Working Paper Series

WORKING PAPER NO. 685

MODELLING NATURAL RESOURCE NEGOTIATIONS; AN APPLICATION TO CALIFORNIA WATER POLICY

by

Gregory D. Adams,

DEPARTMENT OF AGRICULTURAL AND RESOURCE ECONOMICS

BERKELEY

CALIFORNIA AGRICULTURAL EXPERIMENT STATION

University of California

DEPARTMENT OF AGRICULTURAL AND RESOURCE ECONOMICS
DIVISION OF AGRICULTURE AND NATURAL RESOURCES
UNIVERSITY OF CALIFORNIA AT BERKELEY

WORKING PAPER NO. 685

MODELLING NATURAL RESOURCE NEGOTIATIONS; AN APPLICATION TO CALIFORNIA WATER POLICY

by

Gregory D. Adams,

Gordon C. Rausser,

and

Leo K. Simon*

The authors are Graduate Student, Robert Gordon Sproul Distinguished Professor, and Assistant Research Economist, respectively, in the Department of Agricultural and Resource Economics, University of California at Berkeley.

California Agricultural Experiment Station
Giannini Foundation of Agricultural Economics
August, 1993

MODELLING NATURAL RESOURCE NEGOTIATIONS: AN APPLICATION TO CALIFORNIA WATER POLICY

Abstract

Increasing urban and environmental demands for water have led to widespread calls for reform of the existing water allocation system in California. The major interest groups, agricultural water users, urban water agencies and environmental groups, are negotiating over policy reform packages in a non-cooperative, multilateral setting. In this paper we advance a new framework for noncooperative, multilateral bargaining that can be used to conceptualize the negotiation process. In the proposed game theoretic setting, the outcome of the negotiation process depends crucially on the "constitutional" structure of the game. Computer simulations investigate several key issues of the policy debate.

"constitutional" structure of the game: how much input does each group have to the decision making process, what coalitions of groups can adopt and implement proposals, what is the space of issues over which negotiations take place and what happens in the event that the parties fail to reach agreement? As such, the model can be used both by individual participants in a negotiation, to determine optimal strategies, and by sponsoring agencies of negotiations, in order to structure the negotiations in manner that is likely to facilitate an agreement.

A timely example of such negotiations is the "Three Way Water Agreement Process", currently taking place in California. These negotiations, between representatives of urban and agricultural water users and environmental groups have been taking place for the past three years, and have made substantial progress in formulating politically feasible policies to address water allocation disputes in California (Adams). During the course of these negotiations, however, it has become apparent that past (and future) progress has (and will) depend in no small part on the institutional structure of the negotiation process.

We take the California water policy negotiations as a motivating example to demonstrate the noncooperative, multilateral bargaining framework. The multilateral bargaining model is introduced in Section 2. Section 3 presents an overview of the water policy negotiations in California and structures these negotiations in a manner consistent with the model. Section 4 provides quantitative results on selected aspects of the bargaining process. Finally, section

1. Introduction

Natural resource disputes typically involve a large number of individuals or interest groups and a large number of issues. Resolution of these disputes will often involve some form of bargaining; either formally (such as the negotiations leading to the Montreal Protocol on chloroflrocarbon emissions) or in a variety of less formally settings, which may include legislative battles (with the attendant "backroom" negotiations), pre-trail or court monitored settlement negotiations and spontaneous negotiations that arise between some or all of the affected parties.

To date, the economics profession has been largely silent about the optimal structure and strategies for negotiation processes of the type considered here. The vast literature on non-cooperative bargaining is overwhelmingly devoted to models of two players and one issue, and thus of little use when considering more complex, and realistic, negotiation processes. This is unfortunate since unstructured negotiations of this type may well prove not only costly, but unproductive. Accordingly, in this paper we advance a new framework for noncooperative, multilateral bargaining that can be used to conceptualize the negotiation process.

The proposed model is perhaps best thought of as a Computable General Equilibrium model of a negotiation process. In the proposed game theoretic setting, the outcome of the negotiation process depends crucially on the

1. Introduction

Natural resource disputes typically involve a large number of individuals or interest groups and a large number of issues. Resolution of these disputes will often involve some form of bargaining; either formally (such as the negotiations leading to the Montreal Protocol on chloroflrocarbon emissions) or in a variety of less formally settings, which may include legislative battles (with the attendant "backroom" negotiations), pre-trail or court monitored settlement negotiations and spontaneous negotiations that arise between some or all of the affected parties.

To date, the economics profession has been largely silent about the optimal structure and strategies for negotiation processes of the type considered here. The vast literature on non-cooperative bargaining is overwhelmingly devoted to models of two players and one issue, and thus of little use when considering more complex, and realistic, negotiation processes. This is unfortunate since unstructured negotiations of this type may well prove not only costly, but unproductive. Accordingly, in this paper we advance a new framework for noncooperative, multilateral bargaining that can be used to conceptualize the negotiation process.

The proposed model is perhaps best thought of as a Computable General Equilibrium model of a negotiation process. In the proposed game theoretic setting, the outcome of the negotiation process depends crucially on the "constitutional" structure of the game: how much input does each group have to the decision making process, what coalitions of groups can adopt and implement proposals, what is the space of issues over which negotiations take place and what happens in the event that the parties fail to reach agreement? As such, the model can be used both by individual participants in a negotiation, to determine optimal strategies, and by sponsoring agencies of negotiations, in order to structure the negotiations in manner that is likely to facilitate an agreement.

A timely example of such negotiations is the "Three Way Water Agreement Process", currently taking place in California. These negotiations, between representatives of urban and agricultural water users and environmental groups have been taking place for the past three years, and have made substantial progress in formulating politically feasible policies to address water allocation disputes in California (Adams). During the course of these negotiations, however, it has become apparent that past (and future) progress has (and will) depend in no small part on the institutional structure of the negotiation process.

We take the California water policy negotiations as a motivating example to demonstrate the noncooperative, multilateral bargaining framework. The multilateral bargaining model is introduced in Section 2. Section 3 presents an overview of the water policy negotiations in California and structures these negotiations in a manner consistent with the model. Section 4 provides quantitative results on selected aspects of the bargaining process. Finally, section

5 offers conclusions and directions for future research.

2. Multilateral Bargaining Model

The theoretical foundations for this paper are laid out in Rausser and Simon where a noncooperative model of multilateral bargaining is developed. A heuristic description of the model is presented below, while a more formal presentation is included in the Appendix. Readers interested in the complete proofs are referred to the original paper.

The model can be viewed as an extension of the classical Stahl-Rubinstein bargaining game in which two players take turns proposing a division of a "pie" (Stahl, 1972, 1977; Rubinstein). In the classical game, one player proposes a division, which the other can accept or reject. If the division is accepted, the game ends and the division is adopted; if it is rejected, the second player then makes a proposal, which the first player then accepts or rejects; and so on. The Rausser and Simon generalization of this framework incorporates multiple players and multidimensional issue spaces.

In a *multilateral bargaining problem*, there is a finite collection of players who meet together to select a *policy* from some collection of possible alternatives. In addition to these alternatives, there is a *distinguished disagreement policy*, which is imposed by default if players fail to reach agreement. Each player has a utility function defined on the set of possible policies. Players are presumed to be risk averse.

The specification of a multilateral bargaining problem includes a list of *admissible coalitions*. An admissible coalition is interpreted as a subset of the players that can impose a policy decision on the group as a whole. For example, in majority rule decision making, a coalition is defined to be admissible if and only if it contains a majority of the group. Alternatively, a unanimity decision rule implies that the only admissible coalition is the coalition of the whole. More generally, the set of admissible coalitions may have a variety of structures.

A *multilateral-bargaining game* is derived from a multilateral bargaining problem by superimposing upon it a "negotiation process." Specifically, each bargaining game has a finite number of negotiating rounds. A distinction is drawn between odd-numbered rounds of negotiations, called *offer rounds*, and even-numbered rounds, called *response rounds*. In an offer round, each player chooses a *proposal*, consisting of a policy and an admissible coalition. In response rounds, each player specifies an *acceptance set*, indicating the vectors that the player will accept if invited to join a coalition in that round. A *strategy* for a player is a collection of proposals and acceptance sets, one for each round of the game.

