

Incentivizing Collaboration in a Competition

Arunesh Sinha
University of Michigan
Ann Arbor, Michigan
arunesh@umich.edu

Michael P. Wellman
University of Michigan
Ann Arbor, Michigan
wellman@umich.edu

ABSTRACT

Research and design competitions aim to promote innovation or creative production, which are often best achieved through collaboration. The nature of a competition, however, typically necessitates sorting by individual performance. This presents tradeoffs for the competition designer, between incentivizing global performance and distinguishing individual capability. We model this situation in terms of an abstract collaboration game, where individual effort also benefits neighboring agents. We propose a scoring mechanism called LSWM that rewards agents based on localized social welfare. We show that LSWM promotes global performance, in that social optima are equilibria of the mechanism. Moreover, we establish conditions under which the mechanism leads to increased collaboration, and under which it ensures a formally defined distinguishability property. Through experiments, we evaluate the degree of distinguishability achieved whether or not the theoretical conditions identified hold.

KEYWORDS

collaboration incentives; mechanism design; scoring competitions

ACM Reference Format:

Arunesh Sinha and Michael P. Wellman. 2019. Incentivizing Collaboration in a Competition. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019*, IFAAMAS, 9 pages.

1 INTRODUCTION

Research and design competitions [4, 15, 17, 29] have an underlying goal of generating novel solutions or advancing the state-of-the-art in the concerned domain. Collaboration among the competitors can often help in achieving these goals. For example, the best solution in the 2009 Netflix competition was an ensemble of prediction techniques produced by many teams, and the teams negotiated a prize-sharing agreement [28]. Offline negotiation is always an option, but may be difficult especially when competitors are unknown to each other or interact in an ad hoc manner. This raises a natural question: How can beneficial collaboration be internalized in the competition itself, without relying on participant negotiation initiative and ability?

Let us posit that each participant in the competition has some inherent capabilities, and outcomes of the joint effort are described in terms of individual performance for these agents. The designers of the competition have a social objective to optimize the combined performance of the agents. Maximizing the social objective may

require collaboration, but effort devoted to collaboration comes at a cost to the individual agents. The competition takes the form of a contest or championship, where participants are assigned scores and these scores are ultimately used to declare winners and allocate prizes. The task of the competition designer is to find a way to *score the competition*, such that the participants are incentivized to exert appropriate collaboration effort, within the general structure of a competition mechanism. Maximizing global welfare is a standard objective of mechanism design. In our context, this involves promoting collaboration, which may also be considered desirable for its own sake. Another especially important objective for competition events is to differentiate agents based on their capabilities and contributions. That is, the designer cares that the most capable performers and contributors actually win the contest or get the lion's share of prizes.

In pursuing its objectives, the competition designer is limited to specifying a scoring scheme for participants. Unlike more generic mechanism design settings, we assume the designer takes the action space and outcomes as given, and has no ability to compute counterfactual outcomes, that is, what would have happened had the agents behaved differently. Whereas classical mechanism design focuses on incentives for revelation of private information [23], the spirit of a competition is to spur innovation in agent behavior. Thus, it would not make sense to assume the designer could compute optimal behavior itself.

We model the competition domain in terms of an *abstract collaboration game*, where individual utilities directly capture the performance of the player. The game has a graphical structure, such that players can affect performance only of themselves and their neighbors. The key decision for players in our abstraction is how much to collaborate with their neighbors. Collaborative effort improves performance, but at a cost. Agent capability is captured by player-dependent fixed parameters of their utility function, which combines performance and collaboration cost. The competition designer aims to achieve its objectives by defining a scoring function that depends only on the observed performance outcome of the agent behavior and the graph of collaboration.

We claim three main contributions. Our first contribution is the *local social welfare maximizing* (LSWM) scoring mechanism, in which every player receives a score equal to its own individual performance plus the sum over its neighbors'. We prove that maximizing social welfare is a Nash Equilibrium (NE) of this mechanism. Second, we identify conditions under which an increase in player capability (Definition 4.5) leads to an increase in collaboration and in scores. Our third contribution is a definition of *distinguishability* that captures whether a scoring scheme assigns higher scores to players with greater capability, all else equal. We prove that the LSWM mechanism provides distinguishability under *modularity* assumptions on the utility function.

Finally, we conduct simulations to demonstrate that best-response dynamics converges quickly to social optimum in the LSWM mechanism. We also investigate distinguishability over multiple simulated instances of graphical games, showing that distinguishability is prevalent even when the utility function used violates the conditions ensuring the property. All full proofs omitted in the main text are provided in an online appendix.¹

2 A MOTIVATING EXAMPLE DOMAIN

DARPA is a US government research funding agency that regularly employs competitions to spur innovation in high-priority technology areas. For example, their current Spectrum Collaboration Challenge (SC2) offers a total of US\$3.75 million in prize money for the final competition to be held in 2019. The premise of SC2 is that fluid collaboration among intelligent radios can dramatically improve efficiency in use of limited spectrum (frequency range), compared to static spectrum allocations. The challenge is posed as a competition among teams of radio networks moving about in an arena, attempting to accomplish communication goals despite signal interference. The aim of the competition is to spur novel ideas, and at the same time, to demonstrate the benefits of collaboration in avoiding interference, as measured by global communication performance. Inspired by this competition, we address the general abstract problem of scoring a competition in order to identify teams with better capabilities while also encouraging collaboration in service of global performance.

3 MODEL

3.1 Preliminaries: Notation and Definitions

We indicate vector variables by an arrow on top, for instance \vec{x} . The i^{th} element of \vec{x} is x_i and $\vec{x} \geq \vec{y}$ denotes element-wise comparison: $x_i \geq y_i$ for all i . In particular, $\vec{x} \geq 0$ means all elements of \vec{x} are nonnegative. To focus on one element x_i of a vector \vec{x} we write $\vec{x} = (x_i, \vec{x}_{-i})$. Similarly, to focus on a subset $S \subset \{1, \dots, d\}$ of elements we write $\vec{x} = (\vec{x}_S, \vec{x}_{-S})$, where $-S = \{1, \dots, d\} \setminus S$ denotes the complement of S .

3.2 Matches and Measurement Intervals

The unit of play in the competition is a *match* M played over a fixed time interval by a set of players $m \subseteq \{1, \dots, n\}$ with $|m| \leq n$. The players receive a *score* S_i at the end of the match. The match duration is further partitioned into a set T of *measurement intervals* (MIs). Player i receives interval score $S_i(\tau)$ for each MI $\tau \in T$. The overall *match score* is additive over MIs: $S_i(M) = \sum_{\tau \in T} S_i(\tau)$. Next, we model the interaction within each MI as a collaboration game.

