



UvA-DARE (Digital Academic Repository)

Experimenting with new combinations of old ideas

Linde, J.

Publication date
2012

[Link to publication](#)

Citation for published version (APA):

Linde, J. (2012). *Experimenting with new combinations of old ideas*. [Thesis, fully internal, Universiteit van Amsterdam]. Thela Thesis.

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

4. Evolution and Strategies in the Minority Game: A Multi-Round Strategy Experiment *

Should I stay or should I go now?

If I go there will be trouble

An' if I stay it will be double

So come on and let me know!

Lyrics of *Should I Stay or Should I go* The Clash, 1981

4.1. Introduction

Many problems in economics, as well as in other social sciences, center around the competition for (the use of) scarce resources. Often a market institution, or some type of other central agency, serves as a coordination device to allocate these resources, for example through competitive prices. The absence of such an institution, however, may lead to severe coordination problems that are not easy to resolve in a decentralized manner. Such problems may arise when firms need to choose whether or not to enter a new market, or decide which (geographical) market to cater to, or when companies try to be the first to invest in a new technology. Other examples are traders deciding when to buy a rising, or sell a falling stock, commuters choosing a route between their home and the workplace, workers deciding on union membership or high school graduates selecting a college program to enroll into.

*This chapter is based on Linde, Sonnemans and Tuinstra (2011)

In all of these examples, the payoffs (e.g. (future) profits, utility or travel speed) for individual agents crucially depend upon the decisions of the other agents facing exactly the same problem. Certainly in situations where the agents are (nearly) symmetric, the payoffs for successful agents are large and there is not an intuitive focal solution, one can imagine that coordination failure and instability may emerge easily.

The *minority game* provides a very stylized, but intuitively appealing way to model these problems. The minority game has an odd number of players that each simultaneously have to decide between two options, with the players making the minority choice being rewarded, and the others not. The minority game leads to very asymmetric payoffs, making the pure strategy Nash equilibria (any distribution of players across options leading to the largest possible minority) a poor prediction for behavior in the game. The symmetric mixed strategy Nash equilibrium on the other hand (with each player choosing each option with equal probability) may lead to small minorities, and hence suboptimal outcomes. These two observations suggest that the repeated minority game may lead to inherently unstable behavior.

The minority game has been extensively studied by means of a large range of simulation models where different agents use different strategies to play the game repeatedly. These simulation models typically give rise to complex adaptive multi-agent systems, with perpetual fluctuations in the aggregate choices made by agents, but a higher level of coordination on large minorities than in the symmetric mixed strategy Nash equilibrium.

A major drawback of this approach is that the strategies used in these simulation models are selected more or less subjectively by the researchers. Whether decision makers would actually use those strategies is unclear. Since the choice of strategies is a crucial determinant of the dynamic behavior of the game this is highly unsatisfactory. There have been some laboratory experiments on the minority game, but these experiments typically focus on aggregate outcomes. Moreover, due to the relative small number of periods in these experiments and the large strategy space, it is quite difficult to distill the strategies that were used by the participants.

In this paper we introduce an experiment on the repeated (five-player) minority game that employs the strategy method. After gaining some experience with the minority game in the laboratory, participants have to program a strategy. We only impose a few reasonable limitations on the set of strategies they can choose from. A

computer tournament between all submitted strategies then determines a ranking (and the participants that submitted the five highest ranked strategies receive a monetary reward). After providing students with feedback about the performance of their strategy in the computer tournament, they can revise their strategy for the next round (there are five rounds in total – each separated by a week). In this way we elicit explicit strategies from participants, which can subsequently be used for simulation studies with the minority game.

The aim of the paper is twofold. First, we will analyze and classify the strategies that have been submitted by the participants. Second, we will use all strategies submitted in the five rounds (107 unique strategies in total) to run an evolutionary competition, and see which strategies eventually remain.

We find that the strategies submitted lead to aggregate outcomes that are comparable to those under the symmetric mixed strategy Nash equilibrium. Moreover, there is no significant increase in coordination over the five rounds, and learning by participants seems to be limited. The evolutionary competition between the strategies gives four, relatively simple, strategies that remain. Surprisingly, these strategies did not perform very well in the actual experiment. Coordination is enhanced substantially through evolution.

The remainder of the paper is organized as follows. In the next section we will discuss the minority game in more detail and review the computational and experimental literature on this game. We will also discuss a number of related models. Section 4.3 discusses the design of the experiment. In Section 4.4 we analyze the strategies submitted by the participants and in Section 4.5 we use these strategies to establish which of them survives in an evolutionary competition. Section 4.6 summarizes the results.

4.2. The minority game

4.2.1. Definition and relevance

The minority game was introduced by Challet and Zhang (1997) as a stylized version of Arthur's famous *El Farol* bar game (1994). Arthur considers a population of 100 people deciding every Thursday night whether or not to visit the El Farol bar in Santa Fe. This will only be a pleasant experience if at most 60 people are there, otherwise it is too crowded and staying at home would be preferable. Arthur (1994) uses computer

simulations to analyze the interaction of 100 agents. Each agent chooses, from its own set of predictors of the number of attendants, the most accurate one up to that point in time. Aggregate bar attendance turns out to exhibit persistent fluctuations, with average attendance converging to the capacity of the bar, i.e. 60.

The minority game is a symmetric version of the El Farol bar game. There is an odd number of players N , who simultaneously have to choose one of two sides (say *Red* and *Blue*). All players that make the minority choice are rewarded with one ‘point’, the others earn nothing. More specifically, let $s_i = 1$ when player i chooses *Red* and $s_i = 0$ when player i chooses *Blue*. Payoffs for player i are then given by

$$\pi_i(s) = \begin{cases} s_i & \text{when } \sum_{j=1}^N s_j \leq \frac{N-1}{2} \\ 1-s_i & \text{when } \sum_{j=1}^N s_j \geq \frac{N+1}{2} \end{cases} . \quad (4.1)$$

Note that the minority game is one of the simplest games one can think of: there are only two actions to choose from, and only two possible payoffs. Furthermore, the game is symmetric.

The one-shot minority game has many Nash equilibria. First note that any action profile where exactly $\frac{N-1}{2}$ players choose one side constitutes a pure strategy Nash equilibrium. There are $\frac{N!}{\left(\frac{N+1}{2}\right)! \left(\frac{N-1}{2}\right)!}$ of such pure strategy Nash equilibria, which is a substantial number even for moderate values of N . These pure strategy Nash equilibria lead to a very asymmetric distribution of payoffs, with otherwise identical players receiving different payoffs. Moreover, these equilibria are not *strict*: every player in the majority is indifferent between staying in the majority and unilaterally deviating to the minority, which then would become the majority choice. There also exists a symmetric mixed strategy Nash equilibrium, where every player chooses *Red* with probability $\frac{1}{2}$. Expected payoffs in this Nash equilibrium are the same for each player, but aggregate payoffs can easily be smaller than in a pure strategy Nash equilibrium, since there is a positive probability that the minority will be strictly smaller than $(N-1)/2$. Finally, there are infinitely many asymmetric

mixed strategy Nash equilibria. Take for example the 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. Note that this particular type of mixed strategy Nash equilibrium always leads to an efficient outcome.

In this paper we will be primarily interested in the (finitely) repeated minority game, played with a fixed group of players. The multiplicity of (asymmetric) pure strategy Nash equilibria in this symmetric game means that it may be difficult to coordinate on one of those. Moreover, the asymmetry in payoffs and non-strictness of the equilibria make it doubtful that a pure strategy Nash equilibrium, once obtained, will persist very long. It is costless for players in the majority to switch to the other option (although they would prefer other members of the majority to switch), and players in the minority, foreseeing this, may preemptively switch.⁴⁷ The symmetric mixed strategy Nash equilibrium has the obvious disadvantage of giving lower average payoffs than any pure strategy Nash equilibrium. It is therefore unclear what type of behavior to expect in the repeated minority game.

Since its inception the minority game has received quite a lot of attention from physicists, but initially not so much from economists.⁴⁸ There are several reasons for its popularity in physics. First, it is a simple game that allows for studying the interaction of heterogeneous agents as a complex adaptive system. Simulation methods and analytical tools from statistical physics (see e.g. Cavagna et al., 1999, Challet et al., 2000b) have been extensively applied to identify emergent macroscopic properties of these multi-agent systems. Secondly, it has been advanced as a stylized model of a financial market (see e.g. Challet et al. 2000a, 2001), and has become one of the canonical models in the field of ‘econophysics’. In that interpretation the two sides of the minority game correspond to ‘buying’ and ‘selling’ a stock, respectively. If there are more (less) buyers than sellers, the price will be high (low) and sellers (buyers) make a profit.

⁴⁷ 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).

⁴⁸ For example, a search on Web of Knowledge (<http://www.webofknowledge.com>) gives more than 200 published articles with the phrase “minority game” in the title between 1998 and 2011. About 85% of these articles have appeared in physics journals with the rest evenly spread between the fields of computer science, complex systems research and economics.

