Games and Economic Behavior **28,** 155–170 (1999)
Article ID game.1999.0723, available online at http://www.idealibrary.com on IDEAL®

REVIEWS AND COMMENTS

With the intent of stimulating discussion, this section is reserved for book reviews, comments, and letters; your input is welcome. By nature, this material may be subjective, reflecting the opinions of the authors; your responses are therefore encouraged.

Experience from a Course in Game Theory: Pre- and Postclass Problem Sets as a Didactic Device*

The paper summarizes my experience in teaching an undergraduate course in game theory in 1998. Students were required to submit two types of problem sets: preclass problem sets, which served as experiments, and postclass standard problem sets. The separation emphasizes the limited relevance of game theory as a tool for making predictions and giving advice. The paper summarizes the results of 41 experiments which were conducted during the course. It is argued that the crude experimental methods produced results which are not substantially different from those obtained using stricter experimental methods. For further information on the 41 experiments and results, see http://www.princeton.edu/~ ariel/99/gt100.html. *Journal of Economic Literature* Classification Numbers: A2, C7, C9.

1. INTRODUCTION

Teaching game theory to undergraduates has become standard in economics and other social science disciplines. This is great news for game theorists. Academic knowledge is created and circulates within a small circle of researchers for a very long time until the happy moment it enters undergraduate textbooks. In the case of game theory, that moment should not only be a cause for celebration. As game theorists, we have a responsibility for the way it is taught. We are the only people who can control and influence the content of the material taught. In particular, we have a duty to control the message that game theory transmits to the broader community.

It is my impression that most students approach a course in game theory with the belief that game theory is about the way that game situations are

* My sincere thanks to my two excellent research assistants, Yoram Hamo, during the course, and Michael Ornstein, in the writing of this paper. For further information on the 41 experiments and results, see http://www.princeton.edu/~ ariel/99/gt100.html.

0899-8256/99 $30.00

played and that its goal is to predict strategic behavior. They hope the course will provide them with the tools to better play gamelike situations. During the course, they are disappointed with the poor performance of game theory on the descriptive level and with its lack of relevance to practical problems.

The disappointment, of course, leads a student to ask, ‘‘What is the relation between the game theoretic prediction and the real world?’’ My impression is that undergraduate textbooks are vague on this question. This is so probably because we, game theorists, are confused about what the theory is trying to accomplish.

In the past, I held the radical view that undergraduate studies of game theory may negatively influence students. Students may recognize the legitimacy of manipulative considerations. They may come to believe that they need to use mixed strategies. They may tend to become more suspicious and to put less trust in verbal statements. They may adopt dogmatically game theoretic solutions. However, a pilot experiment which I conducted together with a group of graduate students at Tel Aviv University (Gilad Aharanovitz, Kfir Eliaz, Yorma Hamo, Michael Ornstein, Rani Spiegler, and Ehud Yampuler) made me less certain about this position. When we compared the responses of economics students to daily strategic situations before and after a course in game theory, we found little difference before and after the course, though there was a clear correlation between their responses and their second major.

My method of teaching for an undergraduate course in game theory is derived from my views on the relationship between theory and real life. I perceive game theory as the study of a set of considerations used (or to be used) by people in strategic situations. I do not, however, see our models as being in any way constructions or depictions of how individuals actually play gamelike situations and I have never understood how an equilibrium analysis can be used as the basis for a recommendation on how to play real games. My goal as a teacher is to deliver a loud and clear message of separation between game theoretic models and predictions of strategic behavior in real life.

Students in my class were asked to complete two types of assignments. The ‘‘postclass’’ problem sets were standard exercises that can be found in any game theory text. Students were asked to fit games to verbal situations, apply standard game theoretic solution concepts, and investigate them analytically. The answers to these problem sets were categorized as right/wrong.

The new feature here was that students were also requested to respond regularly to preclass problems sets, which were posted weekly on the course Web site. The responses were collected in a log file, which allowed me to enter the class with statistics regarding the results. It was stressed that there are no right/wrong answers to the preclass problems. The results were compared with the standard game theoretic treatment, and it was pointed out that some of the results fit the game theory analysis well while many others did not. Game theory was described as a collection of considerations which could be used or not used.

