Capital Requirements in a Model of Bank Runs: The 2008 Run on Repo

Luis A. Simon*
University of Chicago

This version: March 23, 2020
For latest version click here.

Abstract

Capital requirements involve a trade-off between financial intermediation and financial stability. I analyze this trade-off in a macroeconomic model that allows for systemic bank runs, à la Gertler and Kiyotaki (2015). I show that fixed capital requirements make the economy more prone to runs because they slow down the recovery and reduce welfare compared to the laissez-faire benchmark. On the other hand, appropriately chosen countercyclical capital requirements can increase both financial stability and welfare. To weigh the costs and benefits of this policy, I estimate the probability of a systemic shock to the financial sector from CDS data and find it to be around 0.5% per year prior to the 2007-09 financial crisis. I then show that implementing a countercyclical capital requirement that would have prevented the run on repo markets in 2008 would have cost 3% of steady state bank capital and less than 0.1% in consumption terms.

*Kenneth C. Griffin Department of Economics, University of Chicago. I am deeply grateful to my advisors Harald Uhlig, Veronica Guerrieri, Kinda Hachem, and Randall Kroszner for their guidance and advice.
Email: simonl@uchicago.edu.

1 Introduction

How different would the 2008 Financial Crisis have been if the right policies had been in place? In this paper, I answer this question for one particular market, the repurchase agreement (repo) market. The repo market, which suffered a run in 2008, is a significant source of funding for financial institutions, including, in particular, investment banks. How would the run on repo have been different if the correct capital requirements had been in place? Or equivalently, how can we design policy to minimize the likelihood and severity of a similar run in the future? In this paper, I show that fixed capital requirements would not have prevented the run on repo markets in 2008; indeed they would have actually made the recovery longer. On the other hand, appropriately chosen countercyclical capital requirements could have prevented the run at a small cost to investors and banks.

To explore the role of capital requirements, I use a model of investors who provide short-term funding to banks. Banks use capital more efficiently than investors but they are subject to runs. A run happens when investors decide to withdraw their funding from a bank and the bank cannot repay. A bank run arises because of an illiquidity problem: a bank's assets may be enough to pay back depositors conditional on enough people rolling over their deposits for another period. Because of the illiquidity of assets, if enough investors try to cash out, the price of the banks assets will fall as they try to sell them, causing the bank to become insolvent. In this scenario, fixed capital requirements actually make a run *more* likely because they constrain the bank's ability to lever up, thus preventing the financial sector from recovering quickly after a run. The intuition behind this result is that investors understand that, in the case of a run, banks will face a binding capital constraint and will not be able to recover quickly; hence the crisis will be longer lived. A longer-lived crisis makes investors more likely to run because their likelihood of running depends on how serious the crisis will be. On the other hand, countercyclical capital requirements that allow for higher leverage after a run can increase both welfare and financial stability; they can even get rid of runs altogether. In this case, investors understand that if a crisis happens, it will be short lived because banks will lever up and recover quickly. Notably, welfare and financial stability need not go in the same direction. Lowering leverage to make the financial sector more stable has an associated efficiency cost, and policymakers need to evaluate the trade-off. Does the decrease in the time the economy spends in the crisis states compensate for the fall in consumption during good times? The calibration exercise shows that in the case of the 2008 run on repo, this cost was very small: less than 0.1% of steady state consumption. Hence, regulation to prevent the run would also have increased average utility of investors. Furthermore, the cost in terms of bank capital would also have been relatively small, 3% of bank capital in steady state. Key to weighing the costs and benefits of regulation is understanding how common these crises are. To do this, I estimate the probability of a systemic financial crisis from credit default swaps (CDS) data and find that during good economic times it is on the order of 0.5%.

Since 2008, a booming literature has investigated the cause of the 2008 Financial Crisis and incorporated the shortcomings of the models used by economists on the Crisis' buildup. Much progress has been made in incorporating a richer financial sector into macroeconomic models that are able to explain what we observed during the crisis and the slow recovery that followed. Most of these models have built on the financial accelerator models of Bernanke et al. (1999) and Kiyotaki and Moore (1997). On the regulatory side, the issue of what policies to implement and, in particular, what macroprudential policies should be used to control aggregate risk in the financial sector, have received a great deal of attention. Extensive research has examined capital requirements but only recently have countercyclical requirements been emphasized, as in, for example, Faria-e Castro (2019) and Adrian and Liang (2016).

Countercyclical capital requirements are regularly considered by policymakers. They were introduced in Basel III regulation[1] and added as a policy tool by the Fed with the introduction of the "Countercyclical Capital Buffer (CCyB)" in 2016[2].

[1] In Basel III, the countercyclical capital buffer is introduced as a tool to curb excessive credit growth and, hence, to prevent the buildup of systemic risk. The buffer goes from 0% to 2.5% of risk weighted assets (Bank for International Settlments, 2010).

[2] The policy, announced on September 08, 2016, applies to "banking organizations that are

However, the Fed has not used this tool yet, and it has kept CCyB at 0% since its inception. The rationale behind CCyB is that it should be "turned on" while the economy is still in expansion and financial instabilities are building up. Should a downturn occur, the CCyB is turned back to zero. In this paper, I argue rather than increasing capital requirements before a crisis, they should be decreased after one happens. I show that a credible commitment to this policy is enough to lower the probability of a crisis. Jerome Powell recently said in a press conference: "...I view the level of capital requirements and the level of capital in the system as being about right" Powell (2019). If that is the case, then there is no need to increase the requirements now only to bring them back down if a downturn occurs. Furthermore, large financial institutions participating in the repo market in 2008 were not regulated by the Fed, only by the Securities and Exchange Commission (SEC). In practice, this meant that they were subject to laxer regulations and engaged in what is called *regulatory arbitrage*; wherein off-balance sheet vehicles that disguised these institutions' actual risk and leverage were created; although efforts have been done to correct this problem. More importantly, measuring financial vulnerabilities is not an easy task and we have generally been slow to recognize them. CCyB implicitly assumes that financial vulnerabilities can be detected far in advance, providing regulators with enough time to increase the capital buffer, and giving financial institutions enough time to build their equity. On the other hand, a rule that lowers capital requirements after a run is easier to explain and less costly to implement.

The collapse of the repo markets can be seen as a bank run in the traditional sense of Diamond and Dybvig (1983). However, the run occurred not among depositors lining up outside a bank but investors participating in the repo market. This point has been discussed before, for example, in Gorton and Metrick (2012). Investors in the repo markets–mostly money market mutual funds–stopped rolling over their short-term repo agreements which caused funding problems for institutions relying in this market for funding. Not a lot of data are available from repo market transactions. Two pieces of evidence that show the run are provided in Figure 1. The first panel

subject to the advanced approaches capital rules, generally those with more than $250 billion in assets or $10 billion in on-balance-sheet foreign exposures" (Federal Reserve Board, 2016).

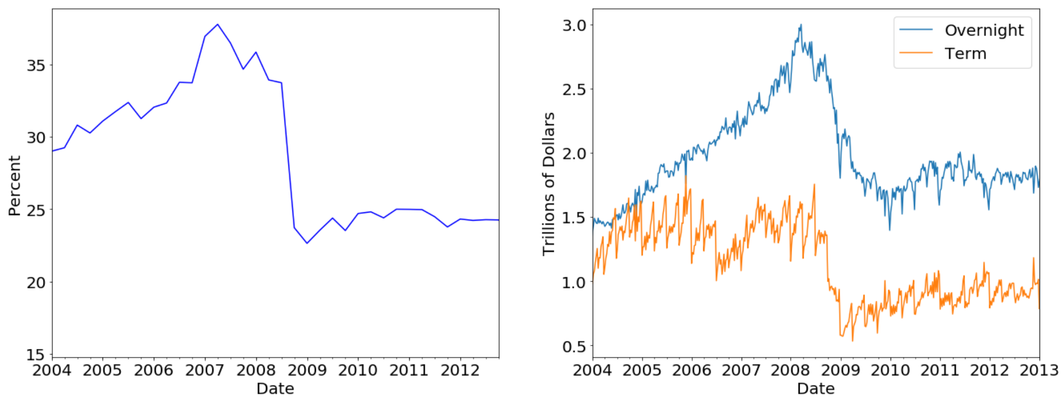

(a) Ratio of Broker-Dealers' Assets to Commercial Banks' Assets.
Source: Fed Flow of Funds.

(b) Weekly repo transactions reported by Primary Dealers.
Source: NY Fed.

Figure 1: Evidence of Repo Market Run in 2008

shows the ratio of broker dealers' assets to commercial banks' assets. The whole financial sector was hit during the crisis but we see that the fall in assets for broker dealers was much larger than the fall suffered by commercial banks; thus the ratio falls dramatically in 2008. The second panel shows transactions reported by primary dealers to the New York Fed. These transactions are a subset of the total which has been estimated to have peaked between 6 and 10 trillion dollars[3]. To get a better sense of the importance of this market, it is worth mentioning that total deposits in the Fed Flow of Funds were around $800 billion dollars in 2008. In the second panel of Figure 1, we again see a large fall in 2008, especially in overnight transactions, which further shows the collapse in short-term funding experienced by financial institutions.

