Economic Experimental Results on Multilateral Bargaining

Glenn W. Harrison, Gordon C. Rausser, and Leo K. Simon

GATT Research Paper 90-GATT 25

Center for Agricultural and Rural Development
Iowa State University
Ames, Iowa 50011

Glenn W. Harrison is with the department of economics, College of Business Administration, University of South Carolina, Columbia. Leo K. Simon and Gordon C. Rausser are members of the department of agricultural and resource economics at the University of California, Berkeley.

This material is based upon work supported by the Cooperative State Research Service, U.S. Department of Agriculture, under Agreement No. 89-38812-4480. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the authors and do not necessary reflect the view of the U.S. Department of Agriculture.

1. INTRODUCTION

In a companion series of papers we have described a new bargaining institution and the design of laboratory experiments to assess its performance (see Rausser and Simon [1991] and Harrison and Simon [1991]). In this paper we describe and evaluate the first series of pilot experiments.

In Section 2 we briefly describe the general Multilateral Bargaining institution once again, so as to make the present discussion self-contained. In Section 3 we run through in some detail the numerical solution to one parameterized version of that model, to familiarize the reader with the structure and workings of the model. In Section 4 we discuss some of the conceptual issues that arise in evaluating the performance of the institution in our experiments. In Section 5 we evaluate the pilot experiments using the procedures developed. Finally, in Section 6 we outline the implications of our results for the next stage of experimentation.

2. THE MULTILATERAL BARGAINING INSTITUTION

The MB institution can be characterized by a model of noncooperative multilateral bargaining with a central player. The model has n+1 players, called the player set. The zero'th player is distinguished from the others and is called the central player. Players 1 through n are peripheral players.

The players participate in a sequential, multilateral bargaining game, similar in spirit to Rubinstein's classic [1982] bilateral game. Their objective in bargaining is to form a coalition, which is just a subset of the player set, and to choose an m-dimensional vector from a set of feasible vectors, called the choice set and assumed to be compact. The choice set may be different for different coalitions.

The central player is distinguished from the others in that she must be included in every coalition. Each player has a utility function defined on the choice set. We assume that utility functions are continuous and strictly quasi-concave.

Problems of this kind are typically formulated as cooperative games. Cooperative game theorists specify some solution concept that satisfies certain appealing properties and then study the set of choices that satisfy the given criterion. Perhaps the most familiar cooperative solution concept is the Core. In the context of the MB institution, a vector x is in the Core if it is feasible for some coalition and if, for every coalition C, there is no feasible vector that is weakly preferred to x by each member of C and strictly preferred by one member.

Noncooperative bargaining theory differs from cooperative game theory in that it attempts to model the actual process of negotiation, rather than just the outcome of the negotiation. A noncooperative model of multilateral bargaining includes an extensive form, which stipulates a particular set of negotiating rules that players must follow.

A natural research program, referred to as the "Nash Program" after Nash [1953], is to study the cooperative and noncooperative versions of a game in conjunction with each other. First one studies a particular cooperative solution concept, then one asks whether the equilibria (usually the subgame perfect equilibria) of some noncooperative model implement the cooperative solutions. Following this approach, we study the relationship between the Core of various bargaining games and the subgame perfect equilibria of our noncooperative version of these games.

The game has a finite number of periods T, each of which is divided into three sub-periods. In the **first** sub-period a player is chosen by Nature to be the proposer. Nature makes it's choice according to a probability distribution over the player set that is prespecified as part of the description of the game. In the **second** sub-period the proposer announces a coalition, of which he must be a member, and a vector that is feasible for that coalition. In the **third** sub-period the remaining members of the proposed coalition each choose whether to accept or reject the proposed vector. If all accept, the game ends. If not, the next period begins and a new proposer is selected. If agreement is not reached by the T'th period then players receive a predetermined disagreement payoff.[1]

A **strategy** for player i specifies the vector that he will announce in each period if selected to be the proposer, as well as a set of vectors that i will accept in each period if he is a member of a coalition announced by some other proposer. A strategy profile is a list of strategies, one for each player. Each strategy profile defines an outcome for the game, which is just a function assigning to each element of the choice set the probability that the game will end with an agreement to select this vector. Note that only a finite number of these probabilities will be positive. Moreover, these positive probabilities need not sum to unity, since the players may never reach an agreement.[2]

A **subgame perfect equilibrium** for a game is a strategy profile with the property that at every sub-period of the game each player's choice is optimal given the strategies specified by the other players. Every T-period game has a subgame perfect equilibrium. Moreover, this equilibrium is generically unique. A striking feature of the model is that there are equilibria in which players fail to agree until the final rounds of bargaining. An

---

[1] The game is well specified whether or not there is a central player. However, the presence of the central player guarantees that the model has a solution. Nonetheless, it can sometimes be instructive to compare our model to the corresponding one in which the central player is excluded (see Harrison and Simon [1991; section 2.3]).

