Pontifícia Universidade Católica
do Rio de Janeiro

João Batista Leal Filho

Understanding Financial and Non-Financial Balance Sheet Recessions

Dissertação de Mestrado

Master's dissertation presented to the Programa de Pós–graduação em Economia, do Departamento de Economia da PUC-Rio in partial fulfillment of the requirements for the degree of Mestre em Economia.

Advisor : Prof. Fernando Mendo
Co-advisor: Prof. Carlos Viana de Carvalho

Rio de Janeiro
September 2025

Pontifícia Universidade Católica
do Rio de Janeiro

João Batista Leal Filho

Understanding Financial and Non-Financial Balance Sheet Recessions

Master's dissertation presented to the Programa de Pós–graduação em Economia da PUC-Rio in partial fulfillment of the requirements for the degree of Mestre em Economia. Approved by the Examination Committee:

Prof. Fernando Mendo
Advisor
Departamento de Economia – PUC-Rio

Prof. Carlos Viana de Carvalho
Co-advisor
Departamento de Economia – PUC-Rio

Prof. Yvan Becard
Departamento de Economia – PUC-Rio

Prof. Alonso Villacorta
Department of Economics – UCSC

Rio de Janeiro, September the 8th, 2025

João Batista Leal Filho

BSc. in Economics, Insper, 2019

Bibliographic data

Leal Filho, João Batista

Understanding Financial and Non-Financial Balance Sheet Recessions / João Batista Leal Filho; advisor: Fernando Mendo; co-advisor: Carlos Viana de Carvalho. – 2025.

71 f: il. color. ; 30 cm

Dissertação (mestrado) - Pontifícia Universidade Católica do Rio de Janeiro, Departamento de Economia, 2025.

Inclui bibliografia

1. Economia – Teses. 2. Flutuações e Ciclos de Negócios. 3. Mercados Financeiros e a Macroeconomia. 4. Crises Financeiras. I. Mendo, Fernando. II. Carvalho, Carlos Viana de. III. Pontifícia Universidade Católica do Rio de Janeiro. Departamento de Economia. IV. Título.

CDD: 330

Acknowledgments

I thank God for His spiritual guidance and blessings throughout this journey. I also thank my parents and sisters for believing in me and supporting my dreams despite all hardships, for showing me how transformative education can be, and for always pushing me to do better.

I thank my advisors, Fernando Mendo and Carlos Carvalho, for their incredible support throughout these years and this project, for teaching me many invaluable lessons, and for inspiring me to pursue an academic career. I also thank the committee members, Alonso Villacorta and Yvan Becard, for their availability and valuable feedback that helped improve this work.

I thank the Economics Department at PUC-Rio for their support during the program. I am grateful to Professors Walter Novaes, Natalie Gimenes, and Leonardo Rezende for their kindness and rigor that helped shape my research interests. I also thank all my friends from PUC-Rio for making this journey more enjoyable.

This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001.

Abstract

Leal Filho, João Batista; Mendo, Fernando (Advisor); Carvalho, Carlos Viana de (Co-Advisor). **Understanding Financial and Non-Financial Balance Sheet Recessions**. Rio de Janeiro, 2025. 71p. Dissertação de Mestrado – Departamento de Economia, Pontifícia Universidade Católica do Rio de Janeiro.

This paper examines the dynamics underlying financial and non-financial balance sheet recessions. We develop a continuous-time model with three agents that captures the distinct mechanisms associated with each type of downturn. Our analysis shows that the effects of capital shocks to firms' balance sheets depend on the capitalization of financial intermediaries. When intermediaries are well capitalized, such shocks trigger fire sales of capital and heightened volatility, leading to severe downturns through the balance sheet channel. By contrast, when intermediaries are poorly capitalized, the same shocks generate deeper contractions in investment and output, as constrained banks reduce credit supply through the lending channel. These results are consistent with some empirical facts and underscores the role of intermediary capitalization in shaping macro-financial outcomes.

Keywords

Business Fluctuations and Cycles; Financial Markets and the Macroeconomy; Financial Crisis.

Resumo

Leal Filho, João Batista; Mendo, Fernando; Carvalho, Carlos Viana de. **Compreendendo Recessões Financeiras e Não Financeiras**. Rio de Janeiro, 2025. 71p. Dissertação de Mestrado – Departamento de Economia, Pontifícia Universidade Católica do Rio de Janeiro.

Este artigo analisa a dinâmica por trás de recessões financeiras e não financeiras. Desenvolvemos um modelo em tempo contínuo com três agentes que captura os distintos mecanismos associados a cada tipo de recessão. Nossa análise mostra que os efeitos de choques de capital nos balanços das firmas dependem da capitalização dos intermediários financeiros. Quando os intermediários estão bem capitalizados, tais choques desencadeiam vendas de capital e aumento da volatilidade, levando a recessões severas por meio do canal do balanço das firmas. Por outro lado, quando os intermediários estão mal capitalizados, os mesmos choques geram contrações mais profundas em investimento e em produção, à medida que bancos reduzem a oferta de crédito. Esses resultados são consistentes com evidências empíricas e destacam o papel central da capitalização dos intermediários nas dinâmicas macroeconômicas e financeiras.

Palavras-chave

Flutuações e Ciclos de Negócios; Mercados Financeiros e a Macroeconomia; Crises Financeiras.

Table of contents

List of figures

List of tables

1
Introduction

This paper examines how the distribution of wealth between financial and non-financial firms shapes the dynamic of different recession episodes. Specifically, we explore the distinctions between financial recessions[1], driven by undercapitalized financial intermediaries, and non-financial recessions, which originate from undercapitalized productive firms. By highlighting the mechanisms through which the capitalization of both financial and non-financial agents[2] affects financial variables and real economic outcomes, our study sheds more light on the forces underlying business cycle fluctuations.

The motivation for this study is illustrated in Figure 1.1, reproduced from Baron et al. (2021). In the left panel, a 30% decline in the equity[3] returns (an equity crash) of financial firms leads to GDP growth persistently below the "no crash" baseline. By contrast, in the right panel, GDP growth eventually returns to the baseline following a comparable equity crash in the non-financial sector. Table 1.1 shows that both financial and non-financial equity crashes are rare events, each with an implied probability of occurrence below 5.5%. However, the likelihood of a recession within two years of a crash in the financial sector is substantially higher than following a non-financial crash or in the absence of a crash.

Figure 1.1: GDP Growth Around Equity Crashes

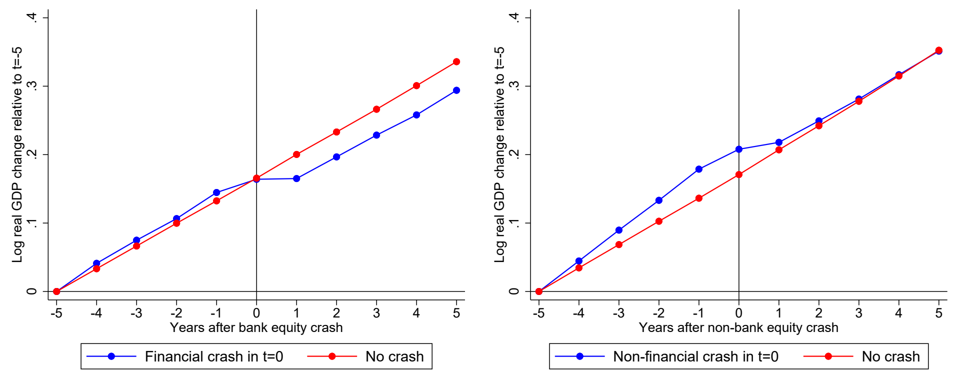

Note: Source: Baron et al. (2021). Equity crash is an event characterized by a drop of at least 30% in the equity value of financial and non-financial firms.

In this paper, we further examine the empirical relationship between financial and non-financial firms capitalization and macroeconomic aggregates.

[1]In this paper, recessions are characterized as events with a negative GDP growth.
[2]We use the terms financial firms, intermediaries, and banks interchangeably. Similarly, we refer to non-financial firms as productive firms/sector or simply firms.
[3]We use the terms equity, wealth, and net worth interchangeably throughout this paper.

We estimate a local projection model with fixed effects, incorporating the equity returns of financial and non-financial firms with an interaction term to capture potential nonlinearities in the transmission of equity crashes to the real economy. Although this approach does not allow for a causal effect, it provides insights into the average behavior of credit-to-GDP, investment-to-GDP, and GDP growth.

Table 1.1: Statistics on Equity Crashes

	Financial Crash	Non-financial Crash	No Crash
Event	95	82	1,593
$P(E)$	5.5%	2.9%	92.6%
Recession|Event	32	12	207
$P(R\|E)$	33.7%	24.0%	13.0%
Financial Crisis|Event	34	18	34
$P(FC\|E)$	35.8%	36.0%	2.1%
N	1720	1720	1720

Note: $P(E)$ = Event/N, $P(R|E)$ = Recession/Event, and $P(FC|E)$ = Financial Crisis/Event, where $P(\cdot)$ is probability. Recession is the number of negative GDP growth in the same and/or in the period following the event. The same is valid for financial crisis. E, R and FC are the number of events, recessions and financial crisis, respectively. Data sources: Jordà et al. (2016) and Baron et al. (2021).

Consistent with the literature, our results show that a 30% equity crash in the productive sector is associated with a sharp but short-lived decline in investment-to-GDP and GDP growth, with no significant change in credit-to-GDP. Firms' ability to continue accessing the credit market support a faster recovery of the economy. We extend this analysis by examining a scenario in which firms and intermediaries simultaneously experience a 30% equity crash. In this case, the shock is associated with a significantly deeper and more persistent downturn: credit-to-GDP, investment-to-GDP, and GDP growth all remain below their "no-crash" paths for an extended period. The severity and persistence of these episodes, linked to intermediary capitalization and constrained credit supply, highlight an amplification mechanism beyond the traditional balance sheet channel described in Bernanke et al. (1999).

Although these empirical findings are suggestive, they do not uncover the mechanisms through which equity crashes propagate to the real economy. We partially address this gap with a continuous-time model in which the net worth of productive firms and financial intermediaries have distinct roles in the dynamics of the economy. Unlike most models in the literature, which often treat the wealth of firms and intermediaries as perfect substitutes[4], our

[4]See, for example, He and Krishnamurthy (2013) and Brunnermeier and Sannikov (2014), where firms and intermediaries are modeled as a single agent ("experts").

framework distinguishes between the two by assigning different productivity levels and discount rates.

Our model comprises three sectors: households, financial intermediaries, and productive firms. All agents hold capital and produce, but firms are specialized and are the most productive agents in the economy. Financial intermediaries can leverage their investments by raising deposits from households, while firms finance their capital acquisitions by issuing bonds to intermediaries. The model also incorporates several financial frictions. First, there is no equity market, implying that all agents must hold all the asset risk in their balance sheets. Second, households have limited market participation, meaning that they do not lend directly to firms, only through intermediaries. Third, intermediaries face a collateral constraint when borrowing from households, making their wealth a key determinant of credit supply to firms.

These frictions give rise to three occasionally binding constraints: a short-sale constraint for households, a short-sale constraint for intermediaries, and a collateral constraint that limits intermediaries' borrowing from households. Combined with two state variables, the wealth shares of firms and intermediaries, these constraints give rise to five attainable equilibrium regions. Each region corresponds to a different macro-financial regime, with recessions classified as either financial or non-financial depending on the capitalization of firms and intermediaries. In this economy, the steady state emerges when the short-sale constraints bind, that is, when firms hold the entire capital stock, but the collateral constraint remains slack.

To the best of our knowledge, this is the first continuous-time model to jointly incorporate three agents, two state variables, and multiple occasionally binding constraints in order to study the macroeconomic implications of shifts in wealth distribution arising from an aggregate capital shock. Solving such a model, however, has significant numerical challenges. To address this, we adapt the probabilistic approach introduced by Huang (2023) to solve high-dimensional continuous-time models with multiple occasionally binding constraints. We demonstrate that this method is capable of handling sharp transitions across equilibrium regions and is well suited to the challenges posed by our framework.

When working with two state variables, ensuring consistency, stability, and monotonicity in the resulting system of partial differential equations (PDEs), as required by standard finite-difference (FD) methods, becomes particularly challenging. These methods frequently fail to converge, especially when multiple occasionally binding constraints segment the state space into different equilibrium regions. In this context, neural networks present a promising alternative. Recent advances in solving PDEs using reinforcement learning have expanded the range of models that are computationally tractable. However, these approaches still face relevant limitations. Like FD methods, they rely on smooth approximations of partial derivatives, which makes them poorly equipped to handle models featuring multiple discontinuities or "kinks" in policy functions, features that naturally arise in environments with financial frictions and sharp transitions across equilibrium regions.

To overcome these challenges, we adapt the probabilistic method to accommodate models with multiple occasionally binding constraints. Rather than solving the system of PDEs directly, this approach reformulates the problem as a system of forward-backward stochastic differential equations (FBSDEs), which is then solved using reinforcement learning. This method offers two primary advantages over traditional analytical (PDE-based) approaches. First, while analytical methods require evaluating the solution over a dense grid and along multiple derivative directions, the probabilistic approach relies on a limited set of conditions derived from simulated Brownian motion paths, substantially reducing both algorithmic complexity and computational burden. Second, unlike analytical methods, the probabilistic framework does not require policy functions to be differentiable across their entire domain, making it particularly well suited for models such as ours that feature sharp transitions between equilibrium regions.

The solution of our model highlights dynamics that distinguish financial from non-financial recessions. By modeling a single aggregate capital shock, affecting both firms and intermediaries, we only capture the economy's response to conditional shocks, i.e., shocks occurring in predetermined regions of the state space, rather than the direct impact of pure equity crashes, as in the empirical analysis. In this setting, when the financial sector is well capitalized, a sequence of capital shocks primarily affecting firms' balance sheets leads the economy into a non-financial recession via the balance sheet channel. In this regime, intermediaries continue to supply credit, as their collateral constraint remains slack. Nevertheless, firms lose a critical source of funding, their own equity, as leverage amplifies the effects of the negative capital shock. As a result, firms are forced to liquidate capital at depressed prices to less productive agents (intermediaries and households), triggering fire sales that increase volatility and deepen the downturn.

Capital shocks in the presence of poorly capitalized intermediaries (financial recessions) are associated with different dynamics. Low-wealth intermediaries face a binding collateral constraint, making them unable to raise deposits from households and limiting the supply of credit to firms. With low wealth, intermediaries also charge a positive spread on firms' bonds, further tightening the financial constraint faced by productive agents. As a result, firms are forced to scale back investment, reduce output, and liquidate capital to less productive agents. A sequence of negative capital shocks in this regime further deteriorates firms' balance sheets, delaying the recovery of firms' net worth, as lower asset holdings and elevated funding costs reduce returns and slow wealth accumulation.

The primary driver behind this dynamic is the lending channel, which links intermediaries' balance sheets to credit supply and the economy. While this mechanism deepens downturns, it simultaneously dampens the balance sheet channel by constraining leverage for both firms and intermediaries, which reduces volatility. Paradoxically, this interaction creates a trade-off between output and financial stability. Poorly capitalized intermediaries operating without a binding collateral constraint are associated with intense fire sales and elevated volatility. In contrast, when the collateral constraint binds, volatility declines because the constraint limits firms' ability to leverage and take risks, thereby reducing the scale of fire sales following capital shocks.

In this setting, the binding constraint promotes a more stable environment by mitigating the balance sheet channel. However, this increased stability comes at a cost: tighter credit conditions lead to more severe downturns. Although we do not undertake a formal welfare analysis, these findings suggest that policies aimed at enhancing financial stability may entail significant economic costs compared with policies focused on strengthening the capitalization of financial intermediaries.

Literature. A large body of research highlights the role of firms' balance sheets in shaping economic and financial dynamics. An early generation of studies examines small, one-time shocks to firms' wealth using linearized models around the steady state[5]. These studies document amplification effects associated with firms' net worth. However, their reliance on local approximations limits their ability to capture the nonlinear dynamics that arise following large shocks that push the economy far from the steady state. A more recent class of continuous-time models has sought to capture global dynamics[6]. Yet, by treating firms and intermediaries as perfectly substitutes, these frameworks largely overlook the role of wealth distribution across firms, intermediaries, and households in shaping macro financial outcomes.

[5]See, for example, Kiyotaki and Moore (1997), Bernanke et al. (1999), Gertler and Kiyotaki (2010), Mendoza (2010), Jermann and Quadrini (2012), Iacoviello (2015), and Guerrieri and Iacoviello (2017).

[6]See, for example, He and Krishnamurthy (2012), He and Krishnamurthy (2013), He and Krishnamurthy (2019), Brunnermeier et al. (2013), Brunnermeier and Sannikov (2014), and Krishnamurthy and Li (2025).

Another line of research highlights the role of financial intermediaries' balance sheets in transmitting shocks to the economy. Empirically, studies such as Gilchrist and Zakrajsek (2012), Schularick and Taylor (2012), Berger and Bouwman (2013), Adrien et al. (2014), Jordà et al. (2016), Muir (2017), Jordà et al. (2018), and Krishnamurthy and Muir (2025) document how fluctuations in credit supply shape the severity of financial crises. In particular, these works find that adverse shocks to intermediaries' net worth raise both the equity risk premium and credit spreads. This reflects not only a deterioration in the fundamentals of firms' assets but also a weakening of intermediaries' balance sheets, which increases the discount rate and raises the marginal cost of funds. A large theoretical literature also emphasizes the central role of intermediaries, particularly during periods of elevated financial stress, including Diamond (1984), Diamond and Rajan (2001), Kashyap et al. (2002), Allen and Gale (2004), Mehran and Thakor (2011), Hu and Varas (2025), and Bolton et al. (2025).

Only a few studies examine how the balance sheets of both financial and non-financial firms jointly influence economic and financial outcomes. Empirically, Baron et al. (2021) show that large declines in bank equity returns are associated with credit contractions and deep and persistent economic downturns. In contrast, drops in non-bank equity returns lead to much smaller contractions in GDP growth and do not significantly affect the credit-to-GDP ratio. Ottonello and Song (2024) analyze the impact of financial firms' earnings announcements on the stock prices of non-financial firms and on aggregate economic activity. They find that negative shocks to bank equity affect non-financial firms' balance sheets and production primarily through the credit supply channel, which is closely tied to the financial health of intermediaries. In this paper, we extend these findings by showing how capital shocks that primarily affect productive firms propagate to the real economy, and how the severity of these shocks depends on the capitalization of the financial sector.

Other papers propose models to study how the balance sheets of firms and intermediaries affect economy's dynamics. Holmstrom and Tirole (1997) develop a static model with a dual moral hazard problem involving firms and intermediaries. In their framework, the borrowing capacity of both agents is constrained by their net worth, which determines the effectiveness of intermediary monitoring. However, their model treats agents' wealth as exogenous and does not capture asset price amplification effects on the real economy. Meh and Moran (2010) embed the Holmstrom and Tirole (1997) framework into a New Keynesian model and show that a negative technology shock reduces bank profitability, impeding their ability to raise funds from households. As a result, banks must finance firms with their own net worth, which is now diminished, allowing the technology shock to propagate through the economy.

