Journal of Economic Behavior and Organization 2 (1981) 153–177. North-Holland

SOCIOLOGICAL VERSUS STRATEGIC FACTORS IN BARGAINING*

Alvin E. ROTH, Michael W.K. MALOUF and J. Keith MURNIGHAN

University of Illinois, Urbana, IL 61801, USA

Received November 1980

Most game-theoretic models of strategic interaction, indeed most economic models of any sort, specify potential outcomes entirely in terms of the preferences of the agents, as captured in their (possibly cardinal) utility functions. The underlying assumption is that the outcome of such interactions is determined entirely by these preferences, together with the strategic possibilities available to the agents. The purpose of this paper is to challenge the adequacy of this assumption in general, by investigating it in the specific context of two-person bargaining. In particular, we consider whether certain experimental results reported earlier can be accounted for strictly in terms of players' preferences and strategic possibilities, and we report a new experimental study designed to answer this question. The results strongly support the conclusion that sociological factors, unrelated to what we normally consider to be the 'economic' parameters of a game, can decisively influence the outcome of bargaining, in a systematic manner.

1. Introduction

Most game-theoretic models of strategic interaction, and indeed, most economic models of any sort, specify potential outcomes entirely in terms of the preferences of the agents, as captured in their (possibly cardinal) utility functions. The underlying assumption is that the outcome of such interactions is determined entirely by these preferences, together with the strategic possibilities available to the agents. The purpose of this paper will be to challenge the adequacy of this assumption in general, by investigating it in the specific context of two-person bargaining. In particular, we will consider whether the experimental results reported in Roth and Malouf (1979) can be accounted for strictly in terms of players' preferences and strategic possibilities, and we will report a new experimental study designed to answer this question.

In *single-person* decision problems an individual's choice is completely modelled by his utility function, the feasible actions, and their consequences. In economic models of *perfect competition* agents are assumed to act as price-takers. Once prices have been determined by the market, the problem

*This work has been supported by National Science Foundation grants no. SOC 78-09928 and no. SES 79-15356 to the University of Illinois. It is also a great pleasure to acknowledge invaluable advice from Ronald M. Harstad of the University of Illinois, and from Charles R. Plott of the California Institute of Technology, and stimulating conversations with the other participants at the 1979 Conference on Experimental Economics, and with Menachem Yaari.

0167-2681/81/0000-0000/$02.50

decomposes into many single-person decision problems. So, in a model of perfect competition, once the feasible actions and their consequences have been specified, the only feature of the model in which there might be room for information other than the preferences of the players is in specifying precisely how prices are determined. (Of course, this might be of considerable importance; e.g., in markets having multiple competitive equilibria.)

As competition becomes less perfect, rational agents presumably act less like price-takers, interaction among agents becomes less like a collection of single-person decision problems, and the assumption that the decision process can be completely modelled without other information than agent preferences becomes less obvious. That is, in game situations, there is no *a priori* reason to believe that structural information about the game together with information about individual player's preferences will be an adequate basis on which to construct a model of rational interaction among the players. Nevertheless, while there have been some attempts to consider the potential effect of 'sociological' information about the players [e.g., information concerning the manner in which players choose among different potentially stable modes of behavior, such as different von Neumann-Morgenstern (1944) solutions], these attempts have all been carried out in the context of models which are defined entirely in terms of the players' preferences and strategic possibilities.[1]

In what follows, we will demonstrate that, at least for the case of two-person bargaining, models defined exclusively in terms of this sort of information are inadequate as descriptive models. (We also argue that, because of the interactive nature of bargaining, such models are consequently unsatisfactory prescriptive models for rational agents.) Our argument will thus be directed not only at identifying shortcomings of specific theories of bargaining which appear in the literature, but at showing that *any* descriptive theory whose predictions are determined exclusively by the players' preferences and strategic possibilities must share these shortcomings. In order to support such a strong conclusion, data from a controlled experiment, conducted under laboratory conditions, will be presented.

This paper consists of six sections. In section 2 we will consider cooperative models of bargaining, and discuss how such models may be tested experimentally. Section 3 will review the experiment of Roth and Malouf (1979), and discuss the manner in which it suggests that these cooperative models fail. Section 4 introduces the class of non-cooperative (strategic) models of bargaining, and discusses how such models are not ruled out by the existing data. Section 5 reports a new experiment, designed

[1]Throughout this paper we will assume we are dealing with situations in which all agents exhibit sufficiently regular individual choice behavior so that their revealed preferences can be modelled by numerical utility functions over the set of possible outcomes. When uncertainty is involved, these are assumed to be expected utility functions. The significance of these assumptions for the empirical parts of the paper will be discussed.

to determine whether the phenomena originally reported by Roth and Malouf (1979) are in fact caused by strategic considerations, and can therefore be modelled by structural information about the game together with information about the players' preferences. It is perhaps appropriate to mention at this point that, before conducting this experiment, we expected that the data would be consistent with a purely strategic theory of bargaining. However, the results clearly indicate that non-strategic, 'sociological' factors play a decisive role.[2] This conclusion and its implications are discussed in section 6.

2. Cooperative models of bargaining

Cooperative games are customarily modelled by specifying the set of feasible utility payoffs attainable by each non-empty subset of players for its members. Following Nash (1950), the two-player bargaining games considered here are modelled by a pair (S, d), where d is a point in the plane, and S is a compact convex subset of the plane which contains d and at least one point x such that $x > d$. The interpretation is that S is the set of feasible expected utility payoffs to the players, any one of which can be achieved if it is agreed to by both players. If no such agreement is reached, then the disagreement point d is the result.

Nash proposed that bargaining between rational players be modelled by means of a function called a *solution*, which selects a feasible outcome for every bargaining game. That is, if we denote the class of all two-player bargaining games by B, a solution is a function $f: B \to R^2$ such that $f(S, d)$ is an element of S. Thus a solution is a model of bargaining which depends only on the information about the underlying game which is contained in the model (S, d).

In order to insure that such a theory of bargaining would depend only on the information about preferences contained in a player's utility function, Nash further proposed that a solution should possess the following property:

Property 1. Independence of equivalent utility representations: if (S, d) and $(\hat{S}, \hat{d})$ are bargaining games such that

$$\hat{S} = \{(a_1x_1 + b_1,\ a_2x_2 + b_2) \mid (x_1, x_2) \in S\},$$

[2]This conclusion may be (at least loosely) comparable to the discovery by cognitive psychologists and researchers in artificial intelligence cf. Bransford and McCarrell (1975) or Minsky (1975) that context plays a crucial role in understanding natural language, and that the meaning of a sentence depends on more than its linguistic structure. (For example, the sentence 'Please press this suit' obviously has a different meaning when you are speaking to a tailor than when you are speaking to a lawyer.) Similarly, we will demonstrate here that the outcome of a game depends on more than its economic structure.