[2] We are just describing the strategy space here. Equilibrium outcomes, to be defined momentarily, will not admit disagreements.

equilibrium outcome is the outcome defined by a subgame perfect equilibrium. Note that since agents may fail to agree at the beginning of the game, the equilibrium outcome need not coincide with the distribution over first period proposals.

Our theoretical analysis concerns the equilibrium outcomes of games with an arbitrarily large number of periods. Accordingly, our bargaining model is defined as a sequence of T-period bargaining games, with T growing to infinity. A solution to our bargaining model is a limit of the equilibrium outcomes of the T-period games.

The first major analytical result for our model is that a solution exists. That is, the outcomes for the T-period games always converge as T grows large. It is here that the central player has a crucial role: when there is no player that is a member of every coalition, T-period outcomes will not in general converge.

A second major result is that, generically, this solution is deterministic. More precisely, there is generically a unique vector x with the property that for every epsilon there exists a T sufficiently large that the agreed upon vector in any game with more than T periods is within epsilon of x with probability one. When such a vector x exists we will refer to it as the solution vector.

Our last major result is that the solution vector is always in the Core of the corresponding cooperative game.

3. A NUMERICAL EXAMPLE

In this section we consider in considerable detail an explicit example that has been solved numerically. As a byproduct of this discussion, the reader will be introduced to our computer algorithm for solving the model. Understanding the algorithm and the solution will help the reader understand the economic logic of the model as well as our approach to its experimental evaluation.

We will refer frequently the computer output which is displayed as Table 1. The first section of the output lists the parameters of the bargaining problem. There are 16 admissible coalitions (numbered from 1 to 16) and six players (numbered from 1 to 6). Each line beginning "Members of coalition number..." is followed by six columns, specifying which players are included in this coalition. For example, coalition #1 consists of players #2, #3, and #5.

The next lines indicate the utility parameters of the five private agents (players 1 through 5). The first five lines give the ideal points, or bliss points, of each player in terms of the horizontal and vertical coordinate that generates the greatest payoff for that player. Thus player #2 has a bliss point of (30, 52), which is to say that she receives the highest possible payoff when the policy values are equal to this. As the policy values deviate from these values, her payoffs decline.

Specifically, the payoff to agent i is a linear function of the Euclidean distance from the ideal point. The intercept of this linear function, denoted α_i, determines the payoff when agent i's ideal point is the chosen policy vector (i.e., when the Euclidean distance from her ideal point is zero). The coefficient of this linear function, denoted β_i, determines the rate at which payoffs decline from the maximum payoff as the Euclidean distance increases. The second set of five numbers in Table 1 describing the utility functions show the values of these two coefficients for each agent.

Each player receives a payoff of zero if there is no agreement.

Each of players #2 through #5 have an equal probability in this game of being asked to make a proposal, but player #1 has 12 times the chance of getting to make a proposal as any of the others (i.e., this player is asked to make the proposal 75% of the time, and each of the other players is asked 6.25% of the time). In the games considered here we do not need to include the government as an active player, hence it has an access probability of zero and is not included in any of the 16 coalitions.

The remainder of the output summarizes the outcome of negotiations in each round of bargaining. Our experiments in this environment run for five rounds. Thus in Table 1 we show the detailed results for each of rounds #1 through #5.

Consider the seven rows of numbers below the statement "Round #1", at the bottom of the table. The first five rows contain nine columns. For $1 \leq i \leq 5$, the first column of row i is the coalition selected by player i in the current round. The second and third columns list the policy vector proposed by i: the second column is the value of the horizontal coordinate, and the third column is the value of the vertical coordinate.

