Strong Price of Anarchy, Utility Games and Coalitional Dynamics

Yoram Bachrach[1], Vasilis Syrgkanis[2,⋆], Éva Tardos[2,⋆⋆], and Milan Vojnović[1]

[1] Microsoft Research Cambridge, UK
yobach,milanv@microsoft.com
[2] Cornell University
{vasilis,eva}@cs.cornell.edu

Abstract. We introduce a framework for studying the effect of cooperation on the quality of outcomes in utility games. Our framework is a coalitional analog of the smoothness framework of non-cooperative games. Coalitional smoothness implies bounds on the strong price of anarchy, the loss of quality of coalitionally stable outcomes. Our coalitional smoothness framework captures existing results bounding the strong price of anarchy of network design games. Moreover, we give novel strong price of anarchy results for any monotone utility-maximization game, showing that if each player's utility is at least his marginal contribution to the welfare, then the strong price of anarchy is at most 2. This captures a broad class of games, including games that have a price of anarchy as high as the number of players. Additionally, we show that in potential games the strong price of anarchy is close to the price of stability, the quality of the best Nash equilibrium.

We also initiate the study of the quality of coalitional out-of-equilibrium outcomes in games. To this end, we define a coalitional version of myopic best-response dynamics, and show that the bound on the strong price of anarchy implied by coalitional smoothness, also extends with small degradation to the average quality of outcomes of the given dynamic.

1 Introduction

We introduce a framework for studying the effect of cooperation on the quality of outcomes in games. In the past decade we have developed a good understanding of the degradation in social welfare in games due to selfish play, quantified by the price of anarchy. There are known tight bounds on the price of anarchy in a range of games from routing, to network design, to various scheduling games. Much less is understood about outcomes of games where players may cooperate.

In many settings players do cooperate, and cooperation can help improve the outcome. The worst possible Nash equilibrium is a very pessimistic prediction

⋆ Work done in part as an intern with Microsoft Research. Supported in part by ONR grant N00014-98-1-0589, NSF grant CCF-0729006, Simons Graduate Fellowship.

⋆⋆ Supported in part by NSF grants CCF-0910940 and CCF-1215994, ONR grant N00014-08-1-0031, a Yahoo! Research Alliance Grant, and a Google Research Grant.

R. Lavi (Ed.): SAGT 2014, LNCS 8768, pp. 218–230, 2014.

of the outcome in games that are not strictly competitive, and where cooperation may improve the utility for all participants. A key issue in understanding cooperative outcomes is the extent to which players can transfer utility among each-other. The two dominant notions of cooperative outcomes considered in the literature are the strong Nash equilibrium of Aumann [5] assuming no utility transfer between players, and the transferable utility notion of the core (see [13] for a survey). Allowing utility transfers between the players leads to an extremely demanding form of equilibrium, a solution is unstable in this sense, if there is a possible joint deviation for a group that improves the total utility of a group, even if this is not improving the utility of every single player. A stable outcome in this sense is automatically socially optimal (otherwise the grand coalition can deviate). An outcome is a strong Nash equilibrium if it is stable subject to coalitional deviations, meaning that no group of players can jointly deviate to improve the solution for every member of the coalition. Strong Nash equilibria do not imply the optimality of the outcome. We identify properties of a game such that a stable outcome cannot be too far from the social optimum. An even less restrictive notion, coalition proof equilibrium, was introduced by Moreno et al [14] requiring extra conditions from a coalitional deviation to be feasible, such as the non-existence of subsequent unilateral deviations. We focus here mainly on strong Nash equilibria and on randomized versions.

The strong price of anarchy was introduced by Andelman et al. [2] and measures the quality degradation of strong Nash equilibria in games. One of the most compelling examples is the class of cost-sharing games, where players choose costly resources and equally share the cost with other users of each resource. Anshelevich et al. [3] showed that the price of anarchy in this class of games with n players can be as bad as n, but showed a tight $H_n = O(\log n)$ bound on the price of stability, the quality loss in the best Nash equilibria compared to the socially optimal solution. While the worst Nash equilibria seems too pessimistic a prediction for the outcome, the best Nash equilibria is potentially too optimistic: while significant cooperation is needed to identify and reach this solution, the stability concept used is that of Nash equilibria, assuming that only individual players can deviate, and not groups. Epstein et al. [8] showed an H_n bound on the the strong price of anarchy, matching the price of stability bound.

