

Bank Runs as Coordination Failures: An Experimental Study^{*}

Rod Garratt

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

Todd Keister

Research and Statistics Group, Federal Reserve Bank of New York
todd.keister@ny.frb.org

October 12, 2006

Abstract

We use experimental methods to investigate what factors contribute to breakdowns in coordination among a bank's depositors. Subjects in our experiment decide whether to leave their money deposited in a bank or withdraw it early; a bank run occurs when there are too many early withdrawals. We explore the effects of adding uncertainty about fundamental withdrawal demand and of changing the number of opportunities subjects have to withdraw. Our results show that (i) bank runs are rare when fundamental withdrawal demand is known but occur frequently when it is stochastic, and (ii) subjects are more likely to withdraw when given multiple opportunities to do so than when presented with a single decision. For the multiple-opportunity case, we evaluate individual withdrawal decisions according to a set of 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.

* We thank the staff of 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 are also grateful to Kelly Bedard, Ted Bergstrom, Gary Charness, Jon Sonstelie, and participants at the 2004 Experimental and Behavioral Economics Workshop in Calgary and the 2005 World Congress of the Econometric Society in London for helpful comments and suggestions. This work was funded by UC-MEXUS-Conacyt grant No. CN-02-77, and was completed while Keister was at the Instituto Tecnológico Autónomo de México (ITAM), whose support is gratefully acknowledged. The views expressed herein are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of New York or the Federal Reserve System.

1 Introduction

Bank runs are an important economic phenomenon. They occurred frequently in the United States before the mid-1930s and continue to occur around the world; recent episodes include a partial run in Russia in 2004 and a run that temporarily closed the banking sector in Argentina in 2001. A sizable theoretical literature has attempted to shed light on the underlying cause of these runs. One of the leading explanations that has been offered is that a run results from a coordination failure. The seminal paper of Diamond and Dybvig (1983) showed how the game played by a bank's depositors naturally has multiple equilibria. In one, the level of withdrawal demand is "normal" and depositors only withdraw their funds if they need to. In the other equilibrium, however, all depositors rush to withdraw because they fear the bank will run out of funds. These actions cause the bank to fail, fulfilling the original beliefs. A bank run can then be interpreted as a switch from the good equilibrium to the bad one.¹

We use experimental methods to test the extent to which breakdowns in coordination can lead to bank runs. Our goal is to determine the plausibility of this explanation in a laboratory setting, as well as to investigate what factors make failures in coordination more or less likely to occur in this context. The subjects in our experiment play the role of depositors in a bank; each chooses 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 each 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 and changing the number of opportunities subjects have to withdraw.

The possibility that some subjects will be forced to withdraw is intended as a proxy for macroeconomic conditions; in bad times, more depositors need to withdraw funds from their banks, which places a "squeeze" on the amount of liquidity available to meet further withdrawals. Models with a random number of forced withdrawals have been studied extensively in the theoretical

¹ There is a sizable literature based on the Diamond-Dybvig model. See, for example, Wallace (1988,1990), Cooper and Ross (1998), Green and Lin (2003), Peck and Shell (2003), and Ennis and Keister (2006). Another explanation of bank runs is based on the release of negative information about the value of the portfolio of an individual bank or of the entire banking system (see, for example, Gorton, 1988, and Saunders and Wilson, 1996). It is unclear which of these explanations is better supported by historical data, and hence investigating the plausibility of the coordination failure explanation in the laboratory is an important exercise.

literature (see Wallace 1988, Green and Lin, 2003, and Peck and Shell, 2003). However, the theory does not offer any guidance as to whether or not depositors would behave differently in this situation, since the set of equilibria is qualitatively the same (both a “run” and a “no-run” equilibrium exist) with and without random forced withdrawals. Our interest is in whether the existence of these shocks, which involve aggregate uncertainty, makes coordination failure more likely.

The answer turns out to be tightly linked to our second treatment variable: the number of opportunities subjects have to withdraw. In half of our experimental sessions, subjects played a simultaneous-move coordination game. Each subject decided whether to withdraw early or to wait, and then the game ended and payoffs were assigned. In the other half of the sessions, however, we gave subjects three opportunities to withdraw before the game ended, informing them of the total number of withdrawals after each opportunity. This treatment gave subjects the option of waiting and observing some information about the actions of others and the aggregate shock before making a final decision. The information they received this way was partial, since they were not told if an observed withdrawal was forced or voluntary. It was also not costless, since the bank could run out of funds while they waited if too many other subjects withdrew early. This treatment adds a realistic feature of banking: depositors have a period of time during which they can choose to withdraw their funds, and they are able to observe some information about the actions of other depositors, for example, by noticing if a line is forming outside the bank. They cannot, however, observe the reason why another depositor is withdrawing (whether she “needs” to withdraw or is panicking); motives for withdrawing are private information.

The theory is silent on how both of our treatment variables should affect play, as the set of equilibria is qualitatively the same under any combination of the number of withdrawal opportunities (multiple or single) and the presence or absence of forced withdrawals. However, our empirical findings are unambiguous. In the absence of forced withdrawals, voluntary withdrawals are rare and subjects effectively coordinate on the no-run outcome. Adding forced withdrawals has a positive, but small, initial effect on voluntary withdrawal rates regardless of the number of withdrawal opportunities. Over time, the interaction of forced withdrawals and multiple withdrawal opportunities leads to high withdrawal rates and almost total bank failures, something that does not happen with forced withdrawals alone.

