Can You Guess the Game You're Playing?[1]

Jörg Oechssler[2] and Burkhard Schipper[3]
Department of Economics
University of Bonn

– revised version –

March 6, 2002

[1] Comments by Uri Gneezy, Steffen Huck, Georg Kirchsteiger, Atanasios Mitropoulos, Wieland Müller, and participants of the ESA meetings in New York are highly appreciated. Suggestions by an anonymous referee and associate editor were also most helpful. We are grateful to the Bonn Graduate School of Economics and the Deutsche Forschungsgemeinschaft, grant OE-198/1/1, for financial support. We also thank Özgür Gürerk for very able programming assistance.

[2] Adenauerallee 24, 53113 Bonn, Germany, oechssler@uni-bonn.de

[3] burkhard.schipper@wiwi.uni-bonn.de

Abstract

Recently there has been much work on learning in games. However, learning usually means “learning about behavior of opponents” rather than “learning about the game” as such. Here we test in an experiment whether players in a repeated encounter can learn the payoff structures of their opponents by rewarding subjects for correct guesses. Our data allows to construct the games that subjects perceive to be playing, the subjective games. We find that subjects often play according to an equilibrium in their subjective game. However, subjective games frequently differ from the games actually played.

Key words: learning, subjective games, experiments
JEL–classification numbers: C72, C92, D83.

Proposed running head: “Can You Guess the Game You're Playing?”

Please send proofs to:
Prof. Jörg Oechssler
Department of Economics
University of Bonn
Adenauerallee 24
53113 Bonn
Germany

1 Introduction

The theoretical and experimental literature on learning in games has substantially increased in recent years (see e.g. Fudenberg and Levine, 1998). By and large, however, this literature is concerned with learning how to play a game rather than with learning about a game. That is, the question of how players perceive a game has rarely been address so far. A normal form game consists of the set of players, the set of possible strategies, and a payoff function for each player. Learning about a game therefore means that players, who have incomplete knowledge about some of these elements, learn about those elements while playing the game.

The purpose of the experiment reported in this paper is to test how players perceive a game which they are only incompletely informed about. As a first step we let the set of players and the set of strategies be commonly known. Furthermore, each player knows his own payoff function. However, he does not know the payoff function of his opponent. After letting fixed pairs of subjects play a given 2×2 game for 15 rounds, we ask subjects to answer a questionnaire about the payoff function of their opponents and reward them for correct guesses. Subsequently, the same pairs play the game for another 5 rounds with high powered incentives.

This allows us to address a number of questions that have previously not been tested experimentally. (1) Do subjects learn to correctly perceive the game? (2) Regardless of how question (1) is answered, do subjects play according to a Nash equilibrium prediction? And (3) what is the relationship between questions (1) and (2)?

We find that subjects are not very successful in guessing the payoff function of their opponents after 15 rounds with the same opponent. Even the pure strategy best reply structure of their opponents seems to be quite difficult to guess. Despite this, subjects are remarkably good in playing according to a Nash equilibrium in the final rounds. Hence, it comes as no surprise that there is no significant relation between the degree to which subjects answered the questionnaire correctly and the instances of Nash play.

Apparently, the games as **perceived** by subjects differ substantially from the **objective** games. Our questionnaire data allow us to construct the **subjective** game as perceived by each player. In our case, the subjective game for player i consists of the commonly known strategy set for both players, the known payoff function of player i, and the perceived payoff function of i's opponent. After constructing the subjective game from the questionnaire data, we can test whether subjects choose an equilibrium strategy in their subjective game. At least for games with a pure strategy equilibrium it turns out that this hypothesis is strongly supported by our data.

Following Kalai and Lehrer (1995), we define a subjective Nash equilibrium as a combination of a strategy profile and a set of beliefs about the payoff function of one's opponent that satisfies two conditions.[1] First, each player optimizes given his beliefs. Second, the beliefs are not contradicted by the actual play. In other words, a strategy combination is a subjective equilibrium if it is an equilibrium in the subjective games of both players. We find that when a subjective Nash equilibrium exists, it is indeed chosen by a large majority of players. Furthermore, when one player perceives a subjective game with two equilibria while the other player perceives a game with just one of those equilibria, the subjective Nash equilibrium provides an interesting new form of equilibrium selection.

The remainder of the paper is structured as follows. In the next section we explain the design of the experiment. Section 3 contains the experimental results and a discussion of subjective games. A short conclusion is given in Section 4. The instructions of the experiment are reprinted in the Appendix.

[1] Related equilibrium concepts were used, among others, by Osborne and Rubinstein (1998), Huck and Sarin (1999), and Sethi (2000). For similar concepts in extensive form games see e.g. Fudenberg and Levine (1993).

2 Experimental design

In a computerized[2] experiment we studied learning behavior in normal form games with incomplete information. Information was incomplete inasmuch as subjects were only told their own payoff matrix but not the payoffs of their opponents. The experiment was divided into three blocks. In each block subjects were matched with one other subject to play the same 2×2 game for 20 rounds.

