Global Games and Demand-Deposit Contracts: An Experimental Study of Bank Runs *

Alexander Klos[†] Finance Center Münster

Norbert Sträter[‡] Finance Center Münster

January 2008

Abstract

This paper experimentally investigates the global games approach applied to the bank run literature. In laboratory scenarios inspired by theoretical models, subjects receive a noisy private signal about the true fundamental state of the banking system. Subjects employ threshold strategies and we find low rates of deviation from threshold strategies. Increasing the repayment rate in the case of an early withdrawal leads to increased thresholds at an individual and aggregate level. However, compared to the theoretical predictions, the reaction to an increase of the repayment rate is less pronounced. Learning, in the sense that thresholds change over time, does not seem to be important.

JEL classification: G21, C92, C72

Keywords: Bank runs, global games, financial intermediation, experimental finance

---

*We would like to thank Timon Albers for his assistance in running the experiments. For helpful comments on earlier versions of this paper, we are indebted to John Hey, Susanne Homölle, Rosemarie Nagel, Andreas Pfingsten, Andreas Trauten, as well as to conference and seminar participants at Tübingen, Münster, the International ESA 2007 in Rome, the annual meeting of the German association for experimental economics (GfeW) 2007 in Goslar, the 2007 International LABSI conference on Political Economy and Public Choice in Siena, the IAREP 2007 Conference in Ljubljana, and the annual meeting of the German Finance Association 2007 in Dresden.

†Finance Center Münster, University of Münster, Universitätsstraße 14 – 16, 48143 Münster, Germany, alexander.klos@wiwi.uni-muenster.de

‡Finance Center Münster, Institut für Kreditwesen, University of Münster, Universitätsstraße 14 – 16, 48143 Münster, Germany, norbert.straeter@wiwi.uni-muenster.de

1 Introduction

The theory of global games (Carlsson and van Damme (1993)) has been useful in modeling various economic applications, such as speculative currency attacks (Morris and Shin (1998)) or the pricing of debt (Morris and Shin (2004)). Morris and Shin (2003) provide an overview of this fast growing literature.

Recently, the theory of global games has also been applied to the bank run literature (see Morris and Shin (2001), Dasgupta (2004), Rochet and Vives (2004), Goldstein and Pauzner (2005)). In these models, depositors receive a noisy private signal about the true fundamentals of the banking sector. This approach seems to be a promising one, especially because of two key features. Firstly, the global games approach can eliminate the multiplicity of equilibria present in classic panic-based bank run models (Diamond and Dybvig (1983)). Therefore, the probability of a bank run can be calculated, which is a useful number when thinking about policy implications. Secondly, the global games approach allows an integration of two related, but different views found in the literature on financial stability. The panic-based view stresses that bank runs are random events, unrelated to changes in the real economy. Panic-based bank runs might even occur when the economic environment is sufficiently strong. In contrast to this, the fundamental view assumes that bank runs are a natural outflow of the business cycle and only occur in connection with negative real shocks. In the global games approach, it is possible that fundamental runs occur, if the fundamentals of the economy – and therefore the signals which agents receive – are very bad. However, in other situations, fundamentals may be good enough so that the prevention of a bank run is desirable from the viewpoint of the depositors, but runs still occur due to strategic uncertainty about others' beliefs. Moderate signals about the fundamental state of the economy may lead to the belief that other agents withdraw with a high probability. In such a situation it is rational to withdraw also. In that sense, moderate expectations about the fundamental state of the financial system can trigger panic-based bank runs.

Our experiment is inspired by the theoretical work of Goldstein and Pauzner (2005) and Morris and Shin (2001). The theoretical work delivers a number of testable predictions. Two of them are the focus of our study:

1. Agents play threshold strategies, i. e. they withdraw their deposits if their signal is below a threshold θ^* and leave their money in the bank for a signal above θ^*. This behavior is partly caused by strategic uncertainty. Agents form beliefs about other players actions based on their private signal. If the fundamental state of the banking sector is bad, agents expect that other depositors withdraw their money early. Given this belief, it is optimal

to withdraw also. The opposite is true for good signals above the fundamental state. For some moderate signal, leaving the money in the bank and withdrawing the money early provide the same utility. This signal θ^* defines the threshold strategy.

2. Increased risk sharing makes banks more vulnerable to bank runs, i. e. a higher repayment rate in the case of early withdrawals increases the threshold θ^*. The intuition is that a higher repayment rate makes withdrawing early more favorable.

However, an empirical test of these predictions with field data is complicated. Empirical work on the stability of financial systems suffers due to the lack of sufficiently detailed data. While there has been a large increasing research on systemic banking crises in recent years, the rarity of banking crises makes causal inference difficult. For example, it is hard to distinguish between crises as the result of long-simmering problems or crises triggered by severe exogenous shocks (Demirgüç-Kunt and Detragiache (2005)). The problem is made worse when the global games approach is considered, as data on the individual depositors' behavior is lacking and every agent's opinion about the stability of the financial system becomes important. One possibility to overcome this problem is the use of experiments as a complementary methodology to field data research.[1] In an experiment, we can control the information of agents. This aspect is difficult to control with data from the field, but it is essential for putting the global games approach to bank runs to test.

Our paper is related to the recently started experimental investigation of bank runs. Schotter and Yorulmazer (2005) investigate in a policy oriented study the factors that affect the severity of fundamental bank runs. In their setup a bank run occurs by any means; they investigate how quickly the depositors withdraw their money. Their results indicate that partial deposit insurance and the existence of insiders may mitigate bank runs. Garratt and Keister (2006) study the conditions that lead to a self-fulfilling bank run and which factors affect its prevalence. One important result is that the pay-off dominant equilibrium is selected unless random forced withdrawals are introduced. Garratt and Keister (2006) investigate static and dynamic decision situations. Players have the option to withdraw their money at different points of time in the latter case, while the static game corresponds to the classic one-period simultaneous move scenario. The paper by Madiès (2006) tests the possibility and the degree of self-fulfilling banking panics. His results indicate that panics are persistent phenomena which can be curbed by a suspension of deposit availability and full deposit coverage.

[1] A recent example of related field evidence is a study by Chen, Goldstein, and Jiang (2007), who examine the role of strategic complementarities in the context of mutual fund outflows.

Our study is also related to the experiment of Heinemann, Nagel, and Ockenfels (2004), who study the theory of global games in the context of currency attacks.[2] They find evidence that players use threshold strategies, and that the observed behavior is close to the global games solution of the game. Duffy and Ochs (2007) extend the framework of Heinemann, Nagel, and Ockenfels (2004) to dynamic games. They force their subjects to use threshold strategies by requesting an explicit cut-off strategy. Duffy and Ochs (2007) find that the thresholds used are similar for static and dynamic games. Shurchkov (2007) investigates the role of a further signal that agents may receive in the second stage of a dynamic game. She finds that subjects learn endogenously (they know that a past attack was unsuccessful) and exogenously (they receive a second signal) as predicted by the theory.

This paper extends the existing literature by studying the predictions of the global games approach to bank run problems. In what follows, the contributions of the paper are discussed in light of the related literature and with respect to the two theoretical predictions cited above.

1. Heinemann, Nagel, and Ockenfels (2004) present evidence in favor of the use of threshold strategies in global games. However, it is not clear if this behavior extends to a global bank run game. From a theoretical perspective, the actions of the players are no longer global strategic complements, as the incentive to withdraw is greatest in the situation where the number of withdrawing agents is just sufficient to cause a bank failure, and not for the case that all players withdraw. Goldstein and Pauzner (2005) show theoretically that even in this case of one-sided strategic complementaries there is only one threshold equilibrium. However, this difference may lead to a situation, where subjects in an experiment no longer play threshold strategies with a single switching point.
   There is also reason to believe that the decision context changes behavior.[3] Numerous studies investigate social preferences experimentally (see e. g. Camerer (2003) for an overview). A threat to the stability of the financial system may cause a social preference for cooperation, which plays no role in the neutral frame chosen by Heinemann, Nagel, and Ockenfels (2004).[4] In agreement with this conjecture, Madiès (2006) reports that there are subjects in his experiments who cooperate in any case, although these subjects report in a questionnaire that they know that this behavior results frequently in a zero payoff.

[2]See e.g. also Cornand (2006) or Cabrales, Nagel, and Armenter (2007) for further experimental studies of global games.

[3]See Loewenstein (1999) for a general discussion of this argument.

[4]The wording of the instructions in the Heinemann, Nagel, and Ockenfels (2004) experiment avoids any reference to currency attacks and describes the problem solely in terms of payoffs (see Heinemann, Nagel, and Ockenfels (2002)).

A further aspect of our experiment is that there is no upper dominance region. An upper dominance region is a parameter region where, independent of the beliefs about other players' actions, no patient depositor withdraws early. Such a region is needed to show theoretically that the threshold equilibrium is the only equilibrium. Therefore, there are multiple equilibra in our experimental setup. One possibility to implement an upper dominance region is to add an external institution that buys the project for signals above some threshold and guarantees a high repayment. A second alternative is to assume that the project yields a sufficiently high repayment rate in period 1 for very good fundamentals. However, as already mentioned, one aim of this experiment is to build up a realistic decision context. We feel that both stories about the upper dominance region are somehow contrived and not a natural element of bank runs for the average subject. A second advantage of this design decision is that the study also serves as a test of the empirical importance of threshold strategies in the presence of multiple equilibra.

2. To the best of our knowledge, there is currently no direct experimental evidence on the role of the repayment rate on the probability of a bank run. However, theory suggests that the repayment rate is the crucial factor to trade off the benefits of risk sharing with the costs of a bank run.

We find that both predictions of the theory are supported by the data. We vary the signals subjects receive in experimental settings inspired by the theoretical models. Threshold strategies are estimated with a standard error model and a logit model. We find low rates of deviation from the threshold strategies. Increasing the repayment rate in the case of early withdrawal leads to increased thresholds at an individual and aggregate level, as predicted by the theory. However, theory also predicts a strong reaction to a change in the repayment rate, while in our experiment only small adjustments are observed. Implications of these results are discussed in the concluding section.

The remainder of the paper is structured as follows. The next section elaborates on our experimental design. Section 3 contains the experimental findings including the statistical analysis. Finally, section 4 draws conclusions and discusses directions for further research.

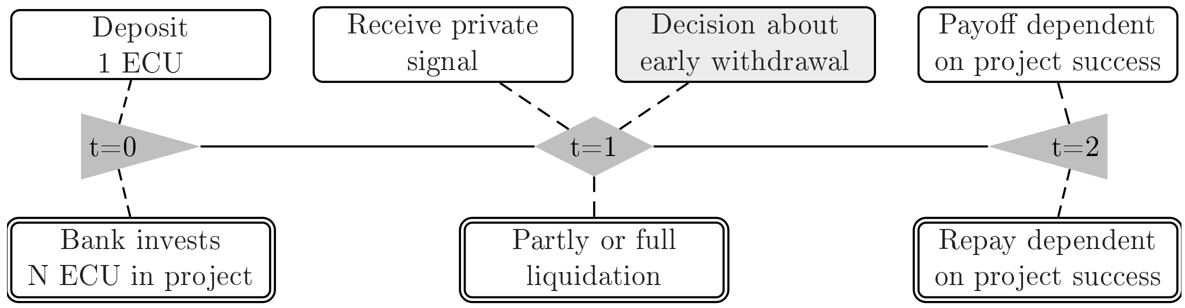

Figure 1: *Timing of events and actions:* The figure illustrates the timing of the different events and actions. The only decision is the binary decision whether or not to withdraw early at time t=1. Events and actions addressing subjects are single framed, events and actions addressing the bank are double framed.