A repo is essentially a short-term collateralized loan. The lender loans funds to the borrower for a specified period of time, most commonly overnight, and the borrower agrees to repay them with interest. At the same time, the borrower puts up collateral for the loan. Treasuries are widely used as collateral but other asset-

[3]Gorton and Metrick (2012) estimate the total size of the repo market to be about $10 trillion, while Copeland et al. (2014) estimate it to be $6.1 trillion.

backed securities such as mortgage-backed securities are also commonly used. For this paper, these are the key elements of the repo market; the collateral used for the loan is going to be a key element in the model. A more comprehensive treatment of the mechanics of the repo market can be found in Copeland et al. (2014) and Wang (2019).

We see the similarities between a traditional bank and other financial institutions when we analyze the flow of funds. I use the term *shadow bank* to refer to financial institutions that are not commercial banks. In particular, I am referring to participants in the repo market, who in this instance, were mostly broker dealers. Figure 2, adapted from Gorton and Metrick (2012), shows the flow of funds in a traditional versus a shadow bank. Both institutions engage in maturity transformation. They are using short-term funding, whether it's coming from depositors or investors, to fund longer-term commitments. A key difference however, is that shadow bank accounts are not insured. Instead, the shadow bank puts up collateral and engages in short-term repo transactions. Given that the basic structure of both types of financial institutions is similar, we can use the same tools to study both. In particular, both types of institutions face the same maturity mismatch problem that can give rise to runs.

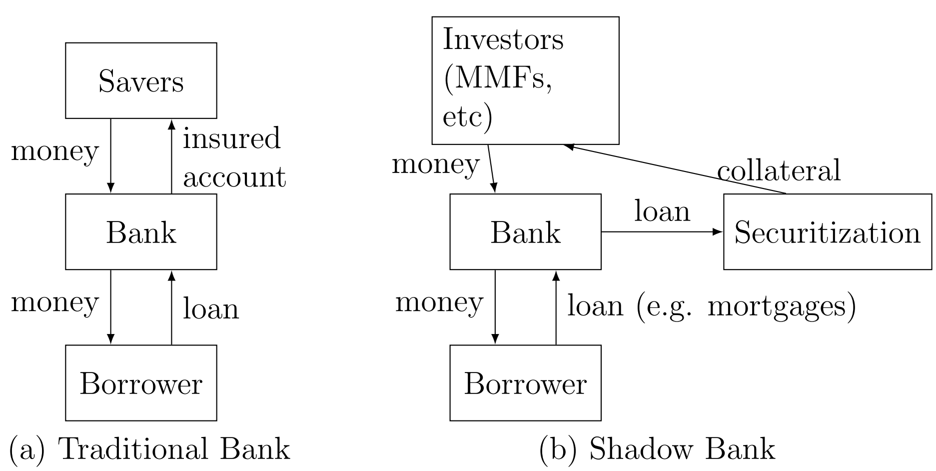

Figure 2: Structure of Traditional vs Shadow Bank.

I explore the effects of policies that could be implemented to prevent future runs, or analogously, to analyze counterfactual scenarios of how the crisis could have unfolded. In particular, to analyze the effect of a leverage constraint, I extend the framework of Gertler and Kiyotaki (2015)–henceforth GK15–and show the benefits of countercyclical policies. To weigh the costs and benefits of policy, policymakers need to understand the likelihood that a crisis will occur. During normal times, they put into place regulations that restrict financial intermediation in the hope that the regulation will reduce the likelihood and severity of crisis states. I use a new measure of systemic risk extracted from CDS that captures the market-implied probability of a financial crisis during the next year. I estimate this probability to be around 0.5% during good economic times. Using this estimate, I make welfare comparisons under counterfactual policies.

The paper is organized as follows. Section 2 reviews the related literature and explains this paper's contributions. Section 3 describes the model. Section 4 shows the calibration and explains how the systemic run probability is calculated. Section 5 describes the policy experiments and compares them to the unconstrained economy. Finally, section 6 concludes.

2 Literature Review

This paper extends the recent literature on macro models with bank runs. Building on the framework of Gertler and Kiyotaki (2015), I add policy tools and show how to pin down the run probability from a microfounded problem. Gertler et al. (2016) analyze the financial crisis through this lens but focus on understanding the effects of having two types of intermediaries: retail and wholesale banks. Faria-e Castro (2019) is also analyzing countercyclical capital requirements but his paper does a positive assessment of the CCyB by the Fed. I complement his analysis by providing a market measure of default probabilities and a microfoundation of the default probability function. Furthermore, Faria-e Castro (2019) recognizes the difficulties in implementing a rule dependent on an unobservable and hard-to-measure object such as financial instabilities; I show that we can get the benefits of countercyclical buffers with only a promise of a change in policy after a crisis, thus avoiding the cost banks face of building up capital before a crisis. To the best of my knowledge, Mendo (2019) provides the only other model in this literature that shows the negative effects of fixed capital requirements. He also argues for state-dependent capital requirements, but in his case, capital requirements are used to prevent the buildup of hidden risks.

The literature on macro models with bank runs essentially combines two thus far separate approaches. On the one hand, financial accelerator models such as the canonical Bernanke et al. (1999) and Kiyotaki and Moore (1997), or more recently Brunnermeier and Sannikov (2014) and He and Krishnamurthy (2013). These models exhibit shock amplification through changes in the balance sheet of financial intermediaries. On the other hand, the bank runs literature started by Diamond and Dybvig (1983), which has mostly focused on stylized two- or three-period models. Goldstein and Pauzner (2005) further extended this literature using global games to pin down the probability of a run based on fundamental shocks. This distinction is key if we want to evaluate costs and benefits of various policies; rather than relying on exogenous beliefs to weigh different states of nature, it allows us to objectively compute expectations over these states. I show how to incorporate the global games approach to bank runs, in the spirit of Goldstein and Pauzner (2005), in a simple way to a larger general equilibrium macroeconomic model.

The literature on capital requirements is extensive and will not be fully reviewed here. Earlier papers by Van den Heuvel (2008) and Diamond and Rajan (2000) focus on the effect of fixed capital requirements. In a more recent example, Begenau and Landvoigt (2018) builds a quantitative model emphasizing the role of the shadow banking sector and in their model countercyclical requirements provide no more gains of than fixed ones. A newer and fast-growing line of research focuses on countercyclical capital requirements. For example, Adrian and Liang (2016) explore the interaction of traditional monetary policy with the trade-off introduced by macroprudential policy on financial stability. Di Tella (2017) shows how a time-varying tax on intermediaries can be similarly used to improve efficiency in the economy. However, few models consider the interaction between capital requirements and bank runs.

Two papers that focus on this theme are Faria-e Castro (2019) and Mendo (2019). Faria-e Castro (2019) focuses on evaluating the Fed's CCyB while Mendo (2019) builds a theoretical model that shows how hidden risks can build up when aggressive stabilization policy is used. Empirically testing the costs and benefits of capital policy is a challenging problem and is recently addressed by Koch et al. (2020), who use historical data to show that countercyclical capital requirements help systemically important institutions weather crises but do not prevent the crises from happening and affecting the rest of the financial sector.

This paper is calibrated to the run on repo markets in 2008, and it makes counterfactual predictions using different policies. The run on repo has been studied by Gorton and Metrick (2012), who argue that it was a central component of the crisis. In contrast, Krishnamurthy et al. (2014) argue that while there was a run on repo, the collapse in asset-backed commercial paper markets was a larger funding disruption for intermediaries. Uhlig (2010) shows how the Financial Crisis had the usual effects of a bank run. Finally, Pozsar et al. (2010) describe the structure of the financial sector in 2008 and discuss the importance of the shadow banking sector.

An important part of the calibration is the probability of observing a systemic run when the economy is not in a crisis. The probability and severity of a crisis are crucial objects in welfare comparisons. The literature on systemic risk has on focused estimating two types of conditional effects. First, Adrian and Brunnermeier (2016) calculate the expected value at risk of the system conditional on an institution being in distress. Second, Acharya et al. (2012) and Brownlees and Engle (2016) ask the opposite question: Conditional on the system being in distress, what is the expected effect on a given institution? However, my model involves an unconditional probability: What is the probability that the system will be in distress, regardless of what happens to any one institution? To address this question, I propose a new method for estimating this probability using a Kalman smoother and CDS data. I propose a simple state space model that allows me to separate default probability into an idiosyncratic and an aggregate component.

