Econometrica, Vol. 62, No. 4 (July, 1994), 795–817

A NONCOOPERATIVE VIEW OF COALITION FORMATION AND THE CORE

BY MOTTY PERRY AND PHILIP J. RENY[1]

A noncooperative implementation of the core is provided for games with transferable utility. The implementation obtained here is meant to reflect the standard motivation for the core as closely as possible. In the model proposed, time is continuous. This idealized treatment of time ensures that there is always time to reject a noncore proposal before it is consumated.

KEYWORDS: Core, implementation, continuous time.

1. INTRODUCTION

THE NOTION OF THE CORE has attracted the interest of economists for more than 100 years. Much of the core's appeal stems from the intuitive and natural story behind it, the story that first motivated F. Y. Edgeworth in 1881. It runs as follows. An allocation which is not the core is regarded as unstable. It is always the case, given such an allocation, that some individuals will be able to form a coalition and obtain an allocation that each strictly prefers.

Thus, the primary motivation for the core is noncooperative in nature. Nonetheless, the core is not a noncooperative solution concept. This is because, in particular, the possibilities for forming coalitions, and making offers and counteroffers, are not explicitly modeled.

In this work, we provide a noncooperative implementation of the core. However, we do not merely implement the core. The nature of the game form employed is designed to reflect the motivating story as accurately as possible. The present results thus provide formal content to the usual intuitive justification for the core. In our view, the core would lose much of its appeal were it not possible to provide such a noncooperative foundation.

The situations covered by our analysis have a payoff structure summarized by a transferable utility game in characteristic function form, which for ease of exposition only, is assumed to be superadditive. This is embedded in a dynamic game in continuous time in which each player bargains over with whom to form

[1] We wish to thank Murali Agastya, Kalyan Chatterjee, Ig Horstmann, Arthur Robson, and Bob Rosenthal for helpful discussions. We are also grateful to three anonymous referees and the editor for numerous comments and suggestions which, in one instance, served to correct an error in the statement of Theorem 3.

The first author gratefully acknowledges support from the Ford Foundation. This work was carried out in part while visiting the Department of Economics at The University of Western Ontario.

The second author gratefully acknowledges support from the Social Sciences and Humanities Research Council of Canada. This work was carried out in part while visiting the Center for Rationality and Interactive Decision Theory and the Department of Economics at the Hebrew University of Jerusalem and the Department of Economics at the University of Wisconsin.

a coalition and how to allocate the coalition's value among its members. The bargaining model employed allows flexibility in the negotiation process. In particular, there is no fixed timing or order of moves. At any positive time, any player can place a proposal on the table. A proposal consists of a coalition together with a suggested allocation of that coalition's value among its members. A coalition is formed if all its members accept the proposal before a new proposal is placed on the table. Once a new proposal is put forward, it automatically rescinds the previous one, so that only one proposal is on the table at any time. After a coalition, S, is formed, any member can choose to leave and consume. In this event, all other members of S must also leave and consume their part of the agreed upon allocation.[2] The remaining players, $N \setminus S$, continue bargaining. Thus, there is no presumption here that the grand coalition must eventually form. Both the ultimate coalition structure and division of the surplus are determined endogenously.

Continuous time is modeled along the lines of Perry and Reny (1993). (For other continuous-time models see Bergin and Macleod (1989), Simon and Stinchcombe (1989), and Stinchcombe (1988).) A key feature of the model is that players can time their proposal so that each member of the coalition to whom it is directed has an opportunity to react. Subgame perfection and stationarity are the other main ingredients. The role of stationarity is its familiar one; it precludes the use of those (nonstationary) strategies that are needed to obtain folk-theorem-like results (see, for instance, the example due to Shaked in Osborne and Rubinstein (1990, Chapter 3)). One senses however, that its connection to the core may run slightly deeper. For implicit in the story justifying the core is the presumption that it is only the blocking proposal itself which matters and not the player who proposed it. Moreover, if rejected, it is understood that the status quo will prevail. Stationarity is one way to formalize these implicit restrictions.

Within this framework it is shown that for (v, N) a superadditive TU game:

(i) If an allocation is not the core of (v, N), then it cannot be supported as a stationary subgame perfect equilibrium (SSPE) outcome.

(ii) If (v, N) is totally balanced (in particular, if it is a market game) then:

(a) an allocation can be supported as an SSPE outcome *if and only if* it is in the core of (v, N), and

(b) if payoffs are discounted over time, then although no SSPE exists, the limit, as $\varepsilon \to 0$, of the set of allocations that can be supported as stationary ε-subgame perfect equilibria coincides with the core of (v, N).

The proof of (i) turns out to be virtually no more than a restatement of the standard motivation for the core. In our view, this is an important feature of the present model. For it serves to place the familiar and intuitive justification for the core on a solid noncooperative foundation. Part (ii)(a) is the implementation

[2] Although this particular convention suits our purpose, so would many others. Any such convention would suffice so long as each member of a formed coalition can guarantee himself at least the payoff associated with the agreed upon allocation.

result, while part (ii)(b) is an expression of the robustness of this result to the introduction of time preference.

An obvious implication of (i) above is that the underlying game, (v, N), must have a nonempty core in order that an SSPE exist. Furthermore, the demands of subgame perfection require that an equilibrium exist in every subgame; in particular, those in which only a subset of players remains. Thus a stronger condition is needed for the existence of an SSPE, namely that the game (v, N) be totally balanced. The total balancedness condition could be dispensed with if, for instance, the game ended once a single coalition had formed. In our view, however, it is more natural to suppose that remaining players (indeed, even members of coalitions that have formed) can if they wish continue to negotiate.

When discounting is present, players willing to accept a proposal wish to do so as soon as possible. Given the continuous time setting, this is ill-defined. No subgame perfect equilibrium (stationary or otherwise) then exists. However, for every positive ε, a stationary ε-subgame perfect equilibrium does exist, and the set of such equilibrium outcomes, for ε small, approximates the core.

Related Literature

There is little doubt that Edgeworth's conjecture, which ultimately produced the core equivalence theorem, firmly and permanently established that the notion of the core has deep connections to economics. Although Edgeworth's conjecture appeared as early as 1881, it was not until Shubik (1959) and Debreu and Scarf (1963) that a general and completely rigorous treatment was provided. The next major development was due to Aumann (1964) and Vind (1964) both of whose introduction of a continuum of consumers serves to provide a remarkably sharp result. Refinements of these ideas and simplifications have continued to appear. A partial list includes Arrow and Hahn (1971), Bewley (1973), Hildenbrand (1974), and Anderson (1978). To a large extent, the relevance of the present paper to economics owes a substantial debt to these earlier contributions.

Implementation theory had its beginnings in the pioneering work of Hurwicz (1972). High on the agenda of research in this field was the search for mechanisms capable of implementing outcomes possessing some kind of efficiency property (from Pareto efficiency to that expressed by the core). Maskin's (1977) seminal contribution provides a general means for implementing a wide class of social choice correspondences. In particular, Maskin's result shows how to (Nash) implement the core. Consequently, as previously mentioned, the purpose of the present paper is not to merely implement the core. We also insist that the implementing mechanism display a natural connection to the ideas of coalition formation and bargaining.

There are by now a number of papers in which a noncooperative game is designed to yield a cooperative solution. See, for instance, Harsanyi (1974), Selten (1981), Binmore (1985), Gul (1989), Hart and Mas-Colell (1992), and Lagunoff (1992). It has become apparent that the order in which players act is important, especially in obtaining efficient (and a fortiori, core) outcomes. Two papers which focus attention on this relationship are Chatterjee et al. (1990) and Moldovanu and Winter (1991). Chatterjee et. al.'s is a model of n-person bargaining with discounting, in which the underlying opportunities are given by a TU game. Among other things, they show that if the TU game is strictly convex, then there is at least one way to order the players so that stationary equilibria are efficient. Furthermore, as the discount factor tends to one, such outcomes converge to core outcomes. Moldovanu and Winter reanalyze the game proposed by Selten (1981). Although the set of stationary subgame perfect equilibria in Selten's game depends on the assumed order of moves, they show that an allocation is in the core (of the underlying NTU game) if and only if for *every possible* order of moves it is an outcome of a stationary subgame perfect equilibrium.

Note that in these last two models, the core is "implemented" (only partially in Chatterjee et. al.) via a *class* of games, one for each order of moves, rather than a single game. In our continuous-time model, the order of moves is not fixed a priori, but is determined in equilibrium. Moreover, every possible order of moves is feasible from the outset. This squares nicely with the results of both Chatterjee et al. and Moldovanu and Winter.

The paper by Kalai, Postlewaite, and Roberts (1979) is closest to ours in spirit. Theirs is a static game in which, simultaneously, each agent announces a coalition to which he wishes to belong and a feasible allocation for that coalition. Any coalition whose members' announcements match, obtains for its members the announced allocation. Others receive their own value. Kalai et. al. show that the set of strong Nash equilibria coincides with the core. However, because strong Nash equilibrium is motivated by the *assumption* that profitable coalitions will form, it cannot be the basis of a complete noncooperative foundation for the core.

In the next section we provide a heuristic description of the model, and show, by means of a few examples, its main features, the roles of the assumptions, and the ideas behind the proofs. In Section 3 a formal description of the model is given. Section 4 contains the main theorems. Time preference is introduced in Section 5. Section 6 provides an explanation of our decision to employ continuous time.

