TI 2023-030/I
Tinbergen Institute Discussion Paper

Bank choice, bank runs, and coordination in the presence of two banks

Jasmina Arifovic[1]
Johan de Jong[2]
Anita Kopányi-Peuker[3]

1 Simon Fraser University

2 Erasmus University Rotterdam and Tinbergen Institute

3 Radboud University Nijmegen, Institute for Management Research

Bank choice, bank runs, and coordination in the presence of two banks*

Jasmina Arifovic† Johan de Jong‡ Anita Kopányi-Peuker§

May 17, 2023

Abstract

We investigate learning in a repeated bank choice game, where agents first choose a bank to deposit in and then decide to withdraw that deposit or not. This game has a single Nash equilibrium in pure strategies, characterized by all agents depositing in the bank that offers the highest return, even though it may be more vulnerable to bankruptcy if some agents withdraw early. We use an individual evolutionary learning algorithm to model under which circumstances and with which beliefs agents can learn the Nash equilibrium in the repeated game and compare the results to an experiment. We find subjects coordinating on the Nash equilibrium under low and medium risk, but efficient coordination fails under high risk (irrespective of whether subjects have full or only partial information).

Keywords: Bank runs, pre-deposit game, individual evolutionary learning algorithm, partial information, experiment.

JEL classification: C63, C92, D90, G40.

*Sadly, Jasmina Arifovic passed away in January 2022. She actively participated in the project until December 2021. At that stage we had designed the experiment, collected the planned data, performed simulations, discussed data-analysis and started writing up the paper. All remaining errors are ours (de Jong and Kopányi-Peuker). We would like to thank William Ho, Jorge Martinez, and Jonathan Puigvert Angulo for research assistance. For comments and suggestions, we would like to thank the participants at 12th Workshop on Theoretical and Experimental Macroeconomics in Berlin, the Maastricht Behavioral and Experimental Economics Symposium, the Experimental Finance conference in Bonn, the Research in Behavioral Finance Conference in Amsterdam, and seminar participants at the Erasmus University Rotterdam, Radboud University Nijmegen and the University of Amsterdam. Financial support from Simon Fraser University and the University of Amsterdam is gratefully acknowledged. IRB approval was obtained from the Ethics Committee at the University of Amsterdam, which Johan de Jong was affiliated with when the project was started. Further declarations of interest: none.

†Simon Fraser University (in memoriam).

‡corresponding author. Erasmus School of Economics, Erasmus University Rotterdam, and Tinbergen Institute. Mailing address: Erasmus School of Economics, Erasmus University Rotterdam, PO Box 1738, 3000 DR Rotterdam, The Netherlands. Email: j.a.dejong@eur.ese.nl

§Radboud University Nijmegen, Institute for Management Research. Mailing address: Nijmegen School of Management, PO Box 9108, 6500HK Nijmegen, The Netherlands. Email: anita.kopanyi-peuker@ru.nl

1 Introduction

Since the seminal paper by Diamond and Dybvig (1983) it has become common to describe bank runs as one of the equilibrium outcomes of a coordination game. In these models patient depositors have to decide to withdraw their deposits (optimal in the bank run equilibrium in which all others also withdraw) or leave the money in (optimal when others of the same type also refrain from withdrawing, the 'good' equilibrium). The pre-deposit game, in which an agent chooses whether or not to deposit money in the bank, is often not explicitly modeled even though it can be an important element of decision making. One can argue whether deposit decisions should be included or not, but the consequences of including it can be large. It imposes strict limits on sunspot solutions (Peck and Shell, 2003) and in the most simple case, in which agents do not receive any signals between the deposit and withdrawal decisions, bank run equilibria could even disappear (de Jong, 2021). As most previous papers focus on withdrawal decisions, research on the pre-deposit game is still scarce.

In this paper we run simulations with an individual evolutionary learning (IEL) algorithm and we conduct laboratory experiments to study learning in the bank choice game, which includes deposit decisions. In the game we allow agents to repeatedly choose the bank they want to deposit in from two available banks, and subsequently decide whether they want to withdraw their deposits, or not. The banks differ in their riskiness and the interest rate they offer. One of the banks promises a higher interest rate (risky bank) but becomes insolvent already with a lower fraction of early withdrawals, while the other offers a lower interest rate, but is less vulnerable to early withdrawals (safe bank). Theoretically there is only one Nash equilibrium in pure strategies in the one-shot game with two different banks (and also in the finitely repeated game): all agents deposit in the risky bank and do not withdraw their money. As a result they all earn the maximum possible payoff in the game. Despite this, de Jong (2021) found that only about half of the subjects play according to this Nash equilibrium strategy in the one-shot game, whereas others mainly deposit in the safe bank. This far-from-equilibrium situation raises questions about how choices would evolve dynamically if agents are allowed to learn. The most important ones are whether those who initially deposit in the safe bank can learn to deposit in the risky bank to increase their earnings, whether risky banks can remain solvent over time, and how the dynamics depend on the riskiness of the bank and on the information agents receive about the history.

To answer these questions we consider two information environments and vary the riskiness of the banks. In the first information setting agents receive information about the number of depositors and withdrawals in both banks (full information) at the end of each round, while in the second they only learn the fraction of agents withdrawing in the bank they deposited in (partial information). Full information is more common in previous bank run experiments, as subjects exactly know with how many other individuals they form a bank, and they also receive information from which they can infer the number of withdrawals in the bank after each round.[1] Given that belief-updating is not straightforward when agents face partial information we decided to investigate the effect of information structure on the dynamics. Furthermore, we consider the partial information environment more realistic, thus more relevant to investigate.

In the second treatment-dimension we vary the vulnerability of the banks, i.e. what fraction of agents is needed so that the bank becomes insolvent. Even though the riskiness of the banks should not make a difference in theory, coordination on the equilibrium and learning might be more difficult in the presence of very risky banks. As soon as the level of coordination required for the bank to remain solvent approaches the level of the 'noise' in the decision making, the coordination cannot be sustained. Moreover, agents may anticipate this noise and choose a different strategy from the beginning. To investigate this question we applied three different bank combinations in our experiment (low risk, medium risk and high risk). IEL includes some randomness in the decision making through its experimentation component. Therefore, the simulations also predict differences in learning for these three different conditions.

Our experimental results reveal only small differences in behavior depending on the information structure, which leads us to conclude that participants can be quite sophisticated in their beliefs-updating. Furthermore, we only observe a consistent failure to coordinate on the Nash equilibrium in sessions with the most vulnerable banks. In the majority of other sessions the decisions converge to the Nash equilibrium rather quickly. This convergence seems faster and more complete when participants receive information about the outcomes in both banks. However, we do not have enough data to assess the significance of this difference. The IEL simulations can predict our experimental results well when we use an initialization that is close to the initial behavior in the experiment, and assume sophisticated beliefs-updating.

Our paper contributes to two main strands of the literature. First, we add to the bank run literature by investigating dynamics in a repeated bank run game with the deposit decision included. Closest to our paper is de Jong (2021) who investigates the one-shot bank choice game with 15 different bank combinations. In our experiment we conduct 3 bank combinations in total in a repeated interaction, and we also vary the information structure. Even though we are not the first to vary bank characteristics or information structures across treatments, or having multiple banks, to the best of our knowledge we are the first investigating these characteristics together with a choice in which bank to deposit

[1] See for example Arifovic et al. (2013), Arifovic et al. (2023), or for a review of the experimental bank run literature Kiss et al. (2022a).

the endowments. Previous literature shows that even though theoretically bank characteristics do not alter possible equilibria, subjects cannot always learn to coordinate on the 'good' equilibrium. In their study of the effect of bank characteristics on withdrawals, Arifovic et al. (2013) introduce the coordination parameter η as the fraction of depositors that need to forego withdrawing ('wait') to make waiting the payoff-maximizing strategy. They consistently find coordination on the good equilibrium for $\eta \leq 0.5$, coordination on the bank run equilibrium for $\eta \geq 0.8$, and mixed results in between. Arifovic and Jiang (2019) and Arifovic et al. (2023) find also similar results, showing that depending on the relative riskiness of the banks subjects can coordinate on the waiting, or can learn to follow a sunspot, or fail to do either of them, and coordinate on always withdrawing from the bank.

The information variation in our experiment is only possible in a multi-bank setup and has not been looked at before. There are experiments in which subjects might not have perfect information about others' withdrawal decision. However, these experiments are mainly conducted on sequential withdrawal decisions, and the information provision concerns observations *within a round*. For examples see Schotter and Yorulmazer (2009), Kiss et al. (2012), Kiss et al. (2014), Davis and Reilly (2016), and Kiss et al. (2022c). In the majority of cases the visibility of others' actions leads to fewer withdrawals.

Finally, as mentioned, we are not the first one to consider a multi-bank setting. Closest to our paper is Shakina (2019) who focusses on the redeposit decision instead of the pre-deposit game. She finds that a redeposit option leads to more withdrawals from banks that cannot receive new depositors and fewer withdrawals from banks that are on the receiving end. Several papers investigated how interbank networks and contagion might affect liquidity and survival of the banks. For examples see Brown et al. (2017), Duffy et al. (2019), and König-Kersting et al. (2022). Note that our paper differs from this strand of the literature, as we do not consider dependency between the fundamentals across banks, and we allow subjects to explicitly choose between the different bank to deposit their initial deposit.

Our bank choice game results theoretically in fewer bankruns. Previous literature investigated also different instruments to reduce the probability of bankruns. Such instruments include deposit insurances (Madies, 2006; Schotter and Yorulmazer, 2009) or the possibility to choose a priority account for the funds (Kiss et al., 2022b). Note however, that these papers differ from our paper, as we do not introduce an additional mechanism or institution, but the presence of two banks is already enough to theoretically eliminate bank runs. For an extensive review of the general experimental bank run literature, see Kiss et al. (2022a).

Second, this paper also contributes to further development of the IEL algorithm. Agents that use the IEL algorithm in their decision making maintain a collection of strategies of which they continuously evaluate the performance. Each period the collection is updated through experimentation and the replacement of poorly performing strategies by better performing ones. IEL has been successfully used to explain results in several experiments, including call markets experiments (Arifovic and Ledyard, 2007; Anufriev et al., 2022), public goods experiments (Arifovic and Ledyard, 2011), and bank run experiments (Arifovic, 2019). However, in these experiments subjects always received enough information to be able to calculate foregone payoffs and evaluate strategies. To use IEL with an environment in which agents only receive partial information (in our case about their own bank), one needs to extend the algorithm with beliefs about the missing pieces of information (here: about what happened in the other bank). We explore several possibilities to do this and find that a variant in which agents form very sophisticated beliefs is closest to the experimental results.

The rest of the paper is organized as follows. In Section 2 we describe the bank choice game. Section 3 presents the simulations in more details together with their results. Section 4 describes the experimental design. In Section 5 we discuss the experimental results and compare them to the simulations. Section 6 concludes.

2 The bank choice game

There are N players who play a repeated version of the bank choice game of de Jong (2021). Each round consist of 3 periods: period 0, 1 and 2. In period 0, players can decide between two banks in which they can deposit their money (the endowment cannot be received as 'cash'). The initial deposit is c. After everyone has chosen a bank, the players enter a typical bank run game in which they are asked whether they want to withdraw their money or not. Important is that the banks to choose from are different. All banks invest the full amount of the endowments entrusted to them, but the return R in period 2 and the liquidation value L of the investment in period 1 differ. A higher return comes at the cost of a lower liquidation value. Therefore, banks with a high R and low L are the more vulnerable or 'risky' banks, while banks with lower R and higher L are 'safer'. We assume that $R_r > R_s > 1 > L_s > L_r$ where r stands for the riskier bank, and s stands for safer bank.

After players make their decisions, earnings in each bank are determined by the payoff function:

$$\pi_{withdraw} = \min\left(\frac{L_k c}{f_k}, c\right), \tag{1}$$

$$\pi_{wait} = \max\left(\frac{\left(1 - \frac{f_k}{L_k}\right) R_k c}{1 - f_k}, 0\right). \tag{2}$$

Table 1: Bank types

	Type 0	Type 1	Type 2	Type 3	Type 4	Type 5
R	3.00	2.67	2.33	2.00	1.67	1.33
L	0.10	0.20	0.40	0.60	0.80	0.93
f^*	0.103	0.135	0.276	0.429	0.615	0.778

Notes: Bank types with the parameter values for return R and liquidation value L, and the fraction of withdrawals f^* at which which waiting and withdrawing yields the same payoff.

Here f_k is the fraction of depositors in bank k who choose to withdraw from the bank. When none of the depositors withdraw $\pi_{wait} > \pi_{withdraw}$ and when all depositors withdraw $\pi_{wait} < \pi_{withdraw}$. There is a single value of f_k for which waiting and withdrawing yields the same payoff, denoted by f_k^*. If the fraction of depositors withdrawing is less than f_k^*, 'wait' becomes payoff dominant. f_k^* is thus a measure of the vulnerability of the 'good' equilibrium in the bank run game. The higher f_k^* is, the less vulnerable the bank is.[2]

In this paper we will use the simplest possible version of the bank choice game as stage game, i.e. there are two banks to choose from and the players will not get any information in between the bank and withdrawal choices. This results in four pure strategies for the players: depositing in the risky bank and waiting, depositing in the safe bank and waiting, depositing in the risky bank and withdrawing, and depositing in the safe bank and withdrawing. This version of the bank choice game has only one Nash equilibrium in pure strategies, characterized by all players depositing in the risky bank and choosing to wait. This is also the single Pareto optimal outcome and it yields the maximum payoff to all participants. Hence the single pure-strategy Nash equilibrium of the repeated game is simply that all players coordinate on this stage game equilibrium in every round. Depending on the parameter choices, the stage game may also have equilibria in mixed strategies. However, in our analysis we focus exclusively on pure strategies and therefore these additional equilibria are mostly beyond the scope of this paper.

The combinations for L and R that we use throughout the paper are listed in Table 1, together with f^*, the fraction of withdrawals resulting in indifference between withdrawing and waiting. Types 1 to 5 were also used in de Jong (2021).[3] We add one extra type, type 0, which is more vulnerable than type 1.

[2]This f_k^* parameter is closely related to the coordination parameter η_k from Arifovic et al. (2013). In fact, $f_k^* = 1 - \eta_k$.

[3]In de Jong (2021) the type number coincided with the number of withdrawals that would drive a bank with 6 depositors to insolvency. The banks we study here will not always have 6 depositors, so the type number does not have any special significance beyond the fact that a higher number corresponds to a safer, less vulnerable bank.

3 Simulation design and results

3.1 Individual Evolutionary Learning

The advantage of IEL compared to other learning models is that it needs relatively few free parameters to describe learning dynamics in repeated games with more than two players relatively well (Arifovic and Ledyard, 2011). In IEL each agent starts round t with a private set of strategies that is inherited from the previous period. This set S^i_{t-1} has a fixed size J and each place j in the set can contain any strategy available to the agent. Strategies in the set are therefore typically not unique. At the end of the round one of the strategies in the set will be selected to be played, but before that the set can undergo changes due to two processes: experimentation and replication.

In the bank choice game we assume no feedback between the deposit and withdrawal decisions. Therefore, agents cannot condition their withdrawal decision on the number of depositors in their bank. As a consequence, strategies naturally contain only two elements: a bank choice and a withdrawal decision. Each element of each strategy in the set is (independently) subject to experimentation with probability ρ. When experimented upon, the element changes. How it changes, depends on the application. In our simulations we limit the strategies available to the agent to only the four pure strategies of the one-shot game and then experimentation simply leads to switching of the choice in the respective element, e.g., risky bank instead of safe bank or waiting instead of withdrawing.

After the experimentation process is complete, all strategies in the set are evaluated in terms of the payoffs that the agent believes they would have generated if they would have been selected in the previous round (calculation of foregone payoffs $\tilde{\pi}$). Agents thus consider a situation in which only their strategy in round $t-1$ is (possibly) different, but the actions of all other agents remain the same. The information that they receive about other agents' actions may or may not be sufficient to calculate the foregone payoffs with certainty. We assume that agents calculate the actual foregone payoffs when they have information about both the number of depositors and the number of withdrawals (full information). In other cases the foregone payoffs depend on the beliefs B^i_{t-1} of the individual agents.