3.3 Collaboration Game Structure

Since scoring is additive over MIs, we focus on a single MI, returning to match-level analysis in Section 5. The interaction in each MI is represented by a *collaboration game* G among the m players. We model G as a *graphical game* [16], that is, structured according to a collaboration graph where the nodes denote players and edges connect players whose actions affect each other's utility. The collaboration graph is chosen by the competition designer. A match is a

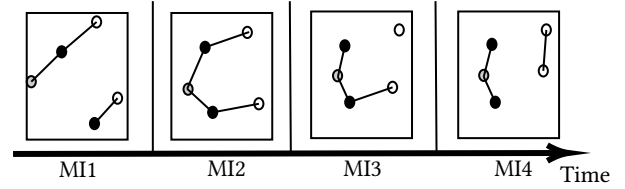


Figure 1: A collaboration graph evolving over a series of measurement intervals. The graph arises from an interaction scenario defined by the designer.

sequence of collaboration games with differing underlying graphs. The sequence of collaboration graphs in a match is generated randomly by the competition designer and the generation process is publicly known. The players know the realization of the sequence of graphs before the start of a match. Figure 1 shows how the collaboration graph can evolve over the course of a match. For example, in the intelligent radio domain, as radios move in prescribed paths, the set of radios they may interfere with naturally varies. Thus, the collaboration graph changes over MIs; the MI is short enough to assume a single collaboration graph in each MI. We elaborate on matches further in Section 5; here we focus on a MI.

A player's action in G is a choice of how much to collaborate. Let $N(i)$ denote the set of neighboring players of i , not including i . Action $a_{i,j} \in [0, 1]$, $j \in N(i)$, denotes the collaboration effort of player i towards player j , higher values corresponding to greater degrees of collaboration. We use \vec{a}_i to denote i 's action vector. The actions of other players that affect i are $a_{j,i}$, $j \in N(i)$. We use $\vec{a}_{N(i)}$ to denote this neighbor action vector. The node occupied by player i is denoted by $v(i)$.

The effect of a player i 's own action and that of its neighbors on its individual performance is represented by a function $E_{v(i)} : [0, 1]^{|N(i)|} \times [0, 1]^{|N(i)|} \rightarrow [0, 1]$. The cost of collaboration is captured by a function $C_{v(i)} : [0, 1]^{|N(i)|} \rightarrow [0, 1]$. The *individual utility* u_i for player i is a combination of collaborative performance and cost of collaboration:

$$u_i(\vec{a}_i, \vec{a}_{N(i)}) = P_i \times E_{v(i)}(\vec{a}_i, \vec{a}_{N(i)}) - c_i \times C_{v(i)}(\vec{a}_i), \quad (1)$$

where P_i and c_i are positive player-specific coefficients that weigh the two components. P_i and c_i can be viewed as representing the inherent *capability* of player i , amplifying or attenuating the efficacy and cost of its collaboration effort. Players can choose their P_i , c_i prior to the start of the competition, but these parameters are fixed once the competition has started. Note that the functions E and C are specific to the position in the graph, and not specific to the players. In particular, when focusing on a single collaboration game with a fixed graph (as is the case in a MI) we drop the subscript for E and C as they are clear from context. Observe that G is a continuous game, with topological action space and utility continuous in actions.

The game G in a MI abstracts away from domain-specific competition details. Individual utility u_i in the collaboration game captures the performance of each player in a MI. For example, in the intelligent radio domain performance might be measured by bytes of data successfully delivered. Similarly, E abstractly captures how players interact. Thus, in the radio domain the static interaction function E captures the effect of interference and C captures the

¹Available on the authors' webpage.

cost of avoiding interference. The players make a number of other radio design choices before the start of the competition, which are captured by P_i, c_i . Hence, the capabilities P_i, c_i are not part of the strategic interaction with other players but are influenced by choice of error-correcting code, power, bandwidth management, etc. It is natural to assume monotonicity in E and C , that is, a positive effect of collaboration on outcomes for a player and its neighbors (through E), but also an increase in cost (through C). We specify this formally in Section 4.3.

3.4 MI Score Design Problem

If rational agents play the collaboration game G in a MI directly, they will be unlikely to collaborate sufficiently, as individual utility (Eq. 1) imposes a cost for collaboration, with much of the benefit going to neighbors. A competition designer aims to correct for this externality through the scoring mechanism. At the same time, a key design goal is to distinguish relative capabilities of the participants in any match. In doing so, the competition designer is limited to only specifying a score function. Specifically, this scoring mechanism design problem is constrained as follows:

- Score designer takes the action and outcome spaces as given.
- Score designer has no knowledge of player capabilities P_i, c_i or the functions E, C . For example, in the intelligent radio domain the P_i, c_i values depend on the secret radio design choices of each player.
- Score designer only observes the final result of the players' interaction, which is the performance of each player as measured by individual utility u_i . For example, in the intelligent radio domain the designer may observe how much traffic was successfully delivered by each player.

As a result, the designer has no ability to mandate a direct mechanism or compute counterfactual outcomes—required for classical mechanisms like VCG [22].² Rather, the target mechanism needs to operate by assigning scores that are a function of the observed outcomes, as measured by individual utilities. Observe that the designer assigns scores, which has no monetary value for the designer, thus, positive payments by the designer is not a concern here.

4 LOCAL SOCIAL WELFARE MAXIMIZATION

The *local social welfare maximizing* (LSWM) mechanism scores player i by the sum of its individual performance (i.e., utility in G) and those of its neighbors: $S_i = u_i + \sum_{j \in N(i)} u_j$. The mechanism induces a new game, where players are assigned scores S_i rather than self performance u_i . We show below that LSWM optimizes social welfare for *any* graph in game G in a MI. Moreover, we analyze the mechanism's ability to distinguish between players of different capabilities.

4.1 Social Welfare Maximization

Our first result is that there are NE action profiles that maximize social welfare for the LSWM induced game.

THEOREM 4.1. *Games induced by the LSWM mechanism have NE that maximize social welfare.*

PROOF. Observe that the game with LSWM is an *exact potential game* [20] with the global social welfare $P(\vec{a}) = \sum_i u_i(\vec{a}_i, \vec{a}_{N(i)})$ as the potential. This can be seen as any change in the action \vec{a}_i of player i changes the score $S_i = u_i + \sum_{j \in N(i)} u_j$ of player i . As the utility u_k of any non-neighbor of i is unaffected by \vec{a}_i , the change in P is also exactly in the terms $u_i + \sum_{j \in N(i)} u_j$. Thus, S_i and P change by the same amount. Then, the proof directly follows from properties of infinite continuous potential games [13] with compact action spaces, mainly that maximum-potential pure strategies are guaranteed to exist and constitute a NE. \square

Note that the above result applies to any graphical game and not just our collaboration-game model. As stated in the proof, LSWM induces a potential game, which has several further advantages. In a potential game, a NE is provably reached by best-response dynamics.³ Further, the NE is unique if the utility functions are differentiable [21]. Moreover, best-response dynamics requires only knowledge of other players' actions and not their utilities. This means that as players learn about the environment and adapt by playing a best response in the LSWM-induced game, they are guaranteed to converge to the NE.