For the case of worst-case Nash equilibria, Roughgarden [16] introduced the framework of *smooth games*, encompassing most price of anarchy bounds. However, such a unified framework does not exist for worst-case coalitionally stable outcomes. We propose a smoothness framework that captures efficiency in most well-established cooperative equilibrium solution concepts such as the strong Nash equilibrium and randomized versions of it. We show how our framework implies existing results on the strong price of anarchy of cost-sharing games and we generate new results on the strong price of anarchy of utility maximization games and of potential games.

The second goal of our paper is to initiate a study of outcomes of dynamic cooperative play and of efficiency of non-equilibrium solution concepts. A key point in the wide applicability of the price of anarchy analysis and of the smoothness framework of [16] is that bounds proved via smoothness automatically extend to coarse correlated equilibria, which are outcomes of no-regret learning by each player [7]. Extending the price of anarchy results to no-regret outcomes is appealing as it is a natural model of player behavior, and no-regret can be achieved via simple strategies.

Studying the efficiency of out-of-equlibrium solution concepts that capture cooperation is even more compelling in the strong price of anarchy analysis, as strong Nash equilibria, unlike for instance mixed Nash equilibria, are not guaranteed to exist, and do not exist in even small and simple cost-sharing games [8]. Hence, we need to identify properties of games that would not only imply approximate efficiency of coalitionally stable equilibrium outcomes, but whose efficiency implications would directly extend with very small degradation even to out-of-equilibrium cooperative dynamic solution concepts.

To this end, we propose a coalitional version of myopic best-response dynamics and we analyze the average welfare of such dynamics in the long run. Our dynamics can be viewed as a coalitional version of sink equilibria proposed by Goemans et al. [9]. More importantly, we show that the efficiency guarantees implied by our coalitional smoothness framework directly extend with small loss to this form of out-of-equilibrium cooperative dynamics. These out-of-equilibrium outcomes always exist and are meaningful even in games that do not admit a strong Nash equilibrium and thereby the direct extension is rather appealing as it provides an efficiency bound that is not conditional on existence.

Our Results. We propose a framework for quantifying the quality of strong Nash equilibria by introducing the notion of coalitional smoothness. We show how coalitional smoothness captures existing results on network design games, we give new results on the strong price of anarchy in utility maximization games, and show that coalitional smoothness in such games implies high social welfare at coalitional sink equilibria, which we define as the out-of-equilibrium myopic behavior as defined by a natural coalitional version of best-response dynamics.

- We define the notion of a (λ, μ)-*coalitionally smooth games* and show that the strong price of anarchy of a (λ, μ)-coalitionally smooth game is bounded by $\lambda/(1+\mu)$ in utility games and $\lambda/(1-\mu)$ in cost minimization games.
- We show that the cost-sharing games of [8] as well as network contribution games [4] studied in the literature are coalitionally smooth.
- We show that in any monotone utility-maximization game, if each player's utility is at least his marginal contribution to the welfare then the strong price of anarchy is at most 2, while the price of anarchy in this class of games can be as high as n. This result complements the results of [22,9] who studied the price of anarchy of utility-maximization games that have submodular social welfare function.
- In potential games, such as the cost-sharing game of [8], the potential minimizer is a Nash equilibrium of high quality. This equilibrium is typically used to bound the price of stability by showing that the social welfare function is similar to the potential function, namely $\lambda \cdot SW(s) \leq \Phi(s) \leq \mu \cdot SW(s)$,

implying a bound of λ/μ on the price of stability. We show that in utility games this condition also implies that the game is (λ, μ)-coalitionally smooth implying a $\frac{\lambda}{1+\mu}$ bound on the strong price of anarchy, and give conditions for a similar bound in cost-minimization games, extending the work of [8].

- Strong price of anarchy bounds via coalitional smoothness also extend to the notions of strong correlated equilibria (see e.g. Moreno et al. [14]) and strong coarse correlated equilibria of [19], which correspond to randomized outcomes where no group of players C has a joint distribution of strategies $\tilde{D}_C$ that each member of the group has regret for. Though there exist games with no strong Nash, that admit such randomized strong equilibria, unfortunately, there are no learning algorithms that guarantee this coalitional no-regret property, and in fact, these concepts may not exist in some games.
- We define a natural *coalitional best response dynamic* and the corresponding *coalitional sink equilibria*, the analog of the notion of myopic sink equilibria introduced by Goemans et al. [9] for coalitional dynamics. While myopic sink equilibria correspond to steady state behavior of the Markov chain defined by iteratively doing random unilateral best respond dynamics, coalitional sink equilibria are the steady state under our coalitional best response dynamic. We do not explicitly model how players choose to transfer utility to each other. However, our dynamic assumes that when a group cooperates, then they can also transfer utility, and hence will choose to optimize the total utility of all group members. We show that in (λ, μ)-coalitionally smooth utility games the social welfare of any coalitional sink equilibrium is at least a $\frac{1}{H_n}\frac{\lambda}{1+\mu}$ fraction of the optimal; extending our analysis of outcomes of coalitional play to games when strong Nash equilibria do not exist.

