On the Structure of Synergies in Cooperative Games

Ariel D. Procaccia
Carnegie Mellon University
arielpro@cs.cmu.edu

Nisarg Shah
Carnegie Mellon University
nkshah@cs.cmu.edu

Max Lee Tucker
Carnegie Mellon University
mltucker@andrew.cmu.edu

Abstract

We investigate synergy, or lack thereof, between agents in cooperative games, building on the popular notion of Shapley value. We think of a pair of agents as synergistic (resp., antagonistic) if the Shapley value of one agent when the other agent participates in a joint effort is higher (resp. lower) than when the other agent does not participate. Our main theoretical result is that any graph specifying synergistic and antagonistic pairs can arise even from a restricted class of cooperative games. We also study the computational complexity of determining whether a given pair of agents is synergistic. Finally, we use the concepts developed in the paper to uncover the structure of synergies in two real-world organizations, the European Union and the International Monetary Fund.

1 Introduction

The European Union (EU) currently consists of 28 European countries, which frequently make joint decisions on policy issues. Decisions to adopt or reject proposals are made through a voting procedure, in which each member country is assigned a weight that is (roughly) proportional to its population; a proposal is accepted if the total weight of all countries that voted in favor is at least a fixed threshold. A country's weight determines the influence it can exert over the outcome, but provides a coarse measure of power. In a classic paper, Shapley (1952) devised a more subtle method of gauging a country's influence; in this context, his notion — known today as the *Shapley value* — informally measures how frequently a country would play a *pivotal* role in the decision making process.

More generally, *cooperative games* are used to model situations where a set of agents need to cooperate in order to accomplish a joint task. Each agent must decide whether or not to participate in the joint effort. The value of the outcome depends on the subset of agents that decide to participate. Thus, each agent's decision of whether to participate influences the final outcome to some extent, giving the agent a certain power. Even in this more general setting, the Shapley value provides a compelling proxy for an agent's power, but it is more commonly interpreted as the *payoff* an agent would receive when the value generated by the grand coalition (which consists of all agents) is split among the agents. For these reasons, the Shapley value is considered to be a useful tool for the design of multiagent systems. Indeed, cooperative games and the Shapley value have been extensively studied in the AI literature (Michalak et al. 2013; Ando 2012; Aziz et al. 2009; Bachrach et al. 2008; Bachrach and Rosenschein 2009; Zick, Skopalik, and Elkind 2011; Bachrach, Parkes, and Rosenschein 2013).

Previous work by economists has established that understanding the *synergies* between agents gives an important new perspective on cooperative games. Indeed, the Nobel-winning economist Myerson (1980) modeled the structure of interactions between agents as "a series of conferences" in which agents would meet "to discuss possible cooperative plans and to sign jointly binding agreements" (Myerson 1980, page 169). He explicitly studied synergies from an axiomatic viewpoint by requiring that for two agents i and j, j's contribution to i is always equal to i's contribution to j (in a formal sense that will be crystallized shortly). Myerson leveraged this new axiom to provide an alternative characterization of the Shapley value.

We too recognize the importance of studying the structure of synergies, but while Myerson (1980) was interested in characterizing allocation rules from constraints on the structure of interactions, our research question in this paper goes in the opposite direction:

> *Which structures of synergies between agents arise when influence is measured using the Shapley value, and how computationally hard is it to uncover these structures?*

1.1 Our Approach and Results

To reason about synergy between a pair of agents (i, j), we look at the Shapley value of agent i in two situations: when agent j commits to joining the coalition, and when j commits to withdrawing from the coalition. If the Shapley value of i is higher in the former case than in the latter, we say that i and j are *synergistic* (agent i "wants" agent j to join the coalition). If the inequality holds in the opposite direction, we say that the two agents are *antagonistic*. We choose to use a definition that is slightly different from the one used by Myerson (1980), for conceptual reasons that we discuss in Section 1.2; but we emphasize here that our theoretical results also hold for Myerson's definition. Like Myerson's notion, our notion of synergy is symmetric. In fact, the difference between the Shapley values in the two scenarios mentioned above is invariant to switching the roles of i and j.

To set up our main theoretical result, we define a *synergy graph* — a complete undirected graph where the vertices are agents and each edge is labeled as synergistic or antagonistic. We show that any synergy graph can arise from a *strictly monotonic* cooperative game. This result implies that the question of whether a pair of agents is synergistic or antagonistic is, in a sense, independent of the relationships between other pairs, thereby giving rise to a rich set of possible synergy structures.

Our second theoretical result asserts that it is computationally hard to check if a given pair of agents is synergistic or antagonistic, even in weighted voting games (WVGs) — perhaps the most popular of all cooperative games.

Finally, we complement our theoretical results with experiments dealing with two real-world WVGs: decision making in the European Union (EU) and in the International Monetary Fund (IMF). We identify surprising structures among synergistic and antagonistic pairs in these domains.