The preclass problem sets served two other purposes: First, they helped the students to concentrate on the examples discussed later in class. It is not a trivial task for students to absorb several games in one class and this method facilitated their understanding of the examples discussed. Second, they provided a cheap and convenient tool for experimentation. I am fully aware of the potential for criticism of this method: no monetary rewards were offered. However, comparisons between the results achieved in class and those received in more standard frameworks show, in my opinion (and I know this may be controversial), insignificant differences.

2. THE EXPERIMENTS

We now move to a summary of the preclass problem sets.

Noncooperative Games

The games in this category were meant to introduce the students to basic strategic considerations. The students' attention was directed to considerations which affect the outcome of a game but are excluded from the game theoretic analysis.

The game "guess the 2/3" (where a subject must announce a number between 0 and 99 with the aim of guessing "the highest integer which is no higher than 2/3 of the average of all the responses") has become a standard tool for demonstrating game theoretic considerations ("I think that they think that . . . ") and pointing out the tension between real-life behavior and analysis. The full analysis of this particular game is not trivial (changing the winning rule to "guess 2/3 of the average of the other players" would actually have been a better exercise in this early stage of the course). The game has a unique Nash equilibrium outcome in which all players choose the number 1. The results fit well with those in the literature (see, for example, Camerer, 1997; Nagel, 1995; and Thaler, 1998), with the difference that I did not allow 0. The winning number was 18. (Nagel's winning number was higher at 24, and Thaler's large experiment with *Financial Times* readers had a result of 13.) Our result is lower than Nagel's due to the fact that 20% of my students (versus almost no one in Nagel's experiment) chose the lowest number. In Thaler's experiment, I suspect that the relatively low number is an outcome of the participants' bias in favor of more sophisticated subjects.

The next game was meant to test whether player 1 is certain that player 2 will take the action which is clearly optimal for player 2:

	A	B
A	5, 5	−100, 4
B	0, 1	0, 0

The outcome (A, A) is the best outcome for both players, and player 2 has no reason to punish player 1 by playing B. Nevertheless, player 1 may be uncertain whether player 2 will employ the correct reasoning. Results: Only 20% of the students did not trust the other player and chose the safe action B. Beard and Beil (1994) tested a similar effect in a two-stage extensive game where player 1 could either take a safe action or give player 2 the option to make a choice. If player 2 makes the irrational choice, player 1 suffers a loss. Though the payoff numbers are different, the results here are in line with those in Experiment 6 in Beard and Beil (1994).

An essential assumption of simultaneous games is that each player makes a move independently of the other. Thus, information about the player who makes the first move is not part of the model. Is this a significant piece of information? In the next game, a student plays the role of player 2 in a battle-of-the-sexes game:

	A	B
A	2, 1	0, 0
B	0, 0	1, 2

The student is told that player 1 moves first, without being informed of player 2's action. Cooper *et al.* (1993) found (with slightly different numbers) that 70% of the players in the role of player 2 chose A, whereas in a standard simultaneous game, the proportion was only 35%. (Note that in their experiment, each player played the game a large number of times). Here, 44% of the students (in the role of player 2) played A. Unfortunately, I did not conduct a standard BoS game to make the proper comparison.

The previous outcome of a similar situation is an additional type of information typically ignored in game theory. In the next game, the situation was formulated to fit the battle of the sexes. Students were asked to imagine that in the last play of a similar situation, they had conceded, in other words, chosen the "inferior" action. Only 24% of the subjects chose to concede again.

The next three problems tested signaling effects in the battle of the sexes. In the first problem, player 2 has just announced that he will play his favorite action. This statement was sufficient for 80% of the students in the role of player 1 to believe the announcement and to play B. In comparison, the effect was even stronger in Cooper *et al*. (1980), where 96% of the subjects played B. (Again, note that the results there are reported for the case where each player played the game a large number of times).

In a second problem, a conversation takes place in which player 1 announces that he will play his favorite action (T), player 2 insists that he will play his favorite action (R), and player 1 remains silent. Is silence interpreted as agreement or as disagreement? In this case, students were asked to predict the outcome of the game. The results: A vast majority (85%) of the subjects predicted the outcome (T, L), implying that they interpreted silence as a confirmation of player 1's first announcement. The closest comparable game is found in Cooper *et al*. (1980), where players had made their moves after a cheap-talk stage in which they made simultaneous announcements. In these cases where one player announced T and the other remained silent, 80% of the outcomes were indeed (T, L).

