

Bank Runs: An Experimental Study*

Rod Garratt
Department of Economics
University of California,
Santa Barbara
garratt@econ.ucsb.edu

Todd Keister
Centro de Investigación
Económica,
ITAM
keister@itam.mx

August 3, 2005

Abstract

We use experimental methods to investigate the extent to which breakdowns in coordination can lead to bank runs. Subjects decide whether to leave money deposited in a bank or withdraw it early; a run occurs when there are too many early withdrawals. We explore the effects of randomly forcing some subjects to withdraw early and varying the number of opportunities subjects have to withdraw. Bank runs occur frequently with forced withdrawals, even if these withdrawals are unlikely to cause the bank to fail. Exposure to bank runs has a much larger effect on future withdrawal behavior when there are multiple withdrawal opportunities than with a single opportunity. We also evaluate individual withdrawal decisions according to simple cutoff rules. We find that the cutoff rule corresponding to the payoff-dominant equilibrium of the game, which involves Bayesian updating of probabilities, explains subject behavior better than other rules. *Journal of Economic Literature* Classification Numbers: Codes: C72, E0

Keywords: Bank run, coordination game, Bayesian updating, experimental

*The authors would like to thank the people at the CASSEL experimental laboratory for the use of their facility and for their assistance in running the experiments. We acknowledge the programming work of Raj Advani. We thank seminar participants at the 2004 Experimental and Behavioral Economics Workshop in Calgary, Alberta. We are also grateful to Kelly Bedard, Ted Bergstrom, Gary Charness, and Jon Sonstelie for helpful comments and suggestions. This work was funded by UC MEXUS-Conacyt grant No. CN-02-77.

Bank runs were a common occurrence in the United States before the mid-1930s and remain an important phenomenon around the world today.¹ A sizable theoretical literature has attempted to shed light on the underlying cause of these runs. Two competing explanations have emerged. In one, a run results from a coordination problem. The seminal paper of Diamond and Dybvig (1983) showed how the game played by a bank's depositors can naturally have multiple equilibria; a bank run can then be interpreted as a switch from a good equilibrium to a bad one.² The second explanation says that a bank run is caused by negative information about the value of the portfolio of individual banks or of the entire banking sector.³ It is unclear which of these explanations is better supported by the data. In U.S. banking history, there were seven instances of system-wide runs on banks between 1864 and 1933, and four of these led to widespread suspension of the convertibility of deposits into currency. In some cases the run coincided with a negative macroeconomic shock, but in many cases there was no discernible change in the value of the assets held by banks experiencing a run. In other words, the historical evidence does not clearly favor either of the theoretical explanations.⁴

We use experimental methods to test the extent to which breakdowns in coordination can lead to bank runs. The subjects in our experiment play the role of depositors in a bank. Each must choose between withdrawing her money early and waiting to withdraw. We begin with a pure coordination game in the spirit of Diamond and Dybvig (1983). If everyone waits to withdraw, they will all receive their initial deposit plus a profit. However, if too many subjects withdraw early, the bank will run out of funds and all remaining depositors will receive nothing. The experiment is designed so that the bank can absorb a certain number of early withdrawals before it becomes unable to meet its obligations to the remaining depositors. We then explore variations of the model that involve randomly forcing some subjects to withdraw early and changing the number of opportunities subjects have to withdraw early.

Adding the possibility that some subjects will be forced to withdraw early complicates the game in interesting ways. First, a subject must contemplate how the possibility that others will be forced to withdraw early affects the expected return to keeping her money deposited. Second, and perhaps more importantly, she must contemplate how other subjects interpret the uncertainty regarding forced withdrawals. Even if a subject understands how the random forced withdrawals affect the game, she may not be confident that others will understand. Hence, a bank run may occur even if everyone understands that the forced withdrawals are very unlikely to bankrupt the bank. This is, in fact, what we observe. In

¹Two recent examples are the partial run on the banking sector in Russia in July 2004 and the run which lead to the temporary closure of the banking sector in Argentina in 2001.

²There is a sizable literature based on the Diamond-Dybvig model. See, for example, Cooper and Ross (1998), Peck and Shell (2003), and Ennis and Keister (2003).

³See, for example, the theoretical work of Allen and Gale (1998) and Goldstein and Pauzner (2005), and the empirical work of Saunders and Wilson (1996) and Gorton (1998).

⁴Ennis (2003) provides a discussion of the evidence both in favor of and against a multiple-equilibrium explanation of bank runs. He also demonstrates that many of the arguments against such an explanation only apply to the most simple models. Richer models with multiple equilibria seem to be wholly consistent with the available evidence.

the experimental sessions without forced withdrawals, there were no bank runs; subjects generally coordinated on the payoff-dominant equilibrium. However, bank runs occurred frequently in the presence of forced withdrawals, even though the probability that forced withdrawals would bankrupt the bank was very low. To highlight this effect, we ran some trials where the maximum possible number of forced withdrawals was below the number required to bankrupt the bank. Even in this setting, bank runs occurred in some instances.

Notice that forced withdrawals need not necessarily increase the likelihood of a bank run. This is because (i) subjects might recognize that forced withdrawals are unlikely to themselves cause the bank to fail and (ii) forced withdrawals can mask voluntary “panic” withdrawals. In other words, subjects may interpret withdrawals observed early on as having been forced and not conclude that a panic is underway. Nevertheless, our experimental data suggests that random forced withdrawals, even when it is unlikely that they will compromise the solvency of the bank, make bank runs more frequent.

We also observe interesting effects when the number of opportunities subjects have to withdraw early is varied. The primary treatment gives subjects three withdrawal opportunities in each trial. That is, each time the game is played, subjects have three chances to withdraw their money before the final payoffs to remaining depositors are determined. After each opportunity, subjects are informed about the number of withdrawals made by other players in previous opportunities, the amount each withdrawing subject received, and the projected payoff to remaining depositors if there are no further withdrawals. We then ran a control treatment that was identical to the primary treatment except that subjects were only given one opportunity to withdraw each trial. The cumulative probability distribution of the total number of forced withdrawals was set to match that of the primary treatment. The null hypothesis is that there is no difference between the withdrawal behavior of subjects in the two treatments. We find that this hypothesis can be rejected. The frequency of withdrawals in the initial trial was similar for the two treatments. However, exposure to bank runs had a much larger effect on future withdrawal behavior with multiple withdrawal opportunities. Overall withdrawal rates were therefore higher in the primary treatment, despite starting off at similar levels.

Using data from the treatment with multiple withdrawal opportunities, we also evaluate withdrawal decisions according to various simple cutoff rules. Interestingly, we find that the cutoff rule corresponding to the payoff-dominant equilibrium of the game, which involves Bayesian updating of the probability of a forced withdrawal, outperforms more naive decision rules in explaining observed subject behavior. One possible explanation is that subjects correctly incorporate information about their own experiences regarding forced withdrawals into their perception of how likely it is that others were forced to withdraw. This has practical implications for how one might expect people to update their perception of the general well-being of the population in response to positive (or negative) personal shocks.

We ran four experimental sessions for this study, each involving 20 subjects who participated in 16 paid trial rounds. The subjects were randomly and anonymously divided into groups of five for each trial, and these groups were reshuffled after every trial. Our initial hope was that subjects would treat each trial independently, since they were playing with new groups. However, as mentioned above, there is evidence that subjects who were exposed

to bank runs early in the session tended to withdraw more often in later trials.⁵ To sort out such learning effects, we report results based on the first trial of each session type separately from the analysis of the full treatment.

While there has been much experimental work on coordination games,⁶ we know of only one other study that has conducted an experimental investigation of bank runs. Schotter and Yorulmazer (2003) study the factors that affect the severity of a bank run. In their setup, the bank is assumed to be insolvent and hence a run is certain to occur. Their interest is in how quickly resources are taken out of the banking system once a crisis is underway, and in how various factors (deposit insurance, asymmetric information, etc.) affect this speed. Our primary focus, in contrast, is on whether or not a run occurs at all. In our setup, there is always an equilibrium where all subjects leave their money in the bank unless they are forced to withdraw. Our interest is in how often play corresponds to this equilibrium compared to the equilibrium where all subjects withdraw their money at the first opportunity. Despite these differences in focus, our results have an important theme in common with those of Schotter and Yorulmazer. Both papers demonstrate that subjects play significantly differently when there are multiple opportunities to withdraw funds than when withdrawing is a one-shot decision. These results indicate that the standard approach of modelling bank runs as a simultaneous-move game (as in Diamond and Dybvig (1983) and many others) may not be the most appropriate one.

