Institute for Empirical Research in Economics
University of Zurich

Working Paper Series
ISSN 1424-0459

Working Paper No. 234

Learning in and about Games

Anke Gerber

September 2006

Learning in and about Games*

Anke Gerber†

September 2006

Abstract

We study learning in finitely repeated 2×2 normal form games, when players have incomplete information about their opponents' payoffs. In a laboratory experiment we investigate whether players (a) learn the game they are playing, (b) learn to predict the behavior of their opponent, and (c) learn to play according to a Nash equilibrium of the repeated game. Our results show that the success in learning the opponent's type depends on the characteristics of the true game. The learning success is much higher for games with pure strategy Nash equilibria than for games with a unique mixed strategy Nash equilibrium, and it is higher for games with symmetric pure strategy Nash equilibria than for games with asymmetric equilibria. Moreover, subjects learn to predict the opponents' behavior very well. However, they rarely play according to a Nash equilibrium and we observe no correlation between equilibrium play and learning about the game.

Keywords: Learning, game theory, incomplete information, experiments.
JEL-Classification: C72, C92, D83.

*I would like to thank Bodo Vogt for his collaboration at an early stage of this project. I am also grateful to Dietmar Fehr, Ernst Fehr, Jörg Oechssler, to the participants of the ESF workshop in Zurich and, in particular, to Philipp Wichardt for valuable comments. Finally, I would like to thank Omar Pennacchio and Sascha Robert-Charrue for the programming of the experiment and Tan Schelling for his support in sorting the experimental data.

†Swiss Banking Institute, University of Zurich, Plattenstr. 32, 8032 Zurich, Switzerland, e-mail: agerber@isb.unizh.ch, tel.: ++41-(0)44-6343708, fax: ++41-(0)44-6344970.

1 Introduction

Many insights in economics are based on an equilibrium analysis. In particular, the applied literature, e.g. in industrial organization, usually derives its results from the assumption that strategic agents behave according to a Nash equilibrium. A crucial question is whether "the house is build on solid grounds," i.e. whether equilibrium behavior can be justified without appealing to agents' introspection and without the critical assumption of common knowledge of rationality.

In an attempt to rationalize equilibrium behavior, a large theoretical literature is concerned with the question whether equilibrium play can be learned by rational or even by less than fully rational agents when they interact repeatedly (for a survey see Fudenberg and Levine, 1999). The learning problem is complicated, if agents not only have to learn to play according to a Nash equilibrium *in* a given game, but they also have to learn *about* the game they are playing, which is the case whenever they have incomplete information about their opponents' payoff functions. Since hardly any game in the real world is played under conditions of complete information, this joint learning problem seems to be the most relevant one. At first sight the theoretical results are encouraging. Under a common prior assumption Jordan (1995) shows that Bayesian learning induces expectations, which converge to the set of Nash equilibria of the repeated game defined by the true types. However, convergence of expectations does not imply convergence of the actual behavior strategies to equilibrium. Convergence of strategies occurs under the "grain of truth" assumption (see Kalai and Lehrer, 1993): If players' beliefs are compatible with the true behavior strategies,[1] then the players' strategies and beliefs are asymptotically approximately Nash with probability one.

However, the theory is silent about how players actually arrive at such compatible beliefs. In the sequel several authors have argued that the grain of truth assumption may be too severe. Nachbar (1997) shows that if players' beliefs have a support satisfying a diversity condition (neutrality), then the players' best re-

[1] Technically this requires that the probability distribution on play paths, induced by the true behavior strategies, is absolutely continuous with respect to each player's belief distribution.

sponses are not in the support of their beliefs. Miller and Sanchirico (1997) argue that generically priors do not satisfy absolute continuity, if they are drawn randomly. Finally, Foster and Young (2001) and Young (2002) show that there is a fundamental tension between rationality and prediction: There are games (with a unique mixed strategy equilibrium) which cannot be learned by rational players, i.e. players almost never learn to predict their opponents' behavior and they almost never come close to playing a Nash equilibrium. The reason is, that there do not exist learning rules that are robust against changes in the players' payoffs. The result is very strong, since the learning rules can be very general and arbitrarily sophisticated. For example, this includes rules, where a player tries to estimate the opponent's payoff function by observing her reactions to different sequences of actions. The only assumption is that players are rational, in the sense that in each period they choose an optimal action given their belief about the probability of future play paths.

In this paper, we study the question of learning in and about games from an empirical point of view. Our objective is to get a detailed picture of how agents perceive a particular strategic situation under incomplete information and how they behave in repeated interaction. Hence, we do not ask *how* agents learn and we do not test the predictions of a particular learning theory. Rather, we ask whether there *is* learning, be it in or about the game. We present the results from a laboratory experiment where subjects played a 2×2 normal form game 50 times with the same partner without knowing the opponent's payoff matrix. Each player could be one of two possible types corresponding to two different payoff matrices. Players only knew their own type and the possible types of their opponents. Since a common prior assumption is hardly ever satisfied in a real world strategic situation, we did not induce any prior about the opponent's type in our experiment. Hence, the subjects in our experiment were facing a strategic situation like the one studied in Kalai and Lehrer (1993), Foster and Young (2001) and Young (2002), where beliefs are entirely endogenous. We had two treatments, one in which no player type had dominant strategies, and one in which both types had dominant strategies. Our experiment is designed such as to give answers to the following questions:

1. Do subjects learn the opponent's type, i.e. do they learn the game they are

playing?

2. Do subjects learn to predict the behavior of their opponent?

3. Do subjects learn to play according to a Nash equilibrium of the repeated game defined by the true types?

In addition we are interested in the following question:

4. Is there any correlation between learning in and about the game? In particular, is learning the opponent's type a necessary condition for learning to play according to a Nash equilibrium?

Concerning question 3, previous experimental research has shown that subjects often do not play according to a Nash equilibrium even under complete information. A prominent example is the finitely repeated prisoner's dilemma, for which a substantial degree of cooperation is observed in all but the last periods of play. Hence, a violation of Nash equilibrium play does not necessarily have any implications for the answers to questions 1 and 2 above. Therefore, in order to find out whether there is learning about the game, it is not sufficient to observe the agents' behavior. Consequently, in each period of play we also elicit the subjects' beliefs concerning the opponent's type and behavior. For their predictions about the opponent's behavior we use the belief elicitation method of Nyarko and Schotter (2002).

Our experimental results show that the success in learning the opponent's type depends on the characteristics of the true game. In general, subjects are successful in learning the game, whenever the true stage game has pure strategy Nash equilibria, and they do not learn the game, whenever there is only a mixed strategy equilibrium. In the latter case subjects on average correctly guessed their opponent's type in 48% of the periods at the end of play, while for games with pure strategy Nash equilibria the percentage of correct guesses is 78%. We also observe heterogeneity in the learning success across games with pure strategy Nash equilibria. The learning success is considerably high for games with symmetric pure strategy Nash equilibria, where subjects on average correctly guessed their opponent's type in 84% of the periods at the end of play. In contrast, if the game has asymmetric equilibria, the average percentage of correct guesses is 73% only.

Overall we find that subjects learn to predict the opponents' behavior very well. The only exception is a game for which we observe the play of mixed strategies, which seems to make the prediction more difficult. Despite the success in learning to predict the opponent's behavior, the subjects play according to a Nash equilibrium in 2 out of 5 different (true) games only. One is a pure coordination game, where subjects coordinate on the efficient equilibrium. The other is a version of the hawk-dove game, where they play according to the mixed strategy Nash equilibrium. For all other games we observe a tendency to coordinate on the efficient, non-equilibrium outcome. There are two explanations for this behavior, which apply to different games in our treatments: In one case subjects have learned the true game, and then show behavior which is consistent with experimental evidence under complete information. In the other case subjects are (partly) ignorant of the true game as in the case of the game which has a mixed strategy equilibrium only, and they show behavior which is consistent with experimental evidence under complete information in the game they perceive to play. Given these results it should come as no surprise that stated beliefs are not a good predictor for actual play, i.e. in general subjects do not play a best-reply to their stated belief. Moreover, we find that stated beliefs are not a better predictor for actual play than are Cournot or fictitious play beliefs. This is in contrast to the findings of Nyarko and Schotter (2002), who study a game with a unique equilibrium in mixed strategies that is played under complete information. Hence, their result does not seem to be robust against changes in the subjects' information or changes in the best-reply structure of the game.

