Expectations and Reputations in Bargaining: An Experimental Study

By ALVIN E. ROTH AND FRANCOISE SCHOUMAKER*

One task of a theory of bargaining is to specify those factors which ultimately resolve the indeterminacy inherent in bargaining. Even if we expect bargainers to reach agreement in the region of individually rational, Pareto optimal contracts—which F. Y. Edgeworth (1881) called the contract curve—it is still necessary to analyze the factors contributing to the selection of a specific agreement.

Game-theoretic models rest upon the assumption that the outcome of bargaining among rational, fully informed agents is determined by the strategic possibilities available to the bargainers, and their preferences as represented by their expected utility functions. Indeed, games in which the players have this information are called games of *complete information*. A long-standing obstacle to testing the descriptive power of these models has been the requirement that the utility of the bargainers be known for each potential agreement. An experimental design permitting this problem to be circumvented in a laboratory setting was introduced in Roth and Michael Malouf (1979). In that experiment, the outcomes of bargaining under controlled conditions were observed to systematically deviate from the predictions of theory. Subsequent experiments, (Roth and Malouf, 1979; Roth, Malouf and J. Keith Murnighan, 1981; Roth and Murnighan, 1982) indicated that the outcome of bargaining is decisively influenced by factors other than the strategic possibilities and preferences of the bargainers, and isolated the cause of the observed deviations from the predictions of classical models, sufficiently so that a specific hypothesis could be proposed to account for them. This paper examines the hypothesis that those missing factors concern the subjective expectations of the bargainers about the behavior of their opponents. The results of the experiment reported here support the hypothesis and demonstrate that distinct, stable, self-fulfilling sets of expectations are compatible with a given bargaining situation as determined by the preferences and strategic possibilities of the bargainers.

I. Review of Three Earlier Experiments

To test theories which depend on the expected utilities of the players, it is desirable to design experiments which allow the participants' utility functions to be determined. A class of games allowing this was introduced in Roth and Malouf. In that experiment, players bargained over the *probability* that they would receive some monetary prize, possibly a different prize for each player. Specifically, they bargained over how to distribute "lottery tickets" to determine the probability that each player would win his personal lottery (i.e., a player who received 40 percent of the lottery tickets would have a 40 percent chance of winning his monetary prize and a 60 percent chance of winning nothing). The rules of the game specified which distributions of lottery tickets were allowable. If no agreement was reached in the allotted time, each player received nothing. We call such games, in which each player has only two possible monetary payoffs, *binary lottery games*.

To interpret the outcomes of a binary lottery game in terms of each player's utility for money, recall that if each player's utility function is normalized so the utility for re-

*Mellon professor of economics, University of Pittsburg, and assistant professor of economics, University of Illinois, respectively. This work has been supported by NSF Grants SES 79-15356 and SOC 78-09928. We acknowledge the help of Michael Barr and David Sides in setting up experimental procedures, and invaluable assistance from our colleagues, Lawrence M. DeBrock and J. Keith Murnighan, in analyzing our data.

ceiving his prize is 1, and the utility for receiving nothing is 0, then the player's utility for any lottery between these two alternatives is the probability of winning the lottery. The set of feasible utility payoffs in such a game equals the set of allowable divisions of lottery tickets. Thus binary lottery games can be used to experimentally test theories of bargaining which depend on the set of feasible utility payoffs. Note that the set of feasible utility payoffs does not depend on the prizes, so a binary lottery game in which the players know the allowable divisions of lottery tickets is a game of complete information, regardless of whether each player also knows the other's prize. The classical models of bargaining that follow the work of John Nash (1950) and depend only on the set of feasible utility payoffs to the players (see Roth, 1979) thus predict that the outcome of a binary lottery game will not depend on whether the players know their opponent's prize.

The experiment of Roth and Malouf was designed to test this hypothesis, among others. Participants played binary lottery games under either *full* or *partial* information. In the full information condition, each player knew his own potential prize and his opponent's potential prize. Under partial information, each player knew only his own prize. Contrary to the predictions of the classical models in the tradition of Nash (1950), the outcomes observed in the two information conditions exhibited dramatic differences: under partial information, outcomes tended to be very close to an equal division of the lottery tickets, while under full information, outcomes showed a pronounced shift toward equal expected payoffs.

