Disruptions in Large-Value Payment Systems: An Experimental Approach*

Klaus Abbink,[a] Ronald Bosman,[b] Ronald Heijmans,[c]
and Frans van Winden[d]
[a]Monash University
[b]VU University
[c]De Nederlandsche Bank
[d]CREED, University of Amsterdam

This experimental study investigates the individual behavior of banks in a large-value payment system. More specifically, we look at (i) the reactions of banks to disruptions in the payment system, (ii) the way in which the history of disruptions affects the behavior of banks (path dependency), and (iii) the effect of more concentration in the payment system (heterogeneous market versus a homogeneous market). The game used in this experiment is a stylized version of a model of Bech and Garratt (2006) in which each bank can choose between paying in the morning (efficient) or in the afternoon (inefficient). The results show that there is significant path dependency in terms of disruption history. Also, the chance of disruption influences the behavior of the participants. Once the system is moving towards the inefficient equilibrium, it does not easily move back to the efficient one. Furthermore, there is a clear leadership effect in the heterogeneous market.

JEL Codes: C92, D70, D78, E58.

1. Introduction

One of the most significant events in the credit crisis of 2008 was that interbank markets became highly stressed. Liquidity in those markets dried up almost completely because banks suddenly became highly uncertain about each other's creditworthiness. In order to

*Author e-mails: klaus.abbink@monash.edu, r.a.j.bosman@vu.nl, ronald.heijmans@dnb.nl, f.a.a.m.vanwinden@uva.nl.

prevent a collapse of the financial system, central banks intervened by injecting massive volumes of liquidity into the financial system. Our paper relates to stress situations in a particular segment of the financial system, namely large-value payment systems (LVPSs), in which banks pay each other large sums of money during the day.[1] Although during the credit crisis such payment systems were in general functioning properly, disruptions can potentially jeopardize the stability of the financial system as a whole.

The terrorist attacks on the World Trade Center (WTC) in 2001 showed that financial systems are vulnerable to wide-scale disruptions of payment systems. The physical damage to property and communication systems made it difficult or even impossible for some banks to execute payments. The impact of the disruption was not limited to the banks that were directly affected. As a result of fewer incoming payments, other banks became reluctant or in some cases even unable to execute payments themselves (due to insufficient available liquidity). As this could have undermined the stability of the financial system as a whole, the Federal Reserve intervened by providing liquidity through the discount window and open market operations.

Because wide-scale disruptions such as in 2001 do not occur very often, there is not much empirical evidence on how financial institutions behave under extreme stress in payment systems. Research has

[1]Historically, the settlement of interbank payments was done through a netting system in which the payments are settled on a net basis once or several times during the settlement day. With the increase of both the number of transactions and the value of these transactions, the settlement risk increased as well. Banks were increasingly concerned about contagion effects in case of unwinding if one participant would not be able to fulfill its obligation at the end of a netting period. To eliminate this settlement risk, central banks typically developed payment systems in which payments are executed at an individual gross basis, so-called real-time gross settlement (RTGS) systems. Payments are settled irrevocably and with finality. The drawback of RTGS systems is that they require more liquidity because payments usually are not synchronized. To smooth the intraday payment flows, central banks provide intraday credit to their banks. This intraday credit is either collateralized (this holds for most countries including European countries) or priced (United States). An example of a large-value payment system is TARGET2, the euro interbank payment system of the Eurosystem which settled daily in 2008 on average EUR 3.126 billion in value with a volume of 348,000 transactions. Over the years both the value and volume have increased significantly.

therefore focused on simulation techniques. For instance, Soramäki, Bech, and Arnold (2007) and Pröpper, van Lelyveld, and Heijmans (2013) investigated interbank payment systems from a network perspective. To study the impact on the (intraday) dynamics of payment flows, in case a bank changes its own behavior or in case of technical problems, simulations are often used. A comprehensive summary of large-value payment system simulations can be found in Heijmans and Heuver (2012). Many of these simulations focus on the question of what the impact is in terms of liquidity, for the other banks in the payment system, if a single participant delays (intentionally or due to technical problems) payments. This changed behavior can be introduced by modifying historical data in order to reflect the stress scenario under investigation, which requires assumptions on the behavior. Another method that gains momentum in payment systems is agent-based modeling. Our paper is linked to this literature.

Arciero et al. (2009), e.g., present an agent-based model of a real-time gross settlement payment system. Their focus is to provide a methodological contribution, showing how the chosen modeling approach can be employed alongside other tools to study the generation and propagation of systemic risk in payment systems. The model studied by Galbiati and Soramäki (2011) looks at how much liquidity should be posted in a collateralized RTGS system. In contrast to our paper, they do not look at disruptions in payment systems. The approach of our paper is to study disruptions in payment systems in an experimental setting. Basically, we let subjects play a stylized game in a laboratory—which is taken to represent a real LVPS in an essential way—and study interventions in this game. Although subjects in our experiment are students rather than professional managers of LVPSs, there is evidence that students do not behave much different from professionals in economic experiments (e.g., Potters and Van Winden 2000 or Fréchette 2015).[2]

[2]A related issue is that subjects in experiments are paid a relatively modest amount of money compared with professionals in the real world. Note, however, that we paid student subjects on average their opportunity costs of time. In this sense the incentives for the subjects are comparable to those outside the laboratory. Furthermore, there is experimental evidence that increasing the financial stakes does not change behavior much in market and coordination experiments, although in some cases the variance is somewhat reduced (see Jenkins et al. 1998; Kocher, Martinsson, and Visser 2008).

We therefore believe that our results are relevant for understanding real LVPSs. Critics of laboratory experimentation in economics argue that people's behavior in the lab is specific to the experimental situation and may be unconnected to their behavior in the field. They therefore question the external validity of the experimental results. Harrison and List (2004) discuss many of these criticisms and provide a survey of their validity. They plead for conducting experiments both in the lab and in the field. In his response to critics, Camerer (2011) reviews a number of lab experiments in various research areas (among which are experiments on pricing and risk-taking) and concludes that overall they generalize well between lab and field. We have reason to assume therefore that our experiment is relevant for understanding real LVPSs.

An advantage of an experiment is that disruptions can be carefully controlled by the experimenter while the behavioral reactions to these disruptions are determined endogenously (in contrast to simulations where such reactions are assumed). To the best of our knowledge, large-value payment systems have not been studied in the laboratory before, which makes this experiment unique in its kind. McAndrews and Rajan (2000), McAndrews and Potter (2002), and Bech and Garratt (2003) argue that banks' decisions in the U.S. payment system Fedwire can essentially be interpreted as a coordination game. As a vehicle of research we therefore use a stylized game theoretical model developed by Bech and Garratt (2006). In most payment systems participants can execute payments throughout the whole business day. In this model, however, a player has to choose either to pay in the morning, which is considered efficient, or to pay in the afternoon, which is inefficient.[3] This game has two equilibria, one of which is efficient.

