AN OUTSIDE OPTION EXPERIMENT*

KEN BINMORE
AVNER SHAKED
JOHN SUTTON

In the economic modeling of bargaining, outside options have often been naively treated by taking them as the disagreement payoffs in an application of the Nash bargaining solution. The paper contrasts this method of predicting outcomes with that obtained from an analysis of optimal strategic behavior in a natural game-theoretic model of the bargaining process. The strategic analysis predicts that the outside options will be *irrelevant* to the final deal unless a bargainer would then go elsewhere. An experiment is reported which indicates that this prediction performs well in comparison with the conventional predictor.

I. INTRODUCTION

The Nash bargaining solution has been widely used as a modeling tool for wage negotiations in applied economics. Recent progress in noncooperative, game-theoretic models of bargaining [Binmore, 1985; Binmore and Dasgupta, 1987; Binmore, Rubinstein, and Wolinsky, 1986; Rubinstein, 1982; Shaked and Sutton, 1984; Sutton, 1986] suggests that some of the modeling problems are not quite so simple as is often assumed. The difficulty considered here is the manner in which the bargainers' outside options are incorporated into the Nash solution.

The Nash bargaining solution [Nash, 1950] is formulated in terms of a set[1] X of utility pairs that represent possible deals on which two bargainers may agree, and a disagreement pair (d_1,d_2) that represents the utilities the bargainers will receive if there is no agreement. The Nash bargaining solution (s_1,s_2) is then the point in X at which the Nash product $(s_1 - d_1)(s_2 - d_2)$ is maximized subject to the constraints $s_1 \geq d_1$ and $s_2 \geq d_2$. In this paper, and in most applications, the agreements amount to sharing a sum M of money (which will not be available without an agreement) between bargainers whose utilities are linear in money. In this case X is a set of the form $\{(x_1,x_2) : x_1 \geq 0,\ x_2 \geq 0,\ x_1 + x_2 \leq M\}$, and $s_i = d_i + (M - d_1 - d_2)/2$ $(i = 1,2)$. The Nash bargaining solution then

*The financial support of STICERD at LSE and of the ESRC; and the excellent research assistance of Yasmin Batliwala, Ami Klin, Nikki Boyce, Maria Herrero, and Carol Van Der Ploeg are acknowledged. Al Roth provided much useful criticism.

1. Which is usually assumed to be convex, closed, bounded above, and comprehensive.

The Quarterly Journal of Economics, November 1989

assigns each player his disagreement payoff plus half what remains from M after the disagreement payoffs have been made.

However, in applications there is often more than one candidate for the disagreement point. One possible candidate is the *impasse point,* by which we mean the utility pair that will result if the bargaining continues forever without agreement being reached or the negotiations being abandoned. We always normalize the impasse point at (0,0). But such an impasse is not the only route that may lead to a failure to agree. Perhaps one or other of the bargainers may unilaterally abandon the negotiations to take up an opportunity elsewhere. Or, perhaps, if agreement is delayed, the opportunity the bargainers are planning to exploit jointly may be lost through the intervention of some random factor outside the bargainers' control. The utility pairs that will arise as a consequence of such breakdowns in the negotiation process provide further candidates for the disagreement point in Nash's solution. In what follows, we shall assume that only one such *breakdown point* (b_1,b_2) is possible and that $b_1 \geq 0$, $b_2 \geq 0$, and $b_1 + b_2 \leq M$. Breakdown will be assumed to be precipitated by one or other of the players leaving the negotiation table for good in order to take up his outside option b_i. The other bargainer is then assumed to follow suit.

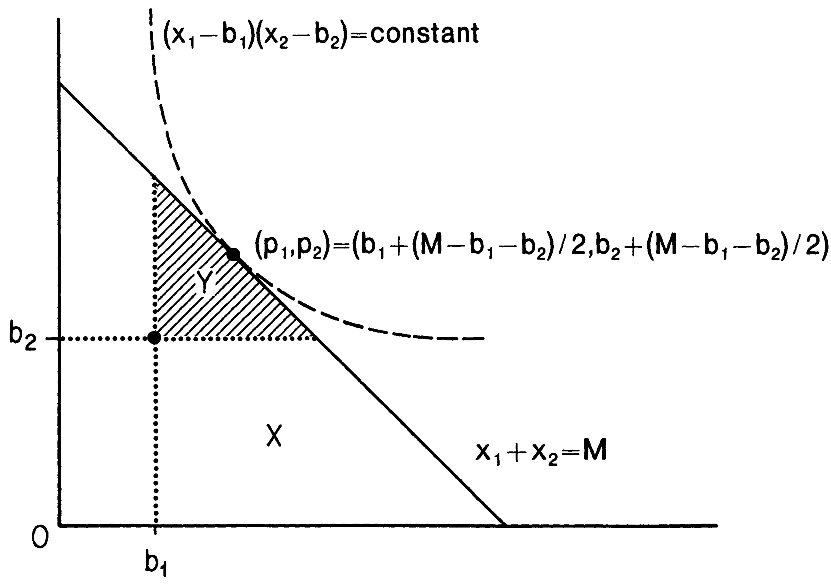

FIGURE I

In wage negotiations it is appropriate to think of the points in X as wage-profit flows, the impasse utilities as income flows during a strike and the outside options as the best income flows available to each side if they cease their partnership altogether.

