4

Bargaining Experiments

Alvin E. Roth

This chapter is divided into two major sections, the better to focus on particular organizing themes that emerge in the literature.*

Section I concerns the terms of agreement observed in bargaining experiments and the factors that cause these to vary. An important debate among bargaining experimenters focuses on the relative roles played by strategic considerations of the kind captured by game-theoretic models compared to sociological or cultural factors which may cause bargainers to focus on certain kinds of agreements in ways that are less sensitive to the strategic features of the bargaining situation. Early experiments uncovered a range of phenomena that were interpreted in almost contradictory ways by different experimenters. As more comprehensive experimental evidence accumulated, there was increasing agreement on the nature of the phenomena to be explained, and new phenomena emerged with wider implications. Some of these implications concern how completely it is possible to observe and control even experimental environments, while others concern the kind of learning that goes on as subjects gain experience in different strategic environments.

Section II concerns disagreements and costly delays, and the factors that lead to inefficiency in bargaining. One of the clearest experimental results, which also accords well with field data, is that a nonnegligible frequency of disagreements is a characteristic of bargaining in virtually all kinds of environments. An opportunity to explore this phenomenon further is presented by the fact that an exception to this generalization is observed in a set of experiments in which subjects interact face to face. A new experiment is discussed, whose results suggest that the cause of this anomaly lies in the special problems of experimental control that arise in relatively uncontrolled face to face encounters. This brief experiment was presented at the 1990 Handbook Workshop in Pittsburgh and has since sparked related experiments that together help illustrate how experiments can be used to test, refine, and/or reject alternative hypotheses. Section II then proceeds to survey the variety of experiments that have been conducted to explore the causes of disagreements and other forms of inefficiency in bargaining. The hypotheses that these experiments test span a wide range of game-theoretic models, and so the results also speak to the larger question of the status of game theory as a descriptive theory of observable behavior. Section II concludes with a discussion of the distribution of agreements over time in bargaining situations in which there is a deadline. This is a subject about which it had proved difficult to gather reliable field data. However, a very clear "deadline effect" was observed in a number of experiments, and this has helped to stimulate theoretical work that in turn suggests directions for further experiments.

Section III concludes with reflections on some of the things these several series of experiments suggest about bargaining, about game theory, about the relationship between observations in the laboratory and in the field, and about the relationship between theory and experiment.

I. Agreements

Different hypotheses about the roles of strategic and cultural factors have been advanced to explain the results of the unstructured bargaining experiments discussed in chapter 1. However, the debate has been brought into clearer focus with experiments designed around more highly structured bargaining. So, after setting the stage with a brief account of how these competing hypotheses have been applied to the unstructured bargaining experiments, I will turn to the main topic of this section, which concerns experiments in which communication between bargainers is limited to making offers, and accepting or rejecting them, with precise rules governing these exchanges.

There is a certain kind of quiet, unfinished drama that plays itself out as this series of experiments unfolds, because many of the investigators approach this subject from almost opposite points of view. To oversimplify a bit, there are the "game theory purists," who seek to test game-theoretic models, and refine them if necessary, and who see this as the best approach for understanding the choices bargainers make in the course of negotiations. And there are the "social norm purists," who seek to describe and understand what they see as binding social constraints that effectively determine bargainers' behavior. (Of course, many investigators are not purists of either sort, but lean towards one of these two very different points of view while being willing to incorporate elements of the other to improve the descriptive power of their hypotheses. And there are bounded rationality versions of these two approaches, as well as other points of view.) We will see how initial experiments in this series gave a common focus to the investigations conducted by these differently inclined investigators, and how subsequent experiments caused them to refine and modify hypotheses. When we get to the present, we will see that these hypotheses show signs of moving towards each other (albeit slowly) and that the nature of the remaining disagreements has been substantially clarified. This is, I think, typical of what we can expect of experiments.

A. Unstructured Bargaining Experiments

The experiments discussed in chapter 1, in which bargaining was conducted in relatively unstructured environments (e.g., with relatively free communication between bargainers) were initially designed to test a particular family of bargaining theories, and went on to investigate the nature of the unpredicted phenomena observed in the initial experiments. It is useful to begin our discussion here by considering some of the alternative hypotheses put forward to explain these phenomena, in order to see why it was natural to pursue the investigation of these hypotheses by turning to experiments in which the individual choices of the bargainers (and not just the outcome of their mutual interaction) could be more clearly observed.

Recall from section III.C of chapter 1 that in the binary lottery game experiments of Roth and Malouf (1979), it was observed that bargainers who were uninformed about the cash value of one anothers' prizes tended to reach agreements in which lottery tickets were divided evenly, while bargainers who were both informed about one anothers' prizes tended to reach agreements in which the bargainer with the smaller prize obtained a higher percentage of the lottery tickets. While the experiment of Roth, Malouf, and Murnighan (1981) supported the view that this shift could not be accounted for entirely in terms of the strategic features of the game, the subsequent experiment of Roth and Murnighan (1982) suggested that the phenomenon was something the bargainers themselves were able to take account of in a strategic way, since the tradeoff between terms of agreement and rates of disagreement in that experiment were observed to coincide closely with what would be expected at a strategic equilibrium.

My colleagues and I tentatively interpreted the data as suggesting that certain agreements became "focal" for reasons that might not be captured by the game-theoretic models (recall the early experiments of Schelling [1957] discussed in chapter 1), but that the existence of these focal agreements was recognized by the bargainers, who incorporated them into their behavior in a strategic, game-theoretic manner. The experiment of Roth and Schoumaker (1983), which showed that the choice of equilibrium could be influenced by manipulating subjects' expectations, provided indirect support for the hypothesis that the effect of the different information conditions on the outcomes observed in Roth and Murnighan (1982) could be accounted for by the way such information changed subjects' expectations. And in Roth (1985) I suggested that parts of the data appeared similar to equilibrium behavior in a coordination game in which the focal points were taken as given. (I will return to this latter study in section II.C.)

A related focal point hypothesis has recently been tested by Mehta, Starmer, and Sugden (1990), who constructed an experiment in which focal points could be manipulated.[1] They observed bargaining situations in which two subjects had to agree on how to divide £10 between them, with each bargainer receiving zero if no agreement were reached. Before the bargaining began, the subjects were dealt four playing cards each, from a deck consisting of eight cards, four aces and four deuces. The subjects were told that all four aces were worth £10, and in order to be paid they must agree to pool their aces and agree how to divide the £10. Since the agreement of both players is required for any money to be received, conventional game-theoretic models treat this as a completely symmetric problem (since only situations in which neither player held all four aces were considered). Yet, although Mehta et al. observed that equal divisions were the modal proposal by holders of one, two, or three aces, they noted that deviations were in the direction of giving more to the bargainer with more aces, with a second mode being a demand of only £2.50 by holders of only a single ace. Their interpretation is that the bargainers use the cards dealt to them as cues to help solve the coordination problem embedded in any bargaining problem, in a manner that causes divisions proportional to cardholdings to join the equal division as a focal agreement.

Other experimenters have interpreted those earlier "focal point" experiments differently. For example, two views that are also very different from each other are expressed by Harrison (1990) and Guth (1988), both of whom consider the experiments of Roth et al. discussed above. Harrison suggests that the appearance of focal points can be explained in entirely game-theoretic terms, since all agreements can arise as strategic equilibria in such games. Guth, on the other hand, proposes that the data is best explained in entirely *non*-game-theoretic terms. He outlines a "behavioral theory of distributive justice," according to which bargainers conclude agreements at what they perceive to be a fair distribution, with information about each others' prizes allowing them to utilize more fundamental notions of fairness. A similar interpretation of these experiments is independently proposed by Foddy (1989). Both Guth and Foddy can be interpreted as proposing that a descriptive theory of bargaining behavior must essentially *be* a theory of what constitute fair distributions of income.[2]

Thus a number of quite different hypotheses about bargaining behavior have been used to organize the data from these unstructured bargaining experiments. Since these hypotheses concern the choices facing the bargainers in the course of negotiations, the experiments that have been conducted to directly investigate them have tended to focus on more structured bargaining situations, in which these choices can be directly observed, and about which these hypotheses therefore make more pointed predictions.[3] We turn now to consider these more structured bargaining experiments.

B. *Sequential Bargaining Experiments*

There has been a good deal of attention, both theoretical and experimental, given to models of two-party bargaining in which time is divided into periods and the opportunity to make an offer alternates between the bargainers. The basic model motivating most of the experiments to be discussed is the following: two bargainers, 1 and 2, alternate making offers over how to divide some amount k of money.

In odd numbered periods t (starting at an initial period $t = 1$) player 1 may propose to player 2 any division $(x_1, x_2) = (x, k - x)$. If player 2 accepts this proposal then the game ends and player 1 receives a utility of $(\delta_1)^{(t-1)}x$ and 2 receives a utility $(\delta_2)^{(t-1)}(k - x)$, where δ_i is a number between 0 and 1 reflecting player i's cost of delay. (That is, a payoff of y dollars to player i at period t gives him the same utility as a payoff of $\delta_i y$ dollars at period $t - 1$.) If player 2 does not accept the offer and if period t is not the final period of the game, then the game proceeds to period $t + 1$, and the roles of the two players are reversed. If an offer made in the last period of the game is refused, then the game ends with each player receiving 0. A game with a maximum number of periods T will be called a T-period game.[4] An observer of such a game will see not only the final outcome, but a sequence of individual choices concerning what offers to make and whether to accept or reject them.

Such a game has many strategic equilibria, but most of these can be thought of as involving an attempt by one of the bargainers to threaten a course of action he would not wish to carry out if his bluff were called. For example, in a two-period game, the player who makes the offer in the first period, player 1, might demand 99 percent of the gains from trade for himself, and threaten that, if player 2 refuses to accept this offer, then in the second period he (player 1) will refuse *any* offer, so that disagreement will result and each player will receive nothing. If this threat is believed, player 2's best response is to accept the 1 percent she is offered in the first period. But the threat implies that, if player 2 rejects the offer in the first period, player 1 will reject offers in the next period that he would then prefer to accept. For this reason such threats may not be credible. The class of equilibria that do not involve such threats are called *subgame perfect*.

A subgame perfect equilibrium can be computed by working backward from the last period. An offer made in period T is an ultimatum, and so at such an equilibrium player i (who will receive 0 if he rejects the offer) will accept any nonnegative offer when payoffs are continuously divisible.[5] So at a subgame perfect equilibrium, player j, who gets to make the proposal in period T, will receive 100 percent of the amount k to be divided, if the game continues to period T. Consequently at period $T - 1$ player j will refuse any offer of less than $(\delta_j)k$ (the present value of what she will get if the game continues to the next period) but accept any offer of more, so that at equilibrium player i receives the share $k - (\delta_j)k$ if the game goes to period $T - 1$, and so at period $T - 2$ he must be offered $(\delta_i)(k - [\delta_j]k)$, and so forth. Working back to period 1 in this way, we can compute the equilibrium division: that is, the amount that the theory predicts player 1 should offer to player 2 at period 1, and player 2 should accept.[6]

The earliest experimental studies of this kind of bargaining reported markedly different results. Their authors drew different conclusions, along the lines of the various hypotheses discussed in section I.A, about the predictive value of perfect equilibrium models of bargaining and about the role that experience, limited foresight, or bargainers' beliefs about fairness might play in explaining their observations. (Questions of fairness arise because in some of these experiments, as in the unstructured bargaining experiments, many observed agreements give both bargainers 50 percent of the available money.) Subsequent experiments brought more agreement on the description of the phenomena to be explained, and the most recent experiments have started to narrow some of the differences in interpreting these phenomena.

1. An Initial Exchange of Views

In each of the following experiments, the predictions tested involved only the ordinal utilities of the bargainers, not their risk posture. Following standard practice in the experimental literature when only ordinal utilities are of concern, in the initial experiments the utility of the bargainers was assumed to be measured by the amount of money they receive (a point I will discuss in detail later).

Guth, Schmittberger, and Schwarz (1982) examined one-period ("ultimatum") bargaining games. Player 1 could propose dividing a fixed sum of k deutsche marks any way he chose, by filling out a form saying "I demand DM x." Player 2 could either accept, in which case player 1 received x and player 2 got $k - x$, or she could reject, in which case each player received 0 for that game. (The subjects were divided into two groups of equal size, with the offer of each player 1 being assigned at random to one of the player 2's, so that no bargainer knew with whom he was bargaining in the other group.)

The perfect equilibrium prediction for such games is that player 1 will ask for and get (essentially) 100 percent of k. However the average demand that players 1 were observed to make was for under 70 percent, both for players playing the game for the first time and for those repeating the game a week later.[7] About 20 percent of offers were rejected. The authors conclude that

> . . . subjects often rely on what they consider a fair or justified result. Furthermore, the ultimatum aspect cannot be completely exploited since subjects do not hesitate to punish if their opponent asks for "too much."[8] (384)

A different conclusion is reached by Binmore, Shaked, and Sutton (1985), who write:

> The work of Guth et al. seems to preclude a predictive role for game theory insofar as bargaining behavior is concerned. Our purpose in this note is to report briefly on an experiment that shows that this conclusion is unwarranted. . . .[9] (1178)

Their experiment studied a two-period bargaining game, in which player 1 makes a proposal of the form $(x, 100 - x)$ to divide 100 pence. If player 2 accepts, this is the result. Otherwise, 2 makes a proposal $(x', 25 - x')$ to divide 25 pence. If player 1 accepts, this is the result; otherwise, each player receives 0. Thus in this game $\delta_1 = \delta_2 = 0.25$, and (since proposals are constrained to be an integer number of pence) at any subgame perfect equilibrium player 1 makes an opening demand x in the range 74–76 pence, and player 2 accepts any opening demand of 74 pence or less. Subjects played a single game, after which player 2 was invited

to play the game again, as player 1. In fact, there was no player 2 in this second game, so only the opening demand was observed.[10]

The modal first demand in the first game was 50 pence, and 15 percent of the first offers were rejected. In the second game (in which only first demands were observed), there was a mode around a first demand near 75 pence. There was thus a clear shift between the two distributions of first demands, in the direction of the equilibrium demand. The authors conclude,

> Our suspicion is that the one-stage ultimatum game is a rather special case, from which it is dangerous to draw general conclusions. In the ultimatum game, the first player might be dissuaded from making an opening demand at, or close to, the "optimum" level, because his opponent would then incur a negligible cost in making an "irrational" rejection. In the two-stage game, these considerations are postponed to the second stage, and so their impact is attenuated. (1180)

Guth and Tietz (1988) responded with an experiment examining two two-stage games with discount factors of 0.9 and 0.1, respectively. So the subgame perfect equilibrium predictions (in percentage terms) for the two cases are (10%, 90%) and (90%, 10%), respectively. They say:

> Our hypothesis is that the consistency of experimental observations and game theoretic predictions observed by Binmore et al. as well as by Fouraker and Siegel is solely due to the moderate relation of equilibrium payoffs which makes the game theoretic solution socially more acceptable.

Subjects played one of the two games twice, each with a randomly chosen other bargainer. Subjects who played the first game as player 1 played the second game as player 2. One difference from the sequential bargaining games discussed above was that disagreement automatically resulted if player 2 rejected an offer from player 1 but made a counterproposal that would give her (player 2) less than player 1 had offered her.[11]

In the first game, the average first demand in games with a discount factor of .1 was 76 percent, and in the second game 67 percent (compared with a perfect equilibrium prediction of 90 percent). For games with a discount factor of 0.9, the average first demand in the first game was 70 percent, and in the second game 59 percent (compared to a predicted first demand of 10 percent).[12] The authors conclude:

> Our main result is that contrary to Binmore, Shaked and Sutton "gamesmenship" is clearly rejected, i.e., the game theoretic solution has nearly no predictive power.

Neelin, Sonnenschein, and Spiegel (1988) also responded to Binmore, Shaked, and Sutton (1985). They reported two experiments involving two-period, three-period, and five-period bargaining games. Neelin et al. observe that the data for all their (2, 3, and 5 period) games are near the perfect equilibrium prediction for two period games. They conclude:

> The strong regularity of the behavior we observed is one of the most noteworthy aspects of our results and lends power to our rejection of both the Stahl/Rubinstein theory and the equal-split model. (829)

In a reply, Binmore, Shaked, and Sutton (1988) declined to attribute the same significance to these results and conjectured that the differences described among these experiments may be due to the differences in experimental procedures employed.

Thus different experimenters reached markedly different conclusions, based on experiments with different parameters and using different procedures.

2. A Larger Experimental Design

Following most of this exchange, Ochs and Roth (1989) conducted an experiment utilizing a larger experimental design, which allowed games with different parameters to be compared under a common set of procedures. They noted that the prior analyses had focused on the accuracy of the perfect equilibrium as a point predictor, that is, on whether the observed outcomes were distributed around the perfect equilibrium division or around some other division of the available money. Their experiment was designed to test the predictive accuracy of some of the *qualitative* predictions of the perfect equilibrium in sequential bargaining, and to detect if changes in the parameters of the game influence the observed outcomes in the predicted direction, even in the case that there might be a systematic error in the point predictions.[13] To this end the experiment was implemented in a way that allowed the discount factors of the two bargainers to be varied independently.[14] In order to compare games like those considered in the earlier experiments, the experimental design allowed comparisons between different combinations of discount factors for games of fixed length, as well as between games of different length for given discount factors. The eight cells of the experiment compare two and three period games using all four combinations of discount factors (δ_1, δ_2), with δ_i equal to 0.4 or 0.6 (see Figure 4.1). The bargainers sat in two rooms and conveyed their offers and responses by filling out a written form. Each bargainer participated in ten bargaining encounters, against a different (anonymous) partner in each round. At the conclusion of each experimental session, one round was chosen at random to be the payoff round, and each bargainer was paid his earnings for that round.[15]

Figure 4.2 displays the following data for each cell of the experiment: (1) the number of bargaining pairs per round; (2) the mean of the observed first period offers to player 2 in each of the ten rounds; (3) the maximum and minimum first period offers in each round; (4) plus and minus two standard errors from the mean offer in each round; (5) the number of first period offers that were rejected in each round. In addition to the data, the perfect equilibrium offer and the equal division offer (which is always $15) are displayed. The offers made in round 10 of each cell represent the behavior of the most experienced bargainers. As Figure 4.2 shows, the subgame perfect equilibrium offer was generally a very poor point

	Two Period		Three Period	
	Chips	Money	Chips	Money
δ_1=.4, δ_2=.4	Cell 1: (59,41) to (61,39)	($17.70, $12.30) to ($18.30,$11.70)	Cell 5: (76,24)	($22.80,$7.20)
δ_1=.6, δ_2=.4	Cell 2: (59,41) to (61,39)	($17.70, $12.30) to ($18.30,$11.70)	Cell 6: (84,16)	($25.20,$4.80)
δ_1=.6, δ_2=.6	Cell 3: (39,61) to (41,59)	($11.70, $18.30) to ($12.30,$17.70)	Cell 7: (77,23) to (76,24)	($23.10,$6.90) to ($22.80,$7.20)
δ_1=.4, δ_2=.6	Cell 4: (39,61) to (41,59)	($11.70, $18.30) to ($12.30,$17.70)	Cell 8: (65,35)	($19.50,$10.50)

Figure 4.1. Experimental design and range of equilibrium predictions. *Source:* Ochs and Roth, 1989.

predictor of the observed outcomes. Cell 1 is the only cell in which the perfect equilibrium offer is within two standard errors of the observed mean. In no other cell does the perfect equilibrium offer fall within plus or minus two standard errors of the estimated population mean.

The perfect equilibrium not only fails as a point predictor of observed behavior, it also fails to account for observed qualitative differences between cells, such as mean first period offers. (As a predictor of the direction of differences in pairwise comparisons of means, the theory does little better than coin flipping.) And while parts of the data appear to be consistent with similar observations made in the earlier experiments, the larger experimental design allows more comparisons to be made, so that observations which, piecewise, appear contradictory, emerge as part of a larger picture. In this regard, the paper notes (379):

> If we had looked only at Cell 1 our conclusions might have been similar to those of Binmore et al., since the data for that cell looks as if after one or two periods of experience, the players settle down to perfect equilibrium proposals. . . . And if we had looked only at Cells 1 and 5, our conclusions might have been similar to those of Neelin et al., since in those two cells both the two and three period games yield observations near the two period predictions. . . . And if we had looked only at cells 5 and 6, we might have concluded, like Guth and Teitz, that the phenomena observed here was closely related to the relatively extreme equilibrium predictions in those cells.

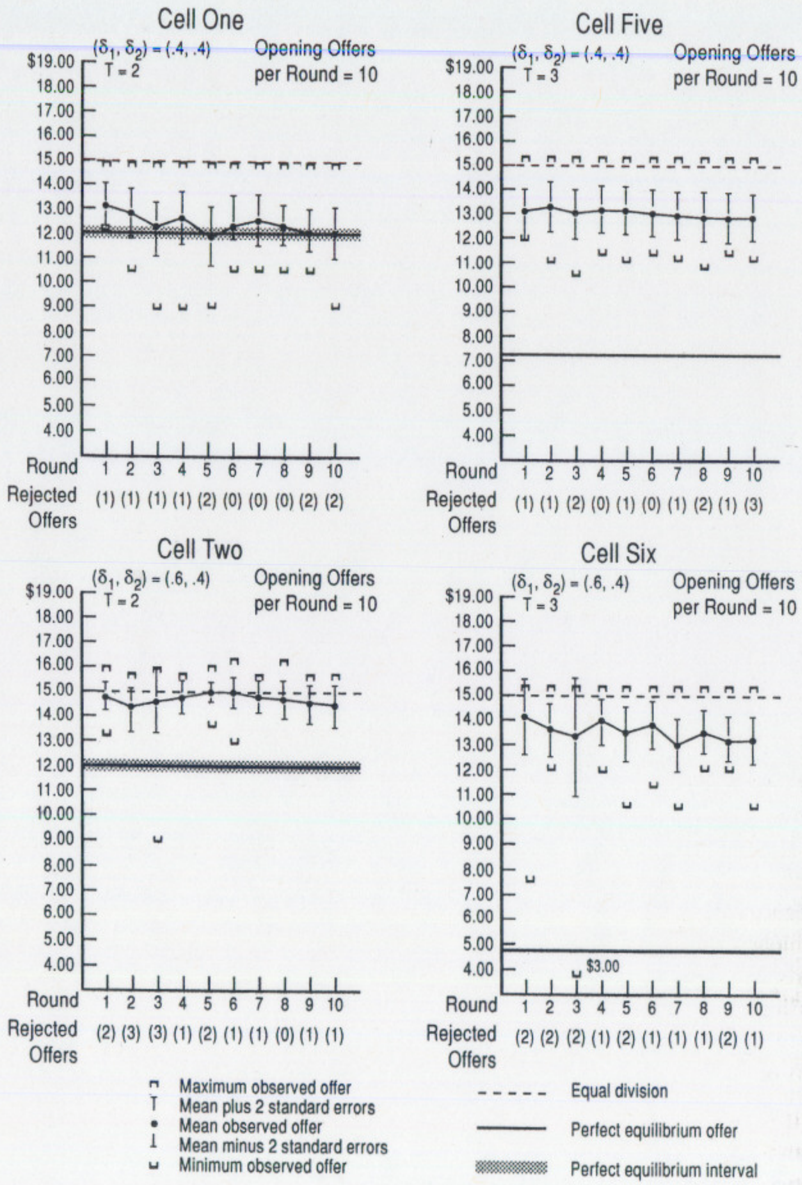

Figure 4.2a. Opening offers to Player 2. *Source:* Ochs and Roth 1989.

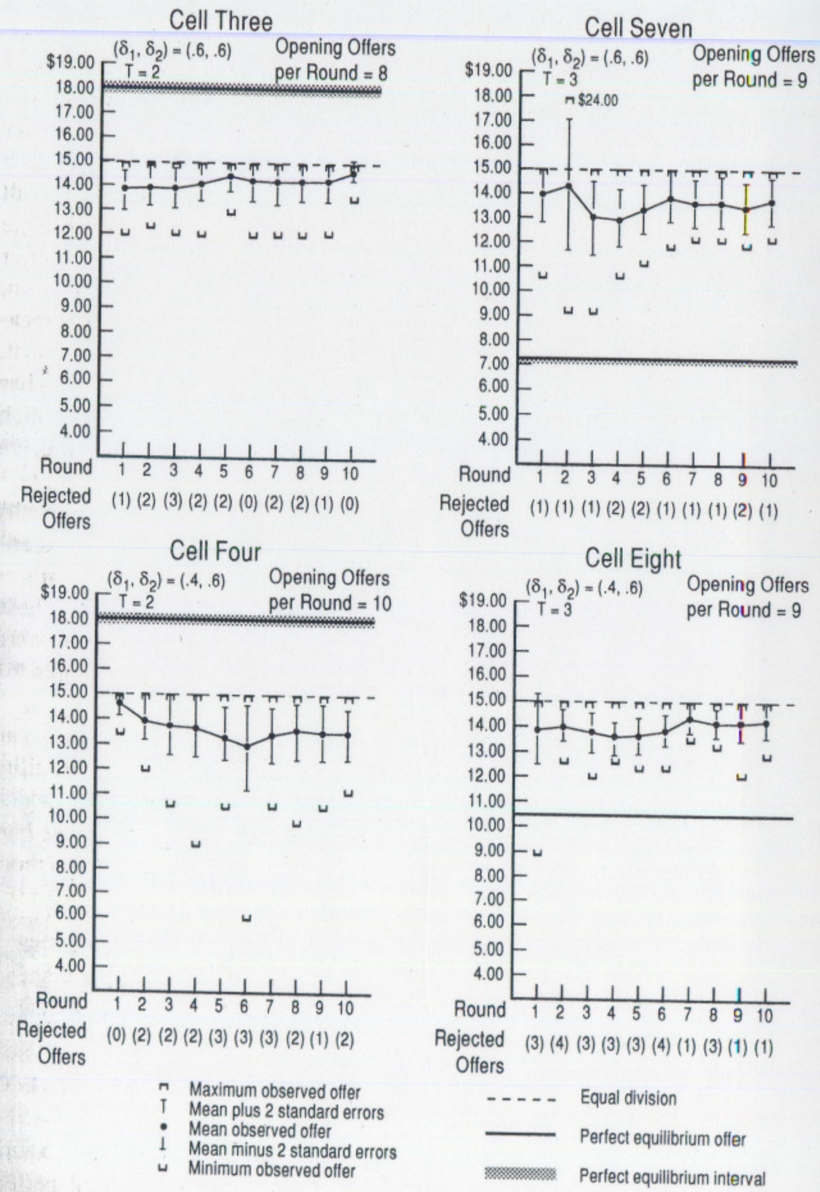

Figure 4.2b. Opening offers to Player 2. *Source:* Ochs and Roth 1989.

Perhaps the most interesting observed regularity for our present purposes concerns what happens when first period offers are rejected, both in this experiment and, as it turns out, in the previous experiments. Approximately 15 percent of first offers met with rejection (including those in games with experienced subjects who had played ten games against different opponents), and of these well over half were followed by counterproposals in which player 2 demanded *less* cash than she had been offered. A significant number of player 2s were rejecting small shares of the relatively large gains available in the first period in favor of large shares of the much smaller gains available in the second period. Since, after player 1 has made a proposal, player 2 is faced with an individual choice problem, we can conclude by revealed preference that these player 2s' utility is *not* measured by their monetary payoff, but must include some nonmonetary component. That is, a player 2 who has made this kind of disadvantageous counteroffer has chosen to make it rather than accepting the offer player 1 had already made, which would have given her (player 2) a higher monetary payoff than she would receive (even) if her own counteroffer is accepted.

Perhaps the most significant observation, though, in view of the diverse results reported in previous experiments, was that when the data of the previous experiments were reanalyzed with this in mind, it turned out that this pattern of rejections and disadvantageous counterproposals was strikingly similar in all of these experiments.[16] Thus the single-period, ultimatum games (in which subjects were observed to reject positive offers) were not a special case in this respect. Table 4.1 summarizes the data.

Ochs and Roth (1989) go on to argue that this and other patterns in the data can plausibly be explained if the unobserved and uncontrolled components of utility in these experiments have to do with subjects' perceptions of "fairness," which involve comparing their share of the available wealth to that of the other bargainer. They note that in most cases agents propose divisions that give them more than half of the proceeds, and say:

> We do not conclude that players "try to be fair." It is enough to suppose that they try to estimate the utilities of the player they are bargaining with, and . . . at least some agents incorporate distributional considerations in their utility functions.

That is, if agents' preferences are such that they will refuse "insultingly low" offers, then this must be taken into account in making offers.[17]

Note that uncontrolled elements in the bargainers' utility in these experiments suggests that none of them can be easily interpreted as tests of perfect equilibrium per se, since to compute a perfect equilibrium we need to know the preferences of the players (and so do they).[18] But the uniformity with which disadvantageous counterproposals have appeared, in contrast to the otherwise quite varied results of these experiments, suggests that bargaining may be an activity that systematically gives bargainers motivations distinct from simple income maximization.

Table 4.1. Disadvantageous Responses

Study	Number of Observations	First-Offer Rejections (%)	Disadvantageous Responses (%)
Guth et al. (1982)	42	19 (8/42)	88 (7/8)[a]
Binmore et al. (1985)	81	15 (12/81)	75 (9/12)
Neelin et al. (1988)	165	14 (23/165)	65 (15/23)
Ochs and Roth (1989)	760	16 (125/760)	81 (101/125)

Source: Ochs and Roth 1989.

[a] One of the rejections was of a (100,0) division, and so was not disadvantageous.

In summary, while many regularities were observed and while the rather different results of previous experiments were mirrored in this experiment as pieces of a larger pattern, the disadvantageous counteroffers observed in this experiment, and then found in previous experiments, suggest that the phenomena being studied have elements that have so far eluded experimental control.

That being the case, it is not surprising that this experiment also has been subject to different interpretations by investigators with different points of view. For example, Thaler (1988), who agrees that this and other evidence suggest that "subjects' utility functions have arguments other than money," writes:

> We have seen that game theory is unsatisfactory as a positive model of behavior. It is also lacking as a prescriptive tool. While none of the subjects in Ochs and Roth's experiments came very close to using the game-theoretic strategies, those who most closely approximated this strategy did not make the most money. (202)

Guth and Tietz (1990) suggest that even less of the traditional apparatus of economic theory can be saved. They write that they "strictly reject" the conclusion of Ochs and Roth that there are uncontrolled elements in the utility functions of the bargainers, since this implies that the bargainers engage in tradeoffs between underlying preferences and strategic considerations. Rather, they favor modeling players as shifting between strategic and equitable considerations in a hierarchical way so that at any point in time players are primarily concerned with one aspect of the problem.

Kennan and Wilson (1993), on the other hand, focusing on the disadvantageous counteroffers, argue that, since bargainers' preferences were not completely controlled in any of the experiments in which Ochs and Roth found disadvantageous counteroffers, the bargainers themselves could not have had common knowledge of one anothers' preferences, and so the most promising models of these phenomena are models of games of incomplete information—i.e., game-theoretic models in which players' uncertainty about one anothers' preferences is explicitly modeled.

Thus this experiment, while it brings some unity to the phenomena that investigators of various persuasions wish to explain, continues to permit quite different interpretations. And it raises new questions. We turn next to consider recent experiments that address some of these.