Ben-Porath and Dekel (1992) provide the background to the next problem. Player 1 is notified that player 2 did not burn money although player 2 could have done so. I doubt if any of the students had in mind the considerations which Ben-Porath and Dekel described; however, the results here are interesting: 46% of the students chose action B, many more than expected without this information. The problem is similar to the one studied in Cooper *et al*. (1980), where the mere existence of an outside option for player 2 (independent of the values of that option) resulted in a stronger tendency for player 1 to yield than the standard BoS.

The next problem in this group is quite different: I attempted to test the students' intuitions as to whether eliminating a player's action can be harmful for that player. The group split more or less equally in their answers. I am not certain that I phrased the problem correctly; however, I incorporate it here since, pedagogically, it is interesting to test intuitions before discussing them formally in class.

The class of coordination games is well suited to experiments. The common finding is that people succeed in coordinating on the salient option (see, for example, Mehta *et al*., 1994). The question now is what are the characteristics of the salient option. Here, subjects were asked to coordinate on one of four alternatives labeled Fiat 97, Fiat 96, Saab 95, and Fiat 97. I wanted to test a conjecture made by Michael Bacharach: When each alternative is described in terms of a number of characteristics, the salient option is the one which is distinctive in most of the characteristics. My own conjecture was that the salient option is the one which is most distinctive from the set of most common alternatives (the Fiats in this experiment). Michael was right!

Zero-Sum Games

The class of zero-sum games is attractive as a teaching device since students are familiar with such games in daily life. Given the sharply defined predictions (in payoff terms) of equilibrium in zero-sum games, comparing results with equilibrium is simple. All the problem sets were given to the students prior to the classroom discussion of the notion of mixed strategies.

The first zero-sum game has a unique mixed-strategy equilibrium in which player 1 plays T with probability 1/3 and player 2 plays L with probability 1/3:

	L	R
T	$2, -2$	$0, 0$
B	$0, 0$	$1, -1$

Quite surprisingly, the results were close to the game theoretic prediction. T was played by 37% of the subjects!

Immediately after responding to this problem, the students were asked to play the game in the role of the column player. In this case, 86% of the students chose the action R, well above the predicted 67%. I suspect that the fact that player 2's payoffs were presented in negative numbers was the main reason for this finding. The results inspired several interesting conjectures. One of the more interesting stated that in the second game, many subjects sought to justify their previous choice.

A similar experiment is found in Fox (1972), though it did not use negative numbers. Fox found that the row player's distribution is close to 50−50 and the column player is concentrated on R.

In the next game, a subject chooses a number in the interval $[0, 100]$, aiming to be as close as possible to his or her opponent, who wants to avoid the subject. The game has many equilibria. Every choice is consistent with some equilibrium. A clear majority of the students chose the middle or the edge points.

Prior to the presentation of the maxmin theorem, students were asked to express their views as to whether "minmax is greater than maxmin" or "maxmin is greater than minmax." The maxmin criterion was presented as a pattern of reasoning in which a player thinks that his opponent will always successfully predict the player's action. The minmax criterion was presented as a pattern of reasoning in which the subject is a magician who always correctly guesses the other player's intentions, and the other player knows it. Though the minmax is never below the maxmin, the responders split almost evenly in their voting. This split demonstrates how difficult and unintuitive is this elementary inequality.

The next two problems were designed to demonstrate systematic deviations from the game theoretic predictions about zero-sum games due to framing effects. The 4-boxes problem is a repetition of an experiment conducted by Rubinstein *et al.* (1996). The subjects were asked to hide a treasure in one of four boxes placed in a row and labeled A, B, A, A. The seeker is able to open only one box. The distribution of answers (16%, 19%, 54%, 11%) is strongly biased toward the central *A* box, avoiding the edges. These results were even more pronounced than those of the original experiment (13%, 31%, 45%, 11%). By the way, the presentation of the results in class created a feeling of real discovery.

The strong tendency to avoid the edges was also obtained in "hide a treasure in a 5×5 table," where a subject hides a treasure in one of the table's 25 boxes. The 64% of the boxes placed at the edges received only 47% of the choices. This is certainly in line with the results of Ayton and Falk (1995), who asked subjects to hide three treasures in such a table. They found avoidance of the edges and strong concentration on boxes B4 and D4 (in our case, the most frequent choice was D4).

