AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

INSTITUTE FOR ECONOMIC RESEARCH

Discussion Paper # 964

An Implementation-theoretic Approach to Non-cooperative Foundations

by

James Bergin and John Duggan

November 1997

Queen's University
Kingston, Ontario, Canada K7L 3N6

An Implementation-theoretic Approach to Non-cooperative Foundations*

James Bergin

Department of Economics

Queen's University

Kingston, ON K7L 3N6

CANADA

John Duggan

Department of Political Science

and Department of Economics

University of Rochester

Rochester, NY 14627

U.S.A.

This Draft: October, 1997

Original Draft: November, 1994

* We have benefited from conversations with L. Corchón, J. Greenberg, A. Mas-Colell, B. Moldovanu, M. Osborne, P. Reny, and J. Swinkels, and from the comments of the associate editor and an anonymous referee. Both authors acknowledge the financial support of the SSHRCC. All errors are our own.

Abstract

This paper reconsiders the literature on non-cooperative foundations of cooperative solutions. The goal of non-cooperative foundations is to provide credible non-cooperative models of negotiation and coalition formation whose equilibrium outcomes agree with a given cooperative solution. Here we argue that this goal is best achieved by explicitly modeling the physical environment and individual preferences, and constructing game forms independent of preferences to implement the cooperative solution. In addition, the game form should reflect salient aspects of negotiation. We propose a general model (called a strategic environment) of the physical environment; we characterize the coalitional functions arising from strategic environments; we demonstrate our approach for the case of the core; and we provide conditions under which core payoffs correspond to payoffs from core outcomes.

1 Introduction.

The cooperative approach in game theory takes an abstract view of social interaction, combining the details of the physical environment and individual preferences in a *coalitional function* V .¹ The process of bargaining and negotiation, through which conflicting individual interests are ultimately resolved, is summarized by a cooperative solution — a mapping that associates sets of payoff vectors to coalitional functions. For example, the core picks payoff vectors such that no subcoalition can separately guarantee more for all its members, while the Shapley value assigns to each individual the average of their marginal contribution to every coalition. Cooperative solutions are often accompanied with informal stories of individual interaction but lack formal non-cooperative models describing the process of negotiation and the outcomes finally agreed upon. The development of such formal models, reconciling the predictions of cooperative solutions with the realities of individual incentives, goes back to Nash (1953) and is commonly referred to as the “Nash program.”

In the work of Nash and much of this literature, models of bargaining and negotiation are built directly on the coalitional function. Typically, negotiation is modeled in terms of proposals and counter proposals over payoffs that coalitions can achieve: a proposal to coalition S is restricted to the payoffs, $V(S)$, achievable for S , and therefore different coalitional functions result in different games. We argue that this approach is unsatisfactory. Because the coalitional function implicitly incorporates the timing of coalitional deviations and punishments, i.e., *effectivity*, this important aspect of negotiation is not explicitly reflected in traditional non-cooperative foundations. Furthermore, because individual utility functions act through the coalitional function, the procedural rules described in non-cooperative foundations technically depend on the preferences of individuals, distinguishing them from the kind of rules written down in a constitution, a legal code, or the charter of a corporation. As a consequence, models based on the coalitional function may be difficult to compare to existing institutions and their support for a cooperative theory may be limited.

The applicability of traditional non-cooperative foundations to problems of institutional design, where a social planner desires the outcomes prescribed by a cooperative theory but does not know individual preferences, is also limited. Here, the planner must construct a set of rules, defined independently of individual preferences (called a *mechanism* or *game form*), that achieves socially desirable outcomes as non-cooperative equilibrium outcomes. The central problem of mechanism design, an unavoidable consequence of the planner’s incomplete information, is to ensure that socially desirable outcomes are achieved by a *fixed* mechanism as individual preferences vary over some predetermined domain. This fundamental constraint significantly restricts the outcomes achievable by such decentralized means,² but it does not arise in the Nash program, because bargaining institutions described in terms of coali-

¹ For a coalition S , $V(S)$ denotes the vectors of payoffs to coalition members achievable (or guaranteeable) by S .

² Maskin (1977) showed that, if the normative goals of the planner are achievable in Nash equilibrium, they must satisfy a strong monotonicity condition. See Abreu and Sen (1990) and Moore and Repullo (1988) for other “preference reversal” restrictions under the subgame perfect equilibrium hypothesis.

tional functions are not fixed — they depend on individual preferences. Without preference information, traditional non-cooperative foundations offer little guidance in the design of institutions to achieve the outcomes of a cooperative solution.

There are cases in which it may matter little which approach one uses. Serrano (1996) argues persuasively that non-cooperative foundations can sometimes be transformed, overcoming the above difficulties, into game forms independent of preferences: if V is derived from a private good economy then the set $V(S)$ of payoff vectors achievable by the coalition S can be replaced by a fixed set of physical allocations, namely, redistributions of the endowments of the coalition’s members. This is true, however, only when the underlying environment is a private good economy, where externalities are not present and effectivity is not an issue. Many interesting environments (see Examples 2-5) do exhibit externalities, and in them the payoffs achievable by a coalition may fail to correspond to a fixed set of allocations or, more generally, physical outcomes. This point, and with it the concomitant inadequacies of the coalitional function, has received attention by Scarf (1971), Rosenthal (1971), and Shapley (1982), among others.³ Moreover, we show that the above-mentioned failure can be critical: in Example 8, we describe a simple environment in which the core is not implementable in Nash or subgame perfect equilibrium, despite the existence of non-cooperative foundations of the core based on the coalitional function.

We therefore propose a direct approach to non-cooperative foundations that disentangles the procedural rules of bargaining from modeling primitives by: (1) explicitly modeling physical outcomes and individual preferences, (2) formulating cooperative solutions in terms of outcomes rather than payoffs (as mappings from the underlying environment to sets of outcomes), and (3) defining rules of negotiation that are independent of preferences, that attempt to capture the salient aspects of actual bargaining, yet are consistent with the cooperative solution of interest, i.e., the equilibrium outcomes of negotiation agree with the outcomes of the cooperative solution over the domain of possible preferences. In particular, a non-cooperative foundation for a cooperative solution should *implement* the solution, now viewed as a social choice correspondence, in terms of an appropriate non-cooperative equilibrium concept. This allows us to compare actual bargaining arrangements to those described by our non-cooperative foun-

³ Shubik (1982) identifies zero-sum games and games with “orthogonal coalitions” as games where the coalition function adequately captures the strategic considerations in the game. A game with orthogonal coalitions is one where “nothing can happen to change a player’s fortune (payoff) unless he himself is party to the action In this case the only threat by outsiders against a coalition is not to belong to it.” But he observes that many games do not fall into this category: “More importantly, many games do not allow the clean separation that we need between questions of strategic optimization and questions of negotiation. Although we can still define the characteristic function for such variable threat situations, we cannot analyze them properly without additional information about the actual rules of play – information that is lost when one passes to the characteristic function.” A similar point is made by Scarf (1971), “In a model where each consumer begins the trading period with a stock of commodities, and has a utility function for final consumption, the utility vectors achievable by a coalition are most naturally taken to be those arising from an arbitrary redistribution of that coalition’s assets.” However, he observes: “If, for example, some of the goods are undesirable and require the use of real resources for their disposal, the players not included in a given coalition may, by their actions, modify the distribution of utilities within the coalition. Similar difficulties arise if external effects in consumption are introduced into the model of exchange, and in many other variations of the neoclassical model as well.” Thus, he concludes: “These examples illustrate the general proposition that the possibilities open to a coalition should perhaps be viewed as derived from a prior specification of the game in its normal form; that is, in terms of the strategic choices open to the individual players, and their evaluations of the outcomes.” Rosenthal (1971) also expresses this viewpoint arguing that the coalition function may be inadequate in its “restricted view of threat possibilities.”

dation, lending (to the extent that they correspond to one another) non-cooperative support to the predictions of a cooperative theory. From the mechanism design perspective, our non-cooperative foundation provides rules of negotiation that give individuals non-cooperative incentives consistent with the cooperative solution of interest.

This “implementation-theoretic” approach is not new: several papers have developed non-cooperative foundations, in the sense proposed here, for various cooperative solutions. Nevertheless, because both approaches involve the construction of games or game forms whose equilibria have specific features, considerable confusion surrounds the relationship between the Nash program and implementation theory. In Section 2, we review the literature on the Nash program and the relevant implementation literature, their connections, and the fundamental difference between them: because preferences are embedded in the V ’s over which a cooperative solution is defined, traditional non-cooperative foundations are necessarily parameterized by preferences. We highlight this difference in Proposition 1, which demonstrates that *from a purely technical perspective* the problem of providing a non-cooperative foundation in the Nash tradition (a collection of parameterized games with the correct equilibrium payoffs) is trivial. This is not true of the implementation problem, where game forms must be defined independently of preferences.

In Section 3, we develop a model of the physical environment, similar in spirit to Debreu’s (1952) generalized games or Ichiishi’s (1981,1993) concept of a *society*, generalizing strategic games and private good economies. We assume that each coalition has a fixed set of conceivable joint plans of action, though we allow for the feasibility of a coalition’s plans to depend on the plans of outsiders. A feasible action correspondence combined with a profile of utility functions for the individuals completes our model and is called a *strategic environment*. Fixing a feasible action correspondence, we define a cooperative solution as a mapping from strategic environments (i.e., profiles of utility functions) to subsets of feasible joint plans representing possible outcomes of social interaction. We define the core in this setting, but, in contrast to work in private good economies, we must take up the issue of effectivity. Following the conventions proposed by Aumann and Peleg (1961), we consider two concepts, α - and β -effectivity, whereby strategic environments are mapped to coalitional functions, and we define two corresponding notions of the core.

