

Negotiation as a Learning Process

JOHN G. CROSS

*Department of Economics
University of Michigan*

This paper presents a discussion of the role of adapting expectations in the bargaining process. Negotiators are characterized as persons who choose bargaining strategies in their attempt to optimize their payoffs from the situation. These strategies are contingent on each party's perception of the strategy of his opponent, and if these perceptions contain errors, expectations will change and this will lead in turn to a modification of each party's strategy choice. The payoff demands and manipulative moves which characterize the bargaining process are seen as combinations of actions which are specified in the original bargaining plans of the parties and of changes in the plans themselves. The influence of the learning process on the settlement point is described as well as some empirical implications of the theory in general.

SOME ALTERNATIVE PERSPECTIVES ON BARGAINING

The most cursory survey of the literature devoted to the bargaining problem is sufficient to impress one with the enormous diversity of opinion on the nature and function of bargaining processes. The breadth of this diversity is partially demonstrated by the variety of titles which have been assigned to papers in this journal. Evidently, one event which we all may agree to call a negotiation may nevertheless be described by several different scientists almost as though it were as many different kinds of event. The comparison and evaluation of existing theories has become a severe challenge because the aspects of bargain-

ing of most concern to a commentator are so often different from those stressed by the formulator. The author who describes the division of an established joint benefit is criticized for neglecting the strategic or manipulative nature of bargaining; the theory of manipulation is rejected because it stresses "dirty tricks" which are inappropriate in the context of a problem in cooperative decision-making; and a model which describes the value of cooperative ventures is easily condemned for overlooking the crucial problem of how the fruits of cooperation are to be divided up. The old story of the blind men describing an elephant could not be more to the point.

This state of affairs has doubtless arisen as a natural consequence of our limited understanding of an extremely complicated phenomenon. Nevertheless, if we are to describe the elephant piecemeal, it is doubly important that each of us take care to define the limits of his own area of concern. In an effort to provide such a focus for this paper, I would like to begin by outlining what I see to be the most significant properties of our common problem, so that I can indicate where in subsequent arguments these properties become relevant.

From statements of negotiators themselves and from descriptions of both professional and journalistic observers, we may distinguish at least four fundamentally different characterizations of the function of bargaining processes.

(1) Most simply, bargaining may be nothing more than a charade. According to this view, the parties have a common (although unstated) understanding of what the final agreement will be, and the elaborate sequence of bids and counterbids, threats, strikes, and other uses of force has as its sole purpose the gratification (or appeasement) of third parties. This interpretation is often applied to labor negotiations in industries which have long histories of experience with unionism. In 1947, Lemuel Boulware, negotiating opposite the Electrical Workers Union on behalf of the General Electric Company, argued that both parties fully understood what the final settlement would be, and maintained that all of the intervening bargaining was nothing more than an expensive fraud.¹ During that negotiation, and subsequently during the 1950s and the early 1960s, General Electric outlined to the union what it expected the agreement to be and refused to move from that position unless the union succeeded in presenting previously unknown "facts" which altered GE's forecast of the course of an ordinary

1. See, for example, Stevens (1963), pp. 34-35.

negotiation. During this period, General Electric suffered a series of long and expensive strikes, and many observers believed these to have been occasioned by the disappointment and hostility which this policy generated among union members who were not granted the concessions from initial positions which had become traditional. Indeed, Boulwarism has often been characterized as an unfair labor practice, not because of any obvious bias in the estimated agreement points (the proposed settlements were often generous), but because it is seen as an attempt to discredit union leaders in the eyes of the membership. In effect, it denies to these leaders the opportunity to appear to their constituents as having squeezed concessions out of reluctant employers, thus resulting in a severe political liability for them.

The characterization of a negotiation as a charade often accompanies a belief that the range of possible agreements is actually very narrow. Suppose that we describe a two-person negotiation in terms of the set of possible outcomes. We define $x_A \equiv x_{a1}, \dots, x_{an}$ as the payoff to party A at the time of settlement, where x_{ai} is the quantity of good i contained in that payoff. Similarly, $x_B \equiv x_{b1}, \dots, x_{bn}$ is the vector of quantities in the payoff to be received by party B. Corresponding to any payoff x_A , there are limits to the available payoff to B, where these limits may be represented by the relations $x_B \leq H(x_A)$. The possible payoffs to each party are given a lower bound by the fact that one always has the option of abandoning the negotiation if the settlement falls below some minimum (which may correspond to the potential return from establishing a cooperative agreement with someone else). In the case of a single good, these boundaries are commonly described in a two-dimensional diagram as in Figure 1.

The settlement point can only appear somewhere in the area abc of Figure 1. There is no reason to expect the length ab to cover a large portion of the function $x_B = H(x_A)$, however. Indeed, returning to the example of wage negotiations, it is often possible for economists to disregard the bargaining process entirely and use ordinary market variables to explain the bulk of empirical variations in wage settlements, suggesting that the indeterminacy left to be resolved through negotiations is so small as to be neglected entirely. Such empirically based suggestions that the range ab is very small lends considerable support to the Boulware view that the elaborate negotiations which seem to take place are only staged for the sake of appearances.

(2) According to a second view, bargaining is a game of chance and skill, something like chess or poker. It may even be played for fun. A




Figure 1

skillful negotiator is someone who makes good moves in the same sense that a chess master makes good moves. One could aptly describe the negotiation as the process whereby a tourist establishes prices to be paid to a middle eastern curio dealer.

This second view of bargaining should place great stress on the rules on which the game is to be played. Any game must be governed by some system of rules, and since the bargaining game is not described by Hoyle or any other acknowledged authority, these rules must be established by tradition, or perhaps even by the parties themselves. Naturally, they may vary from place to place. The sophisticated tourist will learn these rules before entering the curio shop; otherwise he may lose the game merely because he does not know how to play. By inadvertently breaking an unwritten rule, he may even terminate the negotiation altogether if the other party interprets his action as evidence of bad faith and an unwillingness to play the game fairly.