The difference in the occurrence of bank failures across treatments is traced to differences in subjects’ reactions to their exposure to bank runs across treatments. With forced withdrawals,

some banks runs occur even when voluntary withdrawal rates are low. Occasionally, an unfortunate combination of a high realization of the number of forced withdrawals and, possibly, a few voluntary withdrawals depletes the bank's assets. In the treatment with a single withdrawal opportunity, exposure to such a bank run has almost no effect on future withdrawal behavior. As a result, the total number of bank runs observed over time is low and relatively constant. In the treatment with multiple withdrawal opportunities, on the other hand, exposure to a bank run has a "snowballing" effect that leads to future bank failures. These results show that bank runs are more likely to occur in environments where (i) there is significant uncertainty about fundamental withdrawal demand and (ii) depositors receive information about the behavior of other depositors while there is still time to withdraw.²

In order to better understand the forces generating the observed group outcomes in the treatment with multiple withdrawal opportunities, we evaluate individual 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.³ However, as subjects' belief that the payoff dominant equilibrium will be played deteriorates, we naturally see a deterioration of this Bayesian play as well.

While there has been much experimental work on coordination games,⁴ we know of only two other studies that have conducted an experimental investigation of bank runs: Schotter and Yorulmazer (2005) and Madies (2006). Although our experimental design differs substantially from both of these studies, our results have elements in common with each of them. Schotter and Yorulmazer (2005) study the factors that affect the severity of a 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,

² This finding is loosely connected to the theoretical work of Gale (1995), which studies a dynamic game in which players decide when to make an investment. In Gale's setting the coordination problem vanishes as the period length shrinks to zero. This resembles our finding that groups converge to a lower run frequency in the treatment with a single withdrawal opportunity. There is also a connection to the theoretical work of Dutta (2003). He shows that allowing endogenously determined asynchronous moves mitigates the coordination problem. In our treatment with three withdrawal opportunities, subjects can choose when to withdraw but they cannot commit not to withdraw. Hence, the rationale that leads to coordination in the Dutta study does not apply to our games.

³ This aspect of the design was meant to capture whether people update their perception of the general well-being of the population in response to positive (or negative) personal shocks. The experiment was designed so that subjects did not need to update perfectly in order to make the correct choice.

⁴ See, for example, Cooper *et al.* (1990), Van Huyck *et al.* (1990), and the surveys by Ochs (1995) and Cooper (1999).

asymmetric information, etc.) affect this speed. Our primary focus, in contrast, is on whether or not a run occurs at all and what factors affect the prevalence of runs. Despite this difference in focus, the results share an important theme: 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 modeling bank runs as a one-shot, simultaneous-move game (as in Diamond and Dybvig, 1983, and many others) may not be the most appropriate. A recent paper by Madies (2006) does study the prevalence of bank runs in a setting where a no-run equilibrium always exists. Among other findings, the paper shows that bank runs occur less frequently when the bank is more “liquid” in the sense that a larger number of early withdrawals is needed to make the bank insolvent. This relationship is also supported by our results. Beyond this point, however, the papers diverge. Madies (2006) focuses on the effects of partial deposit insurance schemes, while we study the effects of uncertain fundamental withdrawal demand and multiple withdrawal opportunities.

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 1997-8.⁵

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 3 explains the theoretical predictions and presents the results for the case where there is a small chance of a bank failure due to forced withdrawals alone. These results include an econometric analysis of treatment effects, as well as a description of how often bank runs occurred in each of the treatments and how the escalation in the frequency of bank runs differed across the treatments. Section 4 provides a classification of individual decisions according to various cutoff rules. Section 5 explains the theoretical predictions and presents the results for the case where there is no chance of a bank failure due to forced withdrawals. Finally, Section 6 contains some concluding remarks.

⁵ See Boyd *et al.* (2006) for a detailed analysis of the available data on modern banking and financial crises.

2 Experimental Design

We devised a computer-controlled experiment in which subjects play multiple rounds of a withdrawal game with varying strategy sets and payoffs. We ran four experimental sessions, each involving 20 subjects who participated in 19 rounds (3 unpaid and 16 paid). In each round, subjects were randomly and anonymously divided into groups of 5 and were required to make withdrawal decisions. In two of the four sessions, subjects had a single opportunity to withdraw in each round. In the other two sessions subjects had three withdrawal opportunities per round. Payoffs in each round were specific to the play in that round. At the end of each round, feedback was given to all subjects regarding the withdrawal decisions of other players in their group, as well as their resulting payoffs. Groups were reshuffled after every round. Hence, subjects learned about the withdrawal tendencies of the session population, but they did not know the specific tendencies of the members in their group in any round. This design allows subjects to gain experience so that we can look at learning effects, but does not generate group-specific behavior.

We begin by describing the treatment involving a single withdrawal opportunity in each round. The instructions for these sessions are provided in Appendix A.

