Stability via Convexity and LP Duality in OCF games

Yair Zick
School of Physical and Mathematical Sciences
Nanyang Tech. Univ., Singapore
yair0001@ntu.edu.sg

Evangelos Markakis
Department of Informatics
Athens Univ. of Econ. and Business
markakis@gmail.com

Edith Elkind
School of Physical and Mathematical Sciences
Nanyang Tech. Univ., Singapore
eelkind@ntu.edu.sg

Abstract

The *core* is a central solution concept in cooperative game theory, and therefore it is important to know under what conditions the core of a game is guaranteed to be non-empty. Two notions that prove to be very useful in this context are Linear Programming (LP) duality and convexity. In this work, we apply these tools to identify games with overlapping coalitions (OCF games) that admit stable outcomes. We focus on three notions of the core defined in (Chalkiadakis et al. 2010) for such games, namely, the conservative core, the refined core and the optimistic core. First, we show that the conservative core of an OCF game is non-empty if and only if the core of a related classic coalitional game is non-empty. This enables us to improve the result of (Chalkiadakis et al. 2010) by giving a strictly weaker sufficient condition for the non-emptiness of the conservative core. We then use LP duality to characterize OCF games with non-empty refined core; as a corollary, we show that the refined core of a game is non-empty as long as the superadditive cover of its characteristic function is convex. Finally, we identify a large class of OCF games that can be shown to have a non-empty optimistic core using an LP-based argument.

1 Introduction

Cooperative game theory studies settings where a set of players $N = \{1, \dots, n\}$ splits into coalitions in order to generate revenue, which is then divided among the members of each coalition. Different ways of defining acceptable profit-sharing schemes are known as *solution concepts*. A very appealing solution concept is the *core*: this is the set of payoff divisions such that the total payment to every subset of players, $S \subseteq N$, is at least the amount $v(S)$ that this set can make on its own. Arguably, the notion of the core captures our intuition about what constitutes a stable outcome; it is often the case though that a game has an empty core, i.e., there is simply no way to satisfy the demands of all groups of players without breaking the budget. Thus, it is important to identify constraints on the *characteristic function* $v : 2^N \to \mathbb{R}_+$ that ensure core non-emptiness. This question was first addressed by Bondareva (1963) and Shapley (1967) who (independently) gave a characterization of games with a non-empty core, which was based on LP duality (Schrijver 1986); however, their criterion does not typically provide any intuitive structural information about such games. Subsequently, a weaker sufficient condition for core non-emptiness was identified (Shapley 1971): a game has a non-empty core as long as it is *convex*, i.e., a player's marginal contribution to a coalition grows when more agents join it.

The goal of this paper is to extend this line of work to settings where players may form *partial*, or *overlapping*, coalitions. Recall that in classic cooperative games, a player either belongs to a coalition or is not involved in it at all, i.e., when players split into coalitions, they simply form a partition of N. In contrast, in *overlapping coalition formation (OCF) games* (Chalkiadakis et al. 2010; Zick and Elkind 2011) players may divide their efforts among several teams, i.e., form coalitions that overlap. Such *partial coalitions* can be viewed as vectors in $[0, 1]^n$ (where the i-th entry indicates the fraction of resources of the i-th player dedicated to this partial coalition), and an OCF game is specified by a mapping $v : [0, 1]^n \to \mathbb{R}_+$. An outcome of an OCF game is a list of partial coalitions, and, for each coalition $\mathbf{c}$, a payoff vector that describes how the players involved in $\mathbf{c}$ split the amount $v(\mathbf{c})$ they have earned together.

In this model, when a set of players deviates from an existing outcome, it may keep some of its agreements intact, while breaking away from others. Chalkiadakis et al. (2010) and Zick and Elkind (2011) demonstrate that stability in OCF games depends on the other players' reaction to deviation, which leads to a hierarchy of possible definitions for the core. For instance, *conservative* deviators expect nothing from any coalition, regardless of the damage they caused to it; this approach gives rise to the *conservative core*. The smaller *refined core* corresponds to players expecting to get payoffs from the coalitions that they did not withdraw resources from, whereas *optimistic* deviators expect to keep their payoffs, subject to covering the marginal damage caused by their actions; this results in the—even smaller—*optimistic core*. Other notions of the core are also possible and have been studied by Zick and Elkind (2011).

Conditions on the function v that ensure non-emptiness of the conservative core have been formulated in (Chalkiadakis et al. 2010). Specifically, Chalkiadakis et al. propose a notion of convexity for OCF games, which we call *OCF-convexity*, and show that (under mild assumptions on the function v) the conservative core of OCF-convex games is not empty. They also provide a characterization result that is reminiscent of the Bondareva–Shapley theorem. In our work, we first show that the conservative core of an OCF game is non-empty if and only if the core of a related classic coalitional game (without overlaps) is non-empty. Based on this, we give a sufficient condition for the non-emptiness of the conservative core that is strictly weaker than OCF-convexity. This result demonstrates that the additional expressive power of the OCF formalism stems from its ability to describe the relationships between deviators and non-deviators that survive the deviation—an ability that is not afforded by the language of the classic cooperative games.

We then build on the LP-based argument of (Chalkiadakis et al. 2010) to characterize OCF games with a non-empty refined core. This enables us to identify a sufficient condition for the non-emptiness of the refined core: it turns out that it suffices to require that the *superadditive cover* of v is convex. While one can further extend the LP-based line of reasoning to characterize games with non-empty optimistic core, the resulting characterization does not appear to provide useful insights into the structure of such games. Therefore, instead of pursuing this approach, we describe a class of OCF games, which we call *linear bottleneck games* (LBGs), that are guaranteed to have a non-empty optimistic core. These games are motivated by a wide range of settings, including multicommodity flow games, cover games and other combinatorial optimization scenarios. Interestingly, for LBGs, the proof of stability with respect to optimistic deviations proceeds via an application of LP duality. Thus, LP-based techniques prove to be very useful for showing core stability results in OCF games.

