Information and Learning in the Minority Game: A Strategy Experiment

Jona Linde[a], Daniel Gietl[b], Joep Sonnemans[c,d] and Jan Tuinstra[c,d]*

[a] School of Economics and Business, Maastricht University, The Netherlands
[b]Seminar für Wirtschaftspolitik, Ludwig-Maximilians-Universität, München, Germany
[c]Amsterdam School of Economics, University of Amsterdam, The Netherlands
[d] Tinbergen Institute, The Netherlands

* Corresponding author: J.Tuinstra@uva.nl, +31 20 525 4227, P.O. Box 15867, 1001NJ Amsterdam, The Netherlands

March 31, 2016

Abstract: Minority games provide stylized descriptions of decision problems where actions are strategic substitutes. Linde et al. (2014) presents a multi-round strategy-method experiment on the five-person minority game. A remarkable outcome of that experiment is that aggregate efficiency does not increase over the five rounds of the experiment. In the experiment we present in this paper we explore whether the absence of increasing efficiency is due to a lack of information on how to develop better strategies. To examine this we give participants complete information about the syntax and performance of all strategies submitted in the previous round by the other participants, and allow them to choose these strategies for practice simulations. We find that increased information and extended simulation possibilities have a *negative* effect on aggregate efficiency. The reason for this is that participants tend to adjust their strategies in the direction of the winner(s) of the previous round. Strategies therefore become more similar to each other, which reduces efficiency in the minority game.

Keywords: minority game, strategy experiment, learning, information, imitation
JEL codes: C63, C72, C91, D03

Acknowledgements
We thank the Research Priority Area Behavioral Economics of the University of Amsterdam for financial assistance and we are grateful to Ulrich Glogowsky and participants to CEF2014 (Oslo), the 2014 International Symposium on Dynamic Games and Applications (Amsterdam), WEHIA 2015 (Sophia Antipolis), the 2015 Experimental Finance conference (Nijmegen) and a seminar at the University of Technology Sydney for useful comments and suggestions.

1. Introduction

Does more information lead to better decisions? We study this question in a strategy experiment on the minority game. The minority game is a highly stylized representation of games where actions are strategic substitutes, such as congestion games, market entry games and Cournot oligopoly games. Although the rules of the game are simple (an odd number of players simultaneously choose one of two alternatives, and only the players making the minority decision receive a fixed reward), it is not obvious how the game should be played, in particular if it is repeated with a fixed set of players.

The minority game therefore provides a good framework to study which strategies players use and how they adapt these strategies over time and with experience. Linde et al. (2014) presents a five round strategy experiment where in each round participants have to submit a strategy to play the five player minority game for 100 periods. The strategies subsequently play against each other in a computer tournament (with the five participants submitting the most successful strategies for that round receiving monetary prizes), after which participants can revise their strategy for the next round.

Linde et al. (2014) focuses on an evolutionary analysis using the strategies gathered in the experiment and finds that evolution leads to a few surviving strategies and a remarkable high efficiency.[1] However, the results in each of the five rounds of the experiment itself lead to aggregate outcomes consistent with the symmetric mixed strategy Nash equilibrium which is the least efficient equilibrium. Aggregate performance does not increase over the rounds, suggesting participants are not able to improve their strategies over time. This is remarkable because the experimental design gives participants many opportunities to do so: before handing in a strategy for a new round, participants can try out strategies and run simulations against randomly chosen strategies of the previous round. The difficulty is that it is not sufficient to find a strategy that does well against strategies from the previous round, because other

[1]Here efficiency is measured by the average number of points generated by the strategies, which is high if the strategies succeed in coordinating often on outcomes where the minority consists of exactly two players.

participants will also adapt their strategies. Participants are therefore aiming at a 'moving target': they have to predict how the other participants will adapt their strategies and have to respond optimally to those predictions.[2]

The aim of the current paper is to investigate the absence of learning in the strategy experiment of Linde et al. (2014). We run an additional strategy experiment that differs in the information that is provided to the participants. In an evolutionary analysis there is an implicit assumption that actors know the strategies that are used by others and have a tendency to switch from unsuccessful to more successful strategies. That the fast increase in efficiency in the evolutionary analysis is not found over the five rounds of the original experiment may be caused by informational limits in that experiment. The new experiment, named *Information* hereafter, differs from the original experiment (*No Information*) in two ways. First, participants are shown the rank and performance of all submitted strategies of the previous round as well as the syntax of these submitted strategies. Second, before handing in their strategy, the participants have the possibility to run simulations against any strategies of their liking (including successful strategies of other participants from the previous round) and not just against random previous round strategies. We believe these design features facilitate learning: participants may copy elements of strategies that performed well in the previous round. In addition, they may use higher levels of rationality by constructing a strategy that does well against the strategies that did well in the previous round, or even a strategy that does well against strategies that do well against the best strategies from the previous round, etc.

We expected a higher average efficiency in *Information* than in *No Information*, but much to our surprise we find that this additional information leads to a strong *decrease* in average efficiency in rounds 2 and 3 (of the five-round experiment). This remarkable finding is in contrast with the (rather scarce) experimental literature on the repeated minority game, which up to this point has found that information has either no significant effect or a positive impact on efficiency.[3] The reason for our result is that participants tend to imitate the winning strategies from the previous round, or at least move in that direction. This results in

[2] Obviously this reasoning does not stop here: rational participants also have to predict how all other participants predict that all participants adapt their strategies, and so on. Laboratory experiments on the guessing game show that participants typically exhibit one or two of these levels of rationality (see Nagel (1995) and Ho et al. (1998)).

[3] For example, see Chmura and Pitz (2006) and Bottazzi and Devetag (2007).

strategies being too similar to each other, which is detrimental for efficiency as a player is more likely to be successful in the minority game if he or she behaves differently from most other players. Since many strategies are similar, there are only a few different strategies that benefit from them, implying that aggregate efficiency also decreases.

The remainder of the paper is organized as follows. In the next section we discuss the minority game and briefly review the computational and experimental literature on this game. Section 3 discusses the design of the experiment. In Section 4 we present our results. Section 5 provides some concluding remarks and the appendices contain the instructions of the experiment and some additional analyses.

2. The minority game

The minority game is introduced in Challet and Zhang (1997) as a stylized and symmetric version of Arthur's well-known *El Farol* bar game (1994). It involves an odd number of players *N*, who simultaneously have to choose one of two sides (say *Red* and *Blue*). The players that make the minority choice are rewarded with one 'point', the others earn nothing. In particular, if $s_i = 1$ when player $\boldsymbol{i}$ chooses *Red* and $s_i = 0$ when player $\boldsymbol{i}$ chooses *Blue*, payoffs for player $\boldsymbol{i}$ are given by

$$\pi_i(s) = \begin{cases} s_i & when \ \sum_{j=1}^{N} s_j \leq \frac{N-1}{2} \\ 1 - s_i & when \ \sum_{j=1}^{N} s_j \geq \frac{N+1}{2} \end{cases}.$$

Although the minority game is a relatively simple multi-player game (it is symmetric, players can choose from only two actions, which can lead to only one of two possible payoffs) it has many Nash equilibria. In particular, there exist $\binom{N}{(N-1)/2}$ pure strategy Nash equilibria where exactly (N-1)/2 players choose one of the actions, and the remaining players choose the other action. In addition, there exist infinitely many mixed strategy Nash equilibria. One of those is the symmetric mixed strategy Nash equilibrium, where every player chooses *Red* with probability $1/2$, but for example also any action profile where (N-1)/2 players choose *Red* with certainty, (N-1)/2 players choose *Blue* with certainty and the remaining player randomizes with any probability constitutes a mixed strategy Nash equilibrium. Note that none of the equilibria are strict. Moreover, the pure strategy Nash equilibria lead to a very asymmetric distribution of payoffs and many of the mixed strategy Nash equilibria may give rise to inefficiencies because the probability that the minority is smaller than (N-1)/2 is positive.

The minority game is a stylized representation of decision problems where people benefit from behaving in opposition to the majority of the population. It was initially interpreted as a model of speculative financial trading (see e.g. Challet et al. 2000, 2001). In that interpretation one of the two sides of the minority game is seen as buying and the other as selling: buyers earn money when there are few buyers, because this drives down the price, and sellers make a profit when there are few sellers, as they will be able to demand a higher price. This interpretation is often seen as too simplistic, partly because sometimes it is better to belong to the majority in financial trading.

However, the minority game also provides a useful representation of many other important economic decision problems. For example, the minority game is an example of a *congestion game*, where players make use of limited resources (e.g. driving on a road with limited capacity) and payoffs are determined by how many other players use that resource. The minority game is also closely related to the *market entry game* where each of a number of firms has to decide simultaneously whether or not to enter a (new) market, which will only be profitable if not too many firms enter. More generally the minority game is a stylized version of any game where actions are strategic substitutes, such as *Cournot oligopoly* games.

Although the minority game is highly stylized it is not straightforward to predict how the game will be played – particularly when it is repeated many times with the same set of players. There are many different equilibria on which players can coordinate, but even if the players solve this nontrivial coordination problem it is not apparent that the outcome will be sustainable, since payoffs will be very unequal and the players in the majority have an incentive to upset the status quo.[4]

[4] Note that in the repeated minority game there exist pure strategy Nash equilibria where players rotate over the two options in such a way that every player spends the same number of periods in the minority. Total payoffs would then be the same for each player. However, in the absence of the possibility of communication, it seems very hard to coordinate on such an equilibrium, even if the number of players is relatively small. For a folk theorem on the infinitely repeated minority game see Renault et al. (2005).

Following the approach of Arthur (1994) to the El Farol bar game, Challet and Zhang (1997, 1998) investigate the minority game by computer simulations. In these simulations the number of players is large and each players has a fixed set of strategies, randomly drawn from the set of all strategies with a fixed memory M . Such a strategy predicts the next winning side for each possible history of the past M winning sides. Note that these strategies do not use information about the *size* of the minority and that they do not allow for randomization. The agents then use that strategy (from their set of strategies) that, up to that period, would have been the most successful (without taking the effect of that strategy on the outcome into account). Numerical simulations show that the number of agents choosing one side fluctuates around 50%. The higher the volatility of fluctuations (implying that small minorities occur more often) the less efficient is the outcome. Simulations in which the sets of strategies – which are assumed fixed in the standard minority game – evolve under evolutionary pressure show that agents coordinate on aggregate outcomes that are more efficient (see e.g. Li et al. 2000a, 2000b; Sysi-Aho et al. 2005).[5]

The drawback of these computer simulations is that the implemented strategies do not necessarily represent the strategies that human players would use. Laboratory experiments, with paid human subjects, on the minority game may alleviate this concern somewhat. This experimental literature indicates that participants have heterogeneous behavioral rules and generally do not show equilibrium behavior at the individual level, see e.g. Chmura and Güth (2011) and Devetag et al. (2014). The experimental literature on the impact of information in minority games is relatively scarce, but the literature that does exist either finds no effect or a positive effect of information on efficiency. For the 15-player minority game studied in Platkowski and Ramsza (2003), for example, where participants know which side won in each of the previous M periods the value of M (which varies between three and eleven) does not have a significant impact on performance. Chmura and Pitz (2006) run a nine-player minority game with two treatments. In the first treatment, participants only have the information as to whether or not they belonged to the minority in the previous period, while in the second treatment they

[5] The 'econophysics' literature on the minority game is very rich. Challet et al. (2013) even consider the minority game as "more or less solved'" (Challet et al. (2013), p. 13). For much more detailed summaries of the minority game in econophysics see Challet et al. (2013) and Moro (2004). Also, http://www3.unifr.ch/econophysics/ (opened on 17/08/2014) contains many analyses, discussions and extensions of the minority game.

also receive information about how many players chose each side in the last period. The authors find that the additional information in the second treatment improves aggregate efficiency. Finally, Bottazzi and Devetag (2007) have participants play the finitely repeated minority game with stationary groups of five players under different information conditions. They find that information on individual choices within the groups of five has no effect on aggregate performance, but revealing information about more than just the previous round does.

Although these laboratory experiments are quite useful for understanding aggregate behavior of a group of people playing the minority game, the large strategy space makes it difficult to infer individual strategies from the human subjects. In Linde et al. (2014) an experiment on the minority game using the strategy method is presented. That is, participants are explicitly asked to formulate a strategy that specifies how to play the five player minority game for 100 periods, that is, which action to take for every possible decision node of the game.[6] These strategies take part in a computer tournament and participants have the possibility to revise their strategy between rounds. However, they do not seem to learn to improve their strategies over the different rounds and aggregate efficiency in each round of the strategy experiment is close to that of the symmetric mixed strategy Nash equilibrium. On the other hand, an evolutionary competition between all submitted strategies increases average efficiency to a level substantially above that in the experiment itself, which is consistent with models of evolutionary competition in the econophysics literature discussed above. The experiment discussed in this paper is specifically designed to investigate this friction between increased efficiency in the evolutionary competition, and the absence of such an improvement over the five rounds of the experiment in Linde et al. (2014).

[6] The strategy method has been applied before to related games, such as cobweb markets (Sonnemans et al. (2004)), predictions in asset markets (Hommes et al. (2005)), market entry games (Seale and Rapoport (2000)) and the El Farol bar game (Leady (2000)). Brandts and Charness (2011) provide a good overview on possible advantages and disadvantages of the strategy method. The major potential point of criticism of strategy method experiments is the possibility that the strategy method could lead to behaviorally different decisions than the 'direct-respond' method. From a game-theoretic point of view, the methods should not make a difference. Roth (1995), for example, states that "having to submit entire strategies forces subjects to think about each information set in a different way than if they could primarily concentrate on those information sets that arise in the course of the game" (p.323). However, Brandts and Charness (2011) find that the strategy method and the direct-response method lead to qualitatively similar results.

3. Experimental Design

We designed two experiments in which participants have to submit a strategy to play the five-player minority game for 100 periods. Both experiments consist of five rounds, each separated by a week. The first experiment (*No Information*) was conducted in April 2009 and the second experiment (*Information*) in April 2010. Participants in both experiments are students of subsequent cohorts of the so-called "beta-gamma" bachelor program, one of the most challenging programs of the University of Amsterdam and 42 and 43 students participated, respectively.[7] Note that experiment *No Information* is presented and discussed in Linde et al. (2014). Additional information and analyses for that experiment can be found there.