Although the interpretation of the minority game as a model of a financial market may be criticized for being too simple, the minority game is closely related to, and a stylized representation of, many important economic problems. *Congestion games* for example (see Rosenthal, 1973, for a definition and Huberman and Lukose, 1997, for an application to internet congestion) are games where players make use of limited resources and payoffs are determined by how many other players use that resource. In fact, the minority game is a very simple example of a congestion problem with two routes and N users, where each route has a capacity of exactly $(N - 1)/2$ users, and becomes fully congested when more than $(N - 1)/2$ users choose it. An early laboratory experiment on route choice can be found in Iida et al. (1992), more recent experimental evidence is presented in Selten et al. (2007). In these experiments participants have to minimize travel time by repeatedly choosing between two routes, where travel time depends positively (and more gradually than in the minority game) upon the number of users of the route. In equilibrium travel times are the same between the two routes. Both experiments show that aggregate route choices are volatile, fluctuating around the equilibrium, and that there is substantial heterogeneity in participants' behavioral rules. Selten et al. (2007), for example, show that many participants can be classified as either using a 'direct response mode', where a road change follows a bad payoff, or as using a 'contrary response mode', changing routes after a good payoff. Moreover, the number of road changes is negatively correlated with individual payoffs.

Another problem closely related to the minority game is modeled by the *market entry game*. In such a game each of a number of n firms has to decide independently and simultaneously whether to enter a (new) market or not. The payoff for entering depends upon the total number of firms entering and is typically linearly decreasing in that number, e.g. $\pi_E = k + r(c - m)$, where $c < n$ is the capacity of the market, m is the number of entering firms and k and r are positive payoff parameters. Not entering gives payoffs of $\pi_N = k$. In a (pure or symmetric mixed strategy) Nash equilibrium (in expectation) between $c - 1$ and c firms will enter and in such an equilibrium the expected payoff difference between entering and not entering will be small or zero. Coordination in these market entry games has been extensively studied by means of laboratory experiments, see e.g. Sundali et al. (1995), Erev and Rapoport (1998), Rapoport et al. (1998) and Duffy and Hopkins (2005). A robust finding from

this literature is that aggregate behavior is roughly consistent with Nash equilibrium, but that a large variation in strategies can be observed at the individual level, with some subjects always entering, others never entering and yet other subjects conditioning their behavior on the outcome in previous rounds.

Note that the El Farol bar game is in fact a special case of a market entry game, with a payoff function that does not linearly decrease in m but is a step function with a discontinuity exactly when $m = c$. The payoff function of the El Farol bar game is flat everywhere else. Related to this, Zwick and Rapoport (2000) study a market entry experiment where a fixed prize is equally shared between the entrants, provided there is no over-entry. Entrants have to pay an entry fee, and in case of over-entry their payoffs will be negative. Not entering gives a payoff of zero. Just as in the El Farol bar game this market entry game has a discontinuity at $m = c$ and is flat for $m > c$ (payoffs do depend upon m when $m \leq c$, however). In this case the Nash equilibrium does not give a good description of aggregate behavior and, as was the case in the earlier experiments, there is substantial heterogeneity on the individual level.

An important difference between the market entry / El Farol bar games on the one hand and minority / congestion games on the other is that in the former there is always the safe option of not entering, whereas in the latter all alternatives are subject to strategic uncertainty with payoffs of every choice depending on the decisions made by the other agents.⁴⁹ In two earlier papers on coordination games (Ochs, 1990, and Meyer et al., 1992) a safe option was also absent. In these games players have to choose between different markets at which limited resources are available. For laboratory experiments with these coordination games adaptive strategies seem to give a better description of individual behavior than mixed strategy Nash equilibria. Moreover, whenever participants succeed in coordinating on a pure strategy Nash equilibrium, this coordination turns out to be fragile and is typically not maintained in the subsequent periods.

On a more general level the minority game is an abstract version of games where actions are strategic substitutes. Well-known examples of such games are

⁴⁹ Another difference is that the (pure strategy) Nash equilibrium is Pareto efficient in minority and most congestion games, but not in the market entry game. Total payoffs in a market entry game increase when the number of entrants decreases below capacity. The non-entrants then achieve the same payoffs as they do in equilibrium whereas the payoffs of the entrants are higher. However, this type of tacit collusion is typically not observed in market entry experiments, possibly because these experiments often involve rather large groups of about 12 to 20 participants.

cobweb markets (Ezekiel, 1938) and Cournot oligopolies. For example, if most producers in a cobweb market predict next period's price to be higher than the rational expectations (RE) equilibrium price and therefore produce more than the RE equilibrium quantity, the actual market clearing price will be lower than the RE equilibrium price. In such a cobweb market it is therefore better to disagree with the majority prediction. A similar argument holds for Cournot oligopolies. If the other firms on average have high production levels, it is optimal to supply a limited amount, and the other way around.

The minority game is therefore a relevant, although stylized, model for a number of important economic problems. What sets it apart from most other coordination problems is that the pure strategy Nash equilibria lead to Pareto efficient but very asymmetric payoff distributions whereas the symmetric mixed strategy Nash equilibrium equalizes expected payoffs between all players, but those are lower than average payoffs in a pure strategy Nash equilibrium. One might therefore expect less stable outcomes in minority games than in market entry or route choice problems. In the next subsection we review some of the computational and experimental literature that has been done in the past decade on the minority game and that corroborates this conjecture.

4.2.2. Strategies in the minority game: computational and experimental research

In the physics literature the minority game is studied by using computer simulations. In these simulations (see e.g. Challet and Zhang, 1997, 1998) the number of agents is large (typically between $N = 101$ and $N = 1001$) and every player has a fixed set of S strategies (typically $S = 2$, but sometimes higher values of S are used), randomly drawn from the set of all strategies with memory M (typically smaller than 10). Such a strategy maps the history of the past M winning sides into a prediction of the next winning side. The number of different histories is therefore equal to 2^M and since any history can be mapped into one of two sides, the total number of different strategies is 2^{2^M} , a number that increases fast with M (e.g. for $M = 5$ the total number of strategies is already about 4.3×10^9). Note that these strategies do not use information

about the *size* of the minority and that they do not allow for randomization.⁵⁰ Agents collect how well the strategies in their set predict the winning side (but do not consider the effect that a strategy they did not use might have had on the outcome) and in every period choose side according to that strategy, from their set, that is the best predictor up to that period.⁵¹ 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. One of the most celebrated results on the minority game is that of the dependence of ‘cooperation’ on the parameter $\rho = 2^M / N$, for the first time identified by Savit et al. (1999). For small values of ρ , where the number of agents is relatively large compared to the number of possible histories, aggregate behavior in the minority game is dominated by a cycle of period 2 and volatility is higher than under the symmetric mixed strategy Nash equilibrium. However, for moderate values of ρ , volatility drops below that of the symmetric mixed strategy Nash equilibrium, reaching a minimum value at some critical level $\rho = \rho^c$, and increasing towards volatility under the symmetric mixed strategy Nash equilibrium again, when ρ increases beyond that critical value.⁵²

In the typical analysis of the minority game, as discussed above, each agent has a set of strategies from which it chooses one every period. The set of strategies of an individual agent is fixed and randomly drawn from the total set of strategies.⁵³ It therefore contains arbitrary strategies, that may lack any rationale, but nevertheless the agent will hang on to these strategies forever. A number of models have been

⁵⁰ An exception is Johnson et al. (1999) (see discussion below). Cavagna et al. (1999) develop a continuous version of the minority game where instead of making a binary choice each agent submits a ‘bid’ that may lie somewhere in between the two extremes.

⁵¹ Alternatively, in some papers the choice of strategy is assumed to be probabilistic with the probability that a strategy is chosen positively related to its success, for example through a logit specification (see e.g. Cavagna et al., 1999, and Challet et al., 2000b).

⁵² Cavagna (1999) shows numerically that this result is maintained if the actual history is replaced with a ‘fake’ history: what is important for cooperation is that agents react to the *same* information, whether accurate or not. Ho et al. (2005), on the other hand, provide numerical evidence that for large scale minority games with about 3000 agents knowledge of the true history is relevant.

⁵³ In some contributions an agent’s strategies are not drawn independently from the set of all strategies. Challet et al. (2000b), for example, assume that an agent’s second strategy is always chosen such that it is exactly opposite to its first strategy. Yip et al. (2003) consider strategies that are slightly biased to one alternative and show that this improves efficiency. Finally, Wang et al. (2009) consider a minority game with ‘heterogeneous preferences’, meaning that there are agents of different types K , with $K=0,1,\dots,2^M$, where an agent of type K takes the first side for exactly $K/2^M$ (randomly determined) histories.

advanced in which the set of strategies used evolves over time under evolutionary pressure. Li et al. (2000a), for example, consider an evolutionary model where after each generation (consisting of 10.000 periods) all agents are ranked according to their performance in that generation. Of the lowest ranked 20% (or 10% or 40%) of the agents, 50% is randomly selected and replaced with new agents, each endowed with a new and random set of strategies (with the same memory length M). It turns out that this improves efficiency considerably – volatility is now always lower than under the symmetric mixed strategy Nash equilibrium, even for low values of M – although volatility is still a non-monotonic function of ρ . Li et al. (2000b) extend this model by assuming that the strategies of the new agents may have a memory length M that differs from that of the agents they replace. They find that agents with memories below a certain threshold \bar{M} perform very well, whereas agents with longer memories are eventually driven out through evolutionary competition. This result is in sharp contrast to some early evolutionary simulations in Challet and Zhang (1997, 1998), who also establish that evolution substantially increases efficiency, but who show that the average memory length increases through evolution.⁵⁴ Sysi-Aho et al. (2005) use a genetic algorithm to update the strategy set for the poorest performing agents of each generation through crossover between two of the original strategies. They find that fluctuations decrease substantially in such a setting and almost maximal efficiency may be obtained.