It is in recognition that all negotiations contain some elements of this game process, that many of us may be suspicious of Thomas Schelling's (1960) proposed strategies for dealing with what he calls "mixed-motive" situations. Many of his "pre-commitments," threats, and uses of force or randomized force require the introduction of unorthodox elements to circumstances in which negotiation, or some

implicit process very like negotiation is taking place. These unconventional actions may easily be seen as rule-breaking acts and therefore as evidence of bad faith. This interpretation in turn can disrupt the process, delay the settlement, and perhaps even destroy the possibility that an agreement of any kind will ever be reached. Examples of this problem are readily found in the area of labor-management disputes. In particular, first-time negotiations with newly formed unions are often accompanied with great bitterness, strikes, and even violence as tactics are employed which do not enjoy mutual acceptance as fair plays. It is not until after the parties have developed commonly understood (although perhaps unwritten) ground rules that the bitterness and violence subsides. Indeed, once this understanding has been achieved, even strikes may occur in an atmosphere of business as usual.

The point is that it is a mistake to treat negotiation as though it was a special case of warfare. The two may have points in common, but the objective of negotiation is always cooperation. The agreement is the beginning of a relationship rather than the end of one, and the few-holds-barred character of open conflict is simply inappropriate to such a situation. To my mind, the important question is not how force may be used, or why it works, but how it is that certain tactics come to be seen as legitimate moves in a negotiation while others do not.

(3) The aspect of the bargaining process which receives the greatest stress in formal analyses is, of course, its mechanism for dividing the fruits of cooperation among two or more participants. The emphasis is on division rather than cooperation, however, in that the dimensions of the benefits to be allocated to the parties are assumed to have been well established already, and the range of possible settlements is taken to be relatively large. Referring again to Figure 1, each point in the area bounded by abc is beneficial to both parties. Both sides are assumed to have full knowledge as to the elements of this set, with the exception that each party may be unaware of the extent of the benefit (or utility) which an agreement would confer upon the other, and the problem is reduced to that of selecting one agreement point from all the possibilities. Since the elements in the vectors x_A and x_B have been defined so that more is always preferred to less, one can reasonably confine the analysis to the line ab on the boundary.

It is usual for analyses of this problem to take advantage of the full-information assumption, and to express the payoffs in terms of utility functions: $U_A(x_A)$ and $U_B(x_B)$. The line segment ab is then shown as a comparable segment of a utility possibility set. The advantage in

this procedure lies in the reduction of payoff vectors to single-valued utility indices. In effect, no matter how complicated the original situation may be and no matter how many dimensions may be required to define an agreement point, we can state the problem in only one dimension: the selection of a settlement point along the boundary of the utility-possibility set.

Invariably, these utility functions are written on the assumption that each party gains satisfaction only from his own payoff so that he is essentially indifferent as to the payoffs which may be received by anyone else. For the purposes of many kinds of analysis, this assumption seems unobjectionable, but it does weaken any characterization of a negotiation as a game. The satisfaction which one receives from playing a game may be as much a matter of relative as of absolute performance, and if Arthur is pleased because he has driven a hard bargain with Bill, the traditional utility formulation will not be adequate.

(4) Each of the foregoing three interpretations of the bargaining process presumes, more or less, that the set of all possible agreements $x_A \leq H(x_B)$ is known to both parties. It is obvious, however, that in most well-publicized negotiations this is far from the truth, and that in fact much of what is happening centers around a search for mutually beneficial agreements. Indeed, negotiators themselves frequently describe bargaining primarily as a search process. In the face of an imperfect understanding of one another's preferences, each party is exploring a list of issues, some of which are already under discussion and some of which will have to be introduced for pairs of items on which concessions may be profitably exchanged. An issue of small importance to Arthur but of great significance to Bill may be granted to Bill, in exchange for a concession from Bill which is of relatively greater importance to Arthur. There is no reason for the issues to be related in any way other than that they are of mutual concern to the parties. A union may agree to take steps to reduce absenteeism in exchange for an employer's commitment to increase plant security, or the Soviet Union may agree to on-site inspections in exchange for a reduction in some specific American warhead delivery system. Conceivably, the issues need not arise even in the same negotiation. The Soviet Union may make a concession in the Strategic Arms Limitations Talks after closing a large wheat deal with the United States.

In all cases of traded concessions, the division problem is still important in that the terms of trade must be agreed on, but the major effort lies in finding such mutually beneficial exchanges in the first place. The

negotiation as a whole may actually consist of a series of component bargains in each of which (1) the parties discover that two issues afford possible exchanges, (2) terms of trade are settled on, and (3) an agreement is struck (initialled) before the next pair of tradeable issues has even been discovered.

Unfortunately, the literature on negotiation is almost silent on this searching and settling-in-sequence aspect of the process. Most models of bargaining (my own included) stress the full-information, one-dimensional problem, and rely on a paradigm in which the two parties concern themselves with the division of a sum of money, or the determination of an hourly wage. It is particularly distressing to observe that when some counterpart of this paradigm actually arises in practice, the bargaining process as we usually think of it either changes its character or ceases altogether. If it is only the division of an established sum of money which is at issue, most people seem to resort to simple ethical rules such as fair division, or they take advantage of some division rule which has been established in a previous (perhaps even expired) contract. If labor-management negotiations are finally reduced to only one issue (such as wages), the dialogue often seems to stop, mediators and other third parties may be called in to assist, and sometimes strikes occur; all of these being evidence of a failure in the negotiating process. In short, the searching and trading which takes place in a negotiation may be more than simply an interesting dimension of the problem, and may actually be essential to its effective operation.

The searching process may even play a part in solving the division problem. To the extent that the large negotiation is composed of a sequence of smaller parts, the resolution of early issues may set patterns and precedents to be followed by subsequent exchanges among the parties. The disposition of major issues in the final agreement may be conditioned in large part by information and expectations which are established early in the process.