Columns four through eight specify the payoff that each player will earn if the corresponding policy vector is accepted. Thus column four shows that player #1 will receive 90.000, 89.529, 89.397, 89.529 or 89.427 if players 1 through 5, respectively, are selected to be the proposer (and behave optimally).

The sixth row lists the expected payoff for each player conditional on reaching this round of negotiations. It is calculated by simply multiplying the payoff to the agent in that column by the probability that each of the row agents gets to be the proposer. Thus, for player #1, the first payoff listed above is multiplied by 0.75 and the next four payoffs by 0.0625 to obtain the expected payoff in round #1 of 89.868 listed in row 7.

We solve the model by standard dynamic programming techniques, starting from round #5. Our maintained hypothesis is that if no agreement is reached in the last round of negotiations (round #5 here) then each player earns a zero payoff. Consequently, the optimal response for all players except #5 in this round is to propose their globally optimal policy vector. This simply implies that each of these players will propose their ideal point in round #5, which is what we see in Table 1. Since any player except #5 will accept any proposal rather than incur the zero disagreement payoff, the proposer can choose any one of the coalitions excluding player #5 of which she is member. When the proposer is indifferent between coalitions, our computer algorithm chooses the one indexed by the larger number.

Now consider the penultimate round of negotiations, which is round #4 in our Table. A member j of a coalition will accept a policy vector proposed by i in this round if and only if the payoff received by j from the proposal is at least as large as j's **expected** payoff conditional on reaching the next round. For example, player #1 will not accept any proposal in round #4 that does not earn her at least 79.379, since that is her expected payoff from playing in round #5 and she would veto any proposal that gave her less than that.[3]

It follows that to determine her optimal proposal player i must solve a separate nonlinear programming problem for each coalition to which she belongs. This problem ensures that all of the members of the coalition have an incentive not to veto it. In our last example, anybody considering including player #1 in their proposed coalition in round #4 must ensure that the policy proposal generates earnings of at least 79.379 for player #1 (or, to extend the example, 44.829 for player #2, 47.986 for player #3, and so on).

In the current example the policy space has only two dimensions, which we refer to as the horizontal and vertical coordinates. Since each coalition has three or four members, there are two or three "participation constraints" depending on the size of the coalition. In round #2, for example, player #1's

[3] We are implicitly assuming risk-neutrality throughout.

participation is a binding constraint for players #1 and #2. It is not binding for players #3 and #4, since they do not include player #1 in their proposed coalition.

Having solved each of the nonlinear programming problems, player i then picks the coalition that yields her the highest payoff. If the payoff exceeds i's expected payoff conditional on reaching the next round, then i will propose this coalition and the corresponding policy vector. Note that there may be rounds in which member i makes a proposal that is **not** accepted.[4] This does not occur in the numerical example considered here, however.

Consider player #1's choice of coalition in round #5. She chooses coalition #15, consisting of all players except #5. She could have received the same payoff had she chosen coalition #14, which contains the same members as coalition #15 except that player #4 is discarded; it is still a majority coalition. The computer chose coalition #15 simply because the index 15 is larger than the index 14. It is perfectly possible in general that a player can be indifferent in terms of expected payoffs between choosing one coalition or another, even if his policy proposals would differ conditional on either coalition being selected (this is not true for round #5).

The importance of this point will be evident when we come to interpret data from our experiments. It is quite likely that we

---

[4] This can happen for one of two reasons. First, i's best feasible alternative may yield him a lower payoff than his expected payoff conditional on passing to the next round. Second, there may be no proposal available to i that satisfies the necessary participation constraints.

will encounter multiple predictions in the message-space of the experimental game being associated with unique (expected) payoff-space predictions. Thus it will generally be more transparent to test our predictions in terms of how well subjects approximated their maximum possible payoffs rather than how closely their messages correspond to those predicted by our algorithm.

The solution to the MB game in Table 1 is found by allowing the number of rounds to increase until all players make the same policy proposal. Considerable convergence has occurred by round #1, especially considering that we constrain our subjects to report policy values in integers. The solution in this game is the Core outcome (39, 68), which also corresponds to the ideal point of player #1.

Player #1 has a very simple strategy in this game: propose her ideal point whenever asked! Any coalition that does not include player #5 will accept this proposal in any round.