Johnson et al. (1999) take a different approach to study the evolution of agent strategies in the minority game. Each agent collects, for each sequence of winning sides, the outcome that prevailed the last time that sequence obtained. An agent then chooses with probability p (which differs between agents) the same outcome as the last time the sequence appeared and with probability $1 - p$ the other outcome. If the agents aggregate score falls below a certain threshold he changes his strategy by drawing a new probability p from an interval around his old value. Simulations with this model show that segregation emerges: the distribution of the values of p within the population becomes bimodal with peaks at 0 and 1, implying that many agents

⁵⁴ The evolutionary model of Challet and Zhang (1997, 1998) is slightly different, however. After each generation the worst ranked agent is replaced by a clone of the best ranked player, with an additional mutation process where, with some small probability, one of the strategies of the clone is replaced by a random other strategy (in absence of this mutation process efficiency will decrease dramatically through evolution).

always replicate the outcome stored in their memory and many agents always do exactly the opposite.

Clearly, aggregate behavior in minority games crucially depends upon the strategies agents are assumed to use and how these strategies evolve over time. Results on a number of recent laboratory experiments on the minority game may help in understanding which strategies would actually be played by humans. Bottazzi and Devetag (2003, 2007), for example, consider groups of five players playing the minority game for 100 periods, varying the memory length and information provided to the participants, and find that, although aggregate choices are volatile, (allocative) efficiency is higher than in the symmetric mixed strategy Nash equilibrium.⁵⁵ Increasing the memory length has a slightly positive effect on efficiency, but providing additional information does not.⁵⁶ Participants seem to repeat their choices, particularly when they have more information and after a win. Moreover, participants seem to revert to pure strategies towards the end of the experiment.

The failure of the symmetric mixed strategy Nash equilibrium to describe human behavior in the minority game is confirmed in Devetag et al. (2011). They consider a three-player minority game experiment where each player is represented by a team of three participants. Teams are video recorded and their discussion is analyzed to learn about strategies used for playing the minority game. Again, allocative efficiency is higher than in the symmetric mixed strategy Nash equilibrium (but not higher than with individuals instead of teams). Analysis of the video recordings reveals that teams rarely use a randomization strategy and that they tend to focus more on their own past behavior than on other teams over time, in particular when they have been successful.

For most of the experiments on the minority game discussed above it is impossible to determine whether participants randomize or not. Chmura et al. (2010) study a three-player minority game experiment where participants can explicitly use mixed strategies. Moreover, there is random re-matching of groups after each period, which makes the mixed strategy Nash equilibrium a more obvious candidate for

⁵⁵ This is consistent with the findings of Platkowski and Ramsza (2003) who run an experiment in which a group of fifteen students plays the minority game for 200 consecutive periods.

⁵⁶ On the other hand, Chmura and Pitz (2006) show, in a minority game experiment with groups of nine players, a positive effect of adding information about the distribution of individual choices upon efficiency. They also establish a negative correlation between the number of changes of a subject with its cumulative payoff. This negative correlation was also established in simulations with the minority game by Challet and Zhang (1997).

individual behavior. They find that there is considerable heterogeneity in decision rules, and the behavior of only about a quarter of the participants is best described by the symmetric mixed strategy Nash equilibrium.

Other laboratory experiments on the minority game are Wang et al. (2009), who show that a model of fixed rule learning with heterogeneous preferences gives a better description of their experimental minority game data than the standard fixed rule learning model that is typically used in the minority game literature, and Liu et al. (2010) who show that the behavior of fish in a minority game experiment is remarkably similar to that of university staff members: as experience of the subjects increases the variance of decisions first decreases and after a certain point increases again. Finally, Laureti et al. (2004) discuss an interactive web-based experiment where one human agent plays the 95-player minority game against computer agents (with different degrees of sophistication). Humans perform better than computer agents in environments where the latter have a memory up to 4 periods, but more sophisticated computer agents typically outperform the human players.

Although laboratory experiments can shed some light on the type of strategies that humans employ for playing the minority game, it is still difficult to infer exactly the strategies being used, and therefore to draw conclusions on what type of behavior is relevant for minority games. A strategy experiment, where participants have to submit strategies to play the repeated minority game therefore seems appropriate. 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). Also the famous strategy tournament on the repeated prisoner's dilemma in Axelrod (1984) is related to our work. In a recent paper Brandts and Charness (2011) provide an overview of experiments that directly compare the strategy method with the 'direct response' method. They find that in most studies these methods yield qualitatively similar results.

In this paper we use the strategy method to elicit explicitly the strategies used by human players of the minority game. We will analyze those strategies and use them to study evolutionary competition, in order to understand the type of strategies and behavior that is relevant for minority games.

4.3. Design

We designed an experiment in which participants have to submit a strategy to play the five-player minority game for 100 periods. The experiment consists of five rounds, each separated by a week, and took place in April 2009. Participants are students of the so-called “beta-gamma” bachelor program, which is one of the most challenging programs of the University of Amsterdam. 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.

The first stage of the experiment takes place in the CREED laboratory of experimental economics at the University of Amsterdam. The minority game is explained to the participants and they play the minority game two times for 10 periods in two different groups of players. After getting acquainted with the minority game in that way, participants are explained – on a handout and via the computer screen – how to formulate a strategy. Their understanding of formulating strategies is checked by letting them program two verbal strategies using the interface, after which they formulate, test and submit their first strategy.⁵⁷ After a few days all participants receive, by email, the results of the first round. From then on they can login on the website and try out new strategies against the population of strategies of the previous round. Within a week after the laboratory experiment they have to submit their new strategy (which could be identical to their old one) and fill in a short questionnaire. Two days after the deadline they receive the results of the second round. The whole procedure is repeated for another three times. A week after the fifth and final round we explain the goals of the experiment in class, announce the results of the final round and pay out the earnings.

4.3.1. Formulating strategies

Figure 4.1 shows the computer screen where the participants can formulate their strategy.⁵⁸ A strategy has the form of a list of IF-statements that (if the condition is

⁵⁷ The participants can ask the experimenters for further instructions during the initial laboratory experiment. For the later rounds the experimenters were available for assistance via e-mail, although participants made no use of this possibility.

⁵⁸ The experiment is programmed in php/mysql and runs on a (Apache) web server. An English translation of the experiment can be found on www.creedexperiment.nl/minor/english and the

met) returns a number in the interval $[0, 1]$, which is the *probability of changing color*. If a condition in an IF-statement is fulfilled, the subsequent IF-statements are ignored (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). The number of IF-statements is unlimited and strategies can use logical expressions such as AND, OR, (in)equality and negation. In the instructions ample examples were given; see Appendix 1. The strategies can use the history of the last 5 periods, which consists of the outcome in each of these previous periods (i.e. the size of the group) and whether the strategy changed colors in that period or not.

Strategy

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

	Changed	Number of players with your color				
		Win		Lose		
Previous period:	SC[1]	SWin[1]		\$Lose[1]		
		SW1[1]	SW2[1]	\$L3[1]	\$L4[1]	\$L5[1]
2 periods ago:	SC[2]	SWin[2]		\$Lose[2]		
		SW1[2]	SW2[2]	\$L3[2]	\$L4[2]	\$L5[2]
3 periods ago:	SC[3]	SWin[3]		\$Lose[3]		
		SW1[3]	SW2[3]	\$L3[3]	\$L4[3]	\$L5[3]
4 periods ago:	SC[4]	SWin[4]		\$Lose[4]		
		SW1[4]	SW2[4]	\$L3[4]	\$L4[4]	\$L5[4]
5 periods ago:	SC[5]	SWin[5]		\$Lose[5]		
		SW1[5]	SW2[5]	\$L3[5]	\$L4[5]	\$L5[5]

Lost last period with 5 players with my color

```

IF (condition) {
  RETURN number ;
}
ELSE {RETURN 0;}(This is always added: when none of the conditions is met you will not change color)

```

to simulation

Figure 4.1: Computer screen as seen by the participants when they formulate a strategy.

interested reader is invited to formulate a strategy and run simulations with that strategy against actual strategies of our participants.

We have restricted the strategy space in two ways. First, note that the minority game is a symmetric game where the labels of the two sides (red and blue) have no intrinsic meaning. We therefore impose symmetry in our design by letting strategies decide on changing color instead of choosing a color. Although individuals may have a preference for one of the colors, for example preferring winning when choosing blue over winning when choosing red, this limitation seems reasonable: using colors directly would double the number of variables per period.

Another restriction is the length of the history. We took this to be equal to five periods, which we believe gives a sufficient amount of flexibility for participants to develop strategies.⁵⁹ The information about the last five periods is complete and contains whether the strategy made the winning decision, what the distribution of choices is and whether the strategy changed colors in that period (the total number of possible histories is therefore $5^5 \times 2^5 = 100,000$). Note that in the simulation studies of the minority game the information used by strategies only contains the winning sides of the last M periods, but not the distribution of choices. On the other hand, in laboratory experiments on the minority game (e.g. Bottazzi and Devetag, 2007) the full distribution of choices is given in some treatments. Also note that in the bulk of the simulation studies randomization is not allowed, and in almost no laboratory experiments it is observed explicitly (an exception is Chmura et al., 2011). In our design participants can explicitly submit mixed strategies, and in fact about 75% of all submitted strategies make use of randomization.