2. OVERVIEW OF THE MODEL AND EXAMPLES

Consider a situation in which the underlying opportunities can be summarized by a transferable utility game in characteristic function form, (v, N), where $N = (1, 2, \ldots, n)$ is the (finite) set of players and v: $2^N \to \mathbb{R}_+$ associates with every subset, S, of N (or simply, every *coalition*) its value $v(S)$. We shall also assume throughout, for simplicity only, that (v, N) is *superadditive*. That is, for every pair of disjoint coalitions S and T, $v(S) + v(T) \leqslant v(S \cup T)$. Recall then that $x \in \mathbb{R}^n$ is an element of the *core* of (v, N) if for every $S \subseteq N$, $\sum_{i \in S} x_i \geqslant v(S)$, with equality for $S = N$.

The following extensive form describes how coalitions form and associated allocations are achieved. Time is continuous. The game starts at $t = 0$ at which time a player can choose to make a proposal or to be quiet. At every positive time t, a player can choose to make a proposal, accept the current proposal, be quiet, or leave and consume. Note that the choice to leave and consume is formally distinct from the choice to remain quiet. It might then be helpful to think of the choice to leave and consume as being accompanied by an announcement to this effect. We now elaborate further on these possibilities.

A proposal (x, S) by player $i \in N$, consists of a coalition S (where i need not be a member of S), and an allocation $(x_j)_{j \in S}$, the sum of whose components does not exceed $v(S)$. If player i makes a proposal (x, S) at time t, then this proposal is effective so long as no new proposal is made. Thus, only one proposal is dealt with at a time, and once a new proposal is made the previous one is no longer in effect.

A proposal becomes *binding*, when it is accepted by all members of the relevant coalition while it is effective. In particular, the player who made the proposal must, if a member of the relevant coalition, subsequently accept it before it becomes binding. This enables proposals to be made anonymously, the significance of which will become clear in due course. Once a player accepts a proposal, he must remain quiet (i.e., he can neither make a new proposal nor leave to consume) until that proposal becomes binding, or until it is replaced by a new proposal. This restriction serves a purely technical role and is discussed at length in Example 1b below.

Once a proposal, (x, S), becomes binding, we say that *the coalition S has formed and has accepted* (x, S). Any player, i, in S can then choose to continue bargaining (looking forward to a better proposal), or to leave and consume. If he chooses the latter, he obtains the payoff x_i, according to the binding proposal (x, S). When this occurs, every other player, j say, in S must also leave and consume, thereby obtaining a payoff of x_j. Thus, each member, i, of S can ensure himself a payoff of x_i once (x, S) becomes binding. Indeed, any convention with this property—the essence of a *binding* agreement— would suffice for our purposes. So long as no member of S leaves, each can make and accept proposals.

At any time during the course of play then, there may be a number of binding proposals (y^1, S^1), $(y^2, S^2), \ldots, (y^m, S^m)$, say among players who have not left. In order to submit a proposal, (x, S), in the presence of such binding proposals, it is assumed that, for each j, either S and S^j are disjoint or S contains S^j. That is, players are not permitted to make proposals to nonempty strict subsets of coalitions that have already formed. This convention is adopted as a simple means to requiring that *all* coalition members associated with a currently effective binding proposal approve the annulment of their proposal before any new binding proposal supersedes it. Furthermore, we insist that no player be associated with more than one binding proposal at any given time. Thus, if (y^1, S^1) is a binding proposal and (y^2, S^2) is proposed, where S^2 contains S^1, then upon the acceptance of (y^2, S^2) by every member of S^2, the binding proposal, (y^1, S^1), is considered annulled and is stricken from the list of binding proposals, while the proposal (y^2, S^2) is added to the list of binding proposals. That is, binding proposals are annulled if every coalition member agrees to this. Consequently, at all times, the formed coalitions associated with the list of binding proposals are mutually disjoint. Finally, if player i never leaves to consume, his payoff is $d_i \leqslant v(i)$.

It is assumed that for all times t, and after any history up to time t, players are not allowed to choose to make a proposal, accept a proposal, or leave to consume either "just before" or "just after" time t. That is, for every time t, for every history up to t and for every vector of actions available to the players at time t according to the history, there is an $\varepsilon > 0$, such that each player is quiet, according to his strategy, in the open intervals $(t-\varepsilon, t)$ and $(t, t+\varepsilon)$. As is well known, without some such assumption, the continuous-time game is not well-defined. This particular assumption, however, is somewhat more than is needed for that purpose. Indeed, it provides not only the technical service of ensuring well-definedness, it plays a central role in obtaining the results. Note that the ε referred to in the restriction *is not bounded away from zero*. Therefore, this restriction does not preclude players from reacting arbitrarily quickly. For instance, the strategy to accept, at time $t > 0$, any proposal made at time zero that provides you with precisely $t > 0$ units of surplus, and otherwise to remain quiet, remains feasible. Nonetheless, because your reaction time according to this strategy depends upon the proposal made at time zero, it is not bounded away from zero. On the other hand, the strategy to accept, at any time $t > 0$, *any* proposal made at time zero is ruled out by the restriction. The difference between the two is that if an offer is made at time zero, then according to the second strategy, whenever you do happen to accept the proposal, you've accepted it too late. No such difficulty arises under the first strategy.

The solution concept applied to this game is stationary subgame perfect equilibrium (SSPE). That is, stationary strategies constituting a subgame perfect equilibrium (among all strategies, not merely stationary ones). A strategy is stationary if, roughly, after any history, it depends only upon the set of players remaining, the current proposal, the set of players who have accepted it, and the set of binding proposals.[3] Before presenting the formal details of the model, it may be helpful to consider a few examples which illustrate the main ideas.

Example 1: There are two buyers and one seller. The seller (player 1) owns one unit of a divisible good that is of no value to him. The two buyers (players 2 and 3) each value the good at one dollar. A TU game describing this situation is given by:

$$N = \{1,2,3\}, \quad v(i) = v(2,3) = 0 \quad \text{for} \quad i \in N, \quad \text{and}$$
$$v(1,2) = v(1,3) = v(N) = 1.$$

In this game the core consists of a single point, namely $x^* = (1,0,0)$. This

[3] We also allow stationary strategies to depend, in a limited sense, upon time. Indeed, some time dependence is necessary. This is discussed further below.

corresponds to the competitive outcome as the excess of buyers over sellers results in the seller obtaining the entire surplus. We now show (informally) that only the division x^* can be supported as a stationary subgame perfect equilibrium outcome.

Assume, by way of contradiction, that (x_1, x_2, x_3) is a stationary subgame perfect equilibrium outcome with $x_1 < 1$, and assume without loss that $x_2 \leqslant x_3$. Now, by assumption, given the actions of the players at time zero, there is an $\varepsilon > 0$ such that all players are quiet in the interval $(0, \varepsilon)$ according to their strategies. In particular, note that all players remain active up to time ε since no one can leave at time zero.[4] Consider the following deviation by player 1: At time $t_1 = \varepsilon/2$ propose $(x_1 + (1 - x_1 - x_2)/2,\ x_2 + (1 - x_1 - x_2)/2, \{1, 2\})$ and at time $t_2 > t_1$ accept this proposal, where t_2 is such that all players, after 1's deviation and according to their strategies, are quiet in the interval $(t_1, t_2]$. Again, by assumption, such a time t_2 is guaranteed to exist. Consider now the continuation after t_2. Let t_3 be the first time after t_2 at which some player is not quiet according to his strategy. Player 2 can guarantee himself a payoff of $x_2 + (1 - x_1 - x_2)/2$ by accepting the above proposal at any time $t \in (t_2, t_3)$. If, according to his strategy, he is to accept this proposal, then the deviation is successful. But this contradicts x being an equilibrium. So, it must be that player 2, according to his strategy, does not accept 1's proposal. It follows that, according to the equilibrium strategies, at time t_3 someone makes a proposal, say (y, T), and player 2 subsequently obtains a payoff of at least $x_2 + (1 - x_1 - x_2)/2$. Note then that by stationarity, *whenever* (a) all players remain, (b) the current proposal is (y, T), (c) no one has accepted it, and (d) there are no binding proposals, the subsequent outcome, according to the equilibrium strategies, gives player 2 at least $x_2 + (1 - x_1 - x_2)/2$. But this contradicts (x_1, x_2, x_3) being the equilibrium outcome. For player 2 can deviate at any time $t \in (0, \varepsilon)$ by proposing (y, T), thereby rendering (a)–(d) all true and ensuring himself (by stationarity) a payoff exceeding x_2. Note here the importance of the anonymity with which proposals are made. According to stationarity, after any history in which a proposal is made, the continuation must be independent of the player who made it. This is central to player 2's ability to mimic the relevant features of the history satisfying (a)–(d) by proposing (y, T).

The above argument shows that the only candidate for an equilibrium outcome is $x^* = (1, 0, 0)$. We now present (again, informally) a profile of stationary subgame perfect equilibrium strategies yielding the outcome x^*.

If the current time is not an integer, then all players are quiet. Suppose now that the current time is an integer, and player 1 does not belong to any coalition that has formed. If the current proposal is (x^*, N), then each player who has not yet accepted it, does so. If all players have accepted it, then all leave. If the current proposal is not (x^*, N) and player 1 has not accepted it, then player 1

[4] In order to leave, a player must first accept a current proposal. This is so even if a player wishes to consume his own value. In order to consume his own value, player i must first propose $(v(i), \{i\})$, then accept it, and finally leave and consume. Consequently, no player can leave at time zero.

proposes (x^*, N) and the others are quiet. If player 1 has accepted the current proposal, then the others accept it as well.

It remains to describe players' strategies after histories in which player 1 belongs to a coalition that has formed. So, suppose that a coalition, S, containing player 1, has formed and has accepted (y, S). Let $d = 1 - \sum_{i \in S} y_i$ be the amount of surplus "wasted" by the binding proposal (y, S). To obtain each players' action, consult once more the previous paragraph, while ignoring the second half of the second sentence, and replace x^* throughout with $y^d = (y_1 + d, y_2, y_3)$, where $y_i = 0$ if $i \notin S$. In addition, replace the last sentence with: If player 1 has accepted the current proposal, then player 2 (3) either accepts it or proposes (y^d, N) according to whether the current proposal gives him more than y_2 (y_3) or not. These strategies constitute a subgame perfect equilibrium, are stationary, and support the outcome x^*.

