

# Overlapping Coalition Formation

Georgios Chalkiadakis<sup>1</sup>, Edith Elkind<sup>1</sup>, Evangelos Markakis<sup>2</sup>, Nicholas R. Jennings<sup>1</sup>

<sup>1</sup> School of Electronics and Computer Science, University of Southampton, United Kingdom  
{gc2, ee, nrj}@ecs.soton.ac.uk

<sup>2</sup> Centre for Math and Computer Science (CWI) Amsterdam, The Netherlands  
v.markakis@cwi.nl

**Abstract.** In multiagent domains, agents form coalitions to perform tasks. The usual models of cooperative game theory assume that the desired outcome is either the grand coalition or a coalition structure that consists of disjoint coalitions (i.e., a partition of the set of agents). However, in practice an agent may be involved in executing more than one task, and distributing his resources between several (not necessarily disjoint) coalitions. To tackle such scenarios, we introduce a model for cooperative games with overlapping coalitions. We then focus on concepts of stability in this setting. In particular, we define and study a notion of the core, which is a generalization of the corresponding notion in the traditional models of cooperative game theory. Under some quite general conditions, we characterize the elements of core. As a corollary, we also show that any element of the core maximizes the social welfare. We then introduce a concept of balancedness for overlapping coalitional games, and use it to characterize coalition structures that can be extended to elements of the core. Furthermore, we generalize the notion of convexity to our setting, and show that under some natural assumptions convex games have a non-empty core. To the best of our knowledge, this is the first paper to provide a generic model for overlapping coalition formation, along with a theoretical treatment of stability in this setting.

## 1 Introduction

In many settings, groups of agents have to form teams to perform certain tasks. In the game theory literature, this process is known as *coalition formation*. Traditionally, it is assumed that the desired outcome of this process is either the *grand coalition*, i.e., the set of all agents, or a *coalition structure* that consists of disjoint coalitions (i.e., a *partition* of the set of agents). Furthermore, most of this research focuses on *transferable utility* (TU games) in which there is no restriction on how agents can split the total payoff among themselves. In particular, agents from one coalition can make a payment to agents not in that coalition.

While the above assumptions are natural for some settings, there are many scenarios where they are not applicable. This is mainly for two reasons. First, in several scenarios it may only be possible to achieve the best outcome if agents can simultaneously belong to more than one coalition. In such circumstances, agents almost invariably need to distribute their resources between the coalitions in which they participate. Such “overlaps” are natural in a plethora of interesting settings: As a straightforward e-commerce example, consider online trading agents representing individuals or virtual enterprises,

and facing the challenge of forming coalitions and allocating their owners’ capital to a variety of projects simultaneously. Such requirements are also common in systems requiring multirobot coordination, computational grid networks, and sensor networks (see, e.g., [22, 15]). To date, however, there has been essentially no work on overlapping coalition formation, with just a few exceptions which we discuss in Section 3. In particular, we are not aware of any attempt to study coalitional stability in such settings. Second, it may not always be possible to split the value of a coalition with agents that do not even belong to that coalition, i.e., to allow cross-coalition transfers. Indeed, the inability of some of the agents to work together and share payoffs may be one of the primary reasons why the grand coalition does not form, and a particular coalition structure arises (for a detailed discussion, see the work of Aumann and Dreze [4]).

To address the above concerns, we introduce and study a model that explicitly takes *overlapping coalition formation (OCF)* into account. The model is applicable in situations where agents need to allocate different parts of their resources to simultaneously serve different tasks as members of different coalitions. Further, our work departs from the conventional transferable utility framework in several ways. First, there are no inherent superadditivity assumptions in our work, and hence the grand coalition does not always emerge. (Thus, our subsequent definition of the core incorporates coalition structures, unlike most traditional work in economics.) Second, we do not allow cross-coalitional transfers (this is realistic, since an agent not contributing to a coalition should not expect to receive payoff from it). Thus, though we do allow arbitrary transfers *within* coalitions, our games are not, technically speaking, games with fully transferable utility; rather, they can be considered as games with side-payments [4, 19]. Finally, our model can take task (coalitional action) execution explicitly into account; this facilitates possible extensions to tackle coalition formation under uncertainty.<sup>3</sup>

We then explore the stability concept of the *core* for such settings, and provide conditions for its existence. In particular, under some general assumptions, we first provide a characterization for outcomes, i.e., pairs of the form (overlapping coalition structure, imputation), to be in the core. Our proof is based on a graph-theoretic argument, which may be of independent interest. As a corollary of this theorem, we show that any outcome in the core maximizes the social welfare. Second, we characterize coalition structures that admit payoff allocations such that the resulting pair is in the core. This is done by generalizing the Bondareva-Shapley theorem to our setting (note that this theorem does not hold for arbitrary non-transferable utility games). Following that, we extend the notion of convexity in coalitional games to overlapping coalitions, and show that any convex OCF game has a non-empty core. Finally, we provide some natural extensions of our model, and suggest directions for future work.

## 2 Background

In this section, we provide a brief overview of the basic concepts in cooperative game theory regarding non-overlapping coalition structures. Let  $N = \{1, \dots, n\}$  be a set of players (or “agents”). A subset  $C \subseteq N$  is called a *coalition*. A *coalition structure (CS)* (in non-overlapping environments) is a partition of the set of agents.

<sup>3</sup> To ease notation, we only show how to incorporate coalitional actions in the model in Sec. 7.

Under the assumption of *transferable utility*, coalition formation can be abstracted into a fairly simple model. This assumption postulates the existence of a (divisible) commodity (e.g., “money”) that can be freely transferred among players. The role of the *characteristic function* of a *coalitional game with transferable utility (TU-game)* is to specify a single number denoting the worth of a coalition. Formally, a characteristic function  $v : 2^N \mapsto \mathbb{R}$  defines the *value*  $v(C)$  of each coalition  $C$  [32]. A transferable utility game is completely specified by the set of players  $N$  and the characteristic function  $v$ ; we write  $G = (N, v)$ .

While the characteristic function describes the payoffs available to coalitions, it does not prescribe a way of distributing these payoffs. This is captured by the notion of an *imputation*, defined as follows. We say that an *allocation* is a vector of payoffs  $\mathbf{x} = (x_1, \dots, x_n)$  assigning some payoff to each  $j \in N$ . An allocation is *efficient* with respect to a coalition structure  $CS$  if  $\sum_{j \in S} x_j = v(S)$  for all  $S \in CS$ . An allocation  $(x_1, \dots, x_n)$  is called an *imputation* if it is efficient, and satisfies *individual rationality*, i.e.,  $x_j \geq v(\{j\})$  for  $j = 1, \dots, n$ .  $I(CS)$  denotes the set of all imputations for  $CS$ .

