Coalition Formation Processes with Belief Revision among Bounded-Rational Self-Interested Agents*

Fernando Tohmé
Tuomas Sandholm
Email: {tohme, sandholm}@cs.wustl.edu
Department of Computer Science
Washington University
Campus Box 1045
One Brookings Drive
St. Louis, MO 63130
Phone : (314) 935-6160
Fax : (314) 935-7302

Abstract

This paper studies coalition formation among self-interested agents that cannot make sidepayments. We show that α-core stability reduces to analyzing whether some utility profile is maximal for all agents. We also show that strategy profiles that lead to the α-core are a subset of Strong Nash equilibria. This fact carries our α-core-based stability results directly over to two other strategic solution concepts: Nash equilibrium and Coalition-Proof Nash equilibrium.

*This material is based upon work supported by the National Science Foundation under Tuomas Sandholm's CAREER Award IRI-9703122, Grant IRI-9610122, and Grant IIS-9800994. Fernando Tohmé was also supported by the Departamento de Economía of the Universidad Nacional del Sur, in Argentina and by a Fellowship from the National Research Council (CONICET) of Argentina. An early version of this paper appeared in a workshop [41]

The main focus of the paper is to analyze the dynamic process of coalition formation by explicitly modeling the costs of communication and deliberation. We describe an algorithm for sequential action choice where each agent greedily stepwise maximizes its payoff given its beliefs. Conditions are derived under which this process leads to convergence of the agents' beliefs and to a stable coalition structure. We derive these results for the case where the length of the process is exogenously restricted as well as for the case where agents can choose it.

Finally, we show that the outcome of any communication/deliberation process that leads to a stable coalition structure is Pareto-optimal for the original game which does not incorporate communication or deliberation. Conversely, any Pareto-optimal outcome can be supported by a communication/deliberation process that leads to a stable coalition structure.

1 Introduction

In many multiagent settings, self-interested agents—e.g. representing real world companies or individuals—can operate more effectively by forming coalitions and coordinating their activities within each coalition. Therefore, efficient methods for coalition formation are of key importance in multiagent systems. Coalition formation involves three activities: coalition structure generation (partitioning the agents into disjoint coalitions[1]), solving each coalition's (optimization) problem within the coalition, and dividing the value of each coalition among member agents (in case of net cost, this value may be negative).

Coalition formation among self-interested agents has been widely studied in game theory [40, 11, 2, 1, 4, 9]. The main solution concepts are geared toward payoff division among agents in ways that guarantee forms of stability of the coalition structure. Most of these solution concepts focus on the final solution, and usually do not address the dynamic process that leads to the solution. DAI work on coalition formation has introduced protocols for dynamic coalition formation, but the role of strategies and the process itself

[1]This is the usual definition of coalition structures in game theory. For a non-partitional definition see [38].

have not been included in the solution concept. Although the outcomes satisfy different forms of stability, there is often no guarantee that the process itself is stable, i.e. that individual agents would be best off by adhering to the imposed strategies [37, 39]. Also, it is usually assumed that the agents can compute the coalition values exactly [43, 37, 39], but sometimes determining the value of a coalition involves solving an intractable combinatorial optimization problem, e.g. solving how a coalition of dispatch centers should route their pooled vehicles to handle their pooled delivery tasks. Some DAI work has addressed the complexity of coalition value computation by explicitly incorporating computational actions in the solution concept [32, 33]. This allows one to game-theoretically trade off computation cost against solution quality. However, that work does not include protocols for dynamic coalition structure generation (all coalition structures were exhaustively enumerated), nor does it address belief revision.

In this paper we will analyze coalition formation processes from a normative perspective. Each agent's decision to participate in a coalition depends on strategic considerations since the parties are self-interested and evaluate possible agreements based on the advantages they can get from their membership to a given coalition. Since our approach is normative, it is necessary to identify the motivations of each agent. Utilities constitute a natural representation for the differences among agents: utilities are numerical representations of their different preferences.

In some multiagent systems agents can make sidepayments. On the other hand, in many settings it is desirable to be able to handle the interactions without sidepayments. First, the agents might not have (enough) money, or a secure mechanism for transferring money might not exist. Second, most interaction protocols (mechanisms) that use sidepayments are only guaranteed to work if every agent's utility is linear in money, which is not often the case.[2] This paper focuses on games where agents cannot make sidepayments.[3]

Previous research has mostly focused on *superadditive games* [11, 43].[4] Superadditivity means that any pair of coalitions is best off by merging into

[2]This is the case, for example, with the classic Clarke tax mechanism [6, 8, 10, 19]. See also the discussion on modulo mechanisms in [20].

[3]2-agent negotiations in MAS have already often been analyzed in the setting where agents cannot make sidepayments [27].

[4]However, algorithms for coalition structure generation in non-superadditive settings have also been presented [34, 32, 37, 38, 39, 13].

one. Classically it is argued that almost all games are superadditive because, at worst, the agents in a composite coalition can use solutions that they had when they were in separate coalitions.[5] This paper will also focus on superadditive games. However it should be pointed out that not all games need to be superadditive. Superadditivity may be violated, for example, if the setting has anti-cartel penalties or coordination overhead such as communication costs or computation costs which may increase as the number of agents in a coalition increases [32, 33]. Since the agents studied in this paper are self-interested, the appropriate solution concept is one that emphasizes the *stability* of the coalition structure. This means trying to reach agreements where no subgroup of agents is motivated to deviate from the solution. The *core* is one of such solution concept [19]. The core is usually used for games with transferable utility ([40, 11]) but it can be extended to games with nontransferable utility, for example using Aumann's notion of the α-core ([2]) which will be used in this paper as well.

A novel aspect of our approach is the role that belief formation plays in the coalition formation process. This process can be characterized as a sequence of deliberation (computation) and communication actions that the agents take in the dynamic process of coalition formation. Agents' incomplete information leads to standard solution concepts, such as Bayes-Nash equilibrium ([19]) or sequential equilibrium ([16]). Here we will be concerned with a solution concept that combines both coalition stability and incomplete information. The basic idea is that a *belief* of an agent in a given stage of the coalition formation process is a conditional probability distribution on the outcomes, given the previous steps of the process. A coalition structure obtains stability when the beliefs of the agents converge. We show that this coalition structure supports a Pareto optimal outcome. The price paid is tractability: the computation of the optimal coalition formation process can be exponential in the number of agents and in the length of the negotiation process.[6]

[5] A special case of coalition formation where agents cannot make sidepayments is the exchange economy, where parties exchange (multidimensional) endowments. Exchange economies exhibit superadditivity: the best outcome can be reached by the grand coalition consisting of all the agents [21].

[6] Since the problem addressed here is coalition formation with independent, self-interested agents with incomplete information, exponential complexity seems an almost unavoidable consequence. See [26] for a proof of the exponential complexity inherent even

As is common practice [11, 39, 43, 13, 32, 34], we start out by studying coalition formation in *characteristic function games* (CFGs). In such games, the value of each coalition is independent of nonmembers' actions.[7] In general the value of a coalition could depend on nonmembers' actions due to positive and negative externalities (interactions of the agents' solutions). Negative externalities between a coalition and nonmembers are often caused by shared resources. Once nonmembers are using the resource to a certain extent, not enough of that resource is available to agents in the coalition to carry out the planned solution at the minimum cost. Negative externalities can also be caused by conflicting goals. In satisfying their goals, nonmembers may actually move the world further from the coalition's goal state(s) [27]. Positive externalities are often caused by partially overlapping goals. In satisfying their goals, nonmembers may actually move the world closer to the coalition's goal state(s). From there the coalition can reach its goals less expensively than it could have without the actions of nonmembers. General settings with possible externalities can be modeled as *normal form games* (NFGs). CFGs are a strict subset of NFGs. However, many real-world multiagent problems happen to be CFGs [32]. Also, as we show, the claims that we will make using the convenient notation of CFGs directly carry over to NFGs.

The game theoretic context for this work is the Nash program which was originally presented in [23], and which has recently been widely adopted as the analysis method of choice in computational multiagent systems consisting of self-interested agents [31, 36, 35, 14]. The idea is that interactions should be studied by analyzing stable combinations of strategies—one for each agent. The sequential process of coalition formation is similar to the sequential process of bargaining where agents try to reach an agreement by exchanging proposals [28]. However, incomplete information introduces complications that are related to the credibility of statements and cheap talk: noncommital statements can induce a multiplicity of equilibria, called "babbling" equilibria [7]. Negotiation has been proposed as a solution to this problem. According to this idea, negotiation using credible (not necessarily

in the most efficient decentralized allocation processes.

[7]These coalition values may represent the quality of the optimal solution for each coalition's optimization problem, or they may represent the best bounded-rational value that a coalition can get given limited or costly computational resources for solving the problem [32, 33].

truthful) statements would lead to coherent agreements where the agents' beliefs are mutually consistent [22].

Finally, belief revision in this paper is based on a representation of beliefs by means of probability distributions. Bayes theorem offers a tool for updating them. This is a common point shared by our approach and both the game theoretical treatments of incomplete information ([3], [15]) and Bayesian methods in artificial intelligence ([24]).

The rest of the paper is organized as follows. Section 2 introduces the classic game theoretic framework for coalition formation among agents that cannot make sidepayments. Section 3 analyzes outcomes statically with the α-core solution concept. Section 4 shows generally that results derived under the α-core solution concept carry over directly to strategic solution concepts such as the Nash equilibrium, the Strong Nash equilibrium, and the Coalition-Proof Nash equilibrium. Section 5 introduces the dynamic coalition formation process which incorporates deliberation and communication. It shows that stability of the coalition formation process is equivalent to convergence of the agents' beliefs (for both exogenously and endogenously terminated negotiation), and that the outcome is Pareto-optimal.

2 Games and solutions

This section reviews the concept of a coalition game and an approach for defining the value of a coalition (characteristic function) in games where nonmembers' actions affect the value of the coalition, and agents cannot transfer sidepayments. We begin by defining a game.