In Section 4, we extend the results of Shapley and Shubik (1969), Billera and Bixby (1973, 1974), Billera (1974), and Mas-Colell (1975), who characterize the coalitional functions derived from “market games” (i.e., private good economies), thereby establishing the domain of any cooperative theory applicable in this class of environments. We take up a logically prior but, to our knowledge, untouched issue: we establish bounds on the domain of any cooperative theory by characterizing the coalitional functions derived from, or *supported* by, general strategic environments. Moreover, these bounds are tight: we provide complete characterizations of the coalitional functions supported, in either the α - or β -senses, by strategic environments.⁴ The conditions characterizing supportability are quite permissive, strictly

⁴ In fact, strategic games are sufficient to generate all coalitional functions supported by strategic environments — the extra flexibility of our strategic environment model is unnecessary for this purpose.

more so for β -supportability. While these results are of independent interest, they also help elucidate the connections, discussed in Section 2, between implementation-theoretic non-cooperative foundations and traditional ones.

In Section 5, we fix an arbitrary feasible action correspondence and illustrate our approach with implementation-theoretic non-cooperative foundations for the two versions of the core: our game forms subgame perfect implement these versions of the core, imposing no restrictions on individual preferences and without invoking stationarity (which is often done in the work on non-cooperative foundations). Play takes place in continuous time and formalizes the usual core story, with challenging proposals upsetting non-core outcomes. In contrast to traditional non-cooperative foundations, where effectivity is embedded in the coalitional function (and unlike implementation-theoretic non-cooperative foundations in private good economies, where effectivity is moot), our game forms reflect the distinction between α - and β -effectivity through the timing of punishments incurred by deviating coalitions. This dependence is unavoidable in our framework: if different assumptions about effectivity entail distinct cooperative solutions, the rules of bargaining on which those solutions are predicated must be distinct.

Section 6 takes up an important, but overlooked, technical point. If the coalitional function of a strategic environment is calculated and the core solution applied to this coalitional function, the resulting core payoffs may not correspond to the payoffs from core plans of action in the strategic environment. We show, in Example 10, that without typical regularity conditions, e.g., compactness and continuity, on the physical environment a payoff inconsistency may arise: one individual receives a core payoff of zero, despite the fact that the individual can guarantee himself (in the α -effectivity sense) a positive payoff.⁵ We then provide appropriate regularity conditions, for α - and β -effectivity, under which payoff consistency ensured.

2 Implementation and Non-cooperative Foundations.

2.1 Comparison of Approaches.

Non-cooperative foundations traditionally begin with a cooperative solution Ψ specifying a set $\Psi(V)$ of payoffs for each coalitional function V within a given class $\mathcal{V}$. Given V , the objective is to design a game Γ whose equilibrium payoffs coincide with $\Psi(V)$ and which provides a “natural” description of individual interaction, thereby offering some justification of the cooperative theory in terms of self-interested individual behavior. Writing $\Gamma(V)$ for the game corresponding to V , Figure 1 depicts this non-cooperative foundations problem.

⁵ The individual can guarantee a positive payoff, but not one bounded above zero: the other individual in the example can bring the former’s payoff arbitrarily close to zero.

Figure 1

In the literature on the Nash program, authors have gone to great lengths to design games reminiscent of actual bargaining and negotiation. An important distinction of these games is that they are (necessarily) parameterized by coalitional functions, but, as the next proposition demonstrates, if the nature of this parameterization is unrestricted then it is a simple matter to define non-cooperative games with any desired equilibrium payoffs. In what follows, n denotes the number of individuals, indexed by i and j , and $y = (y_i)_{i \in N}$ denotes an n -tuple of payoffs.

Proposition 1 *Assume $n \geq 3$. For every cooperative solution Ψ and every coalitional function V , there exists a strategic game with Nash equilibrium payoffs equal to $\Psi(V)$.*

Proof: Assume for now that $\Psi(V) \neq \emptyset$, pick an arbitrary point $y^* \in \Psi(V)$, let each individual i 's strategy set be $M_i = \Psi(V)$ and let $M = \times_{i \in N} M_i$ with elements $m = (m_1, \dots, m_n)$. Define i 's payoff function $u_i : M \rightarrow \mathbb{R}$ as follows: $u_i(m) = y_i$ if at least $n - 1$ individuals j use strategy $m_j = y$, and $u_i(m) = y_i^*$ otherwise. The Nash equilibrium payoffs of this game are exactly $\Psi(V)$. If $\Psi(V) = \emptyset$, let $M_i = \mathbb{R}$ and $u_i(m) = m_i$. This game has no Nash equilibria, as required. (Alternatively, preferences may be parameterized by the coalitional function and strategy spaces defined independent of preferences.) ■

The games used in the proof of Proposition 1 are not adequate models of social interaction, and we do not suggest that they constitute satisfactory non-cooperative foundations for Ψ . Nevertheless, it is clear that the technical problem confronted in the Nash program is substantially simpler than that of implementation theory, where, because of the inherent difficulty of finding a single game form that works for all preference profiles, naturalness is often secondary.⁶

In this paper, we take as the point of departure a fixed set, A , of conceivable joint plans of action (and a fixed feasible action correspondence, defined in Section 3) and a set $\mathcal{U}$ of possible profiles of utility functions. Each $u = (u_i)_{i \in N} \in \mathcal{U}$ then determines one possible strategic environment in which negotiation might take place. Given a cooperative solution F , formulated as a correspondence from $\mathcal{U}$ to A ,⁷ our approach is to construct a fixed game form, G , ideally capturing the salient aspects of negotiation, but also consistent with F in the following sense: the game form, when played with

⁶ Work with natural mechanisms in implementation includes Jackson, Palfrey, and Srivastava (1994), Dutta, Sen, and Vohra (1995), and Saijo, Tatamitani, and Yamato (1996).

⁷ Implicit in F is some formulation of effectivity. In Section 4, we consider the core correspondences derived from α - and β -effectivity.

preference profile u , entails a strategic form game, say $G(u)$, and we require the equilibrium outcomes of $G(u)$ to coincide with the outcomes, $F(u)$, of the cooperative theory. Thus, as in Figure 1, a collection of games, $\{G(u)\}$, is determined, but the collection of games is generated from a *fixed* game form as preferences of individuals vary, incorporating a restriction — not present in the traditional approach — on how these games can depend on individual preferences.⁸ In sum, the game form G *implements* the correspondence F . Figure 2 depicts this implementation-theoretic approach.

Figure 2

On the one hand, the approach we propose confronts the usual incentive problems in implementation theory; and, on the other, it provides a non-cooperative foundation for the solution F . How are the two approaches connected? Returning to the primitives of the traditional approach, $\mathcal{V}$ and Ψ , we may presume that $\mathcal{V}$ is generated by *some* notion of effectivity applied to *some* class of strategic environments. If we further assume that the physical environment (the actions available to individuals and coalitions) is fixed, so that the members of this class correspond to the possible preferences of individuals, unobservable by a social planner or other outside agent, we get the top arrows of Figure 3, below. In Section 4, we provide necessary and sufficient conditions on $\mathcal{V}$, for α - and β -effectivity, under which this correspondence in fact holds.

Figure 3

Because the correspondence F maps to A , the set of payoffs to individuals at utility profile u is $u(F(u))$. At the same time, applying a given notion of effectivity, u induces a coalitional function V^u . Depending on the definition of F , the structure of the physical environment, and the type of effectivity

⁸ On a positive note, we have also added a source of variability that is not available in the traditional approach: because G does not necessarily depend on u through the coalitional function, we may have $G(u) \neq G(u')$ though V and V' , the coalitional functions corresponding to u and u' , are equal.

employed, the payoffs $u(F(u))$ may or may not match the payoffs $\Psi(V^u)$ of the cooperative solution of interest, i.e., Figure 3 may or may not commute. When it does, equilibrium payoffs from the game $G(u)$ agree with the predictions of Ψ applied to the coalitional function V^u , and the implementation-theoretic non-cooperative foundation extends to the original cooperative solution. In Section 6, we give conditions on the physical environment sufficient for payoff consistency of the core, under α - and β -effectivity, ensuring commutativity of Figure 3 for the core.

2.2 Literature Review.

There is an extensive literature on non-cooperative foundations, following Nash's (1950, 1953) work on bargaining. Hart and Mas-Colell (1994) provide non-cooperative foundations for various cooperative solutions. Starting with an arbitrary coalitional function, they develop a general bargaining model that provides a non-cooperative foundation of the n -person Nash bargaining solution in the pure bargaining case and for the Shapley value in the transferable utility (TU) case. In general, as the probability of breakdown in their bargaining process goes to zero, the stationary subgame perfect equilibrium payoffs of their game converge to the consistent values of Maschler and Owen (1992). Bossert and Tan (1995) develop a multi-stage demand game which yields the egalitarian bargaining solution at every Nash equilibrium.

Recently, the non-cooperative foundations of the core have received substantial attention. Chatterjee, Dutta, Ray, and Sengupta (1993) provide, among other things, a partial foundation for the core of a strictly super-additive TU coalitional function, showing that the limiting payoffs of a sequence of "no delay" stationary subgame perfect equilibria converge to a core payoff. Perry and Reny (1994) consider the core of an arbitrary TU coalitional function, and provide a non-cooperative foundation that formalizes the usual story accompanying the core: negotiation takes place in real time, individuals may make feasible proposals to coalitions, and they may accept a proposal currently on the table. They show that every stationary subgame perfect equilibrium of their game leads to payoffs in the core, and, for balanced TU coalitional functions, every core payoff is supported by a stationary subgame perfect equilibrium. Given a strictly convex TU coalitional function, Serrano (1995) constructs a class of games with subgame perfect equilibrium payoffs matching the core payoffs. Moldovanu and Winter (1995) and Serrano and Vohra (1996, Theorem 2) extend non-cooperative foundations of the core to general coalitional functions. Lagunoff (1994) maintains generality with respect to the non-cooperative game producing core outcomes. Given a coalitional function,⁹ he defines a class of games, each of which yields the core outcomes as subgame perfect equilibria. Okada and Winter (1995) propose axioms isolating a class of non-cooperative games *and* equilibrium concepts (refinements of stationary subgame perfect equilibrium) that produce core payoffs given any super-additive, totally balanced TU coalitional

⁹ Lagunoff's formulation of the coalitional function gives the outcomes (rather than payoffs) achievable by the coalitions. It is most easily interpreted in the context of an economic environment.

function.¹⁰

Several papers have contributed non-cooperative foundations, in the sense we have proposed, by formulating cooperative solutions as social choice correspondences and implementing them. These papers follow Jackson (1992) in using especially simple mechanisms. Moulin (1984) defines the Kalai-Smorodinsky bargaining solution as a social choice correspondence and implements it in subgame perfect equilibrium. Howard (1992) implements the Nash bargaining solution in subgame perfect equilibrium. Conley and Wilkie (1995) define the Nash extension solution for two-person bargaining problems and implement it in subgame perfect equilibrium. Miyagawa (1997) defines the normalized utilitarian bargaining solution as a social choice correspondence and implements it in subgame perfect equilibrium. Einy and Wettstein (1993) implement the core and bargaining set of a private good exchange economy in subgame perfect equilibrium, and Vohra and Serrano (1996, Theorem 1) subgame perfect implement the core of a private ownership economy (a generalization of private and public good production economies).¹¹ The results of Section 5 generalize the core implementation results of the latter two papers by considering general environments (possibly exhibiting externalities) and addressing the issue of effectivity.