Each of the other players have relatively simple strategies as a function of the round that they are in. As already noted, all except player #5 offer their own ideal point in round #5 if negotiations reach that point. In round #4, however, they compromise their offer in the direction of the Core and away from their own ideal point.

Table 2 displays some results from a game that is virtually identical to the one presented in Table 1. The only substantive difference is that player #1 now has a much greater access probability: she now has the same chance of making a proposal as any other player, whereas in Table 1 she had 12 times the chance of any other player.[5]

The effect of this change is to slow down the convergence to an equilibrium. This is particularly visible by looking at the proposals of players #4 and #5 in round #1. In Table 1 we find that player #4 proposed (46.553, 76.479) in round #1, whereas in Table 2 we find that she proposes (52.161, 99.161), which is further away from the equilibrium relative to the ideal point of (62, 109).

The purpose for constructing the variant in Table 2 is to examine more carefully the individual behavior of subjects taking the roles of players #4 and #5. In each of rounds 1 and 2 these players must make non-trivial changes in their proposals, especially the policy values. In Table 1 we note that players #2 and #3 do not change their policy proposals a great deal until later rounds. We wanted to construct an environment in Table 2 in which two players are predicted to make substantial changes in their policy proposals over just two rounds. This enables us to conduct experiments with a shorter horizon (the experiments corresponding to Table 1 ran five rounds; those in Table ran two rounds). This should allow subjects more time to repeat the game

---

[5] There is one further difference, which is completely unimportant for our experiments: players #2 and #3 are transposed relative to Table 1. Each of these players are computer-simulated in our laboratory experiments in the game shown in Table 2.

within a given real-time frame, and hence to learn more about the game.[6]

[6] Of course, we shuffled opponents from game to game so as to mitigate reputation effects developing due to repeated-game strategizing. See Harrison and Simon [1991] for a discussion of this feature of our experimental design.

4. EVALUATING THE INSTITUTION

There are two general questions to be answered before we wade into an evaluation of the laboratory data that has been collected. These questions will allow us to make some sense out of a great deal of information provided in each experiment.

The first question is: "How should we measure the performance of the MB institution?". One of the most important aspects of the previous discussion of the solution to the MB game is that one cannot just look at a unique set of "messages", defined as the coalition and policy choices sent by the agent. There are often several distinct messages that are optimal in the sense of maximizing expected payoff. Similarly, there may also be several distinct messages that may not be optimal, but which generate the same expected payoff.

How are we to decide between these messages? Our computer programs used an arbitrary rule when selecting optimal coalitions (choose the coalition with the smallest index number). Should we say that an agent in our experiments performed differently from our theory if he or she used a different arbitrary rule but still selected a coalition that generated the maximal expected payoff? Clearly not. The upshot of this issue is that we can either keep track of whether or not an agent selected from the optimal set of choices, or we can just describe his or her behavior in terms of the expected payoff from their choices. We believe that the latter approach is the easiest to follow from an expositional perspective, and adopt it here.[7]

The second question is: "What should we compare the performance of the MB institution to?". In other words, how are we to know if the subjects in our experiment performed well or not?

There are three general ways to make such a comparison. The first is relative to some **theoretical model** of behavior, such as the solution discussed in previous sections. The second is relative to some **numerical model** of behavior in which the agents' choices are simulated by "automata" following prescribed strategies. This approach includes the first as a special case, as we shall see. The third approach is relative to some **alternative institution** dealing with the same bargaining problem. An example of such an alternative is the "Committee Institution" described in Harrison and Simon [1990].

Each of these approaches has virtues, and we propose to eventually use them all. In the present study we will consider in detail the first two approaches, since we have not yet conducted the experiments with an alternative institution required to implement the third approach.

---

[7] We are not saying that it is **more** natural to evaluate behavior in payoff-space than in message-space, which is a claim advanced by one of us in a different context (see Harrison [1989; p.749]) and viewed by many experimentalists as too strong. Rather, our present point is that reporting results in message-space would be very messy, since we would have to report a great deal of information. On the other hand, all of this information boils down to a much smaller set of numbers when expressed in terms of payoff-space, facilitating it's communication and interpretation. We will happily provide the complete data from our experiments on request (contact Harrison).

