

Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study

By ALVIN E. ROTH, VESNA PRASNIKAR, MASAHIRO OKUNO-FUJIWARA,
AND SHMUEL ZAMIR*

In an experiment comparing related two-person bargaining and multiperson market environments in Israel, Japan, the United States, and Yugoslavia, market outcomes converged to equilibrium everywhere, and there were no payoff-relevant differences among countries. However, bargaining outcomes were everywhere different from the equilibrium predictions (both in observed agreements and in the substantial frequency of observed disagreements), and substantial differences were observed among countries. Because of the way the experiment was designed, the fact that the market behavior is the same in all countries supports the hypothesis that the differences in bargaining behavior among countries are not due to differences in languages, currencies, or experiments but may tentatively be attributed to cultural differences. (JEL C78, C90, C92)

This paper reports an experiment in which data were collected for a simple one-period bargaining situation (an ultimatum game) and a simple one-period market in four countries: Israel, Japan, the United States, and Yugoslavia. The experiment had three substantive goals: (i) to compare behavior in related bargaining and market environments; (ii) to compare behavior in very different subject pools in order to assess the effect that subject-pool differences may have and to assess how this effect may differ in the bargaining and market environments; and (iii) to use such differences as may be found between subject pools to test and refine hypotheses about the out-of-equilibrium behavior that has frequently been

observed in bargaining games of the kind examined here.

In addition, a major methodological goal of the present investigation was to give us the opportunity to learn from experience how to deal with the formidable problems of experimental design that come to the fore in constructing a multinational experiment, particularly if one of the goals of the experiment is to investigate possible *cultural* differences. These problems include how to control for potential experimental artifacts arising from the different languages in which instructions are given, the different currencies in which subjects are paid, and the different experimenters who conduct the trials in each country. To the extent that these factors can be controlled, different behavior in the different subject pools can cautiously be used as the basis for preliminary conjectures about cultural differences that might account for the different observed behavior.

The two-player bargaining environment we look at is an ultimatum game: one bargainer makes a proposal of how to divide a certain sum of money with another bargainer, who has the opportunity to accept or reject the proposed division. If the second bargainer accepts, each bargainer earns the amount proposed for him by the first bar-

*Roth: Department of Economics, University of Pittsburgh, Pittsburgh, PA 15260; Prasnikar: University of Pittsburgh and University of Ljubljana; Okuno-Fujiwara: University of Tokyo; Zamir: Hebrew University of Jerusalem and University of Pittsburgh.

This work has been partially supported by grants from the National Science Foundation, the Russell Sage Foundation, and the Alfred P. Sloan Foundation. We also thank Roy Navon in Israel, Janez Prasnikar and Lado Rupnik in Yugoslavia, and Koji Ishibashi and Noriyuki Yanagawa in Japan for assistance at critical junctures, and we thank J. Kadane for statistical advice.

gainer, and if the second bargainer rejects, then each bargainer earns zero. To allow us to observe the effects of experience, subjects in the bargaining part of the experiment each participate in ten bargaining sessions against different opponents. Although different pairs of bargainers interact simultaneously, each bargainer learns only the result of his own negotiation.

The multiplayer market environment we examine has a similar structure: multiple buyers (nine, in most sessions) each submit an offer to a single seller to buy an indivisible object worth the same amount to each buyer (and nothing to the seller). The seller has the opportunity to accept or reject the highest price offered. If the seller accepts, then the seller earns the highest price offered, the buyer who made the highest offer (or, in case of ties, a buyer selected by lottery from among those who made the highest offer) receives the difference between the object's value and the price he offered, and all other buyers receive zero. If the seller rejects, then all players receive zero. Each player learns whether a transaction took place and at what price. To allow us to observe the effects of experience, subjects in the market part of the experiment each participate in ten markets, with a changing population of buyers.

In both the market and bargaining environment, the prediction of the unique subgame-perfect equilibrium (under the auxiliary assumption that subjects seek to maximize their monetary payoffs) is that one player will receive all the wealth (or almost all, if payoffs are discrete). To see why this is so, suppose for specificity that the total value of a transaction is \$10 and that offers can be made in units no smaller than \$0.05. In the bargaining game, the assumption of subgame perfectness means that the second bargainer will accept any positive offer, rather than reject it and earn zero. Therefore, at equilibrium, no first bargainer will offer the second more than \$0.05, since even that amount will surely be accepted. Thus, there are two subgame-perfect equilibria: in one of them, the first bargainer offers the second \$0.05, keeping \$9.95 for himself, and the second bargainer

accepts (but would have rejected a proposal in which the first bargainer kept everything for himself). In the second equilibrium, the first bargainer offers zero to the second, keeping everything for himself, and the second bargainer (nevertheless) accepts. These two equilibria become one as the smallest unit of transaction goes to zero, and even with the \$0.05 unit, both equilibria give virtually all the gains from trade to the first bargainer.

The computation of pure-strategy perfect equilibria is almost equally simple in the case of the market game. Here again, the assumption of subgame perfectness means that the seller never rejects the maximum bid when it is positive. Because any buyer who does not submit the maximum bid earns zero with certainty, there cannot be any equilibria at which the high bidder makes a positive profit (by bidding \$9.95 or less) and some other bidder submits a lower bid, since a low bidder could do better by raising his bid to the high bid, which would then give him a positive expected payoff. If the high bid is no greater than \$9.95, all bids must be equal. However, if all bids are equal, they cannot be less than \$9.95, since if they were, then a bidder who raised his bid by \$0.05 would increase his expected payoff because he would win with certainty instead of with probability $\frac{1}{9}$. Thus, the only perfect equilibrium at which the maximum bid is not \$10.00 has all bids equal to \$9.95, so that the seller earns virtually all of the profit. There are also equilibria at which the maximum bid is \$10.00. In fact, *any* distribution of bids in which two or more buyers bid \$10.00 is an equilibrium, since in this case no buyer can earn a positive payoff (even) by changing his bid. Thus, there are many equilibria, but only two equilibrium prices, \$10.00 and \$9.95, so the seller gets (virtually) all the wealth. The situation is the same when we consider perfect equilibria in mixed strategies.¹

¹Consider an equilibrium in which at least one buyer has a positive expected payoff (i.e., in which there is a positive probability that the high bid will be less than \$10.00). Then, at equilibrium every buyer must have a positive expected payoff (since a bid of \$9.95 with

Various experimenters (see Werner Guth and Reinhard Tietz [1990] and Roth [1992] for surveys) have previously noted that observed payoffs in ultimatum bargaining are much less extreme, with the second bargainer reliably receiving significantly more than zero when agreement is reached. In contrast, preliminary evidence from the pilot markets used to assess the viability of the present experiment, which were included in an experiment reported in Prasnikar and Roth (1991),² suggested that, after subjects had acquired a little experience with the market, prices would conform very closely to the perfect-equilibrium prediction that the single seller in each market would receive all the gains from trade. Thus, there was reason to believe that, despite their similar equilibria, these two economic envi-

ronments would yield very different behavior.

The principal patterns of behavior observed in the data from the four countries in which this experiment was run are as follows.

- 1) Regarding the observed market behavior: (a) In every country, the observed market outcomes converge quickly to the perfect equilibrium, and do not deviate from equilibrium once it has been achieved. (In no country was the highest offered price ever rejected in any round, and so the observed outcomes were always Pareto-efficient.) (b) Hence, there are no payoff-relevant differences observed in market behavior between countries.
- 2) Regarding the observed bargaining behavior: (a) In every country, the observed bargaining outcomes are significantly different from the perfect-equilibrium predictions. (Further, in every country, there was a substantial frequency of rejected offers, resulting in Pareto-inefficient outcomes.) (b) However there are substantial differences observed between countries, such as a pronounced shift in the distribution of offers that the first bargainer makes to the second. (The highest offers are made in the United States and Yugoslavia, and the lowest offers are made in Israel, with Japan in the middle. Except for the United States and Yugoslavia, all between-country differences are statistically significant, and between-country differences are bigger than within-country differences among different experimental sessions.) (c) Within every country, the probability that an offer will be rejected is inversely related to the size of the offer (i.e., low offers are rejected more frequently than high offers). However, this pattern does not hold between countries: higher disagreement rates are *not* observed in countries where lower offers are observed. Further, the probability that a given offer is rejected is *lower* in countries where lower offers are observed. (This will allow us to distinguish between

certainty will have a positive expected payoff in this case). Let x be the smallest bid that some buyer makes with positive probability. At equilibrium the bid of x must have a positive expected payoff, so the event that all buyers bid x must have positive probability (since only in this event can a buyer who bids x be the winning bidder). Suppose a buyer changes his mixed strategy by reducing to zero the probability that he bids x and increasing the probability that he bids $x + 0.5$. Then, in the event that all other buyers bid x , he will win with certainty instead of with probability $\frac{1}{9}$, so this increases his expected payoff, provided his expected payoff is positive when he wins with a bid of $x + 0.05$ (i.e., provided x is less than \$9.95). Therefore, in the only equilibrium in which buyers have positive expected payoff, all buyers bid \$9.95 with certainty (and have expected payoff \$0.05/9). As already noted, any strategies are in equilibrium so long as two or more bidders bid \$10.00 with certainty, so even in mixed strategies there are only two equilibrium prices.

Note that, in addition to having extreme perfect-equilibrium distributions, both games have a continuum of other (imperfect) Nash equilibria. Any offer x in the bargaining game can occur at an (imperfect) equilibrium if the second bargainer's strategy is to reject all offers but x , and the same is true in the market game.

Finally, recall that the auxiliary assumption under which all these calculations are made, namely, that subjects are seeking simply to maximize income, has been shown to be questionable in environments such as this (see Jack Ochs and Roth, 1989). We will return to this point when we discuss the results of this experiment.

²See also J. Keith Murnighan and Roth (1980) for an earlier study of this kind of market, conducted under different conditions.

alternative hypotheses about the role that perceptions of fairness may play in this behavior and how it may vary across cultures.)

- 3) Regarding other comparisons: (a) Differences between countries evolved differently in the market and bargaining environments: as subjects gained experience, the between-country differences in market outcomes became smaller (and ultimately vanished), while the differences among bargaining outcomes in different countries grew larger as subjects gained experience. (b) Because of the way our experiment is designed, the pattern of bargaining results and the fact that the observed market behavior is essentially the same in all countries supports the hypothesis that the differences observed in bargaining behavior are not due to differences in languages, currencies, or experimenters but result from other causes.