1.2 Related Work

Myerson's (1980) seminal paper initiated a line of influential work building on his axiom of *balanced contributions* (i.e., the synergy relation is symmetric) to characterize the Shapley value or extensions thereof, such as the Owen value (Owen 1977); see, e.g., the work of Hart and Mas (1989), Calvo et al. (1996), and Grabisch and Roubens (1999). Myerson's notion of synergy is slightly different from the one we employ: he considers the difference between the Shapley value of agent i when agent j has not made any commitments, and the Shapley value of agent i when agent j commits to withdrawing from the coalition (we replace the former term with the Shapley value of i when j commits to joining). A disadvantage of Myerson's definition is that it depends on a comparison of the Shapley values of an agent in two games with different numbers of agents. In some games the Shapley value is inherently biased to be higher in the game with a smaller number of agents, making the comparison of the difference to zero biased. For example, Myerson's definition characterizes almost all pairs of agents in the European Union (EU) and in the International Monetary Fund (IMF) as antagonistic. That said, the theoretical differences between the two notions are minor, and (as noted above) our theoretical results hold for both definitions.

The first attempt to define and quantify synergy dates further back to Owen (1972), whose idea was then rediscovered by Murofushi and Soneda (1993) in the field of fuzzy control systems. Their measure, known as the *interaction index*, calculates the synergistic interaction between two agents by accounting for the surplus generated by their cooperation in all possible coalitions, in a manner similar to ours. Later, Grabisch and Roubens (1999) extended this definition to synergy among more than two agents. These papers have given rise to a body of literature on measuring synergy in fuzzy systems and related fields.

Synergies in cooperative games have also been studied by CS and AI researchers. Deng and Papadimitriou (1994) studied graph-based games in which given synergies between pairs of agents are used to define the game. Conitzer and Sandholm (2006) extended this approach to a more general synergy-based representation. Both papers focus on using given synergies to represent a cooperative game, and study the computational complexity of various cooperative solution concepts in terms of this game representation.

A number of recent publications in the AI literature (Voice, Ramchurn, and Jennings 2012; Bistaffa et al. 2014; Vinyals et al. 2013) have used the term "synergy graph" to denote a *communication graph* that shows connections between agents; its connected subgraphs are the feasible coalitions. In these papers, the feasible coalitions can then be assigned *any real value*. In contrast, the synergy graph in our work measures synergies between agents based on the actual values of coalitions.

2 Preliminaries

A *transferable utility cooperative game* (hereinafter simply referred to as a *game*) $G = (N, v)$ is composed of a set of agents $N = \{1, 2, \ldots, n\}$ (also called the grand coalition) and a characteristic function $v : 2^N \to \mathbb{R}_+ \cup \{0\}$ mapping any coalition (agent subset) to the utility these agents achieve together. For an agent $i \in N$ and a coalition $S \subseteq N$, we denote $S \cup \{i\}$ by $S + i$ and $S \setminus \{i\}$ by $S - i$.

A game is called *simple* if the characteristic function only takes values in $\{0, 1\}$. A game is called *monotonic* (resp. strictly monotonic) if larger coalitions have greater value, i.e., for $S \subsetneq T \subseteq N$, $v(S) \leq v(T)$ (resp. $v(S) < v(T)$).

Weighted Voting Games (WVGs). A WVG $G = (N, \mathbf{w}, t)$ is a simple game where each agent $i \in N$ has a weight $w_i > 0$, forming the weight vector $\mathbf{w}$, and a coalition $S \subseteq N$ has value 1 iff its total weight is at least the threshold t: if $\sum_{i \in S} w_i \geq t$ then $v(S) = 1$, else $v(S) = 0$.

The Shapley Value. The *Shapley value* (Shapley 1952) is a method of measuring power or dividing payoff that is uniquely characterized by four important axioms (Dubey 1975). It measures the contribution of each agent to the grand coalition N by analyzing its *marginal contributions* to various subsets of agents. The marginal contribution of an agent $i \in N$ to a coalition $S \subseteq N - i$ is $v(S + i) - v(S)$.

Now, for a permutation π of agents, let $\Gamma_i^\pi = \{j | \pi(j) < \pi(i)\}$ be the set of agents before i in π. The Shapley value of agent i in game $G = (N, v)$, denoted $\phi_G(i)$, is given by

$$\phi_G(i) = \frac{1}{n!} \sum_{\pi \in S_n} v(\Gamma_i^\pi + i) - v(\Gamma_i^\pi), \tag{1}$$

where S_n is the set of all permutations of length n. For any coalition $S \subseteq N - i$, the number of permutations π where $\Gamma_i^\pi = S$ is exactly $(|S|)! \cdot (n - |S| - 1)!$. Thus, summing over the different coalitions, we get:

$$\phi_G(i) = \frac{1}{n!} \sum_{S \subseteq N - i} |S|! \cdot (n - |S| - 1)! \cdot (v(S + i) - v(S)). \tag{2}$$

3 Theoretical Results

Consider a cooperative game $G = (N, v)$. As described in Section 1, we study synergies by examining the games that are induced by an agent $i \in N$ committing to either participate or not participate in the joint endeavor. In either case, agent i's choice becomes fixed, and agent i is effectively removed from the game, leaving behind a different game among the other agents. We call such removals IN and OUT removals, respectively.