3 General Framework.

In this section, we impose on the physical environment the minimal structure necessary to discuss coalition formation, generalizing many familiar types of economic models. We then extend the conventional notions of effectivity to our general class of environments, and we define two core social choice correspondences. The analysis takes as given a society, denoted N , consisting of individuals $i = 1, \dots, n$. Let $\mathcal{N}$ denote the collection of non-empty subsets of N (coalitions), denoted S . The complement of S is denoted simply $-S$. A *coalitional function*, V , maps each $S \in \mathcal{N}$ to a set $V(S) \subseteq \mathbb{R}^S$ consisting of vectors $y_S = (y_i)_{i \in S}$ such that S can guarantee each member $i \in S$ a payoff of at least y_i .¹² We do not require $V(S)$ to be non-empty. A cooperative solution, generically denoted Ψ , operates on coalitional functions V , assigning a set $\Psi(V) \subseteq \mathbb{R}^N$ of payoffs (possibly empty) to the individuals. The *core* of the coalitional functions V is denoted $C(V)$ and defined as the set of vectors $y \in V(N)$ such that, for all $S \in \mathcal{N}$, there does not exist $z_S \in V(S)$ with $z_S \gg y_S$.¹³

¹⁰ Non-cooperative foundations have also been developed for other solution concepts. Harsanyi (1974) develops a non-cooperative bargaining game that yields stable sets as outcomes; Selten (1981) develops a model of bargaining that produces the “semi-stable” demand vectors; Gul (1989) provides a non-cooperative foundation of the Shapley value in terms of a dynamic matching game; and Serrano (1993) constructs a non-cooperative foundation for the nucleolus.

¹¹ In related work, Wilson (1978) constructs a model of bargaining in a market context that possesses at least one equilibrium outcome in the core, but there may be other (Pareto efficient) equilibria as well. Kalai, Postlewaite, and Roberts (1979) implement the core of an economy but in strong Nash equilibrium, so their approach is not entirely non-cooperative. Bagnoli and Lipman (1989) implement the core of a public good economy in undominated perfect equilibrium.

¹² Formally, we treat $\mathbb{R}^S$ as the set of functions from S to $\mathbb{R}$. Thus, given disjoint coalitions S and S' , along with $y_S \in \mathbb{R}^S$ and $y_{S'} \in \mathbb{R}^{S'}$, it makes sense to write $y_{S \cup S'} = (y_i)_{i \in S \cup S'} \in \mathbb{R}^{S \cup S'}$. We write $y_N \in \mathbb{R}^N$ simply as y . These conventions apply to other vector notation in the paper as well.

¹³ For two vectors y and z in $\mathbb{R}^N$ and $S \in \mathcal{N}$, we write $y_S \gg z_S$ if y is greater than z in every component $i \in S$, and we write $y_S \geq z_S$ if y is at least as great as z in every component $i \in S$.

We suppose that each coalition S has some non-empty set A^S of conceivable *joint plans* of action, denoted a_S , and that individuals have payoffs determined by the plans eventually adopted. The set of conceivable joint plans for the coalition of the whole is denoted simply A , and elements are denoted by a . Given $a_{-S} \in A^{-S}$, let $\varphi_S(a_{-S}) \subseteq A^S$ be the set of joint plans *feasible* for S . The set of plans feasible for the grand coalition, N , is the subset φ_N of A . Call the pair (A, φ) an *environment* if the following conditions are satisfied:

- (a) for all $S \in \mathcal{N}$ and all $a_{-S} \in A^{-S}$, $\varphi_S(a_{-S}) \neq \emptyset$,
- (b) for all $S, S' \in \mathcal{N}$ with $S \cap S' = \emptyset$, $A^S \times A^{S'} \subseteq A^T$, where $T = S \cup S'$,
- (c) for all $S, S' \in \mathcal{N}$ with $S \cap S' = \emptyset$, and for all $a_S \in A^S$, $a_{S'} \in A^{S'}$, and $a_{-T} \in A^{-T}$, $\varphi_S(a_{S'}, a_{-T}) \times \varphi_{S'}(a_S, a_{-T}) \subseteq \varphi_T(a_{-T})$, where $T = S \cup S'$.

Condition (a) guarantees that coalitions always have feasible joint plans. Conditions (b) and (c) merely formalize the notion that independent action is a special case of cooperation. Note that the expression $\varphi_S(a_{S'}, a_{-T})$ in (c) is well-defined, since, by (b), $a_{S'} \in A^{S'}$ and $a_{-T} \in A^{-T}$ implies $(a_{S'}, a_{-T}) \in A^{-S}$.

For each individual i , let $u_i : \varphi_N \rightarrow \mathfrak{R}$ denote a utility function giving the individual's payoffs from feasible joint plans of action; let $u_S = (u_i)_{i \in S}$; and let $u = u_N$. A triple (A, φ, u) is a *strategic environment*. The following examples illustrate the flexibility of the model. Note that, as in Example 5, the set A^S may contain joint plans that are not decomposable into plans for each member of S — in the example, each i 's plans consist only of mixed strategies over the individual's pure strategies, but A^S consists of possibly correlated probability distributions on pure strategy profiles of the members.

Example 1. (Private good economy) Assuming there are k commodities, for each coalition S , let $A^S = \mathfrak{R}_+^{|S| \cdot k}$ denote the set of conceivable joint plans for S , with elements $a_S = (a_i)_{i \in S}$, each $a_i \in \mathfrak{R}^k$. Assume each individual i has an endowment $\omega_i \in \mathfrak{R}_+^k$, and let $X \subseteq \mathfrak{R}^k$ denote an aggregate production set. Define $\varphi_S(a_{-S}) = \{a_S \in A^S \mid \exists x \in X, \sum_{i \in S} a_i = x + \sum_{i \in S} \omega_i\}$, which is independent of a_{-S} . In this example, i 's utility function u_i is independent of a_{-i} and additional restrictions, such as monotonicity or continuity, may be imposed.

Example 2. (Strategic game) Let $(A^i, u_i)_{i=1}^n$ be a game in strategic form. This is a strategic environment in which, for all $S \in \mathcal{N}$, $A^S = \times_{i \in S} A^i$ and $\varphi_S \equiv A^S$ is independent of a_{-S} .¹⁴

Example 3. (Generalized game) Let $(A^i, u_i, \varphi_i)_{i=1}^n$ be a generalized game, where $\varphi_i(a_{-i}) \neq \emptyset$ is the set of strategies available to i , when the other individuals use strategies a_{-i} . This can be viewed as a strategic environment, setting $A^S = \times_{i \in S} A^i$, provided that, for all $S \in \mathcal{N}$ and all $a_{-S} \in A^{-S}$, $\varphi_S(a_{-S}) \equiv \{a_S \in A^S \mid \forall i \in S, a_i \in \varphi_i(a_{-i})\} \neq \emptyset$. In words, for all conceivable joint plans of action for non-members of S , there must be actions for members of S that are individually feasible for each member.

¹⁴ Note that a_S may be an agreement among the members of S as to how to play an extensive form game. When binding commitments are possible, as we implicitly assume, the temporal aspects of a strategic situation may be suppressed.

Example 4. (Splitting a pie) Suppose the individuals must allocate a fixed amount of a transferable good. Here, a plan for i consists of the fraction of the pie that i intends to consume, so $A^i = [0, 1]$. Define $A^S = \times_{i \in S} A^i$ and $\varphi_S(a_{-S}) = \{a_S \in A^S \mid \sum_{i \in S} a_i \leq \max\{0, 1 - \sum_{j \notin S} a_j\}\}$. Utility functions may take any form, though monotonicity is a usual requirement. Thus, “splitting the pie” environments differ from private good economies in that coalitions may feasibly claim any good left by non-members.

Example 5. (Game with correlated strategies) Let X^i be the set of pure strategies of player i , let $X^S = \times_{i \in S} X^i$, and let $U_i : X^N \rightarrow \mathbb{R}$ be individual i ’s (measurable) utility function. Define $A^S = \Delta(X^S)$, the set of probability measures on X^S , and $\varphi_S \equiv A^S$. Given $S, S' \in \mathcal{N}$ with $S \cap S' = \emptyset$, $a_S \in A^S$, and $a_{S'} \in A^{S'}$, associate $(a_S, a_{S'})$ with the product probability measure $a_S \times a_{S'}$ on $X^{S \cup S'}$. Thus, conditions (b) and (c) are satisfied. Given $\mu \in \Delta(X^N)$, the payoff to i is $u_i(\mu) = \int_{X^N} U_i(x) d\mu$.

