

Booms and Systemic Banking Crises*

Frédéric Boissay[†] Fabrice Collard[‡] Frank Smets[§]
European Central Bank University of Bern European Central Bank

This draft: January 22, 2013
First draft: December 2011

Abstract

The empirical literature on systemic banking crises (SBCs) has shown that SBCs are rare events that break out in the midst of credit intensive booms and bring about particularly deep and long-lasting recessions. We attempt to explain these phenomena within a dynamic general equilibrium model featuring a non-trivial banking sector. In the model, banks are heterogeneous with respect to their intermediation skills, which gives rise to an interbank market. Moral hazard and asymmetric information on this market may generate sudden interbank market freezes, SBCs, credit crunches and, ultimately, severe recessions. Simulations of a calibrated version of the model indicate that typical SBCs break out in the midst of a credit boom generated by a sequence of positive supply shocks rather than being the outcome of a big negative wealth shock. We also show that the model can account for the relative severity of recessions with SBCs and their longer duration.

Keywords: Moral Hazard, Asymmetric Information, Lending Boom, Credit Crunch, Systemic Banking Crisis

JEL Class.: E32, E44, G01, G21.

Disclaimer: The views expressed in this paper are our own and should not be interpreted as reflecting the views of the European Central Bank or the Eurosystem.

*We benefited from discussions with F. Alvarez, F. Canova, S. Claessens, R. Cooper, V. Curdia, M. Darracq-Paries, F. Dufourt, H. Degryse, X. Freixas, S. Gilchrist, J. Gomes, P-O. Gourinchas, J. Henry, I. Jaccard, P. Jacquinot, P. Karadi, N. Kiyotaki, L. Laeven, G. Lombardo, K. Nikolov, G. Nuño, G. Saint-Paul, H-S. Shin, J. Stein, G. Ströbl, J. Suarez, M. Trabandt, P. Weil, as well as seminar participants at the ECB, the European University Institute, the University of Munich, CRETE 2012, CEPR conference in St Gallen, Banque de France Conference in Strasbourg, National Bank of Belgium, REDg 2012 Workshop in Madrid, Conference in honor of C. Sims in Princeton, Bocconi University, 3rd Joint French macro workshop in Paris, National Bank of Serbia, University of Bergen, Paris School of Economics. We are particularly indebted to H. Dellas and B. Diba.

[†]European Central Bank, Postfach 16 03 19, 60066 Frankfurt am Main, Germany. email: fred-eric.boissay@ecb.int, URL:<http://www.fboissay.eu>

[‡]Universität Bern, Department of Economics, Schanzenekstrasse 1, Postfach 8573, CH-3001 Bern. email: fabrice.collard@gmail.com, URL:<http://fabcol.free.fr>

[§]European Central Bank, Postfach 16 03 19, 60066 Frankfurt am Main, Germany. email: frank.smets@ecb.int

Non-Technical Summary

Recent empirical research on systemic banking crises (henceforth, SBCs) has highlighted the existence of similar patterns across diverse episodes. SBCs are rare events. Recessions that follow SBC episodes are deeper and longer lasting than other recessions. And, more importantly for the purpose of this paper, SBCs follow credit intensive booms; “banking crises are credit booms gone wrong” (see *e.g.*, Schularick and Taylor, 2012, p. 1032; the notion that banking crises are endogenous and follow prosperous times is also present in Minsky, 1977). Rare, large, adverse financial shocks could possibly account for the first two properties. But they do not seem in line with the fact that the occurrence of an SBC is not random but rather closely linked to credit conditions. So, while most of the existing macro-economic literature on financial crises has focused on understanding and modeling the propagation and the amplification of adverse random shocks, the presence of the third stylized fact mentioned above calls for an alternative approach, which would explain how financial imbalances build up over time and suddenly unravel.

In our model, financial crises take the form of non-linearities in the full equilibrium dynamics of the economy. They result from the pro-cyclicality of bank balance sheets that emanates from interbank market funding. Banks are assumed to be heterogeneous with respect to their intermediation skills, which gives rise to an interbank market. There are two-way relationships between the situation of the interbank market and the real economy. On the one hand, moral hazard and asymmetric information on the interbank market may generate sudden interbank market freezes, SBCs, credit crunches and, ultimately, severe recessions. On the other hand, those frictions become more prevalent as the household accumulates assets in anticipation of severe recessions. The typical run of events leading to a financial crisis is as follows. A sequence of favorable, non permanent, supply shocks hits the economy. The resulting increase in the productivity of capital leads to a demand-driven expansion of credit that pushes the corporate loan rate above steady state. As productivity goes back to trend, firms reduce their demand for credit, whereas households continue to accumulate assets, thus feeding the supply of credit by banks. The credit boom then turns supply-driven and the corporate loan rate goes down, falling below steady state. By giving banks incentives to take more risks or misbehave, too low a corporate loan rate contributes to eroding trust within the banking sector precisely at a time when banks increase in size. When counterparty fears in the interbank market rise too high, the market freezes. Ultimately, the credit boom lowers the resilience of the banking sector to shocks, making systemic crises more likely.

We calibrate the model on the business cycles in the US (post WWII) and the financial cycles in fourteen OECD countries (1870–2008), and assess its quantitative properties. The model reproduces the stylized facts associated with SBCs remarkably well. Most of the time the model behaves like a standard financial accelerator model, but once in a while —on average

every forty years— there is a banking crisis. The larger the credit boom, (i) the higher the probability of an SBC, (ii) the sooner the SBC, and (iii) —once the SBC breaks out— the deeper and the longer the recession. In our simulations, the recessions associated with SBCs are significantly deeper (with a 45% larger output loss) than average recessions. Overall, our results validate the role of supply-driven credit booms leading to credit busts. This result is of particular importance from a policy making perspective as it implies that systemic banking crises are predictable. We indeed use the model to compute the k -step ahead probability of an SBC at any point in time. Fed with actual US data over the period 1960–2011, the model yields remarkably realistic results. For example, the one-year ahead probability of a crisis is essentially zero in the 1960–70s. It jumps up twice during the sample period: in 1982–3, just before the Savings & Loans crisis, and in 2007–9. Although very stylized, our model is thus also able to detect financial imbalances and predict future crises. Finally, we analyse the sensitivity of the above results to the parameters of the model. Among other things, we find that risk averse economies tend to be more prone to crises because more risk averse households typically accumulate more assets during booms, which amplifies credit booms.

1 Introduction

Recent empirical research on systemic banking crises (henceforth, SBCs) has highlighted the existence of similar patterns across diverse episodes (see Reinhart and Rogoff, 2009; Jordà et al., 2011a,b; Claessens et al., 2011; Schularick and Taylor, 2012). SBCs are rare events. Recessions that follow SBC episodes are deeper and longer lasting than other recessions (see Section 2). And, more importantly for the purpose of this paper, SBCs follow credit intensive booms; “banking crises are credit booms gone wrong” (see Schularick and Taylor, 2012, Borio and Drehmann, 2009, and Borio and Lowe, 2002; the notion that banking crises are endogenous and follow prosperous times is also present in Minsky, 1977). Most of the existing macro-economic literature on financial crises has focused on understanding and modeling the propagation and the amplification of random adverse shocks. Indeed, rare, large enough, adverse financial shocks can account for the first two properties (see *e.g.* Gertler and Kiyotaki, 2009). However, by implying that financial crises may break out at any time in the business cycle, they do not seem in line with the fact that the occurrence of an SBC is closely linked to credit conditions (Gorton, 2010, 2012). The third stylized fact therefore calls for an alternative approach.

In this paper, financial crises result from the pro-cyclicality of bank balance sheets that emanates from interbank market funding. During expansions, bank market funding and credit supply increase, pushing down the rates of return on corporate and interbank loans. The lower rates accentuate agency problems in the interbank market that lead to a reduction on market funding and contractions. The larger the credit boom relative to the possibilities for productive use of loans, the larger the fall in interest rates, and the higher the probability of a bank run in —and therefore of a disastrous freeze of— the interbank market.¹ As in Shin (2008) and Hahm et al. (2011), the behavior of banks (credit in our case) during good times sows the seeds of a financial crisis. In our model, banks are heterogeneous in terms of —non-publicly observed— intermediation efficiency. They finance their activities with funds obtained from depositors/shareholders or raised in the interbank market. There exists the usual agency problem in this market as borrowing banks can always divert some of the funds into low return assets that cannot be recovered by the lending banks. The incentives for diversion are stronger for less productive banks and depend on the level of interest rates in the economy. The lower the return on loans, the greater the incentive to engage in fund diversion and hence the greater counterparty risk in the interbank market. The typical run of events leading to a financial crisis is as follows. A sequence of favorable,

¹Our representation of financial crises as market-based bank runs is in line with what happened during the 2007-8 financial crisis (see Uhlig, 2010). Shin (2010, Chap. 8), for example, depicts the demise of Northern Rock —a UK bank— in 2007 as primarily originating from the sudden freezing of the short-term funding market, what he refers to as a “modern bank run”. A traditional, deposit-based, run on the bank took place as well, but it did so one month later, accounted for only 10% of the bank’s fall in total funding, and rapidly stopped because, following the news of the run, the UK authorities pledged 100% deposit guarantees.

non permanent, supply shocks hits the economy. The resulting increase in the productivity of capital leads to a demand-driven expansion of credit that pushes interest rates up. The more efficient banks expand their loan operations by drawing funds from the less efficient banks and market funding in the banking sector as a whole increases. The economy booms. But as the supply shocks run their course, the probability of imminent reversion to average productivity increases. This slows down corporate demand for loans while at the same time inducing households to accumulate savings in order to smooth consumption. Credit expansion becomes supply-driven, putting downward pressure on interest rates. The rate of return on interbank loans declines making the less efficient banks more prone to borrow themselves and divert those funds. As the identity of these banks is not known, counterparty risk in the interbank market goes up, interbank loans decline, and market finance recedes. The stronger the credit expansion during the booming times, the larger the decline in interest rates and the more acute the agency problem in the interbank market. We show that there is a threshold value of interest rates below which the interbank market freezes, corporate credit collapses and the economy tanks. This threshold can be alternatively expressed in terms of the level of banking assets relative to the level of productivity (output) in the economy, that we call the absorption capacity of the banking sector. Supply-driven, excessive credit creation places the economy beyond its absorption capacity, triggering an SBC.

Our work differs from related work on financial crises in several important aspects. In contrast to Shin (2008) and Hahm et al. (2011), whose models are static, ours is a full blown dynamic stochastic general equilibrium (DSGE) model, and thus more suitable for quantitative analysis. Unlike Bernanke et al. (1999), Jermann and Quadrini (2010), Gertler and Karadi (2011), who study the linearized system dynamics around the steady state in models where adverse shocks are amplified by financial market frictions our model analysis characterizes the full equilibrium dynamics inclusive of important and critical non-linearities such as the freezing of interbank markets. This is an important difference because near the steady state our model features a traditional financial accelerator. But away from it (and the large departures from the steady state are the endogenous outcome of a boom-bust endogenous cycle, rather than a big shock) it gives rise to banking crises. Crises are rare but generate particularly large output losses and inefficiencies due to the presence of pecuniary externalities. The models of Bianchi (2009), Bianchi and Mendoza (2010), and Korinek (2010) also exhibit non-linearities and pecuniary externalities but assume that the interest rate is exogenous, so they are at best applicable to small open economies and emerging markets.² Perhaps the models the closest to ours are Brunnermeier and Sannikov (2012) and He and Krishnamurthy (2012). These two models too feature a powerful non-linear amplification mechanism. As in other

²In these models non-linearities are due to occasionally binding constraints, whereas in our case they are due to the economy switching from normal to crisis times. Gertler and Kiyotaki (2012) also develop a model with bank runs and regime switches. While in their model bank runs are deposit-based and unexpected, in ours they are market-based and, more importantly, agents are fully rational: they perfectly know and take into account the probability that runs will occur in the future.

DSGE models with financial frictions, financial crises are the outcomes of adverse exogenous financial shocks (*e.g.* to banks' net worth), whose size ultimately determines the size of the crises. In our case, in contrast, shocks play only a secondary role because crises are related to whether or not financial imbalances (*e.g.* credit boom, ballooning bank balance sheets) have built up in the first place. That is, crises may break out endogenously, even in the absence of negative shocks. Another important feature of our model is that it does not rely on financial shocks to generate banking crises; the technology shock is indeed the only exogenous source of uncertainty.³

We calibrate the model on the business cycles in the US (post-WWII) and the financial cycles in fourteen OECD countries (1870-2008), and assess its quantitative properties. The model reproduces the stylized facts associated with SBCs remarkably well. Most of the time bank assets remain below the threshold for financial crises and the model behaves like a standard financial accelerator model. But once in a while —on average every forty years— there is a banking crisis. The typical banking crisis in our simulations is preceded with a credit boom and brings about both a credit crunch and a recession. On the brink of an SBC, risk averse households accumulate precautionary savings and inadvertently fuel a credit boom, which brings credit creation even closer to the economy's absorption capacity. Our findings are in line with empirical evidence (see Schularick and Taylor, 2011, among others) and validate the role of supply-driven credit booms leading to credit busts. The larger the credit boom, (*i*) the higher the probability of an SBC, (*ii*) the sooner the SBC, and (*iii*) —once the SBC breaks out— the deeper and the longer the recession. In our simulations, a recession associated with SBCs is significantly deeper (with a 45% larger output loss) than the average recession. We use the model to compute the k -step ahead probability of an SBC at any point in time. Fed with actual US data of total factor productivity over the period 1960-2011, the model produces remarkably realistic results. For example, the one-year ahead probability of a crisis is essentially zero in the 60–70s. It jumps up twice during the sample period: in 1982–3, just before the Savings & Loans crisis, and in 2007–9.

The paper proceeds as follows. Section 2 briefly documents key empirical facts about the dynamics of systemic banking crises in 14 OECD countries for the period 1870–2008. Section 3 describes our theoretical framework, the micro-foundations of interbank market freezes and the dynamic implications of such events. Section 4 discusses our calibration strategy and presents our solution method. Section 5 analyses the quantitative implications of the model as well as its performance against the facts documented in Section 2. A last section concludes.

³Another difference with many existing models concerns the modeling of the financial friction. In Bianchi's model, for example, the friction affects the firms and operates through excess credit demand ("over-borrowing"), whereas in our model it operates through excess credit supply. Our model also differs from Brunnermeier and Sannikov's in that it is a discrete time model that we calibrate on US and OECD data, which allows us to confront our model with the data.

2 Key Facts on Systemic Banking Crises

Reinhart and Rogoff (2009), Claessens et al. (2011), Jordà et al. (2011a,b), and Schularick and Taylor (2012) recently documented that SBCs share, despite their variety, a few common regularities. Building upon this earlier work, we briefly describe in this section the key facts on SBCs, against which we will later assess the quantitative properties of our model. To do so we use the historical dataset assembled by Jordà et al. (2011a). This dataset comprises yearly observations for real GDP per capita, total domestic currency loans of banks and banking institutions to non-financial companies and households, banks' total assets, the dates of business cycle peaks, and the dates of banking crises from 1870 to 2008 for 14 OECD developed countries.⁴ A banking crisis is defined as an event during which the financial sector experiences bank runs, sharp increases in default rates accompanied by large losses of capital that result in public intervention, bankruptcy, or the forced merger of major financial institutions (see Laeven and Valencia, 2008). For the purpose of the present paper, we further define as *systemic* those banking crises that are concomitant with a recession, *i.e.* that break out between the peak and the trough of a given business cycle. Jordà et al. use the Bry-Boschan algorithm to date peaks and troughs consistently across countries. We exclude war times and only keep complete business cycles, from peak to peak. After trimming, our sample covers 176 full-length business cycles. The main statistics are reported in Table 1.

Table 1: Statistics on recessions and banking crises

	N. obs.	N	Frequency (%)	Magnitude (%)	Duration (Years)
				from peak to trough	
All banking crises	1,736	78	4.49	–	–
Systemic Banking Crises (SBC)	1,736	42	2.42	–	–
All recessions	1,736	176	10.20	4.86 (5.91)	1.85
Recessions with SBC (A)	1,736	42	23.86	6.74 (6.61)	2.59
Recessions w/o SBC (B)	1,736	134	76.13	4.27 (5.61)	1.61
Test A≠B, p-value (%)	–	–	–	2.61	0.00

Note: the magnitudes reported into parentheses are calculated using the HP-filtered series of output, and are thereby corrected for the underlying trend in output. Following Ravn and Uhlig (2002), we set the parameter of the Hodrick–Prescott filter to 6.25.


Fact #1: Systemic Banking Crises are Rare Events. 78 banking crises can be identified in the sample, which comprises 1,736 observations. The frequency of crises is therefore 4.49%, which means that countries in our sample experience a crisis, on average, every 22 years. Half of those 78 banking crises were systemic. Hence, SBCs are rare events, which occur on average every forty years. In contrast, recessions are much more frequent and occur

⁴The list of country comprises Australia, Canada, Denmark, France, Germany, Great Britain, Italy, Japan, the Netherlands, Norway, Spain, Sweden, Switzerland and the United States of America.

every ten years or so.

Fact #2: Financial recessions are deeper and last longer than other recessions.

