European Economic Review 72 (2014) 39–51

Contents lists available at ScienceDirect

European Economic Review

journal homepage: www.elsevier.com/locate/eer

An experiment on the causes of bank run contagions

Surajeet Chakravarty [a,1], Miguel A. Fonseca [a,*], Todd R. Kaplan [a,b,2]

[a] University of Exeter, United Kingdom
[b] University of Haifa, Israel

ARTICLE INFO

Article history:
Received 21 October 2013
Accepted 8 September 2014
Available online 22 September 2014

JEL classification:
C72
C92
D43

Keywords:
Bank runs
Contagion
Experiments
Multiple equilibria

ABSTRACT

To understand the mechanisms behind bank run contagions, we conduct bank run experiments in a modified Diamond–Dybvig setup with two banks (Left and Right). The banks' liquidity levels are either linked or independent. Left Bank depositors see their bank's liquidity level before deciding. Right Bank depositors only see Left Bank withdrawals before deciding. We find that Left Bank depositors' actions significantly affect Right Bank depositors' behavior, even when liquidities are independent. Furthermore, a panic may be a one-way street: an increase in Left Bank withdrawals can cause a panic run on the Right Bank, but a decrease does not calm depositors.

1. Introduction

Bank runs are important economic phenomena. Over the last decade, we have witnessed visible and traditional bank runs on banks such as Northern Rock, which was the first run on a UK bank in 140 years and Countrywide Financial in the USA. There have been many more non-traditional runs on other financial institutions, such as Bear Sterns, Lehman Brothers, as well as countries—Iceland and Greece being the most high-profile cases. The present paper seeks to understand how bank runs may spread from one economic institution to another (e.g., from Lehman Brothers to AIG; from Greece to Spain). In particular, we ask whether changes in banking fundamentals cause contagions or are pure panics to blame.

Diamond and Dybvig (1983) proposed an influential analysis of bank runs. In their paradigm, a bank run is one of many possible equilibria of the economic system. The driving force for a bank run is the fact that in a fractional reserve system, a bank does not hold enough liquid assets to serve all its customers should they all decide to withdraw their deposits at one given time. Hence, if depositors believe that too many people will withdraw their deposits such that in the future the bank will not have enough money to pay them, then all depositors will withdraw today. This causes a run on the bank, even if the bank is otherwise solvent. This is self-fulfilling because a bank must liquidate its investment portfolio at fire-sale prices in order to meet the unexpected demand today, which hurts its ability to pay tomorrow.[3]

* Corresponding author. Tel.: +44 1392 262584; fax: +44 1392 263242.
E-mail addresses: s.chakravarty@exeter.ac.uk (S. Chakravarty), m.a.fonseca@exeter.ac.uk (M.A. Fonseca), dr@toddkaplan.com (T.R. Kaplan).
[1] Tel.: +44 1392 263419; fax: +44 1392 263242.
[2] Tel.: +44 1392 263237; fax: +44 1392 263242.

[3] There are alternative models in which bank runs are caused by asymmetric information among bank depositors about their banks' fundamentals. In these models, bank runs are caused by depositors' beliefs about solvency of their banks, rather than beliefs about the actions of other depositors. See, for instance, Chari and Jagannathan (1988), Jacklin and Bhattacharya (1988), Calomiris and Kahn (1991), and Chen (1999).

http://dx.doi.org/10.1016/j.euroecorev.2014.09.003
0014-2921/

The same logic may apply to contagions; however, it is important to distinguish between two types: *information-based* contagions and *panic-based* contagions. An information-based contagion occurs when a run on a bank conveys information about the wider financial system: it is a run on one bank caused by a revision of depositors' beliefs about its liquidity or solvency following a run on another bank with related liquidity or solvency. A pure *panic-based* contagion is the case where a run on one bank triggers a run on other banks even though their liquidity and solvency are unrelated. In this case, the behavior of depositors in the first bank should therefore be irrelevant for any revision of beliefs about the liquidity or solvency of the rest of the financial system. We will use this terminology throughout the paper.

An example of the former case was the perceived over-exposure of banks to assets based on sub-prime mortgages during the 2007–2009 financial crisis. A run on an over-exposed bank could conceivably trigger a run on other banks, as it provides the market with a signal about the liquidation value of assets held by the banking sector.[4] On the other hand, we may observe contagions that spread on the basis of pure panics. Friedman and Schwartz (1963) argue that the run on the Bank of the United States in 1930 was not based on fundamentals; yet the run on this bank nevertheless caused a panic on the US banking system, leading to runs on other US banks at the time.

It is difficult to distinguish information-based contagions from panic-based contagions, since historical data does not afford us insight into the beliefs of investors and depositors alike. It is very difficult to ascertain what information investors are responding to, and whether or not the information is spurious. In December 11, 1930, the New York Times reported that the run on the Bank of United States was based on a false rumor spread by a small merchant, a holder of stock in the bank, who claimed that the bank had refused to sell his stock (NYT, 1930). Was this information truthful? We will never know if depositors thought the rumor was true and were withdrawing because of the information; or if they thought the rumor was false, but nonetheless they were anticipating a mass withdrawal by other depositors.

Our paper seeks to answer two questions. Firstly, can a contagion spread by panic alone? Secondly, are there differences in the way pure panic contagions form, develop and subside relative to information-based contagions? These questions are important, as policy designed to prevent and contain an information-based contagion may differ from policy designed to tackle a panic. Making public announcements about banking fundamentals may prove counter-productive, as the recent Northern Rock case highlights.[5]

We seek to answer these questions using experimental data. By abstracting away from the complex reality of financial markets, we gain an insight into how information about banks' liquidity, as well as spurious information, can potentially trigger bank run contagions in a simulated banking system. To this effect, we conduct an experiment in a modified Diamond–Dybvig setup with two banks, Left and Right. Each bank has a mix of impatient depositors, who demand their deposits immediately, and patient depositors, who are willing to withdraw their deposits at a later date. The parameter we manipulate is the liquidation value of the both banks' long-term investment (liquidity). The Left Bank depositors see their own bank's liquidity level and make their withdrawal decisions first. The Right Bank depositors do not know the liquidity level of either bank; however, they do see how many Left Bank withdrawals are made before making their own withdrawal decision.

We consider two treatments: one where both banks' liquidity levels are always the same and another where they are independent of each other. In either treatment, it can be an equilibrium for the Right Bank depositors to imitate (or not) the decisions of the Left Bank depositors. However, we would expect information about Left Bank withdrawals to have a stronger influence on Right Bank depositors' decisions when both banks' liquidity levels are always the same, as this would be an indication of the liquidity level of the Right Bank. In contrast, information about past Right Bank liquidity, as well as past withdrawals on the Right Bank, ought to be more relevant to Right Bank depositors when liquidity levels of the two banks are independent of each other. All the above are plausible mechanisms that drive banking contagions. By studying these factors, we also better understand the processes that determine equilibrium selection in economic systems.

We find that actions taken by depositors in the Left Bank significantly affect Right Bank depositor behavior, especially when the two banks' liquidities are linked. This suggests that the Right Bank depositors use information about Left Bank depositors to update their beliefs about the liquidity of their bank. However, the fact that a similarly positive and significant (though weaker) relationship exists when liquidity levels of both banks are independent of each other means that we cannot rule out the existence of contagion equilibria triggered by 'sunspots', or in our context, pure panic.

