Arbitration and Stability in Cooperative Games

YAIR ZICK
Nanyang Technological University
and
EDITH ELKIND
University of Oxford

We present a formal framework for handling deviation in settings where players divide resources among multiple projects, forming overlapping coalition structures. Having formed such a coalition structure, players share the revenue generated among themselves. Given a profit division, some players may decide that they are underpaid, and deviate from the outcome. The main insight of the work presented in this survey is that when players want to deviate, they must know how the non-deviators would react to their deviation: after the deviation, the deviators may still work with some of the non-deviators, which presents an opportunity for the non-deviators to exert leverage on deviators. We extend the *overlapping coalition formation (OCF)* model of Chalkiadakis et al. [2010] for cooperation with partial coalitions, by introducing *arbitration functions*, a general framework for handling deviation in OCF games. We review some interesting aspects of the model, characterizations of stability in this model, as well as methods for computing stable outcomes.

Categories and Subject Descriptors: J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*; I.2.11 [**Computing Methodologies**]: Artificial Intelligence—*Distributed Artificial Intelligence*; F.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity

General Terms: Algorithms; Economics; Theory

Additional Key Words and Phrases: Overlapping Coalitions, Arbitration Functions, Stability

1. INTRODUCTION

Consider a group of players that possesses commodities of various types, and can generate revenue by combining fractions of their resources. Having divided their resources and generated profits, players must now agree on some reasonable manner in which to split the revenue among themselves. The problem of revenue sharing among collaborative entities is often modeled using *cooperative game theory* [Peleg and Sudhölter 2007]; however, classic cooperative game theory assumes that each player can only join a single coalition, to which it fully allocates its resources, i.e., players form a coalition structure by splitting into disjoint groups. In the setting we describe above, players may join several coalitions, allocating fractions of their resources to several projects.

Chalkiadakis et al. [2010] propose a model for settings where players divide resources among several projects. Formally, given a set of players $N = \{1, \dots, n\}$, a coalition in an *overlapping coalition formation (OCF) game* is a vector $\mathbf{c} \in [0, 1]^n$, where c_i describes the percentage of player i's resources that are committed to the coalition $\mathbf{c}$. The value of a coalition $\mathbf{c}$ is denoted $v(\mathbf{c})$, i.e., the revenue that can be generated if players contribute shares of their resources as per $\mathbf{c}$ is given by a

Authors' addresses: `yair0001@ntu.edu.sg`, `elkind@cs.ox.ac.uk`,

characteristic function $v : [0,1]^n \to \mathbb{R}$. Thus, an *OCF game* is a tuple $\mathcal{G} = \langle N, v\rangle$. This definition generalizes classic cooperative games, where a coalition is a subset of players $S \subseteq N$, and the characteristic function $u : 2^N \to \mathbb{R}$ is from subsets of N to $\mathbb{R}$.

When players divide resources among several coalitions, they form a *coalition structure*. A coalition structure CS is a list of coalitions $(\mathbf{c}_1, \dots, \mathbf{c}_m)$, such that for all $i \in N$, and for all $\mathbf{c} \in CS$, c_i is the amount of resources that player i allocates to the coalition $\mathbf{c}$; naturally, we require that every player contributes at most 100% of his resources, so $\sum_{\mathbf{c}\in CS} c_i \le 1$.

Having formed a coalition structure CS, the players must agree on a way of dividing payoffs from coalitions in CS; such a payoff division, $\mathbf{x}$, consists of a list of vectors $(\mathbf{x}(\mathbf{c}))_{\mathbf{c}\in CS}$, such that $\sum_{i=1}^n x_i(\mathbf{c}) = v(\mathbf{c})$ and if $c_i = 0$ then $x_i(\mathbf{c}) = 0$. That is, the profits of a coalition $\mathbf{c}$ should be distributed within the *support* of $\mathbf{c}$, i.e., the set of players with $c_i > 0$ (denoted $supp(\mathbf{c})$). A payoff division that satisfies these properties is called an *imputation* for CS. A coalition structure-imputation pair $(CS, \mathbf{x})$ is called an *outcome*.

