Discussion Paper No. 666

Negotiation and Arbitration:
A Game-Theoretic Perspective

by

Robert J. Weber

November 1985

Negotiation and Arbitration: A Game-Theoretic Perspective

by

Robert J. Weber
J.L. Kellogg Graduate School of Management
Northwestern University

1. Introduction

In the past few years, a number of texts have been written concerning the "art" of negotiation. The authors of these texts, with little exception, have relied on their personal experiences, and on historical precedent, to derive lists of maxims and "rules" for dealing with others in situations of conflict wherein mutual gain is available through cooperative action. While advice of this kind can be useful, few of the maxims have been given formal (i.e., non-experiential) justification, and there has been little discussion of the domain of applicability of the rules. A purpose of this paper is to provide a linkage between these popular treatises and recent game-theoretic research on bargaining and related issues.

On one level, the game-theoretic perspective involves formal quantitative models of negotiations, and much of the detailed analysis cannot directly be put into practice due to the difficulty in accurately estimating the preferences and beliefs of the involved parties. Fortunately, on another level one can gain much qualitative, applicable insight into real-world problems from an understanding of the principles involved in the formal analysis, and from a study of the various phenomena which arise in simple examples chosen to emphasize different components of more complex real problems. The organization of this paper was chosen to emphasize these principles and phenomena.

The focus of our approach is on interactions involving "rational" parties, who act as expected utility maximizers, who correctly perceive the structure of the "game" they are playing, and who deal with uncertainty according to the laws of probability. Experimentalists in the behavioral sciences have repeatedly shown that very few of their subjects perfectly satisfy these assumptions. Yet violations of the assumptions tend to be consistent in direction across individuals (cf. the Allais paradox, insensitivities to small probabilities, anchoring in the updating of beliefs, etc.). Hence, the study of fictitious rational parties provides both a norm against which actual behavior can be compared, and also a guide to participants in negotiations, as well as third-party intervenors, about potential difficulties which can be anticipated, and at times avoided. Chapter 2 of this paper presents the concept of an "equilibrium" pairing of strategies in a game, together with a discussion of the reason for giving this concept a central focus in the remaining chapters. Chapter 4 develops a precise definition of "strategy" in settings where a party holds private information (about, for example, his own preferences).

Research into bargaining can be classified along two dimensions. One dimension distinguishes problems of "complete information," in which both parties are fully informed concerning the possible outcomes and one another's preferences over those outcomes, from those of "incomplete information," in which some aspects of the situation, typically the individual preferences, are not commonly known to the two parties. The latter case is the one usually faced in actual bargaining; the analysis of the former, simpler case serves to provide perspectives for generalization, and to delineate those issues in the latter which arise purely from the incompleteness of each party's information. One major difference is seen in the avoidability of conflict in the first case, and the inevitability of conflict (at times) in the second. Chapter 3 of this paper deals with negotiations in settings of complete information, and Chapter 4 extends the analysis to settings of incomplete information.

The other dimension distinguishes between studies of the actual mechanics of negotiation, wherein the parties exchange information through discussion of the issues and the making of offers and counter-offers, and the direct study and classification of agreements which could be reached through _some_ set of mechanics. The central part of Chapter 4 deals with the characterization of the range of attainable agreements, while Chapter 5 examines several models which give specific regard to the role of time in the interaction between parties.

In discussing the models and issues which have arisen in research on bargaining, we will also present a perspective on the different roles played by third-party intervenors in conflict situations. While we acknowledge that words such as "mediator" and "arbitrator" carry multi-role connotations in common parlance, we will try to clarify the various roles by providing new, restrictive definitions for mediation, arbitration, regulation, and auditing. Discussion of these roles is interspersed throughout Chapters 3 and 4.

2. _Bargaining, viewed as a noncooperative game_

"Game theory" has traditionally divided its objects of study into "cooperative" and "noncooperative" games. The study of cooperative games begins with a specification of the possible agreements available to two or more parties in settings where their interests conflict, and focuses on the selection of an agreement (or set of agreements) with desired properties (e.g., equity, or stability).

In contrast, the study of noncooperative games focuses directly on the problem of individual strategic choice. The principal objects of investigation are the "equilibrium points" of a game: An equilibrium point is a collection of strategies, one for each player in the game, with the property that each player's specified strategy is optimal for him, given that the other players follow their specified strategies.

Why this focus on equilibria? A "game" can be loosely defined as a situation in which the final outcome for a participant depends not only on his own actions, but also on the actions of others. In order to appropriately choose his own action, the participant must formulate a belief about how the others will act: Presumably, he will then choose his own action as an optimal response to the anticipated actions of the others. If he believes the others to be rational, he must assume that they are going through the same process, that is, that they are formulating beliefs about his action, and choosing their own actions to respond optimally. One of two cases must hold: Either each party correctly formulates his beliefs, in which case the chosen actions form an equilibrium point of the game, or someone errs. Even in this latter case, the equilibria of the game provide standards to which the nature of the error can be compared.

This argument can be put more bluntly - If you are playing a game, and choose to employ a strategy which is not a component of some equilibrium point, then either (1) you are not acting optimally, given your expectations about your opponent's choice of strategy (i.e., _you_ are acting foolishly), or (2) you are expecting your opponent either (a) to act non-optimally, given his expectations about your behavior, or (b) to form his expectations incorrectly (i.e., you are expecting _him_ to act foolishly). Of course, people _do_ act foolishly at times, and both psychological and decision-analytic research on negotiations has been directed towards an understanding of foolish actions and a development of prescriptive rules for the anticipation and exploitation of an opponent's foolishness. But only the game-theoretic perspective provides a view of the "rational" norm, in terms of which foolish actions can be defined, analyzed, and interpreted.

The adjective "noncooperative" is not to be confused with "competitive": In a noncooperative game, mutual gain is frequently available through coordinated actions. A noncooperative game is simply a game in which the strategies of the players are given explicit regard, and in which binding agreements between the players are not permitted. (Later, we will introduce the notion of a "regulator," an intervenor who is given the power to exact penalties upon violators of an agreement. But even in the presence of a regulator, the parties retain full freedom of strategic choice: It is the existence of the penalties which enforces the agreement, by making violation of the agreement more costly than adherence to it.) The various actions available to the parties in the course of negotiations can be explicitly incorporated into the rules of a game; thus, the choice of what to say, and when to say it, becomes a strategic choice, and a negotiation problem, which has a cooperative flavor, can be studied as a noncooperative game.

We give three examples of noncooperative games and their equilibria: The first is trivial, and the latter two are well-known.

Example 1. Two acquaintances are discussing their plans for the evening. Each wishes to go to the opera. If they both attend, they will not only enjoy the opera itself, but also will enjoy one another's company.

Each has available two "strategies": to go, or not to go. For each, "going" is a dominant strategy, preferred to the other strategy no matter what the other individual chooses to do. The only equilibrium pairing of strategies is ("go", "go") (where, by convention, we label one of the two individuals "Player 1", and write his strategy first); clearly, this is the choice of strategies we would expect to observe. Notice that there is no conflict of interest in this problem: There are no two different pairings of strategies, with one pair preferred by one party, and the other pair preferred by the other party.