3. Investigating Observed Regularities

Before trying to explain the pattern of results observed in Ochs and Roth (1989), it is natural to first ask whether this pattern is robust, and which if any aspects of it are sensitive to the particular parameters chosen in that experiment. Several studies address this question.

Weg, Rapoport, and Felsenthal (1990) consider alternating offer bargaining games in which the subjects were not informed precisely how many periods would be allowed, but were given to understand that there would be many. (In fact, bargaining was terminated after twenty periods, and Weg et al. report that only a small fraction of the games they observed went this long.) They first looked at games with discount factors (δ_1, δ_2) of (.9,.5), (.67,.67), and (.5,.9), and then in a second experiment with lower discount factors, of (.5,.17), (.17,.17), and (.17,.5). Of the unpredicted regularities reported by Ochs and Roth (1989), they focused on the frequency of disadvantageous counteroffers, and of offers of equal monetary payoffs and of equal divisions of chips.[19] They report that their results are entirely consistent with the previous observations that the perfect equilibrium is a poor predictor, that there are many offers of equal divisions (predominantly equal monetary divisions), and that many rejections are followed by disadvantageous counteroffers. They do note that the percentage of disadvantageous counteroffers declines as player 2's discount factor rises. This seems natural, since the range of nondisadvantageous counteroffers available to player 2 increases as his discount factor increases.

In another experiment, Rapoport, Weg, and Felsenthal (1990) report a different kind of alternating offer game, in which each player pays a fixed fee for continuing the bargaining another period, rather than having the value of his payoffs diminish by a fixed percentage. They find that the perfect equilibrium predictions perform much better in this kind of game, an observation I shall return to later, in section B.4.[20]

Bolton (1991) reports a comprehensive investigation that begins by considering the replicability of four of the observed regularities enumerated by Ochs and Roth (1989). These are that there was a consistent first mover advantage, that observed mean offers deviated from the perfect equilibrium prediction in the direction of equal division, that a substantial proportion of first period offers were rejected, and that a substantial proportion of rejections were followed by disadvantageous counteroffers. He begins by replicating these observations (see Figure 4.3a,b, cells 1 and 2) for two-period games using discount factors (δ_1, δ_2) of (1/3, 2/3) and (2/3, 1/3), and an initial pie of $12. Like Weg et al. (1990), he observes that the rate of disadvantageous counteroffers is sensitive to the discount factors.[21] He then considers whether these features of the data might be due to the inexperience of the subjects, but finds that when experienced subjects (with experience from earlier sessions) play the game, the observed outcomes exhibit the same phenomena (see Figure 4.3a, cell 4)[22]: the observed agreements have the same means as those of inexperienced subjects, with lower variance, and "the aggregate data on rejections and disadvantageous counteroffers is very similar."

Bolton also tests whether bargainers might be influenced by the fact that they keep the same role from round to round, or that these roles were assigned randomly, by conducting a trial in which bargainers alternate between being player 1 and player 2 for 12 rounds. He finds that the mean agreements observed under these procedures do not differ from those in which bargainers retained the same (randomly assigned) role for all trials.

Having satisfied himself as to the robustness of these observed regularities, Bolton then turns his attention to explaining them, together with the additional regularity noted by Ochs and Roth, that even in two-period games the discount factor of player 1 influences the outcome (in contradiction to the perfect equilibrium prediction when players are simple income maximizers). He proposes a theory of bargainer preferences in which bargainers care not only how much they earn, but what share of the pie they receive. That is, he proposes that bargainers be modeled as having utility functions with two arguments, income and relative share. The idea is that this latter argument is what causes bargaining to deviate from the perfect equilibrium predictions for bargainers who are concerned only with their own income, particularly insofar as it causes bargainers to make rejections and counteroffers that are disadvantageous in terms of income but not in terms of relative share. Bolton further postulates that a bargainer will compare his payoff to that of the person he is bargaining with only when they "share the same pie." In particular, he predicts that when players are paid tournament-style, based on how well they do in comparison with other players in the same position (i.e., other player 2s for a player 2 and other player 1s for a player 1), observed agreements will conform more closely to the perfect equilibrium prediction.

The extent to which this treatment has the predicted effect can be assessed by examining Figures 4.3a and b (cells 5–8), which allow the results of bargaining under a tournament compensation scheme to be compared with those observed when the bargainers are paid in proportion to their share of any agreement reached. Cells 5 and 6 show little evidence of difference (from cells 1 and 2, respectively) when the bargainers are inexperienced, but this changes when experienced bargainers are observed. Comparing cells 7 and 4 we see that when the discount factors are (2/3, 1/3) the observed agreements converge to the perfect equilibrium prediction in the tournament condition, but not in the ordinary payoff condition. However, this seems to be sensitive to the discount factors, as the agreements reached by experienced bargainers in the tournament condition of the (1/3, 2/3) game (cell 8) show no sign of converging to the predicted agreement. So these experiments provide some support for the hypothesis that there are uncontrolled factors in the bargainers' preferences, but leave open many questions about the nature of these factors.

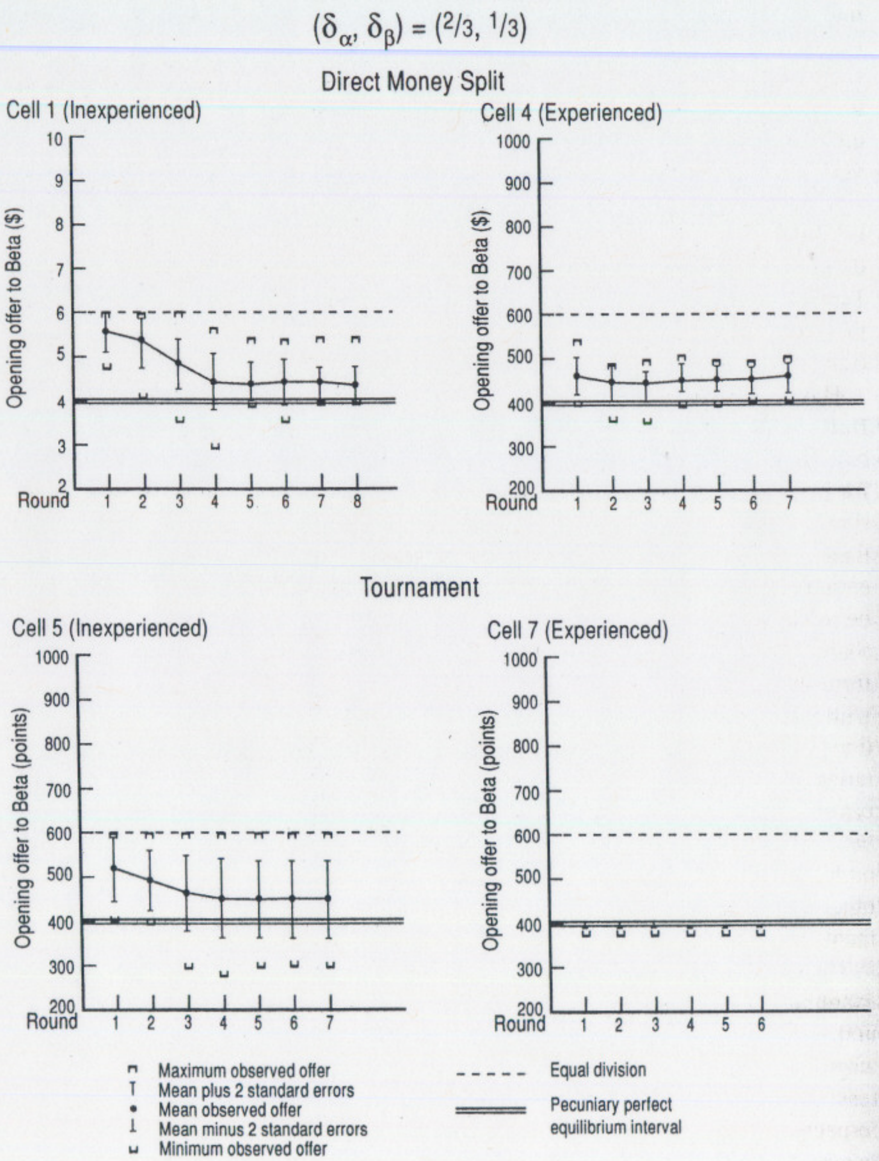

Figure 4.3a. Opening offers with $(\delta_\alpha, \delta_\beta) = (2/3, 1/3)$. *Source:* Bolton 1991.

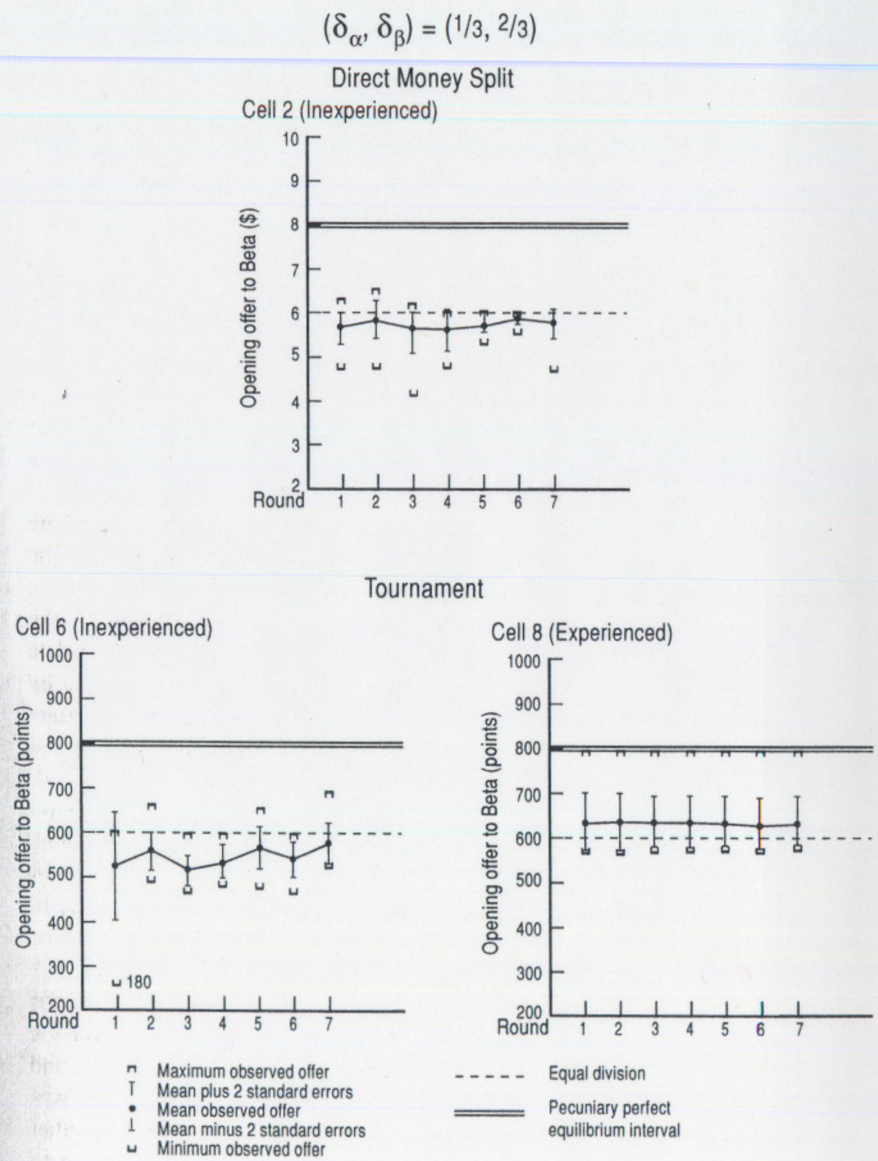

Figure 4.3b. Opening offers with $(\delta_\alpha, \delta_\beta) = (1/3, 2/3)$. *Source:* Bolton 1991.

a. Are Players "Trying to Be Fair"?

One hypothesis about bargainers' preferences that has received some attention is the possibility that they may behave altruistically, at least to the extent that they "try to be fair." This hypothesis cannot be directly tested on data from bargaining games of the kind discussed above, because a player who offers an equal division, for example, may be doing so in order to avoid a negative reaction by the other bargainer.

To explore this hypothesis, Forsythe, Horowitz, Savin, and Sefton (1994) compared ultimatum and "dictator" games. Like an ultimatum game, a dictator game is a two-player game in which player 1 proposes a division of some resource between the two. However, unlike an ultimatum game, in a dictator game player 2 may not reject this proposal (and cause both players to receive zero); the players receive whatever player 1 (the dictator) proposes. So in a dictator game player 1's proposal can be interpreted as a pure expression of his preferences. In the dictator game, but not in the ultimatum game, they observe that the modal offer is the equilibrium offer, at which player 1 offers zero to player 2 (see Figure 4.4, which graphs proposed divisions of \$5). Forsythe et al. conclude that "players are more generous in the ultimatum game than in the dictator game," and so they reject the "fairness hypothesis" as the primary explanation for the generous offers observed. But Forsythe et al. also observe a concentration of offers of equal division in the dictator game, and this is the modal result in the ultimatum game. So the data support the hypothesis that *some* of the subjects may be primarily motivated by considerations of fairness, but that the high concentration of equal division offers observed in the ultimatum game cannot be attributed to a simple desire for equal divisions on the part of players 1.

Bolton (1991) makes a similar observation in the context of sequential bargaining games with more than one period, by considering two period games in which the last period is a dictator game. He observes that player 1s in such games offer more than they do in ordinary two period bargaining games (in which they will have a chance to reject player 2s counteroffer). Thus, in these games also, the player 1s respond to the strategic difference between the two kinds of games.

Of course, if even a few players are substantially influenced by considerations of fairness, this may in some circumstances have a large effect on the strategic environment in which the other players must operate. Kahneman, Knetsch, and Thaler (1986a) created such an environment. After an ultimatum game was played, subjects were told that they had a choice of dividing some money either with another subject, U, who in the previous ultimatum game had chosen to offer an unequal division, or with a subject, E, who had previously chosen to offer an equal division. If they chose U, then the two of them (the chooser and subject U) would each receive \$6, and subject E would receive zero. If they chose E, then the two of them would each receive \$5, and subject U would receive zero. A majority chose E, thus exhibiting a willingness to sacrifice a dollar in order that E rather than U should be paid. (However, among subjects who had themselves made unequal offers in the previous ultimatum game, only a minority chose E.)

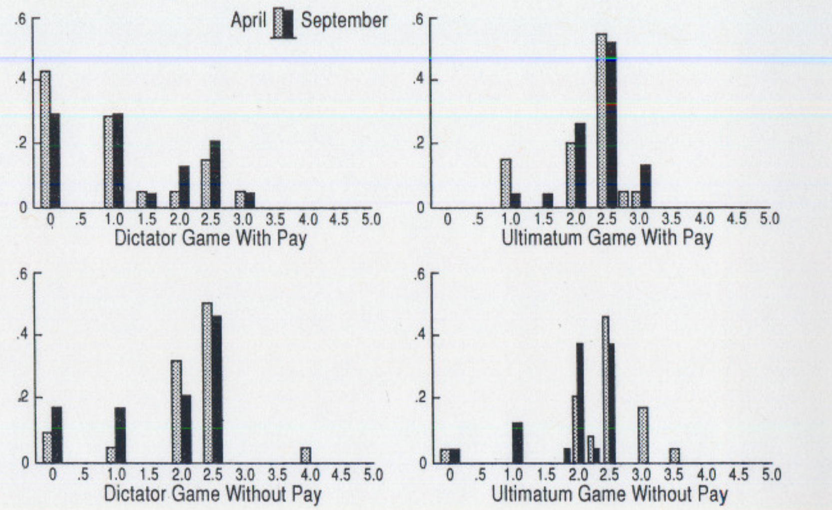

Figure 4.4. Offers in dictator and ultimatum games. *Source:* Forsythe, Horowitz, Savin, and Sefton 1994.

To the extent that some subjects are willing to punish past behavior that they see as unfair (even to others) one can easily imagine how social norms could be created and enforced. For some theoretical literature on the evolution of social norms in environments in which deviators can be punished, see Guth and Yaari (1990a,1990b), Okuno-Fujiwara and Postlewaite (1990), and Kandori (1992). Bolton (1993) and Van Huyck et al. (1992) describe the evolution of social norms in a more biological sense.

But there is a chicken and egg problem here. Although subjects may have clear ideas about what is fair in a variety of circumstances,[23] and although these ideas about fairness may influence the strategic environment, the evidence suggests that subjects adapt their ideas about what is fair in response to their experience, in ways that may be heavily influenced by strategic considerations. That is, although the strategic environment is influenced by ideas about fairness, ideas about fairness are influenced by the strategic environment.

For example, Binmore, Morgan, Shaked, and Sutton (1991) studied two closely related alternating offer sequential bargaining games, which differed only in a relatively subtle way, concerning how the games ended. In "optional breakdown games," players who did not reach agreement in a given period could continue to the next period with the size of the pie reduced by a discount factor δ, unless one of them chose to end the game. If either player chose to end the game, players 1 and 2 received "breakdown" payments α and β respectively, where α is a small fixed percentage of the pie, and β is a larger percentage that changed between games as one of the experimental variables. In "forced breakdown games," players who did not reach agreement in a given period could continue to the next period with probability δ, but with probability $1\text{-}\delta$ the game would end and the players would receive their breakdown payments. In each game, the

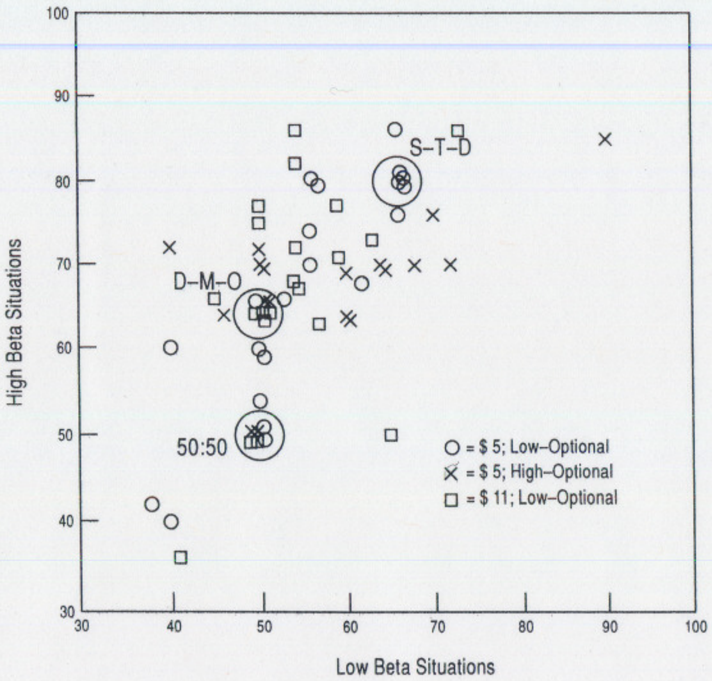

This figure shows the share of the cake for player II proposed as "fair" for low and high β situations by subjects who had experienced *optional* breakdown games.

Figure 4.5a. "Fair" divisions. *Source:* Binmore, Morgan, Shaked, and Sutton 1991.

amount the players could divide by reaching agreement was greater than the sum of their breakdown payments, and β was set to be either 36 percent or 64 percent of the total available pie, while α was equal to 4 percent.

The subgame perfect equilibrium prediction for the optional breakdown game is that the relative sizes of the breakdown payments will not influence the outcome of negotiations unless the breakdown payment of one of the players is larger than his share of the pie at equilibrium of the game in which players' breakdown payments are zero. (The reason is that at a perfect equilibrium in this case no player would choose at any subgame to end the game and receive a payment smaller than his equilibrium payment from continuing, so the breakdown payments are irrelevant.) But the prediction for the forced breakdown game is that agreements should be reached that give a larger share of the pie to the player with the larger breakdown payment, even when that payment is relatively small. (The reason is that in any subgame in which agreement has not been reached, the player with the higher breakdown payment has a higher expected payoff, since there is a positive probability of breakdown.)

With these parameters, the predictions are approximately that player 2's share will be 50 percent in the "low beta" (β = 36 percent) optional breakdown game and 64 percent in the "high beta" optional breakdown game, a pattern

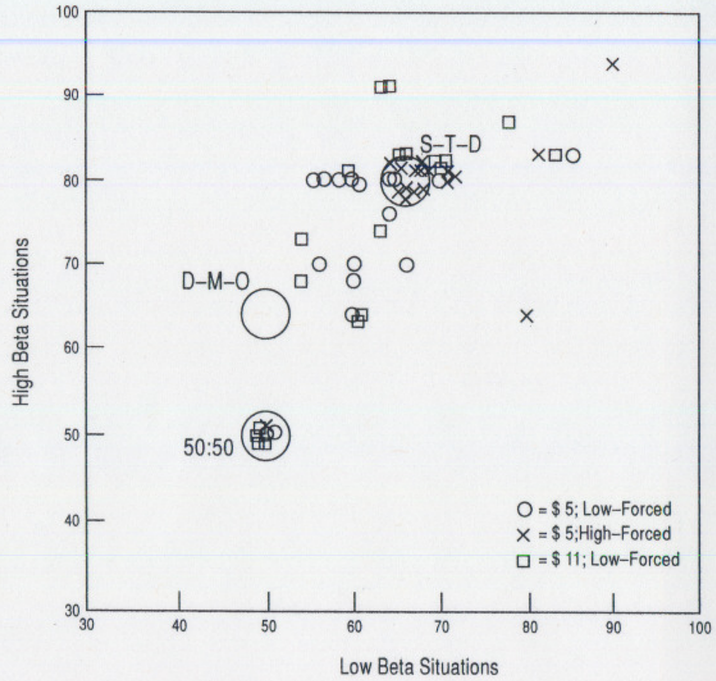

This figure shows the share of the cake for player II proposed as "fair" for low and high β situations by subjects who had experienced *forced* breakdown games.

Figure 4.5b. "Fair" divisions. *Source:* Binmore, Morgan, Shaked, and Sutton 1991.

that Binmore et al. call "deal me out." In the forced breakdown game, the prediction is that the players will each get their breakdown payment plus half the surplus 1 – (α + β), so that player 2 will receive 66 percent in the low beta game and 80 percent in the high beta game, a pattern that Binmore et al. call "split the difference."

Following several plays of each game in which each subject participated as both player 1 and player 2 in one of these games, Binmore et al. presented subjects with a questionnaire asking them to indicate what they thought was a "fair" offer to player 2 for each configuration of breakdown payoffs. (Everyone agreed that 50 percent was fair when breakdown payoffs were equal.) The responses for low and high values of β are graphed in Figure 4.5a for players who experienced the optional breakdown games, and in Figure 4.5b for players who had experienced the forced breakdown games.

It is the difference between the two figures that is especially notable. While some subjects in each kind of game propose that a fair outcome gives player 2 only 50 percent for both low and high beta situations, most subjects think it fair that player 2 should get more than 50 percent in both situations. But while the responses from subjects who had experienced the optional breakdown game are diffuse (Figure 4.5a), the responses of those who had experienced the forced breakdown game are much less so, and are concentrated around the split-the-difference numbers, that give 66 percent to player 2 when β is low and 80 percent when β is high. Thus many subjects' ideas about what constitute "fair" agreements have been influenced by the version of the game they played, and in the forced breakdown game their experience of a strategic environment in which split-the-difference is equilibrium behavior has led to their adopting it as their idea of a fair outcome.[24]

In summary, the evidence from all these sequential bargaining games suggests that some of the away-from-equilibrium behavior (e.g., disadvantageous counteroffers, equal divisions in dictator games) results from bargainers' preferences that concern not only their own income but also their relative share. At the same time, the evidence suggests that much of the away-from-equilibrium behavior does not have such a simple cause, but results instead from strategic considerations (including the anticipation that some offers may not be accepted because of fairness considerations). Despite the heterogeneity of bargainers' motivations, much of the observed behavior contains clear, reproducible regularities (as in Ochs and Roth [1989], Weg et al. [1990], and Bolton [1991]), which indicate that the deviations from the equilibrium predictions reflect systematic features of the bargaining environment. Yet there are also some anomalies (e.g., similar games with different observed behavior) that present the opportunity to design experiments to test different hypotheses about the nature of these systematic features. Some of these are considered next.

4. Pursuing Anomalies

In concentrating on sequential bargaining games in which the players take turns making offers of how to divide a diminishing pie, we have considered a family of games in which the perfect equilibrium and its usual auxiliary assumptions (e.g., that the bargainers are motivated by simple income maximization) yield notably poor predictions. Yet there are other, related games in which the perfect equilibrium predictions perform much better, such as the fixed-cost sequential bargaining games reported by Rapoport et al. (1990). These particular games are difficult to compare with the other sequential bargaining games so far studied, because of the different cost structure they employ, the much greater number of periods they were allowed to run, and the different equilibrium predictions. Yet the similarities between these games are sufficient so that the different success of the equilibrium predictions demands further investigation. We turn next to consider this kind of anomaly.

The simplest of the sequential bargaining games is the ultimatum game, in which each player makes only one decision. The results of the ultimatum game experiments are clear: observed behavior is far from the equilibrium prediction. One way to explore this phenomenon more fully would be to identify a game with closely parallel structure and equilibrium predictions, but in which observed behavior would conform to the equilibrium prediction. Prasnikar and Roth (1989, 1992) identified such a game and conducted an experiment designed to compare it with the ultimatum game.

The game in question was earlier studied by Harrison and Hirshleifer (1989), as one of several games in an experiment concerned with different mechanisms for the provision of public goods. In one of the games, player 1 first proposed a quantity q_1 that he would provide, then player 2 (after being informed of q_1) proposed a quantity q_2 that he would provide, with the quantity q of public good provided being the *maximum* of q_1 and q_2 (the "best shot"). Both players were then paid a "redemption value" based on the quantity q provided; however each player i was charged for the quantity q_i that *he* had provided, at a flat rate of 82¢ per unit (see Table 4.2).

The perfect equilibrium predictions are that player 1 will choose $q_1 = 0$ and player 2 will choose $q_2 = 4$, giving player 1 a profit of \$3.70 and player 2 a profit of \$3.70 − \$3.28 = \$0.42. Harrison and Hirshleifer (1989) conducted an experiment in which best shot games were played under conditions of partial information, in which each player was unaware that his counterpart had the same costs and redemption values. They observed results that were strikingly close to the perfect equilibrium predictions.

One hypothesis is that players' lack of information about each others' payoffs may have disabled whatever countervailing force in favor of more equal distributions of payoffs was at work in the bargaining games reported above. That is, perhaps the reason subjects in the role of player 2 were willing to accept a payoff of \$0.42 was because they were unaware (or unsure) that player 1 was receiving \$3.70, in contrast to the case of ultimatum bargaining games in which such extreme payoff disparities proved to be unacceptable. (Guth's [1988] theory of hierarchical social norms, accessed according to the information available, would presumably account for the results in this way.) This could potentially explain why such a relatively extreme distribution of payoffs was observed in this data, but virtually never in the data from ultimatum games for comparable amounts of money.[25]

The experiment of Prasnikar and Roth (1989, 1992) was designed to investigate both this hypothesis and the hypothesis that the difference between observed behavior in best shot and ultimatum games (despite their similar equilibrium predictions) was due to the different incentives these games gave to players *off* the equilibrium path. To this end, best shot games were examined both under partial information, as in Harrison and Hirshleifer, and under full information, with both players knowing each other's payoffs. In addition, in order that other details of procedure should not complicate the comparisons, a set of ultimatum games were conducted using the same detailed procedures (of recruiting subjects, of transmitting messages, etc.). The best shot games under partial and full information conditions provide a test of the hypothesis that the extreme equilibrium payoffs will only be observed when subjects cannot compare their payoffs, while the comparison of the full information best shot game with the ultimatum game run under the same conditions provides a test of the hypothesis that the structural

Table 4.2. Redemption Values and Expenditure Values for the Best-Shot Games

Redemption Values			Expenditure Values	
Project Level (Units)	Redemption Value of Specifiic Units	Total Redemption Value of All Units	Number of Units You Provide	Cost to You of the Number Units You Provide
0	$0.00	$0.00	0	$0.00
1	1.00	1.00	1	0.82
2	0.95	1.95	2	1.64
3	0.90	2.85	3	2.46
4	0.85	3.70	4	3.28
5	0.80	4.50	5	4.10
6	0.75	5.25	6	4.92
7	0.70	5.95	7	5.74
8	0.65	6.60	8	6.56
9	0.60	7.20	9	7.38
10	0.55	7.75	10	8.20
11	0.50	8.25	11	9.02
12	0.45	8.70	12	9.84
13	0.40	9.10	13	10.66
14	0.35	9.45	14	11.48
15	0.30	9.75	15	12.30
16	0.25	10.00	16	13.12
17	0.20	10.25	17	13.94
18	0.15	10.35	18	14.76
19	0.10	10.45	19	15.58
20	0.05	10.50	20	16.40
21	0.00	10.50	21	21.22

differences in the games make the extreme equilibrium payoffs more likely to be observed in one than in the other. Each subject played only one of the three games, but played it ten times, against a different anonymous opponent in each round. So the experiment also allows the learning that goes on in each game to be compared.

Table 4.3 reports the mean offers x_2 in the ultimatum game, as well as the mean quantities q_1 provided in the sequential best shot games under full and partial information. Recall that the perfect equilibrium prediction is that all these quantities will be zero. The observed means are reported round by round for each game.[26]

In the sequential best shot game under full information the observed means clearly have converged to the equilibrium quantity by the seventh round, after

Table 4.3. Mean Offers by Periods

Periods	Ultimatum Game (mean x_2)[a]	Best Shot, full information game (mean q_1)[b]	Best Shot, partial information game (mean q_1)[c]
1	4.188 (0.329)	1.625 (0.610)	2.700 (0.617)
2	3.825 (0.530)	0.875 (0.482)	2.900 (0.994)
3	3.725 (0.480)	1.125 (0.597)	3.000 (0.848)
4	3.581 (0.438)	0.125 (0.116)	2.100 (0.793)
5	4.231 (0.276)	0.125 (0.116)	2.700 (0.906)
6	4.418 (0.234)	0.125 (0.116)	1.250 (0.605)
7	4.294 (0.166)	0.000 (0.000)	1.100 (0.537)
8	4.531 (0.155)	0.000 (0.000)	0.800 (0.505)
9	4.325 (0.232)	0.000 (0.000)	0.950 (0.567)
10	4.531 (0.155)	0.000 (0.000)	0.700 (0.401)

Source: Prasnikar and Roth 1992.
Note: Values in parentheses are standard errors.
[a] Perfect equilibrium prediction: $x_2 = 0$.
[b] Perfect equilibrium prediction: $q_1 = 0$.
[c] Perfect equilibrium prediction: $q_1 = 0$.

which no player 1 is observed to provide any positive quantity. (And the modal response of players 2 is the equilibrium response of $q_2 = 4$, with 41% of offers $q_1 = 0$ receiving this response overall.) Although the results in the partial information best shot games are significantly different from those in the full information games, the observed means in both best shot games are clearly much closer to zero than are the observed means in the ultimatum games, which are quite similar to the observations for ultimatum games previously reported in the literature. So the best shot game is one in which, even when the players can compare their payoffs, equilibrium payoffs can be observed even though they are extreme.