The idea of replication is that strategies with a higher foregone payoff have a higher probability of becoming more common in the set, pushing out worse performing strategies. In IEL this is implemented in the form of a tournament. For each place m in the new, updated set S^i_t two strategies $s^i_{v,t-1}$ and $s^i_{w,t-1}$ are randomly selected and the one with the highest foregone payoff fills the spot:

$$s^i_{m,t} = \begin{Bmatrix} s^i_{v,t-1} \\ s^i_{w,t-1} \end{Bmatrix} \quad \text{if} \quad \begin{Bmatrix} \tilde{\pi}\left(s^i_{v,t-1}, B^i_{t-1}\right) \geq \tilde{\pi}\left(s^i_{w,t-1}, B^i_{t-1}\right) \\ \tilde{\pi}\left(s^i_{v,t-1}, B^i_{t-1}\right) < \tilde{\pi}\left(s^i_{w,t-1}, B^i_{t-1}\right) \end{Bmatrix}. \tag{3}$$

The final step is to select the strategy to be played in round t. This is determined by a random draw from S_t^i with weights $\omega_{m,t}^i$ equal to the foregone payoffs of the strategies:

$$\omega_{m,t}^i = \frac{\tilde{\pi}\left(s_{m,t-1}^i, B_{t-1}^i\right)}{\sum_{k=1}^{J} \tilde{\pi}\left(s_{k,t-1}^i, B_{t-1}^i\right)}. \tag{4}$$

IEL as outlined above requires assumptions about the set size J, the initialization (the initial strategy set and the actions in round 1, and under some circumstances initial beliefs), and the experimentation rate ρ. Furthermore, we need to specify how beliefs are formed in case the agents are not informed about the number of depositors in each bank and the number of withdrawals. For the set size we choose $J = 100$, in line with the number used in another recent IEL paper Arifovic (2019). As a robustness check we performed simulations with other values for J as well. The differences are minimal. Together with probabilistic choice, the experimentation rate is the main source of noise in the agents' decision making. We therefore use several values to understand the impact of this parameter on the decision dynamics and in particular on convergence. As initialization, rather than choosing an initial strategy set and actions, we choose the probabilities with which the agents select a bank and withdraw in the first round. These probabilities are then used to construct the initial strategy set.[4]

3.2 Beliefs under partial information

In our simulations we consider both cases in which agents know of each bank how many agents deposited there and how many withdrew in the previous round (full information) and cases in which agents only learn about the fraction of withdrawals in the bank in which they had deposited (partial information). In the latter cases the foregone payoffs are mostly determined by beliefs. The choice of what these beliefs should be, is far from trivial. One option is to assume that agents are *naive* in their beliefs and ignore their own influence on the outcome in their own bank. They will thus calculate the foregone payoffs of the alternative action (withdrawing instead of waiting or waiting instead of withdrawing) in their last chosen bank as if that would not change the overall withdrawal fraction in the bank. Beliefs still have to be formed for the other bank, in which the agent did not deposit and about which she does not receive any information in that round. Here, agents could treat depositing in the bank as *neutral*, i.e. yielding the initial endowment, or

[4] The procedure is to randomly determine each strategy in the set independently based on the probabilities we chose for bank choice and withdrawal (the set can only contain pure strategies). This means that the initial strategy sets are actually random variables, which are independently determined for each agent. The chosen strategy in round 1 is one of the strategies in this set (random draw with equal weights).

they could use the last observed withdrawal fraction in that bank and calculate foregone payoffs using that (*last-known*).[5]

Another possibility is to assume that agents are *sophisticated* in their beliefs about the number of depositors and withdrawers in both banks. Sophisticated agents will never hold beliefs that contradict available information. They know that there are I other agents and that the number of depositors can therefore never exceed I. The fraction they observe for their last chosen bank further restricts the number of depositors that are possible. For example, if the fraction of withdrawals is $\frac{1}{3}$, the number of depositors in that bank needs to be a multiple of 3. We next assume that sophisticated agents' beliefs about the number of depositors in the bank chosen in round $t-1$ are as close as possible to the number of depositors they believed the bank had in round $t-2$. This means that if they believed that the number of depositors was 2 in round $t-2$ and they observe a withdrawal fraction of $\frac{1}{3}$ for round $t-1$, they will believe that there are 3 depositors in that bank in period $t-1$ and that one of these depositors withdrew. Beliefs about the number of depositors in the last chosen bank immediately fixes the beliefs about the number of depositors in the other bank as well, but not necessarily the number of withdrawals there. Here we assume that agents will use the fraction of deposits they believed to be withdrawn in round $t-2$, apply this fraction to the new number of depositors believed to be there in round $t-1$ and round the result to the nearest integer.[6] So in total we consider 3 cases in our simulations: naive neutral, naive last-known, and sophisticated.

3.3 Simulation results

Figure 1 condenses the most important results of our simulations, showing information about the average strategies in rounds 26 to 50 of the repeated bank choice game. The experimentation rate ρ was set equal to 0.05, within the range of values used in previous studies. In Appendix B we provide the results for other experimentation rates ($\rho = 0.01$ and $\rho = 0.10$). These are qualitatively similar to the results shown here for $\rho = 0.05$. All data shown is from an average of 100 simulation runs with 12 agents.

There are 8 tables, each showing results for all different combinations of bank types 0 to 5 (details about these types are given in Table 1). The tables on the left represent simulations in which agents start by fully randomizing, with equal probabilities, their bank choice ($f_b = 0.5$) and withdrawal choices from the safe and risky banks ($f_s = 0.5$ and $f_r = 0.5$). It is a default option often chosen to initialize the IEL algorithm. However, the single-shot experiments (de Jong, 2021) already revealed that the number of withdrawals

[5] In case an agent never deposited in a certain bank, their belief about the withdrawal fraction is equal to the value with which a single agent withdraws from that bank in round 1. This is an initialization value and all agents are assumed to know it.

[6] In case two integers are equally close, the agent randomly chooses one of them.

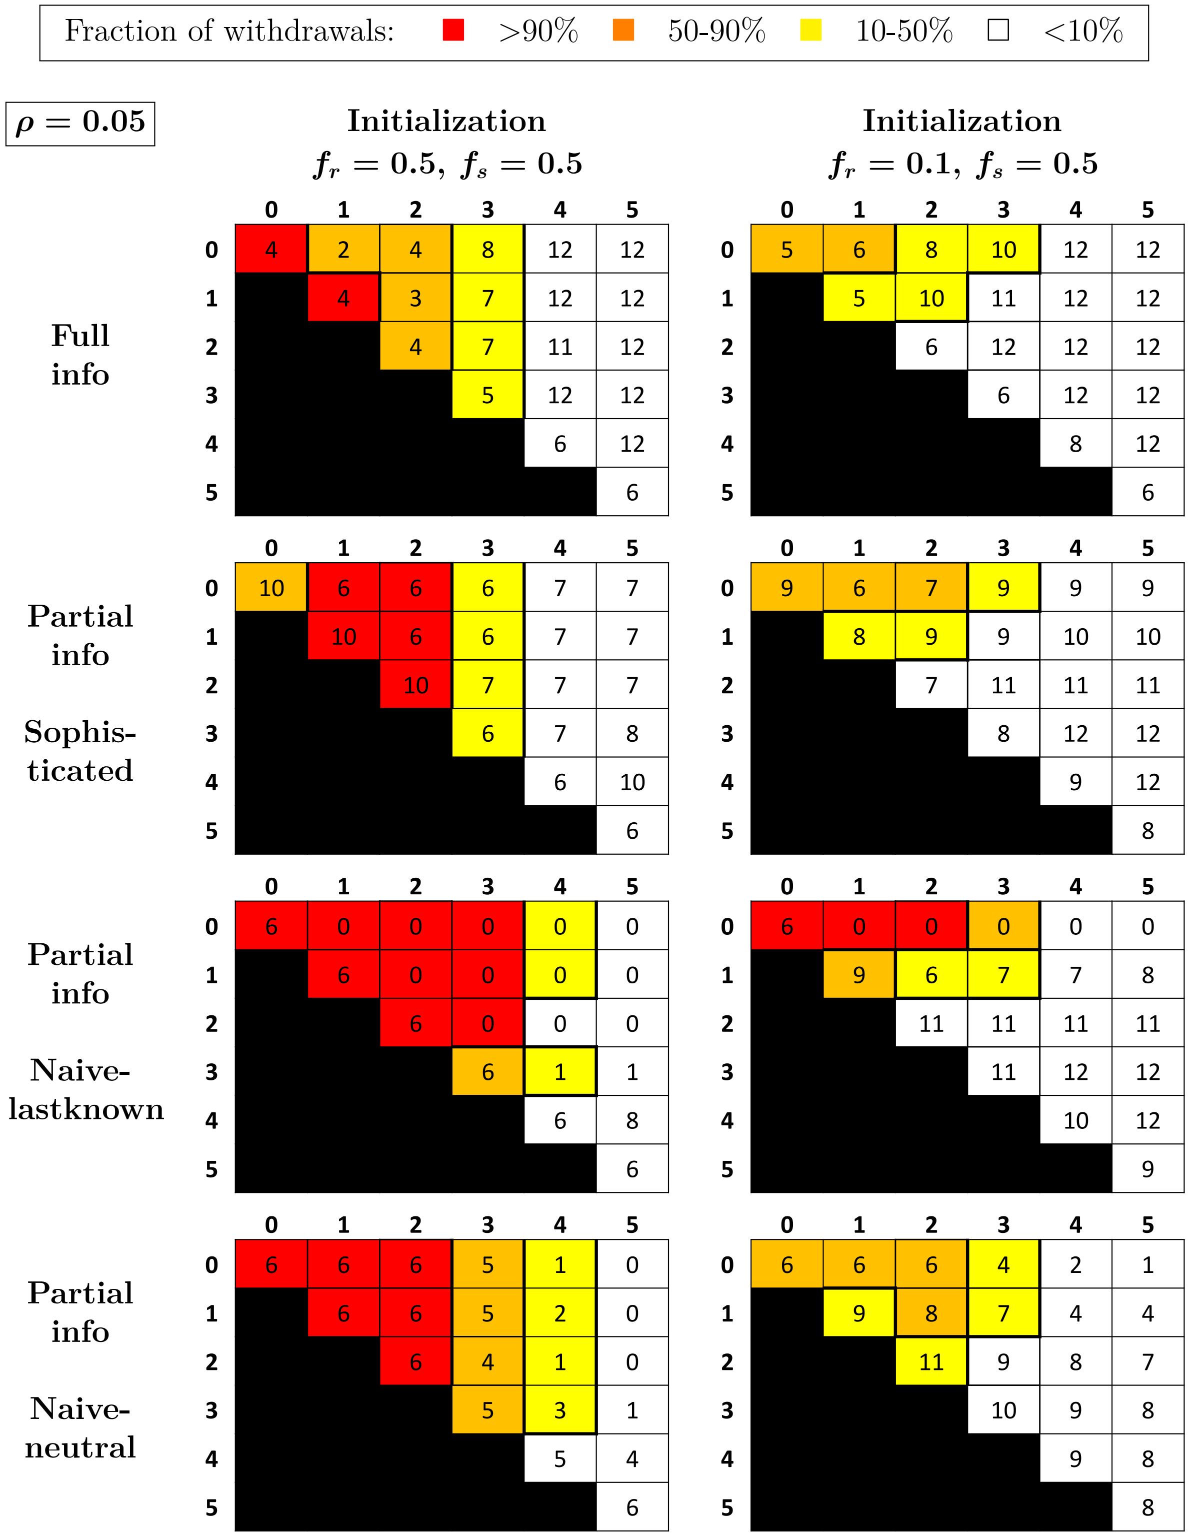

Figure 1: IEL simulation results. Each table shows information about the average strategies in rounds 26 to 50 for all combinations of the 6 bank types defined in Table 1. The color represents the average fraction of withdrawals from both banks and the number represents how many of the 12 agents choose to deposit in the risky bank.

in the risky bank is very low. Therefore we also show, in the tables on the right, simulations in which the agents that start the first round by depositing in the risky bank initially have a low probability of withdrawing ($f_r = 0.1$). The different rows show variations in information and beliefs, with full information simulations in the top row and partial information simulations with the sophisticated, naive-lastknown, and naive-neutral beliefs in the other rows.

The individual cells of the tables each show two pieces of information: the average fraction of agents that withdraw from their bank (shading) and a measure of how many agents deposited in the risky bank (the number). The darkest shading (red) corresponds to more than 90% withdrawals, while no shade (white) means that less than 10% is withdrawn on average. The number in a cell corresponds to the number of depositors in the risky bank. It ranges from 0 (all deposit in the safe bank) to 12 (all deposit in the risky bank). This means that in a white cell with a number 12, agents converge to waiting in the risky bank (the pure-strategy Nash equilibrium) in almost all simulations. Similarly, a red cell with the number 0 means that agents converge to withdrawing in the safe bank in almost all simulations.

A first thing to note is the large impact of the initialization, a sign that path-dependence plays an important role in the dynamics. Regardless of the information or the beliefs, a reduction in the number of withdrawals from the risky bank in the first round helps to prevent runs 25 rounds later. Information and beliefs are also important, particularly when the initialization is not favorable ($f_r = 0.5$). Under those circumstances agents in the IEL simulations only manage to converge to the Nash equilibrium under full information. With partial information there are also few withdrawals for combinations with at least one relatively safe bank (type 4 or 5), but only agents with sophisticated beliefs seem to be able to learn to deposit in a riskier bank (type 0, 1, 2, or 3).

With a favorable (and more realistic) initialization the differences are not nearly as large. Particularly striking is the similarity between the full info simulations and the partial info simulations with sophisticated beliefs. The only difference here seems to be that convergence to the Nash equilibrium is not always complete, either because some agents in the group do not learn it or because some groups do not converge to it (we will come back to this point later). Simulations with naive-lastknown beliefs differ primarily from these two when a bank of type 0 is present (agents never deposit in it if they have an alternative) and in simulations with naive-neutral beliefs the Nash equilibrium choice is much less common. A common element of all tables is that there are more withdrawals when the banks that agents can choose from get riskier.

The majority of the simulations do not result in the (almost certain) convergence to a particular strategy. There are two possible causes of their in-between values. It may be that the individual runs of the simulation simply do not converge to any strategy or there is a lot of variation between the runs, with some runs converging and some others not, or to another strategy. To test this we checked for convergence in every single run using the following criterium for convergence:

1. A minimum of 90% of the agents (which is 11 out of 12 in a simulation with 12 agents) choose the same strategy for two consecutive rounds at some point in the first 50 rounds.
2. In at least 75% of the rounds between the first of the two consecutive rounds mentioned above and the the end of round 50, agents coordinate on this strategy (with a minimum of 90% of the agents choosing this strategy).

We find that it is very common that runs do not converge to any strategy. However, there are also many simulations for which some runs converge to one strategy and other runs to another strategy. The only strategy that runs rarely converge to is withdrawing in the risky bank. Simulations for which only some runs converge are also common.
The convergence tests also provide us with information about how quickly the agents converge to a strategy. For convergence to the Nash equilibrium strategy, both the initialization and the information play a large role. When the fraction of withdrawals in the risky bank starts low ($f_r = 0.1$), convergence often occurs within 4 or 5 rounds. When half of the first-round depositors in the risky bank withdraws, convergence is less common and requires full information. It also takes about twice as long (7 to 10 rounds). With the favorable initialization ($f_r = 0.1$) we also see convergence to the Nash equilibrium strategy when one of the banks is very risky (type 0 or 1) and agents have partial information. Under those circumstances agents need sophisticated beliefs to succeed in coordinating on the Nash equilibrium and it takes a few more rounds (also 7 to 10 rounds).
The simulation results thus reveal that even though theoretically there is a unique pure-strategy Nash equilibrium of the bank choice game, it might not be reached by agents. The convergence to this Nash equilibrium might depend on the relative riskiness of the two banks, the information provided to agents and the initialization. We test the first two factors in an experiment, while the third factor emerges endogenously.

4 Experimental design

4.1 Main design

In the experiment, which we designed to be conducted online, a group of 12 subjects repeatedly play the game discussed in Section 2 for 50 rounds. Subjects start each round with the same initial endowment of 75 points, and face the same decision: first they have to choose between two banks where they want to deposit their endowment. After choosing their bank, they decide whether to wait or withdraw their funds from the chosen bank. Between the bank choice and withdrawal decisions subjects do not receive any information about what others are doing. In particular, they have to make their withdrawal decisions without knowing how many others are in the same bank. Figures 2 and 3 illustrate subjects' decision screens in the two different decisions. After each round subjects receive information about their own payoffs, and - depending on the treatment - they get feedback about either their chosen or both banks. The banks that subjects can choose from are called Round Bank and Square Bank. The Round Bank is always the riskier bank that has a higher potential maximum payoff as well. The Square Bank is the safer bank with a lower potential maximum payoff. In each round subjects face the same two bank types, each chosen from Table 1. One of the two is always a type 1 bank, the type of the other depends on the treatment. In Section 4.2 we discuss the treatments in more detail.