Given a strategic environment (A, φ, u) , there are many ways to define a corresponding coalitional function. We extend the α - and β -representations, defined by Aumann and Peleg (1961) for strategic games, to arbitrary strategic environments. The two representations differ in their outlooks on coalitional power: the α -representation embodies a pessimistic view while the β -representation embodies a more optimistic one. For each coalition S , the joint plans that are always feasible, denoted $A_\alpha^S = \bigcap_{a_{-S} \in A^{-S}} \varphi_S(a_{-S})$, will play an important role in our analysis. Note that condition (c) implies that, for all S and S' with $S \cap S' = \emptyset$, $A_\alpha^S \times A_\alpha^{S'} \subseteq A_\alpha^{S \cup S'}$.

We extend Aumann and Peleg’s definitions as follows. The α -representation of (A, φ, u) , denoted V_α^u , is defined as

$$V_\alpha^u(S) = \{y_S \in \mathbb{R}^S \mid (\exists a_S \in A_\alpha^S)(\forall a_{-S} \in \varphi_{-S}(a_S))(u_S(a_S, a_{-S}) \geq y_S)\},$$

for all $S \in \mathcal{N}$.¹⁵ Thus, a coalition S can guarantee payoffs y_S for its members in the α -sense if there is a joint plan a_S for its members that is always feasible for S and such that, for every feasible joint plan a_{-S} of non-members, the joint plan $a = (a_S, a_{-S})$ gives each member i of S a payoff of at least y_i . The β -representation of (A, φ, u) , denoted V_β^u , is defined as

$$V_\beta^u(S) = \{y_S \in \mathbb{R}^S \mid (\forall a_{-S} \in A_\alpha^{-S})(\exists a_S \in \varphi_S(a_{-S}))(u_S(a_S, a_{-S}) \geq y_S)\}$$

for all $S \in \mathcal{N}$. A coalition S can guarantee payoffs of at least y_S in the β -sense if, no matter which joint plan is decided on by non-members, there is a feasible response for S , constrained by feasibility of a_{-S} , that delivers at least y_S . Clearly, a guarantee in the β -sense is weaker than one in the α -sense.¹⁶

The next example illustrates these constructions in the context of a strategic game.

Example 6. (Effectivity) Suppose $n = 2$ and the environment is such that each individual has two feasible plans of action, independent of the other’s plans: individual 1’s available actions are $\{U, C, D\}$

¹⁵ Note that $(a_S, a_{-S}) \in \varphi_N$ in the expression for $V_\alpha^u(S)$, so $u_S(a_S, a_{-S})$ is well-defined.

¹⁶ It is easily verified that $V_\alpha^u = V_\beta^u$ for private good economies. Generally, $V_\alpha^u(S) \subseteq V_\beta^u(S)$.

and individual 2's are $\{L, M, R\}$. Thus, nine joint plans are feasible for the coalition of the whole. Consider the following two profiles of utility functions.

$$\begin{array}{c}
 u = (u_1, u_2) \\
 \begin{array}{ccc}
 & L & M & R \\
 U & \left(-\frac{2}{3}, \frac{2}{3} \right) & (0, 0) & \left(-\frac{1}{3}, \frac{1}{3} \right) \\
 C & (-1, 1) & (-1, 1) & (1, -1) \\
 D & (-2, -2) & (-1, 1) & (-1, 1)
 \end{array}
 \end{array}
 \qquad
 \begin{array}{c}
 u' = (u'_1, u'_2) \\
 \begin{array}{ccc}
 & L & M & R \\
 U & \left(\frac{2}{3}, -\frac{2}{3} \right) & (1, -1) & (-2, -2) \\
 C & (0, 0) & (1, -1) & (1, -1) \\
 D & \left(\frac{1}{3}, -\frac{1}{3} \right) & (-1, 1) & (1, -1)
 \end{array}
 \end{array}$$

Here, $V_\alpha^u(\{1\}) = (-\infty, -\frac{2}{3}]$, $V_\alpha^u(\{2\}) = (-\infty, 0]$, and $V_\alpha^u(\{1, 2\})$ is the set of payoff vectors dominated by elements of the lefthand matrix; $V_\alpha^{u'}(\{1\}) = (-\infty, 0]$, $V_\alpha^{u'}(\{2\}) = (-\infty, -\frac{2}{3}]$, and $V_\alpha^{u'}(\{1, 2\})$ is the set of payoff vectors dominated by elements of the righthand matrix. Turning to β -effectivity, $V_\beta^u(\{1\}) = (-\infty, -\frac{2}{3}]$, $V_\beta^u(\{2\}) = (-\infty, \frac{2}{3}]$, and $V_\beta^u(\{1, 2\}) = V_\alpha^u(\{1, 2\})$; $V_\beta^{u'}(\{1\}) = (-\infty, \frac{2}{3}]$, $V_\beta^{u'}(\{2\}) = (-\infty, -\frac{2}{3}]$, and $V_\beta^{u'}(\{1, 2\}) = V_\alpha^{u'}(\{1, 2\})$.

Given an environment (A, φ) , we reformulate Ψ as a social choice correspondence, generically denoted F , which maps profiles u to subsets, $F(u)$, of feasible joint plans in A . For example, in the spirit of α -effectivity, we define the α -core social choice correspondence, F_α , as

$$F_\alpha(u) = \{a \in \varphi_N \mid (\nexists S \in \mathcal{N}, \tilde{a}_S \in A_\cap^S)(\forall \tilde{a}_{-S} \in \varphi_{-S}(\tilde{a}_S))(u_S(\tilde{a}_S, \tilde{a}_{-S}) \gg u_S(a))\}$$

A joint plan for the coalition of the whole is in the α -core if no coalition has a joint plan that is always feasible and yields higher payoffs to its members, no matter how non-members react to that plan. The β -core social choice correspondence, F_β , is defined as

$$F_\beta(u) = \{a \in \varphi_N \mid (\nexists S \in \mathcal{N})(\forall \tilde{a}_{-S} \in A_\cap^{-S})(\exists \tilde{a}_S \in \varphi_S(\tilde{a}_{-S}))(u_S(\tilde{a}_S, \tilde{a}_{-S}) \gg u_S(a))\}$$

A joint plan is in the β -core if there is no coalition that, whatever the plans of non-members, can devise a joint plan in response that yields higher payoffs to its members.

In Example 6, it is easily verified that $F_\alpha(u) = \{(U, L), (U, M), (U, R)\}$, $F_\alpha(u') = \{(U, L), (C, L), (D, L)\}$, and $F_\beta(u) = F_\beta(u') = \{(U, L)\}$. Generally, $F_\beta(u) \subseteq F_\alpha(u)$, and for private good economies, $F_\alpha(u) = F_\beta(u)$. In Section 6, we consider the relationship between core payoffs, given by C , and payoffs from core plans, given by F_α or F_β .

Alternatively, we could have defined these social choice correspondences in terms of coalitional functions derived using the appropriate notion of effectivity, that is, we could have defined $F_\alpha(u)$ as the joint plans $a \in \varphi_N$ for which there does not exist $S \in \mathcal{N}$ and $y_S \in V_\alpha^u(S)$ such that $y_S \gg u_S(a)$; and we could have defined $F_\beta(u)$ as those $a \in \varphi_N$ for which there does not exist $S \in \mathcal{N}$ and $y_S \in V_\beta^u(S)$ such that $y_S \gg u_S(a)$. Example 10 in Section 6 demonstrates that these alternative definitions can actually produce different sets of joint plans, though, under typical regularity conditions, they are equivalent to the definitions above.

4 Supporting Coalitional Functions.

In this section, we fully characterize the coalitional functions arising from strategic environments, complementing the results of Shapley and Shubik (1969), Billera and Bixby (1973, 1974), Billera (1974), and Mas-Colell (1975) on the representation of market games. The main result of this line of work is that, under certain ancillary assumptions, a coalitional function arises from (is “supported by”) a private good economic environment if and only if it is totally balanced. Sprumont (1995) considers the coalitional functions arising from public good economies.¹⁷ Though the literature begins with the game-theoretic structure of private and public good economies, the topic of our analysis is, in a sense, more fundamental: When does a coalitional function represent, in the α - or β -senses, a strategic environment? The answer to this question gives us bounds on the domain of any cooperative theory, and, as discussed in Section 2, it helps clarify the connection between implementation-theoretic non-cooperative foundations and traditional ones.

We say a strategic environment (A, φ, u) α -*supports* a coalitional function V if $V = V_\alpha^u$. Thus, V is α -supportable if and only if it represents (in the α -sense) the opportunities for coalitional action in some strategic environment. Similarly, a strategic environment β -*supports* V if $V = V_\beta^u$. The necessary and sufficient conditions for there to exist a strategic environment supporting a given coalitional function are very weak. For example, comprehensiveness, super-additivity, effectiveness, and the weak projection property (defined below) fully characterize the coalitional functions representing a strategic environment in the α -sense.

In fact, our results are stronger than this in two ways. First, our proofs only rely on the structure of strategic games: any coalitional function supported by a strategic environment (generalized game, private good economy, etc.) is supported by a strategic game. Second, we state our results for arbitrary collections of coalitional functions: if each member of the collection satisfies our necessary and sufficient conditions, a *single* environment can be constructed to support the entire collection as individual utility functions are varied.

Several conditions on coalitional functions are immediately necessary for β -supportability: effectiveness, comprehensiveness, and the weak projection property. We say V is *effective* if $V(N) \neq \emptyset$, and it is *comprehensive* if, for all $S \in \mathcal{N}$, all $y_S \in V(S)$, and all $z_S \in \mathbb{R}^S$, $y_S \geq z_S$ implies $z_S \in V(S)$. Effectiveness is often incorporated into the definition of a coalitional function, and comprehensiveness merely captures the notion that if a coalition S can guarantee its members at least y_S and z_S is no higher in any component, then S can guarantee its members at least z_S . V satisfies the *weak projection property* if, for all $S \in \mathcal{N}$ and all $y_S \in V(S)$, there exists $y_{-S} \in \mathbb{R}^{-S}$ such that $y = (y_S, y_{-S}) \in V(N)$. This condition requires that, if a coalition S can guarantee its members payoffs of at least y_S , then the coalition N can also guarantee those individuals at least y_S .