Example 2 (The Prisoners' Dilemma). Two men have been arrested for a minor offense. However, the district attorney is certain (although he has no hard evidence) that they are also responsible for a much more serious crime. He separates the criminals, and offers each the same deal: If neither confesses to the more serious crime (such a confession would implicate both), he will ask for two-year sentences on the lesser offense. If both confess, he will request five-year sentences. But, if only one confesses, that one will go free, and the maximum penalty of the law (an eight-year sentence) will be requested for the other.

In this example, again, each has a dominant strategy: to confess. And again, the strategy pair ("confess", "confess") is the unique equilibrium point of the game. At this equilibrium point, both are worse off than at the outcome of the strategy pair ("don't confess", "don't confess"). Yet, in the absence of any external, enforceable agreement, it must be expected that each will confess.

<u>Example 3 (The Battle of the Sexes).</u> A man (A) and a woman (B) must choose where to spend the weekend. Each can go either to the mountains or to the beach. Each derives pleasure from the other's company, and also from being at his or her more-favored location. However, the man favors the mountains, and the woman, the beach. The payoff matrix below indicates the utility payoffs to each, depending on the ultimate destination chosen by each.

		B	
		mountains	beach
A	mountains	$t_A + d_A, t_B$	d_A, d_B
	beach	$0, 0$	$t_A, t_B + d_B$

(t represents the utility premium for "togetherness", d for "most-favored destination.") In order to focus on the most interesting case, we assume that the togetherness premium for each is greater than his or her destination premium.

In this example, there are three equilibrium points. (1) Both go to the mountains. (2) Both go to the beach. (3) A makes a random decision, going to the mountains with probability $1/2(1 + d_B/t_B)$: B goes to the mountains with probability $1/2(1 - d_A/t_A)$. This last, "mixed-strategy" equilibrium point is inferior for both parties to either of the first two equilibria. But the selection between the first two remains to be decided.

One possibility would be for the two to agree to a coin toss, to select a joint destination. But this merely confounds the problem: How should the coin be weighted? (Obviously, weightings other than 50:50 are possible.) Now, instead of bargaining over the choice of destination, they must choose from among an infinite number of possible weightings.

3. The Nash bargaining model

The first formal game-theoretic analysis of bargaining was presented by John Nash in the early 1950's. He considered situations in which two individuals must choose how to coordinate their actions to mutual advantage, when each is fully aware of the set of potential agreements, and of the preferences of the other over those agreements.

Each party has available a list of actions; any chosen pair of actions yields an outcome. Furthermore, the preferences of each over the possible outcomes, as well as over probabilistic mixtures of outcomes, are commonly-known, and satisfy the standard von Neumann-Morgenstern axioms, i.e., both individuals are expected utility maximizers.

Nash began his analysis by assuming that the bargaining problem under investigation had some pre-specified "conflict" outcome, which would occur in the absence of agreement. For example, in the Battle of the Sexes the natural conflict outcome is for each to go to his or her more-favored destination.

Example 4. Alfred, who is near-broke, possesses a $100 bill. The serial number of the bill, mmddyyyy, happens to be the birthdate of Burton, a wealthy eccentric. Burton wishes to acquire the bill as a keepsake, and would be willing to pay as much as $500 for it. Alfred's utility for money is proportional to the square-root of the amount he holds (i.e., he is risk-averse); Burton's utility for money is linear (i.e., he is risk-neutral). How much should Burton pay Alfred for the bill?

In this example, again, the conflict outcome appears obvious: Alfred spends the $100 bill.

Nash next noted that many different agreements might be "stable," in the sense that, were an agreement reached and further discussion impossible, both parties would voluntarily carry out their roles in the agreement. In the Battle of the Sexes, if the agreement is to use a specific weighted coin to select randomly a joint outcome, then, even after the coin flip, neither party can gain by unilaterally deviating from the agreement and going elsewhere. In the Alfred-Burton example, an agreement on any price between $100 and $500 is an agreement from which neither gains by walking away.

3.1 Mediation

In other cases, there are mutually beneficial, stable agreements which require help from the outside.

Example 5. Consider the following two-person game, in which each party must choose between two actions; the utility payoffs from the various pairs of selections are indicated in the diagram. (A's payoff is the first in each pair.)

		B: left	B: right
A	top	6,6	2,7
A	bottom	7,2	0,0

There are three equilibrium points in this game: (top, right), (bottom, left), and a mixed-strategy equilibrium in which each independently and randomly chooses an action, with his first action being twice as likely to be chosen as his second. The corresponding expected payoffs are (2,7), (7,2), and $(4\frac{2}{3}, 4\frac{2}{3})$. Through the use of a joint randomizing device which, with some triple of agreed-upon probabilities, chooses one of the three equilibria, they can achieve any expected-payoff pair in the small, unshaded triangular region in Figure 1 as the outcome of a stable agreement. (That is, once the device is agreed upon, neither gains by unilaterally failing to carry out his role in the equilibrium point chosen by the device.)

A mediator could help them create other, mutually-preferred stable agreements. For example, they could agree that the mediator would leave the room and flip a fair coin. If it comes up "heads," he will return and privately whisper "top" to A, and "left" to B. If it comes up tails, he will re-flip it: On "heads," he will return and whisper "top" to A and "right" to B ; on "tails," he will whisper "bottom" to A , "left" to B. If each party knows nothing about the mediator's out-of-the-room actions except for his own whispered message, and if he expects the other to obey the mediator's message, he can do no better than to obey his own. This mediated procedure yields them expected payoffs of $(5\frac{1}{4}, 5\frac{1}{4})$. Other stable agreements obtainable through mediation yield expected payoffs in the lower shaded region of the figure.

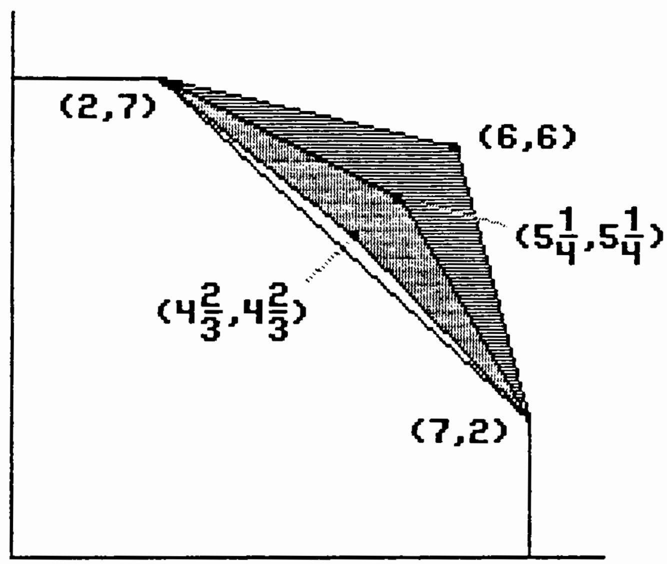

Payoffs attainable through intervention

Figure 1

Generally, our view is that a <u>mediator</u> is an intervenor who can (publicly <u>or</u> privately) receive information from the parties, and transmit information back to them, according to rules upon which the parties themselves have agreed. Under this definition, a mediator can also be delegated the responsibility of carrying out public randomizations, as in the case of a randomized selection of a joint destination in the Battle of the Sexes.

3.2 Regulation

In Example 5, even better payoffs (such as (6,6)) are available to the parties, but no stable agreement allows them to obtain these payoffs. Similarly, in the Prisoners' Dilemma, both would prefer the ("don't confess", "don't confess") outcome to the equilibrium outcome, but neither can expect the other to adhere to an unenforced verbal agreement.