Related Work. The study of efficiency of worst-case Nash equilibria via the price of anarchy was initiated by [12], and has triggered a large body of work. Roughgarden [16] introduced a canonical way of analyzing the price of anarchy by proposing the notion of a (λ, μ)-smooth game and showing that most efficiency proofs can be cast as showing that the game is smooth. Most importantly, [16] showed that any efficiency proven via smoothness arguments directly extends to outcomes of no-regret learning behavior. Recently, similar frameworks have been proposed for games of incomplete information [17,20,21] and games with continuous strategy spaces [18]. However, these frameworks do not take into account coalitional robustness and no canonical way of showing efficiency bounds for coalitional solution concepts existed prior to our work.

The most established coalitionally robust solution concept is that of the strong Nash equilibrium introduced by Aumann [5]. The study of the efficiency of the worst strong Nash equilibrium (strong price of anarchy) was introduced in [2], and follow-up research mostly focused on specific cost minimization games such as network design games [3,1,8]. Our coalitional smoothness framework captures some of the results in this literature and gives a generic condition under which the strong price of anarchy is bounded.

For utility maximization games Vetta [22] defined the class of valid-utility games, which are utility maximization games with a monotone and submodular welfare function and where each player's utility is at least his marginal contribution to the welfare. Vetta [22] showed that every Nash equilibrium of a valid utility game achieves at least half of the optimal welfare. Later these games were analyzed from the perspective of best response dynamics by [9], who introduced the notion of a sink equilibrium (i.e. steady state distribution of the Markov chain defined by best-response dynamics) and showed that for a subclass of valid-utility games the half approximation is achieved after polynomially many rounds, while for the general class, the sink equilibria can have an efficiency that degrades linearly with the number of players. In this paper, we show that without the assumption of submodularity every monotone utility maximization game that satisfies the marginal contribution condition has good strong price of anarchy. Additionally, we define a coalitional version of sink equilibria of [9] and show that for any coalitionally smooth game the efficiency bound at these out-of-equilibrium dynamics degrades from the strong price of anarchy bound by a factor that is only logarithmic in the number of players.

The efficiency of coalitionally robust solution concepts was also studied by Anshelevich et al [4] for a class of contribution games in networks, where pairwise-stable outcomes were analyzed. Most of our theorems imply social welfare bounds for strong Nash equilibria of network contribution games, that hold under much more general assumptions than the ones considered in [4].

The existence of strong Nash equilibria was examined by both game theorists and computer scientists (see e.g. [15,10,8]). [19] show that in singleton congestion games with increasing resource value functions there always exists a strong Nash equilibrium, while [11] show that for decreasing function the set of pure Nash equilibria, which is non-empty, coincides with the set of strong Nash equilibria.

2 Coalitional Smoothness

In this section we introduce the notion of coalitional smoothness and show that it captures the core of a proof on the efficiency of strong Nash equilibria in several games studied in the past, such as network cost sharing games [3,8] as well as in new classes of games that we give, which generalize the well-studied valid-utility games of Vetta [22], by dropping the assumption of submodularity.

For ease of presentation we will present the definition of coalitional smoothness for utility maximization games rather than cost minimization, but the definitions naturally extend to analogous ones for cost minimization. We will consider a standard normal form game among n players. Each player i has a strategy space S_i and a utility $u_i : S_1 \times \ldots \times S_n \to \mathbb{R}_+$. For a subset of players $C \subseteq [n]$ we will denote with $S_C = (S_i)_{i\in C}$ the joint strategy space, with $s_C \in S_C$ a joint strategy profile and with $\Delta(S_C)$ the space of distributions over strategy profiles. We are interested in quantifying the efficiency of coalitional solution concepts with respect to the social welfare, which is defined as the sum of all player utilities: $SW(s) = \sum_{i\in[n]} u_i(s)$. For convenience, we will denote with OPT the maximal social welfare (resp. minimum social cost) achieved among all possible strategy profiles and we will try to upper bound the *price of anarchy*, which is the ratio of the optimal social welfare over the social welfare at any equilibrium in the class of solution concepts that we study (e.g. strong price of anarchy for the case of strong Nash equilibria), or equivalently to lower bound the fraction of the optimal welfare that every equilibrium in the class achieves.