Of course, other classical models describe games in greater detail. The *strategic* form of a game describes not only the feasible utility payoffs, but also the strategies available to the players. In the games described above, strategy choices concern the formulation of messages and proposals during negotiations. Since the strategies available to the players depend on the information they possess, we must consider whether the observed results can be accounted for by the different strategies available to the players in the two information conditions.

The experiment of Roth, Malouf, and Murnighan was designed to address this question and involved binary lottery games with prizes stated in terms of an intermediate commodity. Prizes were expressed in chips having monetary value, and each player played four games under either *high*, *intermediate*, or *low* information. In each condition, each player knew the number of chips in his potential prize and their value, but a player's information about his opponent's prize varied with the information condition. In the high information condition, each player knew the number of chips in his opponent's prize and their value. Under intermediate information, each player knew the number of chips in his opponent's prize, but not their value. Under low information, each player knew neither the number of chips in his opponent's prize, nor their value. In the latter two conditions, players were prevented from communicating the missing information about the prizes.

The experiment took advantage of two kinds of strategic equivalence relations. Binary lottery games with prizes expressed in both chips and money, played in the low information condition of this experiment, are strategically equivalent[1] to binary lottery games with the same monetary prizes expressed in money alone, played in the partial information condition of the previous experiment, because under the rules of the two information conditions, any legal message in one kind of game would be legal in the other. So the strategy sets are the same for both kinds of games, as are the utility functions and the underlying set of alternatives. Also, games expressed in both chips and money, played in the intermediate information condition of this experiment, are strategically equivalent to games expressed in money alone played in the full information condition of

[1]Two games are strategically equivalent if they can both be represented by the same game in strategic form. Thus, any theory of games which depends only on the strategic form of a game yields the same prediction for strategically equivalent games. (Compare Roth, Malouf, and Murnighan, 1981.)

the previous experiment, so long as the monetary values of the two prizes in each money game are in the same proportion as the numbers of chips in the prizes in the corresponding chip game. This is because any legal message in one kind of game can be transformed into a legal message in the other kind of game by substituting references to chips for references to money (or vice versa) in any message concerning the value of the prizes.

If the observed difference between the partial and full information conditions of the previous experiment was due to the players' different strategy sets in the two conditions (the "strategic hypothesis"), then a similar difference should be observed between the low and intermediate information conditions of this experiment. The observed results did not support the strategic hypothesis. The low and high information conditions replicated the partial and full information conditions of the previous experiment, but the outcomes observed in the intermediate information condition did not differ significantly from those in the low information condition. In the intermediate information condition, the observed agreements tended to give both players equal probabilities, regardless of the size of their prize in chips. Thus, information about the artificial commodity, chips, did not affect the outcomes in the same way as did strategically equivalent information about money.

Both the above experiments revealed an effect of information which cannot be explained by existing models. The experiment of Roth and Murnighan was conducted to separate this effect into components resulting from the possession of specific information by specific individuals, and to assess the extent to which the observed behavior can be characterized as equilibrium behavior.[2]

In the two earlier experiments, either both bargainers knew their opponent's prize, or neither bargainer knew his opponent's prize. Also, it was always common knowledge whether the bargainers knew one another's prizes. Information is common knowledge in a game if it is known to all of the players, and if every player knows that all the players know, and that every player knows the others know that *he* knows, and so forth (compare David Lewis, 1969; Robert Aumann, 1976; and Paul Milgrom, 1981). Two bargainers can be thought of as having common knowledge about an event if the event occurs when both of them are present to see it, so that they also see each other seeing it, etc. In these experiments, a set of instructions provides common knowledge to the bargainers if it contains the information that both of them are receiving exactly the same instructions.

Each game of the third experiment was a binary lottery game in which one player had a \$20 prize and the other a \$5 prize. In each of the eight conditions of the experiment, each player knew at least his own prize. The experiment used a 4 (information)×2 (common knowledge) factorial design. The information conditions were: 1) *neither knows* his opponent's prize; 2) the *\$20 player knows* both prizes, but the \$5 player knows only his own prize; 3) the *\$5 player knows* both prizes, but the \$20 player knows only his own prize; and 4) *both players know* both prizes. The second factor made this information common knowledge for half the bargaining pairs, but not common knowledge for the other half.

