



Credit booms, financial crises, and macroprudential policy

Mark Gertler^a, Nobuhiro Kiyotaki^b, Andrea Prestipino^{c,*}



^a New York University, United States of America

^b Princeton University, United States of America

^c Federal Reserve Board, United States of America

ARTICLE INFO

Article history:

Received 15 April 2020

Available online 20 June 2020

JEL classification:

E23

E32

E44

G01

G21

G33

Keywords:

Financial crises

Credit booms

Macroprudential policy

Bank runs

Countercyclical capital buffers

ABSTRACT

We develop a model of banking panics which is consistent with two important features of the data: First, banking crises are usually preceded by credit booms. Second, credit booms often do not result in crises. That is, there are “bad booms” as well as “good booms” in the language of Gorton and Ordóñez (2019). We then consider how the optimal macroprudential policy weighs the benefits of preventing a crisis against the costs of stopping a good boom. We show that countercyclical capital buffers are a critical feature of a successful macroprudential policy.

© 2020 Elsevier Inc. All rights reserved.

1. Introduction

We develop a model of banking panics which is consistent with two important features of the data: First, banking crises are usually preceded by credit booms. Second, credit booms often do not result in crises. That is, there are “bad booms” as well as “good booms” in the language of Gorton and Ordóñez (2019). We then use the model to study macroprudential policy.

We first describe the main facts we think a model of financial crises should capture. Fig. 1 portrays the link between credit growth and financial crises, using data from Krishnamurthy and Muir (2017). The evidence is based on a panel of annual data of industrialized countries, ranging from 1869 to 2018. The authors use the narrative based classification in Jordà et al. (2011) to determine periods in which a country experienced a financial crisis. The figure then plots the average behavior of output, credit growth and credit spreads, around the time a crisis occurs. In each of the three panels, the crisis occurs at year zero. The upper-right panel shows that credit growth on average steadily increases prior to the crisis before declining afterward, as a number of authors have recently noted, e.g. Schularick and Taylor (2012). The bottom panel shows

[☆] We would like to thank the participants of various seminars and conferences, particularly Julianne Begenau, Jonathan Heathcote, Urban Jermann, Vincenzo Quadrini and Alejandro Van der Ghoote for thoughtful comments. The Financial Support of the National Science Foundation, grant SES-1917916, is gratefully acknowledged.

* Corresponding author.

E-mail address: andrea.prestipino@frb.gov (A. Prestipino).

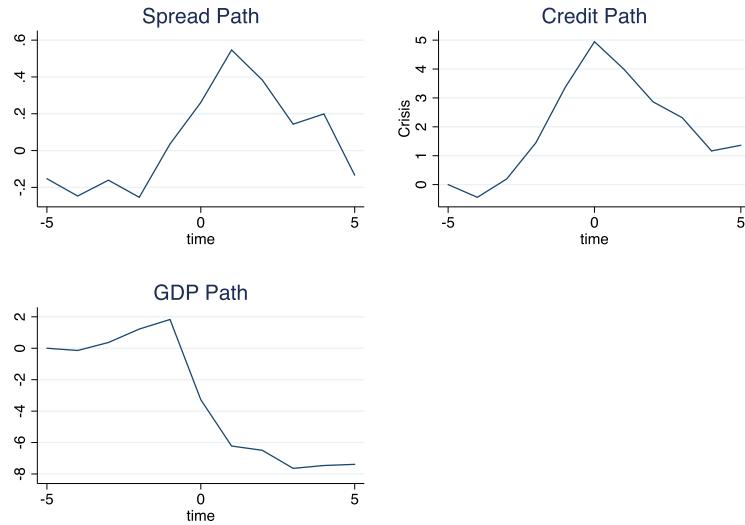


Fig. 1. This figure is from Krishnamurthy and Muir (2017). It plots the behavior of credit spreads, GDP, and the quantity of credit around a financial crisis with the crisis beginning at date 0. GDP and credit are expressed in deviation from (country specific) trend. Spreads are normalized by dividing by the unconditional mean.

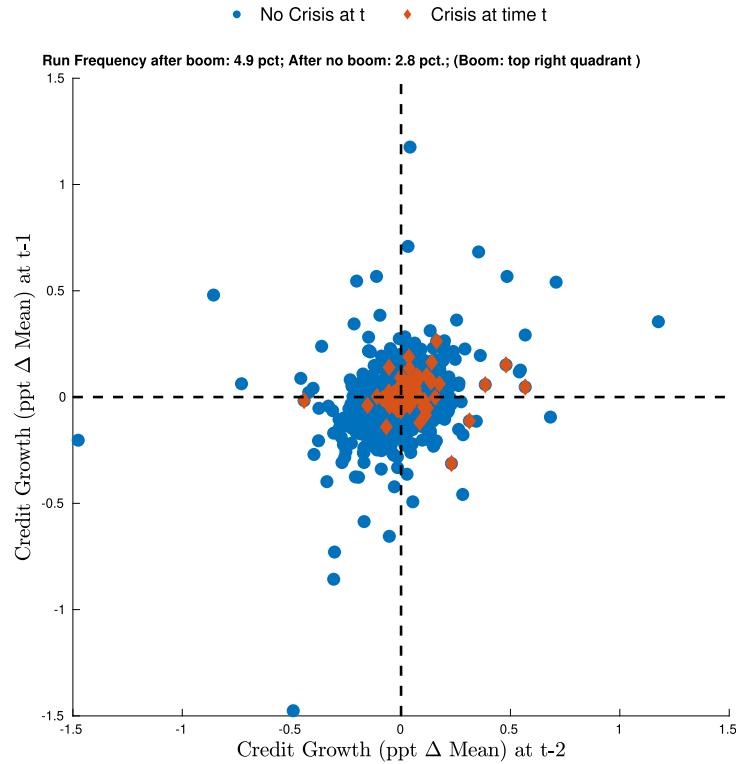


Fig. 2. Credit Booms and Financial Crises. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

that prior to a crisis, GDP growth on average increases relative to trend by roughly two percent, but when the crisis hits it experiences a sharp and persistent decline of nearly eight percent. Finally, as support for the notion that the output contractions reflect financial crises, credit spreads increase on average prior to and during the crisis, before eventually going back to a normal level, as shown in the upper-left panel.

Fig. 2, however, makes clear that high credit growth does not always lead to a crisis, nor is it necessary for a crisis to arise. The data in the figure plots annual demeaned credit growth in a country lagged two years (the horizontal axis) versus

one year (the vertical axis).¹ The red dots are episodes where a country experienced a financial crisis while the blue are instances where a crisis did not occur. If we think of a credit boom as a period in which credit growth is above average for two consecutive years, then the upper right quadrant reflects periods preceded by credit booms. Accordingly, crisis episodes happening after credit booms are all the red dots in the upper right hand quadrant in the figure, while the blue dots in the upper right quadrant are episodes in which a credit boom did not result in a crisis. As the figure shows, more often than not, a credit boom does not result in a financial crisis. Conditional on a credit boom, the probability of a crisis is just 4.9 percent. It is true, however, that a credit boom makes a crisis more likely: conditional on no credit boom, the probability of a crisis drops to 2.8 percent.

Our goal in this paper is to first develop a macroeconomic framework with banking panics that is consistent with the evidence in Figs. 1 and 2, and then to use the model to study regulatory policy. The framework we develop is based on Gertler et al. (2020b), henceforth GKP (2020), which is a standard New Keynesian macro model modified to include banks and banking panics that disrupt real activity. Within that framework, we capture both credit booms preceding crises and the banking collapse and disruption of real activity that follows. In the spirit of Geanakoplos (2010) and Bordalo et al. (2018), the source of the boom is optimistic beliefs by financial intermediaries (or banks in short) about future returns to capital that are eventually disappointed.² This leads to a buildup of bank credit that is funded by an increase in bank leverage, mostly in the form of short term debt. High levels of debt, in turn, make the system vulnerable to a run by increasing the exposure of banks to negative returns on their assets, so that even small negative shocks can trigger system wide runs that result in deep contractions in economic activity. In this latter regard, the model captures the highly nonlinear dimension of financial crises. We use global methods to solve the model numerically in order to characterize these nonlinearities.

There are several differences from our earlier work. First, while in our earlier paper we used a canonical New Keynesian framework with capital accumulation and focused our study on the Great Recession, here we consider a simple endowment economy but allow for recurrent credit booms that may or may not result in banking crises. This allows us to capture the statistical relationship between credit booms and financial crises described above. The presence of good and bad credit booms sets the stage for our study of macroprudential regulation. By restricting financial intermediation, macroprudential policies can prevent the large credit booms that are the root cause of financial crises. However, since the regulator can't tell apart bad credit booms from good ones, attempts at preventing crises will often end up stifling good booms. By matching the relative frequency of good and bad credit booms in the data, our framework allows us to study quantitatively how the optimal policy weighs the benefits of preventing crises against the costs of stopping good booms. We also analyze the features of optimal regulation and show, for example, that countercyclical capital buffers are a critical feature of a successfully designed macroprudential policy.

One final important modeling difference from our earlier work is that we allow for equity injections into the banking sector. In our earlier work we assumed that bank capital was only accumulated via retained earnings. What this implies is that to meet equity capital requirements, the only margin of adjustment is for banks to reduce assets. Allowing for new equity injections introduces a second margin of adjustment. We assume however that at the margin, equity injections are costly. If they were costless, equity finance would become the sole source of funding for banks, eliminating the possibility of runs or any other type of banking instability.³ However, there is a very large literature in finance that argues that equity finance is costly for banks and stresses the important role of debt finance in contexts where agency problems affect the relationship between bank managers and outside investors.⁴ Accordingly, in this paper we assume that equity finance comes at a cost. While we do not explicitly model the frictions that underpin this cost, we discipline its impact on banks funding choices by matching the observed average leverage ratio and equity issuance rate of financial firms. In particular, while the data shows that equity issuance rose during the financial crisis, it remained small as a fraction of total equity. In particular, Fig. 4 shows that the average annual equity issuance of financial firms was one percent of trend equity between 1985 and 2007, and peaked at around 2.4% during 2008–2010.⁵ Accordingly as a check that our parametrization is reasonable, we show that the model implied equity injections after a banking panic are in line with that observed during the recent financial crisis.

Our paper contributes to a large literature that studies the role of financial intermediaries in macroeconomic fluctuations. Much of this literature builds on the conventional financial accelerator model of Bernanke et al. (1998), and Kiyotaki and Moore (1997). While the traditional models were developed to study how procyclical movement in nonfinancial borrowers balance sheets work to amplify and propagate macroeconomic fluctuations, Gertler and Karadi (2011) and Gertler and Kiyotaki (2010) showed how the basic mechanism could be applied to study financial firms as well. One limitation of the original models was that, by studying the local behavior of the economy around a non-stochastic steady state, they could not

¹ Credit growth data is from Jordà et al. (2011). To demean the data we compute for each country separate means of credit growth for the pre-war period and post-war periods.

² Though we use a different belief mechanism from Bordalo et al. (2018), we follow these authors in our earlier work (Gertler et al. (2020b)) by showing that investor forecast errors from our mechanism during the recent boom and bust are consistent with the evidence. See also Boz and Mendoza (2014) and Boissay et al. (2016) for other models that try to capture the boom-bust cycle in credit associated with financial crises.

³ With one hundred percent equity financing, the banks creditors absorb the risk, making the banking system perfectly safe.

⁴ See, for example, Calomiris and Kahn (1991) and Diamond and Rajan (2001).

⁵ This data does not include the government purchase of subordinate debts and preferred stocks through the Troubled Asset Relief Program during the crisis.

capture the non linear dimension of financial crises. To address this limitation, a series of papers have tried to capture the nonlinear dimension of financial crises by exploiting occasionally binding financial constraints, e.g. Mendoza (2010), He and Krishnamurthy (2019) and Brunnermeier and Sannikov (2014). While we also allow for occasionally binding constraints, the main source of non-linearity in our paper is the occurrence of a bank run. As in our earlier work, e.g. Gertler and Kiyotaki (2015), Gertler et al. (2016) and Gertler et al. (2020b), we model bank runs as rollover panics following the Calvo (1988) and Cole and Kehoe (2000) models of sovereign debt crises.⁶ The existence of a bank run equilibrium depends on the health of banks balance sheets. When banks balance sheets are weak, fears of a bank run can become self-fulfilling even in the absence of any negative fundamental shock.⁷ Bank runs, in turn, force banks to liquidate assets at firesale prices, causing a sudden collapse in bank equity, and a deep and prolonged economic contraction.

We also contribute to the growing literature that studies the role of macroprudential regulation in preventing crises. Beginning with Lorenzoni (2008), a lengthy literature has emerged that examines bank regulation in a macroeconomic setting. The main conceptual motive for regulation in this literature is the presence of a pecuniary externality, where individual banks fail to take account of the impact of their risk exposure on the dynamics of asset prices. This work has been both qualitative, e.g. Angeloni and Faia (2013), Jeanne and Korinek (2019), Chari and Kehoe (2016), and quantitative, e.g., Bianchi and Mendoza (2018), Benigno et al. (2013), and Begenau and Landvoigt (2018)). We differ in two main ways. First, since we allow for endogenous nonlinear financial panics that lead to real economic disasters, the main gain from macroprudential policy in our model is reducing the likelihood of one of these disasters. In our view, avoiding such disasters is the primary objective of macroprudential policy in practice. More formally, the externality in our model is that banks fail to account for the impact of their individual risk exposure on the likelihood of a panic. In addition, by modeling credit booms as well as busts and making the distinction between good and bad credit booms, we are able to characterize the tradeoff between reducing the likelihood of banking crises versus stifling good credit booms.

Section 2 develops the baseline model of banking and banking panics. Section 3 introduces beliefs and then numerically illustrates how the model can generate credit booms and busts, including good booms as well as bad booms. Section 4 then analyzes macroprudential policy. The Appendix provides a detailed development of the model and the nonlinear computational algorithm for solving it.

2. Baseline model

The framework is an endowment economy with two goods, consumption C_t and capital K_t . The latter is used to produce consumption goods. We suppose capital is fixed in supply and normalize the total to be unity. The financing of capital takes on one of two forms. First, banks may intermediate the quantity K_t^b . By “intermediate”, we mean that banks issue deposits to households and then use the funds together with their own equity to acquire capital. Second, households directly hold the quantity K_t^h , implying that in the aggregate

$$K_t^b + K_t^h = 1. \quad (1)$$

The division of capital financing between intermediated finance versus direct holding is endogenous and determined in the general equilibrium.

