Liquidity Requirements and the Interbank Loan Market:
An Experimental Investigation

Douglas Davis,[1,2] Oleg Korenok, John Lightle

Virginia Commonwealth University,

Edward Simpson Prescott

Federal Reserve Bank of Cleveland

March 4, 2019

Abstract

We develop a stylized interbank market environment and use it to evaluate with experimental methods the effects of liquidity requirements. Baseline and liquidity-regulated regimes are analyzed in a *simple shock* environment, which features a single idiosyncratic shock, and in a *compound shock* environment, in which the idiosyncratic shock is followed by a randomly occurring second-stage shock. Interbank trading of the illiquid asset follows each shock. In the simple shock environment, we find that liquidity regulations reduce the incidence of bankruptcies, but at a large loss of investment efficiency. In the compound shock environment, liquidity regulations not only impose a loss of investment efficiency but also fail to reduce bankruptcies.

Keywords: Interbank market, liquidity regulation, market experiments

JEL classification: C9, G21

---

[1] *Address for correspondence:* Department of Economics, Virginia Commonwealth University Box 4000, Richmond Virginia 23284-4000. *E-mail*: dddavis@vcu.edu. *Homepage:* https://sites.google.com/a/vcu.edu/dddavis.

[2] Financial assistance from the National Science Foundation (SES 1426980) is gratefully acknowledged. We are also indebted to the journal editor Yuriy Gorodnichenko and two anonymous referees for their useful comments. The usual disclaimer applies. This paper was presented in seminars at the University of Alabama, the University of New Mexico, the University of Pittsburgh, Rice University, the University of Tokyo, the Federal Reserve Bank of Cleveland, and the Federal Reserve Board. It was also presented at the 2017 Society for Experimental Finance Conference in Richmond, Virginia and the 2018 SAET Conference in Taipei, Taiwan. The views expressed in this paper are those of the authors and not necessarily those of the Federal Reserve Bank of Cleveland or the Federal Reserve System.

Liquidity requirements and the interbank loan market: An experimental investigation

1. Introduction

This paper reports a prototype experiment designed to study how an interbank market reallocates liquidity among banks. The experiment studies a multistage game in which banks are subject to random withdrawal shocks. Each bank initially decides how much of its deposits to invest in cash and how much to invest in a high-return, illiquid asset, under the condition that it may subsequently be subject to random withdrawal demands. The bank can meet those demands by distributing cash that it holds and/or by obtaining more cash from selling its illiquid assets in the interbank market. In half the treatments, liquidity requirements – minimum amounts of cash that banks must hold – are imposed on each bank. Investment decisions and the performance of the interbank market with and without the liquidity requirements are compared

Interbank markets are an important form of financial intermediation. Through these markets, banks with excess liquidity lend to banks that need liquidity, a process that reallocates funds and expands the lending capacity of the banking system. During the financial crisis of 2007-2009, these markets became "stressed" in that interest rates on what were usually considered safe, riskless, short-term loans greatly increased and volumes did not increase to meet increases in demand.[3] Furthermore, several large financial institutions, such as Bear Stearns, Lehman Brothers, Northern Rock, and other banks, experienced liquidity problems.

As part of the regulatory response to the crisis, the Basel Committee on Banking Supervision recommended that bank regulators mandate that banks hold a sizable buffer of liquid assets, with the hope that this would prevent the liquidity problems experienced during the

[3] Afonso et al. (2011) analyze the U.S. federal funds market for overnight loans following the Lehman Brothers bankruptcy in September 2008 and find that the daily federal funds rate spiked by a weighted average of 60 basis points. Heider et al. (2015) and Acharya and Merrouche (2013) report similar responses to the crisis in the unsecured euro interbank market and in the U.K. interbank market, respectively.

financial crisis.[4] Economic arguments for liquidity requirements are based on the belief that banks will underprovide liquidity. Bhattacharya and Gale (1987), for example, suggest that banks might free ride off each other's holdings of liquid assets because maintaining these assets is costly, thus creating liquidity shortages in times of aggregate shocks.

Potentially compounding the effects of insufficient liquidity is a propensity for banks to hoard available liquid assets in times of systemic stress. A variety of factors may motivate such hoarding, including precautionary considerations (Gale and Yorulmazer, 2013) and strategic efforts to force rivals to sell assets at fire-sale prices (Diamond and Rajan, 2011). Presumably, by reducing instances of stress, liquidity requirements may ease the negative effects of these incentives. [5]

Basel's liquidity requirements, however, are not without costs. First, required liquidity buffers limit bank lending activity, since banks must hold proportionally more cash and other liquid assets such as Treasury securities.[6] Second, it is possible that liquidity requirements will make a liquidity crisis worse by giving banks an even stronger incentive to hoard liquidity when times are bad. In Gale and Yorulmazer (2013), for example, while liquidity requirements do reduce the probability of a panic, when a crisis does happen the liquidity requirements make the situation worse by increasing banks' incentives to hold liquid assets.

---

[4] Formally, the Basel committee recommended two separate liquidity requirements. The first one is the Liquidity Coverage Ratio (LCR), which requires a bank to hold enough liquid assets to meet expected net cash outflows over a 30-day period. For details on how this ratio is calculated, see Kowalik (2013). The second one is the Net Stable Funding Ratio (NSFR), which is intended to ensure that banks adequately balance the sources and uses of funds over a longer term (one year). In the United States, the LCR has been implemented and applies only to the largest banks, while the NSFR rule has not yet been finalized.

[5] Other market frictions discussed in the literature that may impair interbank market performance in times of stress include counterparty risk (Heider et al. 2015).

[6] See, for example, Kowalik (2013, p. 76). Also, Diamond and Kashyap (2016) analyze a banking model for an institution subject to an LCR requirement and show that the regulated bank must always hold some assets in reserve, even during a crisis, to protect itself from a run, even though doing so inefficiently restricts lending.

Empirically, some evidence suggests that liquidity requirements can make panics worse. In the United States during the National Banking Era (1863-1913), liquidity requirements were one of the main prudential regulatory tools.[7] Despite these requirements, the National Banking Era was characterized by multiple panics, and experts from the era, such as Sprague (1910), believed that reserve requirements – what the liquidity requirements were called -- were partly responsible.[8]

Our experiment design studies both a *simple shock* environment where interbank trade follows an idiosyncratic shock and a *compound shock* environment where, after the initial shock, a second shock probabilistically impacts the banking system. The simple shock environment isolates the effects of a liquidity requirement on the coordination problem associated with interbank trade. The more complex compound shock environment allows insight into banks' propensities to hold liquid assets given some uncertainty about aggregate liquidity needs. In both environments, we examine the capacity of liquidity regulations to improve the stability of an interbank market in terms of reducing the number of bankruptcies, as well as the costs of the regulation in terms of aggregate investment.

2. Related Literature

This paper contributes to a small but growing literature that uses experiments to analyze banking and bank regulation. The largest branch of this literature pertains to experiments

[7] The reserve requirement took the form of requiring a bank to hold assets in the form of specie, Treasury notes, or reserves at other banks of at least 25% of its notes and deposits. The precise requirement depended on whether a bank was a Central Reserve City Bank, a Reserve City Bank, or a Country Bank.

[8] Carlson (2013) describes reserve requirements before, during, and after the National Banking Era. Notably, following the establishment of the Federal Reserve in 1913, the use of reserve requirements remained in use as a tool for promoting bank liquidity. However, the central bank's role as a lender of last resort soon usurped the role of reserve requirements as a provider of liquidity. By the 1930s, reserve requirements, along with open market operations, were seen as tools for implementing credit and monetary policy (Goodfriend and Hargraves, 1983).

examining financial fragility in variations of the Diamond and Dybvig (1983) banking model. (See for example, Davis et al. (2018) and the references therein.) These papers focus on institutional and environmental factors that affect the tendency of depositors in a single bank to ”run” or coordinate on a degenerate equilibrium and, for that reason, are not closely related to the current investigation.

Only a handful of experimental papers examine aspects of interlinked banks. Duffy et al. (2018) implements a variant of the Allen and Gale (2000) model of interconnected bans to examine the extent to which alterations between complete and incomplete network structures affect the susceptibility of a banking system to financial contagion. Choi et al. (2017) examines contagion in more complex partially connected network structures of the type analyzed in Acemoglu et al. (2015). In contrast to our investigation, where we study the process of banks making interbank linkages, these papers impose such linkages exogenously.

The most closely related paper is Davis et al. (2019), (“DKL”), who report an initial experiment examining the efficiency and stability of an interbank loan market. Their experiment design is related to the model by Allen and Gale (2004b), in which banks are impacted by idiosyncratic and systemwide aggregate shocks. DKL find that while the interbank market allows substantial improvements in trading efficiency relative to the autarkic narrow bank solution, investment efficiency remains below maximum sustainable levels in all treatments because of a persistent heterogeneity in portfolio choices by participants within and across periods. Consistent with the predictions of Allen and Gale (2004b), they also observe persistently volatile asset prices and frequent instances of bank losses following interbank exchange.

Here we build on DKL by examining the effects of a liquidity requirement on interbank market performance. Our experiment design differs from that in DKL in three main respects. First, rather than using a double auction to trade assets, we use a trading mechanism where the price is a function of cash needs and supplied cash. Second, we change the shock structure to one in which there are idiosyncratic liquidity shocks in the second stage and, for some of our treatments, a subsequent probabilistically occurring liquidity shock in a third stage. Finally, rather than examining the effects of asset price restrictions on interbank market stability, we evaluate the effects of a liquidity requirement. The first two of these differences allow us to solve for Subgame Perfect Nash Equilibria that are useful as predictive benchmarks and better allow us to identify participants' motivations for liquidity withholding behavior. The remaining difference is part of the motivation for the study.

