# Imitation with Intention and Memory: an Experiment

Astrid Matthey*

January 19, 2009

**Abstract**

The experiment described in this paper analyzes imitation in an individual learning context. It supplements the results obtained for imitation in evolutionary processes.
The paper makes three main contributions. First, it provides a clear distinction between intentional imitation and genuine learning. Second, it shows that players consider more than just last period's performance when choosing imitation examples. This contrasts with the assumption in most theoretical and experimental research. Third, since the design makes imitation explicit, it allows a detailed analysis of imitation behavior, showing that players do not imitate rationally.

JEL classification: D01, D83

Keywords: imitation, learning, experiments

*Max-Planck-Institute of Economics. Corresponding address: matthey@econ.mpg.de
I wish to thank Dorothea Kübler, Jörg Oechssler, Andrew Schotter, participants at the Budapest Workshop on Behavioral Economics, EEA/ESEM Annual Congress 2007 and seminars at the Technical University Berlin and Humboldt University Berlin, as well as two anonymous referees for helpful comments. Financial support by the Deutsche Forschungsgemeinschaft through SFB 649 through part of this research is gratefully acknowledged.

# 1 Introduction

There are two reasons for imitation. First, imitation can be necessary for the development of skills. This applies to tasks that require practice, like eating with chop sticks or walking on one's hands. Second, imitation can be useful to achieve high outcomes, but no information or skill is acquired. For example, buying the same stocks as André Kostolany would not have meant to learn why he picked which stocks, but to profit from his knowledge without acquiring it oneself. In such cases, the necessary information is received *before* the imitation, when one learns what person is worthwhile imitating. Skill is not required.

The experiment described in this paper focuses on the second kind of imitation. It is designed to mimic situations where players make decisions in a changing environment, but cannot learn from imitation in the sense of acquiring skill or information through imitating others. It is particularly relevant under circumstances where players have (or think they have) different levels of information, and can observe others' current behavior.

The results provide clear evidence that people imitate even if they cannot obtain information or acquire skill by doing so. Overall, in 34% of all decisions players intentionally imitate other players. This behavior is neither consistent with genuine learning nor with randomization strategies. The result complements previous experimental research (e.g., Apesteguia et al., 2007; Bosch-Domènech and Vriend, 2003; Selten and Apesteguia, 2002) that focussed on showing whether behavior was *consistent* with an imitation strategy, but could not show whether it was actually driven by imitation motives.

The second point of the paper concerns the choice of imitation examples, that is, whom or what players imitate.[1] In the literature, players are usually assumed to repeat an action that was successful either in the period just before the decision (e.g., Vega-Redondo, 1997; Selten and Ostmann, 2001) or in some period during the time the player can recall (e.g., Alos-Ferrer, 2004; Josephson and Matros, 2004). In experiments, players only have the opportunity to imitate past strategies (e.g., Huck et al., 1999). In contrast, the hypothesis in this experiment is that players imitate other players, rather than repeating particular actions. The main difference between the two concepts is that successful players have long-term strategies that make them successful in changing environments. To distinguish between the two strategies, repeating others' actions is costless in the experiment, while imitating players is costly. The results show that in spite of these costs players frequently chose to imitate other players. The experiment also sheds some light on the use of memory when choosing which player to imitate. Even though some theoretical models account for multi-period memory (e.g., Alos-Ferrer, 2004), the experimental literature so far usually considered the case where only last period's payoffs can be recalled and used in the decision (e.g., Huck et al., 1999, Altavilla et al., 2006).[2] In this experiment, unlimited memory of all previous periods is induced by providing information on all actions and payoffs to all participants

[1] I choose the term "imitation examples" following Selten and Ostmann's (2001) "success examples".

[2] It should be noted that under this assumption it becomes undistinguishable whether actions or players are imitated.

throughout the experiment. Given this setup, the results show that for a majority of players, imitation decisions can best be explained when considering payoffs over more than one period, which contradicts the commonly used assumption.

Third, since imitation is made explicit in the experiment, it can be analyzed in detail. This shows that the observed imitation behavior is neither rational according to the theoretical benchmark, nor is it optimal given player's actual understanding of the game. Rather, the results suggest that players are either overconfident regarding their own performance or have a taste for not imitating, and that they do not learn optimally.

The situations that are mimicked in the experiment, where the environment changes and some players may be better informed than others, are common in real market situations. For example, consider an atomistic market where a new firm enters. This firm may follow the marketing strategy of another firm that proved successful in the market, without understanding yet why this strategy is actually profitable. Assuming that the strategy and payoffs of the successful firm are observable, the new firm does not learn anything from imitation that it could not learn from pure observation. However, by imitating it may ensure high payoffs until it has finally learned to successfully market its products itself.
Similar examples exist for individual decision making. Assume an investor who wants to invest in stocks and does not have much knowledge about the market. But he happens to know that George Soros recently bought shares of the publicly traded company X. By buying such stocks, the investor will not learn anything new about the market, since stock prices are observable. He may nevertheless imitate Soros, believing that this increases his payoffs until he has finally learned to make profitable investment decisions himself. In both cases, the relevant increase in information occurs *before* the imitation. The new firm learns by observation which firms are successful in the market, and the newcomer learns who are successful investors. Practicing skills is not necessary.

In most cases imitation of this kind occurs with some delay or imperfection. This suggests that the profits of the imitator are lower than those of the imitation example. To account for this fact, the experiment introduces an imitation fee. The fee varies between treatments in order to test whether the main effects are robust to different imitation "discounts". Although players quantitatively adjust their behavior to the size of the fee, no qualitative differences between the treatments are found.

Imitation processes have been analyzed theoretically by Vega-Redondo (1997), Schlag (1998, 1999), Selten and Ostmann (2001), and more recently by Fudenberg and Imhof (2006) and Levine and Pesendorfer (2007). In these models, any behavior is defined as imitation which leads to the individual choosing an action which was played last period. Only the exact action the individual should choose is determined differently. While Vega-Redondo and Selten and Ostmann use an imitate-the-best rule, Schlag proposes a proportional rule, where each strategy which yields a higher outcome than one's own is imitated with a probability proportional to the difference in outcomes.
These models' focus on last period's outcomes applies well to evolutionary contexts, where periods can be interpreted as generations, and individuals enter and drop out of the population each period. For individual learning within relatively

short time (either by people or firms), however, such a limited memory seems a restrictive assumption. Alos-Ferrer (2004), Josephson and Matros (2004) and Bergin and Bernhardt (1999) relax this restriction and include multi-period memory in their models. For example, Alos-Ferrer (2004) modifies the imitate-the-best rule, such that players remember the payoffs of $K + 1 > 1$ periods, and imitate the strategy which yielded the highest payoff in memory.

The experimental literature so far followed the evolutionary, no-memory approach of the earlier imitation models. For example, Apesteguia et al. (2007) set up a Cournot market game and test the different predictions that the theoretical models yield for this market. On the aggregate level, different informational settings allow them to distinguish between the models according to the quantities players choose, i.e., the equilibrium the market converges to. On the individual level, the authors find that many players' behavior is consistent with imitation in the majority of periods. However, given the experimental setting, it seems impossible to distinguish explicitly between behavior that looks like imitation but is the result of other strategies, and behavior which is actually driven by imitation motives.[3] This problem equally applies to the experiments of Selten and Apesteguia (2002), Huck et al. (1999), Altavilla et al. (2006). It is described in more detail below.