THE FOCUS OF THIS PAPER

We have described a negotiation as a combination of searching, dividing, game-playing, and fraud. To put all these together into one paper is certainly more than we can accomplish, and, indeed, the emphasis here will remain on the traditional question of division. The

best we will be able to do with respect to these other elements is point out their relevance from time to time in the discussion, leaving, reluctantly, the broader questions to another time.

Since we are to concentrate on the question of division, we will accept the usual assumption that the set of potential agreements (such as the area bounded by abc in Figure 1) is known to both parties, and we will put aside the interesting problem of how the payoff possibility set $x_A \leq H(x_B)$ was discovered in the first place. Moreover, although we do not intend the theory to be so restricted, the diagrams will have to be constructed for the case of only one payoff good.

If we were to ask a witness of a negotiation to describe precisely what he has seen and heard, we would be given a list of dated events. At a time which we might for concreteness call day 1, one party, Arthur, made a public statement as to the firmness of his resolve in the forthcoming negotiation. On day 2, his counterpart, Bill, published a response. On day 3, Arthur made a persuasive case and Bill responded with a threat; an actual offer (which amounts to a demand) was put on the table by Arthur on day 4, and so on. This went on until some date at which the two sides reached agreement.

We may divide these events into two classes: actual payoff demands, and a series of statements, threats, and coercive actions which are not directly related to the payoff, but which are intended to influence the ultimate settlement. These latter we may term manipulative moves. Although they are not a part of the agreement themselves, they do affect the overall value of the negotiation, partly because they may be successful in influencing the settlement point, and partly because they may impose costs of their own. A threat, for example, may be expensive to make and expensive to carry out, and if it were not for its potential impact on the settlement, both parties would prefer that it did not occur. In general, we will treat manipulative moves as potentially costly actions which may change outcomes but which never directly provide positive utility to either party.

Both payoff demands and manipulative actions are dated. We might describe Arthur's behavior for example with a vector $X_A \equiv x_{A1}, x_{A2}, \dots, x_{At^*}$ which represents his payoff demand on day 1, day 2, day 3, and so on until day t^* when the settlement is achieved, and a vector $T_A \equiv t_1, t_2, \dots, t_m$ which represents the particular dates on which Arthur chose to employ the various manipulative moves which are available to him. A value in this vector of $t_i > t^*$ is taken to mean that move i was never used. All of this may be described by a diagram such as Figure 2.




Figure 2

Figure 2 represents Arthur as employing only manipulative moves numbered 1, 2, 4, and 7, his first demand not occurring until well after the negotiation has started, and actually retracting an offer during the course of the bargaining.

The points on Figure 2 represent only those overt demands and moves which are obvious to our disinterested witness. They may actually be a reflection of a much more elaborate and subtle game plan which Arthur has chosen to employ. Arthur may have outlined for himself a whole array of contingency plans which will provide responses to Bill's behavior: "If Bill reduces his demand, then I will reduce mine," or "If Bill makes a threat, then I will reduce my concessions." There may even be random elements in these plans: "If Bill makes a large concession, I will flip a coin in order to determine my own next move." In order to distinguish the behavior we observe from the game plan which may underlie it, we will use the expression r_A to represent the vector of actual observations $[T_A, X_A]$, and a variable S_A to represent a choice of game plans. S_A , in effect, is nothing more than a list of possible r_A -vectors, each of which is conditional upon some pattern of behavior from Bill.

In a similar fashion, we may distinguish a vector r_B of overt demands and manipulative moves made by Bill in the course of the negotiation and define a variable S_B to represent the strategy which underlies this behavior.

It is tempting to regard the values of X_A and X_B as the actual demands of the two parties. On day 3, the demand x_{A3} is on the table, so to speak, and immediate agreement would be possible if Bill were to accede to it. In fact, Bill may be able to achieve immediate settlement on much more favorable terms. The demand x_{A3} may be an exaggeration, a bluff, and both parties may know it. Moreover, as we have already mentioned, established-bargaining situations are filled with conventional, if unwritten, rules of play, and many of these have to do with the end game. Final resolution of a negotiation may be achieved by splitting the difference on all remaining issues, trading the remaining issues in some more or less arbitrary way, or calling in some authoritative final arbiter. Both parties know of these rules and will automatically subtract the value of their influence from an opponent's stated demand. Even though our independent witness may observe a demand of x_{A3} , both Arthur and Bill may know the real demand to be smaller.

A GAME THEORY ANALYSIS

The formal theory of games was originally formulated for the purpose of treating just such a situation as we have described. If the two parties choose game plans (strategies) S_A and S_B , and if these plans are unchanged throughout the course of the negotiation, then knowledge of the rules of play and S_A and S_B themselves would enable one to predict the entire course of events without having to observe the process at all. S_A and S_B together are sufficient to determine the date of settlement, t^* , the nature of the agreement, and the choice and timing of the various manipulative activities.

The value of the total negotiation to each party is determined by three things: (1) the actual payoff received from the settlement, (2) the date of settlement (a settlement long delayed may be worth considerably less than it would have been had it come more promptly), and (3) the costs incurred from the various manipulative actions which have been taken during the course of the negotiation. Since these may all be determined from S_A and S_B together, we may write the utility functions in the form $U_A(S_A, S_B)$ and $U_B(S_A, S_B)$. Based on the assumption that the two parties are fully informed as to the payoff possibilities, the problem becomes the ordinary one of expected utility maximization, and each party will wish to choose a strategy, S , which will maximize his expected payoff given the strategy choice taken by his competitor.

