

Bailouts, Bail-ins and Banking Crises*

Todd Keister
Rutgers University
todd.keister@rutgers.edu

Yuliyana Mitkov
Rutgers University
ymitkov@economics.rutgers.edu

June 11, 2017

We study the interaction between a government’s bailout policy during a banking crisis and individual banks’ willingness to impose losses on (or “bail in”) their investors. Banks in our model hold risky assets and are able to write complete, state-contingent contracts with investors. In the constrained efficient allocation, banks experiencing a loss immediately bail in their investors and this bail-in removes any incentive for investors to run on the bank. In a competitive equilibrium, however, banks may not enact a bail-in if they anticipate being bailed out. In some cases, the decision not to bail in investors provokes a bank run, creating further distortions and leading to even larger bailouts. We ask what macroprudential policies are useful when bailouts crowd out bail-ins.

* Preliminary and incomplete draft. We thank seminar and conference participants at the Federal Reserve Bank of Cleveland, the Federal Reserve Board, de Nederlandsche Bank, Rutgers University, the University of Exeter, the 2016 Oxford Financial Intermediation Theory (OxFIT) Conference, the 2nd Chicago Financial Institutions Conference, the Spring 2017 Midwest Macroeconomics Meetings, and especially Fabio Castiglionesi, Huberto Ennis, Charles Kahn and Joel Shapiro for helpful comments.

1 Introduction

In the years since the financial crisis of 2008 and the associated bailouts of banks and other financial institutions, policy makers in several jurisdictions have drafted rules requiring that these institutions impose losses on (or “bail in”) their investors in any future crisis. These rules aim to both protect taxpayers and change the incentives of banks and investors in a way that makes a crisis less likely. While the specific requirements vary, and are often yet to be finalized, in many cases the bail-in will be triggered by an announcement or action taken by the bank itself. This fact raises the question of what incentives banks will face when deciding whether and when to take actions that bail in their investors. In this paper, we ask how the prospect of being bailed out by the government influences banks’ bail-in decisions and how these decisions, in turn, affect the susceptibility of the banking system to a run by investors.

At one level, the reason why banks and other financial intermediaries sometimes experience runs by their investors is well understood. Banks offer deposit contracts that allow investors to withdraw their funds at face value on demand or at very short notice. During a bank run, investors fear that a combination of real losses and/or heavy withdrawals will leave their bank unable to meet all of its obligations. This belief makes it individually rational for each investor to withdraw her funds at the first opportunity; the ensuing rush to withdraw then guarantees that the bank does indeed fail, justifying investors’ pessimistic beliefs.¹

A key element of this well-known story is that the response to a bank’s losses and/or a run by its investors is *delayed*. In other words, there is a period of time during which a problem clearly exists and investors are rushing to withdraw, but the bank continues to operate as normal. Only when the situation becomes bad enough is some action – freezing deposits, renegotiating obligations, imposing losses on some investors, etc. – taken. This delay tends to deepen the crisis and thereby increase the incentive for investors to withdraw their funds at the earliest opportunity.

From a theoretical perspective, this delayed response to a crisis presents a puzzle. A run on the bank creates a misallocation of resources that makes the bank’s investors as a group worse off. Why do these investors not collectively agree to an alternative arrangement that efficiently allocates whatever losses have occurred while minimizing liquidation and other costs? In particular, why does the banking arrangement not respond more quickly to whatever news leads investors to begin to panic and withdraw their funds?

Most of the literature on bank runs resolves this puzzle using an incomplete-contracts

¹ This basic logic applies not only to commercial banking but also to a wide range of financial intermediation arrangements. See Yorulmazer (2014) for a discussion of several distinct financial intermediation arrangements that experienced run-like episodes during the financial crisis of 2008.

approach.² In particular, it is typically assumed to be impossible to write and/or enforce the type of contracts that would be needed to generate fully state-contingent payments to investors. The classic paper of Diamond and Dybvig (1983), for example, assumes that banks must pay withdrawing investors at face value until the bank has liquidated all of its assets and is completely out of funds. Other contracts – in which, for example, the bank is allowed to impose withdrawal fees when facing a run – are simply not allowed. Even those more recent papers that study more flexible banking arrangements impose some incompleteness of contracts. Peck and Shell (2003), for example, allow a bank to adjust payments to withdrawing investors based on any information it receives. However, the bank is assumed not to observe the realization of a sunspot variable that is available to investors and, in this sense, the ability to make state-contingent payments is still incomplete.³

If the fundamental problem underlying the fragility of banking arrangements is incompleteness of contracts, then an important goal of financial stability policy should be to remove this incompleteness. In other words, a key conclusion of the literature to date is that policy makers should aim to create legal structures under which more fully state-contingent banking contracts become feasible. There has, in fact, been substantial progress in this direction in recent years, including the establishment of orderly resolution mechanisms for large financial institutions and other ways of “bailing in” these institutions’ investors more quickly and more fully than in the past. The reform of money market mutual funds that was adopted in the U.S. in 2014 is a prime example. Under the new rules, certain types of funds are permitted to temporarily prohibit redemptions (called “erecting a gate”) and impose withdrawal fees during periods of high withdrawal demand if doing so is deemed to be in the best interests of the funds’ investors.⁴

In this paper, we ask whether making banking arrangements more fully state contingent – thereby allowing banks increased flexibility to bail in their investors – is sufficient to eliminate the problem of bank runs. To answer this question, we study a model in the tradition of Diamond and Dybvig (1983), but in which banks can freely adjust payments to investors based on any information available to the bank or to its investors. We think of this assumption as capturing an idealized situation in which policy makers’ efforts to improve the

² An important exception is Calomiris and Kahn (1991), in which the ex post misallocation of resources associated with a run is part of a desirable ex ante incentive arrangement to discipline bankers’ behavior. See also Diamond and Rajan (2001).

³ This same approach is taken in a large number of papers that study sunspot-driven bank runs in environments with flexible banking contracts, including Ennis and Keister (2010), Sultanum (2015), Keister (2016), and many others. See Andolfatto et al. (2016) for an interesting model in which the bank does not observe the sunspot state, but can attempt to elicit this information from investors.

⁴ See Ennis (2012) for a discussion of the issues involved in reforming money market mutual funds. There is a growing theoretical literature on bail-ins that we do not survey here; see, for example, Walther and White (2017).

contractual environment have been completely successful. We ask whether and under what conditions bank runs can occur in this idealized environment.

There are two aggregate states in our model and banks face uncertainty about the value of their investments. No banks experience losses in the good aggregate state, but in the bad aggregate state, some banks' assets are impaired. The government is benevolent and taxes agents' endowments in order to provide a public good. If there is a banking crisis, the government can also use these resources to provide bailouts to impaired banks. The government observes the aggregate state but cannot immediately tell which banks have impaired assets and which do not. In addition, the government cannot commit to a bailout plan; instead, the payment made to each bank will be chosen as a best response to the situation at hand. As in Keister (2016), this inability to commit implies that banks in worse financial conditions will receive larger bailout payments, as the government will aim to equalize the marginal utility of consumption across agents to the extent possible.

A bank with impaired assets has fewer resources available to make payments to investors. In an efficient allocation, such a bank would respond by immediately bailing in its investors, reducing all payments so that the loss is evenly shared. When the bank anticipates a government intervention, however, it may have an incentive to delay this response. By instead acting as if its assets were not impaired, the first group of its depositors who withdraw will receive higher payments. The government will eventually learn that the bank's assets are impaired and, at this point, will find the bank to be in worse financial shape as a result of the delayed response. The inability to commit prevents the government from being able to punish the bank at this point; instead, the bank will be given a larger bailout payment as the government aims to raise the consumption levels of its remaining investors. This larger payment then justifies the bank's original decision to delay taking action. In other words, we show that *bailouts delay bail-ins*.

The delay in banks' bail-in decisions has implications at both the aggregate and bank level. The delayed response makes banks with weak fundamentals even worse off and leads the government to make larger bailout payments, at the cost of a lower level of public good provision for everyone. In some cases, the misallocation of resources created by the delay may be large enough to give investors in weak banks an incentive to *run* in an attempt to withdraw before the bail-in is enacted. In these cases, the delayed bail-in *creates* financial fragility.

Our approach has novel implications for the form a banking crisis must take. Models in the tradition of Diamond and Dybvig (1983) typically do not distinguish between a single bank and the banking system; one can often think of the same model as applying equally well to both situations. If the banking system is composed of many banks, such models

predict that there could be a run on a single bank, on a group of banks, or on all banks, depending on how each bank's depositors form their beliefs. In our model, in contrast, there cannot be a run on only one bank, nor can there be a crisis in which only one bank chooses to delay bailing in its investors. If there is only a problem at one bank in our model, the government will choose to provide full deposit insurance, which removes any incentive for investors to run as well as any need for the bank to enact a bail-in. The problems of bank runs and delayed bail-ins can only arise in this model if the underlying losses are sufficiently widespread.

We then analyze possible policy responses to the inefficiencies that arise in the competitive equilibrium. Eliminating bailouts – if possible – would lead banks to immediately bail in their investors when facing losses and would prevent bank runs from occurring in equilibrium. However, it would also eliminate a valuable source of risk sharing and will often lower welfare. We study two policies that can always be used to increase welfare: placing a binding cap on the early payments made by banks and raising additional tax revenue in period 0. We show that the optimal policy combines both of these tools.

The remainder of the paper is organized as follows. The next section describes the economic environment and the actions available to banks, investors, and the government. In Section 3, we derive the constrained efficient allocation of resources in this environment, which is a useful benchmark for what follows. We provide the analysis of equilibrium, including delayed bail-ins and the potential for bank runs, in Section 4. We then discuss possible policy responses in Section 5 before concluding in Section 6.

2 The model

We base our analysis on a version of the Diamond and Dybvig (1983) model with flexible banking contracts and fiscal policy conducted by a government with limited commitment. We introduce idiosyncratic risk to banks' asset holdings and highlight how banks' incentives to react to a loss are influenced by their anticipation of government intervention. In this section, we introduce the agents, preferences, and technologies that characterize the economic environment.

2.1 The environment

Time. There are three time periods, labeled $t = 0, 1, 2$.