4.3.2. Simulations by participants

A novel feature of our design, compared to other strategy method experiments, is that participants can run simulations with a strategy of their own making. Simulations are ran with four randomly drawn strategies (without replacement) from other participants from the previous round.⁶⁰ Since strategies can use a history of up to 5 periods, first 5 random outcomes are drawn. After that, 100 periods are played according to the five

⁵⁹ No more than 17% of all strategies submitted in our experiment uses information from 4 periods ago, and only 13% of the strategies uses information from 5 periods ago, whereas more than 90% uses information of the previous period, about two thirds of the strategies use information from two periods ago and about half of the strategies uses information from three periods ago. Limiting the history to five periods therefore seems to be relatively innocuous.

⁶⁰ 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 q , where $q = 0, 0.1, 0.2, \dots, 0.9, 1$, respectively.

strategies. After each simulation the results of the 100 periods, as well as those of the first five random periods are presented (see Appendix 1). In the presentation the choices of the other four strategies are sorted in each period (first the red and then the blue choices) making it close to impossible to infer what the other strategies in the simulation are. In addition summary statistics are displayed: the total number of points and the number of times the 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. Our approach therefore gives ample opportunities for participants to gain experience with the game and to learn how to play it.

4.3.3. Computer tournament, feedback and earnings

After the deadline a computer tournament with all submitted strategies is run as follows. For every possible combination of five strategies a simulation of 100 periods is done (after five initial randomly selected outcomes), implying that the total number of simulations with about 40 strategies is around 10^8 . Subsequently we determine for each strategy the average number of points it earned in all the simulations it was involved in and use this to rank the strategies, where the first five random outcomes in each simulation are not used in determining the average number of points. After each round, all participants receive an email with a ranking of the strategies and the average number of points earned by each strategy in the simulations.⁶¹ In addition they learn their earnings for that round. In the first four rounds the top five strategies receive 75, 60, 45, 30 and 15 euro, respectively. In the fifth and final round these amounts are doubled (and therefore are 150, 120, 90, 60 and 30 euro). In addition in each round every participant who submits a strategy and fills out a short questionnaire receives 5 euro. One of the questions in the questionnaire is about how confident the participant is about the success of his/her strategy. To have an incentivized check on this question the opportunity is presented to wager the 5 euro. If this option is chosen for a certain round, an extra reward is given when the strategy ends up in the top 5 of

⁶¹ Strategies are identified by the nicknames of participants. It was not possible to observe the strategies used by other participants.

75, 60, 45, 30 and 15 euro, respectively, in that round. Note that with 45 participants the expected value of this option is 5 euro and (risk neutral) participants that expect to have a strategy of above average quality should use this option. Average earnings for the whole experiment were 58.70 euro per participant, ranging from a minimum of 0 euro to a maximum of 260 euro.

4.3.4. Questionnaire and web-server data.

In the laboratory session we administered a short questionnaire about the background of the participants (age, gender, programming experience, etc). After submitting the strategy in each round a few questions about the (formulation of the) strategy are asked: how difficult it was to formulate the strategy, whether they had any problems with the formulation and how confident they are that the strategy will be successful. Finally, the incentivized question about confidence described above was asked. Besides the submitted strategy and the answers to the questionnaire, the web server also saves all actions of the participants: when and how often they log in; which strategies they try; the results of the simulations they run.

4.4. Results from the multi-round strategy experiment

In this section we describe the most important results from the multi-round strategy experiment. In the first subsection we present some results on the aggregate outcomes of the simulations and performance of the individual participants. In the second subsection we will take a closer look at individual strategies, and categorize them by means of cluster analysis.

4.4.1. Aggregate outcomes and performance of participants

For the repeated five-player minority game that we are considering the best outcome is one where the minority consists of two players in every period. This happens in any pure strategy Nash equilibrium (PSNE). However, in the symmetric mixed strategy Nash equilibrium (MSNE), inefficiencies do occur since randomization implies that, with a positive probability, the minority will be smaller than two. In fact, it can be easily checked that in the symmetric MSNE the probability that a minority of two results is 62.50%, whereas the probability of obtaining a minority of 1 (0) is 31.25% (6.25%).

Round	1	2	3	4	5	Symmetric MSNE	PSNE
Number of participants	42	36	34	36	32		
Outcome							
3-2	63.90%	61.66%	64.57%	66.07%	64.32%	62.50%	100%
4-1	30.65%	32.25%	30.26%	29.33%	30.59%	31.25%	0%
5-0	5.45%	6.08%	5.17%	4.60%	5.09%	6.25%	0%
Points							
Average	31.69	31.12	31.88	32.29	31.85	31.25	40
Standard Dev.	1.49	5.17	3.10	3.20	6.36		
Minimum	29.31	21.65	24.98	28.00	18.93		
Maximum	34.68	41.65	36.96	39.87	43.06		
Average change propensity	47.61	38.15	45.99	38.51	40.29		
Pearson correlation change and points	-0.531	0.028	-0.690	-0.493	-0.878		

Table 4.1: Distribution of outcomes and performance of participants over the rounds

As explained above, in each round of the strategy experiment we run a simulation of 100 periods for each possible combination of five submitted strategies. The first round started with 42 participants submitting a strategy; in the subsequent rounds the number of submitted strategies was between 32 and 36.⁶² Rows 3 – 5 of table 4.1 show the distribution of minorities resulting from the simulations with these submitted strategies. These distributions are very similar to those one under the symmetric mixed strategy Nash equilibrium (seventh column of table 4.1). Coordination of the strategies on larger minorities is slightly better than under the symmetric MSNE in most rounds – with the highest level of coordination obtained in round 4 – but slightly worse in round 2. Clearly, coordination in any round is far from that obtained in a pure strategy Nash equilibrium (last column of table 4.1).

It might be argued that the level of coordination is not the appropriate measure to look at, since payments for participants are based upon the relative ranking of the strategies they submit and not on the absolute number of points these strategies generate. However, there is an incentive for participants to maximize their number of points. 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

⁶² Participants that missed a round were allowed to submit in later rounds. Dropping out in a round appears not to be related to success in the previous round: only participating in round 4 is significantly negatively correlated to the rank in round 3 (Wilcoxon rank-sum test with p-value of 0.0543).

that simulation, will hurt their performance relative to the 30 to 35 strategies that are not present in that simulation.

Table 4.1 also presents some results on the performance of participants and the strategies they submitted. The average number of points for the strategies varies from 31.12 and 32.29 between the rounds, which is very close to individual performance in the symmetric MSNE, for which the average number of points is 31.25 in 100 periods (note that the average number of points in a PSNE is 40). In fact, the proportion of strategies with an average number of points larger than 31.25 in round 1 to 5 is given by 23/42, 15/36, 17/34, 19/36 and 17/32, respectively, which is not significantly different from 50% (sign test). The symmetric MSNE therefore does well as a description of aggregate outcomes. It performs poorly at the individual level, however. The dispersion between payoffs generated by strategies, as measured either by the standard deviation of points, or by the range between the minimum and maximum number of points, is considerable, certainly when taking into account that each strategy was involved in more than 30,000 simulations of 100 periods.⁶³ These differences between the rules are therefore structural and there is substantial heterogeneity between the strategies in terms of their performance. This is consistent with the experimental literature on e.g. market entry problems reviewed in Section 2 which features substantial heterogeneity on the individual level, whereas the Nash equilibrium gives a good description of the aggregate outcome. Dispersion is much higher in rounds 2 to 5 than in the first round, suggesting that the strategies are more homogeneous in the first round and that heterogeneity increases after that. Moreover, the fact that the best strategies in round 2 and 5 earn more points than they would in a pure strategy Nash equilibrium suggests that these strategies (and possibly other strategies as well) are able to exploit some of the other strategies.

Strategies may be heterogeneous along different dimensions. For example, the strategies vary from very short (“never change”) to very long.⁶⁴ They also differ

⁶³ Although we simulated all possible combinations of strategies there is still some randomness in the average number of points, due to the first five random periods in each simulation and because strategies may be randomizing themselves. To check whether this randomness has an impact on the outcome we ran all simulations once more. This gives almost identical results: the correlation between the ranks in the two simulations turns out to lie between 0.998 and 1.

⁶⁴ One participant (participant 34) handed in strategies between 35 and 246 IF-statements in rounds 2 to 5. This participant ran about 5000 simulations with a 50%-change strategy and determined, by means of a computer program, for each possible history what the optimal (non-random) response would be. Note that this strategy, although quite creative, responds to strategies from the previous round, not taking into account that those will change as well. This procedure has given rise to one

substantially in how often they change colors. Some strategies never change colors, others change in about 95% of the periods. In all rounds but round 2 performance is negatively related to the propensity to change (Spearman rank correlation p-values smaller than 0.01).⁶⁵ We can see no consistent decrease in this tendency over the rounds, which suggests that participants do not learn that their strategies change too often and adapt them accordingly.