Boissay et al. (2016) develop a model featuring financial frictions between banks and firms, as well as within the interbank market. In their framework, credit crunches[7] can arise from a misallocation of credit between more and less efficient banks, triggering severe downturns. However, their model does not emphasize the role of intermediary capitalization or asset prices. Rampini and Viswanathan (2018) is more closely related to our approach, as it highlights the importance of both firms' and intermediaries' wealth in shaping macroeconomic dynamics. In their model, banks exist due to a collateralization advantage over households, which enables them to issue additional credit. When intermediary wealth declines, the downturn is both severe and persistent, as banks rebuild their balance sheets only gradually, results that closely align with our findings. We depart from their setup by assuming that households cannot lend directly to firms; instead, all credit must be intermediated by banks. Furthermore, our continuous-time framework allows us to capture global dynamics and amplification effects driven by leverage and asset prices.

Elenev et al. (2021) develop a model with financially constrained firms and intermediaries, showing that highly leveraged banks can amplify small credit losses from shocks to firms into full-blown financial crises. Villacorta (2023) assigns a central role to intermediary wealth within a real business cycle framework. In his model, banks face two collateral constraints, when borrowing from households and when lending to firms, with both depending on their monitoring capacity. When firms are poorly capitalized, the economy exhibits the standard financial accelerator mechanism, whereas undercapitalized banks give rise to the bank accelerator mechanism, which is responsible for the greater severity of financial crises.

A major constraint in solving models with more than two state variables and multiple financial frictions lies in the limitations of traditional numerical methods. For instance, d'Avernas et al. (2023) introduce a generalized version of the finite-difference method for two-dimensional problems. While this approach extends the class of models that can be addressed, it does not guarantee convergence, as maintaining consistency, monotonicity, and stability in a system of partial differential equations (PDEs) remains challenging.

Gopalakrishna (2023), Gopalakrishna and Yuntao (2024), and Wu et al. (2024) propose methods that combine deep neural networks with reinforcement learning to solve high-dimensional continuous-time models featuring occasionally binding constraints. These approaches approximate the value function

[7]Financial recessions, in our terminology.

through the universal approximation theorem[8], bypassing the strict requirements imposed by finite-difference methods. In addition to their flexibility in approximating complex policy functions, they scale more efficiently to models with more than two endogenous state variables, making them computationally viable.[9] Similar strategies have been adopted in continuous-time settings by Gao and Wang (2023), Duarte et al. (2024), and Fan et al. (2024).

The probabilistic approach of Huang (2023), adopted here, takes a different route. Rather than solving a system of PDEs directly, the method reformulates the problem as a system of forward-backward stochastic differential equations (FBSDEs) and solves it via reinforcement learning. While the analytical approach requires evaluating the value function at many grid points, a task that becomes increasingly complex with more state variables, the probabilistic method relies only on two conditions derived from simulated Brownian motion paths, which substantially reduces the computational burden.

[8]See Hornik (1991).

[9]For a broader overview of neural network applications in macro-finance, see Fernandez-Villaverde et al. (2024).

2
Empirical Facts

In this section, we document a set of empirical facts regarding the relationship between balance sheets and macroeconomic dynamics. The analysis builds on data from Jordà et al. (2016), which provides measures of investment-to-GDP[1] and financial crisis events, and from Baron et al. (2021), which offers annual data on GDP growth, credit-to-GDP, and equity returns for both financial and non-financial sectors.[2] The combined dataset spans 17 countries over the period 1870 to 2016, yielding a panel of 1,720 country-year observations.

We estimate an empirical model that employs a *continuous* measure of equity returns and includes an interaction term between financial and non-financial equity returns. This specification enables us to capture potential nonlinearities and to distinguish the dynamics of financial versus non-financial recessions. While our analysis does not aim to establish a causal relationship between equity returns and macroeconomic outcomes, it allows us to observe the average response of the economy to changes in net worth. In doing so, we highlight potential channels through which the wealth of financial and non-financial firms may influence macroeconomic fluctuations.

Our empirical approach follows the local projection method from Jordà (2005), and is implemented using the following fixed-effects specification

$$\Delta y_{i,t+h} = \alpha_i + \beta_t^h r_{i,t}^I + \gamma_t^h r_{i,t}^F + \theta_t^h r_{i,t}^I r_{i,t}^F + \Gamma^h X_{i,t} + \varepsilon_{i,t} \tag{2-1}$$

where $\Delta y_{i,t+h}$ denotes the cumulative change in the outcome variable from period t to $t+h$ for country i; $h \in \{0, 1, ..., 6\}$; α_i is a country fixed effect; $r_{i,t}^I$ and $r_{i,t}^F$ denote equity returns for financial and non-financial firms, respectively; and $X_{i,t}$ is a vector of control variables[3]. The interaction term $r_{i,t}^I r_{i,t}^F$ captures potential nonlinearities in the joint response of the economy to changes in equity returns across sectors[4].

The coefficients β_t^h, γ_t^h, and θ_t^h trace the impulse responses over horizon h. Specifically, β_t^h and γ_t^h capture the effects of balance sheet shocks originating

[1] Investment-to-GDP refers to investment by the productive sector (firms) as a share of total output.

[2] Financial equity refers to the market capitalization of the banking sector, while non-financial equity refers to the market capitalization of the productive sector.

[3] We include three-year lags of investment-to-GDP, credit-to-GDP, and GDP growth as control variables.

[4] Year fixed effects are excluded from the baseline specification to allow for cross-country variation in time effects, but are included in robustness checks presented in Tables A1, A2, and A3 in the Appendix.

in the financial and non-financial sectors, respectively. Our analysis extends this framework by incorporating an interaction term, θ_t^h, which captures cross-sectoral amplification effects.

Empirical Fact 1. Equity crashes in non-financial firms are associated with sharp contractions in GDP growth and investment-to-GDP, while credit-to-GDP shows no significant response.

Consistent with Baron et al. (2021), Figure 2.1 shows that a 30% drop in the equity return of non-financial firms is associated with a significant short-run decline in both GDP growth and the investment-to-GDP ratio, but has no clear effect on overall leverage measured by credit-to-GDP. This pattern suggests that, despite a substantial loss of internal funding and a deep downturn, non-financial firms continue to access credit through financial intermediaries, which potentially supports a faster recovery after the crash.

Figure 2.1: Response to a 30% Equity Crash in the Non-Financial Sector

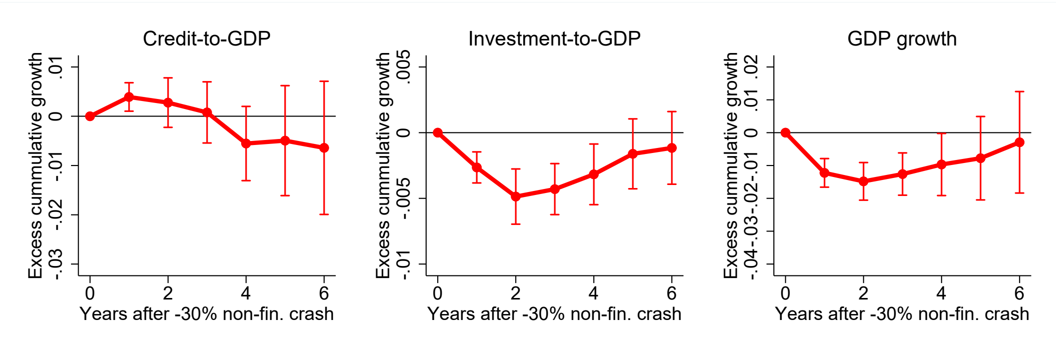

Note: The solid line displays the estimated response to a 30% equity crash in the non-financial sector, conditional on a 0% return in the equity of financial firms.

Empirical Fact 2. Equity crashes in financial firms are associated with sharp and prolonged contractions in credit-to-GDP, investment-to-GDP, and GDP growth.

Also consistent with Baron et al. (2021), Figure 2.2 shows that a 30% drop in the equity returns of financial firms is associated with a severe and persistent contraction in credit-to-GDP, investment-to-GDP, and GDP growth. This pattern suggests that fragility within the financial sector is tied to disruptions in credit markets, which constrain firms' funding capacity and investment, potentially leading to a slower and more prolonged economic recovery.

Empirical Fact 3. Equity crashes in both financial and non-financial firms—i.e., joint crashes—are associated with even more severe and persistent contractions in credit-to-GDP, investment-to-GDP, and GDP growth.

Figure 2.3 presents the response to a 30% equity crash in both financial and non-financial sectors. Since the parameter θ_t^h is negative across all vari-

Figure 2.2: Response to a 30% Equity Crash in the Financial Sector

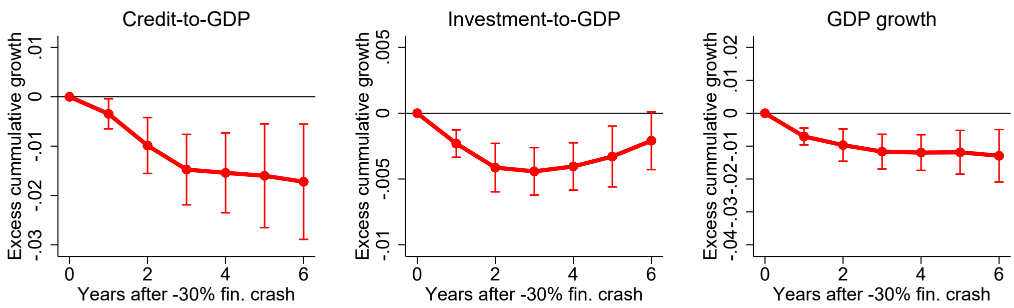

Note: The solid line displays the estimated response to a 30% equity crash in the financial sector, conditional on a 0% return in the equity of non-financial firms.

ables, joint equity crashes are associated with an amplification effect[5]. Consequently, the contractions in credit-to-GDP, investment-to-GDP, and GDP growth are more severe and persistent than those observed in the individual crashes described in the previous two facts (dashed lines). This interdependence underscores the critical role of wealth distribution between firms and intermediaries in the severity of recessions.

Figure 2.3: Response to a Joint 30% Equity Crash

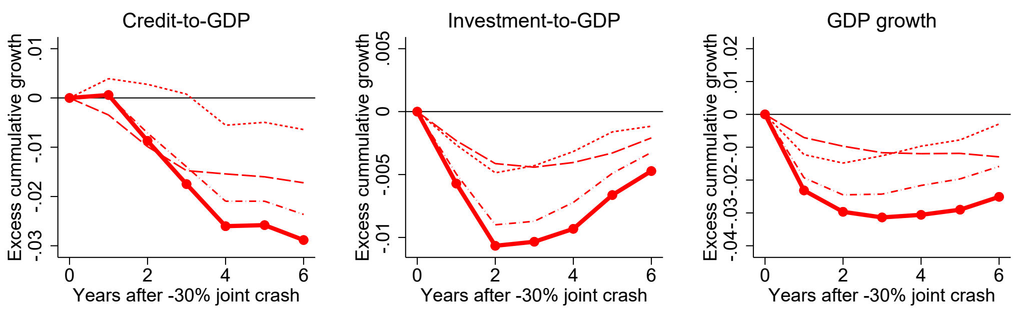

Note: The solid line shows the estimated response to a 30% equity crash in both the financial and non-financial sectors. The long-dashed line corresponds to the response shown in Figure 2.2, while the short-dashed line corresponds to the response in Figure 2.1. The dash-dotted line corresponds to the sum of the long-dashed and short-dashed lines.

Empirical Fact 4. Credit spreads rise sharply following equity crashes in financial firms.

Gilchrist and Zakrajsek (2012) document that adverse shocks to the equity valuations of financial intermediaries reduce their risk-bearing capacity, thereby constraining credit supply, widening the excess bond premium (credit spread), and amplifying contractions in investment and output as firms' access to external funding becomes limited. Baron et al. (2021) further show that

[5]See regression results in Tables A1, A2, and A3 in the Appendix.

corporate and interbank lending spreads rose sharply during episodes of large equity losses in the financial sector between 1870 and 2016.

To capture the mechanisms underlying some of these empirical facts linked to financial and non-financial recessions, we propose a model that explicitly incorporates financial intermediaries, non-financial firms, and credit frictions, jointly determining asset prices, risk premia, and aggregate economic outcomes. The model draws on elements from Kiyotaki and Moore (1997) and Brunnermeier and Sannikov (2014), combining intermediary-based financial frictions with macroeconomic dynamics to analyze how capital shocks interact with the balance sheet of firms and intermediaries, and propagate through asset markets and the broader economy.

3
Model

3.1
Basic Setup

Time is continuous and infinite, and the economy is populated by a continuum of three types of agents: households $h \in \mathbb{H} = [0,1]$, non-financial firms $f \in \mathbb{F} = (1,2]$, and financial firms $i \in \mathbb{I} = (2,3]$. In this model, there is no equity market—only a credit market in which non-financial firms borrow from financial intermediaries, and intermediaries borrow from households. Households are subject to limited market participation and cannot lend directly to non-financial firms[1]. Within this framework, intermediaries play an important role in channeling funds from households to productive firms.

Households consume, lend to intermediaries, and produce, though with lower productivity than non-financial firms. Non-financial firms consume, produce with the highest level of productivity, and finance their operations by borrowing from intermediaries. Financial intermediaries consume, produce at an intermediate level of productivity, lower than that of non-financial firms, and extend credit to firms by raising short-term deposits from households or by using their own net worth. In this model, bank capital influences both the interest rate on loans and the volume of credit available to firms. Figure 3.1 illustrates the balance sheets of the three sectors[2].

Figure 3.1: Balance Sheet Relationship

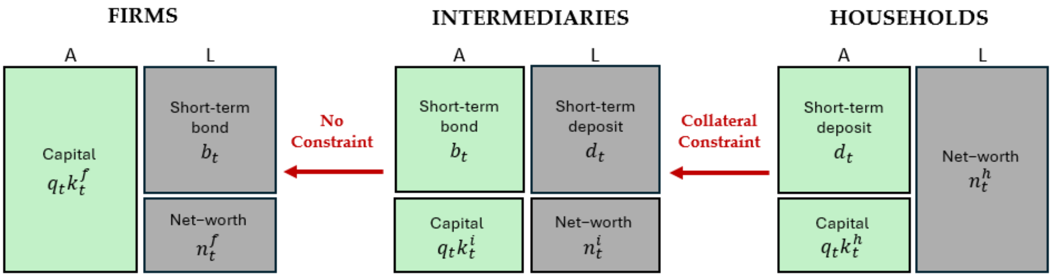

The model incorporates several financial frictions. First, firms are unable to issue equity to external investors, preventing them to sell part of the investment risk in their balance sheets. Second, households face participation

[1]This assumption is inspired by Basak and Cuoco (1998), who develop a framework in which investors face informational frictions and limited participation in equity markets. We adapt this idea by assuming that households also face limited participation in credit markets.

[2]We assume that all agents hold capital and produce to capture the dynamics of fire sales and how each agent's capitalization affects asset prices.

constraints and can only extend funds to firms indirectly via intermediaries. Third, intermediaries are subject to a collateral constraint when borrowing from households, making their net worth a key determinant of their lending capacity. These frictions collectively underscore the role of firms and intermediary balance sheets in shaping economic dynamics.

Production technology. Firms, intermediaries, and households produce the same good in this economy but differ in productivity. Firms are specialized in production and are therefore assumed to be more productive than intermediaries, who in turn are more productive than households. We assume that all agents use the same constant-returns-to-scale production function

$$y_{j,t} = a_j k_{j,t} \tag{3-1}$$

where $j \in \{f, i, h\}$, $y_{j,t}$ denotes total output at time t by agent j, a_j is the productivity parameter satisfying $a_f > a_i > a_h > 0$, and $k_{j,t}$ is the capital input.

Preferences. All agents in the economy are risk-averse and have logarithmic utility functions. Households are assumed to be more patient than intermediaries, who are, in turn, more patient than firms. This difference in time discount rates is important to prevent firms from escaping their financial constraints. The discounted lifetime utility for each agent is then given by

$$\mathbb{E}_0 \left[\int_0^\infty e^{-\rho_j t} \log(c_{j,t}) dt \right] \tag{3-2}$$

where ρ_j denotes the time discount rate satisfying $\rho_f > \rho_i > \rho_h > 0$, and $c_{j,t}$ is the consumption level.

Capital law of motion. Firms, intermediaries, and households hold capital for production and trade with each other. We assume that capital evolves according to a geometric Brownian motion (GBM), subject to an aggregate risk dZ_t that cannot be diversified[3].

$$\frac{dk_{j,t}}{k_{j,t}} = [\Phi(\iota_{j,t}) - \delta]\, dt + \sigma dZ_t \tag{3-3}$$

where $\iota_{j,t}$ is the investment rate per unit of capital, δ is the depreciation rate, and σ is the exogenous volatility. The function $\Phi(\cdot)$ captures capital adjustment cost and satisfies $\Phi(0) = 0$, $\Phi(1) = 1$, $\Phi' > 0$, and $\Phi'' < 0$.

Price of capital. All units of capital are homogeneous and traded at a

[3]Since firms are the most productive agents in the economy, they hold the majority of capital when they are not poorly capitalized and the credit market is not severely impaired. As a result, Brownian shocks impact productive agents more heavily than intermediaries and households.

unique price q_t. We assume that the process for q_t follows a GBM, subject to the same aggregate risk dZ_t

$$\frac{dq_t}{q_t} = \mu_t^q dt + \sigma_t^q dZ_t \tag{3-4}$$

where μ_t^q is the drift and σ_t^q is the volatility loading of the capital price. In this model, both μ_t^q and σ_t^q are determined endogenously.

Return on capital. The return on capital comprises two components: the dividend yield from production and the capital gains from asset price changes.

$$dr_{j,t}^k = \left(\frac{a_j - \iota_{j,t}}{q_t}\right) dt + \frac{d(q_t k_{j,t})}{q_t k_{j,t}}$$

The full expression is given by

$$dr_{j,t}^k = \underbrace{\left[\frac{a_j - \iota_{j,t}}{q_t} + \mu_t^q + \Phi(\iota_{j,t}) - \delta + \sigma\sigma_t^q\right]}_{\mathbb{E}[dr_{j,t}^k]/dt = \mu_{j,t}^k} dt + \underbrace{(\sigma + \sigma_t^q)}_{\sigma_t^k} dZ_t \tag{3-5}$$

Credit friction. Intermediaries can raise funds from households through risk-free, short-term deposits. However, the volume of deposits they can obtain is constrained by the size of their asset holdings[4]. This limitation is captured by the following collateral constraint

$$d_{i,t} \leq \nu^i (b_{i,t} + q_t k_{i,t}) \tag{3-6}$$

where $d_{i,t}$ denotes deposits from households to intermediary i, $b_{i,t}$ represents the intermediary's lending to firms (bonds), and $q_t k_{i,t}$ is the market value of its capital holdings. The parameter $\nu^i \in (0, 1)$ reflects the fraction of the intermediary's assets that can be pledged as collateral.

With the intermediary's balance sheet identity, $b_{i,t} + q_t k_{i,t} = d_{i,t} + n_{i,t}$, we can rewrite the collateral constraint (3-6) as

$$b_{i,t} + q_t k_{i,t} \leq \frac{1}{1-\nu^i} n_{i,t} \quad \Longrightarrow \quad \theta_{i,t}^b + \theta_{i,t}^k \leq \frac{1}{1-\nu^i},$$

where $\theta_{i,t}^b$ and $\theta_{i,t}^k$ denote the shares of intermediary i's wealth invested in firm's bonds and capital, respectively.