3. Experiment Design and Procedures

3.1 Experiment Design. The experiment is based on the incomplete markets model of Gale and Yorulmazer (2013), which in turn is a variant of a frequently used model of the interbank market.[9] We study two distinct regimes, as described below.

3.1.1. The Simple Shock Regime. The simple shock game consists of three stages and is played by eight symmetric and risk-neutral banks. In an initial stage 0, each bank i, $i \in \{1, 2, \ldots, 8\}$, is endowed with \$12 in deposits and chooses a portfolio consisting of cash and assets. Assets can be purchased at a unit price P_o=\$1 in stage 0 and yield a gross return R=\$2 in terminal stage 2, for a net profit of \$1 per asset held to maturity. Assets, however, are illiquid in stage 1 and can be

[9] The basic interbank model, which adapts the banking model by Diamond and Dybvig (1983) to a banking system, was first proposed by Bhattacharya and Gale (1987). Allen and Gale (2004a) provide a welfare analysis of the interbank game. Gale and Yorulmazer (2013) analyze a more complex version of our compound shock regime. Their motivating interest was to identify conditions under which both precautionary and strategic motivations for liquidity withholding arise as subgame perfect equilibrium phenomena, so that they might explore interactions between the competing motivations. Our three-period simple shock regime (in the unregulated baseline condition) is closer in structure to the simplest case in Allen et al. (2009), where the banking system is subject only to an idiosyncratic shock. As will be seen below, in our model, while we can isolate instances of both precautionary and strategic withholding, only precautionary withholding is part of an equilibrium strategy, and that only in the compound shock environment.

converted to cash only by selling them to banks with excess cash.[10] Denote cash holdings for bank i going into stage 1 as c_{i1}, and assets as $12 - c_{i1}$.

At the beginning of stage 1, four randomly selected banks receive a shock of $8, while the remaining four banks receive a shock of $0. Following the shock, banks with excess cash are informed of the aggregate cash deficiency, d_1, and are given the opportunity to supply excess cash that will be used to purchase assets under the condition that assets will be sold at a price of P_1=$1.00 per unit if the supply of cash is less than or equal to the aggregate demand for cash, and P_1= $2.00 otherwise.

Formally, denote $y_i(d_1, c_{i1}) \geq 0$ as the amount of cash bank i makes available in the first stage if it is not hit with a shock and let $Y(d_1) = \sum y_i(d_1, c_{i1})$ be the total cash made available. Then, prices are $P_1 = \begin{cases} \$1 \ if \ Y(d_1) \leq d_1 \\ \$2 \ if \ Y(d_1) > d_1 \end{cases}$.

Following the determination of aggregate available cash, liquidity-deficient banks sell assets, with sales rotated among banks until either $d_1=0$ or $Y(d_1)=0$.[11] Any bank with a cash deficiency following asset sales becomes bankrupt and must pay a $4 liquidation fee. Otherwise,

---

[10] In practice, banks often borrow to meet liquidity needs. Here we follow Allen et al. (2009), by having banks meet withdrawal demands by selling assets rather than through borrowing. Given homogeneous, riskless assets, this simplifying convention does not affect analysis and allows us to avoid the complication of introducing a loan market into the experiment.

[11] If $Y(d_1) \leq d_1$, then the rotation works by sequentially allocating one unit of cash to liquidity-deficient banks until the cash runs out. The rotation order is determined randomly each period, so in the case that cash is insufficient to make all liquidity-deficient banks solvent, those banks that become insolvent are determined randomly. If $Y(d_1) > d_1$, then all liquidity-deficient banks remain solvent. In this case, the sequence of asset purchases is again determined randomly. Unlike the case of insufficient demand, however, the order of purchases is a matter of indifference to banks. Since the purchase price just equals the value of the asset at maturity, banks are indifferent between purchasing assets and holding cash to return to depositors at the end of period 2. In the game variant with a probabilistic second-stage shock, discussed below, the conclusion that given $Y(d_1) > d_1$ banks are indifferent between buying assets at $P_1 = R$ and holding cash may not remain true, since residual cash has the advantage of leaving a bank more prepared to address a second-stage shock. Importantly, however, changes in the relative desirability of purchase decisions off the equilibrium path do not affect banks' incentives to keep $Y(d_1) \leq d_1$.

in stage 2, assets mature, deposits are repaid, and earnings are determined. Figure 1 provides a schema of the simple shock environment. We label this the simple baseline shock game SB.[12]

3.1.1.1. The Trading Mechanism. In practice, interbank trade typically occurs as bilateral over-the-counter (OTC) exchanges, an institutional practice made necessary by the heterogeneous nature of bank portfolios. Moreover, these exchanges take place among banks of asymmetric size and potential access to counterparties. Enriching the design to include these features considerably complicates the analysis and, we fear, incentives in the subsequent environment. For that reason, as a starting point, we use a uniform price trading rule with symmetric banks, as is often seen in the theoretical literature (see, for example, Allen and Gale 2004a, 2004b).

Our uniform pricing alternative, while stylized, allows us to solve for equilibria while eliminating the behavioral complications associated with participants being forced to learn the underlying supply and demand conditions for cash through the price discovery process, as would occur in a double auction. [13] Finally, from a very practical perspective, our sequential mechanism speeds the trading process and allows a substantially larger number of trading periods than would be possible with, for example, a double auction. That said, we recognize that our choice of trading institution may limit the potential applicability of results to some markets and further work to investigate the importance of the trading design is warranted.

3.1.1.2. Baseline Simple Shock Equilibria. To provide reference predictions for experimental results, we solve for the symmetric subgame perfect Nash equilibrium (SPNE) for

[12] A fixed liquidation fee follows the approach taken in Gale and Yorulmazer (2013) (who assumed that the value of the bank in the case of default is 0). It also simplifies both the identification of equilibria and the analysis of results. See DKL (2019) for a roughly parallel design that imposed losses equal to cash deficiencies.

[13] A natural choice of trading mechanism is the double auction, which is well known for its high trading efficiency and robustly competitive performance (See e.g., ch.3 in Davis and Holt, 1993). The double auction, however, also presents significant limitations here. DKL (2019) evaluate an interbank market environment similar to that examined here but use a double auction for asset sales. They find that the sequential nature of contracting in the double auction conflates learning about underlying supply and demand conditions with possible strategic efforts to manipulate the terms of trade.

this game. In the symmetric SPNE each bank holds $c_{i1} = \$4$ as cash. Proofs appear in Appendix A, but the result can be seen intuitively.[14] By holding \$4 in cash and investing the remaining \$8, banks maximize earnings and collectively avoid bankruptcy. In stage 1, the four banks experiencing a shock have a deficiency of \$4 each, making $d_1 = \$16$. The symmetric SPNE strategy is for each unshocked bank to make \$4 available, so $Y(\$16) = \$16$. To see this, first note that given $c_{i1} = \$4$, the expected earnings for each bank is $\pi_i = \frac{1}{2}(\$4) + \frac{1}{2}(\$12) = \$8$. Now consider a deviation by a bank to $c_{i1} < \$4$ in stage 0. In this case, at stage 1, the amount of cash supplied depends on whether the deviating bank is hit by the liquidity shock. If it is, then $d_1 > 16$. The unshocked banks will supply all their cash because, if they did not, they would be passing up a chance to buy assets at P_I=\$1. Consequently, $Y(d_1) = \$16 < d_1$, P_I=\$1, and the deviating bank goes bankrupt and receives a payoff of -\$4. If the deviating bank is not hit by the liquidity shock, then again the banks not hit by the shock supply all their cash because assets are cheap. Consequently, $Y(d_1) = \$12 + c_{i1} < d_1 = \$16$, P_I=\$1, and the deviating bank receives a payoff of $2*(\$12-c_{i1}) + 2 * c_{i1} - \$12 = \$12$. The expected payoff is then $\pi_i = \frac{1}{2}(-\$4) + \frac{1}{2}(\$12) = \$4$. In other words, banks collectively provide exactly the amount of cash needed in the market up to the total amount they have available. Basically, banks have no incentive to save less than \$4 because doing so risks bankruptcy and the gains are no more than what a bank could get by holding 4 units of cash and buying assets in the interbank market. Conversely, banks have no incentive to save more than \$4 because cash cannot be more valuable than an asset, so the expected payoff cannot be higher than in the symmetric SPNE.

[14] Appendix A presents proofs of the symmetric SPNE for this and each of the treatment combinations. In each case, we focus on an SPNE in which actions taken along the equilibrium path are ex-ante symmetric across banks. We also establish in Appendix A that asymmetric substrategies are needed off the equilibrium path, similar to the asymmetric equilibria in the game of chicken.

While we focus on a symmetric SPNE to organize our analysis, there is a wide variety of initial investment decisions that can be supported as asymmetric subgame perfect Nash equilibria. For example, if bank A holds \$5 in cash, bank B holds \$3 in cash, and the rest hold \$4, no profitable deviation is possible because bank A's extra cash just compensates for bank B's cash deficiency. In fact, as shown in Appendix A, an efficient equilibrium exists for any combination of portfolio decisions that yield an aggregate of \$32 in cash. For this reason we focus on aggregate investment outcomes in the results, a focus that captures the propensity of banks to coordinate on any efficient equilibrium.

In an efficient SPNE, neither overinvestment nor liquidity hoarding occurs. Off the equilibrium path, however, individual banks may find it profitable to overinvest if they believe that the rest of the banking system will collectively underinvest. Similarly, off the equilibrium path, banks have a strategic motive to hoard liquidity. For any deviation where $Y(d_1) > d_1$, banks may profitably withhold available cash in order to reduce the supply of cash on the market, so P_1=\$1 rather than P_1=\$2. We classify cash withholding in light of aggregate cash need as '"trategic" hoarding. As will be seen in the results, we see some evidence of such behavior.