Prior to each response round, a proposer is chosen randomly "by nature," according to an exogenously specified vector of *access probabilities*. These i.i.d. probabilities are interpreted as measures of players' relative political "effectiveness"—the higher a player's access weight, the more likely it is that she averse.

The specification of a multilateral bargaining problem includes a list of *admissible coalitions*. An admissible coalition is interpreted as a subset of the players that can impose a policy decision on the group as a whole. For example, in majority rule decision making, a coalition is defined to be admissible if and only if it contains a majority of the group. Alternatively, a unanimity decision rule implies that the only admissible coalition is the coalition of the whole. More generally, the set of admissible coalitions may have a variety of structures.

A *multilateral-bargaining game* is derived from a multilateral bargaining problem by superimposing upon it a "negotiation process." Specifically, each bargaining game has a finite number of negotiating rounds. A distinction is drawn between odd-numbered rounds of negotiations, called *offer rounds*, and even-numbered rounds, called *response rounds*. In an offer round, each player chooses a *proposal*, consisting of a policy and an admissible coalition. In response rounds, each player specifies an *acceptance set*, indicating the vectors that the player will accept if invited to join a coalition in that round. A *strategy* for a player is a collection of proposals and acceptance sets, one for each round of the game.

Prior to each response round, a proposer is chosen randomly "by nature," according to an exogenously specified vector of *access probabilities*. These i.i.d. probabilities are interpreted as measures of players' relative political "effectiveness"—the higher a player's access weight, the more likely it is that she

5 offers conclusions and directions for future research.

2. Multilateral Bargaining Model

The theoretical foundations for this paper are laid out in Rausser and Simon where a noncooperative model of multilateral bargaining is developed. A heuristic description of the model is presented below, while a more formal presentation is included in the Appendix. Readers interested in the complete proofs are referred to the original paper.

The model can be viewed as an extension of the classical Stahl-Rubinstein bargaining game in which two players take turns proposing a division of a "pie" (Stahl, 1972, 1977; Rubinstein). In the classical game, one player proposes a division, which the other can accept or reject. If the division is accepted, the game ends and the division is adopted; if it is rejected, the second player then makes a proposal, which the first player then accepts or rejects; and so on. The Rausser and Simon generalization of this framework incorporates multiple players and multidimensional issue spaces.

In a *multilateral bargaining problem*, there is a finite collection of players who meet together to select a *policy* from some collection of possible alternatives. In addition to these alternatives, there is a *distinguished disagreement policy*, which is imposed by default if players fail to reach agreement. Each player has a utility function defined on the set of possible policies. Players are presumed to be risk will "seize the initiative" in the negotiations. A player's high access might reflect the extent of her political power within the organization or, perhaps, a talent for formulating issues in ways that can lead to workable compromises.

Together with the vector of access probabilities, each profile of strategies uniquely identifies an *outcome*, which is a random variable defined on the set of policies. The outcome is defined as follows. In the first offer round, nature selects some player to be the proposer. If the policy selected by the proposer is approved by each member of the coalition selected by the proposer (i.e., belongs to each member's acceptance set), then this policy is accepted on behalf of the group and negotiations are concluded. If some coalition member rejects the proposed policy, then nature randomizes again to select a proposer for the following offer round and the process is repeated. If the last round of negotiations is reached without agreement having been reached, then the game ends and the "disagreement" policy is implemented by default. Clearly, the procedure just described defines a random variable that assigns positive probability to a finite number of policies.

Having defined strategies and outcomes, the specification of a multilateral bargaining game is completed by defining a solution concept. The standard solution concept for games of this kind is *subgame perfection*. In the present context, however, this concept has no predictive power: for any game in which at least two players are required for agreement, any policy that is weakly preferred by all players to the default outcome can be implemented with certainty as the outcome of subgame perfect equilibrium. Fortunately, almost all of these equilibria violate a natural rationality criterion and can be eliminated by a number of equilibrium refinements. Rausser and Simon adopt a particularly simple refinement, referred to as the *SEDS criterion* (Sequential Elimination of Dominated Strategies). The criterion first eliminates strategies that involve inadmissible (i.e., weakly dominated) play in the final response round.[1] Next, it eliminates strategies that involve inadmissible play in the penultimate round, considering only strategies that survive the first round of elimination. And so on. A profile of strategies that survives this sequence of eliminations is called an *equilibrium* for the game.

There is a simple characterization of the set of equilibrium strategy profiles—in each response round, a player will accept a proposed policy if and only if it generates at least as much utility as her *reservation utility* in that round. That is, the utility she expects to receive if no agreement is reached and play continues into the following round. In each offer round, a player is faced with a two-part problem. For each admissible coalition, she maximizes her utility over the set of policies that provide coalition members with at least their reservation utilities.[2] She then selects a utility-maximal policy from among these maximizers. Under weak conditions, an equilibrium always exists and, moreover, the equilibrium outcomes are generically unique.[3]

A *multilateral bargaining model* is a sequence of multilateral bargaining games, which are all identical except for the number of negotiating rounds, which increases without bound as the sequence progresses. A *solution* to a multilateral bargaining model is any limit of a sequence of equilibrium outcomes for the games in the sequence. A solution will be called *deterministic* if the elements of the limit outcome vector are all identical. Solutions that are not deterministic will be called *stochastic*. When a solution exists, it is interpreted as a proxy for the equilibrium outcome of a bargaining game in which the number of negotiation rounds is finite but arbitrarily large.

Rausser and Simon identify two sets of sufficient conditions for existence of a deterministic solution. The first is that the space of policies for the underlying problem is one-dimensional and that decisions require the consent of a simple majority of the players. When the policy space is multidimensional, it is much more difficult to guarantee the existence of a solution. One relatively straightforward restriction is that there is at least one *essential player*, i.e., a player who is a member of every admissible coalition. For every bargaining problem satisfying this restriction, the derived bargaining model has a deterministic solution.

In the abstract, the latter sufficiency condition is quite restrictive. For example, it clearly conflicts with the formal institutional procedure of decision making by majority rule. However, in a wide variety of collective decision-making contexts, the condition is routinely satisfied. For example, when unanimous agreement is required, each player is essential. In the case of the California water policy debate, it is generally considered that each of the major "players" - agricultural water users, urban water users, and environmental groups - has a *de facto* veto over any major policy reform. Thus, in this context, the proper set of admissible coalitions clearly consists of only unanimous, or quasi-unanimous, coalitions.[4]

To provide some intuition about the internal workings of the model, we will briefly describe a simple example. The example belongs to the class of problems known as *spatial problems*, in which the policy space consists of alternative *locations*. For example, a location could be a site for a public good or, more abstractly, the attributes of some candidate for some office. Each player has a most preferred location, called her *ideal point*. The utility that a player derives from a particular location is a decreasing function of the distance between this location and the player's ideal point. This example shows why a deterministic solution must exist whenever there exists at least one essential player. There are three players. The space of possible locations is represented by the two-dimensional unit simplex, representing a two dimensional space of possible attributes (see Figure 1). Players' ideal points are at the vertices of the triangle. We assume each player is essential, so that the only admissible coalition is the coalition consisting of all three players. Player *i's* access probability will be