We suppose that households are less efficient in evaluating and monitoring capital projects than banks. We capture this notion by assuming that household direct finance entails a management cost $\frac{\alpha}{2} (K_t^h)^2$, which is increasing and convex in the quantity of directly held capital, K_t^h . The increasing marginal managerial cost is meant to capture that a household has limited capacity to manage capital.⁸

In addition to directly holding capital and supplying deposits to banks, we suppose that households are the owners of banks. (Think of households as owning banks that are different from the ones in which they hold deposits.) Accordingly households are the recipients of bank dividend payouts and decide how much equity to inject into banks. In particular, we assume that households can costlessly inject an amount $\bar{\xi}$ of equity in the banking system, but face a convex cost $\frac{\alpha \xi}{2\bar{\xi}} (\xi_t^N - \bar{\xi})^2$ when equity injections ξ_t^N exceed $\bar{\xi}$. We introduce costly equity injections to capture in a simple reduced form way the frictions involved for banks in raising equity.⁹ We then pick the parameters of the cost function to match the empirical properties of equity injections in the banking sector, as shown in Fig. 4.

⁶ This is in contrast with the traditional literature on banking panics originating from Diamond and Dybvig (1983), in which sequential service constraints were key in order to generate bank runs. Our modeling of bank runs as rollover crises seems to capture well the bank runs that were at the heart of the recent financial crisis. See Bernanke et al. (2010), Bernanke (2018) and Gorton (2010).

⁷ Some recent examples where self-fulfilling financial crises can emerge depending on the state of the economy include Benhabib and Wang (2013), Bocola and Lorenzoni (2017), Farhi and Maggi (2018) and Perri and Quadrini (2018). For further attempts to incorporate bank runs in macro models, see Angeloni and Faia (2013), Cooper and Ross (1998), Martin et al. (2014), Robatto (2019) and Uhlig (2010) for example. Adrian et al. (2019) empirically study the possibility of multiple equilibria in the dynamics of financial conditions and GDP in the US.

⁸ We take the quadratic form for convenience since it implies that the marginal managerial cost is linear.

⁹ Jermann and Quadrini (2012) provide a related way to model costs of equity infusion: They suppose the firm faces a quadratic cost of deviating from a positive dividend target. Equity injections are then costly since they involve negative dividend payouts. We model the costs on the household side because it simplifies the algebra within our framework. Another approach is Gertler et al. (2012) in which, everything else equal, agency frictions increase as banks shift funding from short term debt to equity.

Let Z_t be a shock to the flow return on capital and W (for labor income) an endowment of consumption goods that the household receives each period. Since the total supply of capital is fixed at unity, the aggregate resource constraint is given by

$$C_t = Y_t = Z_t + W - \frac{\alpha}{2}(K_t^h)^2 - \frac{\alpha\xi}{2\xi}(\xi_t^N - \bar{\xi})^2, \quad (2)$$

where Z_t obeys the following first order process

$$Z_{t+1} = 1 - \rho + \rho Z_t + \varepsilon_{t+1}. \quad (3)$$

As we will make precise below, we suppose that banks face constraints in borrowing funds from depositors. Bank equity helps reduce these frictions, which accounts for why households will want to inject equity, even if it is costly at the margin, i.e. $\xi_t^N > \bar{\xi}$ in equilibrium. The costs of equity injections, though, work to limit the amount of equity in the banking system. This limit on bank equity in turn helps account for why banks do not intermediate the entire capital stock in equilibrium and instead households hold a fraction, even though direct household finance entails costs, i.e. $K_t^h > 0$.

Note that the model implies that net output declines as the share of bank financing of capital falls because of the direct managerial costs $\frac{\alpha}{2}(K_t^h)^2$. Thus the model implies in a reduced form way that disintermediation leads to a drop in output.¹⁰ A secondary factor contributing to the costs of disintermediation involves the costs of equity issuance. As the share of banking financing of capital declines due to a tightening of credit constraints, the marginal value of bank equity increases, causing equity injections and hence the costs of equity injections to rise. It will turn out, however, that these costs are quite small relative to the household managerial costs for direct finance. Accordingly, it is the managerial costs that largely account for the negative effect of disintermediation on net output.

Finally, it is instructive to compare the rates of the return on bank intermediated capital, R_{t+1}^b , versus that on directly held capital R_{t+1}^h . Let Q_t denote the relative price of capital. Then

$$R_{t+1}^b = \frac{Z_{t+1} + Q_{t+1}}{Q_t},$$

$$R_{t+1}^h = \frac{Z_{t+1} + Q_{t+1}}{Q_t + \alpha K_t^h}.$$

Due to the managerial cost, R_{t+1}^h is less than R_{t+1}^b . Further, this gap widens as households directly hold a larger share of the capital stock, since the marginal managerial cost, αK_t^h , is increasing in K_t^h . The net effect is that in situations where banks shed assets, Q_t must drop sufficiently in order for households to absorb them. In the case of a fire sale, which will arise in the event of a run, Q_t must drop sharply.

2.1. Households

There is a representative household that contains a measure unity of family members. The fraction f of the members are bankers and the fraction $1 - f$ are workers. Each worker receives an endowment (effectively labor income). Each banker manages a financial intermediary and pays dividends to the household. Within the household there is complete consumption insurance. An advantage of this setup is that we introduce financial intermediation but at the same time avoid the complication of heterogeneous households.

The household chooses consumption and saving, as well as the allocation of its portfolio between bank deposits and direct capital holdings. In addition, it can inject new equity into the banking system by providing startup equity to new banks.¹¹

Further, there is turnover: Each period some bankers exit the business and become workers. When they exit, they pay as dividends any residual retained earnings back to the households. An equal number of workers become new bankers. We introduce turnover in banking to give each banker a finite expected horizon. This ensures that banks use leverage to finance assets in the stationary equilibrium. Otherwise, with an infinite horizon, they could over time retain sufficient earnings to purely self finance. In particular, with i.i.d. probability $1 - \sigma$, a banker exits in the subsequent period and with probability σ the banker survives and continues to operate, making a banker's expected horizon equal to $\frac{1}{1-\sigma}$ periods. Each period the exiting bankers are replaced by $(1 - \sigma)f$ workers turned bankers, keeping the total populations of bankers and workers constant.

¹⁰ Gertler et al. (2020b) provide a more realistic description of how a banking collapse leads to an output collapse. In their framework the banking panic leads to a sharp contraction in investment which reduces aggregate demand and output due to nominal rigidities.

¹¹ We assume for simplicity that all equity injections by households are received by new bankers. Because, as we show later, what matters is the equity in the banking system as a whole rather than the distribution across banks, our results are robust to an alternative specification in which equity injections are received by all active banks in an amount that is proportional to their retained earnings.

Each new banker receives a fixed startup transfer from the household, $\frac{\bar{\xi}}{(1-\sigma)I_t}$. Moreover, households can inject additional equity, I_t , into new banks. We assume that these injections entail a quadratic resource cost. In particular, letting $\xi_t^N = \bar{\xi} + I_t$ be the total amount of equity transferred to new bankers, we assume resource costs associated with ξ_t^N of the form

$$f_{\xi}(\xi_t^N) = \begin{cases} \frac{\alpha_{\xi}}{2\bar{\xi}}(\xi_t^N - \bar{\xi})^2 & \xi_t^N \geq \bar{\xi} \\ 0 & \text{otherwise} \end{cases}.$$

As we describe below, the presence of financial market frictions implies that bankers are not able to arbitrage away excess returns on their investment, so that, in equilibrium, the rate of return on their assets is above the interest rate they pay on deposits. Therefore, bankers will always prefer to keep accumulating net worth and only payout dividends when they exit.

Accordingly, for a household with bank equity X_{t-1}^N at $t-1$, the total dividend payouts from bank equity the household receives at t , Π_t , and the total equity it is left with at t , X_t^N , are given by, respectively:

$$\underbrace{\Pi_t = (1-\sigma)X_{t-1}^N R_t^N}_{\text{Dividend at } t} \quad (4)$$

$$\begin{array}{c} X_{t-1}^N \\ \nearrow \\ \underbrace{X_t^N = \sigma X_{t-1}^N R_t^N + \xi_t^N}_{\text{Equity at } t} \\ \searrow \end{array} \quad (5)$$

where R_t^N is the growth rate of bank net worth from $t-1$ to t . Time t dividends are given by the total net worth of the fraction $(1-\sigma)$ of bankers that exit and return to the household. Total bank equity at time t is the sum of total net worth of surviving bankers, $\sigma R_t^N X_{t-1}^N$, and injections into new banks ξ_t^N .

We can now describe the household optimization problem. Let C_t denote consumption, D_t bank deposits and R_t the return on deposits. Then the household chooses $\{C_t, D_t, K_t^h, X_t^N, \xi_t^N\}$ to maximize

$$U_t = E_t \sum_{i=0}^{\infty} \beta^i \ln C_{t+i} \quad (6)$$

subject to

$$C_t + D_t + Q_t K_t^h + \frac{\alpha}{2} (K_t^h)^2 + \xi_t^N + f_{\xi}(\xi_t^N) = W + R_t D_{t-1} + (Z_t + Q_t) K_{t-1}^h + \Pi_t, \quad (7)$$

where dividends, Π_t , and the evolution of equity holdings are given by (4) and (5). The left hand side (LHS) of the budget constraint in equation (7) is the use of funds - consumption and saving in deposit, capital and equity, including the costs of direct finance and equity injection. The right hand side (RHS) is the source of funds - wages, returns on deposit and capital and dividend distribution from retired bankers.

Let $\Lambda_{t,t+1} \equiv \beta \frac{C_t}{C_{t+1}}$ denote the household stochastic discount factor. Then the household's first order conditions for deposits and direct capital holdings are given by:

$$E_t(\Lambda_{t,t+1} R_{t+1}) = 1, \quad (8)$$

$$E_t\left(\Lambda_{t,t+1} \frac{Z_{t+1} + Q_{t+1}}{Q_t + \alpha K_t^h}\right) = 1. \quad (9)$$

Note that the return on deposits R_{t+1} may be risky due to the possibility of default. (See equation (15) below.)

Let ψ_t^h be the multiplier on the equity evolution equation in (5) and $\psi_t^h = \frac{\bar{\psi}_t^h}{U'(C_t)}$ be the multiplier in terms of consumption goods. Then the first order conditions with respect to equity holdings X_t^N and equity injections ξ_t^N are given by, respectively:

$$\psi_t^h = E_t \left[\Lambda_{t,t+1} (1-\sigma + \sigma \psi_{t+1}^h) R_{t+1}^N \right], \quad (10)$$

$$1 + f'_{\xi}(\xi_t^N) \geq \psi_t^h \text{ and } \xi_t^N \geq \bar{\xi}. \quad (11)$$

Note ψ_t^h is the shadow value to the household of having another unit of bank equity in its portfolio. According to equation (10) this shadow value equals the expected discounted return to bank capital, taking into account that the bankers exit

with probability $1 - \sigma$ and continue with probability σ . Equation (11) states that the household adds bank equity to the point where the marginal benefit (the right side) equals the marginal cost of new injections (the left). The first equation in (11) holds with equality and the second with strict inequality, if the shadow value of bank equity to the household exceeds unity.

2.2. Bankers

Bankers fund assets $Q_t k_t^b$ with net worth (or retained earnings) n_t and deposits d_t :

$$Q_t k_t^b = d_t + n_t. \quad (12)$$

Retained earnings n_t are given by the return on bank investments minus debt funding costs:

$$n_t = R_t^b Q_{t-1} k_{t-1}^b - R_t d_{t-1}. \quad (13)$$

In the event of default (either due to a run or insolvency), retained earnings go to zero.

As we discussed earlier, the banker operates on behalf of the household and faces an exit probability $1 - \sigma$. The banker's objective is to maximize the expected present discounted value of dividend payouts to the household. Given the banker faces financial market frictions, which we will introduce shortly, it turns out to be optimal for the banker to delay dividend payouts until exit.¹² Accordingly we can express the banker's objective as:

$$V_t = E_t \{ \Lambda_{t,t+1} [(1 - \sigma)n_{t+1} + \sigma V_{t+1}] \}. \quad (14)$$

There are two additional features critical to generating banking panics. First, deposits are short term and contingent only on the possibility of default. Let \bar{R}_t be the promised deposit rate, p_t the default probability. Then the return on deposits is given by:

$$R_{t+1} = \begin{cases} \bar{R}_t, & \text{with probability } 1 - p_t \\ x_{t+1} \bar{R}_t, & \text{with probability } p_t \end{cases}, \quad (15)$$

where x_{t+1} is the depositor recovery rate at $t + 1$, which equals the ratio of bank assets to its promised deposit obligations as

$$x_{t+1} = \frac{R_{t+1}^b Q_t k_t^b}{\bar{R}_t d_t}. \quad (16)$$

Notice that the recovery rate does not depend upon the place on the queue because we did not impose the sequential service constraint. Notice also that default is possible only if $x_{t+1} < 1$. This consideration will turn out to be important for determining whether an equilibrium with a banking panic can exist, as we discuss shortly.

Second, we introduce an agency problem between a bank and its depositors that limits the bank's ability to obtain funds. Absent such a limit, a financial panic cannot emerge: A panic withdrawal would simply lead the bank to go to the credit market to offset the deposit loss. In particular, we introduce the following moral hazard problem: After the banker borrows funds at t , it may divert the fraction θ of assets for personal use (specifically to pay as dividends to its owner/family). If the bank does not honor its debt, creditors can recover the residual funds and shut the bank down. Recognizing this temptation, rational depositors require that the following incentive constraint be satisfied:

$$\theta Q_t k_t^b \leq V_t. \quad (17)$$

The left side of (17) is the banker's gain from diverting funds while the right hand side is the continuation value V_t from operating honestly.

The bank's decision problem is to choose assets k_t^b , deposits d_t and future net worth n_{t+1} , to maximize the objective (14), subject to the constraints of (12), (17) and the evolution of net worth at $t + 1$ implied by (13). We describe the solution informally and defer a detailed derivation to the Appendix.