One is naturally interested in families of outcomes that satisfy certain desirable properties. Classic cooperative game theory offers several classes of payoff divisions, or *solution concepts*; one such class of payoff divisions is called the *core*. The core of a classic cooperative game is a payoff division such that for every subset $S \subseteq N$, the total payoff to the set S is at least the revenue that S can generate on its own. Alternatively, the core can be though of as the set of all outcomes that are resistant to deviation. A deviation can be thought of as follows: after players partitioned into disjoint coalitions (in a non-OCF game) and divided profits, a set $S \subseteq N$ would want to deviate from the resulting outcome if the members of S could work together, generate a revenue of $u(S)$, and divide this revenue among themselves so that each member of S receives strictly more than what it currently receives. If no such set exists, then the outcome is stable.

Defining stability in OCF games should follow a similar line of reasoning; an outcome $(CS, \mathbf{x})$ is stable if no set $S \subseteq N$ would want to deviate from it. However, as Chalkiadakis et al. [2010] note, deviation in OCF games is a complex matter. Consider a simple exchange market, where a single seller s provides some commodity to two buyers, b_1 and b_2. The seller s has 100kg of sugar; he agrees to sell 40kg to b_1 and 60kg to b_2, and does so with a uniform price of 5$ per kg, for a total revenue of 200$ from b_1 and 300$ from b_2. Suppose that another buyer, b_3, wants 30kg of sugar, offering 7$ per kg. Here, s would want to withdraw 30kg of sugar that were committed to b_1 and b_2, and sell them to b_3. Suppose s withdraws 30kg from b_1 and sells them to b_3; the profitability of this action depends on what happens after the deviation. One option is that nothing happens; b_1 would buy the remaining 10kg of sugar from s. Alternatively, b_1 could refuse to collaborate with s if s deviates, since he feels cheated. In this case, it is still worthwhile for s to switch to working with b_3 (he earns 10$ more), albeit less so. A third possibility is that b_2 *would not work with s either*, for example, if b_1 and b_2 form a cartel. While b_2 was effectively not hurt by s's actions, he may still not wish to work with him, or at the very least threaten to do so in order to deter s from deviating.

To conclude, when assessing the desirability of deviation in OCF games, a deviating set must know how others react to the deviation.

Chalkiadakis et al. [2010] identified this interesting feature of OCF games, and introduced three possible reactions to deviation: the *conservative*, *refined*, and *optimistic* reaction. Under the conservative reaction, S may expect no payoffs from any coalition; like in the non-overlapping case, it assumes that it is “on its own” if it deviates, and assesses the desirability of deviation against the most that it can make on its own. Under the *refined* reaction, S may expect payoff from all coalitions that were not changed by the deviation; if a deviating set does not change some coalition $\mathbf{c}$ when it deviates, it keeps its original payment from $\mathbf{c}$. Finally, under the *optimistic* reaction, S may still receive payoff from a coalition $\mathbf{c}$, if it can reduce its contribution to $\mathbf{c}$ while still paying all agents in $N \setminus S$ the same amount they got from $\mathbf{c}$ under $(CS, \mathbf{x})$; in other words, S may withdraw resources from a coalition, so long as it agrees to assume the damage that its deviation caused.

2. ARBITRATION FUNCTIONS

The three possible reactions to deviation that Chalkiadakis et al. [2010] describe are by no means exhaustive: there can be many reactions to deviation, and they may be quite general in their nature. A reaction that is moderated by a contractual agreement between involved parties can be quite complex, detailing fines and rewards, depending on what changes to the original agreement are made.