The one seller/two buyers example above is sometimes used to discredit the concept of the core. It is, the argument goes, inconceivable that the two buyers will end up with no surplus. They can, for example, tacitly bargain with the seller as a coalition, and this might yield both a strictly positive payoff. Note, however, that for any such payoffs the buyers might secure, there is an offer that the seller both prefers and can make to one of them which, from that one's point of view, is strictly better. And the power of stationarity guarantees that this offer must be accepted.

The example above provides the essentials of the general argument establishing that only core allocations can be supported in equilibrium. Taking this as given, it follows that a necessary condition for the existence of an SSPE is that the core of the TU game be nonempty. But this is not sufficient. Subgame perfection also disciplines behavior subsequent to histories in which *any* subset of players remains. A necessary condition for the existence of an SSPE in such a subgame is that the TU game restricted to this subset of players also have a nonempty core. It thus follows that a necessary condition for the existence of an SSPE is that the game be totally balanced.

Is total balancedness a sufficient condition for existence? What about histories after which some binding proposals are present and all players remain, say? After such a history the underlying game (v, N), is no longer relevant since some player, i say, may be able to guarantee himself more than $v(i)$ by leaving to consume his share of a binding proposal. Indeed, the now relevant TU game differs from (v, N), in that members of formed coalitions can guarantee themselves a payoff equal to their share of the corresponding binding proposal. Now, in order for an SSPE to exist, any candidate continuation equilibrium allocation, x, must be immune to any feasible improvement. Since strict subsets of formed coalitions cannot form, this amounts to x being unblockable by any coalition that is not a strict subset of a formed coalition. When is such an allocation guaranteed to exist? It turns out that for TU games, the assumption of total balancedness is enough to ensure its existence, regardless of the history.[5] The

[5] Indeed, this is the only reason we restrict attention to TU games.

next example serves to illustrate the importance of keeping track, given a history, of the relevant TU game, and how an appropriate continuation equilibrium outcome is constructed.

EXAMPLE 1A: Consider the following attempt at constructing an SSPE for the one seller/two buyers game of Example 1. For every subset of players, S, let $x(S)$ be an element of the core of the game (v, N) restricted to S, and put $x(N) = x^*$. This is possible since (v, N) is totally balanced. For any history, if the set of remaining players, according to that history, is S, then construct the continuation strategies so that the subsequent outcome is $x(S)$.

The above procedure does not produce an SSPE as we now show. Take a history according to which all players remain, and $((.4, .4), \{1, 2\})$ is the only binding proposal. According to the procedure outlined above the subsequent outcome will be $x(N) = x^* = (1, 0, 0)$. But this gives player 2 a continuation payoff of zero, while player 2, by leaving, can obtain a payoff of .4. Clearly, any equilibrium continuation must give player 2 a payoff of at least .4, the amount he can unilaterally guarantee himself by virtue of the existing binding proposal.

The correct procedure for constructing an equilibrium continuation involves strategies that permit the players to agree on an outcome that is unblockable, in the context of the relevant TU game, by coalitions that can feasibly form. It is not difficult to see that, after the history given above, any outcome (x_1, x_2, x_3) with $x_1 + x_2 = 1$ and $x_1, x_2 \geqslant .4$ will suffice.

In general, there may be many binding proposals present after some history, and the continuation outcome may require the joining of two or more previously formed coalitions. Of course, for superadditive games there is never any loss in joining previously formed coalitions. Likewise, there are instances in which some previously formed coalitions *must* be joined in order to achieve a core allocation. Included in the proof of Theorem 2 is a general construction of an appropriate equilibrium continuation outcome, after any history.

The assumption that, upon accepting a proposal, a player must be quiet until either the proposal is replaced by another or until it becomes binding, is somewhat restrictive and so merits some discussion. The example to follow shows that this assumption is needed to ensure the existence of an SSPE. Moreover, it indicates that this assumption plays only a technical, rather than a more substantive role in the analysis.

EXAMPLE 1B: Consider again the one seller/two buyers game of Example 1 and contrary to our assumption allow players, upon accepting a proposal, to make a new proposal if they wish. It is easy to see that this additional flexibility does not affect at all the argument that points other than $x^* = (1, 0, 0)$ cannot be supported as an SSPE allocation. Thus, as before, x^* is the only candidate for an equilibrium allocation. Consider, however, a history according to which, (a) all players remain, (b) there are no binding proposals, (c) the current proposal is $x = (0, .5, .5)$ to the grand coalition, and (d) players 1 and 2 have accepted it. Because x^* is the only candidate for an equilibrium outcome, it follows (by stationarity) that if player 1 now proposes x^* to the grand coalition, his payoff will be 1 and players 2 and 3 will obtain payoffs of zero. However, if player 3 accepts the current proposal, rendering it binding, his payoff will be .5 in the continuation, and player 1's will be zero. Hence, player 1 wishes to replace the current proposal before player 3 accepts it, and player 3 wishes to accept it before player 1 replaces it. With time continuous, there can be no equilibrium in this race to be the first to move after a given time.

To recap, the assumption under discussion is not needed to guarantee that noncore points cannot be supported as equilibrium outcomes. It plays a role only in ensuring that an SSPE supporting a point in the core can be constructed and then only to ensure that irrational (out-of-equilibrium) proposals do not produce histories having no equilibrium continuation.

In the present model, as distinct from most of its predecessors, coalitions are not forced to leave and consume once they form (Gul (1989) is also an exception). Whether to stay and continue negotiating or to leave and consume is endogenously determined by the players' strategies. While this seems to be a natural assumption, it is also necessary, in the present continuous-time context, to ensure the existence of an equilibrium. The final example is designed to illustrate this.

Example 2: Add both an additional buyer and seller to Example 1. Let players 1 and 2 be the sellers and 3, 4, 5 the buyers. Again, each seller has a single divisible good of no value to him, and each buyer values each good alone at one dollar and both goods together at one dollar as well. The unique core point is $x^* = (1,1,0,0,0)$.

Consider now a history according to which there are no binding proposals and the current proposal involves the allocation $x = (0, 1/3, 1/3, 1/3)$ to the coalition $\{2,3,4,5\}$. Suppose also that, according to this history, players 2, 3, and 4, have accepted this proposal. Now, if player 5 also accepts it, rendering it binding, his payoff in the continuation is at least $1/3$. If, however, this proposal is replaced by another before the coalition $\{2,3,4,5\}$ forms, then in the continuation, the resulting division must (by stationarity) be $(1,1,0,0,0)$, the unique core point. It follows that player 5 must accept the proposal and that the coalition $\{2,3,4,5\}$ must form. However, if when the coalition $\{2,3,4,5\}$ forms, each member, contrary to our assumption, must leave and consume, player 1, now alone, will obtain a payoff of zero. But this cannot be an equilibrium, since player 1 can, before player 5 accepts the current proposal, replace it by another, thereby inducing the unique core point in the continuation and obtaining a payoff of 1. So, player 1 wishes to replace the proposal before player 5 accepts it, and player 5 wishes to accept it before player 1 replaces it. Hence, under the assumption that coalitions, once formed, must leave, there is no equilibrium in this game.

On the other hand, if, as we assume, coalitions are not forced to leave, then one can support the following continuation. At the next integer time, time 1 say, player 5 accepts the proposal and the coalition $\{2,3,4,5\}$ forms, but does not leave. At time 2, player 1 proposes the allocation $(1, 0, 1/3, 1/3, 1/3)$ to the grand coalition. At time 3 all players accept this proposal, and at time 4 all leave and consume.

3. THE MODEL

Histories: Let q stand for the choice to be quiet, l for the choice to leave, and a for the choice to accept the current proposal. Let

$$P = \left\{ (x, S) | S \subseteq N,\ x_i \geqslant v(i) \text{ for all } i \in S, \text{ and } \sum_{i \in S} x_i \leqslant v(S) \right\}$$

denote the set of *feasible* proposals.

A *history*, h, up to time t is an n-tuple of functions: $h = (h_1, h_2, \dots, h_n)$, where for each $i = 1, 2, \dots, n$, h_i:$[0, t) \to P \cup \{q, a, l\}$, $h_i^{-1}(P \cup \{a\})$ is a finite set, and $h_i^{-1}(l)$ is either empty or a singleton.

Denote by:

(i) $H(t)$, the set of all histories up to time $t \geqslant 0$, where $H(0) = \varnothing$ denotes the empty history;

(ii) $H = \bigcup_{t=0}^{\infty} H(t)$, the set of all histories;

(iii) $p(h) \in P$, the current proposal according to h;

(iv) $\tau(h)$ for $h \in H(t)$, the amount of time, according to h, that has elapsed up to t since $p(h)$ was proposed;

(v) $N(h) \subseteq N$, the set of players who have not yet left to consume according to h;

(vi) $A(h) \subseteq N(h)$, the set of players who have accepted $p(h)$ according to h; and

(vii) $\Pi(h)$, the set of current binding proposals among players in $N(h)$ according to h. Consequently, $\Pi(h)$ consists of a finite list of proposals whose associated coalitions are mutually disjoint.

We say that player i *has accepted the current proposal* (x, S) *according to* $h \in H(t)$ if the current proposal was made at time $\bar{t} < t$ and $h_i(t') = a$ for some $t' \in (\bar{t}, t)$. If every member of S has accepted (x, S), we say that (x, S) *has been accepted* (by S). Also, denote by $h|t$ the restriction of $h \in H$ to $[0, t)$.