While only one fourth of the recessions we identify involve a banking crisis, these “financial recessions” are on average significantly deeper than other, regular recessions. For instance, we find that the drop in real GDP per capita from peak to trough is 40% bigger during financial recessions (6.74%) than during the average recession (4.86%) (60% deeper than recessions without SBCs), or about 12% when the data are HP-filtered (see Table 1). On average, systemic banking crises also last one year longer. The dynamics of financial recessions too is different: they tend to be preceded by a faster increase in GDP and credit compared with other recessions, as Figure 1 shows. Claessens et al. (2011) report similar patterns based on a shorter data set that includes emerging countries.⁵



Note: The reported % deviations are the average % deviations around the Hodrick–Prescott trend (calculated with a parameter of 6.25). Notice that the implied magnitude of financial recessions in the left chart is about 4.3%, which is lower than that 6.61% reported earlier in Table 1. This discrepancy reflects the fact that the statistics in Table 1 also take into account recessions of more than 6 years. We find similar results when we consider the % deviations of output and credit from their respective linear trends (see the companion technical appendix).

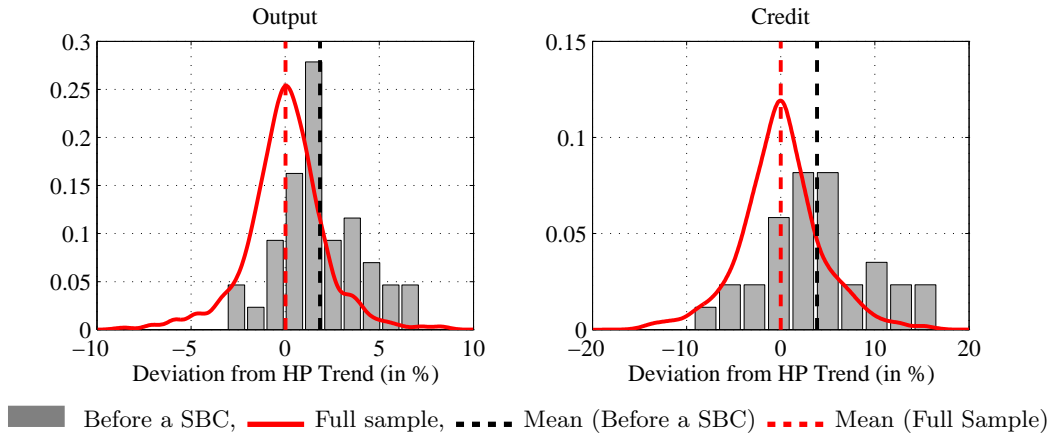
Fact #3: Systemic banking crises break out in the midst of credit intensive booms.

Systemic banking crises do not hit at random (Gorton, 1988). To illustrate this point, Figure 2 reports the empirical distributions of GDP (left panel) and credit (right panel) gaps, as measured by the percentage deviations from a Hodrick–Prescott trend, in the year that precedes a typical systemic banking crisis (histogram). The red line corresponds to the distribution in the full sample, which we use as benchmark. The figure shows that, before a systemic banking crisis both GDP and credit are above trend, with average deviations of

⁵The cross-correlations between credit and output over the sample also show significant differences between normal and crisis times. For example, during regular recessions the maximal correlation between (HP-filtered) credit and output is reached contemporaneously, with $\text{corr}(\text{credit}_t, \text{gdp}_t) = 0.38$. We find similar results for periods outside recessions. In contrast, during financial recessions the maximal correlation is reached with one lag on credit ($\text{corr}(\text{credit}_{t-1}, \text{gdp}_t) = 0.24$), suggesting that in those periods credit *leads* output.

1.8% and 3.8%, respectively. This suggests that crises break out at a particular point in the business cycle, typically in good times, in the midst of a credit boom. A general pattern extensively documented by Reinhart and Rogoff (2009, p. 157).

Figure 2: Distributions of GDP and credit gaps



3 The Model

We consider a closed economy populated with one representative risk averse household, one representative risk neutral competitive firm, and a mass one of heterogeneous, risk neutral, and competitive banks.

3.1 The Representative Firm

The representative firm lives for one period. It produces a homogeneous good that can be either consumed or invested by means of capital, k_t , and labor, h_t , according to a constant returns to scale technology represented by the production function⁶ $F(k_t, h_t; z_t)$ where z_t is the level of total factor productivity (TFP), which is assumed to follow an AR(1) process of the form

$$\log z_t = \rho_z \log z_{t-1} + \varepsilon_t$$

where $|\rho_z| < 1$ and ε_t is an exogenous normally distributed TFP shock with zero mean and standard deviation σ_z that is realized at the beginning of period t . Variations in productivity are the only source of uncertainty and ε_t is realized at the beginning of period t , before the firm decides on its production plan. Capital, k_t , depreciates at rate $\delta \in (0, 1)$. The firm is born with no resources and must borrow k_t from the banks at a gross corporate loan rate R_t at the beginning of the period to be able to achieve production. The corporate loan is repaid at the end of the period. The firm also rents labor services from the household at rate w_t .

⁶The production function is increasing in both inputs, concave and satisfies Inada conditions.

The production plan is decided so as to maximize profits, which are given by

$$\pi_t = F(k_t, h_t; z_t) + (1 - \delta)k_t - R_t k_t - w_t h_t. \quad (1)$$

3.2 The Representative Household

The infinitely lived representative household supplies inelastically one unit of labor per period⁷ and has preferences over the flow of consumption, c_t , which are represented by the utility function:

$$\max_{\{a_{t+\tau+1}, c_{t+\tau}\}_{\tau=0}^{\infty}} \mathbb{E}_t \sum_{\tau=0}^{\infty} \beta^\tau u(c_{t+\tau}), \quad (2)$$

where $u(c_t)$ satisfies the usual regularity conditions,⁸ $\beta \in (0, 1)$ is the psychological discount factor, and $\mathbb{E}_t(\cdot)$ denotes the expectation operator which is taken over $\{\varepsilon_{t+\tau+1}\}_{\tau=0}^{+\infty}$. The household enters period t with assets, a_t , which she deposits in the banking sector and from which she receives a state contingent gross return r_t . There is no friction between the household and the banking sector and —since Modigliani and Miller’s theorem applies— we cannot say anything as to whether a_t is made of bank deposits or bank equity. (This assumption will be relaxed in Section 6.2.⁹) The household earns unit wage w_t from supplying her labor and receives profits π_t from the firm. This income is then used to purchase the consumption good and transfer assets to the next period. Accordingly, the budget constraint is given by

$$c_t + a_{t+1} = r_t a_t + w_t + \pi_t. \quad (3)$$

The saving decision is determined by the standard arbitrage condition

$$u'(c_t) = \beta \mathbb{E}_t(u'(c_{t+1}) r_{t+1}). \quad (4)$$

Notice that, as will become clear shortly, there exists a positive wedge between banks’ gross return on corporate loans (R_t) and the gross return on bank equity/assets (r_t). This wedge is due to inefficiencies in the banking sector.

3.3 The Banking Sector

The banking sector is at the core of the model and plays a non-trivial role because of two specific features. First, banks are heterogeneous with respect to their intermediation technology — some banks are more efficient than others, which potentially gives rise to an interbank

⁷This latter assumption is made for exposition purposes only and will be relaxed in the quantitative analysis (see Sections 4 and 5).

⁸In particular, we have $u'(c) > 0$, $u''(c) < 0$, $u(0) = \infty$ and $u'(\infty) = 0$

⁹To emphasize the fact that the household owns the banks and r_t is state contingent, we will in the meantime interchangeably refer to a_t as bank equity or deposits, and to r_t as the gross return on bank equity/deposits. Note also that we implicitly assume the existence of frictions between the household and the firm that prevent the household from financing the firm directly —this assumption is standard in the macro-literature. In the companion technical appendix of this paper, we also consider the case where the firm finances a fraction of its investment directly through the market.

market. It follows that banks have two types of activities. On the one hand they run traditional banking operations, which consist in collecting deposits/equity from households and lending the funds to the firm. In Shin (2008) and Shin and Shin (2011)’s language, these are “core” activities and, accordingly, bank deposits/equity are banks’ core liabilities. On the other hand, banks also issue interbank claims (“non-core” assets/liabilities) so as to re-allocate assets toward the most efficient banks.¹⁰ Second, the banking sector is subject to both asymmetric information and moral hazard problems, which impair the functioning of the interbank market.

3.3.1 Banks

There is a continuum of one-period,¹¹ risk-neutral, competitive banks that raise deposits/equity a_t from the household at the end of period $t - 1$. At the time they raise deposits/equity, banks are identical and, in particular, have all the same size as they enter period t . At the beginning of period t , each bank draws a random bank-specific intermediation skill. Banks therefore become heterogeneous. Let p denote the bank with ability p , and assume that the p s are distributed over the closed interval $[0, 1]$ with cumulative distribution $\mu(p)$, satisfying $\mu(0) = 0$, $\mu(1) = 1$, $\mu'(p) > 0$. Bank p must pay an intermediation, dead-weight, cost $(1 - p)R_t$ per unit of loan at the end of the period, so that its net return on each loan is pR_t . This cost reflects the bank’s operational costs, for example, the cost of collecting corporate loans or monitoring the firm.¹² As an outside option, banks also have the possibility to invest assets in their own project. This project does not involve any intermediation cost but yields a lower, constant, and exogenous payoff γ per unit of good invested. Such an investment is inefficient, *i.e.* $\gamma < R_t$.¹³ While there are several ways to interpret this outside option, we

¹⁰“The relevant distinction between core- and non-core liabilities can be seen as having to do with whether the claim is held by the ultimate domestic creditors (the domestic household sector). Repos and other claims held by banks on other banks can be regarded as non-core liabilities which are more volatile”, Shin and Shin (2011, p. 3).

¹¹Banks that operate in period t are born at the end of period $t - 1$ and die at the end of period t . We will assume in a moment that banks are heterogeneous and that their types are private information. The assumption of one-period living banks is made to preserve this asymmetry of information over time. An alternative and equivalent approach would be to allow banks to live infinitely and, in order to rule out potential reputation effects, to assume that the types are randomly drawn afresh every period.

¹²This assumption is not crucial but convenient, because in this case bank heterogeneity is immaterial to the representative firm, which always pays its debt irrespective of the bank it borrows from. One could consider several alternative setups without loss of generality. For example, one could assume that there is a continuum of firms and that banks have different monitoring skills, which determine the probability that the projects of the firms they respectively lend to succeed. Typically, the firms borrowing from the skillful banks would then be able to repay their loan in full, while those borrowing from inefficient banks would default. We do not use this setup because we want to confine the inefficiencies within the banking sector and, by doing so, stay the closest possible to the textbook neoclassical model, where firms do not default.

¹³Indeed, if γ were strictly above R_t then banks would not finance the firm and, because of an Inada condition on the production function ($\lim_{k \rightarrow 0} \partial F_k(k, h; z)/\partial k = +\infty$), the marginal productivity of capital would be infinite, hence a contradiction. And the case $\gamma = R_t$ is ruled out by the existence of financial intermediation costs (see below). Notice that, since in the absence of the storage technology unused goods would depreciate at rate δ , the net return of storage is $\gamma - (1 - \delta)$, which we assume is positive.

will refer to it as a storage technology.¹⁴ An important aspect of this assumption is that the funds invested in this outside option cannot be used to finance the firm. This is key for the model to generate credit crunches.

Bank heterogeneity gives rise to an intra-periodic interbank market, where the least efficient banks lend to the most efficient ones at gross rate ρ_t , with $\gamma \leq \rho_t \leq R_t$.¹⁵ Unlike corporate loans, interbank loans do not bear operational costs. Banks take the interbank rate ρ_t and the corporate loan rate R_t as given. Given these rates, bank p decides whether, and how much, it borrows or lends. Hereafter, we will refer to the banks that supply funds on the interbank market as “lenders” and to those that borrow as “borrowers”. Let ϕ_t be the — endogenous and publicly observable — amount borrowed per unit of deposit/equity by a borrower p , with $\phi_t \geq 0$. In the rest of the paper, we will refer to ϕ_t as the “market/interbank funding ratio”, defined as the ratio of market funding (non-core liabilities) to traditional funding (core liabilities). Then bank p ’s gross return on equity/assets is

$$r_t(p) \equiv \max \{ pR_t(1 + \phi_t) - \rho_t\phi_t, \rho_t \}. \quad (5)$$

It is equal to $pR_t(1 + \phi_t) - \rho_t\phi_t$ when bank p borrows $\phi_t a_t$ from other banks at cost ρ_t and lends $(1 + \phi_t)a_t$ to the firm for return pR_t . And it is equal to ρ_t when, instead, bank p does not do financial intermediation and lends to other banks. Bank p chooses to be a borrower when

$$pR_t(1 + \phi_t) - \rho_t\phi_t \geq \rho_t \iff p \geq \bar{p}_t \equiv \frac{\rho_t}{R_t}. \quad (\text{PC})$$

Inequality (PC) is the participation constraint of bank p to the interbank market as borrower, rather than as lender, and pins down the type of the *marginal bank* \bar{p}_t that is indifferent between the two options. Banks with $p < \bar{p}_t$ delegate financial intermediation to more efficient banks with $p \geq \bar{p}_t$. In a frictionless world, all banks with $p < 1$ would lend to the most efficient bank, so that $\bar{p}_t = 1$. This bank would have an infinite market funding ratio ($\phi_t \rightarrow +\infty$) and corner all assets; the economy would then reach the First Best allocation. The presence of two frictions on the interbank market — moral hazard and asymmetric information — prevents the economy from achieving First Best.

Moral Hazard: We assume that the proceeds of the storage technology are not traceable and cannot be seized by creditors. This implies that interbank loan contracts are not enforce-

¹⁴One could assume that the return of this activity varies over time. This would not affect our results as long as the return is not strictly positively correlated with the business cycle and the outside option can still be used as an insurance against adverse aggregate shocks. To fix ideas, one can think of this outside option as an intra-period home production activity or as a safe haven. One could also assume that the household has access to this storage technology. This would not affect our results either. Indeed, since γ is the return that even the worst bank ($p = 0$) can make in any state of the nature, the ex-post return on bank deposits/equity is always above that of storage (*i.e.* $\gamma < r_t$). Hence it would never be optimal for the household to use this technology.

¹⁵The interbank rate is the same for all borrowers, otherwise those that promise the lowest returns would not attract any lender. It has to be the case that $\rho_t \leq R_t$, otherwise no bank would be willing to borrow on the interbank market. Likewise, we have $\rho_t \geq \gamma$, otherwise no banks would be willing to lend.

able and that banks can walk away with the funds raised on the interbank market, without paying the interbank loans. Following the current practice (e.g. Hart, 1995, Burkart and Ellingsen, 2004), we refer to such opportunistic behavior as “cash diversion”. When a bank diverts cash, the proceeds ultimately accrue to the shareholder — *i.e.* the household. The so-diverted cash is stored until the end of the period, and yields the return γ .¹⁶ A bank that diverts $(1 + \phi_t) a_t$ faces a diversion cost proportional to the size of the loan, and can only run away with $(1 + \theta\phi_t) a_t$, for a net payoff of $\gamma (1 + \theta\phi_t) a_t$. Parameter $\theta \in [0, 1]$ reflects the cost of diversion, which is zero when $\theta = 1$ and maximal when $\theta = 0$. From a corporate finance literature viewpoint (e.g. Tirole, 2006), this is a standard moral hazard problem: (i) the gain from diversion increases with ϕ_t , (ii) the opportunity cost of diversion increases with bank efficiency p and (iii) with the corporate loan rate R_t . Features (i) and (ii) imply that efficient banks with “skin in the game” are less inclined to run away than highly leveraged and inefficient banks. Feature (iii) is similar to feature (ii), but in the “time-series” (as opposed to “cross-sectional”) dimension; it implies that banks are more inclined to run away when the return on corporate loans is low. This latter feature captures recent empirical evidence that banks tend to take more credit risk in such a situation (Maddaloni and Peydro, 2011).

Asymmetric Information: Lenders do not observe borrowers’ skills — *i.e.* p is privately known— and therefore do not know borrowers’ private incentives to divert cash. In this context, the loan contracts signed on the interbank market are the same for all banks. Neither ϕ_t nor ρ_t depends on p .¹⁷

By limiting the borrowing capacity of the most efficient bank ($p = 1$), moral hazard will give less efficient banks room to borrow; hence the positive wedge between R_t and r_t . Moral hazard is not enough to generate market freezes, though. For this we also need uncertainty about the quality –and therefore some adverse selection– of borrowers. Hence, both moral hazard and information asymmetry will be necessary to generate SBCs in the model.

Lenders want to deter borrowers from diverting. They can do so by limiting the quantity of funds that borrowers can borrow, so that even the most inefficient banks (*i.e.* those that

¹⁶Two comments are in order here. First, we will soon see that an incentive compatibility constraint will make sure that no bank diverts cash in equilibrium. Hence, cash diversion will be an out-of-equilibrium threat. Second, to be consistent, the return on cash diversion must not be strictly higher than γ . Otherwise, the diversion technology would dominate storage and would then be the relevant outside option for the banks.