In such a context it is conventional to place the disagreement point for the Nash solution at the *breakdown point* (b_1,b_2) as indicated in Figure I. A useful mnemonic for the prediction (p_1,p_2) of the bargaining outcome so generated is *split-the-difference.* The paper contrasts this predictor with a special case of an "outside option principle" derived from an analysis of optimal behavior in a natural game-theoretic model of the bargaining process. The mnemonic used for this special case is *deal-me-out* for reasons to be explained shortly. It selects the bargaining outcome (q_1,q_2) indicated in Figures IIa and IIb. This is the Nash bargaining solution for the set $Y = \{(x_1,x_2)\colon b_1 \leq x_1,\ b_2 \leq x_2,\ x_1 + x_2 \leq M\}$ with the disagreement point at (0,0). Thus, outside options are only used as *constraints* on the range of validity of the Nash bargaining solution. The disagreement point is placed at the impasse point (0,0). With deal-me-out the predicted bargaining outcome is

$$(q_1, q_2) = \begin{cases} (b_1, M - b_1, & \text{if } M/2 < b_1 \\ (M - b_2, b_2), & \text{if } M/2 < b_2 \\ (M/2, M/2), & \text{otherwise,} \end{cases}$$

and so each bargainer gets a half-share of the whole sum of money unless this would assign one bargainer less than his outside option.

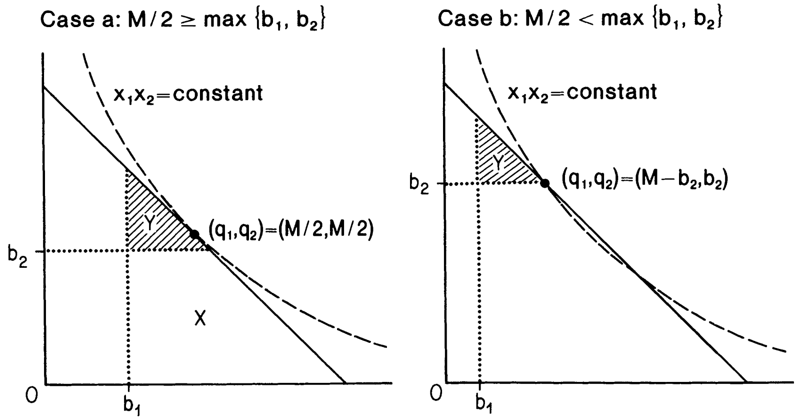

FIGURE II

In the latter case that bargainer receives his outside option, and the other bargainer gets the rest.

The appropriate form of the "outside option principle" is justified formally in Appendix 1 by identifying the unique subgame-perfect equilibrium of a Rubinstein-type bargaining game with alternating offers from which either player can secede after refusing an offer to take up his outside option. Deal-me-out arises when the discount factor of δ in this analysis is approximately one.

Strategically, what is involved is very simple. The attraction of split-the-difference lies in the fact that a larger outside option seems to confer greater bargaining power. But how can a bargainer use his outside option to gain leverage? By threatening to play the deal-me-out card. When is such a threat credible? Only when dealing himself out gives the bargainer a bigger payoff than dealing himself in. It follows that the agreement that would be reached *without* outside options is *immune* to deal-me-out threats, unless the deal assigns one of the bargainers less than he can get elsewhere. The opponent need then only offer him epsilon on top of his outside option to keep him at the table. The theory idealizes epsilon to be zero. In real life, epsilon would need to be chosen sufficiently large not to be dismissed as negligible.

This paper reports the result of an experiment in which anonymous subjects played a Rubinstein-type game with outside options. Deal-me-out predicted the outcomes overwhelmingly better than split-the-difference. If one is willing to believe that the stylized negotiations procedure which constrained our subjects bears a sufficient resemblance to that used in relevant real-life situations, and if one is also willing to believe that the laboratory behavior of our subjects is similarly significant, then our results would seem to *refute* the conventional use of split-the-difference in this context.

Is there any point in such an experimental refutation? Is it not enough to show that the conventional predictor attributes suboptimal behavior to the bargainer? Such naive questions neglect the accumulated evidence that, in laboratory bargaining experiments, subjects seldom take proper account of strategic factors and prefer to settle on deals that are "fair" in some sense (e.g., Güth et al. [1982], Hoffman-Spitzer [1985]). Do the current results not contradict this evidence? In brief, one is not entitled to argue that deal-me-out predicts better *because* it represents optimal behavior. We do not, in fact, believe that our subjects know all about subgame-perfect equilibria and are gifted with the capacity for effortless mental arithmetic. Without extensive opportunities for trial-and-error learning, they can only be anticipated to have a *dim* awareness of the strategic realities. Since our game is very simple, being nearly symmetric when both outside options are zero, it may be that such a dim awareness is enough to generate behavior close to the optimum in strategic terms. But the very symmetry that makes such a scenario plausible simultaneously makes it difficult to distinguish such an explanation of the observed behavior from one which postulates that the subjects are partly motivated by "fair-ness' considerations. The issue of the extent to which "fairness' genuinely motivates subjects in bargaining situations is taken up elsewhere [Binmore, Shaked, and Sutton, 1989]. The current paper is content to establish that deal-me-out predicts better than split-the-difference without committing itself to *why* this should be so.