Subjects' payoffs in each round depend on their own choices, the number of others in their chosen banks, and the number of withdrawals in that bank. Payoffs are determined by the formulas (1) and (2). Subjects are not given the exact formulas, but they are presented with payoff graphs, in which they can see the payoffs for the different actions given the fraction of withdrawals in the chosen bank. Appendix A.3 contains the payoff graphs of all bank types used in the experiment. Given that the number of people in a bank may vary, as it is a decision subjects make, we opted for presenting the payoffs as a function of the fraction of withdrawals. Subjects' total payoff of the bank choice game is their cumulative payoff across the 50 rounds.

For each decision subjects have 1 minute to submit their choices. If they do not submit anything within that time, they do not earn anything in that round. If subjects do not choose a bank in time, then the program skips the withdrawal decision, and subjects proceed to a wait page telling them they have not chosen a bank. Payments for the other subjects are then determined by excluding those who have not chosen a bank. If subjects do not make a withdrawal decision, then they again earn nothing for the given round, and their decision counts as 'wait' for those in the same bank, as they do not withdraw their funds. If a subject does not make decisions for three rounds, they are excluded from the experiment, and will not get paid. This is known to subjects in advance.[7]

[7] As the experiment was conducted online, we needed such an exclusion criterium. Note that subjects still have around 4 minutes to solve connection issues or contact the experimenter should they encounter a problem without being excluded from the experiment, as they have 1 minute for each decision. If subjects do not choose a bank, the waiting page telling them they have not chosen a bank is also displayed for 1 minute (with a 'next' button) to give subjects time to solve possible problems.

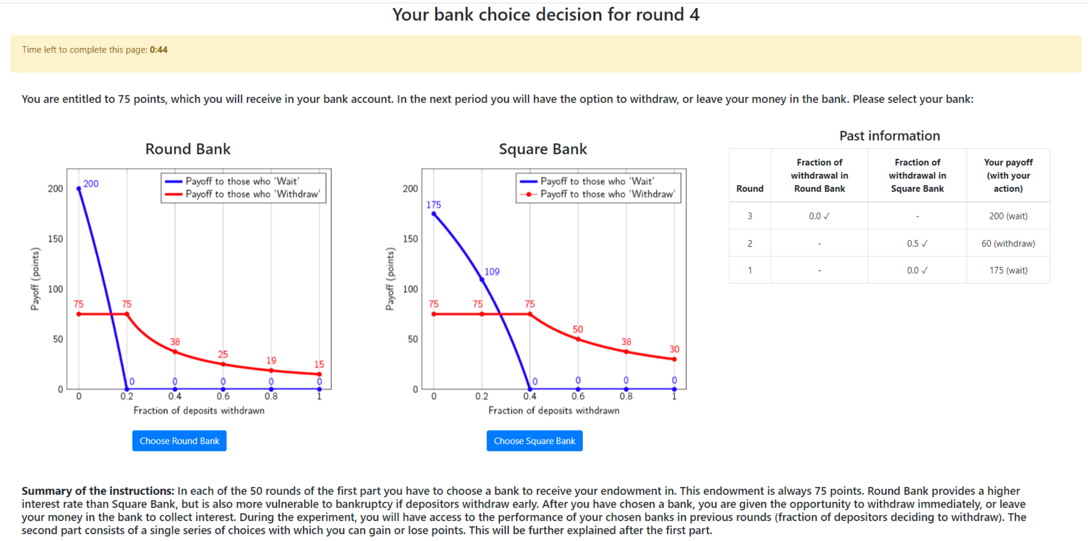

Figure 2: Screenshot of the bank choice decision (partial information)

4.2 Treatments

We implement a 2x3 design, where we vary the information subjects receive as well as the combination of the different banks. In the *partial information* treatments subjects only receive information about what happened in *their chosen bank* in the previous rounds (see the history table on the right hand side of Figure 2). By contrast, in the *full information* treatments subjects receive more detailed information about their chosen bank as well as information about the other bank. In particular, the fraction of withdrawals is presented as a fraction for both banks, where the denominator is the number of depositors in the given bank, and the numerator is the number of withdrawals in the given bank. Subjects' chosen bank is denoted with a tick in the table (see the history table on the right hand side of Figure 3).

On the other dimension we vary the combination of banks used as risky and safer bank. Here we implement three different riskiness level: *low risk*, *medium risk* and *high risk*. Common in all three treatments is that we take bank type 1 from Table 1, and combine it with a different bank. In the low risk treatment the risky bank is a type 1 bank and the safe bank is of type 4, which is a relatively safe type as 80% of the depositors need to withdraw before the bank becomes insolvent. In the medium risk treatment we combine bank type 1 (risky bank) with bank type 2 (safer bank). These two types are close to each other, resulting in moderately high risk. Finally, in high risk bank type 1 is actually the 'safer' bank, and it is combined with the even riskier bank type 0. Note that for this bank type a single withdrawal will already cause insolvency in most cases. Table 2 gives an overview of the different maximum earnings and liquidation values implemented in the experiment.

Based on the simulations and earlier bank run studies with different coordination param-

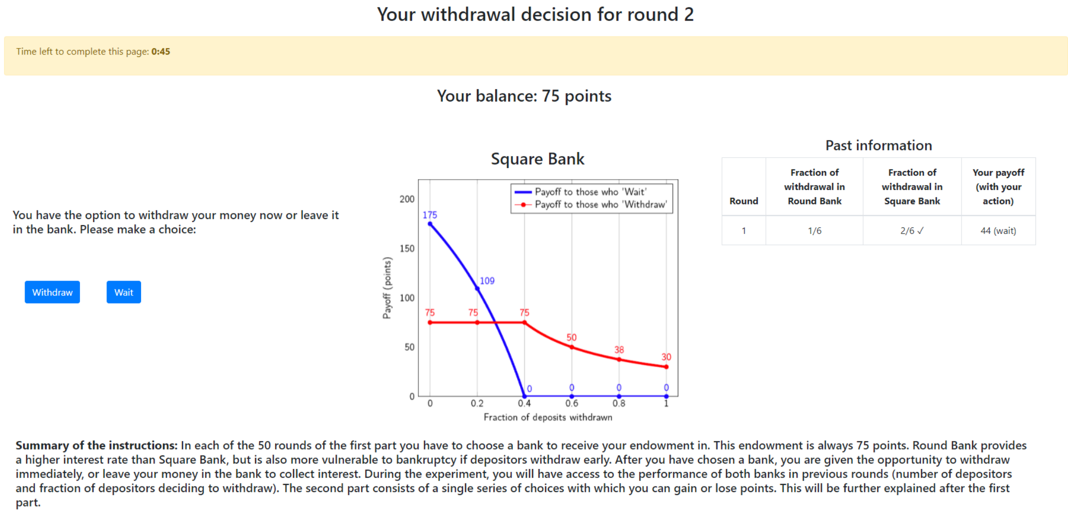

Figure 3: Screenshot of the withdrawal decision (full information)

Table 2: Bank combinations in the experiment

	low risk		medium risk		high risk	
	Round b.	Square b.	Round b.	Square b.	Round b.	Square b.
Bank type	1	4	1	2	0	1
Max. earnings	200	125	200	175	225	200
Liq. value	0.2	0.8	0.2	0.4	0.1	0.2

Notes: Maximum earnings can be reached if all depositors in the given bank decide to wait. It is calculated using the initial endowment of 75. Graphical representation of the payoffs can be found in Appendix A.3.

eters, we expect that the combination of banks affects the ability of subjects to learn the equilibrium of waiting in the risky bank. The safer the safe bank is, the more likely it is that subjects end up in the pure-strategy Nash equilibrium of the game. Looking at the IEL simulations with random initial decisions, the expected outcome in the low risk and the medium risk treatments are different. However, if we look at the initialization in which agents wait with 90% probability after choosing the risky bank, we find that only in simulations of the high risk treatment agents cannot learn the pure-strategy Nash equilibrium. The second initialization, in which agents are more likely to end up in this equilibrium is in line with observations from the one-shot game in de Jong (2021).[8] This leads us to our first hypothesis:

Hypothesis 1 *The safer the safe bank is, the more likely it is that agents can learn the pure-strategy Nash equilibrium, i.e. waiting in the risky bank.*

Next to the variation in bank types, we consider different information settings. In the

[8]A type of forward-induction reasoning could underly the choices observed in that experiment. Intuitively, agents choose the risky bank only if they want to wait, to enjoy higher earnings, as withdrawing is safer from the safe bank.

full information treatments subjects have the best chance to learn to wait in the risky bank. However, information this detailed is hardly observable in real life, so we decided to also investigate learning with just partial information about what happens in the banks. Giving only partial information makes coordination more difficult to subjects, as they need to try out each bank if they want to know how others behave in a particular bank. Therefore we expect similar or less coordination on the pure-strategy Nash equilibrium with partial information. Also, in case participants do converge on using this strategy, it is expected to be slower than under full information. This leads us to our second and third hypotheses:

Hypothesis 2 *Compared to partial information, full information leads to similar or more coordination on waiting in the risky bank.*

Hypothesis 3 *Partial information leads to slower convergence than full information.*

4.3 Procedures

The experiment was programmed in oTree (Chen et al., 2016), and was run as an online experiment with subjects from the University of Amsterdam and Simon Fraser University in Vancouver between May and October 2021. In total 574 subjects participated in 48 sessions. Subjects were mainly students in various fields. Most subjects (46.1%) studied economics or business economics, followed by natural sciences, mathematics, computer science or engineering (13.7%) and social sciences (excluding economics and psychology - 11.3%). 56% of the subjects were female, 43% were males (8 subjects either stated other or did not want to answer). The average age was 21.8 years.[9] None of the subjects participated more than once.

After playing the bank choice game, which was the main part of the experiment, the participants did one additional task (loss aversion elicitation). However, since participants' choices in this task might have been influenced by their experience in the experiment and because the correlation with loss aversion was never intended to be an important part of our investigation, we decided not to analyze this data further. See Appendix A.4 for more details about this task and the instructions.

Sessions took on average one hour with average earnings of 14.90 euros in Amsterdam and 19.90 dollars in Vancouver (including a participation fee of 5 euros, 7 dollars, respectively). Subjects' earnings consist of their cumulative point earnings of the 50 rounds of bank choice game, their earnings from the loss aversion task plus the participation fee. Point earnings from the experiment were exchanged to euros or dollars with 675 points for 1 euro, and 450 points for 1 dollar.

[9]Comparing subjects' gender, age and field of study across treatments with a Kruskal-Wallis test reveals no significant differences across treatments ($p > 0.58$).

In most sessions 12 subjects participated from one of the two locations.[10] We collected 4 groups per treatment per location, resulting in 8 independent observations per treatment. Subjects were one-by-one admitted to a zoom session where we checked their ID, renamed them, and placed them back into the waiting room. Once we prescreened all participants, we admitted all of them in the zoom session, and sent away subjects if more than 12 had shown up. The remaining participants were sent a unique oTree link via private chat. Subjects were only allowed to communicate with the experimenter via the private chat, but not with each other. Subjects read the experimental instructions at their own pace, and had to correctly answer understanding questions before starting the bank choice game. Both the instructions and the understanding questions are reproduced in Appendix A. After everybody correctly answered all questions, subjects played the bank choice game for 50 rounds. After the main game, subjects performed the loss aversion task (these instructions are also found in Appendix A). Finally, subjects filled in a post-experimental questionnaire to provide more information about their strategies, and some demographics.

5 Experimental results

5.1 Coordination on the different outcomes

Figure 4 shows the timeline of the actions that subjects choose in each round. There are four action combinations possible, each corresponding to a pure strategy of the bank choice game: depositing in the risky bank and waiting (red), depositing in the safe bank and waiting (green), depositing in the safe bank and withdrawing (blue), and depositing in the risky bank and withdrawing (black). The rows represent different information structure and source of data: full (Full) or partial (Part) information, and Exp denotes rows with experimental data, whereas IEL denotes rows with the corresponding simulations. For partial information we have three different possible beliefs, sophisticated (Soph), naive-lastknown (Naive-LK) and naive-neutral (Naive-N). The three columns show data on the three different levels of riskiness: low, medium and high risk.

The first thing to note is that in the low risk and medium risk treatments very few participants withdraw. In those treatments most groups converge to playing the Pareto-optimal Nash equilibrium strategy (with sometimes one or two participants deviating by depositing in the safe bank). In the high risk treatments we do not see any sign of convergence to this Nash equilibrium, nor do the participants coordinate on any other action. When we apply the convergence criteria outlined in Section 3.3, we find that none of the high

[10] In three sessions one participant left the experiment during the instructions, in three additional sessions one participant left the experiment during the bankrun game, and in two sessions we started with 11 subjects due to low show-up. The program did not depend on the exact number of subjects, and could handle drop-outs. The 574 subjects include the drop-outs as well.