The following three games were used in Blocks 1, 2 and 3, respectively.

Block 1

	X	Y
X	4,4	0,2
Y	2,0	3,3

Block 2

	X	Y
X	2,2	3,3
Y	0,4	4,0

Block 3

	X	Y
X	3,3	0,4
Y	4,0	2,2

The three games were selected to have as much variety as possible with 2×2 games as they cover all generic types of best reply structures for 2×2 games. We see the study of 2×2 games as a first step towards more complicated games. If subjects have difficulties in learning the payoff structures of such simple games, they will have even more problems with larger games. Furthermore, there is a practical problem. In larger games it would have taken subjects probably too long to sample all cells of the payoff table sufficiently often.[3]

The first game, a variant of the stag–hunt game, is a coordination game with two strict Nash equilibria in pure strategies and one in mixed strategies. The equilibrium (X,X) is payoff dominant, while (Y,Y) is risk dominant. Note also that (Y,Y) results if both players choose their maxmin strategy.[4]

The second game is asymmetric with a unique mixed strategy equilibrium $\{(4/5, 1/5), (1/3, 2/3)\}$. When both players choose their maxmin strategy, (X,X) would be the outcome. Finally, the third game is a prisoners'

[2] We used the software "RatImage" by Abbink and Sadrieh (1995).

[3] Note, however, that in games with a natural order on the strategy space, subjects may coordinate on an equilibrium even if the number of strategies is very large and when subjects have no information about the payoff function (see e.g. Van Huyck et al., 1996).

[4] A player's maxmin payoff is the maximal payoff he can guarantee himself when using only pure strategies. A strategy that ensures this payoff is called a maxmin strategy.

dilemma with Y being the dominant strategy for both players. (Y,Y) is both, the unique Nash equilibrium and the maxmin outcome.

We selected the payoffs in all games to be 0, 2, 3, and 4 in order to have a simple best reply structure without ties. The 0 was chosen instead of 1 such that there is an unambiguous risk dominant equilibrium in the first game.

In each block the same 36 subjects were anonymously matched, either as row– or as column–players, with one other subject. The matching was done such that no subject would meet a prior opponent or someone who was opponent of a prior opponent. This rule was known to subjects. The only feedback subjects received while playing was their own payoff.

In the first 15 rounds subjects were simply paid according to the entries of their payoff matrix. Payoffs were denominated in "Taler". The exchange rate for German Marks ($1T = 0.06$ DM) was known. The last 5 rounds were conducted with high powered incentives as subjects were paid four times the respective entries of their payoff table.

After round 15, subjects were presented by the computer with a questionnaire which consisted of 4 questions regarding the unknown payoff table of their opponent. The questionnaire was announced and explained to subjects in the instructions so that they could use the first 15 rounds for experimentation. The opponent's payoff table was presented as shown below. Subjects were told that the payoffs 0, 2, 3, and 4 would each occur once.

Payoff table of other subject

		others's action X	others's action Y
your action	X	a	b
	Y	c	d

The following four questions were asked in the questionnaire

1. Is a greater than b? yes or no?

2. Is c greater than d? yes or no?

3. Where is the highest payoff? a, b, c, or d?

4. Where is the lowest payoff? a, b, c, or d?

Questions 1 and 2 together should reveal whether subjects understand the (pure strategy) best reply structure of the game, which captures the main strategic aspects of a game. In particular, the best reply structure determines whether the games has one, two, or none pure strategy Nash equilibria. Questions 3 and 4 provide further information about subjects' beliefs with respect to their opponents' payoff matrix, which are important in the context of payoff dominance and maxmin considerations.

Intentionally, we refrained from asking subjects to fill in payoffs in a matrix since this is – at least for non–game theorists – a very unusual exercise, that could result in unreliable data. The format of the questions was intended to make subjects think about the payoff structure of the game in question.

For each correct answer subjects were paid $30T$. Thus subjects' total payoff was the sum of the first 15 rounds in each block with simple incentives plus the sum of the payoffs in the last 5 rounds in each block with quadrupled incentives plus $30T$ for each correct answer. There was no feedback about the questions until the very end of the experiment.

Payoffs were chosen such that the payoffs from the questions and the cumulative payoffs from playing the games were about equally weighted. Low powered incentives in the first 15 rounds were designed to encourage experimentation. The high powered last 5 rounds should reveal whether subjects understood the game. In those five rounds they were able to play without concern for the questionnaire.

The experiment took place in the Bonn Laboratory in May 2000. Subjects were recruited from all over campus. Only about half of the subjects were economics students, few of those had any formal training in game theory. The average payoff in the experiment was DM 29 with sessions lasting about 100 minutes including instruction time. Instructions (see Appendix A) were written on paper and distributed in the beginning of each session. Before we started the experiment subjects had the chance to ask questions about the setup of the experiment.