Most of the conclusions we draw about the susceptibility of these two economic environments to subject-pool differences and the nature of those differences, do not depend on interpreting the causes of observed differences between countries as being cultural in origin. However, many aspects of the design of this experiment were concerned with controlling for the effects of extraneous variables in an experiment conducted in different countries. These and other aspects of the design are described in detail next.

I. Experimental Design

A. Controlling for Between-Country Variables

We first discuss several features of the experimental design which specifically address problems arising from the multinational character of this experiment, namely, the problems of controlling for the effects of different experimenters, different languages, and different currencies. (So far as we are aware, the design issues concerning languages and currencies have not previously

been considered in the manner we propose here.) After discussing these elements of the design, we will return to those features of the design that are particular to the bargaining and market environments that are the focus of this experiment.

Our discussion of those aspects of the experimental design motivated by the multinational character of the experiment will be organized as a statement of a particular problem, followed by the element of the design that addresses this problem.

Problem 1: Experimenter Effects.—Since the experiment involves several experimenters in different locations, between-country differences might arise because of uncontrolled procedural differences or because of uncontrolled personal differences among the experimenters.

Design Solution.—After the procedures were initially designed, each of the experimenters came to Pittsburgh, where they ran (at least) a bargaining session and a market session. The Pittsburgh data were therefore gathered by all of the experimenters before they returned to their home countries to gather the data there.³ In this way, we were able to coordinate the detailed operational procedures among the different experimenters. Also, the Pittsburgh data can be used to detect any pure experimenter effect in the between-country comparisons (i.e., any effect due to personal characteristics of the experimenters), since if these effects exist they will show up not only in the comparisons between countries, but in comparisons of the Pittsburgh sessions conducted by the different experimenters.

Problem 2: Language Effects.—Because the instructions for the experiment are presented in English, Hebrew, Japanese, and Slovenian, systematic differences between countries might be observed because of

³The Yugoslav data were gathered by Prasnikar, who ran the first Pittsburgh sessions with Roth observing. The remaining Pittsburgh data were gathered by Zamir (the Israeli experimenter) and Okuno-Fujiwara (the Japanese experimenter) with Roth and Prasnikar observing.

the way the instructions are translated. (For example, the English words “bargaining,” “negotiating,” and “haggling” are all approximate synonyms but have different connotations which might elicit differences in behavior.)⁴

Design Solutions.—We addressed the problem of language effects both through the way in which the translations were made and, more formally, through the way the instructions for the bargaining and market environments were related. (The instructions in all four languages are available from the first author upon request.)

- 1) Translations: The experimenter responsible for each translation is a national of the country in question who is both linguistically and culturally fluent in American English (all three non-American experimenters had lived for extended periods in the United States). Efforts were made to phrase the English instructions in terms that could be faithfully translated into each of the languages. Aside from avoiding terms with heavy or ambiguous connotations either in English or in translation, this also led to phrasing in less abstract terms than are sometimes used in single-culture experiments. (For example, subjects in bargaining experiments are sometimes instructed that they will be in the position of “player 1” or “player 2,” but this turns out to be difficult to translate into Slovenian without sounding frivolous.)
- 2) Control for translation differences: The instructions for the bargaining and market environments were written in parallel, using the same vocabulary. (For example, in both environments, those subjects who made proposals were referred to as “buyers,” while those who

made acceptances or rejections were termed “sellers.”) If a translation difference is responsible for an observed behavior difference between countries, it should show up in both the market and bargaining data. In particular, the pattern of results that would allow us to be most confident that a between-country difference in bargaining behavior, for example, was not due to translation differences would be if there were no between-country difference observed in the market behavior and if the market and bargaining behavior were also different from each other in each country. That is, suppose we observe a pattern of results of the following sort, in which the data differ between two countries for one of the environments (in this case, the bargaining environment), but not the other:

$$\begin{array}{ccc} \text{market/country 1} & \neq & \text{bargaining/country 1} \\ & = & \neq \\ \text{market/country 2} & \neq & \text{bargaining/country 2} \end{array}$$

If the market data (like the bargaining data) also showed differences between the two countries, or if the market and bargaining data were the same in one of the countries, then we could not be sure that the between-country difference in the bargaining data was not due to some property of the translation. However, if the pattern of results is as above, then we can at least put an upper bound on the effect of the translation: it is not large enough to cause the markets to yield different results in the different countries or to cause the bargaining to yield the same results as the market in one of the countries. This would thus support the hypothesis that the translation is not the cause of the observed difference in the bargaining.⁵

⁴This problem could not have been avoided by presenting the identical instructions in English to English-speaking subjects in each of the countries. Aside from the selection effects of choosing only English-speakers, there is no way to control the different connotations that various English terms and phrases might have to nonnative English-speakers in different countries.

⁵A priori, either environment might have served as a control for the other, but the strong convergence to equilibrium observed in the preliminary trials of the market in the United States made us anticipate that greater between-country differences would be observed in the bargaining. Another common approach for controlling for translation differences in survey research is “back translation,” in which a second translator trans-

Problem 3: Currency Effects.—Because the subjects were paid in dinars, dollars, shekels, and yen, systematic differences between countries might be observed because of the different incentives that the potential payments give to subjects or because of the different numerical scale on which payments are made. (That is, subjects in experiments often tend to choose round numbers [see e.g., Wulf Albers and Gisela Albers, 1983], and these may depend on the units involved, so that subjects proposing prices in dollars might choose different numbers than those dealing in thousands of yen, or hundreds of thousands of dinars.)⁶

Design Solutions.—(a) To assess the extent to which between-country differences might be due to differences in purchasing power, the Pittsburgh data establish a baseline by including sessions in which the potential payoff ranged from \$10 to \$30. In each country, the size of the payoffs was then chosen to give a purchasing power on the high side of \$10. If observed differences between countries fall outside the range of differences due to payoffs observed in Pittsburgh, they are likely to be due to other factors. (b) To control for differences in units, proposed prices in all countries were made in terms of 1,000 tokens, with increments being made in units of five tokens.

We hasten to note that there remain uncontrolled differences between subject pools that might not be regarded as “cultural.” For example, in Israel and Yugoslavia, a much higher percentage of our sample of subjects are army veterans than in the United States or Japan. Therefore, any conclusions about the *causes* of between-country differences have to be circumspect.

B. Other Aspects of the Experimental Design and Procedures

In order that experimental sessions could be easily arranged in multiple locations, a “one-classroom” set of procedures was adopted for each of the two experimental environments. Parallel procedures were used in both environments, so that any observed differences between environments would be attributable to differences in the economic environment, such as the number of buyers or the information available to buyers and sellers.

Subjects who participated in a bargaining session were randomly divided into buyers and sellers and then separated into two rows on opposite sides of the room. After the instructions were read aloud, subjects played a practice round (to verify that everyone understood how to make and respond to proposals) and then ten rounds, changing partners after each round. Buyers made a price proposal by filling out a message form, on which they were identified only by a coded identification number. These message forms were then sorted and distributed to the sellers. The sellers’ responses were returned to the buyers in the same way, so that in any round no buyer knew with which seller he was matched and vice versa, and each bargainer learned only the result of his own negotiation. At the conclusion of the session, one round was chosen at random, and subjects were paid their earnings for that round.⁷

Subjects who participated in a market session were also divided into buyers and sellers but were not separated by category, since there were only two sellers. In each round, the buyers were divided into two markets, A and B. The market in which a buyer was participating in a given round was

lates the instructions back into the original language, so that the two versions can be compared for substantive differences. For our purposes in the present experiment, we felt that this would not be adequate, because the issue was to control for subtle connotations. For a recent market experiment that uses back translation, see Steven Kachelmeier and Mohamed Shehata (1990).

⁶Since the Yugoslav data were collected, a devaluation has reduced currency units by a factor of 10,000.

⁷In all countries except Yugoslavia, where university authorities deemed it inappropriate, subjects were also paid a fixed amount for showing up on time (\$5 in the United States, 10 IS in Israel, and 1,000 yen in Japan, with an additional 500 yen in travel expenses for University of Tokyo students who traveled to Keio University).

indicated on his message form for that round, but no subject knew which (other) buyers were in each market in each round. (The motivations for this arrangement were to prevent the ten rounds from becoming a repeated game among a constant group of buyers and to parallel the bargaining sessions, in which bargaining partners changed each round.) After buyers' proposed prices for a given round were collected, the highest proposed price in each market was posted on the blackboard, together with the identification number of the buyer who had proposed it (or who had been selected at random from among the high proposers). The seller in each market then accepted or rejected this offer. This decision was also posted on the blackboard, and the round ended.

In order that no subject in a market session should know which of the others were sellers and which were buyers, subjects were instructed to fill out forms at each opportunity (so that there would be no point at which only sellers or only buyers were writing). When buyers were recording their proposed prices, sellers were asked to estimate what the high price would be in their market, and while sellers were recording their acceptance or rejection, buyers were asked to estimate the likelihood that the posted high price would be accepted. In order that the procedures in the bargaining sessions would parallel those in the market sessions, buyers and sellers in the bargaining sessions also filled out these additional forms. However, only the offers and acceptances/rejections influenced each subject's payoff from the experiment, and therefore these are the primary data which will be the focus of the analysis.

In each country, a pilot session was conducted with experienced bargainers, recruited from subjects who had completed a bargaining session. All other bargaining and market sessions used subjects who had not previously participated in any other part of this or related experiments. Each session lasted ten rounds, which was announced at the beginning of the session and just prior to the last (tenth) round.

The tenth-round data will therefore be of particular interest, for two reasons. First, it represents the round at which these subjects had acquired the most experience with the game and with the reactions of the other subjects. Second, because the tenth round was the last, the experimental environment gave subjects no incentives extending beyond the play of that round. (In earlier rounds, even though subjects are not engaged in a repeated game with the same players, they may have some incentive to make proposals that will help them gather information about the likely reactions that different prices will elicit, in order to use this information in subsequent rounds.) However, the data from earlier rounds will also be of interest, in order to permit us to investigate the different dynamics by which behavior evolves in the market and bargaining environments in the various subject pools examined.