4.2 Supermodularity: Background and Results

We provide background on supermodularity and some results that are required for the analysis in subsequent sections.

Definition 4.2 ([26]). Let $\vec{x} \vee \vec{x}'$ ($\vec{x} \wedge \vec{x}'$) denote the vector with a component-wise maximum (minimum) of \vec{x} and \vec{x}' . $f : X \rightarrow \mathbb{R}$ is *supermodular* iff

$$f(\vec{x}) + f(\vec{x}') \leq f(\vec{x} \vee \vec{x}') + f(\vec{x} \wedge \vec{x}')$$

whenever $\vec{x} \vee \vec{x}' \in X$ and $\vec{x} \wedge \vec{x}' \in X$.

For our proofs, it is convenient to employ a related condition on *left directional derivatives*, which is defined as $d_{\vec{v}}^- f(x) = \lim_{t \rightarrow 0^-} \frac{f(\vec{x} + t\vec{v}) - f(\vec{x})}{t}$ in the direction \vec{v} . $d_{\vec{v}}^- f$ is known to exist for all convex and concave functions with open convex domain [25]. Note that $d_{\vec{v}}^- f$ is invariant in the magnitude of \vec{v} , that is, $d_{\vec{v}}^- f(x) = d_{\alpha\vec{v}}^- f(x)$ for any $\alpha > 0$. Thus, as a convention we take $\|\vec{v}\| = 1$. For differentiable functions it is known that supermodularity is equivalent to $\frac{\partial f(\vec{x})}{\partial x_i}$ increasing with any other dimension $j \neq i$ for all i . We prove the following result, which can be seen as a generalization of the above for non-differentiable supermodular functions.

THEOREM 4.3. *Let $f : X \rightarrow \mathbb{R}$ be any continuous function for which the left directional derivative exists for all $x \in X$ ($X \subset \mathbb{R}^d$ and open convex) and for all directions \vec{v} . Then the following are equivalent*

- (1) f is supermodular.

²Counterfactual inference is required specifically by the Clarke pivot in VCG. VCG with a zero payment in our context assigns the social welfare as score to every player, which as we argue is unacceptable for a competition due to lack of distinguishability.

³For infinite games the convergence is in the limit when the potential is continuous and strictly concave [13, Thm. 2.2.1], which follows from assumptions introduced below.

- (2) (*monotone left positive directional derivative*) For any \bar{x}, \bar{x}' consider the non-empty subset $S \in \{1, \dots, d\}$ such that $x'_i = x_i$ for $i \in S$ and $x'_i \geq x_i$ otherwise, and any direction \bar{v} such that $v_i \geq 0$ for $i \in S$ and 0 otherwise (with strict inequality for some $j \in S$ since $\|\bar{v}\| = 1$), then it holds that $d_{\bar{v}}^- f(\bar{x}') \geq d_{\bar{v}}^- f(\bar{x})$.

We call the second condition above MLPDD for short. Next, using the above equivalence, we prove a lemma stating that if at any point \bar{x}_0 the left directional derivatives are positive in all positive directions for a strictly concave function, then the function maximizer is more than \bar{x}_0 . The result additionally handles subtle cases of \bar{x}_0 being on the boundary of the feasible region or some dimensions being fixed.

LEMMA 4.4. Consider a continuous strictly concave function $f : Y \rightarrow \mathbb{R}$ with Y open convex and $X = [0, 1]^n \subset Y$. Let $T \subset \{1, \dots, n\}$ be dimensions such that values in T can vary and the values of dimensions in $-T$ are fixed. Let $\bar{x}^0 \in X$ be any point. Further, let $S \subset T$ be a set such that $\bar{x}_{T \setminus S}^0 = 0$ and $\bar{x}_S^0 > 0$. Then $d_{\bar{v}}^- f(\bar{x}^0) \geq 0$ for all \bar{v} with $\bar{v}_S \geq 0$ and $\bar{v}_{-S} = 0$ (strict inequality for at least some $i \in S$ since $\|\bar{v}\| = 1$) and f a supermodular function (or equivalently satisfying the MLPDD property), implies that the maximizer $\bar{x}^* \in X$ of $f(x)$ satisfies $\bar{x}^* \geq \bar{x}^0$.

4.3 Monotone Collaboration

That the LSWM mechanism produces optimal collaboration in equilibrium is the primary sense in which it meets the goal of incentivizing collaboration. In this section we consider conditions that tend to promote mutual collaboration. For this analysis, we require the following technical conditions on the components of individual utility (1):

- (monotone concave E) E is non-decreasing in any argument and concave. That is, increasing collaboration leads to higher E (though could lower overall performance due to C), but with diminishing returns.
- (strictly monotone and strictly convex C) C is a strictly increasing and strictly convex function. Collaboration is costly, and the marginal cost grows with effort level.

Next, we restrict E and $-C$ to be supermodular. As a positively weighted sum of supermodular functions is still supermodular, the score of every player is supermodular and so is the potential P . Under these conditions, we wish to show that improving an individual player's capability leads to across-the-board increases in collaboration. First, we define what it means to improve capability.

Definition 4.5. Suppose that player i 's utility parameters P_i and c_i are replaced by $P_{i'} = (1 + \epsilon)P_i$ and $c_{i'} = (1 - \delta)c_i$, respectively, for $\epsilon, \delta > 0$. We then say that player i improves by (ϵ, δ) , and write $i' >_{\epsilon, \delta} i$.

The above notation naturally extends to comparing capabilities of distinct players i and j , which we use in Section 5. We consider what happens when player i improves by (ϵ, δ) .

THEOREM 4.6. In game G suppose the social-welfare-maximizing NE score is attained at the action profile \bar{a}_k^* , $k \in m$. For any player i , let $i' >_{\epsilon, \delta} i$, and with the improvement suppose the new social-welfare-maximizing NE is \bar{a}'_k . Then, $\bar{a}'_k \geq \bar{a}_k^*$ for all $k \in m$.

PROOF SKETCH. Theorem 1.1.1 in part VII in [12] states that for any maximizer x^* of a concave function f on a closed convex set X satisfies $d_{\bar{v}}^- f(x^*) \geq 0$ for all \bar{v} in the direction $x^* - \bar{x}$ for all $\bar{x} \in X$. The potential P satisfies the requirements of Lemma 4.4. Also, since \bar{a}_k^* for all $k \in m$ maximizes P , from the maximizer property above we get that $d_{\bar{v}}^- P(\bar{a}_k^* \forall k) \geq 0$ all \bar{v} such that $\bar{v}_S \geq 0$ where S is the set $\{(k, l) \mid a_{k,l} > 0\}$ and $\bar{v}_{-S} = 0$. With ϵ, δ improvement the potential function has additional terms in the derivative: $\epsilon P_i d_{\bar{v}}^- E(\bar{a}_j^*, \bar{a}_{N(j)}^*) + \delta d_{\bar{v}}^- c_i C(\bar{a}_j^*)$. Due to monotonic E and C , the additional terms are ≥ 0 for $v \geq 0$. Thus, given the maximizer condition we get that still $d_{\bar{v}}^- P(\bar{a}_k^* \forall k) \geq 0$ for all \bar{v} such that $\bar{v}_S \geq 0$ where S is the set $\{(k, l) \mid a_{k,l} > 0\}$ and $\bar{v}_{-S} = 0$. By Lemma 4.4 this implies a potential-maximizing NE \bar{a}'_k with i' satisfies $\bar{a}'_k \geq \bar{a}_k^*$ for all $k \in m$. \square

The above result is a complete characterization for all NE if the scores are differentiable, since then there is a unique NE. More generally, a similar result holds for any NE.