The results of this experiment permitted three principal conclusions. First, the effect of information on the agreements reached is primarily a function of whether the player with the smaller prize knows both prizes. Second, whether this information is common knowledge influences the frequency of disagreement, with more disagreements occurring in the noncommon knowledge conditions. Third, in the noncommon knowledge conditions, the relationship among the outcomes showed virtually no departure from equilibrium.

Together, these three experiments permit us to speculate fairly specifically on the cause of the observed information effects. The first experiment demonstrated an effect of information about the prizes which could not be accounted for in terms of players' preferences over consequences (lotteries). The second experiment showed that this effect could

[2] We refer here to the familiar *noncooperative* equilibrium of Nash (1951).

not be accounted for by the set of available actions (strategies). The third experiment showed that the effect is consistent with rational (equilibrium) behavior. Thus, if we continue to hypothesize that the players are (approximately) Bayesian utility maximizers, it must be that the effect of information is due to a change in the players' subjective beliefs. For example, information about the prizes, and whether this information is common knowledge, may influence the players' subjective probabilities concerning what agreements are likely to be acceptable to their opponents.

To see how the bargainers' expectations might influence the outcome, consider the following "thought experiment." A randomly selected individual plays some large number of games in which he bargains over how to divide a certain sum of money. Although he doesn't know it, all his opponents are confederates of the experimenter, and they all allow him to obtain, say, 80 percent of the available money. After he has gone through this experience, *you* have the opportunity of bargaining with him on your own behalf (i.e., not as a confederate). His past success is common knowledge. It will be difficult to bargain with him on an equal basis, since he expects (and has every reason to expect) to receive 80 percent of the available money, and since he expects (and has every reason to expect) that you will concede it to him. Suppose the rules of the game are that, after completing any negotiations, the players each separately write down their demands. They receive their demands if they are compatible, and otherwise receive nothing. If this is the only time you will be bargaining with him, the fact that this randomly selected individual now expects to get 80 percent will make it very risky for you to write down a demand of more than 20 percent.

To make precise how such subjective expectations can affect the decisions made by bargainers, consider a simple model of bargaining in a two-stage binary lottery game. In stage one, each individual i makes a demand: he states the probability p_i (of winning his prize) which he wants and offers $1-p_i$ to his opponent. In stage two, each bargainer chooses between repeating his demand or accepting his opponent's offer. An agreement occurs whenever the probabilities demanded add up to no more than 1. No messages can be exchanged.

If, in stage one, the two bargainers' demands add up to no more than 1, an agreement is reached at which each bargainer i wins his prize with probability p_i. If the probabilities in stage one add up to more than 1, the outcome depends on the players' decisions in stage two. If each repeats his demand, a disagreement results and each will have a zero probability of winning his prize. If player i repeats his demand p_i and player j accepts his opponent's offer, there is an agreement on i's terms: i will win his prize with probability p_i and j will win his prize with probability $(1-p_i)$. If both i and j accept the other's offer, i will win his prize with probability $(1-p_j)$ and j will win his prize with probability $(1-p_i)$.

Having both stated their demands in stage one, each bargainer must decide in stage two whether to repeat his demand or accept the other's offer. His expectations as to his opponent's behavior play a crucial role. The hypothesis which the experiment described in the next section is designed to test is that the expectations of the bargainers can be manipulated independently of the strategic possibilities and feasible outcomes of the bargaining situation. This is at odds with the traditional view, perhaps most explicitly stated by John Harsanyi (1977), who considers two-stage bargaining games of essentially this form, and argues that rational players' subjective probabilities are determined by the data of the game.

II. The New Experiment

Agreements in the previous experiments tended to cluster around two divisions of the lottery tickets: one kind of agreement split the lottery tickets equally between the bargainers, the other gave the bargainers equal expected monetary payoffs. This experiment investigates whether, by manipulating the bargainers' expectations, one or the other of these two kinds of agreements can be obtained as a stable equilibrium.

Each player played 25 identical two-stage binary lottery games, of the kind discussed above. Although players were told that they bargained with another individual in each game, each individual in fact played against a programmed opponent (the computer) in the first 15 games, as in the thought experiment. Half the participants had a prize of $40 and half a prize of $10; both players knew both prizes, and players whose prize was $40 always bargained against players whose prize was $10 (each player had the same prize in all 25 games). Subjects were divided into three experimental conditions. The first was a 20–80 condition in which the computer was programmed to promote a 20–80 division of the lottery tickets, which yields equal expected payoffs. The second was a 50–50 condition in which subjects bargained with a computer programmed to promote the equal division of lottery tickets. The third condition was the control: subjects never bargained with the computer but always with other members of that group.