1. *IN Removal:* When agent i commits to joining the coalition, the final coalition contains agent i. Hence, the reduced game among the other agents is $G_i^{IN} = (N - i, v_i^{IN})$ where $v_i^{IN}(S) = v(S + i)$ for all $S \subseteq N - i$.
2. *OUT Removal:* When agent i commits to leaving the coalition, the final coalition does not contain agent i. Hence, the reduced game among the other agents is $G_i^{OUT} = (N - i, v_i^{OUT})$ where $v_i^{OUT}(S) = v(S)$ for all $S \subseteq N - i$.

For any two agents $i, j \in N$, let

$$\Delta_G(i,j) = \phi_{G_i^{IN}}(j) - \phi_{G_i^{OUT}}(j)$$

denote the difference between the Shapley values of agent j in the games arising from agent i joining and agent i leaving. In addition, for two agents $i, j \in N$ and coalition $S \subseteq N - i - j$, we denote

$$v_{ij}(S) = v(S + i + j) + v(S) - v(S + i) - v(S + j).$$

The proof of the following fact is very similar to a proof of Myerson (1980, Lemma 2), and hence it is omitted.

Fact 1 (Symmetry). *In any cooperative game $G = (N, v)$, for any two agents $i, j \in N$,*

$$\Delta_G(i,j) = \Delta_G(j,i) = \sum_{S \subseteq N-i-j} \frac{|S|!(n-|S|-2)!}{(n-1)!} \cdot v_{ij}(S). \tag{3}$$

This fact leads to a natural definition of synergy in cooperative games, which is similar to the definition given by Grabisch and Roubens (1999):

Definition 1. Given a game $G = (N, v)$, two agents $i, j \in N$ are *synergistic* if $\Delta_G(i,j) = \Delta_G(j,i) > 0$, *antagonistic* if $\Delta_G(i,j) = \Delta_G(j,i) < 0$, and *indifferent* if $\Delta_G(i,j) = \Delta_G(j,i) = 0$.

3.1 The Synergy Graph

Our technical point of departure from the literature is the observation that the structure of synergies in the preceding model can be naturally represented via a graph.

Definition 2. Given a cooperative game $G = (N, v)$, its *synergy graph*, denoted $SG(G)$, is the complete *undirected* graph where the agents are the vertices, and there is an edge between every pair of agents (i, j) labelled '$+$' if the pair is synergistic, '$-$' if it is antagonistic, and '0' if it is indifferent.

The undirected nature of the graph is due to Fact 1. We ask: *Which complete undirected graphs with edges labelled $+$, $-$, or 0 can emerge as synergy graphs?* We show that every graph with no indifferences (where the edges are labelled $+$ or $-$) can emerge as a synergy graph, even if the underlying game is restricted to be strictly monotonic.

Theorem 1. *For every complete undirected graph H where each edge is labelled $+$ or $-$, there exists a strictly monotonic cooperative game G such that $SG(G) = H$.*

Proof. We prove the theorem by induction on the number of vertices in H, i.e., the number of agents.

For the base case of a single agent, H has no edges. Thus, any game with a single agent would work. Take $v(\emptyset) = 0$ and $v(N) = 1$. Suppose the result holds for any graph with $n-1$ agents. Now, take a complete undirected graph H with the set of agents N, where $|N| = n$, and each edge labelled either $+$ or $-$. Fix some agent $i \in N$, and consider the graph H_{-i} obtained by removing i and its edges.

By the induction hypothesis there is a strictly monotonic game $G' = (N-i, v')$ with $SG(G') = H_{-i}$. Now, construct a game $G = (N, v)$ as follows. Let $c > 0$ and $\epsilon_j \geq 0$ for all $j \in N$ (we fix them later). For all $S \subseteq N$, define

$$v(S) = \begin{cases} v'(S) & \text{when } i \notin S, \\ v'(S-i) + c & \text{when } i \in S \text{ and } |S| > 2, \\ v'(\emptyset) + c + \epsilon_i & \text{when } S = \{i\}, \\ v'(\{j\}) + c + \epsilon_j & \text{when } S = \{i, j\}, j \in N - i. \end{cases}$$

We show that there exist c and $\{\epsilon_j\}_{j \in N}$ such that the constructed game G satisfies $SG(G) = H$. Let $\epsilon = \max_{j \in N} |\epsilon_j|$. Let $\alpha^* = \frac{1}{2} \cdot \min_{S \subsetneq T \subseteq N-i} v'(T) - v'(S)$. Note that having $\epsilon \leq \alpha^*$ would imply that G is also strictly monotonic. Consider two types of agent pairs.