Although we present our analysis in terms of the classic notion of a run on the banking system, we believe our results also generate insight into other types of financial crises that have occurred around the world in recent years. Investors in Mexican tesobonos in 1994, for example, were very much like the depositors in a bank, each deciding whether to withdraw her investment (by not rolling it over on the due date of the bond) based on her beliefs about the quality of the investment *and* about what other investors would do. We believe, therefore, that the insights gained from the experimental analysis of our simple model can also be helpful for understanding events such as the Mexican crisis of 1994-5 and the crises in East Asia and Russia in 1997-8.⁷

The rest of the paper is organized as follows. The next section describes the experimental design, including the basic game played by subjects and the different treatments applied. Section II explains the theoretical predictions and presents the results of the primary treatment, where subjects have three withdrawal opportunities per trial. These results include a classification of individual decisions according to various cutoff rules. Section III contains the experimental results for the control treatment with one withdrawal opportunity, while Section IV contains an econometric analysis of treatment effects. Section V describes how often bank runs occurred in each of the treatments and shows how the escalation in the frequency of bank runs differed across the treatments. Finally, Section VI contains some concluding remarks.

⁵Of course, this might occur even if subjects ignored the possibility of repeatedly playing against the same subjects as they use their experiences to update their perceptions of behavior in the general population.

⁶See, for example, Cooper *et al.* (1990), Van Huyck *et al.* (1990), and the survey by Ochs (1995).

⁷See Boyd *et al.* (2001) for a detailed analysis of the available data on modern banking and financial crises.

I. Experimental Design

We devised a computer-controlled experiment in which subjects play multiple trials of a coordination game with varying strategy sets and payoffs. In each trial, subjects are randomly divided into groups of five. Each subject begins a trial with one dollar deposited in her group's bank. The bank has promised to pay her \$1.50 if she keeps her deposit in the bank until the end of the trial. However, she can instead request to withdraw her dollar before the end of the trial. Such early withdrawals affect the bank's ability to make the promised payments to the remaining depositors.

The primary treatment had three withdrawal opportunities per trial. At the end of each opportunity, feedback was given to all subjects regarding the number of withdrawals, the payoff given to withdrawers, and projected payoffs to remaining depositors if there were no further withdrawals. The control treatment had a single withdrawal opportunity per trial. Instructions for both treatments are provided in Appendix A.

A. Payoffs

Subjects deal exclusively with payoff tables that describe their payoffs to withdrawing or not withdrawing as a function of the total number of withdrawals (see the instructions provided in Appendix A). The theoretical model that underlies these payoff tables is as follows. The experimental bank in each group begins with assets whose face value totals five dollars (one dollar per depositor). These assets are invested in an illiquid technology that will yield a high return if held to maturity but yields a low return (less than face value) if liquidated to meet early withdrawals. The liquidation value of the assets determines how many depositors can withdraw early before the bank runs out of funds. Once the bank is out of funds, any remaining depositors receive nothing. In each of the treatments, subjects played the game under two different payoff scenarios.

SCENARIO A

Payoff to withdrawers: In this scenario, the bank receives \$0.60 for each dollar of assets liquidated. This implies that the bank can meet up to three early withdrawal requests with full payment of one dollar. Note, however, that three early withdrawals would completely deplete the bank's assets. If the cumulative number of withdrawal requests exceeds three in any withdrawal opportunity, the bank will liquidate all remaining assets and pay each requester in that opportunity an equal share of the proceeds.

Withdrawal requests in each opportunity are treated equally and paid according to the rules stated above. Remaining funds are used to meet withdrawal requests in subsequent opportunities and to pay remaining depositors at the end of the trial. Payoffs thus followed a quasi-sequential service rule: within each opportunity requesters are treated identically, but across opportunities depositors who request to withdraw first are served first.

Payoff to remaining depositors: The return on investment was chosen so that the bank can fulfill its promise to pay remaining depositors \$1.50 as long as there are two or fewer early withdrawals. If there are three or more early withdrawals, the bank's funds are completely depleted and remaining depositors receive nothing. In other words, the return on

the bank's assets was chosen so that depositors who did not withdraw early received either the full promised payment or zero.⁸

SCENARIO B

Payoff to withdrawers: Under the second scenario, the bank receives \$0.80 for each dollar of assets liquidated early. This rate implies that the bank can meet up to four early withdrawal requests with full payment. If all five depositors make withdrawal requests in the first opportunity, the bank liquidates all of its assets and pays \$0.80 to each depositor.

Payoff to remaining depositors: Once again, the bank can afford to pay remaining depositors \$1.50 as long as its assets are not completely depleted by early withdrawals. If there are three or fewer early withdrawals, the remaining depositors will receive \$1.50. If there are four early withdrawals, the single remaining depositor will receive nothing. Note that the higher liquidation value of assets in the scenario B payoffs makes the bank more "robust" to early withdrawals. In this case, an individual depositor will lose money by waiting only if *all* of the other depositors withdraw early.

B. Forced Withdrawals

Some of the trials in each treatment included the possibility of random forced withdrawals. In such trials, a random number of subjects is selected at the beginning of each withdrawal opportunity, and these subjects have a withdrawal request automatically submitted on their behalf. The computer first randomly selects the number of forced withdrawals, and then randomly assigns them to subjects.⁹ At the end of each opportunity, subjects are told the number of withdrawals in their group, but not whether withdrawals by others were forced or voluntary.

In the primary treatment, where there are three withdrawal opportunities, there is a 1/2 probability that one subject will be forced to withdraw and a 1/2 probability that no one will be forced to withdraw in each opportunity. In the control treatment (with a single withdrawal opportunity) there is a 1/8 probability of zero forced withdrawals, a 3/8 probability of one forced withdrawal, a 3/8 probability of two forced withdrawals, and a 1/8 probability of three forced withdrawals. These probabilities were chosen so that the cumulative distribution of forced withdrawals over the course of a trial is the same in both treatments. Of course, in the primary treatment the forced withdrawals occur over time while in the control they occur all at once.

⁸For programming purposes these payoffs (and those described in Scenario B below) were achieved by using net rate of return on the bank's investment of 200%. This high return is required to give saliency to the payoffs. Specifically, we needed to give the bank the ability to absorb some early withdrawals and still deliver a meaningful payoff (\$1.50) to remaining depositors. Subjects do not confront this return directly. Rather they deal only with the resulting payoff tables. Unlike in the Diamond/Dybvig setting, our bank is not a mutual. Any excess bank profits beyond that required to pay subjects \$1.50 each is not distributed.

⁹These random forced withdrawals correspond to the liquidity-preference shocks commonly used in the theoretical literature on bank runs. In particular, our sessions with forced withdrawals correspond to the case of *aggregate* uncertainty about liquidity demand, as studied in Section IV of Diamond and Dybvig (1983) and in Wallace (1990), Peck and Shell (2003), and others.

These probability distributions were also chosen so that the game played by subjects always has a payoff-dominant equilibrium in which no one voluntarily withdraws. In other words, even in the presence of forced withdrawals, the expected payoff to not withdrawing is greater than \$1 *if* all other players do not withdraw voluntarily. Under the scenario A payoffs, the expected payoff to each subject in the primary treatment in this payoff dominant equilibrium is \$1.28. Under the scenario B payoffs, the expected payoff to each subject in the payoff dominant equilibrium is \$1.35.¹⁰ Notice that while these two numbers are similar, there is a qualitative difference between the two scenarios. In scenario A, there is a 1/8 probability that the bank will fail and remaining depositors will get zero even if there are no voluntary withdrawals. In scenario B, on the other hand, the remaining depositors are guaranteed to be paid \$1.50 each if no one voluntarily withdraws.

C. Trial Specification and Ordering

Both treatments had the same number and ordering of trials. We began with the scenario A payoffs and no forced withdrawals. There were two unpaid practice trials (numbered 0 and 1) followed by two paid trials. We then conducted 8 trials (numbered 4 - 11) with random forced withdrawals. After these trials were completed, we paused the experiment and verbally introduced the scenario B payoffs. We then had 1 unpaid practice trial followed by two paid trials (numbers 13-14) using the scenario B payoffs and no forced withdrawals. The purpose of these trials was to familiarize subjects with the new payoffs and to, in some sense, provide a “fresh” start after the scenario A trials. Finally, we conducted 4 trials (numbers 15-18) using scenario B payoffs and random forced withdrawals. This ordering of trials is also listed in the instructions in Appendix A.

Each session therefore had 16 paid trials. Since there were 20 subjects in each session, the control treatment (with a single withdrawal opportunity) gave us 20 individual decisions per trial. The primary treatment (with three withdrawal opportunities) generated up to 60 individual decisions per trial. In terms of group outcomes, of course, we generated four data points per trial in each session, regardless of the treatment type.

We chose to run relatively few trials without forced withdrawals for the following reasons. The purpose of these trials was primarily to set the stage for the trials with forced withdrawals. We hoped to achieve two things: (1) familiarize subjects with the simpler game before complicating it with forced withdrawals, and (2) establish that subjects would play the payoff-dominant equilibrium of this game in the absence of forced withdrawals. We believed that two practice trials followed by two paid trials would be enough to achieve these goals and, as we describe below, this appears to indeed be the case.

