (Gorton and Ordoñez (2014), Gertler et al. (2020), Li (2020)).

Lastly, this chapter 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, but 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 an 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.[13] Phelan and Eslami (2022) uses a finite-state Markov Chain approximation method to overcome the monotonicity. With the advancements in machine learning, recent papers have turned to neural networks 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.[14] The algorithm proposed in this chapter is similar in spirit but also incorporates prior information from the crisis boundary as regularizers and is particularly geared toward 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 points from the state space in an

[13] See Merkel (2020) for a similar technique used to solve a Macro-Finance model in two dimensions.

[14] There is a substantial literature on the deep-learning techniques to solve PDEs in Applied Mathematics, which is covered in the second chapter. For the application of deep learning techniques to solve discrete-time DSGE models, see Azinovic et al. (2019).

informed manner.

1.2 Model

In this section, I present a heterogeneous agent model with stochastic productivity and state-dependent exit rate of experts. There is an infinite horizon economy with a continuum of agents, who are of two types: households and experts. Let $\mathbb{H}$ and $\mathbb{E}$ denote the set of households and experts, respectively. 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.[15] The experts can issue risk-free debt, and obtain a higher return to holding capital as they are more productive than 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, the experts should 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

$$y_{j,t} = a_{j,t} k_{j,t} \qquad j \in \{e,h\} \tag{1.1}$$

where the capital evolves as[16]

$$\frac{dk_{j,t}}{k_{j,t}} = (\Phi(\iota_{j,t}) - \delta)dt + \sigma dZ_t^k \tag{1.2}$$

with $\iota_{j,t}$ as the investment rate, and $\{Z_t \in \mathbb{R}; \mathscr{F}_t, \Omega\}$ is the standard Brownian motions representing the aggregate uncertainty in $(\Omega, \mathbb{P}, \mathscr{F})$. The parameter σ denotes the exogenous volatility of the 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 households and experts. I assume that the investment cost function takes the logarithmic form[17] $\Phi(\iota) = \frac{\log(\kappa\iota+1)}{\kappa}$ 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

$$da_{e,t} = \pi(\hat{a}_e - a_{e,t})dt + \nu \underbrace{(\overline{a}_e - a_{e,t})(a_{e,t} - \underline{a}_e)}_{\sigma_{ae,t}} dZ_t^a \tag{1.3}$$

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

[15] This is due to the homogeneity of preferences. The mass of households and experts is time-varying due to demographic changes and exit. Preferences and assumptions related to exit are explained later.

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

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

always higher than that of the households even though it fluctuates between $\underline{a}_e$ and $\bar{a}_e$.[18] The capital price q_t follows

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

The return process for each type of agent is given by $dR_{j,t} = \frac{d(q_t k_{j,t})}{q_t k_{j,t}} + \frac{(a_{j,t}-\iota_{j,t})k_{j,t}}{q_t k_{j,t}} dt$ where the first component of the R.H.S is capital gain, and the second component is the dividend yield. Note that the dividends are agent-specific due to different productivity rates, and possibly due to different investment rates.[19] The time-varying productivity $a_{e,t}$ is a source of dividend risk for the experts, and therefore Z_t^a acts as a Financial shock in addition to the capital shock. Applying Ito's lemma, we get

$$dR_{j,t} = \underbrace{\left(\mu_t^q + \Phi(\iota_{j,t}) - \delta + \sigma\sigma_t^{q,k} + \varphi\sigma\sigma_t^{q,a} + \frac{a_{j,t}-\iota_{j,t}}{q_t}\right)}_{\mu_{j,t}^R} dt + (\sigma_t^{q,k} + \sigma) dZ_t^k + \sigma_t^{q,a} dZ_t^a \quad (1.4)$$

The aggregate output in the economy is given by $y_t = A_t K_t$, where $K_t = \int_{\mathbb{E}\cup\mathbb{H}} k_{j,t} dj$, and A_t is the aggregate dividend that satisfies

$$A_t = \int_{\mathbb{E}\cup\mathbb{H}} a_{j,t} \frac{k_{j,t}}{K_t} dj$$

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

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

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

$$\frac{d\xi_{j,t}}{\xi_{j,t}} = -r_t dt - \zeta_{j,t}^k dZ_t^k - \zeta_{j,t}^a dZ_t^a \quad (1.5)$$

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.

[18] 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 (1.3).

[19] It turns out that the optimal investment rate is the same for both types of agent since it depends on the capital price and the adjustment cost parameter κ. For now, I assume that the investment rate is agent-specific and show later in (1.13) that it is the same for all agents.

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

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

with

$$f(c_{j,t}, U_{j,t}) = (1-\gamma)\rho U_{j,t}\left(\log(c_{j,t}) - \frac{1}{1-\gamma}\log\Big((1-\gamma)U_{j,t}\Big)\right) \tag{1.6}$$

where γ and ρ are the risk aversion and the discount rate coefficients 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 experts and households that are born each instant respectively. The death risk is not measurable under the filtration generated by the Brownian process $\mathscr{F}_t$ 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 can be thought of as inclusive of the death rate λ_d. I abstract away from the insurance market 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 experts become households, where τ_t is state-dependent. This transition will be taken into consideration in the optimization problem of the agents.[20] This assumption is a parsimonious way to capture bank failures, which are particularly high during financial crises, as seen in Figure (1.1). I assume that households do not have exit, and hence the inflow into the set $\mathbb{E}$ is only through birth, whereas, the inflow into the set $\mathbb{H}$ is due to both birth and migration of experts. Note that the mass of households and experts is time-varying as a result. The experts optimize by maximizing their utility functions, subject to wealth constraints, starting from some initial wealth $w_{e,0}$.[21] Let τ' denote the time at which the experts exit and become households, that is exponentially distributed with the rate τ_t. They solve

$$\begin{aligned} U_{e,t} = \sup_{c_{e,t}, k_{e,t}, \chi_{e,t}} \quad & E_t\left[\int_t^{\tau'} f(c_{e,s}, U_{e,s})ds + U_{h,\tau'}\right] \\ \text{s.t.} \quad \frac{dw_{e,t}}{w_{e,t}} = & \Big(r_t - \frac{c_{e,t}}{w_{e,t}} + \frac{q_t k_{e,t}}{w_{e,t}}(\mu^R_{e,t} - r_t - (1-\chi_{e,t})\epsilon_{h,t})\Big)dt \\ & + \sigma_{w_e,t}\Big((\sigma + \sigma^{q,k}_t)dZ^k_t + \sigma^{q,a}_t dZ^a_t\Big) \end{aligned} \tag{1.7}$$

where $\frac{q_t k_{e,t}}{w_{e,t}}$ and $\chi_{e,t}$ denote the fraction of wealth invested in capital, and the experts' inside equity share, respectively. The experts obtain a continuation utility of $U_{h,\tau'}$ starting from the time of transition into households. While the experts obtain an expected excess return of $\mu^R_{e,t} - r_t$ by investing in the risky asset, they have to pay the outside equity investors $(1-\chi_{e,t})\epsilon_{h,t}$, where $\epsilon_{h,t}$ is the premium demanded by households defined in equation (1.12). Thus, the

[20] Gomez (2019) uses a similar assumption that applies to the leveraged wealthy households, and in Di Tella (2017), a similar exit rate is applied to the intermediaries to generate a non-degenerate stationary distribution. However, they do not model the exit rate as state-dependent. The functional form of τ_t is provided later in (1.22), after constructing of the state space.

[21] 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.1.1.

latter component is netted out from the total expected return on capital investment. The skin-in-the game constraint implies that experts choose the inside equity share $\chi_{e,t} \in [\underline{\chi}, 1]$. On the other hand, households do not issue outside equity implying that $\chi_{h,t} = 1$ always. I write $\chi_{e,t}$ simply as χ_t for notational convenience henceforth. The households solve

$$U_{h,t} = \sup_{c_{h,t}, k_{h,t}} \quad E_t\left[\int_t^\infty f(c_{h,s}, U_{h,s})ds\right] \tag{1.8}$$
$$\text{s.t.} \quad \frac{dw_{h,t}}{w_{h,t}} = \left(r_t - \frac{c_{h,t}}{w_{h,t}} + \frac{q_t k_{h,t}}{w_{h,t}}(\mu^R_{h,t} - r_t)\right)dt + \sigma_{w_h,t}\left((\sigma + \sigma^{q,k}_t)dZ^k_t + \sigma^{q,a}_t dZ^a_t\right)$$

The diffusion terms of the wealth equation are given by

$$\sigma_{w_e,t} = \frac{q_t k_{e,t}}{w_{e,t}}\chi_t \tag{1.9}$$
$$\sigma_{w_h,t} = \frac{q_t k_{h,t}}{w_{h,t}} + (1-\chi_t)\frac{q_t k_{e,t}}{w_{h,t}} \tag{1.10}$$

The experts retain a fraction χ_t of risk in their balance sheet, and hence the fraction of capital invested in the diffusion terms is multiplied by this quantity. The households receive the remaining risk that enters into the second part of equation (1.10). The households face a no-shorting constraint $k_{h,t} \geq 0$. I define

$$\epsilon_{e,t} := \zeta^k_{e,t}(\sigma + \sigma^{q,k}_t) + \zeta^a_{e,t}\sigma^{q,a}_t + \varphi(\zeta^a_{e,t}(\sigma + \sigma^{q,k}_t) + \zeta^k_{e,t}\sigma^{q,a}_t) \tag{1.11}$$
$$\epsilon_{h,t} := \zeta^k_{h,t}(\sigma + \sigma^{q,k}_t) + \zeta^a_{h,t}\sigma^{q,a}_t + \varphi(\zeta^a_{h,t}(\sigma + \sigma^{q,k}_t) + \zeta^k_{h,t}\sigma^{q,a}_t) \tag{1.12}$$

There are two prices of risk for each type of the agent: $\zeta^k_{j,t}$ and $\zeta^a_{j,t}$, corresponding to the capital shock and the productivity shock, respectively. That is, by borrowing in the risk free market at a rate r_t and investing in the risky capital, they obtain the prices of risk $\zeta^k_{j,t}$ and $\zeta^a_{j,t}$. The exit rate of experts does not enter into the individual wealth equation, but it appears in the evolution of aggregated expert wealth as shown in Appendix A.1.1. In fact, there are an infinite number of agents in the economy, but each individual in types $\mathbb{E}$ and $\mathbb{H}$ is 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 1.2.1. A competitive equilibrium is a set of aggregate stochastic processes adapted to the filtration generated by the Brownian motions Z^k_t and Z^a_t. Given an initial distribution of wealth between the experts and households, the processes are prices (q_t, r_t), policy functions $(c_{j,t}, \iota_{j,t}, \psi_t; j \in \{e,h\})$ and net worth $(w_{j,t}; j \in \{e,h\})$, such that

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

Asset pricing conditions The agents choose the optimal rate of investment by maximizing their return on holding capital. That is, $\iota_{j,t}$ solves[22]

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

The optimal investment rate is obtained as

$$\iota^*_{j,t} = \frac{q_t - 1}{\kappa} \tag{1.13}$$

The investment rate is the same for both types of 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 investment and causes output growth to accelerate (since $\Phi'(\cdot) > 0$). The asset pricing relationship for experts is given by[23]

$$\frac{a_{e,t} - \iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu^q_t + \sigma\sigma^{q,k}_t + \varphi\sigma\sigma^{q,a}_t - r_t = \chi_t\epsilon_{e,t} + (1-\chi_t)\epsilon_{h,t} \tag{1.14}$$

where $\epsilon_{j,t}$ is defined in (1.11) and (1.12). The experts will issue the maximum allowed equity ($\chi_t = \underline{\chi}$) if the premium demanded by them is higher than that required by households. The pricing condition of households is given by

$$\frac{a_h - \iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu^q_t + \sigma\sigma^{q,k}_t + \varphi\sigma\sigma^{q,a}_t - r_t \leq \epsilon_{h,t} \tag{1.15}$$

where the equality holds if $\psi_t < 1$. We can combine (1.14) and (1.15) and write the asset pricing condition as

$$\frac{a_{e,t} - a_h}{q_t} \geq \chi_t(\epsilon_{e,t} - \epsilon_{h,t}) \tag{1.16}$$

$$\min\{\chi_t - \underline{\chi}, \quad \epsilon_{e,t} - \epsilon_{h,t}\} = 0 \tag{1.17}$$

Equation (1.16) holds with equality if $\psi_t < 1$. Equation (1.17) states that whenever the risk premium of experts is larger than that of households, experts issue the maximum outside equity (i.e., $\chi_t = \underline{\chi}$). When experts are wealthy enough such that the constraint is no longer binding, the risk premium becomes equal. 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 $a_{e,t}$.[24] The equilibrium conditions map the optimal consumption, investment, capital share, and the capital price to the history of Brownian shocks Z^k_t and Z^a_t through state variables $(z_t, a_{e,t})$ which has a domain denoted by Ω. The

[22] Note that the only component in the expected return that contains investment rate is $\Phi(\iota_{j,t}) - \frac{\iota_{j,t}}{q_t}$.

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

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

wealth share is defined as

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

where $W_{e,t} = \int_{\mathbb{E}} w_{j,t}\,dj$. Moving forward, I write $X_{h,t}$ and $X_{e,t}$ to denote the aggregated quantities $\int_{\mathbb{H}} x_{j,t}\,dj$ and $\int_{\mathbb{E}} x_{j,t}\,dj$ respectively and characterize the model with a representative household and expert.[25]

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

$$\frac{dz_t}{z_t} = \mu_t^z\,dt + \sigma_t^{z,k}\,dZ_t^k + \sigma_t^{z,a}\,dZ_t^a \tag{1.18}$$

where

$$\begin{aligned}
\mu_t^z = {} & \frac{a_{e,t}-\iota_t}{q_t} - \frac{C_{e,t}}{W_{e,t}} + \left(\frac{\chi_t\psi_t}{z_t} - 1\right)\Big((\sigma+\sigma_t^{q,k})(\hat{\zeta}_{e,t}^1 - (\sigma+\sigma_t^{q,k})) + \sigma_t^{q,a}(\hat{\zeta}_{e,t}^2 - \sigma_t^{q,a}) - 2\varphi(\sigma+\sigma_t^{q,k})\sigma_t^{q,a}\Big) \\
& + (1-\chi_t)\big((\sigma+\sigma_t^{q,k})(\hat{\zeta}_{e,t}^1 - \hat{\zeta}_{h,t}^1) + \sigma_t^{q,a}(\hat{\zeta}_{e,t}^2 - \hat{\zeta}_{h,t}^2)\big) + \frac{\lambda_d}{z_t}(\bar{z}-z_t) - \tau(a_{e,t},z_t) \\
\hat{\zeta_{j,t}^1} = {} & \zeta_{j,t}^k + \varphi\zeta_{j,t}^a; \quad j \in \{e,h\} \\
\hat{\zeta_{j,t}^2} = {} & \zeta_{j,t}^a + \varphi\zeta_{j,t}^k; \quad j \in \{e,h\} \\
\sigma_t^{z,k} = {} & \left(\frac{\chi_t\psi_t}{z_t} - 1\right)(\sigma+\sigma_t^{q,k}) \\
\sigma_t^{z,a} = {} & \left(\frac{\chi_t\psi_t}{z_t} - 1\right)\sigma_t^{q,a}
\end{aligned}$$

Proof: See Appendix A.1.1.

The parameters λ_d and $\bar{z}$ denote the death rate and the mean proportion of experts in the economy at each time instant, respectively. The exit rate τ_t, whose functional form is given in the equation (1.22), enters the drift of the wealth share.

1.2.1 Model solution

The solution method is similar to value function iteration, with an inner static loop used to solve 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 the value functions using a deep neural network architecture. The first step solves for equilibrium policies from the value function, which is set to take an arbitrary value at time T. This is analogous to 'policy improvement' in the reinforcement learning literature. In the second step, the neural network solves for the value function at time $T-\Delta t$, taking policies computed in first step as given, which is then used to update policies in the subsequent step. This corresponds to the 'policy evaluation' in the language of reinforcement

[25] That is, since each agent within their respective group are identical, solving for the aggregate agent policies are enough.

learning.[26] The two-step procedure is performed repeatedly until the value function converges. I present and discuss the equilibrium policies, and relegate the methodological details to Appendix A.1.1.

Static decisions and HJB equations: The value function is given by $U_{j,t}$ and the HJB for optimization problem (1.7) can be written as

$$\sup_{C_{j,t},K_{j,t}} f(C_{j,t},U_{j,t}) + E[dU_{j,t}] = 0 \tag{1.19}$$

Homothetic preferences imply that the value function is of the form

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

with the process for the stochastic opportunity set defined as

$$\frac{dJ_{j,t}}{J_{j,t}} = \mu^J_{j,t}dt + \sigma^{J,k}_{j,t}dZ^k_t + \sigma^{J,a}_{j,t}dZ^a_t \tag{1.20}$$

The aggregate wealth dynamics of experts is given by

$$\begin{aligned}\frac{dW_{e,t}}{W_{e,t}} = &\left(r_t - \frac{C_{e,t}}{W_{e,t}} + \frac{q_tK_t}{W_{e,t}}\epsilon_{e,t} - \lambda_d + \frac{\bar{z}\lambda_d}{z_t} - \tau(z_t,a_{e,t})\right)dt \\ &+ \chi_{e,t}\frac{q_tK_t}{W_{e,t}}(\sigma+\sigma^{q,k}_t)dZ^k_t + \chi_{e,t}\frac{q_tK_t}{W_{e,t}}\sigma^{q,a}_t dZ^a_t\end{aligned} \tag{1.21}$$

The terms involving λ_d are due to the birth and death, and $\tau(z_t,a_{e,t})$ is the state dependent exit rate. I assume the following function for the exit rate.

$$\tau_t = \tau_n 1_{\Omega_n}(z_t,a_{e,t}) + \tau_c 1_{\Omega_c}(z_t,a_{e,t}) \tag{1.22}$$

where $\Omega_c = \{(z_t,a_{e,t})|z_t \le z^*(a_e)\}$ is the endogenous region in state space at which the capital gets misallocated, and the economy is in crisis. That is, $z^*(a_e)$ denotes the crisis boundary at which experts find it optimal to fire-sell the capital to households, triggering the financial amplification mechanism. The region $\Omega_n = \{(z_t,a_{e,t})|z_t > z^*(a_e)\}$ corresponds to the normal regime with high output and low risk premium, and $\Omega = \Omega_c \cup \Omega_n$. The parameters (τ_c, τ_n) are

[26] 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.

calibrated to the observed bank default rates. The HJB equation is written as[27]

$$\rho\big[\log\frac{C_{j,t}}{W_{j,t}}-\log J_{j,t}+\log(q_t z_{j,t})\big]+\big(\Phi(\iota)-\delta\big)-\frac{\gamma}{2}\sigma^2+\mu^J_{j,t}-\frac{\gamma}{2}\big((\sigma^{J,k}_{j,t})^2+(\sigma^{J,a}_{j,t})^2+2\varphi\sigma^{J,k}_{j,t}\sigma^{J,a}_{j,t}\big)$$
$$+(1-\gamma)(\sigma\sigma^{J,k}_{j,t}+\varphi\sigma\sigma^{J,a}_{j,t})+1_{j\in\mathbb{E}}\frac{\tau_t}{1-\gamma}\left(\left(\frac{J_{j',t}}{J_{j,t}}\right)^{1-\gamma}-1\right)=0 \tag{1.23}$$

where the last term on the left hand side is due to the exit.[28]

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

$$\hat{C}_{j,t}=\rho \tag{1.24}$$
$$\zeta^k_{e,t}=-(1-\gamma)\sigma^{J,k}_{e,t}+\sigma^{z,k}_t+\sigma^{q,k}_t+\gamma\sigma \tag{1.25}$$
$$\zeta^a_{e,t}=-(1-\gamma)\sigma^{J,a}_{e,t}+\sigma^{z,a}_t+\sigma^{q,a}_t \tag{1.26}$$
$$\zeta^k_{h,t}=-(1-\gamma)\sigma^{J,k}_{h,t}-\frac{z_t}{1-z_t}\sigma^{z,k}_t+\sigma^{q,k}_t+\gamma\sigma \tag{1.27}$$
$$\zeta^a_{h,t}=-(1-\gamma)\sigma^{J,a}_{h,t}-\frac{z_t}{1-z_t}\sigma^{z,a}_t+\sigma^{q,a}_t \tag{1.28}$$

Proof: See Appendix A.1.1.

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 $(J_{j,t},\chi_t,\psi_t,q_t,\sigma^{q,k}_t,\sigma^{q,a}_t)$ which are found by solving for a Markov equilibrium in the state space $\Omega:=\mathbf{z_t}\in(\mathbf{0},\mathbf{1})\times\mathbf{a_{e,t}}\in(\underline{\mathbf{a}}_\mathbf{e},\bar{\mathbf{a}}_\mathbf{e})$.

Definition 1.2.2. A Markov equilibrium in Ω 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 $\hat{C}_e(z_t,a_{e,t}),\hat{C}_h(z_t,a_{e,t}),\psi(z_t,a_{e,t}),\chi_t(z_t,a_{e,t}),\iota_t(z_t,a_{e,t})$, and state variables $\{z_t,a_{e,t}\}$ such that

- $J_{j,t}$ solves the HJB equation and the corresponding policy functions
- Markets clear

$$(\hat{C}_{e,t}z_t+\hat{C}_{h,t}(1-z_t))q_t=\psi_t(a_{e,t}-\iota_t)+(1-\psi_t)(a_h-\iota_t) \tag{1.29}$$
$$\frac{q_tK_{e,t}}{W_{e,t}}z_t+\frac{q_tK_{h,t}}{W_{h,t}}(1-z_t)=1 \tag{1.30}$$

- z_t and $a_{e,t}$ satisfy (1.18) and (1.3) respectively

[27] The value function is conjectured to be a function of aggregate capital K_t, instead of the wealth using the relation $z_t=\frac{W_{e,t}}{q_tK_t}$. Hence, the capital share does not enter the HJB equation directly. See Appendix A.1.1 for further details.

[28] The index j' refers to the other type of agent. That is, for the case of experts, j' refers to households. Note that $z_{j,t}$ equals z_t in the case of experts and $1-z_t$ in the case of households.

Similar to BS2016, there are three regions in the state space that describe the mechanisms of risk-sharing, except that the state space is two-dimensional. In the first region (Ω_c), where z_t is low, the risk premium of experts is high enough such that condition (1.16) holds with equality. In this region, the experts issue maximum allowed equity $1-\underline{\chi}$ to households since their risk premium is high. In the second region, the experts hold all capital in the economy. This corresponds to the case when $\psi = 1$ but the risk premium of experts is still larger than that of households. As a result, they issue the maximum allowed equity (i.e., $\chi_t = \underline{\chi}$). In the third reigon, experts still hold all the capital (i.e., $\psi = 1$) as before, but they now issue outside equity such that $\epsilon_{e,t} = \epsilon_{h,t}$. This is the region where experts are wealthy enough such that the skin-in-the game constraint is no longer binding, and the risk premium of experts and households are equal. The second and third region together form Ω_n.

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

$$||\sigma_t^R||^2 := (\sigma + \sigma_t^{q,k})^2 + (\sigma_t^{q,a})^2 = \frac{\sigma^2 + \left(\frac{\sigma_{ae,t}^2}{q_t}\frac{\partial q_t}{\partial a_{e,t}}\right)^2}{\left(1 - \frac{1}{q_t}\frac{\partial q_t}{\partial z_t} z_t\left(\frac{\psi_t \chi_t}{z_t} - 1\right)\right)^2} \tag{1.31}$$

Proof. See Appendix *A*.1.1.
The first term in the numerator on the R.H.S of equation (1.31) reflects the fundamental volatility while the second term captures the contribution of productivity shocks. There are two effects that drive the total volatility: (a) Since $\frac{\partial q_t}{\partial z_t} > 0$, and $\frac{\psi_t \chi_t}{z_t} \geq 1$ in equilibrium[29] in the crisis region, 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). The equations (1.29), (1.31), and (1.16) are used to solve for $(q_t, \sigma_t^{q,k}, \sigma_t^{q,a}, \chi_t, \psi_t)$. The remaining equilibrium objects can be obtained from these quantities. Appendix A.1.1 explains the solution steps in detail.

1.2.2 Calibration

The calibration strategy follows the standard procedure in the literature where each model parameter is identified with a moment. Table (1.1) presents the list of parameters with the targeted moments.

RBC parameters: The investment cost parameter is calibrated to generate an investment-capital ratio of 10%. The depreciation rate is chosen to match the average investment rate.[30] The conditional risk premium in the model is determined by the productivity gap between households and experts. The predictive regression results from (1.2) estimate the risk premium

[29] The quantity $\frac{\psi_t \chi_t}{z_t}$ is the experts exposure to the investment in risky capital. This quantity is larger than 1 whenever the expected return of experts is greater than that of households, which is the case in crisis region.

[30] The average investment rate in the data is around 14%, with a volatility of 4.7% between the year 1975 to 2015 (He and Krishnamurthy (2019)).

Table 1.1 – Calibrated parameters

	Description	Choice	Target
Technology	Volatility of capital (σ)	0.06	Vol (Risk premium)
	Discount rate (ρ)	0.05	Literature
	Depreciation rate of capital (δ)	0.04	GDP growth
	Investment cost (κ)	3	Investment-capital ratio
	Productivity gap ($\underline{a}_e - a_h$)	0.08	Conditional risk premium
	Correlation of shocks (φ)	0.5	Data
Utility	Risk aversion (γ)	7	Unconditional risk premium
Demographics	Mean proportion of experts ($\bar{z}$)	0.10	Literature
	Turnover (λ_d)	0.03	Literature
Expert Productivity	Mean reversion rate (π)	0.01	Duration of crisis
	Variance (v)	12.5	Data
	Upper level ($\bar{a}_e$)	0.2	Probability of crisis
Exit rate	Normal state (τ_n)	0.055	Literature
	Crisis state (τ_c)	0.66	Data
Friction	Equity retention ($\underline{\chi}$)	0.65	Literature

Note: All values are annualized.

conditional on crises to be around 25%. I choose the gap between expert productivity $\underline{a}_e$ and a_h to target a 25% conditional risk premium. The correlation of shocks to the level of capital and expert productivity is chosen to be 0.5. This is guided by a -0.48 empirical correlation between bank efficiency ratio and log GDP in the period from 1996Q1 to 2020Q4.[31]

Preference parameters: The discount rate is taken to be 5% from the literature (closer to the 4% used in Gertler and Kiyotaki (2010) and Krishnamurthy and Li (2020)). The risk aversion parameter γ is chosen to be 7, which targets an unconditional risk premium of 5%. A larger risk aversion parameter implies a higher risk premium implied by the model, but too large a value can reduce the probability of crisis. I find the value of 7 to be a good balance between the unconditional risk premium and the crisis frequency. The death rate is chosen to be 3%, meaning that experts live on average for 37 years. 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. The fraction of new born agents designated as experts is calibrated to 0.1 following Hansen et al. (2018).

Productivity parameters: The parameter π governs the persistence of productivity and is chosen to target the duration of crisis. The empirical bank efficiency cycle is highly persistent with an AR(1) correlation coefficient of 0.77. The parameter π is calibrated to generate a persistent productivity process. The volatility parameter v is calibrated such that the variance of the simulated productivity process is approximately equal to the empirical bank efficiency variance of 7%. The upper level of productivity $\bar{a}_e$ is chosen to target the probability of crisis.

[31] Bank efficiency ratio is measured as the asset-weighted average of non-operating cost to income ratio for bank holding companies in the US. The data is at quarterly frequency from 1984Q1 till 2020Q4. A higher ratio indicates lower efficiency.

Other parameters: Expert exit rates are parameterized by τ_n and τ_c; these are important in governing the transition into and out of crisis. The baseline rate τ_n is set to 5.5%, comparable to the 6% rate in Krishnamurthy and Li (2020). Empirically, bank default volume is approximately 15 times larger compared to the periods outside of the crisis.[32] Taking into account a recovery rate of 20%, the exit rate during a crisis is around 12 times larger than during the normal period, implying a τ_c of 66%. 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).

Figure (1.2) 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 allows the return volatility to be higher than the fundamental volatility even in the normal regime. When the wealth share of 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 sales. In the crisis zone, experts begin selling capital to households, who always place a lower value on it. 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 the failure to internalize the pecuniary externality by the agents. This is because each individual in the economy takes 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 households will be forced to hold it by market clearing. Since 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 individual agents within the expert group do not take into consideration, which is the increased exit rate when the system enters the crisis region.[33] 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 latter will be gradual.

The jump in prices due to the fire-sale effect can only be explained by the differences in productivity rates in an incomplete market setting and no-shorting constraint. There will be a state space where experts hold all the capital since the risk premium of households is lower than that of experts. In such states, households would desire to hold a negative quantity

[32] This is based on historical data from the FDIC between the year 2001 till 2020. During the period 2008-2011, the total bank default volume is approximately USD 667 billion, whereas outside this period, the total default volume is USD 44 billion.

[33] The aggregate experts' wealth is, however, affected by the exit, which can be seen in equation (1.21).

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 were smaller, they would still desire to hold some positive quantity of capital. This smooths the transition from the normal to the crisis regime.[34]

Figure 1.2 – Model solution

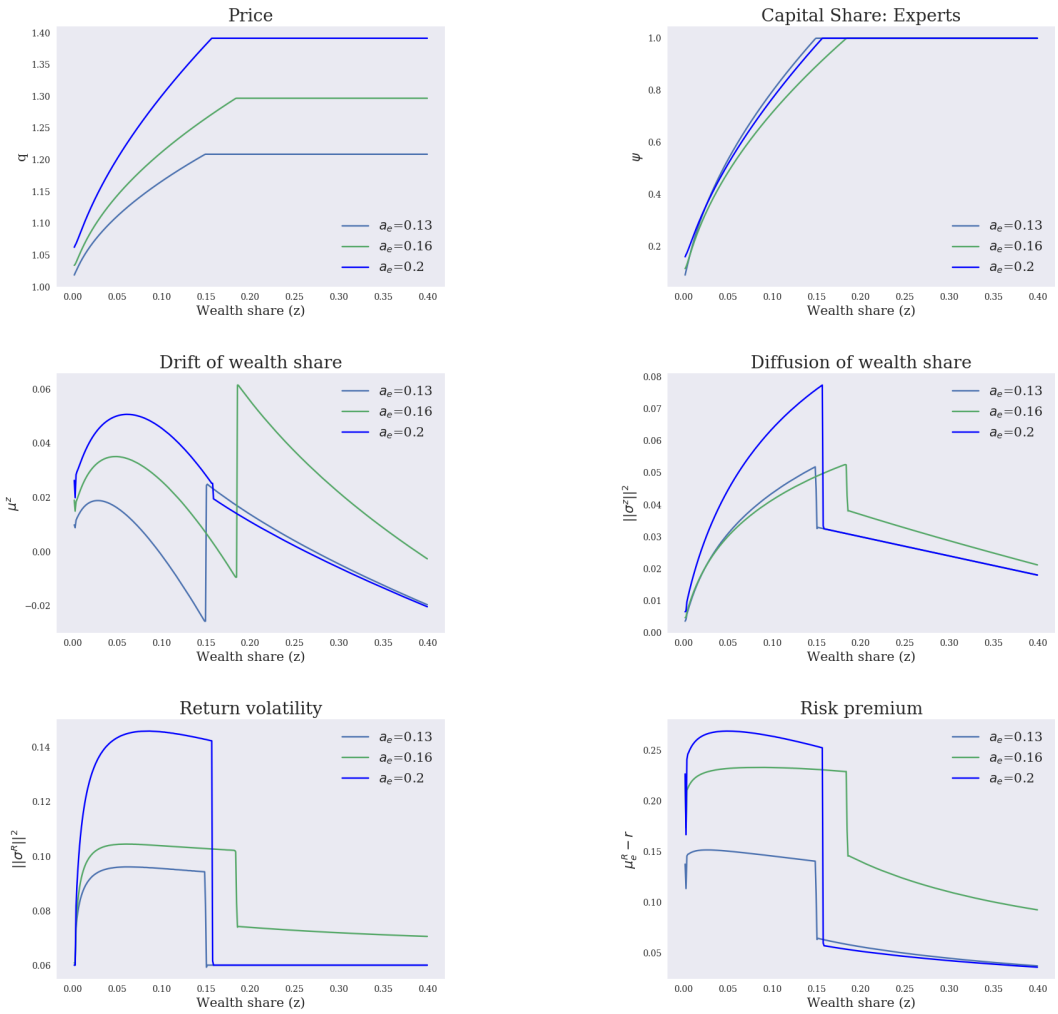

Note: Equilibrium values as functions of the state variable wealth share (z_t) for different values of expert productivity ($a_{e,t}$).

1.3 Quantitative analysis

In this section, I consider 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

[34] 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.

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.

1.3.1 Benchmark model

I assume that the productivity rate of both experts and households is constant such that $a_e > a_h$ holds, and the exit rate is zero. 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 1.2. That is, the output is given by AK technology as in (1.1), with a_e and a_h as the productivity rates of the experts and the households respectively. The evolution of capital is governed by (2.3) as before. Appendix A.1.2 presents the model in detail along with the numerical procedure and the solution.

Comparative statics: Figure (1.3) depicts the risk premium of experts in the benchmark model, as well as the stationary density of expert wealth share.[35] The static comparison from the left-hand side figure in (1.3) shows that as the risk aversion increases, so does the premium on the risky capital for experts. The other equilibrium objects, such as capital price, return volatility, capital share of experts, drift of wealth share, and volatility of wealth share are solved similar to the high-dimensional model in the previous section. The endogenous risk is higher in the crisis region when the risk aversion is lower, but it features a smaller crisis region. Also, changes in the market price of risk induced by varying risk aversion lead to vast differences in the drift of wealth share. This has an impact on how the system transitions into and out of the crisis region.

Stationary distribution: While the left panel in Figure (1.3) 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 a monthly frequency. The simulation maps the Brownian shocks Z_t^k to state variable z_t, which is solved for similarly to the two-dimensional model. 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 assess how well the model captures the salient empirical features of financial crises in the

[35]The parameters used for calibration are shown in Table (*A*.1) in Appendix A.1.

data. I define the crisis moments as follows

Figure 1.3 – Risk premium and stationary distribution

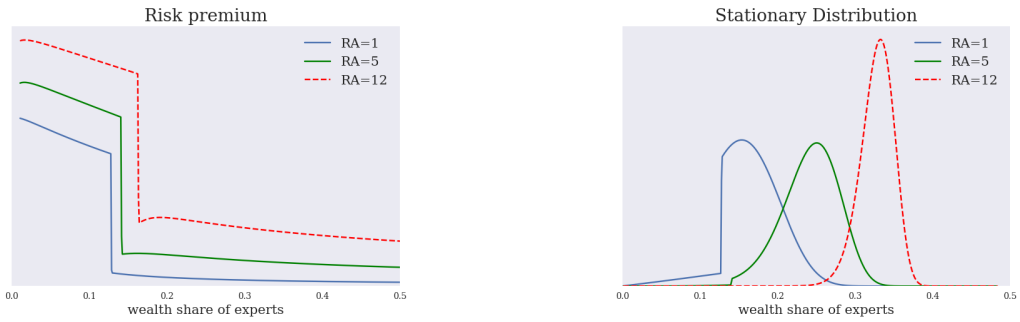

Note: Left panel presents a static comparison of experts risk premium for three different levels of risk aversion. Right panel presents the stationary distribution of expert wealth share for three levels of risk aversion.

1. *Crisis event* is defined as a state where the capital is misallocated to households, and the skin-in-the game constraint is binding. In this state, the risk premium is high, and GDP growth is low, reflecting the empirical nature of the financial crisis. This definition is similar to Maxted (2020), who defines crises as states where the capacity constraint is binding. The *amplification* in my model refers to the moments computed when the economy is in a crisis state.

2. *Probability* of crisis: The proportion of time that the economy spends in a crisis state. Analyzing a large set of advanced economies over several years, Reinhart and Rogoff (2009) estimate this value to be 7% empirically. They define a crisis as a recession accompanied by severe banking panic.

3. *Duration* of crisis: The average amount of months required to recover and revert to normalcy after entering the crisis state. The average length of contraction cycle from NBER recessionary data is around 18 months. The simulated mean duration of the crisis is taken to be the model-implied *persistence*.

The right panel of Figure (1.3) plots the stationary distribution of the wealth share for three different risk aversion levels. As risk aversion increases, the mass of wealth share that lies in the crisis zone diminishes. 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. Looking at the left panel of Figure (1.3), it appears as if increasing risk aversion will not have a drastic impact on the frequency of a crisis since the boundary z^*, the point at which the risk premium jumps and the investment rate falls, moves only slightly to the

right.[36] However, higher risk aversion increases the drift of wealth share a lot and 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. The crisis boundary z^* is far from the stochastic steady state $\hat{z}$ for higher levels of risk aversion, which means that a much longer sequence of negative shocks is required to push the system into the crisis region.[37]

Comparison to Data: While the crisis is well defined and endogenously determined in the model, defining the crisis episodes in the data is a challenge. Reinhart and Rogoff (2009) determine the frequency of crisis states to be around 7% for the advanced economy. This figure is much lower than the 29% of percentage NBER recessionary periods from the year 1874 till today.[38] The stark difference in the frequency between Reinhart and Rogoff (2009) and NBER data is due to the fact that in the former, recessionary periods need to feature severe banking panic to qualify 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 during recessions, where a financial crisis occurs when the wealth share of intermediaries deteriorates 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 7% for the purpose of quantitative calibration. For each z_t simulated from the discretized version of its dynamics, the equilibrium quantities are computed using the mapping given by the equilibrium functions. 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 (1.32).

$$R^e_{t+1} = a + \beta * D_t/P_t + \beta_{rec} * 1_{Rec} * D_t/P_t + \beta_{fin} * 1_{fin} * D_t/P_t + \epsilon_t \tag{1.32}$$

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 (2009). 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. In Table (1.2), regression (I) uses only the dividend yield as a regressor, whereas 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 a monthly frequency from the years 1945 to 2021. The indicator functions 1_{Rec}, and 1_{fin} take a value of 1 in months of NBER non-financial recession and financial recession, respectively. The dummy variable corresponding to the financial crisis is positive

[36]The point z^* denotes the point at which the experts start fire selling the capital to the households, and is defined to be the crisis boundary. Formally, $z^* = sup\{z_t \mid \psi_t < 1\}$ where ψ_t is the share of capital held by the experts.

[37]The stochastic steady state can be defined as $\hat{z} := \{z_t : \mu^z_t(z_t) = 0\}$.

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

and statistically significant as seen in Table (1.2).[39] The R-squared value is also higher when controlling for recession and financial crises, indicating 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 (1.2) and compute the standard deviation to obtain the volatility of the risk premium.

Table 1.2 – Risk premium estimation

	(I)	(II)	(III)
const	-0.02	0.00	0.00
	(0.00)	(0.01)	(0.03)
D_t/P_t	2.13***	1.17**	1.15**
	(3.52)	(1.87)	(0.53)
1_{Rec}		2.21***	1.82***
		(4.15)	(0.60)
1_{fin}			2.18***
			(2.99)
No. of obs.	906	906	906
Adj. R2	0.02	0.04	0.05

Note: The variables 1_{Rec} and 1_{fin} represents recessionary and financial crisis episodes respectively. Recessionary episodes are taken from NBER, and financial crisis periods are taken form Reinhart and Rogoff (2009). The dividend yield and stock return data are from Robert Shiller's website. I use a monthly frequency from the years 1945 to 2021. Model (I) excludes both dummy variables to zero. Model (II) excludes financial crisis dummy but includes recession dummy. Model (III) includes both dummy variables.

Other moments: Table (1.3) presents the data moments that the model aims to match with the methodology to compute them. The benchmark model delivers an unconditional average GDP growth rate of around 2.3% and an investment rate of 7%. An important measure of model success is its ability to capture the observed non-linearity in the data. The GDP growth rate conditional on being in a crisis is around -8%. The empirical 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 5.6%. Note that even though the output of experts and households individually moves 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

[39] This finding is robust to using different time periods such as 1871-2021 (time since Shiller's data is available), and 1914-2021 (since the start of Federal Reserve).