¹⁷To see this, consider a menu of debt contracts $\{\rho_t(\tilde{p}), \phi_t(\tilde{p})\}_{\tilde{p} \in [0,1]}$ intended for the borrowers of types \tilde{p} s, and notice that lenders’ arbitrage across these contracts requires that $\rho_t(\tilde{p}) = \rho_t \forall \tilde{p} \in [0, 1]$. It is easy to see that such a menu of contracts cannot be revealing because any borrower p (*i.e.* with $pR_t > \rho_t$) claiming being of type \tilde{p} would make profit $r_t(\tilde{p} | p) = pR_t + (pR_t - \rho_t) \phi_t(\tilde{p})$ and pick the contract with the highest $\phi_t(\tilde{p})$, independent of its type. It is equally easy to see that there is no revealing menu of equity contracts either. Indeed, consider a menu of equity contracts $\{\eta_t(\tilde{p}), \phi_t(\tilde{p})\}_{\tilde{p} \in [0,1]}$, where $\eta_t(\tilde{p})$ would be the share of retained earnings. Then the net profit of bank p would be $\eta_t(\tilde{p}) (1 + \phi_t(\tilde{p})) pR_t$ and, in equilibrium, this bank would pick the contract that yields the highest $\eta_t(\tilde{p}) (1 + \phi_t(\tilde{p}))$, independently of its own p .

should be lending) have no interest in demanding a loan and diverting it:

$$\gamma(1 + \theta\phi_t) \leq \rho_t. \quad (\text{IC})$$

This incentive compatibility constraint sets a limit to ϕ_t , which can therefore also be interpreted as lenders' funding tolerance, *i.e.* the limit market funding ratio above which a bank refuses to lend or, in Holmström and Tirole's language, the borrower's pledgeable income.¹⁸ The program of bank $p \geq \bar{p}_t$ thus consists in maximizing its return on equity $r_t(p)$ (see (5)) with respect to ϕ_t subject to constraint (IC). Proposition 1 below follows from the fact that, by construction —see (PC), the net return on interbank borrowing is strictly positive for borrowers.¹⁹

Proposition 1 (Optimal Interbank Funding Ratio) *The IC constraint binds at the optimum of the borrowing bank p , which thus exhausts its borrowing capacity: $\phi_t = \frac{\rho_t - \gamma}{\gamma\theta}$.*

The positive relationship between ϕ_t and ρ_t is a critical feature of the interbank funding ratio. When ρ_t increases, the net present value of corporate loans diminishes and only the most efficient banks remain on the demand side of the market. Since these banks have little private incentive to divert, lenders tolerate a higher interbank funding ratio (ϕ_t goes up). This is due to the negative (positive) externality that the marginal bank exerts on the other banks when she enters (leaves) the demand side of the market as, by having higher incentives to run away, she then raises (reduces) lenders' counterparty fears. In the limit case where $\rho_t = \gamma$, there is no demand for interbank loan because borrowers cannot commit to repay. The interbank funding ratio ϕ_t and the type of the marginal bank \bar{p}_t fully describe banks' optimal decisions.

3.3.2 Interbank Market

The equilibrium of the interbank market is characterized by the gross return ρ_t that clears the market. We look for an equilibrium where $\rho_t > \gamma$ so that $\phi_t > 0$ and trade takes place. Since a mass $\mu(\bar{p}_t)$ of banks lend a_t , the aggregate supply of funds is equal to $\mu(\bar{p}_t)a_t$. Since a mass $1 - \mu(\bar{p}_t)$ of banks borrow $\phi_t a_t$, aggregate demand is equal to $(1 - \mu(\bar{p}_t))\phi_t a_t$. The market clears when (using relations (PC) and Proposition 1):

$$\underbrace{\mu\left(\frac{\rho_t}{R_t}\right)}_{\text{supply}} = \underbrace{\left(1 - \mu\left(\frac{\rho_t}{R_t}\right)\right)}_{\text{demand}} \underbrace{\frac{\rho_t - \gamma}{\gamma\theta}}_{\substack{\text{extensive margin} \\ \text{intensive margin}}} \iff R_t = \Psi(\rho_t) \equiv \frac{\rho_t}{\mu^{-1}\left(\frac{\rho_t - \gamma}{\rho_t - \gamma(1 - \theta)}\right)}. \quad (6)$$

¹⁸One could indeed recast the moral hazard problem into a setup *à la* Holmström and Tirole (1997), whereby borrowers may misuse the funds and enjoy private benefits at the expense of their creditors. *Stricto sensu*, the pledgeable income is the highest income that can be pledged without jeopardizing the borrower's incentives, *i.e.* $\rho_t(\rho_t - \gamma)a_t/\gamma\theta$.

¹⁹The proofs of all propositions are reported in Appendix A.

Following the fall in the corporate loan rate, the supply curve shifts to the right while the demand curve shifts to the left. Given the initial equilibrium rate $\rho_t = \rho_E$, demand falls below supply. Market clearing then requires that ρ_t go down, which results in more banks demanding funds (extensive margin). But since the banks that switch from the supply to the demand side are less efficient and have a relatively higher private incentive to divert cash, lenders require borrowers to deleverage. By construction, this intensive margin effect is the strongest when $R_t < \bar{R}$. It follows that, ultimately, aggregate demand decreases and excess supply goes further up. The de-leveraging process feeds itself and goes on until the market freezes, in point A , where $\rho_t = \gamma$.

In point E , where $R_t \geq \bar{R}$, borrowers have enough incentives to finance the firm and an interbank market equilibrium with trade exists; such a situation will be referred to as *normal times*. This equilibrium is *stable* in the sense that, in this point, net aggregate demand is a decreasing function of ρ_t and, following any small perturbation to ρ_t away from ρ_E , a standard Walrasian tâtonnement process brings ρ_t back to ρ_E . For usual cumulative distributions, $\mu(p)$,²¹ another interbank market equilibrium with trade is also possible, in point U . However, we rule it out because it is unstable.

When $R_t < \bar{R}$, things are different: autarky prevails (point A in Figure 3).²² Demand and supply are both equal to zero, and the market clears because (i) borrowers have no pledgeable income ($\phi_t = 0$) and (ii) lenders are indifferent between interbank loans and storage. The marginal bank is then bank $\bar{p}_t = \gamma/R_t$, which is indifferent between financing the firm and using the storage technology. A mass $\mu(\gamma/R_t)$ of banks uses the storage technology, instead of lending to the firm. In the rest of the paper, we will interpret such a situation as a *systemic banking crisis*. This equilibrium is stable because net aggregate demand in this point decreases with ρ_t . Due to strategic complementarities between lenders (see Cooper and John, 1988), the autarkic equilibrium always exists, whatever the values of R_t . (Indeed, no bank has interest in making a loan if no one else does it.) Hence, it also always coexists

²¹Notably for the family distribution $\mu(p) = p^\lambda$ (with $\lambda \geq 0$) that we will be using later in the calibration.

²²Notice that competition rules out the equilibria with rationing that could potentially arise when $R_t < \bar{R}$. To see this, consider a loan contract (ρ_t^*, ϕ_t^*) , with $\rho_t^* > \gamma$, where aggregate supply exceeds aggregate demand, so that some lenders are randomly rationed. In such a situation, there is at least one inefficient bank that is rationed and whose only outside option is storage, with marginal return γ . Assume further that this bank offers a new contract (ρ_t, ϕ_t^*) , with $\gamma < \rho_t < \rho_t^*$, which attracts the whole continuum of borrowers. Although this bank is too small to serve all the borrowers in the continuum, it can still lend a little bit, say $\varepsilon > 0$, to every one of them. Borrowers would then borrow ε at rate ρ_t and $\phi_t^* - \varepsilon$ at rate ρ_t^* . Since the deviating bank is atomistic, one would have $\varepsilon \searrow 0$, and the increase in borrowers' individual payoffs generated by the new contract would be negligible, leaving borrowers' incentives unchanged. From the deviating bank's perspective, in contrast, the marginal return would rise from γ to ρ_t , making it worth to offer the new contract. It follows that, whenever there is excess supply of funds on the interbank market any rationed, inefficient bank can be better off by offering an adequate loan contract. In other terms, rationing is not sustainable as an equilibrium. This contrasts with a world with monopolistic banks, where banks may then have interest in rationing their supply of loans, as in Stiglitz and Weiss (1981), for example.

with the equilibrium with trade whenever the latter exists. In order to rule out potential coordination failures we assume that banks always coordinate on the equilibrium with trade, which is Pareto-dominant, in this case.²³ That is, the interbank market freezes only when there exists no equilibrium with trade.

Based on relations (5) and (6), we can complete the description of the banking sector by deriving the sector's return on equity:

$$r_t \equiv \int_0^1 r_t(p) d\mu(p) = \begin{cases} R_t \int_{\bar{p}_t}^1 p \frac{d\mu(p)}{1-\mu(\bar{p}_t)} , & \text{if an equilibrium with trade exists} \\ R_t \left(\frac{\gamma}{R_t} \mu\left(\frac{\gamma}{R_t}\right) + \int_{\frac{\gamma}{R_t}}^1 p d\mu(p) \right) , & \text{otherwise.} \end{cases} \quad (7)$$

The interpretation of r_t is clear. When the equilibrium with trade exists, inefficient banks delegate financial intermediation to a mass $1 - \mu(\bar{p}_t)$ of efficient banks, each of which therefore lending to the firm a multiple $1 + \phi_t = 1 / (1 - \mu(\bar{p}_t))$ of their initial assets against net return pR_t . In autarky, in contrast, a mass $1 - \mu(\gamma/R_t)$ of the banks make corporate loans, while the remainder use the storage technology. The banking sector is fully efficient when $\gamma \rightarrow 0$, *i.e.* when interbank loan contracts are fully enforceable, as in this case $\bar{R} \rightarrow 0$ and $\bar{p}_t \rightarrow 1$ (the interbank market always exists and only the best bank does the intermediation), and $\phi_t \rightarrow +\infty$ (the best bank is infinitely leveraged). The same is true when $\lim_{p \nearrow 1} \mu(p) = 0$, since in this case there is a mass one of banks with $p = 1$ and banks are homogeneous and all efficient.

3.3.3 Aggregate Supply of Corporate Loans

In normal times banks reallocate their assets through the interbank market, and all assets a_t are ultimately channeled to the firm. In crisis times, in contrast, the interbank market freezes and only the banks with $p \geq \gamma/R_t$ lend to the firm. As a consequence, the banking sector only supplies $(1 - \mu(\gamma/R_t)) a_t$ as corporate loans. Denoting by k_t^s banks' aggregate supply of corporate loans, one thus gets:

$$k_t^s = \begin{cases} a_t , & \text{if an equilibrium with trade exists} \\ \left(1 - \mu\left(\frac{\gamma}{R_t}\right)\right) a_t , & \text{otherwise} \end{cases} . \quad (8)$$

3.4 Recursive Decentralized General Equilibrium

A general equilibrium of the economy is defined as follows.

²³For a discussion on the selection of the Pareto-dominant equilibrium in games with multiple Pareto-rankable Nash equilibria, see Cooper et al. (1990).

Definition 1 (Recursive decentralized general equilibrium) *A decentralized recursive general equilibrium is a sequence of prices $\mathcal{P}_t \equiv \{R_{t+i}, r_{t+i}, \rho_{t+i}, w_{t+i}\}_{i=0}^{\infty}$ and a sequence of quantities $\mathcal{Q}_t \equiv \{c_{t+i}, y_{t+i}, k_{t+i}, h_{t+i}, a_{t+i}\}_{i=0}^{\infty}$ such that for a given sequence of prices, \mathcal{P}_t , the sequence of quantities, \mathcal{Q}_t , solves the optimization problems of the agents, and for a sequence of quantities, \mathcal{Q}_t , the sequence of prices, \mathcal{P}_t , clears the markets.*

In equilibrium, the household supplies one unit of labor, implying that the production level is given by $f(k_t; z_t) \equiv F(k_t, 1; z_t)$ and the marginal efficiency of capital is $f_k(k_t; z_t) \equiv \frac{\partial F(k_t, 1; z_t)}{\partial k_t}$. The market clearing condition on the corporate loan market thus takes the form

$$f_k^{-1}(R_t + \delta - 1; z_t) = \begin{cases} a_t, & \text{if an equilibrium with trade exists} & (a) \\ a_t - \mu \left(\frac{\gamma}{R_t} \right) a_t, & \text{otherwise.} & (b) \end{cases} \quad (9)$$

Relation (9) yields the equilibrium R_t as a function of the two state variables of the model, a_t and z_t . It also points to the two-way relationship that exists between the interbank loan market and the retail corporate loan market. We indeed showed that the way the interbank market operates depends on whether or not $R_t \geq \bar{R}$. Likewise, whether or not the interbank market operates has an impact on the supply of corporate loans and, therefore, on R_t . To solve for the general equilibrium we need to take into account these feedback effects. We proceed in two steps. First, we solve (9a) for R_t under the conjecture that the interbank market equilibrium with trade exists, and then check a posteriori whether indeed $R_t \geq \bar{R}$. In the negative, the interbank market equilibrium with trade cannot emerge, and the interbank market freezes. In this case the equilibrium corporate loan rate is the R_t that solves (9b). Proposition 2 follows.

Proposition 2 (Interbank loan market freeze) *The interbank loan market is at work if and only if $a_t \leq \bar{a}_t \equiv f_k^{-1}(\bar{R} + \delta - 1; z_t)$, and freezes otherwise.*

The threshold \bar{a}_t is the maximum quantity of assets that the banking sector can reallocate efficiently. Above this threshold counterparty fears on the interbank market are so widespread that mistrust prevails and the interbank market freezes. In the rest of the paper we will refer to \bar{a}_t as the absorption capacity of the banking sector. Importantly, Proposition 2 suggests that the ability of the banking sector to re-allocate assets internally ultimately depends on the level of productivity in the real sector, z_t . The more productive the real sector, the more efficient the banking sector ($\partial \bar{a}_t / \partial z_t > 0$). The intuition and mechanics are clear. An increase in total factor productivity raises the demand for capital and the equilibrium corporate loan rate. By raising banks' opportunity cost of storage and cash diversion, the increase in R_t also reduces uncertainty about counterparties' quality, making it less likely for the interbank loan market to freeze. Given a level of assets a_t , there therefore exists a productivity threshold \bar{z}_t below which the interbank market freezes (with $\bar{z}_t \equiv f_k^{-1}(\bar{R} + \delta - 1; a_t)$ and $\partial \bar{z}_t / \partial a_t < 0$).

Overall, our model captures the notion that banks' core liabilities (equity/deposits a_t), which are predetermined, are a stable source of funding whereas non-core liabilities are unstable funding because they are subject to market runs. Proposition 3 below shows how disruptions in the wholesale financial market spill over the retail loan market and trigger a credit crunch.

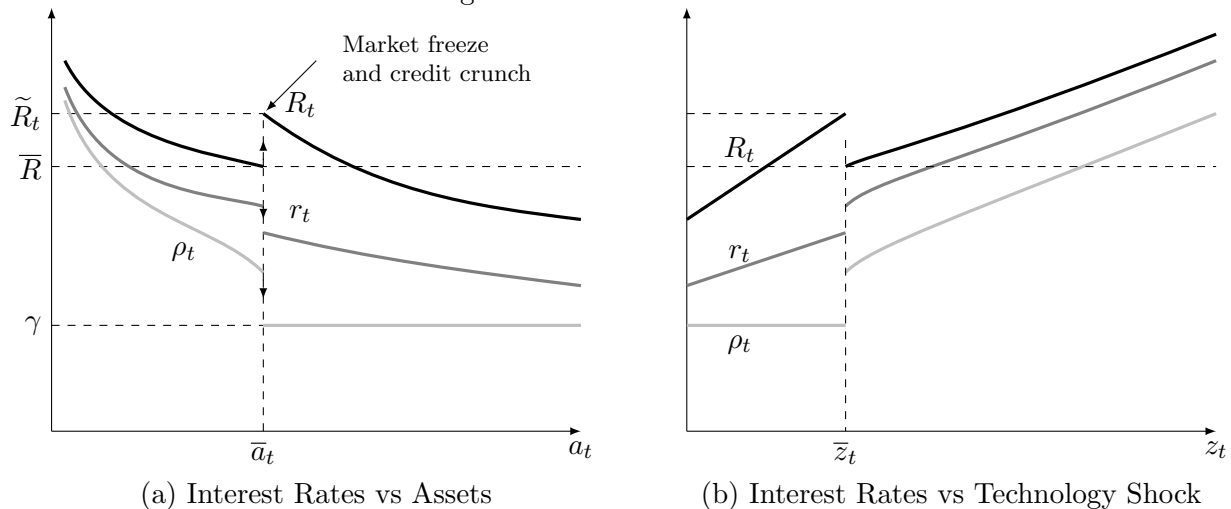
Proposition 3 (Credit crunch) *An interbank market freeze is accompanied with a sudden fall in the supply of corporate loans k_t^s (i.e. given z_t , $\lim_{a_t \searrow \bar{a}_t} k_t^s < \lim_{a_t \nearrow \bar{a}_t} k_t^s$), as well as by a sudden increase in the interest rate spread R_t/r_t (i.e. given z_t , $\lim_{a_t \searrow \bar{a}_t} R_t/r_t > \lim_{a_t \nearrow \bar{a}_t} R_t/r_t$). We will refer to such a situation as a credit crunch.*

Panel (a) of Figure 4 illustrates Proposition 3. It depicts the equilibrium rates, R_t , r_t , and ρ_t as functions of a_t for a given level of productivity z_t . The corporate loan rate monotonically decreases with bank assets almost everywhere, but there is a break for $a_t = \bar{a}_t$, when the level of assets reaches the banking sector's absorption capacity. Above this threshold, a credit crunch occurs and the corporate loan rate suddenly jumps to $\tilde{R}_t \equiv \lim_{a_t \searrow \bar{a}_t} R_t$, with $\tilde{R}_t > \bar{R}$. Notice that, from a *partial* equilibrium perspective, \tilde{R}_t is high enough to restore banks' incentives and reignite the interbank market. But this situation is not sustainable as a *rational expectation general* equilibrium, since by issuing interbank claims banks would be able to raise their supply of corporate loans and, ultimately, R_t would go down below \bar{R} . It follows that the autarkic equilibrium is the only interbank market equilibrium that is consistent with the general equilibrium when $a_t > \bar{a}_t$.²⁴ Panel (b) of Figure 4 depicts the equilibrium rates as functions of z_t for a given level of assets a_t and mirrors Panel (a). It shows that a crisis breaks out as soon as z_t falls below \bar{z}_t . Altogether, these two figures suggest that a SBC may result either from the *endogenous* over-accumulation of assets by the household (Panel (a)), or from an *exogenous* adverse productivity shock (Panel (b)). This variety of banking crises is an important and novel feature of our model, which we will discuss in detail in Sections 4.2 and 5.