An indication of how the best shot games are different from the ultimatum games comes from examining the learning that took place over the course of the

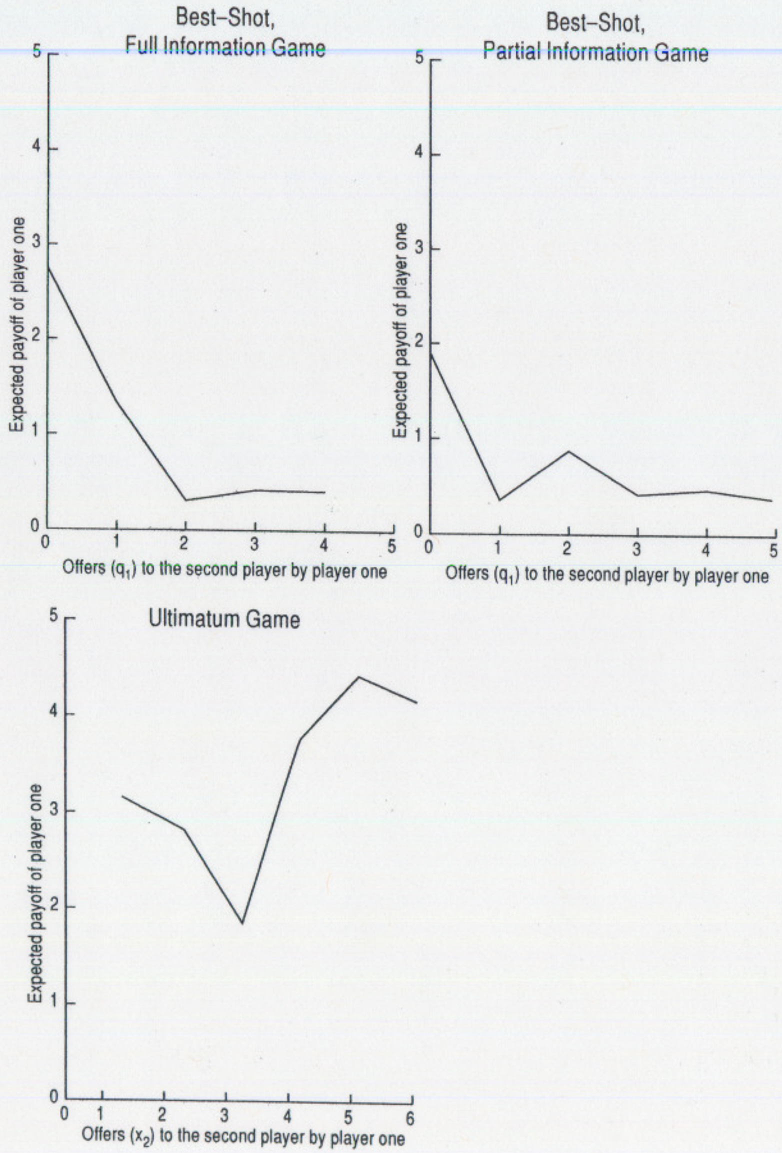

Figure 4.6. Expected payoff of each offer. *Source:* Prasnikar and Roth 1992.

ten rounds each game was played. In the best shot games, half the player 1s in the full information condition and nine out of ten of the player 1s in the partial information condition began by offering positive quantities, but in the face of consistent lack of a positive reply by the player 2s, the number of player 1s offering positive quantities steadily diminished. In contrast, in the ultimatum game, offers by the first player were closest to the equilibrium prediction in the first four rounds, but in the face of steady rejections of low offers, the lowest offer x_2 steadily climbed.

Prasnikar and Roth concluded that the difference between the observed behavior in best shot and ultimatum games, despite the similarity of their equilibrium predictions, is in the off the equilibrium path behavior. This can be assessed by considering how player 2s react when player 1s deviate from the equilibrium prediction—i.e., when they offer $q_1 > 0$ in the best shot games or $x_2 > 0$ in the ultimatum games. The prediction of subgame perfect equilibrium is in all cases that player 1 will maximize his payoff by making the equilibrium offer—i.e., at perfect equilibrium; the predicted response of player 2 is such that a positive offer will yield player 1 a lower payoff than an offer of zero. However, as the graphs in Figures 4.6a, b, and c make clear the best shot games exhibit different behavior in this regard than the ultimatum games. In the best shot games, under both information conditions, the average payoff of player 1s who contributed the equilibrium quantity $q_1 = 0$ is greater than that of player 1s who contributed positive quantities. However, in the ultimatum game, the average payoff to a player 1 who offers player 2 an amount x_2 rises to a maximum for x_2 between 4 and 5 (which is where we observe the mean offer). So in the ultimatum games a player 1 does better as he deviates further from equilibrium, but not in the best shot games. This behavior is comprehensible in both games: in the best shot games, the more player 1 provides, the less incentive player 2 has to provide anything (see Table 4.2), while in the ultimatum games, the more player 1 offers to player 2, the greater his incentive to accept the offer.

To the extent that this explanation identifies the important difference between best shot and ultimatum games, it suggests that whatever role considerations of fairness may play in such games, it is mediated by considerations of strategy of the kind that game theorists are accustomed to studying. But quite a different hypothesis has also been proposed.

a. Distinguishing between Alternative Hypotheses

After an early version of these results was circulated (Prasnikar and Roth 1989), Guth and Tietz (1990) suggested an alternative hypothesis to explain them. They say:

> [E]qual positive contributions in best shot games are obviously inefficient since one of the two contributions is completely useless. If sharing the burden of providing the public good is impossible, fairness considerations cannot be applied. (428)

These comments refine Guth and Tietz's (1988) hypothesis concerning the role played by extreme payoff distributions (recall footnote 25), by adding considerations of convexity and efficiency. In doing so, they raise a clear counterhypothesis to the interpretation given above to the observed differences between best shot and ultimatum games. According to our interpretation, the different off-the-equilibrium-path properties of the two games is responsible for the different observed behavior, despite the comparably unequal payoff distribution at equilibrium. The contrary hypothesis suggested by Guth and Tietz (1990) is that the different observed behavior in the two games is due to the fact that players are concerned with fairness only in the ultimatum games, and that no comparable considerations arise in the best shot games because in those games equality and efficiency are incompatible.

To examine these competing hypotheses, Prasnikar and Roth examined a sequential market game, consisting of one seller and nine buyers. As in the ultimatum games, each buyer offered a price, which if accepted determined the division of $10 between the successful buyer and the seller. (If the seller accepts an offer p from buyer 1, then that buyer earns $10 – p, the seller earns $$p$, and all other buyers earn $0. If the seller rejects all offers, then all players in the market receive $0.)[27] Since the smallest unit in which prices could be quoted was $0.05, there are two subgame perfect equilibrium prices, $9.95 and $10.00. Thus any subgame perfect equilibrium gives (virtually) all the wealth to the seller.[28]

In this game all transactions, not merely equilibrium transactions, are efficient. Thus this game has equilibrium payoff distributions that are as extreme as those of the ultimatum or best shot games, but (like the ultimatum game and unlike the best shot game) it has efficient equal-payoff outcomes.[29] It therefore presents an opportunity to test the conjecture that the observed outcomes of the best shot games were intimately related to the fact that equal payoffs in that game can only be achieved inefficiently.

The observed results do not support this hypothesis. By the fifth round, prices had converged to equilibrium, and all subsequent transactions were at the equilibrium price of $10.[30] (See Table 4.4.) Table 4.4 also makes clear that (except in round 7 in market B) from round 5 on no buyer could have increased his payoff by more than $0.05 by changing his bid. The high bidders in these rounds (who always received zero) were always competing either with another bidder who made the same bid, or one who made a bid that was only $0.05 less.[31] Thus in this game, as in the best shot game and in contrast to the ultimatum game (recall Figures 4.6 a–c), the observed pattern of play is such that agents could not increase their payoff by deviating from the equilibrium prediction. This lends further support to the hypothesis that the off the equilibrium path behavior is of critical importance in understanding the observed behavior in these games. (This hypothesis is also supported by the data from the "infinite horizon" bargaining games of Rapoport, Weg, and Felsenthal [1990] and Weg and Zwick [1991], in which equally extreme equilibrium predictions were achieved with some regularity.) Note that this certainly does not mean that considerations of fairness do not play a role in determining the outcome of the game, but rather that such considerations interact with the strategic features of the game.[32]

Table 4.4. The Highest and Second-Highest Prices in Each of the Markets and the Basic Descriptive Statistics

Period	Market	Highest Price ($)[a]	Second-Highest Price ($)[a]	Mean and SD[b]	Mode[c]	Median	N[d]
1	A	8.90 (1)	8.25 (1)	6.48 (2.52)	8.05	8.05	9
	B	9.90 (1)	8.95 (1)	6.76 (1.84)	5.00	6.50	9
2	A	9.60 (1)	9.00 (1)	6.57 (3.07)	5.00	8.05	9
	B	9.90 (1)	9.00 (2)	6.69 (3.26)	x	8.00	9
3	A	9.85 (1)	9.65 (1)	7.24 (3.24)	x	9.00	9
	B	10.00 (1)	9.95 (1)	8.08 (2.31)	x	9.00	9
4	A	10.00 (2)	9.95 (2)	7.32 (4.00)	x	9.90	9
	B	9.95 (1)	9.90 (1)	7.31 (2.67)	9.00	9.00	9
5	A	10.00 (2)	9.95 (2)	9.14 (1.61)	x	9.90	9
	B	10.00 (2)	9.95 (2)	7.93 (2.76)	x	8.50	9
6	A	10.00 (3)	9.95 (1)	7.21 (3.69)	10.00	9.00	9
	B	10.00 (1)	9.95 (4)	7.81 (3.32)	9.95	9.95	9
7	A	10.00 (1)	9.95 (2)	6.43 (3.28)	x	7.00	9
	B	10.00 (1)	9.60 (1)	5.23 3.07	5.00	5.00	9
8	A	10.00 (2)	9.85 (1)	5.76 (3.74)	x	5.00	9
	B	10.00 (2)	9.85 (1)	5.72 (4.31)	x	7.00	9
9	A	10.00 (1)	9.95 (1)	4.73 (4.11)	x	5.00	9
	B	10.00 (1)	9.95 (1)	5.98 (3.72)	x	5.00	9
10	A	10.00 (2)	9.95 (1)	6.22 (4.23)	x	9.00	9
	B	10.00 (2)	9.95 (1)	6.47 (3.32)	5.00	5.00	9

Source: Prasnikar and Roth 1992.
[a] The number in parentheses is the number of buyers who bid that price.
[b] Numbers in parentheses are standard deviations.
[c] An "x" in the mode column means that there were fewer than three observations at any one price.
[d] N represents the number of buyers in each of the markets.

b. A Cross-Cultural Experiment

It turns out that this behavior is quite robust to changes in subject pools. Roth, Prasnikar, Okuno-Fujiwara, and Zamir (1991) conducted an experiment in which this kind of market game was examined, together with an ultimatum bargaining game, in Jerusalem, Ljubljana, Pittsburgh, and Tokyo.[33] (In each environment, subjects gained experience in ten consecutive transactions with different players, as in Prasnikar and Roth [1992].) Outcomes in the market game converged to equilibrium in all four locations, while outcomes in the ultimatum game remained far from the equilibrium prediction, although there were differences between the bargaining outcomes in different countries. Before discussing these results, it is worth spending a moment on some of the problems of experimental design that arise in conducting an experiment in four countries, namely the problems of controlling for the effects of different experimenters, languages, and currencies. The discussion of these aspects of the experimental design will be organized as a statement of a particular problem, followed by the element of the design that addressed this problem.

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

Design solution Each experimenter came to Pittsburgh and ran (at least) a bargaining session and a market session. The Pittsburgh data were thus gathered by all of the experimenters before they returned to their home countries to gather the data there. In this way we were able to coordinate the detailed operational procedures among the different experimenters. And the Pittsburgh data can be used to detect any effect due to purely personal characteristics of the experimenters, since if these effects were present they would have shown 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 were presented in English, Hebrew, Japanese, and Slovenian, systematic differences between countries might be observed because of the way the instructions are translated. (Consider, for example, the English words "bargaining," "negotiating," and "haggling," which are all approximate synonyms, but whose different connotations might possibly elicit differences in behavior.)[34]

Design solutions The problem of language effects was addressed both through the way in which the translations were made and, more formally, in the way the instructions for the bargaining and market environments were related.

1. *Translations*: The experimenter responsible for each translation was a national of the country in question who is both linguistically and culturally fluent in American English. 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, subjects who made proposals were referred to as "buyers," while those who made acceptances or rejections were termed "sellers.") So 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 we observed—no between-country differences in the market behavior, but differences in the bargaining behavior—at least put an upper bound on the effect of the translation and establish that it is not large enough to cause the markets to yield different results in the different countries. This supports the hypothesis that the translation is not the cause of the observed difference in the bargaining.

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., Albers and 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.[35])

Design solutions First, 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. So 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. Second, 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 5 tokens.

Of course, there remain many uncontrolled differences between subject pools. For example, in Israel and Slovenia a much higher percentage of the sample of subjects were army veterans than in the United States or Japan. So any conclusions about the *causes* of between-country differences have to be circumspect.

Figure 4.7a summarizes the market results from Slovenia, Japan, and Israel, and shows that they are quite similar to those observed in the United States: offers

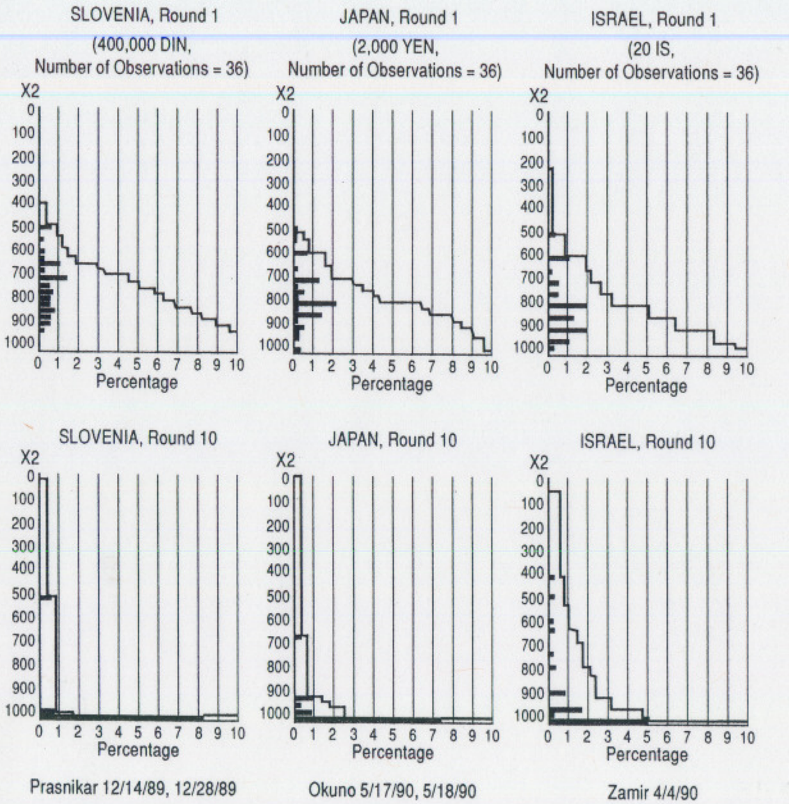

Figure 4.7a. Distribution of market offers in Slovenia, Japan, and Israel. *Source:* Roth, Prasnikar, Okuno-Fujiwara, and Zamir 1991.

in round 1 are diffuse, but by round 10 there is a concentration of offers at the equilibrium prices. So by the tenth round there were no payoff relevant between-country differences observed in these markets.

The situation is different for the ultimatum bargaining game. In each of the countries the modal offer in the first round is the equal division offer of 500. And in none of the countries do the tenth round offers approach the equilibrium prediction of zero. But by the tenth round the distributions of offers are significantly different in different countries. In the United States and Slovenia the modal offer in the tenth round remains at 500, as in the first round. But in Japan, the tenth-round offers exhibit modes at 450 and 400, and in Israel the mode is at 400.[36] Figure 4.7b shows the first- and tenth-round distributions for Slovenia, Japan, and Israel (the distribution observed in the U.S. is similar to that of Slovenia).

These between-country differences offer the opportunity to examine what other features of the bargaining outcomes vary together with the distribution of offers. We can anticipate section II of this chapter by considering how the rate of acceptances and rejections varied between countries. Such a comparison can be made by considering how often the proposal of a given price is accepted. These compar-

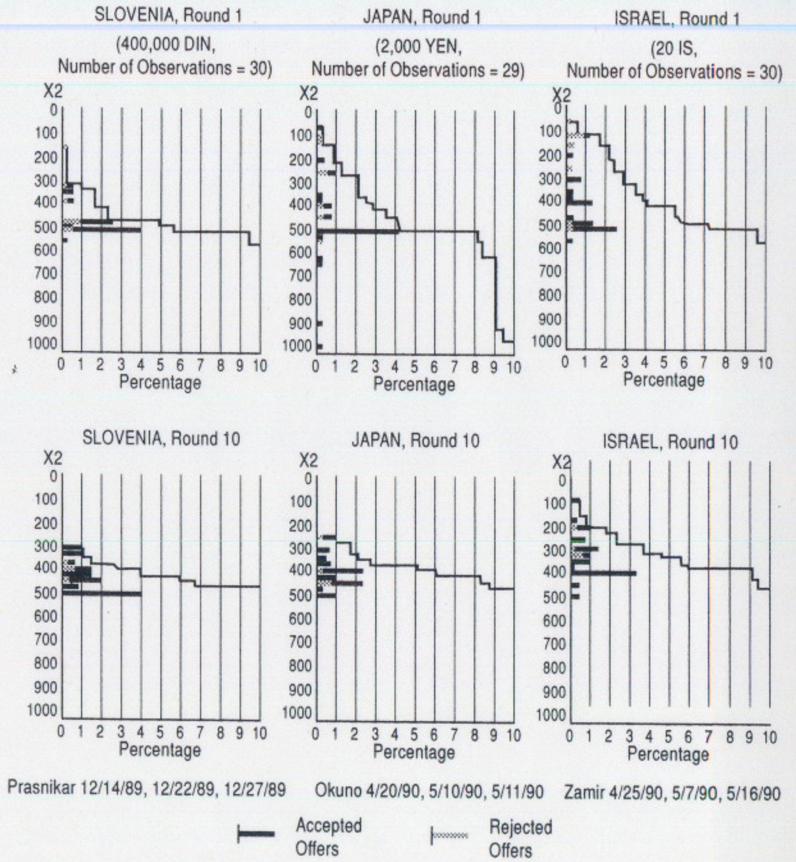

Figure 4.7b. Distribution of bargaining offers in Slovenia, Japan, and Israel. *Source:* Roth, Prasnikar, Okuno-Fujiwara, and Zamir 1991.

isons are slightly complicated by the fact 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. However, the underlying pattern is clear, as is demonstrated by Figure 4.7c. The curves for each country represent the percentage of acceptances for each price that was proposed at least ten times (over all rounds). Each cell of Figure 4.7c 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 that in the United States, Slovenia, and Japan, respectively, while the acceptance rates in Japan are higher than those in the United States and Slovenia. Only in the comparison of the United States and Slovenia (whose distributions of observed offers did not differ significantly) do we have two acceptance rate curves such that one is not higher than the other.

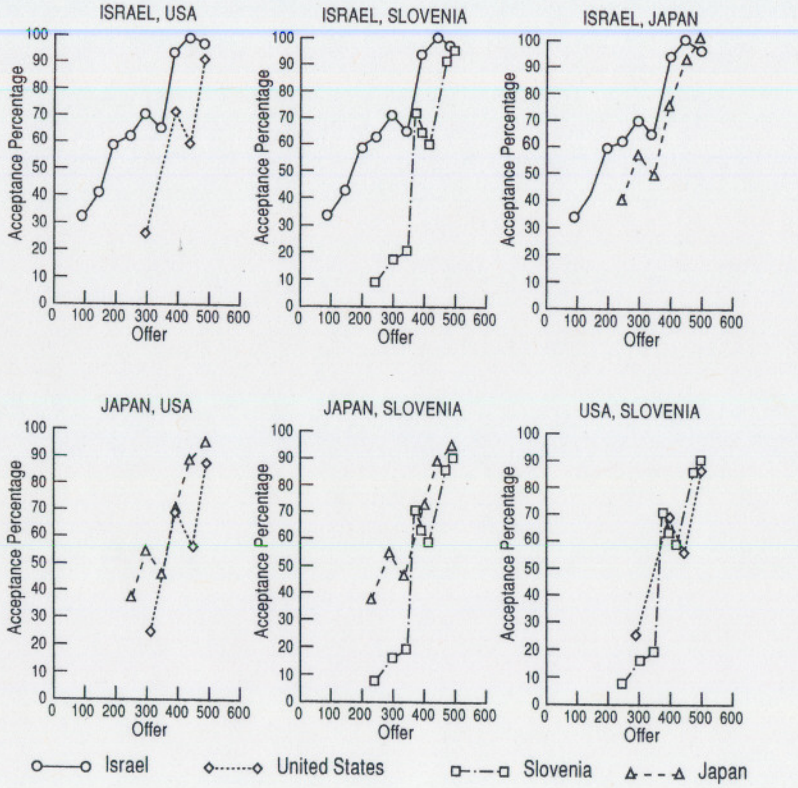

Figure 4.7c. Pairwise comparisons of acceptance rates in bargaining. *Source:* Roth, Prasnikar, Okuno-Fujiwara, and Zamir 1991.

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 is changing from round to round, this is something of a moving target. But Figure 4.7d 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 fifteen times and accepted four times (26.7 percent), so on average the proposal earned $(700 \times .267) = 186.9$, which can be read from the graph for the United States in Figure 4.7d. It is instructive to compare these graphs to the modal offers observed in round 10 in each country. The modal offer in the final round in both the United States and Slovenia is 500, and looking at Figure 4.7d we see that 500 is also the proposed price that maximizes a buyer's average earnings in the United States and Slovenia. The modal offer in the final round in Israel is 400, and we see in Figure 4.7d that here too this is the price that maximizes a buyer's average earnings. And in Japan there are two modal offers in round 10, at 400 and 450, and the latter maximizes a buyer's

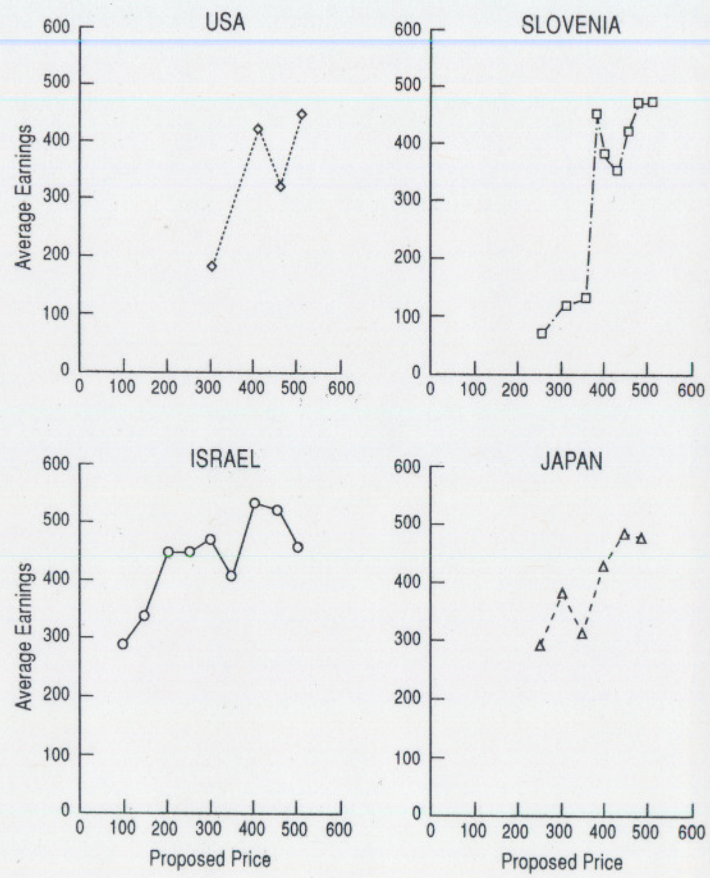

Figure 4.7d. Buyers' earnings in bargaining, by Proposed Price. *Source:* Roth, Prasnikar, Okuno-Fujiwara, and Zamir 1991.

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.)

These results bring us back to the question of the previous section, "Are players trying to be fair?" Taken together, Figures 4.7b and 4.7d suggest that in the ultimatum games the behavior of the player 1s, at least, may be well accounted for by simple income maximization. But the fact that the modal offer in round 10 was in each case the income maximizing offer means that player 1s are able to take into account the behavior of the player 2s, which included (non-income-maximizing) rejections of positive offers even in the final round.

Thus the evidence provides some support for the conjecture that, while efforts to *be* fair do not play a large role in the observed behavior, efforts to avoid being *treated* unfairly oneself may influence the data: that is, the player 2s who reject low positive offers apparently prefer to do so rather than accept an offer that seems to them unreasonably low. To the extent that this is the case, 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. So under this hypothesis, high rates of disagreement would be associated with subject pools in which offers are low. This is *not* what we observe.[37] (Another way to make the point is to note that the question of in which country the bargainers proved to be the toughest is not a well posed question, in the sense that the "toughest" buyers are found in the same place as the least tough sellers.)

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 what varies between subject pools is how much more seems reasonable. To the extent that offers tend towards what is commonly regarded as reasonable and 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 lends 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.

5. Learning and the Role of Experience

Tables 4.3 and 4.4 illustrate that experience plays an important role and that its effect may be different in different games. In the best shot and market games, experience brings the outcome more in line with the equilibrium predictions, but this is not at all the case in the ultimatum games, any more than it was in the two and three period sequential bargaining games studied by Ochs and Roth and summarized in Figure 4.2. And in the four country experiment just discussed, experience brought the market behavior in different countries closer together, while the bargaining behavior became more distinct as the players gained experience (Figures 4.7a,b).

A number of other experiments speak to the role of experience and also to how the process of accumulating experimental evidence allows simple hypotheses to be reformulated and refined. Recall from section I.B.1 that Binmore, Shaked, and Sutton (1985) found that the modal first-period offers in a two-period game were equal divisions the first time the game was played, but that when players 2 were told that they would be players 1 in a second game, their modal first-period offers were at the perfect equilibrium division. They summarized their views at the time as follows (1180):

> While we have considered various possible explanations, the interpretation that we favor is this: subjects, faced with a new problem, simply choose "equal division" as an "obvious" and "acceptable" compromise—an idea familiar from the seminal work of Thomas Schelling (1960). We suspect, on the basis of the present experiments, that such considerations are easily displaced by calculations of strategic advantage, once players fully appreciate the structure of the game.

More recently, the study of Binmore et al. (1991), discussed in section I.B., allowed optional breakdown games played by subjects with several periods of experience to be compared with optional breakdown games played by subjects with only a single period of experience (but with different breakdown payments) reported in Binmore et al. (1989). Binmore et al. (1991) note that this is a case in which the increased experience seems to *decrease* the ability of the equilibrium prediction to organize the data, a phenomenon they characterize as "unlearning." They state (311):

> For the moment, the only safe conclusion would seem to be that, if people are indeed "natural gamesmen" (a view that has been wrongly attributed to us in the past), then experience in this context would appear to lead to some "unlearning" of their game-playing skills.

Of course, different kinds of experience might produce different effects. For example, Binmore et al. (1985) already observed that the simple opportunity to play a game a second time would not be sufficient to cause offers to converge to the equilibrium prediction, since (recall from section I.B.1) Guth et al. (1982) had observed ultimatum games played by players who had experienced a previous game, without finding any difference between the two plays analogous to that observed by Binmore et al. This suggested that perhaps Binmore et al.'s (1985) procedure of having players 2 in the first game become players 1 in the second game gave them an especially appropriate kind of experience for the task of anticipating the strategic situation facing the second player. But recall that when this hypothesis was tested (e.g., by Guth and Tietz [1987], and Bolton [1990]) it was found that reversal of roles in this way did not generally promote convergence to the equilibrium predictions.

Harrison and McCabe (1991) report an experiment designed to help reconsider the question of whether "appropriate" experience might promote convergence to equilibrium behavior. In their experiment, bargainers play a sequence of alternating offer sequential bargaining games, switching back and forth between a three-period game and a two-period game. In the three-period game, the amounts to be divided in the three periods are in the proportions 100, 50, and 25, while in the two-period game the amounts are 50 and 25, so the two-period game is the subgame of the three-period game that arises if the first proposal is rejected. The perfect equilibrium proposal in the two-period game is the equal division (25, 25), so the perfect equilibrium in the three-period game calls for player 1 to propose the division (75, 25). Harrison and McCabe observe that equal divisions are observed immediately in the two-period games, and that in the three-period games average initial proposals move from near (50,50) in the direction of the perfect equilibrium proposal, approaching (70, 30) by the end of fifteen rounds of alternating between the two games.

In interpreting their results, Harrison and McCabe say (13–4):

> Our experiments and those of OR [Ochs and Roth] are the only ones to give subjects any length of experience in terms of more than two repetitions of the game. OR focus on the pure role of experience, and find that the evidence is mixed for two-round games and decidedly negative for game theory in three-round games. . . . In other words, experience *per se* does not appear to be a reliable basis for the formation of common expectations [of achieving the perfect equilibrium].

But they go on to say:

> When one combines experience *and* sequencing the conclusion is a simple one and is perfectly consistent with the predictions of game theory.

Spiegel, Currie, Sonnenschein, and Sen (1990), however, report a related experiment in which subjects play a sequence of five alternating offer sequential bargaining games against different opponents, starting with a one period (ultimatum) game and working their way up to a five-period game. The games have the property that if the first proposal is rejected in any game, the subgame that arises is the game that was played just previously. Thus prior to playing the five-period game players have the experience of playing each of its subgames. Spiegel et al. observe no tendency for proposals in any of their games to approach the perfect equilibrium prediction. This is not too surprising in view of the fact that, as each game is begun, the experience that subjects have of its subgames (starting with the ultimatum game) is different from the perfect equilibrium prediction. But it serves to emphasize that prior experience of a game's subgames does not necessarily promote equilibrium behavior. And the comparison with the experiment of Harrison and McCabe (1991) is illuminating, since in Harrison and McCabe's experiment the two period subgame had a subgame perfect equilibrium that coincides with the "equal split" solution seen so often in experimental data for reasons having nothing to do with subgame perfection. So the different results of the two experiments suggest that it may be necessary to experiment with a range of parameter values before attempting to draw firm conclusions.

In a subsequent experiment with a game derived from the ultimatum game, Harrison and McCabe (1992b) had subjects each simultaneously submit both offers and acceptance/rejection strategies for an ultimatum game, knowing only that they had an equal chance of being assigned the role of player 1 or player 2. (Notice that under this design, *any* symmetric behavior gives both players an equal expected payoff.) In two of the cells of their experiment they attempted to condition the expectations of their subjects with a design loosely modeled on that of Roth and Schoumaker (1983), by exposing the subjects to the play of automated robots playing near equilibrium strategies, and/or by making available other subjects' past histories (including the history of their strategy choices).[38] They report that in the cells in which expectations were conditioned (but not in the unconditioned control condition) the strategies evolve in the direction of the perfect equilibrium.