3 Results

Our experimental design allows us to address four questions, which are taken up in turn in the next four subsections. First, did subjects learn about that part of the game that they were not informed about? Namely, can subjects find out the (pure strategy) best reply structure of the game? Or can they guess where the maximal payoffs of their opponents are located? Secondly, regardless of how question 1 is answered, do subjects play according to a Nash equilibrium? And is there a connection between subjects' answers to the questions and their play in the final, high–powered rounds? Third, can subjects learn how to play the game? For this question we compare our data to two boundedly rational learning theories. Finally, we construct the subjective games that subjects perceive to be playing and ask whether subjects play equilibrium strategies in their subjective games.

3.1 Can they learn the game?

Each subject answered 4 questions about each of the 3 games played in the experiment. One should think that, after playing 15 rounds with the same opponent, subjects should have gotten a pretty good idea about the payoff structure of their opponents. However, on average subjects were not very good in answering the 12 questions. The mean number of correct answers was 6.19 with a standard deviation of 2.56. Thus only about half of the questions were answered correctly. There was, however, substantial variation between subjects, between games, and between questions. For example, there were two subjects (of 36) who answered all 12 questions correctly, while one subject answered none of the questions correctly. A few answers were inconsistent, e.g. when a subject answered that $a > b$ but then designated b to be the maximum payoff.

With regards to the different games there are clear differences too. Table 1 shows the percentage of correctly answered questions in each block, separate for each question.

[Table 1 about here]

In general, subjects rarely got all the questions right. Only between 13.9 and 19.4% of subjects answered all four questions of a particular block correctly. Subjects did somewhat better with the first two questions, which determine whether the pure strategy best reply structures of the games are understood. Between 36.1% and 61.1 % of subjects answered both, question 1 and 2 correctly.

The impression that subjects answered significantly more questions correctly for the coordination game in Block 1 than in both, the mixed strategy game (Block 2) and the prisoners' dilemma (Block 3) was confirmed by a Wilcoxon matched–pairs signed–ranks test (one–sided p-values of 0.015, and 0.048, respectively). There was no significant difference between the number of correct answers to questions in Blocks 2 and 3.[5]

In principle, Nash equilibrium play requires that **both** players perceive the payoff structure correctly. If one considers the number of **pairs** in which both subjects answered questions correctly, results are even less favorable to the hypothesis that subjects can learn the payoff structure of a game easily. Table 2 shows that between 0 and 2 pairs (of 18) were able to answer all questions correctly, between 1 and 7 pairs guessed at least the best reply structure.

[Table 2 about here]

[5] A reason why subjects were not very successful in guessing the best reply structure of the prisoners' dilemma, which has a rather simple structure, might be that subjects, faced with a dominant strategy, were not trying too hard to guess the payoffs of their opponents.

Again one can see that the coordination game in the first block was apparently easier to guess than the other two. In particular, the mixed strategy game in Block 2 seemed to be difficult to understand. Of course, one important aspect here is that game 2 is an asymmetric game.

3.2 Do they play Nash (nevertheless)?

Given subjects' very incomplete understanding of the games, it might come as a surprise that their actual behavior was fairly close to the Nash equilibria of the respective games. While the first 15 rounds of each block were characterized by experimentation as intended (subjects switched among their strategies on average between 3 and 4 times), the final 5 rounds with quadrupled payoffs show remarkable regularities in behavior.

Block 1

- 16 of 18 pairs perfectly coordinated for all 5 rounds.
- 17 of the 18 pairs coordinated for at least 4 of the 5 final rounds on one of the pure strategy equilibria.
- only 4 of those 17 coordinated on the risk dominant equilibrium (Y,Y), whereas 13 coordinated on the payoff dominant equilibrium.

Block 2

The game in Block 2 possesses no equilibrium in pure strategies. Aggregate play of subjects is probably best described by Table 3 which lists the empirical frequencies with which each cell of the matrix was being played and compares those to the mixed strategy equilibrium.

[Table 3 about here]

As can be seen the empirical frequencies are not too far from the theoretical ones with the major deviation being that column–playing subjects chose their first strategy too often.

Block 3

- 8 of 18 pairs perfectly played the Nash equilibrium for all 5 rounds.
- 14 of the 18 pairs played the Nash equilibrium for at least 4 of the 5 final rounds.
- only 1 pair managed to cooperate on (X,X) for 3 rounds with cooperation breaking down in the last two rounds – probably due to endgame effects.

The data we have collected allow us to check for a number of relationships that have hitherto not been studied. First, we can study the relation between the way subjects answered the questions and their observed behavior in the final 5 rounds. Second, how does the pre–question play in the first 15 rounds influence the answering of the questions?

An obvious hypothesis is that subjects (or rather pairs of subjects) who understood the best reply structure of their opponents are better able to play according to a Nash equilibrium than subjects who did not. There is, however, no significant relationship in the data. Consider, for example, the number of pairs of subjects who played a Nash equilibrium in at least 3 of the last 5 rounds. Are those the same subjects who answered the questions correctly? Table 4 shows that this is apparently not the case. In Section 3.4 we will see, however, that there does exist a different interesting relationship between questions and observed play based on "subjective games".

