

NBER WORKING PAPER SERIES

DISSECTING MECHANISMS OF FINANCIAL CRISES:
INTERMEDIATION AND SENTIMENT

Arvind Krishnamurthy
Wenhao Li

Working Paper 27088
<http://www.nber.org/papers/w27088>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2020, Revised February 2024

We thank Matthew Baron, Pedro Bordalo, Eduardo Davila, Maryam Farboodi, Daniel Greenwald, Xavier Gabaix, Paymon Khorrami, Xiaoji Lin, Christopher Hrdlicka, Yueran Ma, Peter Maxted, Guillermo Ordonez, Carolin Pflueger, Alexi Savov, Dejanir Silva, Sanjay R. Singh, Alp Simsek, Laura Veldkamp, Wei Xiong, and seminar participants at USC, Stanford, UW Foster, Cornell, UPenn-Wharton, Virtual Finance Workshop, AFA, AEA, WFA, SFS Cavalcade, IBEFA Summer Meeting, Barcelona Asset Prices, Banco de Portugal, Finance and Macroeconomics Workshop, CICF, China International Conference In Macroeconomics, EFA, CESifo Area Conference on Macro, Money and International Finance, Red Rock Finance Conference, Federal Reserve Board, NY Fed, NBER Asset Pricing, NBER Summer Institute EFCE and Stanford SITE for comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Arvind Krishnamurthy and Wenhao Li. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Dissecting Mechanisms of Financial Crises: Intermediation and Sentiment
Arvind Krishnamurthy and Wenhao Li
NBER Working Paper No. 27088
May 2020, Revised February 2024
JEL No. E7,G01

ABSTRACT

We develop a model of financial crises with both a financial amplification mechanism, via frictional intermediation, and a role for sentiment, via time-varying beliefs about an illiquidity state. We confront the model with data on credit spreads, equity prices, credit, and output across the financial crisis cycle. In particular, we ask the model to match data on the frothy pre-crisis behavior of asset markets and credit, the sharp transition to a crisis where asset values fall, disintermediation occurs and output falls, and the post-crisis period characterized by a slow recovery in output. Our model with the frictional intermediation mechanism and fluctuations in beliefs provides a parsimonious account of the entire crisis cycle. The model with only the frictional intermediation mechanism misses the frothy pre-crisis behavior; fluctuations in beliefs resolve this problem. On the other hand, modeling the belief variation via either a Bayesian or diagnostic model match the broad patterns, with each missing some targets to different extents.

Arvind Krishnamurthy
Stanford Graduate School of Business
Stanford University
655 Knight Way
Stanford, CA 94305
and NBER
akris@stanford.edu

Wenhao Li
University of Southern California
Marshall School of Business
Los Angeles, CA
liwenhao@marshall.usc.edu

1 Introduction

Financial crises have a common character. There is a pre-crisis period that is marked by a runup in credit, leverage, low risk spreads, and an expansion in output. Credit and asset valuations appear “frothy” before a crisis. The transition to the crisis is sharp. There are losses to the financial sector, defaults and bank runs, a jump in risk spreads, and a contraction in credit and output. The aftermath of the crisis is a gradual recovery in credit, output, and fall in risk spreads. These patterns emerge from a large and growing body of research examining financial crisis episodes across countries and time, dating back to the 19th century. See [Bordo et al. \(2001\)](#), [Borio and Lowe \(2002\)](#), [Claessens, Kose and Terrones \(2010\)](#), [Reinhart and Rogoff \(2009a\)](#), [Schularick and Taylor \(2012\)](#), [Jordà, Schularick and Taylor \(2011\)](#), [Laeven and Valencia \(2013\)](#), [Jordà, Schularick and Taylor \(2013\)](#), [Baron and Xiong \(2017\)](#), [Muir \(2017\)](#), [Baron, Verner and Xiong \(2021\)](#), and [Krishnamurthy and Muir \(2024\)](#). This empirical research describes and quantifies these common patterns.

Theoretical research on crises has fallen into two categories. The first emphasizes frictions in financial intermediation that drive an amplification mechanism. The key idea is that the fragility of the financial sector, measured typically as high bank credit growth or low levels of equity capital-to-assets, is an endogenous state variable. An unexpected large-loss event hitting the economy in a state where the financial sector is fragile sets in motion mechanisms whereby the shock is amplified, there is disintermediation, a rise in risk spreads and contraction in output. Recovery takes time, tracking a gradual re-intermediation. The amplification model speaks directly to the transition to crisis and the aftermath of the crisis. See work by [Gertler and Kiyotaki \(2010\)](#), [He and Krishnamurthy \(2013\)](#), [Brunnermeier and Sannikov \(2014\)](#), [He and Krishnamurthy \(2019\)](#), and [Li \(2019\)](#).

The second line of research emphasizes the role of information and beliefs, and harkens back to [Kindelberger \(1978\)](#). There are two key ideas in this research. First, agents experience a period of prosperity and come to believe that risks are low. Second, the crisis is an informational event – a “[Minsky \(1992\)](#) moment” – where risk is reassessed, leading to swings in asset prices, credit, and macroeconomic outcomes. In the work of [Gorton and Ordonez \(2014\)](#) and [Dang, Gorton and Holmström \(2020\)](#), the shift in beliefs occurs because financial sector information is hidden, by design, during prosperous periods, and a crisis is the event when negative information comes to light and agents reassess risks. The shift from no information to information is at the heart of their narrative of crises. The work of [Bordalo, Gennaioli and Shleifer \(2018\)](#) has instead argued that a sharp shift in beliefs in a crisis reflects a change from over-optimistic to over-pessimistic beliefs. Extrapolative expectations are at the heart of their narrative of the belief shift in a crisis.

This paper builds a model that integrates both of these elements, frictional financial intermediation and time-variation in beliefs, into a quantitative macro-finance model. Our

objective is to understand the extent to which these mechanisms can account qualitatively and quantitatively for the macro crisis patterns, and to clarify which elements of these mechanisms are essential. Our model has a financial intermediary sector subject to capital constraints and financed in part by demandable debt. There are two sources of shocks, a Brownian shock to the return on capital and an illiquidity shock where the market for capital assets temporarily freezes up, and debtors refuse to roll over their debts, as in a bank run. In this latter state, sales of capital assets by banks incur a liquidation cost, or alternatively, loans against capital are charged an illiquidity premium. The economy transits through booms and busts driven by the Brownian shock and its impact on the dynamics of real capital and the equity capital of the financial sector. Crises are events where the financial sector intermediates a large fraction of the real capital of the economy, financing this with debt, and the illiquidity shock occurs. In this high bank-credit state, there are runs on banks leading to disintermediation, declines in asset values, and a reduction in output. The illiquidity shock captures a financial panic, such as occurred in both fall 2008 and spring 2020, with differences in macroeconomic outcomes driven in part by differences in financial sector fragility. We also note that our illiquidity shock impacts the economy indirectly via a financial amplification mechanism and not directly via its impact on productivity and output. This approach to modeling leads to endogenous crises in which the financial sector’s fragility is the key factor. The modeling is motivated by our objective to shed light on financial crises such as the 2008 global financial crisis and not on rare consumption disasters such as the 2020 COVID recession. The financial frictions model of our paper is a variant of [Li \(2019\)](#). It draws on ingredients from the recent macro-finance literature on financial crises and intermediation frictions, and particularly [He and Krishnamurthy \(2013\)](#); [Brunnermeier and Sannikov \(2014\)](#); [Gertler and Kiyotaki \(2015\)](#).

Bankers in the economy make decisions based on their beliefs about the likelihood of the illiquidity shock. The illiquidity shock is a Poisson event, the intensity of which follows a hidden two-state Markov process. Agents infer the state and hence the likelihood of the illiquidity shock based on history. A string of no-shock realizations leads them to believe that shocks are unlikely (i.e., the true state is the low-intensity state). A shock occurrence leads them to think that shocks are more likely (i.e., the true state is the high-intensity state). After an extended period with no shocks, banks downplay liquidity risk and increase leverage. The shock triggers a “Minsky moment:” agents’ belief regarding liquidity risk rises and is then amplified and propagated to the macroeconomy depending on the fragility of the financial sector.

We consider two flavors of the learning mechanism, a Bayesian updating process closest to [Moreira and Savov \(2017\)](#) and non-rational updating process, along the lines of [Bordalo, Gennaioli and Shleifer \(2018\)](#), where beliefs over-react to current news.¹

¹The diagnostic updating process is motivated by the work of [Bordalo, Gennaioli and Shleifer \(2018\)](#), and is also related to the models of [Greenwood, Hanson and Jin \(2019\)](#) and [Maxted \(2019\)](#). [Bordalo,](#)

We report three principal results:

1. The model with financial frictions and a time-varying belief process matches the main features of the pre-crisis, crisis, and aftermath.
2. A model with only financial frictions generates the amplification needed to match crisis and post-crisis patterns but fails to match the pre-crisis froth evidence. That is, both a financial friction mechanism and a mechanism involving fluctuations in beliefs are needed to match the crisis cycle evidence.²
3. While belief fluctuations are essential, whether one needs a Bayesian belief process or the non-rational diagnostic process to fit the crisis patterns is more murky. The diagnostic belief model, calibrated to the evidence from [Bordalo, Gennaioli and Shleifer \(2018\)](#), matches the crisis patterns qualitatively. But so does the Bayesian belief model. These two learning variants each fit different dimensions of the data better (and worse), with the diagnostic model’s principal success over the Bayesian model being that the model’s pre-crisis froth is quantitatively closer to the data.

The model has two key state variables: one governing the wealth share of bankers and the other describing agents’ beliefs over the intensity of the illiquidity shock. The wealth-share variable, coupled with financial frictions, governs a financial amplification mechanism studied in prior work. We show that this amplification mechanism helps the model match data on the crisis and its aftermath. In particular, the financial amplification mechanism of the model generates a sharp drop in asset prices, credit, and output. The mean drop in our model is in line with the data, but more telling, the skewness of these variables and their comovement also align with data counterparts. That is, a key feature of financial crises is non-linearity, reflected in a skewed distribution of output declines. The model’s amplification mechanism generates skew in line with that of the data. The model also generates a slow recovery due to the persistence mechanism of financial frictions models.

While the financial frictions wealth-share mechanism is the key to understanding the model’s match of the crisis and aftermath, the belief state variable is needed for the model to match the pre-crisis patterns. If an illiquidity shock has not occurred for some time, agent beliefs drift towards the low likelihood state. Bankers choose to increase leverage as

[Gennaioli and Shleifer \(2018\)](#), [Bordalo et al. \(2019b\)](#), and [Bordalo et al. \(2020\)](#) examine data on survey forecasts of financial and economic variables. They show that these forecasts are hard to square with rational expectations and instead propose a model of diagnostic expectations that matches these data. We use their model and parameterization of diagnostic expectations to study crises. Their survey evidence concerns data that largely varies at business cycle frequencies. We assess how this model of behavior can extend to explaining rare financial crises. In the appendix, we also consider other types of belief distortions, including optimism and pessimism.

²In the appendix, we consider the case of a model without financial frictions but with fluctuations in beliefs over a production disaster. That variation of the model fails to even qualitatively match the association between credit variables and crises.

they are less concerned about liquidity risk. Risk and credit spreads fall and credit grows. From this state, if an illiquidity shock arrives, beliefs jump towards the high likelihood state and banker wealth falls and they endogenously choose less leverage given the heightened liquidity risk. Both forces lead to financial amplification of the shock and persistence as in a crisis. The belief mechanism helps explain why spreads are low and credit is high *before* the crisis. More surprisingly, low spreads and high credit help *predict* a crisis. The reason is that bankers act more risk-tolerant in the pre-crisis period when liquidity risk is low – they drive down spreads/risk premia and increase credit. They also increase leverage, taking liquidity risks, and effectively shift GDP outcomes into tail states. It may be surprising that we find that there are times when crises are more likely and yet risk prices are low and bankers take more leverage. Our model ties these observations together by generating bankers’ willingness to take on illiquidity risk in the pre-crisis period, driven by the beliefs state variable.

We probe this model in two dimensions. First, we find that if the belief intensity is held constant (i.e., no learning mechanism, but illiquidity shocks still occur), the model fails to match the pre-crisis patterns. In such a model, only the banker wealth-share is a state variable. The fragility of the economy to a crisis is summarized by the quantity of bank credit. When this is high, the illiquidity shock triggers a large bank run and triggers a crisis. Thus a crisis is more likely when total bank credit is high. However, this means that forward-looking asset prices will account for the increased fragility as the credit state variable rises. As a result, the model implies that credit spreads will be high in the period before a crisis, contrary to the data. On the other hand, we find that this static-belief model is able to match the data for the crisis and its aftermath, clarifying that the financial amplification mechanism drives these patterns. We also show that this model generates a negative relation between bank credit and equity market excess returns (risk premia), as documented by [Baron and Xiong \(2017\)](#). This occurs in our model because variation in the wealth-share drives variation in bankers’ risk tolerance that generates the required comovement between credit and risk premia. It is worth emphasizing that this result arises in a model with no variation in beliefs.

Second, we compare the Bayesian and diagnostic learning mechanisms. There are differences in the magnitude of the crisis and post-crisis match: the diagnostic model generates too much comovement, relative to data, between bank credit and crisis, for example. But the principal difference is in the amount of pre-crisis froth of these models. The key force in matching the pre-crisis evidence is the sensitivity of the bankers’ endogenous borrowing decision to the true (not perceived) illiquidity state. This sensitivity has the required sign (negative) in both learning mechanisms, but the sensitivity is higher under diagnostic beliefs. In our calibrated models, the Bayesian mechanism gets about two-thirds of the way to matching the data quantitatively, while the diagnostic mechanism brings the froth in the model gets even closer to the data. In the Appendix, we consider other variants of belief

distortions including optimism and pessimism, and show that these models also do better than the static belief model, but not demonstrably better than the Bayesian or diagnostic belief models.

Putting this together, our analysis indicates that a financial amplification mechanism plus a belief mechanism provides a parsimonious account of the main crisis facts. The static belief model fails to match the pre-crisis froth in the data. However, our analysis also indicates that the qualitative patterns of the data do not clearly distinguish the two belief models we consider. Within the bounds of how much one can push the quantitative fit of our parsimonious equilibrium model, both models fit the moments we consider.

The main contribution of this paper is to bridge the recent theoretical work on non-linear macro-finance models (He and Krishnamurthy, 2013; Brunnermeier and Sannikov, 2014; Di Tella, 2017) with the empirical literature on financial crises cited earlier. The models in this theoretical literature feature non-linearities and are solved using global methods. Thus these models are well-suited to characterize the non-linear dynamics in financial crises. But the work thus far is either purely theoretical or aims to match a single crisis-event (e.g., the 2007-2009 financial crisis in He and Krishnamurthy (2019) and Gertler, Kiyotaki and Prestipino (2020)). The empirical crisis literature on the other hand has largely documented systematic patterns in the data rather than assess this data from the standpoint of models. Our paper bridges this gap.³

This paper's objective of matching the boom-bust of the crisis cycle is closest to that of a few other papers that precede ours.⁴ Boissay, Collard and Smets (2016) develop a dynamic model of banking crises that generates the pattern in line with the data that credit booms precede credit market collapse and crises. The key idea in the model is that banks' absorption capacity is reduced during a boom, with the economy potentially hitting a cliff where the credit market collapses. Thus, the probability of a financial crisis rises in the boom. Relative to their analysis, we aim to match the asset market fact that risk premia and credit spreads are low during the boom, which we reconcile with the learning mechanism of our model. Greenwood, Hanson and Jin (2019) and Maxted (2019) construct models

³Another contribution of our paper to the non-linear macro finance work is the model. A major disadvantage of the current models is that they are computationally challenging, and current models restrict attention to one or two-state variables following a Brownian diffusion process. In this paper, we present and solve a model with two state variables and endogenous jumps. Our methodology helps broaden the scope of the literature to encompass richer dynamics with sudden and large disruptions, which are plausibly central to financial crises.

⁴There are other recent macro-finance papers, not explicitly about the boom-bust cycle, but that aim to match crisis facts. Gertler, Kiyotaki and Prestipino (2020) introduces bank runs into a macro-intermediation model. Beliefs, modeled via a sunspot, play a role in driving crisis dynamics. The objective of their paper is to study the 2007-2009 financial crisis rather than disentangling mechanisms underlying the crisis cycle facts. Camous and Van der Ghote (2021) builds on Maxted (2019) and considers diagnostic expectations and financial frictions in a multi-sector model. The model can generate a build-up of instability and a safety trap with low growth. Gopalakrishna (2020) introduces state-dependent bank exit into a quantitative continuous-time macro-finance model and generates a slow recovery in line with empirical evidence.

of the boom-bust crisis cycle with a role for beliefs. [Greenwood, Hanson and Jin \(2019\)](#) present a model where lenders extend credit based on beliefs over the default probabilities of borrowers. There is a feedback between realized default and beliefs regarding default probabilities, similar to the model of this paper, that creates a persistence and amplification mechanism. Like us, their paper aims to match facts on credit growth, credit spreads, and risk premia. But their model is not a full macroeconomic model, and thus does not speak to other macroeconomic data such as output and the conditional distribution of output growth. Their model also does not have an intermediary sector, so it cannot assess the role of intermediary frictions relative to beliefs. Finally, lenders are risk-neutral in their model, so that without diagnostic expectations, risk premia are zero. As a result, their model does not give the Bayesian belief process a chance of explaining the data. [Maxted \(2019\)](#)'s macro-finance model is closer to ours. There is an intermediation sector that is central to crisis dynamics. The paper also considers a full macroeconomic setting, and can thus speak to more macro data. One key point of difference relative to our model is that [Maxted \(2019\)](#)'s diagnostic belief modeling extrapolates the mean growth of capital productivity, whereas in our model, beliefs over a tail illiquidity shock are distorted.⁵ Thus in [Maxted \(2019\)](#), optimism occurs after a period of growth in productivity which can sow the seeds of the crisis, while in our mechanism optimism occurs after a quiet period of no illiquidity shocks, leading endogenously to increased risk-taking. There is evidence that productivity booms are good booms that do not lead to crises ([Gorton and Ordóñez, 2020](#)), and certainly the period before the 2008 financial crisis was marked by a slow-down in productivity growth ([Fernald, 2015](#)).⁶

Finally, this paper also contributes to a larger literature on beliefs and learning in macroeconomics models. Closest to our paper is [Kaplan, Mitman and Violante \(2020\)](#) who dissect the U.S. housing boom-bust cycle around the 2008 crisis to evaluate the role of beliefs and financial constraints in driving the cycle. They conclude that a shift in beliefs during the boom is essential to matching the cycle. Note that we consider banking crises and not housing crises, and broaden our scope to include patterns across many crisis episodes. [Van Nieuwerburgh and Veldkamp \(2006\)](#) show that asymmetry in learning about productivity can generate asymmetries in business cycles. [Simsek \(2013\)](#) explores the interaction of beliefs and credit, building a model where beliefs over upside versus downside payoffs have an asymmetric impact on asset valuations, total credit and fragility of the economy. [Simsek \(2013\)](#) studies the role of belief heterogeneity, which is absent in our model with homogeneous beliefs. Motivated by the slow recovery from the 2008 recession, there is re-

⁵[Ma, Paligorova and Peydro \(2021\)](#) presents survey evidence that assessments by banks over the downside tail, more than the mean of economic forecasts, explain bank lending decisions.

⁶[Farboodi and Kondor \(2020\)](#) present a model of time-varying sentiment that generates a credit cycle that is qualitatively in line with the facts. Sentiment evolves in a Bayesian manner in their model. Thus, like us, they show that the basic facts of the credit cycle can be generated within a Bayesian model. The objective of the paper is different than ours, as their model is not suited to a quantification exercise and does not have an intermediary sector.

search tying learning to slow recoveries. In [Fajgelbaum, Schaal and Taschereau-Dumouchel \(2017\)](#), information flows slowly in times of low activity and uncertainty remains high, discouraging investment. [Liu, Wang and Yang \(2020\)](#) show that the uncertainty and learning about banks' peers can lead to a slow recovery. In [Kozlowski, Veldkamp and Venkateswaran \(2020\)](#), agents learn about the parameters of the economic shock process, and a large negative shock realization as in a deep recession alters agents' estimates of these parameters, leading to a persistent impact of the shock on economic growth. [Bordalo et al. \(2019a\)](#) introduce diagnostic beliefs into a relatively standard real business cycle model. Their model helps to understand the role of diagnostic beliefs in driving business cycles.

The rest of this paper is as follows. In Section 2, we review general patterns of the crisis cycle in the data. In Section 3, we set up a model that combines financial intermediation frictions and beliefs regarding an illiquidity shock. In Section 4, we solve and explain how we calibrate the the model(s). In Section 5, we evaluate the model, explaining its fit and the role of beliefs. We then conclude in Section 6. An appendix follows.

2 The Crisis Cycle

This section reviews broad patterns of the crisis cycle, drawn from the empirical literature on crises. Along the way, we list (numbered below) specific quantitative estimates from the literature which guide our modeling exercise.

What is a financial crisis? [Jordà, Schularick and Taylor \(2011\)](#) state:

In line with the previous studies, we define financial crises as events during which a country's banking sector experiences bank runs, sharp increases in default rates accompanied by large losses of capital that result in public intervention, bankruptcy, or forced merger of financial institutions.

We focus on events, as per the quotation, as financial crises. These events are banking crises and do not necessarily include currency crises or sovereign debt crises unless such events coincide with a banking crisis. [Jordà, Schularick and Taylor \(2011\)](#)'s dating of banking crises is closely related to the approach of [Bordo et al. \(2001\)](#), [Reinhart and Rogoff \(2009a\)](#), and [Laeven and Valencia \(2013\)](#). [Bordo and Meissner \(2016\)](#) discuss the approaches that researchers have taken to crisis-dating as well the drawbacks of different approaches.

1. We target an unconditional annual frequency of financial crises of 4%. In an article written for the Annual Review of Economics, [Taylor \(2015\)](#) reports the historical frequency of financial crises to be 6%. This data point is obtained from a sample of countries in both developing and advanced stages, and covers the period after 1860. The Handbook of

Macroeconomics chapter by [Bordo and Meissner \(2016\)](#) reports numbers in the range of 2 to 4% across the studies by [Bordo et al. \(2001\)](#) and [Reinhart and Rogoff \(2009a\)](#). Another evidence comes from [Jordà, Schularick and Taylor \(2013\)](#), which reports that the average frequency of crises is 3.6% using data from multiple countries. [Sufi and Taylor \(2022\)](#) report an unconditional probability of crises of 2.5%.

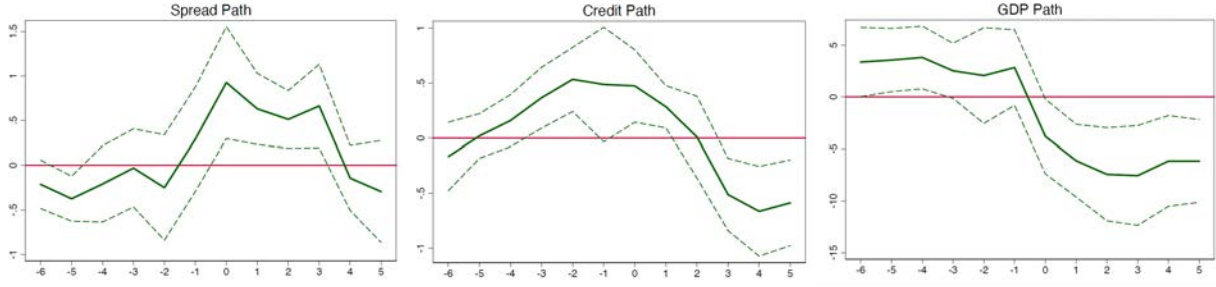


Figure 1: Mean path of credit spread, bank credit, and GDP across a sample of 40 financial crises. Units for spread path are 0.5 means that spreads are 50% above their average for a given country. Units for credit path are that 0.5 indicates that credit/GDP is 0.5σ s above the trend for a given country. Units for GDP path are such that -5 means that GDP is 5% below trend for a given country. Source: Figure 1 of [Krishnamurthy and Muir \(2024\)](#)

Figure 1 plots the mean path of credit spread, credit, and GDP across a sample of 40 international financial crises identified by [Jordà, Schularick and Taylor \(2013\)](#) and [Laeven and Valencia \(2013\)](#) and studied by [Krishnamurthy and Muir \(2024\)](#). Date 0 on the figure corresponds to the date of a financial crisis. The top-left panel plots the path of the mean credit spread, relative to the mean spread for country- i and other countries at time t , from 6-years before the crisis to 5-years after the crisis. The units here are that 0.4 means that spreads are 40% (approximately 0.40σ s) larger than the country's time-series average spread, while -0.2 means that spreads are 0.2σ s below the country's time-series average. The data is annual from 15 countries spanning a period from 1879 to 2013.

We see that spreads run below their average value in the years before the crisis. They rise in the crisis by 0.7σ , going as high as 1σ s over their mean value in the year after the crisis date, before returning over the next 5 years to the mean value. The half-life of the credit spread recovery is 2.5 years in this figure.

The middle panel plots the path of the quantity of bank credit divided by GDP. The credit variable is expressed as the average across-country percentage change in the quantity of credit/GDP from 6-years before the crisis to a given year, after demeaning by the sample growth rate in credit for country- i and normalizing by the standard deviation of credit growth for the country. The value of 0.5 for time 0 means that credit/GDP is 0.5σ s above the country trend. We see that credit grows faster than average in the years leading up to the crisis at time zero. After this point, credit reverses so that by time +5 the variable is substantially below the country average.

The right panel plots GDP, again as an average percentage change from 6-years before the crisis, after demeaning by the sample growth rate in GDP for country- i . GDP is above average trend in the years preceding the crisis. GDP falls below trend in the crisis and remains low up to 5 years after the crisis.

Transition to crisis: A crisis is characterized by a sharp jump in credit spreads, a reversal in the quantity of credit and a decline in GDP. From the data underlying Figure 1:

2. Credit spreads rise by $0.7\sigma_s$ in the first year of the crisis.
3. GDP declines by 9.1%. [Reinhart and Rogoff \(2009b\)](#) report a peak-to-trough decline in GDP across a larger sample of crises of 9.3%. [Jordà, Schularick and Taylor \(2013\)](#) report a 5-year decline in GDP from the date of crisis of around 8%. [Cerra and Saxena \(2008\)](#) report output losses from banking crises of 7.5% with these losses persisting out to 10 years. We will use the 9.1% number in our quantitative exercise.

The rise in credit spreads in the year of the crisis is mirrored in other asset prices. [Reinhart and Rogoff \(2009a\)](#) report that equity prices decline by an average of 55.9% during banking crises. [Muir \(2017\)](#) shows that the dividend-price ratio on the stock market and the excess return on stocks rises during the crisis, indicated a generalized rise in asset market risk premia. [Muir \(2017\)](#) reports that the peak dividend-price ratio in a crisis rises by 43% relative to its pre-crisis level.

Aftermath and severity of crisis:

4. The half-life of the recovery of the credit spread to its mean value is 2.5 years.
5. There is variation in the severity of the crisis. Figure 2 presents data on the variation in the severity of the crisis, as measured by peak-to-trough GDP growth around a financial crisis. The figure reflects significant variation in crisis severity. The data is from a sample of 42 financial crises studied by [Krishnamurthy and Muir \(2024\)](#).
6. The variation in the severity of the crisis is correlated with the increase in spreads measured at the transition into the crisis. [Krishnamurthy and Muir \(2024\)](#) report a coefficient of -2.11 (*s.e.* 0.51) from a regression of 3-year GDP growth following a crisis on the increase in credit spreads from the year before the crisis to the year of a crisis.
7. The variation in the severity of the crisis is correlated with the pre-crisis increased in credit/GDP. [Krishnamurthy and Muir \(2024\)](#) report that one-sigma increase in bank credit/GDP over the 5 years preceding the crisis is correlated with a 2.06% decline in 3-year GDP growth following the crisis. [Mian, Sufi and Verner \(2017\)](#) show that a one-sigma increase in private debt-to-GDP growth over the last 3 years is associated with a 2.1% decline in output over the next 3 years.

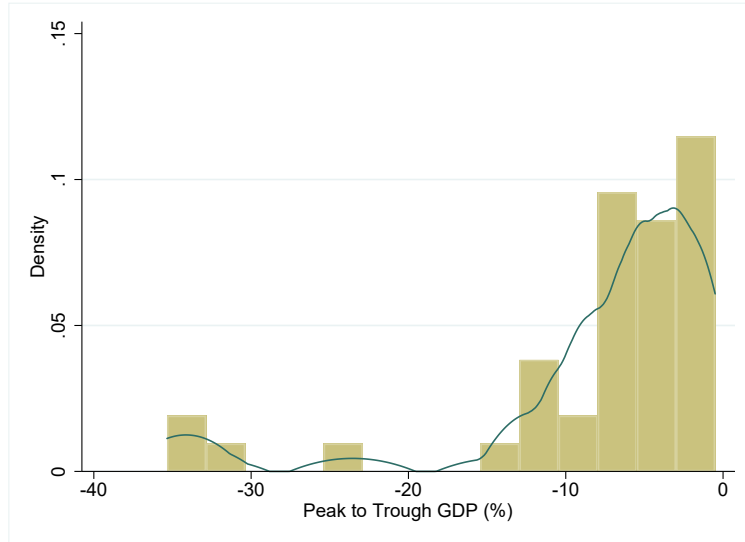


Figure 2: The figure presents a histogram of the peak-to-trough fall in GDP around 42 financial crises studied by [Krishnamurthy and Muir \(2024\)](#).

