Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence

Approximate Equilibrium and Incentivizing Social Coordination

Elliot Anshelevich
Rensselaer Polytechnic Institute
Troy, NY-12180, USA
eanshel@cs.rpi.edu.

Shreyas Sekar
Rensselaer Polytechnic Institute
Troy, NY-12180, USA
sekars@rpi.edu.

Abstract

We study techniques to incentivize self-interested agents to form socially desirable solutions in scenarios where they benefit from mutual coordination. Towards this end, we consider coordination games where agents have different intrinsic preferences but they stand to gain if others choose the same strategy as them. For non-trivial versions of our game, stable solutions like Nash Equilibrium may not exist, or may be socially inefficient even when they do exist. This motivates us to focus on designing efficient algorithms to compute (almost) stable solutions like Approximate Equilibrium that can be realized if agents are provided some additional incentives. Our results apply in many settings like adoption of new products, project selection, and group formation, where a central authority can direct agents towards a strategy but agents may defect if they have better alternatives. We show that for any given instance, we can either compute a high quality approximate equilibrium or a near-optimal solution that can be stabilized by providing small payments to some players. Our results imply that a little influence is necessary in order to ensure that selfish players coordinate and form socially efficient solutions.

1 Introduction

Historically, the term coordination game has been applied to social interactions with positive network externalities. Typically, they are used to represent scenarios like the *Battle of the Sexes* wherein self-interested agents benefit if and only if they choose the same strategy. Such a model, however, does not fully capture real-life situations like the adoption of technologies or opinions and the selection of activities where agents may eschew coordination if their personal preference for an alternative is very strong. For instance, a company may not adhere to common standards if the benefit from using their own proprietary technology far outweighs the gains from coordinating. Bearing this in mind, we consider the following broader interpretation of coordination games as 'a class of games where agents' utilities increase when more people choose the same strategy as them'. Notice that this does not preclude agents from having intrinsic preferences for strategies.

Many social and economic interactions fall within our framework (see Galeotti et al. for specific applications of coordination games) and it is not surprising that the kind of games we are interested in have appeared in various guises throughout literature. Researchers have studied similar kinds of games in several settings including opinion formation (Chierichetti, Kleinberg, and Oren 2013), information sharing (Kleinberg and Ligett 2013), coalition formation (Feldman, Lewin-Eytan, and Naor 2012) and party affiliation (Balcan, Blum, and Mansour 2009), largely focusing on the existence and quality of stable solutions.

Given the significance of social coordination, a natural question that arises is: Do instances of such games result in stable outcomes that are comparable to the social optimum, the solution maximizing social welfare. The *somewhat* negative answer to this question that we provide serves as the starting point for our work as it highlights the need for incentivizing agents to form solutions that are beneficial for society. We attempt to answer the above question by articulating two fundamental drawbacks of coordination games, which naturally lead to the issue of influencing players to form "good" solutions.

1. **Coordination Failures.** Coordination games suffer from Coordination Failures (Cooper 1999) that result in agents becoming trapped in inefficient equilibria despite the existence of high-welfare equilibria. These situations arise when agents settle for less risky alternatives if they anticipate that other agents may not coordinate with them on what are potentially "high-risk, high-reward" solutions. As an example, consider N independent but complementary firms, each with a distinct preferred location, deciding on where to locate. Suppose that each company receives unit utility for choosing their favorite location and one more unit for every additional firm that choose to locate in the same area. Clearly an optimum and stable solution is one where all companies choose the same location. However, the outcome where each chooses their preferred location is also stable as no company could unilaterally deviate and profit. There is a large body of theoretical and experimental evidence (see Kosfeld for a survey) that supports the hypothesis that agents may coordinate on inefficient outcomes even when better alternatives exist.

2. **Non-existence of Equilibrium.** In instances where player relationships are asymmetric (they receive different gains

from coordinating), a Pure Nash Equilibrium[1] may not even exist. The following example illustrates one such simple instance with three players and three strategies.

Example: There are 3 players i_1, i_2, i_3. Player i_j receives a utility of $\sqrt{2}$ if she chooses strategy j and unit utility from strategy $j+1$ (addition here is modulo 3). Also, player i_j receives coordination gains of 1 when player i_{j+1} chooses the same strategy as her. Note that relationships are asymmetric so player i_{j+1} receives no benefit for choosing the same strategy as player i_j.
It is not hard to verify that no Nash equilibrium exists for this instance. The reader is asked to refer to the proof of Proposition 2.5 for details.