Investors. There is a continuum of investors, indexed by $i \in [0, 1]$, in each of a continuum of locations, indexed by $k \in [0, 1]$. Each investor has preferences characterized by

$$U(c_1^{i,k}, c_2^{i,k}, g; \omega^{i,k}) \equiv u(c_1^{i,k} + \omega^{i,k} c_2^{i,k}) + v(g), \quad (1)$$

where $c_t^{i,k}$ denotes the period- t private consumption of investor i in location k and g is the level of the public good, which is available in all locations. The random variable $\omega^{i,k} \in \Omega \equiv \{0, 1\}$ is realized at $t = 1$ and is privately observed by the investor. If $\omega^{i,k} = 0$, she is *impatient* and values private consumption only in period 1, whereas if $\omega^{i,k} = 1$ she values consumption equally in both periods. Each investor will be impatient with a known probability $\pi > 0$, and the fraction of investors who are impatient in each location will also equal π . The functions u and v are assumed to be smooth, strictly increasing, strictly concave and to satisfy the usual Inada conditions. As in Diamond and Dybvig (1983), the function u is assumed to exhibit a coefficient of relative risk aversion that is everywhere greater than one. Each investor is endowed with one unit of an all-purpose good at the beginning of period 0 and nothing in subsequent periods. Investors cannot directly invest their endowments and must instead deposit with a financial intermediary.

Banks. In each location, there is a representative financial intermediary that we refer to as a *bank*.⁵ Each bank accepts deposits in period 0 from investors in its location and invests these funds in a set of ex ante identical projects. A project requires one unit of input at $t = 0$ and offers a gross return of 1 at $t = 1$ or of $R > 1$ at $t = 2$ if it is not impaired. In period 1, however, $\sigma_k \in \Sigma \equiv \{0, \bar{\sigma}\}$ of the projects held by bank k will be revealed to be impaired. An impaired project is worthless: it produces zero return in either period. We will refer to σ_k as the *fundamental state* of bank k . A bank with $\sigma_k = 0$ is said to have *sound* fundamentals, whereas a bank with $\sigma_k = \bar{\sigma}$ is said to have *weak* fundamentals. The realization of σ_k is observed at the beginning of $t = 1$ by the bank's investors, but is not observed by anyone outside of location k .

After investors' preference types and banks' fundamental states are realized, each investor informs her bank whether she wants to withdraw in period 1 or in period 2. The bank observes *all* reports from its investors *before* making any payments to withdrawing investors. Those investors who chose to withdraw in period 1 then begin arriving sequentially at the bank in a randomly-determined order. Investors are isolated from either other during this process and no trade can occur among them; each investor simply consumes the payment she receives from her bank and returns to isolation. As in Wallace (1988) and others, this

⁵ While we use the term "bank" for simplicity, our model should be interpreted as applying to the broad range of financial institutions that engage in maturity transformation.

assumption prevents re-trading opportunities from undermining banks' ability to provide liquidity insurance.

Aggregate uncertainty. The fraction of banks whose assets are impaired depends on the aggregate state of the economy, which is either *good* or *bad*. In the good state, all banks have sound fundamentals. In the bad state, in contrast, a fraction $n \in [0, 1]$ of banks have weak fundamentals and, hence, total losses in the financial system are $n\bar{\sigma}$. The probability of the bad state is denoted q ; we interpret this event as an economic downturn that has differing effects across banks. If we think of the projects in the model as representing loans, for example, then the loans made by some banks are relatively unaffected by the downturn (for simplicity, we assume they are not affected at all), while other banks find they have substantial non-performing loans. Conditional on the bad aggregate state, all banks are equally likely to experience weak fundamentals. The ex-ante probability that a given bank's fundamentals will be weak is, therefore, equal to qn .

The government. The government in our model acts as both a fiscal authority and a banking supervisor. Its objective is to maximize the sum of all investors' expected utilities at all times. The government's only opportunity to raise revenue comes in period 0, when it chooses to tax investors' endowments at rate τ . In period 1, the government will use this revenue to provide the public good and, perhaps, to make transfers (*bailouts*) to banks. The government is unable to commit to the details of the bailout intervention ex-ante, but instead chooses the policy ex post, as a best response to the situation at hand.

The government observes the aggregate state of the economy at the beginning of period but, when the aggregate state is bad, is initially unable to determine *which* banks have weak fundamentals. After a measure $\theta \geq 0$ of investors have withdrawn from each bank, the government observes the idiosyncratic state σ_k of all banks and decides how to allocate its tax revenue between bailout payments to banks and the public good. The parameter θ thus measures how quickly the government can collect bank-specific information during a crisis and respond to this information. Banks that receive a bailout from the government are immediately placed in *resolution* and all subsequent payments made by these banks are chosen by the government. Once the public good has been provided, the government no longer has access to any resources and there will be no further bailouts.

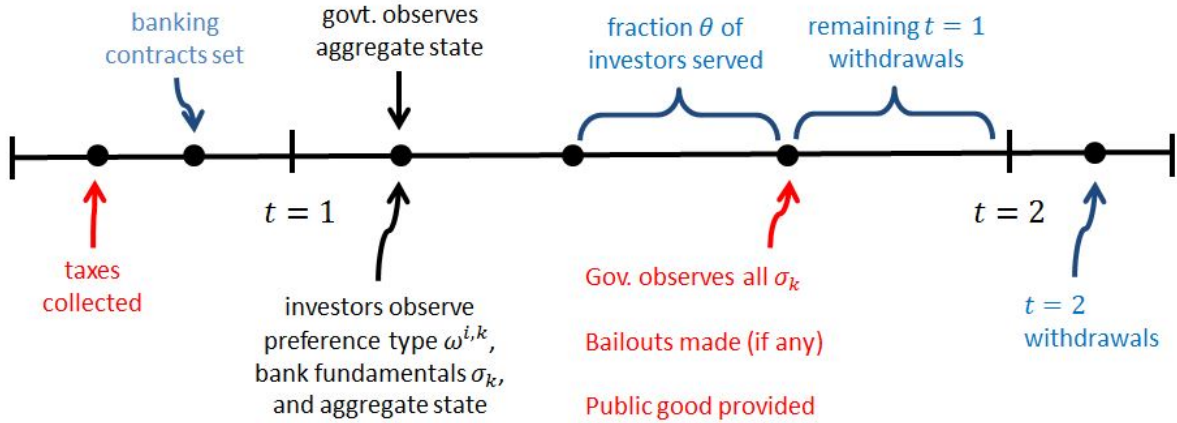


Figure 1: Timeline

2.2 Timeline

The sequence of events is depicted in Figure 1. In period 0, the government chooses the tax rate τ on endowments and investors deposit their after-tax endowment with the bank in their location. At the beginning of period 1, each investor observes her own preference type and the fundamental state of her bank; she then decides whether to withdraw in period 1 or in period 2. Banks observe the choices of their investors and begin making payments to withdrawing investors as they arrive. Once the measure of withdrawals reaches θ , the government observes all banks' fundamental states. At this point, the government may choose to bail out banks with weak fundamentals and places any banks that were bailed out into resolution. After bailout payments are made, all remaining tax revenue is used to provide the public good. Banks that were not bailed out continue to make payments to investors according to their contract, while the remaining payments made by banks in resolution are dictated by the government.

2.3 Discussion

Sequential service. While our model contains many elements that are familiar from the literature on bank runs, there are some key differences. Perhaps most importantly, banks in our model are able to condition payments to *all* investors on the *total* demand for early withdrawal. Green and Lin (2003) refer to this assumption as “the case without sequential service.” This language is potentially confusing when applied to our model: banks still serve withdrawing investors sequentially here. The key point, however, is that a bank is able to observe early withdrawal demand *before* deciding how to allocate resources across agents. By allowing all payments made by the bank to depend on this information, our contract

space is larger than that in most of the bank runs literature. In taking this approach, we aim to capture a contractual environment that is sufficiently rich to eliminate the underlying sources of bank runs that appear in the existing literature.

Costly public insurance. The role of aggregate uncertainty in our model is to force the government to fix a tax plan before knowing the aggregate losses of the banking system. If the government knew in advance how many banks would experience losses, it would collect additional taxes at $t = 0$ for the purpose of providing insurance against this location-specific shock. In fact, given that we assume the government can costlessly raise revenue through lump-sum taxes, it would collect enough revenue to provide complete insurance. Our timing assumption makes providing this insurance costly. If, for example, the probability q of the bad state is close to zero, the government will collect tax revenue equal to the desired level of the public good in the good aggregate state. If the realized state turns out to be bad, the marginal value of public resources will increase, but the government will be unable to raise additional revenue.

Delayed intervention. The assumption that the government observes bank-specific information with a delay is important for our analysis because it implies that some investors can withdraw before the government acts. One can narrowly interpret the parameter θ as measuring the time required to both carry out detailed examinations of banks and implement the legal procedures associated with resolving an insolvent bank. More broadly, however, θ can be thought of as also including a variety of other forces that lead governments to act slowly in the early stages of a crisis. For example, investors who are well-connected politically may use their influence to delay any government intervention until after they have had an opportunity to withdraw. The timing of the intervention might also reflect opaque incentives faced by regulators.⁶ In addition, Brown and Dinc (2005) provide evidence that the timing of a government's intervention in resolving a failed financial institution depends on the electoral cycle. Looking at episodes from 21 major emerging market economies in the 1990s, they find that interventions that would impose large costs on taxpayers and/or would more fully reveal the extent of the crisis were significantly less likely to occur before elections. (See also Rogoff and Sibert, 1988.) The effect of such political factors that delay the policy response

⁶ Kroszner and Strahan (1996) argue that throughout the 1980s the Federal Savings and Loan Insurance Corporation (FSLIC) faced a severe shortage of cash with which to resolve insolvent thrift institutions. This lack of funds forced the FSLIC to practice regulatory forbearance and to delay its explicit intervention in insolvent mutual thrifts in anticipation that the government would eventually supply additional resources. This delay led a large number of insolvent thrift institutions to maximize the value of future government liabilities guarantees (at the taxpayers' expense) by continuing to pay high dividends until the eventual resolution mechanism was put in place.

to a crisis would be captured in our model by an increase in the parameter θ .