In order to implement a numerical model of behavior in the MB institution we must define strategies for our simulated players, or "automata", to follow. A natural way to do this, following Gode and Sunder [1990], is to define "zero intelligence" or "minimal rationality" strategies, so as to provide some lower benchmark to evaluate our human subjects.[8] The idea is to see, by means of computer simulation, what level of performance might be expected of automata playing our games with simple strategies. This provides some way of saying how well our human players did in the institution. Of course, we always have an upper bound on how well they can do, in terms of our model of fully rational players.

Consider four possible strategies that could have been followed in terms of **coalition choice**. In all cases we restrict attention to coalitions that contain a majority of players, since our experiments effectively constrain subjects to these proposals. Whenever some randomization is called for, we shall assume that it is implemented using a uniform probability distribution over the alternatives; extensions to this assumption are immediate, but of

[8] Gode and Sunder [1990] illustrate this approach using the venerable Double Auction market in which buyers and sellers of a perishable commodity can openly make bids or offers, as well as accept any outstanding bid or offer. They define zero-intelligence strategies for buyers (sellers) to be the random selection of a bid between zero (unit cost) and unit valuation (zero), and use a uniform distribution for all random realizations. They find that the Double Auction is **never** less than 75% efficient using these strategies, and is typically around 95% efficient for the demand and supply configurations used in experiments. This indicates that observed experimental efficiencies using human subjects of up to 95% should not be attributed to the rationality of subjects so much as the way in which the Double Auction institution constrains and limits the loss of (aggregate) efficiency due to (individual) irrationality.

minor interest. The strategies are presented in increasing order of their rationality from the perspective of the agent entering it:

- **strategy C0**: select any coalition that includes me;
- **strategy C1**: select any of the coalitions of smallest size that includes me, where "smallest size" still ensures that the coalition have a majority; and
- **strategy C^***: select one of the optimal coalitions.

Note that the last strategy is "the strategy" used in defining our theoretical and numerical solution, as presented earlier.

In terms of the **policy choice** strategies, we have a similar list increasing in the rationality of the strategy:

- **strategy P0**: select any values from the rectangle that just includes the ideal points of all of the agents;
- **strategy P1**: select any values from the convex hull of the ideal points of all of the agents;
- **strategy P2**: select any values from within the convex hull of the ideal points of the players in the proposed coalition; and
- **strategy P^***: select the optimal values for the coalition proposed.

Note that the strategies C^* and P^* together define the fully rational strategy choices which provide an upper bound on the expected payoffs that an agent could receive, assuming that all other players were similarly playing rationally. Strategies C0 and P0 together define what we will refer to as the "zero intelligence" outcome.

It is perfectly possible, and indeed behaviorally plausible, to consider combinations of different levels of rationality in the two sets of strategies. Thus we could envisage agents making poor choices with respect to coalitions (viz., using strategies C0 or C1) and yet making excellent choices with respect to the policy values proposed (i.e., using strategy P^*) **conditional on the coalition choice.** This type of decomposition has been used before in comparable experimental evaluations of dynamic programming environments by Harrison and Morgan [1990].[9]

We appreciate that there are many possible variants on these strategies, but these do seem to span a wide range of possibilities and appear plausible as "minimally rational" strategies. We have deliberately stated these strategies in a time-invariant fashion. This is to provide us with a constant basis for comparison of human choices, rather than a measure that must be defined for each time period of a particular game. Allowing some degree of time-dependence would not be difficult, and would be an obvious extension of our approach (e.g., it might seem plausible that players would put greater weight on policy values closer to their own ideal point as we approach the final period in a game).

---

[9] In that setting the agent was selecting how many job offers to purchase in some time period, and then deciding whether or not to continue searching after looking at the purchased job offers. Behaviorally, the subjects seemed to make excellent decisions as to the number of job offers to request conditional on having decided to search in a given period. Their errors in judgement could be almost entirely attributed to errors in deciding when to stop searching. Our situation is conceptually similar.

We implement this approach by conducting a Monte Carlo simulation of the use of each of the specified strategies. In other words, assuming that we have selected one of the set (C0, C1, C^*) and one of the set (P0, P1, P2, P^*), we would play the T-period game a pre-specified number of times (e.g., 10000 times). This is essential to avoid biasing our results due to small-sample results. Note that each play of the game in turn involves another Monte Carlo game, since we already have one Monte Carlo simulation embedded in the experimental game itself (see Harrison and Simon [1990] for further details). Having run this larger Monte Carlo simulation we have a benchmark expected payoff for each agent for this environment, and this may be directly compared to the actual average payoff that the human agent realized.