Finally, we do not observe any correlation between equilibrium play and learning about the game. The frequency of equilibrium play is neither consistently higher nor consistently lower if players have learned the true game than if they have not learned about the game. If we combine this result with our finding that overall subjects are very successful in learning in and about the game, we see that equilibrium behavior is a bad proxy for learning success.

The papers that are most closely related to our work are Cox, Shachat, and Walker (2001) and Oechssler and Schipper (2003). Cox et al. conducted an experiment in order to test Jordan's (1991) model of Bayesian learning in games of incomplete information. Different from our experiment, they induced a common prior among subjects concerning the distribution of types. Cox et al. find that behavior is consistent with Jordan's Bayesian learning model whenever the true game has a unique equilibrium in pure strategies, and it is inconsistent, whenever the true game has two pure strategy equilibria. Since Cox et al.'s objective was to test the behavioral predictions of Jordan's Bayesian learning model, they did not control for subjects' beliefs concerning the opponent's type and behavior. Hence, their results do not allow any conclusions about the reasons for or the failure of Nash equilibrium play. Our design, in contrast, allows us to compare beliefs with behavior in any single period of play. In particular, we can analyze whether there is any correlation between learning (the true game or the opponents' behavior) and equilibrium play.

Oechssler and Schipper (2003) also conducted an experiment on the repeated play of 2×2 normal form games, but different from Cox, Shachat, and Walker (2001) as well as from our experiment, subjects did not have any information about their opponents' payoff matrix. Subjects played the game repeatedly with the same partner for 20 rounds. After 15 rounds of play subjects where asked about their beliefs concerning the opponent's payoff matrix. Oechssler and Schipper find that subjects are not very successful in learning the true game but that they often play according to an equilibrium in the subjective game they perceive to play. In our experiment we give subjects more time to learn in and about the game (the same game is repeated 50 times). Also, we have control over the subject's beliefs concerning the opponent's type and behavior during the whole course of play, so that we can study learning over time, as well as relate beliefs to behavior in any given period.

The remainder of the paper is organized as follows. In Section 2 we present the design of our experiment. Section 3 reports and discusses the experimental results. Finally, in Section 4 we conclude. Appendix A contains the instructions for the experiment.

2 The Experiment

In this section we give a detailed presentation of the experimental design and procedure. As pointed out in the introduction we are interested in the following research questions:

1. Do players learn the game they are playing?

2. Do they learn to predict their opponents' behavior?

3. Do they learn to play according to a Nash equilibrium of the repeated game?

4. Is there any correlation between learning about the game and equilibrium play?

2.1 Experimental Design

There are two treatments. In each treatment two players are randomly matched to play a 2×2 normal form game over 50 periods. Each player can be one of two possible types, X or Y, which is common knowledge among the players. This gives four possible games per treatment: XX, XY, YX and YY, where XY and YX are identical except for an exchange of the roles as column or row player for type X and Y. Each player only knows her own type and has no information about the opponent's type. In each period players simultaneously choose one of two possible actions, A or B. Players' payoffs are given in Table 1 for Treatment 1 and in Table 2 for Treatment 2. The two treatments cover all best-reply structures in which either none (treatment 1) or both players (treatment 2) have a dominant strategy. Payoffs are given in ECU (Experimental Currency Units). In our experiment 100 ECU were converted to CHF 5 (approximately €3.2). The payoffs were chosen such as to have best-reply structures that are easy to understand and equilibria that are not difficult to compute.

Treatment 1: Game XX is a version of the stag-hunt game. It has two pure strategy Nash equilibria (A, A) and (B, B), where (A, A) is the payoff dominant and (B, B) is the risk dominant equilibrium. This game also has one mixed strategy Nash equilibrium, where each player chooses A with probability $2/3$.

			Column Player				
			Type X			Type Y	
Row Player	Type X		A	B		A	B
		A	4 4	3 0	A	3 4	4 0
		B	0 3	2 2	B	2 3	0 2
	Type Y		A	B		A	B
		A	4 3	3 2	A	3 3	4 2
		B	0 4	2 0	B	2 4	0 0

Table 1: Payoffs in Treatment 1.

Games XY and YX are asymmetric and have a unique mixed strategy Nash equilibrium, where each player chooses A with probability 2/3. Finally, game YY is a version of the hawk-dove game, having two (asymmetric) pure strategy Nash equilibria (A, B) and (B, A) and one mixed strategy Nash equilibrium, where again each player chooses A with probability 2/3.

Treatment 2: In this treatment both players have dominant strategies so that there exists a unique pure strategy Nash equilibrium for all games. For XX the equilibrium is (A, A), for XY it is (A, B), for YX it is (B, A), and finally, for the prisoner's dilemma YY it is (B, B).

Theoretical Prediction

The objective of this paper is to analyze whether agents learn the game they are playing, whether they learn to predict their opponents' behavior and whether they learn to play according to a Nash equilibrium of the repeated game defined by the true types. Hence, our theoretical benchmark is the subgame perfect Nash equilibrium (SPNE) of the finitely repeated game defined by the true types.[2]

[2]Here, we assume that the payoff of the repeated game is the time-average of the payoffs

		Column Player					
		Type X			Type Y		
			A	B		A	B
Row Player	Type X	A	3 3	2 1	A	2 3	3 1
		B	1 2	0 0	B	0 2	1 0
			A	B		A	B
	Type Y	A	3 2	2 0	A	2 2	3 0
		B	1 3	0 1	B	0 3	1 1

Table 2: Payoffs in Treatment 2.

Clearly, any sequence of Nash equilibria of the stage game is a SPNE of the finitely repeated game. The converse, however, is only true if the stage game has a unique Nash equilibrium. In this case, any SPNE of the finitely repeated game consists of the repeated play of the unique Nash equilibrium of the stage game. This gives a unique prediction for games XY/YX in treatment 1 and all games in treatment 2.

If the stage game has multiple Nash equilibria, like games XX and YY in treatment 1, there may be SPNE, where the behavior strategies in some period are no Nash equilibrium of the stage game. Benoit and Krishna (1985) have shown that in the limit, when the number of repetitions goes to infinity, any individually rational and feasible payoff vector can be supported by a SPNE, if every player has at least two different Nash equilibrium payoffs in the stage game.[3] If we apply this limit folk theorem to game XX in treatment 1, we obtain

obtained in each period.

[3]A payoff vector $u \in \mathbb{R}^2$ is *feasible*, if it is in the convex hull of the set of payoffs that are attainable by some pure strategy profile of the stage game, and it is *individually rational*, if $u_i \geq v_i$ for $i = 1, 2$, where v_i is player i's minmax payoff.

that in the limit any payoff in

$$F^{XX} = \{u \in \mathbb{R}^2 \,|\, u_i \geq 2, i = 1, 2,\ 4(u_1 - 3) \leq u_2 \leq 3 + u_1/4\}$$

can be supported as a SPNE of the repeated game (see Figure 1).

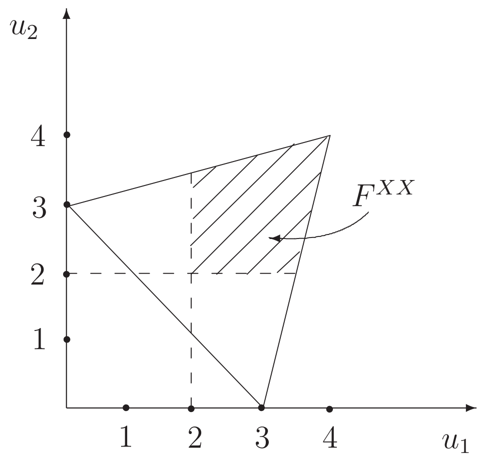

Figure 1: Feasible and individually rational payoffs for game XX in treatment 1.

Similarly, any payoff in

$$F^{YY} = \{u \in \mathbb{R}^2 \,|\, u_i \geq 2, i = 1, 2,\ u_1 + u_2 \leq 6\}$$

can be supported as a SPNE of the repeated game YY in treatment 1 (see Figure 2).