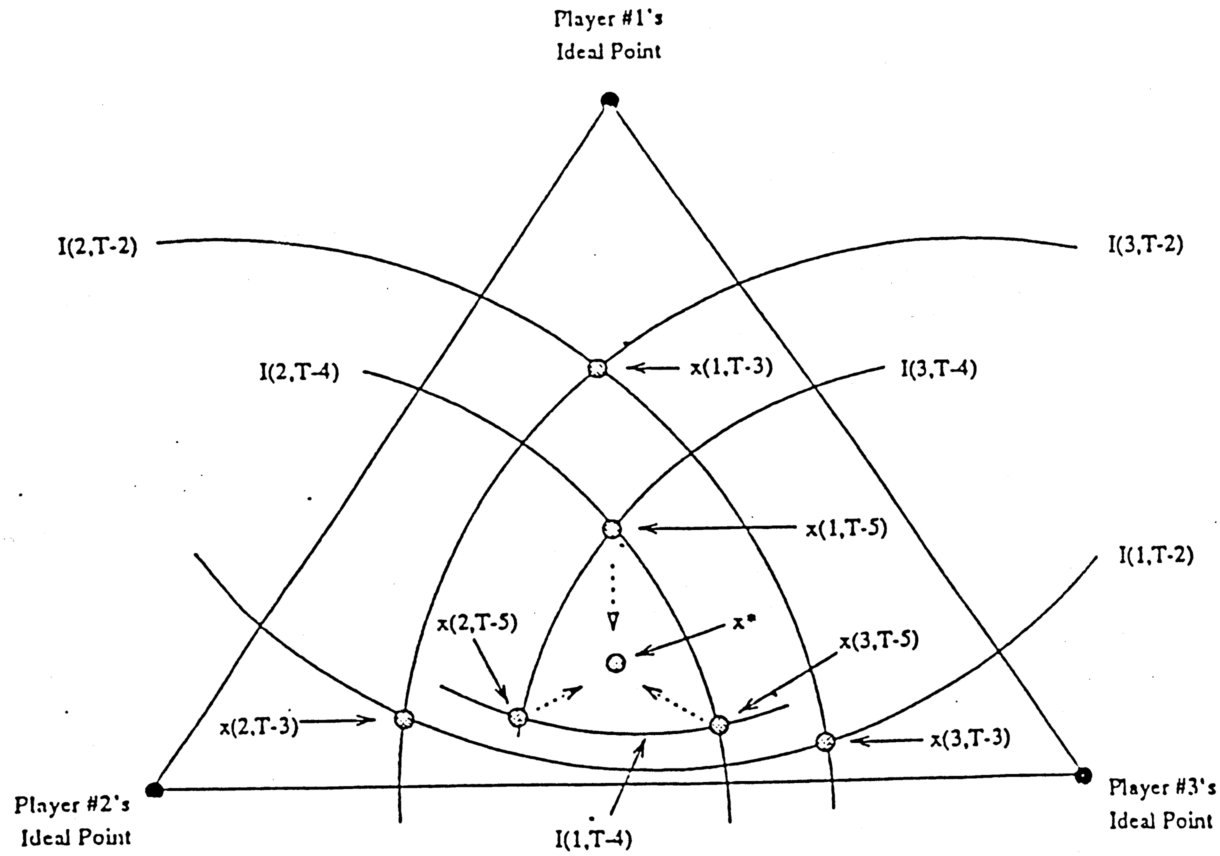

Figure 1: A T-round Three Player Game

denoted by w_i. We will assume that players #2 and #3 are equally powerful, but player #1 is less powerful. Reflecting these assumptions, we set $0 < w_1 < w_2 = w_3$.

We assume that if players fail to agree on a location by round T of the game, the disagreement policy that results yields all players sufficiently low payoffs that any location in the simplex will be considered preferable by all to disagreement. Under this assumption, each player will propose her own ideal point if selected by nature to be the proposer in round T-1, and this proposal will be accepted by the other players in round T. Thus, conditional on entering round T-1, each player faces the following lottery over locations: for each player i, players i's ideal point will be chosen with probability w_i. The line $I(i,T\text{-}2)$ in Figure 1 is the indifference curve corresponding to player i's *reservation utility* in round T-2: any proposal on this line yields player i a utility level that is the certainty equivalent of the lottery she would face if she were to reject the proposal. The line represents player i's *participation constraint* on negotiations in round T-3. Observe that since #1's access probability is lower, the distance between $I(1,T\text{-}2)$ and #1's ideal point is greater than for the other two players.

In round T-3, the set of proposals for which unanimous agreement can be obtained is the region bounded by players' participation constraints. If selected to be the proposer in this round, each player will propose the point in this region closest to her ideal point. Player i's choice is indicated by $x(i,T\text{-}3)$. Thus, the outcome conditional on reaching round T-3 is that $x(i,T\text{-}3)$ will be agreed on with probability w_i. Now consider round T-4. If player i rejects a proposal in this round, she will receive *at least* her reservation utility for round T-2, and strictly more than this with probability w_i. Hence her participation constraint in round T-5 is strictly tighter than in round T-3, and the set of proposals for which agreement can be obtained is strictly smaller. Player i's choice in this round is indicated by $x(i,T\text{-}5)$. Proceeding by backward induction, the limit of such sets is the point denoted by x^*. Hence, if T is sufficiently large, all of the locations proposed in the first round of negotiations will be arbitrarily close to x^*. It follows that in the limit, x^* will be implemented with probability one.

3. The California Water Policy Negotiations

In order to demonstrate how the multilateral bargaining model can be used to analyze complex policy reform processes, we investigate the current water policy negotiations in California. Disputes over the water resources of Western States are well known to anyone familiar with natural resource issues; the contentiousness and intractability of these conflicts is legendary. This is particularly true in California, where a large agricultural industry, a large and rapidly expanding urban population, and a vocal and influential environmental movement have engaged in a constant and increasingly confrontational struggle over water policy issues. In this atmosphere, water policy has become a legal and political battleground.

For the past 3 years, however, a series of unique negotiations have been taking place between the traditionally warring factions on California water policy issues. Representatives from agricultural water agencies, urban water agencies and environmental groups have been meeting on a regular basis to try a forge a consensus based solution to the heated and often acrimonious fight over water policy in California. These negotiations, known popularly as the "Three Way Negotiations", have been occurring outside the context of any specific legislative, regulatory or judicial proceeding and are not sponsored or affiliated with any governmental agency. The negotiations are perhaps best characterized as informal meetings between influential, non-governmental, participants in California water issues. The aim of the negotiations is to break the current policy gridlock by identifying areas of common ground between the participants and formulating innovative policy alternatives that meet these common goals.

A recent detailed case study of these negotiations (Adams) indicates that the structure of the negotiation process itself, as well as the participants and issues involved, has been influential in determining the progress and outcomes of the negotiations. The study also reveals that the participants have spent a good deal of time bargaining over the environment (the formal and informal rules) in which the actual negotiations take place. Drawing from this study, we employ the model described above to investigate several issues that have arisen regarding the structure of the negotiations and the optimal strategies therein. The substantive players and issues are briefly outlined below. Readers interested in a more through treatment of the Three Way Negotiations are referred to the study by Adams.

Major issues that have arisen in the negotiations include the degree to which water will be transferable, the type and level of environmental standards that impact water use, and new infrastructure development. Interest groups active in the negotiations include agricultural water users, urban water users, and environmentalists, and, to a lesser degree, commercial fishermen, and state and federal governments.

Agricultural users currently consume about 85% of the state's developed water supply. These users have historically benefitted from large subsidies in the form of low water prices and have been relatively free of stringent environmental regulations. As a result, this group is content with the status quo and strongly opposes radical reform. It is important to recognize, however, that there is considerable heterogeneity among the members of this group. Important aspects of this heterogeneity include the type of water right the individual farmer holds (surface water, groundwater, date of priority and quality of water) as well as the individuals' location. This heterogeneity will determine whether a farmer may profit from policy reforms such as the transition to a system of marketable water rights, and thus, whether a farmer is likely to support or oppose such a transition.

Urban water users are primarily concerned with the availability of affordable water supplies to support continued urban growth. The value of water, measured as willingness to pay, is much higher in urban use than in agriculture. This group views water markets as the best method of achieving urban water availability. Consequently, urban water user groups are the strongest supporters of unrestricted water markets. Urban users also support new infrastructure development. While urban water users generally oppose strong environmental regulations regarding water use, the high value of water in urban uses tempers this opposition. Urban users are more willing, and able, to pay the costs of meeting environmental standards than are their agricultural counterparts.

Environmental interest groups are primarily concerned with controlling adverse environmental consequences of water use patterns. As such, strong environmental regulations are the primary negotiating objective of this group. Environmentalists also strongly oppose new infrastructure development. Environmental groups have mixed motives with respect to water markets. Transferring water from agricultural to urban use may reduce instream flows and eliminate many incidental wetlands which serve as wildlife habitat. On the other hand, water markets are viewed as an effective method of meeting increasing urban demands without new project development, and transferrable water rights may allow environmental groups to acquire water for environmental purposes. Thus, environmental groups conditionally support water markets but advocate strong environmental constraints on transfers.

The above listing of interests and issues is not meant to be exhaustive but, instead, to represent the major players and policies that have emerged in the negotiations. Analysis of such complex processes has traditionally been beyond the reach of game theoretic models. The model presented in Section 2, however, can yield valuable insights into the process and outcome of such negotiations. To parameterize the model it is necessary to specify the player set, the issue space over which these players negotiate, as well as the preferences of each player over the issue space. The player set will consist of some or all of the interest groups enumerated above, while the issue space will consist of the issues mentioned above. Player preferences can be inferred either from interest group behavior (such as proposals that are actually submitted by each group) or through direct interview procedures. The likely results of the negotiations can then be determined under a variety of constitutional settings.