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 capital prices alone. A drop in the capital price during the crisis period is not large enough to generate the observed drop in the investment rate.[40] The volatility of investment rate implied by the model is close to zero. Overall, the model captures non-linearity in output growth but misses non-linearity in the mean and volatility of investment rates. This result is comparable to HK2019, which has a realistic consumption volatility but an excessively low investment volatility. This calls for future work to match both output and investment dynamics. The mean leverage of the expert sector implied by the model with unitary risk aversion is 3.5, comparable to the empirical leverage of 3.77.[41] The model also features counter-cyclical leverage. Even though the experts fire sell the assets 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 rising leverage. Table (*A*.2) shows that the correlation between the shock and the leverage ranges from -19% to -26% for different risk aversion levels. This matches the empirical correlation of -18% quite well. Overall, for lower risk aversion levels, the model seems to do well in matching the leverage patterns. Lastly, the model does not generate an excessive asset return volatility (Shiller (1981)). The unconditional return volatility is more or less the same as the exogenous capital volatility of 6%, even though the conditional return volatility is large. 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 1.3 – Data moments and methodology

	Mean	Std dev	Data source
Risk premium* (%)	5.5	4.7	Predictive regression
Risk free rate* (%)	4.0	0.9	Amit Goyal's website
GDP growth*(%)	3.3	3.9	FRED
Investment rate (%)	14	4.7	He and Krishnamurthy (2019)
BHC Leverage	3.77	-	He and Krishnamurthy (2019)
Corr(BHC Leverage, GDP cycle)	-0.18	-	He et al. (2017)
Probability of Crisis (%)	7	-	Reinhart and Rogoff (2009)
Duration of Crisis	18-months	-	NBER cycle

Note: The Table presents unconditional mean and standard deviation of key variables in the data along with the methodology to compute the variables. The variables marked with asterisk are estimated using quarterly frequency data between 1950Q1 till 2021Q1. All percentage values are annualized.

Table (1.4) summarizes the ability of the benchmark model to succeed in different aspects. By

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

[41] This number is taken from HK2019.

Table 1.4 – Model success summary

	Quantity of interest	Success level	Comments
Macroeconomic	GDP/Output growth	High	✓
	Investment rate	Low	Low variation and not enough drop in crisis
Experts	Leverage	High	✓
	Cyclicality of leverage	High	✓
Crisis	Probability of crisis	Moderate	Matching prob. of crisis attenuates crisis dynamics
	Duration of crisis	Low	Matching duration attenuates crisis dynamics
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

Note: The model implied moments and probability of crisis is computed by simulating the model at monthly freqnency for 5000 years. All values are annualized.

far, matching the intermediary leverage pattern and the non-linearity in output growth seem to be the strongest suits of the model. For any reasonable parameters in calibration, the model cannot resolve the tension between unconditional risk premium, conditional risk premium, and crisis persistence. The focus of the next section is to provide a resolution to this problem.

1.4 Resolution of the tension between amplification and persistence of crises

In this section, I quantify the model with stochastic productivity and exit of experts, and show that it resolves the tension between persistence and amplification of financial crises and provides reasonable time variation in the prices. The definition of a crisis event, probability, and duration of crisis is similar to the benchmark model.[42] Figure (1.4) plots the stationary

[42] Note that in this two-dimensional model, the crisis boundary is a function of expert productivity. The simulation results show that the economy enters crisis mostly when the productivity is well below its mean. This can also be verified by inspecting the joint density shown in Figure (1.6).

Figure 1.4 – Stationary density of wealth share

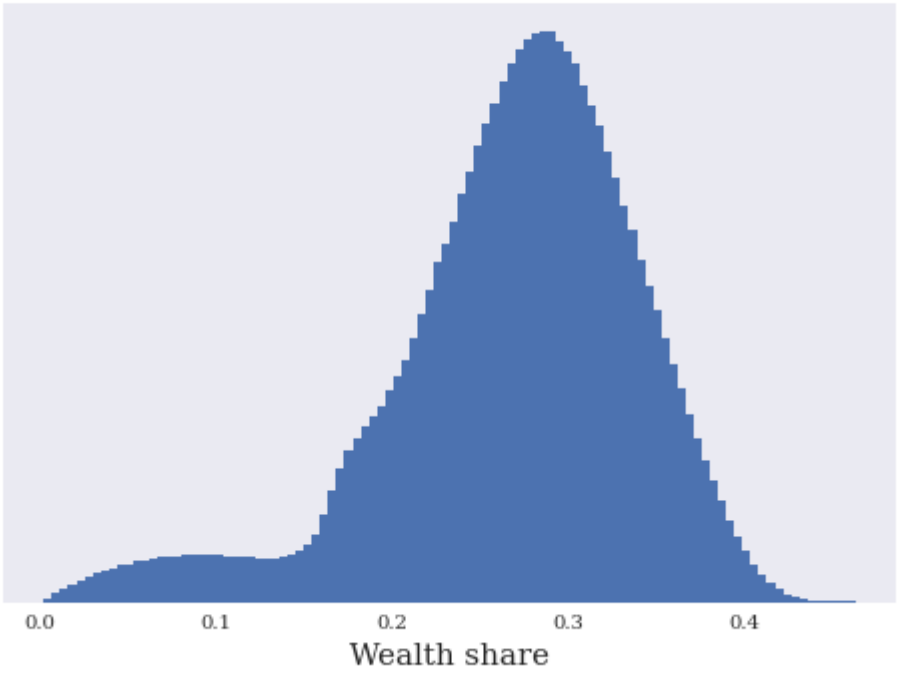

Note: The figure displays stationary marginal density of endogenous wealth share obtained from simulating the model for 5000 years at monthly frequency. The observations are annualized.

marginal distribution of the wealth share obtained through simulation.[43] Table (1.5) presents the average duration of crisis in the benchmark model and the stochastic productivity model and compares them against the data. There is a substantial controversy in the literature regarding the duration of crises (Reinhart and Rogoff (2009)). The NBER reports that the Great Recession started in December 2007 and ended in June 2009, indicating an 18 month duration.[44] 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 my model. The numbers in Table (1.5) can be thought of as the ability of the models to generate the stated duration for a reasonable crisis probability of 7-8%. Both of the benchmark models deliver a duration of crisis that is much lower than observed in the data. The mean duration from my model matches the data quite well although the 10th and 50th percentile values are lower. The parameters used for calibrating my model are shown in Table (1.1).

Figure (1.5) plots the stationary distribution of the wealth share of experts during the time the

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

[44] The average duration of recession in the past 33 cycles from year 1854 to 2020 is around 18 months. Source: https://www.nber.org/cycles.html.

Table 1.5 – Duration of crisis

	Data (NBER)	Benchmark model (RA=1, IES = 1)	Benchmark model (RA=2, IES = 1)	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	6.0	6.5	17.0

Note: Data for computing the empirical duration of crisis is from NBER website. The last three columns presents the model implied duration percentiles obtained from simulating each of the benchmark models for 5000 years at monthly frequency.

system spends in the crisis region. In the benchmark model (left panel), a lot of the mass lies near the crisis boundary of 0.125 compared to the interior region where the wealth share is close to zero. The reason for this is that the benchmark model has only one i.i.d Brownian shock. After a series of negative shocks hit the economy, the system enters a crisis, leading to a sharp increase in the risk premium. Since experts are always leveraged in equilibrium, the risk premium loads positively on the drift of wealth share of experts. Moreover, the assumption of i.i.d Brownian shock implies that a series of negative shocks is often followed by a positive shock. Thus, the experts recapitalize quickly by capturing the high risk premium, leading to short-lived crises. In contrast, the frequency distribution of the wealth share in the crisis region in my model, as shown in the right panel in Figure (1.5), features fatter tails. The economic mechanisms that generate this result rest on three forces. Firstly, negative shocks to the capital impair the net worth of the experts just like in the benchmark model. This is the financial amplification channel that is widely covered in the literature. The second force comes from stochastic productivity. The aggregate banking sector productivity is lower during a crisis state. The key comparative advantage of the experts in my model is that they have a higher productivity rate of operating capital. During bad times, this comparative advantage diminishes.[45] A realistic crisis frequency is obtained even for higher risk aversion levels due to stochastic productivity. With a constant productivity as in the baseline model, the risk-averse experts will always remain wealthy by earning a large premium. Negative shocks to the capital in the stochastic steady state will not be enough to generate realistic crisis events. In my model, negative shocks to the capital also push the experts productivity down, which negatively impacts the risk premium. Hence, a series of negative shocks reduces the premium earned in the normal region and will put downward pressure on the drift of wealth share, eventually causing sufficient deterioration in the net worth of experts to generate crisis events.

The third force is the exit rate of experts, which is higher during bad times. While the overlapping generations capture the demographic changes relating to natural birth and death of

[45] Simulation shows that the crisis zone features both a lower wealth share and a lower productivity of experts as seen in the right panel of Figure (1.6).

agents, the exit rate captures the retirement of the experts. In normal times, experts retire at a rate of 5.5%. When they retire, they don't consume all of their wealth immediately. Instead, they transition into households until death. While the crisis is endogenously determined in my model, as soon as the crisis boundary is hit, the exit rate shoots up from 5.5% to 66%. A higher exit rate during a crisis parsimoniously captures the strikingly large number of bank failures during a financial crisis, as evident in Figure (1.1). The fact that a large fraction of the experts retire and become households means that the proportion of the agents who operate capital more productively is lower in times of distress than in normal times. This has a dominating effect on the drift of wealth share and pushes the economy deeper into the crisis since the drift is negatively affected by exit. The only way for the economy to break out of the crisis is for a remaining smaller proportion of the experts to be more productive again, since higher productivity pushes up the risk premium, enabling the experts to rebuild their wealth.[46] However, the rate at which expert productivity reverts to its mean is low, and this sluggish reversion means that the economy spends a long amount of time in a state of distress until the productivity increases and has a dominant effect on the drift of wealth share. This leads to delayed recovery from the crisis. Once the system is back to normalcy, all capital in the economy is held by the experts, and the financial amplification is shut down.

Table (1.7) compares the moments of key asset pricing and macroeconomic variables between my model and the benchmark model. The unconditional risk premium of 5.0% is comparable to the empirical value of 5.5%, whereas, the benchmark model generates a mere 1.7% premium. Importantly, my model allows for reasonable crisis dynamics by simultaneously generating a high conditional risk premium of 18.2% and long a duration of crisis of 17 months without compromising on the other dimensions. That is, the unconditional mean leverage, GDP growth rate, investment-capital ratio, and correlation between expert leverage and capital shock are comparable to the data.

To further understand the individual roles of stochastic productivity and exit rate in delivering quantitative results, I compare my model against two other benchmark models: a) Model B1, which considers stochastic productivity but without exit, and b) Model B2, which considers constant productivity and state-dependent exit. The trade-offs analyzed in the benchmark model also carry over to model B1. While stochastic productivity helps in generating more time variation in the risk premium compared to the benchmark model, the duration of crisis implied is lower compared to the data, as seen in Table (1.7). Without state-dependent exit, the proportion of experts who recapitalize their balance sheets by earning a large risk premium is high, and the economy recovers quickly from a crisis as a result. Hence, features related to the bankruptcy of intermediaries or bank runs are crucial in explaining the slow recovery from a crisis. This relates to Gertler et al. (2020) who build a quantitative macroeconomic model with bank runs as the main driver behind the 2008 financial crisis. However, incorporating bankruptcy without stochastic productivity is not sufficient to generate realistic crisis dynam-

[46] The consumption-wealth ratio of the agents is constant due to the assumption of a unitary IES. For a non-unitary IES, the consumption-wealth ratio may also also increase due to increased productivity of the experts and contribute positively towards the wealth share of experts.

ics. Table (1.7) presents the results of model B2 that includes state-dependent exit but with constant productivity. When the exit rate is calibrated to empirical default rates, the model B2 leads to a dystopian economy with perennial recession. During bad times, a higher exit rate of experts pushes the economy into crisis. Without the offsetting force of mean-reverting productivity, the effect of the exit rate continues to dominate the drift of wealth share, trapping the economy in a distressed state around 91% of the time.

My model generates a larger drop in the investment rate in the crisis period but falls short of the negative investment rate observed in the data. During the last quarter of 2008, private domestic investment in the United States fell by approximately 8%. The q-theory result in the model ties the investment rate tightly to the capital price. Hence, the capital price needs to fall drastically to generate a fall in the investment rate to the extent that is observed in the data. My model is certainly an improvement over the benchmark in this regard, but more work needs to be done in jointly matching the investment and output dynamics.[47] Lastly, the model implied unconditional volatility of the risk premium is 4.8%, well in line with the empirical value of 4.7% reported in Table (1.6). Overall, my model does a good job of balancing the persistence and the amplification, and delivers a reasonable time variation in the prices.

Table 1.6 – Risk premium moments and probability of crisis

	Data			Benchmark Model (RA=1)		Benchmark Model (RA = 20)	
	All	Recession	Crisis	All	Crisis	All	Crisis
E(Risk premium)	5.5	12.8	25.0	1.7	13.4	5.5	-
Std(Risk premium)	4.7	6.7	7.5	3.1	1.3	0	-
Prob. of Crisis	7.0			6.8		0	

Note: Empirical risk premium moments are computed from the predictive regression (1.32). Probability of crisis is taken from Reinhart and Rogoff (2009). The model implied moments and probability of crisis is computed by simulating the model at monthly freqnency for 5000 years. All values are annualized.

[47] Note that I have assumed a simple logarithmic form to model technological illiquidity following Brunnermeier and Sannikov (2016a). Using other functional forms, for example, as found in Di Tella (2017), also fails to generate a large drop in the investment during crisis periods.

Figure 1.5 – Left tail of the marginal stationary density

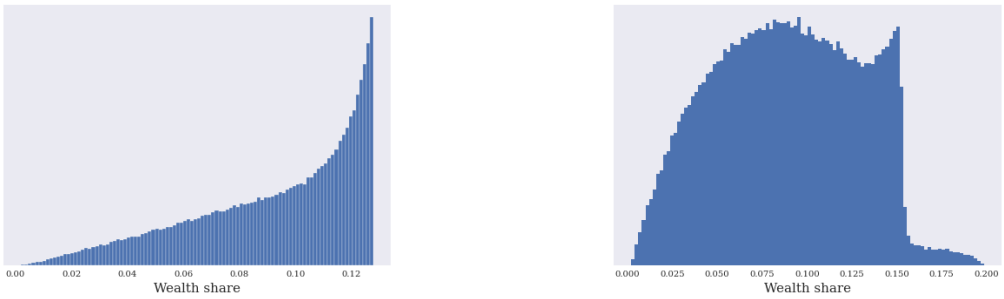

Note: Left panel: Tail of experts wealth share distribution from the benchmark model. Right panel: Tail of experts wealth share distribution from the model with stochastic productivity.

Figure 1.6 – Joint density of state varaibles

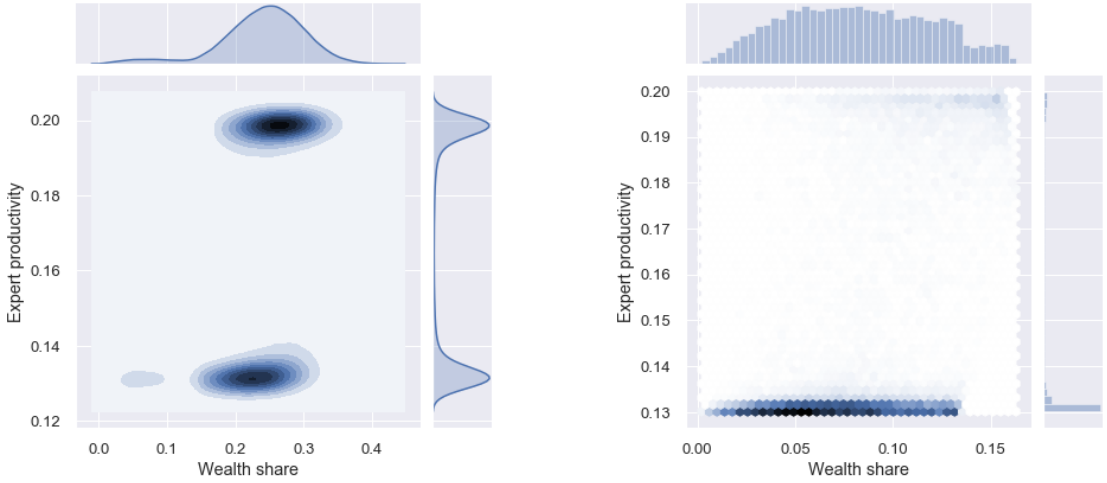

Note: Left panel displays the joint density of wealth share and productivity of experts along with respective marginals in the entire state space. Right panel displays the joint density of wealth share and productivity of experts along with respective marginals in crisis region.

1.5 Conclusion

A financial crisis is characterized by a spike in the risk premium, and a sluggish recovery. Macro-finance models with leveraged intermediaries have trouble explaining these two features simultaneously, especially when they are calibrated to match both conditional and unconditional moments in the data. As a resolution to this puzzle, I have built a macro-finance asset pricing model with intermediaries facing productivity shocks and a state-dependent exit rate. A series of negative capital shocks reduces expert productivity, reflecting the experts' diminishing comparative advantage over households in terms of productivity differential. A simpler model with constant expert productivity and no exit rate cannot simultaneously gen-

Table 1.7 – Summary of key moments

	Data	My model		Benchmark		B1		B2	
	All	All	Crisis	All	Crisis	All	Crisis	All	Crisis
Risk premium (% mean)	5.5	5.0	18.2	1.7	13.4	1.9	14.1	12.7	13.4
Risk premium (% sd)	4.7	4.8	5.2	3.1	2.4	3.5	3.7	2.7	2.9
Investment-capital ratio (% mean)	14	8.7	4.3	7.0	5.6	8.2	7.5	5.5	5.0
BHC Leverage	3.77	4.5	8.8	3.5	5.8	2.7	3.2	3.2	3.5
GDP growth (% mean)	3.2	2.5	-7.1	2.3	-8.0	3.1	-8.2	-6.3	-7.9
Corr(BHC Leverage, GDP)	-0.18	-0.19	-0.01	-0.17	-0.01	-0.15	-0.03	-0.28	-0.05
Probability of Crisis (%)	7.0	8.0		6.8		6.3		90.8	
Duration of Crisis (months)	18.0	17.0		4.0		5.0		13.0	

Note: Comparison of model implied moments. The % values are annualized. The calibrated parameter for my model is given in Table (1.1). The benchmark model does not feature stochastic productivity or exit. The model B1 considers a stochastic productivity but without exit. The model B2 has constant productivity but the experts have a state-dependent exit rate. The calibrated parameters for the benchmark models are given in Table (A.1).

erate amplified and persistent financial crises. There is a trade-off between the risk premium and the probability and duration of crises. I show that auxiliary model features that improve the financial amplification channel dampen the persistence of the crisis.

The richer model with stochastic productivity and a state-dependent exit rate of the experts resolves this tension and quantitatively generates a high risk premium, a large drop in output, decreased financial intermediation, and prolonged distress periods. The twin forces of state-dependent exit and stochastic productivity are at the core of improved dynamics in my model. In particular, a higher exit rate and lower productivity of experts in bad times forces the economy to dip deeper into recession, which eventually revives once productivity mean reverts. The model also generates a large time variation in the risk prices due to the stochastic nature of expert productivity, which is absent in the benchmark model. An interesting avenue for future research is to build a model that endogenously causes variation in the expert productivity, which is an exogenous force in my model. 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 the numerical setup, opening up new avenues to model asset prices with frictions in potentially large dimensions.

2 ALIENs and Continuous-Time Economies

2.1 Introduction

The past decade has seen a surge in macroeconomic and asset pricing models to capture non-linear global dynamics. Models in continuous time offer the ability to capture complex interactions due to their tractability. While the continuous time models characterizing the global dynamics is certainly an improvement to models with linearized solutions, most of the papers resort to smaller state spaces due to computational bottlenecks. The difficulty is amplified when the state variables are endogenous, correlated, and equilibrium quantities exhibit stark non-linearities. For example, D'Avernas and Vandeweyer (2019) show that even in the case of two space dimensions, the standard finite difference method breaks down since it is difficult to preserve the monotonicity of finite difference schemes. Prior literature has addressed this problem by applying an up-winding scheme to the finite difference method, a technique that is borrowed from fluids dynamics, but it is not guaranteed to work in the presence of correlated state variables.[1] Moreover, finite difference methods are not easily scalable, especially when an implicit scheme is employed. This is because an implicit scheme results in a large linear system to be solved, which quickly becomes computationally infeasible when the state space dimension grows.

In this chapter, I propose a distributed deep-learning based technique called Actively Learned and Informed Equilibrium Nets (ALIENs), which can be used to solve a large class of continuous time models in financial economics featuring highly non-linear equilibrium policies and endogenous and correlated state variables with heterogeneous agents. The contribution of the chapter is two-fold. First, I present a general setup of a heterogeneous agent portfolio choice problem and demonstrate how one can solve the model by approximating the Hamilton-Jacobi-Bellman equation using neural networks. I prove that ALIENs with at least one hidden layer offer theoretical convergence to the resulting PDEs when the number of neurons in the hidden layer is large. This result can be thought of as a universal approximation theorem for

[1] See Brunnermeier and Sannikov (2016b),Achdou et al. (2014b),Gomez (2019) etc., for the application of up-winding scheme.

quasi-linear parabolic PDEs that are ubiquitous in continuous time finance and macroeconomic models. Second, I solve two models that feature heterogeneous agents and endogenous state variables. The first model is inspired by Brunnermeier and Sannikov (2016b). It serves as a benchmark to test the neural network approach. I show that the equilibrium policies obtained from a finite difference scheme with up-winding match the solution obtained from ALIENs. The second model adds more shocks to the benchmark model in the spirit of Hansen et al. (2018) and shows how active learning plays a crucial role in ensuring convergence.

The models demonstrated in this chapter feature endogenous jumps in return volatilities and risk prices, and are canonical examples of a macro-finance model with financial frictions that have received substantial attention in the literature. In both the benchmark model and the richer model, the wealth share of the intermediary sector is an endogenous state variable. The value function and other equilibrium policies are a function of the wealth share, which moves endogenously with respect to the underlying shocks in the economy. Many of the heterogeneous agent asset pricing models have a similar structure where the model boils down to solving a system of elliptical PDEs. These PDEs are typically highly non-linear and introduce instabilities when solved using standard methods like the finite difference scheme. Even seemingly small errors get propagated over time and the equilibrium policy functions end up with instabilities similar to 'Gibbs phenomenon' in the spectral theory literature, especially when the policies have jumps.[2] To tackle this problem, it is standard in the literature to convert them into a system of quasi-linear parabolic PDEs by adding a pseudo time dimension. This practice, otherwise known as 'false-transient' method, transforms an elliptical PDE into a parabolic PDE and enables a marching solution (Mallinson and de Vahl Davis (1973)). While most of the asset pricing literature in the past have dealt with parabolic PDEs in one or two space dimensions, solving such PDEs in higher dimensions, remain an open challenge. There is a large literature in the applied mathematics area that deals with PDEs in high dimensions but most of them deal with simple problems that do not feature endogenous state variables with highly non-linear functions as PDE coefficients, that in turn endogenously depend on the policy and the value functions.

Recent applications of deep learning in economics aim to directly approximate the equilibrium functions by simulating the ergodic density. For example, Azinovic et al. (2019) solve discrete time overlapping generations problems by training a deep neural network on simulated data, thereby side-stepping the need for labeled points. They argue that in most cases, the ergodic distribution of state variables lies in a small subspace of the entire state space, and hence one can solve for the equilibrium functions in this small subspace without spending resources and time on part of the state space that does not matter. Simulating from the ergodic density is computationally cheap if a closed form solution for the distribution exists. However, in the cases where state variables are endogenous and depend on the policy and value functions, these functions have to be approximated first before simulating the state variables through a Monte-Carlo procedure. In continuous time macro-finance models, the transformation of

[2] Gibbs phenomenon refers to the instability at jump discontinuous points when Fourier series is used to approximate a function with jump discontinuity. See Kelly (1996) for a nice exposition.

elliptical PDEs into parabolic PDEs already introduces some degree of error. A Monte-Carlo type simulation of state variables using approximated policy and value functions adds further approximation errors. Secondly, typical problems in asset pricing and macro-finance are characterized by high non-linearities in recessionary periods, which are rare occurrences. Therefore, one has to simulate an extraordinarily large number of points in order to have a realistic mass in the state space corresponding to recessions in order to better approximate the equilibrium functions in that subspace and avoid instabilities in future iterations. The framework that I propose takes a different approach. Sparse training points are sampled throughout the state space, but the active learning enforces the sampling to come from the regions with a low tolerance for approximation error from the neural networks. Such a smart-sampling procedure minimizes the error in the initial time iterations so that instabilities in the future iterations are prevented.

I take advantage of recent advancements in high performance computing and offer a parallelized solution using an open-source library developed by Uber called Horovod (see Sergeev and Balso (2018)). Horovod provides a framework for data parallelism where the input is split into mini-batches and transmitted across several nodes, where each node is accompanied by many GPUs and/or CPUs. The model in each node is the same but receives different inputs for training the neural network. Using a Message Passing Interface (MPI) that enables communication across the nodes, the output from each node is averaged using a ring-Allreduce operation. This procedure significantly speeds up the computation time and is similar in spirit to the parallelization scheme in Azinovic et al. (2019). The examples that I have chosen in this chapter come from macro-finance, but the framework can be applied to a variety of problems in finance. For example, a continuous time version of Bansal and Yaron (2004) with mean-reverting long run growth yields a parabolic PDE for the price-dividend ratio. Other problems in asset pricing where this framework is applicable include Wachter (2013), Gârleanu and Panageas (2015), Haddad (2012), Di Tella (2017), Gomez (2019), Di Tella (2019) among others. In macroeconomic models where constraints play an important role, ALIENs offer the possibility to easily encode the economic information from the constraints as regularizers. Importantly, if there are multiple constraints, the framework allows to attach different weights to them depending on their importance in the economic problems. Models where constraints play an important role and can be solved using ALIENs are Achdou et al. (2017), Bolton et al. (2011) etc., among others. While most of these papers are restricted to smaller state spaces due to the curse of dimensionality, ALIENs provides scalability that can be leveraged to solve problems in larger dimensions. For example, Gopalakrishna (2020) shows that a two-dimensional macro-finance model can jointly explain various macroeconomic and asset pricing moments.

2.2 Literature Review

This chapter relates mainly to two strands of literature. First, there has been an explosion of macroeconomic models with financial frictions to characterize the global dynamics after the financial crisis of 2008. The older literature in this area can be traced to Kiyotaki and

Moore (1997) and Bernanke et al. (1998) who study the problem of financial acceleration in a discrete time setting. In the past decade, seminal contributions in this area have come from Brunnermeier and Sannikov (2014a), He and Krishnamurthy (2013), Di Tella (2017), He and Krishnamurthy (2019), and Adrian and Boyarchenko (2012) using continuous time machinery since setting up problems in continuous time offers some degree of tractability that the discrete time models fail to provide. Other related papers in the asset pricing area such as Basak and Cuoco (1998), Basak and Shapiro (2001), Gromb and Vayanos (2002), Gârleanu and Pedersen (2011) etc. study an endowment economy. Given the potential offered by continuous time framework, a second wave of papers emerged capturing more complex dynamics of macroeconomic and financial sector. Drechsler et al. (2018), Di Tella (2017), and Silva (2020) analyze the impact of monetary policy, and Adrian and Boyarchenko (2012), Caballero and Simsek (2017) study macro-prudential policies. Often times, continuous time macro-finance models do not admit full fledged closed form solutions and typically involve solving for a system of partial differential equations (PDEs) to obtain the policy functions.

Hansen et al. (2018) and D'Avernas and Vandeweyer (2019) provide robust solutions to solve the PDEs using some variations of finite difference schemes. Hansen et al. (2018) proposes an implicit finite difference scheme with up-winding and employ parallelization techniques to tackle the problem of solving large linear systems. The technique proposed in this chapter dominates their approach since it offers the advantage of ease in setting up the numerical scheme in two ways. Firstly, my approach allows to accommodate different types of HJB equations with relative ease. For instance, adding a jump term to the capital process will alter the HJB equation and requires setting up a different linear system to solve in the case of an implicit finite difference scheme. This process can be painful depending on the nature of the problem. On the contrary, adding a jump term in my approach requires simply augmenting the regularizer by adding a derivative term, which can easily be accomplished through automatic differentiation. Secondly, adding more features to the existing model and scaling up the dimension again requires setting up a new linear system to solve in the case of an implicit finite difference scheme, whereas, this can be accomplished by simply adding the required higher order partial derivatives to the regularizers in my approach. Moreover, my approach works for any arbitrary time step, whereas, the method in Hansen et al. (2018) suffers from the problem of using appropriate guess for the time step and space dimension step, as they are tightly linked to guarantee convergence. Experimentation shows that for problems with capital misallocation, the time step that needs to be set is very small which significantly decreases the speed of convergence. D'Avernas and Vandeweyer (2019) explores a similar macro-finance model and demonstrate the difficulty in maintaining the monotonicity of finite difference schemes in solving PDEs, and offer a robust solution technique based on Bonnans et al. (2004). The technique involves solving for the right direction to approximate the finite differences so as to preserve the monotonicity which is a necessary condition for convergence. However, the approach that they offer is specific to two space dimensions. On the contrary, ALIENs are easily scalable with minimal effort in coding.

The second related strand of literature is application of machine learning to solve equilibrium

models in economics and finance. The papers that are closer are Duarte (2017) and Fernández-Villaverde et al. (2020) since they also consider equilibrium problems in continuous time. Duarte (2017) encodes policy and value functions with neural networks and performs policy evaluation and policy update in the spirit of reinforcement learning. While Duarte (2017) focusses on representative agent asset pricing models that admit analytical solutions, the problems that I consider are heterogeneous agent macro-finance models with endogenous state variables and jumps in policy functions. These problems do not have closed form solutions and are more complex to solve. Fernández-Villaverde et al. (2020) also employs neural networks to solve a model based on Krusell and Smith (1998) in continuous time. However, it is the law of motion of the aggregate wealth that is solved using neural networks. The value function in their model is solved using traditional finite difference method which is in contrast to ALIENs. Azinovic et al. (2019) considers discrete time economic problems and solves for the equilibrium policy functions using neural networks. They also parallelize their algorithm using Horovod to speed up the performance. While their method is based on simulating from the ergodic density, I consider sparse grid padded with active points for training.

There is a substantial literature in computational physics and applied mathematics to approximate PDEs and HJBs using neural network, starting from Sirignano and Spiliopoulos (2018a) and Raissi et al. (2017). Sirignano and Spiliopoulos (2018a) proposes deep galerkin method to solve PDEs in high dimensions and incorporates monte carlo methods to compute second order derivatives to speed up computation. Raissi et al. (2017) solves canonical two space dimensional problems in computational physics such as Naiver-Stokes and Burgers equations. Han et al. (2018) represents quasi-linear PDEs in the form of forward backward stochastic differential equations and then applies deep neural networks to solve the PDEs. They find that this strategy enables efficient computation of gradients in terms of speed an accuracy. While Raissi et al. (2017), Raissi et al. (2019) etc. use feed-forward neural networks, Sirignano and Spiliopoulos (2018a) finds that advanced architectures like LSTMs offer improved performance. More novel architectures like convolution networks, which are mostly used in image recognition, have also been found to be useful in solving PDEs (Tompson et al. (2017)). In contrast to these papers, the framework that I propose is suited to solve problems in economic and finance. For instance, many problems in macro-finance and asset pricing come with endogenous state variables that are often correlated- a feature that the models in applied mathematics does not typically deal with. Moreover, the non-linearity of the PDEs in the economic models come from the fact that the advection, diffusion, linear, and cross term coefficients of the PDE are endogenously dependent on the equilibrium policy functions. For example, these coefficients in Brunnermeier and Sannikov (2016b) are solved for using a separate Newton-Raphson method since one needs to solve for an algebraic first order differential equation to get these coefficients. These kinds of complications do not arise in the models considered in applied mathematics where more often the coefficients are known constants or simple exogenous functions.[3] Thus, one cannot simply apply the deep learning

[3]The nature of coefficients play an important role in finite difference methods. For example, see Brunner-

tools developed in other areas and hope to solve models in financial economics. Lastly, the usage of sparse and active points in this chapter relates to the literature on adaptive sparse grids that are concerned with a systematic way of generating the state space grid. For example, Brumm and Scheidegger (2017) uses adaptive sparse grids to solve dynamic economic models, whereas Bungartz et al. (2012) solves option pricing models using finite element method.[4]

2.3 General Set-up

In this section, I present a general set-up of a portfolio choice problem with a continuum of agents indexed by j, who have a lifetime recursive utility given by

$$U_{j,t} = E_t\Big[\int_t^\infty f(c_{j,s}, U_{j,s})ds\Big] \tag{2.1}$$

with

$$f(c_{j,t}, U_{j,t}) = \frac{1-\gamma}{1-\frac{1}{\varrho}} U_{j,t}\left[\left(\frac{c_{j,t}}{\big((1-\gamma)U_{j,t}\big)^{1/(1-\gamma)}}\right)^{1-\frac{1}{\varrho}} - \rho\right] \tag{2.2}$$

where ρ, γ, and ϱ are the discount rate, the risk aversion, and the inter-temporal elasticity of substitution (IES) of the agents. I assume that these parameters are the same for all agents in the economy but this is purely for simplicity and can be easily relaxed to introduce further heterogeneity. The agents trade a risky asset, a claim to the dividend denoted by y_t, where $t \in [0,\infty]$, that follows an exogenous process

$$\frac{dy_t}{y_t} = g\,dt + \sigma\,dZ_t \tag{2.3}$$

where Z_t is the standard Brownian motion representing the aggregate uncertainty in $\mathscr{F}_t$, g is the growth rate, and σ is the volatility. The agents also trade in a risk-free security that pays a return r_t that will be determined in the equilibrium. The price of the risky asset q_t is governed by

$$\frac{dq_t}{q_t} = \mu_t\,dt + \sigma_t^R\,dZ_t \tag{2.4}$$

where the drift μ_t and the volatility σ_t^R are endogenous objects to be determined in the equilibrium. The return on the risky asset is given by $dR_t = \frac{a_j(y_t)}{q_t}dt + \frac{dq_t}{q_t}$, where a_j is an agent specific dividend. The net worth of the agents evolve as

$$\frac{dw_{j,t}}{w_{j,t}} = (r_t + \theta_{j,t}\zeta_{j,t} - \hat{c}_{j,t})dt + \theta_{j,t}\sigma_t^R\,dZ_t \tag{2.5}$$

meier and Sannikov (2016b), who use an up-winding scheme in order to deal with the case where the advection coefficients have different signs in different parts of the state space.

[4] These papers propose a systematic way of refining grid, for example by modifying the basis functions as in Brumm and Scheidegger (2017). On the contrary, I use a random sample from the mesh that are padded with active points for training the neural network.

where r_t is the risk-free rate, σ_t^R is the volatility of the return on the risky asset, $\theta_{j,t}$ is the portfolio weight on the risky asset, and $\hat{c}_{j,t}$ is the consumption-wealth ratio. The quantity $a_j(y_t)$ is an agent specific function of the dividends,[5] and $\zeta_{j,t}$ is the price of risk, which may differ across agents due to different dividends. The agents maximize the utility (2.1) subject to the wealth dynamics (2.5). Shorting of the capital by the agents is disallowed. The HJB equation for the optimization problem can be written as

$$\sup_{\hat{c}_{j,t}\theta_{j,t}} f(c_{j,t}, U_{j,t}) + E_t(dU_{j,t}) = 0 \tag{2.6}$$

Due to homothetic preferences, the value function takes the form

$$U_{j,t} = \frac{(J_{j,t} w_{j,t})^{1-\gamma}}{1-\gamma} \tag{2.7}$$

where the stochastic opportunity set $J_{j,t}$ follows the process

$$\frac{dJ_{j,t}}{J_{j,t}} = \mu^J_{j,t} dt + \sigma^J_{j,t} dZ_t \tag{2.8}$$

The equilibrium objects $(\mu^J_{j,t}, \sigma^J_{j,t})$ need to be solved. Applying Ito's lemma to $J_{j,t}$, the HJB equation can be written as

$$\begin{aligned}\frac{\rho}{1-1/\varrho} = \sup_{\hat{c}_{j,t},\theta_{j,t}} & \frac{\hat{c}_{j,t}^{1-1/\varrho}}{1-1/\varrho}\rho J_{j,t}^{1/\varrho-1} + (r_t + \theta_{j,t}\zeta_{j,t} - \hat{c}_{j,t}) \\ & + \mu^J_{j,t} - \frac{\gamma}{2}(\theta^2_{j,t}(\sigma^R_t)^2 + (\sigma^J_{j,t})^2)) - (1-\gamma)\theta_{j,t}\sigma^R_t\sigma^J_{j,t}\end{aligned} \tag{2.9}$$

The optimal quantities $(\hat{c}_{j,t}, \theta_{j,t})$ can be found by maximizing the HJB equation. It then remains to solve for the function $J_{j,t}$ which depends on the state variables $\boldsymbol{x} \in \Omega$. Assuming that the number of state variables is d, applying Ito's lemma to $J_{j,t}$, and equating the drift terms,[6] we have

$$\mu^J J = \sum_{i=1}^{d} \mu^{x_i}(\boldsymbol{x}) \frac{\partial J}{\partial x_i} + \sum_{i,j=1}^{d} b^{i,j}(\boldsymbol{x}) \frac{\partial^2 J}{\partial x_i \partial x_j} \tag{2.10}$$

where $\mu^{x_i}(\boldsymbol{x})$ is the drift of the state variable $x_i \in \boldsymbol{x}$ and $b^{i,j}(\boldsymbol{x}) = \frac{1}{2}\sigma^{x_i}(\boldsymbol{x})\sigma^{x_j}(\boldsymbol{x})$ is the scaled product of volatility of the state variables $\{x_i, x_j\} \in \boldsymbol{x}$. The state variables can be endogenous in which case their drift and the volatility may depend on J. Moreover, the PDE (2.10) is non-linear because the term μ^J depends on J in a highly non-linear fashion. To see that, note that the function μ^J can be obtained from the HJB equation (2.9) after plugging in the optimal

[5] For instance, agents may have a different tax treatment of dividends in which case the net dividend earned will depend on the investor type. In a production economy, some agents may have a lower productivity rate which gives them a lower dividend like in Brunnermeier and Sannikov (2016b).

[6] I drop the agent and the time index in order to avoid cluttering of notations.

$\hat{c}_{j,t}, \theta_{j,t}$. That is, we have

$$\mu^J = \frac{\rho}{1-1/\varrho} - \frac{\hat{c}_{j,t}^{*1-1/\varrho}}{1-1/\varrho}\rho J_{j,t}^{1/\varrho-1} - (r_t + \theta_{j,t}^* \zeta_{j,t} - \hat{c}_{j,t}^*) + \mu_{j,t}^J - \frac{\gamma}{2}(\theta_{j,t}^{*2}(\sigma_t^R)^2 + (\sigma_{j,t}^J)^2)) - (1-\gamma)\theta_{j,t}^* \sigma_t^R \sigma_{j,t}^J \tag{2.11}$$

where $\hat{c}^*$ and θ^* are the optimal consumption-wealth ratio and portfolio choice respectively. Non-linear PDEs are in general difficult to solve and the literature addresses this issue by converting the equation (2.10) into a quasi-linear parabolic PDE by introducing an artificial time derivative. That is, the function μ^J is assumed to be a coefficient whose value is computed based on the value of J from the previous time step. This is similar to Brunnermeier and Sannikov (2016b), who solve for the equilibrium quantities in the static inner loop given the value function, which then gets updated in the outer time loop. This also means that μ^{x_i} and σ^{x_i} do not depend on J and t but only depend on the equilibrium quantities and the state variables themselves. This two-step method, which is standard in the literature, has relevance to the reinforcement learning where the inner static step is called as 'policy evaluation', and the outer time step is called as 'policy update'. Following the tradition, I add a false time derivative to the PDE (2.10) and rewrite it in a general quasi-linear parabolic form at the k-th time iteration as

$$\mathscr{G}[J] := \frac{\partial J}{\partial t} + A\left(\boldsymbol{x}, J, \frac{\partial J}{\partial \boldsymbol{x}}\right) + \frac{1}{2} tr\left[B\left(\boldsymbol{x}, \tilde{J}, \frac{\partial \tilde{J}}{\partial \boldsymbol{x}}\right)\frac{\partial^2 J}{\partial \boldsymbol{x}^2} B\left(\boldsymbol{x}, \tilde{J}, \frac{\partial \tilde{J}}{\partial \boldsymbol{x}}\right)^T\right] = 0 \tag{2.12}$$

$$(t, \boldsymbol{x}) \in [T - k\Delta t, T - (k-1)\Delta t] \times \Omega \tag{2.13}$$

with the boundary conditions

$$J(t, \boldsymbol{x}) = \tilde{J}; \quad \forall (t, \boldsymbol{x}) \in (T - (k-1)\Delta t) \times \Omega \tag{2.14}$$

$$\frac{\partial J(t, \boldsymbol{x})}{\partial \boldsymbol{x}} = 0; \quad \forall (t, \boldsymbol{x}) \in (T - (k-1)\Delta t) \times \partial\Omega \tag{2.15}$$

where in this case, $A\left(\boldsymbol{x}, J, \frac{\partial J}{\partial \boldsymbol{x}}\right) = \sum_{i=1}^d \mu^{x_i}(\tilde{J}) \frac{\partial J}{\partial x_i} - \mu^J J$ and $B\left(\boldsymbol{x}, \tilde{J}, \frac{\partial J}{\partial \boldsymbol{x}}\right) = \sigma^{x_i}(\tilde{J})$. The first boundary condition contains $\tilde{J}$ which is the value obtained in previous time iteration. For example, in kth time step, we solve for the function $J(T - k\Delta t, \boldsymbol{x})$ and $\tilde{J}$ in this case denotes the value $J(T - (k-1)\Delta t, \boldsymbol{x})$. The PDE (2.12) occurs widely in the macro-finance and asset pricing literature including Brunnermeier and Sannikov (2016b), Hansen et al. (2018), Kurlat (2018), Di Tella (2017), Di Tella (2019), Drechsler et al. (2018), Gomez (2019), Krishnamurthy and Li (2020), Li (2020), He and Krishnamurthy (2013), D'Avernas et al. (2019), among others. These papers address different economic problems in varied ways but all of them boil down to solving the PDE (2.12) subjected to some boundary conditions. In general, a closed-form solution to such PDEs do not exist. The literature has so far used finite difference method with up-winding, which works well in smaller dimensions but becomes infeasible as d grows large due to the curse of dimensionality and the difficulty in preserving the monotonicity of the

finite difference scheme.