2 Experimental Design

2.1 Overview and Hypotheses

Our main experimental framework is reflected in Figure 1 which shows the timing of events and actions in our experiment. There are three points in time. In t=0, all subjects deposit 1 ECU (Experimental Currency Unit) in a bank which represents the banking sector as a whole. Subjects who play the role of depositors in our experiment, can either be patient or impatient. Impatient savers only receive utility from consuming in t=1, while for patient savers consumption in t=1 and t=2 is of equal benefit. Subjects do not know a priori whether they are patient or impatient. The ex-ante probability for every subject to be impatient is given by λ. The bank invests the complete deposits collected from N depositors into a project which is expected to yield R ECU per invested ECU at the end of the investment horizon. The investment project can be liquidated before maturity and generates then for each unit of input one unit of output. The success of the investment project depends on the fundamentals of the banking sector. The exact link is established below. At the beginning of t=1, all subjects receive a private signal about the fundamentals. They are also informed if they are an impatient saver and therefore have to withdraw. Next, every patient depositor has to decide if or if not to withdraw early, i. e. to withdraw immediately at t=1. The bank offers for every ECU deposited at t=0 a payment of $r_1 > 1$ ECU in the case of an early withdrawal. These payments are financed by a partial liquidation of the investment project. The bank may run out of funds if too many subjects withdraw early. If the bank is still liquid in period 2, the remaining capital earns the project return R. The profit from the investment project is then divided equally among all subjects who have not withdrawn early.

One important question is how the fundamentals are defined. Two approaches can be distinguished. Morris and Shin (2001) model the project yield R as the uncertain fundamental[5], while Goldstein and Pauzner (2005) fix R and introduce a probability function $p(\theta) : [0,1] \mapsto [0,1]$ instead. The probability function assigns every fundamental θ to a probability $p(\theta)$. $p(\theta)$ is the probability that the investment project is successful and earns the fix rate R. We consider both approaches, and refer to the Morris/Shin approach as model A and to the Goldstein/Pauzner approach as model B. However, pre-tests have shown that subjects find it hard to think about fundamentals θ which determine a success probability $p(\theta)$ according to a complex transformation function. To simplify things, we only consider the case where $p(\theta) = \theta$ for model B, i. e. the uncertain fundamental is exactly the probability of a successful project without further transformation. As mentioned above, the uncertain fundamental θ in model A is the project yield R. In Table 1, we summarize both models in terms of their different payment consequences for subjects. A bank run occurs if $n \geq 1/r_1$, where n is the proportion of early withdrawals relative to possible withdrawals. Only a part of the early withdrawers, which are determined randomly,[6] receive r_1 in this case. Otherwise ($n < 1/r_1$), the bank is still liquid after period 1 and pays all late withdrawers $\frac{(1-nr_1)}{1-n}R$. Note that model B closely resembles the Goldstein/Pauzner model, while there are several differences between Morris and Shin (2001) and model A. Hence, model A should be regarded as a modification of the Goldstein/Pauzner approach with a different uncertain fundamental variable to check the robustness of this factor.

The advantage of the global games approach to the bank run problem is that the multiplicity of equilibria present in classic bank run models can be removed, and depositors play threshold strategies in the unique equilibrium, i. e. they withdraw early if and only if the signal about the fundamental is below some threshold θ^* and wait in other cases. The probability of a bank run can then be calculated on the basis of the ex-ante probability distribution of θ. A central insight of these models is that the degree of risk sharing the banking sector offers, increases the optimal threshold θ^*, and makes the banking sector more vulnerable to runs (see e. g. theorem 2 in Goldstein and Pauzner (2005)). One of the purposes of this paper is to test if the behavior derived by the theory with respect to the repayment rate shown up in a controlled lab experiment. We therefore varied the degree of risk sharing the bank offers (r_1).

[5]See also Dasgupta (2004) and Rochet and Vives (2004) for bank run models in which R depends on a random fundamental variable.

[6]Note that contrary to Diamond and Dybvig (1983) or Goldstein and Pauzner (2005) we do not implement a sequential service constraint for practicability reasons.

Ex-Post Payoffs				
	Model A		Model B	
Withdraw	$n < 1/r_1$	$n \geq 1/r_1$	$n < 1/r_1$	$n \geq 1/r_1$
Early	r_1	$\begin{cases} r_1 & : \frac{1}{nr_1} \\ 0 & : 1 - \frac{1}{nr_1} \end{cases}$	r_1	$\begin{cases} r_1 & : \frac{1}{nr_1} \\ 0 & : 1 - \frac{1}{nr_1} \end{cases}$
Late	$\frac{(1-nr_1)}{1-n}\theta$	0	$\begin{cases} \frac{(1-nr_1)}{1-n}R & : \theta \\ 0 & : 1-\theta \end{cases}$	0

Parameters		
	Model A	Model B
R	uniform distributed betw. 1.3 and 3.0	return always 6
p	success prob always 1	uniform distributed betw. 0 and 1
r_1	1.25 or 1.5	1.25 or 1.5
N	6	6
λ	1/6	1/6
ϵ	0.1	0.1

Table 1: *Payoffs and Parameters:* The table shows the ex-post payoffs for model A and B. Also shown are the parameters chosen for the experiment. θ denotes the fundamental variable. R stands for the project yield in model B, p for the success probability of the investment project, r_1 for the refund in the case of an early withdrawal, N is the number of subjects per group, λ depicts the ex-ante probability of being an impatient saver, and ϵ for the maximum deviation of a signal from the true state of the fundamental.

Reviewing the theory generates a number of hypotheses which follow directly from the intuitions mentioned in the introduction. We focus on two of them:

Hypothesis 1: Subjects use threshold strategies.

Hypothesis 2: Banks become more vulnerable to bank runs when they offer more risk sharing.

2.2 Parameterization

To the best of our knowledge, this experiment is the first that investigates the global games approach applied to bank runs. The parameterization of such experiments is always difficult, but the problem becomes even more critical in our case because of the risk of a ceiling effect. If we choose parameters that are too favorable for a bank run, subjects withdraw their money for nearly every signal. If the experimental situation is too favorable for leaving the money in the bank, we observe no bank runs. The effects of the variation of the risk sharing condition will then be unobservable due to withdrawal rates that are either zero or one hundred percent for every repayment rate r_1. However, for many of the experimental parameters, theory can be used as a guidance. For some other parameters, the designs of Heinemann, Nagel, and Ockenfels

(2004) and of the bank run experiments cited in the introduction are a good starting point. The results of pre-tests are also used to make parameterization easier.[7]

We decide to play our game with N=6 subjects like in Schotter and Yorulmazer (2005). Although bank run games outside the lab are played with a large number of agents, the payoff structure must be simple enough that a subject can quickly understand and use it during the short time period of the experiment. Six subjects seem to be a reasonable compromise.

The ex-ante probability of being an impatient saver is fixed at $\lambda = 1/6$ for both models. So, one subject in every group is forced to withdraw early, leaving us with five analyzable decisions per group.

To make things simple, we choose a uniform distribution for the uncertain fundamentals, i. e. the project yield R for model A and the success probability p for model B. As noted above, we restrict the probability function $p(\theta)$ to be the identity ($p(\theta) = \theta$) for model B. Therefore, a very high project yield R and a low repayment rate r_1 is needed to make the investment project desirable for a risk averse agent. A numerical evaluation of the Goldstein and Pauzner (2005) model suggests that $r_1^{low} = 1.05$, $r_1^{high} = 1.25$, and $R = 6$ may be a reasonable choice for agents with a coefficient of risk aversion slightly above one. However, pre-tests show that for $r_1^{low} = 1.05$ withdrawal rates are relatively low and we therefore risk a ceiling effect under that condition. We choose $r_1^{low} = 1.25$ and $r_1^{high} = 1.5$ for the experiment.

As few as possible parameters should be changed for the parameterization of model A. $r_1^{low} = 1.25$ and $r_1^{high} = 1.5$ are left unchanged. We choose a uniform distribution between $R^{low} = 1.3$ and $R^{high} = 3.0$ for the ex-ante distribution of the fundamental based on our pre-tests. A problem here is that the interval between R^{low} and R^{high} must be broad enough to account for individual heterogeneity. The noise variable ϵ is set to 0.1 in both models.

2.3 Theoretical Prediction of Thresholds

Consider the decision problem of an agent who is not forced to withdraw early in the first period. Depending on her signal s_i, the depositor can infer that the true fundamental value lies between $s_i - \epsilon$ and $s_i + \epsilon$. Because the a-priori probability distributions of the fundamental value and the error term are both uniformly distributed, the a-posteriori probability distribution of the true

[7] A further purpose of the pre-tests was to test the understandability of the instructions. The data of these tests is not included in the paper for several reasons. In the pre-tests, we used a fixed matching design, did not offer monetary incentives, played fewer rounds for more parameters, and had a different subject pool. Subjects were student assistants and the majority of them was familiar with the Diamond and Dybvig (1983) model. All these elements are different from the experiments reported in the paper (see Section 2.5 for further details).

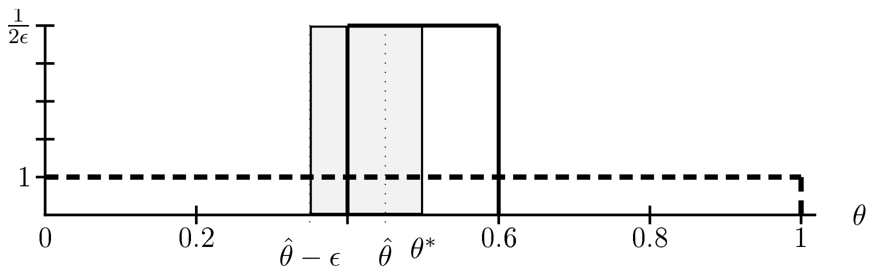

Figure 2: *A-priori and A-posteriori Probability Distributions:* The figure illustrates the probability distribution of the fundamentals before (dashed) and after (solid) the agent receives her signal $s_i = \theta^*$ for model B. θ^* equals 0.5 in this example. Also shown is the probability that one of the other patient agents receives a signal below θ^* if the true fundamental value equals $\hat{\theta}$ (lightgray area).

fundamental value given signal s_i is also uniformly distributed. The optimal action depends now on the belief about other players' actions. Assume that all agents have homogeneous preferences and follow the same threshold strategy which is defined by the threshold θ^*. Playing a threshold strategy implies that the agent is indifferent between withdrawing and waiting for a signal that equals exactly the optimal θ^*. In the case of such a signal $s_i = \theta^*$, the probability that one of the four other patient depositors receives a signal below θ^* if the true fundamental value equals $\hat{\theta}$ is given by $\frac{\theta^*-(\hat{\theta}-\epsilon)}{2\epsilon}$ (see Figure 2).