Related Work Our work investigates conditions for core non-emptiness in the OCF domain, applying LP-based techniques and several notions of convexity for OCF games. Thus, our model is based on the work of Chalkiadakis et al. (2010) and Zick and Elkind (2011). Convexity in cooperative games (Shapley 1971) is a well-explored concept that has been applied to several variants of cooperative games. (Suijs and Borm 1999) define a notion of convexity for stochastic cooperative games, where the value of cooperation is determined by a probability distribution, while (Brânzei, Dimitrov, and Tijs 2003) define convexity for fuzzy cooperative games (Aubin 1981). The use of LP duality for constructing elements in the core dates back to (Shapley and Shubik 1972) and (Owen 1975). Later on, (Deng, Ibaraki, and Nagamochi 1999) provided a general framework for this approach, applicable to various combinatorial optimization games (in particular, games where $v(S)$ is derived by solving a linear program of a specific form).

2 Preliminaries

We now present the model and definitions necessary for our work. Throughout the paper, we use boldface lowercase characters to denote vectors and capital letters for sets. Given a subset S of $\{1, \dots, n\}$, its *indicator vector* $\mathbf{e}^S \in \mathbb{R}^n$ has 1 in its i-th coordinate if $i \in S$ and 0 otherwise.

Cooperative Games with Overlapping Coalitions

A (classic) *cooperative game* is a pair $\mathcal{G} = (N, u)$, where $N = \{1, \dots, n\}$ is the set of *players* and $u : 2^N \to \mathbb{R}_+$ is the *characteristic function*; note that u is defined on subsets of N. In contrast, in a *cooperative games with overlapping coalitions* (overlapping coalition formation, or OCF games) the characteristic function is defined on vectors in $[0, 1]^n$. That is, an OCF game is a pair $\mathcal{G} = (N, v)$, where $N = \{1, \dots, n\}$ is the set of *players* and $v : [0, 1]^n \to \mathbb{R}_+$ is the *characteristic function*; we assume $v(0^n) = 0$. In OCF games, each player possesses a single finite resource; players may collaborate by contributing some fractional amount of their resources to the completion of a task, whose value is determined by v. Thus, a *coalition* in the OCF setting can be identified with a vector $\mathbf{c} \in [0, 1]^n$, whose i-th coordinate, c^i, is the fraction of i's resources devoted to this coalition. If v is not superadditive, it is beneficial for players to divide their resources and form multiple coalitions. In that case, the resulting *coalition structure* $CS = (\mathbf{c}_1, \dots, \mathbf{c}_k)$ generates a total revenue of $\sum_{j=1}^k v(\mathbf{c}_j)$, which we denote by $v(CS)$. Given a coalition structure $CS = (\mathbf{c}_1, \dots, \mathbf{c}_k)$, we write $\mathbf{w}(CS) = \sum_{j=1}^k \mathbf{c}_j$; this vector measures the amount of resources each player devotes to CS. If a coalition $\mathbf{c}$ is listed in CS, we write $\mathbf{c} \in CS$, and if CS' is a sublist of CS, we write $CS' \subseteq CS$.

Clearly, no player can contribute more than 100% of his resources to CS; thus $\mathbf{w}(CS) \leq \mathbf{e}^N$ for any CS. The set of all coalition structures over a subset $S \subseteq N$ is denoted by $\mathcal{CS}(S)$. Given a set S and a coalition structure $CS \in \mathcal{CS}(N)$, we denote by $\mathbf{w}_S(CS)$ the total weight of the members of S in CS; that is, $\mathbf{w}_S(CS)$ equals $\mathbf{w}(CS)$ on all coordinates $i \in S$ and is 0 on all coordinates $i \notin S$.

Players are naturally interested in finding a coalition structure CS^* that maximizes their revenue, i.e., satisfies $v(CS^*) \geq v(CS)$ for all $CS \in \mathcal{CS}(N)$. Given a function $v : [0, 1]^n \to \mathbb{R}_+$, we define $v^* : [0, 1]^n \to \mathbb{R}_+$ as

$$v^*(\mathbf{c}) = \sup\{v(CS) \mid \mathbf{w}(CS) \leq \mathbf{c}\};$$

v^* is called the *superadditive cover* of v. This definition extends the notion of the superadditive cover in the classic setting, introduced by (Aumann and Drèze 1974). Throughout the paper, we assume that v has the *efficient coalition structure property*: for every $\mathbf{c} \in [0, 1]^n$ there is some coalition structure CS such that $\mathbf{w}(CS) \leq \mathbf{c}$ and $v(CS) = v^*(\mathbf{c})$. This can be ensured by relatively mild assumptions on v, such as the ones made in (Chalkiadakis et al. 2010).

Having formed a coalition structure $CS = (\mathbf{c}_1, \dots, \mathbf{c}_k)$, the players must decide how to divide the payoffs. A payoff division for CS is represented by a list of vectors $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_k)$, where $\mathbf{x}_j = (x_j^1, \dots, x_j^n)$ for $j = 1, \dots, k$. We require that under $\mathbf{x}$ the value $v(\mathbf{c}_j)$ is divided among the players in the *support of* $\mathbf{c}_j$, i.e., the set $supp(\mathbf{c}_j) = \{i \in N \mid c_j^i > 0\}$. Also, payoff divisions should be *individually rational*: no player gets less than what he can make on his own. Payoff divisions that satisfy both conditions are called *imputations*. Given a coalition structure CS, we denote the set of imputations for CS by $I(CS)$. Given

$\mathbf{x} \in I(CS)$, the pair $(CS, \mathbf{x})$ is called an *outcome*. We denote by $p_i(CS, \mathbf{x})$ the total payoff of player i under $(CS, \mathbf{x})$: $p_i(CS, \mathbf{x}) = \sum_{j=1}^{k} x_j^i$. We extend this notation to sets by writing $p_S(CS, \mathbf{x}) = \sum_{i \in S} p_i(CS, \mathbf{x})$ for $S \subseteq N$.

We denote by $\mathcal{F}(S)$ the set of all outcomes $(CS, \mathbf{x})$ with $CS \in \mathcal{CS}(S)$. $\mathcal{F}(S)$ consists of all possible ways that players in S can form coalitions and divide the resulting payoffs.

Stability and Arbitration in OCF Games

The *core* of a classic cooperative game (N, u) is the set of all vectors $\mathbf{x} = (x_1, \dots, x_n)$ with $\sum x_i = u(N)$ such that $\sum_{i \in S} x_i \geq u(S)$ for all $S \subseteq N$. This simply means that $\mathbf{x}$ is resistant to set deviations; no set $S \subseteq N$ can improve its payoff by deviating. However, when defining a similar notion for OCF games, one must take into account the fact that in OCF games, sets can deviate in a more complex manner.