3.1.2. The Compound Shock Regime. In order to evaluate the potential amelioratory effects of liquidity requirements in an environment where banks feel compelled to retain cash for precautionary purposes due to some uncertainty about aggregate liquidity needs, we append to the simple shock environment a possible second-stage liquidity shock. To implement this shock, we follow the approach taken in Gale and Yorulmazer (2013). Following the exchange of assets in stage 1, there is a second stage in which, with probability ½, two of the four banks not impacted by the stage 1 shock experience a liquidity shock of $8 each.[15] Which two banks receive the shock is random. In the event of no shock, stage 2 passes to stage 3 and the period ends.

In the event of a stage 2 shock realization a second round of interbank trade follows. Banks with excess cash are shown the aggregate cash deficiency d_2 and submit an amount of cash to make available. As in the stage 1 market, the asset price will equal $1 if the aggregate supply of cash is less than or equal to the cash deficiency, and $2 otherwise. Formally, denote the cash held by bank i going into stage 2 as c_{i2}, the cash bank i makes available upon observing d_2 as $z_i(d_2, c_{i2})$, and the aggregate cash made available in the second stage as $Z(d_2)$. Then, as in the first stage, the price is determined by the relative supply and demand for cash by $P_2 = \begin{cases} \$1 \; if \; Z(d_2) \leq d_2 \\ \$2 \; if \; Z(d_2) > d_2 \end{cases}$.

As in stage 1, following the determination of available cash, liquidity-deficient banks sell assets, with sales rotated among firms until either $d_2=0$ or $Z(d_2)=0$. Any bank with a cash deficiency following the asset exchange becomes bankrupt and must pay a $4 liquidation fee. Otherwise, in stage 3, assets mature, deposits are repaid, and earnings are determined. Figure 2 illustrates this four-stage game, which we call the compound shock (CB) game.

3.1.2.1 Baseline Compound Shock Equilibria. As with the SB game, we analyze the symmetric SPNE to provide a benchmark against which we compare the experimental outcomes. Also, as with the SB game, SPNE support many asymmetric portfolios that yield the same aggregate level of initial investment as the symmetric SPNE. (However, unlike the simple shock environment, not all such asymmetric portfolios are SPNE. See Appendix A). In the CB game,

[15] We follow Gale and Yorulmazer (2013) and expose only initially unshocked banks to a second stage shock. It is possible to generate identical predictions by exposing all banks to a second stage shock of $4, but that shock structure has some undesirable features. Although our choice of this alternative design does not affect the reasoning supporting symmetric SPNE calculations, it reduces the autarkic reference investment level to 0 assets, since banks may be shocked twice. Further, in the symmetric SPNE this alternative shock structure offers no definitive equilibrium quantity prediction following the second-stage shock, since cash needs are a function of the shock realization and may range between $0 and $16. Finally, if banks are exposed to both first- and second-stage shocks, they may earn $0 in the efficient equilibrium, a possibility that may impede coordination on that outcome.

two symmetric SPNE exist. There is a no-exposure equilibrium in which all banks hold cash sufficient to insulate themselves against a bankruptcy in the event of a second shock, and there is an exposure equilibrium, in which all banks choose to stand exposed to bankruptcy in the event of a second shock.

In the no-exposure equilibrium, each bank holds $c_{i1} = \$6$, in cash. The reasoning driving this equilibrium is analogous to that for the SB game and can be seen intuitively. In stage 1, the four banks experiencing a shock each have a deficiency of \$2, so $d_1 = \$8$, and the symmetric prediction is that each bank without a shock makes \$2 available, leaving the unshocked banks with \$4 each. In stage 2, a second shock occurs with a 50 percent probability. In the case of a second shock, the two banks experiencing an \$8 demand for cash each have \$4 and would thus need \$4 more, so $d_2 = \$8$, which matches the \$4 made available by each of the remaining two banks resulting in $Z(\$8) = \$8$. The symmetric SPNE requires the off-path behavior of $Y(d_1) = \min\{d_1, \$8\}$ and $Z(d_2) = \min\{d_2, \sum_i c_{2i}\}$.

Notice that in this equilibrium, there is no general underinvestment. Banks, however, do withhold liquidity for precautionary reasons in stage 1. Given their initial savings decisions, those banks not impacted by the stage-1 shock maintain \$4 in excess cash out of concern for the possibility of a second shock. On the equilibrium path, the trading price of assets is \$1 in both periods, and no bank can increase profits by saving less cash or making more cash available. Banks earn \$1 for every asset held in period 4, and the expected profit in equilibrium is: $\frac{1}{2}(\$4) + \frac{1}{4}(\$8) + \frac{1}{8}(\$12) + \frac{1}{8}(\$4) = \$6$.

In the exposure equilibrium, banks choose to bear the bankruptcy costs of a second shock and hold $c_{i1} = \$4$ in cash, as in the simple shock environment. In this equilibrium, the four banks experiencing a shock in stage 1 have a deficiency of \$4, so $d_1 = \$16$, and the symmetric prediction is that each bank without a shock makes available all their cash, making $Y(d_1) = d_1$, leaving no cash in the banking system. If a second shock occurs in stage 2, the two shocked banks go bankrupt, leaving them each with a loss of $4. In expectation, the payout from the higher investment levels in periods with no second shock compensates for bankruptcies in the event of a second shock, and expected earnings are: $\frac{1}{2}(\$4) + \frac{1}{4}(\$12) + \frac{1}{8}(\$12) + \frac{1}{8}(-\$4) = \$6$, the same as in the no-exposure equilibrium.

Importantly, the exposure equilibrium is not an artifact of our parameter choices but is a general feature of this game. Although increases in bankruptcy costs reduce the expected profitability of the exposure equilibrium, even the relatively high 25% probability of needing cash after a second-stage shock used here damps expected profits only mildly. Quite large bankruptcy costs are needed to eliminate the equilibrium. Given our other parameters, an exposure strategy remains an equilibrium strategy until the cost of bankruptcy exceeds $20.[16]

3.1.3. The Liquidity Requirement. We examine the simple shock and compound shock environments both with and without liquidity requirements. The idea driving liquidity regulations is that banks maintain liquid assets sufficient to cover a short-term liquidity drain. An extreme interpretation of such a requirement in our environment would be to require each bank to hold $8, since in both the simple shock and compound shock regimes banks may need $8. Such an interpretation, however, is uninteresting in that it effectively eliminates the interbank market, which is certainly not a policy objective.

[16] To see this, observe that the most profitable deviation from an {8 asset, $4} portfolio in the exposure equilibrium is the autarkic {4 asset, $8} portfolio, which yields a profit of $4 with certainty each period. Labeling the bankruptcy liquidation cost as B, it follows that the exposure equilibrium will exist until ½(4) + ¼ (12) + 1/8(12) + 1/8 (-B)<4, or B>20. More generally, we observe that in natural contexts the probability of a second shock (some nonidiosyncratic event) is substantially less than ½, making the cost of bankruptcy necessary to eliminate the exposure equilibrium much higher. Labeling p as the probability of a second shock, the pertinent condition becomes ½(4) + ½ (1-p) (12)+ ¼ p (12) + ¼ p (-B)<4. Solving B> 16/p -12. Thus, for example, a second shock probability of 1/8 would imply B>$116.

As an alternative, we require all banks to hold \$4. Although a \$4 requirement does not fully insulate banks from liquidity risk, it does guarantee that the required liquidity in the banking system matches aggregate liquidity needs. Alternatively, viewing the liquidity requirement as a means of stabilizing the banking system in the case of a probabilistically occurring second shock, a \$4 requirement just equals the expected value of the second-stage shock. Finally, to sidestep the rather obvious concern that liquidity requirements undermine stability simply by requiring banks to hold cash that is unavailable in the case it is needed, we allow banks to use required reserves for their own liquidity shocks. Banks, however, may not use required cash holdings to resolve the liquidity needs of other banks.[17]

3.1.3.1 Equilibria with Liquidity Requirements. The restriction that banks must hold cash alters the symmetric SPNE predictions in both the simple and compound shock environments. In the liquidity-restricted simple shock (SL) game, the initial SPNE savings decision is $c_{i1} = \$6$. In period 1, the four shocked banks have a need of \$2, so $d_1 = \$8$, and the remaining four banks have \$6 each, but because the LCR = \$4, they each have only \$2 available, which exactly equals the demand for cash, so $Y(\$8) = \$8$. As in the SB treatment, banks respond to a deviation by providing exactly the amount of cash needed up to the amount they can provide given the LCR. Equilibrium expected profits are $\frac{1}{2}(\$4) + \frac{1}{2}(\$8) = \$6$.

In the liquidity-restricted compound shock (CL) game, there are two SPNE. In the no-exposure equilibrium, initial cash holdings increase to $c_{i1} = \$7$. In stage 1, the four shocked banks need \$1 each, and in equilibrium, the unshocked banks provide exactly \$1, leaving them with \$6.

[17] A prominent concern with the initial implementation of the liquidity requirement imposed by the LCR in the 2010 Basel III Accord was that institutions could not use their required buffer even in the event of stress, and thus would have to maintain excess buffers to address liquidity needs. To ease these concerns, buffer access conditions were clarified with Basel III revisions in 2013. Even with the modifications, however, buffer access is quite clearly limited to instances in which a bank suffers liquidity stress. See Kowalik (2013) for a description.