THEOREM 4.7. In game G suppose a NE score is attained at the action profile \bar{a}_k^* , $k \in m$. For any player i , given $i' >_{\epsilon, \delta} i$, there exists a NE \bar{a}'_k after the improvement such that $\bar{a}'_k \geq \bar{a}_k^*$ for all $k \in m$.

PROOF SKETCH. When i improves we analyze the best-response dynamics starting from \bar{a}_k^* , showing collaboration increases at each step. As stated above, best-response dynamics converges to a NE [13]. We show, via induction, that at any step n the left directional derivatives of the utility of best responding player k w.r.t. the actions of k satisfy the conditions of Lemma 4.4, so that k best responds by increasing its collaboration. Showing this requires invocation of MLPDD; informally since others' collaborations have increased since the last time player k best-responded (by the induction hypothesis) so the left directional derivatives in positive directions also increase by the MLPDD property, which then allows us to invoke Lemma 4.4. \square

Counterexample for monotone collaboration with non-supermodular function: Consider a three-node line graph with the middle player 2's individual utility:

$$0.1(a_{2,1} + a_{2,3}) + a_{1,2} + a_{3,2} - a_{2,1}^2 - a_{2,3}^2 + 1 - \exp(a_{2,1} + a_{2,3})$$

Player 1's individual utility is $P_1(0.1a_{1,2} + a_{2,1}) - a_{1,2}^2$ and player 3's individual utility is $0.1a_{3,2} + a_{2,3} - a_{3,2}^2$. Observe that the term $1 - \exp(a_{2,1} + a_{2,3})$ is not supermodular, which can be inferred by noting that its partial derivative decreases with increase in other dimensions. This term is concave (and monotone in the positive quadrant) as it is the negation of an exponentiation of a convex function. It can be readily checked that all other terms satisfy the monotonicity and (strict) concavity assumptions. Collaboration efforts at the NE for this game are shown in Fig. 2 with varying P_1 , which reveals that collaborations do not increase with P_1 .

4.4 Monotone Scores and Distinguishability

A natural desideratum for any scoring mechanism is that improvement in a player's capability leads to higher scores for that player. We show below how this concept connects to distinguishability. We define monotone scores in a MI as follows:

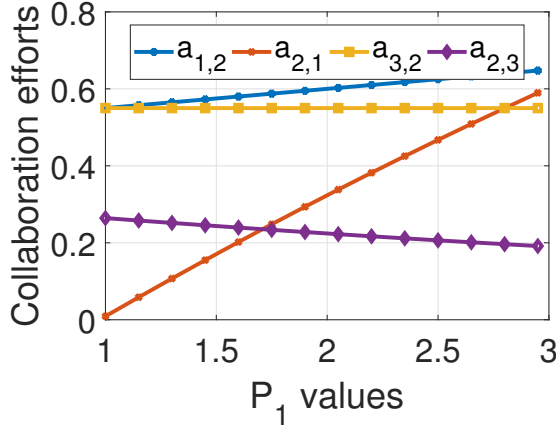


Figure 2: Monotone collaboration counterexample with non-supermodular utility.

Definition 4.8 (Monotone Scores). A scoring mechanism possesses the *monotone scores* property on a MI if $i' >_{\epsilon, \delta} i$ entails that the scores satisfy $S_{i'} \geq S_i + b_i(\epsilon, \delta)$, for some $b_i(\epsilon, \delta) > 0$, when other players are held fixed in the MI. b_i should be monotonically increasing in ϵ and δ .

Next, we focus on obtaining an expression for b_i in Definition 4.8 that is guaranteed by the LSWM mechanism. We show that LSWM provides distinguishability under a modularity assumption, and also show a counterexample when distinguishability is not achieved if the modularity assumption is violated. Specifically, we assume the following modular (or separable) form of the individual utility function:

$$u_i(\vec{a}_i, \vec{a}_{N(i)}) = \sum_{j \in N(i)} P_i \times E_{e(i,j)}(a_{i,j}, a_{j,i}) - c_i \times C_{e(i,j)}(a_{i,j}), \quad (2)$$

where $e(i, j)$ denotes the edge between player i and player j . We adopt the same monotonicity, concavity and supermodularity assumptions for $E_{e(i,j)}$ and $-C_{e(i,j)}$ as we had for E and $-C$ in Section 4.3.

The following readily verifiable observations allow for a structured proof of the next result on distinguishability: Associate the value $W_{i,j} = P_i \times E_{e(i,j)}(a_{i,j}, a_{j,i}) - c_i \times C_{e(i,j)}(a_{i,j}) + P_j \times E_{e(j,i)}(a_{j,i}, a_{i,j}) - c_j \times C_{e(j,i)}(a_{j,i})$ with every edge (i, j) . Then, the potential function (or social welfare) is $\sum_{i,j} W_{i,j}$. Further, $a_{i,j}, a_{j,i}$ affects only the value $W_{i,j}$. Then, the score of player i is $\sum_{j \in N(i)} W_{i,j} + O$, where O are other terms not affected by $a_{i,j}, a_{j,i}$ or P_i, c_i .

THEOREM 4.9. *Assume the collaboration game has modular utility function as defined above. For player i playing the LSWM induced game suppose the potential maximizing NE score is attained at the action profile \vec{a}_k^* for all $k \in m$. Given $i' >_{\epsilon, \delta} i$, i 's score with the improvement at the potential maximizing NE satisfies*

$$S_{i'} \geq S_i + \sum_{j \in N(i)} \epsilon P_i E_{e(i,j)}(\vec{a}_i^*, \vec{a}_{N(i)}^*) + \delta c_i C_{e(i,j)}(\vec{a}_i^*).$$

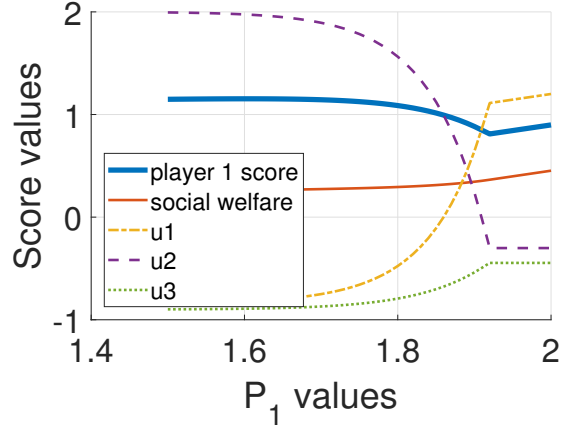


Figure 3: Monotone scores counterexample with non-modular utility.

