

State of the Art: Solution Concepts for Coalitional Games

Simina Brânzei

David R. Cheriton School of Computer Science
University of Waterloo
Waterloo, N2L 3G1, Ontario, Canada
sbranzei@uwaterloo.ca

Abstract

This paper investigates solution concepts for coalitional games. Several solution concepts are characterized, such as the core, Shapley value, bargaining set, stable set, nucleolus, and kernel. We look at recent developments of succinct representations of coalitional games, such as weighted voting games, coalitional resource games, cooperative Boolean games, and marginal contribution nets. Existing solution concepts have prohibitive complexity requirements even for very simple classes of games. We discuss an agenda for finding an equilibrium solution concept that is as appealing as the core, but that is tractable and guaranteed to exist.

Keywords: theory, coalitions, solution concepts, multiagent systems

1. Introduction

Coalitional (or cooperative) games study how groups of self interested players interact to accomplish more together than they could achieve individually. These interactions are modeled using a set of players, a set of actions, and a preference profile over the joint outcomes. While the players have to cooperate to achieve the outcomes that they want, their reasoning is always self-interested and they seek the actions most likely to bring the highest possible payoff (utility) for themselves. Solution concepts describe, in each game, the set of outcomes believed to arise based on the type of reasoning employed by the players. Coalitional solution concepts are often universal, i.e. applicable to all types of coalitional games, and attempt to ensure some form of stability. No matter how good an outcome is for the society at large, it may not be enforced if several players can become rich by deviating. In general, there exists a tension between outcomes that are beneficial for social welfare and outcomes that are stable. In this paper we are mainly concerned with stability.

The rest of the paper is structured as follows. Section 2 mentions the terms and notations useful when discussing solution concepts. Sections 3 through 11 introduce several solution concepts for coalitional games. Section 12 is a survey of recent classes of games and their complexity results with respect to solution concepts. Section 13 contains an agenda for finding an equilibrium concept for coalitional games that is tractable and guaranteed to exist.

2. Preliminaries

Coalitional game theory focuses on what groups of players can achieve. Because of the cooperation required for the players to achieve their goals, there is a need to predict the outcomes of these games with solution concepts that explicitly take the cooperation into account. We first consider the class of transferable utility coalitional games, in which the payoff of a coalition is given and the agents negotiate with each other to divide the payoff.

A *coalitional game with transferable utility* (TU) is a pair (N, v) where N represents the set of players and v is a valuation function $v : 2^N \rightarrow R$ that associates with each coalition $S \subset N$ a real value $v(S)$ that the coalition members can freely redistribute among themselves. The set N is also known as the *grand coalition* and the value $v(\emptyset)$ is zero by default. The function v is also known as a *characteristic function* and the game (N, v) is said to be in *characteristic function form*. A coalition S is a subset of players from N . The players are self-interested and attempt to maximize their own payoff (or utility).

A *coalitional game with nontransferable utility* (NTU) is a pair (N, v) where N is a finite set of players indexed by i and $v : 2^N \rightarrow 2^{R^{|S|}}$ is a map that associates to each coalition $S \subset N$ a subset $v(S) \in R^{|S|}$ of attainable payoff vectors. NTU coalitional games specify exactly what divisions of

payoff are allowed and forbid all the other payoff divisions. The class of TU coalitional games is a subset of the class of NTU coalitional games. A TU coalitional game (N, v) can be specified as the NTU game (N, v) , where $v(S) = \{x \in R^{|N|} \text{ such that } \sum_{i \in S} x_i = v(S) \text{ and } x_j = 0 \text{ if } j \notin S\}$.

We introduce several classes of coalitional games that are useful when discussing solution concepts. A game (N, v) is *superadditive* if for all coalitions $S, T \subset N$ such that $S \cap T = \emptyset$, $v(S \cup T) \geq v(S) + v(T)$. A game (N, v) is *additive* if for all coalitions $S, T \subset N$ such that $S \cap T = \emptyset$, $v(S \cup T) = v(S) + v(T)$. A game (N, v) is *constant-sum* if for any coalition $S \subset N$, $v(S) + v(N \setminus S) = v(N)$. A game (N, v) is *convex* if for all $S, T \subset N$ such that $S \cap T = \emptyset$, the following inequality holds: $v(S \cup T) \geq v(S) + v(T) - v(S \cap T)$. A game (N, v) is *simple* if for all coalitions $S \subset N$, $v(S) \in \{0, 1\}$. A coalition S in a simple game is said to be *winning* if $v(S) = 1$ and *losing* if $v(S) = 0$.

Given a coalitional game (N, v) , a payoff vector x in R^N is *feasible* if $\sum_{i \in N} x_i \leq v(N)$. That is, no payoff structure can promise the players more utility than the worth of the grand coalition N . A payoff structure is *individually rational* if any player receives under that payoff structure at least as much utility as they could obtain on their own. That is, for any player i , the inequality $x_i \geq v(\{i\})$ holds. This is an important requirement in a game since it states that players cannot be forced to accept outcomes that are worse than what they could obtain on their own. The sum $\sum_{i \in S} x_i$ is denoted by $x(S)$. A payoff structure is *coalitionally rational* if for any coalition $S \subset N$ it is the case that $x(S) \geq v(S)$. The *pre-imputation set*, P , is defined as $\{x \in R^N \text{ such that } \sum_{i \in N} x_i = v(N)\}$. The pre-imputations are those payoff structures that are efficient, i.e. distribute completely the payoff of the grand coalition. The *imputation set* I is defined as $\{x \in P \text{ such that } x_i \geq v_i\}$. The imputation set consists of those pre-imputations that respect individual rationality. Given a set of outcomes, a change from one outcome to another that makes at least one player better off without making any other player worse off is called a *Pareto improvement*. The set of outcomes where no Pareto improvements can be made are said to be *Pareto efficient* (or optimal). The *social welfare* is the sum of utilities of all the players for a given outcome. It is not always the case that outcomes that maximize social welfare are also stable or viceversa.

In the game theoretic literature it is usually assumed that the grand coalition forms in equilibrium and the only problem to be solved is the division of payoffs among the players. This assumption is correct in superadditive games, where the addition of any player to a coalition brings positive utility for that coalition. However, it no longer holds in non-superadditive settings, where there may be negative interactions that prevent the players from forming the grand coalition. We will later look at games equipped with *coalition structure*, i.e. possible partitions of the set of players.