2.3.1 Neural network for PDEs

The approximation of function J using a feed-forward deep neural network is done by random sampling of points from the state space making it a mesh-free procedure. This empowers the approach with scalability by alleviating the curse of dimensionality in the PDE time step. Moreover, since the neural network can approximate any type of quasi-linear parabolic PDE, one doesn't need to worry about approximating the derivatives in the right spatial dimensions to preserve monotonicity as in the grid-based finite difference method (D'Avernas and Vandeweyer (2019)).

Feed-forward network:I present a brief introduction to the simple feed-forward neural networks from which the informed neural network is built to solve the PDE (2.12). A single-layer neural network that can approximate J is given by

$$\hat{\boldsymbol{J}}(\boldsymbol{x}|\Theta) = g\Big(\boldsymbol{W}\boldsymbol{x} + \boldsymbol{b}\Big) \tag{2.16}$$

where $\boldsymbol{W}, \boldsymbol{b} \in \Theta$ are the parameters called *weights* and *biases* respectively,[7] and $g(\cdot)$ is the *activation function* which maps the input to the output in a non-linear fashion. The universal approximation theorem (Hornik (1991)) states that any Borel measurable function can be approximated by a feed-forward neural network with a single hidden layer. That is, for any $\epsilon > 0$ and any function $\boldsymbol{J(x)}$ with state variables $\boldsymbol{x} \in I_d$, where I_d is the d-dimensional unit hypercube, the approximation (2.16) satisfies[8]

$$|\hat{\boldsymbol{J}}(\boldsymbol{x}|\boldsymbol{\Theta}) - \boldsymbol{J}(\boldsymbol{x})| < \epsilon \qquad \forall \boldsymbol{x} \in I_d$$

The activation functions can be thought of basis functions but they have very simple functional forms as opposed to complex forms such as Chebychev and Legendre polynomials that are commonly used in the projection methods. One minor shortcoming of the universal approximation theorem is that it does not specify the exact number of neurons required to achieve convergence. That is guided entirely by the empirical procedure. It turns out that a single hidden layer network may not provide a good approximation for the function J and hence it requires to stack multiple hidden layers on top of one another to construct a deep

[7] *Bias* is a terrible terminology that is unfortunately commonplace in the deep learning literature. We can think of b more as a shift parameter.

[8] This activation function $g(\cdot)$ is not dependent on J and has a natural relation with the Galerkin method if we compare (2.16) and (2.33).

neural network. Then, we have

$$\begin{aligned}
\boldsymbol{z_1} &= \boldsymbol{W_1 x} + \boldsymbol{b_1} \\
\boldsymbol{h_1} &= g(\boldsymbol{z_1}) \\
\boldsymbol{z_2} &= \boldsymbol{W_2 h_1} + \boldsymbol{b_2} \\
\boldsymbol{h_2} &= g(\boldsymbol{z_2}) \\
&\vdots \\
\boldsymbol{z_l} &= \boldsymbol{W_l h_{l-1}} + \boldsymbol{b_l} \\
\boldsymbol{\hat{J}} &= \sigma(\boldsymbol{W_{l+1} z_l} + \boldsymbol{b_{l+1}})
\end{aligned}$$

where l is the number of hidden layers and $g(\cdot)$ are the activation functions that remain the same in each hidden layer. The commonly used activation functions in deep-learning literature are Rectified Linear Unit (ReLu), tanh, sigmoid, and shifted-ReLu. I consider tanh activation function for the hidden layers and a sigmoid activation function in the output layer based on superior performance for solving the PDEs. The activation function in hidden layer takes the form

$$g(z) = tanh(z) = \frac{e^z - e^{-z}}{e^z + e^{-z}} \tag{2.17}$$

The sigmoid function is given by

$$\sigma(x) = \frac{1}{1 + e^{-x}} \tag{2.18}$$

Note that the sigmoid function gives values in the range $(0,1)$. This works for the problems considered in this chapter since the stochastic opportunity set J in equilibrium is equal to the consumption-wealth ratio, which lies in the range between 0 and 1. If the opportunity set takes values in the range $(-\infty, +\infty)$, then a linear output layer is recommended. The output from the feed-forward deep neural network $\hat{J}(x|\Theta)$ forms the basis for solving the equation (2.12).

Informed neural nets: The approximation from simple feed-forward network will clearly be poor because it does not encode any information from the PDE. The next logical step is to transform the simple feed-forward network into a more informed network by encoding the economic information into it. I build customized loss-functions that act as regularizers in the neural network optimization. Consider the PDE residual from[9] (2.12)

$$f := \frac{\partial \hat{J}(\Theta)}{\partial t} + A\left(\boldsymbol{x}, \hat{J}(\Theta), \frac{\partial \hat{J}(\Theta)}{\partial \boldsymbol{x}}\right) + \frac{1}{2} tr\left[B\left(\boldsymbol{x}, \tilde{J}, \frac{\partial \tilde{J}}{\partial \boldsymbol{x}}\right) \frac{\partial^2 \hat{J}(\Theta)}{\partial \boldsymbol{x^2}} B\left(\boldsymbol{x}, \tilde{J}, \frac{\partial \tilde{J}}{\partial \boldsymbol{x}}\right)^T\right] \tag{2.19}$$

[9] I denote $\hat{J}(t, \boldsymbol{x} \mid \Theta)$ as $\hat{J}(\Theta)$ for brevity.

where $\tilde{J}$ denotes the value obtained from the previous time iteration.[10] Starting from a simple feed-forward neural network $\hat{J}(\boldsymbol{x}|\Theta)$ that is parameterized by an arbitrary Θ, the goal is to find the optimal Θ^* that ensures $|\hat{J}(t,\boldsymbol{x}|\Theta) - J| < \epsilon$ for all $\epsilon > 0$. Towards this goal, I minimize the loss function that encodes the economic information from the PDE and the inner static loop. The loss function is given by

$$\mathscr{L} = \lambda_f \mathscr{L}_f + \lambda_j \mathscr{L}_j + \lambda_b \mathscr{L}_b + \lambda_c^1 \mathscr{L}_c^1 + \lambda_c^2 \mathscr{L}_c^2 \tag{2.20}$$

where

$$\text{PDE loss} \qquad \mathscr{L}_f = \frac{1}{N_f} \sum_{i=1}^{N_f} |f(\boldsymbol{x}_f^i, t_f^i)|^2 \tag{2.21}$$

$$\text{Bounding loss-1} \qquad \mathscr{L}^j = \frac{1}{N_j} \sum_{i=1}^{N_j} |\hat{J}(\boldsymbol{x}_j^i, t_j^i) - \tilde{J}^i|^2 \tag{2.22}$$

$$\text{Bounding loss-2} \qquad \mathscr{L}^b = \frac{1}{N_b} \sum_{i=1}^{N_b} \left| \frac{\partial \hat{J}}{\partial \boldsymbol{x}_b^i} \right|^2 \tag{2.23}$$

$$\text{Active loss-1} \qquad \mathscr{L}_c^1 = \frac{1}{N_c} \sum_{i=1}^{N_c} |\hat{J}(\boldsymbol{x}_c^i, t_c^i) - \tilde{J}^i|^2 \tag{2.24}$$

$$\text{Active loss-2} \qquad \mathscr{L}_c^2 = \frac{1}{N_c} \sum_{i=1}^{N_c} |f(\boldsymbol{x}_c^i, t_c^i)|^2 \tag{2.25}$$

The points $(\boldsymbol{x}_j^i, \tilde{J}^i)_{i=1}^{N_j}$ and $(\boldsymbol{x}_j^i)_{i=1}^{N_b}$ denote the training sample from the two boundary conditions, $(\boldsymbol{x}_c^i, \tilde{J}^i)_{i=1}^{N_c}$ are the training sample from the active loss region, and $(\boldsymbol{x}_f^i)_{i=1}^{N_f}$, and $(\boldsymbol{x}_c^i)_{i=1}^{N_c}$ correspond to the training samples used to compute the PDE residuals. The loss function is customized to take into account the economic problem at hand. Figure (2.1) presents the architecture of this network. In the appendix (A.2), I prove a convergence theorem along the lines of Sirignano and Spiliopoulos (2018b), related to the neural network approximation of quasi-linear parabolic PDEs of the form (2.12). The L^2 loss $\mathscr{L}$ from (2.20) is a measure of how well the neural network $\hat{J}(t,\boldsymbol{x}|\Theta)$ approximates the function J that solves the PDE (2.12). The goal is to make this loss close to zero.

Theorem 2.3.1. Assume that there exists a unique solution to the PDE (2.19). Suppose we have a generic, infinitely wide neural network that is trained by a gradient decent algorithm to approximate J. Then, this solution also approximately solves the PDE.

Proof: See Appendix A.2.

[10] Note that the volatility and advection coefficient terms are computed based on the value of J from the previous time iteration. Thus, they are simply considered as coefficients.

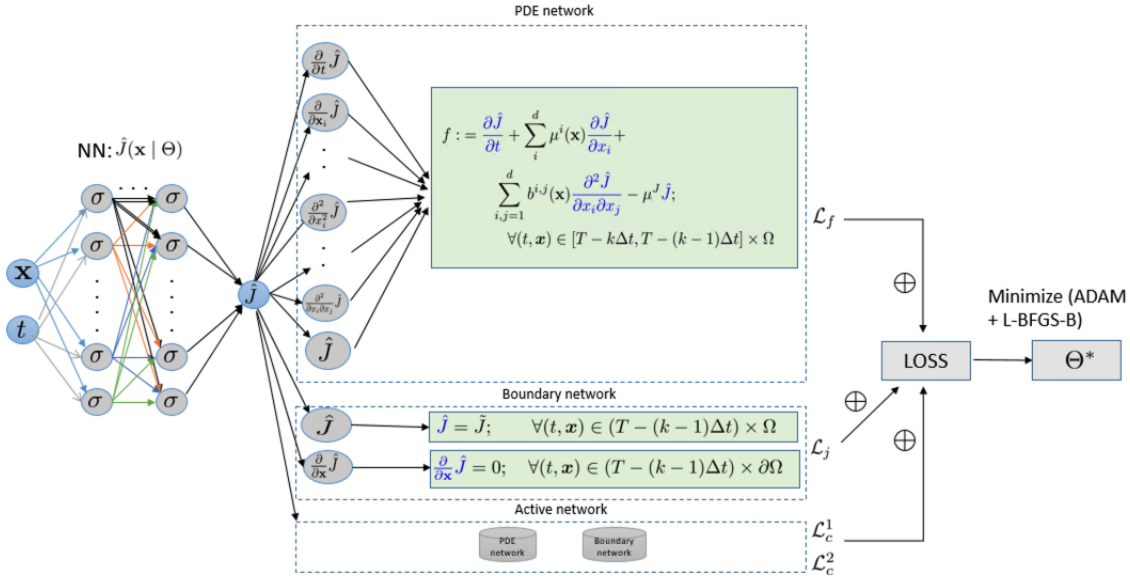

Figure 2.1 – Network architecture

Under the hood

The training samples in this method come entirely from the boundary conditions. Unlike the traditional machine learning paradigm where the model is trained on the training set, cross-validated on the validation test, and tested for accuracy on the test set, the method used in this chapter does not really have training samples in the same sense. The inputs from boundary conditions can be thought of as pseudo-training samples that are used to approximate the function J. As a result, the problem of overfitting, which is ubiquitous in the machine learning paradigm is less relevant here. Moreover, only a fraction of data points are sampled in each PDE step iteration and hence any minor concerns of overfitting is eliminated.[11] The feed-forward architecture shown in Figure (2.1) is simple compared to the more complex ones such as GANs, Autoencoders, and LSTMs used in the asset pricing literature (see Gu et al. (2020), Chen et al. (2019) etc). This raises the question what makes the algorithm succeed in learning the function J in a high dimensional space. The answer lies in the encoding of economic information through customized loss functions that makes the simple feed-forward network more informed. The PDE loss function dictates the neural networks to satisfy the HJB equation such that the residual from the HJB is close to zero in a mean-squared sense. At the same time, the neural networks are also forced to obey the initial/boundary conditions through the bounding loss functions. The goal is then to optimize for the parameters Θ such that the neural network approximation $\hat{J}(\boldsymbol{x}|\Theta)$ minimizes the total loss function. In this optimization problem, the customized loss functions act as *regularizers*.

[11] This is similar to using mini-batches in deep learning to reduce overfitting.

Automatic differentiation: The success of deep learning is largely due to automatic differentiation, which is a computationally efficient way to compute derivatives of potentially non-linear functions. The neural network approximator works by receiving the input $\boldsymbol{x}$ and computing the output $\hat{J}(\boldsymbol{x}|\Theta)$ that is a composition of simple functions. The derivative of $\hat{J}$ with respect to the inputs $\boldsymbol{x}$ can be obtained analytically by repeated applications of the chain rule. The backpropogation algorithm traverses the graph, computes the derivatives of symbolic variables and stores these operations into new nodes in the graph for later use. To compute a higher order derivative, one can simply run the backpropogation again through the extended graph and obtain it easily. While the deep learning literature uses the automatic differentiation in computing the derivatives of loss function with respect to the parameters such as weights and biases, I use it explicitly to take derivatives with respect to the space and time dimensions. This is illustrated in the PDE network in Figure (2.1). The derivatives of function $\hat{J}$ with respect to each space and time dimensions are stored as separate nodes in the graph, which are used to compute higher order derivatives through backpropagation. Formally, the cost of computing $\frac{\partial J}{\partial \boldsymbol{x}}$ is $\mathcal{O}(d) \times \text{cost}(J)$, which is the same as the cost of computing $\frac{\partial^2 J}{\partial^2 \boldsymbol{x}}$ since the backpropagation takes advantage of the first order derivatives stored in the computational graph when applying the chain rule. Thus, the explicit use of automatic differentiation to compute derivatives with respect to the space dimension allows fast computation even in high dimensions. On the contrary, computing such higher order derivatives bears a cost $\mathcal{O}(d^2) \times \text{cost}(J)$ in finite-difference methods, imposing a bottleneck when d is potentially large. Lastly, the separation of the fundamental neural network and the more informed PDE and bounding network allows us to witness the automatic differentiation fully at action, which is the main driver of the learning process.

Optimization: The total loss function is the weighted sum of the loss from PDE, boundary, and crisis network. I use a combination of adaptive momentum (ADAM) and Limited memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS-B) optimizers to solve for Θ^*. While the fundamental network has only four hidden layers, the PDE network adds more layers since it involves gradient computations with respect to the state variables. Hence, the overall network is quite deep, with a highly complex and non-convex loss function. While ADAM optimizer, which is the standard algorithm in deep learning models, is based on first order derivative, L-BFGS-B is a second order method and is empirically found to be effective in solving the problems considered in this chapter. The implementation details are relegated to Appendix A.2.4. The generic algorithm of ALIENs is given in the pseudo-code (1).

Active learning: When the equilibrium policy functions have stark non-linearities, a smart sampling procedure is required to ensure faster convergence. In the case of a regime changing policy function, a subdomain in the state space captures the region where the regime shifts. In such cases, it is crucial to have a better approximation of equilibrium functions in this subdomain so as to avoid instabilities in the future time iterations. If the modeler possesses a prior knowledge of the location of this subdomain, then one can obtain a better approximation

Algorithm 1

procedure ALIENs
  Initialize function J, set tolerance level tol
  while $error > tol$ **do**
    Compute policy functions taking as given J
    Construct a neural network approximation $\hat{J}(\Theta)$
    Construct boundary training points
    Construct active training points
    Compute aggregated loss $\mathscr{L}$
    for epoch=1 to maxEpochs **do**
      Create a mini-batch from training points ▷ Optional
      Minimize loss using ADAM optimizer, ▷ Learning rate = 0.001
    Minimize loss using L-BFGS-F until convergence, update Θ^*
    Compute $error := \max|\hat{J}(\Theta^*) - J|$
    Update $J \leftarrow \hat{J}(\Theta^*)$
    Stop if $error < tol$

by sampling sufficiently large number of training points from this subdomain. However, this information is often not known prior and only gets revealed after the equilibrium policy functions are solved in the first time iteration. Moreover, this subdomain may change from one iteration to next as the policies are evaluated and updated. To tackle this problem, I endow the neural network with an active learning which tracks this subdomain in each iteration by inspecting the equilibrium policy functions. The neural network takes advantage of the fact that the static step precedes the outer time step in an iterative fashion, and actively seeks information about the subdomain at every iteration. Once this subdomain is revealed, active training points are created through random sampling of grid points from this subdomain to construct the loss functions $\mathscr{L}_c^1$ and $\mathscr{L}_c^2$ in (2.20). The practice of seeking new data points to label in supervised learning is called as *active machine learning*, a budding area in the artificial intelligence literature where the model seeks out new inputs to improve the learning in subsequent iterations. In the context of reinforcement learning, the model learns about new state space by actively interacting with the environment. In the models considered in this chapter, the state space is well defined but new grid points are sampled to be used for training the neural network. In doing so, the regularizers $\mathscr{L}_c^1$ and $\mathscr{L}_c^2$ forces the model to learn better in the subdomain since it is costly to have errors in regions of regime shift as they get propagated and amplified in the subsequent iterations. The proposed active learning algorithm alleviates these problems and ensures convergence. The relative weights $(\lambda_f, \lambda_j, \lambda_b, \lambda_c^1, \lambda_c^2)$ control for the importance that each of the components carries in the aggregated loss function.

Let us take a concrete example to understand how active learning works. Consider a two space-dimensional model that generates a capital price as shown in Figure (2.2). For now, we can abstract away from the details of the model which will be explained in Section 2.5 and focus on how active learning is applied. The two state variables are the wealth share (z) and the productivity (a) of expert sector in the economy. The model features non-linear equilibrium

policy where the capital price is high and almost linear when the wealth share of experts is high, and falls as the wealth share drops. The full state space grid shown in Fig (2.2) contains tightly packed points in the wealth share dimension since it is the primary driver of price dynamics in the model. The panels (*c*) and (*d*) displays the sparse grid used in training the neural networks. The blue points span throughout the state space while the green points span the neighborhood of sharp transition. One can notice that this neighborhood changes from iteration 1 in panel (*c*) to iteration 15 in panel (*d*). As non-linearities are slowly introduced into the parabolic PDEs, the dynamics of the equilibrium prices change and the non-linearities in policy functions occur in different parts of the state space. Not only that, the shape of the subdomain may also change, as it does in the example demonstrated. By actively tracking the subdomain and sampling points from the neighbourhood of this subdomain, the sparse training sample is guaranteed to have points from the state space that matter the most to avoid instabilities in future iterations. While it it not costly to use the entire grid when the dimension is small, such a smart sampling method is required in higher dimensions where it is infeasible to use the entire grid to train the neural network.

Distributed learning: The deep learning model in this chapter is much simpler compared to the state-of-the-art models such as GPT-3 model[12]. Having a lighter model allows us to utilize data parallelism instead of model parallelism, where in the former case, the chief worker[13] in the cluster splits the data and distributes it to several workers that share the same model. Once training is completed, the result is averaged across workers and sent to the chief worker that updates the model with the optimal parameters. In cases where the model is too large to store in a single worker, model parallelism is applicable where each worker trains a piece of the model but with the same data. One challenge in data parallelism is that since each worker shares the same model, the parameters need to be synchronized. I utilize Horovod[14], a popular distributed machine learning framework that manages the internal working of gradient aggregation across the workers through *ring-AllReduce* operation. In some sense, Horovod works like a wrapper around the Message Passing Interface that enables communication across different workers. While data parallelism employed in this chapter speeds up computation significantly, it also lends itself nicely to another advantage. As the academic profession moves more towards an open source style research, the data parallelism allows the users to employ big data on models developed by other users. Often, quantitative analysis of economic models requires extensive simulation studies to dissect and understand the underlying mechanisms of the model.[15] Appendix A.2.4 provides the details of Horovod

[12] This is a language model by Brown et al. (2020) from OpenAI that recently gained substantial attention. The model has a total of 175 billion parameters to train.

[13] A worker refers to a node in the cluster that is made up of a number of computer nodes.

[14] See Sergeev and Balso (2018) for details.

[15] See, for example, Gopalakrishna (2020) who uses simulation studies to perform dissection of a macro-finance model.

along with sample code to demonstrate the simplicity in usage and deployment.

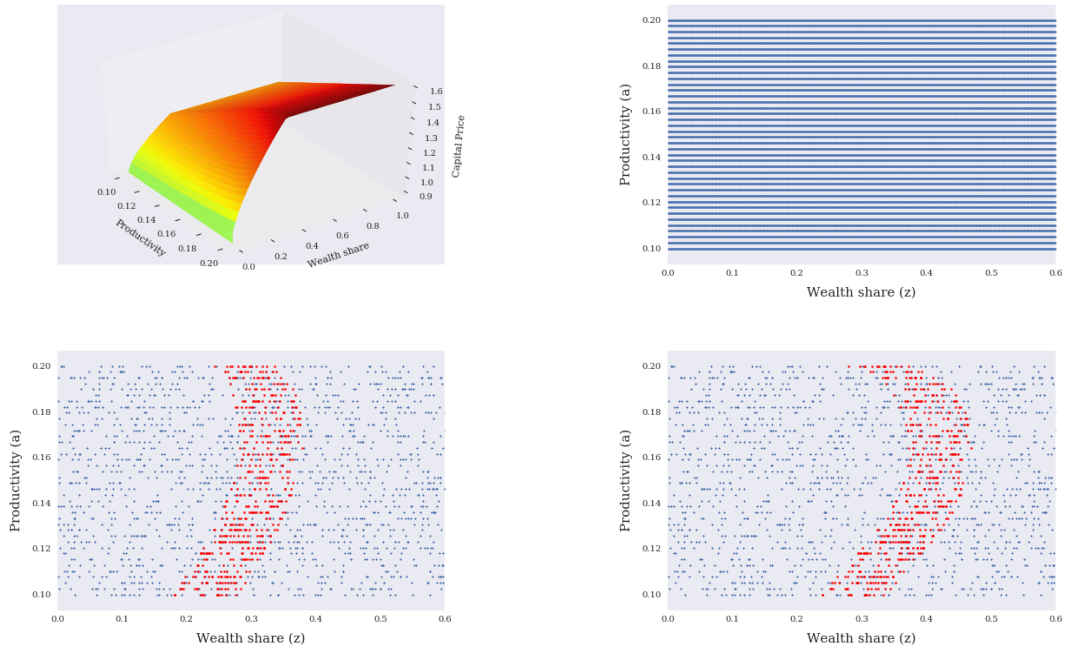

Figure 2.2 – Panel (a) shows Equilibrium capital price. Panel (b) shows full grid with 1000 points in wealth share and 40 points productivity dimension respectively. Panels (c) and (d) shows sparse grid used for learning where points in green represent neighborhood of regime change in the 1^{st} and 15^{th} iteration respectively.

2.4 Benchmark Model

In this section, I consider a one space-dimensional model based on Brunnermeier and Sannikov (2016b) augmented with recursive preferences that will serve as the benchmark. I solve this model using the proposed neural network method with active learning and show that the solution matches the numerical solution obtained from an up-winding finite difference scheme. Since the closed-form solution to the benchmark model is not available, the comparison of the neural network solution is made against the finite difference solution.

2.4.1 Model

I consider a heterogeneous agent economy populated by households (h) and experts (e) who form a set $\mathbb{H}$ and $\mathbb{E}$ respectively. The aggregate capital in the economy is denoted by K_t where $t \in [0,\infty)$ denotes time. Due to homogeneity of preferences, which will be explained later, we can consider a representative household h and expert e for the rest of the model. Both the households and the experts are allowed to hold capital with a no shorting constraint but the households obtain a lower productivity rate (a_h) compared to the experts (a_e) from investing

in the capital. The constraint in the model is such that experts have to retain at least a fraction of the equity in their balance sheet. This is the skin-in-the game constraint that precludes the economy from achieving perfect risk-sharing. The production technology is given by

$$y_{j,t} = a_j k_{j,t}, \qquad j \in \{e, h\}$$

where $k_{j,t}$ is the capital held by agent j whose process is governed by

$$\frac{dk_{j,t}}{k_{j,t}} = (\Phi(\iota_{j,t}) - \delta)dt + \sigma dZ_t^k \tag{2.26}$$

where $\iota_{j,t}$ is the investment rate, δ is the depreciation rate of capital, σ is the volatility of the capital, and $\{Z_t^k \in \mathbb{R}; \mathscr{F}_t, \Omega\}$ are the standard Brownian motions representing the aggregate uncertainty in $(\Omega, \mathbb{P}, \mathscr{F})$. The quantity $\Phi(\cdot)$ is the investment cost function that is concave and has decreasing returns to scale. The aggregate capital in the economy is denoted by K_t. That is, $K_t = \int_{\mathbb{E}} k_{j,t} dj + \int_{\mathbb{H}} k_{j,t} dj$.

Preferences: Each agent has a continuous-time Duffie-Epstein utility given by

$$U_{j,t} = E_t\left[\int_t^\infty f(c_{j,s}, U_{j,s})ds\right] \tag{2.27}$$

where the aggregator takes the form

$$f(c_{j,s}, U_{j,s}) = (1-\gamma)\rho U_{j,t}\left(\log(c_{j,t}) - \frac{1}{1-\gamma}\log\Big((1-\gamma)U_{j,t}\Big)\right) \tag{2.28}$$

The parameter γ captures the risk aversion of agent j, and the inter-temporal elasticity of substitution (IES) is assumed to be 1. The aggregator (2.28) is a limiting case of (A.22).[16] I relegate rest of the model and equilibrium computation to Appendix A.2.3 and go straight to the solution technique. Given the homotheticity of preferences, we can conjecture that the value function is of the form

$$U_{j,t} = \frac{(J_{j,t}K_t)^{1-\gamma}}{1-\gamma} \tag{2.29}$$

where K_t is the aggregate capital, and the stochastic opportunity set $J_{j,t}$ follows the SDE

$$\frac{dJ_{j,t}}{J_{j,t}} = \mu_{j,t}^J dt + \sigma_{j,t}^J dZ_t^k \tag{2.30}$$

where the equilibrium objects $\mu_{j,t}^J$ and $\sigma_{j,t}^J$ will have to be determined. The endogenous state variable of the model is the wealth share of experts defined by

$$z_t = \frac{W_{e,t}}{q_t K_t}$$

[16] This is for simplicity and can be easily relaxed to the case of non-unitary IES.

where $W_{e,t}$ is the aggregated wealth of all experts in the economy. For convenience, the utility (2.29) is written as a function of capital K_t but they could also be written as a function of $W_{j,t}$. By applying Ito's lemma to $J_{j,t}(z_t)$ and equating the drift terms, we have

$$\mu_{j,t}J_{j,t} = \frac{dJ_{j,t}}{dz_t}\mu_t^z + \frac{1}{2}\frac{d^2J_{j,t}}{dz_t^2}(\sigma_t^z)^2 \tag{2.31}$$

where the function $\mu_{j,t}^J$ can be obtained from the HJB equation similar to the general setup in Section 2.3. The solution method involves solving for the equilibrium quantities $(\chi_t, k_{j,t}, q_t, \sigma_t^q)$ using a Newton-Rhaphson method, similar to Brunnermeier and Sannikov (2016b), and then updating the value function $J_{j,t}$ which requires solving (2.31). This equation, when augmented with an artificial time derivative, resembles the quasi-linear parabolic PDE introduced in (2.12). I relegate the details of the Newton-Rhaphson method used to solve for equilibrium quantities to the Appendix and present the methodology to solving (2.31) in the main text.

2.4.2 Traditional methods

The literature has used finite difference method extensively (see Brunnermeier and Sannikov (2016b), Di Tella (2017), Di Tella (2019), Gomez (2019), Hansen et al. (2018), D'Avernas and Vandeweyer (2019) etc.) to solve (2.31). The method works by approximating the derivatives in PDE by discretizing the state space and then solving the discretized problem using an explicit or implicit method with up-winding scheme. I briefly discuss the methodology before presenting the solution. Consider a discretized version of the PDE in (2.12)

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

where $\{i\}_1^{N_z}, \{j\}_1^{N_t}$ denote the space and time dimension grid points respectively, and the derivatives of J are approximated using

$$\begin{aligned}
\frac{\hat{\partial}J_{i,j}}{\hat{\partial}z} &= (\mu_j^x)^+\frac{J_{i+1,j}-J_{i,j}}{\Delta_i} + (\mu_j^x)^-\frac{J_{i,j}-J_{i-1,j}}{\Delta_i} \\
\frac{\hat{\partial}^2 J_{i,j}}{\hat{\partial}z^2} &= \frac{J_{i+1,j}-2J_{i,j}+J_{i-1,j}}{\Delta_i^2} \\
\frac{\hat{\partial}J_{i,j}}{\hat{\partial}t} &= \frac{J_{i,j+1}-J_{i,j}}{\Delta_j}
\end{aligned}$$

where
$$(\mu_j^x)^+ = \begin{cases} \mu_{t,j}^x & \text{if} \quad \mu_{t,j}^x > 0 \\ 0 & \text{if} \quad \text{otherwise} \end{cases} \qquad (\mu_j^x)^- = \begin{cases} \mu_{t,j}^x & \text{if} \quad \mu_{t,j}^x < 0 \\ 0 & \text{if} \quad \text{otherwise} \end{cases} \tag{2.32}$$

The first derivative of J with respect to space dimension is approximated using a forward difference if the advection coefficient is positive, and using a backward difference if the coefficient is negative, as shown in (2.32). This ensures monotonicity of the finite difference

scheme. The method is implicit since $J_{i,j+1}$ is on both the L.H.S and R.H.S of the discretized PDE equation. This presents a system of linear equations that can be solved using traditional methods such as Richardson method. There are two issues that arise when applying a finite difference method in high dimensions. The first is the well-known *curse of dimensionality*. With d space dimensions and 1 time dimension, the mesh size in finite difference method is $\mathcal{O}^{d+1}$ which becomes infeasible to handle if d is large. Moreover, along with the explosion of the mesh size, there arises a need for reduced time step size. The second problem that has received relatively less attention in the literature is the need to preserve monotonicity of the elliptical operator when the state space is large and potentially correlated. In the case of a one-dimensional state space, the monotonicity can be preserved using an up-winding scheme given in (2.32) which approximates the derivatives in the right spatial direction. However, as documented in D'Avernas and Vandeweyer (2019), it is not trivial to preserve monotonicity using an up-winding scheme even in a two space dimensional case since the right direction in which derivatives should be approximated may fall outside the grid.

The second method that is commonly used is a projection method (Judd (1992)) where the function J in (2.12) is approximated using Chebychev polynomials as basis functions. A related technique used extensively in computational engineering is the Galerkin method that finds a linear combination of basis functions that best approximates the PDE. To be precise, suppose the goal is to solve for the PDE $\mathscr{H}(J) = 0$. The approximation for J can be obtained from

$$\tilde{J} = \sum_{i=1}^{n} \alpha_i \phi_i \tag{2.33}$$

where $\{\phi_i\}$ is bases for J. The residual after plugging in the approximation is denoted as

$$\mathscr{R}(\cdot \mid \boldsymbol{\alpha}) = \mathscr{H}(\tilde{J})$$

The coefficients $\boldsymbol{\alpha}$ satisfy the inner product for $j = 1, ..., n$

$$\left\langle \mathscr{R}(\cdot \mid \boldsymbol{\alpha}), \phi_i \right\rangle = 0$$

One has to choose the basis functions carefully since the computation of inner product can be computationally expensive.

2.4.3 Model Solution

I calibrate the benchmark model with parameters in Table (2.1). The goal is to demonstrate that neural network approximation is close to the finite difference approximation, and hence I mostly choose the parameters based on prior literature. The discount rate and depreciation rate of capital is taken from Hansen et al. (2018). The volatility is chosen to be 6% so as to achieve a reasonable variation in the risk prices. Risk aversion parameter for the agents is set to 2, and the IES is set to 1. Productivity parameters, and the equity retention rate are close to

	Description	Symbol	Value
	Volatility of output	σ	0.06
	Discount rate (experts)	ρ	0.05
	Depreciation rate of capital	δ	0.05
Technology/Preferences	Investment cost	κ	10
	Productivity (experts)	a_e	0.15
	Productivity (households)	a_h	0.03
Utility	Risk aversion	γ	2
Friction	Equity retention	$\underline{\chi}$	0.5

Table 2.1 – Calibrated parameters for the benchmark model. All values are annualized.

	Value
No. of hidden layers	4
Hidden units	[30,30,30,30]
Activation function	Tanh (hidden), Sigmoid (output)
Optimizer	ADAM + L-BFGS-B
Learning rate	0.001
Loss function weights$(\lambda_f, \lambda_j, \lambda_b, \lambda_c^1, \lambda_c^2)$	{1,1,0.001,1,1}
Batch size	Full batch

Table 2.2 – Network hyperparameters and architecture.

Brunnermeier and Sannikov (2016b). The network architecture and hyperparameters used to solve the de-coupled system of PDEs (2.31) are provided in the Table (2.2). I use four hidden layers with 30 neurons in each layer. In principle, one could use one hidden layer and more neurons but experimentation shows that it is better to have a deep network instead of a wide network.[17] The total number of points in the inner static step is 1,000 and the total training sample size to solve the PDE is 300. Figure (2.4) present the solution obtained using the finite difference method and the neural network method. They not only look qualitatively the same, but they also match upto the order 1e-3. Figure (2.3) shows the absolute error in value function over time. A comparison with finite difference scheme shows that the neural network method has a larger error drop in each iteration leading to convergence in 20 time steps. Note that the neural network error curve is not as smooth as the finite difference one due to the random sampling from the state space before every PDE time step. However, convergence is achieved in 20 iterations while it takes around 70 iterations for the finite difference method.

[17] It is a common practice in the deep learning literature to have deeper layers instead of wider layers since they have superior performance for a wide range of problems.

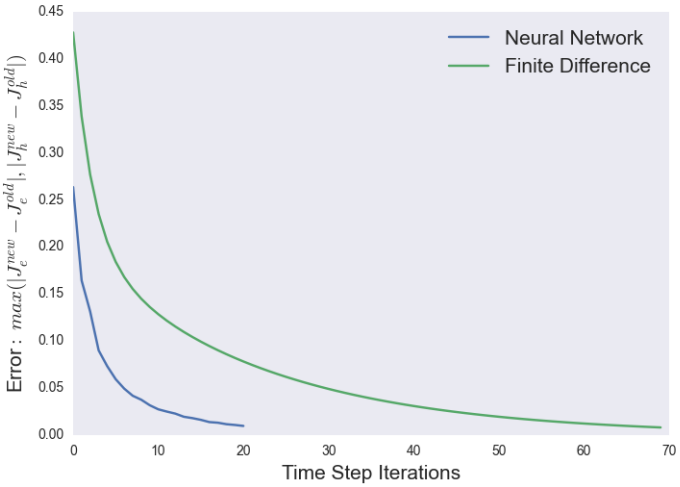

Figure 2.3 – Comparison of value function error in finite difference and neural network method.

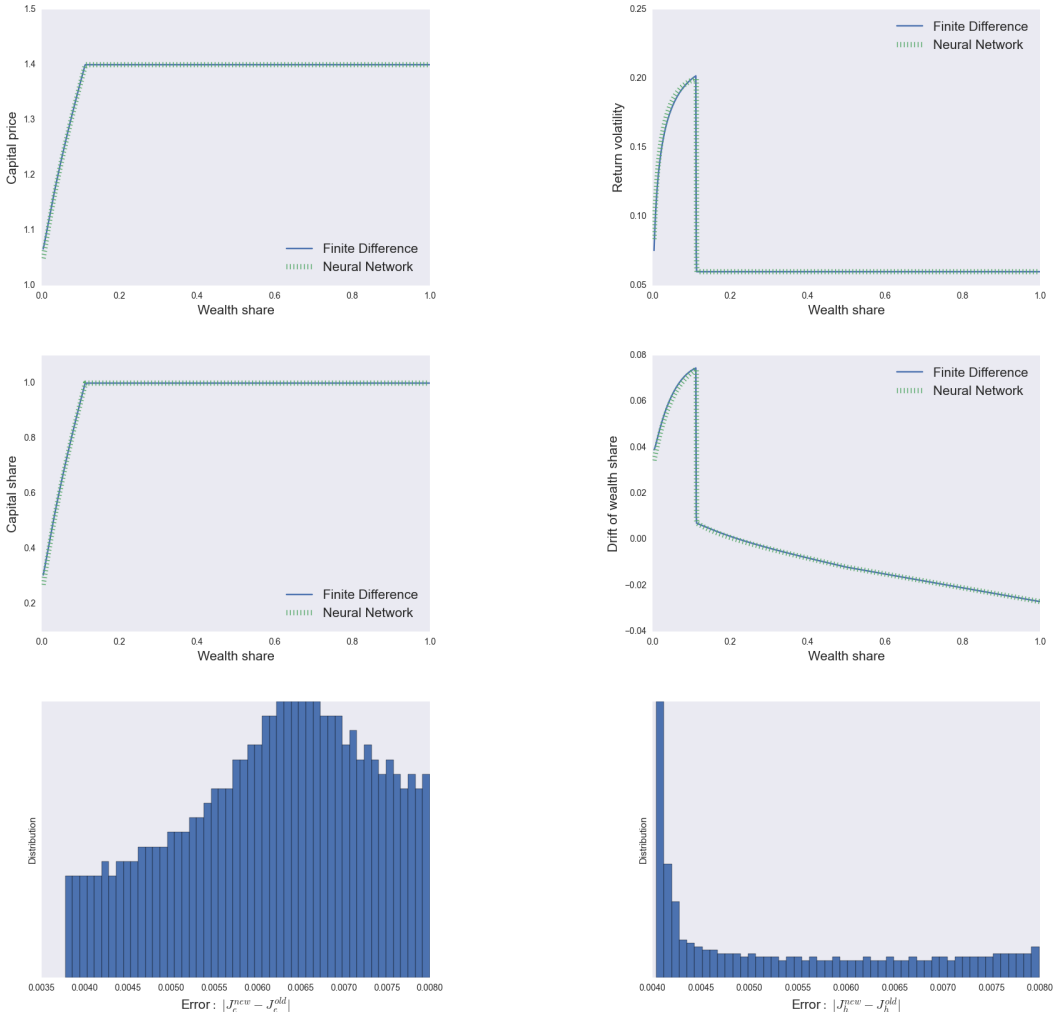

Figure 2.4 – Comparison of equilibrium quantities using finite difference and neural network in the benchmark model. Last two plots correspond to absolute error in functions $\hat{J}_h$ and $\hat{J}_e$.

2.5 Brunnermeier-Sannikov meets Bansal-Yaron

This section considers a heterogeneous agent model similar to the benchmark case but with more complexity. Specifically, I combine the model of Brunnermeier and Sannikov (2016b) and Bansal and Yaron (2004) in continuous time. The economy features heterogeneous agents with experts facing shocks to their productivity and the capital. In addition, the long run growth rate and volatility of capital are subjected to independent Brownian shocks. The state space becomes four dimensional and the PDEs to be solved in the outer time step has five dimensions including the artificial time dimension. I explain the model first and then discuss the results.