The expected utility of an early withdrawing is now the sum of the utilities given that none, one, two, ..., or all $(1-\lambda)N-1$ of the other patient depositors withdraw[8], or more formally

$$\mathrm{EU}_{\mathrm{Withdraw}}(\theta^*) = \sum_{n=0}^{(1-\lambda)N-1} \left[\frac{1}{2\epsilon} \int_{\theta^*-\epsilon}^{\theta^*+\epsilon} B\left(n, (1-\lambda)N-1, \frac{\theta^*-(\theta-\epsilon)}{2\epsilon}\right) \cdot EU(\mathrm{Payoff}_{\mathrm{Withdraw}}|n) \cdot d\theta \right]$$

where $B\left(n, (1-\lambda)N-1, \frac{\theta^*-(\theta-\epsilon)}{2\epsilon}\right)$ is the probability that exactly n of the other $(1-\lambda)N-1$ patient depositors withdraw. B denotes the binominal distribution. The expected utility given that n other patient depositors withdraw depends on the ability of the bank to payoff all early withdrawers. The number of withdrawers in that case is the sum of the n other patient depositors who withdraw, the λN people who have to withdraw early and our agent who decides to withdraw also:

$$EU(\mathrm{Payoff}_{\mathrm{Withdraw}}|n) = \begin{cases} u(r_1) & \text{if } \mathrm{RD}(N/r_1) \geq n+\lambda N+1 \\ \frac{\mathrm{RD}(N/r_1)}{n+\lambda N+1} \cdot u(r_1) + \left(1 - \frac{\mathrm{RD}(N/r_1)}{n+\lambda N+1}\right) \cdot u(0) & \text{if } \mathrm{RD}(N/r_1) < n+\lambda N+1 \end{cases}$$

RD(x) depicits the round down of x.

[8] We assume that λ is chosen in such a way that $(1-\lambda)N-1$ is an integer.

The expected utility of waiting can be calculated analogously and is given by

$$\mathrm{EU}_{\mathrm{Wait}}(\theta^*) = \sum_{n=0}^{(1-\lambda)N-1} \left[\frac{1}{2\epsilon} \int_{\theta^*-\epsilon}^{\theta^*+\epsilon} B\left(n, (1-\lambda)N - 1, \frac{\theta^* - (\theta - \epsilon)}{2\epsilon}\right) \cdot EU(\mathrm{Payoff}_{\mathrm{Wait}}|n) \cdot d\theta \right]$$

The expected utility of waiting given that n of the other four patient despositors withdraw is now different for model A and for model B. For model A, we have

$$EU(\mathrm{Payoff}^A_{\mathrm{Wait}}|n) = \begin{cases} \theta \cdot u\left(\frac{(1-n/N \cdot r_1)}{(1-n/N)} R\right) + (1-\theta) \cdot u(0) & \text{if } \mathrm{RD}(N/r_1) \geq n + \lambda N \\ u(0) & \text{if } \mathrm{RD}(N/r_1) < n + \lambda N \end{cases}$$

and for model B

$$EU(\mathrm{Payoff}^B_{\mathrm{Wait}}|n) = \begin{cases} u\left(\frac{(1-n/N \cdot r_1)}{(1-n/N)} \theta\right) & \text{if } \mathrm{RD}(N/r_1) \geq n + \lambda N \\ u(0) & \text{if } \mathrm{RD}(N/r_1) < n + \lambda N \end{cases}$$

The optimal threshold θ^* is defined by the solution of the equation $\mathrm{EU}_{\mathrm{Withdraw}}(\theta^*) = \mathrm{EU}_{\mathrm{Wait}}(\theta^*)$.

The lack of global strategic complementarities among agents' actions plays a role for $r_1 = 1.5$ in our parameterization. In that case, a partial run with 4 or 5 withdrawers leads to a payoff of zero if the agent waits. Accordingly, the agent who leaves his money at the bank expects zero payoffs in the cases that three or four out of the four other patient agents withdraw. The incentive to withdraw is greater if she expects that three other agents withdraw than if she expects that four other agents withdraw. Note that the lack of global strategic complementarities among agents' actions plays a smaller role for $r_1 = 1.25$, where a partial run leads only to a zero payoff if all other patient agents withdraw. Therefore, if the lack of global strategic complementarities is important for the use of threshold strategies, we should expect a higher degree of deviations from threshold strategies for $r_1 = 1.5$.

Note that the theoretical solution depends on the assumption of homogeneous preferences, risk aversion and the initial wealth of subjects. The following Table 2 gives the predicted threshold for our parameters and different degrees of risk aversion. We assume a power utility function $u(x) = \frac{x^{1-\alpha}-1}{1-\alpha}$ and an initial wealth of 1. Risk neutrality is implied by a risk aversion coefficient of $\alpha = 0$, and for $\alpha = 1$, log utility is considered. As shown in Table 2, optimal thresholds vary widely for different degrees of risk aversion. However, note that for model A and $r_1 = 1.5$, the theory predicts independently from the degree of risk aversion that all agents will withdraw for any signal.

Model A											
α	0	0.2	0.4	0.6	0.8	1.0	1.2	1.4	1.6	1.8	2.0
$r_1 = 1.25$	1.791	1.813	1.837	1.864	1.894	1.927	1.966	2.009	2.059	2.118	2.187
$r_1 = 1.5$	W	W	W	W	W	W	W	W	W	W	W
Model B											
α	0	0.2	0.4	0.6	0.8	1.0	1.2	1.4	1.6	1.8	2.0
$r_1 = 1.25$	0.278	0.320	0.365	0.412	0.461	0.510	0.559	0.608	0.655	0.700	0.743
$r_1 = 1.5$	0.480	0.539	0.601	0.664	0.727	0.791	0.853	0.917	W	W	W

Table 2: *Theoretical Predictions:* The table shows the theoretical predictions for both models and different risk aversion coefficients α. 'W' denotes a situation where withdrawing early is such attractive that withdrawing is favorable for any signal.

2.4 Issues in Experimental Methods

Bank run games are generally modeled as one-shot gambles. Experimental investigations of bank runs therefore face the same problem many game theoretic experiments have: The games have to be played repeatedly to make subjects in a short time period familiar with the decision situation. Some learning is explicitly allowed. On the other side, repeated play of one-shot gambles may change the nature of the game. Strategic considerations like reputation building may confound the results.

For our particular question here, we need to observe how subjects respond to different signals. We also want to allow learning in this relatively complex bank run game. However, strategic considerations due to repeated plays should be minimized. The first consideration for achieving these goals is a random matching design, i. e. every subject is randomly reassigned in a new group after each round and it is guaranteed that a subject will never play again the same opponents. This matching protocol, that has frequently been used in (2x2)-games, is not feasible for an experiment with larger groups. In an experiment with six persons per group played for 40 independent rounds, an unreasonable large number of participants per session is needed. Even at a large university it is not possible to recruit more than a few hundred voluntary subjects.

We therefore choose a random rematching design. Every session consists of 30 participants. At the beginning of each round, subjects are randomly and anonymously divided into five 6-person-groups for every decision. Within each group the participants interact with each other, but there is no interaction between the groups. At the end of each round the participants receive information on the outcome of their group. Information about the other groups is not given. After every round the groups are again randomly rematched. So, rematching with the same person(s) is possible, but the large ratio of session to group size should minimize repeated play effects. In the context of public good games, Botelho, Harrison, Pinto, and Rutström (2005) show that the behavior under a random rematching design approaches the behavior under a random matching design as the session-to-group-size ratio increases. Also note that forced withdrawals are present and they further complicate some kind of strategic signaling.

Such a random rematching design comes at the cost of only a few observations at the session level compared to a fixed matching design. However, we consider the reduction of strategic considerations as more important for our research question. Also, if the effects are strong enough, they should show up in a statistical analysis with only a few observations at the session level.

An alternative to avoid this problem has been proposed by Schotter and Yorulmazer (2005). They increased the payment of subjects in the first round by the factor twenty in a fixed matching design. However, while finding this approach interesting, it is not an alternative for our research question. For the estimation of threshold strategies it is necessary to observe how subjects respond to different signals. If their incentives are higher for the first signals, then thresholds can no longer be reliably estimated.

A further interesting methodological question arises from the design of Heinemann, Nagel, and Ockenfels (2004). In their experimental test of the global games approach to currency attacks, subjects receive ten independent signals in every round. For every signal participants have to decide whether they should attack or not. A more realistic setting would be to present subjects with only one signal in every round. If threshold strategies are still preferred, this result would support the robustness of their use. Additionally, an extension to dynamic global games is obvious if there is only one signal presented at once.

To sum up, theory guided considerations lead to our two different models A and B. Given that for our research questions the most important variable is the repayment rate r_1, we need at least two different values for this variable. In each session the participants are confronted with both values for r_1, but the ordering is varied. With regard to questions on experimental methodology, we decide for a random rematching design with a large session size. Additionally, we vary the number of signals subjects receive in one round to test if separate decisions would produce other results than sessions with joint decisions. Participants in sessions $A_{10}^{1.25/1.5}$ and $B_{10}^{1.25/1.5}$ play eight independent rounds with ten signals for both repayment rates, while participants in sessions $A_{1}^{1.25/1.5}$ and $B_{1}^{1.25/1.5}$ play forty independent rounds with one signal for both repayment rates. These considerations lead to (2x2x2)-design summarized in Table 3. We end up with eight sessions, each consisting of thirty subjects.

Model	Choices per Round	Starts with $r_1^{low} = 1.25$	Starts with $r_1^{high} = 1.5$
A	10	Session $A_{10}^{1.25}$	Session $A_{10}^{1.5}$
A	1	Session $A_{1}^{1.25}$	Session $A_{1}^{1.5}$
B	10	Session $B_{10}^{1.25}$	Session $B_{10}^{1.5}$
B	1	Session $B_{1}^{1.25}$	Session $B_{1}^{1.5}$

Table 3: *Experimental Design:* The table shows the general outline of our design. We test two models: each for different values of the repayment variable and each with a different amount of signals and choices per round.

2.5 Experiment Details

We conducted a computer-controlled experiment. Sessions were run at a PC pool at the University of Münster. No subject could participate in more than one session. In total, 240 undergraduate business and economics students participated, leading to a total of 24,000 analyzable decisions, where no early withdrawal was forced. The experiment was programmed and conducted with the software z-Tree (Fischbacher (2007)).

Participants were individually seated apart and they received written instructions including the relevant payoff tables. Pre-tests showed that payoff tables similar to the ones used in Madiès (2006) are easy to understand (see also the instructions in Appendix A.2). Personally, we acted as the supervisors. We read the instructions aloud and answered all questions. Sessions lasted about 90 minutes including the reading of the instructions. Throughout the experiment, the participants were not allowed to communicate with each other. Before the experiment started, we had a short quiz in order to test whether the subjects understood the payoff tables. At the end of each round participants received information on the outcome of their group. Information about the other groups was not provided.

At the end of the experiment all participants were asked to fill in a questionnaire with their comments regarding the experiment and also about their decision behavior. After finishing the questionnaire the subjects were paid out by converting their total amount of ECU into Euro. The average payment was 16.11 Euro, which is about equivalent to the payout they would receive for two hours of work as a student assistant at the Finance Center (library, computer support, ...).

