

UC Berkeley

CUDARE Working Papers

Title

A noncooperative model of collective decision making: a multilateral bargaining approach

Permalink

<https://escholarship.org/uc/item/8ch813x0>

Authors

Rausser, Gordon C.
Simon, Leo K.

Publication Date

1992-08-01

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

WORKING PAPER NO. 620

**A NONCOOPERATIVE MODEL OF COLLECTIVE DECISION MAKING:
A MULTILATERAL BARGAINING APPROACH**

by

Gordon C. Rausser and Leo K. Simon

California Agricultural Experiment Station
Giannini Foundation of Agricultural Economics
August 1992

**A NONCOOPERATIVE MODEL OF COLLECTIVE DECISION MAKING:
A MULTILATERAL BARGAINING APPROACH. ***

by

Gordon C. Rausser and Leo K. Simon

Department of Agricultural and Resource Economics
230 Giannini Hall
University of California
Berkeley, CA 94720.

August 14, 1992

* The authors are indebted to Elaine Bennett, Eddie Dekel, Glenn Harrison, Bentley MacLeod, Charles Plott and Robert Powell for helpful conversations. Jeremy de Groot provided invaluable research assistance. Greg Adams valiantly labored through the proofs. Research support from the Office of the United States Trade Representative is gratefully acknowledged.

ABSTRACT

This paper extends the Stahl-Rubinstein model of bilateral bargaining to incorporate many players and multidimensional issue spaces. A central feature of our framework is that in each round of negotiations, a proposer is selected randomly. Our bargaining model consists of a sequence of finite-horizon games, in which the horizon increases without bound. A solution to our model is a limit of equilibrium outcomes for the finite-horizon games. A necessary condition for existence of a deterministic solution is that the limit outcome belongs to the core of the underlying bargaining problem. Solutions, if they exist, are generically unique. Two sets of sufficiency conditions for existence are presented. The paper concludes with examples and applications. In particular, we consider bipolar negotiations between two factions, and show that there is a positive relationship between the cohesiveness of one faction relative to the other and its effectiveness in securing the common goals to its members.

This paper proposes a noncooperative model of multilateral bargaining. Our framework 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."¹ After one player has proposed a division, the other can accept or reject the proposal. If the proposal 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. In Stahl's formulation, the game continues for a finite number of rounds; in Rubinstein's extension, the number of rounds is infinite. We propose a generalization of this framework to incorporate multiple players and multidimensional issue spaces. We consider a sequence of games with finite bargaining horizons, and study the limit points of the equilibrium outcomes as the horizon is extended without bound. Departing from the classical approach, we assume that the proposer is chosen randomly "by nature" in each round of bargaining, according to a prespecified vector of "access probabilities."²

The paper focuses on collective decision-making problems. In contrast to the related political science literature, which explicitly models decision-making in formal institutions such as legislatures and committees, our framework is intended to represent in a very stylized way the informal, unstructured negotiations and debates that frequently precede and accompany the formal legislative process. Consider, for example, the current discussions among the formerly Soviet Republics over the fate of the Soviet Union, or the recent negotiations in Canada leading up to the Meech Lake Accord.³ Alternatively, consider negotiations between regional interests within California over, say, the location of a new hydroelectric facility, or between members of an agricultural cooperative over the location of a new processing plant.

Imagine the activity within the White House staff prior to the selection of a nominee for a senior appointment (such as a Supreme Court judgeship). The following scenario might unfold: a number of different senior staff members, including, perhaps, the President himself, are concurrently lobbying each other, each attempting to build support for one particular candidate; somehow, one of the names under consideration is singled out from the others and, in a plenary meeting of the White House staff, attention is focussed exclusively on this candidate. If sufficient support has been generated, the White House will adopt the candidate as its official nominee. Otherwise, the lobbying process will begin again, until agreement is finally reached.

The formalism of our framework conforms rather closely to this informal negotiating process. One aspect of it, however, is difficult to describe analytically: how is one staffer's proposal singled out from the others? In our framework, this problematic issue is "black boxed:" we simply assert that nature selects a player to be the "proposer" in a random way. Presumably, however, there is some relationship between a staffer's status within the organization and the likelihood that her proposal will be singled out for consideration. We operationalize this relationship by assuming that nature's random choice is governed by an exogenously specified vector of *access probabilities*. Players' access probabilities are interpreted as measures of their relative political "effectiveness:" the higher a player's access weight, the more likely it is that she 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.

The scenario spelled out above is intended to be suggestive, but should not be taken too literally. There are several different stories that are consistent, to some degree, with the framework. Alternatively, the framework can be viewed as reduced form of a complicated structural process. In any case, the ultimate usefulness of the framework will be determined by the intuitive "feel" of its predictions and comparative statics properties rather than the extent to which it faithfully mirrors the details of some actual negotiating process.

The paper is organized as follows. Section 1 introduces the model. The formal presentation is contained in sections 2 and 3. Section 4 contains examples and applications. In Appendix A we discuss an important class of problems to which our theorems do not apply. Proofs are gathered together in Appendix B.

¹ Stahl [1972, 1977] and Rubinstein [1982].

² The idea of a random proposer has been explored in several other papers, including Binmore [1987] and Baron-Ferejohn [1989].

³ In Rausser-Simon [1992], we use the framework developed in this paper as a basis for studying the process of privatization in Eastern Europe.

SECTION 1. INTRODUCTION.

Outline of the Framework.

Our game consists of a finite number of *negotiating rounds*. The purpose of negotiations is to select a *policy* from some set of possible alternatives. In odd-numbered rounds, each player chooses a *proposal*, which is a policy, paired with an *admissible coalition*. Between the odd- and even-numbered rounds, one of the players is selected at random to be the *proposer*, according to the prespecified vector of access probabilities. In the even-numbered rounds, each member of the proposer's coalition decides whether to accept or reject the proposer's policy.⁴ The game ends as soon as all coalition members accept a policy. If one member rejects a proposed policy, the players proceed to the next round. If the last round of the game is reached and the players still fail to agree, then the game ends and a prespecified *disagreement outcome* is implemented.

An important parameter of our framework is the set of admissible coalitions. An admissible coalition is interpreted as a subset of the players that has the authority to impose a policy choice on the whole group. For example, in a *majority rule bargaining game*, a coalition is admissible if and only if it contains a majority of players. More generally, the set of admissible coalitions might have a variety of structures. In particular, we will sometimes impose the restriction that at least one player belongs to *every* admissible coalition. Any such player will be referred to as *essential*.

Our equilibrium concept is a refinement of subgame perfection (Selten [1975]). For a bargaining game with a fixed number of bargaining rounds, an *equilibrium outcome* is a probability distribution over the policies that are implemented when players play equilibrium strategies. A *solution* to our model is a limit of equilibrium outcomes, as the number of negotiating rounds increases without bound. The main results in this paper concern the existence of a *deterministic solution*, which is a limit outcome assigning probability one to a single policy. A necessary condition for a policy to be a solution is that the policy belongs to the *core* of the underlying bargaining game, i.e., there exists no admissible coalition whose members all prefer some other policy. Weak conditions guarantee that if a solution exists, it will be unique for generic specifications of players' preferences. We identify two sets of sufficient conditions for existence. If all players are risk averse, then every majority rule bargaining game with a one-dimensional space of policies has a deterministic solution. Alternatively, a deterministic solution exists in general if at least one player is essential. In particular, the latter restriction is satisfied by *unanimity games*, in which the only admissible coalition is the grand coalition, so that each player is essential.⁵

A striking feature of our framework is that even when the bargaining problem is quite complex, the (generically unique) equilibrium solution can easily be computed numerically. Monte Carlo methods can then be applied to investigate the comparative statics properties of the problem. Specifically, our model is solved recursively, by computing a sequence of solutions to straightforward single-person decision problems, until an acceptable degree of convergence has been obtained. Because of its computational tractability, our framework provides a useful analytical approach to examining a wide variety of collective decision-making problems.

Modeling Issues.

A major modeling issue relates to the sufficiency condition that some player be essential. In the abstract, this 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 satisfied *de facto*, even when it is explicitly violated *de jure*. For example, it is difficult to imagine that a candidate could emerge as the White House nominee for a major political appointment without at least the tacit approval of the President. That is, in negotiations with the White House staff, the President would be an essential player. Similarly, in the current negotiations over the future of the Soviet Union, essential status might be conferred upon either Mr Gorbachev, Mr Yeltsin or both. More generally, whenever a group of negotiators has a clearly identified "leader," it may be appropriate to model this player as essential.⁶ Finally, a player might be deemed essential by

⁴ We could have formulated the framework more sparsely, allowing the proposer's coalition to be determined endogenously. Our reasons for requiring players to specify coalitions explicitly are explained in section 2 below.

⁵ The framework as presently formulated has been stripped to its bare essentials. It can be extended in numerous ways without greatly affecting our major conclusions. For example, if players have positive time-discount factors, our sufficiency conditions will no longer guarantee deterministic solutions, though "almost deterministic" solutions will exist if players are sufficiently patient. Another natural way to extend the framework would be to endogenize the determination of access probabilities, by allowing players to "invest in access" during a Cournot-type pregame. This extension would enhance the realism and applicability of the framework, but at the cost of a considerable loss in computational simplicity.

⁶ Conversely, in the absence of leadership, one might expect certain kinds of negotiations to become bogged down; a formal counterpart of

virtue of her role in executing the decisions resulting from the negotiations. For example, in faculty meetings of university departments, the Chairperson will typically have no special voting privileges. Presumably, however, there are certain kinds of policy decisions that will rarely be taken in the face of explicit opposition from the Chair, as the Chair must bear the ultimate responsibility for implementing the policy.

Confusion frequently arises over the relationship between "essentiality" and "access." While essential players will tend in general to have relatively high access probabilities, there is no *necessary* correlation between these two facets of political power. Czechoslovakia's President Havel provides a striking illustration of the distinction. During his country's velvet revolution, Havel's access weight was very high. In the post-revolutionary era, he has acquired essential status, but his "empirical" access probability has undoubtedly declined. Similarly, while President Reagan was clearly an essential player within the Reagan White House, his "revealed" access probability was quite low in the sense that he rarely initiated policy proposals. More generally, in almost every political process, there are many groups that have considerable access but do not participate in the formal decision-making process. In our framework, these participants would be assigned positive access probabilities but would not belong to any admissible coalition. Familiar examples from national politics are "intellectual lobby groups" such as the Brookings Institution, whose access is derived from its members' individual relationships with policy makers, or, in university politics, radical student groups, whose access might be measured by their ability to influence the general climate of opinion.

A second modeling issue arises from our treatment of the time horizon. Since Rubinstein [1982], it has become customary in bargaining theory to assume that the time horizon is infinite. We depart from this custom, and assume that the bargaining horizon is finite but arbitrarily long. A pragmatic justification for this assumption is that the infinite-horizon version of our model has no predictive power: any outcome can be supported as an equilibrium.⁷ More significantly, there are in some circumstances sound modeling grounds for presuming that the horizon is finite. In collective decision-making contexts, impending deadlines can provide a dramatic impetus to compromise: witness the frequency of last-minute resolutions of Congressional deadlocks, and of post-midnight compromises in wage negotiations when strikes are threatened for the following morning. Since finite horizon models are solved by backward induction, attention is inevitably drawn to these "eleventh hour" effects.⁸ Conversely, of course, in an infinite horizon model there is no endgame. Our final argument in support of a finite horizon is again pragmatic. Our model exhibits certain properties that are intuitively appealing and correspond to well-known stylized facts about actual negotiating situations. (See Rausser-Simon [1991] and section 4 below for a preview.) Whatever the "true" explanations are for these facts, the explanations for the properties of our model that mimic them can be traced to players' behavior in the final rounds of negotiations. Thus, our finite horizon assumption can be justified on the grounds that it captures the spirit of some interesting but not well understood phenomena that might otherwise escape attention.

Related Literature.

Until recently, the topic of multilateral bargaining has received surprisingly little attention by noncooperative game theorists. The few papers that have been written focus almost exclusively on various versions of the alternating-offer model. Binmore [1985] considers several alternative extensions of Rubinstein's analysis to the problem of "three player and three cakes:" each pair of players exercises control over the division of a different cake, only one of which can be divided. In unpublished work,⁹ Shaked observed that in any infinite-horizon, alternating-offer, multilateral pure-division problem, if the consent of three players is required for agreement, and if they are not extremely impatient, then any division of the pie can be implemented by subgame perfect equilibrium strategies. The proof follows easily from the following observation: suppose one player proposes an off-the-equilibrium-path division that gives her a positive share of the pie. If players are not too impatient, then at least one of the other two players can be induced to reject this division by the promise of the whole pie in the subgame that will follow if she does so.

this tendency would be an existence failure in our framework.

⁷ The proof of this assertion is sketched in the "Related Literature" subsection below.

⁸ In general, the profession is justifiably skeptical of arguments that involve long and intricate inductive chains. In many instances, however, the problem is mitigated somewhat in our framework because the basic "shape" of the solution is more or less determined after only a few rounds of induction. (This will become clear when we discuss examples in Section 4 below.) This fact may also reassure experimentalists, since there is abundant evidence that experimental subjects seem unable to backward induct much beyond three periods. (See Neelin et al. [1988] and Spiegel et al. [1990]. See, however, Harrison-McCabe [1992] for a dissenting opinion, and Harrison [1991] for a survey.)

⁹ Shaked's result is discussed in Sutton [1986] and Osborne-Rubinstein [1990].

An interesting variant of the alternating-offer model, called the "Proposal-Making Model," has been advanced by Selten [1981]. A player is selected by nature to make the first proposal. She proposes a utility vector, a coalition and a "responder." The responder either accepts or rejects. If she rejects, the responder then proposes a new utility vector, a new coalition and a new responder. If she accepts, the responder designates another member of the coalition as the next responder, and so on until all members of a coalition have agreed to some proposal. This model has been studied extensively in Chatterjee et. al. [1987] and by Bennett and coauthors.¹⁰

Baron-Ferejohn [1989] study a symmetric problem in which n players must divide up a pie, using majority rule. One variant of their model is strikingly similar to ours, yet draws quite different conclusions; players propose divisions of the pie in odd-numbered rounds; nature chooses one of the proposals at random and voting follows in even-numbered rounds. In the two-round version of this model, each proposer keeps slightly more than half of the pie for herself, and distributes a small portion to enough others to obtain a majority vote. In the infinite-horizon version of the game, as usual, virtually any division can be supported as an equilibrium. The two-period outcome, however, is identified as the unique outcome that can be supported by stationary strategies.¹¹ The most important difference between our framework and theirs is that we focus on the limit of finite-horizon outcomes. For generic specifications of players' utilities, the problem posed by Baron-Ferejohn has no solution in our model.¹²

SECTION 2. THE T -ROUND MULTILATERAL BARGAINING GAME.

In our formal presentation, we distinguish between multilateral bargaining problems, games and models. A *multilateral bargaining problem* is, essentially, 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 sequence 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 $I = \{1, \dots, \bar{i}\}$. 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 subset of l -dimensional Euclidean space.

If the policy vector x is selected, player i receives the payoff $u_i(x)$. Of the assumptions we impose on u_i , the only significant one is strict concavity (i.e., players are globally risk averse). In particular, we assume that payoffs are independent of time.¹³

Assumption A2: For each i , $u_i(\cdot)$ is continuous and strictly concave on X^* and satisfies the von-Neumann Morgenstern axioms.¹⁴

To avoid degenerate special cases, we assume that there is a minimal amount of diversity in players' preferences:

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

There is in addition to X a distinguished vector, x^{dfu} , which is called the *disagreement outcome*.¹⁵ If players

¹⁰ Bennett [1991a, 1991b], Bennett and van Damme [1991] Bennett and Houba [1991].

¹¹ Baron-Kalai [1991] show that it can also be isolated by invoking a computational simplicity criterion.

¹² Because of its symmetry, their problem is nongeneric.

¹³ It is straightforward but not particularly illuminating to incorporate time-discounting into the model.

¹⁴ For many applications, the requirement of strict concavity is too strong. For example, if X is the unit simplex, representing players' shares of a dollar, then we would naturally want to allow player i to be indifferent between any two share vectors whose i 'th components are the same. To allow for such preferences, we could assume that for each i , there is some subspace X^i of X such that i is indifferent between any two vectors that differ only on $X - X^i$, and globally risk-averse on X^i . All of the results in the paper hold if Assumption A2 is weakened in this way.

¹⁵ It is convenient to isolate $\{x^{dfu}\}$ from the set X . For example, we can assign x^{dfu} a payoff of negative infinity without violating con-

cannot reach an agreement during the negotiation process then the vector x^{dfit} will be imposed by default. Once again we avoid degenerate special cases by assuming that there is some a negotiable settlement which Pareto dominates the disagreement outcome:

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

Denote by X^* the set $X \cup \{x^{dfit}\}$. We will refer to the vector-valued function, $\mathbf{u} = (u_i)_{i \in I}$ defined on X^* as the *payoff function* for the problem. (Throughout the paper, vectors will be denoted by boldface letters.) Assuming that all other parameters have implicitly been specified, we will denote by $\Gamma(\mathbf{u})$ the bargaining problem with payoff function \mathbf{u} .

The examples discussed in this paper all belong to a class known as *spatial problems*, in which the policy space, X , consists of alternative *locations*. For example, a location could be a site for a public good. More abstractly, a location could be a point in characteristics space, representing, perhaps, the attributes of a candidate for some office. Each player has a most preferred location in X , called her *ideal point*. The vector of players' ideal points will be denoted by $\alpha = (\alpha_i)_{i \in I}$. Letting $d(x, y)$ denote the Euclidean distance between x and y , player i 's utility from a policy x will be a declining function of $d(x, \alpha_i)$.¹⁶ In all of the computational examples, player i 's utility function is assumed to be of the form:

$$u_i(x) = (\gamma_i - d(x, \alpha_i))^{1-\rho_i}; \quad u_i(x^{dfit}) = -\infty, \quad (2.1)$$

where γ_i is a positive constant and $\rho_i \in (0, 1)$ is player i 's *risk aversion coefficient*.

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 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. More generally, the set of admissible coalitions may have a variety of structures. In particular, we will sometimes impose the restriction that one or more players belongs to *every* admissible coalition. In this case, we shall say that the bargaining problem has an *essential player*.

The *core* of a multilateral bargaining problem is defined in the usual way. A policy x can be *blocked* by a coalition C if there exists an alternative policy y such that each member of C strictly prefers y to x . The core is the set of policies that cannot be blocked by any admissible coalition.

The T Round Multilateral Bargaining Game.

A bargaining game is derived from a bargaining problem by superimposing upon it a "negotiation process." We will denote by $\Gamma(\mathbf{u}, T)$ the T -round bargaining game derived from $\Gamma(\mathbf{u})$. We distinguish 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, they specify *acceptance sets*, indicating which vectors they will accept if invited to join a coalition in that round.

For $t \in \{1, 3, \dots, T-1\}$, let $(x_{i,t}, C_{i,t})$ denote player i 's proposal in offer round t , and $A_{i,t+1}$ represent her acceptance set in the following response round. We impose the restriction that acceptance sets must be closed. A *strategy* for player i is a collection of proposals and acceptance sets, $s_i = \left\{ (x_{i,t}, C_{i,t}), A_{i,t+1} \right\}_{t=1,3,\dots,T-1}$. Let S_i denote the set of strategies available to player i . For expositional purposes, we restrict strategies to be *history independent*. That is, players' decisions in round t are independent of the history of moves by nature, and of the history of proposals rejected in previous rounds. As will become apparent below, for generic specifications of players' payoffs this is no more than a notational convenience; for these specifications, all of our results are unchanged, and their proofs are identical, when strategies are history-dependent¹⁷ Of greater consequence is our requirement that acceptance sets can be conditioned neither on the identity of the proposer nor on the composition of the proposed coalition.¹⁸ Both restrictions could be relaxed without affecting the main results, although certain

tinuity.

¹⁶ If $X \subset \mathbb{R}^l$, then $d(x, y) = \left(\sum_{k=1}^l (x_k - y_k)^2 \right)^{1/2}$.

¹⁷ Of course, this statement would not be true if information were incomplete, in which case, information could be revealed as histories unfolded.