It is evident that even in fairly simple coordination games, it may be necessary to guide agents to form desirable solutions. From a central point of view, a high social welfare is the most important requirement, but at the same time it is necessary that selfish agents do not deviate from these centrally promoted solutions. A key algorithmic challenge is therefore, computing stable outcomes with good social welfare that can be formed by providing each agent a small incentive. It is towards this end that we identify approximate equilibria as our primary solution concept.

Approximate Equilibrium and Stability. An α-Approximate Equilibrium is an outcome in which no player can improve their utility by a factor more than α by unilaterally deviating. Observe from the definition that if each player is provided additional benefits equaling a fraction $(\alpha - 1)$ times their original utility, then no player would wish to deviate and the Approximate Equilibrium becomes a Nash Equilibrium. Alternatively, approximate stability corresponds to the addition of a switching cost that captures the inertia players may have in changing strategies unless the added benefit is large enough. In addition to being a simple generalization of Nash Equilibrium, approximate equilibria are also easily implementable or enforceable in natural settings as opposed to non-deterministic generalizations like Mixed Nash Equilibria.

We focus on computing approximate equilibria with high social welfare establishing that although coordination games may not admit Nash Equilibrium, the addition of a relatively small amount of inertia to the game causes stable solutions with high social welfare to exist. We also consider group deviations via approximate strong equilibrium (Feldman and Tamir 2009), computing solutions which are resilient to deviations by sets of players.

Formalizing our Model of Coordination Games. We begin by considering a non-transferable utility game with N players and m distinct strategies. We assume that players have access to all strategies, therefore in any outcome of the game, a player's strategy $s_i \in \{1, \dots, m\}$. Generalizing our examples from the previous section, we not only permit preferred strategies, but allow each player to have asymmetric preferences over the strategy set. Formally, player i derives a utility of w_i^k if she chooses strategy $1 \le k \le m$.

With regards to the coordination aspect, we now propose a framework where the benefits of coordination between two players do not depend on externalities. Specifically, suppose that $w(i,j)$ is the total coordination benefit when players i and j choose the same strategy and that this benefit is divided among the two players. Formally, player i derives a utility of $\gamma_{ij} w(i,j)$ for coordinating with j and player j receives $\gamma_{ji} w(i,j)$. For a given instance of our game, the values $(\gamma_{ij}, \gamma_{ji}) \forall (i,j)$ are fixed and add up to one. So given a strategy vector $\boldsymbol{s} = (s_1, \cdots, s_n)$, the utility of a player i has two components,

$$u_i(\boldsymbol{s}) = w_i^{s_i} + \sum_{s_j = s_i} \gamma_{ij} w(i,j).$$

We parameterize any instance on a factor γ that captures the maximum asymmetry that exists in relationships. Formally $\gamma = \max \frac{\gamma_{ij}}{\gamma_{ji}}$ over all pairs (i,j). We term this parameter, the *Maximum Relationship Imbalance* (MRI). Since we are concerned with the quality of stable solutions, we define a social welfare function $u(\boldsymbol{s})$ to be the sum of player utilities. Mathematically, $u(\boldsymbol{s}) = \sum_i u_i(\boldsymbol{s}) = \sum_i w_i^{s_i} + \sum_{s_i = s_j} w(i,j)$. We define the optimum to be the solution maximizing social welfare and compare the quality of our solutions to the optimum welfare OPT.

Our Contributions. In this work, we consider the following well-motivated question: can we implement solutions with high social welfare by providing each player some incentive to not deviate? Our main results answer this question in the affirmative, and more importantly we show that this is possible for every instance using one of our two incentivizing schemes. First, we present an algorithm based on greedy dynamics to compute a good quality, almost-stable solution.

- (Theorem 3.2) There is a polynomial-time algorithm to compute an α-Approximate Equilibrium ($\alpha \in [1.618, 2]$) with a social welfare that is comparable to the optimum.

An approximate equilibrium corresponds to an easily realizable solution in the presence of either incentives, switching costs or players with inertia. Our second main result considers a complementary notion of stability: the minimum total payment to be provided to players so that they do not deviate from a desired high quality solution. For any given instance, if the algorithm of Theorem 3.2 returns an α-Approximate Equilibrium with social welfare $\rho_\alpha \cdot OPT$, then we show

- (Theorem 3.5) The optimum solution can be stabilized with a total payment of $\frac{\rho_\alpha}{\alpha - 1} OPT$.