These equilibria can be interpreted in a real-life LVPS as follows. Central banks operating an LVPS monitor the functioning of this LVPS closely and encourage banks to pay early (as soon as the obligation is there). One measure central banks have taken to encourage banks to make payments early is to provide collateralized overdrafts. Banks can obtain credit from the central bank without having to pay for this credit. By the end of the day this credit has

[3]This coordination game is known as the stag-hunt game (see also Bech and Garratt 2003).

to be repaid or otherwise the banks do have to pay interest. Furthermore, central banks monitor queues in the LVPS. A payment instruction gets into the queue when the bank has insufficient liquidity available. The payment will be settled when that bank receives sufficient incoming payments or it provides more collateral to be able to receive more (intraday) credit. If a bank systematically delays payments, central banks (may) use moral suasion to “encourage” earlier (timely) payment. In the United Kingdom banks have to live up to so-called throughput guidelines in which banks on average over the month have to settle 50 percent of the value of payments by noon and 75 percent by 2:30 p.m. (see Ball et al. 2011).[4] The reason why central banks monitor their LVPS so closely is that they are afraid that the LVPS will get into a so-called gridlock, i.e., a situation where several payments each await settlement of the others. This might have happened during the attacks on the WTC if the Federal Reserve did not intervene. Given the importance of intraday liquidity, the Bank for International Settlements describes monitoring tools for intraday liquidity management (Basel Committee on Banking Supervision 2013).

The relevance of delaying payments—often called free-riding and in this experiment “paying in the afternoon,” is also addressed by Diehl (2013). He looks at measures of individual banks’ free-riding behavior in the German part of the European large-value payment system. The difficulty with measuring intentional delay from the payment system data is that only the time the payment was sent to the LVPS and when it was settled is known. This is not necessarily the time the payment had to be paid. However, sudden changes (delays) in payment systems are not likely. Diehl (2013) does not find evidence of free-riding behavior.

Heijmans and Heuver (2014) study the behavior of banks in the payments system before and during the credit crisis which started in the summer of 2007. They find that the timing of payments is an important indicator for identifying potential liquidity problems. They show that some banks indeed delay payments. Besides, they find that some banks limit their exposure to a counterparty by

---

[4]As described in Becher, Galbiati, and Tudela (2008), enforcement of the throughput guidelines in CHAPS (Clearing House Automated Payment System) currently relies on peer pressure rather than financial or regulatory sanctions.

setting bilateral limits.[5] These limits were set tighter by some of these banks in case there were clear signals about the deteriorating financial soundness of a certain counterparty or even rumors. Setting such limits is the action of a single bank, based on its own risk perception of the market, and not a coordinated action of all the banks simultaneously. This is in contrast to the LIBOR scandal, which could be seen as a more coordinated action of banks. This scandal arose when it was discovered that banks intentionally manipulated their rates either to profit from trades or to disguise their deteriorating creditworthiness. The loans with manipulated rates are usually settled in LVPSs. However, the number of bank transactions related to money market deals is very limited.[6,7] Besides, the timing of the payment of these "manipulated" loans does not have to be affected in any way by this scandal. Therefore, we leave coordinated action out of the scope of this experiment and focus on the individual decisions of banks.[8]

Our study is closely related to the experimental literature on coordination games. Pure coordination games involve multiple equilibria with the same payoff consequences, provided all players choose the same action. The players' task is to take cues from the environment to identify focal points (Schelling 1960; Mehta, Starmer, and Sugden 1994). More akin to our problem are studies on games with Pareto-ranked equilibria. In these games one equilibrium yields higher payoffs to all players than the other equilibriums yield, such that rational players should select it (Harsanyi and Selten 1988). However, experimental subjects often coordinate on inferior equilibria, in particular when the Pareto-dominant equilibrium is

[5]These limits are the maximum debit position a bank is willing to have to a certain counterparty. Bilateral discussion with commercial banks shows that some banks indeed have sophisticated systems available to set and change these bilateral limits. Banks have the opportunity to set such limits in the European large-value payment system, TARGET2. However, most banks choose to set these limits in their internal systems.

[6]Armantier and Copeland (2012) and Arciero et al. (2016) shed light on the euro area and the American unsecured money market during the crisis, respectively.

[7]Besides money market deals, banks settle many other obligations, such as those on behalf of their own business (not money market related) and their clients.

[8]We have excluded communication between players as a first approach. For our study it is important that banks do not know whether the delay of another bank is due to a technical disruption or strategic behavior.

risky (Van Huyck, Battalio, and Beil 1990, 1991), as is the case in our vehicle of research, or other equilibria are more salient (Abbink and Brandts 2008). For an overview of coordination game experiments, see Devetag and Ortmann (2007). None of the existing studies tackles the problem of random disruptions.

Our main research question is how behavior in the payment system is affected by different random disruptions. We define a disruption as a situation where one or more players are unable to execute a payment in a timely manner—for example, because of an individual technical failure or (temporary) financial problems. In addition, we investigate whether concentration in the interbank payment system—in the sense that players are heterogeneous in terms of their size and, consequently, in the risk they may impose on other players—matters. From an economic point this is relevant because consolidation in the financial sector has led to the emergence of a few very large financial institutions.[9] Real payment systems are often characterized by a few large banks and many smaller ones, which make them look like a heterogeneous market.[10] However, the core of the payment system, comprising large banks that together often have a market share of more than 75 percent, looks more like a homogeneous market.[11] There are also LVPSs that are highly tiered. Tiered systems are characterized by a relatively small number of banks participating in the LVPS. These banks also settle payments on behalf of many other banks. Such a tiered system looks similar to a homogeneous market.[12] This means that payment systems can have characteristics of both types of markets. Finally, this paper investigates whether there is any path dependency, taking into account the history of disruption.

[9]The credit crisis has even enhanced this consolidation process. In the United States, for example, investment banks have typically merged with commercial banks. In general, there is a tendency for weaker banks to be taken over by stronger (larger) banks.

[10]The large-value payment system of the Eurosystem (TARGET2) and the United States (Fedwire) are good examples of heterogeneous markets.

[11]For example, in the Dutch part of the European large-value payment system TARGET2, which consists of fifty credit institutions, the five largest banks account for 79 percent of the total value of outgoing daily payments. The thirty-eight smallest ones only cover 5 percent of this value.

[12]The large-value payment system of the United Kingdom (CHAPS) is a good example of a highly tiered system.

The organization of this paper is as follows. Section 2 describes the experimental design (including the game theoretical model), the procedures used, and the predictions. Section 3 discusses the results. Section 4 goes into some policy issues and provides a conclusion.

2. Experimental Design and Procedures

2.1 Design

Our design is based on a model by Bech and Garratt (2006), which is an n-player liquidity management game. The game envisions an economy with n identical banks, which use a real-time gross settlement system operated by the central bank to settle payments and securities. Banks intend to minimize settlement cost. In this game the business day consists of two periods in which banks can make payments: morning or afternoon. At the beginning of the day banks have a zero balance on their accounts at the central bank. At the start of each business day each bank has a request from customers to pay a customer of each of the other (n – 1) banks an amount of Q as soon as possible. To simplify the model, the bank either processes all n – 1 payments in the morning or in the afternoon. If a bank does not have sufficient funds to execute a payment, it can obtain intraday credit, which is costly and reflected by a fee F. This fee can be avoided by banks by delaying their payments to the afternoon. With this delay, however, there are some social and private costs involved, indicated by D. For example, a delay may displease customers or counterparties, costing banks in terms of potential claims and reputation risk. Also, in the case of operational disruptions, payments might not be settled by the end of the business day. This disruption can either be a failure at the payment system to operate appropriately or a failure at the bank itself. The costs in this case can, for example, be claims as a result of unsettled obligations or loss of reputation. The trade-off between the cost F of paying in the morning and cost D of paying in the afternoon is made by each bank individually. Bech and Garratt (2006) investigate the strategic adjustment banks make in response to temporary disruptions. In particular, they focus on equilibrium selection after the disruption is over.