The first round of each experiment differs from the following four. That round takes place at the CREED laboratory at the University of Amsterdam. First the participants choose a nickname and password so that they can log in to the experiment's website. Subsequently the minority game is explained to the participants and they play the game twice for ten periods in two different groups of players. In the next stage the participants are instructed on a handout and via a computer screen on how to formulate a strategy. To test their understanding of the syntax and the interface, they program two verbal strategies and they then formulate, test and submit their first strategy.[8] Within a couple of days, participants receive feedback on the first round. After receiving feedback they can login in to the website whenever they want and try out strategies of their own making against other strategies. Their definitive strategy for the next round (which can also be the same as in the previous round) has to be submitted within a week after the laboratory experiment. Two days after the deadline the participants receive the results of the second round. This procedure is repeated for rounds 3 to 5. After the fifth round the goals of the experiment are explained, the results of the final round are announced and all earnings of the experiment are paid out.

[7] These students follow courses in the natural sciences as well as the social sciences and they are typically well above average in motivation and capabilities. In particular, their programming experience is substantially higher than that of the average undergraduate student at the University of Amsterdam.

[8] Appendix A contains the English translation of the instructions for both experiments. The participants could ask the experimenters additional questions about, for example, formulating a strategy during the laboratory experiment for the first round or by e-mail for the later rounds.

The difference between the two experiments is that in the *Information* experiment participants are fully informed about the syntax of the strategies submitted by the other participants in the previous round, and that they can use these strategies to simulate against. More specifically: in experiment *No Information* the feedback participants receive after each round consists of the performance of all strategies (by the nickname of the participant) and the ranking of the strategies, whereas in the experiment *Information* participants are additionally informed about the exact strategies submitted by each participant. These strategies are coupled with the nicknames of the participants, so that participants also know how well those strategies performed.

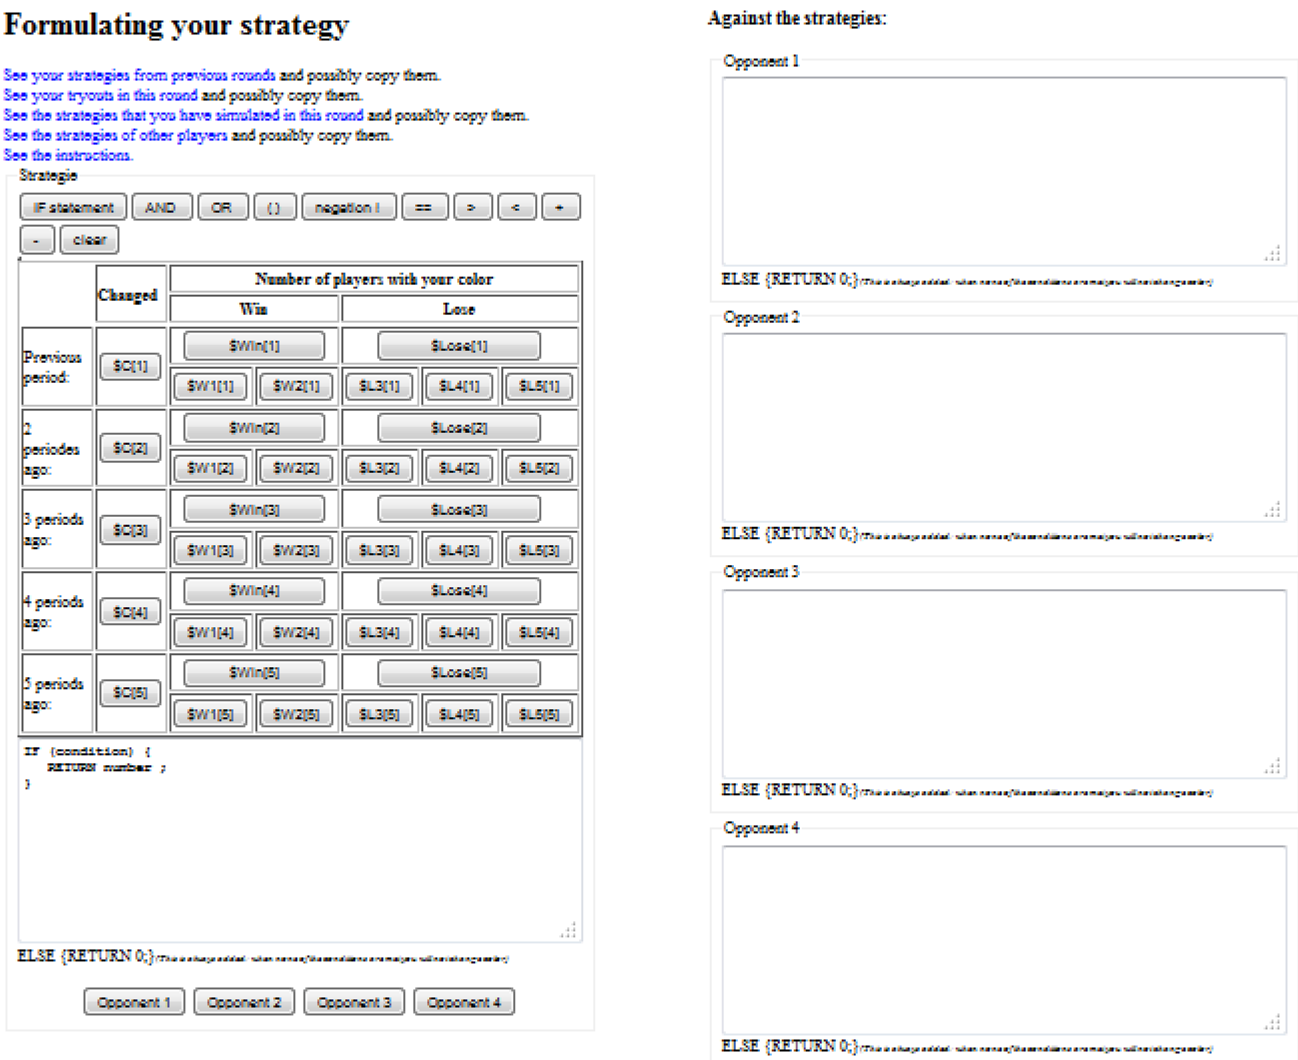

Figure 1: Computer screen as seen by the participants when they formulate a strategy in experiment *Information* (translated from Dutch)

3.1 Formulating a strategy

Figure 1 shows the computer screen as seen by the participants in experiment *Information* when formulating a strategy, for which they use the left-hand side of the screen.[9] A strategy consists of a list of IF-statements, each of which returns a number *p* in the interval [0, 1], provided the condition in the IF-statement is met. The number *p* is the *probability of changing color*. If the condition in an IF-statement is fulfilled, the subsequent IF-statements are ignored (that is, the second and following IF-statements are treated like ELSE IF statements). If none of the conditions are met, the strategy returns 0 (i.e. no change of color). There is no limit to the number of IF-statements a participant can use for his/her strategy.[10] In addition, the strategy can contain logical expressions such as AND, OR, (in)equality and negation. In the instructions ample examples were given. The strategies can use the history of the last five periods, which consists of the outcome in each of these previous periods (i.e. the size of the group that chooses the same color as the participant's strategy) and whether the strategy changed colors in that period or not.[11]

3.2 Simulations by participants

An important feature of our design is that participants can run simulations with a strategy of their own making. In experiment *No Information* simulations are ran with four randomly drawn strategies (without replacement) from other participants from the previous round.[12] Since strategies can use a history of up to five periods, the first five outcomes are randomly drawn. After that, 100 periods are played according to the five strategies. After each simulation the results of the 100 periods, as well as those of the first five random periods are presented.[13] In addition summary statistics are displayed for the 100 periods: the total number of points and the number of times the

---

[9] The experiment is programmed in php/mysql and runs on a (Apache) web server.

[10] The total length of a strategy is capped at the very high number of 1000 characters.

[11] The strategy space is therefore restricted in two ways. First, we reduce the number of variables per period by imposing symmetry between colors: strategies decide on *changing* color instead of *choosing* a color. Second, strategies cannot condition upon information from more than five periods ago. We believe that this still gives a sufficient amount of flexibility for participants to develop strategies.

[12] In the first round no strategies from participants are available. The participants are informed that the strategies they compete against in the simulations they run in the first round are pre-programmed and are not necessarily similar to the strategies the other participants will submit. There are eleven pre-programmed strategies that do not condition on the history of outcomes and change with probability *p*, where $p \in \{0, 0.1, \ldots, 0.9, 1\}$

[13] See Appendix B for an example of the computer screen a participant sees after running a practice simulation in experiment *No Information*.

outcome was in category W1, W2, L3 L4 and L5, respectively, where W1 (W2) represents winning in a group of 1 (2) and L3 (L4, L5) represents losing in a group of 3 (4, 5). Participants can run as many simulations and try as many strategies as they want. They can use these simulations to see how successful their strategy is, but also to check whether their strategy behaves as they intended it to.

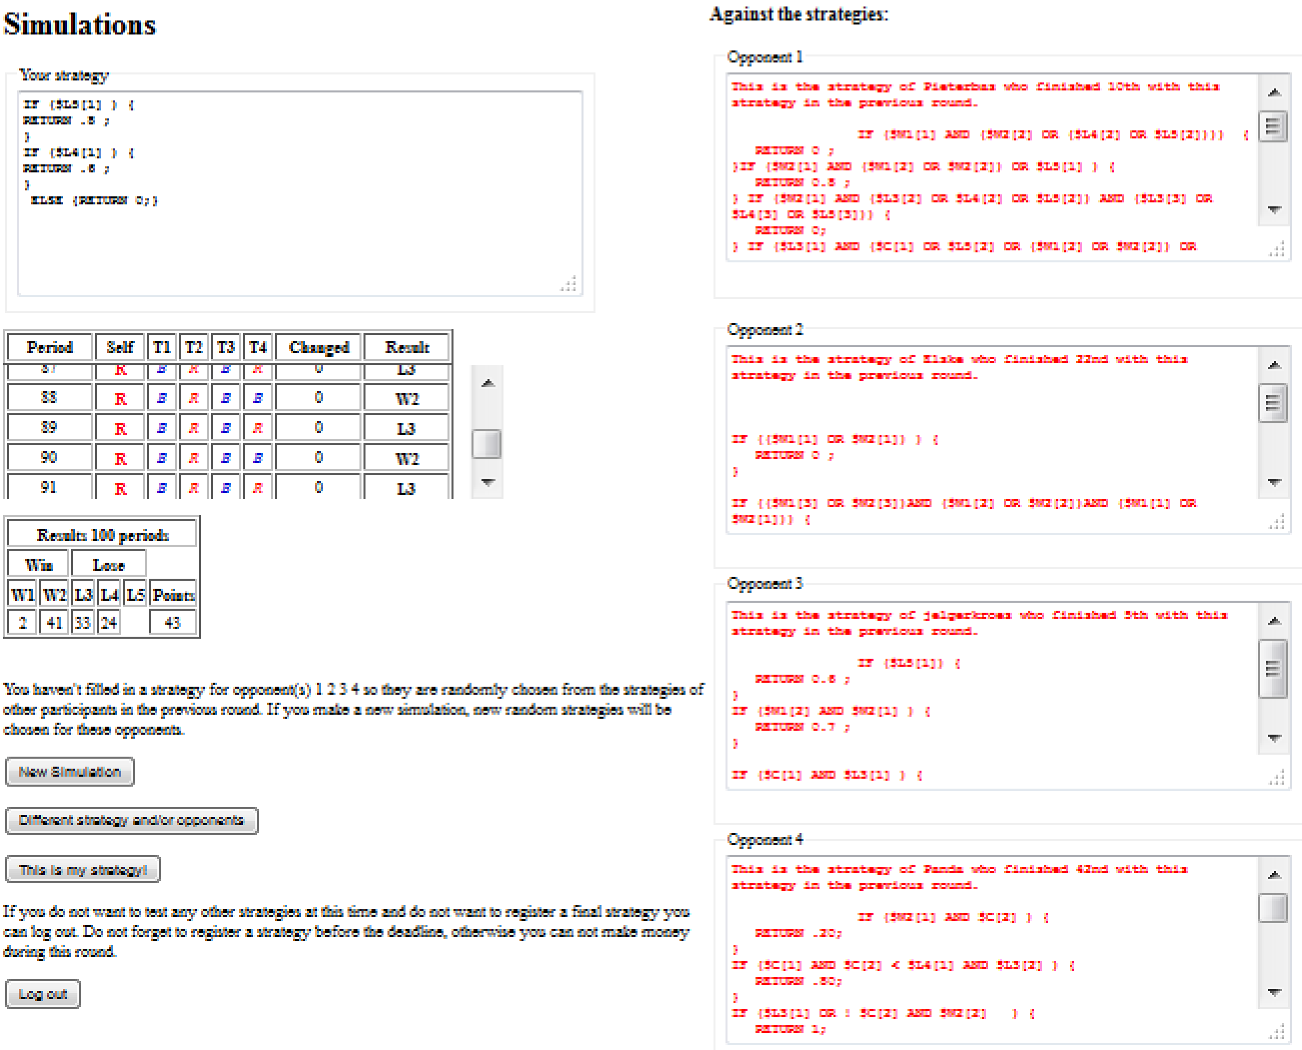

Figure 2: Example of a screen after a strategy is tried out by the participants in experiment *Information* (translated from Dutch).

For experiment *Information*, where participants know exactly which strategies the other participants submitted in the previous round, simulation possibilities are extended. In particular, participants can choose precisely against which opponent strategies they want to simulate their own strategy. They could type or copy and paste their own strategy into the left-hand box and the four opponent strategies in the right-hand boxes of Figure 1. The interface enabled them to easily copy and paste their own strategies from the previous rounds, their practice strategies from this round and the previous round strategies of all participants (they were also shown the ranking of these strategies in the previous round) in any of these five boxes.