Pre-crisis period: In the pre-crisis period, credit markets appear frothy, reflecting low credit spreads and high credit growth. In particular,

8. Conditioning on a crisis at year t , and looking at the 6 years prior to the crisis, [Krishnamurthy and Muir \(2024\)](#) show that credit spreads are 0.45σ s below their country mean (where this country mean is defined to exclude the crisis and 5 years after the crisis).
9. Conditioning on a crisis at year t , credit/GDP in the 5 years before the crisis is 5% above country mean (or 0.5σ s above country mean). The relation between a lending boom and subsequent crisis is well documented in the literature. See [Gourinchas et al. \(2001\)](#), [Schularick and Taylor \(2012\)](#), and [Baron and Xiong \(2017\)](#).

Predicting Crises: There is also evidence that periods of frothy conditions predict and not just precede crises. There are three quantitative estimates that we will aim to match.⁷

10. [Schularick and Taylor \(2012\)](#) find that a one-standard deviation increase in credit growth over the preceding 5 years (= 0.07 in their sample) translates to an increased probability of a financial crisis of 2.11% over the next year.⁸ [Sufi and Taylor \(2022\)](#) report that when previous 5-year credit growth is in its highest sextile, the probability of a crisis quadruples.

⁷[Greenwood et al. \(2020\)](#) present further evidence in line with froth predicting crises. In post-war cross-country data, they document that periods of high credit growth coupled with periods of high returns in the stock market substantially increase the likelihood of a financial crisis. We discuss how our model can speak to their quantitative estimates in Section 5.7.

⁸This comes from Table 2 of their paper which is a Logit specification, and we report the marginal probability from that specification.

11. Conditioning on an episode where credit spreads are below their median value 5 years in a row, [Krishnamurthy and Muir \(2024\)](#) estimate that the conditional probability of a crisis rises by 12.9% over the next 3 years.
12. [Baron and Xiong \(2017\)](#) find that a one-standard deviation increase in credit growth over the preceding 3-years increases the probability of bank equity crash, defined as decline in bank equity by over 30% in the next year, by 5.4%.

3 A Model of Financial Crises with Amplification and Sentiment

In this section, we present a model of financial crises that incorporates both a financial amplification mechanism and a role for sentiment. We fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and assume all stochastic processes are adapted to this space and satisfy the usual conditions. The economy evolves in continuous time. It is populated by a continuum of a unit mass of two classes of agents, households, and bankers. For clarity, aggregate variables are in capital letters, and individual variables are in lower case letters. The basic setup is a variant of [Li \(2019\)](#), which is drawn from [Brunnermeier and Sannikov \(2014\)](#) and [Kiyotaki and Moore \(1997\)](#).

3.1 Agents and Assets

Households maximize expected value of the discounted log utility,

$$\int_0^{\infty} e^{-\rho t} \log(c_t^h) dt \tag{1}$$

and bankers optimize expected value of the same form of discounted log utility,

$$\int_0^{\infty} e^{-\rho t} \log(c_t^b) dt \tag{2}$$

The expectation could be either Bayesian or diagnostic, as we will specify later.

Output is produced by capital. We will simplify by assuming that the capital is held directly by either banks or households. In a richer and more realistic model, the capital will be held and operated by firms that receive loans from banks or households, along the lines of [Holmstrom and Tirole \(1997\)](#). We simplify by collapsing firms into banks, and assuming the banks own the capital.

We assume that credit flowing through banks allows the economy to achieve higher

output and returns to capital. Intermediation is a socially valuable service, and for example, disintermediation in a crisis reduces output. We capture this feature by assuming that banker-operated capital has productivity \bar{A} , which is higher than the household-operated capital productivity of \underline{A} .

The dynamic evolution of productive capital owned by agent $j \in \{\text{banker, household}\}$ is

$$\frac{dk_{j,t}}{k_{j,t}} = \mu_t^K dt - \delta dt + \sigma^K dB_t \quad (3)$$

where the rate of new capital installation μ_t^K is endogenously determined through investment, δ is the exogenous depreciation rate, and σ^K is exogenous capital growth volatility.

Denote the price of productive capital as p_t (i.e., “ q ” in the standard Q-theory). Investment undertaken by an owner, either banker or household, of productive capital is chosen to solve:

$$\max_{\mu_t^K} p_t \mu_t^K - \phi(\mu_t^K),$$

where $\phi(\cdot)$ is an investment adjustment cost:

$$\phi(\mu^K) = \mu^K + \frac{\chi}{2}(\mu^K - \delta)^2. \quad (4)$$

That is, we assume quadratic costs to investment, leading to the q -theory of investment

$$p_t = \phi'(\mu_t^K) \quad \Rightarrow \quad \mu_t^K = \delta + \frac{p_t - 1}{\chi}. \quad (5)$$

The dynamics of capital price p_t is denoted as

$$\frac{dp_t}{p_{t-}} = \mu_t^p dt + \sigma_t^p dB_t - \kappa_{t-}^p dN_t, \quad (6)$$

where μ_t^p , σ_t^p , and κ_{t-}^p are all endogenously determined. The “minus” notation (i.e. p_{t-}) reflects a pre-jump asset price, as will be made clear.

3.2 Financing, Liquidity Risk, and Bank Runs

The Brownian shock dB_t in equation (3) reflects business cycle fluctuations in the effective productivity of capital. We introduce a second shock that we call a “financial illiquidity” shock. We model this as a Poisson shock dN_t that triggers illiquidity and bank runs, and a possible financial crisis if the endogenously determined banking sector fragility is sufficiently high.

Since banker held capital is more productive than household held capital, there is room for an intermediation relationship whereby households provide some funds to bankers to

invest in capital. We assume that the only form of financing is short-term (instantaneous) debt at the [endogenous] interest rate r_t^d . Bankers cannot raise equity, long-term debt, or other forms of financing. When we refer to bank equity, we mean the net-worth of bankers, w_t^b . That is, the financing side of the model is one of inside equity and outside short-term debt. These model simplifications do sweep aside important issues, but we nevertheless go down this path because we aim to build a simple quantitative amplification mechanism and see how well it matches data, rather than explore the micro-foundations of intermediary models.

We assume that in the event of an illiquidity shock, all short-term debt holders run to their own bank and withdraw financing in a coordinated fashion. Raising resources to cover this withdrawal is temporarily costly. That is, asset markets are temporarily illiquid in the illiquidity event. We assume that a cost of α is incurred when capital is liquidated to meet the funding withdrawal during the illiquidity shock. We can think of this cost as liquidation cost or, alternatively, the cost can be mapped into a premium on raising emergency financing from other banks or other households in the economy against the capital. In this latter case, we need to step outside the modeling and interpret the illiquidity event lasting longer than dt . Then, α is proportional to the spread over the riskless rate that the bank pays to obtain funds over the illiquidity episode (if the event lasts dt then a financing spread maps into a cost of order dt). Finally, we assume that the cost is not dissipated but is paid to households proportional to their wealth. This assumption is not essential to the analysis but ensures that the illiquidity shock has no direct impact on output.

In terms of primitives, the occurrence of the liquidity shock leads to a transfer of wealth between bankers and households, where the quantity of the wealth transfer is endogenous to the state of the economy, as we explain in equation (31). There are other financial shocks the literature has explored, such as a tightening of financial constraints ([Jermann and Quadrini, 2012](#)), which can deliver similar dynamics.

The illiquidity shock captures a financial panic, such as occurred in both fall 2008 and spring 2020, with differences in macroeconomic outcomes driven by differences in financial sector fragility. We also note that our illiquidity shock impacts the economy indirectly via a financial amplification mechanism and not directly via its impact on productivity and output as would arise in a rare consumption disasters model ([Barro, 2006](#)). In [Appendix A.9](#) we consider a version of the model with a productivity disaster, but no financial frictions driven amplification mechanism.

Note also that we do not model a [Diamond and Dybvig \(1983\)](#) bank-run game. We simply assume that the shock leads all debtors to pull their funding. It is possible to model the game in detail following [Li \(2019\)](#) whose model is the basis for this paper. However, we learn from that study that the model's positive implications are almost the same with and without the deeper model of the bank-run game.

3.3 Beliefs and Crises

The intensity of the illiquidity shock process dN_t follows a two state continuous-time Markov process, $\tilde{\lambda}_t \in \{\lambda_L, \lambda_H\}$. This intensity changes from λ_L to λ_H at rate $\lambda_{L \rightarrow H}$, and changes from λ_H to λ_L at rate $\lambda_{H \rightarrow L}$. Agents, neither bankers nor households, observe $\tilde{\lambda}_t$. Instead agents infer $\tilde{\lambda}_t$ from observing the history of N_t , i.e., via realizations of the shock process.

We denote the Bayesian expectation as $\lambda_t = E_t[\tilde{\lambda}_t]$. Using Bayes rule,

Lemma 1 (Bayesian Belief Process).

$$d\lambda_t = \begin{pmatrix} (\lambda_L - \lambda_{t-})\lambda_{H \rightarrow L} + (\lambda_H - \lambda_{t-})\lambda_{L \rightarrow H} \\ -(\lambda_{t-} - \lambda_L)(\lambda_H - \lambda_{t-}) \end{pmatrix} dt + \frac{(\lambda_{t-} - \lambda_L)(\lambda_H - \lambda_{t-})}{\lambda_{t-}} dN_t \quad (7)$$

Therefore, if illiquidity occurs, the expected intensity λ_t jumps up. As time goes by, without further illiquidity shocks, the expected intensity λ_t gradually falls.⁹

3.4 Diagnostic Expectations

We also consider a version of our model where agents overweight recent observations motivated by the diagnostic belief model of (Bordalo, Gennaioli and Shleifer, 2018). We adapt their model to our continuous time dynamic equilibrium environment.

Denote the Bayesian belief for the probability of $\tilde{\lambda}_t = \lambda_H$ as π_t , and the diagnostic belief for the probability of $\tilde{\lambda}_t = \lambda_H$ as π_t^θ . Then we define the diagnostic beliefs as

$$\pi_t^\theta = \pi_t \cdot \left(\frac{\pi_t}{E_{t-t_0}[\pi_t]} \right)^\theta \frac{1}{Z_t} \quad (8)$$

$$1 - \pi_t^\theta = (1 - \pi_t) \cdot \left(\frac{1 - \pi_t}{E_{t-t_0}[1 - \pi_t]} \right)^\theta \frac{1}{Z_t} \quad (9)$$

where Z_t is a normalization to ensure that (8) and (9) add up to 1. We call the lag t_0 as the “look-back period,” which is one in the discrete time model of Bordalo, Gennaioli and Shleifer (2018). In our case, the diagnostic beliefs of the process are simply distorted Bayesian beliefs with the benchmark from t_0 time ago. The process π_t^θ features both overreaction and underreaction, depending on the gap between current π_t and past π_{t-t_0} .

⁹In theory, when $\lambda_t \rightarrow \lambda_L$, the drift of $d\lambda_t$ can be positive. The reason is that the underlying intensity process $\tilde{\lambda}_t$ switches between λ_L and λ_H and the average is between the two values, so when λ_t is close to λ_L the dynamics of $\tilde{\lambda}_t$ dominates the information effect and $d\lambda_t$ is positive. However, states with positive $d\lambda_t$ are transient, i.e., λ_t never get back to those states once it drifts outside. In the long run, those states are reached with zero probability and do not matter quantitatively.

Denote the diagnostic belief for the expected intensity of illiquidity shocks as

$$\lambda_t^\theta = E_t^\theta[\tilde{\lambda}_t] \triangleq \pi_t^\theta \lambda_H + (1 - \pi_t^\theta) \lambda_L$$

where E^θ is the expectation with respect to the probability distribution under the diagnostic belief. Then we have the following result:

Lemma 2 (Diagnostic Belief Process). *The diagnostic belief $\lambda_t^\theta = E_t^\theta[\tilde{\lambda}_t]$ is*

$$\lambda_t^\theta = \lambda_L + (\lambda_t - \lambda_L) \frac{(\lambda_H - \lambda_t) + (\lambda_t - \lambda_L)}{\left(\frac{\lambda_t^T - \lambda_L}{\lambda_H - \lambda_t^T} / \frac{\lambda_t - \lambda_L}{\lambda_H - \lambda_t}\right)^\theta (\lambda_H - \lambda_t) + (\lambda_t - \lambda_L)} \quad (10)$$

where $\lambda_t^T = E_{t-T}[\tilde{\lambda}_t]$ is the expected value of $\tilde{\lambda}_t$ under the Bayesian expectation.

In Figure 3(a), we plot the evolution dynamics of the Bayesian and diagnostic belief processes, where the diagnostic belief process is described by (10). We note that when $\theta = 0$, $\lambda_t^\theta = \lambda_t$ so that the diagnostic belief is the same as the Bayesian belief. When θ is above 0, the pre-illiquidity shock belief is lower than the Bayesian belief, and then jumps to a higher level after an illiquidity shock. Right after an illiquidity shock, there is over pessimism. However, after one year, the perceived frequency of the illiquidity shock is below the Bayesian belief so that diagnostic agents are overly optimistic.

In Figure 3(b), we plot the belief λ_t^θ against the underlying state λ_t . In the Bayesian case, this is the 45 degree line. In the diagnostic case, we can see that the diagnostic agent belief is optimistic relative to the true state when λ is low, and pessimistic when λ is high. The diagnostic map is not one-to-one because the diagnostic belief process also depends on the look-back period t_0 (although the dependence is small relative to the dependence on the current state λ_t). In the graph we present the average relationship by fitting a spline to the simulated data. The parameters used in this graph are as described in the calibration section.

In the appendix, we study and present results for other forms of belief distortion, parameterizing the distortion following the logic underlying Figure 3(b). That is, we consider a case of optimism where the agent belief is distorted primarily for the low λ states, as well as a pessimism case where the distortion is for the high λ states. More broadly, our modeling approach is flexible enough to consider other forms of belief distortion.

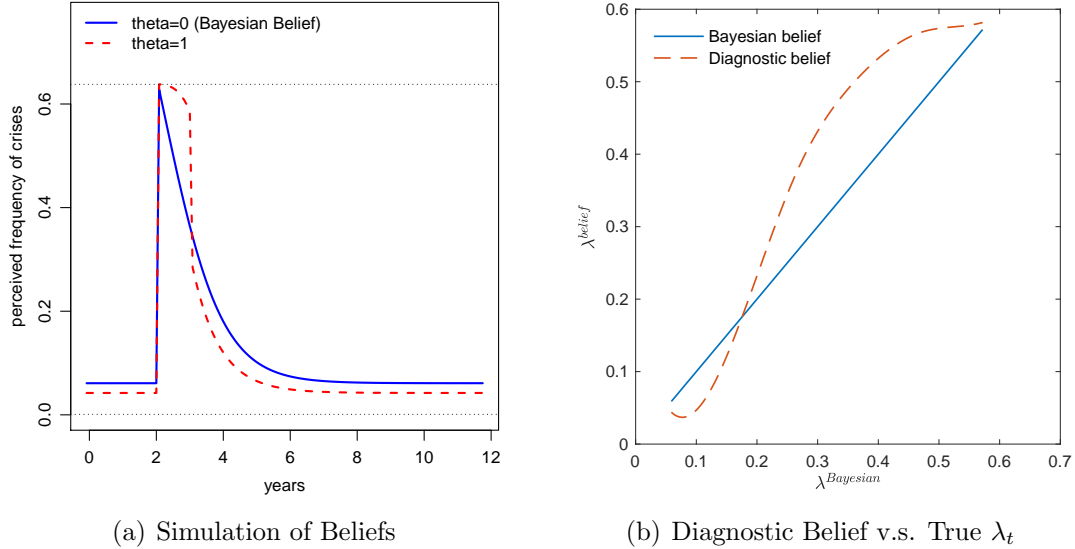


Figure 3: Illustration of the Diagnostic Belief. In panel (a), we simulate the diagnostic belief under different values of θ . The parameter $\theta \geq 0$ means the strength of the behavioral feature of the diagnostic belief, and it becomes the Bayesian belief when $\theta = 0$. Other parameters are set the same as the calibrated diagnostic belief model. The diagnostic belief process is fully described by (10). As θ increases from 0 to 1, the believed frequency of illiquidity shocks in a pre-crisis boom decreases by 30%. In panel (b), we illustrate the average relationship between current Bayesian belief and current diagnostic belief, by fitting a best spline on the simulated data (the mapping is not one-to-one because diagnostic belief also depends on the history of λ_t).

3.5 State Variables and Decisions

We define the total wealth of banks as W_t^b and the total wealth of households as W_t^h . Then we have three state variables. One is the wealth share of bankers, denoted by

$$w_t^b = \frac{W_t^b}{W_t^b + W_t^h}. \quad (11)$$

The second is the expected jump intensity λ_t . The final one is the total productive capital K_t . We construct an equilibrium where all relevant objects scale linearly with K_t . This reduces the computational problem to solving a model with two state variables, w_t and λ_t .

Denote w_t^b as the wealth of a representative banker. Similarly, denote w_t^h as the wealth of a representative household. Let the associated value function be $V^b(w_t^b, w_t, \lambda_t)$ and $V^h(w_t^h, w_t, \lambda_t)$, respectively, at time t . To guarantee a non-degenerate wealth distribution, we assume bankers randomly transit to becoming households at rate η .¹⁰ Bankers take this transition possibility into account in their optimization problems.

In solving their decision problems, agents compute the current state and expectations

¹⁰Without this assumption, the banker, who earns a higher return on capital, will come to own almost all of the wealth of the economy.

over the evolution of the state. In the Bayesian model, these expectations are rational so that the forecast evolution of the state corresponds to the true evolution of the state. In the diagnostic case, the agent's assessment of the current state λ_t^θ is formed via the diagnostic heuristic. But, importantly, the agent does not understand that he is non-rational, nor does he know that any other agent is non-rational, and this is common knowledge.¹¹ As a result, when computing the expectation E_t^θ in the agent's decision problem, the agent forecasts that the law of motion for λ_t^θ is given by the Bayesian (rational) law of motion in (7). Consequently, the policy rules for the agent (consumption and portfolio choices) are mathematically the same function of the state as in the Bayesian case, except that the belief is now the diagnostic belief. When we simulate the model we keep track of both λ^θ and λ . In what follows, we will discuss the policy rules and model solution under the Bayesian belief. The diagnostic model solution comes from replacing λ with λ^θ in the policy functions.

Bankers

Each banker can invest in productive capital and borrow/lend from households or other banks via short-term debt at interest rate r_t^d . Note that short-term debt is riskless even though the price of capital will jump in equilibrium. This is because a forward-looking banker with log utility will never make a portfolio choice that leaves him with negative wealth in any state.

Denote the banker's portfolio choice (as a fraction of the banker's wealth w_t^b) in productive capital as x_t^K . Then the borrowing from household is

$$x_t^d = x_t^K - 1 \quad (12)$$

We will later show that banks always borrow from households and take leverage so we always have $x_t^d > 0$.

Starting from time t , the time that banker will switch to becoming a household is denoted as T , which is exponentially distributed with rate η . A banker with wealth w_t^b solves the problem

$$V^b(w_t^b, w_t, \lambda_t) = \sup_{c_t^b \geq 0, x_t^K \geq 0} E \left[\int_t^T e^{-\rho(s-t)} \log(c_s^b) ds + e^{-\rho T} V^h(w_T^b, w_T) \mid w_t^b, w_t \right], \quad (13)$$

subject to the solvency constraint

$$w_t^b \geq 0. \quad (14)$$

¹¹One could consider a model where a single diagnostic agent thinks they are unbiased but other agents are biased. This is in the spirit of the heterogenous belief models (Simsek, 2013). We conjecture that in such a model the equilibrium impact of belief distortions will be weakened (e.g., bankers will take less leverage if they observe that other bankers are over-levered).

The second part of the objective function is the transition to a household, which changes the continuation value from V^b to V^h . The dynamic bank budget constraint is:

$$\frac{dw_t^b}{w_{t-}^b} = \underbrace{x_{t-}^K (\mu_{t-}^R + \frac{\bar{A}}{p_{t-}}) dt}_{\text{return from capital}} - \underbrace{x_{t-}^d r_{t-}^d dt}_{\text{deposit funding}} - \underbrace{\frac{c_{t-}^b}{w_{t-}^b} dt}_{\text{consumption}} + \underbrace{x_{t-}^K (\sigma^K + \sigma_{t-}^p) dB_t}_{\text{volatility of capital}} - \underbrace{(x_{t-}^K \kappa_{t-}^p + \alpha (x_{t-}^d)^+) dN_t}_{\text{losses in illiquidity shock}},$$

where $(x^d)^+ = \max\{x^d, 0\}$ measures the net borrowing from households, and the “non-dividend” component of capital return is:

$$\mu_t^R = \underbrace{\mu_t^p}_{\text{price appreciation}} - \underbrace{\delta}_{\text{depreciation}} + \underbrace{\sigma^K \sigma_t^p}_{\text{Ito term}} + \underbrace{\mu_t^K - \frac{\phi(\mu_t^K)}{p_t}}_{\text{net investment return}}, \quad (15)$$

In equation (15), the bank obtains returns from capital investment and pays the funding costs to depositors and dividends (i.e., banker consumption) to bank shareholders, subject to the Brownian risks of capital volatility, and losses caused by the liquidity shocks. The return from capital can be classified into a dividend component denoted by \bar{A}/p_t , and a non-dividend component denoted by μ_t^R , which as shown in equation (15) consists of capital price appreciation, capital depreciation, the Ito term on capital volume and price volatility, and finally the net investment returns. During a liquidity shock, the banker suffers from both an exogenous liquidation cost (α) and a valuation drop on their capital holdings (κ_{t-}^p). Note that the net funding withdrawal that has to be fulfilled during an illiquidity episode by selling productive capital is x_{t-}^d .

Households

Each household chooses the consumption rate c_t^h and capital holding y_t^K as a fraction of household wealth for the following objective

$$V^h(w_t^h, w_t, \lambda_t) = \sup_{c_t^h \geq 0, y_t^K \geq 0} E \left[\int_t^\infty e^{-\rho(s-t)} \ln(c_s^h) ds \mid w_t^h, w_t \right], \quad (16)$$

subject to the solvency constraint

$$w_t^h \geq 0, \quad (17)$$

and the budget constraint

$$\frac{dw_t^h}{w_{t-}^h} = \underbrace{y_{t-}^K (\mu_{t-}^R + \frac{A}{p_{t-}}) dt}_{\text{return from capital}} + \underbrace{y_{t-}^d r_{t-}^d dt}_{\text{deposit interest}} - \underbrace{\frac{c_{t-}^h}{w_{t-}^h} dt}_{\text{consumption}} + \underbrace{y_{t-}^K (\sigma^K + \sigma_{t-}^p) dB_t}_{\text{volatility of capital}} - \underbrace{\kappa_{t-}^h dN_t}_{\text{liquidity exposure}} \quad (18)$$

where in the liquidity shock, they also suffer losses on their holdings of capital, but receive a transfer (the exogenous liquidation cost paid by the banker):

$$\kappa_{t-}^h = \underbrace{y_{t-}^K \kappa_{t-}^p}_{\text{valuation drop}} - \underbrace{\alpha(x_{t-}^d)^+}_{\text{transfer}} \frac{w_{t-}}{1 - w_{t-}}. \quad (19)$$

Relative to the bank budget constraint in (15), the household budget constraint (18) differs mainly in two ways: First, households earn a lower dividend return compared to bankers, $\underline{A}/p_t < \bar{A}/p_t$; Second, during the liquidity shock, households provide emergency funding to banks and earn a profit, while bankers lose net worth due to the financing costs. In practice such a profit is likely intermediated by the central bank, which we have omitted in our modeling. Our modeling implies that there is no destruction of wealth in a liquidation shock, so the household financing benefits (per unit of wealth) multiplied by household total wealth $1 - w_t$ equals to the banker financing costs (per unit of wealth) multiplied by banker total wealth w_t . We could alternatively model the liquidation cost as a deadweight loss. Either way, the objective is primarily to ensure that the liquidity shock is purely financial and has no direct impact on aggregate output.

3.6 Equilibrium Definition

Denote the share of capital owned by bankers as

$$\psi_t = \frac{x_t^K W_t^b}{x_t^K W_t^b + y_t^K W_t^h}. \quad (20)$$

Then the aggregate production of consumption goods is

$$Y_t = (\psi_t \bar{A} + (1 - \psi_t) \underline{A}) K_t. \quad (21)$$

Because $\bar{A} > \underline{A}$, output is increasing in ψ_t .

Given that there is no heterogeneity within bankers and within households, we can express the dynamics of aggregate wealth as

$$\frac{dW_t^b}{W_{t-}^b} = \frac{dw_t^b}{w_{t-}^b} - \eta dt \quad (22)$$

$$\frac{dW_t^h}{W_{t-}^h} = \frac{dw_t^h}{w_{t-}^h} + \eta \frac{W_{t-}^b}{W_{t-}^h} dt, \quad (23)$$

where the second terms in both (22) and (23) are due to the transition of bankers to households.

We derive a Markov equilibrium where all choices only depend on the state variables w_t

and λ_t .¹² Let $\hat{c}^b = c^b/w^b$ be the consumption of a representative banker as a fraction of the banker's wealth, and $\hat{c}^h = c^h/w^h$ similarly. The following formalizes the equilibrium definition.

Definition 1 (Equilibrium). *An equilibrium is a set of functions, including the price of capital $p(w_t, \lambda_t)$, bank debt yield $r^d(w_t, \lambda_t)$, household consumption wealth ratio $\hat{c}^h(w_t, \lambda_t)$ and capital holdings $y^K(w_t, \lambda_t)$, banker consumption wealth ratio $\hat{c}^b(w_t, \lambda_t)$ and capital holdings $x^K(w_t, \lambda_t)$, such that*

- *Consumption, investment and portfolio choices are optimal.*
- *Capital good market clears*

$$W_t^b x_t^K + W_t^h y_t^K = p_t K_t. \quad (24)$$

- *The aggregate non-financial wealth of households and banks equal to total value of capital*

$$W_t^b + W_t^h = p_t K_t. \quad (25)$$

- *Consumption goods market clears*

$$\hat{c}_t^b W_t^b + \hat{c}_t^h W_t^h = (\psi_t \bar{A} + (1 - \psi_t) \underline{A}) K_t - i_t K_t. \quad (26)$$

3.7 Leverage, Risk and Liquidity Premia

For an individual bank, the net funding withdrawal that has to be fulfilled during an illiquidity episode by selling productive capital is $(x_t^d)^+ = (x_t^K - 1)^+$. In Appendix A.3, we prove that:

Lemma 3. *In equilibrium, banks always borrow from households and take leverage, i.e.,*

$$x_t^K \geq 1.$$

The statement is true because banks earn higher returns on holding productive capital than households. Thus, we have

$$(x_{t-}^d)^+ = x_{t-}^d. \quad (27)$$

¹²Under diagnostic beliefs, while agents' beliefs are diagnostic they think that theirs and all other agents' beliefs are Bayesian. In other words, the policy functions and state variables are the same as those under Bayesian beliefs. However, because these policy functions are evaluated under diagnostic beliefs, the equilibrium outcomes are different. Furthermore, the dynamics of the state variables are different due to the underlying difference between diagnostic belief and the true process. The solution strategy for the diagnostic belief model is to solve the Bayesian decision rules under Bayesian belief $\theta = 0$, and then apply the same policy functions and simulate the diagnostic model with the diagnostic belief of $\theta > 0$.

Then the bank budget constraint (15) can be rewritten as

$$\frac{dw_t^b}{w_{t-}^b} = -\frac{c_{t-}^b}{w_{t-}^b}dt + r_{t-}^d dt + x_{t-}^K \underbrace{\left(\mu_{t-}^R + \frac{\bar{A}}{p_{t-}} - r_{t-}^d\right)}_{\text{capital excess return}} dt + \underbrace{x_{t-}^K(\sigma^K + \sigma_{t-}^p)}_{\text{Brownian risks}} dB_t - \underbrace{(\alpha x_{t-}^d + x_{t-}^K \kappa_{t-}^p)}_{\text{losses in illiquidity shocks}} dN_t. \quad (28)$$

Denote the non-jump component of wealth growth as \widetilde{dw}_t^b and the jump component as Δw_t^b . With the results in Lemma 3 and the properties of log utility, we write the banker's optimization problem as:

$$\max_{c_{t-}^b, x_{t-}^d, x_{t-}^K} \left\{ \log(c_{t-}^b) + \frac{1}{\rho} \left(E_{t-} \left[\frac{\widetilde{dw}_t^b}{w_{t-}^b} \right] / dt - \frac{1}{2} \left(\frac{\widetilde{dw}_t^b}{w_{t-}^b} \right)^2 / dt + \lambda_{t-} (\log(w_{t-}^b + \Delta w_t^b) - \log(w_{t-}^b)) \right) \right\} \quad (29)$$

The objective has the familiar mean-variance form over the evolution of wealth that comes from log-utility plus an extra component due to the jump illiquidity shock. We note the key quantities that enter into this mean-variance tradeoff: (1) purchasing capital funded by deposits earns the capital excess return; (2) but this return is at the cost of capital price risk (Brownian risk) and the costs in a bank run (losses in an illiquidity shock).