However, if an outside agency can be brought into the situation, and empowered to exact sufficiently high penalties from any deviator, either of these agreements becomes stable. Penalties of 1 or more are sufficient to "stabilize" the agreement (top, left) in Example 5; "Honor among thieves", when backed up by physical retribution, makes a D.A.'s task much more difficult than the Prisoners' Dilemma would suggest.

Formally, we view a regulator as an intervenor to whom the parties may voluntarily grant the ability to force certain actions upon them (through the setting of appropriately-large penalties for noncompliance). In many instances, the civil courts play a regulatory role in enforcing contractual provisions.

The figure for Example 5 illustrates the variety of stable agreements which can be maintained at different levels of communication or intervention. If the parties can only communicate by telephone, they can achieve only the three outcomes corresponding to equilibrium payoffs. If they can meet (to observe the result of joint randomization), they can achieve an expected outcome anywhere within the triangle determined by these three. Through mediation, they can add the lower shaded region to the range of stable agreements. And through regulation, the upper shaded region can also be added.

3.3 Arbitration

Assume that the parties have agreed to use a regulator, in order to expand the set of possible (i.e., stable) agreements to its fullest. There still remains the problem of choosing from among these possible agreements. The parties can do this through open debate, always facing the possibility that they will fail to reach a settlement. Alternatively, they can invite yet another intervenor to enter the picture, and ask him to suggest a particular agreement, on the grounds, for instance, of his perception of "equity." Indeed, if they simultaneously agree to empower a regulator to enforce that suggestion, a settlement is guaranteed. (The regulator-arbitrator combination is what is frequently described as "binding arbitration.")

In order to distinguish between the roles of third-party intervenors in bargaining, we choose to define an arbitrator as an intervenor who is invited to suggest an agreement. Much of the rest of this chapter will focus on the procedures by which an arbitrator might choose his suggestion.

3.4 The Nash solution

Given the multitude of potential agreements, Nash suggested a set of rules (formally, "axioms") which determine a unique suggested agreement for every problem. The rules are stated as conditions an arbitration procedure should satisfy, where by "procedure" we mean a consistent philosophy to be applied across all bargaining problems. First, these rules require that the suggested agreement be feasible, Pareto-efficient (i.e., no alternative feasible agreement should be better for both parties), and individually rational (i.e., the suggested agreement should offer to each party at least as much as he would obtain at the conflict outcome). Second, the suggested settlement should depend on the parties' underlying preferences, and not on the utility functions chosen to represent those preferences. Third, in symmetric situations (that is, situations where the range of feasible agreements is symmetric, and the parties receive equal utility payoffs at the conflict outcome), the suggested agreement should offer equal payoffs to the two parties. Finally (and most controversially), if after an agreement is suggested, it is found that some alternative, unsuggested agreement was in fact not feasible, the original suggestion should still stand (i.e., the procedure should be "independent of irrelevant alternatives").

Nash showed that there is only one agreement-selection procedure which has all of these properties; thus, an arbitrator who accepts these rules as compelling must follow this unique procedure. The procedure selects, in every problem, the agreement which maximizes the product of the parties' respective utility gains from agreement, measured relative to their conflict payoffs.

In the Battle of the Sexes, the selected outcome under the Nash procedure is for the parties to jointly randomize their choice of destination, assigning probability $\frac{1}{2} + \frac{1}{2} \left(\frac{t_B}{d_B} - \frac{t_A}{d_A} \right)$ to the mountains (if this is greater than 1, go to the mountains for certain; if it is less than 0, go to the beach). Notice that the mountains (A's more-favored destination) are selected more frequently when it is B who favors togetherness over destination relatively more than A.

In the Alfred-Burton example, the selected outcome is for Burton to pay Alfred the amount \$x which maximizes $(\sqrt{x} - 10) \cdot (500 - x)$, and hence Burton should pay \$277.78. This is less than the split-the-difference payment of \$300: Alfred's aversion to risk works against him in the arbitrated solution.

3.5 Optimal threats

Having dealt with the question of how to select a final agreement, Nash turned back to the question of how the conflict outcome (which forms, in a sense, the starting point for the arbitrator's considerations) should be identified in situations where the result of disagreement is not obvious (for example, when the parties have available a variety of retributive strategies). He proposed that the parties, *knowing how their dispute will be arbitrated once the conflict point is determined*, simultaneously write down the actions they will take if agreement is not reached, and empower a regulator to force them to carry out these actions in the absence of agreement. Nash then showed that in every case both parties will have *optimal* threatened actions, i.e., threats which leave them optimally positioned for the arbitration stage.

3.6 Summary

In light of the above discussion, we can interpret Nash's approach as separating negotiations into two stages: a threat-making stage, which is strictly competitive (in the sense that each party is attempting to stake as strong a claim as possible prior to the second stage) and determines the conflict outcome, followed by an arbitration stage, in which the gains from agreement are allocated between the parties.

Of course, if the parties both accept the principles presented above, they can determine for themselves the agreement which an arbitrator would suggest, and thus avoid formal arbitration. However, the final agreement might still require a regulator, at least in the form of a judicial system, in order to guarantee that both parties carry through with their responsibilities under the agreement.

It is important to note that the threats made by both parties in the first stage need never be carried out: The procedure always leads to agreement. This will not necessarily be the case when, in the next chapters, we consider bargaining problems in which the parties are not perfectly informed about the situation they face.

[Schelling and others have noted a tactic available in bargaining, even under conditions of complete information, which is not considered in Nash's analysis. One of the parties can attempt to make a preemptive precommitment which changes the set of feasible agreements. For example, in the Battle of the Sexes, one of the parties could make a nonrefundable prepayment on a weekend for two at his or her more-favored destination, and present this action to the other as a <u>fait</u>

accompli. If both should do so, an inferior outcome must result.]

4. Bargaining under uncertainty

The difficulties which can arise when parties hold private information are dramatically illustrated in the following well-known example.

Example 6 (the Akerlof "lemon" problem). An owner of a used car is negotiating with a prospective buyer. The quality of the car is known only to the seller; expressed in terms of the car's value to the seller, the buyer believes it equally likely to be worth any amount between $0 and $500. The buyer, who would utilize the car to a greater extent, would derive 50% more value from its ownership. At what price might a sale take place?

Only if the car is worth less than $x to the seller would he agree to a sale at $x. But then, from the buyer's perspective, given that the seller agrees to accept a price of $x, the expected value of the car to the seller is no more than $x/2, and therefore, its expected value to the buyer is at most $3x/4. Hence, the buyer should refuse to buy the car at any price the seller is willing to accept! (Classroom experiments consistently bear out the empirical validity of this analysis - Subjects argue interminably, but trade never occurs.) Even though both parties know that a mutually advantageous trade exists, trade cannot take place unless someone acts irrationally.

4.1 Auditing

How might the seller and buyer work around this impasse? One possibility is to have a mechanic inspect the car, and provide an appraisal to them. This would convert the problem to one of complete information, amenable to the type of analysis described in Chapter 3.

Another possibility is to write a warranty into the sales contract, providing for payment adjustments after the buyer learns, through use, the quality of the car. Such a contract is actually a spectrum of contingent contracts, each written under the assumption of complete information, one for every possible quality level of the car. (Clearly, a regulator is required to implement a warranteed sale.)