2.5.1 Model

There are two types of agents and the aggregate capital is denoted by K_t with $t \in [0,\infty)$. The capital process is governed by

$$\frac{dk_{j,t}}{k_t} = (\Phi(\iota_{j,t}) + g_t - \delta)dt + \sigma\sqrt{s_t}dZ_t^k \tag{2.34}$$

where $\iota_{j,t}$ is the investment rate, $\Phi(\cdot)$ is the investment cost function, g_t is the time varying growth rate of capital, and $\sigma\sqrt{s_t}$ is the time varying volatility of capital. The production technology is given by

$$y_{j,t} = a_{j,t}k_{j,t} \tag{2.35}$$

where $k_{j,t}$ is the capital held by agent type j. I assume that the productivity of experts is time varying and is governed by the process

$$da_{e,t} = \lambda^a(\hat{a}_e - a_{e,t})dt + \sigma^a(a_{e,t})dZ_t^a \tag{2.36}$$

where λ^a is the mean reversion rate. For simplicity, I assume that the productivity of households is constant $a_{h,t} = a_h$. The growth rate and volatility of capital follows the exogenous processes

$$dg_t = \lambda^g(\hat{g} - g_t)dt + \sigma^g(g_t)dZ_t^g \tag{2.37}$$

$$ds_t = \lambda^s(\hat{s} - s_t)dt + \sigma^s(s_t)dZ_t^s \tag{2.38}$$

where λ^g, λ^s are the mean reversion rates of growth rate and capital processes respectively. I assume that the functions $(\sigma^i(\cdot); i \in \{g, s, a\})$ take the form

$$\sigma^g(\cdot) = (\bar{g} - g_t)(g_t - \underline{g}) \tag{2.39}$$

$$\sigma^s(\cdot) = (\bar{s} - s_t)(s_t - \underline{s}) \tag{2.40}$$

$$\sigma^a(\cdot) = (\bar{a}_e - a_{e,t})(a_{e,t} - \underline{a}_e) \tag{2.41}$$

The assumed functional form above is for convenience in computation and can be modified to resemble a Feller square-root process or an Ornstein-Uhlenbeck process. I impose $\bar{a}_e > \underline{a}_e > a_h$ so that experts productivity is always higher than that of households even thought it is time varying. The Brownian shocks $\{dZ_t^k, dZ_t^g, dZ_t^s, dZ_t^a\}$ have zero cross-correlation for simplicity. Overall, the model can be thought of as a combination of Brunnermeier and Sannikov (2016b) and Bansal and Yaron (2004) with additional productivity shocks. The agents trade the capital which has a price process given by

$$\frac{dq_t}{q_t} = \mu_t^q dt + (\boldsymbol{\sigma_t^q})^T \boldsymbol{dZ_t} \tag{2.42}$$

where the vectors $(\boldsymbol{\sigma^q})^T = \begin{bmatrix} \sigma^{q,k} & \sigma^{q,g} & \sigma^{q,s} & \sigma^{q,a} \end{bmatrix}$ and $\boldsymbol{dZ_t} = \begin{bmatrix} dZ_t^k & dZ_t^g & dZ_t^s & dZ_t^s \end{bmatrix}^T$. The agents can also trade in the risk free security that pays a return r_t. The agents cannot write contracts on the aggregate state of the economy. That is, the investment in capital and the exposure to aggregate shocks $\boldsymbol{Z_t}$ are intertwined. This is for simplicity and can be extended to including a derivative market to hedge the aggregate shocks. This will have an impact on the equilibrium policies and prices but the numerical method to solve the HJB equations will remain the same. Thus, I intertwine the capital holding decisions and the exposure to the aggregate shocks to ease computation of equilibrium policies in the static loop. Since the dividend yield from each unit of the capital held is different for the households and the experts[18], the agent-specific return process follows[19]

$$dR_{j,t} = \left(\frac{a_{j,t} - \iota_{j,t}}{q_t} + \mu_t^q + \Phi(\iota) - \delta + \sigma\sqrt{s_t}\sigma_t^{q,k} + g_t \right) dt + (\boldsymbol{\sigma_t^R})^T \boldsymbol{dZ_t} \tag{2.43}$$

where the vector $(\boldsymbol{\sigma_t^R})^T = \begin{bmatrix} \sigma_t^{q,k} + \sigma\sqrt{s_t} & \sigma_t^{q,g} & \sigma_t^{q,s} & \sigma_t^{q,a} \end{bmatrix}$, and $\boldsymbol{dZ_t}$ is as before. The aggregate output in the economy is $y_t = A_t K_t$ where the aggregate dividend is given by

$$A_t = \int_{\mathbb{H}} a_h \frac{k_{j,t}}{K_t} dj + \int_{\mathbb{E}} a_{e,t} \frac{k_{j,t}}{K_t} dj$$

where $K_t = \int_{\mathbb{H}\cup\mathbb{E}} k_{j,t} dj$ denotes the aggregate capital. Let the share of aggregate capital held by the experts be denoted by ψ_t. That is,

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

The experts are constrained in issuing equity to the households such that they have to retain a fraction $\underline{\chi} \in [0,1]$ of the equity in their balance sheet. They are free to issue the remaining equity to the households who may desire to hold it in equilibrium. The stochastic discount

[18] The investment rate can also be different between the agents, but it turns out that in equilibrium, the optimal investment is tightly linked to the capital price. Since the price is unique, the investment rate $\iota_{j,t}$ is the same for all agents.

[19] The return for agent j is $dR_{j,t} = \frac{(a_{j,t} - \iota_{j,t})k_{j,t}}{q_t k_{j,t}} dt + \frac{d(q_t k_{j,t})}{q_t k_{j,t}}$

factor (SDF) process for each type of agent is given by

$$\frac{\partial \xi_{j,t}}{\xi_{j,t}} = r_t dt - \boldsymbol{\zeta}_{\boldsymbol{j,t}}^T \boldsymbol{dZ_t} \tag{2.44}$$

where the vector $\boldsymbol{\zeta}_{\boldsymbol{j,t}}^T = \begin{bmatrix} \zeta_{j,t}^k & \zeta_{j,t}^g & \zeta_{j,t}^s & \zeta_{j,t}^a \end{bmatrix}$ captures the market prices of risk for each Brownian shock in $\boldsymbol{dZ_t}$. Since both agents trade in the risk-free market, they receive the same return r_t. Following Gopalakrishna (2020), I assume that the experts are subjected to Poisson shocks that will force them to exit the economy and become households. That is, at each time instant dt, a fraction $\tau_t dt$ of the experts transition into households.

Equilibrium: The optimization problem for each agent type j is given by

$$U_{j,t} = \sup_{c_{j,t}, k_{j,t}} E_t \left[\int_t^{T_j} f(c_{j,s}, U_{j,s}) ds + 1_{j \in \mathbb{E}} U_{h,\tau'} \right] \tag{2.45}$$

$$\text{s.t.} \frac{dw_{j,t}}{w_{j,t}} = \left(r_t - \frac{c_{j,t}}{w_{j,t}} + \frac{q_t k_{j,t}}{w_{j,t}} ((\mu_{j,t}^R - r_t) - (1 - \chi_{j,t}) \bar{\epsilon}_{j',t}) \right) dt + \sigma_{w_j,t} (\boldsymbol{\sigma_t^R})^T \boldsymbol{dZ_t} \qquad j \in \{e, h\}$$

where $\bar{\epsilon}_{j,t} = \boldsymbol{\zeta}_{\boldsymbol{j,t}}^T \boldsymbol{\sigma_t^R}$, $T_j = \tau'$ for experts and $T_j = \infty$ for households, and τ' is the time of transition. The agent j earns $\mu_{j,t} - r_t$ from borrowing in the risk free market and investing in the risky capital, but have to pay the outside equity holders a compensation for their risk. Households do not issue outside equity (i.e., $\chi_{h,t} = 1$) and I denote $\chi_{e,t}$ as χ_t for simplicity of notation. The volatility terms in wealth is given by

$$\sigma_{w_e,t} = \frac{q_t k_{e,t}}{w_{e,t}} \chi_t \tag{2.46}$$

$$\sigma_{w_h,t} = \frac{q_t k_{h,t}}{w_{h,t}} + (1 - \chi_t) \frac{q_t w_{e,t}}{w_{h,t}} \tag{2.47}$$

The experts retain the fraction χ_t of risk in their balance sheet and offload the remaining to the households. The agents solve for the optimal consumption $c_{j,t}$ and portfolio holdings $k_{j,t}$ by maximizing the objective function (2.45). The optimal investment rate $\iota_{j,t}$ is found by maximizing the expected return on risky capital. The competitive equilibrium is defined as

Definition 2.5.1. A competitive equilibrium is a set of aggregate stochastic processes adapted to the filtration generated by the Brownian motion $\boldsymbol{Z_t}$. Given an initial distribution of wealth between the experts and the households, the processes are prices (q_t, r_t), policy functions $(c_{j,t}, \iota_{j,t}, k_{j,t}; j \in \{e, h\})$, and net worth $(w_{j,t}; j \in \{e, h\})$, such that

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

I seek a Markov equilibrium where each agents within the same type compute the optimal policies. Let the wealth share experts be defined by

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

where $W_{e,t} = \int_{\mathbb{E}} w_{j,t} dj$ and $K_t = \int_{\mathbb{E}} k_{j,t} dj + \int_{\mathbb{H}} k_{j,t} dj$. The wealth share z_t is the endogenous state variable in the model which moves in response to the other equilibrium objects. The set of exogenous state variables is given by $\{g_t, s_t, a_{e,t}\}$. These four state variables characterize the whole system. The stochastic processes for the exogenous state variables are given in (2.37), (2.38), and (2.36) respectively. The proposition below provides the process for the endogenous state variable.

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

$$\frac{dz_t}{z_t} = \mu_t^z dt + (\boldsymbol{\sigma_t^z})^T \boldsymbol{dZ_t} \tag{2.48}$$

where

$$\mu_t^z = \frac{a_{e,t} - \iota_{e,t}}{q_t} - \frac{C_{e,t}}{W_{e,t}} + \frac{\chi_t \psi_t}{z_t}(\tilde{\boldsymbol{\zeta}}_{\boldsymbol{e,t}}^T \boldsymbol{\sigma_t^R}) + (1-\chi_t)(\boldsymbol{\zeta}_{\boldsymbol{e,t}}^T - \boldsymbol{\zeta}_{\boldsymbol{h,t}}^T)\boldsymbol{\sigma_t^R} - \tau_t$$

$$(\boldsymbol{\sigma_t^z}) := \begin{bmatrix} \sigma_t^{z,k} & \sigma_t^{z,g} & \sigma_t^{z,s} & \sigma_t^{z,a} \end{bmatrix}^T = (\frac{\chi_t \psi_t}{z_t} - 1)\boldsymbol{\sigma_t^R}$$

$$\tilde{\boldsymbol{\zeta}}_{\boldsymbol{e,t}}^T := (\boldsymbol{\sigma_t^R})^T (\boldsymbol{\zeta_{e,t}} - \boldsymbol{\sigma_t^R})$$

Proof: See Appendix A.2.4.
Note that the exit rate τ_t enters the drift of the wealth share.

Asset pricing conditions: The agents maximize the return on capital to obtain the optimal investment rate. That is, they solve

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

I assume a simple investment cost function given by $\Phi(\iota) = \frac{\log(\kappa\iota+1)}{\kappa}$. Then, $\iota_{j,t}^*$ is given by

$$\iota_{j,t}^* = \frac{q_t - 1}{\kappa} \tag{2.49}$$

where κ is an investment cost parameter. Since the optimal investment rate depends only on the capital price, it is not agent-specific. This is a reflection of q-theory result which ties the investment rate and capital price tightly. The asset pricing condition for the experts is given

by[20]

$$\frac{a_{e,t}-\iota_t}{q_t}+\Phi(\iota_t)-\delta+g_t+\mu_t^q+\sigma\sqrt{s_t}\sigma_t^{q,k}-r_t=\chi_t\bar{\boldsymbol{\epsilon}}_{\boldsymbol{e,t}}+(1-\chi_t)\bar{\boldsymbol{\epsilon}}_{\boldsymbol{h,t}} \tag{2.50}$$

where $\bar{\epsilon}_{j,t}=\boldsymbol{\zeta}_{\boldsymbol{j,t}}^T\boldsymbol{\sigma}_{\boldsymbol{t}}^{\boldsymbol{R}}$ and χ_t is the share of inside equity chosen by the experts. For the households, the condition is given by

$$\frac{a_h-\iota_t}{q_t}+\Phi(\iota_t)-\delta+g_t+\mu_t^q+\sigma\sqrt{s_t}\sigma_t^{q,k}-r_t\leq\bar{\boldsymbol{\epsilon}}_{\boldsymbol{h,t}} \tag{2.51}$$

with equality when $\psi_t<1$. We can combine the asset pricing condition for experts and households to get

$$\frac{a_{e,t}-a_h}{q_t}\geq\chi_t(\bar{\epsilon}_{e,t}-\bar{\epsilon}_{h,t}) \tag{2.52}$$

$$\max\{\underline{\chi}-\chi_t,\bar{\epsilon}_{e,t}-\bar{\epsilon}_{h,t}\}=0 \tag{2.53}$$

where (2.52) holds with equality if risk premia of experts is larger than that of the households. In regions of the state space where wealth share is sufficiently large, χ_t is chosen to satisfy $\bar{\epsilon}_{e,t}=\bar{\epsilon}_{h,t}$.

HJB equations: The HJB problem for agent type j can be written as

$$\sup_{c_{j,t},k_{j,t}} f(c_{j,t},U_{j,t})+E[dU_{j,t}]=0 \tag{2.54}$$

The conjecture for value function is of the form

$$U_{j,t}=\frac{(J_{j,t}(z_t,g_t,s_t,a_{e,t})K_t)^{1-\gamma}}{1-\gamma}$$

where $J_{j,t}(\cdot)$ captures the stochastic investment opportunity set. The process for $J_{j,t}$ is given by

$$\frac{\partial J_{j,t}}{J_{j,t}}=\mu_{j,t}^J dt+\boldsymbol{\sigma}_{\boldsymbol{j,t}}^{\boldsymbol{J}}\boldsymbol{dZ_t}$$

where the quantities $\mu_{j,t}^J$ and $\boldsymbol{\sigma}_{\boldsymbol{j,t}}^{\boldsymbol{J}}=\begin{bmatrix}\sigma_{j,t}^k & \sigma_{j,t}^g & \sigma_{j,t}^s & \sigma_{j,t}^a\end{bmatrix}$ are to be solved in equilibrium.

Proposition 5. *The optimal policies are given by*

$$\hat{c}_{j,t}=\rho \tag{2.55}$$

$$\zeta_{j,t}^k=-\sigma_t^{J,k}+\sigma_{j,t}^{z,k}+\sigma_t^{q,k}+\gamma\sigma\sqrt{s_t} \tag{2.56}$$

$$\zeta_{j,t}^i=-\sigma_{j,t}^{J,i}+\sigma_{j,t}^{z,i}+\sigma_t^{q,i}\qquad i\in\{g,s,a\} \tag{2.57}$$

Proof: See Appendix.

[20]This can be derived using a martingale argument as shown in Appendix A.2.4.

Since the IES is set to 1, the consumption-wealth ratio ($\hat{c}_{j,t}$) is equal to the discount rate ρ. The market prices of risk are given up to the other equilibrium objects $(\mu^J_{j,t}, \boldsymbol{\sigma^J_{j,t}}, \boldsymbol{\sigma^R_t}, q_t, \psi_t)$ which are solved in the state space $\boldsymbol{x}_t^T = \begin{bmatrix} z_t & g_t & s_t & a_{e,t} \end{bmatrix}$.

Definition 2.5.2. A Markov equilibrium in $(\mathbf{z_t} \in (\mathbf{0},\mathbf{1}), \mathbf{g_t} \in (\underline{\mathbf{g}}, \bar{\mathbf{g}}), \mathbf{s_t} \in (\underline{\mathbf{s}}, \bar{\mathbf{s}}), \mathbf{a_{e,t}} \in (\underline{\mathbf{a}}_\mathbf{e}, \bar{\mathbf{a}}_\mathbf{e}))$ is a set of adapted processes $q(\boldsymbol{x_t}), r(\boldsymbol{x_t}), J_e(\boldsymbol{x_t}), J_h(\boldsymbol{x_t})$, policy functions $\hat{c}_e(\boldsymbol{x_t}), \hat{c}_h(\boldsymbol{x_t}), \psi(\boldsymbol{x_t})$, and state variables $\boldsymbol{x_t}$ such that

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

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

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

- The state variables $\boldsymbol{x_t}$ satisfy (2.37), (2.38), (2.36), and (A.75).

2.5.2 Numerical method

The model is solved numerically that involves two blocks. The first block is the static inner step that takes $J_{j,t}$ as given and solves for the equilibrium quantities $(q_t, \chi_t, \psi_t, \boldsymbol{\sigma^R_t})$. The other equilibrium objects are computed using these quantities. The second block is the outer time step that takes the equilibrium quantities as given and updates the values of $J_{j,t}$ by solving a de-coupled system of PDEs- one for the expert and one for the household. The inner block is solved using a Newton-Rhaphson method that is computationally fast, and the outer block is solved using ALIENs to overcome the curse of dimensionality problem.

Proposition 6. *The equilibrium objects* $(\psi_t, q_t, \sigma_t^{q,k}, \sigma_{q,g}, \sigma_{q,s}, \sigma_{q,a})$ *are found by solving the differential-algebraic system of equations given by*

$$\begin{aligned} 0 = {} & \frac{a_{e,t} - a_h}{q_t} - \chi_t \left((1-\gamma) \left(\frac{1}{J_{h,t}} \frac{\partial J_{h,t}}{\partial z_t} - \frac{1}{J_{e,t}} \frac{\partial J_{e,t}}{\partial z_t} \right) + \frac{1}{z_t(1-z_t)} \right) (\chi_t \psi_t - z_t) [||\boldsymbol{\sigma^R}||^2] \\ & - \sum_{i \in \{g,s,a\}} \chi_t (1-\gamma) \left(\frac{1}{J_{h,t}} \frac{\partial J_{h,t}}{\partial i_t} - \frac{1}{J_{e,t}} \frac{\partial J_{e,t}}{\partial i_t} \sigma^i \sigma^{q,i} \right) \quad if \quad \psi_t < 1 \end{aligned} \tag{2.58}$$

$$0 = (\rho_e z_t + (1-z_t)\rho_h) q_t - \psi_t (a_{e,t} - \iota_t) + (1-\psi_t)(a_h - \iota_t) \tag{2.59}$$

$$0 = \sigma_t^{q,k} + \sigma \sqrt{s_t} - \frac{\sigma \sqrt{s_t}}{1 - \frac{1}{q_t} \frac{\partial q_t}{\partial z_t} (\chi_t \psi_t - z_t)} \tag{2.60}$$

$$0 = \sigma_t^{q,i} - \frac{\frac{1}{q_t} \frac{\partial q_t}{\partial i_t}}{1 - \frac{1}{q_t} \frac{\partial q_t}{\partial z_t} (\chi_t \psi_t - z_t)} \qquad i \in \{g, s, a\} \tag{2.61}$$

Proof: See Appendix.

The first equation is obtained from the differences in risk premium of the experts and the households from equations (2.50) and (2.51). The second equation comes from the goods market clearing condition and the remaining equations are obtained from writing return volatility objects in terms of other equilibrium quantities. I employ a Newton-Rhaphson method that is computationally fast even in high dimensions. The Newton method is highly sensitive to the initial values provided. To avoid errors, I provide as inputs the equilibrium values from prior points in the grid. I relegate further details to the Appendix. Once the six quantities are found, the other other equilibrium objects ($\boldsymbol{\zeta_{e,t}}, \boldsymbol{\zeta_{h,t}}, \mu_t^z, \boldsymbol{\sigma_t^z}, \mu_{e,t}^J, \mu_{h,t}^J, \boldsymbol{\sigma_{e,t}^J}, \boldsymbol{\sigma_{h,t}^J}$) are easily computed since they are a function of the state variables and the other variables found in the static step.

Proposition 7. *The stochastic opportunity set J is characterized by the PDE*[21]

$$0 = \frac{\partial J}{\partial t} + A\left(\boldsymbol{x}, J, \frac{\partial J}{\partial \boldsymbol{x}}\right) + \frac{1}{2} tr\left[B\left(\boldsymbol{x}, J, \frac{\partial J}{\partial \boldsymbol{x}}\right)\frac{\partial^2 J}{\partial \boldsymbol{x}^2} B\left(\boldsymbol{x}, J, \frac{\partial J}{\partial \boldsymbol{x}}\right)^T\right] \tag{2.62}$$

with the boundary conditions

$$J(\boldsymbol{x}, t) = \tilde{J} \tag{2.63}$$

$$\partial J(\boldsymbol{x}^0, t) = \partial J(\boldsymbol{x}^1, t) = 0 \tag{2.64}$$

where[22]

$$\begin{aligned}
A\left(\boldsymbol{x}, J, \frac{\partial J}{\partial \boldsymbol{x}}\right) &= J\Big(\rho\big(\log\rho - \log J + \log(qz)\big) + \Phi(\iota_t) + g_t - \delta\Big) \\
&\quad - J\left((\gamma-1)\sigma\sigma^{J,k} + \frac{\gamma}{2}||\sigma^J||^2 - 1_{j\in\mathbb{E}}\frac{\tau_t}{1-\gamma}\left(\left(\frac{J_{j,t}}{J_{j',t}}\right)^{1-\gamma} - 1\right)\right) \\
&\quad + \begin{bmatrix}\frac{\partial J}{\partial z} & \frac{\partial J}{\partial g} & \frac{\partial J}{\partial s} & \frac{\partial J}{\partial a_e}\end{bmatrix}\begin{bmatrix} z\mu^z & \mu^g & \mu^s & \mu^a\end{bmatrix}^T \\
B\left(\boldsymbol{x}, J, \frac{\partial J}{\partial \boldsymbol{x}}\right) &= \begin{bmatrix}\sigma^z & \sigma^g & \sigma^s & \sigma^a\end{bmatrix} \\
\frac{\partial^2 J}{\partial \boldsymbol{x}^2} &= \begin{bmatrix}\frac{\partial^2 J}{\partial z^2} & \frac{\partial^2 J}{\partial g^2} & \frac{\partial^2 J}{\partial s^2} & \frac{\partial^2 J}{\partial a_e^2}\end{bmatrix} \\
\boldsymbol{x}^0 &= \{(0,g,s,a),(z,\underline{g},s,a),(z,g,\underline{s},a),(z,g,s,\underline{a}_e)\} \\
\boldsymbol{x}^1 &= \{(1,g,s,a),(z,\bar{g},s,a),(z,g,\bar{s},a),(z,g,s,\bar{a}_e)\}
\end{aligned}$$

Proof: See Appendix.

Once the static step is completed, it remains to solve the partial differential equations given in (2.62). The coefficients of the PDE are computed using the solution form static step. Hence, taking the coefficients as given, the task is to find the function J that solves the above PDE.

[21] I write $J_{j,t}$ simply as J to avoid clustering of notations.

[22] The index j' refers to the other type of agent.

2.5.3 Solution

Figure (2.5) presents the equilibrium quantities for different values of volatility. A higher volatility s_t leads to lower capital price throughout the state space since it reflects an increased macroeconomic uncertainty. Interestingly, a lower volatility decreases the endogenous return volatility in the normal region but increases the return volatility in the crisis region. In Brunnermeier and Sannikov (2014a), the volatility is constant but a static comparison leads to similar result, which they call as the 'volatility paradox'. The paradox refers to the fact that a decrease in exogenous fundamental volatility leads to an increase in endogenous price volatility during crises. This paradox carries over to a more general model of stochastic volatility as well. The volatility has a decreasing effect on the drift of the wealth share. A higher volatility reduces the risk premium and capital price and therefore leads to a lower drift. This has important implications on how the system transitions in and out of crises. A larger volatility means that the system spends longer in crises since the experts rebuild their wealth slowly due to reduced risk premium. In the normal region where experts are wealthy, a higher volatility increases the

risk premium and hence the drift of wealth share is larger.

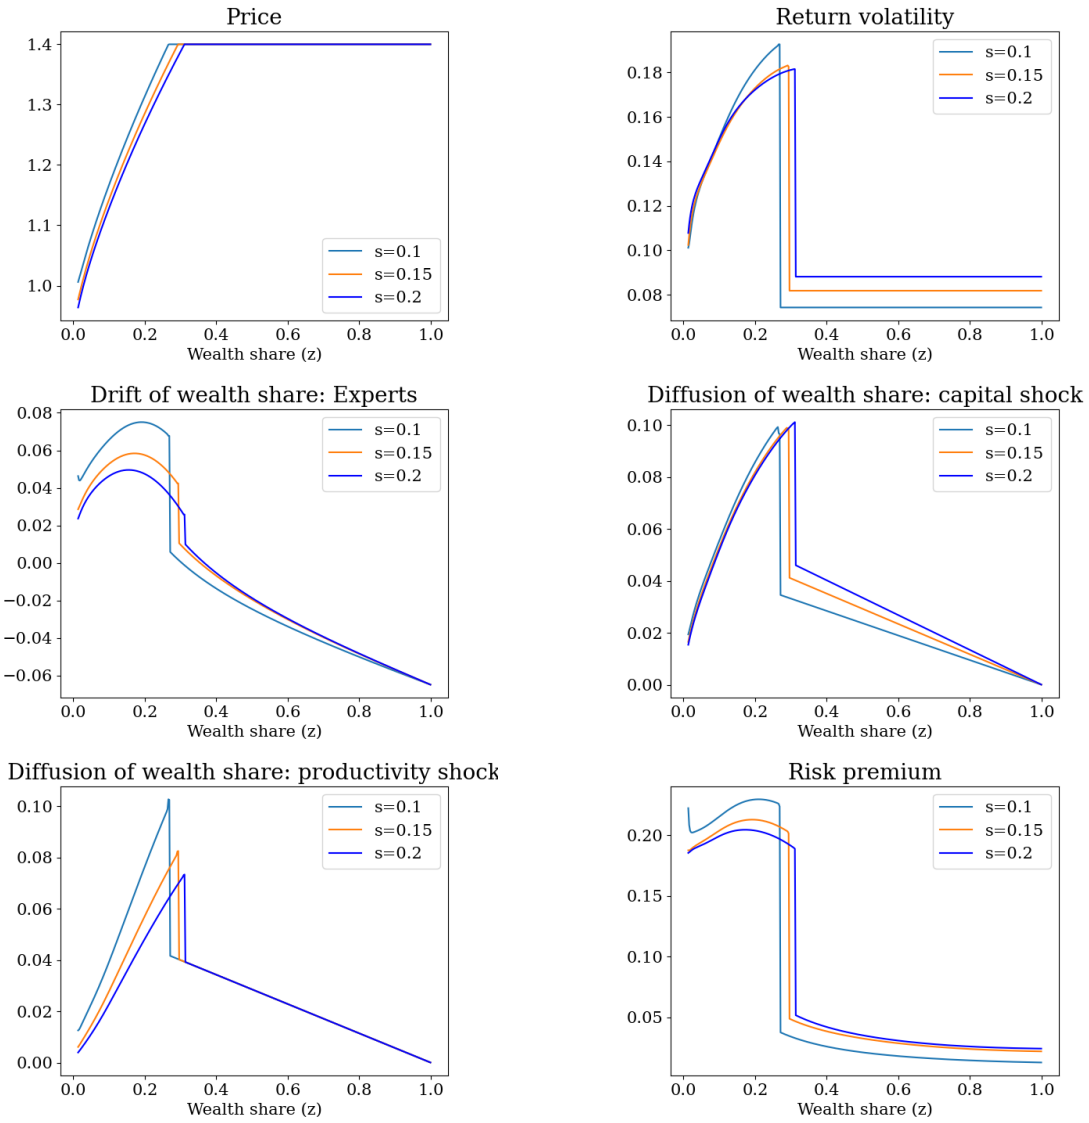

Figure 2.5 – Equilibrium quantities for different volatility values (s_t). Growth rate (g_t) and productivity ($a_{e,t}$) are fixed at respective average values.

Figure (2.6) presents the equilibrium quantities for different values of expert productivity $a_{e,t}$. The time variation induced by productivity on the risk prices are higher than that induced by volatility, since the former directly affects the capital price through the goods market clearing condition. When experts are more productive, the capital price is larger since the aggregate dividend is increasing in the proportion of capital held by the more productive experts. When productivity drops towards the lowest level of 10%, the risk premium goes down decreasing the drift of the wealth share. Hence, less productive experts force the system to spend a longer time in crises. As productivity revers to its upper level of 20%, the drift of the wealth share increases and pushes the system out of crises. Gopalakrishna (2020) shows a similar phenomenon where

the twin forces of stochastic productivity and regime-dependent transition rate of experts help quantitatively explain the crisis dynamics.

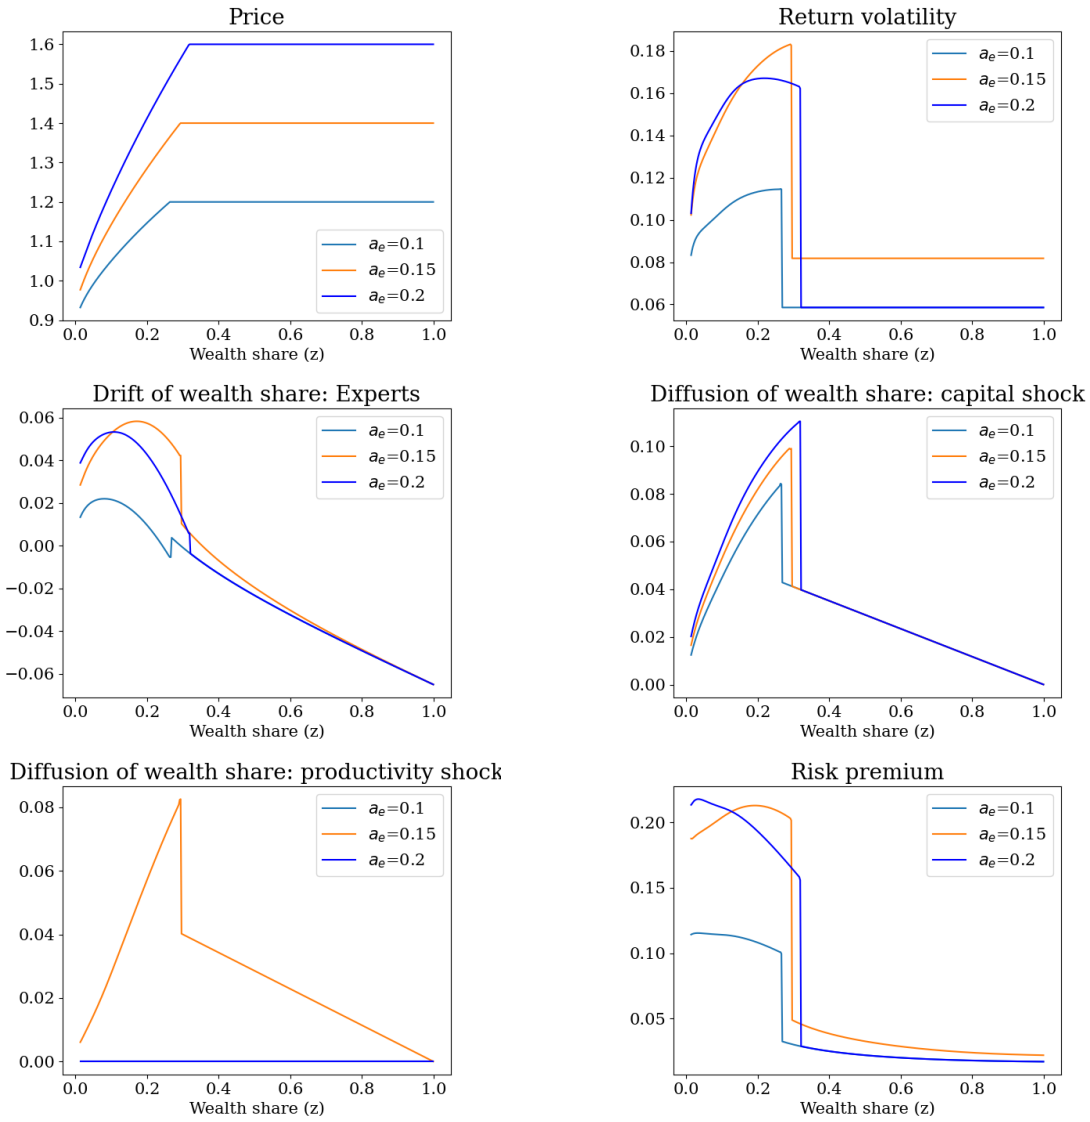

Figure 2.6 – Equilibrium quantities for different productivity values ($a_{e,t}$). Growth rate (g_t) and volatility (s_t) are fixed at respective average values.

2.6 Conclusion

I have developed a new computational technique called Actively Learned and Informed Equilibrium Nets (ALIENs) to solve continuous time economic models featuring heterogeneous agents, occasionally binding constraints, endogenous jumps, and highly non-linear policy functions. The technique relies on solving a system of parabolic differential equations using a collection of deep neural networks by converting them into a sequence of supervised learn-

ing problems. The HJB equation from the agents' optimization problem takes the form of a non-linear elliptical PDE, which is converted into a quasi-linear parabolic type by adding an artificial time derivative and treating the PDE coefficients as known constants. The value function that solves the PDE is approximated using a neural network and is forced to obey the laws that are governed by the economic system. The architecture is split into i) a PDE network that is responsible for the fitted value function to satisfy the partial differential equation, ii) a boundary network that strives to fit the boundary condition of the PDE, and iii) an active network that ensures a better fit in the state space with most economic information. I utilize data parallelism that leverages the computing hardware to significantly speed up the computational time.

I have applied the method to solve a macro-finance benchmark model with an endogenous state variable, non-linear policy functions, and showed that the neural network solution matches the finite difference solution. The second application is to a similar model with capital misallocation, endogenous jumps, and endogenous state variables but in a higher dimension. The value function in this problem is challenging to solve using a traditional finite difference scheme, since endogenous PDE coefficients lead to difficulty in maintaining the monotonicity of the scheme. In addition, the high dimensionality of PDE creates a massive computational bottleneck. ALIENs successfully sidesteps these limitations and ensures convergence by actively tracking the subdomain with most economic information to create informed sample points for training the neural network. Lots of problems in continuous time are formulated in a small dimensional state space due to computational bottlenecks. The ease with which ALIENs can be implemented opens up opportunities in fields as diverse as macroeconomics, asset pricing, and dynamic corporate finance. Moreover, the fact that ALIENs are powered with data parallelism through minimal effort opens up new avenues to perform extensive quantitative analysis that requires experimentation and repeatedly solving models several times.

3 Origins of bank failures

It is widely understood that financial crises stem from disruptions in the credit intermediation activity, and feature a sharp increase in the risk premium of assets. The distress quickly spills over to the real economy causing a downfall in the investment and GDP, as evident during the Global Financial Crisis of 2008. As pointed out by Gertler et al. (2020), all major investment banking activity failed, and the money market fund in the US faced a run during the peak of the crisis. Equity risk premium rose to 17%, private investment rate fell by 10%, and GDP contracted by 8% going into the year 2009 from the last quarter of 2008 in the US. Importantly, there was no big exogenous shock that triggered the crisis. The most common theoretical explanation for this phenomenon is that financial frictions make the intermediary sector fragile, thereby amplifying the shocks to the real economy (See Gertler and Kiyotaki (2010), Brunnermeier and Sannikov (2014b), He and Krishnamurthy (2013), Di Tella (2017), etc.). In this stream of literature, the wealth share of the intermediary sector or the leverage is the key state variable that moves the asset price and the real economy. In good state of the world, intermediary wealth share in the economy is high, and a small shock does not create crisis. However, when the intermediary wealth share is low, a shock of similar magnitude triggers the amplification channel, sending the economy down a path of depressed asset prices, investment, and the GDP. Intermediation activity regains in a sluggish fashion, leading to prolonged recession.

Due to a wave of regulatory measures post 2008 financial crisis, the banks have improved capital ratios compared to the pre-crisis period (Yellen (2017)).[1] Theoretical models of financial intermediation would predict that the systemic risk in the economy is reduced as a result of improved capital ratios making the banking sector safer. However, the empirical evidence in the literature does not fully support this explanation. Sarin and Summers (2016) finds that major banks in the US and around the globe have a higher credit risk than pre-crisis levels, despite having low leverage.[2] If bank leverage alone is sufficient to determine its

[1] The regulatory measures include Dodd Frank legislation, higher capital requirement, and increased stress-testing requirements among others.

[2] Denis Beau, the First Deputy Governor of the Bank of France raises similar concerns at a Central Bank speech on October 2019: "...is the financial system now safe enough? In view of the development of the sources of

likelihood of failure, then banking sector should be safer compared to pre-crisis period given the empirically observed higher capital adequacy ratios. A safer banking sector would imply that the probability of financial crisis is lower compared to the pre-crisis period. However, the odds of bank equity to be wiped out by a magnitude of shock observed during the Global Financial Crisis remains large today despite the high capital ratios of financial institutions (Atkeson et al. (2018)). This calls for a re-examination of the origins of bank failures so as to allow policymakers to set-up preventive measures to avoid financial crisis induced by a weak banking sector.

In this chapter, I empirically investigate the sources of bank failures and find the franchise value of Bank Holding Companies (BHCs) in the US to be a significant determinant of its default probability, in addition to the leverage. Specifically, a higher franchise value is associated with a larger z-score, lower probability of bankruptcy, and lower probabiltiy of seeking assistance from the FDIC to continue operations. This relationship is robust to the inclusion of various control variables that represent size, business model, and capital adequacy ratios of the BHCs. The franchise value captures the present value of future profits that the banks aim to generate as a going concern. It is derived from either the external forces such as regulatory laws, or from the bank-specific factors such as the business model, market power, cost and profit efficiency etc. Figure (3.1) presents the franchise value of BHCs in the US between the year 1996 and 2020. The time series pattern reveals that the franchise value was high during the 1990s and the 2000s until the Global Financial Crisis. This is in part due to anti-competitive legislation by the regulatory authorities such as branching laws that lasted until late 1990s giving banks market power to derive rents from depository activities. Moreover, there is a substantial cross-sectional dispersion in franchise value among the BHCs as seen in Figure (3.1), especially after abolishing the branching laws in the late 1990s. Figure (3.2) contrasts the franchise value for banks that failed during the time period 2002 to 2020, against the franchise value of surviving banks. Strikingly, the failed banks had a lower franchise value in every year compared to the surviving banks.[3] The difference was large in the 1990s and 2000s, narrowing down during the Great Financial Crisis, and then going back up again in the post-crisis period. The stark difference in franchise value between the failed and surviving banks begs the question how does a bank derive the franchise value? I empirically find that the sources of franchise value lie in intermediation cost efficiency, business model such as depository service and degree of domestic presence, rents derived from deposit market power, and rents from government guarantees such as underpriced deposit insurance. In particular, there exists a predictive relationship between intermediation cost efficiency and franchise value: one basis point increase in the intermediation cost leads to a statistically significant drop of around 55 basis points of franchise value. The intermediation cost, which is measured as non-operating expense plus operating expense less deposit rate scaled by total assets, is counter-cyclical and highly persistent. The cyclical component of the cost has a -30% correlation with GDP cycle

vulnerabilities in the financial system, it does not seem to me to be a given that the strengthening we have observed is sufficient and efforts therefore need to be pursued."

[3] Failed banks correspond to those that filed for bankruptcy or required assistance from FDIC to continue its banking operations at any given point during the time sample considered.

between the years 1986 and 2020 that increased in magnitude to -57% in the period 2001 to 2020, with a one quarter lag auto-regressive coefficient of 0.6. Apart from its relation to the franchise value, the intermediation cost is important in its own right since it represents the component of net worth of banks that is unrelated to the core operations. Since the profit from core banking activities and credit provision would likely be low in bad states of the economy, the counter-cyclical nature of intermediation cost is particularly concerning since it would further reduce the net worth of financial institutions.

Figure 3.1 – Franchise value of BHCs in the US between period 2002Q1 and 2020Q4. Franchise value is computed as Market value of equity + Book value of liabilities - (Book value of asset - Goodwill) scaled by (Book value of assets - Goodwill). The data is at quarterly frequency and winsorized at 1% level. The data source is given in Table (3.1).

3.1 Empirical study

In this section, I provide motivating evidence related to the intermediation cost and franchise value of BHCs in the US.