Household problem. Households choose how much to consume $c_{h,t}$, to invest $\iota_{h,t}$, to buy in capital $k_{h,t}$, and to lend to intermediaries $d_{h,t}$. They solve

[4]This constraint can be interpreted as a regulatory requirement, whereby governments mandate that intermediaries back a portion of their deposits with assets.

the following problem

$$
\begin{aligned}
\max_{\left\{c_{h,t},\theta^k_{h,t},\iota_{h,t}\right\}_{t\geq 0}} \quad & \mathbb{E}_0\left[\int_0^\infty e^{-\rho_h t}\log\left(c_{h,t}\right)dt\right] \\
\text{s.t.} \quad & \frac{dn_{h,t}}{n_{h,t}} = -\frac{c_{h,t}}{n_{h,t}}dt + \theta^k_{h,t}dr^{h,k}_t + \theta^d_{h,t}r^d_t dt \\
& \theta^k_{h,t} \geq 0 \\
& n_{h,0} = q_0 k_{h,0}; \quad n_{h,t} > 0; \quad \theta^d_{h,t} = 1 - \theta^k_{h,t}; \quad \forall h, t
\end{aligned}
$$

where $n_{h,t}$ denotes the household's net worth, $\theta^k_{h,t}$ is the share of net worth invested in capital, $dr^{h,k}_t$ is the return on capital given by equation (3-5), $\theta^d_{h,t}$ is the share invested in deposits to intermediaries, and r^d_t is the interest rates on deposits. The constraint $\theta^k_{h,t} \geq 0$ ensures that households cannot short-sale capital.

Intermediary problem. Intermediaries choose how much to consume $c_{i,t}$, to invest $\iota_{i,t}$, to buy in capital $k_{i,t}$, to lend to firms $b_{i,t}$, and to borrow from households, $d_{i,t}$. They solve the following problem

$$
\begin{aligned}
\max_{\left\{c_{i,t},\theta^k_{i,t},\theta^b_{i,t},\iota_{i,t}\right\}_{t\geq 0}} \quad & \mathbb{E}_0\left[\int_0^\infty e^{-\rho_i t}\log\left(c_{i,t}\right)dt\right] \\
\text{s.t.} \quad & \frac{dn_{i,t}}{n_{i,t}} = -\frac{c_{i,t}}{n_{i,t}}dt + \theta^k_{i,t}dr^{i,k}_t + \theta^b_{i,t}r^b_t dt - \theta^d_{i,t}r^d_t dt \\
& 0 \leq 1 - (1-\nu^i)(\theta^b_{i,t} + \theta^k_{i,t}) \\
& \theta^k_{i,t} \geq 0 \\
& n_{i,0} = q_0 k_{i,0}; \quad n_{i,t} > 0; \quad \theta^d_{i,t} = \theta^k_{i,t} + \theta^b_{i,t} - 1; \quad \forall i, t
\end{aligned}
$$

where $n_{i,t}$ denotes the intermediary's net worth, $\theta^b_{i,t}$ and $\theta^d_{i,t}$ represent the shares of net worth allocated to firm loans and deposits, respectively, and r^b_t is the interest rate on bonds from firms.

Firm problem. Firms choose how much to consume $c_{f,t}$, to invest $\iota_{f,t}$, to buy in capital $k_{f,t}$, and to borrow from intermediaries $b_{f,t}$. They solve the following problem

$$
\begin{aligned}
\max_{\left\{c_{f,t},\theta^k_{f,t},\iota_{f,t}\right\}_{t\geq 0}} \quad & \mathbb{E}_0\left[\int_0^\infty e^{-\rho_f t}\log\left(c_{f,t}\right)dt\right] \\
\text{s.t.} \quad & \frac{dn_{f,t}}{n_{f,t}} = -\frac{c_{f,t}}{n_{f,t}}dt + \theta^k_{f,t}dr^{f,k}_t - \theta^b_{f,t}r^b_t dt \\
& n_{f,0} = q_0 k_{f,0}; \quad n_{f,t} > 0; \quad \theta^b_{f,t} = \theta^k_{f,t} - 1; \quad \forall f, t
\end{aligned}
$$

where $n_{f,t}$ denotes the firm's net worth, $\theta^k_{f,t}$ the share of net worth allocated to capital, and $\theta^b_{f,t}$ the share financed through borrowing from intermediaries.

Optimal investment. Since the investment decision is static and time-separable, and assuming a standard logarithmic investment cost function, the optimal investment is determined by solving

$$\max_{\iota_{j,t}} \mathbb{E}[dr_t^{j,k}] = \max_{\iota_{j,t}} \left[\frac{a_j - \iota_{j,t}}{q_t} + \mu_t^q + \frac{1}{\phi} \log\left(1 + \phi\iota_{j,t}\right) - \delta + \sigma\sigma_t^q \right] dt$$

where the parameter ϕ is the elasticity of the investment technology. The solution to this problem yields

$$\iota_{j,t} = \frac{1}{\phi}(q_t - 1), \quad \forall j \Longrightarrow \iota_t = \frac{1}{\phi}(q_t - 1) \tag{3-7}$$

and the investment rate is identical across all agents in the economy.

Equilibrium definition. For a given initial endowment of capital $k_{f,0}, \forall f \in \mathbb{F}$, $k_{i,0}, \forall i \in \mathbb{I}$, and $k_{h,0}, \forall h \in \mathbb{H}$ such that $\int_{\mathbb{F}} k_{f,0}df + \int_{\mathbb{I}} k_{i,0}di + \int_{\mathbb{H}} k_{h,0}dh = k_0$, a competitive equilibrium is characterized by stochastic processes adapted to the filtration generated by Z_t for capital price q_t, deposit interest rate r_t^d, bond interest rate r_t^b, investment rates $(\iota_{f,t}, \iota_{i,t}, \iota_{h,t})$, capital allocations $(k_{f,t}, k_{i,t}, k_{h,t})$, loans $(b_{f,t}, b_{i,t}, d_{i,t}, d_{h,t})$, consumption $(c_{f,t}, c_{i,t}, c_{h,t})$, and net worth $(n_{f,t}, n_{i,t}, n_{h,t})$, such that:

1. Initial net worth satisfy $n_{h,0} = q_0 k_{h,0}$, $n_{i,0} = q_0 k_{i,0}$, and $n_{f,0} = q_0 k_{f,0}$ for $h \in \mathbb{H}$, $i \in \mathbb{I}$, and $f \in \mathbb{F}$, respectively.

2. Taking the processes for q_t, r_t^b, and r_t^d as given, households $h \in \mathbb{H}$, intermediaries $i \in \mathbb{I}$, and firms $f \in \mathbb{F}$ solve their respective optimization problems.

3. Goods, capital, and credit markets clear

$$\begin{aligned}
&\int_{\mathbb{F}} c_{f,t}df + \int_{\mathbb{I}} c_{i,t}di + \int_{\mathbb{H}} c_{h,t}dh = \int_{\mathbb{F}} (a_f - \iota_{f,t})k_{f,t}df + \int_{\mathbb{I}} (a_i - \iota_{i,t})k_{i,t}di + \\
&\int_{\mathbb{H}} (a_h - \iota_{h,t})k_{h,t}dh \\
&\int_{\mathbb{F}} k_{f,t}df + \int_{\mathbb{I}} k_{i,t}di + \int_{\mathbb{H}} k_{h,t}di = k_t \\
&\int_{\mathbb{F}} b_{f,t}df - \int_{\mathbb{I}} b_{i,t}di = 0 \\
&\int_{\mathbb{I}} d_{i,t}di - \int_{\mathbb{H}} d_{h,t}dh = 0
\end{aligned}$$

where k_t denotes aggregate capital, whose law of motion is given by

$$\frac{dk_t}{k_t} = \left[\Phi(\iota_t) - \delta\right] dt + \sigma dZ_t \tag{3-8}$$

3.2 Equilibrium Characterization

Optimal conditions. The optimal portfolio conditions for firms, intermediaries and households are given by[5]

$$\frac{a_f - a_h}{q_t} = r_t^b - r_t^d + \left(\zeta_t^f - \zeta_t^h\right)(\sigma + \sigma_t^q) + \lambda_t^h \tag{3-9}$$

$$\frac{a_f - a_i}{q_t} = \left(\zeta_t^f - \zeta_t^i\right)(\sigma + \sigma_t^q) + \lambda_t^{1,i} \tag{3-10}$$

$$r_t^b - r_t^d = \lambda_t^{2,i}(1 - \nu^i) \tag{3-11}$$

where ζ_t is the price of risk, and λ_t^h, $\lambda_t^{1,i}$, and $\lambda_t^{2,i}$ denote the Lagrange multipliers associated with the short-sale constraint for households, the short-sale constraint for intermediaries, and the collateral constraint, respectively. Note that if the collateral constraint binds, the credit spread in 3-11 is positive[6].

Optimal consumption and price of risk. The optimal consumption and the price of risk are given by

$$c_t^j = \rho_j n_t^j; \quad \text{with } j \in \{f, i, h\} \tag{3-12}$$

$$\zeta_t^j = \theta_t^{j,k}(\sigma + \sigma_t^q); \quad \text{with } j \in \{f, i, h\} \tag{3-13}$$

Note that the agent consume a fixed proportion of his wealth each period. The price of risk captures the premium offered to the agent to hold the risky asset. It is increasing with agents' leverage and asset volatility.

State variables. The state of the economy is fully characterized by two state variables: the wealth share of non-financial firms, $\eta_t^f \equiv n_t^f / n_t$, and the wealth share of intermediaries, $\eta_t^i \equiv n_t^i / n_t$, where $n_t = q_t k_t$ denotes aggregate net worth.

Constraints as function of the state variables. The capital allocations of firms and intermediaries can be expressed as

$$\theta_t^{f,k} = \frac{\kappa_t^f}{\eta_t^f}; \qquad \theta_t^{i,k} = \frac{\kappa_t^i}{\eta_t^i}$$

where $\kappa_t^f \equiv k_t^f / k_t$ and $\kappa_t^i \equiv k_t^i / k_t$. Substituting these expressions into the

[5]Proof is provided in the Appendix.
[6]Further details on the credit spread are provided in Proposition 2 below.

household short-sale constraint, the intermediary short-sale constraint, and the collateral constraint yields the following set of conditions

$$\kappa_t^h \geq 0; \quad \kappa_t^i \geq 0; \quad \kappa_t^f + \kappa_t^i \leq \frac{1}{1-\nu^i}\eta_t^i + \eta_t^f$$

In the collateral constraint equation, when intermediary wealth is sufficiently small, the constraint binds, reducing intermediaries' capacity to extend credit to firms.

Endogenous volatility. The endogenous volatility term σ_t^q can be derived from the capital price function $q(\eta_t^f, \eta_t^i)$.

$$\sigma_t^q = \frac{\partial_f q_t}{q_t}\eta_t^f \sigma_t^{f,\eta} + \frac{\partial_i q_t}{q_t}\eta_t^i \sigma_t^{i,\eta} \tag{3-14}$$

where $\partial_f q_t = \partial q_t / \partial \eta_t^f$ and $\partial_i q_t = \partial q_t / \partial \eta_t^i$. Substituting the expressions for $\sigma_t^{f,\eta}$ and $\sigma_t^{i,\eta}$ into (3-14) yields

$$\sigma + \sigma_t^q = \frac{\sigma}{1 - \frac{\eta_t^f}{q_t}\frac{\partial q_t}{\partial \eta_t^f}\left(\frac{\kappa_t^f - \eta_t^f}{\eta_t^f}\right) - \frac{\eta_t^i}{q_t}\frac{\partial q_t}{\partial \eta_t^i}\left(\frac{\kappa_t^i - \eta_t^i}{\eta_t^i}\right)} \tag{3-15}$$

Equation (3-15) captures the intensity of capital misallocation from firms and intermediaries to less productive households. This measure increases in two key factors: (i) the elasticities of the capital price with respect to the net worth of firms and intermediaries; and (ii) the leverage of non-financial firms and intermediaries, defined as $(\kappa_t^f - \eta_t^f)/\eta_t^f$ and $(\kappa_t^i - \eta_t^i)/\eta_t^i$.

Scaled Output. In this setup, aggregate output is not directly observable. However, we can define output scaled by aggregate capital ($\hat{y}_t = y_t/k_t$) as a proxy for GDP[7].

$$\hat{y}_t = a_f \kappa_t^f + a_i \kappa_t^i + a_h \kappa_t^h \tag{3-16}$$

Since firms are the most productive agents, a reallocation of capital from firms to less productive intermediaries or households reduces scaled output.

Lagrange multipliers and equilibrium regions. In our model, the efficient allocation arises when non-financial firms hold the entire capital stock, that is, when $\lambda_t^h > 0$ and $\lambda_t^{1,i} > 0$. An equilibrium with $\lambda_t^h = 0$ or $\lambda_t^{1,i} = 0$ signals the presence of either a financial or a non-financial recession. In this framework, financial recessions are characterized by downturns in which intermediaries are poorly capitalized and their collateral constraint binds. In

[7] We use the terms "output" and "scaled output" interchangeably in the remainder of this paper.

contrast, non-financial recessions occur when intermediaries are relatively well-capitalized, but non-financial firms are undercapitalized. Table 3.1 outlines each equilibrium region.

Table 3.1: Equilibrium Regions - Definitions

Regime	Region	$\kappa^h \geq 0$	$\kappa^i \geq 0$	Collateral
Non-Financial Rec.	R0	slack	slack	slack
Financial Rec.	R1	slack	slack	binds
Non-Financial Rec.	R2	slack	binds	slack
Financial Rec.	R3	slack	binds	binds
Non-Financial Rec.	R4	binds	slack	slack
Financial Rec.	R5	binds	slack	binds
Expansion	R6	binds	binds	slack
Expansion	R7	binds	binds	binds

Proposition 1. Region R2: $\{\lambda_t^h = 0, \lambda_t^{1,i} > 0, \lambda_t^{2,i} = 0\}$ cannot arise in equilibrium. Intuitively, when intermediaries and households have wealth greater than zero and the collateral constraint is slack, it is never optimal for households to hold capital while intermediaries do not, as intermediaries are strictly more productive than households. Proof is provided in the Appendix.

Proposition 2. The interest rates on bonds and deposits satisfy $r_t^b \geq r_t^d$. When the intermediary collateral constraint is not binding, the two rates are equal: $r_t^b = r_t^d$. Conversely, when the collateral constraint binds, so that households are marginal in pricing deposits to intermediaries, while intermediaries are marginal in pricing bonds to firms, intermediaries charge a premium over their funding cost, implying $r_t^b > r_t^d$. Proof is provided in the Appendix.

Equilibrium conditions. Given the law of motion of the state variables, η_t^f and η_t^i, the dynamics of the economy can be fully characterized by the following system of equations

$$\text{Market Clearing: } 0 = q(\rho_f \eta^f + \rho_i \eta^i + \rho_h \eta^h) - (a_f \kappa^f + a_i \kappa^i + a_h \kappa^h) + \iota$$

$$\text{Portfolio 1: } \frac{a_f - a_h}{q} \geq r_t^b - r_t^d + \left(\frac{\kappa^f}{\eta^f} - \frac{\kappa^h}{\eta^h}\right)(\sigma + \sigma^q)^2, \text{ with equality if } \kappa_t^h > 0$$

$$\text{Portfolio 2: } \frac{a_f - a_i}{q} \geq \left(\frac{\kappa^f}{\eta^f} - \frac{\kappa^i}{\eta^i}\right)(\sigma + \sigma^q)^2, \text{ with equality if } \kappa_t^i > 0$$

$$\text{Portfolio 3: } r_t^b - r_t^d \geq 0, \text{ with equality if } \kappa_t^f + \kappa_t^i < \frac{1}{1-\nu^i}\eta_t^i + \eta_t^f$$

$$\text{Endogenous Volatility: } \sigma + \sigma_t^q = \frac{\sigma}{1 - \frac{\eta_t^f}{q_t}\frac{\partial q_t}{\partial \eta_t^f}\left(\frac{\kappa_t^f - \eta_t^f}{\eta_t^f}\right) - \frac{\eta_t^i}{q_t}\frac{\partial q_t}{\partial \eta_t^i}\left(\frac{\kappa_t^i - \eta_t^i}{\eta_t^i}\right)}$$

where $\eta_t^h = 1 - \eta_t^f - \eta_t^i$ and $\kappa_t^h = 1 - \kappa_t^f - \kappa_t^i$.

4
Numerical Approach

From a standard finite-difference perspective, one possible strategy for solving this system of equations is to implement a generalized Newton-Raphson (NR) method. Mazumder (2016) present the Stone implicit method for solving systems of PDEs. However, applying such methods in models with three occasionally binding constraints and a triangular state space is considerably complex, and convergence is not always guaranteed, as highlighted by d'Avernas et al. (2023).

While deep learning methods have recently gained popularity for solving PDEs[1], their primary limitation lies in capturing transition dynamics across multiple equilibrium regions. Since neural networks do not inherently produce viscosity solutions, their derivative approximations tend to be inaccurate near kink points that separate equilibrium regions[2], a challenge that is particularly relevant in models with financial frictions.

Probabilistic Approach. The probabilistic reinforcement learning (PRL) approach proposed by Huang (2023) addresses these challenges by reformulating the PDEs as a system of forward–backward stochastic differential equations (FBSDEs). This method provides two main advantages over the analytical (PDE-based) approach. First, it circumvents differentiability requirements, making it well suited to models with nonlinear features and occasionally binding constraints. Second, it reduces the solution of the model to computing the value function and an auxiliary variable using neural networks, resulting in a system of two equations with two unknowns that can be solved by simulating paths of Brownian motions[3].

Suppose we want to solve the standard multidimensional HJB equation in continuous time

$$\rho V(\mathbf{x}_t) = u(y_t) + \left(\mathbf{x}_t \mu_t^{\mathbf{x}}\right) V'(\mathbf{x}_t) + \frac{1}{2}\left(\mathbf{x}_t \sigma_t^{\mathbf{x}}\right)^2 V''(\mathbf{x}_t)$$

By the martingale representation theorem, we can rewrite this equation in the

[1]See Fernandez-Villaverde et al. (2024) for a broader review.

[2]See Wu et al. (2024).