We will first formalize the notion of deviation in OCF games. Given a coalition structure $CS = (\mathbf{c}_1, \dots, \mathbf{c}_k)$ and a subset $S \subseteq N$, let $CS|_S = \{\mathbf{c}_j \in CS \mid supp(\mathbf{c}_j) \subseteq S\}$. $CS|_S$ contains all coalitions fully controlled by S, and modifying coalitions outside of $CS|_S$ may affect members of $N \setminus S$. A *deviation* of S from $(CS, \mathbf{x}) \in \mathcal{F}(N)$ can be specified by the amount of resources that each $i \in S$ withdraws from every coalition $\mathbf{c}_j \notin CS|_S$. Having withdrawn resources from coalitions in CS, S can proceed to combine them with the resources it has invested in $CS|_S$ in order to maximize its own revenue. That is, if a set S deviates by withdrawing a weight of $\mathbf{d}_j$ from each coalition $\mathbf{c}_j \notin CS|_S$ (where $\mathbf{d}_j$ is a vector such that $\mathbf{d}_j \leq \mathbf{c}_j$ and $\mathbf{d}_j \leq \mathbf{e}^S$), it can now earn $v^*(\mathbf{w}(CS|_S) + \sum_{\mathbf{c}_j \notin CS|_S} \mathbf{d}_j)$ on its own. Whether this deviation is beneficial for S may depend on whether S is able to keep the payoff from its surviving coalitions with players in $N \setminus S$.

Example 2.1. Consider two players, Alice and Bob. If they both contribute half of their resources, they can complete a task T that is worth \$2; multiple copies of T can be completed. Alice can also use half of her resources to complete a task t, worth \$1. It is thus optimal for Alice and Bob to complete two copies of T together; now they must decide on a payoff division. Suppose they agree that Alice receives \$2 for the first copy of T, but \$0 for the other copy. While Alice is receiving as much as she can make on her own, she may want to deviate from the second copy of T. This, of course, depends on Bob's reaction to Alice's deviation, i.e., on whether Bob decides to punish Alice for her deviation and not let her keep \$2 from the first copy of T.

In Example 2.1, Alice considers deviating from some of the coalitions she participates in. However, her decision crucially depends on Bob's reaction to her deviation. This is pointed out in (Chalkiadakis et al. 2010), who propose three ways in which players react to deviation; (Zick and Elkind 2011) place them within a more general framework of *arbitration functions*. Let ρ_j denote the payoff that S receives from the coalition $\mathbf{c}_j \notin CS|_S$. Under the *conservative* arbitration function, $\rho_j = 0$ for all coalitions. Under the *refined* arbitration function, $\rho_j = \sum_{i \in S} x_j^i$ if $\mathbf{d}_j = 0^n$, and is 0 otherwise; that is, if $\mathbf{c}_j$ is unchanged by S, then S may keep its original payoff from $\mathbf{c}_j$, and it gets nothing if it changes $\mathbf{c}_j$ in any way. Under the *optimistic* arbitration function, $\rho_j = \max\{0, v(\mathbf{c}_j - \mathbf{d}_j) - \sum_{i \notin S} x_j^i\}$; that is, S may withdraw resources from $\mathbf{c}_j$ as long as it pays the players in $N \setminus S$ the same amount that they got under $(CS, \mathbf{x})$. An outcome $(CS, \mathbf{x})$ belongs to *conservative* (resp., *refined*, *optimistic*) core if there does not exists a set S such that each $i \in S$ gets more that $p_i(CS, \mathbf{x})$ by deviating, assuming that players in S form a coalition structure over S and their payoffs from coalitions with players in $N \setminus S$ are given by the conservative (resp., refined, optimistic) arbitration function.

3 Conservative Core

Given an OCF game $\mathcal{G} = (N, v)$, let us define the *discrete superadditive cover* of v as the function $U_v : 2^N \to \mathbb{R}_+$ given by $U_v(S) = v^*(\mathbf{e}^S)$. Note that, unlike v^*, this function is defined on subsets of N rather than arbitrary vectors in $[0, 1]^n$. The function U_v gives rise to a (non-overlapping) superadditive game $\widehat{\mathcal{G}} = (N, U_v)$. We will now argue that the conservative core of $\mathcal{G}$ is non-empty if and only if the core of $\widehat{\mathcal{G}}$ is non-empty. Thus, if one assumes a worst-case reaction to deviation, stability issues in OCF can be analyzed within the non-OCF framework.

Theorem 3.1. *The conservative core of $\mathcal{G}$ is non-empty if and only if the core of $\widehat{\mathcal{G}}$ is non-empty.*

Proof. First, suppose that the conservative core of (N, v) is not empty. As shown in (Chalkiadakis et al. 2010), this means that there is some outcome $(CS, \mathbf{x}) \in \mathcal{F}(N)$ such that for every subset $S \subseteq N$ we have $p_S(CS, \mathbf{x}) \geq v^*(\mathbf{e}^S)$. Set $x_i = p_i(CS, \mathbf{x})$ for each $i \in N$; the payoff division $(x_1, \dots, x_n)$ is in the core of (N, U_v).

For the converse direction, we use a graph-theoretic argument inspired by the one in (Chalkiadakis et al. 2010). Suppose that the core of (N, U_v) contains a vector $(x_1, \dots, x_n)$. By the efficient coalition structure property, there exists a coalition structure $CS \in \mathcal{CS}(N)$ such that $v(CS) = U_v(N)$. By Theorem 1 in (Chalkiadakis et al. 2010), it suffices to show that there exists a $\mathbf{y} \in I(CS)$ such that $p_i(CS, \mathbf{y}) = x_i$ for all $i \in N$. Given an imputation $\mathbf{y} \in I(CS)$, we build a directed graph $\Gamma(\mathbf{y}) = (N, E)$, which contains an edge (i, j) if and only if there exists a $\mathbf{c}_k \in CS$ such that $i, j \in supp(\mathbf{c}_k)$ and $y_k^i > 0$. Hence, an edge (i, j) indicates that i can transfer some payoff to j. We say that a vertex i of $\Gamma(\mathbf{y})$ is *green* if $p_i(CS, \mathbf{y}) > x_i$, *white* if $p_i(CS, \mathbf{y}) = x_i$ and *red* if $p_i(CS, \mathbf{y}) < x_i$.