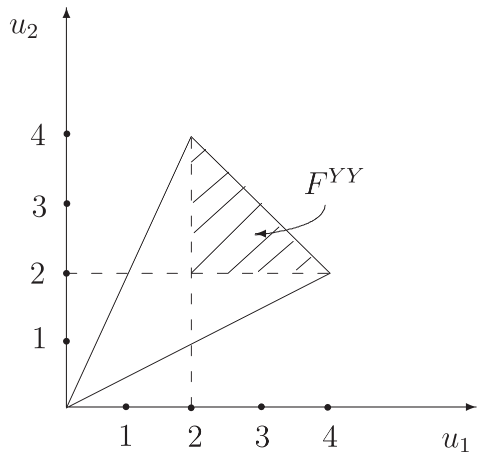

Figure 2: Feasible and individually rational payoffs for game YY in treatment 1.

2.2 Experimental Procedure and Subject Pool

The computerized experiment was conducted in the computer laboratories of the University of Zurich.[4] It lasted for approximately 2 hours including 20 minutes of instructions.[5] A total of 54 students from the University of Zurich and the Swiss Federal Institute of Technology Zurich (ETHZ) participated in the experiment. The subjects were studying several fields, including economics, sociology, law and computer science, and they had not participated in similar game theoretic experiments before. The average payoff of a participant was CHF 39 (approximately €25).

26 students played treatment 1, where we had three blocks with games YY, XX and XY/YX (in the given order), and 28 students played treatment 2, where we had two blocks with games XY and YY (in the given order). Subjects did not know in advance how many blocks they were going to play. In each block subjects were anonymously matched and anonymously assigned the role of the column or

[4]The experiment was programmed and conducted with the software z-Tree (Fischbacher, 1999).

[5]An English translation of the instructions can be found in Appendix A.

row player. Subjects then played the given game for 50 rounds with the same partner. The matching was such that in no block any subject was matched to a prior opponent or to someone who was an opponent of a prior opponent. Subjects were informed that they would get a new partner in each block and that their type and role as a column or row player could change from one block to another.

In order to study whether subjects learn the game they are playing and whether they learn to predict their opponents' behavior, we elicited the corresponding beliefs. To this end in each period subjects had to answer the following questions:

1. Which type is your partner?

If the the answer was correct, the subject received 1 ECU. Otherwise, she received 0 ECU.

2. Which action will be chosen by your partner in this round?
3. Indicate by a number p between 0 (not confident at all) and 100 (fully confident) how confident you are with your answer to question 2.

For their answers to questions 2 and 3 subjects were rewarded according to the following scoring function (cf. Nyarko and Schotter, 2002): If a subject reports confidence p the payoff is

$$1-\left(1-\frac{p}{100}\right)^2,$$

if the opponent chooses the predicted action, and it is

$$1-\left(\frac{p}{100}\right)^2,$$

if the opponent does not choose the predicted action. This scoring function is such that the payoff is maximal (equal to 1), if the opponent's action is correctly predicted with maximal confidence, and it is minimal (equal to 0), if the opponent's action is falsely predicted with maximal confidence. Moreover, the scoring function is such that reporting the true belief is optimal under risk neutrality. Clearly, under risk aversion subjects may find it optimal not to report their true beliefs. For example, always reporting $p = 50$ guarantees a riskless payoff of

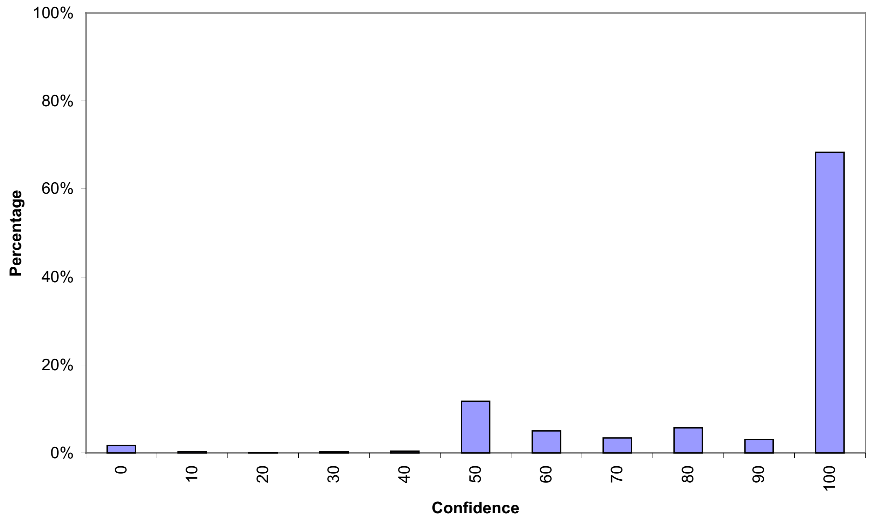

Figure 3: Empirical distribution of the degree of confidence over all blocks and all treatments.

0.75. However, in our experiment we find no evidence for such behavior. Only in 12% of all cases subjects report confidence $p = 50$, while in 68% of all cases they report maximal confidence $p = 100$ (see Figure 3, which gives the empirical distribution of the degree of confidence for our experiment).

The total payoff for a subject is the sum of the payoff from the play of the game and the payoffs for the belief about the opponent's type and the opponent's action in the next period. In each period we displayed on the subject's computer screen her payoff from the play of the game and her payoff for the belief concerning the opponent's strategy in the last period. During the whole experiment subjects did not receive any information about their payoffs for guessing the opponent's type.

One concern with our design could be that belief elicitation may change subjects' behavior. Belief elicitation could focus subjects' attention on these beliefs and away from actual play of the game. For this reason we made sure that the money which could potentially be earned by prediction was small compared to the earnings from the play of the game, but not too small, so that incentives were still sufficiently large. Belief elicitation could also induce agents to rationalize their behavior by stating particular beliefs, so that they think the chosen action is in accordance with the stated beliefs. However, the monetary incentives in our experiment should be sufficient to prevent such behavior. Also, agents could try to (partially) hedge their risk in the play of the game by stating a confidence level, such that the prediction payoff is high if and only if the payoff in the game is low. Finally, and most importantly, subjects could be more inclined to play a best-reply to their beliefs, if we elicit these beliefs since this forces subjects to think more deeply about the strategic situation at hand. Hence, if beliefs are approximately correct, one could expect to observe a higher proportion of equilibrium play with than without belief elicitation. These are serious concerns and any empirical study on beliefs and behavior has to take them into account. The literature on the influence of belief elicitation on behavior is controversial. In their study of a 2×2 game with a unique mixed strategy Nash equilibrium Nyarko and Schotter (2002) do not find any significant change in the subjects' behavior between treatments with and without belief elicitation. In contrast, Rutström and Wilcox (2006) report that belief elicitation changes the play path in a repeated asymmetric matching pennies game in a way such that structural models of belief learning have a better fit under belief elicitation. However, the difference between behavior with and without belief elicitation is only found for players who face strong payoff asymmetries between different actions. Subjects in our experiment do not face such strongly asymmetric payoff opportunities, so given the results in Nyarko and Schotter (2002) and Rutström and Wilcox (2006), we do not expect any critical influence of belief elicitation on behavior. Croson (2000) studies a public goods game and a prisoner's dilemma and finds a significant move towards usage of the dominant strategy under belief elicitation. As we will see later, our data does not give any indication for a high percentage of rational behavior in the case, where the true game is a prisoner's dilemma. Instead, in more than 90% of all cases (excluding the first periods of play) subjects in our experiment use a tit-for-tat strategy under belief elicitation, so the frequency of play of the dominant strategy could hardly be lower. Summarizing, we do not have any reason to suspect that belief elicitation had an influence on behavior in our experiment.

3 Results

In this section we present the results from our experiment. We start by analyzing the question whether subjects learn about their opponent's type.

3.1 Learning about the Game

Table 3 displays the percentage of periods, averaged over all pairs, in which subjects correctly guessed the opponent's type in periods 1-20, 21-40 and 41-50.[6] Looking at the last 10 periods we see that, with the exception of game XY/YX in treatment 1, subjects are very successful in learning about the true game they are playing. Recall that game XY/YX in treatment 1 is the only game that has a unique Nash equilibrium, which is mixed.