D. Method

Subjects: Eighty undergraduate students from the University of California, Los Angeles participated in the experiment. The players participated in four separate sessions, each consisting of 20 subjects. Payoffs were stated in terms of a U.S. dollars. In addition to their earnings in the experiment, players received a \$5.00 show-up fee. Total earnings were \$478.75 and \$445.00 for sessions 1 and 2 of the primary treatment, respectively, and \$517.00

¹⁰These expected values take into account the fact that the individual may herself be forced to withdraw.

and \$495.00 for sessions 1 and 2 of the control treatment. The gap between the highest and lowest payoff was between \$3.25 and \$5.00 in each session.

Procedure: All sessions were conducted in the California Social Science Experimental Laboratory (CASSEL) at UCLA. Players were individually seated in the CASSEL, which consists of 60 networked computer workstations in separate cubicles. Each cubicle contains a computer monitor, keyboard, mouse, and a set of written instructions. The supervisor read the instructions and answered questions to ensure that everyone understood the operation of the computers, game design, and payoff function. Very few questions were asked.

At the beginning of each trial, subjects were randomly and anonymously divided into four groups of five players each. The players were each shown the payoff chart described in the instructions and saw a clock that counted down from thirty seconds. During that time they had the option of clicking on the “Withdraw Now” button in order to place a withdrawal request. If they did not click the “Withdraw Now” button within thirty seconds, their money remained deposited until the next opportunity. After each opportunity, subjects were shown a screen that told them the total number of withdrawals in their group. They were also told how much each withdrawing subject received and the projected payment to remaining depositors. Once all five players clicked “Continue,” the experiment advanced to the next opportunity. At the end of the final opportunity, subjects received information on the outcome of that trial and a report of their cumulative individual earnings for the entire session. Once everyone in the session clicked “Continue,” the experiment advanced to the next trial with newly formed groups.

Once a player withdrew, she had no more decisions to make in that trial. Such players were still updated on the outcome of their own group at the end of each opportunity. If members of a group made sufficient withdrawals to bankrupt the bank, they were informed that the bank was out of money and told to wait until the beginning of the next trial.

At the end of each trial, subjects were only informed of the outcome for the group in which they participated. Information about other groups in the session was not provided. However, it is possible that subjects in a group that went bankrupt might infer whether or not other groups had bankrupted their banks from how long they had to wait between trials. Of course, everyone had to wait the full time if only one group did not bankrupt its bank and hence the information that could be inferred from wait time was not perfect. Any form of communication during the experiment was strictly forbidden.

II. Primary Treatment

In the primary treatment, each of the three withdrawal opportunities represents a different subgame whose payoffs are determined by the number of withdrawals made in the preceding opportunities. The strategies of players in each opportunity can depend on the history of withdrawals to that point. Hence, to model the game we must specify a player’s strategy as a 3-tuple $(s_I, s_{II}(\cdot), s_{III}(\cdot))$, where each component describes the player’s strategy at a particular opportunity. Let W denote the action “withdraw” and N “not withdraw.” Let $n_j \in \{0, 1, 2, 3\}$ denote the (cumulative) number of withdrawal requests in opportunity

j , where $j = I$ or II . Recall that in scenario A the game ends if there are three withdrawal requests, so that in this case using the smaller set $n_j \in \{0, 1, 2\}$ will suffice. A strategy for a player then consists of $s_I \in \{W, N\}$, $s_{II}(n_I) \in \{W, N\}$, and $s_{III}(n_{II}) \in \{W, N\}$ for each possible value of n_j .

A. No Forced Withdrawals

Under scenario A payoffs and when there is no possibility that some subjects will be forced to withdraw, there is a payoff-dominant subgame perfect equilibrium in which each player selects $s_I = N$, $s_{II}(n_I) = N$ if $n_I \leq 2$, and $s_{III}(n_{II}) = N$ if $n_{II} \leq 2$. To see this, consider the decision faced by a player in withdrawal opportunity III. This opportunity is only reached if the number of previous withdrawals is less than 3. In any such subgame (i.e., for any $n_{II} \leq 2$), if a player believes that all other players will follow the strategy above and not withdraw, her payoff will be \$1.00 if she withdraws and \$1.50 if she waits. Hence her optimal strategy will be to also follow the strategy above and not withdraw. Working backward, the same reasoning applies to each of the first two withdrawal opportunities. Hence the above strategy profile does indeed represent a subgame perfect equilibrium. In scenario B, the bank is able to absorb one more early withdrawal before running out of funds. Now, the payoff-dominant subgame-perfect equilibrium of the game without forced withdrawals has each player playing $s_I = N$, $s_{II}(n_I) = N$ if $n_I \leq 3$, and $s_{III}(n_{II}) = N$ if $n_{II} \leq 3$.

Under both payoff scenarios there is also a “banking panic” equilibrium in which all players withdraw in opportunity I. This follows from the fact that withdrawal requests from all but one player will cause the bank to run out of funds in opportunity I and hence the game ends. Notice that this panic equilibrium will exist regardless of the presence or absence of forced withdrawals, or the number of withdrawal opportunities. Also notice that there are no equilibria where all players choose “not withdraw” in the first opportunity and then “withdraw” in a later opportunity. If a player believes that everyone is going to withdraw in, say, opportunity II (and receive a payoff of \$0.60), she has an incentive to deviate by withdrawing in opportunity I and receiving a payoff of \$1. Similar logic rules out other strategy profiles of this sort as equilibria.

The following table summarizes the results of the trials without forced withdrawals under the two different payoff scenarios. As expected, withdrawal rates were low under both payoff scenarios in the absence of forced withdrawals. Moreover, the frequencies show that subjects were more likely to play strategies consistent with the payoff dominant equilibrium of the game in scenario B. This makes sense because under the scenario B payoffs the bank is more robust to early withdrawals; four withdrawals are needed to bankrupt the bank instead of only three.

	<u>Payoff Dominant Equilibrium</u>		<u>Panic Equilibrium</u>	
	Outcomes	Strategies	Outcomes	Strategies
Scenario A	8 of 16 (50%)	72 of 80 (90.0%)	0 of 16 (0%)	4 of 80 (5%)
Scenario B	12 of 16 (75%)	76 of 80 (95.0%)	0 of 16 (0%)	1 of 80 (1.25%)

Table I: Equilibrium outcomes and individual strategies consistent with equilibrium outcomes under payoff scenarios A and B; no forced withdrawals.

B. Forced Withdrawals

When forced withdrawals are added to the primary treatment, there is a 50% chance in each withdrawal opportunity that one player will be selected at random and forced to withdraw. Under scenario A payoffs, there exists a payoff-dominant equilibrium which is qualitatively similar to the one described above. The precise strategy for each player in this equilibrium is now $s_I = N$, $s_{II}(n_I) = N$ if $n_I \leq 1$, and $s_{III}(n_{II}) = N$ if $n_{II} \leq 2$. To see why this is an equilibrium, consider the decision facing a player in withdrawal opportunity III. Recall that this opportunity is only reached if the number of previous withdrawal is less than 3. A player who has not been forced to withdraw must calculate the probability that one of the other players has been forced to withdraw as follows. Suppose there are k remaining depositors after opportunity II. Then the probability that player i should assign to a forced withdrawal having occurred in opportunity III, conditional on her not being forced to withdraw, is

$$\text{Prob}[\text{forced withdrawal} = \text{yes} \mid \text{player } i \text{ forced} = \text{no}] = \frac{k-1}{2k-1}. \quad (1)$$

These probabilities are calculated for different numbers of remaining depositors in table II. The important fact is that the probability a player should assign to a forced withdrawal having occurred, conditional on her not having been forced to withdraw, is always strictly lower than the unconditional probability of one-half.

k	<u>Unconditional</u>	<u>Conditional</u>
5	1/2	4/9
4	1/2	3/7
3	1/2	2/5
2	1/2	1/3

Table II: Conditional and unconditional probabilities of a forced withdrawal

Given the appropriate conditional probability, the expected payoff from waiting, under the belief that no one will voluntarily withdraw, is greater than the expected payoff to withdrawing in each of the possibilities. This is confirmed in Appendix B, which shows payoffs and optimal actions under each withdrawal scenario.¹¹ Hence in the opportunity-III subgame, there is always an equilibrium where players do not voluntarily withdraw.

In opportunity II, when $n_I = 2$ holds, the likelihood of a forced withdrawal by one of the other players (again calculated using (1)) combined with the prospect of a future forced withdrawal in the final opportunity makes withdrawing the optimal action. If, on the other hand, $n_I = 1$ (or 0), then despite considerations of forced withdrawals, it is optimal to not withdraw if one believes there will be no voluntary withdrawals. Finally, if a player believes all other players will follow the strategy given above, her optimal action in the first

¹¹The subjects were not shown anything like the table in Appendix B. During each withdrawal opportunity of the experiment the subjects saw the relevant, updated payoff tables (please see again, the instructions in Appendix A).

withdrawal opportunity is to not withdraw. Therefore, the strategy profile listed above is a subgame perfect equilibrium of the game with forced withdrawals.