With the exception of $p(h)$, $\tau(h)$, and $A(h)$, all entries in the above list are well-defined owing to the fact that but for finitely many points in time, all players are quiet according to any history. To render $p(h)$ well-defined, it suffices to define $p(h) = \varnothing$ if, according to h, one of the following holds: no proposal has been made, or the most recent proposal, (x, S), either has been accepted by S, or coincides with the simultaneous announcement of a distinct proposal. Note that once the proposal (x, S) is accepted by all members of S, it ceases to be the current proposal and is added to the list of binding proposals. The remaining two entries, $\tau(h)$ and $A(h)$, are now taken care of by setting $\tau(h) = 0$ and $A(h) = \varnothing$, whenever $p(h) = \varnothing$.

Payoffs: If a coalition, S, accepts (x, S) and any member of S subsequently leaves the game before accepting a new proposal, then each $i \in S$ is forced to leave and receives the payoff x_i. If player i never leaves and is never forced to leave, he obtains a payoff of $d_i \leqslant v(i)$.

Strategies: A *strategy*, f_i, for player i specifies, for each history, a choice for player i. That is: f_i: $H \rightarrow P \cup \{q, a, l\}$ for all $i \in N$.

We impose the following restrictions upon strategies:

(S1) *If, according to* $h \in H$, *i has accepted the current proposal or i has left to consume, then* $f_i(h) = q$.

(S2) *If, according to* $h \in H$, (x, S) *is a current binding proposal and* $f_i(h) = (y, S')$, *then either* S' *contains* S, *or they are disjoint.*

(S3) *If, according to* $h \in H$, *i is not a member of any formed coalition, then* $f_i(h) \neq l$.

(S4) $\forall i, \forall \bar{t} > 0$, $\forall h \in H(\bar{t})$, $\forall t \in [0, \bar{t})$, $\exists \varepsilon > 0$ *s.t.* $f_i(h|\tau) = q$ $\forall \tau \geqslant 0$ *satisfying* $\tau \in (t - \varepsilon, t + \varepsilon) \setminus \{t\}$.

Conditions (S1)–(S3) merely ensure that the choices made by the players, according to their strategies, are consistent with the rules of the game as previously described. For instance, (S1) ensures that a player is quiet both after leaving to consume and after accepting a proposal until it is either accepted by the associated coalition or replaced by a new proposal. Condition (S2) guarantees that proposals made to any member of a formed coalition must include every member of that coalition, while (S3) ensures that no player can leave without having first accepted a proposal. In particular, in order to leave and consume his own value, any player, i, must first propose $(v(i), \{i\})$, later accept it, and later still leave. In the present continuous-time context this does not restrict in any way a player's freedom to leave and consume. It does, however, ensure that the other players are given ample warning before anyone leaves.

Condition (S4) plays a dual role. Its more substantive one is to ensure that no player wishes to make a proposal, accept a proposal, or leave, as near as possible, yet before or after, any fixed time $t \geqslant 0$. This guarantees, for instance, that if player 1, say, recognizes that according to player 2's strategy, player 2 intends to leave, then player 1 can make a proposal to player 2, with a view to convincing him to stay, at a time during which all others are quiet and player 2 has not yet left. This is clearly important in order to obtain only core outcomes. For it may be that players 1 and 2 must form a coalition to attain a core allocation. This could not be guaranteed were player 1 not ensured of the opportunity to "convince" player 2 to stay by, say, making an appropriate blocking proposal.

As a by-product, condition (S4) also serves a purely technical role. It guarantees that the game is well-defined. To be more precise requires a short discussion. Call $h = (h_1, \ldots, h_n)$ a *path*, if for some $\bar{t} \in (0, \infty]$, (a) $h|t \in H(t)$ for every $t \in [0, \bar{t})$ and (b) $\bar{t}$ is finite iff it is the first time (hence the only time) such that all players have left according to h (i.e. the smallest t' with $h_i(t') = l$ for all $i \in N \backslash N(h|t')$). It is straightforward to determine from every path a unique payoff vector for the players according to the rules described above. Thus, to demonstrate that the game is well-defined, it suffices to show that any strategy-tuple, f, satisfying (S1)–(S4) always induces a unique continuation path. That is, given any $h' \in \mathrm{H}(t')$ it suffices to show that there is a unique path, h, with its associated $\bar{t} \in (0, \infty]$ satisfying (i) $h = h'$ on the domain of h', and (ii) $f(h|t) = h(t)$ for every $t \in [t', \bar{t}]$. That (S4) is (slightly more than) enough to guarantee this is by now rather well-known (see, for instance, Bergin and Macleod (1989), Stinchcombe (1988), and Perry and Reny (1993)). Consequently, (S4) ensures that after any history, any strategy-tuple induces a unique continuation-path which, by definition, takes on the value q ("quiet") for every player for all but finitely many points in any bounded interval of time. In what follows, let $u_i(f|h)$ denote player i's payoff determined from the continuation-path induced by the strategy-tuple f after $h \in H$, and let F_i denote the set of strategies for i satisfying (S1)–(S4).

Solution concept: A strategy profile $\hat{f} = (\hat{f}_1, \hat{f}_2, \ldots, \hat{f}_n)$ is a *stationary subgame perfect equilibrium* (SSPE) if:

(E1) PERFECTION: *For all $i \in N$ and for all $h \in H$,*

$$u_i\left(\left(\hat{f}_i, \hat{f}_{-i}\right)|h\right) \geqslant u_i\left(\left(f_i, \hat{f}_{-i}\right)|h\right) \quad \text{for all} \quad f_i \in F_i.$$

(E2) STATIONARITY: *Players' strategies depend only upon the time that has elapsed since the current proposal was submitted, the current proposal, the set of players who have accepted it, the set of players remaining in the game, and the set of binding proposals among players still in the game. That is, if $h, h' \in H$ satisfy $\tau(h) = \tau(h')$, $p(h) = p(h')$, $A(h) = A(h')$, $N(h) = N(h')$ and $\Pi(h) = \Pi(h')$, then $\hat{f}(h) = \hat{f}(h')$.*

For any history, h, call $(\tau(h), p(h), A(h), N(h), \Pi(h))$ the *state according to h*. Note that the inclusion of τ in the state variable has the effect of allowing stationary strategies to depend, in a limited sense, upon time. Some time dependence is necessary to ensure existence in the present environment. For instance, in the one seller/two buyers example, in order to achieve the unique core point, x^*, as an SSPE outcome, a proposal involving the allocation x^* must be made and then subsequently accepted. However, if a player's strategy instructs him to accept this proposal with no reference to time, then it must be accepted "as soon as possible," yet after, it is made. This violates (S4). To avoid this difficulty, the proposal can instead be accepted, say, one unit of time after it is made. This is compatible with condition (E2).

4. THE THEOREMS

THEOREM 1: *Every stationary subgame perfect equilibrium outcome is in the core of (v, N).*

PROOF: Suppose that x is an SSPE outcome but that it is not in the core. We shall show by way of contradiction that x cannot be supported by an SSPE. So suppose that x can be supported by the SSPE strategy-tuple $f=(f_1,f_2,\dots f_n)$. Clearly $x_i \geqslant v(i)$ for every $i \in N$. Further, since x is not in the core, there is a proposal, (y,S), with $y_i > x_i$ for all $i \in S$. Let $S=\{1,2,\dots,k\}$.

Consider, at time t, any history, $h \in H(t)$, such that (i) $N(h)=N$, (ii) $\Pi(h)=\varnothing$, (iii) players $1,2,\dots,k-1$ have accepted (y,S) according to h, and (iv) $f_i(h)=q$ for all $i=1,2,\dots n$.

CLAIM 1: *According to f, player k will, before any new proposal is made, accept (y,S).*

PROOF OF CLAIM: If, according to f, a proposal, (z,T), is made before player k accepts (y,S) then after it is made, the state variable will be rendered equal to $(0,(z,T),\varnothing,N,\varnothing)$. Moreover, because player k could have accepted the previous proposal, guaranteeing himself a payoff of $y_k > x_k$, but chose instead to allow it to be replaced by (z,T), the continuation must yield k a payoff of say, $w_k \geqslant y_k$. Therefore, by stationarity, whenever the state variable takes on the value $(0,(z,T),\varnothing,N,\varnothing)$, player k's subsequent payoff, according to f, is $w_k \geqslant y_k > x_k$. But this contradicts f being an equilibrium since player k can render the state equal to $(0,(z,T),\varnothing,N,\varnothing)$ by proposing (z,T) near enough time zero ((S4) ensures this). This proves Claim 1.

Next, consider at time t any history $h \in H(t)$ such that (i) $N(h)=N$, (ii) every player $i=1,2,\dots k-2$ has accepted (y,S) according to h, and (iii) $f_i(h)=q$ for $i=1,2,\dots n$.

CLAIM 2: *According to f, the proposal (y,S) will be accepted by players $k-1$ and k (in some order) before any new proposal is made.*

PROOF OF CLAIM 2: Note that player $k-1$ can guarantee himself a payoff of $y_{k-1} > x_{k-1}$ by accepting (y,S) at time t. This is because for some amount of time after $k-1$'s acceptance, by (S4) all players must be quiet. But then at any time during this interval, the history will be such as to induce player k to accept (y,S) in the continuation, by the previous claim. Thus, indeed player $k-1$ can guarantee himself a payoff of y_{k-1}.

Now then, in the continuation after h, suppose by way of contradiction, that (y,S) is not accepted by players k and $k-1$. In this continuation then, either no *new* proposal is made according to f, or one is. In the former case, $k-1$ and k would receive payoffs of d_{k-1} and d_k respectively, contradicting $k-1$'s ability to achieve a payoff of at least $y_{k-1} > x_{k-1} \geqslant d_{k-1}$. And in the latter case, we arrive at a contradiction as in the previous claim. That is, player $k-1$'s payoff subsequent to the new proposal must be at least y_{k-1}. But $k-1$ could then himself make this proposal near enough time zero rendering the state variable such that this subsequent payoff (by stationarity) would be at least $y_{k-1} > x_{k-1}$. This final contradiction establishes Claim 2.