**The non-overlapping core** When rational agents seek to maximize their individual payoffs, the *stability* of the underlying coalition structure becomes critical, as agents might be tempted to abandon agreements in pursuit of further gains for themselves. A structure is stable only if the outcomes attained by the coalitions and the payoff combinations agreed to by the agents satisfy both individual and group rationality. Given this, research in coalition formation has developed several notions of stability, among the strongest and the most well-studied ones being the *core* [18]. Taking coalition structures into account, the core of a TU game is a set of outcomes  $(CS, \mathbf{x})$ ,  $\mathbf{x} \in I(CS)$ , such that no subgroup of agents is motivated to depart from their coalitions in  $CS$ .

**Definition 1.** Let  $CS$  be a coalition structure, and let  $\mathbf{x} \in \mathbb{R}^n$  be an allocation of payoffs to the agents. The core of a game  $(N, v)$  is the set of all pairs  $(CS, \mathbf{x})$  such that  $\mathbf{x} \in I(CS)$  and  $\forall C \subseteq N$  it holds that  $\sum_{j \in C} x_j \geq v(C)$ .

Hence, no coalition would ever “block” the proposal for a core allocation. Unfortunately, the core is a strong notion, and there exist many games where it is empty.

The core definition above is essentially the definition provided in [24] (and is also very similar to the one coined in [17]). If we assume superadditivity of the characteristic function (i.e.,  $v(U \cup T) \geq v(U) + v(T)$  for any disjoint coalitions  $U$  and  $T$ ) then in the definition above we may only consider outcomes where  $CS$  is simply the grand coalition and  $\sum_{j \in N} x_j = v(N)$ . The core definition then becomes the traditional definition that has been used in the vast majority of the economics literature (see [20]).

The environments of interest in our work however are primarily non-superadditive and we will not make any such assumption on the characteristic function. Indeed, there is a plethora of realistic application scenarios where the emergence of the grand coalition is either not guaranteed, might be perceivably harmful, or is plainly impossible (see, e.g., [24, 23]). In addition to such motivations, Aumann and Dreze [4] also provide a thorough and insightful discussion on why coalition structures arise: they put forward a series of arguments on how this might happen, and explain that coalition structures may emerge naturally even in superadditive environments for a variety of reasons.

### 3 Related work

As mentioned in the introduction, very little work exists on overlapping coalition formation settings. Here we discuss some notable exceptions, as well as some related work on the core in the context of non-overlapping coalition structures.

To begin, Shehory and Kraus present a setting for overlapping coalition formation in [28]. In their model, the agents have goals and capabilities, i.e., abilities to execute certain actions. To serve their goals, the agents have to participate in coalitions, to each of which they contribute some of their capabilities, which can thus be thought of as resources. The authors then propose heuristic algorithms that lead to the creation of overlapping coalition structures. However, the authors stop short of addressing the question of the stability of overlapping coalitions. Dang *et al.* [15] also examine heuristic algorithms for overlapping coalition formation to be used in surveillance multi-sensor networks. However, their work does not deal with payoff allocation issues, and does not attack the overlapping coalition formation problem from a game-theoretic perspective.

Conconi and Perroni [13] present a model of international multidimensional policy coordination in a *non-cooperative* setting: Agreement structures between countries can be overlapping, namely a country may participate in multiple agreements, by contributing any number of proposed “elementary strategies” (which can be regarded as being chosen from *discrete* sets of resources) to an agreement. They then introduce an equilibrium concept to describe stability in this setting. In contrast to our work, their setting is non-cooperative, and there is no attempt to globally characterize the set of stable agreements (as we do). While they introduce “negotiation tie-in” restrictions (i.e., requirements that the players must form agreements on multiple issues), they only use these as a tool to describe conditions for the stability of “joint global agreements” (where *all* players participate on agreements over *all* issues). Furthermore, their model can be seen as dealing with discrete rather than continuous resources (in comparison, we deal with continuous resources and our results also hold in the discrete case).

More recently, Albizuri *et al.* [1] presented an extension of Owen’s value [21] (which, in turn, can be thought of as a generalization of the *Shapley value* [25]) to an overlapping coalition formation setting. Specifically, they present an axiomatic characterization of their *configuration value*. Though they show through an example that the stability notion of the “large consistent set” [12] (which is a non-overlapping concept) can be applied to their configuration value, they do not further discuss other solution concepts. Moreover, in contrast to our approach, there is no restriction on the number of coalitions an agent might belong to. In particular, in the model of [1] there is no notion of resources that an agent needs to distribute across the coalitions he belongs to.

Regarding non-overlapping coalition structures as presented in Section 2, Sandholm and Lesser [24] examine the problem of allocating *computational resources* to coalitions. They do not restrict themselves to superadditive settings, but discuss the stability of coalition structures instead. In particular, they introduce a notion of bounded rational core that explicitly takes into account coalition structures. Apt *et al.* [2, 3] also do not restrain themselves to problems where the outcome is the grand coalition only. Instead, they introduce various stability notions for abstract games and discuss simple transformations (e.g., split and merge rules) by which stable partitions of the set of players may emerge. However, they do not consider any extensions to overlapping coalitions.

## 4 Our model

In this section we extend the traditional model of Section 2 to cooperative games with overlapping coalitions. In most scenarios of interest, even if overlapping coalitions are allowed, an agent would not be able to participate in all possible coalitions due to lack of time, cash flow, or energy. To model this, we assume that each agent possesses a certain amount of resources which he can distribute among the coalitions he joins. Without loss of generality, we can make a normalization and assume that each agent has one unit of resource: an agent’s contribution to a coalition is thus given by the fraction of his resources that he allocates to it. We can also think of this as the agent’s “participation level”, or the fraction of time he devotes to a coalition. Of course, an agent may own several types of resources (e.g., time *and* money), and his contribution to a coalition would then be described by a vector rather than a scalar. Our model, and all of our results, extend to this more general setting in a straightforward manner. Nevertheless, for conciseness, we restrict our presentation to the single-resource setting.

As discussed above, in the non-overlapping model a coalition is a subset of agents, and a game is defined by its characteristic function  $v : 2^N \mapsto \mathbb{R}$ , representing the maximum total payoff that a coalition can get. In our setting, an *overlapping (or partial) coalition* is given by a vector  $\mathbf{r} = (r_1, \dots, r_n)$ , where  $r_j$  is the fraction of agent  $j$ ’s resources contributed to this coalition ( $r_j = 0$  means that  $j$  is not a member of the coalition). The *support* of a partial coalition  $\mathbf{r}$  is denoted by  $\text{supp}(\mathbf{r})$  and is defined as  $\text{supp}(\mathbf{r}) = \{j \in N \mid r_j \neq 0\}$ . We can now define the games we will be considering in the rest of this work.

