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

Fernando Tohmé and **Tuomas Sandholm**
e-mail: {tohme, sandholm}@cs.wustl.edu
Department of Computer Science
Washington University
Campus Box 1045
One Brookings Drive
St.Louis, MO 63130

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 the α-core is a subset of strong Nash equilibria. This fact carries our stability results directly over to three strategic solution concepts.

The main focus of the paper is on analyzing 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 maximizes its stepwise payoff given its beliefs. Conditions are derived under which this process leads to convergence of the agents' beliefs and to a stable coalition structure (when the length of the process is exogenously restricted as well as when 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 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.

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 partitioning the agents into disjoint coalitions, solving the coordination problem within each coalition, and dividing the value or cost of each coalition among member agents.

Coalition formation among self-interested agents has been widely studied in game theory [9; 5; 2; 1; 3; 4]. The main solution concepts are geared toward guaranteeing forms of stability of the coalition structure. These concepts focus on the final solution, and they usually do not address the dynamic process that leads to that solution. Recent DAI work on coalition formation has introduced protocols for dynamic coalition formation, but the process and the agents' strategies in that process have not been included in the solution concept. In other words, although the outcomes satisfy different forms of stability, it is often not guaranteed that the process itself is stable or that individual agents should adhere to that process [7; 8]. Also, it is often implicitly assumed that agents can carry out intractable computations [10; 7; 8]. On the other hand, recent DAI work has sometimes addressed the computational limitations by explicitly incorporating computational actions in the solution concept [6]. This allows one to game theoretically trade off computation cost against solution quality. However, that work did not include protocols for dynamic coalition formation, and it did not address belief revision.

This paper studies self-interested agents with a special focus on the sequential deliberation (computation) and communication actions that the agents take in the dynamic process of coalition formation. Section 2 introduces the classic framework of game theoretic coalition formation for agents that cannot make sidepayments. Section 3 analyses 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, among other results, that stability of the coalition formation process is formally equivalent to convergence of the agents' beliefs (for both exogenously and endogenously terminated negotiation), and also that the outcome is

*Fernando Tohmé was supported by the Departamento de Economía in the Universidad Nacional del Sur, in Argentina and by a Fellowship from the National Research Council (CONICET) of Argentina. Tuomas Sandholm was supported by a Washington University, Department of Computer Science, Assistant Professor Startup Package.

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 such that $I = \{1, \ldots, n\}$ is the set of* players, *S_i the set of strategies for each $i \in I$ and $\bar{U} : \prod_{i=1}^n S_i \to R^n$, such that for each $(s_1, \ldots, s_n) \in \prod_{i=1}^n S_i$,*

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

given the individual utilities on strategies: for each i, $u_i : \prod_{i=1}^n S_i \to R$

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)*[1] *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*

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*.[2] 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.

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 [6; 4]. 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) [9; 5; 10; 7; 8]. Characteristic functions are a desirable representation, so one would like to define such mathematical entities also 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. [3] Later in the paper we show that the results that we obtain under the α-assumption carry over to strategic solution concepts 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 [7; 8; 6; 5], 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 3 *Given a game $G = ((S_i)_{i=1}^n, \bar{U})$, with $\bar{U}$ such that its components are non-transferible, we say that the* characteristic function *is*

$$v_\alpha : 2^I \to 2^{R^I}$$

such that for each coalition $T \subseteq I$

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

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

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

Definition 4 *Given a game $G = ((S_i)_{i=1}^n, \bar{U})$, and a $T \subseteq I$, the* set of optimal achievable utilities *for $T = \{j_1, \ldots, j_{|C|}\}$ is the set of $\bar{U}_T$s such that:*

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

and

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

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 follows trivially:

[1] This notion of solution can be easily extended to mixed strategies, replacing each S_i by ΔS_i, the set of probability distributions on S_i.

[2] $s = (s_1, \ldots, s_i, \ldots, s_n)$ is in $\gamma_N(G)$ if for each i and for each $s_i^{'} \neq s_i$, $u_i(s_1, \ldots, s_i^{'}, \ldots, s_n) \leq u_i(s_1, \ldots, s_i, \ldots, s_n)$.

[3] The α-assumption may be impossibly pessimistic. A given nonmember may 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.