4. Illustrative Applications

Here we present some illustrative applications of the multilateral bargaining simulation model. These applications focus on several key areas of the institutional structure in order to illustrate specific aspects of the multi-issue, multilateral negotiation process. Naturally, the examples abstract from the complexity of the general negotiation process, in order to focus attention on the major issues and interests that are central to the policy debate. These examples also highlight the delicate and intricate interrelationships that characterize the type of interactions that take place in complex negotiation processes.

We begin by discussing two experiments in detail, and then comment briefly on a number of other experiments that were conducted. Each of the experiments investigates the effect of a change in institutional structure on the negotiated outcome. In the first experiment the space of policy issues is varied, while the second experiment is concerned with the implications of heterogeneous interest groups. We first provide a brief description of the general experimental procedure. Each experiment is then motivated and described, and the results of the experiment are presented. We present the results of experiments 1 in an informal and intuitive manner. The results of experiment 2 are analyzed in more formal detail.

Each experiment consists of 25 simulations in which one aspect of the bargaining process - referred to as the target variable - is systematically varied. In each simulation, the parameters defining the players' utility functions are randomly selected from a prespecified intervals. These intervals are chosen with our prior, imprecise knowledge about players preferences. (For example, we know environmental groups generally oppose new infrastructure, but we do not know the precise extent to which they will be willing to trade infrastructure for environmental protection.) If our simulations yield similar results over the entire range of utility function parameters that we consider, we can be relatively confident of the robustness of our results. On the other hand, when observed differences in our results can be traced to differences in one parameter or group of parameters, we may be led to form hypotheses about causal relationships within the bargaining process.

In each simulation, we first solve the bargaining model for the initial setting of the target variable. We then successively increment the value of the target variable, each time re-solving the model. Thus, each simulation consists of a family of bargaining models, all identical except for the values of the target variable. By systematically comparing the solutions to the games in each family, we are able to gain considerable insight into the comparative statics effects of the change in the target variable.

Table 1 illustrates how we translate the policy debate, described in Section 3 into a formal bargaining model. First, in order to focus attention on the major players and issues, we limit the formal game to three players (agricultural water users, urban water users and environmentalists) and three issues (degree of transferability of water rights, degree of environmental protection, and new infrastructure development). Each issue is represented as a dimension of a "policy space". Without loss of generality, each dimension is normalized to the unit interval. A specific proposal or policy is represented as a point in this space and players have indirect utility functions, defined directly on the policy space.

Table 1: Parameter Values for a Typical Experiment 1

	Agricultural users (A)		Urban users (U)		Environmentalists (E)	
Variable	Lower bound	Upper bound	Lower bound	Upper bound	Lower bound	Upper bound
$\beta_{i,1}$	0.90	1.00	0.09	1.00	0.00	0.10
$\beta_{i,2}$	0.25	0.35	0.90	1.00	0.50	0.60
$\beta_{i,3}$	0.00	0.10	0.00	1.00	0.09	1.00
$\gamma_{i,1}$	0.90	1.00	0.90	1.00	0.75	0.85
$\gamma_{i,2}$	0.25	0.35	0.90	1.00	0.50	0.06
$\gamma_{i,3}$	0.75	0.85	0.25	0.35	0.90	1.00
ξ_i	-6.00	1.00	-6.00	1.00	-6.00	1.00
ρ_i	0.50	0.50	0.50	0.50	0.50	0.50

These utilities are constant elasticity of substitution functions of the form:

$$u_i(x) = \left(\sum_{k=1}^{n} \gamma_{i,k} [\theta_i - (x_k - \beta_{i,k})^2]^{\xi_i} \right)^{(1-\rho_i)/\xi_i}$$

where x_k represents the setting of the k'th policy variable. The parameter $\beta_{i,k}$ is interpreted as player *i*'s most preferred setting--or ideal point--for the k'th policy variable, while $\gamma_{i,k}$ reflects the relative weight, or importance, that player i attaches to this variable. The *substitutability coefficient* ξ_i determines the curvature of players' indifference surfaces. Finally, ρ_i is a risk aversion factor. The role of θ_i is to ensure that the term inside the square brackets is always positive.

For example, if player *i* is indifferent between any two policy vectors that are equidistant from her ideal points on dimensions 1 and 2, then $\gamma_{i,1} = \gamma_{i,2}$. If player *i*'s indifference surfaces have very little curvature then ξ_i must be close to (but cannot exceed) unity; if the surfaces bend very sharply at any point, then ξ_i must be close to negative infinity.

Consider the parameter intervals presented in Table 1. Player A represents agricultural water users, Player U represents urban water users, and Player E represents environmentalists. The first dimension of the policy space represents the degree of new infrastructure development, the second dimension represents the degree of transferability and the third dimension represents the degree of environmental protection. We are relatively confident, for instance, that environmental groups prefer high levels of environmental protection, while agricultural and urban interests prefer low levels. Thus the players ideal points along this dimension, $\beta_{i,3}$, are constrained to be randomly drawn from the (relatively tight) intervals of [0.9, 1.0] and [0.0, 0.1] for the environmental player, and the urban and agricultural players, respectively. We also know that, while the most preferred policy setting along this dimension is similar for both the agricultural and the urban interests, the issue is much more important to agricultural interests. That is to say, agricultural water users are less willing (and able) to pay for higher levels of environmental improvements. Thus, the relative weight that agricultural interests attach to this issue, $\gamma_{i,3}$, is higher than the weight that urban water users attach to the issue. Finally, note that the intervals for the flexibility, ξ_i, and risk aversion, ρ_i, parameters are equal for all players. This reflects our lack of knowledge about the relative or absolute magnitude of these parameters for the different interest groups.

Experiment 1: Varying the Space of Polices Under Negotiation

This experiment analyzes the effects of restricting the space of policies open to negotiations. When formal negotiation sessions are scheduled, considerable effort may be expended in advance of the actual negotiations to influence the range of issues subject to negotiations. An interest group will often seek to exclude from the negotiating table issues that it opposes but which other groups support. Similarly, a group may seek to exclude discussion of an issue when it is content with the status quo with regards to this issue, while other groups wish to institute reforms. This is a common aspect of pre-negotiations, that is, the setting of the constitutional rules for the negotiations. This experiment challenges the rationality of agenda-setting maneuvers of this kind. If the issue which one group seeks to exclude from the negotiations is the only issue that another group strongly supports, then there may be scope for mutually beneficial compromise only if this issue is placed on the bargaining table. Otherwise, the opposing interest group may choose not to participate in the negotiations, perceiving that it has nothing to gain, or else may participate but negotiate to a second best solution.

An example of such strategic behavior from the water resource negotiations in California is the issue of new infrastructure development. Agricultural groups and, to a lesser extent, urban groups wish to include new infrastructure development as an issue in the negotiations, while environmental interests have generally opposed negotiations on this issue. Since environmental groups have the political and legal means to block any significant new infrastructure project, omitting this issue from the negotiations is equivalent to enforcing a status quo of no new infrastructure development. Some members of the environmental delegation, however, feel that opposition to all new infrastructure development may be counterproductive. While environmental groups can block new infrastructure projects, agricultural and urban groups have the power to block many of water policy goals of the environmental groups. Given this mutual effective veto over other groups goals, some environmental groups argue that negotiation and limited compromise on the issue of infrastructure may be the best strategy.

In experiment 1 we simulated negotiations between three interest groups---agricultural water users (A), urban water users (U), and environmentalists (E)---over three issues---new infrastructure (x_1), degree of transferability of surface water (x_2), and degree of environmental protection (x_3). Consent of all three players was required for agreement: that is, the unique admissible coalition contained all three players.

The experimental procedure was to successively reduce the admissible range of the infrastructure variable, solving the model each time for the increasingly restricted policy space. Formally, the target variable in the experiment is defined to be the range of admissible values that the infrastructure variable (x_1) can take. Initially, x_1 can take any value in the unit interval, representing all possible levels of new infrastructure development. The upper bound on the range of admissible values is then successively reduced by increments of 0.05.