Dictator and Ultimatum Games

The dictator and ultimatum games were used to make the point that monetary and game payoffs are not identical. This is an extremely important point from an educational point of view.

The dictator game illustrates two principal modes of behavior: "taking the entire sum of money" (52%) and "sharing it equally" (35%). In other words, half of the class exhibit preferences which are not purely monetary. The concentration of subjects in the above two modes of behavior is similar to the results of Forsythe *et al.* (1994). They found that even when the players played for real money, only 35% chose to grab the whole sum.

In another version of the dictator game, a player had to decide whether to be generous with another when this generosity would not hurt the player; 81% of the subjects appeared generous.

I think that no other game has been used in more experiments than the ultimatum game. When students had to agree on the partition of 100 shekels, the offers were split into three groups: About 35% of the offers equaled 1, 39% offered 50, and 26% offered a sum between 10 and 40. In comparison, previous experiments such as those of Guth *et al.* (1982) and Forsythe *et al.* (1994) found that with or without the payment of real money, a higher percentage of subjects offered an equal split of the pie, and almost no subjects offered to retain almost all of the money.

When the roles were reversed so that students responded to a hypothetical offer of 10%, 78% accepted the offer. In comparison, Guth *et al*. (1982) found that 60% of the (small number of) subjects rejected offers of 10%, while Roth and Prasnikar (1992) found that offers below 35% were overwhelmingly rejected.

Finally, when students were asked to determine a minimal cutoff point for acceptance when facing an unknown offer, 45% of the students said that they would accept all offers (other than 0) and 35% set the cutoff point at half the divided sum. Previous results (see Harrison and McCabe, 1992) showed a much higher proportion of subjects setting the cutoff point at 50%.

Although there is no difference in the basic modes of behavior appearing among the students and those of the subjects in laboratory experiments, it seems that the students in my class were much more aggressive than the subjects in the previous experiments. This is not a very surprising fact, considering the prevailing mood in Israel (see Roth *et al*., 1991).

Extensive-Form Games

In the first problem in this category, subjects were presented with a situation, described verbally, similar to the one-shot chain store paradox game: A tailor is considering transforming his shop into a minimarket in a location where a grocery store already exists. He is afraid of a possible harmful response from the grocery store, one which will be harmful to the grocery store as well. About 53% of the students recommended that the tailor enter the food market (in line with the subgame perfect equilibrium). In contrast, Schotter *et al*. (1994) found that a much higher percentage of subjects chose to enter. The difference, in my opinion, is due to the fact that the problem here was presented verbally whereas Schotter *et al*. (1994) presented the subjects with an explicit tree which assisted them in the backward inductive reasoning. In fact, when they gave the subjects the opportunity to play the corresponding normal-form game, only 57% of the subjects chose to enter.

The next problem was intended to demonstrate that in a game situation, more rather than less information may be harmful. Students were asked how much they were willing to pay to exchange a play of the battle-of-the-sexes game for a play in a similar game in which the subject would be informed publicly about the other player's move before the subject made his own choice. The three equilibrium payoffs for the BoS version are 20, 10 (pure-equilibrium payoffs), and 6.7 (a mixed-equilibrium payoff), which is also his maxmin value. The value of the only subgame perfect equilibrium of the extensive game which fits the alternative game is 10. Thus, even under the most pessimistic view, a player should not value the offer at more than 3.3. Yet, 56% of the students were ready to pay more than 3.3 and only 26% of the students found the offer valueless.

In the next problem I followed the ideas of Camerer *et al*. (1993), who performed one of the most beautiful experiments I have ever come across. Subjects had to choose the order by which they would expose information in a two-stage extensive game. The payoff numbers were chosen to be complicated (some were negative and had many digits after the decimal point) in order to create the impression that analyzing the game was not a trivial task and required memory. Revealing B first makes the analysis easier. However, only 36% of the subjects analyzed the game from its end, whereas 64% first investigated the content of consequence A. This is definitely in line with the conclusion of Camerer *et al*. (1993) that people tend to analyze an extensive game forward rather than backward.