¹⁸ This last assumption is unlikely to cause serious concern to economists, who tend to insist that the variables in question should not matter. To other social scientists and the world at large, however, this assumption might be regarded as too restrictive. In a model of Middle

equilibrium properties would be affected. A *strategy profile* is a list of strategies, one for each player. Let S denote the set of strategy profiles. A list of strategies for all but one player will be 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 identifies an *outcome*, which is a random variable defined on $X^* = X \cup \{x^{dflt}\}$. The mapping from strategies to outcomes will be referred to as the *outcome function* for the game. In our informal outline of the framework, nature moved between each offer and response round. From a *formal* standpoint, however, the actual sequencing of nature's moves is immaterial, since players' strategies are independent of these moves. Nature simply selects a *proposer sequence*, which is a list of players, one for each offer round, denoted by $\mathbf{t} = (t(1), t(3), \dots, t(T-1)) \in I^{T/2}$. An heuristic interpretation of \mathbf{t} is that for $t \in \{1, 3, \dots, T-1\}$, if negotiations have not already been concluded by the time round t is reached, nature declares that player $t(t)$'s round t proposal will be voted upon in round $t+1$ by the coalition she specifies. For each t , $t(t)$ is an i.i.d. random variable, distributed according to the exogenously specified vector of *access probabilities*, $\mathbf{w} = (w_i)_{i \in I} \gg 0$. (Recall that the magnitude of w_i is interpreted as a measure of player i 's relative "political" or "bargaining power.") Thus, the proposal sequence \mathbf{t} is selected with probability $\omega(\mathbf{t}) = w_{t(1)} \times w_{t(3)} \times \dots \times w_{t(T-1)}$.

The outcome function is a map χ from strategy profiles and "proposer sequences" to policies. Specifically, fix a strategy profile s , where $s_i = (x_{i,t}, C_{i,t}, A_{i,t+1})_{t=1,3,\dots,T-1}$. For each $\mathbf{t} \in I^{T/2}$, a unique policy $\chi(\mathbf{t}, s)$ is defined as follows. If the policy $x_{t(1),1}$ is an element of $A_{j,2}$, for every j in $C_{t(1),1}$, then this vector is *accepted* and negotiations do not proceed beyond the second round. Now suppose that for $t \in \{3, 5, \dots, T-1\}$, the policies proposed in previous offer rounds have all been rejected. If $x_{t(t),t}$ is an element of $A_{j,t+1}$, for every j in the coalition $C_{t(t),t}$, 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.¹⁹

The procedure just described defines a finite-support random variable on X^* . Given a profile s , we denote by $Eu_i(s)$ player i 's expected payoff from the random profile generated by s . That is, $Eu_i(s) = \sum_{\mathbf{t} \in I^{T/2}} \omega(\mathbf{t}) u_i(\chi(\mathbf{t}, s))$.

Similarly, for $t \in \{3, \dots, T+1\}$, $Eu_i(s|t)$ denotes player i 's expected payoff if the profile s is played out starting from round t . We will refer to $Eu_i(s|t)$ as player i 's *reservation utility in round $t-1$* , since this is indeed her expected utility conditional on failure to reach agreement in round $t-1$.

The standard solution concept for these games is *subgame perfection*. Informally, a strategy profile s is subgame perfect if starting from each round of the game, the remaining portion of s_i is optimal for player i , given that players other than i are playing the remaining portions of s_{-i} . 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 default outcome can be implemented with certainty as a subgame perfect equilibrium outcome. For example, the following strategies implement the policy x with certainty. In each offer round, each player proposes x and an arbitrary coalition; in each response round, each player accepts x and no other policy. If x is preferred by all players to x^{dflt} , then these strategies are clearly subgame perfect and x is implemented with probability one.

The equilibria just described violate a natural rationality criterion and can be eliminated by a number of equilibrium refinements. Trembling hand perfection is not sufficiently strong, for the familiar reason that this criterion does not impose sufficient discipline off the equilibrium path. A stronger criterion, such as Myerson's [1978] properness, is needed. In infinite games, however, this criterion involves considerable technicalities.²⁰ To avoid these, we invoke a simpler refinement, which we will call *the SEDS criterion* (Sequential Elimination of Dominated Strategies).²¹ Variants of this criterion are regularly invoked to deal with essentially the same problem as the one that faces us.²²

East negotiations, for example, it would be unfortunate if Israelis were obliged to respond in the same way to any given proposal, regardless of whether it was issued by, say, the U.S. or the P.L.O.

¹⁹ As noted above, there is an equivalent, apparently simpler, specification of the model. Rather than require each player to specify a coalition explicitly, we could endogenize the coalition selection problem by allowing the outcome function to simply count votes. Either way, however, the selection of an optimal coalition is an inescapable task for the proposer, as she solves her maximization problem. Thus, the issue is no more than a notational one and the obvious arguments in favor of explicitness seem to us to justify the additional notational burden.

²⁰ See Simon and Stinchcombe [1991].

²¹ This criterion naturally extends to sequential games the criterion known as Dominance Solvability (Moulin [1979]).

²² See, for example, Baron-Ferejohn [1989], Salant-Goldstein [1990] and Baron-Kalai [1991]. For a rather different application of the same criterion, see Simon-Stinchcombe [1989].

Informally, the procedure begins by eliminating strategies that involve inadmissible (i.e., weakly dominated) play in the final response round. Next, we eliminate strategies that involve inadmissible play in the penultimate round, considering only strategies that survive the first round of elimination. And so on. To define the criterion formally, first declare every strategy for i to be *admissible from round $T+1$* . Now fix $t \leq T$ and assume that for each i , there is an identified set of strategies that are *admissible from round $t+1$* . Define s_i to be *admissible from round t* if (i) it is admissible from round $t+1$ and if there exists no alternative strategy σ_i such that: (ii) σ_i agrees with s_i before t , (iii) σ_i does at least as well as s_i against any subprofile s_{-i} that is admissible from round $t+1$; and (iv) σ_i does strictly better than s_i against some such subprofile. Finally, say that a profile s satisfies the *SEDS criterion* if for each i , s_i is admissible from round one. If s satisfies the SEDS criterion for some bargaining game, we say that s is an *equilibrium* for that game. We will refer to the outcome generated by s as an *equilibrium outcome*.

Results for T -Round Bargaining Games.

Proposition I below characterizes the set of strategy profiles that satisfy the SEDS criterion. Indeed, the characterization theorem provides the basis for our computer algorithm for solving bargaining games numerically. In each round of the game, after strategies that are inadmissible from later rounds have been eliminated, each player is left with a straightforward single-person decision problem. In a 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.²³ In an offer round t , 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 utilities in round $t+1$. She then selects a utility-maximal policy from among these maximizers.

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

- (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 $\bigcup_{C \in \mathcal{C}} \bigcap_{j \in C} \{x \in X : u_j(x) \geq Eu_j(s|t+2)\}$.

The proof of this Proposition depends on two independently useful properties of equilibria, stated in the Lemma below. First, at least two distinct offers are proposed in every offer round. Second, in every offer round there is some policy that yields each player strictly more utility than her reservation utility in the following round.

Lemma I: Let $\Gamma(u)$ be a bargaining problem satisfying Assumptions A1-A4 and let s be an equilibrium for the bargaining game $\Gamma(u, T)$. Then for $t \in \{1, 3, \dots, T-1\}$,

- (a) There are at least two distinct players i and j such that $x_{j,t} \neq x_{i,t}$.
- (b) There exists $x \in X$ such that for all i , $u_i(x) > Eu_i(s|t+2)$.

An obvious corollary of Proposition I (indeed, of Lemma I(b)) is that in every game, agreement is reached immediately with probability one. We can exploit this fact to obtain a convenient, simplified representation of equilibrium outcomes. Given an equilibrium strategy profile s , we denote by $\mathbf{x}(s) = (x_i(s))_{i \in I}$ the vector consisting of the policies proposed by each player in the first round of negotiations. As we have noted, each of these proposals is necessarily accepted. Therefore, $\mathbf{x}(s)$ is a representation of the outcome generated by s . For this reason, we will refer to $\mathbf{x}(s)$ as an *equilibrium outcome vector*. The original outcome can be recovered by combining $\mathbf{x}(s)$ with the access probability vector, w : for each i , $x_i(s)$ is realized with probability $\sum_{\{j : x_j(s) = x_i(s)\}} w_j$.

Another corollary of Proposition I is that we can without loss of generality restrict attention to bargaining problems in which the set of coalitions is minimal in the following sense. Say that a coalition C is *minimal with respect to player i* if there exists no strict subset, C' of C such that the coalition $C' \cup \{i\}$ is admissible.²⁴

²³ By assuming that acceptance sets are closed, we finesse the indeterminacy that arises when a player is indifferent between accepting and rejecting a proposal.

²⁴ This criterion is strictly more stringent than the simpler criterion of (unqualified) minimality, which would be satisfied by any coalition

Corollary I below shows that that player i 's opportunity set is unaffected by the restriction that she must choose only coalitions which are minimal with respect to i . In other words, we lose no generality by assuming that i always chooses coalitions that include herself whenever possible, and exclude as many other players as possible. This fact is of considerable practical value, because when we analyze games numerically, it is obviously important to minimize the number of coalitions for which calculations must be made.

Corollary to Proposition I: Let $\Gamma(\mathbf{u})$ be a bargaining problem satisfying Assumptions A1-A4 and let s be an equilibrium for the bargaining game $\Gamma(\mathbf{u}, T)$. Then there is an equilibrium profile σ for this game which is identical to s with the (possible) exception that in each round, each player i specifies a coalition that is minimal with respect to i .

An immediate implication of Proposition I is that an equilibrium always exists. The pivotal, and by far the most difficult result in the paper is that for generic problems, the equilibrium outcomes for games derived from these problems are unique. Specifically, let \mathbb{W} denote the set of payoff functions on X satisfying Assumptions A2-A4 and endow \mathbb{W} with the sup norm metric.²⁵

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, \mathbb{W}' , of \mathbb{W} such that for each $\mathbf{u}' \in \mathbb{W}'$ and every T , the equilibrium outcome for $\Gamma(\mathbf{u}', T)$ is unique.

The arguments we use to prove uniqueness also imply that in all but exceptional games, all of the above results apply whether or not we restrict strategies to be history independent. The argument is transparent. In each round of any game, players' payoffs and strategic opportunities are independent of anything that has happened in previous rounds. Also, because there is no uncertainty about players' types in the model, there is no payoff-relevant information to be revealed as history unfolds. Now if a player has a unique optimal choice, and this choice is independent of history, the player must act in the same way, regardless of the past history. Finally, in the present context it is generically the case that players' optimal choices are unique in every round.

SECTION 3. THE MULTILATERAL BARGAINING MODEL.

A *multilateral bargaining model* is a sequence of T -round bargaining games, $\{\Gamma(\mathbf{u}, T)\}_{T=2,4,\dots}$ 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.

We define a *solution* to be a limit of a sequence of equilibrium outcomes for the games in the sequence. Since these outcomes are random variables, the natural notion of closeness is the weak-star topology. However, because our sequences of equilibrium outcomes have a special structure, we can simplify matters considerably. It is sufficient simply to identify the pointwise limits of sequences of *equilibrium outcome vectors*. Specifically, suppose that for $\tau = \{2, 4, \dots\}$, s^τ is an equilibrium strategy profile for $\Gamma(\mathbf{u}, \tau)$ and that $\bar{x} = (\bar{x}_i)_{i \in I}$ is a pointwise limit of the sequence $(\mathbf{x}(s^\tau))_{\tau=2,4,\dots}$. We will refer to \bar{x} as a *limit outcome vector*. From our earlier discussion (p. 7), the *outcomes* generated by $(s^\tau)_{\tau=2,4,\dots}$ 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 will be called *deterministic* if the limit outcome has singleton support, or, equivalently, if the elements of the limit outcome vector are all identical. The policy to which a deterministic solution assigns probability one will be referred to as the *solution policy*. 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.

rendered inadmissible by the omission of any player. For example, in a majority rule bargaining problem with five players, the coalition $\{2, 3, 4\}$ is admissible, but is *not* admissible with respect to player #1, since $\{1, 3, 4\}$ is admissible.

²⁵ In the sup norm metric, the distance between two functions is the supremum, taken over all points x in the domain, of the absolute value of the difference between the evaluations of the functions at x .

An *approximate solution* is a sequence of outcome vectors that almost converges. More precisely, the model derived from $\Gamma(\mathbf{u})$ has an ϵ -solution if there exists a policy vector \bar{x} and an even integer T such that for each player i and each even $\tau > T$, the distance between \bar{x}_i the policies proposed by i in the first round of $\Gamma(\mathbf{u}, \tau)$ is no greater than ϵ . Like all approximate equilibrium concepts, the interpretation of approximate solutions in the present context is somewhat problematic from a theoretical standpoint. (For one thing, what constitutes a "good" approximation?) For practical purposes, however, approximate solutions can be virtually as useful as exact solutions as sources of testable hypotheses in the analysis of practical applications. In particular, approximate solutions provide rough-and-ready predictors of the location in policy space of a negotiated agreement. Moreover, since generically this prediction will be unique (more or less), sensible comparative statics questions can be posed. On the other hand, if a "reasonably exact" solution does *not* exist, then the predicted outcome of negotiations will depend in a nontrivial way on the number of negotiation rounds. In this event, little positive or prescriptive significance can be attached to the model's predictions. Nonetheless, existence failures are interesting in the negative sense of indicating inherent instabilities in the negotiating environment. In Appendix A we investigate the kinds of stochastic and approximate solutions that arise in a family of two-dimensional spatial problems.

Results for the Bargaining Model.

A necessary condition for existence of a *deterministic* solution is that the underlying bargaining problem has a nonempty core.

Proposition III: Let $\Gamma(\mathbf{u})$ be a multilateral bargaining problem satisfying assumptions A1-A4. If the multilateral bargaining model derived from this problem has a deterministic solution, then the solution policy belongs to the core of $\Gamma(\mathbf{u})$.

Proof of Proposition III: Assume that x is the solution policy but that there is some policy y and some admissible coalition C such that each member of C strictly prefers y to x . Then there exists $\epsilon > 0$ such that all members of C strictly prefer y to any policy in the ball $B(x, \epsilon)$. For $\tau = 2, 4, \dots$, let s^τ be an equilibrium profile for $\Gamma(\mathbf{u}, \tau)$. For τ sufficiently large, each component of the equilibrium outcome vector $x(s^\tau)$ must be contained in $B(x, \epsilon)$. Thus we have $u_j(y) > u_j(x_j(s^\tau)) > Eu_j(s_j^\tau | 3)$, for every $j \in C$. (The second inequality follows from combining Proposition I(ii) and Lemma I(b).) But this is a contradiction, since by Proposition I(ii), $x_j(s^\tau)$ must be a maximizer of $u_j(\cdot)$ on the set $\bigcap_{j \in C} \{x \in X : u_j(x) \geq Eu_j(s_j^\tau | 3)\}$. \square

Theorems IV and V below 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 are made by majority rule.

Theorem IV: Let $\Gamma(\mathbf{u})$ be a multilateral bargaining model satisfying assumptions A1-A4. If (i) the space of admissible policies, X , is a subset of \mathbf{R}^1 and (ii) a coalition is admissible if and only if it contains strictly more than half of the players in I , then the multilateral bargaining model derived from this problem has a deterministic solution.

When the policy space is multidimensional, it is much more difficult to guarantee convergence. At an abstract level, the task is to identify global stability conditions for a relatively complex, nonlinear stochastic dynamical system. One relatively straightforward way to proceed is to restrict attention to problems in which there is at least one essential player, i.e., a player who is a member of every admissible coalition. The interpretation of this assumption is discussed in detail above (pp. 2-3).

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 that Theorem V is applicable to every problem in which unanimity is required for agreement.

Our final result follows immediately from Theorem II. Solutions, when they exist, are generically unique. Assume that X satisfies assumption A1 and, once again, let \mathbb{W} denote the set of payoff functions on X satisfying Assumptions A2-A4.

Corollary to Theorem II: There is an open, dense subset, \mathbb{W}' , of \mathbb{W} such that for each $u' \in \mathbb{W}'$, if the model derived from the problem $\Gamma(u')$ has a solution, then this solution is unique.

Proof of the Corollary: Suppose that for some $u \in \mathbb{W}$, the model derived from $\Gamma(u)$ has more than one solution. Then necessarily there exists T (in fact, infinitely many T 's) such that the bargaining game $\Gamma(u, T)$ has at least two distinct equilibrium outcomes. But from Theorem II, it follows that the set of all such u 's is contained in the complement of an open, dense subset of \mathbb{W} . \square

Multilateral Bargaining and the Nash Program.

Nash [1953] urged that strategic models and axiomatically derived solution concepts should be studied in conjunction, because "each helps to justify and clarify the other (p. 129)." This dual approach has become known as the "Nash Program." Pursuing this program, Binmore, Rubinstein and Wolinsky [1986] study two strategic models with alternating offers and in each case establish a close relationship between their perfect equilibria and the Nash bargaining solution of the corresponding cooperative game.

More recently, Krishna-Serrano [1991] have extended the Nash Program to the n -player case. Their point of departure is Lensberg's [1988] alternative axiomatization of the multilateral Nash bargaining solution, in which Nash's Independence of Irrelevant Alternatives (IIA) axiom is replaced by Multilateral Stability (MS).²⁶ Oversimplifying slightly, the MS axiom can be paraphrased as follows: if in the solution to a multilateral pie-division problem, player i receives a share of the pie x_i , then in the problem constructed by excluding player i and depleting the total size of the pie by x_i , the remaining players should receive exactly the same portions as they received when i was present. Krishna-Serrano incorporate this axiom into their model in a rather direct way, by allowing individual players to exit from the bargaining table, taking with them the shares of the pie that they have negotiated for themselves. In this way, they are able to reconcile the strategic and axiomatic approaches to multilateral bargaining.

While the two strategic models mentioned above lend plausibility to Nash's axiomatic solution concept, our model presents a challenge to Nash's approach. Specifically, our model violates both the IIA and the MS axioms.²⁷ To see that IIA is violated, compare the solutions to our model when two risk neutral players with equal access are bargaining over the two-dimensional policy spaces $X = \{(x \in \mathbb{R}_+^2: x_1 + x_2 = 1)\}$ versus $Y = \{x \in X: \begin{cases} \epsilon x_1 + x_2 = 2/3(1+\epsilon) & \text{if } x_1 \leq 1/3 \\ x_1 + x_2 = 1 & \text{if } x_1 \geq 1/3 \end{cases}\}$. Assume that in each case, $x^{dfu} = (0, 0)$ while $u_i(x) = x_i$, with $u_i(x^{dfu}) = 0$. By symmetry, the solution in the first case is $(1/2, 1/2)$; for $\epsilon \approx 0$, the solution in the second is approximately $(7/12, 5/12)$. The explanation for the difference is transparent. In the last round of offers, player #2 will propose $(0, 1)$, when the set of alternatives is X , and $(0, 2/3(1+\epsilon))$ when it is Y . Thus in the "eleventh hour" of negotiations, player #2's bargaining position is weaker when bargaining over Y than over X , and this relative weakness is reflected in the corresponding solutions to our model. Thus, in the bargaining environment that we have formulated, the alternatives contained in X but not in Y are by no means strategically irrelevant. It is important to emphasize that the above example has *nothing whatever* to do the fact that players are bargaining over policies rather than utilities: we could, obviously, have defined the spaces X and Y to consist of utility vectors rather than two-dimensional policies.

The violation of MS is of particular interest because, in contrast to the violation of IIA, this axiom is inherently multilateral in nature. Essentially, the MS axiom declares that there can be no "bargaining synergies" between players: the relative bargaining strengths of players j and k must be independent of whether or not player i is present at the bargaining table. In our model, however, such synergies almost always arise, except when the

²⁶ Consider a n -player bargaining problem in which the set of feasible utility vectors is U . Assume that all players receive zero utility in the event of disagreement. The unique utility vector satisfying Nash's four axioms--scale invariance, Pareto optimality, symmetry and independence of irrelevant alternatives--is \bar{u} , defined by $\prod_{i=1}^n \bar{u}_i \geq \prod_{i=1}^n u_i$, for all $u \in U$.

²⁷ The IIA axiom can be loosely paraphrased as follows: suppose that the solution to the bargaining problem is x when the universe of bargaining outcomes is X . For any subset Y that does not contain x , x must be the solution when players are bargaining over $X - Y$.

universe of possible bargaining outcomes is symmetric.²⁸ To see this, consider the three-player unanimity bargaining game in which players are bargaining over the "truncated pie" $V = \{u \in \mathbb{R}_+^3: \sum_{i=1}^3 u_i = 1 \text{ and } u_3 \leq u_1\}$.