In the 20–80 condition, the agent whose prize was $40 bargained with a computer programmed to randomly select a first demand between 75 percent and 80 percent, and to repeat its demand in the second stage. The programmed opponent of the $10 player randomly selected a demand between 20 and 25 percent, and in stage two it accepted any offer giving it at least 20 percent of the lottery tickets. In the 50–50 condition, the programmed opponent of the $40 player randomly selected a first demand between 70 and 75 percent of the lottery tickets and in stage two accepted any offer giving it at least 50 percent. The $10 player bargained with a computer that randomly selected a first demand between 45 and 50 percent of the lottery tickets, and in stage two always repeated its demand.

In trials 16 to 25, subjects in each group bargained with other members of that group. Each game was played with a different anonymous opponent. Bargainers also received additional information, about their opponent's reputation as established in trials 11 through 15. They were told their opponent's first demand, whether he repeated it or accepted his opponent's offer, and which agreement, if any, was reached in each of trials 11 to 15, that is, in the final 5 games played against the computer, (when each player's behavior reflected his experience against a programmed opponent). In trials 16 to 25, every bargainer's experience against the programmed opponents was made common knowledge in this way, as in the thought experiment.[3]

Two elements influence the expectations of a pair of bargainers: their experiences and their reputations. Consider a typical pair in the 20–80 condition, at trial 16. The $40 player has never obtained an agreement in which he received more than 25 percent of the lottery tickets. He knows his opponent is aware of his reputation as established in games 11 to 15. Thus a $10 player, whose own experience has led him to expect that his opponent will capitulate to a demand for equal expected monetary payoffs, has his expectations reinforced when he confronts a $40 player whose reputation indicates that he has, in the past, given in to such demands. Similarly a $40 player, whose experience is that opponents are adamant about an equal division of expected payoffs, will have his expectations confirmed by a reputation indicating that his current opponent has previously behaved in this way.

The experiment is designed to distinguish between two competing hypotheses. The

[3]Players were not informed, until after trial 15, that their reputations would be displayed. So reputation in this experiment serves only as an indicator of a player's past bargaining experience (in periods 11–15). This is in contrast with models in which the players know in advance that they are building a reputation. Robert Rosenthal (1979), and Rosenthal and Henry Landau (1979) consider repeated games with complete information in which players can strategically build and maintain reputations in the course of play, in order to influence future encounters. David Kreps and Robert Wilson (1982b), Milgrom and John Roberts (1982), and Kreps et al. (1982) consider games of incomplete information in which players may seek to play in early encounters in such a way as to build a reputation that will mislead future opponents about their true objective function. However, in the games played in this experiment, the fact that players are not aware that they are building a reputation in periods 11–15, and cannot subsequently alter the reputation established in those periods, removes any possibility that an incentive to alter his reputation can influence a player's bargaining behavior.

classical game-theoretic hypothesis, which states that the outcome of a game can be predicted from the set of feasible utility payoffs and strategic possibilities, implies that, in this experiment, the different experimental conditions should have no continuing effect. Specifically, if this is correct, we would expect to observe that, starting with trial 16, any differences between the two experimental conditions and the control condition would begin to disappear and the outcomes in the three conditions should converge over time, as continued play removes any transient effects due to the initial experience of players in the 20–80 and 50–50 conditions.

But if the expectations have a critical role in determining the outcome, as suggested by the earlier experiments, then we should expect to see divergent outcomes, established in the first 15 trials, persist in a stable fashion in each of the three conditions. We would expect the first condition's mean agreement to be near 20–80 and the second condition's to be near 50–50. The control condition's mean agreement should be somewhere between these two. This would be consistent with the hypothesis that the players' expectations were the uncontrolled factor accounting for the results of the previous experiments.

III. Method

Each participant was seated at a visually isolated terminal of the PLATO computer system at the University of Illinois. Participants received all instructions and conducted all communication through the terminal. There were ten participants in each of the three conditions, which were conducted simultaneously. Subjects were undergraduates. Pretests were run with the same subject pool (but different subjects) to ensure that the instructions were easily understandable.