C. Session Parameters

The basic data in each of the four countries come from three bargaining sessions and two market sessions. In the United States, the value of the object being negotiated or bid for in the bargaining and market sessions was \$10, in Yugoslavia it was 400,000 dinars, in Japan 2,000 yen, and in Israel 20 shekels.⁸ In the United States, an additional session of each of the market and bargaining environments was conducted with a \$30 value, to establish a baseline on the effect of changing the amount of the monetary payoffs.

Recall that all proposed prices were translated into units of 1,000 tokens, with

⁸The Yugoslav data were collected from 14 through 28 December 1989, a period during which there was substantial inflation, and the figure of 400,000 dinars was reached on the basis of a comparison of student wages and a price-index calculation at the beginning of that period. In Israel and Japan, the figures were set on the basis of comparisons of student wages and published figures of purchasing-power parity. In each case, the aim was to choose a figure that would yield a purchasing power slightly on the high side of \$10 in the United States.

the requirement that prices be stated in increments of 5. Thus, in both the market and bargaining environments, there are two perfect-equilibrium prices, differing by 5 tokens. These prices are 995 and 1,000 in the market, and 0 and 5 in the bargaining.⁹ The difference of 5 tokens between the two equilibrium prices is negligible in terms of the payoff to any subject in any of the countries.

Subjects were recruited from the student populations of the University of Pittsburgh, the University of Ljubljana, the Hebrew University, and Keio University and the University of Tokyo. In Pittsburgh, subjects were drawn from undergraduate economics classes and M.B.A. business classes; in Israel, they were drawn from undergraduate first- and second-year economics classes (except for the 16 May 1990 bargaining session, which had a mixture of economics, business and psychology students); in Ljubljana, they were drawn from economics students (the 14 December 1989 bargaining session consisted entirely of first-year students, the other sessions consisted of second-, third-, and fourth-year students mixed together); and in Tokyo, they were drawn from third- and fourth-year economics students at the two universities (mixed together in each session).

In each session, we tried to have 20 subjects, so that there would be nine buyers in each of the two markets operating in each round of each market session and so that in the bargaining sessions each buyer would interact exactly once with each seller. On those occasions when fewer than 20 subjects

reported for a session (in one market session and two bargaining sessions in the United States and in one bargaining session in Japan), the session proceeded with fewer subjects. In the market session, this meant that there were fewer buyers in each market, but in the bargaining session it meant that some buyers and sellers would interact twice in the course of the ten-round session. To prevent such a session from taking on some of the character of repeated play, buyers and sellers were therefore each assigned two coded identification numbers, so that they could not know when they were matched with someone for a second time.

II. Results

A. Market Behavior

Perhaps the single most striking result of this experiment was the remarkably consistent convergence to equilibrium observed in the markets. Recall that two markets, A and B, operated simultaneously in each of the ten rounds of each session. In no nonpractice round¹⁰ was the maximum proposed price ever rejected, and in every session the transaction price in both markets rose to the equilibrium price of either 995 or 1,000. In one session (Israel, 4 April 1990), this double convergence was achieved as early as round 3, and in no session did it occur later than round 7. Furthermore, in no session did the transaction price in either market fall below 995 in any subsequent round.

Since the transaction price in each market is the maximum price offered, the convergence to equilibrium could conceivably be the consequence of one aggressive bidder in each market. However, this is not the case: the markets exhibit a high concentration of bids at or near the equilibrium bids. For example, there is no market session in which fewer than a third of the buyers proposed prices of 995 or 1,000 in round 10 (and in most sessions the percentage is far higher). In every country, over half of the

⁹The decision to use a market with many buyers (i.e., proposers) and a single seller, rather than one with one buyer and many sellers (which would have had equilibrium prices of 0 and 5), was made so as to gather more data on price proposals from each market and in the anticipation that competition might act more forcefully on the active (proposer) side of the market. Prasnikar and Roth (1991) compared bargaining and market games of the kind studied in this paper with an additional two-person game in which the perfect equilibrium gave almost all the wealth to the proposer and in which equilibrium was observed experimentally. The relationships among all three kinds of games will be discussed in Section III.

¹⁰And only once in a practice round.

TABLE 1—THE FIRST TWO HIGHEST PRICES p OFFERED IN EACH OF THE MARKETS AND THE BASIC DESCRIPTIVE STATISTICS

A. Pittsburgh:										
Period	Market	29 June 1989, Prasnikar ($N = 9$)			22 February 1990, Zamir ($N = 7$)			14 March 1990, Okuno-Fujiwara, \$30 ($N = 9$)		
		Highest price p	Second-highest price p	Mean (SD)	Highest price p	Second-highest price p	Mean (SD)	Highest price p	Second-highest price p	Mean (SD)
1	A	800 (1)	775 (1)	575 (198)	800 (1)	700 (2)	586 (165)	900 (1)	885 (1)	772 (152)
	B	900 (3)	800 (1)	581 (338)	805 (1)	800 (1)	629 (153)	850 (1)	805 (1)	664 (151)
2	A	935 (1)	900 (2)	665 (299)	900 (1)	825 (1)	758 (119)	965 (1)	950 (1)	721 (205)
	B	900 (1)	855 (1)	676 (269)	1,000 (1)	900 (1)	635 (262)	930 (1)	880 (1)	684 (252)
3	A	950 (1)	900 (2)	711 (237)	950 (1)	900 (1)	678 (258)	985 (1)	970 (1)	708 (273)
	B	985 (1)	925 (1)	827 (148)	1,000 (1)	950 (1)	778 (213)	970 (1)	950 (1)	753 (171)
4	A	995 (1)	990 (1)	864 (177)	955 (1)	935 (1)	613 (326)	960 (1)	930 (1)	748 (218)
	B	990 (1)	940 (1)	674 (335)	995 (1)	925 (1)	737 (272)	985 (1)	980 (2)	782 (307)
5	A	995 (1)	990 (2)	909 (075)	975 (1)	900 (1)	721 (176)	975 (1)	965 (3)	824 (203)
	B	1,000 (2)	990 (1)	641 (403)	1,000 (1)	995 (1)	925 (146)	995 (1)	990 (1)	799 (193)
6	A	1,000 (1)	950 (1)	581 (376)	995 (1)	990 (1)	822 (216)	990 (2)	970 (1)	879 (161)
	B	995 (4)	980 (1)	868 (325)	925 (1)	750 (1)	539 (298)	995 (1)	970 (1)	782 (199)
7	A	995 (1)	955 (1)	744 (327)	995 (1)	980 (1)	631 (382)	1,000 (1)	995 (2)	858 (151)
	B	995 (4)	990 (1)	834 (325)	995 (1)	955 (1)	700 (298)	1,000 (1)	995 (2)	848 (189)
8	A	995 (3)	990 (1)	758 (377)	995 (1)	990 (1)	796 (221)	995 (2)	960 (1)	811 (191)
	B	1,000 (1)	995 (3)	809 (330)	995 (1)	990 (1)	669 (319)	995 (1)	975 (1)	789 (176)
9	A	1,000 (1)	995 (4)	853 (324)	995 (1)	990 (1)	501 (373)	995 (1)	990 (1)	804 (196)
	B	995 (1)	990 (1)	531 (426)	995 (3)	990 (1)	775 (367)	995 (4)	985 (1)	906 (163)
10	A	1,000 (1)	995 (3)	621 (457)	995 (2)	990 (1)	534 (458)	1,000 (1)	995 (1)	807 (216)
	B	1,000 (1)	995 (2)	741 (342)	1,000 (1)	995 (2)	699 (390)	1,000 (2)	995 (3)	916 (170)

buyers proposed a price of 995 or 1,000 at least once (18 out of 32 buyers in the \$10 sessions in the United States; 31 out of 36 buyers in Yugoslavia; 29 out of 36 buyers in Japan; and 19 out of 36 buyers in Israel).

The market data are summarized in Table 1. For example, looking at Table 1A one sees that in the market conducted in Pitts-

burgh on 29 June 1989, with a transaction value of \$10, there were nine buyers in each market (so there were 20 subjects, two of whom were sellers). In round 1, the high bid in market A was 800, and only one bidder in that market proposed that price, while the second-highest price was 775, also proposed by only one buyer. In the same round, the

TABLE 1—(CONTINUED)

B. Ljubljana:

Period	Market	14 December 1989, Prasnikar ($N = 9$)			28 December 1989, Prasnikar ($N = 9$)		
		Highest price p	Second- highest price p	Mean (SD)	Highest price p	Second- highest price p	Mean (SD)
1	A	875 (2)	840 (1)	730 (113)	920 (1)	890 (2)	795 (97)
	B	825 (1)	820 (1)	664 (158)	870 (1)	835 (1)	736 (89)
2	A	885 (1)	875 (1)	767 (90)	950 (1)	940 (1)	835 (105)
	B	950 (1)	880 (1)	845 (52)	960 (1)	955 (1)	835 (147)
3	A	965 (1)	930 (1)	893 (62)	985 (1)	980 (2)	836 (186)
	B	975 (1)	955 (1)	787 (211)	990 (1)	975 (1)	894 (152)
4	A	1,000 (1)	975 (1)	803 (307)	1,000 (2)	995 (2)	858 (201)
	B	985 (1)	980 (1)	836 (153)	995 (1)	990 (2)	903 (141)
5	A	995 (1)	975 (1)	858 (154)	1,000 (1)	995 (3)	959 (67)
	B	1,000 (1)	995 (1)	833 (316)	1,000 (3)	995 (2)	874 (196)
6	A	995 (3)	990 (1)	839 (319)	1,000 (1)	995 (2)	856 (192)
	B	995 (1)	985 (1)	783 (330)	1,000 (2)	995 (5)	974 (65)
7	A	995 (2)	990 (3)	846 (327)	1,000 (2)	995 (4)	947 (131)
	B	995 (1)	990 (5)	928 (162)	995 (4)	990 (2)	943 (96)
8	A	995 (2)	990 (3)	933 (163)	1,000 (1)	995 (4)	989 (10)
	B	995 (3)	990 (2)	852 (329)	1,000 (2)	995 (5)	995 (4)
9	A	995 (1)	990 (4)	791 (340)	1,000 (1)	995 (5)	994 (3)
	B	995 (3)	990 (4)	988 (8)	995 (6)	990 (2)	991 (7)
10	A	995 (6)	975 (1)	827 (350)	995 (8)	975 (1)	993 (7)
	B	1,000 (1)	995 (7)	941 (165)	1,000 (1)	995 (5)	993 (6)

highest price proposed in market B was 900, and this price was proposed by three different buyers. (Notice that, when proposing their prices, the buyers have no way of knowing which other buyers are in the market with them: a different sorting of buyers could easily have made 900 the transaction price in both markets, or 775 the transaction price in market A.) By round 5, how-

ever, both markets have a transaction price of 995 or 1,000, and the transaction price never drops below 995 in either market in any subsequent round. By round 10, the modal proposed price is 995, with 7 out of 18 buyers proposing prices of 995 or 1,000.

