

Dynamics of Profit-Sharing Games

John Augustine, Ning Chen, Edith Elkind, Angelo Fanelli, Nick Gravin, Dmitry Shiryayev

School of Physical and Mathematical Sciences
Nanyang Technological University, Singapore

Abstract

An important task in the analysis of multiagent systems is to understand how groups of selfish players can form coalitions, i.e., work together in teams. In this paper, we study the dynamics of coalition formation under bounded rationality. We consider settings where each team's profit is given by a concave function, and propose three profit-sharing schemes, each of which is based on the concept of marginal utility. The agents are assumed to be myopic, i.e., they keep changing teams as long as they can increase their payoff by doing so. We study the properties (such as closeness to Nash equilibrium or total profit) of the states that result after a polynomial number of such moves, and prove bounds on the price of anarchy and the price of stability of the corresponding games.

1 Introduction

Cooperation and collaborative task execution are fundamentally important both for human societies and for multiagent systems. Indeed, many tasks are too complicated or resource-consuming to be executed by a single agent, and a collective effort is needed. Such settings are usually modeled using the framework of *cooperative games*, which specify the amount of payoff that each subset of agents can achieve: when the game is played the agents split into teams (coalitions), and the payoff of each team is divided among its members.

The standard framework of cooperative game theory is static, i.e., it does not explain how the players arrive at a particular set of teams and a payoff distribution. However, understanding the dynamics of coalition formation is important for many applications of cooperative games, and game theorists have studied bargaining and coalition formation in cooperative environments (see, e.g. [Chatterjee et al., 1993; Moldovanu and Winter, 1995; Yan, 2003]). Much of this research assumes that the agents are fully rational, i.e., can predict the consequences of their actions and maximize their (expected) utility based on these predictions. However, full rationality is a strong assumption that is unlikely to hold in many real-life scenarios: first, the agents may not have the computational resources to infer their optimal strategies, and second, they may not be sophisticated enough to do so, or lack

information about other players. Such agents may simply respond to their current environment without worrying about the subsequent reaction of other agents; such behavior is said to be *myopic*. Now, coalition formation by computationally limited agents has been studied by a number of researchers in multi-agent systems, starting with the work of [Shehory and Kraus, 1999] and [Sandholm and Lesser, 1997]. However, myopic behavior in coalition formation received relatively little attention in the literature (for some exceptions, see [Dieckmann and Schwalbe, 2002; Chalkiadakis and Boutilier, 2004; Airiau and Sen, 2009]). In contrast, myopic dynamics of non-cooperative games is the subject of a growing body of research (see, e.g. [Fabrikant et al., 2004; Awerbuch et al., 2008; Fanelli et al., 2008]).

In this paper, we merge these streams of research and apply techniques developed in the context of analyzing the dynamics of non-cooperative games to coalition formation settings. In doing so, we depart from the standard model of games with transferable utility, which allows the players in a team to share the payoff arbitrarily: indeed, such flexibility will necessitate a complicated negotiation process whenever a player wants to switch teams. Instead, we consider three payoff models that are based on the concept of marginal utility, i.e., the contribution that the player makes to his current team. Each of the payoff schemes, when combined with a cooperative game, induces a non-cooperative game, whose dynamics can then be studied using the rich set of tools developed for such games in recent years.

We will now describe our payment schemes in more detail. We assume that we are given a concave cooperative game, i.e., the values of the teams are given by a submodular function; the submodularity property means that a player is more useful when he joins a smaller team, and plays an important role in our analysis. In our first scheme, the payment to each agent is given by his marginal utility for his current team; by submodularity, the total payment to the team members never exceeds the team's value. This payment scheme rewards each agent according to the value he creates; we will therefore call these games *Fair Value games*. Our second scheme takes into account the history of the interaction: we keep track of the order in which the players have joined their teams, and pay each agent his marginal contribution to the coalition formed by the players who joined his current team before him. This ensures that the entire payoff of each team is fully distributed among

its members. Moreover, due to the submodularity property a player's payoff never goes down as long as he stays with the same team. This payoff scheme is somewhat reminiscent of the reward schemes employed in industries with strong labor unions; we will therefore refer to these games as *Labor Union games*. Our third scheme can be viewed as a hybrid of the first two: it distributes the team's payoff according to the players' Shapley values, i.e., it pays each player his expected marginal contribution to a coalition formed by its predecessors when players are reordered randomly; the resulting games are called *Shapley games*.

Our Contributions We study the equilibria and dynamics of the three games described above. We establish that all our games admit a Nash equilibrium. Further, we argue that in all these games the total profit of the system in a Nash equilibrium, i.e., the sum of teams' values, is within a factor of 2 from optimal. In addition, for the first two classes of games, a natural dynamic process quickly converges to a state that is almost as good as a Nash equilibrium, and the best Nash equilibrium is in fact an optimal state. We also show that Labor Union games have other desirable properties.