When analyzing the dynamics of bank run contagions, we find evidence which suggests that a banking panic may be a one-way street: when both banks' liquidity levels are independent of each other, an increase in Left Bank withdrawals can cause a panic run on the Right Bank, but a decrease in Left Bank withdrawals cannot calm depositors as effectively.

Changes in the Right Bank's liquidity over time also regulate the likelihood of a run on that bank. Increases in the Right Bank's liquidity level between rounds $t-2$ and $t-1$ lead to increases in withdrawal levels by patient Right Bank depositors in round t and vice-versa, but only significantly in the case where liquidities are independent. That is, in the absence of

[4] Goldstein and Pauzner (2005) analyze this type of contagion effects through a two-bank model where investors get noisy signals about fundamentals about bank 2 after observing aggregate outcomes pertaining to bank 1.

[5] As the *Economist* reported at the time: "Only when the Bank of England said that it would stand by the stricken Northern Rock did depositors start to run for the exit. Attempts by Alistair Darling, the chancellor of the exchequer, to reassure savers served only to lengthen the queues of people outside branches demanding their money. The run did not stop until Mr. Darling gave a taxpayer-backed guarantee on September 17th that, for the time being, all the existing deposits at Northern Rock were safe." (The Economist, 20/09/2007). For a theoretical analysis of the effect on the banking system of revealing information about fundamentals, see Kaplan (2006) and Dang et al. (2012).

actual information about contemporaneous liquidity of their bank, patient Right Bank depositors look at past levels of liquidity (which in our experiment are good predictors of present liquidity) to inform their decision whether or not to withdraw early. This is particularly so in the treatment where the liquidity levels of the Left and Right Banks are independent of each other, and information about past liquidity levels is more salient.

The remainder of the paper is organized as follows. Section 2 contextualizes our paper in the existing literature on bank runs. Section 3 outlines the experimental design and the theoretical predictions. Section 4 presents the empirical results. Section 5 considers implications of our results.

2. Literature

Our paper contributes to both the literature on bank runs and the experimental literature on coordination games. In the former case, the empirical evidence on bank run contagions is scarce, because historically banking contagions are themselves infrequent. The strand of empirical literature focusing on the determinants of bank runs finds that the likelihood of a run on a bank during a crisis is positively correlated with the fundamentals of that bank (Calomiris and Mason, 1997; Schumacher, 2000; Martinez Peria and Schmukler, 2001). Calomiris and Gorton (1991) analyze widespread banking panics in the National Banking period in the U.S. (1863–1913) and find that they were correlated more with stock market declines than agricultural shocks. They also find that some banks failed during the panics primarily due to assets that were weak prior to the panic while others failed primarily due to assets that were weakened by the panic.[6] The failure of a large cooperative bank in India in 2001 generated an interesting case study on the study of bank runs and contagion. Iyer and Puri (2012) study depositor behavior on a bank that had been affected by that failure, and study the institutional determinants of a run on a bank. They find depositor insurance, as well as long-standing bank-depositor relationships, which can effectively mitigate the extent of a run. Iyer and Peydro (2011) study the impact that same failure had on the likelihood of a run on other local banks that had exposure as institutional depositors. They find that banks with high exposure to the failed bank had a higher likelihood of incurring large deposit withdrawals. Banks with weaker fundamentals were also more likely to suffer a run.

The experimental literature on bank runs is both small and very recent (see Dufwenberg, 2012 for a recent survey). This literature has focused on designs which study cases with only one bank. Madies (2006) analyses the possibility and persistence of self-fulfilling bank runs. Schotter and Yorulmazer (2009) find that when there is uncertainty about the rate of return on deposits, the presence of insiders (depositors who know the true rate of return) is welfare enhancing. Garratt and Keister (2009) find that uncertainty regarding the number of impatient depositors increases the likelihood of a bank run; increasing the number of withdrawal opportunities also results in a higher number of bank runs. Kiss et al. (2012,2014) look at the effect of observability of actions on the likelihood of bank runs to emerge. Arifovic et al. (2013) study the likelihood of the emergence of runs as a function of the fraction of the set of depositors who are required to wait in order for the bank to remain solvent. They find that for very low fractions, runs are rare, and for high fractions runs are very frequent. They identify a parameter region for which runs depend on the history of play. Arifovic and Jiang (2014) study the effect of sunspot bank runs by studying the effect of announcements forecasting the number of early withdrawals on the likelihood of runs occurring. They find that subjects react to sunspot announcements in the intermediate region identified by Arifovic et al. (2013) but not in regions where one of the equilibria is stable.[7]

Our paper is related to Brown et al. (2012), who have independently investigated in the lab the determinants of contagion. Like our paper, Brown et al. (2012) conduct experiments with two banks, in which the depositors of one bank make their decisions first, and the depositors of a second bank make their decisions after observing the actions of the first bank depositors. As in our paper, Brown et al. (2012) implement a treatment where there are economic linkages between the two banks, and a treatment where there are no linkages. The two papers differ in four important aspects of the experimental design: the number of depositors in each bank, strategic uncertainty regarding types, the number of repetitions of the game, and the economic variable which links the two banks.

In our experiment, each bank has ten depositors, as opposed to two in Brown et al. (2012). This arguably makes for a more difficult coordination problem for participants to tackle. Furthermore, in our design subjects are randomly allocated types from round to round, which closely follows the Diamond–Dybvig model and adds strategic uncertainty to the experiment, while this feature is absent from Brown et al.'s design, who consider a one-shot game, and in which they elicit beliefs about the returns to the bank's long-term asset as well as the likelihood of the other depositor withdrawing early. While our design did not collect explicit measures of beliefs, it does allow us to study the dynamic effects of changes in liquidation value, and of changes in behavior in the other bank. Finally, as mentioned above, the variable which we allow to vary in our experiment is the liquidation value of the long-term asset of the bank, while Brown et al. allow the rate of return to the long-term asset to vary. Unlike our paper, Brown et al. (2012) only find evidence of contagion when there are explicit economic linkages between the two banks. Combined with our results, this indicates that the mechanisms behind

[6] These weakened assets were caused by fraud and/or asset depreciation. There was also one bank that failed solely to a run with presumably strong fundamentals.

[7] Our paper also contributes to the literature of coordination games with Pareto-ranked equilibria (see Camerer, 2003 and Devetag and Ortmann, 2007 for surveys of the evidence). Perhaps paradoxically, by employing a more complex setup, we are able to shed some light on how beliefs about a particular equilibrium being played are shaped, and how they depend on contextual information, as well as strategically relevant information.

panic-based contagions may be due to fears about short-term liquidity in the banking system rather than the returns on long-term assets, thus vindicating Friedman and Schwartz's view.

3. Theory and experimental design

In this section, we present a simplified version of the Diamond–Dybvig model, which forms the basis of our experimental design and hypothesis. We conclude the section by outlining the experimental procedures.