3 Model

The economy is populated by a continuum with a unit measure of investors and banks. Investors decide whether to make deposits at a bank for a known return or hold capital directly. However, investors are less efficient than banks at managing capital. Absent any other frictions, this inefficiency implies that banks would optimally hold all of the capital. The friction that prevents banks from intermediating all of the capital is moral hazard on the bank side, which constrains their deposit-taking ability.

Investors are infinitely lived and decide how to allocate their resources between consumption, production, and making deposits at a bank. Banks, which are owned by investors, provide deposit services but may suffer runs. Each investor owns a diversified portfolio of banks so all investors receive the same dividend payment at each point in time. This implies that investors are subject only to aggregate risk and not to idiosyncratic bank risk.

There is one unit of capital that does not depreciate and is used to produce the consumption good with linear production technology. The inefficiency of investors to manage capital directly is captured by a convex cost of holding capital. Banks do not face this cost, which is why it is efficient for banks to hold all of the capital stock.

A key feature of the model is the existence of bank runs that happen with an endogenously determined probability. This probability is microfounded using a currency attack model borrowed from the global games literature. In a run, a bank is not able to meet its obligations with depositors, and so it is liquidated and the proceeds are distributed among the depositors. The reason a bank can suffer a run is liquidity mismatch. Banks issue non-contingent deposits to buy capital, but capital may not be perfectly liquid at all times, and it may need to be liquidated at a discount to service depositors. A full description of precisely how the run probability is determined and what happens in a run is found in subsection 3.4. Note that all banks are identical and runs occur on the whole financial sector. Next, I explain in detail each section of the model.

3.1 Investors

Investors are infinitely lived and maximize their expected lifetime utility. They have logarithmic utility and can operate capital to produce goods but do so less efficiently than banks. The function $f(\cdot)$ is the cost they need to pay to operate capital; banks have no such cost. f is increasing and convex. On the production side, investors also receive an endowment every period, W^i.

The representative investor's problem is to maximize his discounted flow of utility subject to his budget constraint. That is,

$$\max_{c_t, d_t, k_t^i} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \log c_t \tag{1}$$

$$\text{s.t.}$$

$$[\text{if no run}]: \; c_t + d_t + q_t k_t^i + f(k_t^i) = W^i + R_t d_{t-1} + (A + q_t) k_{t-1}^i + \Pi_t \tag{2}$$

$$[\text{if run}]: \; c_t^{run} + q_t^{run} k_t^{run,i} + f(k_t^{run,i}) = W^i + x_t R_t d_{t-1} + (A + q_t^{run}) k_{t-1}^i \tag{3}$$

where c_t denotes consumption, q_t is the price of capital, d_t are deposits, R_t is the return on deposits, k_t^i is capital held by investors, A is productivity. Finally, x_t is the recovery rate from depositors in case of a run; it is an endogenous object and will be defined later. Π_t are the cumulative profits from banks that are transferred as a lump sum to the investors if the banks remain solvent. I denote run state variables with the superscript *run*.

Let p_t denote the probability of a run tomorrow and Λ_{t+1} the stochastic discount factor from t to $t+1$. From the investor's problem we can derive the following optimality conditions:

$$1 = (1 - p_t)\Lambda_{t+1} R_t + p_t \Lambda_{t+1}^{run} x_{t+1} R_t \tag{4}$$

$$1 = (1 - p_t)\Lambda_{t+1} \frac{A + q_{t+1}}{q_t + f'(k_t^i)} + p_t \Lambda_{t+1}^{run} \frac{A_{t+1} + q_{t+1}^{run}}{q_t + f'(k_t^i)} \tag{5}$$

These two conditions are the usual asset pricing equations for the return on deposits and capital, respectively. Next, I characterize the banks' problem.

3.2 Banks

Banks are more efficient than investors at using capital; thus in a model without frictions, banks would hold all the capital. However, I assume that there is an agency problem and that banks can run away with a fraction θ of their total assets. Investors internalize this and so will only be willing to deposit up to the point that banks satisfy an incentive compatibility constraint so that it is optimal for the bank to stay in business. This assumption is present in GK15 and in macrofinance literature more generally. It is sometimes referred to as a 'skin-in-the-game' constraint. See, for example, Holmstrom and Tirole (1997) and Kiyotaki and Moore (1997), and more recently, He and Krishnamurthy (2013). He and Krishnamurthy (2011) show how this type of contract arises out of a moral hazard problem.

Banks have an incentive to save until their IC constraint is not binding–that is, until they have enough net worth to take on all the deposits the investors want to provide. To prevent this, I assume that a fraction $1 - \sigma$ of banks exogenously exit each period. To keep the problem stationary, they are replaced by an equal fraction of new banks which get a one-time endowment of W^b to start operating.

Banks are owned by investors, and so they use the investor's stochastic discount factor to discount profits. Gertler et al. (2016) show that as long as there is a positive probability of binding frictions in the future, it is optimal for banks to retain profits until they exit. Hence, we do not need to keep track of interim dividend payments and assume banks only increase their net worth as long as they operate. Intuitively, the fact that we have financial frictions that prevent banks from intermediating more capital even though they are more efficient makes Tobin's Q greater than one.

Banks are also subject to capital requirements in the form of a leverage constraint. The constraint may or may not be binding at different points in time. A bank with net worth n_t faces the following recursive problem:

$$
\begin{aligned}
V_t(n_t) = \max_{k_t^b, d_t} \; & \mathbb{E}_t\left[\Lambda_{t+1}((1-\sigma)n_{t+1} + \sigma V_{t+1}(n_{t+1}))\right] && (6)\\
\text{s.t.} & \\
q_t k_t^b &= n_t + d_t && (7)\\
n_{t+1} &= (A + q_{t+1})k_t^b - R_{t+1}d_t && (8)\\
\theta q_t K_t^b &\leq V_t && (9)\\
\frac{q_t k_t^b}{n_t} &\leq \bar{\phi}_t && (10)
\end{aligned}
$$

The first constraint is simply the accounting identity of the bank's balances sheet relating assets to equity and liabilities. The second one is the equation for the evolution of net worth if there is no run. Next is the incentive compatibility constraint and finally the capital requirement equation. Note that capital requirements are time-varying and they could potentially depend explicitly on other variables in the economy. The net worth of banks grows only in the case of no run. In the case of a run, the bank's net worth goes to 0.

The last two constraints can be rewritten as a single nonlinear one:

$$
\frac{q_t k_t^b}{n_t} \leq \min\left\{\frac{V_t}{\theta n_t}, \bar{\phi}_t\right\} \qquad (11)
$$

As shown by GK15, the constraint will be binding. Intuitively, this is because the return on assets minus the payments to depositors is positive, and so banks want to issue as many deposits as possible.

Note that the bank's problem is linear in net worth and so aggregation is straightforward. Thus we can redefine the problem in terms of choosing leverage as opposed to deposits and capital. This is true as long as the leverage constraint is homogeneous of degree 0 in net worth. Hence, from now on I simply work with the representative bank. Define the value function per unit of capital and leverage, respectively, as:

$$\psi_t \equiv \frac{V_t(n_t)}{n_t} \tag{12}$$

$$\phi \equiv \frac{q_t k_t^b}{n_t} \tag{13}$$

Then the bank's problem becomes

$$\psi_t = \max_{\phi_t} \mathbb{E}_t \left[\Lambda_{t+1}((1-\sigma) + \sigma \psi_{t+1} \frac{n_{t+1}}{n_t} \right] \tag{14}$$

s.t.

$$\frac{n_{t+1}}{n_t} = \phi_t \frac{A_t + q_t}{q_t} - (\phi_t - 1)\bar{R}_t \tag{15}$$

$$\phi_t = \min \left\{ \frac{\psi_t}{\theta}, \bar{\phi}_t \right\} \tag{16}$$

After characterizing the problem of investors and banks, I now proceed to describe the rest of the model's elements.

3.3 Systemic Run Probability

Bank runs, as in Diamond and Dybvig (1983), are a sunspot equilibrium. Runs only depend on the agents' beliefs about the actions of others, but this raises the issue of how to agents form those beliefs. To overcome this issue, Goldstein and Pauzner (2005) use global games and show how to exactly pin down the probability of a bank run when there is fundamental uncertainty. Both of these models, along with most of the literature on bank runs, use very stylized two- or three-period models. The work of Goldstein and Pauzner (2005) cannot be incorporated in a straightforward manner into a larger general equilibrium macroeconomic framework. However, to determine the probability of coordinating on running, I rely on the global games literature, particularly the literature on speculative currency attacks.

Conceptually, a currency attack is similar to a bank run in the sense that if enough investors attack the currency, the currency will depreciate. Likewise, if enough depositors withdraw their money from the bank, the bank will become insolvent. The model for the runs is motivated by a canonical currency attack model from Chamley (2004).