and

$$\hat{d} = (a_1 d_1 + b_1,\ a_2 d_2 + b_2),$$

where a_1, a_2, b_1 and b_2 are numbers such that a_1 and $a_2 > 0$, then

$$f(\hat{S}, \hat{d}) = (a_1 f_1(S, d) + b_1, a_2 f_2(S, d) + b_2).$$

In order to understand the significance of this property, we need to consider the set of underlying alternatives over which the bargaining is conducted. Suppose that two individuals are bargaining over some set of alternatives A, containing some pre-specified disagreement outcome[3] a^*. Then if these individuals have utility functions u_1 and u_2 over the set A, the resulting bargaining game (S, d) is given by

$$S = \{(u_1(a), u_2(a)) \mid a \in A\}, \qquad d = (u_1(a^*), u_2(a^*)). \tag{1}$$

Recall that an individual i's utility function u_i is a real-valued function defined on the set of alternatives A. It is a model of his choice behavior, in the sense that $u_i(a) > u_i(b)$ for two alternatives a and b if and only if he prefers a to b; i.e., if and only if he would choose alternative a when faced with the choice between a and b. Von Neumann and Morgenstern (1944) demonstrated conditions on an individual's preferences which are sufficient so that his choice behavior over risky alternatives is the same as if he were maximizing the expected value of his utility function. Such a utility function is uniquely defined only up to an interval scale, which is to say that the origin (zero point) and unit of the utility function are arbitrary. Thus if u_i is an expected utility function representing individual i's preferences, then another utility function v_i represents the same preferences if and only if $v_i = a_i u_i + b_i$, where a_i is a positive number.

So Property 1 states that if a game $(\hat{S}, \hat{d})$ is derived from (S, d) by transforming the utility functions of the players to equivalent representations of their preferences, then the same transformations applied to the outcome of the game (S, d) should yield the outcome selected in $(\hat{S}, \hat{d})$. That is, if $(\hat{S}, \hat{d})$ is given by

$$\hat{S} = \{(v_1(a), v_2(a)) \mid a \in A\}, \qquad \hat{d} = (v_1(a^*), v_2(a^*)), \tag{1'}$$

where $v_i = a_i u_i + b_i$ for $i = 1, 2$, and if a solution f yields $f(S, d) = (u_1(b), u_2(b))$, then Property 1 requires that $f(\hat{S}, \hat{d}) = (v_1(b), v_2(b))$, i.e., that the payoff

[3] That is, the rules of the game are that any alternative a in the set A will be the outcome of the game if both players agree on it, otherwise a^* will be the outcome. Thus the rules of the game give both players a veto over any outcome other than a^*.

predicted by the solution f should correspond to the same alternative b in both games. Thus Property 1 states that a solution should depend only on those properties of the utility functions which represent the preferences of the players, and not on the arbitrary features of the utility functions.

Nash (1950) also proposed three additional properties which, together with Property 1, characterize a unique solution (which is often referred to as Nash's solution). Other additional properties have subsequently been used to characterize other specific solutions, and the subject has inspired a considerable amount of research [see Roth (1979) for a survey]. However, we will be concerned here with the general class of solutions which are defined on the class B of bargaining games, and which possess Property 1.

A theory of bargaining embodied in such a solution makes two distinct (but related) predictions. First, since the solution depends only on the utility payoffs available to the players, it yields the same prediction for a given game (S,d) no matter how that game arises; e.g., whether the game arises from bargaining over a set of alternatives A by individuals with utility functions u_1 and u_2, or from bargaining over an entirely different set of alternatives A' by individuals with appropriate utility functions. That is, such a solution predicts that bargaining situations which have the same representation (S,d) in utility space will result in the same utility payoffs to the players. Second, if $(\hat{S},\hat{d})$ is related to (S,d) as in the statement of Property 1, then the solution predicts that the utility payoffs resulting from the two games will be related as in Property 1, regardless of whether $(\hat{S},\hat{d})$ differs from (S,d) only by a purely formal transformation [as in eqs. (1) and (1')], or whether the two games have substantive differences, as when they arise from bargaining over different sets of alternatives.

Thus, to the extent that appropriate games can be constructed, experiments can be conducted to test the predictive value of solutions which are defined on the class B, and which are independent of equivalent utility representations. The next section briefly reviews such an experiment, originally reported in Roth and Malouf (1979).

3. The experiment of Roth and Malouf

Since the class of theories considered in the previous section are defined in terms of the players' utilities, experimental tests of such theories must be constructed in a way which permits the players' utilities to be determined. A novel feature of the experiment reported in Roth and Malouf (1979) is that it employed games constructed in a way which permits the utility of the players for each outcome to be determined directly. In order to explain how this was accomplished, it will be helpful to recall precisely what information is contained in an expected utility function.

Consider the case in which the set A of alternatives contains elements a

and c such that a is strictly preferred to c and for any alternative b in A, the player likes a at least as well as b, and b at least as well as c. Then if u is a utility function representing this individual's preferences over the set of alternatives A, it must have the property that $u(a) \geq u(b) \geq u(c)$. Since u is defined only up to an interval scale, we may arbitrarily choose its unit and zero point, and in particular we may take $u(a) = 1$ and $u(c) = 0$. The problem of determining $u(b)$ then becomes the problem of finding the appropriate value between 0 and 1 so that all those lotteries over alternatives that the individual prefers to b have a higher expected utility, and all those lotteries to which b is preferred have a lower expected utility. If we denote by $L(p) = [pa;(1-p)c]$ the lottery that with probability p yields alternative a and with probability $(1-p)$ yields alternative c, then the utility of participating in the lottery $L(p)$ is its expected utility, $pu(a)+(1-p)u(c)=p$. If p is the probability such that the individual is indifferent between b and $L(p)$, then their utilities must be equal, and so, $u(b)=p$. Thus when we say that the utility of alternative b to a given individual is known, we mean that the probability p is known such that the individual is indifferent between having alternative b for certain or having the risky alternative $L(p)$.

3.1. Binary lottery games

Since knowing an individual's expected utility for a given agreement is equivalent to knowing what lottery he or she thinks is as desirable as that agreement, then in a bargaining game in which the feasible agreements are the appropriate kind of lotteries, knowing the utilities of the players at a given agreement is equivalent to simply knowing the lottery they have agreed on. In each game of this experiment, therefore, players bargained over the *probability* that they would receive a certain monetary prize, possibly a different prize for each player.

Specifically, they bargained over how to distribute 'lottery tickets' that would determine the probability that each player would win his or her personal lottery (i.e., a player who received 40% of the lottery tickets would have a 40% chance of winning his monetary prize and a 60% chance of winning nothing). In the event that no agreement was reached in the allotted time, each player received nothing. In other words, a player received his prize only if an agreement was reached on splitting the lottery tickets and if he won the ensuing lottery. Otherwise he received nothing. We will refer to games of this type, in which each player has only two possible monetary payoffs, as *binary lottery games*.