II. The Bargaining Game

This section briefly describes the rules of the game played by our subjects. The manner in which these rules were operationalized is left to the next section.

A "cake" originally worth £7 (approximately $10 at the time) is to be divided between two players if they can agree on how it is to be divided. The bargainers are constrained to employ the following very specific bargaining procedure. Player 1 begins by proposing a division of the cake to player 2. Player 2 then accepts or refuses this proposal. If he refuses, player 2 may then decide not to continue bargaining but to take up an outside option. In the experiment, games were divided into three groups: in group 1, player 2's outside option was zero; in group 2 it was £2; and in group 3 it was £4. For simplicity, player 1's outside option was *always* zero.[2] If player 2 refuses player 1's offer but does not opt out, then all the sums of money mentioned above are reduced by a factor of $\delta < 1$, and a second round of negotiations takes place, just like the first, but with player 2 making an offer to player 1. This procedure continues with the players alternating in being the proposer until either

(a) agreement is reached; or
(b) a player opts out; or
(c) a cutoff point is reached at which the available payoffs have become negligible.[3]

2. Players with a zero outside option were not explicitly reminded of their opportunity to opt out.
3. In practice this cutoff point was never reached.

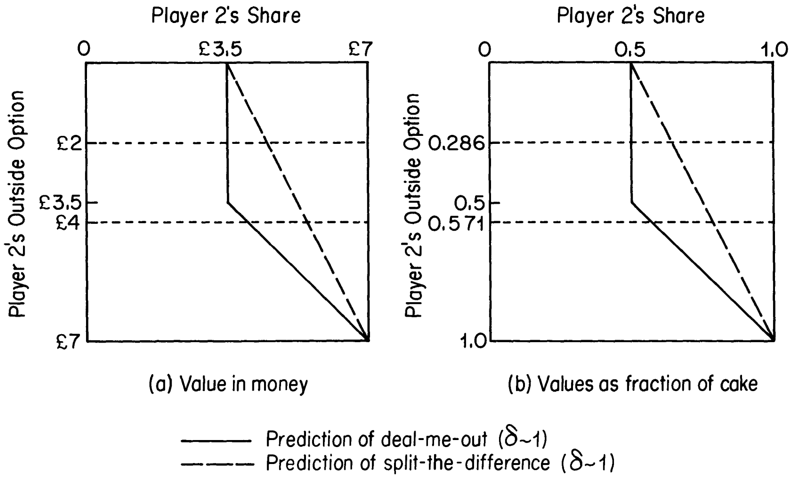

FIGURE III

All this information was known to both players,[4] but much care was taken to ensure that neither player became aware of the real-life identity of his bargaining partner.

Figure III compares the predictions of split-the-difference and deal-me-out. These are appropriate when δ is approximately one. In the experiment, δ was actually taken to be 0.9 so that all sums of money shrank by 10 percent before each new round of negotiation. This explains the slightly different predictions indicated in Figures IV, V, and VI.

Enough information has been provided to appreciate the results of Section IV, but a detailed analysis requires some further comments on the design of the game.

1. The game is based on a bargaining model of Rubinstein [1982] in the belief that this model, with its explicit pattern of offer and counteroffer, captures an essential aspect of real-world bargaining institutions.

2. The game admits an explicit game-theoretic analysis. It has a unique subgame-perfect equilibrium outcome (as proved in Appendix 1). In the zero option case this requires the first player to offer $\delta/(1 + \delta)$ of the cake to the second player and for the second player to accept. Here δ is the players' common discount factor. With $\delta = 0.9$, as in the experiment, the fraction of the cake to be offered is therefore 0.473. Since this is nearly 0.5, the game-theoretic analysis

4. The game is one of perfect information.

therefore leads to an approximately 50:50 split as would, for example, an analysis based on attributing motives of fairness to the players. With a positive outside option for player 2, all remains precisely the same unless $\delta/(1 + \delta)$ of the cake is less than player 2's outside option. If so, then the equilibrium outcome is for player 1 to offer player 2 his or her outside option instead. Since $2 < 0.473 \times 7 = 3.311 < 4$, an equilibrium outcome in groups 1 and 2 of our experiment requires that player 2 gets 0.473 of the cake while, in group 3, player 2 gets just his outside option (which is worth 0.571 of the value of the cake).

3. The above analysis treats money as infinitely divisible. There is also the fact that equilibrium behavior requires specific selections to be made from actions between which a player is indifferent. In particular, players are always indifferent between accepting or refusing an equilibrium offer, but in equilibrium they accept. With a discrete currency the indifference issue can be resolved in theory, since players can always "play safe" by making their offer better than the alternative by an amount equal to the smallest coin available.[5] Of course, this smallest coin will be regarded as "negligible" by most subjects. One must therefore expect to see larger "token" amounts in practice. Rather than commit ourselves to a view on how large such a "token" amount should be taken to be, we increased the size of the experimental cake from £3 in our pilot study to £7 in the main study so that relevant numbers to be compared were always substantially different.