Agents (j,k) where $i \notin \{j,k\}$. Applying Equation (3) for $\Delta_G(j,k)$, and breaking the summation over $S \subseteq N - j - k$ into two parts depending on $i \in S$ or $i \notin S$, we get:

$$\begin{aligned} \Delta_G(j,k) = \sum_{S \subseteq N-i-j-k} & \frac{|S|!(n-|S|-2)!}{(n-1)!} \cdot v_{jk}(S) \\ & + \frac{(|S|+1)!(n-|S|-3)!}{(n-1)!} \cdot v_{jk}(S+i). \end{aligned} \tag{4}$$

Similarly, applying Equation (3) to $\Delta_{G'}(j,k)$,

$$\Delta_{G'}(j,k) = \sum_{S \subseteq N-i-j-k} \frac{|S|!(n-|S|-3)!}{(n-2)!} v_{jk}(S) \tag{5}$$

Now, substituting

$$\frac{|S|!(n-|S|-3)!}{(n-2)!} = \frac{|S|!(n-|S|-2)!}{(n-1)!} + \frac{(|S|+1)!(n-|S|-3)!}{(n-1)!}$$

in Equation (5) and subtracting it from Equation (4),

$$\Delta_G(j,k) - \Delta_{G'}(j,k) = \sum_{S \subseteq N-i-j-k} \frac{(|S|+1)!(n-|S|-3)!}{(n-1)!} [v_{jk}(S+i) - v_{jk}(S)]. \tag{6}$$

Next, we argue that

$$|v_{jk}(S+i) - v_{jk}(S)| \leq 4\epsilon \tag{7}$$

for all $S \subseteq N - i - j - k$. Indeed, note that

$$v_{jk}(S+i) - v_{jk}(S) \\ = [v(S+i+j+k) - v(S+j+k)] + [v(S+i) - v(S)] \\ - [v(S+i+j) - v(S+j)] - [v(S+i+k) - v(S+k)].$$

Since two terms are positive and two are negative, we can subtract c within each term without changing the value of the expression. Further, from our construction, it is clear that $|v(T+i) - v(T) - c| \leq \epsilon$ for all $T \subseteq N - i$. Equation (7) now directly follows. Using Equations (6) and (7),

$$|\Delta_G(j,k) - \Delta_{G'}(j,k)| \\ \leq \sum_{S \subseteq N-i-j-k} \frac{(|S|+1)!(n-|S|-3)!}{(n-1)!} \cdot 4\epsilon \\ = 4\epsilon \cdot \sum_{k=0}^{n-3} \binom{n-3}{k} \frac{(k+1)!(n-k-3)!}{(n-1)!} \\ = 4\epsilon \cdot \sum_{k=0}^{n-3} \frac{k+1}{(n-1)(n-2)} = 2\epsilon.$$

Take $0 < \epsilon^* < \min(\alpha^*, \frac{1}{2} \cdot \min_{j,k \in N-i} |\Delta_{G'}(j,k)|)$ (it exists because the upper bound is positive due to our induction hypothesis); having $0 \leq \epsilon \leq \epsilon^*$ ensures that the signs of $\Delta_G(j,k)$ and $\Delta_{G'}(j,k)$ match for all $j, k \in N - i$, i.e., that the edges of $SG(G)$ not incident on i match the corresponding edges of H, and preserves strict monotonicity.

Agents (i,j) for $j \in N - i$. In this case, note that $v_{ij}(S) \neq 0$ if and only if $S = \emptyset$ (in which case $v_{ij}(S) = \epsilon_j - \epsilon_i$) or $S = \{k\}$ (in which case $v_{ij}(S) = -\epsilon_k$) for some $k \in N - i - j$. Substituting these into Equation (3), we get:

$$\Delta_G(i,j) = \frac{\epsilon_j - \epsilon_i}{n-1} - \frac{\sum_{k \in N-i-j} \epsilon_k}{(n-1)(n-2)} \\ = \frac{(n-2)(\epsilon_j - \epsilon_i) - \sum_{k \in N-i-j} \epsilon_k}{(n-1)(n-2)} \\ = \frac{(n-1)\epsilon_j - (n-2)\epsilon_i - \sum_{k \in N-i} \epsilon_k}{(n-1)(n-2)}, \quad (8)$$

where the last transition follows by adding ϵ_j to the first term and the last term in the numerator. We consider three cases based on the edges in H incident on i. Let n_+ be the number of edges incident on i labelled $+$, and n_- be the number of edges incident on i labelled $-$.

1. $n_+ = 0$: In this case we want $\Delta_G(i,j) < 0$ for all $j \in N - i$. In Equation (8), taking $\epsilon_i = \epsilon^*$ and $\epsilon_j = 0$ for all $j \in N - i$ ensures this.
2. $n_- = 0$: In this case we want $\Delta_G(i,j) > 0$ for all $j \in N - i$. In Equation (8), taking $\epsilon_i = -\epsilon^*$ and $\epsilon_j = 0$ for all $j \in N - i$ ensures this.
3. $n_+ \geq 1, n_- \geq 1$: Set $\epsilon_i = 0$. For each $j \in N - i$, set $\epsilon_j = \epsilon^*/n_+$ if (i,j) is labelled $+$, and $\epsilon_j = -\epsilon^*/n_-$ otherwise. This way, $\sum_{j \in N-i} \epsilon_j = 0$. Thus, from Equation (8), the sign of $\Delta_G(i,j)$ is same as the sign of ϵ_j for each $j \in N - i$, as required.