If participants did not enter a strategy into a right-hand box, a random previous round strategy was used automatically, like in experiment *No Information*. Note that the interface in experiment *Information* is different from that of experiment *No Information*, as in the latter only the left hand side of the screen depicted in Figure 1 is presented. Simulation results are also presented somewhat differently (Figure 2 gives an example for experiment *Information*). In particular, for experiment *Information* the participant sees against which strategies she has simulated (independent of whether the participant chose these strategies herself, or whether they were randomly drawn by the computer program) and she sees the decision for each strategy (where in experiment *No Information* these decisions were ordered by color).

3.3 Computer tournament, earnings and questionnaire

For each round we run a computer tournament with all final strategies. A simulation is run for each possible combination of five strategies. One simulation consists of randomly selected outcomes in the first five periods and 100 consecutive periods where the five strategies are executed. These 100 periods are then used to calculate the points achieved by each of the five strategies in this round. We then calculate the average points earned by each strategy in all simulations and use these to rank the strategies. After each round, all participants receive an email containing the rank and the average points of each participant. In experiment *Information* they could additionally see each participant's strategy (and each strategy's performance) when they logged in. In both experiments participants learn their earnings for that round. Participants that submitted the five best strategies in that round receive €75, €60, €45, €30 and €15, respectively.[14] In addition in each round every participant who submits a strategy and fills out a short questionnaire[15] receives €5.[16]

[14] It might be argued that due to the tournament incentive participants would try to maximize their ranking instead of their total number of points. However, these two different objectives are aligned since strategies that bear a cost in terms of points in order to do relatively well in one particular simulation by making the situation worse for the other four strategies in that simulation, will hurt their performance relative to the 30 to 35 strategies that are not present in that simulation.

[15] The questionnaire contains questions about the background of the participants (like age, gender and programming experience) and questions about the (formulation of the) strategy: how difficult it was to

4. Results

As explained above, the first round of the two experiments took place in the CREED experimental laboratory and started with 42 participants for experiment *No Information* and 43 participants for experiment *Information*. All of these participants submitted a strategy for the first round. In the subsequent four rounds the number of submitted strategies was between 30 and 40 strategies for each experiment (see Table 1). Average earnings for the whole experiment are €58.70 (€58.37) per participant in experiment *No Information* (experiment *Information*), ranging from a minimum of €0 to a maximum of €260 (€200) in experiment *No Information* (experiment *Information*). In Section 4.1 we discuss aggregate outcomes and compare them between experiments and rounds. In Sections 4.2 and 4.3 we analyze participants' simulation and imitation behavior, respectively. Finally, in Section 4.4 we present a cluster analysis of the strategies submitted in experiment *Information* and investigate the interaction between the different clusters, in order to understand the dynamics of the evolution of strategies over the different rounds.

4.1 Aggregate outcomes

For the repeated five-player minority game that we are considering, the most efficient outcome, where the minority consists of exactly two players in every period, is obtained in any pure strategy Nash equilibrium. The symmetric mixed strategy Nash equilibrium, on the other hand, may lead to inefficient outcomes since randomization implies that there is a positive probability that the minority is smaller than two. In fact, in that symmetric mixed strategy Nash equilibrium the probability that exactly two players are in the minority is 62.50%, with the probabilities of minorities of 1 and

---

formulate the strategy, whether they had any problems with the formulation and how confident they are that the strategy will be successful.

[16] For experiment *No Information* there are two slight changes with respect to earnings. First, although payoffs for rounds 1 to 4 are identical, prizes where twice as high in the final round of experiment *No Information* (€150, €120, €90, €60 and €30, respectively). The main motivation for this was to ensure that participation in the experiment remained high. We did not give this extra incentive in experiment *Information* and its effect seems to be small, as the decrease in the number of active participants is comparable for the two experiments. The second difference is that in experiment *No Information* we gave the participants the possibility to wager the €5 they could earn with filling out the questionnaire. If a participant chooses to give up these €5 for a round, the participant would earn an extra reward in that round, provided the strategy ends up in the top five, of €75, €60, €45, €30 or €15 euro, respectively, in that round. This bet was introduced to test participant's confidence in their own strategy. The decision of participants to forego the €5 was only significantly positively correlated with performance of the strategies in round 4 of experiment *No Information* and not in any of the other rounds, see Linde, Sonnemans and Tuinstra (2014), p.85.

0 equal to 31.25% and 6.25%, respectively. It follows that in a pure strategy Nash equilibrium the average number of points per player over 100 periods is equal to 40, whereas the expected average number of points per player is 31.25 in the symmetric mixed strategy Nash equilibrium.[17]

In order to determine the performance of the submitted strategies we run one simulation of 100 periods for each possible combination of five submitted strategies.[18] Table 1 shows, for each round of each experiment, the average number of points earned by the strategies in these simulations.

	No Information					*Information*				
	Round					*Round*				
Points	*1*	*2*	*3*	*4*	*5*	*1*	*2*	*3*	*4*	*5*
Average	31.69	31.12	31.88	32.29	31.85	31.43	29.57	26.32	31.41	32.84
Min.	29.31	21.65	24.98	28.00	18.93	21.35	13.74	8.47	20.64	26.82
Max.	34.68	41.65	36.96	39.87	43.06	39.66	47.61	48.66	47.98	41.92
St.dev.	1.49	5.27	3.10	3.20	6.36	4.54	8.39	12.44	7.12	4.08
N	42	36	34	36	32	43	40	32	35	30

Table 1: Aggregate outcomes and performance of strategies

Several observations stand out from Table 1. First, at the aggregate level, the symmetric mixed strategy Nash equilibrium gives a much better description of the results than the pure strategy Nash equilibrium does: in most of the rounds the average number of points is relatively close to 31.25, and quite far away from 40. Second, in neither experiment were the participants able to improve efficiency over the rounds substantially, so there seems to be limited learning. Third, giving the participants additional information about strategies submitted by the other participants does not improve efficiency. In round 1 information for the participants is the same in the two experiments (the only difference is that participants in experiment *Information* could program their own opponent strategies) and we therefore expect the difference in that round to be limited. However, for three of the next four rounds the average number of points in experiment *Information* is lower than the average number of points in the corresponding round of experiment *No Information*. In particular, the average number of points in rounds 2 and 3 of experiment *Information* are surprisingly low – even substantially below the number of points under the symmetric

[17] We have $(0\times 0.0625+0.2\times 0.3125+0.4\times 0.625)\times 100=31.25$.

[18] That is, if *N* strategies are submitted in a round, we run a total of $\binom{N}{5}$ simulations of 100 periods in that round. Moreover, each submitted strategy is simulated in $\binom{N-1}{4}$ different compositions of five submitted strategies, that is, the number of simulations each strategy participates in lies between 23,751 (round 5, experiment *Information*) and 111,930 (round 1, experiment *Information*).

mixed strategy Nash equilibrium. Finally, although the symmetric mixed strategy Nash equilibrium seems to give a reasonable description of the aggregate outcomes, it does not perform very well as a description of the individual strategies: the dispersion of the average number of points generated by the strategies, as measured either by their standard deviation, or by the range between the minimum and maximum number of points, is much higher than under the symmetric mixed strategy Nash equilibrium.[19] Heterogeneity of the submitted strategies – for example in terms of the complexity or length of the strategy, the length of the history considered and the strategies' propensity to change colors – turns out to be substantial, see Linde et al. (2014) for a discussion. Table 1 suggests that heterogeneity of strategies in experiment *Information* is even higher than in experiment *No Information*, apart from round 5.

In this paper we focus on the second and third observation discussed above: absence of learning over the rounds in both experiments and the fact that additional information about other participants' strategies has a non-positive, or even negative, effect on performance.

Round	***New** strategy does better in **old** environment*	***New** strategy does better in **new** environment*
No information		
2	42.9%	50.0%
3	77.8%	38.9%
4	66.7%	80.0%
5	38.9%	50.0%
No information total	**54.4%**	**53.1%**
Information		
2	80.0%	30.0%
3	87.5%	6.3%
4	69.0%	20.7%
5	59.3%	51.9%
Information total	**75.0%**	**26.6%**

Table 2: The performance of the old and the new strategies in the old (column 2) and the new (column 3) environment. Strategies from participants who handed in the same strategy as in the previous round are dropped from the analysis.

[19] In the symmetric mixed strategy Nash equilibrium there is very little dispersion in average payoffs for participants because most random variation will disappear when each strategy is involved in at least 23,000 simulations of 100 periods.

The impression that there is limited learning over the rounds in both treatments is corroborated by Table 2. For this table we ran, for each of the last four rounds of each experiment, additional simulations to see whether participants are successful in improving their strategy over the rounds. First we ran simulations with each participant's new strategy (for that round) against the strategies of the other participants from the previous round. The second column of Table 2 shows the percentage of new strategies that outperform the strategy that the same participant submitted in the previous round. Note that for experiment *No Information* participants could run simulations with their new strategy against randomly drawn old strategies from the previous round, and one would therefore expect that typically the new strategies outperform the old ones. Surprisingly, however, this turns out to happen in only about half of the cases. For experiment *Information* a much larger fraction of participants (75% over all rounds) submit a strategy for the new round that does better against the strategies from the previous round than their old strategy did. This is not very surprising, however, since in experiment *Information* participants know all submitted strategies of the previous round and also know which of these strategies won that round, and can therefore imitate that strategy. That strategy is then very likely to perform better than the participant's new strategy against the old population of strategies.[20] In fact, the very high percentages of 80.0% and 87.5% in rounds 2 and 3 of experiment *Information* are consistent with a substantial number of participants imitating the winning strategy of the previous round.

Although it is interesting to investigate whether the new strategy does better than the old strategy in the old population of strategies, the task of the participants is to submit a strategy that does well in the new population of strategies. Because many of the strategies change between rounds participants should try to predict how the other strategies change between rounds: strategies that performed superbly in the previous round might perform very badly in the current one. Therefore, in the third column of Table 2 we present the percentage of participants in each round whose new strategy does better against the new population of strategies, than their old strategy would have done in this new population of strategies. If this percentage is high it means that participants are successful in predicting how the other participants

[20] Note, however, that it is not guaranteed to perform better because its performance in the previous round may have been good partly because it was successful in exploiting the strategy it now replaces.

changed their strategy and in developing a strategy that does well in this new population.

For experiment *No Information* the new strategy performs better than the old strategy in about half (53.1%) of the cases, which suggests that participants are not particularly successful in improving their strategy (the only positive exception is round 4 where 80% of the submitted strategies do better than the old strategies would do in the new population of strategies). However, the situation is much worse for experiment *Information* where only 26.6% of the new strategies perform better than the strategy from the previous round against the new strategies. It is particularly alarming that 93.7% of all participants in experiment *Information* who changed their strategy from round 2 to round 3 would have done better in round 3 if they had stuck to their round 2 strategy instead.

The findings presented in Table 2 suggest that participants do not to take sufficiently into account that other participants will change their strategies as well and focus too much on doing well against the strategies from the previous round. In both experiments a strategy that does well against the old population of strategies could be found by simulating against random strategies from the previous round. In addition, in experiment *Information*, participants could also simulate against strategies of their own choice, simply imitate (and maybe slightly adapt) the best strategy of the previous round, or try to find a strategy that performs well against the best strategies of the previous round. However, the results suggest that the extra possibilities given to the participants in experiment *Information* actually lead to strategies that perform worse, at least in rounds 2 and 3.

There may be several reasons for this deterioration in performance. For example, it is possible that strategies of participants who simulated against the best strategies of the previous round (which was only possible in experiment *Information*) are different from (and inferior to) strategies of participants who only simulated against random strategies (possible in both experiments). Alternatively, performance may have been brought down because the most successful strategies of the previous round are imitated (which is also only possible in experiment *Information*). These two explanations will be analyzed in Section 4.2 and Section 4.3, respectively.

However, before we analyze simulation behavior and imitation, we first briefly discuss whether a selection effect has played a role. This might have been the case since the number of participants fluctuates considerably across rounds of the experiment. This selection effect, according to which participants who did not perform well in the previous round may decide not to submit a strategy for the current round, would imply that aggregate efficiency in the later rounds is higher than would have been the case if all participants submitted a strategy. Given that aggregate efficiency does not go up substantially, or even decreases, the selection effect is unlikely to be large. Moreover, we compared, for each of the first four rounds, the performance of the participants who did submit a strategy in the next round, with those who did not and we found hardly any difference in the performance for these two types of participants. [21] We therefore conclude that the selection effect only played a minor role, if any.[22]

4.2 Simulation

In both experiments participants can run practice simulations with the strategies they are considering to submit. In experiment *No Information* these simulations are run against random strategies that the other participants submitted in the previous round (without replacement). In experiment *Information* the simulation possibilities are much broader. In particular, participants can replace one or more of the four randomly drawn strategies from the previous round with strategies of their choice (either of their own making or a specific strategy from the previous round). In this section we analyze to what extent these increased simulation possibilities are used by the participants, and what the impact on the performance of the submitted strategies is. For round 1, participants can only run strategies against preprogrammed strategies (or strategies of their own making in experiment *Information*). Therefore, our analysis focuses on simulation behavior for rounds 2 to 5, which is summarized in Table 3. The rows indicated by "# Simulations" give the total number of simulations run by the 'Simulators' (those participants who did run simulations) including the final

[21] See Appendix C for details.

[22] The selection effect can be studied in strategy method experiments by running simulations with, in addition to the strategies actually submitted for the current round, the last submitted strategy of participants who did not hand in a strategy for this round. This approach is troublesome here. The analysis in the following sections will show that participants in this experiment typically change their strategies substantially between rounds. Therefore, it is also likely that participants not submitting a strategy would have changed their strategy if they would have submitted one for the current round. Using their old strategies instead will then not capture the selection effect.

strategy they submitted. Although in both experiments participants make use of practice simulations, the number of participants running simulations decreases (in absolute as well as in relative terms) over the rounds, from 33 (*No Information*) / 35 (*Information*) in round 2 to only 21 (*No Information*) / 19 (*Information*) in round 5. For experiment *Information* Table 3 also shows whether participants simulated against strategies that were among the five best strategies of the previous round (and thus earned a prize), which we will refer to as Top5 strategies. In round 2, 14 participants simulated at least once against one of the Top5 strategies, but this decreased to five participants in each of the subsequent rounds (corresponding to around one quarter of the participants who simulated). If participants simulate against Top5 strategies, they typically do so against the winner of the previous period. In addition, Table 3 shows the shares of Top5 strategies of all opponent strategies used in the simulations. Random strategies from the previous round make up 78.6% of all opponent strategies, and Top5 strategies 14.8%. The fraction of opponent strategies that is random increases in rounds 4 and 5. We therefore find that relatively few participants in experiment *Information* simulate against Top5 strategies, and that these Top5 shares only make up a small share of all opponent strategies. Moreover, both the fraction of participants simulating at least once against a Top5 strategy, as well as the fraction of opponent strategies that is a Top5 strategy decreases over the rounds.