4. The game-theoretic analysis predicts that agreement will always be reached at the very first opportunity. But, implicit in a noncooperative game theory analysis is the hypothesis that it is common knowledge that the players are rational. In real life even a player who is rational himself might reasonably entertain doubts about the rationality of an anonymous opponent. Delaying agreement might then be worthwhile to provide an opportunity of learning whether the opponent is exploitable. However, even when agreement is not immediate, game theory still provides a prediction of future play *conditional* on no agreement having been reached so far. In odd-numbered periods when player 1 makes the offer, the prediction is just as in (2) above. In even-numbered periods player 2 makes the offer. In groups 1 and 2 the equilibrium outcome then

5. Although this will be only one of many equilibria in the discrete case. However, these all approximate the unique equilibrium of the continuous case provided that the smallest unit of currency is sufficiently small.

gives $1/(1 + \delta) = 0.526$ of the available cake to player 2. In group 3 the equilibrium outcome gives player 2 somewhat more than the current value of his outside option (namely, 0.614 of the value of the currently available cake). It will be noted that a game-theoretic analysis attributes a slight advantage to the player who has the opportunity of making the first proposal.

5. The shrinkage factor $\delta = 0.9$ was chosen with two considerations in mind. The aim was to make the rate of shrinkage fast enough to "blanket" any difference in the "natural" rates of time preference of the subjects but slow enough that the first-mover advantage mentioned in (4) be relatively small. The situation discussed in the introduction is, strictly speaking, the limiting case as $\delta \rightarrow 1-$. In what follows, it should be noted that the deal-me-out prediction is the unique subgame-perfect equilibrium outcome calculated for the *actual* shrinkage factor of $\delta = 0.9$ with the first-mover advantage taken into account.

6. A split-the-difference analysis gives 0.5 of the available cake to player 2 in group 1 games; $0.643 = 4.5/7$ in group 2 games; and $0.786 = 5.5/7$ in group 3 games. Split-the-difference can also be used to predict the bargaining outcome *conditional* on no agreement having been reached so far. Because it is favorable to the split-the-difference predictor, we adopt an interpretation in even-numbered periods which takes account of the fact that player 2 must wait one period before his or her outside option is available again. In groups 1, 2, and 3, respectively, the prediction then is that player 2 gets 0.5, 0.629, and 0.757 of the available cake when making an offer.

7. Deal-me-out is only one of various alternatives to split-the-difference which might be considered. Methodologically, it has a considerable advantage over the other alternatives in that it yields precise and unambiguous predictions, and hence we cannot be accused of altering our rival predictor to suit the data. Given that our rival predictor does better than split-the-difference, we therefore have a sound case for rejecting the latter. But it is not claimed that we necessarily have a good case for rejecting anything else.

III. Experimental Setup

Following pilot studies, 120 subjects were recruited from a wide cross section of LSE students in the social sciences.[6] Students who

6. Including economics, law, demography, social anthropology, politics, management science, sociology, geography, psychology, and computing.

had been exposed to game theory or bargaining models were excluded. Recruitment was carried out from teaching classes and not from a pool of subjects accustomed to psychological experiments. Each student was assigned a "time-slot." To preserve anonymity, the two subjects assigned to the same time-slot were always drawn from different classes. Much care was taken to ensure that subjects had no knowledge of the identity of their opponent either before or after the game.[7] The 60 pairs were partitioned into the three groups itemized in Section II. Group 1 (the "control" group) contained 10 pairs. Groups 2 and 3 each contained 25 pairs.

Subjects were placed in separate rooms before microcomputers linked by a cable. After reading a set of written instructions (Appendix 2), the subjects were "talked through" the instructions again by a research assistant to ensure that they were clearly understood. Reinforcement of the instructions, together with practice in the use of the necessary computer controls, was provided with the help of a video display unit (VDU). The subjects did *not* play a practice game with the computer since we were anxious not to offer cues about what type of play was expected. For the same reason, we were not present in the room ourselves.

The VDU displayed a picture of a rectangular "cake" The player making an offer could divide the cake into two shares by pressing designated keys which moved the dividing line between the share claimed and the share offered up or down. The monetary value of the cake and the value of the share claimed were also displayed. The responding player registered acceptance or rejection of the offer by pressing the Y or N keys accordingly. Players were paid in cash immediately after the game finished.

IV. Results

We report the results using diagrams. The raw data appear in our working paper [Sutton et al., 1985]. The three histograms, Figures IV, V, and VI, group data in bands equal to a 1 percent share of the cake. Offers and agreements are *always* expressed in terms of the amount of the share proposed for, or received by, *player 2*.

Consider Figure V by way of example. Observe that in group 2 games (with player 2's outside option at £2) 11 of the 25 games

7. Thus, subjects could not verify that they had a human opponent. But this is unavoidable if anonymity is to be fully preserved.