Overall, the data suggest that both the kind of prior experience and the kind of game that is experienced influence the way in which behavior changes with experience. While different investigators still maintain different points of view on these matters, in the course of these experiments these differences have been substantially narrowed. In this respect, the progress that has been made in experiments concerned with experience and learning reflects the progress that has been made on understanding other aspects of bargaining behavior, in the course of interchange among experimenters with varying points of view.

This is also a case in which some of the questions left open by the experimental evidence suggest further theoretical work. In this vein, Roth and Erev (1995) explore a family of simple models of adaptive learning (recall the discussion in section III.B of chapter 1).[39] A free variable in this kind of model is the choice of initial conditions, which determine the probabilities with which strategies are played the first time that the game is played, before players have had a chance to start accumulating experience. Roth and Erev find that the learning behavior in the ultimatum game is much more sensitive to the initial conditions than is the behavior in the best shot or market games studied in Prasnikar and Roth (1992) and Roth, Prasnikar, Okuno-Fujiwara, and Zamir (1991). For a wide range of initial conditions, behavior predicted by this model in the best shot and market games converges to the perfect equilibrium prediction. However, this is not the case with the ultimatum game, which over most of this range converges to equilibria at which player 1 offers a nonnegligible percentage of the pie to player 2.[40] Only when the initial conditions are relatively close to the perfect equilibrium does the behavior converge to perfect equilibrium in the ultimatum game.[41]

Roth and Erev go further and simulate the four-country experiment of Roth et al. (1991) using this kind of learning rule. When the initial conditions are as observed in the experiments, the qualitative results of the simulation track those of the experiment. That is, when the initial propensities to play each strategy are estimated from the data observed in the first-period play in Jerusalem, Ljubljana, Pittsburgh, and Tokyo, then the simulated behavior converges to what was observed in the tenth period of the experiment. The market game results quickly approach perfect equilibrium for all four simulations, while in the simulated ultimatum game offers remain at 50 percent when begun with the initial propensities observed in Ljubljana and Pittsburgh, while offers move to 40 percent when begun with the initial propensities observed in Tokyo and Jerusalem. These results thus suggest the conjecture that the differences in ultimatum game results observed in those places may be due primarily to the different behavior of the subjects in the first period. Similarly, when the best-shot game is simulated, starting with the initial propensities observed in the full and partial information conditions of Prasnikar and Roth (1992), the full information simulation approaches perfect equilibrium faster than the partial information simulation, as in the experiment. That is, the simulations suggest that the experimental results are consistent with the hypothesis that subjects used essentially the same learning rules in all games in all locations and that the observed differences in bargaining behavior reflect different initial expectations.[42]

Of course, theoretical results do not begin to prove that this was the case, but they do suggest further experiments to distinguish between this and competing hypotheses. This is what we should expect from the interaction between theory and experiment.[43]

In closing, one aspect of a bargainer's experience that obviously has the potential to affect his future behavior is his experience with rejected proposals and other forms of disagreements and costly delays in reaching agreement. The second part of this chapter concerns such disagreements and delays.

II. Disagreements and Delays

A. *The Frequency of Disagreements and Delays*

One of the facts about bargaining that has struck empirical investigators of all sorts is that a nonnegligible frequency of disagreements and costly delays seems to be ubiquitous. While this would be unsurprising if it occurred only in situations that presented the bargainers with no mutually profitable agreements, most of the evidence suggests that disagreements and costly delays are pervasive even when it is evident that there are gains to be had from agreement. Kennan and Wilson (1990a) summarize this by observing that an element of regret is characteristic of bargaining: they note that most strikes, for example, are eventually settled on terms that could have been reached earlier, without incurring the costs that the strike imposes on all parties.

The game-theoretic models most often used to explain an irreducible frequency of disagreements and delays are models of incomplete information. In these models, bargainers are uncertain about some important features of one anothers' situations, which, if known, would influence the distribution of profits between the parties. The bargainers in these models convey information to one another about their situations by their willingness to risk disagreement or to tolerate delay, in order to influence how the profits are divided.

While these models have much to recommend them, experimental evidence suggests that disagreements are pervasive even in situations (such as those discussed in the first part of this chapter: recall Table 4.1) that eliminate the most obvious potential sources of incomplete information. Forsythe, Kennan, and Sopher (1991), for example, tabulate the data from the different conditions in a number of such experiments as in Table 4.5a.

The frequencies of disagreement and costly delays observed in these experi-

Table 4.5a. Frequency of Disagreements and Delays

Study	Frequency of Inefficient Outcomes (%)
Malouf and Roth (1981)	0–37
Roth and Murnighan (1982)	8–33
Guth, Schmittberger, and Schwarze (1982)	10–29
Binmore, Shaked, and Sutton (1985)	19–42
Neelin, Sonnenschein, and Spiegel (1988)	5–35
Ochs and Roth (1989)	10–29
Forsythe, Kennan, and Sopher (1990a)	19–67

Source: Forsythe, Kennan, and Sopher 1991, 267.

Table 4.5b. Reported Dispute Rates

Study	Data	Rate (%)
Collective Bargaining: Strikes		
Card (1988)	Canadian private sector contracts	22
McConnell (1989)	U.S. private sector contracts	17
Currie and McConnell (1989)	Canadian public sector contracts	13
Collective Bargaining: Arbitrations		
Currie and McConnell (1989)	Canadian public sector contracts	32
Currie (1989)	British Columbia teachers	33
Ashenfelter and Bloom (1984)	New Jersey police	30–49
Ashenfelter et al. (1992)	Arbitration experiments	28–43
Boden (1989)	Worker's compensation	43
Other Types of Negotiations		
Ochs and Roth (1989)	Two-person bargaining games	15
Mnookin et al. (1989)	California child custody	22[a]
White (1989)	Medical malpractice	11[a]

Source: Ashenfelter and Currie 1990, 416.
[a] Percentage of cases that go to court.

ments are not so different from those observed in a variety of field studies. Ashenfelter and Currie[44] summarize the situation as in Table 4.5b, which includes both field studies and experiments.

Other researchers agree with the rough magnitudes of the figures for field studies.[45]

The frequency of disagreement in "complete information" experiments, and the similar rates of disagreement found in other experiments and in field studies, raises some question about whether the incomplete information models are focusing on the underlying causes of disagreement. There are various approaches to answering this, and we will come back to these models in section IIC. However, before discussing the various experiments that have been conducted to explicitly examine the causes of disagreement and the effects of incomplete information, there is an empirical issue that remains to be settled, since reviewing the literature reveals an anomaly that has as yet been only incompletely explored.

While the experiments referred to above consistently report a nonnegligible frequency of disagreement and delay in the same rough range as do various field studies, these experiments have all been conducted so that the bargainers interact anonymously. In contrast, a number of bargaining experiments have been reported in which the bargainers deal with one another face to face, and these tend to report a much lower rate of disagreement. That is, the frequency of disagreements observed under face-to-face bargaining is substantially lower than that observed under anonymous bargaining, and investigators have tended to generalize from these different bodies of data in quite different directions. We consider this anomaly next.

B. Anonymous versus Face-to-Face Bargaining

In the various face-to-face bargaining experiments discussed here, each pair of bargainers engaged in unrestricted conversation over how to divide an available sum of money. In some cases the experimenter sat with the bargainers to monitor the conversation. In some cases there were restrictions on how the money could be divided. With some notable exceptions, these experiments report that observed outcomes are Pareto optimal in almost all cases.[46]

For example, Nydegger and Owen (1975) observed thirty pairs of subjects each play one of three games, in which agreement was required to divide an available pot of money. (In the absence of agreement each bargainer would receive zero.) In each of these games, every pair reached agreement on the outcome that gave the two bargainers equal monetary payoffs.

Similar results were observed in a somewhat different bargaining environment in experiments conducted by Hoffman and Spitzer (1982, 1985). In one of the conditions in their 1982 paper, pairs of subjects were asked to agree on how to divide up to $14, in face-to-face negotiations. However, if no agreement was reached one of them (the "controller," chosen just before negotiations began by the toss of a coin) could simply choose an outcome that would give her up to $12 and the other bargainer nothing. When bargainers negotiated with each other

(i.e., each bargainer received $7), so that the controller settled for a smaller cash payoff than she could have obtained unilaterally. In Hoffman and Spitzer's other observations of two person bargaining, only two out of thirty-two outcomes failed to maximize joint profits (and 14 of these 32 also were equal divisions, giving the controller less than he could obtain unilaterally).[47]

Of course, we have to exercise caution in interpreting these data: since these experiments are different in other ways from the anonymous experiments discussed above, it might be that the much higher levels of efficiency (and equal divisions) are due to something other than the face-to-face conduct of the bargaining.[48] However, in an experiment involving bargaining with incomplete information, Radner and Schotter (1989) reported a careful comparison of face-to-face and anonymous bargaining. They found that face-to-face bargaining captured over 99 percent of the gains from trade in an environment in which anonymous bargaining captured only 92 percent.[49]

In what follows, we consider two very different hypotheses about the causes of these differences and a new experiment to help distinguish between them.

1. Two Hypotheses

a. The Uncontrolled Social Utility Hypothesis

The hypothesis that has motivated many experimenters to conduct experiments under anonymous conditions is that face-to-face interactions call into play all of the social training we are endowed with, and may make it unusually difficult to control preferences. (Ask yourself if you would agree to be very rude to a stranger if I offer to pay you $5.) Under this interpretation, it is difficult to interpret face-to-face bargaining experiments because of the possibility that powerful social motivations that may have little to do with bargaining may be responsible for the observed behavior.[50]

Siegel and Fouraker (1960, 22–3) explained their decision to conduct anonymous bargaining experiments in this way:

> This procedure eliminates certain variables which may well be important in bargaining—variables connected with interpersonal perceptions, prejudices, incompatibilities, etc. It is our belief that such variables should either be systematically studied or controlled in experimentation on bargaining. It cannot be assumed, as has often been done, that such variables may simply be neglected. We have chosen to control these variables at this stage of our research program, with the intention of manipulating and studying them systematically in future studies.

Fouraker and Siegel never did get to the future studies: Their 1963 monograph, published after Siegel's death, also studied anonymous bargaining. But while I am not aware of any subsequent bargaining experiments directed at dissecting the components of face-to-face interaction that contribute to the low incidence of disagreement, there is abundant experimental evidence in the social psychology literature that small differences in the social environment can cause large periments on public goods provision in which manipulations designed to alter individuals' feelings of group identity have substantial effects on the amount of public goods provided. Thus there is indirect evidence that makes plausible the hypothesis that the results of face-to-face bargaining experiments may reflect motivations deriving primarily from uncontrolled aspects of the social environment.

b. The Communication Hypothesis

The counterhypothesis that I think is at least implicit in the work of many experimenters who have focused on face-to-face bargaining is that there are many channels of communication available to face-to-face bargainers that are inevitably eliminated by any procedures that secure anonymity, even if those procedures otherwise allow fairly extensive communication. For example, Radner and Schotter (1989) compared anonymous bargaining in which only numerical bids and offers could be exchanged with face-to-face bargaining in which there was unrestricted verbal communication. They note that they cannot conclude that the high levels of efficiency in face-to-face bargaining were due simply to the fact that the bargainers could talk to each other, since the levels of efficiency they obtain also exceed those reported in Roth and Murnighan (1982), in which bargaining was conducted anonymously via computer terminals, and bargainers could freely exchange typed messages (recall the discussion of this experiment in chapter 1). They raise the question of whether the high levels of efficiency might arise from the channels of communication that face-to-face bargaining makes available *in addition* to the purely linguistic channel.

While I am not aware of any experiments designed to explore this hypothesis in the context of bargaining, there is ample evidence from the large literature on nonverbal communication generally that face-to-face communication employs many channels of communication, including tone of voice, body language, and facial expression. (For example, Grammer et al. [1988] study the "eyebrow flash," which they find serves in a variety of cultures as "a 'social marking tool' which emphasizes the meaning of other facial cues, head movements and even verbal statements.") Thus there is indirect evidence that makes plausible the hypothesis that the low observed levels of disagreement and inefficient agreements in face-to-face bargaining are due to the communication opportunities that such bargaining provides, and that the higher levels of disagreement in anonymous bargaining are due to the restricted channels of communication that are available to anonymous bargainers even when messages are allowed.

2. A New Experiment

Plausible conflicting hypotheses are what the experimental method thrives on, and I take the opportunity here to report briefly a small experiment that takes a step towards investigating the hypotheses just described. Bargaining in an ultimatum game (to divide $10) is compared under three conditions. In the first, baseline condition, bargaining is anonymous, with no communication other than written

Table 4.6. Comparisons of Ultimatum Games Played Anonymously

Condition	Disagreement Frequency (%)	Mean Offers (x_2)	Standard Error	Number of Observations	Percentage of Offers with $x_2 = 5.00$	Percentage of Offers with $x_2 = 5 \pm 0.50$
No communication	33	4.27	1.17	189	31	50
Unrestricted face-to-face communication	4	4.85	0.73	49	75	83
Social face-to-face communication only	6	4.70	0.46	49	39	82

Notes: In order to be comparable to the data in the next two rows, all data in the first row are for the first seven (out of ten) rounds of bargaining. For all ten rounds the corresponding figures are as follows: disagreement frequency is 28%, mean x_2 is 4.30, standard error is 1.05, number of observations is 270, percentage of offers with $x_2 = 5.00$ is 33%, and percentage of offers with $x_2 = 5 \pm 0.50$ is 51%.

offers and acceptances-rejections. In the second condition, there is unrestricted face-to-face communication: prior to making an offer, the buyer and seller have two minutes to discuss the game (or anything else). In the third condition, the buyer and seller have two minutes to converse, but are restricted to "social" conversation; they are required to learn each others' first name and year in school and are *not* allowed to discuss the bargaining game.[51]

If prior observations that face-to-face bargaining yields few disagreements are well founded, the unrestricted face-to-face communication condition should yield fewer disagreements than the anonymous bargaining condition. The comparison of the unrestricted and social communication conditions will then provide a test of the two hypotheses discussed above. If the lower disagreement frequency in the unrestricted communication condition is due to the increased communication, then the substantial decrease in the amount of communication allowed in the social communication condition, which eliminates the most germane verbal communication, should cause the disagreement frequency in that condition to be substantially higher than in the unrestricted communication condition. But if the lower disagreement frequency is due to social pressures arising directly from the face-to-face encounter and unrelated to what the bargainers say to each other about the bargaining, then the disagreement frequency in the social communication condition should be comparable to that in the unrestricted communication condition.

The results are summarized in Table 4.6. The disagreement rate declines from 33 percent in the anonymous condition to 4 percent in the face-to-face with unrestricted communication condition, which conforms with the frequencies of disagreements obtained under those conditions in earlier experiments. But the disagreement frequency when only social communication is allowed is only 6 percent: that is, it does *not* rise appreciably (or significantly).[52] Thus the results cast doubt on the communication hypothesis and support the uncontrolled social utility hypothesis as an explanation of disagreement frequency.

Note that the data suggest that the presence or absence of unrestricted verbal communication may have influenced the distribution of agreements in the two face-to-face conditions, even though it did not influence the disagreement frequency. The percentages of precisely equal divisions in the anonymous and social communication conditions are much lower than in the unrestricted communication condition.[53] But this data set is not large enough to know if this difference is important, especially since when we look at offers that are within fifty cents of an equal division, the two face-to-face conditions resemble each other more than the anonymous condition.

3. Some Further Experiments

After these results were presented at the Handbook Workshop, several of the experimenters who participated conducted follow-up experiments. The first of these to be reported was by Hoffman, McCabe, Shachat, and Smith (1991), at the annual convention of the American Economic Association. They proposed that, since face-to-face interaction among bargainers had such a pronounced effect compared to anonymous bargaining, perhaps a similar effect might be traced to the fact that subjects in these experiments are known to the experimenter. In particular, they proposed to compare games in which (following the usual practice) the experimenters could identify how each subject had behaved with games in which the experimenters would know only how a group of subjects had behaved, and would not know precisely what each subject had done. Their hypothesis was that the failure to observe perfect equilibrium in ultimatum games might be due to the fact that subjects felt they were being observed by the experimenter (e.g., player 1s might feel embarrassed to demand too much for fear of appearing greedy to the experimenter), and that extreme offers, as predicted by the perfect equilibrium, might therefore be forthcoming if experimenters only observed the actions of anonymous subjects.

To test this hypothesis they designed an experiment that examined both ultimatum games and dictator games,[54] and that included an anonymity condition for dictator games in which subjects passed sealed envelopes in which they received their pay without being directly observed by the experimenter.

Hoffman et al. reported that for dictator games conducted under non-anonymity conditions and using the instructions of Forsythe et al (1994), the distribution of offers was similar to that observed by Forsythe et al. (with a mode at the equilibrium offer), while the dictator games in the anonymity condition with the new instructions yielded an even larger mode at the equilibrium offer, although nonequilibrium offers were not extinguished. Their reported ultimatum game results concerned games run under nonanonymity conditions with different instructions. Offers were lower in cells in which the instructions gave the player 1s the "moral authority" to make low offers, but in no cells did offers approach the perfect equilibrium prediction of zero.[55]

On the basis of these results Hoffman et al. concluded that when bargainers were anonymous not only to each other but also to the experimenter, perfect equilibrium offers would more likely be observed. In general, they concluded (Hoffman et al. 1991, 1992):

> The results also emphasize that the argument for the use of anonymity in bargaining experiments as a means of controlling for social influences on preferences has not gone far enough. The presence of the experimenter, as one who knows subjects' bargaining outcomes, is one of the most significant of all treatments for reducing the incidence of self-regarding behavior.

There were a number of reasons to treat these conclusions cautiously, however, in part because of the way the experiment had been conducted and reported. In particular, the results Hoffman et al. attributed to their anonymity condition appeared to be confounded with other possible causes, because the instructions to subjects were radically different in their anonymity and nonanonymity conditions, and because different experimenters conducted the anonymity and nonanonymity conditions. Furthermore, results run under the anonymity condition were reported only for the dictator game and not for the ultimatum game. Thus the conclusions about anonymity rested on differences observed in offers made in dictator games under their anonymity and nonanonymity conditions. The ultimatum games they reported were run under nonanonymity conditions only, and only the proposals made by player 1s were reported—whether they were accepted or rejected by player 2s was not reported. In addition, no hypothesis was offered as to why the lack of subject-experimenter anonymity might inhibit players from making or accepting extreme demands in ultimatum games given that extreme payoffs are not inhibited in other games, such as the best shot games or sequential market games discussed earlier.

To address these concerns, Bolton and Zwick (1992) conducted an experiment that attempted to see whether the conclusions of Hoffman et al. about anonymity could be supported in an experiment that would not be subject to such criticisms. Regarding the difficulty of drawing conclusions about ultimatum games from dictator games (and the importance of reporting disagreement frequencies in ultimatum games) they note:

> Recall that lab ultimatum game investigators consistently report that second movers reject money. It is a key observation. In particular, even if first movers do act to maximize their personal earnings, they will not want to make the perfect equilibrium offer if a large enough proportion of second movers are going to reject it.

And regarding the importance of uniform instructions across cells of the experiment, they note:

> Because the control of information passing between experimenter and subject is crucial, we conducted the entire experiment from a script, referred to as the Experimental Protocol. It provides a thorough description of all remarks made by the experimenter during the experimental session.

The use of a uniform script also reduces most of the potential for introducing uncontrolled variation in conditions when different experimenters conduct different cells of the experiment, but Bolton and Zwick apparently also used the same personnel for each cell of their experiment.

In addition to comparing ultimatum games under anonymous and nonanonymous play, Bolton and Zwick examine a game that has the same move sequence as the ultimatum game and the same incentive structure as the dictator game, which they call the "impunity" game (since player 1s may make low offers "with impunity"). In the impunity game player 1 proposes how to split the pie, and player 2 may accept or reject, but player 1 receives the share he proposed for himself even if player 2 rejects—player 2's acceptance or rejection determines only if player 2 receives the share that player 1 allocated to her. Because ultimatum and impunity games have the same structure, the three cells of their experiment (anonymous and nonanonymous ultimatum games and nonanonymous impunity games) could be conducted from the same script with only very minor modifications.[56]

In examining their data, Bolton and Zwick propose to test two competing hypotheses, which they call the anonymity and punishment hypotheses. By the anonymity hypothesis they mean the conclusion of Hoffman et al. that an important determinant of observed behavior will be whether or not the experimenter can observe the actions chosen by individual subjects in the experiment. In contrast, the punishment hypothesis, motivated by the earlier experimental findings for bargaining games, is that an important determinant of player 2's behavior will be a desire to "punish" player 1 when he makes offers that are "too low" and that the behavior of player 2s will be an important determinant of the behavior of player 1s, who will attempt to maximize their earnings while taking the anticipated response of player 2 into account.

What Bolton and Zwick found is that anonymity made little difference—in the first five rounds of their ultimatum games, offers by player 1s were slightly lower in the nonanonymous condition, and in the last five rounds they were slightly lower in the anonymous condition, and in both conditions behavior was comparable to that observed in previous (nonanonymous) ultimatum game experiments. In contrast, the absence of punishment opportunities had a dramatic effect—in the impunity game virtually 100 percent of the plays in the last five rounds were perfect equilibrium plays, with player 2 receiving the smallest feasible payoff. They conclude that "the punishment hypothesis strongly outperforms the anonymity hypothesis. . . ."[57]

Notice that Bolton and Zwick's impunity game joins a growing list of games—including the best shot and market games discussed earlier—in which experiments have shown that extreme perfect equilibria can be descriptive of observed

behavior. As Bolton and Zwick note, this "provides evidence that first movers act in a self-interested manner even if the experimenter can observe their actions." That is, not only did Bolton and Zwick observe directly that the presence or absence of subject-experimenter anonymity did not influence the frequency of extreme demands, they also observed that extreme demands could be observed with very high frequency even without subject-experimenter anonymity.

Bolton, Katok, and Zwick (forthcoming) conducted a subsequent experiment concerned with dictator games, in order to better analyze whether the effect observed for those games by Hoffman et al. might nevertheless be due to subject-experimenter anonymity or if it was more likely due to the other uncontrolled differences between cells in that experiment. To this end, Bolton, Katok, and Zwick conducted an experiment in which dictator games were examined under different presentation conditions, and with and without anonymity under one of the presentation conditions. They observed that the game was sensitive to presentation effects, and in one of their (nonanonymous) conditions they observed a very large mode at the perfect equilibrium offer. However, they observed essentially no effect due to subject-experimenter anonymity, when games in which the experimenter could not observe which subjects made which offers were compared with other games, presented in the same way, in which the experimenter could make these observations. They conclude (27):

> We find no evidence for the anonymity hypothesis. . . . Comparison of our data with that of previous studies suggests that differences in the context of the game, affected by differences in written directions and independent of experimenter observation, account for the observed differences across dictator studies.

Thus there is no evidence that observation by the experimenter inhibits player 1 in ultimatum games, nor that it is the cause of extreme demands in dictator and impunity games. Rather, the evidence supports the view that the different behavior observed in games with similar perfect equilibrium predictions is due to differences in the games, off the equilibrium path. (The off the equilibrium path difference between ultimatum games and dictator or impunity games is particularly clear, since on the equilibrium path no offer is rejected in any of the games, but only in the ultimatum game is there an off-the-equilibrium-path possibility that an offer will be rejected). The evidence further suggests that the results of ultimatum game experiments show some sensitivity to how the games are presented to the subjects, but that this sensitivity is insufficient to overcome the robust observation that outcomes remain far from the perfect equilibrium prediction.

Two further studies round out the tale. Berg, Dickhaut, and McCabe (1993), in a paper presented at a conference on experimental economics in Amsterdam in 1993, reported an experiment in which a subject decided how much (x) of his $10 show up fee to send to another (anonymous) subject in another room, who would receive three times the amount of money sent by the first subject. The second subject would then decide how much of this amount ($3x) to return to the first.

Perfect equilibrium in this game predicts that no money is sent in either direction. However, the authors found that, under conditions of subject-experimenter anonymity similar to those developed by Bolton and Zwick, nonnegligible sums are sent in both directions. Thus the results of this experiment further disconfirm the conjecture of Hoffman, McCabe, Shachat, and Smith (1991, 1992) that subject-experimenter anonymity would promote perfect equilibrium play.[58]

Finally, Hoffman, McCabe, and Smith (1993), in a paper presented at the Southern Economic Association meetings in November, reported an ultimatum game experiment conducted in the usual way (i.e. without subject-experimenter anonymity), whose chief aim was to see whether a large increase in the monetary incentives would move the results closer to the perfect equilibrium prediction. To this end, they compared ultimatum games played one time where the amount to be divided was either $10 or $100, in units of either $1 or $10, respectively. In addition, they again compared the effects of presenting the game to the subjects by means of different sets of instructions. In this experiment they compared each of two sets of instructions both with $10 games and $100 games so that the separate effects of instructions and monetary incentives could be reliably assessed.

Their experiment confirmed that different sets of instructions could influence the distribution of offers, although in no case did they observe any substantial frequency of perfect equilibrium offers (of $0 or $1 in the $10 games, or of $0 or $10 in the $100 games). Rather, under both sets of instructions, for both $10 and $100 games, offers were predominantly much higher, ranging up to $5 in the $10 games and $50 in the $100 games, and consistent with distributions observed previously.

As to the effect of changing the stakes from $10 to $100, they write that, under both sets of instructions (6), "We cannot reject the hypothesis that the offers are identical with $10 stakes and with $100 stakes."

Thus we see here a series of experiments. one of whose results seems to be that even initially very skeptical investigators are becoming persuaded that the experimental results observed in ultimatum games are not easily displaced artifacts of the experimental methods, but rather represent a very robust phenomenon.

4. Recapitulation of the Methodological Issues

The methodological issues raised by the experiments considered in this section can usefully be thought of at three different levels of generality, concerning the design of individual experiments, concerning the issues of experimental control raised by this whole series of experiments, and concerning the role that series of experiments play in revising hypotheses and resolving discrepancies.

At the level of the individual experiment, it is a familiar observation that it is easiest to reliably observe the effect of a particular variable by changing only that variable. Designing appropriate experiments may require considerable ingenuity, but, as the experiment of Bolton and Zwick shows, a really well-designed experiment can eliminate a great deal of potential confusion.

Notice that the design questions associated with attributing an observed effect to a particular cause may be more complex. For example, while I do not think that anyone can reasonably dispute that subject-experimenter anonymity had at most a negligible effect on the ultimatum games in Bolton and Zwick's experiment (or on dictator games in Bolton, Katok, and Zwick [1993]), it is easy to imagine that some investigator will come up with an alternative explanation of the big difference they observed between ultimatum and impunity games. The reason is that an experimental design intended to distinguish between alternative hypotheses controls for the variables that are relevant for those hypotheses (in this case, the structure of the two games), but may not have controlled for variables relevant to some other (perhaps as yet unstated) hypothesis. Thus there is a strong sense in which "appropriate" experimental designs are creatures of their time—they depend not only on what is being tested, but on what are the most plausible hypotheses at the time they are being tested.

Concerning the issues of experimental control raised by this series of experiments, an analogy may be useful. Chemical reactions are studied in glass vessels rather than in pots made of metal, or even clay, because glass is more inert, less volatile. Before chemists figured this out, they often must have had to deal with anomalous results coming from investigators who used different kinds of pots. Of course, there are materials that are even less volatile than glass, and these may be required to effectively control certain kinds of experiments, but for most purposes glass seems to do fine. In a similar way, face-to-face interaction is a volatile environment in which to conduct economics experiments, and most investigators have followed Siegal and Fouraker in avoiding it and conducting experiments in environments that preserve between-subject anonymity. We have reconfirmed that there is good reason for this. And the experiments of Bolton and Zwick and Bolton, Katok, and Zwick strongly suggest that between-subject anonymity is enough for most purposes, since little if any additional effect was observed due to subject-experimenter anonymity. And just as face-to-face interaction is volatile, so too may be instructions that try to give players the "moral authority" to take certain actions or that in other ways may induce experimenter-demand effects by too clearly indicating the goals of the experimenters.

But while there is ample reason for preferring glass to clay for conducting most experiments, this certainly does not mean that it is not interesting to study the chemistry of clays. Even if the phenomena observed in face-to-face bargaining experiments largely reflect uncontrolled aspects of social interaction, these are worth studying, not only as a means of learning which laboratory procedures are especially volatile, but also as a means of understanding subjects' social perceptions.[59] For example, Burrows and Loomes (1990) use Hoffman and Spitzer's (1985) observation that observed agreements were sensitive to noneconomic aspects of the instructions to motivate an investigation aimed primarily at elucidating subjects' notions of fairness.[60]

Furthermore, the fact that face-to-face bargaining may be difficult to study does not mean that it is unimportant or that its presence or absence may not alter the course of negotiations (even) over stakes much larger than can be studied in the laboratory. Newspaper accounts of bilateral meetings between American presidents and their foreign counterparts are full of analyses of the personal relationships established or not established between the leaders, and there is much anecdotal testimony that such face-to-face meetings may influence negotiations (although not always in the direction of efficiency).[61] Similarly, the fact that many negotiations are conducted through intermediaries (such as lawyers or real estate agents) suggests that many people are prepared to go to considerable expense to avoid face-to-face negotiations, perhaps because they feel the outcome will be different (and more profitable) if the social pressures arising from face-to-face negotiations between principals are avoided.[62] Of course, these phenomena can be explained in other ways as well, which is why it would be useful to study them under controlled conditions.

Finally, since most of the material in this handbook addresses the role that series of experiments play in refining hypotheses and resolving different points of view, let me focus here simply on the *speed* with which this process has proceeded in the present case, in comparison to what we might expect if we had to rely exclusively on field data. In a relatively short period of time, a consensus seems to have (re)emerged about the role played by anonymity between subjects (versus face-to-face interaction); an experiment was reported that suggested to its authors that a similar role might be played by subject-experimenter anonymity; and new experiments were conducted and showed that this was unlikely to be the case. What makes this a relatively fast process is not merely that experimental data can often be gathered more quickly than field data, but that experimental data can be gathered to fit precisely the question of interest (e.g., hypotheses about ultimatum games can be tested on ultimatum games). Furthermore, although investigators with different hypotheses may be inclined to collect different data, and even report different aspects of it, the fact that experimenters all have access to essentially the same data universe, since they can conduct their own experiments, means that investigators are much less dependent on obtaining access to data than in situations where access is limited either by expense or the uniqueness of the field situations being studied.

Returning to substantive matters, the material in this section reconfirms the importance of disagreements in determining the nature of agreements, and the next section considers experiments designed to investigate the causes of disagreements.

C. *Tests of Hypotheses about the Causes of Disagreements and Costly Delays*

The main approaches to modeling disagreements and delays in bargaining can be ordered by how much information they include about the environment, and in how much detail they model the behavior of the bargainers. The simplest models attempt to relate disagreement frequency to features of the bargaining situation such as the shape of the set of possible agreements (and the consequent divergence of interest between the bargainers) or the value of potential agreements and

the costliness of delays in reaching agreements. These models are nonstrategic, in that they do not attempt to model the detailed behavior of the bargainers. A more detailed approach seeks to model the bargaining as a game of complete information, in which the same kinds of features of the bargaining situation, together with rules of bargaining, produce equilibrium behavior on the part of the bargainers that in turn determines the frequency of disagreement. And the most detailed attempts involve modeling bargaining as a game of incomplete information, in which each bargainer takes into account different kinds of private information in determining his own (equilibrium) behavior and in forming his expectations about the behavior of the other bargainer.