3.1. A version of the Diamond–Dybvig model

The Diamond–Dybvig (1983) model (DD) is the basis of our experimental design. In our version of this three-period model, depositors place their money in a bank in period 0 (yesterday) before learning whether they are impatient or patient.[8] When impatient, depositors need to withdraw their money in period 1 (today), as they get relatively very little utility for the money tomorrow (impatient depositors have utility $u(x_1+\alpha \cdot x_2)$ where x_1 is money today, x_2 is money tomorrow, and $0 \le \alpha < 1$). When patient, depositors can wait until period 2 (tomorrow) to withdraw; however, can always withdraw the money today and hold on to it until tomorrow (patient depositors have utility $u(x_1+x_2)$). There is an equal proportion of patient and impatient depositors.

The bank has short-term and long-term investment opportunities for the money. The short-term investment (reserves) returns the exact amount invested. The long-term investment returns an amount $R>1$ tomorrow (but strictly less than $1/\alpha$). However, it is illiquid and returns only $L<1$ today.

The depositors that invested X yesterday have a contract with the bank. They can withdraw their money today and receive X or wait until tomorrow and receive $R \cdot X$ (that is, they can choose between $(x_1, x_2)=(X, 0)$ and $(x_1, x_2)=(0, R \cdot X)$).[9] The bank needs to offer a contract contingent upon withdrawal time, since it does not know which depositors are patient and which are impatient, just the overall fraction. To fulfill this contract, the bank places half its deposits in the short-term investment and half its deposits in the long-term investment.

If all the depositors withdraw the money according to their respective types, then the bank will be able to meet both the demand for cash today and tomorrow. In this case, each depositor has the incentive to indeed withdraw according to his true type. An impatient depositor prefers X today to $R \cdot X$ tomorrow. A patient depositor prefers $R \cdot X$ tomorrow to X today. Hence, all impatient depositors withdrawing today and all patient depositors withdrawing tomorrow is a Nash equilibrium.

While the contract is fulfilled in this Nash equilibrium, in other cases the bank cannot always remain solvent, leading to another Nash equilibrium. In this alternative equilibrium, too many depositors try to withdraw today and the bank is not able to meet the contract tomorrow. For instance, if a fraction $q>1/2$ of depositors withdraw today, then the bank will have to sell part of its long-term asset at the liquidation price. If $\frac{1}{2}L \le q-\frac{1}{2}$, then even if the bank liquidates all of its assets, there will not be enough cash to pay current demand. Waiting until tomorrow will return nothing so even the patient depositors would prefer to withdraw today and receive something rather than wait until tomorrow and receive nothing. This is a bank run equilibrium where everyone withdraws today.[10]

3.2. Experimental design

Our design expanded the DD model by adding another bank, such that we had a Left Bank and a Right Bank. Each bank had 10 depositors, five of whom were patient and the other five were impatient. Every participant took the role of a depositor and stayed with his assigned bank throughout the experiment. In each of the 30 rounds in the experiment, participants had to make a single decision: to withdraw today or to withdraw tomorrow. In every round, the computer randomly assigned participants to one of two types: patient (who are able to wait to withdraw tomorrow) and impatient (who strictly prefer to withdraw today). While impatient depositors had a less important role to play in the experiment, their existence created additional strategic uncertainty regarding patient depositors' decisions.

A bank with strictly more than five depositors withdrawing today faced an excess demand for liquidity and had to sell its long-term investments and receive a rate of return of $L<1$, while waiting until tomorrow yielded a rate of return $R>1$ on assets.

We also modified the original model by allowing each bank to have two possible levels of L. A bank could have high liquidity, $L=0.8$, or it could have low liquidity, $L=0.2$. Each bank's type was determined by a Markov process, where the transition probability was 1/3. This means that there was a two-thirds probability that a bank would maintain its liquidity level in consecutive rounds. The rate of return was constant throughout at $R=1.25$.

[8] Types are equivalent to an idiosyncratic shock to individuals' liquidity needs.

[9] The original DD model also considers an insurance aspect to a bank, in the sense that a sufficiently risk averse depositor is insured against being impatient and receives more than X today and less $R \cdot X$ tomorrow. This is not the focus of our experiment. Hence, we use a parameterization with a contract devoid of this insurance aspect. This is only the optimal contract when depositors have log utility.

[10] There is also a symmetric mixed-strategy Nash equilibrium in which patient depositors withdraw early with some positive probability. This equilibrium is, however, dynamically unstable.

Table 1
Payoffs.

	Total # of other depositors withdrawing today									
	0	1	2	3	4	5	6	7	8	9
Payoffs to impatient depositors										
Low L										
Withdraw today	100	100	100	100	100	100	86	75	67	60
Withdraw tomorrow	50	50	50	50	50	50	0	0	0	0
High L										
Withdraw today	100	100	100	100	100	100	100	100	100	90
Withdraw tomorrow	50	50	50	50	50	50	47	42	31	0
Payoffs to patient depositors										
Low L										
Withdraw today	100	100	100	100	100	100	86	75	67	60
Withdraw tomorrow	125	125	125	125	125	125	0	0	0	0
High L										
Withdraw today	100	100	100	100	100	100	100	100	100	90
Withdraw tomorrow	125	125	125	125	125	125	117	104	78	0

We implemented two distinct treatments. In the first treatment, Independent, the two banks' liquidity levels followed independent Markov processes. In the second treatment, Linked, the two banks' liquidity levels were always the same. Table 1 displays payoffs in a manner similar to that presented to the participants.[11]

In both treatments, Left Bank depositors knew their bank's liquidity level before making their withdrawal decision. Right Bank depositors could only observe the total number of withdrawals on the Left Bank in that round before deciding. They did not know what their bank's liquidity level was in that round. They did however, know what their bank liquidity level was in the previous round, except in the first round of the experiment.

From a strategic point of view, the addition of the Right Bank does not affect the set of equilibria of the Left Bank. This is because the actions of depositors in the second bank carry no payoff consequences to the first bank. In both banks and in either treatment, the bank run and the no-run equilibria are possible. Solution concepts such as sequential equilibrium do not reduce the set of equilibria relative to Nash equilibrium. For example, in either Linked or Independent, Right Bank depositors imitating the actions of Left Bank depositors is a Nash (and sequential) equilibrium. Also, Right Bank depositors ignoring the actions of Left Bank depositors is also a Nash (and sequential) equilibrium.

3.3. Hypotheses

We start by looking at the Left Bank depositors, who are playing a game similar to the DD model. As discussed in the previous subsection, there are multiple equilibria. In some equilibria, patient depositors withdraw tomorrow, and a run on the bank does not occur; in other equilibria, patient depositors withdraw early, and a run on the bank takes place. The Nash equilibrium concept does not rule out any relationship between the liquidity level, L, and the likelihood of a run. Using an evolutionary dynamic process to study the DD model, Temzelidis (1997) states that as banks become more illiquid, the likelihood of a run increases. As such, we should observe more runs when L is low, as opposed to when L is high, which forms our first hypothesis.

Hypothesis 1. The frequency of early withdrawals by patient Left Bank depositors will be higher when the Left Bank's liquidity levels are low.

We turn to the main hypotheses of the paper, which concern the way in which a contagion may spread. Standard theory is unable to guide our understanding of why one equilibrium is played over another; however, behaviorally, one may be more predominant.

Observe that a patient Right Bank depositor may believe that other patient Right Bank depositors will withdraw early if they believe that the Right Bank has low liquidity. Therefore, a run on the Right Bank may be triggered by depositors believing that their bank has a low L. This belief could be formed by observing Left Bank depositors running on their bank.

