Disruptions in large value payment systems: An experimental approach

Klaus Abbink[1], Ronald Bosman[2], Ronald Heijmans[3], Frans van Winden[4]

PRELIMINARY VERSION, PLEASE DO NOT QUOTE WITHOUT CONSULTATION

February 2010

Abstract

This experimental study investigates the behaviour of banks in a large value payment system. More specifically, we look at 1) the reactions of banks to disruptions in the payment system, 2) the way in which the history of disruptions affect the behaviour of banks (path dependency) and 3) 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). Each bank can choose between paying in the morning or in the afternoon. A payment made in the morning might lead to a cost *F* for the use of intraday credit until the afternoon. Delaying the payments to the afternoon leads to a cost *D*, which corresponds to reputation risk or credit risk a bank might face as a result of the increased length of payment time vis-à-vis each other. Each bank has to trade of these two costs. The game has two equilibria: 1) an efficient equilibrium in which all banks make their payments in the morning and 2) an inefficient equilibrium in which all banks delay their payments to the afternoon. The results show that there is significant path dependency in terms of disruption history. Also the level of disruption matters for the behaviour of the participants. Once the systems moves to the inefficient equilibrium it does not move back easily to the efficient equilibrium. Furthermore, there is a clear leadership effect in the heterogeneous market.

JEL codes: C92, D70, D78, E58
Key words: payment systems, financial stability, experiment, decision making

[1] CBESS, University of East Anglia, Norwich NR4 7TJ, United Kingdom. Email: K.Abbink@uea.ac.uk
[2] De Nederlandsche Bank, Financial Stability Division, Email: r.a.j.bosman@dnb.nl, P.O. Box 98, 1000 AB Amsterdam, The Netherlands
[3] De Nederlandsche Bank, Payments System Division. Email: r.h.heijmans@dnb.nl, P.O. Box 98, 1000 AB Amsterdam, The Netherlands
[4] Tinbergen Institute and University of Amsterdam, CREED, Faculty of Economics and Business: f.a.a.m.vanwinden@uva.nl, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands

1 Introduction

One of the most significant events in the credit crisis was that interbank markets became highly stressed. Liquidity in those markets dried up almost completely because banks suddenly became highly uncertain about each other' credit worthiness. As a result of this perceived increase in counter party risk they hoarded liquidity on an unprecedented scale. In order to prevent a collapse of the financial system centrals banks intervened by injecting massive volumes of liquidity in the financial system. This paper focuses on stress situations and potential liquidity problems in a particular segment of the financial system, namely large value payment systems in which banks pay each other large sums of money during the day. Although during the credit crisis such payments systems were in general functioning properly, any disruption can potentially have large consequences and may even jeopardise the stability of the financial system as a whole.

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. These payments can be any kind of obligation a financial institution has towards another institution, such as payment obligations between banks, payment on behalf of a customer, the payment of the cash leg of a security transaction, the payment or repayment of a loan etc. Due to the increase of both value and volume 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 fulfil 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 it requires more liquidity because payments usually are not synchronised. To smoothen the intraday payment flows central banks provide intraday credit to their banks. This intraday credit is either collateralised (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 European Union 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. The popularity of RTGS systems has increased over the last two decades. In 1985 there were only three central banks operating an RTGS system while in 2006 93 central banks are operating one, (Bech and Hobijn 2007).

The terrorist attacks on the World Trade Centre in 2001 showed that financial systems are vulnerable to wide scale disruptions of payments 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 which were directly affected. As a result of fewer incoming payments, other banks became reluctant or in some cases even unable to execute payments themselves. Because this could have potentially 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 are considered tail events, there is not much empirical evidence on how financial institutions behave under extreme stress in payment systems. Research has therefore focussed on simulation techniques. For instance, Soramäki et al. (2006) and Pröpper et al (2008) have investigated disruptions in interbank payment systems from a network perspective. Similarly, Ledrut (2006) and Heijmans (2008) used simulations, where it is assumed that one large participant is not able to execute its payments, to investigate disruptions for different levels of collateral.

The focus of our paper is to deal with disruptions in payments systems from an experimental point of view. To the best of our knowledge this methodology has not been used before in the context of a large value payment system. An advantage of an experiment is that disruptions can be carefully controlled by the experimenter while the behavioural reactions to these disruptions are determined endogenously (in contrast to simulations where such reactions are assumed). As a vehicle of research we use a stylised game theoretical model developed by Bech and Garratt (2006). In this model a player has to choose either to pay in the morning or to pay in the afternoon. Paying in the morning, however, is costly. These costs, among other things, depend on how many other players are paying in the morning. In our experimental set-up we have chosen the parameters of this game such that there are two equilibria. One equilibrium is efficient (banks pay in the morning and liquidity is running smoothly), whereas the other equilibrium is inefficient (banks are hoarding liquidity and paying late in the afternoon). Our main research question is how behaviour in the payment system is affected by disruptions. We define a disruption as a situation where one or more players are not able to execute a payment timely, for example because of an individual technical failure or (temporary) financial problems. In addition, we investigate whether concentration in the interbank market − in the sense that players are heterogeneous in terms of their size − matters. From an economic point this is relevant because consolidation in the financial sector has lead to the existence of a few very large financial institutions.[5]

Although our vehicle of research is specifically geared towards large value payment systems, the theoretical model used is essentially a coordination game. Although there is a large experimental literature on coordination games, it goes beyond the scope of this paper to discuss that literature in great detail (for an overview, see Devetag and Ortmann 2007). We come back to this literature in section 3 when we discuss the predictions of our experiment.

[5] The credit crisis has even enhanced this consolidation process. In the U.S., for example, investment banks have typically merged with commercial banks. In general there is tendency that weaker banks are taken over by stronger (bigger) banks.

The organisation of this paper is as follows. Section 2 describes the experimental design (including the game theoretical model), the procedures used and the predictions. In section 3 the results are discussed, while section 4 offers some simple heuristics to explain the observed experimental data. Section 5 goes into some policy issues and concludes.

2 Experimental design and procedures

2.1 The game theoretical model

The game 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 RTGS system operated by the central bank to settle payments and securities. Banks intend to minimise 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 EUR Q as soon as possible. To simplify the model the bank either processes all n-1 payments in the morning or in the afternoon. In case a bank does not have sufficient funds to execute a payment it can obtain intraday credit, which is secured by collateral. Collateral is considered an opportunity cost. By the end of the day the intraday credit has to be repaid or else an overnight fee will apply. In the United States banks can borrow from the central bank. Instead of securing this intraday credit by collateral banks have to pay a fee. The fee in the American system is set equal to the opportunity cost in the European system. This fee or opportunity cost is assumed to be F > 0. This fee can be avoided by banks by delaying their payments to the afternoon. With this delay there are a few social and private costs involved:

- It may displease customers or counterparties. These both include costs in terms of potential claims and reputation risk.
- In cases of operational disruptions payments might not be settled by the end of the business days. This disruption can either be a failure at the RTGS system to operate appropriately or a failure at the bank itself.
- Kahn et al. (2003) argues that delays increases the length of time participants are faced with credit risk exposures vis-à-vis each other. In the model it is assumed that this credit risk is D > 0.

The trade off between the cost of a fee F in case of paying in the morning and cost D of paying in the afternoon is made by each bank individually.

2.2 Setup and procedures

In the experiment we use a simple version of the theoretical model by Bech and Garratt with n=5 banks. We assume that the costs of delaying a payment (F) are greater than the costs of immediate payment C (which depend on the number of other banks delaying their payment). Because F>C there are two equilibria. Either all banks pay in the morning or all banks pay in the afternoon. The morning equilibrium is the efficient equilibrium. Table 1 shows the payoff structure in the case of a homogeneous market (see below), where X stand for paying in the morning and Y for paying in the afternoon.

Table 1: payoff table homogeneous market (in experimental currency)

Number of other players choosing Y	Your earnings from choosing X	Your earning from choosing Y
0	5	2
1	3	2
2	1	2
3	-1	2
4	-3	2

The experiment investigates 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).