Consider a situation where in an oligopoly market game player 1 sets price $x$ and obtains payoff $a$. Player 2 sets price $y$ and obtains payoff $b > a$. In the next period, player 1 also sets price $y$. In the experiments mentioned above, this would be interpreted as evidence for imitation in individual behavior. This inference, however, is based on the very broad interpretation of "imitation" as any behavior which makes an individual choose an action that was successful in the past. In particular, it does not require an intention to imitate. How, then, does learning occur which is not imitation? If somebody learns by observing others that in chess the player wins who keeps his king until the end, and then tries to keep her king when playing herself - does she imitate? Or did she just learn the rules of the game, and now plays by them?[4]

In this paper, I choose a narrower definition of imitation, in order to be able to distinguish between behavior which is imitating in the colloquial sense of the word (copying, mimicking) on one hand, and behavior which involves observation and genuine learning on the other hand.[5] Imitation in the sense of practicing to acquire skills is not considered.
According to this narrower definition of imitation, players' behavior is imitating if they have an intention to do what somebody else does or did, and do not have sufficient understanding of the situation to choose a successful action themselves. This means that player 1 may truly have wanted to imitate player 2. However, he may as well have learnt the structure of the market from observation in the first period, and then used this knowledge to choose the action with the higher

[3]The authors note that a questionnaire that the participants filled in after the experiment gives hints that some players indeed intended to imitate.

[4]Even trial and error learning could be classified as imitation with this definition, namely imitation of one's own past successful actions.

[5]I define as *genuine learning* all strategies that lead to an expected increase in relevant knowledge about the game, i.e., its structure and rules. This includes complex strategies like Bayesian updating from the observation of the strategies and payoffs of others, but also simple strategies like trial and error.

expected payoff in the second.

Section 2 explains the design of the experiment. The results are presented and discussed in section 3. A brief summary and some concluding comments are provided at the end.

## 2 Experimental design

Participants in the experiment had to choose cells from a table. Each cell returned points. The pattern according to which points were matched to cells changed, but not necessarily in each period. All players knew that this change occurred according to a regular, non-trivial sequence. They also knew that other players might have more or less information than they themselves. This was made explicit in the instructions in order to avoid the implicit assumption that all players are equally informed, which may have been perceived as deception. However, players did not know for sure whether there were others with a different amount of information, nor whether they were better or worse informed than others, nor of what the difference in information consisted. I will call the better informed players "informed", and the less informed players "uninformed".

Informed players saw tables like table 1. They knew that a-cells were always best

| | | |
|---|---|---|
| a1 | b2 | c3 |
| b3 | a3 | a2 |
| c2 | c1 | b1 |

Table 1: Table shown to informed players

(13 to 15 points), b-cells were medium (7 to 9 points) and c-cells were always worst (1 to 3 points). The pattern of cells in the table changed, but players could observe this. However, the ranking of the three cells of a particular letter also changed. In some period, a1 could be best, in another a2, and in yet another a3. The sequence of this change was not known to the informed players but could be learnt during the experiment. (The sequence is described below.) So the task for informed players in each period was to choose an a-cell and find out during the experiment in which period which a-cell returned the highest number of points.[6]

Uninformed players saw table 2 in each period. This table did not contain any

| | | |
|---|---|---|
| a | b | c |
| d | e | f |
| g | h | i |

Table 2: Table shown to uninformed players

payoff-relevant information. Players knew that each cell contained points, and

[6]It was not explicitly stated that they should only choose from a-cells. But it was clearly stated that a-cells were always best.

that the pattern of these points could change every period. But they did not see the a-cells. Hence, it was their task to find out among all cells - instead of just the a-cells - which cells yielded the highest points in which period.

There were only 3 different patterns of how points were assigned to cells. Let these patterns be called X, Y, and Z. The sequence was as follows:

| | |
|---|---|
| period 1-3: | X Y Z |
| period 4-12: | Z X Y X X Z X Y Y |
| period 13-21: | Z X Y X X Z X Y Y |
| period 22-30: | Z X Y X X Z X Y Y |

This means that periods 1 to 3 were unrelated to the rest, and in period 4 a regular sequence with a length of 9 periods started, which was repeated twice.

Instead of choosing a cell themselves, players could name another player whose cell choice was then also valid for them. They had to make this decision before they saw other players' current choices, that is, when making the decision to imitate another player they did not know her cell choice.[7] For example, player 3 might have decided in period 5 to imitate player 7. Player 7 chose, say, cell c and received 14 points. Only afterwards did player 3 learn about the choice of player 7, and the number of points she received.

After each period, all players' cell choices and points were made visible to everybody, and this information about all periods was available on screen throughout the experiment. For consistency, the choices of one type of player were translated into the strategy space of the other. For example, an $a1$ choice of an informed player appeared as a $b$, $e$ etc. choice on the screen of uninformed players, depending on its position in the relevant period. Similarly, a $b$, $f$ etc. choice of an uninformed player appeared as a $b1$, $c3$ etc. choice on the screen of informed players, depending on the positions in the relevant period. This means that all information that was generated during the experiment and was relevant for learning how to obtain a high number of points was available to all players at all times. Imitating another player did not increase the players' information. Rather, since by imitating others players did not generate their own data points (no cell choice), imitation weakly reduced overall information compared to choosing one's own cell.

Players could earn between 1 and 15 points by choosing a cell themselves. If they chose to imitate a player, they received the number of points of the imitated player, minus a fee. There was a "cheap" treatment, in which this fee was one point, and an "expensive" treatment, where the fee was three points. If a player imitated a player who himself imitated another player, the first player received the points of the third player (the one who chose a cell herself), less twice the imitation fee. If this rule yielded a loss, e.g., because two players imitated each other or because a player with a low number of points was imitated, points were set to zero.

The experiment lasted for 30 periods. In each session there were 2 informed and 12 uninformed players (except for one session where for technical reasons there were

[7] The words "imitate" or "imitation" were not used anywhere in the experiment. See also the translated instructions in the appendix.

only 10 uninformed players). Four sessions were conducted for the "cheap" treatment, and three for the "expensive" treatment. In sum, there were 96 subjects, 82 of whom played as uninformed players. 48 uninformed and 8 informed players played in the cheap treatment, while 34 uninformed and 6 informed players played in the expensive treatment. This slightly unbalanced distribution is not a problem since the higher costs in the expensive treatment are not part of the research question. They serve mainly as a robustness check for the results. When averages are displayed in the analysis, the different weights of the two treatments are accounted for. All players played in only one treatment, and treatments differed only in the size of the fee.

Participants were informed that after the 30 periods one period was chosen randomly, of which payoffs were paid out immediately. Points were exchanged into EUR according to the rate 1 point = 0.7 EUR for uninformed players and 1 point = 0.5 EUR for informed players. The different exchange rates were chosen in order to account for the difference in information and give players roughly equal expected payoffs. At the end of the instructions test questions ensured that all participants understood the rules of the experiment.
Participants were students of Humboldt Universität zu Berlin and Technische Universität Berlin, with mostly non-economics majors. The experiments took place at Technische Universität Berlin in June 2006. They were computerized using z-tree (Fischbacher, 2006). Sessions lasted about 60 min. The average amount earned was 9.8 EUR, including a show-up fee.