As before, assume that the disagreement utility vector is zero and that all three players' bargaining attributes are identical. In this asymmetric problem, player #1 appears to have a "natural ally" in player #3 and this is reflected in the equilibrium outcome of the game. In the last round of offers, player #3 proposes the vector $(\frac{1}{2}, 0, \frac{1}{2})$, favoring #1 at the expense of player #2. Thus, once again, in the "eleventh hour" player #3's presence at the bargaining table places #2 at a strategic disadvantage relative to #1, and this weakness is reflected in the solution to the model.²⁹ On the other hand, if #3 were to leave the bargaining table, along with her equilibrium share of the pie, then our model predicts that players #1 and #2 would equally divide the remainder of the pie.

In the study of multilateral bargaining in collective decision-making environments, it is natural to expect bargaining synergies to arise between different players. What are the sources of bargaining synergies? What compromises will emerge as alliances are forged between parties whose interests are interrelated but not coincident? How effective will these alliances be in furthering the common interests of their members? What is the relationship between the "internal" alignment of interests within a given alliance and its "external" effectiveness as it negotiates with other alliances?³⁰ Since bargaining synergies are axiomatized out of existence by the MS criterion, these questions can *only* be addressed in a model that violates MS.

SECTION 4. APPLICATIONS

The main purpose of this section is to illustrate certain properties of our framework and to indicate some problems to which it might be applied. The discussion in this section will be heuristic and informal. For a more systematic and formal approach to comparative statics issues see Rausser-Simon [1991]. We discuss five classes of bargaining problems, labeled A through E. Problem E is a pure-exchange economy. Problems A through D are spatial problems, in which players' preferences satisfy equation (2.1). For expositional purposes, we shall interpret these problems as political in nature, and describe the players as members of some political party. Players whose ideal points lie to the left (resp. right) of the origin will be referred to as the left-wing (resp. right-wing) faction of the party. Locations further along the horizontal axis from the origin denote more extreme political orientations.

Problem A: $X \subset \mathbb{R}^1$; $2n + 1$ players; majority rule.

In this problem, there is an odd number of players, whose ideal points are located along the real line. A coalition is admissible if it contains at least $n+1$ players. It is straightforward to verify that the core is a singleton set consisting of the median player's ideal point. From theorems II and IV, there is a unique, deterministic solution to the derived model; the solution policy is the unique element of the core. There is a striking resemblance between this result and the familiar "median voter theorem" from the political science literature. This theorem states that in a two-candidate election with a one-dimensional issue space, both candidates will locate at the median voter's ideal point. Our result states that a committee consisting of the voters themselves will select the same point.

Problem B: $X \subset \mathbb{R}^1$; $2n$ players; strict majority rule.

This problem is identical to the previous one, except that there is an even number of players, of which strictly more than half are required for agreement. Once again, theorem IV guarantees the existence of a deterministic solution. In this case, the core of the underlying problem is the segment of the real line joining the two median players' ideal points and the solution policy can lie anywhere along this segment. In contrast to the preceding problem, the solution policy is sensitive to all of the parameters in the model, so that interesting comparative statics issues do arise. We will discuss one of the more subtle issues in some detail.

²⁸ A set $U \subset \mathbb{R}_+^n$ is symmetric if for every $u \in U$, and every v obtained by permuting the order of the elements in u then $v \in U$.

²⁹ The properties of the model we have been considering in this and the preceding example are clearly driven by the finiteness of the bargaining horizon.

³⁰ Some of these questions will be addressed in section 4 below, when we consider particular examples. They are the primary focus of Rausser-Simon [1991].

The effect of a shift to the right in the ideal point of the most right-wing player is investigated. Intuitively, this shift can be interpreted as an increase in political extremism. To simplify the analysis, we will impose the following restrictions on the parameter set: (i) all players are equally risk averse; (ii) all members of the same faction have the same access probabilities; (iii) player's ideal points are symmetrically distributed about the origin; (iv) players' ideal points are all distinct. We assign labels to players so that their ideal points are monotone increasing. Thus restriction (iii) states that for $1 \leq i \leq n$, $\alpha_{2n+1-i} = -\alpha_i > 0$.

The increase in α_{2n} has two effects, which we will call the *access effect* and the *risk aversion effect*. The access effect benefits the faction that has greater access; the risk aversion effect benefits the faction containing the extreme player whose ideal point has shifted. For utility functions that satisfy equation (2.1), the latter effect is always very weak relative to the former. Hence if the left-wing of the party has even slightly more access than the the right wing, the solution policy will shift to the left. If the distribution of access is virtually uniform, however, the policy will shift to the right.

The reasoning outlined below applies to any problem in the class identified above. For expositional purposes, however, we will present the arguments in the context of a pair of numerical examples illustrating the two effects. The examples both have six players. In case (i), access is uniformly distributed; in case (ii), it is skewed in favor of the left-wing. The parameters for the simulations are displayed below.

Problem B: Shifting Player #6's Ideal Point to the Right (See Tables 4B).						
	Player #1	Player #2	Player #3	Player #4	Player #5	Player #6
Initial Location:	$\alpha_1 = -4$	$\alpha_2 = -3$	$\alpha_3 = -2$	$\alpha_4 = 2$	$\alpha_5 = 3$	$\alpha_6 = 4$
Perturbed Location:	$\alpha_1 = -4$	$\alpha_2 = -3$	$\alpha_3 = -2$	$\alpha_4 = 2$	$\alpha_5 = 3$	$\alpha_6 = 4.4$
Access-Case (i):	$w_1 = 0.166$	$w_2 = 0.166$	$w_3 = 0.166$	$w_4 = 0.166$	$w_5 = 0.166$	$w_6 = 0.166$
Access-Case (ii):	$w_1 = 0.188$	$w_2 = 0.188$	$w_3 = 0.188$	$w_4 = 0.144$	$w_5 = 0.144$	$w_6 = 0.144$
Risk aversion:	$\rho_1 = 0.2$	$\rho_2 = 0.2$	$\rho_3 = 0.2$	$\rho_4 = 0.2$	$\rho_5 = 0.2$	$\rho_6 = 0.2$
Constant:	$\gamma_1 = 100$	$\gamma_2 = 100$	$\gamma_3 = 100$	$\gamma_4 = 100$	$\gamma_5 = 100$	$\gamma_6 = 100$

Table 4B-(i) compares the last five rounds of negotiations for Case (i). In the last offer round ($T-1$), the shift to the right in player #6's ideal point reduces the other players' reservation utilities in round $T-2$. Because players are all equally risk averse, the effect of this shift is greater for players whose ideal points are further away from α_6 . Now consider the penultimate offer round (round $T-3$). Because admissible coalitions contain at least $n+1$ members, each left-winger (resp. right-winger) must induce one right-winger (resp. left-winger) to accept her proposal. It can be shown that in round $T-3$, the left-wingers all choose player #4 and the right-wingers all choose #3. But as we have seen, player #3's reservation utility in round $T-2$ is lowered more by the shift in α_6 than is player #4's. Therefore, while each left-wing proposal in round $T-3$ shifts to the left, the corresponding rightward shifts in the right-wing proposals are larger. It follows that in round $T-4$, each left-winger's reservation utility is reduced relative to its level in the original model, while each right-wing's reservation utility is either reduced by a lesser amount or, possibly, is increased. The effects of these changes are apparent in Table 4B-(i): compare players' reservation utilities in round $T-4$, and their offers in round $T-5$, of the original and perturbed models. The relative weakness of the left-wing in round $T-5$ is transmitted via backward induction to the first round of negotiations, resulting in a shift to the right in the solution policy.

In Case (ii), the access probabilities of the left-wingers are slightly greater than those of the right-wingers. Table 4B-(ii) illustrates the effects of the shift in α_6 in this case. In round $T-3$, the qualitative effects are the same as in Case (i): the left-wing proposals shift to the left, while the right-wing proposals shift to the right by a greater amount. However, when the asymmetry in access is sufficiently great, the smaller, but more heavily weighted leftward shift dominates the larger but less heavily weighted rightward shifts in the computation of players' expected utilities. Once again, the effects of these changes are apparent in Table 4B-(ii): compare players'

Table 4B-(i): Effect of Shifting Player #6's Ideal Point to the Right in Problem B.
Left and Right Wings Have Equal Access

Rnd	Prpr	x_1	$u_1(\cdot)$	$u_2(\cdot)$	$u_3(\cdot)$	$u_4(\cdot)$	$u_5(\cdot)$	$u_6(\cdot)$
T-1	#1	-4.000	6.310	5.800	5.278	3.031	2.408	1.741
	#2	-3.000	5.800	6.310	5.800	3.624	3.031	2.408
	#3	-2.000	5.278	5.800	6.310	4.193	3.624	3.031
	#4	2.000	3.031	3.624	4.193	6.310	5.800	5.278
	#5	3.000	2.408	3.031	3.624	5.800	6.310	5.800
	#6	4.000	1.741	2.408	3.031	5.278	5.800	6.310
T-2	Eu		4.095	4.495	4.706	4.706	4.495	4.095
T-3	#1	-1.069	4.781	5.314	5.835	4.706*	4.154	3.584
	#2	-1.069	4.781	5.314	5.835	4.706*	4.154	3.584
	#3	-1.069	4.781	5.314	5.835	4.706*	4.154	3.584
	#4	1.069	3.584	4.154	4.706*	5.835	5.314	4.781
	#5	1.069	3.584	4.154	4.706*	5.835	5.314	4.781
	#6	1.069	3.584	4.154	4.706*	5.835	5.314	4.781
T-4	Eu		4.182	4.734	5.270	5.270	4.734	4.182
T-5	#1	-0.019	4.204	4.754	5.288	5.268	4.733	4.182*
	#2	-0.019	4.204	4.754	5.288	5.268	4.733	4.182*
	#3	-0.019	4.204	4.754	5.288	5.268	4.733	4.182*
	#4	0.019	4.182*	4.733	5.268	5.288	4.754	4.204
	#5	0.019	4.182*	4.733	5.268	5.288	4.754	4.204
	#6	0.019	4.182*	4.733	5.268	5.288	4.754	4.204
Perturbed Location Configuration.								
T-1	#1	-4.000	6.310	5.800	5.278	3.031	2.408	1.456
	#2	-3.000	5.800	6.310	5.800	3.624	3.031	2.148
	#3	-2.000	5.278	5.800	6.310	4.193	3.624	2.786
	#4	2.000	3.031	3.624	4.193	6.310	5.800	5.066
	#5	3.000	2.408	3.031	3.624	5.800	6.310	5.592
	#6	4.400	1.456	2.148	2.786	5.066	5.592	6.310
T-2	Eu		4.047	4.452	4.665	4.671	4.461	3.893
T-3	#1	-1.134	4.816	5.349	5.868	4.671*	4.118	3.311
	#2	-1.134	4.816	5.349	5.868	4.671*	4.118	3.311
	#3	-1.134	4.816	5.349	5.868	4.671*	4.118	3.311
	#4	1.144	3.540	4.112	4.665*	5.874	5.354	4.604
	#5	1.144	3.540	4.112	4.665*	5.874	5.354	4.604
	#6	1.144	3.540	4.112	4.665*	5.874	5.354	4.604
T-4	Eu		4.178	4.730	5.267	5.272	4.736	3.957
T-5	#1	-0.018	4.203	4.753	5.288	5.268	4.733	3.957*
	#2	-0.018	4.203	4.753	5.288	5.268	4.733	3.957*
	#3	-0.018	4.203	4.753	5.288	5.268	4.733	3.957*
	#4	0.027	4.178*	4.729	5.264	5.292	4.758	3.983
	#5	0.027	4.178*	4.729	5.264	5.292	4.758	3.983
	#6	0.027	4.178*	4.729	5.264	5.292	4.758	3.983

Table 4B-(ii): Effect of Shifting Player #6's Ideal Point to the Right in Problem B.
Left Wing Has Greater Access

Rnd	Ppr	x_1	$u_1(\cdot)$	$u_2(\cdot)$	$u_3(\cdot)$	$u_4(\cdot)$	$u_5(\cdot)$	$u_6(\cdot)$
T-1	#1	-4.000	6.310	5.800	5.278	3.031	2.408	1.741
	#2	-3.000	5.800	6.310	5.800	3.624	3.031	2.408
	#3	-2.000	5.278	5.800	6.310	4.193	3.624	3.031
	#4	2.000	3.031	3.624	4.193	6.310	5.800	5.278
	#5	3.000	2.408	3.031	3.624	5.800	6.310	5.800
	#6	4.000	1.741	2.408	3.031	5.278	5.800	6.310
T-2	Eu		4.317	4.688	4.848	4.564	4.303	3.873
T-3	#1	-1.330	4.921	5.451	5.969	4.564*	4.008	3.432
	#2	-1.330	4.921	5.451	5.969	4.564*	4.008	3.432
	#3	-1.330	4.921	5.451	5.969	4.564*	4.008	3.432
	#4	0.806	3.736	4.301	4.848*	5.699	5.175	4.638
	#5	0.806	3.736	4.301	4.848*	5.699	5.175	4.638
	#6	0.806	3.736	4.301	4.848*	5.699	5.175	4.638
T-4	Eu		4.406	4.951	5.482	5.058	4.515	3.956
T-5	#1	-0.421	4.427	4.970	5.499	5.055	4.514	3.956*
	#2	-0.421	4.427	4.970	5.499	5.055	4.514	3.956*
	#3	-0.421	4.427	4.970	5.499	5.055	4.514	3.956*
	#4	-0.383	4.406*	4.950	5.479	5.075	4.534	3.977
	#5	-0.383	4.406*	4.950	5.479	5.075	4.534	3.977
	#6	-0.383	4.406*	4.950	5.479	5.075	4.534	3.977
Perturbed Location Configuration.								
T-1	#1	-4.000	6.310	5.800	5.278	3.031	2.408	1.456
	#2	-3.000	5.800	6.310	5.800	3.624	3.031	2.148
	#3	-2.000	5.278	5.800	6.310	4.193	3.624	2.786
	#4	2.000	3.031	3.624	4.193	6.310	5.800	5.066
	#5	3.000	2.408	3.031	3.624	5.800	6.310	5.592
	#6	4.400	1.456	2.148	2.786	5.066	5.592	6.310
T-2	Eu		4.275	4.650	4.813	4.533	4.273	3.663
T-3	#1	-1.386	4.951	5.481	5.998	4.533*	3.976	3.161
	#2	-1.386	4.951	5.481	5.998	4.533*	3.976	3.161
	#3	-1.386	4.951	5.481	5.998	4.533*	3.976	3.161
	#4	0.872	3.698	4.264	4.813*	5.733	5.210	4.455
	#5	0.872	3.698	4.264	4.813*	5.733	5.210	4.455
	#6	0.872	3.698	4.264	4.813*	5.733	5.210	4.455
T-4	Eu		4.406	4.952	5.482	5.055	4.513	3.723
T-5	#1	-0.428	4.431	4.974	5.503	5.051	4.510	3.723*
	#2	-0.428	4.431	4.974	5.503	5.051	4.510	3.723*
	#3	-0.428	4.431	4.974	5.503	5.051	4.510	3.723*
	#4	-0.384	4.406*	4.950	5.480	5.074	4.534	3.749
	#5	-0.384	4.406*	4.950	5.480	5.074	4.534	3.749
	#6	-0.384	4.406*	4.950	5.480	5.074	4.534	3.749

reservation utilities in round $T-4$, and their offers in round $T-5$, of the original and perturbed models. In this case, the right-wingers' reservation utilities fall, while the left-wingers' stay virtually the same. This time, the relative weakness of the right-wing is transmitted to the first round, and the increase in right-wing extremism results in a shift to the left of the solution policy.

Problem C: $X \subset \mathbb{R}^2$; $2n$ players; strict majority rule.

A stylized fact about bipolar negotiations between two factions is that either one of the factions will be more effective in its pursuit of the common objectives of its members, the greater the degree of cohesiveness among its membership. As this example demonstrates, the predictions of our model are consistent with this observation. In spatial problems, a natural measure of the cohesiveness of a faction is the proximity of its members' ideal points to each other. We will show that as the distance between the right-wingers' ideal points is increased, the solution vector shifts to the left.

When the space of policies is two-dimensional, deterministic solutions do not exist in general. However, Theorem IV can be extended to guarantee existence provided that agents' ideal points are confined to an "almost" one-dimensional set. Once again, we illustrate the discussion by a pair of six-player examples, whose parameters are specified below.

Problem C: Reducing the Cohesiveness of the Right-Wing Faction (See Table 4C).						
	Player #1	Player #2	Player #3	Player #4	Player #5	Player #6
Initial Location:	$\alpha_1 = (-9, -1)$	$\alpha_2 = (-9, 0)$	$\alpha_3 = (-9, +1)$	$\alpha_4 = (9, -1)$	$\alpha_5 = (9, 0)$	$\alpha_6 = (9, +1)$
Perturbed Location:	$\alpha_1 = (-9, -1)$	$\alpha_2 = (-9, 0)$	$\alpha_3 = (-9, +1)$	$\alpha_4 = (9, -1)$	$\alpha_5 = (9, 0)$	$\alpha_6 = (9, +1)$
Access:	$w_1 = 0.166$	$w_2 = 0.166$	$w_3 = 0.166$	$w_4 = 0.166$	$w_5 = 0.166$	$w_6 = 0.166$
Risk aversion:	$\rho_1 = 0.5$	$\rho_2 = 0.5$	$\rho_3 = 0.5$	$\rho_4 = 0.5$	$\rho_5 = 0.5$	$\rho_6 = 0.2$
Constant:	$\gamma_1 = 100$	$\gamma_2 = 100$	$\gamma_3 = 100$	$\gamma_4 = 100$	$\gamma_5 = 100$	$\gamma_6 = 100$

Table 4C compares the last four rounds of negotiations for the initial and perturbed locations. In this case, the argument is quite straightforward. In round $T-1$, each player proposes her ideal point. When the ideal points of the right-wingers are dispersed, there is a significant loss in utility for each of them. On the other hand, the vertical shifts in the right-wing proposals are so small relative to the gap between the left- and right-wing locations that the dispersion barely affects the left-wingers at all. (Intuitively, imagine heated disputes between conservatives about the fine details of their ideology which radicals perceive as no more than arcane hair-splitting.) As a result, right-wingers' reservation utilities in round $T-2$ fall, while left-winger's remain almost the same. The effect of this difference, once again, is to shift the solution policy to the left.

Problem D: Three players, $X \subset \mathbb{R}^2$, variable coalition configurations.

In the three preceding problems we assumed that decisions were made by majority rule. In this problem we consider alternative coalition structures. In particular, we consider the effects of declaring one or more players to be essential. The ideal points of the three players are, respectively, $\alpha_1 = (-1, 0)$, $\alpha_2 = (+1, 0)$ and $\alpha_3 = (0, 1)$.