[3]A detailed description of the numerical approach is provided in the Appendix. We validate the PRL method by comparing it to the solution of a simplified version of the model obtained using the Newton–Raphson (NR) method, showing that it accurately captures the dynamics implied by the theoretical framework.

following way

$$V(\mathbf{x}_{t+dt}) = (1+\rho dt)V(\mathbf{x}_t) - u(y_t)dt + \mathbf{z}_t dZ_t$$

where $\mathbf{z}_t = z(\mathbf{x}_t)$ is an unknown function of the state variables and $\mathbf{x}_{t+dt}$ follows the law of motion

$$\mathbf{x}_{t+dt} = \mathbf{x}_t + \mathbf{x}_t \mu_t^{\mathbf{x}} dt + \mathbf{x}_t \sigma_t^{\mathbf{x}} dZ_t$$

The PRL method employs neural networks to approximate the policy functions implied by this HJB equation. Specifically, it solves the following problem

$$\begin{aligned}
\min_{\theta} \quad & \frac{1}{NT}\sum_{i}^{N}\sum_{j}^{T}\left[\hat{f}(\hat{\mathbf{x}}^{i,j}) - (1+\rho\Delta_t^{i,j})\hat{f}(\mathbf{x}^i) + u(\mathbf{x}^i)\Delta_t^{i,j} - \mathbf{z}_t\Delta_z^{i,j}\right] \\
\text{s.t.} \quad & \hat{\mathbf{x}}^{i,j} = \mathbf{x}^i\left(1 + \mu^{\mathbf{x}}\Delta_t^{i,j} + \sigma^{\mathbf{x}}\Delta_z^{i,j}\right) \\
& \hat{y}^{i,j} = y^i\left(1 + \mu^{y}\Delta_t^{i,j} + \sigma^{y}\Delta_z^{i,j}\right)
\end{aligned}$$

where θ denotes the neural network parameters (weights and biases), N is the number of samples drawn uniformly from the state space, and T is the number of time steps. The term $\Delta_t^{i,j} = t_j^i - t_{j-1}^i$ represents the time increment, while $\Delta_z^{i,j} = Z_j^i - Z_{j-1}^i$ corresponds to the Brownian motion increment, simulated from a normal distribution $N(0, \Delta_t^{i,j})$.

Under this method, the solution reduces to finding the coefficient $\mathbf{z}_t$ and the value function $V(\mathbf{x}_t)$, which can be obtained by simulating two paths for the Brownian motion increment dZ_t.

Equilibrium conditions for the probabilistic method. For this method, the equilibrium conditions are characterized by the following system of equations.

$$\begin{aligned}
&\text{Market Clearing: } 0 = q\left(\rho_f\eta^f + \rho_i\eta^i + \rho_h\eta^h\right) - \left(a_f\kappa^f + a_i\kappa^i + a_h\kappa^h\right) + \iota \\
&\text{Portfolio 1: } 0 = \min\left\{\frac{a_f - a_h}{q} - (r_t^b - r_t^d) - \left(\frac{\kappa^f}{\eta^f} - \frac{\kappa^h}{\eta^h}\right)(\sigma + \sigma^q)^2, \kappa_t^h\right\} \\
&\text{Portfolio 2: } 0 = \min\left\{\frac{a_f - a_i}{q} - \left(\frac{\kappa^f}{\eta^f} - \frac{\kappa^i}{\eta^i}\right)(\sigma + \sigma^q)^2, \kappa_t^i\right\} \\
&\text{Portfolio 3: } 0 = \min\left\{r_t^b - r_t^d, \frac{1}{1-\nu^i}\eta_t^i + \eta_t^f - \kappa_t^f - \kappa_t^i\right\} \\
&\text{FSDE 1: } \frac{d\eta_t^f}{\eta_t^f} = \mu_t^{f,\eta}dt + \sigma_t^{f,\eta}dZ_t \\
&\text{FSDE 2: } \frac{d\eta_t^i}{\eta_t^i} = \mu_t^{i,\eta}dt + \sigma_t^{i,\eta}dZ_t
\end{aligned}$$

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

where the drift and volatility terms for the FBSDEs are given by[4]

$$\mu_t^{f,\eta} = -\rho_f + \frac{\kappa_t^f}{\eta_t^f}\left(\frac{a_f - \iota_t}{q_t}\right) + \left(\frac{\kappa_t^f}{\eta_t^f} - 1\right)\left[\mu_t^q + \Phi(\iota_t) - \delta + \sigma\sigma_t^q - r_t^b - (\sigma + \sigma_t^q)^2\right]$$

$$\sigma_t^{f,\eta} = \left(\frac{\kappa_t^f}{\eta_t^f} - 1\right)(\sigma + \sigma_t^q)$$

$$\begin{aligned}\mu_t^{i,\eta} = & -\rho_i + \left(\frac{\kappa_t^f - \eta_t^f}{\eta_t^i}\right)(r_t^b - r_t^d) + \frac{\kappa_t^i}{\eta_t^i}\left(\frac{a_i - \iota_t}{q_t}\right) + \\ & + \left(\frac{\kappa_t^i}{\eta_t^i} - 1\right)\left[\mu_t^q + \Phi(\iota_t) - \delta + \sigma\sigma_t^q - r_t^d - (\sigma + \sigma_t^q)^2\right]\end{aligned}$$

$$\sigma_t^{i,\eta} = \left(\frac{\kappa_t^i}{\eta_t^i} - 1\right)(\sigma + \sigma_t^q)$$

$$\mu_t^q = r_t^b - \frac{a_f - \iota_t}{q_t} - \Phi(\iota_t) + \delta - \sigma\sigma_t^q + \zeta_t^f(\sigma + \sigma_t^q)$$

[4]Proofs are provided in the Appendix. Unlike FD methods, the PRL method does not require explicitly imposing the consistency condition for the endogenous volatility σ_t^q. This difference arises because PRL recovers the policy function for the capital price q_t directly through the BSDE, avoiding the need for numerical derivative approximations in σ_t^q that are generally unstable.

5
Solution

Calibration. The exogenous volatility of capital σ, the elasticity of the investment rate ϕ, the depreciation rate δ, the productivity of firms a_f, the productivity of households a_h, and the discount rate of firms ρ_f are calibrated to match the values used in Brunnermeier and Sannikov (2016). The productivity of intermediaries a_i is assumed to lie between that of firms and households. The discount rates for intermediaries and households are chosen to be lower than that of firms, satisfying the model condition. The full set of calibrated parameters is reported in Table 5.1.

Table 5.1: Calibration

Parameter	Description	Value
ρ_f	Time-discount rate for firms	0.06
ρ_i	Time-discount rate for intermediaries	0.05
ρ_h	Time-discount rate for households	0.04
a_f	Productivity for firms	0.11
a_i	Productivity for intermediaries	0.07
a_h	Productivity for households	0.03
δ	Depreciation rate	0.05
ϕ	Elasticity of investment rate	10
σ	Volatility of capital shock	0.1
ν^i	Collateral constraint parameter	0.7

5.1
Equilibrium Regions

The solution of this model features five equilibrium regions, as illustrated in Figure 5.1 and described in Table 5.2. As established in Proposition 1, region R2 is not attainable in equilibrium. Regions R5 and R7, while theoretically possible, are not observed in the solution. The red circle indicate the stable stochastic steady state[1], identified as the points where the drifts of η^f and η^i are zero. Notably, the steady state lie within region R6, where the short-sale constraints for households and intermediaries are binding, while the collateral constraint remains slack.

Building on this classification, we examine the behavior of the economy within the financial and non-financial recession regions. Figure 5.2 presents top-down views of the main variables, overlaid with the corresponding equilibrium

[1] From this point onward, we use "steady state" to refer to the stable stochastic steady state.

Figure 5.1: Occasionally Binding Constraint Regions

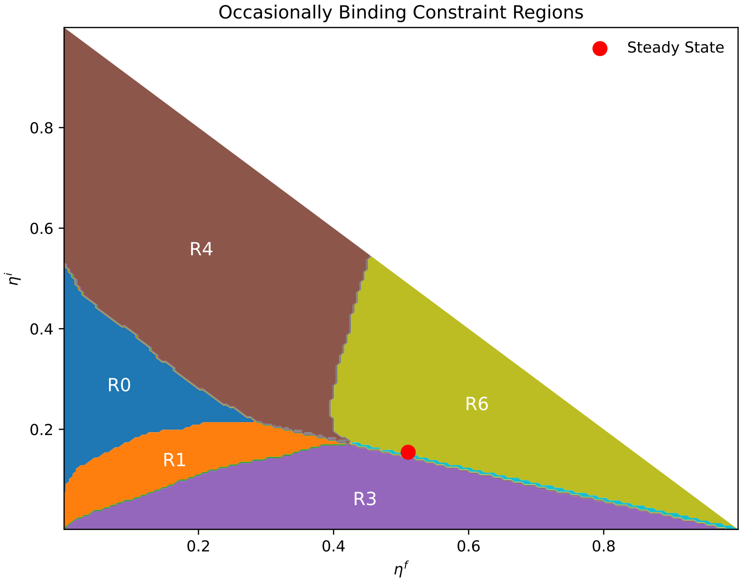

Note: Top-down view of the state space with the equilibrium regions. The borders between regions correspond to points where the constraints bind.

Table 5.2: Observed Equilibrium Regions - Definitions

Regime	**Region**	$\kappa^h \geq 0$	$\kappa^i \geq 0$	**Collateral**
Non-Financial Rec.	R0	slack	slack	slack
Financial Rec.	R1	slack	slack	binds
Financial Rec.	R3	slack	binds	binds
Non-Financial Rec.	R4	binds	slack	slack
Expansion	R6	binds	binds	slack

regions. In the expansionary region (R6), firms hold the entire capital stock, the price of capital remains high and relatively stable, and endogenous volatility is low. In a non-financial recession region (R0 or R4), firms experience a sharp decline in their wealth share, forcing them to liquidate capital to less productive agents. This reallocation depresses the capital price and increases endogenous volatility.

In the financial recession region R1, the collateral constraint binds, limiting the ability of both firms and intermediaries to finance and hold capital on their balance sheets. Consequently, capital shocks in this region have a limited impact on firms' and intermediaries' asset sales to less productive agents, and, by extension, on endogenous volatility. A similar dynamic occurs in region R3, where the collateral constraint is also binding. However, in R3, intermediaries fully exit the capital market due to their diminished wealth. Capital shocks in this region fail to generate high volatility, as firms' asset sales are constrained by their limited capital holdings.

Severe output contractions can occur in both financial and non-financial

Figure 5.2: Equilibrium Regions

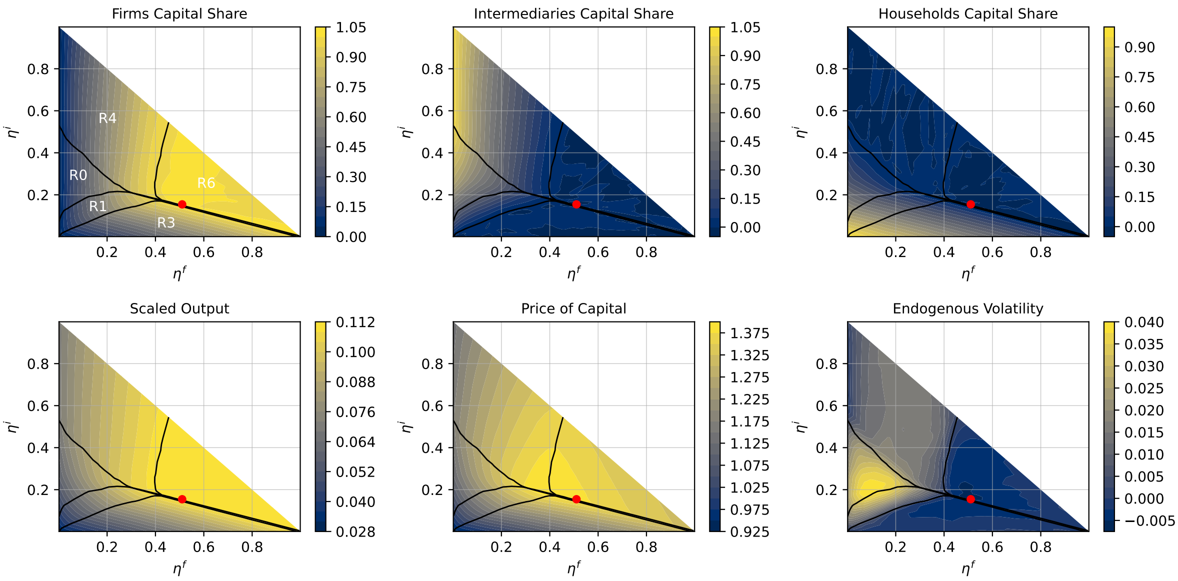

Note: Top-down view of the main variables. The equilibrium regions are delineated by solid black lines.

recessions. Notably, region R3 exhibits the most pronounced contraction, as households hold a large share of capital that firms and intermediaries are unable to retain due to the binding collateral constraint. As a result, downturns associated with financial recessions tend to be particularly severe, whereas non-financial recessions are relatively milder, despite the elevated volatility in those regions. These findings underscore the role of intermediaries' balance sheets in shaping macroeconomic outcomes, as their weak capitalization is associated with deeper recessions.

5.2 Economy Dynamics

Although the system ultimately converges to the steady state, as shown in Figure 5.3, it may remain in certain regions of the state space for extended periods before full convergence is achieved. The speed at which the economy transitions across regions is measured by the expected rate of adjustment[2] depicted in Figure 5.3. A higher speed of adjustment indicates lower persistence and faster convergence.

The speed of recovery in each region reflects key characteristics of the balance sheets of firms and intermediaries. As the most productive agents, firms earn the highest expected return on capital among the three sectors. When

[2]The speed is computed as the Euclidean norm of the velocity vector: $|v| = \sqrt{(\eta^f \mu^{f,\eta})^2 + (\eta^i \mu^{i,\eta})^2}$. This vector reflects the expected rates of change, $d\eta_t^f/dt$ and $d\eta_t^i/dt$. It also captures the average growth of firms' and intermediaries' wealth.

Figure 5.3: Phase Diagram and Speed of Adjustment

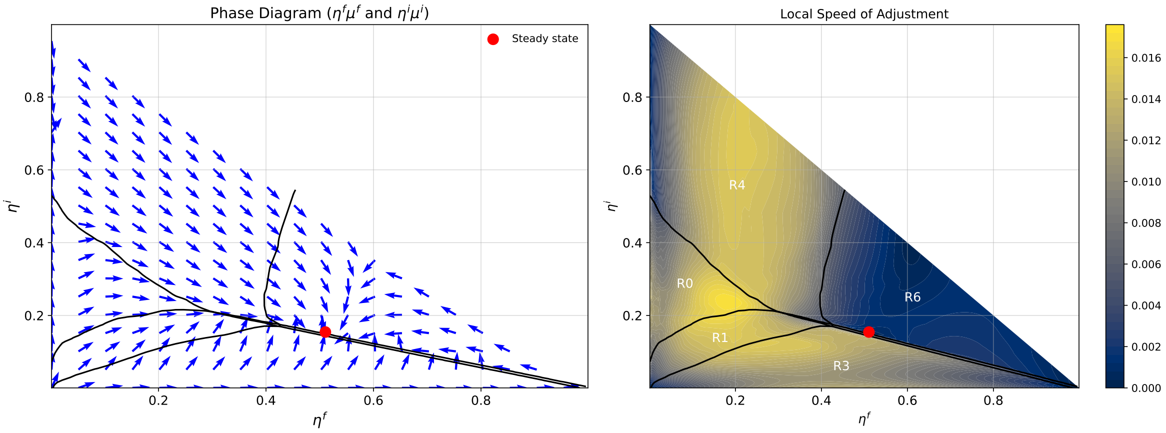

Note: The equilibrium regions are delineated by solid black lines.

intermediaries are well-capitalized but firms are relatively poorly capitalized, the economy enters a non-financial recession region (R0 or R4), where output contractions tend to be short-lived, as firms' balance sheets recover relatively quickly. In these scenarios, firms hold relatively high levels of capital and face elevated equity risk premia (Figure 5.4), as the endogenous volatility and the price of risk rises.

These conditions support a faster rebuilding of net worth by increasing the drift of firms' wealth share, as illustrated in Figure 5.4. However, when firms become severely undercapitalized, that is, when η^f is close to zero, the dynamics change substantially, even though the economy remains in regions R0 or R4. In such cases, recovery is considerably slower, as firms lose a significant portion of their capital by selling it to less productive agents. Consequently, despite the elevated equity premium, the reduced capital holdings severely limit the drift of firms' wealth share, delaying the return to the steady state.

A different dynamic emerges when intermediaries are poorly capitalized. If their collateral constraint binds, shifting the economy into financial recession regions (R1 or R3), intermediaries restrict the supply of credit and raise lending rates to firms by charging a positive spread (Figure 5.4). Although intermediaries benefit from these higher spreads, the downturn becomes deeper and more persistent, as firms face tighter financial constraints on two fronts: reduced access to credit and increased borrowing costs. These pressures lead to a decline in firms' capital holdings, a compression of their equity risk premium, and a reduction in the drift of their wealth share, thereby slowing the pace of recovery.

Figure 5.4: Risk Premium and Credit Spreads

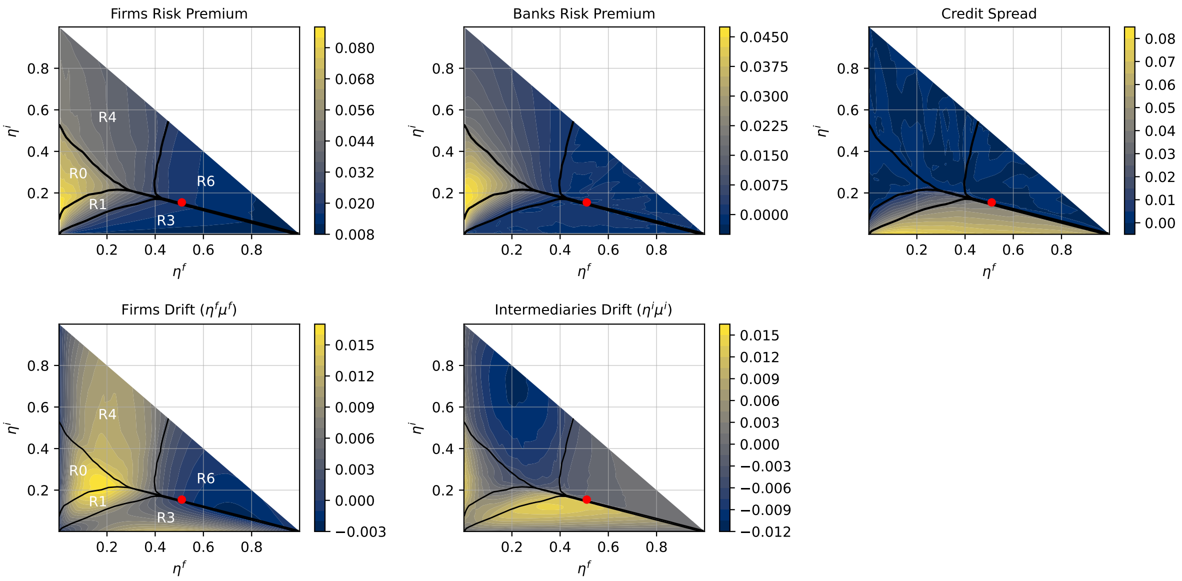

Note: The equilibrium regions are delineated by solid black lines. The red circle denotes the stable stochastic steady state of the economy.