Definition 1 *A* game $G = ((S_i)_{i=1}^n, \bar{U})$ *is defined by the set of agents* $I = \{1, \ldots, n\}$*, the set* S_i *of possible strategies for each agent* $i \in I$*, and the resulting utility profile* $\bar{U} : \prod_{i=1}^n S_i \to \Re^n$*, where for each strategy profile* $(s_1, \ldots, s_n) \in \prod_{i=1}^n S_i$*,*

$$\bar{U}(s_1, \ldots, s_n) = (u_1(s_1, \ldots, s_n), \ldots, u_n(s_1, \ldots, s_n))$$

where $u_i : \prod_{i=1}^n S_i \to \Re$ *is the utility of agent* i*.*

A solution concept defines the reasonable ways that a game can be played by self-interested agents:

Definition 2 *Given a game* $G = ((S_i)_{i=1}^n, \bar{U})$*, a* solution concept *(in pure strategies)*[8] *is a correspondence* $\gamma : G \to \prod_{i=1}^n S_i \cup \emptyset$*, and each* $s = (s_1, \ldots, s_n) \in \gamma(G)$ *is called a solution of* G*.*

The range of a correspondence includes the empty set in order to encompass games that do not have a solution of the type prescribed by the solution concept.

An example of solution concept is given by Nash equilibria: for each game G they are elements of $\gamma_N(G)$, where γ_N is the *Nash correspondence*:

Definition 3 $s = (s_1, \ldots, s_i, \ldots, s_n)$ *is in* $\gamma_N(G)$ *if for each* i *and for each* $s_i^{'} \neq s_i$*,* $u_i(s_1, \ldots, s_i^{'}, \ldots, s_n) \leq u_i(s_1, \ldots, s_i, \ldots, s_n)$*. Each such strategy profile is a* Nash equilibrium*.*

In other words, in a Nash equilibrium no agent is motivated to deviate from its strategy given that the others do not deviate.

Definition 1 characterizes games in terms of the strategies of agents and the corresponding payoffs. These games are said to be in *normal form*. The normal form is a general representation that can be used to model the fact that nonmembers' actions affect the value of the coalition [32, 9]. However, coalition formation has been mostly studied in a strict subset of normal form games—*characteristic function games*—where the value of a coalition does not depend on nonmembers' actions, and it can therefore be represented by a coalition specific characteristic function which provides a payoff for each *coalition* T (i.e. set of agents) [40, 11, 43, 39]. Characteristic functions are a desirable representation, so one would like to define such mathematical entities for normal form games. In such general games, a characteristic function can only be defined by making specific assumptions about nonmembers' strategies. In this paper we follow Aumann's classic approach of making the α*-assumption*, i.e. assuming that nonmembers pick strategies that are worst for the coalition. Each coalition can locally guarantee itself a payoff that is no less than the one prescribed by an analysis under this pessimistic assumption.[9] Later in the paper we show that the results that we obtain

[8] A pure strategy is a strategy that does not involve randomization by the agent. Definition 2 can be easily extended to mixed strategies that involves randomization, by replacing each S_i by ΔS_i, the set of probability distributions on S_i.

[9] The α-assumption may be impossibly pessimistic. A given nonmember can be assumed to pick different strategies when different coalitions are evaluated. This is in contrast with the fact that in any realization, the nonmember can only pick one strategy.

under the α-assumption carry over to strategic solution concepts (Nash equilibrium and its refinements) that can be used directly in normal form games without any assumptions about nonmembers' strategies.

In games where agents can make sidepayments to each other [37, 39, 32, 11], the characteristic function gives the sum of the payoffs of the agents in a coalition. Instead, our analysis focuses on games where agents cannot make sidepayments. In such games, the characteristic function gives a *set* of utility *vectors* that are achievable [2]. This is in order to provide the coalition with a set of alternative utility divisions among member agents. The set contains only Pareto-optimal utility vectors: no agent can be made better off without making some other agent worse off. The next definition formalizes this vector-valued characteristic function under the α-assumption.

Definition 4 *Given a game $G = ((S_i)_{i=1}^n, \bar{U})$, with $\bar{U}$ such that its components are non-transferable, we say that the* characteristic function *is*

$$v_\alpha : 2^I \rightarrow 2^{\Re^I}$$

such that for each coalition $T \subseteq I$

$$v_\alpha(T) \subseteq \Re^I$$

and $v_\alpha(T)$ is the set of optimal achievable utilities *for T.*

The α-assumption comes into play in the definition of these optimal achievable utilities:

Definition 5 *Given a game $G = ((S_i)_{i=1}^n, \bar{U})$, and a $T \subseteq I$, the* set of optimal achievable utilities *for coalition $T = \{j_1, \ldots, j_{|T|}\}$ is the set of utility profiles*

$$\bar{U}_T = (\ldots, \bar{u}_{j_1}, \ldots, \bar{u}_{j_2}, \ldots, \bar{u}_{j_{|T|}}, \ldots)$$

such that

$$\exists s \in \prod_{i=1}^n S_i, \; U(s) = \bar{U}_T$$

and

$$\nexists s^T \in \prod_{j \in T} S_j : \; \forall s^{I-T} \in \prod_{j \notin T} S_j$$

$$U(s^T, s^{I-T}) =$$

$$\bar{U}'_T = (\ldots, \bar{u}'_{j_1}, \ldots, \bar{u}'_{j_2}, \ldots, \bar{u}'_{j_{|T|}}, \ldots)$$

with, for all $j_i \in T$, $\bar{u}'_{j_i} \geq \bar{u}_{j_i}$ *and for at least one* $j^* \in T, \bar{u}'_{j*} > \bar{u}_{j*}$.

The next result of ours follows trivially:

Proposition 1 *For every game* G *in normal form,* v_α *exists.*

Proof. *Suppose that for a game* $G = ((S_i)_{i=1}^n, \bar{U})$, v_α *cannot be defined. So, for at least one coalition* T, $v_\alpha(T)$ *cannot be determined. But that means that a set of* $\bar{U}_T s$ *cannot be defined such that*

$$\nexists s^T \in \prod_{j \in T} S_j : \ \forall s^{I-T} \in \prod_{j \notin T} S_j,$$

$$U(s^T, s^{I-T}) =$$

$$\bar{U}'_T = (\ldots, \bar{u}'_{j_1}, \ldots, \bar{u}'_{j_2}, \ldots, \bar{u}'_{j_{|T|}}, \ldots)$$

and such that, for all $j_i \in T$, $\bar{u}'_{j_i} \geq \bar{u}_{j_i}$ *and for at least one* $j^* \in T, \bar{u}'_{j*} > \bar{u}_{j*}$. *Given this condition, we proceed by evaluating* $\bar{U}$ *for each* $s \in \prod_{i=1}^n S_i$. *If, by hypothesis,* $v_\alpha(T)$ *is not determinate, then* $\bar{U}$ *is not defined for every* s. *Contradiction.* $\square$

To summarize, the normal form notation emphasizes the strategic aspects of the interactions among agents. Each agent's actions are strictly individual even as a member of a coalition. Therefore, a characteristic function has to indicate, for each coalition, how much each of its members gets for joining. Moreover, the characteristic function has to indicate the best combinations of individual payoffs that the coalition can get. Proposition 1 just states that it is possible to define this kind of characteristic function for every strategic game. Once this issue is settled, we can move on to see what kinds of solutions are relevant for this type of coalitional games.

3 The α-core and superadditivity

The α-assumption gives rise to the α-*core* solution concept which defines a stability criterion for the coalition structure. The idea is that strategy profiles that do not have an optimal achievable utility are not candidates for the solution. A vector of joint strategies, is said to be *blocked* by a coalition if its members can be better off by moving to another vector:

Definition 6 *A coalition T* blocks *a strategy profile $s = (s_1, \ldots, s_n)$ if for every $j_i \in T$ there exists an $s' \in \prod_{i=1}^{n} S_i$ such that*

- $\forall j_i$, $u_{j_i}(s') \geq u_{j_i}(s)$ *and for at least one* j_i, $u_{j_i}(s') > u_{j_i}(s)$*; and*
- $(\ldots, u_{j_1}(s'), \ldots, u_{j_2}(s'), \ldots, u_{j_{|T|}}(s'), \ldots) \in v_\alpha(T)$.

The blocking relation defines a particular set of stable joint strategies, the α-*core*. The α-core is the set of joint strategies where no coalition can be formed if its members are better off changing their individual strategies, given that nonmembers pick strategies that are worst for the coalition. In other words, it is the set of joint strategies for which a stable collective agreement can be reached. Formally, the α-core correspondence is defined as follows [2]:

Definition 7 *A strategy profile $s = (s_1, \ldots, s_n)$ is in the α-core $\gamma_{\mathcal{C}_\alpha}$ if there is no coalition T that blocks s.*

As with the Nash correspondence, the α-core correspondence can be empty for some games. The following example demonstrates this.

Example 1 $G = ((S_a, S_b), \bar{U})$*, where the set of players is* $\{a, b\}$

$$S_a \;=\; S_b \;= \{c, nc\}$$

and

$$\bar{U} \;=\; \{(\langle c, nc\rangle, \langle 0, 10\rangle), (\langle c, c\rangle, \langle 5, 5\rangle), (\langle nc, c\rangle, \langle 10, 0\rangle), (\langle nc, nc\rangle, \langle 2, 2\rangle)\}$$

where $(\langle s_a, s_b\rangle, \langle u_a(s_a), u_b(s_b)\rangle)$ is the general form of the elements of $\bar{U}$. This is an instance of the prisoner's dilemma *[18]. The corresponding values of the characteristic function are:*

- $v_\alpha(\{a\}) = \{\langle 10, 0\rangle\}$
- $v_\alpha(\{b\}) = \{\langle 0, 10\rangle\}$
- $v_\alpha(\{a, b\}) = \{\langle 5, 5\rangle\}$

It is easy to see that

- $\{a\}$ *blocks* $\{\langle c, c\rangle, \langle c, nc\rangle, \langle nc, nc\rangle\}$

- *$\{b\}$ blocks $\{\langle c, c\rangle, \langle nc, c\rangle, \langle nc, nc\rangle\}$*
- *$\{a, b\}$ blocks $\{\langle nc, c\rangle, \langle c, nc\rangle, \langle nc, nc\rangle\}$*

Therefore, there is no $\langle s_a, s_b\rangle$ that is not blocked by at least one coalition. In other words, $\gamma_{\mathcal{C}_\alpha}(G)$ is empty.