[Table 4 about here]

How does pre–question play influence the answering of the questions? As expected, subjects used the first 15 rounds mainly for experimentation. As Table 5 shows, pairs of subjects who switched their actions more often, generally were better in answering the questions.

[Table 5 about here]

It is also quite intuitive that one can learn more from a strategy switch when only one player moves at a time ("horizontal" and "vertical" switches). When both players switch at the same time ("diagonal" switch), the result may be confusing, as for Block 1.

3.3 Learning how to play

Additionally to asking whether subjects learn the game they are playing, we may also ask whether subjects learn how to play the game. Given the minimal information and feedback subjects receive about the game, two recently much studied low–information learning theories seem to be applicable, namely, reinforcement learning and payoff assessment learning. We perform the same exercise as in Sarin and Vahid (2001) by comparing the basic reinforcement model of Erev and Roth (1998) with the payoff assessment model of Sarin and Vahid (2001). Since those theories are described in detail in Sarin and Vahid (2001), we will omit this here. Both theories have one parameter which is chosen to minimize the mean squared deviation (MSD) between the data from all three blocks and the average of 200 simulations with the learning models.[6]

[Table 6 about here]

Table 6 lists the MSDs for the reinforcement model (RE) and the Sarin–Vahid model (SV), each for the optimal parameter value (which is reported in the second column). The last column presents the mean MSD for all three blocks. Columns 3 through 5 list MSDs when the model with the overall optimal parameter is applied to a single block only. The models RE* and SV* are like RE and SV except that payoffs in the final 5 rounds of the simulations were counted simply, not fourfold (as in the experiment). The differences between those model variants are negligible, however. As a

[6] All details of the simulations are as in Erev and Roth (1998) and Sarin and Vahid (2001). We thank Farshid Vahid for providing his GAUSS program for the SV model.

comparison we also list the MSD scores for the equilibrium prediction[7] and for random choice (each strategy with probability 1/2).

Comparing the overall MSD scores in Table 6 we see that the RE model performs slightly better than the SV model, though both are not substantially better than simple random play. The equilibrium prediction performs considerably worse. The fact that both learning theories are very similar in their prediction can also be seen when we plot the average proportion of strategy X choices in our 200 simulations against the observed frequencies in the experiment (see Figure 1). Both learning theories track the experimental data reasonably well for Block 2, when there is a unique mixed strategy equilibrium, and to some extent for Block 3, the prisoners' dilemma. In the coordination game of Block 1, however, both learning theories are far off the mark as they cannot mimic the almost perfect coordination of many subjects on the payoff dominant equilibrium. In evaluating the performance of learning theories we should keep in mind, though, that in our experiment subjects were encouraged to use the first 15 rounds for experimentation. Thus, it should not be surprising that random play explains the data quite well for the early rounds. In particular, it may be the case that subjects in the final 5 rounds choose strategies which performed better in the experimentation phase.

[Figure 1 about here]

The last idea can be explored further by investigating the relationship between pre–question play and post–question play. Despite the fact that subjects knew that behavior during the pre–question period partly reflected experimentation on the side of their opponents, we find evidence that subjects chose to a large extent in the final 5 rounds those strategies that were on average more successful in the first 15 rounds. Let x_i and y_i denote subject i's average payoffs received when playing strategy X and Y, respectively.

[7] In Block 1 we selected the payoff dominant equilibrium (X,X) since it had a much lower MSD than the risk dominant equilibrium.

Let f_i denote the relative frequency with which subject i used his strategy X in the final 5 rounds. We estimated the following equation with ordered logit and probit models separately for each block

$$f_i = const. + \beta(x_i - y_i) + \varepsilon.$$

For Blocks 1 and 2, β was highly significant (p-values < 0.01) and positive as expected, which gives support to reinforcement type learning models. However, for Block 3, β is not significant different from zero, which indicates that subjects used the dominant strategy Y in the prisoners' dilemma despite the fact that they sometimes received higher payoffs with strategy X.

3.4 Subjective games

When experimental subjects do not play according to the Nash equilibrium prediction, at least two interpretations are possible. One is that subjects simply do not know how to find a Nash equilibrium or do not want to play according to it. The alternative is that subjects have a wrong perception of the game they are playing and choose a Nash equilibrium of the perceived, subjective game. There are a number of authors who have recently begun to study such alternative approaches (see e.g. Kalai and Lehrer, 1995; Matsushima 1998a,b; Osborne and Rubinstein, 1998; Huck and Sarin, 1999; and Sethi, 2000).