Zick and Elkind [2011] propose a framework that is able to capture such general behavior using a single function that specifies what happens when a set deviates. When a set S deviates from an outcome $(CS, \mathbf{x})$ it specifies how much resources it withdraws from each coalition that is not fully controlled by it, i.e, given a coalition $\mathbf{c} \in CS$ such that $supp(\mathbf{c})$ contains non-S members, the way that S deviates from $\mathbf{c}$ is given by a vector $\delta(\mathbf{c})$ such that $\delta(\mathbf{c}) \leq \mathbf{c}$, and $supp(\delta(\mathbf{c})) \subseteq S$. The first requirement states that S cannot withdraw more from $\mathbf{c}$ than what it has invested to begin with, and the second ensures that only members of S withdraw resources from $\mathbf{c}$. The *arbitration function* [Zick and Elkind 2011] is a function α that, given *i)* an outcome $(CS, \mathbf{x})$ *ii)* a set $S \subseteq N$ *iii)* a deviation δ of S from CS and *iv)* a coalition $\mathbf{c} \in CS$ containing non-S members, assigns a payoff $\alpha(\mathbf{c}) \in \mathbb{R}$. Note that α can be any number in $\mathbb{R}$: it may be negative, i.e., a coalition may fine deviators, and it may be arbitrarily high, i.e., actively rewarding deviation.

We let S use whatever resources it withdrew from CS using δ plus whatever resources are in coalitions it fully controls to generate revenue. This revenue, plus the payoffs to S from coalitions it deviated from, given by $\alpha(\mathbf{c})$, is the total revenue S gets from deviating. Given an outcome $(CS, \mathbf{x})$ and a set S, let $\mathcal{A}^*(CS, \mathbf{x}, S)$ be the most that S can get by deviating from $(CS, \mathbf{x})$. This framework generalizes the reactions to deviation described by Chalkiadakis et al. [2010]. Under the conservative arbitration function, $\alpha_c(\mathbf{c}) \equiv 0$: S always receives nothing from non-deviators. For the refined arbitration function, $\alpha_r(\mathbf{c})$ is the payoff to S from $\mathbf{c}$ under $(CS, \mathbf{x})$ iff $\delta(\mathbf{c}) = 0^n$, i.e., if S does not change a coalition, it is allowed to keep all of its payoffs from that coalition. Finally, for the optimistic arbitration function, we have $\alpha_o(\mathbf{c}) = \sum_{i \in S} x_i(\mathbf{c}) + v(\mathbf{c} - \delta(\mathbf{c})) - v(\mathbf{c})$; that is, $\mathbf{c}$ pays S its original payoff, plus the marginal loss its deviation has caused.

Using arbitration functions, we can easily define other reactions to deviation: the *sensitive* arbitration function allows the deviators to keep payments from coalitions whose *members* were not hurt by the deviation. Formally, $\alpha_s(\mathbf{c}) = \sum_{i \in S} x_i(\mathbf{c})$ iff $\delta(\mathbf{c}') = 0^n$ for all $\mathbf{c}' \in CS$ such that $supp(\mathbf{c}') \cap supp(\mathbf{c}) \cap (N \setminus S) \neq \emptyset$. This example shows that the amount that S receives from a coalition $\mathbf{c}$ need not depend just on the effect that S had on $\mathbf{c}$.

A deviation of S from $(CS, \mathbf{x})$ is called $\mathcal{A}$-*profitable* if S can use the resources it withdrew to generate profits, and divide those profits plus the payoffs from the arbitration function so that every $i \in S$ gets strictly more than $p_i(CS, \mathbf{x})$. Finding $\mathcal{A}$-profitable deviations is a complex task: if, after deviation, S forms a coalition structure CS_d, it must choose an imputation $\mathbf{x}_d \in CS_d$; thus, payoffs to players must satisfy the no side payments rule. Similarly, $\alpha(\mathbf{c})$ can only be divided among players in S who still contribute to $\mathbf{c}$ after the deviation. It is possible that even if $\mathcal{A}^*(CS, \mathbf{x}, S)$ is strictly more than the total payoff to S under $(CS, \mathbf{x})$, there is no way for S to divide revenue from the deviation in such a way that every $i \in S$ is strictly better off. However, Zick and Elkind [2011] show the following result.

THEOREM 2.1 [ZICK AND ELKIND 2011]. *If $\mathcal{A}^*(CS, \mathbf{x}, S)$ is more than the payoff that S receives under $(CS, \mathbf{x})$, then there is a subset $S' \subseteq S$ that can $\mathcal{A}$-profitably deviate.*

Theorem 2.1 implies an equivalence between outcomes where all $S \subseteq N$ are paid at least $\mathcal{A}^*(CS, \mathbf{x}, S)$, and outcomes where no $S \subseteq N$ can $\mathcal{A}$-profitably deviate from $(CS, \mathbf{x})$. Such outcomes are called $\mathcal{A}$-*stable*; the $\mathcal{A}$-*core* of an OCF game $\mathcal{G}$ is the set of all $\mathcal{A}$-stable outcomes.