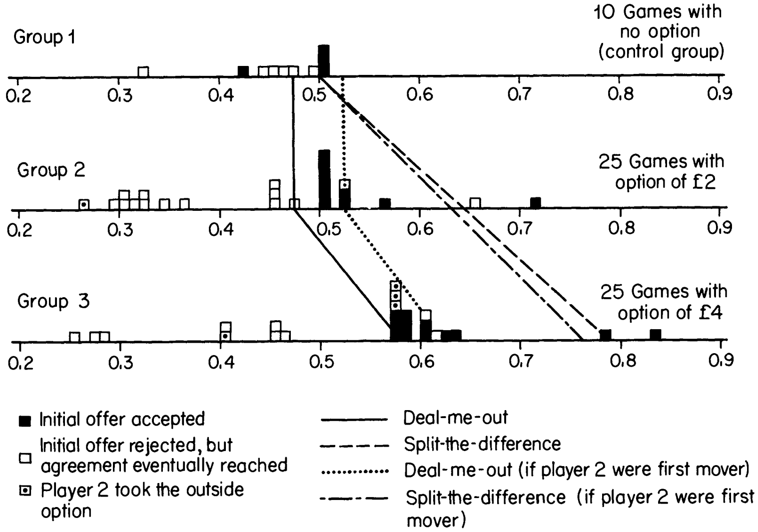

FIGURE IV
The Amount Received by Player 2 as a Fraction of the Original (£7) Cake

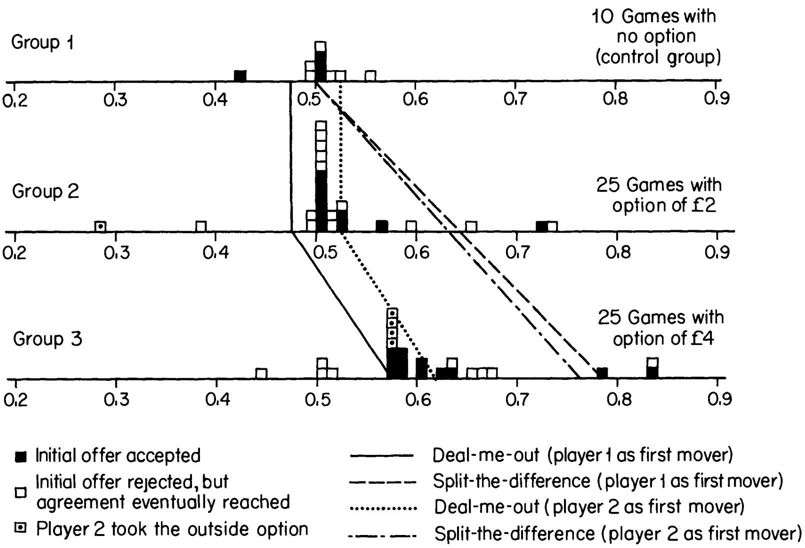

FIGURE V
Player 2's Final Payment as a Fraction of the Cake Available When Bargaining Concluded

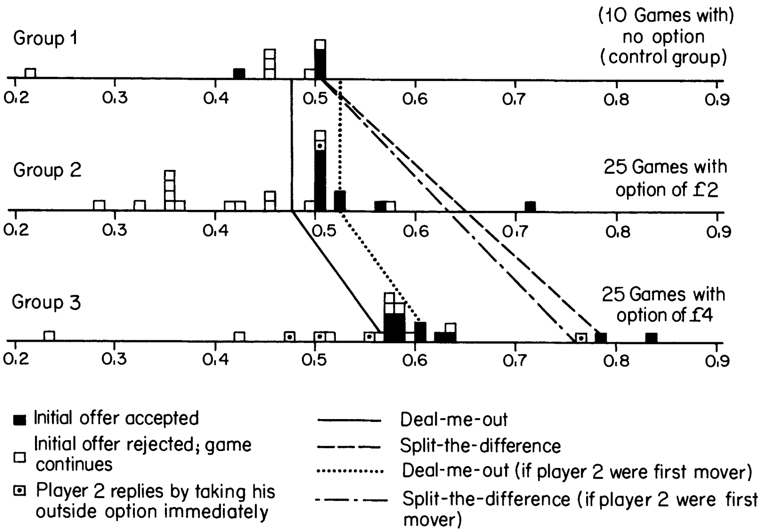

FIGURE VI
Fraction of the Cake Originally Proposed by Player 1 as Player 2's Share

concluded with player 2 receiving a share of between 0.50 and 0.51 of the cake available when the bargaining finished. In 6 of these 11 games agreement was reached immediately. Observe that in group 3 games (with player 2's outside option at £4) 7 of the 25 games concluded with player 2 receiving a share of between 0.57 and 0.58 of the cake available when the bargaining finished. In three of these seven games agreement was reached immediately. In the remaining four games agreement was never reached, since player 2 chose to take his or her outside option rather than continue bargaining.