PROOF. Given the observations above that the potential is a sum of $W_{i,j}$'s, we can only consider the effect of the improvement $i' >_{\epsilon, \delta} i$ on each $W_{i,j}$ separately for all neighbors j of i : $j \in N(i)$. The NE strategy with i yields an additional value $\epsilon P_i E_{e(i,j)}(\vec{a}_i^*, \vec{a}_{N(i)}^*) + \delta c_i C_{e(i,j)}(\vec{a}_i^*)$ for $W_{i,j}$ after i improves, thus, any potential maximizer after the improvement yields a value that is increased by at least $\epsilon P_i E_{e(i,j)}(\vec{a}_i^*, \vec{a}_{N(i)}^*) + \delta c_i C_{e(i,j)}(\vec{a}_i^*)$ for $W_{i,j}$. Next, as observed earlier, since each $W_{i,j}$ is affected only locally by $a_{i,j}, a_{j,i}$, in the new NE any change in $a_{i,j}, a_{j,i}$ does not affect other $W_{k,m}$ (where $(k, m) \notin \{(i, j) \mid j \in N(i)\}$) terms. In fact, the other actions $a_{k,m}$ stay the same as old NE. Also, as the score of i is $\sum_{j \in N(i)} W_{i,j} + O$ where O does not depend on $a_{i,j}, a_{j,i}$ or P_i, c_i and hence O is same in old and new NE, i 's score improves due to improvement in $\sum_{j \in N(i)} W_{i,j}$. This provides the desired result. \square

Counterexample for monotone scores with non-modular function: We provide a counterexample with a supermodular score function that does provide monotone collaboration as proved above. Consider a three-node line graph with the middle player 2's individual score being

$$0.1(a_{2,1} + a_{2,3}) + a_{1,2} + a_{3,2} - a_{2,1}^{1.05} - a_{2,3}^{1.05} - (a_{2,1}^4 + a_{2,3}^4)^{.25}$$

Player 1's individual score is $P_1(0.1a_{1,2} + a_{2,1}) - a_{1,2}^{1.05}$ and player 3's individual score is $0.1a_{3,2} + a_{2,3} - a_{3,2}^{1.05}$. Observe that term $-(a_{2,1}^4 + a_{2,3}^4)^{.25}$ is not modular, but is supermodular, which can be informally inferred by noting that its partial derivative, when it exists, increases in other dimensions within the domain. This term is also concave (and monotone in the positive quadrant) as it is the negation of a norm and norms are convex. It can be readily checked that all other terms satisfy the supermodularity, monotonicity and concavity assumptions. The scores at the NE for this game is shown in Fig. 3 with varying P_1 . As can be seen, the score of player 1 does not increase with increase in its capability P_1 .

Effect of a player's capability improvement on other players' score: The next result shows that the increase in other players' score when i improves depends on the graph structure.

THEOREM 4.10. *Assume the collaboration game has modular utility function. For player i playing the LSWM induced game suppose the potential maximizing NE score is attained at the action profile \bar{a}_k^* for all $k \in m$. Given $i' >_{\epsilon, \delta} i$, player j 's score ($j \neq i$) at the potential maximizing NE satisfies*

- *If player j is more than two hops away from player i , then j 's score stays the same.*
- *If player j is one or two hops away from player i , then j 's score increases but not more than i 's increase.*

PROOF. When i improves, it leads to change in actions $a_{i,k}$ and $a_{k,i}$ for any neighbor k . No other actions change. Then, the utility (not score) of every player two hops or more from i does not change (since no relevant action changed). If player j is more than two hops away, then all its neighbors are at least two hops away and hence unchanged utilities, thus, j 's score (sum of its and neighbors utility) is unchanged.

For j being one hop away, the change for i requires re-maximizing $W_{i,j} = P'_i \times E_{e(i,j)}(a_{i,j}, a_{j,i}) - c'_i \times C_{e(i,j)}(a_{i,j}) + P_j \times E_{e(j,i)}(a_{j,i}, a_{i,j}) - c_j \times C_{e(j,i)}(a_{j,i})$. We claim that this re-maximization leads to increase of both past utilities (part of it that arises from i, j interaction) $P_j \times E_{e(j,i)}(a_{j,i}^*, a_{i,j}^*) - c_j \times C_{e(j,i)}(a_{j,i}^*)$ and $P_i \times E_{e(i,j)}(a_{i,j}^*, a_{j,i}^*) - c_i \times C_{e(i,j)}(a_{i,j}^*)$; this claim is proved in the next paragraph. Thus, j 's utility (part of it that depends on i) increase is not more than the increase in $W_{i,j}$. Thus, j 's increase of score from increase in $W_{i,j}$ and the increased utility of some node k neighbor of both i, j together is less than i 's score increase from $W_{i,j}$ and $W_{i,k}$'s. Hence j 's score increases by less than i 's score.

Now, we prove the claim from the last paragraph. First, let the new maximum be at $a'_{i,j}, a'_{j,i}$. We know from super-modularity (adopting proof of Theorem 4.6 for this simpler scenario) that $a'_{i,j} \geq a_{i,j}^*$ and $a'_{j,i} \geq a_{j,i}^*$. The following function of only $a_{j,i}$ is maximized at $a'_{j,i}$: $f(a_{j,i}) = P'_i \times E_{e(i,j)}(a'_{i,j}, a'_{j,i}) - c'_i \times C_{e(i,j)}(a'_{i,j}) + P_j \times E_{e(j,i)}(a_{j,i}, a'_{i,j}) - c_j \times C_{e(j,i)}(a_{j,i})$. Now, due to monotone super-modular $E_{e(j,i)}$ and $a'_{i,j} \geq a_{i,j}^*$, we know that $E_{e(j,i)}(a_{j,i}^*, a'_{i,j}) \geq E_{e(j,i)}(a_{j,i}^*, a_{i,j}^*)$. Also, as the other terms of the function f are fixed, we have $P_j \times E_{e(j,i)}(a'_{j,i}, a'_{i,j}) - c_j \times C_{e(j,i)}(a'_{j,i}) \geq P_j \times E_{e(j,i)}(a_{j,i}^*, a'_{i,j}) - c_j \times C_{e(j,i)}(a_{j,i}^*) \geq P_j \times E_{e(j,i)}(a_{j,i}^*, a_{i,j}^*) - c_j \times C_{e(j,i)}(a_{j,i}^*)$. Thus, j 's utility increases. Next, the following function of only $a_{i,j}$ is maximized at $a'_{i,j}$: $f(a_{i,j}) = P'_i \times E_{e(i,j)}(a_{i,j}, a'_{j,i}) - c'_i \times C_{e(i,j)}(a_{i,j}) + P_j \times E_{e(j,i)}(a'_{j,i}, a'_{i,j}) - c_j \times C_{e(j,i)}(a'_{j,i})$. Now, due to monotone super-modular $E_{e(i,j)}$ and $a'_{j,i} \geq a_{j,i}^*$, we know that $E_{e(i,j)}(a_{i,j}^*, a'_{j,i}) \geq E_{e(i,j)}(a_{i,j}^*, a_{j,i}^*)$. Also, as the other terms of the function f are fixed, we have $P'_i \times E_{e(i,j)}(a'_{i,j}, a'_{j,i}) - c'_i \times C_{e(i,j)}(a'_{i,j}) \geq P'_i \times E_{e(i,j)}(a_{i,j}^*, a'_{j,i}) - c'_i \times C_{e(i,j)}(a_{i,j}^*) \geq P'_i \times E_{e(i,j)}(a_{i,j}^*, a_{j,i}^*) - c'_i \times C_{e(i,j)}(a_{i,j}^*)$. Further, as $P'_i > P_i, c'_i < c_i$ we get $P'_i \times E_{e(i,j)}(a_{i,j}^*, a_{j,i}^*) - c'_i \times C_{e(i,j)}(a_{i,j}^*) > P_i \times E_{e(i,j)}(a_{i,j}^*, a_{j,i}^*) - c_i \times C_{e(i,j)}(a_{i,j}^*)$. Thus, i 's utility increases.