3. The Core

The core is one of the oldest and best known stability solution concepts in coalitional games. A payoff structure belongs to the core if no group of players can deviate from the current payoff structure and obtain a better utility for themselves by doing so. A payoff vector x is S -feasible if $x(S) = v(S)$, where $x(S) = \sum_{i \in S} x_i$. A payoff vector is feasible if it is N -feasible. Formally, the core of a TU coalitional game is the set $\{x \in P \text{ such that there is no coalition } S \text{ and } S\text{-feasible payoff vector } y \text{ such that } y_i > x_i \text{ for all } i \in S\}$. A slightly different version of the core is given by the set $\{x \in P \text{ such that there is no coalition } S \text{ and payoff structure } y \in v(S) \text{ with } y_i \geq x_i \text{ for all } i \in S, \text{ and } y_j > x_j \text{ for at least one player } j\}$. That is, a coalition S can block if it does not degrade any of its members and strictly improves at least one player. This core is also known as the strong core, since it relaxes the conditions for blocking and is therefore more likely to be empty. The strong core is contained in the "weak" core, previously defined.

The NTU core has a similar definition to the TU core. The core of the NTU game (N, v) is the set $\{x \in v(N) \text{ such that there exists no coalition } S \text{ and payoff structure } y \in v(S) \text{ such that } y_i > x_i \text{ for all } i \in S\}$. A deviating group of players is called a *blocking coalition* (for that allocation). From

the point of rationality, the core is the set of allocations that respect both individual and coalitional rationality.

The following three player majority game has a nonempty core, depending on the value of a parameter. Assume that three players can get together one unit of profit, and two of them can obtain $\alpha \in [0, 1]$ independently of the actions of the third. A single player gets nothing. Formally, the game is the pair (N, v) , where $N = \{1, 2, 3\}$, $v(N) = 1$, $v(S) = \alpha$ for $|S| = 2$, and $v(\{i\}) = 0$ for any player $i \in \{1, 2, 3\}$. The core of the game is given by the set of payoff vectors (x_1, x_2, x_3) such that $x_i \geq 0$, $x(N) = 1$ and $x(S) \geq \alpha$ whenever $|S| = 2$. The core is nonempty if and only if $\alpha \leq 2/3$.

As the examples illustrate, the core can be empty. This is problematic, because in empty core instances there is no way to divide the payoff without having some players deviate. The core can also be too large, case in which it leaves the payoff decision problem open. In both cases, other mechanisms such as contracts or social norms have to be used to enforce an outcome. There exist several approaches for guaranteeing nonemptiness of the core, such as identifying constraints that the game should satisfy, or relaxing the definition of the core itself. Non-additive constant-sum games always have an empty core. On the other hand, convex games always have a nonempty core. More generally, the Bondareva-Shapley theorem gives necessary and sufficient conditions for the core to exist.

A set λ of non-negative weights over 2^N is *balanced* if for any player $i \in N$, $\sum_{i \in S} \lambda(S) = 1$. A game (N, v) is balanced if and only if $v(N) \geq \sum_{S \subset N} \lambda(S) v(S)$ for every balanced set of weights λ . The theorem (Bondareva-Shapley) states that a TU coalitional game has a nonempty core if and only if it is balanced.

To give some intuition into the notion of balancedness, consider that each player has a unit of time, which they can freely distribute among coalitions that they are a member of. For a coalition S to be active for a fraction of time $\lambda(S)$, it must be that all of its members are active in S for that fraction of time. In that case, coalition S obtains a payoff of $\lambda(S)v(S)$. Thus the condition of balanced weights requires that players don't allocate more time than they have. The game is balanced if no feasible allocation of time yields more than the value of the grand coalition $v(N)$.

The Bondareva-Shapley theorem does not completely eliminate the difficulty of recognizing empty core games. In many NTU settings there is no obvious way of using the balancedness condition. Perhaps more importantly, knowing that a game has a nonempty core does not automatically give an algorithm for searching the core elements. Indeed, searching the core by brute force translates into enumerating all game configurations and checking the existence of a blocking coalition for each such configuration. In many games, even clever enumeration algorithms are still exponential because the problem of finding a blocking coalition itself is NP-hard. Another direction for obtaining existence of the core is to refine its definition.

4. Extensions of the Core: Epsilon-Core, Core Cover, Reasonable Set, Weber Set

The *epsilon-core* requires that the players have sufficient incentives in order to deviate. That is, any blocking coalition must guarantee *each* of its players at least ε more payoff than they are currently getting. Equivalently, a payoff vector x is in the epsilon-core of the game (N, v) if and only if for any coalition $S \subset N$, $\sum_{i \in S} x_i \geq v(S) - \varepsilon$. A variation of the epsilon-core can be obtained by requiring that a blocking coalition must guarantee *each* of its players at least ε more payoff than they are currently getting. For $\varepsilon = 0$, the two cores coincide. For ε positive, the epsilon-core is more likely to exist. For example, setting $\varepsilon = \infty$ ensures that *any* configuration is in the epsilon-core, since the cost of deviating is prohibitive. The nonempty epsilon-core obtained for the smallest value of ε is known as the *least* core.

Three sets related to the core are the core cover, the reasonable set, and the Weber set ([5]). All of these sets contain the core as a subset. We first introduce some definitions required for defining them. The *marginal contribution* of a player i to a coalition S is $M_i(S) = v(S) - v(S \setminus \{i\})$. The marginal contribution is the amount by which the value of the coalition S would decrease if player i were to leave the coalition.

The *upper vector* M of a game (N, v) contains at every coordinate i the marginal contribution $M_i(N)$ of the player i to the grand coalition. This value is also known as the utopia payoff. If player i requires more than $M_i(N)$, then it is always better for the other players to throw i out. Given nonempty coalition S and player i , let the *remainder* $R(S, i)$ be the amount which remains for player i if all the other players in S receive their utopia payoffs: $R(S, i) = v(S) - \sum_{j \in S \setminus \{i\}} M_j(N)$. The *lower vector* m of a game (N, v) contains at every coordinate i the value: $m_i = \max_{S: i \in S} R(S, i)$. The value m_i is the minimum right payoff for player i in the grand coalition. This player has the right to ask at least m_i since otherwise he can threaten to form the blocking coalition S , where i would receive m_i by allowing all the other players in S to get their utopia payoffs.