V. COMMENTARY

The inadequacy of split-the-difference as a predictor, as compared with deal-me-out, is clearly exhibited in Figures IV, V, and VI. A feature of the results is the substantial number of failures to agree at the first opportunity. Figures VII and VIII give the full details of the histories of games that lasted at least three rounds. Presumably there would have been more of these games, and with longer histories, if the shrinkage factor $\delta = 0.9$ had been chosen

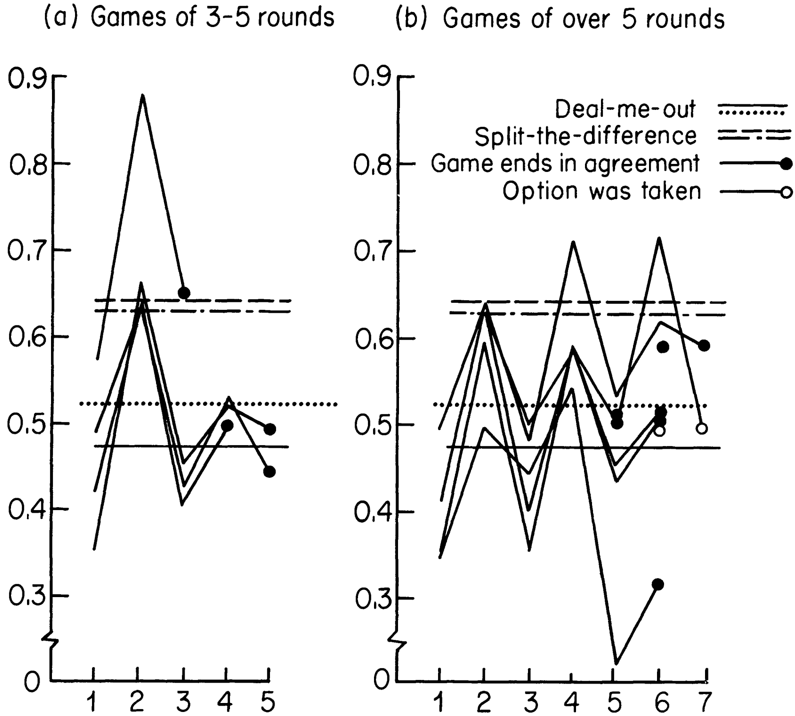

FIGURE VII
Group II (Small Option) Games Involving Disagreement
The Diagram shows the share proposed for player 2 in the successive rounds as a fraction of the cake then available.

closer to one.[8] None of the currently popular bargaining theories predicating perfect information predict disagreement at all, and to this extent, the data are not supportive of any of them. Further research is clearly necessary on this point.

However, results that may be thought to be surprisingly sharp are obtained by examining what player 2 gets once agreement has been reached as a fraction of the *cake available at the time of agreement* (see items (4), (5), and (6) of Section II). These amounts are shown in Figure V. We test the extent to which this impression of sharpness is accurate by asking the following questions:

a. Does the fraction of the available cake obtained by player 2

8. On the other hand, the first-mover advantage would have been diminished. However, the final agreements reached do not support the hypothesis that the first player was able to exploit his first-mover advantage even if he perceived that he had one.

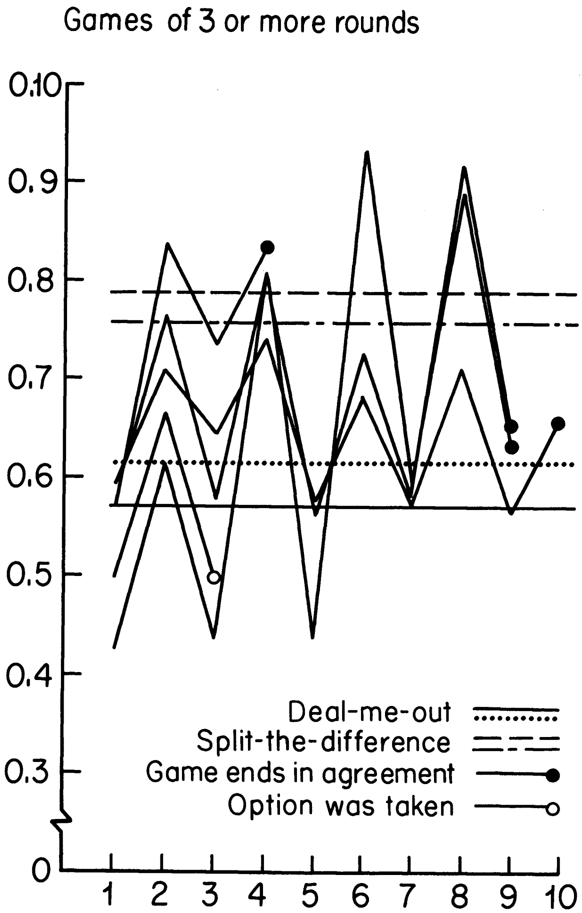

FIGURE VIII
Group 3 (Large Option) Games Involving Disagreement
The Diagram shows the share proposed for player 2 in the successive rounds as a fraction of the cake then available.

in a group 2 game (£2 option) exceed the fraction he obtains in a group 1 game (no option)?

b. Is the fraction of the available cake obtained by player 2 in a group 3 game (£4 option) nearer the split-the-difference fraction than the current value of the outside option?

For question a we tested the null hypothesis that the proportion of agreements in which player 2 gets a fraction strictly exceeding 0.5 is equal in both groups 1 and groups 2. (In view of the clustering at 0.5, this seems a more appropriate criterion than does a test for the equality of the median outcome, which is in fact 0.5 in both groups.) In accordance with our rejection of split-the-difference as a predictor, we observe that the null hypothesis is *not* rejected by a χ^2 test at the 5 percent level.