Two regimes prevail in the economy. The first regime corresponds to normal times, where the interbank market functions well and all the assets available in the economy are used to finance the firm. The other regime corresponds to crisis times, where the interbank market shuts down, preventing the efficient re-allocation of assets. In this case, bank inefficiencies materialize themselves as a widening of the interest rate spread, which is due to the simultaneous increase in the corporate loan rate and the fall in the return on bank asset/equity. Which regime prevails depends on the size of the banks relative to their absorption capacity and on whether bank lending is "excessive". In other words, credit booms may be bad. To tell apart the bad from the good credit booms, we define the probability of a crisis at a n -period horizon as the joint probability that the banking sector's total assets exceed its absorption capacity in $t+n$ (i.e. that $a_{t+n} > \bar{a}_{t+n}$) and not before (i.e. $a_{t+i} \leq \bar{a}_{t+i}$ for $i = 1, \dots, n-1$).

²⁴In other words, condition $R_t \geq \bar{R}$ is necessary but not sufficient to rule out interbank market freezes.

Figure 4: Interest Rates



Definition 2 (Probability of a crisis at a n -period horizon) *Given the data generating process of productivity, the state of the nature at the end of period t (a_{t+1}, z_t) and the optimal asset accumulation rule $a_{t+i+1}(a_{t+i}, z_{t+i})$, the probability that a systemic banking crisis next breaks out in period $t + n$ is $\mathbb{P}(a_{t+1} \leq \bar{a}_{t+1}, \dots, a_{t+n-1} \leq \bar{a}_{t+n-1}, a_{t+n} > \bar{a}_{t+n}) = \int_{\bar{\varepsilon}_{t+1}}^{+\infty} \dots \int_{\bar{\varepsilon}_{t+n-1}}^{+\infty} \int_{-\infty}^{\bar{\varepsilon}_{t+n}} dG(\varepsilon_{t+1}, \dots, \varepsilon_{t+n})$, where $G(\cdot)$ denotes the cumulative of the n -variate Gaussian distribution, $n > 1$, and $\bar{\varepsilon}_{t+i} \equiv \ln \bar{z}_{t+i} - \rho_z \ln z_{t+i-1}$, where \bar{z}_{t+i} is the threshold of productivity in period $t + i$ below which, given the level of financial assets a_{t+i} , a crisis breaks out — i.e. $\bar{z}_{t+i} \equiv f_k^{-1}(\bar{R} + \delta - 1; a_{t+i})$.*

The crisis probability provides an early warning signal of banking crises. Of course, other variables in the model could be used to construct more standard indicators, like the credit to output ratio or the growth rate of credit. But these statistics do not contain as much information about future crises as the crisis probability does, because the latter is fully consistent with general equilibrium effects, agents' rational expectations, and perceived risks. For instance, the ex-ante anticipation of a market freeze leads the household to accumulate assets faster so as to smooth consumption should the market indeed freeze. By doing so, however, the household feeds a supply-driven credit boom, making the crisis more likely ex post. Hence the high crisis probability ex ante.

4 Calibration and Solution of the Model

We assess quantitatively the ability of the model to account for the dynamics of the economy before, during, and after a banking crisis. To do so, we extend the model to the presence of endogenous labor supply decisions.²⁵ The technology is assumed to be represented by

²⁵The equations characterizing the general equilibrium in this case are reported in Appendix B.

the standard constant–returns to scale Cobb–Douglas function, $F(k_t, h_t; z_t) \equiv z_t k_t^\alpha h_t^{1-\alpha}$ with $\alpha \in (0, 1)$. The household is assumed to be endowed with preferences over consumption and leisure that are represented by the following Greenwood et al. (1988) utility function

$$\mathbb{E}_t \sum_{\tau=0}^{\infty} \beta^\tau \frac{1}{1-\sigma} \left(c_{t+\tau} - \vartheta \frac{h_{t+\tau}^{1+v}}{1+v} \right)^{1-\sigma}.$$

It is well known that these preferences wipe out wealth effects in labor supply decisions, and we chose this specification for practical reasons. By yielding a closed form solution for the absorption capacity of banks

$$\bar{a}_t = \Gamma z_t^{\frac{1+v}{v(1-\alpha)}}, \text{ with } \Gamma \equiv \left(\frac{1-\alpha}{\vartheta} \right)^{\frac{1}{v}} \left(\frac{\alpha}{\bar{R} + \delta - 1} \right)^{\frac{v+\alpha}{v(1-\alpha)}},$$

that remains independent of the consumption level, it indeed preserves the sequential character of the resolution of the model. (As we will show later, this assumption plays against our model as the absence of a wealth effect on labor will be one reason why our model overstates the amplitude of banking crises compared with the data.) The absorption capacity now depends in a fundamental way on v , the (inverse of) labor supply elasticity. The larger labor supply elasticity, the more elastic to productivity shocks the absorption capacity of banks. The reason is that, in this case, positive productivity shocks entail a larger increase in hours worked, hence a larger increase in the marginal productivity of capital, in the corporate loan rate, and ultimately in the absorption capacity.

4.1 Calibration

The model is calibrated on a yearly basis, in line with Jordà et al. (2011a)’s database, which we used to document the stylized facts on banking crises in Section 2. We use a rather conventional calibration of the model (see Table 2). The discount factor, β , is set such that the household discounts the future at a 3% rate per annum. The inverse labor supply elasticity is set to $v = 1/3$ which lies within the range of values that are commonly used in the macro literature. The labor disutility parameter, ϑ , is set such that the household would supply one unit of labor in the average steady state. The curvature parameter σ is set to 4.5, which lies within the range of estimated values for this parameter.²⁶ However, given the importance of this parameter for the dynamics of precautionary savings, we will assess the sensitivity of our results to changes in this parameter. The capital elasticity in the production function is set to $\alpha = 0.3$ and capital is assumed to depreciate at a 10% rate per annum ($\delta = 0.1$). The process for the technology shock is estimated using the annual total factor productivity series produced by Fernald (2009) for the US over the post–WWII period. We obtain a persistence parameter $\rho_z = 0.9$ and a standard deviation $\sigma_z = 1.81\%$.

²⁶Mehra and Prescott (1985), for example, impose an upper bound of 10.

The remaining parameters pertain to the banking sector and include the return on storage γ , the cost of diversion θ , and the distribution of banks $\mu(\cdot)$. For tractability reasons we assume that $\mu(p) = p^\lambda$, with $\lambda \in \mathbb{R}^+$. The parameters of the banking sector are calibrated jointly so that in the simulations the model generates (i) a systemic banking crisis every forty years on average, *i.e.* with probability around 2.5% (see Table 1), (ii) an average interest rate spread equal to 1.71%, and (iii) an average corporate loan rate of 4.35%.²⁷ These latter two figures correspond to the lending rate on mid-size business loans for the US between 1990 and 2011, as reported in the US Federal Reserve Bank’s Survey of Terms of Business Lending —see Appendix C.2. We obtain $\gamma = 0.94$, $\lambda = 24$, and $\theta = 0.1$. Based on this calibration, the model generates an average interbank loan rate of 0.86% and an implied threshold for the corporate loan rate of 2.43% (*i.e.*, $\bar{R} = 1.0243$).

Table 2: Calibration

Parameter	Values	
Discount factor	β	1/1.03
Risk aversion	σ	4.500
Frisch elasticity	ν	1/3
Labor disutility	ϑ	0.944
Capital elasticity	α	0.300
Capital depreciation rate	δ	0.100
Standard dev. productivity shock	σ_z	0.018
Persistence of productivity shock	ρ_z	0.900
Bank distribution; $\mu(p) = p^\lambda$	λ	24
Diversion cost	θ	0.1
Storage technology	γ	0.936

4.2 Optimal Asset Accumulation Rule

The model is solved by a collocation method.²⁸ We first discretize the technology shock using the method developed by Tauchen and Hussey (1991), with 31 nodes for the Gauss–Hermite approximation. The decision rule for a_{t+1} is approximated by a function of Chebychev polynomials of the form

$$a_{t+1}(a_t; z_i) = \exp \left(\sum_{j=0}^q \xi_j^N(z_i) T_j(\varphi(a_t)) \right) \mathbb{I}_{a_t \leq \bar{a}_t} + \exp \left(\sum_{j=0}^q \xi_j^C(z_i) T_j(\varphi(a_t)) \right) \mathbb{I}_{a_t > \bar{a}_t},$$

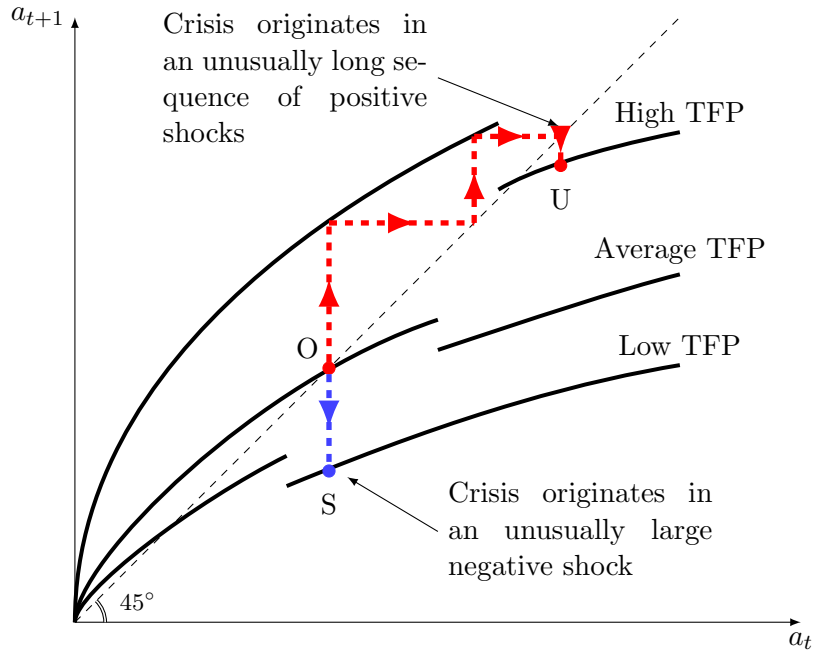
where z_i denotes a particular level of the total factor productivity in the grid. $T_j(\cdot)$ is a Chebychev polynomial of order j , $\xi_j^N(z_i)$ (resp. $\xi_j^C(z_i)$) is the coefficient associated to this

²⁷Because precautionary savings play an important role in our model there is a significant gap between the deterministic and the stochastic steady states. It is therefore more accurate to calibrate the model based the moments calculated over the simulations of the model, rather than on the basis of the deterministic steady state.

²⁸The companion technical appendix provides greater details on this solution method.

polynomial when the economy is in normal times (resp. in a systemic banking crisis) for the value of productivity z_i . We use 15th order Chebychev polynomials ($q=15$). Finally, $\varphi(a_t)$ is a function that maps the level of assets into the interval $(-1,1)$,²⁹ and $\mathbb{I}_{a_t \leq \bar{a}_t}$ (resp. $\mathbb{I}_{a_t > \bar{a}_t}$) is an indicator function that takes value one in normal (resp. crisis) times and zero otherwise. We allow for a discontinuity in the rule at the points when total assets reach the banking sector's absorption capacity—that is, when $a_t = \bar{a}_t$ —because in those points the economy switches regime. (That \bar{a}_t is a known function of state variable z_t greatly simplifies the solution algorithm.) The optimal decision rule $a_{t+1}(a_t; z_t)$ is given by the fixed point solution to the Euler equation (4). The parameters of the approximated decision rule are set such that the Euler residuals are zero at the collocation nodes. We checked the accuracy of our solution using the criteria proposed by Judd (1992).

Figure 5: Optimal Decision Rules



Note: Dashed line: 45° line where $a_{t+1} = a_t$.

Figure 5 illustrates the optimal asset accumulation decision rules for the lowest, the average, and the highest productivity levels in our grid.³⁰ As in the standard, frictionless neoclassical model, the household smooths her consumption over time by accumulating relatively more (less) assets when productivity is above (below) average. But the pace of accumulation is faster than in the frictionless model because of the possibility of crises, which lead the household to also save for precautionary motives. We will discuss this point in the next

²⁹More precisely, we use $\varphi(a_t) = 2 \frac{\log(a_t) - \log(y)}{\log(x) - \log(y)} - 1$, where $x \in \{a_{sup}, \bar{a}_i\}$ and $y \in \{a_{min}, \bar{a}_i\}$. a_{min} and a_{sup} denote the bounds of the interval values for a_t (we use $a_{min} = 0.5$ and $a_{sup} = 20$), z and y denote respectively the upper and the lower bound values we use in normal times and during a systemic banking crisis.

³⁰For illustrative and exposition purposes, we report stylized representations of the decision rules. The solution decision rules for the benchmark calibration are reported in the companion technical appendix.

section.

The optimal decision rules provide a first insight into the dynamics of the model. Starting from the average steady state O, the economy may face a systemic banking crisis for two different reasons. It may first experience a large negative technology shock, which brings the economy down to S. Because this shock instantaneously reduces banks' absorption capacity below the current level of assets, the crisis breaks out on impact. Such a crisis is purely driven by the bad realization of the exogenous shock and the amplification mechanism that the shock triggers. This is the channel usually depicted in the existing literature. In our model, a crisis may occur for another reason. As the economy experiences an unusually long sequence of TFP levels above average, the household has time to accumulate a large stock of assets, which may outgrow banks' absorption capacity, as in point U. In this case the crisis results from the optimal response of the household to positive, as opposed to negative, events. Next section describes such credit–boom led crises in more details.

5 Quantitative Analysis


5.1 The Economy in Normal Times

We start by analyzing the dynamics of the economy in response to a positive one standard deviation productivity shock about the average steady state. Figures 6–7 report the dynamics of our model (plain line) and those as obtained from the basic frictionless neoclassical model (dashed line), which corresponds to our model when $\gamma = 0$, $\theta = 0$, or $\lambda = +\infty$.³¹

Figure 6 reveals that the model behaves very similarly to the standard frictionless model in the face of a small positive shock. In both models, output, consumption, investment and hours worked increase on impact. After one period, the dynamics of the two models depart. The dynamics generated by our model is hump-shaped, whereas the frictionless economy goes back monotonically to the steady state. The hump-shaped pattern comes from a financial accelerator mechanism similar in spirit to that in Bernanke et al. (1999) that amplifies the effects of the shock. Indeed, the corporate loan rate rises on impact (see Figure 7), which mitigates counterparty fears on the interbank loan market and relaxes banks' borrowing constraints. As every borrower raises its market funding, the aggregate demand for interbank loans increases and so does the interbank loan rate. Inefficient banks then switch from the demand to the supply side of the market. This works to raise borrowers' overall quality and reduces the moral hazard problem further. Unlike Bernanke et al. (1999), however,

³¹We report the average of the distribution of the dynamic paths in the economy, as obtained from 100,000 simulations of the model. Notice that, despite the existence of a storage technology, our frictionless model is very similar to –and in effect indistinguishable from– the basic neoclassical growth model. The reason is that for this technology to be operated, one would need an unrealistic, zero-probability, negative TFP shock to occur.