Denote the return to a bank on holding capital as dR_t^b . Then we have the following first order condition for the excess return earned by the banker in purchasing capital funded by deposits:

$$E_{t-}[dR_t^b] - r_{t-}^d = \underbrace{(\sigma^K + \sigma_{t-}^p)^2 x_{t-}^K}_{\text{Brownian risk premium}} + \underbrace{\lambda_{t-}(\alpha + \kappa_{t-}^p) \frac{x_{t-}^K \kappa_{t-}^p + \alpha x_{t-}^d}{1 - x_{t-}^K \kappa_{t-}^p - \alpha x_{t-}^d}}_{\text{liquidity risk premium}} \quad (30)$$

where the first term is the required compensation for taking on Brownian risk. The two sources of Brownian risk are the exogenous capital shock, σ^K , and the endogenous capital price risk, σ_{t-}^p . The second term is a liquidity risk premium. Deposits are subject to run risk, in which case the bank has to sell capital, suffering the exogenous loss of α and the endogenous fire sale loss of κ_{t-}^p . This possible loss requires a compensation, which is the liquidity risk premium.¹³

Equation (30) can also be used to understand the leverage decision of a banker, which is x_t^d . In particular, consider how news that leads the banker to revise upwards his estimate of λ_t will affect leverage. Since purchasing capital funded by runnable deposits exposes the banker to liquidity risk, this higher liquidity risk will lead the banker to take on less leverage. Figure 10 illustrates this negative relationship in our model.

Finally, turning from a given bank's decision problem to the macroeconomy, a key factor

¹³Note that the expected loss under the physical probability is $\lambda_{t-}(\alpha + \kappa_{t-}^p)$. The term for the liquidity risk premium in (30) reflects the risk compensation for being exposed to these losses.

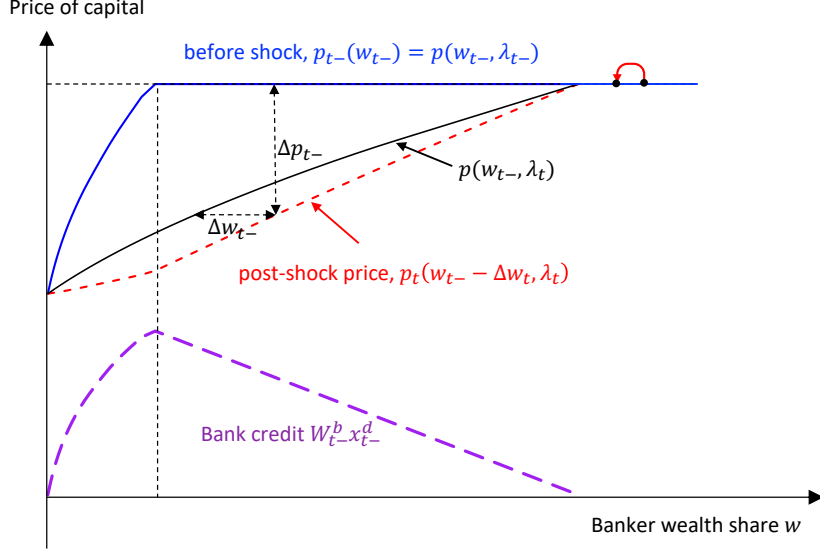


Figure 4: Price of capital and bank credit as a function of w_t , pre- and post- dN_t shock.

in a crisis are the total losses suffered by banks in an illiquidity shock. These losses are proportional to the total debt in the economy, $W_{t-}^b x_{t-}^d$, as well as the exogenous and endogenous liquidation losses, $\alpha + \kappa_{t-}^p$. Thus, the realized losses to the banking sector caused by the bank run is,

$$\text{Loss} = \underbrace{W_{t-}^b x_{t-}^d}_{\text{bank credit}} (\alpha + \kappa_{t-}^p) + W_{t-}^b \kappa_{t-}^p. \quad (31)$$

In the event of an illiquidity shock, W_t^b will jump downwards by approximately this loss. As a result, the capital share ψ_t will fall which can lead to a financial crisis as we next explain.

3.8 State-Dependent Crises

Figure 4 graphs the price of capital in blue as a function the banker's wealth share, w^b , which is one of the state variables in the equilibrium (λ_t is the other state variable). We note that the price of capital is increasing in w^b up to a point and then is flat thereafter. In the increasing portion, both bankers and households own capital. As the wealth share increases, more of the capital is in the bankers' hands, and hence more of the capital produces a higher dividend of \bar{A} . This force leads to a positive relationship between the price of capital and the wealth share. To the right of the dashed line, all of the capital is in the bankers' hands. Now, it will be the case that as the wealth share of bankers rises to the right of the dashed line, the risk premium required by bankers to absorb capital risk falls, which by itself would raise capital prices. However, because of log utility, the interest rate rises to offset the fall in the risk premium, and the net effect on the discount rate is to keep the price of capital constant to the right of the dashed-line.

In the lower part of the figure, we graph total bank borrowing from household which is

$W_{t-}^b x_{t-}^d$, labeled bank-credit. As banker wealth share increases, total bank credit increases as bankers intermediate more of the capital, using borrowing from household. Bank credit is maximal at the dashed vertical line, where banker's hold all of the capital, funded by debt from households. As we go to the right of this point, further increases in wealth mean that banker wealth – equity – rises and hence bank borrowing from household decreases.

The realization of an illiquidity shock has two impacts on equilibrium. First, λ_t rises and the rise in λ_t causes bankers to become more illiquidity averse and de-lever. This causes the price function to shift downwards to the black line on the figure labeled $p(w_{t-}^b, \lambda_t)$. The price reduction is highest around the point where bank credit is highest. Second, the fall in price leads and the exogenous losses due the bank-run lead to a fall in w_t^b of Δw_{t-} , and this loss is amplified as prices are lower at the smaller bank wealth share.

There are three cases of interest. First, if the illiquidity shock occurs when banker wealth share is high – on the right side of the dashed line in the figure – bankers suffer a small exogenous liquidation loss because at high banker wealth/low household wealth, x_{t-}^d is relatively low and hence losses from the run are small. Delevering also has a small impact as bankers own capital with relatively high equity and low debt. The purple dashed-line of bank credit, $W_{t-}^b x_{t-}^d$, shows that bank credit is low for high values of w^b , and hence the loss (see equation (31)) is small. Thus, the post-shock wealth share jumps a small distance to the left, as indicated by the red arrow. But since at this new wealth share, the price of capital is the same as at the old wealth share, there is no endogenous fall in the price of capital and $\kappa_{t-}^p = 0$. Additionally, as output is equal to $(\psi_t \bar{A} + (1 - \psi_t) \underline{A}) K_t$ and $\psi_t = 1$ both before and after the shock in this case, there is no fall in output. Thus, this first case is one of no-amplification of the illiquidity shock.

A second small-amplification case arises when w^b is small. In this case, most capital is held by households, which means the banking sector is small and $W_{t-}^b x_{t-}^d$ is small. The shock triggers delevering of bankers, but given that bankers hold little capital, households also absorb little capital and hence the price reduction (gap from blue to black line) is small. Thus the losses from the illiquidity shock are small, and as a result the absolute fall in w_t^b in this case is small (although large in percentage of w_t^b terms, yet the absolute fall is what matters for macro dynamics). Because the capital price function has some slope in this region, we will have that $\kappa_{t-}^p > 0$, but the effect is minor. We also have that the change in ψ_t and hence the output decline is also small.

The high-amplification case occurs for intermediate values of w^b and high values of bank credit, $W_{t-}^b x_{t-}^d$. In this case, which occurs around the dashed line, delevering triggers a large sale of capital and a large fall in prices. The exogenous loss (proportional to $w_t^b x_t^d$) is high, and leads to a further fall in banker wealth share, which leads to an endogenous fall in the price of capital, which implies further losses to bankers, and so on. The post-shock capital price traces along the red dashed line, reflecting a downward jump in the capital price and

the banker wealth share state variable. The exogenous loss is amplified in this case. We also have that ψ_t falls substantially, causing a large decrease in output. This case is the financial crisis of the model.

Our model thereby captures an amplification mechanism, where the degree is state-dependent. The figure illustrates that a summary variable of the state that measures the extent of amplification is bank credit, $W_{t-}^b x_{t-}^d$. In our simulations, we show that bank credit/GDP forecasts a financial crisis. Not pictured in the figure, but also important to the model's mechanism, is the dependence of amplification on λ_t . We have seen that a low assessment of λ_t leads the banker to increase x_t^d (leverage). Thus, a low λ_t leads the entire curve for bank credit, $W_{t-}^b x_{t-}^d$ in Figure 4, to shift upwards. A useful intuition to help understand our model's results is:

$$\text{Prob of crisis} \propto \underbrace{\text{Bank credit}}_{\uparrow \text{ as } \lambda \downarrow} \times \underbrace{\lambda}_{\text{Prob of liquidity shock}} \quad (32)$$

We return to this relation in Section 5.

3.9 Spreads and Bank Pricing of Liquidity and Credit

In this section, we define spreads that enable us to align the model with data. First, we define the spread on a hypothetical instantaneous loan with interest rate r_t^C and no capital price risk. While there is no price risk when making this loan, we assume it is subject to illiquidity costs of α in the event of a bank run. It is straightforward to show¹⁴ that the spread on this loan relative to the deposit rate is:

$$r_{t-}^C - r_{t-}^d = \frac{\lambda_{t-}}{\underbrace{1 - x_{t-}^K \kappa_{t-}^p - \alpha x_{t-}^d}_{\text{risk-adjusted probability}}} \cdot \underbrace{\alpha}_{\text{liquidity loss}} \quad (33)$$

This object is a pure liquidity spread and reflects banks concern over liquidity risk. We use this spread to help calibrate the unconditional mean intensity of the liquidity shock.

Second, we aim to match the crisis-cycle pattern of bank's credit pricing, which reflect pre-crisis froth, a sharp tightening in the crisis, and a gradual post-crisis recovery. A natural model object that will reflect bank's credit pricing is $E_{t-}[dR_t^b] - r_{t-}^d$ which is banks' required return on holding capital (i.e., loans) over its funding cost. Loosing speaking, $E_{t-}[dR_t^b] - r_{t-}^d$ is a bank's required loan spread. However, the exact historical data we match over the crisis cycle is not loan spreads but credit spreads (see Section 2). There is considerable empirical support for the association between credit spreads and bank lending standards. See [Gilchrist and Zakrajsek \(2012\)](#). We next define a credit spread that is needed to map the model to

¹⁴The derivation detail is in Appendix A.6.

the credit spread data.

We define a zero net-supply defaultable bond, matching the characteristics of the credit spreads in the data presented in Section 2. These defaultable bonds are priced by the banker’s pricing kernel. This last point is worth stressing, as the model-defined credit spread will thus pick-up endogenous variation in bankers’ attitude towards risky lending. We define the credit spread as the yield differential between a risky zero-coupon bond and a zero-coupon safe bond with the same [expected] maturity. We model the default intensity of the bond as related to the intensity of the illiquidity shock, λ_{t-} . In default, the losses to bond holders are affine in the capital price decline κ_{t-}^p . Details on this specification, the bond pricing solution, and the calibration are provided in Appendix A.7.

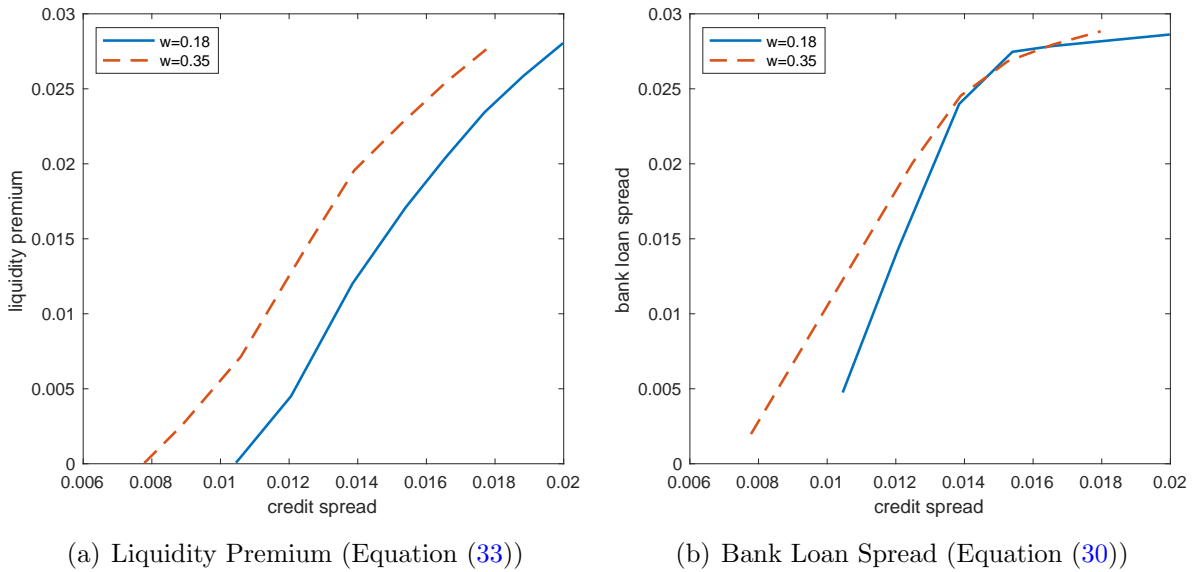


Figure 5: Credit Spread against Bank Loan Spread and Liquidity Premium in the calibrated Bayesian Model. In this figure, we show the relationship between credit spread, liquidity premium, and the bank loan spread. We fix the state w and then trace the relationship among different spreads along the λ dimension.

Figure 5 plots the credit spread in the calibrated model against the liquidity premium, $r_{t-}^C - r_{t-}^d$, and the loan spread $E_{t-}[dR_t^b] - r_{t-}^d$. The variation in the spreads is generated by model’s variation in the state variable λ_t . The figure plots this relation for two different values of w , one at the median, and one at a higher value of w . The upshot from this figure is that all of these spreads move together. We use the liquidity premium and the credit spread as targets for calibration because they have measured data counterparts.

3.10 Solution Methodology and Simulation

The challenge of solving this model comes from both multiple state variables and the endogenous jumps in the state variables. To ensure stability, we use a functional iteration

method that begins with an initial guess of the capital price function $p^{(0)}(w, \lambda)$, and then iterates over the equilibrium equation system to get an updated price $p^{(1)}$. This updating step involves solving a fixed-point problem at each state (w, λ) . Then we iterate until at step k , we have

$$\int_0^1 \int_{\lambda_L}^{\lambda_H} |p^{(k+1)}(w, \lambda) - p^{(k)}(w, \lambda)| d\lambda dw < \varepsilon$$

for a small positive number ε .

We simulate the model at a monthly frequency but analyze simulations at a yearly frequency to be consistent with the data.

- We set the simulation interval as $dt = 1/12$ (a month), and generate the independent Brownian shocks $dB_t \sim \mathcal{N}(0, \sqrt{dt})$, as well as an independent frequency of illiquidity shock process $\tilde{\lambda}_t$. Based on the illiquidity shock process $\tilde{\lambda}_t$, we generate illiquidity shocks dN_t that hits with probability $\tilde{\lambda}_t dt$ for the time interval dt .
- Once shocks are generated, we solve for the dynamics of state variables, including w_t , λ_t , and K_t . For the static belief model, $\lambda_t = \bar{\lambda}$. For the diagnostic belief model, we need to generate λ_t^θ based on λ_t .
- With state variables determined, we generate all other quantities and prices of the model.
- We discard the first one thousand data points of each simulation path collected in this manner. As a result, the initial values do not affect our computed moments. The simulation approximates picking initial conditions from the ergodic distribution of the state variables.
- Finally, we average all of the monthly quantities for a given year to get an annual data set. For prices, we use the first observation of every year.

4 Calibration

In order to map model outputs to data, we need to define a financial crisis. Crises are the events where the growth in bank credit/GDP in a given month falls into the lowest 4% quantile of the distribution of monthly bank credit/GDP growth rates, and there has not been another such event in the previous three years. This latter criterion is to ensure that a longer crisis is still dated as a single crisis, as is done in the empirical literature. This crisis corresponds to a disintermediation event, and in the simulation almost always involves an illiquidity shock and bank run, although as crises are endogenous, not all illiquidity events are crises. We target the 4% number based on fact 1 of Section 2. We also consider a crisis definition based on bank equity crashes, as in [Baron and Xiong \(2017\)](#), in Section 5.9.

We solve and calibrate three variants of the model. The appendix considers further models of distorted beliefs including optimism and pessimism as well as a variant without financial frictions. The variants we study in the main text are:

1. Bayesian (rational) Model: Agents form beliefs over the illiquidity state following Bayes rule, and this belief varies over time.
2. Diagnostic (non-rational) Model: Agents form beliefs over the illiquidity state via diagnostic expectations, and belief varies over time.
3. Static-belief Model: Agents' beliefs are constant.

We apply a combination of external and internal moments to calibrate model parameters. Specifically, we directly set parameter values for those with standard values in the literature. Then we estimate the rest of parameters based on moments chosen to best reflect the economics of those parameters.

Table 1: Externally-Calibrated Parameter Values

Choice	Moment		
δ	Depreciation rate	10%	Literature
ρ	Time discount rate	4%	Literature
χ	Investment adjustment cost	3	Literature
α	Distress illiquidity costs	0.05	Data
θ	Diagnostic belief weight	0.90	Literature

A list of externally-calibrated parameters for the model (not including the credit spread which is given in Appendix A.7) are shown in Table 1. We follow the macroeconomics literature to set annual depreciation rate $\delta = 0.1$ (Gertler and Kiyotaki, 2010), annual time discount rate $\rho = 4\%$ (Gertler and Kiyotaki, 2010), and investment adjustment cost $\chi = 3$ (He and Krishnamurthy, 2019). For the emergency liquidity costs (α), we do not have good data for the historical financial crises to pin these down. From data of the 2008 crisis, the effective liquidation loss is about 0.05, which is the value of $\alpha \cdot \beta$ in Li (2019). Alternatively, we can interpret this liquidation loss as a funding premium. The value of $\alpha = 0.05$ translates to a 10% premium for a illiquidity event that lasts 6 months. Last, in our investigation of beliefs in the model, we choose the diagnostic parameter θ based on the research by Bordalo, Gennaioli and Shleifer (2018), Bordalo et al. (2019b), and Bordalo et al. (2020). These authors estimate θ based on the dynamics of forecasts for financial and economic variables. We set θ equal to 0.9, which is the value used by Bordalo, Gennaioli and Shleifer (2018) and Bordalo et al. (2019b).¹⁵

¹⁵Bordalo et al. (2020) report an estimate of θ of 0.5.

Table 2: Internally-Calibrated Parameters and Targeted Moments

Panel A. Moments				
	Data	Static	Bayesian	Diagnostic
Average liquidity premium	0.90%	0.86%	0.77%	0.70%
Avg credit spread change in crises	70%	9%	56%	55%
Half-life of credit spread recovery (years)	2.5	3.4	2.5	2.5
Output/capital ratio	14%	16%	17%	15%
Avg 3-year output drop in crises	-9.1%	-8.9%	-7.9%	-9.2%
Output growth volatility	3.8%	3.6%	2.9%	3.8%
Average bank leverage	5.0	5.1	5.1	5.0

Panel B. Estimated Parameter Values				
	Parameter	Static	Bayesian	Diagnostic
Avg frequency of liquidity shock	$\bar{\lambda}$	0.072	–	–
High intensity of liquidity shock	λ_H	–	0.561	0.638
Low to high transition	$\lambda_{L \rightarrow H}$	–	0.11	0.11
High to low transition	$\lambda_{H \rightarrow L}$	–	0.47	0.48
Household productivity	A_L	0.12	0.17	0.13
Bank lending advantage	$A_H - A_L$	0.060	0.030	0.024
Volatility of capital growth	σ^K	0.06	0.03	0.03
Banker-household transition rate	η	0.120	0.055	0.034

Then we proceed to estimate other parameters, reported in Table 2 Panel B. The Static Belief model has only one parameter $\bar{\lambda}$ governing the crisis frequency process which is constant over time. There are four parameters governing beliefs in the Bayesian model: λ_H , λ_L , $\lambda_{L \rightarrow H}$, and $\lambda_{H \rightarrow L}$. We note that as long as λ_L is close to zero, the impact of its value is negligible. Therefore, we pick $\lambda_L = 0.001$ directly. The diagnostic model adds θ as one more degree of freedom (the “look-back period” parameter t_0 is set to 1, the implicit value from discrete-time diagnostic belief process such as [Bordalo, Gennaioli and Shleifer \(2018\)](#)). The Bayesian Belief and Diagnostic Belief models are exactly identified, while the Static Belief model has two more moments than parameters. After experimentation with the model, we find that the following moments to be particularly informative for the belief parameters:

- The average liquidity premium will reflect banks’ assessment of liquidity risk, and thus the average value of λ . See equation (33). The spread between P2 rated 3-month commercial paper and 3-month T-bills in data from 1974 to 2018 is 94 basis points. We target a liquidity premium of 90 basis points. [Krishnamurthy and Vissing-Jorgensen \(2015\)](#) estimate the average liquidity premium on long-term Treasury bonds relative to AAA

corporate bonds to be 75 basis points. We focus on a short-term bond in our exercise and thus target a higher spread. Our estimate reported in Panel B implies an average λ for the static model of 0.072, which translates to a liquidity event once over 13.9 years. In the high illiquidity state, the λ_H for the Bayesian and Diagnostic models are a little around 0.6 implying a liquidity event roughly every 1.66 years.

- Credit spread changes during a crisis (fact 2). The spike in the credit spread is $0.7\sigma_s$. This moment helps determine $\lambda_{L \rightarrow H}$, which affects the degree of surprise in beliefs due to the realizations of illiquidity shocks.
- Half-life of credit spread recovery (fact 4). According to [Krishnamurthy and Muir \(2024\)](#), the half-life is 2.5 years. This moment primarily determines $\lambda_{H \rightarrow L}$, since the speed of recovery of beliefs after a illiquidity shock is directly affected by the underlying transition probability.

The parameters \bar{A} , \underline{A} , σ^K , and η govern the output process both in and out of crises. The following targets inform these parameter choices:

- Average output decline during a crisis (fact 3): We target -9.1% as explained in Section 2. This moment is most directly related to the productivity differential $\bar{A} - \underline{A}$.
- Investment to capital ratio: We use the same target as [He and Krishnamurthy \(2019\)](#). This moment mainly affects the average of productivity parameters, \bar{A} and \underline{A} .
- Average output growth volatility: According to Bohn’s historical data, the volatility of real GDP growth from 1791 to 2012 for the U.S. is 4%. This moment mainly affects the capital volatility σ^K .
- We map banks in the model to depository institutions and broker dealers in the flow of funds. Bank equity is defined as total bank assets minus total bank liabilities. Since our model only captures runnable liabilities, we define effective bank liabilities as total liabilities minus insured deposits. Then we calculate bank leverage as (bank equity + effective bank liabilities)/bank equity. Using all data available, we find that bank leverage is approximately 5. This moment disciplines η , the transition rate from bankers to households, which affects the stationary distribution of leverage in the model. For example, setting η very low leads to a stationary distribution where almost all of the wealth is in bankers’ hands and average leverage in equilibrium is very low.

To search for parameter values that best match moments, we need to repeatedly solve the model for a large combination of parameter values. A simple discretization of the parameter space (5 parameters for the benchmark, 7 parameters for the Bayesian and diagnostic models) renders the task computationally infeasible. To resolve this difficulty, we apply the

Smolyak grid method (Judd et al., 2014) to generate a discretized state space. For each version of the model, we follow the estimation procedure:

- Discretize the state space of parameters around their initial values. We pick a discretization level of 3 in the Smolyak discretization. This results in 177 combinations for the static belief model, 241 combinations for the Bayesian model, and 389 combinations for the diagnostic model. Simulate all of these models and collect their moment values.
- Denote the moments in the data as m_1, \dots, m_J , and the moments from the model as $\hat{m}_1, \dots, \hat{m}_J$. From all of the parameter combinations, pick the one that minimizes the objective

$$\sum_{j=1}^J weight_j \frac{|\hat{m}_j - m_j|}{m_j}.$$

Here $weight_j$ reflects the importance of a given target in the estimation. We set the weight for the liquidity premium to be three and the rest of the weights to be one. The average liquidity premium determines the frequency of illiquidity shocks which is a particularly important parameter in the model.

- Once we have picked a set of parameters, we search in a smaller region around this set of parameters and find a new best set of parameters in the smaller region. We iterate the above process until the difference between the optimized objective value between two iterations is below a threshold.

5 Model Evaluation

This section evaluates the models we consider and explains the mechanisms that help the models match the crisis data patterns.

5.1 Targeted Moments

We target means across crises in choosing parameters. Table 2 in the previous section presents the model’s fit in hitting the targets. We re-calibrate the model parameters to best match moments for each version of the model, thus giving each model the best chance to represent the data. Although each version of the model is at least exactly identified (static belief model is overidentified), because the state-space is restricted, we do not fit all of the moments accurately. The static belief model, in particular, misses the spread change in the crisis by a wide margin. It is possible to fit this moment if we increase the exogenous liquidation cost α , but we opt to keep α constant across all of the models to better illustrate the mechanisms underlying the models.

From Table 2, we see that the main difference in parameters between diagnostic and Bayesian model are in λ_H and $A_H - A_L$. The diagnostic agent tends to be more optimistic about the distribution of λ on average, as there is a greater concentration of mass in the distribution within low λ states. Consequently, holding all other factors constant, both the diagnostic agent's mean λ and the average liquidity premium decrease. Then, in order to match the average liquidity premium, the estimation chooses a λ_H of 0.638 rather than 0.561 (even with this higher λ_H , the model still produces an average liquidity premium that is lower than the data). Holding all else constant, a higher value of λ_H will lead to a more significant decline in output and an increase in credit spreads during a crisis. To match the data moments, the estimation reduces $A_H - A_L$ from 0.030 to 0.024. This comparison helps explain one of our main findings: our analysis demonstrates that the Bayesian and diagnostic models yield comparable outcomes. This is because the estimation procedure generates distinct parameter values for these models to match the data moments, subsequently aligning their outputs more closely.

5.2 Average Patterns across Crises

Figure 6 plots the path of the model-generated credit spread, bank credit/GDP and GDP around a crisis at $t = 0$. The credit spread and bank credit variables are plotted in units of standard-deviations from their mean value over the sample. The figure should be compared to the data in Figure 1. We see that the model is able to generate the jump in spreads, contraction in credit, and drop in GDP. For both the Bayesian and Diagnostic model, the magnitudes of the spread spike and GDP decline are also in line with the data. During a crisis, spreads jump about 45% in the model (that is, 0.5σ s) and 70% in the data. As noted above, the magnitude of the spread spike in the static belief model is too small relative to the data. The magnitude of the credit contraction of around 0.6σ s for the Bayesian model is larger than the data counterpart of 0.33σ s. This is likely because in our model all credit is extended via banks, while in the data, there are other intermediaries involved in the credit process. Note that we have not explicitly targeted the credit contraction in the calibration.

[FIGURE 6 HERE]

All of the models match the sharp transition in the crisis, driven by the model's amplification mechanism, and output that is below trend for a sustained period post-crisis. The figures also reveal how the pre-crisis patterns vary across the models. In the years before the crisis, bank credit and GDP are rising while credit spreads are below normal in both dynamic belief models. In the static belief model, spreads are slightly higher than normal, while credit is falling. This contrast points to the need for time-variation in beliefs to fit the data.

Figure 7 plots the annual volatility of the return on capital around crises. Volatility is defined as,

$$\sqrt{(\sigma^K + \sigma_t^p)^2 + \lambda_t(\kappa_t^p + \alpha)^2}, \quad (34)$$

which includes both the Brownian and Poisson shocks. In all three versions of the model, volatility is low before the crisis and high in the crisis. This pattern is qualitatively in keeping with the countercyclical behavior of volatility documented in existing research (Geanakoplos, 2010; Adrian and Brunnermeier, 2016; Danielsson, Valenzuela and Zer, 2018). The absolute level of volatility of the return on capital is low, ranging from 4% to 5.5% in the Bayesian model compared to historical volatility of U.S. stock market returns of 16%. This is a common result in a real business cycle model with moderate investment adjustment costs: quantities adjust rather than prices. Yet the mechanism of our model is via the fluctuations in the equity value of banks (i.e. net worth of bankers), which is a levered claim on capital. The unconditional volatility of that object in the Bayesian model is 15.7% (see Appendix A.8), which is in line with the volatility of U.S. stock market returns. The empirical research on crises has not offered quantitative targets to match for volatility dynamics around crises. Comparing volatility dynamics across the three versions of the model, we see that with belief variation (either Bayesian or diagnostic), volatility drifts down before crises. There is a “calm before the storm” as in the data. In the model without belief variation, volatility drifts up because forward looking agents recognize that risk is rising as the state variable moves closer to the crisis region.