3.1.1 Data

I use a number of data sources for the empirical analysis. The balance sheet and income statement data for the banks are taken from Federal Reserve FR Y-9C reports that are available at quarterly frequency for all BHCs in the US. Table (3.1) shows the variable list pertaining to the FR Y-9C data. Each bank is identified by a unique regulatory identification number

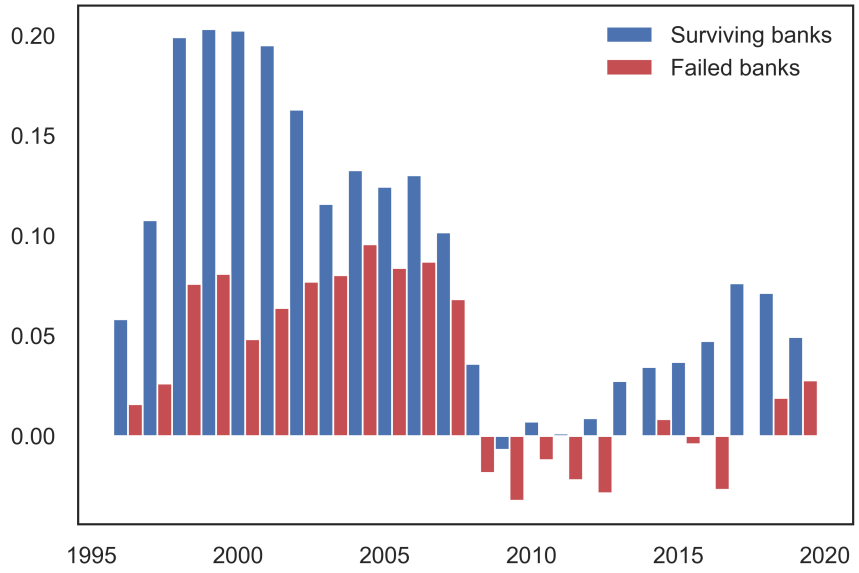

Figure 3.2 – Franchise value of BHCs in the US. Blue bars correspond to banks that did not file for bankruptcy or receive assistance from FDIC, and red bars correspond to the banks that filed for bankruptcy during the sample period. Franchise value is computed as Market value of equity + Book value of liabilities - (Book value of asset - Goodwill) scaled by (Book value of assets - Goodwill). The data is at quarterly frequency. The data source is given in Table (3.1).

(RSSD ID). I merge this with the CRSP universe that contains the market value information for all firms in the US using the PERMCO-RSSD link.[4] The market value is used to calculate the franchise value of banks. List of bank failures are taken from FDIC that includes banks that received assistance transaction from the FDIC.[5] The final data sample is an unbalanced panel of 1,671 publicly traded banks at the Bank Holding Company level at a quarterly frequency from the year 1976 till 2020. Since the data to compute deposit rate is available only from the year 2001 till 2018, I restrict the empirical analysis to this time period. Table (3.2) presents the summary statistics of the variables used in the study. The data to estimate equity risk premium is taken from Welch and Goyal (2008) and Robert Shiller's website.[6]

3.1.2 Franchise Value

Franchise value refers to the present value of future profits that the banks are expected to earn as a going concern. Banks can derive franchise value due to external factors such as the anti-

[4]Source: https://www.newyorkfed.org/research/banking_research/datasets.html.
[5]Source: https://www.fdic.gov/resources/resolutions/bank-failures/failed-bank-list.
[6]Source: http://www.econ.yale.edu/ shiller/data.htm.

Variable	Description	Time coverage
RCON2170	Total Assets	197603- 202012
RCON2200	Total Deposits	197603- 202012
RCON2948	Total Liabilities	197603- 202012
RCFD2146/RCFD3545	Trading assets	198403-202012
RIAD4073	Interest expense	198403-202012
RIAD4093	Non interest expense	198403- 202012
RIAD4000	Operating income	198403-202012
RIAD4074	Interest income	198403- 202012
RCON3163	Goodwill	200103 -201812
RIAD4135	Salaries and employee benefits	200203- 201812
RIAD4217	Premises, furniture etc.	200203- 201812
RIADC232	Amortization and impairment expenses	200203- 201812
RIADC216	Goodwill impairment	200203- 201812
UBPRD488	Risk weighted capital to asset ratio	200103- 202012
RIAD0093	Interest expense savings deposit	200103- 202012
RCONB563	Quarterly average savings deposit	200103- 202012

Table 3.1 – Description of variables and coverage. All variables are taken from FFIEC CDR Call reports except for Risk weighted capital to asset ratio which is taken from FFIEC CDR UBPR Ratios Capital Analysis report.

competitive measures imposed by the regulators. For example, the Douglas Amendment Act of 1956 prohibited banks from acquiring other banks that operated out of the state. Regulation *Q* allowed banks to maintain a low deposit rate regardless of the movement in the Fed funds rate, until it was repealed in last 1970s (Drechsler et al. (2021)). In addition, regulatory laws such as inter-state and intra-state branching restrictions gave banks market power that allowed them to charger lower deposit rates on their liabilities. These restrictions remained in place until late 1990s. The pass-through of market rates to the deposit rates is shown to be about 0.4, far from the value of 1 that one would expect if the markets were perfectly competitive (Drechsler et al. (2021)). As these regulatory restrictions were gradually lifted, the franchise value of banks started to decline (Demsetz et al. (1996)).

Banks can also derive franchise value by possessing superior technology, labor, and capital. Jayaratne and Strahan (1998) document that while the franchise value of all banks declined after anti-competitive measures were removed by the regulators, better-managed banks grew more at the expense of poorly managed banks. To capture the bank-specific factors, I define franchise value as the difference between the market value and the replacement cost following Demsetz et al. (1996). I restrict the sample to banks that are publicly traded and whose share price and number of common equity shares outstanding are available in CRSP database to obtain the market value of equity. I add to this the book value of liabilities from Compustat to obtain a proxy for the market value of banks. The replacement cost is proxied by the difference

	All periods				2001-2018			
	mean	std	min	max	mean	std	min	max
Total Assets (Millions)	0.25	0.76	0.00	6.12	6.12	23.01	0.02	184.93
Deposit/Total assets	0.33	0.12	0.01	0.68	0.39	0.16	0.00	0.74
Trading assets/Total assets	0.00	0.00	0.00	0.01	0.00	0.00	0.00	0.04
Operating income/assets (%)	8.22	2.60	3.13	16.94	6.41	2.43	1.84	21.18
Operating expense/Total assets (%)	3.23	1.95	0.12	7.53	1.62	1.09	0.00	4.36
Non-operating expense/Total assets (%)	3.36	1.47	1.19	11.09	3.06	1.84	0.13	16.10
Leverage	0.91	0.05	0.71	1.06	0.90	0.08	0.29	1.05
CIR (%)	0.37	0.17	0.03	0.69	0.25	0.14	0.01	0.61
Deposit rate (%)	1.10	0.99	0.04	4.40	1.16	1.02	0.01	4.43
ROA (%)	0.82	1.30	-6.23	3.73	0.84	1.29	-6.69	4.50
log(z-score)	4.42	1.04	1.30	6.44	4.81	1.03	1.59	6.63

Table 3.2 – Descriptive Statistics of key variables. All periods refer to time period for which coverage is available, as described in Table (3.1). All variables are winsorized at 1% level. The variables in percentages denote annualized figures.

between book value of assets and goodwill.[7]. To facilitate comparison across banks, I scale the proxy by asset value less goodwill. Figure (A.1) shows the histogram of franchise value pooled across all banks and across all years. The distribution is skewed to the right, with an average value of 5%. From Figure (3.2) that plots the time series average, we see that the franchise value was high until the late 1990s, consistent with the explanation that anti-competitive regulatory measures enabled the banks to use their market power and extract rents from deposits. From the year 1996 till the Global Financial Crisis of 2008, the franchise value of BHCs averaged at 10%, and decreased to an average of 2.7% in the period 2009-2020. Strikingly, the banks that failed or sought assistance from FDIC have a lower franchise value than those that did not, and this pattern is consistent every year from 2001 till 2020 as seen in Figure (3.2). In fact, after the year 2008, the failed banks had a negative franchise value on average until the year 2016. Even during the Great Financial Crisis, the surviving banks managed to have a positive franchise value on average, albeit close to zero. This calls for further investigation on the origins of franchise value and its impact on bank default, which I address next.

3.1.3 Intermediation cost

I define intermediation cost to consist of all non-operating expenses that banks incur. These non-operating expenses pertain to day-to-day administrative tasks, employee compensation, premises and fixed-asset expenses, goodwill impairment, and 'other' non-operating expenses that includes information technology and litigation costs. To facilitate comparison across banks and time, I scale the non-operating expenses to the total assets. The motivation for

[7] This closely follows Demsetz et al. (1996). The motivation for this proxy is that a component of bank's franchise value is the difference between the book value of an asset, and the price that the bank pays to acquire the asset.

doing this is simple- a bank with more number of branches may have higher total costs but may still have lower average cost due to economies of scale and scope. The non-operating expense ratio is quite stable over the past four decades, with an annualized average of 3.8% and standard deviation of 6% in the time sample considered. The increase in cost during the late 1990s could be attributed to a series of bank mergers. For instance, the merger of Wells Fargo with Norwest, SunTrust with Crest, and Bank of America with Nations' Bank increased their costs by at least two times (Jaremski and Sapci (2017)). Secondly, this cost has a correlation of 57% with the deposit rate, and 13% with the Federal funds rate. On the contrary, the operating expense of banks vary a lot over time, and unsurprisingly has a higher correlation with the deposit rate (73%). Around 70% of the bank liabilities are in the form of retail (core) deposits, and thus the overall interest expense moves with the rate on deposit charged by the banks. Any difference between these two can be attributed to the costs involved in monitoring and servicing the loans. Therefore, I construct the intermediation cost measure ($iCost$) by adding the non-operating expense and the residual from netting out the deposit rate from operating expense. Figure (3.3) shows the $iCost$ measure along with the deposit rate. There are two key take-aways from the figure. The first observation is that the $iCost$ measure is *counter-cyclical*, with the cost raising up during the onset of Global Financial Crisis in the year 2008, and dropping down in the subsequent years after the recovery. During crisis, the uncertainty in collateral value and the possibility of borrower default increases, forcing banks to spend more money on monitoring and litigation costs (Ben S. Bernanke (1983)). As the economy recovers out of crisis, these costs go down since the uncertainty over the collateral value declines bringing down the monitoring effort.

Figure (3.4) shows the cyclical patterns of $iCost$ along with the real GDP of the US economy obtained from applying Hodrick-Prescott filter to the time series data. We see that the intermediation the cost rises during the Global Financial Crisis of 2008, and drops sharply after year 2010. The cyclical component of the cost has a correlation of -0.30 with the cyclical component of the US Real GDP between the period 1986Q1 to 2020Q4, that increases in magnitude to -0.57 in the period 2001Q1 to 2020Q4. The second key observation is that the intermediation cost is highly *persistent*. Empirically, after the spike in costs during the Global Financial Crisis, it took around four years for the cost to revert to its pre-crisis level.

Since a major component of the intermediation cost is non-operating expense, I study its individual components by analyzing the FR Y-9C reports that documents five main categories of non-operating expenses. Figure (A.3) shows the time series of the individual components. The first main component is *employee compensation*, that accounts for 51% of the total non-operating expenses during the period 2001-2018. This category includes the bonus compensation, health insurance, and retirement plans. *Premises and fixed-assets* related expenses account for 12% of the total costs. This category includes depreciation, repairs, equipment and furniture, and mortgage interest on real estate. *Goodwill and amortization* expenses are negligible part ($< 0.1\%$) of the overall non-operation expenses. The remaining part, of about 35%, includes information technology costs, legal fees, advertising and stationary expenses, and so on.

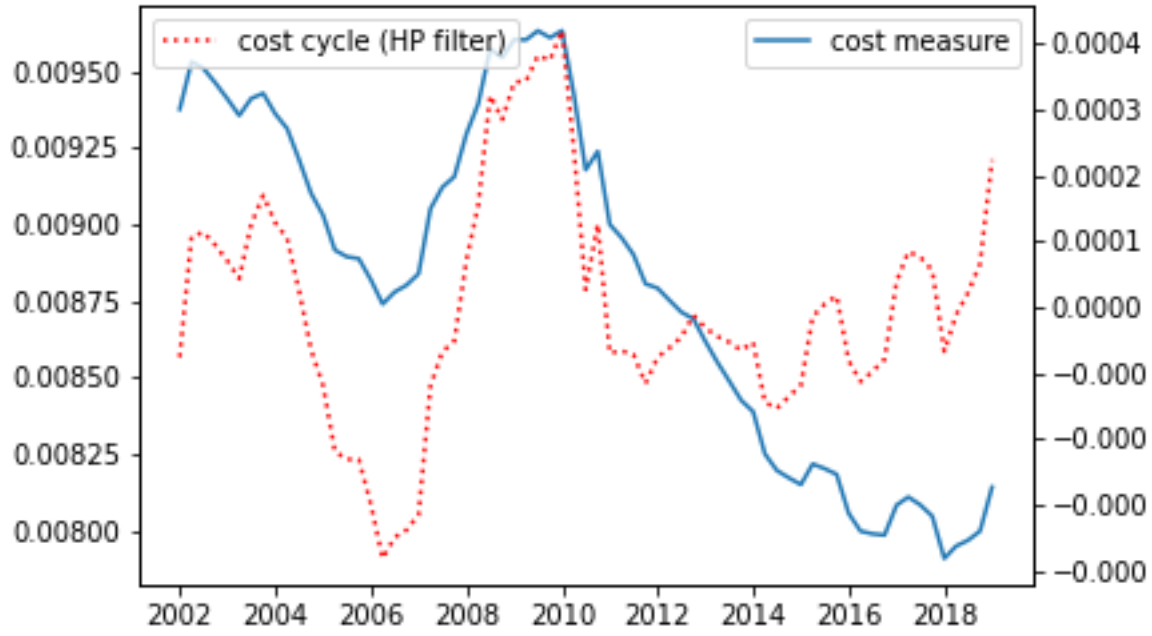

Figure 3.3 – This figure presents the time series of intermediation cost measure ($iCost$) for the banks and the deposit rate. The deposit rate is averaged across all BHCs in the US. The data is at quarterly frequency between the period 1986Q1 to 2020Q4. The data source is given in Table (3.1). The values are annualized and is in percentage terms. The shaded background represent NBER recessionary periods.

3.1.4 Determinants of Franchise value

To understand the determinants of franchise value of BHCs in the US, I run the following regression

$$f_{i,t} = \alpha_t + \beta_1 iCost_{i,t} + \Gamma X_{i,t} + \epsilon_{i,t} \tag{3.1}$$

where $f_{i,t}$ is the franchise value of bank i at time t, $iCost_{i,t}$ is the intermediation cost and $X_{i,t}$ are various controls. I control for time fixed effects in order to exploit only cross-sectional variation across the banks. The results are shown in Table (3.3). In the first specification (I), I regress the franchise value only on log of total assets and find that banks with larger asset size have a higher franchise value, pointing to economies of scale effect. Although the statistical significance is strong, the R-squared is low indicating that bank size may not be the sole determinant of franchise value. In the specification (II), I include as regressors the intermediation cost along with other variables that capture the business model of banks such as the ratio of deposit to total assets, ratio of domestic loans to total assets, and leverage of the bank. The asset size is still a determinant of franchise value and the specification has a higher R-squared value. Importantly, I find statistically significant evidence that the banks

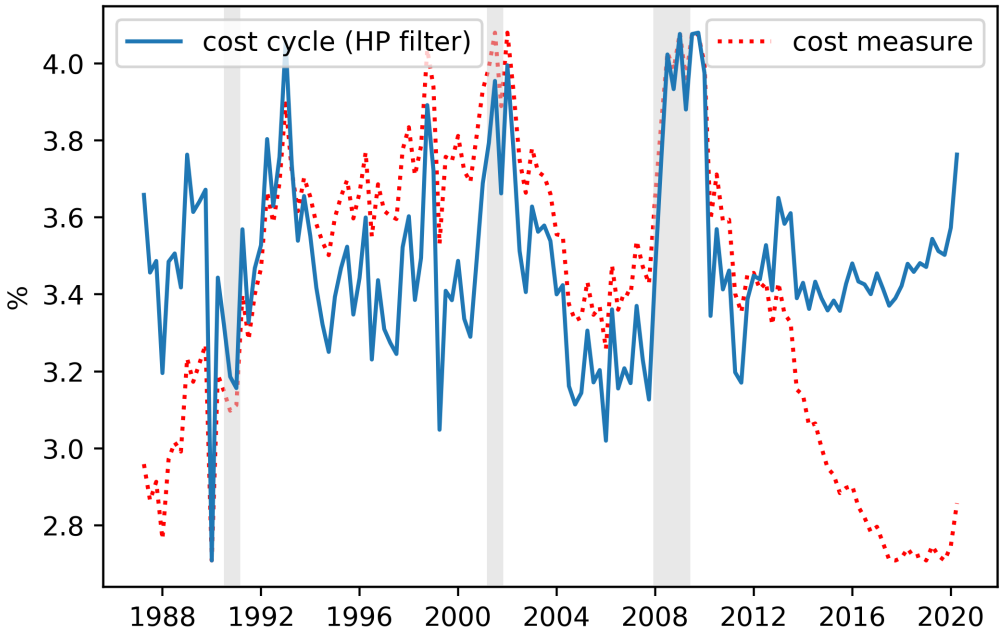

Figure 3.4 – This figure presents $iCost$ along with its cyclical component after applying Hodrick-Prescott filter. The data is at quarterly frequency between the period 1986Q1 to 2020Q4. The data source is given in Table (3.1). The shaded background represent NBER recessionary periods.

with higher intermediation cost have a lower franchise value. In fact, the economic magnitude of this relationship is predominantly strong. A 1 basis point increase in intermediation cost is associated with a 85 basis points drop of the franchise value, which provides evidence for a hypothesis that franchise value of banks originate from cost efficient intermediation of assets. The controls *deposit-asset* ratio and *domestic asset-total asset* ratio capture the business model of banks, and both have statistically significant relationship with the franchise value. A 1% increase in deposit share is associated with an increase in the franchise value of 7.15%. This is consistent with the finding by Egan et al. (2021) who show that deposit productivity is responsible for a majority of the cross-sectional value of banks in th US. The bank leverage, which is a commonly used variable to explain cross-sectional asset prices in intermediary based asset pricing models, is also significant in explaining the franchise value. The control variable *deposit power* is a proxy that captures the rents that banks derive from utilizing their market power as depository institutions.[8] The variable *government guarantee* is a proxy to capture the rent derived from guarantees including the implicit too-big-to-fail related ones

[8]The proxy is computed as $\max(0, r_t^f - r_{i,t}^d)\frac{D_{i,t}}{L_{i,t}}$ where r^f is the risk-free rate, r_i^d is bank i's deposit rate, D_i and L_i is the total deposits and liabilities of the bank i.

from the government.[9] Both of these control variables are positively related to the franchise value, implying that rents from deposit pricing power and goverment guarantees are important determinants of franchise value. The latter conclusion resonates with Atkeson et al. (2018) who find empirically that these guarantees account for a large part of market value of banks in the US.

The specification (III) in Table (3.3) includes a dummy variable that takes a value 1 if a bank i filed for bankruptcy during the time period considered in the sample, and zero otherwise. The coefficient on the dummy is both economically and statistically significant. Compared to the banks that failed, the banks that survived had 3% higher franchise value on average controlling for business model and asset size. The last specification in Table (3.3) includes capital ratio, which is a proxy for risk. This variable refers to the total risk-based capital to risk-weighted assets reported by banks in the FFIEC report. As expected, banks with better capital ratio have a higher franchise value, although both economic and statistical significance is weak.[10] Including the capital ratio as a control does not change the significance of intermediation cost, deposit share, and the failure dummy while the asset size is no longer significant in explaining the franchise value. One should be cautious in deriving causal relationship from the results in Table (3.3) since it could be that the franchise value is the main driver of intermediation cost and deposit share instead of the other way around. To test if there is a predictive relationship, I run the regression of the form

$$f_{i,t} = \alpha_i + \beta_1 iCost_{i,t-1} + \Gamma X_{i,t-1} + \epsilon_{i,t} \tag{3.2}$$

where variables $X_{i,t-1}$ correspond to the lagged values with a lag of one quarter. I use bank fixed effects to exploit only the time variation. By including a variety of control variables to capture the business model, risk, size and leverage, any concerns related to omitted variables are alleviated. The results are shown in Table (A.3). There is a predictive effect of intermediation cost on the franchise value: one basis point increase in cost leads to a drop of around 57 basis points of franchise value. Business model of the banks proxied with share of deposit to assets, and the government guarantees also have a predictive relationship. The main take-away is that the intermediation cost remains to be both statistically and economically significant in driving the franchise value.

3.1.5 Franchise value and bank exit

I formally test the stylized fact from Figure (3.2) that lower franchise value is associated with bank failures by running the following regression

$$\log z_{i,t} = \alpha_t + \beta_1 f_{i,t} + \Gamma X_{i,t} + \epsilon_{i,t} \tag{3.3}$$

[9] The computation of this proxy follows De Nicolo and Zotova (2020) and captures the per-dollar saving realized by the bank in the absence of premium on government guarantees.

[10] The number of observations reduce drastically in this specification since capital ratio is available only at annual frequency (the end of calendar year).

Table 3.3 – This table estimates the determinants of franchise value by regressing $f_{i,t}$ on $iCost/assets$ and various controls. The data is for all listed US Bank Holding companies from the year 2001 till 2018 at quarterly frequency. All variables (except dummy variables) are winsorized at 1% level. Standard errors are robust to heteroskedasticity in errors.

	(1) $f_{i,t}$	(2) $f_{i,t}$	(3) $f_{i,t}$	(4) $f_{i,t}$
logAssets	0.00620***	0.00419***	0.00456***	0.00559***
	(0.000)	(0.000)	(0.000)	(0.000)
iCost/assets		-0.858***	-0.832***	-1.346***
		(0.000)	(0.000)	(0.000)
deposit/assets		0.0715***	0.0653***	0.0527***
		(0.000)	(0.000)	(0.000)
domestic/assets		-0.0783***	-0.0702***	-0.0520***
		(0.000)	(0.000)	(0.000)
trading/assets		-0.405***	-0.308***	-0.181
		(0.000)	(0.001)	(0.288)
interest income/operating income		0.0110	0.0107	0.0336***
		(0.103)	(0.101)	(0.001)
leverage		-0.225***	-0.229***	-0.201***
		(0.000)	(0.000)	(0.000)
deposit power		2.000***	1.777***	2.726***
		(0.000)	(0.000)	(0.002)
govt. guarantee		1.516***	1.512***	1.521***
		(0.000)	(0.000)	(0.000)
failure dummy			-0.0347***	-0.0323***
			(0.000)	(0.000)
capital ratio				0.00170***
				(0.001)
_cons	-0.0419***	-1.242***	-1.240***	-1.338***
	(0.000)	(0.000)	(0.000)	(0.000)
N	23925	23925	23925	5154
Fixed Effects	Time	Time	Time	Time
R^2	0.037	0.475	0.487	0.502

p-values in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

where $z_{i,t}$ is the z-score of the bank i at time t, $f_{i,t}$ is the franchise value, and $X_{i,t}$ are the control variables. The z-score is computed following the literature as[11]

$$z_{i,t} = \frac{(Equity/Asset)_{i,t} + \overline{ROA}_{i,t}}{\sigma(ROA)_{i,t}} \tag{3.4}$$

where the $\overline{ROA}$ and $\sigma(ROA)$ are the rolling 12-quarter mean and standard deviation of the return on assets. A larger z-score value indicates a lower probability of default. Figure (A.4) shows the franchise value of banks with z-score values in the bottom and top quartiles. During the Global Financial Crisis of 2008 and the years thereafter, banks with higher (lower) z-score values had a larger (smaller) franchise value. In some periods during the late 1990s, high risk banks had a higher franchise value. One explanation for this could be that the anti-competitive regulations such as branching laws favored poorly managed banks with a larger probability of default in enforcing market power and deriving franchise value. Once these laws were lifted, cost/profit efficiency through good management practices was the only source of franchise value left for the banks. This enabled the better managed banks to derive a higher franchise value compared to the poorly managed banks, although the overall franchise value trended down due to abolishing of the branching laws.

Specification (I) in Table (3.4) regresses the z-score on assets without any controls. There is evidence of scale effects- that is larger asset size is associated with a lower probability of default. From the specification (II), we see that banks with a larger franchise value is associated with a lower default probability with a coefficient that is economically large. Banks that have a higher share of deposits, in particular foreign deposits, are less risky. Interestingly, once we control for the franchise value and business model of the banks, the scale effect disappears. In fact, conditional on these control variables, the smaller firms are less risky compared to the larger firms in terms of asset size. For robustness, I repeat the analysis by using the data on actual bank failures published by the FDIC. Over the time period 2001 to 2019, a total of 89 banks failed, and 12 banks required assistance from the FDIC in order to prevent bankruptcy. I replace the z-score by a dummy value d_i that takes value 1 if the bank i ever filed for bankruptcy or required assistance in the time period considered, and 0 otherwise. While the z-score is a proxy for the probability of default, this dummy corresponds to the actual distress event, and hence can be used to more accurately estimate the failure probability. Table (A.4) shows the results. The franchise value is statistically significant in explaining the probability of default, along with asset size and deposit share. The result is the same if the dummy variable d_i is replaced with a value 1 if the bank i ever sought assistance from the FDIC to continue operations, and 0 otherwise. Again, caution should be applied in deriving causal relationship between the franchise value and probability of failure. The results rather point to statistical correlations, that is robust to using different proxies for probability of failure.

[11] See Berger et al. (2016), Caiazza et al. (2018), Pino and Sharma (2019), Delis et al. (2014) among others.

Table 3.4 – This table studies the relationship between bank probability of default and franchise value. The probability of default is proxied by z-score computed using (3.4). The data is for all listed US Bank Holding Companies from the year 2001 till 2018 at quarterly frequency. All variables are winsorized at 1% level. Standard errors are robust to heteroskedasticity in errors.

	(1)	(2)	(3)
	$\log z_{i,t}$	$\log z_{i,t}$	$\log z_{i,t}$
logAssets	0.0295***	-0.0304**	-0.0227
	(0.000)	(0.012)	(0.217)
franchise value		3.609***	4.833***
		(0.000)	(0.000)
deposit/assets		0.617***	0.565**
		(0.000)	(0.016)
domestic assets/assets		-0.454**	-0.656
		(0.040)	(0.102)
interest income/operating income		0.0887	0.410
		(0.413)	(0.168)
leverage		-0.500**	1.060***
		(0.025)	(0.007)
capital ratio			0.0124
			(0.370)
_cons	4.184***	5.471***	3.717***
	(0.000)	(0.000)	(0.000)
N	31494	23906	5150
Fixed effect	Time	Time	Time
R^2	0.002	0.049	0.079

p-values in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

3.2 Conclusion

In this chapter, I show that there is a strong empirical support for the franchise value of banks to be related with bank failures. Analyzing a large panel of Bank Holding Companies (BHCs) in the US, I show that increasing franchise value is associated with a larger z-score and lower probability of failure. The BHCs derive franchise value from being intermediation cost efficient, along with the business model, and rents from deposits and government guarantees. Most of the regulation post financial crisis entails improving the capital ratio of banks to reduce systemic risk. While these regulations are certainly in the right direction in improving the balance sheet of intermediaries, the declining franchise value in the present day makes a case against complacency that systemic risk is fully eliminated.

A Appendix

A.1 Appendix to Chapter 1

A.1.1 Model with stochastic productivity

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 (1.5). Since $\xi_e v_t$ is a martingale, the drift equals to zero, which implies

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

The households do not issue outside equity but are exposed to the risk from experts through the equity issuance of the latter. Following similar steps, we get the asset pricing condition for the households as

$$\mu_{h,t}^R - r_t = \epsilon_{h,t}$$

where $\epsilon_{h,t} = \zeta_{h,t}^k(\sigma + \sigma_t^{q,t}) + \zeta_{h,t}^a\sigma_t^{q,a} + \varphi(\zeta_{e,t}^a(\sigma + \sigma_t^{q,k}) + \zeta_{h,t}^k\sigma_t^{q,a})$ ■

Proof of Proposition 1

The law of motion of wealth for the experts and the households are given in the optimization problems (1.7) and (1.8) respectively. Using the law of large numbers to aggregate the wealth of individual household and expert, we get

$$\frac{dW_{h,t}}{W_{h,t}} = \left(r_t - \rho_h - \lambda_d + \theta_{h,t}(\mu^R_{h,t} - r_t) + \frac{(1-\bar{z})\lambda_d}{1-z_t} + \tau_t \frac{W_{e,t}}{W_{h,t}}\right)dt + \theta_{h,t}(\sigma + \sigma^q_t)dZ^k_t + \theta_{h,t}\sigma^a_t dZ^a_t$$
$$\frac{dW_{e,t}}{W_{e,t}} = \left(r_t - \rho_e - \lambda_d + \theta_{e,t}\epsilon_{e,t} + \frac{\bar{z}\lambda_d}{z_t} - \tau_t\right)dt + \theta_{e,t}(\sigma + \sigma^{q,k}_t)dZ^k_t + \theta_{e,t}\sigma^{q,a}_t dZ^a_t$$

where $W_{h,t} = \int_{j\in H} w_{j,t}dj$ and $W_{e,t} = \int_{j\in E} w_{j,t}dj$ denotes aggregated wealth among respective group, $z_t = \frac{W_{e,t}}{W_{h,t}+W_{e,t}} = \frac{W_{e,t}}{q_t K_t}$, and $\theta_{e,t} := \frac{\chi_t \psi_t}{z_t}, \theta_{h,t} := \frac{1-\chi_t\psi_t}{1-z_t}$ from the capital market clearing condition.[1] The terms containing λ_d and $\bar{z}$ are due to the overlapping generations assumption, and the terms with τ_t is due to the exit of the experts. 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[2]

$$\frac{dK_t}{K_t} = (\phi(\iota_t) - \delta)dt + \sigma dZ^k_t$$

Applying Ito's lemma, we get

$$\frac{d(q_t K_t)}{q_t K_t} = (\epsilon_{e,t}\chi_t + (1-\chi_t)\epsilon_{h,t} - \frac{(a_{e,t}-\iota_t)}{q_t} + r_t)dt + (\sigma + \sigma^{q,k}_t)dZ^k_t + \sigma^{q,a}_t dZ^a_t$$
$$\frac{d\langle q_t K_t, q_t K_t\rangle}{(q_t K_t)^2} = ((\sigma^{q,k}_t + \sigma)^2 + (\sigma^{q,a}_t)^2 + 2\varphi(\sigma^{q,k}_t + \sigma)\sigma^{q,a}_t)dt$$
$$\frac{d\langle q_t K_t, W_{e,t}\rangle}{q_t K_t W_{e,t}} = \left(\theta_{e,t}(\sigma^{q,k}_t + \sigma)^2 + \theta_{e,t}(\sigma^{q,a}_t)^2 + 2\varphi(\sigma^{q,k}_t + \sigma)\sigma^{q,a}_t\right)dt$$

and the result follows from here after some algebra. ■

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

[1] Note that $z_t = \frac{W_{e,t}}{q_t K_t}$ and $\psi_t = \frac{K_{e,t}}{K_t}$. Moreover, $\sigma_{w_e,t}(\sigma + \sigma^q_t)z_t + \sigma_{w_h,t}(\sigma + \sigma^q_t)(1-z_t) = (\sigma + \sigma^q_t)$ and similarly for $\sigma^{q,a}_t$. Using these, we can relate $\sigma_{w_j,t}$ to $\theta_{j,t}$.

[2] Since the investment rate is the same for all agents, the evolution of the aggregate capital K_t is the same as the evolution of $k_{j,t}$. To see this, write $\frac{dK_t}{K_t} = \frac{dK_{e,t}}{K_t} + \frac{dK_{h,t}}{K_t} = \psi_t \frac{dK_{e,t}}{K_{e,t}} + (1-\psi_t)\frac{dK_{h,t}}{K_{h,t}}$ and the rest follows from (2.3).

policies and using law of large numbers as

$$\frac{dW_{e,t}}{W_{e,t}} = \big(r_t - \rho - \lambda_d + \frac{\psi_t}{z_t}(\mu^R_{e,t} - r_t) - (1-\chi_t)\frac{\psi_t}{z_t}\epsilon_{h,t} + \frac{\bar{z}\lambda_d}{z_t} - \tau(a_{e,t}, z_t)\big)dt + \frac{\chi_t\psi_t}{z_t}(\sigma + \sigma^{q,k}_t)dZ^k_t + \frac{\chi_t\psi_t}{z_t}\sigma^{q,a}_t dZ^a_t$$

Proof of Proposition 2

The value function conjecture is

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

where $J_{j,t}$ follows the stochastic differential equation $\frac{dJ_{j,t}}{J_{j,t}} = \mu^J_{j,t}dt + \sigma^{J,k}_{j,t}dZ^k_t + \sigma^{J,a}_{j,t}dZ^a_t$ whose drift and volatility needs to be determined in the equilibrium. The HJB equation is given by

$$\sup_{C,K} f(C_{j,t}, U_{j,t}) + E[dU_{j,t}] = 0 \tag{A.1}$$

where $f(C_{j,t}U_{j,t}) = (1-\gamma)\rho U_{j,t}\Big(\log C_{j,t} - \frac{1}{1-\gamma}\log\big((1-\gamma)U_{j,t}\big)\Big)$. The HJB equation is derived directly in terms of the aggregate capital K_t instead of the wealth share z_t. For ease of notation, I will denote the wealth share of the experts and households as $z_{e,t}$ and $z_{h,t}$ respectively but it is to be understood that $z_{e,t} = z_t$ and $z_{h,t} = 1 - z_t$. The value function derivatives are

$$\frac{\partial U_{j,t}}{\partial J_{j,t}} = K_t^{1-\gamma}J_{j,t}^{-\gamma}; \quad \frac{\partial U_{j,t}}{\partial K_t} = J_{j,t}^{1-\gamma}K_t^{-\gamma} \tag{A.2}$$
$$\frac{\partial^2 U_{j,t}}{\partial J^2_{j,t}} = -\gamma K_t^{1-\gamma}J_{j,t}^{-\gamma-1}; \quad \frac{\partial^2 U_{j,t}}{\partial K^2_t} = -\gamma J_{j,t}^{1-\gamma}K_{j,t}^{-(1+\gamma)}; \quad \frac{\partial^2 U_{j,t}}{\partial J_{j,t}\partial K_t} = (1-\gamma)(K_tJ_{j,t})^{-\gamma}$$

Applying Ito's lemma to $U_{j,t}$ and using HJB equation (A.1), we get

$$\begin{aligned}\sup_C \quad & \rho(J_{j,t}K_t)^{1-\gamma}[\log\frac{C_{j,t}}{W_{j,t}} - \log J_{j,t} + \log(q_t z_{j,t})] + (J_{j,t}K_t)^{1-\gamma}(\Phi(\iota) - \delta) \\ & - \frac{\gamma}{2}(J_{j,t}K_t)^{1-\gamma}\sigma^2 + (J_{j,t}K_t)^{1-\gamma}\mu^J_{j,t} - (J_{j,t}K_t)^{1-\gamma}\frac{\gamma}{2}((\sigma^{J,k}_{j,t})^2 + (\sigma^{J,a}_{j,t})^2 + 2\varphi\sigma^{J,k}_{j,t}\sigma^{J,a}_{j,t}) \\ & + (1-\gamma)(J_{j,t}K_t)^{1-\gamma}(\sigma\sigma^{J,k}_{j,t} + \varphi\sigma\sigma^{J,a}_{j,t}) + \tau_t(U_{h,t} - U_{e,t}) = 0\end{aligned} \tag{A.3}$$

Writing the value function expression in terms of the wealth, we have

$$U_{j,t} = \frac{(\tilde{J}_{j,t}W_{j,t})^{1-\gamma}}{1-\gamma}; \qquad f(C_{j,t}, U_{j,t}) = (1-\gamma)\rho U_{j,t}(\log\frac{C_{j,t}}{W_{j,t}} - \tilde{J}_{j,t}) \tag{A.4}$$

where $\tilde{J}_{j,t} = \frac{J_{j,t}}{q_t z_{j,t}}$ and $z_{j,t} = \frac{W_{j,t}}{q_t K_t}$ are used to obtain (A.4). At the optimum, the marginal utilities of wealth and consumption become equal. Therefore,

$$\frac{\partial U_{j,t}}{\partial W_{j,t}} = \frac{\partial f_{j,t}}{\partial C_{j,t}}$$
$$\tilde{J}_{j,t}^{1-\gamma} W_{j,t}^{-\gamma} = (1-\gamma)\rho \frac{U_{j,t}}{C_{j,t}} \implies \frac{C_{j,t}}{W_{j,t}} = \rho$$

This proves the optimal consumption policy. The stochastic discount factor for recursive utility is given by

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

From (A.4), we get

$$\xi_{j,t} = (1-\gamma)\exp\left(\int_0^t \left[(1-\gamma)\rho\left(\log\rho - \tilde{J}_{j,t}\right)\right] ds\right) \frac{U_{j,t}}{W_{j,t}}$$

This implies that $\sigma(\xi_{j,t}) = \sigma\left(\frac{U_{j,t}}{W_{j,t}}\right)$. To compute the R.H.S., we have to obtain $d\left(\frac{U_{j,t}}{W_{j,t}}\right)$. Let $v(J_{j,t}, z_{j,t}, q_t, K_t) := \frac{U_{j,t}}{W_{j,t}}$. Using the derivatives

$$\frac{1}{v}\frac{\partial v}{\partial J_{j,t}} = \frac{1-\gamma}{J_{j,t}}; \quad \frac{1}{v}\frac{\partial v}{\partial z_{j,t}} = -\frac{1}{z_{j,t}}$$
$$\frac{1}{v}\frac{\partial v}{\partial q_t} = -\frac{1}{q_t}; \quad \frac{1}{v}\frac{\partial v}{\partial K_t} = \frac{1-\gamma}{K_t}$$

and applying Ito's lemma, we get

$$\frac{dv}{v} = \underbrace{[\ldots\ldots]}_{\text{drift term}} dt + (1-\gamma)(\sigma_{j,t}^{J,k} dZ_t^k + \sigma_{j,t}^{J,a} dZ_t^a) - (\sigma_{j,t}^{z,k} dZ_t^k + \sigma_{j,t}^{z,a} dZ_t^a) \tag{A.5}$$
$$- ((\sigma + \sigma_t^{q,k}) dZ_t^k + \sigma_t^{q,a} dZ_t^a) + (1-\gamma)\sigma dZ_t^k$$

Applying Ito's lemma to $J_{j,t}(z_t, a_{e,t})$, we have

$$dJ_{j,t} = \frac{\partial J_{j,t}}{\partial z_t} dz_t + \frac{\partial J_{j,t}}{\partial a_{e,t}} da_{e,t} + \frac{1}{2}\frac{\partial^2 J_{j,t}}{\partial z_t^2} d\langle z, z\rangle_t + \frac{1}{2}\frac{\partial J_{j,t}}{\partial a_{e,t}^2} d\langle a_e, a_e\rangle_t$$
$$= (\text{drift terms}) + \frac{\partial J_{j,t}}{\partial z_t} z_t\left(\sigma_t^{z,k} dZ_t^k + \sigma_t^{z,a} dZ_t^a\right) + \frac{\partial J_{j,t}}{\partial a_{e,t}} \sigma_{ae} dZ_t^a$$

Comparing with the SDE (1.20) and matching the diffusion coefficients, we have

$$\sigma_{j,t}^{J,k} J_{j,t} = \frac{\partial J_{j,t}}{\partial z_t} z_t \sigma_t^{z,k} = \frac{\partial J_{j,t}}{\partial z_t} z_t \left(\frac{\chi_t \psi_t}{z_t} - 1\right)(\sigma + \sigma_t^{q,k})$$
$$\sigma_{j,t}^{J,a} J_{j,t} = \frac{\partial J_{j,t}}{\partial a_{e,t}} \sigma_{ae} + \frac{\partial J_{j,t}}{\partial z_t} z_t \sigma_t^{z,a} = \frac{\partial J_{j,t}}{\partial a_{e,t}} \sigma_{ae} + \frac{\partial J_{j,t}}{\partial z_t} z_t \left(\frac{\chi_t \psi_t}{z_t} - 1\right)\sigma_t^{q,a}$$

Collecting the diffusion terms, using $\sigma^{z,i}_{e,t} = \sigma^{z,i}_t, \sigma^{z,i}_{h,t} = -\frac{z_t}{1-z_t}\sigma^{z,i}_t; i \in \{k,a\}$ in equation (A.5), and comparing it to the SDF equation

$$\frac{d\xi_{j,t}}{\xi_{j,t}} = -r_t dt - \zeta^k_{j,t} dZ^k_t - \zeta^a_{j,t} dZ^a_t$$

we get the desired result. ■

Plugging in the optimal consumption-wealth ratio into the HJB equation (*A*.79), we obtain the expressions for $\mu^J_{j,t}$

$$\begin{aligned}\mu^J_{e,t} &= (\gamma-1)(\sigma\sigma^{J,k}_{e,t} + \varphi\sigma\sigma^{J,a}_{e,t}) - (\Phi(\iota_t)-\delta) - \rho\big(\log\rho - \log J_{e,t} + \log(z_t q_t)\big) \\ &\quad + \frac{\gamma}{2}\left((\sigma^{J,k}_{e,t})^2 + (\sigma^{J,a}_{e,t})^2 + 2\varphi\sigma^{J,k}_{e,t}\sigma^{J,a}_{e,t} + \sigma^2\right) - \frac{\tau_t}{1-\gamma}\left(\left(\frac{J_{h,t}}{J_{e,t}}\right)^{1-\gamma} - 1\right) \qquad \text{(A.6)}\\ \mu^J_{h,t} &= (\gamma-1)(\sigma\sigma^{J,k}_{h,t} + \varphi\sigma\sigma^{J,a}_{h,t}) - (\Phi(\iota_t)-\delta) - \rho\big(\log\rho - \log J_{h,t} + \log((1-z_t)q_t)\big) \qquad \text{(A.7)}\\ &\quad + \frac{\gamma}{2}\left((\sigma^{J,k}_{h,t})^2 + (\sigma^{J,a}_{h,t})^2 + 2\varphi\sigma^{J,k}_{h,t}\sigma^{J,a}_{h,t} + \sigma^2\right)\end{aligned}$$

Proof of Proposition 3

Applying Ito's lemma to $q(z_t, a_{e,t})$, we have

$$dq_t = \frac{\partial q_t}{\partial z_t}dz_t + \frac{\partial q_t}{\partial a_{e,t}}da_{e,t} + \frac{1}{2}\frac{\partial^2 q_t}{\partial z_t^2}d\langle z_t, z_t\rangle + \frac{1}{2}\frac{\partial^2 q_t}{\partial a^2_{e,t}}d\langle a_{e,t}, a_{e,t}\rangle + \frac{\partial^2 q_t}{\partial z_t \partial a_{e,t}}d\langle z_t, a_{e,t}\rangle$$

Matching the drift and the volatility terms, we get

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

where $\sigma_{ae,t} = \nu(\bar{a}_e - a_{e,t})(a_{e,t} - \underline{a}_t)$ 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 (1.18) in the above equation and rearranging, we get the result. ■

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

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

The experts will issue maximum outside equity $\underline{\chi}$ whenever their risk premium is larger than that of households. Thus, we can replace χ_t by $\underline{\chi}$ whenever $\psi < 1$. Plugging in the expression for the return processes from (1.4), and using (1.12), (1.11), and Proposition 2, we get

$$\begin{aligned}\frac{a_{e,t} - a_h}{q_t} = \underline{\chi}\Bigg(& (\underline{\chi}\psi_t - z_t)\big((\sigma^{q,k}_t + \sigma)^2 + (\sigma^{q,a}_t)^2 + 2\varphi(\sigma + \sigma^{q,k}_t)\big) \\ & \times\Bigg((1-\gamma)\left(\frac{\partial J_{h,t}}{\partial z_t}\frac{1}{J_{h,t}} - \frac{\partial J_{e,t}}{\partial z_t}\frac{1}{J_{e,t}}\right) + \frac{1}{z_t(1-z_t)}\Bigg) \\ & \quad + (1-\gamma)\left(\frac{\partial J_{h,t}}{\partial a_{h,t}}\frac{1}{J_{h,t}} - \frac{\partial J_{e,t}}{\partial a_{e,t}}\frac{1}{J_{e,t}}\right)\sigma_{ae,t}(\sigma^{q,a}_t + \varphi(\sigma + \sigma^{q,k}_t))\Bigg)\end{aligned} \tag{A.8}$$

The second condition comes from the goods market clearing

$$\rho q_t = \psi_t(a_{e,t} - \iota_t) + (1-\psi_t)(a_h - \iota_t) \tag{A.9}$$

The third and fourth conditions are the return variance components

$$\sigma^{q,k}_t + \sigma = \frac{\sigma}{1 - \frac{1}{q_t}\frac{\partial q_t}{\partial z_t}(\underline{\chi}\psi_t - z_t)} \tag{A.10}$$

$$\sigma^{q,a}_t = \frac{\frac{1}{q_t}\frac{\partial q_t}{\partial a_{e,t}}\sigma_{ae,t}}{1 - \frac{1}{q_t}\frac{\partial q_t}{\partial z_t}(\underline{\chi}\psi_t - z_t)} \tag{A.11}$$

which are partial differential equations solved using a Newton-Raphson scheme. The algorithm is as follows. 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 three following regions in the state space

- $\psi_t < 1$ and $\chi_t = \underline{\chi}$
- $\psi = 1$ and $\chi_t = \underline{\chi}$
- $\psi = 1$ and $\chi_t > \underline{\chi}$

In the first region, the households also hold capital and hence equation (1.15) holds with equality. In this case, the equations (*A*.8), (*A*.9), (A.10),and (A.11) are used to solve for ψ_t, q_t, $(\sigma + \sigma^{q,k}_t)$, and $\sigma^{q,a}_t$. In the second region, the households do not hold capital and hence the equation (1.15) holds with an inequality. In this case, set $\psi_t = 1$, and use (*A*.8),(*A*.10), (*A*.11), and (*A*.9) to solve for χ_t,q_t, $(\sigma + \sigma^{q,k}_t)$, and $\sigma^{q,a}_t$. If $\chi_t < \underline{\chi}$, then set $\chi_t = \underline{\chi}$, otherwise the third region is entered.