A fundamental presumption of the theory under discussion is that delays in settlement reduce the value of agreement. In fact, the mechanism of adjustment to be described in the next section uses this time factor as a major motivating force, but here let us argue only that an agreement indefinitely delayed is regarded by both parties as essentially worthless. This means that the strategy choices are only concerned with dates earlier than some upper-bound K . Under this assumption, it is possible to prove that our specification of the problem guarantees the existence of at least one equilibrium pair of strategies S_A^* , S_B^* where S_A^* maximizes $U_A(S_A, S_B^*)$ and S_B^* maximizes $U_B(S_A^*, S_B)$.²

We certainly do not believe that practicing negotiators are so knowledgeable, so single-minded, or so adept at mathematics as this game-theoretic view suggests; nevertheless, as a point of departure for the analysis of bargaining, this paradigm is very useful. Even though we cannot maintain that the parties actually succeed in discovering their optimal strategies, we can argue that they would like to, and that any party who knew the strategy choice of his competitor and possessed the knowledge and ability necessary to exploit this information would see himself as enjoying a distinct advantage. Moreover, a negotiator with this orientation would be aware of the general properties which might characterize the equilibrium of the idealized model, and he would expect these same properties to be relevant also to his own imperfect efforts to achieve this ideal. There are in fact three characteristics of the game-theoretic equilibrium which are of concern to him and to us:

(1) Corresponding to S_A^* and S_B^* , there will be a pattern of overt demands and manipulative actions r_A^* and r_B^* . The demand components of these vectors are given by X_A^* and X_B^* where the final elements in these vectors, $x_{A_1}^*$ and $x_{B_1}^*$, represent settlement payoffs. These settlement payoffs are the only payoff elements which enter into the utility functions. The others, representing demands which are not met, are irrelevant, and in fact, so long as they do not fall so low as to bring about an inadvertent agreement, they may take on practically any values so far as the payoff utilities of the bargainers are concerned.

2. The strategies S are essentially lists of vectors. S_A , for example, has lists of dates for the use of manipulative moves t_1, \dots, t_m where for every $i, 0 \leq t_i \leq K$, and lists of payoff demand sequences X_A , where each element in a sequence is positive but limited by the conditions of the game. The strategy space is therefore closed, bounded, and convex. If the utility functions are everywhere continuous in all of the elements of S_A and S_B , the possibility of randomized choices within the strategies makes the dependence of Arthur's choice on Bill's strategy selection a continuous relation (and vice versa). The Kakutani Fixed Point theorem may then be used to guarantee a solution.

This does not mean that they are entirely inessential, because they may play an important part in the determination of the equilibrium strategy pair. It does mean, however, that payoff demands made early in a negotiation may bear no direct relationship to the parties' equilibrium-payoff expectations.

(2) The equilibrium is generally not Pareto optimal. Because the manipulative actions require time to be put into place and to have their effects, the settlement date, t^* , is normally substantially greater than zero. This is only to say that the negotiation takes time. However, since manipulative activities are costly in themselves, and delays in agreement always reduce the value of a settlement, values of t^* greater than zero represent a cost, and both parties would be better off if the manipulations never took place and if the settlement payoffs $x_{1t^*}^*$ and $x_{2t^*}^*$ were received immediately at day 0.

Some game theorists have argued that in a full information, cooperative situation, it would be unreasonable to allow for any outcome which is not Pareto optimal. For example, when Nash (1953) introduced the possibility of force and other manipulative actions into his bargaining model, he defined all such actions to be "threats" which influence the outcome, but which need never be carried out, because fully informed negotiators can take them into account without having to bear the cost of actually using them. To follow such a course in practice, however, is clearly beyond the abilities of most negotiators. It would require not only that they have the analytical capacity to solve for the pair S_1^* , S_2^* , but that they be able to implement the implied settlement agreement without actually experiencing many of the events on which that equilibrium depends. On the other hand, negotiators with long experience with one another can be expected to have learned a great deal about one another's bargaining strategies and to have used this knowledge in planning their own behaviors. In recurrent bargaining situations such as the two- or three-year cycles common to labor negotiations, the parties may eventually become so familiar with the situation and with one another that accurate prediction of the final settlement point becomes a realistic possibility. This is just the argument that was described as Boulwareism earlier in this paper. Boulware was arguing, in effect, that the equilibrium pair S_1^* , S_2^* was known to both General Electric and the Union, and that both parties would be better off if they acknowledged that fact and saved themselves the trouble and expense of acting out the game. Of course, the problem in this case was that the union did not see this as a proposal to restore

Pareto optimality, but as a new manipulative move in a more subtle bargaining strategy—and they may have been right.

(3) The equilibrium pair S_A^* , S_B^* is generally not unique. We have already mentioned the fact that from the point of view of the parties' payoff utilities, demands made before the settlement date are irrelevant, and it is likely that a variety of different demand patterns could ultimately lead to the same settlement. Of greater importance is the fact that there may be many different equilibrium strategy pairs which will lead to different settlement points. In the extreme case of a situation in which no manipulative actions are permitted, every payoff pair on the frontier of the utility-possibility set can be shown to be the outcome of a possible equilibrium strategy pair.³ In effect, the game-theoretic model has failed to identify a solution. In practice, the presence of possible manipulative actions may reduce this range of indeterminacy by giving each party some defense against very poor payoffs, but the likelihood of some indeterminacy still remains. One possible way of dealing with this indeterminacy in the theory is to impose more restrictions on the fundamental game-theoretic structure. This is the purpose of arbitration schemes such as those which have been proposed by Nash (1950) and others.⁴ Rather than following this procedure, our theory here will be that negotiators themselves deal with the indeterminacy simply by acting out the process, letting the negotiation come out wherever it will.

IMPERFECT KNOWLEDGE AND LEARNING

We cannot assume that actual bargainers correspond to the idealizations found in the theory of games, but we have argued that they will use such idealizations as models for their own behavior. They seek bargaining strategies which will maximize the expected return to themselves. They recognize that time delay and intervening acts of coercion

3. For example, if settlement occurs whenever $x_A + x_B = Q$, then Arthur may decide to demand some x_A^* at every date until K , when he quits. In the face of this strategy, Bill's utility maximizing choice is to accede and demand x_B^* at all dates short of K where $x_B^* = Q - x_A^*$. Given this strategy Arthur's utility maximizing strategy is still to demand x_A^* at all points in time. Thus we have an equilibrium. Such an equilibrium exists for any x_A^* greater than zero and less than Q .