Now, under what conditions does a stable coalition structure exist? In other words, what are the conditions for non-emptyness of the α-core? In the rest of this section we will show that surprisingly simple conditions are necessary and sufficient for stability. The concept of *superadditivity* will be used to develop an intuition of this phenomenon. Superadditivity implies that any two coalitions are best off merging:

Definition 8 *A game G is* superadditive *if given any two coalitions T_1, T_2, $T_1 \cap T_2 = \emptyset$, $v_\alpha(T_1) \cap v_\alpha(T_2) \subset v_\alpha(T_1 \cup T_2)$.*[10]

The following result of ours establishes an interesting property that relates characteristic functions and superadditivity: *if the intersection of the optimally achievable utilities for all the players is not empty, then the game is superadditive.* This condition on characteristic functions will be later used to discuss stability.

Proposition 2 *For a game G, if $\bigcap_{i \in I} v_\alpha(\{i\}) \neq \emptyset$ then G is superadditive.*

Proof. *We will prove this result by induction on the cardinality of coalitions:*

- *Given T_1, T_2, $T_1 \cap T_2 = \emptyset$, $|T_1| = 1$, $|T_2| = 1$, it is clear that there exist agents $i, j \in I$ such that $T_1 = \{i\}$ and $T_2 = \{j\}$. If $\bar{U}^* = (u_1^*, \ldots, u_n^*) \in \bigcap_{i \in I} v_\alpha(\{i\})$, then in particular $\bar{U}^* \in v_\alpha(T_1) \cap v_\alpha(T_2)$. Suppose $\bar{U}^* \notin v_\alpha(T_1 \cup T_2)$. Then, there exists $s \in \prod_{i=1}^n S_i$ such that $\bar{U}(s) = (\ldots, u_i(s), \ldots, u_j(s), \ldots)$ and $u_i(s) \geq u_i^*$ and $u_j(s) \geq u_j^*$ with strict inequality for one of them, say i. But then, $\bar{U}^* \notin v_\alpha(\{i\})$, contradiction. So, $v_\alpha(T_1) \cap v_\alpha(T_2) \subset v_\alpha(T_1 \cup T_2)$.*

[10] This definition (by Shubik [40]) for games without sidepayments differs technically from the definition of superadditivity for games with sidepayments [40, 11, 32, 43, 37, 39]. However, they are conceptually the same.

- *Assume that* $\bar{U}^* \in v_\alpha(T_1) \cap v_\alpha(T_2) \subset v_\alpha(T_1 \cup T_2)$, *for any pair of coalitions* T_1, T_2, $T_1 \cap T_2 = \emptyset$, $|T_1 \cup T_2| \leq k < n$. *Consider* $i \in I$, $i \notin T_1, i \notin T_2$, *and* $T_1^{'} = T_1 \cup \{i\}$. *Then* $T_1^{'} \cap T_2 = \emptyset$ *and of course* $\bar{U}^* \in v_\alpha(T_1^{'})$ *(by the inductive assumption because* $|T_1^{'}| \leq k$*), so* $\bar{U}^* \in v_\alpha(T_1^{'}) \cap v_\alpha(T_2)$. *Suppose that* $\bar{U}^* \notin v_\alpha(T_1^{'} \cup T_2)$. *Again, this means that exists* $s \in \prod_{i=1}^n S_i$ *such that* $\bar{U}(s) = (\ldots, u_{j_1}(s), \ldots, u_{j_k}(s), \ldots)$, *where* $T_1^{'} \cup T_2 = \{j_1, \ldots, j_k\}$, *and* $u_{j_i}(s) \geq u_{j_i}^*$ *for all* $j_i \in T_1^{'} \cup T_2$, *with strict inequality for one of them, say* j_{i_0}. *Suppose without loss of generality that* $j_{i_0} \in T_1^{'}$, *but then,* $\bar{U}^* \notin v_\alpha(T_1^{'})$. *Contradiction.*

So, $v_\alpha(T_1) \cap v_\alpha(T_2) \subset v_\alpha(T_1 \cup T_2)$ *for any pair of coalitions* T_1, T_2, $T_1 \cap T_2 = \emptyset$, $|T_1 \cup T_2| \leq n$. *That is,* G *is superadditive.* □

To see that superadditivity is a necessary but not a sufficient condition for $\bigcap_{i \in I} v_\alpha(\{i\}) \neq \emptyset$, recall the prisoner's dilemma of Example 1. It is a superadditive game, but $v_\alpha(\{a\})$ and $v_\alpha(\{b\})$ have no element in common.

Our next result relates the condition of the previous proposition to stability of the coalition structure (non-emptyness of the α-core):

Lemma 1 *For a game* G, $\gamma_{\mathcal{C}_\alpha}(G) \neq \emptyset$ *iff* $\bigcap_{T \in 2^I - \emptyset} v_\alpha(T) \neq \emptyset$.

Proof.

- $\rightarrow$) *If* $\gamma_{\mathcal{C}_\alpha}(G) \neq \emptyset$, *then there exists an* s^* *that is not blocked by any coalition. But then, by the definition of blocked joint strategy, it is clear that for each coalition* T, $\bar{U}(s^*) \in v_\alpha(T)$, *and therefore* $s^* \in \bigcap_{T \in 2^I - \{\{\emptyset\}} v_\alpha(T)$
- $\leftarrow$) *If* $\bigcap_{T \in 2^I - \{\emptyset\}} v_\alpha(T) \neq \emptyset$, *then there exists at least one* $\bar{U}^* \in v_\alpha(T)$ *for every possible coalition* T *and therefore an* $s \in \prod_{i=1}^n S_i$ *such that* $\bar{U}(s) = \bar{U}^*$. *So,* s *is not blocked by any coalition and thus* $s \in \gamma_{\mathcal{C}_\alpha}(G)$. □

This lemma is useful for proving the following theorem of ours. The theorem shows that to characterize the stability of the coalition structure in terms of the α-core, only the utilities and the corresponding actions of individual agents are required. One does not need to compare utilities and actions of coalitions.

Theorem 1 *For a game* G, $\gamma_{\mathcal{C}_\alpha}(G) \neq \emptyset$ *iff* $\bigcap_{i \in I} v_\alpha(\{i\}) \neq \emptyset$.

Proof.

- $\rightarrow$) *If $\gamma_{\mathcal{C}_\alpha}(G) \neq \emptyset$, there exists a $s \in \prod_{i=1}^{n} S_i$ such that no coalition blocks it. So for each coalition T, $\bar{U}(s) \in v_\alpha(T)$. In particular for all the coalitions with a single member, $T = \{i\}$. So, $\bar{U}(s) \in \bigcap_{i \in I} v_\alpha(\{i\})$.*

- $\leftarrow$) *By Lemma 1, it is enough to prove that $\bigcap_{T \in 2^I - \{\emptyset\}} v_\alpha(T) \neq \emptyset$. The proof will be by induction on the size of coalitions:*

  - *given that by hypothesis $\exists \bar{U}^* \in \bigcap_{i \in I} v_\alpha(\{i\})$, then for each $i \in I$, $\bar{U}^* \in v_\alpha(\{i\})$*
  - *assume that for each coalition T with $|T| = k < n$, $\bar{U}^* \in v_\alpha(T)$. For any $i \in I, i \notin T$ (by proposition 1):*

    $$\bar{U}^* \in v_\alpha(T) \cap v_\alpha(\{i\}) \subset v_\alpha(T \cup \{i\})$$

    So, $\bar{U}^* \in v_\alpha(T^{'})$, for any $T^{'}$ such that $|T^{'}| = k + 1$. Therefore, $\bar{U}^* \in T^{''}$ for any $T^{''} \in 2^I - \{\emptyset\}$. □

Since, for each agent i, $v_\alpha(\{i\})$ represents that agent's optimal achievable utilities, this result states that the non-emptyness of the α-core is equivalent to the existence of at least one utility vector that is maximal for every agent. This vector, say $\bar{U}$, is Pareto-optimal: there is no other $\bar{U}'$ such that for all i, $\bar{U}_i' \geq \bar{U}_i$, with strict inequality for at least one i.

It follows from Theorem 1 that in games without sidepayments, the coalition structure can be stable only if every possible pair of coalitions is best off merging ([α-core $\neq \emptyset$] $\Rightarrow$ superadditivity). This differs from games with sidepayments: there the core can be nonempty even if the game is not superadditive [32]. On the other hand—as in games with sidepayments—the coalition structure may be unstable even if every pair of coalitions is best off merging (superadditivity $\not\Rightarrow$ [α-core$\neq \emptyset$]).

These results show that the α-core is nonempty only in superadditive games. In all other games there exists at least one coalition that blocks the solution. While this result is interesting *per se*, it could perhaps be interpreted as a criticism of the α-core solution concept itself. However, with self-interested agents with non-transferable utilities, all market-like games are superadditive [21]. Market economies are natural examples of this type of games and electronic economies constitute a natural field of application for the results of this paper.

4 Relationships between axiomatic and strategic solution concepts

In this section we present some new relationships between axiomatic and strategic solution concepts. The importance of these relationships lies in the fact that they allow us to import the other results of this paper (which will be derived for the axiomatic α-core solution concept) directly to strategic solution concepts like Nash equilibrium and its refinements.