On average, subjects correctly guess their opponent's type in 78% of the last 10 periods whenever the true game has a pure strategy Nash equilibrium, while they only correctly guess the opponent's type in 48% of the last 10 periods, if the true game has a mixed strategy Nash equilibrium only. The learning success is highest for the prisoner's dilemma (game YY in treatment 2), where on average subjects correctly guess the opponent's type in 90% of the last 10 periods. If we compare games with symmetric and asymmetric pure strategy Nash equilibria we see that the learning success is slightly higher in the symmetric case: For games with symmetric equilibria the average of correct guesses is 84% (game XX in treatment 1 and game YY in treatment 2), while for games with asymmetric equilibria it is 73% (game YY in treatment 1 and game XY in treatment 2). We performed a Wilcoxon signed ranks test for individual subjects in a treatment to test for the difference in the learning success between the different game types. In treatment 1 the learning success is significantly higher for games with pure strategy Nash equilibria than for the game with a mixed strategy Nash equilibrium only (p-value 0.041 for game XX and 0.026 for game YY, both compared to game XY/YX). In treatment 2 the learning success is significantly higher for the game with a

[6] There is no significant difference between row and column players in those games, where both players are of the same type. Nevertheless, we do not average over player roles, since we consider them separately when testing for the significance in the increase of correct guesses between the beginning and the end of play.

Treatment	Game	Periods 1-20	21-40	41-50
1	XX, Row Player	73%	73%	74%
1	XX, Column Player	70%	77%	81%**
1	XY/YX, Type X	47%	44%	43%
1	XY/YX, Type Y	58%	50%	52%
1	YY, Row Player	57%	77%	82%***
1	YY, Column Player	44%	60%	63%
2	XY, Type X	65%	76%	77%**
2	XY, Type Y	62%	66%	71%
2	YY, Row Player	71%	81%	91%***
2	YY, Column Player	73%	91%	88%***

Table 3: Percentage of correct guesses of the opponent's type. ***, **, * indicates significance at the 1%, 5%, 10% level, respectively, of a Wilcoxon signed ranks test for individual subjects to compare periods 41-50 with periods 1-20.

symmetric pure strategy Nash equilibrium than for the game with an asymmetric Nash equilibrium (p-value 0.076). In treatment 1, however, this difference is not significant.

Since simultaneous correct guessing may be important for the play of a Nash equilibrium, Table 4 reports the percentage of periods, in which both subjects simultaneously correctly guessed the opponent's type. Here, the average of correct guesses is still high for games with symmetric pure strategy Nash equilibria (game XX in treatment 1 and game YY in treatment 2), namely 75% on average, while it is only 56% for games with asymmetric pure strategy Nash equilibria (game YY in treatment 1 and game XY in treatment 2). For the game with a unique Nash equilibrium that is mixed (game XY/YX in treatment 1), the average of

Treatment	Game	Periods 1-20	21-40	41-50
1	XX	57%	65%	70%**
1	XY/YX	27%	28%	28%
1	YY	25%	47%	54%**
2	XY	45%	54%	58%*
2	YY	52%	72%	79%***

Table 4: Percentage of correct guesses by both players. ***, **, * indicates significance at the 1%, 5%, 10% level, respectively, of a Wilcoxon signed ranks test for matched pairs to compare periods 41-50 with periods 1-20.

simultaneous correct guesses is 28% only. From these observations we may expect to see equilibrium play in the case of games with symmetric, pure strategy Nash equilibria, but neither in the case of games with asymmetric, pure strategy Nash equilibria nor in the case of the game with a mixed strategy Nash equilibrium only. However, as we will see later, there is no correlation between correct guessing and equilibrium play.

In all games with pure strategy Nash equilibria there is considerable learning over time: the percentage of periods in which subjects correctly guessed the opponent's type is increasing over time. As indicated in Table 3, the increase is significant (p-value below 0.05) or even highly significant (p-value below 0.01) for at least one type or player role (column or row player) in all games except game XY/YX in treatment 1.[7] A similar result obtains if we look at learning over time with respect to correct guesses by both players (Table 4). At the end of play the percentage of simultaneous correct guesses is significantly higher at the 1% level for game YY in treatment 2, at the 5% level for games XX and YY in treatment 1, and at the 10% level for game XY in treatment 2 (Wilcoxon signed ranks test for matched pairs).

[7]We performed a Wilcoxon signed ranks test for individual subjects.

We conclude that there is significant learning over time but that the final learning success depends on the characteristics of the true game. It is high for games with pure strategy Nash equilibria and it is low for the game having a mixed strategy Nash equilibrium only. Also, the learning success is higher for games with symmetric pure strategy Nash equilibria than for games with asymmetric equilibria.

3.2 Learning in the Game

In the following we analyze the subjects' learning behavior in the game. We start by comparing actual play with the theoretical benchmark given by a SPNE of the repeated game defined by the true types. We then provide a detailed analysis of play and compare our observations with experimental evidence under complete information, whenever such data is available for the games in our experiment.

Do they play according to a SPNE?

Table 5 displays the percentage of periods (averaged over all pairs), in which subjects played according to a pure strategy Nash equilibrium of the true stage game.[8] We find clear evidence for Nash equilibrium play only in case of the coordination game (game XX of treatment 1). Since both players have dominant strategies in treatment 2, a failure to play according to a Nash equilibrium of the stage game implies a clear violation of rationality. Concerning game YY in treatment 1, from the data in Table 5 we cannot judge whether subjects really fail to play according to a SPNE. Recall from Section 2.1 that there are SPNE which are not given by a sequence of pure strategy Nash equilibria of the stage game. Moreover, all games in treatment 1 also have an equilibrium in mixed strategies. We will see later that subjects indeed are playing according to the mixed strategy Nash equilibrium in case of game YY in treatment 1.

If we compare the beginning with the end of play we see that for all games there is an increase in equilibrium play of the true stage game. The question is whether this increase is significant. A Wilcoxon signed ranks test for matched pairs reveals that the increase is significant at the 10% level for game XX in

[8]Game type XY/YX is missing since this game only has a mixed strategy Nash equilibrium.

Treatment	Game	Periods 1-20	21-40	41-50
1	XX	91%	98%	95%*
1	YY	48%	40%	52%
2	XY	47%	48%	61%***
2	YY	34%	24%	46%

Table 5: Percentage of periods with play of a pure strategy Nash equilibrium of the true stage game. ***, **, * indicates significance at the 1%, 5%, 10% level, respectively, of a Wilcoxon signed ranks test for matched pairs to compare periods 41-50 with periods 1-20.

treatment 1 and at the 1% level for game XY in treatment 2. On the contrary, for game YY in both treatments we do not find a significant increase in equilibrium play over time.

Equilibrium play requires that both players know the true game. Hence, one explanation for the failure to play according to an equilibrium of the true game could be the fact that subjects have not learned the game they are playing. The corresponding hypothesis is, that the probability of equilibrium play is higher, if both players know the game than if at least one player does not know the game that is being played. To test this hypothesis, for each game we determine the empirical frequency of equilibrium play in the last 10 periods conditional on (1) both players correctly guessing the opponent's type and (2) at least one player not guessing the opponent's type. As we can see from Figure 4 we do not find much support for the hypothesis that guessing and equilibrium play are positively correlated.[9] Only for game XY in treatment 2 the probability of equilibrium play is considerably higher if both players correctly guessed the opponents type. For

[9]We cannot perform a statistical test for the disaggregated data, since there are not enough groups for which we observe both conditions, i.e. periods in which both subjects correctly guess the opponent's type and periods in which at least one subject does not correctly guess the opponent's type.