While it is difficult to test any of these theories directly (e.g. because they are often concerned with asymptotic results), our questionnaire data give us the opportunity to experimentally study the general idea that players form a subjective game in their mind and play accordingly. This is done by using the questionnaire answers to construct a subjective game whenever possible, i.e. when the answers were consistent. We define a **subjective game** for player i by taking the (objective) payoffs of player i and using i's answers to construct i's belief about j's (pure strategy) best reply structure.[8] We call

[8] For some constellations of answers, j's payoff matrix cannot be recovered in a unique way. However, it is always possible to identify the pure strategy best reply structure.

s_i a **subjective equilibrium strategy** if s_i is a pure Nash equilibrium strategy in i's subjective game.[9]

A strategy profile (s_1, s_2) is a **subjective Nash equilibrium** if it is a Nash equilibrium in the subjective games of both, player 1 and 2.[10] We count a prediction as compatible with the data if the actual play is consistent with the prediction in at least 3 out of the 5 last rounds.[11]

In Block 1 the play of all subjects with consistent answers was compatible with their subjective equilibrium strategies.[12] Furthermore, of the 13 pairs of players with consistent answers, 12 pairs played according to a subjective Nash equilibrium for at least 3 rounds. For example, player 9 perceives a subjective game with a unique Nash equilibrium XX. His opponent, player 10, perceives a game with two pure Nash equilibria, XX and YY. Thus, the unique subjective Nash equilibrium is XX, which is indeed played by both subjects in 4 of the 5 rounds. It seems that the wrong perception of the game works as equilibrium selection device for the underlying coordination game.[13] In fact, this kind of selection device was applicable for 6 out of the 13 pairs, where in 5 out of these 6 cases the payoff dominant equilibrium was selected over the risk dominant equilibrium.

The game in Block 2 seems to be much more difficult for subjects to learn. Only 13 of 35 subjects with consistent answers correctly perceived a subjective game with a unique mixed strategy equilibrium. Of the 22 who perceived a game with at least one pure strategy equilibrium, 17 followed a subjective equilibrium strategy. Interestingly, there was no subjective Nash equilibrium at all, which shows that the perception of the game was quite

[9] If there is no pure strategy equilibrium in the subjective game for some player i, then both pure strategies in the support of the unique mixed strategy equilibrium could be considered to be subjective equilibrium strategies, which, however, would have no predictive power. Therefore, we did not count cases in which a player's subjective game admits only a mixed equilibrium.

[10] Thus, as in Kalai and Lehrer (1995) the beliefs of players must not be contradicted when they play a subjective Nash equilibrium.

[11] This yields a significance level of 10% according to a Binomial test.

[12] Recall that 5 subjects answered in an inconsistent way.

[13] Of course, a subjective equilibrium need not correspond to a Nash equilibrium of the underlying game.

different for the two player types.

In Block 3, 32 of the 35 subjects with consistent answers followed a subjective equilibrium strategy. For 7 pairs of players there was a subjective Nash equilibrium YY, which was actually played by 5 of those pairs in at least 3 rounds. For one pair there was a subjective equilibrium YX, which was, however, played only in one out of five rounds.

The specific theory that is closest to our experimental framework is by Matsushima (1998b) who considers repeated symmetric 2×2 two–person games where players do not know the objective payoff function, $u(s_1, s_2)$.[14] Based on an adaptive learning rule introduced in Matsushima (1998a), the theory predicts that both players choose their maxmin strategy in a class of games satisfying certain conditions. Formally, let $s_1, s_2 \in S$ denote pure strategies and

$$\underline{\nu}(s_1) := \min_{s_2} u(s_1, s_2)$$
$$\overline{\nu}(s_1) := \max_{s_2} u(s_1, s_2)$$

the player's minimal and the maximal payoff for action s_1, respectively. Further, let s_1^* be the unique pure maxmin action such that

$$\underline{\nu}(s_1^*) > \underline{\nu}(s_1), \forall s_1 \neq s_1^*.$$

For 2×2 games that satisfy what Matsushima (1998b) calls the strong property of strategic coordination, namely

$$u(s_1^*, s_1^*) = \overline{\nu}(s_1^*),$$

the learning process assures that players perceive a subjective game in which there is no strategic conflict with respect to implementing the subjectively efficient outcome. In other words, both players choose their maxmin action because they believe (falsely) that this action is strictly dominant for both players and that the resulting outcome is Pareto efficient.

[14] Thus, in contrast to our experimental setting, Matsushima assumes that players do not know even their own payoff matrix. However, the own payoff matrix should not be too difficult to learn in 2×2 games.

The coordination game in Block 1 satisfies the strong property of strategic coordination, whereas the prisoners' dilemma (Block 3) does not. The game of Block 2 does not apply because of asymmetry.