Background information including a brief review of probability theory was presented first. The rules of the game were then introduced. A demand p_i was a number which was the sender's probability of winning his prize; $1-p_i$ was the probability offered to his opponent. As in the previous experiments, probabilities were presented in terms of the division of lottery tickets. The PLATO system computed the expected value of each demand and associated offer for both bargainers. After being made aware of these computations, a bargainer was given the option of cancelling his demand before its transmittal. After both demands were transmitted, stage two began. In stage two, each player had the choice of repeating his own demand or accepting his opponent's offer. When both decisions were made, the bargainers were informed of the outcome.[4] Participants were told that they bargained with a different individual in each game. In both stages, the bargainers were not informed of their opponent's decision until their own decision had been transmitted.

The bargainers in the 20–80 and 50–50 condition were paired with the appropriate programmed opponent for the first 15 trials. The instructions led them to believe they were bargaining with other individuals. The members of the control group were paired with other members of the control group. After trial 15, new instructions appeared on the screen and introduced the notion of reputation. Note again that as a bargainer was establishing his reputation, in trials 11 to 15, he did not know he was doing so. Trial 16 began with each bargainer's screen displaying both his own reputation and his opponent's. Each \$40 player bargained twice with each \$10 player. The value of the prizes was common knowledge.

After game 25, the payoffs were computed as described in the initial instructions: for each bargainer, one of games 1 to 25 was randomly selected; the corresponding lottery was then conducted with the specified prizes and the probabilities agreed upon in that randomly selected game. The players were directed to the monitor who paid them.

IV. Results

Figure 1 shows the mean agreements, by trial, for each of the three conditions for the \$40 player. The figure makes clear that the

[4]There was a time limit of two minutes for the first stage and one minute for the second stage.

TABLE 1—MEAN OUTCOMES: ALL NEGOTIATIONS INCLUDED; TRIALS 16–25

	20–80	Control	50–50	Mean
\$40 Player	22.4_c	30.2_{bc}	40.8_b	31.1
\$10 Player	57.0_a	39.0_b	41.1_b	45.7
Mean	39.7	34.6	40.9	

Note: Cells with common subscripts are *not* significantly different from one another at the .05 level using the Newman-Keuels procedure.

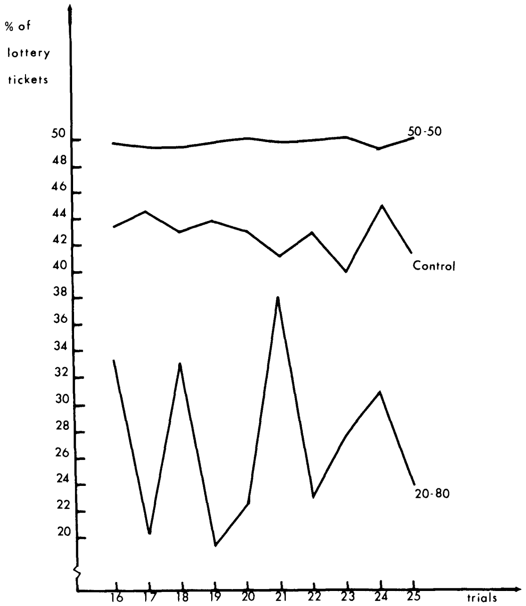

FIGURE 1. AVERAGE PERCENTAGE OF LOTTERY TICKETS OBTAINED BY \$40 PLAYER WHEN AGREEMENT WAS REACHED IN TRIALS 16–25

agreements reached in the three conditions differ markedly.

Preliminary analyses yielded no significant effects for trials for the total set of bargaining outcomes and for the set that excluded disagreements. Thus, the remaining analyses combined the data from the ten bargaining sessions. Analyses of variance for players (high prize vs. low prize) and condition (20–80, 50–50, and control) were conducted for the outcomes over all negotiations (including disagreements), for the outcomes excluding disagreements, and for the first and second offers made by the players. All the

TABLE 2—MEAN AGREEMENTS: DISAGREEMENTS EXCLUDED; TRIALS 16–25

	20–80	Control	50–50
\$40 Player	27.3_e	43.2_d	49.7_c
\$10 Player	69.5_a	55.7_b	50.2_c

Note: See Table 1.

analyses yielded significant effects for condition and for the interaction between players and reputation (F-ratios in each case exceeded 16, with $p < .001$ in each case). The mean outcomes are shown in Tables 1 and 2, along with the results of *post hoc* tests on the significant interactions indicating that, in each case, outcomes in the 20–80 condition were significantly different from the 50–50 or control condition. The agreements that were reached in the control condition (with disagreements excluded) were between those of the other two conditions. The compelling nature of a 50–50 agreement in the 50–50 condition is highlighted by the extremely low variance in the 50–50 condition in contrast with the high variance in the 20–80 condition (see Table 3, which lists all final outcomes).