The *core cover* $CC(N, v)$ is the set of all imputations contained between the *lower vector* and the *upper vector*. That is, $CC(N, v) = \{x \in I \text{ such that } m \leq x \leq M\}$.

The *reasonable set* is defined as: $R(N, v) = \{x \in R^n \text{ such that } v(i) \leq x_i \leq \max_{S: i \in S} [v(S) - v(S \setminus \{i\})]\}$. The reasonable set satisfies three axioms: symmetry, covariance, and monotonicity ([6]). Symmetry states that a reasonable outcome should not differentiate by players' names. Covariance requires that if a game is rescaled by changing the unit and the 0-value payoffs, the outcome should be rescaled accordingly. Monotonicity requires that players with larger marginal contributions receive higher payoffs.

The Weber set is defined using marginal contribution vectors. Let $\pi(N)$ be the set of all permutations $\sigma : N \rightarrow N$. The set $P^\sigma(i) = \{k \in N \text{ such that } \sigma^{-1}(k) < \sigma^{-1}(i)\}$ consists of all the predecessors of i with respect to permutation σ . The *marginal contribution vector* $m^\sigma \in R^n$ with respect to σ has entries $m_i^\sigma = v(P^\sigma(i) \cup \{i\}) - v(P^\sigma(i))$, for each player i . For example, $m_{\sigma(1)}^\sigma = v(\sigma(1))$, $m_{\sigma(2)}^\sigma = v(\sigma(1), \sigma(2)) - v(\sigma(1))$, ..., $m_{\sigma(k)}^\sigma = v(\sigma(1), \dots, \sigma(k)) - v(\sigma(1), \dots, \sigma(k-1))$, for every $k \in N$.

The *Weber set* is the convex hull of all the marginal contribution vectors, corresponding to all the permutations of N . The payoff vector m^σ described above can be computed by allowing the players to enter the coalition in the order $\sigma(1), \dots, \sigma(n)$. Then each player is given the marginal contribution that he creates by entering in this order. The Weber set is characterized axiomatically by additivity, null (or dummy) player, efficiency, and monotonicity ([7]). Additivity requires that the payoff obtained by combining two disjoint coalitions is the sum of their separate worth. A null player has zero marginal contribution in any coalition and he should always receive zero. Efficiency requires that the total gains are distributed. Monotonicity has been introduced in the axioms for the reasonable set. It is interesting to observe that for the class of convex games, the core and the Weber set coincide. The Weber set is closely related to the Shapley value, which we introduce next.

5. The Shapley Value

The Shapley value is another axiomatic solution concept. It is unique among the solution concepts presented so far in that it always exists and contains exactly one point. The axioms characterizing the Shapley value are efficiency, symmetry, null player, and additivity. Again, efficiency requires that the worth of the grand coalition is distributed: $\sum_{i \in N} x_i = v(N)$. Symmetry states that interchangeable players should receive equal payoff, where two players are interchangeable if they have the same marginal contribution to every coalition. The Shapley value is the unique solution concept satisfying these four axioms. The Shapley value of player i in game (N, v) is: $\phi_i = 1/n! \sum_{S \subseteq N \setminus \{i\}} |S|! (|N| - |S| - 1)! [v(S \cup \{i\}) - v(S)]$. Intuitively, the Shapley value captures the average marginal contribution of player i over all the possible orders in which the grand coalition could have been constructed. In convex games, the Shapley value is a member of the core. When the Shapley value is not in the core, it is possible that although the payoff division satisfies good properties, the players may refuse to stay in those configurations. In other words, the Shapley value does not satisfy coalitional rationality and certain groups of players may deviate if it is in their best interest to do so.

Osborne and Rubinstein ([4]) describe an equivalent formulation of the Shapley value in terms of objections and counterobjections. First let (N, v) be a TU coalitional game. For each coalition $S \subset N$, define the subgame (S, v^S) be a TU coalitional game such that $v^S(T) = v(T)$ for any $T \subseteq S$. Let φ be a feasible payoff function. An *objection* of player i against player j to the division x of $v(N)$ could take one of the following forms:

- "Give me more since otherwise I shall leave the game, causing you to obtain $\varphi_j(N \setminus \{i\}, v^{N \setminus \{i\}})$, which is smaller than your current payoff by $x_j - \varphi_j(N \setminus \{i\}, v^{N \setminus \{i\}})$."
- "Give me more since otherwise I shall convince the other players to exclude you from the game, allowing me to obtain $\varphi_i(N \setminus \{j\}, v^{N \setminus \{j\}})$, which is larger than my current payoff by $\varphi_i(N \setminus \{j\}, v^{N \setminus \{j\}}) - x_i$."
- Player j could counteract the first objection by asserting:
- "Indeed, if you leave I will lose something, but if I leave, *you* will lose at least as much: $x_i - \varphi_i(N \setminus \{j\}, v^{N \setminus \{j\}}) \geq x_j - \varphi_j(N \setminus \{i\}, v^{N \setminus \{i\}})$."
- Player j could counteract the second objection by asserting:
- "Indeed, by excluding me you will gain something, but I can exclude *you* and gain at least as much by doing that: $\varphi_j(N \setminus \{i\}, v^{N \setminus \{i\}}) - x_j \geq \varphi_i(N \setminus \{j\}, v^{N \setminus \{j\}}) - x_i$."

The Shapley value is required to satisfy the property that for every objection of any player i against j , there exists a counterobjection of player j . This requirement is equivalent to the *balanced contributions property* which states that for every TU coalitional game (N, v) , $\varphi_i(N, v) - \varphi_i(N \setminus \{j\}, v^{N \setminus \{j\}}) = \varphi_j(N, v) - \varphi_j(N \setminus \{i\}, v^{N \setminus \{i\}})$. The Shapley value is the unique solution concept that satisfies the balanced contributions property.