2.1 Single Withdrawal Opportunity

Subjects began each round with one dollar deposited in their group's bank. At the beginning of a round, subjects were shown a chart that described their payoffs from withdrawing or not withdrawing as a function of the total number of withdrawals. This payoff chart is reproduced in the instructions provided in Appendix A. Subjects had 30 seconds to withdraw their dollar. If they did not, it remained deposited. The experiment had two stages corresponding to different payoff specifications. Payoffs in the first stage of the experiment were defined so that the bank was able to absorb 2 withdrawals (*i.e.*, give two people their dollar back) without defaulting on its promise to pay remaining depositors \$1.50. If one or two subjects placed withdrawal requests, they each got their dollar back and the remaining depositors each were paid \$1.50, as promised. If three subjects placed withdrawal requests, they each got their dollar back, but this required liquidating all of the bank's assets, and hence remaining depositors received nothing. If more than three subjects put in withdrawal requests, the bank's assets were completely liquidated (at a rate of \$0.60 on the dollar) and all requesters were given equal share. Again, any depositors who had not placed a withdrawal request received nothing.

Stage 1 began with four rounds (2 unpaid and 2 paid) without forced withdrawals. These rounds were conducted to (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.⁶ We then conducted eight rounds with random forced withdrawals. In these rounds, the computer first randomly selected the number of forced withdrawals, and then randomly chose that number of subjects and submitted a withdrawal request on their behalf.⁷ The probability distribution over forced withdrawals in stage 1 was chosen so that there was a small probability that forced withdrawals would cause the bank to fail (this will not be the case in stage 2.). In particular, 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. At the end of each round, subjects were told the number of withdrawals in their group, but not whether withdrawals by others were forced or voluntary.

After stage 1 was completed, we conducted a second stage using payoffs that allowed the bank to absorb 3 withdrawals. This was done by setting the liquidation rate at \$0.80, so that each withdrawal caused a smaller reduction in the bank's assets. We conducted 2 rounds (1 unpaid and 1 paid) with the new payoffs without forced withdrawals. Then we conducted 4 paid rounds with forced withdrawals using the same withdrawal probabilities as in the first stage. Notice that in stage 2 there is *no* chance that forced withdrawals alone will cause the bank to fail. Any form of communication during the experiment was strictly forbidden.

2.2 Three Withdrawal Opportunities

The other two sessions followed the same treatment plan described above, but subjects had three opportunities to withdraw their \$1 within each round. (Instructions are provided in Appendix B.). Withdrawal requests in each opportunity were treated equally: subjects got their dollar back, or if there was not enough money to do this at the specified liquidation rate, they equally shared the liquidated value of the bank's assets. Any funds not liquidated during a withdrawal opportunity were used to meet withdrawal requests in subsequent opportunities and to pay remaining depos-

⁶ We believed that two unpaid trials followed by two paid trials would be enough to achieve these goals and, as we describe below, this was indeed the case.

⁷ As mentioned in the introduction, these 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 (1988, 1990), Peck and Shell (2003), and others.

itors at the end of the round. Payoffs thus followed a quasi-sequential service rule: within each opportunity requesters were treated identically, but across opportunities depositors who requested to withdraw first were served first.

After each withdrawal 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. At the end of the final opportunity, subjects received information on their group's outcome for that round, i.e., how many people withdrew, the payoff to withdrawers, the payoff to remaining depositors, and a report of their cumulative individual earnings.

Once a player withdrew, she had no more decisions to make in that round. 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, all members of that group were informed that the bank was out of money and told to wait until the beginning of the next round.⁸

Forced withdrawals occurred over the three withdrawal opportunities. There was a 1/2 probability that one subject would be forced to withdraw and a 1/2 probability that no one would be forced to withdraw in *each opportunity*.⁹ These probabilities were chosen so that the cumulative distribution of forced withdrawals over the course of a round is the same in the sessions with a single withdrawal opportunity. This facilitates a meaningful comparison across treatments.

2.3 Participation and Earnings Summary

Eighty undergraduate students from the University of California, Los Angeles participated in the four sessions of the experiment. In addition to their earnings in the experiment, players received a \$5.00 show-up fee. Total earnings were \$517.00 and \$495.00 for sessions 1 and 2, respectively, of the treatment with a single withdrawal opportunity and \$478.75 and \$445.00 for sessions 1 and 2 of the treatment with three withdrawal opportunities. The gap between the highest and lowest payoff was between \$3.25 and \$5.00 in each session.

⁸ Information about other groups in the session was not provided. However subjects in a group that went bankrupt might infer whether or not other groups had gone bankrupt from how long they had to wait between rounds. Of course, everyone had to wait the full time if only one group did not go bankrupt and, hence, information that could be inferred from the wait time was imperfect.

⁹ Note that because the number of forced withdrawals is independent across opportunities, the observed number of withdrawals in one opportunity contains no information about the number of future forced withdrawals. See Chari and Jagannathan (1988) and Jacklin and Bhattacharya (1988) for models where observed withdrawals do potentially contain information about the state of the nature, which in their models is the (random) return on the bank's assets.

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.

3 Low Risk of Bank Failure from Forced Withdrawals

In this section we examine the first stage of each session, where the bank can absorb two withdrawals without defaulting on its obligations to the remaining depositors. When forced withdrawals are added to this stage, there is a small chance that the forced withdrawals alone will cause the bank to fail. We first describe the equilibria of this game under each treatment scenario: with single or multiple withdrawal opportunities and in the presence or absence of forced withdrawals. We then present the experimental results.