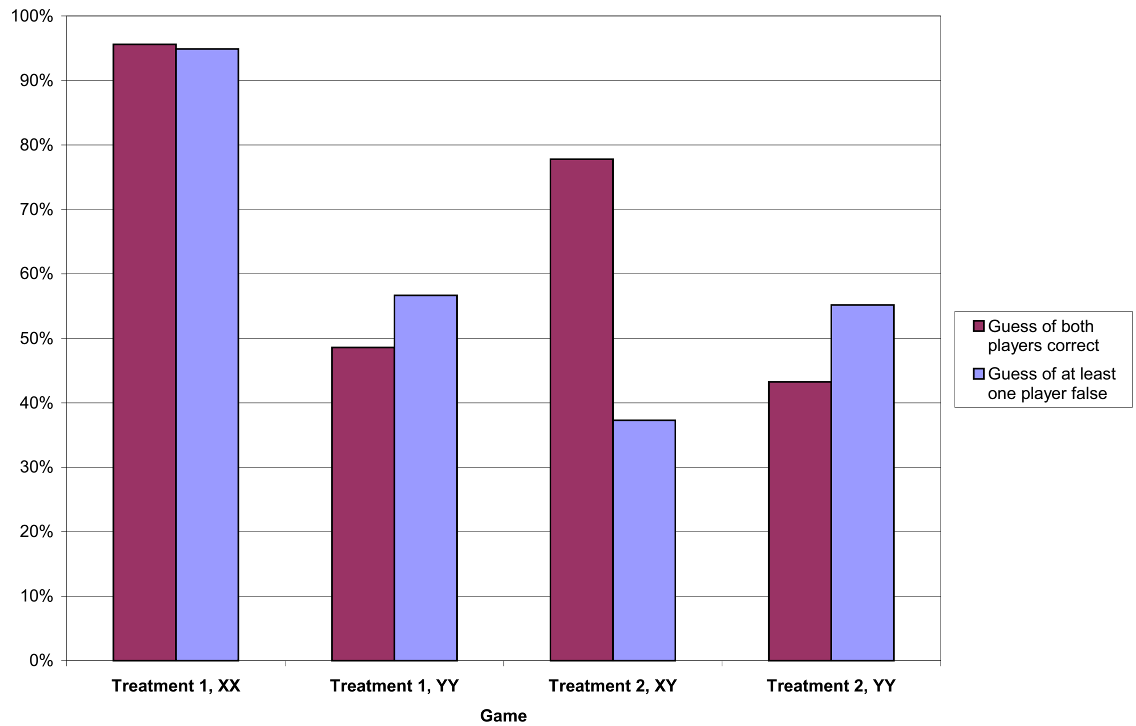

Figure 4: Conditional empirical frequency of equilibrium play in periods 41-50.

game XX in treatment 1 we do not find any correlation between guessing and equilibrium play, while for game YY in both treatments we even find an increase in equilibrium play if at least one player does not correctly guess her opponent's type.

If subjects do not play according to a Nash equilibrium of the true stage game what do they play instead? To answer this question we provide a detailed analysis for each game.

Treatment 1, Game XX:

This game is a symmetric coordination game with two pure strategy Nash equilibria AA and BB, where AA is the payoff dominant and BB is the risk dominant equilibrium. For this game our previous analysis has already shown that subjects learn their opponent's type and that they learn to play according to a pure strategy Nash equilibrium of the stage game. From Table 6 we see that subjects play the payoff dominant equilibrium in 91% of the last 10 periods, while they play the risk dominant equilibrium in only 5% of those periods. Only 1 out of 13 pairs played the risk dominant equilibrium in the majority of periods in the last 10 periods. This seems to be at odds with previous experimental evidence under complete information (see, for example, Cooper et al. 1992), which found much support for the risk dominant equilibrium. However, in these experiments subjects play a sequence of one-shot games with different opponents, so that they may be more concerned with the relative riskiness of the two strategies than with efficiency. In our design, where subjects play against the same partner within one block, they do not face this risk since they have information about the opponent's choices in the past. Hence, they manage to coordinate on the Pareto efficient outcome.

Periods 1-20

	A		B		Sum
A	*89%	(88%)	4%	(5%)	93%
B	5%	(7%)	*2%	(0%)	7%
Sum	95%		5%		100%

Periods 21-40

	A		B		Sum
A	*97%	(96%)	1%	(2%)	98%
B	2%	(2%)	*1%	(0%)	2%
Sum	98%		2%		100%

Periods 41-50

	A		B		Sum
A	*91%	(87%)	2%	(6%)	93%
B	2%	(6%)	*5%	(0%)	7%
Sum	93%		7%		100%

Table 6: **Treatment 1, Game XX**. Empirical distribution of play. Pure strategy Nash equilibria marked with *. Joint distribution determined by empirical marginal distributions in brackets.

Treatment 1, Game XY/YX:

This game is asymmetric and possesses a unique mixed strategy Nash equilibrium, where both players play A with probability 2/3. The unique SPNE of the true repeated game is given by repeated play of the mixed Nash equilibrium of the stage game. As we can see in Table 7, subjects do not play according to this mixed equilibrium but rather coordinate on the efficient outcome AA (73% of periods 41-50), which gives player type X payoff 4 and player type Y payoff 3. Observe that the payoff from the pure strategy Nash equilibrium is 8/3 only. Hence, subject type X plays rational in the true game as well as in the game she perceives to play in more than 50% of the periods (cf. Section 3.1), while Y does not play a best reply to the opponent's strategy. Y plays "smart," however, since his payoff is larger than the payoff he would obtain in the SPNE of the true game. It seems that the subjects' ignorance of the true game facilitates coordination on the efficient outcome:[10] If type Y wrongly believes the opponent is of type Y as well, and if he observes her opponent plays action A, she may conclude that the opponent is willing to sacrifice a payoff of 1 in favor of coordinating on the "fair" and efficient outcome AA in the hawk-dove game. Hence, type Y may reciprocate by playing A as well. Alternatively, since AA can be supported as a SPNE in YY in the limit, type Y's behavior is also consistent with equilibrium play whenever he perceives the true game to be YY.

It is interesting to note an end game effect, though. The empirical frequency of AA is highest in periods 21-40 (84%), while it is considerably smaller in periods 41-50 (73%). A Wilcoxon signed ranks test for matched pairs reveals that the decrease is significant at the 1% level (p-value 0.0039). This end game effect may be explained by the fact that type Y does not find it profitable to reciprocate to the believed opponent's cooperation, if the game gets close to the end.

If we compare these results with empirical evidence under complete information we find significant differences (for an overview see Camerer, 2003). Under complete information the empirical frequencies for games with a unique mixed strategy Nash equilibrium are, in general, not too far from equilibrium. This supports our conjecture that the incomplete information, in particular type Y's

[10]From Section 3.1 we know that subjects do not learn the type of their opponent: the average percentage of correct guesses in the last 10 periods is 48% only.

ignorance of her opponent's type, is responsible for the players' success in coordinating on the efficient outcome.

Periods 1-20

	A		B		Sum
A	68%	(64%)	13%	(17%)	81%
B	11%	(15%)	8%	(4%)	19%
Sum	79%		21%		100%

Periods 21-40

	A		B		Sum
A	84%	(82%)	7%	(8%)	90%
B	7%	(9%)	3%	(1%)	10%
Sum	91%		9%		100%

Periods 41-50

	A		B		Sum
A	73%	(69%)	11%	(15%)	84%
B	9%	(13%)	7%	(3%)	16%
Sum	82%		18%		100%

Table 7: **Treatment 1, Game XY/YX**. Empirical distribution of play (type X in row, type Y in column). Joint distribution determined by empirical marginal distributions in brackets.

Treatment 1, Game YY:

This is the hawk-dove game with two asymmetric pure strategy Nash equilibria AB and BA and one mixed strategy Nash equilibrium, where both players play A with probability 2/3. From Table 8 we see that the empirical frequencies of play get very close to the mixed strategy Nash equilibrium. At the end of the game (periods 41-50) the mixed strategy Nash equilibrium predictions cannot be rejected at the 10% level in a χ^2-test for 8 out of 13 groups. The aggregate frequencies in Table 8 in the middle of play (periods 21-40) are even closer to the mixed strategy Nash equilibrium. However, this is not confirmed if we look at disaggregated data, where the mixed strategy Nash equilibrium predictions cannot be rejected at the 10% level for 6 groups only.