3 The constrained efficient allocation

We begin by studying an allocation that will serve as a useful benchmark in the analysis. Suppose a benevolent planner could control the operations of all banks and the government, as well as investors' withdrawal decisions. This planner observes all of the information available to banks and investors, including each investor's preference type. It faces the same restrictions on fiscal policy as the government; in particular, all tax revenue must be raised at $t = 0$, before the aggregate state is realized. The planner allocates resources to maximize the sum of all investors' utilities.

It is fairly easy to see that the planner will direct all impatient investors to withdraw at $t = 1$, since they do not value later consumption, and will direct all patient investors to withdraw at $t = 2$, since it is less expensive to provide consumption to them after investment has matured. In addition, because investors are risk averse, the planner will choose to treat investors and banks symmetrically. In the good aggregate state, the planner will give a common level of consumption c_{10} in period 1 to all impatient investors and a common level c_{20} in period 2 to all patient investors. (The second subscript indicates that these consumption levels pertain to the good aggregate state, where zero banks have weak fundamentals.) In the bad aggregate state, the planner will give a common consumption profile (c_{1S}, c_{2S}) to investors in all banks with strong fundamentals and a common profile (c_{1W}, c_{2W}) to investors in all banks with weak fundamentals. These consumption levels will be chosen to maximize

$$(1 - q) \{ \pi u(c_{10}) + (1 - \pi) u(c_{20}) + v(\tau) \} \\ + q \left\{ \begin{array}{l} (1 - n) (\pi u(c_{1S}) + (1 - \pi) u(c_{2S})) + n (\pi u(c_{1W}) + (1 - \pi) u(c_{2W})) \\ + v(\tau - (1 - n)b_S - nb_W) \end{array} \right\}.$$

subject to feasibility constraints

$$\pi c_{10} + (1 - \pi) \frac{c_{20}}{R} \leq 1 - \tau \quad (2)$$

$$\pi c_{1S} + (1 - \pi) \frac{c_{2S}}{R} \leq 1 - \tau + b_S \quad (3)$$

$$\pi c_{1W} + (1 - \pi) \frac{c_{2W}}{R} \leq 1 - \tau - \bar{\sigma} + b_W, \quad (4)$$

where b_z denotes the per-investor transfer (or "bailout") given to each bank of type z in the bad aggregate state. These constraints each state that the present value of the consumption given to depositors in a bank must come from the initial deposit $1 - \tau$, minus the loss $\bar{\sigma}$ for

banks with weak fundamentals, plus any bailout received.⁷ The restriction that the planner cannot raise additional tax revenue in period 1 is equivalent to saying that the bailout payments must be non-negative,

$$b_S \geq 0 \quad \text{and} \quad b_W \geq 0. \quad (5)$$

The first-order conditions for the optimal consumption levels can be written as

$$u'(c_{1z}) = Ru'(c_{2z}) = \mu_z \quad \text{for } z = 0, S, W, \quad (6)$$

where μ_z is the Lagrange multiplier on the resource constraint associated with state z normalized by the probability of a bank ending up in that state. The first-order condition for the choice of tax rate τ can be written as

$$(1 - q)v'(\tau) + qv'(\tau - (1 - n)b_S - nb_W) = (1 - q)\mu_0 + q(1 - n)\mu_S + qn\mu_W, \quad (7)$$

which states that the expected marginal value of a unit of public consumption equals the expected marginal value of a unit of private consumption at $t = 0$. The first-order conditions for the bailout payments are

$$v'(\tau - (1 - n)b_S - nb_W) \geq \mu_z, \quad \text{with equality if } b_z > 0, \quad \text{for } z = S, W. \quad (8)$$

If the marginal value of private consumption in some banks were higher than the marginal value of public consumption in the bad aggregate state, the planner would transfer resources to (or “bail out”) these banks until these marginal are equalized. If instead the marginal value of private consumption in a bank is lower than the marginal value of public consumption, the bank will not be bailed out and the constraint in (5) will bind.

The following two propositions characterize the key features of the constrained efficient allocation of resources in our environment. First, the consumption of investors in banks with sound fundamentals is independent of the aggregate state and these banks do not receive bailouts.⁸

⁷ Note that our notation does not allow the planner to make bailout payments in the good aggregate state. This assumption prevents the planner from being able to make tax revenue fully state-contingent by, for example, setting $\tau = 1$ and holding all resources outside of the banking system until the aggregate state is revealed.

⁸ The first part of this result depends on our simplifying assumption that sound banks are completely unaffected by the bad aggregate state, but the second part of the result does not. Even if sound banks were to experience some losses during an economic downturn, the planner would not choose to bail out these banks as long as the losses are small relative to those at weak banks.

Proposition 1. *The constrained efficient allocation satisfies*

$$(c_{10}^*, c_{20}^*) = (c_{1S}^*, c_{2S}^*) \quad \text{and} \quad b_S^* = 0.$$

Given this result, we will drop the (c_{10}, c_{20}) notation in what follows and use (c_{1S}, c_{2S}) to refer to the consumption profile for investors in a bank with sound fundamentals regardless of the aggregate state. Our second result shows that this profile is different from the one assigned by the planner to investors in banks with weak fundamentals.

Proposition 2. *The constrained efficient allocation satisfies*

$$(c_{1S}^*, c_{2S}^*) \gg (c_{1W}^*, c_{2W}^*) \quad \text{and} \quad b_W^* > 0.$$

This result shows that the constrained efficient involves a combination of *bailouts* and *bail-ins* for investors in banks with weak fundamentals. The optimal bailout b_W^* gives investors partial insurance against the risk associated with their bank's losses. However, the consumption of investors in weak banks remains below that of investors in sound banks; this difference can be interpreted as the degree to which the planner "bails in" the investors in weak banks. The efficient level of insurance is only partial in this environment because offering insurance is costly; it requires the planner to collect more tax revenue, which leads to an inefficiently high level of the public good in the good aggregate state.

It is worth pointing out that the constrained-efficient bail-in applies equally to *all* investors in a weak bank, regardless of when they arrive to withdraw. While the desirability of this feature follows immediately from risk aversion, we will see below that it often fails to hold in a decentralized equilibrium. It is also worth noting that the constrained efficient allocation is incentive compatible. The first-order conditions (6) and $R > 1$ imply that $c_{1z}^* < c_{2z}^*$ holds for every state z and, hence, a patient investor always prefers her allocation to that given to an impatient investor (and vice versa).

4 Equilibrium

In this section we begin our investigation of the decentralized economy. Compared to the planner's economy discussed in the previous section, the decentralized economy differs in the following important ways. First, investors' preference types are private information and the banking contract therefore allows investors to choose the period in which they withdraw. Second, each bank is concerned solely with its own investors and takes economy-wide variables, including the level of the public good, as given. Third, there is asymmetric

information between the banks and the government; while the government immediately observes the aggregate state at the beginning of $t = 1$, it must wait for θ withdrawals to take place before observing bank-specific states. The government then makes bailout payments to banks with weak fundamentals and places these bank into resolution. Importantly, the bailout and resolution policies cannot be set ex-ante, but instead are chosen as a best response to the situation at hand.

In this section, we study equilibrium in the withdrawal game played by an individual bank's investors, taking the actions of investors at other banks (and the government) as given. In section 5, we study the joint determination of equilibrium actions across all banks.

4.1 Preliminaries

We begin by reviewing the timeline of events in Figure 1 for the decentralized economy and then provide a general definition of equilibrium.

The tax rate. To simplify the analysis in this section, we assume that the tax rate τ levied by the government in period 0 is set to the value from the constrained efficient allocation, τ^* . We derive equilibrium withdrawal behavior and the equilibrium allocation of resources for this given tax rate. In Section 5, we examine the government's optimal choice of tax rate given the equilibrium outcomes identified in this section.

Banking contracts. In period 0, each bank establishes a contract that specifies how much it will pay to each withdrawing investor as a function of both the bank's fundamental state $\sigma_k \in \{0, \bar{\sigma}\}$ and the fraction $\rho^k \in [\pi, 1]$ of its investors who choose to withdraw early. We allow the government to set an upper bound \bar{c} on the payments made to any investor withdrawing in period 1. One way to justify this upper bound is to assume that while the government cannot dictate the exact terms of the contractual arrangement between a bank and its investors, it is able to impose broad guidelines on the types of contract banks are allowed to offer.

Because investors are risk averse, it will be optimal for a bank to give the same level of consumption to all investors who withdraw in the same period.⁹ Let c_1^k denote the payment made by the bank to each investor who withdraws in period 1. In period 2, the bank

⁹ Keep in mind that our environment is different from that studied by Wallace (1990), Green and Lin (2003), Peck and Shell (2003) and others where the bank gradually learns about the demand for early withdrawal by observing investors' actions as they take place. Here, a bank directly observes total early withdrawal demand before making any payments to investors. It learns no new information as investors sequentially withdraw at $t = 1$ and, therefore, an optimal arrangement will always give the same level of consumption to each of these investors.

will divide its matured investment, plus any bailout payment received, evenly among its remaining depositors. The operation of the bank is, therefore, completely described by the function

$$c_1^k : \{0, \bar{\sigma}\} \times [\pi, 1] \rightarrow [0, \bar{c}]. \quad (9)$$

We refer to the function in (9) as the *banking contract*. There is full commitment to the banking contract in the sense that the plan in (9) will be followed unless the bank is placed into resolution by the government. Each bank's contract is chosen to maximize the expected utility from private consumption of the bank's investors.¹⁰

Bailouts and resolution. After a fraction θ of investors have withdrawn at $t = 1$, the government observes the fundamental state σ_k of each bank and chooses a bailout payment b^k for each bank with weak fundamentals. It then dictates the payments made by these banks to their remaining investors as part of the resolution process. We characterize the government's bailout/resolution policy below.

Withdrawal strategies. An investor's withdrawal decision can depend on both her preference type ω_i^k and the fundamental state of her bank σ_k . (See Figure 1.) A withdrawal strategy for investor i in bank k is, therefore, a mapping:

$$y_k^i : \Omega \times \Sigma \rightarrow \{0, 1\}$$

where $y_k^i = 0$ corresponds to withdrawing in period 1 and $y_k^i = 1$ corresponds to withdrawing in period 2. An investor will always choose to withdraw in period 1 if she is impatient. We introduce the following labels to describe the actions an investor takes in the event she is patient.