3.1 Equilibrium Analysis

In the games with a single withdrawal opportunity there are exactly two pure-strategy Nash equilibria, both with and without forced withdrawals. In the payoff-dominant equilibrium no player voluntarily withdraws. In the “banking panic” equilibrium all players withdraw.

In the games with three withdrawal opportunities the strategy of a player 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 stage 1 the game ends if there are three withdrawal requests, so 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 .

Without forced withdrawals, 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.

With forced withdrawals there is a 50% chance in each withdrawal opportunity that one player will be selected at random and forced to withdraw. There exists a payoff-dominant equilibrium in this game that is similar to the one without forced withdrawals – the only difference is that players have a lower panic threshold for observed withdrawals going into the second opportunity. The 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 verify this, we again start by considering the decision facing a player in withdrawal opportunity III. Recall that this opportunity is only reached if the number of previous withdrawals 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

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

This probability falls from 4/9 to 1/3 as k drops from 5 to 2. However, 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. 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 C, which shows payoffs and optimal actions under each withdrawal scenario.¹⁰ 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

¹⁰ The subjects were not shown anything like the table in Appendix C. During each withdrawal opportunity subjects saw only the relevant, updated payoff tables (see the instructions in Appendices A and B).

withdrawals. Finally, if a player believes that all other players will follow the strategy given above, her optimal action in the first withdrawal opportunity is to not withdraw. Therefore, the strategy profile listed above is a subgame perfect equilibrium of the game with forced withdrawals.

As in the case of a single withdrawal opportunity, there is also a 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, leaving nothing for a player who did not withdraw. Notice that the panic equilibrium will exist regardless of the presence or absence of forced withdrawals. There are no other pure-strategy equilibria.¹¹

3.2 Observed Behavior

A. No Forced Withdrawals. In the rounds with no forced withdrawals, we observed the payoff-dominant Nash equilibrium in 16 out of 16 games in the sessions with a single withdrawal opportunity. In the sessions with 3 withdrawal opportunities the payoff-dominant subgame-perfect equilibrium occurred in only 8 out of 16 games, however 72 out of 80 subjects (90 percent) played strategies consistent with the payoff dominant equilibrium.¹² Hence, while having multiple withdrawal opportunities appears to have a positive effect on withdrawal rates in the absence of forced withdrawals, none of the banks in which withdrawals occurred defaulted on the promise to pay \$1.50 to remaining depositors. In other words, there were no observed bank runs in either treatment.

B. Forced Withdrawals. In the first round with forced withdrawals, there was no significant numerical difference between the withdrawal rates in the treatment with a single withdrawal opportunity and that in the first opportunity of the treatment with three withdrawal opportunities. The frequencies of voluntary withdrawals were 2 out of 29 in the former and 1 out of 36 in the latter. There was, however, a meaningful difference in these withdrawal rates over the subsequent rounds. 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 treatment with multiple withdrawal opportunities. Figure 1 shows percentages of strategies that were consistent with the “panic equilibrium” for each treatment. The differences in the slope

¹¹ Consider, for example, the scenario 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.

¹² All 8 withdrawals occurred in different groups.




Figure 1: Panic equilibrium strategies; Low risk case; 8 rounds with forced withdrawals.

of the trend lines for these two series is significant.¹³ However, we do not draw any conclusions from this observation alone. 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 before reaching conclusions on how the learning effect on voluntary withdrawal rates compared across treatments.

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 rounds 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.¹⁴

¹³ The 95% confidence interval around the slope estimates of the trendline are (0.0016,0.0274) and (0.0299,0.0678) for the treatments with a single and multiple withdrawal opportunities, respectively.

¹⁴ There are, of course, other ways that one could construct the history variable. We would like to emphasize that we

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 7.¹⁵ “Treat Dum” is the treatment dummy, which equals 0 for the treatment with a single withdrawal opportunity and 1 for the treatment with three withdrawal opportunities.

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 I: 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 treatment with a single withdrawal opportunity. 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 two treatments. Due to the numerous interactions of the variables, it is not immediately transparent from Table I 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 treatment with three withdrawal opportunities and is true in all but the first two rounds of the treatment with a single opportunity (in the first two rounds there is a slightly negative marginal effect.) There is also a much more pronounced

did not experiment with alternative specifications of this variable, and hence the reported significance levels are valid.

¹⁵ The values 1 through 7 for the Round variable correspond to rounds 2 through 8 with forced withdrawals in the experiment. The first round with forced withdrawals is not included since the history variable is not defined. It would not make sense to specify the Round variable from 2 to 8 in the Probit.