If a stage-2 shock occurs, the two shocked banks need \$2 each, and the two remaining banks have exactly \$2 available. Similar to the no-exposure equilibrium in the CB game, off-path actions must satisfy $Y(d_1) = \min\{d_1, \$4\}$ and $Z(d_2) = \min\{d_2, \sum_i \max\{c_{i2} - 4, 0\}\}$. Once again, banks have a precautionary motive to hoard cash in period 1 in case they are shocked in period 2. Expected earnings are: $\frac{1}{2}(\$4) + \frac{1}{4}(\$6) + \frac{1}{8}(\$8) + \frac{1}{8}(\$4) = \$5$.

In the exposure equilibrium, initial cash holdings are $c_{i1} = \$6$. Following the stage 1 shock, the four shocked banks each need \$2, which they acquire from the unshocked banks, which then each have \$4. No cash remains available, so in the case of a stage-2 shock, the shocked banks go bankrupt. Expected earnings in the exposure equilibrium are $\frac{1}{2}(\$4) + \frac{1}{4}(\$8) + \frac{1}{8}(\$8) + \frac{1}{8}(-\$4) = \$4.50$.

As with the CB game, the existence of the exposure equilibrium is a fairly general feature of the CL game and exists as long as bankruptcy costs do not exceed \$8.[18] Our parameter selection, however, does give liquidity restrictions a best shot at reducing the incidence of bankruptcies in the compound shock environment. In contrast to the CB game, where expected earnings in the exposure and no-exposure equilibria are the same, at \$6.00, in the CL game, the exposure equilibrium is less attractive in the sense that equilibrium expected earnings are \$4.50, less than the \$5.00 expected earnings available in the no-exposure equilibrium.[19]

[18] Observe that the most profitable deviation from a {6 asset, \$6} portfolio in the exposure equilibrium is the autarkic {4 asset, \$8} portfolio, which yields a profit of \$4 with certainty each period. It follows that the exposure equilibrium will exist until ½(4) + ¼ (8) + 1/8(8) + 1/8 *(-B)*<4, or *B*>8. As with the CB game, reducing second-stage shock probabilities increases the minimum bankruptcy cost necessary to eliminate the exposure. Given a second-stage shock probability of *p*, the exposure equilibrium exists until ½(4) + ½ (1-*p*) (8)+ ¼ *p* (8) + ¼ *p* (-*B*)<4. Solving, *B*> 8/*p* -8. Thus, for example, a second-stage shock probability of 1/8 would imply *B*>\$56.

[19] Although the existence of both exposure and no-exposure equilibria are general features of both baseline and liquidity regulated compound shock regimes, we do not claim liquidity regulation generally reduces expected earnings in the no-exposure equilibrium relative to those in the exposure equilibrium.

[20] Instructions are available in the unpublished Appendix B.

Table 1 summarizes reference equilibrium predictions for the four games that make up our experiment. The benchmark aggregate initial investment predictions for both the simple shock treatments and the no-exposure equilibria in the compound shock treatments define the maximally sustainable investment levels in the sense that these are the levels of investment that banks may collectively sustain without running the risk of bankruptcies. We refer to these as the sustainable equilibria, which we distinguish from the exposure equilibria in the compound shock treatments, where bankruptcy is predicted, and the autarkic allocation, in which each bank fully insulates itself against shocks.

4. Experiment Procedures

The experiment consisted of 12 sessions, three in each of the four treatment cells identified in Table 1 (SB, CB, SL, and SL). In each session, a cohort of 16 participants was randomly seated at visually isolated computer terminals and given a printed set of instructions. A monitor then read the instructions aloud, assisted by a copy projected on a screen at the front of the lab, as participants followed along on their printed copies.[20] Following the instructions, participants completed a short quiz to determine their understanding of the rules. Any errors in the quiz answers as well as all other participant questions were addressed privately by the monitor. Participants were then also anonymously divided into two 8-player markets. As explained to the participants, these markets remained fixed throughout the session, creating two independent markets per session.

Sessions were programmed with the z-Tree software (Fischbacher, 2007) and included a total of 25 trading periods. The first five periods of each session were practice periods conducted in the SB treatment to familiarize participants with the environment. In the first two of these

[20] Instructions are available in the unpublished Appendix B.

practice periods participants were not paid for their decisions, but were free to privately ask any questions they might have about market procedures. These were followed by three additional SB periods for which subjects were paid. After the fifth period, the session was paused and instructions for one of the four treatment sessions were distributed and read aloud.[21] Following a short review and quiz to determine understanding, we conducted two unpaid practice periods in the treatment condition followed by 18 paid periods that constitute experiment results. At the end of the 20th treatment period, the session ended, and participants were privately paid and dismissed one at a time.

Participants were 192 undergraduates enrolled at Virginia Commonwealth University in the spring semester of 2017 and were recruited with the ORSEE recruiting system (Greiner, 2015). Most were upper-level engineering, math, business, and economics students. Sessions lasted 60 to 90 minutes. Lab dollar earnings were converted to U.S. currency at a 1 lab dollar = $0.20 U.S. rate. Earnings ranged from $15.60 to $45.00 and averaged $29.20.

5. Experiment Results

We organize this section by first considering investment decisions and bank stability (the incidence of bankruptcy), which are two aspects of the banking system that are of policy interest. We then combine these results to analyze what determines the number of assets that reach maturity, that is, the number of successful investments. We close this section with a discussion of the way liquidity requirements impact liquidity management decisions across treatments.

[21]One of the four treatments was the SB treatment. In these sessions SB instructions were repeated prior to period 6 even though conditions were the same as in the first five periods.

5.1 Initial Investment Decisions. The lightly shaded bars in Figure 3 report mean initial investment decisions in the simple and compound shock regimes.[22] Notice generally that banks rely heavily on interbank trade in making investment decisions. In all treatments, banks collectively invest at substantially higher rates than is sustainable under the autarky allocation of 32 assets. Furthermore, note that in each environment, liquidity requirements reduce investment levels.

For the simple shock environment, in both the SB and SL treatments, mean initial investments lie below the equilibrium benchmark predictions. This result was largely expected in this environment, since the idiosyncratic shock affects the banking system by a constant amount each period, and in the SPNE, banks use all of their cash.

On the other hand, in the compound shock treatments, unregulated banks exhibit a modest propensity to overinvest relative to the no-exposure equilibrium. In the CB treatment, banks collectively invest about two assets above the no-exposure equilibrium prediction. Nevertheless, investment in the CB treatment remains far below the exposure equilibrium benchmark, suggesting that banks do not coordinate on this alternative equilibrium. Similarly, in the CL treatment, mean investment also exceeds the no-exposure equilibrium benchmark by about two units, suggesting that while the regulations reduce investment activity, they fail to reduce banks' propensities to invest at above sustainable levels.

To more formally evaluate investment activity, we regress initial investment levels against a series of indicator variables D_j, $j \in \{C, L\}$ that delineate the incremental effects of the shock environment (C for compound) and regulatory regime (L for liquidity regulated). Using the SB treatment as the default condition, combinations of these variables allow us to distinguish each of

[22] We report results based on observations pooled across the 18 paid periods of each session. We pooled the observations because we found little evidence that equilibria changed over time within sessions as participants got more experience with the treatments. Time series of investment activity and bankruptcies are reported in Appendix C.2. More formally, tests for differences in investments or bankruptcy rates across session suggest no important adjustment effects. See Appendix C.1, tables C.1.2b, C.1.4b, C.1.5b and C.1.7b.

the remaining three treatment/shock realization conditions. The regression uses a random effects specification. Specifically, we estimate

$$y_{it} = \beta_0 + \beta_L D_L + \beta_C D_C + \beta_{LC} D_L D_C + \varepsilon_i + u_{it} \quad (1)$$

where y_{it} denotes collective initial investment, i identifies a market (1 to 24) and t a period (1 to 18). We cluster data by markets and use a robust (White "sandwich") estimator to control for possible unspecified autocorrelation or heteroscedasticity.

Treatment averages are generated by adding coefficient estimates, and Wald tests are subsequently used to assess differences across treatments.[23] Table 2 summarizes the results. Column (2) lists mean investment levels for each treatment cell. In each case we can reject the null hypothesis that initial investment does not differ from the 32 assets available from the autarky allocation (p<0.01 in each case).

Comparing initial investments with the sustainable equilibrium benchmarks in each treatment, observe that while initial investment decisions across treatments move with these reference benchmarks, banks nevertheless deviate systematically from these predictions. In the simple shock environment, banks collectively under-invest relative to the equilibrium benchmark by 1.92 units in the SB treatment and 2.94 assets in the SL treatment. In the compound shock treatments, banks collectively overinvest relative to the no-exposure benchmark, by 2.15 units in the CB treatment, and 2.34 units in the CL treatment. Nevertheless, while initial investment in the compound shock treatments exceeds the no-exposure benchmark predictions, banks exhibit little tendency to coordinate on exposure equilibria. Mean investment is 13.85 assets below the exposure equilibrium prediction in the CB treatment and 5.66 units below the comparable prediction in the CL treatment (both differences significant at p<0.01).

---

[23] Primary regression results appear as Table C.1.1 in Section C.1 of Appendix C.

The across-treatment differences in the right side of Table 2 highlight the high costs of the liquidity requirement in terms of forgone investment. In the simple shock environment, the liquidity regulations reduce initial investment by 17.02 assets, more than half of the 32-asset gain that interbank trade allows compared with autarky. Similarly, liquidity regulation cuts initial investment by 7.81 assets in the compound shock environment, only slightly less than half the 16-asset gain in investment that interbank trade allows compared with autarky.