First assume that any coalition of two players is admissible, so that no player is essential. In this case, the core of the underlying bargaining problem is clearly empty, so that the model cannot have a deterministic solution. Not surprisingly, the sequence of equilibrium outcomes settles into a cyclic pattern, for reasons very similar to those discussed on p. above. Now assume that player #1 is essential. From Theorem V, this model has a deterministic solution. Since the core of the underlying problem contains exactly one point--player #1's ideal point--it follows from Theorem II that this point is the unique solution policy. Next, consider the unanimity version of this problem,

Table 4C: Reducing the Cohesiveness of the Right Wing in Problem B.
Initial Location Configuration

Rnd	Prpr	x_1	x_2	$u_1(\cdot)$	$u_2(\cdot)$	$u_3(\cdot)$	$u_4(\cdot)$	$u_5(\cdot)$	$u_6(\cdot)$
T-1	#1	-9.000	-1.000	10.000	9.950	9.899	9.055	9.054	9.049
	#2	-9.000	0.000	9.950	10.000	9.950	9.054	9.055	9.054
	#3	-9.000	1.000	9.899	9.950	10.000	9.049	9.054	9.055
	#4	9.000	-1.000	9.055	9.054	9.049	10.000	9.950	9.899
	#5	9.000	0.000	9.054	9.055	9.054	9.950	10.000	9.950
	#6	9.000	1.000	9.049	9.054	9.055	9.899	9.950	10.000
T-2	Eu			9.501	9.510	9.501	9.501	9.510	9.501
T-3	#1	-0.725	-1.000	9.577	9.574	9.565	9.501*	9.499	9.491
	#2	-0.710	0.461	9.570	9.576	9.576	9.496	9.502	9.501*
	#3	-0.725	1.000	9.565	9.574	9.577	9.491	9.499	9.501*
	#4	0.725	-1.000	9.501*	9.499	9.491	9.577	9.574	9.565
	#5	0.710	0.461	9.496	9.502	9.501*	9.570	9.576	9.576
	#6	0.725	1.000	9.491	9.499	9.501*	9.565	9.574	9.577
T-4	Eu			9.533	9.537	9.535	9.533	9.537	9.535
Perturbed Location Configuration.									
T-1	#1	-9.000	-1.000	10.000	9.950	9.899	9.055	9.054	9.046
	#2	-9.000	0.000	9.950	10.000	9.950	9.052	9.055	9.052
	#3	-9.000	1.000	9.899	9.950	10.000	9.046	9.054	9.055
	#4	9.000	-1.500	9.055	9.052	9.046	10.000	9.925	9.849
	#5	9.000	0.000	9.054	9.055	9.054	9.925	10.000	9.925
	#6	9.000	1.500	9.046	9.052	9.055	9.849	9.925	10.000
T-2	Eu			9.501	9.510	9.501	9.488	9.502	9.488
T-3	#1	-0.979	-1.223	9.590	9.586	9.575	9.488*	9.484	9.469
	#2	-0.949	0.671	9.580	9.588	9.589	9.477	9.488	9.488*
	#3	-0.979	1.223	9.575	9.586	9.590	9.469	9.484	9.488*
	#4	0.733	-1.270	9.501*	9.497	9.487	9.578	9.573	9.554
	#5	0.722	0.460	9.496	9.501	9.501*	9.565	9.577	9.574
	#6	0.733	1.270	9.487	9.497	9.501*	9.554	9.573	9.578
T-4	Eu			9.538	9.542	9.540	9.522	9.530	9.525

in which all three players are essential. In this case, the core consists of the convex hull of the three players' ideal points. The solution to the model derived from this problem depends on the entire distribution of bargaining attributes among the three players. The comparative statics properties are predictable. If one player's access probability increased or her risk aversion coefficient decreases, the solution shifts in the direction of that player's ideal point. Finally, suppose that players #1 and #2 are both essential, so that the admissible coalitions are (1, 2) and (1, 2, 3). In this case, the core of the underlying game is the "contract curve" joining the essential players' ideal points, i.e., the line segment $\{(\beta, 0): \beta \in [-1, 1]\}$. If the two essential players have equal access, the solution outcome will be the midpoint of this line, i.e., the origin.

It is instructive to investigate the role that player #3 plays in this configuration. Though players #1 and #2 never invite her to join a coalition, player #3's presence affects the outcome of negotiations, provided her access probability is positive. (Think of #3 as representing a group that is peripheral to the decision-making process, but has the capacity to capture the attention of the general public, and thereby influence the nature of the debate between the major players (cf. our discussion of essentiality and access on pp. 3-3)). To illustrate this, we simulate the effect of a leftward shift in #3's ideal point. The parameters for the illustration are displayed below.

Problem D: Shifting Player #3's Ideal Point to the Left (See Table 4D).			
	Player #1	Player #2	Player #3
Initial Location:	$\alpha_1 = (-1, 0)$	$\alpha_2 = (+1, 0)$	$\alpha_3 = (0, 1)$
Perturbed Location:	$\alpha_1 = (-1, 0)$	$\alpha_2 = (+1, 0)$	$\alpha_3 = (0, -0.05)$
Access:	$w_1 = 0.333$	$w_2 = 0.333$	$w_3 = 0.333$
Risk aversion:	$\rho_1 = 0.2$	$\rho_2 = 0.2$	$\rho_3 = 0.2$
Constant:	$\gamma_1 = 100$	$\gamma_2 = 100$	$\gamma_3 = 100$

Table 4D compares the last six rounds of negotiations for the initial and perturbed locations. In the final offer round ($T-1$), the shift in player #3's proposal benefits player #1 at the expense of #2. In the preceding response round ($T-2$), therefore, #1's reservation utility is higher than initially, while #2's is lower. In the penultimate offer round ($T-3$), there are three changes. Player #1's proposal is closer to α_1 , because #2's reservation utility is lower. Player #2's proposal is further from α_2 , because #1's reservation utility is higher. Finally, player #3's proposal is closer to α_1 and further from α_2 , both because her own ideal point is now closer to α_1 and because of the shifts in the other two players' reservation utilities. All three of these changes benefit #1 at the expense of #2, so that in round $T-4$, #1's reservation utility is higher than initially, while #2's is lower. The effects of these changes are transmitted via backward induction to the first round of negotiations, resulting in a shift to the left in the solution policy.

Problem E: A two-good pure exchange economy with four players.

In this final problem we extend our framework to model negotiations between agents in a pure exchange economy. While the problem is very simple, it extends the preceding analysis in three respects. First, it demonstrates that our framework can be applied a wider class of problems than the ones we have considered thus far. Second, a deterministic solution is obtained even though there is no essential player and the policy space is high dimensional. Third, it extends one of the basic assumptions of the paper, by allowing players' policy choices to depend on the coalitions they select.

There are two commodities and four players. Any subset of these players forms an admissible coalition. Each player has equal access. Players #1 and #2 are each endowed with two units of the first commodity while players #3 and #4 are each endowed with two units of the second. A policy is an allocation $x = (x_{i1}, x_{i2})_{i=1}^4$ satisfying, for $k = 1, 2$, $x_{ik} \geq 0$ and $\sum_{i=1}^4 x_{ik} = 4$. If a player proposes the coalition C , she can propose any allocation in which the players excluded from C are all assigned their initial endowments. Player i 's utility is the Cobb-

Table 4D: Effect of Shifting Player #3's Ideal Point to the Left in Problem D.
Initial Location Configuration

Rnd	Ptpr	x_1	x_2	$u_1(\cdot)$	$u_2(\cdot)$	$u_3(\cdot)$
T-1	#1	-1.000	0.000	66.289	66.056	66.124
	#2	1.000	0.000	66.056	66.289	66.124
	#3	0.000	1.000	66.124	66.124	66.289
T-2	Eu			66.156	66.156	66.179
T-3	#1	-0.138	0.000	66.189	66.156*	66.171
	#2	0.138	0.000	66.156*	66.189	66.171
	#3	0.000	0.543	66.156*	66.156*	66.236
T-4	Eu			66.167	66.167	66.193
T-5	#1	-0.046	0.000	66.178	66.167*	66.172
	#2	0.046	0.000	66.167*	66.178	66.172
	#3	0.000	0.307	66.167*	66.167*	66.208
T-6	Eu			66.171	66.171	66.184
Perturbed Location Configuration						
T-1	#1	-1.000	0.000	66.289	66.056	66.128
	#2	1.000	0.000	66.056	66.289	66.120
	#3	-0.050	1.000	66.128	66.120	66.289
T-2	Eu			66.158	66.155	66.179
T-3	#1	-0.150	0.000	66.190	66.155*	66.172
	#2	0.126	0.000	66.158*	66.187	66.171
	#3	-0.013	0.544	66.158*	66.155*	66.236
T-4	Eu			66.169	66.166	66.193
T-5	#1	-0.057	0.000	66.179	66.166*	66.173
	#2	0.034	0.000	66.169*	66.177	66.172
	#3	-0.012	0.307	66.169*	66.166*	66.208
T-6	Eu			66.172	66.169	66.184

Douglas function $u_i(x) = (x_{i1}x_{i2})^p$, with $p < 1/2$.

The model derived from this problem has a unique deterministic solution. Not surprisingly, since the model is completely symmetric, the solution policy is the symmetric allocation in which each player receives one unit of each commodity. The proof is extremely simple. To reduce notation we set $p=0.25$. It can be established that in each response round, each player has the same reservation utility. For each even integer t , let α_t denote this common reservation utility.

In round $T-1$, each player proposes the grand coalition and selects the allocation in which she receives the the aggregate endowment vector $(4, 4)$. Thus, $\alpha_{T-2} = 1/4(16)^{0.25} = 1/2$. Now fix an odd integer t , and assume that players' reservation utilities in round $t+1$ are all equal to $\alpha_{t+1} < 1$. We will show that in round t , each player selects the grand coalition and, modulo relabeling, the same allocation. Moreover, we will show that the common reservation utility in round $t-1$ is $\alpha_{t-1} \in (1/4(1+3\alpha_{t+1}), 1)$. Consider the options facing player i in round t . Her opportunities in any two-player coalition are clearly dominated by her opportunities in the grand coalition. Moreover, it is straightforward to verify that if she selects any three-player coalition, the best trade she can achieve is $(2 - \sqrt{2}\alpha_{t+1}^2)$ units of the scarce commodity and $2(2 - \sqrt{2}\alpha_{t+1}^2)$ of the other.³¹ If she selects the grand coalition, the best trade she can achieve is $(4 - 3\alpha_{t+1}^2)$ of each commodity. Since $\alpha_{t+1} < 1$, the latter trade yields a higher utility than the former.³² We have thus verified that in round t , each player selects the grand coalition and, modulo relabeling, the same allocation. When player $j \neq i$ makes a proposal, player i receives the utility level α_{t+1} ; when i herself proposes, her utility is $(4 - 3\alpha_{t+1}^2)^{0.25} > 1$. It follows that $\alpha_{t-1} > 1/4(1+3\alpha_{t+1})$. Finally, players' reservation utilities obviously cannot all exceed unity. We have established, therefore, that $\alpha_{t-1} \in (1/4(1+3\alpha_{t+1}), 1)$. This completes our verification of the inductive hypothesis. It follows that for every positive ϵ , if T is sufficiently large then players' common reservation utility in the second round of the T -round game must exceed $(1 - \epsilon)$. But in this case, the policies proposed by each player in the first round must be arbitrarily close to the allocation in which each player receives one unit of each commodity.

³¹ Without loss of generality, consider player #1's opportunities if she selects the coalition $C = \{1, 2, 3\}$. The aggregate endowment is $(4, 2)$. A necessary condition for an optimal allocation is that for $i \in C$, $x_{i1} = 2x_{i2}$. For $i = 2, 3$, $(2x_{i2}^2)^{0.25} = \alpha_{t+1}$, so that $x_{i2} = \sqrt{2}\alpha_{t+1}^2$. Player #1 takes what remains of the second commodity, i.e. $2 - 2\sqrt{2}\alpha_{t+1}^2$, and twice as much of the first commodity.

³² Since $2(2 - \sqrt{2}\alpha_{t+1}^2)^2 < (1.5(2 - \sqrt{2}\alpha_{t+1}^2))^2$, it is sufficient to check that $(3 - \sqrt{4.5}\alpha_{t+1}^2) < (4 - 3\alpha_{t+1}^2)$, i.e., that $(3 - \sqrt{4.5})\alpha_{t+1}^2 < 1$. This inequality clearly holds whenever $\alpha < 1$.

APPENDIX A: TWO-DIMENSIONAL SPATIAL PROBLEMS.

There are, of course, bargaining problems for which neither Theorem IV nor Theorem V applies, either because the policy space is multidimensional or because there is no essential player. This is true of a class of problems that has played an extremely important role in political science theory.³³ These are spatial problems in which the policy space is two-dimensional. The informal discussion below summarizes what can be inferred about this class of problems by applying numerical simulation techniques.

First, for every problem that has a nonempty core, we have been able to compute a deterministic solution for the model derived from that problem. For example, the core is nonempty for every four-person two-dimensional problem with strict majority rule and we have computed solutions to hundreds of corresponding models. Second, the closer a problem is to one with a nonempty core, the more likely it is that the model derived from it will have an exact solution.³⁴ Moreover, if a solution is not exact, it is more likely to be almost exact. Finally, the outcomes implemented by these exact or approximate solutions are more likely to be close to the core of the neighboring problem.

To demonstrate the relationship between the structural characteristics of bargaining problems and the frequency of different solution types for the corresponding models, we report on three Monte Carlo experiments, referred to as experiments A, B and C. In each experiment we sample one hundred bargaining problems. The sample spaces are three increasingly general, parameterized families of five-person spatial problems, ranging from a family in which the core is always "almost nonempty" to one in which only minimal restrictions are imposed. In all three experiments, access probabilities are sampled from the four-dimensional unit simplex, and players' risk aversion coefficients lie on the unit interval. The sample spaces for players' ideal points are displayed in the table below, with $\alpha_i = (\alpha_{i1}, \alpha_{i2})$ denoting player i 's ideal point. In experiments A and B, only α_3 is selected randomly while in experiment C, all five α_i 's are randomly chosen.

Sample Spaces for Players' Ideal Points in Experiments A-C.					
	Player #1	Player #2	Player #3	Player #4	Player #5
A:	$\alpha_1 = (-1, 0)$	$\alpha_2 = (+1, 0)$	$(-.05, -.05) \leq \alpha_3 \leq (.05, .05)$	$\alpha_4 = (0, -1)$	$\alpha_5 = (0, +1)$
B:	$\alpha_1 = (-1, 0)$	$\alpha_2 = (+1, 0)$	$(-0.5, -0.5) \leq \alpha_3 \leq (0.5, 0.5)$	$\alpha_4 = (0, -1)$	$\alpha_5 = (0, +1)$
C:	$(-1, -1) \leq \alpha_1 \leq (1, 1)$	$(-1, -1) \leq \alpha_2 \leq (1, 1)$	$(-1, -1) \leq \alpha_3 \leq (1, 1)$	$(-1, -1) \leq \alpha_4 \leq (1, 1)$	$(-1, -1) \leq \alpha_5 \leq (1, 1)$

It is well known that in experiments A and B, the core is nonempty if and only if player #3's ideal point is located at the origin.³⁵


The results of the three experiments are summarized in the three histograms presented in Table 5.1. In each case, the leftmost column reports the frequency of exact stochastic equilibria. The other columns indicate the frequency of approximate equilibria with different degrees of inexactness: specifically, the height of the bar labelled "from a to b " represents the frequency with which we computed an ϵ -equilibrium with $\epsilon \in (a, b)$. The experimental results are consistent with the qualitative remarks offered above. In particular, in experiment A the likelihood of an exact equilibrium is very high, while virtually all of the approximate equilibria are almost exact. As the class of problems is expanded, the likelihood of an exact equilibrium declines, and the likelihood of a quite inexact solution increases. Of course, because of the methodology used here, these statistics are necessarily subject to certain caveats. In particular, while we have observed stable cycles over thousands of rounds, and inferred from these cycles the existence of approximate solutions, the existence of infinite stable cycles cannot, obviously, be *guaranteed* by numerical methods. Equally obviously, numerical methods cannot *guarantee* the existence of exact

³³ The seminal papers in this literature are Davis and Hinich [1966] and Plott [1967]. For surveys of the literature, see Enelow and Hinich [1984] and Ordeshook [1986].


³⁴ Because our sample spaces are all finite-dimensional, the notion of "close" here is the standard one. Also, the sample spaces in the following discussion are all endowed with Lebesgue measure, and terms such as "more likely" have precise meanings in terms of this measure.

³⁵ See section 4.7 of Ordeshook [1986], and Fiorina-Plott [1978].


Table 5.1: Histograms for Experiments A-C



Experiment A.



Experiment B.



Experiment C.

stochastic solutions.³⁶

Our third observation is that in virtually all of the bargaining problems for which an exact solution was not obtained, the sequence of outcomes settled into a cyclic pattern. It is instructive to investigate the nature of these cycles. Essentially, they arise because as one player alternates between different coalitions, selecting those members whose participation can be obtained most "cheaply," other players' participation "prices" change in response, generating a stable oscillatory pattern of optimal coalition choices. To illustrate this cyclic phenomenon, consider a problem drawn at random in experiment B with the following parameters:

Parameters for the Simulation Displayed in Table 5.2.					
	Player #1	Player #2	Player #3	Player #4	Player #5
Location:	$\alpha_1 = (-1, 0)$	$\alpha_2 = (+1, 0)$	$\alpha_3 = (-0.193, 0.202)$	$\alpha_4 = (0, -1)$	$\alpha_5 = (0, +1)$
Access:	$w_1 = 0.152$	$w_2 = 0.222$	$w_3 = 0.216$	$w_4 = 0.183$	$w_5 = 0.227$
Risk aversion:	$\rho_1 = 0.860$	$\rho_2 = 0.706$	$\rho_3 = 0.200$	$\rho_4 = 0.892$	$\rho_5 = 0.746$
Constant:	$\gamma_1 = 60$	$\gamma_2 = 60$	$\gamma_3 = 60$	$\gamma_4 = 60$	$\gamma_5 = 60$

In Table 5.2, the relevant computations for selected rounds of the 2000 round game are displayed in reverse order.³⁷ Clearly, the sequence of offer vectors settles into a two-period limit cycle.³⁸ For the odd-numbered (offer) rounds, column 2 lists the proposer and column 3 lists the members (in addition to herself) that the proposer invites to form a coalition. Columns 4 and 5 list the policies proposed by each player and 6 through 10 display the utilities that each player derives from each of the proposed policies. An asterisk in a column indicates that the player's participation constraint is binding for the proposer. For the even-numbered (response) rounds, columns 6 through 10 display players' reservation utilities, which are, identically, their expected utilities conditional on reaching the following offer round. Entries that are central to the following discussion are emboldened.

Observe that except for the first component of player #5's proposal, the offers remain relatively similar from round to round. In the fifth round, player #5 proposes the policy $(-0.20, 0.52)$ and the coalition $\{5, 1, 3\}$. Player #1's participation constraint is binding, while player #2's utility from this proposal falls short of her reservation utility in the sixth round. Consequently, #2's reservation utility is *lower* in the fourth round than in the sixth, while for #1 the ordering is reversed. Thus in the third and fifth rounds, the configuration of "participation prices" confronting player #5 is slightly different: in the third, the price of securing #2's agreement is slightly lower relative to the price of securing #1's agreement. The difference is enough to tilt the balance in favor of player #2, and so in the third round, player #5 proposes the policy $(-0.05, 0.48)$ and the coalition $\{5, 2, 3\}$. Player #2's participation constraint is now binding, while player #1's utility from #5's proposal falls short of her reservation utility in the fourth round. Consequently #2's reservation utility is *higher* in the second round than in the fourth, while for #1 the ordering is again reversed. In the first round, the relative prices facing #5 are virtually the same as in the fifth round,³⁹ and she chooses player #1 in preference to #2.

In parameterized families of problems with empty cores, we should not expect to identify conditions that can distinguish models with exact stochastic solutions from those with only approximate solutions. Cyclic patterns arise because the negotiating problem facing players is inherently discontinuous: each player must choose from a finite set of coalitions and as one player switches coalitions, other players' payoffs change discontinuously. Our simulations indicate that exact solutions result whenever players' optimal coalition choices are unchanged from round to round. Conversely, a cyclic pattern emerges whenever at least one player's optimal coalition choice regularly changes. Clearly, it will be extremely difficult to ensure that players' optimal coalition choices remain

³⁶ The exact stochastic solutions that we report are indeed exact to the limits of machine precision, but the tolerances of our computational algorithm are relatively coarse (approximately 10^{-6}).

³⁷ Because the sequence of offers is uniquely determined, the offers made in round #3 of the T -round game are, identically, the initial offers in the $T-2$ -round game, etc.

³⁸ In fact, the cyclic pattern is not quite exact. There are slight differences between the offers in rounds t and $t+4$ that are obscured by rounding.

Table 5.2: A Two-period Limit Cycle.