This equilibrium is consistent with the belief by each player that all other players will not voluntarily withdraw unless it is a dominant strategy to do so. It is a dominant strategy for every player to withdraw in opportunity II if $n_I = 2$, because the possibility of forced withdrawals in opportunities II and III reduces the expected payoff to not withdrawing below the payoff of withdrawing regardless of what the other players do. It is never a dominant strategy to withdraw in opportunity III (remember this opportunity is not reached if $n_{II} > 2$) because even at $n_I = 2$ it is optimal to not withdraw assuming others will not (voluntarily) withdraw.

Switching to scenario B payoffs makes the bank resistant to one more withdrawal and raises the cutoffs in the payoff-dominant equilibrium strategy by one. The payoff-dominant equilibrium has each player playing $s_I = N$, $s_{II}(n_I) = N$ if $n_I \leq 2$, and $s_{III}(n_{II}) = N$ if $n_{II} \leq 3$.¹² As in the case on no forced withdrawals there is, under both payoff scenarios A and B, a panic equilibrium in which everyone withdraws at the first opportunity.¹³

Because of forced withdrawals and learning, aggregate statistics of the observed occurrences of equilibrium in group outcomes over the course of the eight trials are not very meaningful. Hence we conduct our analysis of these trials by examining individual behavior, trial by trial. When designing the experiment, we conjectured that subjects in the primary treatment would follow simple cutoff rules for determining their withdrawal decisions. The simplest such rule would be to withdraw if and only if a certain number of withdrawals have occurred in the previous opportunities. A slightly more elaborate (or “variable”) cutoff rule would factor in the timing of the withdrawal decision. Since the possible number of future forced withdrawals is greater in the earlier opportunities, one might expect cutoff rules of the form: do not withdraw in opportunity I, withdraw in opportunity II if and only if Y or more withdrawals occur in opportunity I, and withdraw in opportunity III if and only if $Y + 1$ or more withdrawals occur in opportunities I and II. Under both payoff scenarios, the strategy corresponding to the payoff-dominant equilibrium is of this type. We now examine the extent to which observed subject behavior is consistent with different variable cutoff rules as well as with the panic strategy of withdrawing in opportunity I.

C. Analysis of Cutoff Rules for Scenario A With Forced Withdrawals

In the payoff-dominant equilibrium described in the previous section (under scenario A payoffs), all players follow a cutoff rule with $Y = 2$. For each subject, we ask whether her observed behavior is consistent with this rule, as well as whether it is consistent with the rule $Y = 1$ and with the panic rule of withdrawing in the first opportunity. Note that an individual’s observed choices can be consistent with more than one rule, depending on the actual decisions she faced.

¹²The expected payoffs to each withdrawal decision used to show that this strategy profile is indeed a subgame perfect equilibrium are presented in Appendix C.

¹³There are also equilibria in which players coordinate their actions based on the realization of the number of forced withdrawals in potentially complicated ways. We do not believe these equilibria to be relevant for explaining subject behavior in the experiments and therefore do not analyze them here.

Interestingly, if a subject were to fail to update her belief regarding the probability of a forced withdrawal when she was not forced to withdraw (i.e., she continued to assign probability 0.5 to this event, instead of using (1)) then she would perceive the panic rule to be a (weakly) dominant strategy and hence we might expect this subject to withdraw in the first opportunity. Why does failure to do Bayesian updating of this probability have such a big effect on the perceived optimal strategy of the subjects? The reason is that it changes the decision the subject makes in opportunity III if she observes two past withdrawals and believes there will be no voluntary withdrawals. In this situation, the subject realizes that one more forced withdrawal will bankrupt the bank. Hence, if she does not withdraw she will receive zero if there is a forced withdrawal and \$1.50 if there is no forced withdrawal. She also realizes that if she withdraws, she will receive \$1.00 if there is no forced withdrawal and \$0.50 if there is a forced withdrawal. Given a chance to make a decision, a Bayesian player observes that she, herself, has not been forced to withdraw and, using equation (1), calculates the probability that one of the other players has been forced to withdraw to be 0.4. Under the belief that others will not withdraw, the expected payoff to not withdrawing is therefore $0.6 * (\$1.50) + 0.4 * (\$0) = \$0.90$, which is greater than the expected payoff to withdrawing of $0.6 * (\$1.00) + 0.4 * (\$0.50) = \$0.80$. Hence she chooses to not withdraw.

A non-Bayesian player, however, regards the probability of a forced withdrawal having occurred to be 0.5 and hence calculates the expected payoff to not withdrawing as being lower than a Bayesian would. In fact, her (perceived) expected payoff to not withdrawing is the same as her (perceived) expected payoff to withdrawing; both are \$0.75. It is thus a (weakly) dominant strategy for her to withdraw. Moreover, we might expect her to do so because the payoff from withdrawing has a lower variance. This decision feeds back to opportunity II and changes her decision under $n_I = 1$ to withdraw, which in turn implies that in opportunity I she prefers to withdraw. Hence, in the absence of Bayesian updating, a player might be expected to follow the panic rule.

This logic indicates that subjects will be much more inclined to withdraw if they do not recognize that their own forced-withdrawal outcome provides information on the outcome of others. In fact, as we demonstrate below, subjects did not tend to withdraw in the situations described above. Behavior, especially in the later withdrawal opportunities, was more consistent with the Bayesian story. Notice that a player did not need to correctly calculate the posterior probability in order to realize that withdrawing immediately is not a dominant strategy. The payoffs were designed so that she only needed to realize that the posterior probability was strictly less than the prior probability of one-half.

We begin by analyzing the data from trial 4, which was the first trial with forced withdrawals. Table III shows frequencies of observed play for three different cutoff rules: the one associated with the payoff-dominant equilibrium ($Y = 2$), an intermediate rule ($Y = 1$), and the one associated with the panic equilibrium. The number of observed strategies in each session is less than 20 (the number of subjects) because of forced withdrawals in the first opportunity. A subject's actions are classified as being consistent with a particular rule if the subject obeys that rule in each decision she faces. Subjects who are forced to withdraw after the first opportunity are classified according to their observed actions prior to that point. The table shows that the cutoff rule from the payoff-dominant equilibrium ($Y = 2$) is

superior to the others in terms of frequencies.

	<u>$Y = 2$</u>	<u>$Y = 1$</u>	<u>Panic</u>	<u>Obs.</u>
Session 1	14 (77.8%)	9 (50.0%)	0 (0.0%)	18
Session 2	14 (77.8%)	9 (50.0%)	1 (5.6%)	18
Combined	28 (77.8%)	18 (50.0%)	1 (2.8%)	36

Table III: Cutoff rule frequencies, scenario A, trial 4

This result is also evident in the data from the later rounds. Figure 1 shows the cutoff rule frequencies for the combined sessions for trials 4-11.¹⁴ The figure reveals two things. First, it shows that the superiority of the $Y = 2$ cutoff rule is not limited to the first trial. Second, it shows that “learning” matters in the experiment, in the sense that the fraction of “panicky” subjects (i.e., subjects who withdrew immediately) increased substantially over time. In fact, this increase almost fully explains the deterioration of the other two rules; most of the drop in subjects following $Y = 1$ and $Y = 2$ can be attributed to subjects changing their withdrawal decision in opportunity I.




Figure 1: Cutoff rule frequencies, scenario A, trials 4-11.

In order to understand why the $Y = 2$ rule is superior to the $Y = 1$ rule for explaining subject behavior, we must examine the differences in observed behavior at the two instances

¹⁴The graph shows pooled data from the two sessions for each trial. The decision to pool the data is justified on the grounds that there is no statistically significant difference (at the 95% level) in the linear relationship between each data set. We wish to point out, however, that visually it appears that the proportion of panicky subjects rose more quickly in session 2 than in session 1.

where these rules differ: in opportunity II with $n_I = 1$ (83 occurrences) and opportunity III with $n_{II} = 2$ (48 occurrences). The $Y = 1$ rule predicts that subjects will withdraw in both these cases, while the $Y = 2$ rule predicts that they will not. In fact, subjects withdrew in these cases only 12% and 27.1% of the time, respectively. Hence, subjects who make it through the first withdrawal opportunity tended not to withdraw in cases predicted by the $Y = 1$ rule.

D. Analysis of Cutoff Rules for Scenario B With Forced Withdrawals

Under the scenario B payoffs with forced withdrawals, the strategy played in the payoff-dominant equilibrium is a variable cutoff rule with $Y = 3$. Table IV lists the frequencies associated with this rule, as well as those for $Y = 2$ and the panic rule, in the first scenario B trial with forced withdrawals (trial 15). The table shows that under the scenario B payoffs, there is no difference in performance between the $Y = 3$ and $Y = 2$ rules; each explains close to 90% of observed behavior. The main reason for the success of these rules is that immediate withdrawals fell to 5.3%, compared to 18.8% in scenario A.