Improving a strategy on the basis of simulations against strategies from the previous round is not straightforward, since the other participants change their strategies as well. The adapted strategy may not perform as well in the new environment of strategies as expected. We consider two criteria that measure whether participants are successful in improving their strategies over the rounds. First, one would expect that the adapted strategy performs better than the old strategy when playing against strategies from the previous period, since the strategy can be tested (without limits) in that environment. Second, the new strategy should do better than the old strategy in the new environment. If it does not, the participant had better refrained from adapting the old strategy. Table 4.2 shows the results: over all rounds the new strategy is an improvement over the old strategy for the old environment in only slightly less than half of the cases. Only for round 4 a clear majority of strategies correspond to improvements over their old versions, when considering the old environment. The comparison of the old and new strategies in the new environment is more promising; but still in about only 60% of the cases the new strategy does better than the old one would have done, with round 4 giving the best results again.

<i>Round</i>	<i>The new strategy would have done better than the old strategy in the old environment</i>	<i>The new strategy does better than the old strategy would have done in the new environment</i>
2	38.89%	61.12%
3	53.57%	44.11%
4	73.33%	88.88%
5	23.33%	46.87%
Total	46.77%	60.87%

Table 4.2: The performance of the old and the new strategies in the old (column 2) and the new (column 3) environment

successful strategy (the winner in round 3), but not in the other rounds partly because the strategies are untidy and prone to mistakes (the strategy for round 2, for example, contained a mistake and never changed).

⁶⁵ This negative correlation between performance and the propensity to change is consistent with the results from computational models and laboratory experiments on the minority game (see Challet and Zhang, 1997, and Chmura and Pitz, 2006, respectively).

Besides the actual submitted strategies the experiment generates a wealth of web-server data, which may shed some light on how people tried to learn. On average participants tried out eight different strategies per round, and ran about 150 simulations with those strategies. More strategies are being tried in the first two rounds and the average number of strategies tried stabilized at around five strategies per participant in each of the last three rounds. We also considered the login behavior of participants in rounds 2-5: on average they logged in between 1 and 2 times in each of those rounds, and were logged in almost two hours in total in those rounds. The effects on the rank of the participant are ambiguous, however. Both the number of strategies used and the number of simulations ran have a positive effect upon the rank of the participant in round 4 only (i.e. higher rank, Spearman rank correlation p-values of 0.0052 and 0.0021, respectively).⁶⁶ Remarkably, the number of strategies used has a *negative* effect on performance in round 2 (Spearman rank correlation p-value of 0.0105). In none of the other rounds there is a significant effect. Also the number of times logged in has no significant effect in any of the rounds, the time logged in only has a *negative* effect on performance in round 2 again (Spearman rank correlation p-value of 0.0111). These results, combined with those from table 4.2, suggest that participants were on average not able to use the possibilities of the minority game website to improve their strategies substantially. On the other hand, some participants did succeed in improving their strategy, as can be seen from the tendency of the best strategies to generate more points over the rounds (see table 4,1).

An analysis of the decisions of participants to forego the fee of five euro in exchange for higher prizes suggests that participants were also unable to accurately predict the performance of their strategy. Participants choose higher prizes over the fee of five euro for about one third of all 180 submitted strategies. This decision was – as to be expected – positively correlated with the answer to the question from the questionnaire how confident the participant was about his strategy in rounds 2, 3, 4 and 5 (Mann-Whitney test p-values of 0.0016, 0.0110, 0.0013 and 0.0017), but remarkably only significantly positively correlated to the performance of the participant in round 4 (Mann-Whitney test p-value of 0.0088).

⁶⁶ We also considered, in the spirit of the analysis presented in table 4.2, the following. Let $S_i(R, T)$ be the score of the strategy of participant i in round R , when playing against the strategies the other participants submitted in round T . We checked whether the differences $S_i(R, R-1) - S_i(R-1, R-1)$ and $S_i(R, R) - S_i(R-1, R)$ are correlated with the number of simulations participants ran in that round. This is only (positively) significant for $S_i(3, 2) - S_i(2, 2)$, (p-value of 0.0059) suggesting that only in round 3 participants, on average, were able to improve their strategies by running more simulations.

Round	Prize	Strategy	Eq. Strat.	Cluster	Evolution
1	1	1-10		5	44.2
	2	1-40		4	13.4
	3	1-31		5	16.8
	4	1-26		5	29.6
	5	1-47		6	500
2	1	2-28	1-28	1	10.4
	2	2-29		1	10.8
	3	2-5	1-5	1	10.8
	4	2-4		5	54.8
	5	2-10	1-10	5	44.2
3	1	3-34		N.A.	N.A.
	2	3-10		5	68.6
	3	3-4	2-4	5	54.8
	4	3-11		5	59.4
	5	3-7	1-46	4	11.2
4	1	4-11	3-11	5	59.4
	2	4-35		5	28.8
	3	4-41		5	17.4
	4	4-4	2-4	5	54.8
	5	4-31		5	18.0
5	1	5-46	1-46	4	11.2
	2	5-38		4	12.2
	3	5-4	2-4	5	54.8
	4	5-31		5	24.4
	5	5-11	3-11	5	59.4

Table 4.3: The winners of the five rounds.

Strategy x-y is the strategy that participant y submitted in round x. The fourth column shows whether the strategy was already submitted in an earlier round (by the same or another participant). The fifth column gives the cluster in which the strategy is classified (see discussion below) and the sixth column gives the number of generations the strategy survives in an evolutionary competition between strategies (averaged over five evolutionary simulations) (see the discussion in the next section).

4.4.2. Classification of strategies

In this section we will study the strategies submitted in the experiment in a bit more detail. Over all the rounds participants submitted 180 strategies. However, only 110 out of those strategies are unique. This is because sometimes participants use the same strategies in two or more different rounds, and sometimes two different participants use the same strategy. Both instances are illustrated in table 4.3, which gives the five winners for each of the five rounds. The fourth column shows that each round contains at least one winner that was first submitted in an earlier period, by the same or a different participant. What is noteworthy that several strategies win a prize in different rounds, e.g. strategies 2-4 (winning in all four rounds in which it participated)

and 3-11 (winning in all three rounds in which it participated). Also, strategy 1-10 – together with its slight variation 3-10 – wins a prize in three rounds.

Obviously, the three participants submitting those strategies perform quite well in the experiment, participant 4 earned 260 euro, participant 10 245 euro and participant 11 235 euro. The only other participant who earned more than 200 euro is participant 46, (245 euro, he/she opted for the additional prize money in the final round, which already has higher stakes, and received the first prize round). All other participants earned 170 euro or less.

One could therefore argue that strategies 2-4, 3-11 and 1-10 (3-10) are the three most successful strategies submitted during the experiment. Table 4.4 gives a description of these strategies. Strategy 2-4 specifies to change only with a positive probability (and under some additional conditions) if the strategy is on the losing side for two consecutive periods. Strategy 3-11 presents a more extreme version of the same principle: it specifies to change (with certainty) only if the strategy lost in four consecutive periods. Finally, strategies 1-10 and 3-10 again tend to change when they lost for two consecutive periods. Moreover, these two strategies have two additional features, not shared with strategies 2-4 and 3-11, namely that they change with probability 0.5 when they lost in a group of five in the previous period, and they also change with a substantial probability (0.75 and 0.5, respectively) when they *won* (in a group of 2) in the previous period, and lost two periods ago in a group of 4. In general, all of these successful strategies are reluctant to change, but also make sure that they will not get stuck in a losing situation forever.

We tried to categorize all submitted strategies in different clusters. Of the 110 unique strategies we exclude the three very long computer generated strategies from participant 34 for the analysis in this and the next section (see footnote 18 for a discussion of those strategies). The main reason for this is that these are not the type of strategies that would typically be used by human decision makers and they are difficult to interpret (another minor practical reason is that these strategies increase computation time considerably).

2-4	If you changed two periods ago and lost two periods ago and you did not change in the previous period and lost in the previous period, change with probability 1 Else if you lost in the last two periods in a group of three and you lost three periods ago, change with probability 0.8 Else if you lost in a group of 4 in the previous period and lost two periods ago but did not change two periods ago, change with probability 0.6
3-11	Only change (with probability 1) when you lost in each of the last four periods
1-10 (3-10)	If you lost in a group of 3 in the last two periods change with probability 0.5 Else if you lost in a group of 4 in the last two periods, change with probability 1 Else if you lost in a group of five in the previous period, change with probability 0.5 Else if you lost in a group of 4 two periods ago and won in a group of 2 in the previous period, change with probability 0.75 (0.5)

Table 4.4: Description of strategies 2-4, 3-11 and 1-10 (3-10).

Note that strategy 3-10 only varies from strategy 1-10 (both submitted by participant 10) in the probability of change in the last IF-statement (0.5 instead of 0.75).

We performed a cluster-analysis with the remaining 107 unique strategies. Figure 4.2 shows the resulting dendrogram. For this analysis we constructed a matrix of distances between strategies, calculated as follows. Strategies can use the history of the last 5 periods (outcome and whether they had changed colors in that period), which gives 100,000 possible histories. For every strategy the decision (probability of change) is calculated for each possible history. The distance between two strategies is then defined as the weighted average absolute difference between these probabilities.⁶⁷ Because not all histories are equally likely (a 5-0 outcome is less likely than a 3-2 outcome) the weights are based upon the distribution that would results from the symmetric mixed strategy Nash equilibrium.⁶⁸ We used the program multidendrograms⁶⁹ to draw the dendrograms, using the algorithm “joint between within” which both tries to minimize the distances within clusters and maximize the distances between clusters (Székely and Rizzo, 2005)

We find six clusters, labeled 1/6, 2/6 etc in figure 4.2, and on a higher level three clusters, labeled 1/3, 2/3 and 3/3. table 4.5 displays for each cluster the most central strategy, that is, the strategy with the minimum average distance to the other

⁶⁷ This is a continuous version of the Hamming distance.