Related Work The games studied in this paper belong to the class of *potential games*, introduced by Monderer and Shapley [1996]. In potential games, any sequence of improvements by players converges to a pure Nash equilibrium. However, the number of steps can be exponential in the description of the game. The complexity of computing (approximate) Nash equilibrium in various subclasses of potential games such as congestion games, cut games, or party affiliation games has received a lot of attention in recent years [Fabrikant et al., 2004; Skopalik and Vöcking, 2008; Bhargat et al., 2010]. Related issues are how long it takes for some form of best response dynamics to reach an equilibrium [Chien and Sinclair, 2007; Ackermann et al., 2008], or how good are the states reached after a polynomial number of steps [Awerbuch et al., 2008; Fanelli et al., 2008].

[Dieckmann and Schwalbe, 2002] study myopic behavior in coalition formation; however, unlike us they assume that the payoffs may be distributed arbitrarily, which leads to very different outcomes; [Chalkiadakis and Boutilier, 2004] extend their approach to incomplete information games. Most similar to our work are recent papers [Gairing and Savani, 2010; Gairing and Savani, 2011], which study the dynamics of a class of cooperative games known as additively separable hedonic games; however, in this paper we consider a much broader class of games.

2 Preliminaries

Non-Cooperative Games A *non-cooperative game* is defined by a tuple $\mathcal{G} = (N, (\Sigma_i)_{i \in N}, (u_i)_{i \in N})$, where $N = \{1, 2, \dots, n\}$ is the set of *players*, Σ_i is the set of (pure) *strategies* of player i , and $u_i : \times_{i \in N} \Sigma_i \rightarrow \mathbb{R}^+ \cup \{0\}$ is the *payoff function* of player i .

Let $\Sigma = \times_{i \in N} \Sigma_i$ be the *strategy profile set* or *state set* of the game, and let $S = (s_1, s_2, \dots, s_n) \in \Sigma$ be a generic state in which each player i chooses strategy $s_i \in \Sigma_i$. Given a strategy profile $S = (s_1, s_2, \dots, s_n)$ and a strategy $s'_i \in \Sigma_i$, let (S_{-i}, s'_i) be the strategy profile obtained from

S by changing the strategy of player i from s_i to s'_i , i.e., $(S_{-i}, s'_i) = (s_1, s_2, \dots, s_{i-1}, s'_i, s_{i+1}, \dots, s_n)$.

Nash Equilibria and Dynamics Given a strategy profile $S = (s_1, s_2, \dots, s_n)$, a strategy $s'_i \in \Sigma_i$ is an *improvement move* for player i if $u_i(S_{-i}, s'_i) > u_i(S)$; further, s'_i is called an α -*improvement move* for i if $u_i(S_{-i}, s'_i) > (1 + \alpha)u_i(S)$, where $\alpha > 0$. A strategy $s_i^b \in \Sigma_i$ is a *best response* for player i in state S if it yields the maximum possible payoff given the strategy choices of the other players, i.e., $u_i(S_{-i}, s_i^b) \geq u_i(S_{-i}, s'_i)$ for any $s'_i \in \Sigma_i$. An α -*best response move* is both an α -improvement and a best response move.

A (pure) *Nash equilibrium* is a strategy profile in which every player plays her best response. Formally, $S = (s_1, s_2, \dots, s_n)$ is a Nash equilibrium if for all $i \in N$ and for any strategy $s'_i \in \Sigma_i$ we have $u_i(S) \geq u_i(S_{-i}, s'_i)$. We denote the set of all (pure) Nash equilibria of a game \mathcal{G} by $\mathcal{NE}(\mathcal{G})$. A profile $S = (s_1, \dots, s_n)$ is called an α -*Nash equilibrium* if no player can improve his payoff by more than a factor of $(1 + \alpha)$ by deviating, i.e., $(1 + \alpha)u_i(S) \geq u_i(S_{-i}, s'_i)$ for any $i \in N$ and any $s'_i \in \Sigma_i$. The set of all α -Nash equilibria of \mathcal{G} is denoted by $\mathcal{NE}^\alpha(\mathcal{G})$. In a *strong Nash equilibrium*, no group of players can improve their payoffs by deviating, i.e., $S = (s_1, \dots, s_n)$ is a strong Nash equilibrium if for all $I \subseteq N$ and any strategy vector $S' = (s'_1, \dots, s'_n)$ such that $s'_i = s_i$ for $i \in N \setminus I$, if $u_i(S') > u_i(S)$ for some $i \in I$, then $u_j(S') < u_j(S)$ for some $j \in I$.