If j is two hops away, from the above result its score increase due to increase in utility of k that is both a neighbor of i and j . As argued above the increase in k 's utility is less than the increase in score of i due to k . Thus, j 's increase in score is also limited to less than i 's increase in score. \square

Theorem 4.10 forms the crux of our argument in Section 5 on match distinguishability. Note that this result does not hold for social welfare scoring, where every player's score increases by the same amount as for the improved player.

5 DISTINGUISHABILITY IN A MATCH

5.1 Match Score Design Problem

As argued above, an important goal of competition design is to distinguish the relative capabilities of participants. In collaborative environments such as ours where the performance of one agent may enhance the score of others and the scoring mechanism is designed to incentivize collaboration, it is particularly important to verify that the differential capability accrues higher match scores to the improved agent. Here we show, under specified conditions, that LSWM achieves distinguishability in a match, in a precise quantitative sense.

Recall that a match is a sequence of $|T|$ number of MIs played by a set m of players. We posit that the match generation process samples a graph randomly (sample an edge with probability p) for each MI. p must be chosen judiciously; a smaller p generates more sparse graphs encouraging distinguishability but not enabling testing the collaborative interaction. For notational ease, we use $\epsilon = \delta$ for the rest of this paper. In particular, if $\epsilon = \delta$ then it can be seen that Thm. 4.9 for a MI implies $S_i(T) \geq (1 + \epsilon)S_i(\tau)$. Further, we normalize the scores in a MI to lie in $[V, 1]$ assuming that every player can achieve at least $V > 0$ score (for any graph).

Definition 5.1 (Match Distinguishability). Given $i >_{\epsilon} j$, match distinguishability means that the match scores satisfy $S_i(M) \geq S_j(M) + g(\epsilon)$ with probability $p(\epsilon)$ for player i and j and any set of other opponents playing in a match M sampled randomly as described above. The functions $g(\cdot), p(\cdot)$ should be both monotonically increasing and $g(0), p(0) \geq 0$.

Also, as LSWM maximizes social welfare in each MI and the score in a match is the sum of scores in every MI within the match, LSWM also maximizes social welfare over a match. Thus, we only need to analyze match distinguishability of LSWM in the sense defined above.

5.2 Match Distinguishability of LSWM

LEMMA 5.2. *If $i >_{\epsilon} j$ and the collaboration game has modular utility function as in Theorem 4.9, then for any match M sampled as described above we get*

$$\mathbb{E}(S_i(M)) \geq \mathbb{E}(S_j(M)) + (1 - p)(1 - p^2)^{|m|-2} \epsilon V |T|$$

PROOF. First, observe that if i and j have same capability then $\mathbb{E}(S_i(M)) = \mathbb{E}(S_j(M))$. Observe that the expected match score is the sum of expected measurement interval scores. Also, expected measurement interval score is just the probability weighted average over all possible graphs. Then, we claim that the probability weighted average of scores for each measurement interval over all possible graphs is same for both players. This follows from the fact that for any graph, switching i and j switches their score (due to same capability) and the resultant graph is still one among all possible graphs occurring with the same probability as the original

one. Thus, the weighted average of scores over all graphs is same for i and j .

Next, if i improves his capability by ϵ , then his score improves by at least ϵV for each possible graph (V is min score across all graphs). Also, if i is connected to j by a path length > 2 or not connect at all, then j 's score does not change when i improves due to the modularity assumption on utility. This follows from the proof of Theorem 4.10 where a player improving affects only utility of neighbors, thereby only affecting scores of players within length 2 path. The probability mass of graphs satisfying any path between i, j length > 2 or not connected is $(1-p)(1-p^2)^{|m|-2}$. $1-p$ arises from no direct edge, and $(1-p^2)^{|m|-2}$ arises from no length two paths. Thus, i 's probability weighted average score over all graphs increase by at least $(1-p)(1-p^2)^{|m|-2}\epsilon V$, which is the score increase for a measurement interval. Adding over all measurement intervals gives the desired results. \square

Lemma 5.2 establishes distinguishability in expected scores. The main result of this section uses the McDiarmid's concentration inequality [5] to also establish distinguishability in actual scores with high probability.

THEOREM 5.3. *If $i >_\epsilon j$ and the collaboration game has modular utility function as in Theorem 4.9, then for any match M sampled as described above we get $S_i(M) \geq S_j(M) + \xi$, with probability $1 - 2 \exp\left(\frac{-2|T|\nu^2}{9}\right)$ for $|T| \geq \frac{3\xi}{\nu}$, where $\nu = (1-p)(1-p^2)^{|m|-2}\epsilon V$.*

6 EXPERIMENTS

Our experiments shed light on (1) convergence rate of best-response dynamics in the LSWM-induced game and (2) variation of distinguishability with sparsity of graphs. We also analyze distinguishability further by an in-depth analysis of the counterexample shown earlier for distinguishability and analyze the effect of out of equilibrium non-collaborative actions of players. Towards that end, we experiment with games with parametrized utilities where the parameters are chosen at random, and the results are averages over the randomly chosen instances. We use parametrized utility function family from our counterexamples to also show experimentally that distinguishability holds on average even though it does not for the specific values of parameters in the counterexample. All experiments were performed using MATLAB on a machine with a 2.8GHz processor and 12GB RAM.