In the first case, the mechanic acts as an auditor; in the second, post-sale observation plays an auditing role. We generally view an auditor as an individual (or procedure) through which information held by one party can be made public. Unless specific mention to the contrary is made, we shall assume throughout the remainder of this paper that auditing is not available, and will consider instead how, through their actions, parties provide information to one another, or to intervenors.

4.2 Games with incomplete information

Beginning in 1965 with research sponsored by the U.S. Arms Control and Disarmament Agency, game theorists and economists have focused substantial effort on attempts to understand bargaining under uncertainty. Most of this research falls into one of two categories: studies of what can conceivably be accomplished by the appropriate choice of a format for negotiations, and studies of what can be expected to occur in the context of some specific format.

Consider a general view of two-party bargaining. The parties make statements, true or false: they bluff, threaten, bluster, and otherwise interact in attempts to convince one another of their respective preferences and constraints. Finally, something happens - either an agreement is reached, or conflict ensues.

Basically, each party, knowing his own preferences, adopts a "private strategy," which specifies how he will act (or respond) at any stage of the negotiations, given what has transpired prior to that stage. (One can view this private strategy as a complete, explicit set of instructions given to an agent who will represent the party in the negotiations.)

An important (and frequently overlooked) consideration in choosing our own private strategy is that the opposing party does not know our own preferences and constraints (i.e., he does not know our "type"), and therefore continually updates his perception of us on the basis of our observed behavior. He does this by assessing the likelihood that we would act the way we have, for each of the possible types of opponent we might be. Therefore, he bases his responses (i.e., portions of his own private strategy) on his assumptions of how each of our possible types would act (i.e. on the private strategies he assumes our various potential types would adopt). It logically follows that, in order to decide upon our own appropriate actions, we must anticipate the conclusions he will draw: We must ask ourselves how we would have acted, had we been any type other than the type we actually are. (The Scottish poet Robert Burns anticipated our need to take this view when he wrote, in his _To a Louse_, "O wad some Pow'r the giftie gie us, to see oursels as others see us! ...")

Example 7 (the Walkenhorst Chemical case). Jack Walkenhorst, a young inventor, is preparing for a court hearing. Lakeland Chemical, a large conglomerate, has filed a patent application on a production process similar to one he has previously patented. If the court validates Lakeland's application and Lakeland begins to compete with Jack, he will suffer substantial short-term losses. However, as a result of his recent research he has an important piece of private information: Another process, significantly different from and much cheaper than either of the two contested processes, is commercially feasible. If Lakeland wins the suit, and engages Jack in competition, they will ultimately lose money, and Jack will eventually recoup his losses.

If Lakeland knew the true situation, they would freely choose to withdraw their patent application. But Jack cannot reveal any details of the new process without jeopardizing the new patent, for which he will not be prepared to apply for another six months. At this point, the outcome of the court case appears to be a toss-up. What can Jack do to improve his situation?

In this example, Jack would like to say to Lakeland, "Believe me - If you pursue the suit, win, and engage me in competition, you will eventually regret it." However, Lakeland cannot know whether Jack truly has something up his sleeve, or is merely bluffing in order to protect his position should he lose the case: that is, they don't know Jack's "type." If the making of this statement would convince them to stay out of competition, then his nonexistent, but potential, "bluffing" type would certainly make the statement. Therefore, Lakeland's perception of the situation will not be changed by Jack's statement: Either of Jack's types (his true type, or his bluffing type) would make it. Consequently, if Lakeland originally considers it unlikely that Jack has the ability to hurt them, his statement will not deter their entry.

The moral of this story is that, when preparing for negotiations, we must not merely focus on the private strategy our actual type will follow: We must also consider which private strategies we would follow, were we any type other than our actual one. One can view the preparation for negotiations as a roundtable discussion among a party and his various alter egos, in which the participants must decide upon the coordinated face they will present to the outside world. Some types might wish to "bluff," i.e., to mimic the private strategy of some other type in hope of persuading the outside world that they are that type. Other types might wish, in turn, to "signal," i.e., to take actions which clearly reveal their actual situation. (Jack Walkenhorst might choose to drop his current suit against Lakeland as a token of faith. If this would convince Lakeland to delay competition, his actual type would gain; if the delay would be of less value to his "bluffing" type than the current 50% chance of winning the suit, then that type would not make the same offer - Dropping the suit is a signal of his true type which Lakeland can believe. Indeed, if Jack is not clever enough to think of this signal, Lakeland (or an intervenor) can suggest it to him. A formal agreement, in which Jack drops the suit in exchange for a six-month delay in Lakeland's entry, works to the advantage of both parties and should be acceptable to both.)

As we have seen, a party's types might find themselves with conflicting desires; some resolution of this internal conflict must be reached before the one "true" type can decide upon his private strategy. An important note, to which we will return shortly, is that the final resolution of the inter-type conflict cannot involve binding agreements across types. Only one type actually exists; the others can't penalize him for breaking any agreement.

In view of the previous considerations, game theorists have chosen to define a strategy for a party in a bargaining environment as a joint specification of private strategies, one for each of his possible types. The private strategy of the true type is implemented; an opponent updates his beliefs about the party on the basis of observed actions, together with that opponent's guess as to the full strategy which was selected. (The standard rule of probability theory used for this updating is known as "Bayes' Rule.")

In a rational world, in which each party considers his opponent's strategic problems as well as his own, it is reasonable to expect that each party will believe his opponent's strategy to be optimal for each of the opponent's types, given the opponent's belief about the party's own choice of strategy. (This is because no type can be compelled by the other types to adopt a non-optimal strategy.) A pairing of such strategies, in which each believes correctly, is formally known as a (Bayesian) equilibrium point of the bargaining "game." The analysis of any specific dispute begins (for a game theorist) with a game-model of the communication and commitment abilities of the parties, and proceeds with a study of the Bayesian equilibria of the game.

Example 8 (dissolving a partnership). Two individuals jointly own a piece of property. They have decided to sever their relationship, and for one of the two to buy the land from the other. Each knows how valuable the land is to him, but is unsure of its worth to the other. They agree that each will write down a bid; the high bidder will keep the land, and pay the amount of his bid to the other.

Assume that each is equally likely to value the land at any level between \$0 and \$1200, and that both know this. Then the unique Bayesian equilibrium point of the bidding game is for each to bid one-third of his own valuation. If, for example, one of them values the land at \$300 and believes the other to be following the indicated equilibrium strategy, then by bidding \$100 he has an expected payoff of $\frac{1}{4}$ \$200 + $\frac{3}{4}$ \$250; he expects to win with probability $\frac{1}{4}$, and when he loses, he expects the other's (winning) bid to be between \$100 and \$400. This private strategy is optimal for him, given his belief about the other's behavior. (Given his belief that his partner will bid a third of the partner's valuation, his own expected payoff, when his valuation is v and he bids b, is $(3b/1200).(v-b) + (1 - 3b/1200).(b+400)/2$. In general, this is maximized by taking $b = v/3$.)

Observe that this bidding arrangement always yields a Pareto-efficient result, i.e., the individual who values the land more highly always ends up in possession of it. Hence, the appropriate choice of a dispute-resolution procedure can, at times, circumvent inefficiencies of the type which arise in the lemon problem.

Note also that an intervenor could suggest the use of this procedure, if the parties found themselves unable to work out an agreement on their own.