The notion of the α-core is axiomatic in that it only characterizes the outcome without a direct reference to strategic behavior. The Nash correspondence is, instead, a strategic solution concept: it is based only on the self-interested strategy choices of agents. Specifically, it analyzes what an agent's best strategy is, given the strategies of others. A strategy profile is in Nash equilibrium if every agent's strategy is a best response to the strategies of the others. Nash equilibrium does not account for the possibility that groups of agents (coalitions) can change their strategies in a coordinated manner. Aumann has introduced a strategic solution concept called the Strong Nash equilibrium to address this issue [1, 4]. A strategy profile is in Strong Nash equilibrium if no subgroup of agents is motivated to change their strategies given that others do not change:

Definition 9 *A strategy profile $s \in \prod_{i=1}^{n} S_i$, in a game G, is a* Strong Nash equilibrium *if for any $T \subseteq I$ and for all $\bar{s}^T \in \prod_{j \in T} S_j$ there exists an $i_0 \in T$ such that $u_{i_0}(s) \geq u_{i_0}(\bar{s}^T, s^{I-T})$.*

This concept gives rise to the Strong Nash correspondence, γ_{SN}, i.e. the set of Strong Nash equilibria. We show a close relationship between the Strong Nash solution concept and the α-core solution concept:

Theorem 2 *For any game G, $\gamma_{\mathcal{C}_\alpha}(G) \subseteq \gamma_{SN}(G)$.*

Proof. *Suppose that $s \in \gamma_{\mathcal{C}_\alpha}(G)$ but $s \notin \gamma_{SN}(G)$. Then, there exist an $T \subseteq I$ and an $\bar{s}^T \in \prod_{j \in T} S_j$ such that for all $j \in C$, $u_j(s) < u_j(\bar{s}^T, s^{I-T})$. That means that $\bar{U}(s)$ is not in $v_\alpha(T)$, so there is a s' such that $\bar{U}(s') \in v_\alpha(T)$ and $u_j(s') \geq u_j(s)$. Contradiction.* □

Often the Strong Nash equilibrium is too strong a solution concept, since in many games no such equilibrium exists. Recently, the *Coalition-Proof*

Nash equilibrium [4] has been suggested as a partial remedy to this problem. This solution concept requires that there is no subgroup that can make a mutually beneficial deviation (keeping the strategies of nonmembers fixed) *in a way that the deviation itself is stable according to the same criterion.* A conceptual problem with this solution concept is that the deviation may be stable within the deviating group, but the solution concept ignores the possibility that some of the agents that deviated may prefer to deviate again with agents that did not originally deviate. Furthermore, even these kinds of solutions do not exist in all games. On the other hand, in games where a solution is stable according to the α-core, the solution is also stable according to the Coalition-Proof Nash equilibrium solution concept. This is because $\gamma_{\mathcal{C}_\alpha}(G) \subseteq \gamma_{CPN}(G)$ (which follows from our result $\gamma_{\mathcal{C}_\alpha}(G) \subseteq \gamma_{SN}(G)$ and the known fact that $\gamma_{SN}(G) \subseteq \gamma_{CPN}(G)$).

We can also relate the Nash equilibrium itself to the α-core (this could alternatively be deduced from Theorem 2 and the fact that $\gamma_{SN}(G) \subseteq \gamma_N(G)$):

Theorem 3 *For any game G, $\gamma_{\mathcal{C}_\alpha}(G) \subset \gamma_N(G)$.*

Proof. *Given $s \in \gamma_{\mathcal{C}_\alpha}(G)$, we will show that s is a Nash equilibrium in pure strategies for G. Suppose not. By Theorem 1 it is enough to consider what happens with single individuals. Then, for a $i_0 \in I$, given the vector $(s_1, \ldots, s_{i_0-1}, s_{i_0+1}, \ldots, s_n)$, the* best response *for i_0 is $s^{'}_{i_0}$ with $u_{i_0}(s_1, \ldots, s^{'}_{i_0}, \ldots, s_n) > u_{i_0}(s)$. But that means that $\{i\}$ blocks s, and therefore $s \notin \gamma_{\mathcal{C}_\alpha}(G)$.Contradiction. This proves that $\gamma_{\mathcal{C}_\alpha}(G) \subseteq \gamma_N(G)$. Example 1 shows that the converse is not true: $\gamma_N(G) = \{\langle nc, nc\rangle\}$ and $\gamma_{\mathcal{C}_\alpha}(G) = \emptyset$. Therefore $\gamma_{\mathcal{C}_\alpha}(G) \subset \gamma_N(G)$.*$\square$

An implication of the results in this section is that the other results of this paper (which are derived for the α-core) carry over directly to analyses that use strategic solution concepts (Nash equilibrium, Coalition-Proof Nash equilibrium or Strong Nash equilibrium). Specifically, any solution that is stable according to the α-core is also stable according to these three solution concepts.

Another implication is that to verify that a strategy profile is in the α-core, one needs to only consider strategy profiles that are Pareto-optimal[11]

[11]The non-emptyness of the α-core is equivalent to the existence of a utility vector $\bar{U}$ which is common to all sets $v_\alpha(\{i\})$ for all agents i. By definition 4, this means that there

and in Nash equilibrium. Alternatively, one can restrict this search to Pareto-optimal Strong Nash equilibria or Pareto-optimal Coalition-Proof Nash equilibria.

5 Bounded rationality in coalition formation

In the previous section it was assumed that deliberation is costless. To relax this assumption we introduce deliberation and communication actions explicitly into the model:

Definition 10 *For each agent i in a game G, let D_i be a set of* deliberation/communication *activities that i can perform in order to choose a strategy s_i to be executed. Each $d_i \in D_i$ is associated with a $C_i(d_i)$, i.e. the cost (for i) of performing the activity d_i.*[12]

In order to avoid unnecessary complications, we assume that $C_i(\cdot)$ can be expressed in the same units as $u_i(\cdot)$. We will not assume any special structure on D_i, except the following:[13]

Definition 11 *For each agent i, we consider its process of* communication/deliberation *$\{a_i^0, a_i^1, \ldots a_i^{t_i}\}$, where $a_i^t \in D_i$, for $t = 0, 1, \ldots, (t_i - 1)$, and $a_i^{t_i} \in S_i$. If $N = max_{i \in I} t_i$, we say that the coalition structure has been formed in N steps. For any i and t such that $t_i \leq t < N$, $a_i^t = a_i^{t_i}$.*

The idea behind this definition is that the agents deliberate and exchange messages until each one decides on a strategy to follow. We also assume that this process is finite and each agent stays commited to its choice once it has reached a decision.

We use a general characterization of the communication/deliberation process without going into the details of how an action leads to another one (e.g.

does not exist a $\bar{U}'$ such that $\bar{U}_i' \geq \bar{U}_i$ for all i, with strict inequality for at least one i. That is, $\bar{U}$ is Pareto-optimal. Therefore one can restrict the search to Pareto-optimal outcomes.

[12] Implicit in this definition is the existence of another level of deliberation, assumed costless, needed for choosing deliberation/communication actions. Although strong, this assumption cannot be dropped without leading to a potential infinite regress [17, 32, 29, 5, 30, 42].

[13] Activities in D_i and strategies in S_i will be called the *actions* of agent i in the process.

how deliberation actions lead to the choice of physical actions). We say that the payoff of agent i in an N-period process is determined as follows:

Definition 12 *If the communication/deliberation process of agent i is $\hat{a}_i = \{a_i^0, \ldots, a_i^{t_i}\}$ ($a_i^{t_i} = s_i$), the payoff is*

$$\rho_i(\hat{a}_i, \ldots, \hat{a}_n) \;=\; u_i(s_1, \ldots, s_i, \ldots, s_n) - (C_i(a_i^0) + \ldots + C_i(a_i^{t_i})) - \bar{C}_i(N - t_i)$$

where $\bar{C}_i > 0$ is the waiting cost*, which is assumed constant per time unit.*

As this definition states, we assume that costs of activities are independent: if the process is $\hat{a}_i = (a_i^1, \ldots, a_i^N)$, its cost $C_i(\hat{a}_i)$ is equal to the sum of the costs of the activities, $C_i(a_i^0) + \ldots + C_i(a_i^{t_i}) + \bar{C}_i(N - t_i)$.

A new game can be defined which explicitly considers the deliberation and communication actions as part of each agent's strategy. This follows the approach of Sandholm and Lesser ([32, 33]) in the sense that such actions are explicitly incorporated into the solution concept. It differs from other DAI approaches to coalition formation where the solution concept (stability criterion) is only applied to the final outcomes [43, 37, 39].

Definition 13 *Given G, $\{D_i\}_{i=1}^n$ and $t > 0$, a new game is defined, $G^t = ((D_i^t \times S_i)_{i=1}^n, P)$, where $P : \prod_{i=1}^n (D_i^t \times S_i) \to \Re^n$ such that for each $\hat{a} \in \prod_{i=1}^n (D_i^t \times S_i)$, $P(\hat{a}) = (\rho_1(\hat{a}_1), \ldots, \rho_n(\hat{a}_n))$.*

The length of the game, t, depends on the available communication/deliberation activities and on the sequential choice of activities. For now we assume that the time limit is given (in Subsection 5.2 we will relax this assumption). To justify this, we suppose that each agent, i, has a degree of impatience, given by a maximum time, t_i, to make a final decision.

In order to maximize payoffs, our self-interested agents engage in negotiations. The final outcomes represent the result of agreements among agents. Coalitions are formed during the negotiations. A particular process generates a stable coalitional structure if the final outcome in the original game (with strategy profile space $\prod_{i=1}^n S_i$) cannot be blocked by a coalition formed in *another* process. Formally,

Definition 14 *A process $\hat{a} = \{a^0, \ldots, a^t\} \in \prod_{i=1}^n (D_i^{t-1} \times S_i)$ leads to a* stable coalition structure *if $a^t \in \prod_{i=1}^n S_i$ cannot be blocked by any coalition formed in another process $\hat{a}' = \{a^{0'}, \ldots, a^{t'}\}$.*

The relationship between stable coalition structures and the α-core is given by the following lemma.

Lemma 2 *If the process* $\hat{a} = \{a^0, \ldots, a^t\} \in \prod_{i=1}^n (D_i^{t-1} \times S_i)$ *is in the* α*-core of game* G^t *then* $\hat{a}$ *leads to a stable coalition structure.*

Proof. *Suppose that there exists a coalition* T *for which there exists an* $s \in \prod_{i=1}^n S_i$ *such that* $u_j(s) \geq u_j(a^t)$ *for* $j \in T$ *and* $u_{j^*}(s) > u_{j^*}(a^t)$ *for* $j^* \in T$*. Then* $\hat{a}' = \{a^{0'}, \ldots, \bar{s}\}$ *is a process in which* T *is formed and obtains* $\bar{s}$*, where* $\bar{s}_j = s_j$ *and* $P_j(\hat{a}') \geq P_j(\hat{a})$*, for* $j \in T$*. Contradiction because* $\hat{a}$ *is in the* α*-core of* G^t*.* □

We can easily restate the notions given in Section 2 in order to find conditions for the stability of the coalition structure. First, we define the characteristic function for G^t, v_{G^t}, via replacing the optimal achievable utilities by the optimal achievable payoffs which incorporate deliberation and communication:

Definition 15 *Given* $G^t = ((D_i^t \times S_i)_{i=1}^n, P)$*, and* $T \subseteq I$*, the set of optimal achievable payoffs for* $T = \{j_1, \ldots, j_{|T|}\}$ *is the set of* $\bar{P}$*s such that*

- *there exists* $\hat{a} \in \prod_{i=1}^n (D_i^t \times S_i)$ *and* $P(\hat{a}) = \bar{P}$*, and*
- $\nexists \hat{a} \in \prod_{i=1}^n (D_i^t \times S_i)$ *such that* $P(\hat{a}) = \bar{P}'$ *with* $\bar{P}'_{j_i} \geq P_{j_i}$ *for all* $j_i \in T$*, and* $\bar{P}'_{j^*} > P_{j^*}$ *for at least one* $j^* \in T$*.*

This means that $\bar{P}$ is an optimal achievable payoff for coalition T if there is no other payoff vector such that the payoff is no worse for any member and it is better for at least one—for all (in particular for the worst) processes that nonmembers can pick. The following proposition shows the analog of Theorem 1 for the game that includes deliberation/communication:

Proposition 3 $\hat{a} \in \prod_{i=1}^n (D_i^{t-1} \times S_i)$ *is such that* $P(\hat{a}) \in v_{G^t}(\{i\})$ *for each* i *iff* $\hat{a}$ *is in* $\gamma_{\mathcal{C}_\alpha}(G^t)$*.*

Proof. *Immediate from Theorem 1.* □

This means that a communication/deliberation-action process in the α-core corresponds to a Pareto-optimal payoff vector. In the rest of this paper we will focus on α-core processes since they are not only Pareto-optimal in G^t but also lead (according to Lemma 2) to an outcome $s \in \prod_{i=1}^{n} S_i$ that cannot be blocked by a coalition formed in any other process.