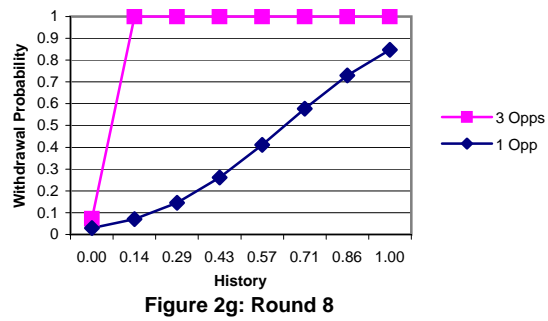


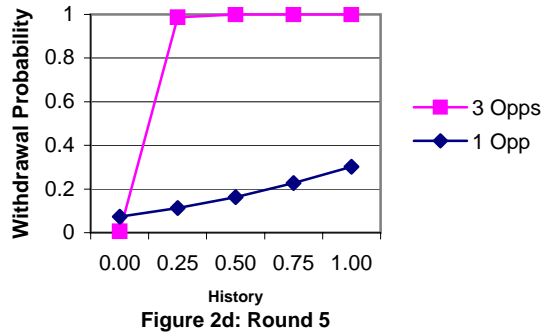
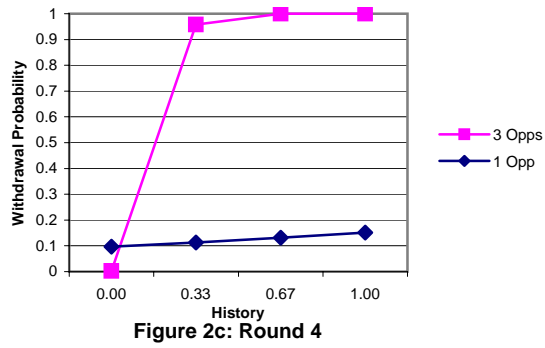
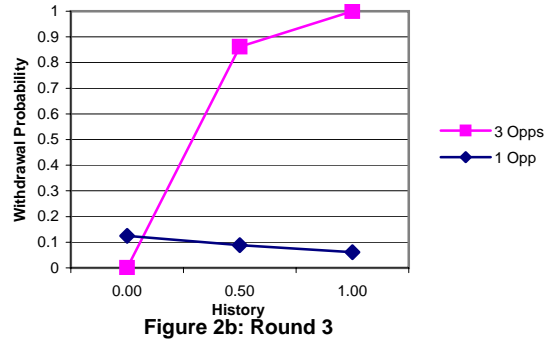
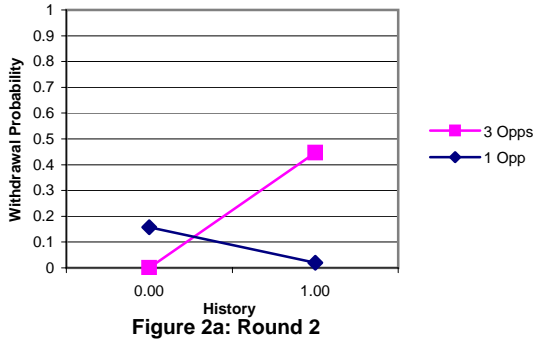
effect in the treatment with multiple withdrawal opportunities than in the treatment with a single withdrawal opportunity. For instance, evaluated at the mean of the independent variables, the estimated marginal effect of an increase in the history variable is .52 in the former treatment and .18 in the latter treatment. The joint test of the coefficients on $\text{Treat Dum} * \text{History}$ and $\text{Treat Dum} * \text{History} * \text{Round}$ 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 treatment with multiple withdrawal opportunities than in the treatment with a single withdrawal opportunity. 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 former treatment and 0.019 in the latter treatment. The joint test of the coefficients on $\text{Treat Dum} * \text{Round}$ and $\text{Treat Dum} * \text{History} * \text{Round}$ has a p value of 0.0262, so this difference is significant at the 5% level.

Figures 2a-g and 3a-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. The effects described above are apparent from the figures.¹⁶ In particular, figures 2c-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 treatment with multiple withdrawal opportunities than in the treatment with a single withdrawal opportunity. Figure 3a shows that holding history constant at 0, voluntary withdrawal probabilities are fairly constant over time, while figure 3c shows that holding history constant at 1, voluntary withdrawal probabilities increase over time. An intermediate case is shown in figure 3b. Note that in figure 3b 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 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.

¹⁶ Some 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.






Figure 3a: History = 0




Figure 3b: History = .5




Figure 3c: History = 1

3.3 Bank Run Analysis

A bank run is defined to have occurred if the bank liquidates all of its assets before the end of a round. No runs occurred without forced withdrawals. Bank runs occurred regularly with forced withdrawals, even though there was only 1/8 probability that forced withdrawal would cause the bank to fail. Figure 4 shows bank run frequencies (out of 8 groups) that occurred in each round of each treatment with forced withdrawals. It is apparent from the figure that frequencies of bank runs rose much more quickly in the treatment with 3 withdrawal opportunities than the treatment with 1 withdrawal opportunity, despite similar occurrences of bank runs in the early rounds. The explanation for this was revealed by the Probit analysis above. While early-round exposure to bank runs was similar across the two treatments, subjects reacted more strongly to exposure to bank runs in the treatment with multiple withdrawal periods (Recall that a 0.1 increase in the bank run history variable increased the probability of withdrawing by 5.2% in the treatment with multiple withdrawal opportunities and only 1.8% in the treatment with a single withdrawal opportunity.) Hence, bank runs had a “snowballing” effect on the withdrawal rates in the treatment with multiple withdrawal opportunities, but not in the treatment with a single withdrawal opportunity.




Figure 4: Bank run frequencies for each treatment, 8 rounds with forced withdrawals.

4 Evidence of Cutoff Rules with 3 Withdrawal Opportunities

When designing the experiment, we conjectured that subjects in the treatment with three withdrawal opportunities 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. The strategy corresponding to the payoff-dominant equilibrium under stage 1 payoffs is of this type with $Y = 2$. We now examine the extent to which observed subject behavior in stage 1 was consistent with two variable cutoff rules, $Y = 1$ and $Y = 2$, as well as with the panic strategy of withdrawing in opportunity I.¹⁷

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.