In summary, interbank exchange motivates banks to invest at levels above those consistent with autarky. In the simple shock treatments, banks exhibit no tendency to invest at above-sustainable equilibrium levels and, in fact, tend to hold too much cash on average. In the compound shock treatments, although banks do exhibit a modest propensity to overinvest relative to the no-exposure equilibrium benchmark predictions, they do not coordinate on exposure equilibria. Finally, the liquidity requirement dramatically reduces initial investment in both the simple shock and compound shock regimes, reducing in each case the potential gains from exchange over autarky that interbank trade allows by roughly 50 percent.

5.2 Bankruptcies. The bars in Figure 4 illustrate the average incidence of bankruptcies. Looking first across the figure generally, observe that bankruptcies occurred frequently. An average of at least one bank suffered bankruptcy every second period in each treatment. Turning to the simple shock environment, notice that despite failing to improve the efficiency of initial investment decisions, liquidity regulations do noticeably reduce the incidence of bankruptcies. The mean incidence of bankruptcies in the SL treatment is considerably lower than in the SB counterpart.

In the compound shock treatment, we distinguish stage 1 and stage 2 bankruptcies with light and dark gray bars, respectively. To coherently reflect the consequences of a second-stage shock realization, stage 2 bankruptcy rates are calculated *only* for those periods where a second-stage shock occurred. Notice first that, as in the simple shock environment, liquidity restrictions in the compound shock treatment do reduce the incidence of first-stage bankruptcies. The regulations, however, do nothing to improve the capacity of banks to respond to a second-stage shock. Although second-stage bankruptcy rates are high in the CB regime, these rates nearly double in the second stage of the CL treatment.

To assess quantitatively the effects of a liquidity requirement on bankruptcies, we run a series of OLS regressions that estimate the number of bankruptcies occurring per period following the first-stage shock, following the second-stage shock, and overall. For first-stage bankruptcies, we regress bankruptcies per period against the indicator variables in equation (1). In the simple shock treatments as well as in those periods in the compound shock treatment where no second-stage shock subsequently occurred, all bankruptcies occur following the first-stage shock, so first-stage bankruptcies are equivalent to total bankruptcies. In those compound shock treatment periods with a realized second-stage shock, we regress both second stage and overall bankruptcies against the first two variables in (1). In all regressions we again cluster the data by markets and use a robust (White "sandwich") estimator to control for possible unspecified autocorrelation or heteroscedasticity.

Table 3 reports treatment averages as linear combinations of coefficients from primary regression results.[24] In the simple shock environment, liquidity-regulated banks suffered 0.35 fewer bankruptcies than did unregulated banks in the baseline treatment (difference significant at $p<.01$). In the compound shock environment the incidence of bankruptcies again falls in the liquidity-regulated regime following the first-stage shock, although the difference across the CB and CL treatments is smaller than in the simple shock environment, at 0.22 fewer bankruptcies,

[24] Primary regression results appear as Table C.1.3 in Appendix C.

and not significantly different from zero.[25] Following the realization of a second-stage shock, however, the incidence of bankruptcies is considerably higher in the CL treatment than in the CB counterpart, at 0.52 banks (a difference significant at $p<.05$). The very high incidence of bankruptcies following the realization of a second-stage shock erases entirely the beneficial effects of liquidity regulation observed in the simple shock environment. In fact, the incidence of bankruptcies in compound shock treatment periods where a second-stage shock occurred is 0.30 banks higher in the CL treatment than in the CB treatment, although this difference is not significant.

In summary, we find that in the simple shock environment, liquidity restrictions do reduce the incidence of bankruptcies. In the compound shock environment, the liquidity regulation may again reduce the incidence of bankruptcies following a first-stage shock, although the effect is weaker than in the simple shock environment. The overall incidence of bankruptcy in the compound shock environment, however, is no lower in the liquidity-regulated treatment than in the baseline, because banks in the CL treatment respond so poorly to the realization of a second shock.

5.3. Mature Investments. The dark gray bars in Figure 3 report the quantity of successful investments, that is, investments that are made and not liquidated due to bankruptcy, in each treatment. We refer to these successful investments as mature investments. Like initial investments, mature investments move with the sustainable equilibrium benchmarks for each treatment. Unlike initial investments, however, mature investments fall below maximum sustainable levels in all treatments, and often by large amounts. As is obvious in Figure 3, the very sizable reductions in initial investments that the liquidity requirements impose dwarf the modest maturity rate improvements that the liquidity requirements generate. On net, in both environments,

---

[25] The difference in bankruptcy rates between the CL and CB regimes just misses significance ($p<.105$).

liquidity restrictions reduce the potential gains from exchange over autarky that interbank trade allows by roughly 50%.[26]

5.4 Liquidity Regulations and Liquidity Management Decisions. The above results raise a series of questions, the most pressing of which regard bankruptcies. First, in the simple shock treatment, what explains the high incidence of bankruptcies when collective initial investments on average are less than the maximum sustainable level? Second, in both the simple shock environment and in the first stage of the compound shock environment, why does the liquidity requirement reduce the incidence of bankruptcies? Third and finally, why are liquidity-regulated banks so frequently less prepared to respond to a second-stage shock than in the counterpart baseline treatment? This subsection addresses these interrelated questions.

Consider first the factors that drive bankruptcies in the simple shock environment. The interbank market is distinct from more standard investment games in that investment decisions, combined with subsequent shocks, create both the supply of and demand for cash. In this context, banks face an equilibrium selection problem analogous to that presented in a multiagent Nash demand game, albeit in less transparent conditions. Banks collectively maximize profits only when they initially invest exactly at the maximally sustainable level and then fully supply available cash following the shock. Collective overinvestment and/or any withholding of available cash yields a collective cash deficiency post-shock, which drives asset prices to their minimum, forcing bankruptcies in the process. Collective underinvestment, with an oversupply of available cash, yields an aggregate cash surplus, driving asset prices to their upper bound and reducing to zero the return on cash acquired through interbank exchange.

[26] A quantitative development of this result appears as Table C.1.5a in Appendix C.

In light of this complex equilibrium selection problem, the efficient solution has little explanatory power despite being an SPNE. The distributions of realized net needs for the simple shock treatments shown in Figure 5 illustrate. Realized "net needs" reflect the aggregate need for reserves following a shock and the subsequent trade of cash for assets. Net needs of zero indicate an efficient outcome, negative net needs indicate an excess of available cash, and positive net needs reflect instances of collective cash deficiencies and the consequent bankruptcies. As can be seen in the figure, in both the SB and SL treatments collective investment decisions are quite disperse and resulted in the efficient solution in no more than 25% of periods.

Coordination failures drive the high incidence of bankruptcies in the simple shock environment. Comparing realized net need distributions across treatments further suggests why liquidity requirements reduce bankruptcies in the simple shock environment. Net need realizations in the SB treatment are considerably more variable and thus yield more periods with larger cash deficiencies. Liquidity restrictions contract the effective investment strategy space for banks from [4, 12] assets to [4, 8] assets.[27] Although the equilibrium selection problem persists in the SL treatment, the reduction in the investment strategy space reduces the range of individual investment decisions and, in this way, the range of collective investment outcomes.[28]

Also note that the distribution of realized net needs for the SB treatment is shifted slightly to the right of that for the SL treatment. The increased variability of investment outcomes may also contribute to this rightward shift, since the high variability may induce increased caution in the

[27] In the experiment sessions, we restricted investment to be at least 4 assets because investing fewer than 4 assets is strictly dominated in expectation, and allowing this option would only generate an added variability induced by potential confusion.

[28] The CDFs for individual investment decisions, shown in Appendix C.3 illustrate the effects of liquidity restrictions on the heterogeneity of individual investment decisions. The net needs charts shown as Figures 5 and 6, which are also useful for other purposes, allude to the same heterogeneity, but aggregate individual decisions into markets and net out the effects of strategic cash withholding.

cash availability decisions on the part of those banks with cash post-shock, which have an interest in keeping the post-shock asset price at $1.

Linear probability estimates of realized cash deficiency incidences in the SB and SL treatments listed in the top two rows of Table 4 provide a summary measure of the effects of liquidity regulation in the simple shock environment.[29] As shown in the table, the banking system experienced an aggregate cash deficiency in 51.85 percent of SB treatment periods, compared with 29.63 percent of SL treatment periods (difference significant at $p<0.01$).

Two factors contribute to the frequency of realized net needs shown in the first column. First, banks often overinvest, which results in cash shortages. Second, banks with cash available tend to withhold it out of a strategically motivated sort of hoarding behavior intended to keep asset prices low.[30] The estimates of initial and withholding-induced net need incidences shown in the second and third columns of Table 4 reveal that, to a large extent, the higher incidence of bankruptcies in the SB treatment is driven by initial investment decisions. As seen in column 2, following initial investment decisions, banks collectively held insufficient reserves in only 14.8 percent of SL periods, compared with 38.89 percent of SB periods (difference significant at $p<0.01$).

Turning to the third column in Table 4, observe that the incidence of increased cash deficiencies due to post-shock withholding is also somewhat higher in the SB treatment (28.70 percent in the SB treatment vs. 21.20 percent in the SL treatment). The difference, however, is

[29] These estimates use the combinations of indicator variables in equation (1) as regressors. As with Tables 2 and 3, the treatment averages reported in Table 4 are generated by assembling coefficient estimates, and Wald tests are subsequently used to assess differences across treatments. Primary regression results appear as Table C.1.6a of Appendix C.

[30] Note that the incidences of initial and withholding-induced cash deficiencies do not sum to the incidence of realized cash deficiencies. A withholding-induced deficiency arises when the total cash deficiency in a period exceeds an initial cash deficiency. Thus, instances of initial and withholding-induced cash deficiencies can occur in the same period.

comparatively small and not statistically significant. Thus, we conclude that in the simple shock environment, the high incidence of bankruptcies is driven primarily by failures to coordinate initial investment decisions. Liquidity restrictions reduce the variability of collective investment outcomes and, as a consequence, reduce the incidence of periods where the cash needs of banks go unsatisfied.