Further, note that the choice of $\{\epsilon_j\}_{j \in N}$ above satisfies $0 \leq \epsilon = \max_{j \in N} |\epsilon_j| \leq \epsilon^*$, which was required to match the edges in H not incident on i and to preserve strict monotonicity. Thus, G is strictly monotonic and $SG(G) = H$, which completes the proof by induction. □

Theorem 1 is reminiscent of McGarvey's Theorem from social choice theory, which states that when several voters express their preferences (as rankings) over various alternatives, and we draw a directed graph where there is an edge from alternative a to alternative b if and only if a majority of the voters prefer a to b, then every complete directed graph (tournament) emerges from some collection of preferences (McGarvey 1953).

We interpret Theorem 1 as saying that the question of whether two agents are synergistic or antagonistic is, in a sense, independent of the signs of the synergies between other pairs. In fact, we conjecture that all complete graphs where each edge is labelled $+$, $-$, or 0 emerge as synergy graphs of strictly monotonic games, that is, each pair of agents can independently be synergistic, antagonistic, *or indifferent*. Our Mathematica simulations show that this is true for up to 7 agents. However, it seems difficult to extend our proof technique to indifferences since equalities are hard to maintain in the induction step.

Naturally, when the class of games under consideration is restricted even further, the set of induced synergy graphs shrinks. For example, *strictly superadditive* games satisfy $v(S \cup T) > v(S) + v(T)$ for any two disjoint coalitions $S, T \subseteq N$. It turns out that in any strictly superadditive game, every vertex must be part of at least one synergistic pair. To see why, note that for any game $G = (N, v)$ and $i \in N$,

$$\sum_{j \in N-i} \Delta_G(i,j) = \sum_{j \in N-i} \left[\phi_{G_i^{IN}}(j) - \phi_{G_i^{OUT}}(j)\right] \\ = \left[v_i^{IN}(N-i) - v_i^{IN}(\emptyset)\right] - \left[v_i^{OUT}(N-i) - v_i^{OUT}(\emptyset)\right] \\ = v(N) - v(N-i) - v(\{i\}) + v(\emptyset).$$

But in a superadditive game, by definition, $v(N) - v(N - i) - v(\{i\}) > 0$, and hence it must be the case that $\Delta_G(i,j) > 0$ for some $j \in N - i$.

3.2 Complexity

Given a game, is it hard to determine if a pair of agents is synergistic, antagonistic, or indifferent? The answer depends on the representation of the game. While it is a trivial exercise to show that the problem is in EXPTIME for general cooperative games with only an oracle for the characteristic function,[1] one may hope for stronger tractability results when the game is constrained to possess a realistic structure.

Specifically, we are interested in weighted voting games (WVGs), which are perhaps the most popular of all cooperative games. They are primarily used to model decision making bodies where agents have different weights, and an action is taken if enough agents agree to it so that

[1]The result is intuitive and the proof is trivial. We omit the proof due to lack of space.

their total weight is at least a fixed threshold. More generally, WVGs model situations where the threshold represents the amount of a resource required to complete a task, and the agent weights represent the amounts of that resource the agents contribute if they decide to participate. It turns out that the problems of determining whether a pair of agents is synergistic and whether it is antagonistic are $\mathcal{NP}$-hard, while the problem of determining if a pair of agents is indifferent is co$\mathcal{NP}$-hard. This is not too surprising, given that the problem of computing the Shapley value in WVGs is $\#\mathcal{P}$-complete (Deng and Papadimitriou 1994; Matsui and Matsui 2000).

Theorem 2. *In WVGs, the problems of determining whether a given pair of agents is synergistic, whether it is antagonistic, and whether it is indifferent are $\mathcal{NP}$-hard, $\mathcal{NP}$-hard, and co$\mathcal{NP}$-hard, respectively.*

Proof. We use a variant of the subset sum problem that we call SUBSET-SUM-EQ. A SUBSET-SUM-EQ instance is given by a set of positive integers $U = \{x_1, \ldots, x_n\}$, and positive integers k and l. The question is to determine if there exists $S \subseteq U$ whose elements sum to exactly k. It is also given that all the solutions are of the given size l, that is, the elements in S sum to exactly k only if $|S| = l$. It can be seen from the reductions given in (Papadimitriou and Steiglitz 1982; Papadimitriou 1994) that SUBSET-SUM-EQ is $\mathcal{NP}$-complete.

We argue that the problem remains $\mathcal{NP}$-complete even if we assume that $l \geq \lceil n/2 \rceil$, because if $l < \lceil n/2 \rceil$ then we simply change k to $\sum_{i=1}^{n} x_i - k$. This gives a polynomial-time reduction from any instance to an instance of this restricted case, because S is a solution of the original instance if and only if $U \setminus S$ is a solution of the new instance.