⁶⁸ 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. This is very close to the numbers in the simulation discussed below, which are 5.63%, 30.87% and 63.50%, respectively.

⁶⁹ <http://deim.urv.cat/~sgomez/multidendrograms.php>, see also Fernández and Gómez (2008).

strategies in that cluster. The most central strategies (CS) give some impression of the kind of strategies in that cluster. CS1/6 and CS2/6 change often, the first typically after a losing period, for the second it is independent of the history of outcomes. The central strategy from cluster 3, CS3/6, is the one shot symmetric mixed strategy Nash equilibrium strategy (and therefore is also independent of the history of outcomes). CS4/6 seems to be a bit peculiar since it changes color after winning. However, this strategy might be quite sensible in an environment with many strategies that have the tendency to change after losing in the last period.⁷⁰ In fact, CS4/6 is the strategy that wins the fifth round (with double prizes) of the strategy experiment (see table 4.3). CS5/6 is the very simple strategy of never changing and CS6/6 only changes (with probability 0.5) after losing in the previous period.

The central strategies suggest that strategies from clusters 1 and 2 change relatively often, and strategies from cluster 5 and 6 relatively little, with the strategies from cluster 3 and 4 somewhere in between. In Section 4.4.1 we saw that, except in round 2, the propensity to change is negatively correlated to performance of the strategies, which suggests that, on average, strategies from clusters 5 and 6 should do well. To a certain extent this conjecture seems to be corroborated by the fourth column of table 4.3, which shows from which clusters the winning strategies in the different rounds originate. Two thirds of the winning strategies are from cluster 5/6, with the other winners typically coming from clusters 4/6 and 6/6, except for the second round where the first three prizes are for strategies from cluster 1. Note that the average number of points was lowest in round 2, and highest in round 4, where all winning strategies come from cluster 5.

⁷⁰ In the terminology of Selten et al. (2007) strategy CS4/6 can be classified as using a ‘contrary response mode’, whereas for example strategy CS1/6 uses a ‘direct response mode’.




Figure 4.2: A cluster analysis of the 107 unique strategies.

Cluster	Most central strategy
1/6	1-5: <i>If you have lost the last period change with probability 0.8; if you win three times in a row change with probability 0.5</i>
2/6	3-36: <i>Always change with probability 0.9</i>
3/6	2-27: <i>Always change with probability 0.5</i>
4/6	1-46: <i>Only change when you won the last period.</i>
5/6	1-34: <i>Never change</i>
6/6	4-6: <i>When you lost the last period change with probability 0.5.</i>

Table 4.5: The most central strategy in each of the six clusters (the strategy with the minimum average distance to the other strategies in that cluster)

To further study these issues we ran 500,000 simulations of a 100-period minority game where, for each simulation, we randomly selected (with replacement) five strategies from the set of 107 unique strategies.⁷¹ For all strategies the average number of points over these simulations and the percentage of changes in these simulations are calculated. The second and third columns of table 4.6 show the average per cluster. The results are consistent with the discussion above. The change propensity is indeed very different between clusters; all differences are statistically significant with the exception of 4/6 versus 1/6 and 3/6. Changes are least frequent in 5/6 and 6/6. Furthermore, the strategies in clusters 5/6 and 6/6 earn more points than the strategies in the other clusters (Mann-Whitney tests, all p-values <0.001) and strategies in cluster 5/6 perform better than those in cluster 6/6 (p-value <0.001).

We have done additional simulations to understand the role that heterogeneity plays in the experiment. First we did simulations in a fully homogeneous setting, i.e. with five identical strategies. The fourth column of table 4.6 shows the averages per cluster. Average earnings in these homogeneous settings appear to be very low, even substantially lower than in the symmetric mixed strategy Nash equilibrium. This suggests that participants designed their strategies to exploit other strategies without taking into account that these strategies may be similar. For example, the central strategy of cluster 4/6 (1-46) “Change only when you won the last period” (i.e. the winner of the final round of the strategy experiment) will always lose after the first period in a homogeneous group. Note however that in the experiment strategies could only meet exact copies of themselves if other participants would submit the same strategy. Strategies from clusters 2/6, 5/6 and 6/6 perform significantly better in

⁷¹ The difference with the simulations that we ran to determine the ranking of strategies in the experiment are that strategies submitted in different rounds can now play against each other; not all combinations of submitted strategies are simulated, and random selection is with replacement.

homogeneous settings than strategies from the other three clusters.⁷² Still, strategies from all clusters benefit from at least some degree of heterogeneity and are not very well adapted to homogeneous environment.

To investigate this issue a bit further, we also ran additional simulations where all strategies were selected from the same cluster. Again, average performance of the strategies from clusters 2/6, 5/6 and 6/6 in these simulations is higher than that of the strategies from the other clusters (who on average still earn less than they would in the symmetric mixed strategy Nash equilibrium). Note that the strategies from cluster 5/6 do quite well when only playing each other (even better than when playing the full population of strategies), whereas strategies from clusters 1/6 and 4/6 seem to be particularly badly equipped to play against strategies that are similar to them. We will get back to this issue when we discuss the evolutionary simulations in the next section.

Cluster	Change (Sd)	Points (Sd)	Earnings in homogeneous simulations (Sd)	Earnings in simulations within cluster (Sd)	Evolution last generation alive (Sd)	N
1/6	54.20 (14.83)	31.20 (1.74)	16.28 (10.71)	20.47 (12.97)	11.27 (1.91)	14
2/6	71.34 (9.23)	30.33 (1.16)	27.97 (7.91)	32.79 (4.70)	9.89 (1.23)	16
3/6	44.66 (12.06)	30.70 (0.59)	23.95 (9.05)	30.31 (2.64)	11.44 (0.74)	33
4/6	46.43 (15.28)	30.80 (1.28)	12.94 (9.68)	23.89 (9.63)	11.72 (0.89)	13
5/6	17.00 (8.18)	34.46 (1.69)	27.59 (9.25)	36.60 (3.90)	105.42 (161.36)	18
6/6	33.99 (10.22)	32.58 (1.23)	25.54 (11.64)	32.38 (3.02)	95.74 (179.97)	13
Tot	44.16 (19.88)	31.59 (1.89)	23.02 (10.68)	29.92 (8.11)	37.27 (97.68)	107
K-W	0.000	0.000	0.000	0.000	0.000	

Table 4.6: Average characteristics of strategies per cluster

Column 2 and 3 display the percentage change and the average number of points per cluster in a simulation with all 107 unique strategies. Column 4 displays the average points in homogenous groups (5 identical strategies) and column 5 displays the average number of points in groups with only strategies from the same cluster. Column 6 displays the average last generation alive in an evolutionary simulation (see Section 5 for a discussion). The last row gives Kruskal-Wallis tests (p -values based on 2-sided tests).

⁷² We also studied whether the clusters differ in other aspects, like complexity, length of history used, etc, but found no consistent differences.

Evolution and Strategies in the Minority Game

In addition to studying success of strategies against all other strategies, or only against strategies from their own cluster (third and fourth column of table 4.6, respectively) we did one final simulation exercise with strategies from only two clusters, in order to learn about the interaction between those clusters. Figure 4.3 shows four of the comparisons between pairs of clusters.⁷³ The left top panel shows the interaction between clusters 1/6 and 4/6. We see that both types of strategies perform well if the majority of strategies are from the other cluster and vice versa. The central strategies for clusters 1/6 and 4/6 (see table 4.5) suggest a reason for this pattern. Strategies from cluster 1/6 typically have a propensity to change after losing while strategies from cluster 4/6 have a propensity to change after winning. The right top panel shows the interaction between clusters 1 and 5. Strategies in both clusters perform quite well (close to 40 points) when only 1 or 2 strategies from cluster 5/6 are present, but if there are 3 or 4 strategies from that cluster all strategies perform bad. The left bottom panel is an example where the composition matters much more for one cluster (6/6) than for the other cluster (4/6). Finally, the right bottom panel shows an example where the performance of both clusters is not very sensitive to the group composition.




Figure 4.3: Interaction between strategies from two clusters. The horizontal axis displays the number of strategies from respectively the blue (cross) and the red (square) cluster; the vertical axis the average number of points.

⁷³ Figures of the other 11 combinations of two clusters are available from the authors upon request.

4.5. Evolutionary competition between submitted strategies

In this section we consider an evolutionary competition between the strategies submitted by the participants to determine which of those strategies will eventually survive. The analysis in Section 4.2 on the different clusters already allows us to formulate some conjectures about the final distribution of surviving strategies. For example, the fifth column of table 4.6 suggests that the evolutionary competition will not eventually result in an environment with only strategies from cluster 1/6 (or only from cluster 3/6 or only from cluster 4/6). On the other hand, Figure 4.3 suggests that it is in principle possible that evolution results in coexistence of strategies from clusters 5/6 and 6/6 (or coexistence of strategies from clusters 1/6 and 4/6).