Informally, this result tells us that if we can provide certain players supplementary utility, then with a finite budget we can stabilize the optimum solution. Given any instance, we first run the algorithm of Theorem 3.2 to compute an α-Approximate Equilibrium with social welfare $\rho_\alpha \cdot$ OPT. If ρ_α is large, then we have an almost-stable solution with high social welfare; else, if ρ_α is small, then with a budget of $\approx \rho_\alpha \cdot$OPT, we can force even the optimum solution to be stable. Together our two theorems imply something much stronger: we can always *either* compute almost-stable solutions with high welfare, *or* directly stabilize a good quality solution with small payments.

[1]We shall henceforth refer to Pure Nash Equilibrium as just Nash Equilibrium

γ	$\alpha = 2$		$\alpha = 1.618$	
	$m = 4$	$m \to \infty$	$m = 4$	$m \to \infty$
1	0.57OPT	0.5OPT	0.424OPT	0.35OPT
2	0.47	0.4	0.37	0.29
10	0.25	0.15	0.18	0.12

Table 1: Performance of our algorithm: The social welfare of our computed solution as a fraction of the optimum welfare for different values of m and γ (MRI).

We obtain tight lower bounds for the social welfare of the solution computed by the algorithm of Theorem 3.2 in terms of γ (Maximum Relationship Imbalance). When γ is not too large, we show that this social welfare is comparable to the optimum. For instance, if relationships are not too asymmetric and $\gamma = 2$, we can compute a 2-Approximate Equilibrium whose social welfare is always at least forty percent of OPT and a 1.618-Approximate Equilibrium whose social welfare is one-third of OPT. If $\gamma < 2$, then we can do much better. Table 1 captures the social welfare of the solution returned by our algorithm for different values of m and γ.

In the process, we also establish that every instance of our Social Coordination Game admits a 1.618-Approximate Equilibrium and that this result is almost tight: $\exists$ instances where no α-Approximate Equilibria exist for any $\alpha < 1.414$. For the special case of $m = 3$ however, we present an algorithm that always returns a 1.414-Approximate Equilibrium. In any coordination game, it is important to also ensure that groups of players do not jointly deviate from the centrally enforced solution. Keeping this in mind, we give an algorithm to compute a 2-Approximate Strong Equilibrium with a social welfare that is at least half of ρ_α for $\alpha = 2$.

Finally, we focus on settings where a central designer may not be able to provide additional incentives to players but can control the parameters of the game. For such cases, we identify a general set of conditions that guarantee the existence of a Nash Equilibrium. These conditions capture a broad sub-class of coordination games where the benefits of coordination among players may be asymmetric but are closely correlated.

Related Work. Our model of social coordination is closely linked to two well known classes of games: non-transferable utility coalition formation and party affiliation. We begin by surveying the substantial literature in both these fields and examine the ties between our model and the games in these frameworks.

Hedonic games model players forming coalitions such that a player's utility depends only on the members of her own coalition. Our model can be embedded in this setting by considering a fixed number of non-anonymous coalitions and a set of players who are anchored, i.e., constrained to join only one particular group. Much of the work in this domain has focused on identifying conditions for the existence of stable solutions (Banerjee, Konishi, and Sönmez 2001; Bogomolnaia and Jackson 2002). It is known, for instance, that if relationships are symmetric then the existence of Nash Equilibrium can be guaranteed by means of a potential function. Augustine et al. (2013) consider a model similar to ours with the coordination benefits being submodular and characterize settings where Nash Equilibrium always exist. Although our games do not admit Nash Equilibrium, our results imply the existence of a stability concept that is slightly weaker than Nash stability for a large class of hedonic games with asymmetry.

Another line of work has focused on quantifying the inefficiency of stable solutions (Brânzei and Larson 2009) and on the computation of stable solutions (Aziz, Brandt, and Seedig 2013; Darmann et al. 2012; Gairing and Savani 2010). Although there have been a number of positive algorithmic results, the focus on approximating both stability and optimality has been limited. With regards to influencing, the recent work on stabilizing desired coalitions via supplementary payments (the Cost of Stability) (Bachrach et al. 2009) is similar to our direct payments technique, albeit in a transferable utility setting.

Party affiliation games are a generalization of pure coordination games where players wish to coordinate with friends and *anti-coordinate* with enemies. On the other hand, in our model the friction is provided by the interplay between a player's individual preference and coordination. Party affiliation games with only two strategies and symmetric relationships have received considerable attention as Nash Equilibrium always exists in these settings although computing it is PLS-Complete (Christodoulou, Mirrokni, and Sidiropoulos 2006). As a positive algorithmic result, Bhalgat et al. (2010) gave a polynomial time algorithm to compute a $(3 + \epsilon)$-Approximate Equilibrium for such games. Even though we look at only one aspect of party affiliation, our model is quite general as we do not impose any restriction on the number of strategies or player relationships.