6. The τ -Value and the Compromise Value

The τ -value was introduced by Tijs ([11]) and is defined on quasi-balanced games. A game (N, v) is *quasi-balanced* if $m \leq M$ and $\sum_{i=1,n} m_i \leq v(N) \leq \sum_{i=1,n} M_i$, where m and M are the lower and upper vectors, respectively. The τ -value for a quasi-balanced game (N, v) is the point in the closed interval $[m, M] \in R^n$ given by $\tau = m + \lambda M$. The parameter $\lambda \in R$ is uniquely chosen such that $\sum_{i \in N} \tau_i = v(N)$. The τ -value is the unique value that satisfies efficiency, restricted proportionality, and the minimal right property([12]) and it always exists. We have already introduced efficiency. The minimal right property states that any player i should receive at least its lower vector value, m_i . The restricted proportionality axiom states that the τ -value is proportional to M if the lower vector is zero ($m_i = 0$ for all $i \in N$). The τ -value satisfies a few more properties ([14]): individual rationality, null player, symmetry, and covariance.

An example of the τ -value is given for the following game. Let (N, v) be such that $N = \{1, 2, 3\}$ and $v(\{1\}) = v(\{2\}) = 0$, $v(\{3\}) = v(\{1, 2\}) = 100$, $v(\{1, 3\}) = 200$, $v(\{2, 3\}) = 300$, $v(N) = 400$. The upper vector for this game is $M = (100, 200, 300)$. The lower vector $m_1 = \max \{ v(\{1\}) - M_2(N), v(\{1, 3\}) - M_3(N), v(N) - M_2(N) - M_3(N) \} = \max \{ 0, -100, -100, -100 \} = 0$. Similarly, it can be verified that $m_2 = 0$ and $m_3 = 100$. Then the τ -value is the vector $\tau = (0, 0, 100) + \lambda(100, 200, 200)$, where λ is chosen such that $\sum_{i \in N} \tau_i = 400$. This gives $\lambda = 3/5$ and so the τ -value is $\tau = (60, 120, 220)$.

The *compromise value* is defined on the class of compromise admissible NTU games([14]). The lower and upper vectors can be defined on NTU games similarly to the TU setting. It can be proved that if (N, v) is an NTU game with nonempty core, then any payoff vector x in the core of the game satisfies the inequality $m \leq x \leq M$. For each nonempty coalition S , the set $dom(S) = \{ x \in R^{|S|} \text{ such that } x < y \text{ for some } y \in v(S) \}$. The payoffs in $dom(S)$ are, in some sense, the best possible for coalition S , since there is no payoff vector y that would make everyone in S happier. An NTU game (N, v) is compromise admissible if $m \leq M$, $m \in v(N)$, and $M \notin dom(N)$. For a compromise admissible NTU game, the compromise value T is the unique point on the line between m and M

which lies in $v(N)$ (i.e. T is allowed by the NTU game specification) and is nearest to the utopia payoff M . Formally, the compromise value is $T = m + \alpha^* (M - m)$, where $\alpha^* = \max \{ \alpha \in [0, 1] \text{ such that } m + \alpha (M - m) \in v(N) \}$. The compromise value is the equivalent formulation of the τ -value on quasi-balanced TU games.

7. The Stable Set

The stable set, also known as the von Neumann-Morgenstern solution, was initially proposed for two player games. The main idea behind this solution concept is that a coalition S that is unhappy with the current payoff structure should be able to object, by threatening to implement credible payoffs which are better for all the members of S .

An imputation x is an *objection* of the coalition S to imputation y if $x_i > y_i$ for all players $i \in S$ and $x(S) \leq v(S)$. In this case, imputation x is said to dominate imputation y via coalition S . Observe that the core is the set of all imputations to which there exists no objection. We introduce two notions of stability that will be used to define the stable set.

A set of imputations Y satisfies *internal stability* if no imputation $y \in Y$ has an objection via some coalition S . A set of imputations Y satisfies *external stability* if for any imputation $z \notin Y$, there exists an objection $y \in Y$. A *stable set* is a set of imputations that satisfy internal and external stability. Each stable set can be interpreted as acceptable behaviours in society. No acceptable behaviour is preferred to another, and for every unacceptable behaviour there exists an acceptable alternative. While a game has only one core, it may have more than one stable set, none of which are subsets of each other. However, the core is a subset of any stable set and if the core is a stable set itself, then it is the only one. The stable set may also be empty.

An example of a game with a stable set is the following three player majority game. Let (N, v) be such that $N = \{1, 2, 3\}$ and $v(S) = 1$ for $|S| \geq 2$ and $v(S) = 0$ otherwise. A stable set for this game is $Y = \{(\frac{1}{2}, \frac{1}{2}, 0), (\frac{1}{2}, 0, \frac{1}{2}), (0, \frac{1}{2}, \frac{1}{2})\}$. Equality is preferred when sharing the available unit.

8. The Bargaining Set

The bargaining set was proposed by Aumann and Maschler and is defined similarly to the Shapley value, in terms of individual objections and counterobjections. Consider a pair (y, S) , where S is a coalition and y is a payoff vector feasible from the point of view of S , i.e. $\sum_{i \in S} x_i \leq v(S)$.

- (y, S) is an objection of player i against j to payoff x if S includes i but not j and $y_k > x_k$ for all $k \in S$. That is, player i can find a coalition S that excludes j , and by doing so guarantees a better payoff for all the players in S , including i .
- (z, T) is a counterobjection to objection (y, S) if z is a payoff vector feasible for T , coalition T includes player j but not i , $z_k \geq x_k$ for all $k \in T$ and $z_k \geq y_k$ for all $k \in S \cap T$. That is, player j can counterbalance i 's objection by finding a coalition T that excludes i , to which j can guarantee a better payoff than they are currently being offered.

The *bargaining set* of a TU game (N, v) is the set of imputations x such that for every objection (y, S) to x of any player i against any player j , there exists a counterobjection (z, T) that j can make. The core is a subset of the bargaining set, and for convex games they coincide. Unlike the core, the bargaining set is always nonempty. In the three player majority game, the bargaining set is the vector $Y = \{(1/3, 1/3, 1/3)\}$. It can be verified that every individual objection has a counterobjection.

There exist several variations on the classical Aumann-Maschler bargaining set. Observe that the core consists of payoff structures to which there exist no objections. However, the core does not assess the credibility of these objections. The bargaining set can be seen as a solution concept that addresses this credibility problem, by stipulating that an objection is justified only when it has no counterobjections. This requirement makes blocking harder and stabilizes the game, ensuring that the bargaining set is always nonempty. Mas-Colell([9]) proposes a modified version of the bargaining set in which the counterobjections must also be verified for their credibility. Dutta *et al*

([8]) generalize this idea to a chain of objections, in which the credibility of each objection is challenged. A terminating object for a chain of objections is an objection to which there exists no counterobjection. Since the number of coalitions is finite, it follows that each chain must eventually terminate. This fact is used to assess the validity of the original objection. The *consistent bargaining set* is the set of payoff structures with no valid objections.