An additional condition, super-additivity, is necessary for α -supportability. V is *super-additive* if, for

¹⁷ See Sprumont (1995) for additional references on this topic.

all $S, S' \in \mathcal{N}$ with $S \cap S' = \emptyset$, $y_S \in V(S)$ and $y_{S'} \in V(S')$ implies $y_{S \cup S'} \in V(S \cup S')$. Super-additivity is not generally necessary for β -supportability, though it follows if V is β -supported by a game with correlated strategies.

An environment (A, φ) α -supports a collection $\mathcal{V}$ of coalitional functions if, for each $V \in \mathcal{V}$, there is a profile u of utility functions such that V is α -supported by (A, φ, u) . β -support is defined similarly. If $\mathcal{V}$ is α -supportable then, clearly, each $V \in \mathcal{V}$ must satisfy effectiveness, comprehensiveness, the weak projection property, and super-additivity. Theorem 1 shows that these conditions characterize the α -supportable collections of coalitional functions.

Theorem 1 *A collection $\mathcal{V}$ of coalitional functions is α -supportable if and only if each $V \in \mathcal{V}$ satisfies effectiveness, comprehensiveness, the weak projection property, and super-additivity.*

Proof: First we show that the α -representation, V_α^u , of a strategic environment (A, φ, u) satisfies these four conditions. Effectiveness follows from (a). Comprehensiveness follows directly from the definition of V_α^u . If $y_S \in V_\alpha^u(S)$ then there exists $a_S \in A_\alpha^S$ such that, for all $a_{-S} \in \varphi_{-S}(a_S)$, $u_S(a) \geq y_S$. Pick some $a_{-S} \in \varphi_{-S}(a_S)$, which is non-empty by condition (a). By condition (c), $(a_S, a_{-S}) \in \varphi_N$. Note that $u(a) \geq (y_S, u_{-S}(a))$. Then $(y_S, u_{-S}(a)) \in V_\alpha^u(N)$, so the weak projection property is satisfied. Checking super-additivity, take S and S' such that $S \cap S' = \emptyset$, and take $y_S \in V_\alpha^u(S)$ and $y_{S'} \in V_\alpha^u(S')$. Let $T = S \cup S'$ and $y_T = (y_S, y_{S'})$. There exists $a'_S \in A_\alpha^S$ such that, for all $a_{-S} \in \varphi_{-S}(a'_S)$, $u_S(a'_S, a_{-S}) \geq y_S$, and similarly for S' . By condition (c), $a'_T = (a'_S, a'_{S'}) \in A_\alpha^T$. Take any $a'_{-T} \in \varphi_{-T}(a'_T)$, and note that $u_T(a') \geq y_T$, as desired.

Next, we show the sufficiency of these four conditions by constructing a strategic game that α -supports $\mathcal{V}$. Let $A^i = \{(S, w, k) \mid S \subseteq N, w \in [0, 1], k \in Z_+\}$, where Z_+ denotes the non-negative integers, with representative element $a_i = (a_i^1, a_i^2, a_i^3)$. For all $V \in \mathcal{V}$ and all S with $V(S) \neq \emptyset$, let $h^{V,S} = (h_i^{V,S})_{i \in S}$ be a function from $[0, 1]$ onto $V(S)$. Given $a = (a_1, \dots, a_n)$, say that S forms at a if $V(S) \neq \emptyset$ and there exists $w \in [0, 1]$ such that, for all $i \in S$, $(a_i^1, a_i^2) = (S, w)$. Note that an individual may belong to at most one coalition forming at a . If S forms at a and $i \in S$, define $u_i(a) = h_i^{V,S}(w)$. Let R be the collection of individuals who belong to no coalition forming at a , and set $S^* = N \setminus R$. By super-additivity, $u_{S^*}(a) \in V(S^*)$. Then let $z_R(a)$ be any R -tuple such that $(u_{S^*}(a), z_R(a)) \in V(N)$, the existence of which is guaranteed by the weak projection property. Let $\bar{a}^3 = \max_i a_i^3$ and, for $i \in R$, define $u_i(a) = \min\{z_i(a), -\bar{a}^3\}$, completing the description of the strategic game. By comprehensiveness, $u(a) \in V(N)$.

Clearly, for all S , $V(S) \subseteq V_\alpha^u(S)$, since S can form and guarantee any payoff in $V(S)$. To see the opposite inclusion, consider a proper coalition $S \neq N$. The coalition S can protect its members from arbitrarily low payoffs only by forming, or by partitioning itself into smaller coalitions, each of which forms. In the first case, the coalition can achieve only payoffs in $V(S)$, and in the second case, this is true by super-additivity. The inclusion holds for N since, for all $a \in A$, $u(a) \in V(N)$. ■

Super-additivity drops out of the conditions characterizing β -supportability.

Theorem 2 A collection $\mathcal{V}$ of coalitional functions is β -supportable if and only if each $V \in \mathcal{V}$ satisfies effectiveness, comprehensiveness, and the weak projection property.

Proof: We first show that the β -representation, V_β^u , of a strategic environment (A, φ, u) satisfies these three conditions. Effectiveness and comprehensiveness again follow directly. If $y_S \in V_\beta^u(S)$ then, for all $a_{-S} \in A_{-S}^-$, there exists $a_S \in \varphi_S(a_{-S})$ such that $u_S(a) \geq y_S$. Using condition (a), pick such an a_{-S} and $a_S \in \varphi_S(a_{-S})$. By condition (b), $(a_S, a_{-S}) \in \varphi_N$. Note that $u(a) \geq (y_S, u_{-S}(a))$, so $(y_S, u_{-S}(a)) \in V_\beta^u(N)$ and the weak projection property is satisfied.

Next, we show the sufficiency of these three conditions by constructing a strategic game that β -supports $\mathcal{V}$. Let $A^i = \{(S, w, k) \mid S \subseteq N, w \in [0, 1], k \in Z_+\}$, with representative element $a_i = (a_i^1, a_i^2, a_i^3)$. For all $V \in \mathcal{V}$ and all S with $V(S) \neq \emptyset$, let $h^{V,S} = (h_i^{V,S})_{i \in S}$ be a function from $[0, 1]$ onto $V(S)$. Given $a = (a_1, \dots, a_n)$, say that S forms at a if: (i) $V(S) \neq \emptyset$, (ii) there exists $w \in [0, 1]$ such that, for all $i \in S$, $(a_i^1, a_i^2) = (S, w)$, and (iii) $\max_{i \in S} a_i^3 > \max_{i \notin S} a_i^3$. Note that at most one coalition can form at a . If S forms at a and $i \in S$, define $u_i(a) = h_i^{V,S}(w)$. Let $R = N \setminus S$, and let $z_R(a)$ be any R -tuple such that $(u_S(a), z_R(a)) \in V(N)$, the existence of which is guaranteed by the weak projection property. Let $\bar{a}^3 = \max_i a_i^3$ and, for $i \in R$, define $u_i(a) = \min\{z_i(a), -\bar{a}^3\}$, completing the description of the strategic game. By comprehensiveness, $u(a) \in V(N)$.

Clearly, for all S , $V(S) \subseteq V_\beta^u(S)$, since given a_{-S} , S can form and secure any payoff in $V(S)$. To see the opposite inclusion, consider a proper coalition $S \neq N$. The coalition S can avoid arbitrarily low payoffs for its members only by forming, in which case the coalition can achieve only payoffs in $V(S)$. The inclusion holds for N since, for all $a \in A$, $u(a) \in V(N)$. $\blacksquare$

It is straightforward to check that, if $V(S) \neq \emptyset$ for all $S \in \mathcal{N}$, super-additivity implies the weak projection property, allowing a simpler statement of Theorem 1. Without the non-emptiness condition, however, the implication need not hold.

A condition stronger than the weak projection property is the *projection property*: if $S \subseteq S'$ and $y_S \in V(S)$ then there exists $z_{S'} \in V(S')$ with $z_S = y_S$. Again, if $V(S)$ is always non-empty then superadditivity implies the projection property. However, this condition is not necessary for supportability in either sense. The next example demonstrates this for the case of β -supportability, maintaining non-emptiness of $V(S)$ for all coalitions.

Example 7. (Projection property) Let $n = 3$ and consider the following strategic game: for each individual i , $A^i = \mathbb{R}$; let $i(a)$ denote the agent with the highest action (ties going to lower indexed agents) at joint plan a , and define

$$u_i(a) = \begin{cases} a_i & \text{if } i = i(a) \\ -\sum_{j=1}^3 |a_j| & \text{else,} \end{cases}$$

for all $a \in A$. Note that, for each i , $V_\beta^u(\{i\}) = \mathbb{R}$. Now, set $S = \{1\}$ and set $S' = \{1, 2\}$. Consider any $(y_1, y_2) \in \mathbb{R}^2$ and the action $a_3 = \max\{|y_1|, |y_2|\} + 1$ for individual 3. The only way S' can secure at

least y_1 for individual 1 is by having him take an action $a_1 \geq a_3$. But then individual 2 receives utility strictly less than y_2 . Therefore, $V_\beta^u(S') = \emptyset$, violating the projection property.

It may seem counterintuitive that both individuals 1 and 2 in the example can, in the β -sense, guarantee any payoff for themselves separately, while together they are apparently powerless. The problem is that individual 1 can guarantee himself a high payoff only by making 2 worse off. And the extent to which 2 must be put out may be arbitrarily great, depending on the action taken by 3.

5 Implementing the Core Social Choice Correspondences

In this section, we fix an arbitrary strategic environment (A, φ, u) , imposing no structure on (A, φ) and placing no restrictions on individual utility functions. We then construct procedural rules governing negotiation that, under the hypothesis of subgame perfect equilibrium, induce behavior resembling the informal story accompanying the core. A common approach in the literature is to consider sequences of games with the sequence described in terms of time between moves going to zero, probability of continuation going to one, or the rate of time discounting between periods going to zero. An alternative is to work directly in continuous time, as do Perry and Reny (1994). Both lines of approach are attempts to provide realistic models of interaction in real time. We choose the latter alternative and construct an extensive game form, denoted G_α , with subgame perfect equilibrium joint plans $F_\alpha(u)$.¹⁸ A minor alteration of the model (regarding the timing of punishments) produces subgame perfect equilibrium joint plans $F_\beta(u)$. Thus, by implementing the α - and β -core social choice correspondences,¹⁹ we provide an implementation-theoretic non-cooperative foundation for the core.