4. See Luce and Raiffa (1957): 121-145.

are inferior to immediate acceptance of the settlement point, although their lack of information as to the nature of that settlement requires that they engage in the process anyway. Finally, they recognize that payoff demands at dates earlier than t^* have very little practical significance so far as the value of the settlement itself is concerned. Instead, presettlement demands may be used as instruments for influencing a competitor's bargaining behavior. (The effectiveness of such instruments will be considered later.)

Suppose that at the start of a negotiation, Arthur has developed some estimate of what Bill's strategy is going to be. We will call this estimate R_B . R_B is subject to some uncertainty depending on the extent of Arthur's experience in negotiating with Bill, and we will represent this uncertainty with an index V_B . V_B may in fact be composed of a vector of terms (e.g., standard deviations) which describe possible errors in each of the dimensions of R_B , but for our purposes it is sufficient to treat it as a single-dimensioned variable which approaches zero as uncertainty is reduced. Given R_B and V_B , Arthur chooses a strategy, S_A , which maximizes the expected benefit which he will receive from the negotiation. Since here we wish to focus on the influence of the variable R_B , we will assume that Arthur does possess the ability to calculate an optimal S_A to correspond to any particular R_B , V_B combination. In a similar fashion, Bill forms an estimate of Arthur's plan, R_A , together with an uncertainty variable V_A , and then selects a strategy S_B which maximizes his own expected utility. Of course, if it should happen that the estimates are correct so that $R_A = S_A$ and $R_B = S_B$, then the parties' strategy choices would already be near equilibrium, and the negotiation would proceed, mechanically, as planned. The two parties might even find some means for short-circuiting the process and jump to an immediate agreement.

We are concerned primarily with cases in which the estimates R_A and R_B are not perfectly accurate. If there are errors in these expectations, this will be discovered over the course of the negotiation, and the parties will feel compelled to revise them. Revisions in R_A and R_B and in V_A and V_B will lead in turn to revisions in strategy choices, and a series of adjustments and readjustments will begin to occur. Naturally, the explicit concessions and overt attempts at outcome manipulation which occur during the course of a negotiation do not reflect only these changes in expectations. The strategies themselves incorporate demand changes and manipulative actions which are dated, and what we observe as the dynamic progress of bargaining may be a mixture of




Figure 3

these planned stages in unchanged strategies and of changes in the strategy choices themselves.

An example of the process which we have in mind is given in Figure 3. Suppose the two parties are to divide a homogenous good which is available in a fixed quantity Q . Arthur has made an estimate, R_B , of Bill's strategy throughout the negotiation and has devised an optimal counterstrategy, S_A , in response. Given S_A and R_B , Arthur expects an agreement at time t_A^e and a settlement share x_A^e . Suppose that R_B is too optimistic, and that Bill's overt demands are following a pattern r_B which is incompatible with that expectation. Eventually, Arthur will have to modify R_B to be more consistent with r_B . On the diagram, this occurs at a time t_1 , when R_B is changed to R'_B . Arthur now must choose a new optimal counterstrategy S'_A where S'_A maximizes $U_A(S_A, R'_B, V_B)$, and where U_A is defined at the point in time t_1 .⁵ Since this learning and

5. Note that the original strategy was chosen at a date $t = 0$ and that S'_A is chosen at a different time t_1 . If the utility function is influenced by the calendar date, this fact must be taken into account. It is also important to consider the possibility that with R_B unchanged, the changing calendar date might lead to a change in S_A . Normally, utility functions are constructed so that this will not occur. This general problem is discussed in Strotz (1956).

adjustment mechanism has taken time to operate, it may be impossible for Arthur to implement S'_A fully because some elements in this strategy would have required application at dates earlier than that at which S'_A was decided on (this possibility is one reason why the presence of uncertainty, V_B , has an influence over Arthur's original strategy choice S_A .) The best Arthur can do now is make a utility-maximizing transition from S_A to S'_A . Arthur's observable bargaining behavior, r_A , is then obtained from S_A , S'_A and the transition between them. On the diagram, $(t_A)'$ and $(X_A)'$ are Arthur's new expected date of settlement and settlement share, respectively.

For the sake of clarity, two important details have been omitted from Figure 3. First, we have represented only the actual payoff demands, and have left off the various manipulative activities which are a part of S_A and r_B . Second, we have mentioned the possibility that explicit payoff demands are routinely reinterpreted by experienced negotiators to take account of various conventional end-game procedures. If it is normal practice to conclude a negotiation by splitting remaining differences or trading-off unresolved issues, then stated demands will always differ from true demands, but in ways which are readily taken into account. The diagram in Figure 3 abstracts from this possibility and is drawn to represent what each party would understand to be the settlement offers.

We intend that the parties in this theory see themselves in symmetric roles. That is, if Arthur finds that his own bargaining behavior is determined in part by adjustments in his own expectations, then he will attribute that same kind of learning to Bill, and conclude that r_B is a reflection both of S_B and of changes in S_B which have occurred as a consequence of adjustments in Bill's expectations. If we may paraphrase the bargaining attitude we mean to represent, Arthur might argue to himself:

- (1) "Bill is too optimistic: he have to learn that S_A is not so favorable to him as he thought. As he does learn, this will be reflected in r_B , and I will be able to get a better settlement than I could now."
- (2) "Even though Bill is learning, it is not going so fast as I expected, and I will have to revise my own expectations downward."

Thus our negotiator is attributing the *same* kind of learning behavior both to himself and to his competitor. There is no suggestion that he expects to behave one way while his competitor behaves another.