Definition 1. For given σ_k , we say investor i in bank k follows:

- (i) the *no-run strategy* if $y_k^i(\omega_k^i, \sigma_k) = \omega_k^i$ for $\omega_k^i \in \{0, 1\}$, and
- (ii) the *run strategy* if $y_k^i(\omega_k^i, \sigma_k) = 0$ for $\omega_k^i \in \{0, 1\}$.

We use y_k to denote the profile of withdrawal strategies for all investors in bank k and y to denote the withdrawal strategies of all investors in the economy. It will often be useful to summarize a profile of withdrawal strategies by the fraction of investors who follow the run

¹⁰ This outcome would obtain, for example, if multiple banks competed for deposits in each location. We use a representative bank in each location only to simplify the presentation.

strategy in that profile, which we denote

$$x_{\sigma_k} \equiv \int_{i \in [0,1]} (1 - y_k^i(\omega_k^i = 1, \sigma_k)) dk.$$

Similarly, we use ρ_k to denote the total demand for early withdrawal from bank k in a given profile, which equals

$$\rho_k = \pi + (1 - \pi)x_k.$$

Allocations. The allocation of private consumption in bank k in a particular state is a specification of how many investors withdraw at $t = 1$ in that state, how much consumption each of these investors receives, and how much consumption each remaining investor receives at $t = 2$. This allocation depends on the banking contract for bank k , the withdrawal strategies of investors in bank k , and the government intervention in bank k (if any). The details of the government intervention, in turn, may depend on the contracts of other banks and the withdrawal strategies of investors in those banks. In general, therefore, the optimal withdrawal behavior for each investor in bank k may depend on the contracts offered by *other* banks and on the withdrawal strategies of investors in *other* banks.

Equilibrium. To study equilibrium withdrawal behavior within a single bank, we fix all banking contracts, the government's intervention policy, and the withdrawal strategies of investors in all other banks, y_{-k} . Together, these items determine the payoffs of what we call the *withdrawal game in bank k* . That is, holding these other items fixed, we can calculate the allocation of private consumption in bank k as a function of the strategies y_k played by that bank's investors. An equilibrium of this game is a profile of strategies for the bank's investors, y_k , such that for every investor i in the bank, y_k^i is a best response to the strategies of the other investors, y_k^{-i} .

An equilibrium of the overall economy is a profile of withdrawal strategies for all investors y^* such that (i) y_k^* is an equilibrium of the withdrawal game in bank k generated by the strategies y_{-k}^* of investors in all other banks, for all k , (ii) the contract in bank k maximizes the expected utility of its investors taking as given the contracts and withdrawal strategies y_{-k}^* of investors in all other banks, for all k , and (iii) the government's bailout and resolution policy maximizes total welfare taking as given all banking contracts and withdrawal strategies y^* . Notice how this definition reflects the timing assumptions depicted in Figure 1. Investors in bank k recognize that their choice of contract will influence equilibrium withdrawal behavior within their own bank but will not affect outcomes at other banks.¹¹ The

¹¹ This result follows, in part, from the assumption that there are a continuum of locations and, hence, the

government's bailout and resolution policies, in contrast, are set after all banking contracts and withdrawal decisions have been made. Because the government cannot commit to the details of these policies *ex ante*, it acts to maximize welfare taking all bank contracts and withdrawal decisions as given.

In the subsections that follow, we derive the properties of the contracts that a bank will use in equilibrium, focusing first on the case where its fundamental state is strong. We then turn to the case where the bank's fundamental state is weak, which requires characterizing the optimal bailout and resolution policies as well. Finally, we then characterize the equilibrium of the entire economy, in which each bank's contract is a best response to all other contracts.

4.2 Banks with sound fundamentals

We assume the government does not give bailouts to banks with sound fundamentals, nor does it place them in resolution.¹² As a result, all investors who chose to withdraw at $t = 1$ receive the contractual amount $c_1^k(0, \rho^k)$, and all investors who chose to withdraw at $t = 2$ receive an even share of the bank's assets, $c_2^k(0, \rho^k)$, which is implicitly defined by

$$p^k c_1(0, \rho^k) + (1 - \rho^k) \frac{c_2(0, \rho^k)}{R} = 1 - \tau. \quad (10)$$

The bank and its investors recognize that ρ^k will result from the equilibrium withdrawal behavior of investors. In particular, if the bank offers a higher payment in period 1 than in period 2, all investors will choose to withdraw early. In other words, equilibrium requires

$$\rho^k = \left\{ \begin{array}{c} \pi \\ \in [\pi, 1] \\ 1 \end{array} \right\} \quad \text{as} \quad c_1(0, \rho^k) \left\{ \begin{array}{c} < \\ = \\ > \end{array} \right\} c_2(0, \rho^k). \quad (11)$$

We refer to (11) as the *implementability constraint*. If a triple (ρ_S, c_{1S}, c_{2S}) satisfy both (10) and (11), then any banking contract with $c_1^k(\sigma_k, \rho_S) = c_{1S}$ for $\sigma_k = 0$ can implement this allocation as an equilibrium of the withdrawal game in bank k , regardless of how the payments $c_1^k(\sigma_k, \rho^k)$ are set for other values of ρ^k . The following result shows that something stronger is true: by choosing these other payments appropriately, the banking contract can be set so that the withdrawal game in bank k has a *unique* equilibrium.

actions taken at one bank have no effect on aggregate variables or on the behavior of the government toward other banks.

¹² Recall from Section 3 that the constrained efficient allocation involves zero bailouts for sound banks. Our assumption here is that the government is able to commit to follow this policy.

Proposition 3. *If the allocation (ρ_S, c_{1S}, c_{2S}) satisfies both (10) and (11), there exists a contract that implements this allocation as the unique equilibrium of the withdrawal game played by a sound bank's investors.*

In light of the above proposition, we can recast the problem of choosing the optimal banking contract as one of directly choosing the allocation (ρ_S, c_{1S}, c_{2S}) to maximize expected utility

$$V_S(\rho_S, c_{1S}) \equiv \rho_S u(c_{1S}) + (1 - \rho_S)u(c_{2S}) \quad (12)$$

subject to the feasibility constraint (10) and the implementability constraint (11) and the restriction $c_1^k \in [0, \bar{c}]$ for all banks. The next result characterizes the solution to this problem.

Proposition 4. *When bank k has sound fundamentals, there is a unique equilibrium of the withdrawal game in the bank associated with the optimal banking contract. The equilibrium allocation (ρ_S, c_{1S}, c_{2S}) satisfies $\rho_S = \pi$ and $c_{1S} = \min\{c_{1S}^*, \bar{c}\}$.*

This result shows that as long as the upper bound \bar{c} is set high enough to allow it, the equilibrium allocation within a sound bank is the same as in the constrained efficient allocation. In other words, resources are always allocated efficiently within a sound bank and investors never run on these banks.

There are many banking contracts that implement the desired allocation, one of which is

$$c_1^k(0, \rho^k) = \begin{cases} \min\{c_{1S}^*, \bar{c}\} \\ 0 \end{cases} \quad \text{if} \quad \begin{cases} \rho^k = \pi \\ \rho^k > \pi \end{cases}. \quad (13)$$

Under this contract, the bank would immediately suspend withdrawals if more than a fraction π of its investors request to withdraw early, saving all of its resources until period 2.¹³ It is easy to verify that waiting to withdraw in period 2 is then the best response for a patient investor to any profile of withdrawal strategies for the other investors. As a result, this contract uniquely implements the desired allocation in the withdrawal game in bank k and a bank run will never occur.

4.3 Banks with weak fundamentals

Characterizing the outcome of the withdrawal game in a weak bank is more complicated because it depends on the government's bailout and resolution policies. Let W denote the set of weak banks,

¹³ The reaction of setting early payments to zero when $\rho_k > \pi$ holds is stronger than needed to eliminate the run equilibrium. It would suffice to set these payments low enough that a patient investor would receive more consumption by waiting to withdraw in period 2. The key point is that the bank can easily structure the contract to prevent a run; the form in equation (10) makes this point in a particularly clean way.

$$W \equiv \{k \in [0, 1] \text{ s.t. } \sigma_k = \bar{\sigma}\}. \quad (14)$$

After the first θ withdrawals have taken place in all banks, the government observes the fundamental state σ_k of each bank. For $k \in W$, the government also observes the bank's current condition: the amount of resources remaining in the bank and the fraction of the bank's remaining investors who are impatient. The government then decides on a bailout payment b^k for each $k \in W$ and places these banks into resolution. We derive the government's best responses by working backward, beginning with the resolution stage.

Resolution. To simplify the presentation, we assume that when a bank is placed into resolution, the government directly observes the preference types of the bank's remaining investors and allocates the bank's resources (including the bailout payment) conditional on these types. One could imagine, for example, the using the court system to evaluate individual's true liquidity needs, as discussed in Ennis and Keister (2009). This assumption is **not** important for our results, however. If instead the government were to offer a new banking contract and have the remaining investors play a withdrawal game based on this new contract, it could choose a contract that yields the outcome we study here as the unique equilibrium of that game.

Let $\hat{\psi}_k$ denote the per-capita level of resources in bank k , including any bailout payment received, after the first θ withdrawals have taken place. Then we have

$$\hat{\psi}_k = \frac{(1 - \tau)(1 - \bar{\sigma}) - \theta c_1^k(\bar{\sigma}, \rho_k) + b_k}{1 - \theta}, \quad (15)$$

where b_k is the per-investor bailout given to bank k . Let $\hat{\rho}^k$ denote the fraction of the bank's remaining investors who are impatient. The allocation of resources for a bank in resolution is chosen to maximize the sum of the utilities for the remaining investors in the bank:

$$\hat{V}(\hat{\psi}_k; \hat{\rho}_k) \equiv \max_{\hat{c}_1^k, \hat{c}_2^k} (1 - \theta) (\hat{\rho}_k u(\hat{c}_1^k) + (1 - \hat{\rho}_k) u(\hat{c}_2^k)) \quad (16)$$

subject to the feasibility constraint

$$\hat{\rho}_k \hat{c}_1^k + (1 - \hat{\rho}_k) \frac{\hat{c}_2^k}{R} \leq \hat{\psi}_k \quad (17)$$

The optimal choice of post-bailout payments is determined by the first order condition

$$u'(\hat{c}_1^k) = R u'(\hat{c}_2^k) = \hat{\mu}(\hat{\psi}_k; \hat{\rho}_k), \quad (18)$$

where $\hat{\mu}$ is the Lagrange multiplier on the resource constraint. Since $R > 1$, this condition implies that a bank in resolution provides more consumption to patient investors withdrawing in period 2 than to the remaining impatient investors who withdraw in period 1.