Generally, the demands reflected the reputations established (see Table 4): in the control condition, most of the offers centered around 50–50. There were 9 disagreements in each of the 50–50 and 80–20 conditions, and 15 disagreements in the control condition. Although these frequencies are not significantly different from one another ($\chi^2(2) = 2.18$), they suggest that the experimental treatment contributed to a reduction in the number of disagreements and an increase in efficiency of the bargaining, as compared to the control condition.

TABLE 3—UNAGGREGATED DATA

	Control		50–50 Condition		20–80 Condition	
Trial	Outcome	Player	Outcome	Player	Outcome	Player
16	44–52	1–6	0–0	1–10	20–80	1–10
	45–55	3–8	50–50	3–2	50–50	3–2
	0–0	5–10	0–0	5–4	0–0	5–4
	0–0	7–2	50–50	7–6	25–75	7–6
	42–58	9–4	49–51	9–8	0–0	9–8
17	47–53	1–10	50–50	1–8	22–62	1–2
	50–50	3–2	53–47	3–10	35–50	3–4
	45–50	5–4	0–0	5–2	30–70	5–6
	40–60	7–6	50–50	7–4	25–60	7–8
	40–60	9–8	45–51	9–6	20–80	9–10
18	50–50	1–2	50–50	1–6	20–79	1–8
	45–55	3–4	0–0	3–8	0–0	3–10
	0–0	5–6	50–50	5–10	30–70	5–2
	40–60	7–8	50–50	7–2	50–50	7–4
	37–63	9–10	48–52	9–4	25–70	9–6
19	45–53	1–4	0–0	1–4	20–80	1–6
	50–50	3–6	50–50	3–6	0–0	3–8
	0–0	5–8	50–50	5–8	12–88	5–10
	35–65	7–10	50–50	7–10	21–79	7–2
	46–50	9–2	49–51	9–2	25–70	9–4
20	41–59	1–4	50–50	1–10	22–65	1–4
	50–50	3–6	50–50	3–2	25–50	3–6
	35–65	5–8	50–50	5–4	25–70	5–8
	0–0	7–10	50–50	7–6	21–79	7–10
	46–50	9–2	0–0	9–8	20–80	9–2
21	34–66	1–10	50–50	1–2	20–80	1–6
	50–50	3–2	0–0	3–4	50–50	3–8
	0–0	5–4	50–50	5–6	0–0	5–10
	45–50	7–6	50–50	7–8	50–50	7–2
	35–65	9–8	49–51	9–10	32–65	9–4
22	50–50	1–2	50–50	1–6	0–0	1–10
	45–55	3–4	50–50	3–8	11–89	3–2
	0–0	5–6	50–50	5–10	33–66	5–4
	0–0	7–8	50–50	7–2	0–0	7–6
	34–66	9–10	49–51	9–4	25–65	9–8
23	40–55	1–6	50–50	1–4	21–78	1–2
	0–0	3–8	0–0	3–6	35–50	3–4
	0–0	5–10	50–50	5–8	30–70	5–6
	0–0	7–2	50–50	7–10	25–75	7–8
	0–0	9–4	50–50	9–2	0–0	9–10
24	0–0	1–8	0–0	1–8	20–80	1–8
	45–55	3–10	50–50	3–10	0–0	3–10
	50–50	5–2	50–50	5–2	24–76	5–2
	45–55	7–4	50–50	7–4	50–50	7–4
	41–58	9–6	47–53	9–6	30–70	9–6
25	50–50	1–2	50–50	1–2	20–80	1–10
	0–0	3–4	50–50	3–4	20–80	3–2
	40–60	5–6	50–50	5–6	30–70	5–4
	0–0	7–8	50–50	7–8	25–75	7–6
	34–66	9–10	49–50	9–10	25–75	9–8

Note: For each condition, for each trial, each line gives the outcomes for the player (with the $40 player's outcome first) and the players' number. Notice that odd numbered players are the $40 ones. For example, the first line in the control condition reads: ($40) player 1 bargained with ($10) player 6 and the outcome was 44 percent of the lottery tickets for 1 and 52 percent of the lottery tickets for player 6.