Informally, we suppose that individuals may either agree to a joint plan of action for the grand coalition or propose an alternative joint plan to a coalition. At most one proposal (in effect, the earliest) is considered by the members of the proposed coalition, any of whom can veto it. If the proposal passes, the members are committed to a joint plan of action, and the remaining individuals decide how to respond. Our analysis highlights a well-known difficulty of the α - and β -cores: supporting core outcomes in equilibrium may require punishment of deviating coalitions that is harmful to the punishers themselves. We can circumvent this problem when there are three or more individuals by assuming punishments are decided by near-unanimity vote; as we show in Example 8, however, the problem is critical if there are only two individuals.

Formally, define G_α as follows. At time $t = 0$, each individual i announces a joint plan $a^i \in A$ for the grand coalition and may make a proposal $(S^i, \tilde{a}_S^i, t_S^i, t^*)$, where $i \in S^i$, $\tilde{a}_S^i \in A_\Omega^S$, and $t_S^i = (t_j^i)_{j \in S}$ associates a time $t_j^i \in (0, 1]$ with each member j of S^i . The time t_j^i is j 's assigned time to vote on i 's

¹⁸ Coherency problems may arise in continuous time models in relating outcomes to strategies. It can be verified that strategy profiles determine outcomes unambiguously in our model, so that these problems do not arise. The issues are discussed at length in Bergin (1992).

¹⁹ See Moore and Repullo (1988) and Abreu and Sen (1990) for general treatments of subgame perfect implementation.

proposal, and $t^* \in [0, 1]$ is the time at which the proposal is put to the grand coalition. A proposal must schedule individuals to vote sequentially, so $t_j^i \neq t_k^i$ for distinct $j, k \in S^i$, and it must satisfy $\max_{j \in S} \{t_j^i\} < t^*$. Thus, a proposal consists of a coalition, a joint plan for the coalition, and a vote schedule.

If no individual makes a proposal, negotiation terminates at $t = 0$ with an outcome determined as follows: it is a if, for at least $n - 1$ individuals i , $a^i = a$; in the absence of such agreement, the outcome is an exogenously determined status quo point $\bar{a} \in A$. Let $g(a^1, \dots, a^n)$ denote the outcome associated with $(a^1, \dots, a^n)$ according to this default rule.

If at least one individual makes a proposal and j is the individual with the earliest scheduled vote, the members of S^j vote sequentially on j 's proposal. If a member accepts the proposal, the vote passes to the next in line. If any member rejects the proposal, negotiation terminates with default outcome $g(a^1, \dots, a^n)$. In case all the members of S^j accept j 's proposal, they are committed to the joint plan $\tilde{a}_{S^j}^j$. Proposals scheduled after j 's are not considered. (If two or more individuals make a proposal and the individual who schedules the earliest vote is not uniquely determined, the outcome is again given by the default rule.)

Following a successful proposal, the actions of the remaining individuals (if there are any) are then determined in a general vote at time t^* . Each individual i announces $\tilde{a}_{-S^j}^i \in \varphi_{-S^j}(\tilde{a}_{S^j}^j)$. If at least $n - 1$ individuals agree on $\tilde{a}_{-S^j}$, negotiation terminates with outcome $(\tilde{a}_{S^j}^j, \tilde{a}_{-S^j})$. In the absence of agreement, negotiation ends with an arbitrary feasible joint plan $(\tilde{a}_{S^j}^j, \hat{a}_{-S^j}) \in \varphi_N$.

Theorem 3 *Assume $n \geq 3$. The joint plans determined by the subgame perfect equilibria of G_α are exactly $F_\alpha(u)$.*

Proof: First, observe that each $a \in F_\alpha(u)$ is a subgame perfect equilibrium outcome at u supported by the following strategy profile. At $t = 0$, each individual i announces $a^i = a$ and no proposal. In sequential voting, individuals vote no unless they have a strict preference for the proposal to pass. To define announcements in the general vote, note that, for any proposal $(S, \tilde{a}_S, t_S)$, there is some $\tilde{a}_{-S} \in \varphi_{-S}(\tilde{a}_S)$ and some $i \in S$ such that $u_i(\tilde{a}_S, \tilde{a}_{-S}) \leq u_i(a)$. In the general vote following announcements of a at $t = 0$ and a proposal by any individual, let all individuals announce $\tilde{a}_{-S}$, ensuring that every proposal will be rejected. Thus, no individual can gain by reporting $\tilde{a}^i \neq a$ at $t = 0$ or making a proposal to a coalition.

Now suppose we have a subgame perfect equilibrium strategy profile leading to an outcome $a \notin F_\alpha(u)$. Thus, there is some coalition S and some $\tilde{a}_S \in A_\cap^S$ such that, for all $\tilde{a}_{-S} \in \varphi_{-S}(\tilde{a}_S)$, $u_S(\tilde{a}_S, \tilde{a}_{-S}) \gg u_S(a)$. Pick any $i \in S$ and consider the following deviation for i : announce a^i as before and propose $(S, \tilde{a}_S, t_S)$, where t_S preempts the proposal process by scheduling the earliest vote. If all other members have accepted i 's proposal, the last individual to vote, say k , effectively chooses between $g(a^1, \dots, a^n)$ and $\tilde{a}$, where $\tilde{a}_{-S}$ is the outcome of the general vote upon passing the proposal. If $u_k(g(a^1, \dots, a^n)) < u_k(\tilde{a}_S, \tilde{a}_{-S})$ then subgame perfection demands that k accept the proposal. If this strict inequality holds for every member of S , each will accept i 's proposal and the outcome of the

deviation is $(\tilde{a}_S, \tilde{a}_{-S})$. Since $u_i(\tilde{a}_S, \tilde{a}_{-S}) > u_i(a)$, this deviation is improving for i , contradicting the supposition that the original strategy profile is a subgame perfect equilibrium.

If there is some member $j \in S$ for whom $u_j(g(a^1, \dots, a^n)) \geq u_j(\tilde{a}_S, \tilde{a}_{-S})$, the individual need not accept i 's proposal, and i 's deviation is not necessarily improving for i . But note that $u_j(g(a^1, \dots, a^n)) > u_j(a)$. In this case, j has an improving deviation: announce a^j as before, make any proposal scheduling a vote before anyone else, and reject it. This changes the outcome from a to $g(a^1, \dots, a^n)$, and contradicts our supposition. ■

A minor modification of G_α gives a non-cooperative foundation for the β -core. A proposal is now simply a pair (S^i, t_S^i) . If no individual makes a proposal or the individual who schedules the earliest vote is not uniquely determined, the outcome is as before. If j schedules the earliest vote, say at t' , then a general vote is held at time $t'/3$, where each i announces $\tilde{a}_{-S^j}^i \in A_\cap^{-S^j}$. If at least $n-1$ individuals agree on $\tilde{a}_{-S^j}$, the non-members of S^j are committed to $\tilde{a}_{-S^j}$. In the absence of agreement, they are committed to an arbitrary joint plan $\hat{a}_{-S^j} \in A_\cap^{-S^j}$. At time $2t'/3$, j announces $\tilde{a}_{S^j}^j \in \varphi_{S^j}(\tilde{a}_{-S^j})$. If no such $\tilde{a}_{S^j}^j$ exists, then the proposal is discarded and the outcome is $g(a^1, \dots, a^n)$. Voting on j 's proposal then proceeds sequentially, as before. If all members accept, negotiation terminates with members of S^j committed to the joint plan $\tilde{a}_{S^j}^j$. Otherwise, negotiation ends with $g(a^1, \dots, a^n)$.

Call this modified game form G_β . Thus, as would be expected, the non-cooperative foundations for the α - and β -cores differ with respect to the timing of punishments and a deviating coalition's commitment to a joint plan. In the case of the α -core, a deviating coalition commits to a plan before punishment is decided, and in the case of the β -core this timing is reversed.

Theorem 4 *Assume $n \geq 3$. The joint plans determined by the subgame perfect equilibria of G_β are exactly $F_\beta(u)$.*

Proof: If $a \in F_\beta(u)$ then, for every challenging coalition S , there exists $\tilde{a}_{-S}$ such that, for all $\tilde{a}_S \in \varphi_S(\tilde{a}_{-S})$ with $\tilde{a}_{-S} \in \varphi_{-S}(\tilde{a}_S)$, there is an $i \in S$ with $u_i(\tilde{a}_S, \tilde{a}_{-S}) \leq u_i(a)$. To obtain a as an equilibrium, let a be proposed by all players at $t = 0$. Following the announcement of a at $t = 0$ and a proposal to coalition S , all individuals announce $\tilde{a}_{-S}$, as defined above, in the general vote. Again, in sequential voting individuals vote no unless they have a strict preference for the proposal to pass. If $a \notin F_\beta(u)$, the argument proceeds as in the proof of Theorem 3. ■

In the discussion above, the number of individuals was assumed to be at least three. When there are just two individuals, the α -core is not generally implementable in subgame perfect equilibrium, as the next example shows.²⁰ In contrast, the distinction between the two and three individual cases is not significant in traditional non-cooperative foundations, where games are parameterized by coalitional functions, and, as evidenced by Proposition 1, difficult technical problems of incentives cannot be addressed.

²⁰ Example 8 works by demonstrating a violation of Moore and Repullo's (1990) condition μ_2 , which is necessary for Nash and subgame perfect implementation when there are only two agents.