Bailouts. In choosing the bailout payments $\{b^k\}$, the government's objective is to maximize the sum of the utilities of all investors in the economy. While bailouts raise the private consumption of investors in weak banks, they lower the provision of the public good, which affects all investors. The government's objective in choosing these payments can be written as

$$\max_{\{b^k\}_{k \in W}} \int_W \hat{V}(\hat{\psi}_k; \hat{\rho}_k) dk + v \left(\tau - \int_W \hat{b}_k dk \right) \quad (19)$$

The first-order condition for this problem is

$$\hat{\mu}(\hat{\psi}_k; \hat{\rho}_k) = v' \left(\tau - \int_W \hat{b}_k dk \right) \text{ for all } k. \quad (20)$$

Notice that the right-hand side of this equation – the marginal utility of public consumption – is independent of k . The optimal bailout policy thus has the feature that the marginal value of resources will be equalized across all weak banks, regardless of their chosen banking contract or the withdrawal behavior of their investors. As a result, all banks in resolution will give a common consumption allocation (\hat{c}_1, \hat{c}_2) to their impatient and patient investors, respectively. These consumption values and the bailout payments b^k will satisfy the resource constraint

$$\hat{\rho}_k \hat{c}_1 + (1 - \hat{\rho}_k) \frac{\hat{c}_2}{R} = \frac{(1 - \tau)(1 - \bar{\sigma}) - \theta c_1^k(\bar{\sigma}, \rho_k) + b_k}{1 - \theta}. \quad (21)$$

Using the fact that (\hat{c}_1, \hat{c}_2) is the same in all weak banks, this constraint shows that the bailout payment made to bank k is increasing in the amount paid out by the bank before being bailed out, $c_1^k(\bar{\sigma}, \rho_k)$. Together with the first-order condition (18), this constraint implies that b^k is increasing in the fraction of bank k 's remaining investors who are impatient, $\hat{\rho}_k$.

Withdrawal behavior. A fraction θ of a weak bank's investors will receive the amount specified by the contract, $c_1^k(\bar{\sigma}, \rho_k)$, before the government intervenes. The bank will then be bailed out and placed into resolution. Its remaining impatient investors will receive \hat{c}_1 and its remaining patient investors will receive \hat{c}_2 , as derived above. A patient investor will choose to withdraw early if the contract sets $c_1^k > \hat{c}_2$ and will choose to wait if $c_1^k < \hat{c}_2$. In other words, the fraction ρ^k of investors who attempt to withdraw from a weak bank at $t = 1$ will

satisfy

$$\rho^k = \left\{ \begin{array}{c} \pi \\ \in [\pi, 1] \\ 1 \end{array} \right\} \text{ if } c_{1W} \left\{ \begin{array}{c} < \\ = \\ > \end{array} \right\} \hat{c}_2. \quad (22)$$

In choosing a contract, the bank recognizes that its investors will behave in accordance with (22), which we refer to as the implementability constraint for weak banks. The next result is the analog of Proposition 3 from the previous section: it shows that any allocation satisfying the implementability constraint can be implemented as the unique equilibrium of the withdrawal game in bank k .

Proposition 5. *If (ρ_W, c_{1W}) satisfy (22), then there exists a banking contract c_1^k that implements this allocation as the unique equilibrium of the withdrawal game played by a weak bank's investors.*

This results allows us to formulate the bank's optimal contract problem as one of directly choosing the allocation (ρ_W, c_{1W}) to maximize

$$V_W(\rho_W, c_{1W}) \equiv \theta u(c_{1W}) + (1 - \theta) [\hat{\rho}_{\bar{\sigma}} u(\hat{c}_1) + (1 - \hat{\rho}_{\bar{\sigma}}) u(\hat{c}_2)] \quad (23)$$

subject to the implementability constraint for weak banks (22) and the relationship

$$\hat{\rho}_{\bar{\sigma}} \equiv \frac{\pi}{1 - \theta} \left(\frac{\rho_W - \theta}{\rho_W} \right). \quad (24)$$

This last expression shows how the fraction of the bank's remaining investors after θ withdrawals are impatient depends on the fraction that initially attempt to withdraw early.

The first term in the objective function in (23) is clearly increasing in the choice of c_{1W} , reflecting the bank's desire to give as much consumption as possible to the investors who withdraw before the bank is placed into resolution. However, the implementability constraint (22) shows that if c_{1W} is set greater than \hat{c}_2 , the bank's investors will run, in which case ρ^k will equal 1. A run on the bank is costly because some early consumption is inefficiently given to patient investors; the fact can be seen by noting that $\hat{\rho}_{\bar{\sigma}}$ is an increasing function of ρ^k and the second term in the objective function (23) is strictly decreasing in $\hat{\rho}_{\bar{\sigma}}$. The next result shows that the solution to the bank's problem takes one of two forms: the early payments will either be set as high as possible, or will be set to the largest value that prevents a bank run.

Proposition 6. *The solution to the program of maximizing (23) subject to (22) and (24) will either set $c_{1W} = \bar{c}$ or $c_{1W} = \hat{c}_2$.*

Which of these two options will be optimal for the bank depends on how low the early payment would need to be set in order to prevent a run. Setting c_{1W} equal to the upper bound \bar{c} will be optimal whenever

$$V_W(\bar{c}, 1) > V_W(\hat{c}_2, \pi). \quad (25)$$

The above inequality implies that the loss to the remaining $1 - \theta$ investors in the bank resulting from keeping payments as high as possible and allowing a run is more than offset by the gain to the first fraction θ to withdraw. By observing that the inequality in (25) is equivalent to

$$u(\bar{c}) - u(\hat{c}_2) > (1 - \pi)(u(\hat{c}_2) - u(\hat{c}_1))$$

the equilibrium of the withdrawal game in a weak bank can be characterized as follows.

Proposition 7. *If bank k has weak fundamentals then:*

(i) *If $u(\bar{c}) - u(\hat{c}_2) < (1 - \pi)(u(\hat{c}_2) - u(\hat{c}_1))$, there is a unique equilibrium of the withdrawal game in bank k associated with the optimal banking contract. The equilibrium allocation has $\rho_W = \pi$ and $c_{1W} = \min\{\bar{c}, \hat{c}_2\}$.*

(ii) *If $u(\bar{c}) - u(\hat{c}_2) > (1 - \pi)(u(\hat{c}_2) - u(\hat{c}_1))$, there is again a unique equilibrium of the withdrawal game in bank k associated with the optimal banking contract. The equilibrium allocation in this case has $\rho_W = 1$ and $c_{1W} = \bar{c}$.*

(iii) *If $u(\bar{c}) - u(\hat{c}_2) = (1 - \pi)(u(\hat{c}_2) - u(\hat{c}_1))$, the withdrawal game in bank k has multiple equilibria, one with $\rho_W = \pi$ and $c_{1W} = \hat{c}_2$ and another with $\rho_W = 1$ and $c_{1W} = \bar{c}$.*

This result shows that, outside of a knife-edge case, the withdrawal game in a weak bank will have a unique equilibrium under the optimal banking contract. In this sense, a bank run in our model is fundamentally different from the type of self-fulfilling run normally studied in the literature based on Diamond and Dybvig (1983). When a bank is in case (ii) of Proposition 7, withdrawing early is a dominant strategy for the bank's investors. In this sense, a bank run in our model does not rely on investors' self-fulfilling beliefs about the actions of other investors in their bank. We show below, however, that the model may still have multiple equilibria because an investor's best response may depend critically on the withdrawal decisions of investors in *other* banks.

4.4 Equilibrium across banks

The preceding sections have investigated the equilibrium outcomes within a given bank, taking the actions of the government and the remaining banks as fixed. We now investigate the properties of the overall equilibrium across all banks, in which both the banking contract

and the withdrawal strategies in each bank are best responses to the actions taking place at other banks.

Constrained inefficiency. We begin by asking whether the equilibrium allocation is constrained efficient. Note that, in order for this allocation to be feasible in the decentralized economy, the upper bound \bar{c} on early payments must be set sufficiently high that sound banks are able to choose c_{1S}^* . For the analysis in this section, we will set $\bar{c} = c_{1S}^*$. Our next result shows that, even though it is feasible, the constrained efficient allocation is never an equilibrium of the decentralized economy.

Proposition 8. *The equilibrium allocation of resources is never constrained efficient.*

The bailout policy creates an incentive for weak banks to set their early payments as high as possible. The only reason a weak bank would voluntarily impose losses on its investors (by setting a payment below $\bar{c} = c_{1S}^*$) is to prevent a run. Note that preventing a run only requires that the payment in period 1 not exceed \hat{c}_2 and, as a result, a weak bank will never set its early payment below this level. In particular, a weak bank will never choose to bail in its investors all the way down to \hat{c}_1 , as occurs in the constrained efficient allocation.

Equilibrium bank runs. In addition to being constrained inefficient, the equilibrium of the full model will, in some cases, involve a run by investors on weak banks.

Proposition 9. *For some parameter values, there exists an equilibrium in which investors run on weak banks. In some cases this equilibrium is unique, but in others it coexists with another equilibrium in which no run occurs.*

In the run equilibrium, all investors in weak banks attempt to withdraw at $t = 1$, that is, the profile of withdrawal strategies has $x_{\bar{\sigma}} = 1$. A fraction θ of these investors successfully withdraw before the government observes $\sigma_k = \bar{\sigma}$ and places the bank into resolution.