# 3 Results

## 3.1 Summary descriptives

Of the 82 uninformed players, only 4 did not imitate at all, i.e., 78 players imitated in at least one of the 30 periods. The maximum number of imitations per subject is 26, the average number is 10.13 (11.6 in the cheap and 8.05 in the expensive treatment, all for uninformed players). The share of uninformed players who imitated per period is shown in figure 1. It shows that for both treatments imitation started off around period 5, was highest around periods 15 to 20, and dropped towards the end without ever falling back below a share of 20%. As was expected, players in the expensive treatment imitated less, showing that imitation is price sensitive. The difference is significant at $p<0.01\%$ (Mann-Whitney ranksum test).

The distribution of the number of imitations per subject is shown in figure 2.

Across all periods in which subjects did not imitate, informed players obtained on average 14.16 points, while uninformed players obtained on average 10.2 points.[8] Note that three of the 14 informed players also imitated a few times (1, 3 and 7 times).

[8] If for some measure there is a treatment effect, and averages over both treatments are displayed, these are averages over *treatments* rather than over individual players. This means that the players of the expensive treatment receive a slightly higher weight, i.o.t. ensure that the two treatments receive equal weight in the average.

Regarding the payoffs from the experiment, uninformed players earned 11.75 points and 10.23 points on average in the cheap and expensive treatment, respectively. Informed players earned an average of 14.23 points (cheap) and 13.89 points (expensive). The difference is not significant. Figure 3 shows the average payoffs players obtained depending on how many times they imitated. It shows a very diverse picture, with high payoffs being obtained for widely varying numbers of imitations. An empirically optimal level of imitation cannot be identified.

Table 3 summarizes the descriptives for the two treatments.

| | cheap | expensive | overall |
|---|---|---|---|
| number of uninformed players | 48 | 34 | 82 |
| number of informed players | 8 | 6 | 14 |
| imitation fee (points) | 1 | 3 | - |
| average number of imitations (out of 30) | 11.6 | 8.05 | 10.13 |
| average share of imitation periods | 39% | 27% | 34% |
| avrg. points from own choice, informed players | 14.35 | 13.91 | 14.16 |
| avrg. points from own choice, uninformed players | 10.48 | 9.91 | 10.2 |
| average payoffs, informed players | 14.23 | 13.89 | 14.08 |
| average payoffs, uninformed players | 11.75 | 10.23 | 10.99 |

Table 3: Summary descriptives for the two treatments

## 3.2 Rational vs. actual play

### 3.2.1 Rational play

The sequence of patterns was unknown to the participants. The only information they had was that there *is* a regular sequence of patterns.

The expected payoff for informed players from random play according to the instructions, i.e., without learning or imitation, was 14 points in each period. For uninformed players, the expected payoff from random play was 8 points. Hence, since the imitation fee was never higher than 3 points, their expected payoff from imitating an informed player dominated the expected payoff from (own) random play in all periods. This means that before uninformed players learned the sequence, imitation was the dominant strategy.

When trying to identify the sequence, players had to wait for the sequence starting to repeat itself. The first repetition of patterns [$X$, $Y$, $Z$] occurred in periods 5 and 6, that repeated the sequence of patterns in periods 1 and 2. In period 7 players learned that this repetition did not identify the sequence. The first repetition over 3 periods occurred in periods 9 to 11, which repeated the patterns in periods 4 to 6. However, in period 12 players again learned that the repetition over 3 periods did not identify the sequence. Finally, the first repetition over 4 periods occurred in periods 13 to 16, which repeated the patterns in periods 4 to 7. From period 17 onwards a player who processed the available information correctly could therefore choose the correct cells himself and avoid the imitation

fee. However, risk averse players who experienced twice that what seemed like the beginning of a repeated sequence actually was not may have chosen to observe more periods before choosing their own cell. The upper limit of a rational but risk averse strategy is period 22, when the sequence starts to repeat itself for the second time. Hence, any imitation after period 22 is irrational. So is choosing a cell oneself as an uninformed player before period 17, apart from the very early periods when no imitation examples have shown up yet, and periods 7 and 12, depending on one's risk attitude.

### 3.2.2 Actual play

As shown above, rational uninformed players would not choose their own cells in periods 8 to 11 and 13 to 16. However, the observed share of imitators in these periods is far below 1. (Periods 1 to 6 are not included in this analysis since they present merely an entry phase.) For the actual share of imitators see again figure 1. On average, in these periods 50% of players in the cheap treatment and 70% in the expensive treatment decided not to imitate.
Starting in period 17 optimal play (imitation vs. own choice) depends on the players' risk attitudes. However, after period 22 all rational players would choose their own cells. This is not the case here. As figure 1 shows, between 43% and 23% of players kept on imitating in these periods. On average, a share of 35% and 30% in the cheap and expensive treatments, respectively, imitated in the periods after period 22. Hence, compared to rational play, players on average first imitated too little (until period 16) and then too much (after period 22).

Consider now individual play. Of the 82 uninformed players, there is exactly one player who played rationally, imitating in all periods between periods 8 and 17, but not afterwards. This player did not imitate in period 7 but in period 12, suggesting that she developed a more cautious approach after the first unsuccessful attempt. She had also understood the sequence when she stopped imitating. Six more players deviated only once or twice from the rational imitation pattern, either imitating in period 23 or not imitating in some period between 8 and 16. However, two of these players did not understand the sequence completely when they stopped imitating.

Only for 22 of the 82 uninformed players (27%) and 7 of the 14 informed players (50%) it can relatively safely be assumed that they learned the sequence of patterns before period 30. These players received the maximum number of points (15) for at least the four last periods. Regarding the overall learning process, figure 4 depicts the number of points that informed and uninformed players obtained from choosing cells themselves. The increase is statistically highly significant for both types, although by design it is limited in size for informed players.

How do actual payoffs compare to the expected payoffs from optimal play? Consider informed players first. Their expected payoff from random play was 14 points per period. This is what they obtained until period 16. Starting from period 17 a rational informed player would have obtained 15 points in each period. (The risk argument is not valid for informed players due to their superior information.) Hence, the expected payoff for rational informed players is 14.47 points per period. Uninformed rational players first had to choose randomly before valid imitation

examples identified themselves. Assume that rational uninformed players would have chosen their own cells for the first five periods before imitation examples were identified, obtaining 8 points per period in expectation. Until period 16 they had then imitated an informed player (except, possibly for periods 7 and 12). This yields expected payoffs of 14 points less the imitation fee (or 7 and 14 points in periods 7 and 12). Between period 17 and 22 optimal play depends on risk attitudes, so I calculate the boundaries: choosing cells optimally oneself yields payoffs of 15 in each period after period 16. Imitating until period 22 and choosing optimally oneself only afterwards yields payoffs of 15 less the imitation fee in periods 17 to 22 and 15 points afterwards. In sum, the expected payoffs for rational uninformed players playing with rational informed players are then in the interval [12.67; 13.1] for players in the cheap treatment and in the interval [11.67; 12.37] for players in the expensive treatment. For rational uninformed players playing with randomizing informed players these intervals change to [12.47; 13.1] in the cheap and [11.47; 12.37] in the expensive treatment.

Actual average payoffs for informed players are 14.08 points. For uninformed players they are 11.75 points in the cheap and 10.23 points in the expensive treatment. All of these values are significantly lower than the lower bound of the theoretical benchmark (t-test, $p<0.01\%$). Hence, all types of players earn significantly less than the theoretical solution would suggest. This has two reasons: first, both types of players learn less and/or slower than rational players would. Second, uninformed (and the the three imitating informed) players show suboptimal patterns of imitation.