Notice that the pattern is very similar both in the other \$10 Pittsburgh market (Table 1A, 22 February 1990), in which 5

TABLE 1—(CONTINUED)

C. Tokyo:

Period	Market	17 May 1990, Okuno-Fujiwara ($N = 9$)			18 May 1990, Okuno-Fujiwara ($N = 9$)		
		Highest price p	Second- highest price p	Mean (SD)	Highest price p	Second- highest price p	Mean (SD)
1	A	900 (1)	850 (2)	764 (132)	875 (1)	855 (1)	778 (77)
	B	990 (1)	950 (1)	735 (180)	910 (1)	850 (1)	784 (71)
2	A	955 (1)	925 (1)	750 (295)	925 (1)	920 (1)	783 (267)
	B	995 (1)	990 (2)	803 (234)	950 (1)	930 (1)	812 (133)
3	A	980 (1)	975 (1)	849 (134)	965 (1)	950 (2)	802 (286)
	B	965 (1)	950 (2)	694 (392)	950 (1)	910 (1)	812 (175)
4	A	975 (2)	955 (1)	933 (44)	975 (1)	970 (1)	895 (122)
	B	995 (1)	985 (2)	968 (19)	975 (1)	960 (2)	886 (135)
5	A	995 (1)	985 (2)	970 (17)	1,000 (1)	990 (1)	976 (15)
	B	995 (1)	985 (1)	928 (79)	980 (2)	975 (2)	919 (121)
6	A	995 (2)	985 (2)	963 (45)	990 (3)	975 (1)	901 (127)
	B	995 (2)	990 (2)	961 (69)	995 (3)	990 (2)	968 (47)
7	A	995 (1)	990 (2)	858 (215)	1,000 (1)	995 (2)	898 (191)
	B	995 (3)	990 (2)	976 (31)	995 (2)	990 (5)	982 (24)
8	A	1,000 (1)	995 (4)	988 (15)	995 (6)	970 (1)	934 (161)
	B	1,000 (1)	995 (3)	909 (175)	1,000 (2)	995 (4)	984 (32)
9	A	995 (6)	950 (1)	906 (175)	1,000 (1)	995 (5)	967 (80)
	B	995 (6)	990 (2)	967 (81)	1,000 (1)	995 (4)	919 (161)
10	A	995 (7)	990 (1)	958 (110)	1,000 (1)	995 (4)	961 (44)
	B	995 (7)	950 (1)	979 (33)	1,000 (2)	995 (3)	880 (328)

out of 14 buyers propose 995 or 1,000 in round 10, and in the Pittsburgh market in which the value of a transaction was \$30 (Table 1A, 14 March 1990), in which 7 out of 18 buyers propose the equilibrium price. Therefore, the change in the scale of the payoffs makes no important difference, as might be expected in a market in which no

buyer is making more than pennies in any of the final rounds.¹¹

¹¹The pattern was very similar in the pilot market session reported and analyzed in Prasnikar and Roth (1991).

TABLE 1—(CONTINUED)

D. Jerusalem:





Period	Market	4 April 1990, 1st session, Zamir ($N = 9$)			4 April 1990, 2nd session, Zamir ($N = 9$)		
		Highest price p	Second- highest price p	Mean (SD)	Highest price p	Second- highest price p	Mean (SD)
1	A	900 (3)	850 (2)	806 (126)	900 (2)	850 (1)	733 (148)
	B	900 (1)	850 (2)	728 (215)	975 (1)	960 (1)	854 (145)
2	A	950 (1)	930 (1)	872 (66)	1,000 (1)	990 (1)	877 (102)
	B	995 (1)	950 (1)	862 (89)	985 (1)	975 (1)	913 (81)
3	A	950 (2)	900 (3)	794 (267)	995 (1)	990 (2)	942 (65)
	B	950 (1)	925 (1)	783 (178)	995 (1)	990 (3)	899 (159)
4	A	965 (1)	950 (2)	846 (145)	1,000 (1)	995 (2)	924 (80)
	B	960 (2)	955 (1)	875 (131)	995 (2)	990 (2)	891 (195)
5	A	960 (2)	955 (2)	931 (39)	1,000 (1)	995 (2)	877 (164)
	B	980 (2)	970 (2)	874 (189)	995 (5)	990 (1)	986 (16)
6	A	1,000 (1)	990 (2)	959 (37)	1,000 (1)	995 (4)	937 (84)
	B	1,000 (1)	980 (1)	919 (74)	995 (4)	975 (1)	945 (67)
7	A	995 (1)	990 (1)	806 (317)	1,000 (2)	995 (2)	920 (161)
	B	995 (1)	990 (1)	911 (157)	995 (4)	985 (1)	802 (310)
8	A	995 (3)	990 (2)	978 (20)	1,000 (2)	995 (2)	843 (305)
	B	995 (1)	980 (1)	869 (185)	1,000 (1)	995 (5)	926 (197)
9	A	1,000 (2)	995 (1)	968 (64)	1,000 (1)	995 (5)	835 (337)
	B	995 (1)	990 (2)	810 (342)	995 (3)	985 (1)	847 (212)
10	A	995 (2)	990 (2)	820 (329)	1,000 (1)	995 (2)	769 (339)
	B	1,000 (1)	995 (3)	964 (41)	1,000 (1)	995 (5)	903 (149)

Notes: The numbers in parentheses under the prices are the numbers of buyers who offered that price. N represents the number of buyers in each market (i.e., in each round there are two sellers and $2N$ buyers, half in market A and half in market B).

Similarly, parts B–D of Table 1 show that by the tenth round the two Yugoslav markets both have 14 out of 18 buyers proposing 995 or 1,000, the two Japanese markets have 14 out of 18 and 10 out of 18 buyers proposing 995 or 1,000, and the two Israeli

markets have 6 out of 18 and 9 out of 18 buyers proposing 995 or 1,000.

Figure 1 shows the distributions and cumulative distributions of proposed prices for rounds 1 and 10 for the two U.S. markets in which the transaction value was \$10 and for



Prasnikar 6/29/89, Zamir 2/22/90

Okuno 3/14/90

FIGURE 1. DISTRIBUTION OF MARKET OFFERS IN THE UNITED STATES




FIGURE 2. DISTRIBUTIONS OF MARKET OFFERS IN YUGOSLAVIA, JAPAN, AND ISRAEL

the one in which it was \$30. (The histogram bars group price intervals of 25. Since price proposals are in increments of five tokens, this means that an interior bar on the graph, such as 500, contains the proposals from 490 through 510, while an endpoint, such as 1,000, contains the proposals of 990, 995, and 1,000.) Figure 2 contains the distributions and cumulative distributions of proposed prices for rounds 1 and 10 in Yugoslavia, Japan, and Israel. The pattern of how proposed prices changed over time is the same for all countries (and for \$10 and \$30 transaction values in the United States). In round 1, no more than 14 percent of the proposals are higher than 910 in any country, and except for a single offer in Japan, there are no proposals in the 990–1,000 range. However, in round 10, at least 39 percent of the proposals are from 990 to 1,000 in every country (and the concentration of proposals in this highest range goes up to over 80 percent). Thus, the shift in proposed prices from the first to tenth rounds in each country is clear.

While there are some detectable differences between the distributions in different countries,¹² we attach little significance to them for two reasons. First, none of these differences influences the payoffs to any agent: in the final rounds of the market sessions, any buyer who proposes a price of less than 900, say, has every reason to expect (correctly) that his earnings will be zero for that round. For this reason, economic theory makes no prediction about what the distribution of low offers will look like. Indeed, recall that, at any equilibrium in which two or more buyers propose a price of 1,000, no further prediction can be made about the price proposals of any other buyers, since these neither influence any payoff nor move the market out of equilibrium.

What Figures 1 and 2 and Table 1 make clear is that, while there is some variation in

transaction prices in round 1, this has disappeared by round 10, when all transactions in all markets in all countries are at equilibrium prices.

B. Bargaining Behavior

The contrast between the bargaining data and the market data is striking: in the bargaining sessions, equilibrium price proposals (of 0 or 5) make up less than 1 percent of the data from any country. In all countries the price proposals made by bargainers are much nearer the middle of the range, and in all countries low offers are rejected at a substantially higher rate than higher offers.

However, despite the gross similarity of the bargaining data from all countries, especially when compared to the market data, there are notable differences between the distributions of bargaining proposals in different countries. The most obvious of these is seen in the different modal proposals (aggregated over all rounds of bargaining). In the United States (for bargaining both over \$10 and \$30) and in Yugoslavia the modal proposal is 500, while in Japan and Israel the modal proposal is 400.

To test formally whether these differences are reliable, we must disaggregate the data by rounds. The reason is that data from different rounds of the same bargaining session are not independent, since the same bargainers are involved (even though they are not paired in the same way). However, the proposals made in a given round by bargainers in different experimental sessions are independent, and so we can look at all the data for each country, round by round.

Before describing the formal tests, we first consider these distributions graphically. Figure 3 presents the round-1 and round-10 distributions and cumulative distributions for the \$10 and \$30 bargaining sessions in the United States, and Figure 4 presents the round-1 and round-10 distributions and cumulative distributions of offers in Yugoslavia, Japan, and Israel. In addition, the figures show the proportion of offers in each interval that were accepted (the black

¹²The Mann-Whitney *U* test reveals significant differences in most comparisons of the distributions, but all of the distributions are very highly concentrated in the highest prices.