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

Proof 1 *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 $s \in \prod_{i=1}^n S_i$, so if it is not determinate, then $\bar{U}$ is not defined for every s. Contradiction $\Box$

3 The α-Core and Superadditivity

Aumann's 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. Given a vector of joint strategies, it is said to be *blocked* by a coalition if its members can be better off by moving to another vector.

Definition 5 *A coalition T blocks a vector of joint strategies $s = (s_1, \ldots, s_n)$ if for every $j_i \in T$ there exists a $s^{'} \in \prod_{i=1}^n S_i$ such that:*

- $\forall j_i$, $u_{j_i}(s^{'}) \geq u_{j_i}(s)$ *and for at least one* $j_0 \in T$, $u_{j_0}(s^{'}) > u_{j_0}(s)$
- $(\ldots, u_{j_1}(s^{'}), \ldots, u_{j_2}(s^{'}), \ldots, u_{j_{|C|}}(s^{'}), \ldots) \in v_\alpha(T)$.

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 such that 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:

Definition 6 *A $s = (s_1, \ldots, s_n)$ is in the α-core $\gamma_{\mathcal{C}_\alpha}$ if there is no coalition T that can block $(s_1, \ldots, s_n)$.*

As with the Nash correspondence, the α-core correspondence can be empty for some games:

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

$$S_a = S_b = \{n, nc\}$$

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. *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.

Under what conditions does a stable coalition structure exist, i.e., what are the conditions for the nonemptyness 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 build an intuition of this phenomenon. Superadditivity implies that any two coalitions are best off merging.

Definition 7 *A game G is* superadditive *if given any two coalitions T_1, T_2, $T_1 \cap T_2 = \emptyset$, $v_\alpha(T_1) \cap v_\alpha(T_2) \subset v_\alpha(T_1 \cup T_2)$.* [4]

We now show an interesting property that relates the characteristic functions and superadditivity. 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 2 *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, $v_\alpha(T_1) \cap v_\alpha(T_2) \subset v_\alpha(T_1 \cup T_2)$.*

[4]Note that this definition by Shubik [9] differs technically (although it is conceptually similar) from the concept of superadditivity in games with sidepayments [9; 5; 6; 10; 7; 8].

- *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 $\bigcap_{i \in I} v_\alpha(\{i\}) \neq \emptyset$, let us revisit the Prisoner's Dilemma of Example 1. It is easy to see that it is a superadditive game, but $v_\alpha(\{a\})$ and $v_\alpha(\{b\})$ have no element in common.

The following result relates the condition of the previous proposition with 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 1
- *$\rightarrow$) 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 result is useful for proving the following theorem. 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, instead of comparing 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 1
- *$\rightarrow$) 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 $i \in I$, $\bar{U}^* \in v_\alpha(\{i\})$*
  - *lets 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 $T^{''} \in 2^I - \{\emptyset\}$ □

Because for each agent i, $v_\alpha(\{i\})$ represents her 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 [6]. 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$\neq \emptyset$]).

4 Relationships between Axiomatic and Strategic Solution Concepts

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

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 [2]:

Definition 8 *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 can show a close relationship between the strong Nash solution concept and the α-core solution concept:

Theorem 2 *For any game G, $\gamma_{\mathcal{C}_\alpha}(G) \subseteq \gamma_{SN}(G)$.* [5]

Proof 2 *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* $\Box$

We can also relate the Nash correspondence and the α-core correspondence (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 3 *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 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 $\gamma_{\mathcal{C}_\alpha}(G) \subset \gamma_N(G)$* $\Box$

One 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 only needs to consider strategy profiles that are Pareto-optimal [6] and in Nash equilibrium. Alternatively, one can restrict this search to Pareto-optimal strong Nash 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 9 *For each agent i in a game G, let D_i be a set of* deliberation-communication *activities that i can perform 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*

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

Definition 10 *For each agent i, we consider her 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, and 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 that each agent stays commited to her choice once she has reached a decision.

We use a very general characterization of the communication-deliberation process, without going into the details of how an action leads to another one (e.g. how deliberation actions lead to the choice of physical actions). This approach has the advantage of providing results that can be applied to any such process. We say that the payoff of agent i in an N-period process is determined as follows:

Definition 11 *If for each i, the communication-deliberation process is $\hat{a}_i = \{a_i^0, \ldots, a_i^{t_i}\}$ $(a_i^{t_i} = s_i)$, the payoff is*

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

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

We assume that costs of activities are independent, so 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, $C_i(a_i^0) + \ldots + C_i(a_i^{t_i}) + \bar{C}_i(N - t_i)$.