[FIGURE 7 HERE]

The variation in risk is directly relevant to the banker’s portfolio decision and thus the equilibrium risk premium. The F.O.C. for the banker gives the risk premium on capital:

$$(\sigma^K + \sigma^p)^2 x_t^K + \lambda \frac{\kappa^p + \alpha}{1 - x^K \kappa^p - \alpha(x_t^K - 1)^+}. \quad (35)$$

Figure 7 plots the Sharpe ratio, which is the risk premium in this equation divided by the volatility in (34). The Sharpe ratio closely matches the countercyclical volatility pattern. The levels are also in keeping with typical estimates in the literature for the U.S. stock market, which center around 0.50. Muir (2017) estimate that the dividend/price ratio rises by 43% at the peak of a financial crisis relative to the pre-crisis average (see Table III of the paper). In the Bayesian belief model the highest Sharpe ratio in the crisis is 56% higher than the pre-crisis average, and in the diagnostic model the same increase is 82%. We also note that the Sharpe ratio dips in the run-up to the crisis in the Bayesian and diagnostic models, but not the static belief model. We develop this point, which reflects pre-crisis “froth,” further below. We will also show that our model can match the empirical association between risk premia and credit growth, as documented by Baron and Xiong (2017).

The success in matching the mean patterns around crises verifies that our model’s mechanisms can speak to the data. In Appendix A.8 we report additional non-crisis moments from the model simulation, including the model’s average Sharpe ratio, risk free rate, and asset price volatilities. We also report the volatility of investment and consumption growth.

5.3 Ergodic Distributions

In Figure 8, we graph the ergodic distributions of the state (w_t, λ_t) for the three models. Underlying movements in w are driven by three forces: the exogenous diffusion shocks to capital shift wealth, creating paths from the center of the distribution to both right and left; paths that go to the left are pushed back to the middle because in low w states, risk premia are high and bankers expected wealth growth is high; the transition rate of bankers into households, η , result in a drift in w of $-\eta w$, which pushes all paths to the left. The result of these forces is a mean-reverting w process and the single-peaked distribution. In the diagnostic and Bayesian models, the realization of a jump leads to a larger adjustment in w relative to the static beliefs model, because agents belief shift from the low illiquidity to the high illiquidity state. As a result, more mass is shifted to low- w states. Broadly, all three of the models generate a similar left-skewed distribution.

[FIGURE 8 HERE]

5.4 Non-targeted Cross-sectional Moments

We next describe the model’s fit in the cross-section of crises, which are non-targeted moments. Within the sample of crises, there are smaller and larger crises. The moments we report measure variation within these crises. We discuss the model’s fit of these non-targeted moments in this section, and delve further into the fit in the next sections. Table 3 summarizes the model’s performance in matching cross-sectional moments.

Panel A reports all the models’ fit on the data patterns on the crisis and its aftermath. First, in the data, episodes where credit spreads increase more are followed by larger output contractions. The first row of Panel A reports these moments from the models and data. All of the model get the signs right, with the Bayesian belief model coming closest to matching the data. Second, crises that are preceded by a run up in bank credit are also more severe crises. The second row of Panel A reports the models’ fit with data on this dimension. In general, all three of the models get the signs right and are in the ballpark of the data, with the Bayesian model closest to matching the data. We note that these crisis facts are well modeled even in the static belief model which only has a financial intermediation mechanism.

Panel B reports that all of the models are able to fit the negative data relationship

between bank credit and risk premia, with the diagnostic model matching the data most closely. The panel is not explicitly about financial crises, but more generally about the relationship between movements in credit and risk premia. In the data, credit growth is negatively correlated with excess equity returns (Baron and Xiong, 2017). Periods of high credit growth are followed by low returns, and periods of low credit growth are followed by high returns. We verify that all of the models we consider deliver this relation, with the diagnostic model matching the data moment most closely. The models match the data via time variation in the supply of risk-bearing capacity. The state variables of the model, such as w , capture variation in the effective risk aversion of the banking sector. When effective risk aversion is low, banks lend more and credit grows, while risk premia are low; the opposite pattern holds when risk aversion is high. This mechanism thus delivers the relation between bank credit and risk premia. The fact that this relationship holds even in the model with static beliefs bears stressing: a sentiment/belief mechanism is not necessary to replicate the credit/risk premia relationship.

Panels C and D consider the pre-crisis patterns where we see divergence across the models. In Panel C, we examine whether the model can reproduce the fact that spreads are below normal before crises. The first row considers the mean pre-crisis spread. Both the diagnostic model and the Bayesian model deliver the below normal spread, while the static belief model delivers an above normal spread. We explain this failure in further detail below. Panel D considers the predictive relationship between measures of credit market excess and subsequent crises. We again see that the static belief model fails to generate a sign in keeping with the data. Both of the belief models succeed in this dimension and provide similar results, although each model does better (and worse) in some dimensions.

[TABLE 3 HERE]

5.5 Mechanism 1: Frictional Intermediation and Credit

Figure 9 graphs the histogram of 3-year GDP growth in crises for all three models. In a model with no financial amplification and only diffusion shocks to AK_t , output growth would be normally distributed. All three models, and especially the static model with only an amplification mechanism, generates the skew in line with the data. Thus, we conclude that the left-skewed output growth distribution in line with the data can be generated by a pure financial amplification mechanism.

[FIGURE 9 HERE]

In the data, the skewness in output growth matches the skewness of the jump in credit spreads in the crisis (fact 7). Panel A in Table 3 evaluates the relationship between the

jump in credit spreads in this model and the fall in GDP. The bottom row of Panel A evaluates the relation between the run-up in bank credit at the start of the crisis and the subsequent severity of the crisis. This is a relation reported by several empirical studies (Jordà, Schularick and Taylor, 2013). As noted all of the variants of the model get the signs right, with the diagnostic and Bayesian model closer to the data counterparts. We have also seen in Panel B of Table 3 that all of the models generate the negative relation between bank credit and excess equity returns (risk premia). The models capture this moment because all of the models embed variation in the supply of risk-bearing capacity that drives both credit and risk premia.

These observations indicate that the frictional intermediation mechanism, which is the only mechanism present in the static belief model, can capture the patterns of the economy in a crisis and its aftermath. Again, it is possible to improve the quantitative fit of the static belief model for the crisis and its aftermath if we allow α to vary across models and be determined via the estimation. We choose not to go down this path because, as we explain next, this static belief model fails to fit the pre-crisis facts even qualitatively. It is likely also possible to bring the belief models closer to the data by adding further realistic features into the model, but we have chosen not to so.

5.6 Mechanism 2: Beliefs and Credit

We report in Table 3 Panel C that the static belief model generates a spread that is higher than normal in the pre-crisis period, contrary to the data. The failure can be understood as follows. The amplification mechanism of the model, which is what drives the response of the economy to the illiquidity shock, is governed by the extent of bank credit. When the credit state is high, a negative shock triggers a large fall in GDP and a crisis. However, since the credit spread is forward-looking, variation in the spread is also driven by the high credit state. The economy is more vulnerable in the high credit state and hence credit spreads are higher when credit is low. As a result, the static belief model generates an above normal spread before a crisis, contrary to the data.

The belief models are able to generate a spread with the right sign of the data.¹⁶ To understand the economics here, consider Figure 10. We graph the policy function of bankers, for both Bayesian and diagnostic models, in choosing borrowing as a function of the true state λ (denoted “rational” in the figure). Bankers in our model lever up to gain high returns on capital, but at the cost of the illiquidity event where they suffer bankruptcy costs from liquidating capital. Thus there is a bankruptcy risk/return tradeoff that drives their borrowing decision. When λ is low, the illiquidity event is less likely, and the banker chooses

¹⁶We report the results of a regression of spreads on a dummy that takes the value of one for the 5 years before a crisis. This regression also includes a control for the 5 years after the crisis so that the pre-crisis dummy indicates the level of spreads relative to non-crises periods.

high borrowing; hence, the negative slope in the curves in the figure. A useful relation to keep in mind is:

$$\text{Severity of crisis} \propto \text{Bank credit} \quad \text{and,} \quad (\text{Credit} \uparrow, \text{Spreads} \downarrow) \text{ as } \lambda \downarrow$$

When λ is low and credit is high, if an illiquidity shock dN_t occurs, then its impact on GDP will be severe and more likely to result in the large GDP decline of a crisis. Finally, when λ is low, spreads are low, as is evident from Figure 5. This endogenous relationship between illiquidity risk and financial fragility generates the low credit spread before crises.

Both the Bayesian and diagnostic model with the calibration of $\theta = 0.9$ generates the low spread. The diagnostic model strengthens the belief mechanism further relative to the Bayesian model. Consider the red dashed curve in Figure 10. We plot the banker's leverage decision as a function of the true lambda – not the agent's perceived diagnostic lambda. Clearly, at lambda of zero, the true and diagnostic lambda are the same. But as lambda becomes larger than zero, the diagnostic agent chooses higher leverage than the Bayesian agent. This is because the banker is overoptimistic and thinks lambda is lower than it actually is. When the true lambda is larger than a threshold, the banker is on average over-pessimistic and thinks lambda is higher than it actually is, thus choosing lower leverage. As a result, the leverage/lambda curve steepens under the diagnostic model generating a lower pre-crisis spread than the Bayesian model.

Quantitatively, [Krishnamurthy and Muir \(2024\)](#) report a range of estimates for the pre-crisis spread from $0.26\sigma_s$ to $0.44\sigma_s$ depending on exact specification. Thus, both the Bayesian and diagnostic model fall within the range of estimates from the data.

[FIGURE 10 HERE]

5.7 Pre-crisis: Predicting a Crisis with High Bank Credit

Next, we quantitatively evaluate the forecasting power of bank credit for crises. Table 3 Panel D, second row, presents the crisis prediction result. For the bank credit regression, [Schularick and Taylor \(2012\)](#) report that a one-sigma increase in bank credit/GDP increases the probability of a crisis over the next year by 2.11%. There are other estimates from the literature that report higher probabilities, as we discuss below.

In the dynamic belief models, we find the variables have the right signs, although the models are somewhat high in terms of magnitudes for the credit quantity, and low for the credit spread variable. The static belief model fails again, generating a sign that is the opposite of the data.

To understand what drives the mechanism in the dynamic belief models, let us return to

the intuition:

$$\text{Prob of crisis} \propto \underbrace{\text{Bank credit}}_{\substack{\uparrow \text{ as } \lambda \downarrow}} \times \underbrace{\lambda}_{\text{Prob of liquidity shock}} \quad (36)$$

There are two competing forces at work. As λ falls, endogenous leverage and bank credit rises, but the probability of the illiquidity shock falls. If the leverage force is stronger, as it is in both versions of the calibrated belief models, we match the data relationship between high credit and higher probabilities of a crisis.

Figure 11 illustrates this further. We plot the density of GDP growth over the next year conditional on the level of credit/GDP today. The red lines correspond to the Bayesian model and the dashed-blue lines correspond to the static-belief case. In panel (a) of the figure, we condition on low bank credit/GDP which as illustrated in Figure 4 is the case when amplification is low. As a result, the economy is faced with moderate volatility of GDP but this volatility is confined to the center of the distribution and there is little mass at the left tail. Next, consider panel (b) where condition on high levels of credit, which can happen when λ is low and the banker is not illiquidity averse. This is a high amplification state. The dotted black vertical line on the figure indicates the cutoff we have used to define a financial crisis. Mass is now pushed from the center of the distribution towards the left-tail crisis states. Effectively, the more risk-tolerant banker is willing to take on more liquidity risk when making decisions. There is less risk at the center of the distribution, but more mass in the tail. As a result, high credit states forecast more left-tail events.

The static belief model has only bank credit as the variable to drive effective risk aversion. With only this state variable driving decisions, the banker chooses leverage in a manner that crises are avoided when w and credit are higher. As shown in Panel C and D of Table 3, the signs on the credit-crisis relationship are the opposite of that in the data. This result reinforces a lesson of our analysis that we do need a model with two state variables to explain the entire crisis cycle.

[FIGURE 11 HERE]

Panels (c) and (d) of Figure 11 plots the distribution of GDP growth over the next year conditional on different levels of credit in the diagnostic model relative to the Bayesian model. We plot the diagnostic model's distribution in green dashed lines and the Bayesian model in red. We can see that the forces that work to generate the relation between froth and crises are similar but stronger in the diagnostic model compared to the Bayesian model. As we go from left to right panel in the figure, the mass in the left tail rises.

[FIGURE 12 HERE]

Figure 12 examines the predictive relation in a different way. In the figure, we plot the banker’s wealth return conditional on high and low values of credit. Recall that our banker has log utility, so the mean and variance of this distribution are the key statistics driving banker utility and the leverage decision. The banker’s wealth volatility is highest in the low credit case (left panel) driven by a significant mass spread between -0.1 and 0.4 at the center of the distribution. Distress and bankruptcy costs are salient to the banker, and thus he chooses low leverage. In the right panel high credit case, the output distribution is tight so that over most of the distribution, there is little distress for the banker. While there is a tail of wealth losses in crisis states, the banker’s decision to take high leverage is largely driven by the tight central peak of the distribution. The banker understands that the typical negative shock will have small effects on his wealth, and is willing to gamble on avoiding the large tail shock. Note also that the banker’s wealth process is different from the economy-wide GDP process, as should be expected in a model where banks drive systemic risk. Banker wealth is more sensitive than GDP to small shocks, and since such shocks are more likely, they are the drivers of the banker’s leverage decision. As a result, the model produces the result that in the Bayesian model, even if illiquidity events are less likely (low λ), crises are more likely.

[FIGURE 13 and 14 HERE]

Figure 13 panel (a) presents the probability of a crisis over the next three years conditional on different quintiles of Bank Credit/GDP. These are marginal probabilities based on a Probit regression of a dummy for the occurrence of a crisis over the next three years on Bank Credit/GDP. The pseudo- R^2 in this regression is 5.7%, which is in line with [Schularick and Taylor \(2012\)](#)’s analysis of historical crises. We note that the probability is near 20% in the highest quintile. In their survey, [Sufi and Taylor \(2022\)](#) report based on historical crises data that when credit growth is in its highest sextile, the probability of a crisis quadruples from the unconditional value of 2.5% to 10% (see Figure 6a). Note that we calibrate our model to an unconditional probability of crises of 4%, so that a quadrupling is in line their estimates. Another estimate in the literature comes from [Greenwood et al. \(2020\)](#) who report a 37.3% probability of a crisis in the next three years conditional on credit growth in its highest quintile as well their measure of an “asset price bubble” in its highest tercile. Note that their predictive regression uses information on both credit growth as well as their bubble indicator, while our regression only uses credit growth. To provide a sense of how much more information an econometrician may be able to extract, possibly uncovering a bubble, in panel (b) of Figure 13, we report the conditional crisis probability now conditioning on both Bank Credit/GDP and the true λ (as opposed to agents’ perceived λ). Here we see crisis probabilities in the highest quintile are 45%. Figure 14 presents the same analysis for the diagnostic model. We note that the probabilities are in the same ballpark as the Bayesian model, consistent with our general finding that these models perform about as well

as each other. In the diagnostic model, the true-lambda crisis probability rises from 45% to 50%, consistent with the diagnostic model’s mechanism that agent beliefs are further from the true λ , although the gain in predictability is modest.

5.8 Pre-crisis: Predicting a Crisis with Low Credit Spread

We next turn to the relation driving froth (low credit spreads) and crises as reflected in the first row Table 3 Panel D. To replicate the spread predictability regressions in [Krishnamurthy and Muir \(2024\)](#), we define “high froth” as a dummy that indicates whether the credit spread is below its median value at time t . In [Krishnamurthy and Muir \(2024\)](#), the froth definition is based on credit spreads being below median over a 5 year period, which is necessary because a crisis in the data is not sharp 0-1 phenomena as in our model (spreads typically rise before the historian-dated crisis). We predict the likelihood of a crisis over the next 5 years in the model, in line with the data moment.

As we will explain, the froth relation holds for the belief models in the parameterization we study, but need not hold generally. Figure 15 draws density plots of next-year GDP growth for the diagnostic, Bayesian and static belief model conditional on different levels of the credit spread. We can see that the static belief model gets the sign of the mass shift wrong. The diagnostic and Bayesian models, on the other hand, succeed in this dimension. We see again that the relative to the Bayesian model, the diagnostic model shifts more mass to the left tail when spreads are low, and leverage is endogenously high. We also see that the larger shift of the diagnostic model brings the coefficient more in line with the data, albeit still too small. See Table 3 Panel D.

[FIGURE 15 HERE]

The logic behind froth is more nuanced than for the high credit relation of the last section. There are two forces driving variation in the credit spread that are salient for understanding the mechanisms: (i) lower λ means less illiquidity events and hence lower spreads; (ii) worse crises mean higher loss-given default (via κ_t^p) and hence higher spreads. If we imagine shutting down effect (ii), then we can understand the froth relation easily from equation (36). Now, if we add back effect (ii), the froth relation is weakened. The reason is that more crises imply larger losses given default and hence higher ex-ante spreads. If we consider the extreme case where $\kappa = 0$, and hence recovery has no fixed component, the regression coefficient on froth falls to 0.96, thus reduced further relative to the data counterpart. The sign of the froth relation depends quantitatively on the exact cyclicity of recoveries in default and thus the relation between λ and spreads. We have calibrated our model to the history of recoveries on BAA bonds in the U.S., as reported by Moodys.

5.9 Bank equity crises

Baron and Xiong (2017) and Baron, Verner and Xiong (2021) define financial crises in terms of large ($<30\%$) declines in bank equity values. They note that many of the crisis patterns documented in the narrative crisis dating literature (e.g., Laeven and Valencia (2013), Jordà, Schularick and Taylor (2011)) hold for this quantitative definition of financial crisis. In this section we define an equity-crash crash as an event where the return on bank equity in a given quarter is below $-X\%$, where X is chosen to yield a frequency of bank equity crashes of 4%. Because crashes can cluster in our simulation, we define the crash-crisis as the first crash that occurs after at least 3 years of no crash-crises. Thus we are effectively defining a crash as a single crisis. In our simulations, $X = 42\%$ for the Bayesian model and 41% in the diagnostic model.

Table 4 Panel A reports the declines in GDP in the 3 years subsequent to the crash-crisis. Baron, Verner and Xiong (2021) report that a crash-crisis is followed by a GDP decline of around 4.5%. Our numbers are larger than theirs. They also consider a definition of crisis which involves a crash and a banking panic. In this event, they show the GDP declines are about 6%. This latter definition is more in line with our model, as a crash almost always occurs with a liquidity shock. We also report the interaction regression, describing how bank credit pre-crisis worsens GDP outcomes in an equity crash. Analogous to our earlier results, bank credit is a vulnerability indicator for GDP declines in a crisis. Note that this is a regression we do in the model, but is not presented in Baron, Verner and Xiong (2021).

Table 4 Panel B presents predictive regressions, analogous to Table III of Baron and Xiong (2017), of bank credit/GDP on the likelihood of an equity crash. Baron and Xiong (2017) report that the the marginal probability of an equity crash rises by around 5.4% (column 7, top row of their table) in response to a one-sigma increase in bank credit/GDP growth. We run this regression in our simulated data and evaluate the change in the probability of an equity crash for a one-sigma increase in bank credit/GDP, evaluated at the mean value of bank credit/GDP. Our model regressions are in line with the data, although the magnitudes are higher than the data regression.

[TABLE 4 HERE]

We have not reported these regressions for the static belief model. That model implies the wrong sign relative to the data. We can see this in Figure 16, which is a plot of the return on bank equity from month t to $t + 1$ if an illiquidity shock occurs against the time t value of bank credit/GDP. Over the entire range, we see that the relation is positive rather than negative.

6 Conclusion

Financial crises have clear regularities. The 2008 global financial crisis was not a unique event. Over the last two decades, researchers have documented a number of common empirical patterns of financial crises. The main contribution of our paper is to apply a model in the class of the recent non-linear macro-finance models (He and Krishnamurthy (2013); Brunnermeier and Sannikov (2014); Di Tella (2017); Gertler, Kiyotaki and Prestipino (2020)) plus a learning mechanism (Moreira and Savov (2017); Bordalo, Gennaioli and Shleifer (2018)) to matching these patterns. We have shown that our model with a financial amplification mechanism plus belief dynamics, either driven by Bayesian or extrapolative expectations, is able to generate patterns on the crisis cycle consistent with the empirical literature on financial crises. The model matches the pre-crisis froth and debt build-up. It matches the sharp transition to a crisis, the left-skewed distribution of output declines and asset price declines, and the slow post-crisis recovery. The quantitative fit of the model does leave room for improvement: the model generates froth pre-crisis, but not as much as the data, and while the association between credit and subsequent crises is positive, it is too strong relative to the data.

Our research also helps to clarify the role of beliefs and learning in matching the crisis cycle. In our model, the crisis is triggered by a “Minsky moment;” a shock that sharply shifts agents’ beliefs regarding liquidity risk and is then amplified and propagated to the macroeconomy depending on aggregate bank credit. The work of Gorton and Ordóñez (2014) and Dang, Gorton and Holmström (2020) argues that such a shift in beliefs occurs because financial sector information is hidden, by design, during normal periods, and a crisis is the event when negative information comes to light. The shift from no-information to information is at the heart of their narrative of crises. The work of Bordalo, Gennaioli and Shleifer (2018) has instead argued that a sharp shift in beliefs in a crisis reflects a change from over-optimistic to over-pessimistic beliefs. Extrapolative expectations are at the heart of their narrative of the belief shift in a crisis. In both of these narratives, the pre-cursor to a crisis is a period where agents’ perceive risk to be low, either because risk is hidden or because it is misperceived. Our work suggests that either of these narratives fit the variation in beliefs over the crisis cycle as needed to match the crisis cycle facts. Indeed, it is likely that other models of belief fluctuation such as Kozłowski, Veldkamp and Venkateswaran (2020) where agents update their models of tail risk based on the realization of tail risk can likely also be used to address the macro crisis-cycle facts.

There are a number of directions to take this research. The crises in our model center around banks and banks’ levered holdings of capital. There is evidence that fluctuations in some forms of capital, such as housing capital and capital in the non-tradable sector, have more explanatory power for financial crises (Jordà, Schularick and Taylor, 2015; Müller and Verner, 2020). There is also evidence that credit growth in banks, firms and households

matter differently for crisis outcomes ([Mian, Sufi and Verner, 2017](#)). Expanding the model to incorporate these elements can shed light on the mechanisms underlying these data patterns. Our analysis has also sidestepped normative issues, and particularly how a planner may act differently across the belief models. The model falls into the class of models where an asset price enters a financial constraint. In these models there is an externality that can motivate policy intervention ([Lorenzoni, 2008](#); [Dávila and Korinek, 2018](#)). In the distorted belief model, we also need to consider whether the planner is omniscient or uses agents' distorted beliefs in computing welfare. This gives rise to additional considerations as outlined in the theoretical analysis of [Dávila and Walther \(2020\)](#). While we do not tackle these issues, given a set of welfare assumptions, our quantitative framework can be utilized to assess the magnitude and efficacy of policy interventions. We leave this topic for future research.

Figures and Tables

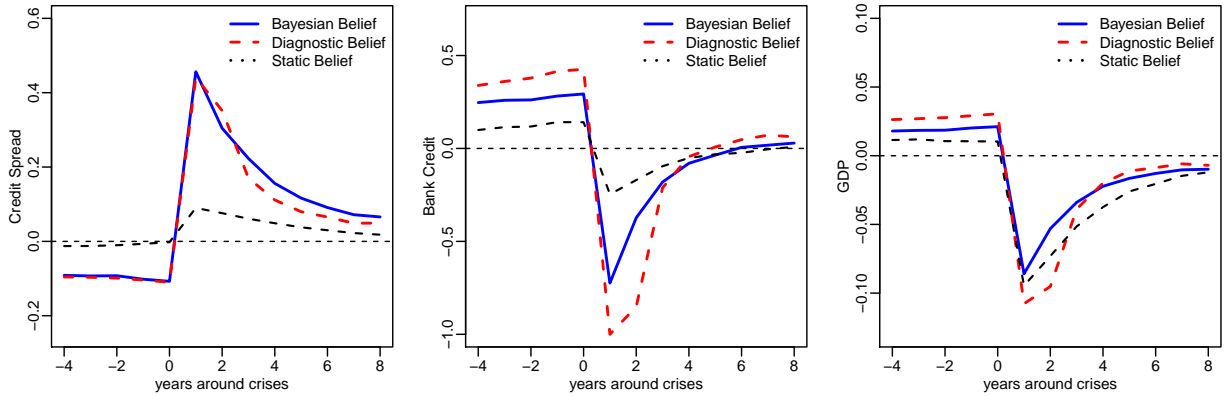
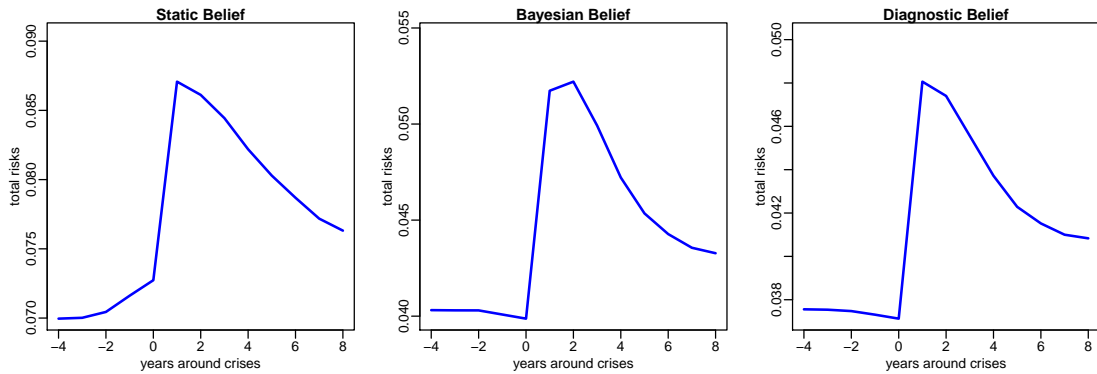
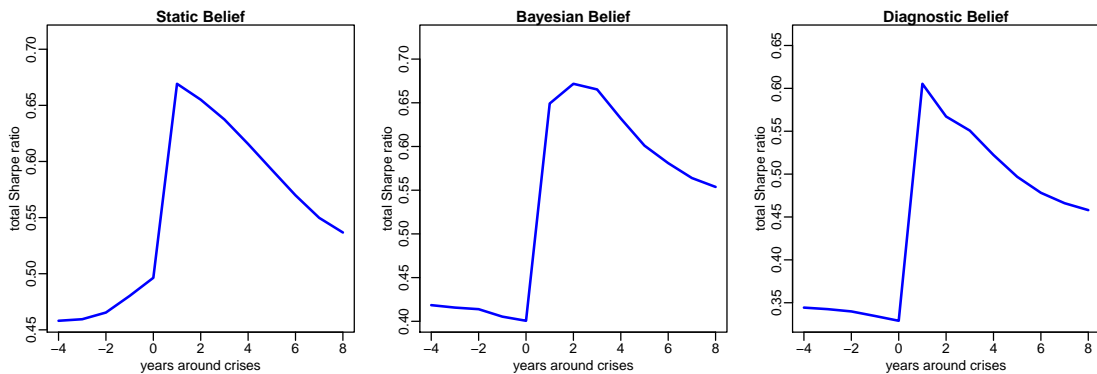


Figure 6: Dynamics of Different Models Around Crises. Credit spread and bank credit are measured as standard-deviations from the mean value. For example, credit spread raising to 0.2 means that it is larger than the value at year 0 by 0.2σ s. GDP is measured in terms of deviation from the long-run mean value.

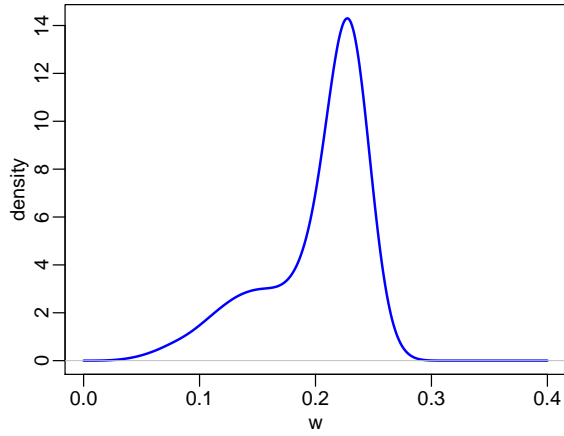


(a) Total Volatility around Crises

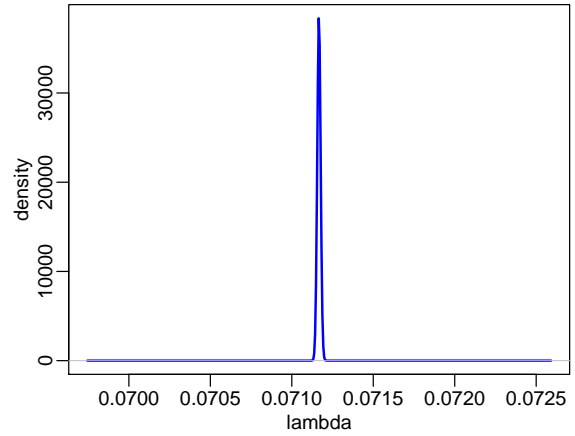


(b) Total Sharpe Ratio around Crises

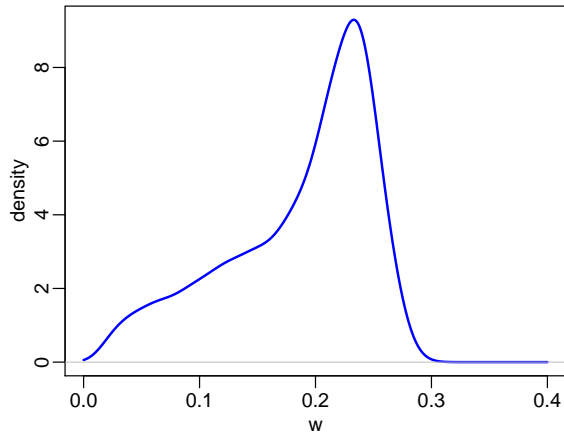
Figure 7: Volatility and Sharpe Ratio. In (a) we plot the total volatility around a crisis, where total volatility is defined in equation (34). In (b) we plot the total Sharpe ratio around a crisis, where total Sharpe ratio is defined as the total risk premium in equation (35) over the total volatility defined in equation (34).