Although participants who run at least one simulation against a Top5 strategy of the previous round typically perform better than the other participants who simulate, and participants who simulate perform slightly better than participants that do not, we find for all rounds that none of these differences is significant at the 5% level.[23]

Summarizing, we conclude that participants do not make substantial use of the increased simulation possibilities in experiment *Information*, and that, if they do, it does not have a significant effect on the performance of the strategy they submit. The differences between aggregate efficiency that we found in Subsection 4.1 (in particular for rounds 2 and 3) can therefore not be explained by the increased simulation possibilities.

---

[23] See Appendix D for details.

	Round 1	*Round 2*	*Round 3*	*Round 4*	*Round 5*	*Total*
No Information						
Participants	42	36	34	36	32	180
Simulators	42	33	23	23	21	142
# Simulations	470	491	354	346	335	1996
Information						
Participants	43	40	32	30	35	180
Simulators	43	35	22	22	19	141
# Simulations	909	661	242	116	152	2090
Against random	79.4%	74.9%	73.6%	90.3%	93.9%	78.6%
Against Top5	--	17.2%	18.4%	7.3%	4.4%	14.8%
# Simulators	--	14	5	5	5	
Against #1	--	6.5%	7.2%	3.7%	2.3%	5.8%
# Simulators	--	13	5	4	3	

Table 3: Simulation Behavior

4.3 Imitation

In this section we investigate whether participants have a tendency to imitate the better strategies from the previous round. In order to compare strategies we use the distance measure introduced in Linde et al. (2014). Note that a strategy attaches a probability of change to each possible history of five periods. Since in each period there can be five outcomes (W1, W2, L3, L4 or L5) and in each period the strategy either changes color or not, there are 5^5 x 2^5 = 100,000 possible histories. For every strategy the probability of changing color is determined for each of these possible histories. The distance $d(x, y)$ between strategies x and y then is defined as the weighted average absolute difference between these probabilities.[24] Because not all histories are equally likely the weights are based upon the distribution that would result from the symmetric mixed strategy Nash equilibrium.[25] We will use the distance measure $d(x, y)$ to determine whether the strategies in round $t + 1$ have become more or less similar to a particular strategy s from round t. More precisely,

[24] This is a continuous version of the Hamming distance.

[25] As discussed in Section 4.1 the symmetric MSNE leads to the outcomes 5-0 in 6.25%, 4-1 in 31.25% and 3-2 in 62.5% of the periods, which is quite close to the distribution in the experiment (see Linde et al. (2014)) and therefore gives a good approximation of the distribution of histories a strategy encounters.

let S_t^k be the set of strategies submitted in round t of experiment $k \in \{NI, I\}$ and define, for each strategy $s \in S_t^k$, the **change in average distance to $\boldsymbol{s}$** as

$$\Delta_t^k(s) = \frac{1}{|S_{t+1}^k|} \sum_{x \in S_{t+1}^k} d(s,x) - \frac{1}{|S_t^k|} \sum_{z \in S_t^k} d(s,z).$$

Note that if $\Delta_t^k(s) > 0$ strategies submitted in round $t+1$ (of experiment k) are on average less similar to strategy $s \in S_t^k$. That is, in that case on average strategies in round $t+1$ shift away from strategy s from round t. On the other hand, if $\Delta_t^k(s) < 0$ strategies submitted in round $t+1$ are more similar to strategy s. The latter is consistent with strategy s from round t being imitated in round $t+1$.

Figures 3 and 4 show the measure $\Delta_t^k(s)$ for each strategy s in rounds 1, 2, 3 and 4, for experiment *No Information* and experiment *Information*, respectively. We order strategies on the horizontal axis by their ranking in the computer tournament of that round (that is, the first strategy on the horizontal axis is the winner of that round).

Figure 3 suggests that there is a weak but positive relationship between the similarity of a strategy to next round's strategies and the rank of that strategy. The reason for this might be that participants who submitted a successful strategy may tend to keep this strategy for the next round, or only adapt it slightly, whereas participants with a strategy that did not perform well are inclined to abandon their strategy altogether, and develop a completely new strategy, possibly by simulating against random strategies from the previous round. This may well lead to a strategy that is more similar to the better strategies of the previous round.

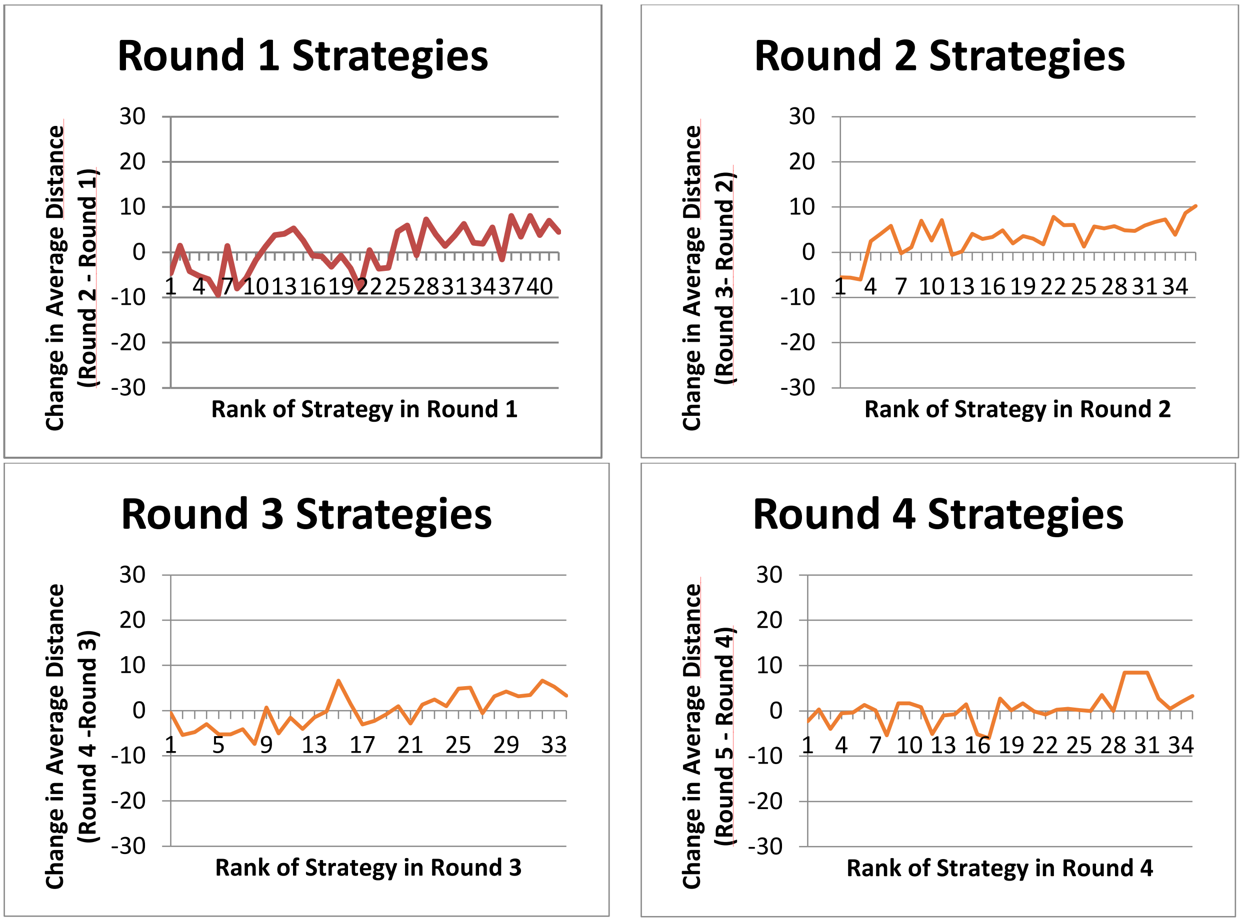

Figure 3: Change in Average Distance in experiment *No Information*.

In contrast to experiment *No Information*, in experiment *Information* participants can observe the strategies of the other participants from the previous round. This facilitates both imitation of successful strategies, and moving away from unsuccessful ones. We therefore expect the relationship between similarity of a strategy to next round's strategies and the rank of that strategy in the current round to be stronger than for experiment *No Information*. This conjecture is confirmed by Figure 4 which indeed suggests a strong positive relationship in all rounds except round 4. Strategies from the previous round that performed poorly (i.e. that have a high rank) have high values of $\Delta_t^I(s)$, which implies that participants shift away from those strategies. In addition, strategies from the previous round that performed well have negative values $\Delta_t^I(s)$, which is consistent with imitation of these strategies. In particular, for each round $\Delta_t^I(s)$ is negative for the two best strategies, and strongly so for the best strategy, and also smaller than the corresponding changes in average distance for experiment *No Information*.

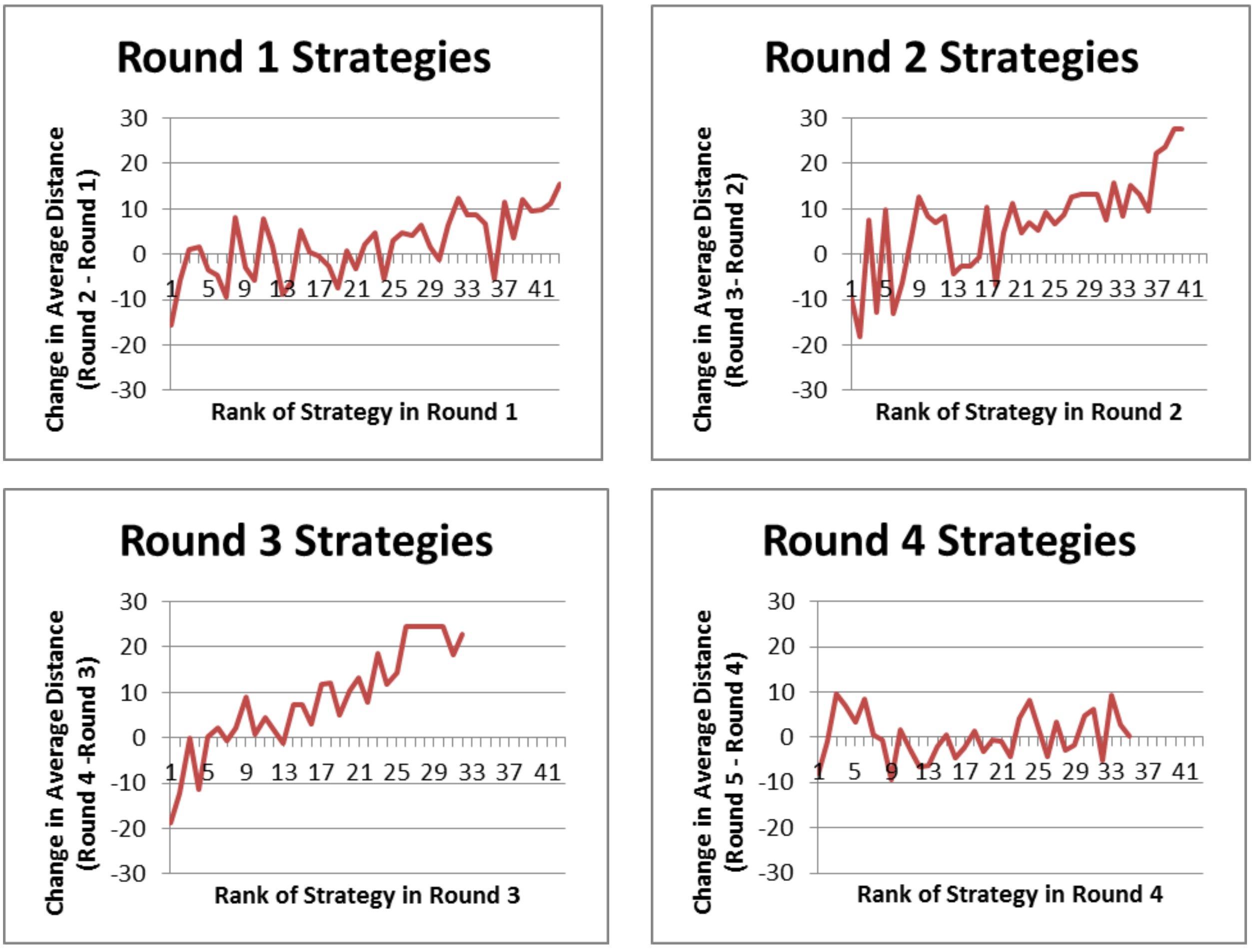

Figure 4: Change in Average Distance in experiment *Information*.

To confirm that the relationship between rank and change in average distance to next round's strategies is stronger for experiment *Information* we ran a regression with the change in average distance, Δ, as dependent variable and the rank of the strategy in the relevant round, r, and a dummy variable D_I which equals 1 for experiment *Information* and 0 for experiment *No Information*, as explanatory variables. The estimated relation is

$$\Delta = -3.279 + 0.233r - 1.837D_I + 0.243D_I \times r \qquad (1)$$

The coefficient on D_I is the only one *not* significant at the 5% level.[26] Equation (1) shows that the relation between the rank and the change in average distance is indeed significantly stronger in experiment *Information*: the slope of the relation between Δ and r equals 0.233 for experiment *No Information* and 0.476 for experiment

[26] The standard errors of the four coefficients are 0.973, 0.044, 1.364 and 0.062, respectively.

Information. Information about other participants' strategies therefore indeed substantially increases participants' tendency to shift away from poorly performing strategies, and move towards the best performing strategies.