Experimental investigations of each of these classes of models are only in their earliest stages so that there are as yet no extended series of closely connected experiments to report. However, in reviewing some of the initial steps that have been taken, it is easy to see how even initial experiments raise questions that suggest further experiments.

1. Nonstrategic Models

A nonstrategic model proposed by Axelrod (1970) was tested in an experiment reported by Malouf and Roth (1981). Axelrod had proposed a measure of what he called "conflict of interest" inherent in a bargaining situation, based (like Nash's solution[63]) on the set of feasible expected utility payoffs available to the bargainers, together with their utilities in the event that no agreement is reached. The idea is that when the set of Pareto optimal agreements consists of a unique point at which both players receive their maximum payoff, as in game 1 of Figure 4.8, there is no conflict of interest between the players. But when the maximum individually rational payoffs of the players are incompatible, some conflict of interest exists. (Because this is a theory based on the independently measured expected utilities of the bargainers, the origin and scale of the units in Figure 4.8 are not important for determining the conflict of interest. Here we take the utilities in case of disagreement to be normalized to (0,0): the experimental implementation of the games will be discussed in a moment.) More formally, consider the rectangle formed by the players' disagreement payoffs and their maximum individually rational payoffs. For the games in Figure 4.8, these are the rectangles whose lower left corner is (0,0), and whose upper right corners are (60, 30) in game 1, (90, 40) in game 2, (90, 50) in game 3, and (90, 90) in game 4. Axelrod proposed that conflict of interest for each game could be measured by the proportion of the associated rectangle that is not in the set of feasible agreements. By this measure, the conflict of interest is 0 for game 1, and it rises to its maximum of .5 for game 4, with games 2 and 3 having intermediate values (of .22 and .27, respectively).

Malouf and Roth (1981) reported an experiment involving the four games in Figure 4.8, to test the hypothesis that the frequency of disagreements would be greater in games with a higher conflict of interest as measured in this way. Since Axelrod's measure was defined in terms of the bargainers' expected utilities, the

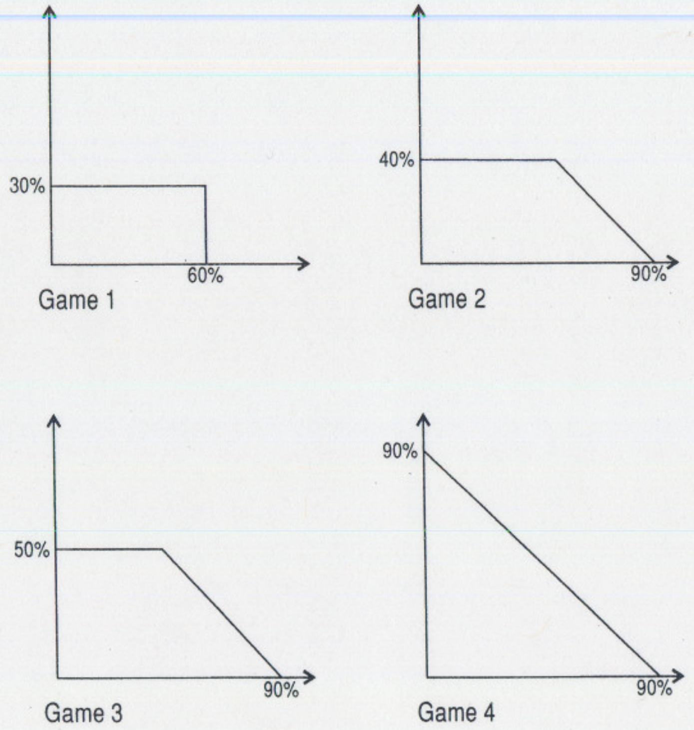

Figure 4.8. Four games with increasing "conflict of interest." *Source:* Roth and Malouf 1981.

games were implemented as binary lottery games with large and small prizes of $10 and $5, respectively.[64] In each game, players divided lottery tickets subject to a restriction on the maximum percentage each player was allowed to receive. (As indicated in Figure 4.8, this maximum was 60 percent for player 1 in game 1 and 90 percent in games 2, 3, and 4, while for player 2 it was 30, 40, 50, and 90% in games 1 through 4.) The principal variables of interest were the mean time to agreement in each game and the percentage of disagreements.

Although the mean times to reach agreement were roughly consistent with the ordering of the games by the conflict of interest measure,[65] the frequencies of disagreement were not. The observed frequencies of disagreement for the four games were 0, 21, 37, and 0 percent, respectively. That is, in game 4 as well as game 1, no disagreements were observed.

Looking at the agreements reached in the four games, an explanation immediately suggests itself. While in games 2 and 3 there is some variance in how the lottery tickets were divided between the two bargainers when agreement was reached, in game 1 almost all of the agreements were for a (60, 30) division, and in game 4 almost all of the agreements were for a (45, 45) division. Thus, despite the incompatibility in the maximum payoffs to the two players in game 4, captured by the high value of the conflict of interest measure for that game, the

symmetry of the game serves to focus the attention of the players on the equal division agreement, just as the lack of conflict of interest in game 1 concentrates the agreements on the unique Pareto optimal agreement. To put it another way, these results suggest that the frequency of disagreements increases with dispersion in the terms of observed agreements. This should not be surprising insofar as the failure to reach agreement presumably reflects that the two bargainers have different expectations about the terms of agreement, and such different expectations are easier to maintain in an environment in which observed agreements are more diverse.

We turn next to examine the experimental evidence in the context of a nonstrategic model that relates the frequency of disagreement to its cost.

Two closely related experiments, by Sopher (1990), and by Forsythe, Kennan, and Sopher (forthcoming), were conducted to test an appealingly simple hypothesis about disagreements formulated by Kennan (1980). Kennan's hypothesis, which was formulated to explain the rate of strikes and settlements in labor management negotiations, is that the incidence and duration of strikes will be a decreasing function of the marginal cost of a strike and an increasing function of the total size of the pie to be divided. Forsythe et al. used a four-cell factorial design in which the pie size ($4.00 and $8.00) and marginal cost ($1.00 and $.50) were varied,[66] in a manner made clear by their description of one cell of their experiment (12):

> Games in this cell consisted of Years which were 8 Months long. In each Year, a "Pot" of $9.60 was available for division between the two bargaining partners. If no agreement was reached in a given Month, the pot was reduced by $9.60/8 = $1.20, and the bargainers received their threat points of $.70 and 0, respectively. Thus, the marginal strike cost is 1.20 – .70 – 0 = $.50, and the total pie size is $.50 × 8 = $4.00."

Bargainers negotiated anonymously via written messages and proposals.

Although both experiments observe a substantial number of periods in which agreement is not reached, the results of these experiments yield little support, on the range of costs studied, for the prediction that strike activity (i.e., periods in which bargainers fail to reach agreement) falls as marginal costs rise. But both experiments report weak support for the proposition that strike activity increases as the size of the pie increases.

Clearly these results suggest directions for future work. The hypothesis that periods of disagreement will decline as the costs of disagreement rise is an appealing one, and it certainly still seems possible that a differently designed experiment (perhaps with larger costs, or with out of pocket costs charged against an initial endowment) might detect such an effect. In the other direction, because of the way the size of the pie was increased, an alternative hypothesis about the increased disagreement frequency observed in the cells with larger pie size is that this increase is related to the larger number of periods over which bargaining took place (e.g., because of the increased incentive in longer games to build a reputation in the early periods as a tough bargainer). The size of the pie and the length f the bargaining horizon are matters that cannot always be distinguished in field tudies, but could be varied independently in the laboratory.

An experiment in which the frequency of disagreements *was* observed to be nversely correlated with the cost of disagreement is reported by Ashenfelter, Currie, Farber, and Spiegel (1992), who considered bargaining followed by binding arbitration in the case of disagreements. In the labor relations literature, the ypothesis that reducing the cost of disagreement (by imposing binding arbitraion) will increase the frequency of disagreement is referred to as the "chilling effect" of arbitration. It has furthermore been hypothesized that disagreements will be more frequent under conventional arbitration (in which the arbitrator chooses the final outcome if the parties fail to agree) than under final offer arbitraion (in which the arbitrator must choose the final offer proposed by one of the wo parties to the dispute if they fail to agree).

A major problem in the design of any experimental test of bargaining in which arbitration plays a role is that arbitration is (at least) a three-party activity, involving two bargainers and an arbitrator. If the arbitrator is to be one of the subjects in the experiment, a critical design feature will be how to design the incentives for the arbitrator, which will influence not only his behavior, but also, therefore, the behavior of the bargainers. Ashenfelter et al. find an elegant solution to this design problem, by incorporating results from the field studies of Ashenfelter and Bloom (1984) and Ashenfelter (1987). These studies suggest that the selection process for arbitrators causes them to behave in such a way that one arbitrator is statistically exchangeable for another. The idea is that, since arbitrators must be acceptable to both parties, arbitrators who are known to favor one side or the other are eliminated. Consequently in Ashenfelter et al.'s experiment, the arbitrator is represented as a random draw from a fixed distribution. The outcome over which the bargainers negotiate is a number between 100 to 500, with one bargainer profiting from higher numbers and the other from lower numbers. (Players had symmetric payoff schedules, but each player knew only his own payoffs.) When bargaining is to be followed by conventional arbitration, the outcome in the case of disagreement is simply the number drawn by the arbitrator. When bargaining is followed by final offer arbitration, the outcome in the case of disagreement is whichever of the bargainers' final offers is closer to the number drawn by the arbitrator.

Data were gathered from fixed pairs of subjects who negotiated with each other for twenty periods. In the first ten periods there was no arbitration; failure to reach agreement resulted in a zero distribution for that period to both parties. Prior to the second ten periods, the bargainers were informed that either conventional or final offer arbitration would determine the outcome in each remaining period, in case no agreement was reached by the bargainers themselves. The bargainers were not informed that the arbitrator was a random distribution, but were given a list, generated from the arbitrator's distribution, of what they were told were the arbitrator's last one hundred decisions. Aside from varying the form of arbitration, the experiment also varied the variance of the arbitrator's distribution in the conventional arbitration conditions.

In all of the experimental conditions, Ashenfelter et al. observed a higher rate of disagreements in the second ten rounds, when disagreement resulted in arbitration, than in the first ten rounds, when disagreement resulted in zero payoffs to both bargainers. They also observed that the frequency of disagreements in the conventional arbitration condition declined as the variance of the arbitrator's distribution increased. Both of these observations are consistent with the hypothesis that disagreement frequency is inversely related to the cost of disagreement, the idea being that higher arbitrator variance increases the cost of arbitration to risk averse bargainers. However, contrary to the hypothesis that final offer arbitration has less of a chilling effect than conventional arbitration, Ashenfelter et al. observed that the dispute rate under final offer arbitration was no lower than under conventional arbitration.

Ashenfelter et al. are careful to note that an alternative explanation of the higher dispute rates observed when arbitration was present is possible, due to a feature of their experimental design that confounds two effects. The source of this confound is that in the first ten periods the bargaining situation was symmetric, with the symmetric agreement being 300 (recall that one party prefers higher numbers, the other lower numbers). However, the second ten periods (the arbitration periods) were not symmetric, since the mean of the arbitrator's distribution was 350, that is, since the arbitrated solutions tended to favor one of the bargainers in case no agreement was reached. Recall from the discussion of Malouf and Roth (1981), earlier in this section, that negotiations in nonsymmetric situations had more disputes than negotiations in symmetric situations. Ashenfelter et al. therefore point out that the asymmetry introduced in their arbitration periods may also account for some of the increase in disagreement frequency that they observe.

Taken together, the work of Malouf and Roth (1981), Sopher (1990), Forsythe et al. (1991), and Ashenfelter et al. (1992) suggest that there may be considerable progress to be made by considering theories of bargaining that depend (only) on the gross features of the bargaining problem. We turn next to consider some experiments motivated by theories that depend on the more detailed strategic structure of the bargaining game.

2. Strategic Models

a. Complete Information Models

Recall that, in the binary lottery bargaining games reviewed in chapter 1, agreements were often bimodal, with one mode at the equal division of lottery tickets and the other at the division yielding equal expected payoffs. Two questions arise in considering how one might try to incorporate this observation into a theory of bargaining. The first is, what causes some potential agreements and not others to become focal points in this way? The second is, given that the bargainers recognize that (for whatever reason) certain potential agreements are especially credible bargaining positions, how will this affect aspects of the bargaining such as the frequency of disagreements? In Roth (1985) I undertook to look at this latter question and to test the predictions of a very simple strategic model.

Table 4.7. A Coordination Model of Disagreement

(50, 50)	$(\frac{1}{2}(50-h), (\frac{1}{2}(150-h))$	$h \geq 50 \geq d_1$
(d_1, d_2)	$(h, 100-h)$	$50 \geq 100-h \geq d_2$

Source: Roth 1985.

In particular, consider a highly structured model of a binary lottery bargaining situation in which two credible bargaining positions are recognized, the (50, 50) division of lottery tickets and another division $(h, 100 - h)$ where $h \geq 50$. If both players agree on one of these two divisions then it is the outcome of the game. If each player holds out for the division more favorable to him, then disagreement results, and if each player is willing to accept the division more favorable to the other, then a compromise agreement is reached whose expected value is the average of the two credible divisions. This game is represented as a two-by-two matrix game in Table 4.7, in which the strategies of the players are simply which of the two divisions to demand, and the disagreement utilities of the players are given by (d_1, d_2).

There are two pure strategy equilibria of this game, each of which results in one of the two credible divisions. There is also a mixed strategy equilibrium, which can be taken as a simple model of disagreement frequency as a function of the distance between the two focal divisions. That is, at the mixed strategy equilibrium, disagreements occur because of coordination failure of the kind proposed by Schelling (1960), and discussed in other contexts in chapter 3. In particular, a disagreement occurs whenever player 1 demands h and player 2 demands 50. At the mixed strategy equilibrium, the probability that 1 demands h is $p = (h - 50)/(150 - h - 2d_2)$ and the probability that player 2 demands 50 is $q = (h - 50)/(h + 50 - 2d_1)$, so the probability of disagreement, pq, is an increasing function of h. That is, as the second focal point becomes more distant from (50, 50), the predicted frequency of disagreement increases. (This hypothesis is, of course, consistent with the more general notion discussed above, in connection with the experiment of Malouf and Roth [1981] that the disagreement frequency is positively associated with the observed dispersion of agreements.)

Table 4.8 considers the disagreement frequency of three earlier experiments that allow a rough test of this prediction, in that they include games with values of h of 50, 75, and 80.[67] The prediction of the mixed strategy equilibrium is that disagreements should be observed with frequencies of 0 percent, 7 percent, and 10 percent respectively for these values of h, and overall the observed frequencies of disagreement are 7 percent, 18 percent, and 25 percent. Thus the observed frequencies of disagreement on this data set move in the direction predicted by this simple coordination model.

Of course, the data in Table 4.8 was assembled from experiments designed for other purposes. It thus seemed worthwhile to conduct an experiment to more specifically test this prediction of the mixed strategy equilibrium model of disagreement frequency. Such an experiment was reported in Roth, Murnighan, and

Table 4.8. Frequency of Disagreement

	$(h, 100-h)$		
Experiment	(50, 50)	(75, 25)	(80, 20)
Roth and Malouf (1979)	2% (1/54)[a]	14% (3/21)[b]	
Roth, Malouf, and Murnighan (1981)	6% (2/32)[c]	20% (6/30)[d]	24% (7/29)[e]
Roth and Murnighan (1982)	11% (7/63)[f]		25% (37/146)[g]
All experiments combined	7% (10/149)	18% (9/51)	25% (44/175)
Prediction of the coordination model (mixed-strategy equilibrium)	0%	7%	10%

Source: Roth 1985.

[a] Games with only a (50,50) focal point in this experiment are all those in the partial-information condition, and games with equal prizes for both bargainers in the full-information condition.

[b] Games with a (75, 25) focal point in this experiment are games 3 and 4 in the full-information condition.

[c] Games with only a (50, 50) focal point in this experiment are all games in the low-information condition.

[d] Games with a (75, 25) focal point in this experiment are games 1 and 3 in the high-information condition.

[e] Games with an (80, 20) focal point in this experiment are games 2 and 4 in the high-information condition.

[f] Games with only a (50, 50) focal point in this experiment are those in which neither player knows both prizes, in the common- and non-common-knowledge conditions.

[g] Games with an (80,20) focal point in this experiment are all those in conditions in which $5 player knows both prizes.

Schoumaker (1988). It involved binary lottery games in which both bargainers had the same low prize (i.e., their payoff in the event of disagreement, or if they lost the lottery resulting from an agreement) of 0, but different high prizes. Games were examined with high prizes for the two bargainers of $10 and $15, $6 and $14, $5 and $20, and $4 and $36. Thus the equal expected value agreements in these four conditions were (60, 40), (70, 30), (80, 20), and (90, 10), respectively. That is, in terms of the model of Roth (1985) the values of h were 60, 70, 80, and 90.

As in the previous binary lottery bargaining experiments the greater the difference between the prizes of the two bargainers, the higher the mean percentage of lottery tickets received by the bargainer with the smaller prize. However, unlike some of the previous experiments, the distribution of agreements was approximately normal rather than bimodal. And, contrary to the predictions of the mixed strategy equilibrium of the coordination game (and contrary to my expectations), the disagreement frequency stayed approximately constant as h varied from 60 to 90. Thus the simple coordination model seems *too* simple to organize the data on disagreement frequencies in this case.

b. Incomplete Information Models

A class of more complex models are models of incomplete information, in which each player may have some private information about his own situation that is unavailable to the other players, while having only probabilistic information about the private information of other players. Following Harsanyi (1967, 1968), models of games of incomplete information proceed by adopting the assumption that (other) players all start with the same prior probability distribution on this private information (which they may update on the basis of their own private information) and that these priors are common knowledge. This is modeled by having the game begin with a probability distribution, known to all the players, which determines each player's private information.[68] Thus players not only have priors over other players' private information, they also know what priors the other players have over their own private information. Strategic models of incomplete information thus include an extra level of detail, since they specify not only the actions and information available to the players in the course of the game, but also their prior probability distributions and information prior to the start of the game.

Experiments seeking to test formal theories of incomplete information must be carefully constructed if they are to meet the assumptions of the theory while controlling for the information and beliefs of the players. The experimental designs discussed below all deal with this by beginning the game with an objective probability distribution known to all of the players, so that the analysis of the experiment can proceed by taking each player's prior probability distribution to be equal to this objective distribution.[69]

Despite these common features of design, the experiments discussed below reflect an interesting diversity of experimental philosophy and theoretical disposition. We will see this both in the design of the experiments—which differ in how closely the experimental environments are related to the models whose predictions are being tested—and in the conclusions drawn from them. I will return to this point later.

Hoggatt and Selten et al. (1978) tested an incomplete information model of bargaining studied by Harsanyi and Selten (1972) and Selten (1975). The game in question involves bargaining over the division of twenty money units with each bargainer having private information about his own cost of reaching agreement, which is either zero or nine units, chosen with equal probability. A bargainer's cost is deducted from his payoff (only) in the event that an agreement is reached. In each of these experiments a money unit was worth ten cents. Some of the flavor of the experiment, and of how the common prior probability distributions of the players were induced, can be given by quoting from the instructions (131–2):

> There are six persons participating in this session and you will play the same bargaining game once against each of the others. In any game two players may divide 20 money units between themselves if they reach agreement. If they reach conflict neither receives any money units. At the beginning of a bargaining game it is decided by a separate random experiment for each player by drawing an "H" or "L" from the bag [which contains two balls, one of each kind] whether he has high or low cost. High cost = 9 money units, and low cost = 0 money units. These costs are deducted from the payments in the event that agreement is reached. . . . You will not know the cost of the other player but you will know your own cost and you also know that the cost of the other player was chosen high or low with equal probability independently from the selection of your costs. In any one game you will not know against which of the other participants you are playing. The other player will find himself in exactly the same general situation.
>
> The bargaining is done via teletype and proceeds in discrete stages. At the first stage the teletype will accept your demand for a share which must be an integer no lower than your cost and not higher than 20. In succeeding stages your demand must not be higher than the demand in the previous stage and no lower than your cost. . . . If a player's move is not completed within [2 minutes] the computer will take the demand of that player in the previous stage. . . .
>
> *Conflict* occurs at any stage for which neither player makes a concession, i.e. both demands remain at the levels set in the previous stage. . . . In case of conflict . . . both players have a net payoff of zero.
>
> *Agreement* is reached should a stage occur in which the sum of both demands is at most 20 money units . . . [in which case] each player gets his demand and then the amount by which the sum of demands falls short of 20 is split evenly.

The authors focus on the predictions of one of the multiple equilibria of the model, which they call the "main representation." Among the predictions at this equilibrium is that whenever two players with high costs play each other (a HH pair) disagreement will result, while for every other combination (HL, LH, and LL pairs) no disagreements will result. The observed results of the experiment were that the disagreement frequency was .729 for HH pairs, .471 for LH and HL pairs, and .097 for LL pairs (compared with predictions of 1.00 for HH and 0 otherwise), so the frequencies have a tendency in the direction of the theory. The authors say (143–144), "This weak tendency in the direction of the theory is not trivial, since a superficial analysis of the game may easily come to the conclusion that there should be no conflict at all, since mutually profitable agreements are possible for each of the type combinations."

However, the authors are less sanguine about the ability of the theory to organize other aspects of the data. For example, agreements are predicted to give the players equal gross payoffs only in the case of two low cost bargainers, while the data shows that for all kinds of bargaining pairs (HH, LH or HL, and LL) the modal agreement gave the players equal gross payoffs of ten money units each (i.e., equal payoffs before the private information bargaining cost was deducted).[70]

An interesting unpredicted feature of the data concerns the effects of the bargainers' risk aversion on the frequency of conflict. (Although the authors assume that the players are all risk neutral in order to simplify their theoretical analysis, they also collect data on the risk aversion of their subjects, by having them choose from a set of lotteries prior to the bargaining experiment.[71] In this preliminary experiment, a subject makes a sequence of A or B choices, with a higher number of A choices signifying a greater propensity to take risks.) A small but significant effect of risk aversion on disagreement frequency appears in the data of the main experiment, with a curious pattern. In the data for LL pairs and LH or HL pairs, an increase in the bargainers' propensity to take risks results in an increased frequency of conflict. However, the data for HH pairs reveals the reverse correlation: an increased propensity to take risks is associated with a reduced frequency of conflict. The authors go on to examine the correlation between risk taking propensity and the number of times bargainers repeat their demands in the course of a bargaining session. (Recall that a bargainer faces a risk that the game will end with disagreement on the next period only if he repeats his previous demand.) What they find is that for LL pairs (in which the stakes are high, since the bargaining costs are low) increased propensity to take risk leads to an increase in repeated demands. But for HH pairs (in which the stakes are lower, since most of the wealth will be consumed by the bargaining costs) the reverse is true. They conclude (158):

> With this result we can now understand the reversal. . . . Risk-takers are attracted by large payoffs and they are more likely to repeat a demand if there is a possibility of a large net payoff. A small additional net payoff does not induce them in the same way to take the risk of losing a small net payoff.
>
> This result illustrates the value of sequential analysis on a large data base. Given the anomaly [of reversals] we were led to search for an explanation by making a finer breakdown of the data.

In an effort to integrate both the predicted and unpredicted results of their experiment, the authors conclude their paper by presenting a "behavioral robot" in order "to produce a complete behavioral representation of modal behavior for the game."[72] Thus while the paper begins with a set of detailed game theoretic predictions about the play of the game, and a careful implementation of experimental conditions corresponding to the model for which the prediction was made, it ends with a non-game-theoretic model of player behavior.

A rather different approach is taken by Forsythe, Kennan, and Sopher (1991), who consider a bargaining game in which only one bargainer is informed about the size of the pie to be divided, while the other bargainer knows only the probability distribution that determines the size of the pie. In particular, the pie can take on one of two values, H ("the good state") and L ("bad state"), with $H > L$, and it takes on the high value H with a known probability p. They write that (253) "a particular goal of this work is to identify predictions that are robust with respect to the simplifying assumptions used in theoretical modeling." To this end, they use as the basis for their theoretical analysis two games—an ultimatum game and a "random dictator" game—which are highly structured and easy to analyze, and then seek to experimentally test the predictions derived from that analysis in a bargaining environment that is much more unstructured and complex, since it involves the free exchange of messages.

In their ultimatum game, the *un*informed player must propose a division of the pie (by specifying how much he demands for himself), which the informed player then accepts or rejects. At a subgame perfect equilibrium the informed player will reject any demand larger than the actual size of the pie (since by doing so he secures a payoff of zero instead of a negative payoff) and accept any demand that is not larger than the pie. So a risk-neutral uninformed player will demand H (i.e., all of the large pie) if $pH > L$, since this gives him an expected payoff of pH. And he will demand L if $pH < L$, which gives him an expected payoff of L.

In the random dictator game both the informed and the uninformed player submit a nonnegative proposal of how much the uninformed player should receive. One of these proposals is then chosen at random and implemented if it is feasible (i.e., if the amount does not exceed the size of the pie). Here also, the strategy of a risk-neutral uninformed player is to demand H if $pH > L$ and to demand L otherwise.

Thus a prediction of both simple models is that there will be disagreement—i.e., both players will receive 0—if and only if $pH > L$ and the pie size is L. The prediction implies that there will never be disagreement in the good state—i.e., when the pie is large, regardless of the probability p.

Forsythe et al. proceeded to test this prediction experimentally, both on a random dictator game of the kind just described and on an unstructured bargaining game in which the subjects were free to exchange written messages as well as offers and acceptances and rejections for ten minutes. They found that in the unstructured bargaining environment the prediction about disagreements was weakly supported, in the sense that disagreements were most frequent in games in which $pH > L$ and the size of the pie was L. But, as in Hoggatt and Selten et al. (1978), the support for the prediction was only weak, in the sense that there were substantial numbers of agreements in this case also, as well as disagreements when the pie size was H.[73] And the division of the pie when agreements were reached did not always conform to the predicted values.

In comparing the behavior observed in the unstructured bargaining with that observed in the random dictator game, Forsythe et al. observe similar behavior in many respects, with the notable exception that disagreements never occurred in the random dictator game when the pie was large, since no player ever demanded more than H. They summarize this comparison as follows (264–5):

> Our results show that the general pattern of the outcomes in the [unstructured] bargaining games was very similar to that of the R[andom] D[ictator] games. Communication did not substantially affect the incidence of strikes in the bad state; however, in the good state, strikes occurred only in the bargaining games, where the informed player could insist that the pie was small and the uninformed player had the right to insist that it was not.

Thus in this experiment increased opportunities for (anonymous written) communication *increased* the frequency of disagreement.[74] Since the random dictator games can be viewed as individual decision problems, the authors analyzed the individual data to see if the subjects are expected income maximizers, as assumed. They conclude that the formal predictions for the random dictator game (and therefore perhaps also for the unstructured bargaining game) failed to be fully descriptive (271) "because there was considerable heterogeneity among our subjects, including sizeable minorities of both risk-averse and risk-loving types, and another minority of altruists."[75] They conclude as follows (271):

> Although we have some encouraging results on the predictive power of the [predictions about disagreement], the results on heterogeneity of preferences and on inconsistency of decisions indicate that much caution is needed in drawing conclusions from behavior in bargaining (and other) experiments. In any game in which the players interact strategically, the theoretical analysis should not begin (as ours did) with the assumption that the players' objective functions are common knowledge.[76]

The next experiment I will discuss, by Radner and Schotter (1989; see also Schotter 1990) reports an incomplete information experiment, using a highly structured sealed bid mechanism. Buyers and sellers would each simultaneously submit a proposed (bid or asked) price, and no trade would result if the buyer's bid was lower than the seller's asked; otherwise the transaction price would be the average of the two. Buyers and sellers had private reservation values drawn from the same distribution.

The single play version of this game, in which the buyer and seller meet only once, has received a good deal of theoretical attention. Chatterjee and Samuelson (1983) studied an equilibrium of the game in which the players' bids are a linear function of their private values. Myerson and Satterthwaite (1983) observed that no equilibrium of the game can achieve one hundred percent efficiency (by achieving trades whenever the buyer's value is higher than the seller's), but that the linear equilibrium studied by Chatterjee and Samuelson achieves maximal efficiency on the set of equilibria.[77] And Leininger, Linhard, and Radner (1989) observed that there are a multitude of other equilibria, with widely varying efficiency. Radner and Schotter write that it is this multiplicity of equilibria ("this theoretical morass") that motivated their experiment. Their idea is that the negative implications of the theoretical multiplicity of equilibria may be tempered, for practical purposes, if subjects can in practice achieve the efficiency of the most efficient, linear equilibrium.

Radner and Schotter adopt an experimental strategy somewhere between that of Hoggatt and Selten et al. (1978) and Forsythe et al. (1991) in terms of how closely the experimental environment they construct conforms to the theoretical model whose predictions are to be tested. In particular, the information and communication available to the players are structured to conform precisely to the model being tested. But whereas the model describes a one-period game, the experiment studies a multiperiod repeated game. They describe their experimental environment as follows (182–3):

> Each seller/buyer drew 15 envelopes from a pile of 500. Each envelope contained a slip of paper with a number written on it. The numbers were generated randomly according to a commonly known probability distribution. If the subject was a buyer, then this random number indicated the value to him of the good being sold in that round. If the subject was a seller, the number represented the cost of producing the good in that round. After observing the realization in the envelope and recording it in their work sheets, subjects then wrote their bids on pieces of paper and handed them to a set of experimental administrators who collected them. When the slips of the buyers and sellers were brought to the front of the room, they were randomly sorted into pairs, each containing the bid of one buyer and one seller. These bids were then compared, and, using the rules of the sealed bid mechanism . . . prices and payoffs were determined. The price and trade results of these transactions were then distributed back to the subjects, and the next round began, which was conducted in an identical manner. In subsequent rounds subjects were paired against the same pair member. . . . Despite the danger of introducing repeated-game elements into what is intended to be a test of a static theory, we felt that this design feature was necessary if the subjects were to successfully select one equilibrium from the multitude defined by the mechanism.

The danger to which Radner and Schotter refer has to do with the fact that the multiperiod game that the subjects in the experiment actually play (since each buyer remains paired with the same seller throughout) has more equilibria than the one-period game, including equilibria which achieve a higher degree of efficiency than any in the one-period game.[78] This presents a special complication in the interpretation of their results, since it turns out that, rather than observing lower efficiency than at the linear equilibrium, they observed *higher* efficiency. That is, when the buyer's value is higher than the seller's, so that trade is profitable, agreements were observed more often than predicted. In the static, one-period case for which the theory was developed, this would be clear indication of non-equilibrium behavior, since the linear equilibrium is maximally efficient on the set of equilibria. However, since this is not the case in the repeated game actually played, the observed behavior could nevertheless be consistent with equilibrium play.