One aspect of our design requires a natural variation on this approach. This is when we have a subset of the players in an experiment computer-simulated. In such cases the strategy used for the simulated players is what we refer to here as (C^*, P^*). To evaluate the benchmark efficiency levels for minimal rationality strategy selection by one automata, we must similarly employ the fully rational strategies for all of the other automata.

5. EXPERIMENTAL RESULTS

5.1 Experimental Design

A large number of pilot experiments have been conducted using subjects drawn from the undergraduate population at the University of California at Berkeley. Table 3 describes the experiments that we report on here.[10] For convenience we will refer to each of the five sessions as A, A', B, B', and C. A session with the same letter employed identical parameters and only differed with respect to the subjects employed.[11] Accordingly we will typically pool such behavior.

From Table 3 we note that sessions A and A' employed the simple environment of Table 2. There were 12 human subjects in session A and 6 subjects in session A'. In each session we had the human subjects randomly assigned to play the roles of Christine and Stephanie, who are players #4 and #5, respectively, in the earlier numerical examples. All of the other players are computer-simulated using the fully-rational strategies defined in Table 2. Players #1 through #3 are referred to here as Lisa, Paula, and Stephanie, respectively, and the government is referred to as Gordon. This use

---

[10] Many others were conducted, but were discarded for one of several reasons (e.g., software failure, opaque instructions, insufficient time to complete more than one or two periods, or a substantive revision to the instructions). This is quite standard in the development of a complex laboratory experiment.

[11] As Table 3 indicates, there are some differences in the actual number of human subjects in each such paired session, implying some differences in the number of replications per period. This difference should not change behavior in any predictable manner, if at all.

of names follows the experimental setup described in Harrison and Simon [1991].

The number of replications of each game shown in Table 3 refers to the number of bargaining groups that were active concurrently in any session. In session A we had 12 human subjects playing two roles, so there were 6 replications of each game. Thus there were 6 human subjects assigned to play Christine and 6 human subjects assigned to play Stephanie. Each human subject retained the same persona from game to game, but played against a randomly selected opponent.

The horizon of the game in sessions A and A' was 2 rounds, following Table 2. We were able to complete 6 games in each of these sessions, allowing us to see how subjects' performance changes with more experience with the institution.

Finally, the last column in Table 3 refers to the stem of the computer file name used to store the detailed results for each session. These names are not quite as cryptic as they may look, but are intended solely to provide interested readers with access to the detailed results if needed.[12]

---

[12] The last numeric characters indicate the day of the month that the experiment was conducted on, the next to last alphabetic character indicates the month that the experiments were conducted in, and the first two characters indicate if this was the first (E1), second (E2), or third (E3) experiment conducted on that day. Thus experiment E2N21 was the second experiment conducted on November 21. The computer file describing the output of this experiment in complete detail is called E2N21.RAW, and is available in "raw ASCII" form so as to facilitate use by researchers. Similar ".RAW" files are available for each of our experiments on request to Harrison.

Sessions B and B' correspond to the environment described in Table 1. These negotiations employed human counterparts to each of the (active) players, and had a horizon of 5 rounds as in Table 1. Pooling over the two sessions we have 5 replications of each experiment, with at least 4 repetitions of each game.

Finally, session C uses the environment of Table 1 and assigns human subjects to play the roles of Jenny and Christine. Three replications of this experiment were obtained.

5.2 Performance Measures

The first task in our analysis of results is to derive the performance measures for each session using the limited rationality strategies described in section 4. This has not yet been completed, but preliminary numerical results are extremely suggestive. In particular, it appears that the MB institution is quite robust to individual irrationality in the sense that it agents do not have to be perfectly rational in order to achieve "good" overall outcomes. These "good" aggregate outcomes are certainly sub-optimal, but they are "close" to being optimal.

This result, if confirmed, constitutes both Good News and Bad News for the experimental evaluation of the MB institution. The Good News is that we do not need to train subjects as intensively as we had anticipated: they can play sub-optimally and still achieve good outcomes. The Bad News is that it will be difficult to motivate subjects to behave in a fully rational manner: if the social pie only increases by a small amount by each individual being more rational, and an individual only gets some fraction of that increased social pie, why should that individual bother exerting himself?