To interpret the set of feasible outcomes of a binary lottery game in terms of each player's utility function for money, recall that if we consider each player's utility function to be normalized so that the utility for receiving his prize is 1, and the utility for receiving nothing is 0, then the player's utility for any lottery between these two alternatives is the probability of winning the lottery.[4] That is, an agreement which gives a player p percent of the lottery tickets gives him a utility of p.[5]

Note that a change in the prizes is therefore equivalent to a change in the scale of the player's utility functions. This makes it possible to use binary lottery games to experimentally investigate the circumstances under which the bargaining process can indeed be described by a solution which is independent of equivalent utility representations.

3.2. Design of the experiment

Each player played four games, in random order, against different opponents. Each player played all four games under one of two information conditions: *full information*, or *partial information*. In the full information condition, each player was informed of the value of his own potential prize and of his opponent's potential prize.[6] In the partial information condition, each player was informed only of the value of his own prize.[7]

Players were seated at isolated computer terminals, and were allowed to communicate freely by teletype, but they were unaware of the identity of their opponents. (The only limitations on free communication were that players were prevented from identifying themselves, or from conveying information about the monetary value of their prizes in the partial

[4] Note that the assumption that a player's preferences are sufficiently regular to be represented by an expected utility function is equivalent, over this simple set of feasible outcomes, to the assumption that the player prefers a higher probability of winning to a lower probability of winning.

[5] Note that we are considering the feasible set of utility payoffs to be defined in terms of the utility function of each player for the lottery which he receives, independently of the bargaining which has taken place to achieve this lottery, and even independently of the lottery which his opponent receives. In doing so, we are taking the point of view that, while the progress of the negotiations may influence the utilities of the bargainers for the agreement eventually reached, the description of any effect which this has on the agreement reached belongs in the model of the bargaining *process*, rather than in the model of the bargaining *situation*. Considerable confusion in the literature has resulted from attempts to interpret bargaining models in terms of the players' utilities for outcomes *after* the bargaining has ended, since no bargaining model can be falsified by experimental evidence if, after an outcome has been chosen, the utilities of the players can be interpreted as having changed in whatever way is necessary to be consistent with the model. In order to have predictive value, bargaining theories must be stated in terms of parameters which can be measured independently of the phenomena which the theories are designed to predict. [A discussion of the case when a player's preferences cannot be interpreted as being independent of his opponent's lottery is found in section 3 of Roth and Rothblum (1980).]

[6] Specifically, the prizes were common knowledge in this condition [cf. Roth and Murnighan (1980)].

[7] Note that in both the full information and in the partial information conditions, the resulting games meet the usual assumption that the game is one of *complete* information: i.e., both players have sufficient information to determine one another's expected utility for every outcome. Of course in the full information condition, the players have additional information, since they know one another's monetary payoffs as well.

information condition.) The bargaining process consisted of the exchange of messages and of (numerical) proposals, and terminated in an agreement when a proposal was accepted or in disagreement if no proposal had been accepted after 12 minutes.[8]

In game 1, no restriction was placed on the percentage of lottery tickets which each player could receive, and both players had the same prize of \$1.00. Game 2 was played with the same prizes as game 1, but one of the players (player 2) was restricted to receive no more than 60% of the lottery tickets. Game 3 was played with the same rules as game 1, but with different prizes for the two players: \$1.25 for player 1, and \$3.75 for player 2. Game 4 was played under the same rules as game 2, with the same prizes as game 3 (see table 1).

Table 1
Prizes and feasible distributions for games 1–4.

Game	Prize for player 1	Prize for player 2	Maximum % allowed for player 1	Maximum % allowed for player 2
1	\$1	\$1	100	100
2	\$1	\$1	100	60
3	\$1.25	\$3.75	100	100
4	\$1.25	\$3.75	100	60

Note that game 1 is related to game 3, and game 2 is related to game 4 by a change in the prizes, and hence by a scale change as in Property 1. So if the bargaining process obeys Property 1, we should observe the same outcomes in these pairs of games. And if the bargaining process depends only on the set of feasible utility payoffs, then we should observe the same outcome for each game under both information conditions, since the set of feasible lotteries (and hence utilities) faced by each player is unaffected by the information conditions.

Denote the difference between the probabilities p_1 and p_2 received by the two players by $D=p_1-p_2$. If (as Nash's solution, for instance, predicts) we were to observe the players divide the lottery tickets equally in these games, so that $p_1=p_2=50\%$, then we would have $D=0$ in these games. On the other hand, if we were to observe the players reach agreements which equalized their expected monetary payoffs, then we would observe $D=0$ for games 1 and 2, and $D=50$ for games 3 and 4 (corresponding to $p_1=75$, $p_2=25$).

[8]The detailed procedures by which these conditions were implemented in practice will be discussed in section 5, since they are essentially the same as those used in the experiment discussed there.

In fact, the observed outcomes in the partial information condition were extremely close to an equal division of the lottery tickets, while the observed outcomes in the full information condition showed a pronounced shift in the direction of equal expected monetary payoffs. (That is, in the full information condition, in the games in which the two players had different prizes, the observed agreements tended to give a higher probability of winning to the player with the smaller prize.) The means and standard deviations of the observed values of D are summarized in table 2.

Table 2
Means and standard deviations for $D = p_1 - p_2$.

Statistic	Game 1	2	3	4
Full information (11 pairs)				
M	0.00	1.91	34.60	21.64
SD	0.00	12.17	19.28	22.48
Partial information (8 pairs)				
M	0.00	−1.32	−2.50	2.50
SD	0.00	8.33	4.63	4.11

Statistical analysis confirms that, in the partial information condition D is not significantly different between games 1 and 3, or between games 2 and 4, while in the full information condition these differences are significant. Also, the outcomes for each of games 3 and 4 are significantly different in the two information conditions.[9]

The observed outcomes in the full information condition thus do not conform to the predictions of Property 1, and their difference from the outcomes observed in the partial information condition cannot be accounted for in terms of the set of feasible utility payoffs. That is, the shift towards equal expected monetary payoffs observed in the full information condition of this experiment and confirmed in subsequent experiments [cf. Roth and Malouf (1980), Malouf (1980), Roth and Murnighan (1980)] cannot be integrated with the results of the partial information condition by a model which depends only on the set of feasible utility payoffs. In the next section, we will explore the possibility that the observed differences between the two information conditions can be accounted for by the fact that the set of feasible negotiation strategies available to the players depends on the nature of the information which they possess.

[9]A more detailed statistical analysis can be found in Roth and Malouf (1979), which also presents the unaggregated data from this experiment.

4. Strategic use of information

Models of the kind we have been considering, which define a game in terms of its feasible utility payoffs, do not include a description of the strategic choices which the players must make to achieve these payoffs. In the experiment reported in the previous section, these strategic choices involve the exchange of both proposals and messages.

Examination of the transcripts of the bargaining encounters in the experiment reveal striking differences in the content of messages between encounters in the two information conditions. While the exchange of messages was vigorous in both information conditions, only in the full information condition could messages contain comparisons of the players' prizes. Not too surprisingly, in the full information condition, players who had smaller prizes than their opponents often persistently demanded more than half of the lottery tickets. Since the results of the bargaining in the full information condition were observed to satisfy this demand while the results in the partial information condition did not, it is certainly plausible to speculate that the difference between the two conditions can be accounted for entirely in terms of the different kinds of strategic choices available to the players in the two conditions. In order to state this hypothesis precisely, we will need to consider a general model of the game which incorporates information about the players' strategies as well as their preferences for the feasible outcomes.