4.3 The revelation principle

There are, of course, numerous other procedures that an intervenor could suggest in order to resolve the dispute in Example 8. Let us consider the (seemingly appalling) general question of what outcomes can result, at equilibrium, from <u>any</u> procedure which might be used to resolve a given dispute.

A simple, yet conceptually deep, type of analysis has become standard. Consider any equilibrium pair of strategies in a particular game. Each party's strategy can be viewed as a book, with each chapter detailing the private strategy of one of that party's types. Given the two actual types, a pairing of the private strategies in the two appropriate chapters will lead to an outcome of the game.

Next, step back from this setting, and picture the two parties in separate rooms, each instructing an agent on how to act on his behalf. Each agent holds in hand the strategy book of his side; all he must be told is which chapter to use. From this new perspective, the two parties can be thought of as playing an "agent-instruction" game, in which the strategy books are prespecified and each must merely tell his agent his type (or, equivalently, point to a chapter in his strategy book). An equilibrium point in this new "chapter-selection" game is for each to tell the truth to his agent. Otherwise, the original strategies could not have been in equilibrium in the original game.

Consequently, <u>anything</u> which can be accomplished at equilibrium through the use of <u>any</u> particular dispute-resolution procedure, can also be accomplished through the use of some other procedure in which the only actions available to the parties are to state their (respective) types, and in which it is in equilibrium for each to truthfully reveal his type.

Example 9. Two parties must share the cost of a public works project; if they cannot agree, the project will not be carried out. The project will cost \$100. Both parties are risk-neutral, and it is known by both that Party A will derive \$90 in benefit from the project. Party B knows the benefit he will receive, but all that is known to A is that there is a 50% chance it is worth \$90, and a 50% chance it is worth only \$30, to B. What possible agreements could they reach?

The revelation principle tells us that any agreement which could be arrived at through any negotiation procedure will be an agreement which could also be achieved in a formally-structured game in which each simply names his type, and each has no incentive to lie. (Since A's type is known to both, only B will actually have a move in this game.)

An outcome of the revelation game will be, most generally, a probability that the project will be carried out, and a sharing of the \$100 cost if it is indeed carried out. Since a different outcome might result from each of the two type-declarations B might make, the full spectrum of possible agreements can be characterized by four parameters: p_H and p_L, the probabilities of project commencement given that B announces his type to be "high" (\$90) or "low" (\$30), and e_H and e_L, the payments to be made by B given his announcement and that the project is carried out.

In order for truth-telling to be optimal (i.e., a best response to A's null action) for B, these parameters must satisfy two incentive constraints:

$$(90 - e_H).p_H \geq (90 - e_L).p_L \qquad \text{(the \$90-type must prefer announcing "H" over "L")}$$

$$(30 - e_H).p_H \leq (30 - e_L).p_L \qquad \text{(The \$30-type must prefer announcing "L" over "H")}$$

Furthermore, in order for A, and for both types of B, to be willing to agree to the procedure, it must satisfy the following <u>participation constraints</u>:

$$\frac{1}{2}\, p_H.(e_H - 10) + \frac{1}{2}\, p_L.(e_L - 10) \geq 0 \qquad \text{(for A to participate)}$$

$$e_H \leq 90 \;,\; e_L \leq 30 \qquad \text{(for both of B's types to participate)}$$

It follows (algebraically) from all this that we must have $p_H \geq p_L$ and $e_H \geq e_L$: that is, when B reports himself to be the \$90-type (in practice, when he acts as if he is that type), the project is more likely to be carried out, and he will be charged a larger share of the cost.

There <u>are</u> agreements which will lead to the project always being done, i.e., agreements with $p_H = p_L = 1$. However, such agreements must have $e_H = e_L \leq 30$, and hence A must bear at least 70% of the cost, independent of B's type. Any alternative agreement which lessens A's burden must have $p_L < 1$, and hence must require that the project is sometimes <u>not</u> carried out. For example, if gains from the project are to be split evenly between A and the announced type of B, then we must have $e_H = 50$ and $e_L = 20$; if the project is to be certainly carried out ($p_H = 1$) when the \$90-type is announced, then p_L can be at most 4/7.

One interpretation of this result is that efficiency and equity are, at times, at least partially incompatible. Only the threat by A of not doing the project can "separate" the two types of B, and this threat is only viable if, when B claims to be his \$30-type, A sometimes actually carries it through.

Example 10 (bilateral trade). One classical type of bargaining problem involves a seller and a buyer, each uncertain of how the other values an object currently held by the seller. Assume that each believes the other to be equally likely to value the object at any amount between \$0 and \$300; each, of course, knows his own valuation. According to the revelation principle, the possible agreements which can result from any choice of negotiation format can be characterized by a pair of functions $p(v_S, v_B)$ (the probability that trade takes place when the seller announces his valuation to be v_S and the buyer announces his to be v_B) and $e(v_S, v_B)$ (the amount to be paid by the buyer when these announcements are made and trade does take place). This pair of functions must satisfy a continuum of incentive constraints: Each buyer or seller type must prefer announcing truthfully to making any other announcement. Furthermore, the functions must satisfy a continuum of participation constraints: Every seller type must expect to be paid at least his valuation when trade takes place, and every buyer type must expect to pay no more than his valuation.

Consider one particular format for arranging a sale. Each writes down a price. If the seller writes a higher price than the buyer, no trade occurs; otherwise, the object is sold at the average of the two amounts. It is simple to show that it is not in equilibrium for both to write truthfully their valuations: If either is truthful, the other can gain by exaggeration. One natural equilibrium pair of strategies is for the seller to write down $2/3\ t_S + 75$, where t_S is his actual valuation, and for the buyer to write down $2/3\ t_B + 25$, where t_B is his valuation. Notice that when the buyer's valuation is only slightly greater than the seller's, trade does not take place; indeed, if the seller's type is greater than 225, there is never a trade.

Consider an alternative mechanism, wherein each writes down an amount (v_S and v_B, respectively), and trade takes place only if $v_S \leq v_B - 75$, at a price of $(v_S+v_B)/3 + 50$.

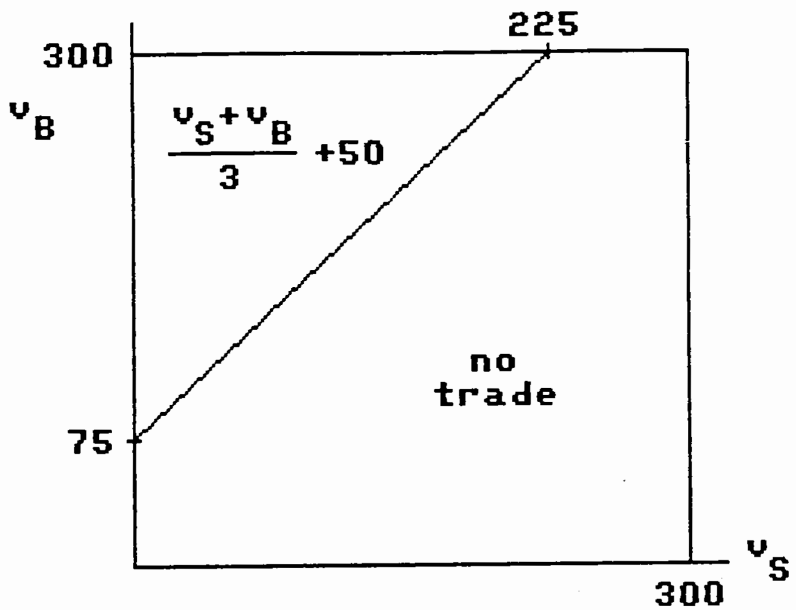

Buyer-Seller mechanism to maximize ex ante expected gains from trade

Figure 2

It can be verified that it _is_ in equilibrium for both to tell the truth; furthermore, every pair of types faces (at equilibrium) the same outcome here as they did in the previous mechanism. This latter mechanism is, in fact, the "revelation game" derived from the former game and equilibrium point using the approach outlined at the beginning of this section. (It is known that the mechanism given here maximizes the traders' *ex ante* (before they learn their types) expected joint gains from trade: if one were organizing a market within which such traders would be forced to deal, this mechanism would be the natural choice.)