For many games (the prisoner's dilemma being an example) a Pareto-optimal payoff can be reached only through the coordinated activity of agents, i.e. the Pareto-optimal outcome is reached only if each agent acts in agreement with the other agents. Let us return to Example 1 to show how communication/deliberation activities allow to reach a coordination on a Pareto-optimal outcome:

Example 2 *Consider again the game G of Example 1, and assume that the communication/deliberation actions at every step are*

$$D_a \;=\; D_b \;=\; \{d, d^{'}, d^{''}\}$$

and the associate costs are

$$C_a(d) = C_b(d) = 0.1$$

$$C_a(d^{'}) = C_b(d^{'}) = 0.1$$
$$C_a(d^{''}) = C_b(d^{''}) = 0.5$$

Say that the actions are abstractly described as follows:

- *d: engage in negotiations*
- *$d^{'}$: reach an agreement*
- *$d^{''}$: sign a contract to enforce the agreement*

Moreover, we assume that engaging in negotiations, d, has as a consequence the realization that if no enforceable agreement is reached, the outcome will be $\langle nc, nc\rangle$.

The sequence $(\langle d,d\rangle, \langle d^{'}, d^{'}\rangle, \langle d^{''}, d^{''}\rangle, \langle c,c\rangle)$ is in the α-core because the payoff for every player, $5-(0.1+0.1+0.5)$, is higher than any other payoff, considering that d is an unavoidable step in the process. If an agent would choose the process $\{d, nc\}$ it would know that the other agent would do the same, and the payoff would be only $2-0.1$.

5.1 Incorporating belief revision

The previous example is very simple, but it shows the sequential nature of an agent's choice of action. This subsection introduces a more sophisticated decision making model for an agent that takes part in coalition formation. This model is used to show results on the joint outcomes and the joint process.

To choose the action that maximizes expected payoff at each step, an agent may need to evaluate the expected payoffs of different actions. We will show when this procedure leads to the formation of a stable coalition structure. To give a mathematical characterization, we introduce the notion of "expected payoff" which is based on each agent's subjective probabilities:

Definition 16 *Given a sequence of actions performed by an agent i, $\hat{a}_i^t = (a_i^0, \ldots, a_i^t) \in D_i^{t+1}$, we say that agent i can define a subjective probability distribution on $\prod_{i=1}^n S_i$ such that $\Delta_i^t(s|a_i^{t+1})$ is the conditional probability of an outcome s, given that the next action is a_i^{t+1}. A probability distribution on the total costs associated with the process to reach $s \in \prod_{i=1}^n S_i$ can be also defined: $\delta_i^t(C_i(s)|a_i^{t+1})$ is the conditional probability of a cost $C_i(s)$, given that the next action is a_i^{t+1}. Then the expected payoff, given that the next action is a_i^{t+1}, is*

$$\bar{\rho}_i^t(a_i^{t+1}) = \sum_{s \in \prod_{i=1}^n S_i} u_i(s)\Delta_i^t(s|a_i^{t+1}) - \sum_{C_i(s),\ s \in \prod_{i=1}^n S_i} C_i(s)\delta_i^t(C_i(s)|a_i^{t+1})$$

An agent can try to maximize its expected payoff in each step, i.e. to choose an $a_i^t \in (D_i \cup S_i)$ that maximizes $\bar{\rho}_i^t(\cdot)$. This is a greedy procedure, and agents that use it may not always converge on a joint solution.

A key element here is the belief formation process that generates the conditional probabilities $\Delta_i^t(s|a_i^{t+1})$ and $\delta_i^t(C_i(s)|a_i^{t+1})$. We assume in the following that $\prod_{i=1}^n S_i$, $\prod_{i=1}^n D_i$ and the length of the process, N, are common knowledge. Agents are assumed to update their beliefs using Bayes Rule [9]. The general decision process is as follows:

Algorithm 1 *At stage $k = 0$ agent i generates a probability distribution μ_i^0 over $(\prod_{i=1}^n D_i)^{N-1} \times \prod_{i=1}^n S_i$, such that for each $\hat{a} \in (\prod_{i=1}^n D_i)^{N-1} \times \prod_{i=1}^n S_i$, $\mu_i^0(\hat{a}) > 0$.*[14],[15]

[14] To choose a_i^0, agent i performs steps $2, \ldots, 6$ with $k + 1 = 0$

[15] Bayesian updating requires positive probability of every process: if Bayes Rule is applied, a process with zero probability cannot get a positive probability thereafter.

For $k = 1 \dots N-1$:

1. *The last common action observed is $a^k = (a_1^k, \dots, a_n^k)$. Then, for every $\bar{a} = (\bar{a}^0, \dots, \bar{a}^N)$, the new distribution μ_i^{k+1} is such that*

$$\mu_i^{k+1}(\bar{a}) \;=\; 0 \quad \textit{if} \;\; \bar{a}^k \;\neq\; a^k$$

$$\textit{and} \;\; \mu_i^{k+1}(\bar{a}) \;=\; \frac{\mu_i^k(\bar{a})}{\sum_{a \in \prod_{i=1}^n D_i)^{N-1} \times \prod_{i=1}^n S_i : \mu_i^{k+1} \neq 0} \mu_i^k(a)} \quad \textit{if} \;\; \bar{a}^k \;=\; a^k$$

2. *For each $a_i^j \in D_i$, $j = 1, \dots, |D_i|$,*

$$\Delta_i^{k+1}(s|a_i^j) \;=\; \sum_{\bar{a}_i^{k+1} = a_i^j, \bar{a}^N = s} \mu_i^{k+1}(\bar{a})$$

3. *For each $a_i^j \in D_i$, $j = 1, \dots, |D_i|$,*

$$\delta_i^{k+1}(s|a_i^j) \;= \mu_i^{k+1}(\bar{a})$$

   such that $\bar{a}_i^{k+1} = a_i^j$, $\bar{a}^N = s$ and $C_i(s) = C_i((\bar{a}_i^0, \dots, \bar{a}_i^N))$

4. *For each $a_i^j \in D_i$, the expected payoff $\bar{\rho}_i^{k+1}(a_i^j)$ is*

$$\bar{\rho}_i^{k+1}(a_i^j) \;=\; \sum_{s \in \prod_{i=1}^n S_i} u_i(s) \Delta_i^{k+1}(s|a_i^j) - \sum_{C_i(s),\; s \in \prod_{i=1}^n S_i} C_i(s) \delta_i^{k+1}(C_i(s)|a_i^j)$$

5. *Agent i selects the action a_i^j that maximizes the expected payoff $\rho_i^{k+1}(a_i^j)$*

6. *The action to perform is the a_i^j found in the previous step.*

In words: each agent considers, at $k = 0$, the set of all possible processes and chooses an action. After observing the actions of the other agents, it discards all processes that at the first stage differ from the observed set of actions. Then it chooses an action according to its new beliefs. The process is repeated until in stage $N-1$ it has to choose a domain action in the space S_i. At each stage, the set of possible processes and therefore the set of possible outcomes is narrowed, until a single outcome is pinpointed.

Algorithm 1 has, in each stage, a worst case complexity of $4\mathcal{B} + |D_i|(|\prod_{i=1}^n S_i|\mathcal{B} + \mathcal{B}^2)$, where $\mathcal{B} = |(\prod_{i=1}^n D_i)^{N-1} \times \prod_{i=1}^n S_i|$ is the number of possible processes. If $D_i = D_j$ for all i and j, the worst case complexity becomes polynomial in $|D_i|^{nN}$. This clearly shows that Algorithms 1 is exponential in the number of agents and the length of negotiations.[16]

When this procedure is performed in conjunction with coordination among agents (in the sense that they happen to follow a process that is in the α-core), they will converge to the belief that a particular outcome s is the most probable one (later we show that s is Pareto-optimal).

Proposition 4 *If $\hat{a} = (a^0, \ldots, a^t, \ldots, a^N)$ is in the α-core (Proposition 4 showed that this means that for each t, $a^t = (a_1^t, \ldots, a_n^t)$ is the vector of optimal decisions) then $\exists m$ such that for $t > m$ there exists an $s \in \prod_{i=1}^n S_i$ which gives the $max_{s \in \prod_{i=1}^n S_i} \Delta_i^{t-1}(s|a_i^t)$ for each i.*

Proof. *If $\hat{a}$ is in the α-core, $P(\hat{a})$ is in $v_{G^N}(\{i\})$ (by Theorem 1). Suppose that for each m there exists $t > m$ such that there does not exist $s \in \prod_{i=1}^n S_i$ which gives the $max_{s \in \prod_{i=1}^n S_i} \Delta_j^{t-1}(s|a_i^t)$. In particular, given $m = N-1$, for $t = N$ there does not exist s giving $max_{s \in \prod_{i=1}^n S_i} \Delta_i^{N-1}(s|a_i^N)$. So, it means that at least one agent will deviate, making another strategy profile more probable. Then, since $a_i^N \in S_i$ it is clear that for at least one agent i^*, $\rho_{i^*}(\hat{a}') > \rho_{i^*}(\hat{a})$, where $a' = (a^{1'}, \ldots, a^{N'})$, $a_j^N = a_j^{N'} = s_j$ for $j \neq i^*$ and $a_{i^*}^{N'} \neq a_{i^*}^N = s_{i^*}$. Contradiction because $P(\hat{a}) \in v_{G^N}(\{i^*\})$.* $\square$

The converse is not true. A process that leads to a stable coalition structure might not be in the α-core. It is intuitive that a stable structure can be formed in a cost-inefficient process. This process could be blocked by another one leading to the same coalition structure, thus preserving stability. Therefore, Proposition 4 only gives a necessary condition for a process to be an element of the α-core. However, this is all we need since the following result shows that a coalition structure is stable if it leads to convergence of beliefs about the strategy profile to be chosen.