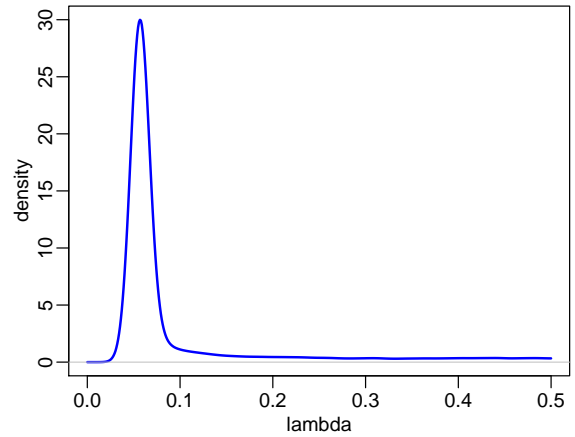
(a) w density of Static Belief model



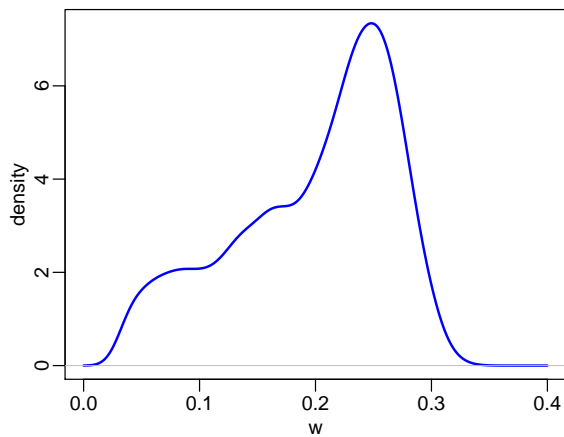
(b) λ density of Static Belief model



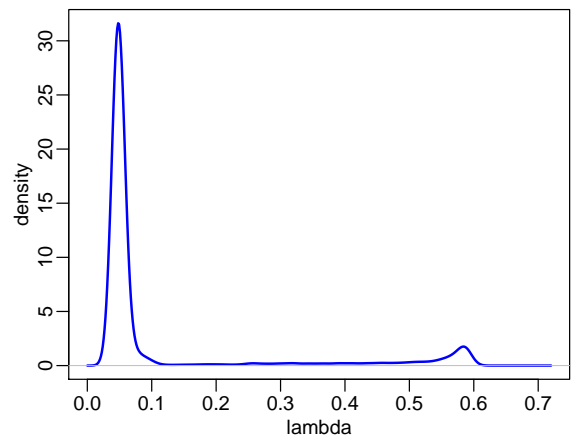
(c) w density of Bayesian model



(d) λ density of Bayesian model



(e) w density of diagnostic model



(f) λ density of diagnostic model

Figure 8: Stationary Distribution of State Variables in Each Model

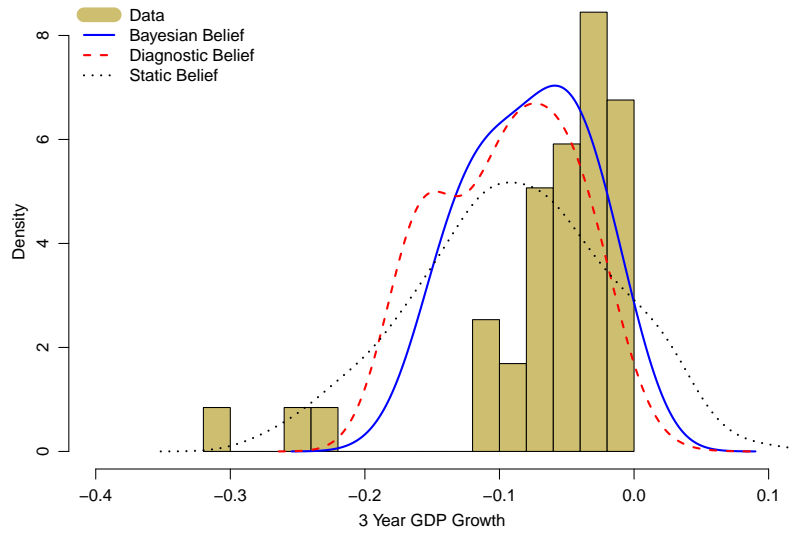


Figure 9: GDP Growth after Crises. We plot the 3-year GDP growth after crises in the model, and peak-to-trough GDP growth in the data around crises.

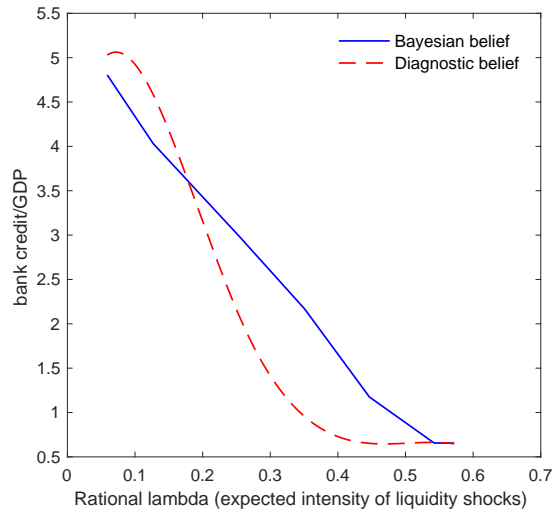
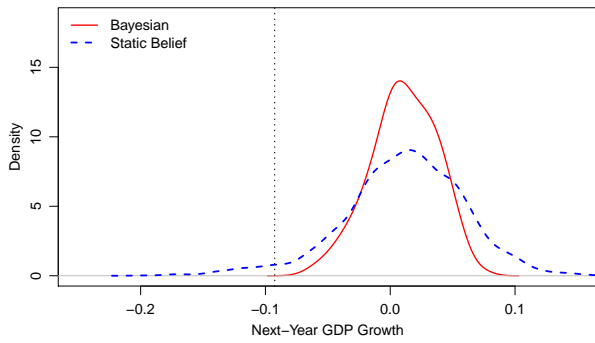
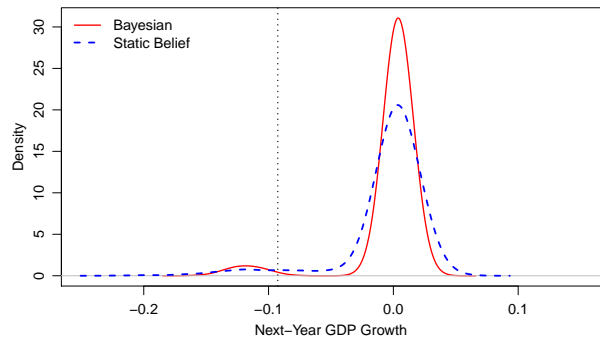


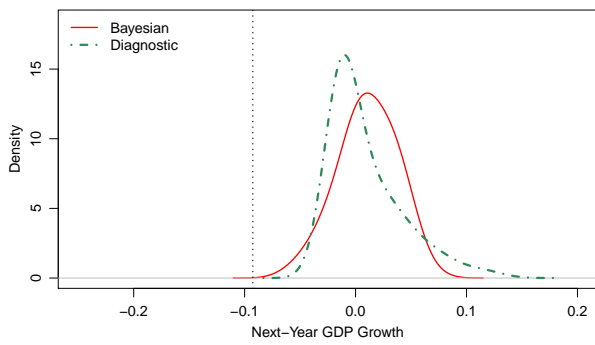
Figure 10: Expected Distress Frequency and Bank Credit/GDP. This figure plots bank credit/GDP as a function of the rational belief λ , given the same state variable w . We simulate the diagnostic model to derive the model-implied relationship between rational λ and the diagnostic belief λ^θ , and show the corresponding bank credit/GDP.



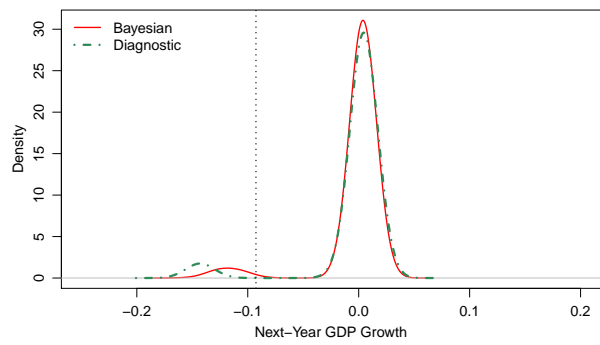
(a) Static Belief v.s. Bayesian: Low Bank Credit



(b) Static Belief v.s. Bayesian: High Bank Credit

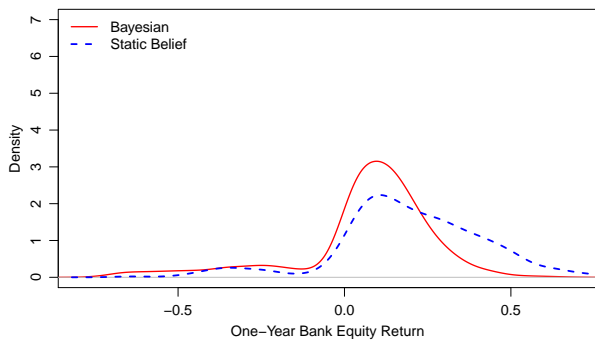


(c) Bayesian v.s. Diagnostic: Low Bank Credit

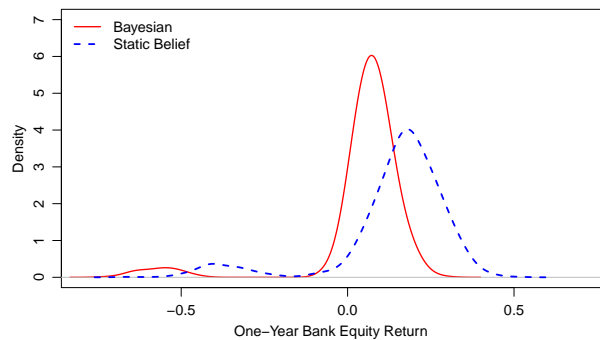


(d) Bayesian v.s. Diagnostic: High Bank Credit

Figure 11: Density of Next-Year GDP Growth Conditional on Bank Credit/GDP. Cutoffs are 30% quantile and 90% quantile of bank credit/GDP.



(a) Static Belief v.s. Bayesian: Low Bank Credit



(b) Static Belief v.s. Bayesian: High Bank Credit

Figure 12: Density of Bank Equity Returns Conditional on Bank Credit/GDP. Cutoffs are 30% quantile and 90% quantile of bank credit/GDP.

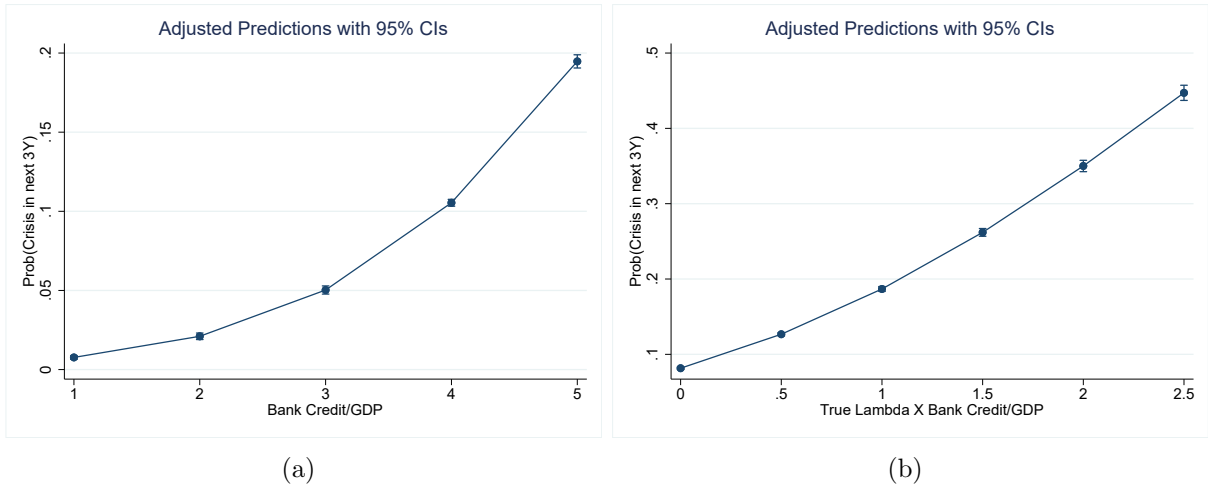


Figure 13: Bayesian Model, Probability of Crisis over next 3 years, by Quintile

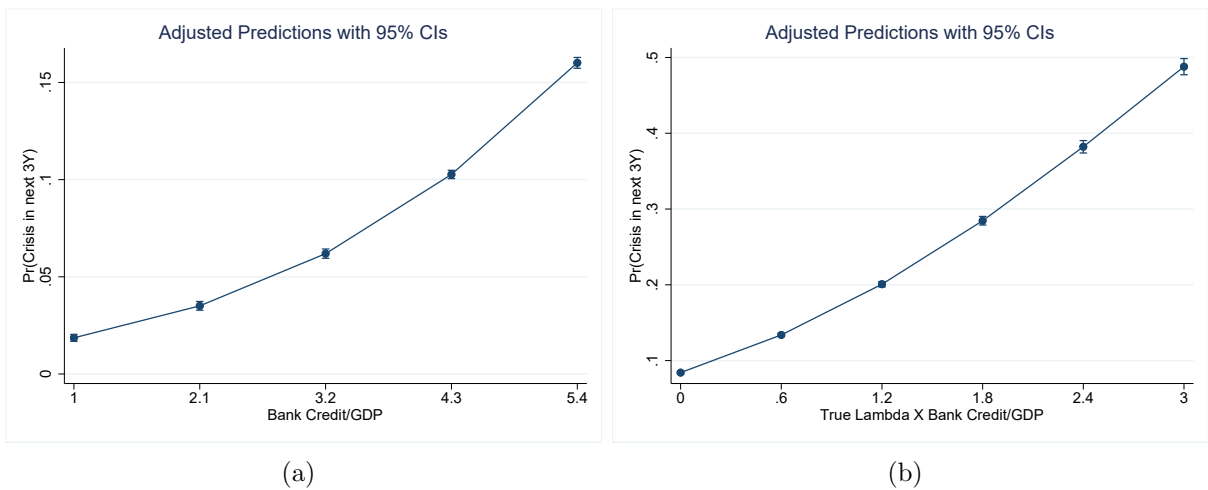
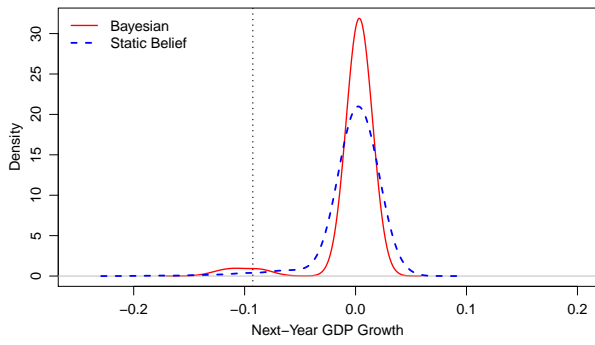
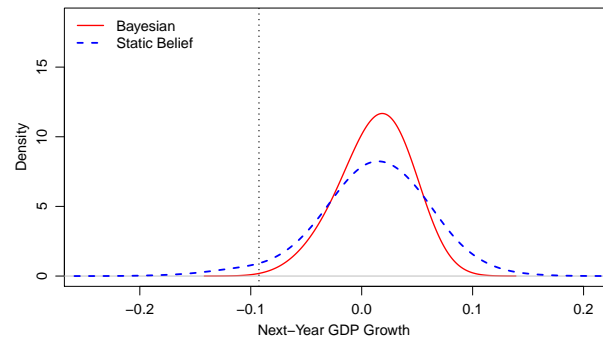


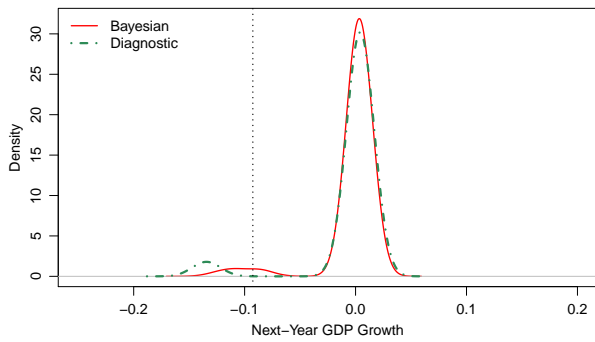
Figure 14: Diagnostic Model, Probability of Crisis over next 3 years, by Quintile



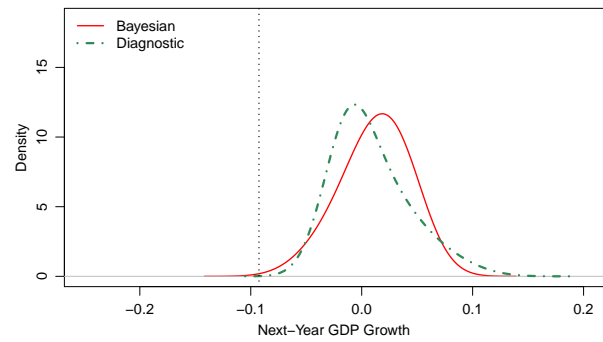
(a) Static Belief v.s. Bayesian: Low Credit Spread



(b) Static Belief v.s. Bayesian: High Credit Spread



(c) Bayesian v.s. Diagnostic: Low Credit Spread



(d) Bayesian v.s. Diagnostic: High Credit Spread

Figure 15: Density of Next-Year GDP Growth in Bayesian and Diagnostic Models Conditional on Credit Spread. Cutoffs are 30% quantile and 90% quantile of bank credit/GDP.

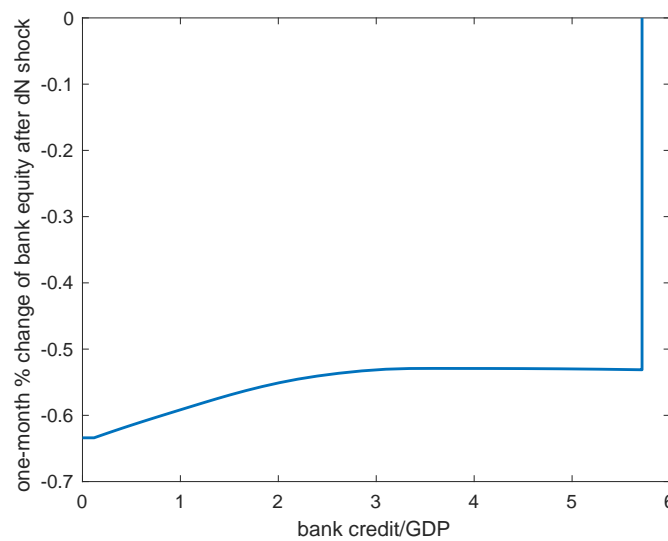


Figure 16: Bank Credit/GDP and Next-Month Bank Equity Returns in the Static-Belief Model

Table 3: Model Simulation and Data: Non-targeted Moments

Panel A: Credit Spread, Bank Credit, and Crisis Severity

	<i>Dependent variable: GDP Growth from t to $t + 3$</i>							
	Static Belief		Bayesian		Diagnostic		Data	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta \text{credit spread}_t * \text{crisis}_t$	-4.88		-2.87		-3.44		-2.11	
$(\frac{\text{bank credit}}{\text{GDP}})_t * \text{crisis}_t$		-0.98		-2.18		-3.49		-2.06
Observations							821	821

Note: Model and data regressions are normalized so that the coefficients reflect the impact of one sigma change in spreads and bank credit/GDP.

Panel B: Bank Credit and Risk Premia

	<i>Dependent variable: Average realized excess return$_{t+1}$</i>						
	Static Belief		Bayesian		Diagnostic		Data
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$(\frac{\text{bank credit}}{\text{GDP}})_t$	-0.01		-0.01		-0.02		-0.02
Observations							867

Note: Model excess return is defined as the return to capital minus the risk-free rate. Data excess return is from Online Appendix of Baron and Xiong (2017) (Table 3, column 1 of Panel B). To ensure comparability, the model return to capital has been normalized to equal the standard deviation of returns reported by Baron and Xiong (2017).

Panel C: Credit Spread Before Crises

	<i>Dependent variable: credit spread$_t$</i>						
	Static Belief		Bayesian		Diagnostic		Data
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
pre-crisis indicator	0.25		-0.25		-0.30		-0.44
Observations							802

Note: regression is: $s_t = \alpha + \beta \cdot 1\{t \text{ is before a crisis}\} + \text{controls}$. For the model, “pre-crisis” is defined as within 1 year before the next crisis. For the data, “pre-crisis” is defined as within 5 years before the next crisis. For both model and data, controls include an indicator of within 5 years after the last crisis. The data regression has more controls such as country fixed effect.

Panel D: Predicting Crises

	<i>Dependent variable: Probability of Crisis</i>							
	Static Belief		Bayesian		Diagnostic		Data	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Froth $_t \rightarrow$ crisis(next 3 years)	-5.94		5.67		7.40		12.90	
$(\frac{\text{bank credit}}{\text{GDP}})_t \rightarrow$ crisis(next year)		0.13		4.05		3.85		2.11
Observations							604	1272

Note: Froth in the model measures if the credit spread is below the median at date t . In the data regression, froth measures if credit spread is below the median over $t - 5$ to t (see [Krishnamurthy and Muir \(2024\)](#)). In both model and data we run a Logit regression of crisis occurring over the next 3 years on the froth measure and report the probability. Bank credit/GDP is the current ratio of bank credit over GDP. The data regression of crisis over the next year on bank credit/GDP is from [Schularick and Taylor \(2012\)](#), and we report the probability of the crisis based on the Logit regression of Table 2.

Table 4: Using Bank Equity Crash to Define a Crisis

Panel A: Crisis, Bank Credit and Severity

	<i>Dependent variable: GDP Growth from t to $t + 3$</i>				
	Bayesian		Diagnostic		Data
	(1)	(2)	(3)	(4)	(5)
crisis_t	-8.24		-10.04		-4.50
$(\frac{\text{bank credit}}{\text{GDP}})_t * \text{crisis}_t$		-2.18		-3.49	
Observations					2548

Note: Model and data regressions are normalized so that the coefficients reflect the impact of one sigma change in spreads, and bank credit/GDP. The coefficient in column (5) is from Table I (column 4) of [Baron, Verner and Xiong \(2021\)](#).

Panel B: Predicting Equity Crashes

	<i>Dependent variable: equity crash from $t + 1$ to $t + 3$</i>		
	Bayesian	Diagnostic	Data
	(1)	(2)	(3)
$(\frac{\text{bank credit}}{\text{GDP}})_t$	9.26	9.60	5.40
Observations			316

Note: The coefficient on Bank Credit/GDP is the sensitivity of crisis probability (%) to a one standard deviation increase in bank credit/GDP. The data regression is from Table III (column 7) of [Baron and Xiong \(2017\)](#).

References

- Adrian, Tobias, and Markus K Brunnermeier.** 2016. “CoVaR.” The American Economic Review, 106(7): 1705.
- Baron, Matthew, and Wei Xiong.** 2017. “Credit expansion and neglected crash risk.” The Quarterly Journal of Economics, 132(2): 713–764.
- Baron, Matthew, Emil Verner, and Wei Xiong.** 2021. “Banking crises without panics.” The Quarterly Journal of Economics, 136(1): 51–113.
- Barro, Robert J.** 2006. “Rare disasters and asset markets in the twentieth century.” The Quarterly Journal of Economics, 121(3): 823–866.
- Boissay, Frédéric, Fabrice Collard, and Frank Smets.** 2016. “Booms and banking crises.” Journal of Political Economy, 124(2): 489–538.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer.** 2018. “Diagnostic expectations and credit cycles.” The Journal of Finance, 73(1): 199–227.
- Bordalo, Pedro, Nicola Gennaioli, Andrei Shleifer, and S Terry.** 2019a. “Real Credit Cycles.” Harvard University.
- Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer.** 2019b. “Diagnostic expectations and stock returns.” The Journal of Finance, 74(6): 2839–2874.
- Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer.** 2020. “Overreaction in macroeconomic expectations.” American Economic Review, 110(9): 2748–82.
- Bordo, Michael, Barry Eichengreen, Daniela Klingebiel, and Maria Soledad Martinez-Peria.** 2001. “Is the crisis problem growing more severe?” Economic policy, 16(32): 52–82.
- Bordo, Michael D, and Christopher M Meissner.** 2016. “Fiscal and financial crises.” In Handbook of macroeconomics. Vol. 2, 355–412. Elsevier.
- Borio, Claudio EV, and Philip William Lowe.** 2002. “Asset prices, financial and monetary stability: exploring the nexus.”
- Brunnermeier, Markus K, and Yuliy Sannikov.** 2014. “A macroeconomic model with a financial sector.” American Economic Review, 104(2): 379–421.
- Camous, Antoine, and Alejandro Van der Ghote.** 2021. “Financial Cycles under Diagnostic Beliefs.” Available at SSRN 3786680.
- Cerra, Valerie, and Sweta Chaman Saxena.** 2008. “Growth Dynamics: The Myth of Economic Recovery.” The American Economic Review, 98(1): 439–457.
- Chen, Long, Pierre Collin-Dufresne, and Robert S Goldstein.** 2008. “On the relation between the credit spread puzzle and the equity premium puzzle.” The Review of Financial Studies, 22(9): 3367–3409.
- Claessens, Stijn, M Ayhan Kose, and Marco E Terrones.** 2010. “The global financial crisis: How similar? How different? How costly?” Journal of Asian Economics, 21(3): 247–264.

- Dang, Tri Vi, Gary Gorton, and Bengt Holmström.** 2020. “The information view of financial crises.” Annual Review of Financial Economics, 12: 39–65.
- Danielsson, Jon, Marcela Valenzuela, and Ilknur Zer.** 2018. “Learning from history: Volatility and financial crises.” The Review of Financial Studies, 31(7): 2774–2805.
- Dávila, Eduardo, and Ansgar Walther.** 2020. “Prudential Policy with Distorted Beliefs.” Available at SSRN 3694722.
- Dávila, Eduardo, and Anton Korinek.** 2018. “Pecuniary externalities in economies with financial frictions.” The Review of Economic Studies, 85(1): 352–395.
- Diamond, Douglas W, and Philip H Dybvig.** 1983. “Bank runs, deposit insurance, and liquidity.” Journal of political economy, 91(3): 401–419.
- Di Tella, Sebastian.** 2017. “Uncertainty shocks and balance sheet recessions.” Journal of Political Economy, 125(6): 2038–2081.
- Fajgelbaum, Pablo D, Edouard Schaal, and Mathieu Taschereau-Dumouchel.** 2017. “Uncertainty traps.” The Quarterly Journal of Economics, 132(4): 1641–1692.
- Farboodi, Maryam, and Péter Kondor.** 2020. “Rational Sentiments and Economic Cycles.”
- Fernald, John G.** 2015. “Productivity and Potential Output before, during, and after the Great Recession.” NBER macroeconomics annual, 29(1): 1–51.
- Gabaix, Xavier.** 2012. “Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance.” The Quarterly Journal of Economics, 127(2): 645–700.
- Geanakoplos, John.** 2010. “The leverage cycle.” NBER macroeconomics annual, 24(1): 1–66.
- Gertler, Mark, and Nobuhiro Kiyotaki.** 2010. “Financial intermediation and credit policy in business cycle analysis.” In Handbook of monetary economics. Vol. 3, 547–599. Elsevier.
- Gertler, Mark, and Nobuhiro Kiyotaki.** 2015. “Banking, liquidity, and bank runs in an infinite horizon economy.” American Economic Review, 105(7): 2011–43.
- Gertler, Mark, Nobuhiro Kiyotaki, and Andrea Prestipino.** 2020. “A macroeconomic model with financial panics.” The Review of Economic Studies, 87(1): 240–288.
- Gilchrist, Simon, and Econ Zakrajsek.** 2012. “Credit Spreads and Business Cycle Fluctuations.” American Economic Review.
- Gopalakrishna, Goutham.** 2020. “A Macro-Finance model with Realistic Crisis Dynamics.” SSRN Electronic Journal.
- Gorton, Gary, and Guillermo Ordonez.** 2014. “Collateral crises.” American Economic Review, 104(2): 343–78.
- Gorton, Gary, and Guillermo Ordonez.** 2020. “Good booms, bad booms.” Journal of the European Economic Association, 18(2): 618–665.