	<u>$Y = 3$</u>	<u>$Y = 2$</u>	<u>Panic</u>	<u>Obs.</u>
Session 1	18 (94.7%)	18 (94.7%)	0 (0.0%)	19
Session 2	16 (84.2%)	16 (84.2%)	2 (10.5%)	19
Combined	34 (89.5%)	34 (89.5%)	2 (5.3%)	38

Table IV: Cutoff rule frequencies, scenario B, trial 15

Under these payoffs there is also very little variation in behavior over time, as shown in figure 2. There is some indication that the proportion of panicky subjects may be on the rise (and that subjects are beginning to abandon the variable cutoff rules), but overall the proportions stay relatively constant.




Figure 2: Cutoff-rule frequencies, scenario B, trials 15-18.

E. Summary of Analysis on Individual Decision Making under Forced Withdrawals

In both payoff scenarios, the cutoff rules that prescribe not withdrawing in the first opportunity explain withdrawal behavior better than the rule associated with the panic equilibrium. Moreover, in scenario A, where the bank is more susceptible to runs and withdrawal rates are higher overall, the cutoff rule consistent with the payoff-dominant equilibrium outperforms the others. The data suggest an increased tendency to withdraw over time. This is apparent in the data for scenario A payoffs and there is some suggestion that this would be the case even under scenario B payoffs, although we did not run enough trials under these payoffs to determine anything conclusive. These results may indicate a drift away from the payoff-dominant equilibrium over time and toward play consistent with the panic equilibrium.

III. Control Treatment

We ran two sessions that were identical to the primary treatment in every respect except that the subjects were only given a single withdrawal opportunity. In order to preserve comparability across treatments, the distribution of the total number of forced withdrawals was set to match that of the primary treatment. In particular, in the single withdrawal opportunity of the control treatment there was a $1/8$ probability of zero forced withdrawals, a $3/8$ probability of one forced withdrawal, a $3/8$ probability of two forced withdrawals, and a $1/8$ probability of three forced withdrawals.

There are exactly two pure strategy Nash equilibria of this game under each of the payoff scenarios, both with and without forced withdrawals. The payoff-dominant equilibrium has all players not withdrawing unless forced to do so. Under the scenario A payoffs, a player receives \$1.50 if she is not forced to withdraw *and* the number of forced withdrawals is two or fewer. A player who is forced to withdraw receives \$1. If there are three forced withdrawals, the two remaining players receive zero. Under the scenario B payoffs, a player always receives \$1.50 if she is not forced to withdraw and \$1 if she is forced to withdraw. The panic equilibrium has all players withdrawing. Under the scenario A payoffs they each receive \$0.60, while under the scenario B payoffs they each receive \$0.80.

A. No Forced Withdrawals

In the trials with no forced withdrawals, we observed the payoff-dominant equilibrium 16 out of 16 times using scenario A payoffs and 15 out of 16 times (there was one withdrawal in the opening round of session 2) using scenario B payoffs. Hence, players in these games played strategies consistent with the payoff-dominant equilibrium almost all of the time, in each of the payoff scenarios. Recall that in the primary treatment with 3 withdrawal opportunities, the payoff-dominant equilibrium occurred only 50% of the time in scenario A and 75% of the time in scenario B in these same trials. Hence having multiple withdrawal opportunities appears to have a positive effect on withdrawal rates in the absence of forced withdrawals.

B. Forced Withdrawals

As in the primary treatment, the occurrence of forced withdrawals and learning dictates that we conduct our analysis of these trials by examining individual behavior on a trial by trial basis. Subjects in the control treatment make a single decision (withdraw or not withdraw), and hence their action in each trial is necessarily consistent with one of the two equilibrium strategies. It is therefore straightforward to check how frequently individual play coincided with each of these strategies in each trial. We begin by looking at behavior in trial 4, the first paid trial with forced withdrawals.

Table V shows frequencies of voluntary withdrawals under each payoff scenario. The percentages identify the fraction of observed play that was consistent with the panic equilibrium. The remaining percentages in each case identify the fraction of observed play that was consistent with the payoff dominant equilibrium. The number of observed strategies in each session is less than 20 (the number of subjects) because of forced withdrawals. As the table shows, all cases have low withdrawal rates in the first trial with forced withdrawals. Because of the smaller strategy set in the control treatment, these numbers are not directly comparable with those for the primary treatment in table III.¹⁵ Nevertheless, a casual inspection of the two tables suggests that the number of withdrawal opportunities has very little impact on the subjects' inclination to play according to the panic equilibrium in the initial trial with forced withdrawals.

	Voluntary Withdrawals	
	Scenario A	Scenario B
Session 1	2 of 14 (14.3%)	0 of 13 (0%)
Session 2	0 of 15 (0%)	2 of 15 (13.3%)
Combined	2 of 29 (6.9%)	2 of 28 (7.1%)

Table V: Trial 4 of the control treatment with forced withdrawals.

The fact that withdrawal rates rose in scenario B of session 2 is somewhat puzzling, but can perhaps be explained by higher withdrawal rates that were observed under scenario A payoffs in session 2, as shown in figure 3. In fact the tendency to panic continued to be higher in the later trials of scenario B payoffs in session 2, as shown in figure 4. The continued higher withdrawal rates in session 2 may have been further influenced by the fact that trials 15 and 16 of session 2 both had groups that realized 3 forced withdrawals, which never happened in session 1. Also, there was a bank failure in one group in trial 16 of session 2, while none occurred in the session 1 trials.

¹⁵Comparing the tables shows that the frequency of play is higher for *both* equilibria in the control treatment (6.9% and 93.1% versus 2.8% and 77.8% for scenario A). This is because observed play necessarily coincided with an equilibrium strategy in the control treatment, which was not true in the primary treatment.




Figure 3: Voluntary withdrawal rates for trials 4-11 of the control treatment.




Figure 4: Voluntary withdrawal rates for trials 15-18 of the control treatment.

In spite of the disparities between sessions 1 and 2 it is still apparent that subjects played strategies consistent with the payoff-dominant equilibrium more frequently than the panic equilibrium and overall withdrawal rates were lower under the scenario B payoffs than under scenario A. However, the issue of whether there is an increased tendency to play the panic equilibrium over time is less clear. We address this issue formally in the next section, where we compare the learning behavior of subjects across the two treatments.

IV. Treatment Comparison

We focus on trials 4 - 11, which had forced withdrawals and scenario A payoffs. In the first of these trials, there was no significant numerical difference between the withdrawal rate in the control treatment and the rate in the first withdrawal opportunity of the primary treatment. From tables III and V, the frequencies of voluntary withdrawals were 2 out of 29 and 1 out of 36, respectively. There was, however, a meaningful difference in the withdrawal rates over the subsequent trials. In particular, the voluntary withdrawal rate (the number of voluntary withdrawals divided by the number of subjects who were not forced to withdraw) rose significantly more quickly over time in the primary treatment.¹⁶ However, we are very reluctant to draw any strong conclusions from this statement. The tendency for some subjects to become more panicky over the course of the experiment might be influenced by their own personal history: how often they see others withdrawing and how many bank runs they observe. All of this contributes to their posterior view of how many players in the population are likely to play the panic strategy. The distribution of personal histories is unique to each experimental session. Even though the parameters are the same across sessions, variability in the outcomes of random forced withdrawals and random matching will produce different individual histories even for identical voluntary withdrawal rates. Hence, it is necessary to control for differences in personal histories in determining which treatment type has a larger learning effect on voluntary withdrawal rates.

Ideally, we would compare groups of individuals with identical histories across treatments, but this requires more data than we are able to obtain. Our approach is instead to construct a summary statistic that reflects an individual's history with respect to exposure to bank runs. The variable, called "history" is defined as the fraction of previous periods, before the current one, in which the subject witnessed a bank run. We want to allow for the possibility that subjects' interpretation of the history variable differs over time; in later trials, values of the history variable contain more information about the withdrawal tendencies of the population. Hence, in the regression analysis that follows we interact the history variable with the round variable.¹⁷

The table below shows the marginal effects from a Probit regression designed to test the null hypothesis that, controlling for differences in personal histories, there is no difference in withdrawal behavior across the two treatments. The dependent variable, "Withdraw," is equal to 1 if the subject voluntarily withdrew at the first opportunity. "Round" is a discrete variable that counts up from 1 to 7.¹⁸ "Treat Dum" is the treatment dummy, which equals 0 for the (control) treatment with a single withdrawal opportunity and 1 for the (primary) treatment with three withdrawal opportunities.

¹⁶The 95% confidence interval around the slope estimates of the trendline are (0.0016,0.0274) and (0.0299,0.0678) for treatments 1 and 2, respectively.

¹⁷There are, of course, other ways that one could construct the history variable. We would like to emphasize that we did not experiment with alternative specifications of this variable, and hence the reported significance levels are valid.

¹⁸Rounds 1 through 7 correspond to trials 5 through 11. The history variable is undefined in trial 4.