Now a new game can be defined which explicitly takes the deliberation and communication actions as part of each agent's strategy. This follows the approach of [6] where such actions are explicitly incorporated into the

[5] Often the strong Nash equilibrium is too strong a solution concept, because in many games no such equilibrium exists. Recently, the *coalition-proof Nash equilibrium* [3] 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 this kinds of solutions do not exist in all games. In games where a solution is stable according to the α-core, the solution is stable according to the coalition-proof Nash equilibrium solution concept also. This is because $\gamma_{\mathcal{C}_\alpha}(G) \subseteq \gamma_{CPN}$ (which follows from our result $\gamma_{\mathcal{C}_\alpha}(G) \subseteq \gamma_{SN}(G)$ and the known fact $\gamma_{SN}(G) \subseteq \gamma_{CPN}(G)$).

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

solution concept. It differs from other DAI approaches to coalition formation where the solution concept only analyzes final outcomes [10; 7; 8].

Definition 12 *From G, $\{D_i\}_{i=1}^n$ and $t > 0$, a new game can be defined, $G^t = ((D_i^t \times S_i)_{i=1}^n, P)$, where $P : \prod_{i=1}^n (D_i^t \times S_i) \to R^n$ 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 critically on the available communication-deliberation activities and on the sequential choice of activities. We assume that the time limit is a given. To justify this, we suppose that each agent i has a degree of impatience, given by a maximum time to make a final decision t_i. In Subsection 5.2 we will relax this assumption.

In order to maximize payoffs, our self-interested agents engage in negotiations. The final outcomes represent the result of agreements among agents. To justify the self-enforcement of these agreements we need a criterion for the stability of the coalition structure:

Definition 13 *A process $\hat{a} = \{a^0, \ldots, a^t\} \in \prod_{i=1}^n (D_i^{t-1} \times S_i)$ defines a* stable coalition structure *if $a^t \in \prod_{i=1}^n S_i$ cannot be blocked (see Section 3) 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 the game G^t then $\hat{a}$ defines a stable coalition structure.*

Proof 2 *Suppose that there exist a coalition T such that exists $s \in \prod_{i=1}^n S_i$ that verifies that $u_j(s) \geq u_j(a^t)$ for $j \in T$ and $u_{j^*}(s) > u_{j^*}(a^t)$ for a $j^* \in T$. 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}, replacing the optimal achievable utilities by the optimal achievable payoffs which incorporate deliberation and communication:

Definition 14 *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 exists $\hat{a} \in \prod_{i=1}^n (D_i^t \times S_i)$ and $P(\hat{a}) = \bar{P}$. Moreover, $\not\exists \hat{a} \in \prod_{i=1}^n (D_i^t \times S_i)$, such that $P(\hat{a}) = \bar{P}'$ with, for all $j_i \in T$, $\bar{P}'_{j_i} \geq P_{j_i}$, and for at least one $j^* \in T$, $\bar{P}'_{j^*} > P_{j^*}$*

This means, again, 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. Now,

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 3 *Immediate from Theorem 1* □

This means that a communication-deliberation-action process in the α-core corresponds to a Pareto-optimal payoff. For many games (the Prisoner's Dilemma being an example) a Pareto-optimal payoff can be reached only through the coordinated activity of agents. Let us give another example:

Example 2 *Consider again the game G of Example 1, and assume that*

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

and

$$C_a(d) = C_b(d) = 0.1$$
$$C_a(d') = C_b(d') = 0.1$$
$$C_a(d'') = C_b(d'') = 0.5$$

where the actions are described as follows:

- *d: evaluate options (we will say that any of the other choices require this deliberation step before they can be chosen)*
- *d': engage in negotiations*
- *d'': reach an enforceable agreement*

Moreover, we assume that the evaluation d has as a consequence the realization that if no enforceable agreement is reached, the outcome will be $\langle nc, nc\rangle$. So, the sequence $(\langle d, d\rangle, \langle d', d'\rangle, \langle d'', d''\rangle, \langle c, c\rangle)$ is in the 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 chooses the process $\{d, nc\}$ she know that the other will do the same, so the payoff will be $2 - 0.1$.