Assume that at the beginning of each period, investors face the following problem: they must decide whether to run or not. If they run, they need to pay a small cost $\gamma > 0$ and they get back the known return $\bar{\kappa}R$. However, if they do not run and there is a run, then they do not get any of their deposits back. Let κ denote the maximum fraction of deposits that a bank can service and δ the fraction of depositors who withdraw. κ is a random variable with mean $\bar{\kappa}$ and is defined below. Under these assumptions, the payoff matrix of an investor is:

	Aggregate State		
		run	no run
Agent	run	$\bar{\kappa}R - \gamma$	$R - \gamma$
	no run	0	R

The net payoffs of running are:

$$u(\kappa, \delta) = \begin{cases} \bar{\kappa}R - \gamma, & \delta \geq \kappa \\ -\gamma, & \delta < \kappa \end{cases} \tag{17}$$

where $\bar{\kappa}R$ corresponds to the recovery on deposits mentioned earlier in the model. The key thing to note here is that payoffs do not depend on the fraction of agents who are withdrawing. I do not require that the liquidation value of banks equals the funds that the investors are receiving. However, in equilibrium it will indeed be the case that I clear this market. This assumption allows me to overcome the nonmonotonicity in payoffs that Goldstein and Pauzner (2005) solve in their model and that makes the global games approach to bank runs more difficult than in other applications.

As I mentioned above, to exactly pin down the run probability we need to have fundamental uncertainty. Hence, assume that κ is a random variable that can be observed with noise by the agents. In particular, assume that

$$\kappa \sim \mathcal{N}(\bar{\kappa}, \sigma_\kappa^2) \tag{18}$$

and each agent i observes a noisy signal s_i such that

$$s_i = \kappa + \varepsilon_i; \quad \varepsilon_i \overset{iid}{\sim} \mathcal{N}(0, \sigma_\varepsilon^2) \tag{19}$$

where the noise in the signals is uncorrelated across agents. With these assumptions, I can find the optimal strategy and pin down the probability of a run. Note that κ is a fraction but I draw it from a distribution with unbounded support. What I am actually doing is using $\max\{\min\{\kappa, 1\}, 0\}$ to ensure that κ is between 0 and 1. Without loss of generality, we can focus on monotone strategies[4]. Slightly abusing notation, let s also denote the cutoff strategy where an investor runs if his signal is below s and does not run if it is above s. The result of this game is summarized in the following theorem:

Theorem 1 (Chamley (2004)) *When $\sigma_\varepsilon \to 0$, the optimal strategy, $s^* \to 1/2$, and the probability of a run is given by:*

$$p = \Phi\left(\frac{1}{2}; \bar{\kappa}, \sigma_\kappa\right)$$

where $\Phi(\cdot; \bar{\kappa}, \sigma_\kappa)$ is the cumulative distribution function of the normal distribution with mean $\bar{\kappa}$ and variance σ_κ^2.

Proof. In the Appendix. ■

As I mentioned before, runs in my model are a sunspot equilibrium and there is no fundamental uncertainty. However, the probability that agents coordinate on the run equilibrium will be derived as if there were. In particular, the probability of a run tomorrow in the model is

[4]The result is shown in Chamley (2004).

$$p_t = \Phi\left(\frac{1}{2}; x_{t+1}, \sigma_\kappa\right) \tag{20}$$

Note that in this specification the probability of a run tomorrow depends negatively on the recovery rate. This result is consistent with the complementarity in the action of the agents found in the global games literature: the larger the losses an investor faces in case of a run, the more likely he is to run.

3.4 Equilibrium

To close the model we need market clearing equations. Market clearing for capital is

$$k_t^i + k_t^b = 1 \tag{21}$$

and the economy-wide resource constraint is

$$c_t + f(k_t^i) = A + W^i + (1-\sigma)W^b \tag{22}$$

Finally, we need to define the recovery rate. For the recovery rate, we assume that if there is a bank run, all of the assets are liquidated and then proportionally distributed among depositors. That is:

$$x_{t+1} = \frac{(A_{t+1} + q_{t+1}^{run})k_t^b}{R_t d_t} \tag{23}$$

where the numerator has the value of the bank's total assets in the run state and the denominator has the total deposits owed by the bank.

In this model, after a run banks will be closed for one period. That is, the run is systemic and not even new banks can operate. Because there are no banks, the economy "resets" after a run; after every run we have a period when no banks operate and investors hold all of the capital. Given that no banks operate for one period, investors have to hold all of the capital for one period. This makes the economy look exactly the same after a run, regardless of how much capital the financial sector had prior to the run. Over time, new banks enter and the financial sector starts growing again. Hence, the state variable in this model is t, or the time elapsed since the last run. Given that, we can define the equilibrium as follows:

Definition 2 (Competitive Equilibrium) *A competitive equilibrium is a set of prices,* $\{q_t, R_t\}_{t\geq 0}$*, quantities,* $\{c_t, d_t, n_t, k_t^i, k_t^b\}_{t\geq 0}$*, recovery rates* $\{x_{t+1}\}_{t\geq 0}$*, and default probabilities* $\{p\}_{t\geq 0}$*, such that:*

- *Given prices, run probabilities, and recovery rates, investors optimize their consumption and saving decisions. That is,* (4) *and* (5) *hold.*
- *Given prices, investors optimize their leverage decision. That is,* (14)-(16) *hold.*
- *Run probability,* p_t*, is given by* (20) *and recovery rate,* x_{t+1}*, by* (23).
- *Markets clear. That is,* (21) *and* (22) *hold.*

4 Calibration

As I have mentioned before, default probability is a key object in the model. To estimate it, I use a Kalman smoother described in the next subsection. The systemic run probability will be one of the moments targeted in the calibration. The rest of the targeted moments are discussed later.

I calibrate the unconstrained economy and then use the same calibration when using the different policy rules. While the repo market was not totally unconstrained, participants in this market regularly engaged in regulatory arbitrage and would move risk off their balance sheet making them seem more stable than they actually were. As shown in GK15, the unconstrained economy is a good approximation of the repo markets prior to the Financial Crisis.

4.1 Systemic Run Probability

There have been several papers that measure systemic risk in the financial sector, especially after the Financial Crisis. The leading approaches to quantifying systemic risk are CoVaR by Adrian and Brunnermeier (2016), SRISK by Brownlees and Engle (2016), and marginal expected shortfall (MES) by Acharya et al. (2012). However, these measure are all trying to identify the systemic importance of a single institution. In other words, they analyze different forms of conditioning. CoVaR, for example, measures the system's value at risk conditional on an institution being in distress. SRISK and MES reverse the conditioning by measuring the expected capital an institution would need conditional on the whole system being in distress. For this paper, what I need is an unconditional probability of a crisis; we are analyzing not the relative importance of one institution versus another but the fragility of the whole system.

In order to measure the systemic default probability, I use the following state space model:

$$\mathbf{p_t} = \mathbf{1}\pi_t + \varepsilon_\mathbf{t} \tag{24}$$

$$\pi_{t+1} = \pi_t + \eta_t \tag{25}$$

where $\mathbf{p_t}$ is the vector of default probabilities for all financial institutions at time t, π_t is the (hidden) systemic default probability, and ε_t, η_t are normally distributed error terms and independent from each other. The model states that the default probability of an institution is the sum of two terms: an aggregate default probability and an idiosyncratic shock.

To obtain an estimate of the run probability for each institution, I use credit default swaps (CDS). The daily pricing data comes from Markit. I use data on all available financial firms that issue dollar-denominated one-year debt in the US and look at the period 2003-2010. I drop institutions missing more than 20% of observations and I end up with 62 institutions. The financial institutions included are listed in Appendix B. To reduce the noise in the estimation, I use weekly averages of the spreads. The dataset also includes reported recovery rates for each institution. Having both the recovery rate and the CDS spread, I can back out default probabilities for each security. To make the observations comparable, I use only no restructuring contracts. Most observations in the dataset are either under no restructuring or modified restructuring, and most issuers trade both instruments trading over my sample period. Figure 3 shows summary statistics of the data and a plots the time series of the quartiles for the default probability.

To compute default probabilities, I start from the following approximation of CDS spreads when the default intensity is assumed to be constant over time:

$$\lambda = \frac{c}{1-x} \tag{26}$$

where λ is the default intensity, c is the CDS spread, and x is the recovery rate. From here, we can derive the default probability as:

$$p_t = e^{\frac{\lambda t}{1-x}} \tag{27}$$

The derivation of the default probability is provided in Appendix B.

I use a Kalman smoother to recover the hidden state. I use an expectation maximization algorithm to estimate the variance of the errors and the initial distribution. The time series of the state is shown in Figure 4. Note that the estimated default probability never goes below 0, even though nothing in the estimation prevents it from doing so. The average of the hidden state from the start of the sample to the first half of 2008 is 0.74%, which is the value we should be targeting in the model. However, during this period, the estimated value of the hidden state had already started its steady increase, which peaked in 2009. The average from the start of the data to before the increase is 0.38%. As a conservative choice, I target a 0.5% steady state default probability in the data.