TABLE 4—REPUTATION[a]

		Control Condition			50–50 Condition			20–80 Condition		
Player	Trial	First Demand	Second Demand	Outcome	First Demand	Second Demand	Outcome	First Demand	Second Demand	Outcome
1	11	62	35	35	50	50	50	20	20	20
	12	55	55	55	51	51	0	20	20	20
	13	55	51	51	50	50	50	23	20	20
	14	47	47	0	51	51	0	23	20	20
	15	55	35	35	50	50	50	23	20	20
2	11	49	49	49	35	51	51	84	79	79
	12	49	49	49	55	52	52	77	77	77
	13	49	49	49	55	52	52	86	78	78
	14	49	49	49	55	50	50	92	79	79
	15	55	50	50	55	53	53	80	80	80
3	11	50	50	50	50	50	50	70	21	21
	12	50	50	0	55	55	0	80	24	24
	13	50	50	50	50	50	50	100	20	20
	14	50	40	40	51	51	0	50	25	25
	15	50	50	50	50	50	50	65	21	21
4	11	68	68	68	55	51	51	80	79	79
	12	62	62	62	55	51	51	65	65	65
	13	85	50	50	80	52	52	100	75	75
	14	65	65	0	80	80	0	100	78	78
	15	56	56	50	80	52	52	100	79	79
5	11	51	32	32	55	55	0	29	29	0
	12	51	51	0	50	50	50	70	70	0
	13	51	51	0	72	27	27	40	40	0
	14	50	40	40	60	29	29	35	35	0
	15	50	45	45	50	50	50	46	20	20
6	11	75	75	0	54	53	53	82	82	0
	12	75	45	45	53	53	0	80	80	80
	13	75	75	0	52	52	52	80	80	80
	14	60	50	50	65	52	52	80	80	80
	15	60	56	56	59	52	52	80	80	80
7	11	90	90	0	50	50	50	27	20	20
	12	50	50	50	50	50	50	25	23	23
	13	55	40	40	51	51	0	27	20	20
	14	75	75	0	50	50	50	26	22	22
	15	55	44	44	50	50	50	30	20	20
8	11	65	65	0	51	51	51	85	78	78
	12	60	60	0	51	51	51	80	74	74
	13	60	60	60	49	49	49	90	78	78
	14	60	60	60	50	50	50	80	80	80
	15	65	65	65	51	51	0	80	79	79
9	11	45	45	0	52	52	0	30	20	20
	12	44	38	38	47	28	28	35	20	20
	13	43	43	0	49	49	49	35	24	24
	14	50	50	50	49	49	49	55	20	20
	15	44	40	40	49	49	49	33	20	20
10	11	65	65	65	54	53	53	82	82	0
	12	70	70	0	54	52	52	80	80	80
	13	76	76	0	52	52	52	80	80	80
	14	66	66	0	53	53	53	81	81	0
	15	69	50	50	54	51	51	80	80	80

[a]Odd numbered players had a $40 prize, even numbered players had a $10 prize.

V. Discussion

The results of this experiment provide strong support for the hypothesis that the outcome of the bargaining is influenced not only by the preferences and strategic options of the bargainers, but also by their expectations. By manipulating these expectations, it proved possible to consistently produce outcomes which differed significantly from those observed when the players' expectations were not manipulated. The fact that there was no significant difference over trials indicates that these expectations were self-reinforcing, so that the outcomes which resulted in the experimental conditions were stable and self-sustaining. Players who expected their opponents to expect equal payoffs (or an equal division of lottery tickets) continued to meet such opponents, and consequently those expectations were reinforced.[5]

These results also lend support to the hypothesis put forward in Section I to explain the results of the previous experiments. Since those experiments involved no artificial manipulation of the bargainers' expectations, it must be that the (common) expectations which bargainers formed when different kinds of information were available resulted in large measure from their previous experience. This suggests that there must be many kinds of potential conflict in which individuals have common expectations which permit them to efficiently reach agreement. The benefits to a society of fostering such common expectations are obvious, since otherwise many bargaining situations would end in disagreement.

The data from the present experiment suggest that individuals entered the experiment with more or less mutually consistent prior expectations about what kinds of agreements would result, and that they updated these expectations in response to their experience in the experiment. These prior expectations, as reflected in the control condition, yielded outcomes closer to 50–50 than to 20–80. It is perhaps for this reason that the outcomes in the 50–50 condition showed so much smaller variance than those in the 20–80 condition.