Figure 6: Impulse Response to a One Standard Deviation Technology Shock (I)



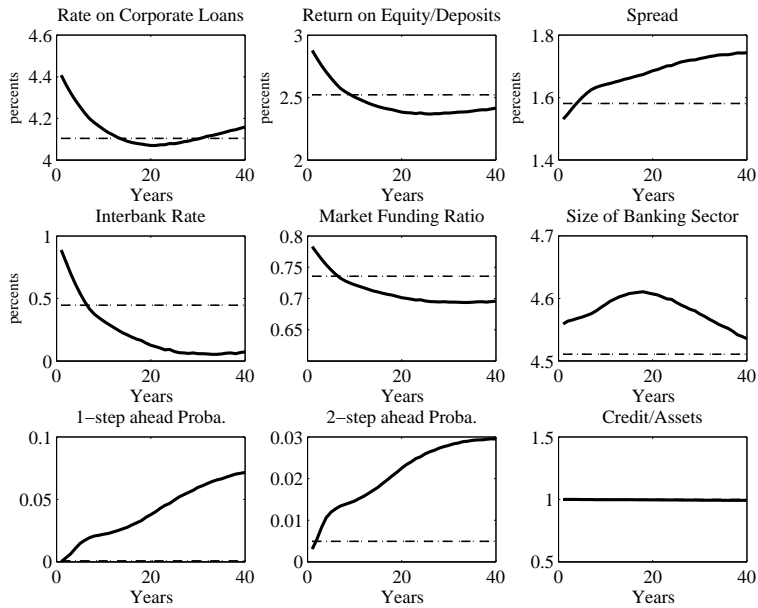
Note: Plain line : Average impulse response function across 100,000 simulations in the model, Dashed line: Average Impulse Response across simulations in the frictionless economy.

the financial accelerator mechanism comes from frictions on the wholesale loan market —not from frictions on the retail loan market. The response of financial variables in Figure 7 makes it clear that the economy does not experience any systemic banking crisis. Even though the corporate loan rate eventually falls below its steady state value as the household accumulates assets, at no point in the dynamics does it fall below \bar{R} (*i.e.* 2.43%). In effect, the positive technology shock does not last long enough to have the household accumulate assets beyond banks' absorption capacity. We obtain similar (mirror) results after a negative one standard deviation productivity shock. Most of the time, the model behaves like a standard financial accelerator model; crises are indeed rare events that occur under specific conditions, as we show in the next section.

5.2 Typical Path to Crisis

The aim of this section is to describe the typical conditions under which systemic banking crises occur. As we already pointed out (see Section 4.2), banking crises may break out in bad as well as in good times in the model. It is therefore not clear which type of shocks (negative/positive, large/small, short/long lived) are the most conducive to crises. Starting from the average steady state (*i.e.* $z_t = 1$), we simulate the model over 500,000 periods, identify the years when a crisis breaks out, and compute the median underlying technological path in the 40 (resp. 20) years that precede (resp. follow) a crisis. This path corresponds to the typical sequence of technology shocks leading to a crisis. We then feed the model with

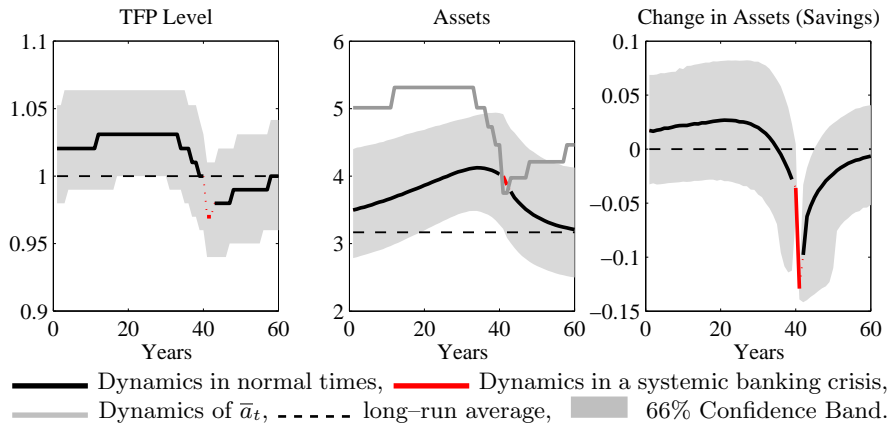
Figure 7: Impulse response to a 1 Standard Deviation Technology Shock (II)



Note: Plain line : Average path across 100,000 simulations in the model, Dashed line: Steady state value.

this sequence of shocks. The left panel of Figure 8 reports the typical path for the technology

Figure 8: Typical path (I)



shock, along with its 66% confidence interval. The red part of the path corresponds to crisis periods, the black one is associated with normal times. One striking result that emerges from this experiment is that the typical banking crisis is preceded by a long period during which total factor productivity is above its mean. In some 20% of the cases, crises even occur at a time when productivity is still above mean. This reveals one important and interesting aspect of the model: the seeds of the crisis lie in productivity being above average for an unusually long time. The reason is that a long period of high productivity gives the household enough time to accumulate assets beyond the banking sector’s absorption capacity (see middle panel in Figure 8). This is a phase where financial imbalances build up, as reflected in the increase

in the one-step ahead probability from an initial 0 to 0.25 few years before the banking crisis in Figure 9.³²

In our model, the typical crisis breaks out in the middle of a credit boom. At the beginning both productivity and the corporate loan rate are above their average steady state, and the credit boom is demand-driven. As productivity goes back to average, the firm reduces its demand for labor, capital, and credit. Following the decrease in her labor income, the household reduces her savings and even dissaves (see right panel in Figure 8). But this reduces only marginally the household's total stock of assets and the overall supply of credit by banks, which remains more than 30% above steady state level. At this point in time the credit boom turns supply-driven and the corporate loan rate goes down below its average steady state (see Figure 9).³³ The fall in the corporate loan rate is followed by a fall in the interbank loan rate, which gives inefficient banks incentives to demand interbank loans. This has two important consequences. First, since more banks finance themselves on the market, the aggregate balance sheet of the banking sector increases in size.³⁴ Second, the fact that more inefficient banks raise funds erodes trust between banks. It follows that the banking sector gets bigger precisely at the time when its absorption capacity diminishes, making it less resilient to adverse shocks. In this context, the rather mild 3% drop of productivity below its average in year 40 is enough to bring the whole sector down. The crisis manifests itself as a sudden 10% fall in the credit to assets ratio, k_t/a_t , and a 30% drop in the size of the banking sector, $a_t + (1 - \mu(\bar{p}_t))\phi_t a_t$. As banks cut their supply of corporate loans, the spread rises from 2% to almost 4%. Interestingly, the rise in the spread is due to the return on equity/deposits falling by more than the corporate loan rate during the crisis, a phenomenon for example observed in 2007–8 in the US.³⁵ The dynamics of the corporate loan rate results from two opposite effects. On the one hand, the crisis brings about a credit crunch that implies, all other things equal, a rise in the corporate loan rate (see Figure 4). On the other hand, the fall in productivity also implies, in the first place, a fall in the firm's demand for corporate loan, and therefore a decrease in the corporate loan rate. In the typical crisis, this latter effect dominates. Ultimately, the corporate loan rate decreases, but by much less than it would otherwise have in normal times following a similar —3%— drop in productivity. After the crisis, all financial variables return back to their steady state levels.

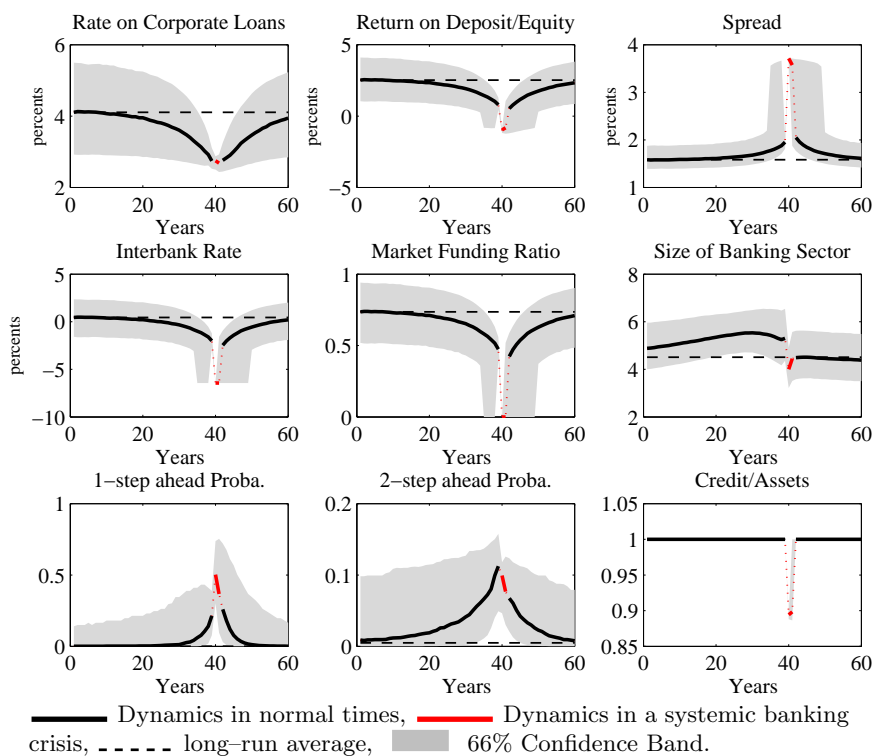
³²Likewise, the two-step ahead probability goes up to 0.15. These probabilities are constructed following Definition 2. Hence, the two-step ahead probability in period t is the probability that a banking crisis will break out in period $t + 2$, conditional on the event that there is no crisis in period $t + 1$.

³³To the best of our knowledge —because of lack of data— there exists no clear historical pattern of the dynamics of interest rates before SBCs. However, the low level of real US interest rates has been shown to be one important factor of the recent crisis (Taylor, 2009). The latter was indeed preceded by particularly low real fed fund rates, with an average of 0.84% over 2000-2007 despite the monetary policy tightening in 2005-2006, against 3.24% over 1986-2000; see Figure C.4 in Appendix C.2.

³⁴On the brink of the typical crisis, the total assets of the banking sector are about 30% above steady state. In the run up to the crisis, borrowing banks' market funding ratio (ϕ_t) goes down so as to maintain private incentives, which is consistent with the pre-2007 crisis developments for the US. But since more banks become leveraged, the leverage of the banking sector as a whole, $(1 - \mu(\bar{p}_t))\phi_t$, increases.

³⁵See Figure C.4 in Appendix C.2.

Figure 9: Typical path (II)



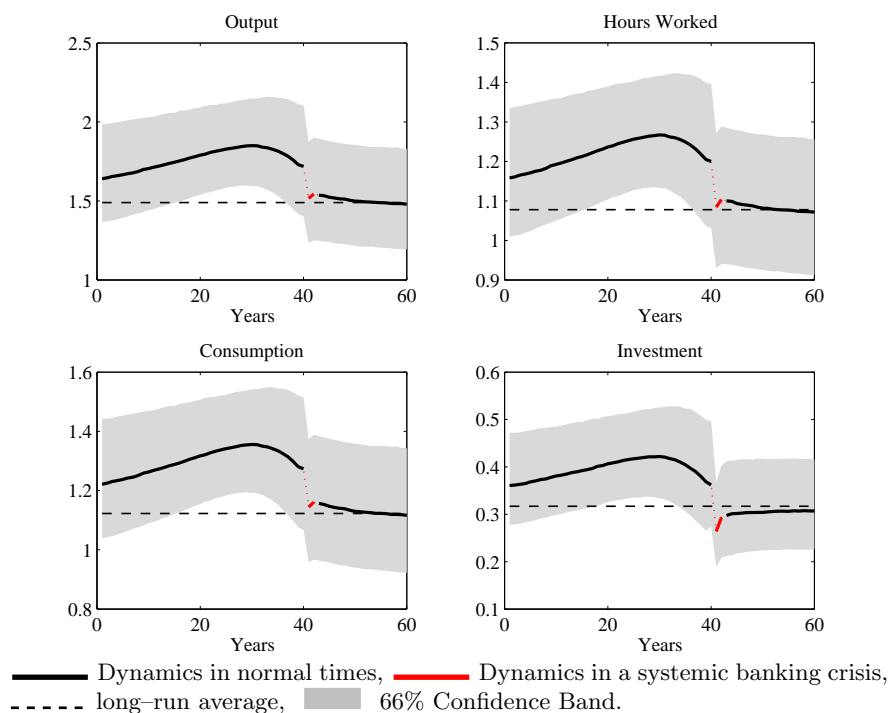
Finally, Figure 10 illustrates the evolution of macroeconomic variables. In the run up to the crisis, the positive wealth effect associated with technological gains leads the household to both consume and invest more. The household also increases her labor supply in order to take advantage of the high real wage, so that the level of output is 15% above its steady state in the year before the crisis. These unusually good times end abruptly with the credit crunch, which triggers a sizable drop in investment, labor, consumption, and output. In the year of the crisis, output falls by 15%. Part of this loss —around 3%— is due to the fall in productivity that takes place as the crisis hits, and would have taken place even in the absence of the crisis as the result of the standard mean-reverting dynamics of productivity; the remaining 12% of the fall in output is attributable to the banking crisis *per se*. Notice that output falls by more than in the data (6.74% —see Table 1), suggesting that financial frictions in the model are relatively severe.³⁶ However, the performance of the model can be improved on that front *(i)* by allowing the firm to issue corporate bonds or equity directly to the household so as to reduce the prevalence of the banking sector in firm financing, or *(ii)* by slowing down the accumulation of assets by the household and curbing the credit booms, e.g. through wealth effects in labor supply decisions or capital adjustment costs.³⁷ Overall, the

³⁶For instance, the assumption that banks live for only one period makes information asymmetries particularly prevalent. One way to mitigate the information problem would be to have banks living several periods with persistent skills. (An extreme case is one where skills are permanent, as in this case banks' profits would perfectly reveal banks' skills and fully dissipate the asymmetric information problem after the first period.) This, however, would be at the price of a lack of tractability of the model.

³⁷On *(i)*: see Shin et al. (2011) for some evidence that substitution effects between corporate bonds and

typical SBC breaks out as productivity returns back to its average level after having remained above it for an unusually long time. This corresponds to a situation where the household has accumulated assets in anticipation of a fall in productivity that took more time than expected to materialize. In effect, one can show that the household over-accumulates assets, compared to what a central planner would, because she does not internalize the effects of her savings decisions on the equilibrium corporate loan rate and, ultimately, on the likelihood of a banking crisis.

Figure 10: Typical path (III)




The above description of the typical path to crises, however, conceals the variety of banking crises, which the model features. In particular, SBCs may also be caused by an exogenous large negative TFP shock. For example, in Appendix C.1 we consider the response of the economy to such a large shock and show that, in this case, the crisis breaks out on impact. A 7% drop from the average steady state of TFP is sufficient to generate a crisis. To get a sense of the frequency of such shock-driven crises with respect to the typical credit boom-led ones, we report in Figure 11 the distributions of the deviations from their average levels of (the log of) TFP, the credit/output ratio, and a measure of financial imbalances in the k th period before an SBC erupts, for $k = 1, 5, 10, 20$.

Panel (a) shows that few SBCs occur after a year where productivity is significantly below its average steady state. (For instance only 10% of SBCs follow a year where TFP is 5% below

bank loans played an important role in mitigating the impact of the 2007–2008 crisis in the US, and De Fiore and Uhlig (2012) for a DSGE model with such effects. On (ii): The difficulty of introducing such wealth effects in our model lies with the numerical resolution of the equilibrium, as in this case the threshold \bar{a}_t would not depend exclusively on state variable z_t anymore.

Figure 11: Distribution of percentage deviations before an SBC



Note: The figure above shows ergodic cumulative distributions. Panels (a) and (b) show $\mathbb{P}(\Delta x_{T-k} \geq \zeta)$ in the case of positive percentage deviations ζ (x-axis, right hand side) from average steady state, and $\mathbb{P}(\Delta x_{T-k} \leq -\zeta)$ in the case of negative deviations $-\zeta$ (x-axis, left hand side), where $\Delta x_{T-k} = \log(x_{T-k}) - \log(\bar{x})$, T is the date of the crisis, \bar{x} is the average steady state of variable x calculated over our 500,000 simulations, and x is the logarithm of TFP or credit/output ratio. Panel (c) shows the probability that the level of assets k periods before the crisis be less than $\zeta\%$ away from the absorption capacity of the banking sector, $\mathbb{P}(\Delta a_{T-k} \leq \zeta)$, with $\Delta a_{T-k} = \log(\bar{a}_{T-k}) - \log(a_{T-k})$.

average.) TFP is moreover almost symmetrically distributed around its average in the year before an SBC, which essentially means that the sign of TFP shocks —or more precisely: whether TFP is below or above its average— is not in itself informative about whether a SBC is about to break out. The skewed distributions of lagged TFP (for $k = 5, 10, 20$) at a medium to long term horizon even suggests that banking crises tend to be preceded by rather long periods of high —and not low— TFP. This is because prolonged sequences of TFP levels above average feed credit booms; and credit booms are conducive to SBCs. Panel (b) shows for instance that the credit to output ratio is 5% above its average steady state in the year that precedes the SBC in 95% of the SBCs. These results point to a general and important property of our model. Namely, the size of financial imbalances that prevail at the time shocks hit the economy are more relevant for the dynamics of the economy than the size and sign of the shocks themselves. Panel (c) illustrates this point more accurately; it shows the distributions of a measure of financial imbalances k periods before SBCs (again, for $k = 5, 10, 20$). We measure financial imbalances by the distance between banks’ absorption capacity and banks’ core liabilities, $\log(\bar{a}_t) - \log(a_t)$. The smaller this distance, the larger the imbalances, and the less resilient to adverse shocks the banking sector. Typically, a distance of less than, say, 5% reflects large financial imbalances. Panel (c) shows that 70% of the SBCs break out in the year after financial imbalances have reached this threshold. This result confirms that most crises in our model break out endogenously, without an adverse exogenous shock happening at the same time.