- Gourinchas, Pierre-Olivier, Rodrigo Valdes, Oscar Landerretche, et al.** 2001. “Lending Booms: Latin America and the World.” Economía Journal, 1(Spring 2001): 47–100.
- Greenwood, Robin, Samuel G Hanson, and Lawrence J Jin.** 2019. “Reflexivity in credit markets.” National Bureau of Economic Research.
- Greenwood, Robin, Samuel G Hanson, Andrei Shleifer, and Jakob Ahm Sørensen.** 2020. “Predictable financial crises.” National Bureau of Economic Research.
- He, Zhiguo, and Arvind Krishnamurthy.** 2013. “Intermediary asset pricing.” American Economic Review, 103(2): 732–70.
- He, Zhiguo, and Arvind Krishnamurthy.** 2019. “A macroeconomic framework for quantifying systemic risk.” American Economic Journal: Macroeconomics, 11(4): 1–37.
- Holmstrom, Bengt, and Jean Tirole.** 1997. “Financial intermediation, loanable funds, and the real sector.” the Quarterly Journal of economics, 112(3): 663–691.
- Jermann, Urban, and Vincenzo Quadrini.** 2012. “Macroeconomic effects of financial shocks.” American Economic Review, 102(1): 238–271.
- Jordà, Òscar, Moritz Schularick, and Alan M Taylor.** 2011. “Financial crises, credit booms, and external imbalances: 140 years of lessons.” IMF Economic Review, 59(2): 340–378.
- Jordà, Òscar, Moritz Schularick, and Alan M Taylor.** 2013. “When credit bites back.” Journal of Money, Credit and Banking, 45(s2): 3–28.
- Jordà, Òscar, Moritz Schularick, and Alan M Taylor.** 2015. “Betting the house.” Journal of international economics, 96: S2–S18.
- Judd, Kenneth L, Lilia Maliar, Serguei Maliar, and Rafael Valero.** 2014. “Smolyak method for solving dynamic economic models: Lagrange interpolation, anisotropic grid and adaptive domain.” Journal of Economic Dynamics and Control, 44: 92–123.
- Kaplan, Greg, Kurt Mitman, and Giovanni L Violante.** 2020. “The housing boom and bust: Model meets evidence.” Journal of Political Economy, 128(9): 3285–3345.
- Kindelberger, Charles P.** 1978. Manias, Panics, and Crashes: A History of Financial Crises.
- Kiyotaki, Nobuhiro, and John Moore.** 1997. “Credit Cycles.” Journal of Political Economy, 105(2): pp. 211–248.
- Kozlowski, Julian, Laura Veldkamp, and Venky Venkateswaran.** 2020. “The tail that wags the economy: Beliefs and persistent stagnation.” Journal of Political Economy, 128(8): 000–000.
- Krishnamurthy, Arvind, and Annette Vissing-Jorgensen.** 2015. “The impact of Treasury supply on financial sector lending and stability.” Journal of Financial Economics, 118(3): 571–600.
- Krishnamurthy, Arvind, and Tyler Muir.** 2024. “How credit cycles across a financial crisis.”

- Laeven, Luc, and Fabian Valencia.** 2013. “Systemic banking crises database.” IMF Economic Review, 61(2): 225–270.
- Liptser, Robert, and Albert N Shiryaev.** 2013. Statistics of Random Processes: II. Applications. Vol. 5, Springer Science & Business Media.
- Liu, Xuewen, Pengfei Wang, and Zhongchao Yang.** 2020. “Delayed Crises and Slow Recoveries.” Available at SSRN 3558492.
- Li, Wenhao.** 2019. “Public Liquidity and Financial Crises.” Available at SSRN 3175101.
- Lorenzoni, Guido.** 2008. “Inefficient credit booms.” The Review of Economic Studies, 75(3): 809–833.
- Maxted, Peter.** 2019. “A Macro-Finance Model with Sentiment.” Tech. rep.
- Ma, Yueran, Teodora Paligorova, and José-Luis Peydro.** 2021. “Expectations and bank lending.” Work. Pap., Chicago Booth Sch. Bus., Chicago Google Scholar Article Location.
- Mian, Atif, Amir Sufi, and Emil Verner.** 2017. “Household debt and business cycles worldwide.” The Quarterly Journal of Economics, 132(4): 1755–1817.
- Minsky, Hyman P.** 1992. “The financial instability hypothesis.” The Jerome Levy Economics Institute Working Paper, , (74).
- Moreira, Alan, and Alexi Savov.** 2017. “The macroeconomics of shadow banking.” The Journal of Finance, 72(6): 2381–2432.
- Muir, Tyler.** 2017. “Financial crises and risk premia.” The Quarterly Journal of Economics, 132(2): 765–809.
- Müller, Karsten, and Emil Verner.** 2020. “Credit Allocation and Macroeconomic Fluctuations.”
- Reinhart, Carmen M., and Kenneth S. Rogoff.** 2009a. “The Aftermath of Financial Crises.” American Economic Review, 99(2): 466–72.
- Reinhart, Carmen M., and Kenneth S. Rogoff.** 2009b. This time is different: Eight centuries of financial folly. Princeton, NJ:Princeton University Press.
- Schularick, Moritz, and Alan M Taylor.** 2012. “Credit booms gone bust: Monetary policy, leverage cycles, and financial crises, 1870-2008.” American Economic Review, 102(2): 1029–61.
- Simsek, Alp.** 2013. “Belief disagreements and collateral constraints.” Econometrica, 81(1): 1–53.
- Sufi, Amir, and Alan M Taylor.** 2022. “Financial crises: A survey.” Handbook of International Economics, 6: 291–340.
- Taylor, Alan M.** 2015. “Credit, financial stability, and the macroeconomy.” Annu. Rev. Econ., 7(1): 309–339.

Van Nieuwerburgh, Stijn, and Laura Veldkamp. 2006. “Learning asymmetries in real business cycles.” Journal of monetary Economics, 53(4): 753–772.

Wachter, Jessica A. 2012. “Can time-varying risk of rare disasters explain aggregate stock market volatility?” The Journal of Finance, forthcoming.

A Model Solutions

A.1 Proof of Lemma 1

We will derive the Bayesian belief process λ_t in two different ways. The first method is by applying the theorem in [Liptser and Shiryaev \(2013\)](#). The second one is by taking the continuous-time limit of a discrete-time process. The reason that we show the second method is because we will use the connection between discrete-time and continuous-time processes to prove the results for the diagnostic belief in [Lemma 2](#).

Method 1

We can represent the Poisson process of bank-run as

$$N_t = \int_0^t \mathbf{1}_{\tilde{\lambda}_s = \lambda_L} dN_t^L + \int_0^t \mathbf{1}_{\tilde{\lambda}_s = \lambda_H} dN_t^H = A_t + M_t$$

where N_t^H and N_t^L are two independent Poisson processes, M_t is a martingale, and A_t is a previsible process

$$A_t = \int_0^t (\mathbf{1}_{\tilde{\lambda}_s = \lambda_L} \lambda_L + \mathbf{1}_{\tilde{\lambda}_s = \lambda_H} \lambda_H) dt$$

Denote $\mathcal{F}_t^N = \sigma\{N_s, 0 \leq s \leq t\}$, $\tilde{\theta} = \mathbf{1}_{\tilde{\lambda}_t = \lambda_H}$, and

$$\theta_t = E[\tilde{\theta}_t | \mathcal{F}_t^N] = P(\tilde{\lambda}_t = \lambda^H | \mathcal{F}_t^N)$$

Then according to Theorem 18.3 of [Liptser and Shiryaev \(2013\)](#), the compensator of N_t that is measurable with respect to \mathcal{F}_t^N is

$$\bar{A}_t = \int_0^t E[(\mathbf{1}_{\tilde{\lambda}_s = \lambda^L} \lambda_L + \mathbf{1}_{\tilde{\lambda}_s = \lambda^H} \lambda_H) | \mathcal{F}_{s-}^N] ds = \int_0^t ((1 - \theta_{s-}) \lambda_L + \theta_{s-} \lambda_H) ds$$

Moreover, the compensator of θ_t is

$$\int_0^t (\mathbf{1}_{\tilde{\lambda}_s = \lambda_H} (-\lambda_{H \rightarrow L}) + \mathbf{1}_{\tilde{\lambda}_s = \lambda_L} \lambda_{L \rightarrow H}) ds$$

and the \mathcal{F}_{t-}^N measurable version is

$$\int_0^t (\theta_{s-} (-\lambda_{H \rightarrow L}) + (1 - \theta_{s-}) \lambda_{L \rightarrow H}) ds$$

Finally, the martingale component of $\tilde{\theta}_t$ is independent from the jumps in N_t . Thus we can

apply Theorem 19.6 of [Liptser and Shiryaev \(2013\)](#) to get

$$\begin{aligned}
d\theta_t &= (\theta_{t-}(-\lambda_{H \rightarrow L}) + (1 - \theta_{t-})\lambda_{L \rightarrow H}) dt + E[\tilde{\lambda}_t(\frac{dA_t}{d\bar{A}_t} - 1)|\mathcal{F}_{t-}^N]d(N_t - \bar{A}_t) \\
&= (\theta_{t-}(-\lambda_{H \rightarrow L}) + (1 - \theta_{t-})\lambda_{L \rightarrow H}) dt + E[\mathbf{1}_{\tilde{\lambda}_t=\lambda_H}(\frac{\mathbf{1}_{\tilde{\lambda}_t=\lambda_L}\lambda_L + \mathbf{1}_{\tilde{\lambda}_t=\lambda_H}\lambda_H}{(1 - \theta_{t-})\lambda_L + \theta_{t-}\lambda_H} - 1)|\mathcal{F}_{t-}^N](dN_t - ((1 - \theta_{t-})\lambda_L + \theta_{t-}\lambda_H)dt) \\
&= (\theta_{t-}(-\lambda_{H \rightarrow L}) + (1 - \theta_{t-})\lambda_{L \rightarrow H}) dt + \frac{\theta_{t-}(1 - \theta_{t-})(\lambda_H - \lambda_L)}{(1 - \theta_{t-})\lambda_L + \theta_{t-}\lambda_H}(dN_t - ((1 - \theta_{t-})\lambda_L + \theta_{t-}\lambda_H)dt) \\
&= (\theta_{t-}(-\lambda_{H \rightarrow L}) + (1 - \theta_{t-})\lambda_{L \rightarrow H} - \theta_{t-}(1 - \theta_{t-})(\lambda_H - \lambda_L)) dt + \frac{\theta_{t-}(1 - \theta_{t-})(\lambda_H - \lambda_L)}{(1 - \theta_{t-})\lambda_L + \theta_{t-}\lambda_H}dN_t
\end{aligned}$$

Denote $\lambda_t = E[\tilde{\lambda}_t|\mathcal{F}_t^N]$. We can get the motion of λ_t from

$$\begin{aligned}
\lambda_t &= E[\mathbf{1}_{\tilde{\lambda}_t=\lambda_H}|\mathcal{F}_t^N]\lambda_H + E[\mathbf{1}_{\tilde{\lambda}_t=\lambda_L}|\mathcal{F}_t^N]\lambda_L \\
\Rightarrow \theta_t &= \frac{\lambda_t - \lambda_L}{\lambda_H - \lambda_L}
\end{aligned}$$

which results in

$$d\lambda_t = \begin{pmatrix} (\lambda_L - \lambda_{t-})\lambda_{H \rightarrow L} + (\lambda_H - \lambda_{t-})\lambda_{L \rightarrow H} \\ -(\lambda_{t-} - \lambda_L)(\lambda_H - \lambda_{t-}) \end{pmatrix} dt + \frac{(\lambda_{t-} - \lambda_L)(\lambda_H - \lambda_{t-})}{\lambda_{t-}}dN_t$$

Method 2

Consider a discrete-time Markov process $\tilde{\lambda}_k$ with two states λ_H and λ_L . We define $\Delta t * \tilde{\lambda}_k$ as the probability of a financial distress shock within a single period. The transition probability from high to low is $\lambda_{H \rightarrow L}\Delta t$, and the transition probability from low to high is $\lambda_{L \rightarrow H}\Delta t$. We note that as $\Delta t \rightarrow 0$, this discrete-time Markov chain converges to the continuous-time Markov chain in our main model.

Agents observe the realizations of financial distress shocks, and update their beliefs. Denote the crash realization process as $N_k \in \{0, 1\}$, and the filtration as $\mathcal{F}_k = \sigma\{N_1, N_2, \dots, N_k\}$. Denote the updated belief at period k as $\lambda_k = \mathbb{E}[\tilde{\lambda}_k|\mathcal{F}_k]$, with $\tilde{\lambda}_k$ the state of the hidden Markov process. In each period, the financial distress shock first realizes, and then the agent updates belief for that period.

Suppose that the belief on the probability at high state λ_H is π_k at period k . Then the relationship between π_k and λ_k is as follows:

$$\lambda_k = \pi_k\lambda_H + (1 - \pi_k)\lambda_L$$

Observing $N_{k+1} = n_k \in \{0, 1\}$, the belief π_{k+1} is

$$\begin{aligned} \pi_{k+1} &= P(\tilde{\lambda}_{k+1} = \lambda_H | N_{k+1} = n_{k+1}, \pi_k) \\ &= \frac{P(N_{k+1} = n_{k+1} | \tilde{\lambda}_{k+1} = \lambda_H, \pi_k) P(\tilde{\lambda}_{k+1} = \lambda_H | \pi_k)}{P(N_{k+1} = n_{k+1} | \tilde{\lambda}_{k+1} = \lambda_H, \pi_k) P(\tilde{\lambda}_{k+1} = \lambda_H | \pi_k) + P(N_{k+1} = n_{k+1} | \tilde{\lambda}_{k+1} = \lambda_L, \pi_k) P(\tilde{\lambda}_{k+1} = \lambda_L | \pi_k)} \end{aligned}$$

Note that the probabilities $P(\tilde{\lambda}_{k+1} = \lambda_H | \pi_k)$ and $P(\tilde{\lambda}_{k+1} = \lambda_L | \pi_k)$ can be calculated from the Markov one-step transition

$$\begin{pmatrix} \pi_k \\ 1 - \pi_k \end{pmatrix}^T \begin{pmatrix} 1 - \lambda_{H \rightarrow L} \Delta t & \lambda_{H \rightarrow L} \Delta t \\ \lambda_{L \rightarrow H} \Delta t & 1 - \lambda_{L \rightarrow H} \Delta t \end{pmatrix} = \begin{pmatrix} \pi_k (1 - \lambda_{H \rightarrow L} \Delta t) + (1 - \pi_k) \lambda_{L \rightarrow H} \Delta t \\ \pi_k \lambda_{H \rightarrow L} \Delta t + (1 - \pi_k) (1 - \lambda_{L \rightarrow H} \Delta t) \end{pmatrix}^T$$

which results in

$$P(\tilde{\lambda}_{k+1} = \lambda_H | \pi_k) = \pi_k (1 - \lambda_{H \rightarrow L} \Delta t) + (1 - \pi_k) \lambda_{L \rightarrow H} \Delta t$$

and

$$P(\tilde{\lambda}_{k+1} = \lambda_L | \pi_k) = \pi_k \lambda_{H \rightarrow L} \Delta t + (1 - \pi_k) (1 - \lambda_{L \rightarrow H} \Delta t)$$

Therefore, the belief π_{k+1} is

$$\pi_{k+1} = \frac{((n_{k+1} \lambda_H \Delta t + (1 - n_{k+1}) (1 - \lambda_H \Delta t)) (\pi_k (1 - \lambda_{H \rightarrow L} \Delta t) + (1 - \pi_k) \lambda_{L \rightarrow H} \Delta t))}{\left(\begin{aligned} &(n_{k+1} \lambda_H \Delta t + (1 - n_{k+1}) (1 - \lambda_H \Delta t)) (\pi_k (1 - \lambda_{H \rightarrow L} \Delta t) + (1 - \pi_k) \lambda_{L \rightarrow H} \Delta t) \\ &+ (n_{k+1} \lambda_L \Delta t + (1 - n_{k+1}) (1 - \lambda_L \Delta t)) (\pi_k \lambda_{H \rightarrow L} \Delta t + (1 - \pi_k) (1 - \lambda_{L \rightarrow H} \Delta t)) \end{aligned} \right)}$$

Now it is easier to separately discuss $n_{k+1} = 0$ and $n_{k+1} = 1$. Suppose that no financial distress shock happens ($n_{k+1} = 0$), then we have

$$\pi_{k+1} = \frac{(1 - \lambda_H \Delta t) (\pi_k (1 - \lambda_{H \rightarrow L} \Delta t) + (1 - \pi_k) \lambda_{L \rightarrow H} \Delta t)}{\left(\begin{aligned} &(1 - \lambda_H \Delta t) (\pi_k (1 - \lambda_{H \rightarrow L} \Delta t) + (1 - \pi_k) \lambda_{L \rightarrow H} \Delta t) \\ &+ (1 - \lambda_L \Delta t) (\pi_k \lambda_{H \rightarrow L} \Delta t + (1 - \pi_k) (1 - \lambda_{L \rightarrow H} \Delta t)) \end{aligned} \right)}$$

Suppose that a financial distress shock happens ($n_{k+1} = 1$), then we have

$$\begin{aligned} \pi_{k+1} &= \frac{\lambda_H \Delta t (\pi_k (1 - \lambda_{H \rightarrow L} \Delta t) + (1 - \pi_k) \lambda_{L \rightarrow H} \Delta t)}{\left(\begin{aligned} &\lambda_H \Delta t (\pi_k (1 - \lambda_{H \rightarrow L} \Delta t) + (1 - \pi_k) \lambda_{L \rightarrow H} \Delta t) \\ &+ \lambda_L \Delta t (\pi_k \lambda_{H \rightarrow L} \Delta t + (1 - \pi_k) (1 - \lambda_{L \rightarrow H} \Delta t)) \end{aligned} \right)} \\ &= \frac{\lambda_H (\pi_k (1 - \lambda_{H \rightarrow L} \Delta t) + (1 - \pi_k) \lambda_{L \rightarrow H} \Delta t)}{\left(\begin{aligned} &\lambda_H (\pi_k (1 - \lambda_{H \rightarrow L} \Delta t) + (1 - \pi_k) \lambda_{L \rightarrow H} \Delta t) \\ &+ \lambda_L (\pi_k \lambda_{H \rightarrow L} \Delta t + (1 - \pi_k) (1 - \lambda_{L \rightarrow H} \Delta t)) \end{aligned} \right)} \end{aligned}$$

Note that taking $\Delta t \rightarrow 0$ will result in $\pi_{k+1} = \pi_k$ when $n_{k+1} = 0$. This is reasonable, because this is like calculating $\mu_t dt$ for the λ_t process in continuous time, which is a small

order term. An appropriate way to derive the time limit is to calculate

$$\begin{aligned}
& \lim_{\Delta t \rightarrow 0} \frac{\pi_{k+1} - \pi_k}{\Delta t} \Big|_{n_{k+1}=0, \mathcal{F}_k} \\
&= \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \left(\begin{array}{c} (1 - \lambda_H \Delta t) (\pi_k (1 - \lambda_{H \rightarrow L} \Delta t) + (1 - \pi_k) \lambda_{L \rightarrow H} \Delta t) \\ -\pi_k (1 - \lambda_H \Delta t) (\pi_k (1 - \lambda_{H \rightarrow L} \Delta t) + (1 - \pi_k) \lambda_{L \rightarrow H} \Delta t) \\ -\pi_k (1 - \lambda_L \Delta t) (\pi_k \lambda_{H \rightarrow L} \Delta t + (1 - \pi_k) (1 - \lambda_{L \rightarrow H} \Delta t)) \end{array} \right) \\
&= \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \left(\begin{array}{c} (1 - \pi_k) (1 - \lambda_H \Delta t) (\pi_k (1 - \lambda_{H \rightarrow L} \Delta t) + (1 - \pi_k) \lambda_{L \rightarrow H} \Delta t) \\ -\pi_k (1 - \lambda_L \Delta t) (\pi_k \lambda_{H \rightarrow L} \Delta t + (1 - \pi_k) (1 - \lambda_{L \rightarrow H} \Delta t)) \end{array} \right) \\
&= \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \left(\begin{array}{c} (1 - \pi_k) (\pi_k - \pi_k \lambda_{H \rightarrow L} \Delta t + (1 - \pi_k) \lambda_{L \rightarrow H} \Delta t - \lambda_H \pi_k \Delta t) \\ -\pi_k (\pi_k \lambda_{H \rightarrow L} \Delta t + (1 - \pi_k) (1 - \lambda_{L \rightarrow H} \Delta t) - \lambda_L (1 - \pi_k) \Delta t) \end{array} \right) \text{ (removing } \Delta t^2 \text{ terms)} \\
&= -\pi_k \lambda_{H \rightarrow L} + (1 - \pi_k) \lambda_{L \rightarrow H} - (\lambda_H - \lambda_L) \pi_k (1 - \pi_k)
\end{aligned}$$

Therefore, we have

$$\lim_{\Delta t \rightarrow 0} \frac{\pi_{k+1} - \pi_k}{\Delta t} \Big|_{n_{k+1}=0, \mathcal{F}_k} = -\pi_k \lambda_{H \rightarrow L} + (1 - \pi_k) \lambda_{L \rightarrow H} - (\lambda_H - \lambda_L) \pi_k (1 - \pi_k) \quad (37)$$

To build an exact connection to λ_k , we can write λ_k in terms of π_k as

$$\pi_k = \frac{\lambda_k - \lambda_L}{\lambda_H - \lambda_L} \quad (38)$$

Then the limit of $\Delta t \rightarrow 0$ expressed with λ_k is

$$\frac{1}{\lambda_H - \lambda_L} \frac{\lambda_{k+1} - \lambda_k}{\Delta t} \Big|_{n_{k+1}=0, \mathcal{F}_k} = -\frac{\lambda_k - \lambda_L}{\lambda_H - \lambda_L} \lambda_{H \rightarrow L} + \frac{\lambda_H - \lambda_k}{\lambda_H - \lambda_L} \lambda_{L \rightarrow H} - (\lambda_H - \lambda_L) \frac{\lambda_k - \lambda_L}{\lambda_H - \lambda_L} \frac{\lambda_H - \lambda_k}{\lambda_H - \lambda_L}$$

which can be simplified as

$$\lim_{\Delta t \rightarrow 0} \frac{\lambda_{k+1} - \lambda_k}{\Delta t} \Big|_{n_{k+1}=0, \mathcal{F}_k} = (\lambda_L - \lambda_k) \lambda_{H \rightarrow L} + (\lambda_H - \lambda_k) \lambda_{L \rightarrow H} - (\lambda_k - \lambda_L) (\lambda_H - \lambda_k) \quad (39)$$

Suppose that a financial distress shock happens ($n_{k+1} = 1$). By taking $\Delta t \rightarrow 0$, the updating is

$$\pi_{k+1} \Big|_{n_{k+1}=1, \mathcal{F}_k} = \frac{\lambda_H \pi_k}{\lambda_H \pi_k + \lambda_L (1 - \pi_k)}$$

Using (38), the updating is

$$\begin{aligned}
\frac{1}{\pi_{k+1}} &= 1 + \frac{\lambda_L}{\lambda_H} \frac{1 - \pi_k}{\pi_k} \\
\lambda_{k+1} &= \frac{\lambda_H (\lambda_k - \lambda_L)}{\lambda_k} + \lambda_L = \frac{(\lambda_H + \lambda_L) \lambda_k - \lambda_H \lambda_L}{\lambda_k}
\end{aligned}$$

which implies

$$\lambda_{k+1} - \lambda_k |_{n_{k+1}=1, \mathcal{F}_k} = \frac{(\lambda_H + \lambda_L)\lambda_k - \lambda_H\lambda_L}{\lambda_k} - \lambda_k = \frac{(\lambda_H - \lambda_k)(\lambda_k - \lambda_L)}{\lambda_k}$$

Finally, we express the above with the continuous-time notation dN_t and dt to get

$$d\lambda_t = \begin{pmatrix} (\lambda_L - \lambda_{t-})\lambda_{H \rightarrow L} + (\lambda_H - \lambda_{t-})\lambda_{L \rightarrow H} \\ -(\lambda_{t-} - \lambda_L)(\lambda_H - \lambda_{t-}) \end{pmatrix} dt + \frac{(\lambda_H - \lambda_{t-})(\lambda_{t-} - \lambda_L)}{\lambda_{t-}} dN_t$$

which is the same as method 1.

A.2 Proof of Lemma 2

To prove Lemma 2, we start with discrete time process and then take the continuous-time limit. The discrete-time distress frequency process $\tilde{\lambda}_t$ is the same as Section A.1. Specifically, the process has two states λ_H and λ_L , with transition probability from high to low as $\lambda_{H \rightarrow L}\Delta t$, and the transition probability from low to high as $\lambda_{L \rightarrow H}\Delta t$. Agents observe the realizations of financial distress shocks, and update their beliefs. Denote the crash realization process as $N_k \in \{0, 1\}$, and the filtration as $\mathcal{F}_k = \sigma\{N_1, N_2, \dots, N_k\}$. Denote the updated belief at period k as $\lambda_k = \mathbb{E}[\tilde{\lambda}_k | \mathcal{F}_k]$, with $\tilde{\lambda}_k$ the state of the hidden Markov process. Also denote the probability $\pi_k = P(\tilde{\lambda}_k = \lambda_H)$, which implies

$$\lambda_k = \pi_k \lambda_H + (1 - \pi_k) \lambda_L$$

We choose the period length Δt so that $T(\Delta t) = t_0/\Delta t$ is an integer, where t_0 is the “look-back period” for the diagnostic belief. Then we denote the reference probability for the diagnostic belief at period k as

$$\pi_k^T = P(\tilde{\lambda}_k = \lambda_H | \pi_{k-T(\Delta t)})$$

We already know from method 2 of Section A.1 that when $\Delta t \rightarrow 0$, the continuous-time limit of the Bayesian belief process results in (7). Our task now is to prove that the discrete-time diagnostic belief process converges to a continuous-time process as in (10). By definition, the diagnostic belief at period k is

$$\begin{aligned} \pi_k^\theta &= \pi_k \cdot \left(\frac{\pi_k}{\pi_k^T}\right)^\theta \frac{1}{Z_k} \\ 1 - \pi_k^\theta &= (1 - \pi_k) \cdot \left(\frac{1 - \pi_k}{1 - \pi_k^T}\right)^\theta \frac{1}{Z_k} \end{aligned}$$

with

$$Z_k = \frac{1}{\pi_k \cdot \left(\frac{\pi_k}{\pi_k^T}\right)^\theta + (1 - \pi_k) \cdot \left(\frac{1 - \pi_k}{1 - \pi_k^T}\right)^\theta}$$

which implies

$$\begin{aligned} \pi_k^\theta &= \pi_k \left(\frac{\pi_k}{\pi_k^T}\right)^\theta \frac{1}{\pi_k \left(\frac{\pi_k}{\pi_k^T}\right)^\theta + (1 - \pi_k) \left(\frac{1 - \pi_k}{1 - \pi_k^T}\right)^\theta} \\ &= \pi_k \frac{1}{\pi_k + (1 - \pi_k) \left(\frac{\pi_k^T}{1 - \pi_k^T} / \frac{\pi_k}{1 - \pi_k}\right)^\theta} \end{aligned}$$

Therefore, if $\pi_k^T < \pi_k$, then $\pi_k^\theta > \pi_k$, leading to an overreaction. Now we can replace the probability with λ_t . Define the expected $\tilde{\lambda}_k$ under the diagnostic belief as λ_k^θ . Then we have

$$\lambda_k^\theta - \lambda_L = (\lambda_k - \lambda_L) \frac{(\lambda_H - \lambda_k) + (\lambda_k - \lambda_L)}{\left(\frac{\lambda_k^T - \lambda_L}{\lambda_H - \lambda_k^T} / \frac{\lambda_k - \lambda_L}{\lambda_H - \lambda_k}\right)^\theta (\lambda_H - \lambda_k) + (\lambda_k - \lambda_L)}$$

where

$$\lambda_k^T = \pi_k^T \lambda_H + (1 - \pi_k^T) \lambda_L$$

The key is to derive π_k^T and λ_k^T under the limit of $\Delta t \rightarrow 0$ while keeping $t = k\Delta t$ constant. Using the probability transition matrix, we get

$$\begin{pmatrix} P(\lambda_k = \lambda_H | \pi_k^T) \\ P(\lambda_k = \lambda_L | \pi_k^T) \end{pmatrix}' = \begin{pmatrix} \pi_{k-T} \\ 1 - \pi_{k-T} \end{pmatrix}' \begin{pmatrix} 1 - \lambda_{H \rightarrow L} \Delta t & \lambda_{H \rightarrow L} \Delta t \\ \lambda_{L \rightarrow H} \Delta t & 1 - \lambda_{L \rightarrow H} \Delta t \end{pmatrix}'$$

where the ' notation denotes transpose of a matrix. The limit of the above expression with $\Delta t \rightarrow 0$ is effectively the transition of a continuous time Markov chain, with rate matrix

$$Q = \begin{pmatrix} -\lambda_{H \rightarrow L} & \lambda_{H \rightarrow L} \\ \lambda_{L \rightarrow H} & -\lambda_{L \rightarrow H} \end{pmatrix}$$