3 Experimental Results

A strategy in our game maps a signal between 1.2 and 3.1 for model A (-0.1 and 1.1 for model B) into a binary decision whether to withdraw or leave the money in the bank. Denote s_i^d with the signal subject i receives in decision d. A threshold strategy $TS(t_i)$ with a threshold of t_i for subject i is therefore defined as

$$TS(t_i) = \begin{cases} \text{Withdrawal (1)} & \text{if } s_i^d < t_i \\ \text{No Withdrawal (0)} & \text{if } s_i^d \geq t_i \end{cases}$$

We use a simple error model as an intuitive starting point for the data analysis.[9] We assume an arbitrary threshold $\hat{t}_i \in [1.2; 3.1]$ for subject i in model A. The threshold defines the strategy $TS(\hat{t}_i)$. Then we classify every decision that is inconsistent with strategy $TS(\hat{t}_i)$, i. e. a withdrawal for a signal above $\hat{t}_i$ and leaving the money in the bank for a signal below $\hat{t}_i$, as an error. The threshold t_i^* that minimizes the percentage of errors defines the threshold strategy used by subject i. Typically, the optimal threshold t_i^* is not unique. In this case, we take the average of the supremum and the infimum of all thresholds that lead to the minimal percentage of erroneous decisions e_i^*. The percentage of erroneous decisions e_i^* for strategy $TS(t_i^*)$ is further used as a measure to judge how consistently subject i used the threshold strategy. The analysis can be done analogous for the aggregate data of a complete session. The left panel of Figure 3 illustrates the estimation procedure for session $A_{10}^{1.5}$ and $r_1 = 1.25$.

The simple error model can be easily modified. Denote with V_0 (V_1) the value associated with the alternative „leave the money until period two“ („withdraw early“). The value for both alternatives depends on the signal because the signal allows inference about other players information. Thus, the utility U_0 (U_1) can be written as $U_0 = V_0(s_j) + \epsilon_0$ ($U_1 = V_1(s_j) + \epsilon_1$) where ϵ_0 (ϵ_1) represents a random element in the evaluation of utility U_0 (U_1) for a given signal s_j. The randomness can either be due to unobserved personal traits, like risk aversion, or is simply due to errors in the value evaluation of an alternative (see e. g. Harless and Camerer (1994) and Hey and Orme (1994) on decision errors). The subjects choose to withdraw conditional on signal s_j if $U_1 > U_0$. The probability that U_1 is greater as U_0 is given by

[9] A similar model has been used by Palfrey and Prisbrey (1996) and Palfrey and Prisbrey (1997) in a public good context.

$$
\begin{aligned}
P(U_1 > U_0) &= P(V_1(s_j) + \epsilon_1 \geq V_0(s_j) + \epsilon_0) \\
&= P(\epsilon_0 - \epsilon_1 \leq V_1(s_j) - V_0(s_j)) \\
&= P(\epsilon \leq V).
\end{aligned}
$$

If the ϵ's, i. e. the differences in the value evaluations, are independently, identically distributed and they satisfy a logistic distribution, then we end up with a logistic regression model for the probability of a withdrawal (see e. g. Greene (2003)). We define the best fitting threshold t_i^* as the signal for which subject i withdraws with a probability of 0.5.

On an individual level, the ϵ's of the logit model represent the partly randomness of a subject's value evaluations. On the session level, the ϵ's also reflect the heterogeneity of preferences within the session. The right panel of Figure 3 illustrates the logit estimation procedure on the session level using the same data as the analysis of the simple error model shown in the left panel.[10]

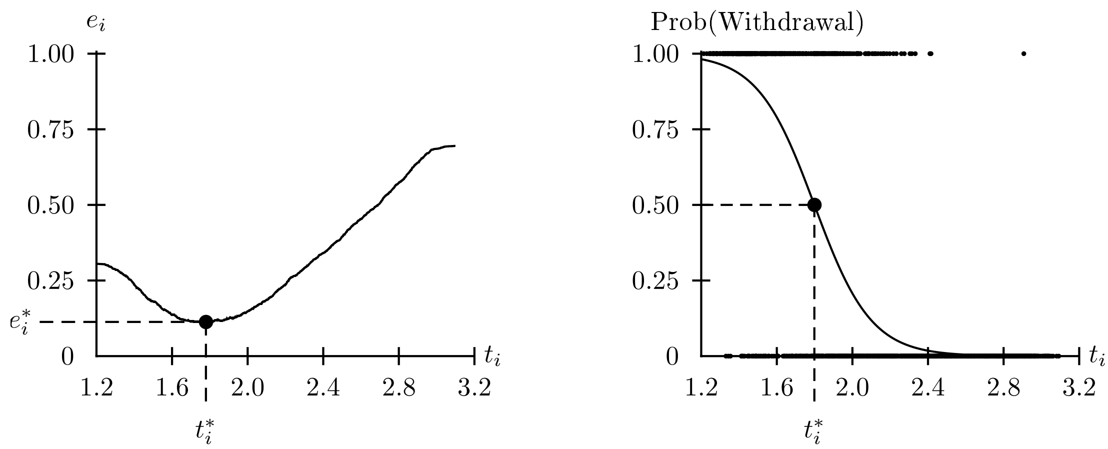

Figure 3: *Estimation of Threshold Strategies:* Both graphs show the estimation of the threshold that fits best the data from session $A_{10}^{1.5}$ for $r_1 = 1.25$. The estimation in the left panel is based on the simple error model, while the estimation in the right panel makes use of a logit regression. The estimated threshold t^* for session $A_{10}^{1.5}$ under the experimental condition $r_1 = 1.25$ is 1.7798 for the simple error model and 1.7992 for the logit model.

In what follows, both approaches, the simple error model and the logit model, will be used to estimate thresholds on the session and on the individual level. Decisions where experiment participants are subject to a forced withdrawal are excluded.

[10]An alternative assumption for the distribution of the ϵ's is a normal distribution. This assumption results in a probit model. We ran probit regressions, but they produce virtually the same results as the logit model and the results are not reported to save space.

3.1 The use of threshold strategies

The use of threshold strategies is analyzed by utilizing the simple error model. For every subject, we apply the methodology outlined in the previous section to estimate personal thresholds t_i^* and the percentage of decisions e_i^* that contradicts the personally optimal threshold strategy $TS(t_i^*)$ of subject i. This method allows us to estimate the degree of usage of threshold strategies directly on the individual basis. Panel I of Table 4 shows the results. The means and medians of error rates e_i^* are surprisingly low, indicating that the vast majority of subjects uses a threshold strategy right from the beginning. We consider hypothesis 1 to be confirmed.

Interestingly, the lowest error rates are observed on average for conditions $A_1^{1.25/1.5}$ and $B_1^{1.25/1.5}$, although threshold strategies are least obvious in these conditions. This result can be interpreted as further evidence for the intuitive appeal of threshold strategies in games with incomplete information about payoffs. One possible criticism of previous experiments (Heinemann, Nagel, and Ockenfels (2004)), which also potentially applies to our sessions $A_{10}^{1.25/1.5}$ and $B_{10}^{1.25/1.5}$, is that the joint decisions for ten given signals causally induces the use of threshold strategies. However, the (even lower) error rates for sessions $A_1^{1.25/1.5}$ and $B_1^{1.25/1.5}$ show that threshold strategies are also common when the experimental interface is less favorable for them.

The lack of strategic complementarities is not important with respect to the wide use of threshold strategies. Recall that the lack of strategic complementarities is more important for $r_1 = 1.5$ than for $r_1 = 1.25$. However, the error rates for both conditions do not reveal any systematic difference. The standard deviations, shown in panel III of Table 4, are also similar for both risk sharing conditions. The lack of strategic complementarities, a fact that considerably complicates theoretical analysis, is of less importance in the lab.

3.2 The role of offered risk sharing

3.2.1 Descriptive statistics

Panel II of Table 4 shows the estimated thresholds on the session level. For both estimation procedures, we obtain higher session thresholds – and therefore an increased probability of a bank run – for the experimental condition with a high degree of risk sharing. Panel III of the same table, which contains the average of estimated individual thresholds, confirms the result.

Table 4: *Experimental Results:* The table summarizes the main experimental results. Panel I shows the average error rates which are estimated with the simple error model outlined in section 3 and provide a measure for how consistent subjects use threshold strategies. The thresholds for the pooled session data – estimated using the simple error model and the logit model – are shown in panel II. Panel III reports the mean and the standard deviation of estimated individual thresholds. Also shown are the results of within Wilcoxon tests for the difference between conditions $r_1 = 1.25$ and $r_1 = 1.5$. Results of OLS regressions based on session thresholds are reported in panel IV. Panel V show the relative frequencies of bank runs. The last panel gives the details of sessionwise probit regression with the actions (0 or 1) as dependent and the signal, a dummy for the risk sharing condition, and the normalized round as explanatory variables. Shown are the marginal effects of every variable with all other variables at their mean. With respect to the dummy variable for for $r_1 = 1.25$, the coefficient gives the discrete change in the probability when the indicator is switched from 0 to 1.

Panel I: Mean of Individual Error Rates – Simple Error Model

	$A_{10}^{1.25}$		$A_{10}^{1.5}$		$A_{1}^{1.25}$		$A_{1}^{1.5}$		$B_{10}^{1.25}$		$B_{10}^{1.5}$		$B_{1}^{1.25}$		$B_{1}^{1.5}$	
	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}
Mean	5.21%	5.02%	4.00%	6.91%	5.36%	3.27%	2.93%	4.36%	5.71%	3.66%	4.46%	7.17%	5.23%	2.51%	3.17%	4.55%
Median	4.48%	4.77%	3.08%	5.20%	4.63%	3.03%	2.78%	3.13%	5.25%	3.10%	2.82%	4.49%	3.52%	1.39%	3.03%	3.13%

Panel II: Estimated Thresholds – Aggregated Data

	$A_{10}^{1.25}$		$A_{10}^{1.5}$		$A_{1}^{1.25}$		$A_{1}^{1.5}$		$B_{10}^{1.25}$		$B_{10}^{1.5}$		$B_{1}^{1.25}$		$B_{1}^{1.5}$	
	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}
Simple Error Model	1.80	2.30	1.78	2.33	1.81	2.03	1.56	2.20	0.44	0.47	0.44	0.50	0.44	0.47	0.44	0.48
Logistic Regressions	1.85	2.25	1.80	2.27	1.75	2.01	1.60	2.18	0.44	0.47	0.45	0.49	0.42	0.49	0.43	0.48

Panel III: Estimated Thresholds – Means and Standard Deviations of Individual Thresholds – 30 Paired Observations per Session

	$A_{10}^{1.25}$		$A_{10}^{1.5}$		$A_{1}^{1.25}$		$A_{1}^{1.5}$		$B_{10}^{1.25}$		$B_{10}^{1.5}$		$B_{1}^{1.25}$		$B_{1}^{1.5}$	
	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}
Simple Error Model	1.85	2.24	1.80	2.25	1.74	2.00	1.59	2.17	0.44	0.47	0.46	0.51	0.43	0.50	0.44	0.47
Standard Deviation	0.17	0.16	0.22	0.24	0.20	0.16	0.15	0.21	0.11	0.11	0.15	0.15	0.14	0.09	0.07	0.10
Wilcoxon Z	-4.720		-4.700		-4.762		-4.782		-2.252		-2.571		-2.376		-2.067	
p-value	(0.0000)***		(0.0000)***		(0.0000)***		(0.0000)***		(0.0243)*		(0.0101)*		(0.0175)*		(0.0387)*	
Logistic Regressions	1.85	2.24	1.80	2.26	1.72	2.00	1.61	2.18	0.44	0.47	0.44	0.49	0.42	0.49	0.43	0.47
Standard Deviation	0.15	0.16	0.22	0.24	0.19	0.16	0.14	0.18	0.11	0.10	0.10	0.09	0.11	0.08	0.08	0.09
Wilcoxon Z	-4.762		-4.762		-4.679		-4.782		-2.602		-2.273		-3.178		-3.013	
p-value	(0.0000)***		(0.0000)***		(0.0000)***		(0.0000)***		(0.0093)**		(0.0230)*		(0.0015)**		(0.0026)**	