5.3 Risk Exposure

While these results describe the behavior of the economy within each equilibrium region, they do not capture how agents adjust their risk exposure in response to capital shocks. The volatility of the capital share, defined in equation 5-1, quantifies the risk exposure by firms, intermediaries, and households, and reflects how risk is redistributed across sectors in the aftermath of such shocks.

$$\kappa_t^j \sigma_t^{j,\kappa} = \left[\frac{\partial \kappa_t^j}{\partial \eta_t^f} (\kappa_t^f - \eta_t^f) + \frac{\partial \kappa_t^j}{\partial \eta_t^i} (\kappa_t^i - \eta_t^i) \right] (\sigma + \sigma_t^q) \tag{5-1}$$

The volatility of capital share depends on three components: (i) the marginal effect of the wealth share on capital share, (ii) the gap between capital holdings and wealth shares—that is, the degree to which capital is financed through external borrowing (leverage)[3], and (iii) endogenous volatility. Higher leverage and greater endogenous volatility increase the agent's risk exposure and, consequently, their capital share volatility. Negative shocks in these regimes generate stronger economic responses, as risk is redistributed more intensively among agents. In such cases, lower capital volatility, i.e., more negative values, indicates that the agent is offloading risk to other agents.

[3]Leverage in this context differs from the leverage term in the endogenous volatility equation (3-15).

Figure 5.5 displays the (negative) volatility loading for the capital share,[4]. In non-financial recession regions R0 and R4, the collateral constraint does not bind, allowing agents to freely reallocate risk through the capital market. In region R4, only firms and intermediaries hold capital. Since firms typically hold a larger share throughout much of this region, they offload risk to intermediaries after negative capital shocks. In region R0, by contrast, all agents hold capital, but both firms and intermediaries are highly leveraged (Figure 5.5), which amplifies risk transfers from these agents to households.

Figure 5.5: Capital Share Volatility

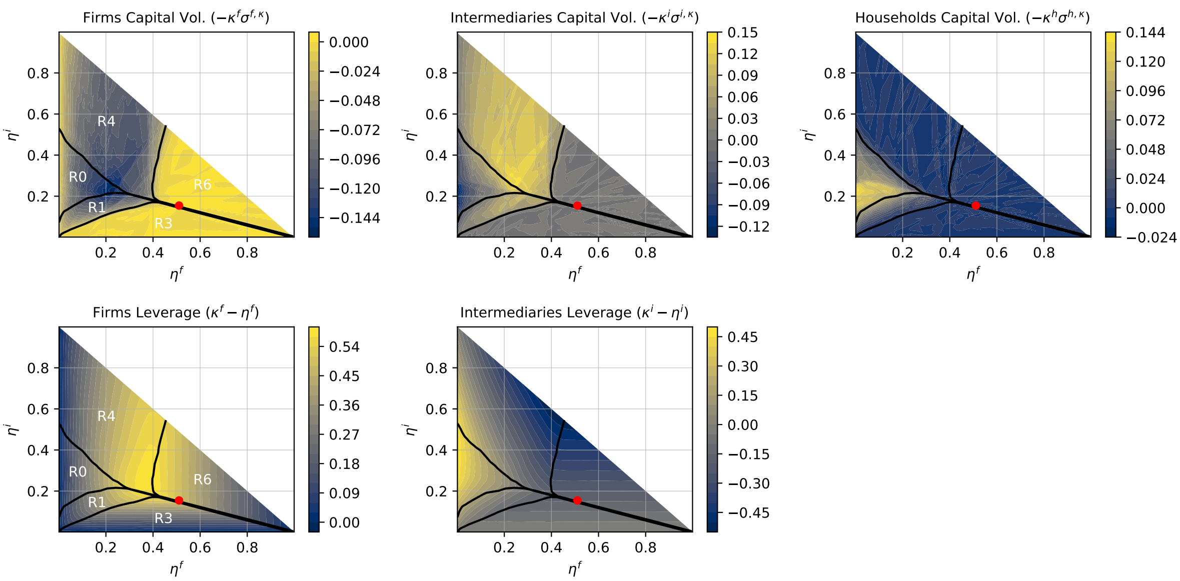

Note: For positive shocks, the color scale should be inverted. The equilibrium regions are delineated by solid black lines. The red circle denotes the stable stochastic steady state of the economy.

In financial recession regions R1 and R3, however, a binding collateral constraint restricts capital reallocation. In region R1, firms face tighter borrowing limits, enabling households to absorb part of the capital stock. As a result, firms, holding less capital, can offload only a limited share of risk to intermediaries and households after negative capital shocks. This friction is even more severe in region R3, where intermediaries hold no capital and are effectively excluded from the risk-sharing process. With firms less leveraged and retaining smaller capital positions than in other regions, their ability to transfer risk to households is further diminished. While this mechanism contributes to greater short-term stability, it also slows down the recovery process.

These results highlight that the leverage of firms and intermediaries is positively associated with both the intensity of risk redistribution across agents

[4]The negative of the volatility loading captures how agents respond to negative capital shocks.

and the degree of instability following capital shocks. Interestingly, they also reveal a paradoxical negative relationship between the severity of downturns and the volatility of capital share, particularly in financial recession regions R1 and R3, where constrained risk-sharing leads to deeper but more stable recessions.

5.4 Economy Response to Shocks

The economy's response to capital shocks can be analyzed by constructing impulse response functions (IRFs) under the assumption that the economy is continuously hit by a negative Brownian shock at the 5% quantile of the standard normal distribution, i.e., $dZ_t = -1.645\sqrt{dt}$, over a three-month horizon. We consider two scenarios: (1) shocks originating in the non-financial recession region R4, with initial state $(\eta^f, \eta^i) = (0.30, 0.30)$; and (2) shocks originating in the financial recession region R3, characterized by $(\eta^f, \eta^i) = (0.30, 0.10)$[5]. For comparison, we also analyze the economy's "expected path" from each of these points back to the steady state in the absence of the deterministic shocks.

Shocks in a Non-Financial Recession Region. A sequence of shocks in the non-financial recession region R4, where both firms and intermediaries hold capital, negatively affects both types of agents, but firms more severely. Figure 5.6 shows the behavior of firms' and intermediaries' wealth share.

Figure 5.6: IRFs in a Non-Financial Recession Region

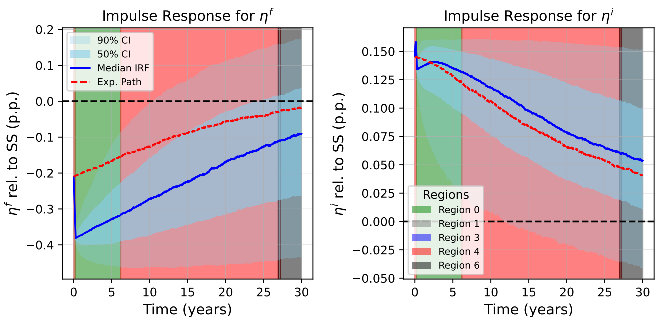

Note: The blue line shows the median IRF following a sequence of negative capital shocks, while the blue shaded areas represent the 50% and 90% confidence intervals. The red dashed line ("expected path") represents the IRF in the absence of deterministic shocks. The shaded background corresponds to the equilibrium regions defined in Figure 5.1.

Figure 5.7 illustrates the response of the main variables. When intermediaries are well-capitalized but the economy is hit by a sequence of negative capital shocks starting at R4, firms, who hold the majority of capital, are the agents most directly affected. As a result, the economy transitions into a deeper

[5] Further details on the implementation are provided in the Appendix.

Figure 5.7: IRFs in a Non-Financial Recession Region

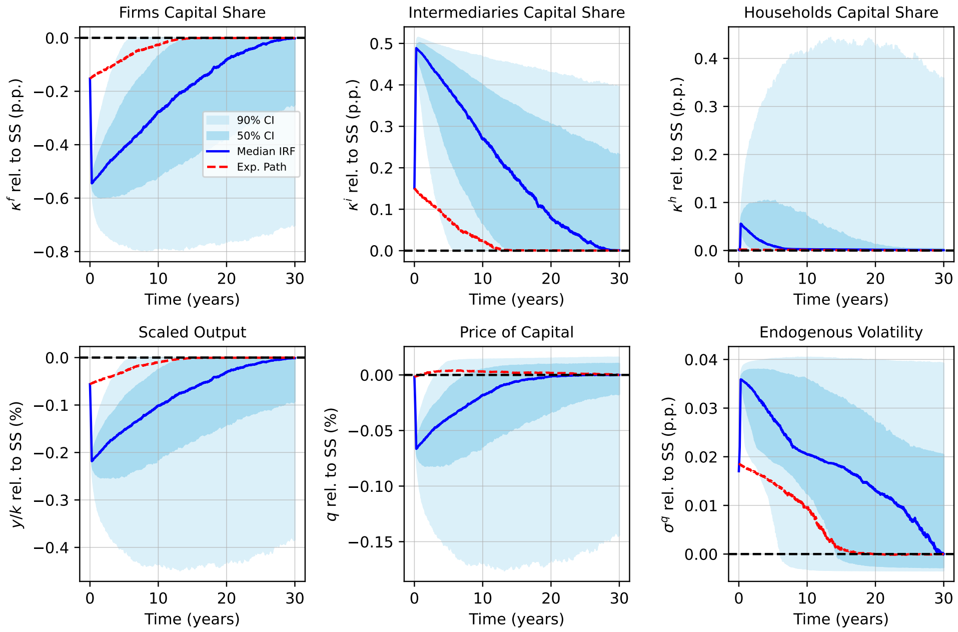

Note: The blue line shows the median IRF following a sequence of negative capital shocks, while the blue shaded areas represent the 50% and 90% confidence intervals. The red dashed line ("expected path") represents the IRF in the absence of deterministic shocks.

non-financial recession at R0, where firms primarily sell capital to intermediaries and, to a lesser extent, to households. Because firms operate with high leverage, these asset sales trigger a sharp erosion of their net worth, amplifying the effect of the initial capital shock.

The resulting wealth losses further constrain firms' ability to retain capital, prompting additional sales to less productive agents. This fire-sale dynamic depresses capital prices and magnifies subsequent wealth declines. The combination of high leverage, forced asset sales, and falling capital prices generates heightened endogenous volatility and places the economy in an unstable equilibrium. In this scenario, scaled output falls sharply relative to its steady-state level, reflecting the inefficient reallocation of capital from firms to intermediaries and households. We refer to the mechanism driving this dynamic as the balance sheet channel.

After the sequence of negative shocks, the economy gradually converges back to the steady state. Firms' capital share, having fallen by nearly 50 p.p. from its steady-state level, recovers slowly. This sluggish adjustment keeps output and capital prices depressed for an extended period, while endogenous volatility remains elevated relative to its steady-state equilibrium. The dynamics of the economy are broadly similar to the case without the sequence of shocks, suggesting that much of the recession's severity is determined by the region in which the economy starts.

Shocks in a Financial Recession Region. When intermediaries are poorly capitalized, the economy exhibits markedly different dynamics compared to the previous exercise. In this case, both firms and intermediaries start with wealth shares well below their steady-state levels, and capital is held only by firms and households. As shown in Figure 5.8, the positive impact of capital shocks on intermediaries' balance sheets is insufficient to restore the economy to its steady state[6]. Instead, the economy remains trapped in the financial recession regions, initially at R3 and subsequently at R1, for an extended period before gradually passing through the non-financial recession region R4 and returning to the steady-state region R6.

Figure 5.8: IRFs in a Financial Recession Region

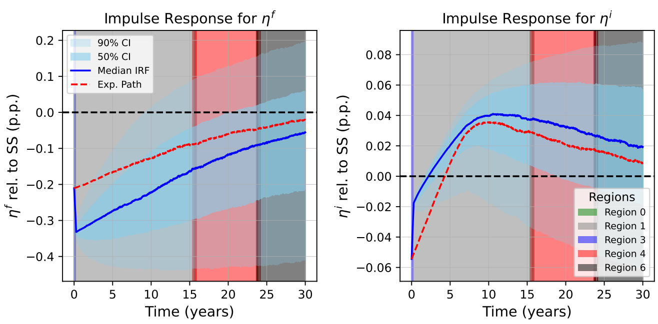

Note: The blue line shows the median IRF following a sequence of negative capital shocks, while the blue shaded areas represent the 50% and 90% confidence intervals. The red dashed line ("expected path") represents the IRF in the absence of deterministic shocks. The shaded background indicates the equilibrium regions defined in Figure 5.1.

This dynamic is illustrated in the variables shown in Figure 5.9. The economy begins in a distressed equilibrium, where capital is held only by firms and households. The sequence of capital shocks primarily affects firms, which still hold a significant share of capital, forcing them to sell part of their holdings to intermediaries as the economy transitions to region R1. In this shift, wealthier banks are able to acquire a limited portion of assets from constrained firms, which, together with the relatively high capital held by households, depresses asset prices. However, because firms and intermediaries are not highly leveraged due to the binding collateral constraint, the balance sheet channel is dampened, and the resulting increase in endogenous volatility from capital sales is smaller than that observed in the previous exercise.

Under these conditions, poorly capitalized intermediaries push the economy into a severe downturn by sharply constraining firms' funding capacity, as

[6]From the definition of intermediaries' wealth-share volatility, $\sigma_t^{i,\eta} = (\kappa_t^i/\eta_t^i - 1)(\sigma + \sigma_t^q)$, it follows that when $\kappa^i = 0$, $\sigma^{i,\eta}$ is negative. Consequently, negative shocks in these regions increase η^i.

Figure 5.9: IRFs in a Financial Recession Region

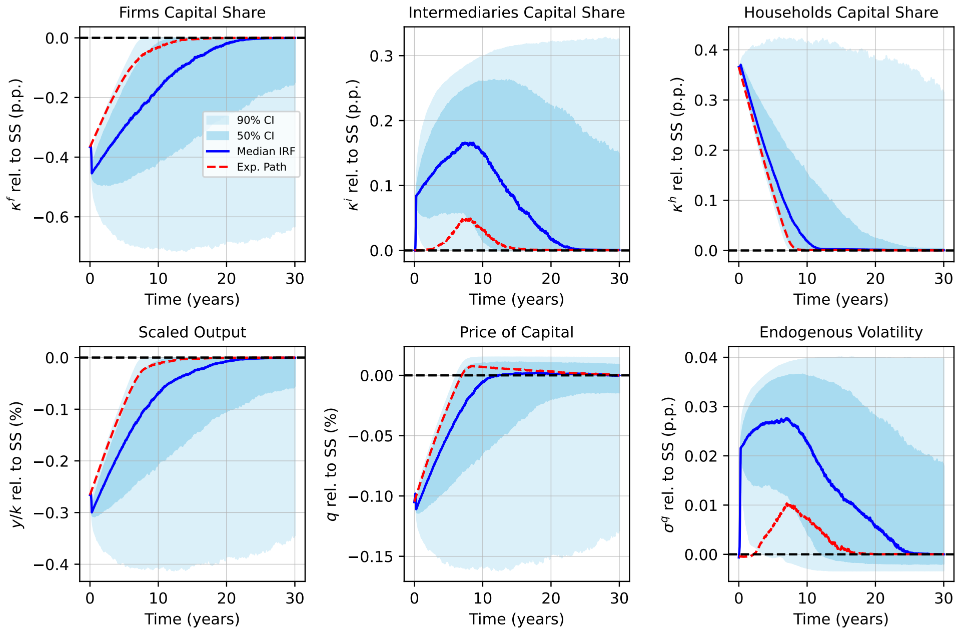

Note: The blue line shows the median IRF following a sequence of negative capital shocks, while the blue shaded areas represent the 50% and 90% confidence intervals. The red dashed line ("expected path") represents the IRF in the absence of deterministic shocks.

Figure 5.10: IRFs in a Financial Recession Region

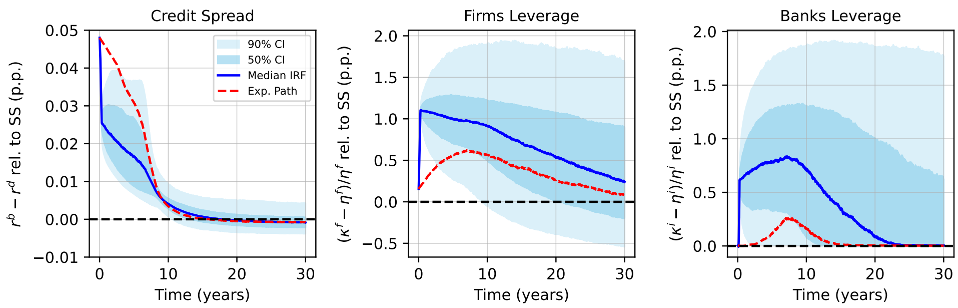

Note: The blue line shows the median IRF following a sequence of negative capital shocks, while the blue shaded areas represent the 50% and 90% confidence intervals. The red dashed line ("expected path") represents the IRF in the absence of deterministic shocks.

both their equity base and access to credit are impaired. With limited leverage, firms are forced to reduce their capital holdings by selling the majority of their assets to households, further deepening the recession. This contraction in credit supply amplifies the decline in output, producing a more severe downturn than in non-financial recessions. We refer to this amplification mechanism as the lending channel, which directly links the severity of financial recessions to the strength of intermediaries' balance sheets.

Similar to the previous exercise, the return to the steady state is gradual. However, the adjustment of intermediaries' capital share and endogenous volatility is non-monotonic. As noted earlier, the sequence of capital shocks increases intermediaries' wealth share, reducing credit spreads (Figure 5.10) and boosting their demand for capital. As the shocks dissipate, firms steadily expand their capital holdings, while intermediaries, benefiting from a higher capital share, rising asset prices, and consequently greater wealth, gradually purchase capital from households.

Around period 10, households sell all the capital in their balance sheets, after which intermediaries cease buying assets. As the credit market returns to normality with rising intermediaries' wealth shares, firms, facing a greater incentive to hold capital due to their higher productivity, begin to increase their holdings by acquiring capital from intermediaries. Endogenous volatility follows this dynamic, rising (falling) as the leverage of both firms and intermediaries increases (decreases).

Output and Financial Stability Although the poor capitalization of intermediaries introduces a powerful amplification mechanism (the lending channel) that deepens downturns, it simultaneously dampens the balance sheet channel, which is the primary source of financial instability during non-financial recessions. This interaction reveals a key trade-off between economic growth and financial stability: tighter financial constraints on intermediaries may intensify the severity and duration of recessions but also limit excessive risk-taking and reduce financial market volatility.

Policies aimed at increasing the equity of financial intermediaries, commonly discussed in the context of wealth crashes, can generate two opposing effects. On one hand, such policies may accelerate the economy's recovery following shocks to firms' balance sheets by allowing intermediaries to mobilize household savings and extend credit to firms. On another hand, they may also push the economy into a more volatile regime, as increased credit provision enables greater capital accumulation by firms, raising leverage and amplifying fire sales in response to adverse shocks. These findings suggest that policies prioritizing financial stability may, paradoxically, lead to worse economic outcomes than those directly targeting the capitalization of intermediaries.

6
Model Limitations and Extensions

This paper sheds light on the dynamics underlying financial and non-financial recessions. Consistent with empirical observations, it demonstrates that financial recessions tend to be more severe than non-financial recessions due to the poor capitalization of intermediaries. However, the model has limitations in capturing the longer persistence of these events[1] and other dynamics that remain open for further investigation. By modeling a single capital shock that affects both firms and intermediaries simultaneously, the model does not allow the economy to transition from the steady state into a financial recession region through capital shocks alone. While it enables analysis of how the economy evolves once it begins in a financial recession region, it fails to capture the effects of capital shocks that affect only intermediaries. This limitation reduces the model's ability to explain the distinct mechanisms underlying the empirical differences between financial and non-financial recessions.