4.4 Classification of strategies

The analysis from Subsection 4.3 suggests that participants tend to imitate the best strategies from the previous round and shift away from the worst strategies, in particular in experiment *Information*. In this subsection we use a cluster analysis of the 124 unique strategies submitted in that experiment to understand the impact that imitation may have had on the performance of the strategies in the different rounds of the experiment.[27] We consider the matrix of weighted distances between all unique strategies of experiment *Information* and used the program "multidendrograms" [28] to draw the dendrograms, using the clustering algorithm "WARD". The horizontal axis of Figure 5 lists all unique strategies and the vertical axis shows a measure of distances between the strategies within a cluster. For the analysis in this paper, we use the five clusters highlighted in Figure 5.

[27] Strategies were defined as duplicates if the two compared strategies decide to change with the exact same probability in all possible situations. For a cluster analysis of the strategies in experiment *No Information*, see Linde, Sonnemans and Tuinstra (2014)

[28] http://deim.urv.cat/~sgomez/multidendrograms.php (opened on 17/08/2014). See also Fernández and Gómez (2008).

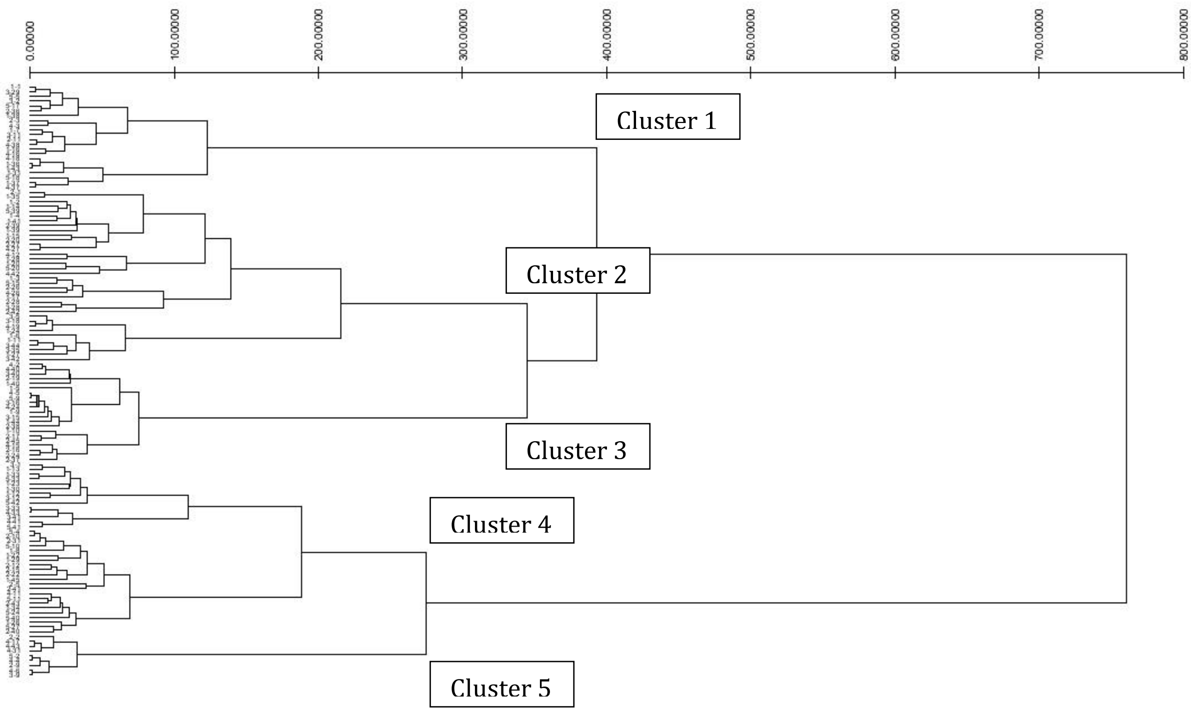

Figure 5: Cluster analysis of all 124 unique strategies in treatment I.

Table 4 gives some characteristics of the five different clusters. Strategies in **Cluster 1** generally change with a relatively high probability (but not certainty) after losing, and typically only change with a very low (or zero) probability when winning in one of the previous periods. **Cluster 2** strategies change sides very often, and this cluster includes the strategy that always changes after losing in the previous period. Strategies in **Cluster 3** roughly resemble the symmetric mixed strategy Nash Equilibrium. **Cluster 4** contains strategies that rarely change. These strategies often consider a lot of periods, winning and losing situations a well as whether they have previously changed. However, they generally either assign a small change probability if a condition is met or conditions for changing are relatively seldom met. **Cluster 5** is the smallest cluster and contains strategies that generally change with a high probability if they won in the last period (or sometimes in one of the last two periods). Losing and changes in previous periods barely lead to changes.

	Cluster 1	*Cluster 2*	*Cluster 3*	*Cluster 4*	*Cluster 5*	*Unique Strat.*
# of strategies	22	36	21	36	9	124
Randomization	91%	67%	100%	83%	67%	81%
# of periods considered	2.32	2.58	2.33	3.08	1.44	2.56
Consider winning situations	55%	92%	95%	78%	100%	82%
Consider losing situations	100%	94%	86%	92%	22%	88%
Consider whether you changed	27%	61%	19%	64%	22%	46%
Average Change Propensity	45.30 (10.12)	65.45 (15.89)	50.09 (8.01)	22.65 (11.32)	25.00 (6.51)	41.74 (23.00)
Description	Lose→ Probably (but not certainly) Shift	Change a lot. Including: Lose→ Certainly Shift	Similar to MSNE	Hardly Change	Win→ Shift with high prob. or certainty	
Most central strategy	If you lost the last period change with *p=0.6*	If you won in a group of 1 in the last period or lost in a group of 3 in the last period *p=1*; else if change with *p=0.75*	If you won in a group of 2 in each of the last two periods change with *p=0.6*; else if change with *p=0.5*	Never Change	If you won in the last period change with *p=0.8*	

Table 4: Characteristics of the clusters in experiment *Information*.

Our next step is to run additional simulations with all unique strategies of experiment *Information* to understand how the performance of clusters is influenced by strategies of the own or other clusters. Table 5 summarizes the simulation results.[1]

	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5
Within Sim.	26.77	31.92	30.45	35.57	0.59
Homo. Sim.	23.89	22.43	22.56	26.74	0.71
Cluster 1	−	−		+	+
Cluster 2					+
Cluster 3					
Cluster 4	+			+	
Cluster 5	+	+		−	−

Table 5: Interaction within and between clusters.

The second row of Table 5 (Within Simulations) shows the average number of points strategies from a cluster earn when playing only against other strategies from their own cluster and the third row (Homogenous Simulations) shows the average number of points when strategies (from a certain cluster) play against four copies of themselves.

These simulations show that cluster 4 strategies do relatively well against strategies of their own cluster, but that strategies from cluster 1 and, in particular, cluster 5 do quite badly against strategies from their own cluster. The remarkably bad performance of cluster 5 strategies is caused by the fact that these strategies switch when winning but rarely switch when losing, implying that they get locked into a five player majority quickly. All clusters perform poorly in Homogenous Simulations (in particular cluster 5 again, for the same reason as above), which shows that strategies are not very well equipped to play against themselves, and are instead designed to exploit other strategies.

Rows 4 to 8 of Table 5 indicate how strategies from different clusters interact with each other. In particular, it shows the effect of increasing the number of strategies from the cluster given in the column on the average number of points

[1] See Appendix E for more details on the underlying simulations.

earned by strategies from the cluster in the row. A + (-) means that there is a positive (negative) effect, a blank entry that the effect is small.

The table shows, for example, that an increase in the number of strategies from cluster 1 in a simulation decreases, on average, the number of points cluster 1 strategies obtain (which is consistent with Within Simulations). However, strategies from cluster 4 and particularly those of cluster 5 benefit from the increase of cluster 1 strategies. For cluster 5 this makes intuitive sense: cluster 5 strategies' win-shift approach only pays off if enough other strategies shift when they are losing (which is what cluster 1 strategies often do). For the same reason cluster 5 strategies benefit from the increase of cluster 2 strategies.

In turn, cluster 1 strategies are hurt by an increase of cluster 2 strategies, because the lose-shift approach of cluster 1 strategies does not pay off if many strategies change after losing. Similarly, it makes intuitive sense that this approach will be successful in the presence of more cluster 4 or cluster 5 strategies, which change rarely, or change after winning, respectively. Cluster 4 strategies benefit from meeting more strategies from their own cluster, but strategies from cluster 5 are hurt by strategies from cluster 4 and by strategies from their own cluster, since the strategies from both clusters rarely shift when losing.

Now that we understand how performance of strategies is generally affected by strategies from the other clusters we can investigate the dynamics of strategy evolution over the different rounds of experiment *Information*. For this we consider all 180 strategies submitted in this experiment (and not just those that are unique). Table 6 shows how the distribution of strategies over the five clusters changes from round to round. In round 1, the majority of submitted strategies (25 out of 43) are from clusters 1 and 2. We know from Table 5 that strategies from cluster 5 benefit from this, and indeed the winner of the first round is the only strategy from cluster 5 that was submitted.

The number of cluster 5 strategies then increases to four in round 2. These four cluster 5 strategies in round 2 rank poorly and make up the last four in round 2 with a very low average number of points of 14.31.[2] At the same time, the number of cluster 1 strategies decreases substantially from round 1 to round 2, but they rank very

[2] The winner from round 1 handed in exactly the same strategy in round 2 but ended up second to last in round 2, earning almost 26 points less than in round 1.

high in the second round, with all three strategies among the prize winners of round 2 and a point average (44.34) that even outperforms the average pay off in pure strategy Nash equilibria. Moreover, the number of cluster 2 strategies also decreases significantly, while round 2 contains more cluster 3 and 4 strategies than round 1.

Round	# and (øPoints) Cluster 1	# and (øPoints) Cluster 2	# and (øPoints) Cluster 3	# and (øPoints) Cluster 4	# and (øPoints) Cluster 5	# and (øPoints) All Strat.	Cluster of Winner	Strategy of Winner
1	9 (30.89)	16 (29.39)	5 (30.66)	12 (34.20)	1 (39.66)	43 (31.43)	5	If you won the previous period always change;
2	3 (44.34)	9 (30.74)	9 (30.09)	15 (29.68)	4 (14.31)	40 (29.57)	1	If you lost the previous period change with probability 0.5;
3	8 (27.32)	12 (15.90)	5 (36.55)	6 (33.61)	1 (48.66)	32 (26.32)	5	If you won the previous period change with probability 0.5;
4	6 (32.73)	8 (29.26)	6 (31.07)	9 (29.60)	6 (35.99)	35 (31.41)	1	If you lost the two previous periods change with probability 0.7;
5	7 (33.71)	5 (33.55)	3 (34.52)	13 (31.62)	2 (33.39)	30 (32.84)	4	*
Total	33 (32.18)	50 (26.79)	28 (32.03)	55 (31.54)	14 (30.59)	180		

* If you won the previous period change with probability 0.0625; else if you won the second last period change with probability 0.125; else if you won the third last period change with probability 0.25; else if you won the fourth last period change with probability 0.5; else if you won the fifth last period change with certainty.

Table 6: The number of strategies and average points in each cluster in each round (column 2 to 7) and the cluster of the winner in each round (column 8)

These results are consistent with the simulation results presented in Table 5. Strategies from cluster 5 are very sensitive to changes in cluster composition, benefiting strongly from clusters 1 and 2 and suffering heavily from clusters 4 and 5. The immense drop in performance of cluster 5 from round 1 to round 2 can therefore be attributed to the decrease in the amount of strategies that shift after losing (cluster 1 and cluster 2) and the increase of strategies that do not shift after losing (cluster 4 and cluster 5). Therefore, cluster 5 strategies generally tend to shift into a majority after being in the minority in the previous period and then get stuck in this majority as many of the strategies in round 2 seldom shift after losing.

Aggregate performance in round 2 is dragged down by the strategies from cluster 5. To corroborate this we run simulations with all strategies of round 2 except the four cluster 5 strategies.[3] Even though cluster 1 strategies would be hurt by the absence of cluster 5 strategies[4] aggregate performance in round 2 would still be 30.93 points (which is close to the average number points under the symmetric mixed strategy Nash equilibrium and in line with most of the other rounds in both experiments). The bad performance in round 2 of experiment *Information* is therefore, to a substantial extent, due to the increase in cluster 5 strategies.

The development of round 2 to 3 roughly mirrors that of round 1 to 2, with more participants using strategies from clusters 1 and 2, and shifting away from cluster 5. As we know from Table 5, this is detrimental for strategies from cluster 1 and benefits strategies from cluster 5 and indeed the only cluster 5 strategy performs exceptionally well and wins round 3. The increase in cluster 2 strategies, and the associated severe drop in performance of these strategies is also quite remarkable. In addition, the low aggregate performance in round 3 can partly be explained by the special feature of round 3 that five completely identical strategies (from cluster 2) are handed in.[5] This strategy always (and only) changes after losing in the previous period.[6] Note that this strategy is exactly the same as the winner in round 2 except for the higher change probability (1 instead of 0.5). Additionally, there is one almost identical strategy.[7] Therefore, strategies from cluster 2 in round 3 generally have a stronger tendency to change (with certainty) after losing the previous period and not to change as much after winning, compared to cluster 2 strategies in other rounds. Redoing the simulations without the strategies from cluster 2 would increase

[3] More precisely, using only these 36 strategies, we run one simulation of 100 periods for each possible combination of five strategies. The interested reader can see the same kind of simulations with (individually) excluding cluster 1, 2, 3 and 4 from round 2 in Appendix E.

[4] Cluster 1 strategies would earn on average 38.54 points in the absence of the four cluster 5 strategies, instead of the average number of points of 44.34 they earned in the simulations for determining the ranks of the strategies.

[5] Usually very few identical strategies are handed in in each round. The strategy “never change” was used three times in round 2 and three times in round 5. Some strategies were used twice in one round and most strategies were used just once in one round.

[6] These five strategies performed extremely poorly in round 3, only achieving 8.8 points each.

[7] This strategy was: always change after losing the previous period, else change with probability 0.2 when winning in a group of 2 in the previous period, otherwise never change. The strategy earned only 8.5 points in round 3.

aggregate performance from 26.32 to 32.21 points, which is in line with the results of most of the other rounds again.