Panel IV: Estimated Thresholds – Simple Error Model – Analysis on the Session Level – 8 Observations per Model

	Panel IVa: Model A (Adj. R^2: 0.8169)	Panel IVb: Model B (Adj. R^2: 0.7981)
Constant	2.1288*** (32.31, p<0.000)	0.48375*** (63.97, p<0.000)
Dummy for $r_1 = 1.25$	-0.4275** (-5.62, p=0.002)	-0.0475** (-5.44, p=0.003)
Dummy for Ordering	0.0975 (1.28, p=0.256)	-0.0025 (-0.29, p=0.786)

Panel V: Summary Statistics on Total Bank Runs

	$A_{10}^{1.25}$		$A_{10}^{1.5}$		$A_{1}^{1.25}$		$A_{1}^{1.5}$		$B_{10}^{1.25}$		$B_{10}^{1.5}$		$B_{1}^{1.25}$		$B_{1}^{1.5}$	
	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}	r_1^{low}	r_1^{high}
BR	27.00%	48.00%	29.25%	56.25%	19.00%	33.50%	16.00%	44.50%	37.00%	38.75%	44.00%	51.00%	33.00%	48.50%	46.00%	49.00%

Panel VI: Stability of Thresholds – Influence of the Round

	$A_{10}^{1.25}$	$A_{10}^{1.5}$	$A_{1}^{1.25}$	$A_{1}^{1.5}$	$B_{10}^{1.25}$	$B_{10}^{1.5}$	$B_{1}^{1.25}$	$B_{1}^{1.5}$
Signal	-1.455	-1.174	-0.925	-0.769	-2.542	-2.010	-2.992	-3.051
z-statistic (p-value)	-35.55 (0.000)***	-36.76 (0.000)***	-24.20 (0.000)***	-22.57 (0.000)***	-36.42 (0.000)***	-39.89 (0.000)***	-24.17 (0.000)***	-23.91 (0.000)***
Dummy for $r_1 = 1.25$	-0.519	-0.513	-0.221	-0.475	-0.107	-0.076	-0.199	-0.162
t-statistic (p-value)	-20.71 (0.000)***	-22.18 (0.000)***	-9.83 (0.000)***	-17.62 (0.000)***	-4.57 (0.000)***	-3.52 (0.000)***	-5.43 (0.000)***	-4.33 (0.000)***
Round	-0.002	0.005	0.001	-0.001	-0.013	0.001	0.001	-0.002
t-statistic (p-value)	-0.33 (0.738)	1.08 (0.278)	1.58 (0.115)	-1.07 (0.286)	-2.50 (0.012)*	0.30 (0.760)	0.56 (0.578)	-1.52 (0.128)

Note, that higher thresholds for higher repayment rates can be observed in between- and within-session comparisons. This result contradicts the view that within-designs point subjects to the differences the researcher investigates, and are therefore not as reliable as between-subjects experiments (see Camerer (2003), page 41 – 42, for a brief discussion).

To test the statistical significance of the results, we take two different approaches. In light of the discussion of different matching protocols in section 2, the most conservative way to analyze the data is to analyze estimated thresholds on the session level.

3.2.2 Evidence on the session level

We have two observations per session (TS), one for $r_1 = 1.25$ and one for $r_1 = 1.5$, resulting in a total of 16 observations, eight for model A and eight for model B. We run linear regressions separately for both models to infer the influence factors of session thresholds. Dummy variables for the degree of risk sharing (RD) and the ordering of risk sharing conditions (OD) are used as explanatory variables.

$$\mathrm{TS}_i = \mathrm{const} + \beta_1 \cdot \mathrm{RD} + \beta_2 \cdot \mathrm{OD} + \epsilon_i$$

It turns out that the dummy for a low degree of risk sharing is significantly negative for both models as predicted by the theory. The ordering of risk sharing conditions does not have a significant impact. Details of the regression results are given in panel IV of Table 4.

We also report the percentage of bank runs in the first row of panel V of Table 4. For model A, bank runs occur in 22.8% of all cases for r_1^{low}, but in 45.6% of all cases for r_1^{high}. More bank runs under a higher repayment rate r_1 are also observed for model B (r_1^{low}: 40%; r_1^{high}: 46.8%).

3.2.3 Evidence on the individual level

Stronger evidence for the key role of offered risk sharing can be provided by the analysis of thresholds estimated on the individual level. As already mentioned above, we estimate an individual threshold for decisions under the conditions $r_1 = 1.25$ and $r_1 = 1.5$. Therefore, we have two matched observations per subject. We apply a non-parametric Wilcoxon test, separately for every session, consequently based on thirty matched observations of individual thresholds. Panel III of Table 4 shows the results. For every session, the difference in estimated thresholds is significant. As the analysis on the session level has already suggested, differences in model B are smaller than in model A. Overall, the evidence suggest that hypothesis 2 is confirmed.

3.3 Stability of Thresholds

Learning in experimental games is a popular area of research (see e. g. Camerer (2003) for a review). In our experimental game, learning might have changed the thresholds over time. However, estimated thresholds are virtually unaffected by the ordering of risk sharing conditions (see the regression results in panel IV of Table 4). If thresholds move much over time, an order effect should have become evident.

We further estimated thresholds with the logit model separately for rounds 1 – 2, 3 – 4, 5 – 6, and 7 – 8 for the sessions $A_{10}^{1.25/1.5}$ and $B_{10}^{1.25/1.5}$, and rounds 1 – 10, 11 – 20, 21 – 30, and 31 – 40 for the sessions $A_{1}^{1.25/1.5}$ and $B_{1}^{1.25/1.5}$. Figure 4 reports on the results. Thresholds for subperiods are remarkably stable over time. No systematic change in played thresholds can be noticed with the naked eye.

To analyze the stability of thresholds numerically, we run sessionwise probit regressions (see panel VI of Table 4).[11] The binary decision of subjects is the dependent variable. Explanatory variables are the signals, a dummy for the risk sharing condition, and the round in the experiment. The round variable in the regression is a normalized variable in the sense that the variable is set to one when the risk sharing condition (r_1) has changed.

The sign of the round variable varies unsystematically from session to session and is only significant in one case. We conclude that thresholds on the session level are stable over time.

3.4 Panel Data Analysis

Observations from different subjects in experimental games are not independent. One way to deal with this problem is to analyze data only on the session level. So far, we have followed this basis. Individual behavior is only analyzed within a session and data is only pooled for across-session analysis in the form of session outcomes. A second way to address the issue is the use of panel econometrics that explicitly takes into account the dependence of observations. In this subsection, we use the second way.

[11] Garratt and Keister (2006) conduct a similar probit analysis. However, they focus on the interaction of treatment and time effects, while we prefer to analyze the time effects sessionwise. Note that the regression specification assumes that all subjects are homogeneous. In subsection 3.4 we relax this (and other) assumptions and present results of panel regressions.

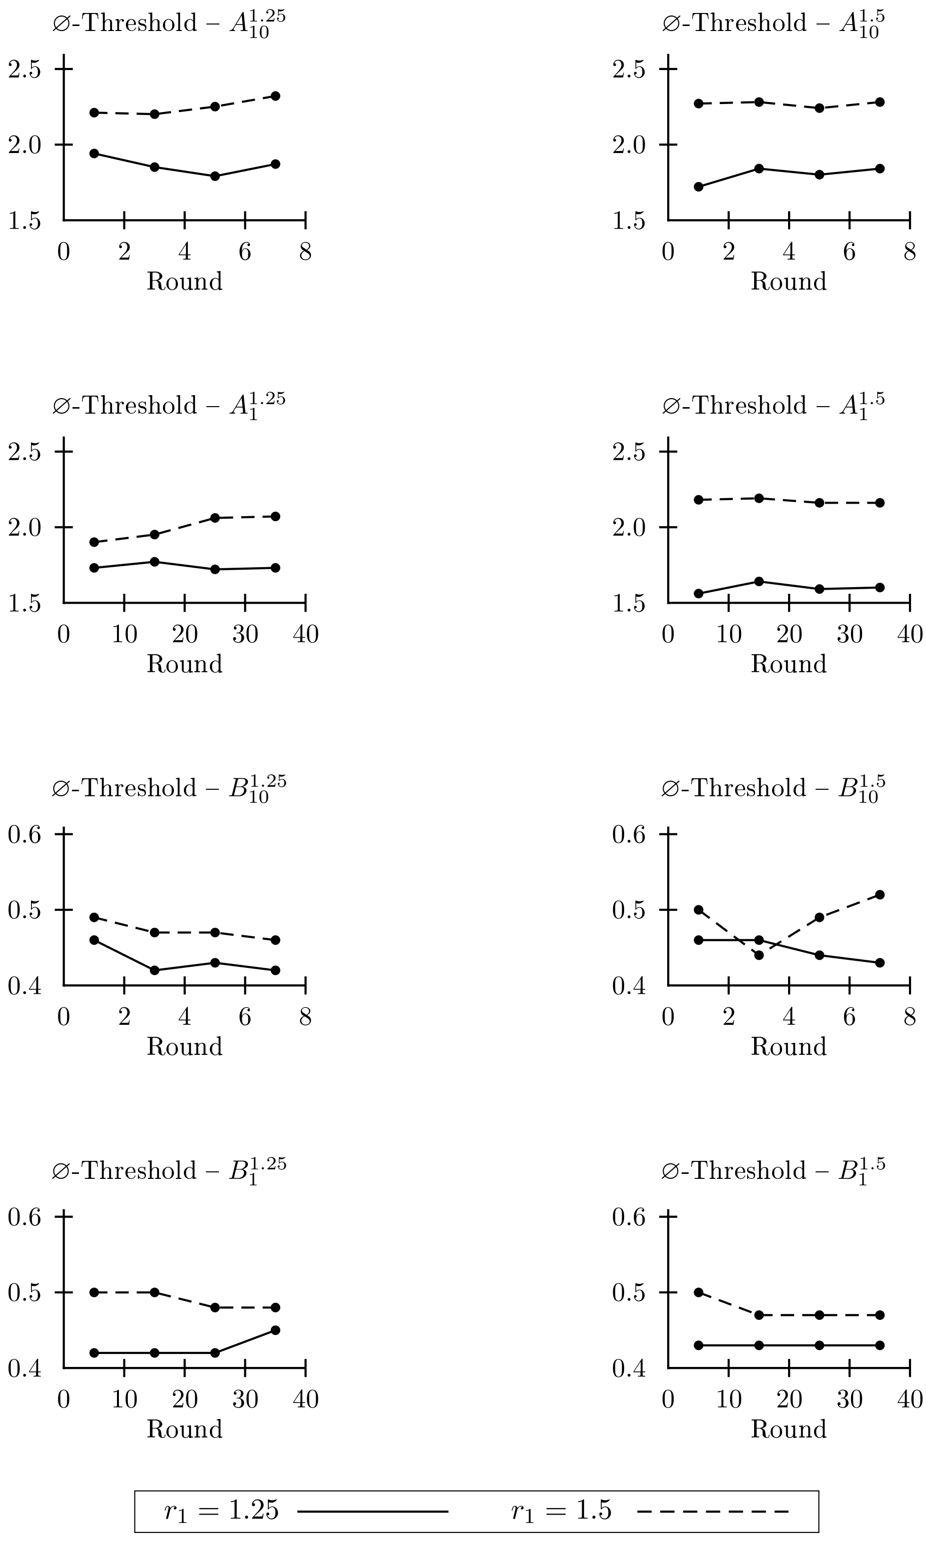

Figure 4: *Thresholds over Time:* The figure shows how the session thresholds evolve over time. For the sessions $A_{10}^{1.25}$, $A_{10}^{1.5}$, $B_{10}^{1.25}$, and $B_{10}^{1.5}$, we estimated thresholds by the use of the logit model for the rounds 1 – 2, 3 – 4, 5 – 6, and 7 – 8. For the sessions $A_{1}^{1.25}$, $A_{1}^{1.5}$, $B_{1}^{1.25}$, and $B_{1}^{1.5}$, we estimated thresholds for the rounds 1 – 10, 11 – 20, 21 – 30, and 31 – 40.