### 3.2.3 Optimal imitation given actual learning

The last section has shown that players did on average not play rationally. This section analyzes whether they played optimally given their actual knowledge. A player's decision to imitate in a certain period is optimal if her expected payoff from imitation exceeds her expected payoff from choosing a cell herself, where the latter depends on her true rather than theoretical understanding of the sequence. Since this understanding cannot be observed, the actually obtained points are taken as an (ex-post) proxy. The proxy used for the expected payoff from imitation is the maximum number of average points that a player in the same session earned in the periods prior to the imitation period, minus the imitation fee (1 or 3 points in the cheap and expensive treatment, respectively). The number of points that are included in the proxy is varied to allow for differences in players' use of memory.

On average, uninformed players earned significantly less than the expected profit from imitation when they chose a cell themselves. The value varies depending on how many periods prior to the decision are included in the proxy. The loss decreases when more periods are included (see figure 5). This result shows that on average players stopped imitating too early, i.e., when their actual understanding of the sequence was insufficient and they should have continued to imitate. There are (at least) two possible explanations: first, players may have been overconfident regarding their understanding of the sequence, and, second, they may have a taste for choosing a cell themselves. The decreasing loss when more periods are considered in the proxy is an indication that players use more than one period memory when choosing imitation examples.

After the last period in which they imitated, uninformed players in the cheap treatment earned on average 13.1 points, while those in the expensive treatment earned on average 11.96 points. This shows that on average players stopped imitating before they fully learned the sequence. The difference between treatments is lower than the difference in imitation fees (2 points), suggesting that imitation in the expensive treatment ends relatively later. Players in the expensive treatment seem not to react "linearly" to the higher imitation fee.

## 3.3 Imitation behavior in detail

### 3.3.1 Starting and ending imitation

Given that imitation behavior is neither rational nor optimal given players actual understanding of the sequence, what triggers imitation? In the period just before an uninformed player (re-)started imitation, she earned on average 6.45 points in the cheap and 7.61 points in the expensive treatment. This difference is significant ($p<0.01\%$), but not as large as the difference in imitation fees. The average difference in points in the two periods prior to imitation is -4.21 points in the cheap and -2.71 points in the expensive treatment. Hence, on average, a drop in points from choosing a cell oneself preceded the decision to imitate.

The average number of points in the period right after a player stopped imitating is 10.6 in the cheap and 10.3 in the expensive treatment. The difference is not significant. This means that the performance after players decided to stop imitating is the same across treatments and, as expected, higher than right before they started imitating. However, it also shows again that many players stopped imitating before they fully understood the sequence of patterns.

Real imitation "mistakes" occur very rarely in the data, which suggests that players chose imitation examples carefully until the last periods. There are only four cases where players were imitated who had at most 9 points in the period prior to the imitation (of 973 instances of imitation overall). There are an additional five cases where players were imitated after a "random hit", i.e., after obtaining at least 13 points in the period just prior to the imitation but no more than 9 points in the period before.

### 3.3.2 Whom do players imitate?

Whom do players imitate if they do not make their own cell choice? Figure 6 depicts the share of mistakes generated by a particular imitation rule. This is the share of imitators who imitated a player who did not obtain the maximum score according to a particular rule, i.e., a particular length of memory, across all sessions and periods. The first column of the graph shows the share of imitators who imitated a player that did not have the maximum number of points in the period prior to the imitation decision. The second column shows the share of imitators who imitated a player that did not have the maximum number of average points over the two periods prior to the imitation decision. The third column shows the share of imitators who imitated a player that did not have the maximum number of average points over the three periods prior to the imitation decision and so on.

It shows that at this highly aggregated level, the rule "imitate the player with the maximum number of points last period" returns the best fit, but rules which consider several more periods do not fare much worse.

The same question can be analyzed in a slightly different way, asking for the share of imitators for whom a certain imitation rule returns the best fit. Consider as an example a player who imitated 15 times. If for 8 of these decisions the player that was imitated had the highest number of points in the period prior to the imitation, for 6 decisions he had the highest average number of points in the two periods prior to the imitation, and for 12 decisions he had the highest average number of points in the three periods prior to the imitation, the rule "imitate the player who had the highest average number of points in the previous three periods" would be counted as best fit for this particular player. If several rules returned the same number of fitting decisions, they would be counted as best fit with a weight relative to the overall number of best fitting rules (e.g., for three best fitting rules, each is counted with one third). Note that the analysis only includes periods 12 to 30, since averages over up to 11 periods are considered.

The results over all players are summarized in figure 7. It shows that there is no clearly dominating rule. The very short-sighted rules (1 to 2 periods) and the very far-sighted rule have a slight advantage over the medium-sighted rules, but all of the considered rules provide the best fit for at least 5% of the players. Note that when even longer averages are included in the analysis, their fit is lower than for the average over 11 periods. Since longer averages lead to less periods being analyzable, they have been omitted.

It should be noted that the design is not optimally suited for the analysis of the use of memory. In many periods, the measures that are used in the above analysis are highly correlated. This means that the player who obtained the highest number of points in the previous period often also obtained the highest average number of points over the 2, 3 etc. previous periods. This reduces the ability of the analysis to discriminate between the different rules, and inhibits the test of more involved imitation rules. What can safely be inferred from the results, however, is that "imitate last period's winner" is not the only or even dominating rule that players use. Accordingly, they lend support to models which consider more than one period memory.

### 3.3.3 Switching

Considering only the periods in which a certain player imitates, i.e., disregarding intermediate periods of own cell choice, players imitated the same imitation example for an average of 7.97 periods in the cheap and 5.46 periods in the expensive treatment.

Of the 973 instances of imitation, 95 were direct switches. This means that players chose different imitation examples in consecutive periods. In the optimal case, this should take place whenever the expected payoff from imitation is higher for some other player than for the one that was imitated so far.

The average change in profit that the switches induced is 0.99 points (median 1 point), i.e., on average the decision to choose a new imitation example was rational

ex-post. The distribution of payoff changes due to switches is shown in figure 8. When taking means over all switching decisions of an individual player the picture remains similar: the average and median player increased their payoff as a result of switching (see figure 9 for the distribution).

Analyzing single switching decisions shows that there are 3 cases where players switched to imitating a player who did not obtain the maximum number of points in the previous period, 8 cases where players switched to imitating a player who did obtain the maximum number of points in the previous period, but not the maximum average number of points over the last 2, 3 and 4 periods, 5 cases where the imitated player did obtain the maximum average number of points only in the previous two periods and 4 cases where the imitated player did obtain the maximum average number of points only in the previous three periods. This shows that overall, switching was not arbitrary but occurred when the previous imitation example repeatedly obtained less points than some other player.

# 4 Summary and conclusion

The experiment described in this paper analyzes imitation behavior in a simple setup of individual choice. The merit of the experiment is to analyze imitation behavior explicitly and obtain clear evidence for players' intention to imitate.

The main results can be summarized as follows. First, players show a considerable extent of intentional imitation that cannot be explained with genuine learning strategies. On one hand, this supports the assumption of earlier studies that players do imitate other players. On the other, it shows a clear distinction between imitation and optimal play as the result of learning. This distinction could not be made in earlier experiments.

Second, the analysis indicates that when choosing imitation examples, a majority of players considers the performance of other players over more than one period. This may not seem overly surprising in the current experiment. However, it contradicts a general assumption made in most of the theoretical and experimental literature on imitation, showing that it's applicability is limited.