The last three problems regard bargaining situations:

Discounting. Students were asked to predict the outcome of a bargaining session between two bargainers possessing identical characteristics except that one is more impatient than the other. Results: 51% of the students predicted an equal split. The other subjects were equally divided as to whether the more impatient person would get more or less than half of the pie. Thus, the results do not confirm the intuition that people evaluate impatience as a negative factor in bargaining. This result is in line with the results of Ochs and Roth (1989), who demonstrated the negligible effect of the different discount rates on bargaining outcomes.

Being a proposer or being a responder. Game theoretic models suggest that being a proposer provides a strategic advantage over being a responder. However, students considered the role of responder more attractive.

Reputation. A seller has refused several offers made by the subject. How does the subject interpret the rejections? Results: 62% of the students did not find the refusal informative, but almost all the rest considered the refusal to be an indication that the value of the item is higher than was initially thought.

Finite-Horizon Games

The family of finite-horizon games has been used mainly as a tool in teaching the meaning of strategy in extensive games.

The first game in this class was the 100-period centipede game. The results: Very few students (about 10%) chose to stop the game immediately, as Nash equilibrium predicts; 49% of the students chose to never stop; and 22% chose to stop the game at the last or the penultimate opportunity. The closest comparison to these results comes from Nagel and Tang (1998), who found (in a six-period centipede game) that almost no subjects were following the Nash equilibrium strategies and that the vast majority of subjects were stopping two or three periods before the end of the game. (See also McKelvey and Palfrey, 1992).

Not all problems were intended to refute the game theoretic predictions. In a game with a structure similar to that of the centipede game, the subject was the first in a sequence of players to decide whether to stop or to pass the game to the next player in line, with payoffs that made all players pass the unique subgame perfect equilibrium. Results: About 62% of the students indeed chose pass.

In two other experiments, students were asked to play a four-period repetition of the prisoner's dilemma game and of the battle of the sexes. Students were asked to specify their strategies as plans of action; that is, they were not asked to specify actions following histories which contradict their own plans. The objective was to emphasize the contrast between the formal concept of a strategy and the intuitive notion of a strategy as a plan of action (see Rubinstein, 1991).

In the repeated PD, only 34% of the subjects chose C in the first period. The results contain a large number of strategies, many of which were difficult to interpret (in retrospect, it would have been better to ask the students to describe their strategies in words, as well). About half of the students chose to play constant D, 5% of the students chose to play constant C, and 5% chose the tit-for-tat strategy. The closest previous experiment is that of Selten and Stoecker (1986) (however, the results are difficult to compare).

In the repeated BoS, only 10% of the subjects started the game by playing the less favorable action. Once again, the results contained a large variety of strategies. Two strategies were most frequent: 22% of the students started the game by playing the more favorable action and continuing to play the best response against their opponent's last played action; 12% of the students chose the strategy "play the favorable action unless, in the past, the opponent played his favorable action in a strict majority of the periods."

Randomization

In class I presented various interpretations of mixed strategies (see Osborne and Rubinstein, 1994, Chap. 3):

- The naive interpretation—a player chooses a random device such as a roulette wheel.
- The purification idea—a player's behavior depends deterministically on unobserved factors.

- The beliefs interpretation—a player's mixed strategy is what other players think about a player's behavior.

- The large population interpretation—a mixed strategy is a distribution of the modes of behavior displayed by a large population of players who are matched randomly in order to play the game.

Adopting either of the first two interpretations led to a discussion of random behavior. My aim was to demonstrate to the students that when people choose random behavior, they create patterns of behavior which are not so random.

In one experiment students were asked to choose an integer between 1 and 9. In Simon (1971) and in Simon and Primavera (1972), the number 7 was clearly the most frequent choice (33% and 24% of the subjects, respectively, chose 7 from among the numbers $0, 1, \ldots, 9$). Here, the conjecture that people overchoose 7 was confirmed. In fact, 7 and 5 were the most frequent choices at 17% each.

In another experiment, students were given the following task: "randomly choose 4 of the integers $\{1, 2, \ldots, 8\}$." Three patterns were observed in the results:

1. Of the 70 possible answers, only 2 were chosen by more than three students: seven students chose "1234" and six chose "1357."

2. Although 21% of the possible sequences do not include either 1 or 2, only 6% of the students actually excluded both 1 and 2 from their chosen sequence.

3. The number 6 was chosen by 30% of the students, far less than all other numbers. Even when we exclude the students who chose the sequence "1234" or "1357," the proportion of answers which did not include 6 was well above 50%. One explanation is that subjects who chose three numbers from 1, 2, 3, 4, 5 felt they must also choose one of the last two numbers; thus, they skipped 6.