Other models of coordination in strategic settings. There has been a renewed interest in characterizing the effect of several parameters on the kind of equilibrium outcomes that emerge in coordination games. These include the the cost of forming links (Goyal and Vega-Redondo 2005; Jackson and Watts 2002), network structure (Chwe 2000), level of interaction (Morris 2000) and incomplete information (Galeotti et al. 2010). The literature on coordination games is too vast and the reader is asked to refer to the paper by Galeotti et al. for a detailed survey.

2 Preliminaries and Warm-up Results

In this section, we address some fundamental questions regarding stability and optimality in our Social Coordination Game (SCG) in order to gain a better understanding of our model. We then present algorithms to compute stable solutions for special classes of games where players only have two and three strategies to choose from respectively. We begin by casting our game in graph theoretic framework to compare our problem to existing optimization problems.

Social Coordination as a Network Game. We can view our model as a game played on a complete graph $G = (V, E)$ where the nodes include the players and m additional anchored nodes (which are constrained to choose only one strategy). Each directed edge (i, j) has a weight $\gamma_{ij} w(i, j)$, the utility player i derives from coordinating with j. Now, it is not hard to see that the problem of maximizing social welfare is equivalent to the problem of dividing the graph into m clusters to maximize the weight inside the clusters. The latter problem was referred to in (Langberg, Rabani, and Swamy 2006) as the *Multiway Uncut* Problem and shown to be NP-Hard. The hardness result therefore, extends to our problem of maximizing social welfare.

Proposition 2.1 *For $m > 2$, computing OPT for an instance of the Social Coordination Game is NP-Hard.*

Langberg et al. also exhibited a 0.8535 approximation algorithm for the Multiway Uncut problem on undirected graphs, which we now extend to the directed version via a simple reduction. The reader is asked to refer to the full version of this paper for detailed proofs of all our results (Anshelevich and Sekar 2014).

Proposition 2.2 *There exists a polynomial-time algorithm to compute a solution of the SCG such that its social welfare is at least a fraction 0.8535 of OPT.*

However, in settings with selfish players, it becomes imperative to compute solutions which are not only comparable to OPT but also ensure individual stability. We first consider the most natural stability concept in such games: Nash Equilibrium, and formalize our intuition about *Coordination Failures* from the previous section. That is, we show that even when Nash Equilibrium exist, all stable solutions may have a social welfare that is only a small fraction of OPT.

Proposition 2.3 *If Nash Equilibrium exists, both the Price of Anarchy(PoA) and the Price of Stability(PoS) for the SCG can be as large as m, the number of strategies.*

Recall that the PoA is the ratio of the social welfare of the optimum to that of the worst Nash Equilibrium and the PoS is the same for the best Nash Equilibrium. Proposition 2.3 also tells us that if we are able to limit the number of choices available to agents, then all stable solutions exhibit high social welfare. Of particular interest to us in this regard is the special case when players only have two available options ($m = 2$). For this case, we are able to guarantee the existence of a Nash Equilibrium by actually constructing one for any given instance. We present the following algorithm to compute a Nash Equilibrium when $m = 2$ that is at least half as good as OPT.

> **Algorithm 1:** Pick a strategy and allow all players who want to deviate from this strategy to perform best-response until no player wants to deviate. Now allow all players who want to deviate from the other strategy to perform best-response.

Proposition 2.4 *The above algorithm returns a Nash Equilibrium from any given starting state such that its social welfare is at least half of OPT.*

Finally, before moving on to our main results, we also briefly consider the case of three strategies, i.e., $m = 3$. We have already established the non-existence of Nash Equilibrium for this case. Instead, we give an algorithm to compute a 1.414-Approximate Equilibrium for every single instance with three strategies. This factor of 1.414 is tight: there exist instances where for any $\alpha < \sqrt{2}$, no α-Approximate Equilibrium exists.

Proposition 2.5 *The following algorithm returns a $\sqrt{2}$-Approximate Equilibrium for all instances with three strategies.*

> "Run Algorithm 1 for any two strategies (say 1 and 2). Now allow any player whose utility improves by at least a $\sqrt{2}$ factor to deviate to 3. Every time a player deviates to strategy 3, run Algorithm 1 to ensure that players in 1 and 2 are stable w.r.t each other."