As robustness checks, I use other variations to the state space model, and the results are very similar. I repeated the same analysis in logarithms so that the default probability cannot be negative by construction. I also used contracts written under modified restructuring (as opposed to no restructuring), and, again, the results

	1 Year Spread (%)	1 Year Default Prob. (%)	Recovery Rate (%)
Mean	2.54	3.11	38.58
Std	10.3	8.33	4.45
Min	0.02	0.03	5.0
25%	0.10	0.16	39.3
50%	0.30	0.49	40.0
75%	1.14	1.83	40.0
Max	294.0	97.8	68.1
N	27634	27603	36613

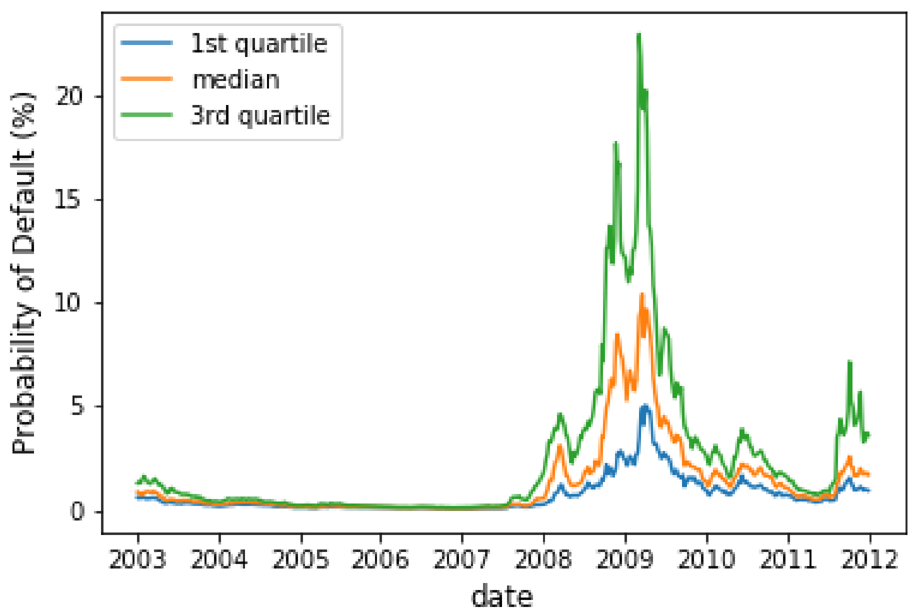

Figure 3: CDS Summary Statistics, 2002-2012. Source: Markit.

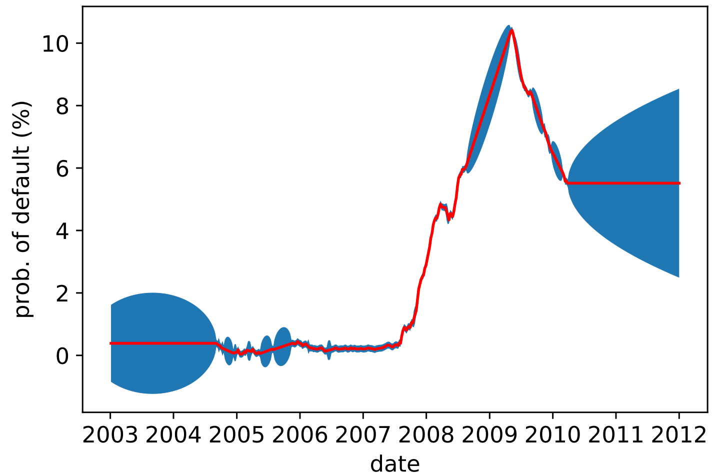

Figure 4: Kalman smoother hidden state with 95% confidence intervals.

do not change significantly. I omit these figures from the paper. Finally, I also do principal component analysis and interpret the first component as the systemic default probability. The first component explains around 65% of the variance in the data. However, the interpretation of the first component from PCA is not as straightforward as the result from the Kalman smoother and we do not have a model to help us interpret what this component is detecting. The first component is also much noisier than the result of the Kalman smoother, and it goes above 1 after 2008. However, the pre-crisis are always between 0.01 and 0.07 which is consistent with the results of the Kalman smoother. The first principal component and the share of variance explained by each component are documented in Appendix C.

It is important to note that from financial markets data we are calculating risk neutral probabilities. To move from risk-neutral probabilities to physical probabilities I use the the state price density implied by the model. A standard result in finance (see, for e.g., Cochrane (2009)) tells us that we can define the state price density for an economy with states, s, by:

$$\pi_t^*(s) = \frac{\Lambda_{t+1}(s)}{\mathbb{E}(\Lambda_{t+1})}\pi_t(s) \tag{28}$$

where $\pi_t(s)$ denotes the physical probability of state s tomorrow, $\pi_t^*(s)$ the risk-neutral probability, and Λ is the stochastic discount factor of investors. Given that the model has only two states, run and no run, we can invert the above equation to find an expression for the physical probability of a run tomorrow, p_t, given our estimate of a risk neutral run probability, p^*:

$$p_t = \frac{p_t^* u'(c^{run})}{p_t^* u'(c^{run}) + (1 - p_t^*)u'(c_{t+1})} \tag{29}$$

This correction turns out to have a very small effect given the rest of the calibration. Using the stochastic discount factor implied by the model, the steady state physical run probability implied by a 0.5% risk-neutral probability is 0.49%.

Now I show the calibration for the rest of the model's parameters.

4.2 Other Parameters

I calibrate the model without policy and then use that calibration to test the counterfactual policies. The model is calibrated to a quarterly frequency, so I use $\beta = 0.99$, which is standard in the literature and implies a risk-free rate of approximately 4% in steady state. Here, due to the run probability the steady state risk free rate is marginally higher. Besides the run probability described in the last section, I target the fall in consumption after a run. I choose 3.0%, which is the year-on-year fall in GDP at the bottom of the crisis, in the second quarter of 2009. It is difficult to tease out exactly how much of the fall in GDP is due to the run on repo markets and how much is due to to other causes, such as the fall in house prices that started in 2007 or the collapse of the asset-backed commercial paper market in 2008. At the same time, the Fed stimulus plus the Treasury's Troubled Asset Relief Program (TARP) had an effect in the opposite direction. I do not take a stance on the relative shares and instead just focus only on matching the bottom of the fall in GDP. I also target a steady state leverage for banks of 14, in line with the estimate of Begenau and Landvoigt (2018) for the shadow banking sector. Finally, I also target a 100 basis point difference in expected annual return from bank capital over the risk-free rate, consistent with Gertler and Kiyotaki (2015). Table 1 shows all of the values for the parameters.

Parameter	Symbol	Value
Investor cost function	$f(k) = \frac{\alpha}{2}k^2$	$\alpha = 0.015$
Share of assets bank diverges	θ	0.2
Bank survival probability	σ	0.95
TFP	A	0.015
Investor endowment	W^i	0.23
Bank endowment	W^b	2.4×10^{-4}
Discount factor	β	0.99
Variance of run probability cdf	σ_κ	0.4

Table 1: Parameters of the Model.

The solution method is similar to GK15 and Christiano et al. (2016), and the reader is referred to them for a detailed description of the numerical procedure. The model has to be solved globally for the run state, the transition, and the steady state because the steady state objects of the model depend on investors beliefs about what happens after a run. I use the term *steady state* loosely; I mean where the economy converges to if a run never happens, even though the run probability is always positive. I assume the economy converges to the steady state after T periods for a large T and then solve jointly for the transition and the steady state. It is important to note that investors have rational expectations about both policy and the transition dynamics after a run.

5 Results

I run a series of policy experiments to demonstrate the mechanisms present in the model. I keep TFP fixed throughout to show only the effect of capital requirements.

First, I introduce a fixed capital requirement. This fixed requirement will not bind in steady state but will only bind after a run. Second, I analyze an economy in which capital requirements depend on the time that has passed since the last run. In this countercyclical policy, the economy will be unconstrained for some time after a run but will have binding requirements in steady state. The experiment assumes that the economy suffers a run in period 0 and plots the recovery back to steady state, assuming that no other run happens, even if the run probability is positive.

5.1 Fixed Capital Requirements

In the first policy experiment I analyze the effect of fixed capital requirements. I use a fixed cap on leverage. Note that in this experiment the rule will not be binding in steady state, yet it will affect steady state values. Agents have rational expectations and bank runs are correctly anticipated; hence the agents' steady state decisions depend on what would happen should the run materialize.

Adding a fixed capital requirement has a counterintuitive effect; runs become more likely and are longer lived. This is so because after a run, banks increase their leverage until their equity recovers, but with the cap in place they take longer to do so. Hence, with the fixed requirement in place, runs are longer lived and consequently, agents have higher probability of running.