For question b we observe that the fractions of the available cake obtained by player 2 in group 3 games fall into three classes. Four points lie well below 0.571 (which is the value of the outside option to player 2 as a fraction of the available cake when player 1 is proposing). These outcomes are incompatible with either deal-me-out or split-the-difference, however loosely defined. Of the remaining 21 points, 18 are closer to deal-me-out (0.571 for agreement in odd-numbered periods). The null hypothesis—that the fraction of the relevant population generating outcomes closer to deal-me-out is less than one half—is rejected at the 5 percent level by the present data.

Finally, it should be noted that the intuition for split-the-difference is not without some support from the data. The counterproposals made by player 2s who had refused the opening proposal in Figure VII (showing group 2 games) cluster around the split-the-difference level. However, most of these counterproposals were refused.

VI. Conclusion

Split-the-difference has been widely used to predict the outcome of wage negotiations in applied economics. The theoretical foundations for this predictor have been questioned in Binmore [1985] and Shaked and Sutton [1984]. The current paper provides experimental support for these doubts.

Some care is necessary in evaluating the implications of the results. They are not immediately relevant if incomplete information or reputation effects are important. Nor are they relevant if breakdown may occur through random events outside the control of the bargainer (see Appendix 1). Even where they are relevant, it must be remembered that social benefit, for example, may not only be a factor in determining a worker's outside option: it may also be a factor in determining the size of the available "cake."

Finally, it cannot properly be argued that the results demonstrate that our subjects were motivated largely by enlightened self-interest. They were clearly unwilling to settle for less than their outside option, but within this constraint a "fairness" explanation of their behavior is consistent with the data. Our latest experimental study [Binmore, Shaked, and Sutton, 1989] bears on these issues. We shall only observe here that the results of the new study support the rejection of split-the-difference under conditions comparable to those of the current paper.

Appendix 1: The Unique Subgame-Perfect Equilibrium Outcome

In this appendix we begin by offering a formal demonstration that the infinite horizon version of the bargaining game of Section II has a unique subgame-perfect equilibrium outcome. We also compute the equilibrium payoffs. (The necessary argument is only sketched in Binmore [1985], while Shaked and Sutton [1984] is unpublished.) We assume a common discount factor δ $(0 < \delta < 1)$ and let player 2's outside option be s $(0 \leq s \leq 1)$. Without loss of generality the original size of the cake is taken to be 1.

Let m_1 and M_1 be the infimum and supremum of equilibrium payoffs to player 1 in the game. Let m_2 and M_2 be the infimum and supremum of equilibrium payoffs to player 2 in the companion game in which it is player 2 who moves first. We claim that the following inequalities hold:

$$m_1 \geq 1 - \max\{\delta M_2, s\} \tag{1}$$

$$1 - M_1 \geq \max\{m_2, s\} \tag{2}$$

$$m_2 \geq 1 - \delta M_1 \tag{3}$$

$$1 - M_2 \geq \delta m_1. \tag{4}$$

Inequality (1) follows from the fact that, in equilibrium, player 2 *must* accept any opening offer y with $y > \max\{\delta M_2, s\}$ because the right-hand side is the *most* that player 2 can get from refusing. Thus, in equilibrium player 1 cannot get less than x, where $x < 1 - \max\{\delta M_2, s\}$, because he can always guarantee x by making x his opening demand. Inequality (2) follows from the fact that, in equilibrium, player 2 must get at least z, for each $z < \max\{\delta m_2, s\}$, because z can be guaranteed by refusing player 1's opening offer. Hence player 1 can get at most $1 - a$ in equilibrium. Inequalities (3) and (4) are just the same, but with the roles of players 1 and 2 reversed, and $s = 0$.

We distinguish three cases: (a) $s \leq \delta m_2$; (b) $\delta m_2 < s < \delta M_2$; and (c) $\delta M_2 \leq s$. Case (a) leads immediately to the conclusion that $1/(1 + \delta) \leq m_1 \leq M_1 \leq 1/(1 + \delta)$ and $1/(1 + \delta) \leq m_2 \leq M_2 \leq 1/1(+ \delta)$. Thus, in case (a), $m_1 = M_1 = m_2 = M_2 = 1/(1 + \delta)$. The same argument applied in case (b) yields the contradiction $1/(1 + \delta) < m_2 \leq M_2 \leq 1/(1 + \delta)$. In case (c) the conclusion is that $1 - s \leq m_1 \leq M_1 \leq 1 - s$ and $1 - \delta(1 - s) \leq m_2 \leq M_2 \leq 1 - \delta(1 - s)$. Thus, $m_1 = M_1 = 1 - s$, and $m_2 = M_2 = 1 - \delta(1 - s)$. Using the computed values of m_2 and M_2, it only remains to observe that case (a) occurs when $s \leq \delta/(1 + \delta)$ and case (c) occurs when $s \geq \delta/(1 + \delta)$.

This shows that, if subgame-perfect equilibria exist, then they generate a unique outcome. Existence, however, is trivial. Each player always demands his equilibrium payoff when proposing and accepts his equilibrium payoff (or more) when responding. Section I of the paper describes the limiting case as $\delta \rightarrow 1-$.