Consider next the effect of liquidity requirements in the compound shock environment. As in the simple shock environment, banks again face an equilibrium selection problem, albeit one that is more complicated. In making initial investment decisions, banks must consider not only possible cash needs following a first-stage shock, but cash needs following a possible second-stage shock as well. Moreover, liquidity restrictions further increase the necessary sophistication of banks' investment strategies. In the CB treatment, coming to appreciate that banks must boost cash buffers to address a possible second-stage shock requires a relatively straightforward assessment: The additional \$2 needed in the symmetric SPNE is simply each bank's expected liability in the case of an aftershock. In the CL regime, however, the reasoning supporting the equilibrium adjustment is still more involved. Here, because banks must maintain \$4 in reserves, they must on average hold an additional \$3, a total that can be determined only by reasoning recursively from cash needs in the second stage. Given the added complexity of equilibrium play in the CL treatment, it is unsurprising to observe weaker conformance with equilibrium behavior.

The distributions of realized net needs for the compound shock treatments, shown as Figure 6, illustrate clearly the behavioral consequences of the additional complexity that liquidity requirements create. In the first stage, realized net needs are again more heterogeneous in the unregulated CB treatment, as was the case for the baseline treatment in the simple shock environment. In the second stage, however, CL banks generally find themselves with cash buffers insufficient to respond to a second-stage shock.[31]

Regression estimates of incidences of net cash deficiency for the compound shock treatments, shown in the middle and bottom panels of Table 4, summarize these effects overall and allow some insight into the relative importance of hoarding behavior in the compound shock environment.[32] Following the first-stage shock, the 38.9% incidence of realized cash deficiencies in the CB treatment exceeds the comparable 31.5% incidence in the CL treatment (although unlike the simple shock environment the difference is not significant.) Separating the factors driving realized first-stage cash deficiencies in the compound shock environment reveals a different response to an idiosyncratic shock among the liquidity-regulated banks relative to that observed in the simple shock environment. Unlike the unregulated simple shock environment, the incidence of initial cash deficiencies in the CB treatment (at 55.6%) was 10.1 percentage points *lower* than in the CL counterpart (at 65.7%) although this difference was not statistically significant. However, withholding-induced cash deficiencies occurred in 13.89 percent of CB periods, compared with only 5.10% of CL periods, and this difference was statistically significant..

Thus, in the first stage liquidity-regulated banks compounded a (statistically insignificant) increased propensity to fail to maintain the cash buffers needed to keep the banking system solvent in the event of a second-stage shock with a (statistically significant) increased propensity to make too much of their limited available cash.[33] Combined, these propensities left the liquidity-

[31] The bottom panel of Figure 6 illustrates realized net needs only for periods in which a second-stage shock is realized. Absent a second-stage shock, net needs are by definition zero.

[32] Parallel to the data used in generating the lower panel of Figure 6, linear probability estimates, for the second stage of the compound shock treatments, limit observations to periods in which a second-stage shock occurred, and estimate only the effect of switching from baseline to the liquidity-regulated regime.

[33] The incidences of initial cash deficiencies for the compound shock treatments listed in column 2 of Table 4 reflect the frequency of periods in which banks collectively held buffers insufficient to respond to both an initial and a second-stage shock (for example, $16 in the CB treatment and $8 in the CL treatment). The incidences of withholding-induced deficiencies reported in column (3) indicate periods in which banks with sufficient second-

regulated banking system markedly less prepared for a second-stage shock. As shown in the bottom rows of Table 4, the 79.80% incidence of realized aggregate cash deficiencies in the CL treatment is nearly 27 percentage points higher than the 53.0% aggregate cash deficiency rate in the CB treatment. The difference is driven entirely by the cash deficiency at the beginning of the second stage. In the CB treatment, the banking system held insufficient cash to address a second-stage shock in 40.15 percent of periods, compared with a 75.12% rate in CL treatment periods. Second-stage hoarding behavior also occurred in both the CB and the CL treatment periods; however, the incidences of withholding-induced shortages are virtually identical in each treatment (25.00% in CB vs. 25.90% in CL).

In summary, in the compound shock environment, liquidity requirements increase the difficulty of identifying the equilibrium investment strategy. In the first stage, liquidity-regulated banks frequently invest too much initially in the sense that they hold cash buffers too small to cope with the event of a second-stage shock. Subsequently, liquidity-regulated banks often hoard less following the initial shock. Combined, these tendencies reduce the incidence of stage 1 bankruptcies, but at the cost of chronically deficient buffers for the possible stage 2 shock. As a consequence, liquidity-regulated banks experience markedly higher bankruptcies following the occurrence of a second-stage shock.

6. Conclusions

This paper reports an experiment conducted to examine the capacity of a liquidity requirement to improve the stability of an interbank market, as well as the costs in terms of forgone liquidity transformation. We find that a liquidity requirement does, to some extent, reduce the

stage buffers created or exacerbated deficiencies in their first-stage solvency needs ($8 in the CB treatment and $4 in the CL treatment) by withholding cash.

incidence of bankruptcies in response to a simple nonprobabilistically occurring idiosyncratic shock. The requirement, however, is less effective when there is some uncertainty in the banking system about aggregate cash needs. Banks remain persistently unprepared to respond to a probabilistically occurring second-stage shock. Moreover we find that liquidity requirements impose a very high cost in terms of forgone investment. On net, liquidity requirements reduce by about half the potential gains from exchange that interbank trade allows compared with autarky.

Despite the costs of liquidity requirements in our experiment, we do not view our results as a cost-benefit analysis of liquidity regulations. The streamlined environment studied here is hardly a realistic description of existing interbank markets. The number of banks is small, the liquidity shocks are very large, the money supply is fixed, and there is no central bank. Moreover, we force interbank markets to operate by asset sales rather than loans and impose a stylized uniform price trading mechanism.

Instead, we view our experiment as providing two things. First, it identifies qualitative factors such as heterogeneous investment behavior, forgone investment opportunities, challenges in preparing for a large liquidity shock, and coordination problems in reallocating liquidity that could impact the evaluation of liquidity regulation. Second, it is an initial step toward designing market experiments that are useful for analyzing interbank markets. Future experimental work could build on our experiments to incorporate features such as more banks, different shock patterns, interbank network features as in Babus and Kondor (2018) and Craig and Ma (2018), or bilateral lending as in Afonso and Lagos (2015).

References

Acemoglu, D., Ozdaglar, A., Tabaz-Saleh, A. 2015. Systemic Risk and Stability in Financial Networks. American Economic Review 105(2), 564-608.

Acharya V., Merrouche, O. 2013. Precautionary hoarding of liquidity and interbank markets: evidence from subprime crisis. Review of Finance 17(1), 107-160.

Afonso, G., Kovner, A., Schoar, A. 2011. Stressed, not frozen: the federal funds market in the financial crises. Journal of Finance 66(4), 1109-1139.

Afonso, G., Lagos, R. 2015. Trade dynamics in the market for federal funds. Econometrica 83(1), 263-313.

Allen, F., Carletti, E., Gale, D. 2009. Interbank market liquidity and central bank intervention. Journal of Monetary Economics 56(5), 639–652.

Allen, F., Gale, D. 2004a. Financial intermediaries and markets. Econometrica 72(4), 1023-1061.

Allen, F., Gale, D. 2004b. Financial fragility, liquidity and asset prices. Journal of the European Economic Association 2(6), 1015-1048.

Allen, F., Gale, D. 2000. Financial Contagion. Journal of Political Economy 108 (1), 1-33.

Babus, A., Kondor, P. 2018. Trading and information diffusion in over-the-counter markets. Paul Woolley Centre Working Paper No. 57.

Bhattacharya, S., Gale, D. 1987. Preference shocks, liquidity and central bank policy. In Barnett, W., Singleton, K. (Eds.), New Approaches to Monetary Economics. Cambridge University Press, Cambridge.

Carlson, M., 2013. "Lessons from the historical use of reserve requirements in the United States to promote bank liquidity." FEDS Working Paper 2013-11.

Choi, S., Gallo, E., Wallace, B. 2017. Financial contagion in networks: A market experiment. Manuscript. Seoul National University.

Craig, B. and Ma, Y. 2018. Intermediation in the Interbank Lending Market. Manuscript, Stanford Graduate School of Business.

Davis, D., Holt, C. 1993. Experimental economics. Princeton University Press: Princeton.

Davis, D., Korenok, O. Lightle, J., 2019. An experimental examination of the interbank loan market. Experimental Economics, forthcoming.

Davis, D., Korenok, O. Lightle, J., 2018. Liquidity regulation, banking history and financial fragility: an experimental examination. Manuscript, Virginia Commonwealth University.

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

Diamond, D., Kashyap, A. 2016. Liquidity requirements, liquidity choices and financial stability." In Taylor, J., Woodford, M. (Eds.), Handbook of Macroeconomics, vol. 2A. North-Holland: Amsterdam, 2263-2304.

Diamond, D., Rajan, R. 2011. Fear of fire sales, illiquidity seeking and credit freezes. Quarterly Journal of Economics, 126 (2), 557–591.

Duffy, J., Karadimitropoulou, A., Parravano, M. 2019. Financial contagion in the laboratory: Does network structure matter? Journal of Money Credit and Banking (forthcoming)

Fischbacher, U. 2007. "z-Tree: Zurich toolbox for ready-made economic experiments. Experimental Economics, 10 (2), 171-178.

Gale, D., Yorulmazer, T. 2013. Liquidity hoarding. Theoretical Economics 8 (2), 1555-1561.