Example 8. (α -core not implementable) Suppose $n = 2$ and the environment is such that each individual has two feasible plans of action, independent of the other's plans: individual 1's actions are $\{U, D\}$ and individual 2's are $\{L, R\}$. Thus, four joint plans are feasible for the coalition of the whole. Consider the following two profiles of utility functions.

$$\begin{array}{cc}
 u = (u_1, u_2) & u' = (u'_1, u'_2) \\
 \begin{array}{cc}
 & \begin{array}{cc} L & R \end{array} \\
 \begin{array}{c} U \\ D \end{array} & \begin{pmatrix} (0, 0) & (-1, 1) \\ (-2, 2) & (1, -1) \end{pmatrix}
 \end{array} & \begin{array}{cc}
 & \begin{array}{cc} L & R \end{array} \\
 \begin{array}{c} U \\ D \end{array} & \begin{pmatrix} (2, -2) & (0, 0) \\ (-1, 1) & (1, -1) \end{pmatrix}
 \end{array}
 \end{array}$$

It can be verified that $(U, L) \in F_\alpha(u)$ and $(U, R) \in F_\alpha(u')$. If F_α is implementable in subgame perfect equilibrium then there exists an extensive form mechanism and strategy pairs $s = (s_1, s_2)$ and $s' = (s'_1, s'_2)$ such that s is a Nash equilibrium at u with outcome (U, L) and s' is a Nash equilibrium at u' with outcome (U, R) . Since (U, L) is a Nash equilibrium outcome at u , any deviation by individual 2 leads to an outcome in $\{(U, L), (D, R)\}$. In particular, (s_1, s'_2) yields an outcome from $\{(U, L), (D, R)\}$. Similarly, since (U, R) is a Nash equilibrium outcome at u' , any deviation by individual 1 gives an outcome in $\{(D, L), (U, R)\}$. In particular, (s_1, s'_2) yields an outcome in $\{(D, L), (U, R)\}$, a contradiction.

The same argument used in Example 8 to prove that F_α is not implementable can be applied in Example 6 to show that F_α is not implementable, though F_β , which is constant, is trivially Nash implementable over the restricted domain consisting of u and u' . This demonstrates the important (but often implicit) role played by effectivity in the definition of a cooperative solution.

The next example shows that we cannot avoid Example 8 by restricting traditional non-cooperative foundations to “good” coalitional functions, ones arising from strategic environments where the core is subgame perfect implementable: implementability of the core cannot be inferred from a given collection of coalitional functions. In the example, we define an environment and two profiles of utility functions such that (1) individual utility functions generate the same coalitional functions (under α -effectivity) as in Example 8, and (2) the α -core is Nash implementable on the pair of profiles.

Example 9. (α -core is implementable) Let individual 1's actions be $\{U, C, D\}$ and individual 2's be $\{L, M, R\}$, and consider the following two profiles.

$$\begin{array}{cc}
 u = (u_1, u_2) & u' = (u'_1, u'_2) \\
 \begin{array}{ccc}
 & \begin{array}{ccc} L & M & R \end{array} \\
 \begin{array}{c} U \\ C \\ D \end{array} & \begin{pmatrix} (0, 0) & (-1, 1) & (1, -1) \\ (0, 0) & (-3, -3) & (-3, -3) \\ (-2, 2) & (-3, -3) & (-3, -3) \end{pmatrix}
 \end{array} & \begin{array}{ccc}
 & \begin{array}{ccc} L & M & R \end{array} \\
 \begin{array}{c} U \\ C \\ D \end{array} & \begin{pmatrix} (0, 0) & (0, 0) & (2, -2) \\ (1, -1) & (-3, 3) & (-3, -3) \\ (-1, 1) & (-3, -3) & (-3, -3) \end{pmatrix}
 \end{array}
 \end{array}$$

Here, $F_\alpha(u) = F_\alpha(u') = \{(U, L), (C, L), (U, M)\}$. Thus, the α -core is constant on these two profiles, and it is Nash implementable by the obvious mechanism: both individuals announce an outcome in $\{(U, L), (C, L), (U, M)\}$; if they announce the same pair then that is the outcome; and if they announce different pairs, the outcome is (D, R) .

6 Payoff Consistency of the Core.

A natural question, related to the commutativity of Figure 3 in Section 2, is the following: Do payoffs from the core plans, $u(F_\alpha(u))$, match the core payoffs, $C(V_\alpha^u)$ (and likewise for β -effectivity)? There are two ways in which the answer may be “no,” one definitional and the other more substantive. First, by convention, the coalitional function is defined to be comprehensive, so the payoff frontier of $V(N)$ may contain line segments and $C(V_\alpha^u)$ may be infinite, even if the underlying space of actions is finite — so that in general it cannot be that $u(F_\alpha(u)) = C(V_\alpha^u)$. Therefore, we need to consider the comprehensive hull of $u(F_\alpha(u))$, $L(u(F_\alpha(u)))$, where $L(Y) = \{z \in \mathbb{R}^N \mid (\exists y \in Y)(y \geq z)\}$ denotes the comprehensive hull of $Y \subseteq \mathbb{R}^N$. With this notation, our equivalence requirement is that $u(F_\alpha(u)) \subseteq C(V_\alpha^u)$ (as before) and $C(V_\alpha^u) \subseteq L(u(F_\alpha(u)))$. The second issue is technical: as the following example shows, without some rather weak regularity conditions on the strategic environment, payoff consistency may not obtain.

Example 10.A. (Payoff discontinuity) Let $n = 2$ and consider the following strategic game: $A^1 = A^2 = [0, 1]$, $u_1(a) = 1 - a_2$, $a_2 < 1$; $u_1(a) = 1$, $a_2 = 1$, and $u_2(a) = 0$. In this case, $V_\alpha^u(\{1\}) = V_\alpha^u(\{2\}) = (-\infty, 0]$ and $V_\alpha^u(1, 2) = (-\infty, 1] \times (-\infty, 0]$, so that $(0, 0) \in C(V_\alpha^u)$, although there is no $a \in A$ with $u(a) = (0, 0)$.

Example 10.B. (Non-compactness) Let $n = 3$, and consider the following strategic game: $A^1 = \{0, 1\}$, and A^2 and A^3 equal the positive integers; payoffs are given by

$$\begin{aligned} u_1(a) &= \frac{a_1}{a_2 a_3} \\ u_2(a) &= u_3(a) = -a_1 \end{aligned}$$

So, $V_\alpha^u(\{1\}) = (-\infty, 0]$ and $V_\alpha^u(\{1, 2, 3\}) = \{(y_1, y_2, y_3) \mid y_1, y_2, y_3 \leq 0 \text{ or } y_1 \leq 1, y_2, y_3 \leq 1\}$. For this example $(0, 0, 0) \in C(V_\alpha^u)$, though $(0, 0, 0) \notin u(F_\alpha(u))$: individual 1 can guarantee himself a positive payoff by choosing $a_1 = 1$.

The previous example shows that under α -effectivity, lower semi-continuity and compactness are needed for general equivalence. The next theorem asserts that they are sufficient as well.

Theorem 5 *In general, $u(F_\alpha(u)) \subseteq C(V_\alpha^u)$. Conversely, $C(V_\alpha^u) \subseteq L(u(F_\alpha(u)))$ if (i) for all $S \in \mathcal{N}$, φ_S is compact-valued, and (ii) for all $i \in N$ and all $a_i \in A^i$, $u_i(a_i, \cdot)$ is lower semi-continuous on A^{-i} .*

Proof: Take any $a \in F_\alpha(u)$. If it is not true that $u(a) \in C(V_\alpha^u)$ then there exists S and $y_S \in V_\alpha^u(S)$ such that $y_S \gg u_S(a)$. Since $y_S \in V_\alpha^u(S)$, there exists $\tilde{a}_S \in A_\cap^S$ such that, for all $\tilde{a}_{-S} \in \varphi_{-S}(\tilde{a}_S)$, $u_S(\tilde{a}_S, \tilde{a}_{-S}) \geq y_S \gg u_S(a)$. But then $a \notin F_\alpha(u)$, a contradiction.

Now take $y \in C(V_\alpha^u)$. Thus, $y \in \{z \in V_\alpha^u(N) \mid (\forall S \in \mathcal{N})(\nexists z'_S \in V_\alpha^u(S))(z'_S \gg z_S)\}$. Since $y \in V_\alpha^u(N)$, there exists $a \in \varphi_N$ such that $u(a) \geq y$. Clearly, $u(a) \in C(V_\alpha^u)$, and we claim that $a \in F_\alpha(u)$. Otherwise, there exists a coalition S and joint plan $\tilde{a}_S \in A_\cap^S$ such that, for all $\tilde{a}_{-S} \in \varphi_{-S}(\tilde{a}_S)$, $u_S(\tilde{a}_S, \tilde{a}_{-S}) \gg u_S(a)$. Note that a solution, (a_{-S}^*, j) , to

$$\min_{\tilde{a}_{-S} \in \varphi_{-S}(\tilde{a}_S), i \in S} u_i(\tilde{a}_S, \tilde{a}_{-S}) - u_i(a)$$

exists by our assumptions. Let $k = u_j(\tilde{a}_S, a_{-S}^*) - u_j(a)$, and define $y_S \in \mathbb{R}^S$ by $y_i = u_i(a) + k$ for all $i \in S$. Then $y_S \in V_\alpha^u(S)$ and $y_S \gg u_S(a)$ imply $u(a) \notin C(V_\alpha^u)$, a contradiction. $\blacksquare$

In view of the Example 10, the conditions of the theorem are tight: if one is violated, there may be core payoffs that do not correspond to any α -core plans. Equivalence for β -effectivity can be obtained under similar assumptions, though now lower hemi-continuity of the feasibility correspondence plays a crucial role.