The questions of experimental design that this raises are worth a digression. Although economic theories of equilibrium so far have little to say about *equilibration*—i.e. about how equilibrium might be achieved—experimenters still have to think about it. (For example, few people would be persuaded to reject an otherwise plausible theory of equilibrium on the basis of a single play of a game by inexperienced subjects.) For this reason it is increasingly common in experimental tests of equilibrium predictions to give the subjects an opportunity to gain some experience with the game. The most common way to handle this is to have the subjects play the game many times, but against different opponents each time, to preserve the one-period nature of the games being played.[79] (In the present case, such a design would have avoided the difficulties in interpreting the observed efficiency of the bargaining results.) But when there are multiple equilibria it might be that different groups of subjects will drift to different equilibria, and that the ability of the experimenter to observe this would be sacrificed if each subject gains his experience from the whole subject pool. It was apparently an hypothesis of this sort that led Radner and Schotter to their repeated game design. The implications of this for interpreting the results would have been different had the results been different or had the issue of the maximum efficiency achievable at equilibrium not been at the center of the hypotheses to be tested.

Returning to the behavior observed by Radner and Schotter, much of it appeared to be linear; that is, subjects' prices were linear functions of their private values in each period. But there was less "shading" of buyers' bids than predicted at the linear equilibrium, and consequently more agreements were reached than predicted. (This linearity appeared to diminish as subjects gained experience, however.) The experiment also included several cells with different rules. In one of these, payoffs were by binary lotteries, to control for unobserved risk aversion that might be a factor in causing buyers to shade their bids less than the equilibrium prediction. If so, the prediction for the binary lottery games is that the "correct" amount of shading would be observed. The results were not consistent with this hypothesis. Similarly, the one-period equilibrium prediction that efficiency would be lower if an ultimatum game replaced the split-the-difference mechanism was not confirmed in a cell in which the ultimatum game was played. But in these games also, the fact that repeated games rather than single-period games were observed suggests that further experimentation might clarify what was observed.

Another experiment in which greater efficiency was observed than was predicted is reported by Rapoport, Erev, and Zwick (forthcoming). They studied a multiperiod game in which a seller negotiates with a buyer over the price of an indivisible good. It is common knowledge that the good has zero value to the seller, but its value to the buyer, v, is known only to the buyer: the seller knows only that v was drawn from a deck of 101 cards numbered 0 through 100. In each period $t = 0,1,\ldots$, the seller (the *un*informed player) sets a price, which the buyer accepts or rejects. If the buyer accepts price p in period t, then the game ends and the seller's payoff is $\delta^t p$ and the buyer's payoff is $\delta^t(v - p)$, where δ is a (common) discount factor between 0 and 1. If the buyer rejects, the game continues to the next period, and the seller again sets a price, unless the quantity $\delta^t(v - p)$ has become smaller than \$1, in which case the game is terminated by the experimenter.[80] Subjects played the game eighteen times, against changing opponents, as both buyers and sellers. Each subject played under three different discount factors δ; 0.9, 0.66, and 0.33.

This game is modeled on the infinite horizon game examined by Fudenberg, Levine, and Tirole (1985), who identified a generically unique sequential equilibrium path when the (continuous) distribution from which the buyer's value is drawn has a support that strictly exceeds the seller's value.[81] In this equilibrium, the price set by the seller declines in each period, in a nonlinear way, from the initial price that the seller chooses. And this initial price is highest when the discount factor is lowest. At equilibrium the game ends with agreement in finitely many periods, but not generally in the first period. Thus agreements are predicted to be inefficient, in the sense that costly delays are a part of the equilibrium.

Rapoport et al. observed that prices did decline monotonically, as predicted, but found the other aspects of the equilibrium predictions to be less descriptive. First, the initial prices set by sellers increased as the discount factor increased, contrary to prediction. Second, as sellers gained experience, they adopted a strategy of setting linearly declining prices. Finally, buyers tended to accept sooner than predicted. As in the experiment of Radner and Schotter, the deviations from equilibrium resulted in a higher than predicted efficiency of bargaining, which in this case means that agreements were reached sooner than predicted, and so suffered less discounting.

The final incomplete information experiment I will discuss, by Mitzkewitz and Nagel (1993), uses an innovative design to explore ultimatum games in a manner that also sheds some light on the behavior observed in complete information ultimatum games. They explored two different kinds of ultimatum games. In each of them, the amount to be divided (the "size of the cake") was first determined by the roll of a die to be an integer amount between 1 and 6. (All payoffs were counted in an artificial currency called "thalers," with one thaler worth 1.20 DM. Proposals could only be made in units of 0.5 thalers, which the authors remark was the price of a cup of coffee in the student cafeteria at the University of Bonn.) Only the proposer (player A) was informed of the size of the cake, while the accepter/rejecter, player B, knew only that the probabilities of each size from 1 to 6 were equal. Subjects played one of two different ultimatum games, an "offer game" or a "demand game," eight times against a different, anonymous other player each time. A given subject always played the same game and was always in the same position (A or B).

In the "offer game," after the die is thrown and player A is informed of the result, he makes an offer to player B. An offer may be any multiple of 0.5 that does not exceed the size of the cake: for example, if the die comes up 2, there are five feasible offers: 0, 0.5, 1, 1.5, and 2. If player B rejects the offer, both players receive zero. If player B accepts, then she (player B) earns the amount offered, while A earns the actual value of the cake minus the amount offered. So after player B hears the offer, she knows precisely what she will earn if she accepts, but she does not know what player A will earn (except if she is offered 5.5 or 6.0, in which case she can deduce that the cake was of size 6).

The demand game proceeds like the offer game, except that after player A is informed of the size of the cake, he communicates to player B a demand (of what he, player A, will receive) rather than an offer. If B rejects the demand, both players receive zero. If she accepts the demand, then player A earns what he demanded, while B earns the actual value of the cake minus what A demanded. So after player B hears the demand, she knows what player A will earn if she accepts his proposal, but she doesn't know what she will earn herself (unless the demand is 5.5 or 6.0, in which case she can deduce that she will receive 0.5 or 0.0, respectively).

While these two games each have numerous equilibria, the sequential equilibria all involve player B receiving either 0.0 or 0.5, that is, either zero or the smallest monetary unit. The equilibria at which B receives 0.0 are weak, in the sense that B is indifferent between accepting and rejecting the proposal at those equilibria, so the authors concentrate on the sequential equilibria that are strict on the equilibrium path—i.e., which give both players a positive incentive not to deviate. In both games, therefore, the prediction is that player B will receive 0.5, the smallest monetary unit, regardless of the size of the cake, and that player A will receive the rest. Note that *no* disagreements are predicted (unlike the ultimatum game whose equilibrium predictions were studied by Forsythe et al., in which the *un*informed player made the proposal).

The prediction that player B will receive no more than the smallest unit (and that no disagreements will occur) is of course familiar from the case of ultimatum games with complete information, where we have observed that the experimental results are far from the equilibrium prediction. Mitzkewitz and Nagel point out that if simple envy is the reason small positive offers are rejected in complete information games, then we might expect different behavior in these incomplete information games, in which player B cannot directly compare the payoffs to the two players. This is particularly so in the offer game, in which player B never learns the payoff of player A.

Following Selten (1967), the authors employ the "strategy method" in their experimental design. Both players are required to submit complete *strategies* for the game, *before* the die is thrown and player A is informed of the size of the cake. That is, player A is required to submit in advance the offer or demand he will make, depending on the size of the cake, for each of the six possible outcomes of the toss of the die. And player B is required to indicate whether she will accept or reject each of the thirteen possible offers or demands from 0.0 to 6.0.

While the data that they gather in this way is quite complex, a sense of the behavior of player As can be gotten from considering the modal offer for each possible cake size (pooled over all eight rounds). In the offer game, for cake sizes (1, 2, 3, 4, 5, 6), the modal *offers* are (0.5, 1, 1.5, 2, 2, 2).

That is, the modal behavior in the offer game is for player As to offer half the cake until it gets to be size 4, and to continue offering 2 for cakes of size 5 and 6. The mean percentage of the cake which player A proposes to keep for himself rises as the cake size increases. The authors note that in this weak sense, the data conforms to the prediction of the strict sequential equilibrium.

By contrast in the demand game, the modal *demands* for cake sizes 1 to 6 are (1, 2, 3, 3, 3, 3). That is, the modal behavior in the demand game is for player As to demand all of the cake until it reaches size 3, and to offer the remainder to player B for cakes of size 4, 5, and 6.

A sense of the player B's behavior can be gotten from considering the mean frequencies with which different offers or demands were accepted. A particularly illuminating comparison is between offers of 0.5 in the offer game and demands of 5.5 in the demand game, since in both cases player B is sure to receive 0.5. In the offer game, this is accepted 51 percent of the time, while in the demand game (when it means player A will get 5.5) it is accepted only 24 percent of the time—i.e. less than half as often. (The authors characterize this as "resistance to visible unfairness" on the part of players B in the demand game.) In the offer game, the

rate of acceptance rises to 96 percent with offers of 2, and 99 percent with offers of 2.5, while in the demand game the acceptance rate does not rise to over 90 percent until the demand drops to 1.5, and it does not reach 99 percent until the demand is only 0.5. So here also there is a considerable difference between the two games. But in both games there is a substantial frequency of disagreement, contrary to the prediction.

In analyzing the data round by round, the authors observe another difference between the offer and demand games. In the demand games, but not in the offer games, it may happen that a player B accepts a proposal, expecting to receive a positive payoff, but instead receives zero. The authors observe that, following such a "failure by accepting," player Bs are more inclined to reject subsequent demands, but that no parallel pattern is present in the offer games.

Thus in the offer game, player As seem reluctant to offer too little, while in the demand game they hesitate to demand too much. Since player B's judgment of what is "too little" or "too much" must be made in ignorance of the actual size of the cake, the difference between the offer game and the demand game is marked.[82] The differences observed between the offer and demand games, which both have the same strict sequential equilibrium predictions when subjects are assumed to care only about their own payoffs, provide further support for the proposition that a descriptive theory of bargaining behavior will need to take account of more complex kinds of preferences.

In summary, these experiments suggest that there remains considerable room for improvement in our understanding of the causes of disagreement and delay. Even in this group of experiments designed to test the most complex and subtle of the strategic theories of bargaining, rates of disagreement are observed that are both higher and lower than predicted in various circumstances. More generally, in each of these five incomplete information experiments, the equilibrium predictions capture at least some of the qualitative features of the data, but fall considerably short of being perfect predictors. Speaking of the "weak tendency towards the equilibrium outcome" observed in their offer games, Mitzkewitz and Nagel write (42), "We are curious whether believers in the descriptive relevance of game theory find this result encouraging or disappointing. . . ." The same could be asked about the results of each of these incomplete information experiments and, indeed, of many of the experiments discussed in this volume.

In this regard, while I have focused more on the results of these experiments than on the broadest conclusions their authors draw from them about the status of game theory as a descriptive theory, I would be remiss not to mention that these experimenters too express the full range of opinions on this subject. (Among these incomplete information experimenters, Radner and Schotter have been the most optimistic in their conclusions about the extent to which the data they observe conforms to equilibrium predictions for the game they studied, and, based on qualitatively similar evidence, Rapoport, Erev, and Zwick have been perhaps the least optimistic.) While all of these experiments have shown the predictive value of some of the qualitative predictions of some equilibria of the game (or of a closely related game), and the failure of others, different investigators assign these different importance. In this respect the situation in the emerging experimental study of bargaining under incomplete information mirrors the tests of game theoretic models generally.

There is, however, something special about theories of incomplete information that may continue to make their experimental evaluation not only especially difficult, but also especially susceptible to controversy. I have already referred to the fact that theories of incomplete information present special problems of experimental control, since they depend on the beliefs that players maintain. These problems of control are complicated by the fact that very small changes in beliefs—sometimes even arbitrarily small changes in beliefs about events that are predicted not to occur—can support very different equilibria. Part of the appeal of such models in the theoretical literature is precisely that they can be used to account for some observed behavior in terms of essentially unobservable parameters. But to the extent that the predictions derived from these models depend on those aspects of the experimental environment that are hardest to control or to observe, tests of these models will present continuing challenges to experimental design. And the interpretation of experimental tests will likely leave room for differences of opinion about whether or not the relevant features have been adequately controlled.

c. A Digression on the Strategy Method

This is a good place to pause and consider some of the advantages and disadvantages of the strategy method—i.e. of simultaneously asking all players for strategies (decisions at every information set) rather than observing each player's choices only at those information sets that arise in the course of a play of the game.[83] The obvious disadvantage is that it removes from experimental observation the possible effects of the timing of decisions in the course of the game. Thus, for example, in ultimatum games played by the strategy method, it will not be possible to observe any effects that may be due to the accepter/rejecter making her decision *after* the proposer has made his decision, knowing what has been proposed.[84]

The equally obvious advantage to collecting full strategies from the participants is that it allows the experimenter to acquire data on all information sets of the game, not just those that are actually reached in the course of the game. Thus, for example, in the experiment of Mitzkewitz and Nagel discussed above, it was possible to observe acceptance and rejection decisions for all offers, not merely those that were actually made. (In comparison, recall Figure 4.7c, discussed earlier in connection with Roth et al. [1991], in which reliable data on acceptances and rejections could only be gathered from offers that had been made sufficiently often.) Furthermore, as Mitzkewitz and Nagel note, observing subjects' entire strategies, rather than just the moves that occur in the game, may give insight into their motivation. Thus a subject in their offer game who offers half of a small cake might be thought to be "trying to be fair," but if we observe that the same subject would have offered less than half of a large cake, we may have reason to reconsider this hypothesis.

Finally, a difference that is not obviously either an advantage or a disadvantage, but which may be a cause of different outcomes for games played by the strategy method versus those played in the ordinary manner, is that having to submit entire strategies forces subjects to think about each information set in a different way than if they could primarily concentrate on those information sets that arise in the course of the game. This is a similar point to that raised in the debate about why certain sorts of "structured" experience may have different effects than simple experience with the ordinary play of a game, as discussed earlier.

In summary, if a game has many information sets, then changing from ordinary play of the game, in which subjects may make decisions at different times and with varying information, to having subjects simultaneously make all potential decisions at the same time, amounts to a significant change in the game itself. Formally, a game that has many information sets when played in the ordinary manner is transformed by the strategy method to a game in which each player has only a single information set. However, this is a change that leaves many game theoretic predictions unchanged, since it is equivalent to going from the game represented by the extensive form to the game represented by the strategic ("normal") form. There is thus room for experiments focused on determining for which kinds of games there may be significant differences in observed behavior when the strategy method is used. The results of such experiments have the potential not only to illuminate an important issue of experimental methodology, but also to point to domains in which the transformation from the extensive to the strategic form representations of a game may conceal important features of the game.

D. Deadlines

The previous section discussed experiments designed to test the predictions of particular theories. The present section discusses a class of phenomena in which experimental observation preceded the theories that have now been proposed. The phenomena in question concern the distribution of agreements over time in bargaining environments in which there is a deadline. This is a topic that has been widely discussed in an anecdotal way, but which turns out not to have been subjected to a great deal of systematic study.

Roth, Murnighan, and Schoumaker (1988) analyzed the distribution of agreements over time from four bargaining experiments, three of which were previously published experiments designed to test hypotheses unrelated to the timing of agreements.[85] In each of these experiments, unstructured anonymous bargaining was permitted to proceed between nine and twelve minutes, via terminals in a computer laboratory that automatically recorded the time of each agreement. Players could transmit English language messages to each other via terminals, monitored to prevent violations of anonymity, and sometimes subject to other restrictions. If agreement was not reached in the specified time, each player received zero. The last three minutes of bargaining were marked by a "time remaining" clock on the screen.

Although these experiments displayed considerable variation in the terms of agreements that were reached, and in the frequency of disagreements, there were substantial similarities in their distribution of agreements over time. In all experiments there was a high concentration of agreements near the deadline.

Overall, slightly less than half of all agreements were observed in the final thirty seconds of bargaining. Of those agreements reached in the final thirty seconds, approximately half were reached in the final five seconds. And of those agreements reached in the final five seconds, approximately half were reached in the final second.

Figures 4.9a and 4.9b display the distribution of agreements over time in all observations from the fourth new experiment reported in Roth et al. (1988).[86] The figures are typical of the distributions observed in all four experiments. In this experiment, bargainers engaged in binary lottery games in which the two bargainers had different prizes.[87] In the four experimental conditions, the prizes were $10 and $15; $6 and $14; $5 and $20; and $4 and $36. (So the distributions of lottery tickets that would equalize bargainers' expected incomes were (60 percent, 40 percent); (70 percent, 30 percent); (80 percent, 20 percent); and (90 percent, 10 percent) respectively.) Both bargainers were informed about the value of both prizes, and this was common knowledge. As in previous experiments of this kind, the percentage of lottery tickets obtained by the bargainer with the lower prize was observed to increase significantly across conditions as the low prize decreased. But although the observed agreements were different, neither the timing of agreements nor the variance of agreement times responded to the differences in the bargainers' prize values. That is, the distribution of agreements over time varied much less than the agreements themselves, and showed high concentrations of agreements near the deadline in all four conditions.

Roth et al. summarize the role of laboratory experimentation in this exploratory way as follows (806):

> Since last-minute agreements are widely believed to occur frequently in naturally occurring negotiations, it may be helpful to state clearly just what it is that laboratory investigations have to contribute to the study of deadline phenomena. First, while there is a great deal of anecdotal information about the frequency of "eleventh hour" agreements in naturally occurring negotiations, it has proved difficult to collect reliable data. Second, being able to study deadline phenomena in the laboratory will enable us to distinguish between alternative hypotheses in a way that the study of field data does not permit.[88] Third, while the distribution of agreements over time is one of the clearest phenomena observed in bargaining experiments to date, none of the presently available theoretical models of bargaining is able to account simultaneously for the distribution of agreements over time together with the observed patterns of agreements and substantial observed frequency of disagreements, so these results suggest clear directions for further theoretical work.

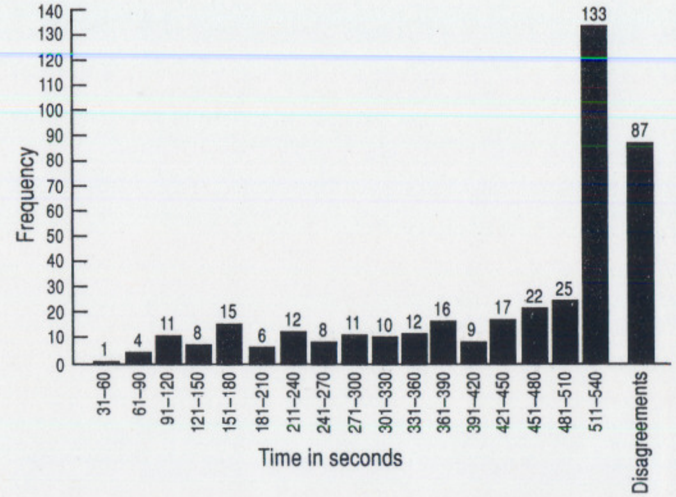

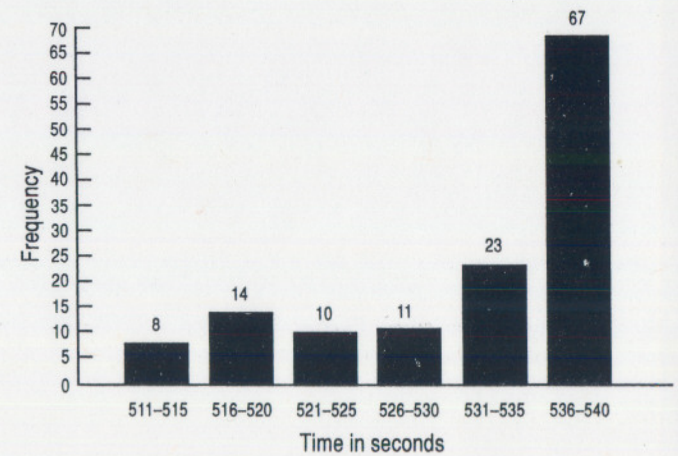

Figure 4.9. (a) Frequency of agreements and disagreements. (b) Frequency of agreements received in the last 30 seconds of bargaining. *Source:* Roth, Murnighan, and Schoumaker, 1988.

The deadline effect has indeed subsequently proved to be a fruitful subject for theoretical models. Regarding the description of the deadline effect itself, Kennan and Wilson (1993) observe that the data reported in Roth et al. (1988) can be roughly described with an "accelerated failure time model." That is, the increasing agreement rate as time expires has the property that whatever the rate is during the first x percent of the available time is roughly the same as the rate during the first x percent of the remaining $(1 - x)$ percent of the time. In such a model the time remaining can be thought of as a rescaling of the whole bargaining interval. For example, they note that Roth et al. reported 621 agreements from bargaining sessions with a nine-minute deadline, and that 11 percent of these agreements came in the first three minutes, that is, in the first third of the time available. They then observe that 13 percent of the remaining 552 agreements came in the first third of the remaining six minutes, and that the figure for the first third of the last four minutes is again around 11 percent. They note (95) that "this does not explain why the agreement rate in the initial three minutes was 11 percent, but given any initial agreement rate, the rescaling argument explains the deadline effect."

A number of strategic models have also been proposed to explain the deadline effect, in varying degrees of approximation and completeness, and in several different environments. (So it is probably more accurate to speak of these papers as exploring multiple related deadline effect*s*, some with a more stylized relation to empirical evidence than others.) These models fall into two main categories. As in the earlier discussion of game-theoretic models of disagreement, these can be usefully separated into models with complete information and with incomplete information. The two classes of models suggest two distinct causes of deadline effects and thus suggest avenues for further experimentation.

In particular, two of the complete information models can be thought of as stylized "timing models," in which a deadline effect arises as bargainers jockey to get into the position of the proposer in an ultimatum game, by trying to delay until they can make the last effective offer. Fershtman and Seidmann (1993) model deadline effects in a multiperiod sequential bargaining model in which the player who will propose in each period is chosen by lottery. In their model, bargainers can make "endogenous commitments": once a bargainer rejects an offer he is committed never to take a worse offer. Equilibrium behavior in their model depends on the discount factor: if it is low, agreement is reached in the first period, but if it is high enough, then the game will end in the last period with the proposer receiving all the surplus. (Before the end, no serious offers are made in this case, since these would be rejected and—via the endogenous commitment—improve the strategic position of the player who had rejected a positive offer.)

Ma and Manove (1993) look at a complete information model in which time is continuous and in which there are random delays in transmitting and responding to offers. Players alternate offers in their model, and a player whose turn it is to make an offer can prevent the other player from taking the initiative by delaying his offer. At equilibrium player 1 waits until a critical moment near the end before making an offer, which is accepted if it arrives near enough to the deadline, but rejected if it arrives too early, in which case a counteroffer is made.

In contrast, incomplete information models such as those by Hart (1989), Ponsati-Obiols (1989), Cramton and Tracy (1992), and Spier (1992) present deadline effects as arising from attempts by bargainers to show that their private information puts them in a powerful position. Only by engaging in costly delays can they distinguish themselves from bargainers in weaker positions. Ponsati-Obiols and Spier model deadlines after which no further bargaining can take place, while Hart and Cramton and Tracy consider models in which bargaining may proceed after the deadline, but becomes significantly more costly. All of these models yield equilibria in which a substantial concentration of agreements occur around the deadline.[89]

Each of these strategic models generates testable hypotheses. Because each of the models is formulated for a different environment, tests of their precise predictions would involve an array of experimental environments as well. However, to the extent that each model is meant to identify features of bargaining that contribute to deadline effects in a variety of bargaining environments, it becomes possible to contemplate experimental tests designed to better discern and compare the extent of such contributions. In particular, theories of incomplete information suggest that deadline effects are associated with the efforts of some players to distinguish themselves on the basis of their private information, while timing models suggest that the detailed mechanics of sending and receiving offers may play a large role. Both of these effects are predicted to occur in ways that will make them hard to control (even) in the laboratory, but because the effects are different it should be possible to design experiments that help to distinguish between them.

Deadline effects also offer a window through which to examine the effect of the bargaining that takes place well before the deadline. That is to say, even if the majority of agreements take place in the final minute of a ten-minute bargaining session, for example, presumably the bargaining that occurs before the final minute influences those agreements that are reached very near the deadline. It thus seems plausible to expect a different pattern of agreements and disagreements in the final minute of a two-minute bargaining session than of a ten-minute session.[90] The nature of this difference would shed some light on the effect of the bargaining that takes place long before the deadline.

III. Concluding Remarks

We did not need experiments to tell us that bargaining is subtle and complex, but experiments have given us insights into these subtleties and complexities that would have been difficult to obtain in any other way. In these concluding remarks I will not try to comprehensively summarize the experimental results presented in this chapter; they defy easy summary, since they deal with many issues and suggest further experiments in many directions. Instead, I will try to select a few results to illustrate the kinds of things we have learned from these series of experiments. First I will consider what kinds of things we have learned about bargaining and about game-theoretic models of bargaining. Then I will briefly reflect on what we have learned about experiments more generally and about the relationship between theory and experiment.

One of the most general things that experiments demonstrate is that subjects adjust their behavior as they gain experience and learn about the game they are playing and the behavior of other subjects. This is an observation that is certainly not confined to bargaining experiments and suggests that theoretical work on learning and dynamic adjustment and adaptation may eventually prove very fruitful.

One of the particular things bargaining experiments suggest is that bargainers may be concerned with more than their own payoffs in evaluating the outcome of bargaining. On a purely methodological level, this illustrates how difficult it is, even in the laboratory, to gain complete control over the experimental environment. This means that experimental tests of some kinds of theories may require considerable perseverance and ingenuity. And this will become more relevant the more the theories being tested rely on features of the environment that are difficult to control, or on observations that are difficult to make. (The incomplete information theories of disagreement frequency come immediately to mind.)

At the same time, the degree of control available in the laboratory makes possible observations that could not be made as clearly, if at all, in field data. Experiments have proved in this way to be a powerful instrument for bringing to light overlooked regularities and for investigating hypotheses about their causes. For example, the prevalence of equal divisions in the experimental data led naturally to hypotheses that these might be intimately related to notions of fairness. Further evidence, however, made clear that whatever the extent to which notions of fairness may play a role in determining the outcome of bargaining, it is not the case that bargainers are primarily trying to *be* fair. (If they were, they would not have produced the consistent first mover advantages observed in Figure 4.2, or the difference between ultimatum games and dictator games observed in Figure 4.4, or the modal offers consistent with income maximization observed when the tenth round offers in Figure 4.7b are compared with the revenue curves in Figure 4.7d.)

This is not to say that some notions of fairness may not play an important role—recall the frequency with which disadvantageous responses are observed (e.g., in Table 4.1) when bargainers are offered what they may regard as inequitable divisions. But notions of fairness are labile and appear to respond to strategic considerations (recall Figures 4.5a and 4.5b comparing players' assessments of fair outcomes after playing optional or forced breakdown games, and recall the extreme payoff disparities that proved to be acceptable in best-shot and market games, e.g., in Table 4.3 and Table 4.4 or Figure 4.7a).

What does this tell us about the game-theoretic models that the sequential bargaining experiments I concentrated on in this chapter were primarily designed to test? If we take the position that the assumption that players are concerned only with their own payoffs is a critical part of these models, then the evidence provides little support for their general predictive power (although in some games, such as the best shot, market, and impunity games, they work just fine, an observation that raises theoretical questions of its own). But if we take the position that the models only yield predictions after we have accurately determined the preferences of the players, then we have to be more cautious. It is in just this sense that I would not want to claim that the results for ultimatum games provide a test of subgame perfectness per se.

On the one hand, the frequent observation of disagreements and disadvantageous responses generally is consistent with the hypothesis that, although bargainers are concerned primarily with their own payoffs, a wider range of threats are credible than are captured in refinements such as subgame perfection. Under this interpretation, less emphasis on equilibrium refinement and more emphasis on the effects of multiple equilibria might be a productive direction in which to look for theories with descriptive power. On the other hand, the observation that we see the simple equilibrium predictions supported in some games but not in others is consistent with the hypothesis that players have more complex preferences and that the outcome of a game is not only influenced by the preferences of the players, but also influences them. Under this interpretation, it would be productive to redirect some of the theoretical effort now spent on solving games and spend relatively more time on learning how to model as games the situations we wish to study. These two approaches are not mutually exclusive of course, but they illustrate how different hypotheses suggest different theoretical directions, just as they suggest different additional experiments.

Because these experiments have been designed to test particular theories, it is natural to evaluate them in terms of those theories. But (particularly because the experiments have uncovered unpredicted phenomena) it is also natural to ask to what extent the phenomena observed in the laboratory are likely to generalize to the wider world. It seems to me that there are two different reasons to entertain some healthy skepticism. The first reason is that the environments we explore in the laboratory are quite simple and artificial, precisely because they are designed to provide controlled tests of particular hypotheses. So bargaining outside of the laboratory virtually always takes place in more complex environments. Consequently, some of the phenomena that appear important in the laboratory may have much diminished importance in naturally occurring negotiations, and phenomena that have no opportunity to emerge in the laboratory may assume much more importance.

But there is also cause for optimism that many of the phenomena that appear important in laboratory negotiations are also important outside of the laboratory. To mention just a few reasons, consider that the ubiquity of disagreements is reflected in both experimental and field data (recall Table 4.5), as is the deadline effect (recall Figure 4.9 and the deadline effect in contract renegotiations noted in field data by Cramton and Tracy). Similarly, a concern for equity certainly appears to play a role in many large negotiations, such as when executive bonuses become an issue in salary negotiations with unionized manufacturing workers.

The second cause for skepticism about the generalizability of experimental results has to do with the scale of rewards that it is feasible to offer in most experiments. For example, what can we conclude from the results of ultimatum bargaining games for thirty dollars, about the likely results if the same games were played for a million dollars? In particular, how much would *you* offer if given the opportunity to propose an ultimatum division to an anonymous person picked at random from some well specified population? Let me suggest that the answer would likely depend on the population and might be different if your opposite number is picked from a group of business tycoons than from a group of manual laborers. Right away, this suggests a departure from the simple predictions of the theory, which is that you should make the same offer—of no more than a penny—regardless of the amount you are dividing and who you are dividing it with. And suppose it should turn out (perhaps because you are risk averse) that the utility maximizing offer when dividing a million dollars is in the neighborhood of a hundred thousand dollars, while the utility maximizing offer when dividing a billion dollars is only a million dollars. Should we draw much comfort from the fact that, in percentage terms, we are approaching the simple predictions of the theory? I think there would not be much comfort to draw if that is how the data from these immensely costly experiments turned out, since the fact that the percentage offered moves towards zero does not negate the fact that, contrary to the simple prediction, the amount it is sensible to offer is not negligible.

That is, while there is no reason to think that the percentages we see in the laboratory when the stakes are small are universal constants that will be observed for stakes of any size, neither is there any reason to suppose that the unpredicted phenomena uncovered in the laboratory would disappear as the stakes surpassed some threshold. Comparing the results of existing experiments already allows us to observe similar phenomena as the scale of rewards changes.