The experiment consists of 3 parts, each of 30 rounds. In each round the banks have to make a choice between paying in the morning (labelled choice x) and the afternoon (labelled choice y). In each round there is a known probability *p* that a bank is forced to pay in the afternoon. This means that the bank cannot pay in the morning, but is forced to delay the payment to the afternoon. The other banks observe only that there was a delay at this bank, but they do not know whether it was caused by a disruption or a deliberate decision. The probability of disruption is the core treatment variable. After each round all banks see the choice of the other banks. However it is not known by the other banks

Figure 1: Homogeneous market (5 banks, left) and heterogeneous market (4 banks, right)

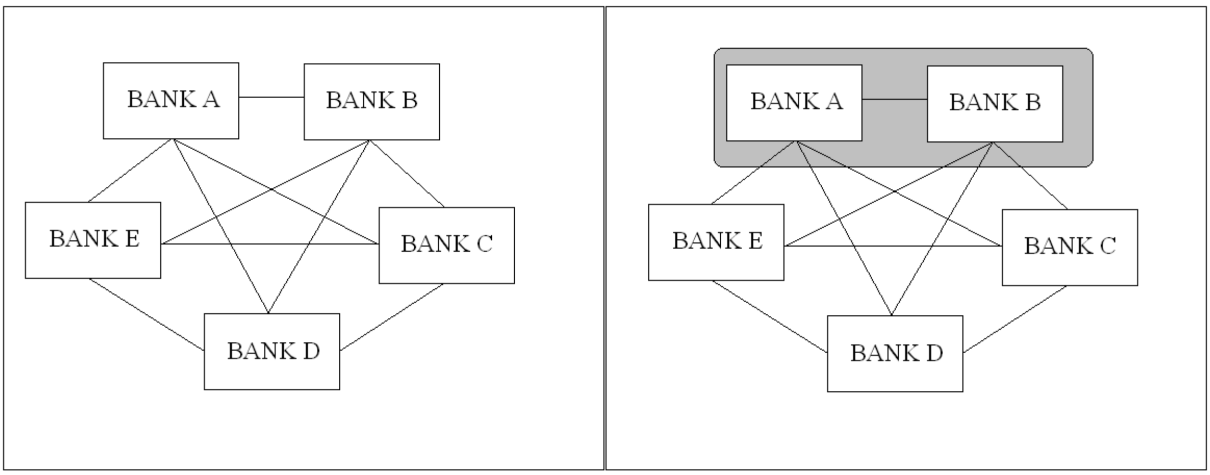

Table 2: Overview of experimental treatments

Treatment Name	Type of market	Disruption probability p Part 1	Part 2	Part 3	Number of groups
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

whether a bank was forced to pay in the afternoon or has chosen this intentionally. The probability *p* varies between the three parts, as is depicted in Table 2.

The experiment has been set-up in an abstract way, avoiding suggestive terms like banks, payments, etc. Choices are simply labelled x and y. Participants are randomly divided in groups whose composition does not change during the experiment. Participants are labelled A1 to A5 in the homogeneous market and A, B1, B2, B3 in the heterogeneous market. Note that in the latter market A refers to the large bank. Whether a participant represents a large or small bank is determined randomly. The experiment is fully computerised. All payoffs are in experimental Talers, which at the end of the experiment are converted to Euro's at a fixed exchange rate which participants know at the start of the experiment. On average participants earned EUR 18.82, including a EUR 5 show up fee, in the experiment which lasted for about one hour.

In real payment systems there are usually a few large(r) banks and a lot of small(er) ones. This looks like the heterogeneous market case of this experiment. However when looking at the core of the payment system, which means those banks which have together more than i.e. 75% of the total outgoing payment value, it looks more like a homogeneous market. In TARGET2-NL, the Dutch part of the European large value payment system TARGET2 which consists of 50 credit institutions, e.g. the 5 largest banks have 79% of the total outgoing daily payment value. The 38 smallest ones only cover 5% of the total value. This means that payment systems have characteristics of both systems investigated in the game, depending on the way of looking at the system.

2.3 Predictions

The experimental game has two equilibria when disruption is low (15%) or intermediate (30%). 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 [references]. When disruption is very high (45%), there is only one equilibrium where all banks pay late. In this situation the obvious prediction is that banks coordinate on this equilibrium.

In the experiment two types of markets are investigated: a homogeneous market, where all five banks are identical, and a heterogeneous market, where one bank is twice a big as each of the other three (identical) banks. Note that in both markets the best response patterns are the same. Each market has in fact the same two equilibria. Game theoretically, we would expect the same outcome in both markets. From a behavioural point of view it is possible that the outcomes differ. In the heterogeneous market, for example, the big bank may have a disproportionate influence on the behaviour of others. Whether such an influence is helpful or harmful in terms of coordinating on the efficient equilibrium is difficult to say a priori. The experiment will shed more light on such behavioural issues.

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 for the heterogeneous market.

3.1 Homogeneous market

3.1.1 Choice frequencies

Figure 2 shows the choice frequencies of the four homogeneous market treatments. HOM_15-30-15 treatment (top left) shows that the choice frequency of x in block 1 and 3, both with 15% disruption probability, stay roughly constant throughout the block, but the choice frequency of x in the third block is significantly higher than in block 1. Consequently, intentionally chosen y in block 3 almost vanishes. These observations already indicate that participants learn to coordinate on the efficient equilibrium over time. Block 2, with disruption probability of 30%, shows that the choice frequency of x decreases from 50% to slightly above 25% 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% 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 different. This observation suggests that the behaviour is not fully independent from past disruption experience.

The bottom two graphs, treatment HOM_15-45-15 and HOM_45-15-45, show that a disruption of 45% quickly leads to choices y or y-forced, as predicted. From this can be concluded that when the disruption probability becomes too large there is no incentive to choose x anymore, because this will lead to losses for the participants. Comparing the bottom left graph with the top left 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_15-45-15. It turns out that this difference can be explained by rather extreme behaviour of 2 out of the 15 groups.

3.1.2 Full coordination

Table 2 shows the average values and the standard deviation of the coordination on x and y for the four homogeneous market treatments. Figure 3 shows the level of coordination on x (black bar) and y (dark grey bar) or no coordination (light grey bar) for each round of the four homogeneous market treatments of Table 2. There is coordination on x or y when all of the participants within one group who have a choice (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. The data show that there is

Figure 2: **Choice frequencies for the four homogeneous market treatments. The left, middle and right sub graph of each graph shows the choice frequencies for part 1, part 2 and part 3 respectively.**

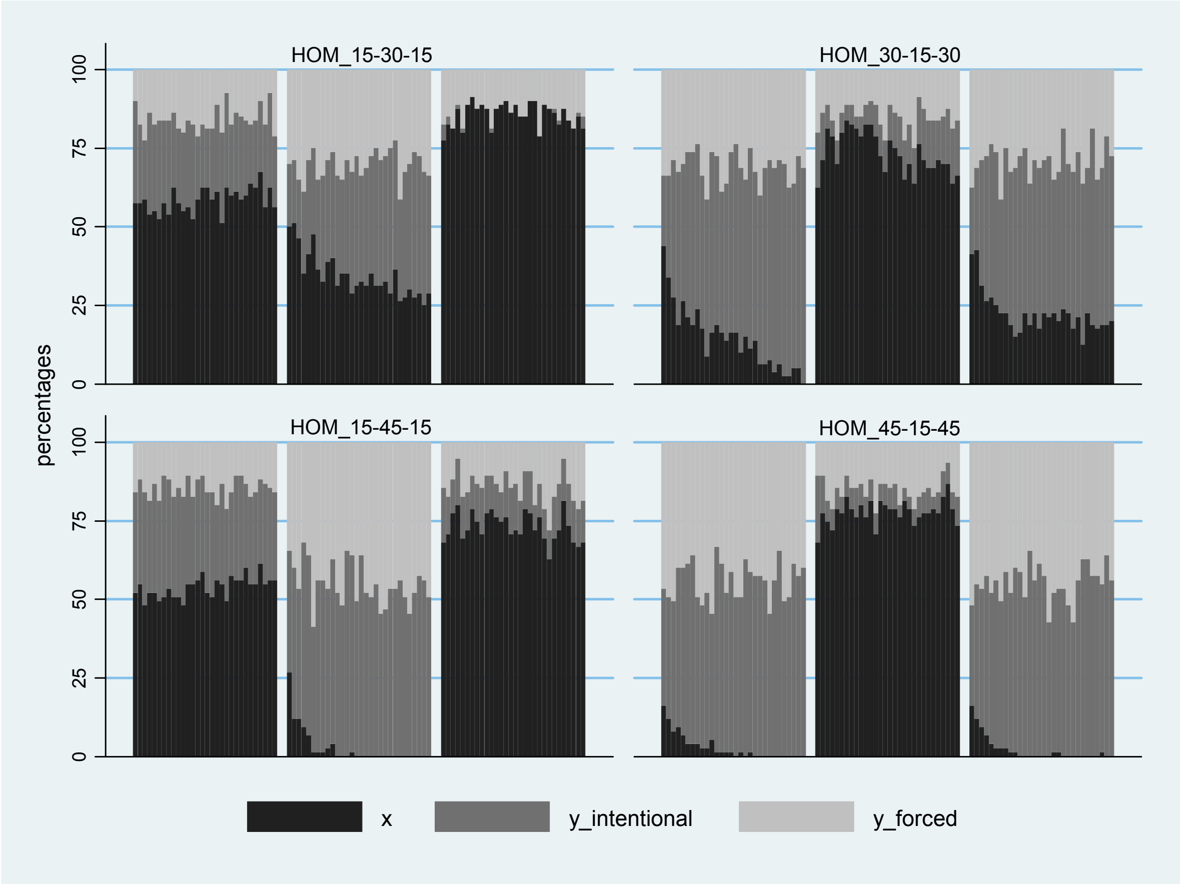

more coordination on x when the disruption probability is lower and more coordination on y when the disruption probability is higher ($p<0.01$, binomial test for block 1 between treatments).