Let $\Delta_i(S)$ be the improvement in the player's payoff if he performs his best response, i.e., $\Delta_i(S) = u_i(S_{-i}, s_i^b) - u_i(S)$, where s_i^b is the best response of player i in state S . For any $Z \subseteq N$ let $\Delta_Z(S) = \sum_{i \in Z} \Delta_i(S)$, and let $\Delta(S) = \Delta_N(S)$. A *Nash dynamic* (respectively, α -*Nash dynamic*) is any sequence of best response (respectively, α -best response) moves. A *basic Nash dynamic* (respectively, *basic α -Nash dynamic*) is any Nash dynamic (respectively, α -Nash dynamic) such that at each state S the player i that makes a move has the maximum absolute improvement, i.e., $i \in \arg \max_{j \in N} \Delta_j(S)$.

Price of Anarchy Given a game \mathcal{G} with a set of states Σ , and a function $f : \Sigma \rightarrow \mathbb{R}^+ \cup \{0\}$, we write $\text{OPT}_f(\mathcal{G}) = \max_{S \in \Sigma} f(S)$. The *price of anarchy* $\text{PoA}_f(\mathcal{G})$ and the *price of stability* $\text{PoS}_f(\mathcal{G})$ of a game \mathcal{G} with respect to a function f are, respectively, the worst-case ratio and the best-case ratio between the value of f in a Nash equilibrium and $\text{OPT}_f(\mathcal{G})$, i.e., $\text{PoA}_f(\mathcal{G}) = \max_{S \in \mathcal{NE}(\mathcal{G})} \frac{\text{OPT}_f(\mathcal{G})}{f(S)}$, $\text{PoS}_f(\mathcal{G}) = \min_{S \in \mathcal{NE}(\mathcal{G})} \frac{\text{OPT}_f(\mathcal{G})}{f(S)}$. The *strong price of anarchy* and the *strong price of stability* are defined similarly; the only difference is that the maximum (respectively, minimum) is taken over all strong Nash equilibria. Further, the α -*price of anarchy* $\text{PoA}_f^\alpha(\mathcal{G})$ of a game \mathcal{G} with respect to f is defined as $\text{PoA}_f^\alpha(\mathcal{G}) = \max_{S \in \mathcal{NE}^\alpha(\mathcal{G})} \frac{\text{OPT}_f(\mathcal{G})}{f(S)}$; the α -*price of stability* $\text{PoS}_f^\alpha(\mathcal{G})$ can be defined similarly. Originally, these notions were defined with respect to the social welfare function, i.e., $f = \sum_{i \in N} u_i(S)$. However, we give a more general definition since in the setting of this paper it is natural to use a different function f . We omit the index f when the function f is clear from the context.

Potential Games A non-cooperative game \mathcal{G} is called a *potential game* if there is a function $\Phi : \Sigma \rightarrow \mathbb{N}$ such that for any state S and any improvement move s'_i of a player i in S we have $\Phi(S_{-i}, s'_i) - \Phi(S) > 0$; the function Φ is called the *potential function* of \mathcal{G} . The game \mathcal{G} is called an *exact potential game* if we have $\Phi(S_{-i}, s'_i) - \Phi(S) = u_i(S_{-i}, s'_i) - u_i(S)$. It is known that any potential game has a pure Nash equilibrium [Monderer and Shapley, 1996; Rosenthal, 1973].

Cooperative Games A *cooperative game* $G = (N, v)$ is given by a set of *players* N and a *characteristic function* $v : 2^N \rightarrow \mathbb{R}^+ \cup \{0\}$ that for each set $I \subseteq N$ specifies the profit that the players in I can earn by working together. We assume that $v(\emptyset) = 0$. A *coalition structure* over N is a partition of players in N , i.e., a collection of sets I_1, \dots, I_k such that (i) $I_i \subseteq N$ for $i = 1, \dots, k$; (ii) $I_i \cap I_j = \emptyset$ for all $i < j \leq k$; and (iii) $\bigcup_{j=1}^k I_j = N$. A game $G = (N, v)$ is called *monotone* if v is *non-decreasing*, i.e., $v(I) \leq v(J)$ for any $I \subset J \subseteq N$. Further, G is called *concave* if v is *submodular*, i.e., for any $I \subset J \subseteq N$ and any $i \in N \setminus J$ we have $v(I \cup \{i\}) - v(I) \geq v(J \cup \{i\}) - v(J)$. Informally, in a concave game a player is more useful when he joins a smaller coalition.

3 Perfect β -Nice Games

In this section, we define the class of perfect β -nice games (our definition is inspired by [Awerbuch et al., 2008], but differs from the one given there), and prove a number of results for such games. Later, we will show that many of the profit-sharing games considered in the paper belong to this class.

Definition 1. A *potential game* \mathcal{G} with a *potential function* Φ is called *perfect with respect to a function* $f : \Sigma \rightarrow \mathbb{R}^+ \cup \{0\}$ if for any state S it holds that $f(S) \geq \sum_{i \in N} u_i(S)$, and, moreover, for any improvement move s'_i of player i we have

$$f(S_{-i}, s'_i) - f(S) \geq \Phi(S_{-i}, s'_i) - \Phi(S) \geq u_i(S_{-i}, s'_i) - u_i(S).$$

Also, a game \mathcal{G} is called *β -nice with respect to f* if for every state S we have $f(S) \geq \sum_{i \in N} u_i(S)$ and $\beta \cdot f(S) + \Delta(S) \geq \text{OPT}_f(\mathcal{G})$.