Now, take an instance $I = (U = \{x_1, \ldots, x_n\}, k, l)$ of SUBSET-SUM-EQ, and let T denote the number of solutions to I. Construct a WVG $G = (N, \mathbf{w}, t)$, where $|N| = n + 2$, as follows. For $1 \leq i \leq n$, let $w_i = 2x_i$. Let $w_{n+1} = 2W + 2$, and $w_{n+2} = 1$, where $W = \sum_{i=1}^{n} x_i$. Set the threshold $t = 2W+2+2k+1$. We are now interested in $\Delta_G(n+1, n+2)$.

In the game G_{n+1}^{IN} obtained by the IN removal of agent $n+1$ from G, the threshold is $2k+1$. Further, agent $n+2$ has a positive marginal contribution to a coalition S if and only if $w(S) = 2k$. Every such coalition corresponds to a solution to I. Hence,

$$\phi_{G_{n+1}^{IN}}(n+2) = \frac{l!(n-l)!}{(n+1)!} \cdot T.$$

In the game G_{n+1}^{OUT} obtained by the OUT removal of agent $n+1$ from G, the threshold is $2W + 2 + 2k + 1$, so every coalition has value 0. This implies that $\phi_{G_{n+1}^{OUT}}(n+2) = 0$. Therefore, $\Delta_G(n+1, n+2) = \phi_{G_{n+1}^{IN}}(n+2)$; this is positive or zero if and only if T is positive or zero, respectively. Hence, the problems of determining whether a given pair of agents is synergistic, and whether it is indifferent, are $\mathcal{NP}$-hard and co$\mathcal{NP}$-hard, respectively.

For the problem of determining whether a given pair of agents is antagonistic, we use a similar, but different reduction. Take an instance $I = (U = \{x_1, \ldots, x_n\}, k, l)$ of SUBSET-SUM-EQ, and let T denote the number of solutions to I. Construct a WVG $G = (N, \mathbf{w}, t)$, where $|N| = n + 2$, as follows. For $1 \leq i \leq n$, let $w_i = 2x_i$. Let $w_{n+1} = 2k+1$, and $w_{n+2} = 1$. Set the threshold $t = 2k+1$. We are still interested in $\Delta_G(n+1, n+2)$.

In the game G_{n+1}^{IN} obtained by the IN removal of agent $n+1$ from G, the threshold is 0, so every coalition has value 1. Thus, $\phi_{G_{n+1}^{IN}}(n+2) = 0$. In the game G_{n+1}^{OUT} obtained by the OUT removal of agent $n+1$ from G, the threshold is still $2k+1$. Further, agent $n+2$ has a positive marginal contribution to a coalition S if and only if $w(S) = 2k$. Every such coalition corresponds to a solution to I. Hence,

$$\phi_{G_{n+1}^{OUT}}(n+2) = \frac{l!(n-l)!}{(n+1)!} \cdot T.$$

Thus, $\Delta_G(n+1, n+2) = -\phi_{G_{n+1}^{OUT}}(n+2)$; this is negative if and only if T is positive. Hence, the problem of checking if a given pair of agents is antagonistic is $\mathcal{NP}$-hard. □

4 Experimental Results

In this section we conduct computational experiments on two real-world WVGs, representing decision making in the European Union (EU) and in the International Monetary Fund (IMF). The purpose is twofold: (i) to show that the synergy structure of real-world WVGs can be computed, despite Theorem 2, and (ii) to tease out insights about synergies in these high-stakes games, with possibly wider applications.

The EU game consists of 28 agents with weights varying from 3 to 29 (total weight 352), and a quota of 260 (Edward and Lane 2013). The IMF game consists of 128 agents with weights varying from 0.03 to 16.75 (total weight 100). For most policy decisions, the IMF uses simple majority (50% quota), while some decisions require supermajority quotas of 70% and 85% (Weiss 2012). Since the IMF game is computationally intensive, we perform our experiment using only the simple majority quota. We used the dynamic programming algorithm of (Matsui and Matsui 2000; Bachrach and Shah 2013) for computing the Shapley values in WVGs.

We use "heat maps" to represent synergy and antagonism in any WVG. A heat map of a WVG is a square image where on both axes agents are sorted in increasing order of their weights, from top to bottom and from left to right. Thus, the entry in row i and column j represents the synergy or antagonism between the agent with the i'th lowest weight and the agent with the j'th lowest weight. In the *plain heat map*, a cell is colored blue (dark gray in grayscale) if the corresponding pair of agents is synergistic, and colored red (light gray in grayscale) if it is antagonistic. In the *gradient heat map*, the colors have varying intensity, which reflect the *magnitude of synergy or antagonism* between various pairs. Note that both heat maps are symmetric with respect to the diagonal due to Fact 1.