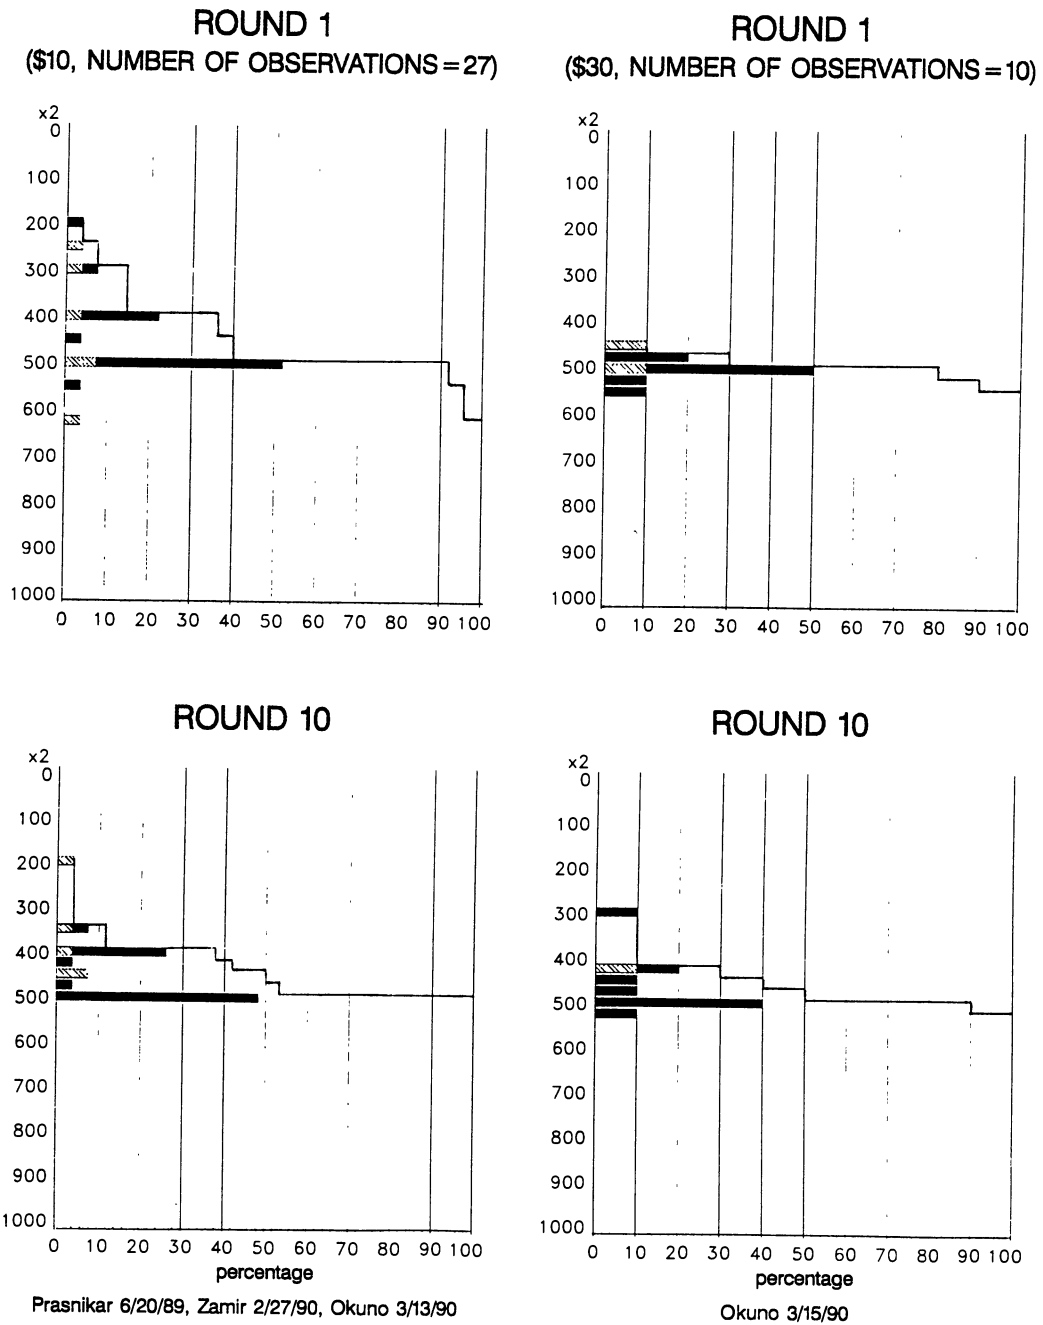


FIGURE 3. DISTRIBUTION OF BARGAINING OFFERS IN THE UNITED STATES (SOLID BARS = ACCEPTED OFFERS; STRIPED BARS = REJECTED OFFERS)




FIGURE 4. DISTRIBUTIONS OF BARGAINING OFFERS IN YUGOSLAVIA, JAPAN, AND ISRAEL (SOLID BARS = ACCEPTED OFFERS; STRIPED BARS = REJECTED OFFERS)

region) and rejected (the lighter, striped region).

In choosing the appropriate statistical test to compare the distributions between countries, we had to take into account the following features of these empirical distributions:

- 1) The distributions are highly asymmetric, with few observations higher than the mode. The distributions are clearly not normal and in fact fail the Kolmogorov-Smirnov test for normality.¹³
- 2) The sample sizes we are dealing with are relatively small. In each session involving 20 bargainers, there are 100 offers made by the 10 buyers in 10 rounds. However, because of the dependence between offers made by the same buyer in different rounds, we compare data round by round, which means that we base our tests on samples of size 10 when comparing two sessions and on samples of size 30 when we compare two countries with three sessions in each.

Guided by these considerations we used the nonparametric Mann-Whitney *U* test based on ranks (see e.g., Sidney Siegel, 1956). The idea of the test is as follows. Assume we are comparing two samples from variables *X* (offers in population 1) and *Y* (offers in population 2). Line up all the observations of both variables, from smallest to largest. The *U* statistic is then the number of times an *X* observation precedes a *Y* observation (with some adjustment for ties). If *U* is very small, this should be an indication that the distribution of *X* is "higher" than the distribution of *Y*. The distribution of *U* is known when *X* and *Y* have the same distribution, and therefore we can test the null hypothesis that *X* and *Y* have the same

distribution against the alternative hypothesis that the two distributions are different.

The *U* statistics also provide an estimate of the magnitude of the difference between the two groups. A measure for this difference is $P(X > Y)$; that is, the probability that a random observation from *X* will be higher than a random observation from *Y*. If *X* and *Y* have the same distribution, then clearly $P(X > Y) = P(Y > X)$, and the higher the distribution of *X* compared to that of *Y*, the higher will be $P(X > Y)$ minus $P(Y > X)$. With the *U* statistics and a routine counting of the number of ties in the two samples of offers, we estimated $P(X > Y)$, $P(Y > X)$, and $P(X = Y)$. These estimates have the following probabilistic interpretation: assume that we draw two random offers *x* and *y* from the given samples of offers in group 1 and 2, respectively; our estimates are then the probabilities $P(x > y)$, $P(y > x)$, and $P(x = y)$, respectively.

Table 2 presents the results of the comparisons of the round-10 bargaining propos-

TABLE 2—MANN-WHITNEY TEST BETWEEN COUNTRIES: POOLED ROUND-10 BARGAINING

Country	Country		
	United States	Yugoslavia	Japan
Yugoslavia	0.35 0.45 (0.51)		
Japan	0.26 0.62 (0.02)*	0.29 0.59 (0.04)*	
Israel	0.10 0.79 (0.00)*	0.11 0.81 (0.00)*	0.21 0.69 (0.00)*

Notes: In each cell, the first number is the probability that an offer from the row country exceeds an offer from the column country. The second number is the probability that an offer from the column country exceeds an offer from the row country. The third number (in parentheses) is the significance level for the two-sided test of the null hypothesis that the offers are drawn from identical distributions (i.e., the probability that a difference as extreme as the observed difference would arise from two identical distributions). Cells in which this probability is less than 0.05 are marked with an asterisk.

¹³In each of the four countries, the pooled sample of all offers failed the Kolmogorov-Smirnov test at the 0.01 significance level. The same test by round in each country rejects the normal-distribution hypothesis in 27 out of 40 cases (the exceptions are eight rounds in Israel and five in Japan, and even there the maximum significance level was less than 0.4).

TABLE 3—MANN-WHITNEY *U* TEST, BY SESSION: NUMBER OF DIFFERENCES AT SIGNIFICANCE LEVEL 0.05 FOR ROUND-10 BARGAINING

Country	Country			
	United States	Yugoslavia	Japan	Israel
United States	0/6			
Yugoslavia	3/12	2/3		
Japan	5/12	3/9	0/3	
Israel	11/12	6/9	6/9	0/3

Notes: Each diagonal cell indicates how many significant differences were found when the round-10 bargaining data were compared for each pair of sessions in that country. Each off-diagonal cell indicates how many significant differences were found when the round-10 data from each session in the row country were compared with the round-10 data from each session in the column country. Since there are four sessions in the United States (one for each experimenter plus a \$30 session) and three in each of the other countries, there are six pairwise comparisons within the United States and three within each other country, and there are 12 pairwise comparisons between the United States and each other country and nine pairwise comparisons between each other pair of countries.

als in all countries. The observed distributions are significantly different between every pair of countries except the United States and Yugoslavia.

This raises the question of whether these differences can be accounted for by currency or experimenter effects. As discussed earlier, tests for this possibility were built into the experimental design by having the three \$10 bargaining sessions in Pittsburgh run by the three different experimenters and by having a fourth, \$30, Pittsburgh bargaining session. If the distributions of offers are responding to changes in the value of a transaction, this should show up as a difference between the \$30 bargaining session and the other sessions, while if there is an effect due to one of the experimenters, this should show up in comparisons involving the Pittsburgh session run by that experimenter. Contrary to either of these hypotheses, the first cell of Table 3 reports that there are no significant differences among the six pairwise comparisons of these four sessions.

When we compare the data for a given round from different sessions in this way, we are of course looking at smaller samples than when we pooled the data for a given round from all sessions in each country, as in Table 2. Therefore, one issue we need to

consider is whether the lack of significant differences in the Pittsburgh sessions, in contrast to the significant differences between countries, might be due to the small sample sizes. Table 3 addresses this by reporting the results of pairwise comparisons between the unpooled round-10 bargaining data from all sessions in all countries. Note first that, except in Yugoslavia where one of the sessions was significantly different from the other two, there are no significant differences between sessions from the same country. The one Yugoslav session that differs from the others also differs from three of the U.S. sessions. Otherwise, the between-country differences by sessions mirror the between-country differences pooled across sessions.