4.4 Incentive-efficiency

It would seem foolish for an arbitrator to suggest a mechanism for dispute resolution which left every type no better off, and some types worse off, than some other mechanism would have. In other words, the suggested mechanism should be *incentive-efficient*, i.e., efficient subject to the incentive constraints. A generalization of Nash's solution to the complete-information bargaining problem should therefore select a particular incentive-efficient mechanism.

In Example 9, the incentive-efficient agreements all have $p_H = 1$. Furthermore, if $e_L \geq 10$, then $p_L = (90 - e_H)/(90 - e_L)$.

In Example 10, the mechanism presented can be shown to be one of the incentive-efficient mechanisms. Furthermore, no incentive-efficient mechanism is *ex post* Pareto-efficient: It is impossible to arrange for advantageous trades to always take place.

4.5 Equity and durability

Although parties usually enter negotiations with a primary objective of reaching an agreement advantageous to themselves, much of the ensuing discussion between the parties concerns the "fairness" of different proposed agreements.

From where do the parties obtain their notions of what is fair? Certainly, there are commonly accepted principles which are culturally based. "The greatest good for the greatest number." "From each according to his ability; to each according to his needs," and "Whatever can be obtained from the sweat of the brow" are examples of such principles; clearly, they stand somewhat in contradiction to each other.

Sometimes, precedent plays a role in perceptions of fairness. Labor negotiations typically take the previous contract as a starting point, and each party will argue that a concession on one issue "should" be matched by an opposing concession on another. At other times, a neutral third party will be asked to resolve a dispute in terms of his external view of equity: The parties will submit their dispute to binding arbitration.

The Nash solution in settings of complete information was derived from a list of desired properties, at least two of which (individual rationality and symmetry) were directly concerned with equity. Furthermore, the use of threat-making to establish the original conflict outcome carries with it a notion of equity: Those who will suffer relatively more if agreement is <u>not</u> reached, receive relatively less from the agreement which <u>is</u> reached. (In Example 4, Alfred receives less than half of the monetary gains available from trade with Burton.)

Recently, Myerson has proposed a generalization of the Nash solution to bargaining games with incomplete information. His approach gives explicit regard to the inter-type competition we have previously discussed.

Example 10 (continued). Assume that the mechanism described earlier (each names a price; trade takes place if the named prices are compatible, at the average of the two named prices; each follows his specified equilibrium strategy) is proposed to two traders. At the time of the proposal, each trader of course knows his own type. If the seller's valuation is greater than \$225, he will naturally make an objection. For example, if his valuation is \$250, he expects no trade to take place if this mechanism is used. Instead, he could commit himself to a first-and-final offer of some higher amount, say, \$275. Although the buyer may be antagonized by this action, if his valuation is greater than \$275 and he believes the seller's commitment then there is some chance that he will accept the offer. Similarly, if the buyer's valuation is less than \$75, he will object to the use of this mechanism.

Even when the seller's valuation is less than \$225, and he agrees to use the proposed mechanism, his mere agreement reveals something about him - namely, that his valuation is less than \$225. (Otherwise, the buyer would expect him to object.) With this extra information, the buyer might choose to reject the proposed mechanism, and put additional pressure on the seller.

Consequently, this mechanism is not "durable," in the sense that either it will not be accepted by at least one party, or it will be accepted and the equilibrium will not persist, since the parties' beliefs about each other will be changed by their acceptances.

In response to this difficulty, Myerson has proposed a "neutral" bargaining solution, which takes into account the inter-type conflict which could upset a proposed mechanism.

For the buyer-seller problem we are considering, the neutral bargaining solution is summarized by Figure 3: Each trader

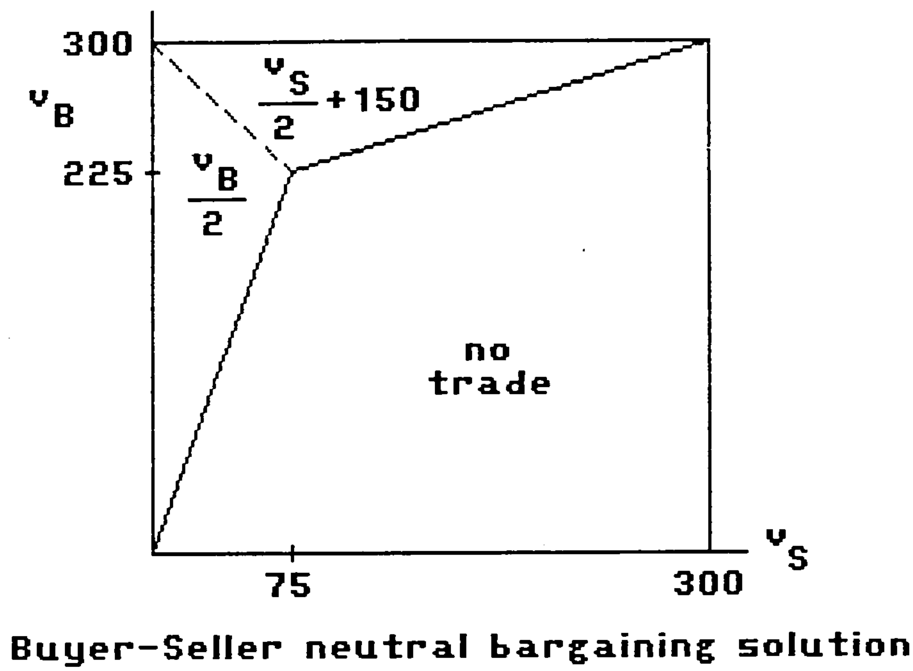

Figure 3

names his valuation; trade takes place if the named valuations (v_S for the seller, and v_B for the buyer) lie in one of the two triangles, at the price indicated within the triangle. It can be shown that this mechanism is both incentive-compatible (it is optimal for each to tell the truth, given that the other does so) and incentive-efficient. Notice that trade takes place more often for the extreme types than under the previous mechanism, and slightly less often for the intermediate types. This must be so, in order for the extreme types to not wish to "upset" the mechanism.

5. The role of time

In the previous chapter, we explored the range of settlements which could be obtained at equilibrium through <u>any</u> mechanism for negotiations, and we found that all such settlements could be obtained through the use of a one-stage simultaneous-type-revelation procedure.

However, in actuality negotiations typically take place over time, with the parties gradually revealing themselves through their actions. We shall look in turn at negotiations over time when there is no information to be revealed, at protracted negotiations during which the parties accrue payoffs which depend on their stage-to-stage actions, and finally at negotiations in which information is revealed over time, but the only payoff comes at the conclusion.