Figure 5 shows the path of the economy after it is hit by a run at time $t = 0$ and then does not suffer another run even though the run probability remains positive. Assuming no other run happens, I plot the path of the economy back to steady state. The figure has 4 lines that show the unconstrained economy and three different levels for the leverage cap. The first one is such that steady state leverage falls to 13.3, a value that will be comparable to the experiment with countercyclical requirements. The other values of the leverage cap are such that steady state leverage falls to 10 and 8, respectively. These last two numbers are in line with average values of leverage for commercial banks. Note that the actual cap is significantly above the steady state value. The constraint only binds for a few periods after a run and not in steady state. Nevertheless, the cap is enough to lower steady state leverage.

Investors internalize the fact that if a run happens, banks will be constrained and take a long time to recover, hence, investors lower their deposits even in good times and steady state leverage falls. Note also that the run probability is higher both in steady state and after the run with the policy in place. Even for the looser cap of 10, the run probability in steady state more than doubles. The run probability depends on the severity of the crisis and because the financial sector will take longer to recover, the run probability is higher. Even though the steady state does not change significantly, the economy takes much longer to recover: consumption with the policy in place remains lower at every point in time and the tighter the constraint, the lower consumption will be in the transition. Consumption in the steady state is marginally different and steady state utility does not change significantly.

Clearly, a fixed capital requirement makes this economy strictly worse off: runs are more likely to occur and are longer lived, meanwhile, steady state consumption does not change significantly.

5.2 Countercyclical Capital Requirements

The next experiment involves using a countercyclical capital requirement. Under this policy, leverage is capped in steady state but if there is a run then leverage will be unconstrained for t^* periods. This policy is correctly anticipated by all agents. The exact number of periods after which the constraint is restored is chosen arbitrarily and I use 60 for the results in this section, but the idea is that banks will be allowed to operate unconstrained for a long period of time to allow them to recover their net worth. Results change very little by setting the time after which the policy start binding at 40 or 80 periods and they can be found in Appendix D. The policy rule is:

$$\bar{\phi}_t = \begin{cases} \infty & , t < t^* = 60 \\ \bar{\phi} & , t \geq t^* = 60 \end{cases} \tag{30}$$

Figure 6 shows the unconstrained economy and three different rules. The first rule caps leverage at 13.7, close to but below the steady state value of 13.9. The

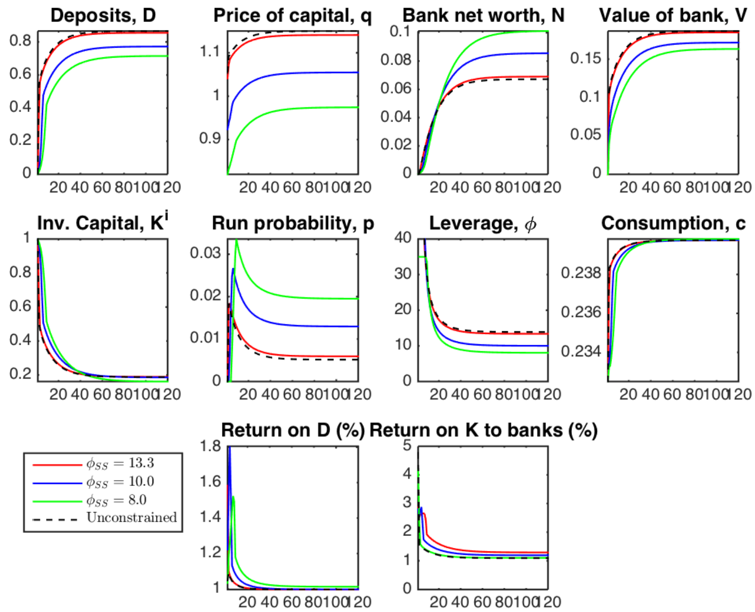

Figure 5: Fixed Capital Requirements for Different Rules vs Unconstrained Economy (GK15).

second rule is set at 13.4, this is the highest value at which the run probability falls to zero. And the third rule is set at 13.0. Note that the run probability is slightly higher with the policy right after the run happens and this continures until the leverage cap is restored. However, the steady state run probability is lower in steady state, and, in two cases, it is exactly zero. That is, the transition shown is actually an off-equilibrium path. The fact that agents believe that the policymaker will remove the leverage constraint in case of a run is enough to get rid of runs completely. In anticipation of the policy, banks increase their leverage a few periods before the policy sets in and will jump discreetly to the constraint in the period it is introduced. This adjustment causes a jump in capital holdings by investors, which, due to the efficiency cost, causes a discrete fall in consumption when the policy sets in. Another effect of the policy is that the price of capital falls less after the run hits; this happens due to the lower run probability in the future, because the price of capital is, as usual, the expected discounted flow of dividends. Under the 13.4 rule, the price of capital falls 6.7% compared to 8.3% in the unconstrained economy. Furthermore, the change in financial intermediation in the steady state is small, steady state bank capital is only 3.0% lower than in the unconstrained economy.

Note that consumption is lower with the countercyclical policy in place. This is the trade-off mentioned before, the leverage constraint decreases financial intermediation which has an efficiency cost and results in lower steady state consumption and utility. Once the run probability hits zero, there are no further gains of decreasing leverage. Runs are not possible anymore and lowering the constraint only decreases the size of the financial sector which makes the economy less efficient. However, in the calibrated model, this fall in consumption is very small. Under the 13.4 rule that makes the run probability fall to zero, the fall in consumption is less than 0.1%.

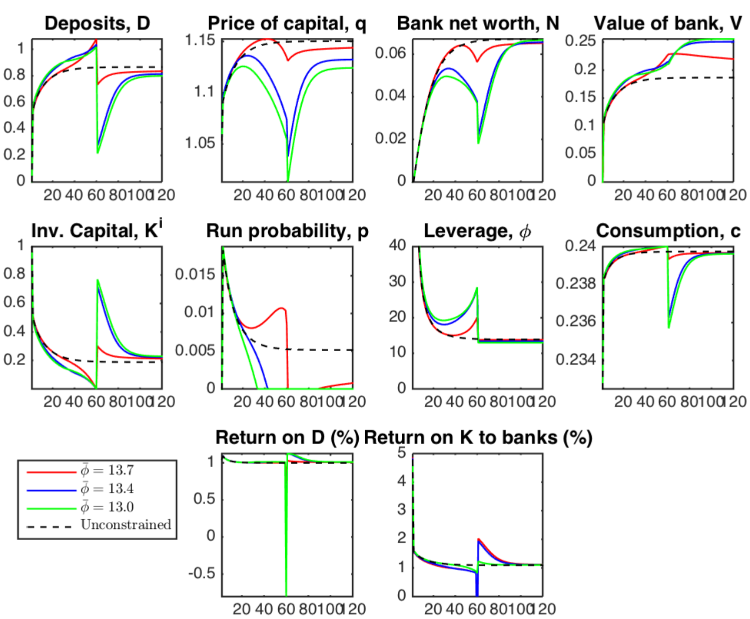

Figure 6: Countercyclical Capital Requirements for Different Rules vs Unconstrained Economy (GK15).

For the next exercise and the welfare comparisons, I will only use the constraint that sets leverage to 13.4 in steady state, the largest one with zero run probability in steady state. Figure 7 shows a different experiment, a simulation of the economy under the policy and under laissez-faire. Note that consumption is constant but lower with the policy in place, since the run probability goes to zero. I numerically calculate utility and it is higher with the policy so agents would indeed prefer that the government implement it. In this case, the countercyclical capital requirements are welfare improving. It is important to note that this is not always the case and the result of the trade-off between lower consumption in steady state and more financial stability need not always yield higher utility with the policy in place. Under different parameter combinations agents may prefer higher consumption in steady state along with higher probability of a crisis.

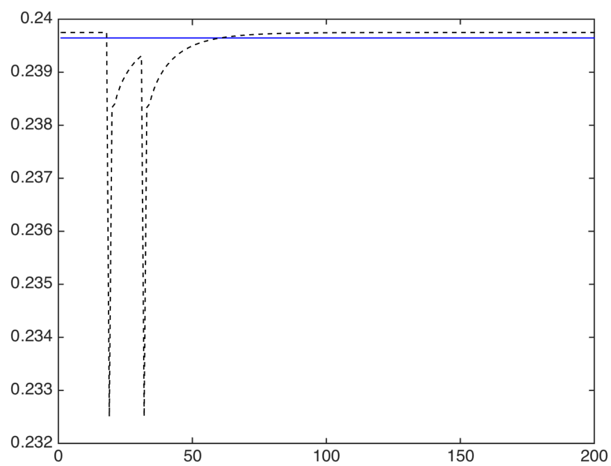

Figure 7: Consumption. With Leverage Constraint, Unconstrained (GK15).