So far we have concentrated on the round-10 data. Looking back at the comparison of rounds 1 and 10 given in Figures 3 and 4, we see that the differences between countries appear to increase from round 1 to round 10: the modal offer in round 1 is 500 in every country, whereas by round 10 the modal offer has shifted to 400 in Israel and Japan (which has a second mode of 450), while in the United States and Yugoslavia the round-10 mode remains at 500. Table 4 confirms this impression and shows that the differences between the distribu-

TABLE 4—PROBABILITY MEASURE Δp OF DIFFERENCE BETWEEN COUNTRIES' BARGAINING, BY ROUNDS

Round	US-YU	US-JA	US-IS	YU-JA	YU-IS	JA-IS
1	0.099	0.051	0.365	-0.038	0.271	0.264
2	0.042	0.330	0.446	0.313	0.429	0.076
3	-0.217	-0.022	0.221	0.139	0.438	0.192
4	-0.153	-0.031	0.178	0.122	0.329	0.207
5	0.132	0.169	0.505	0.015	0.422	0.364
6	0.332	0.301	0.575	-0.014	0.364	0.364
7	0.091	0.243	0.530	0.160	0.486	0.368
8	0.133	0.321	0.604	0.218	0.527	0.331
9	-0.054	0.161	0.496	0.259	0.590	0.377
10	0.099	0.360	0.695	0.302	0.703	0.479

Notes: US = United States, YU = Yugoslavia, JA = Japan, and IS = Israel.

tions of proposals by bargainers in different countries, as measured by $\Delta p = P\{x > y\} - P\{y > x\}$, tend to *increase* as the bargainers gain more experience. Although this increase is not monotonic, the difference in the last round is greater than the difference in the first round in every comparison, except between the United States and Yugoslavia, where it is constant (which is consistent with our general finding of no significant differences between those two countries' data).

This is confirmed by a simple regression model $\Delta p = \beta_0 + \beta n + \varepsilon$ which gives an estimate of β that is positive in all six comparisons. In the test of the hypothesis $\beta = 0$ against $\beta > 0$, the three comparisons involv-

ing Israel are significant at conventional levels (0.001 in two cases and 0.01 in one), and the comparisons of Japan with the United States and Yugoslavia have significance levels of 0.055 and 0.095, respectively. For the United States and Yugoslavia, β is not significantly different from zero ($P = 0.26$). (All tests passed the Kolmogorov-Smirnov test for the normality of the residuals.)

Another perspective on the differences between bargaining behavior in different countries comes from looking not at proposed prices, but at acceptances and rejections. We first look at overall disagreement rates in each country (i.e., the percentage of all offers that are rejected, without conditioning on the offer). Table 5 presents these

TABLE 5—REJECTION FREQUENCIES IN BARGAINING, BY ROUND AND COUNTRY

Round	United States		Yugoslavia		Japan		Israel	
	Frequency	Percentage	Frequency	Percentage	Frequency	Percentage	Frequency	Percentage
1	6	22	8	27	7	24	8	27
2	7	26	6	20	11	38	6	20
3	10	37	7	23	10	34	8	27
4	12	44	9	30	10	34	8	27
5	12	44	10	33	9	31	7	23
6	7	26	10	33	10	34	8	27
7	7	26	11	37	4	14	9	30
8	3	11	10	33	9	31	6	20
9	7	26	9	30	8	28	3	10
10	5	19	7	23	4	14	4	13
Total:	76	28	87	29	67	22	83	28

figures for rounds 1–10. The clear pattern within each country, that higher offers are accepted more frequently than lower offers (see Fig. 5),¹⁴ is not mirrored at all when we compare countries where offers are high (the United States and Yugoslavia) with those where they are low (Japan and Israel). Over all rounds, the disagreement rates, which are 28 percent, 29 percent, 22 percent, and 28 percent, respectively, are certainly not rising. At round 10, the disagreement rates of 19 percent, 23 percent, 14 percent, and 13 percent are actually *lower* for the two low-offer countries. (However, we can only speculate whether the relationship among these last-round rates is robust, since the disagreement rates fluctuate so widely between rounds.)

A more detailed comparison of acceptances and rejections between countries can be made by considering how often the proposal of a given price is accepted. These comparisons are slightly complicated by the facts that the number of proposals of a given price is different in different countries and that observed rates of acceptance fluctuate widely for offers that were observed only rarely.¹⁵ However, the underlying pattern is clear, as is demonstrated by Figure 5. The curves for each country represent the percentage of acceptances for each price that was proposed at least 10 times (over all rounds). Each cell of Figure 5 compares the resulting curves for a pair of countries, and these comparisons mirror those concerning the distribution of proposals. In each case, the country with the lower distribution of offered prices has a higher rate of acceptance for each proposed price. Thus, we see that the acceptance rate in Israel for each

offer is higher than the corresponding rates in the United States, Yugoslavia, and Japan, while the acceptance rates in Japan are higher than those in the United States and Yugoslavia. Only in the comparison of the United States and Yugoslavia do we have two acceptance-rate curves such that the one that begins consistently lower ends consistently higher.

Given that different offers are accepted with different probabilities, it is natural to ask, for each country, what is the expected payoff to a buyer from making a particular offer. Since the behavior of the bargainers changes from round to round, this is something of a moving target. Nonetheless, Figure 6 presents the curves based on the pooled data from all rounds in each country for all offers that were made at least ten times. Thus, for example, if a buyer proposes a price of 300, he will earn 700 if it is accepted and 0 if it is rejected. In the United States, the price 300 was proposed 15 times and accepted four times (26.7 percent), so on average the proposal earned $(700 \times 0.267) = 186.9$, which can be read from the graph for the United States in Figure 6. It is instructive to compare these graphs to the modal offers observed in round 10 in each country (in Figs. 3 and 4). The modal offer in the final round in both the United States and Yugoslavia is 500, which is also the proposed price that maximizes a buyer's average earnings in these countries. The modal offer in the final round in Israel is 400; here too, this is the price that maximizes a buyer's average earnings. Finally, in Japan there are two modal offers in round 10, 400 and 450, and the latter maximizes a buyer's average earnings. Thus, by round 10, the buyers seem to be adapting to the experience of the prior rounds in a manner roughly consistent with simple income-maximization. (The same cannot be said of the sellers, who continue to reject low positive offers.) Of course, since we have observed that subjects are changing their behavior as they gain experience, the round-10 offers may not meet with the same average responses as in the earlier rounds, and to the extent that this is the case, there is reason

¹⁴The increasing acceptance rate of higher offers within each country is not completely monotonic, and in fact a small "kink" is visible in each country's acceptance-rate curve, which may possibly indicate a small bias related to whether offers are made in round numbers of tokens.

¹⁵For example, if a particular price is proposed only once, the percentage of acceptances will be either 100 percent or 0 percent.

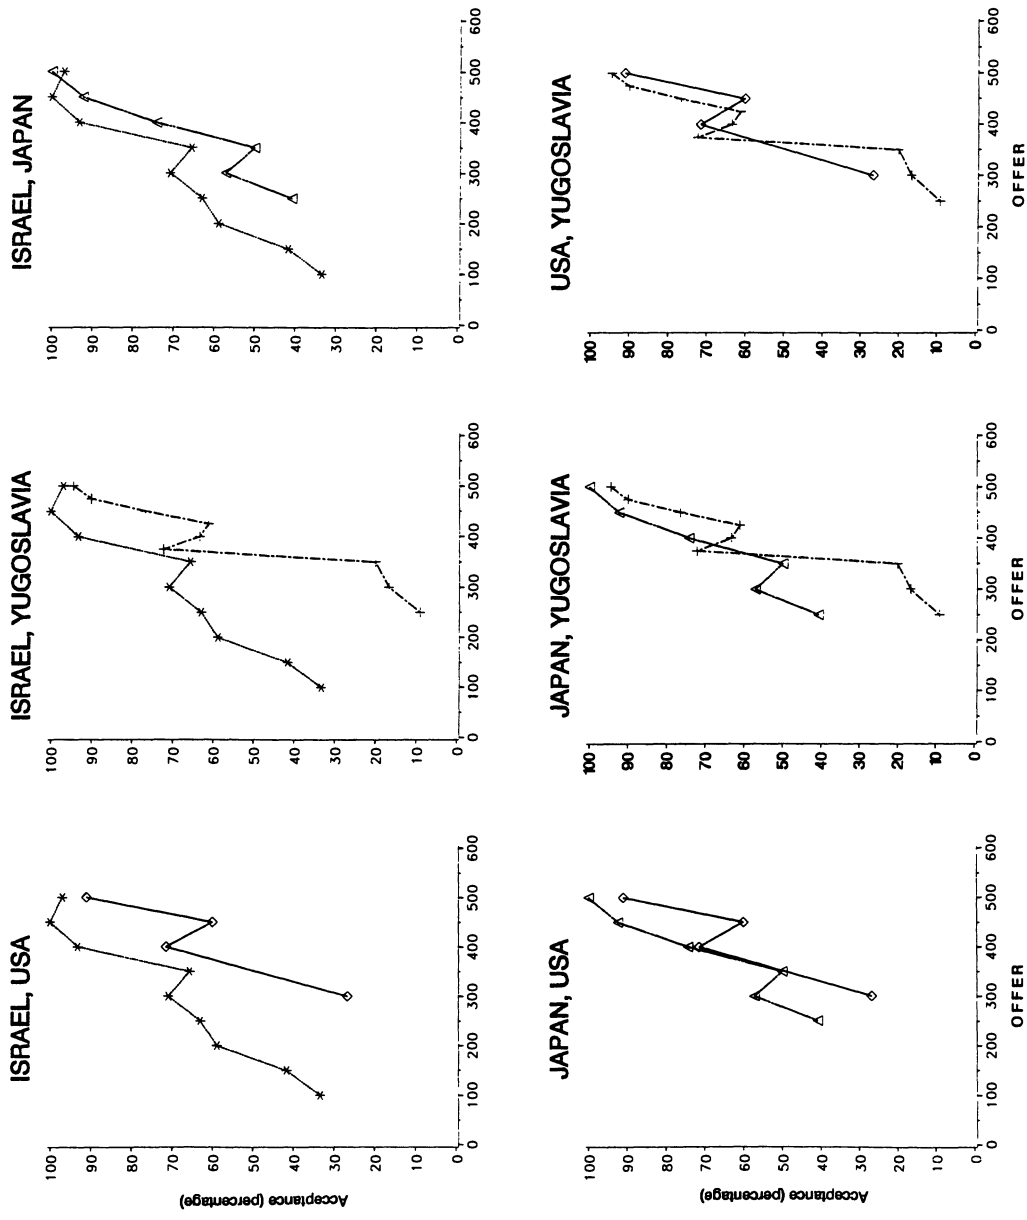


FIGURE 5. PAIRWISE COMPARISONS OF ACCEPTANCE RATES IN BARGAINING SESSIONS (POOLED DATA OF 10 ROUNDS; * = ISRAEL, \diamond = UNITED STATES, + = YUGOSLAVIA, Δ = JAPAN)