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":

Definition 15 *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,* *such that $\delta_i^t(C_i(s)|a_i^{t+1})$ is the conditional probability of a cost $C_i(s)$* [7], *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 her expected payoff in each step, i.e. to choose a $a_i^t \in (D_i \cup S_i)$ 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. But when this procedure is performed in conjunction with coordination among agents (in the sense that they agree on 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 $t > m$ exists an $s \in \prod_{i=1}^n S_i$ that gives the $max_{s\in\prod_{i=1}^n S_i}\Delta_i^{t-1}(s|a_i^t)$ for each i.*

Proof 4 *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 exists a $t > m$ such that it does not exist an $s \in \prod_{i=1}^n S_i$ that 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 $t = N$ it does not exist a s giving the $max_{s\in\prod_{i=1}^n S_i}\Delta_i^{N-1}(s|a_i^N)$. If so, it means that at least an agent will deviate, making another strategy profile more probable. Then, as $a_i^N \in S_i$ it is clear that for at least an 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 $P(\hat{a}) \in v_{G^N}(\{i^*\})$* □

The converse is not true. A process that leads to a stable coalition structure may 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 a convergence of beliefs about the strategy profile to be chosen.

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

[7] $C_i(s)$ is the average cost of a process with the form $\hat{a} = (a_i^0, \ldots, s_i)$, incorporating also an estimate of the length of time to reach s

Proof 4
- $\rightarrow$) *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 a $j^* \in T$, $u_{j^*}(a^{N'}) \geq u_{j^*}(a^N)$ 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 exists a $t > m$ 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 a i^* and a $s^* \neq a^N$ 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* □

This result can be generalized as follows:

Theorem 5 *In G^N, $\hat{a} = (a^1, \ldots, a^N)$ defines a stable coalition structure iff there exist an m such that for each i and for all $t > m$, 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).*

Proof 5
- $\rightarrow$) *Assume that for all m exists $i \in I$ such that forall $a_i^t \in D_i \cup S_i$ 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 results that 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.*
- $\leftarrow$) *suppose that exists a coalition $T \subseteq I$ that can be formed in a another process $\hat{a}'$. Therefore for at least one $i \in T$, 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 the greedy process of Definition 15 is guaranteed to lead to a stable coalition structure. Sometimes the *process* that leads to this coalition structure is stable according to the α-core and sometimes not.

5.2 Deliberation-Communication Processes of Different Lengths

The previous result is 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, the previous result 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 $max_{i\in I}t_i$. If not, a condition on the convergence of beliefs can be given. The following result shows also that every convergent process (in the sense that agents agree in their beliefs about the final outcome), defines a stable coalition structure in an *endogenously defined* timing. In other words, there always exists a process that provides the outcome on which the agents agree, and the length of the process is finite:

Theorem 6 *If there exists an m such that for each i and 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), then there exists an N such that $a^N = s^*$.*

Proof 6 *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, 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 $a^N = s^*$* $\Box$

5.3 Comparing 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-frontier of the original game G). Conversely, any Pareto-optimal outcome can be supported by a process that converges to a stable coalition structure. [8] 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 $\hat{a} = (a^1, \ldots, a^N)$ generates 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 7

- $\to$) *Suppose that 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* $\Box$

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 agent. This property is closely related with 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 defines 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 maximization of stepwise 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 (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.

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.

[8] There is an analogy between this result and the *Folk* theorems for repeated games [4]. Both show that the outcome of a process lies in a particular region of the strategy space: above the minimax point in the Folk result and the Pareto-frontier here.

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 Side Payments, *Transactions of the American Mathematical Society* 98(3), 1961.

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

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

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

[6] Sandholm, T. and Lesser, V., Coalition Formation among Bounded Rational Agents, *14 IJCAI*:662-669, 1995.

[7] Shechory, O. and Kraus, S., Task Allocation via Coalition Formation among Autonomous Agents, *14 IJCAI*:655-661, 1995.

[8] Shechory, O. and Kraus, S., A Kernel-oriented Model for Coalition-formation in General Environments: Implementation and Results,*AAAI-96*:134-140, 1996.

[9] 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.

[10] Zlotkin, G. and Rosenschein, J., Coalition, Cryptography and Stability: Mechanisms for Coalition Formation in Task Oriented Domains, *AAAI-94*:432-437, 1994.