Consider a strategic model of a binary lottery game involving players with utility functions u_1 and u_2 defined on some set A of alternatives (lotteries), one of which will be the outcome of the game. Players 1 and 2 have strategy sets S_1 and S_2, respectively, and associated with every pair of strategy choices $s=(s_1, s_2)$ is an outcome $O(s)=a$ contained in A. That is, the outcome of the game is determined by the outcome function O from the set $S_1 \times S_2$ to A: i.e., the rules of the game are that each player i chooses a strategy s_i from his strategy set S_i, and the combined choices of the players determine the outcome a in A, which results in the utility payoff vector $(u_1(a), u_2(a))$. We will refer to such a model as the *expanded*[10] *strategic form* of a game; i.e., a game G in expanded strategic form consists of the elements $G=(S_1, S_2, O, A, u_1, u_2)$. We will adopt the convention that in a binary lottery game in expanded strategic form, the utility functions are normalized

[10]The (unexpanded) strategic form of a game consists of the strategy sets S_1 and S_2 and *payoff functions* Π_1 and Π_2 such that, for any strategy pair $s=(s_1, s_2)$, $\Pi_i(s)=u_i(O(s))$ for $i=1, 2$. That is, the (unexpanded) strategic form of a game represents the actual outcomes of the game only in terms of the utilities of the players. Since we will be interested in distinguishing the strategy choices of the players from the set of resulting outcomes (over which the players' utility functions are defined), we will use the expanded strategic form of the game, rather than collapsing the functions u_i and O into the single 'payoff function' Π_i. Although our concern here is with binary lottery games, we phrase the discussion in terms of this general model in order to make clear how the issues discussed here apply to arbitrary games.

so that each player's utility for a lottery is equal to the probability it gives him of winning his prize. (This means that when we compare different games whose sets of alternatives involve lotteries over different prizes, we will be comparing games defined in terms of different utility functions.)

Modelled in this way, all information about the players' preferences is contained in the utility functions u_1 and u_2 and the set A of alternatives on which these preferences are defined, while the strategy sets S_1, S_2, and the outcome function O contain the 'structural' information about the game, which tells us how the players actions are translated into outcomes. That is, u_1, u_2, and A model the players' objectives, while the strategy sets S_1 and S_2 together with the outcome function O model the mechanism which the game provides for resolving these (different) objectives. If we have a theory of games which predicts the outcome of a game in terms of the players' preferences and the structure of the game mechanism, then two games which have the same relationship between strategy choices, outcomes, and preferences will yield the same prediction. Formally, we can express this as follows.

Let $G=(S_1, S_2, O, A, u_1, u_2)$ and $\hat{G}=(\hat{S}_1, \hat{S}_2, \hat{O}, \hat{A}, \hat{u}_1, \hat{u}_2)$ be two games in expanded strategic form. Then G and $\hat{G}$ are defined to be *strategically equivalent*[11] if there exist transformations T_1 and T_2 such that, for $i=1$, 2, $T_i: S_i \to \hat{S}_i$ is one-to-one and onto, and for every strategy pair $s=(s_1, s_2)$ in $S_1 \times S_2, u_i(O(s)) = \hat{u}_i(\hat{O}(\hat{s}))$ for $i=1$, 2, where $\hat{s}=(T_1(s_1), T_2(s_2))$ is the image of s under $T=(T_1, T_2)$. Thus the transformations T_1 and T_2 can be regarded essentially as relabellings of the strategy sets S_1 and S_2, and the outcome function $\hat{O}$ acts on the relabelled strategy sets $\hat{S}_1$ and $\hat{S}_2$ in the same way that O acts on S_1 and S_2.

A model of the bargaining process which depends only on the preferences of the players and the structure of their strategic possibilities would be one which made the same predictions for any two strategically equivalent games G and $\hat{G}$. For instance, if we let E denote the set of all bargaining games in expanded strategic form, then a parallel to a solution of the kind considered in section 2 would be a function g: $E \to R^2$. That is, a solution g for games in expanded strategic form selects a feasible utility payoff $g(G)$ for every bargaining game G in the class E. Such a solution g depends only on the preferences of the players and the structure of the game if it obeys the following property:

Property 2. Invariance with respect to strategic equivalence: If G and $\hat{G}$ are strategically equivalent games, then $g(G)=g(\hat{G})$.

Note that, although for simplicity we are considering solutions g which select a payoff corresponding to a unique outcome, the extension of Property

[11] The notion of strategic equivalence takes a somewhat different form when applied to other kinds of game models

2 to the case of solution concepts which select sets of payoffs corresponding to more than one outcome is straightforward. For example, the set of all Nash equilibrium payoffs[12] of the game is invariant with respect to strategic equivalence since, if $s=(s_1,s_2)$ is any equilibrium strategy pair yielding the payoff vector $(u_1(O(s)), u_2(O(s)))$ in the game G, then $\hat{s}=(T_1(s), T_2(s))$ is an equilibrium strategy pair yielding the same payoffs in a strategically equivalent game $\hat{G}$. So, for instance, a theory of games which merely specified that the outcome of a game would be an equilibrium outcome would be invariant with respect to strategic equivalence.

Of course, a theory of games which selects a *unique* equilibrium payoff of each game, for instance, may or may not be invariant with respect to strategic equivalence, depending on whether it depends only on the structural information of the game, or on information about strategies and outcomes which is not preserved in going from one game to a strategically equivalent game.[13] The experiment reported in the next section is designed to determine whether the bargaining process observed in the experiment of Roth and Malouf (1979) is invariant with respect to strategic equivalence.

Since we are interested in bargaining situations which allow extensive communication between the participants, we will not attempt to model their strategy sets explicitly.[14] Instead, we will argue that certain games are strategically equivalent by demonstrating the isomorphism between them. By observing strategically equivalent games, we will seek to determine whether the differences observed between the full and partial information conditions are due to the structural properties of the game, or whether they are due to other factors.

5. A new experiment

The experiment reported in this section involves binary lottery games (cf. section 3) whose prizes are stated in terms of an intermediate commodity. Each bargainer was told that the prizes would be expressed in 'chips' having monetary value, and each player played four games under one of three

[12]A pair of strategy choices $s=(s_1,s_2)$ in a game G is a Nash (1951) equilibrium if s_1 is player 1's best response against player 2's choice of s_2, and s_2 is player 2's best response to player 1's choice of s_1 [i.e., if $u_1(O(s))\geq u_1(O(t_1,s_2))$ for all t_1 in S_1, and if $u_2(O(s))\geq u_2(O(s_1,t_2))$ for all t_2 in S_2]. If s is an equilibrium pair of strategies, the resulting outcome $O(s)$ is an equilibrium outcome, and $(u_1(O(s)), u_2(O(s)))$ is an equilibrium payoff.