The intuition behind coalitional smoothness is that it requires from the game to admit a good strategy profile such that if enough players coalitionally deviate to this strategy from any state with low social welfare then they achieve a good fraction of the optimal welfare. Specifically, it imposes that if we order the players arbitrarily and consider only the coalitional deviations of all the suffixes of this order, then the sum of utilities of the first player in each of the suffixes, after the coalitional deviation of the suffix, is at least a λ fraction of the optimal welfare or else μ times the current social welfare is at least a λ fraction of the optimal.

Definition 1 (Coalitional Smoothness). *A utility maximization game is (λ, μ)-coalitionally smooth if there exists a strategy profile s^* such that for any strategy profile s and for any permutation π of the players:*

$$\textstyle\sum_{i=1}^n u_i(s^*_{N_{\pi(i)}}, s_{-N_{\pi(i)}}) \geq \lambda \cdot OPT - \mu \cdot SW(s) \quad (1)$$

where $N_{\pi(i)} = \{j \in [n] : \pi(j) \geq \pi(i)\}$ is the set of all players succeeding i in the permutation and (s_{N_t}, s_{-N_t}) is the strategy profile where all players in $i \in N_t$ play s^*_i and all other players play s. [1]

We now formally define the notion of a strong Nash equilibrium introduced by Aumann [5] and show that coalitional smoothness implies high efficiency at every strong Nash equilibrium of a game.

Definition 2 (Strong Nash Equilibrium). *A strategy profile s is a strong Nash equilibrium if for any coalition $C \subseteq [n]$ and for any coalitional strategy $s_C \in S_C$, there exists a player $i \in C$ such that: $u_i(s) \geq u_i(s_C, s_{-C})$.*

Theorem 3. *If a game is (λ, μ)-coalitionally smooth for some $\lambda, \mu \geq 0$ then every strong Nash equilibrium has social welfare at least $\frac{\lambda}{1+\mu}$ of the optimal.* [2]

Proof. Let s be strong Nash equilibrium strategy profile and let s^* be the optimal strategy profile. If all players coalitionally deviate to s^* then, by the definition of a strong Nash equilibrium, there is a player i who is blocking the deviation, i.e. $u_i(s) \geq u_i(s^*)$. Without loss of generality, reorder the players such that this is player 1. Similarly, if players $\{2, \ldots, n\}$ deviate to playing their strategy in s^* then there exists some player, obviously different than 1 who is blocking the deviation. Without loss of generality, by reordering we can assume that this player is 2. Using similar reasoning we can reorder the players such that if players $\{i, \ldots, n\}$ deviate to their strategy in the optimal strategy profile x^* then player

[1] In the case of cost minimization games we would require: $\sum_{i=1}^n c_i(s^*_{N_{\pi(i)}}, s_{-N_{\pi(i)}}) \leq \lambda \cdot SC(s^*) + \mu \cdot SC(s)$.

[2] In cost-minimization games (λ, μ)-coalitional smoothness for $\lambda \geq 0$, $\mu \leq 1$ implies that the social cost at a strong Nash is at most $\frac{\lambda}{1-\mu}$ of the minimum cost.

i is the one blocking the deviation. That is player i's utility at the strong Nash equilibrium is at least his utility in the deviating strategy profile. Thus under this order $\forall k \in N$: $u_i(s) \geq u_i(s^*_{N_k}, s_{-N_k})$. Summing over all players and using the coalitional smoothness property for the above order we get the result:

$$SW(s) = \sum_{i=1}^{N} u_i(s) \geq \sum_{i=1}^{N} u_i(s^*_{N_i}, s_{-N_i}) \geq \lambda SW(s^*) - \mu SW(s)$$

□

Extension to Randomized Solution Concepts. Similar to smoothness, coalitional smoothness also implies efficiency bounds even for randomized coalition-proof solution concepts. Adapting randomized solution concepts such as correlated equilibria so as to make them robust to coalitional deviations is not as straightforward as in the case of unilateral stability. This is mainly due to information considerations. One such concept is that of strong correlated equilibria (see e.g. Moreno and Wooders [14]) and it's relaxation, the strong coarse correlated equilibrium (see e.g. Rozenfeld et al. [19]).