[16]This seems to be a discouraging result. However, it is perhaps less so in light of the fact that in economic processes where agents have independent preferences, exponential complexity in the number of agents or in the length of processes is pervasive [26].

Theorem 4 *For G^N, a process $\hat{a} = (a^0, \dots, a^N)$ leads to a stable coalition structure iff $\exists m$ such that $\forall t > m$ $\exists s \in \prod_{i=1}^n S_i$ that gives the $max_{s \in \prod_{i=1}^n S_i} \Delta_i^{t-1}(s|a_i^t)$ for each i.*

Proof.

- $\rightarrow$) *Trivial: if a coalition T formed in a stage $t \leq m$ can block $\hat{a}^N$ it means that there exists $\hat{a}^{N'} \in \prod_{i=1}^n S_i$, such that $u_j(a^{N'}) \geq u_j(a^N)$ for $j \in T$ and for one $j^* \in T$, $u_{j^*}(a^{N'}) \geq u_{j^*}(a^N)$ and $a^{N'}$ maximizes $\bar{\rho}_j^{N-1}(\cdot)$. Contradiction because the optimal decision in N is a^N.*

- $\leftarrow$) *Suppose that for all m there exists $t > m$ such that there is no $s \in \prod_{i=1}^n S_i$ that gives the $max_{s \in \prod_{i=1}^n S_i} \Delta_i^{t-1}(s|a_i^t)$ for each i. If so, for $m = N - 1$, there exists i^* and $s^* \neq a^N$ verifying that $\bar{\rho}_{i^*}^{N-1}(s^*) > \bar{\rho}_{i^*}^{N-1}(a^N)$, but then, $\{i^*\}$ is a coalition that blocks a^N. Contradiction.* $\square$

Alternatively, this result can be stated in a more detailed way because utilities and costs are independent:

Theorem 5 *In G^N, a process $\hat{a} = (a^0, \dots, a^N)$ leads to a stable coalition structure iff $\exists m$ such that $\forall t > m$ and for each i, there exists a pair $(a^N, a_i^t) \in (\prod_{i=1}^n S_i) \times D_i$ such that $u_i(\cdot)\Delta_i^{t-1}(\cdot|\cdot)$ and $C_i(\cdot)\delta_i^{t-1}(C_i(\cdot)|\cdot)$ achieve a maximum and a minimum (respectively) in (a^N, a_i^t).*[17]

Proof.

- $\rightarrow$) *Assume that for all m there exists $i \in I$ such that for all $a_i^t \in D_i \cup S_i$ there exists $t > m$ for which $u_i(\cdot)\Delta_i^{t-1}(\cdot|\cdot)$ and $C_i(\cdot)\delta_i^{t-1}(C_i(\cdot)|\cdot)$ do not achieve a maximum and a minimum in (a^N, a_i^t). Taking $m = N - 1$ it follows that there exists $s \in \prod_{i=1}^n S_i$ such that $u_i(s)\Delta_i^{N-1}(s|s_i) > u_i(a^N)\Delta_i^{N-1}(a^N|a_i^N)$ and $C_i(s)\delta_i^{N-1}(C_i(s)|a_i^N) < C_i(a^N)\delta_i^{N-1}(C_i(a^N)|a_i^N)$. Therefore i can block a^N. Contradiction.*

[17] A trivial example where this conditions can be fulfilled is that of a game G with a unique Nash-equilibrium that is also Pareto-optimal: each individual action chosen will be one that leads to that maximal outcome and will therefore minimize costs. There are other examples as well.

- $\leftarrow$) *suppose that there exists a coalition* $T \subseteq I$ *that can be formed in another process* $\hat{a}'$. *Therefore for at least one* $i \in T$, *there exists* $s \in \prod_{i=1}^{n} S_i$ $(s_i = a_i^{N'})$*such that* $u_i(s)\Delta_i^{N-1}(s|s_i) > u_i(a^N)\Delta_i^{N-1}(a^N|a_i^N)$ *and* $C_i(s)\delta_i^{N-1}(C_i(s)|s_i) < C_i(a^N)\delta_i^{N-1}(C_i(a^N)|s_i)$. *Contradiction because for* (a^N, a_i^N) $u_i(\cdot)\Delta_i^{t-1}(\cdot|\cdot)$ *and* $C_i(\cdot)\delta_i^{t-1}(C_i(\cdot)|\cdot)$ *achieve a maximum and a minimum.* □

Put together, Theorem 5 gives the conditions under which Algorithm 1 leads to a stable coalition structure. Sometimes the *process* that leads to this coalition structure is stable according to the α-core and sometimes not. In either case, *stability is equivalent to the coincidence of beliefs about the process*. That is, if the negotiation does not lead the agents to share the same expectations about the final outcome, the result will not be supported by a stable coalition structure.

To see how this works, let us revisit Example 2, this time incorporating belief revision:

Example 3 *Consider the game* G *of Example 2, and assume that each agent can choose among eight possible sequences of actions (each sequence has at most four stages):*

- $I \equiv (d, d', d'', c)$
- $II \equiv (d, d', d'', nc)$
- $III \equiv (d, d'', c)$
- $IV \equiv (d, d'', nc)$
- $V \equiv (d', d'', c)$
- $VI \equiv (d', d'', nc)$
- $VII \equiv (d'', c)$
- $VIII \equiv (d'', nc)$

We assume that agents a and b have symmetric prior probability distributions. Specifically, for $i = a, b$, the distribution μ_i^0 is as follows:[18]

- sequences where the first action is d: $\mu_i^0(I) = 0.512$; $\mu_i^0(III) = 0.128$; $\mu_i^0(II) = 0.128$; $\mu_i^0(IV) = 0.032$.
- sequences where the first action is d': $\mu_i^0(V) = 0.08$; $\mu_i^0(VII) = 0.02$.
- sequences where the first action is d'': $\mu_i^0(VII) = 0.08$; $\mu_i^0(VIII) = 0.02$.

It is immediate that

- $\Delta_i^1(\langle c, c\rangle, d) = \mu_i^0(I) + \mu_i^0(III) = 0.64$
- $\Delta_i^1(\langle nc, nc\rangle, d) = \mu_i^0(II) + \mu_i^0(IV) = 0.16$
- $\Delta_i^1(\langle c, c\rangle, d') = \mu_i^0(V) = 0.08$
- $\Delta_i^1(\langle nc, nc\rangle, d') = \mu_i^0(VI) = 0.02$
- $\Delta_i^1(\langle c, c\rangle, d'') = \mu_i^0(VII) = 0.08$
- $\Delta_i^1(\langle nc, nc\rangle, d'') = \mu_i^0(VIII) = 0.02$

Let the costs for $i = a, b$ be

- $C_i(\langle c, c\rangle) = 0.7$ *with probability* 0.512*(if i chooses sequence I)*
- $C_i(\langle c, c\rangle) = 0.6$ *with probability* 0.128*(if i chooses sequence III)*

[18] We assume that all other possible cases have probability 0. This is contrary to Algorithm 1, but it simplifies the exposition.

- $C_i(\langle c, c\rangle) = 0.6$ *with probability* 0.08 *(if i chooses sequence V)*

- $C_i(\langle c, c\rangle) = 0.5$ *with probability* 0.08 *(if i chooses sequence VII)*

Note that $C_i(\langle c, c\rangle)$ varies depending on the communication/deliberation process that leads to the domain actions $\langle c, c\rangle$. We can also include a "breaching cost" into the model, say 0.6. This means that if the agents agree on a particular profile of domain strategies, but some agent deviates, then that agent has to pay the breaching cost. With the breaching cost included we have:

- $C_i(\langle nc, nc\rangle) = 1.3$ *with probability* 0.128 *(if i chooses sequence II)*

- $C_i(\langle nc, nc\rangle) = 1.2$ *with probability* 0.032 *(if i chooses sequence IV)*

- $C_i(\langle nc, nc\rangle) = 1.2$ *with probability* 0.02 *(if i chooses sequence VI)*

- $C_i(\langle nc, nc\rangle) = 1.1$ *with probability* 0.02 *(if i chooses sequence $VIII$)*

Therefore, the agent will have the following beliefs about costs:

- $\delta_i^1(C_i(\langle c, c\rangle)|d) = 0.64$ *(choosing sequences I or III)*

- $\delta_i^1(C_i(\langle nc, nc\rangle)|d) = 0.16$ *(choosing sequences II or IV)*

- $\delta_i^1(C_i(\langle c, c\rangle)|d^{'}) = 0.08$ *(choosing sequence V)*

- $\delta_i^1(C_i(\langle nc, nc\rangle)|d^{'}) = 0.02$ *(choosing sequence VI)*

- $\delta_i^1(C_i(\langle c, c\rangle)|d^{''}) = 0.08$ *(choosing sequence VII)*

- $\delta_i^1(C_i(\langle nc, nc\rangle)|d^{''}) = 0.02$ *(choosing sequence $VIII$)*