3 General Social Coordination Games

In the previous sections, we showed that Nash Equilibrium may not exist for Social Coordination Games and even when it does, its quality may not be close to the optimum. This motivates us to relax our notion of stability and consider approximately stable solutions in order to obtain better guarantees on the social welfare. Our first main result in this section is a quadratic algorithm to compute an α-Approximate Equilibrium for $\alpha \in [\phi, 2]$ ($\phi \approx 1.618$) and a 2-Approximate Strong Equilibrium with good social welfare. As we showed that α-Approximate Equilibria for $\alpha < 1.414$ do not exist for general SCGs, our stability results are nearly tight.Recall that in an α-Approximate Equilibrium, no player can deviate from one strategy to another and increase her utility by more than α times her previous utility.

We obtain tight bounds on the social welfare of our solution. Using this, we characterize the near-linear trade-off between stability and optimality by showing how the social welfare decreases when we decrease the stability factor from $\alpha = 2$ to $\alpha = \phi$ (See Figure 1(a)). Our bounds for social welfare depend on γ, the Maximum Relationship Imbalance (MRI), that measures the maximum asymmetry over all relationships. We explicitly use this factor γ in our welfare as often real-life relationships are bounded in their asymmetry; if one player receives large gains from coordinating with another, then the second player derives at least some fraction of that benefit. Observe that $\gamma = 1$ denotes a special case of interest wherein the rewards of all relationships are shared equally and as γ increases, reward sharing becomes more and more asymmetric.

We begin by presenting a simple algorithm called *one-shot α-BR*, which allows each player to play her best-response at most once as long as the improvement in utility is by a factor $\alpha \geq 1$. We then make use of this algorithm and propose a hybrid procedure that takes the best solution returned by two separate instances of the algorithm, in order to compute a solution which is both approximately stable and is a good approximation to OPT. First, we show the following lemma.

Algorithm 1 One-shot α-BR

Require: A starting strategy k_0.

1: Begin with all players in the initial strategy k_0.
2: While there exists a player whose current strategy is k_0, who can improve her utility by at least a factor α by deviating to another strategy, allow her to perform best-response.

Lemma 3.1 *The* One-shot α-BR *algorithm returns a* $\left(\max\left(\alpha, \frac{1}{\alpha}+1\right)\right)$*-Approximate Equilibrium, for any starting strategy* k_0.

That is, for $\alpha \in [\phi, 2]$, the *One-Shot α-BR* algorithm returns an α-Approximate Equilibrium. We are now in a position to show our main result. For the purpose of the following algorithm, we use k^* to denote the strategy that players have the maximum preference for, $k^* = \arg\max_k(\sum_i w_i^k)$. We also use A_T to denote the total utility that all players derive from their preferred strategy ($A_T = \sum_i \max_k(w_i^k)$) and P_T the maximum coordination benefit ($P_T = \sum_{(i,j)} w(i,j)$).

Theorem 3.2 *The following algorithm returns an α-Approximate Equilibrium for $\alpha \in [\phi, 2]$ whose social welfare is approximately at least a fraction $\approx \max(\frac{\alpha}{\gamma+3}, \frac{1}{m})$ of the optimum social welfare.*

> **Algorithm:** "For a given α, run *One-shot α-BR* and *One-shot $\frac{1}{\alpha-1}$-BR* with k^* as the starting strategy. Let the returned solutions be $\boldsymbol{s_1}$ and $\boldsymbol{s_2}$. Return the solution among these two with greater social welfare."

(Proof Sketch) By Lemma 3.1, we are guaranteed that both $\boldsymbol{s_1}$ and $\boldsymbol{s_2}$ are α-Approximate equilibria as $\frac{1}{\frac{1}{\alpha-1}} + 1 = \alpha$. Moving on to social welfare, since each player is allowed to perform best-response exactly once, the total utility of $\boldsymbol{s_1}$ is at least the total utility players derive from their preferred strategy (discounted by α), that is $\frac{A_T}{\alpha}$. Similarly, the algorithm results in a solution whose social welfare is not too far from that of the starting state. It is also not hard to see that the starting state (all players choosing k^*) has a welfare of at least $\frac{A_T}{m} + P_T$. We now have two lower bounds on the social welfare of our solution. Using the fact that OPT is no more than $A_T + P_T$, and finding the point where the maximum of the two lower bounds is minimized, we get our final welfare. ■

Discussion. We attempt to break-down the dependence of social welfare on different parameters. First, notice from the approximate formula $\frac{\alpha}{\gamma+3}$ that the social welfare increases almost linearly with α. As is usually the case, sacrificing a little individual stability results in greater overall well being. At the same time, we see that social welfare decreases when γ becomes larger, i.e., when relationships become more asymmetric. Therefore, for a designer, there are two possible measures to ensure socially efficient outcomes: **(i)** imposing a higher switching cost on agents, **(ii)** splitting the rewards of coordination almost equally.