Third, the results show that imitation is not fully rational. On average, players imitate too little in the early periods and too much in the later periods, indicating insufficient learning. Even given their actual understanding of the game, players imitate too little on average, indicating either overconfidence or a taste for choosing cells oneself. Accordingly, for both types of players and in both treatments, average payoffs are significantly lower than the theoretical benchmarks.

When interpreting the results of the experiment, it should be noted that it tests imitation in an individual learning context. Some of the results may not or not to full extent apply to other contexts, e.g., the oligopoly markets studied in earlier experiments on imitation. The experiment was primarily designed to clearly distinguish imitation from genuine learning, and to test some of the assumptions that are often made in studies of imitation, e.g., on the use of memory. Although the design allows a detailed analysis of some of the aspects of imitation, more research is needed to fully understand the factors that govern the imitation decision and the choice of imitation examples.

# Appendix A

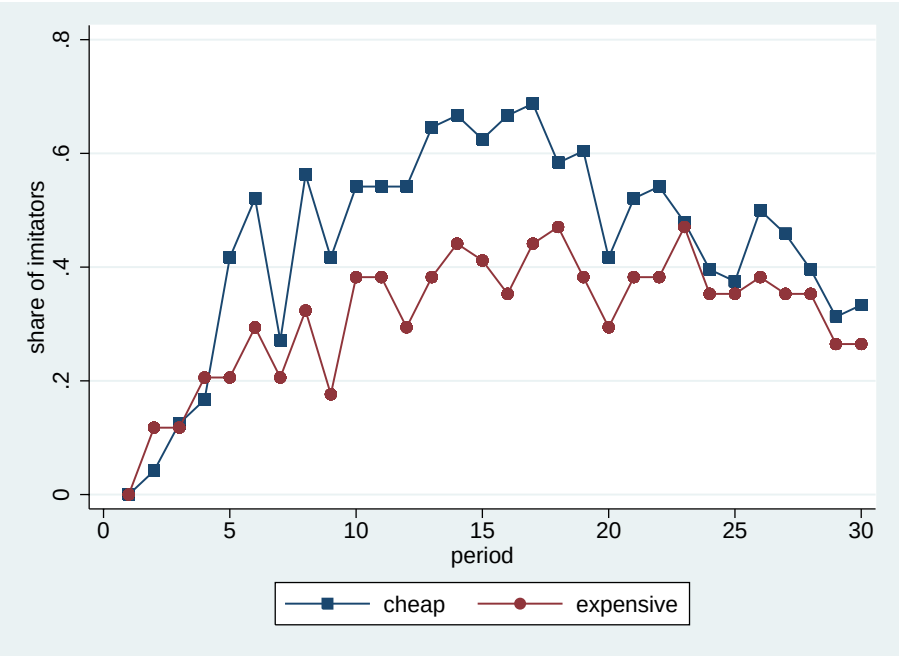


Figure 1: Share of imitations of uninformed players per period.

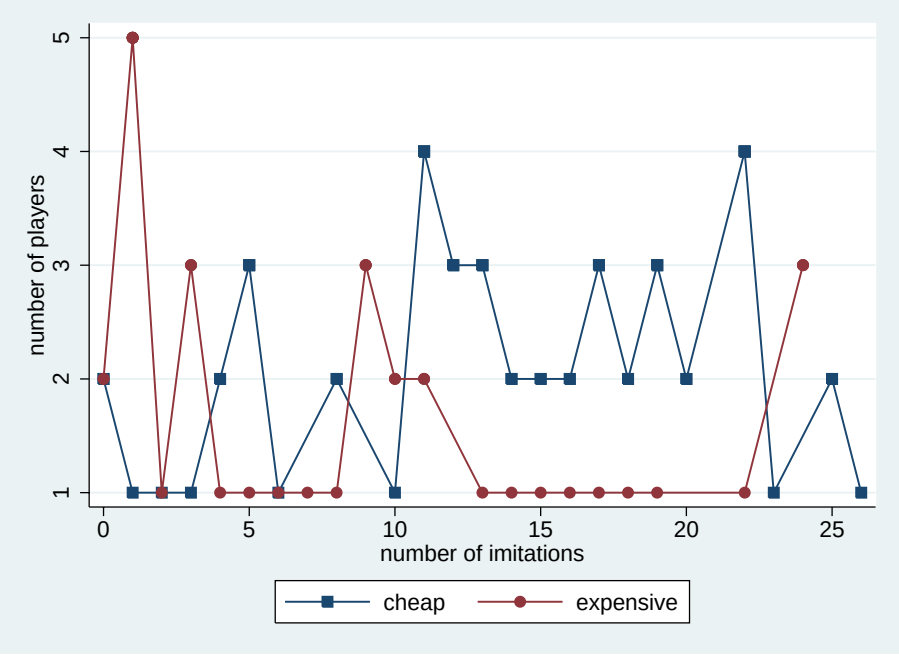


Figure 2: Number of players who imitated a certain number of times.

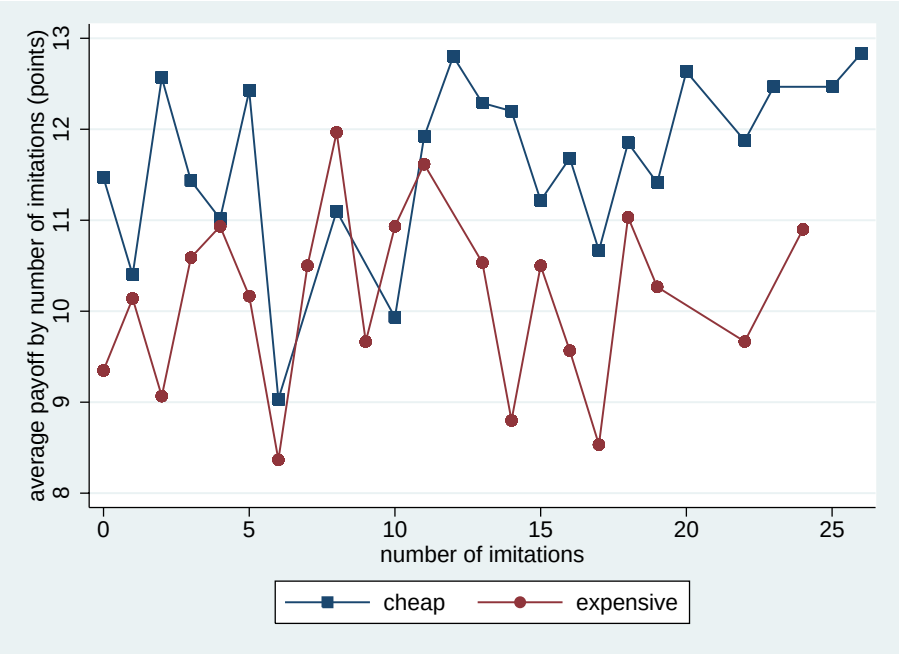


Figure 3: Average payoff per period and player, depending on the number of imitations.