Let $F : I(CS) \to \mathbb{R}_+$ be defined as follows:

$$F(\mathbf{z}) = \sum_{i=1}^{n} \min\{0, x_i - p_i(CS, \mathbf{z})\}.$$

F is a continuous function over a compact set, and thus attains its minimal value on $I(CS)$ at some point $\mathbf{y} \in I(CS)$. Among all minima of F, pick one for which the number of white vertices in $\Gamma(\mathbf{y})$ is minimal, and denote it by $\mathbf{y}_0$. Set $\Gamma = \Gamma(\mathbf{y}_0)$ and let $p_i = p_i(CS, \mathbf{y}_0)$.

Observe that if no vertex of Γ is red, we are done. Thus, suppose that Γ has a red vertex; since $\sum_{i \in N} p_i = v(CS) =$

$U_v(N) = \sum_{i\in N} x_i$, this implies that Γ also has a green vertex. Now, suppose that there is an edge connecting a green vertex i with a red or white vertex j. Then we can modify $\mathbf{y}_0$ by making i transfer some payoff $\varepsilon < p_i - x_i$ to j. This transfer will not alter i's contribution to F. Further, if j was white, it now becomes green, and if j was red, this lowers its contribution to F. In both cases, we get a contradiction with our choice of $\mathbf{y}_0$. Thus, no such edge exists in Γ, i.e., green vertices receive no payoff from coalitions they form with white or red vertices. Let us denote the non-green vertices by NG and the green vertices by G. The argument above shows that $p_G(CS, \mathbf{y}_0) = p_G(CS|_G, \mathbf{y}_0) = v(CS|_G) \le v^*(\mathbf{e}^G)$. On the other hand, $p_G(CS, \mathbf{y}_0) > \sum_{i\in G} x_i \ge v^*(\mathbf{e}^G)$, which is a contradiction unless $G = \emptyset$. But then the set of red vertices is empty, too, so we are done. □

In fact, the proof of Theorem 3.1 shows a stronger claim: for every stable payoff vector $\mathbf{p} = (p_1, \dots, p_n)$ for U_v and every optimal coalition structure CS, there is an outcome $(CS, \mathbf{x})$ such that $p_i(CS, \mathbf{x}) = p_i$ for all $i \in N$. This is an extension of a result in (Aumann and Drèze 1974), where it is shown that the core of a non-OCF game $\mathcal{G} = (N, u)$ with coalition structures is the space $\mathcal{CS} \times \mathcal{I}$, where $\mathcal{CS}$ is the set of all optimal coalition structures and $\mathcal{I}$ is the set of all stable payoff divisions for the superadditive cover of u.

Convexity and Non-emptiness of the Conservative Core

Recall that the core of a classic coalitional game $\mathcal{G} = (N, u)$ is non-empty as long as u is *supermodular*[1], i.e., for all $S \subseteq T \subseteq N$ and all $R \subseteq N \setminus T$, we have $u(S \cup R) - u(S) \le u(T \cup R) - u(T)$. Combining this fact with Theorem 3.1, we get a sufficient condition for the non-emptiness of the conservative core.

Corollary 3.2. *Consider an OCF game $\mathcal{G} = (N, v)$. If U_v is supermodular, then the conservative core of $\mathcal{G}$ is non-empty.*

Corollary 3.2 can be compared with the sufficient condition for the non-emptiness of the conservative core given in (Chalkiadakis et al. 2010). Specifically, (Chalkiadakis et al. 2010) define a notion of convexity for OCF games and show that OCF games that are convex in this sense have a non-empty conservative core. We will now argue that the notion of convexity defined in (Chalkiadakis et al. 2010) implies supermodularity of the discrete superadditive cover, while the converse is not true. Hence, Corollary 3.2 strengthens the result of (Chalkiadakis et al. 2010): there are OCF games that can be shown to have a non-empty conservative core using the former, but not the latter. We start by reproducing the definition of convexity given in (Chalkiadakis et al. 2010).

Definition 3.3 (OCF-convexity (Chalkiadakis et al. 2010))**.** An OCF game $\mathcal{G} = (N, v)$ is called *OCF-convex* if for every $S \subseteq T \subseteq N$ and every $R \subseteq N \setminus T$ the following condition holds: given outcomes $(CS_S, \mathbf{x}_S) \in \mathcal{F}(S)$, $(CS_T, \mathbf{x}_T) \in \mathcal{F}(T)$ and $(CS_{S\cup R}, \mathbf{x}_{S\cup R}) \in \mathcal{F}(S \cup R)$ such that $p_j(CS_{S\cup R}, \mathbf{x}_{S\cup R}) \ge p_j(CS_S, \mathbf{x}_S)$ for all $j \in S$, there is an outcome $(CS_{T\cup R}, \mathbf{x}_{T\cup R}) \in \mathcal{F}(T \cup R)$ such that

1. $p_j(CS_{S\cup R}, \mathbf{x}_{S\cup R}) \le p_j(CS_{T\cup R}, \mathbf{x}_{T\cup R}), \forall j \in S \cup R$;
2. $p_i(CS_T, \mathbf{x}_T) \le p_i(CS_{T\cup R}, \mathbf{x}_{T\cup R}), \forall i \in T$.

Definition 3.3 can be interpreted as follows. Given sets $S \subseteq T \subseteq N$ and a set $R \subseteq N \setminus T$, suppose that S offers R to form a coalition structure together in a way that appeals to all players in S. Then a larger set $T \supseteq S$ will always be able to offer at least as good a deal to $S \cup R$ without shortchanging its own members.

We will now show that if the OCF game (N, v) with $v : [0, 1]^n \to \mathbb{R}_+$ is OCF-convex, then $U_v : 2^N \to \mathbb{R}_+$ is supermodular. We will use the following lemma (we omit the proof, which is based on a graph-coloring argument similar to the one in the proof of Theorem 3.1).