Rnd	Prpr	Coal	x_1	x_2	$u_1(\cdot)$	$u_2(\cdot)$	$u_3(\cdot)$	$u_4(\cdot)$	$u_5(\cdot)$
1000	Eu				1.77076	3.31871	26.29755	1.55071	2.82409
999	#1	{3,4}	-0.479479	0.078169	1.77245	3.31335	26.29755*	1.55071*	2.82227
	#2	{4,5}	0.263466	0.150188	1.76932	3.32546	26.24555	1.55071*	2.82409*
	#3	{1,5}	-0.193165	0.201892	1.77117	3.31787	26.40391	1.55060	2.82492
	#4	{1,2}	-0.120128	-0.300296	1.77076*	3.31871*	26.22866	1.55204	2.81902
	#5	{2,3}	-0.054990	0.481519	1.77021	3.31871*	26.29755*	1.54985	2.82856
998	Eu				1.77066	3.31921	26.29644	1.55073	2.82408
997	#1	{3,4}	-0.478789	0.068909	1.77245	3.31337	26.29644*	1.55073*	2.82218
	#2	{4,5}	0.248006	0.144683	1.76939	3.32523	26.25072	1.55073*	2.82408*
	#3	{1,5}	-0.193165	0.201892	1.77117	3.31787	26.40391	1.55060	2.82492
	#4	{1,2}	-0.091547	-0.289190	1.77066*	3.31921*	26.23077	1.55202	2.81918
	#5	{1,3}	-0.198921	0.516905	1.77066*	3.31628	26.29644*	1.54971	2.82855
6	Eu				1.77066	3.31921	26.29644	1.55073	2.82408
5	#1	{3,4}	-0.478789	0.068909	1.77245	3.31337	26.29644*	1.55073*	2.82218
	#2	{4,5}	0.248006	0.144683	1.76939	3.32523	26.25072	1.55073*	2.82408*
	#3	{1,5}	-0.193165	0.201892	1.77117	3.31787	26.40391	1.55060	2.82492
	#4	{1,2}	-0.091547	-0.289190	1.77066*	3.31921*	26.23077	1.55202	2.81918
	#5	{1,3}	-0.198921	0.516905	1.77066*	3.31628	26.29644*	1.54971	2.82855
4	Eu				1.77076	3.31871	26.29755	1.55071	2.82409
3	#1	{3,4}	-0.479479	0.078169	1.77245	3.31335	26.29755*	1.55071*	2.82227
	#2	{4,5}	0.263466	0.150188	1.76932	3.32546	26.24555	1.55071*	2.82409*
	#3	{1,5}	-0.193165	0.201892	1.77117	3.31787	26.40391	1.55060	2.82492
	#4	{1,2}	-0.120128	-0.300296	1.77076*	3.31871*	26.22866	1.55204	2.81902
	#5	{2,3}	-0.054990	0.481519	1.77021	3.31871*	26.29755*	1.54985	2.82856
2	Eu				1.77066	3.31921	26.29644	1.55073	2.82408
1	#1	{3,4}	-0.478789	0.068909	1.77245	3.31337	26.29644*	1.55073*	2.82218
	#2	{4,5}	0.248006	0.144683	1.76939	3.32523	26.25072	1.55073*	2.82408*
	#3	{1,5}	-0.193165	0.201892	1.77117	3.31787	26.40391	1.55060	2.82492
	#4	{1,2}	-0.091547	-0.289190	1.77066*	3.31921*	26.23077	1.55202	2.81918
	#5	{1,3}	-0.198921	0.516905	1.77066*	3.31628	26.29644*	1.54971	2.82855

constant, and hence to guarantee existence of an exact stochastic solution. On the other hand, our Monte Carlo experiments do suggest that there is a predictable relationship between the structure of a given family of problems and the relative likelihood of exact, almost exact and inexact solutions for the corresponding models.

The preceding discussion raises a wide variety of questions such as: What is the precise relationship between the dynamical system we have been investigating and static concepts such as the core? Is there an necessary relationship between the likelihood of a nonempty approximate core and the frequency of an exact or almost exact solution; conversely, are there families of problems for which even approximate cores are usually empty, yet exact or approximate solutions to the corresponding models arise as frequently, say, as in experiment A? Why are the cyclic patterns we have observed so prevalent, rather than, say, some kind of chaotic behavior? To what extent are the various observations reported above robust with respect to alternative functional forms?

REFERENCES.

- Baron, D. and J. Ferejohn (1989), "Bargaining in Legislatures," *American Political Science Review*, **84**, pp. 1181-1206.
- Baron, D. and E. Kalai (1991), "Dividing a Cake by Majority: The Simplest Equilibria," Unpublished manuscript.
- Bennett E. (1991a), "Three Approaches to Bargaining in NTU Games," in *Strategic Bargaining*, R. Selten (ed.), Berlin: Springer Verlag.
- Bennett E. (1991b), "Multilateral Bargaining Problems," Unpublished manuscript.
- Bennett E. and Houba, H. (1991), "Bargaining Among Three Players," Unpublished manuscript.
- Bennett E. and Van Damme, E. (1991), "Demand, Commitment, Bargaining: the Case of Apex Games," in *Strategic Bargaining*, R. Selten (ed.), Berlin: Springer Verlag.
- Binmore, K (1985), "Bargaining and Coalitions," in *Game Theoretic Models of Bargaining*, A. E. Roth (ed.), Cambridge: Cambridge University Press.
- Binmore, K. (1987), "Nash Bargaining Theory II," in *The Economics of Bargaining*, K. Binmore, and P. Dasgupta (ed.), Oxford: Blackwell.
- Binmore, K., A. Rubinstein and A. Wolinsky (1986), "The Nash Bargaining Solution in Economic Modelling," *Rand Journal of Economics*, **17**, 176-188.
- Binmore, K. and M. Osborne (1990), "Sequential Elimination of Dominated Strategies in the Bargaining Game of Alternating Offers," Unpublished manuscript.
- Chatterjee, K., B. Dutta, D. Ray and D. Sengupta (1987), "A Noncooperative Theory of Coalitional Bargaining," Unpublished manuscript.
- Davis, O.A. and M.J. Hinich (1966), "A Mathematical Model of Policy Formation in a Democratic Society," in *Mathematical Applications in Political Science, II*, J.L. Bernd (ed.), Dallas: Southern Methodist University Press.
- Enelow, J. M. and M. J. Hinich (1984), *A Spatial Theory of Elections* New York: Cambridge University Press.
- Fiorina, M. and C. Plott (1978), "Committee Decisions under Majority Rule: An Experimental Study," *American Political Science Review*, **72**, pp. 575-598.
- Harrison, G. (1991), "Multilateral Bargaining in Economics Experiments: A Survey," Unpublished manuscript, University of South Carolina, Department of Economics.
- Harrison, G. and K. McCabe (1992), "Testing Bargaining Theory in Experiments," in *Research in Experimental Economics (volume 5)*, R. M. Isaac (ed.) Greenwich: JAI Press, forthcoming.
- Krishna, V. and R. Serrano (1991), "Multilateral Bargaining," Harvard Business School Working Paper #91-026.
- Lensberg, T. (1988), "Stability and the Nash Solution," *Journal of Economic Theory*, **45**, pp. 330-341.
- Moulin, H. (1979), "Dominance Solvable Voting Schemes," *Econometrica*, **47**, pp. 1337-1351.

- Myerson, R. (1987), "Refinements of the Nash Equilibrium Concept," *International Journal of Game Theory*, 7, pp. 73-80.
- Nash, J.F. (1953), "Two-Person Cooperative Games," *Econometrica*, 21, 128-140.
- Neelin, J., H. Sonnenschein and M. Spiegel (1988), "A Further Test of Noncooperative Bargaining Theory," *American Economic Review*, 78, pp. 824-836.
- Ordeshook, P.C. (1986), *Game Theory and Political Theory: an Introduction*, New York: Cambridge University Press.
- Osborne, M.J. and A. Rubinstein (1990), *Bargaining and Markets*, San Diego: Academic Press.
- Plott, C. (1967), "A Notion of Equilibrium and its Possibility under Majority Rule," *American Economic Review*, 57, pp. 787-806.
- Rausser, G. and L. Simon (1991), "Burden Sharing and Public Good Investments in Policy Reform," Unpublished manuscript, University of California at Berkeley.
- Rausser, G. and L. Simon (1992), "The Political Economy of Transition in Eastern Europe: Packaging Enterprises for Privatization," in *The Emergence of Market Economies in Eastern Europe*, C. Clague and G. Rausser (eds.), forthcoming.
- Rubinstein, A. (1982), "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50, 97-109.
- Salant, S. and E. Goldstein (1990), "Predicting Committee Behavior in Majority Rule Voting Experiments," *Rand Journal of Economics*, 20, 293-313.
- Selten, R. (1975), "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Form Games," *International Journal of Game Theory*, 4, pp. 25-55.
- Selten, R. (1981), "A Noncooperative Model of Characteristic Function Bargaining," in *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*, V. Boehm and H. Nachtamp (eds.), Wissenschaftsverlag Bibliographisches Institut Mannheim, Wein - Zurich, pp. 131-151.
- Simon, L. and M. Stinchcombe (1989), "Extensive Form Games in Continuous Time: Pure Strategies," *Econometrica*, 57,
- Simon, L. and M. Stinchcombe (1991), "Equilibrium Refinements in Infinite Games: The Compact and Continuous Case," University of California San Diego Working Paper #91-22.
- Spiegel, M., J. Neelin, H. Sonnenschein and A. Sen (1990), "Fairness and Strategic Behavior in Two-Person, Alternating Offer Games: Results from Bargaining Experiments," Unpublished manuscript, University of California, Los Angeles.
- Sutton, J (1986), "Non-Cooperative Bargaining Theory: An Introduction," *Review of Economic Studies*, 53, 709-724.
- 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*, ed. by R. Henn and O. Moeschlin, Lecture Notes in Economics and Mathematical Systems No. 141. Berlin: Springer-Verlag.

APPENDIX A: TWO-DIMENSIONAL SPATIAL PROBLEMS.

There are, of course, bargaining problems for which neither Theorem IV nor Theorem V applies, either because the policy space is multidimensional or because there is no essential player. This is true of a class of problems that has played an extremely important role in political science theory.³³ These are spatial problems in which the policy space is two-dimensional. The informal discussion below summarizes what can be inferred about this class of problems by applying numerical simulation techniques.

First, for every problem that has a nonempty core, we have been able to compute a deterministic solution for the model derived from that problem. For example, the core is nonempty for every four-person two-dimensional problem with strict majority rule and we have computed solutions to hundreds of corresponding models. Second, the closer a problem is to one with a nonempty core, the more likely it is that the model derived from it will have an exact solution.³⁴ Moreover, if a solution is not exact, it is more likely to be almost exact. Finally, the outcomes implemented by these exact or approximate solutions are more likely to be close to the core of the neighboring problem.

To demonstrate the relationship between the structural characteristics of bargaining problems and the frequency of different solution types for the corresponding models, we report on three Monte Carlo experiments, referred to as experiments A, B and C. In each experiment we sample one hundred bargaining problems. The sample spaces are three increasingly general, parameterized families of five-person spatial problems, ranging from a family in which the core is always "almost nonempty" to one in which only minimal restrictions are imposed. In all three experiments, access probabilities are sampled from the four-dimensional unit simplex, and players' risk aversion coefficients lie on the unit interval. The sample spaces for players' ideal points are displayed in the table below, with $\alpha_i = (\alpha_{i1}, \alpha_{i2})$ denoting player i 's ideal point. In experiments A and B, only α_3 is selected randomly while in experiment C, all five α_i 's are randomly chosen.

Sample Spaces for Players' Ideal Points in Experiments A-C.					
	Player #1	Player #2	Player #3	Player #4	Player #5
A:	$\alpha_1 = (-1,0)$	$\alpha_2 = (+1,0)$	$(-.05,-.05) \leq \alpha_3 \leq (.05,.05)$	$\alpha_4 = (0,-1)$	$\alpha_5 = (0,+1)$
B:	$\alpha_1 = (-1,0)$	$\alpha_2 = (+1,0)$	$(-0.5,-0.5) \leq \alpha_3 \leq (0.5,0.5)$	$\alpha_4 = (0,-1)$	$\alpha_5 = (0,+1)$
C:	$(-1,-1) \leq \alpha_1 \leq (1,1)$	$(-1,-1) \leq \alpha_2 \leq (1,1)$	$(-1,-1) \leq \alpha_3 \leq (1,1)$	$(-1,-1) \leq \alpha_4 \leq (1,1)$	$(-1,-1) \leq \alpha_5 \leq (1,1)$

It is well known that in experiments A and B, the core is nonempty if and only if player #3's ideal point is located at the origin.³⁵

The results of the three experiments are summarized in the three histograms presented in Table 5.1. In each case, the leftmost column reports the frequency of exact stochastic equilibria. The other columns indicate the frequency of approximate equilibria with different degrees of inexactness: specifically, the height of the bar labelled "from a to b " represents the frequency with which we computed an ϵ -equilibrium with $\epsilon \in (a, b]$. The experimental results are consistent with the qualitative remarks offered above. In particular, in experiment A the likelihood of an exact equilibrium is very high, while virtually all of the approximate equilibria are almost exact. As the class of problems is expanded, the likelihood of an exact equilibrium declines, and the likelihood of a quite inexact solution increases. Of course, because of the methodology used here, these statistics are necessarily subject to certain caveats. In particular, while we have observed stable cycles over thousands of rounds, and inferred from

³³ The seminal papers in this literature are Davis and Hinich [1966] and Plott [1967]. For surveys of the literature, see Enelow and Hinich [1984] and Ordeshook [1986].

³⁴ Because our sample spaces are all finite-dimensional, the notion of "close" here is the standard one. Also, the sample spaces in the following discussion are all endowed with Lebesgue measure, and terms such as "more likely" have precise meanings in terms of this measure.

³⁵ See section 4.7 of Ordeshook [1986], and Fiorina-Plott [1978].

Table 5.2: A Two-period Limit Cycle.

Rnd	Prpr	Coal	x_1	x_2	$u_1(\cdot)$	$u_2(\cdot)$	$u_3(\cdot)$	$u_4(\cdot)$	$u_5(\cdot)$
1000	Eu				1.77076	3.31871	26.29755	1.55071	2.82409
999	#1	{3,4}	-0.479479	0.078169	1.77245	3.31335	26.29755*	1.55071*	2.82227
	#2	{4,5}	0.263466	0.150188	1.76932	3.32546	26.24555	1.55071*	2.82409*
	#3	{1,5}	-0.193165	0.201892	1.77117	3.31787	26.40391	1.55060	2.82492
	#4	{1,2}	-0.120128	-0.300296	1.77076*	3.31871*	26.22866	1.55204	2.81902
	#5	{2,3}	-0.054990	0.481519	1.77021	3.31871*	26.29755*	1.54985	2.82856
998	Eu				1.77066	3.31921	26.29644	1.55073	2.82408
997	#1	{3,4}	-0.478789	0.068909	1.77245	3.31337	26.29644*	1.55073*	2.82218
	#2	{4,5}	0.248006	0.144683	1.76939	3.32523	26.25072	1.55073*	2.82408*
	#3	{1,5}	-0.193165	0.201892	1.77117	3.31787	26.40391	1.55060	2.82492
	#4	{1,2}	-0.091547	-0.289190	1.77066*	3.31921*	26.23077	1.55202	2.81918
	#5	{1,3}	-0.198921	0.516905	1.77066*	3.31628	26.29644*	1.54971	2.82855
⋮									
6	Eu				1.77066	3.31921	26.29644	1.55073	2.82408
5	#1	{3,4}	-0.478789	0.068909	1.77245	3.31337	26.29644*	1.55073*	2.82218
	#2	{4,5}	0.248006	0.144683	1.76939	3.32523	26.25072	1.55073*	2.82408*
	#3	{1,5}	-0.193165	0.201892	1.77117	3.31787	26.40391	1.55060	2.82492
	#4	{1,2}	-0.091547	-0.289190	1.77066*	3.31921*	26.23077	1.55202	2.81918
	#5	{1,3}	-0.198921	0.516905	1.77066*	3.31628	26.29644*	1.54971	2.82855
4	Eu				1.77076	3.31871	26.29755	1.55071	2.82409
3	#1	{3,4}	-0.479479	0.078169	1.77245	3.31335	26.29755*	1.55071*	2.82227
	#2	{4,5}	0.263466	0.150188	1.76932	3.32546	26.24555	1.55071*	2.82409*
	#3	{1,5}	-0.193165	0.201892	1.77117	3.31787	26.40391	1.55060	2.82492
	#4	{1,2}	-0.120128	-0.300296	1.77076*	3.31871*	26.22866	1.55204	2.81902
	#5	{2,3}	-0.054990	0.481519	1.77021	3.31871*	26.29755*	1.54985	2.82856
2	Eu				1.77066	3.31921	26.29644	1.55073	2.82408
1	#1	{3,4}	-0.478789	0.068909	1.77245	3.31337	26.29644*	1.55073*	2.82218
	#2	{4,5}	0.248006	0.144683	1.76939	3.32523	26.25072	1.55073*	2.82408*
	#3	{1,5}	-0.193165	0.201892	1.77117	3.31787	26.40391	1.55060	2.82492
	#4	{1,2}	-0.091547	-0.289190	1.77066*	3.31921*	26.23077	1.55202	2.81918
	#5	{1,3}	-0.198921	0.516905	1.77066*	3.31628	26.29644*	1.54971	2.82855

round to round. Conversely, a cyclic pattern emerges whenever at least one player's optimal coalition choice regularly changes. Clearly, it will be extremely difficult to ensure that players' optimal coalition choices remain constant, and hence to guarantee existence of an exact stochastic solution. On the other hand, our Monte Carlo experiments do suggest that there is a predictable relationship between the structure of a given family of problems and the relative likelihood of exact, almost exact and inexact solutions for the corresponding models.

The preceding discussion raises a wide variety of questions such as: What is the precise relationship between the dynamical system we have been investigating and static concepts such as the core? Is there an necessary relationship between the likelihood of a nonempty approximate core and the frequency of an exact or almost exact solution; conversely, are there families of problems for which even approximate cores are usually empty, yet exact or approximate solutions to the corresponding models arise as frequently, say, as in experiment A? Why are the cyclic patterns we have observed so prevalent, rather than, say, some kind of chaotic behavior? To what extent are the various observations reported above robust with respect to alternative functional forms?

APPENDIX B: PROOFS.

Proof of Proposition I and Lemma I: The proofs of Proposition I and Lemma I are interwoven. We first establish part (i) of the proposition for $t = T$. Consider a policy vector $\bar{x} \in X$ such that $u_i(\bar{x}) < u_i(x^{df\mu})$. Clearly, if (1) round $T-1$ is reached, (2) some player proposes \bar{x} and (3) i has the deciding vote, then i does strictly worse if she accepts \bar{x} than if she rejects it. Similarly, for x such that $u_i(x) > u_i(x^{df\mu})$, i does strictly worse if she rejects x than if she accepts it. Moreover, in either case, conditions (1)-(3) are indeed satisfied if each $j \neq i$ plays as follows: $A_{j,T} = X$; $x_{j,T-1} = \bar{x}$; and for each $t \in \{2, 4, \dots, T-2\}$, $A_{j,t} = \emptyset$. This establishes that if $s_i \in S_{i,T}$, then i 's acceptance set in the last period must contain the set $\{x \in X: u_i(x) > u_i(x^{df\mu})\}$ and exclude the set $\{x \in X: u_i(x) < u_i(x^{df\mu})\}$. To complete the proof of part (i), observe that acceptance sets are required to be closed.

We now prove parts (a) and (b) of the Lemma, for $t = T-1$. Let $J = \{j \in I: \bar{x}_{j,T-1} \in \partial \bar{A}_{j,T-1}\}$.⁴⁰ If J is empty, then $\bar{x}_{i,T-1} \in \text{interior}(\bigcap_{j \in \bar{C}_{i,T-1}} \bar{A}_{j,T-1})$ and part (a) follows immediately from Assumption A3. Assume therefore, that J is nonempty. We will show that for all $j \in J$, $\bar{x}_{j,T-1} \neq \bar{x}_{i,T-1}$. It follows from part (i) of Proposition I that for all $j \in J$, $u_j(\bar{x}_{i,T-1}) = u_j(x^{df\mu})$. From assumption A4, however, there exists \bar{x} such that $u_{j'}(\bar{x}) > u_{j'}(x^{df\mu})$, for all $j' \in I$. Since any coalition of players must accept \bar{x} if it is proposed, it follows that for any player $j \in J$, $u_j(\bar{x}_{j,T-1}) \geq u_j(\bar{x}) > u_j(\bar{x}_{i,T-1})$, verifying that as claimed, $\bar{x}_{i,T-1} \neq \bar{x}_{j,T-1}$.

For $t = T-1$, part (b) of the Lemma is an immediate implication of Assumption A4. As noted above, for every player i , the vector \bar{x} identified by Assumption A4 will be accepted by all players and yields i a strictly higher payoff than $Eu_i(\bar{s} | T)$.

We now return to the proposition, to prove part (ii) for $t = T-1$. After elimination of weakly dominated strategies in round T , player j is left with a unique admissible choice in round T : the acceptance set $\{x \in X: u_j(x) \geq u_j(x^{df\mu})\}$. Part (ii) now follows immediately from this fact and part (b) of the Lemma with $t = T-1$.