From the bank balance sheet condition (12) and the evolution of the net worth (13), the rate of return on bank net worth is given by

$$\frac{n_{t+1}}{n_t} = R_{t+1}^N = \left(R_{t+1}^b - R_{t+1} \right) \frac{Q_t k_t^b}{n_t} + R_{t+1}. \quad (18)$$

The first term in the RHS shows how the bank can use leverage, $\frac{Q_t k_t^b}{n_t} > 1$, to amplify its return on net worth whenever the return on its assets exceed the deposit rate, i.e. when excess returns $(R_{t+1}^b - R_{t+1})$ are positive. The second term is the rate

¹² See the Appendix.

of return on deposits (which the bank can save by having an additional unit of net worth). Note that if the bank had no leverage (i.e., $\frac{Q_t k^b}{n_t} = 1$), then the rate of return on bank equity would simply equal the rate of return on capital held by banks, R_{t+1}^b . In the case where (appropriately discounted) excess returns are positive, banks will want to boost their profit margins by increasing borrowing.

The incentive constraint (17), on other hand, limits the ability of banks to increase their borrowing. To see this, let ψ_t^b be the bank franchise value per unit of its net worth - Tobin's Q ratio, or the shadow value of bank net worth:

$$\psi_t^b = \frac{V_t}{n_t}. \quad (19)$$

Using (18) and (19) to substitute for n_{t+1} , V_t and V_{t+1} in (14), we get an expression for the shadow value of equity as the discounted expected return on bank equity:

$$\psi_t^b = E_t \left[\Lambda_{t,t+1} (1 - \sigma + \sigma \psi_{t+1}^b) R_{t+1}^N \right]. \quad (20)$$

Then, combining equations (17) and (19) yields the following endogenous capital requirement, κ_t :

$$\kappa_t \equiv \frac{n_t}{Q_t k_t^b} \geq \frac{\theta}{\psi_t^b}. \quad (21)$$

As we discuss in the appendix the shadow value of net worth is independent of bank specific characteristics and it exceeds unity, i.e. $\psi_t^b > 1$.¹³ Therefore, given that θ is strictly between zero and one, the capital requirement κ_t lies strictly between zero and unity as well.

Notice that the shadow value of net worth ψ_t^b is increasing in risk-adjusted expected excess returns, because the return on bank equity R_{t+1}^N , in the RHS of (20), is increasing in excess returns of bank assets over deposit $R_{t+1}^b - R_{t+1}$.¹⁴ When the marginal risk-adjusted expected excess returns are positive, the bank would like to increase its leverage multiple, $\frac{Q_t k^b}{n_t}$, as much as possible. Equation (21) however implies that banks will be limited in the amount of leverage they can take. According to (21), the required bank equity - asset ratio is increasing in the diversion rate θ and decreasing in the shadow value of net worth ψ_t^b . A rise in θ increases the bank's temptation to divert assets, everything else equal. To satisfy the incentive constraint the bank must reduce deposits, leading it to scale back assets relative to net worth. Conversely, an increase in ψ_t^b raises the franchise value V_t reducing the bank's temptation to divert. As a result, the bank can satisfy the incentive constraint with a smaller capital asset ratio.

There are three implications of (21) that are relevant to the analysis of runs that follows. First, the bank cannot operate with $n_t \leq 0$. A bank with zero or negative net worth can never satisfy the incentive constraint: It will always want to divert the proceeds from any deposits it issues. It turns out that the inability of the bank to operate with zero net worth is critical for the existence of a bank run equilibrium, as we describe shortly.

Second, the required capital ratio $\frac{\theta}{\psi_t^b}$ varies inversely with ψ_t^b . Thus, the endogenous capital requirement is relaxed in periods when ψ_t^b rises, allowing banks to operate with lower capital ratios in these instances. Since ψ_t^b depends positively on $E_t (R_{t+1}^b - R_{t+1})$, periods in which expected excess returns on bank assets rise are also periods in which banks capital ratios are low.¹⁵ The significance for our purposes, is that the probability of a run equilibrium increases when banks capital ratios are low. In our experiments below we explore two possible ways that expected excess returns increase causing bank capital ratios to decline and the bank run probability to rise: negative fundamental shocks and positive belief shocks.

Finally, since κ_t does not depend on individual bank's characteristics, banks portfolio choices are homogeneous in bank net worth and the aggregate demand for capital by banks is simply

$$Q_t K_t^b = \frac{1}{\kappa_t} N_t, \quad (22)$$

¹³ Notice that in principle the return on bank net worth R_{t+1}^N in equation (20) could depend on bank specific portfolio choices. The appendix shows that in practice all banks choose the same leverage and hence have the same rate of return on net worth.

¹⁴ The risk adjusted expected excess return is defined as

$$E_t [\Lambda_{t,t+1} (1 - \sigma + \sigma \psi_{t+1}^b) (R_{t+1}^b - R_{t+1})].$$

¹⁵ In the data, net worth of our model corresponds to the mark-to-market difference between assets and liabilities of the bank balance sheet. It is different from the book value often used in the official report, which is slow in reacting to market conditions. Also bank assets here are securities and loans to the non-financial sector, which exclude those to other financial intermediaries. In the data, the net mark-to-market capital ratio of the financial intermediation sector - the ratio of net worth of the aggregate financial intermediaries to the securities and loans to the nonfinancial sector - tends to move procyclically, even though the gross capital ratio - the ratio of net worth to the book value total assets (including securities and loans to the other intermediaries) of some individual intermediaries may move procyclically. While Adrian and Shin (2010) show book leverage, i.e. the inverse of book capital ratio, is procyclical for investment bankers, He et al. (2010) and He et al. (2017) show that market leverage is countercyclical, in line with our model prediction of procyclical capital ratios.

where N_t is total bank net worth.¹⁶ Hence, in what follows, we only use the portfolio choices K_t^b and D_t of a representative bank with net worth N_t .

2.3. Bank runs

Within our framework, a bank run is a rollover panic, similar to the Cole and Kehoe (2000) model of self-fulfilling debt crisis. In particular, a self-fulfilling bank run equilibrium (rollover crisis) exists under the following circumstances: An arbitrary depositor believes that if other households do not roll over their deposits, the depositor will lose money by rolling over. This condition is met if banks' net worth goes to zero in the event of the run. As we discussed earlier, banks with zero net worth cannot operate. Because they cannot credibly promise not to abscond with deposits, any household who lends money to banks in the wake of the run will lose money.

The timing of events is as follows: At the start of $t+1$, depositors decide whether to roll over deposits. If a run equilibrium exists at $t+1$, they may choose not to roll over. If the panic happens, banks liquidate capital and sell to households at the liquidation price Q_{t+1}^* . Because households are less efficient at holding capital, Q_{t+1}^* will lie below the normal market price Q_{t+1} . Depositors then get back a fraction of the promised return, depending on the recovery rate x_{t+1} as defined in equation (16).

For computational simplicity as well as realism, we assume that new banks do not enter during the period of the panic: They wait until the next period when the run has stopped. Thus we have

$$\xi_t^N = 0, \text{ if there is a run at } t. \quad (23)$$

As discussed, the run equilibrium exists if bank net worth goes to zero in the event of the panic. This will be the case if the depositor recovery rate is less than unity. It follows that the run equilibrium exists at $t+1$ if the liquidation value of bank assets is less than the promised obligation of deposits:

$$(Q_{t+1}^* + Z_{t+1})K_t^b < \bar{R}_t D_t, \quad (24)$$

which is the same as the condition $x_{t+1} < 1$. The liquidation price in turn is given by the household's first order condition for capital holding,

$$Q_{t+1}^* = E_{t+1} [\Lambda_{t+1, t+2} (Z_{t+2} + Q_{t+2})] - \alpha K_{t+1}^h, \quad (25)$$

evaluated at $K_{t+1}^h = 1$.

Let ι_{t+1} be a sunspot variable that takes on a value of unity if the sunspot occurs and zero otherwise. Then a run occurs at $t+1$ if (i) condition (24) is met, and (ii) $\iota_{t+1} = 1$. In order to not introduce any exogenous cyclicalities into the likelihood of a banking panic, we assume the sunspot appears with fixed probability κ^s . Then, letting Z_{t+1}^R be the threshold value of Z_{t+1} below which a run is possible, the probability of a run p_t^R is given by

$$p_t^R = Pr\{Z_{t+1} < Z_{t+1}^R\} \cdot \kappa^s, \quad (26)$$

where Z_{t+1}^R is the value of productivity at which banks are just able to pay their deposit obligations even if prices drop to their liquidation value $Q_{t+1}^*(Z_{t+1}^R)$:

$$Q_{t+1}^*(Z_{t+1}^R) + Z_{t+1}^R = \frac{D_t \bar{R}_t}{K_t^b}. \quad (27)$$

Equations (26) and (27) suggest two forces that can raise the likelihood of a run equilibrium existing. The first is bad luck: a sequence of negative shocks to the productivity of capital can increase the likelihood that Z_{t+1} will fall below the threshold value Z_{t+1}^R . The second is banks' financial fragility, measured by the ratio of the deposit obligation to the book value of capital, $\frac{D_t \bar{R}_t}{K_t^b}$. A rise in leverage increases Z_{t+1}^R , raising the likelihood that Z_{t+1} will be below Z_{t+1}^R . To foreshadow, a credit boom will increase the banking sectors' exposure to panics by increasing leverage.

2.4. Aggregation and equilibrium

If there is no run at time t , the aggregate net worth of active banks is given by the net worth of surviving bankers from $t-1$ plus new net worth injected by households into new banks:

$$N_t = \sigma \left[(Z_t + Q_t) K_{t-1}^b - D_{t-1} R_t \right] + \xi_t^N. \quad (28)$$

¹⁶ When the constraint is binding, equation (21) holds with equality so that κ_t only depends on ψ_t^b and hence it is independent of individual bank's net worth n_t . When the constraint is not binding, κ_t will be pinned down by an arbitrage condition that expected discount excess returns equal zero (where the discount factor takes into account that the constraint might bind in the future). The arbitrage condition also depends on aggregate variables only so that it still does not depend on individual bank's net worth. See Appendix for details.

Notice that it is possible that, even without a bank run the realization of productivity could be so low that the banks are forced to default. In this case, equations (15) and (28) imply that aggregate net worth is simply given by ξ_t^N .¹⁷

Using the expression for the rate return on bank equity and for leverage, (18), (22), (23), yields:

$$N_t = X_t^N = \begin{cases} \sigma R_t^N N_{t-1} + \xi_t^N & \text{if there is no run at } t \\ 0 & \text{if there is a run at } t \end{cases},$$

where

$$R_t^N = (R_t^b - R_t) \frac{1}{k_{t-1}} + R_t. \quad (29)$$

Here we see the rate of return on bank equity R_t^N is increasing in the excess return weighted by the bank leverage multiple at $t-1$, the inverse of the bank equity - asset ratio.

2.5. Summary of key model ingredients

We conclude this section by summarizing the key features of the model that lead to the possibility of (i) banking distress affecting the real economy and, in the extreme, (ii) outright banking panics that precipitate sharp persistent contractions in real activity.

First, to have financial intermediation matter to real activity, banks must be more efficient than households (or more generally non-experts) in evaluating and monitoring capital.¹⁸ Otherwise a contraction in banking would not matter: Households would simply absorb the capital in their respective portfolios without any effect on real activity, and the price and rate of return on capital would remain unchanged. In particular, we capture banks' advantage by assuming that households face a management cost that is increasing and convex in the amount of assets that they hold directly in their portfolios. The convex cost captures the idea that households face capacity constraints in managing capital. It also implies that following a decline in bank intermediation, because of households limited ability to absorb capital, the price of capital declines and the rate of return increases. In the limiting case of a panic, which features a firesale of bank assets to households, the price declines sharply.

In our simple endowment economy model, the shift from bank intermediated to direct household finance reduces net output due to the rise in managerial costs. But we view this mechanism as a convenient way to model the real effects of disintermediation in an endowment economy. In Gertler et al. (2020b), which features a small scale macro model with banking, the decline in asset prices and increase in excess returns following a banking panic leads to a decline in investment and real activity. Also, while here a bank run causes the entire banking sector to be wiped out, in Gertler et al. (2016) we show that the same model of banking instability can be adapted to characterize runs on a sub-sector of the broad financial industry. Hence, one should think of the banks in our model as the most exposed and vulnerable sub-sector of the financial industry, corresponding for instance to the shadow banking sector in the run up to the Great Recession.

Second, to account for why banks do not intermediate the entire capital stock, we suppose that there is an agency problem which limits their ability to borrow in credit markets. In particular, we assume that banks fund assets with short term debt and equity (or net worth). Further, it is costly for banks to raise equity. Because of the agency problem, the amount of short term debt a bank can issue depends positively on its net worth. The greater the fraction of assets a bank funds with equity, the larger is the relative stake that it has in its asset holdings, making it more likely it will not mishandle depositor funds. This in turn induces creditors to lend more to the bank as its equity increases. The endogenous equity constraint gives banks the incentive to raise equity in order to relax the constraint. As we have discussed, though, banks rely heavily on leverage, even when they are under financial distress. We capture this behavior by assuming that it is costly for households to inject equity into the banking system. Further, we parametrize this cost function so that the model is consistent with both the trend and cycle in new issues of bank equity.

Finally, what makes banks financial fragile and vulnerable to runs? Here there are two key considerations. First, from the incentive constraint that falls out of the agency problem, banks cannot operate with zero net worth. With zero net worth, they will be unable to raise funds: Any rational creditor will realize that a bank with zero net worth has no stake in its' portfolio and will thus only mishandle funds. Second, banks rely heavily on short term debt. These two considerations open up the possibility of a self-fulfilling roll over panic. In such a panic, depositors collectively decide not to roll over their deposits to banks, forcing banks to liquidate assets. An equilibrium with a rollover crisis exists if the panic forces bank net worth to zero. In this instance, no individual creditor has the incentive to deviate from the group and supply the bank with credit, since the bank has zero net worth. Finally, we note that bank net worth goes to zero during a panic when the depositor recovery rate is less than unity; that is, when the liquidation value of bank assets lies below the face value of the banks deposit obligations. It follows that a run equilibrium is more likely to exist when bank leverage is high. As we discuss shortly, our belief driven credit booms are situations in which bank leverage increases. In this respect, the credit

¹⁷ See Appendix for a characterization of the probability of insolvency without runs.

¹⁸ In GKP (2020) we allow for the more realistic possibility that banks have an advantage in evaluating and monitoring only a subset of capital investments. That is, we suppose that households face convex managerial costs only when the share of capital they hold exceeds a threshold ratio.