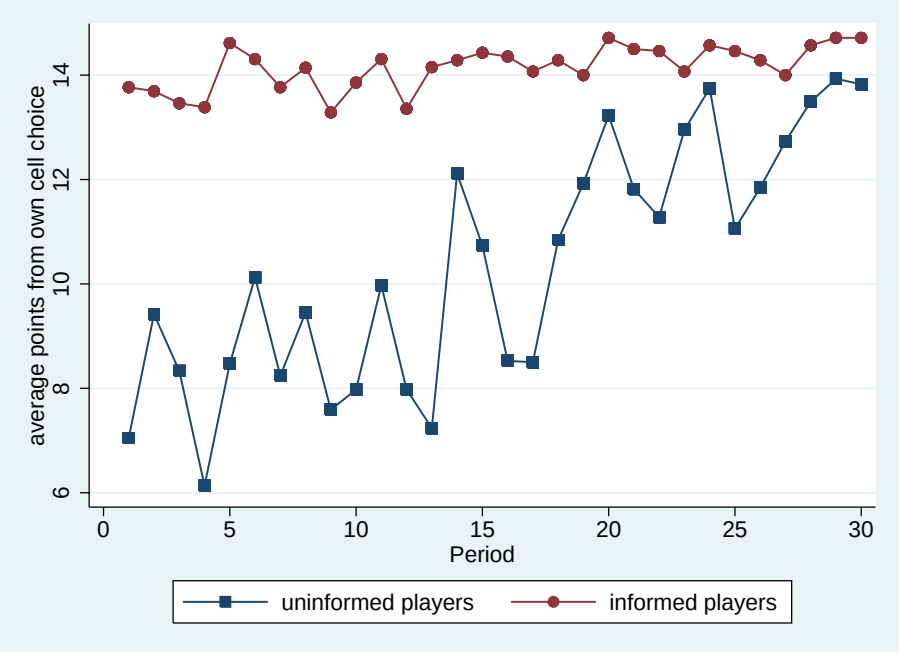


Figure 4: Average points obtained from choosing a cell oneself.

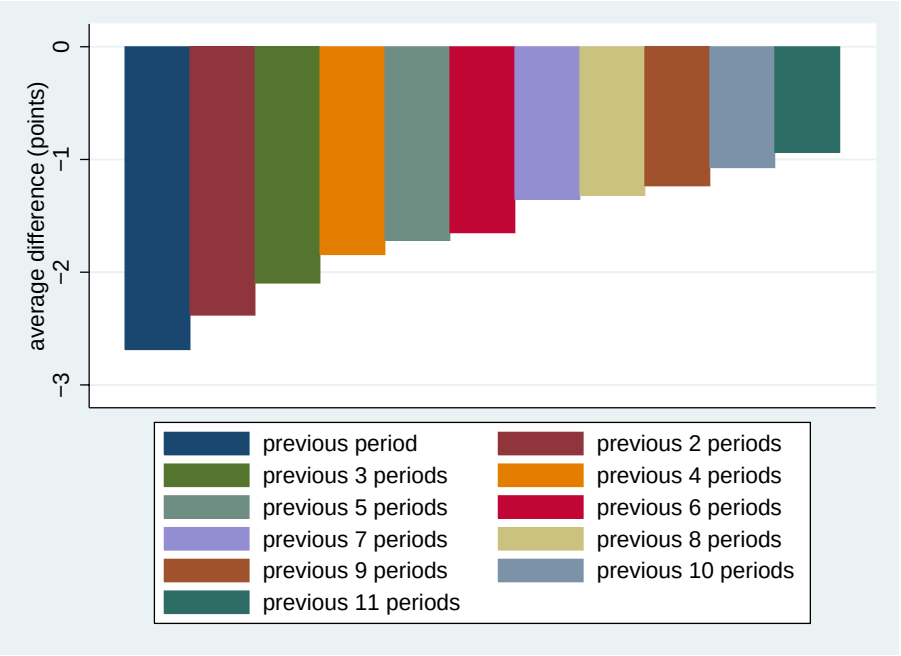


Figure 5: Average difference between points from own cell choice and expected points from imitation per subject, using proxies including only the previous period (first column) up to the 11 previous periods (last column).

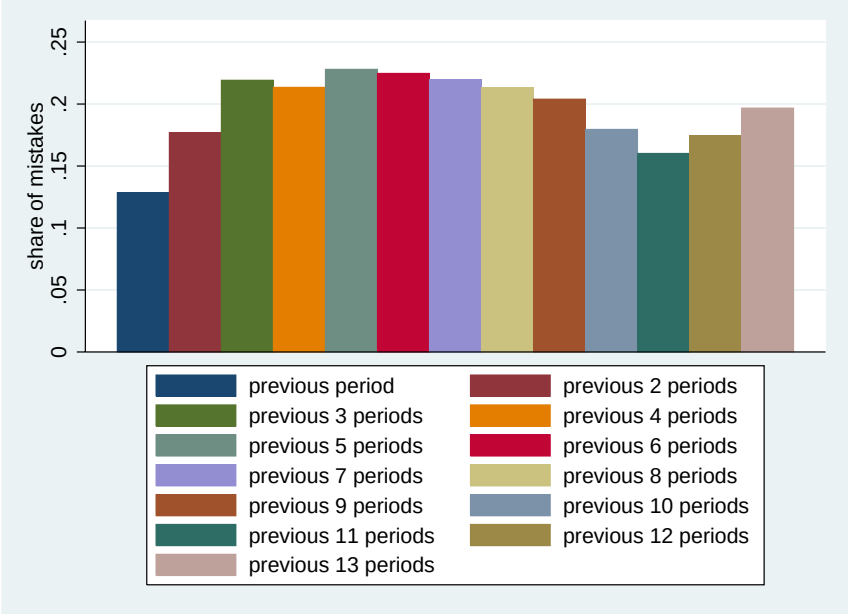


Figure 6: Share of imitators who imitate a player who wasn't an imitation example according to a particular imitation rule. Rules that are considered are "Imitate player with maximum points in the previous period" (first column), "Imitate player with maximum average points in the previous 2 periods" (second column) etc. until "Imitate player with maximum average points in the previous 13 periods" (last column).

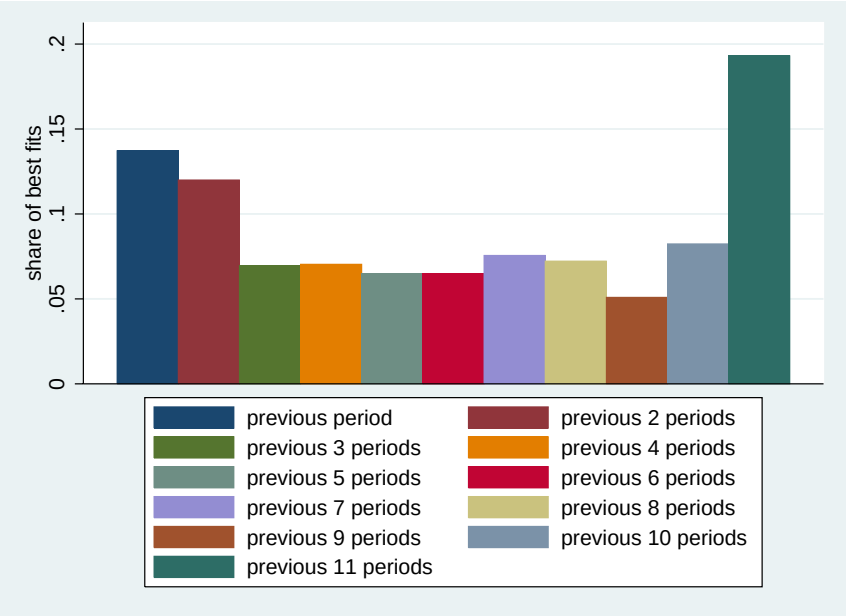


Figure 7: Share of imitators for whom a particular imitation rule returned the best fit. Rules that are considered are "Imitate player with maximum points in the previous period" (first column), "Imitate player with maximum average points in the previous 2 periods" etc. until "Imitate player with maximum average points in the previous 11 periods" (last column).