[13]That is, such a theory might not select the 'same' equilibrium in two strategically equivalent games if it depends on information about the strategy sets which isn't preserved by strategic equivalence.

[14]These strategy sets are infinite, involving as they do the choice not only of the content of individual messages, but also their timing. If we wished to be able to compute equilibria of a bargaining game in strategic form, we would need to confine our attention to games in which the strategy sets are much more restricted. Some interesting results obtained using this alternative approach are contained in recent papers by Rubinstein (1980) and Binmore (1980).

information conditions: *high information*, *intermediate information*, or *low information*. In each of the three conditions, each player knew the number of chips in his potential prize and their monetary value, but the information each player was given about his opponent's prize varied with the information condition. In the high information condition, each player was informed of both the number of chips in his opponent's potential prize and their monetary value. In the intermediate information condition, each player was informed of the number of chips in his opponent's potential prize, but not of their monetary value. In the low information condition, neither player was informed of either the number of chips in his opponent's potential prize, or of their value. In the latter two conditions, players were prevented from communicating the missing information about the prizes (see the detailed description of methods below).

The four games (which were played in random order) are summarized in table 3. Note that the games are counterbalanced in the sense that, in two of the games, the player with the higher number of chips also has a higher value per chip (and hence a higher value prize), while in the other two games, the player with the higher number of chips has a lower value per chip and a lower value prize.

Table 3
Games.

	Player 1			Player 2		
	Number of chips	Value per chip	Value of prize	Number of chips	Value per chip	Value of prize
Game 1	60	\$0.05	\$3.00	20	\$0.45	\$9.00
Game 2	80	\$0.03	\$2.40	240	\$0.04	\$9.60
Game 3	100	\$0.09	\$9.00	300	\$0.01	\$3.00
Game 4	150	\$0.08	\$12.00	50	\$0.06	\$3.00

The experiment has been designed to take advantage of two kinds of strategic equivalence relations. First, note that binary lottery games whose prizes are expressed in both chips and money, played in the low information condition of this experiment, are strategically equivalent to binary lottery games with the same monetary prizes whose prizes are expressed in money alone, played in the partial information condition of the previous experiments. This follows from the fact that, under the rules of the low and partial information conditions, any message which is legal for one kind of game would be a legal message for the other, and so the strategy sets are the same for both kinds of games, as are the utility functions and the underlying set of alternatives.

Second, chip games[15] played under the intermediate information condition of this experiment are strategically equivalent to money games played under the full information condition of the previous experiments, so long as the monetary values of the two prizes in each money game are in the same proportion as the numbers of chips in the prizes in the corresponding chip game. This follows from the fact that any legal message in one kind of game can be transformed into a legal message in the other kind of game by substituting references to chips for references to money (or vice versa) in any message concerning the value of the prizes. And since the relative value (in money or chips) of the prizes is the same in corresponding games, the outcome function preserves the necessary equivalence; e.g., an agreement by the players that they should receive equal expected values (in money or chips) will yield the same probabilities (and hence the same utilities) to the players in both kinds of games.

Having established these equivalence relations, we can now proceed to formulate two distinct sets of predictions concerning the results of this new experiment, depending on which of two competing hypotheses we believe account for the experimental results summarized in section 3. These hypotheses and the corresponding predictions can be stated as follows.

5. . The strategic hypothesis

As discussed in section 4, this hypothesis states that the shift towards equal expected monetary payoffs in the full information condition as compared to the partial information condition, observed in the previous experiments, is due to the different strategy sets available to the players in the two conditions. Consequently, this hypothesis predicts that a similar shift will be observed in the *intermediate* information condition as compared to the *low* information condition of this experiment, since the chip games played under these two conditions in this experiment are strategically equivalent, respectively, to games played under the full and partial information conditions of the previous experiments. Specifically, the prediction is that games played in the low information condition of this experiment will lead to agreements in which the two players receive approximately equal probabilities of winning their prizes, while games played in the intermediate information condition will lead to agreements in which the player with the smaller number of chips will receive a significantly higher probability of winning his prize than will his opponent. Thus the prediction is that the observed values of D will not deviate significantly from zero in any of the games in the low information condition, while in the intermediate

[15]For the sake of brevity, binary lottery games will sometimes be referred to as *chip games* or *money games*, depending on whether the prizes are expressed in chips as well as money, or in money alone.

information condition, D should decrease significantly in games 1 and 4, and increase significantly in games 2 and 3. (Notice that the strategic hypothesis makes no prediction about the high information condition of this experiment, since the games in this condition are not strategically equivalent to any games played in the previous experiments. However, if the strategic hypothesis is correct, the observations in this condition will illuminate the interaction between negotiation strategies concerning chips and those concerning money because of the way the games are counterbalanced.)

5.2. The sociological hypothesis

This hypothesis seeks to account for the phenomena reported in section 3 in terms of social conventions which exist among the bargainers. The underlying idea is that in conflict situations involving a wide range of rational potential agreements, social conventions may serve to make some arguments and demands more credible than others. Thus this hypothesis views the low variance observed in the partial information condition of the previous experiments as evidence that the agreement at which both players have an equal chance of winning their prizes is supported by a social norm which inclines both players to believe that their opponent may be unwilling to accept less. The shift towards equal expected monetary payoffs which was observed in the full information condition is viewed as evidence that when information about the monetary value of the prizes is available, the agreement giving the players equal expected payoffs is also supported by such a convention, and so the bargaining focuses on resolving the difference between two credible positions.[16]

By 'social conventions', we mean customs or beliefs which are commonly shared by the members of a particular society. In order to be commonly shared, such conventions must necessarily be concerned with familiar quantities. By stating the prizes in terms of an unfamiliar artificial commodity ('chips') which conveys no information about more familiar quantities such as the value of a given prize or a player's probability of winning it, this new experiment seeks to introduce a quantity about which no social conventions apply. The sociological hypothesis predicts, therefore, that information about the number of chips in each prize will not affect the bargaining. Specifically, this hypothesis predicts that the low and high information conditions of this experiment will replicate the partial and full information conditions of the previous experiments, respectively, and that the intermediate information condition will not differ significantly from the low information condition. Thus the prediction is that D will not differ significantly from zero in any of the games in the low or intermediate

[16]Informal analysis of the transcripts from the bargaining observed in previous experiments lends support to this hypothesis.

information conditions, but that in the high information condition D will significantly increase in games 1 and 2, and decrease in games 3 and 4.

The most pointed difference between the predictions of the two hypotheses is therefore in their predictions for the intermediate information condition of the current experiment. The strategic hypothesis predicts that the intermediate information condition of this experiment will resemble the full information condition of the previous experiments, and exhibit a pronounced shift away from agreements at which the players receive equal probabilities. The sociological hypothesis predicts instead that the intermediate information condition will resemble the partial information condition of the previous experiments, with the players receiving equal probabilities.

Before reporting the results of this experiment, we describe below the methods by which it was implemented. These methods are substantially the same as those employed in the experiment of Roth and Malouf (1979).