Withdraw	dF/dx	Std. Err.	z	P > z	x-bar
Round	-.0278256	.0177075	-1.59	0.111	4.0287
History	-.3234865	.2052944	-1.60	0.110	.373478
Treat Dum*	-.6982604	.1801047	-3.03	0.002	.565121
Treat Dum * Round	.0951046	.0327419	2.41	0.016	2.29801
Treat Dum * History	1.021354	.2449298	3.43	0.001	.253711
History * Round	.1245327	.0515573	2.51	0.012	1.68433
Treat Dum * History * Round	-.168838	.0599233	-2.53	0.011	1.1479
obs. P	.183223			Number of obs:	453
pred. P	.10986 (at x-bar)			Pseudo R ² :	.1915

*dF/dx is for discrete change of dummy variable from 0 to 1
z and P >|z| are the test of the underlying coefficient being 0.

Table VI: Results of Probit analysis.

A joint test of “Round” and “History * Round” has a p value of 0.0406, suggesting that history and round effects are significant in the control treatment. Moreover, we reject the null hypothesis that treatment type does not matter (a joint test of the variables “Treat Dum,” “Treat Dum * Round,” “Treat Dum * History,” and “Treat Dum * History * Round” has a p value of 0.0094). Hence, we conclude that the round and history variables impact withdrawal probabilities differently in the primary treatment than in the control. Due to the numerous interactions of the variables, it is not immediately transparent from table VI how each of the variables impacts the withdrawal probability. We expected that higher values for the history variable would translate into higher withdrawal probabilities. This is always true in the primary treatment and is true in all but the first two rounds of the control, where there is a slightly negative marginal effect. For rounds 3 and up, there is a positive relationship between history and withdrawals in the control. There is also a much more pronounced effect in the primary treatment than in the control. For instance, evaluated at the mean of the independent variables, the estimated marginal effect of an increase in the history variable is .52 in the primary treatment and .18 in the control. The joint test of the coefficients on Treat Dum * History and Treat Dum * History * Round has a p value of 0.0014, so this difference is significant at the 1% level.

As mentioned earlier, we also allowed for the possibility that subjects would react differently to the same level of the history variable over time. In fact, the analysis shows that people react more strongly to this variable over time, and that this effect is far more pronounced in the primary treatment than in the control. Starting from the mean of the independent variables, the estimated marginal effect of a unit increase in the round variable is 0.051 in the primary treatment and 0.019 in the control. The joint test of the coefficients on Treat Dum * Round and Treat Dum * History * Round has a p value of 0.0262, so this difference is significant at the 5% level.

Figures 5a-g and 6a-c illustrate these predictions for each possible history holding the round fixed, and for each round holding the history fixed (at 0, 0.5 and 1), respectively. In each case, separate predictions are charted for each treatment.

Figure 5a: Round 1




Figure 5b: Round 2




Figure 5c: Round 3




Figure 5d: Round 4




Figure 5e: Round 5




Figure 5f: Round 6




Figure 5g: Round 7




Figure 5: Estimated withdrawal probabilities for each round.

Figure 6a: History = 0




Figure 6b: History = .5




Figure 6c: History = 1




Figure 6: Withdrawal probabilities for selected histories.

The effects described above are apparent from the figures.¹⁹ In particular, figures 5c-g show that holding round constant at 3 or higher, voluntary withdrawal probabilities are positively related to the history variable and that the effect is much more prominent in the primary treatment than in the control. Figure 6a shows that holding history constant at 0, voluntary withdrawal probabilities are fairly constant over time, while figure 6c shows that holding history constant at 1, voluntary withdrawal probabilities increase over time. An intermediate case is shown in figure 6b. Note that in figure 6b all odd-numbered rounds are necessarily out-of-sample predictions, as 0.5 is not a possible history.

The combined implication of these figures is that we expect to see similar frequencies of

¹⁹Some caution is required in interpreting the figures, as some of the data points plotted in the figures are based on situations with low sample sizes.

bank runs early on in both treatment environments, but the frequency of observed bank runs should grow at a faster rate over time in the treatment with multiple withdrawal opportunities. This is very close to what we see in the data, as the next section illustrates.

V. Bank Run Analysis

A bank run is defined to have occurred if the bank liquidates all of its assets before the end of a trial. In scenario A this occurs if there are 3 or more withdrawals, while in scenario B this occurs if there are 4 or more withdrawals. No runs occurred without forced withdrawals in either scenario A or B. Bank runs occurred regularly (48% of the time) under scenario A payoffs with forced withdrawals, even though there was only 1/8 probability that forced withdrawal would cause the bank to fail. There were a small number of bank runs under scenario B payoffs with forced withdrawals (one occurrence in session 2 of the control treatment and one occurrence in the final round of each of the two sessions of the primary treatment) despite there being zero probability that forced withdrawals will cause the bank to fail.

Figure 7 shows bank run frequencies (out of 8 groups) that occurred in each trial of each treatment under scenario A payoffs with forced withdrawals. It is apparent from the figure that frequencies of bank runs rose much more quickly in the primary treatment (with 3 withdrawal opportunities) than the control (with 1 withdrawal opportunity), despite similar occurrences of bank runs in the early trials. This offers a compelling reason why the proportion of panicky subjects rose more quickly in the treatment with multiple withdrawal opportunities: the subjects in that treatment were exposed to more bank runs.




Figure 7: Bank run frequencies for each treatment, trials 4-11

VI. Conclusions

The key findings of this study as it relates to breakdowns in coordination and the occurrence of bank runs are that (1) forced withdrawals frequently lead to the occurrence of bank runs even when they do not present a likely threat to the solvency of the bank, and (2) it matters whether subjects are given multiple withdrawal opportunities (with feedback) or a single opportunity. Elaborating on point (2), subjects were equally likely to withdraw at the first opportunity in the first trial of both treatments, but exposure to bank runs had a greater (positive) effect on future withdrawals in the treatment with multiple withdrawal opportunities.²⁰ Of course, some caution is required in extrapolating these results to consumer behavior during financial crises. There does seem to be very strong evidence that the ability of people to coordinate on the payoff-dominant equilibrium is sensitive to forced withdrawals, which are meant to imitate aggregate uncertainty about individual-level shocks that occur during times of financial crises. However, the repeated play aspect of the experiment is somewhat artificial, and this is where the strongest treatment effect was observed. Nevertheless, it is interesting from a theoretical perspective to note that the standard approach of modelling bank runs using a simultaneous-move game may not be the most appropriate one. Moreover, the results suggest that in countries where people have a history of exposure to financial crises, withdrawal behavior might depend on the system in place for providing withdrawal opportunities and on the informational flow regarding the withdrawal activity of others.

Another key finding concerns the analysis of withdrawal behavior in the treatment with multiple withdrawal opportunities. We tested various cutoff rules for characterizing individual decisions. The experiment was specifically designed to differentiate between two such rules, one of which was consistent with Bayesian updating and one of which was not. We found that the cutoff rule associated with Bayesian updating outperformed the other rules. This effect disappeared over time, however, as increased rates of immediate withdrawals eventually made the panic rule – withdraw immediately – a superior predictor. As mentioned above, this does not suggest that subjects learned not to be Bayesian. Rather, it suggests they stopped believing that others would play their part of the payoff-dominant equilibrium.

References

- Allen, Franklin and Douglas Gale, 1998, Optimal financial crises, *Journal of Finance* 53, 1245-1284
- Boyd, John H., Sungkyu Kwak, and Bruce D. Smith (in press) The real output losses associated with modern banking crises, *Journal of Money, Credit, and Banking*.
- Cooper, Russell, Douglas V. DeJong, Robert Forsythe, and Thomas W. Ross, 1990, Selection criteria in coordination games: Some experimental results, *American Economic Review*

²⁰These findings are based on rounds with forced withdrawals and scenario A payoffs. The the rounds with scenario B payoffs or no forced withdrawals had very little variation in individual withdrawal behavior within or across treatments; very few individuals withdrew in these rounds.