- 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

$$\inf_{z \to 0^+} q_t = \frac{a_h \kappa + 1}{\rho \kappa + 1} \tag{A.12}$$

- For iterations $i > 1, \forall j$, use the discretized versions of the equations $(A.10)$ and $(A.11)$

$$(\sigma^{q,k} + \sigma)_{i,j} = \sigma \left(1 - \frac{1}{q_{i,j}} \left(\frac{q_{i,j} - q_{i-1,j}}{\Delta_i} z_i (\frac{\psi_{i,j}}{z_i} - 1)\right)\right)^{-1} \tag{A.13}$$

$$(\sigma^{q,a})_{i,j} = \left(\frac{q_{i,j} - q_{i,j-1}}{\Delta_j} \sigma_{ae,j}\right) \left(1 - \frac{1}{q_{i,j}} \left(\frac{q_{i,j} - q_{i-1,j}}{\Delta_i} z_i (\frac{\psi_{i,j}}{z_i} - 1)\right)\right)^{-1} \tag{A.14}$$

along with the equations $(A.8)$, and $(A.9)$ to solve for $q_{i,j}, \psi_{i,j}, (\sigma + \sigma^q)_{i,j}, (\sigma^{q,a})_{i,j}$.[3] Note that in this region, $\chi_t = \underline{\chi}$ since the risk premium of experts is larger than that of households. The 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 (A.8), $(A.9)$, $(A.13)$, and $(A.14)$ to solve for $\chi_{i,j}$, $q_{i,j}$, $(\sigma + \sigma^{q,k})_{i,j}$ and $(\sigma^{q,a})_{i,j}$. If $\chi_{i,j} < \underline{\chi}$, set $\chi^*_{i,j} = \underline{\chi}$, otherwise set $\chi^*_{i,j} = \chi_{i,j}$. When $\chi_{i,j} > \underline{\chi}$, the system is in the third region where all capital is held by the experts ($\psi_{i,j} = 1$), and risk is perfectly shared between the experts and the households by setting $\epsilon_{e,t} = \epsilon_{h,t}$. The value of χ^*_t is obtained such that $\chi^*_t = \underset{\chi}{argsolve} \quad \epsilon_{e,t} - \epsilon_{h,t} = 0$. Since the premiums $\epsilon_{e,t}, \epsilon_{h,t}$ depend on the χ_t, I iterate between these two quantities until $|\chi^{new}_t - \chi^{old}_t| < tol$ for some tolerance level.

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

$$\begin{aligned} \mu^J_{j,t} J_{j,t} = \frac{\partial J_{j,t}}{\partial t} + \frac{\partial J_{j,t}}{\partial z_t} \mu^z_t + \frac{\partial J_{j,t}}{\partial a_{e,t}} \mu^a_t + \frac{1}{2} \frac{\partial^2 J_{j,t}}{\partial z_t^2} \left((\sigma^{z,k}_{j,t})^2 + (\sigma^{z,a}_{j,t})^2 + 2\varphi \sigma^{z,k}_{j,t} \sigma^{z,a}_{j,t}\right) + \frac{1}{2} \frac{\partial^2 J_{j,t}}{\partial a_{e,t}^2} \sigma^2_{ae,t} \\ + \frac{\partial^2 J_{j,t}}{\partial z_t \partial a_{a,e}} \left(z_t \sigma^{z,k}_{j,t} \sigma_{ae,t} \varphi + \sigma_a \sigma^{z,a}_{j,t}\right) \end{aligned} \tag{A.15}$$

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 $(A.6)$. The PDEs are solved using the neural network method explained in the second chapter. Using the updated function $J_{j,t}$, the static step is performed again. The procedure is repeated until the function $J_{j,t}$ converges upto a pre-specified tolerance level.

[3] 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.

Table A.1 – Calibration: Benchmark model

	Parameter	Benchmark Model	Model B1	Model B2	Target
Technology	Volatility (σ)	0.06	0.06	0.06	Volatility of risk premium
	Discount rate (experts) (ρ_e)	0.06	0.05	0.05	Literature
	Discount rate (households) (ρ_h)	0.04	0.05	0.05	Literature
	Depreciation rate (δ)	0.02	0.1	0.1	GDP growth
	Investment cost (κ)	5	5	5	Investment-capital ratio
	Expert Productivity (a_e)	0.11	0.2	0.15	Conditional risk premium
	Household Productivity (a_h)	0.03	0.02	-0.03	Consumption-output ratio
	Correlation of shocks (φ)	-	0.5	-	Data
Preference	Utility parameters (γ)	2	5	5	Unconditional risk premium
Demographics	Mean expert mass ($\bar{z}$)	0.1	0.1	0.1	Literature
	Turnover (λ_d)	0.03	0.001	0.001	Literature
Expert productivity	Mean reversion rate (π)	-	0.01	-	Duration of crisis
	Variance (ν)	-	4.2	-	Data
Friction	Equity retention ($\underline{\chi}$)	0.5	0.95	0.95	Literature

Note: Calibrated parameters for the benchmark models along with the target. The benchmark model does not feature stochastic productivity or exit rate. The model B1 considers stochastic productivity but without exit. The model B2 has constant productivity but the experts have a state-dependent exit rate.

A.1.2 Benchmark model

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, the return process can be written as

$$dR_{j,t} = \underbrace{\left(\frac{a_j - \iota_{j,t}}{q_t} + \Phi(\iota_{j,t}) - \delta + \mu_t^q + \sigma\sigma_t^q\right)}_{\mu_{j,t}^R} dt + (\sigma + \sigma_t^q) dZ_t^k \tag{A.16}$$

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

$$\frac{d\xi_{j,t}}{\xi_{j,t}} = -r_t dt - \zeta_{j,t} dZ_t^k \tag{A.17}$$

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[4]

$$\frac{\frac{a_e-\iota_t}{q_t}+\Phi(\iota_t)-\delta+\mu_t^q+\sigma\sigma_{q,t}-r_t}{\sigma+\sigma_{q,t}}=\chi_t\zeta_{e,t}+(1-\chi_t)\zeta_{h,t} \tag{A.18}$$

$$\frac{\frac{a_h-\iota_t}{q_t}+\Phi(\iota_t)-\delta++\mu_t^q+\sigma\sigma_{q,t}-r_t}{\sigma+\sigma_{q,t}}\leq\zeta_{h,t} \tag{A.19}$$

The equality holds in (A.72) if the households own some amount of capital ($\psi_t<1$). The optimal investment rate is the same as before and is given in (1.13). The agents solve

$$\sup_{c_{j,t},\chi_{j,t},k_{j,t}} E_t\left[\int_t^\infty f(c_{j,s},U_{j,s})ds\right] \tag{A.20}$$
$$\text{s.t.}\frac{dw_{j,t}}{w_{j,t}}=(r_t-\frac{c_{j,t}}{w_{j,t}}+\frac{q_tk_{j,t}}{w_{j,t}}(\mu_{j,t}^R-r_t-(1-\chi_{j,t})(\sigma+\sigma_t^q)\zeta_{j',t})dt+\sigma_{w_j,t}(\sigma+\sigma_t^q)dZ_t^k$$

where the aggregator $f(c_{j,s},U_{j,s})$ is given in (1.6) and the index j' denotes the other type of agent. The households do not issue outside equity and hence $\chi_{h,t}=1$. On the other hand, the experts issue outside equity but are constrained to hold at least a fraction $\underline{\chi}$ of equity in their balance sheet. Thus, $\chi_{e,t}\in[\underline{\chi},1]$. Moving forward, I denote $\chi_{e,t}$ as simply χ_t for notation brevity. The expressions for $\sigma_{w_j,t}$ is the same as in the stochastic productivity model given in (1.9) and (1.10). 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. The wealth share is defined as

$$z_t=\frac{W_{e,t}}{W_{e,t}+W_{h,t}}=\frac{W_{e,t}}{q_tK_t}$$

where $W_{e,t}=\int_{\mathbb{E}}w_{j,t}dj$ and $K_t=\int_{\mathbb{E}}k_{j,t}dj+\int_{\mathbb{H}}k_{j,t}dj$. Moving forward, I denote $X_{e,t}$ to mean $\int_{\mathbb{E}}x_{j,t}dj$, and similarly for the households.

Proposition 8. *The law of motion of the wealth share of experts is given by*

$$\frac{dz_t}{z_t}=\mu_t^zdt+\sigma_t^zdZ_t^k \tag{A.21}$$

where

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

Proof: The law of motion of wealth for the households and the experts are given by equation $(A.68)$. Using the law of large numbers to aggregate the wealth of individual household and

[4]This can be proved using the Martingale argument similar to the model with stochastic productivity. See Appendix A.1.2 for the proof.

expert, we get

$$\frac{dW_{h,t}}{W_{h,t}} = \left(r_t - \frac{C_{h,t}}{W_{h,t}} - \lambda_d + \frac{1-\chi_t\psi_t}{1-z_t}(\mu^R_{h,t} - r_t) + \frac{(1-\bar{z})\lambda_d}{1-z_t}\right)dt + \frac{1-\chi_t\psi_t}{1-z_t}(\sigma + \sigma^q_t)dZ_t$$
$$\frac{dW_{e,t}}{W_{e,t}} = \left(r_t - \frac{C_{e,t}}{W_{e,t}} - \lambda_d + \frac{\chi_t\psi_t}{z_t}\zeta_{e,t}(\sigma + \sigma^q_t) + \frac{\bar{z}\lambda_d}{z_t}\right)dt + \frac{\chi_t\psi_t}{z_t}(\sigma + \sigma^q_t)dZ_t$$

where $W_{h,t} = \int_{j\in H} w_{j,t}dj$ and $W_{e,t} = \int_{j\in E} w_{j,t}dj$ denotes the aggregated wealth among respective group. Similar to the stochastic productivity model, the volatility terms $\frac{\chi_t\psi_t}{z_t}(\sigma + \sigma^q_t)$ and $\frac{1-\chi_t\psi_t}{1-z_t}(\sigma + \sigma^q_t)$ can be derived using the definitions of z_t, ψ_t and the market clearing condition $\sigma_{w_e,t}z_t(\sigma + \sigma^q_t) + \sigma_{w_h,t}(1-z_t)(\sigma + \sigma^q_t) = (\sigma + \sigma^q_t)$. 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

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

and the result follows from here after some algebra. ■

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 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.

Asset pricing conditions

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

$$dr^v_t = (\mu^R_{e,t} - (1-\chi_t)\epsilon_{h,t})dt + \chi_t(\sigma^{q,k}_t + \sigma)dZ^k_t$$

where $\epsilon_{h,t} = \zeta_{h,t}(\sigma^q_t + \sigma)$. 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 $\frac{dv_t}{v_t} = dr^v_t$, and

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

where $\epsilon_{e,t} = \zeta_{e,t}(\sigma + \sigma_t^q)$, and $\xi_{e,t}$ follows the process in (A.67). Since $\xi_e \nu_t$ is a martingale, the drift equals to zero, which implies $\mu_{e,t}^R - r_t = \chi_t \epsilon_{e,t} + (1-\chi_t)\epsilon_{h,t}$ It follows similarly for the households with the difference that since they do not issue outside equity, their asset pricing condition is $\mu_{h,t}^R - r_t = \epsilon_{h,t}$ ■

While the quantitative analysis of the benchmark model in main text assumes that agents have recursive utility and IES=1, I present and solve the model for a broader range of preference specifications. I consider four different types of utility functions. Let

$$f(c_{j,s}, U_{j,s}) = \begin{cases} \rho_j \log(c_{j,t}) - \rho_j U_{j,t} & \text{if} \quad \gamma_j = 1, \varrho_j = 1 \\ \frac{c_{j,t}^{1-\gamma_j}}{1-\gamma_j} - \rho_j U_{j,t} & \text{if} \quad \gamma_j = \varrho_j^{-1} \neq 1 \\ (1-\gamma_j)\rho_j U_{j,t}\left(\log(c_{j,t}) - \frac{1}{1-\gamma_j}\log\left((1-\gamma_j)U_{j,t}\right)\right) & \text{if} \quad \gamma_j \neq 1, \varrho_j = 1 \\ \frac{1-\gamma_j}{1-\frac{1}{\varrho_j}} U_{j,t}\left[\left(\frac{c_{j,t}}{\left((1-\gamma_j)U_{j,t}\right)^{1/(1-\gamma_j)}}\right)^{1-\frac{1}{\varrho_j}} - \rho_j\right] & \text{if} \quad \gamma_j \neq 1, \varrho_j \neq 1 \end{cases} \tag{A.22}$$

I allow for preference heterogeneity in risk aversion and discount rate for generality. I solve for a Markov equilibrium in the state variable $z_t \in (0,1)$ for a representative household and expert by aggregating all agents within their respective group.

Proposition 9. *The optimal consumption policy and price of risk are given by*

$$\hat{C}_{e,t} = \begin{cases} \rho_e & \textit{if} \quad (\log \quad \textit{or} \quad \textit{Recursive (IES=1)}) \\ J_{e,t}^{-1/\gamma_e}(z_t q_t)^{\frac{1-\gamma_e}{\gamma_e}} & \textit{if} \quad \textit{CRRA} \\ \frac{J_{e,t}^{\frac{1-\varrho_j}{1-\gamma_e}}}{(z_t q_t)^{1-\varrho_j}} & \textit{if} \quad \textit{Recursive (IES} \neq 1\textit{)} \end{cases} \tag{A.23}$$

$$\hat{C}_{h,t} = \begin{cases} \rho_h & \textit{if} \quad (\textit{log} \quad \textit{or} \quad \textit{Recursive (IES=1)}) \\ J_{h,t}^{-1/\gamma_h}((1-z_t)q_t)^{\frac{1-\gamma_h}{\gamma_h}} & \textit{if} \quad \textit{CRRA} \\ \frac{J_{h,t}^{\frac{1-\varrho_j}{1-\gamma_h}}}{((1-z_t)q_t)^{1-\varrho_j}} & \textit{if} \quad \textit{Recursive (IES} \neq 1\textit{)} \end{cases} \tag{A.24}$$

$$\zeta_{e,t} = \begin{cases} \frac{\chi_t \psi_t}{z_t}(\sigma + \sigma_t^q) & \textit{if} \quad \textit{log} \\ -\sigma_{e,t}^J + \sigma_t^z + \sigma_t^q + \gamma_e \sigma & \textit{if} \quad (\textit{CRRA} \quad \textit{or} \quad \textit{Recursive}) \end{cases} \tag{A.25}$$

$$\zeta_{h,t} = \begin{cases} \frac{(1-\chi_t\psi_t)}{1-z_t}(\sigma + \sigma_t^q) & \textit{if} \quad \textit{log} \\ -\sigma_{h,t}^J - \frac{z_t}{1-z_t}\sigma_t^z + \sigma_t^q + \gamma_h \sigma & \textit{if} \quad (\textit{CRRA} \quad \textit{or} \quad \textit{Recursive}) \end{cases} \tag{A.26}$$

Proof: The HJB equation is given by

$$\sup_{c,K} f(c_{j,t}, U_{j,t}) + E[dU_{j,t}] = 0 \tag{A.27}$$

I consider three cases of utility functions separately.

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

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

and where the second equality follows from $z_t = \frac{W_{e,t}}{q_t K_t} = 1 - \frac{W_{h,t}}{q_t K_t}$. Also, $f(C_{j,t}, U_{j,t}) = \rho_j \log(C_{j,t}) - \rho_j U_{j,t}$. The value function derivatives are

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

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

$$\sup_{C,\theta_{j,t}} \rho_j \log C_{j,t} - \rho(\log W_{j,t} + \tilde{J}_{j,t}) + r_t - \frac{C_{j,t}}{W_{j,t}} + \theta_{j,t}(\sigma + \sigma_t^q)\zeta_{j,t} - \frac{1}{2}\theta_{j,t}^2(\sigma + \sigma_t^q)^2 + \mu_t^{\tilde{J}} = 0$$

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

$$\hat{c}_{j,t} = \rho_j \tag{A.28}$$

$$\zeta_{e,t} = \frac{\chi_t \psi_t}{z_t}(\sigma + \sigma_t^q) \tag{A.29}$$

$$\zeta_{h,t} = \frac{1 - \chi_t \psi_t}{1 - z_t}(\sigma + \sigma_t^q) \tag{A.30}$$

(b) CRRA Utility The value function conjecture is

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

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

$$\begin{aligned} &\frac{\partial U_{j,t}}{\partial J_{j,t}} = \frac{K_t^{1-\gamma_j}}{1-\gamma_j}; \quad \frac{\partial U_{j,t}}{\partial K_t} = J_{j,t} K_t^{-\gamma_j} \\ &\frac{\partial^2 U_{j,t}}{\partial J_{j,t}^2} = 0; \quad \frac{\partial^2 U_{j,t}}{\partial K_t^2} = -\gamma_j J_{j,t} K_t^{-(1+\gamma_j)}; \quad \frac{\partial^2 U_{j,t}}{\partial J_{j,t} \partial K_t} = K_t^{-\gamma_j} \end{aligned} \tag{A.31}$$

Applying Ito's lemma and using HJB, we get

$$\begin{aligned} \sup_{C,K} \quad &-\rho\frac{J_{j,t}K_t^{1-\gamma_j}}{1-\gamma_j} + \frac{C_t^{1-\gamma_j}}{1-\gamma_j} + \frac{J_{j,t}K_t^{1-\gamma_j}}{1-\gamma_j}\mu_{j,t}^J + J_{j,t}K_t^{1-\gamma_j}(\Phi(\iota_t) - \delta) \\ &- \sigma^2\frac{\gamma_j}{2}J_{j,t}K_t^{1-\gamma_j} + J_{j,t}K_t^{1-\gamma_j}\sigma\sigma_{j,t}^J = 0 \end{aligned} \tag{A.32}$$

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

$$\frac{C_{e,t}}{W_{e,t}} = \frac{(z_t q_t)^{\frac{1-\gamma_e}{\gamma_e}}}{J_{e,t}^{\frac{1}{\gamma_e}}}; \qquad \frac{C_{h,t}}{W_{h,t}} = \frac{((1-z_t)q_t)^{\frac{1-\gamma_h}{\gamma_h}}}{J_{h,t}^{\frac{1}{\gamma_h}}} \tag{A.33}$$

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

$$\zeta_{e,t} = \gamma_e \sigma^{c}_{e,t} = -\sigma^{J}_{e,t} + \sigma^{z}_{t} + \sigma^{q}_{t} + \gamma_e \sigma; \quad \zeta_{h,t} = \gamma_h \sigma^{c}_{h,t} = -\sigma^{J}_{h,t} - \frac{z_t}{1-z_t}\sigma^{z}_{t} + \sigma^{q}_{t} + \gamma_h \sigma \tag{A.34}$$

Plugging in the optimal consumption-wealth ratio from (*A*.33) into HJB equation (*A*.32), we get the expressions for $\mu^{J}_{j,t}$

$$\mu^{J}_{e,t} = \rho_e - \frac{(z_t q_t)^{\frac{1-\gamma_e}{\gamma_e}}}{J_{e,t}^{1/\gamma_e}} - (1-\gamma_e)\left(\Phi(\iota_t) - \delta - \frac{\gamma_e}{2}\sigma^2 + \sigma^{J}_{e,t}\sigma\right) \tag{A.35}$$

$$\mu^{J}_{h,t} = \rho_e - \frac{((1-z_t)q_t)^{\frac{1-\gamma_h}{\gamma_h}}}{J_{h,t}^{1/\gamma_h}} - (1-\gamma_h)\left(\Phi(\iota_t) - \delta - \frac{\gamma_h}{2}\sigma^2 + \sigma^{J}_{h,t}\sigma\right) \tag{A.36}$$

(c) Recursive Utility (IES=1) The value function conjecture is the same as that of CRRA utility, and $f(C_{j,t}U_{j,t}) = (1-\gamma_j)\rho_j U_{j,t}\left(\log C_{j,t} - \frac{1}{1-\gamma_j}\log\left((1-\gamma_j)U_{j,t}\right)\right)$. Plugging in the conjecture for

value function in HJB equation $(A.1)$ and applying Ito's lemma[5], we get

$$\begin{aligned}\sup_{C,K} \quad & \rho J_{j,t} K_t^{1-\gamma_j} [\log \frac{C_{j,t}}{W_{j,t}} - \frac{1}{1-\gamma_j} \log J_{j,t} + \log(q_t z_t)] + J_{j,t} \frac{K_t^{1-\gamma_j}}{1-\gamma_j} \mu_{j,t}^J \\ & + J_{j,t} K_t^{1-\gamma_j} (\Phi(\iota_t) - \delta) - J_{j,t} K_t^{1-\gamma_j} \frac{1}{2} \gamma_j \sigma^2 + J_{j,t} K_t^{1-\gamma_j} \sigma \sigma_{j,t}^J = 0\end{aligned} \tag{A.37}$$

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

$$\begin{aligned}\frac{\partial U_{j,t}}{\partial W_{j,t}} &= \frac{\partial f}{\partial C_{j,t}} \\ \tilde{J}_{j,t} W_{j,t}^{-\gamma_j} = (1-\gamma_j) \rho_j \frac{U_{j,t}}{C_{j,t}} &\implies \frac{C_{j,t}}{W_{j,t}} = \rho_j\end{aligned}$$

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

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

This implies that $\sigma(\xi_{j,t}) = \sigma\left(\frac{U_{j,t}}{W_{j,t}} \right)$. 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 $(A.37)$, we obtain the expressions for $\mu_{j,t}^J$

$$\mu_{e,t}^J = (\gamma_e - 1)\big(\rho_e \log \rho_e + \log(q_t z_t) \big) + \rho_e \log J_{e,t} - (1-\gamma_e)(\Phi(\iota_t) - \delta - \frac{\gamma_e}{2} \sigma^2 + \sigma \sigma_{e,t}^J) \tag{A.38}$$

$$\mu_{h,t}^J = (\gamma_h - 1)\big(\rho_h \log \rho_h + \log(q_t (1 - z_t)) \big) + \rho_h \log J_{h,t} - (1-\gamma_h)(\Phi(\iota_t) - \delta - \frac{\gamma_h}{2} \sigma^2 + \sigma \sigma_{h,t}^J) \tag{A.39}$$

[5] The value function derivatives are the same as in the CRRA case given by $(A.31)$.

(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

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

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

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

where $J_{j,t}$ follows the stochastic differential equation $\frac{dJ_{j,t}}{J_{j,t}} = \mu^J_{j,t}dt + \sigma^J_{j,t}dZ_t$ whose drift and volatility needs to be determined in the equilibrium.[6]

As before, 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

$$\sup_{c,K} \frac{1}{1-\frac{1}{\varrho_j}}\left(\frac{C_{j,t}^{1-\frac{1}{\varrho_j}}}{J_{j,t}^{\frac{1-\frac{1}{\varrho_j}}{1-\gamma_j}} K_t^{1-\frac{1}{\varrho_j}}} - \rho_j\right) J_{j,t}K_t^{1-\gamma_j} + \frac{J_{j,t}K_t^{1-\gamma_j}}{1-\gamma}\mu^J_{j,t} + J_{j,t}K_t^{1-\gamma_j}(\Phi(\iota_t)-\delta) \tag{A.40}$$
$$-\sigma^2\frac{\gamma_j}{2}J_{j,t}K_t^{1-\gamma_j} + J_{j,t}K_t^{1-\gamma_j}\sigma\sigma^J_{j,t} = 0$$

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 $q_tK_t = \frac{W_{e,t}}{z_t} = \frac{W_{h,t}}{1-z_t}$, we have

$$\frac{\partial f_{e,t}}{\partial C_{e,t}} = C_{e,t}^{-\frac{1}{\varrho_e}} J_{e,t}^{\frac{\frac{1}{\varrho_e}-\gamma_e}{1-\gamma_e}} (z_tq_t)^{\gamma_j-\frac{1}{\varrho_e}}$$
$$\frac{\partial f_{h,t}}{\partial C_{h,t}} = C_{h,t}^{-\frac{1}{\varrho_h}} J_{h,t}^{\frac{\frac{1}{\varrho_h}-\gamma_h}{1-\gamma_h}} ((1-z_t)q_t)^{\gamma_j-\frac{1}{\varrho_h}}$$
$$\frac{\partial U_{e,t}}{\partial W_{e,t}} = \frac{J_{e,t}}{(z_tq_t)^{1-\gamma_e}} W_{e,t}^{1-\gamma_e}$$
$$\frac{\partial U_{h,t}}{\partial W_{h,t}} = \frac{J_{h,t}}{((1-z_t)q_t)^{1-\gamma_h}} W_{h,t}^{1-\gamma_h}$$

[6] Since the value function conjecture is the same as in CRRA case, the value function derivatives are given by (*A*.31).

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

$$\frac{C_{e,t}}{W_{e,t}} = \frac{J_{e,t}^{\frac{1-\varrho_e}{1-\gamma_e}}}{(z_t q_t)^{1-\varrho_e}}; \qquad \frac{C_{h,t}}{W_{h,t}} = \frac{J_{h,t}^{\frac{1-\varrho_h}{1-\gamma_h}}}{((1-z_t)q_t)^{1-\varrho_h}} \tag{A.41}$$

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 $\sigma(\xi_{j,t}) = \sigma\left(\frac{U_{j,t}}{W_{j,t}}\right)$. 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.

$$\zeta_{e,t} = -\sigma_{e,t}^J + \sigma_t^z + \sigma_t^q + \gamma_e\sigma \tag{A.42}$$

$$\zeta_{h,t} = -\sigma_{h,t}^J - \frac{z_t}{1-z_t}\sigma_t^z + \sigma_t^q + \gamma_h\sigma \tag{A.43}$$

Substituting the consumption-wealth ratio into the HJB equation (*A*.40), we the expression for $\mu_{j,t}^J$

$$\mu_{e,t}^J = \frac{(\gamma_e-1)}{1-\frac{1}{\varrho_e}}\left((q_t z_t)^{\varrho_e-1}J_{e,t}^{\frac{1-\varrho_e}{1-\gamma_e}} - \rho_e\right) - (1-\gamma_e)(\Phi(\iota_t) - \delta - \frac{\gamma_e}{2}\sigma^2 + \sigma\sigma_{e,t}^J) \tag{A.44}$$

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

	All	Crisis	Normal	All	Crisis	Normal	All	Crisis	Normal	All	Crisis	Normal
E[leverage]	3.5	5.8	3.32	2.08	4.36	2.08	1.57		1.57	1.29		1.29
E[inv. rate]	7.00%	5.6%	7.10%	5.80%	5.60%	5.80%	5.50%		5.50%	5.30%		5.30%
E[risk premia]	1.70%	13.40%	0.82%	2.70%	16.50%	2.70%	4.50%		4.50%	8.00%		8.00%
E[return volatility]	6.20%	15.90%	5.70%	5.80%	14.30%	5.80%	5.80%		5.80%	5.90%		5.90%
E[GDP growth rate]	2.30%	-8.0%	3.07%	2.10%	-10.70%	2.10%	1.90%		1.90%	1.80%		1.80%
Std[inv. rate]	0.10%	1.30%	0.22%	0.13%	0.23%	0.12%	0.08%		0.08%	0.04%		0.04%
Std[risk premia]	3.10%	2.40%	3.5%	0.43%	0.56%	0.21%	0.17%		0.17%	0.10%		0.10%
Corr(leverage,shock)	-0.27	-0.04	-0.24	-0.19	0.29	-0.19	-0.21		-0.21	-0.26		-0.26
Prof. of crisis	6.80%			0.10%			0.01%			0.00%		
Risk aversion	1			5			10			20		

Table A.2 – Benchmark model implied moments for different risk aversion levels with parameters from Table (*A*.1).

A.2 Appendix to Chapter 2

A.2.1 Proof of (2.3.1)

The proof can be done in two steps.

Step-1 First, I refer to the result that infinitely wide neural nets trained by the gradient descent are kernel regressions (Jacot et al. (2018)). Consider a recursive formulation of the neural network introduced in the section (2.3.1).

$$h^{(l)}(\boldsymbol{x}) = \frac{1}{\sqrt{d_l}} \boldsymbol{W}^{(l)} \boldsymbol{z}^{(l)} + \boldsymbol{b}^l \in \mathbb{R}^{d_{l+1}}$$
$$\boldsymbol{z}^{(l)}(\boldsymbol{x}) = \sigma(\boldsymbol{W}^{(l-1)} \boldsymbol{h}^{(l-1)}(\boldsymbol{x}) + \boldsymbol{b}^{(l-1)})$$

for $l = 1, ..., L$ where the parameters $\boldsymbol{W}^{(l)} \in \mathbb{R}^{d_{l+1} \times d_l}$ and $\boldsymbol{b}^{(h)} \in \mathbb{R}^{d_{l+1}}$ are the weights and biases respectively, and $\sigma(\cdot)$ is as defined in (2.3.1). The final layer of the neural network is given by

$$J(\boldsymbol{x}, \boldsymbol{\theta}) = \boldsymbol{h}^{(L)}(\boldsymbol{x}) = \frac{1}{\sqrt{d_L}} \boldsymbol{W}^{(L)} \boldsymbol{z}^{(L)} + \boldsymbol{b}^L \tag{A.45}$$

where the parameters $\boldsymbol{W}^{(L)} \in \mathbb{R}^{1 \times d_l}$ and $\boldsymbol{b}^{(L)} \in \mathbb{R}$ are the weights and biases of the final layer in the neural network. The parameter vector $\boldsymbol{\theta} = \{\boldsymbol{W}^{(l)}, \boldsymbol{b}^{(l)}\}_{l=0}^{L}$ combines the parameters of all layers in the network and is initialized with i.i.d standard normal random values. Then, in the sequential limit of the width of hidden layers $l_1, l_2, ..., l_L \to \infty$, the co-ordinates of the hidden layers converges asymptotically to an $i.i.d$ Gaussian process with covariance $\Sigma^{l-1} : \mathbb{R}^{d_{l-1}} \times \mathbb{R}^{d_{l-1}} \to \mathbb{R}$ that can be defined recursively as folllows.

$$\Sigma^{(0)} = \boldsymbol{x}^T \boldsymbol{x}' + 1$$
$$\Lambda^{(h)}(\boldsymbol{x}, \boldsymbol{x}') = \begin{pmatrix} \Sigma^{(l-1)}(\boldsymbol{x}, \boldsymbol{x}) & \Sigma^{(l-1)}(\boldsymbol{x}, \boldsymbol{x}') \\ \Sigma^{(l-1)}(\boldsymbol{x}', \boldsymbol{x}) & \Sigma^{(l-1)}(\boldsymbol{x}', \boldsymbol{x}') \end{pmatrix} \in \mathbb{R}^{2 \times 2},$$
$$\Sigma^{(h)}(\boldsymbol{x}, \boldsymbol{x}') = \mathbb{E}_{u,v}[\sigma(u)\sigma(v)] + 1$$

for $l = 1, ..., L$, where $(u, v) \sim N(0, \Lambda^{(1)})$, where $\boldsymbol{x}$ and $\boldsymbol{x}'$ are two different input samples. Let us define

$$\dot{\Sigma}(\boldsymbol{x}, \boldsymbol{x}') = \mathbb{E}_{(u,v)}[\dot{\sigma}(u)\dot{\sigma}(v)] \tag{A.46}$$

where $\dot{\sigma}$ is the derivative of the activation function. The neural tangent kernel (NTK, henceforth) is defined as

$$K_t(\boldsymbol{x}, \boldsymbol{x}') = \left\langle \frac{\partial J(\boldsymbol{x}, \boldsymbol{\theta}(t)}{\partial \boldsymbol{\theta}}, \frac{\partial J(\boldsymbol{x}', \boldsymbol{\theta}(t)}{\partial \boldsymbol{\theta}} \right\rangle \tag{A.47}$$

As shown in Jacot et al. (2018), this kernel converges in probability to a deterministic kernel $\Theta^L(\boldsymbol{x}, \boldsymbol{x'})$ as the width of hidden layers go to infinity, for all $0 \leq t \leq T$. Thus, a sufficiently wide deep neural network resembles a kernel regression with deterministic kernel.

Step-2 Next step is to show that the NTK of well-posed PDE defined in (2.12) also converges to a deterministic kernel with random parameter initialization. Without loss of generality, let us consider a simplified PDE that takes the form

$$\mathscr{G}[J](\boldsymbol{x}) = 0 \qquad x \in \Omega \tag{A.48}$$

$$J(\boldsymbol{x}) = g(\boldsymbol{x}) \qquad x \in \partial\Omega \tag{A.49}$$

where $\mathscr{G}$ is a differential operator. Note that this formulation includes time dependent PDE where t is one of the dimensions in $\boldsymbol{x}$. The loss function after approximating J with a deep neural network object is given by

$$\mathscr{L}(\boldsymbol{\theta}) = \frac{1}{2}\sum_{i=1}^{N_b} |J(\boldsymbol{x}_b^i) - g(\boldsymbol{x}_b^i)|^2 + \frac{1}{2}\sum_{i=1}^{N_f} \mathscr{G}[J](\boldsymbol{x}_f^i) \tag{A.50}$$

where N_f and N_b are the batch sizes for the training samples to approximate the PDE residual and boundary condition respectively. Assume that the training samples are drawn randomly at each iteration of the gradient descent. Let the continuous-time gradient flow of the PDE approximation using a neural network trained with an infinitesimally small learning rate be defined as

$$\frac{d\boldsymbol{\theta}}{dt} = -\nabla\mathscr{L}(\boldsymbol{\theta}) \tag{A.51}$$

where $\mathscr{L}$ is the loss function. Let[7] $\mathscr{G}J(t) = \mathscr{G}[J](\boldsymbol{x}, \boldsymbol{\theta}); \forall (t,x) \in \Omega_T$ and $J(t) = J(\boldsymbol{x}, \boldsymbol{\theta}(t)); \forall (t,x) \in \Omega_{T'}$.

Lemma A.2.1. Given the training points to train the neural network that approximates the PDE system and the gradient flow given in (A.51), $J(t)$ and $\mathscr{G}(t)$ obey the following evolution

$$\begin{bmatrix} \frac{dJ(\boldsymbol{x}_b, \boldsymbol{\theta}(t))}{dt} \\ \frac{d\mathscr{G}J(\boldsymbol{x}_f, \boldsymbol{\theta}(t))}{dt} \end{bmatrix} = -\begin{bmatrix} \boldsymbol{K}_{jj}(t) & \boldsymbol{K}_{jf}(t) \\ \boldsymbol{K}_{fj}(t) & \boldsymbol{K}_{ff}(t) \end{bmatrix} \cdot \begin{bmatrix} J(\boldsymbol{x}_b, \boldsymbol{\theta}(t)) - g(\boldsymbol{x}_b) \\ \mathscr{G}J(\boldsymbol{x}_f, \boldsymbol{\theta}(t)) \end{bmatrix} \tag{A.52}$$

where $\boldsymbol{K}_{fj}(t) = \boldsymbol{K}_{jf}(t)^T$ and $\boldsymbol{K}_{jj}(t) \in \mathbb{R}^{N_b \times N_b}$, $\boldsymbol{K}_{jf}(t) \in \mathbb{R}^{N_b \times N_f}$, and $\boldsymbol{K}_{ff}(t) \in \mathbb{R}^{N_r \times N_r}$, whose

[7] I define $\Omega_T := [T - k\delta t, T - (k-1)\delta t] \times \Omega$, and $\Omega_{T'} := (T - (k-1)\delta t) \times \Omega$ for simplicity.