A decomposition reveals that the two eigenvalues of this matrix are 0 and $-(a + b)$, where $a = \lambda_{H \rightarrow L}$ and $b = \lambda_{L \rightarrow H}$. The associated eigenvector formed matrix is

$$\bar{Q} = \begin{pmatrix} 1 & -a \\ 1 & b \end{pmatrix}$$

with the inverse

$$\bar{Q}^{-1} = \frac{1}{a + b} \begin{pmatrix} b & a \\ -1 & 1 \end{pmatrix}$$

Then we can decompose

$$Q = \bar{Q} \begin{pmatrix} 0 & \\ & -(a + b) \end{pmatrix} \bar{Q}^{-1}$$

Then the transition for t units of time is

$$\bar{Q} \begin{pmatrix} 1 & \\ & e^{-(a+b)t} \end{pmatrix} \bar{Q}^{-1} = \frac{1}{a+b} \begin{pmatrix} b + ae^{-(a+b)t} & a - be^{-(a+b)t} \\ b - be^{-(a+b)t} & a + be^{-(a+b)t} \end{pmatrix}$$

Using the t notation ($t = k * \Delta t$), and taking the limit $\Delta t \rightarrow 0$ while keeping t unchanged, we have

$$\begin{aligned} \lim_{\Delta t \rightarrow 0} \begin{pmatrix} P(\lambda_k = \lambda_H | \pi_k^T) \\ P(\lambda_k = \lambda_L | \pi_k^T) \end{pmatrix}^T &= \begin{pmatrix} P(\lambda_t = \lambda_H | \pi_{t-t_0}) \\ P(\lambda_t = \lambda_L | \pi_{t-t_0}) \end{pmatrix}^T \\ &= \begin{pmatrix} \pi_{t-t_0} \\ 1 - \pi_{t-t_0} \end{pmatrix}^T \frac{1}{a+b} \begin{pmatrix} b + ae^{-(a+b)t_0} & a - be^{-(a+b)t_0} \\ b - be^{-(a+b)t_0} & a + be^{-(a+b)t_0} \end{pmatrix} \\ &\triangleq \begin{pmatrix} a_H \pi_{t-t_0} + a_L (1 - \pi_{t-t_0}) \\ b_H \pi_{t-t_0} + b_L (1 - \pi_{t-t_0}) \end{pmatrix}^T \end{aligned}$$

where

$$\begin{pmatrix} a_H & b_H \\ a_L & b_L \end{pmatrix} = \frac{1}{a+b} \begin{pmatrix} b + ae^{-(a+b)t_0} & a - ae^{-(a+b)t_0} \\ b - be^{-(a+b)t_0} & a + be^{-(a+b)t_0} \end{pmatrix} \quad (40)$$

Therefore, the intensity process follows

$$\lambda_t^\theta - \lambda_L = (\lambda_t - \lambda_L) \frac{(\lambda_H - \lambda_t) + (\lambda_t - \lambda_L)}{\left(\frac{\lambda_t^T - \lambda_L}{\lambda_H - \lambda_t^T} / \frac{\lambda_t - \lambda_L}{\lambda_H - \lambda_t} \right)^\theta (\lambda_H - \lambda_t) + (\lambda_t - \lambda_L)} \quad (41)$$

where

$$\lambda_t^T - \lambda_L = a_H (\lambda_{t-t_0} - \lambda_L) + a_L (\lambda_H - \lambda_{t-t_0}) \quad (42)$$

$$\lambda_H - \lambda_t^T = b_H (\lambda_{t-t_0} - \lambda_L) + b_L (\lambda_H - \lambda_{t-t_0}) \quad (43)$$

When the total transition rates $a + b$ are low, we have $a_H \approx 1$, $a_L \approx 0$, $b_H \approx 0$, and $b_L \approx 1$. Then we have $\lambda_t^T \approx \lambda_{t-t_0}$. When $\lambda_t^T > \lambda_t$, i.e., the likelihood of a crisis is decreasing, then the subjective probability is even lower, with $\lambda_t^\theta < \lambda_t$. When $\lambda_t^T < \lambda_t$, i.e., the likelihood of a crisis is increasing, then the subjective probability is even higher, with $\lambda_t^\theta > \lambda_t$. These predictions are perfectly consistent with the spirit of the diagnostic expectations. The extent of such extrapolation is larger as θ becomes larger, and we have $\lambda_t^\theta = \lambda_t$ when $\theta = 0$.

A.3 Proof of Lemma 3

To save on notation, we omit the subscripts t and $t-$.

Suppose that in equilibrium, $x^K < 1$. This implies that $(x^d)^+ = (x^K - 1)^+ = 0$, which

leads to the following first order conditions for households and bankers:

$$\mu^R + \frac{\bar{A}}{p} - r^d = (\sigma^K + \sigma^p)^2 x^K + \lambda \kappa^p \frac{1}{1 - x^K \kappa^p}$$

$$\mu^R + \frac{\underline{A}}{p} - r^d = (\sigma^K + \sigma^p)^2 y^K + \lambda \kappa^p \frac{1}{1 - y^K \kappa^p}$$

Subtracting the above two equations, we obtain

$$\frac{\bar{A} - \underline{A}}{p} = \left((\sigma^K + \sigma^p)^2 + \frac{\lambda(\kappa^p)^2}{(1 - x^K \kappa^p)(1 - y^K \kappa^p)} \right) (x^K - y^K) \quad (44)$$

The first bracket on the right hand side is always positive, since the nonnegative wealth constraint implies $x^K \kappa^p < 1$ and $y^K \kappa^p < 1$. However, from market-clearing conditions (24) and (25),

$$w x^K + (1 - w) y^K = 1$$

Under the assumption of $x^K < 1$, we must have

$$y^K > x^K$$

which implies that the right-hand side of (44) should be negative. This is a contradiction since the left-hand side of (44) is positive.

Importantly, all of the above derivations go through regardless of whether we use the Bayesian belief or the diagnostic belief, as long as bankers and households have the same belief.

In summary, we have $x^K \geq 1$ in equilibrium. In other words, bankers borrow from households in the debt market.

A.4 First-Order Conditions

In this section, we derive bank and household first-order conditions. To save on notation, we omit the subscripts t and $t-$.

From equation (27) and a bank's optimization problem in (29) and (28), we obtain the first-order condition over x^K ,

$$\mu^R + \frac{\bar{A}}{p} - r^d = (\sigma^K + \sigma^p)^2 x^K + \lambda \frac{\kappa^p + \alpha}{1 - x^K \kappa^p - \alpha \Delta x} \quad (45)$$

As a result, the excess return on capital is

$$\begin{aligned} E[dR^b] &= \mu^R + \frac{\bar{A}}{p} - r^d - \lambda(\kappa^p + \alpha) \\ &= (\sigma^K + \sigma^p)^2 x^K + \lambda(\alpha + \kappa^p) \frac{x^K \kappa^p + \alpha x^d}{1 - x^K \kappa^p - \alpha x^d} \end{aligned} \quad (46)$$

Using the HJB equation for households combined with household budget dynamics in (18), we obtain the first-order condition over y^K as

$$\mu^R + \frac{\underline{A}}{p} - r^d \leq (\sigma^K + \sigma^p)^2 y^K + \lambda \frac{\kappa^p}{1 - \kappa^h}, \text{ with equality if } y^K > 0 \quad (47)$$

In equation (47), the left hand side is the yield spread on productive capital over bank debt, while the right hand side includes the risk-adjusted losses of productive capital in liquidity shocks. When the yield spread is lower than the cost, households do not hold productive capital and set $y^K = 0$.

Combining (47) and (45), we have

$$\frac{\bar{A} - \underline{A}}{p} \geq (\sigma^K + \sigma^p)^2 (x^K - y^K) + \lambda \frac{\kappa^p + \alpha}{1 - x^K \kappa^p - \alpha \Delta x} - \lambda \frac{\kappa^p}{1 - \kappa^h}$$

where the equality holds when $y^K > 0$.

A.5 Equilibrium Solutions

With log utility, the optimal consumption rule is $\dot{c}^b = \dot{c}^h = \rho$. Then we simplify the equilibrium conditions into the following equations:

$$\rho = \frac{\psi A^H + (1 - \psi) A^L - i}{p}. \quad (48)$$

$$x^K w + y^K (1 - w) = 1. \quad (49)$$

$$\psi = \frac{x^K w}{x^K w + y^K (1 - w)} = x^K w, \quad (50)$$

Next, we derive the dynamics of state variables. We apply Ito's lemma on the definition

of wealth share in (11) and get the dynamics of w as

$$\begin{aligned} \frac{dw}{w} &\triangleq \mu^w dt + \sigma^w dB - \kappa^w dN \\ &= (1-w) \left(\mu^b - \mu^h + (\sigma^h)^2 - \sigma^b \sigma^h - w(\sigma^b - \sigma^h)^2 - \eta \right) dt \\ &\quad + (1-w)(\sigma^b - \sigma^h)dB - (1-w) \frac{1 - \frac{1-\kappa^b}{1-\kappa^h}}{1 + w(\frac{1-\kappa^b}{1-\kappa^h} - 1)} dN. \end{aligned} \quad (51)$$

where all variables in the right hand side should have subscripts $t-$ which we omit. Then we can apply Ito's lemma on price function $p(w)$ to get

$$\begin{cases} \mu^p = p_w w \mu^w + \frac{1}{2} p_{ww} (w \sigma^w)^2 + p_\lambda \mu^\lambda(\lambda) \\ \sigma^p = p_w w (1-w) (\sigma^b - \sigma^h) \\ \kappa^p = 1 - p\left(w \frac{1-\kappa^b}{1-\kappa^h - w(\kappa^b - \kappa^h)}, \lambda\right) / p(w, \lambda). \end{cases} \quad (52)$$

To fully characterize the economy, we also need to know the dynamics of aggregate capital quantity K (although all policy functions are scalable with respect to K). Denote the Ito process for K as

$$\frac{dK}{K} = \mu^K dt - \delta dt + \sigma^K dB, \quad (53)$$

We collect the system of equations for jumps from (6), (15), and (19) as follows:

$$\begin{cases} \kappa^b = x^K \kappa^p + \alpha x^d \\ \kappa^h = y^K \kappa^p - \alpha x^d \frac{w}{1-w} \\ \kappa^p = 1 - p\left(w \frac{1-\kappa^b}{1-\kappa^h - w(\kappa^b - \kappa^h)}, \lambda + \kappa^\lambda(\lambda)\right) / p(w, \lambda) \end{cases} \quad (54)$$

From (15), (18), and (52), we collect the exposure to Brownian shocks as

$$\begin{cases} \sigma^p = p_w w (1-w) (\sigma^b - \sigma^h) \\ \sigma^h = y^K (\sigma^K + \sigma^p) \\ \sigma^b = x^K (\sigma^K + \sigma^p). \end{cases} \quad (55)$$

Diagnostic Beliefs

We solve the model with diagnostic beliefs as follows. As households act as if their beliefs are the true ones, their policy functions are the same as the model with Bayesian beliefs. However, the true (physical) frequency of jumps differ from that of the agents' beliefs. There are two steps to clear the market during a jump with diagnostic beliefs:

- First, the agents interpret λ_t^θ as the Bayesian belief. After a crisis shock dN_t , the market price of capital switches to the level under this "Bayesian belief".

- The realization of belief, however, is different from the Bayesian expectation, because of the diagnostic belief formation. Now additional price adjustment is needed to clear the market under the diagnostic belief.

A.6 Loan Spread

To price the loan spread, denote the bank holding of this risk-free but illiquid loan as x_t^C . Then the equivalent bank optimization problem in (29) will have two additional terms involving x_t^C :

$$\dots + x_{t-}^C (r_{t-}^C - r_{t-}^d) dt - x_{t-}^C \alpha dN_t$$

which implies the FOC on x_{t-}^C as

$$r_{t-}^C - r_{t-}^d = \frac{\lambda_{t-}}{1 - x_{t-}^K \kappa_{t-}^p - \alpha x_{t-}^d - \alpha x_{t-}^C} \alpha.$$

Since the illiquid loan is in zero supply, we use the existing bank SDF to price the loan and we have $x_{t-}^C = 0$ and

$$r_{t-}^C - r_{t-}^d = \frac{\lambda_{t-}}{1 - x_{t-}^K \kappa_{t-}^p - \alpha x_{t-}^d} \alpha.$$

A.7 Credit Spread

In this section, we define the credit spread used in the calibration, derive the jump differential equation for the credit spread and provide the solution methodology.

Define τ as the expected maturity of the bond. We assume that the bond matures based on the realizations of a Poisson event with intensity $1/\tau$. This modeling allows for a simple recursive formulation for bond pricing. Moreover, we suppose that a fraction of the maturity events result in default, while another fraction result in full repayment. In particular, we assume that a bond matures in two cases: (1) conditional on the financial illiquidity dN_t shock, the bond matures with probability π ; (2) conditional on another independent Poisson process dN_t^τ (with intensity λ_t^τ), the bond matures with probability 1. The two intensities sum up to a fixed number, i.e.,

$$\pi \lambda_t + \lambda_t^\tau = 1/\tau \tag{56}$$

where τ can be interpreted as the maturity of the bond. We can see that

$$1/\tau \geq \pi \lambda^H$$

and therefore,

$$\tau \leq \frac{1}{\pi \lambda^H}$$

which is the maximum maturity of bonds that we can define with this method.

Each risky bond has a face value of 1. One unit value of a risky asset is continuously posted to back this risky bond, i.e., the bond is fully collateralized if the bond matures as long as there is no jump in the value of the risky asset. If dN_t hits when the bond matures, the underlying risky asset's value jumps downwards by $m \cdot \kappa_{t-}^p + \hat{\kappa}_0$. The first term varies with economic conditions. It contains capital price drop κ_{t-}^p , and a multiplier m that measures the exposure of the collateral to capital price decline. The second term here a constant “baseline” loss given default. If maturity occurs with no illiquidity event, we assume that the bond pays back in full. Thus, the loss function upon maturity for the risky bond is

$$\hat{\kappa}_t = (m \cdot \kappa_{t-}^p + \hat{\kappa}_0)dN_t \quad (57)$$

This structure gives a time-varying default probability. Specifically, when a bond matures, the probability of default is

$$\frac{\pi \lambda_t}{\pi \lambda_t + \lambda_t^\tau} = \tau \pi \lambda_t \quad (58)$$

Therefore, the unconditional probability of default is $\tau \pi \bar{\lambda}$, where $\bar{\lambda}$ is the unconditional average of the expected illiquidity frequency.

Denote the current market value of this risky bond, priced using the banker's pricing kernel, as $v_t = v(w_t, \lambda_t)$, and the value of the safe bond as \bar{v}_t . Then we define the credit spread as

$$\mathcal{S}_t(p_{t_0}) = \frac{1}{\tau} \log(1/v_t) - \frac{1}{\tau} \log(1/\bar{v}_t) \quad (59)$$

We expect $\mathcal{S}_t \geq 0$, given that risky bonds may default, and default occurs in high marginal utility states. Solving for this credit spread involves solving an endogenous jump equation with second-order derivatives.

HJB Equations

From Ito's lemma, we have

$$\begin{aligned} dv(w, \lambda) &= \frac{\partial v(w, \lambda)}{\partial w} (w \mu^w dt + w \sigma^w dB_t) + \frac{1}{2} \frac{\partial^2 v(w, \lambda)}{\partial w^2} w^2 (\sigma^w)^2 dt \\ &+ \frac{\partial v(w, \lambda)}{\partial \lambda} \mu^\lambda(\lambda) dt + (v(w + \Delta w, \lambda + \Delta \lambda) - v(w, \lambda)) dN_t \end{aligned}$$

Denote

$$\frac{dv(w, \lambda)}{v(w, \lambda)} = \mu^v dt + \sigma^v dB_t - \kappa^v dN_t$$

Matching the coefficients, we have

$$v(w, \lambda)\mu^v = \frac{\partial v(w, \lambda)}{\partial w} w \mu^w + \frac{1}{2} \frac{\partial^2 v(w, \lambda)}{\partial w^2} w^2 (\sigma^w)^2 + \frac{\partial v(w, \lambda)}{\partial \lambda} \mu^\lambda(\lambda)$$

$$v(w, \lambda)\sigma^v = \frac{\partial v(w, \lambda)}{\partial w} w \sigma^w$$

$$v(w, \lambda)\kappa^v = v(w, \lambda) - v(w + \Delta w, \lambda + \Delta \lambda)$$

From banker's perspective, the optimization problem is

$$\frac{dw_t^b}{w_t^b} = \dots + x_{t-}^v \left(\frac{dv_t}{v_{t-}} - \frac{v_{t-} - (1 - \hat{\kappa}_t)}{v_{t-}} \xi_t dN_t - \kappa_{t-}^v (1 - \xi_t) dN_t + \frac{v_{t-} - (1 - \hat{\kappa}_t)}{v_{t-}} dN_t^\tau \right)$$

with $\lambda_t^\tau = 1/\tau - \pi\lambda_t$, $\xi_t \in \{0, 1\}$, $P(\xi_t = 1) = \pi$, and $\{\xi_t\}$ is an i.i.d. process that is independent from everything else. The jump κ_{t-}^v is the amount of decline of bond price upon the distress shock if the bond does not mature during the financial distress shock.

Rewriting the above and omitting the time subscripts, we have

$$\frac{dw^b}{w^b} = \left(r^f + x^K (\mu^R + \frac{A^H}{p} - r^f) + x^d (r^f - r^d) + x^v (\mu^v - r^f) - \rho \right) dt$$

$$+ (x^K (\sigma^K + \sigma^p) + x^v \sigma^v) dB_t - (x^K \kappa^p + \alpha x^d + x^v \xi \frac{v - (1 - \kappa^p - \hat{\kappa}_0)}{v} + x^v (1 - \xi) \kappa^v) dN_t - x^v \frac{v - 1}{v} dN_t^\tau$$

where I have omitted the subscripts t and $t-$ for simplicity. To solve the price of the safe bond \bar{v} , we can simply replace the notation v with \bar{v} , and set the term κ^p and $\hat{\kappa}^0$ both to zero.

The first order condition over x^v is

$$\mu^v - r^f - \lambda\pi \frac{\frac{v - (1 - \kappa^p - \hat{\kappa}^0)}{v}}{1 - (x^K \kappa^p + \alpha x^d + x^v \frac{v - (1 - \kappa^p - \hat{\kappa}^0)}{v})} - \lambda(1 - \pi) \frac{\kappa^v}{1 - (x^K \kappa^p + \alpha x^d + x^v \kappa^v)} - \lambda^\tau \frac{\frac{v-1}{v}}{1 + x^v \frac{v-1}{v}}$$

$$- \underbrace{(\sigma^v)^2 x^v}_{\text{compensation for change in risk - bearing capacity}} - \underbrace{x^K \sigma^v (\sigma^K + \sigma^p)}_{\text{compensation for covariance}} = 0$$

Given that in equilibrium $x^v = 0$, we have

$$\mu^v - r^f = \lambda\pi \frac{1}{1 - \kappa^b} \frac{v - (1 - \kappa^p - \hat{\kappa}^0)}{v} + \lambda(1 - \pi) \frac{1}{1 - \kappa^b} \kappa^v + \lambda^\tau \frac{v - 1}{v} + x^K \sigma^v (\sigma^K + \sigma^p)$$

with

$$\lambda^\tau = \frac{1}{\tau} - \pi\lambda$$

Therefore, the excess return has three components: (1) the compensation for losses during a distress shock, (2) the compensation for losses (negative losses mean positive benefits) in a maturity event without distress shock, and (3) the compensation for exposure to the

volatility risk dB_t , where the price of risk is $x^K(\sigma^K + \sigma^p)$. This equation together with the matched coefficients form an HJB equation for the value of bonds,

$$\begin{aligned} \frac{\partial v}{\partial w} w \mu^w + \frac{1}{2} \frac{\partial^2 v}{\partial w^2} w^2 (\sigma^w)^2 + \frac{\partial v}{\partial \lambda} \mu^\lambda - r^f v = x^K (\sigma^K + \sigma^p) \frac{\partial v}{\partial w} w \sigma^w \\ + \lambda \pi \frac{1}{1 - \kappa^b} (v - (1 - \kappa^p - \hat{\kappa}^0)) + \lambda (1 - \pi) \frac{1}{1 - \kappa^b} \kappa^v v + \lambda^\tau (v - 1) \end{aligned} \quad (60)$$

Solution Methods

We will use the ‘‘false time derivative’’ method, by introducing a time dependence of v . Define such a function as $\tilde{v}(w, \lambda, t)$. Following a similar derivation as (60), we can get the HJB equation for \tilde{v} as

$$\begin{aligned} \frac{\partial \tilde{v}}{\partial t} = \lambda \pi \frac{1}{1 - \kappa^b} (v - (1 - \kappa^p - \hat{\kappa}^0)) + \lambda (1 - \pi) \frac{1}{1 - \kappa^b} \kappa^v v + \lambda^\tau (v - 1) \\ + x^K (\sigma^K + \sigma^p) \frac{\partial \tilde{v}}{\partial w} w \sigma^w + r^f \tilde{v} - \left(\frac{\partial \tilde{v}}{\partial w} w \mu^w + \frac{1}{2} \frac{\partial^2 \tilde{v}}{\partial w^2} w^2 (\sigma^w)^2 + \frac{\partial \tilde{v}}{\partial \lambda} \mu^\lambda \right) \end{aligned}$$

We can start with a function \tilde{v} that satisfies $\tilde{v}(0, \lambda, T) = v(0, \lambda)$, and $\tilde{v}(1, \lambda, T) = v(1, \lambda)$, and has linear interpolation in other regions. By taking T large enough, we are going to have convergence before t reaches 0, i.e., two iterations have close to zero differences. Denote the converged solution as $\tilde{v}(w, \lambda, 0)$. From the property of convergence, we must have $\partial \tilde{v}(w, \lambda, t) / \partial t|_{t=0} = 0$. As a result, $\tilde{v}(w, \lambda, 0)$ satisfies the original PDE of $v(w, \lambda)$, which implies that $v(w, \lambda) = \tilde{v}(w, \lambda, 0)$.

Next, we show how to solve the boundary conditions at $w = 0$ and $w = 1$.

Boundary Conditions

We note that $w = 0$ and $w = 1$ are two absorbing boundaries. At both $w = 0$ and $w = 1$, we have $p = \underline{p}$ or \bar{p} forever, and $\mu^w = \sigma^w = \kappa^p = 0$. Thus, we can simplify the HJB equation (60) into

$$\begin{aligned} \frac{\partial v(w, \lambda)}{\partial \lambda} \mu^\lambda(\lambda) - r^f(w, \lambda) v(w, \lambda) = \lambda \pi \frac{1}{1 - \kappa^b(w, \lambda)} (v(w, \lambda) - (1 - \hat{\kappa}^0)) \\ + \lambda (1 - \pi) \frac{1}{1 - \kappa^b(w, \lambda)} \kappa^v(w, \lambda) v(w, \lambda) + \lambda^\tau(\lambda) (v(w, \lambda) - 1), \quad w \in \{0, 1\} \end{aligned} \quad (61)$$

Suppose that $\kappa^v = 0$ when $\lambda = \lambda^*$ (defined as $\mu^\lambda(\lambda^*) = 0$). Then we get

$$v^{(0)}(w, \lambda^*) = \frac{\lambda^* \pi \frac{1}{1 - \kappa^b(w, \lambda^*)} (1 - \hat{\kappa}^0) + \lambda^{\tau}(\lambda^*)}{\lambda^* \pi \frac{1}{1 - \kappa^b(w, \lambda^*)} + r^f(w, \lambda^*) + \lambda^{\tau}(\lambda^*)}, \quad w \in \{0, 1\}$$

Denote the value function at iteration k as $v^{(k)}(w, \lambda)$. Then for $w = 1$ or $w = 0$, the

algorithm works as follows:

- Step k: Solve for the jump $\kappa^v v = v(w, \lambda) - v(w + \delta w, \lambda + \delta \lambda)$ using $v = v^{(k)}$. Denote this value as $\Delta v^{(k)}$. With such jump solved, we translate the jump equation (61) into an ODE of $v(w, \lambda), w \in \{0, 1\}$ as a function of λ . The ODE solution starts with the initial value $v(w, \lambda^*) = v^{(k)}(w, \lambda^*), w \in \{0, 1\}$. Solve this ODE and denote the solution as $v^{(k+1)}$.
- Stop if

$$\int_{\lambda_L}^{\lambda_H} |v^{(k+1)}(w, \lambda) - v^{(k)}(w, \lambda)| d\lambda < \varepsilon, \quad w \in \{0, 1\}$$

for a small $\varepsilon > 0$.

Finally, we notice that once the $\lambda = \lambda^*$, it will not go up or down unless there is a dN_t shock. Once we know the jump component, we can solve $v(w, \lambda^*)$ along the w dimension as an ODE. The ODE is

$$\frac{\partial^2 v}{\partial w^2} = \frac{\left(\lambda^* \pi \frac{1}{1-\kappa^b} (v - (1 - \kappa^p - \hat{\kappa}^0)) + \lambda(1 - \pi) \frac{1}{1-\kappa^b} \kappa^v v \right.}{\frac{1}{2} w^2 (\sigma^w)^2} \left. + \lambda^\tau (v - 1) + x^K (\sigma^K + \sigma^p) \frac{\partial v}{\partial w} w \sigma^w + r^f v - \frac{\partial v}{\partial w} w \mu^w \right)$$

for $w \neq 0, 1$.

Credit Spread Calibration

Table 5 summarizes the credit spread calibration.

Table 5: Calibrated Parameters for the Credit Spread Construction

Parameters	Choice	Moment
τ	Risky bond maturity	7 Years
π	Maturing probability in illiquidity	0.31
$mE_{\text{crises}}[\kappa_t^p] -$ $mE_{\text{non-crisis}}[\kappa_t^p]$	Additional loss in crises	0.1
$mE_{\text{crises}}[\kappa_t^p] + \hat{\kappa}_0$	Baseline default loss	0.55

- In our baseline calibration, we target the an average maturity of $\tau = 7$ years, which is the average maturity of bonds used in [Krishnamurthy and Muir \(2024\)](#).
- According to [Chen, Collin-Dufresne and Goldstein \(2008\)](#), the 10-year BAA (AAA) default rate is 4.89% (0.63%). The difference in their default rates is 4.26%. We use 4% as

our target. In the model, the default rate is

$$\pi \bar{\lambda} = 0.04$$

where $\bar{\lambda}$ is the average frequency of financial illiquidity, which is 12.8% according to our calibration. Therefore, we have $\pi = 0.31$.

- The total loss given default is $m \cdot \kappa_t^p + \hat{\kappa}^0$ if a illiquidity shock dN_t hits, where κ_t^p is the percentage decline of capital price p_t during a crisis shock. The price jump component κ_t^p is large during crises but close to zero otherwise. We calibrate the loss given default to that of BAA bonds, which from Moodys data has been 55% on average over the last three decades and rose by 10% during the 2008 crisis. As a result, we set m so that $m \cdot \kappa_t^p$ during crises is 10% larger than other defaults. Then we set the average of losses during default to 55% to get $\hat{\kappa}_0$.

Finally, we should note that we define our spread measures in units of standard-deviation differences relative to the unconditional mean value of the credit spread. This is what [Krishnamurthy and Muir \(2024\)](#) do in their empirical work. As a result of this normalization, the results are relatively insensitive to the exact values of the credit-spread calibration.

A.8 Additional Moments

Table 6: Additional Moments

Volatilities	Static	Bayesian	Diagnostic
vol(liquidity premium)	0.6%	0.6%	0.9%
vol(risk-free rate)	0.3%	0.7%	0.7%
vol(annual equity return)	18.2%	15.7%	15.9%
vol(investment yearly growth rate)	4.2%	3.4%	4.4%
vol(consumption growth return)	2.2%	1.6%	1.8%
mean(risk free rate)	3.1%	7.4%	3.7%
average total Sharpe ratio of capital return	0.49	0.48	0.41

Table 6 reports other asset pricing moments from our model. We report the volatility of asset prices and returns. Based on U.S. data from 1934 to 2020, the volatility of the liquidity premium in the 0.37 and the volatility of the risk-free rate in the U.S. is 3.5%. The volatility of banking-sector equity yearly return in the U.S. from 1973 to 2020 is 26%, somewhat higher than the model counterparts. The volatility of consumption growth is in line with historical data, while that of investment growth is smaller relative to the historical data.

One notable difference from the table is that the risk free rate is much higher in the Bayesian model. Digging into this further, we find that this result is because the calibrated

productivity parameters A^H and A^L are higher in the Bayesian model than the other two models so that the growth out output is higher leading to the higher risk free rate.

The average Sharpe ratio on the capital return in all three models are in line with historical Sharpe ratios for the U.S. equity market.

A.9 Model without Financial Frictions

In this section, we discuss model results when we shut off financial frictions. We show that this variation of model fails to match most of the empirical facts, so that the financial friction is needed to explain the crisis cycle facts.