We can bound the price of anarchy of a β -nice game by β .

Lemma 1. For any $f : \Sigma \rightarrow \mathbb{R}^+ \cup \{0\}$ and any game \mathcal{G} that is β -nice w.r.t. f we have $\text{PoA}_f(\mathcal{G}) \leq \beta$.

Lemma 1 extends to α -price of anarchy for any $\alpha \geq 0$.

Lemma 2. For any $f : \Sigma \rightarrow \mathbb{R}^+ \cup \{0\}$, any $\alpha \geq 0$, and any game \mathcal{G} that is β -nice w.r.t. f we have $\text{PoA}_f^\alpha(\mathcal{G}) \leq \alpha + \beta$.

The next theorem states that after a polynomial number of steps, for every perfect β -nice potential game the basic Nash dynamic reaches a state whose relative quality (with respect to f) is close to our bound on the price of anarchy.

Theorem 1. Consider any function $f : \Sigma \rightarrow \mathbb{R}^+ \cup \{0\}$ and any game \mathcal{G} that is perfect β -nice with respect to f . For any $\epsilon > 0$ the basic Nash dynamic converges to a state S^F with $f(S^F) \geq \frac{\text{OPT}_f(\mathcal{G})}{\beta} (1 - \epsilon)$ in at most $\lceil \frac{n}{\beta} \ln \frac{1}{\epsilon} \rceil$ steps, starting from any initial state.

4 Profit-Sharing Games

We will now describe three non-cooperative games that can be constructed from any monotone concave cooperative game. Each of our games can be described by a triple $\mathcal{G} = (N, v, M)$, where (N, v) is a monotone concave cooperative game with $N = \{1, \dots, n\}$, and $M = \{1, \dots, m\}$ is a set of m parties; we require $m \leq n$. All three games considered in this section model the setting where the players in N form a coalition structure over N that consists of m coalitions. Thus, each player needs to choose exactly one party from M , i.e., for each $i \in N$ we have $\Sigma_i = M$. In some cases (see Section 4.2), we also allow players to be unaffiliated. To model this, we expand the set of strategies by setting $\Sigma_i = M \cup \{0\}$. Intuitively, the parties correspond to different companies, and the players correspond to the potential employees of these companies; we desire to assign employees to companies so as to maximize the total productivity.

In two of our games (see Section 4.1 and Section 4.3), a state of the game is completely described by the assignment of the players to the parties, i.e., we can write $S = (s_1, \dots, s_n)$, where $s_i \in M$ for all $i \in N$. Alternatively, we can specify a state of the game by providing a partition of the set N into m components Q_1, \dots, Q_m , where Q_j is the set of all players that chose party j , i.e., we can write $S = (Q_1, \dots, Q_m)$; we will use both forms of notation throughout the paper. In the game described in Section 4.2, the state of the game depends not only on which parties the players chose, but also on the order in which they joined the party; we postpone the formal description of this model till Section 4.2. In all three models, each player's payoff is based on the concept of marginal utility; however, in different models this idea is instantiated in different ways.

An important parameter of a state $S = (Q_1, \dots, Q_m)$ in each of these games is its *total profit* $\text{tp}(S) = \sum_{j \in M} v(Q_j)$. While for the games defined in Section 4.2 and Section 4.3, the total profit coincides with the social welfare, for the game described in Section 4.1 this is not necessarily the case. As we are interested in finding the most efficient partition of players into teams, we consider the total profit of a state a more relevant quantity than its social welfare. Hence, in what follows, we consider the price of anarchy and the price of stability with respect to the total profit, i.e., we have $\text{OPT}(\mathcal{G}) = \text{OPT}_{\text{tp}}(\mathcal{G})$, $\text{PoA}(\mathcal{G}) = \text{PoA}_{\text{tp}}(\mathcal{G})$, $\text{PoS}(\mathcal{G}) = \text{PoS}_{\text{tp}}(\mathcal{G})$.

All of our results generalize to the setting where each party $j \in M$ is associated with a different non-decreasing modular profit function $v_j : 2^N \rightarrow \mathbb{R}^+ \cup \{0\}$, i.e., different companies possess different technologies, and therefore may have different levels of productivity. Formally, any such game is given by a tuple $\mathcal{G} = (N, v_1, \dots, v_m, M)$, where $M = \{1, \dots, m\}$, and for each $j \in M$ the function v_j is a non-decreasing submodular function $v_j : 2^N \rightarrow \mathbb{R}^+ \cup \{0\}$, where $v_j(\emptyset) = 0$. In this case, the total profit function in a state $S = (Q_1, \dots, Q_m)$ is $\text{tp}(S) = \sum_{j \in M} v_j(Q_j)$. In what follows, we present our results for this more general setting.