Hypothesis 2. The fraction of early withdrawals by patient Right Bank depositors will be correlated with the total number of early withdrawals on the Left Bank. Likewise, changes in the number of early withdrawals in the Left Bank will be positively correlated with withdrawals by patient Right Bank depositors. These correlations will be higher in Linked than in Independent.

[11] Our payoffs in case of excess early demand equal the expected payoffs rather than being based on a sequential service constraint. This was done to facilitate participants' understanding of the task. Note that the terms 'patient' and 'impatient' were replaced with 'type-A' and 'type-B'. Likewise, 'L' was called 'Reserves', and 'depositors' were called 'customers'. See Web Appendix for a copy of the instructions.

The first part of Hypothesis 2 is tested by examining the correlation between behavior of patient Right Bank depositors with total number of withdrawals in Left Bank in both treatments. Note that it is possible to observe a correlation between Left Bank withdrawals and Right Bank withdrawals in level terms without observing any effect in terms of changes. All one would require for this to be the case is for different sessions (which proxy markets) to have different initial conditions and remain at their respective states. Although there would be a correlation between behavior across banks, that correlation is driven by variation across markets, rather than an adjustment process within a market. We believe that any correlation would not be driven by such a variation and thus we hypothesize that changes the number of withdrawals in one bank (i.e., the start or the end of a run in that bank) affect the change in the likelihood of a run in another bank.

Whether Left Bank depositor behavior conveys information to Right Bank depositors can be determined by comparing the aforementioned correlations in Linked to that in Independent. If indeed the correlations are solely driven by a revision of beliefs about the liquidity of their bank, then we should observe a positive correlation in Linked but not in Independent. This would constitute a *information-based* contagion. If we find a positive correlation in the latter case, this would be evidence supporting pure *panic-based* contagions. The last part of Hypothesis 2 conjectures that there would exist an *information-based* contagion.

The previous two hypotheses concerned how a bank run can start and spread contemporaneously. We can also look at how a run on a bank propagates over time. In our experiment, the level of liquidity of a given bank follows a Markov process, thus, for Right Bank depositors, the level of L in the previous round is informative about the level of L in the current round and may affect the likelihood of a run. Alternatively, a patient Right Bank depositor may believe that other patient Right Bank depositors will withdraw early if there was a run on the Right Bank in the previous round. This leads to our next Hypothesis.

Hypothesis 3. The likelihood of an early withdrawal by patient Right Bank depositors will be correlated with both the Right Bank's liquidity in the previous round and the total number of early withdrawals on the Right Bank in the previous round.

The first part of Hypothesis 3 is based on the liquidity level of the bank, while the second part concerns a panic mechanism of propagation: a run now triggers a run in the future, even though the liquidity of the bank may since have changed. To understand if any, or both of the two is at work, we need to estimate the likelihood of an early withdrawal as a function of the level of past liquidity of the bank, as well as the number of past withdrawals on the same bank in the previous round. If only the former is a significant predictor of behavior, then only the liquidity level drives the persistence of a bank run; if the latter is also a significant predictor of current depositor behavior, then we have evidence for the existence of panic propagation mechanisms.

3.4. Experimental procedures

We provided written instruction sets (see Web Appendix), which informed participants of all the features of the market. We generated six independent sessions for each treatment (Independent and Linked). Each session had 20 participants, who interacted with each other for the duration of the experiment. There were 30 rounds in the experiment. At the beginning of the experiment, each participant was assigned to a bank (Left or Right), and remained a depositor of that bank for the whole experiment. In each round, each participant was randomly assigned a depositor type, A or B (corresponding to patient or impatient depositor), for his bank.

Participants sat at a booth which did not allow visual or verbal communication and interacted via a computer terminal. At the end of each round, participants were reminded about their own decision, and were told what the level of reserves their bank had that round (which is L in our model), as well as how many withdrawals were made either today or tomorrow at their bank.

The participants' payment was the sum of their payoffs from three rounds, which were randomly picked by the computer – this was done to avoid income effects. At the end of the experiment, participants filled in a socio-demographic questionnaire before being paid and leaving the lab. Each session lasted on average for 90 min. A total of 240 undergraduate students from a variety of backgrounds participated in our experiments. No one participated in more than one session and no one had participated in similar experiments before. The sessions took place in March and October 2011. The average payment was £13.15 ($20.66).[12]

4. Experimental results

We begin by analyzing the effect of bank liquidity on the fraction of depositors who withdraw early. Impatient depositors, as predicted, almost always withdrew early, regardless of the level of liquidity of their bank.[13] There are, however, significantly more early withdrawals by Left Bank patient depositors when $L=0.2$ than when $L=0.8$: a fraction of 0.71 in the former case vs. a fraction of 0.15 in the latter ($p<0.001$, Mann–Whitney test (MW) for independent samples).[14]

[12] The software was programmed in Z-Tree (Fischbacher, 2007) and we used the recruitment software ORSEE (Greiner, 2004).

[13] The frequency of early withdrawals for impatient Left Bank depositors was 99% when liquidity levels were low; 98% when liquidity levels were high, and 96% overall.

[14] In this part of the analysis we pool data on both treatments's Left Banks since their behavior should not be affected. Indeed that is the case for both levels of liquidity ($L=0.2$: $p=0.69$; $L=0.8$, $p=0.38$, MW). Whenever performing tests using non-parametric statistics, we use session-level averages, as is common practice; p denotes p-values on hypotheses.

Table 2
Frequency of withdrawals by Left Bank depositors as a function of L.

Linked+Independent	Total Left Bank withdrawals							
	0–3	4	5	6	7	8	9	10
$L=0.2$	–	0	1	7	29	49	60	45
$L=0.8$	–	4	85	53	21	5	1	0
$\Pr(L=0.2)$	–	0	0.01	0.12	0.58	0.91	0.98	1.00

Table 3
Fraction of withdrawals by Patient Right Bank depositors as a function of total Left Bank withdrawals and Right Bank liquidity in $t-1$.

	Total Left Bank withdrawals							
	0–3	4	5	6	7	8	9	10
Linked								
$L^{RB}_{t-1}=0.2$	–	0.30	0.26	0.46	0.51	0.64	0.63	0.76
N	–	2	13	10	14	19	19	14
$L^{RB}_{t-1}=0.8$	–	0.20	0.22	0.23	0.24	0.53	0.68	0.33
N	–	1	36	16	5	26	16	3
Total	–	0.27	0.23	0.31	0.45	0.61	0.65	0.68
N	–	3	50	27	21	27	35	17
Independent								
$L^{RB}_{t-1}=0.2$	–	–	0.57	0.60	0.68	0.65	0.65	0.76
N	–	–	15	15	15	17	12	21
$L^{RB}_{t-1}=0.8$	–	0.20	0.17	0.32	0.46	0.38	0.29	0.57
N	–	1	19	17	13	9	13	7
Total	–	0.20	0.34	0.45	0.57	0.56	0.45	0.71
N	–	1	36	33	29	27	26	28

Table 2 breaks down the frequency of early withdrawals over the 30 periods as a function of L, in combined data from both treatments.[15] Table 2 also calculates the probability of being in the state where $L=0.2$, conditional on observing a given level of withdrawals. While this statistic would not be available to subjects during the experiment, it is a good proxy of how informative the Left Bank depositors' behavior is regarding the level of L in the Left Bank over the course of the experiment.