The panel logit analysis should principally serve as an explicit test of learning at an individual level, in contrast to the previous subsection, where we analyze learning at the session level. We therefore include the unnormalized round in which the decision was made as an explanatory variable into the regression. Given the differences in our models (A and B) and the differences in the total number of decisions, we can only pool the data from two different sessions together for one panel analysis. That means for model A, we can consider the data from sessions $A_{10}^{1.25}$ and $A_{10}^{1.5}$ together, but have to run separate regressions for sessions $A_1^{1.25}$ and $A_1^{1.5}$, because subjects have made fewer decisions in the latter case and the time structure of the panel is therefore not identical.

Specifically, we estimate the following one-way fixed effect model for four subsets of our data. x_{it} is the withdrawal decision of subject i in decision t and s_{it} is the corresponding signal. The variable RD equals 1 if $r_1 = 1.25$ and 0 otherwise. t depicts the round.

$$x_{it} = \text{const} + \beta_1 \cdot s_{it} + \beta_2 \cdot \text{RD} + \beta_3 \cdot t + u_{it}$$
$$\text{with } u_{it} = \alpha_i + e_{it}$$

Table 5 illustrates the results. In accordance with evidence that has already been given, the signal is the most important predictor of behavior. Also, the dummy for a low degree of risk sharing $r_1 = 1.25$ is negative and significant. There is some degree of learning present in model A, but not in model B. The probability of a withdrawal slightly declines over time. However, the decline is not strong enough to cause an economically significant effect on session thresholds as has been shown above.

Sessions $A_{10}^{1.25}$ and $A_{10}^{1.5}$ (n=8,000)				Sessions $A_1^{1.25}$ and $A_1^{1.5}$ (n=4,000)			
		z	p-value			z	p-value
Signal s_{it}	-8.004	-40.90	(0.000)***	Signal s_{it}	-8.857	-26.44	(0.000)***
r_1	-3.398	-29.25	(0.000)***	r_1	-3.711	-19.98	(0.000)***
Round	-0.020	-2.18	(0.029)*	Round	-0.025	-8.66	(0.000)***
Sessions $B_{10}^{1.25}$ and $B_{10}^{1.5}$ (n=8,000)				**Sessions $B_1^{1.25}$ and $B_1^{1.5}$ (n=4,000)**			
		z	p-value			z	p-value
Signal s_{it}	-12.781	-42.89	(0.000)***	Signal s_{it}	-17.978	-26.15	(0.000)***
r_1	-0.467	-5.72	(0.000)***	r_1	-1.023	-7.24	(0.000)***
Round	-0.006	-0.73	(0.473)	Round	0.002	0.71	(0.476)

Table 5: *Panel Regressions:* The table reports on the panel logit regressions. The dependent variable is the withdrawal decision x_{it} of subject i for decision t. The dummy for r_1 equals 1 if $r_1 = 1.25$ and 0 otherwise. Shown are the coefficients, the z-statistics and the corresponding p-values.

3.5 Comparing Estimated and Predicted Thresholds

If we consider the estimated thresholds on the aggregate level based on the simple error model, the mean of the estimated threshold for model A (model B) increases from 1.74 (0.44) for $r_1 = 1.25$ to 2.22 (0.48) for $r_1 = 1.5$. Using the logit estimation method and/or the individual thresholds alters these magnitudes only marginally. Comparing the estimates with the theoretical predictions in Table 2 reveals that – independent of risk aversion – a stronger effect of increased risk sharing would have been expected.

One possible explanation may be that subjects have been too optimistic for the high risk sharing condition ($r_1 = 1.5$), e. g. they assume that much more agents leave their money at the bank compared to the theoretical prediction. Given the uniform distribution of the fundamentals and the error term, each possible number of withdrawers between one and six should occur with equal probability of 1/6 in an $\pm\epsilon$-environment around the estimated session threshold. Note that the analysis in the previous subsections reveals that there is no economically important trend in the used threshold. Therefore, a look at the empirical distributions can show if there are indeed more scenarios that lead to a lower number of withdrawers around the optimal session threshold.[12] The tables in Appendix A.1 show that the relative frequency of the number of withdrawers in an interval between $\theta^* - \epsilon$ and $\theta^* + \epsilon$, with θ^* equal to the estimated session threshold, is not equal to 1/6 for all possible number of withdrawers. This result could be easily explained by differences in risk aversion, and therefore differences in individually used thresholds. However, there is no evidence that the number of withdrawers around the estimated session thresholds are systematically biased toward a smaller number of withdrawers. It is therefore highly unlikely that the rather small reaction to the repayment rate can be explained by systematically biased expectations.

4 Discussion

Our results suggest that the global games approach applied to bank runs leads to experimentally valid predictions with respect to the comparative statics. Both hypotheses are supported by the data: Our subjects use threshold strategies and increased risk sharing leads to increased thresholds at the individual and the aggregate level.

[12] An alternative would be to change the experimental design and simply ask for expectations. The disadvantage of such a direct elicitation is that the simple fact of asking such questions may alter behavior.

However, the rather small reaction of thresholds to changes in the repayment rates is at odds with the theory. This result is important with respect to the socially optimal repayment rate. One major advantage of the global games approach is the ability to trade off the benefits of high repayment rates (increased risk sharing) with its costs (increased probability of a bank run). Such a social optimization requires valid predictions about how people react to changes in the repayment rates. At least in our experimental setup, this assumption is questionable. As a result the optimal r_1 for the banking system as a whole is difficult to determinate. A simple calibration based on past experience would not work, because depositors do not react *quantitatively* as predicted by the theory to a change in the repayment rate.

Note that the optimal thresholds depend only on the local properties of the payoff functions near the optimal θ^*. This feature of the theoretical solution is a result of the distributional assumption of the error term. We conjecture that subjects do not fully realize this fact. Different assumptions about the distribution of the error term, e. g. a normal distribution, may therefore result in a better performance of the theoretical point predictions by reducing the gap between predicted session thresholds for $r_1 = 1.25$ and $r_1 = 1.5$. Further research may explore if our conjecture is correct.

Another curiosity is the fact that threshold strategies are played even if there is no interior solution according to the global games approach. This situation occurs for model A under the high risk sharing condition $r_1^{high} = 1.5$. In this case, withdrawing early leads to a higher utility for every possible signal. An explanation may be that subjects use threshold strategies as a heuristic or as a best response to the heuristics of others. Further experiments may use a set-up in which several heuristics should lead to qualitatively different responses of rational players. In such a framework, one may be able to assess the importance of heuristics and rational responses.

The question arises if the global games approach is successful in organizing the data, because it correctly predicts the comparative statics of results, or unsuccessful, because it fails to predict the magnitude of a repayment rate variation. Note that there are many reasons why the concrete thresholds in an experimental scenario may deviate from the theoretical prediction outlined above. The most obvious reason is unobservable individual differences, most notably risk aversion.[13] Also, the theoretical prediction depends on the assumption that all agents have ho-

[13]We have tried to assess the individual risk aversion using the method of Holt and Laury (2002). In an unreported analysis, we regress individual thresholds on risk aversion for every session. Risk aversion does not systematically explain variation in individual threshold strategies. Note that there are several possible reasons for these results. We do not pay subjects according to their performance for the risk aversion task and subjects may therefore not take the task serious enough. A competing explanation is that risk preferences depend on the decision context and are not a stable personal trait. Measures of risk aversion elicited from one task are therefore not useful for predicting behavior in other tasks.

mogeneous preferences. The crucial step of building expectations about other players actions is different if subjects assume heterogeneous preferences (see Hellwig (2002) for a discussion of this point). Further theoretical and experimental work may profitably analyze to what extend heterogeneity and beliefs about heterogeneity may affect economic outcomes in global games. It is of course possible that more sophisticated theoretical models can capture the sluggish reaction to a change in the repayment rate observed in this experiment.

Further research (theoretical, experimental and - if possible - using field data) beyond the early models is also needed. Possibly the most natural extension of our design is to include a stochastic number of forced withdrawals, that is a stochastic λ. Garratt and Keister (2006) report an increased number of panic-based runs if random forced withdrawals are introduced. A similar result is expected in the global games context, but this modification may also have a negative impact on the wide use of threshold strategies. A further addition is to consider the public information case and to contrast them with the global games approach. We refrained from this modification in our study, because an infinitely precise signal is in our view unrealistic. Presumably, depositors will not judge the fundamental state of the banking sector in the exactly same way. Instead, they will build their own expectations based on the public available information. Their opinions – the signals in our language – are therefore highly correlated, but not identical, and there is uncertainty about the opinions – the signals – of other agents. However, the public information case can be explored in further research.

With regard to policy implications a further extension of our experimental design is eligible. An analysis of factors which are also investigated in the other experimental studies of bank runs, like the introduction of deposit insurance, insiders or the suspension of convertibility, seems to be promising.

The growing theoretical literature on global games is increasingly interested in dynamic global games and their applications (see, e. g. Chamley (2003), Heidhues and Melissas (2006), Dasgupta (2007), Angeletos, Hellwig, and Pavan (2007)), i. e. theory extends beyond one-shot horizons and investigates dynamic decision situations. The experimental conditions A_1 and B_1, where threshold strategies are as common as in conditions with ten decisions presented at once, provide a natural starting point for the investigation of dynamic global games.[14] An interesting question for further research is the inspection of a dynamic global bank run game. Duffy and Ochs (2007) present evidence that it does not make a significant difference if a one-shot or a repeated game is considered in the currency attack context of Heinemann, Nagel, and Ockenfels (2004).

[14]See Shurchkov (2007); see also Duffy and Ochs (2007) for a different approach.

However, Schotter and Yorulmazer (2005) show that the dynamic perspective is crucial for bank run problems. Therefore, it is an interesting, yet unanswered question regarding what happens in a dynamic global bank run game. Furthermore, there is a need for further experimental evidence on dynamic global games with respect to other applications (see e. g. Brunnermeier and Morgan (2006)).