Essentially, a strong correlated equilibrium is a distribution over strategy profiles, such that if players are recommended strategies based on this distribution, then there exists no coalitional (randomized) deviation under which every member of the coalition is strictly better off. The coalitional deviation can depend on the recommendation that each member of the coalition received. Thus implicitly it is assumed that if players commit to a coalition ex-ante, then after receiving their recommendations on which strategy to play, they share it publicly among the players in the coalition and decide on a joint deviation. The strong coarse correlated equilibrium is a relaxation of the above concept where the coalitional deviation is independent of the recommendations. We defer a formal definition of the two concepts and a more elaborate discussion to the full version [6].

Non-Submodular Monotone Utility Games. Consider a utility maximization game in which every player has an s_i^{out} strategy, corresponding to the player not entering the game. Further assume that the game is monotone with respect to participation, i.e. no player can decrease the social welfare by entering the game: $\forall i \in [n], \forall s : SW(s) \geq SW(s_i^{out}, s_{-i})$. We show that the coalitional smoothness of such a game is captured exactly by the proportion of the marginal contribution to the social welfare that a player is guaranteed to get as utility.

Theorem 4. *Any monotone utility maximization game is guaranteed to be (γ, γ)-coalitionally smooth, if each player is guaranteed at least a γ fraction of his marginal contribution to the social welfare:*

$$\forall s : u_i(s) \geq \gamma \left(SW(s) - SW(s_i^{out}, s_{-i})\right) \tag{2}$$

Proof. Consider any order of the players and let s^* be the strategy profile that maximizes the social welfare. By the marginal contribution property we have:

$$\sum_{i=1}^{n} u_i(s^*_{N_i}, s_{-N_i}) \geq \gamma \cdot \sum_{i=1}^{n} \left(SW\left(s^*_{N_i}, s_{-N_i}\right) - SW\left(s_i^{out}, s_{N^*_{i+1}}, s_{-N_i}\right)\right)$$

In addition, by the monotonicity assumption the social welfare can only increase when a player enters the game with any strategy:

$$SW(s_i^{out}, s^*_{N_{i+1}}, s_{-N_i}) \le SW(s_k, s^*_{N_{i+1}}, s_{-N_i}) = SW(s^*_{N_{i+1}}, s_{-N_{i+1}})$$

Combining the above inequalities we get a telescoping sum that yields the desired property:

$$\begin{aligned}\textstyle\sum_{i=1}^n u_i(s^*_{N_i}, s_{-N_i}) \ge&\ \gamma\cdot\textstyle\sum_{i=1}^n \Big(SW(s^*_{N_i}, s_{-N_i}) - SW(s_{N^*_{i+1}}, s_{-N_{i+1}})\Big)\\ \ge&\ \gamma\cdot SW(s^*) - \gamma\cdot SW(s) = \gamma\cdot \mathrm{OPT} - \gamma\cdot SW(s)\end{aligned}$$

Which is exactly the (γ,γ)-coalitional smoothness property we wanted. □

This latter result complements Vetta's [22] results on valid-utility games. A valid-utility game is a monotone utility-maximization game with the following additional structural property: strategies of the players can be viewed as sets of some ground set of elements and the social welfare can be viewed as a monotone *submodular* set function on the union of the chosen strategies. As re-interpreted by Roughgarden [16], Vetta showed that in any monotone utility-maximization game with a submodular welfare function, if each player receives a γ fraction of their marginal contribution to the welfare, then the game is (γ,γ)-smooth implying that every Nash equilibrium achieves a $\frac{\gamma}{\gamma+1}$ fraction of the optimal welfare. In the absence of submodularity there are easy examples where the worst Nash equilibrium doesn't achieve any constant fraction of the optimal welfare, despite satisfying the marginal contribution condition. However, our result shows that even in the absence of submodularity every such game will be (γ,γ)-coalitionally smooth, implying that every strong Nash equilibrium will achieve a $\frac{\gamma}{\gamma+1}$ fraction of the welfare.

It is important to note that the approximate marginal contribution condition and the submodularity condition are very orthogonal ones. For instance, it is possible that a game satisfies the approximate marginal contribution condition for some constant, but is not submodular or even approximately submodular under existing definitions of approximate submodularity. In the full version of the paper, we give a class of welfare sharing games, where our efficiency theorem applies to give constant bounds on the string price of anarchy, whilst the price of anarchy is unbounded due to the non-submodularity of the social welfare.