¹⁷ Note that an individual’s observed choices can be consistent with more than one rule, depending on the actual decisions she faced, hence reported frequencies within sessions will exceed the total number of subjects.

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 the first rounds with forced withdrawals. Table II 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 II: Cutoff rule frequencies, first round with forced withdrawals

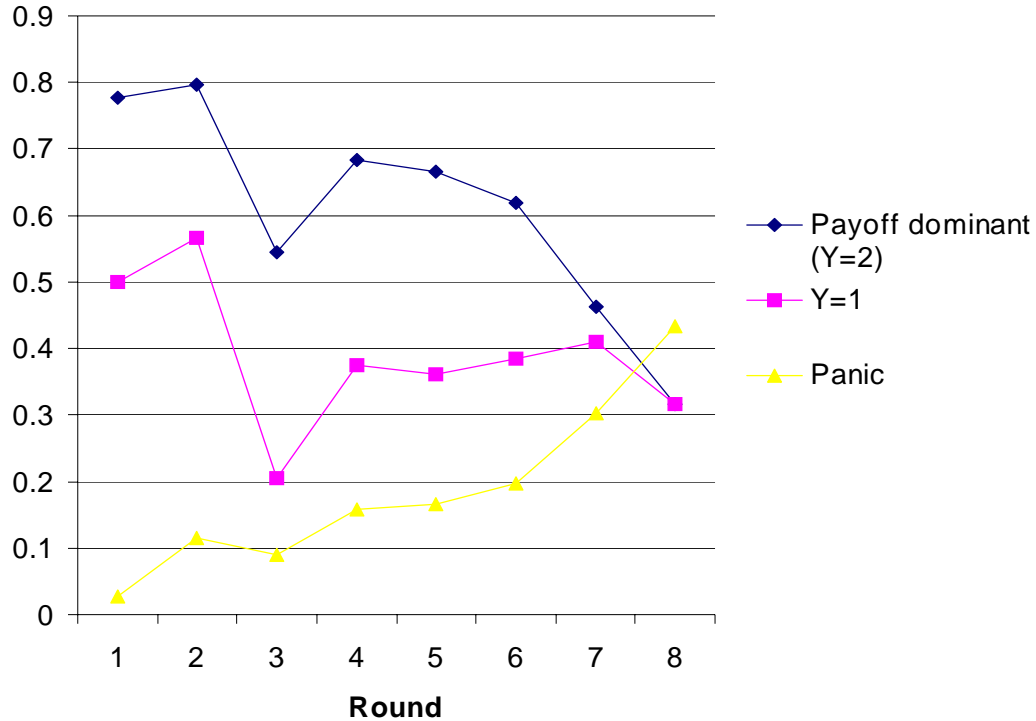


Figure 5: Cutoff rule frequencies, Rounds 1-8 with forced withdrawals.

This result is also evident in the data from the later rounds. Figure 5 shows the cutoff-rule frequencies for the combined sessions for the eight rounds with forced withdrawals.¹⁸ The figure reveals two things. First, it shows that the superiority of the $Y = 2$ cutoff rule is not limited to the first round. 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.

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 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%

¹⁸ The 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.

and 27.1% of the time, respectively. Hence, subjects who make it through the first withdrawal opportunity tended not to withdraw as predicted by the $Y = 1$ rule.

5 No Risk of Bank Failures from Forced Withdrawals

We now examine the second stage of each session, where the bank can withstand three withdrawals without defaulting on its obligations to the remaining depositors. In this case, an individual depositor will lose money by not withdrawing only if *all* of the other depositors withdraw. This change qualitatively affects the impact of forced withdrawals. In stage 1, there was a 1/8 probability that forced withdrawals alone would cause the bank to fail, leaving all remaining depositors with nothing. Here, there is *no* chance that the bank will fail if there are no voluntary withdrawals. In addition, the bank can now meet up to four early withdrawal requests with full payment to the withdrawers.

Overall withdrawal rates were substantially lower in the second stage than in the first, meaning that there is less variation in individual behavior both within and across treatments. Moreover, we ran fewer rounds in the second stage, in part because the lower withdrawal rate made each round more expensive to run. As a result, our results in this section are not as precise as those in Section 3. Nevertheless, as we describe below, the results are revealing and are entirely consistent with those for the first stage.

5.1 Equilibrium Analysis

In the treatment with a single withdrawal opportunity, there continue to be two Nash equilibria: a payoff dominant equilibrium in which no one voluntarily withdraws and a panic equilibrium in which everyone withdraws. In the treatment with three withdrawal opportunities, 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 $n_{II} \leq 3$. As in the case of stage 1 payoffs described in Section 3, adding forced withdrawals makes subjects resistant to one less withdrawal going into opportunity II, and corresponds to the $Y = 3$ cutoff rule.¹⁹ As in the case of no forced withdrawals, there is also a panic equilibrium in which everyone withdraws at the first opportunity.

¹⁹ The calculations needed to demonstrate that this strategy profile is indeed a subgame perfect equilibrium are presented in Appendix D.