The result in Proposition 9 is established in Figure 2 which depicts the type of equilibria that arise for different combinations of the parameters n , the fraction of weak banks, and $\bar{\sigma}$, the loss in each of them. The figure uses the utility function¹⁴

$$u(c_1^{i,k} + \omega^{i,k} c_2^{i,k}) = \frac{\left(c_1^{i,k} + \omega^{i,k} c_2^{i,k}\right)^{1-\gamma} - 1}{1-\gamma} \quad \text{and} \quad v(g) = \delta \frac{g^{1-\gamma}}{1-\gamma}. \quad (26)$$

For parameter combinations in the dark region in the lower-left part of the graph, there is a unique equilibrium of the model and the allocation in this equilibrium does not involve a bank

¹⁴ The other parameters of the model are set to $R = 1.5$, $\pi = 0.5$, $\gamma = 5$, $\delta = 0.5$, $q = 0.05$ and $\theta = 0.5$. The tax rate τ is set to its constrained efficient value from section 3.




Figure 2: Equilibrium with a bank run

run. When the losses $\bar{\sigma}$ suffered by a weak bank are small and/or few banks experience these losses, the process of resolving these banks has a relatively small cost for the government. When this cost is small, the government remains in good fiscal condition and will choose to make bailout payments that lead to relatively high consumption levels (\hat{c}_1, \hat{c}_2) for the remaining investors in banks placed into resolution. This fact, in turn, makes running in an attempt to withdraw before the government intervenes less attractive for patient investors in a weak bank. As a result, a unique equilibrium exists and all patient investors wait until $t = 2$ to withdraw.

In the unshaded region in the upper-right portion of the figure, in contrast, both the number of banks experiencing a loss and the amount lost by each of these banks are significant. In this case, the government's budget constraint will be substantially impacted by its desire to bail out weak banks in a crisis. As the marginal value of public resources rises, the bailout and resolution process will lead to lower consumption levels (\hat{c}_1, \hat{c}_2) for the remaining investors in these banks. When \hat{c}_2 is low enough, the equilibrium within a weak bank k will involve a run by patient investors, as shown in Proposition 7. The overall equilibrium in this region is again unique, but the (larger) losses on weak banks' asset are now compounded by the additional liquidation of assets and misallocation of resources created by the run.

Multiple equilibria. In the grey region in Figure 2, both of the equilibria described above exist. The fact that multiple equilibria exist in this region is particularly interesting in light of Proposition 7, which showed that the equilibrium of the withdrawal game within each

bank is unique except for in a knife-edge case. The multiplicity of equilibria illustrated in Figure 2 arises because of an externality in payoffs *across* weak banks. When a run occurs at other weak banks, this event causes more investment to be liquidated and leads to larger bailouts at those banks. The larger bailouts place greater strain on the government's budget constraint and lead – all else being equal – to a smaller bailout at bank k . In the lighter-shaded region in Figure 2, this smaller bailout lowers the consumption levels (\hat{c}_1, \hat{c}_2) enough to make running a best response for the patient investors in bank k . In other words, in our model there is a strategic complementarity in the withdrawal decisions of investors across banks. The usual strategic complementarity that appears in models in the Diamond-Dybvig tradition – which arises between investors within a bank – is eliminated by the more flexible banking contracts. However, the government's bailout and resolution policy introduces this new complementarity in actions across banks, which creates the region of multiple equilibria in Figure 2.

It is worth emphasizing that a run on bank k lowers the welfare of the bank's investors in much the same way as in the existing literature. Holding fixed the bailout payment it receives, a bank's investors would be strictly better off if there were no run on the bank. Moreover, the bank has contractual tools that would allow it to prevent the run. The problem, however, is that preventing the run requires decreasing the payment given to the first θ investors who withdraw, and this action would decrease the bailout payment the bank receives. Instead, in this equilibrium, the bank's investors choose to tolerate the run as a side effect of the plan that maximizes the level of payments the bank is able to make to its investors before the government intervenes.

Runs cannot be based on sunspots alone. The pattern in Figure 2 suggests that the run equilibrium does not exist for $\bar{\sigma}$ equal or sufficiently close to zero. The next proposition shows that this property holds more generally.

Proposition 10. *Given other parameter values, there exists $\bar{\sigma}^* > 0$ such that a bank run does not occur in equilibrium for any $\bar{\sigma} < \bar{\sigma}^*$.*

According to Proposition 10, a run in this environment cannot occur unless the fundamental was large enough. The reason is as follows. In the bad aggregate state, a sound bank would deliver higher utility to its investors compared to a weak bank. At the same time – conditional on weak banks *not* experiencing a run – the utility gain associated with being sound is approaching zero as $\bar{\sigma} \rightarrow 0$. If the run equilibrium exists for $\bar{\sigma}$ close to zero, then almost all of the losses in the weak banks will be generated solely by the run from their investors. In this case, an individual weak bank can always do better by deviating from

equilibrium play and implementing a contract with strong suspension clause both when the bank is sound and when the bank is weak. Indeed, since $\bar{\sigma} \approx 0$, by preventing a run, bank- k also ensures that its investors receive almost identical welfare as the investors in sound banks which in turn is strictly higher than the welfare in the remaining weak banks (since those weak banks experience a run). This reasoning establishes that the run equilibrium cannot exist for $\bar{\sigma}$ equal or sufficiently close to zero since an individual weak bank strictly gains by implementing a run-proof contract.

Runs must be systemic. Unlike other models of financial fragility, a bank run in our model cannot be an isolated event in the sense of occurring at a single bank in our model or even a small group of banks. The pattern in Figure 2 also suggests that a run can only occur when the number of weak banks is sufficiently large. The following proposition formalizes this result.

Proposition 11. *Given other parameter values, there exists $n^* > 0$ such that a bank run does not occur in equilibrium for any $n < n^*$.*

If the number of affected banks is small, the associated losses will have a small impact on the government’s budget constraint. If the government remains in good fiscal condition, the bailout policy it will choose ex post treats weak banks generously, leaving their patient investors with no incentive to run.

Observe how the environment we study is different from most of the literature on banking crisis. Here, runs will not occur in equilibrium unless there were real shocks $\bar{\sigma} > 0$, these shocks were large enough (Proposition 10) and sufficiently wide spread (Proposition 11). An implication of Propositions 10 and 11 is that a run cannot be based on sunspots alone, regardless of how many banks may have experienced the “*bad sunspot*”. Finally, it is important to stress that Propositions 10 and 11 are *not* stating that sunspots cannot play a role in our framework. Instead, a necessary condition for a sunspot state to affect equilibrium outcomes is that n and $\bar{\sigma}$ must both be positive and to belong to the region of parameters where both the run and the no-run equilibrium co-exist (Proposition 9). In this case, one can introduce a sunspot that would serve to coordinate investors’ withdrawal decisions on one or the other equilibrium. In this sense, a bank run would still be *self-fulfilling*. As stressed in the section above on multiple equilibria, however, the logic here is different from a standard Diamond and Dybvig model, since in our environment the investors in a given weak bank would choose to run if and only if they expect the investors in the other weak banks to run.

Discussion. A number of recent legislative changes aim to promote financial stability by endowing financial intermediaries with increased contractual flexibility, which would allow them to react as soon as they start to experience distress. For example, “gates” and withdrawal fees in money market mutual funds, swing pricing in the mutual fund industry more generally, and the new bail-in rules in the US, Europe and elsewhere can all be interpreted as giving intermediaries the opportunity – but not necessarily the obligation – of imposing losses on all (or subset) of their investors if this is deemed desirable for the long term health of the institution. The hope of these legislative reforms is that these new “bail-in options” would not only be effective in mitigating fragility (or even preventing runs entirely), but in addition, would eliminate the need for taxpayers to finance a bailout or at least drastically reduce the cost of government’s interventions.

For instance, one important purpose of the recent reforms to money market mutual funds in the U.S. is to reduce investors’ incentive to redeem quickly and ahead of others (i.e. to run) when the fund is in distress. The imposition of fees and gates must be approved by a fund’s board of directors, who are to use these tools only if this is determined to be in the best interest of their shareholders. Notice that from the perspective of our model, withdrawal fees and “gates” can be captured as setting lower payments in weak banks. Our results suggest that, in an environment characterized by limited commitment, asymmetric information and bailouts, these bail-in options may not be used and thus be ineffective in promoting financial stability. In the next section, we examine ways a policy maker might reduce the inefficiencies that arise in the competitive equilibrium in our model and promote financial stability.

5 Macprudential policy

Given that the equilibrium allocation studied above is always constrained inefficient and, in addition, may involve a welfare-decreasing run by investors on weak banks, it is natural to ask what types of prudential policy would be useful in this environment. In this section, we study three such policies: restricting the early payments made by banks, increasing government revenue, and eliminating bailouts. While none of these policies leads to the constrained efficient allocation derived in Section 3, each is capable of raising welfare in some situations.

5.1 Restricting early payments

In Section 4, we set the cap on early payments \bar{c} equal to c_{1S}^* , the payment made by sound banks in the constrained efficient allocation. This approach ensured that the constrained efficient allocation from Section 3 was feasible in the decentralized economy. We have es-


tablished in Proposition 8, however, that the decentralized equilibrium is never constrained efficient. The reason is that weak banks would never voluntarily impose a bail-in on their investors to the full extent required in the constrained efficient allocation. In this section, we will treat the cap on early payments as a *policy instrument* and allow the government to impose a potentially binding cap, that is $\bar{c} < c_{1S}^*$. Specifically, the government selects \bar{c} in period 0 in order to maximize the sum of investors' expected utilities. The main result in this subsection is the following.

Proposition 12. *For some parameter values, it is optimal for the government to impose a binding cap, that is $\bar{c}^* < c_{1S}^*$. Moreover, in the region of the parameter where the cap is binding, \bar{c}^* is a decreasing function of n , $\bar{\sigma}$, θ and q .*


Recall that the inefficiency in our environment arises because of the (socially) excessive levels of maturity transformation undertaken by weak banks before being bailed out by the government. By selecting the cap \bar{c} below the payment weak banks would choose in the absence of this cap, the government will bring their levels of maturity transformation closer to the socially desirable level. At the same time, a binding cap would also impose a cost on the sound banks since they will be prevented from setting c_{1S}^* in period 1. Proposition 12 shows that it can nevertheless be optimal for the government to impose a binding cap $\bar{c} < c_{1S}^*$ since the negative effect on the sound banks will be more than offset by the lower misallocation of resources associated with the weak banks.