Table 2 clarifies three important aspects concerning the behavior of Left Bank depositors. Firstly, when $L=0.2$, the likelihood of there being a run on the Left Bank (i.e., seven or more depositors withdrawing early) is quite high: 183/191, or 96%. In particular, conditional on observing eight or more depositors withdrawing early, the probability of L being equal to 0.2 is 96.25%. Secondly, when $L=0.8$, the likelihood of there *not* being a run on the bank (i.e., six or less depositors withdrawing early) is quite high: 142/171 or 84%. The probability of L being equal to 0.8 conditional on observing six or less early withdrawals is equal to 95%. In other words, extreme events, such as either very few early withdrawals or very many early withdrawals by patient Left Bank depositors, constitute a very accurate signal of the current level of L on the Left Bank. Finally, it is important to note that the probability of L being equal to 0.2 monotonically increases in the number of total withdrawals in the Left Bank. Since these depositors knew their bank's liquidity levels before deciding, their behavior is an accurate signal of the level L in the Left Bank.[16] This is our first result.

Results 1. *Patient Left Bank depositors run more often when their bank's liquidity levels are low.*

Given that Right Bank depositors know the total number of early withdrawals on the Left Bank before they make their decision, it is pertinent to examine what fraction of Patient Right Bank depositors withdrew early, conditional on the total number of early withdrawals by Left Bank depositors. Table 3 shows that this relationship is positive in both treatments. In the Linked treatment, Spearman's ρ is 0.59 ($p<0.01$), while in the Independent treatment, Spearman's ρ is 0.32 ($p<0.01$).

[15] The picture is similar looking at Linked and Independent separate.

[16] It is also possible that past withdrawal levels may affect the likelihood of a run occurring on the Left Bank. To test for this, we conducted the following random effects probit regression: $w^{LB}_{i,t}=I\{\beta_1 L^{LB}_t+\beta_2 W^{LB}_{t-1}+\beta_3 W^{LB}_{t-2}+\beta_4 W^{LB}_{t-3}+\beta_5 Round+u_i+e_{i,t}\}$, where $w^{LB}_{i,t}$ equals one if patient Left Bank depositor i withdrew early in period t, L^{LB}_t denotes the Left Bank's level of L in round t, and W^{LB}_{t-x} denotes the total early withdrawals on the Left Bank in period $t-x$. As expected, the coefficient on L^{LB}_t is as expected large and significant, -1.800 ($p<0.001$). While the coefficient on W^{LB}_{t-1} is small and significant ($0.087, p<0.001$), the other lags on W^{LB} are both small and either marginally significant or not significant at all: $W^{LB}_{t-2}: 0.016, p=0.553$; $W^{LB}_{t-3}: 0.040, p=0.085$. In other words, current levels of liquidity matter more than past withdrawals to patient depositors when they are informed of their bank's liquidity. Note this takes into account that L^{LB}_t can vary by 0.6 and W^{LB} can vary by 5.

Table 4
Estimates from random-effects probit regression on the determinants of patient Right Bank depositors' withdrawals.

	(1)	(2)
W_t^{LB}	0.139***	
	(0.041)	
$W_t^{LB} \leq 6$		−0.633***
		(0.173)
$W_t^{LB} \geq 8$		−0.109
		(0.128)
L_{t-1}^{RB}	−0.794***	−0.804***
	(0.118)	(0.125)
W_{t-1}^{RB}	0.119***	0.112***
	(0.043)	(0.039)
W_t^{LB} × LINKED	0.152***	
	(0.055)	
$W_t^{LB} \leq 6$ × LINKED		0.008
		(0.254)
$W_t^{LB} \geq 8$ × LINKED		0.560***
		(0.210)
L_{t-1}^{RB} × LINKED	0.574***	0.646***
	(0.168)	(0.176)
W_{t-1}^{RB} × LINKED	−0.045	−0.015
	(0.068)	(0.057)
Round	0.027***	0.032***
	(0.008)	(0.007)
Round × LINKED	−0.020*	−0.022**
	(0.012)	(0.011)
LINKED	−1.411**	−0.843
	(0.605)	(0.575)
Constant	−1.078**	0.213
	(0.479)	(0.354)
ρ	0.256	0.254
	(0.043)	(0.039)
N	1740	1740

Bootstrapped standard errors at market level in parentheses.
*** Significance at 1% level.
** Significance at 5% level.
* Significance at 10% level.

The effect of an extra early withdrawal in the Left Bank may be more salient when $L=0.2$ in the previous period than when $L=0.8$, particularly in the LINKED treatment. To examine this, we conditioned the information on the Right Bank's L in $t-1$. Past Right Bank liquidity makes depositors to react differently to withdrawal behavior by Left Bank depositors: when L_{t-1}^{RB} is low, an increase from five to six early withdrawals in the Left Bank leads to a jump in withdrawals in the Right Bank, while when L_{t-1} is high, we observe a similar jump only when early withdrawals on the Left Bank go from seven to eight. In the INDEPENDENT treatment, we do not observe clear jumps in likelihood of early withdrawal as a function of Right Bank behavior, although a positive correlation remains ($L_{t-1}^{RB}=0.2$, Spearman's $\rho=0.226, p=0.028$; $L_{t-1}^{RB}=0.8$, Spearman's $\rho=0.327, p=0.003$).

We have now established that the past level of liquidity of the Right Bank, as well as information about the behavior of the Left Bank's depositors, is correlated with the Right Bank depositors' decisions. It is therefore important to understand each relationship, while statistically controlling for the effect of the other. Table 4 reports results of random effects probit regressions using the withdrawal decision by patient Right Bank depositor i in period t as the dependent variable. The first specification we use is

$$w_{it}^{RB} = I\{\beta_1(1+X)+\beta_2(1+X)\times \text{LINKED}+u_i+\varepsilon_{it}>0\} \quad (1)$$

where w_{it}^{RB} equals one if patient Right Bank depositor i withdrew early in period t. $I\{\cdot\}$ is an indicator function which is equal to one if the left-hand side of the inequality is positive and takes a value of zero otherwise; X is a vector of regressors, which includes the following variables: W_{t-1}^{RB}, the total number of withdrawals on the Right Bank in the previous round; L_{t-1}^{RB}, the liquidity level of the Right Bank in the previous round; W_t^{LB}, the total number of withdrawals on the Left Bank in the current period; *Round* is a time trend, as well as their respective interactions which the dummy variable LINKED which equals one if the observation comes from the LINKED treatment and zero otherwise.[17]

[17] It is worth noting that the effect of an extra early withdrawal in the Left Bank may be more salient when $L=0.2$ in the previous period, than when $L=0.8$, as per the analysis in Table 3. We considered econometric specifications with interaction terms between W_t^{LB} and L_{t-1}^{RB} but those terms did not yield significant coefficients. As such we opted not to include them in the analysis.

The estimates from this model are summarized in column (1) in Table 4. As Table 2 established, particular ranges of withdrawal levels provide very accurate signals about the current level of L of the Left Bank. In particular, the likelihood of low L is close to one when total withdrawals are more than or equal to eight, and close to zero when total withdrawals are six or fewer. As such, it is possible that patient Right Bank depositors' withdrawals may change as a function of whether Left Bank withdrawals are below or above these thresholds. To test for this, we estimated a different specification where we replaced W_t^{LB} with two dummy variables: $W_t^{LB} \leq 6$, which takes the value of one if six or fewer Left Bank depositors withdrew early in period t; and $W_t^{LB} \geq 8$, which equals one if eight or more Left Bank depositors withdrew early in period t. The omitted category corresponds to the case where seven Left Bank depositors withdrew early.[18] The estimates from this model are summarized in column (2) in Table 4.