Let us begin with the IMF game. Figures 1(a) and 1(b) show the plain and gradient heat maps, respectively, of the IMF game. From the plain heat map, it is evident that a very high weight country is synergistic with all other countries (the blue patches at the bottom and at the right), except for other very high weight countries (the red region at the bottom-right corner). The gradient map reveals additional structure: The strongest antagonism exists between pairs of countries with very high weight, while the strongest synergy exists between pairs involving one very high weight country and one slightly lower (but still high) weight country.

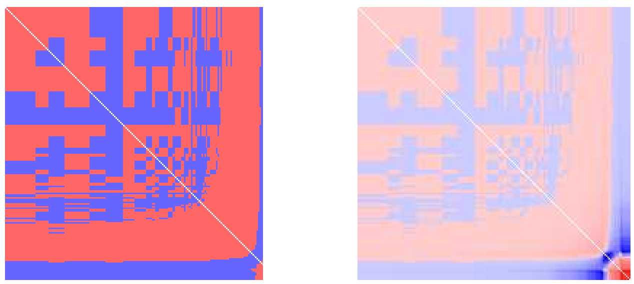

(a) Plain heat map (b) Gradient heat map

Figure 1: Synergies in the IMF game.

The EU game is less computationally intensive, so we were able to render heat maps with various quotas (as a percentage of the total weight, which is 352), in order to tease out additional structure. This is motivated by the fact that many decision making bodies use different quotas for decisions with varying importance. Figures 2(a), 2(b), 2(d), and 2(c) show the plain heat maps of the EU game with quotas 20%, 40%, 60%, and 80%, respectively.

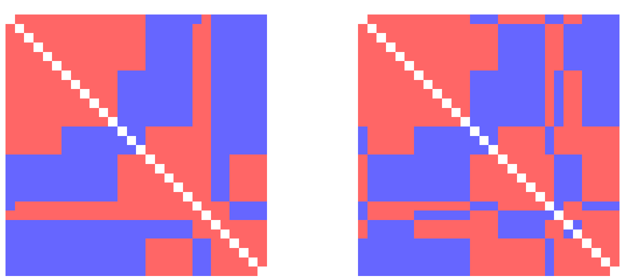

(a) 20% (b) 40%

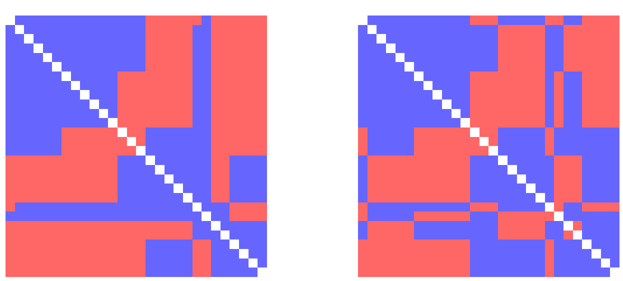

(c) 80% (d) 60%

Figure 2: The EU game with various quotas.

Looking at Figure 2, a pattern becomes immediately apparent: The figures for quota $x\%$ and $(100 - x)\%$ (for $x = 20, 40$) are complements of each other, that is, a pair is synergistic in one of them if and only if it is antagonistic in the other. Extending this pattern to 50% quota, one would expect all pairs of countries to be indifferent, because this quota is the complement of itself. However, this is not the case. Moreover, we found that a quota of $92 = 352 - 260$ does not generate a heat map that complements the heat map generated by the quota 260, which is the one used in practice. Note that the complement quota of 92 has a real-world interpretation: It is the weight required to vote down a proposal.

These observations have uncovered a deeper connection between a WVG with threshold $x\%$ of the total weight and the same WVG but with threshold $(100 - x)\%$ of the total weight. It is possible to show, using Equation (3) and algebraic manipulations, that the two are indeed complements of each other *if no coalition in the game has total weight equal to the threshold*. This is indeed the case in the EU game with the quota set to 20%, 40%, 60%, or 80% (since all of them generate non-integral quotas). Hence, these games satisfy the sufficient condition. However, when the quota is equal to 50%, i.e., 176 votes, there are coalitions with total weight exactly equal to the quota. This is also the case when the quota is 260. Further, we observed that there are coalitions with total weight equal to the real-world threshold of 61 in each and every composition of the Israeli Knesset since the country was founded in 1948 — 18 compositions in total! In general, since real-world WVGs use integer weights, and include many agents with very small weights, it is common to find coalitions whose total weight is exactly the threshold. In conclusion, the EU game with various quotas exhibits a particularly interesting structure of synergies that is surely common to many WVGs, but other real-world games (including the EU game coupled with its actual quota) do not satisfy the same sufficient condition and may give rise to different, possibly more intricate, synergy structures.

5 Discussion

In this paper, we have studied the structure of synergies in general cooperative games. Building on the symmetric synergy notion of Myerson (1980), we have established the existence of rich synergy structures (Theorem 1), but have also shown that these structures are theoretically hard to compute (Theorem 2). Nevertheless, in our computational experiments we were able to compute the synergy structures of real-world WVGs, leading to additional insights.