5.1 Complete information

We have already seen that there is no impediment to parties reaching a Pareto-efficient outcome when there is no uncertainty. If we assume that bargaining takes place over time, and that there is a cost associated with delay in reaching an agreement, then the only efficient behavior must involve agreement being reached essentially immediately.

<u>Example 11 (the Rubinstein offer-counteroffer model).</u> Consider a seller, holding an object worth nothing to him, and a buyer who values the object at \$300. Both valuations are known to both parties. They negotiate through the exchange of offers: First the seller proposes a sale price, and the buyer either accepts that price, or rejects it. If he rejects it, then he follows with a counter-proposal, stating another

price. The seller can either accept this new price, or follow with yet another proposal. And so on...

Assume that one unit of time passes between any rejection and the subsequent proposal. One model of the cost of delay in reaching an agreement involves discounting the payoffs of both parties in the final agreement. Specifically, let d_S and d_B be discount rates between 0 and 1, and assume the payoff to the seller if the final agreement is a price of p in the t-th stage is $d_S^{t-1}.p$, while the payoff to the buyer is $d_B^{t-1}.(300-p)$.

There are many equilibrium pairings of strategies in this bargaining game. For example, the seller can ask for \$290 in every stage, accepting a counter-offer only if it is at least \$290, and the buyer can accept any price at or below \$290, making a counter-offer of \$1 whenever the seller asks for more than \$290. (Recall that a strategy for a party must specify his action in any situation he might face.) This pairing leads to a sale at \$290 in the first stage: neither can do any better, as long as the other holds to his own specified strategy. Similarly, there are other equilibria which yield immediate sales at <u>any</u> price between \$0 and \$300.

However, the specified strategies call for foolish actions from the parties in certain circumstances. For example, suppose the seller opens with a proposal of \$290, and the buyer rejects this offer, making a counter-offer of \$289. (While this will not happen if they follow the specified strategies, the seller must be prepared for this possibility.) If \$289 is greater than d_S.\$300, it would be foolish for the seller to reject this counter-offer: He cannot expect to improve his lot by continuing the game through another stage. Hence, the specified pair of strategies calls for non-optimal actions in some "subgames" which arise off the equilibrium path.

An equilibrium point is said to be <u>perfect</u> if the parties' strategies call for optimal behavior in every subgame which might arise. Remarkably, if we restrict our attention to perfect equilibria in this sequential game, we find that there is only one. Let

$$p_S = \frac{300\ (1-d_B)}{1 - d_B d_S}\ , \text{ and } \quad p_B = p_S d_S\ .$$

The unique perfect equilibrium has the seller asking for p_S initially, and at every subsequent stage accepting any price of p_B or more. He rejects any lower counter-offer, again asking for p_S. The buyer accepts any price of p_S or less, rejecting higher prices and counter-offering a price of p_B. In this equilibrium, the sale takes place in the first stage, at a price of p_S.

If the discount factors are equal, and very close to 1 (i.e., if the interval between successive stages is quite short, so the cost of delayed agreement is small), then p_S will be very near \$150. More generally, if the parties face any bargaining problem without uncertainty, and the proposals and counter-proposals consist of feasible agreements, then the unique perfect equilibrium outcome is immediate settlement on an agreement which, when $p_S = p_B$ and both are near 1, is near the Nash solution to the bargaining problem. (This is a "limit" result.) One might view this as further validation of the Nash solution as a "natural" result of negotiations.

Short of the limit, the sale price is somewhat above \$150. This is because it is to the seller's advantage to move first:

if the buyer were given the first move, the price would be somewhat below $150.

Note that p_S is increasing in d_S, and decreasing in d_B. The less costly it is to either party to wait, the better off he is in the perfect equilibrium outcome. This accords well with common perception: In negotiations, patience is a virtue.

The parameters d_S and d_B can be given an alternative interpretation. Assume that there is no cost to waiting, but that there is a probability of $1-d_B$ that the seller will walk away from the negotiations any time he makes an offer which is rejected, and a probability of $1-d_S$ that the buyer will walk away any time one of <u>his</u> offers is rejected. The same, unique perfect equilibrium persists in this new setting: The more likely a party is to walk away if one of his offers is rejected (and the more likely his opponent <u>perceives</u> his departure to be), the better off he is at the equilibrium outcome.

5.2 <u>Repeated games with incomplete information</u>

Over the past twenty years, many researchers have studied the repeated play of a game, when the players' interests are strictly opposed and each holds private information. The most-commonly-studied model is one in which the parties are uncertain of the payoff structure of the game; the actions of both parties are publicly revealed at the end of each period, but neither side learns the payoffs; payoffs accumulate over time. While this is not too accurate a model of bargaining (where there usually is an end to the negotiations, and the terminal payoff is much more important than intermediate payoffs), still, investigation of this model provides insight into the way a party's opponent can learn about him from observation of his actions, and therefore, about how the party's actions should be chosen.

A principal result is that optimal strategies typically involve a single initial reference to the information a party holds, followed by period-to-period moves which depend only on the outcome of that single reference. An analogue of this in actual negotiations is the initial briefing of a representative, at which time he is given only the information the party he represents is prepared to reveal in the course of the negotiations; subsequently, no further information is revealed to him (and hence, his choice of actions during the negotiations can reveal no more than the information he is given at the briefing).

Recently, Hart has extended this analysis to games with private information on one side, and gains available to the players through cooperative actions. Hart's result is that, when mutual gains are available, it is frequently necessary to partially brief a representative, send him to the negotiating table, and (depending on the course of the negotiations) to periodically recall the agent for further briefings. This work provides some insight into the process observed in international arms control negotiations.

5.3 Bargaining with incomplete information

Attempts to extend the Rubinstein offer-counteroffer analysis (with time-discounted payoffs) to bargaining problems with incomplete information have met with difficulties. The most successful approach to date is that of Grossman and Perry:

the difficulty they encounter is that such games have many equilibria, and the natural analogue of the "perfection" argument in the Rubinstein model is not clear. An active area of current research is the development of equilibrium selection techniques, to be applied in order to obtain a single "special" equilibrium of such games. However, the predictive appeal of such models is questionable, given the numerous ways in which individuals could attempt to affect the selection process.

In general, negotiations which unfold over time vary in many dimensions: the nature of the information initially held by the parties, the channels of communication, the costs of delayed agreement or conflict, the types of settlements which are feasible - differences in these dimensions lead to problems requiring substantially different analytical approaches. We give here a simple example which has been used by various authors as a model of courtship behavior, primitive tribal customs, military escalation, and strikes. In this example, there is competition for an indivisible prize. Each party knows his own valuation of the prize, communication is limited to observation of the other's intransigence, costs of delayed settlement are opportunity losses, which accrue to both parties linearly over time, and the only possible agreements require total concession by one party.

Example 12 (The War of Attrition). Two parties compete for possession of a prize. The two private valuations are independent draws from a commonly-known distribution, and each party knows only his own valuation. The parties face one another passively, suffering a constant loss per unit time. Competition ends when one party withdraws; each pays his accrued loss (the same amount for each), and the remaining party claims the prize.