5.2 Observed Behavior

A. No Forced Withdrawals. There were very few voluntary withdrawals in either treatment. The payoff dominant equilibrium was observed in 15 out of 16 games (there was one withdrawal in the opening round of session 2) in the treatment with a single withdrawal opportunity. It was observed in 12 out of 16 games in the treatment with 3 withdrawal opportunities, and 76 out of 80 subjects (95 percent) played strategies consistent with the payoff dominant equilibrium. As in the first stage, having multiple withdrawal opportunities has a positive effect on withdrawal rates in the absence of forced withdrawals, but the difference is not significant. None of the banks in which withdrawals occurred defaulted on the promise to pay \$1.50 to remaining depositors. Hence there were no observed instances of bank failures going into the rounds with forced withdrawals.

B. Forced Withdrawals.

Voluntary withdrawal rates were low in both treatments and there was no significant variation across rounds in subjects' tendency to play strategies consistent with one equilibrium or the other. The lower two lines in Figure 6 compare the voluntary withdrawal rate in sessions with a single withdrawal opportunity to that in the first opportunity of the sessions with three withdrawal opportunities. The rates in the figure represent the percentage of subjects whose play was consistent with the panic equilibrium in each of these treatments.²⁰ With multiple withdrawal opportunities, additional voluntary withdrawals occurred after the first opportunity. However, these withdrawals were, for the most part, consistent with the payoff-dominant equilibrium strategy, which is the cut-off rule $Y = 3$. The top line in figure 6 shows that the majority of play in the multiple-opportunity case was indeed consistent with the $Y = 3$ rule.

A single bank run occurred in the sessions with one withdrawal opportunity, while two bank runs occurred in the sessions with multiple withdrawal opportunities. Keeping in mind the small sample sizes, we interpret these results as follows. First, bank runs do occur in the presence of forced withdrawals, even when there is no chance that these withdrawals alone will exhaust the bank's funds. Second, the data are at least consistent with the results from Section 3, which showed bank runs to be more likely in the case of multiple withdrawal opportunities.

²⁰ The percentages are based on the number of subjects that were eligible to withdraw in each round (i.e., were not forced). The actual frequencies are 2 of 26, 4 of 30, 3 of 31 and 2 of 30 for the case of 1 opportunity and 2 of 38, 1 of 37, 2 of 39, and 3 of 34 for the case of 3 opportunities.

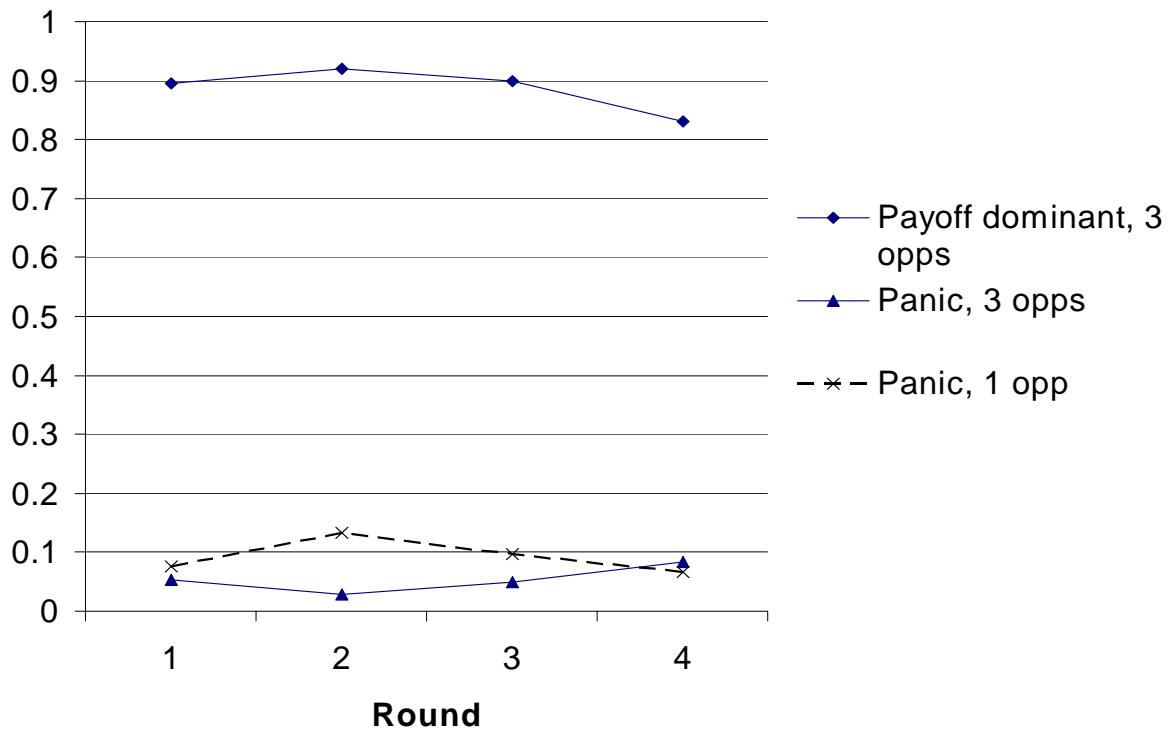


Figure 6: Equilibrium strategies; No risk case.