The predictions for the game in Block 1 are that players (1) come to play the risk dominant Nash equilibrium (Y,Y) and (2) perceive strategy Y as a dominant strategy for both players. Neither prediction is supported by our data. In the last 5 rounds subjects chose in 64 out of 90 times the payoff dominant equilibrium, only 22 times they coordinated on the risk dominant equilibrium. It is even more informative to consider the answers to the questionnaire. According to Matsushima's theory players should conceive the game as

	X	Y
X	a, a	$c, 2$
Y	$2, c$	$3, 3$

with $a < 2$ and $c < 3$. Yet, 23 of our subjects considered a to be the maximum payoff of their opponent and 28 subjects reported that $a > 2$. Furthermore, 27 subjects reported their belief that $c > 3$.

For the prisoners' dilemma in Block 3 Matsushima's theory performs better. The theory predicts that the Nash equilibrium of the objective game (Y,Y) is the unique Nash equilibrium of the subjective game and that subjects believe the resulting payoff to be Pareto efficient. The first prediction agrees with our data as YY was chosen in 72 out of 90 times in the last 5 rounds. The second prediction was moderately successful as 18 of 35 subjects with consistent answers believed that the payoff of YY was Pareto efficient.

To summarize, our data yield only moderate support for Matsushima's theory. However, the general idea of subjective equilibria seems to be strongly supported by our data.

4 Conclusion

The purpose of this experiment was to test how players perceive a game which they are only incompletely informed about. In particular, we wanted to know whether players are able to guess the payoff structure of their opponents after repeated encounters. We see this experiment as a first step toward an investigation on "learning about games" rather than "learning in games". It should be complementary to another interesting recent experimental approach which aims to observe subjects' reasoning process based on their lookup pattern in game trees or payoff matrices (see Camerer et al., 1993 and Costa–Gomes et al., 2001).

Our results indicate that subjects were not very good in guessing the payoff structure of their opponents in 2×2 games, although the results varied across games. Nevertheless, play was often quite close to a Nash equilibrium. Two low–information learning theories, reinforcement learning (Erev and Roth, 1998) and payoff assessment learning (Sarin and Vahid, 2001) are able to explain the experimental data reasonably well, except for the coordination game data.

A unique feature of our questionnaire is that it allowed us to reconstruct subjective games as perceived by subjects. We found that most subjects followed an equilibrium strategy of "their" subjective game. Furthermore, in some instances a subjective Nash equilibrium could be constructed which provided an equilibrium selection device in case several Nash equilibria existed in the underlying game.

Future work may extend our experiment in several dimensions. Most interesting in our view would be to consider other aspects of a game about which players are incompletely informed, like their own payoff, the strategy sets, or the number of players. Conceivable is also to run the current experiment with larger games. However, given that subjects had trouble guessing the payoff structure of 2×2 games, we are not too optimistic that they would succeed in larger games in any reasonable time span.

A Translation of instructions

Welcome to our experiment! Please take your time to read the entire instructions carefully! During the next one and a half hours you can make some money by making various decisions at a computer. Please do not speak with other participants during the experiment. If you have any questions regarding the procedure, please refer quietly to the experimenter.

1. Timing

The experiment has 3 blocks with 20 rounds each. In each block you are randomly matched by the computer with another participant. After each block you are matched with a new partner, who is different from any of your previous ones.

Each block consists of 20 rounds. In every round, both participants simultaneously select among two possible actions: action "X" or action "Y".

During a block the payoff rules that determine the payoffs are constant. At the beginning of a new block the payoff rules are changed. After round 15 in each block you will be asked to answer a short questionnaire. You will be rewarded for correct answers.

After answering the questionnaire the computer starts rounds 16 to 20 during which all your payoffs are counted with quadrupled value. The new block starts after round 20.

[Figure 1 about here]

2. Payoffs in a given round

You and the other participant have two actions available. Hence, there are 4 possible combinations of actions (XX, YX, XY, YY). Your payoff in each round is depicted in a payoff table. The possible payoffs can take the values 0, 2, 3 or 4 and are denominated in "Taler".

The table below is an example for a payoff table (in this example the entries are changed):

Example of your payoff table

		others's action	
		X	Y
your action	X	0	2
	Y	3	4

For instance, if you choose action X and the other participant chooses Y, then your payoff is 2. The payoffs of the other participants are determined by a similar table, which can, however, differ from yours. The other participant only knows his payoff table, and you know only yours. The computer will always show your currently valid payoff table.

In every round you choose with your mouse or keyboard your action "X" or "Y" and confirm. After the other participant chooses his action, the computer informs you about your payoff and the next round is started.

3. The questionnaire

After round 15 in each block the computer asks you 4 questions about the payoff table of the other participant which is unknown to you. The other participant's payoff table can be written as follows:

Payoff table of other subject

		others's action	
		X	Y
your action	X	a	b
	Y	c	d

You should find out about a, b, c and d. Those payoffs are 0, 2, 3 or 4, once each. The following questions are posed to you:

a) Is a greater than b? yes or no?

b) Is c greater than d? yes or no?

c) Where is the highest payoff? a b c d

d) Where is the lowest payoff? a b c d

You should answer those questions carefully because for each correct answer you receive 30 Taler.