Bennet ([16]) introduces the aspiration bargaining set. The aspiration approach to coalition formation asks the following question: Given a payoff vector x , when can player i view his component x_i as potentially attainable? Player i could view x_i as potentially attainable if there is a coalition S containing i which can afford x , that is $x(S) \leq v(S)$. A payoff vector x in which each player i can view their payoff x_i as potentially attainable is called an anticipation. Formally, payoff vector x is an *anticipation* for the game (N, v) if for every player $i \in N$, there exists a coalition S such that $x(S) \leq v(S)$. Anticipations are not always rational outcomes. For example, in some games an anticipation can be the zero vector, in which every player receives payoff $x_i = 0$. However, the self interest of the players is likely to lead to higher anticipations whenever possible. Consider the situation during negotiations where some coalition S has a surplus of payoff after each player received x_i , i.e. $v(S) > x(S)$. It would be reasonable to expect that the payoff anticipations would increase until there was no longer a surplus, and anticipations with no surplus are called *aspirations*.

Anticipation x is an aspiration if for every anticipation vector y and coalition S that satisfies $y_j > x_j$ and $y_i \geq x_i$, for all i and some j in S , it follows that $y(S) > v(S)$. Not all aspirations are reasonable. If given payoff structure x and two players $i, j \in S$ such that every coalition $S \ni i$ that can realize x also contains j , then player i needs j to realize payoff x_i . However, if j can obtain x_j without player i , then i is in a vulnerable position with respect to j , and player j can argue that they should receive more payoff. An aspiration in which no player is vulnerable is called a bargaining aspiration, and the set of all bargaining aspirations is called the *aspiration bargaining set*. The aspiration bargaining set can also be defined for NTU games ([17]).

9. The Kernel

The kernel was introduced by Davis and Maschler. Given an imputation x of the TU coalitional game (N, v) , let $e(S, x) = v(S) - x(S)$ be the *excess* of the coalition S at x . If the excess is positive, then coalition S must sacrifice the amount $e(S, x)$ to have imputation x implemented. If the excess is negative, then $e(S, x)$ is the bonus that coalition S receives when x is implemented. The kernel is defined in terms of objections and counterobjections as follows:

- Player i objects against j in imputation x by pointing out there is coalition S that includes i but excludes j , and $x_j > v(\{j\})$. From player i 's point of view, coalition S is required to sacrifice too much in the absence of player j .
- Player j can counterobject to i with coalition T that includes j but excludes i , and $e(T, x) \geq e(S, x)$. Player j 's counterargument is the existence of a coalition T that excludes player i and sacrifices equally or more.

The kernel is the set of imputations such that for any objection of any player i against any player j , there exists a counterobjection that j can bring. The kernel can be alternatively defined using the maximum excess. For any two players i and j , let $s_{ij}(x)$ be the maximum excess of any coalition containing i but not j : $s_{ij}(x) = \max_{S \in C} \{e(S, x) \text{ where } i \in S \text{ and } j \notin S\}$. Then the kernel is the set of imputations x such that for any two players i and j , either $s_{ji}(x) \geq s_{ij}(x)$ or $x_j = v(\{j\})$. The requirement $x_j = v(\{j\})$ states that no-one cannot object against a player that merely receives their individually rational outcome. The kernel of the three player majority game is $\{(1/3, 1/3, 1/3)\}$.

10. The Nucleolus

The nucleolus was introduced by Schmeidler. The excess $e(S, x)$ can be interpreted as the strength of the complaint of coalition S against imputation x . The higher the complaint, the more unhappy coalition S is with imputation x and the more loudly it protests against x . The nucleolus attempts to minimize the complaints under the budget constraints (the feasibility of x). The

minimization is done by resolving the loudest complaint first, then the next loudest, until finishing all the complaints. The resulting imputation is a lexicographic minimum of all the complaints. The nucleolus of any TU coalitional game is nonempty and a subset of the kernel. Since the kernel is a subset of the bargaining set, the nucleolus is also included in the bargaining set. Moreover, the nucleolus is a singleton.

An equivalent definition of the nucleolus can be formulated in terms of objections and counterobjections. A pair (S, y) consisting of a coalition S and imputation y is an objection to imputation x if $e(S, x) > e(S, y)$. Equivalently, $y(S) > x(S)$, implying that coalition S objects to imputation x because it receives a better payoff under imputation y . A coalition T is a counter-objection to the objection (S, y) if $e(T, y) > e(T, x)$ (i.e. $x(T) > y(T)$) and $e(T, y) \geq e(S, x)$. Coalition T receives a better payoff under imputation y and T has a higher complaint under y than S under x .

11. The Shapley-Shubik and Banzhaf Power Index

The *Banzhaf power index* was introduced by John Banzhaf III for the purpose of analyzing voting systems. The main idea behind the Banzhaf power index is that the voting power of a player comes from their ability to change the outcomes of elections. Let (N, v) be a simple coalitional game. For each player i , the Banzhaf index is defined as the proportion of all voting outcomes in which the result would be different if i changed their vote. Intuitively, the player is *pivotal* for those outcomes. A player that is pivotal for any winning coalition is a *dictator*. The Banzhaf index satisfies symmetry and the dummy player axioms. If the game is normalized, it also satisfies efficiency. The Banzhaf index is not additive. Formally, the Banzhaf index of player i is $\beta_i = 1/2^{n-1} \sum_{i \in S \subseteq N} [v(S) - v(S \setminus \{i\})]$.

The *Shapley-Shubik power index* applies the Shapley value to simple games. The players are assumed to vote in order. A player i is pivotal for an ordering if the respective coalition was losing before i was introduced, but winning afterwards. The Shapley-Shubik index is the fraction of all $n!$ orderings in which the player is pivotal. Formally, the Shapley-Shubik index of player i is $\varphi_i = \sum_{i \in S \subseteq N} 1/|N|! (|S| - 1)! (|N| - |S|)! [v(S) - v(S \setminus \{i\})]$. Both the Banzhaf and Shapley-Shubik indices are weighted marginal contributions.