A possible extension to address these limitations is to model distinct capital shocks for firms and intermediaries. Introducing two separate shocks that affect each agent differently could capture the effect of balance sheet crashes in both agents separately. In this framework, firms might anticipate the risk of intermediaries becoming poorly capitalized and constraining the supply of credit, which would reduce their funding capacity, and, as a hedging strategy, sell part of their capital to intermediaries in the steady state. Such a portfolio reallocation would amplify the relevance of intermediary-specific shocks, potentially pushing the economy into a financial recession region. We would also expect endogenous volatility to increase more than in the baseline model, as intermediaries would not only face the lending channel but also reinforce the balance sheet channel through fire sales of capital following negative shocks. Capturing this dynamic, however, may require moving beyond the current framework.

[1] We provide a more direct comparison between the model predictions and proxies for the variables used in the empirical analysis in the Appendix.

7
Conclusion

This paper shows that negative capital shocks to the balance sheets of firms and intermediaries have distinct and significant implications for financial variables and aggregate dynamics, depending on the region of the state space in which the economy starts. In particular, we demonstrate that the distribution of wealth across sectors plays a central role in shaping macro-financial outcomes. The model replicates empirical patterns related to the dynamics of both financial and non-financial recessions.

We find that non-financial recessions are characterized by severe downturns, primarily driven by the balance sheet channel operating through firms. However, when intermediaries are poorly capitalized and collateral-constrained, features of financial recessions, negative shocks lead to deeper contractions in output via the lending channel. This result stems from constrained credit markets, limited asset holdings by firms, and slow wealth accumulation after the shock. Notably, our findings reveal a trade-off between output and financial stability: while collateral-constrained intermediaries amplify the depth of recessions, they also promote a more stable financial environment by limiting excessive risk-taking in credit markets.

These results suggest that understanding the mechanisms behind both financial and non-financial recessions requires analyzing not only the balance sheets of firms and intermediaries separately but also their interaction. Policymakers should pay close attention to the capitalization of financial intermediaries, even when shocks primarily affect the non-financial sector, as well-capitalized intermediaries can play an important role in the speed of recovery.

Bibliography

Adrien, T., Erkko, E., and Muir, T. (2014). Financial intermediaries and the cross-section of asset returns. *The Journal of Finance*, 69(6):2557–2596.

Allen, F. and Gale, D. (2004). Financial intermediaries and markets. *Econometrica*, 72(4):1023–1061.

Barles, G. and Souganidis, P. E. (1991). Convergence of approximations schemes for fully nonlinear second order equations. *Asymptotic Analysis*, 4:271–283.

Baron, M., Verner, E., and Xiong, W. (2021). Banking crisis without panics. *The Quarterly Journal of Economics*, 136(1):51–113.

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

Berger, A. N. and Bouwman, C. H. S. (2013). How does capital affect bank performance during financial crisis. *Journal of Financial Economics*, 109(1):146–176.

Bernanke, B., Getler, M., and Gilchrist, S. (1999). The financial accelerator in a quantitative business cycle framework. *Handbook of Macroeconomics*, 1(C):1341–1393.

Boissay, F., Collard, F., and Smets, F. (2016). Booms and banking crisis. *Journal of Political Economy*, 124(2):489–538.

Bolton, P., Li, Y., Wang, N., and Yang, J. (2025). Dynamic banking and the value of deposits. *The Journal of Finance*. Forthcoming.

Brunnermeier, M. K., Eisenbach, T. M., and Sannikov, Y. (2013). Macroeconomics with financial frictions: A survey. *Advances in Economics and Econometrics*, 10:3–94.

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

Brunnermeier, M. K. and Sannikov, Y. (2016). Macro, money, and finance: A continuous-time approach. *Handbook of Macroeconomics*, 2:1497–1545.

d'Avernas, A., Petersen, D., and Vandeweyer, Q. (2023). A solution method for continuous-time models. Working Paper.

Diamond, D. W. (1984). Financial intermediation and delegated monitoring. *The Review of Economic Studies*, 51(3):393–414.

Diamond, D. W. and Rajan, R. G. (2001). Liquidity risks, liquidity creation, and financial fragility: A theory of banking. *Journal of Political Economy*, 109(2):287–327.

Duarte, V., Duarte, D., and Silva, D. H. (2024). Machine learning for continuous-time finance. *The Review of Financial Studies*, 37(11):3217–3271.

Elenev, V., Landvoigt, T., and Nieuwerburgh, S. V. (2021). A macroeconomic model with financially constrained producers and intermediaries. *Econometrica*, 89(3):1361–1418.

Fan, B., Qiao, E., Anran, J., G., Z., Li, W., and Lu, L. (2024). Deep learning for solving and estimating dynamic macro-finance models. *Computational Economics*, 65:3885–3921.

Fernandez-Villaverde, J., Nuno, G., and Perla, J. (2024). Taming the curse of dimensionality: Quantitative economics with deep learning. NBER, Working Paper.

Gao, W. and Wang, C. (2023). Active learning based sampling for high-dimensional nonlinear partial differential equations. *Journal of Computational Physics*, 475:111848.

Gertler, M. and Kiyotaki, N. (2010). Financial intermediation and credit policy in business cycle analysis. *Handbook of Monetary Economics*, 3:547–599.

Gilchrist, S. and Zakrajsek, E. (2012). Credit spreads and business cycle fluctuations. *American Economic Review*, 102(4):1692–1720.

Gopalakrishna, G. (2023). A macro-finance model with realistic crisis dynamics. Working Paper.

Gopalakrishna, G. and Yuntao, W. (2024). Aliens for continuous time economies. Working Paper.

Guerrieri, L. and Iacoviello, M. (2017). Collateral constraint and macroeconomic asymmetries. *Journal of Monetary Economics*, 90:28–49.

He, Z. and Krishnamurthy, A. (2012). A model of capital and crisis. *The Review of Economic Studies*, 79(2):735–777.

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.

Holmstrom, B. and Tirole, J. (1997). Financial intermediation, loanable funds, and the real sector. *The Quarterly Journal of Economics*, 122(3):663–691.

Hornik, K. (1991). Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(2):251–257.

Hu, Y. and Varas, F. (2025). Intermediary financing without commitment. *Journal of Financial Economics*, 167:104025.

Huang, J. (2023). A probabilistic solution to high-dimensional continuous-time macro and finance models. Working Paper.

Iacoviello, M. (2015). Financial business cycles. *Review of Economic Dynamics*, 18(1):140–163.

Jermann, U. and Quadrini, V. (2012). Macroeconomic effects of financial shocks. *American Economic Review*, 102(1):238–271.

Jordà, O. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182.

Jordà, O., Schularick, M., and Taylor, A. (2016). Macrofinancial history and the new business cycle facts. NBER, Macroeconomics Annual.

Jordà, O., Schularick, M., Taylor, A., and Ward, F. (2018). Global financial cycles and risk premiums. NBER, Working Paper.

Kashyap, A. K., Rajan, R. G., and Stein, J. C. (2002). Banks as liquidity providers: An explanation for the coexistence of lending and deposit-taking. *The Journal of Finance*, 57(1):33–73.

Kiyotaki, N. and Moore, J. (1997). Credit cycles. *Journal of Political Economy*, 105(2):211–248.

Krishnamurthy, A. and Li, W. (2025). Dissecting the mechanisms of financial crisis: Intermediation and sentiment. *Journal of Political Economy*, 133(3):935–985.

Krishnamurthy, A. and Muir, T. (2025). How credit cycles across a financial crisis. *The Journal of Finance*, 80(3):1339–1378.

Mazumder, S. (2016). *Numerical Methods for Partial Differential Equations: Finite Difference and Finite Volume Methods*. Elsevier.

Meh, C. A. and Moran, K. (2010). The role of bank capital in the propagation of shocks. *Journal of Economic Dynamics and Control*, 34(3):555–576.

Mehran, H. and Thakor, A. (2011). Bank capital and value in the cross-section. *The Review of Financial Studies*, 24(4):1019–1067.

Mendoza, E. (2010). Sudden stops, financial crisis, and leverage. *American Economic Review*, 100(5):1941–1966.

Muir, T. (2017). Financial crisis and risk premia. *The Quarterly Journal of Economics*, 132(2):765–809.

Ottonello, P. and Song, W. (2024). Financial intermediaries and the macroeconomy: Evidence from a high-frequency identification. Working Paper.

Petersen, P. and Zech, J. (2025). Mathematical theory of deep learning. arXiv, Working Paper.

Rampini, A. and Viswanathan, S. (2018). Financial intermediary capital. *The Review of Economic Studies*, 86(1):413–455.

Schularick, M. and Taylor, A. M. (2012). Credit booms gone bust: Monetary policy, leverage cycles, and financial crisis, 1870-2008. *American Economic Review*, 102(2):1029–1061.

Villacorta, A. (2023). Business cycles and the balance sheets of the financial and non-financial sectors. ESRB, Working Paper.

Wu, Y., Guo, J., Golapakrishna, G., and Poulos, Z. (2024). Deep-macrofin: Informed equilibrium neural network for continuous-time economic models. arXiv, Working Paper.

A
Local Projection Estimations

Tables A1, A2, and A3 present the estimation results for model (2-1). All regressions include the control variables specified in the model.

Table A1: Equity Returns x Credit-to-GDP

	Credit-to-GDP$_{t,t+1}$			Credit-to-GDP$_{t,t+3}$			Credit-to-GDP$_{t,t+6}$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
r_t^I	0.012*	0.012*	0.015**	0.050***	0.049***	0.051***	0.061**	0.057**	0.065**
	(0.006)	(0.006)	(0.006)	(0.013)	(0.014)	(0.014)	(0.021)	(0.024)	(0.024)
r_t^F	-0.013**	-0.013**	-0.016**	-0.003	-0.003	-0.012	0.021	0.021	0.022
	(0.005)	(0.006)	(0.007)	(0.010)	(0.013)	(0.012)	(0.021)	(0.027)	(0.024)
$r_t^I \cdot r_t^F$	0.002	0.002	-0.002	-0.038**	-0.039*	-0.056***	-0.064***	-0.058**	-0.087***
	(0.006)	(0.007)	(0.009)	(0.017)	(0.020)	(0.019)	(0.016)	(0.021)	(0.022)
Observations	1284	1284	1283	1284	1284	1283	1168	1168	1167
R^2	0.268	0.261	0.301	0.215	0.213	0.289	0.126	0.138	0.246
Country Fixed Effects	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Time Fixed Effects	No	No	Yes	No	No	Yes	No	No	Yes

Standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A2: Equity Returns x Investment-to-GDP

	Investment-to-GDP$_{t,t+1}$			Investment-to-GDP$_{t,t+3}$			Investment-to-GDP$_{t,t+6}$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
r_t^I	0.008***	0.008***	0.005**	0.015***	0.015***	0.010**	0.008*	0.007	0.006
	(0.002)	(0.002)	(0.002)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.006)
r_t^F	0.008***	0.009***	0.006**	0.013***	0.014***	0.015***	0.003	0.004	0.006
	(0.002)	(0.002)	(0.003)	(0.003)	(0.004)	(0.003)	(0.004)	(0.006)	(0.005)
$r_t^I \cdot r_t^F$	-0.009**	-0.009*	-0.004	-0.019***	-0.018**	-0.018***	-0.019***	-0.016*	-0.011
	(0.004)	(0.005)	(0.004)	(0.006)	(0.007)	(0.006)	(0.006)	(0.008)	(0.009)
Observations	1288	1288	1287	1288	1288	1287	1172	1172	1171
R^2	0.164	0.160	0.216	0.166	0.165	0.266	0.135	0.142	0.230
Country Fixed Effects	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Time Fixed Effects	No	No	Yes	No	No	Yes	No	No	Yes

Standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A3: Equity Returns x GDP Growth

	GDP Growth$_{t,t+1}$			GDP Growth$_{t,t+3}$			GDP Growth$_{t,t+6}$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
r_t^I	0.022***	0.024***	0.021**	0.032***	0.039***	0.039**	0.027**	0.043**	0.042**
	(0.005)	(0.005)	(0.008)	(0.011)	(0.011)	(0.014)	(0.012)	(0.016)	(0.015)
r_t^F	0.041***	0.041***	0.035***	0.046***	0.042***	0.053***	0.023	0.010	0.030
	(0.009)	(0.009)	(0.008)	(0.011)	(0.013)	(0.012)	(0.025)	(0.031)	(0.023)
$r_t^I \cdot r_t^F$	-0.036***	-0.043***	-0.032***	-0.063***	-0.079***	-0.083***	-0.068***	-0.103***	-0.089**
	(0.008)	(0.008)	(0.007)	(0.014)	(0.017)	(0.020)	(0.020)	(0.022)	(0.032)
Observations	1288	1288	1287	1288	1288	1287	1172	1172	1171
R^2	0.269	0.279	0.431	0.209	0.247	0.518	0.173	0.251	0.555
Country Fixed Effects	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Time Fixed Effects	No	No	Yes	No	No	Yes	No	No	Yes

Standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B
Optimality Conditions

Firms, intermediaries, and households take the price of capital q_t, the bond interest rate r_t^b, and the deposit interest rate r_t^d as given when choosing their consumption and portfolio allocations. The Hamiltonian for each agent type is given by

Households:

$$\max_{c^h,\theta^{h,k}} e^{-\rho_h t}\log(c_t^h) + \xi_t^h n_t^h \left[-\frac{c_t^h}{n_t^h} + \theta_t^{h,k}\mu_t^{h,k} + (1-\theta_t^{h,k})r_t^d + \lambda_t^h\theta_t^{h,k}\right] - \\ - \xi_t^h n_t^h \zeta_t^h \theta_t^{h,k}(\sigma+\sigma_t^q)$$

Intermediaries:

$$\max_{c^i,\theta^{i,k},\theta^{i,b}} e^{-\rho_i t}\log(c_t^i) + \xi_t^i n_t^i \Bigg[-\frac{c_t^i}{n_t^i} + \theta_t^{i,k}\mu_t^{i,k} + \theta_t^{i,b}r_t^b + (1-\theta_t^{i,k}-\theta_t^{i,b})r_t^d + \lambda_t^{1,i}\theta_t^{i,k} + \\ + \lambda_t^{2,i}\left[1-(1-\nu^i)(\theta_t^{i,b}+\theta_t^{i,k})\right]\Bigg] - \xi_t^i n_t^i \zeta_t^i \theta_t^{i,k}(\sigma+\sigma_t^q)$$

Firms:

$$\max_{c^f,\theta^{f,k}} e^{-\rho_f t}\log(c_t^f) + \xi_t^f n_t^f \left[-\frac{c_t^f}{n_t^f} + \theta_t^{f,k}\mu_t^{f,k} + (1-\theta_t^{f,k})r_t^b\right] - \xi_t^f n_t^f \zeta_t^f \theta_t^{f,k}(\sigma+\sigma_t^q)$$

where λ_t^h, $\lambda_t^{1,i}$, and $\lambda_t^{2,i}$ denote the Lagrange multipliers associated with the short-sale constraint for households ($\theta_t^{h,k} \geq 0$), the short-sale constraint for intermediaries ($\theta_t^{i,k} \geq 0$), and the collateral constraint (3-6), respectively.

The first-order conditions for consumption and portfolio are given by

$$c_t^j\colon\ \frac{e^{-\rho_j t}}{c_t^j} = \xi_t^j; \quad \text{with } j \in \{f, i, h\} \tag{B-1}$$

$$\theta_t^{h,k}\colon\ \frac{a_h - \iota_t}{q_t} + \mu_t^q + \Phi(\iota_t) - \delta + \sigma\sigma_t^q - r_t^d = \zeta_t^h(\sigma+\sigma_t^q) - \lambda_t^h \tag{B-2}$$

$$\theta_t^{i,k}\colon\ \frac{a_i - \iota_t}{q_t} + \mu_t^q + \Phi(\iota_t) - \delta + \sigma\sigma_t^q - r_t^d = \zeta_t^i(\sigma+\sigma_t^q) - \lambda_t^{1,i} + \lambda_t^{2,i}(1-\nu^i) \tag{B-3}$$

$$\theta_t^{i,b}\colon\ r_t^b - r_t^d = \lambda_t^{2,i}(1-\nu^i) \tag{B-4}$$

$$\theta_t^{f,k}\colon\ \frac{a_f - \iota_t}{q_t} + \mu_t^q + \Phi(\iota_t) - \delta + \sigma\sigma_t^q - r_t^b = \zeta_t^f(\sigma + \sigma_t^q) \tag{B-5}$$

where ζ_t is the price of risk, and ξ_t is the stochastic discount factor (SDF).

Optimal consumption and price of risk. The law of motion of the SDF follows a backward stochastic differential equation given by

$$d\xi_t^j = -\frac{\partial \mathcal{H}^j}{\partial n_t^j}dt - \xi_t^j\zeta_t^j dZ_t \quad \text{for } j \in \{f, i, h\} \tag{B-6}$$

For firms, we conjecture that optimal consumption takes the form $c_t^f = g_f n_t^f$[1]. Applying Itô's Lemma to the FOC for c_t^f (B-1) yields

$$\frac{d\xi_t^f}{\xi_t^f} = -\rho_f dt - \frac{dn_t^f}{n_t^f} + \left(\frac{dn_t^f}{n_t^f}\right)^2$$

Substituting the law of motion for firms' net worth into the expression above, we obtain

$$\frac{d\xi_t^f}{\xi_t^f} = \underbrace{\left[-\rho_f + g_f - \theta_t^{f,k}\mu_t^{f,k} - \left(1 - \theta_t^{f,k}\right) r_t^b + \left(\theta_t^{f,k}\right)^2 (\sigma + \sigma_t^q)^2\right]}_{\mu_t^{f,\xi}} dt -$$
$$- \theta_t^{f,k}(\sigma + \sigma_t^q)dZ_t$$

To be consistent with equation (B-6), the drift of the stochastic discount factor must satisfy

$$\xi_t^f \mu_t^{f,\xi} = -\frac{\partial \mathcal{H}^f}{\partial n_t^f} = -\xi_t^f \left[\theta_t^{f,k}\mu_t^{f,k} + \left(1 - \theta_t^{f,k}\right) r_t^b - \zeta_t^f \theta_t^{f,k}(\sigma + \sigma_t^q)\right]$$

$$\zeta_t^f = \theta_t^{f,k}\left(\sigma + \sigma_t^q\right)$$

Substituting the expression for ζ_t^f into the previous equation and matching it with the drift of the law of motion of ξ_t^f, we obtain $\rho_f = g_f$. Hence, the optimal consumption for firms is given by

$$c_t^f = \rho_f n_t^f$$

Equivalently, the optimal consumption and the price of risk for other

[1]We focus on the case of firms, as the household and intermediary cases follow a similar procedure.

agents are given by

$$c_t^j = \rho_j n_t^j; \quad \text{with } j \in \{i, h\}$$

$$\zeta_t^j = \theta_t^{j,k}(\sigma + \sigma_t^q); \quad \text{with } j \in \{i, h\}$$

C
Law of Motion of the Net Worth Share

Applying Itô's product rule in the definition of η_t^j

$$\frac{d\eta_t^j}{\eta_t^j} = \frac{dn_t^j}{n_t^j} - \frac{dn_t}{n_t} + \left(\frac{dn_t}{n_t}\right)^2 - \frac{dn_t^j}{n_t^j}\frac{dn_t}{n_t}$$

where n_t is the aggregate net worth given by $n_t = q_t k_t$.