Now fix $t = \{2, 4, \dots, T-2\}$ and assume that part (i) of the Proposition has been proved for round $t+2$ while part (ii) of the Proposition and parts (a) and (b) of the Lemma have been proved for round $t+1$. Part (i) of the Proposition can now be proved for t , using exactly the same argument as we used for $t = T$. Now consider parts (a) and (b) of the Lemma, for round $t-1$. If round $t+1$ of the game is reached, then the vector of offers $(\bar{x}_{i,t+1})_{i \in I}$ will be proposed and accepted. Let $Ex_{t+1} = \sum_{i \in I} w_i \bar{x}_{i,t+1}$. Because the offers in this round are not all identical, it follows from the strict concavity payoffs that $u_{j'}(Ex_{t+1}) > Eu_{j'}(\bar{s} | t+2)$, for every $j' \in I$. Now repeat the argument proving parts (a) and (b) for $t = T-1$, but replace \bar{x} with Ex_{t+1} . Finally, part (ii) of the proposition for round $t-1$ can be proved by exactly the same argument that was used to prove part (ii) for round $T-1$. \square

⁴⁰ Given a set X , the symbol " ∂X " denotes the boundary of X .

Proof of the Corollary to Theorem I: Let $\bar{s} = (\bar{s}_i)_{i \in I}$ be an equilibrium profile for the game $\Gamma(u, T)$, where $\bar{s}_i = (\bar{x}_{i,t}, \bar{C}_{i,t}, \bar{A}_{i,t+1})_{t=1,3,\dots,T-1}$. Suppose that for some i and $t \in \{1, 3, \dots, T-1\}$, $\bar{C}_{i,t}$ is not minimal with respect to i . Then there exists $C' \subset \bar{C}_{i,t}$, $C' \neq \bar{C}_{i,t}$, such that $C' \cup \{i\}$ is admissible. Thus, $\bigcap_{j \in \bar{C}_{i,t}} \bar{A}_{j,t+1} \subset \bigcap_{j \in C'} \bar{A}_{j,t+1}$ while by Proposition I, $\max\{u_i(\cdot): x \in \bigcap_{j \in \bar{C}_{i,t}} \bar{A}_{j,t+1}\} \geq \max\{u_i(\cdot): x \in \bigcap_{j \in C'} \bar{A}_{j,t+1}\}$. Since $u_i(\cdot)$ is strictly concave, the maximizers on the two constraint sets must coincide. Moreover, from Lemma I(b) and Proposition I, $u_i(\bar{x}_{i,t}) \in \text{interior}(\bar{A}_{i,t+1})$, so that $\bar{x}_{i,t}$ is also a maximizer on $\bigcap_{j \in C' \cup \{i\}} \bar{A}_{j,t+1}$. Thus the profile \bar{s} remains an equilibrium after substituting the coalition C' for $\bar{C}_{i,t}$. \square

Proof of Theorem II: While the existence result is immediate, the proof of uniqueness is extremely intricate. Accordingly, we precede it with an heuristic guide. Recall that in each offer round, player i solves a two-part maximization problem. She first considers each admissible coalition in turn and maximizes her utility subject to the condition that all members of that coalition must accept her choice. For each coalition, our strict concavity conditions guarantee a unique optimal choice. She then chooses a utility-maximal policy from among these maximizers. To guarantee that a game has a unique equilibrium, it is sufficient to ensure that for each player in each round, there is a unique solution to the second stage of her maximization problem. As usual, we start from the end of the game and work backwards. In round $T-1$, we accomplish this for generic games simply by increasing slightly each player's utility on a small neighborhood of one of her optimal choices. In round $t-1$, for $t < T$, however, the problem is much more delicate; to obtain uniqueness in this round, we must locally perturb player's utilities without interfering with any of the adjustments we have already made. Our approach is to arrange things, whenever possible, in round $t+1$ so that the offers players make in round $t-1$ are distinct from all of the offers they make in later rounds. When things can be arranged in this way, we can simply perturb players' utilities on neighborhoods of their $t-1$ round offers, without affecting any of our previous perturbations. It is not always possible, however, to ensure that offers in different rounds are always distinct: there is an open set of utility functions for which at least one player repeatedly proposes her ideal point. This fact dramatically complicates the proof. Fortunately, however, there are only finitely many exceptional cases that must be dealt with; because we have a continuum of degrees of freedom, we can take care of these cases through an intricate process of "anticipatory planning."

Definitions: Fix a twice continuously differentiable function $v \in \mathbb{U}$. Let $\alpha(v) = \{\alpha_i(v)\}_{i \in I}$ denote the vector of players' ideal points; that is, for each i , $\alpha_i(v)$ globally maximizes $v_i(\cdot)$ on X . When confusion will not result, we will use boldface lowercase letters to denote both *vectors* and the *sets* corresponding to these vectors.⁴¹ For $i \in I$ and $x \in X$, let $Hv_i(x)$ denote the Hessian matrix of v_i evaluated at x . Since X is compact, and v is strictly concave, there exists $\eta(v) > 0$ such that for all i and all unit length vectors $\beta \in \mathbb{R}^{\bar{i}}$, $\beta' Hv_i(x) \beta < -\eta(v)$. Now, for each pair of vectors $y \in X^{\bar{i}}$ and $\varepsilon \in \mathbb{R}_+^{\bar{i}}$, and scalar $\delta > 0$, define

⁴¹ For example, $\alpha(v)$ will sometimes denote the set $\bigcup_i \alpha_i(v)$, and sometimes the ordered \bar{i} -tuple $\{\alpha_i(v)\}_{i \in I}$.

$$\Psi^v(y, \varepsilon, \delta) = \{\psi: X \rightarrow \mathbb{R}_+^I: \psi(x^{dfu}) = 0 \text{ and for all } i,$$

for all $x \in X$ and all unit length vectors $\beta \in \mathbb{R}^I, \beta'H\psi(x)\beta < \eta(v) ;$

if $\varepsilon_i = 0$ then $\psi_i(\cdot) \equiv 0;$

if $\varepsilon_i > 0$ then $\psi_i(y_i) = \varepsilon_i > \psi_i(y'_i)$, for all $y'_i \neq y_i$

$\psi_i(\cdot) = 0$ on $X \sim B(y, \delta)$.

Observe that for all $\psi \in \Psi^v(\cdot, \cdot, \cdot)$, $v + \psi \in \mathbb{W}$. For the remainder of the proof, we will assume without further comment that the symbols t, T and τ denote *even* integers. For each $t \leq T$, define the set $V(t)$ as follows.

$$V(t) = \{v \in \mathbb{W} : \forall \tau \in [t, T),$$

each player has a unique optimal policy choice in round $\tau+1$ of $\Gamma(v, T)\}$

In the definitions that follow, we always presume that $v \in V(t)$. Let $z(v, T+1) = (x^{dfu}, \dots, x^{dfu})$ and for $\tau \in [t, T)$, let $z(v, \tau+1)$ denote the vector of optimal choices for players in round $\tau+1$ of $\Gamma(v, T)$. Let $Z(v, t) = \bigcup_{\tau \in [t, T)} z(v, \tau+1)$. (Observe that for all v , $Z(v, T) = \emptyset$). Let $\pi(v, \tau) = \sum_j w_j v(z_j(v, \tau+1))$ denote the vector

of players' reservation utilities in round τ . For each i and $\tau \in [t, T)$, let $L_i(v, \tau) = \{y \in X: v_i(y) = \pi_i(v, \tau)\}$ denote the "lower" boundary of player i 's acceptance set in round τ of $\Gamma(v, T)$. Let $A(v, \tau) = \bigcap_{j \in CC \in \mathbb{C}} \text{co}(L_j(v, \tau))$

denote the set of policies that will be accepted by some coalition in round τ .⁴² For each i , let $M_i(v, t-1) = \{y \in X: y \text{ maximizes } v_i(\cdot) \text{ on } A(v, t)\}$. Observe that for $v \in V(t)$, $M_i(v, t-1)$ is necessarily a finite set, but in general may contain more than one element. Also, observe that

$$\text{for all } y \in M_i(v, t-1), \text{ either } y = \alpha_i(v) \text{ or there exists } j \neq i \text{ such that } y \in L_j(v, t). \quad (\text{B.II.1})$$

To see that (B.II.1) is true, recall that as an immediate implication of Lemma I(b), each player's optimal offer must yield her more utility than her reservation utility in the following round. Thus, $M_i(v, t-1) \cap L_i(v, t)$ must be empty. Therefore, if $y \in M_i(v, t-1)$ but $y \notin \bigcup_{j \neq i} L_j(v, t)$, then y must be an interior point of $A(v, t)$ and hence a global maximizer of $v_i(\cdot)$ on X . For $t' \geq t$, $v^t \in V(t)$ and $v^{t'} \in V(t')$, we will say that v^t and $v^{t'}$ are *strategically equivalent from t'* if the following three conditions are satisfied: (a) $\alpha(v^t) = \alpha(v^{t'})$; and for all $\tau \geq t'$, (b) $z(v^t, \tau+1) = z(v^{t'}, \tau+1)$; and (c) $v^t(z(v^t, \tau+1)) = v^{t'}(z(v^{t'}, \tau+1))$. Observe that strategic equivalence is transitive in the following sense: if for $t < t' < t''$, v^t and $v^{t'}$ are strategically equivalent from t'' while v^t and $v^{t'}$ are strategically equivalent from t' , then v^t and $v^{t'}$ are strategically equivalent from t'' .

The construction that follows relates to the "intricate process of anticipatory planning" referred to above. Let I_{-i} denote the set $I \sim \{i\}$. For each $\pi \in \mathbb{R}$ and nonempty set $J \in I_{-i}$, define the affine function $\theta_i(\pi, J)$ by

$$\theta_i(\pi, J) = \left[w_i v_i(\alpha_i(v)) + \sum_{i \in J \cup \{i\}} w_i v_i(\alpha_i(v)) \right] + \sum_{i \in J} w_i \pi.$$

Now, for each scalar $\pi > 0$, define $P_i^0(v, \pi) = \{\pi\}$ and for each even integer $k > 0$, define

⁴² "co(Y)" denotes the convex hull of Y .

$$P_i^k(v, \pi) = \bigcup_{\pi \in \Pi_i^{k-2}(v)} \bigcup_{\emptyset \neq J \subset I_{-i}} \theta_i(\pi', J). \text{ Obviously,}$$

$$\text{for any finite set } \Omega \subset \mathbb{R}, \text{ the set of } \pi\text{'s for which } P_i^k(v, \pi) \cap v_i(\Omega) \neq \emptyset \text{ is finite;} \quad (\text{B.II.2a})$$

Also, since $\sum_{i \in J} w_i \in (0, 1)$, for all $\emptyset \neq J \subset I_{-i}$, it follows that for all even k , $\pi' \in P_i^{k-2}(v, \pi) \cap P_i^k(v, \pi)$ only if π' is a fixed point one of some member of a fixed finite set of affine functions, each of which has at most one fixed point. Hence,

$$\text{for all even } \kappa > k \geq 0, \text{ the set of } \pi\text{'s for which } P_i^k(v, \pi) \cap P_i^\kappa(v, \pi) \neq \emptyset \text{ is finite.} \quad (\text{B.II.2b})$$

Now for each even integer $k \in [0, t]$, let $\Pi_i^k(v, t) = P_i^k(v, \pi_i(v, t))$. $\Pi_i^k(v, t)$ has the following interpretation. For each $\tau \in (t-k, t]$, suppose that in round $\tau-1$, player i proposes her ideal point, at least one player $j \neq i$ proposes a policy that yields i her reservation utility in round τ , and the remaining players propose their ideal points. In this event, player i 's reservation utility in round $t-k$ must belong to the finite set $\Pi_i^k(v, t)$. The relevance of this arcane fact will eventually become apparent. For now, we will simply assert that the following fact plays a critical role in the proof.

$$\text{For all } v \in V(t-2), \text{ either for each } k \in [0, t-2], \Pi_i^k(v, t-2) \subset \Pi_i^{k+2}(v, t). \quad (\text{B.II.3})$$

$$\text{or } z(v, t-1) = \alpha(v)$$

$$\text{or there exists } t^1 \in I \text{ and } t^2 \neq i \text{ s.t. } z_{t^1}(v, t-1) \in L_{t^2}(v, t) - L_i(v, t)$$

To see that (B.II.3) is true, observe first that from (B.II.1), if neither the second nor third conditions are satisfied, then, necessarily, $z_i(v, t-1) = \alpha_i(v)$, while for some $\emptyset \neq J \subset I_{-i}$, $z_t(v, t-1) \in L_i(v, t)$, for $t \in J$, and otherwise, $z_t(v, t-1) = \alpha_t(v)$. In this case, player i 's reservation utility in round $t-2$, $\pi_i(v, t-2) = \theta_i(\pi_i(v, t), J) \in \Pi_i^0(v, t-2) \subset \Pi_i^2(v, t)$. Moreover, by construction, if $\Pi_i^{k-2}(v, t-2) \subset \Pi_i^k(v, t)$, for some $k \in (0, t-2]$, then $\Pi_i^k(v, t-2) \subset \Pi_i^{k+2}(v, t)$.

Finally, let $\Omega \subset \mathbb{R}$ be an arbitrary finite set and define the set $U(t; \Omega)$ as follows:

$$U(t; \Omega) = \{v \in \mathbb{W} : \text{for all } i \text{ and all } k \in [0, t], \Pi_i^k(v, t) \cap v_i(\Omega) = \emptyset\}^{43}$$

$$\text{for all even } \kappa \in (k, t], \Pi_i^k(v, t) \cap \Pi_i^\kappa(v, t) = \emptyset$$

Observe that $\Omega' \subset \Omega$ implies $U(t; \Omega) \subset U(t; \Omega')$. Also observe that that

$$\Omega' \subset \Omega \text{ implies } U(t; \Omega) \subset U(t; \Omega'). \quad (\text{B.II.4a})$$

$$\text{for all } v \in U(t; Z(v, t)) \text{ and all } i \in I, L_i(v, t) \cap Z(v, t) = \emptyset. \quad (\text{B.II.4b})$$

(B.II.4a) is immediate. (B.II.4b) is equivalent to the statement that $\Pi_i^0(v, t) \cap v_i(Z(v, t)) = \emptyset$. An immediate consequence of (B.II.1) and (B.II.4) is that for any finite set Ω that contains $Z(v, t)$,

$$\text{for all } v \in U(t; \Omega), \text{ if } y \in M_i(v, t-1) \text{ and } y \neq \alpha_i(v), \text{ then } y \notin Z(v, t). \quad (\text{B.II.5})$$

⁴³ For a real-valued function f , " $f(Y)$ " denotes the image of the set Y under the map f , i.e., $\{f(y) \in \mathbb{R} : y \in Y\}$.

We now present three lemmas, from which the proof of the Theorem will follow (relatively) easily.

Lemma II.1: Fix $t \leq T$, $x \in X^T$ and $v \in V(t) \cap U(t; \Omega)$, for some finite set Ω that contains $Z(v, t) \cup \alpha(v)$. There exists $\bar{\delta} > 0$ and $\gamma > 0$ such that for all $\delta \in (0, \bar{\delta})$ if $\varepsilon \in B(0, \gamma)$ with $\varepsilon_i = 0$ whenever $x_i \in \Omega$, then for all $\psi \in \Psi^v(x, \varepsilon, \delta)$, v and $v + \psi$ are strategically equivalent from t .

Proof of Lemma II.1: Pick $\bar{\delta} > 0$ sufficiently small that for all i : (δ -i) for all i such that $x_i \notin Z(v, t)$, $B(x_i, \delta) \cap B(Z(v, t), \delta) = \emptyset$ and (δ -ii) for all $\tau \geq t$, $x_i \notin A(v, \tau)$ implies $B(x_i, \delta) \cap A(v, \tau) = \emptyset$. Pick $\gamma > 0$ such that: (γ -i) for all $\tau > t$ and all i such that $x_i \notin Z(v, t)$, $x_i \in A(v, \tau)$ implies that $v_i(B(x_i, \delta)) < v_i(z_i(v, \tau-1)) - \gamma$ and (γ -ii) for all $\varepsilon \in B(0, \gamma)$, $\Psi^v(x, \varepsilon, \delta)$ is nonempty.⁴⁴ Now pick $\delta \in (0, \bar{\delta})$ and $\varepsilon \in B(0, \gamma)$ with $\varepsilon_i > 0$ whenever $x_i \notin Z(v, t)$ and pick $\psi \in \Psi^v(x, \varepsilon, \delta)$. Observe that by definition $z(v, T+1) \equiv z(v+\psi, T+1) \equiv (x^{df^t}, \dots, x^{df^t})$ and, by assumption, $\psi(x^{df^t}) = \psi(\alpha(v)) = 0$. Hence v and $v + \psi$ are strategically equivalent from T . Now fix $\tau \in (t, T]$ and assume that v and $v + \psi$ are strategically equivalent from τ . We will show that they are strategically equivalent from $\tau-2$.

Fix $i \in I$. Together with our assumption that $x_j \notin Z(v, t)$ whenever $\varepsilon_j > 0$ condition (δ -ii) implies that there exists a nbd N_i of $z_i(v, \tau-1)$ such that for each j , $\psi_j(\cdot) = 0$ on N_i . Moreover, since strategic equivalence from τ implies that $\psi(z(v, \tau+1)) = 0$, we have $\pi(v+\psi, \tau) = \pi(v, \tau)$. Thus, for each j and $y \in N_i$, $v_j(y) \geq \pi_j(v, \tau)$ if and only if $(v_j+\psi_j)(y) \geq \pi_j(v+\psi, \tau)$. It follows that $A(v+\psi, \tau) \cap N_i = A(v, \tau) \cap N_i$ and so $z_i(v, \tau-1) \in A(v+\psi, \tau)$. Moreover, since $\psi(\cdot)$ is nonnegative it follows that for all j and all $y \in A(v, \tau)$, $(v_j+\psi_j)(y) \geq \pi_j(v, \tau) = \pi_j(v+\psi, \tau)$. Thus $A(v+\psi, \tau) \subset A(v, \tau)$. We now have two cases to consider. First, suppose that $x_i \notin A(v, \tau)$. In this case, condition (δ -ii) implies that $B(x_i, \delta) \cap A(v, \tau) = \emptyset$ so that $\psi_i(\cdot) = 0$ on $A(v, \tau)$. Second, suppose that $x_i \in A(v, \tau)$. In this case, $\psi_i(\cdot) = 0$ on the set $A(v+\psi, \tau) - B(x_i, \delta)$, while condition (γ -i) and the definition of ψ_i imply that for all $y \in B(x_i, \delta)$, $v_i(z_i(v, \tau-1)) - v_i(y) > \gamma > \psi_i(y)$. This establishes that for all $z_i(v, \tau-1) \neq y \in A(v+\psi, \tau)$, $(v_i+\psi_i)(z_i(v, \tau-1)) > (v_i+\psi_i)(y)$ and proves that $z_i(v+\psi, \tau-1) = z_i(v, \tau-1)$. It now follows immediately from the definition of δ that $\psi(z(v, \tau-2)) = 0$. \square

Lemma II.2: Fix $t \leq T$ and $v \in V(t) \cap U(t; \Omega)$, for some finite set Ω that contains $Z(v, t) \cup \alpha(v)$. Fix $x \in X^T$ such that for each i , $x_i \in M_i(v, t-1)$. There exists $\gamma > 0$ and $\bar{\delta} > 0$ such that for all $\delta \in (0, \bar{\delta})$ if $\varepsilon \in B(0, \gamma)$ with $\varepsilon_i = 0$ whenever $x_i = \alpha_i(v)$, then for all $\psi \in \Psi^v(x, \varepsilon, \delta)$, v and $v + \psi$ are strategically equivalent from t and $v+\psi \in V(t-2)$.

Proof of Lemma II.2: Pick $\gamma > 0$ and $\bar{\delta} > 0$ sufficiently small that the conclusion of Lemma II.1 holds for $(t, v, x, \gamma, \bar{\delta})$. We can assume w.l.o.g. that $\bar{\delta}$ is sufficiently small that the following conditions are satisfied in addition: (δ -iii) for all $x_j \neq x_i$, $B(x_j, \delta) \cap B(x_i, \delta) = \emptyset$; and (δ -iv) if $x_i \in \text{int}(\text{co}(L_i(v, t)))$, then $B(x_i, \delta) \subset \text{co}(L_i(v, t))$. Pick $\delta \in (0, \bar{\delta})$ and $\varepsilon \in B(0, \gamma)$ with $\varepsilon_i = 0$ whenever $x_i = \alpha_i(v)$ and pick

⁴⁴ A γ satisfying (γ -i) exists because by assumption, $z_i(v, \tau-1) \neq x_i$ is the unique maximizer of $v_i(\cdot)$ on $A(v, \tau)$.