Proceeding inductively with analogous Claims $3, 4, \dots, k-1$ establishes that for any time t and any history $h \in H(t)$ such that (i) $N(h) = N$, (ii) player 1 has accepted (y, S) according to h, and (iii) $f_i(h) = q$ for $i = 1, 2, \dots, n$, the continuation according to $(f_1, f_2, \dots, f_n)$ involves the coalition, S, accepting (y, S).

But this contradicts x being an SSPE since player 1 can always offer (y, S) and then accept it sufficiently near time zero (again, by (S4)) so that precisely such a history is generated, thereby ensuring himself a payoff of $y_1 > x_1$. *Q.E.D.*

REMARK 1: Note how the line of proof employed above follows the story that is commonly used to motivate the notion of the core. A proposal whose associated allocation is not in the core cannot be supported since a blocking coalition is able to coordinate a mutually preferred proposal *before* the candidate proposal is accepted.

REMARK 2: An implication of Theorem 1 is that (v, N) having a nonempty core is necessary for the existence of an SSPE. Indeed, since the arguments above apply after any history in which only a subset of the players remain, the game (v, N) restricted to any subset of players must also have a nonempty core. Consequently, if an SSPE is to exist, (v, N) must be totally balanced. We now establish that this condition is, in fact, sufficient.

THEOREM 2: *If (v, N) is totally balanced, then any element of its core can be supported as an SSPE outcome.*

The proof of Theorem 2 is constructive and it is contained in the Appendix.

REMARK 3: Taken together, Theorems 1 and 2 provide a core implementation result for totally balanced TU games.

REMARK 4: Suppose that the underlying TU game is generated by an exchange economy in which consumers' preferences are either (i) identical, continuous, convex, and homothetic, or (ii) quasilinear. If there are sufficiently many consumers, the core equivalence theorem then implies that the outcome of our game must be nearly competitive.

5. TIME PREFERENCE

The analysis above relies on the absence of time preference. Indeed with time continuous, the presence of strict time preference precludes existence since players would wish to accept a proposal at the earliest possible time after it is made. However, no such earliest time is available. There is, therefore, a tension between the rigors of perfect rationality and time preference in this continuous-time framework. A natural resolution of this tension is to nonetheless include time preference and to consider instead, ε-equilibria.

With the inclusion of time preference, an outcome must specify not only the division of the surplus among the players, but also the time at which consumption takes place. Consequently, an outcome is a $2n$-tuple, $(x_i, t_i)_{i=1}^n$, where x_i denotes player i's portion of the surplus and t_i denotes the time that player i consumes it. We restrict attention to the case in which players discount the future through a discount factor $\delta_i \in (0,1)$. Consequently, the outcome $(x_i, t_i)_{i=1}^n$ yields player i a payoff of $x_i(\delta_i)^{t_i}$. In this setting we assume that $d_i = 0$ for all i.

We say that a feasible allocation x belongs to the *ε-core* of the game (v, N) if for every coalition $S \subseteq N$, whenever $\Sigma_{i \in S} y_i \leqslant v(S)$, it is the case that $y_j \leqslant x_j + \varepsilon$ for some $j \in S$. We also define a profile of stationary strategies f^* to be an $\varepsilon - \text{SSPE}$ if for every history $h \in H$ and for all i, $u_i(f^* | h) + \varepsilon \geqslant u_i((f_i, f_{-i}^*) | h)$, for all $f_i \in F_i$.

THEOREM 3:[6] *If payoffs are discounted and (v, N) is totally balanced, then as $\varepsilon \to 0$, the set of ε-SSPE allocations converges (in Hausdorff distance) to the set of core allocations. Moreover, the time at which any fixed amount of positive consumption takes place converges to zero.*

PROOF: See the Appendix.

REMARK 5: The result expressed in Theorem 3 should be contrasted with what one would expect were discounting introduced in a discrete-time setting. Together with stationarity, discounting and discrete time would, by virtue of Rubinstein-like arguments, not allow one to support every core allocation as an $\varepsilon - \text{SSPE}$. And this would be so for all $\varepsilon > 0$ small enough. Maintaining the presence of continuous time is critical in order that the results be robust to the inclusion of time preference.

REMARK 6: Because players can react arbitrarily quickly when time is continuous, the conclusions of Theorem 3 are valid for all positive discount factors less than one.

6. WHY CONTINUOUS TIME?

Our decision to employ a continuous-time framework is related to our desire to implement the core in a *natural* manner. Since what is natural to one may not be to another, we cannot possibly *prove* that a continuous-time model is necessary to implement the core in a natural way. We seek only to describe why we are otherwise unable to provide such an implementation.

[6] We are grateful to an anonymous referee for pointing out an error in an earlier version of this Theorem.

In our view, a natural model of the story motivating the core must be dynamic in the sense that players must not be forced to move simultaneously. In addition, they must be given the opportunity to make *proposals*, namely offers to specific coalitions which also describe the division of the surplus among members of the coalition. Moreover, in order for a proposal to become binding, it must first be accepted by (at least) each of the players to whom it was directed.

Under these restrictions, the central drawback of the standard discrete-time approach can be uncovered by focusing attention on but one modeling decision: When a proposal is made, who has an opportunity to reject it? In our view, there are two natural possibilities: (i) all players or (ii) only those to whom the proposal is directed. However, in discrete time, neither of these options serves to provide a natural implementation of the core. Giving all players an opportunity to reject any proposal endows each with the ability to hold out for any prespecified payoff. The core then has no chance. But when some players do not have the opportunity to reject a proposal, and it happens that other players they need to obtain their core payoff do not need them in return, those other players may form a coalition, leaving the former players without a hope of obtaining their part of the core allocation. For every sufficiently rich discrete-time model we considered, the above argument could be applied to conclude that the core could not be implemented via that model.

How does continuous time help? In our continuous-time model, all players have the opportunity to reject the equilibrium proposal simply by making another proposal before it is accepted by each coalition member to whom it is directed; there is always a moment in time available to do so before the equilibrium agreement is consumated. However, in continuous time, this does not imply that each player can hold out by continually rejecting until he obtains his desired amount as in the discrete-time case. Indeed, players in our continuous-time model cannot reject a proposal "as soon as possible," yet after it is made. According to their strategy, they must be quiet for some positive (although potentially arbitrarily small) amount of time after a proposal is made; and this is ample time for those to whom the offer is directed to accept it.

Thus, in continuous time, it is possible to give each player the ability to reject any candidate equilibrium *without* simultaneously empowering each player to dictate the outcome by employing a strategy that undermines any proposal incompatible with it. We are unable to accomplish this in a natural way in discrete time.

APPENDIX

The proof of Theorem 2 below relies on two lemmas, Lemmas 1 and 2. These are stated and proved immediately following the proof of Theorem 2.

THEOREM 2: *If (v, N) is totally balanced, then any element of its core can be supported as an SSPE outcome.*

Before providing a proof, we remind the reader of the feature highlighted by Example 1a, namely, the need to keep track of the relevant TU game. When there are binding proposals present, each player can guarantee himself the payoff associated with his part of the binding proposal. Also, as shown in Example 2, some players may be needed in order to achieve a core allocation and so must be convinced to stay despite having reached a previous agreement. Moreover, if a proposal is one that has been accepted by some coalition members and will be accepted by the others, then it must not be the case that players outside the coalition are hurt by this (otherwise they will try to replace the proposal before it is accepted, while the coalition members will try to accept it before it is replaced).

Two functions, $z(h)$ and $\hat{z}(h)$, are defined below with all of these requirements in mind. Both provide for any history, h, a vector of payoffs to the remaining players. The strategies constructed in the proof employ these functions as follows. If there is a current proposal according to h, then $z(h)$ is the equilibrium payoff vector if the proposal is rejected, while $\hat{z}(h)$ is the equilibrium payoff vector if it is accepted. If there is no current proposal, then the two are equal and both provide the equilibrium payoff vector.

These functions have the following essential properties. For every player, i say, who is a member of a formed coalition, $z_i(h)$ is at least as large as the payoff guaranteed him by the binding agreement (Lemma 1 part (a)). Indeed, no feasible coalition can block $z(h)$ (this follows from Lemma 1 part (b)). Finally, if there is a current proposal and player i is not a member of the coalition to whom it is directed, then $z_i(h) = \hat{z}_i(h)$. Therefore, such outside players have no stake in whether or not the current proposal is accepted. Consequently, a race to eliminate a proposal before it is accepted can always be avoided.

PROOF: We provide SSPE strategies supporting an outcome, x, in the core of (v, N). For each nonempty $S \subset N$ let (v, S) denote the TU game with player set S and characteristic function v restricted to S. Choose for each such $S \neq N$ an element, $x(S)$, from the core of (v, S), and set $x(N) = x$. This is possible because (v, N) is totally balanced.