**Definition 2.** An OCF-game  $G$  with player set  $N = \{1, \dots, n\}$  is given by a function  $v : [0, 1]^n \rightarrow \mathbb{R}$ , where  $v(0^n) = 0$ .

Function  $v$  maps each partial coalition  $\mathbf{r}$  to the corresponding payoff. We denote this game by  $G = (N, v)$ , or, if  $N$  is clear from the context, simply by  $v$ . Clearly, a “classic” coalition  $S$  can now be represented as the vector  $e^S$ , where  $(e^S)_j = 1$  for  $j \in S$  and 0 otherwise. In the economics literature, these are sometimes called *crisp* coalitions, whereas coalitions of the form  $(r_1, \dots, r_n)$  with at least one  $r_j$  in  $(0, 1)$  are referred to as *fuzzy*<sup>4</sup> coalitions [7]. We will avoid this term in this work as in computer science the term “fuzzy” refers to other concepts. We instead refer to coalitions of this kind as *partial* coalitions, or simply coalitions.

In most scenarios of interest,  $v$  is *monotone*, i.e., satisfies  $v(\mathbf{r}) \geq v(\mathbf{r}')$  for any  $\mathbf{r}, \mathbf{r}'$  such that  $r_j \geq r'_j$  for all  $j = 1, \dots, n$ . Note that if  $v$  is monotone, we have  $v(\mathbf{r}) \geq 0$  for any  $\mathbf{r} \in [0, 1]^n$ , since we set  $v(0, \dots, 0) = 0$ . In what follows, we will explicitly indicate which of our results rely on the monotonicity of  $v$ .

We now need to specify what the outcomes of an OCF-game are. In the non-overlapping setting, an outcome is a pair  $(CS, \mathbf{x})$ , where  $CS$  is a partition on  $N$  and  $\mathbf{x}$  is an imputation for  $CS$ . To extend this definition to our scenario, we start by introducing the notion of a coalition structure with overlapping coalitions:

**Definition 3.** For a set of agents  $T \subseteq N$ , a coalition structure on  $T$  is a finite list of vectors (partial coalitions)  $CS_T = (\mathbf{r}^1, \dots, \mathbf{r}^k)$  that satisfies (i)  $\mathbf{r}^i \in [0, 1]^n$ ; (ii)

<sup>4</sup> At the end of this section we discuss how our work differs from existing work on fuzzy games.

$\text{supp}(\mathbf{r}^i) \subseteq T$  for all  $i = 1, \dots, k$ ; and (iii)  $\sum_{i=1}^k r_j^i \leq 1$  for all  $j \in T$ . We will refer to  $k$  as the size of the coalition structure  $CS_T$  and write  $|CS_T| = k$ . Also,  $CS_T$  denotes the set of all coalition structures on  $T$ .

In the definition above, each  $\mathbf{r}^i = (r_1^i, r_2^i, \dots, r_n^i)$  corresponds to some partial coalition ( $r_j^i$  being the fraction of the resources that agent  $j$  contributes to  $\mathbf{r}^i$ ). The constraints state that every agent from  $T$  distributes at most one unit of his resources between the various coalitions he participates in (those may include the singleton coalition). This allows coalitions to be overlapping. Note that the coalition structure is a list rather than a set, i.e., it can contain two or more identical partial coalitions. Observe also that an agent is not required to allocate all of his resources, i.e., it can be the case that  $\sum_{i=1, \dots, k} r_j^i < 1$ . However, under monotonicity, we can assume that for each agent  $j$  we have  $\sum_{i=1}^k r_j^i = 1$  (i.e., a coalition structure is a fractional partition of the agents).

We should note here that the introduction of overlapping coalition structures imposes some new technical challenges. While in the non-overlapping setting the number of different coalition structures is finite, in our setting there can be infinitely many different partial coalitions, and hence infinitely many coalition structures. In particular, this implies that it is impossible to find the social welfare-maximizing coalition structure by enumerating all candidate solutions (in fact the maximum may not even be attained). In contrast, in a non-OCF setting this approach is possible—though, in general, infeasible.

We now extend the definition of  $v$  to coalition structures by setting  $v(CS) = \sum_{\mathbf{r} \in CS} v(\mathbf{r})$ . Furthermore, for any  $S \subseteq N$  we define  $v^*(S) = \sup_{CS \in CS_S} v(CS)$ . Intuitively,  $v^*(S)$  is the least upper bound on the value that the members of  $S$  can achieve by forming a coalition structure. We say that  $v$  is *bounded* if  $v^*(N) < \infty$ ; for most games of interest,  $v$  is likely to be bounded.

As in our setting the agents will not necessarily form the grand coalition, we will be interested in arguing about coalition structures from  $CS_N$ . The coalition structure will impose restrictions on admissible ways of distributing the gains: a payoff vector corresponds to an imputation *iff* it is obtained by distributing the value of each coalition:

**Definition 4.** Given a coalition structure  $CS \in CS_N$ ,  $|CS| = k$ , an imputation for  $CS$  is a  $k$ -tuple  $\mathbf{x} = (\mathbf{x}^1, \dots, \mathbf{x}^k)$ , where  $\mathbf{x}^i \in \mathbb{R}^n$  for  $i = 1, \dots, k$ , such that

- (Payoff Distribution) for every partial coalition  $\mathbf{r}^i \in CS$  we have  $\sum_{j=1}^n x_j^i = v(\mathbf{r}^i)$  and  $r_j^i = 0$  implies  $x_j^i = 0$ ;
- (Individual Rationality) the total payoff of agent  $j$  is at least as big as what he can achieve on his own:  $\sum_{i=1}^k x_j^i \geq v^*(\{j\})$ .

The set of all imputations for  $CS$  is denoted by  $I(CS)$ . Notice that in Definition 4, the profit from a task assigned to a partial coalition is only distributed among agents involved in executing it. Thus, no transfers of that payoff are allowed to outsiders. Now, the set of outcomes that is of interest to us is the set of *feasible agreements*:

**Definition 5.** A feasible agreement (or an outcome) for a set of agents  $J \subseteq N$  is a tuple  $(CS, \mathbf{x})$  where  $CS \in CS_J$ ,  $|CS| = k$  for some  $k \in \mathbb{N}$ , and  $\mathbf{x} = (\mathbf{x}^1, \dots, \mathbf{x}^k) \in I(CS)$ . We denote the set of all feasible agreements for  $J$  by  $\mathcal{F}(J)$ .