We start by looking at the effect of withdrawals by Left Bank depositors on patient Right Bank depositor behavior. Our first model assumes a linear relationship between the two variables. We find a positive and significant coefficient on W_t^{LB} in both treatments. The larger coefficient is, as expected, in the Linked treatment, and the difference is statistically significant ($z=2.27, p=0.023$). Our second model instead assumes a non-linear relationship between Left Bank withdrawals and patient Right Bank depositor behavior. The coefficient on $W_t^{LB} \leq 6$ is negative and significant, which indicates that when patient Right Bank depositors are less likely to withdraw early when they observe six or fewer withdrawals in the Left Bank (which is a good predictor of the Left Bank's L being equal to 0.2), relative to when they observe seven early withdrawals in the Left Bank. The interaction dummy $W_t^{LB} \leq 6 \times$ Linked is not statistically significant ($z=0.03, p=0.975$), which indicates that there is no statistical difference in the likelihood of early withdrawal when subjects observe six or fewer withdrawals between Linked and Independent. Therefore, when patient Right Bank depositors observe an accurate signal that the Left Bank's L is low, they are less likely to withdraw early than when they observe an inaccurate signal about the Left Bank's level of L. In contrast, the coefficient on $W_t^{LB} \geq 8$ is not significant ($z=-0.85, p=0.393$). In other words, the likelihood of early withdrawal by patient Right Bank depositors in the Independent treatment does not significantly change when they observe eight or more withdrawals. The opposite is true in the Linked treatment: the coefficient on $W_t^{LB} \geq 8 \times$ Linked is positive and significant ($z=2.67, p=0.008$). In other words, patient Right Bank depositors are more likely to withdraw early conditional on observing eight or more early withdrawals (which is a good predictor of the Left Bank's L being equal to 0.8) in the Left Bank in Linked than in Independent. This is our second result.

Results 2. *Patient Right Bank depositors are more likely to withdraw early, the higher the total number of early withdrawals by Left Bank depositors. This result is stronger in the* Linked *treatment.*

This result supports Hypothesis 2, in that Left Bank depositor behavior influences Right Bank depositors, particularly when it conveys information which can be used to update beliefs about liquidity levels. However, the fact that this relationship is significant in the Independent treatment means that we cannot rule out pure panics as potential causes of bank run contagions. This relationship in Independent can be compared to the sunspot equilibrium found in Arifovic and Jiang (2014). However in our paper, the sunspots are created endogenously; also, rather than the message being about a run in their *own* bank, as in Arifovic and Jiang's experiment, the signal in our setup is about a run occurring on the *other* bank, which may trigger a run in their bank.

We now turn to the effect of past liquidity levels in the Right Bank on current depositor behavior. In Independent, we see a negative and significant effect of the Right Bank's liquidity level in the previous round on the level of withdrawals in the current round. This effect is significantly smaller in Linked, and not significantly different from zero ($L_{t-1}^{RB}+L_{t-1}^{RB}\times$ Linked $=0, z=2.40, p=0.122$): patient Right Bank depositors are influenced by past liquidity conditions in their own bank in Independent but not in Linked.

Results 3. *In* Independent, *Patient Right Bank depositors are more likely to withdraw early if the liquidity level of their bank in the previous round was low.*

We finalize this analysis by looking at the persistence of bank runs. Will patient Right bank depositors be more likely to withdraw early if total early withdrawals on the Right Bank in the previous round were high? We find a positive and significant coefficient on W_{t-1}^{RB} in both specifications, indicating some persistence in bank runs in Independent. The coefficient on the interaction between W_{t-1}^{RB} and Linked is negative but not statistically significant, indicating no difference between the two treatments. This is our next result.[19]

Results 4. *Patient Right Bank depositors are more likely to withdraw the higher the total number of withdrawals on their bank was in the previous round.*

[18] The choice of threshold level of withdrawals was made as a function of the accuracy of the signal. However, the results are robust to selecting different levels of withdrawals for the two dummies.

[19] It is possible that subjects may have a concern about the history of withdrawal levels, rather than just the withdrawal levels in the preceding round. To check for whether this is the case, we augmented regression (1) in Table 4 by adding W_{t-2}^{RB}, W_{t-3}^{RB} and their interactions with Linked. The coefficients on the other variables on this augmented regression retained their signs and significances, and the coefficient on W_{t-2}^{RB} was equal to 0.066 ($p=0.073$), while $W_{t-2}^{RB}\times$ Linked had a coefficient of -0.080 ($p=0.115$). The coefficient on W_{t-3}^{RB} was equal to 0.071 ($p=0.123$), while $W_{t-3}^{RB}\times$ Linked had a coefficient of -0.038 ($p=0.530$). In other words, lagged withdrawal levels by Right Bank depositors in $t-2$ had a much smaller effect than that in $t-1$, and virtually no effect in $t-3$.

We now focus on how patient Right Bank depositors react to changes in market conditions. In particular, we analyze how changes in early withdrawals by patient Right Bank depositors are affected by changes in the number of early withdrawals from the previous round to the current round in the Left Bank, as well as changes in the liquidity of the Right Bank from two rounds ago to the previous round. To do so, we report a series of random effects least squares regressions. The dependent variable is the change in the proportion of early withdrawals by patient Right Bank depositors. We used aggregated data as opposed to individual-level data because participants were randomly assigned a role (patient or impatient) in every round. As such, on average, half of the time participants who were patient depositors in one round were impatient depositors in the following round.

We consider two econometric specifications, which we describe in turn. The first specification is

$$\Delta W_{it}^{RB} = \beta_0 + \beta_1 \Delta W^{LB} + \beta_2(\Delta L = 0(high)) + \beta_3(\Delta L > 0) + \beta_4(\Delta L < 0) + \beta_5 Round + \alpha_i + \varepsilon_{it}, \tag{2}$$

and has as regressors ΔW^{LB}, the change in total withdrawals by Left Bank depositors, in addition to dummies for positive and negative changes in the Right Bank's liquidity in the previous round, $(\Delta L > 0)$, $(\Delta L < 0)$, a dummy for no change in L when L was already high ($\Delta L = 0$ (*high*)), as well as a time trend (*Round*).

The results from this specification are in column (3) in Table 5. The coefficient on (ΔW^{LB}) is positive and highly significant for both Independent and Linked. An increase in the number of early withdrawals in the Left Bank leads to an increase in early withdrawals in the Right Bank, although the effect is significantly higher in Linked than Independent ($(\Delta W^{LB}) \times$ Linked $= 0.438, p < 0.01$.) The coefficient on ($\Delta L = 0$ (*high*)) is non-significant in both treatments, suggesting no difference relative to the default category ($\Delta L = 0$ (*low*)). The coefficient on $(\Delta L > 0)$ is negative and highly significant, which means that an increase in liquidity levels is correlated with a decrease in the number of withdrawals by patient Right Bank depositors; there is no difference in effect size between treatments: $(\Delta L > 0) \times$Linked $= 0.069, p = 0.433$. On the other hand, the coefficient on $(\Delta L < 0)$ is positive in both treatments, but significantly different from zero only in Independent.