Table 1

Calibration of baseline parameters.

Parameters	Description	Value	Target	Model
Calibrated Parameters				
θ	Share of Divertible Assets	0.23	Capital Ratios = 10 pct	$E(\kappa) = 10$ pct
σ	Banker Survival Rate	0.935	Quarterly Spread = 50 bpts	$E(R^b - R) = 48$ bpts
ξ	Startup Equity	1 pct of N^{SS}	HH Share of Intermediation = .5	$K^h = 0.49$
α_ξ	Equity Injections Costs	0.001	Average Issuance rate = 1 pct	$E \frac{\xi^N}{N^{SS}} = 1.1$ pct
α	HH Intermediation Costs	0.00625	Output Drop During Run = 6 pct	$\frac{Y^* - Y^{SS}}{Y^{SS}} = 6.4$ pct
χ^s	Sunspot Probability	0.125	Avg Yearly Frequency of Runs = 3.7 pct	$4 \cdot Ep^R = 3.6$ pct
$\sigma(\epsilon^Z)$	Std Dev of Z Innovation	0.01	Std Dev of U.S. Output = 1.9 pct	$\sigma(Y) = 1.9$ pct
Fixed Parameters				
β	Impatience	0.99	-	-
ρ^Z	Serial Correlation of Z	0.95	-	-
W	HH Endowment	$2 \cdot Z$	-	-

boom raises the likelihood of a bank run and can lead the system to a panic even in the absence of negative fundamental shocks to banks' assets.

The Appendix provides a detailed description of the equilibrium equations.

3. Credit booms and busts: a numerical illustration

We now show via numerical simulation how the model can generate credit booms and busts consistent with the evidence presented in Figs. 1 and 2. For expositional reasons, we first start with the bust phase of a crisis. That is, we consider a model where fundamental shocks are the outside force that drives the economy into a crisis zone where runs can occur. Here the idea is to illustrate how the model can generate a financial collapse which has spillover effects for the real economy.

We first describe how we calibrate our model. Then we illustrate how, starting with a banking system that is "safe", i.e. not susceptible to runs, a series of negative shocks can weaken bank balance sheets, moving the economy to a crisis zone where a financial collapse can occur. We then introduce our belief mechanism and show how it can generate a credit boom that may or may not lead to a bust.

3.1. Calibration

Table 1 shows the parameter values used in our experiments together with the calibration targets. There are ten parameters. Four are reasonably standard: including the discount factor β , the serial correlation of the capital productivity shock, ρ_Z , the standard deviation of this shock σ and the household "labor" endowment W . We set β at 0.99, a standard value in the literature. We choose a similarly conventional value for $\rho = 0.95$. We pick σ_Z so that the model produces a standard deviation of output equal to 1.9 percent, consistent with the evidence. Finally, we set W equal to twice the size of steady state capital income Z to capture the idea that on average the labor share is twice the capital share.

Six parameters govern the financial sector and are nonstandard. They include: the fraction of assets banker can divert θ , the banker survival rate σ , the parameter governing marginal household direct financing costs α , the new bankers endowment ξ , the parameters governing costs of equity injections, α_ξ , and the sunspot probability χ^s . We choose these parameters to hit the following six targets: (a) The average bank equity - asset ratio κ equals 0.1; (b) An average annual spread between the return on bank assets R^b and the deposit rate R of two hundred basis points¹⁹; (c) The average household share of asset holding equals one half; (d) An average annual run probability of 3.7 percent (roughly, one every twenty-five years); (e) An output contraction during a bank run of ten percent on average, consistent with the evidence from Krishnamurthy and Muir (2017); (f) An average ratio of bank equity injections and trend financial equity of 1 percent.

3.2. A run driven by fundamental shocks

Before introducing a belief mechanism that can generate credit booms and busts, we first illustrate how the model can generate a nonlinear financial crisis with fundamental shocks as the underlying driving force. Under our parametrization, a run equilibrium does not exist in the risk adjusted steady state. We accordingly suppose that at time 1, there is a negative innovation to productivity just large enough to move the economy into a crisis zone, i.e., an environment where a run equilibrium exists. Intuitively a large negative productivity shock can open up the possibility of a run by (i) reducing bank net worth and hence increasing bank leverage and (ii) reducing the liquidation price of bank assets.

The solid line in the upper left panel of Fig. 3 displays the path of the productivity shock. The diamond on the vertical axis is the threshold value of the productivity shock, Z_{t+1}^R , below which a run equilibrium exists at $t + 1$. The threshold

¹⁹ See Philippon (2015).

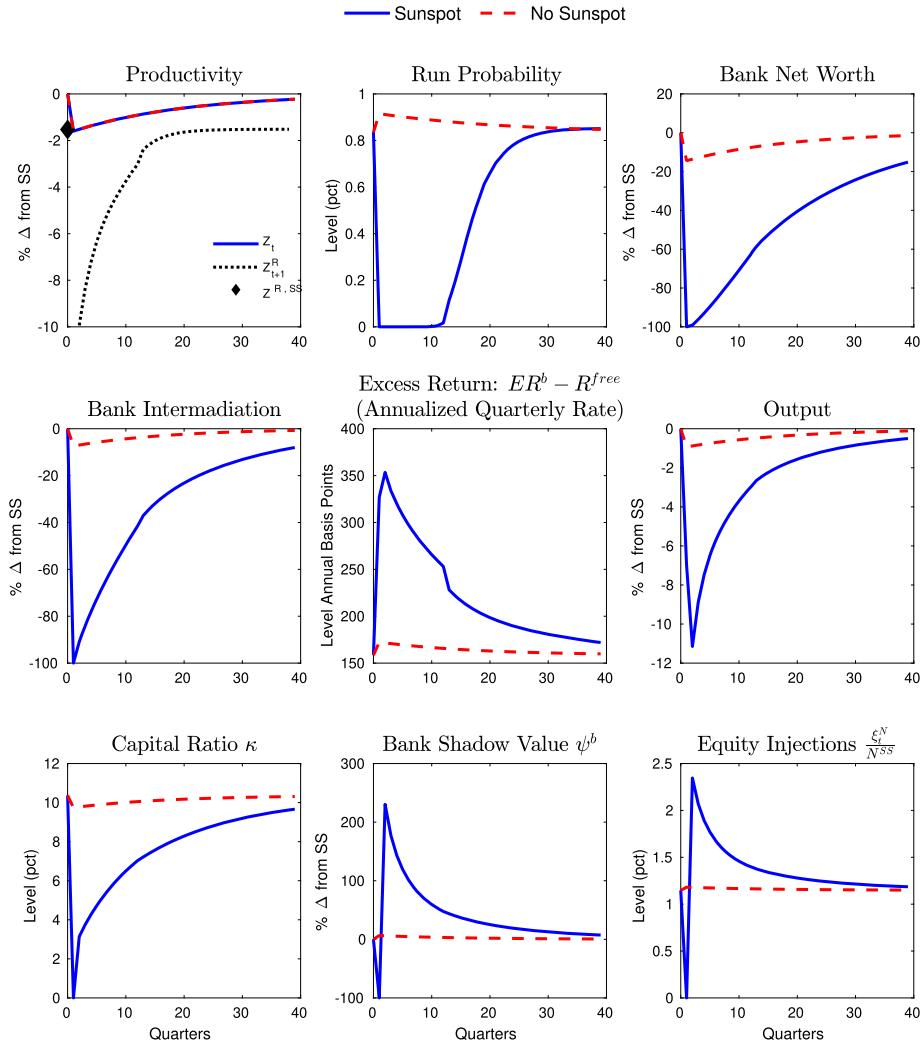


Fig. 3. Run after a large negative shock.

is almost two standard deviations below the risk adjusted steady state value of Z_t . As the panel illustrates the shock puts Z_{t+1} just below the threshold Z_{t+1}^R . Moving forward through time, the dotted line gives the crisis zone threshold for Z_{t+i} for each value of $i > 1$ after the run has occurred.²⁰

Given the economy reaches the crisis zone in period 1, we suppose there is a run, i.e. the sunspot appears and households do not rollover deposits. The solid line in each of the remaining panels gives the response of the economy in the case of the run. For comparison the dashed line shows the response for the case where the sunspot is not observed and hence the run does not occur. The run leads to a firesale of bank assets, causing bank net worth and bank intermediation to go to zero. Because it is costly for households to absorb the assets, the spread between the expected return on bank assets and the risk free rate jumps more than three hundred basis points, causing the shadow value of bank equity to more than triple. The disintermediation of bank assets leads to a sharp drop in output of more than ten percent. The figure makes clear the nonlinear aspect of the crisis. Absent the panic, output only drops less than one percent. In the wake of the run, the level of bank net worth slowly recovers as new banks enter and households increase equity injections in the financial sector in response to the sharp rise in the shadow value of bank equity. However, given that injecting equity is costly, the share of assets intermediated by banks recovers only slowly and so does output.

As discussed above, the assumption that equity injections in the financial sector are costly is key in order for financial frictions to have a bite and for banking panics to be possible. Fig. 4 shows that while we calibrated our cost function to

²⁰ If the run does not happen the threshold reverts back to its mean, as can be deduced by the behavior of the run probability and inspection of equation (26).

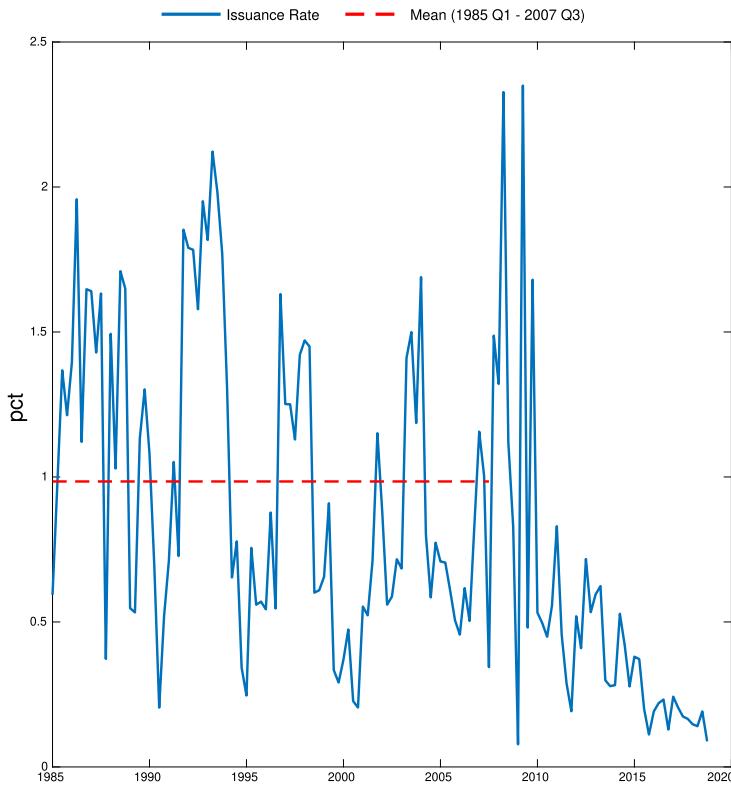


Fig. 4. Financial firms equity issuance as a fraction of trend equity.

match the average level of equity injections over time, our model predictions about the increase in equity injections after a crisis captures quite well the observed market response during the recent financial crisis.

3.3. News driven optimism and credit booms

One of the major weaknesses of the model of bank runs driven by fundamental shocks is that financial crises often occur without major productivity shocks, as in the recent Global Financial Crisis. To address this, we now extend the model to allow for credit booms, building on our earlier work, GKP (2020). In that framework, news that bankers receive about the possibility of improved fundamentals lead to a credit buildup. However, because the improved fundamentals do not materialize, the high leverage pushes the economy into a crisis zone where a banking panic is possible. Here we allow for the possibility that the credit booms can lead to good as well as bad outcomes. Good outcomes are possible either because the improved fundamentals arise or because, even if they don't, depositors do not coordinate on the run equilibrium and the panic never materializes. In this latter case, the credit boom raises the share of intermediated finance, which is expansionary even if the improved fundamentals do not arise. In the end, our goal is to match the data presented earlier in Fig. 2, which shows that, while high credit growth makes a crisis more likely, it typically does not lead to a crisis. Conversely, crises can occur in the absence of large credit growth.

Following GKP (2020), we model beliefs by considering a variant of a "news" shock. Under the standard formulation, at time t , individuals suddenly learn with certainty that a fundamental disturbance of a given size will occur j periods in the future. We relax this assumption in two ways. First, we assume that there is a probability the shock may not occur. Second, we assume that rather than having a single date in the future when the shock can occur, there is a probability distribution over a number of possible dates. As time passes without the occurrence of the shock, individuals update their priors on these various possibilities. We also assume that only bankers, who are the experts at managing assets, have optimistic beliefs. In fact, it is the relative optimism of bankers compared to households that generates the vulnerability of the financial system.²¹

²¹ As we describe in Appendix, we assume that households are aware that bankers became optimistic but do not change their beliefs about the productivity of capital, i.e. they do not believe the news. This allows us to have diverse beliefs without having households extract information from prices. A similar assumption is made for the same reason, for instance, in Cogley and Sargent (2009). Because households know bankers are more optimistic, they understand that there is less danger for bankers to divert their assets and loose their franchise. This allows bankers to raise their leverage multiple.

Table 2
Calibration of news shocks.

Parameters	Description	Value
$\mu(t^B)$	Expected time of Z boom	10.5 Quarters ahead
$\sigma(t^B)$	Standard Deviation of Prior	2 Quarters
T	News Horizon	21 Quarters
B	Size of Productivity Boom	$2 \cdot \sigma(\epsilon^Z)$
\bar{P}_0^B	Banker Prob. that Boom will happen	0.999
\bar{P}_0^{TRUE}	True Prob. that Boom will happen	0.5
χ^n	Prob. of Receiving News	0.02

In particular, with some fixed probability χ^n , at time t^N bankers receive news that there may be a high return on capital in the form of a large positive capital productivity shock. But they do not know for sure (i) whether the shock will occur and (ii) conditional on the shock arriving, when it will occur. If the shock is realized at some time $\tau > t^N$, it takes the form of a one time impulse to the capital productivity shock process of size $B > 0$. Formally, the news bankers receive is that the capital productivity will follow the process

$$Z_\tau = 1 - \rho_\xi + \rho_\xi Z_{\tau-1} + \epsilon_\tau + \tilde{B}_\tau \text{ for } \tau > t^N$$

where $\tilde{B}_\tau = B$ if the large shock realizes at τ , and $\tilde{B}_\tau = 0$ otherwise. Given the capital productivity shock is serially correlated, there will be a persistent effect of B . However, given it is a one time shock, if it occurs, there will be no subsequent realizations of this impulse. In contrast to our earlier paper, though, we will allow for recurrent (though infrequent) news shocks as we describe below.