These conjectures await final numerical results, which will be to hand within a few weeks.

5.3 Evaluation

TO BE WRITTEN.

6. CONCLUSIONS

Although we are yet to complete the evaluation of our pilot experimental results, several conclusions as to the next round of experiments seem warranted.

The first conclusion is that there may be little to gain by conducting experiments in which one human subject plays against computer-simulated (fully rational) opponents. If limited rationality suffices to achieve good social outcomes, then the most interesting environment to examine is one in which **all players are human**. This will provide the best opportunity to see if some player can exploit the limited rationality of the other players.[13] It will also provide the most natural basis for comparison with alternative institutions that might not be so robust to acts of individual irrationality (e.g., committee institutions).

The second conclusion is that we should attempt to see if **experience** makes a difference with respect to the ability of subjects to exploit the limited rationality of other players. It could be that human subjects behave in a limited rational manner when all are inexperienced, but that they try to behave in a strategic manner when they have had some experience. Could this attempt to behave strategically cause the MB institution to generate poor bargaining outcomes?

---

[13] Of course, we could computer-simulate limited-rationality behavior.

The third conclusion is that we should see if these results carry over into **more realistic environments**. We can attempt to estimate spatial preferences that correspond to field bargaining environments (e.g., the GATT negotiations) and see if our early results carry over to them.

REFERENCES

Gode, Dhananjay, and Sunder, Shyam, "Efficiency of Double Auctions with Weakly Rational Traders: Markets as a Partial Substitute for Individual Rationality", **Working Paper #1990-15**, Graduate School of Industrial Administration, Carnegie-Mellon University, May 1990.

Harrison, Glenn W., "Theory and Misbehavior of First-Price Auctions", **American Economic Review**, v.79, September 1989, pp.749-762.

Harrison, Glenn W., and Morgan, Peter, "Search Intensity in Experiments", **The Economic Journal**, v.100, June 1990, pp.478-486.

Harrison, Glenn W., and Simon, Leo K., "Experimental Assessment of a Multilateral Bargaining Institution", **Unpublished Manuscript**, Department of Economics, University of South Carolina, January 1991.

Nash, John F., "Two-Person Cooperative Games", **Econometrica**, v.21, 1953, pp.128-140.

Rausser, Gordon C., and Simon, Leo K., "A Noncooperative Model of Multilateral Bargaining", **Unpublished Manuscript**, Department of Agricultural and Resource Economics, University of California at Berkeley, January 1991.

Rubinstein, Ariel, "Perfect Equilibrium in a Bargaining Model", **Econometrica**, v.50, January 1982, pp.97-109.

Table 1
A Numerical Example

Members of coalition number 1:	OUT	IN	IN	IN	OUT
Members of coalition number 2:	OUT	IN	IN	OUT	IN
Members of coalition number 3:	OUT	IN	OUT	IN	IN
Members of coalition number 4:	OUT	OUT	IN	IN	IN
Members of coalition number 5:	IN	IN	IN	OUT	OUT
Members of coalition number 6:	IN	IN	OUT	IN	OUT
Members of coalition number 7:	IN	IN	OUT	OUT	IN
Members of coalition number 8:	IN	OUT	IN	IN	OUT
Members of coalition number 9:	IN	OUT	IN	OUT	IN
Members of coalition number 10:	IN	OUT	OUT	IN	IN
Members of coalition number 11:	OUT	IN	IN	IN	IN
Members of coalition number 12:	IN	OUT	IN	IN	IN
Members of coalition number 13:	IN	IN	OUT	IN	IN
Members of coalition number 14:	IN	IN	IN	OUT	IN
Members of coalition number 15:	IN	IN	IN	IN	OUT
Members of coalition number 16:	IN	IN	IN	IN	IN

Ideal point of player number 1:	39.000	68.000
Ideal point of player number 2:	30.000	52.000
Ideal point of player number 3:	25.000	72.000
Ideal point of player number 4:	62.000	109.000
Ideal point of player number 5:	165.000	32.000

Utility coefficients of player number 1: (alpha, beta)	90.000	1.000
Utility coefficients of player number 2: (alpha, beta)	70.000	1.000
Utility coefficients of player number 3: (alpha, beta)	70.000	1.000
Utility coefficients of player number 4: (alpha, beta)	90.000	1.000
Utility coefficients of player number 5: (alpha, beta)	110.000	1.000