Our algorithm provides good guarantees when γ is not too large. For instance, when the benefits from coordination are off by at most a factor of 2 ($\gamma \leq 2$), our algorithm returns a 1.618-Approximate Equilibrium that is almost one third of OPT even if the number of available alternatives is arbitrarily large. Contrast this with the result in Proposition 2.3 that states as $m \to \infty$, the social welfare of equilibria becomes infinitely worse off than the optimum. Finally, the social welfare is bounded by $\approx \frac{1}{m}$, which indicates that even when γ is large, our solutions are at least as good as that of the Nash Equilibrium (when they exist). Figure 1(b) illustrates how the social welfare drops when relationships become more asymmetric. We also remark here that our algorithm provides particularly good guarantees on social welfare when the setting is symmetric ($\gamma = 1$).

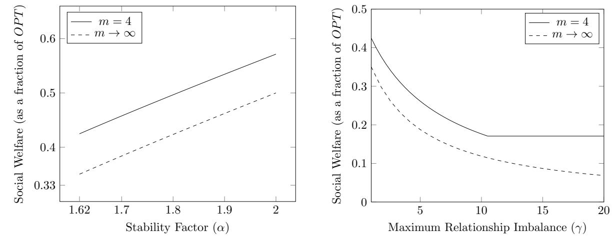

(a) Symmetric SCG($\gamma = 1$ case): Trade-off between welfare and stability.

(b) Asymmetric SCG: How the asymmetry of the graph affects welfare when $\alpha = 1.618$.

Figure 1: Social Welfare as a function of α and γ.

Corollary 3.3 *For the special case when coordination benefits are split equally ($\gamma = 1$), our algorithm returns a* 2*-Approximate Equilibrium that is more half than as good as the optimum and a* 1.618*-Approximate Equilibrium that is at least a fraction* 0.35 *of OPT.*

Strong Equilibrium. In any game involving coordination, it becomes important to consider Strong Nash stability as coalitions of players may be able to cooperate when switching to different strategies. For our SCG, we focus on computing solutions where for any group of players who can deviate and improve their utility, at least one player's utility increases by a factor no greater than α. We term this α-Approximate Strong Equilibrium. Our main computational result here is that running the *One-shot α-BR* algorithm returns a $(\alpha + 1)$-Approximate Strong Equilibrium.

Proposition 3.4 *For $\alpha \geq 1$, the* One-shot α-BR *returns a* $(1+\alpha)$*-Approximate Strong Equilibrium.*

We remark here that for $\alpha = 1$, we get a 2-Approximate Strong Equilibrium whose social welfare is at least half of that mentioned in Theorem 3.2. Therefore, our *One-shot α-BR* algorithm also offers resilience against group deviations.

Direct Payments as Incentives

So far, we have looked at approximate equilibria for our game, which become fully stable when there are some incentives or when there is limited inertia in switching. We now consider a more explicit incentivizing technique where a central authority can provide arbitrary payments to any subset of players in order to enforce a desired solution. Ideally, this desired solution is the social optimum although we later consider other high welfare solutions which are easily computable. In such situations, there are often budgets constraints which determine the total payments that can be made and we wish to obtain bounds on the minimum budget required for any given instance.

We now define our incentive scheme: we are interested in 'stabilizing' a given solution $\boldsymbol{s}$ by providing each player i a payment of $\nu_i OPT$, where $0 \leq \nu_i \leq 1$, such that her total utility is now $u_i(\boldsymbol{s}) + \nu_i \cdot OPT$. More precisely, players are provided this additional utility if and only if they stick to their prescribed strategy under $\boldsymbol{s}$. A solution is said to be successfully stabilized if after the additional payments, $\forall i$, no player wishes to deviate from strategy s_i. Our goal is to bound ν, the total payment that is required as a fraction of the optimum welfare, i.e., $\nu \cdot OPT = \sum_i \nu_i \cdot OPT$. Now suppose for that for a given instance of the coordination game, the α-Approximate Equilibrium returned by our Algorithm of Theorem 3.2 has a social welfare of ρ_α·OPT for $\rho_\alpha \leq 1$. Our second main result is the following.