4. The last 5 rounds of a block

After answering the questions, the computer starts the rounds 16 to 20 of the same block. The payoff tables remain the same as in the first 15 rounds. However, your payoff will be quadruplicated by the computer from round 16 to 20. After round 20, the block ends. You are matched with a new participant and the payoff tables are changed.

5. Total payoffs

Your total payoff in Taler consists of the payoffs in round 1 through 15 of all 3 blocks, plus the quadruple payoffs in round 16 through 20 in all 3 blocks, plus 30 Taler for each correctly answered question. During the experiment the computer informs you about the payoffs you earned in all rounds. However, you are told the payoffs for correct answers only at the end of the experiment. At the end of the experiment your total payoff is paid to you under consideration of an exchange rate of 1 Taler = 0.06 DM. Thank you for your participation.

References

[1] Abbink, K. and Sadrieh, A. (1995). "RatImage: Research Assistance Toolbox for Computer–Aided Human Behaviour Experiments", SFB 303 Discussion Paper B-325, University of Bonn.

[2] Camerer, C., Johnson, E., Rymon, T., and Sen, S. (1993). "Cognition and Framing in Sequential Bargaining for Gains and Losses," in Binmore, K., Kirman, A., and Tani, P. (eds.). Frontiers of Game Theory, Cambridge, M.A.: The MIT Press.

[3] Costa–Gomes, M., Crawford, V., and Broseta, B. (2001). "Cognition and Behavior in Normal–Form Games: An Experimental Study", Econometrica 69, 1193-1235.

[4] Erev, I. and Roth, A. (1998). "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria", American Economic Review **88**, 848-881.

[5] Fudenberg, D. and Levine, D. (1993). "Self–Confirming Equilibrium", Econometrica 61, 523–545.

[6] Fudenberg, D. and Levine, D. (1998). The Theory of Learning in Games, Cambridge, M.A.: The MIT Press.

[7] Huck, S. and Sarin, R. (1999). "Players with Limited Memory", mimeo, University College London.

[8] Kalai, E. and Lehrer, E. (1995). "Subjective Games and Equilibria", Games and Economic Behavior **8**, 123-163.

[9] Matsushima, H. (1998a). "Learning about Stochastic Payoff Structure", mimeo, University of Tokyo.

[10] Matsushima, H. (1998b). "Towards a Theory of Subjective Games", mimeo, University of Tokyo.

[11] Osborne, M. J. and Rubinstein, A. (1998). "Games with Procedurally Rational Players, American Economic Review **88**, 834-847.

[12] Sarin, R. and Vahid, F. (2001). "Predicting How People Play Games: A Simple Dynamic Model of Choice", Games and Economic Behavior **34**, 104-122.

[13] Sethi, R. (2000). "Stability of Equilibria with Procedurally Rational Players", Games and Economic Behavior **32**, 85-104.

[14] Van Huyck, J., R. Battalio, and Rankin, F. (1996). "Selection Dynamics and Adaptive Behavior without Much Information", mimeo Texas A&M University.

Table 1: Percentage of correctly answered questions

Question	Block 1	Block 2	Block 3
Q1	77.8 %	58.3 %	66.7 %
Q2	75.0 %	50.0 %	63.9 %
Q3	63.8 %	27.8 %	30.6 %
Q4	27.8 %	44.4 %	33.3 %
Q1 & Q2	61.1 %	36.1 %	47.2 %
all four questions correct	19.4 %	16.7 %	13.9 %
inconsistent answers	3.5 %	0.7 %	0.7 %

Table 2: Pairs of subjects that correctly answered questions

Question	Block 1	Block 2	Block 3
Q1 & Q2	38.9 %	5.6 %	22.2 %
all four questions correct	11.1 %	0.0 %	5.6 %

Table 3: Empirical and theoretical frequencies, Block 2

	X	Y	$\sum$
X	29 (24)	37 (48)	66 (72)
Y	14 (6)	10 (12)	24 (18)
$\sum$	43 (30)	47 (60)	

Note: The theoretical frequencies are shown in parenthesis

Table 4: Questionnaire answers vs. play in final rounds

	Block 1 at least 3 Nash eq.		Block 3 at least 3 Nash eq.	
	yes	no	yes	no
Q1 & Q2 correct	7	0	3	1
Q1 & Q2 incorrect	10	1	12	2

Table 5: Average number of strategy switches

	vertical & horizontal		diagonal	
	Q1 & Q2 correct	incorrect	correct	incorrect
Block 1	4.86	4.64	0.57	1.27
Block 2	4.08	2.48	0.77	0.30
Block 3	3.35	1.84	0.53	0.21

Table 6: MSD scores ($\times 100$) for the RE and the SV model

model	opt. parameter	Block 1	Block 2	Block 3	mean MDS
RE	$s(1) = 16$	7.37	1.26	3.41	4.01
RE*	$s(1) = 13$	7.35	1.18	3.40	3.98
SV	$\lambda = .041$	6.76	1.57	4.09	4.14
SV*	$\lambda = .054$	7.29	1.53	3.69	4.17
equilibrium		16.50	6.96	8.00	10.49
random		5.42	1.90	7.78	5.03

Note: The parameter in the RE models is the initial propensity of each strategy ($s(1)/2$). In the SV models, λ is the relative weight with which current payoffs enter the assessments. Initial assessments were chosen from a uniform distribution on $[0, 4]$. Optimal parameters were found via a grid search to minimize the mean MSD on all blocks.