3. CHARACTERIZING $\mathcal{A}$-STABILITY IN OCF GAMES

How can we decide if an OCF game admits a stable outcome? In classic cooperative games, this question is answered by the Bondareva–Shapley theorem [Bondareva 1963; Shapley 1967]. Briefly, a collection of weights $(\delta_S)_{S \subseteq N}$ is called *balanced* if for all $i \in N$, $\sum_{S: i \in S} \delta_S = 1$, and $\delta_S \geq 0$ for all $S \subseteq N$. Bondareva [1963] and Shapley [1967] show that a cooperative game $\mathcal{G} = \langle N, u \rangle$ has a non-empty core iff for any balanced collection of weights, $\sum_{S \subseteq N} \delta_S v(S) \leq opt(\mathcal{G})$, where $opt(\mathcal{G})$ is the value of an optimal coalition structure in $\mathcal{G}$. Zick et al. [2012] show that OCF games with $\mathcal{A}$-stable outcomes admit a similar characterization.

Sometimes, stability of an OCF game can be derived from the stability of a related classic game. Specifically, given an OCF game $\mathcal{G} = \langle N, v \rangle$, the *discrete superadditive cover* of $\mathcal{G}$ is a classic game $\bar{\mathcal{G}} = \langle N, U_v \rangle$ where $U_v(S)$ is the most that the members of S can make using only their resources. Zick et al. [2012] show that the conservative core of $\mathcal{G}$ is essentially equivalent to the (classic) core of the *discrete superadditive cover* of $\mathcal{G}$.

THEOREM 3.1 [ZICK ET AL. 2012]. *Given an optimal coalition structure CS and a payoff division $\mathbf{p}$ in the core of $\bar{\mathcal{G}} = \langle N, U_v \rangle$ there exists an imputation $\mathbf{x} \in I(CS)$ such that each player $i \in N$ receives a total payoff of p_i under $(CS, \mathbf{x})$.*

Theorem 3.1 immediately implies that if the discrete superadditive cover of $\mathcal{G}$ is convex [Shapley 1971], then the conservative core of $\mathcal{G}$ is not empty. Further, all optimal coalition structures are equally easy to stabilize: if CS, CS' are two optimal coalition structures and the conservative core of $\mathcal{G}$ is not empty, there exist imputations $\mathbf{x} \in I(CS)$ and $\mathbf{x}' \in I(CS')$ such that $(CS, \mathbf{x})$ and $(CS', \mathbf{x}')$ are both in the conservative core, and the payoff to any player $i \in N$ is the same under $(CS, \mathbf{x})$ and $(CS', \mathbf{x}')$.

In contrast, for the refined core the latter property does not hold: it is possible that even if CS is an optimal coalition structure and the refined core is not empty, there is no $\mathbf{x} \in I(CS)$ such that $(CS, \mathbf{x})$ is in the refined core. This indicates that the refined core is considerably more complex than the conservative core. Zick et al. [2012] characterize refined-stable outcomes in a manner similar to the Bondareva–Shapley theorem. Using this characterization, one can show a sufficient condition for the non-emptiness of the refined core of a game.

THEOREM 3.2. *Given an OCF game $\mathcal{G} = \langle N, v \rangle$, if v^* is homogeneous of degree $k \geq 1$, then the refined core of $\mathcal{G}$ is not empty.*

In fact, one can show that when v^* is homogeneous of degree $k \geq 1$, any optimal coalition structure can be stabilized with respect to the refined arbitration function.