Excess randomization is also well documented in the literature, under the name "matching probabilities." Following an idea I was working on with the late Amos Tversky, students had to guess the second major of five randomly chosen students who study economics in a double-major program. Though the distribution of the second major was not given to the students, one could expect that they had some idea, and, in any case, maximization of the probabilities to win the prize should have led them to choose the major they believe is the most frequent among the five guesses. The results clearly demonstrate an excess use of randomization. Only 28% of the students repeated the same guess five times. All the others diversified their answers and included at least three different choices in their list of guesses. The results correspond well with those of Loewenstein and Read (1995), who demonstrated a strong tendency toward diversification when subjects had to choose a sequence of three items even though one item was viewed by them as superior to the others.

In the next experiment, students were asked to choose one of four possibilities, which yield prizes with probabilities of 21%, 27%, 32%, and 20%, respectively. The students were explicitly allowed to randomize. In this case, 69% of the students chose the most promising option without randomizing and 28% of the students chose to randomize. This question was presented during the first class of the course. When we presented the question to another group of students just after their exam in a similar game theory course, the proportion of randomizers was about the same.

Failures

A course in game theory is not a course on rational behavior. However, I consider it important to demonstrate the limits of rational behavior to the students.

The problem of count the number of F's was distributed this year on the Internet (I do not know who initiated this beautiful problem). A subject was asked to count the number of F's in a 90-letter text. The question was given, first of all, for fun but it was also supposed to remind the students that people often make systematic mistakes. Counting the number of F's in an 81-letter text is supposed to be a trivial task, but only 36% of the students did it right. The most common answer, 3, was given by a quarter of the students.

The next problem is the simplest problem yet devised which demonstrates the winner's-curse phenomenon. A student has to bid for an object; the student will only receive it if he offers more than its real value, a number which is distributed uniformly in the interval $[0, 1000]$. The bidder will then be able to sell the item for 150% of its real value. The problem was studied first (with and without real money) by Samuelson and Bazerman (1985). The results here were not significantly different (with the exception that 18% of the students answered 1000). Only 12% of the students gave the optimal offer, 0. A large proportion of subjects (24%) chose the bid 500 and 42% bid in the range 500–750.

Another puzzle, which has been widely discussed in the last few years, is the exchange of envelopes. Two positive numbers, one twice the size of the other, are placed in two sealed envelopes. The subject randomly receives one of the envelopes and another person tries to persuade the subject to exchange the envelopes with the claim: "if the number x is in your envelope, the expected value of the other envelope is $3x/2$." This problem was offered to the students at a late stage in the course followed by a postclass problem where the students had to understand the Brams and Kilgour (1995) game theoretic treatment of the problem. Only 22% of the students expressed willingness to exchange envelopes.

Some other problems dealt with the basic assumptions of the VNM theory of decision making under uncertainty.

A variant of the Allais paradox (originating in Kahneman and Tversky (1979)) was presented to the students. Students had to choose between two lotteries: one which yields $4000 with probability 0.2 and a second which yields $3000 with probability 0.25. The results: 72% of the students chose the first lottery, very close to the results of Kahneman and Tversky (1979), where 65% of the subjects chose the first lottery. Assuming that the subjects would prefer the certain $3000 to a lottery which yields $4000 with probability 0.8, these results conflict with expected utility theory.

A basic principle of rationality in decision making under uncertainty is the ‘‘sure thing’’ principle: If action D is better than action C under two exclusive circumstances, then D is also better than C when the decision maker knows that only one of the two circumstances is true. Shafir and Tversky (1992) showed that more people choose to cooperate in the prisoner’s dilemma than in either of the two cases in which they are told that the other player had cooperated or defected. Here, the problems were given to the students in the order ‘‘player 2 has made up his mind to cooperate,’’ ‘‘player 2 has made up his mind to defect,’’ and a regular prisoner’s dilemma. The results were in line with those of Shafir and Tversky (1992): Only 9% cooperated when the other player did so, only 4% cooperated when the other player defected, and 16% cooperated when they did not know what the other player chose (3%, 16%, and 37%) were the corresponding figures in Shafir and Tversky (1992), with similar though not identical payoffs).