Lemma 3.4. *Given any two sets $S \subseteq T \subseteq N$, and coalition structures $CS_T \in \mathcal{CS}(T)$, $CS_S \in \mathcal{CS}(S)$ such that $v(CS_T) = v^*(\mathbf{e}^T)$, $v(CS_S) = v^*(\mathbf{e}^S)$, there are imputations $\mathbf{x}_1 \in I(CS_S)$ and $\mathbf{y}_1 \in I(CS_T)$ such that $p_i(CS_S, \mathbf{x}_1) = p_i(CS_T, \mathbf{y}_1)$ for all $i \in S$.*

Lemma 3.4 simply states that given two socially optimal coalition structures over two subsets of agents such that $S \subseteq T \subseteq N$, it is possible to divide payoffs in such a way that agents in S receive the same payoffs in both coalition structures. Using this fact, we can now prove our initial claim.

Theorem 3.5. *If a game $\mathcal{G} = (N, v)$ is OCF-convex, then U_v is supermodular.*

Proof. We need to show that for all $S \subseteq T \subseteq N$ and all $R \subseteq N \setminus T$, we have $U_v(S \cup R) - U_v(S) \le U_v(T \cup R) - U_v(T)$. Set $S' = S \cup R$, $T' = T \cup R$, and consider coalition structures $CS_S \in \mathcal{CS}(S)$, $CS_{S'} \in \mathcal{CS}(S')$ such that $v(CS_S) = v^*(\mathbf{e}^S)$, $v(CS_{S'}) = v^*(\mathbf{e}^{S'})$. By Lemma 3.4, there are outcomes $(CS_S, \mathbf{x}_S)$ and $(CS_{S'}, \mathbf{x}_{S'})$ such that $p_i(CS_{S'}, \mathbf{x}_{S'}) = p_i(CS_S, \mathbf{x}_S)$ for all $i \in S$; thus, the total payoff to R from $(CS_{S'}, \mathbf{x}_{S'})$ is $v^*(\mathbf{e}^{S'}) - v^*(\mathbf{e}^S)$.

Consider an outcome $(CS_T, \mathbf{x}_T) \in \mathcal{F}(T)$ such that $v(CS_T) = v^*(\mathbf{e}^T)$; since $\mathcal{G}$ is OCF-convex, there is an outcome $(CS_{T'}, \mathbf{x}_{T'})$ that is better for all members of S' than $(CS_{S'}, \mathbf{x}_{S'})$, and also pays T a total of at least $v^*(\mathbf{e}^T)$. The payoff to R under $(CS_{T'}, \mathbf{x}_{T'})$ is $v(CS_{T'}) - p_T(CS_{T'}, \mathbf{x}_{T'})$. We have $v(CS_{T'}) \le v^*(\mathbf{e}^{T'})$, $p_T(CS_{T'}, \mathbf{x}_{T'}) \ge v^*(\mathbf{e}^T)$, so the payoff to R under $(CS_{T'}, \mathbf{x}_{T'})$ is at most $v^*(\mathbf{e}^{T'}) - v^*(\mathbf{e}^T)$. OCF convexity of $\mathcal{G}$ now implies $v^*(\mathbf{e}^{S'}) - v^*(\mathbf{e}^S) \le v^*(\mathbf{e}^{T'}) - v^*(\mathbf{e}^T)$. □

However, the converse of Theorem 3.5 is not true: supermodularity of U_v does not imply that $\mathcal{G}$ is OCF-convex.

Example 3.6. Consider a game $\mathcal{G} = (N, v)$ where $N = \{1, 2, 3\}$ and v is defined as follows: $v(1, 1, 0) = 5$, $v(0, 1, 1) = 4$, $v(1, 0, .5) = 7$, $v(0, 1, .5) = 3$, $v(0, 0, 1) = 1$, $v(\mathbf{c}) = 0$ for any other partial coalition $\mathbf{c}$.

It is easy to check that U_v is convex. However, $\mathcal{G}$ is not OCF-convex. Indeed, set $S = \{2\}$, $T = \{2, 3\}$, $R = \{1\}$, with $CS_S = ((0, 1, 0))$, $CS_T = ((0, 1, 1))$, $CS_{S\cup R} =$

[1] Supermodularity, for functions defined over subsets, is sometimes referred to as convexity. We prefer to use the term supermodularity to avoid confusion with the notion of convexity for functions defined on $[0, 1]^n$ used in the next section.

$((1,1,0))$. Assume that the players in T and $S \cup R$ share the payoffs according to $\mathbf{x}_T = ((0,2,2))$ and $\mathbf{x}_{S\cup R} = ((1,4,0))$, respectively. For $\mathcal{G}$ to be OCF-convex, there has to exist a coalition structure over CS_N where player 2 earns at least 4, player 1 earns at least 1, and player 3 earns at least 2. However, this is clearly impossible.

4 Refined Core

Zick and Elkind (2011) provide the following characterization of outcomes in the refined core: given a coalition structure CS, there is an imputation $\mathbf{x} \in I(CS)$ such that $(CS, \mathbf{x})$ is in the refined core if and only if $p_S(CS', \mathbf{x}) \geq v^*(\mathbf{w}_S(CS'))$ for every $S \subseteq N$ and every coalition structure $CS' \subseteq CS$ containing $CS|_S$. Using this fact, we now give a characterization of refined core non-emptiness.

Theorem 4.1. *The refined core of an OCF game (N, v) is not empty if and only if there exists a coalition structure $CS \in \mathcal{CS}(N)$ with $v(CS) = v^*(\mathbf{e}^N)$ such that*

$$\sum_{S\subseteq N} \sum_{CS|_S\subseteq CS'\subseteq CS} \delta_{S,CS'} v^*(\mathbf{w}_S(CS')) \leq v^*(\mathbf{e}^N)$$

for every collection of weights $(\delta_{S,CS'})$ such that $\sum_{S:i\in S}\sum_{CS|_S\subseteq CS'\subseteq CS} \delta_{S,CS'} = 1$.