Proposition 12 is established in Figure 3 which shows welfare in period 0 as a function of \bar{c} .¹⁵ The parameter values used for panel (a) are: $R = 1.5$, $\pi = \theta = 0.5$, $\delta = 0.5$, $\gamma = 6$, $q = 0.05$, $n = 0.2$ and $\bar{\sigma} = 0.2$. The parameters for panel (b) are the same, but with larger real losses, $\bar{\sigma} = 0.5$. When the economy has multiple equilibria, we assume that the equilibrium with a bank run is selected; the results would be qualitatively unchanged if we used some other equilibrium selection rule. We see that the optimal choice of \bar{c} - represented by the blue vertical line in each panel - is below c_{1S}^* in both panels of the figure (where c_{1S}^* is represented by the dashed red line).

¹⁵ For some parameter values weak banks choose to prevent a run by setting $c_{1W} = \hat{c}_2 < c_{1S}^*$ (see Proposition 7). In this case the optimal choice of \bar{c} would either be $\bar{c} < \hat{c}_2$ or the cap would be chosen to be non-binding at $\bar{c} = c_{1S}^*$. For the rest of the discussion in this and the next section, we restrict attention to parameter values such that the government chooses to impose a binding cap.



(a)



(b)

Figure 3: The optimal choice of \bar{c} .

The kink in welfare in each panel shows that the run equilibrium does not exist for \bar{c} below a certain level. In other words, a cap which is sufficiently low would also have a *prudential* effect by preventing the run equilibrium and in some cases the optimal choice of \bar{c} eliminates the run equilibrium as shown on panel (b) in Figure 3. Recall that a run in this environment is the result of an *externality* in payoffs across weak banks. That is, while it is optimal for weak bank k to keep its payoffs as high as possible, if all other weak banks behave in this way, then bank k ends up receiving a lower bailout payment, which in turn leads to a lower consumption allocation for the remaining investors in bank k . If this

payoff externality is sufficiently severe, patient investors will choose to run on the bank in an attempt to withdraw before the government's intervention. One way of reducing this payoff externality is by lowering the upper limit on set of contractually allowable payoffs in period 1. As shown in Figure 3, in some cases the government will be willing to introduce a sizable distortion in the sound banks in order to avoid the additional misallocation of resources resulting from a bank run.¹⁶

In addition, Proposition 12 shows that the government would choose an even lower \bar{c} if the misallocation resources stemming from the behavior of the weak banks is larger. This misallocation will be proportional to the fraction of weak banks n , the size of the losses in each weak bank $\bar{\sigma}$, the fraction of withdrawals which must take place before the government learns who are the weak banks θ and the probability of the bad aggregate state q . Also, notice that if the probability of the bad aggregate state q is larger then, other things being equal, the government will be more willing to impose a binding cap.

An even better policy in this environment is to potentially impose the cap **only** in the bad aggregate state. That is, if the aggregate is good, then all banks are sound and the cap is never binding. On the other hand, in the bad aggregate state the government can choose to impose a binding cap. Notice that such a policy will be feasible since the government observes the aggregate state. For the rest of the discussion in this and the next section, we allow the government to follow this more flexible policy and impose the cap only conditional on the bad aggregate state. This policy can be interpreted as imposing restrictions on the dividends paid out by all banks during a period of financial stress. Alternatively, one can think of \bar{c} as a contingent equity with a systemic trigger – if the aggregate state is bad (the systemic event) then all banks must bail-in their investors, even though the government realizes that not all banks are weak.

Ideally, the government would impose a cap only on the weak banks since they are the one responsible for the misallocation of the resources in the economy. Such a policy, however, is not feasible since the government does not observe which banks experienced real losses (at least not until a fraction θ of the investors have withdrawn). In this case, a weak bank would initially claim to be sound in order to be exempt from the cap and be able to give higher payouts before being placed in resolution. Given this limitation, the government has no choice but to follow a less refined approach where the cap is imposed on **all** banks.

¹⁶ Our assumption that \bar{c} is fixed in period 0 is important here. If the government were unable to commit to \bar{c} and could change its value after withdrawal decisions have been made, a time-inconsistency problem would arise as the government would often prefer to set a higher cap. Our model indicates that finding a way to commit to \bar{c} by, for example, passing regulations which embed systemic triggers or similar measures is important for financial stability. This logic is similar to that in Ennis and Keister (2009), which shows how the inability to commit to suspend payments can create financial fragility.

5.2 Increasing the tax rate

Previously the tax rate τ was set to its level in the constrained efficient allocation, τ^* . We relax this assumption in the current section and allow the amount of taxes collected in period 0 to serve as another macroprudential tool. Specifically, in period 0, the government is choosing both the cap on early payments in the bad aggregate state \bar{c} and the tax rate τ . Figure 4 shows what equilibria exist for various combinations of \bar{c} and τ . The parameter values are $R = 1.5$, $\pi = \theta = 0.5$, $\delta = 0.5$, $\gamma = 6$, $q = 0.05$, $n = 0.5$ and $\bar{\sigma} = 0.2$. The black region shows combinations of \bar{c} and τ where the equilibrium is unique and the allocation does not involve a run. The shaded region shows combinations of \bar{c} and τ where both equilibria exist – one where all weak banks experience a run and one where there is no run. Finally, the unshaded region shows combinations of \bar{c} and τ where the equilibrium is unique and such that all weak banks experience a run.

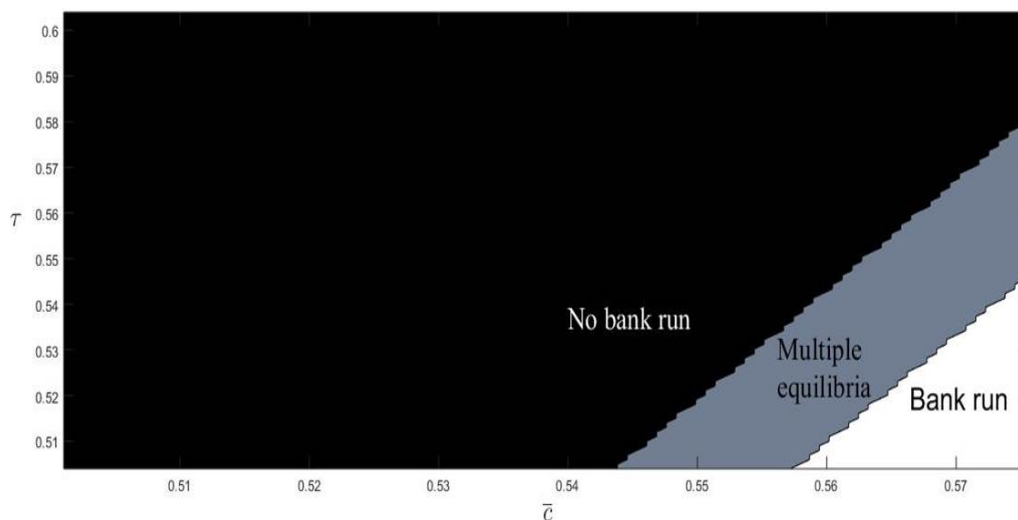


Figure 4: Fragility for different combinations of \bar{c} and τ .

Note that the policy tools \bar{c} and τ serve as partial substitutes in promoting financial stability. That is, starting in the bank-run region, the government can eliminate the run equilibrium by either collecting more taxes (and thus ensuring higher ex-post bailouts) or lowering \bar{c} (and thus ensuring that patient investors have no incentive to run on weak banks). One advantage of \bar{c} over τ is that the former will be imposed only conditional on the bad aggregate state, whereas the later must be set before the aggregate state has been realized. In general, the government would find it optimal to operate on both policy margins, that is, by setting τ above its constrained efficient level and by setting \bar{c} below c_{1S}^* .¹⁷ The reason

¹⁷For example with the parameters used to construct Figure 4, the optimal choice for the limit on early

for that is not hard to see: increasing τ above its constrained efficient level of τ^* has the benefit of not distorting the allocation in sound banks, but results in an inefficiently low level of private investment in the good aggregate state. At the same time, lowering \bar{c} below c_{1S}^* would not affect the level of private investment, but will distort the allocation in sound banks whenever the aggregate state is bad.

In period 0, the government would choose τ and \bar{c} in order to maximize the sum of investors expected utilities. We have the following result.

Proposition 13. *If $q > 0$, the optimal choice of the tax rate in the decentralized economy is above its constrained efficient counterpart. That is, $\tau^D > \tau^*$.*

Compared to the constrained efficient allocation, weak banks in the decentralized economy always choose to keep their early payments at an exceedingly high level (from a social perspective) and therefore will have fewer remaining resources available when bailout payments are made. This fact, in turn, would end up placing greater strain on the government whose objective is to use the available tax revenues τ both to support private consumption through the bailout transfers to weak banks and to provide the public good. Collecting more taxes in period 0 is one way to relax the budget constraint of the government in the bad aggregate state.

In addition, the government has an added *prudential* motive to collect more taxes in period 0. In particular, a larger tax rate τ implies that the government would have more resources and provide larger bailouts to weak banks. Patient investors, in turn, anticipate that the bailout transfers will be more generous and - provided that these transfers are sufficiently high - will choose not to run on the weak banks. Thus, increasing τ would eventually have a prudential effect by eliminating the run equilibrium. At the same time, the opportunity cost of collecting more taxes is that fewer resources will be placed in the more productive private technology. Hence the government's willingness to increase the tax rate above its constrained efficient level depends on the ex-ante probability of the bad aggregate state q . The larger q is the more willing is the government to increase the tax rate above its constrained efficient level.