4.1 Fair Value Games

In our first model, the utility $u_i(S)$ of a player i in a state $S = (Q_1, \dots, Q_m)$ is given by i 's marginal contribution to

the coalition he belongs to, i.e., if $i \in Q_j$, we set $u_i(S) = v_j(S) - v_j(S \setminus \{i\})$. As this payment scheme rewards each player according to the value he creates, we will refer to this type of games as *Fair Value games*. Observe that since the functions v_j are assumed to be submodular, we have $\sum_{i \in Q_j} u_i(S) \leq v_j(Q_j)$ for all $j \in M$, i.e., the total payment to the employees of a company never exceeds the profit of the company. Moreover, it may be the case that the profit of a company is strictly greater than the amount it pays to its employees; we can think of the difference between the two quantities as the owner's/shareholders' value. Consequently, in these games the total profit of all parties may differ from the social welfare, as defined in Section 2.

We will now argue that Fair Value games have a number of desirable properties. In particular, any such game is a potential game, and therefore has a pure Nash equilibrium.

Theorem 2. *Every Fair Value game \mathcal{G} is a perfect 2-nice exact potential game w.r.t. the total profit function.*

Combining Theorem 2, Lemmas 1 and 2 and Theorem 1, we obtain the following corollaries.

Corollary 1. *For every Fair Value game \mathcal{G} and every $\alpha \geq 0$ we have $\text{PoA}^\alpha(\mathcal{G}) \leq 2 + \alpha$. In particular, $\text{PoA}(\mathcal{G}) \leq 2$.*

Corollary 2. *For every Fair Value game \mathcal{G} and any $\epsilon > 0$, the basic Nash dynamic converges to a state S^F with total profit $\text{tp}(S^F) \geq \frac{\text{OPT}(\mathcal{G})}{2}(1 - \epsilon)$ in at most $\lceil \frac{n}{2} \ln \frac{1}{\epsilon} \rceil$ steps,*

Corollary 2 generalizes to basic α -Nash dynamic; we omit the exact statement of this result due to space constraints.

Since every Fair Value game is an exact potential game with the potential function given by the total profit, any profit-maximizing state is necessarily a Nash equilibrium. This implies the following proposition.

Proposition 1. *$\text{PoS}(\mathcal{G}) = 1$ for any Fair Value game \mathcal{G} .*

4.2 Labor Union Games

In Fair Value games, the player's payoff only depends on his current marginal value to the enterprise, i.e., one's salary may go down as the company expands. However, in many real-life settings, this is not the case. For instance, in many industries, especially ones that are highly unionized, an employee that has spent many years working for the company typically receives a higher salary than a new hire with the same set of skills. Our second class of games, which we will refer to as *Labor Union games*, aims to model this type of settings. Specifically, in this class of games, we modify the notion of state so as to take into account the order in which the players have joined their respective parties; the payment to each player is then determined by his marginal utility for the coalition formed by his predecessors. The submodularity property guarantees that a player's payoff never goes down as long as he stays with the same party.

Formally, in a Labor Union game \mathcal{G} that corresponds to a tuple (N, v_1, \dots, v_m, M) , we allow the players to be unaffiliated, i.e., for each $i \in N$ we set $\Sigma_i = M \cup \{0\}$. If player i plays strategy 0, we set his payoff to be 0 irrespective of the other players' strategies. A state of \mathcal{G} is given by a tuple $\mathcal{P} = (P_1, \dots, P_m)$, where P_j is the sequence of players

in party j , ordered according to their arrival time. As before, the profit of party j is given by the function v_j ; note that the value of v_j does not depend on the order in which the players join j . The payoff of each player, however, is dependent on their position in the affiliation order. Specifically, for a player $i \in P_j$, let $P_j(i)$ be the set of players that appear in P_j before i . Player i 's payoff is then defined as $u_i(\mathcal{P}) = v_j(P_j(i) \cup \{i\}) - v_j(P_j(i))$.

We remark that, technically speaking, Labor Union games are not non-cooperative games. Rather, each state of a Labor Union game induces a non-cooperative game as described above; after any player makes a move, the induced non-cooperative game changes. Abusing terminology, we will say that a state \mathcal{P} of a Labor Union game \mathcal{G} is a Nash equilibrium if for each player $i \in N$ staying with his current party is a best response in the induced game; all other notions that were defined for non-cooperative games in Section 2, as well as the results in Section 3, can be extended to Labor Union games in a similar manner.

We now state two fundamental properties of our model.

- **Guaranteed payoff:** Consider two players i and i' in P_j . Suppose i' moves to another party. The payoff of player i will not decrease. Indeed, if i' succeeds i in the sequence P_j , then by definition, i 's payoff is unchanged. If i' precedes i in P_j , then, since v_j is non-decreasing and submodular, i 's payoff will not decrease; it may, however, increase.
- **Full payoff distribution:** The sum of the payoffs of players within a party j is a telescopic sum that evaluates to $v_j(P_j)$. Therefore, the total profit $\text{tp}(\mathcal{P}) = \sum_{j \in M} v_j(P_j)$ in a state \mathcal{P} equals to the social welfare in this state. In other words, in Labor Union games, the profit of each enterprise is distributed among its employees, without creating any value for the owners/shareholders.