The payoff  $p_j$  of an agent  $j$  under a feasible agreement  $(CS, \mathbf{x})$  is  $p_j(CS, \mathbf{x}) = \sum_{i=1}^k x_j^i$ . We write  $\mathbf{p}(CS, \mathbf{x})$  to denote the vector  $(p_1(CS, \mathbf{x}), \dots, p_n(CS, \mathbf{x}))$ .

**Definition of the core** As explained in Section 2, no group of agents should be able to profitably deviate from a configuration in the core. Hence, any definition of the core has to depend on the notion of permissible deviations used. In this section we consider a fairly straightforward such notion.

In the non-overlapping setting with coalition structures, a pair (coalition structure, imputation) is in the core if no set of agents can jointly profit by deviating and forming a coalition on their own (Def. 1). Similarly, in the overlapping coalitions scenario, a deviating group of agents  $S$  may also want to break its obligations to other agents, i.e., withdraw or reduce its contribution to coalitions that contain agents in  $N \setminus S$ . As the deviators do not take into account the interests of other agents, they cannot expect to receive payoffs from the coalitions that contain non-deviating agents and therefore might have been hurt by the deviation. We formalize this approach as follows.

**Definition 6 (The overlapping core).** *A tuple  $(CS, \mathbf{x})$  is in the core of an OCF-game  $G = (N, v)$  (we write  $(CS, \mathbf{x}) \in \text{core}(G)$ ), if for any set of agents  $J \subseteq N$ , any coalition structure  $CS_J$  on  $J$ , and any imputation  $\mathbf{y} \in I(CS_J)$ , we have  $p_j(CS_J, \mathbf{y}) \leq p_j(CS, \mathbf{x})$  for some agent  $j \in J$ .*

It is easy to see that when restricted to non-overlapping coalitions, this definition collapses to the traditional definition of the core.

Finally, we point out here the differences between our concept of the core and the *Aubin core* [7], which is the main solution concept in fuzzy games. An outcome in the core of an OCF game is not necessarily the grand coalition, but it can be an (overlapping) coalition structure. In contrast, in the Aubin core, the imputations and core elements are always the grand coalition along with a split of the payoff  $v(N)$ . Furthermore, in our definition a core outcome needs to be stable against any deviation of a set  $S$  to a (possibly overlapping) coalition structure. In the Aubin core, an outcome need only be stable against a deviation to a partial (“fuzzy”) coalition, but not necessarily against deviations to a coalition structure. In short, the formation of coalition structures (overlapping or not) is not addressed in the fuzzy games literature.

## 5 Core characterization

In the previous section, we introduced a definition of the core for overlapping coalition formation games. We now proceed to provide a characterization of the set of outcomes in the core: essentially, an outcome is in the core if and only if under this outcome the total payments to each subset of agents match the maximum value that can be achieved by this subset. Our proof relies on several technical restrictions on the function  $v$  that defines the game. In particular, we require  $v$  to be continuous, monotone and bounded (observe that if a game is monotone and bounded, then  $v^*(S) < \infty$  for any  $S \subseteq N$ ), as well as to satisfy another natural restriction defined later. These assumptions allow us to avoid some pathological situations that may arise in our model at its generality, such as the supremum  $v^*(N)$  being unachievable (e.g., if  $v$  is strictly concave in one of its arguments, it can be the case that no finite coalition structure can achieve  $v^*(N)$ ).

Specifically, we say that a game  $(N, v)$  is  $U$ -finite if for any  $(CS, \mathbf{x})$  such that  $|CS| > U$  and  $\mathbf{x} \in I(CS)$ , there exists a  $(CS', \mathbf{y})$  such that  $|CS'| \leq U$ ,  $\mathbf{y} \in I(CS')$ , and  $p_j(CS, \mathbf{x}) \leq p_j(CS', \mathbf{y})$  for all  $j = 1, \dots, n$  (i.e., for any outcome  $(CS, \mathbf{x})$  with more than  $U$  coalitions there exists another outcome  $(CS', \mathbf{y})$  with at most  $U$  coalitions that is weakly preferred to  $(CS, \mathbf{x})$  by all agents). When this condition holds, we can assume that all coalition structures that arise in a game consist of at most  $U$  partial coalitions. This is a natural restriction in many practical scenarios, as it might be difficult for agents to maintain a very complicated collaboration pattern. It holds when, e.g., there is a bound on the number of partial coalitions each agent can be involved in. Another natural example is provided by a class of games where for any two partial coalitions  $\mathbf{r}, \mathbf{r}'$  such that  $\text{supp}(\mathbf{r}) = \text{supp}(\mathbf{r}')$  and  $r_j + r'_j \leq 1$  for any  $j = 1, \dots, n$ , we have  $v(\mathbf{r} + \mathbf{r}') \geq v(\mathbf{r}) + v(\mathbf{r}')$ . Note that in such games we can assume that no coalition structure contains two partial coalitions with the same support  $S$ , as it is at least as profitable for the players in  $S$  to merge these partial coalitions. (However, notice that this does not imply superadditivity, nor does it mean that the grand coalition necessarily emerges, as the criterion above refers only to coalitions with identical support.) Hence, any such game is  $2^n$ -finite.

*Remark 1.* Note that in all of our results  $U$  can also be a function of  $n$  (as long as  $U(n) < \infty$ ). Alternatively, instead of imposing the condition of  $U$ -finiteness on  $v(\cdot)$ , we could restrict the set of allowed outcomes (or potential deviations) to coalition structures with at most  $U$  partial coalitions. All of our results hold under this model as well.

We now state and prove the first of our main results.

**Theorem 1.** *Given a game  $(N, v)$ , where  $v$  is monotone, continuous, bounded, and  $U$ -finite for some  $U \in \mathbb{N}$ , an outcome  $(CS, \mathbf{x})$  is in the core of  $(N, v)$  iff for all  $S \subseteq N$*

$$\sum_{j \in S} p_j(CS, \mathbf{x}) \geq v^*(S). \quad (1)$$

**Proof :** For the “if” direction, suppose that  $(CS, \mathbf{x})$  satisfies  $\sum_{j \in S} p_j(CS, \mathbf{x}) \geq v^*(S)$  for all  $S \subseteq N$ . Assume for the sake of contradiction that  $(CS, \mathbf{x})$  is not in the core, i.e., there exists a set  $S$ , a coalition structure  $CS_S \in \mathcal{CS}_S$  and an imputation  $\mathbf{y} \in I(CS_S)$  such that  $p_j(CS_S, \mathbf{y}) > p_j(CS, \mathbf{x})$  for all  $j \in S$ . Then we have  $v(CS_S) = \sum_{j \in S} p_j(CS_S, \mathbf{y}) > \sum_{j \in S} p_j(CS, \mathbf{x}) = v^*(S)$ , a contradiction with the way  $v^*(S)$  was defined.