6 Conclusions

In this paper, I study the effect of capital requirements in an economy that is subject to bank runs, and I show that fixed capital requirements decrease welfare because they increase the probability and length of a run. I also show how countercyclical capital requirements can improve financial stability and that a tight enough constraint will eliminate runs. Furthermore, I show that a credible commitment to countercyclical policies can cause significant gains in stability although there is an associated efficiency cost. I calibrate the model using CDS data to estimate the implied default probability of financial institutions. In the calibrated model, the costs associated with a countercyclical policy in 2008 would have been small and increased the average utility of investors. Steady state consumption would have fallen by less than 0.1% and bank capital would have fallen by 3%.

References

Acharya, V., Engle, R., and Richardson, M. (2012). Capital shortfall: A new approach to ranking and regulating systemic risks. *American Economic Review*, 102(3):59–64.

Adrian, T. and Brunnermeier, M. K. (2016). Covar. *American Economic Review*, 106(7):1705–41.

Adrian, T. and Liang, N. (2016). Monetary policy, financial conditions, and financial stability.

Bank for International Settlments (2010). Basel iii: A global regulatory framework for more resilient banks and banking systems. Technical report, Bank for International Settlments.

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

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

Brownlees, C. and Engle, R. F. (2016). Srisk: A conditional capital shortfall measure of systemic risk. *The Review of Financial Studies*, 30(1):48–79.

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

Chamley, C. P. (2004). *Rational herds: Economic models of social learning*. Cambridge University Press.

Christiano, L., Dalgic, H., and Wen, X. (2016). Discussion slides on gertler-kiyotaki: Banking, liquidity and bank runs in an infinite horizon economy aer (2015).

Cochrane, J. H. (2009). *Asset pricing: Revised edition*. Princeton university press.

Copeland, A., Martin, A., and Walker, M. (2014). Repo runs: Evidence from the tri-party repo market. *The Journal of Finance*, 69(6):2343–2380.

Di Tella, S. (2017). Optimal regulation of financial intermediaries. Technical report, National Bureau of Economic Research.

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

Diamond, D. W. and Rajan, R. G. (2000). A theory of bank capital. *The Journal of Finance*, 55(6):2431–2465.

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

Federal Reserve Board (2016). Press release: Federal reserve board approves final policy statement detailing framework for setting countercyclical capital buffer.

Gertler, M. and Kiyotaki, N. (2015). Banking, liquidity, and bank runs in an infinite horizon economy. *American Economic Review*, 105(7):2011–43.

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

Goldstein, I. and Pauzner, A. (2005). Demand–deposit contracts and the probability of bank runs. *the Journal of Finance*, 60(3):1293–1327.

Gorton, G. and Metrick, A. (2012). Securitized banking and the run on repo. *Journal of Financial economics*, 104(3):425–451.

He, Z. and Krishnamurthy, A. (2011). A model of capital and crises. *The Review of Economic Studies*, page rdr036.

He, Z. and Krishnamurthy, A. (2013). Intermediary asset pricing. *The American Economic Review*, 103(2):732–770.

Holmstrom, B. and Tirole, J. (1997). Financial intermediation, loanable funds, and the real sector. *the Quarterly Journal of economics*, 112(3):663–691.

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

Koch, C., Richardson, G., and Van Horn, P. (2020). Countercyclical capital buffers: A cautionary tale. Technical report, National Bureau of Economic Research.

Krishnamurthy, A., Nagel, S., and Orlov, D. (2014). Sizing up repo. *The Journal of Finance*, 69(6):2381–2417.

Mendo, F. (2019). Risk to control risk. Technical report, Technical report, Princeton University.

Powell, J. (2019). Transcript of chair powell's press conference july 31, 2019.

Pozsar, Z., Adrian, T., Ashcraft, A., and Boesky, H. (2010). Shadow banking. *New York*, 458(458):3–9.

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

Van den Heuvel, S. J. (2008). The welfare cost of bank capital requirements. *Journal of Monetary Economics*, 55(2):298–320.

Wang, Y. (2019). Counterparty risk and repo runs.

A Proof of Theorem 1 (Run Probability Equilibrium)

The proof proceeds in steps. I first find the optimal strategy of the game and then proceed to take the limit as the variance of the noise disappears. Given the complementarity in actions, that is, the more people run the more likely the run is to be successful, I consider symmetric, monotone strategies only, and in particular, cutoff strategies.

Slightly abusing notation, let s be the cutoff point of a strategy where an agent runs if his signal is smaller than s and doesn't run otherwise.

Let $F(\cdot)$ denote the c.d.f. of a standard normal distribution, then define

$$\alpha(s) = F\left(\frac{s - \alpha(s)}{\sigma_\varepsilon}\right)$$

We now establish two properties of $\alpha(s)$ that will be used later. First, by symmetry of the normal distribution around 0, note that $\alpha(1/2) = 1/2$. Second, the derivative of $\alpha(s)$ satisfies

$$0 < \alpha'(s) < \frac{f(0)}{f(0) + \sigma_\varepsilon}$$

where $f(\cdot)$ is the p.d.f of a standard normal distribution. To see this, note that the derivative of α is

$$\alpha'(s) = f\left(\frac{s - \alpha(s)}{\sigma_\varepsilon}\right)\frac{1}{\sigma_\varepsilon}\left(1 - \alpha'(s)\right)$$

$$\iff \alpha'(s) = \frac{f\left(\frac{s-\alpha(s)}{\sigma_\varepsilon}\right)}{f\left(\frac{s-\alpha(s)}{\sigma_\varepsilon}\right) + \sigma_\varepsilon}$$

which is clearly greater than 0 since $f > 0$ always. Note that f reaches a maximum at 0 so the result follows.

Lemma 3 *Given a strategy s, the run is successful* $\iff \kappa \leq \alpha(s)$

Proof. ($\Longrightarrow$)First, assume the run is successful. Then, $\kappa \leq \delta$.

Note that $s \mid \kappa \sim \mathcal{N}(\kappa, \sigma_\varepsilon^2)$. And since δ denotes the fraction of agents who withdraw, then

$$\kappa \leq \delta = F\left(\frac{s-\kappa}{\sigma_\varepsilon}\right)$$

and the cutoff is given by

$$\kappa^* \leq \delta = F\left(\frac{s-\kappa^*}{\sigma_\varepsilon}\right)$$

Defining $\alpha(s) \equiv \kappa^*$ the result follows.

($\Longleftarrow$) Assume $\kappa \leq \alpha(s)$. Given s, the fraction who withdraw is

$$F\left(\frac{s-\kappa}{\sigma_\varepsilon}\right)$$

If $\kappa = \alpha(s)$, then the run is successful and we have that

$$\kappa = \alpha(s) = F\left(\frac{s-\alpha(s)}{\sigma_\varepsilon}\right)$$

Now we need to show that it is also true if $\kappa < \alpha(s)$. Note that it is true if $0 < \alpha'(s) < 1$ which was shown before. ■

We also need to establish the posterior distribution of the state κ after the agent observes his signal. For that, notation is simpler if we work with precisions instead of variances. Denote the precision of a random variable y by $\rho_y \equiv 1/\sigma_y^2$.

Lemma 4 *The posterior distribution of κ after an agent receives his signal is $\kappa \mid s \sim \mathcal{N}(m(s), \rho^{-1})$ with*

$$m(s) \equiv \frac{\rho_\kappa}{\rho}\bar{\kappa} + \frac{\rho_\varepsilon}{\rho}s$$

$$\rho \equiv \rho_\kappa + \rho_\varepsilon$$

Proof. We know that

$$s \mid \kappa \sim \mathcal{N}(\kappa, \sigma_\varepsilon^2)$$
$$\kappa \sim \mathcal{N}(\bar{\kappa}, \sigma_\kappa^2)$$

and by Bayes' rule:

$$f(\kappa \mid s) \propto f(s \mid \kappa) f(\kappa)$$

Then,

$$\begin{aligned} f(\kappa \mid s) &\propto \frac{1}{\sqrt{2\pi\sigma_\varepsilon^2}} \exp\left\{-\frac{1}{2}\frac{(s-\kappa)^2}{\sigma_\varepsilon^2}\right\} \frac{1}{\sqrt{2\pi\sigma_\kappa^2}} \exp\left\{-\frac{1}{2}\frac{(\kappa-\bar{\kappa})^2}{\sigma_\kappa^2}\right\} \\ &\propto \exp\left\{-\frac{\rho}{2}(\kappa - m(s))^2\right\} \end{aligned}$$

after some algebra. Note that the last expression is the kernel of a normal distribution with the mean and variance we wanted. ■

Define the net cumulative value function (CVF) as

$$\begin{aligned} V_N(s) &= \mathbb{E}(u(\kappa, s) \mid s) \\ &= \int_{-\infty}^{\infty} u(\kappa, s) f(\kappa \mid s) d\kappa \\ &= \bar{\kappa} R \int_{-\infty}^{\phi(s)} f(\kappa \mid s) d\kappa - \gamma \\ &= \bar{\kappa} R F((\alpha(s) - m(s))\sqrt{\rho}) - \gamma \end{aligned}$$

A strategy s^* is a *strong rational expectations equilibrium (SREE)* as defined by Chamley (2004) if $V_N(s^*) = 0$.