Proof Sketch. Fix a coalition structure $CS = (\mathbf{c}_1, \ldots, \mathbf{c}_k)$, and consider the following linear program:

$$\begin{aligned} \text{min:} \quad & \sum_{j=1}^{k}\sum_{i=1}^{n} x_j^i & (1)\\ \text{s.t.} \quad & \sum_{i\in supp(\mathbf{c}_j)} x_j^i \geq v(\mathbf{c}_j) \qquad \forall j \in \{1,\ldots,k\} \\ & p_S(CS', \mathbf{x}) \geq v^*(\mathbf{w}_S(CS')), \end{aligned}$$

where the last constraint holds for all $S \subseteq N$ and all CS' such that $CS|_S \subseteq CS' \subseteq CS$. Note that we need not impose the constraints $x_j^i \geq 0$, as these are implied by the stability constraints. If the optimal solution to (1) equals $v^*(\mathbf{e}^N)$, then the refined core contains an outcome of the form $(CS, \mathbf{x})$. The dual of (1) is

$$\begin{aligned} \text{max:} \quad & \sum_{j=1}^{k} r_j v(\mathbf{c}_j) + \sum_{\substack{S\subseteq N\\ CS|_S\subseteq CS'\subseteq CS}} \delta_{S,CS'} v^*(\mathbf{w}_S(CS')) & (2)\\ \text{s.t.} \quad & r_j + \sum_{\substack{S:i\in S\\ CS|_S\subseteq CS'\subseteq CS}} \delta_{S,CS'} = 1 \\ & \delta_{S,CS'}, r_j \geq 0, \end{aligned}$$

with a constraint for each coalition $\mathbf{c}_j \in CS$ and each $i \in supp(\mathbf{c}_j)$. Observe that the dual constraints are equalities since x_j^i are unconstrained in (1). We can assume that $v(CS) = v^*(\mathbf{e}^N)$: otherwise CS is clearly unstable, as N can profitably deviate.

Let $\Gamma = (V, E)$ be a graph such that $V = \{1, \ldots, k\}$ and $(j, j') \in E$ if and only if $supp(\mathbf{c}_j) \cap supp(\mathbf{c}_{j'}) \neq \emptyset$. Note that if $(j, j') \in E$, then $r_j = r_{j'}$ for any feasible solution to (2). We partition CS into $C_1, \ldots, C_m$ according to the connected components of Γ. For each $\ell \in [m]$, we set $S_\ell = \bigcup_{\mathbf{c}\in C_\ell} supp(\mathbf{c})$. Also, we set $r_\ell = r_j$ for some $\mathbf{c}_j \in C_\ell$; note that $r_j = r_{j'}$ for all j, j' such that $\mathbf{c}_j, \mathbf{c}'_j \in C_\ell$. Since $v(CS) = v(\mathbf{e}^N)$, $v(C_\ell) = v^*(\mathbf{e}^{S_\ell})$. Thus, (2) becomes

$$\begin{aligned} \text{max:} \quad & \sum_{\ell=1}^{m} r_\ell v^*(\mathbf{e}^{S_\ell}) + \sum_{\substack{S\subseteq N\\ CS|_S\subseteq CS'\subseteq CS}} \delta_{S,CS'} v^*(\mathbf{w}_S(CS')) & (3)\\ \text{s.t.} \quad & r_\ell + \sum_{\substack{S:i\in S\\ CS|_S\subseteq CS'\subseteq CS}} \delta_{S,CS'} = 1 \qquad \forall i \in S_\ell \end{aligned}$$

Finally, observe that any optimal solution to (3) remains optimal (and feasible) if we move all the weight from r_ℓ to $\delta_{S_\ell,CS}$, i.e., increment $\delta_{S_\ell,CS}$ by r_ℓ and set r_ℓ to 0. Thus, we can ignore the variables r_ℓ. Moreover, $v^*(\mathbf{e}^N)$ is an optimal solution to (3) if and only if the value of each feasible solution is at most $v^*(\mathbf{e}^N)$: indeed, setting $\delta_{N,CS}$ to 1, and the rest of the $\delta_{S,CS'}$ to 0 gives a value of $v^*(\mathbf{e}^N)$. □

Theorem 4.1 can be interpreted as follows: given an optimal coalition structure CS, we denote by $\delta_{S,CS'}$ the probability that S will deviate from CS by withdrawing resources from CS'. If the expected social welfare from any such (randomized) deviation is at most $v^*(\mathbf{e}^N)$, then CS can be stabilized with respect to the refined arbitration function.

As a corollary of Theorem 4.1 we obtain that the refined core is non-empty as long as v^* is convex.

Corollary 4.2. *if v^* is convex, then the refined core of (N, v) is not empty.*

In fact, the proof of Corollary 4.2 (omitted due to space constraints) shows a stronger claim: if v^* is convex, then any optimal coalition structure CS admits an imputation $\mathbf{x} \in I(CS)$ such that $(CS, \mathbf{x})$ is in the refined core. This is not the case in general; there exist games where some optimal coalition structures cannot be stabilized, while others can be.

Example 4.3. Consider the following three-player game with four types of tasks. Task t_1 can be completed by player 1 alone, requires all of his resources, and is worth 5. Task t_{12} requires 50% of both player 1 and player 2's resources and is worth 10. Task T_{12} requires all of the resources from both players 1 and 2 and is worth 20. Finally, task t_{23} requires all of player 3's resources and 50% of player 2's and is worth 9. There are two optimal coalition structures: one where players 1 and 2 complete t_{12} twice (which we denote by CS), and one where players 1 and 2 complete T_{12} once (which we denote by CS'). Simply put, it is best for players 1 and 2 to work together and earn a total of 20 (by completing t_{12} twice or T_{12} once), while player 3 works alone and makes 0.

The coalition structure CS cannot be stabilized with respect to the refined arbitration function. Indeed under any imputation $\mathbf{x}$ such that $(CS, \mathbf{x})$ is in the refined core player 2 would have to get at least 8 for each copy of t_{12}. However, this means that player 1 gets at most 4 from working with player 2, while he can get 5 by working alone. In contrast, the outcome where CS' is formed and players 1 and 2 share the value of T_{12} equally is in the refined core.

5 Linear Bottleneck Games and the Optimistic Core

The LP-based approach of Section 4 can be extended to characterize the stricter notion of the optimistic core.

Indeed, an optimistic deviation of a set $S \subseteq N$ from $(CS, \mathbf{x})$ can be described by

(a) the list of coalitions $CS|_S \subseteq CS' \subseteq CS$ that S fully withdraws from, and

(b) the amount of resources each $i \in S$ withdraws from each coalition in $CS \setminus CS'$.