Ethical Values

Does game theory affect the ethical attitudes of students concerning behavior in strategic situations? The suspicion may arise that game theory intensifies self motivations, strengthens manipulative attitudes, reduces the importance of ethical considerations, and so forth. In collaboration with a group of students, Gilad Aharanovitz, Kfir Eliaz, Yoram Hamo, Michael Ornstein, Rani Spiegler, and Ehud Yampuler, we gave a series of questions to the students at the first meeting of the course. We compared the results with those of a similar group of students who had just completed a similar course given by another teacher. The results did not show any clear difference between students’ behavior before and after taking the course.

We still feel that more experiments should be conducted on the subject. In the meantime, let it suffice to present the results of two of our problems.

In order to examine the tendency of students to behave manipulatively in elections, they were presented with some hypothetical election situations in which their candidates were doomed to lose but where they could increase chances of their second-best choice. Indeed, 79% of the students were prepared not to vote for their favorite candidate in order to help their second-best choice to win. (For a related experiment see Eckel and Holt, 1989).

In another problem, students were placed in the role of auto dealers who had offered a price and then received information that the potential buyer was ready to pay more than they had offered the buyer. They then had to decide whether or not to raise their price. Here, 55% of the students stated that they would not raise the price.

3. A SHORT COMMENT ON EXPERIMENTAL METHODS

The main purpose of this paper was to summarize my experience in teaching an undergraduate course in game theory. However, in retrospect, I feel that the experience presents an opportunity to evaluate the experimental methods used in game theory. Researchers are split into two camps: Some create careful laboratory environments and pay the subjects monetary rewards for their performance in the experiment; others ask subjects to fill out questionnaires requiring them to speculate on hypothetical situations.

I fall into the second category. My impression is that the results are as significant as those obtained under more sterile conditions, in the sense that the same modes of behavior appear in both sets of results. If we were interested in obtaining precise statistics regarding the appearance of those modes of behavior in the general population, then both methods would be deeply flawed since our subjects are never chosen from random representative samples. In cases where the previous results differ quantitatively (as they do in the dictator and ultimatum games, for example), the distribution of modes of behavior is clearly affected by culture, education, and personal characteristics; hence, there is no reason to expect uniform results. I would therefore like to stress my doubts as to the necessity of laboratory conditions and the use of real money in experimental game theory.

REFERENCES

Ayton, P., and Falk, R. (1995). "Subjective Randomness in Hide-and-Seek Games," in *Books of Abstracts of the 15th Bi-annual Conference on Subjective Probability, Utility and Decision-Making*, p. 37.

Beard, T., and Beil, R. Jr. (1994). "Do People Rely on the Self-Interested Maximization of Others? An Experimental Test," *Manag. Sci.* **40**, 252–262.

Ben-Porath, E., and Dekel, E. (1992). "Signaling Future Actions and the Potential for Self-Sacrifice," *J. Econ. Theory* **57**, 36–51.

Brams, S., and Kilgour, M. (1995). "The Box Problem: To Switch or Not to Switch," *Math. Magazine* **68**, 27–34.

Camerer, C., Johnson, E., Rymon, T., and Sen, S. (1993). "Cognition and Framing in Sequential Bargaining for Gains and Losses," in *Frontiers of Game Theory* (K. Binmore, A. Kirman, and P. Tani, Eds.), pp. 27–48. Boston: MIT Press.

Camerer, C. (1997). "Progress in Behavioral Game Theory," *J. Econ. Perspect.* **11**, 167–188.

Cooper, R., DeJong, D. V., Forsythe, R., and Ross, T. (1980). "Alternative Institutions for Resolving Coordination Problems: Experimental Evidence on Forward Induction and Preplay Communication," in *Experiment in Non-cooperative Oligopoly* (J. W. Friedman and A. C. Hoggat, Eds.), pp. 129–146. Stamford, CT: JAI Press.

Cooper, R., DeJong D. V., Forsythe, R., and Ross, T. (1993). "Forward Induction in the Battle of the Sexes Games," *Amer. Econ. Rev.* **83**, 1303–1316.

Eckel, C., and Holt, C. (1989). "Strategic Voting in Agenda-Controlled Committee Experiments," *Amer. Econ. Rev.* **79**, 763–773.

Forsythe, R., Horowitz, J., Savin, N., and Sefton, M. (1994). "Fairness in Simple Bargaining Experiments," *Games Econ. Behav.* **6**, 347–369.