There are also methodological implications. First, Heinemann, Nagel, and Ockenfels (2004) find indirect evidence of the use of threshold strategies in global games. A more direct test would involve an analysis of actions for different signals from the same subject, as noted by Duffy and Ochs (2007). We conduct such a test by using the simple error model and find support for the use of threshold strategies on the individual level. Therefore, the conclusion of Heinemann, Nagel, and Ockenfels (2004) is robust with respect to our different analysis method.

An interesting methodological result is that learning is of minor importance, i. e. thresholds are stable over time. There are at least two possible reasons for this result. Firstly, as noted in the design section, we choose a large session size in a random rematching design, therefore minimizing any strategic considerations resulting from the repeated plays of the one-shot gamble. Secondly, by introducing forced withdrawals our set-up further complicated strategic signaling and coordination. People do not learn equilibrium in our bank run game. They play it right from the beginning.

A Appendix

A.1 Empirical Distributions Around the Optimal Threshold

The tables in the following subsections show the relative frequency of the number of withdrawers when the true fundamental state lies within the interval $[\theta^{agg} - 0.1; \theta^{agg} + 0.1]$. θ^{agg} denotes the estimated threshold on the session level (aggregated data). Z is the number of six-person-groups which received a signal between $\theta^{agg} - 0.1$ and $\theta^{agg} + 0.1$. SEM stands for the simple error model and LR for the logit regressions.

A.1.1 Model A

Error Model	θ^{agg}	1	2	3	4	5	6	Z
Sessions $A_{10}^{1.25} - r_1 = 1.25$								
SEM	1.80	9.09%	13.64%	22.73%	27.27%	20.45%	6.82%	44
LR	1.85	13.64%	25.00%	22.73%	18.18%	15.91%	4.55%	44
Sessions $A_{10}^{1.25} - r_1 = 1.5$								
SEM	2.30	2.50%	22.50%	40.00%	15.00%	12.50%	7.50%	40
LR	2.25	0.00%	13.89%	47.22%	13.89%	13.89%	11.11%	36
Sessions $A_{10}^{1.5} - r_1 = 1.25$								
SEM	1.78	2.56%	17.95%	35.90%	20.51%	12.82%	10.26%	39
LR	1.8	4.65%	18.60%	34.88%	20.93%	11.63%	9.30%	43
Sessions $A_{10}^{1.5} - r_1 = 1.5$								
SEM	2.33	13.73%	25.49%	23.53%	9.80%	19.61%	7.84%	51
LR	2.27	8.33%	20.83%	16.67%	16.67%	25.00%	12.50%	48
Sessions $A_{1}^{1.25} - r_1 = 1.25$								
SEM	1.81	20.00%	12.00%	40.00%	20.00%	4.00%	4.00%	25
LR	1.75	11.54%	7.69%	38.46%	23.08%	15.38%	3.85%	26
Sessions $A_{1}^{1.25} - r_1 = 1.5$								
SEM	2.03	4.35%	34.78%	21.74%	30.43%	4.35%	4.35%	23
LR	2.01	4.76%	33.33%	19.05%	33.33%	4.76%	4.76%	21
Sessions $A_{1}^{1.5} - r_1 = 1.25$								
SEM	1.56	0.00%	18.18%	18.18%	27.27%	27.27%	9.09%	22
LR	1.60	4.17%	20.83%	29.17%	12.50%	25.00%	8.33%	24
Sessions $A_{1}^{1.5} - r_1 = 1.5$								
SEM	2.20	19.05%	14.29%	28.57%	14.29%	19.05%	4.76%	21
LR	2.18	9.09%	9.09%	27.27%	27.27%	18.18%	9.09%	22

A.1.2 Model B

Error Model	θ^{agg}	1	2	3	4	5	6	Z
Sessions $B_{10}^{1.25}$ – $r_1 = 1.25$								
SEM	0.44	10.59%	15.29%	16.47%	24.71%	23.53%	9.41%	85
LR	0.44	10.59%	15.29%	16.47%	24.71%	23.53%	9.41%	85
Sessions $B_{10}^{1.25}$ – $r_1 = 1.5$								
SEM	0.47	5.63%	23.94%	21.13%	18.31%	22.54%	8.45%	71
LR	0.47	5.63%	23.94%	21.13%	18.31%	22.54%	8.45%	71
Sessions $B_{10}^{1.5}$ – $r_1 = 1.25$								
SEM	0.44	8.82%	17.65%	29.41%	30.88%	7.35%	5.88%	68
LR	0.45	8.96%	19.40%	31.34%	28.36%	5.97%	5.97%	67
Sessions $B_{10}^{1.5}$ – $r_1 = 1.5$								
SEM	0.50	6.94%	25.00%	18.06%	19.44%	22.22%	8.33%	72
LR	0.49	6.67%	25.33%	17.33%	20.00%	21.33%	9.33%	75
Sessions $B_{1}^{1.25}$ – $r_1 = 1.25$								
SEM	0.44	14.29%	17.14%	17.14%	37.14%	8.57%	5.71%	35
LR	0.42	10.81%	16.22%	24.32%	29.73%	8.11%	10.81%	37
Sessions $B_{1}^{1.25}$ – $r_1 = 1.5$								
SEM	0.47	8.11%	18.92%	16.22%	18.92%	16.22%	21.62%	37
LR	0.49	10.26%	23.08%	20.51%	17.95%	12.82%	15.38%	39
Sessions $B_{1}^{1.5}$ – $r_1 = 1.25$								
SEM	0.44	2.50%	25.00%	15.00%	17.50%	22.50%	17.50%	40
LR	0.43	2.38%	23.81%	14.29%	19.05%	23.81%	16.67%	42
Sessions $B_{1}^{1.5}$ – $r_1 = 1.5$								
SEM	0.48	9.76%	17.07%	17.07%	31.71%	12.20%	12.20%	41
LR	0.48	9.76%	17.07%	17.07%	31.71%	12.20%	12.20%	41

A.2 Instructions

Instructions given to the participants varied according to different sessions. Here, we include an English translation of the instructions for session $A_{10}^{1.25}$, i.e. the session started with $r_1 = 1.25$ and shifted to $r_1 = 1.5$. For the other sessions instructions were adapted accordingly.

Introduction

Thank you for participating in this economic experiment. The principle aim of this experiment is to analyze your decision-making behavior in groups. Your profits depend on your decisions but also on the decisions of the other participants in the experiment. You will make all your decisions at your workstation in our lab and to assist you in your decision process you will receive some printouts with helpful tables. This experiment is financed by our University.

General Information

You are one of 30 participants in this experiment and the rules are the same for all participants.

The experiment consists of 8 independent rounds, which are played by five 6-person-groups. At the beginning of each new round five new 6-person-groups will be formed out of the 30 participants. Within each group the participants interact with each other, but there is no interaction between the groups, each group is completely separate.

In order to conduct the experiment in a manageable time-frame we ask you to make your decisions in every round in the suggested processing time of 3 minutes. During each round a timer on your screen will inform you about the time remaining.

A round is finished after every participant has made their decisions. In a one minute information phase each participant is informed about her results in the previous round. Following this, the next round will start.

A quick summary of the rules

In every round each participant must make 10 decisions, all of which are similar in structure. Each participant will make a total of 80 decisions (8 rounds with 10 decisions).

Each of you has one ECU (Experimental Currency Unit), which you placed into a bank in time period 0. The bank invests your and your group members' deposits in an investment project. The investment project will last for 2 periods. After the two periods (t=2) the bank can realize project earnings R, which vary between 1.30 and 3.00 ECU. If the bank has to partly liquidate the project in period 1, it receives no project earnings for the liquidated part. It only receives the invested capital back.

However, the bank promises a fixed payment of 1.25 ECU to all withdrawing participants who have already requested their money in period 1. Therefore, it has to liquidate parts of the investment project to settle these claims. Participants who leave their money until period 2 will receive a payment which depends on the money remaining in the project and the project earnings R.

Your decision situation in detail

In this experiment we ask you to decide at which point of time you would like to withdraw money; that means if you would like to withdraw money in period one or two. The chronological structure is summarized in the following illustration.

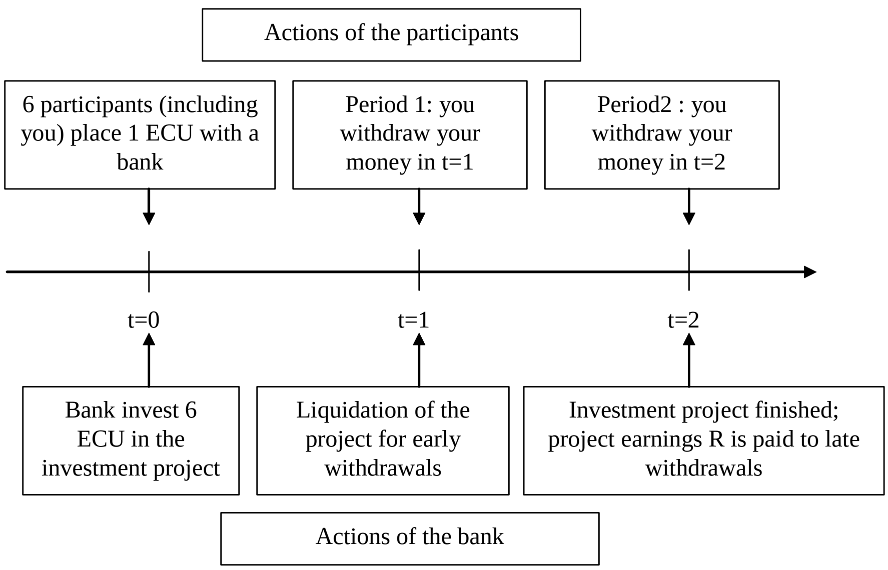

You might ask yourself how you can determine your payment if you leave your money at the bank until period 2. The following rule may help you: If some of the participants withdraw their money in the first period, the bank is forced to liquidate parts of the investment project to meet these expenses. If we call the number of participants who claim their money back in period one n, then n*1.25 ECU is claimed back. Due to this the bank retains (6-n*1.25) ECU of the originally invested amount (6 ECU), which realizes the project earnings of R ECU for each ECU remaining. In the second period the bank is liquidated and the generated proceeds are distributed among the remaining participants.

Example: If two of the six participants in a group claim a repayment in period one, each receives 1.25 ECU. Thus the bank has to liquidate 2.50 (2*1.25) in period one and 3.50 ECU (6-2.50) remains in the investment project. Assuming the project earnings amount to R=3, the bank possesses 10.50 ECU (3*3.50) at the end of the investment project. These 10.50 ECU are shared equally among the remaining four participants, so that each receives 2.63 ECU.

In the case that five or six participants demand their deposits in period one, the bank goes bankrupt. We call this situation a bank run. For example, if five participants claim their money back in period one, the bank has to pay out 6.25 ECU (5*1.25). Since this amount exceeds 6 ECU, the bank can no longer meet the withdrawal demands and the bank goes bankrupt. It can only pay out four participants with 1.25 ECU for which it expends 5 ECU (4*1.25). Among the five participants who wanted to withdraw their money four are selected by chance and they receive their payment of 1.25 ECU. The remaining amount of 1 ECU (6-4*1.25) is lost and the last (fifth) participant receives no payment. If all six participants claim their money, again only four participants receive 1.25 ECU. Two participants receive no payment. To sum up, if there are more than four participants who claim back their deposits, it is decided randomly which four participants receive a payment and which participants receive no payment.