Table 5
Random effects least squares estimation of changes in early withdrawals.

	(3)	(4)
ΔW^{LB}	0.592***	
	(0.109)	
$\Delta W^{LB} > 0$		0.154***
		(0.056)
$\Delta W^{LB} < 0$		−0.051
		(0.056)
$\Delta L = 0$ (high)	−0.052	−0.058
	(0.054)	(0.057)
$\Delta L > 0$	−0.238***	−0.229***
	(0.059)	(0.063)
$\Delta L < 0$	0.246***	0.248***
	(0.059)	(0.063)
$\Delta W^{LB} \times$ Linked	0.438***	
	(0.163)	
$\Delta W^{LB} > 0 \times$ Linked		0.036
		(0.078)
$\Delta W^{LB} < 0 \times$ Linked		−0.119
		(0.080)
$\Delta L = 0$ (high) $\times$Linked	0.066	0.150*
	(0.078)	(0.082)
$\Delta L > 0 \times$ Linked	0.069	0.165*
	(0.089)	(0.092)
$\Delta L < 0 \times$ Linked	−0.177**	−0.171*
	(0.087)	(0.093)
Linked	−0.052	−0.070
	(0.082)	(0.100)
Round	−0.002	−0.001
	(0.003)	(0.003)
Round $\times$ Linked	0.003	0.003
	(0.004)	(0.004)
Constant	0.052	0.012
	(0.057)	(0.071)
Groups, Observations	12, 28	12, 28
R^2	0.35	0.26

Standard errors in parentheses.
*** Significance at 1% level.
** Significance at 5% level.
* Significance at 10% level.

Furthermore the difference in coefficients between the two conditions is significant ($(\Delta L<0)$ $\times$LINKED $=-0.177$, $p=0.041$). In other words, the effect of a drop in Right Bank liquidity on Right Bank withdrawals is only significant in INDEPENDENT. Finally, we do not observe any time trend on either treatment.

The second econometric specification considers the sign of changes in the number of withdrawals in the Left Bank, rather than the size of the effect. This specification permits us to infer whether or not increases in Left Bank withdrawals have a different qualitative effect than decreases in Left Bank withdrawals. The new specification is

$$\Delta W_{it}^{RB}=\beta_0+\beta_1(\Delta W^{LB}>0)+\beta_1(\Delta W^{LB}<0)+\beta_2(\Delta L=0(high))+\beta_3(\Delta L>0)+\beta_4(\Delta L<0)+\beta_5 Round+\alpha_i+\varepsilon_{it}, \quad (3)$$

which includes dummy variables for increases and decreases in total withdrawals in the Left Bank, $(\Delta W^{LB}>0)$ and $(\Delta W^{LB}<0)$, respectively. The omitted category is no change in withdrawals. The results from this regression are displayed in column (4) in Table 5. We find positive and significant coefficients on $(\Delta W^{LB}>0)$ in both treatments, with no statistical difference between the two ($(\Delta W^{LB}>0)\times$ LINKED $=0.036, p=0.649$.) We find negative coefficients in $(\Delta W^{LB}<0)$ in both treatments, though only significant in LINKED. An increase in L leads to a decrease in withdrawals by patient Right Bank depositors, though only significantly so in INDEPENDENT. Likewise a decrease in L leads to an increase in withdrawals by patient Right Bank depositors, though again only significantly so in INDEPENDENT. We do not observe any time trend. We summarize the findings from this analysis below.

Results 5.

(i) *A rise in total Left Bank withdrawals leads to an increase in withdrawals by patient Right Bank depositors. However, there is no significant change when there is a drop in total Left Bank withdrawals in* INDEPENDENT.

(ii) *A rise (fall) in Right Bank liquidity levels between rounds* $t-2$ and $t-1$ *leads to a fall (rise) in withdrawal levels by patient Right Bank depositors in rounds t, particularly in* INDEPENDENT.

There are two possible explanations as to why behavior in the Left Bank influences behavior in the Right Bank: depositors in the Right Bank may simply be imitating behavior of Left Bank depositors or they may instead believe that other Right Bank depositors are imitating Left Bank depositor behavior, even though they themselves would not necessarily do so. These two mechanisms are closely related to the notions of primary and secondary saliency put forward by Mehta et al. (1994). Unfortunately, these mechanisms predict the same behavior in our experiment.

Note we also examined individual behavior of Right Bank Depositors by classifying individuals according to a type that closest resembled their behavior. We observed a higher proportion participants who only responded to Left Bank depositor behavior in LINKED (41%) than in INDEPENDENT (17%). In place of this, we saw an increase in those that always withdrew (23% in INDEPENDENT vs. 11% in LINKED) and an increase in those that myopically withdrew early whenever their bank's previous liquidity was low (23% in INDEPENDENT vs. 16% in LINKED). For the complete analysis, see Web Appendix.

5. Discussion

The theoretical literature on bank runs distinguishes two main causes of bank runs and banking contagions: insolvency or illiquidity. From an empirical point of view, the former is easier to detect, as evidence will be present in the balance sheets. The latter is more difficult to analyze, as it is driven by beliefs about the bank's short-term liquidity, as well as beliefs about the behavior of other depositors. Not only are real bank runs rare, but even when they do occur, it is impossible to gauge depositors' beliefs about banking liquidity levels, as well as beliefs about other depositors' actions. We tackle this difficulty by using experimental methods and simplifying the problem faced by real depositors to its core: a coordination problem among depositors. In this environment, the role of depositor beliefs – both about liquidity levels and about what other depositors will do – is crucial in determining which action depositors take, and in turn which equilibrium is selected.

In the absence of direct information about short-term liquidity, we find evidence that depositors do base their beliefs about actions at least partially on what a different set of depositors did at the same bank the previous period; however, our key finding is that liquidity is strongly correlated not only with the likelihood of a run on a bank, but also with the likelihood of contagion spreading to a separate bank. We identify three mechanisms through which short-term liquidity affects runs. The first is the contemporaneous effect of liquidity when liquidity levels are known. There is a clear relationship between liquidity and the likelihood of a run. While the no-run equilibrium is always Pareto-superior to the run equilibrium, its riskiness is higher when the bank's liquidity is low since in this case if a single patient depositor withdraws early, all depositors who withdraw later will receive a payoff of zero. This indicates the importance of off-equilibrium payoffs in determining the likelihood of players picking a particular equilibrium.

The second and third mechanisms concern the formation of beliefs about liquidity when that information is not known. In our experiment, the fact that banks' liquidity levels follow a Markov process means that when current liquidity is unknown, participants can partially infer it from the level of liquidity in the previous round. This indicates that if depositors anchor their beliefs about current liquidity on past liquidity, a bank run could potentially persist over time even when liquidity levels no longer support the existence of such an equilibrium, as per the first mechanism. This mechanism is particularly salient in the treatment in which both banks' liquidity levels are independent of each other since past liquidity is the best predictor of current liquidity.