Theorem 3.5 *For any given instance, the optimum solution can be stabilized by providing direct payments to players such that the total payment is at most a fraction* $\nu \leq \frac{\rho_\alpha}{(\alpha-1)}$ *times the social welfare of OPT.*

The above result indicates that if we run our algorithm from Theorem 3.2 and get a solution with low social welfare (ρ_α is small), then we can stabilize the optimum solution with a total payment that is approximately equal to ρ_α·OPT. We previously observed that when both γ and m are large, the solution returned by our algorithm may not be too efficient. In such cases, we can do much better by implementing the optimal solution and providing small incentives to players.

(Proof Sketch) We give intuition for why the fractional incentive ν is no more than $2\frac{\rho_\alpha}{\alpha-1}$ here. Some additional optimization gives us a better bound. Let s^* be the optimum solution and k_i denote player i's best-response. Then, the minimum incentive that we must provide player i to not deviate is $u_i(k_i, s^*_{-i}) - u_i(s^*)$. The above term has two components, added utility due to intrinsic preference($A(X)$), and that due to coordination($P(X)$). We claim that $A(s^*) \geq P(X)$, where $A(s^*)$ is the intrinsic utility in OPT. If this were not true then the solution where all players choose k^* would have a utility greater than $P(X) + P(s^*) > A(s^*) + P(s^*)$, which is a contradiction. We also observe that both $(\alpha - 1)A(X)$ and $(\alpha - 1)A(s^*)$ are less than $\rho_\alpha \cdot OPT$, which means $A(X) + P(X) \leq A(X) + A(s^*) \leq 2\frac{\rho_\alpha}{\alpha-1}$. ■

The bottom line for a central enforcer is that a little incentivizing goes a long way in ensuring socially efficient outcomes. For every given instance, we can either apply the algorithm of Theorem 3.2 to compute a good quality approximate equilibrium or enforce even the optimum solution with a small total budget, thereby always ensuring high welfare, stable outcomes. Unfortunately, computing the optimum is NP-hard. Therefore, we consider the 0.8535 approximation to OPT from Proposition 2.2 and show that it can be stabilized with the same total payments.

Corollary 3.6 *There exists a solution computable in polynomial time whose social welfare is at least* 0.8535 *times OPT, which can be stabilized by providing total payments* ν *no greater than* $\frac{\rho_\alpha}{\alpha-1}$ *times OPT.*

Incentivizing without payments. Both the influencing techniques that we have looked at at so far involve providing additional utility to players in the form of payments or other equivalent incentives. For completeness, we also consider settings where the central designer may not be able to provide any direct incentives but can exert some control over the parameters of the game. If the players cannot be guided towards good solutions, then it is important that natural game play (best-response dynamics) results in stable solutions. Identifying special classes of games where best-response converges to a Nash Equilibrium is therefore a problem of considerable interest. Towards this end, we illustrate a condition that guarantees that the benefits of coordination received by a single player are correlated across relationships and then show that this is a sufficient condition for the existence of a potential function.

(Correlated Coordination Condition) A given instance of the SCG is said to satisfy this condition if $\exists$ a vector $\boldsymbol{\gamma} = (\gamma_1, \cdots, \gamma_N)$ such that $\forall (i,j)$ with $w(i,j) > 0$, we have $\gamma_{ij} = \frac{\gamma_i}{\gamma_i+\gamma_j}$ and $\gamma_{ji} = \frac{\gamma_j}{\gamma_i+\gamma_j}$.

The γ_i associated with each player defines the level of influence that she commands over her relationships. For any two players i and j, if $\gamma_i \geq \gamma_j$, then player i receives a greater benefit due to coordinating with player j. Notice that this definition does not impose any restriction on how asymmetric relationships can be. We now claim that games which obey this condition admit an ordinal (inexact) potential function (Monderer and Shapley 1996) that ensures that best-response always converges to an equilibrium.

Theorem 3.7 *Social Coordination Games which obey the Correlated Coordination Condition admit a potential function. Therefore, best-response dynamics always converge to a Nash Equilibrium in such games.*

(Proof Sketch) The following is an ordinal potential for any SCG that obeys the Correlated Coordination Condition.

$$\Phi(\boldsymbol{s}) = \sum_i \frac{w_i^{s_i}}{\gamma_i} + \sum_{s_i = s_j} \frac{w(i,j)}{\gamma_i + \gamma_j}.$$

It is not hard to verify that every better-response move results in an increase in this function and vice-versa.■

4 Conclusions and Future Work

In this work, we take a first step towards influencing players in games where they benefit from mutual coordination. At a high level, we showed that with limited incentives, selfish players can always be made to form good solutions, i.e., we can either compute an almost-stable solution with high welfare or provide small payments to implement the optimal solution. From an algorithmic perspective, we present an efficient algorithm to compute a 1.618-Approximate Equilibrium. A natural direction forward would be to model more complex coordinations between players with the benefit being supermodular in the set of players coordinating. Some preliminary results on the existence and computation of stable solutions for SCGs with supermodularities are presented in the full version (Anshelevich and Sekar 2014).