We postulate that the law of motion for the aggregate net worth is given by $dn_t = n_t \mu_t^n dt + n_t \sigma_t^n dZ_t$. Applying Itô's lemma in $q_t k_t = n_t$, yields

$$\frac{dn_t}{n_t} = \frac{d(q_t k_t)}{q_t k_t} = \left[\mu_t^q + \Phi(\iota_t) - \delta + \sigma\sigma_t^q\right] dt + (\sigma + \sigma_t^q)\, dZ_t$$

For agent $j \in \{f, i, h\}$, the law of motion for η_t^j is given by

$$\frac{d\eta_t^j}{\eta_t^j} = \underbrace{\left[\mu_t^{j,n} - \mu_t^n - (\sigma_t^{j,n} - \sigma_t^n)\sigma_t^n\right]}_{\mu_t^{j,\eta}} dt + \underbrace{(\sigma_t^{j,n} - \sigma_t^n)}_{\sigma^{j,\eta}} dZ_t$$

Firms law of motion. The drift $\mu_t^{f,n}$ and the volatility $\sigma_t^{f,n}$ terms are given by

$$\mu_t^{f,n} = -\rho_f + \frac{\kappa_t^f}{\eta_t^f}\left(\frac{a_f - \iota_t}{q_t} + \mu_t^q + \Phi(\iota_t) - \delta + \sigma\sigma_t^q\right) + \left(1 - \frac{\kappa_t^f}{\eta_t^f}\right) r_t^b$$

$$\sigma_t^{f,n} = \frac{\kappa_t^f}{\eta_t^f}(\sigma + \sigma_t^q)$$

Replacing these conditions into the law of motion for η_t^f yields

$$\mu_t^{f,\eta} = -\rho_f + \frac{\kappa_t^f}{\eta_t^f}\left(\frac{a_f - \iota_t}{q_t}\right) + \left(\frac{\kappa_t^f}{\eta_t^f} - 1\right)\left[\mu_t^q + \Phi(\iota_t) - \delta + \sigma\sigma_t^q - r_t^b - (\sigma + \sigma_t^q)^2\right]$$

$$\sigma_t^{f,\eta} = \left(\frac{\kappa_t^f}{\eta_t^f} - 1\right)(\sigma + \sigma_t^q)$$

Intermediaries law of motion. The drift $\mu_t^{i,n}$ and the volatility $\sigma_t^{i,n}$ terms are given by

$$\mu_t^{i,n} = -\rho_i + \frac{\kappa_t^i}{\eta_t^i}\left(\frac{a_i - \iota_t}{q_t} + \mu_t^q + \Phi(\iota_t) - \delta + \sigma\sigma_t^q\right) + \left(\frac{\kappa_t^f - \eta_t^f}{\eta_t^i}\right) r_t^b + \\ + \left(1 - \frac{\kappa_t^f - \eta_t^f}{\eta_t^i} - \frac{\kappa_t^i}{\eta_t^i}\right) r_t^d$$

$$\sigma_t^{i,n} = \frac{\kappa_t^i}{\eta_t^i}(\sigma + \sigma_t^q)$$

Replacing these conditions into the law of motion for η_t^i yields

$$\mu_t^{i,\eta} = -\rho_i + \left(\frac{\kappa_t^f - \eta_t^f}{\eta_t^i}\right)(r_t^b - r_t^d) + \frac{\kappa_t^i}{\eta_t^i}\left(\frac{a_i - \iota_t}{q_t}\right) + \\ + \left(\frac{\kappa_t^i}{\eta_t^i} - 1\right)\left[\mu_t^q + \Phi(\iota_t) - \delta + \sigma\sigma_t^q - r_t^d - (\sigma + \sigma_t^q)^2\right]$$

$$\sigma_t^{i,\eta} = \left(\frac{\kappa_t^i}{\eta_t^i} - 1\right)(\sigma + \sigma_t^q)$$

D
Proof of Proposition 1

Proposition 1. Region R2: $\{\lambda_t^h = 0, \lambda_t^{1,i} > 0, \lambda_t^{2,i} = 0\}$ cannot arise in equilibrium. Intuitively, when the collateral constraint is slack, it is never optimal for households to hold capital while intermediaries do not, as intermediaries are strictly more productive than households.

Proof. Suppose, for the sake of contradiction, that such an allocation occurs in equilibrium. By definition of the region, the collateral constraint is slack, i.e. $\lambda_t^{2,i} = 0$, intermediaries hold no capital, i.e., $\lambda_t^{i,1} > 0$, and households hold a strictly positive share of capital, i.e. $\lambda_t^h = 0$.

Consider the first-order conditions for the capital portfolios of intermediaries (B-3) and households (B-2). From the Lagrange conditions given above, these expressions can be written as

$$\frac{a_i - \iota_t}{q_t} + \mu_t^q + \Phi(\iota_t) - \delta + \sigma\sigma_t^q - r_t^d < 0$$

$$\frac{a_h - \iota_t}{q_t} + \mu_t^q + \Phi(\iota_t) - \delta + \sigma\sigma_t^q - r_t^d = \theta_t^{h,k}(\sigma + \sigma_t^q)^2$$

Subtracting the second equation from the first yields

$$\frac{a_i - a_h}{q_t} < -\theta_t^{h,k}(\sigma + \sigma_t^q)^2$$

Since $\theta_t^{h,k} \geq 0$ by definition, the right-hand side is strictly negative. Hence, the inequality implies that the left-hand side is also strictly negative. However, by assumption, $a_i > a_h$ and $q_t > 0$, so the left-hand side must be strictly positive. This contradiction implies that such an allocation cannot occur in equilibrium.

E
Proof of Proposition 2

Proposition 2. The interest rates on bonds and deposits satisfy $r_t^b \geq r_t^d$. When the intermediary collateral constraint is not binding, the two rates are equal: $r_t^b = r_t^d$. Conversely, when the collateral constraint binds, so that households are marginal in pricing deposits to intermediaries, while intermediaries are marginal in pricing bonds to firms, intermediaries charge a premium over their funding cost, implying $r_t^b > r_t^d$.

Proof. The result follows directly from the intermediary's first-order condition for bond holdings, given by equation (B-4)

$$r_t^b - r_t^d - \lambda_t^{2,i}(1 - \nu^i) = 0$$

When the collateral constraint is slack, i.e., $\lambda_t^{2,i} = 0$, intermediaries equate the marginal benefit and marginal cost of holding bonds and raising deposits. This implies

$$r_t^b = r_t^d.$$

However, when the collateral constraint binds, i.e., $\lambda_t^{2,i} > 0$, the shadow value of intermediary wealth is strictly positive, reflecting the scarcity of internal funding. In this case, intermediaries require a premium on bond returns to compensate for the constraint, yielding

$$r_t^b > r_t^d.$$

This occurs because, under a binding collateral constraint, households are marginal in pricing deposits, since they are the direct providers of risk-free funding, while intermediaries are marginal in pricing bonds, which they supply to firms. As a result, the wedge between bond and deposit rates increases, as greater tightness in the collateral constraint raises the cost of financial intermediation.

F
Equilibrium Regions

Below, we provide the economic intuition underlying each equilibrium region in this economy.

1. $\lambda_t^h = 0$. In this scenario, the household short-sale constraint is slack, i.e. they hold a positive share of capital ($\theta_t^{h,k} > 0$).

   (a) **Region R0:** $\{\lambda_t^{1,i} = 0, \lambda_t^{2,i} = 0\}$. In this region, intermediaries also hold a positive share of capital, and their collateral constraint is slack. The economy is in a non-financial recession, characterized by poorly capitalized non-financial firms. The deterioration of the non-financial sector's balance sheet is sufficiently severe that firms' demand for external financing falls to levels that do not fully utilize intermediaries' lending capacity. As a result, the collateral constraint faced by intermediaries does not bind.

   (b) **Region R1:** $\{\lambda_t^{1,i} = 0, \lambda_t^{2,i} > 0\}$. In this region, intermediaries also hold capital, but their collateral constraint binds. This scenario reflects a financial recession in which both the financial and non-financial sectors are poorly capitalized, with intermediaries in a more fragile position than in region R0. As a result, intermediaries are unable to fully meet firms' demand for credit. The binding collateral constraint restricts credit supply, thereby limiting firms' investment and production.

   (c) **Region R2:** $\{\lambda_t^{1,i} > 0, \lambda_t^{2,i} = 0\}$. This region is not attainable in equilibrium, as stated in Proposition 2.

   (d) **Region R3:** $\{\lambda_t^{1,i} > 0, \lambda_t^{2,i} > 0\}$. In this region, intermediaries are fully excluded from capital holdings, and their collateral constraint binds tightly. Capital is held exclusively by firms and households. The economy is in a financial recession, and intermediary net worth is so severely depleted that they are forced to fully liquidate their capital holdings to households. Since firms are also financially constrained and unable to absorb this capital, a substantial share of assets is inefficiently allocated to less productive households.

2. $\lambda_t^h > 0$. In this scenario, the household short-sale constraint is binding, i.e. they do not hold capital in their balance sheets ($\theta_t^{h,k} = 0$).

(a) **Region R4:** $\{\lambda_t^{1,i} = 0, \lambda_t^{2,i} = 0\}$. Intermediaries hold capital, and their collateral constraint is slack. This region corresponds to a non-financial recession, characterized by poorly capitalized non-financial firms but relatively strong balance sheets among intermediaries. In this scenario, intermediaries are able to absorb the capital previously held by households and continue supplying credit to firms without facing funding constraints.

(b) **Region R5:** $\{\lambda_t^{1,i} = 0, \lambda_t^{2,i} > 0\}$. In this region, intermediaries hold capital, but their collateral constraint is binding. This scenario reflects a financial recession in which non-financial firms are relatively well-capitalized, enabling them to absorb some of the capital previously held by households. However, the weak balance sheets of intermediaries trigger the collateral constraint, restricting the supply of credit to firms and thereby leading to a downturn.

(c) **Region R6:** $\{\lambda_t^{1,i} > 0, \lambda_t^{2,i} = 0\}$. Intermediaries do not hold capital, and their collateral constraint is slack. This region represents a stable equilibrium in which non-financial firms hold all the capital. Although intermediaries are relatively poorly capitalized, they do not face binding funding constraints, as firms' wealth is sufficient to satisfy their demand for funds. This allocation corresponds to the first-best outcome in the economy.

(d) **Region R7:** $\{\lambda_t^{1,i} > 0, \lambda_t^{2,i} > 0\}$. Intermediaries do not hold capital, and their collateral constraint binds. In this scenario, non-financial firms are sufficiently well-capitalized to absorb all available capital, while intermediaries are constrained limiting the supply of credit. The impact of the credit constrain on aggregate outcomes is limited, as firms' high net worth allows them to self-finance investment.

G
Numerical Approach

Neural networks. Neural networks can be used to approximate policy functions in macro-finance problems. Let $f(\mathbf{x}) : \mathbb{R}^n \to \mathbb{R}$ denote the target function and $\hat{f}(\mathbf{x})$ its neural network approximation. The architecture illustrated in Figure G1 consists of an input layer (blue), two hidden layers (green), and an output layer (red). Aggregation in layer i is given by $a_i(\mathbf{x}) = \mathbf{b}_i + \mathbf{w}_i h_{i-1}(\mathbf{x})$, where $\mathbf{b}_i$ is the bias vector, $\mathbf{w}_i$ is the weight matrix (i.e., the learnable parameters), and $h_i(\mathbf{x}) = \sigma[a_i(\mathbf{x})]$ denotes the activation function (e.g. tanh, ReLU, SiLU, GeLU, sin, etc.). The input to the network is $\mathbf{x}$, and the final output is recursively constructed as

$$\hat{f}(\mathbf{x}) = O\left\{\mathbf{b}_2 + \mathbf{w}_2\left(\sigma\left[\mathbf{b}_1 + \mathbf{w}_1 h_0(\mathbf{x})\right]\right)\right\}$$

where O denotes the activation function used in the output layer, typically chosen to be linear, softplus, or ReLU, depending on whether the output is unrestricted or constrained to be positive.

Figure G1: Deep Neural Network

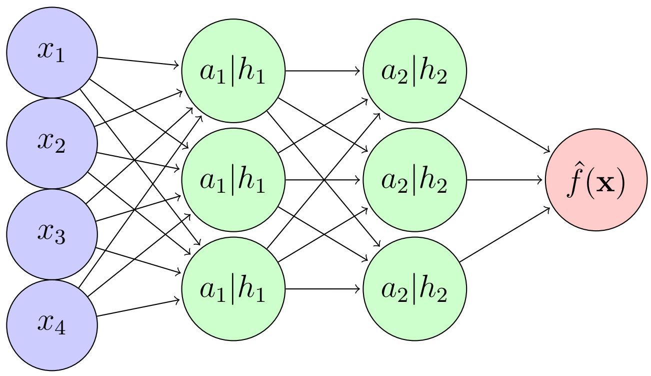

To approximate $f(\mathbf{x})$, we train the neural network by minimizing a composite loss function constructed from the residuals associated with the model's conditions. These residuals may originate from the Hamilton-Jacobi-Bellman (HJB) equation, equilibrium conditions, boundary conditions, market-clearing constraints, or other structural features of the model. The training objective for the neural network can thus be expressed as

$$\min_{\theta} \sum_{i} \lambda_i \mathcal{L}_i(\theta, \mathbf{x}) \tag{G-1}$$

where $\mathbf{x}$ denotes training data sampled from a discretized grid over the state space, θ represents the neural network parameters (i.e., weights and biases), and λ_i are user-defined weights assigned to each component of the loss function.

Each loss component $\mathcal{L}_i$ is typically specified as a mean squared error function

$$\mathcal{L}_i(\theta, \mathbf{x}) = \frac{1}{N}\sum_{n=1}^{N}\left[\hat{f}(\theta, \mathbf{x}) - f(\theta, \mathbf{x})\right]^2$$

which penalizes deviations between the neural network approximation $\hat{f}(\theta, \mathbf{x})$ and the true function $f(\theta, \mathbf{x})$ at each sampled point.

For the approximation to be effective, certain theoretical conditions must be met. The universal approximation theorem, establishes that a feedforward neural network with at least one hidden layer and a non-polynomial activation function can approximate any continuous function on a compact subset of $\mathbb{R}^n$, provided the network has a sufficiently large number of neurons. Consequently, neural networks can, in principle, approximate any continuous policy function arising from a broad class of macro-finance models.[1]

Value function example. To illustrate how the analytical and probabilistic methods operate in the context of neural network applications, consider the value function associated with a continuous-time optimization problem defined over a multi-dimensional state vector $\mathbf{x}_t$

$$V(\mathbf{x}_t) = \max_{\{y_t\}_{t\geq 0}} \mathbb{E}\left[\int_0^\infty e^{-\rho t} u(y_t) dt\right] \tag{G-2}$$

To simplify the exposition, assume y_t is set at its optimal. Then, the associated stochastic HJB equation takes the form

$$\rho V(\mathbf{x}_t)dt = u(y_t)dt + \mathbb{E}_t[dV(\mathbf{x}_t)] \tag{G-3}$$

where the law of motion for the state variable $\mathbf{x}_t$ follows a geometric Brownian motion

$$\frac{d\mathbf{x}_t}{\mathbf{x}_t} = \mu_t^{\mathbf{x}} dt + \sigma_t^{\mathbf{x}} dZ_t \tag{G-4}$$

G.1 Analytical Approach

The analytical approach involves solving models through a system of PDEs. While widely employed, these methods encounter several limitations

[1] A generalized version of the universal approximation theorem (see Petersen and Zech (2025)) extends this result to functions with discontinuities, including jump processes. This broadens the range of macro-finance models whose policy functions can, in theory, be approximated using neural networks.

when applied to nonlinear and high-dimensional models.

These limitations can be analyzed through the deterministic version of the HJB equation (G-3).

$$\rho V(\mathbf{x}_t) = u(y_t) + (\mathbf{x}_t\mu_t^{\mathbf{x}})\, V'(\mathbf{x}_t) + \frac{1}{2}\,(\mathbf{x}_t\sigma_t^{\mathbf{x}})^2\, V''(\mathbf{x}_t) \qquad \text{(G-5)}$$

where $V'(\cdot)$ and $V''(\cdot)$ denote the first and second derivatives of the value function.

In standard numerical procedures, these derivatives are approximated using finite differences. As shown by Barles and Souganidis (1991), convergence to the viscosity solution requires that the numerical scheme satisfy three key properties: consistency, stability, and monotonicity. However, d'Avernas et al. (2023) demonstrate that even generalized FD schemes may fail in models with as few as two state variables when any of these conditions are violated in parts of the state space. As a result, finite-difference methods often prove unsuitable for solving high-dimensional or highly nonlinear models, as they tend to become unstable and computationally intractable.

In contrast, deep learning techniques provide a more scalable and flexible framework for solving high-dimensional PDEs. Recent contributions by Gopalakrishna and Yuntao (2024) and Wu et al. (2024) show that neural networks can approximate solutions to equation (G-5), as well as agents' policy functions, without strictly adhering to the finite-difference criteria of Barles and Souganidis (1991). These methods instead exploit the universal approximation theorem for neural networks.

Despite these advantages, the analytical reinforcement learning approach has limitations. First, while it can handle high-dimensional state spaces, the computational cost increases rapidly with model complexity. Second, its reliance on the differentiability of policy functions limits its applicability to models with non-smooth features, such as kinks or regime changes. In particular, in the presence of occasionally binding constraints, it often becomes necessary to train separate neural networks for each equilibrium region to maintain accuracy and stability.

G.2
Probabilistic Approach

The PRL approach offers a compelling alternative to the limitations of the analytical method. As demonstrated by Huang (2023), systems of PDEs can be reformulated as coupled FBSDEs that characterize the law of motion of the state variables and a key endogenous variables. This reformulation enables a more efficient solution to high-dimensional models that feature multiple occasionally binding constraints and strong nonlinearities, without requiring the differentiability assumption that often pose challenges in the analytical approach.

By the martingale representation theorem, equation (G-3) can be written as

$$V(\mathbf{x}_{t+dt}) = (1 + \rho dt)V(\mathbf{x}_t) - u(y_t)dt + \mathbf{z}_t dZ_t \tag{G-6}$$

where $\mathbf{z}_t = z(\mathbf{x}_t)$ is an unknown function of the state variables, and $\mathbf{x}_{t+dt}$ follows the law of motion

$$\mathbf{x}_{t+dt} = \mathbf{x}_t + \mathbf{x}_t \mu_t^{\mathbf{x}} dt + \mathbf{x}_t \sigma_t^{\mathbf{x}} dZ_t$$

as in equation (G-4).

In this formulation, solving the HJB equation reduces to jointly determining the coefficient vector $\mathbf{z}_t$ and the value function $V(\mathbf{x}_t)$ by simulating two sample paths of the Brownian increment dZ_t. This transforms the high-dimensional PDE problem into a system of two equations with two unknowns, independent of the dimensionality of the state space.