For the “only if” direction, consider an outcome  $(CS, \mathbf{x})$  that does not satisfy (1); we will show that  $(CS, \mathbf{x})$  is not in the core. To begin, set  $\mathbf{p} = \mathbf{p}(CS, \mathbf{x})$ , and assume  $\sum_{j \in S} p_j < v^*(S)$  for some  $S \subseteq N$ . To show that  $(CS, \mathbf{x})$  is not in the core, we will construct a set  $S'$ , a coalition structure  $CS_{S'} \in \mathcal{CS}_{S'}$  and an imputation  $\mathbf{y} \in I(CS_{S'})$  such that  $p_j(CS_{S'}, \mathbf{y}) > p_j$  for all  $j \in S'$ . Fix a set  $S$  that satisfies  $\sum_{j \in S} p_j < v^*(S)$ . Choose  $\varepsilon$  small enough so that  $\sum_{j \in S} p_j < v^*(S) - \varepsilon$ , and let  $\mathcal{CS}_S^\varepsilon = \{CS_S \in \mathcal{CS}_S \mid v(CS_S) \geq v^*(S) - \varepsilon\}$ . By definition of  $v^*(S)$ , there is an infinite sequence of coalition structures  $CS^{(t)}$  that satisfies  $\lim_{t \rightarrow \infty} v(CS^{(t)}) = v^*(S)$ , so the set  $\mathcal{CS}_S^\varepsilon$  is non-empty. Given a coalition structure  $CS_S \in \mathcal{CS}_S$ , an imputation  $\mathbf{y} \in I(CS_S)$  and a respective payoff vector  $\mathbf{q} = \mathbf{p}(CS_S, \mathbf{y})$ , define the *total loss*  $TL(CS_S, \mathbf{q})$  of  $(CS_S, \mathbf{q})$

as  $\sum_{j:p_j > q_j} (p_j - q_j)$ . Set  $TL_{\min} = \inf\{TL(CS_S, \mathbf{q}) \mid CS_S \in \mathcal{CS}_S^\varepsilon, \mathbf{y} \in I(CS_S), \mathbf{q} = \mathbf{p}(CS_S, \mathbf{y})\}$ . First, we prove that there exists a coalition structure  $CS \in \mathcal{CS}_S^\varepsilon$  and an imputation  $\mathbf{y} \in I(CS_S)$  that achieve the total loss of  $TL_{\min}$  (the proof is omitted).

**Lemma 1.** *Under the theorem's conditions, there exists a  $CS_S \in \mathcal{CS}_S^\varepsilon$ , an imputation  $\mathbf{y} \in I(CS_S)$  and a payoff vector  $\mathbf{q} = \mathbf{p}(CS_S, \mathbf{y})$  s.t.  $TL(CS_S, \mathbf{q}) = TL_{\min}$ .*

Let  $(CS_S, \mathbf{y})$  be an outcome that satisfies  $v(CS_S) \geq v^*(S) - \varepsilon$ ,  $TL(CS_S, \mathbf{p}(CS_S, \mathbf{y})) = TL_{\min}$ , whose existence is guaranteed by Lemma 1. Set  $\mathbf{q} = \mathbf{p}(CS_S, \mathbf{y})$ . Let us now construct a directed graph  $\Gamma$  whose vertices are the agents and there is an edge from  $j$  to  $i$  if there exists a coalition in  $CS_S$  containing both  $j$  and  $i$  such that under  $\mathbf{y}$ , agent  $j$  gets a non-zero payoff from that coalition, i.e., for some  $\mathbf{r}^k \in CS_S$  we have  $r_j^k, r_i^k > 0$  and  $y_j^k > 0$ . Observe that if there is an edge  $(j, i)$  in  $\Gamma$ , we can change  $\mathbf{y}^k$  by increasing the payoff to  $i$  by a small enough  $\delta$  and decreasing the payoff to  $j$  by the same value of  $\delta$  without violating the constraints, i.e., we have  $\mathbf{z} = (z^1, \dots, z^t) \in I(CS_S)$ , where  $z^l = y^l$  for  $l \neq k$  and  $z^k = (y_1^k, \dots, y_j^k - \delta, \dots, y_i^k + \delta, \dots, y_n^k)$ . Now, color all vertices of  $\Gamma$  as follows: a vertex  $j$  is red if the agent  $j$  is underpaid under  $\mathbf{y}$ , i.e.,  $q_j < p_j$ , white if  $j$  is indifferent, i.e.,  $q_j = p_j$ , and green if he is overpaid, i.e.,  $q_j > p_j$ . As  $\sum_{j \in S} p_j < v^*(S) - \varepsilon$  and  $\sum_{j \in S} q_j = v(CS_S) \geq v^*(S) - \varepsilon$ , the graph contains at least one green vertex. As argued above, if there is a path from a green vertex  $j$  to a red vertex  $i$ , we can transfer a small amount of payoff from  $j$  to  $i$  and hence decrease the total loss, which is a contradiction with our choice of  $(CS_S, \mathbf{y})$ . Hence, given an arbitrary green vertex  $j$ , the set of all vertices reachable from  $j$  in the graph, which we denote by  $R(j)$ , can only contain green or white vertices.

We would now like to argue that the agents in  $R(j)$  can successfully deviate from  $(CS, \mathbf{x})$ . Indeed, let  $CS'$  be the coalition structure that consists of the coalitions that the agents in  $R(j)$  form among themselves in  $CS_S$ . Clearly, the value of  $CS'$  is equal to the total value of the coalitions formed by these agents in  $CS_S$ . Note also that under  $(CS_S, \mathbf{y})$ , the agents in  $R(j)$  do not get any payoffs from coalitions that involve agents not in  $R(j)$ . Indeed, suppose that an  $i \in R(j)$  gets a non-zero payoff from a coalition that involves an agent  $k \notin R(j)$ . Then in  $\Gamma$  there is an edge from  $i$  to  $k$ , a contradiction with how  $R(j)$  was constructed. In other words, in  $CS_S$ , the payoffs that the agents in  $R(j)$  get come only from the coalitions that they form among themselves, and yet these agents are all green or white, i.e., each of them is doing no worse than what he was doing under  $CS$ , and some of them (in particular, agent  $j$ ) are doing strictly better. To finish the proof, let the agents in  $R(j)$  distribute the payoffs in the same way as in  $(CS_S, \mathbf{y})$ , except that player  $j$  transfers a small fraction of his payoffs to each of the white players in  $R(j)$  (this is possible by construction). The last step ensures that each agent in  $R(j)$  is strictly better off than in  $(CS, \mathbf{x})$ . This demonstrates that  $(CS, \mathbf{x})$  is not in the core, as required.  $\square$

*Remark 2.* Note that we did not have to make use of the additional restrictions we imposed on  $v$  to prove the “if” direction of the theorem (these are used in the proof of Lemma 1). Hence, this implication holds for an arbitrary  $G$ .