To meaningfully discuss a model without financial frictions but with belief variation, we need to consider an economy where the liquidity shock reduces aggregate resources (and is not just a transfer shock as in our main model), such as the models of consumption disasters (e.g., Wachter (2012); Gabaix (2012)). Consider this case. Since we have a production economy, we directly specify the losses in production in a disaster. Suppose that the dynamic evolution of productive capital owned by agent $j \in \{\text{banker, household}\}$ is

$$\frac{dk_{j,t}}{k_{j,t}} = \mu_t^K dt - \delta dt + \sigma^K dB_t - \kappa_t^K dN_t \quad (62)$$

and the belief regarding the disaster shock dN_t follows the same process as our main model. The key departure from our main model is that the dN_t shock is now a direct shock to the efficiency units of capital, instead of a financial shock that indirectly affects output. Agents in the economy believe that the actual shock frequency is $\lambda_t^{\text{belief}}$, which is a function of the underlying rational belief $\lambda_t^{\text{Bayesian}}$. We assume $\kappa_t^K > 0$ and follows an independent distribution from other variables. As we there are no financial frictions, banks can issue equity to households and perfectly share risks. Since bankers earn a higher return on capital, households optimally only hold bank equity and earn dividends equal to \bar{A} on capital instead of \underline{A} . Given identical log preferences and no financial constraints, we can aggregate across agents and study an economy populated by only bankers.

Define the total wealth as W_t . The consumption-good market clearing implies

$$\rho W_t = \bar{A}K_t - i_t K_t.$$

Wealth is the total value of capital:

$$p_t K_t = W_t.$$

The above two equations together imply

$$\rho p_t = \bar{A} - i_t$$

Since $i_t = \phi(\mu_t^K) = \phi(\delta + (p_t - 1)/\chi)$, we obtain

$$\rho p_t = \bar{A} - \phi(\delta + (p_t - 1)/\chi).$$

which implies that the value of capital is constant, regardless of the fluctuations in the frequency of disaster shocks dN_t . The main reason is that although the risk premium increases when the frequency of dN_t rises, the risk-free rate endogenously falls, and thus the net effect on asset price cancels out (the exact offset is due to log utility).

We note that since banks and households receive the same return on assets, the ratio between banker wealth to household wealth is constant and equal to w_0 . Thus, the economy only has two state variables, λ_t^{belief} and K_t , while asset prices and allocations are only driven by λ_t^{belief} . Furthermore, bank lending to firms (holding of capital) is always $w_0 K_t$, and since the output is always $\bar{A} K_t$, bank credit/GDP ratio is constant, w_0/\bar{A} . As a result, this model cannot match any of the empirical facts about the association between bank credit and macroeconomic outcomes.

Next, we define crises and analyze cross-section severity of crises. To allow for richness of the model, we assume that the fall in capital quantity, κ_t^K , upon a Poisson shock is drawn from the same distribution $F^K(\cdot)$ independently each time. Since total output is $\bar{A} K_t$, the output drop is equal to the change in capital κ_t^K . We define a crisis in the same way as the main paper, i.e., a dN_t shock that causes the fall in output to be below the 4% quantile. An economy has a crisis when $\kappa^K > \bar{\kappa}$ where $\bar{\kappa}$ solves $\bar{\lambda}(1 - F^K(\bar{\kappa})) = 0.04$ and $\bar{\lambda}$ is the average intensity of Poisson shocks. It is easy to prove that the crisis indicator $1\{\kappa_t^K > \bar{\kappa}\}dN_t$ is independent of all information before the crisis, including quantities and prices. Consequently, in this model, crises are not predictable with the credit variables. Furthermore, the cross-sectional severity of crises, $\kappa_t^K 1\{\kappa_t^K > \bar{\kappa}\}$ can also not be predicted using credit information from before the crisis.

Nevertheless, the credit spread in this economy will vary with λ_t^{belief} , since losses on debt claims on firms will be related to the dN_t shock. Denote the credit spread as $s(w_t, \lambda_t^{belief})$. We can easily prove that the model does not feature ‘‘froth’’ (i.e., low credit spread) before crises, since conditional on crises, the expected credit spread is the same due to independence between (κ_t^K, dN_t) and the belief λ_t^{belief} . Formally,

$$E[s(w_0, \lambda_{t-}^{belief}) | dN_t = 1, \kappa_t^K > \bar{\kappa}] = E[s(w_0, \lambda_{t-}^{belief})].$$

Intuitively, crisis severity or likelihood does not have an endogenous relationship with belief

or “froth” before crises, and thus this model is unable to generate the empirical pattern.

We can also prove that a low credit spread predict less rather more likely crises. Formally, conditional on credit spread below the mean value \bar{s} , the expectation of a crisis event is:

$$\begin{aligned} & E[\mathbf{1}\{dN_t = 1, \kappa_t^K > \bar{\kappa}\} | s(w_0, \lambda_{t-}^{belief}) < \bar{s}] \\ & = (1 - F(\bar{\kappa})) E[\lambda_{t-}^{belief} | s(w_0, \lambda_{t-}^{belief}) < \bar{s}] \end{aligned}$$

Given that credit spread $s(w_0, \lambda_{t-}^{belief})$ increases in λ_{t-}^{belief} , we find that

$$E[\lambda_{t-}^{belief} | s(w_0, \lambda_{t-}^{belief}) > \bar{s}] > E[\lambda_{t-}^{belief} | s(w_0, \lambda_{t-}^{belief}) < \bar{s}]$$

Thus, low credit spreads indicates that crises are less likely, instead of more likely.

Furthermore, we can also show that the spike of the credit spread in a crisis is unrelated to crisis severity. This is because

$$\begin{aligned} & E[s(w_0, \lambda_t^{belief}) - s(w_0, \lambda_{t-}^{belief}) | dN_t = 1, \kappa_t^K > \bar{\kappa}] \\ & = E[s(w_0, \lambda_{t-}^{belief} + \kappa_{t-}^{belief}) | dN_t = 1, \kappa_t^K > \bar{\kappa}] - E[s(w_0, \lambda_{t-}^{belief})] \\ & = E[s(w_0, \lambda_{t-}^{belief} + \kappa_{t-}^{belief})] - E[s(w_0, \lambda_{t-}^{belief})] \end{aligned}$$

where the second equality is due to the independence between κ_t^K and the frequency of crises.

In summary, a model without financial frictions but with a production (and thus consumption) disaster can generate a time-varying credit spread. However, it cannot generate the following main facts about crises:

- Conditional on crises, high bank credit predicts more severe crises (since bank credit/GDP is constant in the model).
- Bank credit/GDP predicts excess returns.
- Conditional on crises, credit spread before a crisis below its unconditional average.
- Low credit spread predicts more frequent crises; in this model the opposite will occur.

B General Beliefs

The solution method of our model offers an avenue to evaluate alternative belief processes beyond the Bayesian and diagnostic models of the main text. That is, as long as agents think they are Bayesian, the model’s policy rules are the same as that of the Bayesian model.

As with our analysis of the diagnostic model, we simply need to simulate beliefs under an alternative belief process to evaluate the impact of alternative beliefs.

The more general way of formulating beliefs is to write

$$d\lambda_t^{belief} = A_t dt + B_t dN_t$$

where A_t and B_t are adapted stochastic processes.

In the case of Bayesian beliefs, A_t and B_t can be written as linear functions of λ_t , capturing all of the history dependence. In the case of diagnostic beliefs both A and B are functions of λ_t and λ_{t-T} where T is the reference point for diagnostic expectations and λ_t and λ_{t-T} are the Bayesian beliefs.

We will proceed in the following steps:

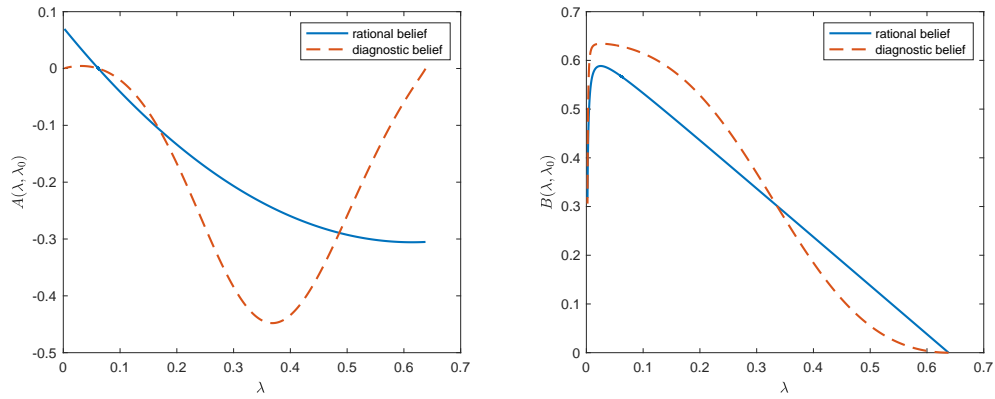
- First, we show that the historical reference point λ_{t-T} is quantitatively not important for the diagnostic belief model. The model's performance can be summarized via mapping from λ_t to λ_t^{belief} .
- Second, we illustrate how belief distortion in different states λ_t affect key model moments.
- Finally, we separate the diagnostic belief into an optimism portion and a pessimism portion, and we separately calibrate these two versions and show quantitative results under these different forms of beliefs.

First, we write the diagnostic belief as

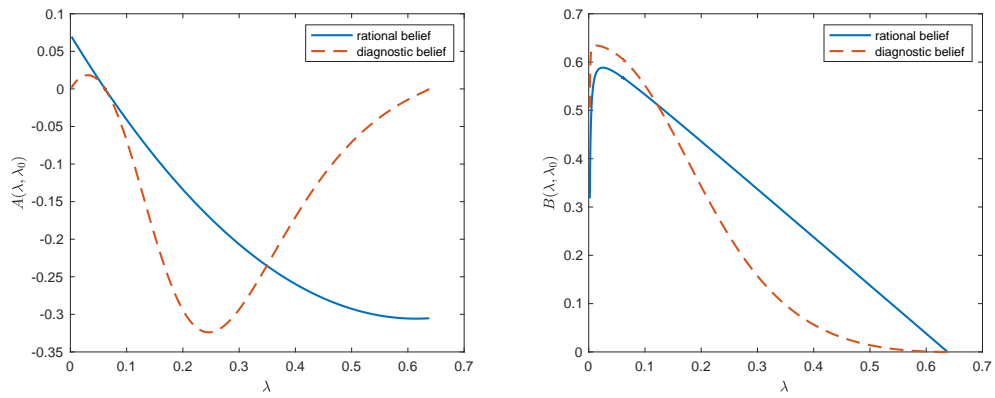
$$d\lambda_t^{belief} = A(\lambda_t, \lambda_{t-T})dt + B(\lambda_t, \lambda_{t-T})dN_t$$

where λ_t is the rational belief process. Then we plot the functions $A(\lambda, \lambda_0)$ and $B(\lambda, \lambda_0)$ – where since t does not matter, our notation is that λ_0 is the reference point – as functions of λ in both the Bayesian case and the diagnostic belief case. In Figure 17, we show three values of λ_0 : high λ_0 (90% quantile of the stationary distribution), average λ_0 (mean of the stationary distribution), and low λ_0 (10% quantile of the stationary distribution). We find that the basic properties of the drift and jump functions are quantitatively similar across all scenarios. In Figure 18, we plot $\lambda_t^{diagnostic}(\lambda_t, \lambda_{t-T})$ as a function of $\lambda_t = \lambda_t^{Bayesian}$ for three values of $\lambda_0 = \lambda_{t-T}$. Again, we find that the functions are very similar across different values of λ_0 .

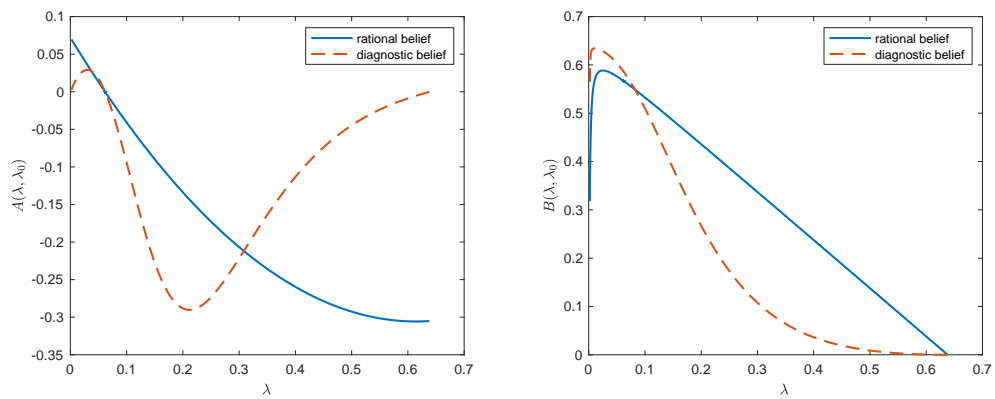
The above results indicate that we can ignore the dependence on λ_{t-T} without much loss. Furthermore, we note that directly specifying the distorted behavioral belief $\lambda^{belief}(\lambda^{Bayesian})$ as a function of the Bayesian belief is a convenient way to understand the impact of the distorted belief.



(a) High λ_0 (90% quantile of stationary distribution)

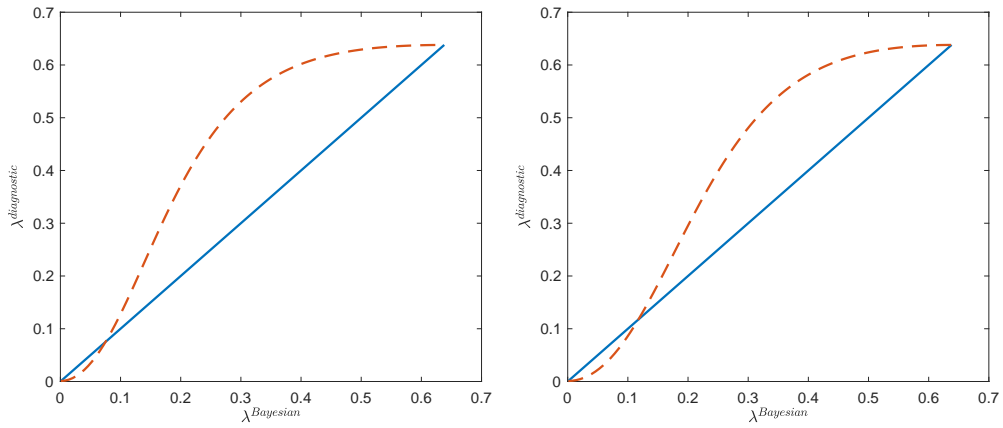


(b) Average λ_0 (average of the stationary distribution)



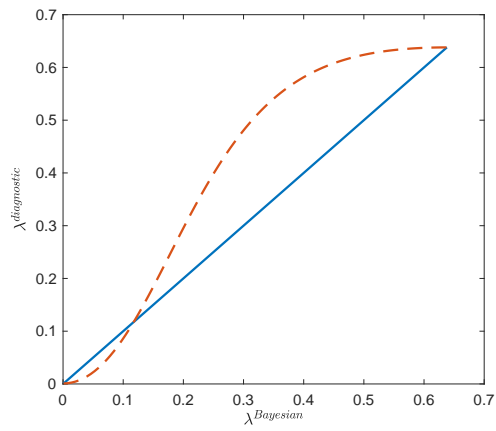
(c) Low λ_0 (10% quantile of stationary distribution)

Figure 17: Drift and Jump of Bayesian and Diagnostic Belief Process, with Reference $\lambda_0 = \lambda_{t-T}$.



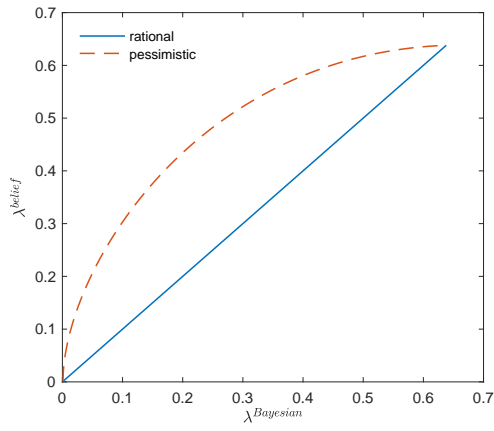
(a) $\lambda_0 = 0.04$

(b) $\lambda_0 = 0.1$

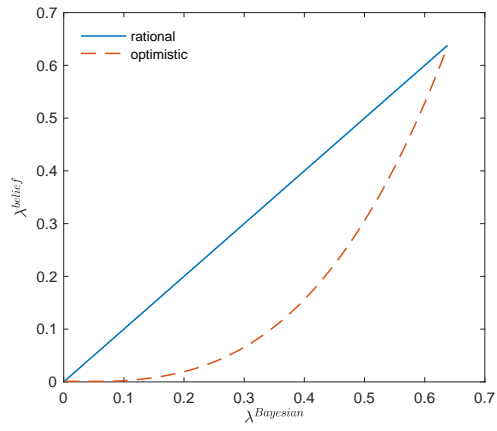


(c) $\lambda_0 = 0.44$

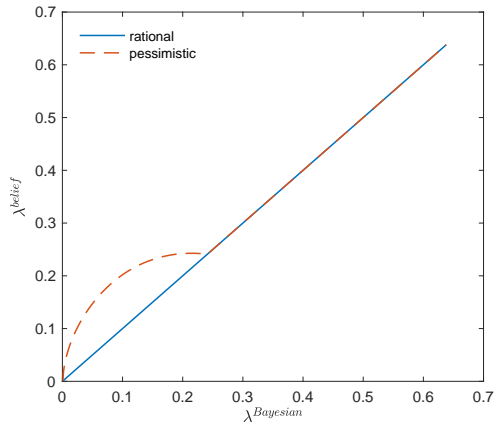
Figure 18: Diagnostic Belief $\lambda_t^{Diagnostic}$ as a Function of $\lambda_t^{Bayesian}$, with Reference Belief $\lambda_0 = \lambda_{t-T}$



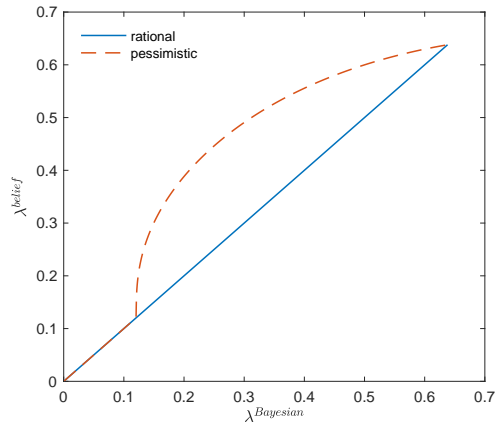
(a) Case 1



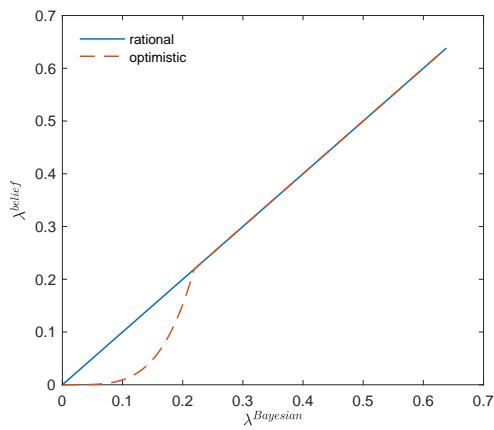
(b) Case 2



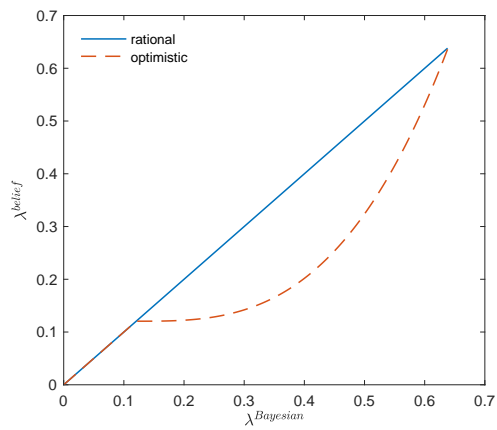
(c) Case 3



(d) Case 4



(e) Case 5



(f) Case 6

Figure 19: General Behavioral Belief

Next, we consider different cases of optimism and pessimism in different regions of the state space, $\lambda_t = \lambda^{Bayesian}$.

1. Pessimism for all $\lambda_t \in (\lambda_L, \lambda_H)$.
2. Optimism for all $\lambda_t \in (\lambda_L, \lambda_H)$.
3. Pessimism for low λ_t , but Bayesian otherwise.
4. Pessimism for high λ_t , but Bayesian otherwise.
5. Optimism for low λ_t , but Bayesian otherwise.
6. Optimism for high λ_t , but Bayesian otherwise.

We plot the functions in each case in Figure 19. Then we simulate the model under each case, keeping all other parameters of the model as the same. Results are reported in Table 7. Note that these results are comparative statics and the model is not recalibrated in each case. Our key findings are as follows:

- Average liquidity premium is sensitive to whether on average the behavioral belief is optimistic or pessimistic.
- Average output/capital ratio is insensitive to beliefs.
- Average output drop in a crisis is directly related to the difference of λ^{belief} between when the underlying λ is high and low, not the absolute amount of optimism or pessimism.
- Bank leverage is highly related to whether on average the behavioral belief is optimistic or pessimistic.
- Low pre-crisis credit spread is directly affected by the difference of λ^{belief} between high λ and low λ , not the absolute amount of optimism or pessimism. For example, case 1, 2, 4, and 5 all feature large differences in λ^{belief} for high and low λ , and thus a very low credit spread before crises. Case 3 and 6 have small differences in λ^{belief} for high and low λ , and thus not a very low credit spread before crises.

In summary, there are two aspects that are important for the quantitative results:

1. On average, whether the belief is optimistic or pessimistic.
2. The difference between λ^{belief} when the underlying λ is high v.s. when the underlying λ is low.

Table 7: Model Moments and Different Cases of Alternative Beliefs

Moments	Bayesian	Alternative Beliefs					
		case 1	case 2	case 3	case 4	case 5	case 6
Avg liquidity premium (%)	0.77	1.76	0.39	1.28	0.98	0.56	0.72
Avg credit spread change in crises	56%	0.39	0.45	0.38	0.50	0.53	0.37
Half-life of credit spread (yrs)	2.5	3.25	1.17	1.42	3.25	2.33	0.92
Output/capital ratio	17%	0.15	0.15	0.15	0.15	0.15	0.15
Avg 3-year output drop in crises	-0.08	-0.08	-0.09	-0.05	-0.12	-0.11	-0.05
Output growth volatility	0.03	0.03	0.03	0.03	0.03	0.04	0.03
Avg bank leverage	5.1	2.5	7.5	2.8	4.4	7.3	4.2
Lower Credit Spread	-0.25	-0.28	-0.29	0.04	-0.31	-0.30	-0.03

Thus far we have discussed these effects as a comparative static exercise. An important lesson from our analysis is that recalibrating the model to match the same targets tends to lessen the quantitative differences between the models. We next discuss this effect focusing on two of the cases we have presented: case 4 and case 5.

We define an optimistic belief as

$$\lambda_t^{optimistic} = \min\{\lambda_t, \lambda^{diagnostic}(\lambda_t, \lambda_0)\}$$

and a pessimistic belief as

$$\lambda_t^{pessimistic} = \max\{\lambda_t, \lambda^{diagnostic}(\lambda_t, \lambda_0)\}$$

where for simplicity we use the average value of λ_t in the calibrated Bayesian model as λ_0 , so the comparison with the Bayesian model does not involve a change of reference point for the behavioral belief. We note that since the recalibration will change the parameters λ_H , $\lambda_{H \rightarrow L}$ etc., the optimistic belief and pessimistic belief mappings will adjust accordingly. Figure 20 plots these cases from the model simulation.

Table 8 presents the model's fit of moment targets and Table 9 reports the resulting estimates for the parameter values. The last two columns consider the optimism and pessimism cases. From the top panel of the table, we note that the models match the targets as well as the other belief variants.

We can see the impact of optimism when considering the true frequency of the liquidity shocks. The optimistic agent's beliefs are that λ is lower than the true λ in the states where λ is low. That by itself will lead to too low an average liquidity premium. Thus the model recalibrates and sets λ_H higher (row 3 of the lower panel) in order to match the average liquidity premium. As a result, the true frequency of liquidity shock rises substantially. By itself, this would mean that crises would lead to a large fall in GDP, larger than the

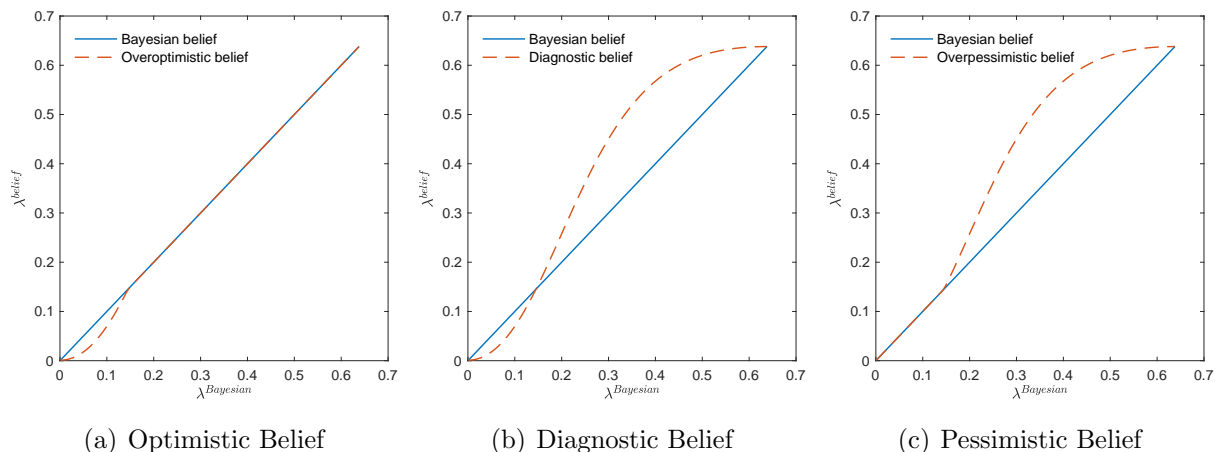


Figure 20: Functional forms of Different Beliefs Derived from Diagnostic Beliefs

Table 8: Calibration Moments for Alternative Beliefs

	Data	Bayesian	Optimistic	Pessimistic
1. Average liquidity premium	0.90%	0.77%	0.88%	0.92%
2. Avg credit spread change in crises	70%	56%	61%	56%
3. Half-life of credit spread recovery (years)	2.5	2.5	2.0	2.6
4. Output/capital ratio	14%	17%	14%	14%
5. Avg 3-year output drop in crises	-9.1%	-7.9%	-10.1%	-9.1%
6. Output growth volatility	3.8%	2.9%	3.4%	2.6%
7. Average bank leverage	5.0	5.1	4.8	4.9

Table 9: Calibrated Parameter Values for Alternative Beliefs

Explanation	Parameter	Bayesian	Optimistic	Pessimistic
1. Household productivity	A_L	0.17	0.14	0.13
2. Bank lending advantage	$A_H - A_L$	0.03	0.02	0.02
3. High intensity of liquidity shock	λ_H	0.56	0.84	0.77
4. Low to high transition	$\lambda_{L \rightarrow H}$	0.11	0.11	0.11
5. High to low transition	$\lambda_{H \rightarrow L}$	0.47	0.40	0.64
6. Volatility of capital growth	σ^K	0.03	0.03	0.03
7. Banker-household transition rate	η	0.05	0.03	0.04

calibration target of row 5 in the top row. The calibration then adjusts the parameter $A_H - A_L$ which is closely related to the output decline in the crisis, reducing this gap and thereby matching the output decline in the crisis. A similar logic applies to the pessimism case, although here the model hits the liquidity premium target by increasing the transition rate $\lambda_{H \rightarrow L}$ so that the model spends less time in the high λ pessimistic states.

Finally, we replicate two of the model simulations linking bank credit and subsequent outcomes. In both cases we see that the belief variants of optimism and pessimism produce results that are similar to the Bayesian and diagnostic models in terms of matching the data.

Table 10: GDP Growth and Credit Spread (General Beliefs)

	<i>Dependent variable: GDP Growth from t to $t + 3$</i>					
	Bayesian		Overoptimistic		Overpessimistic	
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \text{credit spread}_t * \text{crisis}_t$	-2.87		-3.61		-3.08	
$(\frac{\text{bank credit}}{\text{GDP}})_t * \text{crisis}_t$		-2.18		-3.79		-3.53
R^2	0.12	0.12	0.13	0.13	0.16	0.16

Note: Model and data regressions are normalized so that the coefficients reflect the impact of one sigma change in spreads, and bank credit/GDP.

Table 11: Bank Credit Predicting Capital Excess Returns (General Beliefs)

<i>Dependent variable:</i>			
Average realized excess return t_{+1}			
	(1) Rational	(2) Optimistic	(3) Pessimistic
$(\frac{\text{bank credit}}{\text{GDP}})_t$	-0.01	-0.01	-0.01
Observations			
R ²	0.01	0.01	0.01

Note: Model excess return is defined as the return to capital minus the risk-free rate. Data excess return is the excess equity index return from Online Appendix Table 3 of Baron and Xiong (2017). To ensure comparability, the model return to capital has been normalized to equal the standard deviation of returns reported by Baron and Xiong (2017).