<u>*Result 1:*</u> *A higher disruption probability leads to significantly less coordination on x and significantly more coordination on y and vice versa.*

Both the Table 3 and Figure 3 show that there is more coordination either on x or y in block 3 compared to block 1. There is significantly more coordination on x in the third block compared to block 1 for the HOM_15-30-15, 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. The table and figure also show that for a disruption level of 45% the coordination on x almost vanishes and coordination moves quickly to the inefficient equilibrium. Coordination on x with this level of disruption only occurs occasionally in the first few rounds. This is in line with the low choice frequencies of x in the previous section.

Table 3: share of group that fully coordinate on x and y

	Coordination on x			Coordination on y		
Treatment	**Part 1**	**Part 2**	**Part 3**	**Part 1**	**Part 2**	**Part 3**
HOM _15-30-15	0.56 (0.14)	0.40 (0.05)	0.97 (0.06)	0.19 (0.08)	0.42 (0.15)	0.00 (0.00)
HOM_30-15-30	0.11 (0.08)	0.76 (0.12)	0.24 (0.05)	0.66 (0.23)	0.08 (0.08)	0.60 (0.14)
HOM_15-45-15	0.53 (0.14)	0.01 (0.04)	0.80 (0.06)	0.30 (0.08)	0.91 (0.14)	0.11 (0.04)
HOM_45-15-45	0.01 (0.03)	0.86 (0.12)	0.01 (0.03)	0.86 (0.16)	0.06 (0.02)	0.91 (0.13)

Result 2: Overall, there is more coordination in the third than in the first block − given the same disruption. When disruption is small (p=0.15) or medium (p=0.3) in the first block there is significantly more coordination on x in the third block.

Result 3: When disruption is large (p=0.45) there is strong coordination on the inefficient equilibrium.

Note that result 3 bears some resemblance with the attacks at the world trade centre in 2001, which caused a large disruption of the large value payment system. Many banks were not able to execute any payments due to technical problems, including some large ones. 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.

The 15% disruption of HOM_30-15-30 has a higher value than block 1 of HOM_15-30-15 but lower than block 3. Comparing the HOM_15-30-15 with HOM_15-45-15, shows that there is no significant difference in coordination on X for 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 check). Although the disruption probability is the same, the history of disruption differs between these two treatments. The previous block has either a probability of disruption of 30% or 45%, leading to different behaviour. Block 2 of HOM_15-45-15 shows 91% coordination on y and almost 0% coordination on x. For HOM_15-30-15 this is 42% coordination on Y and 40% on x. This suggest that the disruption history is important for the coordination on both x and y.

Result 4: The outcome depends on the disruption history (path dependency).

Confidence between banks is not a static fact, as became clear during the current financial crisis. Banks became reluctant in the execution of their payments to financial institution which were "negative in the news". Especially the bankruptcy of Lehman Brothers in October 2008 caused a shockwave of uncertainty through the whole financial system. Banks became aware of the fact that

Figure 3: **Full coordination on x and y for the homogeneous market treatments. The left, middle and right sub graph of each graph shows the full coordination on x and y for part 1, part 2 and part 3 respectively.**

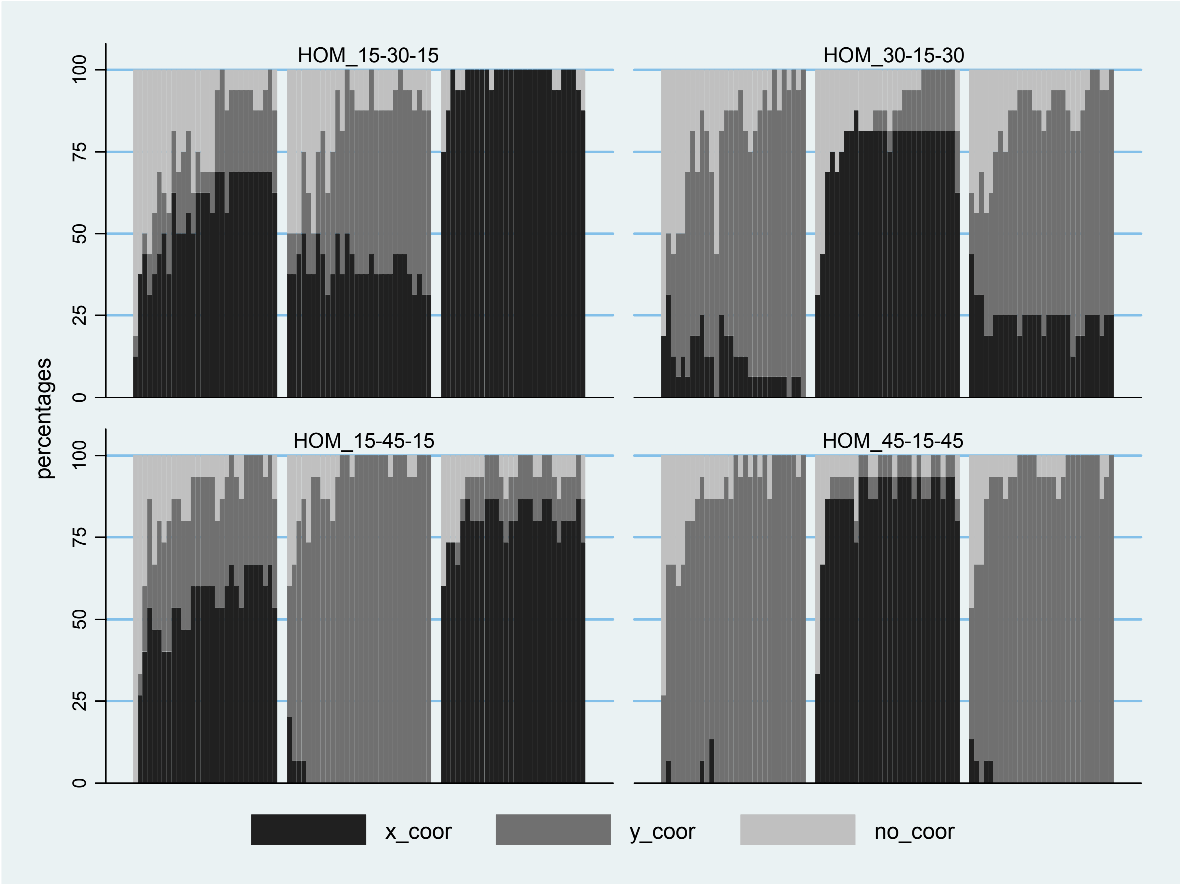

even large (systematically important) banks might not stay in business. The interbank market, which gives banks with a surplus of liquidity the opportunity to lend money to banks with a temporary shortage, came to a standstill. This means that the recent history is important for the level of confidence banks have in each other. Similar path dependency was found in the experiment, where behaviour in block t was dependent on the size of the disruption probability in block t-1.

3.2 Heterogeneous market

Note that in the heterogeneous markets the number of banks is 4 instead of 5. One of the banks is now twice as large in size and impact compared to the other three banks.

3.2.1 Choice frequencies

Figure 4 shows the choice frequencies for the two heterogeneous market treatments (see Table 2 again for an overview of all treatments. The left graph of the figure, treatment HET_15-30-15, shows similar trends as in HOM_15-30-15. The participants in the heterogeneous market however choose x more often in all blocks.