$\psi \in \Psi^v(x, \varepsilon, \delta)$. We first establish that the hypothesis in Lemma II.1 is satisfied. From (B.II.1), if $\varepsilon_i > 0$ then there exists $j \neq i$ such that $x_i \in L_j(v, t)$, i.e., $v_j(x_i) = \pi_j(v, t) \in \Pi_j^0(v, t)$. Since $v \in U(t; Z(v, t))$, $\Pi_j^0(v, t) \cap v_j(Z(v, t)) = \emptyset$. Therefore, $x_i \notin Z(v, t)$. It now follows from Lemma II.1 that for each $\psi \in \Psi^v(x, \varepsilon, \delta)$ v and $v + \psi$ are strategically equivalent from t .

We next establish that for each $i \in I$, $x_i \in A(v + \psi, t)$. Since $x_i \in M_i(v, t-1)$, there exists a coalition $C \in \mathcal{C}$ such that $x_i \in \bigcap_{j \in C} \text{co}(L_j(v, t))$. Suppose that for $j \in C$, $x_j = x_i$. It follows from Lemma I(b) that $x_i \in \text{int}(\text{co}(L_j(v, t)))$; condition $(\delta\text{-iv})$ now implies that $\psi_j(L_j(v, t)) = 0$. Thus, $L_j(v, t) = L_j(v + \psi, t)$, and so $x_i \in \text{co}(L_j(v + \psi))$. Now suppose that for $j \in C$, $x_j \neq x_i$. In this case, condition $(\delta\text{-iii})$ implies that $\psi_j(\cdot) = 0$ on a nbd N_j of x_i . Therefore, $L_j(v + \psi, t) \cap N_j = L_j(v, t) \cap N_j$. We have established, therefore, that $x_i \in \bigcap_{j \in C} \text{co}(L_j(v + \psi, t)) \subset A(v + \psi, t)$. We now show that x_i is the unique element of $M_i(v + \psi, t-1)$. Since $\psi(\cdot)$ is nonnegative, $(v_j + \psi_j)(y) \geq \pi_j(v, t) = \pi_j(v + \psi, t)$, for all j and all $y \in A(v, t)$. Thus $A(v + \psi, t) \subset A(v, t)$. First suppose that $x_i \neq \alpha_i(v)$. In this case $\psi_i(\cdot)$ attains a unique maximum at x_i . Moreover, x_i maximizes $v_i(\cdot)$ on $A(v, t)$. Therefore, $(v_i + \psi_i)(\cdot)$ attains a unique maximum on $A(v + \psi, t)$ at x_i . If $x_i = \alpha_i(v)$, then x_i is the unique global maximizer of $v_i(\cdot)$ and $\psi_i(\cdot) \equiv 0$. Once again, therefore, $(v_i + \psi_i)(\cdot)$ attains a unique maximum on $A(v + \psi, t)$ at x_i . This completes the proof that $v + \psi \in V(t-2)$ \square

Lemma II.3: Fix $t \leq T$ and $v \in V(t-2) \cap U(t; \Omega)$, for some finite set Ω that contains $Z(v, t) \cup \alpha(v)$. If $z(v, t-1) \neq \alpha(v)$, then for all $\bar{\gamma} > 0$, there exists $x \in X^{\bar{\gamma}}$, $\gamma \in (0, \bar{\gamma})$, $\varepsilon \in B(0, \gamma)$, $\delta > 0$ and $\phi \in \Psi^v(x, \varepsilon, \delta)$ such that v and $v + \phi$ are strategically equivalent from t , $z(v, t-1) = z(v + \phi, t-1)$ and $v + \phi \in V(t-2) \cap U(t-2; \Omega \cup z(v, t-1)) \subset U(t-2; Z(v + \phi, t-2) \cup \alpha(v + \phi))$.

Proof of Lemma II.3: Define \bar{I} as follows:

$$\bar{I} = \{i \in I: \text{there exists } i^1 \in I \text{ and } i^2 \neq i \text{ s.t. } z_{i^1}(v, t-1) \in L_{i^2}(v, t) - L_i(v, t)\}.$$

For each $i \in \bar{I}$, pick $i^1(i)$ and $i^2(i)$ such that $z_{i^1(i)}(v, t-1) \in L_{i^2(i)}(v, t) - L_i(v, t)$. (Note that, possibly, $i^1(i) = i$.) Define $x \in X^{\bar{\gamma}}$ as follows: for $i \in \bar{I}$, $x_i = z_{i^1(i)}(v, t-1)$; otherwise, pick x_i arbitrarily. Pick $\delta > 0$ sufficiently small that: $(\delta\text{-iii})$ for all i and all $x_j \neq z_i(v, t-1)$, $B(x_j, \delta) \cap B(z_i(v, t-1), \delta) = \emptyset$; and $(\delta\text{-iv})$ if $x_j \in \text{int}(\text{co}(L_j(v, t)))$, then $B(x_j, \delta) \subset \text{co}(L_j(v, t))$. Pick $\gamma \in (0, \bar{\gamma})$ sufficiently small that $(\gamma\text{-ii})$ $\Psi(x, \varepsilon, \delta)$ is nonempty, for all $\varepsilon \in B(0, \gamma)$ and $(\gamma\text{-iii})$ if $x_i \neq z_i(v, t-1)$, then $v_i(B(x_i, \delta)) < v_i(z_i(v, t-1)) - \gamma$. (Such a $\gamma > 0$ exists because, by assumption, $x_i = z_{i^1(i)}(v, t-1) \in A(v, t)$, while $z_i(v, t-1) \neq x_i$ is the unique maximizer of $v_i(\cdot)$ on $A(v, t)$.) Finally, pick $\varepsilon \in B(0, \gamma)$ as follows: for $i \notin \bar{I}$, set $\varepsilon_i = 0$; For $i \in \bar{I}$, define $q_i(\varepsilon_i) = \pi_i(v, t) + \varepsilon_i \sum_{\{i^k: i^k = i\}} w_{i^k}$ and pick $\varepsilon_i \in (0, \gamma)$ such that for all even $k \in [0, t-2]$, $P_i^k(v, q_i(\varepsilon_i)) \cap v_i(Z(v, t-2) \cup \alpha(v)) = \emptyset$, while for all even $\kappa \in (k, t-2]$, $P_i^k(v, q_i(\varepsilon_i)) \cap P_i^\kappa(v, q_i(\varepsilon_i)) = \emptyset$. Statements (B.II.2a) and (B.II.2b) imply that these conditions are satisfied for all but finitely many ε_i 's. Since $v \in U(t; \Omega)$ and Ω contains $Z(v, t)$, $v_{i^2(i)}(x_i) = \pi_{i^2(i)}(v, t) \notin v_{i^2(i)}(Z(v, t))$, so that $x_i \notin Z(v, t)$, for each i such that $\varepsilon_i > 0$. It follows from Lemma II.1 that when γ and δ are sufficiently small, then for each $\phi \in \Psi^v(x, \varepsilon, \delta)$, v and $v + \phi$ are strategically equivalent from

t . We will assume that γ and δ are indeed sufficiently small and pick $\phi \in \Psi^{\gamma}(x, \varepsilon, \delta)$.

We first establish that $v+\phi \in V(t-2)$, with $z(v+\phi, t-1) = z(v, t-1)$. The argument replicates almost exactly the corresponding argument in the proof of Lemma II.2. We begin by showing that for each $i \in I$, $z_i(v, t-1) \in A(v+\phi, t)$. Pick $C \in \mathbb{C}$ such that $z_i(v, t-1) \in \bigcap_{j \in C} \text{co}(L_j(v, t))$ and consider $i \neq j \in C$. If $j \notin \bar{I}$, then $\phi_j(\cdot) \equiv 0$, so that, trivially, $z_i(v, t-1) \in \text{co}(L_j(v+\phi, t))$. Assume, therefore, that $j \in \bar{I}$. If $x_j = z_j(v, t-1)$, then by construction $x_j = z_{i, z_j}(v, t-1) \notin L_j(v, t-1)$. Therefore, $z_i(v, t-1) \in \text{int}(\text{co}(L_j(v, t)))$ and condition $(\delta\text{-iv})$ implies that $\phi_j(L_j(v, t)) = 0$. Therefore, $L_j(v, t) = L_j(v+\phi, t)$, so that $z_i(v, t-1) \in \text{co}(L_j(v+\phi, t))$. If $x_j \neq z_j(v, t-1)$, then condition $(\delta\text{-iii})$ implies that $\phi_j(\cdot) = 0$ on a nbd N_j of $z_j(v, t-1)$. Therefore, $L_j(v+\phi, t) \cap N_j = L_j(v, t) \cap N_j$, so that once again $z_i(v, t-1) \in \text{co}(L_j(v+\phi, t))$. We have established, therefore, that $x_i \in \bigcap_{j \in C} \text{co}(L_j(v+\phi, t)) \subset A(v+\phi, t)$.

We now show that $z_i(v, t-1)$ is the unique element of $M_i(v+\phi, t-1)$. Replicate the reasoning in the proof of Lemma II.2 to establish that $A(v+\phi, t) \subset A(v, t)$. If $z_i(v, t-1) = x_i$, then x_i is the unique maximizer of both $v_i(\cdot)$ and $\phi_i(\cdot)$ on $A(v+\phi, t)$. If $x_i \neq z_i(v, t-1)$, then $\phi_i(\cdot) = 0$ on the set $A(v+\phi, t) - B(x_i, \delta)$, while condition $(\gamma\text{-iii})$ and the definition of ϕ_i imply that for all $y \in B(x_i, \delta)$, $v_i(z_i(v, t-1)) - v_i(y) > \gamma > \phi_i(y)$. This establishes that for all $z_i(v, t-1) \neq y \in A(v+\phi, t)$, $(v_i+\phi_i)(z_i(v, t-1)) > (v_i+\phi_i)(y)$ and proves that $z_i(v+\phi, t-1) = z_i(v, t-1)$.

To establish that $v + \phi \in U(t-2; Z(v+\phi, t-2) \cup \alpha(v+\phi))$, we must show that for each i and all $k \in [0, t-2]$,

$$\Pi_i^k(v+\phi, t-2) \cap v_i(U(t-2; Z(v+\phi, t-2) \cup \alpha(v+\phi))) = \emptyset; \quad (\text{B.II.6a})$$

$$\text{and for all even } \kappa \in (k, t-2], \Pi_i^k(v+\phi, t-2) \cap \Pi_i^\kappa(v+\phi, t-2) = \emptyset. \quad (\text{B.II.6b})$$

First, note that as a consequence of $(\delta\text{-iii})$, $\phi_i(x_j) = 0$, for all j such that $x_j \neq x_i = z_{i, z_j}(v, t-1)$. Therefore, for each i :

$$\begin{aligned} \pi_i(v+\phi, t-2) &\equiv \sum_i w_i (v_i + \phi_i)(z_i(v+\phi, t-1)) = \sum_i w_i (v_i + \phi_i)(z_i(v, t-1)) \\ &= \sum_i w_i v_i(z_i(v, t-1)) + \varepsilon_i \sum_{\{ix_i=x_i\}} w_i = \pi_i(v, t-2) + \varepsilon_i \sum_{\{ix_i=x_i\}} w_i \\ &\equiv q_i(\varepsilon_i). \end{aligned}$$

Thus for all i , $\Pi_i^k(v+\phi, t-2) \equiv P_i^k(q_i(\varepsilon_i))$. For $i \in \bar{I}$, conditions (B.II.6a) and (B.II.6b) hold by construction of ε_i . For $i \notin \bar{I}$, $\varepsilon_i = 0$ so that $\pi_i(v+\phi, t-2) = q_i(\varepsilon_i) = \pi_i(v, t-2)$. Applying statement (B.II.3) and the fact that $z(v, t-1) \neq \alpha(v)$, $\Pi_i^{k-2}(v+\phi, t-2) \equiv \Pi_i^{k-2}(v, t-2) \subset \Pi_i^k(v, t)$, for each $k \in [2, t]$. Because $v \in U(t; \Omega)$ and Ω contains $Z(v, t) \cup \alpha(v)$, it follows that for all even k and $\kappa \in (k, t-2]$, $\Pi_i^k(v+\phi, t-2) \cap \Pi_i^\kappa(v+\phi, t-2) = \emptyset$ and $\Pi_i^k(v+\phi, t-2) \cap v_i(Z(v, t) \cup \alpha(v)) = \emptyset$. To establish that $v+\phi \in U(t-2; Z(v+\phi, t-2) \cup \alpha(v+\phi))$, the only remaining condition to check is that

$$\text{for all } k \in [0, t-2], \Pi_i^k(v+\phi, t-2) \cap v_i(z(v, t-1)) = \emptyset. \quad (\text{B.II.7})$$

Observe that for all $i \in \bar{I}$, $z_i(v, t-1) = \alpha_i(v)$, while for $j \neq i$, either $z_j(v, t-1) \in L_i(v, t)$, in which case $v_i(z_j(v, t-1)) = \pi_i(v, t) \in \Pi_i^0(v, t)$, or $z_j(v, t-1) = \alpha_j(v)$. In the first instance, (B.II.7) follows from the fact that for $k \in [2, t]$, $\Pi_i^{k-2}(v+\phi, t-2) \cap \Pi_i^0(v, t) \subset \Pi_i^k(v, t) \cap \Pi_i^0(v, t) = \emptyset$. In the second instance, (B.II.7)

follows from the fact that for $k \in [2, t]$, $\Pi_i^{k-2}(v+\phi, t-2) \cap v_i(\alpha(v)) \subset \Pi_i^k(v, t) \cap v_i(\alpha(v)) = \emptyset$. \square

We can now prove Theorem II. Define the set U^* as follows

$$U^* = \{u = (u_i)_{i \in I} \in \mathbb{W} : \text{there exists } T, \text{ an equilibrium s for } \Gamma(u, T), i \in I \text{ and } x \neq x' \text{ such that}$$

$$\text{for some } 1 \leq t \leq T, \text{ both } x \text{ and } x' \text{ maximize } u_i^n(\cdot) \text{ on } \bigcup_{c \in C} \bigcap_{j \in C} \{y \in X : u_j^n(y) \geq Eu_j^n(s | t+2)\}.$$

To prove Theorem II, it is sufficient to show that the closure of U^* has an empty interior. Pick $u^* \in cl(U^*)$ and a sequence (u^n) in U^* such that for every n , $u^n \in B(u^*, n^{-1})$. We will construct a sequence of continuous functions, (v^n) , such that for every n , $v^n \in B(u^n, n^{-1})$, so that (v^n) converges to u^* . We will show that for sufficiently large n , the v^n 's satisfy assumptions A2-A4 and so belong to \mathbb{W} , but do not belong to U^* . The existence of such a sequence will establish that the closure of U^* has an empty interior. We now fix n , drop the n superscript, and replace it with a T , so that u^n becomes u^T . There are two cases to consider

Case I: $\alpha(u^T) \in A(u^T, T)$. In this case, clearly, $z(u^T, T-1) = \alpha(u^T)$, so that $u^T \in V(T-2)$. For each $i \in I$, define $q_i(\varepsilon_i) = \sum_i w_i u_i^T(\alpha_i(u^T)) + \varepsilon_i$ and pick $\varepsilon_i \in (0, (nT)^{-1})$ such that for all even $k \in [0, T-2]$,

$P_i^k(u^T, q_i(\varepsilon_i)) \cap u^T(\alpha(u^T)) = \emptyset$, while for all even $\kappa \in (k, t-2]$, $P_i^k(u^T, q_i(\varepsilon_i)) \cap P_i^\kappa(u^T, q_i(\varepsilon_i)) = \emptyset$. Now define $\phi^T(\cdot) \equiv \varepsilon$ and $u^{T-2} = u^T + \phi^T$. Since $Z(u^{T-2}, T-2) = \alpha(u^{T-2})$, it follows that $u^{T-2} \in V(T-2) \cap U(T-2; Z(u^{T-2}, T-2) \cup \alpha(u^{T-2}))$. Now fix $t < T$ and assume that u^t has been defined such that u^t and u^{t+2} are strategically equivalent from $t+2$, while $u^t \in V(t) \cap U(t; Z(u^t, t) \cup \alpha(u^t))$ (This condition is certainly satisfied for $t = T-2$.) We will construct $u^{t-2} \in B(u^t, 2(nT)^{-1})$ such that u^{t-2} and u^t are strategically equivalent from t while $u^{t-2} \in V(t-2) \cap U(t-2; Z(u^{t-2}, t-2) \cup \alpha(u^{t-2}))$. There are two cases to consider.

Case I(a): For some i , $\alpha_i(u) \in A(u, t)$. Pick $x \in X^T$ such that for each i , $x_i \in M_i(u^t, t-1)$. Applying Lemma II.2, there exists $\delta > 0$, $\varepsilon \in B(0, (nT)^{-1})$ with $\varepsilon_i = 0$ whenever $x_i = \alpha_i(u^t)$ and $\psi^t \in \Psi^{u^t}(x, \varepsilon, \delta)$ such that u^t and $u^t + \psi^t$ are strategically equivalent from t and $u^t + \psi^t \in V(t-2)$. From Lemma II.3, there exists $x \in X^T$, $\gamma \in (0, nT^{-1})$, $\varepsilon \in B(0, \gamma)$, $\delta > 0$ and $\phi^t \in \Psi^{u^t + \psi^t}(x, \varepsilon, \delta)$ such that $u^t + \psi^t$ and $u^t + \psi^t + \phi^t$ are strategically equivalent from t and $u^t + \psi^t + \phi^t \in V(t-2) \cap U(t-2; Z(u^t, t-2) \cup \alpha(u^t))$. Set $u^{t-2} = u^t + \psi^t + \phi^t$ and observe that indeed $u^t \in B(u^{t+2}, 2(nT)^{-1})$, while u^{t-2} and u^t are strategically equivalent from t .

Case I(b): $\alpha(u^t) \subset A(u^t, t)$. In this case, set $u^{t-2} = u^t$. Because strategic equivalence is transitive in the sense defined above, $z(u^{t-2}, t-1) = \alpha(u^{t-2}) = \alpha(u^{T-2}) = z(u^{T-2}, T-1)$ and $\pi(u^{t-2}, t-2) = \pi(u^{T-2}, T-2)$. Clearly, in this case, we can set $u^{t-k} = u^t$, for each $k \in [4, t]$, and observe that $z(u^{t-k}, t-k+1) = z(u^{T-k}, T-k+1)$.

In either case (a) or (b), we can now define $v^n = u^0$. Observe that $v^n \in B(u, n^{-1}) \cap V(0)$, so that $v^n \notin U^*$. Let $\bar{\Omega} = Z(v^n, 0)$.

Case II: $\alpha(u^T) \notin A(u^T, T)$. In this case, set $v^T = u^T$. We first pick $\varepsilon \in B(0, (nT)^{-1})$ such that for even each $k \in [0, T]$, $P_i^k(v^T(x^{dfk}) + \varepsilon) \cap v_i^T(\bar{\Omega} \cup \alpha(v^T)) = \emptyset$, while for all even $\kappa \in (k, T-2]$, $P_i^k(v^T(x^{dfk}) + \varepsilon) \cap P_i^\kappa(v^T(x^{dfk}) + \varepsilon) = \emptyset$. Now define ϕ^T by $\phi^T(X) = 0$, while $\phi^T(x^{dfk}) = \varepsilon$. Define $v^{T-2} = u^T + \phi^T$ and proceed exactly as in case I, with the following exceptions: everywhere replace u 's with v 's; replace

$U(t; Z(u^t, t) \cup \alpha(u^t))$ with $U(t; \bar{\Omega} \cup Z(v^t, t) \cup \alpha(v^t))$; if case (b) applies for some t , then define $v^{t-2} = v^t$ and observe that $z(v^{t-2}, t-1) = \alpha(v^{t-2}) = \alpha(u^{T-2}) = z(u^{T-2}, T-1)$ and $\pi(u^{t-2}, t-2) = \pi(v^{T-2}, T-2)$. Once again, in this case, we can set $v^{t-k} = v^t$, for each $k \in [4, t]$, and observe that $z(v^{t-k}, t-k+1) = z(u^{T-k}, T-k+1)$ \square .