These results have several implications for the development of the theory of bargaining. The most striking is that it may be necessary to incorporate the expectations of the bargainers into any description (or definition) of equilibrium outcomes, and that there may in general be multiple equilibria supported by different sets of mutually consistent expectations. Some models consistent with this suggestion have already been explored, and the notion that agents' beliefs play a role in determining outcomes is not a new one.[6] However, the results of this experiment suggest that, because of the role which agents' beliefs play in determining the outcome of bargaining, it may be necessary to look to models in which at least some elements of the expectations which players bring to the bargaining will be empirically determined exogenous variables.

[5]The experimental conditions can be interpreted as having changed the *strategic risk posture* of the players, as measured by the certain payoff which they would regard as equivalent to the (risky) opportunity to engage in the bargaining. The notion of strategic risk posture (introduced in Roth, 1977a,b), and studied in the context of bargaining games in Roth (1978, 1979) plays a role parallel to ordinary risk posture in determining an individual's utility for engaging in a game. The role of ordinary risk posture in bargaining models is explored in Roth and Uriel Rothblum (1982).

[6]An interesting paper in which a general definition of equilibrium in games is proposed which explicitly involves certain beliefs of the players is Kreps and Wilson (1982a).

REFERENCES

Aumann, Robert J., "Agreeing to Disagree," *The Annals of Statistics*, November 1976, *4*, 1236–39.

Edgeworth, F. Y., *Mathematical Psychics*, London: Routledge and Kegan Paul, 1881.

Harsanyi, John C., *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*, Cambridge: Cambridge University Press, 1977.

Kreps, David M. et al., "Rational Cooperation in the Finitely-Repeated Prisoners' Dilemma," *Journal of Economic Theory*, August 1982, *27*, 245–52.

______ **and Wilson, Robert**, (1982a) "Sequential Equilibria," *Econometrica*, September

1982, *50*, 863–94.

______ **and** ______, (1982b) "Reputation and Imperfect Information," *Journal of Economic Theory*, August 1982, *27*, 253–79.

Lewis, David K., *Convention: A Philosophical Study*, Cambridge: Harvard University Press, 1969.

Milgrom, Paul, "An Axiomatic Characterization of Common Knowledge," *Econometrica*, January 1981, *49*, 219–21.

______ **and Roberts, John,** "Predation, Reputation, and Entry Deterrence," *Journal of Economic Theory*, August 1982, *27*, 280–312.

Nash, John F., "The Bargaining Problem," *Econometrica*, April 1950, *18*, 155–62.

______, "Non Cooperative Games," *Annals of Mathematics*, September 1951, *54*, 286–95.

Rosenthal, Robert W., "Sequences of Games With Varying Opponents," *Econometrica*, November 1979, *47*, 1353–66.

______ **and Landau, Henry J.,** "A Game-Theoretic Analysis of Bargaining With Reputations," *Journal of Mathematical Psychology*, December 1979, *20*, 233–55.

Roth, Alvin E., (1977a) "The Shapley Value as a von Neumann-Morgenstern Utility Function," *Econometrica*, April 1977, *45*, 657–64.

______, (1977b) "Bargaining Ability, The Utility of Playing a Game, and Models of Coalition Formation," *Journal of Mathematical Psychology*, October 1977, *16*, 153–60.

______, "The Nash Solution and the Utility of Bargaining," *Econometrica*, May 1978, *46*, 587–94.

______, *Axiomatic Models of Bargaining*, Berlin: Springer, 1979.

______ **and Malouf, Michael W. K.,** "Game-Theoretic Models and the Role of Information in Bargaining," *Psychological Review*, November 1979, *86*, 574–94.

______, ______, **and Murnighan, J. Keith,** "Sociological versus Strategic Factors in Bargaining," *Journal of Economic Behavior and Organization*, June 1981, *2*, 153–77.

______ **and Murnighan, J. Keith,** "The Role of Information in Bargaining: An Experimental Study," *Econometrica*, September 1982, *50*, 1123–42.

______ **and Rothblum, Uriel G.,** "Risk Aversion and Nash's Solution for Bargaining Games with Risky Outcomes," *Econometrica*, May 1982, *50*, 639–47.