12. Games and Representations

The standard definition of coalitional games in characteristic function form is exponential in the number of agents. For a game (N, v) , the characteristic function v specifies the payoff for each of the $2^{|N|}$ possible coalitions of N . If these coalition values are not related to each other in a meaningful way, a search algorithm has to examine all the possible coalitions, and thus none of the solution concepts discussed so far can be computed in polynomial time in the number of players. There has been much recent work in the computer science community proposing new classes of coalitional games that are succinct, yet provide a rich enough structure to have interesting coalitional properties.

Deng and Papadimitriou ([23]) suggest using computational complexity as a criterion for judging whether a solution concept is appropriate or not. If a solution concept is to be useful, then the complexity of determining the outcomes predicted by that solution concept should not be too great. As we already mentioned, if the characteristic function randomly assigns values to coalitions, then searching for solution concepts is clearly exponential in the number of players. The question becomes more interesting when the representation is polynomial. Deng and Papadimitriou propose the following game. Given an undirected graph $G = (N, E)$, let $w_{i,j}$ be the weight on each edge (i, j) . Each node of the graph represents a player and a coalition S can guarantee for its players the weight of the subgraph induced by the corresponding nodes: $v(S) = \sum_{(i, j) \subseteq S} w_{i,j}$. If all the weights are nonnegative, the game is convex. The Shapley value for this game can be computed in polynomial time and it is half the sum of edge weights adjacent to each node. The core is easy to find when the game is convex, but intractable otherwise. The kernel, the nucleolus, the epsilon-core, and the bargaining set are also intractable. The stable set is believed to be undecidable.

Weighted voting games assign to each player i a weight w_i and define an overall game quota q . A coalition S is winning if and only if the sum of weights of its players exceeds the quota. Thus $v(S) = 1$ if $\sum_{i \in S} w_i \geq q$, and $v(S) = 0$, otherwise. The game is denoted by $(q, w_1, \dots, w_n)$. Computing the Shapley value is $\#P$ for these games ([24]). There exists a pseudopolynomial time algorithm by dynamic programming. However, approximating the Shapley value within a constant factor is intractable. Similar hardness results exist for the core, the least core, and the nucleolus. Determining emptiness of the core can be done in polynomial time, but finding imputations in the least core or the nucleolus is NP-hard. Pseudopolynomial algorithms exist for these solution concepts ([24], [25]). Chalkiadakis *et al* ([30]) introduce *coalitional games with beliefs*, a generalization of coalitional games to environments where players have private beliefs about the capabilities of the other players. They introduce a notion of the core for games with beliefs, with and without coalition structure. They give hardness results and analyze weighted voting games with beliefs. Introducing beliefs in weighted voting games adds a layer of complexity and the core-related questions become even harder to answer.

Qualitative coalitional games have been introduced by Wooldridge and Dunne ([26]). The players are assumed to have goals that they want to achieve, and are happy when they are in a coalition where they can accomplish at least some of those goals. A qualitative coalitional game is specified by a set N of players and a collection $\{G_i\}$ of sets of goals, where G_i represents the goals that player i wants to achieve, and is drawn from a set G of overall goals. In addition, a characteristic function $v : 2^N \rightarrow 2^{P(G)}$, where $P(G)$ is the set of all possible subsets of the goals, and $v(S)$ denotes the goals that coalition S could achieve in each of the possible ways that its players could cooperate. Wooldridge and Dunne propose a representation based on propositional logic, which is complete, but not always succinct. Checking nonemptiness of the core is D_1^P complete.

Coalitional resource games ([27]) are a variation of qualitative coalitional games in which the game is given by a set of players, a set of goals for each player, an endowment function that specifies the resources of each player, and a requirement function that indicates, for each goal, the resources that have to be spent in order to achieve that goal. Resource games are a strict subset of qualitative coalitional games and can always be represented succinctly. A coalition is successful if it can cooperate in such a way to achieve at least one goal for each of its members. A resource is *necessary* if exploiting that resource is required for every successful coalition. Deciding the success of a coalition and whether a resource is necessary are NP-complete and co-NP-complete, respectively. Dunne *et al* ([28]) investigate further the complexity of coalitional resource games and show that checking nonemptiness of the core is NP-hard. They show that checking stability is exponential in the number of goals, but polynomial in the number of players and resources, and thus feasible for games with bounded number of goals. They also give a negotiation protocol that guarantees several attractive properties.

Coalitional skill games have been introduced by Bachrach and Rosenschein ([29]). A coalitional skill game is specified by a set of players, a set of tasks, and a set of skills. Each agent is endowed with a set of skills, and every task requires certain skills. A coalition can perform a task if every skill required to perform it is owned by some player in that coalition. The characteristic function of a coalition is the value of the tasks that can be performed by that coalition. The authors examine complexity questions such as determining the value of a coalition, checking nonemptiness of the core and finding it, computing the Shapley value and Banzhaf power index. The questions about the core can be answered in polynomial time for simpler versions of coalitional skill games, but are intractable otherwise. Computing the Shapley value is NP-hard and the Banzhaf power index is $\#P$ -complete.

Cooperative boolean games are a compact family of coalitional games ([30]). Each player has a goal which is represented as a formula of boolean variables. A player has unique control over a subset of variables and can freely set their values. However, the action of setting variable values is costly. The game specifies a cost function for each of these actions and the players aim to achieve their goals while minimizing the costs. Cost minimization can only be accomplished by

cooperating. Search algorithms for the core and the stable set are investigated and turn out to be intractable.

There exists a rich literature on *hedonic games* ([32]), where each player's payoff is entirely determined by the identity of the other members of their coalition. Consider a set N of players. A coalition partition $\Pi = \{S_k\}_{k=1,K}$ is a disjoint set of coalitions that exhaustively partitions N . Each player has preferences over the coalition structures which are completely determined by the coalition that the player belongs to. A hedonic game is a pair $(N, >)$ of players and preferences. Bogomolnaia and Jackson study Pareto efficiency and the existence of several stability properties, such as the core, individual and Nash stability. Elkind and Wooldridge ([33]) study a succinct, rule-based representation for hedonic games and give hardness results for the complexity of core related questions under this representation. Brânzei and Larson ([34]) look at a subset of hedonic games that can be modeled by a graph for which they characterize the maximal welfare partitions and compute a bound for the cost of stability.