In the same vein as Theorem 1, one may ask which synergy graphs can arise from WVGs. Our Mathematica simulations show that even for 3 agents, there exist synergy graphs that do not arise from any WVG. However, when edge labels are restricted to just $+$ and $-$ (i.e., indifferences are not allowed), all synergy graphs with a small number of agents do arise from WVGs. We believe that this is true for any number of agents, but formally establishing this conjecture remains an open problem.

6 Acknowledgments

Procaccia and Shah were partially supported by the NSF under grants IIS-1350598 and CCF-1215883.

References

Ando, K. 2012. Computation of the Shapley value of minimum cost spanning tree games: #P-hardness and polynomial cases. *Japan Journal of Industrial and Applied Mathematics* 29(3):385–400.

Aziz, H.; Lachish, O.; Paterson, M.; and Savani, R. 2009. Power indices in spanning connectivity games. In *Proceedings of the 5th International Conference on Algorithmic Aspects in Information and Management (AAIM)*, 55–67.

Bachrach, Y., and Rosenschein, J. S. 2009. Power in threshold network flow games. *Autonomous Agents and Multi-Agent Systems* 18(1):106–132.

Bachrach, Y., and Shah, N. 2013. Reliability weighted voting games. In *Proceedings of the 6th International Symposium on Algorithmic Game Theory (SAGT)*, 38–49.

Bachrach, Y.; Markakis, E.; Procaccia, A. D.; Rosenschein, J. S.; and Saberi, A. 2008. Approximating power indices. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 943–950.

Bachrach, Y.; Parkes, D. C.; and Rosenschein, J. S. 2013. Computing cooperative solution concepts in coalitional skill games. *Artificial Intelligence* 204:1–21.

Bistaffa, F.; Farinelli, A.; Cerquides, J.; Rodriguez-Aguilar, J. A.; and Ramchurn, S. 2014. Anytime coalition structure generation on synergy graphs. In *Proceedings of the 13th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. Forthcoming.

Calvo, E.; Lasaga, J. J.; and Winter, E. 1996. The principle of balanced contributions and hierarchies of cooperation. *Mathematical Social Sciences* 31:171–182.

Conitzer, V., and Sandholm, T. 2006. Complexity of constructing solutions in the core based on synergies among coalitions. *Artificial Intelligence* 170(6):607–619.

Deng, X., and Papadimitriou, C. H. 1994. On the complexity of cooperative solution concepts. *Mathematics of Operations Research* 19(2):257–266.

Dubey, P. 1975. On the uniqueness of the Shapley value. *International Journal of Game Theory* 4(3):131–139.

Edward, D., and Lane, R. 2013. *Edward and Lane on European Union Law*. Edward Elgar Publishing Limited.

Grabisch, M., and Roubens, M. 1999. An axiomatic approach to the concept of interaction among players in cooperative games. *International Journal of Game Theory* 28:547–565.

Hart, S., and Mas-Colell, A. 1989. Potential, value, and consistency. *Econometrica* 57(3):589–614.

Matsui, T., and Matsui, Y. 2000. A survey of algorithms for calculating power indices of weighted majority games. *Journal of the Operations Research Society of Japan* 43:71–86.

McGarvey, D. C. 1953. A theorem on the construction of voting paradoxes. *Econometrica* 21:608–610.

Michalak, T. P.; Rahwan, T.; Szczepanski, P. L.; Skibski, O.; Narayanam, R.; Jennings, N. R.; and Wooldridge, M. J. 2013. Computational analysis of connectivity games with applications to the investigation of terrorist networks. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI)*, 293–301.

Murofushi, T., and Soneda, S. 1993. Techniques for reading fuzzy measures (iii): interaction index. In *Proceedings of the 9th Fuzzy Systems Symposium (FSS)*, 693–696.

Myerson, R. B. 1980. Conference structures and fair allocation rules. *International Journal of Game Theory* 9(3):169–182.

Owen, G. 1972. Multilinear extensions of games. *Management Science* 18(5-part-2):64–79.

Owen, G. 1977. Values of games with a priori unions. In Henn, R., and Moeschlin, O., eds., *Essays in Mathematical Economics and Game Theory*. Springer. 76–88.

Papadimitriou, C. H., and Steiglitz, K. 1982. *Combinatorial optimization: algorithms and complexity*. Prentice-Hall, Inc.

Papadimitriou, C. H. 1994. *Computational complexity*. Addison-Wesley.

Shapley, L. S. 1952. A value for n-person games. Technical report, DTIC Document.

Vinyals, M.; Voice, T.; Ramchurn, S. D. D.; and Jennings, N. R. 2013. A hierarchical dynamic programming algorithm for optimal coalition structure generation. arXiv:1310.6704.

Voice, T.; Ramchurn, S. D.; and Jennings, N. R. 2012. On coalition formation with sparse synergies. In *Proceedings of the 11th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 223–230.

Weiss, M. 2012. *International Monetary Fund: Background and Issues for Congress*. CreateSpace Independent Publishing Platform.

Zick, Y.; Skopalik, A.; and Elkind, E. 2011. The Shapley value as a function of the quota in weighted voting games. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, 490–496.