Given the awareness which our theory attributes to the bargainers, it is natural to expect them to attempt to manipulate one another's

expectations. However, that same awareness would lead them to expect such attempts from their competitors and to discount behavior which might be so interpreted. In my view, the sequence of moves and counter-moves or bluffs and counterbluffs which this might produce will have as its major effect an increase in the values of the uncertainty variables V_A and V_B . Increases in V_A and V_B will in turn have two consequences. First, as already noted, strategy choices are affected by uncertainty: large amounts of uncertainty may, for example, encourage very large, initial payoff demands as a kind of insurance against making an unnecessarily generous offer. Thus, bluffing may induce counterbluffing. Second, by reducing the confidence which one party may place in his own estimate of his competitor's behavior, large values of V_A and V_B will slow the learning process and increase the persistence of initial expectations. We have already discussed the fact that payoff demands which are made before the settlement date have no significance apart from their impact on the dynamic progress of the bargaining. If Arthur plans to settle at time t_A^c for payoff x_A^c , a very large preliminary demand may still be tried, for purely strategic reasons, without endangering that settlement. However, he will discount, quite properly, all such demands which might be put forward by Bill, recognizing them to be similarly costless bluffs. Thus, the information value of these early demands will be seen to be negligible, and the learning process will proceed slowly at best.

As the anticipated settlement dates t_A^c and t_B^c approach, uncertainty will decline and payoff demands will become more reliable indicators of expectations. In the extreme case, the date t_A^c might actually arrive before Arthur learns of his overoptimism. At t_A^c , however, Arthur will try for his expected payoff x_A^c , and when he is rebuffed, then it will be obvious to him that R_B was in error, and adjustments in it will take place. Furthermore, in trying to obtain x_A^c , Arthur will make an unambiguous statement of his own demands, and so long as Bill recognizes it as such, he can use this information to improve his own estimate R_A , whether or not the date t_B^c has arrived. We expect this pattern of bluff and attempted manipulation to be a typical one: early in a negotiation, a great deal of strategic maneuvering may occur, uncertainty will be high, and as a consequence, the parties will be reluctant to draw inferences concerning the future course of the bargaining. As anticipated settlement dates approach, however, the parties will get down to business, and the information flows will become much more reliable.

THE BARGAINING PROCESS

The model, as we have outlined it, is very general and would require several further restrictions before any explicit solutions could be obtained. We would have to specify the forms of the utility functions, the dependence of strategy choices on uncertainty, the nature of the learning process, and the rules which determine the availability of various manipulative actions. Nevertheless, even from this general model, we may draw some conclusions regarding the relationship between payoff demands and expectations, and hence about the character of the dynamic course of the bargaining.

Suppose that we describe as the general case of bargaining a situation in which each party entertains overoptimistic expectations. (If expectations are realistic in the sense that $R_A = S_A$ and $R_B = S_B$, the negotiation is reduced to a charade, and if one party is extremely pessimistic, the negotiation may reach a settlement before the learning mechanism has any effect at all.) In this case, the diagram in Figure 3 and a similar one drawn from the point of view of Bill can be used to describe the course of negotiation.

The learning mechanism establishes a dynamic interaction between the two parties' behaviors. Arthur's strategy S_A determines the current course of r_A , which is used by Bill in the formation of R_A . In response to R_A , Bill selects strategy S_B which determines the course of r_B , and this in turn is the basis for Arthur's estimate R_B . It is important to bear in mind that as a consequence of learning, S_A and R_B may be continually changing, so that what is learned in the form of R_A or R_B is a composite of a sequence of strategies rather than a single one. From either negotiator's point of view, there is no way to distinguish changing strategies from the mechanisms of an invariant one. Arthur may observe r_B , but he cannot know whether r_B is the outgrowth of learning (he hopes that it is), or the manifestations of some subtle strategy of Bill's.

The choice of bargaining strategy S_A is motivated by a desire to maximize $U_A(S_A, S_B)$. This utility value is not dependent solely on the settlement payoff, however, for it also reflects the costs of time delay before a settlement is reached, and the losses which may be suffered because of various attempts to use force or coercion. The selection S_A reflects a trade-off between these opposing values. If the cost of time delay is small, S_A will reflect a willingness to wait in the interests of achieving a large settlement; if the cost is large, S_A will reflect a willingness to sacrifice some payoff for the sake of early agreement. Thus

if Arthur views his prospects optimistically (Bill is seen to be conceding rapidly, or to be making small demands), the anticipated time cost of negotiating is reduced (because the delay will be shorter) and S_A will be designed to achieve a larger payoff at settlement time. If R_B is seen to include the use of force, even if the date for that force has not yet arrived, s_A will be designed to reduce the cost of that force, partly through the use of countermeasures, but also partly through reductions in the final payoff expectations.

Even in this most general model, a number of conclusions may be drawn immediately.

- (1) The negotiation will take more time than either party initially expected.
- (2) High time costs, or vulnerability to coercion will reduce a party's settlement payoff.
- (3) The use of force or coercion which increases the cost of delays in agreement will reduce the duration of the negotiation.
- (4) If a party's learning rate is high, for whatever reason, the duration of the negotiation will be reduced.
- (5) If a party's learning rate is high, for whatever reason, that party will receive a smaller payoff at settlement time than he would otherwise.
- (6) There is a central tendency in the bargaining process which reduces any asymmetries in the parties' expectations which are not reflections of differences in learning rates, utility functions, or in the availability of means for altering the outcome through the use of force.

Some of these six conclusions are not very startling. Conclusion 1 is obtained directly from our discussion of Figure 3 and is no more than a reflection of our description of the general case of bargaining as one in which both parties are overoptimistic. Conclusions 2 and 3 are equally straightforward and are obtained from our description of the utility function. If, due to high bargaining costs, a party lowers his payoff expectations whatever the value of R_B , then, even with the learning process, r_A as drawn in Figure 3 will be lower at every point in time, and Arthur's final payoff will be reduced. Conclusion 4 is an obvious consequence of the fact that a high learning rate will lead to a more rapid downward shift in expectations, which will lead in turn to a more rapid downward fall in r_A . Conclusion 5 is obtained from the same observation. If Arthur finds that his estimate of Bill's behavior was overoptimistic, and if he reacts strongly to this discovery, r_A will reflect a substantial decline in expectations. If, on the other hand, Bill is slow to react to similar information, r_B will not include such a decline, and the agreement will occur after Arthur has made most of the concessions.