The PRL framework employs neural networks to approximate the unknown functions implied by the equilibrium conditions. The residuals from these conditions are then incorporated into the loss function, similarly to the analytical approach. However, unlike the analytical method, PRL minimizes a loss function derived from equation (G-6), subject to a system of FBSDEs.

$$\begin{aligned}
\min_{\theta} \quad & \frac{1}{NT}\sum_{i}^{N}\sum_{j}^{T}\left[\hat{f}(\hat{\mathbf{x}}^{i,j}) - (1+\rho\Delta_t^{i,j})\hat{f}(\mathbf{x}^i) + u(\mathbf{x}^i)\Delta_t^{i,j} - z_t\Delta_z^{i,j}\right] \\
\text{s.t.} \quad & \hat{\mathbf{x}}^{i,j} = \mathbf{x}^i\left(1 + \mu^{\mathbf{x}}\Delta_t^{i,j} + \sigma^{\mathbf{x}}\Delta_z^{i,j}\right) \\
& \hat{y}^{i,j} = y^i\left(1 + \mu^{y}\Delta_t^{i,j} + \sigma^{y}\Delta_z^{i,j}\right)
\end{aligned}$$

where θ denotes the neural network parameters (weights and biases), N is the number of samples drawn uniformly from the state space, and T is the number of time steps. The term $\Delta_t^{i,j} = t_j^i - t_{j-1}^i$ represents the time increment, while $\Delta_z^{i,j} = Z_j^i - Z_{j-1}^i$ corresponds to the Brownian motion increment, sampled from $N(0, \Delta_t^{i,j})$. The system of FBSDEs is embedded through the law of motion for $\mathbf{x}$ and y. To improve the accuracy of the solution, it is sufficient to increase either N and/or T.

Forward stochastic differential equations (FSDEs). FSDEs describe how past actions influence today's decisions. They typically model the evolution of variables such as wealth, net worth shares, and other state variables for which initial conditions are specified. In the example above, equation (G-4) represents an FSDE.

Backward stochastic differential equations (BSDEs). BSDEs describe how expectations about future variables influence today's decisions. They typically model the evolution of variables such as value functions, prices, and other variables for which terminal conditions are specified. In the example above, the law of motion for y_t represents a BSDE.

Despite its flexibility in handling a broad class of macro-finance models, the probabilistic approach remains computationally intensive, especially when large values of N and/or T are needed to capture strong nonlinearities or multiple occasionally binding constraints. To mitigate this burden, all neural network training is conducted on a T4 GPU from Google Colab.

H
Additional Results

H.1
Global Solution

Figure H1 presents the model's key variables as functions of the two state variables. The price of capital increases with the wealth share of both firms and intermediaries, as these agents expand their capital holdings either through new investment or by trading capital with households and each other. As firms and intermediaries accumulate wealth, endogenous volatility declines, reflecting lower leverage within these sectors and reduced capital price elasticity. Output also rises with higher capital holdings by firms, the most productive agents in the economy.

Figure H1: PRL Solution

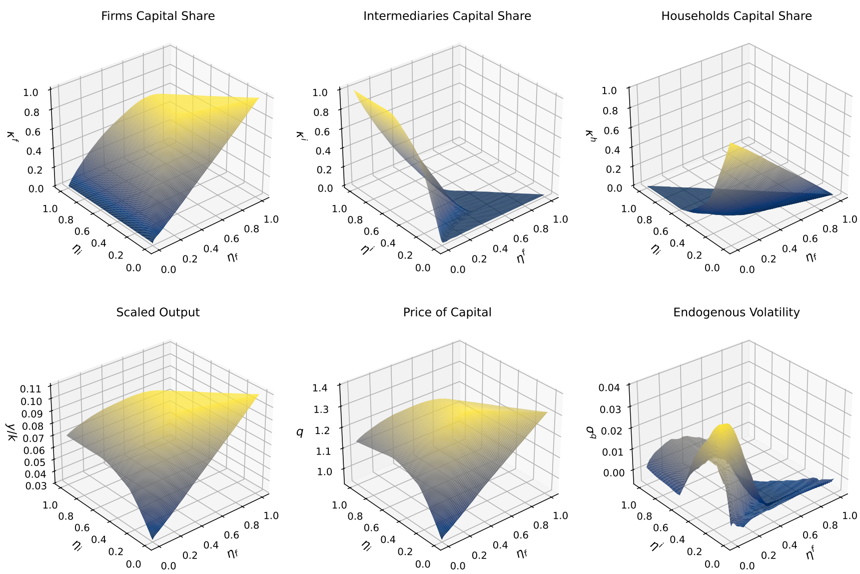

When the wealth of firms or intermediaries declines, they sell a substantial portion of their capital holdings to households. This reallocation leads to a decline in output, as capital shifts toward less productive agents. In this setting, endogenous volatility tends to rise as the price of capital falls. However, the dynamics of volatility depend critically on the capitalization of intermediaries: as intermediary wealth falls, volatility initially increases but eventually declines if the deterioration in intermediary balance sheets continues.

H.2
Unit Shocks

The economy's response to capital shocks across the state space can also be analyzed by simulating unit Brownian shocks, where $dZ_t = +1$ represents a positive shock and $dZ_t = -1$ a negative shock.

The left panel of Figure H2 presents the resulting vector field, illustrating both the direction of risk redistribution and the magnitude of adjustment across the state space. Shorter and thinner arrows indicate slower responses, implying that the speed of adjustment after a sequence of negative capital shocks increases with the leverage and exposure of firms and intermediaries. Prolonged sequences of negative shocks drive firms and/or intermediaries to transfer risk to less productive agents, pushing the economy toward the points $(\eta^f, \eta^i) = (0, 1)$ or $(\eta^f, \eta^i) = (0, 0)$[1]. Importantly, regardless of the region the economy enters after a sequence of shocks, it eventually returns to the steady state, as shown in the phase diagram in the right panel.

Figure H2: Economy Response to Unit Shocks and Phase Diagram

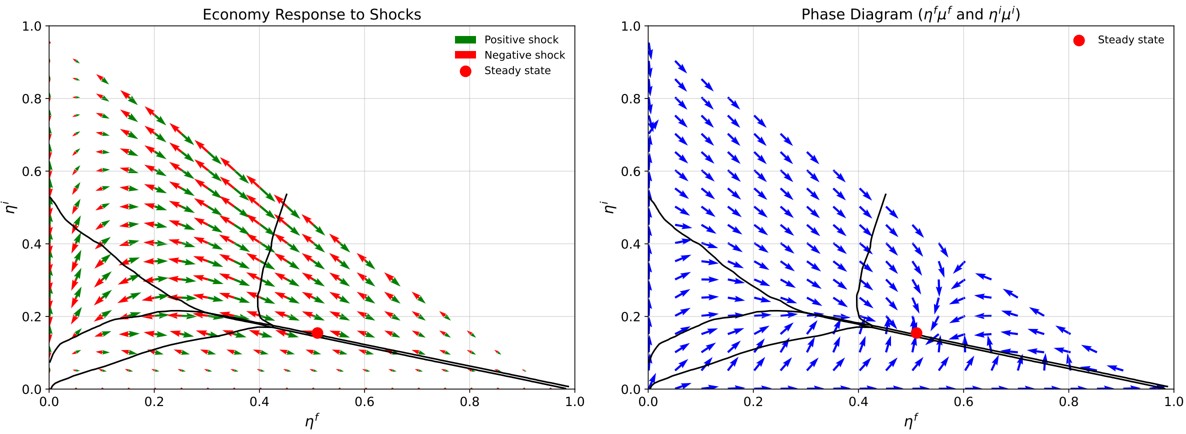

Note: The left panel illustrates the economy's response to unit shocks—positive or negative—at each point in the state space. The direction of the arrows indicates the trajectory of the economy following the shock, while the length and width of the arrows reflect the speed of adjustment. The equilibrium regions are delineated by solid black lines. The right panel presents the phase diagram of the drift in firms' and intermediaries' wealth shares, illustrating the expected paths of the state variables.

Given the high level of risk exposure of firms and intermediaries, shocks in the non-financial recession region R4 are typically short-lived. Yet, a prolonged sequence of adverse shocks can drive the economy to states where firms become severely undercapitalized and hold only a small share of the capital stock. In such cases, recovery tends to be sluggish. As firms' wealth share and capital holdings erode and intermediaries emerge as the primary holders of capital,

[1]Conversely, a sequence of positive shocks drives the economy toward $(\eta^f, \eta^i) = (1, 0)$, where firms accumulate all the wealth.

subsequent shocks compel intermediaries to offload risk onto households, thereby amplifying both the severity and persistence of the downturn in this region.

Financial recession regions exhibit dynamics similar to those of deep non-financial recessions. Once the economy enters these regions, it tends to become trapped, with recovery occurring only gradually. Two key factors contribute to this dynamic. First, because firms and intermediaries hold smaller shares of capital and operate with lower leverage, their volatility loadings are small[2]. As a result, shocks induce only modest changes in their wealth shares. Second, the drift in firms' wealth share is also relatively small, implying that their balance sheets recover slowly following adverse shocks. Together, these dynamics lead the economy to remain in financial recession regions for prolonged periods.

[2]The wealth-share volatility is given by $\sigma_t^{j,\eta} = (\kappa_t^j/\eta_t^j - 1)(\sigma + \sigma_t^q)$, which increases with leverage and endogenous volatility.

I
Evaluating the Probabilistic Method - Two-State Model

I.1
Probabilistic Reinforcement Learning Algorithm

Neural network architecture. We set up the neural network to approximate the functions for the deposit rate r^d, the bond rate r^b, the price of capital q, the capital share of firms κ^f, and the endogenous volatility σ^q. The network consists of four hidden layers with 128 neurons each. We use a sample size of $N = 50,000$ for the state variables, a time step of $\Delta t = 0.005$, and a time horizon of $T = 20$.

We use the ReLU activation function in the hidden layers. In the output layer, we apply the ReLU function to q and κ^f, which are strictly positive by construction, and a linear activation to r^b, r^d, and σ^q.

Initialization. To initialize the network's weights and biases, we conduct a pre-training procedure using the following configuration: epochs $= [300, 300, 300]$, learning rates $= [10^{-4}, 10^{-5}, 10^{-6}]$, and batch sizes $= [512, 512, 512]$. The pre-training proceeds in three successive phases, each with a decreasing learning rate, allowing the network to gradually refine its parameter estimates. The implementation proceeds as follows.

1. Compute $\tilde{r}^{i,d}$, $\tilde{r}^{i,b}$, $\tilde{q}^i$, $\kappa^{i,f}$, and $\tilde{\sigma}^q$ in the neural network.

2. Set target values for the predicted variables based on the solution to the one-state model[1]: $r^{d*} = 0.03$, $r^{b*} = 0.03$, $q^* = 1.3$, $\kappa^{f*} = 0.5$, and $\sigma^{q*} = 0.1$.

3. Using the targets from step 1, compute the average loss function

$$\mathcal{L} = \frac{1}{N}\sum_{i=1}^{N}\Bigg[\left(\tilde{r}^{i,d} - r^{d*}\right)^2 + \left(\tilde{r}^{i,b} - r^{b*}\right)^2 + \left(\tilde{q}^i - q^*\right)^2 + \left(\tilde{\kappa}^{i,f} - \kappa^{f*}\right)^2 + \\ + \left(\tilde{\sigma}^q - \sigma^{q*}\right)^2\Bigg]$$

4. Upon completing all training epochs, save the resulting weights and biases for use in the initialization of the full model.

Training. We train the network on the full model using the following configuration: epochs $= [500, 500, 500]$, learning rates $= [1\text{e-}4, 1\text{e-}5, 1\text{e-}6]$, and

[1]The one-state model follows Brunnermeier and Sannikov (2016); its solution is presented in the Appendix.

batch sizes $= [512, 512, 512]$. After training, the resulting weights and biases are saved to generate the final solution. The complete implementation proceeds as follows.

1. Compute $\hat{r}^d$, $\hat{r}^b$, $\hat{q}$, $\hat{\kappa}^f$, and $\hat{\sigma}^q$ in the neural network.

2. Set $r^d = \hat{r}^d$, $r^b = \hat{r}^b$, $q = \hat{q}$, $\kappa^f = \hat{\kappa}^f$, and $\sigma^q = \hat{\sigma}^q$ to compute κ^i using the market clearing condition.

$$\kappa^i = \frac{1}{\phi(a_i - a_h)} \left[q(1 + \phi\hat{\rho}) - 1 - \phi(a_f - a_h)\kappa_f - \phi a_h\right]$$

   where $\hat{\rho} = \eta^f \rho_f + \eta^i \rho_i + \eta^h \rho_h$.

3. Using the results from step 1, compute the errors based on the system.

$$\begin{aligned}
\text{e1} &= \min\left\{\frac{a_f - a_h}{q} - (r^b_t - r^d_t) - \left(\frac{\kappa^f}{\eta^f} - \frac{\kappa^h}{\eta^h}\right)(\sigma + \sigma^q)^2, \kappa^h_t\right\} \\
\text{e2} &= \min\left\{\frac{a_f - a_i}{q} - \left(\frac{\kappa^f}{\eta^f} - \frac{\kappa^i}{\eta^i}\right)(\sigma + \sigma^q)^2, \kappa^i_t\right\} \\
\text{e3} &= \min\left\{r^b_t - r^d_t, \frac{1}{1 - \nu^i}\eta^i_t + \eta^f_t - \kappa^f_t\right\}
\end{aligned}$$

4. Ensure that the endogenous volatility is within its proper bounds.

$$\text{e4} = \text{ReLU}(-\sigma^q) \text{ for } \eta^f < 0.01$$

   where $\text{ReLU}(x) = \max\{0, x\}^2$.

5. Compute the drift of the price of capital μ^q, the state variable drift $\mu^{f,\eta}$ and $\mu^{i,\eta}$, and volatility $\sigma^{f,\eta}$ and $\sigma^{i,\eta}$.

6. Update the price of capital and the state variable using the system of FBSDE.

$$\begin{aligned}
\eta^f_{t+\Delta t} &= \eta^f_t \left[1 + \mu^{f,\eta}_t \Delta t + \sigma^{f,\eta}_t (Z_{t+\Delta t} - Z_t)\right] \\
\eta^i_{t+\Delta t} &= \eta^i_t \left[1 + \mu^{i,\eta}_t \Delta t + \sigma^{i,\eta}_t (Z_{t+\Delta t} - Z_t)\right] \\
q_{t+\Delta t} &= q_t \left[1 + \mu^q_t \Delta t + \sigma^q_t (Z_{t+\Delta t} - Z_t)\right]
\end{aligned}$$

   where $Z_{t+\Delta t} - Z_t \sim \mathcal{N}(0, \Delta t)$.

7. Compute a new value for $\{\tilde{r}^d, \tilde{r}^b, \tilde{q}, \tilde{\kappa}^f, \tilde{\sigma}^q\}$ using the new $\eta^f_{t+\Delta t}$ and $\eta^i_{t+\Delta t}$ as inputs of the neural network.

[2] This constraint prevents the neural network from producing solutions with negative endogenous volatility when η^f is close to zero.

8. Compute the average loss function given by

$$\mathcal{L} = \frac{1}{NM}\sum_{i=1}^{N}\sum_{j=1}^{M}\left[\left(\tilde{q}^{i,j} - q_{t+\Delta t}^{i,j}\right)^2 + \left(\frac{\tilde{q}^{i,j} - q_{t+\Delta t}^{i,j}}{q_{t+\Delta t}^{i,j}}\right)^2 + \left(\mathrm{e}1^{i,j}\right)^2 + \right.$$
$$\left. + \left(\mathrm{e}2^{i,j}\right)^2 + \left(\mathrm{e}3^{i,j}\right)^2 + \left(\mathrm{e}4^{i,j}\right)^2\right]$$

9. Use $\{\tilde{r}^d, \tilde{r}^b, q_{t+\Delta t}, \tilde{\kappa}^f, \tilde{\sigma}^q\}$ as the new values for the next iteration.

10. After completing the iteration, save weights and bias for the full model.

I.2 Solution

The accuracy of the probabilistic method can be analyzed by comparing the solution of the two-state model when $\nu^i \approx 1$ and $\eta^i \approx 0$. This case converges to the model proposed by Brunnermeier and Sannikov (2016), with only two agents, experts (firms + intermediaries) and households, and no collateral constraint.

In Figures I1 and I2, the dashed gray line represents the NR solution of the one-state model, while the solid red line corresponds to the PRL solution of the limiting case of the two-state model. The PRL approach closely approximates both the capital price function and the share of capital held by firms and intermediaries. Some discrepancies arise around the transition from the crisis regime to the normal regime in the endogenous volatility; nevertheless, the PRL solution still produces policy functions that capture the model's key dynamics. Table I1 reports the accuracy of the probabilistic method, showing that the mean squared error is on the order of 10^{-5} or smaller for all variables.

Table I1: Mean Squared Error

Mean Squared Error				
q	σ^q	κ^e	$\eta\mu^\eta$	$\eta\sigma^\eta$
$1.3 \cdot 10^{-5}$	$2.3 \cdot 10^{-5}$	$2.2 \cdot 10^{-5}$	$1.6 \cdot 10^{-6}$	$7.0 \cdot 10^{-6}$

Figure I1: Probabilistic Reinforcement Learning x Newton-Raphson

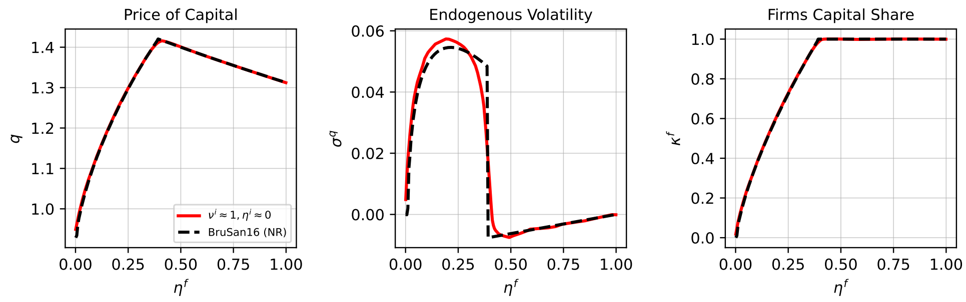

Note: Network setup: hidden layers = [128, 128, 128, 128]; epoch = [300, 300, 300]; learning rate = [1e4, 1e-5, 1e-6]; batch size = [512, 512, 512], T = 0.1, nt = 20, nsample = 50 000.

Figure I2: Probabilistic Reinforcement Learning x Newton-Raphson

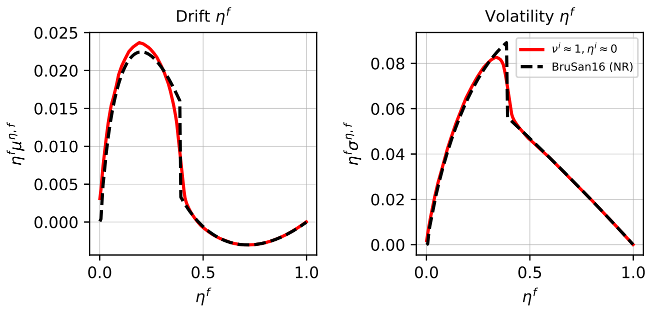

Note: Network setup: hidden layers = [128, 128, 128, 128]; epoch = [300, 300, 300]; learning rate = [1e4, 1e-5, 1e-6]; batch size = [512, 512, 512], T = 0.1, nt = 20, nsample = 50 000.