5.3. *Methods*

Each participant was seated at a visually isolated terminal of a computer-assisted instruction system developed at the University of Illinois, called PLATO, whose features include advanced graphic displays and interactive capability. The experiment was conducted in a room containing over 70 terminals, most of which were occupied at any given time by students uninvolved in this experiment. Participants were seated by the experimenter in order of their arrival at scattered terminals throughout the room, and for the remainder of the experiment they received all of their instructions, and conducted all communication, through the terminal.

The subjects were drawn from an introductory business administration course taken primarily by college juniors. Pretests were run with the same subject pool to make sure that the instructions to participants were clear and easily understandable.

Background information including a brief review of probability theory was presented first. The main tools of the bargaining were then introduced: these consisted of sending messages or sending proposals. A proposal was a pair of numbers, the first of which was the sender's probability of receiving his/her prize and the second was the receiver's probability. The use of the computer enabled any asymmetry in the presentation to be avoided. PLATO also computed the expected number of chips and the expected monetary value of each proposal and displayed the proposal on a graph of the feasible region. After being made aware of those computations, consistent with the information condition,[17] the bargainer was

[17]In each information condition, a bargainer saw displayed the expected value in money and chips which he would receive from any proposal he made or received. In the high information condition he also saw the same computations concerning his opponent. In the intermediate

given the option of cancelling the proposal before its transmittal. Proposals were said to be binding on the sender, and an agreement was reached whenever one of the bargainers returned a proposal identical to the one he had just received.

Messages were not binding. Instead, they were used to transmit any thoughts which the bargainers wanted to convey to each other. To insure anonymity, the monitor intercepted any messages that revealed the identity of the players. In the low information condition and in the intermediate information condition, the monitor also intercepted messages containing restricted information about the available prizes. The intercepted message was returned to the sender with a heading indicating the reason for such action.

To verify their understanding of the basic notions the subjects were given some drills followed by a simulated bargaining session with the computer. As soon as all the participants finished this portion of the experiment, they were paired at random and the bargaining started.

At the end of 10 minutes or when agreement was reached (whichever came first), the subjects were informed of the results of that game and were asked to wait until all the other bargainers were finished. For the subsequent game there were new random pairings, and the bargaining resumed. The cycle continued until all four games were completed. At no point in the experiment were the players aware of what the other participants were doing, or of the identity of their opponents.

The bargaining process consisted of the exchange of messages and proposals, and participants were instructed that 'your objective should be to maximize your own earnings by taking advantage of the special features of each session'. Only if the bargainers reached agreement on what percentage of the 'lottery tickets' each would receive were they allowed the opportunity to participate in the lottery for the particular game being played. All transactions were automatically recorded.

The lotteries were held after all four games were completed, and each player was informed of the outcomes and the amount of his winnings. A brief explanation of the purpose of the experiment was then given, and the subjects were offered the opportunity to make comments, ask questions etc., and were directed to the monitor who paid them.

5.4. Results

Difference scores, $D = p_1 - p_2$, were again the major dependent variable in

information condition, expected monetary value computations were not displayed for his opponent, and in the low information condition, neither was information concerning the number of chips in his opponent's prize. That is, in each condition, only computations which could be made with the available information were displayed.

the analyses. The independent variables, the four games and the three information conditions, were treated as between factors in a 4 × 4 completely crossed analysis of variance.[18] Two sets of analyses were run, one including all bargaining outcomes (including disagreements), the other including only outcomes where an agreement was reached. The findings for the two analyses were almost identical. The analysis for only the agreements (disagreements excluded) revealed a significant effect for games [$F(3,114) = 14.36$, $p < 0.0001$] and a significant games by information interaction [$F(6,114) = 12.92$, $p < 0.0001$]. The means for the effects are shown in table 4 and fig. 1: the unaggregated data are shown in table 5. The main effect for games appears to be a function of the high value of the prize for player 2 in games 1 and 2 (resulting in large negative scores for D in the full information condition) and the high value of the prize for player 1 in games 3 and 4 (resulting in large

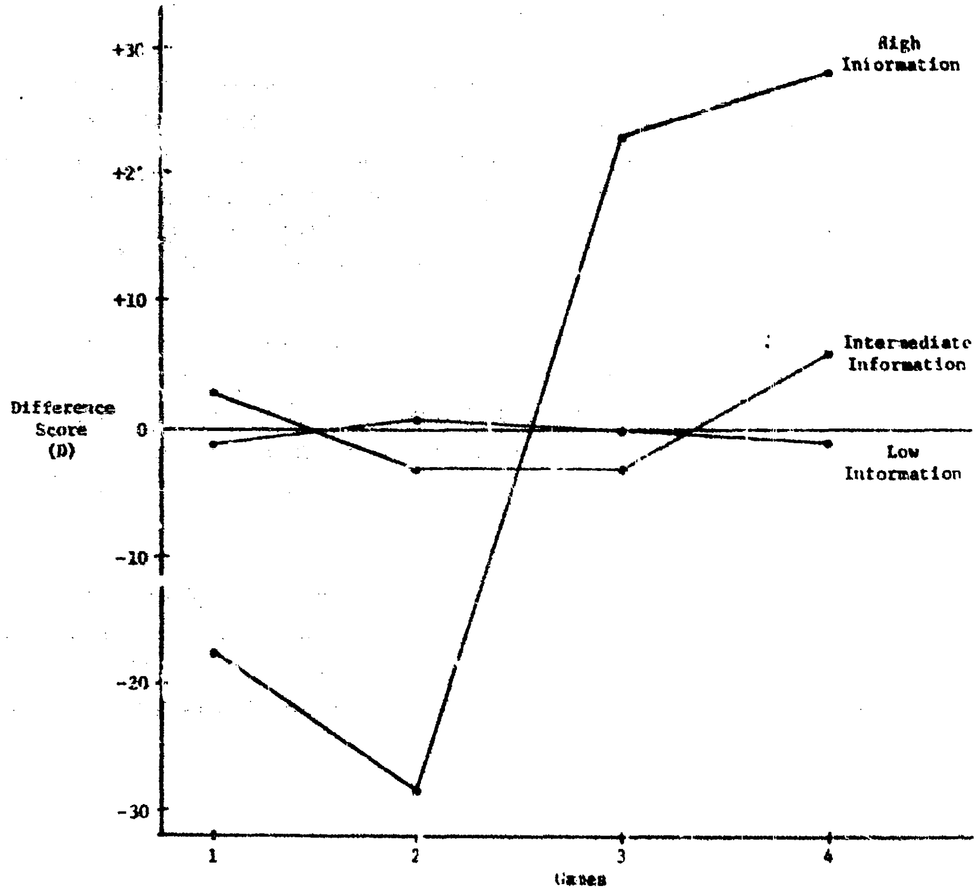

Fig. 1. Means for the games by information interaction.