For the best-response dynamics we chose a supermodular function with score functions as

$$E(\vec{a}_i, \vec{a}_{N(i)}) = \sum_j w_{i,j}(a_{j,i} + f * a_{i,j}) \text{ and}$$

$$C(\vec{a}_i) = \left(\sum_j a_{i,j}^K\right)^{1/K} + \sum_j c_{i,j} a_{i,j}^L,$$

with $w_{i,j}, c_{i,j} \geq 0$. The graph is chosen with 0.5 probability of each edge existing. The parameters $w_{i,j}, c_{i,j}, f, K, L$ are sampled at random from the value ranges $(0, 1), (0, 1), (0, 0.1), (1, 7), (1, 2)$ respectively. The performance level and cost of collaboration parameters P_i, c_i are fixed to 1. The best-response dynamics was run until no strategies changed, with a 10^{-4} tolerance. As shown in Fig. 4, the number of rounds grows with increasing graph size.

Our next set of experiments build on the counterexample for distinguishability presented earlier. We examine what happens in the case different graphs are sampled with the same kind of utility in that counterexample. Towards that end, we choose the utility of each player to be

$$0.1P_i(1 + \sum_{j \in N(i)} a_{i,j}) - \sum_{j \in N(i)} a_{i,j}^{1.05} - 1(|N(i)| > 1) \left(\sum_{j \in N(i)} a_{i,j}^4\right)^{0.25},$$

where 1 is the indicator function. This utility imposes a l_4 norm cost only when more than one neighbor is present and has the same utility as for the counterexample against distinguishability. We perform the experiment with $P_1 = 1, P_2 = 2, P_3 = 3$ (so that player 3 deserves highest scores due to highest capabilities). We sample 50 graphs for each fixed number of edges and Fig. 5 shows the percentage change in scores of players as measured against player 2's scores. As can be seen, a complete graph provides no distinguishability whereas distinguishability increases on average in percentage variation with fewer edges.

Next, we conduct the same experiment as done in Fig. 5, but make player 3 a non-collaborator, that is player 3 plays an out of equilibrium action to not collaborate with any other player and other players continue to use their collaboration strategy. Fig. 6 shows the results where it can be seen that when collaboration matters (that is, with 1 or 2 edges) player 3 loses a significant amount of score in spite of possessing the best capability. As expected, when collaboration is not required (that is, no edges in the graph) player 3 performs the best.

Finally, we show how the total match scores evolve with increasing MI averaged over 30 randomly chosen match runs (probability of edge is 0.5), for the all-collaborating setup used in Fig. 5 as well as player 3 non-collaborator setup used in Fig. 6. Following the result of Fig. 6, this result reinforces the claim that higher scores in a match arise from both capability of player and playing the equilibrium action.

7 RELATED WORK

Potential games and graphical games have been studied extensively [23]. Games that are both potential and graphical have also been explored [4, 7, 27, 27]. Recent work [2, 24] relates such games to Markov random fields. We contribute to this literature by providing a mechanism to convert a graphical game to a social-welfare-maximizing graphical potential game while providing the distinguishability property.

Collaboration has been studied extensively in the framework of cooperative or coalitional games [8]. Supermodularity has also been widely used, in the form of supermodular set functions, in cooperative games [11, 30] and in modeling social networks [14]. Our work is set within the non-cooperative game framework, under the assumption that neither the competition designer nor participants have the means to enforce coalition agreements. Supermodular games [18] (games with continuous supermodular utility) have also been extensively studied in non-cooperative settings, particularly for modeling economic situations with complementary goods. Our supermodularity and modularity conditions reveal the nature of utility functions that enable providing the properties we desire.

Private provisioning of public goods is an area of research where a consumer benefits from provisioning of public goods by players

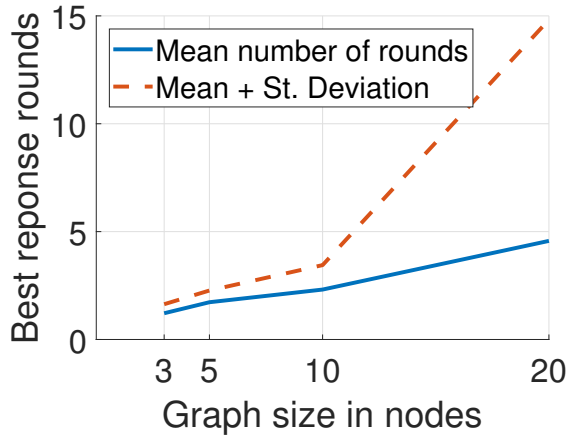


Figure 4: Number of best-response rounds to converge to NE of the collaboration game with varying number of nodes.

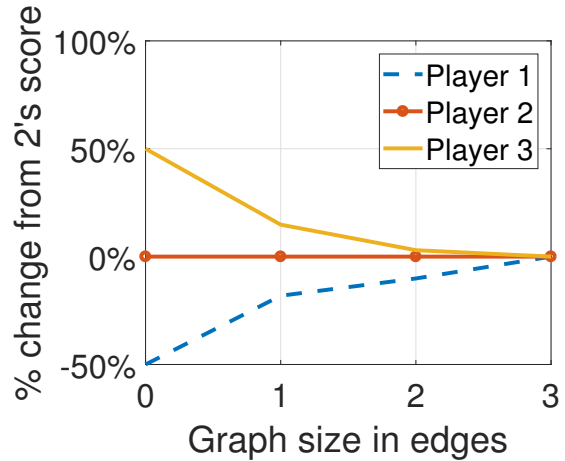


Figure 5: Distinguishability with graph structure with three players using player 2's score as baseline. Three edges is a complete graph.

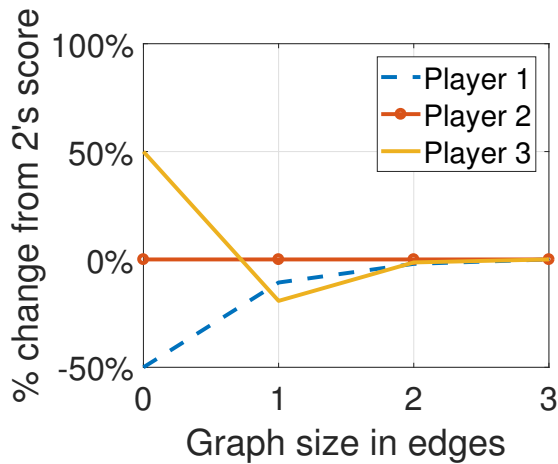


Figure 6: Distinguishability with the same setup as Fig. 5 but player 3 never collaborates, and suffers with one or two edges.

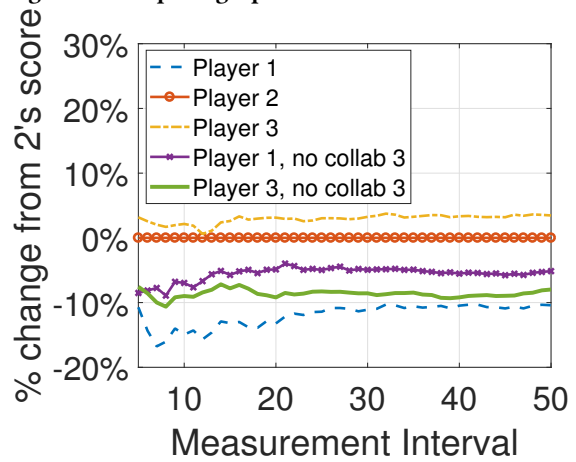


Figure 7: Match score with increasing MIs, using player 2 as baseline for all collaborating and player 3 not collaborating.