When they receive the news at t^N , bankers' prior probability that a shock will eventually occur is given by \bar{P} . Conditional on the shock happening, the future date when it will happen, $t^N + \tau \in \{t^N + 1, t^N + 2, \dots, t^N + T\}$, is random. In particular τ is distributed according to a probability mass function ζ_τ which we assume to be a discrete approximation of a truncated normal with mean μ^B , standard deviation σ^B , and support $[1, T]$. Thus at date t^N , the probability that the shock happens at $t^N + \tau$ is given by

$$\Pr_{t^N}(\tilde{B}_{t^N+\tau} = B) = \begin{cases} \bar{P} \cdot \zeta_\tau, & \text{for } \tau = 1, 2, \dots, T \\ 0, & \text{for } \tau > T \end{cases}.$$

As long as no shock is observed for ν quarters, i.e. until date $t = t^N + \nu$ ($\geq t^N + 1$), bankers update their beliefs using Bayes rule:

$$\begin{aligned} \Pr_{t^N+\nu}(\tilde{B}_{t^N+\tau} = B) &= \frac{\bar{P} \cdot \zeta_\tau}{1 - \sum_{j=1}^{\nu} \bar{P} \cdot \zeta_j} \\ &= \frac{\sum_{j=\nu+1}^T \bar{P} \cdot \zeta_j}{1 - \sum_{j=1}^{\nu} \bar{P} \cdot \zeta_j} \cdot \frac{\zeta_\tau}{\sum_{j=\nu+1}^T \zeta_j}, \end{aligned} \quad (30)$$

for $\tau = \nu + 1, \dots, T$, and $\Pr_{t^N+\nu}(\tilde{B}_{t^N+\tau} = B) = 0$ for $\tau > T$. The first term in the last line is the posterior probability of the shock ever happening, which we denote by \bar{P}_t and which is decreasing with t . The second term is the probability that the shock realizes at $t^N + \tau$ conditional on the shock eventually happening. The latter is increasing with ν until $\nu = T - 1$, before becoming zero.

Observe that the process will generate a burst of optimism that will eventually fade if the good news is not realized. Early on, bankers will steadily raise their forecasts of the near term return on capital as they approach the date where, a priori, the shock is most likely to occur. As time passes without the realization of the shock, bankers' become less certain it will ever occur: The optimism proceeds to vanish.

We now illustrate how with the belief mechanism just described, the model generates a boom/bust scenario. Table 2 describes our calibration of the belief process. We assume that bankers receive the optimistic news ten quarters in advance of the prior on the most likely date the boom in fundamentals is likely to occur. Our empirical motivation is the housing boom which began in early 2005 and peaked roughly ten quarters later. Accordingly we pick the mean of the conditional distribution of τ , μ^B , so that the prior on when the shock is most likely to occur is ten quarters after the receipt of the news. We pick the standard deviation σ^B to ensure that by six quarters after the conditional mean, if the shock has not occurred, bankers' will give up hope that it will ever occur.²² Next we set the size of the impulse B to equal a two standard

²² Given our discrete approximation of the normal distribution, a choice of σ^B translates into a maximum numbers of periods within which the shock can occur.

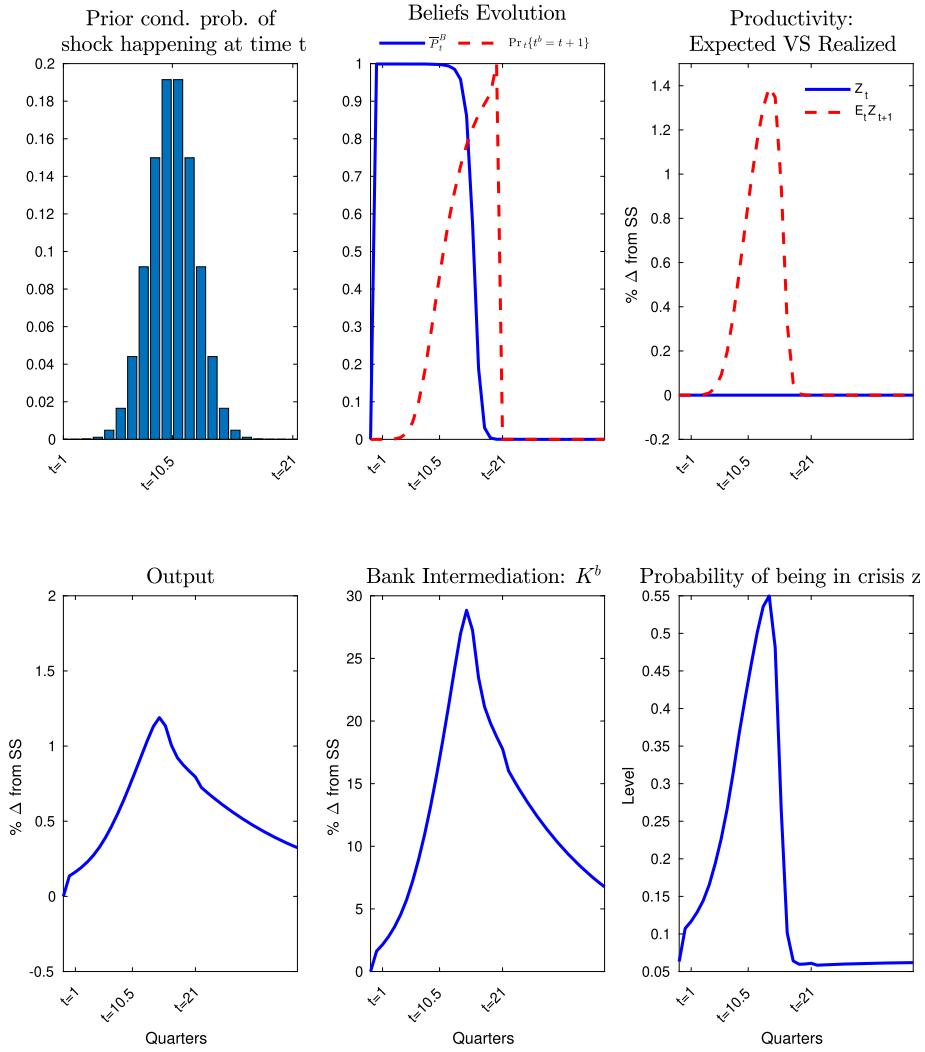


Fig. 5. Belief dynamics and credit booms.

deviation shock, that is, a shock which is unusually large but not beyond the realm of possibility.²³ Finally, we pick the prior probability that the shock will even occur \bar{P} , to ensure that economy reaches the crisis zone six quarters after the conditional mean without any fundamental shocks.

Fig. 5 characterizes the dynamics of beliefs and the credit boom that can emerge absent any fundamental shocks. The top-left panel gives the prior distribution for the time the shock will happen, conditional on it happening, i.e. $\{\zeta_{t^N+i}\}_{i=1}^T$. The top-middle panel then illustrates the ingredients bankers use to forecast the shock. The blue line gives the probability the shock will eventually happen, \bar{P}_t . When the news is received at $t = 1$, the probability jumps to its prior value near unity. Time passing without the shock occurring leads bankers' to reduce this probability. The optimism fades rapidly as time passes the conditional mean, the most likely time the shock was expected to occur. The dashed red line then gives the probability the shock will occur in the subsequent period, conditional on it eventually happening. Notice that this conditional probability equals unity at date $t^N + T - 1$ when the next period is the last possible date for the shock to occur, see equation (30) with $v = T - 1$. The estimate that the shock will occur in the subsequent period is then the product of the blue and red lines.

We choose our belief process to capture the idea that, once bankers become optimistic, their optimism is relatively resilient to disappointing news until the later stages of the boom. The relative optimism of bankers during the early stage of the boom leads to a shift in the allocation of capital finance from household to bank portfolios. Further, the way banks

²³ Note that the prior probability that the shock will occur, \bar{P}_{t^N} , and the size of the shock when it occurs, B , only influence the expected capital productivity through their product $\bar{P}_{t^N} \cdot B$.

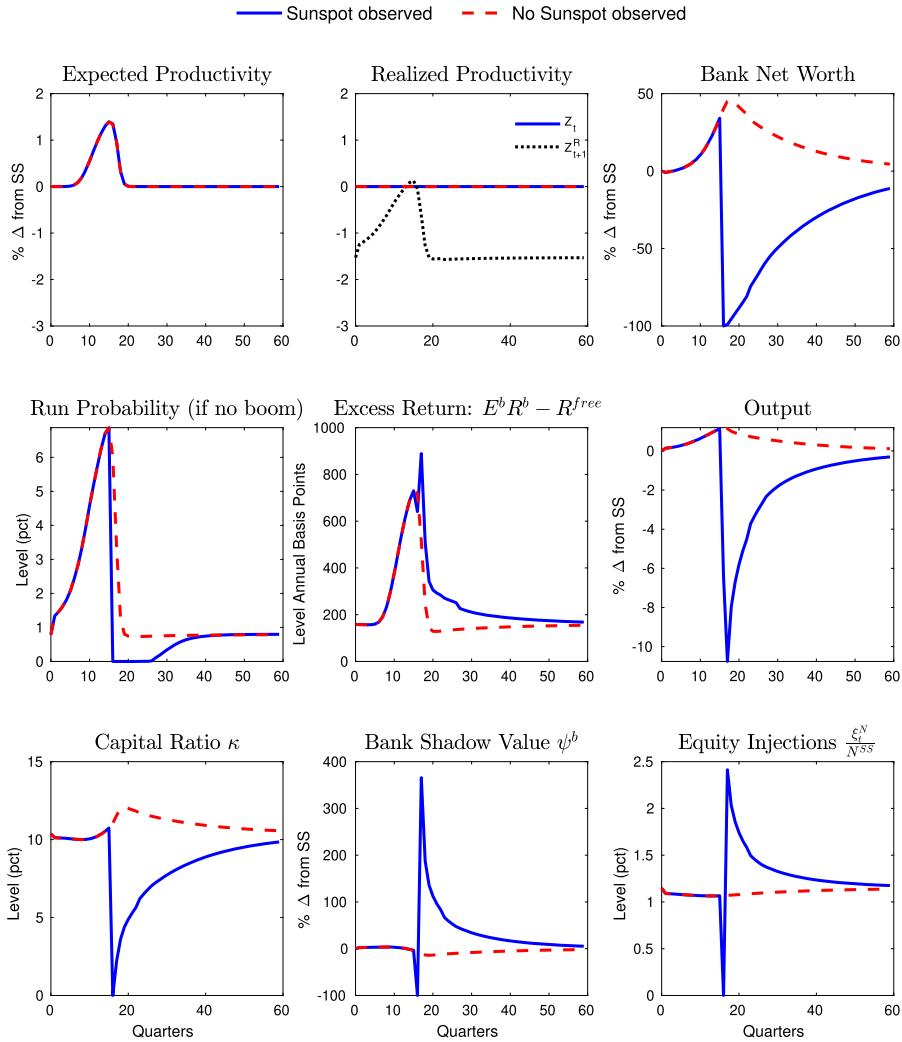


Fig. 6. A bank run after a credit boom.

expand their asset holding is by increasing their borrowing. As a result, bank leverage increases, which in turn increases the vulnerability of the banking system to a panic.

To illustrate the boom/bust nature of beliefs, the top-right panel portrays one quarter ahead forecast of the productivity shock (the dashed red line). After receiving the news at $t = 1$, optimism steadily builds. However, as time continues to pass without a large productivity improvement happening, the optimism fades. Note that throughout the boom and bust in beliefs, the true fundamental shock (the blue line), is unchanged. Thus, there is serial correlation in the forecast errors of the capital productivity shock.

The bottom-left panel shows the response of output to the news. The increase in bankers' optimism leads bankers to expect higher returns on assets which induces a rise in bank intermediation and, in turn, an increase in output of nearly one percent. There is however a nontrivial buildup of debt as bankers fund the twenty five percent increase in assets mostly by issuing deposits in the bottom-middle panel. The bank capital ratio (equity to assets) in fact declines as bankers' optimism raises their perceived shadow value of net worth ψ_t^b , relaxing the incentive constraint.²⁴ (See equation (21).) The increase in leverage raises the probability the economy moves into a crisis zone where a run is possible, as the bottom-right panel shows. In this regard, the boom lays the seeds of the bust.

We now illustrate how a wave of optimism can generate a credit boom that leads to a banking panic. Fig. 6 illustrates the experiment. The news of a possible improvement in fundamentals is received in period 1. The prior probability distribution is as described in the previous figure. The top-left panel is the one period ahead forecast of capital productivity. Expected productivity increases as the economy approaches the prior conditional mean. However, because the productivity boom is

²⁴ As discussed by Gertler et al. (2016), there were additional factors contributing to the leverage buildup, including financial innovation. For simplicity we abstract from these factors and note only that including them would increase the debt buildup further and the resulting degree of fragility.

not realized, the expected productivity begins to decline in the later periods. As just described, bankers' optimism leads to an overall increase in bank assets funded by a rise in bank leverage, which moves the economy into a crisis zone. In the top-middle panel, the solid line is realized productivity, which is unchanged throughout. As before, the dotted line is the threshold value for the capital productivity shock, Z_{t+1}^R , below which a run equilibrium exists. As the panel makes clear the news shock moves the economy steadily toward a crisis zone, which it reaches roughly three years later.²⁵

Here we illustrate a case in which, once the economy reaches the crisis zone, the sunspot appears and a rollover panic ensues. The difference from the earlier case is that we do not require a fundamental shock to move the economy to a crisis zone, so we do without it. Overall, the effect of the banking crisis is very similar to the case without the debt boom. The contraction in output in terms of both amplitude and persistence is similar to the case of the fundamentals driven panic. As before the spread between the expected rate of return on bank assets and the risk-free rate increases prior to and during the crisis (in the middle-middle panel), again consistent with the evidence presented in the introduction. One important difference is that the wave of optimism generates a credit boom prior to the crisis, also consistent with the facts we presented earlier. Finally, as shown in the bottom-right panel, despite the increase in fragility of the banking sector households do not start injecting equity until after the crisis occurs. This is because the increase in fragility in this case is by an excessive optimism of financial intermediaries that is not shared by households. Accordingly, households' expectations of future bank excess returns do not increase with those of bankers, leading their subjective probability of a crisis to rise. As a result, households' desire to hold bank equity slightly declines before the crisis occurs.