It is easily verifiable that Theorem 1 holds in the non-overlapping case with coalition structures as well. The result is trivial to prove in that setting, as each agent's payoffs

come from just one coalition; in contrast, we had to use more involved combinatorial arguments for transferring payoffs among agents. We also get the following interesting result as a corollary:

**Corollary 1.** *By setting  $S = N$  in the statement of Theorem 1, we conclude that any outcome in the core maximizes the social welfare.*

**Characterizing the Core Coalition Structures** In Theorem 1, we saw a necessary and sufficient condition for a tuple  $(CS, \mathbf{x})$  to belong to the core. Now, suppose that we are only given a structure  $CS = (\mathbf{r}^1, \dots, \mathbf{r}^k)$  and we want to check whether there exists *some* payoff allocation  $\mathbf{x}$  such that  $(CS, \mathbf{x})$  belongs to the core. Below we characterize the set of coalition structures  $CS$  that admit payoff allocations  $\mathbf{x}$  such that the corresponding tuple  $(CS, \mathbf{x})$  belongs to the core. Our characterization can be seen as a generalization of the notion of *balancedness* in the context of overlapping coalition formation. In the classic setting, the analogous question is “when does the grand coalition admit a payoff allocation in the core”, answered by Bondareva and Shapley [6, 26]. Before we proceed to our result, we define balancedness w.r.t. to a coalition structure:

**Definition 7.** *Fix a coalition structure  $CS = (\mathbf{r}^1, \dots, \mathbf{r}^k)$ ,  $k \in \mathbb{N}$ , and let  $K = \{1, \dots, k\}$ . A collection of numbers  $\{\lambda_S\}_{S \subseteq N}$ ,  $\{\mu_i\}_{i \in K}$  is called balanced w.r.t. the given coalition structure  $CS$ , if and only if  $\lambda_S \geq 0$  for all  $S$ , and  $\sum_{S: j \in S} \lambda_S + \mu_i = 1$  for all  $i \in K$ ,  $j \in \text{supp}(\mathbf{r}^i)$ .*

**Definition 8.** *A game is called balanced w.r.t. a coalition structure  $CS = (\mathbf{r}^1, \dots, \mathbf{r}^k)$  if and only if for every collection  $\{\lambda_S\}_{S \subseteq N}$ ,  $\{\mu_i\}_{i \in K}$ , which is balanced w.r.t.  $CS$ , it holds  $\sum_S \lambda_S v^*(S) + \sum_{i=1}^k \mu_i v(\mathbf{r}^i) \leq v^*(N)$ .*

The proof of the following theorem (omitted due to space constraints) is based on LP-Duality, and relies on the characterization result of Theorem 1; furthermore, the proof illustrates that the condition of balancedness introduced above arises rather naturally.

**Theorem 2.** *Let  $(N, v)$  be an OCF-game, where  $v$  is monotone, continuous, bounded, and  $U$ -finite for some  $U \in \mathbb{N}$  and consider a coalition structure  $CS = (\mathbf{r}^1, \dots, \mathbf{r}^k)$ , for some  $k \in \mathbb{N}$ . There exists an imputation  $\mathbf{x}$  such that  $(CS, \mathbf{x})$  belongs to the core if and only if the game is balanced w.r.t.  $CS$ .*

*Remark 3.* In the traditional superadditive setting, the condition of balancedness is a bit simpler and more intuitive. In our setting, the characterization leads to a slightly more complicated expression, essentially due to the fact that the linear program that describes core allocations for each coalition structure requires a larger set of constraints.

## 6 Convex OCF games have a non-empty core

In this section, we first generalize the notion of convexity to OCF-games and then proceed to show that it provides a sufficient condition for the non-emptiness of the core.

Recall that for classical TU-games convexity means that for  $R \subseteq N$  and  $S \subset T \subseteq N \setminus R$  it holds that  $v(S \cup R) - v(S) \leq v(T \cup R) - v(T)$ . Thus, convexity in the classic

TU-games setting means that it is more useful for a coalition  $R$  to join a larger coalition than a smaller one. We now apply this intuition to our setting:

**Definition 9.** An OCF-game  $G = (N, v)$  is convex if for each  $R \subseteq N$  and  $S \subset T \subseteq N \setminus R$  the following condition holds: for any  $(CS^S, \mathbf{x}^S) \in \mathcal{F}(S)$ , any  $(CS^T, \mathbf{x}^T) \in \mathcal{F}(T)$ , and any  $(CS^{S \cup R}, \mathbf{x}^{S \cup R}) \in \mathcal{F}(S \cup R)$  that satisfies  $p_j(CS^{S \cup R}, \mathbf{x}^{S \cup R}) \geq p_j(CS^S, \mathbf{x}^S) \forall j \in S$ , there exists an outcome  $(CS^{T \cup R}, \mathbf{x}^{T \cup R}) \in \mathcal{F}(T \cup R)$  s.t.

$$\begin{aligned} p_j(CS^{T \cup R}, \mathbf{x}^{T \cup R}) &\geq p_j(CS^T, \mathbf{x}^T) \quad \forall j \in T, \text{ and} \\ p_j(CS^{T \cup R}, \mathbf{x}^{T \cup R}) &\geq p_j(CS^{S \cup R}, \mathbf{x}^{S \cup R}) \quad \forall j \in R. \end{aligned}$$

This definition is similar in flavour to that in [29], where a generalization of convexity is defined in the context of stochastic cooperative games. The intuition behind this definition is as follows: Consider two fixed agreements, one on  $S$  and one on  $T$  respectively. Any time that there is a feasible agreement on  $S \cup R$  that the members of  $S$  do not object to compared to their own agreement (i.e., all members of  $S$  are weakly better off than in their previous agreement), then there is a feasible agreement on  $T \cup R$  such that (i) the members of  $T$  do not object to this agreement, compared to the previous agreement on  $T$  and (ii) the members of  $R$  weakly prefer this agreement to the agreement on  $S \cup R$ .

We now show that convexity is a sufficient condition for the non-emptiness of the core in analogy to the classic result on convex TU-games [27]. Here we only give an outline of the proof.

**Theorem 3.** If an OCF-game  $G = (N, v)$  is convex and  $v$  is continuous, bounded, monotone and  $U$ -finite for some  $U \in \mathbb{N}$ , then the core of this game is not empty.