Goodhart, C., 2008. The regulatory response to the financial crisis. Journal of Financial Stability, 4 (4), 351-358.

Goodfriend, M., Hargraves, M. 1983. A Historical Assessment of the Rationales and Functions of Reserve Requirements. Federal Reserve Bank of Richmond Economic Review (March/April), 3-21.

Greiner, B. 2015. Subject pool recruitment procedures: organizing experiments with ORSEE. Journal of the Economic Science Association, 1(1), 114-125.
Heider, F., Hoerova, M., Holthausen, C. 2015. Liquidity hoarding and interbank market spreads: the role of counterparty risk. Journal of Financial Economics, 118 (2), 336–354.
Kowalik, M. 2013. Basel liquidity regulation: was it improved with the 2013 revisions? Federal Reserve Bank of Kansas City Economic Review (Second Quarter): 65-87.
Sprague, O. 1910. History of Crises under the National Banking System. U.S. Government Printing Office: Washington.

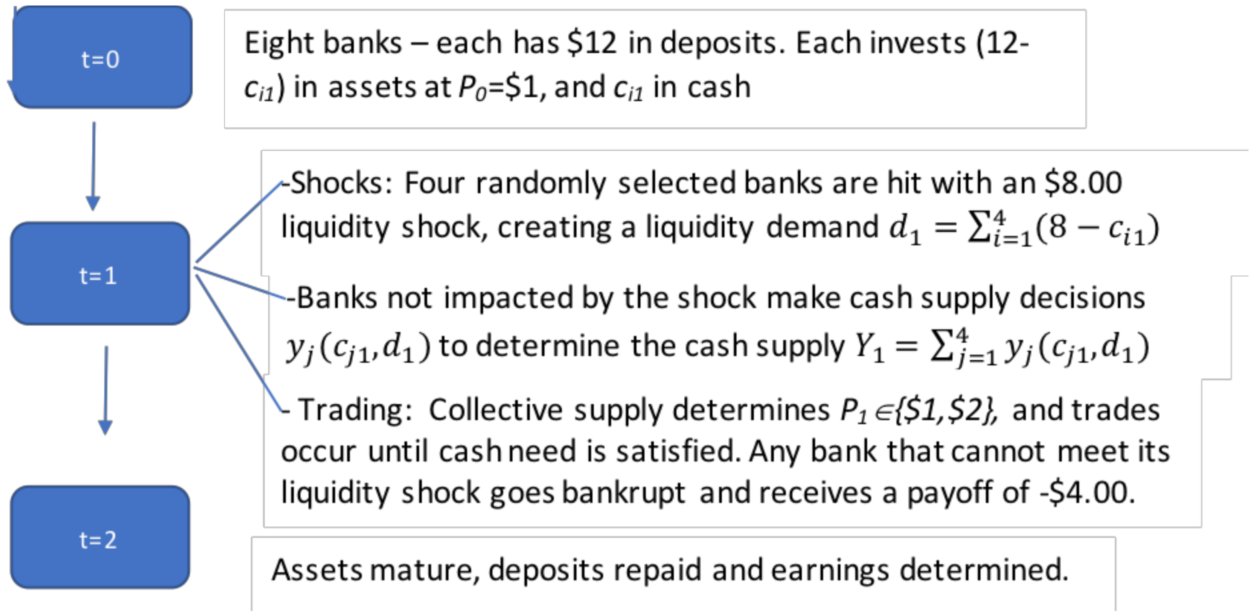

Figure 1. Illustration of the simple shock game.

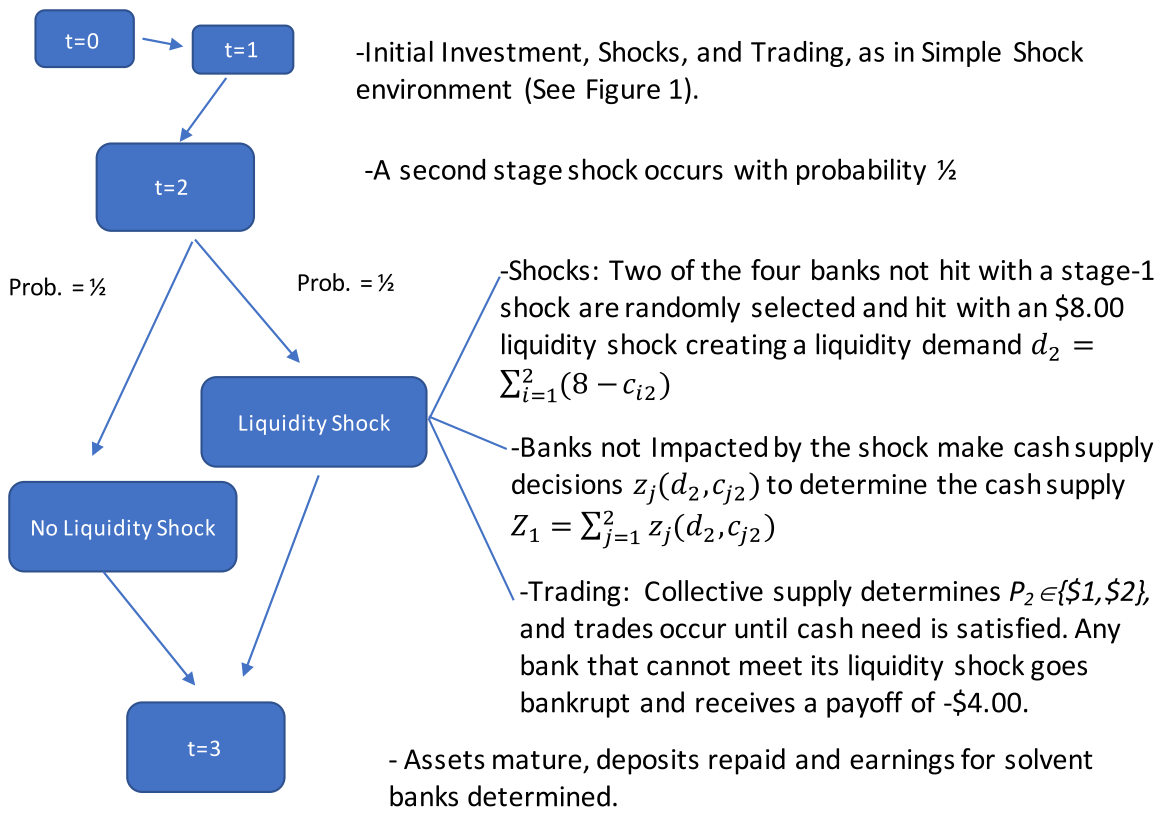

Figure 2. Illustration of the compound shock game.

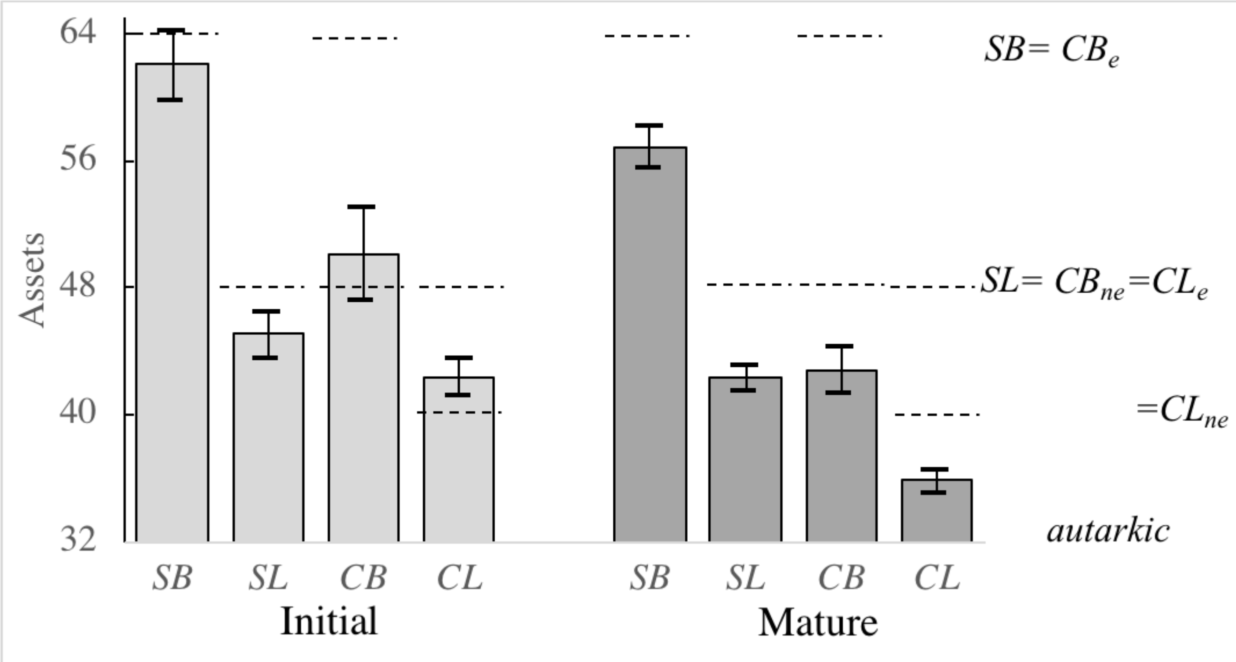

Figure 3. Mean investment levels for each treatment. 'Initial' refers to mean initial investment decisions. 'Mature' refers to the mean level of investments that were not liquidated due to bankruptcy. Dashed lines highlight benchmark SPNE investment levels, and the thin bars illustrate 95% confidence bands about mean investments. The autarkic outcome of 32 units in each treatment refers to an allocation in which each bank invests 4 assets and does not need to use interbank trade to meet withdrawal shocks