We next illustrate that, consistent with the earlier evidence we presented, it is possible to have a credit boom that does not lead to a crisis. There are two possible reasons for why. First, the positive fundamental shock actually materializes. Second, the shock does not materialize but the panic doesn't arise even though the economy is in a crisis zone because depositors do not coordinate on the bad equilibrium (i.e. the sunspot doesn't appear). Fig. 7 displays both cases. As in the previous experiment, bankers receive positive news at time 1. The solid lines portray the case where the large productivity improvement materializes as bankers expected. In this case the expected jump in productivity arises in period 10, the peak of the conditional prior mean. The runup to period ten is identical to the case where a panic occurs, as portrayed in the previous figure. However, the realization of the productivity improvement leads to an increase in output (in the middle-right panel), which moves the economy out of the crisis zone as the top-middle panel shows. The dashed lines are the case where the boom never occurs but a panic still does not arise because the sunspot does not appear. There is in fact a rise in output, though smaller than in the case where the productivity boom is realized. The source of the rise in output is the optimism that gives rise to an increased in the share of capital intermediated by banks.²⁶

Thus far we have characterized single episodes of credit booms and displayed circumstances where they may or may not lead to a bank run. As a prelude to analyzing macroprudential regulation, we next consider recurrent credit booms and busts. Our goal is to match the evidence on the link between credit growth and the frequency of financial panics described in Fig. 2. We assume that the probability of receiving news x^n is equal to 2 percent per quarter, which corresponds to once every twelve and a half years on average. Further, once news is received, there is no additional news realization until the current process has played out, i.e. there is no news from $t^N + 1$ until either $t^N + T$ or the period in which the boom actually happens.²⁷

We suppose the true probability the boom actually happens is fifty percent conditional on bankers receiving the news. We capture the idea that bankers are optimistic by supposing that upon receiving the news, they have a strong prior probability of .999 that the boom will happen. Given that credit booms are relatively infrequent it is not unreasonable to suppose that bankers have not had enough experience to learn the true probability of good realizations. Alternatively, think of the high prior as capturing a "This Time is Different" mentality.²⁸

We simulate the model and then record the relation between the occurrence of a crisis in a given year and credit growth in the two preceding years in Fig. 8. The left panel shows the data from Jordà et al. (2011) as in Fig. 2. The right panel is the simulation result of the model. The model does a reasonable job of capturing that, as in the data, crises are more likely following a sustained period of positive credit growth. Within the model, conditional on positive credit growth in the prior two consecutive years, a crisis occurs 4.9 percent of the time just as in the data. Runs without credit booms are a bit more frequent in the model than in the data, i.e. 3.2 against 2.8, but overall the predictive power of credit booms for banking crises as captured by the odds ratio of bank runs with and without a boom is in line with the empirical counterpart: 1.5 in the model against 1.8 in the data. One difference though is that credit growth in the model is less persistent than in the data.

²⁵ The run threshold in the figure is the one associated with the case in which the run does not happen. This facilitates comparison with Fig. 7 and helps to visualize that the threshold crosses the realized level of productivity 13 quarters after the shock, making the run possible. The behavior of the threshold after a run is as depicted in Fig. 3.

²⁶ Interestingly, without the realization of productivity improvement, the bank net worth increases more than the case of the realization, as long as there is no run, because the excess return on bank asset is larger.

²⁷ Thus the unconditional mean arrival rate of news is lower than two percent per quarter.

²⁸ See Reinhart and Rogoff (2009).

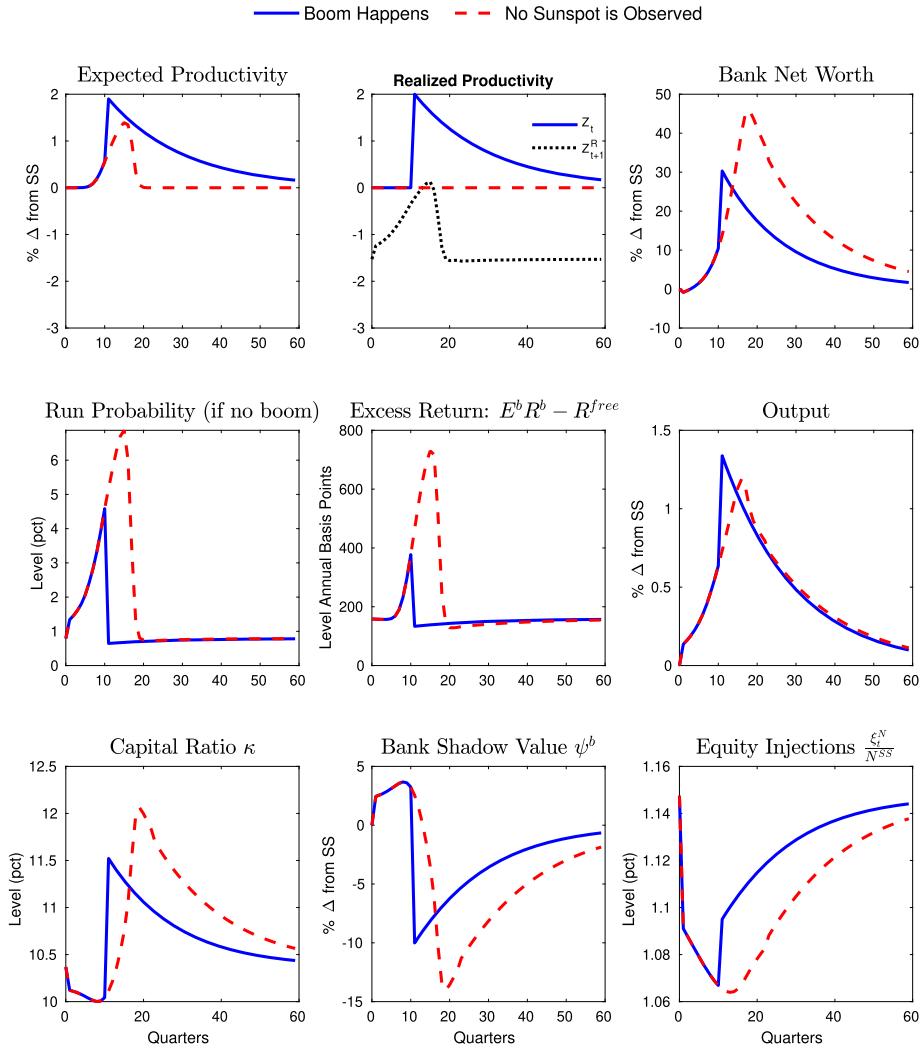


Fig. 7. Good booms.

4. Macroprudential regulation

Within our framework, the decentralized banking equilibrium is inefficient for two reasons. First, as in Lorenzoni (2008), there is a pecuniary externality that leads banks to fail to internalize the impact of their leverage decisions on the behavior of the price of capital. Second, banks also fail to internalize the impact of their leverage decisions on the likelihood of a panic - we call this externality "run externality". The inefficiency of the decentralized equilibrium provides a rationale for macroprudential policy. Indeed, it turns out that the quantitative gains from macroprudential policy in our framework are associated with the run externality.

We consider a macroprudential regulator that sets a time varying bank capital requirement $\bar{\kappa}_t$. This implies that the relevant capital requirement for banks, κ_t , is now the maximum between the regulatory requirement, $\bar{\kappa}_t$, and the market imposed capital requirement κ_t^m , given by equation (21). That is,

$$\kappa_t = \max(\bar{\kappa}_t, \kappa_t^m), \quad (31)$$

with $\kappa_t^m = \theta / \psi_t^b$.²⁹

We consider a simple policy rule for bank capital requirements that allows for a countercyclical buffer. Let \bar{N} be a threshold value of net worth in the banking system above which the capital requirement is set at the "normal value" $\bar{\kappa}$.

²⁹ The presence of regulation implies that the equilibrium value of ψ_t^b will be different in the regulated economy as we explain below.

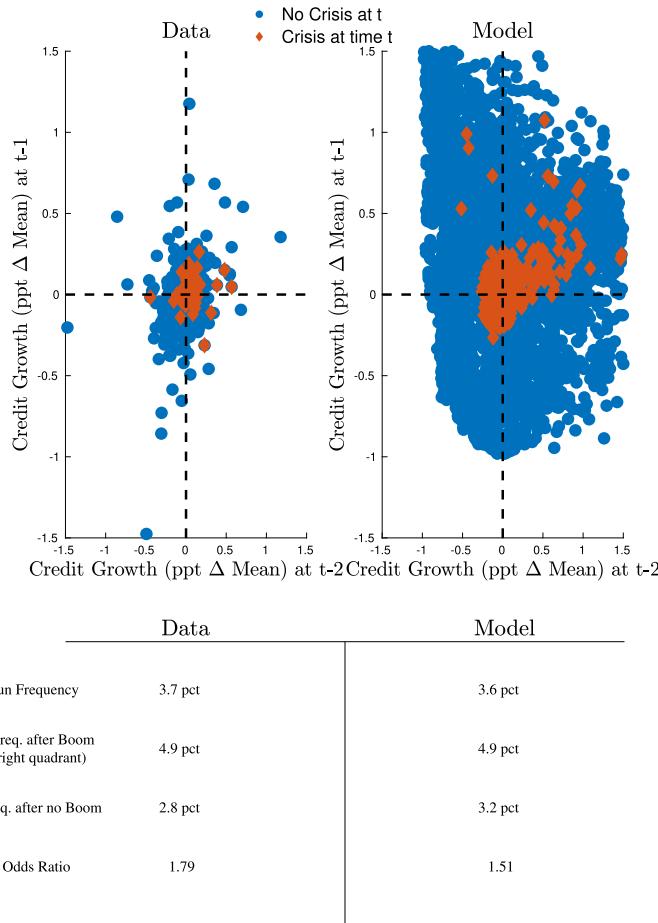


Fig. 8. Good and bad booms in the model and in the data.

When bank net worth falls below N , the requirement is relaxed.³⁰ We assume for simplicity the regulatory requirement goes to zero. In this instance the market requirement κ_t^m will apply.

We restrict policy to be determined by the simple rule

$$\bar{\kappa}_t = \begin{cases} \bar{\kappa} & N_t \geq \bar{N} \\ 0 & N_t < \bar{N} \end{cases}.$$

We look for $(\bar{\kappa}, \bar{N})$ that maximize welfare, which we take to be the unconditional expected utility of the representative household. Note that the rule allows for a countercyclical capital buffer, since the capital requirement is relaxed when aggregate bank net worth drops below the threshold N .³¹

Fig. 9 shows the market determined capital requirement in the unregulated equilibrium. At the value of equity in the risk adjusted steady state N_{SS}^{DE} , the capital requirement is ten percent. As bank net worth falls below the risk adjusted steady state the market capital requirement falls as well. With low bank net worth, bank credit availability is lower, implying high excess returns to bank assets. The high excess returns are associated with a high shadow value of bank net worth, which relaxes the incentive constraint permitting greater leverage and hence leads to a lower market determined capital requirement. Conversely, as net worth goes above steady state, excess returns fall which tightens capital requirements.

Fig. 10 then compares the optimal regulatory capital requirements in the solid line with ones arising in the unregulated equilibrium in the dashed line. The threshold \bar{N} lies below the risk adjusted steady state value N_{SS}^{DE} . When net worth falls below \bar{N} , the regulatory requirement falls to zero. Conversely, when it goes above \bar{N} , the requirement goes to twelve percent,

³⁰ For computational reasons, the criterion to relax the capital requirement is based on the level of bank net worth. However, given the very high correlation between bank net worth and output in the model, similar results would follow if the criterion for relaxing the capital requirement was based on output.

³¹ See Kashyap et al. (2004) for a discussion of countercyclical capital requirements. Faria-e Castro (2019) is another paper that studies countercyclical capital buffers in a model with bank runs.

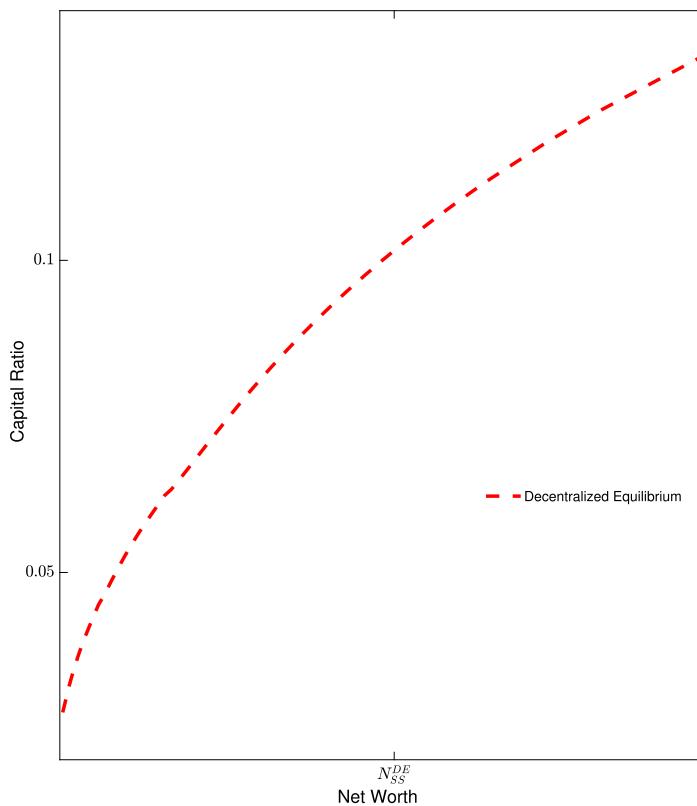


Fig. 9. Equilibrium capital ratios in the decentralized economy.