Figure 4

The central tendency described in Conclusion 6 is an important property in that it reduces the dependence of the outcome on the unspecified initial expectations of the two parties. It comes about because of interactions between the two learning mechanisms. Suppose that the two negotiators are approximately equal in learning ability, but that Arthur enters the negotiation with expectations which are substantially more optimistic than those of Bill. Compared to Bill, Arthur has more to learn so to speak, and as time passes, Arthur will have to make several adjustments in S_A . These same adjustments in S_A , however, will be reflected in an r_A which is quite favorable from Bill's point of view and which will therefore slow adjustments in S_B . Arthur's optimism and the discovery of it force him to modify his behavior, and these modifications in turn interfere with changes in Bill's expectations, even though these, also, may be too optimistic. Thus the degrees of optimism of the two parties are inclined to converge.

An example of this process is shown in Figure 4. For the sake of a simple case, let us represent the parties' expectations with only the two variables x_A^e and x_B^e . Rates of change in expectations are represented by \dot{x}_A^e and \dot{x}_B^e . According to our proposition that learning is responsible for these changes, \dot{x}_A^e and \dot{x}_B^e are interdependent. Bill's expectations affect Bill's own bargaining behavior, but this behavior in turn determines Arthur's expectations. We can write this relationship as $\dot{x}_A^e =$

$F(\dot{x}_B^c)$, and without being too specific as to the precise form of this function, recognize that \dot{x}_A^c and \dot{x}_B^c vary inversely with one another. The more rapidly Bill learns, the larger is \dot{x}_B^c , hence the more justified is Arthur's optimism, and the smaller is \dot{x}_A^c . Similarly, \dot{x}_B^c is dependent on Arthur's bargaining behavior and hence on \dot{x}_A^c , and we will write this relation as $\dot{x}_B^c = G(\dot{x}_A^c)$. These two functions are represented in Figure 4, where $-\dot{x}_A^c$ and $-\dot{x}_B^c$ have been put on the axes (with the minus signs, the axes may be thought of as representing rates of concession). Point C on the Figure represents an equilibrium in that the two concession rates are mutually consistent, and other pairs of \dot{x}_A^c and \dot{x}_B^c will tend to converge on C (at least in stable cases). Beginning at any point such as D, the sequence $\dot{x}_A^c = F(\dot{x}_B^c)$, $\dot{x}_B^c = G(\dot{x}_A^c)$, . . . follows the dashed path to C.

Except for the restrictions used to draw Figure 4, we have described a theory rather than a model, and in order to provide explicit solutions to particular bargaining situations we would have to provide specific forms for the relations which are involved. In fact, a great many models (and hence a great many empirical possibilities) can be represented along the lines of the theory which has been described. One such example is my own earlier bargaining model (Cross, 1965) in which each party's strategy was simply to demand the expected settlement payoff at each point in time ($x_{At} = \dot{x}_A^c$ for all $t \leq t_A^c$). In this model, R_A and R_B were estimates of \dot{x}_A^c and \dot{x}_B^c respectively, and the learning function made changes in these estimates' linear functions of the differences ($R_A - \dot{x}_A^c$) and ($R_B - \dot{x}_B^c$). This particular formulation was quite properly criticized for the restrictiveness of these assumptions, particularly for the unrealistic specification of the strategy choices. In the general theory, it is always in the interest of a negotiator to conceal changes in his own expectations (because such changes are encouraging to the other party and slow his learning), and a strategy choice known to have the form $x_{At} = x_A^c$ for all $t \leq t_A^c$ maximizes the information content of every stated demand. Coddington (1966) has provided a model at another extreme in which each party's strategy is to demand some arbitrarily large payoff until the expected settlement date at which time the true settlement offer is made. Strategies of somewhat more realism than either of these can be introduced without great difficulty. For example, Arthur may plan to demand $x_{At} = x_A^c + P(t_A^c - t)$, where P is a positive constant, so that some concessions can occur whether x_A^c is changing or not. Such a strategy combines an appearance of cooperative negotiation with a useful element of concealment. Perhaps the

most realistic representation of a strategy would be to make demands which can be traded away for concessions from the other party, making explicit the possibility that the parties' strategy choices contain conditional elements. Unfortunately, the mathematics of such a model would be far more difficult than those associated with our more naive formulations.

SYMMETRY

Central to our theory is the proposition that negotiators are similar in their understanding and insight into the bargaining process and into one another's behavior. If Arthur finds that his expectations are too optimistic, he may at least hope that Bill is also having to revise his own. If Arthur uses some strategic bluffs in an attempt to influence Bill's expectations, he expects some similar attempts from Bill. Under these circumstances, as we have already argued, bluffing activities may not have a significant effect on the outcome. Inflated demands will be expected and discounted, and their most important consequence will be a reduction in the parties' rates of learning rather than alteration of the ultimate settlement.

In fact, a bargaining pattern may develop in which each party makes initially inflated demands himself, disregards the demands of his competitor, and then gradually relinquishes these exaggerations, expecting his competitor to do the same. This may become such an established routine that it would actually be dangerous to do anything else. The concessions which are made possible by the initial exaggerations may become an unwritten rule of the game, and any party who abandons this course may then be seen to be negotiating in bad faith with the result that the negotiation is seriously disrupted. Earlier in this paper, we suggested this to be a possible explanation for the hostility which Boulware's approach apparently engendered in negotiations between General Electric and its employees.