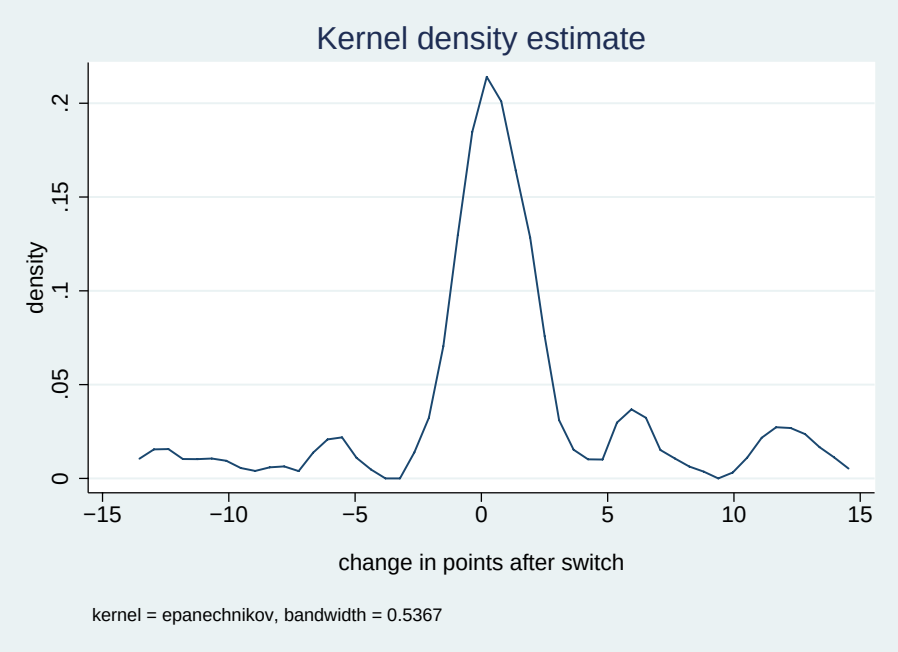


Figure 8: Distribution of changes in points after switching, over all switching decisions.

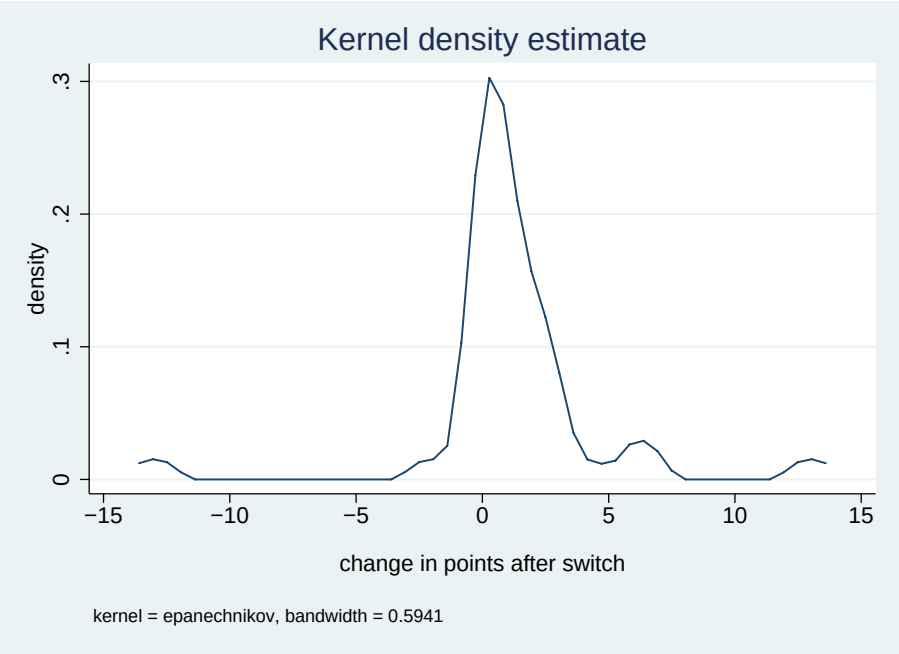


Figure 9: Distribution of mean changes in points after switching per subject, over all subjects.

# Appendix B

This is the translated version of the instructions informed players received:

## Instructions

The experiment you will now participate in is part of a research project financed by the Deutsche Forschungsgemeinschaft (DFG). It aims at analyzing economic decision making.

In the experiment you can earn a considerable amount of money, which depends on your decisions. Accordingly, it is important that you read the instructions carefully.

Please note that these instructions are only for your use. You are not allowed to pass on any information to other participants. Similarly, during the entire experiment it is not permitted to talk to other participants. If you have a question, please raise your hand. We will than come to you and answer your question. Please do not ask your question allowed. If you do not comply with these rules, we have to stop the experiment.

General information

The experiment consists of 30 periods. In each period you make a decision. According to these decisions you receive points. At the end of the experiment, these points are exchanged into Euro and paid out in cash. The course of the experiment, your decisions and the payoffs are explained in detail in what follows.

Decision

In each period a table with nine cells appears on your screen, which may look like this:

| | | |
|---|---|---|
| a1 | b2 | c3 |
| b3 | a3 | a2 |
| c2 | c1 | b1 |

Cells are always labelled with a1-a3, b1-b3, c1-c3. In each period you can choose one of those cells and receive points for your decision.

Cells that are labelled with the letter *a* return the highest number of points. cells that are labelled with the letter *b* return medium numbers of points. cells that are labelled with the letter *c* return the lowest number of points. However, cells that are labelled with the same letter can still return different numbers of points. For example, cell a1 can return a higher or a lower number of points than cells a2 and a3 etc.

The pattern of the table, i.e., the points returned by the cells, changes. This change occurs according to a certain rhythm. The logic of the letters - a for the highest points, b for medium points, c for the lowest points - is preserved by the change, but the relation of points within the letters can change. Hence, in one period cell b1 may return a higher number of points than cell b2, while in the next period b2 returns a higher number of points than b1.

When you have decided for a cell, please type it in under "Your cell choice" and press "OK". The number of points returned by the cells then appears on your screen and is added to your account. The choice of cells happens just once every round, it cannot me amended later on.

In each period, the distribution of points over the different cells is the same for all participants. It is, however, possible, that other participants have more or less information about this distribution than you do yourself.

Instead of choosing a cell yourself, you can also choose one of the other participants, whose cell choice is then also valid for you. To do so, you simply type in the number of the player you have chosen under "Your chosen player". You then receive the same number of points as this player, less one point as "fee". You only learn the cell choice of this player after you have chosen him. Note that you can only choose a cell OR a player, not both. Just as for the cell choice, the choice of a player is binding. Even if you are not happy with the cell choice of this player afterwards, you cannot change your number of points anymore.

If you want to choose a cell yourself, you simply leave the box "Your chosen player" empty.