[3, 9]. Our work relates to some recent work on private provisioning of public goods on networks, in which a consumer interact within a fixed network structure and benefits only from their direct neighbors' provisions [1, 6]. In particular, these works assume knowledge of the private provisioning of each player to decide any payments for the players. The income redistribution mechanisms also restricts the payments as, unlike free scores in our case, it is not desirable for the designer to incur a cost. Thus, the income redistribution mechanisms have very different characteristics from our LSWM mechanism. Moreover, these works have no notion of capability of players and ensuring higher overall utility of players' with higher capability is not a concern.

There is a large body of work on mechanism design [23] including constrained design [10] as well as design on graphs [19]. As described above, however, our competition scoring problem differs considerably from the standard mechanism design setting.

8 SUMMARY

We studied the problem of designing scoring mechanisms to promote collaboration in a competition. Our proposed LSWM scoring mechanism maximizes social welfare in equilibrium, while preserving distinguishability under certain conditions. The key idea of LSWM is to exploit locality in the underlying interaction so that only relevant counterparts are brought into the incentive scheme. We supported the design through theoretical and experimental analyses. Some of the ideas presented in this work apply more broadly to graphical games and characterizations of continuous supermodularity.

Acknowledgement. This work was supported in part by the DARPA Spectrum Challenge project and MURI grant W911NF-18-1-0208 from the US Army Research Office.

REFERENCES

- [1] Nizar Allouch. 2015. On the private provision of public goods on networks. *Journal of Economic Theory* 157 (2015), 527 – 552.
- [2] Yakov Babichenko and Omer Tamuz. 2016. Graphical potential games. *Journal of Economic Theory* 163 (2016), 889–899.
- [3] B. Douglas Bernheim. 1986. On the voluntary and involuntary provision of public goods. *American Economic Review* (1986), 789–793.
- [4] Kostas Bimpikis, Shayan Ehsani, and Rahmi Ilklic. 2014. Cournot competition in networked markets (abstract). In *15th ACM Conference on Economics and Computation*. 733.
- [5] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. 2013. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press.
- [6] Yann Bramoullé and Rachel Kranton. 2007. Public goods in networks. *Journal of Economic Theory* 135, 1 (2007), 478 – 494.
- [7] Yann Bramoullé, Rachel Kranton, and Martin D’amours. 2014. Strategic interaction and networks. *American Economic Review* 104, 3 (2014), 898–930.
- [8] Georgios Chalkiadakis, Edith Elkind, and Michael Wooldridge. 2011. *Computational Aspects of Cooperative Game Theory*. Morgan & Claypool.
- [9] Richard Cornes and Roger Hartley. 2007. Aggregative public good games. *Journal of Public Economic Theory* 9, 2 (2007), 201–219.
- [10] Jason Hartline. 2008. *Lectures on Frugal Mechanism Design*. Technical Report. Northwestern University.
- [11] Daisuke Hatano and Yuichi Yoshida. 2017. Computing least cores of supermodular cooperative games. In *31st AAAI Conference on Artificial Intelligence*. 551–557.
- [12] Jean-Baptiste Hiriart-Urruty and Claude Lemaréchal. 2013. *Convex Analysis and Minimization Algorithms I: Fundamentals*. Vol. 305. Springer.
- [13] Zhengao Huang. 2002. *Fictitious Play for Games with a Continuum of Strategies*. Ph.D. Dissertation. State University of New York at Stony Brook.
- [14] Ayumi Igarashi. 2017. Supermodular games on social networks. In *16th International Conference on Autonomous Agents and Multi-Agent Systems*. 299–307.
- [15] Patrick R. Jordan and Michael P. Wellman. 2010. Designing the Ad Auctions game for the Trading Agent Competition. In *Agent-Mediated Electronic Commerce: Designing Trading Strategies and Mechanisms for Electronic Markets*, Esther David, Enrico Gerding, David Sarne, and Onn Shehory (Eds.). Lecture Notes in Business Information Processing, Vol. 59. Springer-Verlag.
- [16] Michael Kearns, Michael L. Littman, and Satinder Singh. 2001. Graphical models for game theory. In *17th Conference on Uncertainty in Artificial Intelligence*. Seattle, 253–260.
- [17] Wolfgang Ketter, John Collins, and Prashant Reddy. 2013. Power TAC: A competitive economic simulation of the smart grid. *Energy Economics* 39 (2013), 262–270.
- [18] Jonathan Levin. 2003. *Supermodular games*. Technical Report. Lectures Notes, Department of Economics, Stanford University.
- [19] Bin Li, Dong Hao, Dengji Zhao, and Tao Zhou. 2017. Mechanism design in social networks. In *31st AAAI Conference on Artificial Intelligence*. 586–592.
- [20] Dov Monderer and Lloyd S. Shapley. 1996. Potential Games. *Games and Economic Behavior* 14 (1996), 124–143.
- [21] Abraham Neyman. 1997. Correlated equilibrium and potential games. *International Journal of Game Theory* 26, 2 (1997), 223–227.
- [22] Noam Nisan. 2007. Introduction to mechanism design (for computer scientists). See [23], 209–241.
- [23] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani (Eds.). 2007. *Algorithmic Game Theory*. Cambridge University Press.
- [24] Luis E Ortiz. 2015. Graphical potential games. *arXiv preprint arXiv:1505.01539* (2015).
- [25] Robert R Phelps. 1993. *Convex Functions, Monotone Operators and Differentiability* (second ed.). Lecture Notes in Mathematics, Vol. 1364. Springer.
- [26] David Simchi-Levi, Xin Chen, and Julien Bramel. 2014. *The Logic of Logistics: Theory, Algorithms, and Applications for Logistics Management* (third ed.). Springer, Chapter Convexity and supermodularity, 15–44.
- [27] Cem Tekin, Mingyan Liu, Richard Southwell, Jianwei Huang, and Sahand Haji Ali Ahmad. 2012. Atomic congestion games on graphs and their applications in networking. *IEEE/ACM Transactions on Networking* 20, 5 (2012), 1541–1552.
- [28] Eliot van Buskirk. 2009. Now the Netflix prize was won. *Wired* (Sept. 2009).
- [29] Michael P. Wellman, Amy Greenwald, and Peter Stone. 2007. *Autonomous Bidding Agents: Strategies and Lessons from the Trading Agent Competition*. MIT Press.
- [30] Yair Zick, Evangelos Markakis, and Edith Elkind. 2012. Stability via Convexity and LP Duality in OCF games. In *26th AAAI Conference on Artificial Intelligence*. 1506–1512.