In summary, the main conclusions of this experiment are:

Result 1: Initially, reducing the upper bound on infrastructure development has no effect *on the outcome of negotiations. Once the bound is sufficiently small, a further reduction increases the level of environmental quality and the utility of the environmentalists, at the expense of the other two groups. Eventually, however, still further reductions reverse these positive effects, reducing the level of environmental quality and the utility of the environmentalists. The utilities of the other two groups continue to fall.*

Several aspects of these results warrant particular attention. First suppose that the only two available options are either to include or to exclude the issue of infrastructure investment in the negotiating process. In this case, all three groups benefit from its inclusion. When infrastructure is included, the environmental group can concede a little on infrastructure in return for concessions by the agricultural group on environmental concerns. When infrastructure is excluded, however, there is no issue in the policy space that the agricultural group particularly favors and potential gains from trade are substantially reduced.

Now suppose that the range of admissible infrastructure values is a variable aspect of the negotiating framework. For example, suppose that when the agenda for the negotiations is specified, opportunities are provided for the discussion of moderate scale infrastructure investments, such as improvements of existing facilities (improvements to bay conveyance facilities), while discussion of projects of a larger scale (new dams) is foreclosed. Our results suggest that excluding from discussion the highest end of the infrastructure range will have no effect on negotiations. The reason is that, because of the nature of the default option in the experiment, the environmentalists would veto any proposals involving very high infrastructure levels, even in the last round of negotiations. Thus, the high end of the range is entirely irrelevant to the negotiation process.[6] Further restricting the range of admissible infrastructure values does affect the negotiated outcome, however.

The basic intuition for these results is quite straightforward. Environmentalists benefit from small reductions in the maximum admissible level of infrastructure development, because these reductions weaken the bargaining positions of the urban and agricultural users. For large reductions, however, the constraint on infrastructure is binding on the environmentalists as well. As the bargaining proceeds, the environmentalists will find themselves at a "corner solution": they would prefer to concede along the infrastructure dimension in exchange for more environmental protection, but are unable to do so because of the exogenously imposed constraint. Gains to trade are sacrificed and all parties are made worse off.

Experiment 2: Coalition Breaking and the Degree of Preference Heterogeneity

Whenever negotiations take place between interest groups, each of which represents a diverse constituency, there will be differences of opinion among the members of each "negotiating team". For example, the group we identified as the "agricultural interest group" is in fact a loosely knit conglomeration of subgroups, each of which has a quite distinct perspective on the water policy debate. To the extent that these subgroups are represented at the bargaining table, it may be more appropriate to view the "agricultural interest group" as an alliance rather than a monolithic entity. Questions immediately arise concerning the relationship between the internal structure of each alliance and its performance within the negotiations.

From a policy-making perspective, it may be critically important to understand the nature of inter-alliance dynamics. For example, if some interest group is perceived as obstructing the reform process, then one strategic option open to policy makers might be to "divide and conquer" the coalition that this group represents. If the above hypothesis is true, then one may wish to structure the negotiations, or design policy proposals, in such a way as to exacerbate existing tensions within the coalition. As we shall see, our results call the wisdom of such a strategy into question.

In this experiment, we investigate the connection between the performance of the alliance and the degree of congruence between its members' positions on the various issues. Specifically, we test the natural hypothesis that the alliance will perform more effectively, the more homogenous are the preferences of its members. We focus on transferability, an issues on which different agricultural interests take widely diverging positions. Users with valuable surface water rights are expected to be more supportive of transferability than those without such rights. Since agricultural groups as a whole have generally opposed increased transferability of water rights, identifying those agricultural water users who will most benefit from increased transferability may be a productive strategy for those interest groups favoring more liberal transfer policies.

To simplify the experiment, we assume that the agricultural alliance consists of two subgroups, A and B, each of which is represented by a distinct player in the game. A natural measure of homogeneity is the euclidean distance between each players ideal point in the policy space. The experiment involves successively moving A's ideal point further away from B's.

In the experiment there are four players and two admissible coalitions. Player E is the environmental interest group, player U is the urban user group, and A and B are the agricultural users. The two admissible coalitions consist of players E, U and either player A or player B. That is, implementation of an agreement requires "quasi-unanimous" approval. The utility functions for players A and B are identical except for the locations of their ideal points. The intervals from which parameter values were randomly generated are specified in Table 1.

We consider the comparative statics effect of increasing the variable $\beta_{B,2}$, which is player B's ideal point along the transferability axis. Note that in all the experiments, the corresponding variable for player A, $\beta_{A,2}$, is held constant and never exceeds $\beta_{B,2}$. The interpretation is that player B stands to gain more from the formation of a market for water transfers.

Two versions of the experiment were considered. The only difference is the relationship between the variables $\beta_{A,1}$ and $\beta_{B,2}$; i.e., the two players' ideal points along the infrastructure axis. In version (i), we restrict $\beta_{A,1}$ to be slightly greater than $\beta_{B,1}$. In version (ii), $\beta_{A,1}$ is slightly less than $\beta_{B,1}$. Either version seems equally plausible as a description of reality. In each case, player B is presumed to represent the subgroup of agricultural users who are potential suppliers within the proposed water market, i.e., those with superior endowments of historical water rights.

In version (i), ($\beta_{A,1} > \beta_{B,1}$ while $\beta_{A,2} < \beta_{B,2}$) B's weaker preference for infrastructure might be due to concerns that infrastructure development would increase aggregate water supply and dilute the potential rents to the owners of the existing supply. In version (ii), ($\beta_{A,1} < \beta_{B,1}$ and $\beta_{A,2} < \beta_{B,2}$) B's greater preference for infrastructure might be due to cost concerns: infrastructure development in the form of building new canals and lining old ones would reduce the unit cost of transporting water and hence increase profitability.

Intuitively, the distinction between the parameter values in the two versions seems insignificant. In either case, the increase in $\beta_{B,2}$ increases the euclidean distance between A's and B's ideal points, and would seem to drive a wedge between the interests of the two agricultural subgroups. Since these two subgroups compete with each other to represent the agricultural interests, differentiating their interests would appear to weaken the bargaining power of each of them, and so shift the negotiated solution in a direction that benefitted the other two groups at the expense of the agricultural alliance. In fact, the two versions yield diametrically opposite results.

In version (i), the effect of increasing $\beta_{B,2}$ is to increase all three of the policy variables in virtually every element of the sample. The effect on players' utilities is not quite as clear cut: generally, the effect of the change is to decrease the utilities of player A and increase the utilities of players E and U. (For player B, comparisons of utility levels are meaningless because B's utility function is actually changing as the experiment proceeds.) In version (ii) all these signs are reversed: all three variables policy variables fall as $\beta_{B,2}$ increases; the utility of player A increases and those of players E and U decrease.

The major results are:

Result 2: If player A prefers less infrastructure than player B, an increase in player B's preference for transferability leads players A and B to act in ways that are increasingly congruent, in spite of the fact that their ideal points are further apart. That is, as player B's preference for transferability increases, the offers proposed by players A and B move closer and closer together in almost every round of the negotiations. This increased cohesiveness within the "agricultural alliance" results in changes in the policy variables that benefit both members of the alliance. On the other hand, if player A prefers more infrastructure than player B, then an increase in player B's preference for transferability leads the two players to act in ways that are less congruent, resulting in a degradation of the alliance's performance.

The explanation for this result is that in version (ii), the increase in $\beta_{B,2}$ leads players A and B to act in ways that are increasingly congruent, in spite of the fact that their preferences are becoming increasingly disparate. Specifically, in almost every round of negotiations, the offers proposed by players A and B become closer and closer together as $\beta_{B,2}$ increases. Recall that offers are points in the euclidian 3-space, as are players most preferred points (the vectors β_i). As $\beta_{B,2}$ increases, the most preferred points of players A and B become more distant, their preferences are more dissimilar, yet the optimal negotiation proposals become more similar. Informally, this increased cohesiveness within the "agricultural alliance" leads to changes in the policy variables that benefit both groups. In version (i) on the other hand, the increase in $\beta_{B,2}$ leads players A and B to acts in ways that are increasingly disparate, degrading the performance of both.