Finally, the possible expected payoffs in the first iteration are, for $i = a, b$:

$$\begin{aligned}\bar{\rho}_i^1(d) &= u_i(\langle c, c\rangle)\Delta_i^1(\langle c, c\rangle, d) + u_i(\langle nc, nc\rangle)\Delta_i^1(\langle nc, nc\rangle, d)\\ &- C_i(\langle c, c\rangle)\delta_i^1(C_i(\langle c, c\rangle)|d) - C_i(\langle nc, nc\rangle)\delta_i^1(C_i(\langle nc, nc\rangle)|d)\\ &= 5 \times 0.64 + 2 \times 0.16 - (0.7 + 0.6) \times 0.64 - (1.3 + 1.2) \times 0.16 = 2.288\end{aligned}$$

$$\begin{aligned}\bar{\rho}_i^1(d^{'}) &= u_i(\langle c, c\rangle)\Delta_i^1(\langle c, c\rangle, d^{'}) + u_i(\langle nc, nc\rangle)\Delta_i^1(\langle nc, nc\rangle, d^{'})\\ &- C_i(\langle c, c\rangle)\delta_i^1(C_i(\langle c, c\rangle)|d^{'}) - C_i(\langle nc, nc\rangle)\delta_i^1(C_i(\langle nc, nc\rangle)|d^{'})\\ &= 5 \times 0.08 + 2 \times 0.02 - 0.6 \times 0.08 - 1.2 \times 0.02 = 0.368\end{aligned}$$

$$\begin{aligned}\bar{\rho}_i^1(d^{''}) &= u_i(\langle c, c\rangle)\Delta_i^1(\langle c, c\rangle, d^{''}) + u_i(\langle nc, nc\rangle)\Delta_i^1(\langle nc, nc\rangle, d^{''})\\ &- C_i(\langle c, c\rangle)\delta_i^1(C_i(\langle c, c\rangle)|d^{''}) - C_i(\langle nc, nc\rangle)\delta_i^1(C_i(\langle nc, nc\rangle)|d^{''})\\ &= 5 \times 0.08 + 2 \times 0.02 - 0.5 \times 0.08 - 1.1 \times 0.02 = 0.378\end{aligned}$$

So, according to Algorithm 1, both agents choose action d. Therefore only sequences I to IV are still possible. Probabilities are re-evaluated using the fact that the total probability of sequences beginning with d was $\mu_i^0(I) + \mu_i^0(II) + \mu_i^0(III) + \mu_i^0(IV) = 0.8$, thus assigning probability 0 to processes V, VI, VII and $VIII$:

- $\mu_i^1(I) = \frac{\mu_i^0(I)}{0.8} = \frac{0.512}{0.8} = 0.64$
- $\mu_i^1(III) = \frac{\mu_i^0(III)}{0.8} = \frac{0.128}{0.8} = 0.16$
- $\mu_i^1(II) = \frac{\mu_i^0(II)}{0.8} = \frac{0.128}{0.8} = 0.16$
- $\mu_i^1(IV) = \frac{\mu_i^0(IV)}{0.8} = \frac{0.032}{0.8} = 0.04.$

Therefore, the beliefs about the final (domain) strategies are

- $\Delta_i^2(\langle c, c\rangle, d^{'}) = \mu_i^1(I) = 0.64$

- $\Delta_i^2(\langle nc, nc\rangle, d^{'}) = \mu_i^1(II) = 0.16$

- $\Delta_i^2(\langle c, c\rangle, d^{''}) = \mu_i^1(III) = 0.16$

- $\Delta_i^2(\langle nc, nc\rangle, d^{''}) = \mu_i^1(IV) = 0.04$

and the beliefs about the corresponding costs are

- $\delta_i^2(C_i(\langle c, c\rangle)|d^{'}) = \mu_i^1(I) = 0.64$

- $\delta_i^2(C_i(\langle nc, nc\rangle)|d^{'}) = \mu_i^1(II) = 0.16$

- $\delta_i^2(C_i(\langle c, c\rangle)|d^{''}) = \mu_i^1(III) = 0.16$

- $\delta_i^2(C_i(\langle nc, nc\rangle)|d^{''}) = \mu_i^1(IV) = 0.04$

The expected payoffs are:

$$\begin{aligned}
\bar{\rho}_i^2(d^{'}) &= u_i(\langle c, c\rangle)\Delta_i^2(\langle c, c\rangle, d^{'}) + u_i(\langle nc, nc\rangle)\Delta_i^2(\langle nc, nc\rangle, d^{'}) \\
&- C_i(\langle c, c\rangle)\delta_i^2(C_i(\langle c, c\rangle)|d^{'}) - C_i(\langle nc, nc\rangle)\delta_i^2(C_i(\langle nc, nc\rangle)|d^{'}) \\
&= 5 \times 0.64 + 2 \times 0.16 - 0.7 \times 0.64 - 1.3 \times 0.16 = 2.864
\end{aligned}$$

$$\begin{aligned}
\bar{\rho}_i^2(d^{''}) &= u_i(\langle c, c\rangle)\Delta_i^2(\langle c, c\rangle, d^{''}) + u_i(\langle nc, nc\rangle)\Delta_a^i(\langle nc, nc\rangle, d^{''}) \\
&- C_i(\langle c, c\rangle)\delta_i^2(C_i(\langle c, c\rangle)|d^{''}) - C_i(\langle nc, nc\rangle)\delta_i^2(C_i(\langle nc, nc\rangle)|d^{''}) \\
&= 5 \times 0.16 + 2 \times 0.04 - 0.7 \times 0.16 - 1.3 \times 0.04 = 0.716
\end{aligned}$$

Now, both agents will choose $d^{'}$. Therefore only sequences I and II remain feasible. Since the only possibility is to choose $d^{''}$, we can directly proceed to the fourth iteration of Algorithm 1. The probabilities are (since $\mu_i^1(I) + \mu_i^1(II) = 0.8$)

- $\mu_i^3(I) = \mu_i^2(I) = \frac{\mu_i^1(I)}{0.8} = \frac{0.64}{0.8} = 0.8$

- $\mu_i^3(II) = \mu_i^2(II) = \frac{\mu_i^1(II)}{0.8} = \frac{0.16}{0.8} = 0.2.$

Therefore the beliefs become:

- $\Delta_i^4(\langle c, c\rangle, c) = \mu_i^3(I) = 0.8$
- $\Delta_i^4(\langle nc, nc\rangle, nc) = \mu_i^3(II) = 0.2$
- $\delta_i^4(C_i(\langle c, c\rangle)|c) = \mu_i^3(I) = 0.8$
- $\delta_i^4(C_i(\langle nc, nc\rangle)|nc) = \mu_i^3(II) = 0.2.$

Now the expected payoffs are:

$$\begin{aligned}\bar{\rho}_i^4(c) &= u_i(\langle c, c\rangle)\Delta_i^4(\langle c, c\rangle, c) - C_i(\langle c, c\rangle)\delta_i^4(C_i(\langle c, c\rangle)|c) \\ &= 5 \times 0.8 - 0.7 \times 0.8 = 3.44\end{aligned}$$

$$\begin{aligned}\bar{\rho}_i^4(nc) &= u_i(\langle nc, nc\rangle)\Delta_i^4(\langle nc, nc\rangle, nc) - C_i(\langle nc, nc\rangle)\delta_i^4(C_i(\langle nc, nc\rangle)|nc) \\ &= 2 \times 0.2 - 1.3 \times 0.2 = 0.14\end{aligned}$$

So c is the chosen domain strategy for each agent. Therefore, the process $I \times I$ (each agent choosing sequence I) is stable according to Theorem 4: $\Delta_i^1(\langle c, c\rangle|d), \Delta_i^2(\langle c, c\rangle|d')$ and $\Delta_i^3(\langle c, c\rangle|d'')$ were maximal among the conditional probabilities of strategies in G given actions taken from D.

To summarize, even in the presence of uncertainty, if beliefs at the beginning of the deliberation/communication processes are "reasonable", the agents will converge to a process in the α-core. As Example 3 shows, Algorithm 1 leads to a shared belief among the agents. This shared point of view supports a stable process—the same one as the one leading to cooperation in Example 1.

5.2 Deliberation/communication processes of different lengths

The results of the previous section are highly dependent on the length of the process: two processes $\hat{a}$ and $\hat{a}'$ are comparable only if their lengths are the same. If not, Theorems 4 and 5 cannot be applied. If we maintain that the degree of impatience of each agent, t_i, is given beforehand, it is clear that the game has a definite length $max_{i \in I} t_i$. Even if not, a condition on the convergence of beliefs can be given. The following result shows that every convergent process (in the sense that agents agree in their beliefs about the final outcome), leads to a stable coalition structure in an *endogenously defined* timing. In other words, there always exists a process that provides the outcome on which agents agree, and the length of the process is finite even if it is not given exogenously. This result is independent of the belief updating mechanism used by the agents.

Theorem 6 *If* $\exists m$ *such that* $\forall i$ *and* $\forall t > m$, $u_i(\cdot)\Delta_i^{t-1}(\cdot|\cdot)$ *and* $C_i(\cdot)\delta_i^{t-1}(C_i(\cdot)|\cdot)$ *have a maximum and a minimum (respectively) in a pair* (s^*, a_i^t)*, then there exists a finite* N *such that* $a^N = s^*$*.*

Proof. *Given that for each* i*, for each* $t > m$, $u_i(\cdot)\Delta_i^{t-1}(\cdot|\cdot)$ *and* $C_i(\cdot)\delta_i^{t-1}(C_i(\cdot)|\cdot)$ *have a maximum and a minimum (respectively) in a pair* (s^*, a_i^t)*, we know that for every* i*, there exists* t_i *such that* $s_i^* = a_i^{t_i}$ *(otherwise the process* $\hat{a}_i$ *is infinite and* $C_i(s^*) \to \infty$ *and therefore* $\rho(\hat{a}_i) \to -\infty$*). For* $t > t_i$*,* a_i^t *is the action of waiting for the decision of the other agents. Then, taking* $N = max_i t_i$*, we see that* $a^N = s^*$*.* □

What this result indicates is simply that if a negotiation has to end successfully, some agents have to wait until the rest of the negotiators make their decisions. Therefore, an externally given end-time is unnecessary in successful negotiations (as well as in unsuccessful ones).

5.3 Relating outcomes of rational and bounded rational agents

Communication/deliberation processes were introduced in order to describe the dynamic formation of coalition structures in the original game G. Theorem 7 shows that the outcome of a process that converges to a stable coalition structure is Pareto-optimal (part of the Pareto or efficiency frontier in the original game G). Conversely, any Pareto-optimal outcome can be supported by a process that converges to a stable coalition structure.[19] Formally, the relationship between outcomes in a stable coalition structure formed in a communication/deliberation process, and outcomes achieved by perfectly rational agents in the original game G is as follows:

Theorem 7 *A process $\hat{a} = (a^1, \ldots, a^N)$ leads to a stable coalition structure iff there does not exists $s \in \prod_{i=1}^n S_i$, such that $u_i(s) \geq u_i(a^N)$ for all i and $u_{i^*}(s) > u_{i^*}(a^N)$ for at least one i^*.*

Proof.