Example:

You see the following table:

| | | |
|---|---|---|
| a1 | b2 | c3 |
| b3 | a3 | a2 |
| c2 | c1 | b1 |

Case 1: You choose cell a1 and receive 12 points. (This number of points is fictitious. It has no relation to the true number of points in the experiment.)
Case 2: You decide to follow the choice of player 3. In the box "Your chosen player" you type "3". This player chooses cell a1. He receives 12 points, and you receive 12-1=11 points.

After each period cell choices and points of all players are displayed. The first row shows the number of the period, the following rows show the cell choices and points of all players. F1 denotes the cell choice of player 1, P1 his number of points. F2 denotes the cell choice of player 2, P2 his number of points etc. The cell choices and points of all players are visible on screen throughout the entire experiment. If a player chose another player rather than a cell, the box of is cell choice stays empty.

After all 30 periods are played one period is drawn randomly, which is then relevant for payoffs. The number of points that you received in this period is then exchanged into EUR according to the following rate:

$$2 \text{ points} = 1 \text{ EUR}$$

In addition you receive 3 EUR for your participation. The sum of show up fee and the payoff from your decisions is then paid out in cash immediately.

Questions

Please answer the following questions to ensure that you have understood the instructions.

1. Player 1 chooses a cell with the label b, player 2 chooses a cell with the label c. Who receives the higher number of points?

2. Player 1 chooses a cell with the label a. Player 2 also chooses a cell with the label a. Who receives the higher number of points?

3. Player 1 decides to choose the cell that player 2 chooses. Player 2 receives 10 points. How many points does player 1 receive?

4. In period 4, player 1 chooses an a cell in the upper right corner of the table. In period 5, he again chooses the cell in the upper right corner. Does he receive the same number of points?

This is the translated version of the instructions uninformed players received:

# Instructions

The experiment you will now participate in is part of a research project financed by the Deutsche Forschungsgemeinschaft (DFG). It aims at analyzing economic decision making.

In the experiment you can earn a considerable amount of money, which depends on your decisions. Accordingly, it is important that you read the instructions carefully.

Please note that these instructions are only for your use. You are not allowed to pass on any information to other participants. Similarly, during the entire experiment it is not permitted to talk to other participants. If you have a question, please raise your hand. We will than come to you and answer your question. Please do not ask your question allowed. If you do not comply with these rules, we have to stop the experiment.

General information

The experiment consists of 30 periods. In each period you make a decision. According to these decisions you receive points. At the end of the experiment, these points are exchanged into Euro and paid out in cash. The course of the experiment, your decisions and the payoffs are explained in detail in what follows.

Decision

In each period the following table appears on your screen:

| | | |
|---|---|---|
| a | b | c |
| d | e | f |
| g | h | i |

cells are labelled with the letters $a$ to $i$. In each period you can choose one of those cells and receive points for this. The letters do not contain any information about the numbers of points that are assigned to these cells. They simply serve to improve clarity. The points that are assigned to these cells change. However, this change occurs according to a regular pattern.

When you have decided for a cell, please type it in under "Your cell choice" and press "OK". The number of points returned by the cells then appears on your screen and is added to your account. The choice of cells happens just once every round, it cannot me amended later on.

In each period, the distribution of points over the different cells is the same for all participants. It is, however, possible, that other participants have more or less information about this distribution than you do yourself.

Instead of choosing a cell yourself, you can also choose one of the other participants, whose cell choice is then also valid for you. To do so, you simply type in the number of the player you have chosen under "Your chosen player". You then receive the same number of points as this player, less one point as "fee". You only

learn the cell choice of this player after you have chosen him. Note that you can only choose a cell OR a player, not both. Just as for the cell choice, the choice of a player is binding. Even if you are not happy with the cell choice of this player afterwards, you cannot change your number of points anymore.

If you want to choose a cell yourself, you simply leave the box "Your chosen player" empty.

Example:

You see this table:

| | | |
|---|---|---|
| a | b | c |
| d | e | f |
| g | h | i |

Case 1: You choose cell a and receive 12 points. (This number of points is fictitious. It has no relation to the true number of points in the experiment.)
Case 2: You decide to follow the choice of player 2. In the box "Your chosen player" you type "3". This player chooses cell a. He receives 12 points, and you receive 12-1=11 points.

After each period cell choices and points of all players are displayed. The first row shows the number of the period, the following rows show the cell choices and points of all players. F1 denotes the cell choice of player 1, P1 his number of points. F2 denotes the cell choice of player 2, P2 his number of points etc. The cell choices and points of all players are visible on screen throughout the entire experiment. If a player chose another player rather than a cell, the box of is cell choice stays empty.

After all 30 periods are played one period is drawn randomly, which is then relevant for payoffs. The number of points that you received in this period is then exchanged into EUR according to the following rate:

$$1.5 \text{ points} = 1 \text{ EUR}$$

In addition you receive 3 EUR for your participation. The sum of show up fee and the payoff from your decisions is then paid out in cash immediately.

Questions

Please answer the following questions to ensure that you have understood the instructions.

1. Player 1 chooses a cell with the label e, player 2 chooses a cell with the label g. Who receives the higher number of points?

3. Player 1 decides to choose the cell that player 2 chooses. Player 2 receives 10 points. How many points does player 1 receive?

4. In period 4, player 1 chooses cell c. In period 5, he again chooses cell c. Does he receive the same number of points?

## References


Alos-Ferrer, C. (2004), "Cournot versus Walras in dynamic oligopolies with memory", International Journal of Industrial Organization, 22, 193-217.

Apesteguia, J., Huck, S., and Oechssler, J. (2007), "Imitation - Theory and Experimental Evidence", Journal of Economic Theory, 136(1), 217-235.

Altavilla, C., Luini L., and Sbriglia, P. (2006), "Social Learning in Market Games", Journal of Economic Behavior and Organization, 61(4), 632-652.

Bergin, J. and Bernhardt, D. (1999) "Comparative Dynamics" Working Papers 981, Queen's University, Department of Economics.

Bosch-Domènech, A., and Vriend, N.J (2003), "Imitation of Successful Behaviour in Cournot Markets", Economic Journal, 113, 495-524.

Fischbacher, U., (2006): z-Tree: Zurich Toolbox for Ready-made Economic Experiments, Experimental Economics 10(2), 171-178.

Fudenberg, D. and Imhof, L. (2006), "Imitation processes with small mutations", Journal of Economic Theory, 131(1), 251-262.

Huck, S., Normann, H.T., and Oechssler, J. (1999), "Learning in Cournot Oligopoly: An Experiment", Economic Journal, 109, C80- C95.

Josephson, J., and Matros, A. (2004), "Stochastic imitation in finite games", Games and Economic Behavior, 49, 244-259.

Levine, D., and Pesendorfer, W. (2007), "The evolution of cooperation through imitation", Games and Economic Behavior, 58(2), 293-315.

Schlag, K. (1998), "Why Imitate, and If So, How? A Boundedly Rational Approach to Multi-armed Bandits", Journal of Economic Theory 78, 130-56.

Schlag, K. (1999), Which One Should I Imitate?, Journal of Mathematical Economics, 31, 493-522.

Selten, R. and Apesteguia, J. (2002), "Experimentally Observed Imitation and Cooperation in Price Competition on the Circle", Bonn Econ Discussion Paper 19/2002.

Selten, R. and Ostmann, A. (2001), "Imitation Equilibrium", Homo Oeconomicus, 43, 111-149.

Vega-Redondo, F. (1997), "The Evolution of Walrasian Behavior", Econometrica, 65, 375-384.