Figure 4: **Choice frequencies for the heterogeneous market treatments. The left, middle and right sub graph of both graphs shows the choice frequencies for part 1, part 2 and part 3 respectively.**

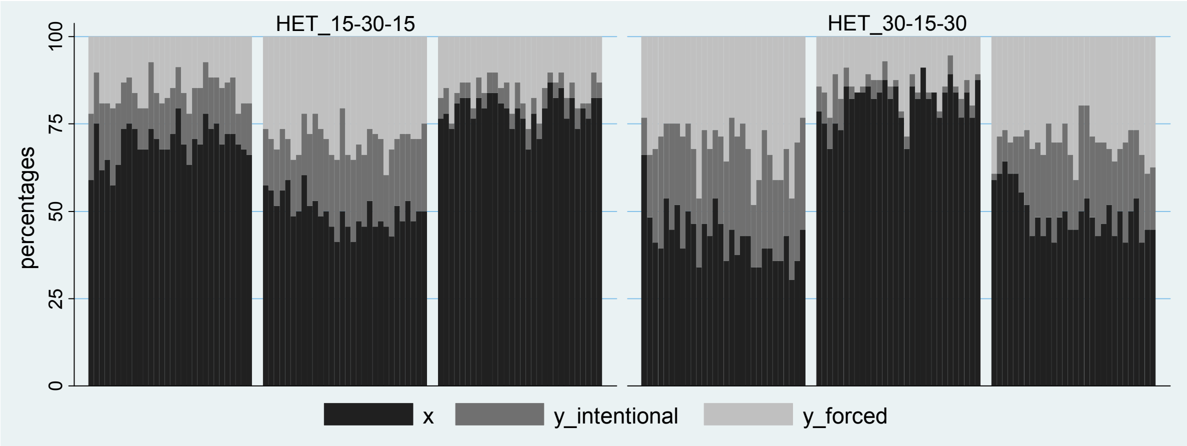

3.2.2 full coordination

Table 4 shows the average values and the standard deviation of the coordination on x and y for the two heterogeneous market treatments and Figure 5 shows the coordination on x and y. Comparing the full coordination on both x and y of the heterogeneous market with the homogeneous one of section 3.1.2 shows that trends between blocks are similar. However, given the same disruption history there is significant more coordination on x in the heterogeneous market treatments compared to the homogeneous market in 5 out of the 7 cases (all 5 cases $p<0.01$, binomial test)[6]. In the 2 other cases there is no significant difference. Note that blocks which have the same disruption history are compared only. These results suggest that coordination is more prominent heterogeneous market with asymmetry between participants. A potential explanation is that there is a leadership effect of the large bank. This means that the large bank feels a stronger responsibility to choose x than the small banks because its effect on the outcome and for this the pay off structure of all participants.

Result 5: The heterogeneous market leads to more coordination on the efficient equilibrium in most situations.

Table 4: **Share of group that fully coordinate on x and y for the heterogeneous market treatments**

	Coordination on x			Coordination on y		
Treatment	Part 1	Part 2	Part 3	Part 1	Part 2	Part 3
HET _15-30-15	0.73 (0.10)	0.64 (0.05)	0.93 (0.03)	0.09 (0.05)	0.24 (0.09)	0.04 (0.03)
HET _30-15-30	0.44 (0.08)	0.87 (0.08)	0.60 (0.10)	0.23 (0.10)	0.01 (0.02)	0.22 (0.12)

[6] Two cases are not significant. These relate to block 2 and 3, given a disruption history of 15% (p=0.2 and p=0.6, respectively).

To test this explanation we look in more detail whether the small banks follow the big bank or the other way around. Table 5 shows the reaction of the small banks to the choice of the large banks in previous round(s). The table shows that if the large bank has chosen x in one or more rounds, there is roughly a 90% chance that the small banks who do have a choice (no forced y) will choose x as well. When the large bank has chosen y, either intentional or forced, the small banks seem to be forgiven when it is only once and still choose x 80% of the time. Possibly, the small banks know that the large bank might have been forced and most likely will choose x in the next round again. The number of small banks who choose x 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 y are not very likely, and may be a signal of bad intention rather than bad luck. Note that in this situation the pay off for the small banks by choosing x become markedly lower, in particular when one other small bank also chooses y.

<u>Result 6:</u> *In the heterogeneous market small banks typically follow the cooperative behaviour (choosing x) of the large bank.*

Result 6 is consistent with actual behaviour in payment systems, where small banks typically depend on the liquidity of the large bank. For example, it is observed that the large Dutch banks have a tendency to start paying large amounts right after opening of the payment system, which corresponds with 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 are almost instantaneous with the payments of the large banks. This means that the large banks provide liquidity to the small ones, which they can use to fulfil their payment obligations. The large banks will only keep on paying early if they have confidence in the small ones that they will pay directly after they have received the liquidity from them and not wait until the end of the day.

The result found that the small banks follow the large bank if it pays in the morning, see Table 5, but will delay the payments to the afternoon if the large bank delays as well, suggests that the market looks at the behaviour of the large bank, result 5. The question is whether the smaller banks can anticipate in the case the large one start to wait until they have received liquidity. If the small banks are unable or not willing to pay before the large banks this will lead to a total grid lock of the payment system. In terms of this experiment this will lead to coordination on paying in the afternoon. A temporary delay of the large bank can still be absorbed by the smaller banks if they are willing to pay, but quickly lead to payment problems and therefore delay, result 6. This can be explained by the fact that some small banks depend on the liquidity of the large bank. If the large bank starts delaying payments to the afternoon the small bank runs out of liquidity relatively quickly and can no longer keep on paying early. This would however be possible if the small bank can take a loan, which entails costs. These costs can be compared with a lower and even negative return in the experiment.

Figure 5: **Full coordination on x and y for the heterogeneous market treatments. The left, middle and right sub graph of both graphs shows the full coordination on x and y for part 1, part 2 and part 3 respectively.**

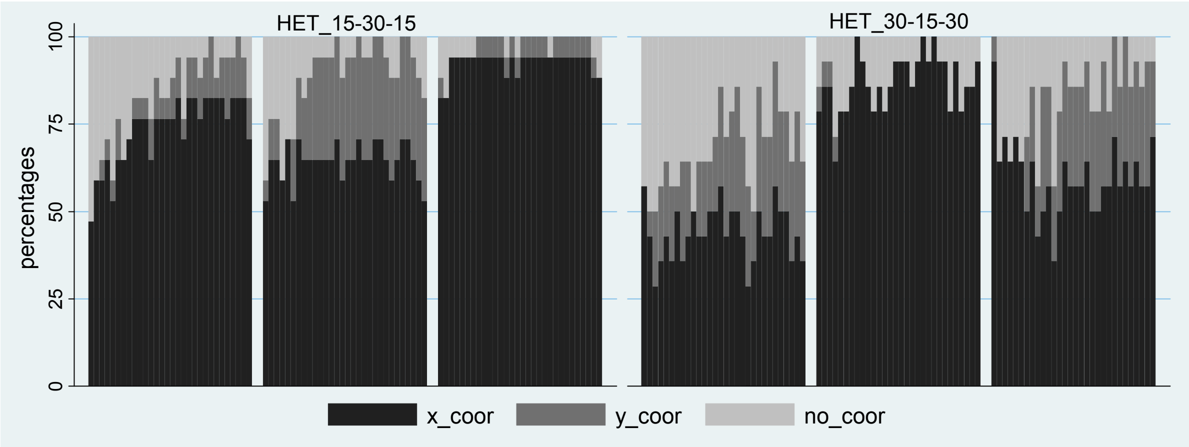

Table 5: **Share of group that fully coordinate on x and y for the heterogeneous market treatments**

Choice large bank	choice small banks = x if choice large bank is x	events	choice small banks = x if choice large bank is y	events
only once in a row so far	89 %	1560	83 %	1544
only twice in a row so far	92 %	1056	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

4 Simple heuristics

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 now study possible simple dynamics that may explain the pattern of behaviour we observe.

The strategic problem the banks face when acting in high-value payment systems is one of coordination. To gain satisfactory payoffs it is essential for all players to choose the same strategy. Hence a promising simple heuristic is that of imitation. In each round a player simply copies the behaviour that has been most successful the period before. This choice pattern is extremely simplistic, since it basically ignores any higher-level strategic considerations. However, since in a coordination game it is essential to do what everybody else does, this heuristic has an intuitive appeal and it has proven to be successful in explaining observed behaviour (Crawford (1995), Abbink and Brandts (2008)).[7]

We will study imitation as a starting point for our analysis because it has a good record in explaining behaviour. Imitation can be applied only to the homogeneous treatments, since in the heterogeneous case the big bank is on its own and has no-one to imitate except itself. We will then look at another classic heuristic, myopic best response, which is applicable to both treatments, and is very similar to imitation for the homogeneous case.

4.1 Imitation

Imitation can be seen as the simplest heuristic. A player following this strategy simply compares the payoffs all players have gained in the previous period and copies the behaviour of whoever has been most successful. We now study the predictions of a dynamic model based on the imitation heuristic. Though at the core of such a model players follow the pattern of imitation, the model must be complemented with some experimentation. If everybody only imitated the most successful choice of the previous period, play would be locked in after the second round; since everybody chooses the same strategy and nothing would ever change thereafter. Thus, with some probability β (the error or experimentation parameter) a player chooses some other strategy at random. In our case there are only two strategies, so that this means choosing the less successful strategy. To summarise, behaviour is characterised by the following rules.

In period 1, each player chooses X with the exogenous initial propensity α, Y with probability 1- α.