Cluster 2 strategies are not only detrimental to aggregate performance in round 2, they also hurt themselves.[8] For each of the twelve cluster 2 strategies in round 3 separately, we calculate the average points earned in simulations of 100 periods with all combinations of four round 3 strategies from the other four clusters. The average number of points generated by the round 2 strategies over these simulation is 21.55 points, which – although still low – is an improvement of 5.65 points compared to the performance in the experiment. If we repeat this analysis, but this time also exclude all cluster 1 strategies the average of all cluster 2 strategies rises to 31.92 points. This shows that cluster 2 strategies were substantially hurt by the presence of cluster 1 and cluster 2 strategies in round 3.

In round 4, participants again shift away from strategies that did poorly (cluster 1 and 2) and relatively many participants move to the cluster of the previous round winner (cluster 5). The point average is a lot more even across clusters and becomes even more balanced in round 5.

6. Concluding remarks

In this paper we analyzed the impact of information about other players' strategies in a strategy method experiment on the minority game. Participants have to submit strategies to play the five-player minority game repeatedly. In experiment *No Information* (also see Linde et al. (2014)) participants could simulate their intended strategies against random strategies submitted by other participants in the previous round. In experiment *Information* they obtain in addition precise information about the syntax and performance of the strategies submitted by the other participants in the previous round. Moreover, in that second experiment they could choose against which strategies to simulate and might, for example select successful strategies from the previous round as opponent strategies.

Our conjecture was that aggregate efficiency increases with this additional information and simulation possibilities. Participants are better able to predict the strategies they will play against in the next round and can use the simulation

[8] Even if we exclude cluster 1, cluster 2 strategies only achieve a very poor average of 16.67 and with the exclusion of any other cluster, the cluster 2 performance is even worse than its 15.90 points achieved in round 3.

environment to construct strategies that perform better against these predicted strategies. The improved strategies might lead to efficiency levels closer to that of the evolutionary simulation in Linde et al. (2014).

However, aggregate efficiency is not higher in experiment *Information* than it is in experiment *No Information*, and it is actually substantially lower in rounds 2 and 3. The reason for this is that the information provided to the participants induces them to move in the direction of the best-performing strategies of the previous round and away from the strategies that did not perform well in that round. This behavior does not take into account that other participants change their strategies as well.

As a consequence strategies become too similar, which is detrimental to aggregate performance. In particular, our cluster analysis of the submitted strategies in experiment *Information* suggests that the tension between *lose-shift* strategies (strategies that change color in the minority game (with a high probability) after losing, i.e. being in the majority) and *win-shift* strategies (which change color after winning in the minority game) is responsible for the low aggregate performance. Both lose-shift and win-shift strategies are able to exploit strategies of the other type quite well, but suffer when they meet too many strategies of their own type. When there are many lose-shift strategies and few win-shift strategies, the latter type of strategy performs better (which happens in round 1 of experiment *Information*). However, because all strategies and their performance are public information, this leads to an increase in win-shift strategies and a decrease in lose-shift strategies. This in turn improves performance of the latter at the expense of the former, and so on. This mechanism underlies the results in rounds 2 and 3 (where it works in the opposite direction) of experiment *Information*. It also explains why in that experiment for so many participants the strategy from the previous round would have done better against the current round strategies than the strategy actually used in the current round.

Our results on learning are mixed. Participants do not exploit the information and the simulation environment optimally to improve their strategies. Indeed, imitation seems to impede efficiency in rounds 2 and 3. However, already in round 4 participants have learned that imitation in this environment is not necessarily beneficial. As a consequence in rounds 4 and 5 of experiment *Information* aggregate efficiency is quite similar to that in experiment *No Information*. That additional information is not always helpful has been shown in other settings as well. Offerman et al. (2002), for example, show that giving participants additional information about other participants' earnings in a Cournot oligopoly experiment lowers their payoffs.

The strategic environment that we are considering puts a penalty on imitation. By definition in the minority game players want to be in the minority, and – to a certain extent – this also holds for the type of strategy they choose (note that only for cluster 4 strategies payoffs increase when they meet more strategies from their own cluster). It might therefore be interesting to run the same type of strategy experiment in a strategic environment where imitation tends to be rewarded.

References

Arthur, W.B. (1994). Inductive reasoning and bounded rationality. *American Economic Review* 84, 406-411.

Bottazzi, G. and G. Devetag (2007). Competition and coordination in experimental minority games. *Journal of Evolutionary Economics* 17, 241-275.

Brandts, J. and G. Charness (2011). The strategy versus the direct response method: a first survey of experimental comparisons. *Experimental Economics* 14, 375-398.

Challet, D., M. Marsili, and Y.-C. Zhang (2000). Modelling market mechanism with minority game. *Physica A* 276, 284-315.

Challet, D., M. Marsili, and Y.-C. Zhang (2001). Stylized facts of financial market and market crashes in minority games. *Physica A* 294, 514-524.

Challet, D., Marsili, M. and Y.-C. Zhang (2013). Minority Games: Interacting Agents in Financial Markets. *Oxford University Press*, Oxford.

Challet, D., M. Marsilli, and R. Zecchina (2000). Statistical mechanics of systems with heterogeneous agents: Minority games. *Physical Review Letters* 84, 1824-1827.

Challet, D. and Y.-C. Zhang (1997). Emergence of cooperation and organization in an evolutionary game. *Physica A* 246, 407-418.

Challet, D. and Y.-C. Zhang (1998). On the minority game: Analytical and numerical studies. *Physica A* 256, 514-532.

Chmura, T. and W. Güth (2011). The Minority of Three-Game: An Experimental and Theoretical Analysis. *Games* 2011, 2, 333-354.

Chmura, T. and T. Pitz (2006). Succesful strategies in repeated minority games. *Physica A* 363, 477-480.

Devetag, G., F. Pancotto and T. Brenner (2014). The minority game unpacked: Coordination and competition in a team-based experiment. *Journal of Evolutionary Economics* 24, 761-797.

Fernández, A. and S. Gómez (2008). Solving non-uniqueness in agglomerative hierarchical clustering using multidendrograms. *Journal of Classification* 25, 43-65.

Ho, T.H., C. Camerer and K. Weigelt (1998). Iterated dominance and iterated best response in experimental “p-beauty contests.” *American Economic Review* 88, 947-969.

Hommes, C., J. Sonnemans, J. Tuinstra and H. van de Velden (2005). A strategy experiment in dynamic asset pricing. *Journal of Economic Dynamics and Control* 29, 823-843.

Leady, J.R. (2000). Learning in coordination games utilizing categorized information: Experimental evidence. *Unpublished Memo University of Michigan.*

Linde, J. J. Sonnemans and J. Tuinstra (2014). Strategies and Evolution in the Minority Game: A Multi-Round Strategy Experiment. *Games and Economic Behavior* 86, 77–95.

Li, Y., R. Riolo and R. Savit (2000a). Evolution in minority games. (I). Games with a fixed strategy space. *Physica A* 276, 234-264.

Li, Y., R. Riolo and R. Savit (2000b). Evolution in minority games. (II). Games with variable strategy spaces. *Physica A* 276, 265-283.

Moro, E (2004). The Minority Game: an introductory guide. In: *Advances in Condensed Matter and Statistical Physics*. E. Korutcheva and R. Cuerno, eds. Nova Science Publishers, New York.

Nagel, R. (1995). Unravelling in guessing games: An experimental study. *American Economic Review* 85, 1313-1326.

Offerman, T., J. Potters and J. Sonnemans (2002). Imitation and belief learning in an oligopoly experiment. *Review of Economic Studies* 69, 973-997.

Platkowski, T. and M. Ramsza (2003). Playing minority game. *Physica A* 323, 726-734.

Renault, J., S. Scarlatti and M. Scarsini (2005). A folk theorem for minority games. *Games and Economic Behavior* 53, 208-230.

Roth, A. (1995). Bargaining experiments. In J. Kagel & A. Roth (Eds.), *Handbook of experimental economics*. Princeton: Princeton University Press.

Seale, D.A. and A. Rapoport (2000). Elicitation of strategy profiles in large group coordination games. *Experimental Economics* 3, 153-179.

Sonnemans, J., C. Hommes, J. Tuinstra and H. van de Velden (2004). The instability of a heterogeneous cobweb economy: A strategy experiment on expectation formation. *Journal of Economic Behavior and Organization* 54, 453-481.

Sysi-Aho, M., A. Chakraborti and K. Kaski (2005). Searching for good strategies in adaptive minority games. *Physical Review E* 69, 036125

Appendix A: Experimental instructions (translated from Dutch)

Text in CAPITALS are comments on, but not part of, the instructions. Parts in green are only in the instructions for experiment *No Information*, parts in purple are only in the instructions for experiment *Information*. Everything else is in the instructions of both experiments.

Game instructions

The game is played with 5 players. All these players choose between red or blue. Players who selected the color selected by the **smallest** number of players earn one point. The other players earn nothing. In the experiment the decision isn´t made directly by you, but by a strategy devised by you. . This strategy decides when you change color and when you don't. How this exactly works is explained below.

We will first play the game for a number of rounds where you do take the decision yourself. The first round you choose red or blue. In later rounds you choose whether to change color or not.

PARTICIPANTS THEN TWICE PLAYED THE MINORITY GAME FOR 10 ROUNDS AFTER WHICH THEY RECEIVED THE FOLLOWING INSTRUCTIONS

Instructions experiment

Strategies

Before you decided each round whether to change color or not. From now on you will formulate a strategy which in each situation decides for you whether to change color or not. **You specify in which situations you want to change color**. You hand in this strategy in the form of a computer code. How this works is explained below.

Conditions

We use computer code consisting of so called "IF statements" that look like this: "IF (condition) { RETURN number ;"} ".
With these you can determine when you will change or not. Your strategy can consist of multiple if statements.

Condition

The condition in your if statement is either **true** or **false**. In the condition you can use the history of the previous 5 rounds: per period the number of players with your color (including you) and whether you changed color. The table below shows the codes for these events. During the experiment you can make strategies by clicking on the required codes, so you do not need to learn the codes by heart.
For constructing your conditions you can use arguments. ***These arguments are: and/or (OR), and (AND), negation (!), equality (==), smaller than (<), larger than (>), brackets ()***. In the experiment, you can use these arguments by clicking on them. To use the arguments ==, > and < you should view the events as variables which have the value 1 if they are true and 0 if they are false. You can add or subtract conditions using + and -.

IF statement | AND | OR | ( ) | negation ! | == | > | < | + | - | clear

	Changed	Number of players with your color				
		Win		Lose		
Previous round:	$C[1]	$Win[1]		$Lose[1]		
		$W1[1]	$W2[1]	$L3[1]	$L4[1]	$L5[1]
2 rounds ago:	$C[2]	$Win[2]		$Lose[2]		
		$W1[2]	$W2[2]	$L3[2]	$L4[2]	$L5[2]
3 rounds ago:	$C[3]	$Win[3]		$Lose[3]		
		$W1[3]	$W2[3]	$L3[3]	$L4[3]	$L5[3]
4 rounds ago:	$C[4]	$Win[4]		$Lose[4]		
		$W1[4]	$W2[4]	$L3[4]	$L4[4]	$L5[4]
5 rounds ago:	$C[5]	$Win[5]		$Lose[5]		
		$W1[5]	$W2[5]	$L3[5]	$L4[5]	$L5[5]

(This is an example, you cannot yet click on anything.)
Below you will find a number of examples of IF statements. These are only examples and not necessarily smart strategies.

Example 1 (OR argument)
IF ($W1[2] OR $C[5])
means "if I won 2 periods ago with 1 player (including myself) choosing my color **and/or** if I changed color 5 periods ago."

Example 2 (AND argument and negation !)
IF ($L3[4] AND ! $C[2])
means "if I lost 4 periods ago with 3 players (including myself) choosing my color **and** I did not change 2 periods ago."

Example 3 (inequality >)
IF ($C[1]+$C[2]+$C[3] > $W2[1]+$W2[2]+$W2[3])
means "if I in the previous 3 periods changed color **more often** than I won with 2 players (including myself) choosing my color in those same periods."

Example 4 (equality == and negation !)
IF ($C[3] == ! $W2[5])
means "if I changed 3 periods ago **and** I did not win with 2 players (including myself) choosing my color 5 periods ago **or** if I did not change 3 periods ago **and** I did win with 2 players (including myself) choosing my color 5 periods ago."

Number
Your IF statement always ends with "{ RETURN number; }". The number you fill out here determines what happens if your condition is **true**. A 1 means you will change color, a 0 that you will not and a number between 0 and 1 means that you will change color with that probability.

Example 5 (number between 0 an 1)
IF ($L5[4]== $W2[1]) {

RETURN 0.64;
}
means "if I lost 4 periods ago with 5 players (including myself) choosing my color **and** I won 1 period
ago with 2 players (including myself) choosing my color, **or** both are not true, I will change color with
a probability of 64%. "

You can also use 1 as a condition. 1 means "always true".

Example 6 (condition that is always true: 1)
IF (1) {
RETURN 0.4;
}
means "independent of the history I will change color with a probability of 40%."

Strategy
Your strategy can consist of multiple IF statements. In that case the statements are reviewed in the order in which you wrote them down. If a condition in an IF statement is true, subsequent IF statements are ignored. (For those with programming experience: they can be considered ELSEIF statements). If none of your IF statements is fulfilled it is assumed that you will not change color.

Example 7 (multiple IF statements)
IF ($C[5] AND $W2[1]) {
RETURN 1;
}
IF ($L3[2]) {
RETURN 0.5;
}
means "if I changed color 5 periods ago **and** I won with 2 players (including myself) choosing my color in the previous period, I will change color. If that is not true, but I have lost with 3 players (including myself) choosing my color 2 periods ago I will change color with a 50% probability. Otherwise I do not change."
Example 8 (multiple IF statements)
IF ($W1[4] OR ($C[1] AND $L4[3])) {
RETURN 0.2;
}
IF (1) {
RETURN 0.7;
}
means "if I won with 1 player (including myself) choosing my color 4 periods ago, **and/or** I changed color in the previous period and I have lost with 4 players (including myself) choosing my color 3 periods ago, than I change color with a probability of 20%. In all other cases I change color with a probability of 70%."