Given the collection of binding proposals $\Pi = \{(y^1, S^1), (y^2, S^2), \dots, (y^m, S^m)\}$, and $S \subset N$ such that for each k, either S contains S^k or they are disjoint, define for every $i \in S$, $z_i(\Pi, S)$ as follows (recalling that the sets S^k must be mutually disjoint):

$$z_i(\Pi, S) = \begin{cases} y_i^k + \dfrac{\sum_{j \in S^k} \left(x_j(S) - y_j^k\right)}{|S^k|}, & \text{if } i \in S^k, \\ x_i(S), \quad \text{if } i \in S \setminus \bigcup_{k=1}^{m} S^k. \end{cases}$$

If according to $h \in H$, the set of binding proposals is $\Pi(h) = \Pi$ above, and the current proposal is $p(h) = (y, S)$, then permuting the indices on the elements of $\Pi(h)$ if necessary, there is an $r \leqslant m$ such that S contains S^j for $j = 1, \dots, r$ and $S \cap S^j = \varnothing$ for $j = r+1, \dots, m$. Define $\hat{\Pi}(h)$ by $\hat{\Pi}(h) = \{(y, S), (y^{r+1}, S^{r+1}), \dots, (y^m, S^m)\}$. Note that if (y, S) were accepted by the members of S, the set of binding proposals would change from $\Pi(h)$ to $\hat{\Pi}(h)$. To simplify notation, $z_i(\Pi(h), N(h))$ and $z_i(\hat{\Pi}(h), N(h))$ will be denoted by $z_i(h)$ and $\hat{z}_i(h)$, respectively. Note that both z_i and $\hat{z}_i$ depend only upon $p(h)$, $\Pi(h)$, and $N(h)$.

In the strategies to follow, $z_i(h)$ will be player i's payoff subsequent to h if the current proposal, according to h, is rejected, while $\hat{z}_i(h)$ will be player i's payoff if the current proposal is accepted. Consequently, it is enough to compare these two values in deciding whether or not to accept the current proposal.

We now provide the desired strategies. For every $t > 0$, $h \in H(t)$ and $i \in N$,

(i) if $\tau(h)$ is not a positive integer, then $f_i(h) = q$;

(ii) if $\tau(h)$ is a positive integer:

(a) if $p(h) = (y, S)$,[7] and $\hat{z}_j(h) \geqslant z_j(h)$ for all $j \in S \setminus A(h)$, then

$$f_i(h) = \begin{cases} a, & \text{if } i \in S \setminus A(h), \\ q, & \text{otherwise;} \end{cases}$$

[7] Note then that $S \setminus A(h) \neq \varnothing$ since once a proposal is accepted, it is no longer the current proposal.

(b) otherwise[8]

$$f_i(h) = \begin{cases} l, & \text{if } i \in N(h) \backslash A(h) \text{ and } \Pi(h) = \{(z(h), N(h))\}, \\ (z(h), N(h)), & \text{if } i \in N(h) \backslash A(h) \text{ and } \Pi(h) \neq \{(z(h), N(h))\}, \\ q, & \text{if } i \in A(h). \end{cases}$$

These strategies are stationary since they depend only upon $\tau(h)$, $p(h)$, $A(h)$, $N(h)$, and $\Pi(h)$. In addition, the path generated by $f = (f_1, \dots f_n)$ has all players propose (x, N) at time 1, accept (x, N) at time 2, and leave at time 3. At all other times the players are quiet according to this path.Therefore f yields the outcome x. (For a more formal treatment, see Lemma 2 and put $h = \varnothing$ there.) It remains to check that f constitutes a subgame perfect equilibrium.

Note that if $f_i(h) = l$, then according to the above definition of f_i, it must be the case that $\Pi(h) = \{(z(h), N(h))\}$. Hence, according to f, every player in $N(h) \backslash A(h)$ chooses to leave as well and so f_i is clearly optimal given such a history. So, given $t \geqslant 0$, consider $h \in H(t)$ such that $i \notin A(h)$ and $f_j(h) \neq l$ for all $j \in N(h)$. It suffices to show that, given such a history, f_i is best for i if the others employ f_{-i}.

Let f_i' denote another strategy for i. Together, f_i' and f_{-i} generate a unique continuation path h', subsequent to h. Since (by part (a) of Lemma 1) $z_i(h) \geqslant v(i)$, and because f_i yields i at least a payoff of $z_i(h)$, i does no better if, according to h', i never leaves. So suppose that, according to h', i leaves at time t' with the proposal (y', S'), and so receives a payoff of y_i'. Since (y', S') must have been accepted by each member of S' before time t', we have $(y', S') \in \Pi(h'|t')$ so that by Lemma 1 $z_i(h'|t') \geqslant y_i'$.

We consider three exhaustive, although not mutually exclusive, cases:

Case A: $A(h) = \varnothing$.

Case B: $p(h) = (w, T)$, and either $\hat{z}_i(h) \leq z_i(h)$ or for some $j \in T \backslash A(h)$ it is the case that $\hat{z}_j(h) < z_j(h)$.

Case C: $p(h) = (w, T)$, $\hat{z}_i(h) > z_i(h)$ and for all $j \in T \backslash A(h)$ it is the case that $\hat{z}_j(h) \geqslant z_j(h)$.

Note that A covers all instances in which $p(h) = \varnothing$ (and more) and B and C cover all other cases. Since the argument applying to cases A and B have a common component, we treat these two cases together for a time, then separating them when necessary.

Cases A or B: By employing f_i, player i will obtain a payoff of $z_i(h)$. This is straightforward for Case B and follows from Lemma 2 for Case A. So, suppose by way of contradiction that $y_i' > z_i(h)$.

Let $t^* = \inf\{\hat{t} \in [t, t'] | z_i(h) < z_i(h'|\hat{t})\}$. We claim that $z_i(h) \geqslant z_i(h'|t^*)$. If $t = t^*$, then this follows since $h'|t = h$. So, suppose that $t^* > t$ and that the inequality fails. Then $z_i(h) < z_i(h'|t^*) = z_i(\Pi(h'|t^*), N(h'|t^*))$. Now by (S4), according to h', every player is quiet for times sufficiently near t^*. But this implies that for some $\varepsilon > 0$ and small enough, $\Pi(h'|t^* - \varepsilon) = \Pi(h'|t^*)$ and $N(h'|t^* - \varepsilon) = N(h'|t^*)$. Hence, $z_i(h'|t^* - \varepsilon) = z_i(h'|t^*) > z_i(h)$ contradicting the definition of t^*. We conclude that indeed $z_i(h) \geqslant z_i(h'|t^*)$, as claimed. Note then that $t^* < t'$.

Therefore, for some sequence of positive numbers $\{\varepsilon_n\}$ converging to zero, we have $z_i(h'|t^* + \varepsilon_n) > z_i(h) \geqslant z_i(h'|t^*)$. But this can happen, again by (S4), only if according to h', at time t^* either the current proposal is accepted by a coalition containing i, or someone leaves the game. The latter is not possible since i leaves at $t' > t^*$ and according to their strategies if some $j \neq i$ leaves at $t^* < t'$, then *all* $j \neq i$ leave at t^*. But this would mean that at time t', player i is alone and so $y_i' = v(i) \leqslant z_i(h)$, a contradiction. Hence, according to h', at time t^* the current proposal $p(h'|t^*) = (y, S)$ with $i \in S$, is accepted.

Now, because $\hat{z}_i(h'|t^*)$ is precisely the resulting value of $z_i(\cdot)$ were the current proposal accepted, and because by (S4) all players are quiet for times near enough t^*, we have $\hat{z}_i(h'|t^*) = z_i(h'|t^* + \varepsilon_n)$ for n large enough. But $z_i(h'|t^* + \varepsilon_n) > z_i(h'|t^*)$ then implies that $\hat{z}_i(h'|t^*) > z_i(h'|t^*)$. Since by part (c) of Lemma 1, $\sum_{j \in S} \hat{z}_j(h'|t^*) = \sum_{j \in S} z_j(h'|t^*)$, there must, since $i \in S$, be a $k \in S$ such that $\hat{z}_k(h'|t^*) < z_k(h'|t^*)$. We now continue from this point with Case A only. The continuation of the argument for Case B will be provided immediately after.

[8] When h, by virtue of (b), renders $f_i(h) = (z(h), N(h))$, it must be ensured that $(z(h), N(h))$ is a feasible proposal. This is the only place where transferable utility is used. To see that, in fact, $(z(h), N(h))$ is a feasible proposal for *every* history, h, note that $\sum_{i \in N(h)} z_i(h) = \sum_{i \in N(h)} x_i(N(h)) = v(N(h))$, where the first equality follows by definition of $z(\cdot)$, and the second follows because $x(N(h))$ is in the core of $(v, N(h))$. Also, $z_i(h) \geqslant v(i)$ for all $i \in N(h)$, since $z_i(h) \geqslant y_i^k \geqslant v(i)$ if $i \in S^k$ (the first inequality by Lemma 1 and the second because (y^k, S^k) is feasible); and $z_i(h) = x_i(N(h)) \geqslant v(i)$ if $i \in N(h) \backslash \cup_{k=1}^m S^k$ (the inequality holds because $x(N(h))$ is in the core of $(v, N(h))$).

Case A: $A(h)=\varnothing$: Now, because at time t^* the proposal (y,S) is accepted by all members of S who have not yet accepted it, and $A(h)=\varnothing$, there is a time $t_k \in [t,t^*]$ such that player k according to h' accepts (y,S) at time t_k. But since $k \neq i$ is playing according to f_k, it must then be the case that $\hat{z}_k(h'|t_k) \geqslant z_k(h'|t_k)$. But since the current proposal, the collection of binding proposals, and the players remaining in the game are unchanged from t_k to t^* we have $\hat{z}_k(h'|t_k)=\hat{z}_k(h'|t^*)$ and $z_k(h'|t_k)=z_k(h'|t^*)$. But this contradicts $\hat{z}_k(h'|t^*)<z_k(h'|t^*)$. We conclude that $z_i(h) \geqslant y_i'$. This completes the argument for Case A. We now pick up where we left off and conclude the argument for Case B.

Case B: $p(h)=(w,T)$ and either $\hat{z}_i(h) \leqslant z_i(h)$, or for some $j \neq i\ j \in T \backslash A(h)$, it is the case that $\hat{z}_j(h)<z_j(h)$.