5.3 Eliminating bailouts

Suppose that in period 0 the government can commit to a strict *no-bailouts* policy. That is, there will be no bailouts, even after the government discovers the bank-specific states. Note that such a rule requires commitment, since ex-post the government would find it optimal

payments is $\bar{c} = 0.5349$ (whereas $c_{1S}^* = 0.5749$) and the optimal choice for the tax rate is $\tau^{DE} = 0.5047$ (compared to $\tau^* = 0.5040$). This combination of τ and \bar{c} ensures that the run equilibrium does not exist.

to bail out weak banks not only because they might experience a run, but also because they have sustained real losses. The entire discussion in this section is therefore predicated on the ability of the government to pre-commit not to bail out weak banks. We can show that a no-bailout policy will prevent bank runs, but in many cases will lower welfare.


Proposition 14. *If the government can credibly commit not to bail out weak banks, then bank runs will not occur in equilibrium. For some parameters imposing a no-bailout rule would raise welfare, while for other parameters welfare would decrease.*

According to Proposition 14, the first implication of a (credible) no-bailout policy is that it will promote financial stability since runs can no longer occur as part of the equilibrium in the decentralized economy. The reason is that in our environment weak banks have all the information and contractual flexibility necessary to prevent runs and, in the absence of bailouts, will have no incentive to delay fully bailing-in their investors as soon as they have sustained real losses or realize that a run is underway. Proposition 14 also establishes that a no-bailout rule might improve ex-ante welfare for some parameter values. However, for other parameters eliminating bailouts would, in fact, result in inferior outcomes. The reason is that ex-post bailouts deliver a socially valuable transfer of tax revenue from the public sector to those banks that had sustained losses and thus ensuring a better allocation of resources between the investors in the weak banks and the public good. Given that banks can experience runs in addition to sustaining fundamental losses, such an ex-post co-insurance mechanism between the private and the public sector can be socially valuable even if it leads to moral hazard (as it invariably does in our environment).


Proposition 14 is demonstrated in panel (a) and panel (b) in Figure 5, where the y -axis in each panel plots the ex-ante welfare associated with the bailouts economy (the dashed red line) and with the no-bailouts economy (the black line). Focus first on panel (a) where the x -axis plots the real losses $\bar{\sigma}$ sustained by weak banks in the bad aggregate state.¹⁸ We can see from the figure that the two economies deliver almost identical welfare for relatively small values of $\bar{\sigma}$. However, as $\bar{\sigma}$ is increasing beyond approximately 0.04 the bailouts economy starts to perform better until it becomes fragile at around 0.1 (hence the kink in the dashed red line). However, for sufficiently high levels of the real losses, the positive effect of ex-post bailouts (in terms of the improved allocation of resources between the public and the private good in the bad aggregate state) is relatively large. Thus permitting bailouts delivers higher ex-ante welfare, even if weak banks experience a run in the bailouts economy. Next, focus on panel (b) on the figure where the x -axis plots the fraction of banks suffering real losses,

¹⁸ The remaining parameters for panel (a) in Figure 5 are $R = 1, 5$, $\pi = 0.5$, $\gamma = 5$, $\delta = 0.5$, $\theta = 0.5$, $q = 0.05$ and $n = 0.6$. The remaining parameters for panel (b) are $R = 1, 5$, $\pi = 0.5$, $\gamma = 5$, $\delta = 0.5$, $\theta = 0.5$, $q = 0.05$ and $\bar{\sigma} = 0.1$.

n . We can see that for sufficiently high values of n prohibiting bailouts would *both* promote stability and increase ex-ante welfare. The reason for this result is the following: holding fixed the level of real losses $\bar{\sigma}$, the benefit of promoting stability will be high if the fraction of banks the can experience a run n is high. As panel (b) on figure 5 shows, in such cases, it can it be desirable to pre-commit to a strict no-bailouts policy, even at the cost of forsaking the socially valuable ex-post function of the government’s bailouts.



(a)



(b)

Figure 5: Welfare in the no-bailouts economy.

To summarize: in section 5 we have considered three macroprudential policies designed

to address the misallocation of resources generated by the weak banks and the potential fragility of the decentralized economy: (i) restricting early payments, (ii) increasing the tax rate and (iii) eliminating government bailouts. A no-bailout policy ensures that runs do not occur but would also dispense with the socially valuable ex-post insurance function of the government's bailouts intervention. The effect on ex-ante welfare is thus ambiguous and depends on parameter values. In addition, a no-bailout policy might be plagued by government credibility issues. On the other hand, if bailouts are anticipated, then we showed how the government would - in general - choose to combine (i) and (ii) by both increasing the tax rate beyond the level in the constrained efficient economy and by imposing a cap on early payments in the bad aggregate state. With policies (i) and (ii) at the government's disposal, runs would sometimes be prevented and the misallocation of resources would be always be reduced - but never completely eliminated.

6 Conclusion

A necessary ingredient for a bank run to occur in the the bank in question be slow to react to the surge in withdrawal demand. This slow reaction is what leads investors to anticipate that the future payments made by the bank will be smaller and, hence, gives them an incentive to try to withdraw before the reaction comes. In the previous literature, the primary factors behind this slow response have been exogenously imposed rather than derived endogenously as part of the equilibrium outcome. Specifically, banks' failure to respond in a timely manner has been justified by assuming that either (i) contracts are rigid and therefore cannot be ex-post altered to deal with a run, or (ii) banks were unable to respond efficiently to a run because they were (at least initially) unaware that the run was actually taking place.

In contrast, we have presented a model of banking and government interventions where (i) banks maximize the utilities of their investors (i.e. there are no agency costs), (ii) contracts can be made fully state contingent, and (iii) banks always have sufficient information to respond in a timely and effective way to an incipient run. In common with the existing literature, a bank run in our setting can occur only when the bank's reaction to the run is delayed. However, the delayed reaction in our model is the endogenous choice of the bank, acting in the best interests of its investors. We show that this framework has a number of interesting implications. For example, banks will not have an incentive to use bail-in options and similar measures to impose discretionary losses on their investors when they anticipate to be bailed out by the government later on. As a result, the new bail-in rules might turn out to be not as effective in promoting financial stability as they were originally expected to be. We addressed some of the possible approaches to fix these weaknesses of the "bail-in

rules” and concluded that they are unlikely to solve the problem of bank runs on their own due to a combination of asymmetric information and the policy maker’s lack commitment.

References

- [1] Allen, F. and Gale, D. (2004) “Financial fragility, liquidity, and asset prices” *Journal of the European Economic Association*, **2**, 1015-1048.
- [2] Andolfatto, D., Nosal, E., and Sultanum, B. (2016) “Preventing bank runs,” *Theoretical Economics*, forthcoming.
- [3] Bianchi, J. (2012), “Efficient bailouts?” NBER Working Paper No. 18587.
- [4] Boyd, J.H., Chang, C. and Smith, B.D. (2002) “Deposit insurance: A reconsideration” *Journal of Monetary Economics*, **49**, 1235-1260.
- [5] Brown, C. and Dinç, S, (2005), The politics of bank failures: Evidence from emerging markets, *The Quarterly Journal of Economics*, 120, issue 4, p. 1413-1444
- [6] Calomiris, C.W., and Kahn, C.M. (1991) “The role of demandable debt in structuring optimal banking arrangements” *American Economic Review*, **81**, 497-513.
- [7] Chari, V.V. and Kehoe, P.J. (2016) “Bailouts, time inconsistency, and optimal regulation: A macroeconomic view” *American Economic Review*, **106**, 2458-93.
- [8] Cooper, R. and Kempf, H. (2013) “Deposit insurance and bank liquidation without commitment: Can we sleep well?” NBER Working Paper No. 19132.
- [9] Diamond, D.W. and Dybvig, P.H. (1983) “Bank runs, deposit insurance, and liquidity” *Journal of Political Economy*, **91**, 401-419.
- [10] Diamond, D.W. and Rajan, R.G. (2001) “Liquidity risk, liquidity creation and financial fragility: A theory of banking” *Journal of Political Economy*, **109**, 287-327.
- [11] Ennis, H.M. (2012) “Some theoretical considerations regarding net asset values for money market funds” Federal Reserve Bank of Richmond *Quarterly Review*, **98**, 231-254.
- [12] Ennis, H.M. and Keister, T. (2009) “Bank runs and institutions: The perils of intervention” *American Economic Review*, **99**, 1588-1607.

- [13] Ennis, H.M. and Keister, T. (2010) “Banking panics and policy responses” *Journal of Monetary Economics*, **57**, 404-419.
- [14] Farhi, E. and Tirole, J. (2012) “Collective moral hazard, maturity mismatch and systemic bailouts” *American Economic Review*, **102**, 60-93.
- [15] Goldstein, I. and Pauzner, A. (2005) “Demand deposit contracts and the probability of bank runs” *Journal of Finance*, **60**, 1293-1327.
- [16] Green, E.J. (2010) “Bailouts” Federal Reserve Bank of Richmond *Economic Quarterly*, **96**,11-32.
- [17] Green, E.J. and Lin, P. (2003) “Implementing efficient allocations in a model of financial intermediation” *Journal of Economic Theory*, **109**, 1-23.
- [18] Keister T. (2016) “Bailouts and Financial Fragility” *Review of Economic Studies*, **83**, 704–736.
- [19] Kroszner, R. and Strahan, P. 1996. “Regulatory incentives and the thrift crisis: Dividends, mutual-to-stock conversions, and financial distress.” *Journal of Finance* **51**, 1285–320.
- [20] Martin, A. (2006) “Liquidity provision vs. deposit insurance: preventing bank panics without moral hazard” *Economic Theory*, **28**, 197–211.
- [21] Peck, J. and Shell, K. (2003) “Equilibrium bank runs” *Journal of Political Economy*, **111**, 103-123.
- [22] Rogoff, K. and Sibert, A (1988) “Elections and macroeconomic policy cycles” *Review of Economic Studies* **55**, Jan. 1988, 1-16
- [23] Wallace, N. (1988) “Another attempt to explain an illiquid banking system: the Diamond and Dybvig model with sequential service taken seriously” *Minneapolis Fed Quarterly Review*, **12** (Fall), 3-16.
- [24] Wallace, N. (1990) “A banking model in which partial suspension is best” *Minneapolis Fed Quarterly Review*, **14** (Fall), 11-23.
- [25] Walther, A. and White, L. (2017) “Optimal bank resolution” mimeo, April.