80, 218-233

- Cooper, Russell and Thomas W. Ross, 1998, Bank runs: liquidity costs and investment distortions, *Journal of Monetary Economics* 41, 27-38
- Diamond, Douglas W. and Philip H. Dybvig, 1983, Bank runs, deposit insurance, and liquidity, *Journal of Political Economy* 91, 401-419
- Ennis, Huberto M., 2003, Economic fundamentals and bank runs, *Federal Reserve Bank of Richmond Economic Quarterly* 89, 55-71
- Ennis, Huberto M. and Todd Keister (in press) Bank runs and investment decisions revisited, *Journal of Monetary Economics*.
- Goldstein, Itay and Ady Pauzner, 2005, Demand-deposit contracts and the probability of bank runs, *Journal of Finance* 60, 1293-1327
- Gorton, Gary, 1988, Banking panics and business cycles, *Oxford Economic Papers* 40, 751-781
- Ochs, Jack, 1995, Coordination problems, in A.E. Roth and J.H. Kagel, eds.: *Handbook of Experimental Economics* (Princeton University Press)
- Peck, James and Karl Shell, 2003, Equilibrium bank runs, *Journal of Political Economy* 111, 103-123
- Saunders, Anthony and Berry Wilson, 1996, Contagious bank runs: Evidence from the 1929-1933 period, *Journal of Financial Intermediation* 5, 409-423
- Schotter, Andrew and Tanju Yorulmazer, 2003, On the severity of bank runs: An experimental study, mimeo.
- Van Huyck, John B., Raymond C. Battalio, and Richard O. Beil, 1990, Tacit coordination games, strategic uncertainty, and coordination failure, *American Economic Review* 80, 234-248
- Wallace, Neil, 1990, A banking model in which partial suspension is best, *Federal Reserve Bank of Minneapolis Quarterly Review* 14, 11-23

Bank Deposit Experiment INSTRUCTIONS

This experiment has been designed to study decision-making behavior in groups. If you follow the instructions carefully and make good decisions, you may earn a considerable amount of money. The participants may earn different amounts of money in this experiment because each participant's earnings are based partly on his/her decisions and partly on the decisions of the other group members. The money you earn will be paid to you, in cash, at the end of the experiment. Therefore, it is important that you do your best. A research foundation has contributed the money to conduct this study.

Description of the Task

You and four other people each have \$1 deposited in an experimental account. You must decide whether to request to withdraw your \$1 at any one of THREE withdrawal opportunities you will be given, or leave it deposited in the account.

How much money you will receive if you make a withdrawal request or if you leave your money deposited depends on the withdrawal decisions of the other four people in the experiment. This is explained below. Withdrawal opportunities are numbered using roman numerals I through III.

Withdrawal Opportunity I. Below is a chart that you can use to figure out the payoffs associated with your withdrawal decision in Withdrawal Opportunity I. You will see this chart on your computer screen when the experiment begins. Remember, how much you receive if you make a withdrawal request or how much you earn by leaving your money deposited depends on how many other people place withdrawal requests. The chart gives you payoffs for all the possible numbers of requests. The word "hypothetical" is used in the chart because you do not know how many withdrawal requests will be made when you make your decision. If TWO or fewer withdrawal requests are made then each requester will receive \$1 and each remaining depositor will have a projected payment of \$1.50. The projected payment is the amount each remaining depositor will receive if there are no more withdrawal requests in the remaining two withdrawal opportunities. If there are future withdrawals, remaining depositors might get less, as the following charts will show. If THREE or more withdrawal requests are made then each requester will receive \$1 or less as shown in the chart, and the remaining depositors will get \$0.

Hypothetical number of new withdrawal requests	Amount each requester would receive	Projected payment to each remaining depositor
0	not applicable	\$1.50
1	\$1	\$1.50
2	\$1	\$1.50
3	\$1	\$0
4	\$0.75	\$0
5	\$0.60	not applicable

Withdrawal Opportunity II. The payoff chart for Withdrawal Opportunity II depends on the number of withdrawal requests made in Withdrawal Opportunity I. Below is the payoff chart that would apply if 1 withdrawal request were made during Withdrawal Opportunity I. Now the highest possible number of new requests is 4, so the chart has 1 less row than before. The projected payment assuming ONE or fewer withdrawal requests is \$1.50. However, now if there are TWO or more withdrawal requests remaining depositors get \$0. The amount each requester receives is \$1 for up to two new withdrawals and less than that for more than two withdrawals, as shown in the chart.

Hypothetical number of new withdrawal requests	Amount each requester would receive	Projected payment to each remaining depositor
0	not applicable	\$1.50
1	\$1	\$1.50
2	\$1	\$0
3	\$0.67	\$0
4	\$0.50	not applicable

Withdrawal Opportunity III. The payoff chart for Withdrawal Opportunity III depends on the number of withdrawal requests made in withdrawal opportunities I and II. At the beginning of Withdrawal Opportunity III you will again see a payoff table that reflects the previous withdrawals and shows the projected payments corresponding to any additional withdrawals. Now, since this is the last withdrawal opportunity, the projected payments corresponding to each hypothetical number of new withdrawal requests will be the actual payments.

Where do these payoffs come from? It is not important that you fully understand how the numbers are determined. However, the underlying story is that the account manager has invested the \$15 from the experimental account in assets that cannot be converted to cash before the end of the trial without paying a penalty. The dollar amounts you see reflect the ability of the account manager to meet her obligations of paying requesting individuals \$1 (if possible) during the withdrawal opportunities and up to \$1.50 to remaining depositors at the end of the trial.

Procedure

You will perform the task described above numerous times. Each time is called a trial. Each trial is completely separate. That is, you will start each trial with \$1 in the experimental account. You will keep the money you earn in every trial. At the end of each trial, your earnings for that trial and your total earnings will appear on your computer screen.

You do not play with the same people each trial. New groups of five are formed randomly every trial out of the twenty people participating in the experiment.

At the beginning of each withdrawal opportunity, you will be shown a screen similar to the pictorial representation below.

Trial A1		
Withdrawal Opportunity I		
Payoff Table		
Hypothetical number of new withdrawal requests	Amount each requester would receive	Projected payment to each remaining depositor
0	not applicable	\$1.50
1	\$1	\$1.50
2	\$1	\$1.50
3	\$1	\$0
4	\$0.75	\$0
5	\$0.60	not applicable

Time remaining in Withdrawal Opportunity I: **30** seconds

To make a withdrawal request click the "Withdraw Now" button at the bottom of the page before time expires. If you do not click the "Withdraw Now" button before the time expires your money will remain deposited.

Withdraw Now

You can make a withdrawal request by clicking the “Withdraw Now” button before time expires. At each withdrawal opportunity you will be given 30 seconds to make your decision. If you do not click the “Withdrawal Now” button your money will remain deposited and you will either advance to the next withdrawal opportunity. If it is Withdrawal Opportunity III the trial will end and you will receive the payoff to remaining depositors.

At the end of each withdrawal opportunity you will see a summary that lists the number of new withdrawal requests that were placed during that opportunity, the amount each requester received, the number of remaining depositors, and the projected payment to remaining

depositors. A pictorial representation of a possible summary following Withdrawal Opportunity I is provided below.

Trial A1

Opportunity	New requests	Amount received	Remaining depositors	Projected payment to each remaining depositor
I	1	\$1.00	4	\$1.50

Withdrawal Opportunity I is over.

Continue

The other people in the experiment will also view the same screens.

At the end of each trial you will see a summary that lists your earning for the trial and your cumulative earnings for the experiment (not including the show-up fee).

Trial Variations

There are two types of trials. Type A trials are played as described above. Type B trials involve a randomly determined number of forced withdrawals each withdrawal opportunity. The specific rules for the type B trials will be reviewed as these trials are reached during the experiment.

Payment at the End of the Session

You will participate in a maximum of 18 trials. The first two trials will be unpaid practice trials. At the end of the entire experiment, the supervisor will pay you your earnings in cash.

Please remember, communicating with other people during the experiment is strictly forbidden.

Thank you for your participation.

TRIAL SUMMARY

Trial A1-A2: no forced withdrawals (unpaid)

Trials A3-A4: no forced withdrawals

Trials B1-B8: At the beginning of EACH withdrawal opportunity there is a 50% chance that one of the remaining depositors will be forced to withdraw. When they occur, forced withdrawals are randomly assigned.

New payoffs

Hypothetical number of new withdrawal requests	Amount each requester would receive	Projected payment to each remaining depositor
0	not applicable	\$1.50
1	\$1	\$1.50
2	\$1	\$1.50
3	\$1	\$1.50
4	\$1	\$0
5	\$0.80	not applicable

Trial A5: no forced withdrawals (unpaid)

Trials A6-A7: no forced withdrawals

Trials B9-B12: These trials use the same rules as B1-B8, but with new payoffs. At the beginning of EACH withdrawal opportunity there is a 50% chance that one of the remaining depositors will be forced to withdraw. Once again, any forced withdrawals are randomly assigned.

Bank Deposit Experiment INSTRUCTIONS

This experiment has been designed to study decision-making behavior in groups. If you follow the instructions carefully and make good decisions, you may earn a considerable amount of money. The participants may earn different amounts of money in this experiment because each participant's earnings are based partly on his/her decisions and partly on the decisions of other group members. The money you earn will be paid to you, in cash, at the end of the experiment. Therefore, it is important that you do your best. A research foundation has contributed the money to conduct this study.

Description of the Task

You and four other people each have 1 experimental dollar deposited in an experimental bank. You must decide whether to request to withdraw your \$1 or leave it deposited. At the end of the experiment we will be pay you \$0.50 US for each experimental dollar you earn during the experiment.

How much money you will receive depends on your own decision and on the decisions of the other four people in your group. This is explained below.