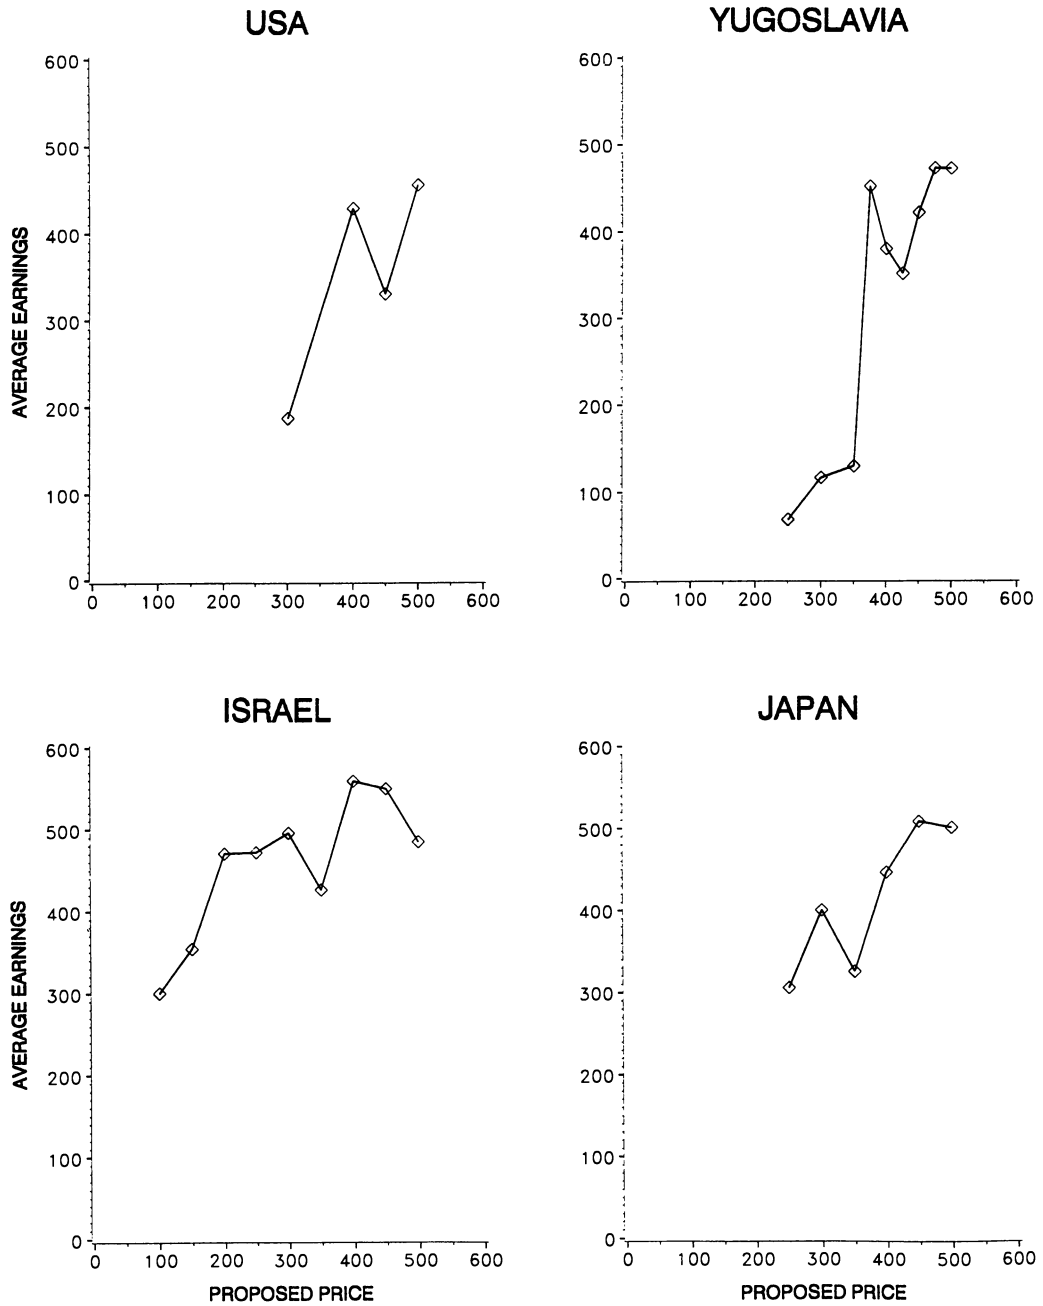


FIGURE 6. BUYERS' EARNINGS IN BARGAINING SESSIONS, BY PROPOSED PRICES (POOLED DATA OF 10 ROUNDS)

to believe that the process has not yet converged by round 10.¹⁶

III. Discussion

Both the market and bargaining environments chosen for this experiment have extreme perfect-equilibrium predictions, in which one player receives all the benefit from the transaction. Nevertheless, the market sessions exhibited a vigorous convergence to equilibrium that was robust to subject-pool differences and transaction values.¹⁷

¹⁶In an ideal world, we would at this point compare the experienced-bargainer sessions we conducted in each country. Instead, we have a cautionary tale about the difficulty of maintaining common procedures (and consequently experimental control) in an experiment conducted in four countries. Despite our efforts to remain in almost constant contact by electronic mail when the experimental sessions were in progress, in this aspect of the experiment our coordination failed. In Yugoslavia, the experienced subjects had all taken part in the same previous bargaining session (which turns out to have been the one Yugoslav bargaining session whose results were different from the other two), while in the other three countries, experienced bargainers were recruited from all three prior sessions. Furthermore, in Yugoslavia and the United States, the final experienced bargaining encounters occurred in the 20th round, while in Israel and Japan the sessions ended after the 10th round (with a subsequent experienced session in Japan lasting for 20 rounds). These procedural inconsistencies, together with the fact that we have many fewer experienced sessions than inexperienced sessions, make the comparisons of these groups less reliable than we would like. Consequently, we will simply note that the pooled data from all rounds of the experienced bargaining in each country are consistent with the between-country differences we have reported above. The experienced data also suggest that the evolution of behavior noted above from round 1 through round 10 may continue to evolve. However, a further exploration of these issues must await more data.

¹⁷The strength of the forces pushing buyers to propose prices that gave them zero profit can perhaps be better appreciated by noting that buyers in the market sessions appeared to find the experience quite frustrating. It was not uncommon, when subjects were being individually paid at the end of each session, for someone to ask why all the other buyers had behaved so foolishly. (When asked about his own bidding behavior, such a buyer would typically respond that the actions of the others had forced him to bid high as well.)

In contrast to the market sessions, the bargaining sessions did not show any tendency toward the equilibrium prediction in any of the subject pools. Furthermore, there were clear differences in the outcomes of bargaining in different subject pools. These differences cannot be attributed simply to variations among inexperienced subjects, since they grew larger from round 1 to round 10, as the bargainers gained experience with the game and with each other.

The out-of-equilibrium behavior in the bargaining game is consistent with the behavior that has been uniformly observed by experimenters who have looked at ultimatum-bargaining games of this kind, starting with the paper by Guth et al. (1982). Prominent in the discussion of this phenomenon has been the idea that bargainers' conceptions of fairness might be an important explanatory variable, particularly insofar as such conceptions might explain the propensity of bargainers in the position of the second mover to reject positive offers (see Guth and Tietz [1990], Prasnikar and Roth [1991], and Roth [1992] for discussions of this literature). The relationship observed in this experiment between offers and acceptance rates in different subject pools can help distinguish between alternative hypotheses about how ideas about the fairness (or "reasonableness") of different proposals might account for these subject-pool differences.

One hypothesis is that the different subject pools share a common idea about what constitutes a fair or reasonable proposal (an obvious candidate is the fifty-fifty proposal of 500) and that the difference among subject pools is in something like their aggressiveness or "toughness." In this view, buyers in more aggressive subject pools would be more inclined to take advantage of their first-mover position to try to obtain more for themselves than might be considered fair. That is, such a buyer would recognize that a fifty-fifty split is "fair," but would seek to take more. However, if aggressiveness is a property of the subject pool, the sellers would share it and would presumably be less inclined to accept unfair offers than

less aggressive sellers in other subject pools. Under this hypothesis, high rates of disagreement would be associated with subject pools in which offers are low. This is not what we observe.¹⁸

Instead, the subject pools where offers are low (Japan and Israel) do not exhibit any higher rates of disagreement than the high-offer subject pools. This suggests that what varies between subject pools is not a property like aggressiveness or toughness, but rather the perception of what constitutes a reasonable offer under the circumstances. That is, suppose that in all subject pools it seems reasonable for the first mover to ask for more than half the profit from the transaction and that what varies between subject pools is how much more seems reasonable. To the extent that offers tend toward what is commonly regarded as reasonable, and assuming that offers regarded as reasonable are accepted, there would be no reason to expect disagreement rates to vary between subject pools, even when offers do. Our data thus lend some support to the hypothesis that the subject-pool differences observed in this experiment are related to different expectations about what constitutes an acceptable offer, rather than different propensities to trespass on a shared notion of what constitutes such an offer.