Round #5							
#15	39.000	68.000	90.000	51.642	55.440	42.989	-21.042
#15	30.000	52.000	71.642	70.000	49.384	24.632	-26.473
#15	25.000	72.000	75.440	49.384	70.000	37.674	-35.602
#15	62.000	109.000	42.989	4.632	17.674	90.000	-18.600
#10	126.269	45.996	0.000	-26.456	-34.554	0.000	68.817
Expected payoffs:			79.379	44.829	47.986	41.761	-16.523
Round #4							
#15	39.000	68.000	90.000	51.642	55.440	42.989	-21.042
#5	33.793	58.743	79.379	62.263	54.092	32.369	-23.905
#5	28.788	70.918	79.379	51.044	66.060	39.470	-31.663
#4	46.553	76.479	78.645	40.450	47.986	53.997	-16.523
#3	53.287	61.555	74.327	44.829	39.846	41.761	-5.557
Expected payoffs:			86.983	51.143	54.579	42.717	-20.634
Round #3							
#15	39.000	68.000	90.000	51.642	55.440	42.989	-21.042
#2	38.472	64.495	86.456	54.904	54.579	39.659	-20.634
#1	35.005	70.180	85.449	51.143	59.831	42.717	-25.486
#4	40.403	71.255	86.456	48.114	54.579	46.513	-20.634
#3	42.915	65.740	85.480	51.143	51.023	42.717	-16.662
Expected payoffs:			88.990	51.563	55.331	42.967	-20.995
Round #2							
#15	39.000	68.000	90.000	51.642	55.440	42.989	-21.042
#2	38.699	66.754	88.718	52.872	55.331	41.755	-20.995
#1	37.564	68.813	88.350	51.563	57.038	42.967	-22.647
#4	39.403	69.217	88.718	50.383	55.331	44.247	-20.995
#3	40.416	67.213	88.380	51.563	53.858	42.967	-19.465
Expected payoffs:			89.635	51.631	55.427	42.988	-21.038
Round #1							
#15	39.000	68.000	90.000	51.642	55.440	42.989	-21.042
#2	38.875	67.545	89.529	52.099	55.427	42.532	-21.038
#1	38.475	68.297	89.397	51.631	56.025	42.988	-21.628
#4	39.134	68.452	89.529	51.182	55.427	43.449	-21.038
#3	39.500	67.722	89.427	51.631	54.882	42.988	-20.484
Expected payoffs:			89.868	51.641	55.440	42.989	-21.043

Table 2
Another Numerical Example

Round #2							
#15	39.000	68.000	90.000	55.440	51.642	42.989	-21.042
#15	25.000	72.000	75.440	70.000	49.384	37.674	-35.602
#15	30.000	52.000	71.642	49.384	70.000	24.632	-26.473
#15	62.000	109.000	42.989	17.674	4.632	90.000	-18.600
#10	126.269	45.996	0.000	-34.554	-26.456	0.000	68.817
Expected payoffs:			56.014	31.589	29.841	39.059	-6.580
Round #1							
#15	39.000	68.000	90.000	55.440	51.642	42.989	-21.042
#5	25.000	72.000	75.440	70.000	49.384	37.674	-35.602
#5	30.000	52.000	71.642	49.384	70.000	24.632	-26.473
#6	52.161	99.161	56.174	31.589	17.892	76.085	-21.314
#10	71.774	59.005	56.014	21.455	27.643	39.059	12.941
Expected payoffs:			69.854	45.574	43.312	44.088	-18.298

Table 3
Experimental Design

Session	Environment	# Subjects	Human Players	Simulated Players	Replic-ations	Horizon	Repet-itions	File Name
A	Table 2	8	Christine Stephanie	Lisa Paula Jenny	4	2	6	E1N21
A'	Table 2	12	Christine Stephanie	Lisa Paula Jenny	6	2	6	E2N21
B	Table 1	10	Lisa Paula Jenny Christine Stephanie		2	5	4	E1N19
B'	Table 1	15	Lisa Paula Jenny Christine Stephanie		3	5	5	E2N20
C	Table 1	6	Jenny Christine	Lisa Paula Stephanie	3	5	4	E1N20