Network flow games are a class of simple games on a directed graph ([35]). The edges of the graph have capacities corresponding to the bandwidth of the edge. This game models many real situations, such as a network of connected computers, where the edges are the network connections and the weights the number of bits per second that can be transmitted through that connection. Given a network N and a set of players S , $N \downarrow S$ denotes the network flow that is completely owned by the players in S . For two points P_1 and P_2 in the network and an amount of flow (bandwidth) b that has to be ensured between the two points, a simple game can be defined as follows. Let the characteristic function v be such that $v(S) = 1$ if $N \downarrow S$ allows a flow of b between P_1 and P_2 , and $v(S) = 0$ otherwise. Determining the Shapley value is NP-hard unless restrictions are imposed on the network. Bachrach and Rosenschein ([36]) show that computing the Banzhaf power index in network flow games is also intractable.

Jeong and Shoham ([37]) propose a fully expressive representation based on rules that describe the marginal contributions of the players. *Marginal contribution nets* are succinct for some natural classes of games such as recommendation games, but can be exponential in worst case. Marginal contribution nets are used to develop an efficient algorithm for the Shapley value, and an algorithm for determining core membership and core emptiness that is exponential only in the treewidth of the net. The synergy representation ([1], Ch 12) is complete for superadditive games. The idea is that the game is stored as a pair (N, s) , where N is a set of players and s is a function mapping each coalition C to a value $s(C)$. This value is the synergy generated by coalition C when its members work together. It is important to note that only strictly positive synergies are included in the game specification. Conitzer and Sandholm ([38]) propose a concise approach in which games are decomposed in several issues, such that each issue has its own characteristic function. A multi-issue representation is a collection of coalitional games (issues) $(N_1, v_1), \dots, (N_k, v_k)$ which together form the game (N, v) . The set of players is $N = N_1 \cup \dots \cup N_k$ and the characteristic function is $v(S) = \sum_{i=1,k} v_i(N_i \cap S)$ for every coalition $S \subseteq N$. The multi-issue representation is universal and allows computation of the Shapley value in linear time in the size of the input. Algorithms for the core, such as checking core membership remain NP-hard.

13. The Search for New Solution Concepts

Solution concepts have been introduced as tools that can describe and predict the outcomes of coalitional games. As we have seen, many game representations have been recently emerging, but unfortunately the algorithms for computing their solution concepts have prohibitive runtime complexity.

The core, perhaps the most convincing equilibrium concept for coalitional games, has some major problems. Finding the core is NP-hard even for very simple classes of games. In addition, the core can be either empty or too large, cases in which other social rules must be used to select an appropriate outcome. The core extensions, such as the epsilon-core and the least-core are less convincing. The epsilon-core may be appropriate in some settings, where there exist costs

associated with deviating. However, there is no reason to believe that the utilities of the players are always comparable and that they will incur the same epsilon penalty for deviating every time they try to do so. The epsilon-core can also be empty. The least core is always nonempty, but the smallest epsilon constraint seems somewhat artificial. It is not clear that the cost of blocking can indeed be set up at the value required for the epsilon-core to exist. In addition, finding the least core is at least as hard as finding the core. The stable set always exists, but is often too large to make any meaningful predictions. NP-hardness results also exist about the nucleolus, stable set, bargaining set, kernel in a variety of games.

The problems illustrated with solution concepts are not unique to coalitional games. Pure Nash equilibrium, the most influential solution concept in non-cooperative game theory, suffers from similar drawbacks. Pure Nash equilibrium does not always exist. Mixed Nash equilibrium is guaranteed to exist, but is not very realistic, since it requires the players to randomize between equally attractive options with the sole purpose of inducing a similar behaviour in their opponent. Moreover, Nash equilibrium is PPAD-complete, thus very likely intractable ([2]). Correlated equilibrium, a generalization of Nash, is not necessarily a good replacement, as it requires a complicated machinery with a trusted randomizer.

Fabrikant and Papadimitriou ([18]) and Papadimitriou ([19]) propose an agenda for finding new game theoretic solution concepts. The solution concept should be a natural, compelling, convincing, and realistic model of behaviour. It should be tractable and guaranteed to exist in order to have universal applicability. There have been several attempts of finding solution concepts to replace Nash equilibrium. Among these, we mention sink equilibrium, unit recall equilibrium, CURB sets, and iterated regret minimization. While this agenda is mostly geared towards noncooperative game theory, we believe that coalitional solution concepts exhibit similar problems and that it would be useful to find a solution concept with these desiderata in mind. In particular, we would like an equilibrium solution concept for coalitional games that is as compelling as the core, but satisfies universal existence and ensures good complexity. It would be interesting to know if requiring individual rationality, tractability, universal existence, and a different notion of group rationality is feasible. An impossibility result for a satisfactory concept of stability in hedonic games has been given by Barbera and Gerber ([22]). A short list of axioms, namely universal existence, symmetry, Pareto optimality, and self-consistency cannot be satisfied by any solution concept. Their counterexample parallels Arrow's impossibility theorem in voting.

The core has also been criticized for the extremely myopic behaviour of the players, who disregard the consequences of their actions. It is possible that by deviating, a player would guide the coalition structure to an equilibrium with possibly disastrous payoff for that player. The stable set has a similar problem. Greenberg ([10]) shows that the stable set requires optimism from the members of a deviating coalition, since they would proceed with deviating if at least one of the final outcomes would benefit that coalition. Greenberg proposes a version of conservative stability, in which players only proceed with a deviation if every possible final outcome makes its members better off. Diamantoudi and Xue ([20]) investigate farsightedness in hedonic games.

The classic definition of the core does not take into account the fact that players in natural scenarios are computationally bounded, and may not be able to find a blocking coalition. A possible interpretation of boundedness is that blocking coalitions cannot be larger than a parameter k . For $k = 1$, a player can only check if their individual rationality is respected. For $k = n$, a player can verify all the possible coalitions when searching for a better payoff. Sandholm and Lesser ([21]) develop a normative, application- and protocol-independent theory of coalition formation for computationally bounded players and introduce the bounded-rational core. The bounded-rational core is equivalent to the classical core concept when players have complete search algorithms and computation is costless. There still exist games where the bounded-rational core is empty.