In our experiment we use a simple version of the theoretical model by Bech and Garratt. Because $F \geq D$, there are two equilibria in pure strategies, assuming each bank maximizes its own earnings. Either all banks pay in the morning or all banks pay in the afternoon. The morning equilibrium is the efficient equilibrium.[13] In each of the several rounds of the experiment, the banks have to make a choice between paying in the morning (labeled choice X) and paying in the afternoon (labeled choice Y). In each round, furthermore, there is a known probability p that a bank is forced to pay in the afternoon.[14] This means that the bank cannot pay in the morning but is forced to delay payment to the afternoon. The other banks observe that there was a delay at this bank, but they do not know whether it was caused by a disruption (a forced Y) or a deliberate decision. The probability of disruption is the core parameter of the experiment. After each round, all banks see the choice of the other banks. However, it is not known by the other banks whether a bank was forced to pay in the afternoon or chose to do so intentionally.

The experiment consists of three parts, each consisting of thirty rounds.[15] The probability p varies between the three parts. Instructions for each part were only provided when the respective part began. From a game-theoretic point of view, our design is a finitely repeated version of the (deterministic) game considered by Bech and Garratt. In our setup this should not change the equilibria. Repetition causes additional room for equilibria in games that have Pareto-improving non-equilibrium outcomes in the stage game (Benoit and Krishna 1985). Players can then use equilibrium selection towards the end as a means to enforce more profitable play in the supergame. However, this is not the case in our experiment, as there is no

[13] See proposition 1 of Bech and Garratt (2006).

[14] The inability of many banks during the attacks at the WTC can be interpreted as the forced delayed payment.

[15] A round in the experiment can be interpreted as a business day. Having thirty rounds in one block is a period of constant probability that banks cannot pay early. A (steep) increase of this probability could be seen as a period, e.g., during the attacks at the WTC, in which it took some banks several days to solve the technical difficulties they faced. In a pilot we also investigated a disruption probability of 0 percent. This leads to X choices only. An increase to a 15 percent probability of disruption in the second block of thirty rounds also leads to X choices only—provided that a player has a choice. The pilot showed that there will be consistent coordination on X when the disruption probability is 0 percent.

Figure 1. Two Types of Markets

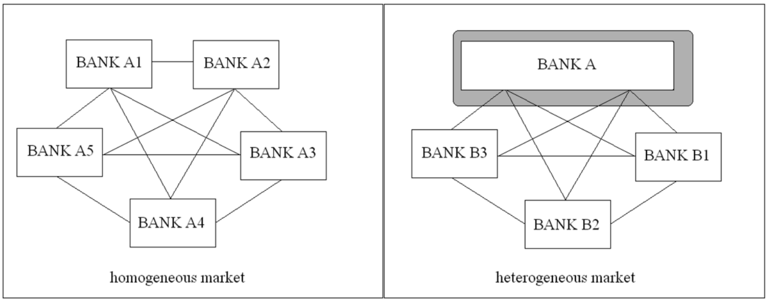

Table 1. Overview of Experimental Treatments

Treatment Name	Type of Market	Disruption Probability p			Number of Groups
		Part 1	Part 2	Part 3	
HOM_15-30-15	Homogeneous	15%	30%	15%	16
HOM_30-15-30	Homogeneous	30%	15%	30%	16
HOM_15-45-15	Homogeneous	15%	45%	15%	15
HOM_45-15-45	Homogeneous	45%	15%	45%	15
HET_15-30-15	Heterogeneous	15%	30%	15%	17
HET_30-15-30	Heterogeneous	30%	15%	30%	14

strategy combination that could improve upon the outcome in the efficient equilibrium.

The experiment investigates the impact of the disruption probability in two types of markets: a homogeneous market and a heterogeneous market. The homogeneous market represents a market in which all banks are identical both in size and impact (n = 5). The heterogeneous market case, on the other hand, constitutes a market in which one bank is twice as large as the other banks, thus making and receiving twice as many payments (n = 4). Conceptually, one can see the heterogeneous market as the homogeneous market where two identical (small) banks have merged; see figure 1. Table 1 provides an overview of the different treatments investigated in the experiment. Instructions and computer screenshots are presented in the appendix.

Table 2. Earnings Table of Homogeneous Market (in experimental currency)

Number of Other Players Choosing X	Number of Other Players Choosing Y	Your Earnings from Choosing X	Your Earnings from Choosing Y
4	0	5	2
3	1	3	2
2	2	1	2
1	3	−1	2
0	4	−3	2

Table 2 shows the earnings in the case of a homogeneous market with five identical banks, where X stands for paying in the morning and Y for paying in the afternoon. Earnings are determined by a fixed payoff of 5, while F = 2 and D = 3/4.[16]

2.2 Procedures

The experiment was run with undergraduate students of the University of Amsterdam using the CREED laboratory. Upon arrival, participants were randomly seated in the laboratory. Subsequently, the instructions for the experiment were given. Students could only participate in the experiment once.

The computerized experiment was set up in an abstract way, avoiding suggestive terms like banks. Choices were simply labeled X and Y. Forced choices were indicated by Y_f on the computer screen of participants. Participants were randomly divided into groups whose composition did not change during the experiment. Participants were labeled A1 to A5 in the homogeneous market and A, B1, B2,

[16] In the case of paying in the afternoon, earnings equal $-(n - 1) \cdot D + 5$, with n being the total number of banks. Earnings if the bank instead chooses paying in the morning equal $-(n - 1 - |Si|m) \cdot F + 5$, where $|Si|m$ denotes the number of other banks paying in the morning. The heterogeneous market case follows straightforwardly and is therefore left out, to save space (see the instructions in the appendix for details). We only notice here that, due to the concentration, the banks run a differential risk in the heterogeneous case. For example, for a small bank only two Y choosers (instead of three) can make the choice of Y optimal now.

and B3 in the heterogeneous market.[17] Note that in the latter market A refers to the large bank (see figure 1). Whether a participant represented a large or a small bank was determined randomly. All payoffs were in experimental talers, which at the end of the experiment were converted into euros at a fixed exchange rate known to the participants. Each experiment took approximately one hour and the average earnings were EUR 18.82, including a show-up fee of EUR 5. In total, 434 students participated in the experiment.

2.3 Predictions

The experimental game has two equilibria in pure strategies when the probability of disruption is "low" (15 percent) or "intermediate" (30 percent). In the first equilibrium, all banks pay in the morning. In the second equilibrium, all banks defer their payment to the afternoon. Note that the first equilibrium is efficient. In this equilibrium all banks are better off than in the second equilibrium. So, one would expect that banks would try to coordinate on this equilibrium. The efficient equilibrium, however, is risky in the sense that paying in the morning is costly when two or more banks decide to defer their payment to the afternoon. Whether or not banks will coordinate on the efficient equilibrium depends, among other things, on their risk attitude. Experimental research shows that in coordination games where the efficient equilibrium is risk-dominated by other equilibria, the efficient equilibrium need not be the obvious outcome (e.g., van Huyck, Battalio, and Beil 1990). When the chance of disruption is "high" (45 percent), there is only one equilibrium, where all banks pay late. In this situation the obvious prediction is that banks coordinate on this equilibrium.[18]

[17]We apply the "unitary player assumption" here, which is common in theoretical and experimental economics. The experimental evidence of a behavioral equivalence of group behavior and individual behavior is very mixed, also regarding markets (see, e.g., Cox and Hayne 2006; Bornstein et al. 2008; Raab and Schipper 2009; Ambrus, Greiner, and Pathak 2013). Bosman, Hennig-Schmidt, and van Winden (2006) discuss eight factors that may account for the mixed evidence. As a first approach, it seems justified therefore to make this assumption.