5 Acknowledgements

This work was supported in part by NSF awards CCF-0914782, CCF-1101495, CNS-1017932, and CNS-1218374.

References

Anshelevich, E., and Sekar, S. 2014. Approximate equilibrium and incentivizing social coordination. *arXiv preprint arXiv:1404.4718*.

Augustine, J.; Chen, N.; Elkind, E.; Fanelli, A.; Gravin, N.; and Shiryaev, D. 2013. Dynamics of profit-sharing games. *Internet Mathematics* (just-accepted).

Aziz, H.; Brandt, F.; and Seedig, H. G. 2013. Computing desirable partitions in additively separable hedonic games. *Artificial Intelligence* 195:316–334.

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

Balcan, M.-F.; Blum, A.; and Mansour, Y. 2009. Improved equilibria via public service advertising. In *Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 728–737.

Banerjee, S.; Konishi, H.; and Sönmez, T. 2001. Core in a simple coalition formation game. *Social Choice and Welfare* 18(1):135–153.

Bhalgat, A.; Chakraborty, T.; and Khanna, S. 2010. Approximating pure nash equilibrium in cut, party affiliation, and satisfiability games. In *Proceedings of the 11th ACM conference on Electronic Commerce (EC)*, 73–82. ACM.

Bogomolnaia, A., and Jackson, M. O. 2002. The stability of hedonic coalition structures. *Games and Economic Behavior* 38(2):201–230.

Brânzei, S., and Larson, K. 2009. Coalitional affinity games and the stability gap. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, 79–84.

Chierichetti, F.; Kleinberg, J.; and Oren, S. 2013. On discrete preferences and coordination. In *Proceedings of the fourteenth ACM conference on Electronic Commerce (EC)*, 233–250. ACM.

Christodoulou, G.; Mirrokni, V. S.; and Sidiropoulos, A. 2006. Convergence and approximation in potential games. In *Proceedings of the 18th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, 349–360. Springer.

Chwe, M. S.-Y. 2000. Communication and coordination in social networks. *The Review of Economic Studies* 67(1):1–16.

Cooper, R. 1999. *Coordination games*. Cambridge University Press.

Darmann, A.; Elkind, E.; Kurz, S.; Lang, J.; Schauer, J.; and Woeginger, G. 2012. Group activity selection problem. In *Proceedings of the 8th Intl. Workshop Internet & Network Economics (WINE)*. Springer. 156–169.

Feldman, M., and Tamir, T. 2009. Approximate strong equilibrium in job scheduling games. *Journal of Artificial Intelligence Research* 36(1):387–414.

Feldman, M.; Lewin-Eytan, L.; and Naor, J. S. 2012. Hedonic clustering games. In *Proceedinbgs of the 24th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 267–276. ACM.

Gairing, M., and Savani, R. 2010. Computing stable outcomes in hedonic games. In *Proceedings of the 3rd International Symposium on Algorithmic Game Theory (SAGT)*. Springer. 174–185.

Galeotti, A.; Goyal, S.; Jackson, M. O.; Vega-Redondo, F.; and Yariv, L. 2010. Network games. *The Review of Economic Studies* 77(1):218–244.

Goyal, S., and Vega-Redondo, F. 2005. Network formation and social coordination. *Games and Economic Behavior* 50(2):178–207.

Jackson, M. O., and Watts, A. 2002. On the formation of interaction networks in social coordination games. *Games and Economic Behavior* 41(2):265–291.

Kleinberg, J., and Ligett, K. 2013. Information-sharing in social networks. *Games and Economic Behavior* 82:702–716.

Kosfeld, M. 2004. Economic networks in the laboratory: A survey. *Review of Network Economics* 3(1).

Langberg, M.; Rabani, Y.; and Swamy, C. 2006. Approximation algorithms for graph homomorphism problems. In *Proceedings of the 9th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APRPOX)*, 176–187. Springer.

Monderer, D., and Shapley, L. S. 1996. Potential games. *Games and economic behavior* 14(1):124–143.

Morris, S. 2000. Contagion. *The Review of Economic Studies* 67(1):57–78.