4. COMPUTING $\mathcal{A}$-STABLE OUTCOMES

There is a well-established body of literature studying computational aspects of cooperative games (see [Chalkiadakis et al. 2011]). Chalkiadakis et al. [2010] study some computational issues in OCF games, focusing their attention on a class of OCF games called *threshold task games*, and conservative core stability. Zick et al. [2012] study computational aspects of general classes of OCF games. They show that, while the problem of finding $\mathcal{A}$-stable outcomes is (rather unsurprisingly) NP-hard, it is possible to find $\mathcal{A}$-stable outcomes in polynomial time if players are limited in their interactions: *i*) their resources are integer weights that are polynomial in n; *ii*) they are not allowed to form large coalitions; and *iii*) their interactions are simple in structure (they form a tree, or, more generally, have bounded treewidth). Zick et al. [2012] also observe that the complexity of finding $\mathcal{A}$-stable outcomes is highly dependent on the structure of $\mathcal{A}$; if $\mathcal{A}$ takes on too complex a structure (for example, if it is given by a complicated legal contract), it is NP-hard to decide if an outcome is $\mathcal{A}$-stable. In other words, if we assume bounded rationality, then a complex arbitration function is an effective barrier to deviation.

Finally, Zick et al. [2012] study a class of games (linear bottleneck games, or LBGs) which are guaranteed to have a non-empty optimistic core; this class is a generalization of multicommodity flow games, and captures several fractional combinatorial optimization problems (such as fractional matching markets, fractional graph covering, etc.).

5. DISCUSSION

We believe that the main contribution of the works described in this survey is the introduction of a new paradigm for studying strategic interactions among rational agents. Non-deviators can also be strategic, and use their leverage in order to enforce certain outcomes. Enforcement can be either in the form of hefty penalties for deviation, or via exacting high computational costs on deviators. Non-deviator reaction has been implicitly studied before—every strategic interaction must assume something about the behavior of non-deviators—and has recently received attention outside the framework of OCF games. Brânzei et al. [2013] study different reactions to deviations in stable matching, and Ackerman and Brânzei [2012] study reactions to deviations in collaboration networks, and the Nash equilibria that result. Other research fields could also benefit from reexamining the way non-deviators react to deviations in strategic settings.

In this survey, we present a general framework for handling deviation in OCF games, and discuss algorithmic and game-theoretic properties of the resulting model. Our work can be extended in several interesting ways. First, while exact algorithms for computing solution concepts in OCF games have been studied, approximately $\mathcal{A}$-stable outcomes are also of interest; this is akin to the cost of stability in classic cooperative games [Bachrach et al. 2009]. Second, while we identify properties of OCF games that ensure stability for a given arbitration function, one can alternatively fix an OCF game $\mathcal{G}$, and identify arbitration functions that ensure that $\mathcal{G}$ is stable.

REFERENCES

Ackerman, M. and Brânzei, S. 2012. Research quality, fairness, and authorship order. *CoRR abs/1208.3391.*

Bachrach, Y., Elkind, E., Meir, R., Pasechnik, D., Zuckerman, M., Rothe, J., and Rosenschein, J. 2009. The cost of stability in coalitional games. In *Proceedings of the 2nd International Symposium on Algorithmic Game Theory (SAGT-09).* 122–134.

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

Brânzei, S., Michalak, T., Rahwan, T., Larson, K., and Jennings, N. R. 2013. Matchings with externalities and attitudes. In *Proceedings of the 12th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-13).* 295–302.

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.

Chalkiadakis, G., Elkind, E., and Wooldridge, M. 2011. *Computational Aspects of Cooperative Game Theory.* Morgan and Claypool.

Peleg, B. and Sudhölter, P. 2007. *Introduction to the Theory of Cooperative Games*, Second ed. Theory and Decision Library. Series C: Game Theory, Mathematical Programming and Operations Research, vol. 34. Springer, Berlin.

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.

Zick, Y., Chalkiadakis, G., and Elkind, E. 2012. Overlapping coalition formation games: Charting the tractability frontier. In *Proceedings of the 11th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-12).* 787–794.

Zick, Y. and Elkind, E. 2011. Arbitrators in overlapping coalition formation games. In *Proceedings of the 10th International Joint Conference on Autonomous agents and Multiagent Systems (AAMAS-11).* 55–62.

Zick, Y., Markakis, E., and Elkind, E. 2012. Stability via convexity and LP-duality in OCF games. In *Proceedings of the 26th AAAI Conference on AI (AAAI-12).* 1506–1512.