During the experiment you can either click on all the codes you may need while making a strategy or write them down yourself. You can also cut (ctrl x), copy (ctrl c), paste (ctrl v) and undo things (ctrl z), or redo things that you undid (ctrl y).

Simulations and results
Participants' earnings in every round depended on the place of their strategy in the ranking of strategies In order to determine a ranking of strategies all possible combinations of strategies are considered in a simulation. Each simulation starts with 5 rounds where each player chooses red or blue with equal chance. This way a random history is created. Then 100 rounds are played with the same combination of strategies. For the history it is assumed that you didn´t change color in the first round. The 5 random rounds don´t count towards a strategy's score.
For each simulation the number of points scored by each strategy is recorded. The final score is the average score over all simulations a strategy was involved in. On this basis a ranking is determined.

In the table below you can see the earnings in each round.

	Earnings in each round
Best strategy	€75
Second place	€60
Third place	€45
Fourth place	€30
Fifth place	€15
All other strategies	€0

Using this ranking earnings were determined according to the following table:

	Rounds 1, 2, 3 and 4	Round 5
Best strategy	€75	€150
Second place	€60	€120
Third place	€45	€90
Fourth place	€30	€60
Fifth place	€15	€30
All other strategies	€0	€0

First strategy
In a moment, when you have formulated your strategy, you can check it. There is then a simulation performed with four different pre-programmed strategies as your opponent strategies. Note that for determining the rankings you play against strategies of other participants. The performance of your strategy against the preprogrammed strategies have nothing to say about how your strategy performs in the final simulations.

New strategy
After each round you receive the results by email. In your email you will find the webaddress to adjust your strategy for the next round. After each round you can see how your strategy performs by seeing what happens when your strategy plays against four random strategies from **the previous round**. You can do this as often as you want. Then you can change your strategy. You can also try out new strategies against random strategies from the previous round or against strategies that you consider yourself, for example because you think others are going to use this strategy in the

next round. Below is a picture of the page where you can run simulations. On the left you see a box above the buttons with the codes that you can use to formulate a strategy. The box below is your strategy. By clicking on the buttons below this box you can choose the strategies that your opponents are using. You can also cut a strategy (ctrl-x) and paste (ctrl-v) in the boxes for opponents that you see on the right.

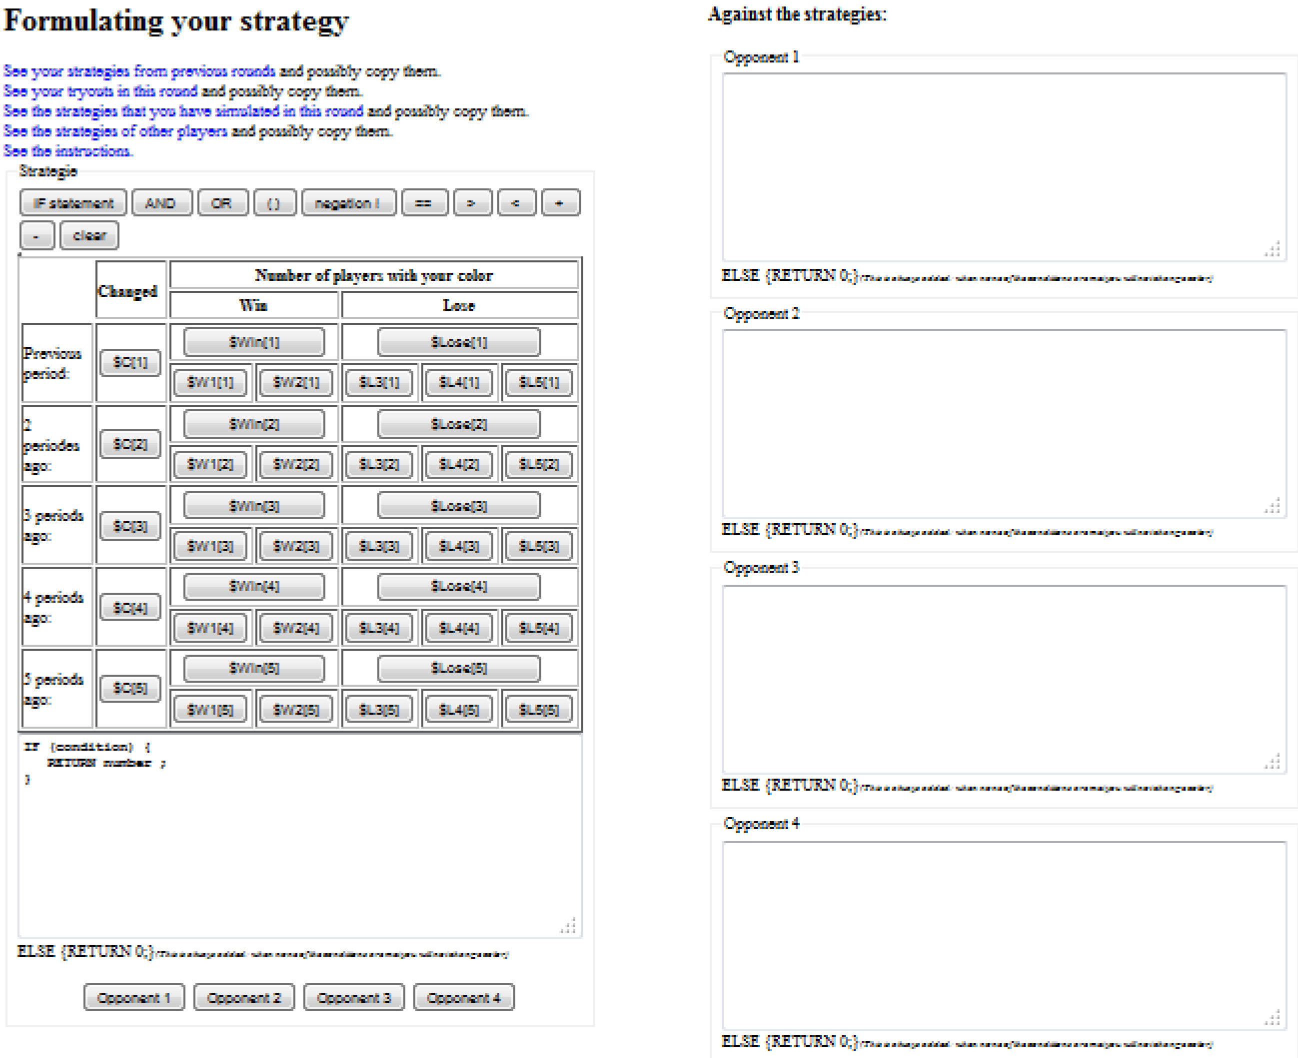

If you don’t enter any opponent strategies you play against random strategies from the previous round. You will still get to see against what strategies you play. You can also fill in some opponent strategies yourself, the other opponent strategies will then be chosen at random by the computer.

In the first round, you play against pre-programmed strategies, if you don’t enter an opponent strategy. You do not get to see these pre-programmed strategies. Above the buttons that you use to formulate a strategy, you can see links that you can click on. If you click on them, a pop-up containing the described strategies appears. These strategies can be used to formulate a new strategy or your opponent strategies in a simulation. You can use strategies by selecting them with the mouse and using copy (ctrl-c) and paste. You can then also edit those strategies. These links appear only if there are strategies that conform to their description. In the first round, for example, there aren’t any links to strategies from previous rounds yet. When you are satisfied with your new strategy you hand it in. You can also hand in your old strategy. **If you do not hand in any strategy, you cannot make any money.**

FOR EXPERIMENT *No Information* THE PAGE WHERE PARTICIPANTS COULD ENTER THEIR STRATEGY LOOKED AS FOLLOWS

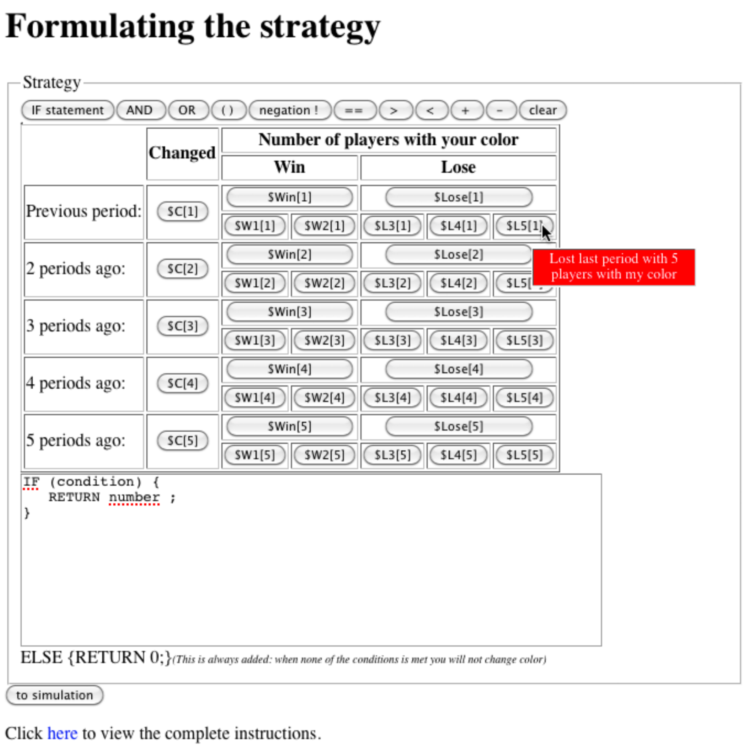

Questionnaire

After handing in a strategy you will be presented with a questionnaire. For each round in which you fill out the questionnaire you will receive 5 euros.

Appendix B: Examples of simulation screens

Example of screen seen by participant after running practice simulation in experiment

Information

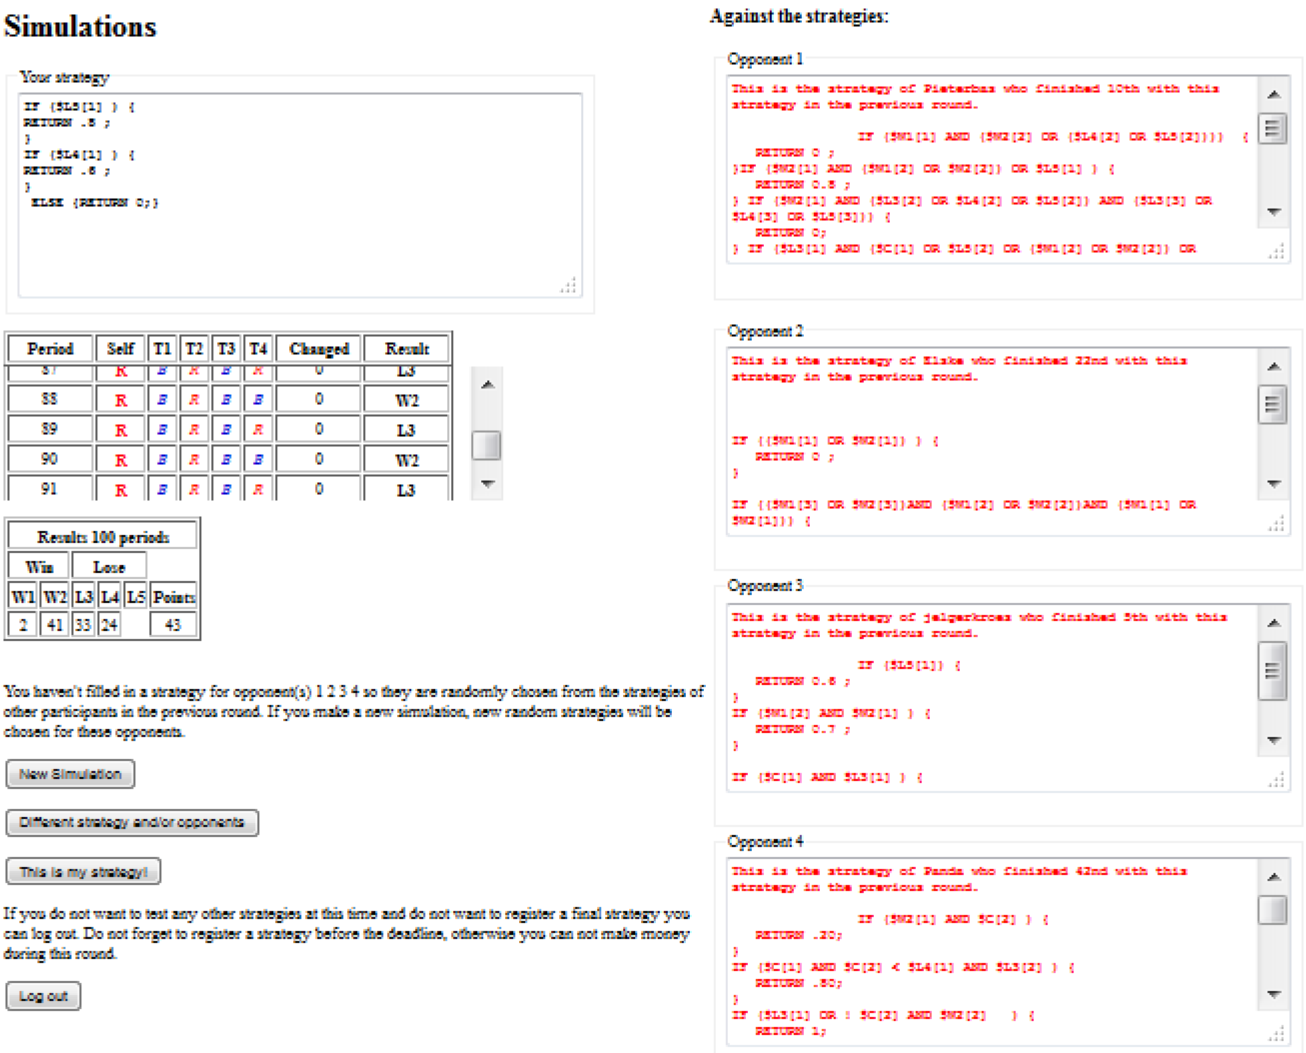

Example of screen seen by participant after running practice simulation in experiment *No Information*:

Simulations

with strategy:

```
IF ($L4[1] ) {
RETURN .2 ;
}
IF ($L5[1] ) {
RETURN .3 ;
}
ELSE {RETURN 0;}
```

Period	Self	Others	Changed	Result
101	B	RRRB	0	W2
102	B	RRBB	0	L3
103	B	RRBB	0	L3
104	B	RRBB	0	L3
105	B	RRRR	0	W1

Results 100 periods					
Win		Lose			
W1	W2	L3	L4	L5	Points
4	26	38	31	1	30

New simulation

Other strategy

This is my strategy! (To submit the definite strategy for that period, not working in the demo)

If you do not want to test any other strategies at this time and do not want to register a final strategy you can log out. Do not forget to register a strategy before the deadline, otherwise you can not make money during this round.