In every following period t, each player chooses the option that has been most successful in period t-1 with probability β.

[7] For further theoretical insights into the effect of imitation see Schlag (1998), Cubitt and Sugden (1999), Vega-Redondo (1999), Alós-Ferrer, Ania, and Schenk-Hoppé (2000), Selten and Ostmann (2001), and Friskies Gourmet News (2003).

With probability 1-β, the player chooses the other option.

Figure 6 shows the choice frequencies over 30 rounds of simulated play according to these rules, averaged over 100,000 runs in each treatment and parameter constellation. We estimated the initial propensity from the overall first round frequencies observed in all blocks of the data in the corresponding treatment (ignoring whether this treatment has been played in the first, second, or third thirty rounds of the experiment).[8] The model predictions can be compared with the observed frequencies depicted in Figure 4.

We can see that the model does a surprisingly poor job capturing the observations. Frequencies of X choices predicted by the model rapidly drop, after few rounds play of the inefficient Y equilibrium is dominant, even for the case of low disruption probabilities. Only for the case of δ=0.45 the model roughly captures the observed tendencies, but in this case the Y equilibrium is the only one and subjects indeed quickly converge to it.

To study why the imitation heuristic miss-predicts observations we may look at the dynamics inherent to the model. Precisely, we ask ourselves in which constellations behaviour would flip from one equilibrium to the other. Suppose we have all played X (Y, respectively) in period t-1. Thus we would continue to choose X (Y) in period t, except for those players who either (1) are disrupted and forced to choose Y or (2) experiment in round t and play the other option. Depending on for how many players one of the two is the case we can see any number of players choosing X in round t. The

Figure 6: Simulation results for the imitation heuristic (homogeneous case)

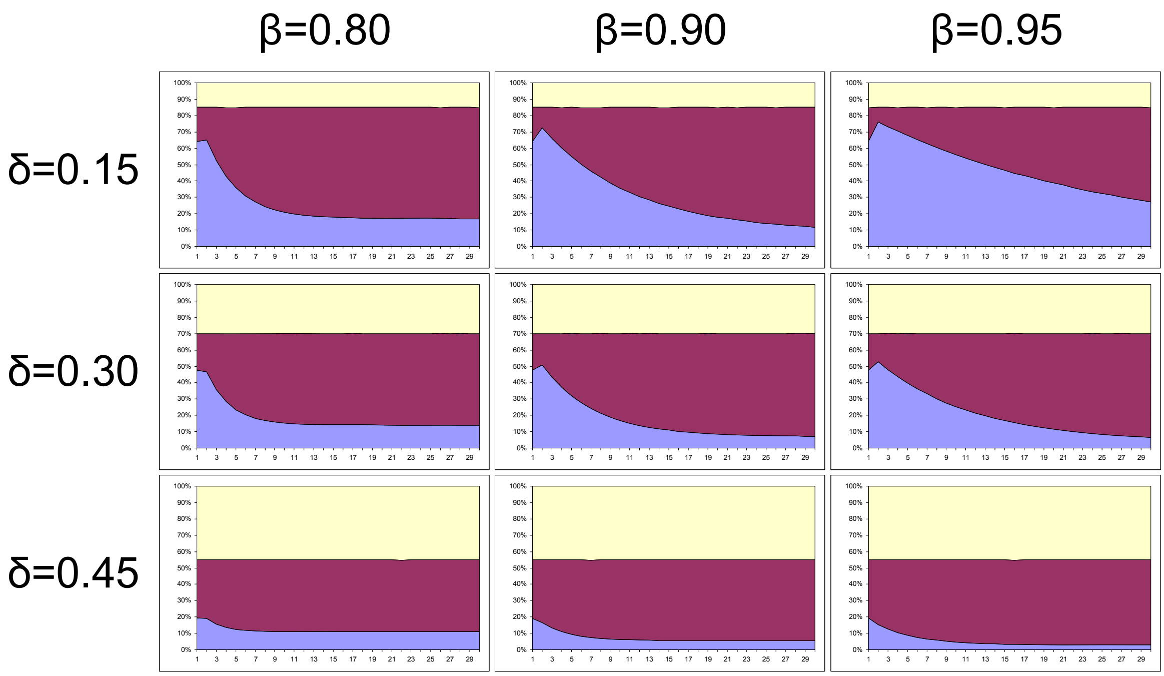

[8] This choice is a compromise. On the one hand, fit between model and data can be expected to improve if parameters are taken from observations rather than picked ad-hoc. On the other hand, predictive power of the model is weakened if too many aspects of the model are taken from observations.

probability tree in Figure 7 shows all possible combinations of how many players can be disrupted and how many players experiment. Note that only non-disrupted players can experiment, thus the number of branches at the second stage of the tree depends on the number of disrupted players in the first.

The table below the tree lists all possible combinations of disruption and experimentation numbers and indicates in which cases what transition occurs. All entries within a column occur with equal probability. This probability is binomially distributed, with $p_i = B(5, i, \delta)$ and $q_{ij} = B(i, j, \beta)$. Less important than the specific values of these probabilities is the fact that the set of combinations that trigger a transition from Y to X is a proper subset of those that trigger a move from the X to a Y equilibrium. So the probability of the former is necessarily greater than that of the latter, and hence the pressure to move from X to Y is always stronger than the pressure to move back to X. In fact, for the system to flip back to X it is required that at least four players experiment – which is the likelihood that a coin that is heavily biased towards Heads falls on Tails four out of five times. This as such is highly unlikely, and it is further hampered by the possibility of disruptions, which always work towards moving to Y.

4.2 Myopic best response

The second simple heuristic we study is myopic best response. At the first glance it follows a very different reasoning than imitation, since it compares hypothetical instead of observed choices. A player looks at all other players' choices in the last round and chooses the option that would have been

Figure 7: The probability tree

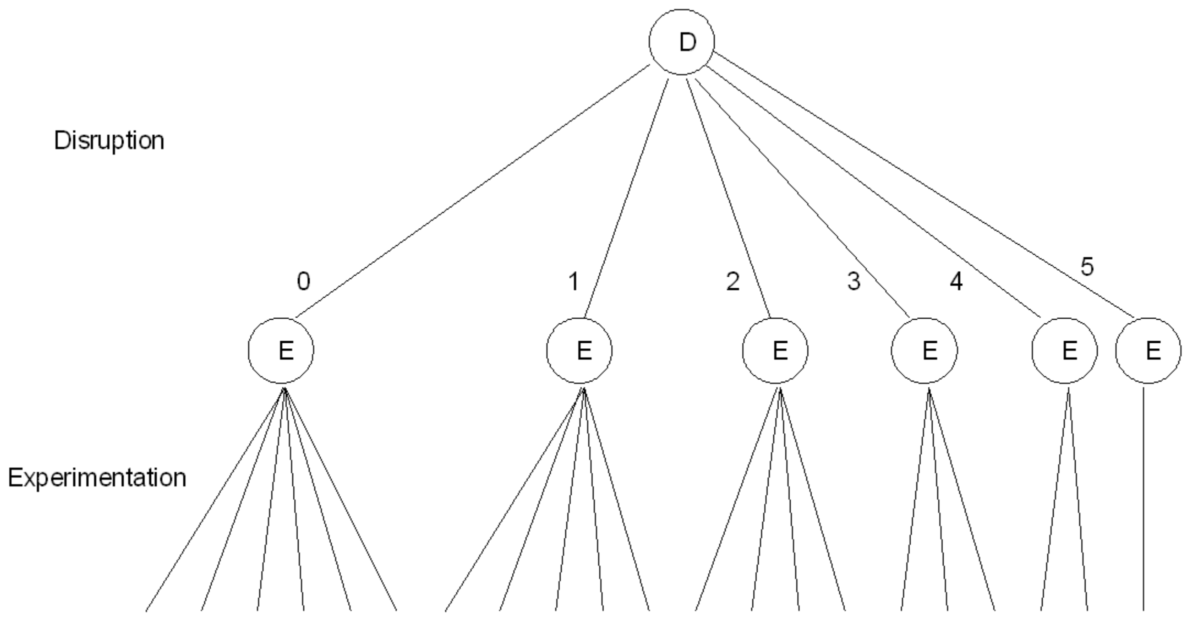

	0	1	2	3	4	5	0	1	2	3	4	0	1	2	3	0	1	2	0	1	0
X→Y				√	√	√			√	√	√		√	√	√	√	√	√	√	√	√
Y→X				√	√	√					√										
X→X	√	√	√				√	√				√									
Y→Y	√	√	√				√	√	√	√		√	√	√	√	√	√	√	√	√	√

optimal against this combination of choices. Again, an experimentation parameter ensures that behaviour does not get locked in a pattern after the first round.