$(m,n)-th$ entry is given by

$$(\boldsymbol{K}_{jj})_{mn}(t) = \left\langle \frac{dJ(\boldsymbol{x}_b^i,\boldsymbol{\theta}(t))}{d\boldsymbol{\theta}}, \frac{dJ(\boldsymbol{x}_b^i,\boldsymbol{\theta}(t))}{d\boldsymbol{\theta}} \right\rangle$$
$$(\boldsymbol{K}_{jf})_{mn}(t) = \left\langle \frac{dJ(\boldsymbol{x}_b^i,\boldsymbol{\theta}(t))}{d\boldsymbol{\theta}}, \frac{d\mathcal{G}J(\boldsymbol{x}_f^i,\boldsymbol{\theta}(t))}{d\boldsymbol{\theta}} \right\rangle$$
$$(\boldsymbol{K}_{ff})_{mn}(t) = \left\langle \frac{d\mathcal{G}J(\boldsymbol{x}_f^i,\boldsymbol{\theta}(t))}{d\boldsymbol{\theta}}, \frac{d\mathcal{G}J(\boldsymbol{x}_f^i,\boldsymbol{\theta}(t))}{d\boldsymbol{\theta}} \right\rangle$$

Remark: The notation $\langle\cdot\rangle$ denotes the inner product, and the matrices defined above are all positive semi-definite. Let $\boldsymbol{K}(t) := \begin{bmatrix} \boldsymbol{K}_{jj}(t) & \boldsymbol{K}_{jf}(t) \\ \boldsymbol{K}_{fj}(t) & \boldsymbol{K}_{ff}(t) \end{bmatrix}$.

Theorem A.2.2. For the PDE (A.48) approximated with one hidden layer initialized with random values and trained using a gradient descent with infinitesimally small learning rate, we have the following

1. The NTK $\boldsymbol{K}(t)$ of the PDE converges in probability to a deterministic kernel in the infinite width limit.
$$\boldsymbol{K}(0) = \begin{bmatrix} \boldsymbol{K}_{jj}(0) & \boldsymbol{K}_{jf}(0) \\ \boldsymbol{K}_{fj}(0) & \boldsymbol{K}_{ff}(0) \end{bmatrix} \to \boldsymbol{K}^*$$

2. Assume that $\boldsymbol{\Theta}$ is uniformly bounded $\forall\, t$, and there exists a constant $K > 0$ such that
$$\int_0^T |\sum_{i=1}^{N_b} \big(J(x_b^i,\boldsymbol{\theta}(\tau)) - g(x_b^i)\big)| d\tau \le K$$
$$\int_0^T |\sum_{i=1}^{N_f} \mathcal{G}J(x_f^i,\boldsymbol{\theta}(\tau))| d\tau \le K$$
and the activation function σ is smooth with $|\sigma^{(k)}| \le C$ for $0 \le K \le 4$ where $\sigma^{(k)}$ denotes the $k-th$ order derivative of the σ. Then, in the limit $N \to \infty$,
$$\sup_{t\in[0,T]} ||\boldsymbol{K}(t) - \boldsymbol{K}(0)||_2 = 0 \tag{A.53}$$

Proof: The proof can be found in Wang et al. (2020).

The final step is to show that as the sample size increases, the outcome of the kernel regression recovers the true function fully everywhere. This can be done following Mei et al. (2021).

A.2.2 Approximation error

I consider a special case of quasi-linear parabolic PDE of the form[8]

$$\mathscr{G}[J](t,\boldsymbol{x}) := \frac{\partial J}{\partial t} + \sum_{j}^{d} a(t,\boldsymbol{x})\frac{\partial J}{\partial x_j} + \sum_{j}^{d} b(t,\boldsymbol{x})\frac{\partial^2 J}{\partial x_i \partial x_j} + \hat{A}(t,\boldsymbol{x},J(t,\boldsymbol{x}),\nabla J(t,\boldsymbol{x})) = 0; \qquad \forall (t,\boldsymbol{x}) \in \Omega_T \tag{A.54}$$

with the boundary conditions

$$J(t,\boldsymbol{x}) = J_0(t,\boldsymbol{x}); \quad \forall (t,\boldsymbol{x}) \in \Omega \tag{A.55}$$

$$J(t,\boldsymbol{x}) = J_0(t,\boldsymbol{x}); \quad \forall (t,\boldsymbol{x}) \in \Omega_c^1 \tag{A.56}$$

$$\mathscr{G}[J](t,\boldsymbol{x}) = 0; \qquad \forall (t,\boldsymbol{x}) \in \Omega_c^2 \tag{A.57}$$

$$\frac{\partial J(t,\boldsymbol{x})}{\partial \boldsymbol{x}} = 0; \quad \forall (t,\boldsymbol{x}) \in \partial\Omega_T \tag{A.58}$$

where $\Omega_c^1 \subset \Omega$ and $\Omega_c^2 \subset \Omega$ are the active subdomains where the boundary condition and PDE are satisfied respectively. Note that the functions $a(t,\boldsymbol{x})$ and $b(t,\boldsymbol{x})$ are bounded and do not depend on the function J. This is a valid assumption since the PDE in (2.12) is solved using value function iteration where the coefficients are determined in the inner static loop and hence do not depend on the function J that needs to be solved.[9] I also assume that the PDE (A.54) has a unique solution $J(t,\boldsymbol{x}) \in \mathscr{C}^{1,2}(\Omega_T)$ with its derivatives uniformly bounded. I prove the Theorem (A.2.4) for this special class of quasi-linear parabolic PDE, that is common in the macro-finance literature (see Achdou et al. (2017)), using a single hidden layer neural network with n number of neurons. Consider the class

$$\check{\mathscr{C}}^n(\sigma) = \left\{ h(t,\boldsymbol{x}) : \mathbb{R}^{1+d} \to \mathbb{R} \mid h(t,\boldsymbol{x}) = \sum_{j=1}^{n} \beta_j \sigma\big(\alpha_j t + \sum_{i=1}^{d} w_{i,j} x_j + b_i\big) \right\} \tag{A.59}$$

such that $\check{\mathscr{C}}(\sigma) = \cup_{n\geq 1}\check{\mathscr{C}}^n(\sigma)$, σ is the activation function, and $(\alpha_j, w_{i,j}, b_i)$ are the parameters of the neural network. The PDE (A.54) approximated with $\hat{J}(t,\boldsymbol{x}|\Theta) \in \check{\mathscr{C}}(\sigma)$ has the total loss given by

$$\begin{aligned}\mathscr{L}(\hat{J}) &= \|\mathscr{G}[J](t,\boldsymbol{x})\|^2_{\Omega_T} + \left\| \frac{\partial \hat{J}(t,\boldsymbol{x}|\Theta)}{\partial \boldsymbol{x}} \right\|^2_{\partial\Omega_T} + \left\| \hat{J}(t,\boldsymbol{x}|\Theta) - J_0 \right\|^2_{\Omega} + \\ &\quad \left\| \hat{J}(t,\boldsymbol{x}|\Theta) - J_0 \right\|^2_{\Omega_c^1} + \|\mathscr{G}[J](t,\boldsymbol{x})\|^2_{\Omega_c^2} \\ &= \|\mathscr{G}[J](t,\boldsymbol{x})\|^2_{\Omega_{C,T}} + \left\| \frac{\partial \hat{J}(t,\boldsymbol{x}|\Theta)}{\partial \boldsymbol{x}} \right\|^2_{\partial\Omega_T} + \left\| \hat{J}(t,\boldsymbol{x}|\Theta) - J_0 \right\|^2_{\Omega_C}\end{aligned} \tag{A.60}$$

[8] Note that by defining $\Omega_T := [T-(k-1)\Delta T, T-k\Delta T]\times\Omega$, $\Omega_C := (T-(k-1)\Delta t)\times\Omega_c$ and $\partial\Omega_T := (T-(k-1)\Delta t)\times\partial\Omega$, we recover the PDE (2.12) at kth time iteration. Also J_0 is mapped to $\tilde{J}$ in (2.12).

[9] In fact, $a(\cdot)$ and $b(\cdot)$ are treated as coefficients but I prove for a general case where they could be functions of the state variables.

where $\Omega_{C,T} := \Omega_c^2 \cup \Omega_T$ and $\Omega_C = \Omega \cup \Omega_c^1$

$$
\begin{aligned}
&= \left\| \mathscr{G}[\hat{J}](t,\boldsymbol{x}|\Theta) - \mathscr{G}[J](t,\boldsymbol{x}) \right\|_{\Omega_{C,T}}^2 + \left\| \frac{\partial \hat{J}(t,\boldsymbol{x}|\Theta)}{\partial \boldsymbol{x}} \right\|_{\partial\Omega_T}^2 + \left\| \hat{J}(t,\boldsymbol{x}|\Theta) - J_0 \right\|_{\Omega_C}^2 \\
&\leq \int_{\Omega_{C,T}} \left| \frac{\partial J(\boldsymbol{x},t)}{\partial t} - \frac{\partial \hat{J}(t,\boldsymbol{x}|\Theta)}{\partial t} \right|^2 d\mu_1 \\
&+ \int_{\Omega_{C,T}} \left| \sum_{i=1}^{d} a(t,\boldsymbol{x}) \Big(\frac{\partial J(t,\boldsymbol{x})}{\partial x_i} - \frac{\partial \hat{J}(t,\boldsymbol{x})}{\partial x_i} \Big) \right|^2 d\mu_1 \\
&+ \int_{\Omega_{C,T}} \left| \sum_{i,j=1}^{d} b(t,\boldsymbol{x}) \Big(\frac{\partial^2 J(t,\boldsymbol{x})}{\partial x_i x_j} - \frac{\partial^2 \hat{J}(t,\boldsymbol{x})}{\partial x_i x_j} \Big) \right|^2 d\mu_1 \\
&+ \int_{\Omega_{C,T}} \left| \hat{A}(t,\boldsymbol{x},J(t,\boldsymbol{x}),\nabla J(t,\boldsymbol{x})) - \hat{A}(t,\boldsymbol{x},\hat{J}(t,\boldsymbol{x}),\nabla \hat{J}(t,\boldsymbol{x})) \right|^2 d\mu_1 \\
&+ \int_{\Omega_C} \left| \hat{J}(t,\boldsymbol{x}|\Theta) - J_0(t,\boldsymbol{x}) \right|^2 d\mu_2 + \int_{\partial\Omega_T} \left| \frac{\partial \hat{J}(t,\boldsymbol{x}|\Theta)}{\partial \boldsymbol{x}} \right| d\mu_3
\end{aligned} \tag{A.61}
$$

Definition A.2.1. A subset $S \subset \mathscr{C}^m(\Omega)$ is uniformly m-dense on compacts of $C^m(\Omega)$ if $\forall f \in \mathscr{C}^m(\Omega)$, for all compact subsets X of Ω and $\forall \epsilon > 0$, $\exists g(f,X,\epsilon) \in S$ such that

$$||f - g||_{m,X} < \epsilon$$

Lemma A.2.3. There exists a function $\hat{J} \in \mathscr{C}^n(\sigma)$ defined in (A.59) such that for all $\epsilon > 0$, we have

$$
\begin{aligned}
\sup_{(t,x)\in\Omega_T} |J(t,\boldsymbol{x}) - \hat{J}(t,\boldsymbol{x}|\Theta)| + \sup_{(t,x)\in\Omega_T} \left| \frac{\partial J(t,\boldsymbol{x})}{\partial t} - \frac{\partial \hat{J}(t,\boldsymbol{x}|\Theta)}{\partial t} \right| \\
+ \max_{|i|\leq 2} \sup_{(t,\boldsymbol{x})\in\Omega_T} \left| \frac{\partial^{(i)} J(t,\boldsymbol{x})}{\partial \boldsymbol{x}^{(i)}} - \frac{\partial^{(i)} \hat{J}(t,\boldsymbol{x}|\Theta)}{\partial \boldsymbol{x}^i} \right| < \epsilon
\end{aligned} \tag{A.62}
$$

Proof: This lemma is a direct consequence of Theorem 3 in Hornik (1991) with $m = 2$. The theorem states that if $\sigma \in \mathscr{C}^m(\mathbb{R}^{1+d})$ is a non-constant and bounded function, then $\check{C}(\sigma)$ is uniformly m-dense on compacts in $\mathscr{C}^m(\mathbb{R}^{1+d})$ ■

Lemma A.2.4. Consider a class of neural networks $\check{\mathscr{C}}^n$ with one hidden layer and n number of neurons. Let us denote $\hat{J}^n(t,\boldsymbol{x}|\Theta)$ as the neural network approximation of the function J that solves (2.12), and let $\mathscr{L}$ be the loss function given in (2.20). Then, under certain conditions,

$$
\begin{aligned}
&\exists \hat{J}^n(t,\boldsymbol{x}|\Theta) \in \check{\mathscr{C}}^n \text{such that} \\
&\mathscr{L}(\hat{J}^n(t,\boldsymbol{x}|\Theta)) \to 0 \quad \text{as} \quad n \to \infty
\end{aligned} \tag{A.63}
$$

Theorem A.2.5. Let $\check{C}(\sigma)$ be given by $(A.59)$ with $\sigma(\cdot)$ a non-constant, bounded function, and

let (μ_1, μ_2, μ_3) be the measures with support in $(\Omega_{C,T}, \Omega_C, \partial\Omega_T)$ respectively. Assume that $\hat{A}(t,x,y,z)$ is Lipschitz continuous with Lipschitz constants growing at most polynomially in y and z. Then, $\forall \epsilon > 0$, $\exists K > 0$ such that there exists a function $\hat{J}(t,x|\Theta) \in \check{C}(\sigma)$ satisfying

$$\mathscr{L}(\hat{J}) < K\epsilon^2$$

Proof: If $\hat{A}(t,x,y,z)$ is Lipschitz continuous with Lipschitz constants growing at most polynomially in y and z, we have

$$\begin{aligned}
&\left|\hat{A}(t,x,\hat{J}(t,\boldsymbol{x}|\Theta),\nabla\hat{J}(t,\boldsymbol{x}|\Theta)) - \hat{A}(t,x,J(t,\boldsymbol{x}),\nabla J(t,\boldsymbol{x}))\right| \\
&\qquad \leq \left(\left|\hat{J}(t,\boldsymbol{x}|\Theta)\right|^{a_1/2} + \left|\nabla\hat{J}(t,\boldsymbol{x}|\Theta)\right|^{a_2/2} + |J(t,\boldsymbol{x})|^{a_3/2} + |\nabla J(t,\boldsymbol{x})|^{a_4/2}\right) \\
&\qquad \times \left(\left|\hat{J}(t,\boldsymbol{x}|\Theta) - J(t,\boldsymbol{x})\right| + \left|\nabla\hat{J}(t,\boldsymbol{x}|\Theta) - \nabla J(t,\boldsymbol{x})\right|\right) \qquad \text{(A.64)}
\end{aligned}$$

for some constants $0 < \{a_i\}_{i=1}^4 < \infty$ This condition will be crucial in proving convergence as shown below. Applying Young's inequality with exponents $p_1 = p_2 = 2$, we get

$$\begin{aligned}
&\int_{\Omega_T} \left|\hat{A}(t,\boldsymbol{x},\hat{J}(t,\boldsymbol{x}|\Theta),\nabla\hat{J}(t,\boldsymbol{x}|\Theta)) - \hat{A}(t,\boldsymbol{x},J(t,\boldsymbol{x}),\nabla J(t,\boldsymbol{x})\right|^2 d\mu_1 \\
&\qquad \leq 2\int_{\Omega_T} \left(\left|\hat{J}(t,\boldsymbol{x}|\Theta)\right|^{a_1} + \left|\nabla\hat{J}(t,\boldsymbol{x}|\Theta)\right|^{a_2} + |J(t,\boldsymbol{x})|^{a_3} + |\nabla J(t,\boldsymbol{x})|^{a_4}\right) \\
&\qquad \times \left(\left|\hat{J}(t,\boldsymbol{x}|\Theta) - J(t,\boldsymbol{x})\right| + \left|\nabla\hat{J}(t,\boldsymbol{x}|\Theta) - \nabla J(t,\boldsymbol{x})\right|\right) d\mu_1
\end{aligned}$$

Applying Hölder inequality $\|fg\|_1 \leq \|f\|_p \|g\|_q$ for some constants $p, q \in [1,\infty)$ with $\frac{1}{p} + \frac{1}{q} = 1$, we have

$$\begin{aligned}
&\int_{\Omega_{C,T}} \left|\hat{A}(t,\boldsymbol{x},\hat{J}(t,\boldsymbol{x}|\Theta),\nabla\hat{J}(t,\boldsymbol{x}|\Theta)) - \hat{A}(t,\boldsymbol{x},J(t,\boldsymbol{x}),\nabla J(t,\boldsymbol{x})\right|^2 d\mu_1 \\
&\leq 2\left(\int_{\Omega_{C,T}} \left(\left|\hat{J}(t,\boldsymbol{x}|\Theta)\right|^{a_1} + \left|\nabla\hat{J}(t,\boldsymbol{x}|\Theta)\right|^{a_2} + |J(t,\boldsymbol{x})|^{a_3} + |\nabla J(t,\boldsymbol{x})|^{a_4}\right)^p d\mu_1\right)^{1/p} \\
&\quad \left(\int_{\Omega_{C,T}} \left(\left|\hat{J}(t,\boldsymbol{x}|\Theta) - J(t,\boldsymbol{x})\right|^2 + \left|\nabla\hat{J}(t,\boldsymbol{x}|\Theta) - \nabla J(t,\boldsymbol{x})\right|^2\right)^q d\mu_1\right)^{1/q} \\
&\leq C_1\left(\int_{\Omega_{C,T}} \left(\left|\hat{J}(\Theta) - J(t,\boldsymbol{x})\right|^{a_1} + \left|\nabla\hat{J}(\Theta) - \nabla J(t,\boldsymbol{x})\right|^{a2} + |J(t,\boldsymbol{x})|^{a_1\wedge a_3} + |\nabla J(t,\boldsymbol{x})|^{a_2\wedge a_4)}\right)^p d\mu_1\right)^{1/p} \\
&\quad \times \left(\int_{\Omega_{C,T}} \left(\left|\hat{J}(t,\boldsymbol{x}|\Theta) - J(t,\boldsymbol{x})\right|^2 + \left|\nabla\hat{J}(t,\boldsymbol{x}|\Theta) - \nabla J(t,\boldsymbol{x})\right|^2\right)^q d\mu_1\right)^{1/q} \\
&\leq C_1\left(\epsilon^{a_1} + \epsilon^{a_2} + \sup_{\Omega_{C,T}} |J|^{a_1\wedge a_3} + \sup_{\Omega_{C,T}} |\nabla J|^{a_2\wedge a_4}\right)\mu_1(\Omega_{C,T})^{1/p}\left(2\epsilon^2\mu_1(\Omega_{C,T})^{1/q}\right) \\
&\leq C_2\epsilon^2
\end{aligned}$$

where $C_1, C_2 < \infty$ are some constants that may depend on ϵ, and $\hat{J}(t,\boldsymbol{x} \mid \Theta)$ is abbreviated as $\hat{J}(\Theta)$ for brevity in some places. The condition (A.62) is used in the second last inequality. By Fubini Theorem,

$$\begin{aligned}\int_{\Omega_{C,T}} \left| \sum_{i=1}^{d} a(t,\boldsymbol{x}) \left(\frac{\partial J(t,\boldsymbol{x})}{\partial x_i} - \frac{\partial \hat{J}(t,\boldsymbol{x})}{\partial x_i} \right) \right|^2 d\mu_1 &\leq \sum_{i=1}^{d} \int_{\Omega_{C,T}} a(t,\boldsymbol{x})^2 \left| \left(\frac{\partial J(t,\boldsymbol{x})}{\partial x_i} - \frac{\partial \hat{J}(t,\boldsymbol{x})}{\partial x_i} \right) \right|^2 d\mu_1 \\ &\leq \sum_{j=1}^{d} \int_{\Omega_{C,T}} A^2 \epsilon^2 d\mu_1 = dA^2 \mu_1(\Omega_{C,T}) \epsilon^2\end{aligned}$$

where the constant $A < \infty$ bounds the function $a(t,\boldsymbol{x})$. Similarly,

$$\begin{aligned}\int_{\Omega_{C,T}} \left| \sum_{i,j=1}^{d} b(t,\boldsymbol{x}) \left(\frac{\partial^2 J(t,\boldsymbol{x})}{\partial x_j x_i} - \frac{\partial^2 \hat{J}(t,\boldsymbol{x})}{\partial x_j x_i} \right) \right|^2 d\mu_1 &\leq \sum_{i=1}^{d} \int_{\Omega_{C,T}} b(t,\boldsymbol{x})^2 \left| \left(\frac{\partial^2 J(t,\boldsymbol{x})}{\partial x_j x_i} - \frac{\partial^2 \hat{J}(t,\boldsymbol{x})}{\partial x_j x_i} \right) \right|^2 d\mu_1 \\ &\leq \sum_{j=1}^{d} \int_{\Omega_{C,T}} B^2 \epsilon^2 d\mu_1 = dB^2 \mu_1(\Omega_{C,T}) \epsilon^2\end{aligned}$$

where the constant $B < \infty$ bounds the function $b(t,\boldsymbol{x})$. Finally, from the condition (A.62), we have

$$\begin{aligned}\int_{\Omega_{C,T}} \left| \frac{\hat{J}(t,\boldsymbol{x}|\Theta)}{\partial t} - \frac{J(t,\boldsymbol{x})}{\partial t} \right|^2 d\mu_1 &\leq \epsilon^2 \mu_1(\Omega_{C,T}) \\ \int_{\Omega_C} \left| \hat{J}(t,\boldsymbol{x}|\Theta) - J_0(t,\boldsymbol{x}) \right|^2 d\mu_2 &\leq \epsilon^2 \mu_2(\Omega_C) \\ \int_{\partial\Omega_T} \left| \frac{\hat{J}(t,\boldsymbol{x}|\Theta)}{\partial \boldsymbol{x}} \right|^2 d\mu_3 &\leq \epsilon^2 \mu_3(\partial\Omega_T)\end{aligned}$$

Putting it together, we have for some constant K

$$\mathscr{L} = \epsilon^2 \left(C_2 + dA^2 \mu_1(\Omega_{C,T}) + dB^2 \mu_1(\Omega_{C,T}) + \mu_1(\Omega_{C,T}) + \mu_2(\Omega_C) + \mu_3(\partial\Omega_T) \right) = K\epsilon^2 \tag{A.65}$$

■

A.2.3 Benchmark model

Model set up

The return on capital held by each type of agent is given by

$$dR_{j,t} = \frac{d(q_t k_{j,t})}{q_t k_{j,t}} + \frac{(a_j - \iota_{j,t})}{q_t} dt$$

where q_t is the price of each unit of capital that follows the process

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

The terms μ_t^q, and σ_t^q are endogenously determined in equilibrium. Using this dynamics for the price, the return process can be written as

$$dR_{j,t} = \underbrace{\left(\frac{a_j - \iota_{j,t}}{q_t} + \Phi(\iota_{j,t}) - \delta + \mu_t^q + \sigma\sigma_t^q\right)}_{\mu_{j,t}^R} dt + (\sigma + \sigma_t^q) dZ_t \tag{A.66}$$

Experts and households trade the capital and the experts are allowed to issue some outside equity. However, they have a skin in the game constraint: i.e, they have to hold at least a fraction $\underline{\chi} \in [0,1]$ of equity in their balance sheet. In addition to the risky capital, the agents also trade a risk free asset that pays a return r_t. Since the markets are not complete, there is no unique stochastic discount factor (SDF). Let $\xi_{e,t}$ and $\xi_{h,t}$ denote the SDF of experts and households respectively. Then, the process for SDF is given as

$$\frac{d\xi_{j,t}}{\xi_{j,t}} = -r_t dt - \zeta_{j,t} dZ_t \tag{A.67}$$

where, $\zeta_{j,t}$ is the market price of risk. Since both agents invest in the risk-free asset, the drift of the SDF process is the same for all agents. The aggregate output in the economy is given by

$$y_t = A_t K_t$$

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

$$A_t = \int_{\mathbb{H}} a_h \frac{k_{j,t}}{K_t} dj + \int_{\mathbb{E}} a_e \frac{k_{j,t}}{K_t} dj$$

Let the capital share held by expert sector be denoted by

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

Equilibrium: The agents optimize by maximising their respective utility functions, subject to the wealth constraints starting from some initial wealth $w_{j,0}$. They solve

$$\sup_{c_{j,t},k_{j,t}} \quad E_t\left[\int_t^\infty f(c_{j,s}, U_{j,s}) ds\right] \tag{A.68}$$

$$\text{s.t.} \frac{dw_{j,t}}{w_{j,t}} = (r_t - \frac{c_{j,t}}{w_{j,t}} + \frac{q_t k_{j,t}}{w_{j,t}}\big((\mu_{j,t}^R - r_t) - (1-\chi_{j,t})\zeta_{j',t}\big) dt + \sigma_{w_j,t}(\sigma + \sigma_t^q) dZ_t \qquad j \in \{e,h\}$$

where $\chi_{j,t}$ denotes the skin-in-the game constraint of agent j, and j' denotes the other type of agent. The aggregator $f(c_{j,s}, U_{j,s})$ is given in equation $(A.22)$. By borrowing in the risk free market at a rate r_t and investing in risky capital, the agents obtain the market price of risk $\zeta_{j,t}$ less the compensation to outside equity holders. Note that since the agents retain only the fraction $\chi_{j,t}$ of risk in their balance sheet, the diffusion terms in wealth equation are given by

$$\sigma_{w_e,t} = \frac{q_t k_{e,t}}{w_{e,t}} \chi_t \tag{A.69}$$

$$\sigma_{w_h,t} = \frac{q_t k_{h,t}}{w_{h,t}} + (1-\chi_t)\frac{q_t w_{e,t}}{w_{h,t}} \tag{A.70}$$

The households do not issue outside equity and therefore $\chi_{h,t} = 1$. For the simplicity of notation, I denote $\chi_{e,t}$ as χ_t henceforth. The asset pricing conditions for the experts and the households are given by[10]

$$\frac{\frac{a_e - \iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_{q,t} - r_t}{\sigma + \sigma_{q,t}} = \chi_t \zeta_{e,t} + (1-\chi_t)\zeta_{h,t} \tag{A.71}$$

$$\frac{\frac{a_h - \iota_t}{q_t} + \Phi(\iota_t) - \delta + + \mu_t^q + \sigma\sigma_t^a - r_t}{\sigma + \sigma_t^q} \leq \zeta_{h,t} \tag{A.72}$$

When the risk premium demanded by the experts is large, they will sell maximum allowed equity to the households. Since the households do not issue outside equity, their asset pricing condition is simpler. The equality holds in (A.72) if the households own some amount of capital ($\psi_t < 1$). Combining the asset pricing conditions, we have

$$\frac{a_e - a_h}{q_t} \geq \chi_t(\zeta_{e,t} - \zeta_{h,t}) \tag{A.73}$$

$$\min\{\chi_t - \underline{\chi}, \zeta_{e,t} - \zeta_{h,t}\} = 0 \tag{A.74}$$

The equality in (A.73) holds when both the experts and the households hold capital. In this region, the experts issue maximum allowed equity ($\chi_t = \underline{\chi}$) as dictated by the condition (A.74). This is typically when the wealth share is low and the economy is in crisis state. The second region is when the premium of experts is still higher than that of the households but all capital is held by the experts and the economy is out of the crisis state. Since the premium of experts is higher, the experts issue maximum equity in this region as well. In the third region, there is perfect risk sharing where the premium of both type of agent becomes equal and χ_t is chosen to be equal to the wealth share z_t.

Solving the model: There are in fact an infinite number of agents in the economy but each individual in type $\mathbb{E}$ and $\mathbb{H}$ are identical and have the same preferences. Therefore, one can seek an equilibrium in which all agents in the same group take the same policy decisions. The

[10] This can be proved using the Martingale argument. See Appendix A.1.2 for the proof.

system can be summarized with only one state variable: the wealth share of the experts which is sufficient to characterize the wealth distribution agents. It is defined as

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

where $W_{e,t} = \int_{\mathbb{E}} w_{j,t} dj$ and $K_t = \int_{\mathbb{E}} k_{j,t} dj + \int_{\mathbb{H}} k_{j,t} dj$. Moving forward, we write $X_{h,t}$ and $X_{e,t}$ to denote the aggregated quantity $\int_{\mathbb{H}} x_{j,t} dj$ and $\int_{\mathbb{E}} x_{j,t} dj$ respectively.

Proposition 10. *The law of motion of the wealth share of experts is given by*

$$\frac{dz_t}{z_t} = \mu_t^z dt + \sigma_t^z dZ_t \tag{A.75}$$

where

$$\mu_t^z = \frac{a_e - \iota_t}{q_t} - \frac{C_{e,t}}{W_{e,t}} + (\frac{\chi_t \psi_t}{z_t} - 1)(\sigma + \sigma_{q,t})(\zeta_{e,t} - (\sigma + \sigma_t^q)) + (1-\chi_t)(\sigma + \sigma_t^q)(\zeta_{e,t} - \zeta_{h,t})$$
$$\sigma_t^z = (\frac{\chi_t \psi_t}{z_t} - 1)(\sigma + \sigma_t^q)$$

Proof: The law of motion of agents aggregated by their type is given by

$$\frac{dW_{h,t}}{W_{h,t}} = \Big(r_t - \frac{C_{h,t}}{W_{h,t}} + \frac{1-\chi_t\psi_t}{1-z_t}(\mu_{h,t}^R - r_t)\Big)dt + \frac{1-\chi_t\psi_t}{1-z_t}(\sigma + \sigma_t^q)dZ_t^k$$
$$\frac{dW_{e,t}}{W_{e,t}} = \Big(r_t - \frac{C_{e,t}}{W_{e,t}} + \frac{\chi_t\psi_t}{z_t}(\mu_{e,t}^R - r_t)\Big)dt + \frac{\chi_t\psi_t}{z_t}(\sigma + \sigma_t^q)dZ_t^k$$

where $W_{h,t} = \int_{j\in\mathbb{H}} w_{j,t} dj$ and similarly for the experts, and the expressions $\frac{q_t K_{e,t}}{W_{e,t}} = \frac{\psi_t}{z_t}$ and $\frac{q_t K_{h,t}}{W_{h,t}} = \frac{1-\psi_t}{1-z_t}$ are used along with the definition of z_t. ■

A.2.4 Brunnermeier-Sannikov meets Bansal-Yaron

Proof of asset pricing conditions

Consider the problem of experts first. The expected return earned from investing in the risky capital is given by

$$dr_t^v = (\mu_{e,t}^R - (1-\chi_t)\bar{\epsilon}_{h,t})dt + \chi_t(\boldsymbol{\sigma_t^R})^T \boldsymbol{dZ_t}$$

where χ_t is the experts' inside equity share. The experts have to pay the outside equity holders $(1-\chi_t)\bar{\epsilon}_{h,t}$ from the expected return that they earn, and hence this part is netted out from the drift. Since they are only exposed to a fraction χ_t of their total investment in the risky capital, the diffusion terms are multiplied by this fraction. For an investment of \$1 in the risky capital, the value of the investment strategy is given by

$$\frac{d(\xi_{e,t} v_t)}{\xi_{e,t} v_t} = (-r_t + \mu_{e,t}^R - (1-\chi_t)\bar{\epsilon}_{h,t} - \chi_t \bar{\epsilon}_{e,t})dt + \text{difusion terms}$$

where the $\bar{\epsilon}_{e,t} = \boldsymbol{\zeta}_{\boldsymbol{e,t}}^T \boldsymbol{\sigma}_{\boldsymbol{t}}^{\boldsymbol{R}}$ and $\xi_{e,t}$ is the SDF process given by (A.67). Since the stochastic process ν_t is a martingale, the drift term should be zero. This gives us

$$\mu_{e,t}^R - t_t = \chi_t \bar{\epsilon}_{e,t} + (1-\chi_t)\bar{\epsilon}_{h,t}$$

The asset pricing condition for the households follows in a similar fashion except that they do not issue outside equity, and hence we arrive at

$$\mu_{h,t}^R - r_t = \bar{\epsilon}_{h,t}$$

where $\bar{\epsilon}_{h,t} = \boldsymbol{\zeta}_{\boldsymbol{h,t}}^T \boldsymbol{\sigma}_{\boldsymbol{t}}^{\boldsymbol{R}}$ ■

Proof of Proposition 8

The law of motion of the wealth obtained by aggregating wealth and using law of large numbers is given by

$$\frac{dW_{h,t}}{W_{h,t}} = \big(r_t - \hat{C}_{h,t} + \theta_{h,t}(\mu_{h,t}^R - r_t) + \tau_t \frac{W_{e,t}}{W_{h,t}}\big)dt + \theta_{h,t}(\boldsymbol{\sigma}_{\boldsymbol{t}}^{\boldsymbol{R}})^T \boldsymbol{dZ_t} \tag{A.76}$$

$$\frac{dW_{e,t}}{W_{e,t}} = \big(r_t - \hat{C}_{e,t} + \theta_{e,t}\bar{\epsilon}_{e,t} - \tau_t\big)dt + \theta_{e,t}(\boldsymbol{\sigma}_{\boldsymbol{t}}^{\boldsymbol{R}})^T \boldsymbol{dZ_t} \tag{A.77}$$

where $W_{h,t} = \int_{\mathbb{H}} w_{j,t} dj$ and $W_{e,t} = \int_{\mathbb{E}} w_{j,t} dj$ are the aggregate wealth within the respective class. The wealth share of the experts is defined to be $z_t = \frac{W_{e,t}}{W_{e,t}+W_{h,t}}$. Also, the variables $\theta_{e,t} := \frac{\chi_t \psi_t}{z_t}$ and by market clearing, $\theta_{h,t} = \frac{1-\chi_t\psi_t}{1-z_t}$

Applying Ito's lemma, we get

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

where

$$\frac{d(q_t k_t)}{q_t k_t} = (\bar{\epsilon}_{e,t}(\sigma + \sigma_t^q) - \frac{(a_{e,t} - \iota_t)}{q_t} + r_t)dt + (\boldsymbol{\sigma}_{\boldsymbol{t}}^{\boldsymbol{R}})^T \boldsymbol{dZ_t}$$

$$\frac{d\langle q_t k_t, q_t k_t\rangle}{(q_t k_t)^2} = ||\boldsymbol{\sigma}_{\boldsymbol{t}}^{\boldsymbol{R}}||^2 dt$$

$$\frac{d\langle q_t k_t, w_{e,t}\rangle}{q_t k_t w_{e,t}} = \big(\theta_{e,t}||\boldsymbol{\sigma}_{\boldsymbol{t}}^{\boldsymbol{R}}||^2\big)dt$$

We get the desired result after few steps of algebra. ■

Proof of Proposition 5

The conjecture for the value function is

$$U_{j,t} = \frac{(J_{j,t}(\boldsymbol{x})K_t)^{1-\gamma}}{1-\gamma} \tag{A.78}$$

where K_t is the aggregate capital, and the the stochastic opportunity set $J_{j,t}$ satisfies the equation

$$\frac{dJ_{j,t}}{J_{j,t}} = \mu^J_{j,t}dt + (\boldsymbol{\sigma^J_t})^T \boldsymbol{dZ_t}$$

The objects $\mu^J_{j,t}$ and $\boldsymbol{\sigma^R_t}$ needs to be determined in equilibrium. Applying Ito's lemma to $U_{j,t}$ and using the HJB equation $\sup_{c,k} f(c_{j,t}, U_{j,t}) + E[dU_{j,t}] = 0$, with

$$f(c_{j,t}, U_{j,t}) = (1-\gamma)\rho U_{j,t}\left(\log c_{j,t} - \frac{1}{1-\gamma}\log\big((1-\gamma)U_{j,t}\big)\right)$$

we get

$$\begin{aligned}\sup_C \quad & \rho(J_{j,t}K_t)^{1-\gamma}[\log\frac{C_{j,t}}{W_{j,t}} - \log J_{j,t} + \log(q_t z_{j,t})] + \frac{(J_{j,t}K_t)^{1-\gamma}}{1-\gamma}\mu^J_{j,t} + (J_{j,t}K_t)^{1-\gamma}(\Phi(\iota_t - \delta)) \\ & - (J_{j,t}K_t)^{1-\gamma}\frac{1}{2}\gamma\sigma^2 + (1-\gamma)(J_{j,t}K_t)^{1-\gamma}\sigma\sigma^{J,k}_{j,t} - \frac{\gamma}{2}(J_{j,t}K_t)^{1-\gamma}||\sigma^J||^2 + \tau_t(U_{h,t} - U_{e,t}) = 0\end{aligned} \tag{A.79}$$

By envelope condition, the marginal utilities of wealth and consumption should equal at the optimum. Since $\tilde{J} = \frac{J}{qz}$, we can rewrite

$$U_{j,t} = \frac{(\tilde{J}_{j,t}W_{j,t})^{1-\gamma}}{1-\gamma}; \qquad f(C_{j,t}, U_{j,t}) = (1-\gamma)\rho U_{j,t}(log\frac{C_{j,t}}{W_{j,t}} - \tilde{J}_{j,t}) \tag{A.80}$$

Using this, we have

$$\frac{\partial U_{j,t}}{\partial W_{j,t}} = \frac{\partial f_{j,t}}{\partial C_{j,t}} \implies (\tilde{J}_{j,t}W_{j,t})^{-\gamma} = (1-\gamma)\rho\frac{U_{j,t}}{C_{j,t}} \implies \frac{C_{j,t}}{W_{j,t}} = \rho$$

That is, the optimal consumption-wealth ratio is equal to the discount rate. The SDF for recursive utility is expression as

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

Utilizing (A.80), we get

$$\xi_{j,t} = exp\left(\int_0^t \left[(1-\gamma)\rho(\log\rho - \tilde{J}_{j,t})\right]ds\right)\frac{U_{j,t}}{W_{j,t}}$$

Thus, the volatility of the SDF is equal to the volatility of the quantity $\frac{U}{W}$. Let us define $v(J,\boldsymbol{x}) := \frac{U}{W}$. Using Ito's lemma, equating the coefficients of volatility terms to the volatility of SDF equation (A.67), we get the result. ■

Proof of Proposition 6

The first equation (2.58) comes from plugging in the risk premium from (5) in the asset pricing condition (A.73). The volatility of the opportunity set σ^J in the risk premium is expressed in terms of the partial derivative of J with respect to the state variables. This can be easily derived using Ito's lemma and comparing the diffusion terms of (2.30) and $dJ(z,\boldsymbol{x})$. The second equation (2.59) comes from the capital market clearing condition and using $z_t = \frac{W_t}{q_t K_t}$, $\psi_t = \frac{K_{e,t}}{q_t K_t}$. To derive (2.60) and (2.61), first apply Ito's lemma to $q(z_t,\boldsymbol{x})$ to get

$$\begin{aligned} dq(\boldsymbol{x}) &= \frac{\partial q}{\partial \boldsymbol{x}}\boldsymbol{dx} + \frac{1}{2}\frac{\partial^2 q}{\partial^2 \boldsymbol{x}^2}d\langle \boldsymbol{x},\boldsymbol{x}\rangle \\ &= \text{drift terms} + \frac{\partial q}{\partial \boldsymbol{x}}\boldsymbol{\sigma^x} \end{aligned}$$

Matching the volatility terms with the capital price equation (2.4), we get the result. ■

Proof of Proposition 7

Applying Ito's lemma to $J(t,\boldsymbol{x})$, we get

$$\begin{aligned} dJ(\boldsymbol{x}) &= \frac{\partial J}{\partial \boldsymbol{x}}d\boldsymbol{x} + \frac{1}{2}\frac{\partial^2 J}{\partial \boldsymbol{x}^2}d\langle \boldsymbol{x},\boldsymbol{x}\rangle)dt \\ &= \left(\frac{\partial J}{\partial \boldsymbol{x}}\boldsymbol{\mu^x} + \frac{1}{2}\frac{\partial^2 J}{\partial \boldsymbol{x}^2}(\boldsymbol{\sigma^x})^2\right)dt + \text{volatility terms} \end{aligned}$$

Comparing this with the equation (2.30) and matching the drift terms, we get the expression

$$J\mu^J = \frac{\partial J}{\partial \boldsymbol{x}}\boldsymbol{\mu^x} + \frac{1}{2}\frac{\partial^2 J}{\partial \boldsymbol{x}^2}(\boldsymbol{\sigma^x})^2$$

Adding the fase time-derivative, it remains to derive the quantity μ^J which can be obtained from the HJB equation (A.79). The term $A\left(\boldsymbol{x}, J, \frac{\partial J}{\partial \boldsymbol{x}}\right)$ includes both μ^J, as well as the diffusion term $\frac{\partial J}{\partial \boldsymbol{x}}\boldsymbol{\mu^x}$. The term $B\left(\boldsymbol{x}, J, \frac{\partial J}{\partial \boldsymbol{x}}\right)$ represents the diffusion terms $(\boldsymbol{\sigma^x})^2$. This proves the proposition. ■

Implementation details

Data efficiency: One of the main advantages of using neural network to fit the PDE is data efficiency. For the benchmark model, I use 1000 grid points for space dimension in the inner Newton-Raphson method. When it comes to training the neural network, I randomly sample 300 points in each iteration. At kth iteration, the function to be solved is denoted by $J(T-k\Delta t, z)$ whose economic behavior is governed by the given PDE in the domain $[T-k\Delta t, T-(k-1)\Delta t]\times\Omega_z$. I randomly sample time points from the range $[T-k\Delta t, T-(k-1)\Delta t]$ in order to avoid errors in computing gradients with respect to time dimension. The 300 grid points include these time points as well. Figure () presents the sparse grid used for training. The task is to solve for the function $J(T-k\Delta t, z)$ such that the PDE is respected in the domain $[T-k\Delta t, T-(k-1)\Delta t]$ along with bounding conditions in the domain $(T-(k-1)\Delta t)\times\Omega$ and $(T-(k-1)\Delta t)\times\partial\Omega$.