Log out

Appendix C: The selection effect

We compared the performance of participants in the current round that did not submit a strategy in the subsequent round ("Dropouts") to subjects that did submit in the subsequent round ("Non Dropouts"). If Dropouts are primarily low-performing (high-performing) subjects, it is likely that the development of aggregate performance is upward (downward) biased due to the selection effect. One would expect Dropouts to perform worse on average, as low-performing subjects could be inclined not to hand in a strategy in the following round (e.g. because they are frustrated or think they do not have a chance to win). However, it can also be argued that participants who have already won money in the current round might be satisfied with their winnings and/or have altruistic preferences and do not submit a strategy in the following round because they want others to win. Table C.1 shows that the difference in performance between Dropouts and Non Dropouts is very small, suggesting that there was no meaningful selection effect in experiment *Information*.

	Round 1	Round 2	Round 3	Round 4
Dropouts				
#	3	8	3	8
Avg. Points	**31.31**	**28.88**	**25.37**	**31.93**
Standard Deviation	0.96	8.20	17.02	6.73
Non Dropouts				
#	40	32	29	27
Avg. Points	**31.44**	**29.75**	**26.42**	**31.25**
Standard Deviation	4.65	8.30	11.64	7.09

Table C.1: Number and performance of Dropouts and Non Dropouts. Average points are those earned in the current round (the round shown in the first row).

Appendix D: Simulation behavior

Table D.1 shows the average ranks of participants, participants that simulated and participants that simulated at least once against at least one Top5 strategy in experiment *Information* in a specific round. Typically simulators performed better (that is, have a lower rank) than participants who did not simulate (except in round 3) and simulators that simulated against a Top5 strategy performed better than simulators that did not (except in round 5) but none of the differences are significant.

	Average Rank Round 2	Average Rank Round 3	Average Rank Round 4	Average Rank Round 5
All Participants	**20.50** (11.54)	**16.50** (9.23)	**18.00** (10.10)	**15.50** (8.66)
Simulators	**20.17** (12.03)	**17.50** (9.95)	**17.05** (9.85)	**15.11** (8.56)
Non-Simulators	**22.80** (6.88)	**14.30** (6.91)	**19.62** (10.31)	**16.18** (8.77)
Top5 Simulators	**17.43** (11.24)	**13.80** (10.26)	**16.20** (8.91)	**20.75** (9.04)
All Participants without Top5 Simulators	**22.15** (11.37)	**17.00** (8.94)	**18.48** (10.14)	**14.69** (8.42)
Simulators without Top5 Simulators	**22.00** (12.19)	**18.59** (9.59)	**17.67** (9.93)	**13.60** (7.97)

* Standard Deviation in brackets

Table D.1: Simulation behavior and average rank

We also computed, using the distance measure discussed in Subsection 4.2, the distance of final strategies of Top5 Simulators to other Top5 Simulator strategies and to all other participants, respectively. Table D.2 shows that strategies of Top5 simulators have a similar distance to strategies of other Top5 simulators as to strategies of non-Top5 Simulators, indicating that these two kinds of strategies are quite similar.

	Average Distance Round 2	Average Distance Round 3	Average Distance Round 4	Average Distance Round 5
Other Top5 Simulators	**46.48** (4.77)	**38.07** (4.60)	**46.77** (3.17)	**42.18** (5.78)
All Participants without Top5 Simulators	**43.89** (9.91)	**36.66** (5.93)	**43.52** (4.61)	**43.32** (9.19)

* Standard Deviation between average distances of each strategy in brackets.
** The lower the distance the more similar the strategies.

Table D.2: Average distances of strategies of participants that at least once simulated against the Top5 of the previous round to other such strategies (row 1) and to all strategies of participants that never simulated against the Top5 of the previous round (row 2).

Appendix E: Cluster analysis

Table 5 from the main text is based upon Tables E.1-E.6 below.

For Within Simulations in Table E.1, we use all unique strategies of that cluster and run one simulation of 100 periods for each possible combination of five strategies. For Homogenous Simulations in Table E.1, we run 10,000 simulations of 100 periods with 5 identical strategies for each unique strategy and then calculate the average points of strategies in each cluster. We use 10,000 simulations per strategy (instead of just 1) in order to eliminate (most of) the randomness.

	Within Simulation	Standard Deviation	Homogenous Simulation	Standard Deviation
Cluster 1	**26.77**	6.04	**23.89**	9.02
Cluster 2	**31.92**	7.01	**22.43**	11.49
Cluster 3	**30.45**	2.21	**22.56**	11.84
Cluster 4	**35.57**	6.37	**26.74**	10.85
Cluster 5	**0.59**	1.16	**0.71**	1.54

Table E.1: Within Cluster Simulations and simulations against identical strategies

Table E.2 is constructed in the following way. For each column in the table, we draw 100,000 combinations of five strategies using a Mersenne Twister random number generator [9] and run one simulation of 100 periods for each combination. The difference between the columns is the way the combinations of five strategies are drawn. In the second column of Table E.2 we exclude all cluster 1 strategies and randomly draw five strategies from all other unique strategies. In the third column, for each combination, exactly one strategy is randomly drawn from all unique cluster 1 strategies and the other four strategies are randomly drawn from all other unique strategies. Similarly, a combination of five strategies in column 4 (5, 6) consists of two (three, four) randomly drawn strategies from all unique cluster 1 strategies and three (two, one) randomly drawn strategies from all other unique strategies.[10] Tables

[9] http://www.math.sci.hiroshima-u.ac.jp/~m-mat/MT/emt.html (opened on 17/08/2014).

[10] Whenever one or more strategies are drawn from all other clusters, each strategy has the same likelihood to be chosen. For example, in Table E.2 there are 36 unique strategies from cluster 2 (22 from cluster 3, 36 from cluster 4 and 9 from cluster 5). It is therefore much more likely that a cluster 2 strategy is drawn from all other clusters than a cluster 5 strategy, simply because there are more unique cluster 2 strategies. This feature has the advantage that the environment in Table E.2 is quite similar to the environment of the experiment, because clusters that were more frequently used in the experiment

E.3 to E.6 are constructed the same way as Table E.2 but focus on the effect of cluster 2, 3, 4 and 5 on all clusters. For each combination of five strategies in Tables E.2 to E.6 the strategies are drawn with replacement. Note that the last columns, which analyze the performances when four strategies from one cluster face one strategy from all other strategies, are of lesser importance as this situation happens relatively infrequently in the experiment.

	No Cluster 1 Strategy	One Cluster 1 Strategy	Two Cluster 1 Strategies	Three Cluster 1 Strategies	Four Cluster 1 Strategies
Cluster 1	**x**	**38.23**	**31.38**	**27.03**	**26.23**
	x	2.48	2.26	3.10	4.52
Cluster 2	**30.95**	**28.88**	**28.45**	**30.53**	**33.59**
	2.64	1.81	5.55	9.79	13.85
Cluster 3	**30.48**	**30.43**	**32.61**	**37.11**	**41.51**
	1.98	0.99	2.90	5.23	7.72
Cluster 4	**33.64**	**33.32**	**35.15**	**38.73**	**40.56**
	3.62	2.26	2.84	3.89	4.83
Cluster 5	**21.89**	**30.14**	**40.06**	**50.77**	**57.51**
	0.68	0.46	1.31	2.96	5.28
Total	**31.00**	**32.34**	**32.16**	**30.86**	**28.94**
	4.22	3.81	4.99	8.89	11.82

Table E.2: Effect of Cluster 1. Points in bold, standard deviations below.

	No Cluster 2 Strategy	One Cluster 2 Strategy	Two Cluster 2 Strategies	Three Cluster 2 Strategies	Four Cluster 2 Strategies
Cluster 1	**36.68**	**33.57**	**32.43**	**30.48**	**27.80**
	2.82	2.26	2.38	2.37	2.51
Cluster 2	**X**	**28.59**	**29.87**	**30.57**	**31.09**
	X	2.40	1.34	2.53	4.19
Cluster 3	**29.19**	**30.84**	**31.94**	**32.91**	**34.38**
	1.88	0.86	1.57	2.64	4.75
Cluster 4	**34.47**	**34.11**	**33.26**	**32.33**	**30.31**
	3.66	2.52	2.08	2.04	2.46
Cluster 5	**23.91**	**28.07**	**30.22**	**32.85**	**37.34**
	0.77	0.44	0.95	1.99	3.75
Total	**32.69**	**31.79**	**31.40**	**31.16**	**31.15**
	4.94	3.27	2.28	2.57	4.42

Table E.3: Effect of Cluster 2. Points in bold, standard deviations below.

are generally also more often used as a strategy from all other clusters in the simulations for Table E.2-E.6.

	No Cluster 3 Strategy	One Cluster 3 Strategy	Two Cluster 3 Strategies	Three Cluster 3 Strategies	Four Cluster 3 Strategies
Cluster 1	**33.55**	**33.81**	**34.96**	**36.20**	**37.33**
	3.14	1.78	1.23	1.35	1.82
Cluster 2	**29.28**	**29.99**	**31.19**	**32.28**	**33.12**
	2.49	1.11	2.02	3.26	4.30
Cluster 3	**x**	**31.30**	**30.51**	**30.31**	**30.28**
	x	1.43	0.77	1.15	1.72
Cluster 4	**34.48**	**33.47**	**32.54**	**31.59**	**30.40**
	3.03	2.29	1.75	1.63	1.96
Cluster 5	**28.87**	**27.52**	**26.01**	**24.13**	**22.33**
	0.68	0.47	0.68	0.86	0.71
Total	**31.98**	**31.71**	**31.42**	**31.05**	**30.64**
	0.35	2.58	2.69	3.55	4.59

Table E.4: Effect of Cluster 3. Points in bold, standard deviations below.

	No Cluster 4 Strategy	One Cluster 4 Strategy	Two Cluster 4 Strategies	Three Cluster 4 Strategies	Four Cluster 4 Strategies
Cluster 1	**31.23**	**34.08**	**35.07**	**37.31**	**39.09**
	1.27	2.29	3.10	4.17	5.79
Cluster 2	**30.85**	**30.27**	**28.87**	**27.14**	**22.97**
	2.13	1.44	1.62	2.59	4.18
Cluster 3	**31.93**	**31.22**	**30.08**	**28.29**	**24.83**
	1.34	0.94	1.23	2.30	3.68
Cluster 4	**x**	**31.58**	**34.19**	**34.79**	**35.73**
	x	1.39	2.55	3.29	4.49
Cluster 5	**31.43**	**28.40**	**26.28**	**21.34**	**16.50**
	1.15	0.50	0.28	1.08	1.29
Total	**31.26**	**31.33**	**31.94**	**32.63**	**33.93**
	1.73	2.14	3.65	5.72	8.67

Table E.5: Effect of Cluster 4. Points in bold, standard deviations below.

	No Cluster 5 Strategy	One Cluster 5 Strategy	Two Cluster 5 Strategies	Three Cluster 5 Strategies	Four Cluster 5 Strategies
Cluster 1	**31.06**	**39.39**	**60.21**	**87.73**	**88.86**
	2.64	2.52	6.50	11.86	19.67
Cluster 2	**29.49**	**30.15**	**35.52**	**45.25**	**49.10**
	2.76	4.55	13.31	20.36	27.72
Cluster 3	**31.59**	**28.35**	**31.38**	**43.88**	**46.75**
	1.37	3.11	8.59	17.24	26.95
Cluster 4	**34.55**	**31.17**	**34.66**	**62.78**	**66.77**
	2.42	4.57	14.53	25.73	30.36
Cluster 5	**x**	**32.94**	**15.29**	**1.97**	**0.29**
	x	0.56	1.01	1.56	0.25
Total	**31.76**	**32.12**	**29.65**	**24.64**	**12.56**
	3.17	5.30	16.43	29.12	34.05

Table E.6: Effect of Cluster 5. Points in bold, standard deviations below.

Tables E.7 and E.8 show simulation results for round 2 when excluding all cluster 5 strategies and for round 3 when excluding all strategies from a given cluster, respectively.

	Average Cluster 1	Average Cluster 2	Average Cluster 3	Average Cluster 4	Average All Strategies
Average Points	38.54	30.23	29.46	30.72	30.93
Standard Deviation	(2.14)	(3.08)	(3.32)	(4.73)	(4.51)
Effect of exclusion*	-5.80	-0.51	-0.63	1.04	1.36

*Effect of exclusion shows the average performance after exclusion of cluster 5 minus the average performance with all round 2 strategies.

Table E.7: Performances of all round 2 strategies after excluding all cluster 5 strategies of round 2

	All Round 3 Strategies	Without Cluster 1	Without Cluster 2	Without Cluster 3	Without Cluster 4	Without Cluster 5
Cluster 1	**27.32**		2.48	0.45	-1.08	-0.25
Cluster 2	**15.90**	0.77		-1.41	-0.22	-0.65
Cluster 3	**36.55**	0.29	-6.18		2.16	0.87
Cluster 4	**33.61**	-0.99	2.08	-0.08		0.25
Cluster 5	**48.66**	-1.66	-8.79	2.99	5.08	
Aggregate Performance	**26.32**	-0.20	5.89	-2.29	-1.50	-0.84

Table E.8: Performance of all round 3 strategies (column 2) and effect on performance when a single cluster is excluded from round 3 (column 3 to 7)

.