We model the evolutionary competition between strategies as follows. In the first generation every strategy i has the same *weight* $w(i,1) = \frac{1}{N} = \frac{1}{107}$. In every generation g we run 2000 simulations of 100-period minority games. In each of these games five strategies are randomly selected (with replacement), where the probability of selecting strategy i equals its weight $w(i,g)$. For each strategy i we determine the average number of points it earned, averaged over all simulations that it was part of. We denote this average by $P(i,g)$. We also determine the average number of points earned by all strategies, averaged over all simulations, and denote this by $M(g)$. After each generation of 2000 simulations the weights of the different strategies are updated on the basis of how well they did as compared to the whole population of strategies. This updating is formalized as follows: $\bar{w}(i,g+1) = (1 + \lambda [P(i,g) - M(g)])w(i,g)$ where λ is a positive parameter which measures *selection pressure*. Note that if a strategy performs better than the average strategy in a generation, its weight increases ($\bar{w}(i,g+1) > w(g,i)$). If $\bar{w}(i,g+1) < 0$, which happens if a strategy performs much worse than average, its weight is set to $\bar{w}(i,g+1) = 0$ and the strategy becomes extinct. The same thing happens if a strategy was selected in none of the 2000 simulations of a generation (which is only likely to happen if its weight is very small to begin with).

The final weights are then determined as $w(i,g+1) = \frac{\bar{w}(i,g+1)}{\sum_j \bar{w}(j,g+1)}$ to make sure that

the weights sum up to one again.

Evolution and Strategies in the Minority Game




Figure 4.4: Evolutionary analyses with a starting population of 107 strategies. On the horizontal axis the generation (after generation 50 in steps of 10), on the vertical axis the percentage of the population, averaged over five evolutionary simulations.

We ran this evolutionary simulation five different times with the same 107 strategies and for 500 generations, with the parameter λ set equal to 0.05.⁷⁴ Figure 4.4 shows how the weight of each of the 107 strategies, averaged over the five simulations, evolves over the 500 generations. The outcome of this evolutionary simulation is very robust: in all five simulations the same four strategies survive for 500 generations with more or less the same weights in the last generation. In only one of the simulations another strategy (2-32) survives for the first 500 generations, but its weight in generation 500 is very small (0.11% of the population), and it seems likely that this strategy would die out if the evolutionary simulation would run longer.

⁷⁴ A low value of λ leads to a very slow evolution and long simulation times, while a large value increases the role of ‘bad luck’. The parameter value we chose is relatively low: even the worst performing strategies take at least 8 generations to die out. Moreover, for higher values of λ some weights may become negative, which does not happen in our case.




Figure 4.5: Development of average points earned and propensity to change during the evolutionary analysis

Generations are on the horizontal axis, the left vertical axis (green line) is the average percentage that a strategy changes and the right vertical axis (red line) the average number of points of a strategy in a game of 100 periods, again averaged over the five evolutionary simulations.

Although there is little evidence that, over the course of the strategy experiment, coordination on the optimal outcome increases (see the discussion in Section 4.1) coordination grows considerable, on average, in the evolutionary simulations.⁷⁵

Figure 4.5 shows, for the first 50 generations, the average number of points and the average percentage that strategies in the population change color (these levels remain more or less constant after the first 50 generations). The average number of points clearly increases over the generations, from about 31.5 (which is very close to the average number of points under the symmetric mixed strategy Nash equilibrium, which equals 31.25) to 38.5 points (which approaches the maximal possible average number of points of 40 that is obtained under any pure strategy Nash equilibrium). After 50 generations the efficient distribution (with two players in the minority) occurs in 92.66% of the periods, whereas an inefficient outcome with one (zero)

⁷⁵ This is reminiscent of the result from evolutionary simulations in computational models of the minority game, where efficiency also becomes quite large (see e.g. Li et al, 2000ab, Sysi-Aho et al., 2005).

player in the minority occurs in 6.97% (0.37%) of the periods, and these numbers remain more or less constant from then on (for the first generation these numbers are 63.63%, 30.85% and 5.52%). Particularly the worst outcome (with all players making the same decision) is quite rare now, occurring only once every 270 periods, as opposed to about every 18 periods in the mixed strategy Nash equilibrium.

Figure 4.5 suggests that the substantial increase in coordination is directly related to a decrease in the percentage with which strategies in the population change colors, which goes from about 45% when all 107 strategies are present in the population, to around 8% when only the four survivors remain. This relationship between change and number of points may be caused by the fact that predictability, and thereby coordination, increases in an environment with strategies that rarely change, provided that at least some of these strategies condition on the history.

Let us now study in some more detail which strategies survive the evolutionary competition, and what they look like. The survivors, all from clusters 5/6 and 6/6 are described in table 4.7. It is interesting to compare the survivors of the evolutionary competition with the winners of the five rounds in the strategy experiment, given in table 4.3. Several things are noteworthy. Out of the four survivors of the evolutionary competition, only strategy 1-47, the clear winner of the evolutionary competition, won a prize in the strategy experiment (fifth place in the first round). In the last column in table 4.3 we indicated the average number of generations the winners of the strategy experiment survive in the evolutionary competition. Most winners become extinct within the first 20 generations (with the winners from clusters 1/6 and 4/6 dying out first).⁷⁶ The best performing strategies in the strategy experiment (strategies 2-4, 1-10 (3-10) and 3-11) survive on average for about 50 generations and are therefore also relatively successful in the evolutionary competition.⁷⁷ Other successful strategies in the evolutionary competition are 2-32, that becomes extinct only after 349 generations, strategy 4-1 (143 generations) and 5-10 (66 generations).

⁷⁶ Strategy 2-27, which is the central strategy of cluster 3/6 and corresponds to the symmetric mixed strategy Nash equilibrium, does not perform very well. It does not win in any round of the strategy experiment, and it becomes extinct in each of the five evolutionary simulations in generation 15 or earlier.

⁷⁷ Note that strategy 3-10 survives substantially longer than strategy 1-10 (on average 68.6 versus 44.2 generations). The only difference between the two strategies is that the former changes color with probability 0.5 (instead of probability 0.75) if the strategy lost in a group of four two periods ago and won in a group of two in the previous period (see table 4.4). This suggests that the last condition of strategies 1-10 and 3-10 diminishes the evolutionary fitness of these strategies.

1-47	If you won the last period: don't change Else if you lost the last period in a group of 3 or 4: change with probability 0.2 Else if you lost one or both of the periods -2 and -3: change with probability 0.6
2-35	Only change (with probability 1) when you lost the last period in a group of 5
1-32	Only change (with probability 1) when you lost the last period in a group of 4 or if you won in period -3 in a group of 1.
1-34	Never change

Table 4.7: The four surviving strategies after 500 generations.

Apparently, strategies 2-4, 1-10 (3-10) and 3-11 perform quite well in an environment where there are still many other strategies, possibly by exploiting the less successful strategies. However, when those less successful strategies disappear, these three winners cannot feed off them anymore and are overtaken by the strategies that eventually win the evolutionary competition and that were not that successful when all strategies were still present. To understand this phenomenon better table 4.7 shows the four strategies that survive 500 generations of evolutionary competition.

Strategy 1-47 is the clear winner of the evolutionary competition, with a proportion of the population of more than 60%. This strategy does not disturb a winning situation, is not too eager to change after losing (and does not change if it lost in the previous period in a group of five, implicitly predicting that enough other strategies will change in that situation). Finally, it prevents staying in a losing situation forever. In contrast, 1-34 is more conservative and never changes. This strategy can never become a dominating strategy in a population because it does relatively bad when meeting itself (in a homogeneous group its expected earnings are 31.25 points, see table 4.8). In fact, it earns less than the other surviving strategies in last generations; which indicates it is still decreasing in strength. We have run simulations for all possible group compositions with the four survivors and solved for the population equilibrium by finding a distribution for which all active strategies earn the same average number of points (this equilibrium is given in the sixth column of table 4.8). We find that strategy 1-34 would eventually die out, and the population fraction of strategy 2-35 (change only when the last period was lost in a group of 5) would increase. The average number of points earned in the population equilibrium is slightly higher than that in generation 500 (38.47 versus 38.40).

Strategy	Cluster	Population proportion after 500 generations	Earnings	Earnings Homogeneous	Population equilibrium	% Change
1-47	6/6	61.5%	38.43	36.8	60.1%	14.8%
2-35	5/6	14.6%	38.98	31.3	24.5%	4.1%
1-32	6/6	13.5%	38.24	21.7	15.4%	28.8%
1-34	5/6	10.4%	37.76	31.3	0%	0%
Overall			38.40		38.47	

Table 4.8: Characteristics of the four surviving strategies.

Cluster refers to section 4.4.2 and figure 4.2. The population proportion (earnings) are the average in generations 491-500 in the 5 evolutionary simulations. The earnings solo are the average earnings of that strategy in homogenous groups. The population equilibrium displays the population proportions in the equilibrium is which all strategies have the same expected earnings. The final column displays the percentage of periods the strategy changes color.