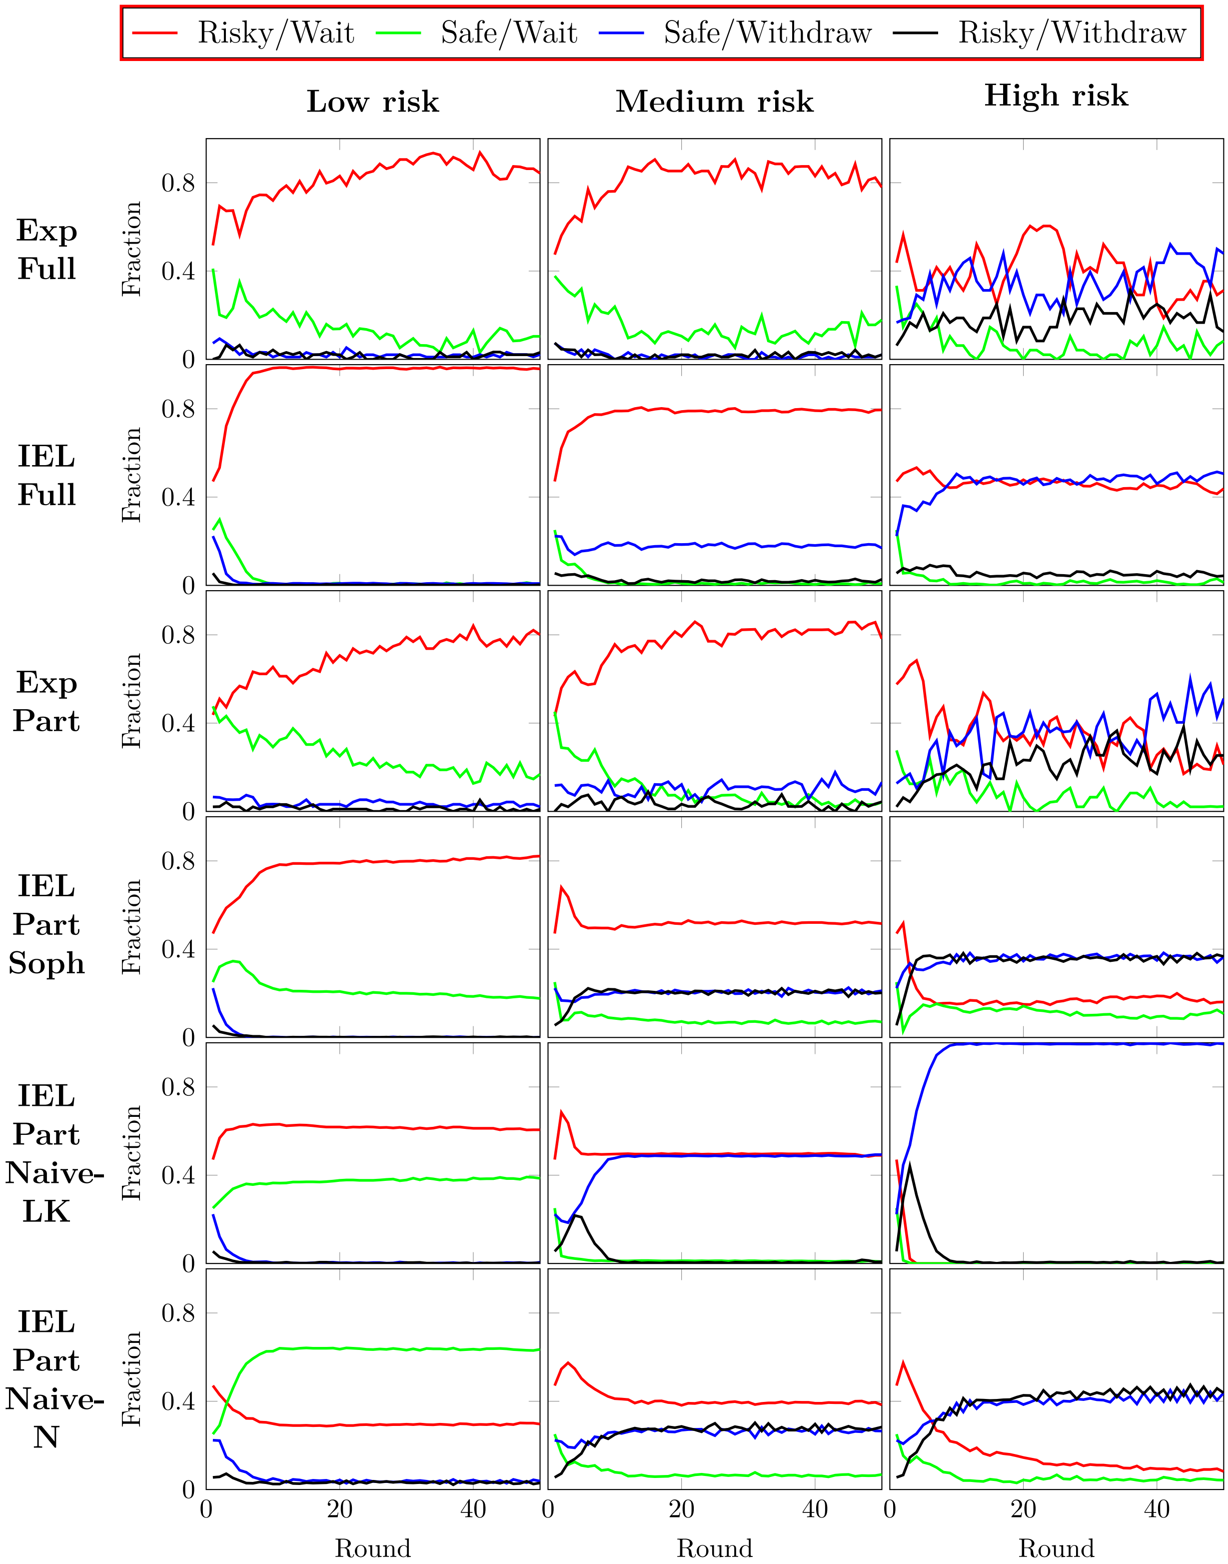

Figure 4: Strategies in the experiment and predicted by IEL

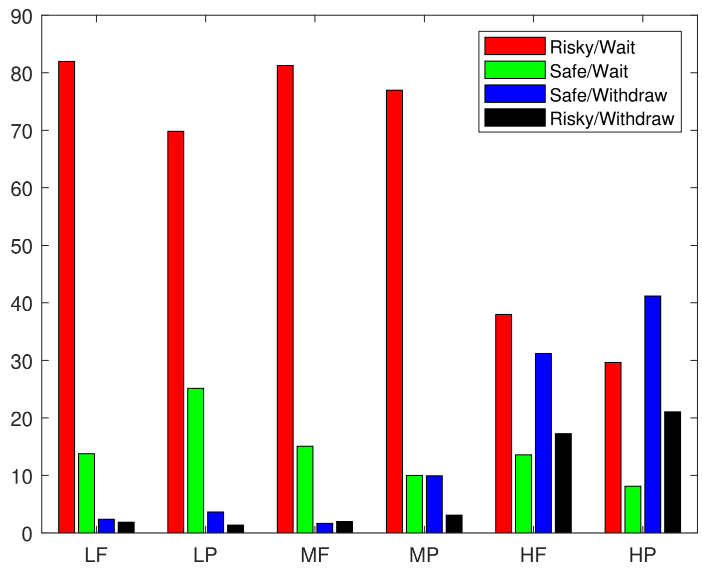

Figure 5: Percentages of chosen actions per treatment in the whole experiment

risk sessions converges to any of the four pure strategies in 50 periods, compared with 6 out of 8 (low risk, full information), 2 out of 8 (low risk, partial information), 5 out of 8 (medium risk, full information), and 4 out of 8 (medium risk, partial information) in the other treatments. In Appendix C we provide the data for each of the individual experimental sessions. Second, the IEL simulations describe the experimental data relatively well when agents have sophisticated beliefs under partial information, or possess full information. Naive beliefs provide a much worse fit.

As we can see from Figure 4, the dynamics seem to be different for high risk compared to the two other risk types. However, the differences seem to be less pronounced looking across information structures, even though there seems to be a slight advantage of having full information for coordination. Figure 5 shows the aggregate behavior in all treatments for the whole experiment. The figure shows the same strategies as in Figure 4, but as a percentage of all chosen strategies in the given treatment. It confirms the previously mentioned result that subjects under high risk behave differently compared to the other two risk treatments, but the effect of information is less pronounced. For all the analyses presented in this section we pool the data across the two locations because we are interested in the treatment effects in general, and the purpose of the experiment is not to investigate location differences. We do find some differences though and we present those in Appendix D.

To test the differences seen in Figures 4 and 5, we investigate treatment effects by means of multinomial panel logit regressions. Table 3 presents the results of these regressions with the chosen actions as dependent variable and individuals as panels. The base action

Table 3: Multinomial panel logit regressions − all data

Action: (base: Risky/Wait)	Safe/Wait	Safe/Withdraw	Risky/Withdraw
Withdrawal fraction in chosen bank$_{t-1}$	2.435***	3.998***	2.757***
	(0.406)	(0.305)	(0.307)
Full information	−0.057	−0.802*	−0.281
	(0.393)	(0.342)	(0.329)
Medium risk	−0.616	0.008	0.177
	(0.511)	(0.605)	(0.548)
High risk	−0.382	3.312***	3.305***
	(0.557)	(0.389)	(0.425)
Round	−0.045***	−0.009	−0.007
	(0.008)	(0.005)	(0.005)
Constant	−1.488**	−4.823***	−5.050***
	(0.512)	(0.451)	(0.441)
# of observations	27,792		
# of panels	571		
Log Likelihood	−16,491.349		

Notes: ***: significant on 0.1%-level,**: significant on 1%-level, *: significant on 5%-level. The columns show the estimated multinomial logit coefficients with individuals as panels. The reference level is the Nash equilibrium action: choosing the risky bank and wait. Full information is 1 for the full info treatments, 0 otherwise. Medium risk is 1 for the medium risk treatments, 0 otherwise. High risk is 1 for the high risk treatments, 0 otherwise. Standard errors are clustered on the session level, and are in brackets.

in the multinomial logit is waiting in the risky bank (a Nash equilibrium strategy). The different columns show the effect of the independent variables on the probability of a given action *relative* to this Nash equilibrium action. The independent variables are a dummy for full information (1 for full information, 0 for partial information), a dummy for medium risk (1 for the medium risk treatments, 0 otherwise), a dummy for high risk (1 for the high risk treatments, 0 otherwise), round number and – to control for path dependency – the withdrawal fraction in the *last chosen bank* from the last period. This information is available in both information treatments, and given the group interaction, it seems to be a more relevant piece of information than subjects' own previous decision.[11] First turning our attention to the relative risk, we find no significant treatment effects between low risk and medium risk (see coefficients 'medium risk' in Table 3). The relative likelihood of choosing any particular action compared to the Nash equilibrium action is not different in these two treatments. However, when we look at high risk, we see that the relative likelihood of choosing an action with withdrawal (columns 2 and 3) compared to choosing the Nash equilibrium action increases with high risk. In the high risk treatments subjects coordinate less often on the Nash equilibrium action, and end up more often

[11] The inclusion or exclusion of the last observed withdrawal fraction does not substantially changes the results on treatment effects (see Table E.1 in Appendix E).

in withdrawing their funds from the bank compared to the low risk treatments. These observations are consistent with Figures 4 and 5 showing a more pronounced difference between high risk and the other two treatments, and no substantial difference between low and medium risk.[12] These observations lead us to our first result, which relates to Hypothesis 1.

Result 1 *We do not see a strict monotonic relationship between the level of coordination on the pure-strategy Nash equilibrium and the riskiness of the safest bank. Subjects in low and medium risk treatments coordinate more often on this Nash equilibrium without a clear ordering between the treatments. High risk leads to a significant shift from the Nash equilibrium action to actions involving withdrawing from the chosen bank.*

Next we look at the effect of the information structure. Table 3 shows that for most actions the information structure does not have an effect. The only significant result we obtain is for the action withdrawing from the safe bank, for which the odds-ratio relative to the Nash equilibrium action decreases under full information. However, this result is not observed when we include the interaction terms in the regressions (see Table E.2 in Appendix E) or when we run the same multinomial regression per riskiness level (see Table E.3 in Appendix E).[13] Therefore, even though we see a significant treatment effect in Table 3, we are cautious in interpreting it and we draw the following conclusion, related to our second hypothesis.

Result 2 *On average, full information does not lead to more coordination on waiting in the risky bank.*

Before turning to the speed of convergence, we conclude that high risk has the strongest effect of all treatment dimensions. It is important to note that path dependency indeed plays a large role in our experiment, as expected (remember, IEL is also dependent on the initialization). In fact, the coefficient of the previous-period withdrawal fraction is highly significant for all three strategies. The more subjects withdrew in a given round, the less likely agents choose the Nash equilibrium action in the following round. Note that a direct comparison between the magnitude of this coefficient and the treatment effect is less useful, as treatments are dummy variables, whereas the withdrawing fraction is a variable taking values between 0 and 1. The impact of an additional subject choosing to withdraw depends on the number of depositors in the given bank.

[12] The results are robust to the inclusion of the interaction terms between treatments and / or controls for demographics (see Table E.2 in Appendix E).

[13] Figure 5 suggests more coordination on the pure-strategy Nash equilibrium under full information than under to partial information. Running a panel logit model with binary outcome (1 - Nash equilibrium action, 0 - otherwise) does not reveal any information effect, but confirms the effect of high risk. The results of the logit models are available upon request.

Table 4: Comparison of first rounds of convergence

	$\boldsymbol{t_c}$ full	$\boldsymbol{t_c}$ partial	$\bar{t}_c$ full	$\bar{t}_c$ partial	full vs partial
low risk	9, 11, 3, 4, 19, 26	7, 5	12	6	0.64
medium risk	12, 5, 4, 5, 15	13, 27, 8, 5	8	13	0.40
high risk	-	-	-	-	1

To test our hypothesis on the speed of convergence, we check for each experimental session if convergence took place (according to our criterion in Section 3.3) and, for those that did, record the first round of convergence. Table 4 displays the lists of first convergence rounds ($\boldsymbol{t_c}$) for all treatments, their averages ($\bar{t}_c$), and the p-value of a Mann-Whitney U test comparing the results of full and partial information. Although the averages differ considerably from each other, we do not find any significant effects due to the low number of partial information sessions in which we observe convergence. Therefore we cannot draw any definite conclusions about whether the amount of information provided to the agents affects the speed of convergence or not.

Result 3 *We do not find a significant effect of the information structure (full or partial information) on the speed of convergence to any particular action.*

5.2 First round decisions

To get more insights in the observed dynamics, we turn to the beginning of the experiment. Did subjects decide differently in the first round already, or did they learn to converge to a given outcome later? Figure 6 shows which actions subjects choose in the first round of the experiment. First note that in the first round the distributions of strategies look more similar across treatments, and especially across risk levels, than in the case in which we include all the rounds.[14] By combining the strategies involving the choice of the same bank (columns 1 and 4, and columns 2 and 3) we can see that in all treatments around half of the subjects (43-55%) decided to choose the risky bank regardless of the information and the relative riskiness of the two banks. Furthermore, we observe much less withdrawing from the banks than waiting. This is even more pronounced when we compare withdrawal rates in the two bank types: the action 'Safe bank and withdraw' is chosen more often than the action 'Risky bank and withdraw'. This is in line with the

[14] Given that in the first round all individual decisions are independent, we also provide test-statistics based on the two-sided Kolmogorov-Smirnov test on pairwise comparisons. None of the pairwise comparisons between treatments (by either keeping the information structure or the risk-level fixed) results in significant differences between the distributions (the smallest $p-$value is 0.199 for the test between LF and HF).

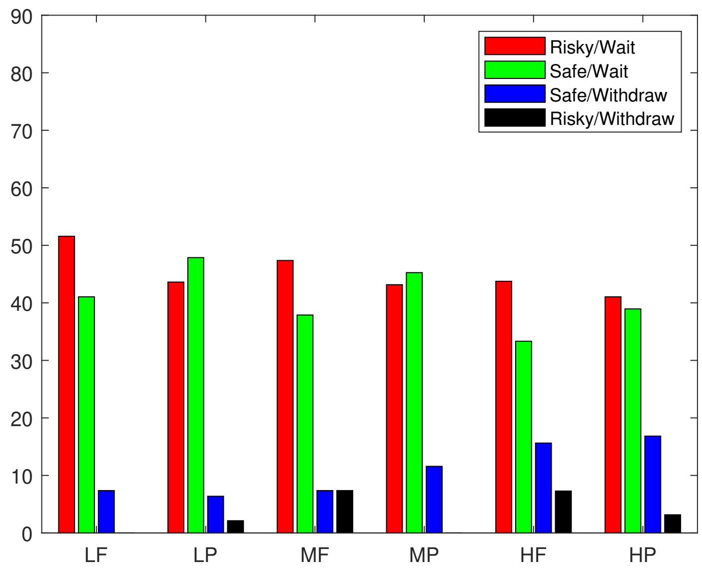

Figure 6: Percentages of actions per treatment in the first round

intuitive forward induction argument of more withdrawal from the safe bank, and it is also in line with the findings of de Jong (2021).[15]

Even though we do not see many significant treatment differences in the first round decisions, the resulting outcomes of these decisions are already different. While on average banks do not become insolvent with these average withdrawal rates under low risk, and medium risk, this is not the case for high risk. For high risk most banks start with some withdrawal which can be detrimental for later rounds, as insolvency was more likely to occur.

In the non-parametric tests we performed pairwise comparisons with subsamples of the data that only contained the two treatments we compared to each other. To further investigate the general treatment effects, we pooled all the data, and looked at multinomial logit regressions as well. This way we investigate the treatment effects also from a different angle. Table 5 reports the results of this logit regression. We estimate how the different treatments affect the log-odds relatively to the base category. This reference category is again the pure-strategy Nash equilibrium choice: choosing the risky bank and wait. Here we only focus on the effect of riskiness, because information only affects differences in the feedback which is not yet available before the first round.[16] The results show

[15]We formally test the withdrawal rate across bank types keeping the treatment fixed. In line with the forward induction argument, we use one-sided proportion tests to test whether withdrawing rate in the risky bank is lower than in the safe bank. The difference is significant on the 1%-level for medium and high risk under partial information and for low risk under full information ($p = 0.001$, $p = 0.003$ and $p = 0.002$, resp.), and significant on the 5%-level for high risk under full information ($p = 0.02$).

[16]Further robustness checks showing similar effects (including dummy for full information, as well as controls for demographics) are relegated to Table E.4 in Appendix E.

Table 5: Multinomial logit regression for the first round choices

Action: (base: Risky/Wait)	Safe/Wait	Safe/Withdraw	Risky/Withdraw
Medium risk	-0.016	0.371	1.300
	(0.218)	(0.394)	(0.817)
High risk	-0.091	0.974^{**}	1.715^{*}
	(0.223)	(0.365)	(0.790)
Constant	-0.069	-1.935^{***}	-3.807^{***}
	(0.152)	(0.297)	(0.716)
# of observations	570		
Log Likelihood	-607.415		

Notes: ***: significant on 0.1%-level,**: significant on 1%-level, *: significant on 5%-level. The columns show the estimated multinomial logit coefficients. The reference level is the Nash-equilibrium action: choosing the risky bank and wait. Medium risk is 1 for the medium risk treatments, 0 otherwise. High risk is 1 for the high risk treatments, 0 otherwise. Heteroskedasticity-robust standard errors are in brackets.

that medium risk does not have a significant effect on the relative probabilities to choose among the different strategies. However, high risk increases the probability of choosing an action with withdrawal relative to making the pure-strategy Nash equilibrium choice (see columns 2 and 3 in Table 5). These findings are consistent with Figure 6, and could be observed in the whole experiment as well, with a stronger significance.

6 Conclusion

Bank runs are relatively rare phenomena and there are few people that are subject to them more than once or twice in their lifetimes. This has led some authors to focus primarily on single-shot games or first-round decisions (e.g. Schotter and Yorulmazer, 2009 or de Jong, 2021). However, there are two reasons why learning can be important for studying bank runs. The first is that the bank run setting is only one specific example of a coordination problem. Coordination problems in general are quite common. People may learn from earlier (mis)coordination and apply this also when they are facing a decision to withdraw from a bank or not. The second reason is that as long as depositors manage to avoid bank runs, the coordination problem stays intact. This may be less relevant in a study that aims to look at the parameters that (consistently) trigger bank runs, but does become relevant when we broaden the analysis, for example by including deposit decisions. Unlike in many bank run experiments, the (fewer) withdrawals in the single-shot bank choice game, in which subjects make both a deposit and a withdrawal decision, rarely cause banks to become insolvent (de Jong, 2021). This is mainly because those who are reluctant to bear the strategic uncertainty that comes with the 'wait' decision initially avoid depositing in risky banks. About half the subjects leave a considerable amount of money on the table with this choice. This initial situation is therefore also unlikely to persist and when depositors move from one bank to another (riskier) bank, it is natural for all depositors to reevaluate also the withdrawal decision again.