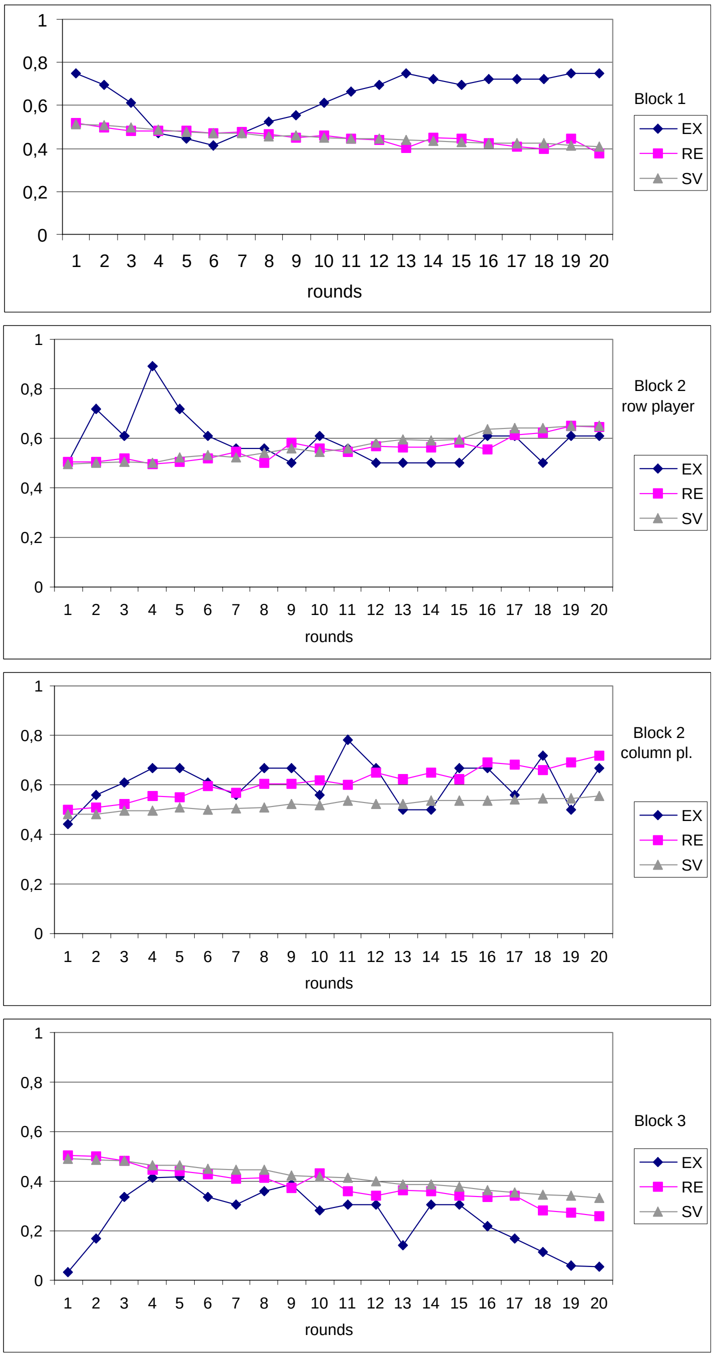

Figure 1: Proportion of X choices in the data (EX), and for 200 simulations of reinforcement learning (RE), and payoff assessment learning (SV).

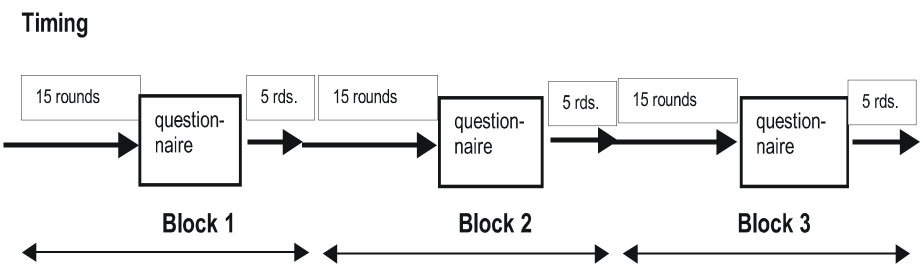
Timing
15 rounds
question-
naire
5 rds.
15 rounds
question-
naire
5 rds.
15 rounds
question-
naire
5 rds.
Block 1
Block 2
Block 3

Figure 2: