

Keynote lecture

How to play (3×3) -games. A strategy method experiment

Reinhard Selten,^a Klaus Abbink,^b Joachim Buchta,^a and
Abdolkarim Sadrieh^{c,*}

^a *Laboratorium für experimentelle Wirtschaftsforschung, Universität Bonn, Adenauerallee 24-42,
53113 Bonn, Germany*

^b *School of Economics, The University of Nottingham, University Park, Nottingham, NG7 2RD, UK*

^c *Department of Economics and CentER, Tilburg University, PO Box 90153, 5000 LE Tilburg, The Netherlands*

Received 4 January 2001

Abstract

Using the strategy method (Selten, 1967, *Beiträge zur experimentellen Wirtschaftsforschung*, pp. 136–168) we elicit subjects' strategies for playing any 2-person (3×3) -game with integer payoffs between 0 and 99. In each of 5 tournaments, every strategy pair plays 500 000 games. The frequency of pure strategy equilibrium play increases from 51% in the first to 74% in the last tournament, with the equilibria that maximize joint payoff being preferred when multiple exist. For games without pure equilibria, strategies are typically based on elements of the *best-reply cascade*: MAP (maximize the expected payoff against uniformly randomizing opponents), BR-MAP (best reply to MAP), and BR-BR-MAP (best reply to BR-MAP).

© 2003 Elsevier Inc. All rights reserved.

JEL classification: C78; C91; C92; D82

Keywords: 2-person games; Equilibrium selection; Experimental economics

1. Introduction

Suppose you have to play a (3×3) -game against somebody intelligent and educated, but not necessarily an expert in game theory. How should you proceed to select one

* Corresponding author.

E-mail addresses: rselten@wiwi.uni-bonn.de (R. Selten), klaus.abbink@nottingham.ac.uk (K. Abbink), sadrieh@kub.nl (A. Sadrieh).

of your strategies? Standard non-cooperative game theory offers the notion of a Nash equilibrium (Nash, 1951) in pure or mixed strategies. This suggests that in a game with a uniquely determined equilibrium, one should choose the equilibrium strategy. However, a (3×3) -game may have many equilibria. A theory of equilibrium selection, e.g., the one by Harsanyi and Selten (1988), could be used to pick one of them, but nobody would expect that it has descriptive relevance for situations as the one described above.

In this study we are less interested in the task of playing a particular (3×3) -game, but rather in the much more demanding one of developing a general method for playing such games. For this purpose, we make use of the strategy method (Selten, 1967). This means that in our experiment subjects had to write programs determining the choice of a strategy for every (3×3) -game with randomly generated integer payoffs between 0 and 99, including the boundaries. The strategy method has been applied successfully in a variety of other contexts, e.g., repeated prisoner's dilemma games (Axelrod, 1984), oligopolies (Keser, 1991; Selten et al., 1997), two-person bargaining (Kuon, 1994); public goods games (Keser and Gardner, 1999).

After having gained experience with spontaneous play against anonymous opponents, our subjects were involved in five rounds of strategy programming, each of which was followed by a computer tournament with 500 000 randomly generated (3×3) -games. Extensive feedback about the tournament results led to considerable learning from one programming round to the next.

One striking result concerned games with pure strategy equilibria (roughly 80% of the 500 000). In these games, the percentage of pure strategy equilibrium plays markedly increased from 51% in the first tournament to 74% in the fifth one. At the end, the pure equilibrium play resulted in 98% of all games with only one pure equilibrium. In games with more than one pure equilibrium, a clear tendency towards the selection of the one with the maximum joint payoff was observed.

A different picture emerged for the way in which the strategy programs deal with games without pure strategy equilibria. The subjects do not consider mixed strategy equilibria in their programs. The strategy programs are based on much simpler strategic concepts combined in various ways. The simplest one is MAP, maximal average payoff, the strategy which maximizes the sum of the three payoffs obtainable against the possible choices of the other player. BR-MAP, the best reply to MAP, and BR-BR-MAP, the best reply to BR-MAP, are also important ingredients of the strategy programs. MAP, BR-MAP, and BR-BR-MAP form a hierarchy that we refer to as the *best-reply cascade*.

The best-reply cascade is similar to Nagel's *steps of reasoning* (1993, 1995) and to Stahl's *levels of smartness* (1993). In Stahl's theory, each player gives a best response to a mixed population of opponents using lower levels of smartness. Stahl and Wilson (1995) conducted experiments in which 48 subjects played 12 symmetric (3×3) -games once. Nine of the twelve games had a uniquely determined symmetric pure strategy equilibrium and three of them had no pure strategy equilibria. In seven of the nine games with a uniquely determined symmetric equilibrium the symmetric equilibrium strategy was the most frequent choice. They interpreted their results in terms of a logit choice model involving five types of players and three levels of smartness.

While the instances of the best-reply cascade we find in our subjects' strategy programs are supportive of the basic idea of Stahl and Wilson (1995), because they can be interpreted

as special extreme cases of the smartness levels, other aspects of our data are not in agreement with their theory. The model of Stahl and Wilson (1995) does not distinguish between games with and without pure strategy equilibria. However, almost all of our subjects made such a distinction in their final strategies. In the games with pure strategy equilibria, an overwhelmingly majority of subjects explicitly chose an equilibrium strategy. The strategy prescriptions for games without pure strategy equilibria make use of the best-reply cascade, but show a great diversity of concepts. An interpretation in terms of a Bayesian optimization involving probability distributions over opponents of lower smartness, often is not possible. Of course, when comparing our results to those of Stahl and Wilson (1995), one should keep in mind that our strategy method experiment created a more favorable environment for subjects' long-term learning process than the spontaneous play observed by Stahl and Wilson (1995).

2. The experiment

The subjects were 36 third and fourth year students of economics. The experiment took the form of a student seminar, lasting for a whole winter semester. (The instructions are available upon request.) The subjects were divided into two groups of 18 each without interaction between groups. The meetings of groups alternated. In each week only one of the groups met. The subjects first played (3×3) -games spontaneously against anonymous opponents, randomly matched with other group members, and only later entered the strategy programming rounds.

The games used in the experiment had integer payoffs between 0 and 99, including the boundaries. Each of the 18 payoffs was determined by independent, random draws with equal probabilities for all 100 possibilities. This random generation was applied both to the games in the spontaneous game playing sessions and to the game base of the tournaments. However, the same game base, consisting of 500 000 games, was used in each tournament.

In an introduction, the task and the random generation process were explained to the participants. Then the subjects were involved in spontaneous play up to the middle of the second session. The second part of this session was devoted to the explanation of the techniques of preparing the programs. The subjects had to draw flow charts describing their strategies and to supplement them by a verbal documentation explaining the intention and underlying reasoning. They were required to use a standard flow chart representation containing ordinary mathematical tools only and a standardized notation for references to payoffs and actions. (In order to avoid confusion, we use the word *action* for a strategy of a (3×3) -game and reserve the word *strategy* for a method applicable to all (3×3) -games with payoffs as described above.)

Each flow chart was translated to a computer program by one of the authors. Within each group each program was matched against each of the 17 other programs in each of both roles. In this way each program was involved in 17 million games. There were altogether 153 million plays in each tournament.

The strategy programs had no memory. This means that they could not store results of previous games played in the tournament and make choices dependent on them. The only data available to the program was the structure of the game to be played. Therefore,

supergame strategies could not be implemented. Consequently, the games played by the programs were strictly one-shot, even though so many of them were played. Since supergame effects are excluded, this setup is very favorable for the examination of non-cooperative theories for one-shot games. Nevertheless, learning from one tournament to the next is possible.

The subjects had the opportunity to use the laboratory terminals for the preparation of their strategies. A special *trainer* software was available that permitted observing choices made by their own strategy programs for arbitrary bimatrices. In this way, the subjects could check whether their programs did what they were intended to do. Subjects only received coaching concerning the technical aspects of their programs, not concerning the contents. Especially, they did not receive any training concerning game theoretic concepts.

It sometimes happened, mainly in the first programming round, and very rarely later, that a strategy left the choice for some games unspecified due to incomplete case distinctions. For such games, this participant received zero payoffs. Payoffs for specified choices were averages of the payoffs against all opponents with specified choices (it never happened that all opponents had unspecified choices). Payoffs for the whole tournament were sums of all these averages. The percentages of unspecified choices aggregated over both groups were 3.3% for the first simulation, 0.5% for the second, 0.08% for the third, 0.05% for the fourth, and 0% for the last tournament, in which no unspecified choices occurred.

After each tournament, except after the last one, the participants received feedback about their results in 1000 randomly selected games in one role ordered according to increasing success. For each of these games the participants could inspect the bimatrices and the distribution of choices of the other participants on the other side as well as their own choice. This made it relatively easy for the subjects to see where their strategies needed to be improved. Each participant could print out and take home up to 40 of these games with the information described above. The measurement of success in a particular role of a particular game was based upon the total payoff against the other 17 participants. The *maximal possible payoff* is the best payoff which could have been obtained. The measure of success was the *best-reply payoff ratio*, defined as the actual payoff divided by the maximal possible payoff. In the same way, a best-reply payoff ratio can also be formed for a strategy as a whole.

After a tournament the participants also received feedback about the average payoff and the average best-reply payoff ratio of every strategy within their own group. These were a strategy's per play averages over the 17 million plays it had been involved in. The information enabled participants to compare themselves with the other group members. The subjects were also informed on the number of their own unspecified choices.

The subjects were not motivated by monetary payoffs, but by grades. They were told that grades were based on absolute success in the last tournament, measured by the best-reply payoff ratio (50%), on the quality of strategy documentation (25%), and active involvement (25%). Emphasis was put on the importance of absolute success, rather than relative success. We told the subjects that it was not the ranking, but the absolute best-reply payoff ratio, no matter what distribution of grades would result. The motivation by final grades seemed to be very effective. Judging by the elaborateness of the flow charts provided by the subjects, they must have spent an enormous amount of effort on the task.

3. Development of payoffs and best-reply payoff ratios

Figure 1 shows the development of average payoffs and average best-reply payoff ratios over the five tournaments separately for both groups. These are the per play averages over all strategies and all plays of the games in each tournament (18 times 17 million). The scales of the two variables, on the left- and the right-hand side of the figure, are unrelated.

The development over the tournament must be seen in the light of some benchmarks for average payoffs and average best-reply payoff ratios as listed in Table 1. Let us first look at a population of randomizers choosing each action with equal probability. In this population, the average payoff is 49.5 and the average best-reply ratio is 66.5% for our game base. We can see that already in the first tournament the average performance is much better.

Approximate upper bounds can be derived by simulation of a hypothetical population of joint-payoff maximizers. In this population, every player chooses the first action with a joint payoff maximum field. This results in joint payoff maximization in all cases in which only one field exists with a maximal payoff sum. In other, relatively rare cases, a lack of coordination may lead to a lower payoff. In this population, the average payoff is 79.3 and the average best-reply ratio is 95.0%. Of course this upper bound presupposes a cooperation which would be inadequate in the experimental situation. The players are motivated by their own payoffs and not by the joint payoffs and, therefore, it is not surprising that the average payoffs obtained in the tournaments remain considerably below the approximate upper bound even in the last tournament. For the same reason, the best-reply ratios reach higher levels in the last tournaments than in the population of joint payoff maximizers, even if they start lower.

Another benchmark is supplied by a population of MAP players (MAP is the abbreviation for “maximal average payoff”). A MAP player always chooses an action that maximizes the sum of the three payoffs obtainable against the three possible choices of the opponent. If there are several such actions, each of them is taken with the same probability. The average payoff for a population of MAP-players is 63.7 and the average best-reply

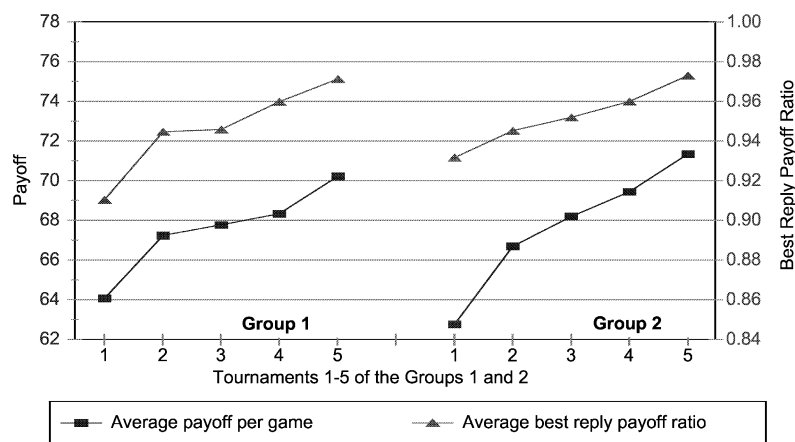


Fig. 1. Development of average payoffs and average best-reply payoff ratios.

Table 1
Benchmark and tournament results

	Average payoff	Average best-reply payoff ratio
Randomizers	49.5	66.5%
Joint-payoff maximizers	79.3	95.0%
MAP players	63.7	85.5%
First tournament	63.4	92.1%
Second tournament	67.0	94.5%
Third tournament	68.0	94.9%
Fourth tournament	68.9	96.0%
Fifth tournament	70.8	97.2%

ratio is 85.5%. It can be seen that these average payoffs are near to the ones of groups 1 and 2 in the first tournament. However, the best-reply ratios in the last tournament are much higher than that of the MAP-population.

In the light of the benchmarks, it can be seen that there are considerable improvements of average payoffs and best-reply payoff ratios over the five tournaments in both groups. All four sequences depicted in Fig. 1 are monotonously increasing. Due to the experimental setup, learning is only possible from one tournament to the next. Learning seems to take place to an astonishing degree within the sequence of only five tournaments. It is worthwhile to note that this learning from tournament to tournament is of a different nature than other kinds of learning that are often suggested in the context of experimental games, like adaptive learning. Within the complex task of specifying a general strategy for *all* games under consideration, learning cannot be just a consequence of simple reinforcement, but is rather the result of a thorough analysis of the strategic environment of the previous tournament.

4. A remark on the statistical evaluation

Strictly speaking, there are only two independent observations in our experiment. In view of the interaction within a subject group, each of the two subject groups is only one observation. Nevertheless, some statistical tests can be applied in order to reject the null hypothesis of independence amongst successive tournament results. Thus, as shown in Fig. 1, the result that average payoffs monotonously increase from one tournament to the next can be tested against the null hypothesis that the distributions of average payoffs are identical and independent in each tournament. This can be done by counting the number of inversions within the sequences of average payoffs for both groups, i.e., the number of times in which a higher value precedes a lower one.

The number of inversions in both groups together can serve as a test statistic. Table 2 shows the cumulative probabilities for the lower tail of the distribution of the sum of the number of inversions in both groups under the null hypothesis. Since this distribution is symmetric, the table informs us both about one-sided and two-sided levels of significance. We can see that up to three inversions one obtains a two-sided level of significance of less than 5%.

Table 2
Cumulative probabilities for the lower tail of the distribution of the sum of inversions

Sum of inversion numbers for both groups	0	1	2	3	4	5	6
Cumulative probabilities	0.00007	0.00063	0.00299	0.01007	0.02681	0.05972	0.11535

In Fig. 1, all sequences are monotonously increasing. The null hypothesis can be rejected for the development of the average payoffs and the best reply payoff ratios on the significance level 0.01%. This strongly supports the idea that there is learning from one round to the next.

5. Structural properties of the game base

The task of choosing a strategy is more difficult in some games than in others. Therefore, it is useful to classify the games in the game base according to strategic criteria. A very important characteristic is the number of pure strategy equilibria. Another aspect of game structure is the presence or absence of strictly dominated actions on at least one side. In this respect, three cases are of special interest. The first case is that in which at least one of the players has a strictly dominating strategy, i.e., an action that strictly dominates both other actions. Such games are very easy to analyze. A broader class is that of dominance solvable games. In these games a unique “solution” can be found by iterative elimination of strictly dominated strategies. However, not every game with a dominating strategy of one player is dominance solvable, since it may happen that the other player has several best replies to it. Finally, a still broader class contains all games with at least one strictly dominated strategy on at least one side. The dominance solvable games also belong to this class.

Figure 2 shows the percentages of games with none, one, two, three, or four pure strategy equilibria in the game base and also the percentages of games in each of the dominance related classes that were just described. Since the probability of drawing games with four or more pure strategy equilibria is extremely small, our randomly drawn game base only contained eleven games with four and no games with more than four pure strategy equilibria. We do not report on the games with four pure strategy equilibria in the rest of this paper, because the results for these games are practically identical to the results for the games with three pure strategy equilibria and because the sample of eleven games seems too small to allow an independent examination.

Only 20.4% of the games have no pure strategy equilibrium. 57.3% of the games have exactly one pure strategy equilibrium. Among these are the 42% that are dominance solvable. Thus, in 73.3% of all games with only one pure strategy equilibrium, the equilibrium can be found by an iterative elimination of strictly dominated strategies. It is interesting to see that 76.5% of the games can be reduced by the elimination of at least one strictly dominated strategy.

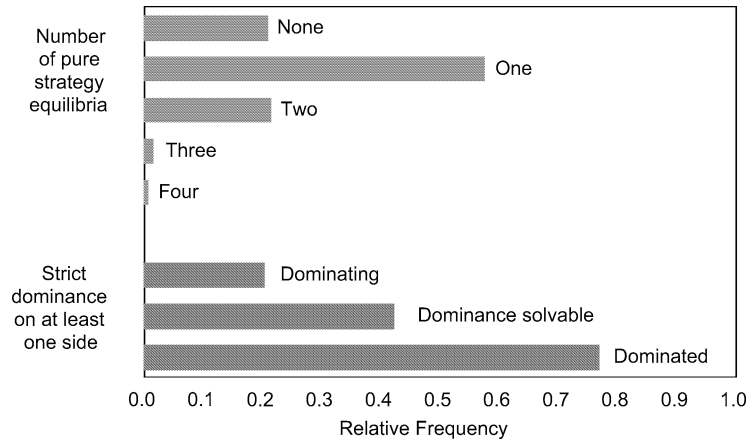


Fig. 2. Structural properties of the game base.

6. Last tournament performance in different types of games

Figure 3 shows last tournament average payoffs of both groups compared with the maximal cooperative payoffs. The term maximal cooperative payoffs refers to the simulation results for joint-payoff maximizers, who choose the first action with a maximum joint-payoff field (see Section 8). It can be seen that in the case of only one pure strategy equilibrium the last tournament results come surprisingly close to the cooperative payoffs benchmark. In the case of no pure strategy equilibrium, the difference to the benchmark is much greater.

It is interesting that the benchmark payoff monotonously increases with the number of pure strategy equilibria. The observed group payoffs are lowest for games without pure

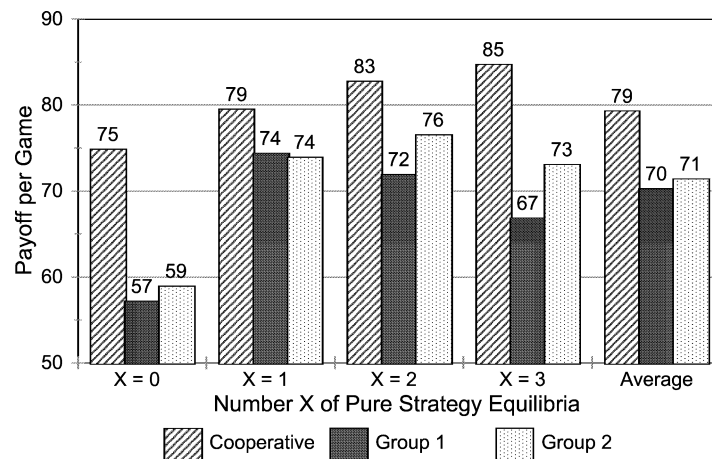


Fig. 3. Maximal cooperative payoffs and average last tournament payoffs.

strategy equilibria. There seems to be little difference between games with one, two, or three equilibria. However, the difference of observed payoffs to the cooperative payoffs increases for both groups as the number of equilibria moves from one to three. Presumably this is due to coordination difficulties caused by the presence of more than one pure strategy equilibrium.

7. Learning to play pure strategy equilibrium

Figure 4 shows the relative frequency of pure strategy equilibrium plays of both groups separately for the games with one, two, or three equilibria. It can be seen that in games with one pure strategy equilibrium, the frequency of equilibrium plays monotonously increases from relatively low levels (63.2% for group 1 and 74.5% for group 2) in the first tournament to surprisingly high levels (98.9% for group 1 and 96.5% for group 2) in the last tournament. This result is significant at the 0.01% level, two-sided (see Table 2).

A similar tendency is observed in games with two pure strategy equilibria, even if it starts and ends on markedly lower levels. Here, too, we see monotonous sequences for both groups connected with a significance level of less than 0.01%, two-sided (see Table 2). In the case of games with three pure equilibria, the sum of inversion numbers is 1. This indicates a level of significance of less than 0.5%, two-sided.

The results shown in Fig. 4 strongly support the idea that the subjects learn to coordinate more and more on pure strategy equilibrium plays. The coordination can be seen most strongly in games with only one pure strategy equilibrium, but to some extent also in games with two and even with three pure strategy equilibria.

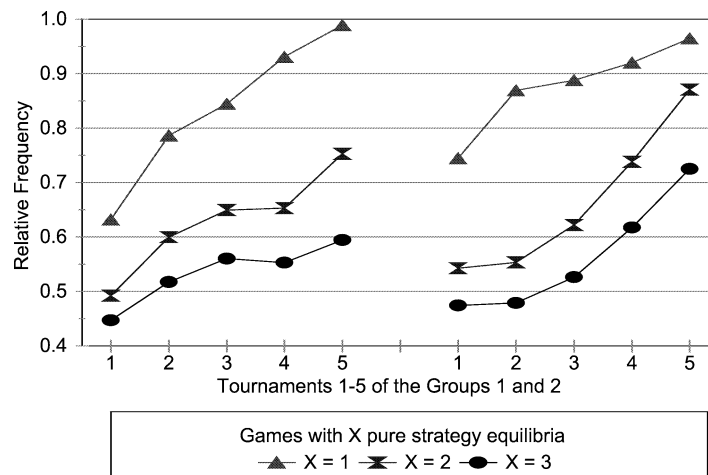


Fig. 4. Relative frequency of pure strategy equilibrium plays.

8. Coordination at the joint-maximum equilibrium

As we have seen in the preceding section, the percentage of equilibrium plays in games with several pure strategy equilibria tends to increase from tournament to tournament. It seems to be the case that the subjects learn to coordinate on the joint-maximum equilibrium, i.e. the pure strategy equilibrium with the highest sum of payoffs. Of course, there are some games that have several joint-maximum pure strategy equilibria. With the exception of these rare cases, the joint-payoff maximum selection criterium, emerges as a simple coordination device. The criterion is often explicitly programmed in the last tournament strategies.

Figure 5 shows the relative frequencies of joint-maximum equilibrium plays as a fraction of all pure strategy equilibrium plays in games with two and three pure strategy equilibria. It can be seen that group 2 shows a high degree of coordination at the joint-maximum equilibrium in the last two tournaments. In group 1, coordination strongly occurs only in the last tournament, but less markedly than in the last two tournaments of group 2.

There is no evidence for a gradually increasing development towards the joint-maximum criterium for equilibrium selection. The sum of inversion numbers is 5 for games with two pure strategy equilibria and 8 for games with three pure strategy equilibria. As we can see in Table 2, these results are not significant on the 10% level, two-sided. It seems to be the case that the acceptance of the joint-maximum equilibrium selection criterium is a development of later tournaments only. Casual inspection of the program codes supports this impression. However, answering the question whether a particular program makes use of this selection criterium or not is not as simple as one might think. Particularly in earlier tournaments the programs are often written in a roundabout way that makes it difficult to see what the guiding ideas are. Therefore, we restricted the statistical evaluation of the structure of the programs to the final tournament.

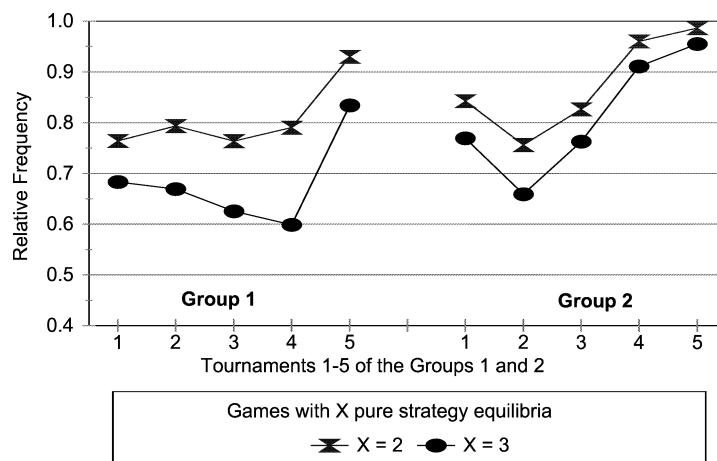


Fig. 5. Relative frequency of joint-maximum equilibrium plays as a fraction of all pure strategy equilibrium plays.

9. Choices of weakly dominated actions

From the point of view of normative game theory, one should not expect the use of weakly dominated actions. At least, one should expect a decrease in the frequency of weakly dominated action choices from tournament to tournament. Figure 6 shows the relative frequencies of weakly dominated action choices within all cases in which such choices were possible for both groups in different classes of game situations. Of special interest is the case in which the only equilibrium action is weakly dominated. In this case, the frequency of weakly dominated action choices tends to increase, rather than to decrease from tournament to tournament, even if in the case of group 2 the increase is not monotonous. The sum of inversion numbers is 1, which is connected to a significance level of 0.5%, two-sided (see Table 2).

Even in games in which a weakly dominated action was only one of several equilibrium actions, a non-negligible frequency of such weakly dominated choices was observed. In this case, the sum of the inversion numbers is 2, which is significant on the level of 1%, two-sided. It can also be seen that weakly dominated non-equilibrium actions are very rarely chosen. In both groups, this frequency is well below 1% in the last tournament.

The increasing tendency to use weakly dominated equilibrium strategies is due to the fact that the subjects learned to play the equilibrium action more and more in games with only one pure strategy equilibrium. The concept of pure strategy equilibrium was important in advanced tournament strategies, but dominance played almost no role.

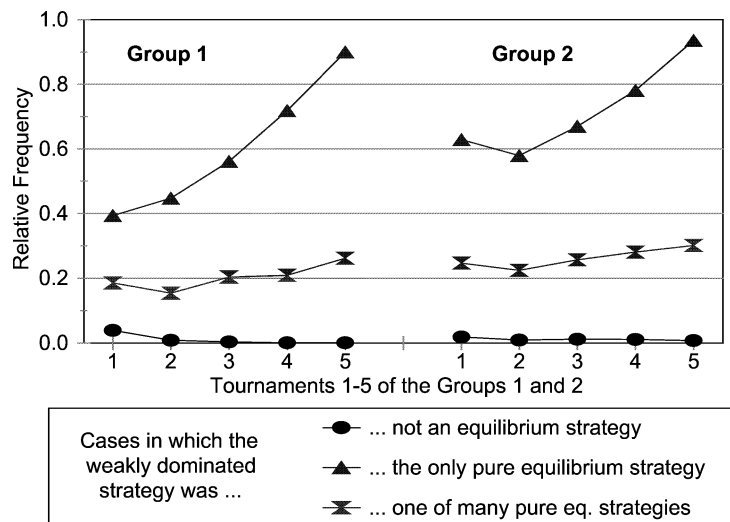


Fig. 6. Choices of weakly dominated actions as a fraction of possible cases.

10. How to play (3 × 3)-games with pure equilibria

In the case of games with pure equilibria, our experimental results indicate a clear answer to the question of how to play (3 × 3)-games: *Choose the action connected to the joint-maximum equilibrium*. Of course, this rule cannot be applied in the rare exceptions without a unique joint-maximum equilibrium, but otherwise, it provides a surprisingly simple coordination device. Both groups, which did not interact, independently developed a tendency towards this rule. Admittedly, a sample of two is not conclusive, but seems nevertheless suggestive. It would be desirable to replicate our experiment to obtain more reliable conclusions. Nevertheless, in view of the great effort required for a strategy study, it seems to us, that it is worthwhile to report our results and to draw some tentative conclusions, even if no statistical significance can be claimed. We hope that our example encourages similar work by others.

The tendency towards the joint-maximum equilibrium may be partly due to the fact that, in the tournaments, each participant played each game in both roles. Thereby, the situation is symmetrized. Any two participants, who choose joint-maximum equilibrium actions wherever possible, receive the same aggregate payoffs against each other over all such games. It is quite possible that something else emerges in a strategy study that avoids this kind of symmetry.

11. Development of average best-reply payoff ratios

Figure 7 shows the development of average best-reply payoff ratios over the five tournaments, separately for both groups and for games with zero, one, two, or three pure strategy equilibria. These are per play averages over all subjects and all plays of the corresponding games in each group. The tendency towards the use of equilibrium strategies results in very high best-reply payoff ratios in the case of only one pure equilibrium. Due to the coordination at the joint-maximum equilibrium, even in the cases of two or three pure equilibria, quite high best-reply payoff ratios are achieved in the last tournaments. It is remarkable that even in the games with no pure equilibria, a quite high best-reply payoff ratio of more than 0.9 is reached by both groups in the last tournament.

In all cases, we can observe an increasing tendency of the average best-reply payoff ratio. The highest number of inversions is three, which occurs for the games with three pure equilibria and corresponds to a cumulative probability of 0.01007 (see Table 2). In the cases of two pure equilibria and no equilibria there are only two inversions. For the games with only one pure strategy equilibrium we find no inversions at all. Obviously, the null-hypothesis that the increase of the best-reply payoff ratios is a random effect can be rejected at significance levels below 2% in all cases. This is a clear indication of learning from one tournament to the next. The global result on the development of the payoff ratios is largely confirmed for the classes of games with different numbers of pure equilibria separately.




Fig. 7. Development of average best-reply payoff ratios.

12. The problem of playing (3 × 3)-games without pure equilibria

As we have seen, strategic choice in games with pure strategy equilibria becomes more and more dominated by the equilibrium concept and the joint payoff maximum selection criterion. In the case of games without pure strategy equilibria, we do not observe a tendency towards a similarly clear answer to the question how to play such games. Nevertheless, some principles emerge that seem to determine the strategy construction of successful participants.

In the introduction, we already mentioned the concepts MAP, BR-MAP, and BR-BR-MAP. The meaning of these abbreviations is explained once again at the bottom of Table 3. MAP can be viewed as the best reply to a zero-intelligence opponent, who chooses each of his three actions with the probability 1/3. BR-MAP is the best reply to MAP and BR-BR-MAP is the best reply to BR-MAP. Obviously, the three strategies form a hierarchy of successive best replies. In the following we refer to this hierarchy as the *best reply cascade*.

The *best reply cascade* is similar to Nagel's *steps of reasoning* (1993, 1995) with the only difference that in her case a higher step was not exactly the best reply to the next lower one, since subjects tended to ignore the influence of their own actions on the optimal choice. The best reply cascade is also similar to Stahl's *levels of smartness* (1993), with the distinction that in Stahl's theory, players use best replies to a probability distribution over opponents of all lower smartness levels. The strategies of the best reply cascade can formally be interpreted as extreme special cases of strategies used in Stahl's theory.

Table 3 describes the essential features of the three most successful strategies in each of the two groups. In this description the treatment of minor details is omitted in order to present a clearer picture. For example, it is not described how a participant choosing BR-BR-MAP treats cases with several actions with the properties BR-MAP or BR-BR-MAP.

It can be seen that with the exception of one strategy (group 2, rank 3) all strategies presented in Table 3 make use of the concepts of the best reply cascade. In three cases (group 1, ranks 2–4) a member of the best reply cascade, namely BR-BR-MAP, was

Table 3
Essential features of the most successful last tournament strategies

Group 1, Rank 1 Essential feature	Payoff 60.9 Assign the probabilities 12/17, 2.5/17, and 2.5/17 to opponent's BR-MAP and to each of the other action. Choose best reply to this mixed strategy.	Best reply payoff ratio 0.97
Group 1, Ranks 2–4 Essential feature	Payoff 59.1 Choose BR-BR-MAP.	Best reply payoff ratio 0.94
Group 2, Rank 1 Essential feature	Payoff 63.0 Form matrix of both players' MAP and BR-MAP. Choose row with maximal average payoff in this matrix.	Best reply payoff ratio 0.98
Group 2, Rank 2 Essential feature	Payoff 62.9 Form matrix of opponent's MAP and BR-MAP and own BR-MAP and BR-BR-MAP. Choose row with maximal average payoff in this matrix.	Best reply payoff ratio 0.97
Group 2, Rank 3 Essential feature	Payoff 61.7 Assign probabilities 0.6, 0.3, 0.1 to opponent's actions with highest, middle, and lowest average payoffs for the opponent. Choose best reply to this mixed strategy.	Best reply payoff ratio 0.96
Abbreviation	Explanation	
MAP	Action with the maximal average payoff (best reply to zero-intelligence randomizers)	
BR-MAP	Best reply to opponent's MAP	
BR-BR-MAP	Best reply to opponent's BR-MAP	

specified as the last tournament substrategy for (3×3) -games without pure strategy equilibria.

Two strategies (group 1, rank 1 and group 2, rank 3) optimize against distributions over the opponent's actions. In the first case, the weight 12/17 is assigned to BR-MAP and the remaining 5/17 are split evenly between the two other actions. In the second case, the weighting of the opponent's actions is not guided by the best reply cascade, but by the order of the opponent's actions with respect to his own average payoffs. In this order the actions receive probabilities 0.6, 0.3, and 0.1.

Two strategies (group 2, ranks 1 and 2) first make use of best reply cascade concept in order to reduce the (3×3) -game to a (2×2) -game and then choose the action with the maximum own row payoff sum in the reduced game. In the first case, the reduced matrix is formed by MAP and BR-MAP of both players. In the second case, the reduced matrix is formed by the opponent's MAP and BR-MAP and the own BR-MAP and BR-BR-MAP. In a way, the MAP-principle is applied again on a higher level to the reduced matrix.

It is interesting that the most successful last tournament strategies are quite simple. Less successful strategies are sometimes much more complicated. Six of the seven most successful strategies make use of the best reply cascade concepts, MAP, BR-MAP, and BR-BR-MAP, in some way.

13. Frequently used strategy components in games without pure strategy equilibria

Table 4 presents an overview over the strategy components used by more than one final tournament strategy. The components are listed in the order of their frequencies. Obviously, the best reply cascade concepts appear quite often. The list also contains some components

Table 4
Frequently used strategy components in games without pure strategy equilibria

Components used by more than one final tournament strategy	Frequency ^a	
	Group 1	Group 2
BR-MAP	6	9
MAP	3	6
opponent's mixed action	3	2
BR-BR-MAP	3	1
BR-maximin	3	1
opponent's MAP (excluding BR-MAP)	2	2
maximin	2	2
(iterative) elimination of dominated actions	3	1
opponent's BR-MAP (excluding BR-BR-MAP)	2	1
maximum joint payoff	0	3
maximum sum of joint payoffs in a row	0	2

MAP: Action with the maximal average payoff (best reply to zero-intelligence player).

BR-MAP: Best reply to opponent's MAP.

BR-BR-MAP: Best reply to opponent's BR-MAP.

^a Strategies may involve more than one component.

not used by the most successful strategies. Amongst these are maximin and BR-maximin (the best reply to the opponent's maximin). The frequency of these components is comparable to that of the less frequent best reply cascade concepts. The same is true for the elimination of dominated strategies (iterative or not). The two least frequent components in the list are the maximization of the joint payoff and the maximization of the sum of joint payoffs in a row.

Even in the last tournament 3 of the 36 strategies still made some use of the maximum joint payoff criterium. This component was more frequent in earlier tournaments, but declined with experience. As we have seen in Section 3, a population guided by this principle alone obtains a surprisingly high best reply ratio of 95% for all games (see Table 1). Even for the games without pure strategy equilibria the corresponding figure is still 89%. This indicates that the maximization of joint payoff may not be a bad heuristic, even if one is only interested in one's own payoff, at least as long as the use of the criterium is widespread in the population. However, the players learn by experience that equilibrium considerations in the case of games with pure strategy equilibria and best reply cascade concepts in the games without them are more successful ingredients of strategies.

14. Features of choices in games without pure strategy equilibria

It is of interest to analyze the question, how the importance of the best reply cascade concepts, MAP, BR-MAP, and BR-BR-MAP, changes from tournament to tournament in games without pure strategy equilibria. In the following, we explore this, not by evaluating the strategy programs, but by looking at the actions actually chosen by these programs in the tournaments.

Figure 8 shows the relative frequencies of actions chosen with the features MAP, BR-MAP, BR-BR-MAP, or none of these. It may happen that a choice fits more than one

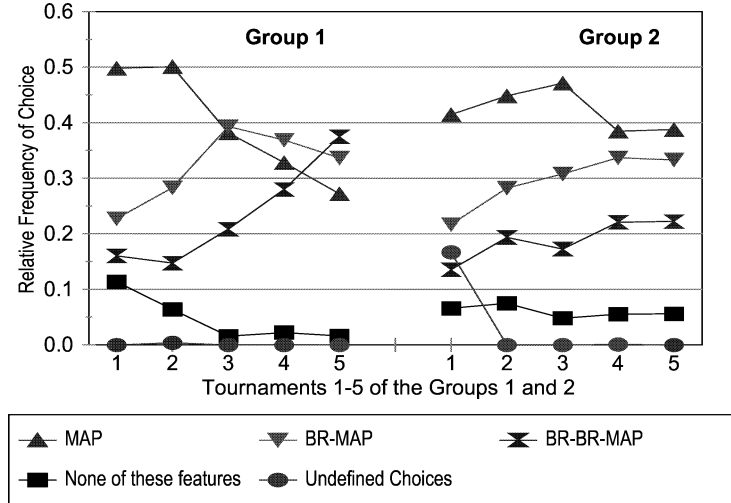


Fig. 8. Features of strategy choices in games without pure strategy equilibria.

of these categories. If a choice belongs to k categories, each is credited with $1/k$ in the counting (where $k = 1, 2, 3$).

In some cases, strategies were incomplete, so that action choices were undefined. The relative frequencies for these undefined cases are also shown in Fig. 8. Undefined action choices occurred with non-negligible relative frequency only in the first tournament of group 2 (about 16%). The next highest relative frequencies of undefined choices were 0.4% (group 1, tournament 2), 0.1% (group 2, tournament 4), and 0.0008% (group 1, tournament 1). In all other tournaments there were no undefined choices.

The count of action choice features provides an impression of the relative importance of the best reply cascade concepts, which in some sense is more informative than the information shown in Table 4. Instead of a mere indication of whether the corresponding concept is explicitly used in the strategy program or not, we receive a quantitative measure for the importance of each feature in actual play.

It can be seen that in the beginning MAP is the most frequently observed feature. In group 2, this is true for all tournaments, whereas in group 1, the importance of MAP is surpassed by BR-MAP in tournament 3. Finally, both MAP and BR-MAP are surpassed by BR-BR-MAP in tournament 5. The development in group 1 conveys the impression of a movement in the direction of higher levels of sophistication in the best reply cascade, first from MAP to BR-MAP and then both from MAP and BR-MAP to BR-BR-MAP. Group 2 does not show a similarly clear pattern. Here we also see a decline of MAP and a rise of BR-MAP as well as BR-BR-MAP, but not yet a decline of BR-MAP towards the end.

It is interesting to see that in group 1 the number of choices with none of the three features decreases up to tournament 3 and then stays on a very low level, whereas in group 2 this frequency remains at about 5% throughout all five tournaments. This also suggests a slower rise of sophistication in the use of best reply cascade concepts in group 2.

15. Average payoffs obtained by selected strategies in games without pure strategy equilibria

Further insights into the significance of the concepts of the best reply cascade can be obtained by looking at a fictitious situation in which MAP, BR-MAP, and BR-BR-MAP are the only three permissible strategies for games without pure strategy equilibria. Figure 9 shows the bimatrix obtained by restricting the strategy space in the way described above for the game base used in the tournaments. The entries represent average payoffs per play. This game has no pure strategy equilibrium. There is only one equilibrium in mixed strategies, which is indicated in Fig. 9.

The best reply relationships amongst the three pure strategies are shown in Fig. 10. BR-MAP is the best reply to MAP and BR-BR-MAP is the best reply to BR-MAP. This is not surprising, since it follows by the definition of the concept. However, it is remarkable that MAP is the best reply to BR-BR-MAP. The three best reply cascade concepts form a cycle. In fact, this may explain why higher order constructs, such as BR-BR-BR-MAP,

	MAP	BR-MAP	BR-BR-MAP
MAP	62 62	45 76	57 52
BR-MAP	76 45	56 56	36 72
BR-BR-MAP	52 57	72 36	50 50
probability in mixed strategy equilibrium	.381	.240	.379

Fig. 9. The game with strategies restricted to MAP, BR-MAP, and BR-BR-MAP.




Fig. 10. Best-Reply scheme for the game in Fig. 9.

did not evolve with experience, in the same way the concept of BR-BR-MAP had emerged after the appearance of BR-MAP.

It is not inconceivable that in an experiment with many more tournaments, behavior would converge to something similar to the mixed equilibrium in Fig. 9. However, in view of the developments shown in Fig. 8, this can be no more than a vague speculation. Yet, it is interesting to compare the mixed strategy equilibrium payoff with the observed average payoff of the last tournament:

- mixed strategy equilibrium payoff 56.025,
- observed average payoff of the last tournament 58.025.

Obviously, the two payoffs are not far apart. This seems to indicate that the analysis of the restricted game of Fig. 9 may have some predictive power for the behavior in games without pure strategy equilibria.

16. Concluding remarks

Our strategy study has led to a clear result for games with pure equilibrium points. In games with pure strategy equilibria, the subjects learned to adjust their programs in such a way that at the end almost always a pure equilibrium strategy was played. Moreover, in games with more than one pure equilibrium, there was a clear tendency to coordinate on the equilibrium with the maximum joint payoff. This shows that under favorable conditions the game theoretic notion of a pure strategy equilibrium is strongly supported by the observed behavior of sophisticated players.

Of course the tendency towards pure strategy equilibrium plays was much less pronounced in the beginning. The game theoretic equilibrium notion is not naturally present in the minds of most inexperienced subjects. But, as far as pure strategy equilibrium is concerned, it is clearly learned in repeated tournaments.

The behavior in games without pure strategy equilibria is strongly influenced by the best-reply cascade of MAP (strategy with maximum average payoff), BR-MAP (best-reply to MAP), and BR-BR-MAP (best-reply to BR-MAP). These concepts are often used by the strategy programs combined in different ways. The best-reply cascade seems to offer a natural approach to strategy planning for games without pure strategy equilibria. It is maybe worth noticing that the best-reply cascade is not prolonged beyond BR-BR-MAP.

Acknowledgments

We are indebted to Bodo Theissen for valuable programming assistance. We thank the audiences at the ESA Annual Meeting, 1998, Mannheim, and at the First World Congress of the Game Theoretic Society, 2000, Bilbao, as well as two anonymous referees for helpful comments and suggestions. We gratefully acknowledge financial support from the *Deutsche Forschungsgemeinschaft* through the *Sonderforschungsbereich 303*, from the European Union through the TMR program ENDEAR (FMRX-CT98-0238) and through the Marie Curie Fellowship HPMF-CT-199-00312 (Sadrieh), and from the *Land Nordrhein-Westfalen*.

References

- Axelrod, R., 1984. *The Evolution of Cooperation*. Basic Books, New York.
- Harsanyi, J.C., Selten, R., 1988. *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge, MA.
- Keser, C., 1991. Experimental Duopoly Markets with Demand Inertia. In: *Lecture Notes in Econom. and Math. Systems*, Vol. 391. Springer-Verlag, Berlin.
- Keser, C., Gardner, R., 1999. Strategic behavior of experienced subjects in a common pool resource game. *Int. J. Game Theory* 28, 241–252.
- Kuon, B., 1994. Two-Person Bargaining Experiments with Incomplete Information. In: *Lecture Notes in Econom. and Math. Systems*, Vol. 412. Springer-Verlag, Berlin.
- Nagel, R., 1993. Experimental Results on Interactive Competitive Guessing, SFB 303 Discussion Paper B-236, University of Bonn.
- Nagel, R., 1995. Unraveling in guessing games: An experimental study. *Amer. Econ. Rev.* 85, 1313–1326.
- Nash, J.F., 1951. Non-cooperative games. *Ann. of Math.* 54, 286–295.
- Selten, R., 1967. Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopolexperiments. In: Sauermann, H. (Ed.), *Beiträge zur experimentellen Wirtschaftsforschung*, Vol. I. Mohr, Tübingen, pp. 136–168.
- Selten, R., Mitzkewitz, M., Uhlich, G.R., 1997. Duopoly strategies programmed by experienced players. *Econometrica* 65, 517–555.
- Stahl, D.O., 1993. Evolution of smart_n players. *Games Econ. Behav.* 5, 604–617.
- Stahl, D.O., Wilson, P.W., 1995. On players' models of other players: Theory and experimental evidence. *Games Econ. Behav.* 10, 218–254.