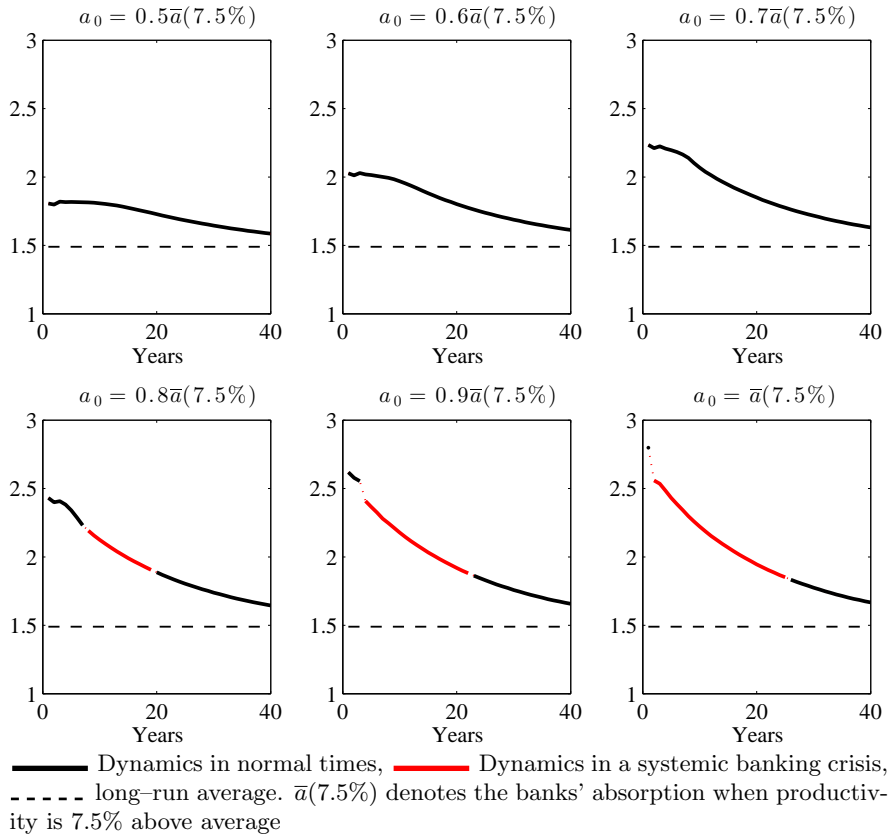
The above discussion prompts the question why, in our model, crises are more likely to occur in good, rather than in bad, times. This is due to the asymmetric effects of permanent income mechanisms on financial stability over the business cycle. Bad times in the model are typically times where TFP is low and the household dis-saves. By lowering the TFP threshold \bar{z}_t , below which a crisis breaks out, the fall in savings makes crises less likely. Hence, in bad times, the dynamics of savings tends to stabilize the financial sector. In good times, in contrast, TFP is high and the household accumulates assets, which by raising \bar{z}_t makes crises more likely. In this case the dynamics of savings tends to destabilize the financial sector. This asymmetric effect of savings is the basic reason why credit–boom led crises are so prevalent in our model. Figure 16 in next section also illustrates this point.

5.3 The Role of Financial Imbalances

This section focuses on the effects of financial imbalances on the frequency, the duration, and the amplitude of banking crises. As above, we measure financial imbalances as the distance between \bar{a}_t and a_t . The experiment compares the transition dynamics of the economy toward its average steady state as the initial asset position of the agent varies. Building upon the previous experiment, which showed that crises are most likely to occur after a productivity

boom, we impose that productivity is initially 7.5% above average in all the simulations. We simulate the economy starting from six different initial levels of assets ranging from 0% (large financial imbalances) to 50% (small imbalances) below absorption capacity.³⁸ Figures 12 and 13 report the median adjustment dynamics of output and the credit to assets ratio across 500,000 simulations.³⁹

Figure 12: Output dynamics: Sensitivity to Initial Conditions



We find that the closer the banking sector to its absorption capacity, the more likely and the sooner the crisis. No crisis occurs so long the initial asset position of the economy is at least 30% below the absorption capacity. Financial imbalances pave the way to crises. Interestingly, initial conditions not only affect the probability and the timing of a crisis, but also its duration and amplitude. Figure 13 thus shows that the length of the credit crunch raises from 12 to 19 periods as the initial level of assets goes from 20% to 10% below absorption capacity. In that case the initial drop in output raises from 1.85% to 5.88% (see Figure 12).

Figure 14 illustrates the banking sector's resilience to shocks by reporting the proportion of

³⁸We do not take a stance on the reasons why the economy reached these particular levels of assets in the first place, as this is not relevant here.

³⁹The mean path would be the usual representation. However this representation would average out the effects of the financial crises, which do not break out along every one of the 500,000 simulated paths. This is the reason why, in the context of our model the median path is more informative.

Figure 13: Credit/Assets Dynamics: Sensitivity to Initial Conditions





Figure 14: Frequency of SBCs: Sensitivity to Initial Conditions



This figure reports the evolution of the frequency of SBCs during the transition toward the average steady state.

paths featuring SBCs along the transition dynamics. In the economy with small imbalances (upper left panel), this proportion does not exceed 20% at any point in time. This suggests that, for this economy to experience a SBC, extreme TFP levels (and therefore extreme sequences of TFP shocks) in the 20% tail of the TFP distribution must occur. In contrast, the economy with the largest imbalances (lower right panel) is much more fragile and unlikely to escape a crisis, irrespective of the size of the shocks that hit it. In this economy more than 80% of the trajectories indeed feature a SBC.

The Role of Permanent Income Mechanisms: Financial imbalances are brought about by the household’s rational expectations and permanent income behaviors —consumption smoothing and precautionary savings. As imbalances build up, the representative household realizes that the likelihood of a crisis increases, and accumulates assets to smooth her consumption profile. By doing so, however, the household inadvertently exacerbates the imbalances that may precipitate the crisis. To illustrate and assess quantitatively the importance of saving behaviors in the model, we replicated the above experiment using a “Solow” version of our model, where the saving rate is constant and the household neither smooths consumption nor accumulates precautionary savings.⁴⁰ We set the saving rate, s , such that the economy reaches the same average steady state as in our benchmark economy ($s = 0.22$). For comparison purposes, we also set the initial level of productivity 7.5% above average in all the simulations, and only let the initial level of assets vary. Because in this version of the model only large imbalances lead to a crisis, we only report in Figure 15 the results for initial levels of asset greater than $0.8\bar{a}(7.5\%)$.

Panel (a) of Figure 15 shows that the household accumulates assets more slowly in the absence (plain line) than in the presence (dashed line) of permanent income mechanisms. The upshot is that SBCs are less frequent, shorter, and milder in the Solow version of the model. When the economy is started 10% below absorption capacity (middle panel), for instance, the crisis takes more time to materialize and entails a smaller credit crunch, with the credit to assets ratio falling by 10.3%, against 11.1% in the benchmark. In this case output drops by 3 percentage points less and the crisis is 11 years shorter than in the benchmark.

More generally, financial imbalances do not build up as much, and therefore do not play as big a role on financial stability in the absence of permanent income mechanisms. To illustrate this point, we derive the typical path to crisis in the Solow economy and find that the typical crisis is shock-driven, as opposed to credit-boom driven (compare Figures 16 and 8).⁴¹ This result highlights the importance of saving dynamics to explain why, in our model, systemic

⁴⁰See the companion technical appendix for a detailed description of this version of the model.

⁴¹This result is also reflected in the difference between the optimal decision rules in our benchmark model and those in the Solow model, which all cross the 45 degree line in the normal times regime. For a comparison between the optimal decision rules, see the companion technical appendix.

Figure 15: The Role of Permanent Income Mechanisms

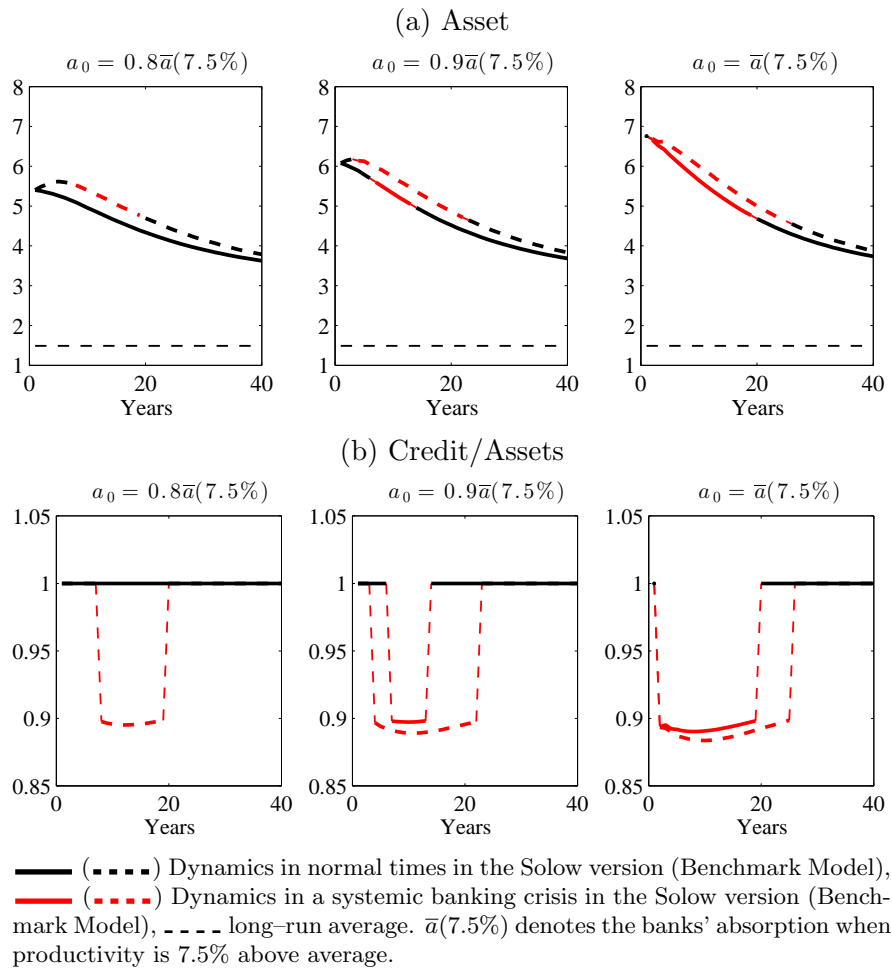



Figure 16: Typical path: Solow model



banking crises break out in the midst of credit intensive booms (Fact #3).

5.4 Assessment of the Model

We now assess the ability of the model to account for the main stylized facts on systemic banking crises (see Section 2). As a first step, we simulate the model and generate long time series (500,000 periods) for output and credit. We then identify all recessions. A “recession” in the simulations starts when the economy experiences growth below a threshold that guarantees that, on average, the model generates recessions 10% of the time, as in the actual data (see Table 1). We further identify as financial recessions those recessions that feature a systemic banking crisis. Table 3 reports summary recession statistics similar to those reported in Table 1.

Table 3: Statistics on recessions and banking crises

	Frequency (%)	Magnitude (%) from peak to trough	Duration (Years)
Systemic Banking Crises (SBC)	2.69	–	–
All recessions	10.00	12.08 (7.30)	2.08
Recessions with SBC	13.00	17.87 (10.50)	2.62
Recessions w/o SBC	87.00	10.04 (6.73)	1.90


Note: As in Table 1, the magnitudes reported into parentheses are calculated using the HP-filtered series of output, and are thereby corrected for the underlying trend in output. Notice that the average amplitude (duration) of financial recessions is slightly higher, with a 17.87% fall in output (2.62 years), than along the typical path in Figure 10, where it is around 15% (2 years). These differences are due to the non-linearities in the model, which imply that the average response of the economy to shocks (as reported above) differs from the response of the economy to the average shock (as reported in Figure 10).

On average, the duration of a recession featuring a banking crisis is 26% larger than that of other recessions. It also replicates the observation that recessions with SBCs are deeper than the average recession (48%). The model, however, overall over-estimates the amplitude and duration of recessions (compare Table 3 with Table 1). As discussed in section 5.2, this result is due to (i) our assumptions that firm financing is fully intermediated and (ii) the absence of wealth effects on labor supply decisions. Since the model was calibrated so as to obtain systemic banking crises every forty years in the simulations, the model replicates —by construction— Fact #1 (*i.e.* crises are rare events). Hence, we focus on Fact #2 and Fact #3.

Fact #2: Financial recessions are deeper and last longer than other recessions. Figure 17 replicates Figure 1 and reports the average dynamics of the cyclical component of output and credit around financial recessions (dark line) and the other recessions (red

line). The striking observation from the figure is its similarity with the actual data. The

Figure 17: Dynamics of GDP and Credit gaps around Recessions




Note: For the sake of consistency, we treated the simulated series of output and credit in the same way as we treated the actual data in Section 2. Output and credit are first expressed in logarithm and the cyclical component is obtained using the HP-filter with $\lambda = 6.25$. Thus, as in Figure 1, the reported % deviations are the average % deviations around a Hodrick–Prescott trend. The implied magnitude of financial recessions (the left chart) is about 7.1%, which is lower than that 10.50% reported in Table 3. This discrepancy reflects the fact that the statistics in Table 3 also take into account recessions of more than 6 years. We find similar results when we consider the % deviations of output and credit from their respective linear trends (see the companion technical appendix).

model is able to replicate the facts that financial recessions particularly deep and that both output and credit are much larger (5.2% and 3.4% above trend, respectively) before financial recessions than before regular recessions. In particular, as in the data, the credit gap drops to -1% during a SBC, much more than in a regular recession. The phase that follows the credit crunch also presents a U-shape pattern that reflects the greater persistence of financial recessions. Financial recessions last longer because they tend to occur in good times, when TFP is above mean and likely to go down. The exogenous fall in TFP that tends to follow upon the typical financial recession works to delay the recovery.

Fact #3: Systemic banking crises break out in the midst of credit intensive booms. Figure 18 replicates Figure 2 and reports the ergodic distributions of output and credit gaps in the full sample (red line) and in the year preceding a systemic banking crisis (histogram). This figure too bears a lot of similarities with its empirical counterpart. Both output and credit gaps are above trend in the year that precedes a banking crisis. One may argue that this result is not so surprising, since recessions (either in the data or in the model) should —by construction— follow expansions. However, we have seen with Figure 17 that the output and credit gaps are both much higher before a SBC than before a regular recession, and we find similar results for the distributions of GDP and credit growth rates (not reported). In other terms, these results are not driven by the way we filter the series.

Figure 18: Distributions of GDP and credit gaps




As a last experiment, we want to illustrate the prediction properties of our model. To do so, we compute the probabilities that a SBC breaks out at a one- and a two-year horizon —as defined in Definition 2— and compare them with the actual US data over the period 1960-2011. The US experienced two banking crises during this period, both within the last three decades: the Savings & Loans crisis in 1984, and the subprime crisis in 2008. Jordà et al. (2011a) do not count the 2000 dotcom bubble bust as a *banking* crisis. We are interested in whether our model is able to predict these events. In the model, the crisis probability in year t ($t=1960, \dots, 2011$) is fully determined by the level of productivity, z_t , and the level of assets in the banking sector, a_t . We derive this probability from the model’s internal dynamics that we feed with Fernald (2009)’s series of total factor productivity.⁴² Figure 19 reports the implied k -step ahead probability of a crisis for $k=1,2$ (left panel) as well as the underlying series of total factor productivity (right panel).

Considering its stylized character and the limited amount of information that it uses, the model predicts the US crisis episodes remarkably well. The one-step ahead probability (left panel) is essentially zero until the early 80s, spikes to 40% within three years, in 1981-1983, just before the Savings & Loans crisis, and remains relatively high throughout the rest of the period. While the probability is still 30% the year before the dotcom bubble bust, it goes down to 15% in 2001-2005 and picks up again as of 2007. At 20% the 2007 crisis probability is still relatively high by historical standards, and further increases in 2008-10, as the Great Recession materializes. The two-step ahead probability follows a similar pattern, with a one year lead. To a large extent, these results reflect the relationship in the model between financial stability and total factor productivity (right panel), which persistently stayed 10% above trend until the early 80s and has fluctuated below trend since then. We find similar

⁴²We detrended this series beforehand in order to be consistent with our stationary setup. More precisely, we linearly detrended the logarithm of Fernald’s “naive” productivity series over the sample 1947–2011. We use the whole sample to use as much information as available. We then use the so-detrended series to construct the level of productivity, z_t , for the period 1960–2011. For the initial value of bank assets, a_{1960} , we take the steady state level of assets associated with z_{1960} in the model.