The third mechanism concerns the updating of beliefs about one's bank based on the behavior of depositors in another bank. When both banks have the same liquidity and depositors believe that bank runs are more likely when liquidity levels are low, a run by informed depositors in one bank provides a signal that may trigger a run by uninformed depositors in another bank. We find evidence for these information-based contagions. We also find evidence that banking contagions can be caused by panic. We observe that a run on one bank increases likelihood of depositors of another bank running even though both banks' liquidities are independent of each other. The behavior of depositors in the first bank is a meaningless signal but is not ignored.

Distinguishing between these two types of contagion matters because they display different dynamics. When bank liquidities are linked, the level of withdrawals in the Left Bank acts as a coordination device for Right Bank depositors. As such, runs on the latter bank are as easy to start as to stop. However, panic-based contagions are harder to stop once started. In the absence of a reliable signal, depositors may not be able to coordinate on the no-run equilibrium and, as such, panic-based contagions may be more persistent than information-based ones.

This suggests that there is value not only in reinforcing banking inter-linkages for their value in diversifying risk (Allen and Gale, 2000), but also in making those linkages common knowledge. This is because avoiding the spread of contagion can then be achieved by focusing on its origin, as opposed to panic-based contagions, which may require action throughout the financial system in order to be quelled. This is particularly relevant, as we are dealing with bank runs that are liquidity based, which are always economically inefficient.

An additional source of systemic risk to the banking system could be aggregate uncertainty about liquidity shocks to depositors (the proportion of impatient depositors). We leave this issue for future research.

Acknowledgments

We are very grateful to Tim Miller for his outstanding help in programming the software and running the sessions. Financial support from the University of Exeter Business School is also gratefully acknowledged. We also thank Mardi Dungey, Dan Friedman, Ro'i Zultan, seminar participants at Bar Ilan University, Ben-Gurion University, Copenhagen Business School, University of Düsseldorf, Hebrew University, University of Exeter, University of Guelph, University of Haifa, Queen Mary University of London, University of Queensland, the University of Tasmania, the 2012 ESA International Meeting at NYU, 2012 ESA American Meeting at U. Arizona, the 2014 RES Conference, as well as two anonymous referees and an Associate Editor for helpful comments and discussions. All remaining errors are ours alone.

Appendix A. Supplementary data

Supplementary data associated with this paper can be found in the online version at http://dx.doi.org/10.1016/j.euroecorev.2014.09.003.

References

Allen, F., Gale, D., 2000. Financial contagion. J. Polit. Econ. 108 (1), 1–33.
Arifovic, J., Jiang, J., 2014. Do Sunspots Matter? Evidence from an Experimental Study of Bank Runs. Bank of Canada. Working Paper 2014-12.
Arifovic, J., Jiang, J., Xu, Y., 2013. Experimental evidence of bank runs as pure coordination failures. J. Econ. Dyn. Control 37 (12), 2446–2465.
Brown, M., Trautmann, S., Vlahu, R., 2012. Contagious Bank Runs: Experimental Evidence. DNB Working Paper No. 363.
Calomiris, C.W., Gorton, G., 1991. The origins of banking panics: Models, facts and bank regulation. In: Glenn Hubbard, R. (Ed.), Financial markets and financial crises, University of Chicago Press, Chicago, pp. 109–173.
Calomiris, C., Kahn, C., 1991. The role of demandable debt in structuring optimal banking arrangements. Am. Econ. Rev. 81 (3), 497–513.
Calomiris, C., Mason, J., 1997. Contagion and bank failures during the great depression: the june 1932 chicago banking panic. Am. Econ. Rev. 87 (5), 863–883.
Camerer, C.F., 2003. Behavioral Game Theory: Experiments in Strategic Interaction. Russell Sage Foundation, New York.
Chari, V., Jagannathan, R., 1988. Banking panics, information, and rational expectations equilibrium. J. Finance 43 (3), 749–761.
Chen, Y., 1999. Banking panics: the role of the first-come, first-served rule and information externalities. J. Polit. Econ. 107 (5), 946–968.
Dang, T., Gorton, G., Holmstrom, B., 2012. Ignorance, Debt and Financial Crises. Mimeo.
Devetag, G., Ortmann, A., 2007. When and why? A critical survey on coordination failure in the laboratory. Exper. Econ. 10 (3), 331–344.
Diamond, D., Dybvig, P., 1983. Bank runs, deposit insurance, and liquidity. J. Polit. Econ. 20, 401–419.
Dufwenberg, M., 2012. Banking on Experiments? Mimeo.
Fischbacher, U., 2007. z-tree: Zurich toolbox for ready-made economic experiments. Exper. Econ. 10 (2), 171–178.
Friedman, M., Schwartz, A., 1963. A Monetary History of the United States, 1867–1960. Princeton University Press, Princeton.
Garratt, R., Keister, T., 2009. Bank runs as coordination failures: an experimental study. J. Econ. Behav. Organ. 71 (2), 300–317.
Goldstein, I., Pauzner, A., 2005. Demand-deposit contracts and the probability of bank runs. J. Finance 60 (3), 1293–1327.
Greiner, B., 2004. An Online Recruitment System for Economic Experiments. Discussion paper, MPRA, University Library of Munich, Germany, http://econpapers.repec.org/RePEc:pra:mprapa:13513.
Iyer, R., Peydro, J.-L., 2011. Interbank contagion at work: evidence from a natural experiment. Rev. Financ. Stud. 24 (4), 1337–1377.
Iyer, R., Puri, M., 2012. Understanding bank runs: the importance of depositor-bank relationships and networks. Am. Econ. Rev. 102 (4), 1414–1445.
Jacklin, C., Bhattacharya, S., 1988. Distinguishing panics and information-based bank runs: welfare and policy implications. J. Polit. Econ. 96 (3), 568–592.
Kaplan, T., 2006. Why banks should keep secrets. Econ. Theory 27, 341–357.
Kiss, H., Rodriguez-Lara, L., Garcia, A., 2014. Do social networks prevent or promote bank runs? J. Econ. Behav. Organ. 101, 87–99.

Kiss, H., Rodriguez-Lara, I., Garcia, A., 2012. On the effects of observability and deposit insurance on bank runs: an experimental study. J. Money Credit Bank 44 (8), 1651–1665.
Madies, P., 2006. An experimental exploration of self-fulfilling banking panics: their occurrence, persistence, and prevention. J. Bus. 79 (4), 1831.
Martinez Peria, M., Schmukler, S., 2001. Do depositors punish banks for bad behavior? Market discipline, deposit insurance, and banking crises. J. Finance 56, 1029–1051.
Mehta, J., Starmer, C., Sugden, R., 1994. The nature of salience: an experimental investigation of pure coordination games. Am. Econ. Rev. 84 (3), 658–673.
New York Times, December 11, 1930. False Rumor Leads to Trouble at Bank: Branches of Bank of United States in the Bronx Meet All Withdrawal Demands.
Schotter, A., Yorulmazer, T., 2009. On the dynamics and severity of bank runs: an experimental study. J. Financ. Intermediation 18 (2), 217–241.
Schumacher, L., 2000. Bank runs and currency run in a system without a safety net: Argentina and the “Tequilla” shock. J. Monet. Econ. 46, 257–277.
Temzelidis, T., 1997. Evolution, coordination, and banking panics. J. Monet. Econ. 40, 163–183.
The Economist, 2007. Britain's Bank Run: The Bank That Failed. URL: ⟨http://www.economist.com/node/9832838⟩ (last accessed: 26/09/2012).