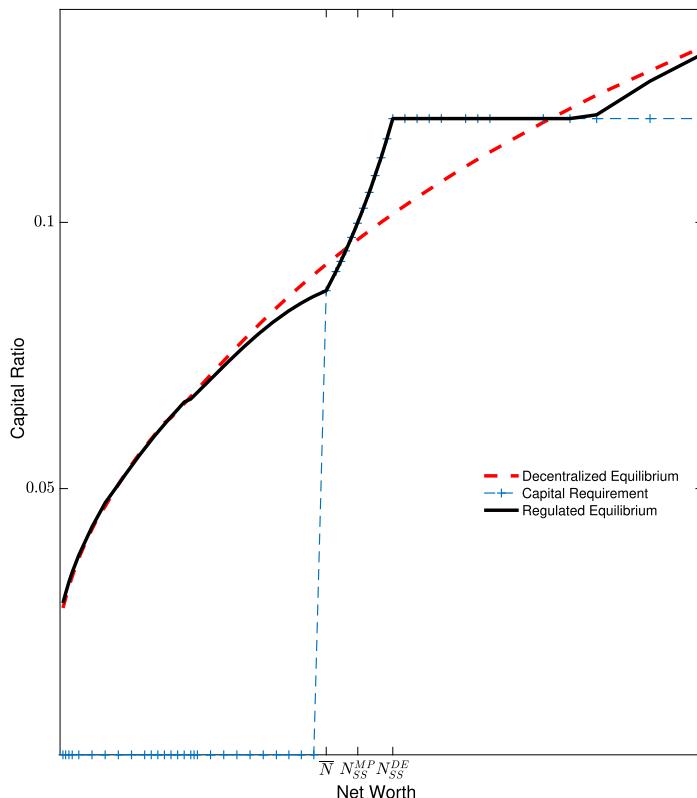


Fig. 10. Equilibrium capital ratios in the regulated economy.

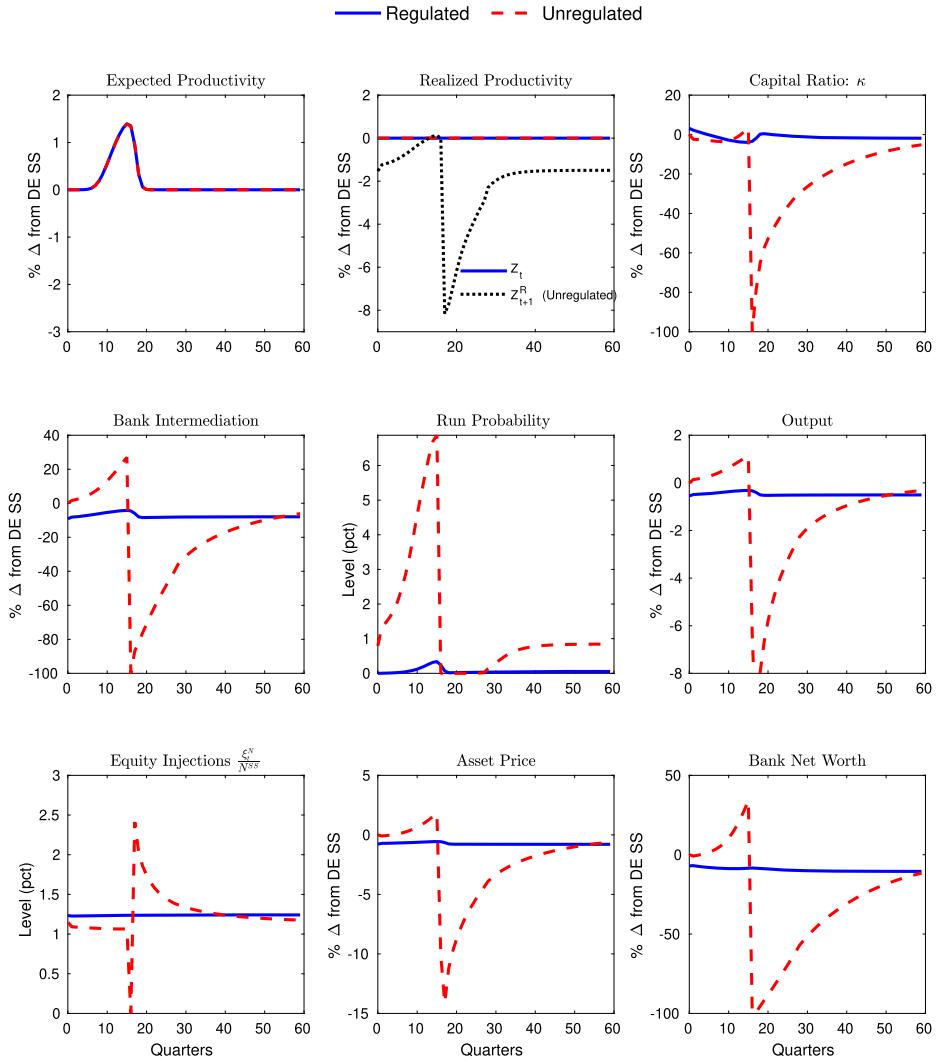


Fig. 11. Avoiding runs with macroprudential regulation.

which is above the steady state requirement for the unregulated equilibrium. For computational reasons, we smooth out the increase as N rises above \bar{N} .

Fig. 10 shows the pattern of capital requirements for the regulated equilibrium. Regulatory capital requirements are binding for intermediate levels of net worth. When bank net worth is very low, $\bar{\kappa}_t$ drops to 0 so that market requirements become binding. When net worth is high enough, the induced decline in excess returns causes market determined capital requirements to exceed $\bar{\kappa}$.

Note that as bank net worth is just below the threshold where capital requirements are binding, the market determined requirement for the regulated economy actually falls below the capital requirement for the unregulated case. This is because the shadow value of bank net worth is higher in the regulated economy than in the unregulated economy. Intuitively, when regulatory requirements are binding, the shadow value of net worth in the regulated economy is higher than in the unregulated equilibrium since the run probability is lower and excess returns on bank assets are higher due to the anticipated regulation in future. This in turn has a positive impact on the shadow value of net worth when banks are close to the regulatory threshold since they will eventually move to the region where the regulatory requirements applies.³²

We next analyze how the optimal macroprudential policy affects behavior. In Fig. 11 we consider a optimism driven credit boom of the type that leads to a banking panic. The dotted line portrays the credit boom and bust that occurs in the unregulated equilibrium. The solid line is the behavior with the macroprudential policy put in place. For comparability, we suppose the economy begins in the unregulated equilibrium, so that in both cases the initial state is the risk adjusted

³² See Van der Ghote (2018) for a similar argument for a coordinating monetary and financial regulation policies.

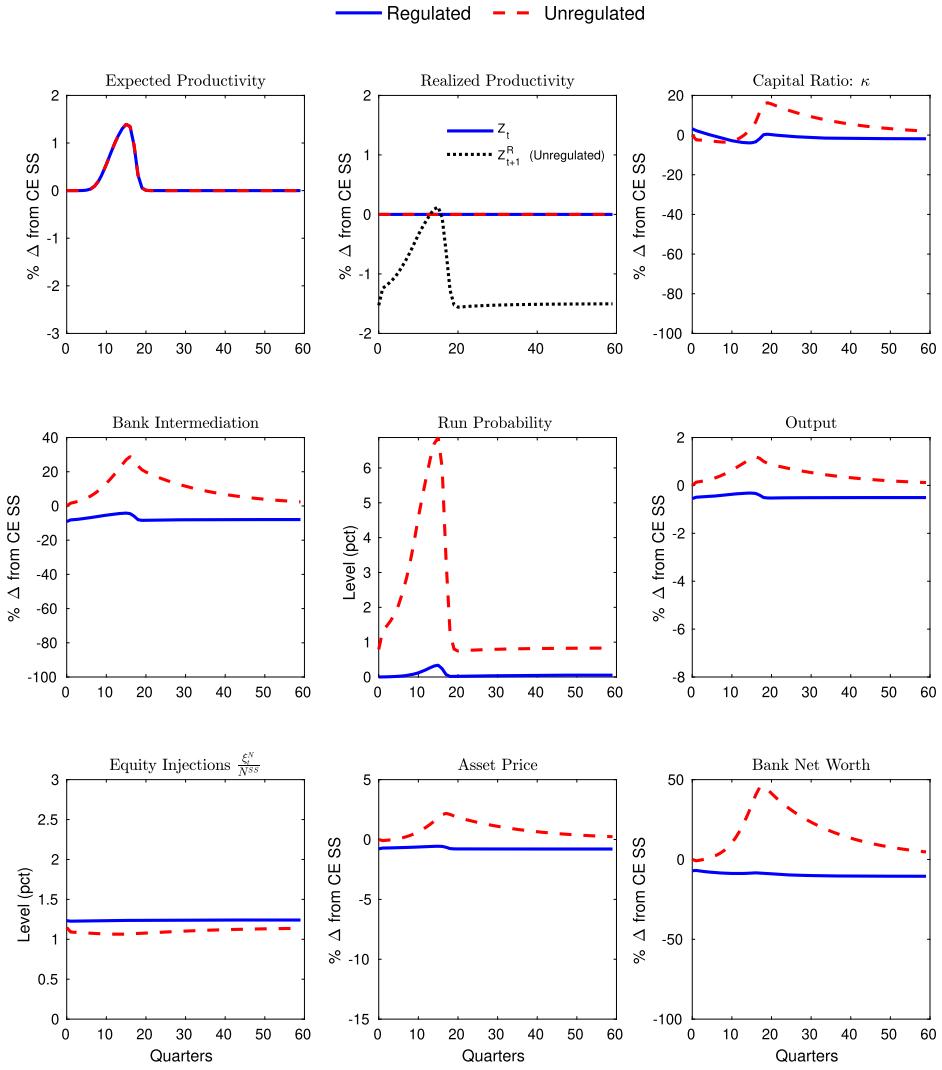


Fig. 12. Stifling good booms with macroprudential regulation.

steady state of the unregulated economy. The macroprudential policy is then imposed at time 0. The tightening of capital requirements produces an initial drop in bank intermediation. As in the unregulated equilibrium, the optimism wave which fails to be validated by a productivity boom leads to an increase in the run probability. But this increase is far more modest than in the unregulated equilibrium. Absent any large negative shock to fundamentals, the economy never enters a crisis zone. The regulation avoids a panic in this case. The cost is that output growth is muted during the optimism phase.

In Fig. 12 we consider a case where the credit boom is a false alarm. We consider the example where the fundamental does not materialize but the panic still does not occur (i.e., the sunspot does not appear). In this case the unregulated economy would produce a modest output boom. Thus, in this instance, the unregulated economy yields a better outcome. The same would be true for the case where the productivity boom is realized. Accordingly, the gain from macroprudential regulation is reducing the likelihood of a costly banking panic. This gain of course must be weighed against the cost of constraining the economy during credit booms that are false alarms.

Fig. 13 shows how macroprudential policy affects the distribution of output and welfare. By preventing boom bust cycles in credit as well as good booms, macroprudential policy induces a much less variable distribution of output while having only negligible effects on average output. This stabilization properties however have non-negligible effects on welfare as the policy is effective in reducing the probability of the large and persistent drops in output associated with bank runs.

The overall effects of the optimal macroprudential policy on output, the run probability and welfare are reported in the middle column of Table 3, which also reports the behavior of the decentralized economy in the left column. Macroprudential policy cuts the quarterly run probability more than half, to 0.4 percent from 0.9. The capital requirements lead to a reduction in quarterly output of 0.6 percent during periods without a banking crises. However, because the likelihood of costly banking

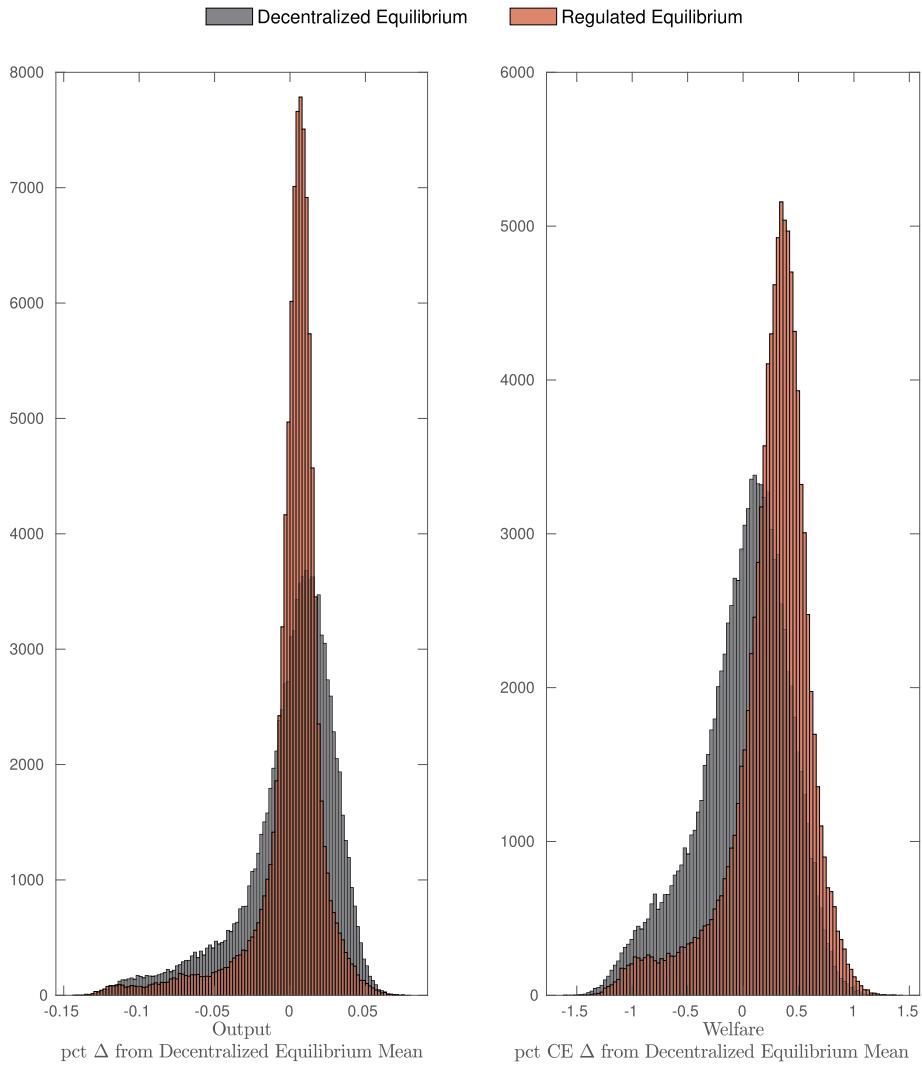


Fig. 13. Distribution of output and welfare: decentralized and regulated economy.