**Proof Sketch:** To prove the theorem, we explicitly construct an outcome  $(CS, \mathbf{x})$ , and show that it belongs to the core. The construction proceeds in rounds. First let  $p_1$  be the maximum payoff that agent 1 can achieve on his own. In round 2, we choose an agreement  $(CS_2, x_2)$  on  $\{1, 2\}$  that maximizes agent 2's payoff subject to the constraint that agent 1 receives at least  $p_1$ . Then in round 3 we pick an agreement  $(CS_3, x_3)$  that maximizes the payoff of agent 3 subject to the constraint that the other 2 agents are not worse off compared to the previous round. We continue in this manner till we reach an agreement  $(CS_n, x_n)$  in round  $n$ . We then prove that  $(CS_n, x_n)$  is in the core by using induction and exploiting the game's convexity property.  $\square$

In the traditional setting, if a game is represented using oracle access for  $v(S)$ , there is a trivial algorithm for computing an element of the core in convex games. Indeed, one can set the payoff vector to be the vector of the marginal contributions of the agents for an arbitrary permutation of the set of agents. In our setting, our proof does yield a procedure for constructing an element of the core, however not a polynomial time one. Our procedure requires solving a series of optimization questions, which for arbitrary convex games are NP-hard. In the future, we would like to find classes of convex games where our proof yields a polynomial time algorithm. In particular, looking at our proof, this would be true for games in which we can solve in polynomial time the following problem: Given a set of agents  $S \subseteq N$ , a feasible agreement on  $S$ ,  $(CS, \mathbf{x})$ , and an agent  $k \notin S$ , find a feasible agreement  $(CS', \mathbf{y})$  on  $S \cup \{k\}$  that maximizes  $p_k(CS', \mathbf{y})$  subject to the constraints:  $p_j(CS', \mathbf{y}) \geq p_j(CS, \mathbf{x})$ .

## 7 Conclusions, extensions, and future work

In this paper we introduced a model of cooperative games that allows for overlapping coalitions and takes into account the need for resource allocation. In doing so, we generalize the usual models where either the grand coalition is the only desirable outcome or the outcomes are required to be partitions of the set of agents. Given our model, we defined and studied a notion of the core which is a generalization of the core in the traditional models of cooperative game theory. Under some quite general conditions, we provided a characterization for an outcome—that is, a (coalition structure, imputation) pair—to belong to the core. We also showed that any outcome in the core maximizes the social welfare. Further, we introduced balancedness for OCF games, defined balanced OCF games, and showed that a coalition structure  $CS$  admits an imputation  $x$  so that  $(CS, x)$  is in the core if and only if the game is balanced. Finally, we extended the notion of convexity to our setting and showed that convex games have a non-empty core. This is one of the very first attempts to provide a theoretical treatment of overlapping coalition formation, and, to the best of our knowledge, the first to present a generic model for overlapping coalition formation and study stability in a thorough manner.

**Extensions** In order to not overload notation, we avoided modeling *coalitional actions* in our presentation so far. However, in realistic environments coalitions are formed to execute tasks, which can be represented as coalitional actions. This is easily incorporated in our model, as follows: A coalition is allowed to select an action from a (usually finite) action space  $\mathcal{A}$ . Without loss of generality, we assume that each coalition can undertake any action in  $\mathcal{A}$ .<sup>5</sup> The value of a coalition is then determined by the resource contribution levels of its members *and* the action selected. Therefore, the characteristic function in our setting is then defined on  $(r, a)$  pairs, where  $r = (r_1, \dots, r_n)$  is a vector of resources, and  $a \in \mathcal{A}$  is an action. All of our definitions and results generalize readily to the situation where each coalition has a choice of actions (simply put, our presentation so far corresponds to a situation where each coalition had exactly one action available to it).

Another extension we have examined has to do with modeling the available resources. For ease of presentation it was assumed throughout the paper that there exists only one type of (continuous) resource. Nevertheless, all of our results still hold if we assume multiple types of resources. Moreover, we have also studied a “discrete” OCF setting, with agent resources taking values in a finite set (i.e., an agent cannot contribute an arbitrary percentage of his resources to a coalition). With discrete resources, the number of possible coalition structures is now finite (as a coalition in our setting is a collection of resources—see Section 4). All of our definitions and theorems carry through in this setting with minor differences in the arguments used in the proofs.

Finally, we have also investigated alternative notions of deviations and concepts of the core. Specifically, we have so far assumed that deviators cannot rely on getting any payoffs from their coalitions with non-deviators: in this sense, the deviators can be seen as “self-reliant”, and their form of deviation as “malicious” (since when deviating they pay no consideration to the harm inflicted on their partners). However, one can

---

<sup>5</sup> The situation where this is not the case can be modeled by setting the value of the respective (coalition, action) pair to 0.

consider a more “benign” form of deviation, where deviating teams do not break their obligations to non-deviators (possibly due to restrictions imposed by the multiagent system architecture), but instead have to rearrange resources already used in coalitions with other deviating agents. In this case, it is natural to assume that the deviating agents get to keep their share of the profit from the coalitions with non-deviators. For lack of space, we refer further discussion of these issues to an extended version of this paper.

**Future work** There exist many exciting open questions for future work. In particular, it would be interesting to investigate the alternative notion of stability proposed above. We also plan to study the computational complexity of the core. Even in superadditive settings (where the coalition structure is simply the grand coalition  $N$ ), computing an allocation in the core or checking if the core is non-empty are NP-hard problems [11, 31, 16, 14]. In the absence of superadditivity, there are even stronger lower bounds on the complexity of the problem [23]. Hence we can only hope to identify special classes of games where we can have efficient algorithms for computing core allocations. As noted earlier, an element of the core in convex games can be computed in the traditional setting simply by taking the vector of the marginal contributions of the agents for an arbitrary permutation of the set of agents. In our setting, even though our proof yields a procedure for constructing an element of the core, it requires solving a series of optimization questions, which for arbitrary convex games are NP-hard. It would be desirable to find classes of convex games where our proof yields a polynomial time algorithm.

We are also interested in finding processes that lead to the core in not necessarily convex games; though *randomized algorithms* such as the ones presented in [17] and [9] could trivially extend to the overlapping setting, they would be of little practical value here due to the huge space of potential overlapping configurations. Therefore, we are interested in finding ways to exploit known game structure to prune the search space for potential stable configurations. Another subject of future research is extending our model to allow for infinite coalition structures. Furthermore, it would be definitely interesting to establish links between outcomes in the core and outcomes of bargaining equilibria in overlapping coalitional bargaining games.