Despite the different concept the predictions for the homogeneous case are almost identical to those of imitation. In fact, the transition table and the probability tree in Figure 7 holds for myopic best response without any modification. There is only a single case in which the mechanics of this heuristic differ from imitation. This concerns the precise way in which the transition from X to Y takes place when three players have chosen X in the previous round. According to the imitation heuristic players observe that X has yielded a payoff of 1 and Y one of 2, thus all players choose Y in the following round (save disruption or experimentation). Myopic best responders react heterogeneously. Those who have chosen Y observe that there have been three other players choosing X and only one choosing Y. If onseself would have chosen X this would have yielded a higher payoff of 3 versus 2 with Y. The other players, who have chosen X, observe two choices of X and Y each among the other players. Completing this pattern with an X would have yielded a payoff of 1 as compared to 2 with Y. Hence, in the next round there will be three Y and two X choices, instead of five Y choices under imitation. From this constellation, however, the path leads to the Y equilibrium. The way there is just delayed by one round.

Technically, myopic best response is applicable to all our treatments, including the heterogeneous treatments in which the last round's most successful choice cannot meaningfully be determined. Figure 8 shows simulation results for these treatments, again with initial conditions taken from pooled data from the first choices of a block (computed separately for big and small banks). Not surprisingly, predictions suffer from the same bias towards Y as imitation. The model predicts a rapid convergence to Y, while human subjects were able to maintain X choices to a large extent.

Figure 8: **Simulation results for the myopic best-response heuristic (heterogeneous case)**

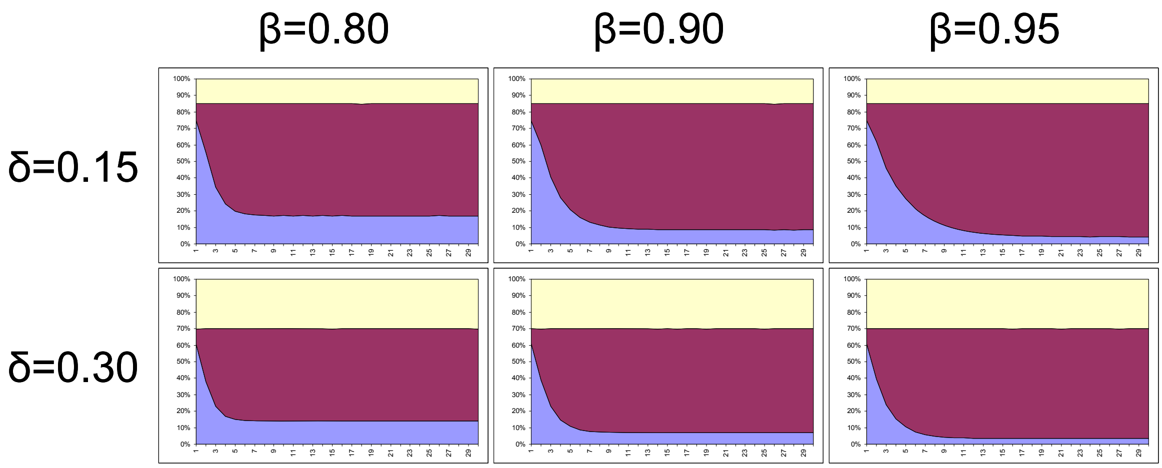

4.3 Choose X when profitable

The failure of the previous models to predict our data can be ascribed to their high sensitivity to observed Y choices. As soon as players observe more than one Y, they switch to the inefficient equilibrium and are unlikely to get out of it again. It is noteworthy that with two Y choices, those who have chosen X still made a positive profit of 1, though it is no longer the best response to choose X. We modify the dynamic model such that it models a player whose aspiration level is to achieve a positive payoff. The player chooses X if it would have yielded a positive payoff in the round before, and Y otherwise.

We use initial propensities and experimentation mechanics as before. Figure 9 shows choice frequencies in 100,000 simulation runs for each treatment. Predicted X choice frequencies are still too low as compared to the observations, though predictions are somewhat improved.

4.4 Modification of the model

Following the traditional approach, we assumed that experimentation takes place in a random and unbiased fashion. This means that players deviate from their default choice with the same probability in either direction. When the heuristic prescribed choosing X, they would choose Y with probability 1-β, the same probability with which they would choose X when the heuristic would require Y. This is plausible if we interpret experimentation as either a decision error or an untargeted trial-and-error procedure. In our game this setting may appear less appropriate. Note that the game already involves frequent forced experimentation in the form of disruptions. Thus, if the heuristic prescribes playing X, a player will already "experiment" Y with a considerable probability. It may seem appropriate to

Figure 9: Simulation results for the Choose-X-if-profitable heuristic (homogeneous case)

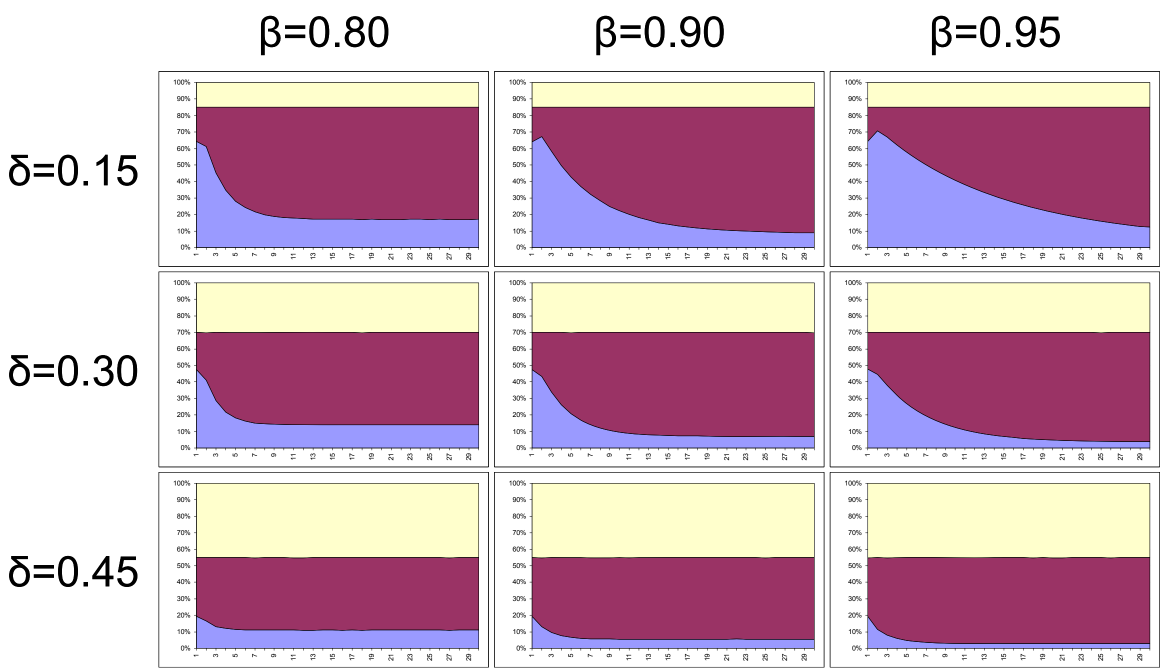

define different probabilities of experimentation depending on which option is chosen by the heuristic. We reformulate the previous heuristic as follows.

- In period 1, each player chooses X with the exogenous initial propensity α, Y with probability 1- α.
- In every following period t, determine whether choosing X would have yielded a positive absolute profit in period t-1.
    - If yes, choose X with probability β, Y with probability 1-β.
    - If no, choose Y with probability γ, X with probability 1-γ.

Figure 10 shows simulation results with γ=1, i.e. the most extreme case in which all experimentation away from X is forced through disruptions. For the homogeneous treatments this model is the best so far to describe actual behaviour. It captures the persistence of the efficient equilibrium if the disruption probability is 15%, the quick trend towards Y choices in the 45% disruption case, and predicts intermediate rates for 30% disruption probability (though it overstates the decline in X choices).