Network Cost-Sharing Games. In this section we analyze the well-studied class of cost sharing games [3], using the coalitional smoothness property. The game is defined by a set of resources R each associated with a cost c_r. Each player's strategy space S_i is a set of subsets of R. The cost of each resource is shared equally among all players that use the resource and a players total cost is the sum of his cost-shares on the resources that he uses. If we denote with $n_r(s)$ the number of players using resource r under strategy profile s, then: $c_i(s) = \sum_{r\in s_i} \frac{c_r}{n_r(s)}$.

Epstein et al [8] showed that every strong Nash equilibrium of the above class of games has social cost at most H_n times the optimal, where H_n is the n-th harmonic number. Here we re-interpret that result as showing that these games are $(H_n, 0)$-coalitionally smooth. In the next section we show that the analysis of [8] can be applied to a more broad class of potential games, showing a strong connection between the price of stability and the strong price of anarchy.

Theorem 5 ([8]). *Cost sharing games are $(H_n, 0)$-coalitionally smooth.*

3 Best Nash vs. Worst Strong Nash Equilibrium

Strong Nash equilibria are a subset of Nash equilibria, so in games when strong Nash equilibria exist, the strong price of anarchy cannot be better than the price of stability (the quality of best Nash). In this section we show that in potential games these two notions are surprisingly close. We show that through the lens of coalitional smoothness there is a strong connection between the analysis of the efficiency of the worst strong Nash equilibria and the dominant analysis of the best Nash equilibria in potential games. A game admits a potential function if there exists a common function $\Phi(s)$ for all players, such that a player's difference in utility from a unilateral deviation is equal to difference in the potential:

$$u_i(s'_i, s_{-i}) - u_i(s) = \Phi(s'_i, s_{-i}) - \Phi(s) \tag{3}$$

A large amount of recent work in the algorithmic game theory literature has focused on the analysis of the efficiency of the best Nash equilibrium (price of stability). For the case of potential games the dominant way of analysing the price of stability is the Potential Method: suppose that the potential function is (λ, μ)-*close* to the social welfare, in the sense that $\lambda \cdot SW(s) \le \Phi(s) \le \mu \cdot SW(s)$, for some parameters $\lambda, \mu \ge 0$. Then the best Nash equilibrium achieves at least $\frac{\lambda}{\mu}$ of the optimal social welfare. The proof relies on the simple fact that the potential maximizer is always a Nash equilibrium and by the (λ, μ) property it's easy to see that the potential maximizer has social welfare that is the above fraction of the optimal social welfare.

The following theorems show that for such potential games the price of stability is very close to the strong price of anarchy, i.e. the implied quality of the best Nash equilibrium is close to the quality of the worst strong Nash equilibrium.

Theorem 6. *In a utility-maximization potential game with non-negative utilities, if the potential is (λ, μ)-close to the social welfare then the game is (λ, μ)-coalitionally smooth, implying that every strong Nash equilibrium achieves at least $\frac{\lambda}{1+\mu}$ of the optimal social welfare.*

Proof. Consider an arbitrary order of the players and some strategy profile s. By the definition of the potential function and the fact that utilities are non-negative, we have

$$\begin{aligned} u_i(s^*_{N_i}, s_{-N_i}) &= \Phi(s^*_{N_i}, s_{-N_i}) - \Phi(s^*_{N_{i+1}}, s_{-N_{i+1}}) + u_i(s^*_{N_{i+1}}, s_{-N_{i+1}}) \\ &\ge \Phi(s^*_{N_i}, s_{-N_i}) - \Phi(s^*_{N_{i+1}}, s_{-N_{i+1}}) \end{aligned}$$

Combining with our assumption on the relation between potential and social welfare we obtain the coalitional smoothness property:

$$\begin{aligned}\textstyle\sum_{i=1}^n u_i(s^*_{N_i}, s_{-N_i}) \geq & \textstyle\sum_{i=1}^n \left(\Phi(s^*_{N_i}, s_{-N_i}) - \Phi(s^*_{N_{i+1}}, s_{-N_{i+1}})\right) \\ = & \ \Phi(s^*) - \Phi(s) \geq \lambda \cdot SW(s^*) - \mu \cdot SW(s)\end{aligned}$$

□

The (λ, μ)-closeness property of the potential function does not imply smoothness of the game according to the standard definition of smoothness [16] and hence a price of anarchy bound. It does so only if the potential is a submodular function and by following a similar analysis as in the case of valid utility games as we show in the full version. Such a property for instance, holds in utility congestion games with decreasing resource utilities. However, Theorem 6 does not require submodularity of the potential.