6 Conclusions

Our experiment and the analysis of the resulting data generate several key insights into breakdowns in coordination and the occurrence of bank runs. First, there is strong evidence that the ability of people to coordinate on the payoff-dominant equilibrium is sensitive to the presence of aggregate uncertainty about fundamental withdrawal demand, even when this uncertainty alone poses little or no threat to the solvency of the bank. The random forced withdrawals in our model mimic the type of uncertainty that is likely to be present under unfavorable macroeconomic conditions or in times of financial distress. Of course, some caution is required in extrapolating these results to consumer behavior during financial crises. The repeated play aspect of the experiment is somewhat artificial, and this is where the strongest treatment effect was observed. Nevertheless, the finding that individuals coordinate on the payoff-dominant equilibrium less frequently in the presence of aggregate uncertainty is interesting and potentially important.

Second, it matters whether subjects are given multiple opportunities to withdraw (with feedback) or a single opportunity. Initially, subjects were equally likely to withdraw at the first opportunity in both treatments, but exposure to bank runs had a greater (positive) effect on future

withdrawals in the treatment with multiple withdrawal opportunities. This result indicates that the standard approach of modelling bank runs using a one-shot, 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.

Finally, the analysis of withdrawal behavior in the treatment with multiple withdrawal opportunities generates insight into individual decision making in a dynamic environment with uncertainty. We tested various cutoff rules for characterizing individual decisions in this setting. 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 rule. This effect disappeared over time, as increased rates of immediate withdrawals eventually made the panic rule – withdraw immediately – a superior predictor. This does not suggest that subjects stopped updating, however. Rather, it suggests they stopped believing that others would play their part of the payoff-dominant equilibrium, as coordination broke down and bank runs became more prevalent.

References

- Boyd, J. H., S. Kwak, and B. D. Smith, “The Real Output Losses Associated with Modern Banking Crises,” *Journal of Money, Credit, and Banking* 37 (2006) 977-999.
- Chari, V.V. and R. Jagannathan, “Banking Panics, Information, and Rational Expectations Equilibrium,” *Journal of Finance* 43 (1988), 749-761.
- Cooper, R., 1999, *Coordination Games: Complementarities and Macroeconomics*, Cambridge University Press, Cambridge, UK.
- Cooper, R., D. V. DeJong, R. Forsythe, and T. W. Ross, “Selection Criteria in Coordination Games: Some Experimental Results,” *American Economic Review* 80 (1990), 218-233.
- Cooper, Russell and Thomas W. Ross, “Bank Runs: Liquidity Costs and Investment Distortions,” *Journal of Monetary Economics* 41 (1998), 27-38.
- Diamond, D. W. and P. H. Dybvig, “Bank Runs, Deposit Insurance, and Liquidity,” *Journal of Political Economy* 91 (1983), 401-419.

- Dutta, Prajit K., "Coordination Need Not be a Problem," mimeo, 2003.
- Ennis, H. M. and T. Keister, "Bank Runs and Investment Decisions Revisited," *Journal of Monetary Economics* 53 (2006) 217-232.
- Gale, D., "Dynamic Coordination Games," *Economic Theory* 5 (1995), 1-18.
- Gorton, G., "Banking Panics and Business Cycles," *Oxford Economic Papers* 40 (1988), 751-781.
- Green, E. J. and P. Lin, "Implementing efficient allocations in a model of financial intermediation," *Journal of Economic Theory* 109 (2003), 1-23.
- Jacklin, C.J. and S. Bhattacharya, "Distinguishing Panics and Information-Based Bank Runs: Welfare and Policy Implications," *Journal of Political Economy* 96 (1988), 568-592.
- Madies, P., "An Experimental Exploration of Self-Fulfilling Banking Panics: Their Occurrence, Persistence, and Prevention," *Journal of Business* 79 (2006), 1831-1866.
- Ochs, J., 1995, "Coordination Problems," in A.E. Roth and J.H. Kagel, eds.: *Handbook of Experimental Economics* (Princeton University Press).
- Peck, J. and K. Shell, "Equilibrium Bank Runs," *Journal of Political Economy* 111 (2003), 103-123.
- Saunders, A. and B. Wilson, "Contagious Bank Runs: Evidence from the 1929-1933 Period," *Journal of Financial Intermediation* 5 (1996), 409-423.
- Schotter, A. and T. Yorulmazer, "On the Severity of Bank Runs: An Experimental Study," mimeo, 2005.
- Van Huyck, J. B., R. C. Battalio, and R. O. Beil, "Tacit Coordination Games, Strategic Uncertainty, and Coordination failure," *American Economic Review* 80 (1990), 234-248.
- Wallace, N., "Another Attempt to Explain an Illiquid Banking System: The Diamond and Dybvig Model with Sequential Service Taken Seriously," *Federal Reserve Bank of Minneapolis Quarterly Review* 12 (1988), 3-16.
- Wallace, N., "A Banking Model in Which Partial Suspension is Best," *Federal Reserve Bank of Minneapolis Quarterly Review* 14 (1990), 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 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 dollar deposited in an experimental bank. You must decide whether to request to withdraw your \$1 or leave it deposited.

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	\$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. 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 20 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.

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 \$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 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.

Appendix C

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 Action
		# of forced withdrawals	Expected payoff	# of forced withdrawals	Expected payoff	# of forced withdrawals	Expected 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 Action
		# of forced withdrawals	Expected payoff	# of forced withdrawals	Expected payoff	# of forced withdrawals	Expected 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 Action
		# of forced withdrawals	Expected payoff	# of forced withdrawals	Expected payoff	# of forced withdrawals	Expected 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 D

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.