Proof of Theorem IV: We begin by introducing some further notation. Define the mappings $G_i(\cdot)$ and $U_i(\cdot)$ on \mathbf{R} by, for each $\mathbf{x} = (x_j)_{j \in I}$, $G_i(\mathbf{x}) = \{y \in \mathbf{R}: u_i(y) \geq \min_j u_i(x_j)\}$ and $U_i(\mathbf{x}) = \{y \in \mathbf{R}: u_i(y) \geq \sum_j w_j u_i(x_j)\}$. Given a closed set $Y \subset \mathbf{R}$, let $l(Y)$ and $h(Y)$ denote, respectively, the minimal and maximal elements of Y . (Treating \mathbf{x} as a set, we will sometimes refer to $l(\mathbf{x})$ and $h(\mathbf{x})$.) Finally, for each i and proposal profile \mathbf{x} , let $\underline{\beta}_i(\mathbf{x}) = l(G_i(\mathbf{x}))$ and $\bar{\beta}_i(\mathbf{x}) = h(G_i(\mathbf{x}))$. The proof relies on the following Lemma.

Lemma IV.1: For each $\varepsilon > 0$, there exists $\delta > 0$ such that for every $i \in I$ and $\mathbf{x} = (x_i)_{i \in I} \subset X$:

if $h(\mathbf{x}) - l(\mathbf{x}) > \varepsilon$ and either (a) $\alpha_i = x_i$ or (b) $\alpha_i \notin (l(\mathbf{x}), h(\mathbf{x}))$, then

$$U_i(\mathbf{x}) \subset [l(G_i(\mathbf{x})) + \delta, h(G_i(\mathbf{x})) - \delta].$$

Proof of Lemma IV.1: If the Lemma were false, we could find $\varepsilon > 0$, and for every n , a vector \mathbf{x}^n in X and $i \in I$ such that either condition (a) or (b) above is satisfied for i while $U_i(\mathbf{x}^n) \not\subset [l(G_i(\mathbf{x}^n)) + n^{-1}, h(G_i(\mathbf{x}^n)) - n^{-1}]$. Pick a convergent subsequence of the \mathbf{x}^n 's, again indexed by n , such that for some fixed player i , either condition (a) is satisfied for every n or condition (b) is satisfied for every n . Let $\bar{\mathbf{x}}$ be the limit of the subsequence. Clearly $h(G_i(\bar{\mathbf{x}})) - l(G_i(\bar{\mathbf{x}})) \geq \varepsilon$ while there exists $\bar{y} \in U_i(\bar{\mathbf{x}})$ such that either $\bar{y} \leq l(G_i(\bar{\mathbf{x}}))$ or $\bar{y} \geq h(G_i(\bar{\mathbf{x}}))$. Since u_i is strictly concave, it follows that $u_i(\bar{y}) \leq \min_j u_i(\bar{x}_j)$.

First assume that for the identified player i , condition (b) holds for every n . We can assume without loss of generality that for every n , $\alpha_i \geq h(\mathbf{x}^n)$ so that $\alpha_i \geq h(\bar{\mathbf{x}})$. Because u_i is strictly concave, $u_i(h(\bar{\mathbf{x}})) > \sum_j w_j u_i(\bar{x}_j) > u_i(l(\bar{\mathbf{x}}))$. But by assumption, $u_i(\bar{y}) \leq u_i(l(\bar{\mathbf{x}}))$, contradicting the fact that $\bar{y} \in U_i(\bar{\mathbf{x}})$. Next, assume that for this i , condition (a) holds for every n , so that $\alpha_i = \bar{x}_i$. If $u_i(h(\bar{\mathbf{x}})) \neq u_i(l(\bar{\mathbf{x}}))$ then the preceding argument can be applied again. Assume therefore that $u_i(h(\bar{\mathbf{x}})) = u_i(l(\bar{\mathbf{x}}))$. By strict concavity, $u_i(l(\bar{\mathbf{x}})) < u_i(y)$, for each $y \in (l(\bar{\mathbf{x}}), h(\bar{\mathbf{x}}))$. Moreover, by assumption $\alpha_i = \bar{x}_i \in (l(\bar{\mathbf{x}}), h(\bar{\mathbf{x}}))$. Therefore, once again, $u_i(\bar{y}) \leq u_i(l(\bar{\mathbf{x}})) < \sum_j w_j u_i(\bar{x}_j)$, contradicting the fact that $\bar{y} \in U_i(\bar{\mathbf{x}})$. \square .

We can now proceed with the proof of the theorem. The concavity of u_i implies that $u_i(E\mathbf{x}) > \sum_j w_j u_i(x_j)$, for every i , so that for every policy vector \mathbf{x} , $\bigcap_{i \in I} U_i(\mathbf{x})$ is nonempty. It follows immediately that $\bigcup_{C \in \mathcal{C}} \bigcap_{i \in C} U_i(\mathbf{x})$ is a convex set.

Fix a particular equilibrium profile \mathbf{s} , and for $t \in \{1, 3, \dots, T-1\}$, let \mathbf{x}_t denote the profile of policy vectors proposed in round t . Note that from Theorem II, player i 's acceptance set in round $t \in \{2, 4, \dots, T-2\}$ must be $U_i(\mathbf{x}_{t+1})$. Thus, in round $t \in \{1, 3, \dots, T-3\}$, the set of policy vectors that will be acceptable to some coalition in round t is given by $\bigcup_{C \in \mathcal{C}} \bigcap_{i \in C} U_i(\mathbf{x}_{t+2})$. Since this set is convex, it follows that if $\alpha_i \in (l(\mathbf{x}_{t+2}), h(\mathbf{x}_{t+2}))$, for some i , then if i proposes α_i , it will be accepted by some coalition. We have established, then, that for each i ,

if $\alpha_i \in (l(x_{t+2}), h(x_{t+2}))$ then $x_{i,t+2} = \alpha_i$.

so that the hypothesis of Lemma IV.1 is satisfied.

Let $\underline{v}(t, \cdot)$ and $\bar{v}(t, \cdot)$ be alternative enumerations of I such that for $1 \leq k < \bar{i}$,⁴⁵ $\underline{\beta}_{\underline{v}(t, k)}(x_t) \leq \underline{\beta}_{\underline{v}(t, k+1)}(x_t)$, while $\bar{\beta}_{\bar{v}(t, k)}(x_t) \geq \bar{\beta}_{\bar{v}(t, k+1)}(x_t)$. Next, define \bar{i} to be the smallest integer strictly larger than $\bar{i}/2$ and define $I_t = \{\underline{v}(t, 1), \dots, \underline{v}(t, \bar{i})\}$ and $\bar{I}_t = \{\bar{v}(t, 1), \dots, \bar{v}(t, \bar{i})\}$. Observe that for each t , $I_t \cap \bar{I}_t \neq \emptyset$ and for each $\tau \neq t$, $I_t \cap \bar{I}_\tau \neq \emptyset$. Set $\underline{\beta}_t = \underline{\beta}_{\underline{v}(t, \bar{i})}(x_t)$ and $\bar{\beta}_t = \bar{\beta}_{\bar{v}(t, \bar{i})}(x_t)$. Thus, a policy vector y is contained in $[\underline{\beta}_t, \bar{\beta}_t]$ if and only if for a strict majority of the players in I , y is weakly preferred to the *least* preferred element of x_t . Specifically, every $y \in [\underline{\beta}_t, Ex_t]$, is weakly preferred to $h(x_t)$ by every $i \in I_t$, while every $y \in [Ex_t, \bar{\beta}_t]$, is weakly preferred to $l(x_t)$ by every $i \in \bar{I}_t$.

We are now ready to proceed with the proof of the theorem. For $t \in \{1, 3, \dots, T-3\}$, it is clearly true that

$$x_t \subset \bigcup_{C \in \mathcal{C}} \bigcap_{i \in C} U_i(x_{t+2}) \quad (\text{B.IV.1})$$

Now, fix $\varepsilon > 0$ and choose $\delta > 0$ for which the conclusion of Lemma IV.1 applies. We will show that if \bar{T} is sufficiently large, then for $T > \bar{T}$, the solution for the T -round game will be contained in an interval of length no greater than ε . Specifically, we have shown (B.IV.1) that for each t , $x_t \subset [\underline{\beta}_{t+2}, \bar{\beta}_{t+2}]$. We will show that when $h(x_{t+2}) - l(x_{t+2})$ exceeds ε , the interval $[\underline{\beta}_t, \bar{\beta}_t]$ will be contained in $[\underline{\beta}_{t+2}, \bar{\beta}_{t+2}]$, but its length will be shorter by at least δ . This fact will establish the theorem.

It follows from the first inclusion of B.IV.1 that for all i , there are at least \bar{i} players j such that $x_{i,t} \in U_j(x_{t+2})$. If $h(x_{t+2}) - l(x_{t+2}) > \varepsilon$, then, Lemma IV.1 implies that $x_{i,t} > \underline{\beta}_{t+2} + \delta$ while $x_{i,t} < \bar{\beta}_{t+2} - \delta$. Summarizing, we have established that for each $t \in \{1, 3, \dots, T-3\}$

$$\text{if } h(x_{t+2}) - l(x_{t+2}) > \varepsilon, \text{ then } x_t \subset [\underline{\beta}_{t+2} + \delta, \bar{\beta}_{t+2} - \delta]. \quad (\text{B.IV.2})$$

The next step in the proof is to show that for each $t \in \{1, 3, \dots, T-1\}$,

$$\text{either } \underline{\beta}_t = l(x_t) \text{ or } \bar{\beta}_t = h(x_t). \quad (\text{B.IV.3})$$

To see this, observe that for each i , $\underline{\beta}_i(x_t) \leq l(x_t)$ while $\bar{\beta}_i(x_t) \geq h(x_t)$. Moreover, because payoffs are concave, at most one of these inequalities can be strict for any i . Thus if $\underline{\beta}_i(x_t) \leq \underline{\beta}_t < l(x_t)$, for each $i \in I_t$, then $\bar{\beta}_i(x_t) = \bar{\beta}_t = h(x_t)$, for $i \in I_t \cap \bar{I}_t$. Since $I_t \cap \bar{I}_t$ is nonempty, this establishes that (B.IV.3) is true. We will now assume (without loss of generality) that $\bar{\beta}_t = h(x_t)$, and rewrite B.IV.2 as

$$\bar{\beta}_t \subset [\underline{\beta}_{t+2} + \delta, \bar{\beta}_{t+2} - \delta]. \quad (\text{B.IV.2'})$$

To complete the proof of the theorem, we need to show that $\underline{\beta}_t \geq \underline{\beta}_{t+2}$. To see this, observe first that for each $i \in I_t$,

$$u_i(\underline{\beta}_t) \geq u_i(\underline{\beta}_t(x_t)) \geq u_i(\bar{\beta}_t) > u_i(\bar{\beta}_{t+2}). \quad (\text{B.IV.4})$$

The second inequality holds because $\bar{\beta}_t = h(x_t)$; the third because $\bar{\beta}_{t+2} > \bar{\beta}_t$. We now have two cases to consider. First assume that $\bar{\beta}_{t+2} = h(x_{t+2})$. In this case, B.IV.4 implies that $\underline{\beta}_t \in G_i(x_{t+2})$, for $i \in I_{t+2}$, so that, immediately,

⁴⁵ Recall that the set of players I has \bar{i} elements.

$\beta_i \geq \beta_{i+2}$. Second assume that $\beta_{i+2} \in I(x_{i+2})$. In this case, $u_i(\bar{\beta}_{i+2}) \geq u_i(\bar{\beta}_i(x_{i+2})) \geq u_i(\beta_{i+2})$ for $i \in I_{i+2}$. But from B.IV.4, $u_i(\beta_i) > u_i(\bar{\beta}_{i+2})$, for each $i \in I_j$. Since $I_j \cap I_{i+2}$ is nonempty, there exists i such that

$$u_i(\beta_i) > u_i(\bar{\beta}_{i+2}) \geq u_i(\beta_{i+2}).$$

Since u_i is concave and $\bar{\beta}_{i+2} > \beta_{i+2}$, it now follows that $\beta_i > \beta_{i+2}$. \square

Proof of Theorem V: The proof uses the following lemma repeatedly.

Lemma V.1: Fix $\epsilon > 0$, an integer k , and a strictly positive probability vector $p \in \Delta^{k-1}$. There exists $\delta > 0$ such that for each i and $y = (y_{\nu})_{\nu=1}^k \subset X$, $\text{diam}(y) \geq \epsilon$ implies $u_i(p \cdot y) - \sum_{\nu} p_{\nu} u_i(y_{\nu}) \geq \delta$.

Proof of Lemma V.1: If the Lemma were false, then we could find $\epsilon > 0$, $i \in I$ and a sequence of vectors, $\{y^n\}$ in X , such that for each n , $\text{diam}(y^n) \geq \epsilon$ and $u_i(p \cdot y^n) - \sum_{\nu} p_{\nu} u_i(y_{\nu}^n) < n^{-1}$. Since X is compact the sequence $\{y^n\}$ has a convergent subsequence. Let \bar{y} be the limit of this subsequence. Since u_i is continuous, $u_i(p \cdot \bar{y}) \leq \sum_{\nu} p_{\nu} u_i(\bar{y}_{\nu})$. Moreover, the diameter of \bar{y} is at least ϵ . However, since the vector p is strictly positive, $p \cdot \bar{y}$ is contained in the relative interior of the convex hull of \bar{y} . But this contradicts the assumption that u_i is strictly concave. \square

We now proceed with the proof of the theorem. Let (x^T) denote the sequence of outcomes corresponding to a nested sequence of equilibrium strategy profiles for the T -round games. Assume that player #1 is an essential player. For each T , let $\theta^T = Eu_1(x^T)$.

Step #1: The sequence (θ^T) is a strictly increasing, Cauchy sequence.

Proof of Step #1: Fix an even integer T . Since player #1 is essential, each player's policy proposal in round #1 of the $T+2$ -round game must yield player #1 a payoff of at least θ^T . Moreover, from Lemma I(a) player #1's own proposal yields a payoff strictly exceeding θ^T . This establishes that the sequence is strictly increasing. Because u_i is continuous and X is compact, $u_i(\cdot)$ is bounded on X . Hence the sequence is Cauchy.

Step #2: For all positive ϵ , there exists a \bar{T} such that for each $T > \bar{T}$, $\text{diam}(x^T) < \epsilon$.

Proof of Step #2: Suppose to the contrary that there exists a subsequence, (x^n) , of (x^T) such that for each n , $\text{diam}(x^n) \geq \epsilon$. From Lemma V.1, there exists $\delta > 0$ such that for each n , and each i $u_i(w \cdot x^n) - \sum_{j \in I} w_j u_i(x_j^n) \geq \delta$. It follows that for each n , player #1's own proposal in round #1 of the $n+2$ -round game must yield a payoff that exceeds θ^n by at least δ . Thus, for each n , $\theta^{n+2} \geq \theta^n + w_1 \delta$. But this contradicts Step #1.

Step #3: The limit of any convergent subsequence of (x^T) is a singleton profile (\bar{y}) such that $u_1(\bar{y}) = \bar{\theta} = \lim_T \theta^T$. Moreover, a convergent subsequence exists.

Proof of Step #3: The first statement follows immediately from Steps #1 and #2. The second follows from the fact that X is compact.

Step #4: If $\{y\}$ is the limit of a convergent subsequence of $\{x^T\}$, then y belongs to the core of the underlying game. Moreover, there are at most \bar{i} distinct limits of convergent subsequences.

Proof of Step #4: The first sentence follows from an argument identical to the proof of Theorem II. Assume that there are \bar{k} distinct limits of convergent subsequences, $\{y^1, \dots, y^{\bar{k}}\}$. From Step #3, $u_1(y^\kappa) = \bar{\theta}$, for each κ , so that for any $k \neq \kappa$, $\frac{1}{2}y^\kappa + \frac{1}{2}y^k$ yields player #1 a strictly higher payoff than either. Moreover, for each κ , since y^κ belongs to the core, it cannot be Pareto dominated; thus, there must exist $i(\kappa) > 1$ such that $U_{i(\kappa)}(y^\kappa) \cap U_1(y^\kappa)$ has an empty interior. Suppose that $i(\kappa) = i(k) = i$, for $\kappa \neq k$. Since u_i and u_1 are both strictly concave, then $\frac{1}{2}y^\kappa + \frac{1}{2}y^k$ must yield player i a higher payoff than *either* y^κ or y^k . But this means that either $U_i(y^\kappa) \cap U_1(y^\kappa)$ or $U_i(y^k) \cap U_1(y^k)$ has a nonempty interior. We have established, then, that $\kappa \neq k$ implies $i(\kappa) \neq i(k)$ and hence that $\bar{k} \leq \bar{i}$.

Step #5: For every $\epsilon > 0$ there exists \bar{T} such that for $T > \bar{T}$, $\text{diam}(x^{T+2} \cup x^T) < \epsilon$.

Proof of Step #5: Suppose to the contrary that there exists $\epsilon > 0$ and a subsequence $\{x^{T^n}\}_{n=1}^\infty$, such that for each n , $\text{diam}(x^{T^n+2} \cup x^{T^n}) > 3\epsilon$. From Step #2, we can pick \bar{n} sufficiently large that for $T > T^{\bar{n}}$, the diameter of x^T is less than ϵ . Clearly, for such T , the distance between any point in the convex hull of x^T and any point in the convex hull of x^{T+2} must be at least ϵ . Pick $\bar{\delta} > 0$ such that the conclusion of Lemma V.1 holds for this ϵ , with $k = 2$, $p = (\frac{1}{2}, \frac{1}{2})$ and $\delta = 3\bar{\delta}$; Thus, for $T > T^{\bar{n}}$, we have for each player i ,

$$u_i(\frac{1}{2}w \cdot (x^T + x^{T+2})) - \frac{1}{2} \sum_j w_j (u_i(x_j^{T^n} + u_i(x_j^{T^n+2})) \geq 3\bar{\delta}. \quad (\text{B.V.1})$$

Next, using Step 1 and, once again, Step 2, pick $n > \bar{n}$ sufficiently large that for each i $\theta^{T^n+2} - \theta^{T^n} < \delta/2$ and for each $T \geq T^n$, $\text{diam}(u_i(x^T)) < \bar{\delta}$. Let $\bar{x} = \frac{1}{2}w \cdot (x^{T^n} + x^{T^n+2})$. Now, in the first round of the (T^n+2) -round game, player #1 proposes $x_1^{T^n+2}$ to some coalition C . We claim that \bar{x} will be accepted by each player in C and is strictly preferred by player #1 to $x_1^{T^n+2}$. This contradicts the hypothesis that $x_1^{T^n+2}$ is player #1's best alternative at this point of the game, and hence establishes Step #5. For $i \in C$, we have

$$\sum_j w_j u_i(x_j^{T^n+2}) \geq u_i(x_1^{T^n+2}) - \delta \geq \sum_j w_j u_i(x_j^{T^n}) - \delta \quad (\text{B.V.2})$$

The first inequality follows from our choice of T^n ; the second uses the condition for acceptance by i of $x_1^{T^n+2}$. Combining (B.V.1) and (B.V.2) yields $u_i(\bar{x}) > \sum_j w_j u_i(x_j^{T^n}) + 2\delta$, for each $i \in C$. On the other hand for player #1, we have

$$u_1(x_1^{T^n+2}) \leq \sum_j w_j u_1(x_j^{T^n+2}) + \delta \equiv \theta^{T^n+2} + \delta \leq \theta^{T^n} + 2\delta \equiv \sum_j w_j u_1(x_j^{T^n}) + 2\delta \quad (\text{B.V.3})$$

Both inequalities follow from our choice of T^n . Combining (B.V.1) and (B.V.3) yields $u_i(\bar{x}) > u_1(x_1^{T^n+2}) + \delta$ which establishes the claim above.

Step #6: The sequence (x^T) has a (unique) limit point.

Proof of Step #6: Let Y denote the intersection of $u_1^{-1}(\bar{\theta})$ and the core. From Step #4, Y is a finite set. If Y is a singleton set, then Step #6 follows immediately. Assume, therefore, that Y contains two distinct elements and choose $\epsilon > 0$ such that any two elements of Y are separated by at least 3ϵ . From Step #4, we can pick \bar{T} such that for every $T \geq \bar{T}$, $x^T \in B(Y, \epsilon)$. Thus, there is a unique policy $\bar{y} \in Y$ such that $x^{\bar{T}} \in B(\bar{y}, \epsilon)$. Moreover, from Step #5, there exists $\bar{\bar{T}} > \bar{T}$ such that for every $T > \bar{\bar{T}}$, $x^{T+2} \subset B(x^T, \epsilon)$. It now follows from the two previous sentences that for every $T > \bar{\bar{T}}$, $x^T \in B(\bar{y}, \epsilon)$. This establishes Step #6 and completes the proof of the Theorem. \square