The guaranteed payoff property distinguishes the Labor Union games from the Fair Value games, where a player who maintains his affiliation to a party might not be rewarded, but may rather see a reduction in his payoff as other players move to join his party. This, of course, may incentivize him to shift his affiliation as well, leading to a vicious cycle of moves. In contrast, in Labor Union games, a player is guaranteed that his payoff will not decrease if he maintains his affiliation to a party. This suggests that in Labor Union games stability may be easier to achieve. In what follows, we will see that this is indeed the case.

We will first show that Labor Union games are perfect 2-nice with respect to the total profit (or, equivalently, social welfare); this will allow us to apply the machinery developed in Section 3. Abusing notation, let $\Delta_i(\mathcal{P})$ denote the improvement in the payoff of player i if he performs a best response move from \mathcal{P} , and let $\Delta(\mathcal{P}) = \sum_{i \in N} \Delta_i(\mathcal{P})$.

Proposition 2. *Any Labor Union game \mathcal{G} is a perfect 2-nice game with respect to the total profit function.*

As in the case of Fair Value games, Proposition 2 allows us to bound the price of anarchy of any Labor Union game, and the time needed to converge to a state with "good" total profit; again, Corollary 4 generalizes to basic α -Nash dynamic.

Corollary 3. For every Labor Union game \mathcal{G} and every $\alpha \geq 0$ we have $\text{PoA}^\alpha(\mathcal{G}) \leq 2 + \alpha$. In particular, $\text{PoA}(\mathcal{G}) \leq 2$.

Corollary 4. For every Labor Union game \mathcal{G} and any $\epsilon > 0$, the basic Nash dynamic converges to a state S^F with total profit $\text{tp}(S^F) \geq \frac{\text{OPT}(\mathcal{G})}{2}(1 - \epsilon)$ in at most $\lceil \frac{n}{2} \ln \frac{1}{\epsilon} \rceil$ steps, from any initial state.

Let $\mathcal{O}(\mathcal{G}) = (O_1, \dots, O_m)$ be a state that maximizes the total profit in a game \mathcal{G} , and let $\text{OPT}(\mathcal{G}) = \text{tp}(\mathcal{O}(\mathcal{G}))$. As in the case of Fair Value games, it is not hard to see that $\mathcal{O}(\mathcal{G})$ is a Nash equilibrium, i.e., $\text{PoS}(\mathcal{G}) = 1$. In fact, in Labor Union games, no coalition of players can profitably deviate from $\mathcal{O}(\mathcal{G})$ regardless of the order in which they deviate: by the guaranteed payoff property, any such deviation would not harm the non-deviators and thus would lead to a state whose total profit exceeds that of $\mathcal{O}(\mathcal{G})$, a contradiction. This implies the following result.

Proposition 3. In any Labor Union game \mathcal{G} , $\mathcal{O}(\mathcal{G})$ is a strong Nash equilibrium, i.e., the strong price of stability is 1.

Moreover, for Labor Union games under certain dynamics and certain initial states one can guarantee convergence to α -Nash equilibrium or even Nash equilibrium.

Proposition 4. Consider any Labor Union game $\mathcal{G} = (N, v_1, \dots, v_m, M)$ such that $v_j(I) \geq 1$ for any $j \in M$ and any $I \in 2^N \setminus \{\emptyset\}$. For any such \mathcal{G} , the α -Nash dynamic starting from any state in which all players are affiliated with some party converges to an α -Nash equilibrium in $O(\frac{n \log W}{\log(1+\alpha)})$ steps, where W is the maximum payoff that any player can achieve.

Proof. After each move in the α -Nash dynamic, a player improves her payoff by a factor of $1 + \alpha$, and the guaranteed payoff property ensures that payoffs of other players are unaffected. So, if a player starts with a payoff of at least 1, she will reach a payoff of W after $O(\frac{\log W}{\log(1+\alpha)})$ steps. Therefore, in $O(\frac{n \log W}{\log(1+\alpha)})$ steps, we are guaranteed to reach an α -Nash equilibrium. \square

Proposition 5. Suppose a Labor Union game \mathcal{G} with n players starts at a state in which every player is unaffiliated. Then, in exactly n steps of the Nash dynamic, the system will reach a Nash equilibrium.