- $\rightarrow$) *Suppose that there exists $s \in \prod_{i=1}^n S_i$, $u_i(s) \geq u_i(a^N)$ for all i and at least for one i^*, $u_{i^*}(s) > u_{i^*}(a^N)$. Then, another process $\hat{a}' = (a'^1, \ldots, a'^N)$ can be generated, $a'^N = s$. Contradiction.*
- $\leftarrow$) *Suppose that there exists another process $\hat{a}' = (a'^1, \ldots, a'^N)$, $a'^N = s$, such that for at least one i^*, $u_{i^*}(s) > u_{i^*}(a^N)$. Contradiction.* $\square$

This result can be interpreted positively or negatively. The negative aspect is that stability is not enough to select a single outcome. Seen positively, Theorem 7 says that any stable process will lead to an efficient result and, conversely, that any efficient outcome can be attained by means of a stable communication/deliberation process.

6 Conclusions

We analyzed the problem of coalition formation in games without sidepayments. First, the α-core solution concept was reviewed in the context of games in which agents are perfectly rational. We showed that a solution is in the α-core if the corresponding utility profile is Pareto-optimal, i.e. an individual utility cannot be improved without diminishing the utility of another

[19]There is an analogy between this result and the *folk theorems* for repeated games [9]. Both show that the outcome of a process lies in a particular region of the payoffs space: above the minimax point in the folk theorems and the Pareto-frontier here.

agent. This property is closely related to superadditivity, a property indicating that shared optimal achievable utilities for different coalitions remain optimal achievable utilities for the union of the coalitions. Superadditivity implies that any two coalitions are best off by merging.

Next we explored the relationships between axiomatic and strategic solution concepts. We showed that any solution that is stable according to the α-core corresponds to a Strong Nash equilibrium (and to a Coalition-Proof Nash equilibrium and a Nash equilibrium). This allows us to study games with the α-core solution concept while our positive stability results carry over directly to these three strategic equilibrium-based solution concepts. This also allows one to confine the search for stable α-core solutions to the space of Pareto-efficient Strong Nash equilibria (or Coalition-Proof Nash equilibria or Nash equilibria).

For bounded rational agents we showed that the α-core solution concept provides clues about the properties of the deliberation/communication processes that lead to stable coalition structures. Specifically, we showed that a process leads to a stable coalition structure if its outcome cannot be blocked by a coalition formed in another process of the same length.

We characterized the communication/deliberation process as a greedy stepwise maximization of expected payoff where deliberation and communication actions incur costs. We showed that when agents agree to a process that is in the α-core, this greedy algorithm leads to convergence of the agents' beliefs in a finite number of steps. We also showed that the convergence of beliefs implies that the final outcome is stable. This holds when the protocol length is exogenously restricted as well as when agents can endogenously decide the length. More general mathematical conditions for such stability were also derived, and their meaning seems inescapable: stability can only be achieved when all the agents share their beliefs about the final outcome of the negotiation.

Finally, we showed that the outcome of any communication/deliberation process that leads to a stable coalition structure is Pareto-optimal for the original game that does not incorporate communication or deliberation. Conversely, any Pareto-optimal outcome can be supported by a communication/deliberation process that leads to a stable coalition structure.

References

[1] Aumann, R., Acceptable Points in General Cooperative n-Person Games, in Tucker,A. and Luce,D. (eds.) *Contributions to the Theory of Games IV*, Princeton University Press, Princeton 1959.

[2] Aumann, R., The Core of a Cooperative Game without Sidepayments, *Transactions of the American Mathematical Society* **98**: 539-552, 1961.

[3] Aumann, R., Subjectivity and Correlation in Randomized Strategies, *Journal of Mathematical Economics* **1**: 67-96, 1974.

[4] Bernheim, B., Peleg, B., and Whinston, M., Coalition-Proof Nash Equilibria: I Concepts, *Journal of Economic Theory* **42**: 1-12, 1987.

[5] Boddy, M., and Dean, T., Deliberation Scheduling for Problem Solving in Time-Constrained Environments, *Artificial Intelligence* **67**: 245-285, 1994.

[6] Clarke, E., Multipart pricing of public goods, *Public Choice* **11**: 17-33, 1971.

[7] Crawford, V., and Sobel, J., Strategic Information Transmission, *Econometrica* **50**: 579-594, 1982.

[8] Ephrati, E., and Rosenschein, J., Deriving Consensus in Multiagent Systems, *Artificial Intelligence* **87**: 21-74, 1996.

[9] Fudenberg, D., and Tirole, J., *Game Theory*, MIT Press, Cambridge 1991.

[10] Groves, T., and Ledyard, J., Optimal allocation of public goods: a solution to the 'free rider' problem, *Econometrica* **45**: 783-809, 1977.

[11] Kahan, J. and Rapoport, A., *Theories of Coalition Formation*, Lawrence Erlbaum Associates, 1984.

[12] Kalakota, R., and Whinston, A.B., *Frontiers of Electronic Commerce*, Addison-Wesley, N.Y. 1996.

[13] Ketchpel, S., Forming Coalitions in the Face of Uncertain Rewards, *Proceedings of the National Conference on Artificial Intelligence*: 414-419, 1994.

[14] Kraus, S., Wilkenfeld, J., and Zlotkin, G., Multiagent Negotiation Under Time Constraints, *Artificial Intelligence*, **75**: 297–345, 1995.

[15] Kreps, D., and Wilson, R., Sequential Equilibria, *Econometrica* **50**: 863-894, 1982.

[16] Kreps, D., *A Course in Microeconomic Theory*, Princeton University Press, Princeton 1990.

[17] Lipman, B., How to Decide to Decide How to Decide How to... : Modeling Limited Rationality, *Econometrica* **59**: 1105-1126, 1991.

[18] Luce, R. D., and Raiffa, H., *Games and Decisions*, John Wiley and Sons, New York 1957.

[19] Mas-Colell, A., Whinston, M., and Green, J., *Microeconomic Theory*, Oxford Economic Press, Oxford 1995.

[20] Moore, J. Implementation, Contracts, and Renegotiation in Environments with Complete Information, in *Advances in Economic Theory* vol I, Laffont, J-J. (ed), Cambridge University Press, Cambridge 1992.

[21] Moulin, H., *Cooperative Microeconomics: a Game Theoretical Introduction*, Princeton University Press, Princeton 1995.

[22] Myerson, R., Credible Negotiation Statements and Coherent Plans, *Journal of Economic Theory* **48**: 264-303, 1989.

[23] Nash, J., Noncooperative Games, *Annals of Mathematics* **54**: 289-295, 1951.

[24] Pearl, J., *Probabilistic Reasoning in Intelligent Systems*, Morgan-Kaufmann, San Mateo CA, 1988.

[25] Peleg, B., Axiomatizations of the Core, in Aumann, R., and Hart, S. (eds.) *Handbook of Game Theory vol. I*, North-Holland, Amsterdam 1992.

[26] Reiter, S., and Simon, C., Decentralized Processes for Finding Equilibrium, *Journal of Economic Theory* **56**: 400-425, 1992.

[27] Rosenschein, J. S., and Zlotkin, G., *Rules of Encounter*, MIT Press, Cambridge 1994.

[28] Rubinstein, A., Perfect Equilibria in a Bargaining Model, *Econometrica* **50**: 97-109, 1982.

[29] Russell, S., and Wefald, E., *Do the Right Thing: Studies in Limited Rationality*, MIT Press, Cambridge 1991.

[30] Sandholm, T., and Lesser, V., Utility-Based Termination of Anytime Algorithms, *Proceedings of the ECAI Workshop on Decision Theory for DAI Applications*: 88-99, 1994.

[31] Sandholm, T., Unenforced Ecommerce Transactions, *IEEE Internet Computing* **1**: 47-54, 1997.

[32] Sandholm, T., and Lesser, V., Coalitions among Computationally Bounded Agents, *Artificial Intelligence*, Special Issue on Economic Principles of Multiagent Systems **94**: 99-137, 1997.

[33] Sandholm, T., and Lesser, V., Coalition Formation among Bounded Rational Agents", *Fourtheenth International Joint Conference on Artificial Intelligence (IJCAI-95)*: 662-669, 1995.

[34] Sandholm, T., Larson, K., Andersson, M., Shehory, O. and Tohmé, F., Anytime Coalition Structure Generation with Worst Case Guarantees, *Proceedings of the Fifteenth National Conference on Artificial Inteligence (AAAI-98)*: 46-53, 1998.

[35] Sandholm, T., and Lesser, V., Advantages of a Leveled Commitment Contracting Protocol, *Proceedings of the Thirteenth National Conference on Artificial Inteligence (AAAI-96)*: 126-133, 1996.

[36] Sandholm, T., *Negotiation among Self-Interested Computationally Limited Agents*, Ph.D. thesis, Dept. of Computer Science, University of Massachusetts, Amherst 1996.

[37] Shehory, O., and Kraus, S., Task Allocation via Coalition Formation among Autonomous Agents, *Fourtheenth International Joint Conference on Artificial Intelligence (IJCAI-95)*: 655-661, 1995.

[38] Shehory, O., and Kraus, S., Methods for Task Allocation via Agent Coalition Formation, *Artificial Intelligence* **101**: 165-200, 1998.

[39] Shehory, O., and Kraus, S., A Kernel-oriented Model for Coalition-formation in General Environments: Implementation and Results, *Proceedings of the Thirteenth National Conference on Artificial Inteligence (AAAI-96)*: 134-140, 1996.

[40] Shubik, M., Game Theory Models and Methods in Political Economics, in Arrow, K., and Intriligator, M. (eds.) *Handbook of Mathematical Economics vol.I*, North-Holland, Amsterdam 1981.

[41] Tohmé, F., and Sandholm, T., Coalition Formation Processes with Belief Revision among Bounded Rational Self-Interested Agents, *Fifteenth International Joint Conference on Artificial Intelligence (IJCAI-97), Workshop on Social Interaction and Communityware*: 43-51, Nagoya, Japan, August 25, 1997.

[42] Zilberstein, S., and Russell, S., Optimal Composition of Real-Time Systems, *Artificial Intelligence* **82**: 181-213, 1996.

[43] Zlotkin, G., and Rosenschein, J., Coalition, Cryptography and Stability: Mechanisms for Coalition Formation in Task Oriented Domains, *Proceedings of the Twelfth National Conference on Artificial Inteligence (AAAI-94)*: 432-437, 1994.