The source of this striking disparity becomes readily apparent after inspection of Figure 2 below. A key observation is that in almost all rounds of the negotiations, when players A and B make proposals the participation constraints for players U and E are both binding; i.e. the offers A and B make are constrained by the acceptance sets of both player U and E. Mathematically, each of these constraints is a two-dimensional manifold in $\mathbb{R}^3$, the intersection of both constraints is thus a one-dimensional manifold. Thus, players A and B have only one degree of freedom when optimizing subject to U's and E's constraints: once one of the three policy variables is chosen, the values of the other two are uniquely determined.

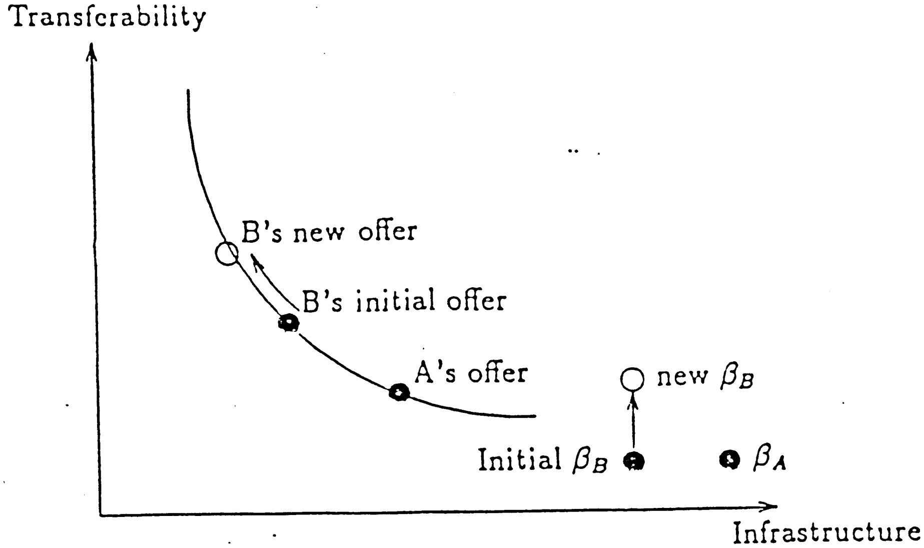

Version (i)

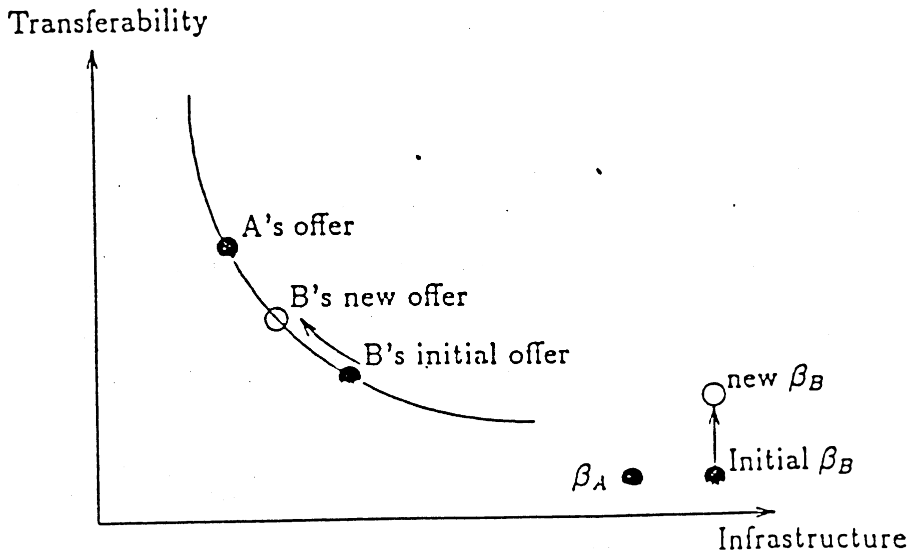

Version (ii)

Figure 2: Relationship between Optimal Offers and Ideal Points in Experiment 2

For a fixed constraint set, let $x_2(x_1)$ and $x_3(x_1)$ denote, respectively, the values of the second variable (transferability) and third variable (environmental quality) once the value of the first variable (infrastructure) is chosen. For the parameter ranges specified in Table 2 above, both $\partial x_2(x_1)/\partial x_1$ and $\partial x_3(x_1)/\partial x_1$ are negative.

Figure 2, illustrates the projection of a typical constraint manifold onto $\mathbb{R}^2$ (the first and second component of x). Moving southeast along the curve, the suppressed values of x_3 increases. For the relevant interval of x_1's, the values of $x_3(\cdot)$ do not exceed the ideal values, $\beta_{A,3}$ and $\beta_{B,3}$, for both types of agricultural user. Thus in each version of the experiment, as either player A or player B moves southeast down the curve, she moves closer to her ideal levels for both infrastructure and transferability, but further away from her ideal point for environmental protection. Each player's utility along the constraint surface is maximized at the point at which the sum of the gains from moving closer to her ideal point in the first two dimensions is just offset by the loss from moving further away from her ideal point in the third dimension.

First consider version (i) of the experiment. In this case, $\beta_{A,1}$ lies to the right of $\beta_{B,1}$ while $\beta_{A,3} \approx \beta_{B,3}$ and, initially, $\beta_{A,2} = \beta_{B,2}$. We have restricted all the other parameters of the two players' utilities to be approximately equal. Thus at any given point along the curve, the marginal gain to player A from moving southeast along the curve must exceed the marginal gain to player B, while the marginal cost is approximately the same. Hence, at the initial value of $\beta_{B,2}$, player A's optimal choice must lie to the southeast of player B's choice. The effect of increasing $\beta_{B,2}$ is to reduce the marginal gain that B obtains from moving southeast along the curve, so that as $\beta_{B,2}$ increases, player B's optimal choice moves to the northwest, i.e., further away from player A's optimum, which, of course, is unaffected by the change in $\beta_{B,2}$.

In version (ii), everything is the same as in version (i) except that $\beta_{A,1}$ lies to the left of $\beta_{B,1}$. Hence, at the beginning of each simulation, when $\beta_{B,2}$ is set at its initial value, the relative positions of the two players' optima are reversed: player B's choice lies to the southeast of player A's. The effect of increasing $\beta_{B,2}$, however, is the same: player B's optimal choice again moves to the northwest. In this case, however, the optimum moves closer to A's choice, which once again remains constant as B's changes. Summarizing, for the particular context in which the two agricultural players find themselves, their actions in this version become more congruent as the disparity between their utility functions increases.

This experiment dramatically demonstrates the complexities of multi-issue, multi-party bargaining. Players' behavior depends on the complex interactions and constraints imposed by both the other players' behavior and the institutional structure under which negotiations take place. In particular, the results of this experiment challenge the following assertion, that at first seems to obvious to question: an alliance will fare more poorly in negotiations as the individual interests of the alliance members become more disparate. Bargaining, like politics, may make for strange bedfellows, and predicting who these bedfellows will be is not a trivial exercise.

Other Experimental Results:

We have preformed several experimental investigations regarding various other aspects of the water policy negotiations in California. While space constraints do not permit a detailed exposition of these experiments, we briefly review some of the results here.

One experiment involved the implications of a change in the disagreement outcome; the policy that would be enforced if no agreement was reached within the negotiation process. This experiment simulated the effect of recent legislative and judicial actions which have dramatically altered the status quo distribution of rights and responsibilities regarding water use in the State. These changes in the status quo have generally favored the position of environmentalists and urban users relative to agricultural users.

As would be expected, this change in the disagreement outcome toward the most preferred point of the environmental interest group benefitted the performance of this group in the negotiations. As the disagreement policy becomes more appealing to the environmentalists, this group becomes more able to credibly commit to abandoning the negotiations unless other groups concede. The negotiated level of environmental protection rises, the level of new infrastructure falls, and the level of transferability rises slightly. The credible threat works to the advantage of the environmentalists and to the detriment of the agricultural and urban water users.

Another experiment involved changing the access probabilities of various players in the negotiations. An increase in a players access probability may reflect increased influence by that player on political decision processes. The inclusion of group representatives on state boards governing resource use, or the election of political candidates supportive of group goals are possible examples of how a groups access probability may increase. Not surprisingly, increasing a groups access probability improves it's performance in the negotiations and leads to negotiated polices closer to the groups ideal position. A less obvious result is that increasing the access probability of one player benefits other players who have similar preferences to that player. Political power may not simply be having influence over policy decisions, but also having preferences similar to those that do have such influence.