4.3 Shapley Games

In our third class of games, which we call *Shapley games*, the players' payoffs are determined in a way that is inspired by the definition of the Shapley value [Shapley, 1953]. Like in Fair Value games, a state of a Shapley game is fully described by the partition of the players into parties. Given a state $S = (Q_1, \dots, Q_m)$ and a player $i \in Q_j$, we define player i 's payoff as

$$u_i(S) = \sum_{Q \subseteq Q_j \setminus \{i\}} \frac{|Q|!(|Q_j| - |Q| - 1)!}{|Q_j|!} (v_j(Q \cup \{i\}) - v_j(Q)).$$

Intuitively, the payment to each player can be viewed as his average payment in the Labor Union model, where the average is taken over all possible orderings of the players in

the party. This immediately implies $\sum_{i \in Q_j} u_i(S) = v_j(Q_j)$. Thus, Shapley games share features with both the Fair Value games and the Labor Union games. Like Fair Value games, the order in which the players join the party is unimportant. Moreover, if all payoff functions are additive, i.e., we have $u_i(S \cup \{j\}) - u_i(S) = u_i(\{j\})$ for any $i \in N$ and any $S \subseteq N \setminus \{i\}$, then the respective Shapley game coincides with the Fair Value game that corresponds to (N, v_1, \dots, v_m, M) . On the other hand, similarly to the Labor Union games, the entire profit of each party is distributed among its members. We will first show that any Shapley game is an exact potential game and hence admits a Nash equilibrium in pure strategies.

Theorem 3. Any Shapley game $\mathcal{G} = (N, v_1, \dots, v_m, M)$, is an exact potential game with the potential function given by

$$\Phi(S) = \sum_{j \in M} \sum_{Q \subseteq Q_j} \frac{(|Q| - 1)! (|Q_j| - |Q|)!}{|Q_j|!} v_j(Q).$$

Just like in other profit-sharing games, the price of anarchy in Shapley games is bounded by 2.

Theorem 4. In any Shapley game $\mathcal{G} = (N, v_1, \dots, v_m, M)$ with $|N| = n$, we have $\text{PoA}(\mathcal{G}) \leq 2 - \frac{1}{n}$.

The following claim shows that the bound given in Theorem 4 is almost tight.

Proposition 6. For any $n \geq 3$, there exists a Shapley game $\mathcal{G} = (N, v_1, v_2, M)$ with $|N| = n$ and $|M| = 2$ such that $\text{PoA}(\mathcal{G}) = 2 - \frac{2}{n+1}$ and $\text{PoS}(\mathcal{G}) = 2 - \frac{4}{n+1}$.

5 Cut Games and Profit Sharing Games

We will now describe a family of succinctly representable profit-sharing games that can be described in terms of undirected weighted graphs. It turns out that while two well-studied classes of games on such graphs do not induce profit-sharing games, a "hybrid" approach does.

In the classic *cut games* [Schäffer and Yannakakis, 1991], players are the vertices of a weighted graph $G = (N, E)$. The state of the game is a partition of players into two parties, and the payoff of each player is the sum of the weights of cut edges that are incident on him. A cut game naturally corresponds to a coalitional game with the set of players N , where the value of a coalition $S \subseteq N$ equals to the weight of the cut induced by S and $N \setminus S$. However, this game is not monotone, so it does not induce a profit-sharing game, as defined in Section 4. In *induced subgraph games* [Deng and Papadimitriou, 1994], the value of a coalition S equals to the total weight of all edges that have both endpoints in S ; while these games are monotone, they are not concave.

Finally, consider a game where the value of a coalition $S \subseteq N$ equals the total weight of all edges incident on vertices in S , i.e., both internal edges of S (as in induced subgraph games) and the edges leaving S (as in cut games). It is not hard to see that this game is both monotone and concave, and hence induces a profit-sharing game as described in Section 4. We will now explain how to compute players' payoffs in the corresponding Fair Value games, Labor Union games and Shapley games. Consider a state of the game with two parties S and $N \setminus S$ and a player $i \in S$. Let A (resp., B)

denote the total weight of edges incident on i that connect i to a predecessor (resp., successor) within S , and let C be the total weight of the cut edges incident on i . Then i 's payoff can be computed as follows:

Fair Value Games: The payoff of i is given by $\frac{A+B}{2} + C$. Intuitively, an edge between two players represents a shared skill, and i 's unique skills within a coalition are weighted more toward his payoff than his shared skills.

Labor Union Games: The payoff of i is given by $B + C$. Intuitively, i 's payoff reflects the unique skills that i possessed when he joined the party. Players who share skills with i , but join after i , will not get any payoff for those shared skills.

Shapley Games: The payoff of i is given by $\frac{A+B}{2} + C$, just as in Fair Value games.

One can see that this interpretation easily extends to multiple parties and hyperedges. We also note that many of the notions that we have discussed are naturally meaningful in this variant of the cut game: for instance, an optimal state for $m = 2$ is a configuration in which the weighted cut size is maximized.