Finally, the incorporation of actions in our model allows for the investigation of action stochasticity and, more generally, uncertainty in an OCF setting. For instance, a coalitional action can be associated with a distribution over possible payoff outcomes resulting from its execution. This poses challenges to study such models from both a theoretical and a practical standpoint, since the introduction of uncertainty leads to several intricacies not readily resolved by the use of “deterministic” concepts and models, as the work of Suijs *et al.* [29, 30], Blankenburg *et al.* [5], and Chalkiadakis *et al.* [8–10] demonstrates. On a related note, enriching our model description so as to capture type uncertainty [8–10] would allow for the ready translation of uncertainty regarding the types (capabilities) of players to coalitional value uncertainty, while also capturing the potential stochasticity of actions’ outcomes.

**Acknowledgements** We would like to thank the participants of the 13th Coalition Theory Network workshop, and especially Tomasz Michalak, for valuable feedback regarding these ideas. This research was undertaken as part of the ALADDIN (Autonomous Learning Agents for Decentralised Data and Information Networks) project, which is

jointly funded by a BAE Systems and EPSRC strategic partnership (EP/C548051/1). Edith Elkind was also supported by ESRC under grant ES/F035845/1.

## References

1. M. Albizuri, J. Aurrecochea, and J. Zarzuelo. Configuration values: Extensions of the coalitional Owen value. *Games and Economic Behavior*, 57:1–17, 2006.
2. K. Apt and T. Radzik. Stable Partitions in coalitional games, 2006. Work.Paper, <http://arxiv.org/abs/cs.GT/0605132>.
3. K. Apt and A. Witzel. A Generic Approach to Coalition Formation, 2007. Work.Paper, <http://arxiv.org/abs/0709.0435>.
4. R. Aumann and J. Dreze. Cooperative Games with Coalition Structures. *International Journal of Game Theory*, 3(4):217–237, 1974.
5. B. Blankenburg, M. Klusch, and O. Shehory. Fuzzy Kernel-Stable Coalitions Between Rational Agents. In *Proceedings of the Second International Conference on Autonomous Agents and Multiagent Systems (AAMAS'03)*, 2003.
6. O. Bondareva. Some Applications of Linear Programming Methods to the Theory of Cooperative Games (in Russian). *Problemy Kibernetiki*, 10:119–139, 1963.
7. R. Branzei, D. Dimitrov, and S. Tijs. *Models in cooperative game theory*. Springer, 2005.
8. G. Chalkiadakis. *A Bayesian Approach to Multiagent Reinforcement Learning and Coalition Formation under Uncertainty*. PhD thesis, Department of Computer Science, University of Toronto, Toronto, 2007.
9. G. Chalkiadakis and C. Boutilier. Bayesian Reinforcement Learning for Coalition Formation Under Uncertainty. In *Proceedings of the 3rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS'04)*, 2004.
10. G. Chalkiadakis, E. Markakis, and C. Boutilier. Coalition Formation under Uncertainty: Bargaining Equilibria and the Bayesian Core Stability Concept. In *Proceedings of AAMAS'07*, 2007.
11. V. Chvatal. Rational behavior and computational complexity, 1978. Tech. Rep. SOCS-78.9, School of Computer Science, McGill University, Montreal.
12. M. Chwe. Farsighted Coalitional Stability. *Journal of Economic Theory*, 63:299–325, 1994.
13. P. Conconi and C. Perroni. Issue Linkage and Issue Tie-In in Multilateral Negotiations, 2001. CESifo Work.Paper 601.
14. V. Conitzer and T. Sandholm. Complexity of Determining Non-Emptiness of the Core. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI-03)*, 2003.
15. V. D. Dang, R. K. Dash, A. Rogers, and N. R. Jennings. Overlapping coalition formation for efficient data fusion in multi-sensor networks. In *Proceedings of the 21st National Conference on AI (AAAI-06)*, pages 635–640, 2006.
16. X. Deng and C. Papadimitriou. On the complexity of cooperative solution concepts. *Mathematics of Operation Research*, 19:257–266, 1994.
17. T. Dieckmann and U. Schwalbe. Dynamic Coalition Formation and the Core, 1998. Economics Department Working Paper Series, Department of Economics, National University of Ireland - Maynooth.
18. D. Gillies. *Some Theorems on n-Person Games*. PhD thesis, Department of Mathematics, Princeton University, 1953.
19. J. Greenberg and A. Kats. Unilateral Transfers and Pareto Optimality. *Econometrica*, 48(3):777–779, Apr. 1980.
20. M. Osborne and A. Rubinstein. *A course in game theory*. MIT Press, 1994.
21. G. Owen. Values of games with a priori unions. In R. Herrn and O. Moschlin, editors, *Lecture Notes in Economics and Mathematical Systems: Essays in Honor of Oskar Morgenstern*. Springer-Verlag, 1977.
22. J. Patel, W. Teacy, N. R. Jennings, M. Luck, S. Chalmers, N. Oren, T. Norman, A. Preece, P. Gray, G. Shercliff, P. Stockreisser, J. Shao, W. Gray, N. Fiddian, and S. Thompson. Agent-based virtual organisations for the Grid. *Multiagent and Grid Systems*, 1(4):237–249, 2005.
23. T. Sandholm, K. Larson, M. Andersson, O. Shehory, and F. Tohme. Coalition Structure Generation with Worst Case Guarantees. *Artificial Intelligence*, 111(1–2):209–238, 1999.
24. T. Sandholm and V. Lesser. Coalitions Among Computationally Bounded Agents. *Artif. Intelligence*, 94(1), 1997.
25. L. Shapley. A Value for n-Person Games. In H. Kuhn and A. Tucker, editors, *Contributions to the Theory of Games II*, pages 307–317. Princeton University Press, Princeton, 1953.
26. L. Shapley. On Balanced Sets and Cores. *Naval Research Logistics Quarterly*, 14:453–460, 1967.
27. L. Shapley. Cores of Convex Games. *International Journal of Game Theory*, 1:11–26, 1971.
28. O. Shehory and S. Kraus. Formation of overlapping coalitions for precedence-ordered task-execution among autonomous agents. In *Proc. of the 2nd Intern. Conference on Multi-Agent Systems (ICMAS-96)*, pages 330–337, 1996.
29. J. Suijs and P. Borm. Stochastic cooperative games: superadditivity, convexity and certainty equivalents. *Journal of Games and Economic Behavior*, 27:331–345, 1999.
30. J. Suijs, P. Borm, A. D. Wagenaere, and S. Tijs. Cooperative games with stochastic payoffs. *European Journal of Operational Research*, 113:193–205, 1999.
31. A. Tanin. On the core of network synthesis games. *Mathematical Programming*, 50:123–135, 1991.
32. J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.