A final experiment involved varying the structure of the admissible coalition. This experiment, similar in spirit to experiment 2 above, was also concerned with interest group heterogeneity. As mentioned above, interest groups rarely are homogenous entities; more often they represent a spectrum of agents with similar, although by no means, identical interests. If the legal/institutional environment is such that consent from all interest groups, but not all members of all interest groups, is required, then more heterogenous coalitions would be expected to fare more poorly in the negotiations. Furthermore, as the percentage of interest group members whose consent is required to reach agreement declines, the detrimental effect of group heterogeneity may be expected to increase.

For example, suppose an interest group is comprised of three distinct sub-groups. Under a strict unanimity rule, the consent of all three sub-groups would be required for implementation of any agreement. If the consent of only two of the sub-groups were required, this would correspond to broadening of the list of admissible coalitions; any coalition comprising 2 of the three sub-groups would be admissible. Requiring only one of the sub-groups support would broaden the list of admissible coalitions still.

Broadening the list of admissible coalitions (by requiring support from fewer sub-groups) exacerbates the detrimental effect that preference heterogeneity has on interest group performance. When only one sub-groups support is required, all sub-groups compete with each other for coalition membership. Usually it is the sub-group with preferences most similar to other interest groups that ends up representing the heterogenous group. The utility of the excluded sub-groups suffers as, surprisingly, does the utility of the included sub-group. In competing for coalition membership, all sub-groups modify their negotiation stances to accommodate the views of other interest groups and attract invitation into the ruling coalition. This competition has all sub-groups, including the sub-group eventually included, accepting less from the negotiations than they would under a strict majority rule.

5. Conclusion

Many current natural resource issues are characterized by the type of non-cooperative, multilateral bargaining modelled in this paper. Crafting institutional processes and negotiation strategies to achieve sustainable policy reform amidst these types of conflicts is a daunting task. To date, there has been little formal analysis relating to the design and implementation of these reform processes. We believe the model presented here allows such analysis. The outcome of multidimensional, multilateral negotiations depends crucially on the constitutional structure over which the negotiations take place, as well as the preferences and internal structure of the participating interest groups. These myriad complex factors interact to produce outcomes that are not always *a priori* obvious. Careful modelling of the particular bargaining situation, such as presented here, can help us understand, and hence shape, these processes.

Application of the model to the California water policy *negotiations* yields several insights. The first experiment indicates that environmental objectives may be served by allowing *more* project development. Agricultural interests are often able to block environmental policies to which they are opposed, while environmental groups hold similar veto power over new projects. When this is the case, the required consent of the agricultural coalition for higher environmental standards can only be obtained by consent, on the part of the environmental interest group, to limited new infrastructure development. This analysis suggests that both the environmental and agricultural groups could benefit from "gains from trade" and thus mutual compromise may be able to break the current regulatory deadlock that exists in California water policy.

The second experiment demonstrates the effect that interest group heterogeneity can have on the performance of that group in negotiations. Members of interest groups often have similar, but not identical preferences, particularly when there are several issues under negotiation. If policy reforms can be implemented without the acceptance of all members of a particular group, then targeted proposals by other interest groups may effectively split members of this opposing group. In this case, reform may be achieved without the need to placate the most extreme members of the opposing coalition.

Exploiting heterogeneity in multi-issue, multi-party negotiations can be a delicate task, however, as this experiment indicates. The degree to which intra-coalition heterogeneity will affect the negotiating strategy and performance of a group depends on the relationships between the preferences of the groups members on all issues under negotiation and on the bargaining strategies of other groups. If two sub-groups of a blocking coalition have slightly different preferences over one issue, then targeting proposals which exploit preference heterogeneity among the sub-groups on another issue may actually make the groups adopt negotiating strategies that are more congruent. This experiment demonstrates the complexity of the policy design process and the need for rigorous and detailed modelling of the type presented here.

Footnotes

1. Several other, more sophisticated, equilibrium refinements (such as Myerson's properness or Kohlberg-Merten's strategic stability) would also suffice. On the other hand, the less rigorous approach taken by many authors, of requiring players to accept any offer that is strictly preferred to the status quo, would also suffice.

2. This set is necessarily nonempty if players are risk averse, since each one strictly prefers the expected outcome of the lottery in the following offer round to the lottery itself.

3. Formally, fix a game form and a universe of possible utility functions, endowed with the "sup norm" metric. For an open, dense subset of these functions, the derived multilateral bargaining games have unique equilibrium outcomes.

4. By quasi-unanimous coalition we mean any coalition that has support from at least some (pre-specified) minimum number of the members of all interest groups.

5. While the argument presented below is not a general one, it does indicate the flavor of the general proof.

6. Had we considered a default option that was less favorable to the environmentalists, this result would of course no longer hold.

References

Adams, Gregory D. 1993. "The Three Way Water Agreement Process: A Consensual Approach To Water Policy." Natural Heritage Institute, San Francisco, California

The Economist. "The First Commodity." London and New York. p. 11, March 28-April 3, 1992.

Kohlberg, E. and J. F. Mertens. 1986 "On the Strategic Stability of Equilibria." *Econometrica* Vol. 54, pp. 1003-1038.

Myerson, P. 1978. "Refinements of the Nash Equilibrium Concept." *International Journal of Game Theory* Vol. 7, pp. 73-80.

Rausser, Gordon C., and Leo K. Simon. 1991. "A Noncooperative Model of Collective Decision Making: A Multilateral Bargaining Approach." Working Paper No. 618, Department of Agricultural and Resource Economics, University of California at Berkeley.

Rubinstein, A. 1982. "Perfect Equilibrium in a Bargaining Model." *Econometrica*, Vol. 50, pp. 97-109.

Stahl, I. 1972. "Bargaining Theory." Stockholm: Stockholm School of Economics.

Stahl, I. 1977. "An N-Person Bargaining Game in the Extensive Form." In *Mathematical Economics and Game Theory*, R. Henn and O. Moeschlin, eds. Lecture Notes in Economics and Mathematical Systems No. 141. Berlin: Springer-Verlag.

Appendix

The Multilateral Bargaining Model

A distinction is drawn between multilateral bargaining problems, games and models. A *multilateral bargaining problem* is a game in the sense used by cooperative game theorists. Each bargaining problem gives rise to a family of noncooperative, finite extensive form *multilateral bargaining games* that are identical except for the number of negotiating rounds. A *multilateral bargaining model* consists of a family of T-round bargaining games derived from a common bargaining problem, in which T increases without bound.

The Underlying Multilateral Bargaining Problem

There is a finite set of players, denoted by $I = \{1, \ldots, \bar{i}\}$. The representative player will be denoted by i. The players meet together to select a *policy* from some set X, of *possible alternatives*.

Assumption A1: X is a convex, compact l-dimensional Euclidian space.

If the policy vector x is selected, player i receives the payoff $u_i(x)$.

Assumption A2: For each i, $u_i(\cdot)$ is continuous and strictly concave on X and satisfies the von-Neumann Morgenstern axioms.

Assumption 2 implies players are globally risk averse and payoffs are independent of time. It is straightforward but not particularly illuminating to incorporate time discounting into the model.

To avoid degenerate special cases, a minimal amount of diversity in players preferences is assumed.

Assumption A3: For $i \neq j$, the maximizers of $u_i(\cdot)$ and $u_j(\cdot)$ are distinct.

There is in addition to X a distinguished vector, x^{dflt}, called the *disagreement outcome*. If players cannot reach agreement during the negotiation process, then the vector x^{dflt} will be imposed by default. It is convenient to isolate $\{x^{dflt}\}$ from the set X. For example, we can assign x^{dflt} a payoff of negative infinity without violating continuity.

Once again to avoid special degenerate cases, a negotiable settlement which Pareto dominates the disagreement outcome is assumed to exist.

Assumption A4: There exists $x \in X$ such that for each i, $u_i(x) > u_i(x^{dflt})$.