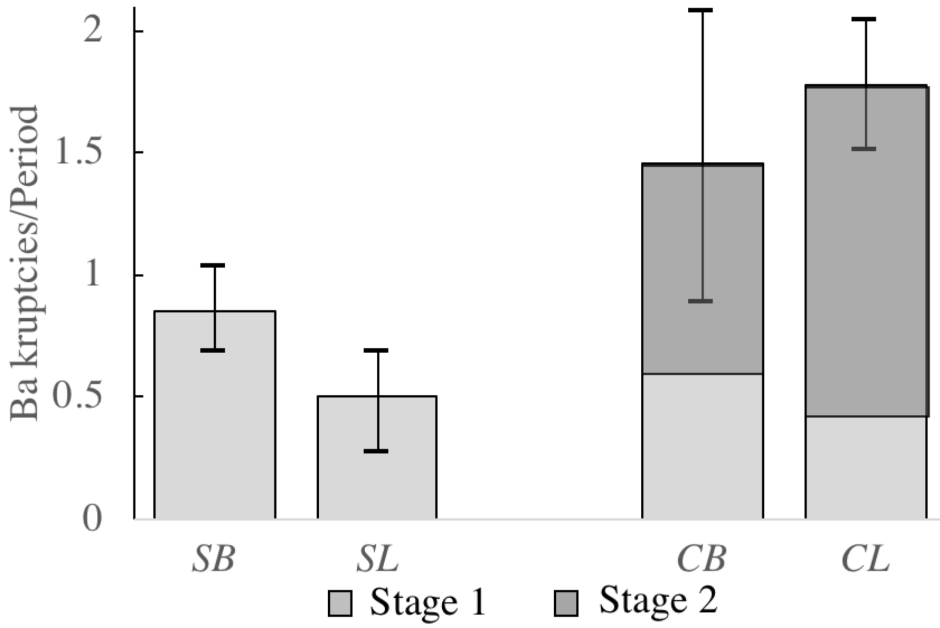

Figure 4. Bankruptcy rates for all idiosyncratic shock periods and for stage 2 periods where a second-stage shock occurred. The thin lines about the top of each bar illustrate the 95% confidence interval about the mean overall bankruptcy rate for a treatment.

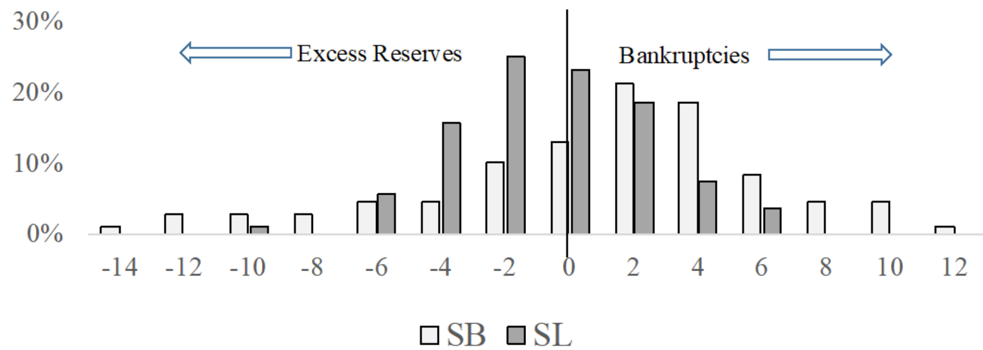

Figure 5. Distribution of realized net cash needs in the simple shock environment

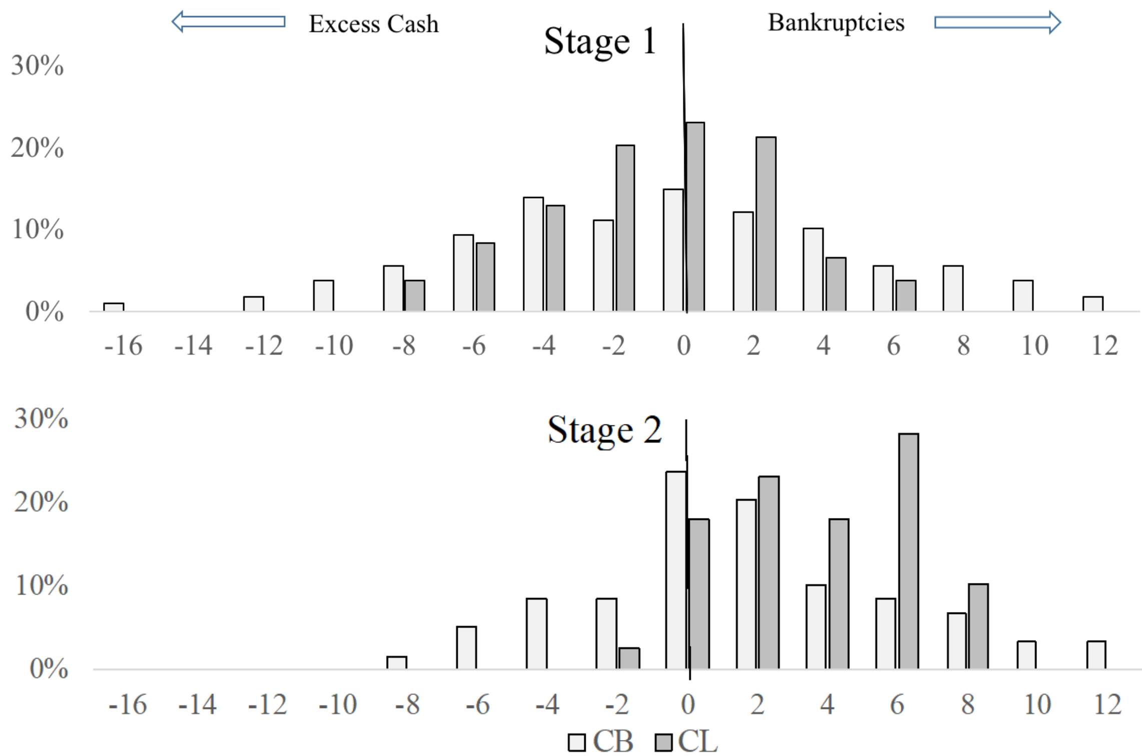

Figure 6. Distribution of net realized cash needs in the complex environment for both stages.

Treatment Equilibrium	Period 0 Cash	Expected Payoff	Aggregate Investment
SB	$4.00	$8.00	64
CB_{ne}	$6.00	$6.00	48
CB_{e}	$4.00	$6.00	64
SL	$6.00	$6.00	48
CL_{ne}	$7.00	$5.00	40
CL_{e}	$6.00	$4.50	48
Autarkic (all treatments)	$8.00	$4.00	32

Table 1. Reference equilibrium predictions. Each row in the table summarizes individual cash holdings and expected payoffs along with aggregate investment for SPNE listed in the left column of the table. For the compound shock environment, the *ne* subscript refers to the no-exposure SPNE and the *e* subscript refers to the exposure SPNE.

(1) Treatment	(2) $\overline{inv}$	(3a) i_{ne}	(3b) $\overline{inv}$- i_{ne}	(4a) i_e	(4b) $\overline{inv}$- i_e	(5) *SB-SL*	(6) *CB-CL*
SB	62.08	64	-1.92^{*}				
SL	45.06	48	-2.94^{***}			17.02^{***}	
CB	50.15	48	2.15^{**}	64	-13.85^{***}		
CL	42.34	40	2.34^{***}	48	-5.66^{***}		7.81^{***}

Table 2. Initial investments. $\overline{inv}$ denotes mean initial investment, i_{ne} denotes aggregate investment in the no-exposure SPNE for the SB, SL, CB and CL treatments, i_e denotes aggregate investment in the exposure SPNE for the CB and CL treatments.
Note: *, **, and *** denote rejections of the null that the listed difference equals zero, p<0.10, 0.05, and 0.01, respectively. In all treatments, $\overline{inv}$ exceeds the autarkic level of 32 units at p<0.01.

			Periods with a 2nd Stage Shock			
(1) Treatment	(2) 1st Stage		(3) 2nd Stage		(4) Total†	
	Rate	*B* vs. *L*	Rate	*B* vs. *L*	Rate	*B* vs. *L*
Simple Shock Environment						
SB	0.85					
SL	0.50	0.35***				
Compound Shock Environment						
CB	0.68		0.86		1.49	
CL	0.46	0.22	1.38	-0.52**	1.79	-0.30

Table 3. Average bankruptcies per period. Note: *, **, *** , indicates that the bankruptcy rate in the Baseline treatment significantly differs from the rate for the comparable Liquidity Requirement treatment at $p<.10$, $p<.05$, and $p<.01$, respectively.† Total bankruptcy rates in column (4) do not exactly decompose into entries in columns (2) and (4) because the estimates in (2) are based on all compound shock periods, while (3) and (4) are based only on periods with a second-stage shock realization.

	Realized	Initial	Withholding Induced
Simple Shock Environment			
SB	51.85%	38.89%	28.70%
SL	29.63%***	14.8%***	21.20%
Compound Shock Environment 1st Stage			
CB	38.90%	55.60%	13.89%
CL	31.50%	65.70%	5.50%*
Compound Shock Environment 2nd Stage			
CB	53.00%	40.15%	25.00%
CL	79.8%**	75.12%***	25.90%

Table 4. Incidences of aggregate cash deficiencies. Note that: *, **, and *** indicate rejection of the null hypothesis that the incidences of cash deficiencies for the Baseline and Liquidity Requirement treatments do not significantly differ, at $p<0.10$, 0.05, and 0.01, respectively.