One application of Theorem 6 is in the context of network contribution games [4]. In a network contribution game each player corresponds to a node in a network. Each edge corresponds to a "friendship" between the connecting nodes or more generally some joint venture. Each player has a budget of effort that he strategically distributes among his friendships. Each friendship e between two players i and j, has a value $v_e(x_i, x_j)$ that corresponds to the value produced as a function of the efforts put into it by the two players. This value is equally split among the two players. It is easy to see that in such a game the social welfare is the total value produced, while the potential is equal to half of the social welfare. Thus, by applying Theorem 6 we get that for arbitrary "friendship" value functions $v_e(\cdot, \cdot)$ the game is $\left(\frac{1}{2}, \frac{1}{2}\right)$-coalitionally smooth and hence every strong Nash equilibrium achieves at least $\frac{1}{3}$ of the optimal welfare. In contrast, observe that Nash equilibria can have unbounded inefficiency,[3] and the game is not (λ, μ)-smooth under the unilateral notion of smoothness for any λ, μ.

For settings where a player can only have non-negative externalities on the utilities of other players by entering the game, a much stronger connection can be drawn. More concretely, a utility maximization game has non-negative externalities if for any strategy profile s and for any pair of players i, j: $u_i(s) \geq u_i(s_j^{out}, s_{-j})$.[4] The s_i^{out} strategy is not required to be a valid strategy that the player can actually pick, but rather a hypothetical strategy, requiring the property that the cost of the player in that strategy is 0, and the cost functions and the potential are extended appropriately such that the potential function property is maintained even in this augmented strategy space and the potential when all players have left the game is 0: $\Phi(s^{out}) = 0$. For instance, every congestion game has the above property if s_i^{out} is defined as the empty set of resources.

[3] Consider a line of four nodes (A, B, C, D). Each player has budget 1. Edges (A, B) and (C, D) have constant value of 1, while edge (B, C) has a huge value H, if both players put all their budget on it and 0 o.w.. Players B, C placing their budget on their alternative friendships is a Nash equilibrium, but not a strong Nash equilibrium.

[4] A cost-minimization game has non-negative externalities if $c_i(s) \leq c_i(s_j^{out}, s_{-j})$.

Theorem 7. *A utility-maximization potential game with only positive externalities and such that $\Phi(s) \geq \lambda \cdot SW(s)$ is $(\lambda, 0)$-coalitionally smooth. Similarly, a cost-minimization, potential game with only positive externalities and such that $\Phi(s) \leq \lambda \cdot SC(s)$ is $(\lambda, 0)$-coalitionally smooth.*

In the context of cost-minimization, one well-studied example of such a setting is that of network cost-sharing games and the $log(n)$ strong price of anarchy result of [8] is a special instance of Theorem 7. For utility-maximization games, one example is that of network contribution games under the restriction that friendship value functions $v_e(\cdot, \cdot)$ are increasing in both coordinates. Under this restriction applying Theorem 7 we get the improved bound that every strong Nash equilibrium achieves at least 1/2 of the optimal social welfare.

4 Coalitional Best-Response Dynamics

In this section we initiate the study of efficiency of dynamic coalitional behavior. We show that if a utility game is (λ, μ)-coalitionally smooth then this implies an efficiency guarantee for out of equilibrium dynamic behavior in a certain best-response like dynamic. This is particularly interesting for games that do not admit a strong Nash equilibrium, but where coalitional deviations are bound to occur. Our approach is similar in spirit to the notion of myopic sink equilibria introduced by Goemans et al. [9]. Myopic sink equilibria correspond to steady state behavior of the Markov chain defined by iteratively doing random unilateral best respond dynamics. However, such a notion does not capture settings where players can communicate and at each step perform coalitional deviations.

We introduce a version of coalitional best-response dynamics, that allows for coalitional deviations at each time step, giving more probability to small coalitions. In our dynamic, at each step a selected group is chosen to cooperate. We assume that when a group cooperates, then they can also transfer utility, and hence will choose to optimize the total utility of all group members. Then we analyze the social welfare of the steady states arising in the long run as we perform coalitional best response dynamics for a long period. Similar to [9] we will refer to these steady states as coalitional sink equilibria. Similar to sink equilibria that are a way of studying games whose best response dynamics might not converge to a pure Nash equilibrium or even games that do not admit a pure Nash equilibrium, coalitional sink equilibria are an interesting alternative for analyzing efficiency in games that do not admit a strong Nash equilibrium, which admittedly is even more rare than the pure Nash equilibrium.