The final part of the proof consists on showing that as $\sigma_\varepsilon \to 0$, $s^* \to 1/2$ and that this is the only solution. To do that, we will show that V_N is a decreasing function and that is equal to 0 at only one point, s^*.

First we show that $V_N(s)$ is a decreasing function of s. Taking the derivative we note that:

$$\begin{aligned}\frac{\partial V_N}{\partial s} &= Af((\alpha(s) - m(s))\sqrt{\rho})\sqrt{\rho}(\alpha'(s) - m'(s)) < 0 \\ &\iff \alpha'(s) - m'(s) < 0\end{aligned}$$

And we also showed before that

$$\alpha'(s) < \frac{f(0)}{f(0) + \sigma_\varepsilon}$$

so it is enough to show that

$$\frac{f(0)}{f(0) + \sigma_\varepsilon} < m'(s)$$

Rearranging the above inequality we see that it holds if and only if the following holds

$$\sigma_\varepsilon < \frac{\sigma_\kappa}{f(0)}$$

So we can take σ_ε small enough to satisfy the above inequality. Hence, if σ_ε is small enough, we are guaranteed to have a decreasing function V_N.

To show $V_N(1/2) = 0$ in the limit, we note that $\lim_{\sigma_\varepsilon \to 0} m(s) = s$. And since 1/2 is a fixed point of α, we have that $\lim_{\sigma_\varepsilon \to 0} \alpha(1/2) - m(1/2) = 0$

Since we know that the difference $\alpha(s) - m(s)$ is decreasing and it is equal to 0 at 1/2 (in the limit), then it must be positive for $s < 1/2$ and negative otherwise. Since $\lim_{\sigma_\varepsilon \to 0} \rho = \infty$, we have that:

$$\lim_{\sigma_\varepsilon \to 0} F(s) = \begin{cases} 1, & s < 1/2 \\ 0, & s > 1/2 \end{cases}$$

So that $F(s)$ is step function crossing the y axis at 1/2. Then, $V_N(1/2) \to 0$ as $\sigma_\varepsilon \to 0$ as was to be shown.

Hence, the ex ante probability of a run is the probability that $\kappa < 1/2$, that is:

$$P(\kappa \leq 1/2) = F\left(\frac{1/2 - \bar{\kappa}}{\sigma_\kappa}\right)$$

as we claimed. ■

B Credit Default Swaps

B.1 List of CDS Issuers Included in the Sample

American Express Company
Amern Express Credit Corporation
American Financial Group
AIG
Aon
Associates Corp of North America
Berkshire Hathaway
Boeing Capital
Allstate
Ambac Financial Group
Ambac Assurance Corporation
Avalon Bay Communities
Caterpillar Financial Services
Citigroup
CNA Financial
Capital One
Chubb
EOP Operating Ltd Partnership
ERP Operating Ltd Partnership
Countrywide Home Loans
Franklin Templeton Investments
GE Capital
Goldman Sachs
GATX
The Hartford
Healthcare Realty Trust
International Lease Finance Corp
JPMorgan Chase & Co
KeyCorp
Kimco Realty
Health Care REIT
Heller Financial
iStar
Legg Mason
MBIA Inc
MBIA Insurance Corp
MGI
Mack-Cali Realty
Marsh & McLennan
Merrill Lynch & Co
MetLife
Liberty Mutual
Lincoln National
Loews Corporation
Morgan Stanley
National Rural Utilities
PMI
ProLogis
Prudential
Radian Group
Radian Asset Assurance
Sallie Mae
Safeco Corp
Charles Schwab
Simon Property Group
Simon Property Group LP
SunTrust Banks
Sears Roebuck Acceptance Corp
US Bancorp
Vornado Realty Trust
Wells Fargo
Philip Morris Capital

B.2 Approximation to CDS Spread

We start from equation (26). Using language and well-known results from survival analysis, we want to find the *lifetime distribution function*, defined as:

$$F(t) = P(t^* \leq t) \tag{31}$$

that is, that the time at which default happens, denoted t^*, is before some time t. And the default intensity–or hazard rate in survival analysis terms–is defined as the probability of default per unit of time conditional on no earlier default:

$$\lambda(t) \equiv \lim_{dt \to 0} \frac{P(t \leq t^* \leq t + dt)}{dt(1 - F(t))} \tag{32}$$

$$= \frac{F'(t)}{1 - F(t)} \tag{33}$$

solving the differential equation we have that:

$$F(t) = 1 - e^{-\int_0^t \lambda(u) du} \tag{34}$$

And with a constant default intensity λ and using our approximation, for a one year CDS we have that:

$$F(1) = 1 - e^{-\lambda} = 1 - e^{-\frac{c}{1-x}} \tag{35}$$

which allows us to calculate the default probability given the spread and the recovery rate.

C Principal Components Analysis Figures

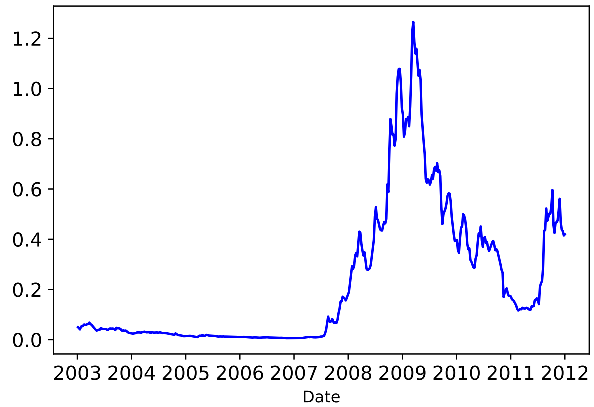

(a) Time Series of First Component.

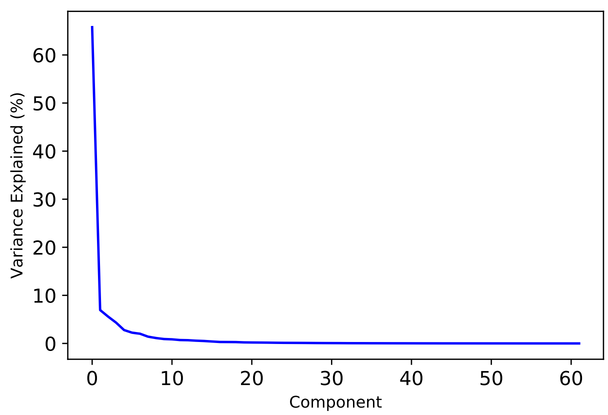

(b) Variance Explained by Each Component.

Figure 8: Principal Components Analysis Results

D Alternative Countercyclical Rules

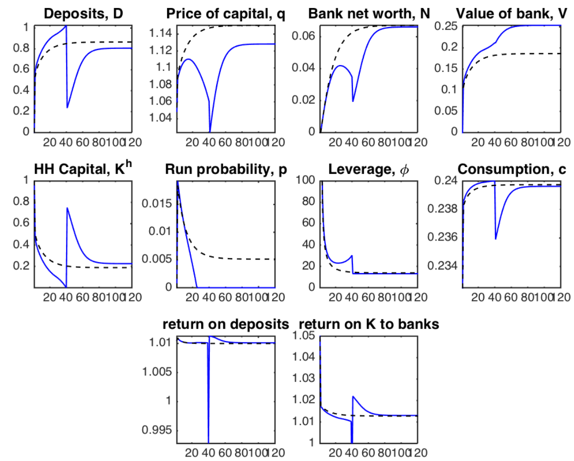

(a) Rule Starts after 40 Periods.

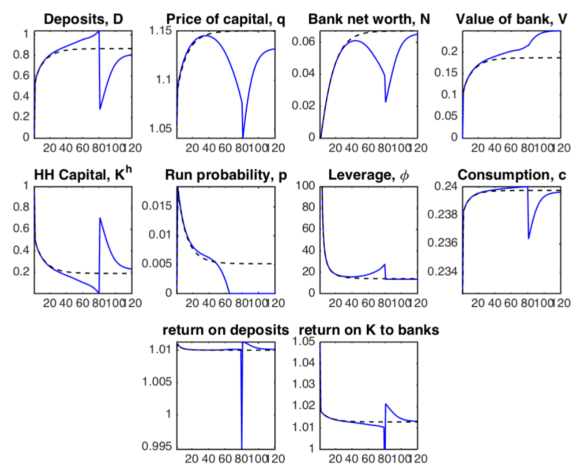

(b) Rule Starts after 80 Periods.

Figure 9: With Leverage Constraint, No Policy (GK15). The leverage constraint used is the highest such that run probability goes to zero.