Figure 19: k -step ahead Probabilities of a Financial Crisis ($k=1,2$)



Note: The vertical thin dashed lines correspond to the 1984 Savings & Loans, the 2000 dotcom and 2008 crises.

results when we use Fernald’s TFP series corrected for the rate of factor utilization, or when we detrend the series of TFP using a break in the trend to account for a structural US productivity slowdown in the mid 70s (see the companion technical appendix).

6 Discussion

6.1 Sensitivity Analysis

We now turn to the sensitivity of the properties of the model to the parameters. We simulate the model for 500,000 periods, compute the means of some key quantities across these simulations, and compare the results with our benchmark calibration (see Table 4).

Risk Averse Economies Are Prone to Crises: We first vary the utility curvature parameter σ from our benchmark 4.5 to values 2 and 10, therefore changing the degree of risk aversion of the household. By making the household more willing to accumulate assets for precautionary motives, *ceteris paribus*, the increase in σ works to raise the quantity of assets banks have to process without affecting banks’ absorption capacity — leaving banks more exposed to adverse shocks. The probability of a crisis is thus higher than in the benchmark (5.4% *versus* 2.7%). In other words, the risk averse economy is paradoxically more prone to systemic banking crises. It also experiences deeper and longer crises than the benchmark economy, with output falling by 1.1 percentage point more from peak to trough and crises lasting 1.4 year longer. The main reason is that, by accumulating more assets, the economy builds up larger imbalances that make it difficult to escape crises once they occur. Accordingly, the banking sector of the risk averse economy is also less efficient, with an interest rate spread of 2.09%, against 1.71% in the benchmark. In contrast, less risk averse economies are more resilient to systemic banking crises. When $\sigma = 2$, for instance, SBCs occur only with a

Table 4: Sensitivity Analysis

	Benchmark	σ		ν		θ	λ	σ_z	ρ_z
		2	10	0.25	1	0.20	35	0.02	0.95
<i>Returns:</i>									
Interbank (ρ)	0.86	1.04	0.23	0.84	0.83	0.40	1.34	0.89	0.72
Corporate (R)	4.35	4.55	3.70	4.28	4.41	5.50	3.70	4.32	4.29
deposit/equity (r)	2.64	2.96	1.61	2.52	2.80	2.61	2.67	2.55	2.59
Spread ($R - r$)	1.71	1.59	2.09	1.77	1.60	7.2.89	1.03	1.77	1.70
\bar{R}	2.43	2.43	2.43	2.43	2.43	4.83	0.41	2.43	2.43
<i>Crises:</i>									
Probability	2.69	1.20	5.43	3.31	0.99	7.34	0.16	3.35	1.90
Duration	2.62	1.89	4.08	2.90	1.97	5.06	1.87	2.82	2.92
Amplitude	17.87	17.30	19.00	19.94	11.96	16.90	15.80	19.36	16.08

Note: All numbers are averages over a long simulation of 500,000 periods and, except for durations, are expressed in percents. In the case where the persistence of the technology shock is raised to $\rho_z = 0.95$, the standard deviation of the innovation was rescaled so as to maintain constant the volatility of TFP.

1.20% probability.

Elastic Labor Supply Increases the Probability of a Crisis: We vary the (inverse of the) labor supply Frisch elasticity and consider values $\nu = 1/4$ and $\nu = 1$, the latter being rather standard in the DSGE literature (see Christiano et al., 2005). The Frisch elasticity affects the speed of accumulation of assets. In good times, for example, the household typically works more and, therefore, accumulates assets faster when $\nu = 1/4$ (high elasticity) than when $\nu = 1$ (low elasticity). Since financial imbalances build up faster in the high elasticity economy, this economy is also more prone to crises: when ν is reduced from 1 to 1/4, the probability of a crisis increases from 0.99% to 3.31%. Moreover, when the crisis breaks out the impact of the credit crunch on output is magnified by the larger response of hours worked; hence crises are also more severe and longer lasting in economies with higher labor elasticity.

Contract Enforceability and Bank Efficiency Improve Financial Stability: The third column of the table reports statistics for an economy where the cost of diversion is lower than in the benchmark. The increase in θ from 0.1 to 0.2 works, in the first place, to aggravate the moral hazard problem between banks. Banks have to deleverage in order to keep issuing interbank claims. Of course, in the general equilibrium there are some counterbalancing effects. For example, since the banking sector is less efficient the return on bank deposit/equity goes down, the household reduces her savings, banks reduce their supply of credit to the firm, and the equilibrium corporate loan rate tends to increase, which works to restore banks' incentives. But such effects are of second order. Overall, the crisis probability jumps from 2.69% to 7.34%.

A change in the distribution of banks has qualitatively similar effects. In the fourth column we consider an economy where $\lambda = 35$, implying a high concentration of the distribution of banks towards the top. The banking sector as a whole is more efficient than in the benchmark. Since efficient banks have lower incentives to divert cash, the moral hazard problem is less stringent and counterparty fears recede. Lenders tolerate a higher market funding ratio, aggregate demand and the interbank rate rise, which crowds the less efficient banks out of the demand side of the interbank loan market. As a result, the crisis probability drops from 2.7% to 0.16%. When they occur banking crises are however slightly more pronounced and more protracted. This is because they break out after a longer build-up phase, upon larger imbalances.

Higher Uncertainty is Conducive to Crises: As the only source of uncertainty, technology shocks play a crucial role in the model. In the last two columns in Table 4 we consider two experiments. First, we increase the volatility of the shock by raising σ_z from 0.018 to 0.02, leaving its persistence unchanged. The consequences are straightforward: the household accumulates more assets for precautionary motives, the corporate loan rate decreases with respect to benchmark, and the financial sector is more fragile. Next, we increase the persistence of the shock by raising ρ_z shock from 0.9 to 0.95, leaving its volatility unchanged. This change exerts two opposite effects on financial stability. On the one hand, following a positive productivity shock the household does not accumulate assets as fast as in the benchmark. Hence financial imbalances build up more slowly. On the other hand, the productivity level is more likely to remain high for longer, implying that imbalances accumulate over a longer period of time and, therefore, grow larger. In our calibration, the first effect dominates, and the probability of a crisis overall decreases.

6.2 Bank Leverage, Bank Defaults

One caveat of our model is that it abstracts from the dynamics of bank leverage and bank defaults, two elements yet important from a macro-prudential perspective. The reason is that, absent any friction between the banks and the household, the deposit to equity ratio is indeterminate. In this section we introduce a friction that permits us to pin this ratio down and to analyze the evolution of bank leverage and bank defaults along the typical path to crisis. For simplicity, we will keep a framework where the bank deposit to equity ratio does not have any impact on aggregate dynamics, so that the typical path to crisis is unchanged. To derive this ratio we make two additional assumptions. First, we assume that banks can collect deposits, which we define as the risk-free asset of the economy. Since deposits must be risk-free, there is a limit to the quantity of deposits that a given bank can collect. This limit is related to the value of the bank's assets in the worst possible state of the nature, *i.e.* γa_{t+1} for a bank born in t . (The worst possible state of the nature for a bank is one

where the bank is inefficient and there is a crisis.) Second, we assume that banks are run by risk-neutral managers, whose objective is to maximize their banks' expected return on equity. That managers do not discount risk—or have a different discount factor than the shareholder's—is a standard assumption that creates a wedge between the managers and the shareholder's objectives. While in equilibrium the shareholder will be indifferent between equity and deposits, managers will in contrast be willing to raise as much deposits as possible. To model defaults, we further assume that the shareholder has a *partial* liability, in the sense that she must recapitalize the banks that are unable to repay their deposits up to a fraction ξ of the banks' assets, with $\xi > 0$. The shareholder's liability thus amounts to ξa_{t+1} . Under this assumption, a manager is able to raise deposits up to $(\gamma + \xi)a_{t+1}$, which is more than what her bank will be able to repay, should a bad enough state of the nature materialize. We will say in this case that the bank defaults.⁴³ Let r_{t+1}^d be the risk-free (non-state contingent) gross return on deposits d_{t+1} and $r_{t+1}^e(p)$ be the *ex post* return on bank p 's equity e_{t+1} at the end of period $t + 1$ with, by definition, $a_{t+1} \equiv d_{t+1} + e_{t+1}$ and

$$r_{t+1}^e(p) \equiv \frac{r_{t+1}(p)a_{t+1} - r_{t+1}^d d_{t+1}}{e_{t+1}}, \quad (10)$$

where $r_{t+1}(p)$ is defined in (5). Since bank p defaults when $r_{t+1}^e(p) < 0$ the type \tilde{p}_{t+1} of the marginal bank that defaults at the end of period $t + 1$ is given by

$$\tilde{p}_{t+1} = r_{t+1}^{-1} \left(r_{t+1}^d \frac{d_{t+1}}{a_{t+1}} \right) \mathbb{I}_{\tilde{p}_{t+1} \leq r_{t+1}^{-1} \left(r_{t+1}^d \frac{d_{t+1}}{a_{t+1}} \right)}, \quad (11)$$

where $\mathbb{I}_{\tilde{p}_{t+1} \leq r_{t+1}^{-1} \left(r_{t+1}^d \frac{d_{t+1}}{a_{t+1}} \right)}$ is an indicator function that takes value one if $\tilde{p}_{t+1} \leq r_{t+1}^{-1} \left(r_{t+1}^d \frac{d_{t+1}}{a_{t+1}} \right)$ and zero otherwise. The household's no-arbitrage condition between deposits and equity then requires that

$$\mathbb{E}_t(u'(c_{t+1})r_{t+1}^e) = r_{t+1}^d \mathbb{E}_t(u'(c_{t+1})), \quad (12)$$

where $r_{t+1}^e \equiv \int_0^1 r_{t+1}^e(p) d\mu(p)$. This condition relies on r_{t+1}^d being a risk-free rate determined at the end of period t , yields r_{t+1}^d so that the household is indifferent between deposits and equity in equilibrium, and implies that $r_{t+1}^d < \mathbb{E}_t(r_{t+1}^e)$. Risk-neutral managers choose the

⁴³An alternative way to introduce bank defaults could be to assume that a deposit insurance scheme is in place, that makes it possible for the managers to raise more risk-free deposits. The difference with our setup is that deposit insurance is a subsidy to the banking sector, which would then have to be accounted for in the Euler equation (4). In the companion technical appendix of the paper, we show that deposit insurance would affect the dynamics of savings and the typical path to crisis, but only marginally. The constraint that a given bank's deposits are limited by the bank's equity capital (see constraint 14 below) has indeed very little impact on the aggregate dynamics in our model. The reason is that banks have the possibility to issue equity on the market so as to raise their deposits and, ultimately, increase their supply of corporate loans. That is, bank capital is not scarce. In this respect, our model stands in contrast with other models, like Gertler and Kiyotaki (2009), where banks are not allowed to issue equity. This restriction is very severe because it implies banks can only accumulate net worth through retained earnings. And in these models banks typically never live long enough to accumulate enough net worth out of financing constraints. As a result, bank capital is scarce and, by limiting the quantity of credit that banks can supply, it affects the aggregate dynamics. This is not the case here.

deposit to equity ratio so as to maximize their respective banks' expected returns on equity⁴⁴

$$\max_{d_{t+1}/e_{t+1}} \beta \mathbb{E}_t(r_{t+1}) + \beta \left(\mathbb{E}_t(r_{t+1}) - r_{t+1}^d \right) \frac{d_{t+1}}{e_{t+1}} \quad (13)$$

subject to the feasibility constraint that deposits must be risk-free,

$$d_{t+1} r_{t+1}^d \leq (\gamma + \xi) a_{t+1}. \quad (14)$$

Since (from (10) and (12)) $\mathbb{E}_t(r_{t+1}) > r_{t+1}^d$ it is clear that managers want to raise as much deposits as possible and that, in equilibrium, constraint (14) binds. Hence,

$$\frac{d_{t+1}}{e_{t+1}} = \frac{\gamma + \xi}{\max(r_{t+1}^d - \gamma - \xi, 0^+)}. \quad (15)$$

One can now easily derive borrowing (lending) banks' optimal leverage ℓ_{t+1}^b (ℓ_{t+1}^l) as

$$\ell_{t+1}^b = \frac{d_{t+1} + \phi_{t+1}(d_{t+1} + e_{t+1})}{e_{t+1}} \text{ and } \ell_{t+1}^l = \frac{d_{t+1}}{e_{t+1}}, \quad (16)$$

as well as the mass Δ_{t+1} of the banks that default as

$$\Delta_{t+1} = \mu(\tilde{p}_{t+1}). \quad (17)$$

Bank leverage depends on both the market funding ratio and the deposit to equity ratio. It is easy to see from Proposition 1 and relation (15) that these two ratios move in opposite directions. On the one hand the market funding ratio increases with ρ_{t+1} . On the other hand, the deposit to equity ratio decreases with r_{t+1}^d , as bank managers can afford raising more deposits when deposits become cheaper. Since ρ_{t+1} and r_{t+1}^d move together, it is not clear how leverage moves along the typical path. To investigate this, we set $\xi = 0.4\%$ so that in the simulations 1% of the banks default on average every year; with this calibration bank leverage is 16.13 on average (*i.e.* the bank capital ratio is around 6%), no bank defaults in normal times, and 12.5% of the banks default in crisis times (see Table 5). The evolution of

Table 5: Leverage and Default

Leverage			Default		
Borrowing Banks	Lending Banks	Aggregate	Normal	Crisis	Average
19.82	11.41	16.13	0	12.55%	0.90%

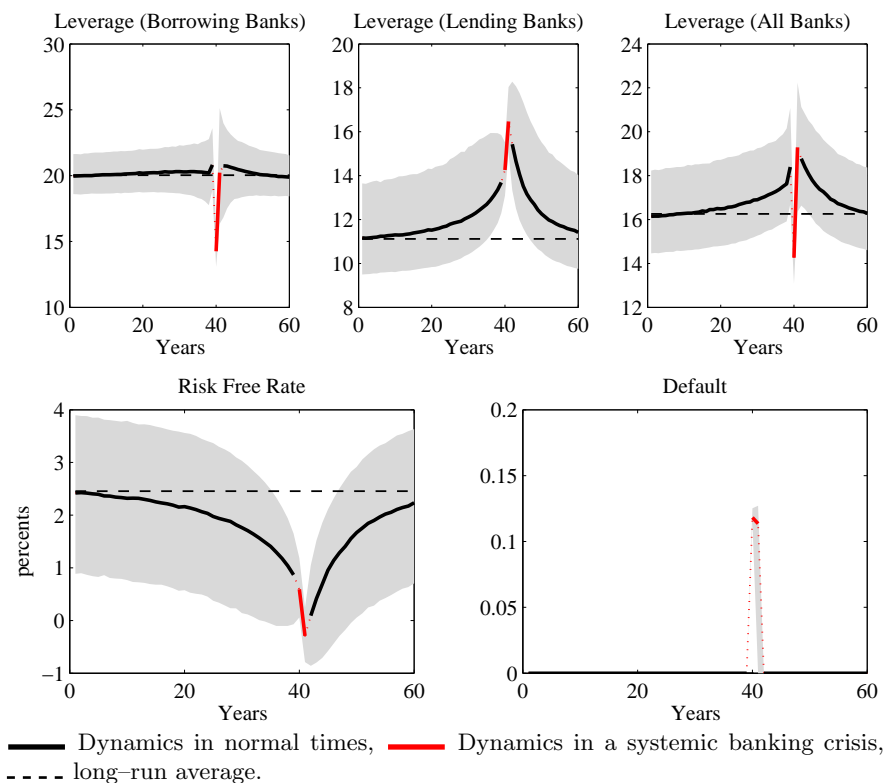
Note: All numbers are averages over a long simulation of 500,000 periods.

leverage, defaults, and the deposit rate along the typical path to crisis is reported in Figure 20.⁴⁵ In the run-up to the typical crisis, bank leverage overall increases for all banks, be

⁴⁴More precisely, bank managers choose d_{t+1}/e_{t+1} to maximize $\beta \mathbb{E}_t(r_{t+1}^e)$ subject to the identities $r_{t+1} a_{t+1} \equiv r_{t+1}^d d_{t+1} + r_{t+1}^e e_{t+1}$ and $a_{t+1} \equiv d_{t+1} + e_{t+1}$. If managers' incentives were aligned with the banks' shareholder, then they would instead maximize $\beta \mathbb{E}_t \left(\frac{u'(c_{t+1})}{u'(c_t)} r_{t+1}^e \right)$ and be indifferent between equity and deposits.

⁴⁵We focus on these variables because, the household being indifferent between equity and deposits, the deposit to equity ratio does not have any effect on the macroeconomic dynamics and all the other variables of the model follow the same path as in Figures 8–10.

Figure 20: Typical Path: Leverage and Default



they lenders or borrowers. For borrowers, in particular, the rise in the deposit to equity ratio offsets the diminution of the market funding ratio (see Figure 9); hence, even for those banks leverage is slightly pro-cyclical in normal times. During the crisis, in contrast, the evolution of leverage differs across banks, depending on whether or not they rely on market funding. While the leverage of borrowing banks (which finance themselves through the market) falls during the crisis, that of lending banks (which only finance themselves through equity and deposits) increases. The reason is that, in our model, bank equity is more volatile than retail deposits but less volatile than interbank loans. This result highlights the importance of the (un)stability of banks' various sources of funding in determining the dynamics of leverage. Finally, the lower right panel in Figure 20 shows that a significant mass of banks need to be recapitalized during systemic banking crises.