Milgrom and Weber have studied this situation, and obtained the following results. (1) There is a unique symmetric equilibrium point for this game. (2) At equilibrium, there is a positive probability that competition continues for so long that both the loser <u>and</u> the winner suffer net losses. (Of course, at equilibrium each has a non-negative <u>expected</u> profit; otherwise, one of them could gain in expectation by withdrawing immediately.) (3) If it is commonly known that both parties have exactly the same valuation, then at equilibrium both have expected payoffs of 0 (i.e., on average, they will "compete away" the entire value of the prize). (4) The greater the likely difference in valuations, the greater the expected payoffs to both. (Or, as the French say, "Vive la difference!") (5) The longer competition endures, the more likely it is to continue. (The outlook for settlement grows steadily bleaker over time. In the strike interpretation, this result provides some justification for a policy of delayed government intervention.)

6. <u>Summary</u>

Game-theoretic studies of bargaining problems (as well as of other types of conflict situations) have helped to clarify our understanding of the role of private information and the nature of strategies in competitive environments. These studies have also provided a formal structure for the investigation of issues involving efficiency and equity, and have helped to delineate the roles played by intervenors (and the limitations intervenors must acknowledge).

6.1 Research prospects

One active area of research deals with equilibrium point selection and classification; a goal of this research is to explain why some types of equilibrium behavior are observed more frequently than others. Another area concerns the evaluation and comparison of alternative frameworks for dispute resolution: For example, one might ask how different rules for the allocation of court costs and legal fees in civil suits affect pre-trial settlement behavior. Several technical problems remain to be solved before current analytical techniques can be extended to cover problems involving multiple dimensions of uncertainty, or multiple negotiation stages.

Most of this paper has discussed problems of two-party bargaining. There is a rich history of game-theoretic research into multi-party issues, but most of this work has focused on issues of stability (the core, bargaining set, von Neumann-Morgenstern solution, and the like) and equity (the Shapley value). Little is yet known about the dynamics of coalition formation (and dissolution) in multi-party negotiations.

6.2 Lessons to be remembered

In conclusion, what are the principal messages game theory has to deliver to practitioners? For negotiators, there are two: Realize that, when you hold private information, it is important to consider what actions you would have taken, had your information been different than it actually is. Carefully consider what strategy you expect your opponent to follow, and whether this expectation is justified. (In particular, if your own strategy, together with your opponent's, does not form an equilibrium point, ask yourself why. Are you expecting him to be less clever than you?)

For intervenors: Be aware of incentive constraints, and their role in occasionally leaving the parties in an irreconcilable position of conflict. (Strikes, for example, are always non-optimal <u>ex post</u>; still, the threat of a strike is an essential component of the labor-management bargaining process. The intensity of the threat can reveal useful information, but only if there is a chance that it will have to be carried out.) Provide opportunities for "safe" revelation of information to you; seek means for auditing statements made to you, or for conditioning the contract on future information or behavior.

References

G. Akerlof [1970], "The Market for Lemons: Qualitative Uncertainty and the Market Mechanism," Quarterly Journal of Economics 84, 488-500.

V.P. Crawford [1985], "Efficient and Durable Decision Rules: A Reformulation," Econometrica 53, 817-835.

S.J. Grossman and M. Perry [1984], "Sequential Bargaining under Asymmetric Information," discussion paper, University of Chicago.

J.C. Harsanyi [1967-8], "Games with Incomplete Information Played by 'Bayesian' Players," Management Science 14, 159-182, 320-334, 486-502.

S. Hart [1985], "Non-Zero-Sum Two-Person Repeated Games with Incomplete Information," Mathematics of Operations Research 10, 117-153.

B. Holmstrom and R.B. Myerson [1983], "Efficient and Durable Decision Rules with Incomplete Information," Econometrica 51, 1799-1819.

D. Kreps, "Walkenhorst Chemical," case study, Stanford University Graduate School of Business.

R.D. Luce and H. Raiffa [1957], Games and Decisions, Wiley.

P.R. Milgrom and R.J. Weber [1983], "Distributional Strategies for Games with Incomplete Information," discussion paper, Northwestern University, to appear in Mathematics of Operations Research.

R.B. Myerson [1979], "Incentive Compatibility and the Bargaining Problem," Econometrica 47, 61-73.

R.B. Myerson [1984], "Two-Person Bargaining Problems with Incomplete Information," Econometrica 52, 461-487.

R.B. Myerson [1985a], "Bayesian Equilibrium and Incentive Compatibility: An Introduction," discussion paper, Northwestern University, to appear in Social Goals and Social Organization, edited by L. Hurwicz, D. Schmeidler, and H. Sonnenschein, Cambridge University Press.

R.B. Myerson [1985b], "Analysis of Two-Person Bargaining Games with Incomplete Information," in Roth [1985].

R.B. Myerson [1985c], "Negotiation in Games: A Theoretical Overview," discussion paper, Northwestern University.

R.B. Myerson and M.A. Satterthwaite [1983], "Efficient Mechanisms for Bilateral Trading," Journal of Economic Theory 29, 265-281.

J.F. Nash [1950], "The Bargaining Problem," Econometrica 18, 155-162.

J.F. Nash [1951], "Noncooperative Games," Annals of Mathematics 54, 289-295.

J.F. Nash [1953], "Two-Person Cooperative Games," Econometrica 21, 128-40.

H. Raiffa [1982], The Art and Science of Negotiation, Harvard University Press.

A.E. Roth (ed.) [1985], Game-Theoretic Models of Bargaining, Cambridge University Press.

A. Rubinstein [1982], "Perfect Equilibrium in a Bargaining Model," Econometrica 50, 97-109.

A. Rubinstein [1985], "Choice of Conjectures in a Bargaining Game with Incomplete Information," in Roth [1985].

T.C. Shelling [1960], The Strategy of Conflict, Harvard University Press.

M. Shubik [1982], Game Theory in the Social Sciences, MIT Press.

S. Sorin [1980], "An Introduction to Two-Person Zero-Sum Repeated Games with Incomplete Information," IMSSS Technical Report No. 312, Stanford University.

Bibliographic Notes

Primary sources for general information on the tools and models of game theory are Luce and Raiffa [1957] and Shubik [1982]; Shelling [1960] provides interesting commentary and additional perspectives.

Raiffa [1982] presents a mixture of decision-analytic and game-theoretic views on negotiation; Roth [1985] collects a number of survey and research papers covering the focus of current research activity.

Acknowledgements

The preparation of this paper was supported, in part, by the National Institute for Dispute Resolution. Comments from David Lax, Jeanne Brett, Max Bazerman, William Samuelson, and James Sebenius, together with extensive discussions with Roger Myerson, contributed to the paper's final form; however, the author bears full responsibility for any inaccuracies or personal views appearing herein. Part of the writing was done while the author was a visiting faculty member at the Graduate Institute of Business Administration, in Chulalongkorn University, Bangkok, Thailand.

Teaching materials

The NIDR Program on Professional Education has supported the development of two packages of classroom materials on dispute resolution.

One package, "The Manager as Negotiator and Dispute Resolver," includes four experiential exercises and detailed teaching notes which develop views of both negotiation and third-party intervention from the perspectives of organizational behavior, organizational design, and managerial psychology.

The other package, "The Manager as Negotiator: Teaching Materials in Dispute Resolution for Decision Analysis and Economics Courses," includes three role-playing cases, teaching notes, a collection of annotated examples suitable for inclusion in decision analysis or microeconomics courses, and this paper.

Faculty interested in introducing or developing classroom material on dispute resolution are encouraged to contact NIDR. Further information, and order forms for the two packages, can be obtained by writing:

Robert Jones
National Institute for Dispute Resolution
1901 L Street NW, Suite 600
Washington, DC 20036
(202-466-4764)