Split-the-difference can also emerge from a noncooperative analysis under suitable conditions. To see this, suppose that the game we have just studied is modified so that outside options are no longer available but that, after each refusal of a proposal, a breakdown in communications occurs with probability π, resulting in the payoff pair $(0,s)$ regardless of any desire the players may have to continue negotiating. For simplicity, we take $\delta = 1$. The inequalities of the preceding analysis are replaced by

$$m_1 \geq 1 - \{(1 - \pi)M_2 + \pi s\} \tag{5}$$

$$1 - M_1 \geq (1 - \pi)m_2 + \pi s \tag{6}$$

$$m_2 \geq 1 - (1 - \pi)M_1 \tag{7}$$

$$1 - M_2 \geq (1 - \pi)m_1, \tag{8}$$

from which it follows that $m_1 = M_1 = (1 - s)/(2 - \pi)$ and $m_2 = M_2 = \{1 + (1 - \pi)s\}/(2 - \pi)$. The limiting case $\pi \rightarrow 0+$ is split-the-difference.

In the finite horizon cases uniqueness is immediate, but the computation of the equilibrium payoffs is tedious. These converge to the infinite horizon payoffs as the horizon is allowed to recede to infinity.

Appendix 2: Instructions to Subjects

The following instruction sheet was given to subjects filling the role of Player 1 in those games where Player 2 had an outside option that was initially worth £4. The instruction sheets given to players in other conditions were similar to this. Having read these instructions, subjects were "talked through" them by an assistant, and then the rules were explained again by means of a display on the VDU.

INSTRUCTIONS TO PLAYER 1

The aim of this exercise is to examine how people behave in bargaining situations.

You will be asked to divide a cake (worth a certain sum of money) between yourself and an opponent.

The initial value of the cake is £7.00.

At certain times, your opponent can, if he/she wishes, "opt out," and be paid a certain sum (initially £4.00); if he/she does this, you will receive nothing.

You do not have any such "outside option."

As bargaining continues over time, these values will be reduced, in a manner to be explained below.

Incidentally, it was decided at random before you came in, who would have the "outside option."

The way bargaining will proceed is as follows:

You will make your opponent an offer of some share of the cake. Your opponent can do one of 3 things:

1. Accept your offer in which case the game ends. And you and your opponent each receive the agreed amount.
2. Your opponent can decide to "opt out" of the game in which case he/she will be paid £4.00 and you will receive nothing.
3. Reject the offer—in which case the cake shrinks by 10 percent *and so does the outside option.* Now it becomes your opponent's turn to make you an offer.

The cake is now worth £6.30. Your opponent makes you an offer. You can do one of two things:

1. Accept your opponent's offer, in which case the game ends and you and your opponent receive the agreed amount.
2. Reject your opponent's offer, in which case the cake, *and the option,* shrink by a further 10 percent and it becomes your turn once again to make your opponent an offer.

The game continues in this way, with the sums of money shrinking by 10 percent following each rejection, until an agreement is reached.

All this information is known to your opponent.

A computer demonstration now follows.

UNIVERSITY OF MICHIGAN AND LONDON SCHOOL OF ECONOMICS

REFERENCES

Binmore, K., "Bargaining and Coalitions," in *Game Theoretic Models of Bargaining,* A. Roth, ed. (Cambridge: Cambridge University Press, 1985) (previously ICERD discussion paper 83/71).

——, and P. Dasgupta, eds., *Economics of Bargaining* (Oxford: Basil Blackwell, 1987).

——, A. Rubinstein, and A. Wolinsky, "The Nash Bargaining Solution in Economic Modeling," *Rand Journal of Economics,* XVII (1986), 176–88.

——, A. Shaked, and J. Sutton, "Testing Noncooperative Bargaining Theory: A Preliminary Study," *American Economic Review,* LXXV (1985), 1178–80 (previously ICERD discussion paper 84/99).

——, ——, and ——, "Testing Noncooperative Bargaining Theory: A Reply," *American Economic Review,* forthcoming.

——,——, and ——, "Do People Exploit Their Bargaining Power?" CREST discussion paper, University of Michigan, 1984.

Guth, W., R. Schmittberger, and B. Schwarze, "An Experimental Analysis of Ultimatum Bargaining," *Journal of Economic Behavior and Organization,* III (1982), 367–88.

Hoffman, E., and M. Spitzer, "Entitlements, Rights and Fairness: An Experimental Examination of Subjects' Concepts of Distributive Justice," *Journal of Legal Studies,* XIV (1985), 259–97.

Nash, J., "The Bargaining Problem," *Econometrica,* XVIII (1950), 155–62.

Rubinstein, A., "Perfect Equilibrium in a Bargaining Model," *Econometrica, L* (1982), 97–110.

Shaked, A., and J. Sutton, "The Semi-Walrasian Economy," ICERD discussion paper 84/98, London School of Economics, 1984.

Sutton, J., "Noncooperative Bargaining Theory—An Introduction," *Review of Economic Studies,* LIII (1986), 709–24.

——., A. Shaked, and K. Binmore, "An Outside Option Experiment," ICERD discussion paper 85/124, London School of Economics, 1985.