The payoff to S from this deviation is the profit S can make by using the resources it has withdrawn while absorbing the damage it has caused to the coalitions in $CS \setminus CS'$. To show that $(CS, \mathbf{x})$ is in the optimistic core, we need to argue that every deviation of this form is not profitable. Note that there are infinitely many ways to partially withdraw resources from coalitions in $CS \setminus CS'$. Thus, if we want to describe outcomes in the optimistic core by a linear program, we will have to specify infinitely many constraints. However, it is easy to show that many of these constraints are redundant, so an LP duality-based characterization similar to the one for the refined core can be derived.

Unfortunately, this characterization does not appear to provide useful insights into the structure of the optimistic core, nor does it lead to simple convexity-like conditions for the optimistic core non-emptiness. Thus, in this section we pursue a different approach. Namely, we define a large class of OCF games that is motivated by combinatorial optimization scenarios, and prove that these games always have a non-empty optimistic core. Moreover, we show that for games in this class an optimal coalition structure can be found using linear programming, and the dual LP solution can be used to find an imputation in the optimistic core. Our results in this section build on prior work in classic cooperative game theory, where dual solutions have been used to derive explicit payoff divisions that guarantee core stability (Deng, Ibaraki, and Nagamochi 1999; Jain and Mahdian 2007; Markakis and Saberi 2005).

Definition 5.1. A *Linear Bottleneck Game* $\mathcal{G} = (N, \omega, \mathcal{T})$ is given by a set of players $N = \{1, \ldots, n\}$, a list $\omega = (\omega^1, \ldots, \omega^n)$ of players' *weights*, and a list of *tasks* $\mathcal{T} = (T_1, \ldots, T_m)$, where each task T_j is associated with a set of players A_j who are needed to complete it, as well as a *value* $\pi_j \in \mathbb{R}_+$. We assume that $A_j \neq A_{j'}$ for $j \neq j'$, and for each $i \in N$ there is a task $T_k \in \mathcal{T}$ with $A_k = \{i\}$. The characteristic function of this game is defined as follows: given a partial coalition $\mathbf{c} \in [0,1]^n$, we set

$$v(\mathbf{c}) = \begin{cases} \pi_j \cdot \min\limits_{i \in A_j} c^i \omega^i & \text{if } supp(\mathbf{c}) = A_j \text{ for some } j \in [m] \\ 0 & \text{otherwise.} \end{cases}$$

These games are linear in the sense that the payoff earned by a partial coalition scales linearly with the smallest contribution to this coalition. The assumption that $A_j \neq A_{j'}$ for $j \neq j'$ ensures that the characteristic function is well-defined; since each player can work on his own (possibly earning a payoff of 0), all resources are used.

LBGs can be used to describe a variety of settings, including, e.g., *multicommodity flow games* (Vazirani 2001; Markakis and Saberi 2005). Briefly, in multicommodity flow games pairs of vertices in a network want to send and receive flow, which has to be transmitted by edges of the network. This setting can be modeled by a linear bottleneck game, where both vertices and edges are players. The weight of an edge player is the capacity of his edge, while the weight of a vertex player is the amount of commodity he possesses. We omit a formal description due to space constraints.

Observe that given an optimal coalition structure CS for a linear bottleneck game, we can assume without loss of generality that for every $\mathbf{c}$ in CS and every $i, k \in A_j$ we have $c^i \omega^i = c^k \omega^k$: investing more weight than one's team members does not increase the payoff from the task, so a player might as well use this weight to work alone. Also, it can be assumed that CS contains at most one coalition $\mathbf{c}$ with $supp(\mathbf{c}) = A_j$ for each $j = 1, \ldots, m$: if $supp(\mathbf{c}) = supp(\mathbf{d}) = A_j$, then $v(\mathbf{c} + \mathbf{d}) \geq v(\mathbf{c}) + v(\mathbf{d})$, so two coalitions with the same support can be merged. Thus, we can assume that in an optimal coalition structure each A_j forms at most one coalition $\mathbf{c}_j$. This coalition can be identified with the weight that each member of A_j invests in T_j, denoted W_j, i.e., $W_j = c_j^i \omega^i$ for every $i \in A_j$. Thus, an optimal coalition structure can be described by a list $W_1, \ldots W_m$, indicating how much weight is allocated to each task.

We can now write a linear program that finds an optimal coalition structure for an LBG $\mathcal{G} = (N, \omega, \mathcal{T})$:

$$\begin{aligned} \max: \quad & \textstyle\sum_{j=1}^m W_j \pi_j \\ \text{s.t.} \quad & \textstyle\sum_{j : i \in A_j} W_j \leq \omega^i \quad \forall i \in N \\ & W_j \geq 0 \qquad\qquad \forall j \in [m] \end{aligned} \tag{4}$$

The dual of LP (4) is

$$\begin{aligned} \min: \quad & \textstyle\sum_{i=1}^n \gamma^i \omega^i \\ \text{s.t.} \quad & \textstyle\sum_{i \in A_j} \gamma^i \geq \pi_j \quad \forall j \in [m] \\ & \gamma^i \geq 0 \qquad\quad \forall i \in N \end{aligned} \tag{5}$$

Let $\widehat{W_1}, \ldots, \widehat{W_m}$ and $\widehat{\gamma^1}, \ldots, \widehat{\gamma^n}$ be optimal solutions to (4) and (5), respectively. Note that optimal solutions to these linear programs exist since LP (4) is feasible and bounded. Let CS be the coalition structure that corresponds to $\widehat{W_1}, \ldots, \widehat{W_m}$. We construct a payoff vector $\mathbf{x}$ for CS as follows: for every $j = 1, \ldots, m$ we set $x_j^i = \widehat{\gamma^i} \widehat{W_j}$ if $i \in A_j$, and $x_j^i = 0$ otherwise. In words, each player i has some "bargaining power" $\widehat{\gamma^i}$, and is paid for each task he works on in proportion to his bargaining power. Note that both CS and $\mathbf{x}$ can be computed efficiently from the description of the game. We will now show that $\mathbf{x}$ is an imputation for CS, and, moreover, $(CS, \mathbf{x})$ is in the optimistic core.