The following illustration summarizes the payment scheme:

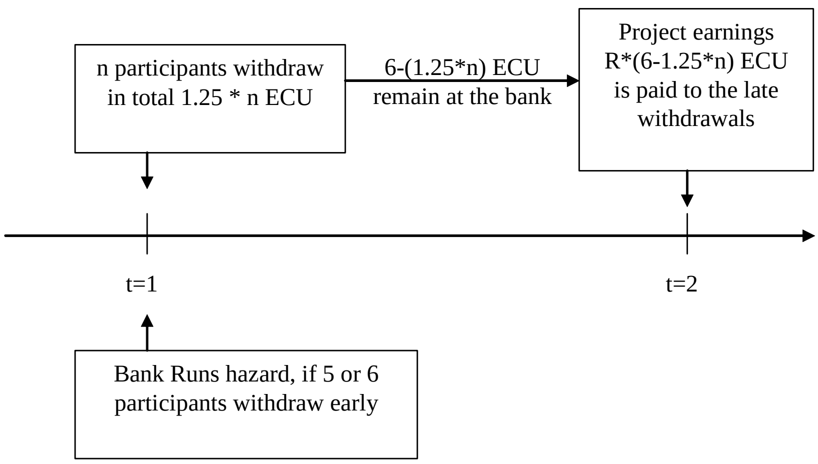

If you decide to withdraw your money in period one, your credit entry depends on how many participants also would like to withdraw their money in period one. If less than five participants decide to claim back their money in period one and you will definitely have 1.25 ECU transferred into your account.

If you decide to withdraw your money in period two, your credit entry depends on how many participants already claimed back their money in period 1 and how high the project earnings R are. The lower the number of participants who chose period one and the higher the project earnings R, the more credit you will receive in period two. If there are more than four participants claiming their money back in period one, you will definitely receive zero ECU in period two for sure.

Forced withdrawals

In each decision situation there is exactly one participant in each 6-person-group who has to claim back their money early. Thus this participant has to choose period one. The following message will appear on your screen: "As a result of an urgent need for money you are forced to withdraw your money. Hence you automatically have chosen period one." In each decision the odds are six to one (one in each 6-persons-group) that you will be forced to withdraw.

Calculation spreadsheets

Below is a table that will help you to figure out the payoffs associated with your withdrawal decision assuming that the project earnings are R=3. The chart gives you payoffs for all the possible numbers of requests. Please take a moment to make sure that you completely understand the spreadsheet.

Number of participants wishing to withdraw		Period 1			Period 2		
In period 1	In period 2	Number of participants who can be satisfied	Number of participants who are not satisfied	Individual amount of the withdrawal	Number of participants who can be satisfied	Number of participants who are not satisfied	Individual amount of the withdrawal
1	5	1	0	1.25	5	0	2.85
2	4	2	0	1.25	4	0	2.63
3	3	3	0	1.25	3	0	2.25
4	2	4	0	1.25	2	0	1.50
5	1	4	1	1.25 or 0	0	1	0
6	0	4	2	1.25 or 0	0	0	0

The handouts which you received as well as these instructions should aid your decision process. They contain tables which show your payments for different values of R.

The unknown project earnings R

As previously mentioned the success in period two also depends on the project earnings R. R can vary between 1.30 and 3.00. Every number between 1.30 and 3.00 has the same probability to be drawn, so that for instance the value R=1.78 will appear on average in every 171st decision situation (from 1.30 to 3.00 there are 171 different values for R). When you make your decision you do not know the amount of the project earnings R.

However, for each situation each participant will receive a hint for the unknown earnings R. This hint corresponds to a number which lies in a range between R -0.10 and R + 0.10. All numbers in this interval have the same probability to be drawn. Participants' hints are independent of the hints which the other participants receive. If for example in one decision situation, R would be 1.78, every participant will get a different signal which lies between 1.68 and 1.88.

Thus within a complete round you get 10 independent signals.

Please note

Every decision is completely independent of other decisions. That means for each decision you start again with one ECU in your bank account and have to decide to withdraw in period one or two. The money you earned at the end of each round belongs to you and is credited to your account. In each round five new 6-person-groups will be randomly formed out of the 30 participants of the experiment.

Information after each round

After each round every participant will be informed about the outcomes of the previous round. For all ten decisions you will be informed about:

(1) Your hint for R
(2) The real value of R
(3) Your chosen decision
(4) How many participants decided to withdraw early
(5) The payment you received
After a predetermined time limit the next round will start.

If you wish you can leave the information round before the end by pushing the grey OK button. If you do this, there is no chance to return to the previous information. A new round starts when you get the following message: „Please make your decision now".

Quiz and test round

To ensure that you understand these instructions and have a feeling for the decision situation, we will ask you to answer three easy questions correctly before you start with the real experiment. At the beginning of the experiment you will play one test round which should clarify the experiment to you.

Variation of the experiment

After you have made all your decisions (eight rounds each with ten decisions), we will repeat the experiment with a small variation. We will inform you about the variation after we have played the first eight rounds.

Questionnaire

After completing the whole experiment we will ask you to fill out a questionnaire that allows you to give us feedback concerning the experiment.

Payment

At the end of the experiment the ECUs that you earned will be converted into Euros and paid in cash. One ECU corresponds to 0.07 Euro or seven Eurocent.

Questions about the instructions

If you still have any questions about the instructions please ask now. Please do not ask afterwards and we ask you not to speak during the experiment. Thank you very much for following these instructions and for participating in this experiment.

References

Angeletos, G.-M., C. Hellwig, and A. Pavan (2007): "Dynamic Global Games of Regime Change: Learning, Multiplicity, and Timing of Attacks," *Econometrica*, 75, 711 – 756.

Botelho, A., G. W. Harrison, L. Pinto, and E. E. Rutström (2005): "Testing static game theory with dynamic experiments: A case study of public goods," Working Paper, University of Florida.

Brunnermeier, M. K., and J. Morgan (2006): "Clock Games: Theory and Experiments," Working Paper, Princeton University.

Cabrales, A., R. Nagel, and R. Armenter (2007): "Equilibrium Selection Through Incomplete Information in Coordination Games: An Experimental Study," *Experimental Economics*, 10, forthcoming.

Camerer, C. (2003): *Behavioral Game Theory: Experiments on Strategic Interaction*. Princeton: Princeton University Press.

Carlsson, H., and E. van Damme (1993): "Global Games and Equilibrium Selection," *Econometrica*, 61, 989 – 1018.

Chamley, C. (2003): "Dynamic Speculative Attacks," *American Economic Review*, 93, 603 – 621.

Chen, Q., I. Goldstein, and W. Jiang (2007): "Payoff Complementarities and Financial Fragility: Evidence from Mutual Fund Outflows," Working Paper, Duke University.

Cornand, C. (2006): "Speculative Attacks and Informational Structure: An Experimental Study," *Review of International Economics*, 14, 797 – 817.

Dasgupta, A. (2004): "Financial Contagion through Capital Connections: A Model of the Origin and Spread of Bank Panics," *Journal of the European Economic Association*, 2, 1049 – 1084.

——— (2007): "Coordination and Delay in Global Games," *Journal of Economic Theory*, 134, 195 – 225.

Demirgüç-Kunt, A., and E. Detragiache (2005): "Cross-Country Empirical Studies of Systemic Bank Distress: A Survey," IMF Working Paper, International Monetary Fund.

Diamond, D. W., and P. H. Dybvig (1983): "Bank Runs, Deposit Insurance, and Liquidity," *Journal of Political Economy*, 91, 401 – 419.

Duffy, J., and J. Ochs (2007): "Equilibrium Selection in Entry Games: An Experimental Study," Working Paper, University of Pittsburgh.

Fischbacher, U. (2007): "z-Tree: Zurich Toolbox for Readymade Economic Experiments," *Experimental Economics*, 10, 171 – 178.

Garratt, R., and T. Keister (2006): "Bank Runs as Coordination Failures: An Experimental Study," Working Paper, University of California.

Goldstein, I., and A. Pauzner (2005): "Demand-Deposit Contracts and the Probability of Bank Runs," *Journal of Finance*, 60, 1293 – 1327.

Greene, W. H. (2003): *Econometric Analysis*. Upper Saddle River: Prentice Hall.

Harless, D., and C. F. Camerer (1994): "The Predictive Utility of Generalized Expected Utility Theories," *Econometrica*, 62, 1251 – 1290.

Heidhues, P., and N. Melissas (2006): "Equilibria in a Dynamic Global Game: The Role of Cohort Effects," *Economic Theory*, 28, 531 – 557.

Heinemann, F., R. Nagel, and P. Ockenfels (2002): "Speculative Attacks and Financial Architecture: Experimental Analysis of Coordination Games with Public and Private Information," Working Paper, University of Frankfurt.

——— (2004): "The Theory of Global Games on Test: Experimental Analysis of Coordination Games with Public and Private Information," *Econometrica*, 72, 1583 – 1599.

Hellwig, C. (2002): "Imperfect Common Knowledge of Preferences in Global Coordination Games," Working Paper, UCLA.

Hey, J. D., and C. D. Orme (1994): "Investigating Generalizations of Expected Utility Theory Using Experimental Data," *Econometrica*, 62, 1291 – 1326.

Holt, C., and S. Laury (2002): "Risk Aversion and Incentive Effects," *American Economic*

Review, 92, 1644 – 1655.

Loewenstein, G. (1999): "Experimental Economics from the Vantage-Point of Behavioural Economics," *Economic Journal*, 109, 23 – 34.

Madiès, P. (2006): "An Experimental Exploration of Self-Fulfilling Banking Panics: Their Occurrence, Persistence and Prevention," *Journal of Business*, 79, 1831 – 1866.

Morris, S., and H. S. Shin (1998): "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks," *American Economic Review*, 88, 587 – 597.

——— (2001): "Rethinking Multiple Equilibria in Macroeconomic Modelling," *NBER Macroeconomics Annual*, 15, 139 – 161.

——— (2003): "Global Games: Theory and Applications," in *Advances in Economics and Econometrics (Proceedings of the Eighth World Congress of the Econometric Society)*, ed. by M. Dewatripont, L. P. Hansen, and S. J. Turnovsky. Cambridge: Cambridge University Press, pp. 56 – 114.

——— (2004): "Coordination Risk and the Price of Debt," *European Economic Review*, 48, 133 – 153.

Palfrey, T. R., and J. R. Prisbrey (1996): "Altruism, Reputation, and Noise in Linear Public Goods Experiments," *Journal of Public Economics*, 61, 409 – 427.

——— (1997): "Anomalous Behavior in Linear Public Goods Experiments: How Much and Why?," *American Economic Review*, 87, 829 – 846.

Rochet, J.-C., and X. Vives (2004): "Coordination Failures and the Lender of Last Resort: Was Bagehot Right After All?," *Journal of the European Economic Association*, 2, 1116 – 1147.

Schotter, A., and T. Yorulmazer (2005): "On the Severity of Bank Runs: An Experimental Study," Working Paper, New York University.

Shurchkov, O. I. (2007): "Coordination and Learning in Dynamic Global Games: Experimental Evidence," Working Paper, MIT.