Recall that in around 46% of the last 10 periods at least one subject in a pair does not guess her opponent's type, i.e. believes the true game is XY, respectively YX. In this game the unique Nash equilibrium is mixed. Therefore, we may explain the observed strategies of the players as equilibrium strategies in the true as well as in the subjective game they perceive to play.[11] Moreover, previous experiments (see Camerer, 2003) have shown that in case of a game with a unique Nash equilibrium that is mixed, the empirical frequencies are, in general, not too far from the equilibrium strategies. Hence, our data is consistent with the experimental evidence under complete information, if we take into account the fact that in 46% of the periods at least one subjects perceives to play game XY, respectively YX.

[11] Using the terminology of Kalai and Lehrer (1995), subjects are playing according to a *subjective equilibrium*.

Periods 1-20

	A		B		Sum
A	39%	(40%)	*24%	(23%)	63%
B	*24%	(24%)	13%	(14%)	37%
Sum	63%		37%		100%

Periods 21-40

	A		B		Sum
A	45%	(45%)	*23%	(23%)	68%
B	*20%	(21%)	11%	(11%)	32%
Sum	66%		34%		100%

Periods 41-50

	A		B		Sum
A	33%	(35%)	*29%	(27%)	62%
B	*23%	(21%)	15%	(17%)	38%
Sum	56%		44%		100%

Table 8: **Treatment 1, Game YY**. Empirical distribution of play. Pure strategy Nash equilibrium marked with *. Joint distribution determined by empirical marginal distributions in brackets.

Treatment 2, Game XY:

This game is asymmetric with AB being the unique Nash equilibrium in dominant strategies. As we can see in Table 9 player type X indeed plays her dominant strategy most of the time (93%) at the end of the game. Player type Y on the contrary plays his dominant strategy in approximately 2/3 of the cases only. Again we can explain this finding by the subjects' ignorance of their opponents' type. From our previous analysis recall that in 42% of the last 10 periods at least one subject in a pair does not correctly guess the opponent's type. If type Y wrongly believes his opponent is of type Y as well, so that they are playing the prisoner's dilemma, he may want to cooperate (i.e. play A), if his opponent is cooperating as well, which in fact she does in 93% of the cases. If, on the other hand, subject type Y correctly believes his opponent is of type X, then from his perspective there is no reason to cooperate. On the contrary, for type X there is no conflict between payoff maximization and efficiency. Playing the dominant strategy A leads to the efficient outcome independent of the opponent's type, whenever the opponent plays A as well. Hence, we expect type X to play A independent of her beliefs about the opponent's type, while we expect type Y to play A if he believes his opponent is of type Y as well and B if he believes the opponent is of type X.

This conjecture is confirmed by Table 10, which gives the conditional empirical frequencies for playing A conditional on believing the opponent is of type X, respectively Y. As we can see, type X always plays her dominant strategy A with a high probability (100%, if she believes the opponent is of type X as well and 91%, if she believes the opponent is of type Y). Type Y, however, plays A with a small probability (22%), if he believes his opponent is of type X, and with a higher probability (60%), if he believes his opponent is of type Y.[12]

We again observe an end game effect: type Y's degree of cooperation, i.e. the empirical frequency of play of action A is much higher in periods 21-40 (43%) than at the end of the game (33%). The decrease is significant at the 5% level (p-value 0.0273).

[12] We cannot determine the significance of the difference in conditional play of the action A, since we do not have enough pairs, for which we observe the choice of action A under both conditions, i.e. conditional on guessing the opponent is of type X as well as conditional on guessing the opponent is of type Y.

Periods 1-20

	A		B		Sum
A	32%	(27%)	*47%	(52%)	79%
B	3%	(7%)	18%	(14%)	21%
Sum	35%		65%		100%

Periods 21-40

	A		B		Sum
A	42%	(38%)	*48%	(51%)	89%
B	1%	(5%)	10%	(6%)	11%
Sum	43%		57%		100%

Periods 41-50

	A		B		Sum
A	32%	(31%)	*61%	(62%)	93%
B	1%	(2%)	6%	(5%)	7%
Sum	33%		67%		100%

Table 9: **Treatment 2, Game XY**. Empirical distribution of play (type X in row, Type Y in column). Pure strategy Nash equilibrium marked with *. Joint distribution determined by empirical marginal distributions in brackets.

		Type	
		X	**Y**
Guessing	**X**	100%	22%
	Y	91%	60%

Table 10: **Treatment 2, Game XY**. Conditional empirical frequency of action A, conditional on guessing the opponent is of type X, respectively Y (both types, periods 41-50).

Treatment 2, Game YY:

This game is a prisoner's dilemma with BB being the unique Nash equilibrium in dominant strategies and AA being the efficient outcome. As we can see in Table 11 in the middle of the game the Nash equilibrium is played only in 24% of the cases, while cooperation is achieved in 63% of the cases. There is a clear end-game effect here: the degree of cooperation is much lower in the end, where we observe the efficient outcome in only 37% of the cases while the Nash equilibrium is being played in 46% of the cases. The decrease is highly significant (p-value 0.0093). Recall from our previous analysis (Tables 3 and 4) that players know the true game, so the result is consistent with the extensive experimental evidence on the prisoner's dilemma under complete information (for early work see Lave, 1962, and Rapoport and Chammah, 1965).

It is important to note that the results for the prisoner's dilemma differ from those for all other games in our experiment with respect to the correlation between the subjects' choices of actions: As we can see in Table 11, the joint distribution determined from the observed empirical marginal frequencies of play is quite different from the actual empirical distribution, i.e. subjects choices are not independent from each other. A more detailed analysis of the strategies chosen by the subjects reveals that they are playing according to a "tit-for-tat" strategy most of the time (see Table 12), i.e. they play A (cooperate) if the opponent played A in the previous period and they play B otherwise. This explains the large degree

Periods 1-20

	A		B		Sum
A	36%	(26%)	15%	(25%)	51%
B	15%	(25%)	*34%	(24%)	49%
Sum	50%		50%		100%

Periods 21-40

	A		B		Sum
A	63%	(48%)	6%	(21%)	69%
B	7%	(22%)	*24%	(9%)	31%
Sum	70%		30%		100%

Periods 41-50

	A		B		Sum
A	37%	(21%)	11%	(28%)	49%
B	6%	(22%)	*46%	(29%)	51%
Sum	43%		57%		100%

Table 11: **Treatment 2, Game YY**. Empirical distribution of play. Pure strategy Nash equilibrium marked with *. Joint distribution determined by empirical marginal distributions in brackets.

	Row Player	**Column Player**
Periods 1-20	75%	71%
Periods 21-40	91%	92%
Periods 41-50	93%	92%

Table 12: **Treatment 2, Game YY**. Empirical frequency of the "tit-for-tat" strategy.

of correlation between the actions chosen by both subjects as well as the large degree of cooperation we observe.

Our previous analysis has provided answers to questions 1, 3 and 4, that we posed in the beginning (see p. 7). It remains to consider question 2:

Do they learn to predict the opponent's behavior?

We now analyze whether the subjects in our experiment learn to predict their opponent's behavior. To this end we determine the average payoff a subject obtained for predicting the opponent's behavior. Table 13 presents the corresponding payoffs averaged over all subjects of the same type. Recall that the maximum prediction payoff is 1 and that a player can guarantee himself a risk-free payoff of 0.75, if he predicts an arbitrary action for his opponent and states the confidence level $p = 50$. Hence, a payoff of 0.75 or below indicates that the subject has not learned anything about the opponent's behavior yet, while a payoff close to 1 indicates that the subject has learned to predict the opponent's behavior.