Theorem 6 *In general, $u(F_\beta(u)) \subseteq C(V_\beta^u)$. Conversely, $C(V_\beta^u) \subseteq L(u(F_\beta(u)))$ if (i) for all $S \in \mathcal{N}$, φ_S is a lower hemi-continuous correspondence with compact values, and (ii) for all $i \in N$, u_i is upper semi-continuous on A .*

Proof: Take any $a \in F_\beta(u)$. We need to show that $u(a) \in C(V_\beta^u)$. If not, there exists $S \in \mathcal{N}$ and $y_S \in V_\beta^u(S)$ such that $y_S \gg u_S(a)$. This implies that for all $\tilde{a}_{-S} \in A_{\bar{N}}^{-S}$ there exists $\tilde{a}_S \in \varphi_S(\tilde{a}_{-S})$ such that $u_S(\tilde{a}_S, \tilde{a}_{-S}) \geq y_S \gg u_S(a)$. But then $a \notin F_\beta(u)$, a contradiction.

Now take $y \in C(V_\beta^u)$. Since $y \in V_\beta^u(N)$, there exists $a \in \varphi_N$ such that $u(a) \geq y$. Clearly, $u(a) \in C(V_\beta^u)$. We claim that $a \in F_\beta(u)$. Otherwise, there exists $S \in \mathcal{N}$ such that, for all $\tilde{a}_{-S} \in A_{\bar{N}}^{-S}$, there exists $\tilde{a}_S \in \varphi_S(\tilde{a}_{-S})$ such that $u_S(\tilde{a}_S, \tilde{a}_{-S}) \gg u_S(a)$. Suppose

$$\inf_{\tilde{a}_{-S} \in A_{\bar{N}}^{-S}} \sup_{i \in N} \sup_{\tilde{a}_S \in \varphi_S(\tilde{a}_{-S})} u_i(\tilde{a}_S, \tilde{a}_{-S}) - u_i(a) \leq 0.$$

Take any sequence $\{a_{-S}^k\}$ satisfying $a_{-S}^k \in A_{\bar{N}}^{-S}$ and

$$\sup_{\tilde{a}_S \in \varphi_S(\tilde{a}_{-S}^k)} u_i(\tilde{a}_S, \tilde{a}_{-S}^k) - u_i(a) \leq \frac{1}{k}$$

for all k . (Since N is finite, we fix i without loss of generality.) By our assumptions, this sequence has some limit point $\tilde{a}_{-S} \in A_{\bar{N}}^{-S}$. By supposition, however, there is some $\tilde{a}_S \in \varphi_S(\tilde{a}_{-S})$ such that $u_i(\tilde{a}_S, \tilde{a}_{-S}) - u_i(a) > 0$. Since φ_S is lower hemi-continuous, there is a sequence $\{a_S^k\}$ with $a_S^k \in \varphi_S(a_{-S}^k)$ for all k such that $(a_S^k, a_{-S}^k) \rightarrow (\tilde{a}_S, \tilde{a}_{-S})$. Since each u_i is upper semi-continuous, $\liminf_k u_S(a_S^k, a_{-S}^k) \geq u_i(\tilde{a}_S, \tilde{a}_{-S}) - u_i(a) > 0$, contrary to our choice of $\{a_{-S}^k\}$. Therefore, $\inf \sup u_i(\tilde{a}_S, \tilde{a}_{-S}) - u_i(a) = k$ for some $k > 0$. Defining $y_S \in \mathbb{R}^S$ by $y_i = u_i(a) + k/2$ for all $i \in S$, we have $y_S \in V_\beta^u(S)$. Then $y_S \gg u_S(a)$ implies $u(a) \notin C(V_\beta^u)$, a contradiction. $\blacksquare$

References

- [1] Abreu, D. and A. Sen (1990) "Subgame Perfect Implementation: A Necessary and Almost Sufficient Condition," *Journal of Economic Theory*, 50: 285-299.
- [2] Aumann, R. (1961) "The Core of a Cooperative Game without Side Payments," *Transactions of the American Mathematical Society*, 98: 539-552.
- [3] Aumann, R. and B. Peleg (1960) "Von Neumann-Morgenstern Solutions to Cooperative Games without Side Payments," *Bulletin of the American Mathematical Society*, 66: 173-179.

- [4] Bagnoli, M. and B. Lipman (1989) "Provision of Public Goods: Fully Implementing the Core through Private Contributions," *Review of Economic Studies*, 56: 583-601.
- [5] Bergin, J. (1992) "A Model of Strategic Behavior in Repeated Games," *Journal of Mathematical Economics*, 21:113-153.
- [6] Billera, L. (1974) "On Games without Side Payments Arising from a General Class of Markets," *Journal of Mathematical Economics*, 1: 129-139.
- [7] Billera, L. and R. Bixby (1973) "A Characterization of Polyhedral Market Games," *International Journal of Game Theory*, 2: 253-261.
- [8] Billera, L. and R. Bixby (1974) "Market Representations of n-person Games," *Bulletin of the American Mathematical Society*, 80: 522-526.
- [9] Bossert, W. and G. Tan (1995) "A Strategic Justification of the Egalitarian Bargaining Solution," *Social Choice and Welfare*, 12: 29-41.
- [10] Chatterjee, K., B. Dutta, D. Ray, and K. Sengupta (1993) "A Non-cooperative Theory of Coalitional Bargaining," *Review of Economic Studies*, 60: 463-477.
- [11] Debreu, G. (1952) "A Social Equilibrium Existence Theorem," *Proceedings of the National Academy of Sciences of the United States*, 38: 886-893.
- [12] Dutta, B., A. Sen, and R. Vohra (1995) "Nash Implementation through Elementary Mechanisms in Economic Environments," *Economic Design*, 1: 173-204.
- [13] Einy, E. and D. Wettstein, (1993) "A Non-cooperative Interpretation of Bargaining Sets and the Core," mimeo.
- [14] Gul, F. (1989) "Bargaining Foundations of Shapley Value," *Econometrica*, 57: 81-95.
- [15] Harsanyi, J. (1974) "An Equilibrium-Point Interpretation of Stable Sets and a Proposed Alternative Definition," *Management Science*, 20: 1472-1495.
- [16] Hart, S. and A. MasColell (1994) "Bargaining and Value," *Econometrica*, 64: 357-380.
- [17] Howard, J. (1992) "A Social Choice Rule and Its Implementation in Perfect Equilibrium," *Journal of Economic Theory*, 56: 142-159.
- [18] Ichiishi, T. (1993) *The Cooperative Nature of the Firm*, New York: Cambridge University Press.
- [19] Ichiishi, T. (1981) "A Social Coalitional Equilibrium Existence Lemma," *Econometrica*, 49: 369-377.
- [20] Jackson, M. (1992) "Implementation in Undominated Strategies: A Look at Bounded Mechanisms," *Review of Economic Studies*, 59: 757-776.
- [21] Jackson, M., T. Palrey, and S. Srivastava (1994) "Undominated Nash Implementation in Bounded Mechanisms," *Games and Economic Behavior*, 6: 474-501.
- [22] Kalai, E., A. Postlewaite, and J. Roberts (1979) "A Group Incentive Compatible Mechanism Yielding Core Allocations," *Journal of Economic Theory*, 20: 13-22.
- [23] Lagunoff, R. (1994) "A Simple Non-cooperative Core Story," *Games and Economic Behavior*, 7: 54-61.
- [24] Maschler, M. and G. Owen (1992) "The Consistent Shapley Value for Games without Side Payments," in R. Selten, ed., *Rational Interaction*, Berlin: Springer-Verlag.
- [25] Maskin, E. (1977) "Nash Equilibrium and Welfare Optimality," mimeo.
- [26] Miyagawa, E. (1997) "Implementation of the Normalized Utilitarian Bargaining Solution," mimeo.

- [27] Moldovanu, B. and E. Winter (1995) "Order Independent Equilibria," *Games and Economic Behavior*, 9: 21-34.
- [28] Moore, J. and R. Repullo (1988) "Subgame Perfect Implementation," *Econometrica*, 56: 1191-1220.
- [29] Moore, J. and R. Repullo (1990) "Nash Implementation: A Full Characterization," *Econometrica*, 58: 1083-1100.
- [30] Moulin, H. (1984) "Implementing the Kalai-Smorodinsky Bargaining Solution," *Journal of Economic Theory*, 33: 32-45.
- [31] Nash, J. (1950) "The Bargaining Problem," *Econometrica*, 18: 155-162.
- [32] Nash, J. (1953) "Two-person Cooperative Games," *Econometrica*, 21:128-140.
- [33] Okada, A. and E. Winter (1995) "A Non-cooperative Axiomatization of the Core," Kyoto Institute of Economic Research, DP 421.
- [34] Perry, M. and P. Reny (1994) "A Non-cooperative View of Coalition Formation and the Core," *Econometrica*, 62: 795-817.
- [35] Rosenthal, R. (1972) "Cooperative Games in Effectiveness Form," *Journal of Economic theory*, 5: 88-101.
- [36] Saijo, T., Y. Tatamitani, and T. Yamato (1996) "Toward Natural Implementation," *International Economic Review*, 37: 949-980.
- [37] Scarf, H. (1971) "On the Existence of a Cooperative Solution for a General Class of N-Person Games," *Journal of Economic theory*, 3: 169-181.
- [38] Selten, R. (1981) "A Non-cooperative Model of Characteristic Function Bargaining," in V. Bohm and H.H. Nachtkamp, ed., *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*, Bibliographisches Institut: Mannheim.
- [39] Serrano, R. (1993) "Non-cooperative Implementation of the Nucleolus: The 3-Player Case," *International Journal of Game Theory*, 22: 345-357.
- [40] Serrano, R. (1995) "A Market to Implement the Core," *Journal of Economic Theory*, 67: 285-294.
- [41] Serrano, R. (1996) "A Comment on the Nash Program and the Theory of Implementation," Universitat Pompeu Fabra, Economics Working Paper 161.
- [42] Serrano, R. and R. Vohra (1996) "Non-cooperative Implementation of the Core," *Social Choice and Welfare*, forthcoming.
- [43] Shubik, M. (1982), *Game Theory in the Social Sciences, Concepts and Solutions*, MIT Press: Cambridge, MA.
- [44] Sprumont, Y. (1995) "On the Game-theoretic Structure of Public-Good Economies," mimeo, CRDE and Université de Montréal.
- [45] Wilson, R. (1987) "Competitive Exchange," *Econometrica*, 46: 577-585.