This brings us to the question of whether such differences can be attributed to *cultural* differences between subject pools, where for an operational definition of "cultural" we mean differences that cannot be attributed to variables other than the nation in which the data were gathered. As we have already indicated, there are uncontrolled differences in subject pools (such as differences in military service) that must make any such attribution speculative. However, the experiment was designed to control for those variables that seemed to

us to be potentially most troublesome, namely, currency, experimenter, and language effects. To the extent that the experimental controls were adequate, the results indicate that the subject-pool differences cannot be attributed to any of these variables.¹⁹ Consequently, we offer the conjecture that the observed subject-pool differences are cultural in character. Such a conjecture must stand or fall on the repeatability and robustness of these results and on the extent to which similar differences among these countries can be observed in related economic environments. In this connection, and in view of the difficulty of controlling for between-country variables, laboratory experimentation seems to us to offer the possibility of focusing on some kinds of cultural differences in behavior that cannot be studied in any other way.²⁰

Finally, we consider what implications the results of this experiment have for the ongoing assessment of the extent to which dif-

¹⁹Recall, however, that the control for translation differences built into our design provides only an upper bound on how great an effect might be due to linguistic factors. In particular, if the differences observed in bargaining behavior are due to translation differences, they must be due to differences too small to have had an effect on the market behavior. Given the robust convergence to equilibrium observed in the market, one might conjecture that it would take large translation differences indeed to affect market behavior, and in this case there is a possibility that the observed bargaining differences might nevertheless be due to differences in the translations, rather than to subject-pool differences. While such a possibility is not entirely ruled out by the data, we are skeptical that the bargaining differences are primarily linguistic in origin.

²⁰At the same time, to the extent that experiments control for extraneous variables by eliminating much of the natural context in which negotiations may take place, there are aspects of cultural differences in bargaining behavior that cannot be studied in the laboratory. For example, differences such as how and when negotiations begin and end or how disagreement is expressed may involve important cultural differences that can only be observed in the natural context of negotiations; and of course, critical features of a negotiating environment, such as the legal framework in which negotiations take place, are important between-country variables that influence the outcome of bargaining in natural contexts but are deliberately excluded here.

¹⁸Another way to make more or less the same point about the "toughness hypothesis" is to note that we are sometimes asked in which country the bargainers proved to be the toughest. Our data suggest that this is not a well-posed question, in the sense that the "toughest" buyers are found in the same place as the least tough sellers.

ferent game-theoretic predictions may be descriptive of observed behavior. Diverse opinions on this subject have been expressed in the experimental literature.

Our evidence lends little support to the view that perfect-equilibrium predictions are not at all descriptive of observable behavior or to the view that they are only descriptive when they are not extreme.²¹ Equilibrium prices are clearly reached in the markets we study, even though the equilibrium is so extreme that the buyers who find themselves proposing these prices earn nothing or next to nothing.

At the same time, the failure of observed behavior in the bargaining games even to approach the equilibrium prediction (and in particular the readiness of sellers in that game to earn zero by rejecting offers that would give them positive earnings) raises questions about the auxiliary assumption under which the equilibrium predictions were made, namely, that the players are attempting to maximize their earnings.²² However, if players are not attempting to maximize their earnings, then why do the equilibrium predictions made under that assumption for the market games do so well? Preliminary discussions with various investigators in this area suggest at least two possible explanations. One is that the observed bargaining behavior is dominated by concerns about fairness which are context-dependent and do not arise in the market environment. Another is that whatever non-monetary concerns enter bargainers' preferences do so in both environments, but the competitive pressure toward equilibrium in

the market overwhelms any such factors in players' preferences.²³

Whether or not nonmonetary factors play a role in either or both environments, the results of this experiment lend strong support to the hypothesis that the different outcomes observed in these two environments result from different behavior *away* from the equilibrium. This helps explain the relation between the equilibrium predictions and the observed bargaining and market behavior. To see why this is so, we need to compare these two games once again, both from the point of view of the equilibrium predictions and the observed behavior.

From the point of view of the equilibrium predictions, the two games are similar in that both predictions give one player 0, but they are dissimilar in that it is the buyer in the market game who is predicted to get 0, while the buyer in the bargaining game is predicted to get 1,000. This dissimilarity largely disappears when we look at observed behavior. In the market sessions, a buyer who proposes the equilibrium price certainly will earn 0; but a buyer who proposes the equilibrium price in the bargaining will earn 0 with very high probability, because

²¹See, for example, Matthew Spiegel et al. (1990), who review evidence from sequential bargaining games from a variety of subject pools, for a recent suggestion that the degree of inequality in the equilibrium payoff division may be the decisive element in determining its descriptive ability.

²²See Ochs and Roth (1989) for a discussion of the consistency of this kind of disadvantageous rejection among a number of bargaining experiments; see Gary Bolton (1991) for a model in which a player's utility depends both on absolute and relative earnings and for a carefully conducted series of experimental tests of that model.

²³These competitive pressures need not be due to simple income-maximization. For example, consider a hypothetical buyer whose preference for equality is such that his first-choice outcome would be to have all buyers submit identical bids of \$5 (or \$1) and who bids accordingly in the first two rounds. When he sees how high the actual transaction price is, he becomes annoyed with the other buyers, and (with the same motivation that would have caused him to express his displeasure by rejecting too small an offer if he were a seller in the ultimatum game) he decides to become the high bidder in round 3 in order to deprive other buyers of the benefits of what he sees as their unreasonable behavior. The point in considering such a hypothetical buyer is to observe that in the market game his nonmonetary preferences cause him to behave in a manner indistinguishable from an income-maximizer, while in the ultimatum game his preferences lead away from the equilibrium predicted for income-maximizers. The difference lies not in the preferences, or in the "social norms" elicited by the game which these preferences may reflect, but in how such preferences interact in the different games and in the outcome that emerges. (The above example is from Prasnikar and Roth [1991].)

low offers were accepted only with low probability (Figs. 5, 6). Despite the similarity of actual payoffs at the equilibrium offer, the payoffs away from the equilibrium proposals are quite different. In the market sessions, a buyer who consistently proposes a price of 500, say, will earn 0 in every round; but a buyer who consistently proposes a price of 500 in the bargaining sessions will with high probability reach an agreement in every round and will receive \$5 (or \$15, 10 shekels, 1,000 yen, or 200,000 dinars, depending on which session he is in). Thus, buyers in both environments earn little or nothing when they make equilibrium proposals, but in the markets we observed, buyers also earned nothing when they made nonequilibrium proposals, whereas buyers in the bargaining sessions could maximize their earnings by moving substantially away from the equilibrium proposal (see Fig. 6).

This conclusion is similar to that reached in Prasnikar and Roth (1991), on the basis of comparisons between some of the U.S. bargaining data considered here with another two-player game in which the first player was (also) a proposer and the second an acceptor/rejecter.²⁴ In that game, as in the bargaining game, the equilibrium prediction gave almost nothing to the second player, but unlike the bargaining game, first players who deviated from equilibrium were not rewarded. After players gained experience with this game, the observed behavior converged to equilibrium.

In all these games, the behavior of the first mover is well accounted for by applying standard game-theoretic analysis, together with the (usual) assumption that first movers are income-maximizers, to the empirically

observed behavior of the second mover. However, in the bargaining game, the behavior of the second movers (i.e., the sellers who refuse positive offers) cannot be accounted for by a standard game-theoretic model built on the usual auxiliary assumption of income-maximization. Therefore, Figure 6 suggests that the buyers in the bargaining games adapt to the "nonstandard" behavior of the sellers in a "standard" game-theoretic way. Thus, these data suggest to us that, while the problem of developing descriptively powerful theory for games of this sort does not call for anything like the wholesale abandonment of the apparatus of game theory, neither is it likely that game-theoretic analysis unaided by empirical observation will lead to reliable models of behavior.

REFERENCES

- Albers, Wulf and Albers, Gisela, "On the Prominence Structure of the Decimal System," in R. W. Scholz, ed., *Decision Making under Uncertainty*, Amsterdam: Elsevier, 1983, pp. 271-87.
- Bolton, Gary, "A Comparative Model of Bargaining: Theory and Evidence," *American Economic Review*, December 1991, 81, 1096-1136.
- Guth, Werner, Schmittberger, R. and Schwarz, B., "An Experimental Analysis of Ultimatum Bargaining," *Journal of Economic Behavior and Organization*, December 1982, 3, 367-88.
- _____ and Tietz, Reinhard, "Ultimatum Bargaining Behavior: A Survey and Comparison of Experimental Results," *Journal of Economic Psychology*, September 1990, 11, 417-49.
- Harrison, Glenn W. and Hirshleifer, J., "An Experimental Evaluation of Weakest Link/Best Shot Models of Public Goods," *Journal of Political Economy*, February 1989, 97, 201-25.
- Kachelmeier, Steven J. and Shehata, Mohamed, "Culture and Competition: A Laboratory Market Comparison Between the People's Republic of China, the United States, and Canada," mimeo, University of Texas

²⁴Glenn Harrison and Jack Hirshleifer (1989) earlier studied the same game as a model of public-goods provision and called it the "best-shot" game. The rules are that player 1 proposes the amount of public good he will provide and then player 2 does, with the amount provided being the *maximum* provided by either player. However, each player must pay for the full amount that he agreed to provide. Consequently, both players have an incentive to free-ride, and at perfect equilibrium the first player does so.

- at Austin, 1990.
- Murnighan, J. Keith and Roth, Alvin E.**, "The Effect of Group Size and Communication Availability on Coalition Bargaining in a Veto Game," *Journal of Personality and Social Psychology*, January 1980, 39, 92-103.
- Ochs, Jack and Roth, Alvin E.**, "An Experimental Study of Sequential Bargaining," *American Economic Review*, June 1989, 79, 355-84.
- Prasnikar, Vesna and Roth, Alvin E.**, "Considerations of Fairness and Strategy: Experimental Data from Sequential Games," *Quarterly Journal of Economics*, 1991, forthcoming.
- Roth, Alvin E.**, "Bargaining Experiments," in J. Kagel and A. E. Roth, eds., *Handbook of Experimental Economics*, Princeton: Princeton University Press, 1992, forthcoming.
- Siegel, Sidney**, *Nonparametric Statistics for the Behavioral Sciences*, New York: McGraw-Hill, 1956.
- Spiegel, Matthew, Currie, Janet, Sonnenschein, Hugo and Sen, Arunava**, "First-Mover Advantage and the Division of Surplus in Two-Person, Alternating-Offer Games: Results from Bargaining Experiments," mimeo, Columbia University, 1990.