Withdrawal Decision. You will see the chart below on your computer screen when the experiment begins. How much you receive if you make a withdrawal request or how much you earn by leaving your money deposited depends on how many other people in your group place withdrawal requests. The chart lists the payoffs for all the possible numbers of requests. The word "hypothetical" is used in the chart because you do not know how many withdrawal requests will be made when you make your decision. If ONE or TWO withdrawal requests are made, each requester will receive \$1 and the remaining depositors will get \$1.50. If THREE or more withdrawal requests are made, each requester will receive \$1 or less, as shown in the chart, and the remaining depositors will get \$0.

Hypothetical number of new withdrawal requests	Amount each requester would receive	Payment to each remaining depositor
0	not applicable	\$1.50
1	\$1	\$1.50
2	\$1	\$1.50
3	\$1	\$0
4	\$0.75	\$0
5	\$0.60	not applicable

Where do these payoffs come from? It is not important that you fully understand how the numbers are determined. However, the underlying story is that the account manager has invested the \$5 from the experimental account in assets that cannot be converted to cash before the end of the trial without paying a penalty. The dollar amounts you see reflect the ability of the account manager to meet her obligations of paying requesting individuals \$1 (if possible) during the withdrawal opportunity and up to \$1.50 to remaining depositors at the end of the trial.

Procedure

You will perform the task described above numerous times. Each time is called a trial. Each trial is completely separate. That is, you will start each trial with \$1 in the experimental bank. You will keep the money you earn in every trial. At the end of each trial, your earnings for that trial and your total earnings will appear on your computer screen.

You do not play with the same people each trial. New groups of five are formed randomly every trial from the twenty people participating in the experiment.

At the beginning of each trial, you will be shown a screen similar to the pictorial representation below. The title “Withdrawal Opportunity I” will be on your screen, suggesting that there might be additional withdrawal opportunities (i.e., II, III, etc.). This is not the case. There is only ONE withdrawal opportunity per trial.

Trial A1
Withdrawal Opportunity I

Payoff Table

Hypothetical number of new withdrawal requests	Amount each requester would receive	Projected payment to each remaining depositor
0	not applicable	\$1.50
1	\$1	\$1.50
2	\$1	\$1.50
3	\$1	\$1.50
4	\$0.75	\$0
5	\$0.60	not applicable

Time remaining in Withdrawal Opportunity I: **30** seconds

To make a withdrawal request click the "Withdraw Now" button at the bottom of the page before time expires. If you do not click the "Withdraw Now" button before the time expires your money will remain deposited.

Withdraw Now

You can make a withdrawal request by clicking the “Withdraw Now” button before time expires. You will be given 30 seconds to make your decision. If you do not click the “Withdrawal Now” button, your money will remain deposited.

At the end of each trial you will see a summary that lists the number of withdrawal requests that were placed during the trial, the amount each requester received, the number of remaining depositors, and the payment to remaining depositors. A pictorial representation of a possible summary following Trial A1 is provided below.

Trial A1

Opportunity	Withdrawal requests	Amount received	Remaining depositors	Payment to each remaining depositor
I	1	\$1.00	4	\$1.50

Withdrawal Opportunity I is over.

Continue

The other people in the experiment will also view the same screens.

You will also see a summary that lists your earning for the trial and your cumulative earnings for the experiment (not including the show-up fee).

Trial Variations

There are two types of trials. Type A trials are played as described above. In Type B trials, some people may be randomly chosen and forced to withdraw. The specific rules for the type B trials will be discussed as these trials are reached during the experiment.

Payment at the End of the Session

You will participate in a maximum of 31 trials. The first two trials will be unpaid practice trials. At the end of the entire experiment, the supervisor will pay your earnings to you in cash.

Please remember, communicating with other people during the experiment is strictly forbidden.

Thank you for your participation.

TRIAL SUMMARY

Trail A1-A2: no forced withdrawals (unpaid)

Trails A3-A4: no forced withdrawals

Trials B1-B8: At the beginning of EACH trail there is a $\frac{3}{8}$ chance that one depositor will be forced to withdraw, a $\frac{3}{8}$ chance that two depositors will be forced to withdraw, and a $\frac{1}{8}$ chance that three depositors will be forced to withdraw. When they occur, forced withdrawals are randomly assigned.

New payoffs

Hypothetical number of new withdrawal requests	Amount each requester would receive	Payment to each remaining depositor
0	not applicable	\$1.50
1	\$1	\$1.50
2	\$1	\$1.50
3	\$1	\$1.50
4	\$1	\$0
5	\$0.80	Not applicable

Trial A5: no forced withdrawals (unpaid)

Trials A6-A7: no forced withdrawals

Trials B9-B12: These trials use the same rules as B1-B8, but with new payoffs. At the beginning of EACH trail there is a $\frac{3}{8}$ chance that one depositor will be forced to withdraw, a $\frac{3}{8}$ chance that two depositors will be forced to withdraw, and a $\frac{1}{8}$ chance that three depositors will be forced to withdraw. Once again, any forced withdrawals are randomly assigned.

Appendix B

Optimal Actions with Bayesian Updating, Liquidation Value = .6.

Withdrawal Opportunity III								
Number of previous withdrawals	Posterior prob of a forced withdrawal in this opportunity	Expected payoff from not withdrawing			Expected payoff from withdrawing			Optimal <u>Action</u>
		# of forced withdrawals		Expected	# of forced withdrawals		Expected	
		0	1	payoff	0	1	payoff	
0	0.44	1.50	1.50	1.50	1.00	1.00	1.00	not withdraw
1	0.43	1.50	1.50	1.50	1.00	1.00	1.00	not withdraw
2	0.40	1.50	0.00	0.90	1.00	0.50	0.80	not withdraw

Withdrawal Opportunity II								
Number of previous withdrawals	Posterior prob of a forced withdrawal in this opportunity	Expected payoff from not withdrawing			Expected payoff from withdrawing			Optimal <u>Action</u>
		# of forced withdrawals		Expected	# of forced withdrawals		Expected	
		0	1	payoff	0	1	payoff	
0	0.44	1.45	1.44	1.44	1.00	1.00	1.00	not withdraw
1	0.43	1.44	0.92	1.21	1.00	1.00	1.00	not withdraw
2	0.40	0.92	0.00	0.55	1.00	0.50	0.80	withdraw

Withdrawal Opportunity I								
Number of previous withdrawals	Posterior prob of a forced withdrawal in this opportunity	Expected payoff from not withdrawing			Expected payoff from withdrawing			Optimal <u>Action</u>
		# of forced withdrawals		Expected	# of forced withdrawals		Expected	
		0	1	payoff	0	1	payoff	
0	0.44	1.40	1.19	1.31	1.00	1.00	1.00	not withdraw

Game ends if there are 3 withdrawals. Players factor in probability of being forced to withdraw in future rounds when calculating expected payoff to not withdrawing. Optimal actions are determined under the assumption that all other agents do not withdraw unless withdrawing is a dominant strategy. Situations that are shaded necessarily represent off-equilibrium behavior.

Appendix C

Optimal Actions with Bayesian Updating, Liquidation Value = .8.

Withdrawal Opportunity III								
Number of previous withdrawals	Posterior prob of a forced withdrawal in this opportunity	Expected payoff from not withdrawing			Expected payoff from withdrawing			Optimal Action
		# of forced withdrawals		Expected payoff	# of forced withdrawals		Expected payoff	
		0	1		0	1		
0	0.44	1.50	1.50	1.50	1.00	1.00	1.00	not withdraw
1	0.43	1.50	1.50	1.50	1.00	1.00	1.00	not withdraw
2	0.40	1.50	1.50	1.50	1.00	1.00	1.00	not withdraw
3	0.33	1.50	0.00	1.00	1.00	0.50	0.83	not withdraw

Withdrawal Opportunity II								
Number of previous withdrawals	Posterior prob of a forced withdrawal this period	Expected payoff from not withdrawing			Expected payoff from withdrawing			Optimal Action
		# of forced withdrawals		Expected payoff	# of forced withdrawals		Expected payoff	
		0	1		0	1		
0	0.44	1.45	1.44	1.44	1.00	1.00	1.00	not withdraw
1	0.43	1.44	1.42	1.43	1.00	1.00	1.00	not withdraw
2	0.40	1.42	1.00	1.25	1.00	1.00	1.00	not withdraw
3	0.33	1.00	0.00	0.67	1.00	0.50	0.83	withdraw

Withdrawal opportunity I								
Number of previous withdrawals	Posterior prob. # of forced w/d this period = 1	Expected payoff from not withdrawing			Expected payoff from withdrawing			Optimal Action
		# of forced withdrawals		Expected payoff	# of forced withdrawals		Expected payoff	
		0	1		0	1		
0	0.44	1.40	1.38	1.39	1.00	1.00	1.00	not withdraw

Game ends if there are 4 withdrawals. Players factor in probability of being forced to withdraw in future rounds when calculating expected payoff to not withdrawing. Optimal actions are determined under the assumption that all other agents do not withdraw unless withdrawing is a dominant strategy. Situations that are shaded necessarily represent off-equilibrium behavior.