Our coalitional dynamics are as follows: At each iteration a coalition is picked at random from a distribution that favors coalitions of smaller size. Specifically, first a coalition size k is picked inversely proportional to the size and then a coalition of size k is picked uniformly at random. Next, the picked coalition deviates to the joint strategy profile that maximizes the total utility of the coalition, conditional on the current strategy of every player outside of the coalition.

Theorem 8. *If a utility maximization game with non-negative utilities is* (λ, μ)*-coalitionally smooth then the expected social welfare at every coalitional sink equilibrium is at least* $\frac{1}{H_n}\frac{\lambda}{1+\mu}$ *of the optimal.*

If we picked the coalitional size k, according to some probability distribution with a density $p(k)$ satisfying $p(k) \geq \frac{1}{c}\frac{1}{H_n \cdot k}$, then by the analysis of the above theorem we get that every coalitional sink equilibrium achieves welfare at least $\frac{1}{c \cdot H_n}\frac{\lambda}{1+\mu}$ of the optimal.

Remark 1. The Markov chain defined by the coalitional best-response dynamics might take long time to converge to a steady state. However, our analysis shows a stronger statement: at any iteration T if we take the empirical distribution defined by the best-response play up till time T, then the expected welfare of this empirical distribution is at least $\frac{T-1}{2T}\frac{\lambda}{H_n+\mu}$ of the optimal welfare.

References

1. Albers, S.: On the value of coordination in network design. In: SODA (2008)
2. Andelman, N., Feldman, M., Mansour, Y.: Strong price of anarchy. Games and Economic Behavior 65(2), 289–317 (2009)
3. Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, E., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. In: FOCS, pp. 295–304 (2004)
4. Anshelevich, E., Hoefer, M.: Contribution Games in Networks. Algorithmica, 1–37 (2011)
5. Aumann, R.J.: Acceptable points in general cooperative N-person games. In: Luce, R.D., Tucker, A.W. (eds.) Contribution to the theory of game IV, Annals of Mathematical Study 40, pp. 287–324. University Press (1959)
6. Bachrach, Y., Syrgkanis, V., Tardos, É., Vojnovic, M.: Strong price of anarchy and coalitional dynamics. CoRR, abs/1307.2537 (2013)
7. Blum, A., Mansour, Y.: Learning, Regret Minimization and Equilibria. Camb. Univ. Press (2007)
8. Epstein, A., Feldman, M., Mansour, Y.: Strong equilibrium in cost sharing connection games. Games and Economic Behavior 67(1), 51–68 (2009)
9. Goemans, M., Mirrokni, V., Vetta, A.: Sink equilibria and convergence. In: FOCS, pp. 142–154 (2005)
10. Harks, T., Klimm, M., Möhring, R.H.: Strong nash equilibria in games with the lexicographical improvement property. In: Leonardi, S. (ed.) WINE 2009. LNCS, vol. 5929, pp. 463–470. Springer, Heidelberg (2009)
11. Holzman, R., Law-Yone, N.: Strong equilibrium in congestion games. Games and Economic Behavior 21(1–2), 85–101 (1997)
12. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999)
13. Maschler, M.: The bargaining set, kernel, and nucleolus. In: Handbook of Game Theory with Economic Applications, vol. 1, ch.18, pp. 591–667. Elsevier (1992)
14. Moreno, D., Wooders, J.: Coalition-proof equilibrium. Games and Economic Behavior 17(1), 80–112 (1996)
15. Nessah, R., Tian, G.: On the existence of strong nash equilibria. Working Papers 2009-ECO-06, IESEG School of Management (2009)

16. Roughgarden, T.: Intrinsic robustness of the price of anarchy. In: STOC (2009)
17. Roughgarden, T.: The price of anarchy in games of incomplete information. In: ACM EC (2012)
18. Roughgarden, T., Schoppmann, F.: Local smoothness and the price of anarchy in atomic splittable congestion games. In: SODA (2011)
19. Rozenfeld, O., Tennenholtz, M.: Strong and correlated strong equilibria in monotone congestion games. In: Spirakis, P., Mavronicolas, M., Kontogiannis, S. (eds.) WINE 2006. LNCS, vol. 4286, pp. 74–86. Springer, Heidelberg (2006)
20. Syrgkanis, V.: Bayesian Games and the Smoothness Framework. ArXiv e-prints (March 2012)
21. Syrgkanis, V., Tardos, E.: Composable and efficient mechanisms. In: STOC (2013)
22. Vetta, A.: Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions (2002)