Claim: *According to h', the proposal $p(h)=(w,T)$ is either rejected by some member of T, or player i replaces it by another proposal before it is accepted.*

Proof of Claim: Since (y,S) is proposed at time t^* according to h', it suffices to show that according to h', $(w,T)=p(h)$ is not accepted by all members of T. If for some $j \neq i$, $j \in T \backslash A(h)$ and $\hat{z}_j(h)<z_j(h)$, then according to f_j player j will not accept $p(h)$. Since, while the current proposal remains $p(h)=(w,T)$, $\hat{z}_j(\cdot)$ and $z_j(\cdot)$ remain constant (recall that the collection of remaining players is constant from t to t' according to h'), player j will not, according to f_j, accept it. The only other possibility is that $\hat{z}_i(h) \leqslant z_i(h)$ and for $j \neq i$ and $j \in T \backslash A(h)$ $\hat{z}_j(h) \geqslant z_j(h)$. But in this case, because $p(h)$ will be accepted by all remaining $j \in T \backslash A(h)$ according to f_j, player i can play so that $p(h)$ is accepted. If he does so, then just after it is accepted, there will be a period of time during which all players are quiet, the value of z_i will be $\hat{z}_i(h)$, and during which there will be no current proposal. Since by Case A it is optimal from this point on for i to play according to f_i, his resulting payoff, y_i', will (by Lemma 2) be equal to $\hat{z}_i(h) \leqslant z_i(h)$, a contradiction. This proves the Claim.

By the Claim, the proposal (y,S) must then be made at time t or later according to h'. In addition, (y,S) is accepted by every member of S at or before time t^*. Hence, there is a time $t_k \in [t,t^*]$ such that player $k \neq i$ (that player satisfying $\hat{z}_k(h'|t^*)<z_k(h|t^*)$ from the argument above applying to both Cases A and B), according to h' accepts (y,S) at time t_k. But as in the last part of the argument in Case A, this is inconsistent with player k's strategy, f_k. This completes the argument for Case B.

Case C: $p(h)=(w,T)$, $\hat{z}_i(h)>z_i(h)$, and for every $j \in T \backslash A(h)$, $\hat{z}_j(h) \geqslant z_j(h)$. If all players play according to f, then i obtains the payoff $\hat{z}_i(h)$. We wish to show that the strategy f_i' does not yield i a payoff above $\hat{z}_i(h)$.

Now, all $j \in T \backslash A(h)$ other than i (recall that $i \notin A(h)$ and note that $\hat{z}_i(h)>z_i(h)$ implies $i \in T$) will, according to f_j, accept the proposal (w,T) if it is not replaced by a new proposal. Also, all $j \notin T$ will, according to f_j, be quiet until (w,T) is accepted or replaced by a new proposal. So, if i accepts (w,T) before making a new proposal, then (w,T) will be accepted by all members of T. Subsequently, there will (for a period of time) be no current proposal. By Case *A*, it is then optimal for i to employ f_i and thereby obtain $\hat{z}_i(h)$, the new value of z_i. Thus, i cannot improve upon f_i by accepting (w,T) before making a new proposal.

So, suppose that before i accepts (w,T), i makes a new proposal, (w',T'). For a period of time subsequent to i's proposal the current proposal will be (w',T'), no player will have accepted it, and the value of z_i will remain equal to $z_i(h)$. Since, during this period, no player has accepted the current proposal, the analysis of Case A indicates that i can do no better than obtain a payoff of $z_i(h)$. Since $z_i(h)<\hat{z}_i(h)$ we conclude that i cannot improve upon f_i by making a proposal instead of accepting (w,T).

The only remaining option is for i to leave before making a proposal and before accepting (w,T). But this yields i a payoff no greater than $z_i(h)$ regardless of when i leaves since the others, according to their strategies, either are quiet or accept the current proposal. But $z_i(h) \leqslant \hat{z}_i(h)$ and so this is no better than employing f_i either. This completes the argument for Case C. *Q.E.D.*

Lemma 1: *Consider the function $z_i(\cdot)$ as defined in the proof of Theorem 2. For every $h \in H$, if $p(h)=(y,S)$ and $\Pi(h)=\{(y^1,S^1),\ldots,(y^m,S^m)\}$, then*

(a) $z_i(h) \geqslant y_i^k$ *for every $k=1,\ldots,m$ and $i \in S^k$,*

(b) $\sum_{i \in S^k} z_i(h)=\sum_{i \in S^k} x_i(N(h))$ *for every $k=1,\ldots,m$, and*

(c) $\sum_{i \in S} z_i(h)=\sum_{i \in S} \hat{z}_i(h)$.

PROOF: (a) For every k, $\Sigma_{i \in S^k} y_i^k \leqslant v(S^k)$ since (y^k, S^k) is a feasible proposal. In addition, because $x(N(h))$ is in the core of $(v, N(h))$ and $S^k \subseteq N(h)$ we have $\Sigma_{i \in S^k} x_i(N(h)) \geqslant v(S^k)$. Hence $\Sigma_{i \in S^k}(x_i(N(h)) - y_i^k) \geqslant 0$. This, together with the definition of $z_i(h)$ yields (a).

(b) This follows trivially from the definition of $z_i(h)$.

(c) Direct calculation, using the definitions of z_i and $\hat{z}_i$, establishes that both sums are, in fact, equal to $\Sigma_{i \in S} x_i(N(h))$. *Q.E.D.*

LEMMA 2: *For any $t \geqslant 0$ and $h \in H(t)$, if $A(h) = \varnothing$, then the outcome generated by $(f_1, \ldots, f_n)$ after h is $z(h)$.*

PROOF: Let t_1 be the smallest time at least as large as t such that $t_1 - \tau(h)$ is an integer. According to f, all players are quiet between t and t_1. Call the history so far generated $h_1 \in H(t_1)$. Now if according to h_1, $f_i(h_1) = l$, then the outcome will clearly be $z(h_1)$ and this is equal to $z(h)$. If $f_i(h_1) = (z(h_1), N(h_1))$, then between t_1 and $t_2 = t_1 + 1$, all players will be quiet according to f. Call the history generated this way $h_2 \in H(t_2)$. Since $p(h_2) = (z(h_1), N(h_1))$, we have $\hat{z}(h_2) = z(h_1)$. To see this, note that since no player has left between t and t_2, $N(h) = N(h_1) = N(h_2)$ and by definition:

$$\hat{z}_i(h_2) = z_i(h_1) + \frac{\sum_{j \in N(h_1)} \left(x_j(N(h_1)) - z_j(h_1) \right)}{|N(h_1)|} \quad \text{for all } i \in N(h_1).$$

But now use part (b) of Lemma 1 to conclude that $\hat{z}_i(h_2) = z_i(h_1)$. Since no offer has been accepted between t and t_2 according to h_2, we have $z_i(h_2) = z_i(h_1) = z_i(h)$ for every $i \in N(h)$. Hence $\hat{z}_i(h_2) = z_i(h_2)$. Therefore, given h_2, according to f, all players in $N(h) = N(h_2)$ accept $p(h_2) = (z(h), N(h))$. Finally, at time $t_3 = t_2 + 1$ all players, having been quiet between t_2 and t_3 leave according to f, yielding the outcome $z(h)$.

Finally, suppose that $p(h_1) = (y, S)$ and that $f_i(h_1) = a$ for some $i \in S$. Since $A(h) = A(h_1) = \varnothing$, this means that $\hat{z}_i(h_1) \geqslant z_i(h_1)$ for every $i \in S$, and $f_i(h_1) = q$ for all other i. Moreover, by part (c) of Lemma 1 it must then be the case that $\hat{z}_i(h_1) = z_i(h_1)$ for all $i \in S$. Since for all $i \in N(h_1) \setminus S$, $\hat{z}_i(h_1) = z_i(h_1)$, we then have $\hat{z}_i(h_1) = z_i(h_1)$ for all $i \in N(h_1)$. According to f, all players are quiet until time $t_2 = t_1 + 1$. Let $h_2 \in H(t_2)$ be the history so far generated. Accordingly, $p(h_2) = \varnothing$ so that $f_i(h_2) = (z(h_2), N(h_2))$. But for every $i \in N(h_2) = N(h)$, $\hat{z}_i(h_2) = \hat{z}_i(h_1) = z(h)$, so that $f_i(h_2) = (z(h), N(h))$. As before, in one unit of time all players accept $(z(h), N(h))$ and in another unit of time all players leave. *Q.E.D.*

THEOREM 3: *If payoffs are discounted and (v, N) is totally balanced, then as $\varepsilon \to 0$, the set of ε-SSPE allocations converges (in Hausdorff distance) to the set of core allocations. Moreover, the time at which any fixed amount of positive consumption takes place converges to zero.*

PROOF: The proof is split into two parts.

Part I. We first show that any core allocation can be supported as an ε-SSPE with consumption taking place arbitrarily quickly. Indeed, to see this, it suffices to employ the strategies provided in the proof of Theorem 2 with the following modifications. Let $T_\Delta = \{\Delta, 2\Delta, 3\Delta, \ldots\}$ for $\Delta > 0$. Replace in (i) the phrase "if $\tau(h)$ is not a positive integer" by "if $\tau(h) \notin T_\Delta$, and replace in (ii) the phrase "if $\tau(h)$ is a positive integer" by "if $\tau(h) \in T_\Delta$." For Δ close enough to zero, the resulting strategies will constitute an ε-SSPE, with all players leaving and consuming at time 3Δ.

Part II. For each $\varepsilon > 0$, let f^ε denote an ε-SSPE supporting the outcome $(x_i^\varepsilon, t_i^\varepsilon)_{i=1}^n$, and let $(x_i, t_i)_{i=1}^n$ denote one of the limit points of the sequence of outcomes as $\varepsilon \to 0$.[9] By Part I, it is enough to show that (*i*) x is in the core of (v, N); and (ii) $t_i = 0$ for all i with $x_i > 0$.