7 Conclusion

We develop a simple macroeconomic model that accounts for key features of systemic banking crises and helps understand why, historically, banking crises *(i)* occur in unusually good times and *(ii)* bring about the most severe recessions. The primary cause of systemic banking crises in the model is the accumulation of assets by the household in anticipation of future adverse shocks. When long enough, this asset accumulation process generates financial imbalances

(credit boom, disproportionately large bank balance sheets), which, in the best case, lower the resilience of the banking sector to shocks and, in the worst case, lead endogenously to a crisis. In this context, adverse random shocks themselves play a secondary role, as what matters the most is the conditions under which these shocks occur. This last result is of particular importance from a policy making perspective as it suggests that systemic banking crises are predictable. We show indeed that, although very stylized, our model is able to detect financial imbalances and predict future crises.

References

- Bernanke, B., M. Gertler and S. Gilchrist, 1999, "The Financial Accelerator in a Quantitative Business Cycle Framework," *Handbook of Macroeconomics*, Volume 1, Chapter 11, J.B. Taylor and M. Woodford eds, Elsevier.
- Bianchi, J., 2011, "Overborrowing and Systemic Externalities in the Business Cycle," *American Economic Review*, 101(7), 3400–3426.
- Bianchi, J. and E. G. Mendoza, 2010, "Overborrowing, Financial Crises and 'Macroprudential' Policy," NBER Working Paper No. 16091.
- Borio, C. and M. Drehmann, 2009, "Assesing the risk of banking crises: Revisited", *BIS Quarterly Review*, March, 29–46.
- Borio, C. and P. Lowe, 2002, "Assessing the risk of banking crises", *BIS Quarterly Review*, December, 43–54.
- Brunnermeier, M. and L. Pedersen, 2009, "Market Liquidity and Funding Liquidity," *Review of Financial Studies*, 22 (6): 2201–2238.
- Brunnermeier, M. and Y. Sannikov, 2012, "A Macroeconomic Model with a Financial Sector," mimeo, Princeton University.
- Burkart, M. and T. Ellingsen, 2004, "In-Kind Finance: A Theory of Trade Credit," *American Economic Review*, 94(3): 569–590.
- Christiano, L., M. Eichenbaum and C. Evans, 2005,, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113(1), 1–45.
- Claessens, S., A. Kose and M. Terrones, 2008, "How Do Business and Financial Cycles Interact?," IMF working Paper No. WP/11/88.
- Claessens, S., A. Kose and M. Terrones, 2011, "What Happens During Recessions, Crunches and Busts?," IMF working Paper No. WP/08/274.
- Cooper, R. and A. John, 1988, "Coordinating Coordination Failures in Keynesian Models," *Quarterly Journal of Economics*, (103)3, August, 441–463.
- Cooper, R., D. DeJong, R. Forsythe, and T. Ross, 1990, "Selection Criteria in Coordination Games: Some Experimental Results," *American Economic Review*, vol. 80, No. 1, March, 218–233.
- De Fiore, F. and H. Uhlig, 2012, "Corporate Debt Structure and the Financial Crisis," mimeo, European Central Bank.
- Fernald, J., 2009, "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity," mimeo, Federal Reserve Bank of San Francisco.
- Gertler, M. and P. Karadi, 2011, "A Model of Unconventional Monetary Policy', *Journal of Monetary Economics*, forthcoming.
- Gertler, M. and N. Kiyotaki, 2009, "Financial Intermediation and Credit Policy in Business Cycle Analysis," Mimeo, Princeton University.
- Gertler, M. and N. Kiyotaki, 2012, "Banking, Liquidity and Bank Runs in an Infinite Horizon Economy," Mimeo, Princeton University.

- Gorton, G., 1988, "Banking Panics and Business Cycles," *Oxford Economic Papers*, New Series, vol. 40, No. 4, 751–781.
- Gorton, G., 2010, Interview with "The Region," December, downloadable from <http://www.minneapolisfed.org/pubs/region/10-12/gorton.pdf>.
- Gorton, G., 2012, "Some Reflections on the Recent Financial Crisis," NBER Working Paper No 18397.
- Greenwood, J., Z. Hercowitz and G. Huffman, 1988, "Investment, Capacity Utilization, and the Real Business Cycle," *American Economic Review*, vol 78, No. 3, 402–417.
- Hart, O., 1995, "Firms, contracts, and Financial Structure," Oxford University Press.
- Hahm, J-H., H-S. Shin and K. Shin, 2011, "Non-Core Bank Liabilities and Financial Vulnerability," mimeo, Princeton University.
- He, Z. and A. Krishnamurthy, 2012, "A Macroeconomic Framework for Quantifying Systemic Risk," mimeo, University of Chicago.
- Holmström, B. and J. Tirole, 1997, "Financial Intermediation, Loanable Funds, and the Real Sector," *Quarterly Journal of Economics*, vol. CXII, No 3, August, 663–691.
- Jermann, U. and V. Quadrini, 2010, "Macroeconomic Effects of Financial Shocks", *American Economic Review*, forthcoming.
- Jordà, O., M. Schularick and A. Taylor, 2011a, "Financial Crises, Credit Booms and External Imbalances: 140 Years of Lessons," *IMF Economic Review*, Vol. 59, 340–378.
- Jordà, O., M. Schularick and A. Taylor, 2011b, "When Credit Bites Back: Leverage, Business Cycles, and Crises," NBER Working Paper No.17621.
- Judd, K. L. 1992, "Projection methods for solving aggregate growth models," *Journal of Economic Theory*, Vol. 58, 410-52.
- Korinek, A., 2010, "Systemic Risk-Taking: Accelerator Effects, Externalities, and Regulatory Responses," mimeo, University of Maryland.
- Laeven, L. and F. Valencia, 2008, "Systemic Banking Crises: A New Database," IMF Working Paper No 08224.
- Maddaloni, A. and J-L. Peydro, 2011, "Bank Risk-Taking, Securitization, Supervision, and Low Interest Rates: Evidence from the Euro Area and US Lending Standards," *Review of Financial Studies*, Vol. 24, 2121–65, 2011.
- Mehra, R. and E. Prescott, 1985, "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 145–161.
- Mendoza, E., 2010, "Sudden Stops, Financial Crises and Leverage," *American Economic Review*, Vol. 100(5), 1941–66.
- Mendoza, E. and K., Smith, 2005, "Quantitative Implications of a Debt-Deflation Theory of Sudden Stops and Asset Prices," *Journal of International Economics*, Vol. 70(1), 82–114.
- Minsky, H., 1977, "A Theory of Systemic Fragility," in E. Altman and A. Sametz, eds., *Financial Crises*.
- Ravn, M. and H. Uhlig, 2002, "On Adjusting the Hodrick-Prescott Filter for the Frequency of Observations," *Review of Economics and Statistics*, May, Vol. 84, No. 2, 371–376.

Reinhart, C. and K. Rogoff, 2009, “This Time is Different — Eight Centuries of Financial Folly,” Princeton University Press.

Schularick, M. and A. Taylor, 2011, “Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870-2008,” *American Economic Review*, Vol. 102, 1029–1061.

Shin, H-S., 2008, “Securitisation and Financial Stability,” mimeo, Princeton University.

Shin, H-S., 2010, “Risk and Liquidity,” Clarendon Lecture in Finance, Oxford University Press.

Shin, H-S., and K. Shin., 2011, “pro-cyclicality and Monetary Aggregates,” NBER Working Paper No 16836.

Stiglitz, J. and A. Weiss, 1981, “Credit Rationing in Markets with Imperfect Information,” *American Economic Review*, vol. 71, No. 3 (June), 393–410.

Tauchen, G. and R. Hussey, 1991, “Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models,” *Econometrica*, vol. 59, No. 2, 371–96.

Taylor, J., 2009, “The Financial Crisis and the Policy Responses: An Empirical Analysis of What Went Wrong,” *Critical Review*, vol. 21, 341–364.

Tirole, J., 2006, “The Theory of Corporate Finance,” Princeton University Press.

Uhlig, H., 2010, “A Model of Systemic Bank Run,” *Journal of Monetary Economics*, vol. 57, 78–96.

— APPENDIX —

A Proofs of Propositions

Proposition 1: The program of a borrowing bank writes

$$\begin{aligned} \max_{\phi_t} \quad & pR_t(1 + \phi_t) - \rho_t\phi_t \\ \text{s.t.} \quad & pR_t(1 + \phi_t) - \rho_t\phi_t \geq \rho_t \\ & \gamma(1 + \theta\phi_t) \leq \rho_t \end{aligned}$$

The participation constraint indicates that only banks with ability $p \geq \bar{p}_t \equiv \rho_t/R_t$ will borrow. We focus on this segment of the market. The problem simplifies to

$$\begin{aligned} \max_{\phi_t} \quad & pR_t(1 + \phi_t) - \rho_t\phi_t \\ \text{s.t.} \quad & \gamma(1 + \theta\phi_t) \leq \rho_t \end{aligned}$$

for $p \geq \bar{p}_t$. Let us denote by λ the Lagrange multiplier associated to the incentive constraint, the first order conditions are then

$$\begin{aligned} pR_t - \rho_t &= \gamma\theta\lambda \\ \lambda(\rho_t - \gamma(1 + \theta\phi_t)) &= 0 \end{aligned}$$

It follows that $\lambda > 0$ for all $p > \bar{p}_t$. The result follows. ■

Proposition 2: The market clearing condition (9a) in normal times together with the optimal demand for capital yields the normal times equilibrium corporate loan rate $R_t = f_k(a_t; z_t) + 1 - \delta$. The interbank market is at work if and only if this rate is above \bar{R} , *i.e.*, if and only if $a_t \leq \bar{a}_t \equiv f_k^{-1}(\bar{R} + \delta - 1; z_t)$. ■

Proposition 3: The proof proceeds in two steps.

- i) From (9) and Proposition 2, we show that the corporate loan rate increases during a credit crunch:

$$\begin{aligned} \lim_{a_t \nearrow \bar{a}_t} R_t + \delta - 1 &= \lim_{a_t \nearrow \bar{a}_t} f_k(a_t; z_t) = \lim_{a_t \searrow \bar{a}_t} f_k(a_t; z_t) \\ &< \lim_{a_t \searrow \bar{a}_t} f_k((1 - \mu(\gamma/R_t))a_t; z_t) = \lim_{a_t \searrow \bar{a}_t} R_t + \delta - 1 \end{aligned}$$

- ii) From (7), we know that $\lim_{a_t \nearrow \bar{a}_t} r_t/R_t = \int_{\bar{p}_t}^1 p \frac{d\mu(p)}{1 - \mu(\bar{p}_t)}$, which conditional expectation monotonically increases in \bar{p}_t . Using point (i) and the inequality $\rho_t \geq \gamma$, we have that $\lim_{a_t \nearrow \bar{a}_t} \bar{p}_t > \lim_{a_t \nearrow \bar{a}_t} \gamma/R_t > \lim_{a_t \searrow \bar{a}_t} \gamma/R_t$. Hence, $\lim_{a_t \nearrow \bar{a}_t} r_t/R_t > \lim_{a_t \searrow \bar{a}_t} \int_{\gamma/R_t}^1 p \frac{d\mu(p)}{1 - \mu(\gamma/R_t)}$. Finally, we know from (7) that

$$\lim_{a_t \searrow \bar{a}_t} r_t/R_t = \lim_{a_t \searrow \bar{a}_t} \frac{\gamma}{R_t} \mu\left(\frac{\gamma}{R_t}\right) + \int_{\gamma/R_t}^1 p d\mu(p) < \lim_{a_t \searrow \bar{a}_t} \int_{\gamma/R_t}^1 p \frac{d\mu(p)}{1 - \mu(\gamma/R_t)},$$

which implies $\lim_{a_t \nearrow \bar{a}_t} r_t/R_t > \lim_{a_t \searrow \bar{a}_t} r_t/R_t$. The result follows. ■

B Equations of the Model with endogenous Labour Supply

1. $y_t = z_t k_t^\alpha h_t^{1-\alpha} + (\gamma + \delta - 1)(a_t - k_t)$
2. $R_t = \alpha k_t^{\frac{-v(1-\alpha)}{v+\alpha}} z_t^{\frac{1+v}{v+\alpha}} \left(\frac{1-\alpha}{\vartheta} \right)^{\frac{1-\alpha}{v+\alpha}} + 1 - \delta$
3. $\left(c_t - \vartheta \frac{h_t^{1+v}}{1+v} \right)^{-\sigma} = \beta \mathbb{E}_t \left[\left(c_{t+1} - \vartheta \frac{h_{t+1}^{1+v}}{1+v} \right)^{-\sigma} r_{t+1} \right]$
4. $h_t = \left(\frac{(1-\alpha)z_t}{\vartheta} \right)^{\frac{1}{v+\alpha}} k_t^{\frac{\alpha}{v+\alpha}}$
5. $\bar{a}_t \equiv ((1-\alpha)/\vartheta)^{\frac{1}{v}} (\alpha/(\bar{R} + \delta - 1))^{\frac{v+\alpha}{v(1-\alpha)}} z_t^{\frac{1+v}{v(1-\alpha)}}$
6. $i_t = a_{t+1} - (1-\delta)a_t$

If $a_t \leq \bar{a}_t$ (normal times)

- 7a. $k_t = a_t$
- 8a. $\frac{r_t}{R_t} = \int_{\bar{p}_t}^1 p \frac{d\mu(p)}{1-\mu(\bar{p}_t)}$
- 9a. $\bar{p}_t = \frac{\rho_t}{R_t}$
- 10a. $R_t = \frac{\rho_t}{\mu^{-1} \left(\frac{\rho_t - \gamma}{\rho_t - (1-\theta)\gamma} \right)}$, with $\rho_t > \bar{\rho}$
- 11a. $y_t = c_t + i_t + (R_t - r_t)a_t$

If $a_t > \bar{a}_t$ (crisis times)

- 7b. $k_t = a_t - \mu(\gamma/R_t)a_t$
- 8b. $\frac{r_t}{R_t} = \frac{\gamma}{R_t} \mu(\gamma/R_t) + \int_{\gamma/R_t}^1 p d\mu(p)$
- 9b. $\bar{p}_t = \gamma/R_t$
- 10b. $\rho_t = \gamma$
- 11b. $y_t = c_t + i_t + (R_t - r_t)a_t - (R_t - \gamma)(a_t - k_t)$

Few comments are in order here. (i) In autarky assets $a_t - k_t$ are stored for return γ . Since capital depreciates at rate δ , the value added of storage is $\gamma + \delta - 1$, as reflected in equation 1. (Remember that $\gamma + \delta - 1 > 0$.) (ii) The good market clearing conditions 11a and 11b are derived from Walras' law and the agents' budget constraints. Summing up the household and the firm's budget constraints (1) and (3), one gets: $c_t + a_{t+1} = z_t k_t^\alpha h_t^{1-\alpha} + (1-\delta)k_t + r_t a_t - R_t k_t$, which simplifies to $c_t + i_t = y_t - (\gamma + \delta - 1)(a_t - k_t) - (1-\delta)(k_t - a_t) + r_t a_t - R_t k_t \iff y_t = c_t + i_t + (R_t - r_t)a_t - (R_t - \gamma)(a_t - k_t)$. (iii) The last two terms of this accounting identity correspond to the aggregate cost of financial frictions. In autarky, for instance, this cost is, by definition, equal to $\int_{\gamma/R_t}^1 (1-p) R_t a_t d\mu(p)$, which indeed simplifies to (using 7b and 8b): $(1 - \mu(\gamma/R_t)) R_t a_t - a_t \int_{\gamma/R_t}^1 p R_t d\mu(p) = R_t k_t - r_t a_t + \gamma(k_t - a_t) = (R_t - r_t)a_t - (R_t - \gamma)(a_t - k_t)$.

C Additional Figures

C.1 Response to a Large Negative Shock

The economy is started from the average steady state and is subjected to a 10% drop in total factor productivity in the initial period. As in other DSGE models featuring financial frictions, the model is able to generate a financial crisis from a large negative shock, provided that the shock is large enough. In our case, a 7% drop from average steady state is sufficient to generate crises.

Figure C.1: Response to a large drop in TFP (I)




Figure C.2: Response to a large drop in TFP (II)






Figure C.3: Response to a large drop in TFP (III)




C.2 Evolution of US Interest Rates and Spreads

Figure C.4: Evolution of Various Corporate Loan Spreads

(a) Spread: Corporate loan rates - Federal Fund Rate



(b) Underlying Real Corporate Loan Rates



The effective federal fund rate is from the Federal Reserve Economic Database (<http://research.stlouisfed.org/fred2/series/FEDFUNDS?cid=118>). The corporate loan rates are from the Federal reserve (<http://www.federalreserve.gov/releases/e2/e2chart.htm>). Nominal rates are deflated using the annualized quarter on quarter rate of growth of the GDP implicit deflator (<http://research.stlouisfed.org/fred2/series/GDPDEF?cid=21>).