In this paper we study dynamics and learning in the bank choice game both with a (repeated) bank choice experiment and with simulations. For the experiment we used the most basic version of the game. Subjects can choose between a 'risky' and 'safe' bank to deposit an initial endowment in and after making this choice they have to decide whether to withdraw that deposit or not. This is repeated for 50 rounds with the same two banks. As in the earlier single-shot study, both banks receive a similar share of depositors in the first round. In most sessions the initial withdrawals are also quite low. At this point we see few significant differences between treatments with different levels of riskiness of the two banks. However, this changes in subsequent periods. In sessions which combine low- or moderate-risk banks with high-risk banks, the majority of depositor groups converge on playing the single pure-strategy Nash equilibrium of the game: deposit in risky bank and not withdrawing. However, in sessions with only high-risk banks subjects do not coordinate on any action, not even after 50 rounds.

In our experiment we also varied the amount of information that was shared with subjects at the end of each round. In the full information treatments subjects could see how many depositors each bank had and how many of those withdrew. In the partial information treatments they were only informed about the fraction of withdrawals in the bank that they deposited in. In the experiment we only noticed a minor impact of information on the subjects' decisions. However, this treatment variation has large consequences for the simulations with the individual evolutionary algorithm, because the partial information condition required us to extend the algorithm with beliefs. We tried several possible ways of introducing beliefs and found that the only way in which we can reproduce the experimental results is by assuming that agents are very sophisticated in their belief-updating. Our simulations also suggest that learning is faster under full information. Unfortunately we do not have enough data points to assess the validity of this claim.

Apart from the insights we gain on belief-updating, our results may also have implications for the external validity of other experiments under full information. With the latter we mean in this case that subjects have enough information to calculate the payoffs they would have received if they would have chosen differently in the last round. These experiments are quite common. Although one may question if this is a realistic feature, our results show that there do not have to be substantial differences between results in full and partial information environments.

Our paper is the first paper investigating the dynamics and learning in a bank choice game, where agents can choose where to deposit their initial endowment and then subsequently decide whether they withdraw their money or not. As a first step we chose to study the effects of the relative riskiness of the two banks and the information structure in a static environment. This leaves many directions for further investigation. One possibility is to study an environment in which the agents receive information in between the bank choice and withdrawal decisions, for example about the size of the bank. This opens up new strategies as it allows agents to condition their actions on the information they receive. Another is to let the interest rates and riskiness of the banks evolve endogenously, also depending on bank size for example. This adds a level of reality that we would be interested in, but have not been able to combine with our main goals in this paper. We leave these directions for future research.

References

Anufriev, M., J. Arifovic, J. Ledyard, and V. Panchenko (2022): "The role of information in a continuous double auction: An experiment and learning model," *Journal of Economic Dynamics and Control*, 141, 104387.

Arifovic, J. (2019): "Evolution of sunspot like behavior in the agent based economies of bank runs," *Journal of Evolutionary Economics*, 29, 365–389.

Arifovic, J., C. Hommes, A. Kopányi-Peuker, and I. Salle (2023): "Ten isn't large! Group size and coordination in a large-scale experiment," *American Economic Journal: Microeconomics*, 15, 580–617.

Arifovic, J. and J. H. Jiang (2019): "Strategic uncertainty and the power of extrinsic signals–evidence from an experimental study of bank runs," *Journal of Economic Behavior & Organization*, 167, 1–17.

Arifovic, J., J. H. Jiang, and Y. Xu (2013): "Experimental evidence of bank runs as pure coordination failures," *Journal of Economic Dynamics and Control*, 37, 2446–2465.

Arifovic, J. and J. Ledyard (2007): "Call market book information and efficiency," *Journal of Economic Dynamics and Control*, 31, 1971–2000.

——— (2011): "A behavioral model for mechanism design: Individual evolutionary learning," *Journal of Economic Behavior & Organization*, 78, 374–395.

Brown, M., S. T. Trautmann, and R. Vlahu (2017): "Understanding bank-run contagion," *Management Science*, 63, 2272–2282.

Chen, D. L., M. Schonger, and C. Wickens (2016): "oTree—An open-source platform for laboratory, online, and field experiments," *Journal of Behavioral and Experimental Finance*, 9, 88 – 97.

Davis, D. D. and R. J. Reilly (2016): "On freezing depositor funds at financially distressed banks: An experimental analysis," *Journal of Money, Credit and Banking*, 48, 989–1017.

de Jong, J. (2021): "The initial deposit decision and the occurrence of bank runs," Working Paper.

Diamond, D. W. and P. H. Dybvig (1983): "Bank runs, deposit insurance, and liquidity," *Journal of Political Economy*, 91, 401–419.

Duffy, J., A. Karadimitropoulou, and M. Parravano (2019): "Financial contagion in the laboratory: Does network structure matter?" *Journal of Money, Credit and Banking*, 51, 1097–1136.

Kiss, H. J., I. Rodriguez-Lara, and A. Rosa-García (2012): "On the effects of deposit insurance and observability on bank runs: an experimental study," *Journal of Money, Credit and Banking*, 44, 1651–1665.

——— (2014): "Do women panic more than men? An experimental study of financial decisions," *Journal of Behavioral and Experimental Economics*, 52, 40–51.

——— (2016): "Think twice before running! Bank runs and cognitive abilities," *Journal of Behavioral and Experimental Economics*, 64, 12–19.

——— (2018): "Panic bank runs," *Economics Letters*, 162, 146–149.

Kiss, H. J., I. Rodriguez-Lara, and A. Rosa-Garcia (2022a): "Experimental bank runs," *Handbook of Experimental Finance*, 347.

——— (2022b): "Preventing (panic) bank runs," *Journal of Behavioral and Experimental Finance*, 35, 100697.

——— (2022c): "Who withdraws first? Line formation during bank runs," *Journal of Banking & Finance*, 140, 106491.

König-Kersting, C., S. T. Trautmann, and R. Vlahu (2022): "Bank instability: Interbank linkages and the role of disclosure," *Journal of Banking & Finance*, 134, 106353.

Madies, P. (2006): "An experimental exploration of self-fulfilling banking panics: Their occurrence, persistence, and prevention," *The Journal of Business*, 79, 1831–1866.

Peck, J. and K. Shell (2003): "Equilibrium bank runs," *Journal of Political Economy*, 111, 103–123.

Schotter, A. and T. Yorulmazer (2009): "On the dynamics and severity of bank runs: An experimental study," *Journal of Financial Intermediation*, 18, 217–241.

Shakina, E. (2019): "Bank runs as a coordination problem within a two-bank set-up: Who will survive?" *Economics Letters*, 177, 85–88.

Trautmann, S. T. and R. Vlahu (2013): "Strategic loan defaults and coordination: An experimental analysis," *Journal of Banking & Finance*, 37, 747–760.

A Instructions

In this Appendix we reproduce the experimental instructions. In Section A.1 the instructions for the bank choice game in treatment low risk, partial info is given for both locations. In Section A.2 we list the differences between the different treatments to give a complete view of all instructions. Section A.3 graphically presents the 4 different bank types we used in the experiment. Finally, in Section A.4 the instructions for the loss aversion task is presented. These instructions were identical in all bank choice treatments.

A.1 Instructions for treatment LP

This subsection presents the instructions for the treatment low risk, partial information. The differences in locations only consist of a different exchange rate, and the different information needed to pay subjects out. These differences are added in italics in brackets in the text. Bold fonts were also bold in the experiment. The correct answers for the understanding questions are added in italics in brackets after each question.

PAGE 1

Welcome!

Today you participate in an experiment on economic decision making. In this experiment we use **anonymization** to ensure that neither the experimenters, nor the other participants know who is behind the decisions that you take. To be able to pay you, we will ask you for an [(Vancouver) ***email address*** / (Amsterdam) ***IBAN-number***], but not your name. Also, this information will be removed from the data set once your payment is processed.

Note that only those who **fully complete** the experiment **will get paid**. Therefore, if you experience errors or connection problems of any kind, contact the experimenters via the Zoom chat or, if that is not possible, send an email to [*email Vancouver* / *email Amsterdam*].

This page is followed by a page with instructions. Please read these instructions carefully and answer the five-question quiz on the following page in order to proceed to the experiment. On the quiz page, you can revisit the instructions by scrolling downwards. Once all participants have answered all quiz questions correctly, the experiment will start. At the end of the experiment, you will be asked to fill out a short questionnaire.

By clicking the button below, you consent to the collection of your participation data for the sole purpose of research. Thank you for your participation and good luck!

PAGE 2

Introduction

This experiment consists of two parts. The first part has 50 rounds in which you can earn points. The second part consists of a single page with choices with which you can win more points or lose some. When both parts are finished the total number of points that you earned is converted into [*dollars / euros*] at a rate of [*450 points for 1 dollar / 675 points for 1 euro*]. How much you will earn exactly depends on your choices and the choices of other participants, but it can add up to a considerable amount. Additionally, you will receive a [*$7.00 / €5.00*] participation fee.

General information

In each round of the first part you have to make two choices. First you select the bank in which you want to receive your endowment of 75 points. You can choose between two options: Round Bank and Square Bank. In Round Bank you can potentially earn more interest, but it is also more vulnerable when people withdraw early. In one of the sections below we will go into more detail about this tradeoff.

When everyone selected a bank, you are given the choice to either withdraw your endowment immediately, or wait and collect interest. If few people withdraw, you earn more by leaving the money in the bank. However, if many withdraw, the bank will not be able to fully pay out all depositors and those who decided to wait, lose their money.

In the next round you again start with an endowment of 75 points. Any money that you earned in the previous round is set aside to be paid out at the end of the experiment.

The depositors

You are participating in this experiment with 11 other participants. All receive the same endowment, are given the same information, and face the same choices as you are confronted with. Upon choosing a bank, some may decide to deposit at Round Bank and some others at Square Bank. Your payoff will only depend on your decision and the decisions of those who chose to deposit in the same bank. The decisions of those who deposited in the other bank will not affect your earnings in that round. At the time that you have to make the decision to withdraw or not, you do not know exactly how many others deposited in the same bank. At the end of the round you learn what fraction of the total depositors of your bank chose to withdraw their deposits.

The banks

The tradeoff between the interest rate and the bank's vulnerability is best explained using a graph. Below you find how much a depositor receives when choosing to withdraw (red curve) and choosing to wait (blue curve) as a function of the fraction of total deposits that is withdrawn in the particular bank in that period. The graph on the left is for Round Bank and the graph on the right for Square Bank.

Round bank **Square bank**

[Figures A.1b and A.1d appear here next to each other in the instructions.]

Let's start with Round Bank. When none of the depositors in Round Bank withdraws, they all receive 200 points. However, when a fraction of 0.2 of the depositors withdraws (that is 1 out of 5 depositors), those who withdraw receive 75 points and those who wait receive nothing. When more than a fraction of 0.2 of the depositors withdraws, those who withdraw receive less than 75 points, while those who wait still receive nothing.

In Square Bank the maximum earnings when everyone waits, are lower: 125 points. However, when a fraction of 0.2 of the depositors withdraws those who wait still receive 117 points. Only when a fraction of 0.8 of the depositors withdraws (4 out of 5 depositors), those who decided to wait receive nothing. So, although the maximum earnings in Square Bank are considerably lower, it is also considerably less vulnerable to early withdrawals.

Decision time

In each round you have enough but limited time to make your decisions. You have 1 minute to choose your bank, and then again 1 minute to choose whether you wait or withdraw your deposit. If you don't make at least one of your decisions on time, you earn nothing for the given round. You are not counted then as a depositor for the others either. If you don't make the bank choice decision on time, you will not see the screen for the second decision. If you fail to make at least one of your decisions in 3 rounds of the experiment, you cannot continue with the experiment, and you will not be paid at all.

Information

The graphs with the payoffs in the two banks will be visible on the screen when you have to choose a bank to deposit in. Next to the graphs a table with the past performance of the banks will be provided. It shows the fraction of depositors withdrawing in your bank. You will only see this information for one of the banks for each round: the one you chose in the given round. In case a bank attracted no depositors in a particular round, the fraction of withdrawing depositors cannot be calculated. In that case the fraction is represented by a dash ('-'). The last column of the table shows how many points you earned in that round.

Contact

During the experiment you can always contact the experimenter using the Zoom chat function. The experimenter will first try to help you via chat. If the problem cannot be resolved that way, you will be invited into a break-out room where full Zoom functionality can be enabled (including screen sharing if necessary). If for any reason you cannot contact us via Zoom, you can send an email to [*email Vancouver / email Amsterdam*].

Summary

In each of the 50 rounds of the first part you have to choose a bank to receive your endowment in. This endowment is always 75 points. Round Bank provides a higher interest rate than Square Bank, but is also more vulnerable to bankruptcy if depositors withdraw early. After you have chosen a bank, you are given the opportunity to withdraw immediately, or leave your money in the bank to collect interest. During the experiment, you will have access to the performance of your chosen banks in previous rounds (fraction of depositors deciding to withdraw). The second part consists of a single series of choices with which you can gain or lose points. This will be further explained after the first part.

On the next screen you are asked to answer some questions to test your understanding of the experiment.

PAGE 3

Understanding Questions

1. Suppose you end period 1 with earnings of 150 points. How many points can you

deposit in the bank in period 2? [*Answer: 75*]

2. Suppose you do not choose a bank in a given period. What are your earnings in points in that period? [*Answer: 0*]

3. For this question you have to use the graphs provided in the instructions below. Suppose you choose the Square Bank. Next, you decide to wait, but some other depositors withdraw. The withdrawing depositors constitute a fraction of 0.2 of the total number of depositors in Square Bank in that round. What are your earnings in points in this period? [*Answer: 117*]

4. For this question you have to use the graphs provided in the instructions below. Suppose you choose the Round Bank with 4 other depositors. The earnings graph associated with Round Bank is shown above. You and another depositor decide to withdraw the deposit. What are your earnings in points in this period? [*Answer: 38*]

5. Do you receive information about the number of depositors in your chosen bank?

   a. No [*correct answer*]

   b. Yes

A.2 Differences in other treatments

In this subsection we detail out the differences between our treatments by listing the changes compared to the instructions in Section A.1. We divide the differences into different risk levels and different information.

Different risk levels: For the different risk levels the following texts change:

- For medium risk we use bank types 1 (Figure A.1b) and 2 (Figure A.1c) for Round and Square bank, respectively. For high risk we use bank types 0 (Figure A.1a) and 1 (Figure A.1b) for Round and Square bank, respectively.

- The numbers in the paragraph 'Let's start with Round Bank...' change in high risk. The maximum earnings of 200 becomes 225, but the threshold fraction decreases from 0.2 to 0.1 (thus 1 out of 10 depositors).

- The numbers in the paragraph 'In Square Bank' changes for both other risk levels. The maximum earnings are 175 and 200 points in medium, and high risk, respectively. The example fraction of depositors in the second sentence matches the threshold for the Round Bank, thus changes to 0.1 in high risk from 0.2. The payoff corresponding to this fraction is 109 in medium risk, and 111 in high risk. The threshold fraction of withdrawals of 0.8 (4 out of 5 depositors) changes to 0.4 (2 out of 5 depositors) in medium risk and to 0.2 (2 out of 10 depositors) in high risk.

- The correct answer to Understanding question nr. 3 changes from '117' to '109' in medium risk, and to '0' in high risk.
- The correct answer to Understanding question nr. 4 changes from '38' to '19' in high risk.

Different information: In the treatments with full information there are the following differences:

- Under 'Information' two sentences, starting by 'It shows the fraction...' change to 'It shows the fraction of depositors withdrawing in your bank by indicating the number of withdrawals and the number of depositors in the bank (e.g. 2/5 means that 2 out of 5 depositors withdrew their money from the bank). This information will be given for both banks, also the one in which you did not deposit in that particular round.'
- Under 'Summary' the sentence starting with 'During the experiment,...' changes to 'During the experiment, you will have access to the performance of both banks in previous rounds (number of depositors, and fraction of depositors deciding to withdraw).'
- The correct answer to Understanding question nr. 5 ('Do you receive information...?') changes from 'No' to 'Yes'.

A.3 Bank types

Figure A.1 shows the payoffs of the different actions in the different bank types as the function of fraction of withdrawals.

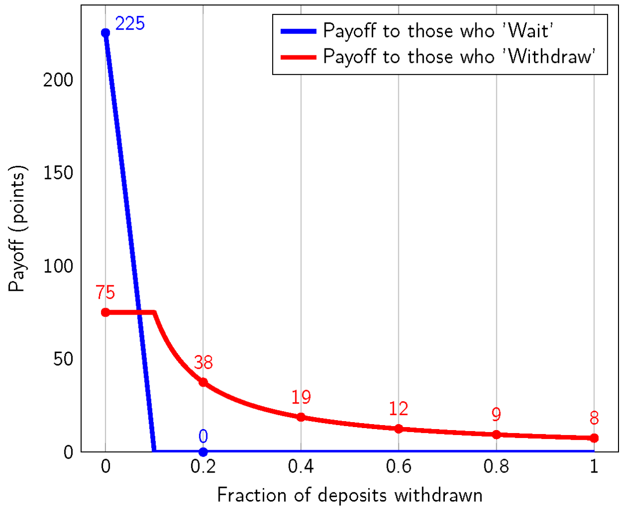

(a) **Type 0**

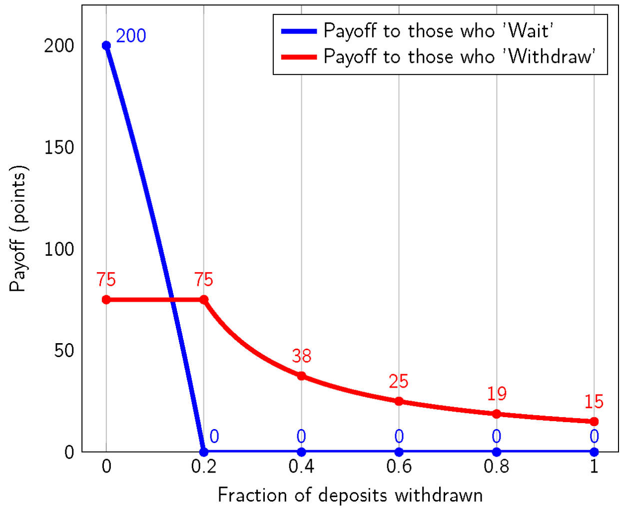

(b) **Type 1**

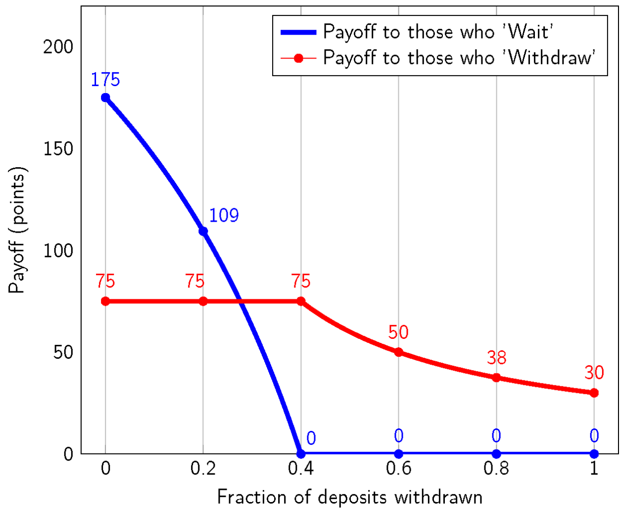

(c) **Type 2**

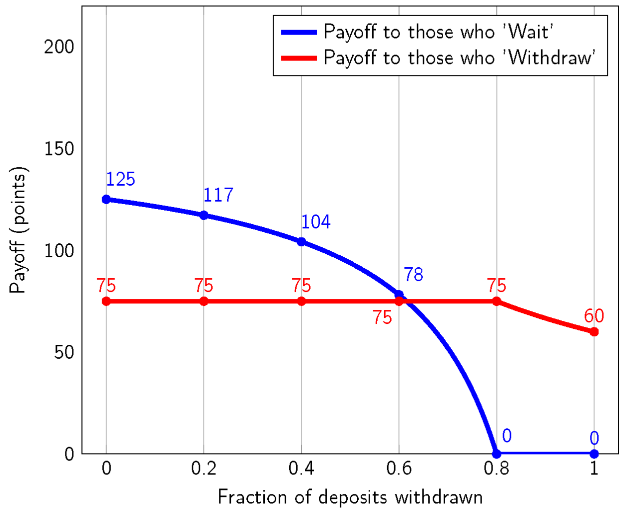

(d) **Type 4**

Figure A.1: Bank types used in the experiment

A.4 Loss aversion task

In this subsection we discuss the loss aversion task and reproduce its instructions. We implemented a loss aversion task after the 50 rounds of the main experimental task as previous studies show some correlation between loss aversion and decisions in bankrun games (Trautmann and Vlahu, 2013; Kiss et al., 2018). By contrast, this correlation was not found for risk aversion (Kiss et al., 2014, 2016, 2018; Shakina, 2019). Before the bank choice game subjects are aware that there is a second task they need to perform, but

they do not know the nature of that task until they finish the first part. The second task consists of 6 lotteries, and subjects need to choose whether they want to play that particular lottery or not. After all subjects make their decisions, one of the lotteries and its realization is randomly drawn. Subjects who chose to play that lottery receive the corresponding earnings (which can be positive or negative), whereas subjects who decided not to play the chosen lottery receive nothing. In all 6 lotteries the probability of winning 1350 points is 50%. With 50% probability however subjects can lose money. The lowest absolute loss is 150 points, and this increases to 1650 in steps of 300. For the exact payoffs see the table below in the instruction. We do not impose a time limit on this task.
The instructions for this task was the same in all treatments. The only difference in location was the conversion rate, denoted by italics in brackets below. In the table below radio buttons were presented for subjects to make their decisions.

Instructions for the loss aversion task
Below you find a series of 6 lotteries. In each of them there is a 50% chance to lose points and a 50% chance to gain points. In the end, one of these lotteries will be randomly selected (with equal probabilities). You can indicate for each lottery if you would like to play this lottery if it is selected, or not. If you chose to play a lottery and it is selected, the lottery is played and you will gain or lose points. If you chose not to play that particular lottery, nothing will happen. As a reminder, the conversion rate is [*450 points for 1 dollar / 675 points for 1 euro*].

50%	50%	Accept to play?
-150	1350	Yes / No
-450	1350	Yes / No
-750	1350	Yes / No
-1050	1350	Yes / No
-1350	1350	Yes / No
-1650	1350	Yes / No

B Simulation results for different levels of experimentation

This Appendix presents further simulation results with experimentation rates $\rho = 0.01$ and $\rho = 0.10$. The results are qualitatively similar to the results presented in Section 3.3 with $\rho = 0.05$.

Fraction of withdrawals: ■ >90% ■ 50-90% ■ 10-50% □ <10%

$\rho = 0.01$

Full info

Initialization $f_r = 0.5,\ f_s = 0.5$

	0	1	2	3	4	5
0	1	2	2	9	12	12
1		2	3	7	12	12
2			2	7	11	12
3				5	12	12
4					5	12
5						6

Initialization $f_r = 0.1,\ f_s = 0.5$

	0	1	2	3	4	5
0	5	7	10	11	12	12
1		7	11	12	12	12
2			8	12	12	12
3				8	12	12
4					11	12
5						8

Partial info Sophis-ticated

Initialization $f_r = 0.5,\ f_s = 0.5$

	0	1	2	3	4	5
0	11	5	7	7	7	7
1		11	6	6	7	6
2			10	6	7	6
3				6	6	7
4					6	10
5						6

Initialization $f_r = 0.1,\ f_s = 0.5$

	0	1	2	3	4	5
0	11	7	8	9	9	10
1		10	9	10	9	9
2			10	11	11	12
3				10	12	12
4					11	12
5						10

Partial info Naive-lastknown

Initialization $f_r = 0.5,\ f_s = 0.5$

	0	1	2	3	4	5
0	6	0	0	0	0	0
1		6	0	0	0	0
2			6	0	0	0
3				6	1	1
4					5	8
5						5

Initialization $f_r = 0.1,\ f_s = 0.5$

	0	1	2	3	4	5
0	5	0	0	0	1	1
1		10	7	8	8	8
2			11	11	11	11
3				12	12	12
4					11	12
5						11

Partial info Naive-neutral

Initialization $f_r = 0.5,\ f_s = 0.5$

	0	1	2	3	4	5
0	6	7	7	6	1	0
1		6	7	6	1	0
2			6	5	2	0
3				5	2	1
4					5	4
5						6

Initialization $f_r = 0.1,\ f_s = 0.5$

	0	1	2	3	4	5
0	8	8	8	7	4	3
1		10	10	9	5	5
2			11	10	8	8
3				11	9	8
4					9	8
5						8

Figure B.1: Simulation results with $\rho = 0.01$. Each table shows information about the average strategies in rounds 26 to 50 for all combinations of the 6 bank types defined in Table 1. The color represents the average fraction of withdrawals from both banks and the number represents how many of the 12 agents choose to deposit in the risky bank.

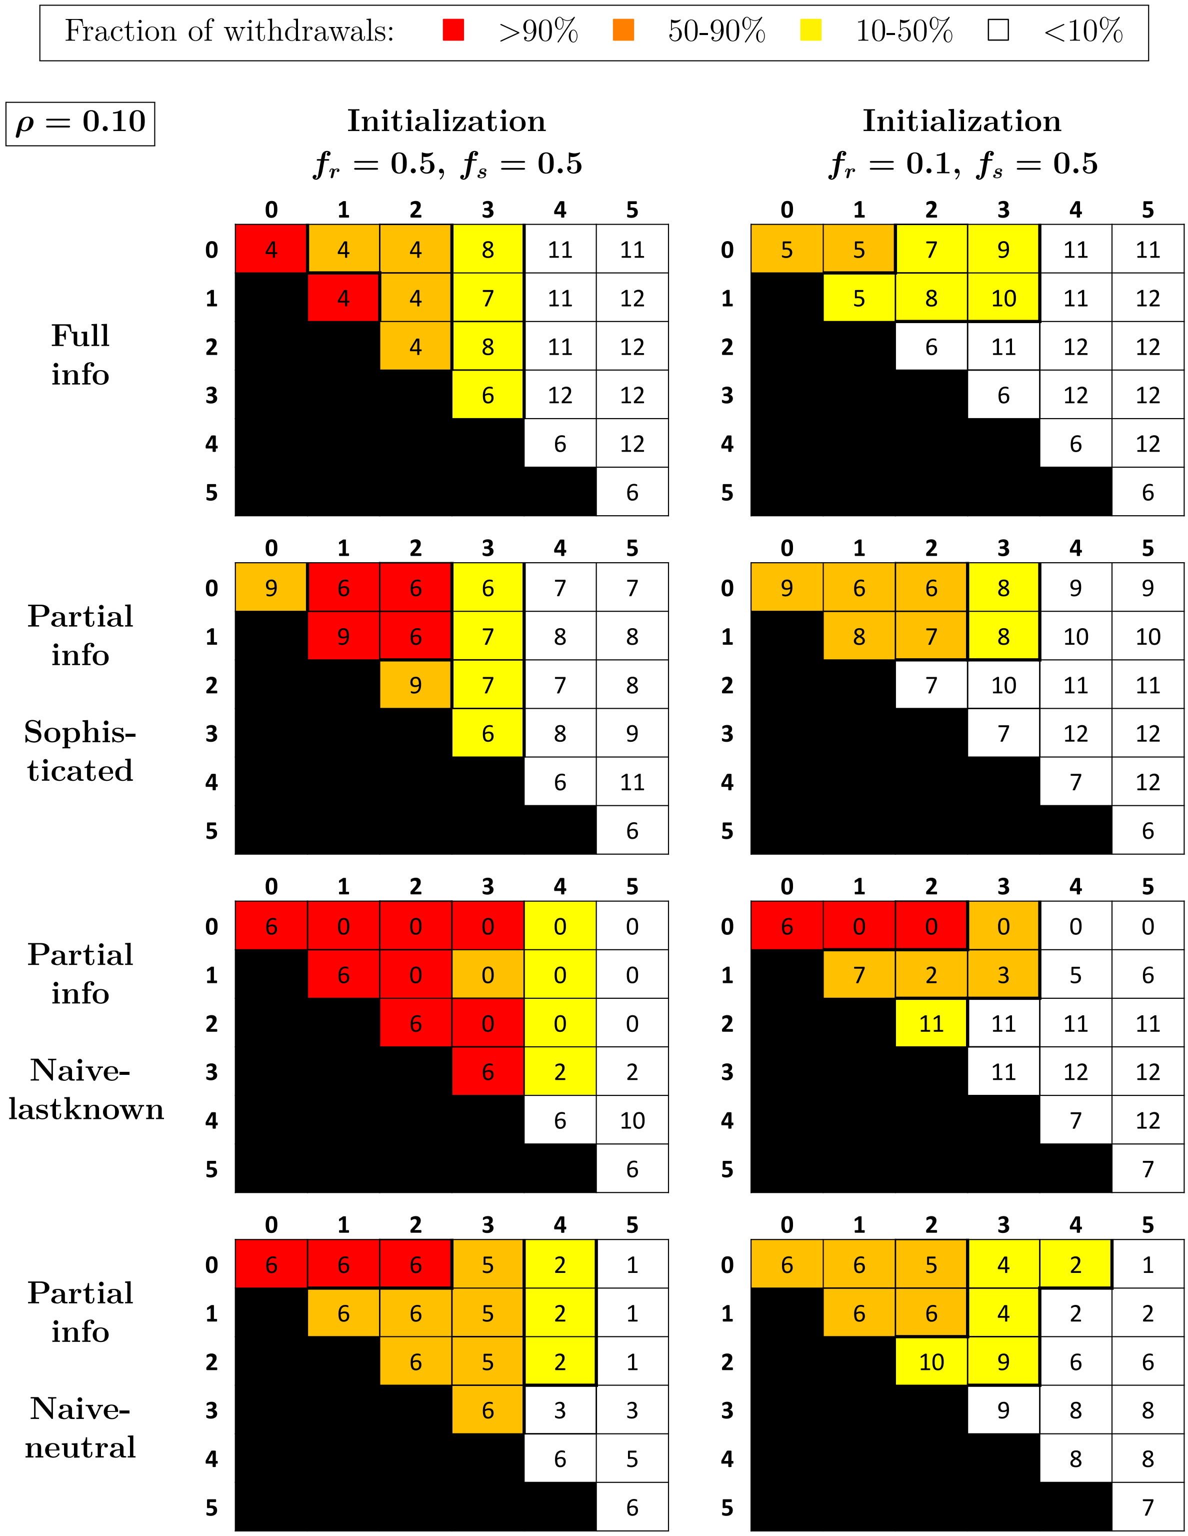

Figure B.2: Simulation results with $\rho = 0.10$. Each table shows information about the average strategies in rounds 26 to 50 for all combinations of the 6 bank types defined in Table 1. The color represents the average fraction of withdrawals from both banks and the number represents how many of the 12 agents choose to deposit in the risky bank.

C Data per session

In this Appendix we provide plots of average strategies over rounds for each individual session. Figure C.1 shows these plots for the treatments with full information and Figure C.2 contains the plots for the partial info treatments. The four strategies are again represented by different colors: red for depositing in the risky bank and waiting, green for depositing in the safe bank and waiting, blue for depositing in the safe bank and withdrawing, and black for depositing in the risky bank and withdrawing.

D Location effect

In this Appendix we discuss the data in the two locations, Amsterdam and Vancouver. First, Figure D.1 gives an overview over the aggregated data per treatment and location. Figure D.2 presents the percentages of strategies chosen over all periods per location and treatment. From the figures we can see that coordination seems to be quicker in Amsterdam than in Vancouver, but the qualitative results with respect to treatment difference are largely the same. High risk reduces coordination on the Nash equilibrium in both locations, whereas information structure seems to have a lesser effect. Given that we are more interested in the aggregated treatment effects (our hypotheses compare treatments) we pooled the data from the two locations for the main analyses. In this appendix we discuss the data per location.

As we can see from Figures D.1 and D.2 comparing full information and partial information within a location does not show very different behavior for most of the cases. Partial information leads to a bit less choice of the Nash equilibrium for both locations and for all treatments. The only exception is the medium risk treatments in Vancouver, where there is a slightly higher percentage of waiting in the risky bank under partial information. In the Nash equilibrium choices, we see the highest difference between low-risk full and partial information treatments in Vancouver, where in both treatments subjects learn to wait, but under full information they wait more often in the risky bank. Looking across riskiness levels, we find no substantial differences between low and medium risk for three out of the 4 cases. Under partial information the low-risk and medium risk treatments seem to differ more in Vancouver than in Amsterdam. However, comparing low and medium risk to high risk, we always see the same relationship: subjects withdraw more often under high risk irrespective of the information structure and the location.

Running the same multinomial logit regression on a sample split per location as in Table 3 in Section 5 shows the same treatment effects for both locations. Table D.1 shows the regression results confirming the same treatment effects as we have for the whole data. The information structure does not have a significant effect on the actions chosen, whereas

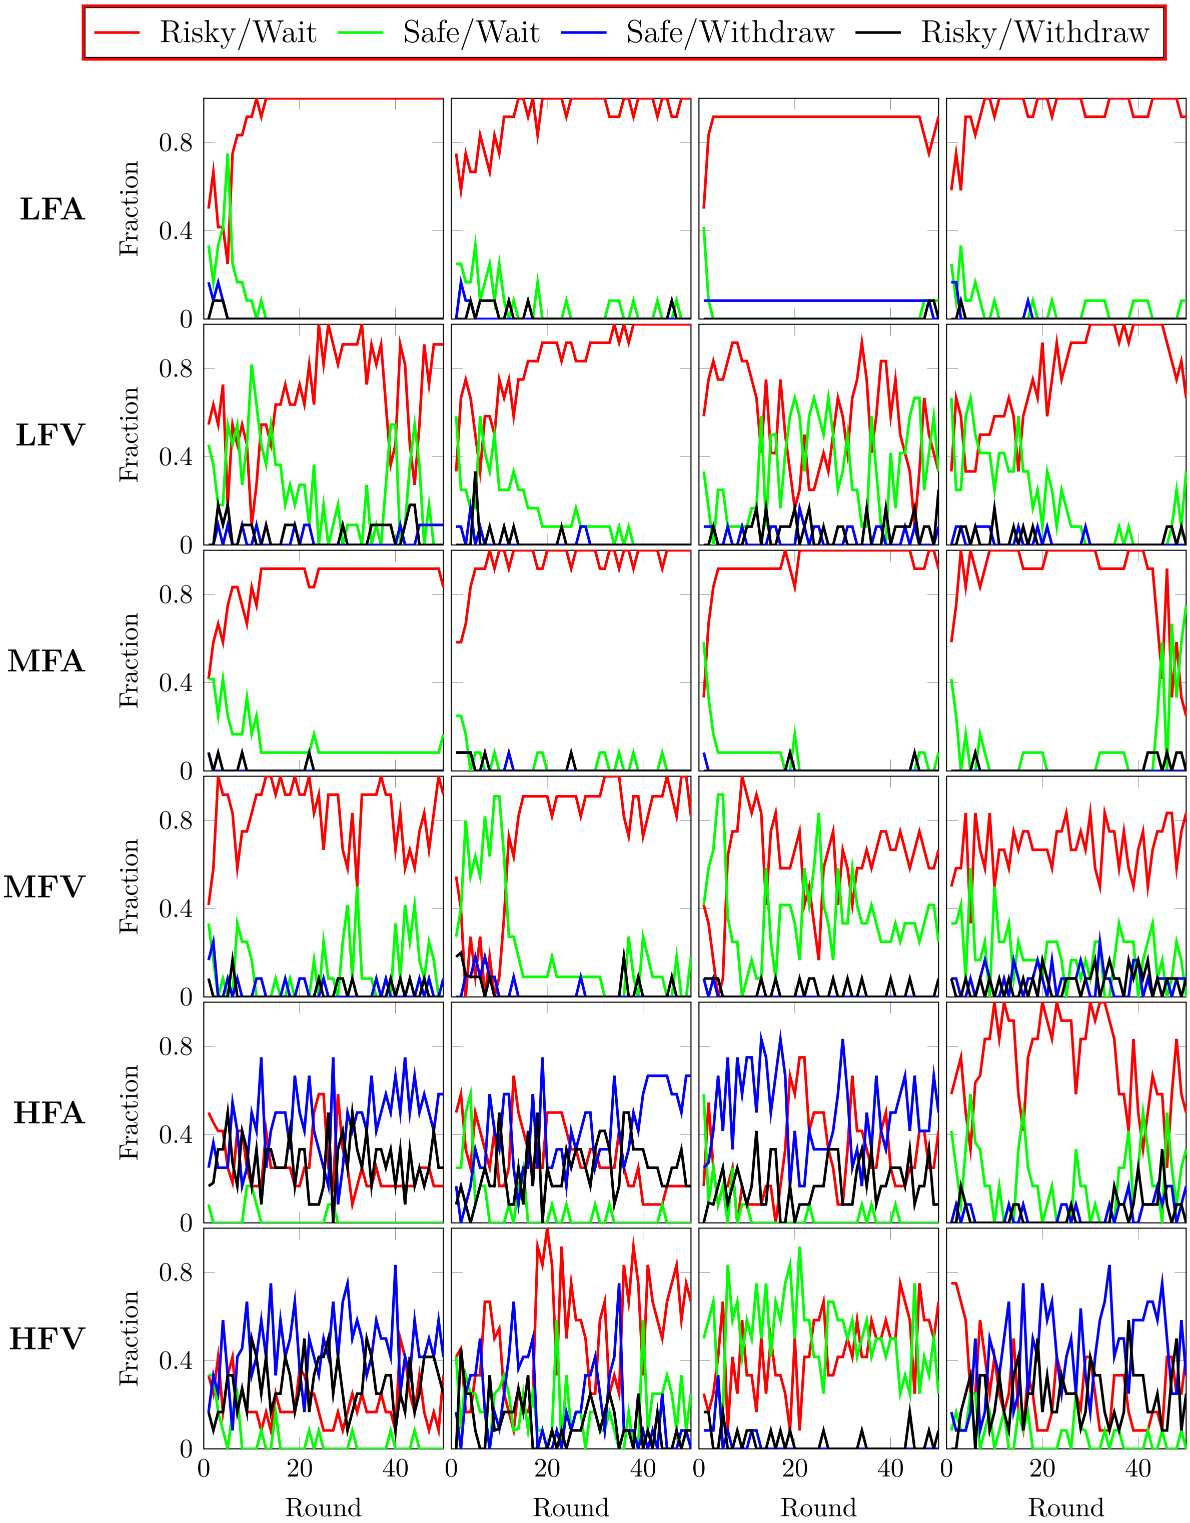

Figure C.1: Strategies per group under full information. The first letter, L, M, or H, indicates whether the combination of banks was low risk, medium risk, or high risk and the last letter, A or V, is used to distinguish Amsterdam and Vancouver sessions. All treatments in this figure are full information treatments (hence the middle letter F).

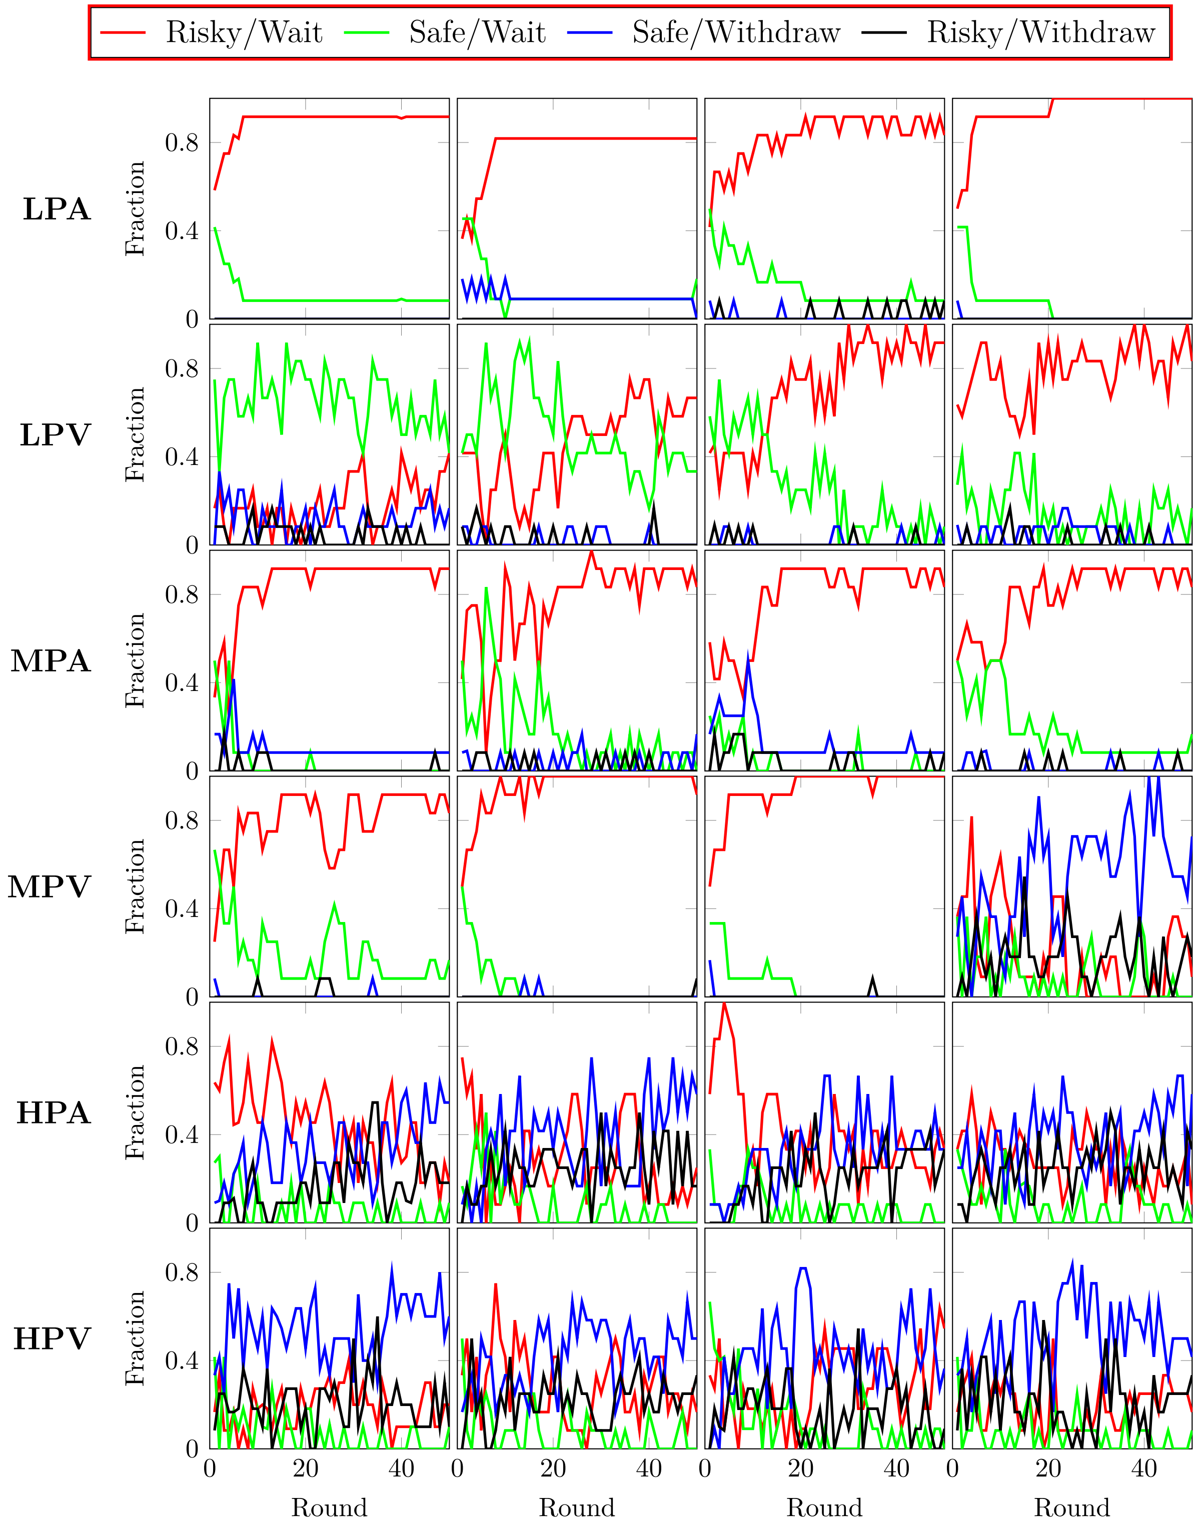

Figure C.2: Strategies per group under partial information. The first letter, L, M, or H, indicates whether the combination of banks was low risk, medium risk, or high risk and the last letter, A or V, is used to distinguish Amsterdam and Vancouver sessions. All treatments in this figure are partial information treatments (hence the middle letter P).

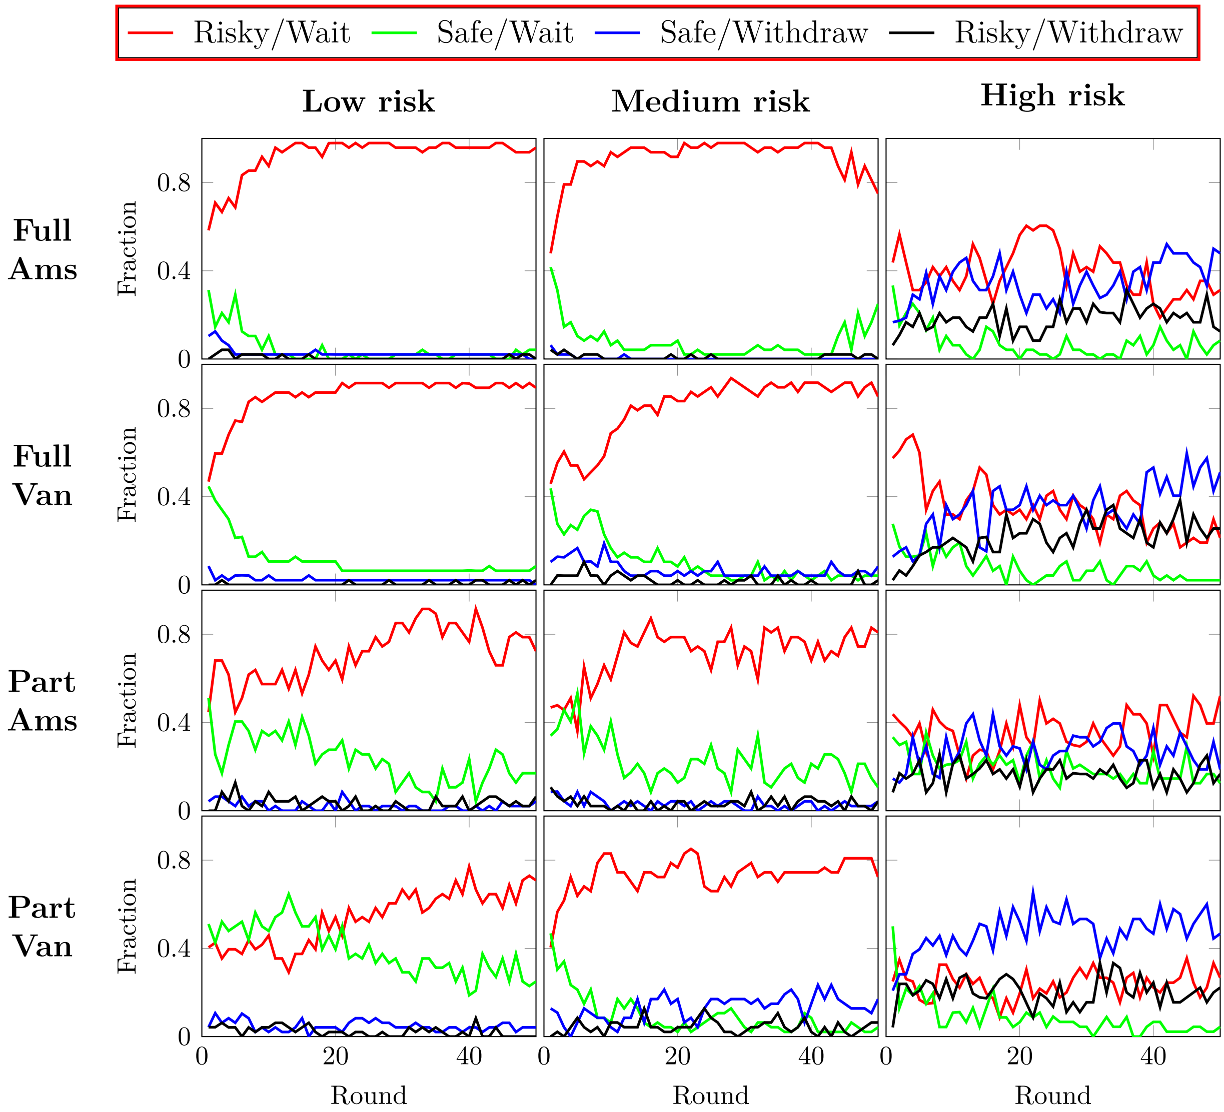

Figure D.1: Strategies in the experiment in Amsterdam and Vancouver. The first two rows correspond to treatments with full, and the third and forth row correspond to treatments with partial information.

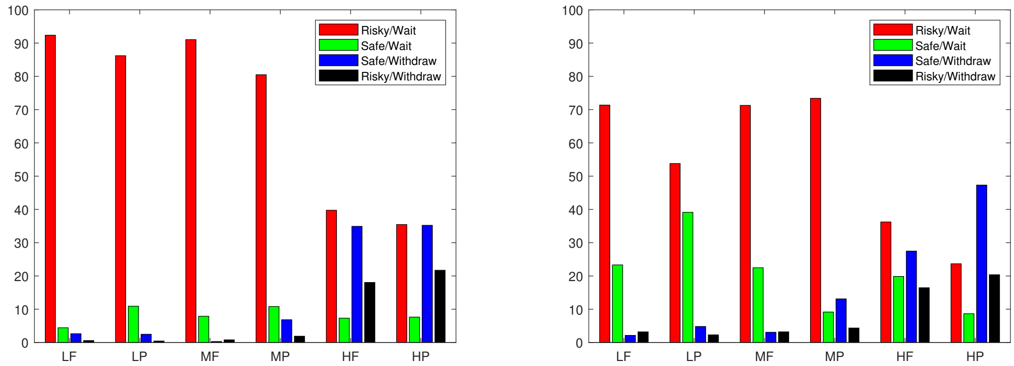

Figure D.2: Percentages of strategies per location and treatment in Amsterdam (left panel) and in Vancouver (right panel)

high risk increases the relative likelihood of choosing an action with withdrawal compared to the Nash equilibrium action.

Looking at the first round decisions, we see different patterns in the two locations. The left panel of Figure D.3 shows the actions chosen in Amsterdam in the first round, whereas the right panel shows the first round choices in Vancouver. In Vancouver subject seem to more often choose the safe bank, but the general pattern is the same: a substantial fraction of people chooses to wait in the first round in both locations. Looking at treatment effects, we can see that those are different in the two locations. Table D.2 shows the multinomial regression results for both locations in the first round. In Amsterdam the only treatment effect occurs for the action 'Withdraw in the risky bank': for both riskiness level it is more likely than the Nash equilibrium action. This highly significant effect is likely due to the fact that under low risk nobody chose this action in the first round. In Vancouver we only observe an effect of high risk, but not of medium risk: subjects in the high risk treatment are more likely to choose to withdraw from the safe bank compared to the Nash equilibrium action. Even though we find these differences in the first round decisions, the learning in the end is not different across location, and the treatment effects considering the whole experiment are similar in directions and significance.

E Regression results

Table E.1 shows multinomial regression results without controlling for path dependency. Table E.2 shows different specifications for the multinomial logit regressions with always controlling for path dependency over the course of the experiment. Table E.3 restricts the analysis for subsamples based on riskiness. Table E.4 shows robustness of the results of the first round multinomial logit regressions.

Table D.1: Multinomial panel logit regressions − all data per location

Panel A: Amsterdam			
Action:			
(base: Risky/Wait)	Safe/Wait	Safe/Withdraw	Risky/Withdraw
Withdrawal fraction in chosen bank$_{t-1}$	3.034***	4.434***	2.960***
	(0.672)	(0.590)	(0.622)
Full information	−0.586	−0.770	−0.345
	(0.347)	(0.400)	(0.351)
Medium risk	0.600	0.819	1.029
	(0.398)	(0.669)	(0.606)
High risk	0.663	4.179***	4.797***
	(0.411)	(0.565)	(0.689)
Round	−0.060**	−0.006	0.001
	(0.018)	(0.009)	(0.009)
Constant	−2.713***	−6.126***	−6.883***
	(0.445)	(0.599)	(0.679)
# of observations	13,984		
# of panels	286		
Log Likelihood	−6,484.075		
Panel B: Vancouver			
Action:			
(base: Risky/Wait)	Safe/Wait	Safe/Withdraw	Risky/Withdraw
Withdrawal fraction in chosen bank$_{t-1}$	2.023***	3.558***	2.488***
	(0.450)	(0.343)	(0.284)
Full information	0.334	−0.860	−0.272
	(0.503)	(0.494)	(0.445)
Medium risk	−1.466*	−0.396	−0.321
	(0.655)	(0.886)	(0.687)
High risk	−1.154	2.761***	2.427***
	(0.730)	(0.424)	(0.392)
Round	−0.040***	−0.011	−0.015*
	(0.010)	(0.006)	(0.0065)
Constant	−0.323	−3.805***	−3.776***
	(0.642)	(0.504)	(0.402)
# of observations	13,808		
# of panels	285		
Log Likelihood	−9,882.398		

Notes: ***: significant on 0.1%-level,**: significant on 1%-level, *: significant on 5%-level. The columns show the estimated multinomial logit coefficients with individuals as panels. The reference level is the Nash-equilibrium action: choosing the risky bank and wait. Full information is 1 for the full info treatments, 0 otherwise. Medium risk is 1 for the medium risk treatments, 0 otherwise. High risk is 1 for the high risk treatments, 0 otherwise. Standard errors are clustered on the session level, and are in brackets. Panel A uses data from Amsterdam, whereas Panel B uses data from Vancouver.

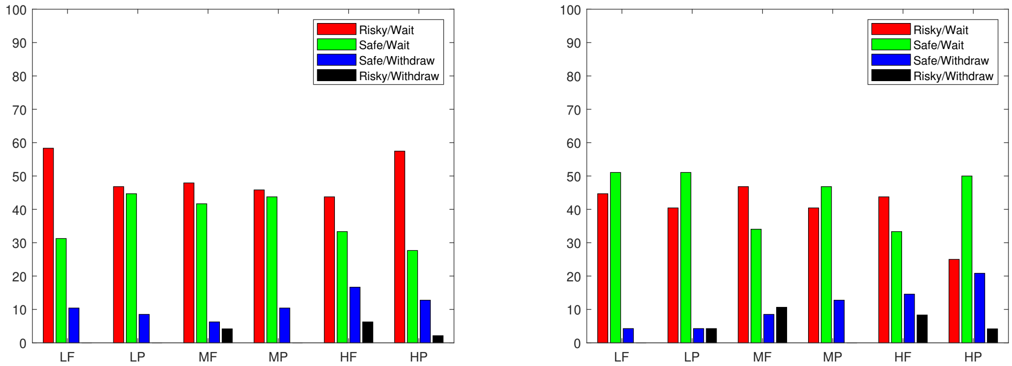

Figure D.3: Percentages of strategies per location and treatment in Amsterdam (left panel) and in Vancouver (right panel) in the first round

Table D.2: Multinomial logit regressions – first round per location

Panel A: Amsterdam			
Action:			
(base: Risky/Wait)	Safe/Wait	Safe/Withdraw	Risky/Withdraw
Medium risk	0.235	−0.012	13.377***
	(0.308)	(0.528)	(0.739)
High risk	−0.175	0.483	14.006***
	(0.322)	(0.473)	(0.541)
Constant	−0.329	−1.715***	−16.491***
	(0.219)	(0.363)	(0.148)
# of observations	286		
Log Likelihood	−291.492		
Panel B: Vancouver			
Action:			
(base: Risky/Wait)	Safe/Wait	Safe/Withdraw	Risky/Withdraw
Medium risk	−0.258	0.892	0.892
	(0.311)	(0.633)	(0.867)
High risk	0.010	1.639**	1.291
	(0.319)	(0.604)	(0.851)
Constant	0.182	−2.303***	−2.996***
	(0.214)	(0.525)	(0.726)
# of observations	284		
Log Likelihood	−308.600		

Notes: ***: significant on 0.1%-level, **: significant on 1%-level, *: significant on 5%-level. The columns show the estimated multinomial logit coefficients. The reference level is the Nash-equilibrium action: choosing the risky bank and wait. Medium risk is 1 for the medium risk treatments, 0 otherwise. High risk is 1 for the high risk treatments, 0 otherwise. Robust standard errors are in brackets. Panel A uses data from Amsterdam, whereas Panel B uses data from Vancouver.

Table E.1: Multinomial panel logit regressions − all data, no control for path dependency

Action: (base: Risky/Wait)	Safe/Wait	Safe/Withdraw	Risky/Withdraw
Full information	−0.269	−1.212**	−0.452
	(0.370)	(0.411)	(0.420)
Medium risk	−0.469	0.295	0.344
	(0.498)	(0.568)	(0.608)
High risk	0.582	5.602***	4.763***
	(0.465)	(0.452)	(0.463)
Round	−0.048***	0.006	0.001
	(0.009)	(0.008)	(0.008)
Constant	−1.026*	−5.223***	−5.330***
	(0.485)	(0.478)	(0.496)
# of observations	28,399		
# of panels	571		
Log Likelihood	−17,986.865		

Notes: ***: significant on 0.1%-level,**: significant on 1%-level, *: significant on 5%-level. The columns show the estimated multinomial logit coefficients with individuals as panels. The reference level is the Nash-equilibrium action: choosing the risky bank and wait. Full information is 1 for the full info treatments, 0 otherwise. Medium risk is 1 for the medium risk treatments, 0 otherwise. High risk is 1 for the high risk treatments, 0 otherwise. Standard errors are clustered on the session level, and are in brackets.

Table E.2: Multinomial panel logit regressions for all decisions - robustness checks

Action: (base: (Risky/Wait)	Safe / Wait	Safe / Withdraw	Risky / Withdraw	Safe / Wait	Safe / Withdraw	Risky / Withdraw	Safe / Wait	Safe / Withdraw	Risky / Withdraw
Withdrawal fraction	2.443***	3.994***	2.751***	2.396***	3.985***	2.751***	2.404***	3.981***	2.744***
in chosen bank$_{t-1}$	(0.407)	(0.303)	(0.304)	(0.394)	(0.309)	(0.302)	(0.397)	(0.307)	(0.300)
Full information	−0.950	−0.453	0.395	−0.106	−0.795*	−0.250	−0.932	−0.452	0.364
	(0.841)	(0.724)	(0.733)	(0.339)	(0.330)	(0.307)	(0.620)	(0.640)	(0.585)
Medium risk	−1.422	0.454	0.536	−0.555	0.069	0.131	−1.311	0.536	0.582
	(0.746)	(0.951)	(0.916)	(0.420)	(0.571)	(0.474)	(0.689)	(0.923)	(0.832)
High risk	−0.943	3.398***	3.798***	−0.312	3.375***	3.414***	−0.816	3.439***	3.842***
	(0.727)	(0.665)	(0.617)	(0.462)	(0.352)	(0.382)	(0.649)	(0.606)	(0.529)
Full info * medium	1.599	−1.043	−0.861				1.506	−1.103	−0.855
	(1.000)	(1.146)	(1.099)				(0.816)	(1.044)	(0.931)
Full info * high	1.107	−0.175	−0.918				1.000	−0.129	−0.787
	(0.990)	(0.780)	(0.804)				(0.812)	(0.719)	(0.702)
Round	−0.045***	−0.008	−0.007	−0.045***	−0.009	−0.008	−0.045***	−0.009	−0.007
	(0.008)	(0.005)	(0.005)	(0.008)	(0.005)	(0.006)	(0.008)	(0.005)	(0.006)
Age				0.029	−0.059	−0.103***	0.027	−0.058	−0.101***
				(0.028)	(0.033)	(0.022)	(0.027)	(0.033)	(0.022)
Male				−0.396*	−0.538**	−0.392*	−0.397*	−0.541**	−0.379*
				(0.189)	(0.174)	(0.158)	(0.186)	(0.178)	(0.163)
Amsterdam				−1.688***	−0.839*	−1.005**	−1.665***	−0.865*	−1.012**
				(0.346)	(0.347)	(0.321)	(0.331)	(0.345)	(0.322)
Constant	−1.030	−4.993***	−5.424***	−1.123	−2.896***	−2.146***	−0.676	−3.060**	−2.520***
	(0.670)	(0.673)	(0.640)	(0.809)	(0.741)	(0.528)	(0.863)	(0.889)	(0.584)
# of observations	27,792			27,716			27,716		
# of panels	571			567			567		
Log Likelihood	−16,482.25			−16,344.547			−16,335.198		

Notes: ***: significant on 0.1%-level, **: significant on 1%-level, *: significant on 5%-level. The columns show the estimated multinomial logit coefficients with individuals as panels. The reference level is the Nash-equilibrium action: choosing the risky bank and wait. Full information is 1 for the full info treatments, 0 otherwise. Medium risk is 1 for the medium risk treatments, 0 otherwise. High risk is 1 for the high risk treatments, 0 otherwise. Full info * medium and full info * high are interaction terms between the above-mentioned variables. Age is the subject's age in years, Male is 1 for males, and 0 otherwise. Amsterdam is 1 for individuals participating in Amsterdam, 0 otherwise. Standard errors are clustered on the session level, and are in brackets.

Table E.3: Multinomial logit regressions for the risk treatments separately

Treatment:	Low risk			Medium risk			High risk		
Action: (base: (Risky/Wait)	Safe / Wait	Safe / Withdraw	Risky / Withdraw	Safe / Wait	Safe / Withdraw	Risky / Withdraw	Safe / Wait	Safe / Withdraw	Risky / Withdraw
Withdrawal fraction	4.430***	5.493***	4.547**	4.302***	5.347***	4.383***	1.130**	3.387***	2.093***
in chosen bank$_{t-1}$	(1.236)	(1.083)	(1.323)	(0.784)	(0.778)	(0.791)	(0.388)	(0.264)	(0.236)
Full information	−1.022	−0.400	0.470	0.788	−1.542	−0.332	0.064	−0.580	−0.536
	(0.836)	(0.573)	(0.732)	(0.554)	(0.785)	(0.746)	(0.439)	(0.337)	(0.365)
Round	−0.053**	−0.056***	−0.036*	−0.047**	−0.021	−0.020	−0.028***	−0.006	0.007
	(0.017)	(0.016)	(0.016)	(0.017)	(0.016)	(0.011)	(0.007)	(0.005)	(0.006)
Constant	−1.094	−5.317***	−5.405***	−2.648***	−4.913***	−5.177***	−1.554***	−1.684***	−1.670***
	(0.691)	(0.684)	(0.782)	(0.437)	(0.749)	(0.766)	(0.315)	(0.231)	(0.231)
# of observations	9,301			9,296			9,195		
# of panels	190			190			191		
Log Likelihood	−3,739.092			−3,709.315			−8,861.004		

Notes: ***: significant on 0.1%-level, **: significant on 1%-level, *: significant on 5%-level. The columns show the estimated multinomial logit coefficients with individuals as panels. The reference level is the Nash-equilibrium action: choosing the risky bank and wait. Full information is 1 for the full info treatments, 0 otherwise. Standard errors are clustered on the session level, and are in brackets.

Table E.4: Logit regressions for the first round decisions

Action: (base: (Risky/Wait)	Safe / Wait	Safe / Withdraw	Risky / Withdraw	Safe / Wait	Safe / Withdraw	Risky / Withdraw	Safe / Wait	Safe / Withdraw	Risky / Withdraw
Full information	−0.273	−0.241	0.921				−0.262	−0.195	0.933
	(0.182)	(0.286)	(0.542)				(0.184)	(0.287)	(0.545)
Medium risk	−0.018	0.367	1.305	−-0.034	0.382	1.300	−0.037	0.380	1.314
	(0.218)	(0.394)	(0.824)	(0.220)	(0.398)	(0.817)	(0.220)	(0.398)	(0.828)
High risk	−0.094	0.972**	1.722*	−0.147	0.968**	1.726*	−0.149	0.965**	1.740*
	(0.224)	(0.365)	(0.799)	(0.227)	(0.371)	(0.792)	(0.228)	(0.372)	(0.807)
Age				0.046	−0.025	0.016	0.047	−0.024	0.005
				(0.031)	(0.043)	(0.044)	(0.031)	(0.042)	(0.045)
Male				−0.325	0.076	0.044	−0.324	0.077	0.055
				(0.186)	(0.288)	(0.485)	(0.187)	(0.288)	(0.485)
Amsterdam				−0.415*	−0.201	−1.023*	−0.417*	−0.200	−1.025*
				(0.184)	(0.287)	(0.513)	(0.184)	(0.287)	(0.510)
Constant	0.0676	−1.813***	−4.400***	−0.740	−1.310	−3.729**	−0.635	−1.234	−4.103**
	(0.177)	(0.321)	(0.994)	(0.708)	(0.967)	(1.159)	(0.702)	(0.973)	(1.292)
Observations	570			566			566		
Log Likelihood	−603.917			−594.474			−591.157		

Notes: ***: significant on 0.1%-level, **: significant on 1%-level, *: significant on 5%-level. The columns show the estimated multinomial logit coefficients. The reference level is the Nash-equilibrium action: choosing the risky bank and wait. Full information is 1 for the full info treatments, 0 otherwise. Medium risk is 1 for the medium risk treatments, 0 otherwise. High risk is 1 for the high risk treatments, 0 otherwise. Age is the subject's age in years, Male is 1 for males, and 0 otherwise. Amsterdam is 1 for individuals participating in Amsterdam, 0 otherwise. Heteroskedasticity-robust standard errors are in brackets.