Proof of (i): Suppose, to the contrary, that x is not in the core. Consequently, there is a proposal (y, S) and an $\eta > 0$ with $y_i > x_i + \eta$ for all $i \in S$. Choose now $\varepsilon > 0$, smaller than η so that $y_i > x_i^\varepsilon + \eta$ for all $i \in S$. Since the proof so closely parallels that of Theorem 1, we will consider only

[9] The sequence $\{x_i^\varepsilon\}$ has finite limit points because in any ε-SSPE, for all i, $0 \leqslant x_i^\varepsilon \leqslant v(N)$. However, the sequence $\{t_i^\varepsilon\}$ need not have finite limit points (consider the game $v(S) = 0$ for all S). Consequently, we allow $t_i = +\infty$ as a limit point here.

the case in which S contains but two players, 1 and 2, say. Consider then the following deviation from f^{ε} by player 1. Propose $(y,\{1,2\})$ at time $\eta_1 > 0$ where, according to f^{ε}, all players are quiet in the time interval $(0, 2\eta_1)$. In addition, accept it at time $\eta_2 > \eta_1$ where, according to f^{ε} and subsequent to the initial deviation, all players are quiet in the time interval $(\eta_1, 2\eta_2)$. Now, if η_1 and η_2 are small enough, then player 2 must in the continuation accept this proposal. Why? First, note that if accepted by player 2 soon enough, it can yield player 2 a payoff arbitrarily close to $y_2(\delta_2)^{\eta_2}$ which, for η_2 small enough, can be made arbitrarily close to $y_2 > x_2^{\varepsilon} + \eta$. Consequently, for η_1 and η_2 small enough, player 2, by accepting this proposal soon enough can obtain a payoff in excess of $x_2^{\varepsilon} + \varepsilon$, since $\varepsilon < \eta$. Therefore, if player 2 does not accept it, player 2 must be otherwise guaranteed a payoff in excess of $x_2^{\varepsilon} + \varepsilon$. But this can only come about if player 2 replaces the proposal, $(y,\{1,2\})$, with a new one. However, by stationarity, player 2 could have made this new proposal and achieved the higher payoff whether or not player 1 deviated. But this contradicts the hypothesis that x^{ε} is an ε-SSPE outcome. Hence, indeed player 2 must accept the proposal $(y,\{1,2\})$ made and accepted by player 1. It remains to argue that player 2 accepts this proposal soon enough so that it is a profitable deviation for player 1. But this is obvious so long as ε is small enough because player 2 (strictly) discounts the future. That is, for any $\varepsilon > 0$ and associated $\eta_1, \eta_2 > 0$ described above, if ε is small enough, then player 2 must accept $(y,\{1,2\})$ arbitrarily near η_2. Consequently, if we choose $\varepsilon, \eta_1, \eta_2 > 0$ small enough, player 1's payoff from the deviation can be made arbitrarily close to $y_1 > x_1^{\varepsilon} + \eta > x_1^{\varepsilon}(\delta_1)^{t_1^{\varepsilon}} + \varepsilon$. But the second inequality yields the final contradiction; namely that the deviation provides player 1 with a payoff in excess of ε above his alleged ε-SSPE payoff. This completes the proof of (i).

Proof of (ii): Suppose that for some player, j say, $x_j > 0$ and $t_j > 0$. Now, if $x_j \leqslant v(j)$, then for ε small enough, $x_j^{\varepsilon}(\delta_j)^{t_j^{\varepsilon}} + \varepsilon < v(j)$ so that j can unilaterally improve his payoff by more than ε by proposing $(v(j),\{j\})$, accepting it and leaving close enough to time zero. Consequently, it must be the case that $x_j > v(j)$. Since t_j^{ε} is the time that player j consumes, player j must leave, according to f^{ε}, at time t_j^{ε} (finite) with the binding proposal $(x_{T_{\varepsilon}}, T_{\varepsilon})$, where T_{ε} is some subset of players and $x_{T_{\varepsilon}}^{\varepsilon}$ is the projection of x^{ε} onto T_{ε}. Since there are but finitely many coalitions, we may assume that T_{ε} does not depend upon ε. We thus write simply T and x_T. Choose $a > 0$ small enough so that the proposal (z^a, T) where $z_i^a = x_i + 2a$ for $i \neq j$, $z_j^a = x_j - 2a(|T| - 1)$, is (1) feasible; and (2) $z_i^a > x_i(\delta_i)^{t_i} + a$ for all $i \in T$. Note that satisfying (1) is possible because $x_j > v(j)$, so that $z_j^a > v(j)$ for a small enough. That (2) can also be satisfied follows trivially for $i \neq j$, and for $i = j$ because both x_j and t_j are strictly positive. Now, for $\varepsilon < a$, consider the following deviation from f^{ε} by player j. At time $\eta > 0$, propose (z^a, T), where η is chosen so that, according to f^{ε}, all players are quiet in $(0, 2\eta)$. For $\eta > 0$ and small enough, if accepted by all members of T soon enough, each i in T receives a payoff arbitrarily close to $z_i^a > x_i(\delta_i)^{t_i} + a > x_i^{\varepsilon}(\delta_i)^{t_i^{\varepsilon}} + \varepsilon$ for ε small enough. By an argument similar to that employed in the proof of (i), each i in T accepts this proposal and indeed, so long as ε and η are small enough, it is accepted as close as we wish to time zero. Consequently, for ε small enough, this deviation is successful, a contradiction. This establishes (ii) and completes the proof. *Q.E.D.*

Dept. of Economics, Mount Scopus, Hebrew University, Jerusalem, 91905 Israel

and

Dept. of Economics, University of Western Ontario, London, Ontario, N6A-5C2 Canada

Manuscript received October, 1991; final revision received September, 1993.

REFERENCES

ANDERSON, R. M. (1978): "An Elementary Core Equivalence Theorem," *Econometrica*, 46, 1483–1488.

ARROW, K., AND F. HAHN (1971): *General Competitive Analysis*. San Fransisco: Holden-Day

AUMANN, R. J. (1964): "Markets with a Continuum of Traders," *Econometrica*, 34, 1–7.

BERGIN, J., AND MACLEOD (1989): "Renegotiation Proof Equilibria in Continuous Time Games," Discussion Paper #753, Department of Economics, Queens University.

BEWLEY, T. F. (1973): "Edgeworth's Conjecture," *Econometrica*, 43, 425–454.

Binmore, K. G. (1985): "Bargaining and Coalitions," in *Game Theoretic Models of Bargaining*, ed. by Alvin E. Roth. Cambridge: Cambridge University Press.

Chatterjee, K., B. Dutta, D. Ray, and K. Sengupta (1990): "A Noncooperative Theory of Coalitional Bargaining," Working Paper, Pennsylvania State University.

Debreu, G., and H. Scarf (1963): "A Limit Theorem on the Core of an Economy," *International Economic Review*, 236–246.

Edgeworth, F. Y. (1881): *Mathematical Psychics*. London: Kegan Paul.

Gul, F. (1989): "Bargaining Foundations of the Shapley Value," *Econometrica*, 57, 81–95.

Harsanyi, J. C. (1974): "An Equilibrium-Point Interpretation of Stable Sets and a Proposed Alternative Definition," *Management Science*, 20, 1472–1495.

Hart, S., and A. Mas-Colell (1992): "A Model of n-person Non-Cooperative Bargaining," Discussion Paper #10, Center for Rationality and Interactive Decision Theory, The Hebrew University of Jerusalem.

Hildenbrand, W. (1974): *Core and Equilibrium of a Large Economy*. Princeton: Princeton University Press.

Hurwicz, L. (1972): "On Informationally Decentralized Systems," in Decision and Organization, ed. by R. Radner and C. B. McGuire. Amsterdam: North-Holland.

Kalai, E., E. Postwaite, and J. Roberts (1979): "A Group Incentive Compatible Mechanism Yielding Core Allocations," *Journal of Economic Theory*, 20, 13–22.

Lagunoff, R. D. (1992): "Some Simple Noncooperative Core Stories," mimeo, University of Pennsylvania.

Maskin, E. (1977): "Nash Equilibrium and Welfare Optimality," mimeo, MIT.

Maskin, E., and J. Tirole (1988): "A Theory of Dynamic Oligopoly II, Price Competition," *Econometrica*, 56, 571–600.

Moldovanu, B., and E. Winter (1991): "Order Independent Equilibria," Working Paper, University of Bonn.

Osborne, M., and A. Rubinstein (1990): *Bargaining and Markets*. San Diego: Academic Press.

Perry, M., and P. Reny (1993): "A Non-Cooperative Bargaining Model with Strategically Timed Offers," *Journal of Economic Theory*, 59, 50–77.

Selten, R. (1981): "A Noncooperative Model of Characteristic Function Bargaining," in *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*, ed. by Bohm and H. H. Nachtkamp. Mannheim: Bibliographisches Institut, pp. 131–151.

Shubik, M. (1959): "Edgeworth Market Games," in *Contributions to the Theory of Games*, ed. by A. W. Tucker and R. D. Luce. Princeton: Princeton University Press, pp. 267–268.

Simon, L., and M. Stinchcombe (1989): "Extensive Form Games in Continuous Time, Part I: Pure Strategies," *Econometrica*, 57, 1171–1214.

Stinchcombe, M. (1988): "Maximal Strategy Sets for Continuous-Time Game Theory," forthcoming in *Econometrica*.

Vind, K. (1964): "Edgeworth Allocations in an Exchange Economy with Many Traders," *International Economic Review*, 5, 165–177.