Table 3
Effects of macroprudential regulation.

	Decentralized Equilibrium	Optimal Regulation ($\bar{\kappa} = 0.12$; $\bar{N} = 0.8 * N_{SS}^{DE}$)	Fixed Capital Requirements ($\bar{\kappa} = 0.12$; $\bar{N} = 0$)
Run Frequency	0.9 pct	0.4 pct	0.4 pct
AVG Output Cond No Run	0 pct	-0.6 pct	-0.7 pct
Δ from Decentralized Economy			
AVG Output	0 pct	0.1 pct	-0.9 pct
Δ from Decentralized Economy			
Welfare Gain	0 pct	0.25 pct	-0.77 pct
Δ Permanent Consumption			

panics is reduced, average output is 0.1 percent higher.³³ The average output with regulation is slightly higher than without regulation because the loss of output due to banking crisis is deep and persistent even though the crises are rare and the average output is lower during booms. Combined with the reduction in the variance and left skewness of the output distribution, this delivers an increase in welfare of 0.25 percentage points of steady state consumption per period. Note that this is a very conservative estimate since we are using log preferences with a coefficient of relative risk aversion of unity.

³³ In Gertler et al. (2020a), we elaborate on the point that the main gains from macroprudential policy come from reducing the likelihood of costly panics.

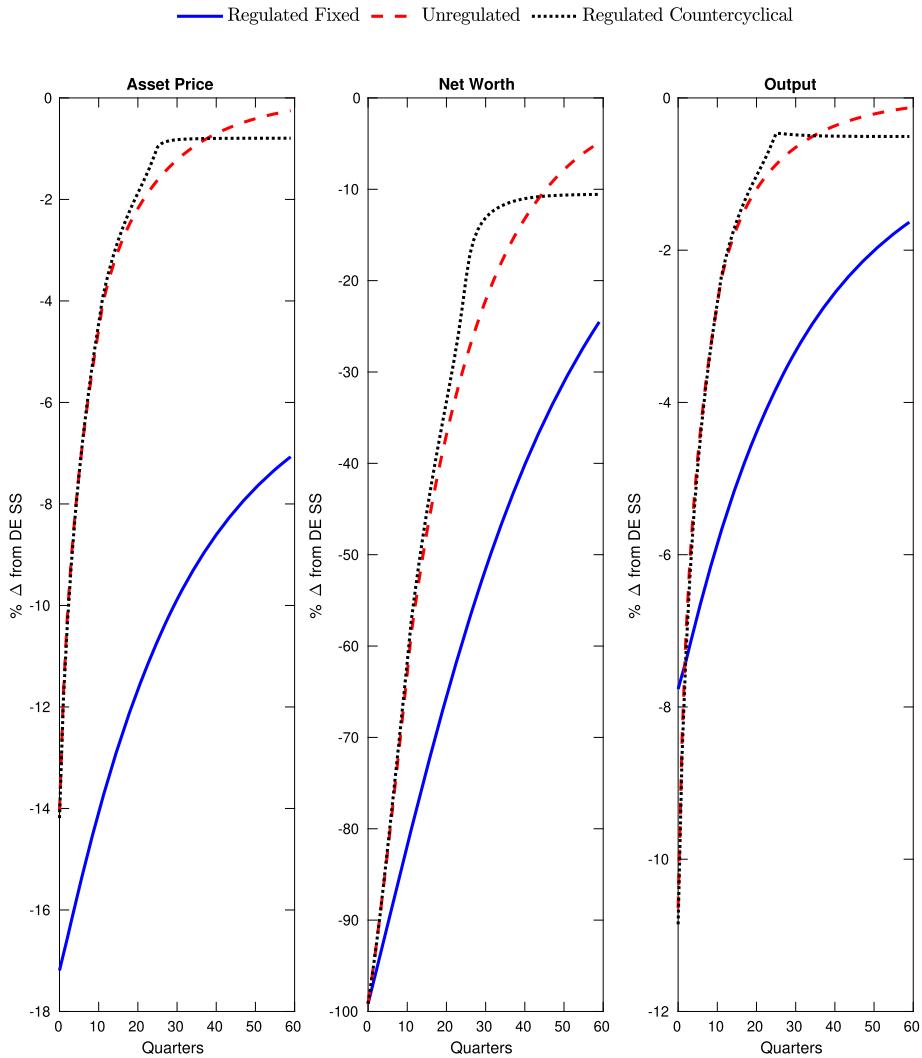


Fig. 14. Recovery from a run: the role of countercyclical buffers.

The last column in Table 3 portrays the case where we eliminate the countercyclical capital buffer and instead assume that regulatory capital requirements are uniform over the cycle. This policy has the same effect on the run probability as the optimal countercyclical policy, but this reduction in the run probability comes at a much higher cost in terms of output which ends up being almost one percent below the unregulated equilibrium on average. The net effect is that the policy produces a welfare loss of about three quarters percent of steady state consumption each quarter.

Fig. 14 illustrates why not relaxing the capital requirement in bad times has harmful effects. Under the optimal policy (the dotted line), relaxing capital requirements allows banks greater freedom to issue deposits to invest in high excess return assets after the crisis at date 0. This in turn allows banks to build their equity base at a faster pace, returning the economy to normal. By contrast, if capital requirements are rigid and not relaxed after the crisis (the solid line), banks build equity at a much slower pace, implying a more protracted period of low output.

5. Concluding remarks

We develop a simple quantitative model of credit booms and busts. The framework is consistent with the evidence that credit booms tend to lead crises, but most of the time a boom does not lead to a bust. The model also replicates other key features of financial crises, including increasing credit spreads and sharply contracting output. Importantly, the model captures the nonlinear dimension of financial crises. Much of the time, the economy operates in a “safe zone” with a banking system that is financially strong and not susceptible to a run. However, a belief driven credit boom or a series of bad fundamental shocks can raise bank leverage ratios, making the system vulnerable to runs. These runs, further have

costly effects on the real economy. Because the model is highly nonlinear, we use global methods to solve it numerically, as discussed in the appendix.

We then use the framework to study macroprudential policy. The particular policy we consider is a capital requirement that limits bank leverage. The primary goal of this policy is to reduce the likelihood of a disastrous financial collapse. Because in our model, as in the data, credit booms could be good as well as bad, regulators face a tradeoff between reducing the likelihood of crisis versus stifling a good credit boom. We consider a simple regulatory policy that allows for a countercyclical capital buffer. We then solve for the parameters of the rule that maximize welfare. We find that the regulatory policy indeed improves welfare mainly by reducing the frequency of costly financial panics. Further, the countercyclical buffer is important. Not relaxing capital requirements in a crisis has the effect of amplifying the downturn, thus reducing welfare.

There are several directions for new research. Limits on banks' ability to raise equity capital plays a key role. It constrains their ability to raise funds and opens up the possibility that they can become vulnerable to panics. We relied on a reduced form function to capture costs of capital injections that was consistent with the evidence on new equity issuance. However, a deeper understanding of these costs would be desirable. Similarly, that banks rely heavily on short term non-contingent debt plays a key role in making them occasionally susceptible to panics. A deeper treatment of this issue is also in order. Finally, our model blurs the distinction between commercial and shadow banks. Of course, any regulation of commercial banks will affect the allocation of funds between commercial and shadow banks (e.g. Begenau and Landvoigt (2018)). Adding in this consideration is an important topic for future research.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.red.2020.06.004>.

References

Adrian, T., Shin, H.S., 2010. The changing nature of financial intermediation and the financial crisis of 2007–2009. *Annual Review of Economics* 2 (1), 603–618.

Adrian, T., Boyarchenko, N., Giannone, D., 2019. Multimodality in macro-financial dynamics. FRB of New York Staff Report (903).

Angeloni, I., Faia, E., 2013. Capital regulation and monetary policy with fragile banks. *Journal of Monetary Economics* 60 (3), 311–324.

Begenau, J., Landvoigt, T., 2018. Financial regulation in a quantitative model of the modern banking system. Available at SSRN 2748206.

Benhabib, J., Wang, P., 2013. Financial constraints, endogenous markups, and self-fulfilling equilibria. *Journal of Monetary Economics* 60 (7), 789–805.

Benigno, G., Chen, H., Otrok, C., Rebucci, A., Young, E.R., 2013. Financial crises and macro-prudential policies. *Journal of International Economics* 89 (2), 453–470.

Bernanke, B., 2018. The real effects of the financial crisis. *Brookings Papers on Economic Activity* 20.

Bernanke, B., Gertler, M., Gilchrist, S., 1998. The financial accelerator in a quantitative business cycle framework. Technical report. National Bureau of Economic Research.

Bernanke, B.S., et al., 2010. Causes of the recent financial and economic crisis: testimony before the financial crisis inquiry commission, September 2, 2010. Technical report. Board of Governors of the Federal Reserve System (US).

Bianchi, J., Mendoza, E.G., 2018. Optimal time-consistent macroprudential policy. *Journal of Political Economy* 126 (2), 588–634.

Boccola, L., Lorenzoni, G., 2017. A model of financial crises in open economies. Technical report. Working Paper.

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

Bordalo, P., Gennaioli, N., Shleifer, A., 2018. Diagnostic expectations and credit cycles. *The Journal of Finance* 73 (1), 199–227.

Boz, E., Mendoza, E.G., 2014. Financial innovation, the discovery of risk, and the US credit crisis. *Journal of Monetary Economics* 62, 1–22.

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

Calomiris, C.W., Kahn, C.M., 1991. The role of demandable debt in structuring optimal banking arrangements. *The American Economic Review*, 497–513.

Calvo, G.A., 1988. Servicing the public debt: the role of expectations. *The American Economic Review*, 647–661.

Chari, V.V., Kehoe, P.J., 2016. Bailouts, time inconsistency, and optimal regulation: a macroeconomic view. *The American Economic Review* 106 (9), 2458–2493.

Cogley, T., Sargent, T.J., 2009. Diverse beliefs, survival and the market price of risk. *The Economic Journal* 119 (536), 354–376.

Cole, H.L., Kehoe, T.J., 2000. Self-fulfilling debt crises. *The Review of Economic Studies* 67 (1), 91–116.

Cooper, R., Ross, T.W., 1998. Bank runs: liquidity costs and investment distortions. *Journal of Monetary Economics* 41 (1), 27–38.

Diamond, D.W., Dybvig, P.H., 1983. Bank runs, deposit insurance, and liquidity. *Journal of Political Economy* 91 (3), 401–419.

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

Farhi, E., Maggi, M., 2018. A model of the international monetary system. *The Quarterly Journal of Economics* 133 (1), 295–355.

Faria-e Castro, M., 2019. A quantitative analysis of countercyclical capital buffers. FRB St. Louis Working Paper (2019-8).

Geanakoplos, J., 2010. The leverage cycle. *NBER Macroeconomics Annual* 2009 24.

Gertler, M., Karadi, P., 2011. A model of unconventional monetary policy. *Journal of Monetary Economics* 58 (1), 17–34.

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

Gertler, M., Kiyotaki, N., 2015. Banking, liquidity, and bank runs in an infinite horizon economy. *The American Economic Review* 105 (7), 2011–2043.

Gertler, M., Kiyotaki, N., Queralto, A., 2012. Financial crises, bank risk exposure and government financial policy. *Journal of Monetary Economics* 59, S17–S34.

Gertler, M., Kiyotaki, N., Prestipino, A., 2016. Wholesale banking and bank runs in macroeconomic modeling of financial crises. In: *Handbook of Macroeconomics*, vol. 2. Elsevier, pp. 1345–1425.

Gertler, M., Kiyotaki, N., Prestipino, A., 2020a. Banking panics as endogenous disasters and the welfare gains from macroprudential policy. *The American Economic Review: Papers and Proceedings* 110, 463–469.

Gertler, M., Kiyotaki, N., Prestipino, A., 2020b. A macroeconomic model with financial panics. *The Review of Economic Studies* 87 (1), 240–288.

Gorton, G., Ordóñez, G., 2019. Good booms, bad booms. *Journal of the European Economic Association* jyv058.

Gorton, G.B., 2010. *Slapped by the Invisible Hand: The Panic of 2007*. Oxford University Press.

He, Z., Krishnamurthy, A., 2019. A macroeconomic framework for quantifying systemic risk. *American Economic Journal: Macroeconomics* 11 (4), 1–37.

He, Z., Khang, I.G., Krishnamurthy, A., 2010. Balance sheet adjustments during the 2008 crisis. *IMF Economic Reviews* 58 (1), 118–156.

He, Z., Kelly, B., Manela, A., 2017. Intermediary asset pricing: new evidence from many asset classes. *Journal of Financial Economics* 126 (1), 1–35.

Jeanne, O., Korinek, A., 2019. Managing credit booms and busts: a pigouvian taxation approach. *Journal of Monetary Economics* 107, 2–17.

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

Jordà, Ò., Schularick, M., Taylor, A.M., 2011. Financial crises, credit booms, and external imbalances: 140 years of lessons. *IMF Economic Reviews* 59 (2), 340–378.

Kashyap, A.K., Stein, J.C., et al., 2004. Cyclical implications of the Basel II capital standards. *FRB Chicago Economic Perspectives* 28 (1), 18–33.

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

Krishnamurthy, A., Muir, T., 2017. How credit cycles across a financial crisis. Mimeo.

Lorenzoni, G., 2008. Inefficient credit booms. *The Review of Economic Studies* 75 (3), 809–833.

Martin, A., Skeie, D., Thadden, E.-L.v., 2014. Repo runs. *The Review of Financial Studies* 27 (4), 957–989.

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

Perri, F., Quadrini, V., 2018. International recessions. *The American Economic Review* 108 (4–5), 935–984.

Philippon, T., 2015. Has the US finance industry become less efficient? On the theory and measurement of financial intermediation. *The American Economic Review* 105 (4), 1408–1438.

Reinhart, C.M., Rogoff, K.S., 2009. This Time Is Different: Eight Centuries of Financial Folly. Princeton University Press.

Robatto, R., 2019. Systemic banking panics, liquidity risk, and monetary policy. *Review of Economic Dynamics* 34, 20–42.

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

Uhlig, H., 2010. A model of a systemic bank run. *Journal of Monetary Economics* 57 (1), 78–96.

Van der Ghote, A., 2018. Coordinating monetary and financial regulatory policies. Technical report. ECB Working Paper.