Fox, J. (1972). "The Learning of Strategies in a Simple Two-Person Zero-Sum Game Without a Saddlepoint," *Behav. Sci.* **17**, 300–308.

Guth, W., Schmitteberger, R., and Schwarze, B. (1982). "An Experimental Analysis of Ultimatum Bargaining," *J. Econ. Behav. Organization* **3**, 367–388.

Harison, G., and McCabe, K. (1992). Expectations and Fairness in a Simple Bargaining Experiment, mimeo.

Kahneman, D., and Tversky, A. (1979). "Prospect Theory: An Analysis of Decision under Risk," *Econometrica* **47**, 263–292.

Loewenstein, G., and Read, D. (1995). "Diversification Bias: Explaining the Discrepancy in Variety Seeking between Combined and Separated Choices," *J. Experimental Psychology* **1**, 34–49.

McKelvey, R., and Palfrey, R. (1992). "An Experimental Study of the Centipede Game," *Econometrica* **60**, 803–836.

Mehta, J., Starmer, C., and Sugden, R. (1994). "The Nature of Salience: An Experimental Investigation of Pure Coordination Games," *Amer. Econ. Rev.* **84**, 658–673.

Nagel, R. (1995). "Unraveling in Guessing Games: An Experimental Study," *Amer. Econ. Rev.* **85**, 1313–1326.

Nagel, R., and Tang, F. (1998). "Experimental Results on the Centipede Game in Normal Form: An Investigation on Learning," *J. Math. Psychology* **42**, 356–384.

Ochs, J., and Roth, E. (1989). "An Experimental Study of Sequential Bargaining," *Amer. Econ. Rev.* **79**, 355–384.

Osborne, M., and Rubinstein, A. (1994). *A Course in Game Theory*, Cambridge, MA: MIT Press.

Roth, E., Prasnikar, V., Okuno-Fujiwara, M., and Zamir, S. (1991). "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh and Tokyo: An Experimental Study," *Amer. Econ. Rev.* **81**, 1068–1095.

Roth, E., and Prasnikar, V. (1992). "Considerations of Fairness and Strategy: Experimental Data from Sequential Games," *Quart. J. Econom.* **107**, 865–888.

Rubinstein, A. (1991). ''Comments on the Interpretation of Game Theory,'' *Econometrica* **59**, 904–921.

Rubinstein, A., Tversky, A., and Heller, D. (1996). ''Naive Strategies in Zero-Sum Games,'' in *Understanding Strategic Interaction—Essays in Honor of Reinhard Selten* (W. Guth *et al*., Eds.), pp. 394–402. Berlin/New York: Springer-Verlag.

Samuelson, W., and Bazerman, M. (1985). ''The Winner's Curse in Bilateral Negotiations,'' in *Research in Experimental Economics* Vol. 3, V. L. Smith, Ed. pp. 105–137. Greenwich, CT: JAI Press.

Schotter, A., Weilgelt, K., and Wilson, C. (1994). ''A Laboratory Investigation of Multiperson Rationality and Presentation Effects,'' *Games Econ. Behav*. **6**, 445–468.

Selten, R., and Stoecker, R. (1986). ''End Behavior in Sequences of Finite Prisoner's Dilemma Supergames,'' *J. Econ. Behav. Organization* **7**, 47–70.

Shafir, E., and Tversky, A. (1992). ''Thinking through Uncertainty: Nonconsequential Reasoning and Choice,'' *Cognitive Psychol*. **24**, 449–474.

Simon, W. (1971). ''Number and Color Responses of Some College Students: Preliminary Evidence for 'Blue Seven Phenomenon','' *Perceptual Motor Skills* **84**, 658–673.

Simon, W., and Primavera, L. (1972). ''Investigation of the 'Blue Seven Phenomenon' in Elementary and High School Children,'' *Psychol. Rep*. **31**, 128–130.

Thaler, R. (1998). ''Giving Markets a Human Dimension,'' *Mastering Finance in Financial Times* **6**, 4.

Ariel Rubinstein

School of Economics
Tel Aviv University
Tel Aviv, Israel 69978
and
Department of Economics
Princeton University
Princeton, New Jersey 08544

E-mail: rariel@post.tau.ac.il
ariel@princeton.edu