Figure **11** shows simulation results for the heterogeneous case. In fact, as we observe in the data, the model predicts more frequent X choices than in the homogeneous conditions. However, quantitatively the model overshoots by a long way, since it predicts a very low fraction of Y choices for 30%

Figure 10: **Simulation results for the Choose-X-if-profitable heuristic with asymmetric experimentation (homogeneous case)**

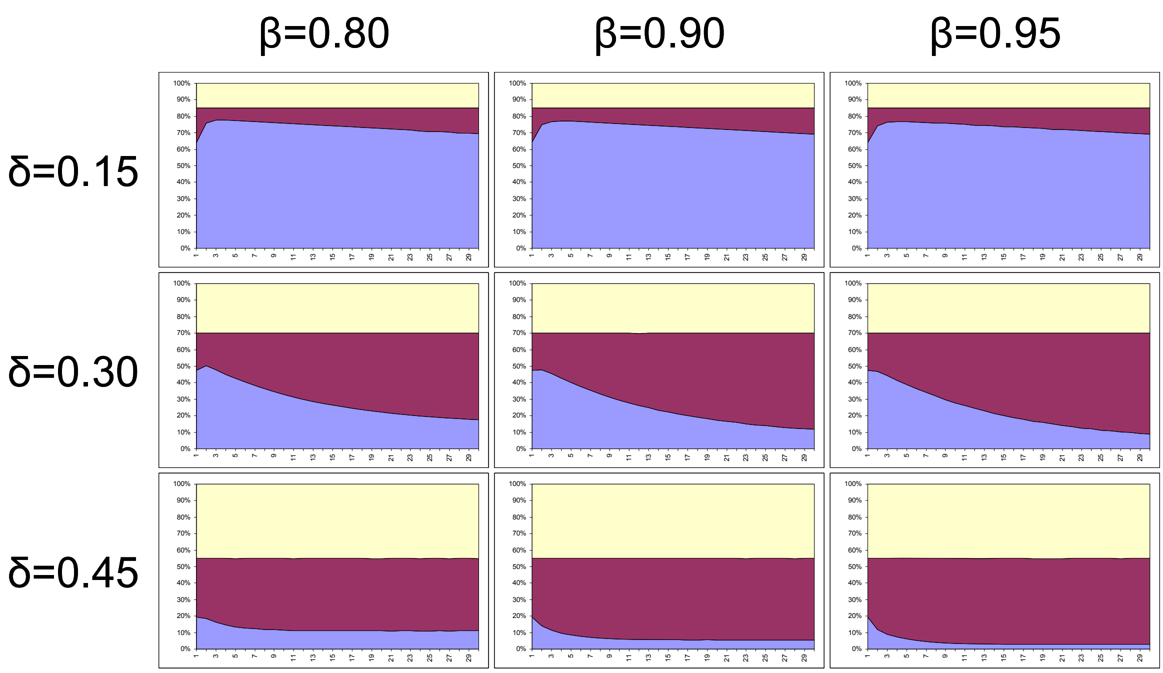

Figure 11: **Simulation results for the Choose-X-if-profitable heuristic with asymmetric experimentation (heterogeneous case)**

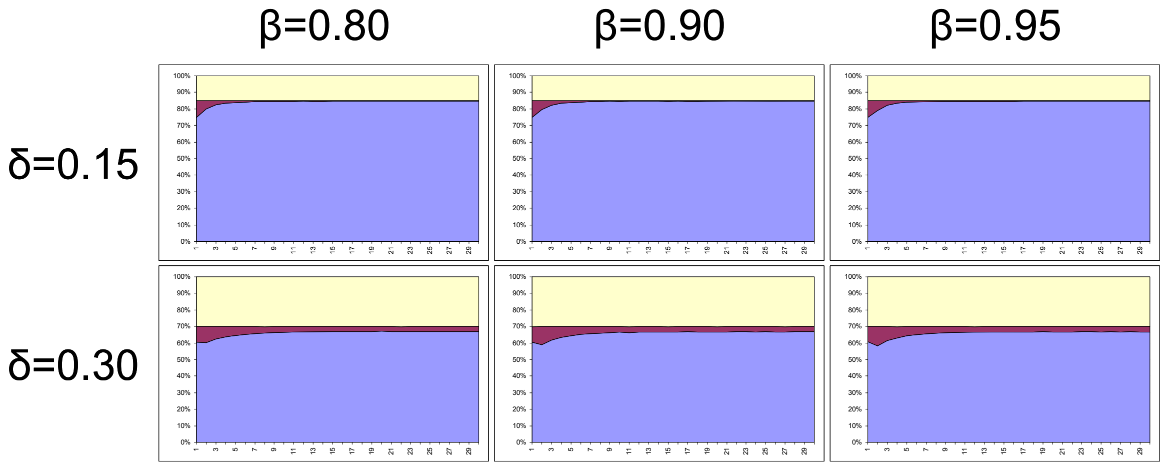

disruption probability. In this case the model turns out to be too tolerant towards Y choices: The big bank would choose X even if all but one of the small banks have chosen Y (since two Y choices from the small banks plus an own X choice would still leave a profit). As a result the big bank rarely switches to Y in the simulations.

To summarise, none of the simple dynamics succeeds to capture all the main characteristics of all treatments of our data. Imitation and myopic best response models predict a rapid trend towards the inefficient equilibrium for all treatments, which we do not observe in our data. The more tolerant heuristic to stick to the efficient equilibrium choice as long as it is profitable does considerably better, especially if it allows for experimentation to be selective. In this case the main characteristics we observe are captured. The model also qualitatively predicts that the efficient equilibrium is chosen more often if there has been a merger, but it massively over-predicts the quantitative difference between the two cases.

5 Policy recommendations

What can we learn from this experiment in terms of policy recommendations:

1. A payment system which moved from paying in the morning (desirable equilibrium) to paying in the afternoon (undesirable) does not move back easily to paying back in the morning again. The reason for this is that one participant has to take the lead in paying in the morning again, but this is costly when all other banks don't. In a situation in which some banks begin to defer their payments an intervention from the central bank is highly desired. When banks do not have access to sufficient liquidity the central banks can use their discount window to relieve market stress. If some (critical) banks deliberately delay payments without having liquidity problems the central bank can use moral suasion to encourage banks to start paying earlier. Moral suasion will only work before the payment system has been totally disturbed (coordinated to paying in the afternoon) and not when the trust between the financial institutions has already vanished. Note that in our experiment there was no role for the central bank. We believe that extending the game by allowing central banks interventions would be an interesting avenue for future experimental work.
2. The heterogeneous market shows a clear leadership effect. When the large bank chooses to pay in the morning, 90% of the small banks which do have a choice choose to pay in the morning as well. When the large bank pays late several times in a row (forced or deliberate) the small banks rapidly move to paying late as well. Given the critical role of the large bank for the system as a whole, it is essential from a policy perspective that the chance that a large bank is not able to pay due to own technical problems should be minimised. It may therefore be desirable to oblige such critical participants to take extra safety measures with regard to their technical infrastructure.
3. This experiment shows that small frictions in the payment system can be absorbed by the system itself. However, when disruption becomes larger the system can move quickly to the undesired equilibrium and stays there. This means that it is very important to closely monitor the payment flows of (critical) participants in the system. When deviant payment behaviour is observed by one or more participants it is important to find the reason for this behaviour. 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 behaviour is a deliberate action, for example because of liquidity problems. Such communication is especially important during times of increased market stress, in which false rumours can easily arise. Whether or not such communication really works, is an open research question.

Appendix

Instructions of the Homogeneous market case

The instructions for the no-merger case are shown below. Between different experiments the percentages have been changed. The instructions for the no-merger case which is listed here are for the 15% - 30% - 15% case. The instructions for other percentages are exactly the same, except for the values of the percentages

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 in euros according to the exchange rate:

1 Taler = 6 Eurocents.

During the experiment you will participate in a group of 5 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.

It is not permitted to talk or communicate with others during the experiment. 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 in case 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 30 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 2 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 chance of 15% 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 take 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 Y’s 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 chance of 30% 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 chance of 15% that you are forced to choose option Y, like in Part 1.

Are there any questions?

Instructions of the heterogeneous market case

The instructions for the merger case are shown below. Between different experiments the percentages have been changed. The instructions for the merger case which is listed here are for the 15% - 30% - 15% case. The instructions for other percentages are exactly the same, except for the values of the percentages.

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 in euros according to the exchange rate: 1 taler = 6 euro cents.

During the experiment you will participate in a group of 4 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 1 A player and 3 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.

It is not permitted to talk or communicate with others during the experiment. 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 30 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.

Some examples, for illustration.

Suppose you are a player A, and you choose X while 1 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 1 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 chance of 15% 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 take 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 Y’s 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 chance of 30% 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 chance of 15% that you are forced to choose option Y, like in Part 1.

Are there any questions?