6 Conclusions and Future Work

We have studied the dynamics of coalition formation under marginal contribution-based profit division schemes. We have introduced three classes of non-cooperative games that can be constructed from any concave cooperative game. We have shown that all three profit distribution schemes considered in this paper have desirable properties: all three games admit a Nash equilibrium, and even the worst Nash equilibrium is within a factor of 2 from the optimal configuration. In addition, for Fair Value games and Labor Union games a natural dynamic process quickly converges to a state with a fairly high total profit. Thus, when rules for sharing the payoff are fixed in advance, we can expect a system composed of bounded-rational selfish players to quickly converge to an acceptable set of teams.

Of course, the picture given by our results is far from complete; rather, our work should be seen as a first step towards understanding the behavior of myopic selfish agents in coalition formation settings. In particular, our results seem to suggest that keeping track of the history of the game and distributing payoffs in a way that respects players "seniority" leads to better stability properties; it would be interesting to see if this observation is true in practice, and whether it generalizes to other settings, such as congestion games.

Acknowledgements: this work was supported by a NAP grant, NRF (Singapore) under grant NRF RF2009-08, and A-STAR SINGA graduate fellowship.

References

[Ackermann et al., 2008] Ackermann, H., Röglin, H., and Vöcking, B. (2008). On the impact of combinatorial structure on congestion games. *J. ACM*, 55(6).

[Airiau and Sen, 2009] Airiau, S. and Sen, S. (2009). A fair payoff distribution for myopic rational agents. In *AAMAS*, pages 1305–1306.

[Awerbuch et al., 2008] Awerbuch, B., Azar, Y., Epstein, A., Mirokni, V. S., and Skopalik, A. (2008). Fast convergence to nearly optimal solutions in potential games. In *EC*, pages 264–273.

[Bhalgat et al., 2010] Bhalgat, A., Chakraborty, T., and Khanna, S. (2010). Approximating pure Nash equilibrium in cut, party affiliation, and satisfiability games. In *EC*, pages 73–82.

[Chalkiadakis and Boutilier, 2004] Chalkiadakis, G. and Boutilier, C. (2004). Bayesian reinforcement learning for coalition formation under uncertainty. In *AAMAS*, pages 1090–1097.

[Chatterjee et al., 1993] Chatterjee, K., Dutta, B., Ray, D., and Sen-gupta, K. (1993). A noncooperative theory of coalitional bargaining. *Review of Economic Studies*, 60(2):463–77.

[Chien and Sinclair, 2007] Chien, S. and Sinclair, A. (2007). Convergence to approximate Nash equilibria in congestion games. In *SODA*, pages 169–178.

[Deng and Papadimitriou, 1994] Deng, X. and Papadimitriou, C. H. (1994). On the complexity of cooperative solution concepts. *Math. Oper. Res.*, 19:257–266.

[Dieckmann and Schwalbe, 2002] Dieckmann, T. and Schwalbe, U. (2002). Dynamic coalition formation and the core. *Journal of Economic Behavior and Organization*, 49:363–380.

[Fabrikant et al., 2004] Fabrikant, A., Papadimitriou, C. H., and Talwar, K. (2004). The complexity of pure Nash equilibria. In *STOC*, pages 604–612.

[Fanelli et al., 2008] Fanelli, A., Flammini, M., and Moscardelli, L. (2008). The speed of convergence in congestion games under best-response dynamics. In *ICALP (1)*, pages 796–807.

[Gairing and Savani, 2010] Gairing, M. and Savani, R. (2010). Computing stable outcomes in hedonic games. In *SAGT*, pages 174–185.

[Gairing and Savani, 2011] Gairing, M. and Savani, R. (2011). Computing stable outcomes in hedonic games with voting-based deviations. In *AAMAS*, to appear.

[Moldovanu and Winter, 1995] Moldovanu, B. and Winter, E. (1995). Order independent equilibria. *Games and Economic Behavior*, 9(1):21–34.

[Monderer and Shapley, 1996] Monderer, D. and Shapley, L. S. (1996). Potential games. *Games and Economic Behavior*, 14(1):124–143.

[Rosenthal, 1973] Rosenthal, R. W. (1973). A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67.

[Sandholm and Lesser, 1997] Sandholm, T. and Lesser, V. (1997). Coalitions among computationally bounded agents. *Artif. Intell.*, 94:99–137.

[Schäffer and Yannakakis, 1991] Schäffer, A. A. and Yannakakis, M. (1991). Simple local search problems that are hard to solve. *SIAM J. Comput.*, 20(1):56–87.

[Shapley, 1953] Shapley, L. S. (1953). A value for n-person games. *Contributions to the theory of games*, 2:307–317.

[Shehory and Kraus, 1999] Shehory, O. and Kraus, S. (1999). Feasible formation of coalitions among autonomous agents in non-superadditive environments. *Comp. Intell.*, 15:218–251.

[Skopalik and Vöcking, 2008] Skopalik, A. and Vöcking, B. (2008). Inapproximability of pure Nash equilibria. In *STOC*, pages 355–364.

[Yan, 2003] Yan, H. (2003). Noncooperative selection of the core. *International Journal of Game Theory*, 31:527–540.