As we can see, for most games subjects start with a low prediction payoff close to 0.75 or even below (the only exception is game XX in treatment 1). The prediction payoff then increases in the middle of play and gets as high as 0.9 for most games except game YY in treatment 1. As indicated in the table, the increase in the prediction payoff between periods 1-20 and 21-40 is significant (p-

Treatment	Game	Periods		
		1-20	**21-40**	**41-50**
1	XX, Row Player	0.93	0.98**	0.94
1	XX, Column Player	0.89	0.98***	0.94
1	XY/YX, Type X	0.81	0.91**	0.89
1	XY/YX, Type Y	0.85	0.90	0.92
1	YY, Row Player	0.71	0.75	0.70
1	YY, Column Player	0.74	0.77	0.74
2	XY, Type X	0.76	0.86***	0.87
2	XY, Type Y	0.81	0.93***	0.91
2	YY, Row Player	0.76	0.91***	0.85
2	YY, Column Player	0.78	0.91***	0.89

Table 13: Average payoff from predicting the opponent's behavior. ***, **, * indicates significance at the 1%, 5%, 10% level, respectively, of a Wilcoxon signed ranks test for individual subjects to compare periods 21-40 with periods 1-20.

value below 0.05) or even highly significant (p-value below 0.01) for most games.[13] Hence, overall subjects learn to predict the opponent's behavior very well. In the end of the game the prediction payoff then decreases again. For game XY/YX in treatment 1 and both games in treatment 2 this can be explained by an end game effect: For these games we found that the behavior in the end of play differs from behavior in the middle of play. 10 periods of play in the end may not be enough to learn to predict the changed behavior of the opponent.

Concerning the low prediction payoff in game YY in treatment 1 recall that we observe play of the mixed strategy Nash equilibrium in this game and that a correct prediction of the mixed strategy (play A with probability 2/3) yields a

[13]We performed a Wilcoxon signed ranks test for individual subjects.

prediction payoff $7/9 \approx 0.78$. Hence, it could be that subjects correctly predict the opponent to play according to the mixed strategy equilibrium. However, the individual prediction data reveals that this is not the case. Rather, the majority of subjects reports either confidence $p = 100$, indicating that they predict the opponent to play a pure strategy, or $p = 50$ indicating that they are completely uncertain about the opponent's strategy.[14]

Do they play a best-reply to some belief?

Finally, we analyze whether subjects myopically best respond to some belief. In a recent experiment Nyarko and Schotter (2002) have shown that subjects' behavior can be explained much better by best response play to their stated beliefs than by best response play to some empirical beliefs, like Cournot or fictitious play beliefs. Here, the Cournot belief assigns probability 1 to the opponent's last action, while fictitious play beliefs are given by the empirical frequency with which the opponent has chosen each action in all previous periods. In our experiment, we can distinguish between a best-reply to Cournot, fictitious play or stated beliefs in treatment 1 only. All games in treatment 2 have dominant strategies for both players, so that the best-reply is independent of the belief concerning the opponent's play. Moreover, our previous analysis has already shown that subjects do not necessarily play their dominant strategy in treatment 2, so that actual play cannot be explained by myopic best-reply behavior.

Therefore, we only study treatment 1. For each game we determine the percentage of data that can be explained by the assumption that subjects play a myopic best-reply to Cournot, fictitious play or their stated beliefs. The latter beliefs are obtained from the action subjects predict for the opponent and from their stated confidence level. As we can see in Table 14, stated beliefs are not a better predictor for actual play than are Cournot or fictitious play beliefs.

Subjects' behavior in the coordination game XX can be explained very well by myopic best-reply behavior. This is consistent with our finding that subjects

[14] A subject who reports a confidence $p = 50$ may also predict the opponent to choose both actions with probability 0.5. However, given the data it seems more likely that the subject is completely uncertain about the opponent's strategy and hence chooses $p = 50$ in order to receive a riskfree prediction payoff of 0.75.

Treatment	Game	Cournot Beliefs	Fictitious Play Beliefs	Stated Beliefs
1	XX	96%	94%	93%
1	XY, Type X	87%	80%	85%
1	XY, Type Y	19%	24%	24%
1	YY	43%	51%	50%

Table 14: Percentage of data (all periods), that can be explained by Cournot, fictitious play or stated beliefs.

coordinate on the efficient equilibrium in this game. Also the behavior of type X in game XY/YX can be explained very well by myopic best-reply behavior. This is not the case for type Y in game XY/YX. Again this is consistent with our previous finding concerning the actions chosen by both types in this game. Finally, subjects' behavior in game YY cannot be explained by myopic best-reply behavior. Recall that the majority of pairs is playing according to the mixed strategy Nash equilibrium in this game but that subjects fail to predict that the opponent is playing according to the mixed strategy equilibrium.

Hence, our data does not confirm the finding of Nyarko and Schotter (2002), that stated beliefs are a much better predictor for actual play than are Cournot or fictitious play beliefs. Since Nyarko and Schotter study a particular 2×2 normal form game having a unique mixed strategy equilibrium and subjects in their experiment have complete information about the game that is being played, we conclude that their findings are not necessarily robust against changes in the subjects' information or against changes in the best-reply structure of the game.

4 Conclusion

In the real world there is hardly any strategic situation, where agents have complete information about their opponents' preferences. Nevertheless, many economic models assume complete information. Moreover, it is often assumed that agents have correct beliefs concerning their opponents' behavior and that they are rational, i.e. play a best-reply to their beliefs. Taken together this leads to the prediction of Nash equilibrium play in the "true" game. In this paper we have provided an empirical test of the question, whether the assumption of complete information and equilibrium play can be rationalized by learning in repeated interaction.

Our results show that there is significant learning about the game over time, but that the final learning success depends on the characteristics of the true stage game. The learning success is high, whenever the stage game has pure strategy Nash equilibria, and low, whenever the unique stage game equilibrium is in mixed strategies. This confirms theoretical results showing that learning is particularly difficult in games having mixed strategy Nash equilibria only (see Foster and Young, 2001, and Young, 2002). Among games with pure strategy Nash equilibria, those with symmetric equilibria are learned best. In particular, the prisoner's dilemma is learned very well, while the learning success is considerably lower for another game in our experiment, which has an asymmetric Nash equilibrium in dominant strategies. This may come as a surprise because one could expect that games with dominant strategies are generally learned very well, since play of the dominant strategy immediately reveals the opponent's type. However, our experiment shows that subjects do not necessarily play their dominant strategy, so that their actions do not reveal their type.

Moreover, our results show that subjects learn to predict their opponent's behavior very well. However, with the exception of the coordination game and the hawk-dove game subjects do not play according to a subgame perfect Nash equilibrium (SPNE) of the true repeated game. Hence, the reason for the failure to play according to a SPNE is not, that subjects have a wrong belief about the opponent's behavior, but that they behave irrationally: they often do not play a best-reply to their beliefs. In this way for several games they manage to coordinate on the efficient, non-equilibrium outcome and achieve a higher payoff than in a Nash equilibrium of the true game. In some games the cooperative behavior is facilitated by the subjects' ignorance of the true game: Having a false belief about the opponent's type, a subject may interpret the opponent's action as cooperative behavior and hence reciprocate by cooperation, while in fact the opponent is not cooperating but merely playing her dominant strategy.

Our experiment shows that equilibrium play is a bad proxy for learning success. Only from stated beliefs but not from observed behavior we can conclude whether agents have learned *about* the true game, and whether they have learned to predict their opponents' behavior, i.e. have learned *in* the game. Further empirical research is needed in order to understand why learning apparently is easier for games with symmetric than for games with asymmetric Nash equilibria, as well as why the learning success is so low for games with mixed strategy Nash equilibria.

A Instructions

Following is an English translation of the instructions for the experiment as they were given to the participants.

Instructions of the Game

Welcome to our experiment! You are going to make various decisions at a computer terminal. Your payoff depends on your performance. Hence, you should read the following instructions very carefully. In case you have questions, get in touch with the instructor. Observe that you are not allowed to talk to the other participants during the whole experiment.

Overview of the Game

The game consists of several blocks. In each block you play a game with another participant over 50 rounds. The game is described in the following.

Sequence of Actions within a Block

At the beginning of a block you are randomly matched to another participant with whom you play the following game for 50 rounds.

First you are told, whether you are a "row player" or a "column player." Your partner then is assigned the other role. In addition you and your partner are assigned one of two possible types, X or Y. The game you are playing is given by the combination of the types of the row and column player (see Table 1).[15] If both players are of type X, you play the game in the upper left corner. If the row player is of type X and the column player is of type Y, you play the game in the upper right corner. If the row player is of type Y and the column player is of type X, you play the game in the lower left corner. If both players are of type Y, you play the game in the lower right corner.

In each round you and your partner choose one of two possible actions, action A or action B. Your payoff depends on your own action and on the action chosen by your partner. Hence, there are four possible combinations of actions (AA, AB, BA, BB). In the corresponding fields you find the payoffs for the row player

[15] We display only the table for treatment 1.