[18]The data treated each of the games played by each of the subjects as a between rather than within factor. Although every player played each of the four games, they always faced a different opponent. Treatment of the data in this fashion makes statistical tests more conservative.

positive score for D in the full information condition). The interaction was analyzed further by assessing the simple main effects of information within each game. These results, also shown in table 4, indicate that the players' outcomes significantly differed from one another (and from zero) in each of the four games. Post hoc tests for these effects indicated that the high information conditions differed significantly from the low and intermediate information conditions in each game. No other differences were found.

Table 4
Mean difference scores (D) for the games by information interaction (and the associated simple effects).

	Information			Simple effects		
Game	Low	Intermediate	High	F	df	$p<$
1	-0.71_b	2.86_b	-17.50_a	7.17	2.30	0.01
2	1.00_b	-2.85_b	-28.67_a	14.40	2.30	0.001
3	0.00_b	-3.20_b	22.42_a	10.49	2.25	0.001
4	-0.71_b	6.36_b	27.60_a	10.02	2.28	0.001

[a]Cells with common subscripts, within each game, are not significantly different from one another ($\alpha = 0.05$) using the Newman–Keuls procedure.

Additional analyses, combining the data from the current study with the data from the Roth and Malouf (1979) study, tested the differences predicted by the strategic and sociological hypotheses. Both hypotheses suggest that the data from the partial information condition in the Roth and Malouf (1979) study and the data from the low information condition in the current study will be equivalent and will not differ from equal probability outcomes (i.e., $D=0$). Using games as a four level factor[19] and the data from the two studies as a second factor in a between effects analysis of variance, the results showed no significant differences between the two studies or among the four games. The data (see tables 2 and 4) show almost no departure from $D=0$. Thus, when players have no information about their opponent's payoffs, equal probability outcomes predominated in both studies.

The strategic hypothesis predicts that the outcomes in the intermediate information condition in the present study should be similar to the outcomes observed in the full information condition in the Roth and Malouf (1979) study. In other words, the movement away from an equal probability outcome observed in the Roth and Malouf (1979) study is predicted to be observed again in the intermediate information conditions. The sociological hypothesis, on the other hand, predicts that the Roth and Malouf (1979) data from the full information condition will be similar to the data in the

[19]The four games in the two studies are not comparable; this factor was included only to increase the ability to diagnose the cause of any significant effects that might have resulted.

Table 5
Summary of final agreements

		Game 1			Game 2			Game 3			Game 4		
		Player			Player			Player			Player		
Information	Group	1	2	D	1	2	D	1	2	D	1	2	D
High	4	70	30	40	70	30	40	25	75	−50	40	60	−20
		0	0	0	0	0	0	37	63	−26	33	67	−34
		70	30	40	67	33	34	28.5	71.5	−43	25	75	−50
		0	0	0	75	25	50	35	65	−30	51	49	2
	5	50	50	0	50	50	0	50	50	0	0	0	0
		55	45	10	50	50	0	50	50	0	50	50	0
		50	50	0	75	25	50	40	60	−20	0	0	0
		60	40	20	50	50	0	45	55	−10	25	75	−50
	8	60	40	20	65	35	30	40	60	−20	0	0	0
		70	30	40	0	0	0	0	0	0	0	0	0
		0	0	0	80	20	60	0	0	0	100	0	100
	9	60	40	20	50	50	0	40	60	−20	30	70	−40
		60	40	20	60	40	20	50	50	0	35	65	−30
		50	50	0	0	0	0	0	0	0	50	50	0
		50	50	0	80	20	60	25	75	−50	23	77	−54

Inter-mediate	3	50	50	0	50	50	0	50	50	0	50	50	0
		49	51	−2	50	50	0	50	50	0	47	53	−6
		50	50	0	45	50	−5	45	55	−10	30	65	−35
	6	50	50	0	50	50	0	60	40	20	50	50	0
		50	50	0	50	50	0	50	50	0	50	50	0
		45	55	−10	0	0	0	50	50	0	40	60	−20
		46	54	−8	50	50	0	0	0	0	45	55	−10
	7	40	60	−20	52	48	4	50	50	0	50	50	0
		50	50	0	42	58	−16	0	0	0	50	50	0
		0	0	0	52	48	4	0	0	0	50	50	0
		53	47	6	55	45	10	49	51	−2	43	57	−14
	10	50	50	0	50	50	0	52	48	4	50	50	0
		55	45	10	70	30	40	0	0	0	0	0	0
		52	48	4	0	0	0	60	40	20	48	52	−4
		40	60	−20	50	50	0	0	0	0	50	50	0
Low	1	50	50	0	51	49	2	50	50	0	51	49	2
		55	45	10	49	51	−2	50	50	0	50	50	0
		50	50	0	50	50	0	50	50	0	50	50	0
		50	50	0	51	49	2	0	0	0	49	51	−2
	2	47.5	52.5	−5	50	50	0	50	50	0	50	50	0
		50	50	0	50	50	0	50	50	0	47.5	52.5	−5
		0	0	0	45	55	−10	50	50	0	0	0	0
		50	50	0	50	50	0	50	50	0	55	45	10

high information condition of this experiment, and that the intermediate information condition of this experiment should yield results that are not significantly different from the partial information condition of the previous study. Prior to statistical analysis, the data from the current study were transformed to control for the differences between the games. In particular, equal expected monetary value outcomes in games 2 and 4 were 80–20; equal expected value outcomes in games 1 and 3 and in the unequal payoff games of the previous study were 75–25. Thus, the data in games 2 and 4 were transformed; the analyses compared the proportions of movement from equal probability splits toward equal expected value splits in the two studies.

The results are very clear. For only the agreements, a test including the data from the unequal payoff games in the full information condition of Roth and Malouf (1979) and the intermediate information condition from the present study indicates strong differences in the outcomes: $F(1,68)=23.56$, $p<0.0001$; the test comparing the same Roth and Malouf (1979) data and the high information condition of this experiment reveals almost no difference: $F(1,63)<1$, ns. Consulting tables 2 and 4 clearly show the similarity of the data for the Roth and Malouf full information condition and the data from the full information condition in this study. In addition, the simple main effects analysis of the current data also show the marked differences between the full and intermediate information conditions. Thus, the sociological hypothesis is strongly supported.

6. Discussion

The results of this experiment provide strong support for the sociological hypothesis outlined in section 5, and clearly demonstrate that the opportunity to strategically employ arguments concerning the monetary prizes has a markedly different effect than the opportunity to employ strategically equivalent arguments concerning the value of the prizes as expressed in terms of the artificial commodity, chips. Interestingly, this difference does not seem to be due to an unwillingness of the bargainers to advance arguments in terms of chips; informal analysis of the transcripts reveals striking similarities among the messages in the intermediate and high information conditions and in the full information condition of Roth and Malouf (1979).