Theorem 5.2. *Let $\mathcal{G} = (N, \omega, \mathcal{T})$ be a linear bottleneck game, and let CS and $\mathbf{x}$ be the coalition structure and the payoff vector constructed above. Then $\mathbf{x} \in I(CS)$ and $(CS, \mathbf{x})$ is in the optimistic core of $\mathcal{G}$.*

Proof Sketch. To see that $\mathbf{x}$ satisfies coalitional efficiency, note that the sum of payoffs from task T_j is $\sum_{i \in A_j} x_j^i =$

$\sum_{i \in A_j} \widehat{\gamma^i}\widehat{W_j} = \widehat{W_j} \sum_{i \in A_j} \widehat{\gamma^i}$. As $\widehat{\gamma^1}, \dots, \widehat{\gamma^n}$ is an optimal solution to (5), we have either $\sum_{i \in A_j} \widehat{\gamma^i} = \pi_j$ or $\widehat{W_j} = 0$. Thus, for any task T_j that is actually executed (i.e., $\widehat{W_j} > 0$), its total payoff $\pi_j \widehat{W_j}$ is shared by players in A_j. We will now give an outline of the proof that $(CS, \mathbf{x})$ is in the optimistic core; this will also imply that $\mathbf{x}$ is individually rational. Consider a deviation from $(CS, \mathbf{x})$ by a set S. This deviation can be described by (a) a list of tasks that S abandons completely, and (b) the amount of weight that players in S withdraw from all other tasks. By deviating, S loses all of the payoff it was getting from tasks in (a), and has to assume the marginal damage for each task in (b). This loss needs to be compared to the payoff that S earns by optimally using the withdrawn resources. The latter is given by a linear program that corresponds to an LBG on S. Finally, we consider the dual of this linear program, observe that $\widehat{\gamma^1}, \dots, \widehat{\gamma^n}$ is a feasible solution for it, and use LP duality. □

As shown in (Zick and Elkind 2011), the optimistic core is contained in any other arbitrated core. Thus, Theorem 5.2 implies that the arbitrated core of an LBG is not empty for any arbitration function.

Representing multicommodity flow games as LBGs leads to an exponential blowup in representation size (there is a task for each path in the network); however, the construction above can be used to find an outcome in the optimistic core of a multicommodity flow game in polynomial time, since the resulting primal LP admits an efficient separation oracle.

6 Conclusions and Future Work

One of the main contributions of this paper is the equivalence between the conservative core of an OCF game and the core of its discrete superadditive cover. While one might argue that this makes the notion of conservative core redundant, we do not think that this is the case, especially if one adopts an algorithmic perspective. Indeed, the proof of Theorem 3.1 is non-constructive, and does not provide an efficient way of finding an outcome in the conservative core of (N, v) given an outcome in the core of (N, U_v). Further, computing U_v may be a difficult task in itself. Nevertheless, it is fair to say that to utilize the full expressive power of the arbitrated OCF model one needs to consider arbitration functions that are more permissive than the conservative arbitrator.

Throughout the paper, we focus on the three arbitration functions described in (Chalkiadakis et al. 2010); however, it would be interesting to see what Bondareva–Shapley-like conditions can be formulated for other arbitration functions. An initial observation is that for general arbitration functions it may be impossible to describe the conditions for core non-emptiness by a linear program, since the arbitration function itself may be non-linear.

Our work examines three different notions of convexity for OCF games (namely, OCF-convexity, supermodularity of U_v, and convexity of v^*) and partially describes their relationship. However, the study of convexity and its importance for stability of OCF games is far from complete. In particular, it is not clear whether OCF-convexity implies convexity of v^*. Also, it would be useful to find a convexity condition that implies optimistic core non-emptiness.

Finally, we analyze Linear Bottleneck Games from an OCF perspective. We believe that viewing LBGs—and the myriad fractional combinatorial optimization games that they represent—as OCF games, provides a more faithful picture of the interactions than can be obtained within the classic coalitional games model. Our strong positive result showing that any such game has a non-empty optimistic core indicates that even the most stringent version of OCF stability can be achieved in real-life settings.

Acknowledgements. The first author was supported by SINGA graduate fellowship. The second author was partially supported by the project AGT of the action THALIS (co-financed by the EU and Greek national funds). The third author is supported by National Research Foundation (Singapore) under Research Fellowship NRF2009-08.

References

Aubin, J. 1981. Cooperative fuzzy games. *Mathematics of Operations Research* 6(1):1–13.

Aumann, R., and Drèze, J. 1974. Cooperative games with coalition structures. *International Journal of Game Theory* 3:217–237.

Bondareva, O. 1963. Some applications of linear programming methods to the theory of cooperative games. *Problemy kibernetiki* 10:119–139.

Brânzei, R.; Dimitrov, D.; and Tijs, S. 2003. Convex fuzzy games and participation monotonic allocation schemes. *Fuzzy Sets and Systems* 139:267–281.

Chalkiadakis, G.; Elkind, E.; Markakis, E.; Polukarov, M.; and Jennings, N. 2010. Cooperative games with overlapping coalitions. *Journal of Artificial Intelligence Research* 39:179–216.

Deng, X.; Ibaraki, T.; and Nagamochi, H. 1999. Algorithmic aspects of the core of combinatorial optimization games. *Mathematics of Operations Research* 24(3):751–766.

Jain, K., and Mahdian, M. 2007. Cost sharing. In Nisan, N.; Roughgarden, T.; Tardas, E.; and Vazirani, V., eds., *Algorithmic Game Theory*. Cambridge University Press. chapter 15, 383–408.

Markakis, E., and Saberi, A. 2005. On the core of the multicommodity flow game. *Decision support systems* 39(1):3–10.

Owen, G. 1975. On the core of linear production games. *Mathematical Programming* 9:358–370.

Schrijver, A. 1986. *Theory of Linear and Integer Programming*. Wiley, Chichester.

Shapley, L., and Shubik, M. 1972. The assignment game 1: The core. *International Journal of Game Theory* 1:111–130.

Shapley, L. 1967. On balanced sets and cores. *Naval Research Logistics Quarterly* 14(4):453–460.

Shapley, L. 1971. Cores of convex games. *International Journal of Game Theory* 1:11–26.

Suijs, J., and Borm, P. 1999. Stochastic cooperative games: Superadditivity, convexity, and certainty equivalents. *Games and Economic Behavior* 27(2):331–345.

Vazirani, V. 2001. *Approximation algorithms*. Springer Verlag.

Zick, Y., and Elkind, E. 2011. Arbitrators in overlapping coalition formation games. In *AAMAS 2011*, 55–62.