But it goes without saying that, in speculating about the outcomes of experiments that have not been conducted, I am just speculating. In this respect the situation may resemble that which faces chemists and chemical engineers. While the basic theories of chemistry apply to reactions across an enormous range of scales and while the phenomena observed in test tubes allow these theories to be tested and refined, and provide the basis for our understanding of most chemical phenomena, it is nevertheless true that new phenomena emerge as reactions are scaled up from the lab bench to the pilot plant, and from the pilot plant to commercial production. So, while I think there is every reason to believe that the phenomena that appear to be important in the economics laboratory will remain important outside of the laboratory, I do not doubt that some laboratory observations will prove of more general importance than others. To distinguish among them will require other kinds of empirical work as well.

Aside from what we have learned about bargaining and bargaining theory, the series of experiments discussed in this chapter also say something about the relationship of experiments and theory. One of the first things that strikes me, which may be a reflection of how early in the history of experimental economics we still are, is that the boundaries between experimental and other kinds of work are still very permeable. One measure of this is how often in the present chapter I have had the occasion to refer to experimental work by famous economists who are far better known for their nonexperimental work. (For example, from A to Z—Orley Ashenfelter, Ken Binmore, Roy Radner, Reinhard Selten, Hugo Sonnenschein, Manny Yaari, and Shmuel Zamir—are all Fellows of the Econometric Society who fall into this category.) This reflects the naturally close relationship between theory and experiment, particularly when the object of the experiment is to test a well-formulated theoretical proposition.

That being the case, it is worthwhile to note that there is also a certain degree of natural *separation* between theory and experiment. Although some experimental results lead quickly to new theory (recall the discussion of the deadline effect), it is perhaps more common to see experiments lead first to other experiments. For example, the ultimatum results of Guth et al. (1982) led to a series of experiments intended to test the robustness of the unpredicted phenomena they observed. Only when a clear pattern of related regularities was observed (recall the experiment of Ochs and Roth [1989]) did we start to see new game-theoretic models proposed and tested, as in the work of Bolton (1991, 1993).[91] Because the tasks of elucidating regularities and explaining them involve different kinds of effort, it should come as no surprise that the theoretical and experimental literatures will often proceed with their own agendas, on separate but intersecting paths, particularly when the object of the experiments is to explore unpredicted regularities.

Notes

*This chapter has been revised in response to comments from the participants in the 1990 Handbook Workshop in Pittsburgh and from readers of subsequent versions including particularly Gary Bolton, Ido Erev, Bob Forsythe, Werner Guth, Glenn Harrison, Charlie Holt, John Kagel, Michael Mitzkewitz, Rosemarie Nagel, Jack Ochs, Vesna Prasnikar, and Rami Zwick.

1. They write: "This experiment is complementary with experiments by Roth and Malouf (1979) and Roth and Murnighan (1982), and addresses the same basic issue."
2. Guth highlights the differences in our interpretations as follows (1988, 709–10):

> The experimental results of Roth and Malouf show that expected monetary payoffs dominate winning probabilities as a reward standard. But since expected monetary payoffs can only be equilibrated if both prizes are known, the dominant reward standard cannot be used if this prerequisite is not given. . . . This interpretation is supported by another study (Roth and Murnighan, 1982) showing that the shift toward equal expected monetary rewards is mainly caused by the fact that the player with the smaller prize is informed about both prizes. . . .
>
> "It seems justified to say that the behavioral theory of distributive justice offers an intuitively convincing and straightforward explanation for the experimental results of Roth and Malouf contradicting the most fundamental game theoretic axioms. In our view this explanation is more convincing than the approach of Roth and Schoumaker (1983). . . .
>
> "What Roth (1985) calls the focal point phenomenon is in our view just the problem of deciding between two reward standards differing in their prerequisites. One can only wonder why Roth and his coauthors do not even consider the explanation offered by the behavioral theory of distributive justice (the first version of this paper, finished in 1983, was strongly influenced by discussions with Alvin E. Roth). Probably the main reason is that this would mean to finally give up the illusion that people can meet the requirements of normative decision theory.

3. Since this is a handbook, let me emphasize the methodological point. The design of an experiment is intimately related to the kinds of hypotheses it is intended to test. The unstructured binary lottery experiments were designed to test hypotheses generated by theories

phrased in terms of bargainers' risk aversion. The more structured experiments discussed next are designed to investigate hypotheses concerning individual behavior in the course of bargaining.

4. Much of the recent theoretical work using this kind of model follows the treatment by Ariel Rubinstein (1982) of the infinite horizon case. An exploration of various aspects of the finite horizon case is given by Ingolf Stahl (1972). For a survey, see Osborne and Rubinstein (1990).
5. If payoffs are discrete so that offers can only be made to the nearest penny, for example, then there are subgame perfect equilibria at which *i* refuses to take 0 but accepts the smallest positive offer—e.g., one cent.
6. When payoffs are continuous this equilibrium division is unique, so perfect equilibrium in a two-period game calls for player 1 to offer player 2 the amount $\delta_2 k$ in the first period (and demand $k - \delta_2 k$ for himself), while in a three-period game player 1 offers player 2 the amount $\delta_2(k - \delta_1 k)$ in the first period, and demands $k - \delta_2(k - \delta_1 k)$ for himself.
7. Each subject played a single game in each session.
8. Kahneman, Knetsch, and Thaler (1986a) report an ultimatum game experiment that focuses more precisely on subjects' willingness to "punish" what they perceive as unfair behavior. This experiment will be discussed in section B.3.
9. They add: "This does not mean that our results are inconsistent with those of Guth et al. Under similar conditions, we obtain similar results. Moreover our full results would seem to refute the more obvious rationalizations of the behavior observed by Guth et al. as 'optimising with complex motivations.' Instead, our results indicate that this behavior is not stable in the sense that it can be easily displaced by simple optimizing behavior, once small changes are made in the playing conditions."
10. In the second game the subject now in the role of player 1 had no opportunity to observe that no player 2 was present, since in this experiment the two bargainers sat at computer terminals in different rooms.
11. Note that this rule makes the games more like ultimatum games, since some demands of player 1 (e.g., demands of less than 90 percent in games with discount factor of .1) can only be rejected at the cost of disagreement.
12. Similar results are observed in two period games with even more extreme equilibrium predictions by Weg and Smith (1992).
13. Recall that the unstructured bargaining experiments discussed in chapter 1 found support for the qualitative predictions about risk aversion made by theories whose point predictions were systematically in error.
14. Each of the earlier experiments was designed to correspond to the case that the players have equal discount factors, i.e., $\delta_1 = \delta_2 = \delta$, with the costliness of delay implemented by making the amount of money being divided in period $t + 1$ equal to δ times the amount available at period *t*. Since half the cells of the experimental design of Ochs and Roth require different discount rates for the two bargainers, the discounting could not be implemented in this way. Instead, in each period, the commodity to be divided consisted of 100 "chips." In period 1 of each game, each chip was worth \$0.30 to each bargainer. In period 2, each chip was worth δ_1(\$0.30) to player 1 and δ_2(\$0.30) to player 2, and in period 3 of the three period games each chip was worth $(\delta_1)^2$(\$0.30) and $(\delta_2)^2$(\$0.30) respectively. That is, the rate at which subjects were paid for each of the 100 chips that they might receive depended on their discount rate and the period in which agreement was reached.
15. Having only one payoff round helps control for income effects.
16. When the necessary data from these earlier experiments were not contained in published accounts, they were readily available from the working papers circulated by the authors. That this is a good experimental practice cannot be overemphasized, since the easy availability of data permits just these sorts of comparisons. And there is a special place in heaven for any journal editor who permits unaggregated data to be published. (Of course, most of these places remain vacant.)

17. Kravitz and Gunto (1992) investigate this hypothesis by conducting an ultimatum game experiment in which insulting or accommodating messages (prepared by the experimenters) are presented along with the offers (which are also prepared by the experimenters, although subjects are led to believe that other subjects have sent the offers). Holding offers constant, they observe a higher rate of rejection of offers accompanied by insulting messages. They remark (80) that "the effects of [the messages] illustrate the importance of nonstrategic factors in economic behavior." They also report a roleplaying experiment (in which all prizes are hypothetical) and a questionnaire study in which subjects are asked to estimate the rejection rates for various offers. From these they conclude that the primary motivation leading to offers of equal division is the fear that lower offers are likely to be rejected. Apart from its bearing on the matter at hand, this experiment also allows me to note that experiments like this one, which are primarily addressed to an audience of psychologists, often tend to have a different style from those in the economics literature, even when they build upon the same prior experiments and explore similar hypotheses. In this connection, see also Loewenstein, Thompson, and Bazerman (1989), who estimate a utility function for distributions of income in bargaining outcomes, based on data in which subjects rate how satisfied they would be with various outcomes to a hypothetical bargaining situation. They reach the conclusion, similar to that reached on the basis of some of the quite different experiments surveyed in this chapter, that subjects dislike inequalities in which they receive the smaller share much more than they dislike inequalities as such; that is, their concerns seem to focus more on not being at a disadvantage than on being fair. Although I cannot begin to explore it here, it is interesting to note how different, although complementary, styles of research may develop in different disciplines for reasons that are dictated not merely by different choices of problems or even different theoretical dispositions, but also by the history and sociology of the disciplines.
18. However, Ochs and Roth (1989) do report consistency across subgames; for example, the pattern of offers and responses observed in the second period of those three period games in which the first offer is rejected resembles the pattern observed in the first period offers and responses. This could be interpreted as indirect evidence supporting the subgame perfectness hypothesis with respect to the unobserved preferences.
19. These latter two kinds of equal divisions were different from each other even in the first round of their experiment because (92) "in contrast to the procedure used by Ochs and Roth, in the present study discounting commenced on round 1," so that bargainers who split the chips equally nevertheless have different earnings if they have different discount factors.
20. Similar results are found in a subsequent study by Weg and Zwick (1991). An earlier study using this kind of cost structure viewed from a different theoretical framework is reported in Contini (1968).
21. For both sets of discount factors he observes that about 20% of the first offers are rejected, with about 85% of these rejections followed by disadvantageous counterproposals when the discount factors were (2/3, 1/3), and about 20% when the discount factors were (1/3, 2/3). (Daughety [1993] presents a model of utilities in which envy plays a role, in which disadvantageous offers are predicted to occur more often when player 2's discount factor is low.) Zwick, Rapoport, and Howard (1992) make the related observation that the rate of disadvantageous counteroffers is sensitive to the probability of termination.
22. Garcia and Roth (in preparation) get similar replication for both inexperienced and experienced Ss in two period games with discount factors (.4,.4) and (.6,.6) with a \$10 initial pie.
23. See Yaari and Bar Hillel (1984), Kahneman, Knetsch, and Thaler (1986b), and Blinder and Choi (1990) for studies that emphasize this point in nonbargaining contexts.
24. For related experiments, see Zwick, Rapoport, and Howard (1992), Weg, Zwick, and Rapoport (forthcoming), and Kahn and Murnighan (1993). Kahn and Murnighan look at a large experimental design, varying the existence of an outside option, its size, the probability of termination following a rejection, the discount factor, and which player is the first mover. (They refer to the resulting experiment as a "Noah's ark $2 \times 2 \times 2 \times 2 \times 2$ design") They find

some qualitative support for game-theoretic predictions, but only weakly, and observe that, contrary to the equilibrium prediction, players often exercised their outside option. In summation, they state "Previous demand and ultimatum game research . . . also contributes to the idea that game theory does not predict at all well in these situations. But these studies often concluded that concerns for fairness or altruism might explain the results. . . . But the data here find no support for fairness or altruism either. Instead, these data make a strong case for Ochs and Roth's (1989) hypothesis that players focus on a minimally acceptable offer. . . ." Weg et al. (1992) interpret their results as being more favorable to game-theoretic predictions, although they surmise on the basis of their results that some of Binmore et al.'s (1991) observations may be sensitive to the choice of parameters.

25. Guth and Tietz (1988) write (113): "Our hypothesis is that the consistency of experimental observations and game theoretic predictions observed by Binmore et al . . . is solely due to the moderate relation of equilibrium payoffs which makes the game theoretic solution socially more acceptable." They note that Binmore et al. (1985) examined two-period bargaining games whose equilibrium prediction was for payoffs in the ratio 3:1. In their own experiment, Guth and Tietz employed equilibrium payoff ratios of 9:1. So the equilibrium payoff ratio in these best shot games is virtually identical to those in the bargaining games discussed by Guth and Tietz, since \$3.70/\$0.42 = 8.8.

26. It is worth pausing to consider some of the statistical issues that arise in formally analyzing this kind of experimental data. The fact that each subject played ten consecutive games means that the data from different periods of the same game cannot be assumed to be independent. And not only autocorrelation, but also potential learning effects (diminishing variance by periods) raise questions that need to be addressed in analyzing the data. There remains considerable room for improvement in econometric methods and tests to address these issues. In Prasnikar and Roth (1992), we approached them as follows. Let $y_{it} = \mu_t + \epsilon_{it}$, where i indexes individuals and t indexes periods. Consider the following error structure:

$$(*)\ \epsilon_{it} = \rho\, \epsilon_{it-1} + u_{it}, \qquad E(u^2_{it}) = \sigma^2_t$$

and $E(\epsilon_{it}, \epsilon_{jt}') = 0$ if $i \neq j$. To test whether σ^2_t is constant across t, we used the Breusch-Pagan (score) test. The test statistics are 87.59 for the full information game, 17.95 for the partial information game, and 27.48 for the ultimatum game. Since the critical value is $\chi^2(0.95; 9) = 16.90$, this indicates the presence of heteroscedasticity. We corrected for the presence of heteroscedasticity using White's (1980) consistent estimator of Σ. To test for autocorrelation, we estimated ρ in the above equation while imposing the constraint $\sigma^2_t = \sigma^2$. The estimates of ρ are 0.247 (standard error = 0.109) for the full information game, 0.644 (standard error = 0.076) for the partial information game, and 0.694 (standard error = 0.081) for the ultimatum game. Thus we also found evidence of positive autocorrelation. A test of the joint null hypothesis of no heteroscedasticity and no autocorrelation produced a test statistics of 21.43 which is greater then the critical $\chi^2(0.95; 10) = 18.30$).

27. Twenty subjects participated, each playing ten rounds. In each round two markets, A and B, operated simultaneously, and buyers were switched between the markets from round to round so that the composition of the markets was not the same in any two rounds. In each round every buyer submitted a price, and the maximum price in each market was reported to the seller in that market, who could accept or reject it. The transactions were then made public (by being recorded on a blackboard). Successful buyers were identified only by anonymous identification numbers. If more than one buyer offered the maximum price (and it was accepted) then one of those buyers would be chosen at random to complete the transaction.

28. The computation of pure strategy perfect equilibria is straightforward. 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. So if the high bid is no greater than \$9.95, all bids must be equal. But 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 since he would win with certainty instead of with probability 1/9. So 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 at 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. So there are many equilibria, but only two equilibrium prices, \$10.00 and \$9.95. And the situation is the same when we consider perfect equilibria in mixed strategies.

29. There are two kinds of equal-payoff outcomes: if all buyers offer a price of \$1, every player has an expected payoff of \$1, and if all buyers offer a price of \$5, the successful buyer will have the same payoff as the seller.

30. In the subsequent experiments with this game discussed next, the transaction price has sometimes settled down at \$9.95 (i.e., at the other equilibrium price).

31. It is noteworthy that the high bids were not submitted by a small proportion of the buyers (in which case we might have supposed that the high bidders were unrepresentative of the buyer population). *Half* of the buyers (9 out of 18) submitted at least one bid of \$10, and in the last period 6 out of 18 buyers submitted bids of \$9.95 or \$10.00.

32. Regarding the sequential market game, Prasnikar and Roth (1992, 885) emphasize this point as follows:

> Note that we are not claiming that the dynamics that led to equilibrium in the later rounds of this game are necessarily due to simple income maximization, although it would be surprising if this did not play some role. To be clear about what we mean, it may be useful to speculate a little, beyond the evidence, about buyers' motivations. Consider a hypothetical buyer whose preference for equality is such that his very 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 *this* game his non-monetary 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.

33. At the time the experiment was conducted, Ljubljana, which is the capital of Slovenia, was a part of Yugoslavia.

34. This 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.

35. After the Ljubljana data were collected, a devaluation reduced Yugoslav currency units by a factor of 10,000.

36. The observed distributions are significantly different for every pair of countries except the United States and Slovenia, and the between country differences are larger than the differences between groups within a given country. (Because the distributions are highly asymmetric, the statistical test used is the Mann-Whitney U test, which is based on the rank of each observation in the sample distribution.) Other observations of bargaining behavior in different subject pools are found in Spiegel, Currie, Sonnenschein, and Sen (1990), Carter and Irons (1991), Kagel, Kim, and Moser (1992), and Eckel and Grossman (1992a,1992b). Carter and Irons report some subject pool differences between economics and psychology students in

one period play, while Kagel, Kim, and Moser find no such differences in an experiment with repeated play against different opponents (although the experiment of Kagel et al. is primarily designed to investigate the effects of information differences in bargaining, so there are respects in which it and the experiment of Carter and Irons are difficult to compare). Eckel and Grossman report gender differences.

37. In the tenth round, 19 percent of offers were rejected in the United States, 23 percent in Slovenia, 14 percent in Japan, and 13 percent in Israel.

38. Note how the formal structure of this game is very different from that of an ultimatum game, in which players know which role they will play, and move sequentially, rather than simultaneously. For example, because players move simultaneously, this game has no subgames, unlike the ultimatum game. This difference is magnified in the repeated game: in a repeated ultimatum game, players can receive experience on past plays of the game, for example, on how player 2's reacted to offers which were made, but not on how they would have reacted to other offers that were not made. That players have this information in the present design vastly increases the number of information sets, and thus the strategy sets, of the players. How issues like these influence the design of experiments will be briefly discussed later, in the section on the strategy method. The experiments referred to in that section share with this one the feature that subjects simultaneously select entire strategies. But Harrison and McCabe's design differs from the typical "strategy method" experiment in which subjects only get feedback on the actual play of the game, rather than on the entire strategies chosen by other subjects. Thus the change in the repeated game under Harrison and McCabe's design is more substantial than that discussed later.

39. In this family of models, which can be applied to games with finite pure strategy sets, each player initially has some propensity to play his ith pure strategy, given by some real number q_i, and the probability that he plays his ith pure strategy the first time he plays the game is $q_i/\Sigma q_j$, where the sum is over all his pure strategies j. If the ith pure strategy is played at stage k and the player receives a payoff of x, then the propensity to play strategy i is updated according to the payoff received (e.g., by replacing q_i with $q_i + x$), so that pure strategies that have been played and have met with success tend over time to be played with greater frequency than those that have met with less success.

40. The reason is that the propensity to make very low offers falls more quickly than the propensity to accept very low offers rises. This is because the difference between accepting and rejecting a very low offer is small and thus has only modest impact on the propensities of player 2 to reject small offers, while the difference for player 1 between having a very low offer rejected, and earning zero, or having a moderately low offer accepted, and consequently earning more than half the pie, is much larger, and more quickly encourages player 1 to abandon very low offers in favor of somewhat larger ones. Once player 1 makes very low offers less often, there is even less pressure on player 2s to learn not to reject them, and so on.

41. See also Gale, Binmore, and Samuelson (forthcoming) for some simulations of ultimatum game play using replicator dynamics motivated by biological evolution.

42. Let me hasten to add that there is no reason to think that subjects use the particular simple learning rule just described. In fact, there are good reasons to think that they do not, for example, the convergence observed in the simulations takes many more iterations than the ten iterations needed to produce the same behavior in the experiments. The fact that the simple simulated learning rule may be very different from those used by the experimental subjects, but both sets of rules produce similar intermediate term outcomes, suggests that the phenomena discussed here may be quite robust—i.e., that very different learning rules will converge to perfect equilibrium in the best shot and market games, but will not converge to perfect equilibrium in the ultimatum game.

43. And of course, different theoretical developments suggest different ways to analyze experimental data. In just such a way, Fudenberg and Levine (1993c) reanalyze the best shot data of Prasnikar and Roth (1992) and the ultimatum data of Roth et al. (1991), based on the learning-based equilibrium notions set out in Fudenberg and Levine (1993a,1993b).

44. A bio/bibliographic note: Janet Currie is the former Janet Neelin of Neelin et al. (1988).

45. See, e.g., Kennan and Wilson (1990b), and Card (1990) on labor disputes. And in a study of legal disputes arising from private antitrust litigation, Salop and White (1988, Table 1.9) report a disagreement rate (interpreted as a lack of either a settlement or a dismissal) of about 25 percent, on a sample of almost 2,000 cases, and they cite other studies that find comparable rates. Kennan and Wilson (1993) further observe that disagreement rates in legal disputes substantially underestimate the inefficiencies associated with such disputes, and they cite studies indicating that the sum of the attorneys' fees often exceed the amount collected by successful plaintiffs. And Salop and White (1988, 43) estimate that in their sample "the litigation costs of settled cases were 70 to 80 percent of those of fully litigated cases." A similar argument can be made about the level of inefficiency in labor agreements. (At the Summer School on Bargaining held by the Institute for Advanced Studies of the Hebrew University of Jerusalem in June, 1990, Ken Arrow cited featherbedding agreements reached in the transportation and printing industries as examples of how disagreement rates in labor negotiations underestimate the inefficiency rate.) Of course, in contrast to experiments, in field studies it may be difficult to determine the set of efficient agreements.

46. One exception is Rapoport, Frenkel, and Perner (1977), who employed bargaining games presented in a matrix format in which Pareto optimal mixtures were not transparent and who observed significant departures from Pareto optimality.

47. Harrison and McKee (1985) argue that it is difficult to interpret the experiment of Hoffman and Spitzer (1982) as they intended, since the fact that controllers settled for a smaller payoff than they could have taken for themselves indicates that their monetary payoffs did not serve to experimentally control their preferences, which were therefore uncontrolled. Using somewhat larger payoffs, Harrison and McKee report a lower frequency of equal divisions, but the percentage of Pareto optimal agreements remained comparably high.

In their 1985 paper, Hoffman and Spitzer report that in similar experiments in which the position of "controller" was allocated to the winner of a simple game and in which the instructions to the participants gave the controller "moral authority" to claim his prize unilaterally, the frequency of equal splits was reduced, while the frequency of Pareto optimal agreements remained high (91 percent overall). See Shogren (1992) for a related experiment using binary lottery games.

48. For example, Harrison (1992), in discussing a subsequent paper (Harrison, Hoffman, Rutstrom, and Spitzer 1987) in this stream of work, notes that the subjects had a special motivation to reach what they regarded as the "right" answer. He says (13), "In this case there was something other than financial motivation at work; it should be noted that their Professor, the experimenter, was present in the room, albeit silent and impassive." Precisely to avoid such uncontrolled sources of motivation, many experimenters routinely exclude their own students from participating in the experiments they conduct.

49. And there are a number of fairly close "between experiment" comparisons that support this conclusion. For example, Binmore, Shaked, and Sutton (1989) report an experiment in which, as in Hoffman and Spitzer (1982), one of the bargainers can unilaterally give himself a certain minimum payoff. In their experiment, in which bargaining is conducted anonymously, the player with the outside option receives no less than his outside option, and there are significant numbers of rejected offers, in sharp contrast to the results of Hoffman and Spitzer. And in another pair of closely parallel experiments, Roth and Malouf (1982) observed fewer equal splits and more disagreements in anonymous bargaining than were observed by Nydegger and Owen (1975) in face-to-face bargaining.

50. In the introduction to their 1950 study "Social Pressures in Informal Groups," Festinger, Schachter, and Back note (4) that "much of the pressure to conformity undoubtedly comes from the smaller groups within a society to which individuals belong. These pressures exist as group standards of the face-to-face group and are only sometimes formalized and made very explicit. Their enforcement depends more on relatively subtle influences and indirect pressures although these are frequently very powerful."

51. The data for the anonymous bargaining condition were collected as part of the Prasnikar and Roth (1992) study already discussed in section I.B.4. The data for the two face-to-face conditions were also collected by Prasnikar and Roth, from the same subject pool, using the same instructions modified only to accommodate the new rules of communication. Twenty-eight subjects were recruited and randomly divided into the buyer/seller positions in the two conditions so that there were seven buyers and seven sellers in each of the two conditions. In each condition, each buyer bargained once with each seller over the division of $10. One of the seven bargaining rounds was then chosen at random to be the payoff round.
52. But note again the statistical problems that arise even in the analysis of such a small data set as this, because of the lack of independence between different rounds of the same game. Recall the discussion in footnote 26 of statistics in connection with the similar experimental design of Prasnikar and Roth (1992).
53. In the unrestricted communication condition, bargainers could (and did) sometimes state that they would not accept anything less than an equal share. Formal theoretical models of how such "cheap talk" (i.e., not backed up by any formal method of commitment) might influence behavior are beginning to appear in the literature (see, e.g., Crawford 1990 or Farrell 1987). For an experiment directly motivated by such models, see Cooper, DeJong, Forsythe, and Ross (1989) (and cf. chapter 3).
54. Recall from our discussion of the experiment of Forsythe et al. (1994) that a dictator game differs from an ultimatum game in that player 2 may not accept or reject player 1's offer—whatever division player 1 proposes is the outcome of the game. So in a dictator game, unlike an ultimatum game, player 1 need not be concerned with the possibility that player 2 will reject his proposal.
55. The modal offer in a cell with Forsythe et al.'s ultimatum instructions was 50 percent, while in various cells in which the instructions sought to encourage low offers the mode was 40 percent, with a second mode at 30 percent in one cell.
56. A particularly elegant feature of Bolton and Zwick's design is the way in which the anonymous and nonanonymous conditions and the ultimatum game and impunity game conditions were made closely comparable. In the anonymity condition, player 1 began the game by putting a box, corresponding to one of the feasible (discrete) offers, in a mailbag for transmittal to player 2. Player 2 began the game with a *pair* of boxes corresponding to each possible offer, one corresponding to rejection of the offer, and one corresponding to acceptance. Player 2's boxes contained envelopes with the payoffs, and player 2 responded to player 1's offer by unsealing one of the two boxes, taking out his own envelope, and putting the box and player 1's payoff in the mailbag to be transmitted to player 1. At the end of this process player 2 once again had in his possession a pair of boxes corresponding to each possible offer (since the one sent to him by player 1 substituted for the one he opened and sent back) so that the experimenter could verify that only one box had been opened (the boxes all had seals) without knowing what offer had been sent or whether it had been accepted or rejected. At the end of a round, all remaining boxes were put in a trash bag, and it is by examining the discarded boxes after the experiment that the experimenters were able to know what offers and acceptance/rejections had been made during the round, without knowing which subjects had made them. For the nonanonymity conditions, the only difference was that instead of boxes filled with cash, the players received cards and, to be paid, had to turn in the cards to the experimenter, who would therefore know (from their cards) what had transpired.
57. Both Bolton and Zwick's and Hoffman et al.'s results were presented together at a conference on experimental economics conducted at the University of Amsterdam in September 1992. Hoffman et al. (1992) responded to the earlier criticism by including the disagreement data from the nonanonymous ultimatum games reported. The disagreement data did not support the conclusion that behavior in the "moral entitlement" conditions was moving in the direction of perfect equilibrium play, because a nonnegligible percentage of the lower offers that were elicited in these conditions had in fact been rejected. Thus the results of these two investigations are not in fact quite as different as they seemed when the disagreement data were not available.

58. Two interesting experiments that show that similar kinds of "reciprocal fairness" are not extinguished by market environments are reported in Fehr, Kirchsteiger, and Riedl (1993a,1993b).
59. Questions about volatile laboratory environments are an area in which it may be profitable for experimental economists to examine the psychology literature: the earlier mentioned work of Robyn Dawes on public goods comes to mind in this regard, as does the bargaining experiment of Deutsch and Kotik (1978). For a review of the social psychology literature on negotiation, see Thompson (1990), for the organizational behavior literature, see Neale and Bazerman (1991).
60. Marwell and Schmitt (1968), in a paper titled "Are 'Trivial' Games the Most Interesting Psychologically?" argued that laboratory games in which the underlying economic motivations are small relative to other motivations may be especially good instruments for studying these other motivations.
61. For example, the origins of the Cuban missile crisis have sometimes been attributed to Nikīta Kruschev's impression in his first face-to-face summit meeting with John Kennedy that Kennedy was indecisive and would not respond to the stationing of missiles.
62. And there is experimental evidence that negotiators who are acting as agents behave differently than when they are acting as principals: see, for example, Lamm (1978) and the references he cites. (Lamm finds that negotiators elected to be representatives of groups have more disagreements than those bargaining on their own behalf.) Similarly, Shogren (1989) reports that face-to-face negotiations proceed differently when subjects are acting as part of a team rather than as individuals, and Schotter, Snyder, and Zheng (1992) report that principals act somewhat as if they are submitting sealed bids when they give instructions to agents who will bargain on their behalf (with a consequent loss of efficiency). Of course, different kinds of inefficiencies (from the point of view of the principals) are introduced when both sides hire agents: see Ashenfelter and Bloom (1990) for some estimates of the extent to which the decision to hire lawyers in certain kinds of negotiations constitutes a prisoner's dilemma.
63. Recall the discussion of bargaining in chapter 1.
64. Recall the discussion of binary lottery games in sections II.C.2 and II.F.4 of chapter 1. In a binary lottery game, subjects are paid in lottery tickets which determine their probability of winning one of two monetary prizes, so that each subject's expected utility can be taken to be his probability of winning the larger of his two prizes.
65. Because of the high variance in the time taken to reach agreement in the cases in which agreement was reached, the only significant differences were between game 1, which had the shortest mean time to agreement, and game 4, which had the longest.
66. The doubling of pie size was achieved by doubling the number of periods of the game. Forsythe et al. also consider the effect of negotiations over long-term versus short-term contracts.
67. In games in which players knew the value of one another's prizes, the potentially focal division (h, $100 - h$) was taken to be the equal expected value division, while (50,50) was taken to be the only focal agreement for games in which the bargainers had equal prizes, and for games in which they did not know one another's prizes. Since each of the experiments summarized in the table was a binary lottery game, in the event of disagreement each bargainer had a probability of 0 of winning his prize, so $(d_1, d_2) = (0,0)$.
68. In game-theoretic terminology, this means that the game of incomplete information (in which some aspects of the game tree are not common knowledge) is *modeled* as a game of *imperfect* information, in which the game tree is common knowledge, but players have different information about an initial chance move.
69. That is, the experiments are designed to implement the games of imperfect information used to model situations of incomplete information.
70. This mode was most pronounced for the HH pairs, where 12 out of 13 observed agreements gave each player a gross payoff of 10 (and the other agreement had gross payoffs of 9.5 and 10.5). And although equal agreements were the mode in the LL case, where equal agreements

are predicted (14 out of 56 LL agreements were (10,10) splits, with the next most commonly observed agreement, a (9,11) split, being observed 7 times), 75 percent of the observations (42 out of 56) were *not* equal splits. About these, the authors say (145):

> It is interesting to look at the reasons for the occurrence of so many cases of LL-agreements where one player received more than the other. One may be tempted to think that the player with the higher payoff achieves this result by some kind of bluffing behavior which involves repetitions of demand in order to convey the impression that he is a type H player. Actually in 25 of the 42 cases of LL-agreements with unequal payoffs the player with the higher agreement payoff did not repeat his demand even once. Obviously in these cases the other player either had a lower initial demand or he lowered his demand more quickly. The player with the higher agreement payoff did not have to do anything special in order to get the higher payoff. It just happened to him that the other behaved in a "soft" way.