In each of these conditions, the apparently disadvantaged player — i.e., the player whose prize consisted of fewer chips in the intermediate information condition, or the player whose prize had the smaller monetary value in the full and high information conditions — frequently argued that he should receive a larger probability of winning than his opponent, in compensation for his smaller prize, and claimed that he would insist on doing so. The frequent response of the apparently advantaged player was that a fifty-fifty division of the lottery tickets looked reasonable to him, and that he would take nothing less. But as the results of the experiment showed, this potential standoff was resolved differently in the different conditions. In the intermediate information condition, the player with the smaller number of chips tended to back off from his demand for a higher probability and accept an equal probability of winning, while in the high information condition, and in the full information condition of Roth and Malouf (1979), the player with the higher valued prize tended to back off from his demand for an equal probability and accept a smaller probability of winning (cf., table 4).

In view of the fact that 'disadvantaged' bargainers were so successful in obtaining higher probabilities in the high information condition by employing arguments concerning money, and that they employed strategically equivalent arguments concerning chips in the intermediate information condition, it is all the more surprising that arguments concerning chips had no statistically significant effect on the mean observed agreements. Of course, in the intermediate information condition there was a very small tendency observed in each of the four games for the player with fewer chips to get a higher probability of winning (cf. fig. 1). But, as the figure makes clear, even if this should prove to be a reliable effect, it is an order of magnitude smaller than the corresponding effect observed in the high information condition which resulted in players with a smaller monetary prize receiving a higher probability. Thus the difference between the outcomes observed in the high information condition and those observed in the low information condition cannot be accounted for by models constructed entirely in terms of the feasible utility payoffs and strategy sets of the players. Instead, the outcomes depend, to a significant degree, on the sociological content of the shared information and the feasible messages.

Further examination of the transcripts sheds some light on this sociological content. In a high proportion of the bargaining encounters, notions of equity and fairness were invoked by the bargainers in support of their positions. These notions were invoked strategically,[20] presumably to enhance the credibility of a bargainer's demand. Viewed in this way, the results of the experiment suggest that the reason strategically equivalent arguments did not have the same effect in different information conditions is that different notions of equity need not be equally credible. Specifically, the results suggest that equalizing the probability that each player will win his prize is a more credible notion of equity than equalizing each players' expected payoff in chips, but not in money. Thus information about the

[20]Thus, for example, in the intermediate information condition, a player who suggested that the fairest agreement is to equalize each player's expected number of chips was invariably a player who had a smaller prize in chips than his opponent.

monetary prizes affects the bargaining in a way which it is not affected by information about the prizes in terms of chips.

Of course, these remarks about equity involve some speculation beyond the conclusions established in the last section by statistical analysis of the data. Our intention is merely to suggest an explanatory model for the sociological phenomenon which the data demonstrates so clearly. To formally investigate this particular explanation, it would be necessary to directly investigate (among other things) the degree to which information about a bargaining situation affects people's idea of what constitutes an equitable settlement. As it happens, a study along these lines (conducted independently of our work) reported in Yaari and Bar-Hillel (1980), supports the hypothesis that equity notions are highly sensitive to certain kinds of non-economic information.

The influence of sociological factors on the outcome of bargaining can be observed in non-experimental settings as well. For instance, when an American President issues voluntary wage guidelines as an anti-inflation measure, he does so in order to influence the outcome of industrial wage negotiations. Even though voluntary guidelines need not change any of the economic factors underlying such negotiations, they might influence the outcome by enhancing the credibility of, say, a management proposal which conformed to the guidelines. Similarly, wage settlements reached in one industry may affect the outcome of negotiations in another, even when there is no direct economic influence of wages in one industry on activity in the other.

In conclusion, we have shown that sociological factors — i.e., factors unrelated to what we normally consider to be the 'economic' parameters of a game — can decisively influence the outcome of bargaining.[21] This has both encouraging and discouraging implications for the prospect of developing theories of bargaining.

On the discouraging side, these results suggest that it is unlikely that a general theory of rational behavior in bargaining can be constructed on a purely deductive basis. While there certainly are situations in which strategic factors constrain the range of potential outcomes sufficiently to permit useful conclusions to be drawn from first principles, in many other situations (as in the bargaining games considered here) deduction based simply on the rationality of the players leaves a wide range of potential outcomes (e.g., the contract curve in a bargaining game is the entire set of individually rational, Pareto optimal outcomes). Of course, the discouraging impact of these results is tempered by the fact that, at least since the time of Edgeworth

[21]Schelling (1960), who reached a similar conclusion on intuitive grounds, suggested that experimental methods would be needed to further pursue such matters. We obviously agree. That is, since the sociological context in which bargaining is conducted plays such an important role in determining the outcome, any theory of bargaining — whether descriptive or prescriptive — must necessarily include empirical content, rather than being purely deductive in nature.

(1881), attempts to develop deductive theories of bargaining based primarily on the consequences of individual rationality have met with only limited success.

So the encouraging side of these results is that they suggest an approach which may lead to more successful bargaining theories. Specifically, if certain kinds of sociological information can be incorporated into a theory of bargaining, it may be possible to eliminate some of the indeterminacy which cannot be resolved by theories which depend on purely economic information. We hope to have more to say on this subject in future papers.

References

Binmore, Ken, 1980, Nash bargaining theory II, Mimeo. (London School of Economics, London).

Bransford, John D. and Nancy S. McCarrell, 1975, A sketch of a cognitive approach to comprehension: Some thoughts about understanding what it means to comprehend, in: W.B. Weimer and D.S. Palermo, eds., Cognition and the symbolic process (Erlbaum Associates, NJ).

Edgeworth, F.Y., 1881, Mathematical psychics (Routledge and Kegan Paul, London).

Malouf, Michael W.K., 1980, An experimental investigation of two-person bargaining, Unpublished Ph.D. Dissertation (University of Illinois, Urbana, IL).

Minsky, Marvin, 1975, A framework for representing knowledge, in: P.H. Winston, ed., The psychology of computer vision (McGraw-Hill, New York).

Nash, John F., 1950, The bargaining problem, Econometrica 28, 155–162.

Nash, John F., 1951, Noncooperative games, Annals of Mathematics 54, 286–295.

Roth, Alvin E., 1979, Axiomatic models of bargaining (Springer, Berlin and New York).

Roth, Alvin E. and Michael W.K. Malouf, 1979, Game-theoretic models and the role of information in bargaining, Psychological Review 86, 574–594.

Roth, Alvin E. and Michael W.K. Malouf, 1980, Scale changes and shared information in bargaining: An experimental study, Mimeo.

Roth, Alvin E. and J. Keith Murnighan, 1980, Common knowledge and deniability in bargaining, Manuscript.

Roth, Alvin E. and Uriel G. Rothblum, 1980, Risk aversion and Nash's solution for bargaining games with risky outcomes, Econometrica, forthcoming.

Rubinstein, Ariel, 1980, Perfect equilibrium in a bargaining model, Mimeo. (London School of Economics, London).

Schelling, Thomas C., 1960, The strategy of conflict (Harvard University Press, Cambridge, MA).

Von Neumann, John and Oskar Morgenstern, 1944, Theory of games and economic behavior (Princeton University Press, Princeton, NJ) 2nd ed. 1947, 3rd ed. 1953.

Yaari, Menachem and M. Bar-Hillel, 1980, Judgements on justice, forthcoming.