Denote by X^{*} the set $X \cup \{x^{dflt}\}$. Refer to the vector valued function $\mathbf{u} = (u_i)_{i \in I}$ defined on X^{*} as the *payoff function* for the bargaining problem. (Vectors are denoted by boldface letters throughout the model.) Assuming that all other parameters have been implicitly specified, $\Gamma(\mathbf{u})$ will denote the bargaining problem with payoff function **u**

The specification of a multilateral bargaining problem includes a list of admissible coalitions, $\mathbb{C}$ with representative element C. An admissible coalition is interpreted as a subset of players that can impose a policy decision on the group as a whole. One or more players may belong to *every* admissible coalition; in this case the bargaining problem has an *essential player.*

The T Round Multilateral Bargaining Game

A bargaining game is derived from a bargaining problem by superimposing a "negotiation process". Denote by $\Gamma(\mathbf{u},T)$ the T-round bargaining game derived from $\Gamma(\mathbf{u})$. A distinction is drawn between odd-numbered rounds of negotiations, called *offer rounds*, and even-numbered rounds, called *response rounds*. In offer rounds, players choose *proposals*, consisting of policies from X and coalitions from $\mathbb{C}$. In response rounds, players specify *acceptance sets*, indicating the vectors they will accept if invited to join a coalition in that round.

For $t \in \{1,3,...,T\text{-}1\}$, let $(x_{i,t},C_{i,t})$ denote player i's proposal in round t, and $A_{i,t+1}$ represent her acceptance set in the following response round. Acceptance sets are restricted to be closed. A *strategy* for player i is a collection of proposals and acceptance sets, $s_i = [\ (x_{i,t},C_{i,t})\ ,\ A_{i,t+1}\ ]_{t=1,3,...T\text{-}1}$. Let S_i denote the set of strategies available to player i. A *strategy profile* is a list of strategies, one for each player. Let S denote the list of strategy profiles. A list of strategies for all but one player is called a *subprofile*. Let $S_{-i} = \prod_{j \neq i} S_j$ denote the set of subprofiles that omit player i, with representative element s_{-i}.

Each profile of strategies uniquely defines an *outcome*, which is a random variable defined on $X^{*} = X \cup \{x^{dflt}\}$. Prior to the game, nature selects a *proposer sequence*, a list of players, one for each offer round, denoted by $\iota = (\iota(1), \iota(3),...,\iota(T\text{-}1)) \in I^{T/2}$. For each t, $\iota(t)$ is an i.i.d. random variable, distributed according to the exogenously specified vector of *access probabilities*, $\mathrm{w} = (\mathrm{w}_i)_{i \in I} \gg 0$. thus, proposal sequence ι is selected with probability $\omega(\boldsymbol{\eta}) = w_{\eta(1)} x w_{\eta(3)} x \ldots x w_{\eta(T-1)}$.

The outcome function is a map χ from strategy profiles and "proposer sequences" to policies. Specifically, fix a strategy profile s, where $s_i = (x_{i,t},\ C_{i,t},\ A_{i,t+1})_{t=1,3,...T\text{-}1}$. For each $\iota \in I^{T/2}$, a unique policy $\chi(\iota,s)$ is defined as follows. If the policy $x_{\iota(1),1}$ is an element of $A_{j,2}$ for every j in $C_{\iota(1),1}$, then this vector is *accepted* and negotiations do not proceed beyond the second round. Now suppose that for t $\in$ $\{3,5,...,T\text{-}1\}$, the policies proposed in previous offer rounds have all been rejected. If $x_{\iota(t),t}$ is an element of $A_{j,t+1}$ for every j in coalition $C_{\iota(t),t}$, then this vector is accepted and negotiations do not proceed beyond the $t+1$'th round. If agreement is not reached by round T, then the vector x^{dflt} is selected by default.

The procedure just described defines a finite-support random variable on X^*. Given a profile s, let $Eu_i(s)$ denote player i's expected payoff from the random profile generated by s. That is, $Eu_i(s) = \sum_{\iota \in I(T/2)} \omega(\iota) u_i(\chi(\iota,s))$. Similarly, for $t \in \{3,...,T+1\}$, $Eu_i(s|t)$ denotes player i's expected payoff if the profile s is played out starting from round t. $Eu_i(s|t)$ is referred to as player i's *reservation utility* in round t-1; this is her expected utility conditional on failure to reach agreement in round t-1.

The standard solution concept for games of this type is *subgame perfection*. In the present context, this concept has no predictive power: for any game in which at least two players are required for agreement, any policy that is weakly preferred by all players to the disagreement outcome can be implemented with certainty as a subgame perfect equilibrium. Most of these equilibria violate a natural rationality criterion and can be eliminated by a number of equilibrium refinements, such as Myerson's properness, or by the standard method in the literature of requiring players to accept any offer that is strictly preferred to the status quo.

Results for *T*-Round Bargaining Games

Proposition I: Let $\Gamma(\mathbf{u})$ be a bargaining problem satisfying assumptions A1-A4. Then s is an equilibrium for the bargaining game if and only if for each i and each $t \in \{1,3,...,T\text{-}1\}$:

(i) $A_{i,t+1} = \{x \in X : u_i(x) \geq Eu_i(s|t+2)\}$.

(ii) $x_{i,t} \in A_{j,t}$, for all $j \in C_{i,t}$ and $x_{i,t}$ maximizes $u_i(\cdot)$ on the set $\cup_{C \in \mathcal{C}} \cap_{j \in C} \{x \in X: u_j(x) \geq Eu_j(s|t+2)\}$.

Proposition I says that in each round of the game, after strategies that are inadmissible in later rounds have been eliminated, each player is left with a straightforward decision single-person problem. In a response round, a player will accept a proposed policy if and only if it generates as much utility as her reservation utility in that round. In an offer round, a player is faced with a two part problem. For each admissible coalition, she maximizes her utility subject to the constraint that other coalition members must receive at least their reservation utility in the following response round. She then selects a utility maximal policy from among these maximizers.

An immediate implication of Proposition I is that an equilibrium always exists. For generic problems, the equilibrium outcomes for games derived from these problems are unique. Specifically, let $\tilde{\upsilon}$ denote the set of payoff functions on X satisfying Assumptions A2-A4 and endow $\tilde{\upsilon}$ with the sup norm metric.

Theorem II: Let $\Gamma(\mathbf{u})$ be a bargaining problem satisfying Assumptions A1-A4. Then for every even integer T, the derived game $\Gamma(\mathbf{u},T)$ has an equilibrium. Moreover, there is an open dense subset, $\bar{\upsilon}'$ of $\bar{\upsilon}$ such that for each $\mathbf{u}' \in \bar{\upsilon}'$ and every T, the equilibrium outcome for $\Gamma(\mathbf{u}',T)$ is unique.

The Multilateral Bargaining Model

A *multilateral bargaining model* is a sequence of *multilateral bargaining games*, $\{\Gamma(\mathbf{u},T)\}_{T=2,4...}$ in which T increases without bound. The games in the sequence are all derived from the same underlying bargaining problem; the only difference between them is the number of negotiating rounds. A *solution* is a limit to the sequence of equilibrium outcomes for games in the sequence. It is sufficient to identify the pointwise limits of sequences as *equilibrium outcome vectors*. Specifically, suppose that for $\tau = \{2,4,....\}$, s^{τ} is an equilibrium profile for $\Gamma(\mathbf{u},\tau)$ and that $\bar{x} = (\bar{x}_i)_{i\in I}$ is a pointwise limit of the sequence $(x(s^{\tau}))_{\tau=2,4,...}$, have the following weak-star limit:for each i, $\bar{x}_i$ is realized with probability $\sum_{(j:\, \bar{x}_j = \bar{x}_i)} w_j$.

A solution is called *deterministic* if the limit outcome has singleton support, or, equivalently, if the elements of the limit outcome vector are all identical.Solutions that are not deterministic are called *stochastic*. The policy to which a deterministic solution assigns probability one is referred to as the *solution policy*. When a solution exists, it is interpreted as a proxy for the equilibrium outcome of a bargaining game in which the number of negotiation rounds is finite but arbitrarily large.

In general, a deterministic solution will not exist. One relatively straightforward way to ensure the existence of a deterministic solution is to restrict attention to problems in which there is a least one essential player, i.e., a player who is a member of every admissible coalition.

Theorem V: Let $\Gamma(\mathbf{u})$ be a multilateral bargaining problem satisfying Assumptions A1-A4. If the problem has at least one essential player, then the multilateral bargaining model derived from this problem has a deterministic solution.

Note the Theorem V applies to every problem in which unanimity is required for agreement.