71. Recall that the method of controlling for risk aversion via binary lottery games had not yet been introduced at the time this experiment was conducted.
72. In a subsequent paper, Hoggatt, Brandstatter, and Blatman (1978) refer to such robots as "*Selten* robots." (See also Hoggatt [1969] for an earlier exploration of robots in experimental work.)
73. Like Hoggatt and Selten et al. (1978), Forsythe et al. conclude that the predictions of the incomplete information model about the frequency of disagreements capture important features of the observed behavior, particularly in view of the base level rate of disagreements observed even in complete information experiments (recall Table 4.5a).
74. Recall our discussion of the communication hypothesis versus the uncontrolled social utility hypothesis in connection with the *lower* frequency of disagreement observed by Radner and Schotter (1989) when increased communication was also accompanied by face-to-face interaction in the environment they studied. (Radner and Schotter's environment is further discussed next.)
75. The conclusions about risk aversion are only indirect inferences, because, as the authors note, they have attempted to neither measure nor control for risk aversion. In this respect they note (268):

> If the subjects were all selfish expected-utility maximizers . . . the heterogeneity in risk attitudes could be eliminated by using the ingenious binary lottery procedure introduced by Roth and Malouf (1979). We chose not to use this procedure, mainly because it introduces considerable additional complexity in an already complicated experimental environment. In addition, the procedure works only under assumptions which are implausible in our context: that each subject acts selfishly and obeys the compound lottery axiom.

In this regard, Forsythe and his colleagues seem to be conforming to an increasingly common practice in experimental work by attempting to introduce careful and elaborate controls for risk aversion only in experiments in which risk aversion is thought (at the outset) to be a major factor influencing behavior. (See in this connection the papers by Cooper, DeJong, Forsythe, and Ross [1989, 1990], both of which use the binary lottery procedure.) In the context of the present experiment Forsythe et al. are also careful to warn of the problems associated with subjects who may be concerned about the distribution of payoffs rather than merely with their own payoffs, since for such subjects the outcomes will no longer be binary even when binary lottery games are employed. (Recall the discussion of binary lottery games in chapter 1.)

76. Note the similarity to the conclusions of Kennan and Wilson (1993) discussed in the beginning of this chapter in connection with the "complete information" environments considered by Ochs and Roth (1989).
77. A pair of strategies that would achieve 100 percent efficiency would be for each player to state his reservation price. But this would not be an equilibrium, since when there are substantial gains from trade (i.e., when the buyer's reservation price is substantially higher than the seller's) then the buyer could have increased his profit by stating a lower price, and the seller could have increased his profit by stating a higher price. But when both players "shade" their bids in this way, they will miss those trades that could have occurred, for small profit, when the buyer's reservation price is only slightly higher than the seller's. So there is a tradeoff between efficiency and the maximization of expected profit that occurs at equilibrium.

78. The theoretical argument is roughly that high degrees of efficiency can be achieved in the many period game via an equilibrium in which deviation is deterred by the threat of adopting a highly inefficient equilibrium in the final periods. Precisely the degree of efficiency that can be achieved depends on both the length of the game and the distribution from which values are drawn, since the players can only detect deviations statistically.
79. Recall the discussion of learning in section I.
80. Subjects were not informed of the termination rule in case the payoffs became too small, and the authors report that fewer than 6 percent of the games were terminated for this reason.
81. Thus the experimental game approximates the game from which the predictions are derived, but differs in having a finite stopping rule, a discrete distribution of buyer values, and a positive probability that the buyer's value equals the seller's.
82. A similarly designed incomplete information experiment that focuses on the effect of varying the uncertainty about the amount available to be divided is reported by Rapoport, Sundali, and Potter (1992). Some related observations are made by Straub and Murnighan (1992) and by Croson (1992), who conduct ultimatum game experiments in which the amount to be divided may be either known or unknown by player 2 (but in which no attempt was made to induce a probability distribution on the unknown amount).
83. In much of the German experimental literature following Selten (1967) the term "strategy method" is reserved for a set of procedures in which subjects first gain experience playing the game in the usual way, before being asked to submit strategies, and have the opportunity to revise those strategies based on further experience.
84. Since both players know the timing structure of the game, this could have an effect not only on the accepter/rejecter but also on the proposer. This might be true for purely behavioral reasons not captured in formal game-theoretic models of idealized rationality (e.g., if subjects are unable to adhere to commitments they would wish to make before knowing what proposal will be made), but also for reasons reflected in formal game-theoretic models. For example, the notion of subgame perfect equilibrium is lost in the transition from the extensive to the strategic form of the game, since there are no subgames in a game in which players state their strategies simultaneously.
85. The previously published experiments were Roth, Malouf, and Murnighan (1981), Roth and Murnighan (1982), and Murnighan, Roth, and Schoumaker (1988), all of which are briefly described in chapter 1. Roth et al. (1988, 808) caution that "although all this data comes from laboratory experiments, there is a sense in which it is not all fully 'experimental data,' since the experiments were mostly designed for purposes other than to test specific hypotheses about agreement times."
86. This experiment was briefly discussed above as a test of the complete information model of disagreement frequency proposed in Roth (1985).
87. The general procedures were the same as those in Roth and Murnighan (1982), discussed in chapter 1.
88. "For example, labor negotiators often attribute a tendency to reach agreements just before contracts expire to the difficulty of selling any agreement to a diverse constituency if there is still time for continuing negotiations. However the deadline effect observed in our laboratory environment cannot be attributed to this, since each bargainer is bargaining strictly on his own behalf."
89. Cramton and Tracy (1992) also present new empirical evidence from field data concerning deadline effects. In data concerning 5,002 contracts culled from publications of the Bureau of Labor Statistics and the Bureau of National Affairs, they observe a sharp peak in settlement rates just *after* the deadline created by the expiration of the existing contract.

90. I don't know of any experiments that directly address this point. However, Coursey (1982) compared single play bargaining with a 1 hour deadline to repeated bargaining with a 128 second deadline and found more disagreements in the latter condition. Since most disagreements were observed in early periods of the repeated bargaining, when bargainers have the most incentives to engage in reputation building (since they bargain repeatedly with the same partner), it would be premature to draw any conclusions about the influence of the deadlines in this experiment.
91. See also Rabin (1993) for some theoretical work on fairness motivated by experimental results.

Bibliography

Albers, Wulf, and Gisela Albers. 1983. On the prominence structure of the decimal system. In *Decision making under uncertainty*, R.W. Scholz, editor. Amsterdam: Elsevier. 271–87.

Ashenfelter, Orley. 1987. A model of arbitral behavior. *American Economic Review*, 77: 342–46.

Ashenfelter, Orley, and David Bloom. 1984. Models of arbitrator behavior: Theory and evidence. *American Economic Review* 74:111–25.

———. 1990. Lawyers as agents of the devil in a prisoner's dilemma game. Mimeo.

Ashenfelter, Orley, and Janet Currie. 1990. Negotiator behavior and the occurrence of disputes. *American Economic Review, Papers and Proceedings* 80:416–20.

Ashenfelter, Orley, Janet Currie, Henry S. Farber, and Matthew Spiegel. 1992. An experimental comparison of dispute rates in alternative arbitration systems. *Econometrica* 60: 1407–33.

Axelrod, Robert. 1970. *Conflict of interest*. Chicago: Markham.

Berg, Joyce, John Dickhaut, and Kevin McCabe. 1993. Trust, reciprocity, and social Norms. Mimeo. University of Minnesota.

Binmore, Ken, Peter Morgan, Avner Shaked, and John Sutton. 1991. Do people exploit their bargaining power?: An experimental study. *Games and Economic Behavior* 3:295–322.

Binmore, Ken, Avner Shaked, and John Sutton. 1985. Testing noncooperative bargaining theory: A preliminary study. *American Economic Review* 75:1178–80.

———. 1988. A further test of noncooperative bargaining theory: Reply. *American Economic Review* 78:837–9.

———. 1989. An outside option experiment. *Quarterly Journal of Economics* 104:753–70.

Blinder, Alan S., and Don H. Choi. 1990. A shred of evidence on theories of wage stickiness. *Quarterly Journal of Economics* 1003–15.

Boden, Leslie. 1989. Dispute resolution in worker's compensation: The role of adversarial experts. Mimeo. Boston University.

Bolton, Gary. 1991. A comparative model of bargaining: Theory and evidence. *American Economic Review* 81:1096–1136.

Bolton, Gary E. 1993. The rationality of splitting equally. Working Paper. Penn State University.

Bolton, Gary E., Elena Katok, and Rami Zwick. Forthcoming. Dictator game giving: Rules of fairness versus random acts of kindness. *International Journal of Game Theory*.

Bolton, Gary E., and Rami Zwick. Forthcoming. Anonymity versus punishment in ultimatum bargaining. *Games and economic behavior*.

Burrows, Paul, and Graham Loomes. 1990. The impact of fairness on bargaining behavior. Mimeo. Center for Experimental Economics. University of York.

Card, David. 1988. Strikes and wages: A test of a signalling model. Mimeo. Princeton University.

Card, David. 1990. Strikes and bargaining: A survey of the recent empirical literature. *American Economic Review* 80:410–15.

Carter, John R., and Michael D. Irons. 1991. Are economists different, and if so, why? *Journal of Economic Perspectives* 5:171–77.

Chatterjee, Kalyan, and William Samuelson. 1983. Bargaining under incomplete information. *Operations Research* 31:835–51.

Contini, Bruno. 1968. The value of time in bargaining negotiations: Some experimental evidence. *American Economic Review* 58:374–93.

Cooper, Russell, Douglas V. DeJong, Robert Forsythe, and Thomas W. Ross. 1989. Communication in the battle of the sexes game: Some experimental results. *RAND Journal of Economics* 20:568–87.

———. 1990. Selection criteria in coordination games: Some experimental results. *American Economic Review* 80:218–33.

Coursey, Don L. 1982. Bilateral bargaining, Pareto optimality, and the empirical frequency of impasse. *Journal of Economic Behavior and Organization* 3:243–59.

Cramton, Peter C., and Joseph S. Tracy. 1992. Strikes and delays in wage bargaining: Theory and data. *American Economic Review* 82:100–21.

Crawford, Vincent P. 1990. Explicit communication and bargaining outcomes. *American Economic Review, Papers and Proceedings* 80:213–19.

Croson, Rachel T. A. 1992. Information in ultimatum games: An experimental study. Mimeo. Department of Economics. Harvard University.

Currie, Janet. 1989. Who uses interest arbitration? The case of British Columbia's teachers 1947–1981. *Industrial and Labor Relations Review* 42:363–79.

Currie, Janet, and Sheena McConnell. 1989. Strikes and arbitration in the public sector: Can legislation reduce dispute costs? Mimeo. UCLA.

Daughety, Andrew F. 1993. Socially-influenced choice: Equity considerations in models of consumer choice and in games. Working paper #93–01. Department of Economics, University of Iowa.

Dawes, Robyn M. 1991. Social dilemmas, economic self-interest, and evolutionary theory. In *Recent research in psychology: Frontiers of Mathematical Psychology: essays in honor of Clyde Coombs*, D. R. Brown and J.E.K. Smith, editors, New York: Springer-Verlag.

Deutsch, Morton, and Paul Kotik. 1978. Altruism and bargaining. In *Bargaining Behavior*, Contributions to Experimental Economics 7, Heinz Sauermann, editor, Tubingen: J.C.B. Mohr. 20–40.

Eckel, Catherine C., and Philip Grossman. 1992a. Chivalry and solidarity in ultimatum games. Working paper. Virginia Polytechnic Institute.

———. 1992b. The Price of Fairness: Gender Differences in Punishment Games. Working paper. Virginia Polytechnic Institute.

Farrell, Joseph. 1987. Cheap talk, coordination, and entry. *RAND Journal of Economics* 18: 34–9.

Fehr, Ernst, Georg Kirchsteiger, and Arno Riedl. 1993a. Does fairness prevent market clearing? An experimental investigation. *Quarterly Journal of Economics* 108:437–59.

———. 1993b. Gift exchange and ultimatums in experimental markets. Mimeo. University of Technology. Austria.

Fershtman, Chaim, and Daniel J. Seidmann. Forthcoming. Deadline effects and inefficient delay in bargaining with endogenous commitment. *Journal of Economic Theory*.

Festinger, Leon, Stanley Schachter, and Kurt Back. 1950. *Social pressures in informal groups: A study of human factors in housing*. Stanford, Calif.: Stanford University Press.

Foddy, Margaret. 1989. Information control as a bargaining tactic in social exchange. *Advances in Group Processes* 6:139–78.

Forsythe, Robert, John Kennan, and Barry Sopher. Forthcoming. Dividing a shrinking pie: An experimental study of strikes in bargaining games with complete information. In *Research in experimental economics*, R. Mark Issac (editor). Greenwich, Conn.: JAI Press.

———. 1991. An experimental analysis of bargaining and strikes with one sided private information. *American Economic Review*. 81:253–78.

Forsythe, Robert, Joel L. Horowitz, N. E. Savin, and Martin Sefton. 1994. Fairness in simple bargaining experiments. *Games and economic behavior*.6:347–69.

Fouraker, Lawrence E., and Sidney Siegel. 1963. *Bargaining behavior*. New York: McGraw Hill.

Fudenberg, Drew, and David K. Levine. 1993a. Self-confirming equilibrium. *Econometrica* 61:523–46.

———. 1993b. Steady state learning and Nash equilibrium. *Econometrica* 61:547–73.

———. 1993c. How irrational are subjects In extensive-form games? Discussion paper 14. Harvard Institute of Economic Research, November.

Fudenberg, Drew, David K. Levine, and Jean Tirole. 1985. Infinite-horizon models of bargaining with one-sided incomplete information. In *Game-theoretic models of bargaining*, A. E. Roth, editor, Cambridge: Cambridge University Press. 73–98.

Gale, John, Kenneth G. Binmore, and Larry Samuelson. Forthcoming. Learning to be imperfect: The ultimatum game. *Games and economic behavior*.

Grammer, Karl, Wulf Schiefenhovel, Margret Schleidt, Beatrice Lorenz, and Irenaus Eibl-Eibesfeldt. 1988. Patterns on the face: The eyebrow flash in crosscultural comparison. *Ethology* 77:279–99.

Guth, Werner. 1988. On the behavioral approach to distributive justice: A theoretical and experimental investigation. In *Applied Behavioral Economics*, S. Maital, editor, New York: New York University Press, 2:703–17.

Guth, Werner, R. Schmittberger, and B. Schwarz. 1982. An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior and Organization* 3:367–88.

Guth, Werner, and Reinhard Tietz. 1988. Ultimatum bargaining for a shrinking cake—An experimental analysis. In *Bounded Rational Behavior in Experimental Games and Markets*, R. Tietz, W. Albers, R. Selten, editors, Springer: Berlin.

———. 1990. Ultimatum bargaining behavior: A survey and comparison of experimental results. *Journal of Economic Psychology* 11:417–49.

Guth, Werner, and Menahem Yaari. 1990a. An evolutionary approach to explaining reciprocal behavior in a simple strategic game. Mimeo.

———. 1990b. Incomplete information about reciprocal incentives: An evolutionary approach to explaining cooperative behavior. Mimeo.

Harrison, Glenn W. 1990. Rational expectations and experimental methods. In *Rational Expectations and Efficiency in Futures Markets*, B. A. Goss, editor, London: Routledge.

Harrison, Glenn W. 1992. Flat payoff functions and the experimentalists: Reply to the critics, Mss., February, University of South Carolina.

Harrison, Glenn W., and J. Hirshleifer. 1989. An Experimental Evaluation of Weakest Link/Best Shot Models of Public Goods. *Journal of Political Economy*. 97:201–225.

Harrison, Glenn W., Elizabeth Hoffman, E. E. Rustrom, and Matthew L. Spitzer. 1987. Coasian solutions to the externality problem in experimental markets. *Economic Journal* 97: 388–402.

Harrison, Glenn W, and Kevin A. McCabe. 1991. Testing non-cooperative bargaining theory in experiments. Working paper B-91–03. College of Business Administration, University of South Carolina. Forthcoming in *Research in Experimental Economics, vol. 5*, R. Mark Isaac, editor, Greenwich, Conn.: JAI Press.

———.1992b. Expectations and fairness in a simple bargaining experiment. Working paper B-92–10. September, University of South Carolina.

Harrison, Glenn W. and Michael McKee. 1985. Experimental evaluation of the Coase theorem. *Journal of Law and Economics* 28:653–70.

Harsanyi, John C. 1967–1968. Games with incomplete information played by 'Bayesian' players. Parts I-III. *Management Science* 14:159–82, 320–34, 486–502.

Harsanyi, John C., and Reinhard Selten. 1972. A generalized Nash solution for two-person bargaining games with incomplete information. *Management Science* 18:80–106.

Hart, Oliver. 1989. Bargaining and strikes. *Quarterly Journal of Economics* 104:25–44.

Hoffman, Elizabeth, Kevin McCabe, Keith Shachat, and Vernon L. Smith. 1991. Preferences, property rights and anonymity in bargaining games. Mimeo. December, University of Arizona. rev. June 1992.

Hoffman, Elizabeth, Kevin McCabe, and Vernon L. Smith. 1993. On property rights and the monetary stakes in ultimatum game bargaining. Mimeo. University of Arizona.

Hoffman, E., and M. L. Spitzer. 1982. The Coase theorem: Some experimental tests. *Journal of Law and Economics* 25:73–98.

———. 1985. Entitlements, rights, and fairness: An experimental examination of subjects' concepts of distributive justice. *Journal of Legal Studies* 14. 2:259–97.

Hoggatt, Austin C. 1969. Response of paid student subjects to differential behaviour of robots in bifurcated duopoly games. *Review of Economic Studies* 36:417–32.

Hoggatt, Austin C., Hermann Brandstatter, and Peter Blatman. 1978. Robots as instrumental functions in the study of bargaining behavior. In *Bargaining Behavior*, Contributions to Experimental Economics 7, H. Sauermann, editor, Tubingen: J.C.B. Mohr. 179–210.

Hoggatt, Austin C., Reinhard Selten, David Crockett, Shlomo Gill, and Jeff Moore. 1978. Bargaining experiments with incomplete information. In *Bargaining Behavior*, Contributions to Experimental Economics 7, H. Sauermann, editor, Tubingen: J.C.B. Mohr. 127–78.

Kagel, John H., Chung Kim, and Donald Moser. 1992. Ultimatum games with asymmetric information and asymmetric payoffs. Mimeo. University of Pittsburgh.

Kahn, Lawrence M., and J. Keith Murnighan. 1993. A general experiment on bargaining in demand games with outside options. *American Economic Review* 83:1260–80.

Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. 1986a. Fairness and the assumptions of economics. *Journal of Business* 59, no. 4, pt. 2:S285–S300.

———. 1986b. Fairness as a constraint on profit seeking: Entitlements in the market. *American Economic Review* 76:728–41.

Kandori, Michihiro. 1992. Social norms and community enforcement. *Review of Economic Studies* 59:63–80.

Kennan, John. 1980. Pareto optimality and the economics of strike duration. *Journal of Labor Research* 1:77–93.

Kennan, John, and Robert Wilson. 1990a. Theories of bargaining delays. *Science* 7 September, vol. 249: 1124–8.

———. 1990b. Can strategic bargaining models explain collective bargaining data? *American Economic Review, Papers and Proceedings*. May vol. 80: 405–09.

———. 1993. Bargaining with private information. *Journal of Economic Literature* 31:45–104.

Kravitz, David A., and Samuel Gunto. 1992. Decisions and perceptions of recipients in ultimatum bargaining games. *Journal of Socio-Economics* 21:65–84.

Lamm, Helmut. 1978. Group-related influences on negotiation behavior: Two-person negotiation as a function of representation and election. In *Bargaining Behavior*. Contributions to Experimental Economics 7, Heinz Sauermann, editor, Tubingen: J.C.B. Mohr. 284–309.

Leininger, W., P. Linhard, and R. Radner 1989. Equilibria of the sealed-bid mechanism for bargaining with incomplete information. *Journal of Economic Theory* 48:63–106.

Loewenstein, George F., Leigh Thompson, and Max H. Bazerman. 1989. Social utility and decision making in interpersonal contexts. *Journal of Personality and Social Psychology* 57:426–41.

Ma, Ching-to Albert, and Michael Manove. 1993. Bargaining with deadlines and imperfect player control. *Econometrica* 61:1313–39.

Malouf, Michael W. K., and Alvin E. Roth. 1981. Disagreement in bargaining: An experimental study. *Journal of Conflict Resolution* 25:329–48.

Marwell, Gerald, and David R. Schmitt. 1968. Are "trivial" games the most interesting psychologically? *Behavioral Science* 13:125–8.

McConnell, Sheena. 1989. Strikes, wages and private information. *American Economic Review* 79:801–15.

Mehta, Judith, Chris Starmer, and Robert Sugden. 1990. Focal points in bargaining: An experimental investigation. Mimeo. Economics Research Centre, University of East Anglia.

Mitzkewitz, Michael, and Rosemarie Nagel. 1993. Envy, greed and anticipation in ultimatum games with incomplete information: An experimental study. Discussion paper B-181. March. Universitat Bonn. *International Journal of Game Theory* 22:171–98.

Mnookin, Robert et al. 1989. Private ordering revisited: What custodial arrangements are parents negotiating? Mimeo. Stanford University.

Murnighan, J. Keith, Alvin E. Roth, and Francoise Schoumaker. 1988. Risk aversion in bargaining: An experimental study. *Journal of Risk and Uncertainty* 1:101–24.

Myerson, Roger, and Mark Satterthwaite. 1983. Efficient mechanisms for bilateral trading. *Journal of Economic Theory* 29:265–81.

Neale, Margaret A., and Max H. Bazerman. 1991. *Cognition and rationality in negotiation*. New York: The Free Press.

Neelin, Janet, Hugo Sonnenschein, and Matthew Spiegel. 1988. A further test of noncooperative bargaining theory: Comment. *American Economic Review* 78:824–36.

Nydegger, Rudy V., and Guillermo Owen. 1975. Two person bargaining: An experimental test of the Nash axioms. *International Journal of Game Theory* 3:239–349.

Ochs, Jack, and Alvin E. Roth. 1989. An experimental study of sequential bargaining. *American Economic Review* 79:355–84.

Okuno-Fujiwara, Masahiro, and Andrew Postlewaite. 1990. Social norm in random matching games. Mimeo.

Osborne, Martin J., and Ariel Rubinstein. 1990. *Bargaining and markets*. San Diego: Academic Press.

Ponsati-Obiols, Clara. 1989. The deadline effect in bargaining: A theoretical note. Mimeo. Bellcore.

Prasnikar, Vesna, and Alvin E. Roth. 1989. Perceptions of fairness and considerations of strategy in bargaining: Some Experimental Data. Preliminary draft.

———. 1992. Considerations of fairness and strategy: Experimental data from sequential games. *Quarterly Journal of Economics* August:865–88.

Rabin, Matthew. 1993. Incorporating fairness into game theory and economics. *American Economic Review* 83:1281–1302.

Radner, Roy, and Andrew Schotter. 1989. The sealed bid mechanism: An experimental study. *Journal of Economic Theory* 48:179–220.

Rapoport, Amnon, Ido Erev, and Rami Zwick. Forthcoming. An experimental study of buyer-seller negotiation with one-sided incomplete information and time discounting. *Management Science*.

Rapoport, Amnon, James A. Sundali, and Richard E. Potter. 1992. Ultimatum games with incomplete information: Effects of the variability of the pie size. Mimeo. University of Arizona.

Rapoport, Amnon, Eythan Weg, and Dan S. Felsenthal. 1990. Effects of fixed costs in two-person sequential bargaining. *Theory and Decision* 28:47–72.

Rapoport, Anatol, Oded Frenkel, and Josef Perner. 1977. Experiments with cooperative 2 × 2 games. *Theory and Decision* 8:67–92.

Roth, Alvin E. 1985. Toward a focal-point theory of bargaining. In *Game-theoretic models of bargaining*, A. E. Roth, editor, Cambridge: Cambridge University Press. 259–68.

Roth, Alvin E., and Ido Erev. Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior* 8 (January): 164–212. (Special Issue on the Nobel Symposium on Game Theory, June 18–20, Bjorkborn, Sweden.)

Roth, Alvin E., and Michael W. K. Malouf. 1979. Game-theoretic models and the role of information in bargaining. *Psychological Review* 86:574–94.

———. 1982. Scale changes and shared information in bargaining: An experimental study. *Mathematical Social Sciences* 3:157–77.

Roth, Alvin E., Michael W. K. Malouf, and J. Keith Murnighan. 1981. Sociological versus strategic factors in bargaining. *Journal of Economic Behavior and Organization* 2: 153–77.

Roth, Alvin E., and J. Keith Murnighan. 1982. The role of information in bargaining: An experimental study. *Econometrica* 50:1123–42.

Roth, Alvin E., J. Keith Murnighan, and Francoise Schoumaker. 1988. The deadline effect in bargaining: Some experimental evidence. *American Economic Review* 78:806–23.

Roth, Alvin E., Vesna Prasnikar, Masahiro Okuno-Fujiwara, and Shmuel Zamir. 1991. Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An experimental study. *American Economic Review* 81:1068–95.

Roth, Alvin E., and Francoise Schoumaker. 1983. Expectations and reputations in bargaining: An experimental study. *American Economic Review* 73:362–72.

Rubinstein, Ariel. 1982. Perfect equilibrium in a bargaining model. *Econometrica* 50:97–109.

Salop, Steven C., and Lawrence J. White. 1988. Private antitrust litigation: An introduction and framework. In *Private antitrust litigation*, L. J. White, editor, Cambridge: MIT Press.

Schelling, Thomas C. 1957. Bargaining, communication, and limited war. *Journal of Conflict Resolution* 1:19–36.

Schelling, Thomas C. 1960. *The strategy of conflict*. Cambridge: Harvard University Press.

Schotter, Andrew. 1990. Bad and good news about the sealed-bid mechanism: Some experimental results. *American Economic Review, Papers and Proceedings* 80:220–26.

Schotter, Andrew, Blaine Snyder, and Wei Zheng. 1992. Bargaining through agents: An experimental Study. Mimeo. New York University. October.

Selten, Reinhard. 1967. Die strategiemethode zur erforschung des eigeschrankt rationalen verhaltens im rahmen eines oligopolexperiments. In *Beitrage zur experimentellen wirtschaftsforschung*, H. Sauermann, editor, Tubingen: J.C.B. Mohr. 136–68.

Selten, Reinhard. 1975. Bargaining under incomplete information: A numerical example. In *Dynamische wirtschaftsanalyse*, O. Becker and R. Richter, editors, Tubingen: J.C.B. Mohr. 203–32.

Shogren, Jason F. 1989. Fairness in bargaining requires a context: An experimental examination of loyalty. *Economics Letters* 31:319–23.

Shogren, Jason F. 1992. An experiment on Coasian bargaining over ex ante lotteries and ex post rewards. *Journal of Economic Behavior and Organization*, Special Issue on Experimental Economics, 17:153–69.

Siegel, Sidney, and Lawrence E. Fouraker. 1960. *Bargaining and group decision making: Experiments in bilateral monopoly*. New York: McGraw-Hill.

Sopher, Barry. 1990. Bargaining and the joint-cost theory of strikes: An experimental study. *Journal of Labour Economics* 8:48–74.

Spiegel, Matthew, Janet Currie, Hugo Sonnenschein, and Arunava Sen. 1990. First-mover advantage and the division of surplus in two-person, alternating-offer games: Results from bargaining experiments. Mimeo.

Spier, Kathryn. 1992. The dynamics of pretrial negotiations. *Review of Economic Studies* 59: 93–108.

Stahl, Ingolf. 1972. *Bargaining theory*. Stockholm: Economic Research Institute.

Straub, Paul G., and J. Keith Murnighan. 1992. An experimental investigation of ultimatum games: Information, fairness, expectations, and lowest acceptable offers. Mimeo. December 28, University of Illinois.

Thaler, Richard H. 1988. The ultimatum game. *Journal of Economic Perspectives* 2: 195–206.

Thompson, Leigh. 1990. Negotiation behavior and outcomes: Empirical evidence and theoretical Issues. *Psychological Bulletin* 3:515–32.

Van Huyck, John, Raymond Battalio, Sondip Mathur, Andreas Ortmann, and Patsy Van Huyck. 1992. On the origin of convention: Evidence from symmetric bargaining games. Working paper. Department of Economics. Texas A&M University.

Weg, Eythan, Amnon Rapoport, and Dan S. Felsenthal. 1990. Two-person bargaining behavior in fixed discounting factors games with infinite horizon. *Games and Economic Behavior* 2:76–95.

Weg, Eythan, and Vernon Smith. 1993. On the failure to induce meager offers in ultimatum games. *Journal of Economic Psychology* 14:17–32.

Weg, Eythan, and Rami Zwick. 1991. On the robustness of perfect equilibrium in fixed cost sequential bargaining under an isomorphic transformation. *Economics Letters* 36:21–4.

Weg, Eythan, Rami Zwick, and Amnon Rapoport. Forthcoming. Bargaining in uncertain environments: A systematic distortion of perfect equilibrium demands. *Games and economic behavior*.

White, Michelle. 1989. Dispute prevention and dispute resolution in medical malpractice. Mimeo. University of Michigan.

Yaari, Menahem E., and Maya Bar-Hillel. 1984. On dividing justly. *Social Choice and Welfare* 1:1–24.

Zwick, Rami, Amnon Rapoport, and John C. Howard. 1992. Two person sequential bargaining behavior with exogenous breakdown. *Theory and Decision* 32:241–68.

5

Industrial Organization: A Survey of Laboratory Research

Charles A. Holt

I. Overview

Despite the contrast between the relative simplicity of the laboratory and the complexity of most naturally occurring markets, there is a well-established tradition of experimental research in the field of industrial organization (IO).* Indeed, the first market experiment resulted from Edward Chamberlin's conjectures about the imperfect nature of competition. This chapter will survey the extensive experimental literature that is motivated by IO issues, beginning in sections II and III with the story of how economists initially became interested in experimentation, and with a discussion of the potential usefulness of laboratory techniques in this area. Section IV contains a review of some procedural issues that arise in the conduct of market experiments. One pervasive theme is the importance of the rules and informational conditions of the laboratory *market institution*. Section V contains descriptions and comparisons of the trading institutions that are used in the study of industrial organization issues. The four substantive sections that follow are organized around traditional topics: monopoly regulation and potential entry, concentration and market power, conditions that facilitate cooperation, and product differentiation.

One set of issues to be considered is the extent to which a monopoly seller can exercise market power in the laboratory, if this power can be protected with predatory pricing, and whether contestability and decentralized regulatory mechanisms mitigate this power. A second set of issues is based on the usefulness of simple concentration measures: do they predict supra-competitive pricing, or it is necessary to consider more subtle notions of market power? Contracts, trading rules, communications and other factors that do and do not seem to facilitate cooperation in laboratory situations are discussed. In experiments with interrelated markets, the discussion covers failures (due to asymmetric quality information) and successes where competition generates efficient coordination across markets).