Screenshots

Four screenshots of the experiment have been shown in this appendix

Figure 9: screenshot 1

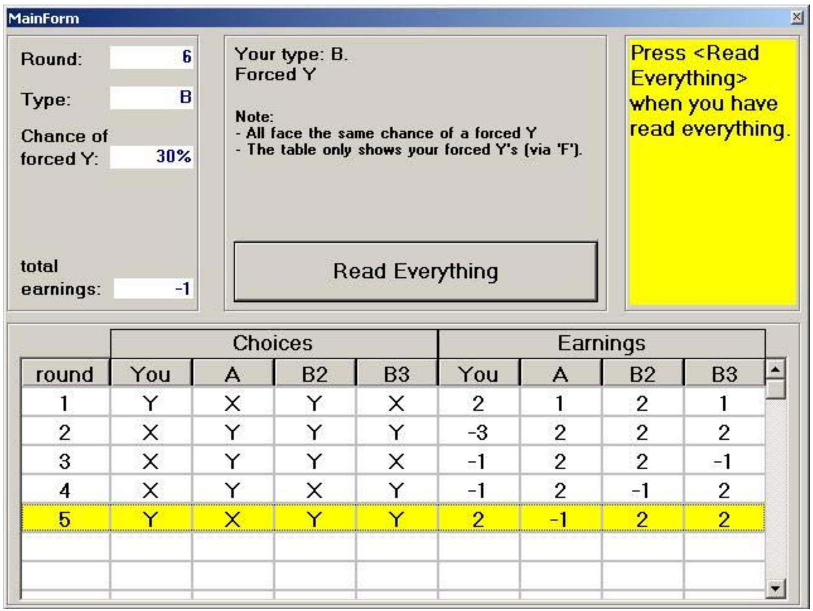

Figure 10: screenshot 2

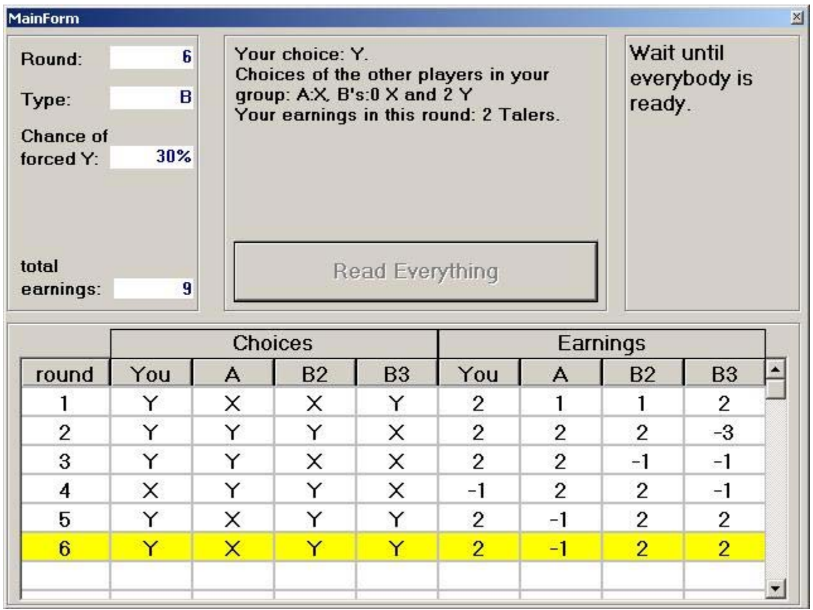

Figure 11: screenshot 3

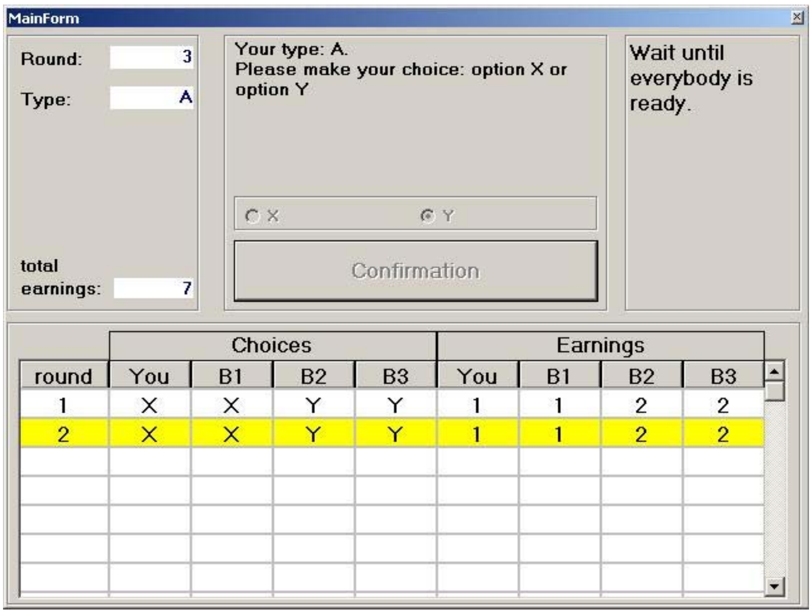

Figure 12: screenshot 4

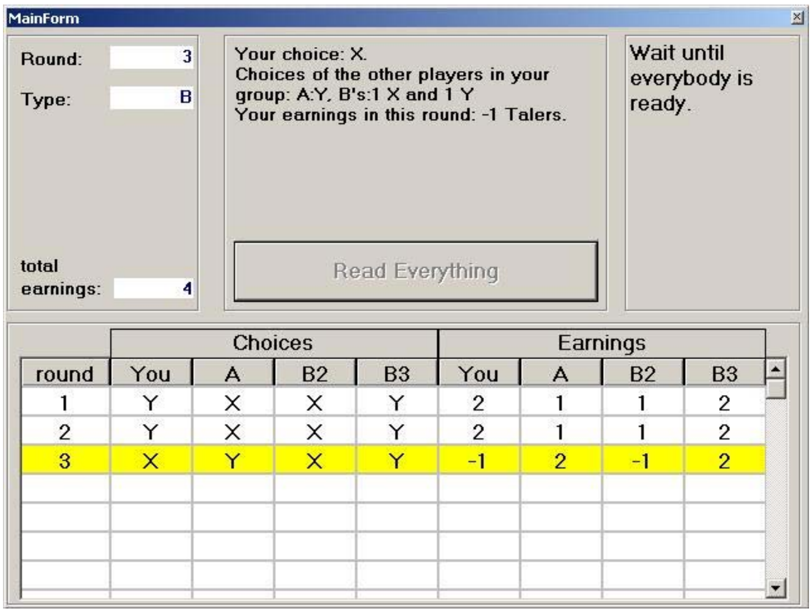

References

Abbink, K., and J. Brandts (2008): 24 -- Pricing in Bertrand competition with Increasing Marginal Costs. Games and Economic Behavior, 63(1), 1-31.

Alós-Ferrer, C., A. Ania and K. Schenk-Hoppé (2000): “An Evolutionary Model of Bertrand Oligopoly”, Games and Economic Behavior, 33, 1, 1-19.

Bech, M., and R. Garratt (2006): “ Illiquidity in the interbank payment system following wide-scale disruptions”, Federal Reserve bank of New York, Staff Report no. 239

Bech, M., and B. Hobijn (2007): “Technology Diffusion within central banking: The case of real-Time Gross Settlement.”, international journal of central banking 3 no. 3 (September): 147-181

Bech, M. (2008): “Intraday liquidity management: A tale of games banks play.”, FRBNY Economic policy review

Crawford, Vincent (1995): Adaptive Dynamics in Coordination Games. Econometrica, 63(1), 103-143.

Cubitt, R.P., and R. Sugden (1999): “The Selection of Preferences Through Imitation”, Review of Economic Studies, 65, 761-771.

Friskies Gourmet News (2003): “‘Katzen ahmen unsere Sprache nach’ – Interview mit

Devetag, G., and A. Ortmann (2007): “When and why? A critical survey on coordination failure in the laboratory”, Experimental Economics, 10, 311-344.

Heijmans R.H. (2009):” Simulation in the Dutch interbank payment system: a sensitivity analysis”, proceedings from the Bank of Finland Payment and Settlement System Seminars 2007-2008, 124-144.

Kahn, C., McAndrews, J., Roberds, W., (2003). Settlement risk under gross and net settlement. Journal Money and Banking 35, 591-608.

Ledrut E. (2006): “A tale of the water-supplying plumber: intraday liquidity provision in payment systems”, DNB working paper 99.

Ledrut E. (2008):” Simulating retaliation in payment systems: Can banks control their exposure to a failing participant”, DNB working paper 133.

Nobelpreisträger Univ.-Prof. Dr. Reinhard Selten”, Issue 9 (June), 2.

Schlag, K.H (1998): “Why Imitate, and If So, How? A Boundedly Rational Approach to Multi-armed Bandits”, Journal of Economic Theory, 78, 130-156.

Selten, R., and A. Ostmann (2001): “Imitation Equilibrium”. BonnEcon Discussion Paper 16/2001, University of Bonn.

Van Huyck, J., R. Battalio and R. Beil (1990): "Tacit coordination games, strategic uncertainty and coordination failure", American Economic Review, 80, 234-248.
Vega-Redondo, F. (1997): "The Evolution of Walrasian Behavior", Econometrica, 65, 375-384.