Apart from these unwritten laws and the occasional need to impress constituents with a public display of bargaining prowess, there is reason to believe that misrepresentation of one's expectations is not always in one's best interest. If one wishes an overoptimistic competitor to learn of his error, the best procedure is to provide him with good evidence that he is wrong, and a reputation for bluff is not conducive to this end. Indeed, if we were to attempt to devise a model of optimal

bluffing, we would quickly be driven to a logical dilemma. Suppose that Arthur recognizes that through the learning mechanism, Bill's behavior is influenced by Arthur's own bargaining tactics because R_A is affected by r_A and r_A is in part a consequence of S_A . Arthur represents his estimate of the nature of this dependence with a function $R_B = f(S_A)$ where he may let the form of the function $f(\)$ incorporate the uncertainty variables. Arthur's optimal bargaining strategy is now one which maximizes $U_A(S_A, f(S_A))$. By the symmetry of our model, however, Bill will devise a function which represents the dependence of R_A on Bill's own behavior, $R_A = g(S_B)$, and choose a strategy, S_B , which maximizes $U_B(g(S_B), S_B)$. This behavior, however, violates the premises on which the function $f(\)$ was defined. In fact, the functions $f(\)$ and $g(\)$ are inconsistent at any level of analysis. Since the function $g(\)$ depends on Arthur's choice of behavior, the function $f(\)$ must incorporate this dependence. By symmetry, $g(\)$ incorporates a similar dependence, and $f(\)$, again, is misspecified. We are confronted by an infinite regress.

We could, of course, escape by abandoning the assumption of symmetry between the parties. If Arthur recognizes the dependence of Bill's expectations on his own strategy, but this insight is not shared by Bill, then Arthur has a distinct advantage. His bluffs work, and the settlement moves to his benefit. There is room here to distinguish a more skillful from a less skillful negotiator. The theory is not otherwise altered, however. The function $f(\)$ is an estimate which may be too optimistic, for example, in that Arthur may expect his bluffs to work better than they do, and Arthur may have to learn to revise his expectations downward in the usual fashion. Moreover, as Bill becomes more experienced, Arthur's bluffs will become ineffective and the operation of the process will become symmetric again.

A FINAL NOTE

The theory which has been outlined here reflect this author's conviction that expectations and learning play a central role in bargaining. The introduction of expectations has often been criticized, however, on the grounds that it reduces the empirical usefulness of the theory by employing variables which are not directly observable by any third party. Of course, any theory may be better than none, and we might further argue that since virtually all bargaining models already extant

rely heavily on utility functions, the empirical testability of the theory is compromised very little by the addition of an expectations function. Nevertheless, empirical usefulness is the objective of any theory, and it is worthwhile to extend our original list of six implications of the theory into a few practical areas.

We begin with the most obvious:

- (7) Parties with access to potent threats will do well (in terms of settlement payoff), as will those who are relatively unconcerned about the timing of an agreement.

Such implications as these are already widely accepted, apparently simply because they appeal strongly to the intuition. They explain such beliefs as: large strike funds will operate to a union's advantage in wage negotiations or that a large stockpile of finished products will benefit an employer.

Less obvious are conclusions which apply to the dynamics of the bargaining process.

- (8) Very large initial payoff demands may worsen a party's settlement payoff.

This implication is drawn from the observation that large initial demands will require, at some time or other, large apparent concessions which will slow the other party's learning.

- (9) Political negotiations will take more time than economic negotiations.

This may sometimes be the case simply because the passage of time is of less concern to the parties, but it may be expected also because political issues are less well defined than are economic issues, and the presence of easily quantifiable variables contributes to rapid learning.

- (10) Multidimensional negotiations may actually come to settlement earlier than unidimensional ones.

In this paper, we have not been able to examine the searching component of multidimensional negotiations. We would like to note, however, that if issues in a negotiation are settled in sequence, one exchange being initialled before the next is seriously considered, then the flow of information between the negotiators is greatly increased. It becomes possible for each party to gauge the expecta-

tions, values, and intransigence of the other with much more confidence than is possible if the only available information is to be found in early and undoubtedly exaggerated payoff demands. Adjustments in expectations are correspondingly accelerated and settlement may arrive sooner. Unfortunately, we cannot assert this conclusion unambiguously because a large number of issues naturally require more time just for their definition. On the other hand, it is possible that our case could be put even more strongly: the persistence of bluffing and distrust in a unidimensional negotiation may block information flows so completely that the learning process ceases to operate entirely thus forcing the parties to turn to outside assistance through mediation (which is designed to improve communication in a negotiation) or arbitration (which is designed to determine a settlement).

- (11) Even in the presence of an equilibrating tendency in the learning process, initial expectations will be reflected to some degree in the final settlement, and these initial expectations are themselves determined by observable variables.

If profits are high, union members are likely to look to forthcoming wage negotiations with optimism, and the settlement will reflect that optimism. For this reason, employers are often concerned not to appear too prosperous when the current contract expires. If the United States makes large trade concessions to the Common Market, the Soviet Union will be inclined to view upcoming grain deals with more optimism, and this will operate to their benefit in the final settlement, whether or not the United States intends to be generous. In short, the settlement of a negotiation will move to the benefit of party A if similar or parallel negotiations are commonly seen to have been favorably settled by parties in the same position.

Some empirical testing of these conclusions has already taken place. Siegel and Fouraker (1960) have provided experimental support for the conclusions that high time costs accelerate the bargaining process and that optimistic expectations (or aspirations) will delay a settlement, and they have also found the expected tendency for a party with the most reliable basis for learning to lose somewhat in terms of settlement payoff. These results have been reproduced by Contini (1968) in a similar set of experiments. Outside the arena of controlled experimentation, Ashenfelter and Johnson (1969) have investigated the expectations hypothesis in the form of a suggestion that factors which might generally increase labor union expectations will delay wage

settlements. Their results conform to our prediction and Siegel and Fouraker's experimental conclusion on this point, but the study is weakened by the fact that no use is made of any formal bargaining theory beyond a naive adaptation of the Hicks (1932) model, so that one is unable to draw any inferences regarding the settlement point itself. Moreover, these authors use the occurrence of strikes as their measure of time delay without any acknowledgement that as manipulative actions, strikes may be determined by other forces as well. Nevertheless, this study is an encouraging indication that the use of expectations in the theory need not take it outside the realm of applicability, and it is to be hoped that some of the other implications of bargaining theory will soon be subjected to similar testing.

John G. Cross is Professor of Economics at the University of Michigan and an Associate Research Economist at the Mental Health Research Institute. He is the author of a book on The Economics of Bargaining and of several articles which apply learning processes to the theory of economic behavior in a number of areas.