Another criticism of the core is with respect to its fragility. In any game configuration, it is sufficient that one blocking coalition exists for that configuration to automatically be labeled unstable. This assumption seems quite strong because of its combinatorial operator. In social

settings, if we view riots or revolutions as instances of blocking coalitions, it is not realistic to assume that players will exhaustively search the possible groups in society in order to find a coalition to protest with. Rather, what seems to be happening in these situations is that there exists enough unhappiness in that society and, as a result, potential blocking coalitions are numerous and players can easily find a deviating group.

Acknowledgements

I would like to thank professor Kate Larson for the excellent suggestions and feedback on this paper.

14. References

- [1] Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, 2009.
- [2] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. *The Complexity of Computing a Nash Equilibrium*. Proceedings of STOC, 2006.
- [3] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [4] M. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
- [5] R. Brânzei, D. Dimitrov, and S. Tijs. *Models in Cooperative Game Theory*. Springer Verlag, 2005.
- [6] L. A. Gerard-Varet and S. Zamir. *Remarks on the Reasonable Set of Outcomes in a General Coalition Function Form Game*. International Journal of Game Theory, Volume 16, Issue 2, 1987.
- [7] J. Derks, H. Haller, and H. Peters. *The selectope for cooperative games*. International Journal of Game Theory, 29:23–38, 2000.
- [8] B. Dutta, D. Ray, K. Sengupta, R. Vohra. *A consistent bargaining set*. Journal of Economic Theory, 1989, 49:93-112
- [9] A. Mas-Colell. *An equivalence theorem for a bargaining set*. Journal of Mathematical Economics, 1989, 18:129-138
- [10] J. Greenberg. *The theory of social situations: An alternative game theoretic approach*. Cambridge University Press, 1990.
- [11] S. H. Tijs. *Bounds for the core and the τ -value*. In O. Moeschlin and D. Pallaschke, Eds., *Game Theory and Mathematical Economics*, North-Holland, Amsterdam, 1981.
- [12] S. H. Tijs. *An axiomatization of the τ -value*. Math. Social Sciences 12, 9—20, 1987.
- [13] J. M. Bilbao, A. Jimenez-Losada, E. Lebron, S. H. Tijs. *The τ -value for games on matroids*, 2008. <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.77.7808> .
- [14] S. Tijs and G. J. Otten. *Compromise Values in Cooperative Game Theory*, 1991. Top 1 1-51.
- [15] S. Hart. and A. Mas-Colell, editors. *Classical Cooperative Theory I: Core-Like Concepts*. Springer-Verlag, 1997, 35-42.
- [16] E. Bennet. *The Aspiration Approach to Predicting Coalition Formation and Payoff Distribution in Sidepayment Games*. International Journal of Game Theory, Vol 12, 1: 1-28, 1983.
- [17] E. Bennet and W. R. Zame. *Bargaining in Cooperative Games*. International Journal of Game Theory, Vol 17, 4: 279-300, 1988.
- [18] A. Fabrikant and C. H. Papadimitriou. *The complexity of game dynamics: BGP oscillations, sink equilibria, and beyond*. SODA, 2008.
- [19] C. H. Papadimitriou. *The Search for Equilibrium Concepts*. SAGT, 2008.
- [20] E. Diamantoudi and L. Xue. *Farsighted stability in hedonic games*. Soc Choice Welfare, 21:39-61, 2003.
- [21] T. Sandholm and V. Lesser. *Coalitions among Computationally Bounded Agents*. Artificial Intelligence, 94(1), 99-137, Special issue on Economic Principles of Multiagent Systems.
- [22] S. Barbera and A. Gerber. *A note on the impossibility of a satisfactory concept of stability for coalition formation games*. Economics Letters, 95:85-90, 2007.

- [23] X. Deng and C. H. Papadimitriou. *On the Complexity of Cooperative Solution Concepts*. Mathematics of Operations Research, 19:257-266, 1994.
- [24] E. Elkind, L. A. Goldberg, P. Goldberg, and M. Wooldridge. *Computational Complexity of Weighted Threshold Games*. AAAI, 2007.
- [25] E. Elkind and D. Pasechnik. *Computing the nucleolus of weighted voting games*. SODA, 2009.
- [26] M. Wooldridge and P. E. Dunne. *On the computational complexity of qualitative coalitional games*. Artificial Intelligence, Volume 158, 1: 27-73, 2004.
- [27] M. Wooldridge and P. E. Dunne. *On the computational complexity of coalitional resource games*. Artificial Intelligence, Vol 170, 10: 835-871, 2006.
- [28] P. E. Dunne, S. Kraus, E. Manisterski, M. Wooldridge. *Solving coalitional resource games*. Artificial Intelligence, Vol 174, 1: 20-50, 2010.
- [29] Y. Bachrach and J. S. Rosenschein. *Coalitional Skill Games*. AAMAS, 2008.
- [30] G. Chalkiadakis, E. Elkind, and N. R. Jennings. *Simple Coalitional Games with Beliefs*. IJCAI, 2009.
- [31] P. E. Dunne, W. Hoek, S. Kraus, M. Wooldridge. *Cooperative Boolean Games*. AAMAS, 2008.
- [32] A. Bogomolnaia and M. O. Jackson. *The Stability of Hedonic Coalition Structures*. Games and Economic Behaviour, 38: 201-230, 2002.
- [33] E. Elkind and M. Wooldridge. *Hedonic Coalition Nets*. AAMAS 2009.
- [34] S. Branzei and K. Larson. *Coalitional Affinity Games and the Stability Gap*. IJCAI, 2009.
- [35] M. Wooldridge. *An Introduction to Multiagent Systems* (second edition). John Wiley & Sons, 2009.
- [36] Y. Bachrach and J. S. Rosenschein. *Computing the Banzhaf Power Index in Network Flow Games*. AAMAS, 2007.
- [37] S. Jeong and Y. Shoham. *Marginal Contribution Nets: A Compact Representation Scheme for Coalitional Games*. Electronic Commerce, 2005.
- [38] V. Conitzer and T. Sandholm. *Computing Shapley Values, Manipulating Value Division Schemes, and Checking Core Membership in Multi-Issue Domains*. AAAI, 2004.