Simplicity in coding: I rely on Tensorflow, a popular library developed by Google to copmute derivatives in an efficient way. The method $tf.gradients$ computes the required symbolic partial derivatives, which are then collected to form the PDE loss residual. In a similar fashion, the bounding network and active network can computed. An important thing to note is that the module $tf.gradients$ creates a computational graph and does not perform the calculation yet. Once all networks in ALIEN are built, one can start the tensoflow session which then initiates the computation of gradients. This allows us to build the model first and then distribute the data efficiently as demonstrated later. The code snipper (A.1) shows how to approximate the function J using a neural network. The inputs are the state variables along with the weights and biases which are the parameters of the neural network. Before training begins, the weights are initialized using Xavier initialization as explained earlier. The snippet (A.2) demonstrates the computation of PDE network. The inputs are the function $\hat{J}$ approximated using the neural network along with advection, diffusion, and linear terms which are PDE coefficients.[11] The gradients of the approximated function $\hat{J}$ with respect to the state variables are computed using automatic differentiation through the tensorflow module $tf.gradients$. Note that automatic differentiation is commonly used in machine learning to obtain derivatives of functions with respect to the neural network parameters. Here, I utilize automatic differentiation to obtain derivatives with respect to the state variables. Apart from this difference, the gradient computation is standard. One can notice that the coding effort involved in computing the derivatives is very minimal.

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

```

Listing A.1 – Approximating J using a neural network: Benchmark model

```
def f(z,t):
  #compute fundamental network Jhat
```

[11] The advection, diffusion, and linear term coefficients are calculated before starting the training process and hence they can be simply passed as inputs into the neural network algorithm.

```
  J = J(z,t)
  #compute first partial derivatives
  J_t = tf.gradients(J,t)[0]
  J_z = tf.gradients(J,z)[0]
  #compute second partial derivatibes
  J_zz = tf.gradients(J_z,z)[0]
  #compute PDE residual
  f = J_t + advection * J_z + diffusion * J_zz - linearTerm * J
  return f
```

Listing A.2 – Constructing a regularizer: Benchmark model

While the benchmark model can be solved using traditional methods discussed earlier, it is not trivial to extend these methods to solve models in higher dimensions. In finite difference schemes, it is not only problematic to maintain the monotonicity but also difficult to code especially in the case of implicit schemes. For example, Hansen et al. (2018) solves a collection of nested macro-finance models in 3 dimensions which involves setting up dimension-specific matrices to solve the PDEs. It is not only difficult to code but also requires high performance computing libraries such as Paradiso (specific to C++) to solve large linear systems that show up in the implicit finite difference scheme. In contract, the framework proposed in this paper involves less coding effort in scaling to higher dimensions. To appreciate the simplicity, I demonstrate sample codes from the capital misallocation model considered in Section 2.5. In code snippet (A.3), the function J is approximated using the neural network. The inputs are 4 state variables and 1 time dimension, along with weights and biases which are the parameters of the network. As before, the weights go through Xavier initialization before the learning begins. The snippet (A.4) constructs the regularizer corresponding to the PDE network. The inputs are approximated function $\hat{J}$ along with the PDE coefficients that are known. Using automatic differentiation, one can easily obtain the partial derivatives and compute the PDE residual. It only takes a few additional lines of coding to move from a 1 dimensional model to 4 dimensional model. In contrast, it is not at all trivial to move easily to higher dimensions using finite difference schemes. The ease in implementation shifts the burden from the modeler to Tensorflow thereby freeing up time to focus on more important issues from an economic standpoint.

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

```

Listing A.3 – Approximating J using a neural network: 4D model

```
def net_f(z,g,s,a,t):
        #compute fundamental network Jhat
        J = J(z,g,s,a,t)
        #compute first partial derivatives
        J_z, J_g = tf.gradients(J,z)[0], tf.gradients(J,g)[0]
        J_s, J_a = tf.gradients(J,s)[0], tf.gradients(J,a)[0]
        J_t = tf.gradients(J,t)[0]
        #compute second partial derivatives
        J_zz = tf.gradients(J_z,z)[0]
        J_gg = tf.gradients(J_g,g)[0]
        J_ss = tf.gradients(J_s,s)[0]
        J_aa = tf.gradients(J_a,a)[0]
        J_zg = tf.gradients(J_z,g)[0]
        J_zs = tf.gradients(J_z,s)[0]
```

```
        J_za = tf.gradients(J_z,a)[0]
        #compute PDE residual
        f =  J_t + diffusion_z * J_zz +  diffusion_g * J_gg + \
                   diffusion_s * J_ss + diffusion_a * J_aa + \
                   advection_z * J_z + advection_g * J_g + \
                   advection_s * J_s + advection_a * J_a + \
                   cross_term_zg * J_zg + cross_term_zs * J_zs + \
                   cross_term_za * J_za -  linearTerm * J
        return f
```

Listing A.4 – Constructing a regularizer: 4D model

Distributed learning: The concept of distributed learning is not new and entails utilization of multiple workers or GPUs to speed up computation. Specifically, data parallelism works by dividing up the data into pieces (or mini-batches in the language of machine learning) and sending to different workers that will run the data through the same model. Algorithm 2 presents a simple data parallelism approach. There are a few bottlenecks presented by this procedure. First, it requires the user to employ a library that can communicate across workers. Secondly, and more importantly, the communication overhead resulting from the cross-worker interface may be significant thereby defeating the purpose of using a distributed algorithm. For example, Sergeev and Balso (2018) finds that roughly half of the computational resources are lost due to the overhead when they train a big data model on 128 GPUs. The reason for such heavy overhead is that the default way of communicating across workers is through a parameter sharing approach, where each node assumes the role of either a worker or a parameter server. The role of worker is to train the model, and the parameter server aggregates the gradients. The user is left to decide the optimal ratio of parameter server to worker. A small ratio leads to a large computational bottleneck, and a large ratio leads to communication overhead. Andrew Gibiansky (2017) at Baidu presents an inter-worker communication algorithm that bypasses these problems. It is based on ring-AllReduce, where each worker communicates with its two neighbours only in a ring-like fashion for a total of $2*(N-1)$ times. During the first $N-1$ communications, each worker sends its data to the neighbours, and receives the data from the neighbors to store it in the buffer. In the next $N-1$ communications, the workers receive the data from the neighbours and update the buffer. This algorithm is shown to the optimal one for the utilization of the bandwidth provided the buffer is large enough.

Algorithm 2

procedure DISTRIBUTED ALGORITHM
    Assign a chief worker
    Divide data by number of workers
    while $worker < N$ **do** ▷ N is number of workers
        Assign model to current worker
        Run the data through the model
        Compute gradients
        Send gradients to chief worker
    Average gradients from multiple workers
    Update the model

Horovod: Sergeev and Balso (2018) leverages the advantages of ring-AllReduce algorithm and combines it with Tensorflow to build libraries that facilities easy implementation of distributed learning. The user only has to add a few lines to code to enable a hybrid parallelization of the deep learning algorithm. Algorithm 3 presents the pseudo-code for implementing Horovod.

Algorithm 3

procedure HOROVOD
    Initialize Horovod
    Assign a GPU to each tensorflow process
    Start a tensorflow session
    Split data based on number of workers
    Build deep learning model and set up loss functions
    Wrap the optimizer with Horovod optimizer ▷ To average gradients
    Initiate Tensorflow session to train the deep learning model
    Broadcast variables from chief worker to all other workers ▷ This makes sure that all workers have the same initial parameters in the model.
    Train the model until convergence

The code snippet (A.5) presents the implementation of Horovod to ALIENs. Each line in the pseudo-code (3) can be implemented simply by calling a module in Horovod library as seen in the code snippet. This again shifts the burden from the modeler to the library that not only frees up the modeler's time but also eliminates the necessity to deal with inter-worker communication issues that are rarely of interest to an economist.

```
def J():
  ...
def net_f():
  ...

hvd.init() #initialize Horovod
config = tf.ConfigProto() #pin GPUs to processes
config.gpu_options.visible_device_list = str(hvd.local_rank()) #assign
  chief worker
```

```
config.gpu_options.allow_growth = True #enable GPU
sess= tf.Session(config=config) #Configure tensorflow
if hvd.rank()==0:
  ... #assign a piece of data to chief worker
else:
  while hvd.rank() < hvd.size():
    ... #assign a piece of data to each worker

def build_model():
  #initialize parameters using Xavier initialization
  #parametrize the function J using J()
  #buld loss function using net_f()
  #set up tensorflow optimizer in the variable name opt
  optimizer = hvd.DistributedOptimizer(opt)
  #minimize loss
  #initialize Tensorflow session
  bcast = hvd.broadcast_global_variables(0) #Broadcast parameters to all
  workers
  sess.run(bcast)
  #train the deep learning model
```

Listing A.5 – ALIENs using Horovod

A.3 Appendix to Chapter 3

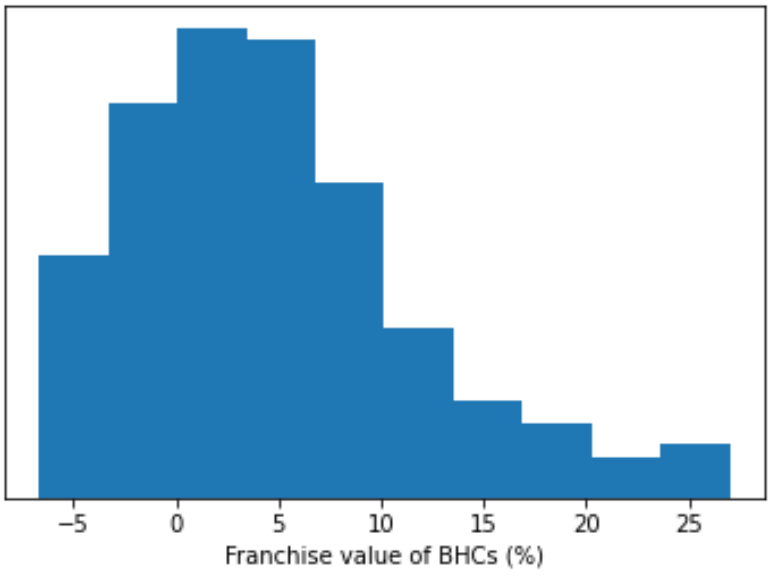

Figure A.1 – The figure presents the histogram of franchise value of all listed BHCs in the US between the year 2001 and 2018. The franchise value is computed as Market value of equity + Book value of liabilities - (Book value of asset - Goodwill) scaled by (Book value of assets - Goodwill). The data source is given in Table (3.1).

Table A.3 – This table estimates the determinants of franchise value by regressing $f_{i,t}$ on $iCost/assets$ and various controls where all regressors are lagged by one quarter. The data is for all listed US Bank Holding companies from the year 2001 till 2018 at quarterly frequency. All variables are winsorized at 1% level. Standard errors are robust to heteroskedasticity in errors.

	(1)	(2)	(3)
	$f_{i,t}$	$f_{i,t}$	$f_{i,t}$
logAssets	0.0033**	0.001***	0.001**
	(0.00)	(0.00)	(0.038)
iCost/assets		-0.57***	-0.56**
		(0.00)	(0.02)
deposits/assets		0.016***	-0.002
		(0.00)	(0.70)
domestic assets/assets		-0.024***	-0.03***
		(0.00)	(0.00)
leverage		-0.04***	-0.03 *
		(0.00)	(0.08)
deposit power		1.42***	2.37***
		(0.00)	(0.00)
govt. guarantee		0.54***	0.50***
		(0.00)	(0.00)
capital ratio			0.0002
			(0.34)
N	22,009	22,009	5,138
Fixed Effects	Bank	Bank	Bank
R^2	0.73	0.79	0.79

p-values in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A.4 – This table studies the relationship between bank probability of default and franchise value. The probability of default is proxied by a dummy variable d_i that takes value 1 if the bank i filed for bankruptcy in the time period considered. The data is for all listed US Bank Holding Companies from the year 2001 till 2018 at quarterly frequency. All variables are winsorized at 1% level. Standard errors are robust to heteroskedasticity in errors.

	(1)	(2)	(3)
	d_i	d_i	d_i
log Assets	0.00696***	0.0118***	0.0151***
	(0.000)	(0.000)	(0.000)
franchise value		-0.448***	-0.472***
		(0.000)	(0.004)
deposit/assets		-0.139***	-0.168*
		(0.002)	(0.051)
domestic/assets		0.0968**	0.156*
		(0.035)	(0.063)
trading assets/assets		2.648***	1.484*
		(0.000)	(0.095)
interest income/operating income		-0.00640	-0.0245
		(0.622)	(0.738)
leverage		-0.111***	-0.191**
		(0.000)	(0.014)
capital ratio			-0.00102
			(0.640)
_cons	-0.0718***	-0.0315	-0.0142
	(0.000)	(0.518)	(0.907)
N	31516	23925	5154
Fixed Effect	Time	Time	Time
R^2	0.005	0.052	0.052

p-values in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A.5 – This table studies the relationship between bank probability of default and franchise value. The probability of default is proxied by a dummy variable a_i that takes value 1 if the bank i sought assistance from FDIC to continue as ongoing concern in the time period considered. The data is for all listed US Bank Holding Companies from the year 2001 till 2018 at quarterly frequency. All variables are winsorized at 1% level. Standard errors are robust to heteroskedasticity in errors.

	(1)	(2)	(3)
	a_i	a_i	a_i
log Assets	0.00832***	0.00565***	0.0106***
	(0.000)	(0.000)	(0.000)
franchise value		-0.329***	-0.266***
		(0.000)	(0.000)
deposit/assets		0.0396***	0.00958
		(0.000)	(0.554)
foreign/assets		-0.0987***	-0.0404**
		(0.000)	(0.038)
trading/assets		3.305***	2.384***
		(0.000)	(0.004)
interest income/operating income		0.0228	0.0801***
		(0.137)	(0.001)
leverage		-0.180***	-0.115**
		(0.000)	(0.019)
capital ratio			0.00305***
			(0.000)
_cons	-0.108***	0.148***	-0.0980
	(0.000)	(0.000)	(0.232)
N	31516	23925	5154
Fixed Effect	Time	Time	Time
R^2	0.021	0.096	0.091

p-values in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

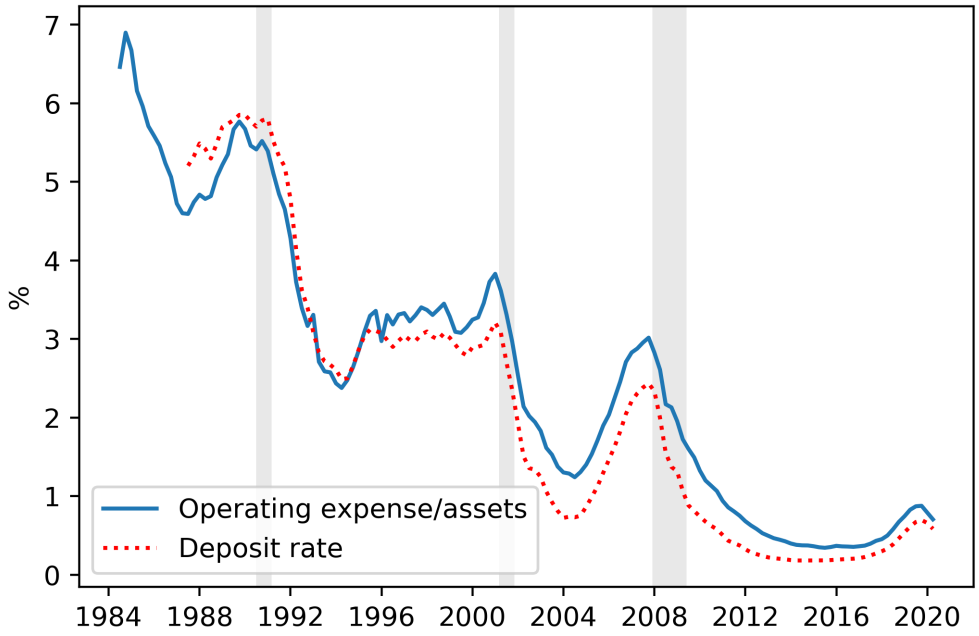

Figure A.2 – This figure presents the time series of operating expense to asset ratio for the banks and the deposit rate, both averaged across all BHCs in the US. The data is at quarterly frequency between the period 1986Q1 to 2020Q4. The data source is given in Table (3.1). The values are annualized and is in percentage terms. The shaded background represent NBER recessionary periods.

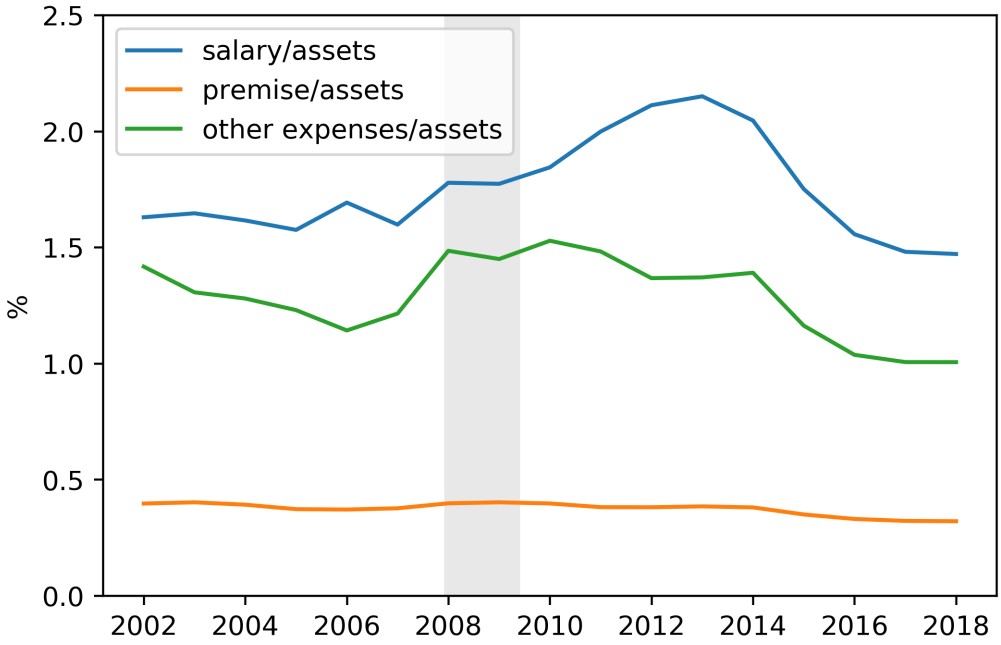

Figure A.3 – This figure presents the time series of individual components of operating expenses scaled by total assets for the banks and the deposit rate, both averaged across all BHCs in the US. The data is at quarterly frequency between the period 1986Q1 to 2018Q4. The data source is given in Table (3.1). The values are annualized and is in percentage terms. The shaded background represent NBER recessionary periods.

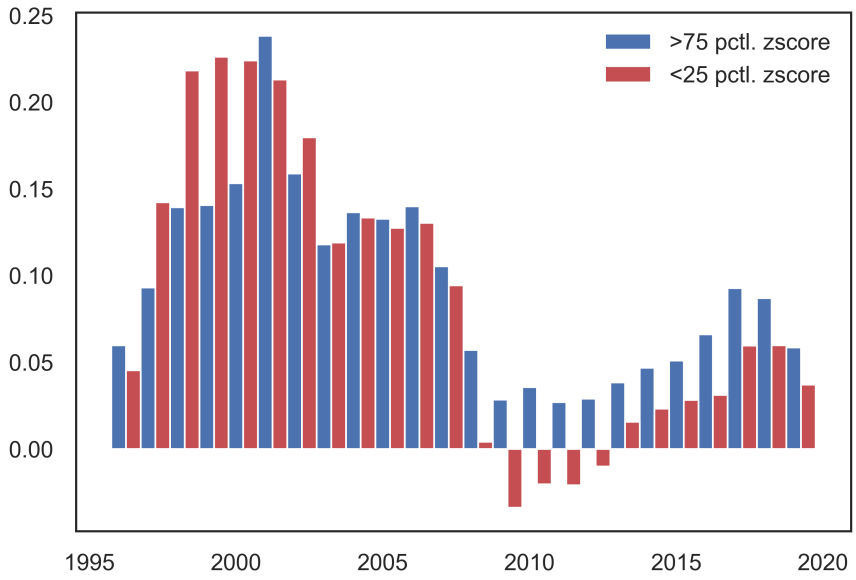

Figure A.4 – Franchise value of BHCs in the US. Blue bars correspond to banks with z-score value in top 25-percentile, and red bars correspond to the banks with z-score in bottom 25-percentile. Franchise value is computed as Market value of equity + Book value of liabilities - (Book value of asset - Goodwill) scaled by (Book value of assets - Goodwill). The data is at quarterly frequency. The data source is given in Table (3.1).

Bibliography

Achdou, Y., Buera, F. J., Lasry, J.-M., Lions, P.-L., and Moll, B. (2014a). Partial differential equation models in macroeconomics. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 372(2028).

Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L., and Moll, B. (2014b). Heterogeneous agent models in continuous time. *Preprint*.

Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L., and Moll, B. (2017). Income and wealth distribution in macroeconomics: A continuous-time approach. *NBER Working Paper Series*.

Adrian, T. and Boyarchenko, N. (2012). Intermediary Leverage Cycles and Financial Stability. *SSRN Electronic Journal*, (August 2012).

Adrian, T., Etula, E., and Muir, T. (2014). Financial Intermediaries and the Cross-Section of Asset Returns. *Journal of Finance*, 69(6):2557–2596.

Andrew Gibiansky (2017). Bringing HPC techniques to deep learning. *http://research.baidu.com/bringing-hpc-techniques-deep-learning. [Online; accessed 10-June-2019]*.

Atkeson, A. G., D'avernas, A., Eisfeldt, A. L., Weill, P.-O., and Hall, B. (2018). Government Guarantees and the Valuation of American Banks. *NBER Working Paper 24706*.

Azinovic, M., Gaegauf, L., and Scheidegger, S. (2019). Deep Equilibrium Nets. *SSRN Electronic Journal*.

Bansal, R. and Yaron, A. (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance*, 59(4):1481–1509.

Baron, M. and Xiong, W. (2017). Credit expansion and neglected crash risk. *Quarterly Journal of Economics*, 132(2):713–764.

Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics*, 121(3):823–866.

Basak, S. and Cuoco, D. (1998). An equilibrium model with restricted stock market participation. *Review of Financial Studies*.

Basak, S. and Shapiro, A. (2001). Value-at-risk-based risk management: Optimal policies and asset prices. *Review of Financial Studies*.

Ben S. Bernanke (1983). Nonmonetary Effects of the Financial Crisis in the Propagation of the Great Depression on JSTOR. *The American Economic Review Vol. 73, No. 3 (Jun., 1983), pp. 257-276 (20 pages)*, 73(3):257–276.

Berger, A. N., Bouwman, C. H., Kick, T., and Schaeck, K. (2016). Bank liquidity creation following regulatory interventions and capital support. *Journal of Financial Intermediation*, 26:115–141.

Bernanke, B. S., Bernanke, B. S., Gertler, M., Gertler, M., Gilchrist, S., and Gilchrist, S. (1998). The Financial Accelerator in a Business Cycle Framework. *NBER Working Paper*.

Bigio, S. and D'Avernas, A. (2021). Financial Risk Capacity. *American Economic Journal: Macroeconomics*, 13(4):142–81.

Bolton, P., Chen, H., and Wang, N. (2011). A Unified Theory of Tobin'sq, Corporate Investment, Financing, and Risk Management. *Journal of Finance*.

Bonnans, J. F., Ottenwaelter, É., and Zidani, H. (2004). A fast algorithm for the two dimensional HJB equation of stochastic control. *Mathematical Modelling and Numerical Analysis*, 38(4):723–735.

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. (2020). Language models are few-shot learners.

Brumm, J. and Scheidegger, S. (2017). Using Adaptive Sparse Grids to Solve High-Dimensional Dynamic Models. *Econometrica*.

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

Brunnermeier, M. K. and Sannikov, Y. (2014b). A Macroeconomic Model with a Financial Sector. *American Economic Review*, 104(2):379–421.

Brunnermeier, M. K. and Sannikov, Y. (2016a). Macro, Money, and Finance: A Continuous-Time Approach. In *Handbook of Macroeconomics*.

Brunnermeier, M. K. and Sannikov, Y. (2016b). Macro, Money, and Finance: A Continuous-Time Approach. In *Handbook of Macroeconomics*.

Bungartz, H. J., Heinecke, A., Pflüger, D., and Schraufstetter, S. (2012). Option pricing with a direct adaptive sparse grid approach. In *Journal of Computational and Applied Mathematics*.

Caballero, R. J. and Simsek, A. (2017). A Risk-Centric Model of Demand Recessions and Macroprudential Policy. *SSRN Electronic Journal*.

Caiazza, S., Cotugno, M., Fiordelisi, F., and Stefanelli, V. (2018). The spillover effect of enforcement actions on bank risk-taking. *Journal of Banking and Finance*, 91:146–159.

Campbell, J. Y. and Cochrane, J. H. (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–251.

Chen, L., Pelger, M., and Zhu, J. (2019). Deep Learning in Asset Pricing. *SSRN Electronic Journal*.

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

D'Avernas, A., Vandeweyer, Q., and Darracq Paries, M. (2019). Unconventional Monetary Policy and Funding Liquidity Risk. *SSRN Electronic Journal*.

De Nicolo, G. and Zotova, V. (2020). Bank Risk and Bank Rents: The Franchise Value Hypothesis Reconsidered. *SSRN Electronic Journal*.

Delis, M. D., Hasan, I., and Tsionas, E. G. (2014). The risk of financial intermediaries. *Journal of Banking and Finance*.

Demsetz, R. S., Saidenberg, M. R., and Strahan, P. E. (1996). Banks with Something to Lose: The Disciplinary Role of Franchise Value. *ECONOMIC POLICY REVIEW*.

Di Tella, S. (2017). Uncertainty shocks and balance sheet recessions. *Journal of Political Economy*, 125(6):2038–2081.

Di Tella, S. (2019). Optimal regulation of financial intermediaries. *American Economic Review*.

Dindo, P., Modena, A., and Pelizzon, L. (2022). Risk pooling, intermediation efficiency, and the business cycle. *Journal of Economic Dynamics and Control*, page 104500.

Drechsler, I., Savov, A., and Schnabl, P. (2018). A Model of Monetary Policy and Risk Premia. *Journal of Finance*.

Drechsler, I., Savov, A., and Schnabl, P. (2021). Banking on Deposits: Maturity Transformation without Interest Rate Risk. *The Journal of Finance*, 76(3):1091–1143.

Duarte, V. (2017). Machine Learning for Continuous-Time Economics.

Egan, M., Lewellen, S., and Sunderam, A. (2021). The Cross-Section of Bank Value. *The Review of Financial Studies*.

Eisfeldt, A. L., Lustig, H., and Zhang, L. (2017). Complex Asset Markets. *NBER Working Papers*, 123(5):1177–1200.

Feng, G. and Serletis, A. (2010). 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.

Fernández-Villaverde, J., Hurtado, S., and Nuno, G. (2020). Financial Frictions and the Wealth Distribution. *SSRN Electronic Journal*.

Gârleanu, N. and Panageas, S. (2015). Young, old, conservative, and bold: The implications of heterogeneity and finite lives for asset pricing. *Journal of Political Economy*, 123(3):670–685.

Gârleanu, N. and Pedersen, L. H. (2011). Margin-based asset pricing and deviations from the law of one price. *Review of Financial Studies*.

Gertler, M. and Kiyotaki, N. (2010). Financial Intermediation and Credit Policy in Business Cycle Analysis. *Handbook of Monetary Economics*, 3(C):547–599.

Gertler, M., Kiyotaki, N., and Prestipino, A. (2020). A Macroeconomic Model with Financial Panics. *The Review of Economic Studies*, 87(1):240–288.

Glorot, X. and Bengio, Y. (2010). Understanding the difficulty of training deep feedforward neural networks. In *Journal of Machine Learning Research*.

Gomez, M. (2019). Asset Prices and Wealth Inequality. *Working Paper*.

Gopalakrishna, G. (2020). A Macro-Finance model with Realistic Crisis Dynamics. *SSRN Electronic Journal*.

Gorton, G. and Ordoñez, G. (2014). Collateral Crises. *American Economic Review*, 104(2):343–78.

Gorton, G. and Ordoñez, G. (2020). Good Booms, Bad Booms. *Journal of the European Economic Association*, 18(2):618–665.

Gromb, D. and Vayanos, D. (2002). Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics*.

Gu, S., Kelly, B., and Xiu, D. (2020). Empirical Asset Pricing via Machine Learning. *The Review of Financial Studies*, 33(5):2223–2273.

Haddad, V. (2012). Concentrated Ownership and Equilibrium Asset Prices. *SSRN Electronic Journal*.

Han, J., Jentzen, A., and Weinan, E. (2018). Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences of the United States of America*.

Hansen, L. P., Khorrami, P., and Tourre, F. (2018). Comparative Valuation Dynamics in Models with Financing Restrictions. *Working Paper*.

He, Z., Kelly, B., and Manela, A. (2017). Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics*, 126(1):1–35.

He, Z. and Krishnamurthy, A. (2013). Intermediary asset pricing. *American Economic Review*, 103(2):732–770.

He, Z. and Krishnamurthy, A. (2019). A macroeconomic framework for quantifying systemic risk. *American Economic Journal: Macroeconomics*, 11(4):1–37.

Hornik, K. (1991). Approximation capabilities of multilayer feedforward networks. *Neural Networks*.

Hornik, K., Stinchcombe, M., and White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*.

Hughes, J. P., Mester, L. J., and Moon, C. G. (2001). 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.

Ikeda, D. and Kurozumi, T. (2019). Slow Post-financial Crisis Recovery and Monetary Policy. *American Economic Journal: Macroeconomics*, 11(4):82–112.

Jacot, A., Gabriel, F., and Hongler, C. (2018). Neural Tangent Kernel: Convergence and Generalization in Neural Networks. *Advances in Neural Information Processing Systems*, 2018-December:8571–8580.

Jaremski, M. and Sapci, A. (2017). Understanding the Cyclical Nature of Financial Intermediation Costs. *Southern Economic Journal*.

Jayaratne, J. and Strahan, P. E. (1998). Entry restrictions, industry evolution, and dynamic efficiency: Evidence from commercial banking. *Journal of Law and Economics*.

Judd, K. L. (1992). Projection methods for solving aggregate growth models. *Journal of Economic Theory*.

Kelly, S. E. (1996). Gibbs phenomenon for wavelets. *Applied and Computational Harmonic Analysis*.

Khorrami, P. (2016). Entry and Slow-Moving Capital: Using Asset Markets to Infer the Costs of Risk Concentration.

Kiyotaki, N. and Moore, J. (1997). Credit cycles. *Journal of Political Economy*.

Krishnamurthy, A. and Li, W. (2020). Dissecting Mechanisms of Financial Crises: Intermediation and Sentiment. *SSRN Electronic Journal*.

Krusell, P. and Smith, A. A. (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*.

Kurlat, P. (2018). How I Learned to Stop Worrying and Love Fire Sales. *National Bureau of Economic Research*.

Li, W. (2020). Public liquidity, bank runs, and financial crises. *SSRN Electronic Journal*.

Mallinson, G. D. and de Vahl Davis, G. (1973). The method of the false transient for the solution of coupled elliptic equations. *Journal of Computational Physics*.

Maxted, P. (2020). A Macro-Finance Model with Sentiment. *SSRN Working Paper*.

Mei, S., Misiakiewicz, T., and Montanari, A. (2021). Generalization error of random features and kernel methods: hypercontractivity and kernel matrix concentration. *Applied and Computational Harmonic Analysis*, 59:3–84.

Merkel, S. (2020). The Macro Implications of Narrow Banking: Financial Stability versus Growth. *Working Paper*.

Moll, B. (2014). Productivity losses from financial frictions: Can self-financing undo capital misallocation? *American Economic Review*, 104(10):3186–3221.

Moreira, A. and Savov, A. (2017). The Macroeconomics of Shadow Banking. *Journal of Finance*, 72(6):2381–2432.

Muir, T. (2017). Financial crises and risk premia. *Quarterly Journal of Economics*, 132(2):765–809.

Phelan, G. (2016). Financial Intermediation, Leverage, and Macroeconomic Instability. *American Economic Journal: Macroeconomics*, 8(4):199–224.

Phelan, T. and Eslami, K. (2022). Applications of Markov chain approximation methods to optimal control problems in economics. *Journal of Economic Dynamics and Control*, 143:104437.

Pino, G. and Sharma, S. C. (2019). On the Contagion Effect in the US Banking Sector. *Journal of Money, Credit and Banking*.

Raissi, M., Perdikaris, P., and Karniadakis, G. E. (2017). Physics informed deep learning (Part II): Data-driven discovery of nonlinear partial differential equations.

Raissi, M., Perdikaris, P., and Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*.

Reinhart, C. M. and Rogoff, K. S. (2009). *This time is different: Eight centuries of financial folly*.

Sarin, N. and Summers, L. H. (2016). Understanding bank risk through market measures. *Brookings Papers on Economic Activity*.

Sergeev, A. and Balso, M. D. (2018). Horovod: Fast and easy distributed deep learning in tensorflow.

Shiller, R. J. (1981). Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends? *American Economic Review*, 71(3):421–436.

Shleifer, A. and Vishny, R. (2011). Fire sales in finance and macroeconomics. *Journal of Economic Perspectives*, 25(1):29–48.

Silva, D. (2020). The Risk Channel of Unconventional Monetary Policy. *University of Illinois at Urbana-Champaign Working Paper*.

Sirignano, J. and Spiliopoulos, K. (2018a). DGM: A deep learning algorithm for solving partial differential equations. *Journal of Computational Physics*.

Sirignano, J. and Spiliopoulos, K. (2018b). DGM: A deep learning algorithm for solving partial differential equations. *Journal of Computational Physics*.

Tompson, J., Schlachter, K., Sprechmann, P., and Perlin, K. (2017). Accelerating eulerian fluid simulation with convolutional networks. In *34th International Conference on Machine Learning, ICML 2017*.

Wachter, J. A. (2013). Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility? *Journal of Finance*.

Wang, S., Yu, X., and Perdikaris, P. (2020). When and why PINNs fail to train: A neural tangent kernel perspective. *Journal of Computational Physics*, 449.

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

Yellen, J. (2017). Financial Stability a Decade after the Onset of the Crisis: a speech at \"Fostering a Dynamic Global Recovery,\" a symposium sponsored by the Federal Reserve Bank of Kansas City, Jackson Hole, Wyoming, August 25, 2017.

Goutham Gopalakrishna

goutham.gopalakrishna@epfl.ch ∘ Personal Website ∘ +1 609-436-8385 (or) +41 779243829

Office Contact Information

EPFL CDM SFI EXTRA 128 (Extranef UNIL),
Quartier UNIL-Dorigny
CH-1015, Lausanne
Switzerland

Graduate Studies

Swiss Finance Institute at École Polytechnique Fédérale de Lausanne (SFI-EPFL)
2017-present

PhD Candidate in Finance

Expected Completion Date: June 2023

References

Professor Pierre Collin-Dufresne (main advisor)
College of Management of Technology
École Polytechnique Fédérale de Lausanne (SFI)
+41 21 693 01 36
pierre.collin-dufresne@epfl.ch

Professor Markus Brunnermeier
Department of Economics
Princeton University
+1 609-258-4811
markus@princeton.edu

Professor Julien Hugonnier
College of Management of Technology
École Polytechnique Fédérale de Lausanne (SFI)
+41 21 693 01 14
julien.hugonnier@epfl.ch

Princeton University *Sep 2022 - Feb 2023*

Visiting Student Research Collaborator

Bendheim Center For Finance

Host: Professor Markus Brunnermeier

Fields

Primary Macro-Finance, Financial Intermediation

Secondary Machine Learning, Asset Pricing

Research Affiliations

CESifo Distinguished Affiliate, Munich, Germany *July 2021-Present*

Job Market Paper

"A Macro-Finance Model with Realistic Crisis Dynamics." 2022.

CESifo Distinguished Affiliate Award 2021
Runner up for European Systemic Risk Board 2021 Ieke van den Burg Prize

Financial recessions are typically characterized by a large risk premium and a slow recovery. However, macro-finance models have trouble quantitatively explaining these empirical features, especially when they are calibrated to simultaneously match both the observed unconditional and conditional macroeconomic and asset pricing moments. In this paper, I build a macro-finance model that quantitatively explains the salient features of a financial crisis, such as a large drop in output, a spike in the risk premium, reduced financial intermediation, and a long duration of economic distress. The model has leveraged intermediaries with stochastic productivity and a state-dependent exit rate that governs the transition into and out of a crisis. A model without these two features suffers from a trade-off between the amplification and persistence of crisis. I show that my model resolves this tension and generates realistic crisis dynamics.

Presentations (*in-person): Princeton University Finance Seminar Series* (2022), AFA poster (2022), SFI Job Market Workshop (2022), CESifo conference on Macro, Money, and International Finance (2021), RiskLab/BoF/ESRB Conference (2021), Paris December Meetings (2021), DGF German Finance Association Innsbruck* (2021), Econometric Society Meetings (2021; North America, Europe, Asia, Australia), AFFI PhD session (2021), AEFIN Ph.D. Mentoring Day (2021), Day-Ahead Workshop on Financial Regulation poster Zurich* (2021), Workshop on Macroeconomic Research Carcow (2021), Money Macro and Finance Society Conference (2021), Miami Winter Research Conference on Machine Learning and Business (2021), New Zeland Finance Conference (2021), SFI Gerzensee Research Days (2021), UNIL/EPFL Brown Bag (2020).

Working Papers

1. "ALIENs and Continuous Time Economies." May 2021.

   Presentations: Princeton University (2022), SFI-UZH Computational Finance seminar (2021), EUI Artificial Intelligence seminar (2021).

2. "Intermediaries with something to lose: On the origins, and consequences of bank failures." 2022. (Draft available upon request.)

   Presentations: 20th Macro Finance Society PhD session (2022), CESifo Conference on Macro, Money, and International Finance (2022), EPFL-UNIL PhD seminar (2022), SFI-UZH Computational Finance seminar (2022).

Work in Progress

1. "Capital Strucure Dynamics with Active Debt and Equity Management." With Julien Hugonnier and Erwan Morellec.
2. "Supply Chain Finance and Firm Capital Structure." With Claudio Tebaldi and Laura Bottazzi.
3. "Interest Rate Uncertainty and Public Debt Dynamics." With Andrea Modena.

Teaching

EPFL, Switzerland	Game Theory and Strategic Decisions (Undergraduate; TA Fall 2021)
	Optimization Methods (Graduate; TA Fall 2021)
	Financial Big Data (Graduate; TA Fall 2018-2020)
	Financial Applications of Blockchain (Graduate; TA Fall 2018-2020)
	Academic Supervisor for Executive MBA (Fall 2020-2021)
IFMR GSB, India	Computational Finance (MBA; Visiting Lecturer Fall 2021)
University of Bologna, Italy	Mathematics (Undergraduate; TA Fall 2016)
	Corporate Finance (Undergraduate; TA Fall 2016)
	Asset Pricing (Graduate; TA Spring 2017)
	Computational Tools (Undergraduate; TA Spring 2017)
	Mathematical Economics (Graduate; Spring 2017)

Honors, Scholarships, Fellowships, and Grants

CESifo Distinguished Affiliate Award, 2021, worth EUR 1,000	*2021*
Swiss Finance Institute PhD Fellowship worth CHF 30,000	*2017-2018*
University of Bologna Merit Scholarship worth EUR 22,000	*2015-2017*

Business Experience

Moody's Analytics Knowledge Services *2012-2015*
Quantitative Researcher

Hospira healthcare *2011-2012*
Quant Executive: Finance and Supply Chain

Prior Education

Anna University, Chennai, India *2009*
Bachelor of Engineering, Computer Science

University of Bologna, Italy *2017*
Master of Economics

Languages

Tamizh (Native), English (Fluent), French (Basic), Hindi (Conversational)

Last updated: October 2022