Let us now compare the strategies of the survivors of the evolutionary competition (table 4.7) with those of the best performing strategies in the strategy experiment (table 4.4). In general, the latter set of strategies shows a high level of inertia: these strategies typically change colors after they lost at least two consecutive periods, and not even always then. This turns out to be quite profitable in an environment where many strategies change color immediately after losing. However, these strategies are too reluctant to change in an environment where evolutionary competition has driven out the strategies that change often. In such an environment strategies that change a bit quicker (like the first three strategies in table 4.7) do better, and therefore eventually win the evolutionary competition.

To test robustness 25 additional evolutionary simulations are run, each with a random selection of 50 of the original 107 strategies.⁷⁸ In these simulations between 2 and 7 strategies survive up to generation 500. In 24 of the 25 simulations all survivors are from the clusters 5/6 and 6/6, the exception is the only one with 7 survivors where there is one survivor from cluster 3/6 and one from cluster 4/6, with the other five from clusters 5/6 and 6/6. The four surviving strategies from the original simulation all do very well: 1-47, 1-32 and 1-34 are survivors in all simulations in which they are involved; 2-35 survives in 10 of the 13 (77%) simulations in which it is involved. All

⁷⁸ The only other difference with the original evolutionary simulations is that now in a generation only 1000, instead of 2000 minority games (each of 100 periods) are simulated. Since the number of strategies is less than half the number for the original evolutionary simulation, this implies that (at least in the beginning of the evolutionary simulation, when all weights are equal) each strategy is involved on average in the same number of simulations as in the original evolutionary simulation.

other strategies perform worse with the exception of strategy 2-32, which is very similar to 1-32 and survives in 10 out of 12 (83%) of the simulations which it was involved in. In these 25 additional simulations the average number of points increases and the average propensity to change decreases over generations, like in figure 4.5 for the original evolutionary simulations. The average number of points in generation 491-500 is slightly less than in the original simulations; 37.12 versus 38.40 and the average propensity to change is higher 12% versus 8%.

The evolutionary simulations discussed this far all suggest that as evolution proceeds the surviving strategies will change color less, on average. This, of course, also has to do with the environment: in an environment where other strategies change more often than not, a good strategy might be also to change more often. To illustrate this point we ran 5 simulations without the strategies of clusters 5/6 and 6/6. In the first generation the average number of changes is about 52%. This increases to about 85% in generation 50 (see figure 4.6). Average earnings also increase: to an impressive 39 points around generation 20 and finally stabilize around 37.8 points. A final remarkable aspect of these simulations is that it takes a longer time to weed out bad strategies: strategies become extinct in generation 15 (instead of generation 7).




Figure 4.6: Evolutionary simulations with strategies of cluster 5/6 and 6/6 excluded. On the left axis change (green line) and on the right axis (red) earnings.

4.6. Conclusion

In this paper we used an internet based strategy method experiment to explicitly elicit the strategies employed in the minority game. Participants could try out their strategies on the minority game website by simulating them against strategies submitted by the other participants in the previous round. This allowed them to adapt their strategies in the direction they believe will be successful. We believe this is not only a novel aspect of the experimental design, but also relevant for many applications where decision makers, for example traders in financial markets, have the possibility to employ technological tools to try to improve their decisions.

We find that the aggregate behavior of the submitted strategies is close to that implied by the symmetric mixed strategy Nash equilibrium. However, there is considerable heterogeneity between the submitted strategies. Remarkably, participants do not structurally succeed in improving their strategies over the five rounds, and hence the amount of learning seems to be limited, although the minority game website provides ample opportunities to learn. We aim to examine the causes of participant's inability to improve performance in future experiments.

A cluster analysis revealed that the submitted strategies can be divided in six distinct groups. The central strategies in each cluster give an idea of the types of strategies in each group. We find both lose shift (1/6) and win shift (4/6) types of strategies. The most central strategy in another cluster (6/6) also only changes after losing, but with a low probability. The central strategies of the three other clusters are independent of the history and either change very often (1/1), half the time as in the symmetric mixed strategy Nash equilibrium (3/6), or never (6/6). Success of a strategy is negatively correlated with how often the strategy changes sides. The most successful strategies come from cluster 5/6 and have a tendency to only change sides after being on the losing side for two consecutive periods or more.

The experiment shows that the restrictions placed on strategies in many evolutionary simulations of the minority game do not allow for the kind of strategies people actually use. Importantly we find that people use mixed strategies and condition their actions on more detailed histories than just winning or losing. We therefore perform an evolutionary simulation with the 107 unique strategies gathered in the experiment.

Although the most successful strategies from the actual strategy experiment perform relatively well in the evolutionary competition, they do die out eventually, to the benefit of some other strategies that survive the evolutionary competition and that are slightly less reluctant to change. The intuition for this is that the initially successful strategies do quite well in an environment where all submitted strategies are present, because they profit from strategies that change often. When those strategies become extinct, however, the initially successful strategies are not very well suited to the new strategic environment and other strategies take over. Evolutionary competition leads to a fast and dramatic improvement in coordination.

In conclusion we find that there is substantial heterogeneity in the strategies people use in the minority game and many strategies condition on quite specific histories and use randomization. Participants on average fail to improve their strategies between rounds despite the possibilities offered by the minority game website. It appears that as a result there is very little coordination and aggregate outcomes resemble the mixed strategy equilibrium. If bad strategies are allowed to go extinct in an evolutionary simulation the remaining strategies do achieve much higher levels of coordination. The surviving strategies are reluctant to change, thereby enhancing stability, but they do change, with a small probability, when a bad situation lasts too long. This appears to be the key to their coordination.

4.7. Appendix A: Experimental Instructions

(Translated from Dutch. Original Dutch instruction available upon request)

The minority game

The minority game is played with 5 players, each of which chooses either **red** or **blue**. Players who selected the color selected by the **smallest** number of players earn one point. Other players earn nothing. Then a new round is started and everyone decides to change color or not. The game is repeated a large number of rounds with the same player.

In the experiment the decision isn't made directly by you, but by a strategy devised by you. How this exactly 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 round the number of players with your color (including you) and whether you changed color can be used. The table

Evolution and Strategies in the Minority Game

below shows the codes for these events. In construction your conditions you can use arguments. **These arguments are: and/or (OR), and (AND), negation (!), equality (==), smaller than (<), larger than (>), brackets ().** 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 - (this is an example, you can do 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 getal ; }". 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 and 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 one 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, then 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 strategies 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. Using this ranking earnings were determined according to the following table:

Evolution and Strategies in the Minority Game

	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

Testing strategies

After writing a strategy you can test it in simulations against four random strategies from round 5 of the experiment. In the experiment subjects could test their strategies in simulations against four random strategies from the previous round. In the first round they could try their strategies against a set of preprogrammed strategies. Subjects were told this could only help them to determine whether their strategy worked as intended, not whether it was a good strategy. The reason for this, it was explained is that when the ranking is determined their strategy will be playing against strategies made by other players. Like the subjects in the experiment you can try as many different strategies and as many simulations as you want.

Example of a screen to formulate a strategy (translation from Dutch). The buttons had context dependent help (showed in the red box).

Formulating the strategy

Strategy

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

	Changed	Number of players with your color				
		Win		Lose		
Previous period:	SC(1)	\$Win[1]		\$Lose[1]		
		\$W1[1]	\$W2[1]	\$L3[1]	\$L4[1]	\$L5[1]
2 periods ago:	SC(2)	\$Win[2]		\$Lose[2]		
		\$W1[2]	\$W2[2]	\$L3[2]	\$L4[2]	\$L5[2]
3 periods ago:	SC(3)	\$Win[3]		\$Lose[3]		
		\$W1[3]	\$W2[3]	\$L3[3]	\$L4[3]	\$L5[3]
4 periods ago:	SC(4)	\$Win[4]		\$Lose[4]		
		\$W1[4]	\$W2[4]	\$L3[4]	\$L4[4]	\$L5[4]
5 periods ago:	SC(5)	\$Win[5]		\$Lose[5]		
		\$W1[5]	\$W2[5]	\$L3[5]	\$L4[5]	\$L5[5]

Lost last period with 5 players with my color

```

IF (condition) {
  RETURN number;
}

ELSE {RETURN 0;}(This is always added: when none of the conditions is met you will not change color)

```

to simulation

Click [here](#) to view the complete instructions.

Example of a screen after a strategy is tried out by the participant

Simulations

with strategy:

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

Period	Self	Others	Changed	Result
101	B	<i>RRRB</i>	0	W2
102	B	<i>RRBB</i>	0	L3
103	B	<i>RRBB</i>	0	L3
104	B	<i>RRBB</i>	0	L3
105	B	<i>RRRR</i>	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

Example feedback of round by email (translated from Dutch)

Dear [name participant]

The simulations for round [round number] are finished and these are the results:
Your strategy finished on place [rank] and your earnings in this round are [earnings].

The general results

Rank	Login name	Average number of points
1	Chung-Lin KWA!	39.873644
2	gemer92	38.148873
3	Kees	37.256570
4	Witchy	36.612166
....		
35	A-town	27.996524
36	capital P	27.888102

After logging in on www.creedexperiment.nl/minor/login.php you can run simulations with your strategy of round [round number] and 4 other strategies also from round [round number]. After that you can try out new strategies against 4 strategies of round [round number] and after that submit your final strategy for round [round number + 1].

Best regards,

Jona Linde, Joep Sonnemans en Jan Tuinstra