Table 1

			Column Player				
			Type X			Type Y	
Row Player	Type X		A	B		A	B
		A	4 4	3 0	A	3 4	4 0
		B	0 3	2 2	B	2 3	0 2
	Type Y		A	B		A	B
		A	4 3	3 2	A	3 3	4 2
		B	0 4	2 0	B	2 4	0 0

as the first, lower and blue number, and the payoff for the column player as the second, higher and orange number.

Example:

If both players are of type X, you play the game in the upper left corner. If both players choose action A, the row player gets a payoff of 4 and the column player gets a payoff of 4. If the row player chooses A and the column player chooses B, the row player gets a payoff of 0 and the column player gets a payoff of 3. If the row player chooses B and the column player chooses A, the row player gets a payoff of 3 and the column player gets a payoff of 0. If both players choose action B, the row player gets a payoff of 2 and the column player gets a payoff of 2.

One difficulty is the fact that you are not informed about the type of your partner and your partner is not informed about your type. The game is determined but neither you nor your partner exactly know it. If you are the row player and of type X, the only thing you know is that one of the two upper games is being played, but you do not know which one. If you are the row player and of type Y, the only thing you know is that one of the two lower games is being played. If you are the column player and of type X, the only thing you know is that one of the two left games is being played. If you are the column player and of type Y, the only thing you know is that one of the two right games is being played.

In addition to your actions, in each round you have to answer questions, which refer to the game you are playing and to the action chosen by your partner in the given round. For correct answers you receive a positive payoff.

Payoffs

Payoffs for the chosen action

The payoffs are given in ECU (Experimental Currency Units): 1 ECU corresponds to 0.05 CHF.

Payoffs for answers to the questions

In addition to choosing an action, in each round you have to answer three questions.

1. Which type is your partner?

If your answer is correct, you receive 1 ECU. Otherwise, you receive 0 ECU.

2. Which action will be chosen by your partner in this round?

3. Indicate by a number p between 0 (not confident at all) and 100 (fully confident) how confident you are with your answer to question 2.

If your answer is p, and your partner chooses the action you predicted, your payoff is

$$1 - \left(1 - \frac{p}{100}\right)^2 \text{ ECU.}$$

If your partner chooses the other action, which you did not predict, your payoff is

$$1 - \left(\frac{p}{100}\right)^2 \text{ ECU.}$$

In the following figure you can see your payoff depending on your confidence p and the action chosen by your partner.

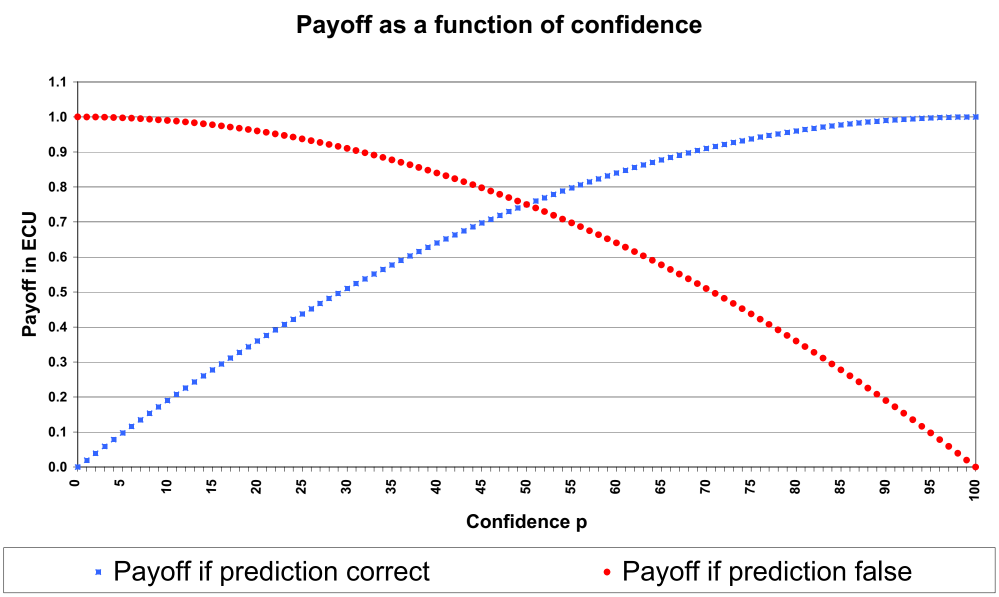

The payoffs from your chosen action and the answers to the questions are summed up over all rounds, converted to Swiss Francs, and are payed to you at the end of the experiment.

Your Information

In each round you are informed about your payoff from the play of the game in the previous round and about the payoff for your answer to question 3. In addition you are informed about the sum of your payoffs for the play of the game and your answers to question 3 in all previous round. At no time you are informed about whether your answers to the question concerning your partner's type were correct or false. Hence, you only see one part of your total payoff on your computer screen.

New Block

After 50 rounds you are randomly matched to a new partner and the game starts again. Thereby, your type may change and you may be assigned a new role as row or column player.

References

Benoit, J.-P., and V. Krishna (1985): "Finitely Repeated Games," *Econometrica*, 53, 905–922.

Camerer, C. F. (2003): *Behavioral Game Theory. Experiments in Strategic Interaction.* Princeton University Press, Princeton, New Jersey.

Cooper, R. W., D. V. DeJong, R. Forsythe, and T. W. Ross (1992): "Forward Induction in Coordination Games," *Economics Letters*, 40, 167–172.

Cox, J. C., J. Shachat, and M. Walker (2001): "An Experiment to Evaluate Bayesian Learning of Nash Equilibrium Play," *Games and Economic Behavior*, 34, 11–33.

Croson, R. (2000): "Thinking like a Game Theorist: Factors Affecting the Frequency of Equilibrium Play," *Journal of Economic Behavior & Organization*, 41, 299–314.

Fischbacher, U. (1999): "z-Tree. Toolbox for Readymade Economic Experiments," IEW Working Paper 21, University of Zurich.

Foster, J., and H. P. Young (2001): "On the Impossibility of Predicting the Behavior of Rational Agents," *Proceedings of the National Academy of Sciences of the USA*, 98, 12848–12853.

Fudenberg, D., and D. K. Levine (1999): *The Theory of Learning in Games.* MIT Press, Cambridge, Massachusetts.

Jordan, J. S. (1991): "Bayesian Learning in Normal Form Games," *Games and Economic Behavior*, 3, 60–81.

——— (1995): "Bayesian Learning in Repeated Games," *Games and Economic Behavior*, 9, 8–20.

Kalai, E., and E. Lehrer (1993): "Rational Learning Leads to Nash Equilibrium," *Econometrica*, 61, 1019–1045.

——— (1995): “Subjective Games and Equilibria,” *Games and Economic Behavior*, 8, 123–163.

Lave, L. B. (1962): “An Empirical Approach to the Prisoner’s Dilemma Game,” *Quarterly Journal of Economics*, 76, 424–436.

Miller, R. I., and C. W. Sanchirico (1997): “Almost Everybody Disagrees Almost All the Time: The Genericity of Weakly-Merging Nowhere,” Department of Economics Discussion Paper Series 9697/25, Columbia University.

Nachbar, J. H. (1997): “Prediction, Optimization, and Learning in Repeated Games,” *Econometrica*, 65, 275–309.

Nyarko, Y., and A. Schotter (2002): “An Experimental Study of Belief Learning Using Elicited Beliefs,” *Econometrica*, 70, 971–1005.

Oechssler, J., and B. Schipper (2003): “Can You Guess the Game You are Playing?,” *Games and Economic Behavior*, 43, 137–152.

Rapoport, A., and A. M. Chammah (1965): *Prisoner’s dilemma: A Study in Conflict and Cooperation.* University of Michigan Press, Ann Arbor, Michigan.

Rutström, E. E., and N. T. Wilcox (2006): “Stated Beliefs Versus Empirical Beliefs: A Methodological Inquiry and Empirical Test,” mimeo, Unviversity of Houston.

Young, P. (2002): “Bounded Rationality and Learning. On the Limits to Rational Learning,” *European Economic Review*, 46, 791–799.