[18]The finitely repeated version of our (Bech and Garratt) game makes multiple equilibria possible because of the presence of two equilibria—an efficient one (all pay in the morning) and an inefficient one (all pay in the afternoon); see Benoit

The homogeneous and heterogeneous markets in fact have the same two equilibria. From a standard game theoretical point of view, we would expect the same outcome in both markets. From a behavioral point of view it is possible that the outcomes differ. In the heterogeneous market, for example, the large bank may have a disproportionate influence on the behavior of others. Whether such an influence is helpful or harmful in terms of coordinating on the efficient equilibrium is difficult to say a priori, and the experiment will shed more light on such behavioral issues. Finally, we investigate whether there is any path dependency and how this relates to the probability of the disruptions.

3. Results

This section describes the results of the different experimental treatments. We look at plain-choice frequencies and a measure that captures the degree of coordination, called "full coordination" (the situation where participants make the same choice, given that a participant is not forced to "choose" Y). Section 3.1 describes the results for the homogeneous market and section 3.2 describes those for the heterogeneous market.

3.1 Homogeneous Market

3.1.1 Choice Frequencies

We take a first shot at the data by simply looking at the choice frequencies of the four homogeneous market treatments, as depicted in figure 1. In figure 2, the HOM_15-30-15 treatment (top left) shows that the choice frequency of X in parts 1 and 3, both with 15 percent disruption probability, does not change much throughout each block of rounds. However, the choice frequency of X in block 3 is higher than in block 1. In fact, intentionally chosen Y in block 3 almost vanishes. These observations suggest that participants learn

and Krishna (1985). Any string of efficient/morning and inefficient/afternoon outcomes can be sustained as equilibrium. However, players cannot reach any higher payoff than that obtained in the efficient (morning) outcome, which makes these additional equilibria not particularly interesting, in contrast to what will happen in the experiment.

Figure 2. Choice Frequencies (homogeneous markets)

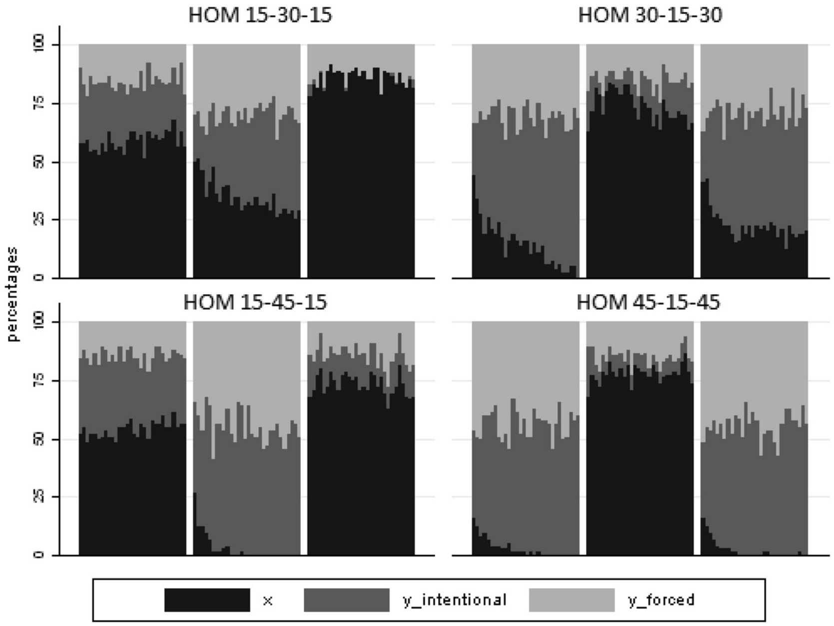

Note: The panels of each graph show the choice frequencies per round for the respective parts of the experimental treatment.

to coordinate on the efficient equilibrium over time. Block 2, with a disruption probability of 30 percent, shows that the choice frequency of X decreases from 50 percent to slightly above 25 percent and the choice frequency of intentional Y increases. The results for the reversed-order treatment, HOM_30-15-30 (top right), show a similar pattern for the 30 percent blocks, but a stronger decrease of choice frequency X within the blocks is observed—making the overall choice frequencies of X when p = 30% lower in the reversed-order treatment. This observation suggests that behavior is not fully independent from past disruption experience.

The bottom two graphs, referring to treatments HOM_15-45-15 and HOM_45-15-45, show that a disruption probability of 45 percent quickly leads to choices Y or Y-forced, as predicted. From this it can be concluded that when the disruption probability becomes too large, there is no incentive to choose X anymore, because this

Table 3. Fraction of Groups Fully Coordinating on X or Y (homogeneous markets)

Treatment	Coordination on X			Coordination on Y		
	Part 1	Part 2	Part 3	Part 1	Part 2	Part 3
HOM_15-30-15	0.56	0.40	0.97	0.19	0.42	0.00
	(0.14)	(0.05)	(0.06)	(0.08)	(0.15)	(0.00)
HOM_30-15-30	0.11	0.76	0.24	0.66	0.08	0.60
	(0.08)	(0.12)	(0.05)	(0.23)	(0.08)	(0.14)
HOM_15-45-15	0.53	0.01	0.80	0.30	0.91	0.11
	(0.14)	(0.04)	(0.06)	(0.08)	(0.14)	(0.04)
HOM_45-15-45	0.01	0.86	0.01	0.86	0.06	0.91
	(0.03)	(0.12)	(0.03)	(0.16)	(0.02)	(0.13)

Note: Standard deviation is shown in parentheses.

will lead to losses for the participants. Comparing the bottom left graph with the top left graph shows that the increasing trend in X choices in going from block 1 to 3 is similar. However, in block 3 of HOM_15-45-15, the increase in X appears less strong than in block 3 of HOM_15-30-15.

3.1.2 Frequency of Full Coordination

Table 3 shows the average fraction of the groups that fully coordinate on X and Y for the four homogeneous market treatments. There is full coordination on X or Y when all of the participants within one group who have a choice (i.e., who are not forced to choose Y) choose X or Y, respectively. There has to be at least one participant who has a choice in order to get full coordination on X or Y. Figure 3 shows the level of coordination on X (black bar) and Y (dark grey bar), and the absence of coordination (light grey bar), for each round of the four treatments. The data show that there is more coordination on X when the disruption probability is low ($p = 15\%$) and more coordination on Y when the disruption probability is intermediate or high (30 percent or 45 percent, respectively) ($p < 0.01$, binomial test for block 1 between treatments). In the context of a payment system, this suggests that larger disruptions are associated with less efficiency.

Figure 3. Percentage of Groups Fully Coordinating on X or Y (homogeneous markets)

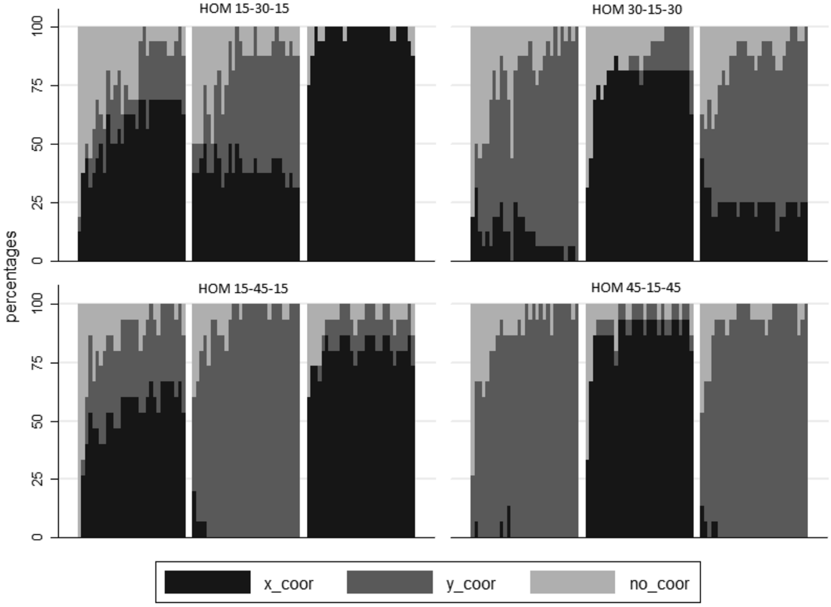

Note: The panels of each graph show the full coordination per round for the respective parts of the experimental treatment.

RESULT 1. *A higher disruption probability leads to less coordination on X and more coordination on Y.*

Both table 3 and figure 3 show that there is more coordination either on X or Y in block 3 compared with block 1. There is significantly more coordination on X in the third block compared with block 1 for the HOM_15-30-15, the HOM_30-15-30, and the HOM_15-45-15 treatments and less coordination on Y (all $p < 0.01$, binomial test). Participants thus learn to coordinate on the efficient equilibrium, which is even speeded up if there is a prior disruption chance of 30 percent or 45 percent. The table and figure also show that for a disruption probability of 45 percent, coordination on X almost vanishes and quickly moves to the inefficient equilibrium. Coordination on X only occurs occasionally in the first few rounds. This is in line with the low-choice frequencies of X in that case, presented in the previous subsection. In the context of a payment system, this means that if the disruption is very likely, there is no incentive anymore to pay as soon as possible. This situation seems similar to the one of the attacks on the WTC in 2001 when many banks, including some large ones, were not able to execute payments due to technical problems. Some banks were reluctant to execute any payment, even though they were able to, because they did not know the impact of the attacks on the stability of the financial system. Understandably, these events threatened to move the payments system to the inefficient equilibrium, which was a reason for the authorities to intervene.

Figure 3 further shows that the developments within blocks are somewhat monotonic, especially the blocks with 15 percent disruption probability. With respect to disruption probability of 30 percent and 45 percent, there is a (bit of a) decreasing trend for X. This suggests that once a trend has been established in the payments system, it is unlikely to reverse. The 15 percent disruption probability of HOM_30-15-30 shows a higher level of coordination on X than block 1 of HOM_15-30-15 but a lower level when compared with block 3. Comparing HOM_15-30-15 with HOM_15-45-15 shows that there is no significant difference in coordination on X in block 1. Block 3 of these two treatments, however, shows some differences, with significantly more coordination on X in HOM_15-30-15 ($p < 0.01$, binomial test). Although the disruption probability is the same, the history of disruption exposure differs between these two treatments. The previous block has either a probability of disruption of 30 percent or 45 percent, leading to different behavior. Block 2 of HOM_15-45-15 shows 91 percent coordination on Y and almost 0 percent coordination on X. For HOM_15-30-15 this is 42 percent coordination on Y and 40 percent on X. This suggests that the disruption history is important for the coordination on both X and Y. In terms of payment systems, this means that the payment behavior of banks depends on history.

RESULT 2. *Overall, there is more coordination in the third part of the experiment than in the first part, given the same disruption probability. If the chance of disruption is low ($p = 0.15$) or intermediate ($p = 0.3$) in the first part, there is more coordination on X in the* *third part. If the disruption probability is high ($p = 0.45$), there is strong coordination on the inefficient equilibrium.*

Result 3. *There is evidence of path dependency, as the outcome depends on the disruption history.*

Confidence between banks is not a static fact, as became clear during the current financial crisis. Banks became reluctant to execute their payments to financial institutions that were "negative in the news." Especially the bankruptcy of Lehman Brothers in October 2008 caused a shock wave of uncertainty through the whole financial system. Banks became aware of the fact that even large (systemically important) banks might not stay in business. The interbank money market, which gives banks with a surplus of liquidity the opportunity to lend money to banks with a temporary shortage, came to a standstill (see, e.g., Armantier and Copeland 2012 and Arciero et al. 2016 for the unsecured money market in the euro area and the United States, respectively). This indicates that recent history is important for the level of confidence banks have in each other, as our experimental result suggests.

3.2 Heterogeneous Market

Recall that in the heterogeneous markets the number of banks is four instead of five. One of the banks is now twice as large in size and impact compared with the other three banks.

3.2.1 Choice Frequencies

Again we take a first shot at the data by looking at plain-choice frequencies in the two heterogeneous markets (see figure 4 and table 3). The left graph of the figure, concerning treatment HET_15-30-15, shows similar trends as in HOM_15-30-15 (figure 2). However, across the different blocks, participants in the heterogeneous markets show a tendency to choose X more often.

3.2.2 Frequency of Full Coordination

Table 4 shows the average fraction of the groups that fully coordinate on X or Y in the two heterogeneous market treatments, while

Figure 4. Choice Frequencies (heterogeneous markets)

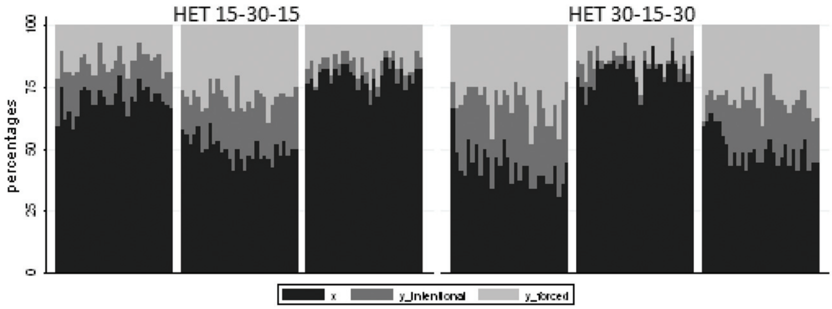

Note: The panels of each graph show the choice frequencies per round for the respective parts of the experimental treatment.

Table 4. Fraction of Groups Fully Coordinating on X or Y (heterogeneous markets)

Treatment	Coordination on X			Coordination on Y		
	Part 1	Part 2	Part 3	Part 1	Part 2	Part 3
HET_15-30-15	0.73	0.64	0.93	0.09	0.24	0.04
	(0.10)	(0.05)	(0.03)	(0.05)	(0.09)	(0.03)
HET_30-15-30	0.44	0.87	0.60	0.23	0.01	0.22
	(0.08)	(0.08)	(0.10)	(0.10)	(0.02)	(0.12)

Note: Standard deviation is shown in parentheses.

figure 5 shows the coordination over the rounds. Comparing the full coordination on both X and Y of the heterogeneous market with the corresponding homogeneous one of section 3.1.2 shows that trends between blocks are similar. However, given the same immediate disruption history, there is significantly more coordination on X in the heterogeneous market treatments than in the homogeneous market in five out of the seven possible cases with the same immediate disruption history (all five cases $p < 0.01$, binomial test).[19] In the two

[19]All blocks of treatment 1 and 2 can be compared with treatment 5 and 6, respectively. The first block of treatment 3 can be compared with the first block of treatment 5. Two cases are not significant. These relate to blocks 2 and 3, given an immediate disruption history of 15 percent ($p = 0.2$ and $p = 0.6$, respectively).

Figure 5. Percentage of Groups Fully Coordinating on X or Y (heterogeneous markets)

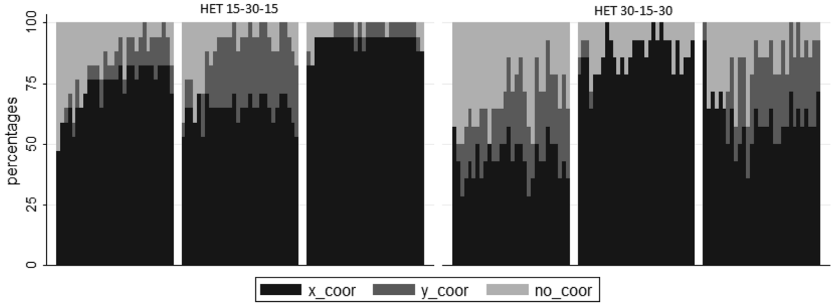

Note: The panels of each graph show the full coordination per round for the respective parts of the experimental treatment.

other cases, there is no significant difference. Note that only blocks that have the same disruption history are compared. These results suggest that coordination is more prominent in a heterogeneous market with asymmetry between participants. A potential explanation is that there is a leadership effect of the large bank, which may feel more responsible than the small banks to choose X because of its relatively large effect on the earnings of all participants. In terms of payment systems, this suggests that a system that consists of one (or perhaps a few) large banks and many small(er) banks will lead to more efficiency than a payment system in which banks are more similar in size.

RESULT 4. *The heterogeneous market leads to more coordination on the efficient equilibrium in most situations characterized by the same immediate disruption history.*

To shed more light on this explanation, we look in more detail at whether the small banks follow the large bank or the other way around in both the 15 percent and 30 percent disruption probability cases. Table 5 shows the reaction of the small banks to the choice of the large bank in previous round(s). The table shows that if the large bank chose X in one or more consecutive rounds, counting from the previous round backwards, there is roughly a 90 percent chance that small banks with a choice (no forced Y) will choose X as well.

Table 5. Leadership of Large Bank—Percentage of Small Banks Following the Large Bank

Choice of Large Bank	Choice of Small Banks = X if Choice of Large Banks Is X	Events	Choice of Small Banks = X if Choice of Large Bank Is Y	Events
Only Once in a Row so Far	89%	1,560	83%	1,544
Only Twice in a Row so Far	92%	1,056	63%	516
Only Three Times in a Row so Far	94%	808	39%	216
Only Four Times in a Row so Far	92%	592	20%	144

Note: "So far" means from the previous round counted backwards.

If the large bank has chosen Y, either intentionally or forcedly, the small banks seem to ignore this when it is only once, as they still choose X 83 percent of the time in that case. Possibly, the small banks will reason that the large bank might have been forced and most likely will choose X in the next round again. The number of small banks choosing X quickly drops if the large bank chooses Y more than once in a row. This can be explained by the fact that two or more forced Ys are not very likely and may be a signal of bad intention rather than bad luck. Note from the payoff table in the appendix that in this situation the payoff for the small banks from choosing X becomes markedly lower, in particular when one other small bank also chooses Y. Large banks in an LVPS do generally have more liquidity available at their account than smaller ones, as a result of the reserve maintenance requirements. This liquidity can intraday be used to make payments. In the case of the European LVPS, TARGET2, many banks have sent in payment instructions before the opening of the system (7:00 a.m.). Banks with sufficient liquidity available will settle all payments immediately, while the ones not having sufficient liquidity will have to wait until sufficient liquidity has been received. However, if these large banks are not able to pay or if they delay payments, this may hamper the smooth functioning of the payment system.

The fact that small banks follow the larger bank is consistent with actual behavior in payment systems, where small banks typically depend on the liquidity of the large bank. For example, it is observed in the Netherlands that large banks have a tendency to start paying large amounts right after opening of the payment system, which corresponds to paying in the morning in terms of our experimental game. The smaller banks usually follow immediately after that. This can still be considered as "paying in the morning" because these payments follow almost instantaneously after the payments of the large banks. This means that the large banks provide liquidity to the small ones, which the latter can use to fulfill their payment obligations.

3.3 Dynamics

Our results show that when the probability of disruptions is moderate, subjects typically achieve a high level of coordination on the efficient equilibrium, while with higher probabilities results are more mixed. We have studied several behavioral rules (heuristics) to explain the observed dynamic behavioral patterns. Most prominently, we investigated the "myopic best response" rule which is an application of the adjustment process that Bech and Garratt propose (section 3 of their paper). However, in addition, we have looked at simply copying the behavior of whoever was most successful ("imitation") as an even simpler and (in the literature on behavioral adjustment) often-considered heuristic, as well as a more general behavioral rule where participants choose X as long as it is profitable. As none of these heuristics turns out to be very successful in capturing our data, we do not include them in this paper and refer to section 4 of Abbink et al. (2010) for a detailed analysis and graphical representation of the outcomes.

4. Conclusions

In this paper we used a stylized coordination game of Bech and Garratt (2006) to experimentally study bank behavior in a large-value payment system that is hindered by disruptions. We draw the following conclusions.

First, once behavior moves in the direction of coordination on the inefficient equilibrium, it is not likely that behavior moves back to the efficient equilibrium. The reason for this is that one player has to take the lead in going for the efficient equilibrium, but this is costly if other players do not follow suit. In the context of a payment system, these findings suggest that once a trend has been established, it is unlikely to reverse. In a situation where some banks begin to defer their payments, an intervention from the central bank may be highly desired or even necessary. When banks do not have access to sufficient liquidity—i.e., they are forced to go for the inefficient equilibrium—central banks can use their discount window to relieve market stress, which happened during the attacks at the WTC. If some (critical) banks deliberately delay payments without having liquidity problems, the central bank can use its authority to encourage banks to start paying earlier (cf. Chaudhuri, Schotter, and Sopher 2009, who study the role of advice in coordination games). Such moral suasion only works though if the payment system has not been disturbed totally (i.e., coordinated fully on the inefficient equilibrium). So, once coordination failures start emerging, central banks need to react quickly; otherwise, trust between banks might have fully vanished and coordination on the "good" equilibrium becomes highly unlikely. Note that in our experimental study there was no role for the central bank. We believe that extending the game by allowing central bank interventions would be an interesting avenue for future experimental work.

Second, coordination on the efficient equilibrium turns out to be easier in a heterogeneous market where there is a clear leader in terms of size. If such a leader goes for the efficient equilibrium, 90 percent of smaller players who have a choice follow the leader. If the leader is not cooperative for several rounds in a row (forcedly or deliberately), the smaller players rapidly move to such a strategy as well. Given the critical role of the large player for the system as a whole, it is essential from a payment system perspective to minimize the chance that large banks are not able to execute payments due to their own technical problems. It may therefore be desirable to oblige such critical participants to take extra safety measures with regard to their technical infrastructure.

Finally, our experiment shows that small frictions in coordination games can be absorbed easily and need not jeopardize the stability of the efficient equilibrium (cf. the 15 percent disruption cases). However, when friction becomes larger, the system can move quickly to the undesired equilibrium and stay there. In the context of payment systems, this suggests that it is very important to closely monitor the payment flows of (critical) participants in the system. If deviant payment behavior is observed by one or more participants, it is important to find the reason for this behavior. If the cause is a technical problem of one participant, the other participants in the payment system should be informed about the incident. In this way it may be avoided that the other participants falsely conclude that the deviant behavior is a deliberate action, for example, related to liquidity considerations. Such communication is especially important during times of market stress, when false rumors can easily arise. Since we did not study communication in our experiment, it is an open research question whether this could work or not.

Appendix

Instructions for the Homogeneous Market Case

The instructions for the homogeneous market case are shown below. Between the different experimental treatments, only the percentages change. The instructions listed here are for the 15%-30%-15% case.

Instructions

Welcome to this experiment. The experiment consists of three parts in which you will have to make decisions. In each part it is possible to earn money. How much you earn depends on your own decisions and on the decisions of other participants in the experiment. At the end of the experiment a show-up fee of 5 euros plus your total earnings during the experiment will be paid to you in cash. Payments are confidential; we will not inform any of the other participants. In the experiment, all earnings will be expressed in talers, which will be converted to euros according to the exchange rate: 1 taler = 6 Eurocents.

During the experiment you will participate in a group of five players. You will be matched with the same players throughout the experiment. These other players in your group will be labeled P2, P3, P4, and P5.You will not be informed of who the other players are, nor will they be informed of your identity.

Talking or communicating with others during the experiment is not permitted. If you have a question, please raise your hand and we will come to your desk to answer it.

Warning: In this experiment you can avoid making any loss (negative earnings). However, note that if you end up with a loss, it will be charged against your show-up fee.

We start now with the instructions for part 1, which have been distributed also on paper. The instructions for the other two parts will be given when they start.

Instructions: Part 1

This part consists of thirty rounds. In each round you and the other four players in your group will have to choose one of two options: X or Y. Your earnings in a round depend on your choice and on the choices of the other four players, in the following manner:

- If you choose Y, your earnings are 2 talers regardless of the choices of the others.
- If you choose X, your earnings depend on how many of the other players choose Y.

Your exact earnings in talers from choosing X or Y, for a given number of other players choosing Y, are listed in the following table. This earnings table is the same for all players.

Number of Other Players Choosing Y	**Your Earnings from Choosing X**	**Your Earnings from Choosing Y**
0	5	2
1	3	2
2	1	2
3	−1	2
4	−3	2

For example, if two other players choose Y, then your earnings from choosing X will be 1, while your earnings from choosing Y would be 2.

Forced Y. Note, however, that you may not be free to choose your preferred option. In each round, each of you will face a 15 percent chance that you are forced to choose option Y. We will call this a "forced Y."

Whether or not a player is forced to choose Y is randomly determined by the computer for each player separately and independently from the other players. Further, a forced Y does not depend on what happened in previous rounds.

On the computer screen where you make your decision, you will be reminded of this chance of a forced Y, for your convenience. Furthermore, in the table at the bottom of that screen (showing past decisions and earnings), your forced Ys are indicated in the column showing your choices with an "F." Note that you will not be informed of other players' forced Y choices.

You are now kindly requested to do a few exercises on the computer to make you fully familiar with the earnings table. In these exercises you cannot earn any money. Thereafter, we will start with part 1.

Please raise your hand if you have any question. We will then come over to your table to answer your question.

Instructions Part 2

Part 2 is exactly the same as part 1, except for one modification: In each round, each of you will now face a 30 percent chance that you are forced to choose option Y. Are there any questions?

Instructions Part 3

Part 3 is exactly the same as part 2, except for one modification: In each round, each of you will now face a 15 percent chance that you are forced to choose option Y, as in part 1. Are there any questions?

Instructions for the Heterogeneous Case

The instructions for the heterogeneous market case are shown below. Again, between the different experimental treatments only the percentages change. The instructions listed here are for the 15%-30%-15% case.

Instructions

Welcome to this experiment. The experiment consists of three parts in which you will have to make decisions. In each part it is possible to earn money. How much you earn depends on your own decisions and on the decisions of other participants in the experiment. At the end of the experiment a show-up fee of 5 euros plus your total earnings during the experiment will be paid to you in cash. Payments are confidential; we will not inform any of the other participants. In the experiment, all earnings will be expressed in talers, which will be converted to euros according to the exchange rate: 1 taler = 6 Eurocents.

During the experiment you will participate in a group of four players. You will be matched with the same players throughout the experiment. There are two types of players: A and B. The difference is related to the consequences of their decisions, as will be explained below. In fact, there will be one A player and three B players in your group. If you happen to be player A, then the others are B players, who will be labeled B1, B2, and B3. If you are a B player, then the other players in your group comprise a player A and two other B players, denoted as B2 and B3. You will learn your type when Part 1 starts; it will stay the same during the whole experiment. Because we have pre-assigned a type to each table, you have drawn your type yourself when you selected a table number in the reception room. You will not be informed of who the other players are, nor will they be informed of your identity.

Talking or communicating with others during the experiment is not permitted. If you have a question, please raise your hand and we will come to your table to answer it.

We start now with the instructions for part 1, which have been distributed also on paper. The instructions for the other two parts will be given when they start.

Instructions Part 1

First of all, note that your type (A or B) will be shown at the upper-left part of your computer screen, below a window showing the round number.

This part consists of thirty rounds. In each round you and the other three players in your group will have to choose one of two options: X or Y. Your earnings in a round depend on your type (A or B), your choice, and the choices of the other three players, in the following manner:

- If you choose Y, your earnings are 2, regardless of your type and the choices of the others.
- If you choose X, your earnings depend on your type and on how many of the other players choose Y.

Your exact earnings from choosing X or Y, given your type and the Y choices of the other players in your group, are listed in the following tables for, respectively, player A and a B player.

The following are some examples, for illustration.

Suppose you are a player A, and you choose X while one of the other players chooses Y. Then the upper table shows that your earnings will be 3.

Alternatively, suppose you are a B player, and you choose X while one of the other players chooses Y. Then it depends on whether this other player choosing Y is a player A or another B player. If it is player A, then the lower table shows that your earnings are 1, while your earnings are 3 if it is a B player. Thus, player A has a larger impact on your earnings than a B player.

Player A

Your Choice	Number of B Players Choosing Y	Your Earnings
X	0	5
X	1	3
X	2	1
X	3	−1
Y	0	2
Y	1	2
Y	2	2
Y	3	2

Player B

Player A's Choice	Number of Other B Players Choosing Y	Your Earnings from Choosing X	Your Earnings from Choosing Y
X	0	5	2
X	1	3	2
X	2	1	2
Y	0	1	2
Y	1	−1	2
Y	2	−3	2

Forced Y. Note, however, that you may not be free to choose your preferred option. In each round, each of you will face a 15 percent chance that you are forced to choose option Y. We will call this a "forced Y."

Whether or not a player is forced to choose Y is randomly determined by the computer for each player separately and independently from the other players. Further, a forced Y does not depend on what happened in previous rounds.

On the computer screen where you make your decision, you will be reminded of this chance of a forced Y, for your convenience. Furthermore, in the table at the bottom of that screen (showing past decisions and earnings), your forced Ys are indicated in the column showing your choices with an "F." Note that you will not be informed of other players' forced Y choices.

You are now kindly requested to do a few exercises on the computer to make you fully familiar with the earnings table. In these exercises you cannot earn any money. Thereafter, we will start with part 1.

Please raise your hand if you have any question. We will then come over to your table to answer your question.

Instructions Part 2

Part 2 is exactly the same as part 1, except for one modification: In each round, each of you will now face a 30 percent chance that you are forced to choose option Y. Are there any questions?

Instructions Part 3

Part 3 is exactly the same as part 2, except for one modification: In each round, each of you will now face a 15 percent chance that you are forced to choose option Y, as in part 1. Are there any questions?

References

Abbink, K., R. Bosman, R. Heijmans, and F. Van Winden. 2010. "Disruptions in Large Value Payment Systems: An Experimental Approach." DNB Working Paper No. 263.

Abbink, K., and J. Brandts. 2008. "Pricing in Bertrand Competition with Increasing Marginal Costs." *Games and Economic Behavior* 63 (1): 1–31.

Ambrus, A., B. Greiner, and P. Pathak. 2013. "How Individual Preferences Get Aggregated in Groups — An Experimental Study." Working Paper.

Arciero, L., C. Biancotti, L. D'Aurizo, and C. Impenna. 2009. "Exploring Agent-Based Methods for the Analysis of Payment Systems: A Crisis Model for StarLogo TNG." *Journal of Artificial Societies and Social Simulations* 12 (1).

Arciero, L., R. Heijmans, R. Heuver, M. Massarenti, C. Picillo, and F. Vacirca. 2016. "How to Measure the Unsecured Money Market: The Eurosystem's Implementation and Validation using TARGET2 Data." *International Journal of Central Banking* 12 (1): 247–80.

Armantier, O., and A. M. Copeland. 2012. "Assessing the Quality of 'Furfine-Based' Algorithms." Staff Report No. 575, Federal Reserve Bank of New York.

Ball, A., D. Denbee, M. Manning, and A. Wetherilt. 2011. "Intraday Liquidity: Risk and Regulation." Financial Stability Paper No. 11, Bank of England.

Basel Committee on Banking Supervision. 2013. "Monitoring Tools for Intraday Liquidity Management." Bank for International Settlements (April).

Bech, M., and R. Garratt. 2003. "The Intraday Liquidity Management Game." *Journal of Economic Theory* 109 (2): 198–210.

———. 2006. "Illiquidity in the Interbank Payment System Following Wide-Scale Disruptions." Staff Report No. 239, Federal Reserve Bank of New York.

Becher, C., M. Galbiati, and M. Tudela. 2008. "The Timing and Funding of CHAPS Sterling Payments." *Economic Policy Review* (Federal Reserve Bank of New York) 14 (2, September): 113–33.

Benoit, J., and V. Krishna. 1985. "Finitely Repeated Games." *Econometrica* 53 (4): 905–22.

Bornstein, G., T. Krugler, D. Budescu, and R. Selton. 2008. "Repeated Price Competition between Individuals and between Teams." *Journal of Economic Behavior and Organization* 66 (3–4): 808–21.

Bosman, R., H. Hennig-Schmidt, and F. van Winden. 2006. "Exploring Group Decision Making in a Power-to-Take Experiment." *Experimental Economics* 9 (1): 35–51.

Camerer, F. 2011. "The Promise and Success of Lab-Field Generalizability in Experimental Economics: A Critical Reply to Levitt and List." In *Handbook of Experimental Economic Methodology*, ed. G. R. Fréchette and A. Schotter (chapter 14). Oxford University Press.

Chaudhuri, A., A. Schotter, and B. Sopher. 2009. "Talking Ourselves to Efficiency: Coordination in Inter-generational Minimum Effort Games with Private, Almost Common and Common Knowledge of Advice." *Economic Journal* 119 (534): 91–122.

Cox, J., and S. Hayne. 2006. "Barking up the Right Tree: Are Small Groups Rational Agents?" *Experimental Economics* 9 (3): 209–22.

Devetag, G., and A. Ortmann. 2007. "When and Why? A Critical Survey on Coordination Failure in the Laboratory." *Experimental Economics* 10 (3): 311–44.

Diehl, M. 2013. "Measuring Free Riding in Large-Value Payment Systems: The Case of TARGET2." *Journal of Financial Market Infrastructures* 1 (3): 31–53.

Fréchette, G. 2015. "Laboratory Experiments: Professionals versus Students." In *Handbook of Experimental Economic Methodology*, ed. G. R. Fréchette and A. Schotter, 36–90 (chapter 17). Oxford University Press.

Galbiati, M., and K. Soramäki. 2011. "An Agent-based Model of Payment Systems." *Journal of Economic Dynamics and Control* 35 (6): 859–75.

Harrison, G., and J. List. 2004. "Field Experiments." *Journal of Economic Literature* 42 (4): 1009–55.

Harsanyi, J., and R. Selten. 1988. *A General Theory of Equilibrium Selection in Games.* MIT Press.

Heijmans, R., and R. Heuver. 2012. "Preparing Simulations in Large Value Payment Systems using Historical Data." In *Simulation in Computational Finance and Economics: Tools and Emerging Applications*, ed. B. Alexandrove-Kabadjova, S. Martinez-Jaramillo, A. Garcia-Almanza, and E. Tsang, 46–68 (chapter 3). IGI Global.

———. 2014. "Is this Bank Ill? The Diagnosis of Doctor TARGET2." *Journal of Financial Market Infrastructures* 2 (3): 31–53.

Jenkins, G. D., Jr., A. Mitra, N. Gupta, and J. D. Shaw. 1998. "Are Financial Incentives Related to Performance? A Meta-Analytic Review of Empirical Research." *Journal of Applied Psychology* 83 (5): 777–87.

Kocher, M., P. Martinsson, and M. Visser. 2008. "Does Stake Size Matter for Cooperation and Punishment?" *Economics Letters* 99 (3): 508–11.

McAndrews, J., and S. Potter. 2002. "Liquidity Effects of the Events of September 11, 2001." *Economic Policy Review* (Federal Reserve Bank of New York) 8 (2, November): 59–79.

McAndrews, J., and S. Rajan. 2000. "The Timing and Funding of Fedwire Funds Transfers." *Economic Policy Review* (Federal Reserve Bank of New York) 6 (2, July): 17–32.

Mehta, J., C. Starmer, and R. Sugden. 1994. "The Nature of Salience: An Experimental Investigation of Pure Coordination Games." *American Economic Review* 84 (3): 658–73.

Potters, J., and F. van Winden. 2000. "Professionals and Students in a Lobbying Experiment: Professional Rules of Conduct and Subject Surrogacy." *Journal of Economic Behavior and Organization* 43 (4): 499–522.

Pröpper, M., I. van Lelyveld, and R. Heijmans. 2013. "Network Dynamics of TOP Payments." *Journal of Financial Market Infrastructures* 1 (3): 3–29.

Raab, P., and B. Schipper. 2009. "Cournot Competition between Teams: An Experimental Study." *Journal of Economic Behavior and Organization* 72 (2): 691–702.

Schelling, T. C. 1960. *The Strategy of Conflict*. Harvard University Press.

Soramäki, K., M. Bech, and J. Arnold. 2007. "The Topology of Interbank Payment Flows." *Physica A* 379 (1): 317–33.

Van Huyck, J., R. Battalio, and R. Beil. 1990. "Tacit Coordination Games, Strategic Uncertainty and Coordination Failure." *American Economic Review* 80 (1): 234–48.

———. 1991. "Strategic Uncertainty, Equilibrium Selection, and Coordination Failure in Average Opinion Games." *Quarterly Journal of Economics* 106 (3): 885–910.